

ITL Bulletin

ADVISING USERS ON INFORMATION TECHNOLOGY

ACQUIRING AND DEPLOYING INTRUSION DETECTION SYSTEMS

Intrusion detection is a topic extensively discussed in the popular press, IT trade articles, scholarly journals, and security newsletters. Intrusion detection systems (IDSs) hold great promise for deterring or mitigating the damage caused by "hacking" or breaking into sensitive IT systems. Thus, IDSs are at the forefront of many organizations' attention and now their checkbooks. Careful planning, phased deployments, and specialized training are among the steps that can be taken to minimize the potential for costly deployments, poor returns on investment, or high maintenance costs while maximizing the security benefit to an organization. This *ITL Bulletin* provides basic information about IDSs to help organizations avoid common pitfalls in acquiring, deploying, and maintaining IDSs. The topics covered are:

- A definition of intrusion detection,
- Reasons to acquire IDSs,
- Types of IDSs,
- IDS monitoring approaches,
- IDS event analysis approaches,
- IDSs that automatically respond to attacks,
- Tools that complement IDSs,
- Limitations of IDSs,
- Deployment of IDSs, and
- The future of IDSs.

A Definition of Intrusion Detection

Intrusion detection is the process of detecting unauthorized use of, or attack upon, a computer or network. IDSs are software or hardware systems that detect such misuse. IDSs can detect attempts to compromise the confidentiality, integrity, and availability of a computer or network. The attacks can come from

attackers on the Internet, *authorized* insiders who misuse the privileges given them, and *unauthorized* insiders who attempt to gain unauthorized privileges.

Reasons to Acquire IDSs

Intrusion detection capabilities are rapidly becoming necessary additions to every large organization's security infrastructure. The question for security professionals should not be whether to use intrusion detection, but which features and capabilities to use. However, one must still justify the purchase of an IDS. There are at least three good reasons to justify the acquisition of IDSs: to detect attacks and other security violations that cannot be prevented, to prevent attackers from probing a network, and to document the intrusion threat to an organization.

Detecting attacks that cannot be prevented

Attackers, using well-known techniques, can penetrate many networks. This often happens when known vulnerabilities in the network cannot be fixed. For instance, in many legacy systems, the operating systems cannot be updated. In updateable systems, administrators may not have or take the time to install all the necessary patches in a large number of hosts. In addition, it is usually not possible to perfectly map an organization's computer use policy to its access control mechanisms and thus authorized users often can perform unauthorized actions. Users may also demand network services and protocols that are known to be flawed and subject to attack. Although, ideally, we would fix all vulnerabilities, this is seldom possible. Therefore, an excellent

Continued on page 2

ITL Bulletins are published by the Information Technology Laboratory (ITL) of the National Institute of Standards and Technology (NIST). Each bulletin presents an in-depth discussion of a single topic of significant interest to the information systems community. **Bulletins are issued on an as-needed basis** and are available from ITL Publications, National Institute of Standards and Technology, 100 Bureau Drive, Stop 8900, Gaithersburg, MD 20899-8900, telephone (301) 975-2832. To be placed on a mailing list to receive future bulletins, send your name, organization, and address to this office.

Bulletins issued since May 1998

- *A Comparison of Year 2000 Solutions*, May 1998
- *Training for Information Technology Security: Evaluating the Effectiveness of Results-based Learning*, June 1998
- *Cryptography Standards and Infrastructures for the Twenty-first Century*, September 1998
- *Common Criteria: Launching the International Standard*, November 1998
- *What Is Year 2000 Compliance?*, December 1998
- *Secure Web-based Access to High Performance Computing Resources*, January 1999
- *Enhancements to Data Encryption and Digital Signature Federal Standards*, February 1999
- *Measurement and Standards for Computational Science and Engineering*, March 1999
- *Guide for Developing Security Plans for Information Technology Systems*, April 1999
- *Computer Attacks: What They Are and How to Defend Against Them*, May 1999
- *The Advanced Encryption Standard: A Status Report*, August 1999
- *Securing Web Servers*, September 1999

approach for protecting a network may be to use an IDS to detect when an attacker has penetrated a system using an uncorrectable flaw. It is better at least to know that a system has been penetrated so that administrators can perform damage control and recovery than not to know that the system has been penetrated.

Preventing attackers from probing a network

A computer or network without an IDS may allow attackers to leisurely and without retribution explore its weaknesses. If a single, known vulnerability exists in such a network, a determined attacker will eventually find and exploit it. The same network with an IDS installed is a much more formidable challenge to an attacker. Although the attacker may continue to probe the network for weaknesses, the IDS should detect these attempts, may block these attempts, and can alert security personnel who can take appropriate action.

