INTRUSION.COM

Security solutions for a .com world™

Maximizing the Value of Network Intrusion Detection

A corporate white paper from the product management group of Intrusion.com

Table of Contents

Primer 3
Different Types of Intrusion Detection Systems 3
Intrusion Detection Architecture 4
Focus on Network Intrusion Detection 5

Maximizing the Value of Network Intrusion Detection Systems 7

Policy First, Solutions Second 8

Total Time to Deployment 10

Tuned for Success 12

Security = Visibility + Control 13

Intrusion Detection isn’t Just for Security Anymore 15

Summary 16

Primer

Security Policies and Intrusion Detection Systems

Every effort to secure a network starts with policy decisions. These decisions may be simply a hallway discussion or scribbles on a restaurant napkin – or they can be a well thought out corporate security policy. Like all things, the more preparation that goes into the security policy, the more cost effectively and quickly it can be deployed.

Security policies address information assets or mission critical systems that need protecting. Starting with creating a perimeter to separate the private enterprise network from the Internet cloud, the security policy will generally specify deploying progressively sophisticated technologies to increase visibility and control over assets and systems to ensure the success of the enterprise.

Each technology will also have a security policy specific to its use. The security policy for intrusion detection systems (IDS) will define what information assets are to be protected, what type of IDS is needed, where it will be placed, what type of attacks the IDS will detect, and the type of response or alert which will be provided when a specific attack is identified.

Different types of Intrusion Detection Systems

Network Intrusion Detection Systems (NIDS) analyze network traffic for attacks that exploit the connections between computers and the data that can be accessed via a network connection. NIDS can detect the broadest range of attacks on corporate information assets which may include Denial of Service (DoS or DDoS for Distributed-DoS) attacks which are aimed at stopping the enterprise or its customer's from accessing corporate IT assets. The role of the network IDS is to flag and sometimes stop an attack before it gets to information assets or causes damage. NIDS are effective for monitoring both inbound and outbound network traffic.

> "Like a trained basset hound sniffing for drugs in luggage at the airport, network IDS products attempt to detect any malicious packets among the billions that travel the wires of your company's network every year."
>
> *"Security Watch", Infoworld.com Dec 2000 (Scambray/McClure)*

Host-Based Intrusion Detection Systems (HIDS) monitor specific files, logs and registry settings on a single PC and can alert on any access, modification, deletion and copying of the monitored object. The role of a HIDS is to flag any tampering with a specific PC and can automatically replace altered files when changed to ensure data integrity.

A derivation of HIDS is centralized-host-based intrusion detection (CHIDS) that serves the same purpose but does the analysis centrally by sending monitored files, logs and registry settings to the manager for analysis. The primary difference between these systems is as follows.

- CHIDS is more secure because it sends all the needed information off the host so that if the host is compromised, the alerting and forensic analysis can still take place. The tradeoff is that centralized analysis requires substantially more network bandwidth to move the data to the manager.

- HIDS makes policy compliance decisions locally and only sends alerts to the manager when warranted. This uses substantially less network bandwidth. The shortcoming of HIDS is that if the host is compromised there is no alert or forensic data to determine what happened or what was lost.

Hybrid Intrusion Detection Systems complement HIDS technology with the ability to monitor the network traffic coming in or out of a specific host. This is very different than NIDS technology that monitors all network traffic.

Intrusion Detection Architecture

All intrusion detection systems are based on a multi-tier architecture of a detection technology, a data analysis and configuration management layer and the user console or graphical user interface (GUI). When used by an individual on a single host, all three layers of the system may reside on the same PC. In enterprise or managed service deployments, each layer of the intrusion detection system is generally deployed separately to facilitate operations, ensure performance and support organizational workflow.

Detection technologies vary by the different types of intrusion detection systems.

- **Sensors** (sometimes called engines or probes) are deployable software or appliance-based technologies that allow network intrusion detection systems to monitor the mass of traffic on high-speed networks. Sensors are placed in specific locations at the network perimeter or within the network fabric. Sensors are processor-intensive devices and generally require their own PC or appliance to function correctly. The sensor analyzes all network traffic, looking for evidence of intrusion, and then reports the information to a centrally located manager following the parameters of the network IDS policy.

