



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at <http://www.giac.org/registration/gsec>

An Evolution in Security: Intrusion Prevention

GIAC Essentials Security Essentials Certificate (GSEC)

Practical Assignment
Version 1.4b

Dar-Ning Kung

Submitted October 2, 2003

Abstract

An enterprise, even one with business critical operations, by necessity must be connected to the Internet, accepting the risks along with the benefits. Security Administrators and Network Administrators face a number of challenges as they attempt to mitigate risks while maintaining the benefits conferred. Commonly deployed firewalls and routers with access control lists do not provide sufficient protection against increasingly sophisticated cyber attacks. Intrusion detection systems (IDSs) have become vital to detecting and alerting security administrators about these attacks, although IDS must still evolve to play a more proactive role in the actual defense of networks.

The evolution of IDS is towards Intrusion prevention system (IPS). The principal selling points of IPS are its proactive protection against attacks, a shorter cycle time to remediation, and claims of reduced personnel labor/cost that increase the return on investment. This paper presents a discussion about the following topics: What are the problems that drive the need for IDS and for IPS; how are host-based and network-based IDS different; how do host-based and network-based IPS differ; and how should IPS be implemented. This paper also examines if the claims made for IPS are valid and warrant an investment in the use of IPS within the enterprise.

What are the problems that drive the need for IDS and IPS?

Today, all organizations connected to the Internet face increasing risks and threats originating from cyberspace and from within the internal network [2]. Computer security professionals are working diligently to secure their enterprise against data loss and destruction. Increasingly sophisticated, rapid moving and automated attacks require a corresponding increase in a sophisticated, automated and rapid security response.

Information security departments must balance the need to protect the host operating environments, applications, and critical data against attacks with the openness required by business objectives. Hackers exist in a parallel universe and seek to take advantage of these same systems's connectedness and any vulnerability that can be exploited. Security professionals are under pressure to rapidly learn new IT security strategies, technologies and best practices to protect information and to achieve the objectives of the organization's mission. Meanwhile, hackers are also learning new strategies so as to create new attacks, generate updated worms and viruses, and use an evolving tool set to develop new malicious attacks.

To defend against the hackers who attempt to exploit systems, defense-in-depth provides a solid foundation upon which to defend the confidentiality, integrity and availability of the information and systems that are the crown jewels of daily operations [7, 8, 9]. Perimeter devices, such as routers and firewalls, are

important to protecting organization's perimeter; however, firewalls may not be able to combat new and clever attacks, which require the addition of security devices such as IDS or IPS. Traditional firewalls are limited to Layer 3 or Layer 4 inspection and cannot detect attacks at the application level that are contained within the packet payload. IDS can detect types of attacks that exploit Layer 2 (media access control) through Layer 7 (application) vulnerabilities [1], but detection is not enough. Preventative measures are required, which drives the evolution of network architectures to incorporate Intrusion Prevention Systems.

Host-based and Network-based IDS: Uses and Placement

Intrusion detection systems have two categories: network-based IDS (NIDS) is used to identify possible network intrusions and host-based IDS (HIDS), which is resident on a computer or a server. Network-based IDS is used to examine packets that traverse the network, whereas a host-based IDS is used to examine changes that may be made to a particular host.

While NIDS examines network packets, its primary goal is to determine if those packets contain dangerous payloads and more sophisticated NIDS also examine traffic for patterns that may indicate malicious behavior. In order to examine network traffic, a NIDS uses one or more of the following configurations:

- A passive tap, such as a Shomiti tap¹, where network traffic can be sent to the NIDS using a full duplex Ethernet network link without interrupting the traffic flow.
- Span mode², where the network traffic can be sent to NIDS via a switch's mirroring port. Span mode allows an IDS administrator to direct specific traffic to an IDS sensor by specifying a particular set of VLAN or switch ports.
- In-line mode³, which positions the IDS directly in the data path. In-line mode is required for an IPS implementation (which will be discussed in more detail below), because the network traffic actually traverses the IDS device; the IDS can block or drop malicious attacks.
- Multi-port mode⁴, where an IDS uses multiple sensors to examine network traffic at multiple network points throughout the network. This mode is useful when examining different network segments (e.g., before and after a firewall) and is very important for implementing event correlation analysis.

Malicious payloads are discovered by checking packets against pre-defined signatures, much the same way that an anti-virus product checks for known bad payloads. Using signature matching the IDS examines and validates packets against pre-defined illegitimate values specified in the vendor's signatures. A second packet checking mechanism uses protocol detection, which looks for protocol ambiguities, violations and atypical activity [11]. Another, anomaly detection, checking for patterns of malicious behavior, is harder to detect, but this

is usually done by baselining the normal network patterns during a learning period. When the IDS detects statistically significant changes to the traffic patterns previously defined as “normal,” an alert is triggered. Anomaly detection is particularly useful for detecting Denial of Service (DoS) attacks⁵. In a DoS attack the traffic behavior shows a sudden and significant spike compared to previously defined “acceptable” traffic behaviors and thresholds⁶. If a NIDS is to effectively detect the significant majority of known and new/unknown attacks it needs to provide multiple methods of detection for full attack coverage [10].

