

An Introduction to Intrusion Detection --- *&* ASSESSMENT

for System and Network
Security Management



The security assurance company

A GartnerGroup Affiliate

Participating Vendors:



Prepared by Rebecca Bace from Infidel, Inc. for ICSA, Inc.



Executive Summary

Systems and networks are subject to electronic attacks. The increasingly frequent attacks on Internet-visible systems are attempts to breach information security requirements for protection of data. Vulnerability-assessment tools check systems and networks for system problems and configuration errors that represent security vulnerabilities. Intrusion-detection systems collect information from a variety of vantage points within computer systems and networks and analyze this information for symptoms of security breaches. Both intrusion-detection and vulnerability-assessment technologies allow organizations to protect themselves from losses associated with network security problems.

The market for intrusion-detection products, driven by reports of steadily increasing computer security breaches, has grown from \$40 million in 1997 to \$100 million in 1998. Intrusion-detection is the logical complement to network firewalls, extending the security management capabilities of system administrators to include security audit, monitoring, attack recognition, and response.

Intrusion detection systems perform a variety of functions:

- Monitoring and analysis of user and system activity
- Auditing of system configurations and vulnerabilities
- Assessing the integrity of critical system and data files
- Recognition of activity patterns reflecting known attacks
- Statistical analysis for abnormal activity patterns
- Operating-system audit-trail management, with recognition of user activity reflecting policy violations

Benefits of intrusion-detection and vulnerability-assessment products include the following:

- Improving integrity of other parts of the information security infrastructure
- Improved system monitoring

An Introduction to Intrusion Detection and Assessment

- Tracing user activity from the point of entry to point of exit or impact
- Recognizing and reporting alterations to data files
- Spotting errors of system configuration and sometimes correcting them
- Recognizing specific types of attack and alerting appropriate staff for defensive responses
- Keeping system management personnel up to date on recent corrections to programs
- Allowing non-expert staff to contribute to system security
- Providing guidelines in establishing information-security policies

Unrealistic expectations about intrusion-detection and vulnerability assessment products must be corrected: these products are not silver bullets and they

- cannot compensate for weak identification and authentication mechanisms
- cannot conduct investigations of attacks without human intervention
- cannot intuit the contents of your organizational security policy
- cannot compensate for weaknesses in network protocols
- cannot compensate for problems in the quality or integrity of information the system provides
- cannot analyze all of the traffic on a busy network
- cannot always deal with problems involving packet-level attacks
- cannot deal with modern network hardware and features

Table of Contents

INTRODUCTION	9
• Definitions	9
• About ICSA's Intrusion Detection Systems Consortium	10
• About This White Paper Series	11
 INTRUSION DETECTION OVERVIEW	 11
• What Is Intrusion Detection?	11
• Vulnerability Assessment and Intrusion Detection	12
• Products Can Be Successfully Deployed in Operational Environments	12
• Table of Technology Features	13
 WHERE INTRUSION DETECTION SYSTEMS, FILE INTEGRITY AND VULNERABILITY ASSESSMENT PRODUCTS FIT IN NETWORK SECURITY MANAGEMENT	 14
• Network Security Management	14
• The Security Hierarchy	14
• Why Firewalls aren't enough	14
• Who Guards the Guard? – Trust and Intrusion Detection	15
• System Security Management – a Process View	15
 DEBUNKING MARKETING HYPE – WHAT INTRUSION DETECTION SYSTEMS AND RELATED TECHNOLOGIES CAN AND CANNOT DO	 16
• Realistic benefits	16
- They can lend a greater degree of integrity to the rest of your security infrastructure.	16
- They can make sense of often obtuse system information sources, telling you what's really happening on your systems.	17
- They can trace user activity from the point of entry to point of exit or impact	17
- They can recognize and report alterations to data files	17

- They can spot errors of your system configuration that have security implications, sometimes correcting them if the user wishes	17
- They can recognize when your system appears to be subject to a particular attack.	17
- They can relieve your system management staff of the task of monitoring the Internet searching for the latest hacker attacks.	17
- They can make the security management of your systems by non-expert staff possible.	18
- They can provide guidelines that assist you in the vital step of establishing a security policy for your computing assets.	18
• Unrealistic expectations	18
- They are not silver bullets	18
- They cannot compensate for weak identification and authentication mechanisms	18
- They cannot conduct investigations of attacks without human intervention	18
- They cannot intuit the contents of your organizational security policy.	19
- They cannot compensate for weaknesses in network protocols	19
- They cannot compensate for problems in the quality or integrity of information the system provides	19
- They cannot analyze all of the traffic on a busy network	19
- They cannot always deal with problems involving packet-level attacks	19
- They cannot deal with modern network hardware and features	20
CASE STUDIES FOR INTRUSION DETECTION AND RELATED PRODUCTS	20
Case 1: Integrity Analysis	20
Case 2: Vulnerability Assessment	20
Case 3: Host-based Intrusion Detection	21
FREQUENTLY ASKED QUESTIONS	22
About Intrusion Detection	22
• What is an Intrusion Detection System?	22
• What does it do?	22
• But we already have a firewall – why do we need an intrusion detection system, too?	22

• What can an intrusion detection system catch that a firewall can't?	22
• We've invested in a lot of security devices for our network resources:	23
About Vulnerability Assessment Products	23
• What are Vulnerability Assessment Products?	23
• How do they work?	23
• What is the value added in Vulnerability Assessment Products?	23
INTRUSION DETECTION AND VULNERABILITY ASSESSMENT:	
TECHNICAL CONCEPTS AND DEFINITIONS	23
Intrusion Detection	23
Descriptors for Intrusion Detection Systems Features and Functions	24
• Monitoring Approach	24
• Timing of Information Collection and Analysis	26
• Location of Analysis	27
• Types of Analysis	28
• Responses to Detection of Misuse or Attack	29
• Management Functions and Deployment Issues	30
• System Integrity	31
• Other Features	31
Vulnerability Assessment	31
• Introduction	31
• Assessment approach	32
• Location of Analysis	34
• Reporting	34
• Deployment	34
• Responses	34
• Management Functions	35
• System Integrity	35

SUMMARY AND CONCLUSION	35
Wide range of goals for product users	36
Developments in Other Security Product Lines Will Increase The Importance of Intrusion Detection	36
Capabilities for Intrusion Detection Products are improving	36
GLOSSARY	36
FOR FURTHER READING	38
ABOUT THE AUTHOR	38

INTRODUCTION

Intrusion detection systems help computer systems prepare for and deal with attacks. They collect information from a variety of vantage points within computer systems and networks, and analyze this information for symptoms of security problems. Vulnerability Assessment systems check systems and networks for system problems and configuration errors that represent security vulnerabilities. Both intrusion detection and vulnerability assessment technologies allow organizations to protect themselves from losses associated with network security problems.

