

Preventing and Detecting Insider Attacks Using IDS

Nathan Einwechter 2002-03-20

Preventing and Detecting Insider Attacks Using IDS

by *Nathan Einwechter*

last updated March 20, 2002

A Typical Insider Attack - the Disgruntled Employee

Shortly after lunch break, an employee angrily strides out of his supervisor's office, down two rows of desks, and into a single cubicle. He slumps down into his chair and releases an exasperated sigh, as he runs his hands through his hair in disappointment.

The raise he thought he was in for has been turned down. He slowly stands up, peering over the cubicle walls to survey the area for other employees. But the area is deserted as most people are out enjoying lunch. Sitting back down, he turns to his computer console, goes to the command line and brings nmap to life against the company's accounting systems. The console displays accounting's SQL server. A few keystrokes later, the employee has edited a few columns in the database, giving himself the raise he had longed for.

The Problems With Insider Attacks

This noisy attack could have easily been detected if the network administrators had utilised an internal IDS system. An IDS could have not only detected the attack, it could also have allowed administrators to learn of vulnerable services, and figure out who the attacker really was.

Insider Attacks are an unusual type of threat. Unlike external attacks, the intruder is someone who has been entrusted with authorized access to the network. In fact, the attacker requires access in order to fulfil their obligations to the victim organization. Furthermore, they often have a substantial amount of knowledge about the network architecture, including where their targeted files or systems are located. Because many organizations' security is focussed on protecting the perimeter of the network, little attention is paid to what is occurring within the system. As a result, insider attacks may not be discovered for months after the attack, long enough for the perpetrator to get off scot-free.

Although the vast majority of attention is given to protecting against external threats, insider

attacks are a serious, and common, threat. According to a CPI/FBI survey, 59% of companies surveyed said they have had one or more attacks reported internally. Almost 8% of those companies reported 60 or more internal incidents. (The survey is available upon request at: <http://www.gocsi.com/forms/fbi/pdf.html>.) The main issues that need to be addressed in preventing and detecting insider attacks include: what the basic problems of insider attacks are, how IDS systems can help solve this problem, and finally how an internal IDS system should be deployed using various IDS technologies.

Protecting Against Insider Attacks

Because of this extra knowledge of internal system procedures, computer and server configurations, and policy, the inside attacker has an advantage over the system administrators. It's hard enough to keep external attackers out; keeping authorized users from manipulating their access to the system to exploit sensitive information makes the administrator's job that much harder.

Not only must security personnel combat the inside attackers' privileged knowledge of a specific network, but they have few attack prevention and detection products with which to do so. This is particularly true of intrusion detection systems (IDS), most of which are designed and implemented to detect external threats. However, this does not have to be the case. As security admins are starting to realize, IDSs can be a valuable tool in the effort to prevent and detect insider attacks.

How Can an IDS Protect Against Insider Attacks?

As their name implies, intrusion detection systems are designed to detect attacks. However, configuring an IDS to detect internal attacks can be difficult. Part of the challenge lies in creating a good rule set for the internal IDS. The reason the rule set needs to be different is due to the fact that different network users require a different amount of access to different services, servers, and systems for their work. The rule set of the internal IDS system should be created so that all the static of employees' day-to-day work activities, such as accessing various services and servers, does not trigger attack warnings, and only the important information is reported. This important information would include detected activities that users do not require for their daily work, as well as any other glaringly obvious attacks such as an nmap, or queso scan.

The logging and reporting of attacks by the internal IDS systems can be used to do much more than detect specific, isolated, and unrelated attacks. By combining the data from all internal IDS systems, system administrators can identify attack trends and patterns. Once attack trends and patterns are identified, the admins will be more able to identify any network users who pose a threat to network security, have been exhibiting any malicious network behaviour, or who are doing anything that is against company policy in general. Once these users have been identified, the proper action can be taken to prevent any successful intrusions or the continuance of the activity.

