

# Next Generation Intrusion Detection in High-Speed Networks

Security Technologies for Advanced  
Network Counter-Insurgence

## Table of Contents

Introduction .....	2
The Importance of Intrusion Protection .....	3
Firewalls .....	3
Intrusion Detection .....	4
The Limits of Network-Based Sensors .....	5
Too Much Traffic to Watch Well .....	5
The Inability to See All The Traffic .....	5
Fail-Open Architecture .....	6
Inability to Evaluate Impact of Suspect Packets .....	6
Not Enough Information .....	6
Failure to Detect Certain Attack Types .....	7
Too Many False Positives .....	8
Requirements for Next-Generation Solutions .....	9
Pro-Active Preventive Probing via “Scanners” .....	9
Planning A Strategy .....	10
Defect Management .....	10
Real-Time Host-Based Monitoring .....	11
Monitoring The Data Stream .....	12
Second Generation Anomaly Analysis .....	12
Centralized & Collaborative Aspects .....	13
The Concept of Decoy “Sting” Servers .....	13
“Virtual” Network Forensics .....	14
Sacrificial Hosts .....	14
Network Associates’ CyberCop Intrusion Protection Family .....	15
Conclusion .....	16

The information in this guide has been provided by Networks Associates Technology, Inc. To the best knowledge of Network Associates, these companies offer the types of products described. These companies are solely responsible for their software, distribution, and support services. Network Associates disclaims any and all liabilities for and makes no warranties, expressed or implied, with respect to these products, including, without limitation, the implied warranties of merchantability and fitness for a particular purpose. Distribution of these products, or information concerning these products, does not constitute Network Associates’ endorsement of the products, the companies, or support services. Product information is subject to change without notice.

## Introduction

As companies interconnect more of their mission-critical computers, internally and to the Internet, the dangers of attacks by intruders become increasingly important, and the scale of potential damage also rises. One relatively new technology now being deployed as an essential security tool in today's enterprise networks is Intrusion Detection.

In today's increasingly interconnected world, it has become extremely difficult to detect security breaches without the deployment of an intrusion detection system (IDS). There are several different methods for detecting unauthorized or attempted access currently on the market. However, none of these first-generation products, are flexible enough to fully address the high speeds and modern topologies present in most large networks today.

The purpose of this paper is twofold—to briefly describe the inherent limits of these first-generation Network-based Intrusion Detection Systems (NIDS), relative to the network architecture they are asked to protect, and to introduce the next generation of intrusion detection technology.

## The Importance of Intrusion Protection

One unfortunate, but inevitable, consequence of the increasingly rapid adoption of computer, networking and Internet technologies is the likelihood that there will be more security vulnerabilities in your company's network and computing infrastructure. Between configuration settings, policy decisions, software patches, and the sheer number of devices in most corporate networks, it is simply not practical to maintain network security today without some form of intrusion protection. In terms of security-related bugs alone an average of twenty to forty new vulnerabilities in commonly used networking and computer products are discovered each month.

At the same time, the tremendous opportunities of e-business ensure that this problem will not go away. Companies are connecting their networks and computers to the Internet for communication, access, commerce and other activities. All these Internet activities are legitimate and necessary for business, but Internet connectivity also increases the potential risk of these systems for unauthorized access and other attacks.

Meanwhile, network attackers are becoming both more numerous, more capable, and more organized. Unfortunately, increasingly inexpensive hardware and easy to replicate "cookbook" invasion techniques have actually lowered the bar for trying to "crack" networks and computers. The tools to probe and attack networks are often automated and GUI-driven, and lists of common weaknesses are readily available via the Internet. As a result, any would-be intruder can now obtain powerful hacking tools and information in a matter of hours.

### Firewalls

Firewalls are essential to any network security system. They are, however, only part of the solution. If a firewall is overly restrictive, it may interfere with the performance or ability to connect for legitimate users. This forces users to either accept constraints on the company's ability to do business, or results in some users seeking alternative solutions which may pose far greater security risks than going through the firewall. Most networks are also full of unintentional "back doors" around the firewall (e.g., forgotten dial-in modems, X.25 connections to legacy systems, somebody running an unsecured proxy server, etc.). More than half of all breaches today originate from someone already legitimately behind the firewall.