- **Agents** are deployable software installed on a particular PC in a host-based intrusion detection system. Agent software generally has a small footprint and uses very little processing power. The agent's function is to monitor specific files or logs on the host and reporting to a central manager if and when these particular files are accessed, modified, deleted or copied according to the host-based security policy. Agents are considered intelligent software as they determine policy compliance on the host and then only report breaks in the security policy.

- **Hybrid agents** combine the functionality of a host-based agent with network based sensor technology that is limited to analyzing only the network traffic addressed to the specific host where the hybrid agent is installed. A hybrid agent's footprint is generally larger because of the additional functionality. The processor utilization of the hybrid agent is much greater than a host-based agent because of the continual processing of network traffic for the host.

- **Collectors** are like agents in that they are lightweight software applications that reside on the host, similar to agents. The primary difference is that collectors are considered dumb devices because they do not make any decision at the host level. A collector's function is to harvest log, registry and file information from the host and to forward all of it to a central manager as soon as the entry occurs. The central manager does all analysis and decision-making for policy compliance. Most applications that use collectors are considered centralized, host-based intrusion detection systems.

The **manager** layer is responsible for accepting inputs from the deployed detection technologies and storing, analyzing and correlating the data for higher level intrusion detection. The manager is also the configuration and policy repository for the intrusion detection system. The manager uses some type of data and configuration store and applies ease-of-use, data mining and system management features to make the large amount of data provided by intrusion detection systems usable information for security policy enforcement and IT policy decision support. The manager is generally installed in a data center or server room with other server platforms that warrant physical protection like automated back-up, halon fire-protection and uninterruptible power supplies (UPS).

The operator of the intrusion detection system will interface with the system via a **console** (sometimes called GUI or UI) that is generally installed on a PC in the network operations center (NOC) or on the security professional's primary PC. The console's primary function is to make the monitoring and reporting of the system as intuitive and flexible as possible – thereby increasing the value of the information provided by the system.

Focus on Network Intrusion Detection

When considering intrusion detection technologies there are many existing models that can help determine where to start. Gartner Group's Total Cost of Ownership (TCO), Giga's Total Economic Impact (TEI) and Hurwitz's Return on Opportunity (ROO) are all valid models that basically say that the initial cost of the technology should be the smallest part of the consideration. In considering the total picture of what it will cost, the benefits you will receive and the cost of management and administration, network intrusion detection is the logical first step in intrusion detection deployment.

Network IDS …

- Provides the broadest coverage because it can be placed on the primary junctures within the network – so instead of placing a host-based or hybrid agent on each server in the DMZ, you can place a single network sensor within the DMZ.

- Is much easier to deploy as the sensor is a standalone PC and many NIDS are now available as appliances, further reducing the total time-to- and cost-of-deployment. Host and Hybrid IDSes require agent installation on all protected hosts, each installation requiring regular updating.

- Provides far greater detail into the nature of network traffic for decision support and policy definition. Host and Hybrid IDSes only provide information on the specific PC or server, and HIDS doesn't provide any information as to the traffic.

- The NIDS market is very mature and has already had a renaissance to develop some of the most sophisticated security technology in the industry. HIDS are still fairly simply technologies with the majority of improvement coming from solving the problems of deploying and managing agents. Hybrid technologies are still very new and offer little insight into deployment and management strategies.

- Many network intrusion detection systems can block an attack to stop a hacker from gaining access. These blocking tactics vary in efficiency but can respond to an attacker in real-time, thereby limiting the possible damage of an attack. Because HIDS are based on log file analysis, they are always after-the-fact. Hybrid systems can respond in real-time, but this protection is only offered to the individual host.

- NIDS can interact with other perimeter technologies to strengthen the enterprise perimeter. NIDS leverages the existing investment in routing and firewall technologies by dynamically updating other perimeter policies to respond to threats in real-time. Host-based systems are bound to single devices and cannot support existing investments in perimeter security technologies.

