

Intrusion Detection Systems

BUYER'S GUIDE

Sponsors:

BindView

Clicknet

CyberSafe

IDC

ISS

NAI

ODS

Tripwire

icsalabs
A DIVISION OF ICSA.NET

Sponsors



Table of Contents

Part I	Background & Technology	5
	<i>Introduction</i>	5
	ICSA.net's Intrusion Detection Systems Consortium	6
	<i>Policy is Key</i>	8
	Management Direction	8
	Policy Development	9
	<i>Technology Overview</i>	13
	Intrusion Detection Systems	14
	Network IDS	15
	Host IDS	17
	File Integrity Checkers	18
	Vulnerability Scanners	20
	Network Vulnerability Scanner	20
	Host Vulnerability Scanner	22
	Technical Summary	23
	<i>Debunking Marketing Hype</i>	24
	What Intrusion Detection Systems and Related Technologies Can and Cannot Do	24
	<i>Guidelines for Selecting Products</i>	26
	Management Issues	26
	Technical Issues	27
	Web Site Issues	27
	IDS Product Decision Trees	28
	<i>Practical Integration Issues</i>	31
	Sensor Placement for a Network IDS	31
	What about outside the firewall?	32
	Host integration for Host IDS	32
	Alarm Configuration	32
	Integration Schedule	33
	<i>Case Studies for Intrusion Detection and Related Products</i>	34
	<i>IDC Report on Market Share and Industry Growth</i>	36
	Introduction	36
	The Market	36
	Conclusion	39
	<i>Appendix A - Incident Response Planning</i>	40
	Policy and Plans	40
	Incident Handling Resources	43
	Sample Incident Reporting Form	44
	<i>Appendix B - Frequently Asked Questions About Intrusion Detection</i>	46
	<i>Appendix C - Glossary</i>	48
	<i>For Further Reading</i>	51
Part II	Product Information	52

ICSA.net industry Guides are available at special quantity rates for use in corporate training programs. For more information, please visit www.icsa.net.

Products or brand names used in this Guide may be trade names or trademarks. Where we believe that there may be proprietary claims to such trade names or trademarks, the name has been used with an initial capital or it has been capitalized in the style used by the name claimant. Regardless of the capitalization used, all such names have been used in an editorial manner without any intent to convey endorsement of or other affiliation with the name claimant. Neither the author nor the publisher intends to express any judgement as to the validity or legal status of any such proprietary claims.

Information contained in this work has been obtained by ICSA.net from sources believed to be reliable. However, neither ICSA.net nor its authors guarantee the accuracy or completeness of any information published herein and neither ICSA.net nor its authors shall be responsible for any errors, omissions, or damages arising out of use of this information. This work is published with the understanding that ICSA.net and its authors are supplying information but are not attempting to render engineering or other professional services. If such services are required, the assistance of appropriate professional help should be sought.

part 1
background
& technology

Introduction

Last year, a home or business in the United States was broken into every 11 minutes. Based on information coming from various response teams, during the last year, a computer was attacked or broken into more than once per second. These statistics may be only the tip of the iceberg as companies develop the ability to identify and track break-ins. The means for identifying and tracking break-ins is called “intrusion detection.”

Intrusion detection systems (IDS), which have long been a topic for theoretical research and development, are gaining mainstream popularity as companies move more of their critical business interactions to the Internet. An intrusion detection system can provide advance knowledge of attacks or intrusion attempts by detecting an intruder’s actions. In this respect, intrusion detection systems are a powerful tool in the organization’s fight to keep its computing resources secure.

This guide will describe the primary categories of intrusion detection technology and provide some guidance on how to select the right tools. Although some may consider IDS tools just another check box in the audit compliance guideline, they should be planned as an integral part of the organizational network. As computer networks become ever more complex, intrusion detection will take on a greater role in the organization.

In his book on the topic, Edward Amoroso defines the term intrusion detection as: “The process of identifying and responding to malicious activity targeted at computing and networking resources.” An intrusion detection system attempts to uncover behavior or configurations that are, indicate, or could lead to malicious activity. Unfortunately, such a broad definition would lead you to believe that nearly anything, including papers describing system hardening, constitute intrusion detection systems. We must further constrain the definition.

An intrusion detection system (IDS) is composed of hardware and software elements that work together to find unexpected events that may indicate an attack will happen, is happening, or has happened. Note that we must think in all three tenses; some products warn in advance that an attack may take place, some warn as they notice an attack in progress, and some warn when they notice the aftereffects of the attack.

Purpose

The purpose of this guide is to provide a well-grounded framework within which readers can make appropriate purchase and implementation decisions. In it, we will:

- Discuss the dependencies and externalities often encountered when implementing an IDS in a production environment
- Describe how the tools are linked to corporate incident handling objectives
- Set expectations for IDS products accordingly
- Propose a process for evaluating and selecting an intrusion detection system. The process begins with policy and moves through evaluation, election, and implementation of various types of intrusion detection tools.
- Provide a good deal of common sense advice about intrusion detection and incident handling
- Examine many of the common myths of the industry
- Offer case studies to describe tools in action

For anyone interested in the intrusion detection market space, the IDC Market Report section in Part II summarizes an extensive market research report from International Data Corp.

With the wealth of information contained here, we hope that this guide will become a means of making IDS an effective part of any information security program.

ICSA.net's Intrusion Detection Systems Consortium

The Intrusion Detection Buyer's Guide is a project produced on behalf of the ICSA.net Intrusion Detection Systems Consortium (IDSC). ICSA.net formed the consortium in 1998 to provide product developers an open forum within which they could work towards common goals. These goals include educating end-users, influencing industry standards, and maintaining product and marketing integrity.

Members meet on a quarterly basis and participate in ongoing discussions and cooperative projects such as this white paper. Membership is open to any commercial developer of intrusion detection and vulnerability assessment products. See <http://www.icsa.net> for details.

As agreed to by the IDSC members,

“The mission of the IDSC is to facilitate the adoption of intrusion detection products by defining common terminology, increasing market awareness, maintaining product integrity and influencing industry standards.”

IDSC members as of December 1, 1999 include:

AXENT Technologies, Inc.
Bindview Corporation
Clicknet Software
Computer Associates, Inc.
CyberSafe Corporation
IBM
Internet Security Systems, Inc
Network Associates, Inc.
Network Security Wizards
ODS Networks
Qwest Communications International Inc.
Tripwire, Inc.

Policy Is Key

An IDS is only useful in the context of a business or operations process. By defining the policy first, we can increase our chances for fielding the right IDS architecture.

The first consideration when contemplating an IDS is, “how will we respond in the event of an intrusion?” With the advent of modern intrusion detection systems, we are able to measure the attack-resistance of our networks every microsecond. We can know within minutes if we are being probed, hacked, infected, or misused. And we can be alerted with e-mail, pages, or phone calls day or night if any one of the 40 million members of the international Internet decides to make our network the next target of unauthorized access.

We say “can” because that is the extreme far end of the intrusion detection spectrum. Too often, people install all of the features that “can” be used and forget to configure tools to meet their own requirements for what “must” be done. This Chapter is intended to help consumers of IDS technology decide what “must” be done in order to select tools that meet those requirements.

Management Direction

Prior to defining a policy for intrusion detection, managers should consider talking to executives of the corporation and determining what level of interest they have in intrusion detection. Surprisingly, many executives believe that current systems provide ample information for detecting and tracking incidents on corporate networks. Many others believe that any monitoring on a corporate network amounts to “snooping” and should be discouraged. In this case, it might be better to ask executives if they would rather find out about an intrusion from employees or from the local newspaper. Both beliefs should be addressed prior to defining the policy.

Managers should also discuss the idea of apprehending intruders with executive management. If executives believe that prosecution of computer criminals is unnecessary or inappropriate, they may feel that the use of an IDS may be unnecessary. In such a situation, intrusion detection may still be valuable for security awareness, but incident response becomes simple. Every time an intrusion attempt is detected from outside the organization, simply cut off communications with the source of the attack. If an attack is detected from inside the organization, identify the culprit and turn the issue over to Human Resources.

Policy Development

The idea of an intrusion is usually associated with some form of response. Therefore, rather than preparing an intrusion detection policy, per se, organizations should integrate the need for an IDS into an overall incident handling policy. The elements of a good incident handling policy have the following components:

- Statement of scope
- Acceptable computer and network use
- Detection and Reporting requirements
- Responsibility for responding to incidents
- Responsibility for managing incident response

Statement of scope

The incident handling policy should define what an incident is and what types of incidents are covered under the policy. If the organization has a separate virus intervention policy, that may be excluded from the incident handling policy. Some organizations integrate the Emergency Operations portion of the Disaster Recovery policy into the incident handling policy.

It is important to define what an incident is. One example of a definition is: “The act of violating one of the information protection policies or standards adopted by the corporation.” Although this is a broad definition, it provides a basis for a more complete definition later. Security managers must define what an incident is and what forms of incidents should yield a response. Typically, things like unauthorized network access, unauthorized changes to system hardware or software, denial of service, and unauthorized release of proprietary data via computer are considered incidents.

Next, the scope should define whether incidents are limited to actions that affect the corporate network, or networks of subsidiaries. This provides additional information to intrusion detection system designers.

Finally, scope should address the difference between incidents that originate inside the organization (internal incidents) or incidents that originate outside the organization (external incidents). In certain cases, an incident may be clearly defined as either originating from an internal or an external source. Frequently, external incidents have the potential for involving law enforcement. Therefore, if there is any need to differentiate for the purposes of a policy, it should be explained here.

Acceptable computer and network use

If corporate policies have not defined acceptable use elsewhere, the incident response policy should do so. Acceptable use defines types of activity on the network that are legal, but may be unauthorized. Moreover, many of the laws governing computer crime in the United States base their definition of crime on the concept of “unauthorized access.” Without clear policies regarding unauthorized access, intruders may be more difficult to prosecute.

The corporation should describe authorized types of computer use regarding the Internet. Most often, human resources will be able to assist in identifying web browsing and file downloads that should be considered unauthorized. These policies can assist the IDS designer in defining criteria for alarms.

Detection and Reporting Requirements

Detection and Reporting will define the character of an Incident Response plan more than any other section. Effectively written requirements can reduce the number of incidents experienced by the organization and improve the handling of each, and will directly affect the selection and use of any IDS products. Whether detection is automatic or manual, it must follow the other criteria in the incident handling policy.

Manual detection methods usually involve users who notice abnormal activity. Rather than define a specific approach for each user, management should define and broadcast a simple method for contacting the incident response team. The Institute for Internal Auditors performed a survey that found over 50% of internal bank frauds were noticed and reported by employees via routine fraud reporting mechanisms. Any organization concerned about fraud should examine processes for handling reports from users. The process should include some level of technical assistance to determine if the problem is actually an incident or merely a system malfunction.

Outside the realm of manual detection, we have automated detection. Automated detection may come from an IDS system or from some reporting mechanism on another platform. Regardless of the source, the basic process for handling these events must be specified in the policy document.

The detection process may involve more than one kind of technology and the policy for detection may differ, depending on the type of system used. For example, if the system is signature based, the response process should immediately investigate the reason for the alarm, with a view towards unauthorized use of a system. If the system is integrity or statistically based, the response may be an issue for configuration management, rather than for the security group.

Rather than defining a single policy for intrusion detection, we recommend evaluating the information found in the Technology Overview of this Guide. After reviewing the kinds of technology available and the problems they attempt to solve, the organization should be able to determine the most appropriate detection policy for its networks.

Reporting should cover how incidents are brought to the attention of the incident response team. The reporting process should be fast and simple. If end users are expected to report incidents, they should know how and where to report the incident at all times.

Any organization that can effectively process an initial report and route it correctly to the team is a candidate for the central notification point. Some organizations route incident reports through the help desk. If the help desk is staffed for extended hours support, this can be an effective method of entering the first report. In other organizations, the security guard desk is tasked with receiving initial reports. Still others use a pager system so that employees can contact a member of the team directly. Regardless of where the incident is first reported, it should be reported properly. This often requires an interview with the user reporting the incident.

To obtain adequate information for each incident report, a form is useful. The form should be completed by the recipient of the initial incident report. It should include the name of the person noticing the problem (if applicable) along with any available tracking information (program affected, Network Address of machine, time, date, behavior during the incident, warning messages, etc.). Appendix A on Incident Handling contains a copy of the CERT Incident Reporting Form.

Responsibility for responding to incidents

Incidents frequently involve more than one organization in a company. Therefore, the roles and responsibilities for responding should be defined early. In some organizations, this involves policy and routine coordination meetings. In others, the policy is negotiated so that each member organization understands their role.