## Intrusion Detection

“Intrusion detection” is a type of network security that, as the name implies, attempts to detect, identify and isolate attempts to “intrude” or make inappropriate, unauthorized use of computers. Attacks originate either via an external network connection, or from within your own organization. Target systems are usually server or workstation systems, however attackers may also focus on network devices such as hubs, routers and switches.

An intrusion detection system (IDS) helps identify the fact that attacks are occurring. It may also be able to detect attacks that other security components don't see and help collect forensic evidence which can be used to identify intruders. (Simply knowing that they can be identified may deter some attackers.) Current network IDS products use a predominantly passive approach to collecting data via protocol analysis garnered by watching traffic on the network. Each network IDS (there may be many IDS products on a company's network, one for each segment that the company wants to monitor) attaches to, and monitors the traffic on specific network segments. The IDS gets copies of its segment's traffic to inspect by “listening in promiscuous mode” and having its network interface card bring in a copy of every packet it sees. The IDS examines these packets, and attempts to determine whether they represent an intrusion attempt. It does this by seeing if the contents of the packet contain the “signature” of a known attack method, that is, whether it contains a string of characters that matches a specified pattern, or otherwise fits rules that define known attack methods.

Intrusion detection products are based on the assumption that an intruder can be detected through an examination of network traffic and of various system events such as CPU utilization, system calls, user location, and various file activities. Network sensors and system monitors convert observed events into chronologically sorted records of system activities. Called “audit trails,” these records are analyzed by IDS products for unusual or suspect behavior. IDS approaches include:

- **Signature-Based Intrusion Detection**—Signature-Based Intrusion Detection is based on the assumption that intrusion attempts can be characterized by the comparison of user activities against a database of known attacks that lead to compromised system states. Most commercial intrusion detection products perform signature-based intrusion detection against properties that initiate rules when audit records or system status information begin to indicate illegal activity. These predefined rules typically look for high-level state change patterns observed in the audit data compared to predefined penetration state change scenarios. In general, a signature can be concerned with a process or an event.

- **Statistical-Based Intrusion Detection**—Statistical-Based Intrusion Detection systems seek to identify abusive behavior by noting and analyzing audit data that deviates from a predicted norm. Statistical-Based Intrusion Detection systems are based on the premise that intrusions can be detected by inspecting a system's audit trail data for “out of the ordinary” activity, and that an intruder's behavior will be noticeably different than that of a legitimate user. Any sequence of system events deviating from the expected profile by a significant amount is flagged as a potential intrusion attempt.

## The Limits of Network-Based Sensors

Today's first generation network intrusion detection systems (NIDS) were designed to protect networks by attempting to watch *all traffic on a network* for signs of attack. Although this network-based approach to intrusion detection showed initial promise, it is now running up against some critical limitations. Additionally, within the past year, some new types of attack techniques have been identified which a network-based IDS products simply can't detect. (See “Insertion, Evasion and Denial of Service: Eluding Network Intrusion Detection”, Ptacek and Newsham, Secure Networks, Inc. 1998) The following is based on the current limitations of network sensor-based products, and the methodology of detection employed which is reliant on an inherently flawed approach to the problem.

### Too Much Traffic to Watch Well

As networks get faster, network-based IDS products simply cannot keep pace. Examining the contents of each packet, and seeing if it matches any of the known signatures and rules, takes time and consumes resources. A network IDS can examine all of the traffic on a 10Mbps LAN and check it for up to a dozen signatures. However, many of today's LANs are running at far higher speeds—50, 100 or more megabits per second. And network speeds keep growing faster than the technology for high-speed packet signature analysis. Today's network IDS can't reliably watch more than 10-20Mbps of traffic; you run the risk of it losing data, or being unable to watch for more than just a few signatures. Although NIDS technology will keep improving, and speeding up, the speed of networks to watch will continue growing as well.

### The Inability to See All The Traffic

Many companies are beginning to use “switched Ethernet” technology to architect their Local Area Networks. As network intrusion detection technology cannot reliably watch this traffic on the wire, another method employed is to connect the NIDS to the spanning port of the switch. This unfortunately, is also not the answer. Implementing an NIDS on the switch will considerably degrade the packet throughput (as the hub has to wait for the spanning port

to catch up) before sensing the packet. This means that implementing a network IDS attached to the spanning port of a switched hub would defeat the purpose of a switched, high speed network in the first place.