Further, the logs provided by IDS systems can allow the system administrators an audit trail in case there are in fact any successful intrusions. Identified attack trends and patterns can also allow system administrators to see where people are trying to attack against the most. This would allow them to identify any possible security holes, or policy oversights, as well as any servers on the network that have a higher risk of being attacked, and thus allow them to know which systems to keep security tighter on.

Deploying IDS to Combat Insider Attacks

A combination of IDS systems should be used to detect insider attacks. The systems that can be deployed to assist in combating against insider attacks include network intrusion detection systems (NIDS), network node intrusion detection systems (NNIDS), host-based intrusion detection systems (HIDS), anomaly-based intrusion detection systems, and the analytical powers of the distributed intrusion detection system (dIDS). These systems each have their uses within the network, along with certain advantages and disadvantages, all of which shall be discussed. The use of network taps to allow some of these systems to operate will also be covered, as well as general security guidelines to follow with regards to deploying the various IDS systems.

Network Intrusion Detection Systems (NIDS)

NIDS systems can be used as a broader detection tool, to detect attacks against a number of networked systems within its particular network segment. This type of system provides the greatest scope of monitoring, and would be best suited for a general IDS system that covers non-critical systems or as a secondary IDS for critical systems. A good example of where a NIDS might be deployed (when protecting against internal threats) is right between a division router, and that division's actual systems. By doing this, any attacks against any system in that

division would be detected and reported to network administration. NIDSs may also be deployed on switches, hubs, or any other point where multiple systems are networked together, usually through the use of a network tap.

Network Node Intrusion Detection Systems (NNIDS)

NNIDS systems are ideally suited to be on critical systems, such as database servers and backup servers. This IDS system detects attacks only against the network node on which it is installed; it does not worry about any other attacks that may be occurring on other parts of the network. This limits the scope of the NNIDS, but allows extra detection abilities for mission critical systems.

Host-Based Intrusion Detection Systems (HIDS)

HIDS systems are less concerned with actually detecting attacks from a network/protocol perspective; instead, they continually look at system logs, critical system files, and other resources that may be monitored for any suspicious activity such as critical file modifications, or suspicious patterns of activity.

Some of the specific things a HIDS can monitor include event logs, IDS logs, system files, and the windows registry. When it monitors the system files and windows registry, it creates and stores a snapshot of the last known "clean" system. It then compares this clean snapshot against the current state of the system to detect any modified files, etc. If it detects any modifications, or suspicious activity in the logs, it simply alerts the administrators to the changes, and appropriate action can then be taken. While the HIDS doesn't differentiate between internal and external attacks, it will notify the system administrator of an unauthorized file change that, if conducted by an inside attacker, will be detected more rapidly than without the HIDS.

HIDS are usually installed on critical workstations, and servers that require the extra layer of protection that is on top of the regular IDS system installed.

Anomaly-Based Intrusion Detection Systems

Anomaly-based intrusion detection systems are a relatively new idea. In combating an internal threat, the idea behind an anomaly-based IDS is to establish a baseline of "normal" activity by

what types of traffic are going across the network destined to specific systems, or originating from specific systems and in what amounts in normal working conditions. Any deviance from that baseline in either traffic type, or amount could then be detected and considered a potential incident.

Anomaly-based intrusion detection systems are becoming more and more important in protecting networks from insider attacks. This is largely because they solve the difficulty of allowing certain users access to certain systems but not others. The anomaly-based IDS solves this by only detecting things out of the normal base line for that user, thus circumventing this problem without a lot of analytical time that would normally be used to filter out the static, or normal traffic from the attack logs of other IDS systems.