While there are places in a security policy for all types of intrusion detection, network intrusion detection offers a logical starting point. For the broadest impact on network security, the shortest-time-to-deployment and the most amount of network management information for the security dollar spent, network intrusion detection provides an accelerated path to showing the value and benefit, both security-wise and financially for a security investment.

Maximizing the Value of Network Intrusion Detection Systems

The beginning of making the most of an investment is starting with the right investment. In this environment of savvy marketing and thickly fielded markets, it can be difficult to analyze and compare performance metrics. The greater challenge could be making sure that the metrics being proposed are valid.

The Right Performance Metrics

The most popular metric for NIDS is number of packets dropped – or how many packets the NIDS was not able to analyze. But, that's not the reason a NIDS is purchased. If routers and firewalls, being inline devices, drop packets – they block traffic, and for these devices the metric is appropriate.

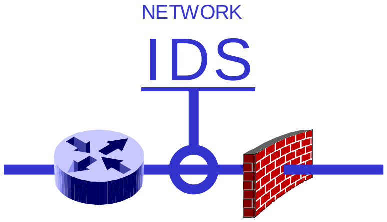

Network intrusion detection systems cannot be categorized with routers and firewalls because they are radically different in the following two ways.

- Network intrusion detection systems are passive devices and connect to the network with a "T" connection so it cannot block traffic.

- Network intrusion detection systems are not access control devices so dropped packets cannot result in blocked transmissions.

The key to using the right metric is being certain about what you are buying the network IDS to do. The goal of the NIDS is to correctly identify an attack REGARDLESS of the complexities of network saturation.

The correct metric for intrusion detection systems is "attack detection" at the varying levels of network saturation, for 10Mbs, 100Mbs and Gigabit segments.

The Value of Flexibility

In the mid-90's, many analyst houses were trying to find the right calculation to determine if one product was more cost effective than another over a period of time. The first of these metrics was Total Cost of Ownership (TCO) – which attempted to expose the cost of management, administration and various end user costs as a factor above the acquisition price. TCO was quickly followed by another model, Total Economic Impact (TEI). TEI went a step further to try and expose that the total cost of IT technology needed to be balanced against the benefit received from the technology and the flexibility of the technology, filtered by risk.

Especially in the security field, the benefits of flexibility are paramount. Networks are growing and changing at an amazing pace easily outstripping the ability of the IT department to hire and train personnel. Additionally, the pace at which new threats are being identified and exploited exacerbates the issues of network and information asset protection. To be effective, NIDS must accommodate the ever-changing environment of the enterprise network.

Flexibility in a network intrusion detection system can be measured in three main areas: customization, deployment and management.

- **Customization** allows the security professional to adapt the IDS policy to the uniqueness of the enterprise network. Customization increases the value of the data being produced by the IDS to provide highly valuable information about the usage of the network.

- **Deployment** flexibility allows for homogenous data to be collected from disparate network segments. Fast Ethernet, Gigabit segments, software deployments on existing hardware and appliance based solutions allow IT and security professionals to spend appropriately for their network needs.

- **Management** scalability makes it possible to leverage security investments. Network intelligence is enhanced with standardized information from myriad network segments in a single location or hierarchical architecture.

Long-term value recognition is highly dependent on the long-term viability of the solution. Forklift upgrades, total replacement of a system, are the bane of enterprise executives an indicator of unrealized value or a mistaken technology acquisition. Flexibility is the assurance that the technology you are looking at will fit the current enterprise as well as the enterprise of the future.

Policy First, Solutions Second

In the early days of network security it was common for someone in the IT or MIS department to buy a firewall and learn about network security from the firewall’s manual. This resulted in one of two possible outcomes – either the firewall was configured to be totally innocuous to ensure it wouldn’t cause trouble for the network, or the administrator used the firewall to its fullest potential only to realize the real need was something other than the firewall, in which they now have a substantial time and financial investment. The net result for either case: no value.

The companies who made successful investments in firewalls started with a strategy as to what they wanted the firewall to block – and then sought out the solution that provided the closest match to their requirements.

Security policies are goals the company has for their security investments and is the first step to gaining the maximum value of any security technology investment.

- Firewall's can end up neutered by an ineffectual policy, when not configured according to the set objectives of a security policy.