### Fail-Open Architecture

Certain types of network sensor-based security systems; when they fail (due to overload, crashes, or Denial of Service attacks) leave the network they were guarding “open”, often without notification of the problem to the central console. The only other option for a sensor is to “fail-closed” impacting a company’s essential network services until the sensor is brought back “on-line”. Fail-closed architectures are possible on host-based systems (a stipulation in firewall technology for example), if the authentication module of the OS integrates with the host-based intrusion detection system and ensures that the last set of “rules” implemented stay in effect until the administrator resets the system locally. This is, in effect a “personal firewall” for critical systems requiring the utmost security and data integrity.

### Inability to Evaluate Impact of Suspect Packets

A network-based IDS can’t predict whether a given destination machine will see a suspect packet, (e.g., an IP packet with a bad UDP checksum) and, if seen, whether it would be processed as expected by the network IDS. This means the IDS must inspect all packets, which in turn may overload it. Obviously in network communications, packets can be sent unreliably or duplicates may be sent. When network protocols such as IP receive duplicate data, the IDS must choose either the old or new data. Different implementations of IP from different OS vendors make different decisions about this issue.

### Not Enough Information

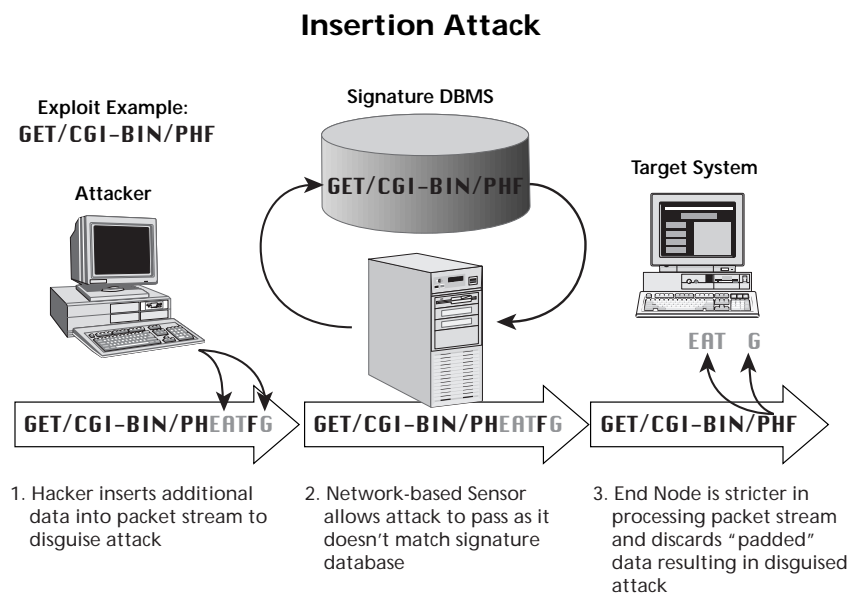
A network IDS can’t predict the implication of a packet just by looking at it. It also has to have information about the network segments, the end systems, etc., none of which is provided by simple packet capture. Trying to provide this information may be too much work for your organization, and trying to include this information in high-speed attack monitoring may not be possible

Also, because an NIDS is usually on a dedicated machine, rather than on one that it is watching, differences between the NIDS host and protected hosts in hardware, network drivers, etc., can lead to discrepancies in what may work as an attack.

## Failure to Detect Certain Attack Types

Perhaps most significantly, as has been shown recently (Ptacek/Newsham); there are at least twenty-six different techniques for executing a single attack that may elude the detection of first-generation network IDS products. Three classes of attacks have been found which exploit the fundamental nature of network IDS-based on IP and TCP protocol analysis. These attacks either evade signature recognition, or consume enough resources to disrupt/disable the network IDS. In tests, all network IDS products on the market today proved vulnerable to each type of attack. The implication is that, *short of some fundamental redesign, today's network ID systems cannot claim to offer significant intrusion protection for your network.*

For example, one form of insertion attack inserts extra characters into the packet stream, which keeps the contents of the packets from matching an attack signature. E.g., data is sent one character per packet, and the "exploit" string "GET /cgi-bin/phf" is masked by inserting padding characters, so the network IDS sees a pattern such as "GET /cgi-bin/pheatf". The end node discards the padded data due to processing and the assumption that the additional characters are caused by inherent noise on the line, resulting in the target system "recreating" the original exploit string.