For the purposes of an incident handling policy, the following organizations should be involved:

- Information Security
 - Networking (routers, gateways, WAN systems, DHCP services)
 - Systems Administration (OS management and maintenance)
 - Legal
 - Physical or Industrial Security
 - Public or Investor Relations
 - Internal Audit
 - Human Resources
-

Not all of these organizations will play a role in every incident, but each will play a role in developing and approving the policy. The policy should assign management responsibility to each organization for any task that primarily involves their organization. Moreover, the interface between each organization should be defined so that communications can be handled efficiently.

At this point, the corporation should include an important designation of responsibility. By specifically assigning responsibility for incident response, management can make the response team agents of the corporation. This is a necessary step to protect members of the team from legal trouble after a response. Virus writers and hackers have often copyrighted their works in a vain attempt to deter investigators. Even though copyrights do not apply to illegal activities, there may be a time when computer criminals attempt to sue the investigators who caught them. By asserting agency for incident response personnel, the corporation is clearly committing to a defense for the people who help it respond to incidents. Contractors who are authorized to participate in incident response activities should be covered in the same way.

Responsibility for managing incident response

In any emergency activity, someone should be designated as the responsible manager. This individual will coordinate actions during the incident and will coordinate reporting as well. When assigning management responsibility it is important to remember that most incidents fall outside of working hours. Therefore, the responsible manager may be the first person available via telephone rather than the same person for every incident. Regardless of who is responsible, the response manager must ensure that all appropriate organizations are informed of progress in a timely fashion.

Technology Overview

The goal of an intrusion detection system is to provide an indication of a potential or real attack. An attack or intrusion is a transient event, whereas a vulnerability represents an exposure, which carries the potential for an attack or intrusion. The difference between an attack and a vulnerability, then, is that an attack exists at a particular time, while a vulnerability exists independently of the time of observation. Another way to think of this is that an attack is an attempt to exploit a vulnerability (or, in some cases, a perceived vulnerability). This leads us to categorize various types of intrusion detection systems.

Figure 1 demonstrates the difference between vulnerability scanners and intrusion detection systems. Vulnerability scanning is less time critical than intrusion detection. Subsequently, the deployment of each technology can vary inside organizations. Figure 2 maps IDS types onto the technology landscape of Figure 1.

There are five different categories of IDS covered in this guide. Not all of these categories represent “classical intrusion detection” but they play a role in the overall goal of detecting or preventing intrusions on a corporate network. The categories are:

- Network Based Intrusion Detection System
- Host Based Intrusion Detection System
- File Integrity Checker
- Network Vulnerability Scanner
- Host Vulnerability Scanner

TECHNOLOGY LANDSCAPE

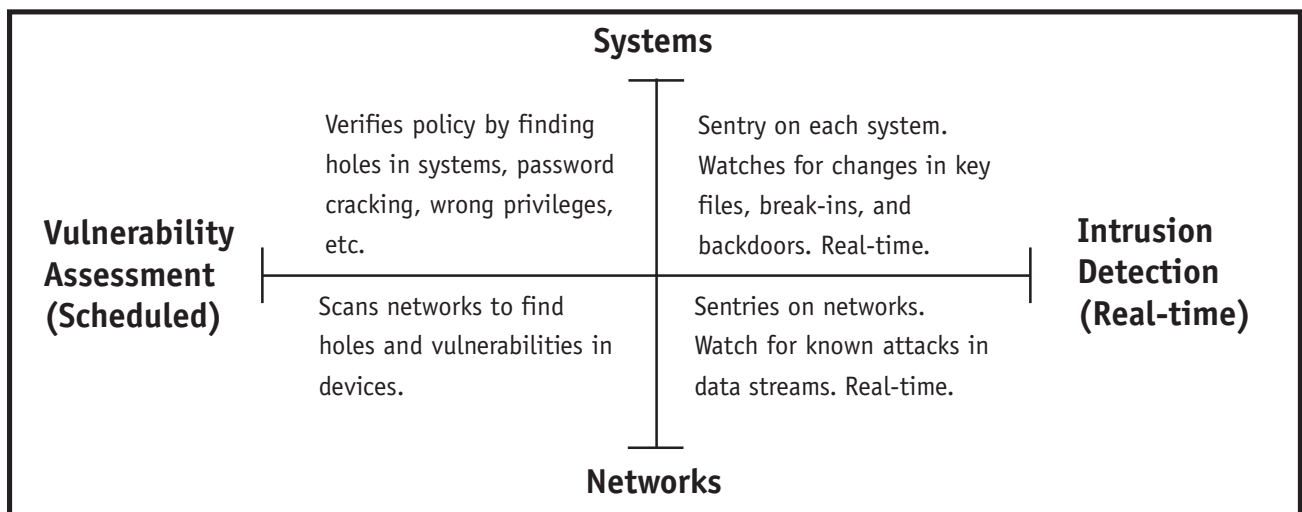


Figure 1. As shown here, IDS products can be categorized as either preventive or responsive. They can also be categorized according to their emphasis on either system or network scanning.

The IDS tools covered in this guide fall into two technology categories: intrusion detection systems and vulnerability scanners. We can further decompose these two categories into host and network-based systems. As shown in Figure 2, vulnerability scanners can be run at any time because we assume that a vulnerability exists until repaired. An intrusion, on the other hand, exploits a specific vulnerability and must be detected as soon as possible after it starts. For this reason, intrusion detection tools must run more frequently than vulnerability scanners. This is why most IDS vendors attempt to make their intrusion detection tools work in real-time.

TECHNOLOGY LANDSCAPE

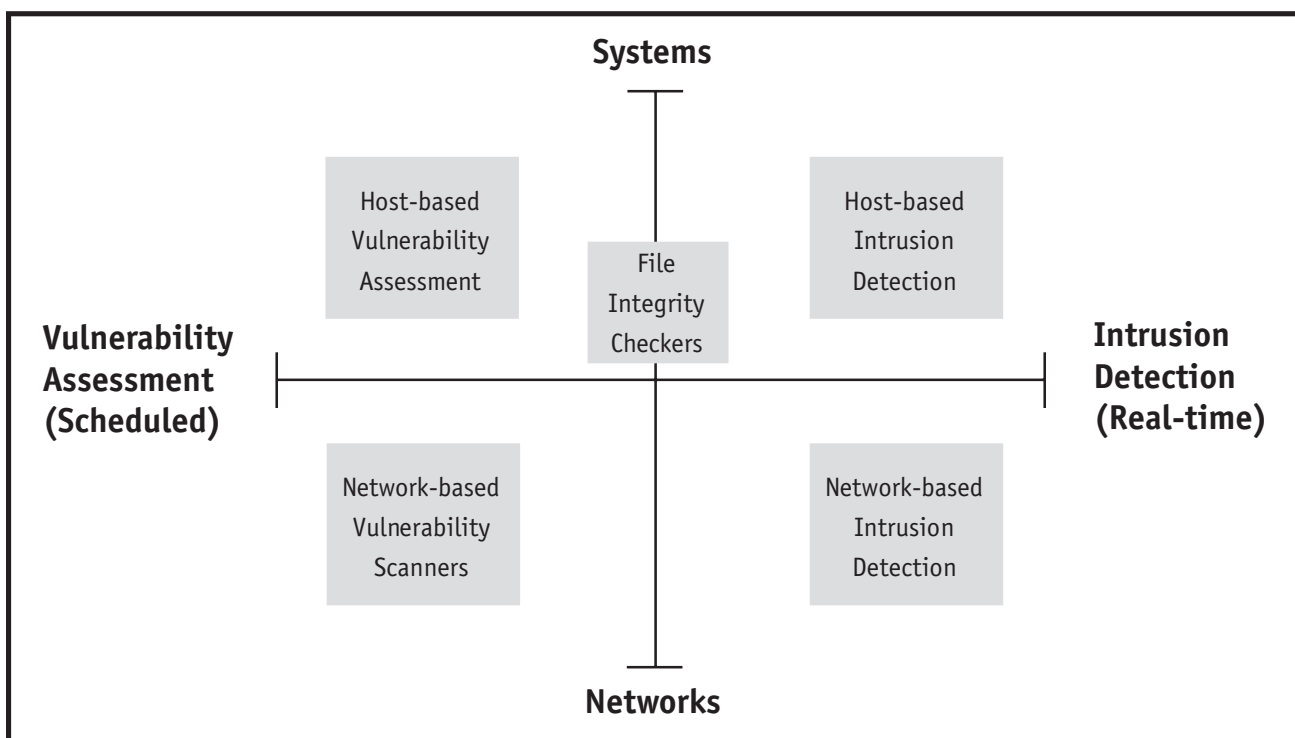


Figure 2. This maps IDS technologies to the matrix, showing the difference in focus for each technology.

Intrusion Detection Systems

An intrusion detection system (IDS) examines system or network activity to find possible intrusions or attacks. Intrusion detection systems are either network-based or host-based; vendors are only beginning to integrate the two technologies.

Network based intrusion detection systems are most common, and examine passing network traffic for signs of intrusion (See Figure 3). Host-based systems look at user and process activity on the local machine for signs of intrusion. Since each type has specific strengths and weaknesses, we will cover each type of tool in the following sections.

You might ask, “how does an IDS determine what is suspicious?” This is a good question. There are, generally speaking, three kinds of commercially available analysis engines:

- Event or Signature-based Analysis
- Statistical Analysis
- Adaptive Systems

The event, or signature-based, systems function much like the anti-virus software with which most people are familiar. The vendor produces a list of patterns that it deems to be suspicious or indicative of an attack; the IDS merely scans the environment looking for a match to the known patterns. The IDS can then respond by taking a user-defined action, sending an alert, or performing additional logging. This is the most common kind of intrusion detection system.

A statistical analysis system builds statistical models of the environment, such as the average length of a telnet session, then looks for deviations from “normal”. After over 10 years of government research, some products are just beginning to incorporate this technology into marketable products.

The adaptive systems start with generalized rules for the environment, then learn, or adapt to, local conditions that would otherwise be unusual. After the initial learning period, the system understands how people interact with the environment, and then warn operators about unusual activities. There is a considerable amount of active research in this area.

You should keep in mind that any IDS will both miss some kinds of suspicious activity (false negatives) and signal alarms when there is nothing wrong (false positives). This is why organizations must have a strong human process that interacts with the IDS to evaluate the operating environment. The machine intelligence of most intrusion detection systems is still evolving, though current research is working to improve this.

Remember, when reading these sections, that the discussion deals with generalizations. Each specific product has strengths and weaknesses, and some tools use multiple technologies to accomplish their goals. For example, a system may use both signature and statistical logic.

Network IDS

The network IDS usually has two logical components: the sensor and the management station. The sensor sits on a network segment, monitoring it for suspicious traffic. The management station receives alarms from the sensor(s) and displays them to an operator.

The sensors are usually dedicated systems that exist only to monitor the network. They have an network interface in promiscuous mode, which means they receive all network traffic, not just that destined for their IP address, and they capture passing network traffic for analysis. If they detect something that looks unusual, they pass it back to the analysis station.

The analysis station can display the alarms or do additional analysis. Some displays are simply an interface to a network management tool, like HP Openview, but some are custom GUIs designed to help the operator analyze the problem.