- An IDS will flood the administrator with so much data, much of it false-positives, it will become unusable.

Maximizing the value of network intrusion detection starts with a security policy that defines the metrics of success. The intrusion detection security policy needs to define the following three key points.

1. **What are you buying a network intrusion detection system to do?** This is a simple but very important starting point. Without a firm goal of what you want an IDS to do it will be impossible to determine if the deployment is successful – following the Taoist adage, if you don't know where you're going any road will lead you there. Responses to this question often fall into five main categories.
    a. Tell me who is attacking me – when placed at the perimeter of the enterprise, outside of the firewall, a NIDS can very effectively expose the nature of attacks launched toward the enterprise.

    b. How am I being attacked – if someone gained root access to your server, you need to know how they did it so you can stop it from occurring again. Capable NIDS can provide attack signature descriptions and many will also show keystroke-by-keystroke playback to help security professionals not fall victim to an attack twice.

    c. Tell me who on the inside is a threat – when placed within the network the NIDS will expose threats from authorized users, being benign employees, disgruntled employees or current or past contractors. NIDS sensors may be placed in mission critical areas like the DMZ or at the perimeter on the inside of the firewall.

    d. Help mitigate threat within specific network segments – NIDS sensors may be deployed within specific network segments which are more sensitive or previously identified as need special attention to contain risk. Circumstantial or suspected threat can be validated or mitigated by a subnet deployed NIDS.

    e. How do I get forensic evidence for prosecution – the information from a network intrusion detection system provides a unique view of the network with standardized data from multiple points across the enterprise for event correlation and activity tracking.

2. **What are you willing to miss with your NIDS deployment?** Whenever a sensor is deployed in the network there are specific threats being sought, and at the same time specific threats that are not. Reducing the number of items a NIDS is looking for will increase its performance. When deploying a sensor on the outside of the firewall, analysis of outbound traffic can be categorized for non-analysis or ignored all together. When placed within a specific subnet, targeting a NIDS sensor for certain types of threats and excluding other traffic reduces the amount of forensic and real time data an operator needs to analyze.

3. **How long will the specific network intrusion detection sensor be in place?** Like the military, deciding to use small, mobile mortars or large, high-power cannons, security professionals must determine what performance metric are important for each type of NIDS. There are some NIDS sensors that once deployed become a permanent part of the enterprise defenses. Subnet deployed sensors may be moved on a quarterly basis as needed. Determining ahead of time how long a sensor must be in place, the issues surrounding scalability, speed, form factor and deployability manifest themselves.

By making policy decisions around these three primary points, the total picture of network security becomes clearer with network intrusion detection filling a specific role in the execution of an overall enterprise security policy.

Total Time to Deployment

Nothing kills the perceived value of a project faster than not being able to quickly realize the benefits after a major investment.

One of the primary benefits of network intrusion detection is that it is deployed at a single point and protects complete network segments. Reducing the time and effort that goes into protecting that single network segment increases the perceived value of the IDS investment. To maximize the value of intrusion detection, a common strategy is to use appliances to reduce the total time to deployment – the time between acquisition and realization of benefits.

Appliances are task specific computers that vary in input/output (I/O) options and processing speed and memory. In general, appliances fall into three main categories: desktop, rack-mountable and chassis.

Desktop appliances are small, low profile devices built to be placed within the workplace, though they may also be placed on a shelf in a rack. Desktop appliances usually do not have I/O or keyboard-video-mouse

(KVM) capabilities as they are meant to have no user intervention and are to be remotely managed. Because of their size and remote management capabilities, desktop appliances are particularly useful for internal subnets. Desktop IDS appliances can be easily placed on a desk or hidden within a work environment and connected to the network segment that warrants additional security, monitoring or auditing. In a managed IDS service offering, the desktop appliance offers an easily deployable IDS sensor that has limited risk of onsite molestation, due to the lack of I/O and KVM.

Rack-mountable appliances are usually 1U or 2U high and generally offer I/O and KVM. Rack-mountable appliances are common to server rooms in enterprises and data-centers for managed service providers. These appliances are meant to be highly accessible, expandable and upgradeable. In contrast to desktop appliances that are really task specific, rack-mountable appliances are basically PCs and offer the same features of PC with the additional value add of the appliance vendor's management software and OS hardening.