This document explains how intrusion detection and vulnerability assessment products fit into the overall framework of security products. It includes case histories outlining scenarios in which the products have been used by customer organizations. Finally, the concepts and definitions section provides information about product features, explaining why they represent effective countermeasures to hacking and misuse.

Protecting critical information systems and networks is a complex operation, with many tradeoffs and considerations. The effectiveness of any security solution strategy depends on selecting the right products with the right combination of features for the system environment one wishes to protect. In this document, we provide the information one needs in order to be a savvy consumer in the areas of intrusion detection and vulnerability assessment.

Definitions

It is important the reader understand the following terms used in this paper:

Network Security is the property of computer systems and networks that specifies that the systems

in question and their elements can be trusted to act as expected in safeguarding their owners' and users' information. The goals of security include confidentiality (ensuring only authorized users can read or copy a given file or object), *control* (only authorized users can decide when to allow access to information), *integrity* (only authorized users can alter or delete a given file or object), *authenticity* (correctness of attribution or description), *availability* (no unauthorized user can deny authorized users timely access to files or other system resources), and *utility* (fitness for a specified purpose).

Intrusion Detection systems collect information from a variety of system and network sources, then analyze the information for signs of intrusion (attacks coming from outside the organization) and misuse (attacks originating inside the organization.)

Vulnerability Assessment (*scanners*) performs rigorous examinations of systems in order to locate problems that represent *security vulnerabilities*.

Security vulnerabilities are features or errors in system software or configuration that increase the likelihood of damage from attackers, accidents or errors.

Security Policy is the statement of an organization's posture towards security. It states what an organization considers to be valuable, and specifies how the things of value are to be protected. In practical use, security *policies* are coarse grained (*i.e.*, generalized statements that apply to the organization as a whole) and drive finer-grained *procedures*, *guidelines*, and *practices*, which specify how the policy is to be implemented at group, office, network, and system, and user levels.

Security infrastructure is the complement of measures, ranging from policy, procedures, and practices to technologies and products that represent an organization's security initiative. The goals of security counter-measures are to *detect* problems, to *delay* damage and to *mitigate* the effects of error and attack. From this perspective, vulnerability assessment and intrusion detection are necessary parts of the security infrastructure but do not, by themselves, represent a complete security infrastructure.

About ICSA's Intrusion Detection Systems Consortium

ICSA formed the Intrusion Detection Systems Consortium (IDSC) 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 and services. See <http://www.icsa.net/services/consortia/intrusion>.

IDSC Mission Statement:

"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."¹

IDSC members as of March 22, 1999 include:

AXENT Technologies, Inc.
(www.axent.com)

BindView Development Corporation
(www.bindview.com)

Centrax Corporation
(www.centraxcorp.com)

IBM
(www.ers.ibm.com)

Internet Security Systems, Inc.
(www.iss.net)

Memco Software Inc.
(www.abirnet.com)

Network Associates, Inc.
(www.nai.com)

Qwest Communications International, Inc.
(www.qwest.com)

Security Dynamics, Inc.
(www.securitydynamics.com)

Tripwire Security Systems, Inc.
(www.tripwiresecurity.com)

¹ Source: ICSA

About This White Paper Series

This is the first of a series of white papers on topics relating to intrusion detection products. These documents will help users and potential users to become familiar with intrusion detection and vulnerability assessment so that they can select those products that best meet their needs.

INTRUSION DETECTION OVERVIEW

Intrusion detection is an important security technology market. According to industry estimates, the market for intrusion detection products has grown from \$40 million in 1997² to \$100 million in 1998³. This market growth is driven by reports of steadily increasing computer security breaches (22% rise from 1996 to 1998, with \$136 million in associated losses, according to a leading survey⁴). Intrusion detection is considered by many to be the logical complement to network firewalls, extending the security management capabilities of system administrators to include security audit, monitoring, attack recognition, and response.

In this paper, we provide an overview of this security technology, expanding the traditional view of intrusion detection systems to include vulnerability assessment. These technologies play a vital role in modern system security management.

What Is Intrusion Detection?

Intrusion detection systems help computer systems prepare for and deal with attacks. They accomplish this goal by collecting information from a variety of system and network sources, then analyzing the information for symptoms of security prob-

lems. In some cases, intrusion detection systems allow the user to specify real-time responses to the violations.

Intrusion detection systems perform a variety of functions:

- Monitoring and analysis of user and system activity
- Auditing of system configurations and vulnerabilities
- Assessing the integrity of critical system and data files
- Recognition of activity patterns reflecting known attacks
- Statistical analysis for abnormal activity patterns
- Operating system audit trail management, with recognition of user activity reflecting policy violations

Some systems provide additional features, including:

- Automatic installation of vendor-provided software patches
- Installation and operation of decoy servers to record information about intruders.

The combination of these features allows system managers to more easily handle the monitoring, audit, and assessment of their systems and networks. This ongoing assessment and audit activity is a necessary part of sound security management practice.

¹ Source: Yankee Group

² Source: Aberdeen Group

³ Source: "Third Annual CSII/FBI Computer Crime and Security Survey", Computer Security Institute, March, 1998.

Vulnerability Assessment and Intrusion Detection

Vulnerability assessment products (also known as *scanners*) perform rigorous examinations of systems in order to determine weaknesses that might allow security violations. These products use two strategies for performing these examinations. First, *passive*, host-based mechanisms inspect system configuration files for unwise settings, system password files for weak passwords, and other system objects for security policy violations. These checks are followed, in most cases, by *active*, network-based assessment, which reenact common intrusion scripts, recording system responses to the scripts.

The results of vulnerability assessment tools represent a snapshot of system security at a point in time. Although these systems *cannot* reliably detect an attack in progress, they *can* determine that an attack is possible, and furthermore, they *can sometimes* determine that an attack has occurred. Because they offer benefits that are similar to those provided by intrusion detection systems, we include them in the sphere of intrusion detection technologies and products.

Products Can Be Successfully Deployed in Operational Environments

The objective of intrusion detection and vulnerability assessment is to make complex, tedious, and sometimes virtually impossible system security management functions possible for those who are not security experts. Products are therefore designed with user-friendly interfaces that assist system administrators in their installation, configuration, and use. Most products include information about the problems they discover, including how to correct these problems, and serve as valuable guidance for those whom need to improve

their security skills. Many vendors provide consulting and integration services to assist customers in successfully using their products to achieve their security goals.

Table of Technology Features

This table outlines the features for Intrusion Detection and Vulnerability Assessment systems, outlining the strengths and weaknesses of each. The specifics of each feature are explained in greater detail in the Section *Intrusion Detection and Vulnerability Assessment: Technical Concepts and Definitions*.