Similarly, since TCP/IP reassembles data streams using sequencing numbers in the packets, one "evasion" technique is for an attacker's packets to be sent out of sequence. So, again, what the network Intrusion Detection Sensor sees doesn't look like an attack... but when these packets are reassembled by the target system, they contain the original attack.

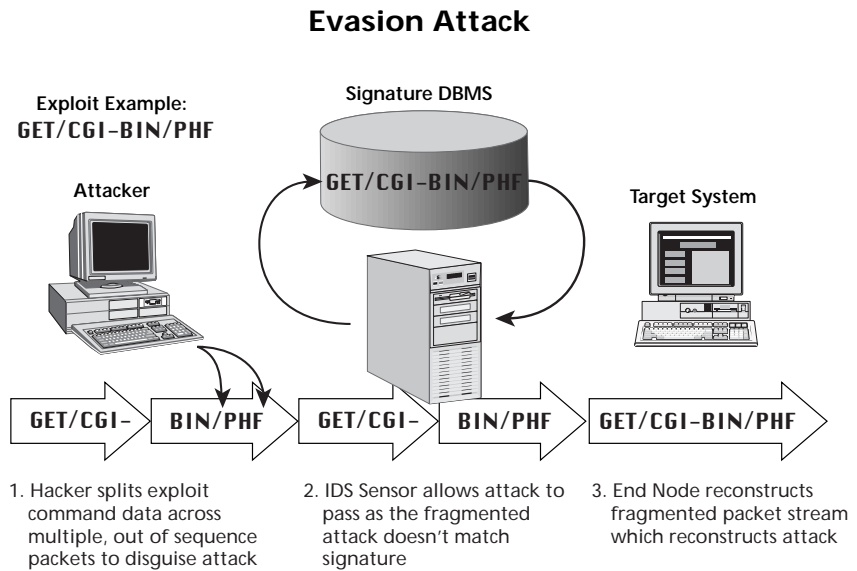
FIGURE 1

depicts a step by step insertion attack: The TCP/IP connection is "hijacked," the packet sequence determined, and the attack inserted into the data stream.



FIGURE 2

outlines the steps taken to produce an evasion attack which also circumvents a Network-based Intrusion Detection sensor.



Additionally, the paper by Ptacek and Newsham identified numerous Denial of Service (DoS) attacks against either the network IDS sensor itself, or a host that cannot be detected or prevented by a network IDS. DoS attacks are intended to compromise a device's availability, by making them too busy, crashing them, etc. Common network DoS attacks include mail bombs, ping floods, and attacking known software bugs. This is not just an academic statement; reports show these types of attacks with increasing frequency. This fact alone means that traditional firewalls performing packet analysis using rules and patterns are no longer sufficient to ensure safety from network-based attacks.

#### Too Many False Positives

Still another problem with NIDS products is that they are prone to "false positives" (perceived threats that appear to the NIDS to be real, but are just normal data transactions). It's easy for an NIDS, for example, to produce an alert when a ping of death attack occurs. If the NIDS produces an alert every time it sees any kind of ping, however, it may produce the opposite effect. In fact, knowledgeable invaders have been known to create a "boy who cried wolf" scenario by generating so many alerts that appear to be false positives that network administrators simply filter out or ignore these alerts, possibly allowing serious attacks to go unnoticed. There is currently no benchmark for evaluating the signal-to-noise ratio produced by NIDS products.

## Requirements for Next-Generation Solutions

The implications of new network architectures, the limits of network-based intrusion detection systems, and the capabilities of new threats are clear: *companies need new intrusion detection tools.*

These new tools need to be:

- Compliant with high-speed, switched networks
- Compliant with encrypted or VPN traffic
- Deployable selectively, based on system vulnerability/priority
- Able to detect attacks that network-based intrusion detection miss.

Given the potential impact of successful attacks, the number of possible attacks, and the on-going discovery of new attracts and vulnerabilities, there are other intrusion detection-related capabilities which companies should also consider, such as:

- The ability to test, *in advance*, one's own network and computer devices, to determine whether security holes exist, and begin fixing them before intruders attack them
- Some way to divert intruders who can't be kept out to less valuable areas of the corporate network.

## Pro-Active Preventive Probing via "Scanners"

Just as we know specific information about computer viruses, there is a similar body of knowledge regarding known security vulnerabilities and attack methods, many of which are utilized by would-be network intruders. These weakness may be as simple as whether the password on a field service account has been left unset, or as complex as what level of a certain type of attack a network-based IDS can sustain before failing. Categories of typical vulnerabilities include misconfigurations, policy violations, and software updates (e.g., bug fixes, new features).