Chassis appliances are many-Us high and allow the insertion of modular PCs, usually called "blades" or "cards," that are specifically built to share power and sometimes communications across the common back plane of the chassis. Chassis appliances are generally used for high-availability options and have features like hot swappable, redundant power supplies, cooling and memory.

The selection of an appliance is specific to the type of deployment planned. Desktop appliances are the newest form factor and provide the widest range of options for placement within the network. Desktop appliances are ideal for opening up remote and branch offices and departmental subnets to the security and network knowledge NIDS bring. Rack-mountable systems are the most common type of appliances. Maximum value in a rack-mount appliance is also maximum density. With rack space at a premium in most data centers and server rooms, there is no need to accept anything but a 1U high appliance to provide the level of throughput needed. Competent vendors are now offering both 100Mbps and Gigabit NIDS appliances in 1U rack-mountable form factors. Chassis-based appliances are the ultimate for large-scale deployments and data center manageability. Chassis' simplify clustering and site-specific or customer specific hardware and application allocation.

Vendor-provided appliances are generally covered by the vendors support programs and offer the shortest replacement times and highest levels of service coverage. In most cases, the appliance is treated as a black box, where opening the appliance voids the warranty. The plus side of this "hands-off" approach is that when the appliance is no longer responding, it is replaced avoiding a long and usually difficult repair cycle. The primary value in the appliance is that it comes ready to deploy, thereby providing the benefits of IDS in the shortest time possible, thus increasing the perceived value of the IDS.

Tuned for Success

Be it a Steinway, a Porsche or network intrusion detection system, without the ability to tune for specific goals dramatically limits the value and the success of the performance. When high-performance is the primary value driver, the ability to make the device work specifically to the needs of the environment, operator and application are germane. But this self-evident axiom has escaped many intrusion detection systems.

For a network intrusion detection system, there are two ways to detect intrusion: string matching and context analysis.

String matching, also known as network grepping in the Unix world, matches a series of characters to identify a threat. String matching offers tremendous speed in identification and additional signatures are easy to create and customize. The negative aspect of string matching is that it is prone to false-positives.

Context Analysis, also known as signature analysis, builds on string matching by allowing a scripted analysis of the context of the string thereby dramatically reducing false positives. The negative aspect of context analysis is speed; it simply takes longer to detect attacks in context.

Tuning allows the security professional to balance the needs of performance against the needs of accuracy. Additionally, without the ability to extend to signatures to properly identify unique network characteristics leads to false-positives and the resulting loss of productivity and reduction of information value. To maximize the value of the NIDS, the security professional must have the ability to extend and customize both string matching and context analysis signatures.

A higher level of performance tuning comes in global filters. Global filters make it possible at the sensor and console level to make decisions about what will be analyzed before a signature set or monitoring is applied.

- **At the console**: global filters come in the form of data sorting, where the information is still captured but not shown in the monitor. This facilitates monitoring and makes the operator's job much easier, as only relevant information is presented and can be made specific to operator preferences or vocational tasking.

- **At the sensor**: global filters allow exclusion of data by MAC or IP address, protocol or port. In addition to the inherent load balancing features of global filters at the sensor, this allows the enterprise to deploy a standard set of attack signatures to all sensors while other tuning characteristics are handled by sensor configuration – simplifying the task of signature maintenance.

Scenario

The enterprise is being recruited heavily by the competition; HR wants to know who is responding and sending out resumes.

String Matching Solution

A string-match is quick created to find the strings "resume", "qualifications" and "CV".

Alerts are now produced on each of the strings with a majority of false positives because it alerts when HR receives a resume, when HR forwards it to a manager, when HR sends out replies, when the manager sends it to an interview group – and when an employee send their resume out.

Context Analysis Solution

A script is written to find the strings listed above but only when they are outbound and not coming from HR.

Alerts are now showing only valid issues and creating the required value. It took longer to write the script but the information is now valuable.

A security policy is the first step to maximizing the value of network intrusion detection but the ability to effectively tune to system to the specific needs of the policy, the enterprise and the operator is paramount to assuring the value of IDS.