Documenting the threat

It is important to verify that a network is under attack or likely to be attacked to justify spending money for securing the network. Furthermore, it is important to understand the frequency and characteristics of attacks in order to understand what security measures are appropriate

Who we are

The Information Technology Laboratory (ITL) is a major research component of the National Institute of Standards and Technology (NIST) of the Technology Administration, U.S. Department of Commerce. We develop tests and measurement methods, reference data, proof-of-concept implementations, and technical analyses that help to advance the development and use of new information technology. We seek to overcome barriers to the efficient use of information technology, and to make systems more interoperable, easily usable, scalable, and secure than they are today.

for the network. IDSs can itemize, characterize, and verify the threat from both outside and inside attacks, thereby providing a sound foundation for computer security expenditures. Using IDSs in this manner is important, since many people mistakenly believe that no one (outsiders or insiders) would be interested in breaking into their networks.

Types of IDSs

There are several types of IDSs available today, characterized by different monitoring and analysis approaches. Each has distinct uses, advantages, and disadvantages. IDSs can monitor events at three different levels: network, host, and application. IDSs can analyze these events using two techniques: signature detection and anomaly detection. Some IDSs also have the ability to automatically respond to the detected attacks. These variations are discussed in the following sections.

IDS Monitoring Approaches

One way to delineate IDSs is to look at what they monitor. Some IDSs listen on network backbones and analyze network packets to find attackers. Other IDSs reside on the hosts that they are defending and monitor the operating system for signs of intrusion. Still others monitor individual applications.

Network-Based IDSs

Network-based IDSs, currently the most common type of commercial product offering, detect attacks by capturing and analyzing network packets. Listening on a network backbone, a single network-based IDS can monitor a large amount of information. Network-based IDSs usually consist of a set of single-purpose hosts that "sniff" or capture network traffic in various parts of a network and report attacks to a single management console. Since no other applications run on the hosts used by a network-based IDS, they can be secured against attack. Many of them even have "stealth" modes, which make it extremely difficult for

an attacker to detect their presence and locate them.

Advantages:

- A few well-placed network-based IDSs can monitor a large network.
- The deployment of network-based IDSs has little impact upon an existing network. The network-based IDSs are usually passive devices that listen on a network wire without interfering with the normal operation of a network. Thus, it is usually easy to retrofit a network to include network-based IDSs with a minimal installation effort.
- Network-based IDSs can be made very secure against attack and even made invisible to many attackers.

Disadvantages:

- Network-based IDSs may have difficulty processing all packets in a large or busy network and, therefore, may fail to recognize an attack launched during periods of high traffic. Some vendors are attempting to solve this problem by implementing IDSs completely in hardware, which is much faster. The need to analyze packets quickly also forces vendors to try and detect attacks with as little computing resources as possible, which may reduce detection effectiveness.
- Many of the advantages of network-based IDSs do not always apply to more modern switch-based networks. Switches can subdivide networks into many small segments (usually one fast Ethernet wire per host) and can provide dedicated links between hosts serviced by the same switch. Most switches do not provide universal monitoring ports and this reduces the monitoring range of a network-based IDS sensor to a single host. In switches that do provide such monitoring ports, often the single port cannot mirror all traffic traversing the switch.
- Network-based IDSs cannot analyze encrypted information. This increasingly will become a problem as use of encryption becomes more popular both by organizations and by attackers.

- Most network-based IDSs do not report whether or not an attack was successful, they only report that an attack was initiated. After a detected attack, administrators must manually investigate each attacked host to determine whether or not the hosts were penetrated.

Host-Based IDSs

Host-based IDSs operate by analyzing the activity on a particular computer. As such, they must collect information from the host they are monitoring. This allows an IDS to analyze activities on the host at a very fine granularity and to determine exactly which processes and users are performing malicious activities on the operating system. Some host-based IDSs simplify the management of a set of hosts by having management functions and attack reports centralized at a single security console. Others generate messages that are compatible with network management systems.

Advantages:

- Host-based IDSs can detect attacks that are not detectable by a network-based IDS since they have a view of events 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 monitored host.
- Host-based IDSs can operate in switched networks.

ITL Bulletins Via E-Mail

We now offer the option of delivering your ITL Bulletins in ASCII format directly to your e-mail address. To subscribe to this service, send an e-mail message to listproc@nist.gov with the message **subscribe itl-bulletin**, and your proper name, e.g., John Doe. For instructions on using listproc, send a message to listproc@nist.gov with the message **HELP**. To have the bulletin sent to an e-mail address other than the From address, contact the ITL editor at 301-975-2832 or elizabeth.lennon@nist.gov.