Technology Features of Intrusion Detection and Vulnerability Assessment

TYPE OF SYSTEM		Intrusion Detection								Vulnerability Assessment				
SYSTEM DESIGN FEATURES		Monitoring Approach				Timing of Analysis		Type of Analysis		Targets and strategies				
WHAT CAN IT DO?		Application –Based	Host-based	Target-based	Network-based	Integrated	Batch/Interval Mode	Real Time	Signature Analysis	Integrity Analysis	Statistical Analysis	Host-based (passive)	Network-based (active)	Password Assessment
“D” = Detects														
“P” = Prevents														
“R” = Repairs														
“S” = Supports														
Type	Examples of Security Problems													
Confidentiality	Unauthorized access to files and system resources		D			D			D			P	P	P
	Violation of corporate system use policies	D	D			D			D			P	P	P
	Violation of corporate security policies	D	D	D	D	D			D	D		P	P	
	Weak or nonexistent passwords	D	D			D			D			D		D
Integrity	Placement of trojan horses and malicious software		D	P		D	D	D	D	P		P	P	
	Presence of trojan horses and malicious software			D		D			D	D				
	Network service-based attacks				D	D		D	D				P	
	CGI-based attacks	D				D			D			P	P	
Availability	Denial of service attacks				D	D		D	D				P	
	Failure or misconfiguration of firewalls	D			D				D			P	P	
	Attacks occurring over encrypted networks	D	D			D			D					
	Unusual activity or variations from normal use patterns										D			
Other	Errors in system or network configuration		D						D			D,P,R	D,P,R	
	Liability exposure associated with attackers using organizational resources to attack others	P	P	P	P	P	P	P	P	P	P	P	P	P
	Post-incident damage assessment	S	S	S	S	S	S	S		S		S	S	S


WHERE INTRUSION DETECTION SYSTEMS, FILE INTEGRITY AND VULNERABILITY ASSESSMENT PRODUCTS FIT IN NETWORK SECURITY MANAGEMENT

Network Security Management

Network Security Management is a process in which one establishes and maintains policies, procedures, and practices required for protecting networked information system assets. Intrusion Detection and Vulnerability Assessment products provide capabilities needed as part of sound Network Security Management practice.

The Security Hierarchy


The following diagram outlines an Information Security Hierarchy. It outlines the security measures that comprise the foundation for any security technology.



Note that in order to achieve enterprise-wide security results, Layers 1-3 must exist in order for the technologies and products in Layer 4 to be effective.

Figure 2 illustrates the existing elements of the information security market. Intrusion detection and vulnerability assessment fit in the sector labeled “Assessment.”

Figure 2 – Information Security Market⁶



Source: GartnerGroup

Note that in monitoring and auditing the rest of the products in the matrix, intrusion detection support all of the goals.

Why Firewalls aren't enough

A common question is how intrusion detection complements firewalls. One way of characterizing the difference is provided by classifying security violation by *source*—whether they come from outside the organization's network or from within. Firewalls act re as a barrier between corporate (internal) networks and the outside world (Internet), and filter incoming traffic according to a security policy.

This is a valuable function and would be sufficient protection were it not for these facts:

1. Not all access to the Internet occurs through the firewall.

⁵ Source: Gartner Group, Conference Presentation, May, 1997.

⁶ Source: Gartner Group, Conference Presentation, May, 1997.

Users, for a variety of reasons ranging from naiveté to impatience, sometimes set up unauthorized modem connections between their systems connected to the internal network and outside Internet access providers or other avenues to the Internet. The firewall cannot mitigate risk associated with connections it never sees.

2. Not all threat originates outside the firewall.

A vast majority of loss due to security incidents is traced to insiders. Again, the firewall only sees traffic at the boundaries between the internal network and the Internet. If the traffic reflecting security breaches never flows past the firewall, it cannot see the problems.

As more organizations utilize strong encryption to secure files and public network connections, the focus of adversaries will shift to those places in the network in which the information of interest is not as likely to be protected: the internal network. Intrusion detection systems are the only part of the infrastructure that is privy to the traffic on the internal network. Therefore, they will become even more important as security infrastructures evolve.

3. Firewalls are subject to attack themselves

Attacks and strategies for circumventing firewalls have been widely publicized since the first firewalls were fielded. A common attack strategy is to utilize *tunneling* to bypass firewall protections. Tunneling is the practice of encapsulating a message in one protocol (that might be blocked by firewall filters) inside a second message⁷.

Who Guards the Guard? – Trust and Intrusion Detection

Another area of discussion when considering the value of intrusion detection systems is the need to monitor the rest of the security infrastructure. Firewalls, identification and authentication (I & A) products, access control products, virtual private networks, encryption products, and virus scanners all perform functions essential to system security. Given their vital roles, however, they are also prime targets of attack by adversaries. On a less sinister note, they are also managed by mere mortals, and therefore subject to human error. Be it due to misconfiguration, outright failure, or attack, the failure of any of these components of the security infrastructure jeopardizes the security of the systems they protect.


By monitoring the event logs generated by these systems, as well as monitoring the system activities for signs of attack, intrusion detection systems provide an added measure of integrity to the rest of the security infrastructure. Vulnerability assessment products also allow system management to test new configurations of the security infrastructure for flaws and omissions that might lead to problems.

System Security Management – A Process View

Securing systems is not a point fix. It is an ongoing process targeting a dynamic environment in which new threats arise daily. Figure 3 shows one view of security management.

⁷ Bellovin, Steven M., and Cheswick, William R., *Firewalls and Internet Security, Repelling the Wily Hacker*, 1994, Addison-Wesley Publishing Company, p 76.

Figure 3 – A process view of system security management



Prevention covers those proactive measures taken by organizations to mitigate risks to their system security. Much of the classic, government-sponsored work in computer security addresses this area by focusing on the design and implementation of more secure operating systems and applications software. Also covered in “Prevention” includes security policy formation, encryption, strong identification and authentication, and firewalls.

Functions in the detection phase are primarily provided by intrusion detection systems, although virus scanners also fall into this category. As pictured in the diagram, detection involves monitoring the targeted system(s), analyzing the information gathered for problems, then, based on the system settings, responding to the problems, reporting the problems, or both.

The results of the detection process drive the other two stages of managing security, investigating problems that are discovered, documenting the cause of the problem, and either correcting the problem or devising a means of dealing with it should it occur again. A common vision for future intrusion detection systems is that of performing these last two stages automatically, or else performing the functions internal to detection so well that the need for the last two stages is virtually eliminated.

The combination of investigation and diagnosis/resolution phases is often called *Incident Response* or *Incident Handling*. Organizations should specify policies, procedures, and practices to address this area, as it does the rest of security.

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

They CAN lend a greater degree of integrity to the rest of your security infrastructure.

Intrusion detection systems, because they monitor the operation of firewalls, encrypting routers, key management servers and files critical to other security mechanisms, 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.

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, even to expert system administrators and security officers. Intrusion-detection systems allow administrators and managers to tune, organize, and comprehend what these information sources tell them, often revealing problems before loss occurs.