Given that this knowledge exists, it makes sense for a company to focus *first* on these known vulnerabilities by regularly testing its firewalls, routers, hosts and other devices using a firewall auditing tool capable of testing invalid packet throughput. An good example of this would be the tracer packet firewall tests included in Network Associates' CyberCop Scanner product. With these vulnerabilities identified, the company can then make decisions what to do, rather than wait for attackers to identify and take advantage of weaknesses.

## Planning A Strategy

Using a proactive scanning tool and a database of vulnerabilities to test against, an organization can conduct its own routine checks of routers, firewalls, network devices, operating systems, web servers, applications, workstations, and even network intrusion detection systems. This provides a company with a more proactive approach to security, rather than simply waiting for network intruders to locate and take advantage of weak points.

Network Associates currently documents over 600 tests that can be run by such a scanning engine to perform security testing, including the 26 recently-discovered techniques to attack first-generation network intrusion detection systems. In our experience, it takes only a matter of minutes to conduct the full set of tests against a single device. By running multiple scanner engines in parallel, tests of larger networks and larger numbers of hosts can be done even more efficiently. It may make sense to approach scanning as a multi-step process, especially given that your company is probably bringing in new machines and upgrading or replacing existing ones on a weekly, daily or even hourly basis. Your company can then decide what to fix, based on factors including how difficult/expensive it is to resolve the problem, how vulnerable the system is otherwise, and how important the system/application/data is.

## Defect Management

Of course, to be meaningful at the enterprise level, testing needs to be conducted on a regular basis. Tests should be initiated after fixes, patches or other countermeasures are applied, after any meaningful change or upgrade is done to hardware or software, as new vulnerabilities are discovered, and on a scheduled basis, e.g. monthly or quarterly. This on-going testing should be viewed as one aspect of the defect management process that many companies have for mission-critical resources and services, including their networks as well as manufacturing facilities, delivery, and support. Assuming the database is comprehensive, scanning, followed by applying the resolution advice, will close a significant percentage of security vulnerabilities, and give your company a good sense of what isn't or can't be fixed. This in turn helps identify where to apply more resources.

Ideally, you will also be able to customize the scanner's vulnerability database, and any testing scripts for scanners, to reflect not only your company's equipment but also its resolution policies and practices, and those of any consulting partners involved in security management.




FIGURE 3

Regular proactive scanning is always the first step in intrusion protection. Real-time monitoring should be used to watch vulnerabilities that cannot practically be closed.

Additionally, there is always the likelihood of security vulnerabilities that scanning may catch but you won't be able to close; e.g. design issues, misuse, unsecured architectures, etc. Since these potential weaknesses cannot be "closed," these must be dealt with in other ways. In many cases, the unsecured resource will be one you can't simply remove. (Figure 3) These are good places to add *host-based* monitoring, to make sure that attack attempts are detected should they occur.

## Real-Time Host-Based Monitoring

If you want to know what is happening from moment to moment to the systems on your network, there is no substitute for monitoring them *directly*. In addition to the many shortcomings of network-based IDS products, some attack behavior can only be seen from inside the host and viewed directly from the OS and kernel rather than via a network session. *The best solution is to deploy lightweight intrusion detection agents directly on each of the computers you want to watch.*

Because the sensor is inside the host, it can also observe events and system behaviors, including some which are difficult or impossible to see from a telnet, login or shell session. This in turn makes it possible for the sensor to watch and analyze events like login events, and also watch system behavior. By comparing these against a database of rules, such as IP addresses or protocol requests to block, and thresholds, such as a number of failed login attempts, the sensor can identify possible intrusion attempts. The response may be any combination of logging, alerting, or other activities.

Host-based monitoring software, can easily “listen to” 100% of the traffic coming into the device. By looking at the data stream emerging from the kernel and the protocol stack, the monitor is, by definition, viewing exactly what the host is receiving. Any masking techniques such as insertion, padding, fragmentation, or out-of-sequence delivery, which would evade a network-based IDS can be easily caught by a host-based IDS. A good host-based IDS will actually contain several “engines” which have the ability to communicate state information between each other. One examining the packet stream for attack signatures, the second watching the log file for misuse activities, and a third one watching system behavior for anomalies based on thresholds of system metric, while the overall system digests all of the events and detects very high level signatures. A host-based monitoring installation like this can be implemented with a reasonably small footprint, using low CPU resources and minimal memory space.