Disadvantages:

- The collection mechanisms must usually be installed and maintained on every host to be monitored.
- Since portions of these systems reside on the host being attacked, host-based IDSs may be attacked and disabled by a clever attacker.
- Host-based IDSs are not well suited for detecting network scans of all hosts in a network since the IDS at each host only sees the network packets that the host receives.
- Host-based IDSs often have difficulty detecting and operating in the face of denial-of-service attacks.
- Host-based IDSs use the computing resources of the hosts they are monitoring.

Application-Based IDSs

Application-based IDSs monitor the events transpiring within an application. Often application-based IDSs detect attacks by analyzing the application's log files. By interfacing with an application directly and having significant domain or application knowledge, application-based IDSs are more likely to have a more discerning or fine-grained view of suspicious activity in the application.

Advantages:

- Application-based IDSs can monitor activity at a very fine granularity, which often allows them to track unauthorized activity to individual users.
- Application-based IDSs can often work in encrypted environments, since they interface with the application that may be performing encryption.

Disadvantages:

- Application-based IDSs may be more vulnerable than host-based IDSs to being attacked and disabled since they run as an application on the host they are monitoring.

The distinction between an application-based IDS and a host-based IDS is not always clear, so for the remainder of this bulletin, we will refer to both as host-based IDSs.

IDS Event Analysis Approaches

There are two primary approaches to analyzing events to detect attacks: signature detection and anomaly detection. Signature detection is the primary technique used by most commercial systems; however, anomaly detection is the subject of much research and is used in a limited form by a number of IDSs.

Signature-Based IDSs

Signature-based detection looks for activity that matches a predefined set of events that uniquely describe a known attack. Signature-based IDSs thus must be specifically programmed to detect each known attack. This technique is extremely effective and is the primary method used in commercial products for detecting attacks.

Advantages:

- Signature-based IDSs are very effective at detecting attacks without generating an overwhelming number of false alarms.

Disadvantages:

- Signature-based IDSs must be programmed to detect each attack and thus must be constantly updated with signatures of new attacks.
- Many signature-based IDSs have narrowly defined signatures that prevent them from detecting variants of common attacks.

Anomaly-Based IDSs

Anomaly-based IDSs find attacks by identifying unusual behavior (anomalies) on a host or network. They function on the observation that some attackers behave differently than "normal" users and thus can be detected by systems that identify these differences. Anomaly-based IDSs establish a baseline of normal behavior by profiling particular users or network connections and then statistically measure when monitored activity deviates from the norm. Unfortunately, these IDSs often produce a large number of false alarms, since normal user and

network behavior can vary wildly. Despite this weakness, researchers assert that anomaly-based IDSs are able to detect never-before-seen attacks, unlike signature-based IDSs that rely on analysis of past attacks. Although some commercial IDSs include restricted forms of anomaly detection, few, if any, rely solely on this technology. However, anomaly detection remains an active intrusion detection research area.

Advantages:

- Anomaly-based IDSs detect unusual behavior and thus have the ability to detect attacks without having to be specifically programmed to detect them.

Disadvantages:

- Anomaly detection approaches usually produce a large number of false alarms due to the unpredictable nature of users and networks.
- Anomaly detection approaches often require extensive "training sets" of system event records in order to characterize normal behavior patterns.

IDSs that Automatically Respond to Attacks

Since human administrators are not always available when an attack occurs, some IDSs can be configured to automatically respond to attacks. The simplest form of automated response is active notification. Upon detection of an attack, an IDS may e-mail or page an administrator. A more active response is to stop an attack in progress and then block future accesses by the attacker. Typically, IDSs do not have the ability to block a particular person, but instead block Internet Protocol (IP) addresses from which an attacker is operating. It is very difficult to automatically stop a determined and knowledgeable attacker, but IDSs often can deter expert attackers or stop novice hackers by using the following capabilities:

- Cutting TCP connections by injecting reset packets into the attacker's connections to the target of the attack,

- Reconfiguring routers and firewalls to block packets from the attacker's location (IP address or site),
- Reconfiguring routers and firewalls to block the protocols being used by an attacker, and
- In extreme situations, reconfiguring routers and firewalls to sever all connections using particular network interfaces.

A more aggressive way to respond to an attacker is to launch attacks against or attempt to actively gain information about the attacker's host or site. However, this response can be extremely dangerous for an organization since it may be illegal and cause damage to innocent Internet users. It is even more dangerous to allow IDSs to automatically launch these attacks, but limited automated "strike-back" strategies are sometimes used for critical systems. Obtain legal advice before pursuing any of these options.