They CAN trace user activity from the point of entry to point of exit or impact

Intrusion-detection systems offer improvements over perimeter protections such as firewalls. Expert attackers can often penetrate firewalls; therefore, the ability to correlate activity corresponding to a particular user is critical to improving security.

They CAN recognize and report alterations to data files

Putting Trojan Horses in critical system files is a standard attack technique. Similarly, the alteration of critical information files to mask illegal activity, damage reputations, or commit fraud is common. 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.

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.

These products offer extensive vendor support and turnkey design so that even novice security personnel can look for hundreds of problems by pushing a button. Some of these product offerings even offer automated fixes for the problems uncovered.

They CAN recognize when your system appears to be subject to a particular attack.

Vulnerability assessment products also allow the user of a system to quickly determine what attacks should be of concern to that system. Again, strong vendor support allows novice security personnel to reenact scores of hacker attacks against their system, automatically recording the results of these attack attempts. These products also provide a valuable sanity check for those installing and setting up new security infrastructures. It is far better for a system manager to determine that his firewall is incorrectly configured immediately than to discover this after an attacker has successfully penetrated it.

They CAN relieve your system management staff of the task of monitoring the Internet searching for the latest hacker attacks.

Many intrusion detection and assessment tools come with extensive attack signature databases against which they match information from your system. The firms developing these products have expert staffs that monitor the Internet and other sources for reports and other information about new hacker attack tools and techniques. They then use this information to develop new signatures that are provided to customers for download from web sites, downloaded to customers via encrypted e-mail messages, or both.

They CAN make the security management of your systems by non-expert staff possible.

Some intrusion detection and assessment tools offer those with no security expertise the ability to manage security-relevant features of your systems from a user-friendly interface. These are window-based, point and click screens that step users through setup and configuration in a logical, readily understood fashion.

They CAN provide guidelines that assist you in the vital step of establishing a security policy for your computing assets.

Many intrusion detection and assessment products are part of comprehensive security suites that include security policy building tools. These provide you easy-to-understand guidance in building your security policy, prompting you for information and answers that allow you to articulate goals and guidelines for the use of your computer systems.

Unrealistic expectations

They are not silver 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.

They CANNOT compensate for weak identification and authentication mechanisms

Although leading-edge research in intrusion detection asserts that sophisticated statistical analysis of user behavior can assist in identification of a particular person by observing their system activity, this fact is far from demonstrated. Therefore, we must still rely on other means of identification and authentication of users. This is best accomplished by strong authentication mechanisms (including token-based or biometric schemes and one-time passwords). A security infrastructure that includes strong I&A *and* intrusion detection is stronger than one containing only one or the other.

They CANNOT conduct investigations of attacks without human intervention

In very secure environments, incidents happen. In order to minimize the occurrence of incidents (and the possibility of resulting damage) 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. In some cases, especially those involving a dedicated attacker, finding the attacker, then pursuing criminal charges against the attacker is the only way to make the attacks cease. However, the intrusion-detection system is not capable of identifying the person at the other end of the connection without human intervention. The best that it can do is identify the IP address of the system that served as the attacker’s point of entry—the rest is up to a human incident handler.

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. These functions can not only spot the high-school hacker executing the “teardrop” attack against your file server, but also spot the programmer accessing the payroll system after hours. However, this policy compliance checking can exist only if there is a security policy to serve as a template for constructing detection signatures.

They CANNOT compensate for weaknesses in network protocols

TCP/IP and many other network protocols do not perform strong authentication of host source/destination addresses. This means that the source address that is reflected in the packets carrying an attack does not necessarily correspond to the *real* source of the attack. It is difficult to identify who is attacking one's system; it is very difficult to prove the identity of an attacker in a court of law—for example, in civil or criminal legal processes.

They CANNOT compensate for problems in the quality or integrity of information the system provides

In other words, “garbage in garbage out” still applies. System information sources are mined from a variety of points within the system. Despite the best efforts on the part of system vendors, many of these sources are software-based; as such, the data are subject to alteration by attackers. Many hacker tools (for example “cloak” and “zap”) explicitly target system logs, selectively erasing records corresponding to the time of the attack and covering the intruders' tracks. This argues for the value of integrated, sometimes redundant, informa-

tion sources; each additional source increases the possibility of obtaining information not corrupted by an attacker.

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. In fact, vendors themselves characterized the maximum bandwidth at which they had demonstrated their products to operate without loss with 100% analysis coverage at 65 MBPS.

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 IDSs' 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.

They CANNOT deal with modern network hardware and features

Dealing with fragmented packets can also be problematic. This problem has serious ramifications when one considers modern high-speed ATM networks that use packet fragmentation as a means of optimizing bandwidth. Other problems associated with advances in network technologies include the effect of switched networks on packet-capture-based network intrusion detection systems. As the effect of switched networks is to establish a network segment for each host, the range of coverage for a network intrusion system is reduced to a single host. This problem can be mitigated in those switches offering monitoring ports or spanning capability; however, these features are not universal in current equipment.

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.

The 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 also 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 2 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 3 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.

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 cen-

tralize 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 e-mail, 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.

About Vulnerability Assessment Products

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 at-

tacks, 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.

INTRUSION DETECTION AND VULNERABILITY ASSESSMENT: TECHNICAL CONCEPTS AND DEFINITIONS

The following terms explain the main concepts in intrusion detection, and will help to standardize the terminology and description of evolving products.

Intrusion Detection

Intrusion Detection Systems are security management tools that:

- Collect information from a variety of system sources,
- Analyze that information for patterns reflecting misuse or unusual activity,

- In some cases, automatically respond to detected activity, and
- Report the outcome of the detection process.

Descriptors for Intrusion Detection Systems Features and Functions

Monitoring Approach

Application-based

Application-based intrusion detection sensors collect information at the application level. Examples of application-level include logs generated by database management software, web servers, or firewalls. With the proliferation of Web-based electric commerce, security will increasingly focus on interactions between users and application programs and data.

Advantages of application-level monitoring:

- This approach allows targeting of finer-grained activities on the system (e.g. one can monitor for a user utilizing a particular application feature.)

Disadvantages:

- Applications-layer vulnerabilities can undermine the integrity of application-based monitoring and detection approaches.

Host-based

Host-based intrusion detection agents (also called *sensors*) collect information reflecting the activity that occurs on a particular system. This information is sometimes in the form of operating-system audit trails. It can also include system logs, other logs generated by operating system processes, and contents of system objects not reflected in the standard operating system audit and logging mechanisms.