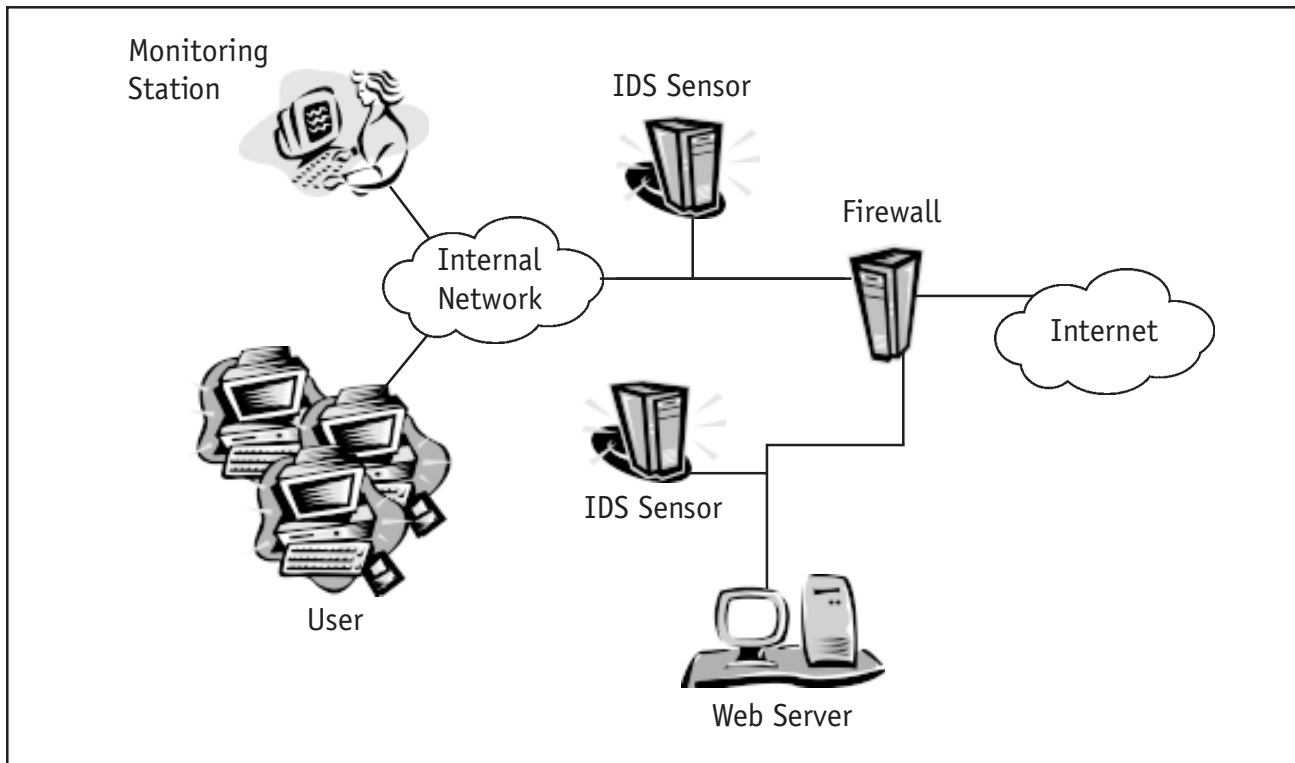


Figure 3. This diagram shows placement of a traditional network based IDS with two sensors on separate network segments that communicate with a monitoring station on the internal network.

Strengths

The network intrusion detection systems can detect some of the attacks that use the network. They are good for detecting access without authority or some kinds of access in excess of authority.

A network-based IDS does not require modification of production servers or hosts. This is an advantage because production servers frequently have close operating tolerances for CPU, I/O, and disk capacity; installing additional software may exceed the systems capacities.

The IDS is not on a critical path for any production services or processes because a network-based IDS does not act as a router or other critical device. System failure does not have a significant impact on the business. A side benefit of this is that you are likely to encounter less resistance from other people within your organization; the risk to existing critical processes is lower with a network system than with a host system.

Network-based IDS systems tend to be more self-contained than host-based systems. They run on a dedicated system that is simple to install; merely unbox the device, do some remedial configuration, and plug it into your network in a location that permits it to monitor sensitive traffic.

Weaknesses

A network based IDS, on the other hand, only examines network traffic on the segment to which it is directly connected, but it cannot detect an attack that travels through a different network segment. This problem—localized vision—is particularly endemic in a switched ethernet environment. The problem may require that an organization purchase many sensors in order to meet their network coverage goals. Since each sensor costs money, broad coverage with network IDS sensors can become prohibitively expensive.

Network intrusion detection systems tend to use signature analysis in order to meet performance requirements. This will detect common programmed attacks from external sources, but it is inadequate for detecting more complex information threats. This requires a more robust ability to examine the environment.

A network intrusion detection system may need to communicate large volumes of data back to the central analysis system. Sometimes that means that any given monitored packet generates a larger amount of analysis traffic. Many such systems use aggressive data-reduction processes to reduce the amount of communicated traffic. They also push much of the decision-making processes out into the sensor itself and use the central station as a status display or communications center, rather than for actual analysis. The disadvantage of this is that it provides very little coordination amongst sensors; any given sensor is unaware that another has detected an attack. Such a system cannot normally detect synergistic or complex attacks.

A network-based IDS may have a difficult time handling attacks within encrypted sessions. Fortunately, there are very few attacks that take place within an encrypted traffic session, other than attacks against weak web servers. This will become more of an issue as organizations transition to IPv6.

Host IDS

The host-based IDS looks for signs of intrusion on the local host system. These frequently use the host system's audit and logging mechanism as a source of information for analysis. They look for unusual activity that is confined to the local host such as logins, improper file access, unapproved privilege escalation, or alterations on system privileges. This IDS architecture generally uses rule-based engines for analyzing activity; an example of such a rule might be, "superuser privilege can only be attained through the su command." Therefore successive login attempts to the *root* account might be considered an attack.

Strengths

A host-based IDS can be an extremely powerful tool for analyzing a possible attack. For example, it can sometimes tell exactly what the attacker did, which commands he ran, what files he opened, and what system calls he executed, rather than just a rather vague accusation that an he attempted to execute a dangerous command. A host-based IDS usually provides much more detailed and relevant information than a network-based IDS.

Host-based systems tend to have lower false positive rates than do network-based systems. This happens because the range of commands executed on a specific host are much more focused than the types of traffic flowing across a network. This property can reduce the complexity of host based analysis engines.

Host based systems can be used in environments where broad intrusion detection is not needed, or where the bandwidth is not available for sensor-to-analysis communications. Host based systems can be completely self-contained. This also allows host-based systems to run, in some cases, from read-only media; this prevents the attackers from disabling the IDS.

Finally, a host-based system may be less risky to configure with an active response, such as terminating a service or logging off an offending user. A host based system is more difficult to spoof into restricting access from legitimate sources.

Weaknesses

Host-based systems require installation on the particular device that you wish to protect. If, for example, you have a human resources server, and you want to protect it, you have to install the IDS on that server. As mentioned earlier, this can pose capacity problems. In some cases, it can even pose security problems, since the security personnel may not ordinarily have access to the server in question.

Another problem associated with host-based systems is that they tend to rely on the innate logging and monitoring capabilities of the server. If the server isn't configured to do adequate logging and monitoring, you have to change the configuration of, possibly, a production machine, which is a tremendous change management problem. How do you adequately predict the results of adding that logging capability?

Host-based systems are relatively expensive. Many organizations do not have the financial resources to protect entire network segments using host-based systems. These organizations must very carefully choose which systems to protect. This can leave wide gaps in ID coverage, because, for example, an attacker on a neighboring but unprotected system can sniff authentication information or other sensitive material from the network.

Finally, host-based systems suffer, to an even greater degree, from local-vision restrictions. They are almost totally ignorant of the network environment. Thus, the analysis time required to evaluate damage from a potential intrusion increases linearly with the number of protected hosts; i.e., if a human takes time t to investigate an incident on one system, it will take $2t$ to investigate two systems, $3t$ to investigate three systems, and so forth.

File Integrity Checkers

A file integrity checker examines the files on a computer to determine whether they have been altered since the last time the integrity checker was run. The integrity checker keeps a database of hash values for each file. Each time the checker runs, it recalculates the hash value and compares it to the stored value. If the hash values are different, the file has changed. If the values are not different, the file has not changed.

A hash function is a mathematical process for reducing the sequence of bytes in a file to a fixed-length number. The same file will always produce the same hash value and any change in the file is supposed to produce a different value. Unlike encryption, a hash is a one-way function; you cannot produce the original source file from the hash value.

Some hashes are more secure than others. Very secure hashes, which meet special mathematical requirements, are called cryptographically secure hashes. One of the requirements for a cryptographically secure hash is that it is very hard (the technical term is computationally infeasible) to calculate a collision— two inputs that hash to the same number. That means that even if an attacker changes a file, he cannot change the file in a way that fools the integrity checker into thinking the file has not been changed.

Strengths

It is computationally infeasible to defeat the mathematics in an integrity checker. This makes it a very, very strong tool for detecting changes to files on a computer. It is so strong, in fact, that it is one of the most important tools you can use to detect misuse of computer systems.

Integrity checkers can be configured to watch everything on the system, or only important files. They are extremely flexible.

Once attackers compromise a system, they like to do two things. First, they like to cover up their tracks, which means that they will alter system binaries, libraries, or log files to hide the fact that they are or have been on the system. Second, they will make changes to ensure they will have continued access to the system. A properly configured file integrity checker will detect both activities.

Weaknesses

Integrity checkers rely on data stored on the local computers. Like log files, this data (the database of hash values) is vulnerable to modification on the system. Say an attacker gains superuser privileges on your system. The attacker can find the integrity checker, make any hostile changes to the system, and re-run the integrity checker to recreate the database of hash values. When the system administrator runs the tool, it will not report any changes to the system.

One way around this problem is to keep your database on read-only media such as a writable CD. This works for mostly stable and unchanging systems, but for most production machines it is incredibly inconvenient.

The integrity checker must be configured for each system. Usually this is a time-consuming and complicated task. If the operating system is not good at enforcing system integrity, installation becomes even more complex.

Once the checker is configured, it must be run frequently. Depending on the operating system, simple changes or normal operation can report from tens to thousands of changes. For example, an integrity checker run just before upgrading MS-Outlook on a Windows NT system, then just after the installation, reported over 1800 changes.

Finally, integrity checkers consume a considerable number of system resources.” They can chew CPU, memory, and disk space. Many administrators will not want to run integrity checkers frequently. This limits their functionality, because a checker run once a month will report so many changes that a real attack has a good chance of going unnoticed.

Vulnerability Scanners

A vulnerability scanner differs from an intrusion detection system, as mentioned earlier, in that the vulnerability scanner looks for static configurations and the IDS looks for transient misuse or abnormalities. A vulnerability scanner may look for a known NFS vulnerability by examining the available services and configuration on a remote system. An IDS, handling the same vulnerability, would only report the existence of the vulnerability when an attacker attempted to exploit it.

Vulnerability scanners, whether network or host scanners, give the organization the opportunity to fix problems before they arise, rather than reacting to an intrusion or misuse that is already in progress. An intrusion detection system detects intrusions in progress, while a vulnerability scanner allows the organization to prevent the intrusion in the first place. Vulnerability scanners may be helpful in organizations without a good incident response capability.

Network Vulnerability Scanner

A network vulnerability scanner operates remotely by examining the network interface on a remote system. It will look for vulnerable services running on that remote machine, and report on a possible vulnerability. For example, it is well-known that `rex` is a weak service; a network vulnerability scanner will attempt to connect to the `rex` service on the target system. If the connection succeeds, the scanner will report a `rex` vulnerability.

Since a network vulnerability scanner can be run from a single machine on the network, it can be installed without impacting the configuration management of other machines. Frequently, these scanners are used by auditors and security groups because they can provide an “outsider’s view” of security holes in a computer or network.

Strengths

Network vulnerability scanners can report on a variety of target architectures. Some work with routers, some with Unix systems, and some with NT or other Windows platforms.

Network vulnerability scanners are, in general, very easy to install and begin using (from a technical standpoint). Unlike host-based systems, which usually require software installation or reconfiguration, a network-based system can be dropped into place on a network. Simply plug the interface into the switch and boot up the machine.

The GUI with a network system tends to be quite intuitive, which means that junior personnel can monitor the system and call more senior analysts if something unusual crops up.

Weaknesses

Network vulnerability scanners are almost exclusively signature-based systems. Like a signature-based IDS, a signature-based vulnerability scanner can only detect those vulnerabilities it is programmed to recognize. If a new vulnerability comes into play, as they frequently do, there is a window of opportunity for the attacker before the vendor updates the signatures (and the customer downloads and installs the new signatures). If the vulnerability remains closely held, systems can remain vulnerable to attack for long periods of time.

If customers are as negligent with their vulnerability scanner signatures as history shows they are with virus signatures, then many organizations will be vulnerable to attack, even though they run vulnerability scanners at regular intervals. A recent analysis showed that 90 percent of web servers running IIS are still vulnerable to a well-documented and very serious security vulnerability, for which the vendor has produced a patch and a security advisory. A vulnerability scanner can only point out possible problems; the organization must still fix them.

Another potential problem with a network vulnerability scanner is that the output still requires skilled interpretation. Every environment has different operating requirements and different security vulnerabilities. In fact, the concept of “vulnerability” embodies other loosely-defined concepts, such as risk, threat, acceptability, and expected attacker skills. Since each of these varies with each organization, the degree to which a particular configuration represents a “vulnerability” also varies with each organization.

When the vulnerability scanner reports a particular vulnerability, the organization’s network or operations personnel must evaluate that report within the context of the organization’s operating environment. The vulnerability may not pose an unacceptable risk in that organization’s environment, or the risk may be forced upon the organization by business requirements. This may seem silly in the context of a security discussion, but in the real-world security concerns frequently fall prey to business justifications.

Vulnerability scanners have been known to take down the target system. Some IP implementations are not robust enough to handle many simultaneous connections, or IP packets with unusual flag combinations. The traffic generated by an aggressive port scan, for example, can sometimes crash a machine.