Tools that Complement IDSs

Several tools exist that complement IDSs and are often labeled as IDSs by vendors since they perform similar functions. This section discusses these tools and how they can enhance an organization's intrusion detection capability.

Honey Pot and Padded Cell Systems

Several novel additions to the intrusion detection product line recently hit the market and it is important to understand how these products differ from traditional IDSs. *Honey pots* are decoy systems that attempt to lure an attacker away from critical systems. These systems are filled with information that is seemingly valuable but which has actually been fabricated and which would not be accessed by an honest user. Thus, when access to the honey pot is detected, there is a high likelihood that it is an attacker. Monitors and event loggers on the honey pot detect these unauthorized accesses and collect information about an attacker's activities. The purpose of the honey pot is to divert an attacker from accessing critical systems, collect information about the

attacker's activity, and encourage the attacker to stay on the system long enough for administrators to respond.

Padded cells take a different approach. Instead of trying to attract attackers with tempting data, a padded cell waits for a traditional IDS to detect an attacker. The attacker is then seamlessly transferred to a special padded cell host. The attacker may not realize anything has happened, but the attacker is now in a simulated environment where no harm can be caused. Like the honey pot, this simulated environment can be filled with interesting data to convince an attacker that the attack is going according to plan. Padded cells offer unique opportunities to monitor the actions of an attacker. IDS researchers have used padded cell and honey pot systems since the late 1980s, but until recently no commercial products have been available.

Advantages:

- Attackers can be diverted to system targets that they cannot damage.
- Administrators can be given time to decide how to respond to an attacker.
- Attackers' actions can be more easily monitored, with results used to improve system protections.
- Honey pots may be effective at catching insiders who are snooping around a network.

Disadvantages:

- Honey pots and padded cells have not yet been shown to be widely useful security technologies.
- An expert attacker, once diverted into a decoy system, may become angry and launch a more hostile attack against an organization's systems.
- A high level of expertise is needed for administrators and security managers in order to use these systems.
- The legal implications of using such devices are not well defined.

Vulnerability Assessment Tools

Vulnerability assessment tools determine when a network or host is vulnerable to known attacks. Since this activity is related to actually detect-

ing attacks, these tools are sometimes referred to as intrusion detection tools. They come in two varieties: passive and active. Passive vulnerability assessment tools scan the host on which they reside for insecure configurations, software versions known to contain exploitable flaws, and weak passwords. Active assessment tools reside on a single host and scan a network looking for vulnerable hosts. The tool sends a variety of network packets at target hosts, and from the responses, the tool can determine the server and operating system software on each host. In addition, it can identify specific versions of software and determine the presence or absence of security-related patches. The active assessment tool compares this information with a library of software version numbers known to be insecure and determines if the hosts are vulnerable to known attacks.

Limitations of IDSs

Current intrusion detection products have limitations that one must be aware of before undertaking an IDS deployment.

- Despite vendor claims, most IDSs do not scale well as enterprise-wide solutions. The problems include the lack of sufficient integration with other security tools and sophisticated network management systems, the inability of IDSs to assess and visualize enterprise-level threats, and the inability of organizations to investigate the large number of alarms generated by hundreds or thousands of IDS sensors.
- Many IDSs create a large number of false positives that waste administrators' time and may even initiate damaging automated responses.
- While almost all IDSs are marketed as "real time" systems, during heavy network or host activity, an IDS may take several minutes before reporting and automatically responding to an attack.
- IDSs usually cannot detect newly published attacks or variants of existing attacks. This can be a seri-

ous problem as 30-40 new computer attacks are posted on the Web every month. An attacker may simply wait for a new attack to be posted and then quickly penetrate a target network.

- IDSs' automated responses are often ineffective against sophisticated attackers. They usually stop novice hackers but, improperly configured, can hurt a network by interrupting legitimate network traffic.
- IDSs must be monitored by skilled computer security personnel in order to achieve maximum benefits and to understand the significance of what the IDS detects.
- IDS maintenance and monitoring can use a substantial amount of personnel resources.
- Many IDSs are not failsafe; that is, they are not well protected from attack or subversion.
- Many IDSs do not have user interfaces that allow users to spot cooperative or coordinated attacks.
- IDSs cannot be used in isolation, but must be part of a framework of computer security measures. For a list of such measures, see the May 1999 *ITL Bulletin* entitled "Computer Attacks: What They Are and How to Defend Against Them." (See below.)