The typical placement of NIDS sensors is just behind the perimeter firewall⁷ and also just in front of servers that provide business critical functions. By placing the NIDS just behind the firewall it is possible to see if a malicious traffic has made it past the firewall. By placing the NIDS just in front of business critical servers any insider-generated malicious traffic may be detected. Some security engineers may also choose to install a NIDS sensor in front of the perimeter firewall to detect DoS attacks, however this sensor will receive a very large amount of noisy traffic, referred to as “doorknob rattling⁸,” that is stopped by the perimeter firewall. Sensor placement outside the perimeter does little to tell security administrators about the malicious traffic within your network.

Another tool in the IDS arsenal is the use of host-based IDS (HIDS). Besides being resident directly on host computers, a HIDS plays a different role from a NIDS⁹. A HIDS primarily inspects the host computer system's configuration files, detecting unauthorized changes to key files and settings that indicate changes or policy violations. For example, password files are monitored for unauthorized changes and key system areas are checked to detect permissions that may indicate policy violations. When an unauthorized change is detected (such to a registry setting or file permission), an alert is sent to a security administrator for further investigation. A more sophisticated HIDS can also analyze the activity of its host system in a finely granular manner and determine exactly which processes and users might be behaving in a way that signifies possible malicious intent. By its very nature, detecting system configurations as opposed to traffic traversing the network, a HIDS is able to detect successful attack attempts that are not suitably detected by a NIDS. A HIDS is also better able to handle encrypted information, which while encrypted during transit across the network, is decrypted once on the monitored target host¹⁰.

The typical placement of HIDS sensors is on those machines that provide for the crucial functions of an organization, its “crown jewels,” machines such as mail servers, database servers, and file servers. From a security standpoint, HIDSs are also employed to be the guardians of the guardians, placing sensors on firewalls, access control servers and log collection servers.

Challenges to the Utility of IDS

Since there is a broad range of attack types it is best to employ a broad range of detection strategies. Most importantly, the repertoire of attacks keeps changing, meaning that for an IDS to be effective it must be able to detect new attacks as well as existing attacks.

The typical way to make an IDS aware of new attacks is to provide it with new information about the attacks, and this is usually done through the IDS vendor providing new signatures. However, an IDS that is solely using signatures for detection is only able to detect attacks for which signatures have been created. It will not be able to detect new/unknown (“zero-day”) attacks for which there is no signature¹¹. Because of the constant stream of newly discovered vulnerabilities it is necessary for vendors to rapidly provide updated signatures for the IDS engine. For the highest utility, the dissemination process should be automatic and should not interrupt IDS operations. Vendors who do not provide updates quickly will force security administrators to write their own signatures, which can severely impact administrators who are already juggling heavy workloads.

Encryption, such as SSL and SSH, also increases the difficulties for IDS to accurately detect malicious attacks. Encrypted traffic is opaque to an IDS¹², unless the IDS is working in tandem with an SSL acceleration proxy device, which forces traffic to be decrypted at the SSL acceleration proxy in order to be read. This increases both the latency and the cost, and even then may be foiled if an encryption tunnel is created that does not use the proxy.

Detecting new attacks using traffic anomaly detection requires a relatively extensive time in learning mode, examining a large number of packets over time to set a baseline. Subsequent detection then requires that a significant number of packets be reassembled in sequence and examined for anomalies. Protocol anomaly detection also requires packet assembly. The latency this creates is often not suitable for high-bandwidth networks, particularly those using applications that are sensitive to delay. Packet reassembly and analysis has to occur at wire-speed, which is a difficult technical feat within high-speed (multi-megabit) networks, making processing speed a critical challenge with a large volume of network traffic¹³. If the NIDS doesn't have the processing speed for large traffic analysis, the packets may be dropped causing an attack to miss being detected. The problem of latency and processing power was a significant issue in first generation IDSs, which often consisted of software running on a PC or workstation. Later generations of devices have countered this threat by developing purpose-built appliances with task-specific operating systems and custom-built Application Specific Integrated Circuits (ASICs) to accelerate the packet processing. Additionally, using load-balancing with multiple NIDSs in an active/active configuration for parallel processing the network traffic also helps to avoid processing bottlenecks.

Detection accuracy is another major challenge. The common complaint about IDSs is that they generate too many false positives, mistakenly identifying benign

packets as attacks¹⁴. Security administrators must often validate the veracity of the attack, sorting through false positives, making for a labor intensive validation process. This requirement for human cross-checking both degrades the value of IDS implementation and delays the response, especially since going through alerts and logs often occurs after the event in question. If an IDS normally requires 24x7 monitoring the associated labor costs skyrocket and the effectiveness of the IDS diminishes the longer necessary countermeasures are delayed.