Finally, network vulnerability scanners tend to contain a huge amount of vulnerability data. If anyone ever breaks into the scanner system, compromise of most other machines on the network can become child’s play. Protect the scanner to prevent unauthorized use of scanning data.

Host Vulnerability Scanner

A host vulnerability scanner differs from a network vulnerability scanner in that it is confined entirely to the local operating system. A network vulnerability scanner requires the target machine be accessible from the network in order for it to operate; a host vulnerability scanner does not.

Host vulnerability scanners are software packages that are installed on particular operating systems. Once the software is installed it can be configured to run at any time of the day or night. Usually the scans performed by this type of tool are scheduled to run at a low priority to reduce the impact of the scan on production work.

Strengths

Host vulnerability scanners tend to be much more tuned and accurate for a given operating system. They can frequently tell the user which patches to apply to fix identified vulnerabilities, while network scanners sometimes only provide general guidelines. When considering which product to purchase, look at sample reports for your specific operating systems to determine how much information is contained in the reports. The depth and accuracy of reporting should be a selection criterion.

Host vulnerability scanners do not consume network bandwidth when they run. All processing is restricted to the local host system.

Host vulnerability scanners are not as likely as network scanners to cause the IP stack on the target system to hang. Some operating systems have weak or poorly implemented IP stacks associated with their network interfaces; sending a high volume of traffic (as a portscanner would do) or unusual TCP header flags may cause the interface to hang, which would necessitate a reboot.

Host vulnerability scanners are less likely to be used against you by a successful intruder. A hacker who breaks a system and finds a network vulnerability scanner, or reports from such a scanner, is likely to use the tool or the reports to attack other systems within your organization. A host based tool provides much less useful information for extending an attack; i.e., it is better compartmentalized than a network-based tool.

Weaknesses

Host vulnerability scanners are, again, signature based. They look for known-dangerous system configurations and report on a cookbook approach to mitigating those particular threats. Local system procedures and operating requirements may require flexibility in finding and applying solutions to any given vulnerability.

Installing host vulnerability scanners requires the cooperation of system administrators. Since the software usually runs with privilege, the system administrator for each machine should agree to the purpose and configuration of the tool. In many organizations, this can be a difficult coordination task.

As with any host-based system, attackers can modify the vulnerability scanner so it does not report on the vulnerabilities the attackers wish to exploit. There is an important corollary here—the vulnerability scanner cannot protect a system that was compromised when the scanner was installed.

Technical Summary

The various types of IDS tools provide a wide range of capabilities for identifying or analyzing potential threats to computer networks. The type of analysis and the time between analysis and intervention will control decisions regarding type of tool to select. Certain organizations may find that they have a need for each type of product in a complementary security implementation on large networks.

The common thread throughout these five categories is “human capabilities.” The organization must have expertise in-house (or hire a service) that can evaluate the legitimacy of any given intrusion warning. Any of these products can produce significant false positive rates; this requires investigative expertise within the organization.

Debunking Marketing Hype

What Intrusion Detection Systems and Related Technologies Can and Cannot Do

Every new market suffers from exaggeration and misconception. Some of the claims made in marketing materials are reasonable and others are misleading. Herewith, a primer on how to read intrusion detection marketing literature.

Realistic benefits

1. They can lend a greater degree of integrity to the rest of your security infrastructure. Intrusion detection systems provide additional layers of protection to a secured system. The strategy of a system attacker will often include attacking or otherwise nullifying security devices protecting the intended target. Intrusion detection systems can recognize these first hallmarks of attack, and potentially respond to them, mitigating damage. In addition, when these devices fail, due to configuration, attack, or user error, intrusion detection systems can recognize the problem and notify the right people.
 2. They can make sense of often obtuse system information sources, telling you what's really happening on your systems. Operating system audit trails and other system logs are a treasure trove of information about what's going on internal to your systems. They are also often incomprehensible. Intrusion detection systems allow administrators and managers to tune, organize, and comprehend what these information sources tell them, often revealing problems before loss occurs.
 3. They can trace user activity from the point of entry to point of exit or impact. In the unlikely event that an intruder gets past a perimeter defense device, such as a firewall, an IDS provides a way to catch his actions. In addition, an IDS can detect the actions of a "bad guy" already on the inside—an attack from an insider or via a previously unknown entry path to the network.
 4. They can recognize and report alterations to critical system and data files. File integrity assessment tools utilize strong cryptographic checksums to render these files tamper-evident and, in the case of a problem, quickly ascertain the extent of damage.
 5. They can spot errors of your system configuration that have security implications, sometimes correcting them if the user wishes. Vulnerability assessment products allow consistent auditing and diagnosis of system configuration settings that might cause security problems.
-

6. They can recognize when your system appears to be vulnerable to particular attacks. Vulnerability assessment products also allow the administrator of a system to quickly determine what attacks should be of concern.

Unrealistic expectations

1. They are not magic bullets. Security is a complex area with myriad possibilities and difficulties. In networks, it is also a “weakest link” phenomenon—i.e., it only takes one vulnerability on one machine to allow an adversary to gain entry and potentially wreak havoc on the entire network. The time it takes for this to occur is also minuscule. There are no magic solutions to network security problems, and intrusion detection products are no exception to this rule. However, as part of a comprehensive security management they can play a vital role in protecting your systems.
 2. They cannot compensate for weak identification and authentication mechanisms. We must still rely on strong identification and authentication of users. A security infrastructure that includes strong I&A and intrusion detection is stronger than one containing only one or the other.
 3. They cannot conduct investigations of attacks without human intervention. One must perform incident handling. One must investigate the attacks, determine, where possible, the responsible party, then diagnose and correct the vulnerability that allowed the problem to occur, reporting the attack and particulars to authorities where required.
 4. They cannot intuit the contents of your organizational security policy. Intrusion detection expert systems increase in value when they are allowed to function as both hacker/burglar alarms and policy-compliance engines.
 5. They cannot compensate for problems in the quality or integrity of information the system provides. In other words, “garbage in garbage out” still applies.
 6. They cannot analyze all of the traffic on a busy network. Network-based intrusion detection is capable of monitoring traffic of a network, but only to a point. First, given the vantage-point of network-based intrusion detection sources that rely on network adapters set to promiscuous mode, not all packets are visible to the systems. Second, as traffic levels rise, the associated processing load required to keep up becomes prohibitive and the analysis engine either falls behind or fails.
 7. They cannot always deal with problems involving packet-level attacks. There are weaknesses in packet-capture-based network intrusion detection systems. The heart of the vulnerabilities involves the difference between the IDS’s interpretation of the outcome of a network transaction (based on its reconstruction of the network session) and the destination node for that network session’s actual handling of the transaction. Therefore, a knowledgeable adversary can send series of fragmented and otherwise doctored packets that elude detection, but launch attacks on the destination node. Worse yet, an adversary can use this sort of packet manipulation to accomplish a denial of service attack on the IDS itself by overflowing memory allocated for incoming packet queues.
-

Guidelines for Selecting Products

Technical capabilities are certainly an important part of the purchase decision, but they are not the only part. There are a number of organizational and environmental considerations when fielding an IDS. Many of these issues answer the question “Should we field an intrusion detection system?” which should be asked and reasonably answered before the question, “Which IDS should we field?”

This section provides three separate decision trees to assist with the decision making process. These trees are intended to provide basic guidance on the decision making process for an IDS purchase. Begin by examining some of the decision points on the Management tree. Then proceed to the Technical tree to examine some of the issues related to use of an IDS on an internal network. Finally, if you have a mission critical web server we recommend evaluating Web Site issues in Step 3. You may find that you require more than one type of IDS product. If so, this section may help you decide which types of IDS meet the majority of your information security needs.

Management Issues

The first step in evaluating an intrusion detection project is to determine the extent of the organization's support for this project. Remember that the purpose of an intrusion detection project is not just to detect possible intrusions, but to “do something” about them. Therefore, intrusion detection must be considered as part of a larger incident response process. Management must decide what type of response is appropriate and how to resolve any identified incidents.

In addition, incident response to an intrusion can involve a significant number of legal, public affairs, networking, and systems administration personnel, as well as lost productivity and system downtime. The organization should fully understand these issues before embarking on an intrusion detection project. The extent of organizational support will have an effect on the final product selection; systems that are hard to implement and field will require more management patience than systems that are essentially drop-in boxes. The latter, however, will require better human analysis efforts to interpret and detect false positives.

The relative strength of configuration management processes within the organization will directly affect both the selection and efficiency of any host-based system and, to a lesser degree, any network-based system. In some organizations, any user on the system can alter the system's configuration. They may, for example, alter the `/etc/hosts.equiv` file, which could appear to be a hostile act. Once the change is noted, it must be investigated before the organization can continue to have faith in the system. Organizations with strong CM processes, however, can feel comfortable that any unrecognized changes are highly suspicious.

Technical Issues

Technical issues abound when installing an intrusion detection system. Typically, the most serious technical issue is dealing with the additional network traffic that can be generated by an IDS. Another issue is creating a process for interpreting results from the IDS. Organizations must decide how to handle the interpreted results and how to integrate those results into a plan for response and correction.

The first step in system planning is to identify critical information resources, both for storage and communications. In other words, make a list of what you are trying to protect. In some cases, this might be a web server in your perimeter DMZ. In other cases, it might be the firewall configuration that permits traffic to the DMZ. In still others, it might be the mainframe or back end database systems that store critical customer information. It might even be, to some extent, your company's image (if you are worried about defaced web sites). As you build the list of critical resources, diagram the systems and networks to provide a visual representation of the environment and placement of IDS tools. This diagram should be compared with official network maps to ensure consistency with the overall network architecture.

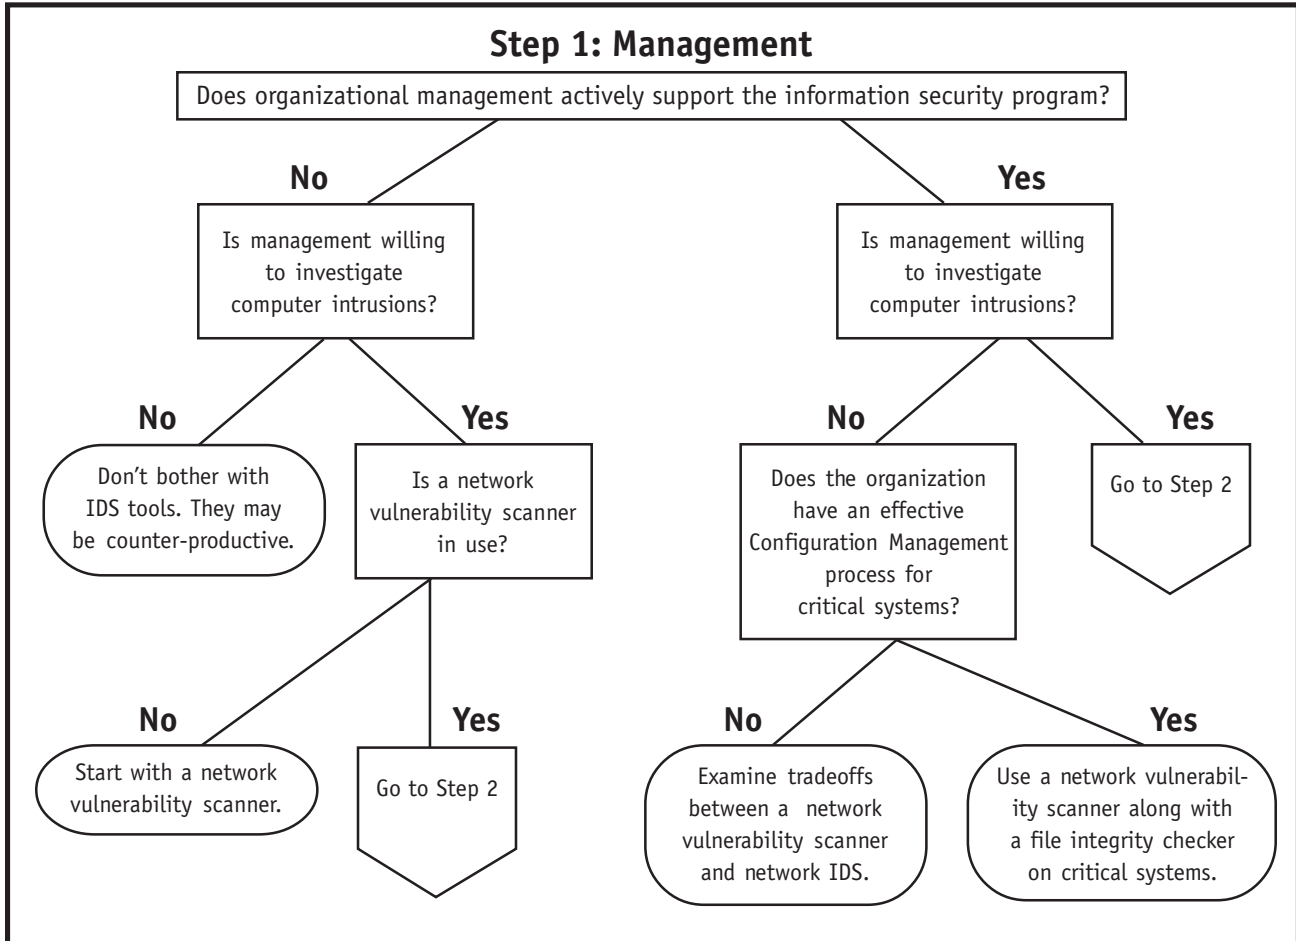
Once critical information resources have been identified, evaluate your ability to maintain the security of those resources. If you have no means of configuration management on the web site, do you really have the capability to maintain the integrity of the server? If there is no priority in the organization for repairing known vulnerabilities on critical systems, will detection software on those systems be valuable?