Advantages:

- Systems can monitor information access in terms of “who accessed what”
- Systems can map problem activities to a specific user id
- Systems can track behavior changes associated with misuse
- Systems can operate in encrypted environments
- Systems can operate in switched network environments
- Systems can distribute the load associated with monitoring across available hosts on large networks, thereby cutting deployment costs

Disadvantages:

- Network activity is not visible to host-based detectors
- Running audit mechanisms can incur additional resource overhead
- When audit trails are used as data sources, they can take up significant storage
- Operating system vulnerabilities can undermine the integrity of host-based agents and analyzers
- Host-based agents must be more platform-specific, which adds to deployment costs
- Management and deployment costs associated with host-based systems are usually greater than in other approaches

Target-Based Approaches

Integrity analysis (see section 3.2.4.3) enables one to implement a focused and effective monitoring strategy for systems in which data integrity and process integrity are of primary concern. This approach monitors specific files, system objects and

system object attributes for change, looking at the *outcome* of attack processes rather than the *details* of the attack processes. Some systems use *checksums* (computations whose value depends on the original constitution of the system object) to detect breaches of integrity.

Advantages:

- Because it does not depend on historical records of behavior, integrity analysis may detect intrusions that other methodologies do not;
- This approach allows reliable detection of both placement and presence of attacks that modify the system (e.g., Trojan horses);
- Because its footprints and intrusiveness are low, this approach can be useful for monitoring systems with modest processing or communications bandwidth;
- This approach is effective for determining which files need to be replaced in order to recover a system, rather than reinstalling everything from the original source or backup, as is often done.

Disadvantages:

- Depending on the number of files, system objects and object attributes for which checksums are computed, this approach may still levy an appreciable processing load on low-end systems;
- The approach is not well suited to real-time detection processes, as it monitors for the outcome of attacks, not for the attacks themselves while they are in progress.

Network-based

Network-based intrusion detection sensors collect information from the network itself. This information is usually gathered by packet sniffing, using network interfaces set in promiscuous mode;

however, some agents are integrated in network hardware devices.

Advantages:

- The data come without any special requirements for auditing or logging mechanisms; in most cases collection of network data occurs with the configuration of a network interface card.
- The insertion of a network-level agent does not affect existing data sources.
- Network-level agents can monitor and detect network attacks. (e.g., SYN flood and packet storm attacks).

Disadvantages:

- Although some network-based systems can infer from network traffic what is happening on hosts, they cannot tell the outcome of commands executed on the host. This is an issue in detection, when distinguishing between user error and malfeasance.
- Network-based agents cannot scan protocols or content if network traffic is encrypted.
- Network-based monitoring and intrusion detection becomes more difficult on modern switched networks. Switched networks establish a network segment for each host; therefore, network-based monitors are reduced to monitoring a single host. Network switches that support a monitoring or scanning port can at least partially mitigate this issue.
- Current network-based monitoring approaches cannot handle high-speed networks.

Integrated approaches

Some intrusion detection products combine application, host, and network-based sensors.

Advantages:

- As agents at applications, host, and network levels are used, the system can target activity at any or all levels.
- It is easier to see patterns of attacks over time and across the network space; this is of value in damage assessment and system recovery; it also aids in investigating the incident and pursuing legal remedies (e.g. criminal prosecutions).

Disadvantages:

- There are no industry standards with regards to interoperability of intrusion detection components; therefore it is difficult or impossible to integrate components from different vendors.
- Integrated systems are more difficult to manage and deploy.

Timing of Information Collection and Analysis

Once the *location(s)* of intrusion detection system agents are established, the *timing* of the information collection and analysis are of interest.

Batch or Interval Oriented

In batch-oriented (also called *interval-oriented*) approaches, operating-system audit mechanisms or other host-based agents log event information to files and the intrusion detection system periodically analyzes these files for signs of intrusion or misuse.

Advantages:

- They are well suited to environments in which threat levels are low and single-attack loss potentials high (e.g., financial institutions). In these environments, users are often more interested in establishing accountability for problems

than immediately responding to suspected incidents. In this situation, batch-oriented analysis will likely be combined with other investigative process in order to identify the person responsible for the incident and support criminal prosecution for the incident.

- Batch mode analysis schemes impose less processing load on systems than real-time analysis, especially when collection intervals are short and data volumes are therefore low.
- Batch-oriented collection and analysis of information are particularly well suited to organizations in which system and personnel resources are limited. Organizations that have no full-time security personnel may find that real-time alarms generated by intrusion detection systems are seldom used. In such circumstances, it makes little sense to tolerate the processing load associated with real-time analysis and alarms.
- Attacks on computer systems often involve repetitive attacks on the same targets. For example, an attacker may enter a system via a password-grabbing attack, then install a Trojan horse “back door” in order to return later and continue the attack. Batch-mode analysis can usually recognize such *attack signatures*.
- Many current legal practices relating to computer evidence were established with batch-mode collection and manual analysis in mind. Therefore, it may be easier to submit system logs collected and processed in batch mode as evidence.

Disadvantages of batch-mode analysis include the following:

- Users will seldom see incidents before they are complete.

- Therefore, there is virtually no possibility of actively countering incidents as they happen in an attempt to minimize damage.
- The aggregation of information for batch-mode analysis consumes more disk storage on the analysis system. This can result in huge amounts of data for enterprise networks.

Real Time

Real time systems provide information collection, analysis, and reporting (with possible responses) on a continuous basis. The term “real-time” is used here as in process-control systems; that is, the detection process happens quickly enough to hinder the attack. Note that while this definition applies to systems that take milliseconds to perform analysis, it can also describe systems that are slower. Real-time systems provide a variety of real-time alarms (many support off-site alarming mechanisms such as email, pagers, and telephone messaging), as well as automatic responses to attacks. Typical responses range from simple notification to increasing the sensitivity of the monitoring, terminating the network connection from the source of the attack or changing system settings to limit damage.

Advantages:

- Depending on the speed of the analysis, attacks may be detected quickly enough to allow system administrators to interrupt them;
- Depending on the speed and sensitivity of the analysis, system administrators may be able to perform incident handling (leading to recovery of system operations) more quickly;
- In cases that occur on systems where legal remedies are available, system administrators may be able to collect information that allows

more effective identification and prosecution of intruders.

Disadvantages:

- They tend to consume more memory and processing resource on the analysis system than *post facto* systems;
- There are serious legal issues associated with automated responses that attempt to harm the attacking systems, a feature associated with some real-time systems;
- Configuration of real-time systems is critical; a badly formed signature can generate so many false alarms that a real attack goes unnoticed.

Location of Analysis

As in sensors, analysis functions can reside at host-level, at network-level, or both. Performing analysis strictly at the host level has the advantage of minimizing network load. However, it has the disadvantage of not allowing the detection of broad scale attacks targeting a network of machines (for instance, an attacker sequentially hopping through a network performing brute force password guessing against each host).

Consolidating raw data and performing analysis strictly at the network level (in the case of systems with sensors at both host and network levels) offer the capability to detect attacks that involve more than one host on the network. The disadvantage to this approach is that the network load associated with transferring raw host-level information to the analysis engine can be crippling.