If an organization's management team understands the importance of implementing IDS and seeks to maximize its return on investment, NIDS may be the preferable choice because the solution covers the organization network instead of just one host or system. And if the requirements for organizational security (meaning the integrity, availability and confidentiality of its assets) justify the costs, an organization will choose a NIDS solution for its network and HIDS solutions for its mission critical systems. However, there are issues with IDS that must be addressed, one of the most important being the time it takes to respond to attacks. In an attempt to solve this problem the security market turns to the emerging art of intrusion prevention.

The Need for Intrusion Prevention Systems

If intrusion detection is the equivalent of radar, then intrusion prevention is the unmanned aerial drone directed by that same radar. A paramount requirement for effective use of either is accuracy. A timely response is important to good defense, but a fast and indiscriminant one will ultimately end up being more damaging than a slow but careful response. An automated intrusion prevention system must provide equal assurances that only malicious traffic is stopped and that only legitimate traffic is passed, or else the consequences to operations will end up being negative.

A properly configured and properly performing IDS detects malicious attack attempts and quickly informs the IDS administrator to take mitigating actions. Ideally the IDS should help the administrator to quickly analyze the nature of the attack and also provide some guidance action(s) to take to minimize any damage that might be caused by the attack. However, in most cases, the process is not that straightforward and more often attack prevention requires a high degree of coordination between the IDS administrator and network system administrator to work together to both validate the attack and to mitigate its effects. After coordinating among staff and doing an analysis, the outcome may be to change the firewall rule set to block the source IP of the attack attempt. However, for a medium to large enterprise network, an IDS administrator may not be able to effectively handle monitoring the alerts, validating them, and taking immediate remediation actions. The attacks may already have reached the targeted victim and caused damage well before the analysis is completed. Add to this the cost associated with a 7x24 IDS monitoring staff along with the need to have that staff

take actions quickly when a large number of attacks occur within a short period of time and security administrators end up with an expensive proposition. Even then, the monitoring staff may not be able to respond in a timely manner during the crisis.

Enter the IPS: In order to solve these problems it is necessary to consider a network device that can precisely block the malicious packets, either via firewall auto-blocking or by dropping the packets directly within the IPS device itself. An advanced generation of IDS products works with firewall APIs (such as Check Point's OPSEC¹⁵) to instruct the firewall to issue TCP resets or to add firewall rules to temporarily (or permanently) block attacks. And a new generation of purpose-built IPS appliances is being designed to perform attack filtration themselves, at a more granular level than most firewalls can. As intrusion prevention systems emerge from development labs they are filling a niche as natural solutions for the deficiencies of the IDS. IPS performs the functions of an IDS and adds more targeted firewall-like operations. IPS vendors claim it can help maintain business continuity through active prevention, and reduce the cycle time between attack detection and mitigation.

Evolving from intrusion detection to intrusion protection does have its benefits, though they may not be as substantial as the IPS vendors claim. IPS can provide additional network protection against malicious attacks; free up IT resources, including security administration; and reduce security management cost and financial losses caused by successful attacks. Adding an intrusion prevention system to an existing security network infrastructure provides additional protection to internal networks and strengthens the defense-in-depth arsenal. Effective network-based IPS (NIPS) helps to prevent attacks from reaching the target host and host-based IPS (HIPS) helps prevent the host from being compromised by an attack that may pass beyond the NIPS.

Host-based and Network-based IPS: Uses and Placement

Like the IDS market, the IPS market is divided into two classes of product: host-based intrusion prevention and network-based intrusion prevention. A host-based IPS is installed on a host to monitor and deter malicious activity on the host. And a network-based IPS is packaged as an appliance situated topologically behind the perimeter firewall, working to detect and prevent harmful inbound TCP/IP activity.

From an architectural standpoint the host-based and network-based IPSs are similarly structured; incorporating a sensor module, a rules engine, and a reporting module. While the skeletal architecture of the two IPS classes is similar, their operational paradigms are quite dissimilar. Host-based IPS, like its IDS counterpart, resides on the host server or workstation, but incorporates many features found in personal firewalls and even some features of anti-virus applications. HIPS examines incoming and outgoing traffic while also watching

for suspicious behavior on the host itself; blocking system hijack attempts, checking for activity representative of trojan horse applications, worms and other destructive threats. Behavior blocking identifies patterns of operations that appear to be consistent with the destructive goals of a virus, e.g., browsing all the entries in a directory followed by opening each discovered file for write access or perhaps deleting each file. The host-based IPS extends this paradigm; examining registry operations, examining access to objects, and stopping certain library function calls that may attempt to exploit buffer overflow vulnerabilities. Host-based IPSs, such as Network Associates' McAfee Entercpt [6] and Cisco Security Agent (formerly Okena StormWatch Agent) [3] are installed on a Microsoft *Windows* host—no UNIX flavors are supported—and monitor and deter malicious activity on the host.