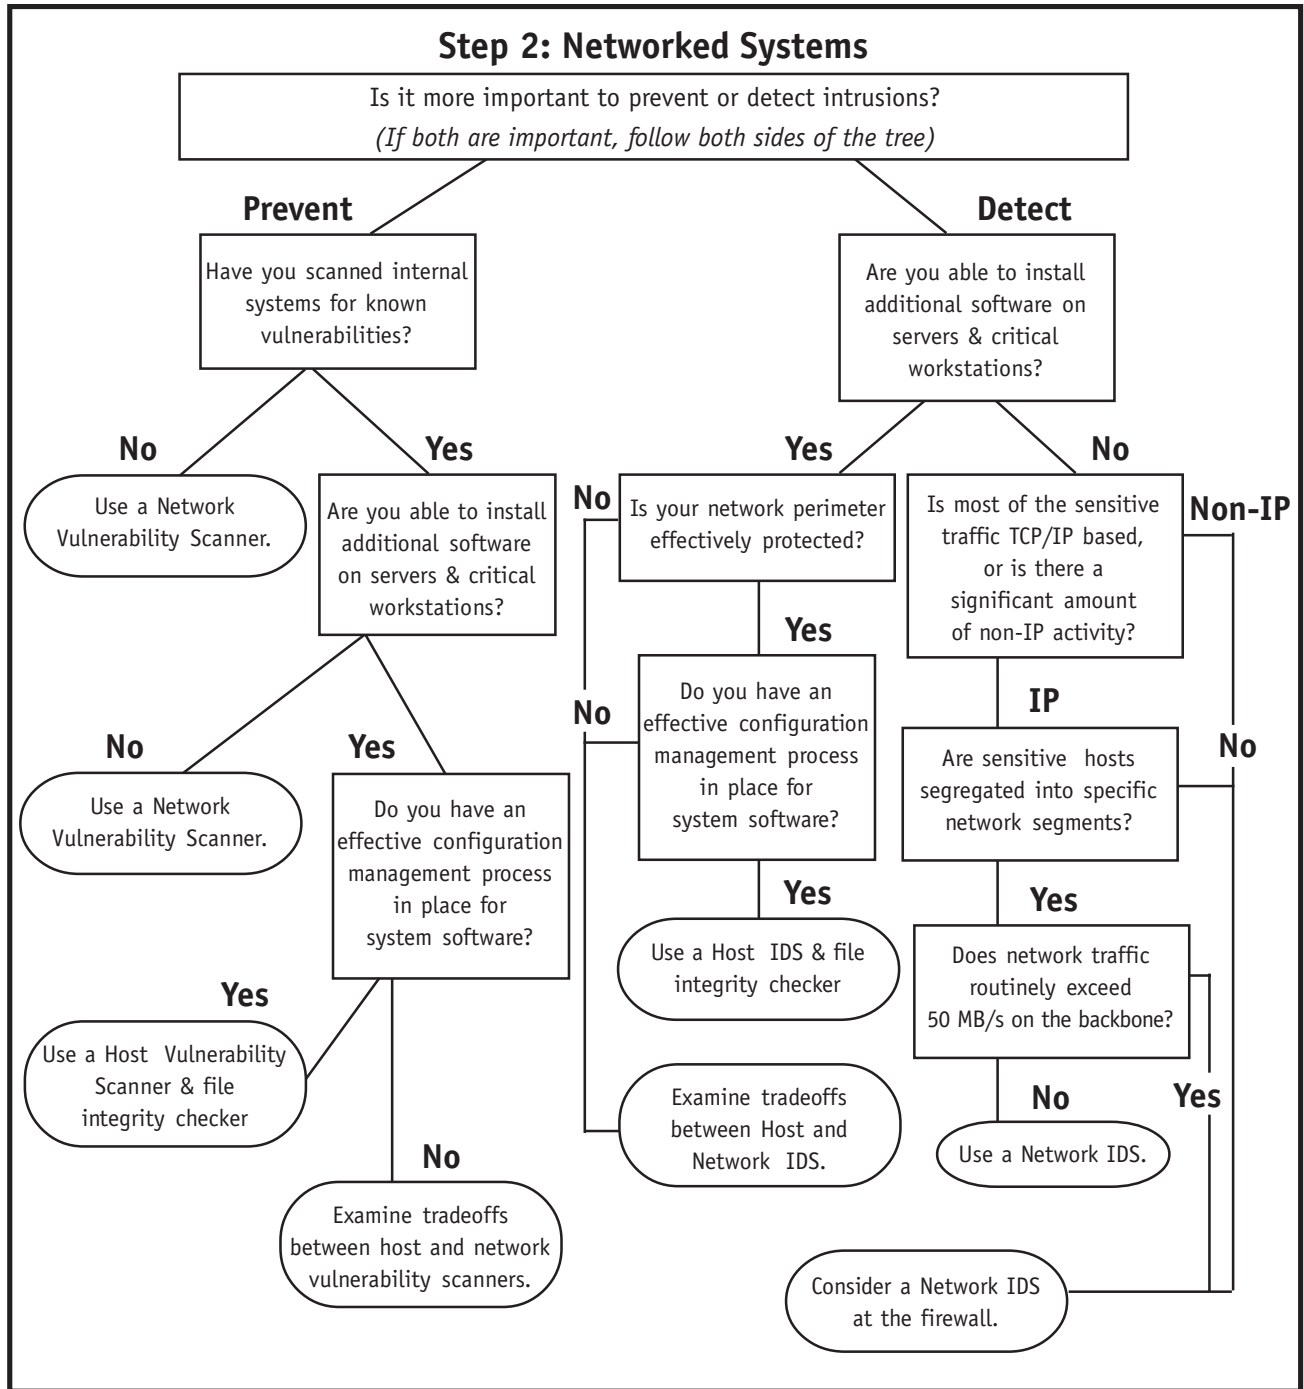
Access to machines is also a critical factor. Host based intrusion detection systems require more attention than network based intrusion detection. Since a host based IDS requires routine privileged access to the OS of the affected machine, a host based IDS implementation requires significant support from the group that maintains those machines. This single factor is the most important consideration when installing a host based IDS.

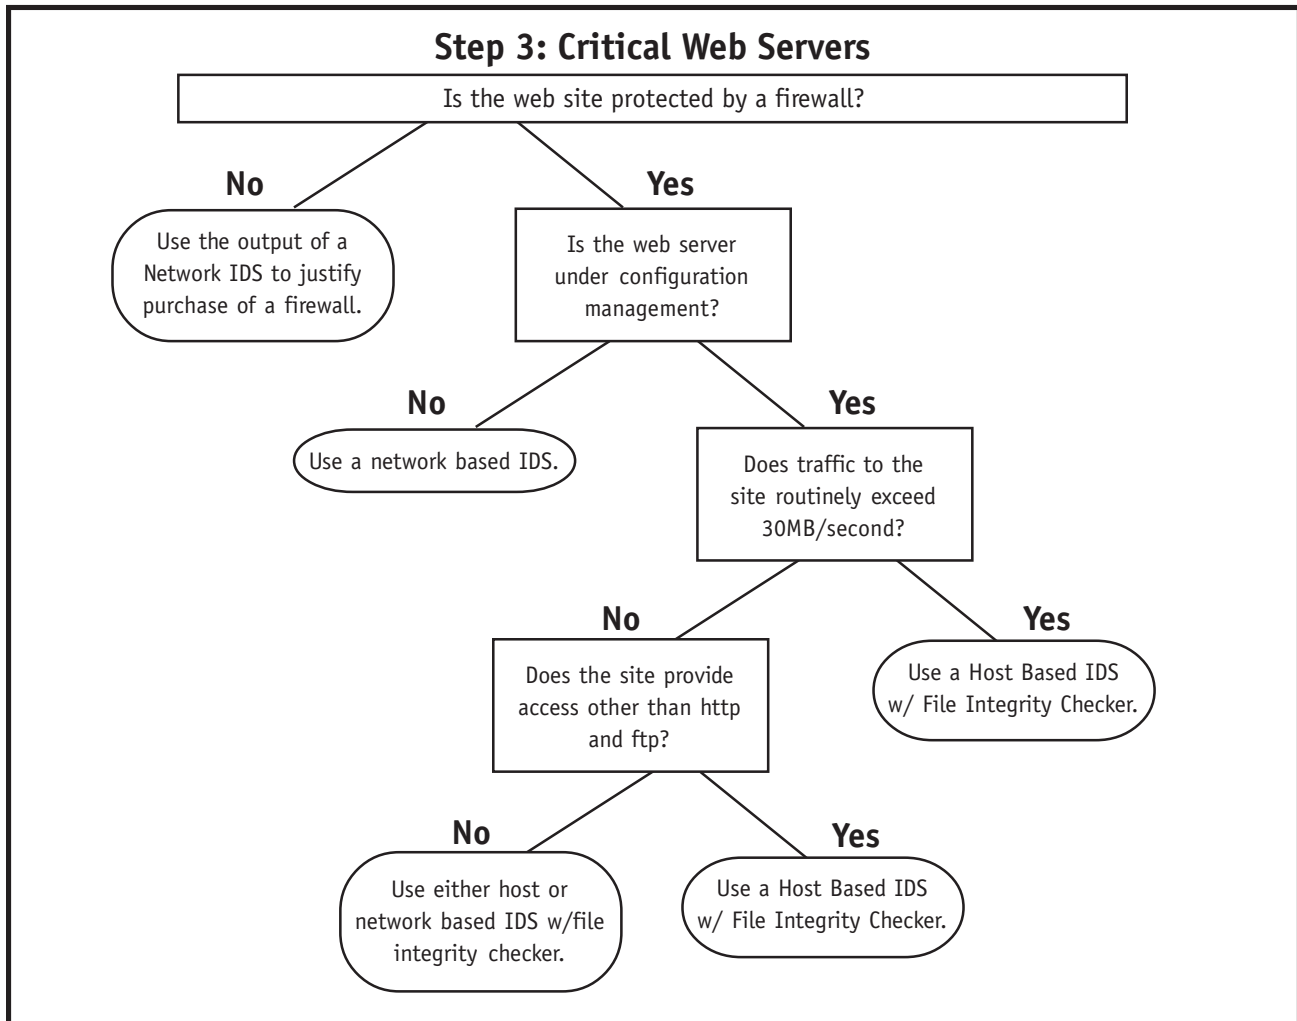
Web Site Issues

Installing an IDS for a web server or DMZ presents special considerations. Since web servers are more susceptible to manipulation, they should be given special consideration by the IDS team. Transactional web servers are more vulnerable than strictly informational web servers, since the range of interaction is larger. Consider using more than one type of IDS in these environments, because they are more sensitive and configuration management for web servers is usually tighter than Configuration Management for internal systems.

IDS PRODUCT DECISION TREES







Practical Integration Issues

Installing or integrating an IDS with other tools in the security arsenal requires some extra planning. If you have been following other ICSA.net guides, you should know about firewall configuration issues and network topology security considerations. This section should continue where those documents left off and help you to avoid common pitfalls when installing your IDS.

To avoid duplicating material presented elsewhere, this section assumes you have an internal intrusion response policy and the skilled personnel necessary to monitor the IDS. It assumes you have a firewall. It assumes your perimeter system topography follows prudent security principles.

Sensor Placement for a Network IDS

If you are deploying a network IDS, you should decide in advance where to place the monitoring sensors. This will depend significantly on what kind of intrusion or attempted intrusion you are trying to detect.

Start by creating a detailed network diagram, if you don't already have one. A network diagram can be invaluable to IDS planning. When looking at the diagram, evaluate key network choke points or collections of systems that are sensitive to business operations. A well prepared diagram may provide intrinsic clues to the right location for IDS sensors.