As in sensor placement, the optimal strategy for performing analysis of logs is one in which analysis is done at both host and network levels. The analysis done at the host level can be simple or extensive depending on the nature of the sensor

information generated in that host or the signature against which the information is matched. The network-level analysis can take the results from the host-level analysis and use it to detect signs of network-wide attack or suspicious behavior without incurring as heavy a network load. Furthermore, in larger networks, this sort of approach can be applied hierarchically. That is, groups of hosts can report to a network analysis engine, which in turn reports its results to another analysis engine that collects results from a number of other network analysis engines and so on. This hierarchical structure lets intrusion-detection products succeed even in larger organizations.

Types of Analysis

Signature analysis

Signatures are patterns corresponding to known attacks or misuses of systems. They may be simple (character string matching looking for a single term or command) or complex (security state transition written as a formal mathematical expression). In general a signature can be concerned with a process (the execution of a particular command) or an outcome (the acquisition of a root shell.)

Signature analysis is pattern matching of system settings and user activities against a database of known attacks. Most commercial intrusion detection products perform signature analysis against a vendor-supplied database of known attacks. Additional signatures specified by the customer can also be added as part of the intrusion detection system configuration process. Most vendors also include periodic updates of signature databases as part of software maintenance agreements.

One advantage of signature analysis is that it allows sensors to collect a more tightly targeted set of system data, thereby reducing system overhead.

Unless signature databases are unusually large (say hundreds of thousands or millions of complex signatures), signature analysis is usually more efficient than statistical analysis due to the absence of floating point computations.

Statistical analysis

Statistical analysis finds deviations from normal patterns of behavior. This feature, common in research settings, is found in few commercial intrusion detection products. Statistical profiles are created for system objects (e.g., users, files, directories, devices, etc.) by measuring various attributes of normal use (e.g., number of accesses, number of times an operation fails, time of day, etc.). Mean frequencies and measures of variability are calculated for each type of normal usage. Possible intrusions are signaled when observed values fall outside the normal range. For example, statistical analysis might signal an unusual event if an accountant who had never previously logged into the network outside the hours of 8 AM to 6 PM were to access the system at 2 AM.

The advantages of statistical analysis are:

- The system may detect heretofore unknown attacks;
- Statistical methods may allow one to detect more complex attacks, such as those that occur over extended periods.

Disadvantages of statistical analysis (at this time) are:

- It is relatively easy for an adversary to trick the detector into accepting attack activity as normal by gradually varying behavior over time;
- The possibility of false alarms is much greater in statistical detectors;
- Statistical detectors do not deal well with changes in user activities (e.g., when the man-

ager assumes the duties of a subordinate in an emergency). This rigidity can be a problem in organizations where change is frequent. This can result in both false alarms and false negatives (missed attacks).

Integrity analysis

Integrity analysis focuses on whether some aspect of a file or object has been altered. This often includes file and directory attributes, content and data streams. Integrity analysis often utilizes strong cryptographic mechanisms, called *message digest* (or *hash*) *algorithms*, which can recognize even subtle changes.

Advantages:

- Any successful attack where files were altered, network packet grabbers were left behind, or rootkits were deployed will be detected regardless of whether or not the attack was detected by signature or statistical analysis.

Disadvantages:

- Because current implementations tend to work in batch mode, they are not conducive to real-time response.

Responses to Detection of Misuse or Attack

Some network-based intrusion detection systems permit one to specify a desired reaction to a detected problem. This feature has captured the imagination of many in the security management arena, especially as the frequency of *denial-of-service attacks* (saturation of system resources) has increased.

Alter the Environment

A typical response to a detected network attack is to take steps to alter the environment of the system under attack. This alteration can consist of terminating the connection used by the attacker and reconfiguring network devices to block further

access to the site from the same source address. The response mechanisms are intended to allow system administrators to take an active role within their authority to minimize damage associated with a detected attack.

Although it is a popular topic of discussion, striking back by attacking the source is ill advised at this point. TCP/IP, the basis for Internet communications, allows spoofing of packet source addressing; therefore, retaliation against the putative source of an attack might in fact damage an innocent party whose IP address had been forged for the attack.

Another valuable feature of intrusion detection systems is to drill down into information sources by setting agents and audit mechanisms to collect more information about the connection in question. This can also include collecting information that allows playback of attacks. This response allows the system administrator to collect information that supports more accurate judgements about the intent of the attacker. It also allows collection of information that might assist law enforcement or other investigators in identifying those responsible for the attack.

Validation

Knowledgeable attackers will often attempt to target the intrusion detection sensors or the analysis engine. In this case, a validation response, in which the sensors and/or analysis engine are queried in order to determine whether they continue to work properly, is suitable.

Real Time Notification

Finally, most real-time systems allow a system administrator to select a variety of alarm mechanisms to notify responsible parties of detected attacks. The alarms can notify key personnel by email or pager messages sent instantaneously with information about the problem. A message to the

system console is standard, and many systems allow a variety of visual and auditory signals as part of the alarm.

Management Functions and Deployment Issues

Customers need flexibility in adapting intrusion-detection systems to their own environments. The following features help to tailor these products to specific needs.

Configuration

No two organizations are the same. Each has a different set of security and management concerns driving security policy, a different set of hardware and software platforms included in their systems environment, a different set of users or a different set of operational policies. Therefore, the first issue facing a customer who acquires an intrusion detection system is the installation and setup of the system.

Many products, especially those designed for Windows NT environments, are shipped with clear, concise directions and installation scripts included. However, configuring these products is still an involved process. Information that customers must enter range from the IP addresses of the systems protected by the product to the sorts of security violations or system activities that the products are to detect and report. This is when a clear, current set of site security policy, procedures, and practices pays off handsomely.

Audit Subsystem Management

Products that include host-level agents typically use operating-system audit mechanisms. These products offer improved user interfaces to the operating-system audit controls, allowing users to specify what information is collected and how it is collected.

Reporting

One of the benefits of intrusion detection systems is the demonstration of due diligence in system security management practice. A key to demonstrating this due diligence (e.g., to upper management, internal auditors and regulatory personnel) is to document the findings of intrusion-detection products over a particular time interval.

Most intrusion-detection products have the ability to easily generate reports; many offer the capability to export report data to databases for subsequent analysis and archiving. Many offer multiple report formats (e.g., hard copy, screen, and HTML), with features allowing the user to report different layers of detail depending on the intended recipient of the report.

Control

Once the intrusion detection product is configured to the system environment, the next issue is actually running the system. Rudimentary controls include starting and stopping the system, establishing the schedule at which certain activities should take place, and specifying how alarms should be handled. In the control function, another critical issue in intrusion detection products is addressed: the security and reliability of the intrusion detection system itself. One way of addressing this is to require authentication before the system responds to control or configuration commands. This reduces the risk of an adversary gaining access to the system and shutting it down.