Security = Visibility + Control

Security is not a deliverable. Security products provide two primary benefits: visibility and control. And, it is the combination of these two benefits that make an enterprise's private computer network secure.

- **Visibility**: the ability to see and understand the nature of the network and the traffic on the network. Visibility is paramount to decision making.
- **Control**: the ability to affect network traffic including access to the network or parts thereof. Control is paramount to enforcement.

Maximizing the value of intrusion detection must also embrace IDSes role in the overall security fabric.

Routers were the first security devices in the network to establish a perimeter. The router is the IP equivalent of the early PBX on the telephone network – a small microprocessor that simply matches addresses to allow access. With little decision processing capability, the router requires highly educated operators and a lot of time to effectively enforce security policy. Beyond throughput and the access control list (ACL), the router is not effective for visibility and is totally dependent on the operator's ability to maintain the ACL to be effective for control.

Taking security to the next step, the firewall was created to add a layer of intelligence to the router. Firewalls brought a rules-based-policy to the ACL. A rules-based-policy allows the operator to create rules that automatically enforce the security policy using qualifiers such as IP address, time, protocol, and direction to specify actions like blocking, logging, alerting or allowing to pass. Firewalls added the needed intelligence to the ACL function of perimeter security.

Both the router and firewall are limited to being able to enforce an ACL. Since they are control-only devices, neither provides enough information to the security professionals to determine if the ACL or the rules-based-policy makes the enterprise more secure. This is like a security guard being able to tell if someone walking past him has a badge or not – but cannot tell if that person is supposed to have a badge or if the person is using their badge to walk out of the building with the company jewels.

Intrusion detection systems were created to add a new level of visibility into the nature and characteristics of the network. IDSes expose the packets within the data stream to identify threats from authorized users, back-door attacks and hackers who have thwarted the perimeter defenses.

Best-of-class IDSes fully decode protocols to completely expose packets and their contents. Using large databases of attack signatures tuned to the specific environment, the IDS then alerts security professionals to suspected threats. Network intrusion detection systems can also block specific data streams that are suspect. Unlike perimeter defenses that completely block users from access – IDSes are more discriminating, minimizing the risk of blocking authorized users who have made a benign mistake. IDSes can also be configured to "shun" the perimeter defenses – updating the ACL to respond to suspected source or destination threats. While possible, shunning is rarely used because of the denial of service potential created with an automated ACL adjustment.

Perhaps the greatest value network intrusion detection provides is the information about the use and usage of the network. This information can be used to:

- Increase the value and efficacy of the perimeter defenses,
- Produces hard evidence for the altering of the enterprise security policy, and
- Provides decision support for network management.

NIDS give security professionals an unsurpassed view into the traffic on their networks. The information from the intrusion detection system removes much of the guesswork from security professionals and feeds the routers, firewalls and other ACL-based security devices with the hard-core information that makes these devices more effective at establishing enterprise security. Combining the visibility of the IDS throughout the network with the control of perimeter defenses makes the enterprise more secure.

Security policies become self-evident with NIDS exposing the threats that exist within the network in real time and forensically. Rather than needing to hypothesize about what should be in the security policy and guess at the possible use and misuse of the network, network IDSes provide most of the data required to build reality based security policies and make changes to that policy based on promise, not paranoia.

The information from the NIDS goes beyond security. Maximizing the value of intrusion detection reaches throughout IT and MIS in helping all IT departments with a greater understanding of the use and usage of the network. While a great deal of security information is lost outside of the security department, intrusion detection systems can be used to help all departments do their jobs more effectively.

Intrusion Detection isn't Just for Security Anymore

Scenario

A small high-tech company was close to being acquired by a giant in their industry.

Though they had signed confidentiality documents, one of the executives sent an email to a few of his staff regarding the impending acquisition. Seeing that the email was from his boss and therefore reputable, one of the employees forwarded the email to some of his family members. The emails were not malicious, simply the exuberance of two people looking forward to their upcoming success and sharing it with people they thought they could trust.

A family member was at a party with an employee of the industry giant who overheard discussion of the acquisition.