If the IDS is going to monitor a web server for penetrations, then the most useful position for the sensor will be on the DMZ segment with the web server. This assumes, of course, that your web server is in a DMZ segment, rather than outside or inside the firewall (neither of which is a particularly good idea). If attackers compromise the server, the IDS has the best chance of detecting either the original penetration or the resulting activity originating from the compromised host.

If the IDS is going to monitor for intrusions targeting internal servers, such as DNS servers or mail servers, the best place for a sensor is just inside the firewall on the segment that connects the firewall to the internal network. The logic behind this is that the firewall will prevent the vast majority of attacks aimed at the organization, and that regular monitoring of firewall logs will identify them. The IDS on the internal segment will detect some of those attacks that manage to get through the firewall. This is called "defense in depth."

Some organizations will want to use the IDS to monitor internal resources such as a sensitive collection of machines or a specific department or physical location. In that case, the most logical place for the IDS sensor will be on the choke point between those systems and the rest of the internal network.

What about outside the firewall?

Note that we have not discussed placing a sensor outside the firewall. Some sources recommend or advocate an IDS sensor outside the firewall to monitor attacks from the Internet.

Let's look at that idea in an operational environment to see what it means.

The firewall itself should log any attacks it stops (assuming you have a logging firewall and read the logs). What added value does the IDS provide in this scenario? If you aren't looking at the firewall logs, why would you pay attention to the IDS? If you are, you'll see the attacks stopped at that point. The IDS is redundant. It probably isn't a productive use of organizational money and time to install and monitor an IDS outside the firewall.

That being said, if you still want to count or evaluate attacks that are stopped at the firewall, an IDS can simplify the process of tracking and classifying attacks. The data should still be in the firewall logs, but the IDS may improve the readability or clarity of the information.

Host integration for Host IDS

If you plan to use a host-based system, you should have an adequate testing and familiarization phase. This allows the operators and analysts to become familiar with the operation of that particular piece of software.

The IDS should be installed on a development system well in advance of planned installation on a production system. Even on a quiescent system, some files will change regularly (for example, the audit files), so the IDS will report some changes. Some host-based systems, as an additional example, will report when a user process alters the system password file. This would happen if an intruder added an account. It also happens, however, when a user changes his or her password. The IDS analyst needs time to become familiar with the correct operation of each system, so that he or she can properly diagnose deviations from "normal" alarms.

Keep in mind, when using a host-based system, that it should be monitored frequently. This means twice a day, at least. If an attacker has superuser access to the system, he or she can alter the IDS or the IDS database to suppress alarms. If the IDS writes to a file, the attacker can simply edit the output file. In other words, always be suspicious that something may have altered the configuration of the IDS.

Alarm Configuration

IDSs come with a configurable alarm levels. Some will integrate with network management stations, some allow paging, some send e-mail, and some can interoperate with firewalls to shut down all traffic from the network that originated the attack.

Be very cautious about using these features. In fact, for the first month or two, turn off all alarms. Only look at the output from the system to see what it is detecting. All IDSs have, as discussed above, some level of false positives; that level can be as high as 80 or 90 percent of reported alarms. You need to be familiar with your particular system before you start turning on alarms.

Alarm misconfiguration, or over-aggressive response to alarms, can lead to an organizational decision to turn off the IDS. Richard Marcinko, a former US Navy SEAL, tells about throwing rabbits over fences into areas protected by motion sensors. When the guard force got tired of responding to alarms (only to find rabbits munching on the lawn), Marcinko's team was free to pass through that area, knowing the alarms would be ignored. Hackers know that IDS installations are monitored by humans, and that humans have human failings. They know that if they trigger alarm after alarm, the people monitoring the system will stop paying attention.

Likewise, if the IDS is configured to instruct the firewall to deny all traffic from "attacking" networks, the hackers can easily exploit this. Someone sufficiently motivated or malicious could use this against the organization by spoofing attacks from the organization's business partners or well-known sites (Yahoo, AltaVista, Amazon, CNN, Microsoft, etc) so that the firewall will deny inbound traffic from those sites, including e-mail and web site traffic.

Remember, the IDS is not security's saving grace, it is only a tool (and a fairly unintelligent one at that).

Integration Schedule

Install one sensor at a time. Don't rush the installation in order to roll out the IDS capability in a short time span. It takes a certain amount of time for the administrators and analysts to gain familiarity with the peculiarities of a given system or network point, and the peculiarities may not be the same from point to point. A sensor in a DMZ may see a given set of behaviors, while a sensor on the internal network may see another set of behaviors, with a very small intersection.

It is crucially important that the staff assigned to monitor the IDS be adequately familiar with each device in the configuration.

Case Studies for Intrusion Detection and Related Products

Case 1: Integrity Analysis

In 1996, one of the early online web-based stock trading sites was placed in full operation, and was infiltrated by an outside attacker. The trading system consisted of approximately twenty web servers connected to a central database server. When the system manager realized that an attacker was on the loose inside the firewall, and was actively logging into the server, there was an understandable amount of alarm.

In situations like this, damage containment should be the first priority. However, in this case, shutting down or disconnecting all the web servers from the Internet was not an acceptable option. First, doing so would constitute a “trading halt” event, and would cause the corporation to be fined in 15-minute increments by the SEC. Second, the damage to reputation caused by a shutdown would be extremely high, as would the damage associated with the possibility of word leaking out that an intruder had successfully broken into the system.

Because the system manager had already deployed a product utilizing Integrity analysis, it was possible to ascertain quickly which machines were compromised and to determine the scope of the infiltration. The customer computed that they saved about 260 hours of system administration time, in a case where each minute was valued at an extreme premium. Time is critical when an attacker is on the loose in your network.

This story ends happily. Only a fraction of the machines were compromised, and were promptly shut down. The database server was found to be intact, which allowed the web site continue functioning on the remaining web servers. The system administration team conducted damage eradication and recovery at a more leisurely pace.

Case 2: Vulnerability Assessment

A consulting company that does network design, security assessment and integration services is frequently called in when a company is initially establishing a network, restructuring an existing one or adding new and complex capabilities. In the words of their President, “Many companies do not realize that when Windows NT is installed ‘out of the box,’ it’s designed to be wide open to allow for flexible network implementations. And it’s pretty difficult to get a global picture of your environment, because you have to go through a lengthy process of ‘machine by machine’, or ‘share by share’, or ‘domain by domain.’ They simply do not have the training, background and expertise to know what specific rights and permissions to turn off.

“We use a vulnerability assessment product combined with a network management product to help uncover information about user rights, permissions, account access, account restrictions, and users that have easily-guessed passwords.

“One eye-opening experience we found at a customer site was where someone with user privileges granted themselves administrator rights. When we ran a user access report we found a user who had used a hack to make himself an administrator. To make matters worse, the account was active, and it belonged to a former employee that had been gone for two months.

“It would have taken us forever to find this situation because it is extremely time consuming to manually check each and every user account for security violations. But it is much easier with a vulnerability assessment product where information across an entire enterprise can be consolidated into one single report.”

Case 3: Host-based Intrusion Detection

In December of 1998, a medium size California bank decided that they needed better control of their internal security. They needed both consistency in their security configurations as well as monitoring for suspicious behaviors from authorized users inside the system. They selected a host-based intrusion detection tool that also provided host-based assessment.

When agents were deployed to 10 servers and a handful of workstations. After installation, an audit policy was deployed that reduced the amount of data collected to a reasonable level, and a detection policy was established that matched the objective of monitoring for anomalous behavior. The security officer then used the assessment capabilities to bring all the servers up to a consistent level of security configuration that was acceptable to the security officer.

Within 24 hours of beginning monitoring the security officer observed irregular usage of two administrative accounts. They were being used to read mail and edit documents during regular working hours. The security policy specified that administrative accounts were only to be used for tasks requiring administrative privilege and were not to be used for daily activities such as reading mail. The employees who were using their admin accounts were reprimanded and the activity stopped. Within 48 hours of monitoring the security officer observed an unauthorized account using the backup software. The immediate security risk was that the backup software had privilege to read every file on the system bypassing all access control. The security officer called the account owner and quickly determined that the backup software had been installed under the wrong account making this powerful software vulnerable to compromise. The software was re-installed under a better-protected account.

Within 72 hours of monitoring the security officer observed regular account logins from a set of three accounts at 1:30 AM, 2:30 AM, and 3:30 AM. All the indications were that this was an automated program using these three accounts to login at the same time everyday. By using the data forensics capabilities of the intrusion detection tool the security officer looked back over the last 3 days to determine other accesses and executions by these accounts during these times. The next effort was to talk to the account owners to determine if they had knowledge of programs under their control during this time. Through a combination of analyzing the data and interviewing the end-users, it was determined to be MAPI interactive logons for mail. This pattern is now recognized as authorized.

IDS Market Share and Industry Growth

Excerpted from “Plugging the Holes In eCommerce: The Market For Intrusion Detection and Vulnerability Assessment Software, 1999-2003” by Abner Germanow, International Data Corporation. Used by permission.

Introduction

IDC believes mission-critical networks must maintain a clear view of potential vulnerabilities as well as keep vigil over networks and systems. A fair number of enterprises that are conscious of the value of their networks have installed firewalls in an attempt to keep malicious entities out of the network. Even though firewalls are a core component of any security architecture, network managers cannot rely only on the firewall.

IDC defines the IdnA market as those tools that can detect intrusions or analyze vulnerabilities that might permit intrusions on a computer network. IdnA tools are further broken down into network based tools and host based tools. To date, vulnerability assessment tools hold the majority of revenues in this market. However, as the IdnA market expands, vulnerability assessment will take second stage to intrusion detection.

The Market

The market for IDnA products was \$136 million in 1998 and will grow to \$978 million by 2003. In 1998, leading vendors included AXENT, ISS, and Network Associates. Current vulnerability assessment products are strengthening quickly, but customers are demanding less complex products with simpler management interfaces. Vendors have made great strides in improving usability in an effort to offer products that can be operated by mere mortals rather than MIT PhD's in computer science. Intrusion detection systems have farther to go down this path than vulnerability assessment tools before either technology is able to fully penetrate the corporate mass market.

Table 1
Worldwide Intrusion Detection & Vulnerability Assessment Software
Revenue, Growth, & Share, 1997-2003

	1997	1998	1999	2000	2001	2002	2003	1999-2003 CAGR (%)
Vulnerability Assessment (\$M)	37.7	91.0	162.2	253.5	346.9	406.1	449.8	29.0
Growth (%)	-	141	78	56	37	17	11	
Share (%)	65	67	62	58	54	49	46	
Intrusion Detection (\$M)	20.3	45.3	99.4	183.5	295.5	422.7	518.1	51.8
Growth (%)	-	123	120	85	61	43	25	
Share (%)	35	33	38	42	46	51	54	
Total (\$M)	58.0	136.3	261.7	437.0	642.4	828.7	977.9	39.0
Growth (%)	-	135	92	67	47	29	18	

Key Assumptions:

- The market for intrusion detection and vulnerability assessment software products will exhibit very similar development to the firewall market (refer to Figure 2).
- Intrusion detection products produce ongoing results, while vulnerability assessment products expose overall inadequacies. IDC expects network managers to increasingly purchase intrusion detection products that showcase their skills rather than vulnerability assessment products that expose faults.
- The network is the business, or at least its importance skyrocketed over the last four years. The ability to quickly identify and fix vulnerabilities as well as identify and stop intruders is gaining momentum. As the network's value to the business increases, spending on protecting that asset will increase accordingly.
- Network and systems managers must continually accomplish more with less. As the amount of skill and time required to commit to intrusion detection and vulnerability assessment products decreases, potential customers will increase.
- Intrusion detection products operate independently of other systems today. Integration with other security products such as firewalls, routers, and systems and network management products is paramount to continued growth. Integration will produce easier-to-install and more robust products whose value will be determined by large numbers of plug-and-play integration points that will provide comprehensive coverage, reaction, and investigation.
- Vulnerability assessment products will find integration points with network and systems management platforms; however, in order to serve as an independent quality control agent, these products will remain functionally separate.

Messages in the Data:

- Intrusion detection products are becoming the lead point products and will lead the overall market in revenue within three years.
- Intrusion detection products will experience very high rates of growth, especially during the latter part of the forecast period, when product maturity will enable sales into the full range of corporate environments.
- Vulnerability assessment products will begin to take a supporting role to intrusion detection but still experience high growth rates.

Source: International Data Corporation, 1999

The worldwide software market for IDnA software products was \$136 million in 1998. IDC estimates that this figure represents 135% growth over the 1997 market levels. In 1998, vulnerability assessment products garnered the lion's share of the market at \$91 million, or 67% of the overall market. Intrusion detection products represented 33% of the market at \$45 million.