Proof of Validity

In some cases, intrusion-detection systems are used to ensure the operation of other parts of the security infrastructure (e.g., firewalls). In this proof of validity, the intrusion-detection system analyses information from both inside and outside the coverage area of the security mechanism in question, then compares results. The mechanism is proven valid when the intrusion-detection

system isolates evidence that an attack (sensed on the outside) is blocked by the mechanism (therefore not sensed on the inside). Vulnerability-assessment products are often used as part of this validation process and they function in synergy with intrusion-detection systems.

System Integrity

Given the role of intrusion-detection systems and the sensitivity of the information they sometimes contain, system designers have devoted considerable thought to the measures needed to protect the system itself. A standard strategy of attackers is to determine what security mechanisms are in place and then take steps to nullify or circumvent them. One can therefore assume that intrusion-detection systems will operate in a hostile threat environment. Consequently, many features are included in systems to minimize the chances that the system will be successfully defeated, or worse yet, will be used as a vehicle of attack.

Some vendors provide embedded license mechanisms to assure that only legitimate customers of the vendor can utilize the product. This reduces the risk that if the software is successfully stolen from the vendor or customer, the adversary could use it to monitor other machines or probe them for vulnerabilities.

Another protection strategy utilizes strong encryption to secure communications between sensors, analysis engines, and control consoles. This lowers the risk that an adversary might spoof the sensor output for a particular system in order to mask attack activity.

Some vendors use digital signature and message digest algorithms to protect signature database updates from tampering by adversaries. This allows time-sensitive distribution of new attack signatures to customers without the risk of corruption.

Other Features

Some vendors offer *decoy* server software to allow more accurate characterization of the threat levels for a customer's system environment. A decoy server is just that—a server that has no other purpose than attracting hacker attention. It is equipped with sensitive agents that collect information about the hacker's location, the path of the attack and the substance of the attack. The decoy collects and logs this information to a secure location. Some decoy servers also provide features that create “jail” environments to which hackers are redirected—environments in which their attacks cannot damage operational systems.

Vulnerability Assessment

Introduction

Vulnerability-assessment products (also known as *scanners*) are security management tools that:

- Conduct exhaustive checks of systems in an attempt to locate exposures to security vulnerabilities;
- Report the number, nature, and severity of these exposures
- Allow a system administrator to determine the security status of a system at a particular time
- Allow security auditors to determine the effectiveness of an organization's system security administration
- In some cases, once an incident occurs, allow investigators to determine the avenue of entry for an intruder or attacker

Vulnerability assessment products complement intrusion-detection systems: they allow system administrators to be *proactive* in securing their systems by finding and closing security holes before

attackers can use them. Intrusion-detection systems are by nature *reactive*: they monitor for attackers targeting systems in hopes of interrupting the attacks before the system is damaged.

Assessment approach

Application-based Assessment

Application-based assessment uses passive, non-invasive techniques to check settings and configurations within application packages for errors known to have security ramifications.

Host-based Assessment

Host-based assessment uses *passive*, non-invasive techniques to check system settings and configurations for errors known to cause security problems. These checks typically encompass system internals and include things such as file permissions and ownership settings and whether operating-system bug patches have been applied.

Most vulnerability-assessment products perform password analysis as part of their assessment. Password analysis consists of running *password crackers* against password files, utilizing a well-known attack in order to quickly locate weak, non-existent, or otherwise flawed passwords.

Advantages:

- It yields a very accurate, host-specific picture of security holes;
- It catches security holes that aren't exposed during a network-based assessment.

Disadvantages:

- The assessment methods are platform-specific and thus require precise configuration for each type of host used by the organization
- Deployment and update often require much more effort than in network-based assessment

Target-based Assessment

Target-based assessment (also known as *file integrity assessment*) uses *passive*, non-invasive techniques to check the integrity of system and data files as well as system objects and their attributes (e.g., hidden data streams, databases, and registry keys). Target-based assessment products use cryptographic checksums (message-digest algorithms) to make tampering evident for critical systems objects and files. Message-digest algorithms are based on hash functions, which possess the property that extremely subtle changes in the input to the function produce large differences in the result. This means that a change in a data stream fed to a message digest algorithm produces a huge change in the checksum generated by the algorithm. These algorithms are cryptographically strong; i.e., given a particular output value, it is practically impossible to come up with another input to the algorithm that will produce an identical output. This eliminates a common attack against relatively simple CRC (cyclic redundancy code) checksums in which hackers mask alterations to files by altering the content of the file so that the same checksum is generated for both the original and the tampered file. ⁸

Target-based assessment products run in a closed loop, processing files, system objects, and system object attributes to generate checksums; they then compare them to previous checksums, looking for changes. When a change is detected, the product sends a message to the intrusion-detection system that records the problem with a time stamp corresponding to the probable time of alteration. This process can provide a one-record trigger for an intruder alert or it can serve as a milestone for an investigator performing a trace of the events leading to the alteration.

Network-based Assessment

Network-based vulnerability assessment uses *active*, invasive techniques to determine whether a given system is vulnerable to a set of attacks. In network-based assessment, a variety of attack scenarios are reenacted against the target system(s), and results analyzed in order to determine the system's vulnerability to attack. In some cases, network assessment is used to scan for network-specific problems (e.g., port scanning.)

Network-based vulnerability assessment is often used for penetration testing (specifically, testing a firewall) and security auditing.

Advantages:

- It finds security holes on a variety of platforms and systems
- Because it is not as platform dependent as host-based vulnerability assessment, it is easy to deploy quickly
- As it does not assume host-level access, it is easier to deploy from a political point of view

Disadvantages:

- As it does not consider platform-specific vulnerabilities, it is often less accurate than host-based assessment
- It can affect network operations and performance

Integrated Assessment

Integrated vulnerability assessment combines both active, network-based assessment with passive, host-based assessment techniques, often combining them with a centralized management function. We note here that Windows NT environments do not recognize as crisp a policy boundary between host and network-based access.

Advantage:

- It combines the host-based advantages of improved identification of platform-specific vulnerabilities with the network-based capabilities to identify problems across wide ranges of affected systems and networks.

Disadvantage:

- The effort required to deploy and maintain the combined assessment engines is greater.

Location of Analysis

Collecting data is the first step in vulnerability assessment; data analysis is the second. In large complex network installations, it is helpful to organize vulnerability assessment using a console-agent architecture. This architecture is particularly helpful where networks are heterogeneous, i.e., with a wide range of operating system platforms.

Advantages:

- Centralized architectures can tailor agents to specific operating system platforms, and vary the coverage and rigor of assessments based on the threat environment
- Distributed architectures allow one to scan across network or NT policy domains

Disadvantages:

- Distributed architectures require additional remote privileges on the networks scanned

⁸ Garfinkel, Simson, and Spafford, Gene, *Practical UNIX and Internet Security, Second Edition*, Sebastopol, CA, O'Reilly and Associates, 1996.