Network-based IPS, again like its IDS counterpart, resides inline to perimeter traffic, just behind the perimeter firewall and perhaps also in front of servers that provide business critical functions. By placing the NIPS just behind the firewall it is possible to examine and block malicious traffic that has made it past the firewall, but in a method that is significantly more granular than the firewall. Often firewalls are limited to stopping an entire class of traffic from a source IP, whereas the NIPS can stop only that traffic that is known to contain a malicious payload or is engaging in known bad behavior. In practice, this would allow legitimate SQL queries to pass through while blocking SQL Slammer attempts from traversing the internal network¹⁶. Some security engineers may also choose to install a NIPS device in front of the perimeter firewall to protect against DoS attacks and reduce the load on the firewall, however this is an expensive proposition just for reducing the DoS loads on the firewall, which is a task most likely already being handled by load balancers residing in front of the firewall¹⁷. Network-based IPSs, such as Network Associates' McAfee IntruShield [5] and Netscreen's Netscreen-IDP [4] provide network intrusion detection and network intrusion protection.

Challenges to the Utility of IPS

As stated previously, the use of signatures for detection has its drawbacks and NIPS does have the same reliance on signatures. The workarounds are the same: vendors must provide signatures in as timely a manner as possible. Vendors may improve this process by developing a method to collect a nonymous event information in aggregate from its clients, allowing it to be alerted to the fact that an unknown packet type is being seen across its customer base, alerting it on a global scale that there may be a problem worthy of investigating.

Encryption likewise poses a challenge to NIPS. Not only can an attacker bypass the signature checking via the use of polymorphic viral code, but by trivially encrypting the payload beforehand and incorporating a mini-engine that decrypts the payload before branching, the IPS may not be able to examine the payload. There are emerging standards for cryptographically signed code, which could

help the NIPS to distinguish between acceptable and non-acceptable code, but this is unpopular both because of the overhead it introduces to system development and because of its negative impact on open and interoperable systems.