To ensure no chance of being implicated in insider trading the industry giant promptly removed their bid for the small company.

Though there was no prosecution to follow, the small company was certainly at risk of legal action.

Solution

A network intrusion detection system can be easily tuned to alert senior executives when names, phrases or figures cross the network. This scenario could have been mitigated at the first incidence and the employees educated as to the risks they were opening up the company to by discussing this sensitive topic. Additionally, sophisticated IDSes could have terminated outbound traffic which had information regarding the sensitive topic and not originating from the President's office.

Maximizing the value of any security product means making the value apply to the business – not just the security department. While it sounds like hyperbole, the pat answer of "business enablement" won't cut it.

Network intrusion detection systems can provide information about the way employees are using the network … making it possible to address one of the most costly factors in business today: mitigating employee risk. Especially in today's employment environment, finding/training/retaining employees is one of the most challenging and costly aspects of business.

The cost factors for at-risk employees start with the lost productivity of the employee and others people the employee affects. The enterprise then suffers from the void the employee leaves. Recruiting new talent has tremendous opportunity cost for the hiring manager and future peers in addition to the cost of advertising, placement fees and relocation. Lastly is the cost of ramp-up time and training of the new employee.

NIDS can be used to identify at-risk employees and help to create an environment that supports employees and helps them avoid situations that could put them at risk. Harassment, divulging sensitive information and inappropriate use of corporate resources can all be stopped or addressed before they become firing offenses. Additionally, employees who are looking for another job can be engaged and potentially saved before there is a competing offer on the table or a resignation letter written.

One of the primary benefits of the visibility provided by a network intrusion detection system is seeing what is passing through the network. NIDS make it easy to find the information that can protect your human resources.

- Find outbound traffic that would indicate who is looking for a job
- Find inbound traffic from recruiters trying to locate candidates
- Find traffic about proprietary or confidential corporate information
- Find harassing or abusing traffic on the network
- Find misuse of corporate IT assets for humor, shopping, stock trading or social intercourse

By being able to identify certain types of traffic when they can be addressed and employees can be salvaged, a tremendous cost and corporate risk can be mitigated.

Human resource issues are ubiquitous among enterprises today. The business case for NIDS will be unique for each enterprise and can reach beyond the benefits of mitigating employee risk and attrition.

Summary

To maximize the value of network intrusion detection, a proper foundation needs to be laid including choosing the right metrics for product evaluation and success. Equally important is creating a policy that clearly identifies what the enterprise is trying to accomplish with intrusion detection so that the right devices may be deployed in the right scenario to achieve the goal.

Appliances will make the most difference in being able to quickly realize the benefits of an intrusion detection purchase – thereby increasing the perceived value and enhancing the enterprise's security. Long term, appliances reduce the total cost of system ownership maximizing the value of the initial investment over the life of the product.

The ability to tune the NIDS to the unique characteristics of the enterprise network is the number one element contributing to the reduction of false positives and the resulting increase in the value of the information provided by the network intrusion detection system.

Increasing the value of the network intrusion detection system is clearly on two fronts – making it more valuable to the security professional and making it more valuable to the enterprise. The security professional's value drivers are listed above with reduced total cost of ownership features of the system. Increasing the value of the system to the enterprise is where IDS projects go from "another IT burden" to a "business enabler."

- NIDS information provides decision support for security policy creation for ACL devices at the perimeter and within the network.

- NIDS provided network intelligence makes security policies fact based and helps remove the conspiracy theory stigma to which many security policies fall victim.

- NIDS network intelligence and network-usage intelligence can be used by other departments within the enterprise to address more quantifiable business enablement issues. Providing quantifiable data for network management to reduce the total cost of network ownership and providing whole-network monitoring of employee risk mitigation are front-line issues with most enterprises.

NIDS is very different than perimeter defense like firewalls, VPN or anti-virus gateways. To maximize the value of this uniquely powerful tool and the exceptional visibility it provides into the use and usage of the enterprise network, the security professional needs to promote and implement network intrusion detection as a network intelligence device, not just a security device.

To contact the author please email Ryon Packer, executive director of marketing for Intrusion.com at rpacker@intrusion.com.