### Monitoring The Data Stream

There are currently several hundred “signatures” of attacks at the packet level; the host monitor can check for any or all of these, and respond according. For example, if the sensor detects what appears to be a Denial of Service (DoS) attack, or a UDP (Unreliable Datagram Protocol) based attack, it can cause all this traffic to be discarded. Additionally, host monitoring can also be used to provide filtering security such as restricting connections in or out based on rules, whether users are authenticated, what applications they are trying to access, etc.

### Second-Generation Anomaly Analysis

Host-based IDS products can also use far more efficient intrusion detection techniques such as second-generation “anomaly analysis.” An anomaly is an event, or threshold, that differs from what was expected, or what has been typical. These behaviors and thresholds are defined as heuristics (rules); the rules may be defined externally by network administrators. If the sensor is sophisticated enough, it may also develop part of its behavioral database by watching users over a period of time and charting what it believes to be normal. If the charted behavior is deviated from enough, this can trigger an alert. Anomaly analysis at this level is virtually impossible for network-based IDS products trying to watch everything on the wire.

One example of a rule might be, “User A has logged in from Point B for the past six weeks, and done these things.” If it appears that User A is now logging in from Point C, and trying to do different things, or has failed to login seven times in a row, within minutes, this

would be identified as an anomaly. A more complex rule might be, “If events A, B, and C happen, don’t worry, but if A, B, C, and D happen, notify the network administrator.” An alerting response can be any of a number of activities, ranging from turning on logging to setting off an alarm, sending a message to a pager, sending e-mail, doing an SNMP alert, etc.

### Centralized & Collaborative Aspects

This distributed approach has additional advantages, namely that it doesn’t depend on communications with or from a central console to react to potential threats. Although host-based monitors can and will operate independent of each other or any central console, there are some centralized and collaborative ways they can also work:

- **Receiving Database Updates**—New patterns and rules can be distributed over the network, on a regular and ad hoc basis, enabling a host IDS to watch for new attack patterns or methods as they are identified.
- **Sharing Event and Action Information**—Host monitors can also exchange information with each other, and accept this as new rules. For example, information about possible attacks detected, or actions taken such as locking out a user account that attacks appear to be being launched from.

Network Associates CyberCop Monitor, for example, employs a central console to which the individual monitors and other systems report. A coalescing function on the report database watches for, and prevents, repetition in the log files. This prevents having hundreds/thousands of entries in the log files that would skew all charts and graphs possibly leading to an incorrect assumption regarding the security of the network. As you can see, host-base sensors make it possible to monitor today’s higher levels of traffic for attack attempts, and to see attacks which may elude or take out a network IDS, as well as detect many types of attacks that a network IDS would never be able to detect.

### The Concept of Decoy “Sting” Servers

Even with proactive vulnerability scanning and second-generation host-based monitoring, some companies want to look for signs of potentially dangerous activity before it becomes a problem.

For such companies, it may make sense to arrange for some zero-value resource which will attract attackers, diverting them from more valuable, sensitive systems, and, ideally, making it possible to collect data which may help identify them for prosecution. It would also be a great benefit in identifying the method they used to gain access so it can be closed off.

One classic technique for diverting and identifying intruders who are insiders, already inside the firewall, or external users who do succeed in penetrating the firewall, is to set up a “sting” area as a decoy. The concept of “sting” or decoy systems within a trusted environment is not new. In his book “The Cuckoo’s Egg” (Doubleday, 1995) astrophysicist Cliff Stoll talks about how he employed a version of this technique, using “sting” data to help detect and identify international network spies.

#### “Virtual” Network Forensics

Rather than dedicate lots of machines to this purpose, it’s possible to set up a “sting host” which, to the network and intruders on it, appears to be a number of host computers and network devices, such as Cisco routers, NT or Solaris servers, on one or more equally non-existent network segments. These fake systems will have entries in your DNS (Domain Name System) tables, and will typically have “tempting,” real world names like “Accounting-1,” “Payroll-7,” etc. The fake systems should, of course, give the appearance of being very secure, which in turn implies that they house something valuable. The sting area can only be found through intrusive attacks; by definition, anybody attempting to access these devices is doing something inappropriate. Only someone deliberately looking for a way to gain unauthorized network access would find, and try to access, these apparent network devices and computers.