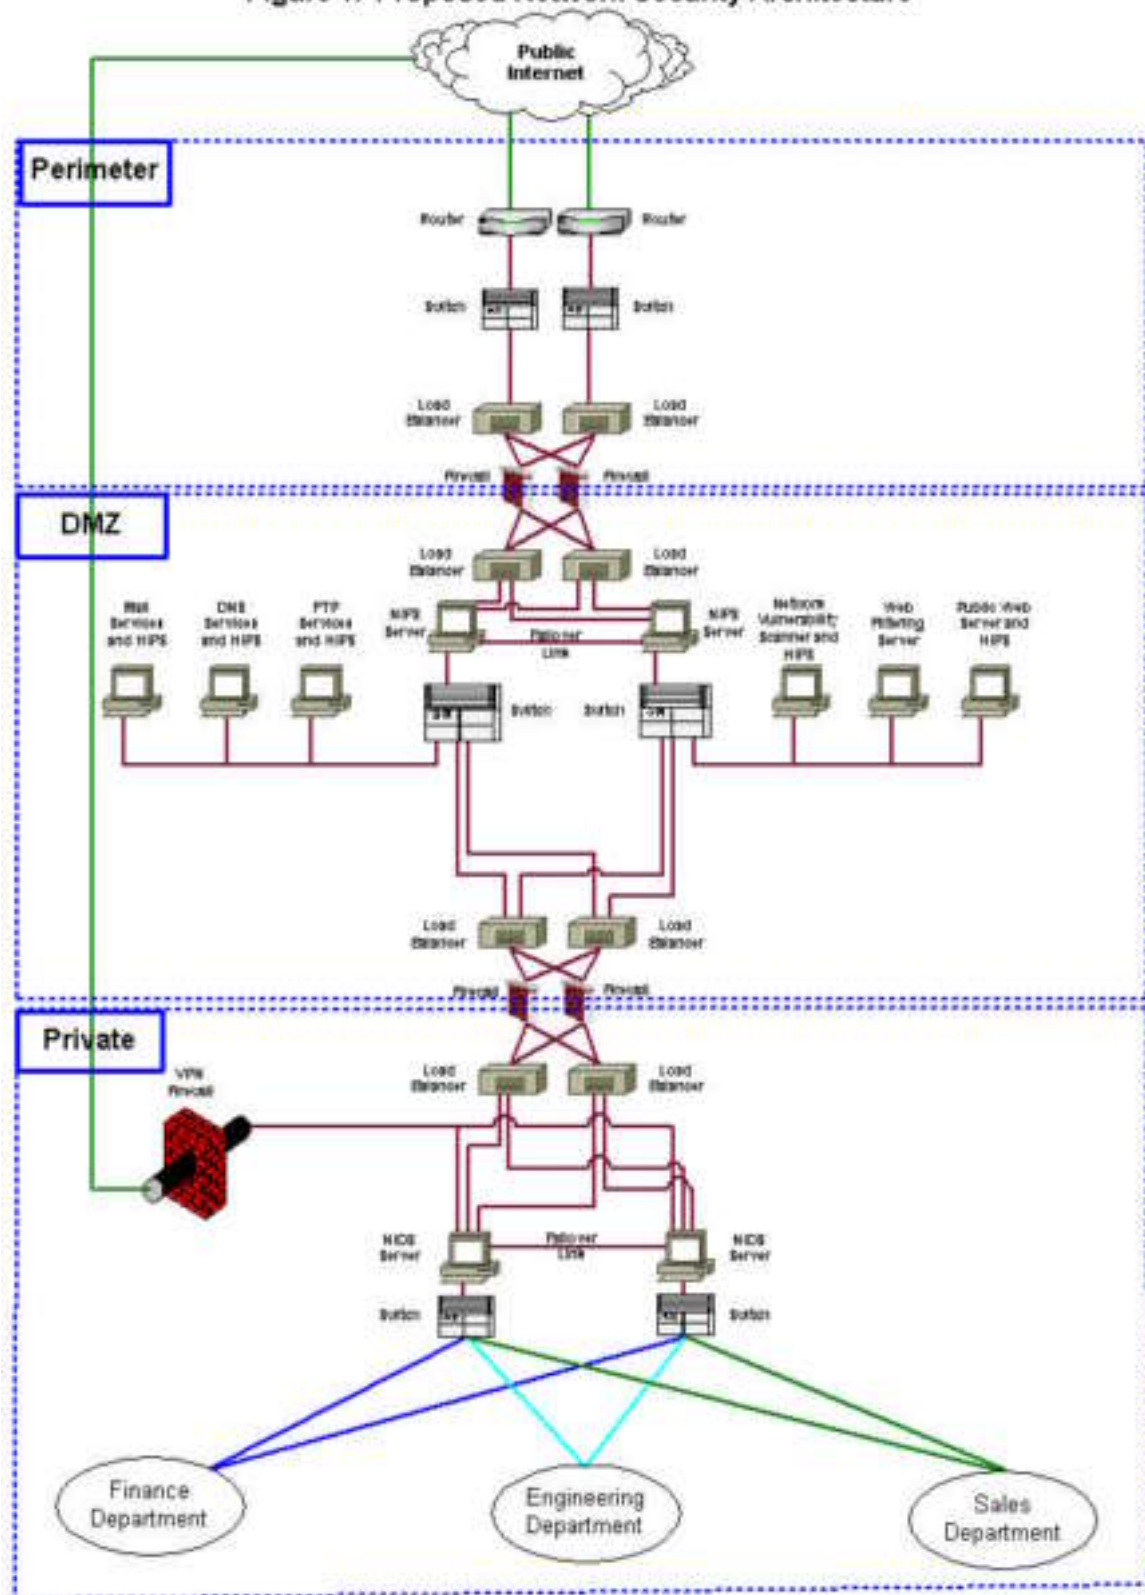
Anomaly detection relies on having some definition of allowed behavior and then noting when observed behaviors differ. Anomaly detection systems monitor networks for two primary criteria: characteristic deviation and statistical deviation. Characteristic deviations tend to be more qualitative. For example, “this host does not normally transfer files outside of the company.” Statistical deviations, on the other hand, tend to be more quantitative. For example, “This site’s ICMP traffic never exceeds 10% of capacity.” A simplistic implementation of protocol anomaly detection may only look for a small number of known problematic conditions, such as overlong buffers, while a more in-depth implementation may evaluate all data for compliance. The tradeoff will be in speed because the more detailed an implementation, the more comparisons it must perform at each stage. NIPS devices are just starting to emerge that can work at multi-gigabit speeds, but these vendor claims have not been substantiated through testing with the more in-depth evaluation features activated and working on a realistic traffic mix.

Host-base IPSs have challenges of their own: Because they augment the discretionary access controls offered by the native OS by implementing a limited form of type enforcement, (e.g., ensuring that the authorized Microsoft SQL Server accesses DBMS-related files and folders only, while the Microsoft IIS engine confines its accesses to hypertext documents and their bedfellows) they require a good deal of semantic understanding as to how an application operates. Consequently, they are not readily extensible to new classes of applications. And since the host-based IPS is shouldering work that should be the responsibility of the application, it also incurs the performance penalty associated with checking string lengths.

Implementing IPS in an Existing Network

Deploying IPS has its own set of challenges, over and above those challenges faced with deploying IDS. A sophisticated network security architecture must be tailored to the needs of the specific environment, though the most common is the 3-tier architecture: (a perimeter zone, a demilitarized zone (DMZ), and the private (internal) network zone). Figure 1 shows a typical three-tiered network security infrastructure. The layered approach has the following advantages: access to external services at the DMZ does not impact the security of the private network, which is protected by a firewall; public services are isolated in the DMZ; and the internal network has two layers of firewall protection. In addition, the externally accessible servers in the DMZ prevent a compromised server from analyzing the traffic to/from the internal network, limiting the potential for damage to the internal network.