Deployment of IDSs

Intrusion detection technology is a necessary addition to every large organization's computer security framework. However, given the weaknesses found in some of today's products, and the relatively limited security skill level of most system administrators, careful planning, preparation, prototyping, testing, and specialized training are critical steps for an effective IDS deployment.

NIST suggests performing a thorough requirements analysis, carefully selecting the intrusion detection strategy and solution that is compatible with the organization's network infrastructure, policies, and resource level. Organizations should consider a staged deployment of

IDSs to gain experience and to ascertain how many monitoring and maintenance resources they will require. There is a large variance in the resource requirements for each type of IDS. IDSs require significant preparation and ongoing human interaction. Organizations must have appropriate security policies, plans, and procedures in place so that personnel will know how to react to the many and varied alarms IDSs will produce.

We recommend consideration of a combination of network-based IDSs and host-based IDSs to protect an enterprise-wide network. First deploy network-based IDSs since they are usually the simplest to install and maintain; then follow up by defending critical servers with host-based IDSs. Honey pots should be used judiciously and only by organizations with a highly skilled technical staff that are willing to experiment with leading-edge technology. Padded cells are currently unavailable except as research prototypes.

Deploying Network-Based IDSs

There are many options for placing a network-based IDS and different advantages for each location:

Location: Behind each external firewall

Advantage: Sees attacks that are penetrating the network's perimeter defenses from the outside world.

Location: In front of an external firewall

Advantage: Proves that attacks from the Internet are regularly launched against the network.

Location: On major network backbones

Advantage: Detects unauthorized activity by those within a network and monitors a large amount of a network's traffic.

Location: On critical subnets
 Advantage: Detects attacks on critical resources.

Deploying Host-Based IDSs

Once an organization has deployed network-based IDSs, host-based IDSs can offer an additional level of protection. However, it can be time-consuming to install host-based IDSs on every host in an enterprise. Therefore, it is often preferable to begin by installing host-based IDSs only on critical servers. This placement will decrease the overall deployment costs and allow limited personnel to focus on alarms generated from the most important hosts. Once the operation and maintenance of host-based IDSs is routine, more security-conscious organizations may consider installing host-based IDSs on the majority of their hosts. In this case, purchase host-based systems that have an easy-to-use centralized management and reporting function since the management of alerts from a large set of hosts can be daunting.

The Future of IDSs

The IDS research field has been active since around 1985, but the wide-scale commercial use of IDSs

did not start until about 1996. In 1998, sales of IDS tools reached \$100 million. By all estimates, the market for IDS tools should continue to grow strongly. From this history, it is apparent that while the IDS research field is maturing, commercial IDSs are still in their formative years. Some commercial IDSs have received negative publicity due to their large number of false positives, overwhelming numbers of attack reports, lack of scalability, and lack of integration with enterprise management systems. However, given that commercial IDSs are still evolving rapidly, we believe that these issues will be addressed quickly. The development of IDS products is likely to parallel that of anti-virus software. Early anti-virus software created false alarms on many normal user actions and did not catch all known viruses. However, over time, anti-virus software progressed to its current state, in which few users notice that it is running and they have confidence that it detects all known viruses.

For More Information

Acquiring, deploying, and maintaining an IDS is a complex task. Fortunately, many excellent resources in the form of books and seminars exist

to guide the public on IDS technology. Several free IDS resources are available:

For an overview of IDSs and their capabilities, read the white paper "An Introduction to Intrusion Detection Assessment for System and Network Security Management" at <http://www.icsa.net/services/consortia/intrusion/intrusion.pdf>.

For a survey of commercially available IDSs that allows one to easily compare features, read the "Intrusion Detection System Product Survey" published by the Los Alamos National Laboratory and found at <http://lib-www.lanl.gov/la-pubs/00416750.pdf>.

Information on the computer attacks that IDSs detect can be found in the May 1999 *ITL Bulletin* entitled "Computer Attacks: What They Are and How to Defend Against Them," available at <http://www.nist.gov/itl/lab/bulletns/cslbull1.htm>.

NOTE: Any mention of commercial products is for information only; it does not imply recommendation or endorsement by the National Institute of Standards and Technology nor does it imply that the products mentioned are necessarily the best available for the purpose.

Address Service Requested

Penalty for Private Use \$300

Official Business

Gaithersburg, MD 20899-8900
 100 Bureau Drive, Stop 8900

National Institute of Standards and Technology

U.S. DEPARTMENT OF COMMERCE

PRSRV STD
 POSTAGE & FEES PAID
 NIST
 PERMIT NUMBER G195