A decoy network can be used to identify inappropriate internal users, who are already inside the firewall. It can also help identify the more elite attackers who are able to penetrate the firewalls, evade network ID systems and possibly compromise host-based IDS products. The sting server has the capability of generating alerts directed to network administrators or other personnel that a potential attack is underway, and begin logging events of the apparent intruder’s activities.

#### Sacrificial Hosts

If you want to go even one step further, sting systems can also be set to redirect intruders to “sacrificial hosts” that are actual computers with dummy data. When an intruder begins accessing one of these sacrificial hosts, logging can be turned on to record evidence of attempted tampering such as attempts to manipulate or destroy data.

## Network Associates' CyberCop Intrusion Protection Family

To address these new challenges and requirements in intrusion detection, Network Associates offers a complete line of intrusion protection products:

- **CyberCop Scanner**—CyberCop Scanner proactively examines computer systems and network devices for security vulnerabilities in enterprise network environments. CyberCop Scanner enables security professionals to test NT and UNIX workstations, servers, hubs, switches, and includes Network Associates' unique tracer packet firewall test to provide thorough perimeter audits of firewalls and routers. Also included in CyberCop Scanner is CASL (Custom Audit Scripting Language), a custom toolkit developed by Network Associates to enhance system security and simplify the creation of complex vulnerability tests. Using a simple GUI interface, security specialists may create unique simulated attacks against network devices to discover new vulnerabilities within their enterprise. Report options include executive summaries, drill-down detail reports, and field resolution advice. Most importantly, CyberCop Scanner uses ground-breaking AutoUpdate technology to keep the engine, resolution and vulnerability database current.
- **CyberCop Monitor**—CyberCop Monitor is a "Next-Generation" host-based Intrusion Detection tool that provides real-time packet analysis, system event misuse, and anomaly analysis to detect, and respond to intruders. CyberCop Monitor's unique architecture is compatible with high-speed and switched network environments providing a complete network monitoring solution across today's diverse network topologies.
- **CyberCop Sting**—CyberCop Sting provides a unique extension to traditional intrusion detection methods by creating a virtual network of decoy routers and servers on a "sacrificial" host. The Sting server is used to discover would-be hackers, and log attack efforts to help determine their origin, whether they originate from outside, or even inside the network environment. CyberCop Sting provides vital evidence collection to catch unauthorized users without putting production systems and data at risk.

The Network Associates CyberCop Intrusion Protection product line also includes:

- Central console software for reporting data, and remote agent configuration
- Auditing tests to assess the your intrusion systems
- Vulnerability database editor, allowing custom settings and resolution advice for each scan test



- Monthly updates from NAI Labs, the research division of Network Associates Inc. are automatically incorporated into the various engines using our AutoUpdate feature, including:
  - Latest advances in intrusion detection technology
  - Signatures of new attacks
  - Information about new software patches and other fixes.

## Conclusion

Attacks against network assets continue to rise dramatically. Sophisticated tools developed by experienced hackers are now distributed freely across the Internet. These tools allow complicated attacks to be staged by relatively inexperienced hobbyist hackers. At the same time, the e-business revolution demands that companies continue connecting mission-critical systems to the Internet.

As we have seen, first-generation network intrusion detection sensors are simply not living up to their initial promise. Although network IDS products are not the “silver-bullet” they were once considered to be, today’s next-generation multi-tier approach to intrusion protection can be a tremendously valuable complement to existing security systems such as firewalls, authentication, encryption, and virus protection.



Who's watching your network

For more information on products, services, and support,  
contact your authorized Network Associates sales representative

---

CORPORATE HEADQUARTERS

3965 Freedom Circle  
Santa Clara, CA 95054  
TEL: (408) 988-3832\*\*  
FAX: (408) 970-9727

*\*\*Call for additional Worldwide Sales Offices*

Visit our Website



<http://www.nai.com>

\* Network Associates, CyberCop, PGP and Total Network Security are registered trademarks of Network Associates, Inc. and/or its affiliated companies in the U.S. and/or other countries.  
All other registered and unregistered trademarks in this document are the sole property of their respective owners. All specifications may be changed without notice.

©1999 Networks Associates Technology, Inc. All rights reserved.

6-ACT-IDS-001 5/99