Figure 1. Proposed Network Security Architecture



Here, the router determines the next network point to which a packet should be forwarded toward its destination, it also serves as a first line of defense by dropping packets that violate the access control lists. The firewalls work closely with the routers, stopping packets that violate the firewall rule set.

Behind the firewall is the network-based IPS. This device should not be placed at such a critical juncture within the network without some significant thought given to how it is deployed, which is in four distinct phases. Network Associates IntruShield products can perform both NIDS (in tap or span mode) and NIPS functions (using in-line mode). The vendor suggests a progressive IPS deployment starting with (1) detection but no prevention (in-line mode as an IDS); (2) in-line detection with no prevention (in-line mode); (3) detection and selective prevention (in-line mode); (4) and finally detection with broad prevention (in-line mode)¹⁸. Security administrators can develop confidence in the device as well as adjust policies and configuration parameters at each phase to avoid potential missteps caused by IDS/IPS implementation throughout the whole process. When installing an IPS, it is necessary for security administrators to understand the pitfalls and consequences, because a solution that inadvertently blocks legitimate traffic will become instantly unpopular.

Using this tiered architecture, a company might have avoided the malicious RPC attacks that occurred late in the summer of 2003. It is still vitally necessary to patch systems, but IPS can detect and filter in real-time the payloads containing the Blaster and Nachi worms while allowing critical traffic to use the RPC ports. Since these worms were particularly fast moving, there was little chance to protect unpatched hosts (unless the appropriate protections had been put into place prior to the worm being unleashed). Securing the perimeter is crucial to protecting the overall organization's IT assets. A firewall provides a false sense of security without IDS and/or IPS because many attacks otherwise slip through the routers and firewalls, since routers and firewalls do not detect malicious payloads. Moreover, firewalls typically filter based on the layer two and three (source IP), but not based on what is in the payload at layer seven.

The IPS augments the firewall by being able to block or drop an attack packet, all subsequent packets for the session, or initiate a TCP reset within the IPS itself when in an in-line mode. IPS may also reconfigure a firewall rule set to block offending traffic using a common firewall API such as OPSEC¹⁹. Working together in tandem with the firewall, and with each other, NIPS and HIPS can detect and protect against suspicious activity that occurs at the perimeter of the network and on business critical hosts.

While worms get most of the attention, IPS can also help thwart more targeted attacks as well. The fundamental axiom of network-based IPS is that every network-based attack begins with a reconnaissance phase that performs a TCP and UDP port scan and/or executes canned checks for widely publicized

vulnerabilities. First, assuming that an attacker has the advantage of operating at leisure, he/she may attempt to disguise a port scan amid legitimate network traffic by not only jumbling the port sequence or by scanning discrete subsequences, but also transacting authorized operations (e.g., downloads from a public Web server) between sporadic port scans. The NIPS selection criteria should include the ability to detect these out of sequence or “low and slow” scans. If it does not, the security design should consider augmenting the NIPS with an event correlation engine, which can bring together disparate events and cast subtle patterns in a stark light. Second, the majority of damage perpetrated by an attack is localized to the hosts where the vulnerable files reside. Malicious intent cannot necessarily be discovered through the analysis of the network-layer and transport-layer traffic while remaining ignorant of application-layer directives. So consider HIPS for those important hosts.

Part of the implementation IPS plan is to determine how the NIPS interoperates with your current environment. Consider if the policy that the IPS offers differs from the policy that the firewall administrator views and uses on the GUI console. If so, this is a cause for concern. These policies must be aligned. Aligning policies is a greater challenge when the organizational security policy is enforced not by a single firewall, but by a constellation of multiple firewalls (possibly from different vendors) that cooperate to form a logical unit. Test if there a sufficient method for the firewall and the IPS to communicate with each other and determine what rules are added or removed over time. This will help with the problem of divergent policies over time. Since policies are not static there is the risk that after a year or two of operation both devices may be enforcing divergent or conflicting policies. Ideally, there should be a common method for viewing, changing or managing the policies across multiple machines and multiple classes of security devices.

The problem that dogs IPS the most (and may continue to do so) is the issue of Crossover Error Rate (CER); that is, the point where failures to detect malicious activity (“misses” or in the lingo: False Rejection Rate Type I Errors) are balanced with false alarms (False Acceptance Rate, or a Type II Errors)²⁰. If an IPS mishandles legitimate packets by dropping them, it inadvertently participates in a successful DoS attacks caused by the IPS. There is no easy solution, but selecting an IPS that employs multiple types of detection can lower but not reduce the CER. By lab testing the prospective IPS device in your own lab with your actual traffic, you will have a chance to tweak your policy and improve the accuracy of your results. Caveat Emptor: the reported accuracy for IDS and IPS devices touted by vendors is as useful as sticking your thumb into the air to determine wind-speed. Vendor test cases may be carefully designed and probably do not cover all the aspects of the network traffic patterns of your particular organization. It is questionable, at best, to claim near - 100% accuracy considering the dynamic behavior of unknown attacks and newly uncovered host vulnerabilities.