Reporting

Reporting in vulnerability assessment holds the key to understanding and rectifying security holes. Reporting provides the opportunity to document the security health of the systems scanned, to publish problems to an appropriate level of management so that resources and responsibilities are assigned to fix them, and to educate everyone in the organization about the importance of system security and how to achieve it. Options provided include variable reporting formats (with HTML offering the ability to selectively “drill down” to a finer level of detail as desired) and levels of detail, providing different amounts of background information about the vulnerabilities and associated fixes.

Deployment

Although it is easy to understand the requirements driving vulnerability-assessment products, and easier still to understand how the products might be used to support an organizational security strategy, perhaps the most critical features in selecting a product are those regarding the deployment of that product in an operational environment.

- Most products provide user-friendly installation features, supported by automated scripts and strong technical support.
- Configuration options for products vary widely. Look for features such as network mapping, menu-driven configuration of security checks and network coverage, and on-screen help mechanisms.
- Most products allow system administrators to set up schedules for scanning, an option that allows them to schedule assessments for hours of low system utilization (e.g. outside business hours).

- Regular updates to the security check database is critical as new security holes are discovered every day. Update processes that are automatic and data-driven minimize the time required to incorporate new updates.
- Look for products that provide support of routine, repeatable scanning. Some support this with *differential scan* capabilities, which allow users to automatically compare results of successive scans, pointing out problems and inconsistencies that surface.
- Vulnerability-assessment products can provide features that should be used only with caution (e.g., security checks addressing network denial-of-service attacks). These checks, in replicating the attacks, can crash targets. Products should inform users of this problem when they select the denial-of-service checks.

Responses

Once the user has run vulnerability-assessment tools and spotted vulnerabilities, the user can specify responses. The response options provided include the following:

- Alarm mechanisms allow the system to send real-time alerts via a variety of means (e.g., SNMP, pager, email, etc.) of high-risk vulnerabilities that have been discovered.
- Report mechanisms allow the system to generate organized reports itemizing the results of vulnerability assessments.
- Some products have the ability to respond to detected security holes by actively closing them (by either amending file configurations or settings or else applying security patches) rather than simply reporting their existence.

Management Functions

As in intrusion detection, vulnerability-assessment products have various management functions:

Exporting data in a variety of formats (HTML, Crystal Reports, ODBC, MDB, etc.) allows system administrators and managers to utilize a variety of reporting tools to further analyze the results of the vulnerability assessment.

Network mapping makes it much easier to specify which hosts are to be scanned. With network mapping, one can do this selection by point and click selection of targets. Without it, manually entering all the addresses of hosts to be scanned can be an arduous, time-consuming process.

The capability to tailor the coverage of an assessment to a target is an important management function. This might include the ability to configure which checks runs against which targets, to add custom user-defined checks, and to configure certain parameters for individual checks.

System Integrity

As in intrusion detection, there are special security considerations associated with the design, deployment, and maintenance of vulnerability assessment products.

- Protection issues: The database of security checks must be protected, so that it does not become a primer for attackers. This can be accomplished by a variety of strategies; encryption of contents is perhaps the most common.

When encryption is used, however, U.S. government export control policy for encryption technologies can affect those measures available for products fielded outside the country.

- As new attacks surface daily, product vendors must provide means for customers to update

the lists of security checks performed by vulnerability assessment products. This update process must, itself, be protected. In distributed architectures, the communications between console and agent must be protected, and using cryptographic techniques may provide this protection.

- As vulnerability assessment systems can themselves be used by attackers to identify targets, there must be countermeasures to prevent this malicious use. These measures can include the broadcast of the identification of the source address of the scanning host to the target, and strong licensing mechanisms that limit the coverage of the scanner.

SUMMARY AND CONCLUSION

Wide range of goals for product users

Users of intrusion detection products span public and private institutions, running the gamut of industries. The goals realized by users of intrusion detection systems include:

- Support of internal audit
- Control of liability exposure
- Incident handling and investigative support
- Improved damage assessment and recovery
- Improved security management process
- Discovery of new problems/issues before damage occurs
- Documentation of compliance with legal and statutory requirements
- Recovery of systems suffering security violations

Developments in Other Security Product Lines will Increase the Importance of Intrusion Detection

Encryption is growing in popularity and products including encryption features are becoming ubiquitous. As more organizations utilize these products to secure their data as it travels over public networks, adversaries will adapt their attack strategies to accommodate this. The predictable outcome is that attacks will shift to those areas in which data is not encrypted: the internal network.

At the same time, corporate employment practices will continue to focus on outsourcing, strategic partnerships with other organizations, and telecommuting. All of these typically involve remote access to the internal network, thereby expanding the security perimeter of the organization to areas not physically protected.

Intrusion detection systems are the only part of the IDS/Firewall protection infrastructure privy to the traffic on the internal network. Therefore, they will become even more important as security infrastructures evolve.

Capabilities for Intrusion Detection Products are improving

The capabilities for intrusion detection are growing, as new products enter the marketplace, and existing organizations expand their product offerings to allow additional sensor inputs, improved analysis techniques, and more extensive signature databases.

Thanks to government and military interest in Information Warfare, of which Intrusion Detection is a vital defensive component, funding of research efforts has skyrocketed, with no end in sight. This increased activity will result in enhanced understanding of the intrusion detection process and

new features in future products. Plans are afoot to embed intrusion detection products as standard components of major governmental and financial networks.

As intrusion detection remains an active research area, look for future products to implement new techniques for managing data and detecting scenarios of interest. Also look for additional products that function at application level and that interoperate with network management platforms. Finally, look for product features that are integrated into a bevy of special purpose devices, ranging from bandwidth management products to “black box” plug-ins for targeted environments.

GLOSSARY

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.

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.

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.

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.

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.

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.

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.

SENSOR - The ID component that periodically collects data from the data source. Also known as **AGENT**.*

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**.

* In many existing ID systems, the sensor and the analyzer are part of the same component.

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.

[Format shows **Author, initial** (year). Title. Publication (date) volume (number): page (article length)]

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 simple 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

ABOUT THE AUTHOR

Rebecca Bace is the President of Infidel, Inc., a security consulting firm headquartered in Scotts Valley, California. Ms. Bace has worked as a manager and researcher in intrusion detection and network security since 1989, when she became the leader of the National Security Agency's Computer Misuse and Anomaly Detection Research program. This program funded many of the early research efforts in intrusion detection and vulnerability assessment, including the work performed at the University of California, Davis, and Purdue University. In 1996, Ms. Bace went to Los Alamos National Laboratory as Deputy Security Officer for the Computing, Information, and Communications Division. In 1998 she formed Infidel, a strategic consulting practice specializing in intrusion detection and network security. She holds the M.E.S. in Electrical Engineering from Loyola College. Ms. Bace is currently writing a book on the history, theory, and practice of intrusion detection, to be published in late 1999.