Like most products related to enabling secure commerce, the outlook for this market is extremely positive. The demand for network and systems security is rising alongside the growth in e-commerce and e-business. IDC expects this market to almost double in 1999 to reach \$262 million. The market will continue to reach \$978 million in 2003 at a CAGR of 39%.

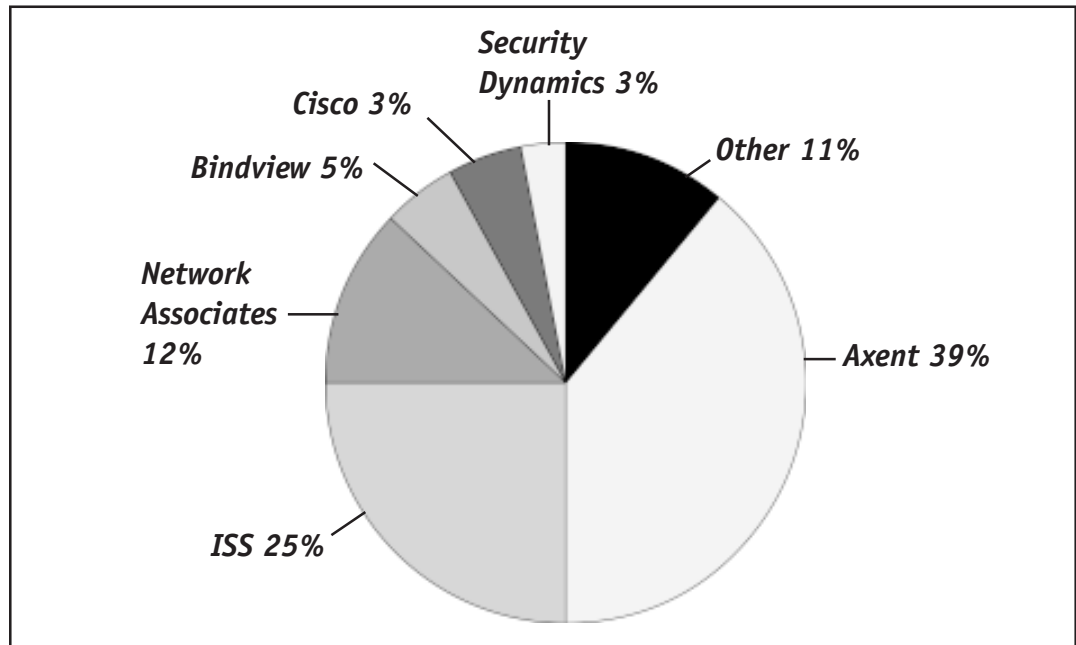
What is the primary driver for this market? In a word, complexity. Networks and systems are invariably patched and tied together in ways vendors never intended their products to be used. Proprietary and highly customized applications abound, and throughout it all, buggy software of various versions. Despite this environment, a business must use the network to reach out to almost every entity with which it maintains a relationship.

The network is no longer about bits, bytes, and email. The network is about establishing, maintaining, and enhancing trust in every relationship the business owns. The network is the business. There is a wide variety of ways to foster trust through various networked applications, but there are far more ways to violate trust. No company wants to be featured on the front page of today's paper talking about how its Web site failed, or worse, explaining to customers, stock holders, and Wall Street why the company's cash flow and profitability went down along with its Web site.

IDC expects the intrusion detection segment of the IDnA market to grow from \$50 million in 1998 to \$431 million in 2003 at a CAGR of about 44%. Growth is expected to be derived from very security-conscious enterprise customers, application service providers that must preserve the integrity of their customer's applications, and the traditional government and military sectors.

Due to the high number of false positives that old products and some current products produce, successful implementations of intrusion detection products can still require a fair amount of associated labor costs. Integrating the intrusion system with classic network and systems management platforms beyond simply sending alerts off to the proper console will enable customers to manage their systems without significant retraining. Enabling existing employees to manage more without significantly extending responsibilities or requiring large amounts of training will lower another barrier to entry for intrusion detection.

WORLDWIDE INTRUSION DETECTION AND VULNERABILITY ASSESSMENT SOFTWARE REVENUE SHARES BY VENDOR, 1998



Conclusion

Not only are customers faced with waves of complexity on all fronts, they are also struggling to support their environments with stable or decreasing resources. Vendors that are able to deliver products that are easy to use, are easy to install, and deliver on the promise of solidifying the electronic foundation of business will be well rewarded.

The infrastructure on which the network-based business will rely requires a watertight foundation that can be easily monitored and assessed. This ease of assessment will set a level of trust in the network that will determine the business value of the applications business line managers rely on. Fostering and maintaining this trust in the network is the key value for IDnA products. Components of the network infrastructure, from operating systems to servers to management applications, will continue to evolve and present a steady stream of vulnerabilities.

As the network becomes easier for the user, the back-end becomes increasingly complex. e-Commerce activities require regular assessment and constant monitoring, giving IDnA products the ability to grow with high-value networks long into the future.

Appendix A - Incident Response Planning

Any organization that intends to detect intrusions should have a plan for handling them. This falls under the category of Incident Response. Failure to plan for incident response can have a devastating effect on any resulting investigation.

Let's use an example to understand the importance of incident response planning. The Acme Widget Company help desk receives a telephone call from a user saying that her virus detection software detected a virus in an E-mail from a colleague. Is this an incident? Probably not.

Then another employee reports a virus in an E-mail just after the virus scanner on the Internet E-mail proxy server reports removing a virus from an outbound E-mail message. Is this an incident? Maybe.

At this point, the help desk should be able to determine if the incident response team should be notified. If this level of virus activity is very unusual, the first member of the team might be called in to make a decision about further action. If viruses are not considered incidents, there may be no need to activate an incident response team.

The most important principle to keep in mind when writing an incident response plan is that incidents are seldom easy to identify or handle. Planners should emphasize flexible guidelines that enable response personnel to understand problems and make decisions in rapid succession. Like snowflakes, no two incidents are exactly alike.

Policy and Plans

The first step for incident response planning is to examine (or develop) a policy. The plan should be driven by management's expectations, as described in the policy. As shown above, the policy should describe types of incidents and the requirements for handling incidents. The policy describes *what* should be done. The Incident Response Plan should describe *how* incidents should be handled.

Incident Response Plans range from the simple to the extremely complex. At a minimum, plans should describe basic procedures for reporting and responding to any incident defined in the Incident Handling Policy. Plans should specify who responds and how response personnel interact. Plans should be kept up to date so that any new member of the response team can find contact information or refer to procedures.

One candidate outline for an Incident Response Plan might look like this:

- 1) Response Team Overview
 - a) Define the scope and membership of the team. Outline references to policies and functions within the corporation. Establish the basis for the Incident Response Function
- 2) Incident Definition
 - a) Explain types of incidents and appropriate responses. Clearly define what types of incidents should be handled.

- 3) Incident Reporting Process
 - a) Describe the process for reporting incidents.
 - b) Define the process for receiving reports of incidents. Include forms as attachments to the document.
 - c) Define the reporting and escalation hierarchy.
 - d) Define specific procedures for handling incident reports after the incident has occurred.
 - e) Define procedures for classifying and storing incident related information.
 - f) Outline employee responsibilities for the security of information.
 - g) Specify corporate requirements for interacting with outside organizations.
- 4) Incident Management Process
 - a) Provide a high level description of the overall incident response process.
 - b) Describe duties of each member of the incident response team
 - c) Describe documentation procedures
 - d) Detailed Incident Response Process
 - i) Receiving Reports
 - ii) Initial Notification
 - iii) Initial priorities for response personnel
 - iv) Containment activities
 - v) Escalation procedures
 - vi) Investigation procedures
 - vii) Evidence preservation
 - viii) Remediation procedures
 - ix) Returning to normal operations
- 5) Notification trees and contact lists
 - a) Describe primary and alternate means for contacting each member in the notification tree.
 - b) List telephone, FAX, and E-mail information for each member of the tree.
- 6) Forms
 - a) Incident Reporting Form
 - b) Incident Tracking record format

THE FOLLOWING CHECKLIST IS INTENDED TO ASSIST IN PLANNING AN INCIDENT RESPONSE CAPABILITY

1	Does the plan implement requirements of the policy?	
	Does the plan clearly define roles and responsibilities during the response process?	
	Does the plan appropriately recognize the role of legal counsel during and after an incident?	
	Does the plan describe the process for initiating and ultimately resolving an incident?	
	Does the plan cover initial notification via normal problem reporting processes (e.g. via the help desk)?	
	Does the plan contain current primary and alternate telephone numbers for all response personnel?	
	Does the plan address communications with business partners and vendors?	
	Does the plan address protection of incident related information from unauthorized access?	
	Does the plan provide a mechanism for communicating among response team personnel?	
	Does the plan address methods for interacting with the media?	
	Does the plan address methods for communicating with affected employees?	
	Does the plan specify requirements for communicating with law enforcement personnel?	
	Does the plan provide for alternate methods of communication in case the primary method becomes unavailable?	
	Does the plan address decision-making processes for dealing with large-scale system or data corruption?	
	Does the plan address procedures for restoring damaged systems?	

When building an Incident Response don't plan to address all possible contingencies. Rather than attempt to define a policy for each contingency, it may be easier to define an approach for making decisions about certain types of problems. For example, if the web site is down for more than 3 hours, a back-up site should be brought online. By defining the timeframe in advance, response personnel are less likely to postpone difficult decisions during the heat of the moment.

Overall, the Incident Response plan should provide enough planning and prior coordination to simplify the response process when it occurs. If response personnel know each other and understand the process for handling an incident, successful resolution is a much more likely event.

Incident Handling Resources

Several web-based resources are available for anyone planning an incident response capability. These resources define incident reporting and handling procedures from a variety of viewpoints. Since most of the information is freely available, we have only provided pointers in this document.

Forum of Incident Response and Security Teams (FIRST)

FIRST is a group of incident response teams originally sponsored jointly by CERT-CC, CIAC, and NIST. FIRST teams routinely participate in meetings to enhance their capabilities and provide information via a web site at NIST. The site can be reached at: <http://www.first.org>

The Computer Emergency Response Team Coordination Center (CERT-CC)

CERT-CC is sponsored by DARPA and operated by the Software Engineering Institute at Carnegie Mellon University. CERT-CC was one of the first incident handling organizations to respond to a broad range of incidents around the Internet. It was the first incident response team with a general responsibility for incidents on the Internet. CERT-CC provides information on reporting and assistance in handling incidents of all kinds. CERT-CC also published bulletins describing known vulnerabilities to computer systems. This site can be reached at: <http://www.cert.org>

The Federal Computer Incident Response Capability (FedCIRC)

FedCIRC is sponsored by the National Infrastructure Protection Center and the US Department of Justice. Its responsibility includes US Federal government computers and networks. FedCIRC aggregates information from several federal agencies and, subsequently, can provide significant statistical data on attack vectors from the Internet. This site can be reached at: <http://www.fedcirc.gov>

Department of Energy Computer Incident Advisory Capability (CIAC)

Another of the oldest incident handling teams in the US, this organization primarily supports the Department of Energy. It also provides notices to the international community regarding new threats to computer systems. This site can be reached at: <http://ciac.llnl.gov>

The Hoaxes Page

One type of incident that does not receive enough attention from new incident handlers is the hoax. Hoaxes take various forms, but most report on a problem that doesn't really exist. Many users who are new to the Internet pass these hoaxes along, creating a type of chain letter that needlessly consumes bandwidth on a network. This site debunks the vast majority of Internet hoaxes and provides response personnel with an authoritative means of investigating unusual reports. Several organizations offer this service. They can be found at:

<http://www.icsa.net/services/consortia/anti-virus/alerthoax.shtml> - ICSA.net Hoaxes Page

<http://www.datafellows.com/news/hoax.htm> - DataFellows Hoaxes Page

<http://www.av.ibm.com/BreakingNews/HypeAlert/> - IBM Antivirus Hoaxes Page

<http://www.symantec.com/avcenter/hoax.html> - Symantec Antivirus Research Center Hoaxes Page

Other Incident Handling Web Sites

In addition to these sites, many others deal with the topic of incident handling.

<http://www.coast.cs.purdue.edu> - No security professional should be without a bookmark on this site. It is sponsored by the computer security laboratory at Purdue University.

<http://www.nasirc.nasa.gov> - This site provides incident handling information to organizations around the National Aeronautics and Space Administration.

Sample Incident Reporting Form

version 5.1

July 1999

CERT(R) Coordination Center Incident Reporting Form

CERT/CC has developed the following form in an effort to gather incident information. If you believe you are involved in an incident, we would appreciate your completing the form below. If you do not believe you are involved in an incident, but have a question, send email to:

cert@cert.org

Note that our policy is to keep any information specific to your site confidential unless we receive your permission to release that information.