The Hard and the Soft Benefits of IPS

The IPS is not a panacea and cannot provide a full range of comprehensive coverage and 100% reliable security device. However, IPS can reduce the load for security administrators to examine all alerts from IDS while it provides necessary blocking 'indisputable real attacks' in real time and an automatic fashion that already is a cost-saving factor against any potential incident response performed by administrators. In addition, both IDS and IPS can provide detection and protection 24x7 since most organizations may not be able to afford skilled personnel around clock for monitoring network traffic and analyzing security log files. Furthermore, statistical reports from IDS can provide a powerful and convincing result for security staff to demonstrate to top managers the needs for better IT security, and request for adequate security funding and resources to secure the organization business operations and to safeguard IT assets. IPS cannot be (and may never be) used for completely replacing security administrators.

After understanding the limitations of IDSs and IPSs, security administrators require careful planning, preparation, testing, evaluating network traffic patterns, and progressively deployment them. Because of the complexity of IPS, any rush implementation may create considerably negative impact on legitimate network traffic. As most security professionals agree, security is not "out the box," and this is particularly true for IDS and IPS deployment. The devices have to learn normal traffic patterns (inbound and outbound) as well as behaviors of the organization's applications and available services. Security administrators also need to provide appropriate security policies and procedures to react to various alerts based on available resources.

Finally, security event management and correlation tools are useful for security administrators to analyze security alerts and incidents by examining and correlating information from various security devices on the network or host log files against known vulnerabilities in real-time. The iterative interaction with the tools can help security administrators reduce false positives and explore false negatives using the feedback provided by real-time event correlation.

Conclusion

IPS has been advertised as the "next big thing" to meet the demanding and imperative needs for security administrators – switching from reactive security devices such as IDS to a proactive protection approach. However, like many new technologies, a high degree of skepticism arises regarding vendor statements regarding accuracy and performance that sound too good to be true. The decision for choosing a specific product should make use of testing with a real-world traffic mix that parallels your organization's usage and load. Additionally, make use of independent product comparisons and the experiences of those brave "bleeding edge" organizations that have experimented with the

technology you are considering. Most importantly, a progressive IPS implementation based on the confidence level of security administrators is usually the most reasonable approach.

With early detections, administrators can respond accordingly and mitigate risks caused by in-progress attacks. With proper configuring and tuning of the IPS to reduce false positives, the mechanism can avoid attack-related costs of data loss and service interruptions, loss of productivity, negative business impacts, and financial losses caused by successful attacks, making for a good return on investment. The IDS is akin to a radar system that can spot incoming attacks, while an IPS is like an anti-missile-missile system that uses that same radar to stop incoming threats before they reach their intended targets.

It is crucial for security administrators to recognize the limits of IPS before implementation because IPS may not be able to precisely analyze and respond properly against all known and/or unknown first-time malicious attacks. It is unrealistic to expect 100 percent accuracy for IDS or IPS. There will always be the need for security administrators to monitor and analyze the potential false positives and false negatives, and continuously tune the policies or update signatures used by IDS or IPS. Another important fact is that IPS cannot be used as a single mechanism against every malicious attack. Without recognizing the limits of IDS and IPS, they may provide a false sense of security and lead to inadequate readiness for the dangerous attacks with unknown scope and complexity that wait just over the horizon.

© SANS Institute 2004,

References

- [1] Bace, R., and Mell, P. Intrusion Detection System, NIST Special Publication SP 800-31, National Institute of Standards and Technology, November 2001.
- [2] CERT Coordination Center, "CERT/CC Statistics 1988-2003", URL: http://www.cert.org/stats/cert_stats.html, (October 2, 2003).
- [3] Cisco. "Cisco Security Agent - Introduction". URL: <http://www.cisco.com/en/US/products/sw/secursw/ps5057/index.html>, (October 2, 2003).
- [4] Netscreen. "McAfee System Protection, Host Intrusion Prevention". URL: <http://www.netscreen.com/products/idp/index.jsp>, (October 2, 2003).
- [5] Network Associates. "McAfee Network Protection, Network IDS Sensors". URL: http://www.nai.com/us/products/sniffer/network_intrusion_prevention/network_ids_sensors/category.htm, (October 2, 2003).
- [6] Network Associates. "McAfee System Protection, Host Intrusion Prevention". URL: http://www.networkassociates.com/us/products/mcafee/host_intrusion_prevention/category.htm, (October 2, 2003).
- [7] SANS Institute, Course Material: "SANS Security Leadership: Part 2, Defense in Depth – In Depth", 2002.
- [8] SANS Institute, Course Material: "SANS Security Essentials with CISSP CBK", Version 2.1, 2003.
- [9] SANS Institute, Reading Room Documents, URL: <http://rr.sans.org/index.php>, (October 2, 2003).
- [10] Skoudis, Edward, "Sneaking Past IDS", URL: <http://infosecurymag.techtarget.com/2002/jul/sneaking.shtml>. (October 2, 2003).
- [11] Tanase, Matt, "The Great IDS Debate: Signature Analysis Versus Protocol Analysis", URL: <http://www.securityfocus.com/infocus/1663>, (October 2, 2003).