Anomaly based IDS are usually deployed in the same locations that a NIDS would be, which is to say, switches, hubs, or any other point where multiple systems are networked together

Distributed Intrusion Detection Systems (dIDS)

Many system administrators find it difficult to review the data from all of the networks IDS systems. On a large network with an understaffed IT department and a large number of IDS logs, there are not enough hours in the day to review all the information that may be generated. This problem, however, can be taken care of by implementing a dIDS system. dIDS systems, in their most basic form, collect and aggregate attack logs from multiple IDS and firewall devices. This allows system administrators to view attack information in an aggregated form at a centralized location. This reduces the time needed to review the log files and allows the administrators to have a broader view of attack trends and patterns across the network, thus achieving the goal of identifying attack trends and patterns as described earlier, in a simple manner. More information on the subject can be found by reading the SecurityFocus article [An Introduction to Distributed Intrusion Detection Systems](#).

The dIDS system helps prevent and detect insider attacks by considerably shortening the amount of time system administrators require to review logs files, and identify attack trends and patterns. By reducing the amount of time required to review log files and identify attack trends and patterns, insider attacks will be discovered quicker than most conventional methods, as well as allow the system administrators to identify possible future attacks before they happen.

Use of Network Taps

On the networks of today, the use of hubs is becoming less of an option due to efficiency problems. Thus, most networks are using switches in the place of the aged hubs. This is done to increase the efficiency and speed of the network. In order to work effectively, network intrusion detection systems and anomaly-based IDS systems must see the largest number of systems possible. Since a network switch basically creates a whole different network segment for each port connected to it, the use of a NIDS or anomaly-based IDS becomes an issue.

To get around this, some people have suggested a transition over to NNIDS for a majority of the intrusion detection capabilities. This, however, has its disadvantages, as it becomes costly both financially, to administer a network consisting completely of NNIDS systems.

As an alternative to the costly implementation of NNIDS systems, the NIDS and anomaly-based IDS systems can both be implemented onto the switched network using network taps on all the switches. Network taps on switches allow the system using the tap to view all traffic being sent or received by the switch. Since all the traffic is viewable by the system on this tap, the use of a NIDS or an anomaly-based IDS is again viable. This solution is significantly less costly, and also opens the option for administrators to do network or protocol analysis to find any networking or programming issues within the network. Thus, it solves two problems at once.

IDS Security

The security of the actual IDS systems and the attack logs they produce is often overlooked. The reason IDS security itself should be taken into account is to prevent the tampering of attack logs, which become critical when attempting to discover how an intrusion occurred and who committed the intrusion. The first thing to consider in increasing the security of IDS information is the use of remote logging. When using a dIDS system, this is a given, as the log information is sent to a centralised server for aggregation and analysis. When a dIDS system is not being used however, a centralised logging server should be created to receive the attack logs from the various IDS and firewall systems that are reporting attacks. The remote logging server should be as secure as possible, with the least possible number of services running. It should be dedicated to nothing but receiving these logs. Doing this prevents the possibility of log file modification, and thus a destruction of vital information.

The second IDS security concern that must be looked at is methods of administrative access. Users should be sure to have a good secure (non-default) administrators account and password

on every device. Also make sure that all other accounts have been removed from the device. All services that are not required by the device to be running and reporting attack information should also be disabled. SNMP access should be disabled if not required, or the default community names changed to a more secure phrase.

Conclusion

Although insider attacks pose some unique challenges for security administrators, they can be easily detected by various types of IDS systems. By utilising these systems, attacks can not only be detected, they can also be properly investigated by identifying attack trends and patterns. The IDS systems that allow us to accomplish these goals must also be protected against attacks as well, to prevent the corruption of attack data. It is only through identifying attack trends and patterns, and keeping logs un-corrupted that insider attacks can be thwarted from the IDS part of the security spectrum.

[Nathan Einwechter](#) is currently a core member of [BCN Group](#), developing a proposed national cyber defense system. He also has worked as a System Developer/Incident Analyst with [myNetWatchman](#), and as a Senior Research Scientist/Head of Research and Development for [Fate Research Labs](#).

[Privacy Statement](#)

Copyright 2006, SecurityFocus