Return this form to:

cert@cert.org

If you are unable to email this form, please send it by FAX. The CERT/CC FAX number is:

+ 1 412 268 6989

We would appreciate any feedback or comments you have on this Incident Reporting Form. Please send your comments to:

cert@cert.org

Submit this form to: cert@cert.org

If you are unable to send email, fax this form to: +1 412 268 6989

Your contact information

name:
email address...:
telephone number:
other.....:

Affected Machine(s)

(duplicate for each host)
hostname and IP.:
time zone.....:

Source(s) of the Attack

(duplicate for each host)
hostname or IP.:
timezone.....:
been in contact?:

Description of the incident

(Include dates, methods of intrusion, intruder tools involved, software versions and patch levels, intruder tool output, details of vulnerabilities exploited, source of attack, or any other relevant information.)



Appendix B - Frequently Asked Questions About Intrusion Detection

What is an Intrusion Detection System?

An intrusion detection system monitors computer systems, looking for signs of intrusion (unauthorized users) or misuse (authorized users overstepping their bounds).

What does it do?

Intrusion detection systems monitor a variety of information sources from systems, analyzing this information in a variety of ways. The first, most common, is that it compares this information to large databases of attack signatures, each reflecting an attempt to bypass or nullify security protections. The second is that it looks for problems related to authorized users overstepping their permissions (e.g., a shipping clerk searching executive payroll records). Finally, some intrusion detection systems perform statistical analysis on the information, looking for patterns of abnormal activity that might not fall into the prior two categories (e.g., accesses that occur at strange times, or an unusual number of failed logins.)

But we already have a firewall—why do we need an intrusion detection system, too?

The firewall is the security equivalent of a security fence around your property and the guard post at the front gate. It can keep the most unsavory of characters out, but cannot necessarily tell what is going on inside the compound. Intrusion detection systems are the equivalent of multi-sensor video monitoring and burglar alarm systems. They centralize this information, analyze it for patterns of suspicious behavior in much the same way a guard at a monitoring post watches the feeds from security cameras, and in some cases, deals with problems they detect. Most loss due to computer security incidents is still due to insider abuse. Intrusion detection systems, not firewalls, are capable of detecting this category of security violation. Perhaps more importantly, firewalls are subject to circumvention by a variety of well-known attacks.

What can an intrusion detection system catch that a firewall can't?

Firewalls are subject to many attacks. The two considered most worrisome are tunneling attacks and application-based attacks.

Tunneling attacks arise due to a property of network protocols. Firewalls filter packets, and make pass/block decisions based on the network protocol. Rules typically check a database to determine whether a particular protocol is allowed, if so, the packet is allowed to pass. This represents a problem when an attacker masks traffic that should be screened by the firewall by encapsulating it within packets corresponding to another network protocol.

Application-based attacks refer to the practice of exploiting vulnerabilities in applications by sending packets that communicate directly with those applications. Therefore, one could exploit a problem with Web software by sending an HTTP command that exercises a buffer overflow in the web application. If the firewall is configured to pass HTTP traffic, the packet containing the attack will pass.

We've invested in a lot of security devices for our network resources: we have token-based Identification and Authentication, require our employees to encrypt their email, have firewalls, require users to generate good passwords and change them often—why do we need intrusion detection, too?

Even when you have a great existing security infrastructure, you still need the added assurance intrusion detection systems provide. No matter how well designed the security point products, they are still subject to failure, due to hardware or software anomalies or user problems. Users sometime nullify the protection afforded by the products by disabling or bypassing them. Intrusion detection systems, because they are capable of monitoring messages from the other pieces of the security infrastructure, are able to detect when failure occurs. In some cases, they can tell you what happens until someone can restore them to service.

What are Vulnerability Assessment Products?

Vulnerability Assessment Products, also known as “Vulnerability Scanners,” are software products that perform security audits on systems, searching for signs that the systems being scanned are vulnerable to certain systems attacks.

How do they work?

Vulnerability Assessment Products take two approaches to locating and reporting security vulnerabilities. The first approach, a “passive” scan, inspects system settings such as file permissions, ownership of critical files, path settings, etc., for configurations that experience has shown lead to security problems. The second approach, an “active” scan, actually reenacts a series of known hacker attacks, recording the results of the attacks. Some products also perform password cracking on password files in order to discover bad/weak passwords that might be easily guessed by hackers. Finally, the products record their findings in a result screen and in a report mechanism.

What is the value added in Vulnerability Assessment Products?

Vulnerability Assessment Products are a valuable part of any organization's system security management program. They allow system managers to baseline the security of a new system. They allow periodic security audits to determine the security health of a system at a given time. Many of them provide the ability to perform “differential analysis” by archiving the results of scans, then comparing subsequent scans to the archives, reporting when new vulnerabilities or unexpected changes appear.

Appendix C - Glossary

ACCOUNTING - A mechanism, usually built into a computer operating system, for tracking how many resources a user consumes. Frequently confused with auditing.

ACTIVITY - Instantiations of the data source that are identified by the analyzer as being of interest to the security administrator. Examples of this include (but are not limited to) network sessions, user activity, and application events. Activity can range from extremely serious occurrences (such as an unequivocally malicious attack) to less serious occurrences (such as unusual user activity that's worth a further look).

AGENT - The ID component that periodically collects data from the data source, sometimes performing some analysis or organization of the data. Also known as **SENSOR**.

ANALYZER - The ID component that analyzes the data collected by the sensor for signs of unauthorized or undesired activity or for events that might be of interest to the security administrator.*

AUDIT LOG - The log of system events and activities generated by the operating system.

CERTIFICATE - A computer-based record which at least (1) identifies the certification authority issuing it, (2) names or identifies its subscriber, (3) contains the subscriber's public key, and (4) is digitally signed by the certification authority issuing it. Also known as a "digital certificate".

CERTIFICATION AUTHORITY - A person who issues a certificate. Related Terms: Some documents use the term "certificate issuer" to refer to what we call a "certification authority". The two terms are closely synonymous.

DATA SOURCE - The raw information that an intrusion detection system uses to detect unauthorized or undesired activity. Common data sources include (but are not limited to) raw network packets, operating system audit logs, application audit logs, and system-generated checksum data.

ETHERNET - A physical network protocol for transmitting information across copper wires. Ethernet network segments are restricted to distances normally less than 415 meters and utilize a packet oriented message transfer protocol. Ethernet is the most popular physical network topology in use today.

EVENT - A notification from an analyzer to the security administrator a signature has triggered. An event typically contains information about the activity that triggered the signature, as well as the specifics of the occurrence.

FILE ASSESSMENT - A technology in which message digest hashing algorithms are used to render files and directories tamper evident.

FIREWALL - A computer or router (or combination thereof) configured to permit or deny specific kinds of traffic through it. Usually used to protect a network from potentially hostile outside networks; intranetwork firewalls, however, are becoming more popular. Available in a variety of strengths and reliability.

FIRST - The Forum of Incident Response and Security Teams. A collection of incident response teams originally sponsored by the National Institute for Standards and Technology.

INCIDENT HANDLING - The part of the Security Management Process concerning the investigation and resolution of security incidents that occur and are detected. Also known as INCIDENT RESPONSE.

INTRUSION DETECTION - The technology concerned with monitoring computer systems in order to recognize signs of intrusions or policy violations.

IDS MANAGER - The ID component from which the security administrator manages the various components of the ID system. Management functions typically include (but are not limited to) sensor configuration, analyzer configuration, event notification management, data consolidation, and reporting.

IP SNIFFING - Unauthorized monitoring of directly connected IP traffic. Most IP traffic is sent cleartext (unencrypted), so it is possible to see what passes by on the network. For example, sniffers can capture the login and password pairs from telnet sessions. One of the most significant causes of break-ins from the Internet.

IP SPOOFING - Misrepresenting an IP address to gain access to a remote system. Also known as "IP masquerading".

MESSAGE DIGEST ALGORITHMS - Specialized cryptographic algorithms that are used to render files tamper-evident. The nature of message digest algorithms dictates that if an input data file is changed in any way, the checksum that is calculated from that data file value calculated will change. Furthermore, a small change in the input data file will result in a large difference in the result.

PROMISCUOUS MODE - IP interfaces can be configured to listen to all network traffic, not just that traffic destined for its particular IP address. Sniffers and network based intrusion detection systems require that the computer's IP interface be set to promiscuous mode.

RESPONSE - The actions that an analyzer takes when a signature is triggered. Sending an event notification to the security administrator is a very common response. Other responses include (but are not limited to) logging the activity, recording the raw data (from the data source) that caused the signature to trigger, terminating a network, user, or application session, or altering network or system access controls.

SCANNING - The technology concerned with scanning computer systems and networks in order to find security vulnerabilities. Also known as VULNERABILITY ASSESSMENT.

SECURITY ADMINISTRATOR - The human with responsibility for the successful deployment and operation of the intrusion detection system. This person may ultimately be charged with responsibility for the defense of the network. In some organizations, the security administrator is associated with the network or systems administration groups. In other organizations, it's an independent position.

SECURITY VULNERABILITY - a feature or error in system software or configuration that increase the likelihood of damage from attackers, accidents or errors.

SENSOR - The ID component that periodically collects data from the data source. Also known as AGENT. (In many existing ID systems, the sensor and the analyzer are part of the same component.)

SIGNATURE - A rule used by the analyzer to identify interesting activity to the security administrator. Signatures are the mechanism by which ID systems detect intrusions.

SYSTEM LOG - The log of system events and activities, generated by a system process. The system log is typically at a greater degree of abstraction than the operating system audit log.

VULNERABILITY ASSESSMENT - The technology concerned with scanning computer systems and networks in order to find security vulnerabilities. Also known as SCANNING.

For Further Reading

- The following list of recent articles about intrusion detection and vulnerability assessment is derived from the Computer Select Database for March 1999. For further information about this valuable resource, call The Gale Group at 800-848-1472 in the United States or Canada.

Delmonico, D. (1998). Detect Network Intruders Before They Wreak Havoc. *InternetWeek* (Oct 5, 1998) (735):38(1)

Scambray, J. & S. McClure (1998). Digital sentries. *InfoWorld* (May 4, 1998) 20(18):1(8) + extensive product comparisons in several articles.

Hurwicz, M. (1998). Cracker tracking: tighter security with intrusion detection. *Byte* (May 1998) 23(5):112C(5)

Null, C. (1998). Covering your assets, electronically. *LAN Times* (April 27, 1998) 15(9):44(2)

Santalesa, R. (1998). Extra-Strength Vulnerability Detection. *Windows Sources* (April 1998) 6(4):10(1)

- The Following books were found by doing a search on Amazon.com (www.amazon.com)

Intrusion Detection: An Introduction to Internet Surveillance, Correlation, Traps, Trace Back, and Response, Edward G. Amoroso(Preface) / Paperback / Published 1999 218 pages 1 edition (February 1999)
Intrusion Net Books; ISBN: 0966670078

Intrusion Detection: Network Security Beyond the Firewall, Terry Escamilla / Paperback / Published 1998 416 pages (October 1998)
John Wiley & Sons; ISBN: 0471290009

Network Intrusion Detection: An Analyst's Handbook, Stephen Northcutt / Paperback / Published 1999, 267 Pages
New Riders Publishing; ISBN 0735708681

part 2
*product
information*

The ICSA.net Intrusion Detection Buyer's Guide is produced in two parts. Part One discusses IDS background issues and technology. Part Two presents specific information on over 30 different Intrusion Detection and Vulnerability Assessment products. The product descriptions are presented in a structured format that enables vendors to describe the distinguishing features of their products, while also allowing potential customers to compare and contrast product features.

The ICSA.net Intrusion Detection Buyers Guide is available in several formats:

- The ICSA.net web site is the primary distribution vehicle for the guide. It has both parts One and Two, and will always contain the most up-to-date information. You can view it for free at http://www.icsa.net/html/communities/ids/buyers_guide/index.shtml.
- The ICSA.net web site also has a downloadable PDF file of Part One.
- A hardcopy bound version has a color cover and contains the text from Part One.
- A CD-ROM version is available as a standalone piece, and is also included with selected hardcopy versions. It contains both Parts One and Two.

Finally, a special word of thanks is due to the following sponsors of the ICSA.net Intrusion Detection Buyer's Guide, without whose help it would not have been produced:

BindView Corporation
Clicknet Software
CyberSafe Corporation
International Data Corporation
Internet Security Systems, Inc.
Network Associates, Inc.
ODS Networks
Tripwire, Inc.

We trust you will agree that sponsorship of this project is a clear and commendable indication of the commitment that these companies have made to a fair and open market for quality products.