Citations

- 1 "Century Tap Family." <http://www.shomiti.net/shomiti/century-tap.html>
- 2 "Understanding Spanning-Tree Protocol."
http://www.cisco.com/univercd/cc/td/doc/product/rtrmgmt/sw_ntman/cwsmain/cwsi2/cwsiug2/vlan2/stpapp.htm
- 3 "Snort Inline Mode." http://winfingerprint.sourceforge.net/presentations/honeynet_project-mexico2003.ppt
- 4 Barisani, Andrea. "Multiple Snort Sensors HOWTO." <http://www.infis.univ.trieste.it/~lcars/ids/> (16 Oct 2002)
- 5 "IntruShield sensors tightly integrate signature and anomaly detection techniques and DoS detection."
http://www.networkassociates.com/us/products/sniffer/network_intrusion_prevention/network_ids_sensors/4000.htm
- 6 Ptacek, Thomas. "Insertion, Evasion, and Denial of Service: Eluding Network Intrusion Detection."
<http://secinf.net/info/ids/idspaper/idspaper.html> (16 Oct 2002)
- 7 Cisco Secure IDS Sensor Deployment, http://networking.earthweb.com/netsecur/article.php/10952_981221_2
- 8 Doorknob rattling noun: Probing a computer that is connected to the Internet to see if it has any vulnerabilities that can be exploited. Search Word Spy with "doorknob rattling", <http://www.wordspy.com/words/doorknobrattling.asp> (6 Nov 2000)
- 9 Zirkle, Laurie. "What is host-based intrusion detection?" http://www.sans.org/resources/idfaq/host_based.php
- 10 "Network- vs. Host-based Intrusion Detection." Internet Security Systems
http://documents.iss.net/whitepapers/nvh_ids.pdf (2 Oct 1998)
- 11 Yee, Andre. "The Intelligent IDS: Next Generation Network Intrusion Management Revealed." http://www.forum-intrusion.com/The_Intelligent_IDS.pdf (July 2003)
- 12 Marshall, Geoff. "If a firewall is your first line of defense then an IDS should be your second, "NIDS cannot examine encrypted traffic." http://www.scmagazine.com/scmagazine/2003_04/test_02/ (April 2003)
- 13 "Capacity Verification for High Speed Network IDS."
http://www.cisco.com/en/US/products/hw/vpndevc/ps4077/prod_technical_reference09186a0080124525.html
- 14 Patton, Samuel; Yuric, William; Doss, David. "An Achilles' Heel in Signature-Based IDS: Squealing False Positives in SNORT." http://www.raid-symposium.org/raid2001/papers/patton_yurcik_doss_raid2001.pdf
See also: "Intrusion detection systems: Reducing network security risk."
<http://www.zdnetindia.com/print.html?iElementId=79197> (3 April 2003)
- 15 OPSEC (Open Platform for Security) is the industry's open, multi-vendor security framework. With over 350 partners, OPSEC guarantees customers the broadest choice of best-of-breed integrated applications and deployment platforms.
<http://www.opsec.com/>

16 Parker, Laura. "TIPPINGPOINT TECHNOLOGIES IMMEDIATELY INOCULATES CUSTOMERS AGAINST SQL SLAMMER ATTACKS." http://www.tippingpoint.com/news_events/pdf/SQLSlammer_012703.pdf (27 January 2003)

17 Yerxa, Gregory. "Firewall & Load-Balancer: Perfect Union?" <http://www.networkcomputing.com/1102/1102ws1.html> (7 February 2000)

18 Network Associates, McAfee IntruShield White Paper "Intrusion Prevention: Myths, Challenges" (April 2003)
http://www.nai.com/us/products/sniffer/product_lit.htm

19 INTRUSION DETECTION SYSTEM SOLUTION "IDS can drop the packet, terminate the session, reconfigure access control lists (ACLs) on routers and switches, or dynamically modify the firewall policy to "shun" the intruder."
http://www.cisco.com/en/US/netsol/ns110/ns170/ns171/ns292/networking_solutions_customer_profile09186a008017f0ac.html

20 IDS: Crossover Error Rate (WAS "Intrusion Prevention")" <http://lists.insecure.org/lists/focus-ids/2002/Dec/0028.html>
see also Krause, Micki & Tipton, Harold. CHARACTERISTICS OF BIOMETRIC SYSTEMS "False Reject Rate ... False Accept Rate." Handbook of Information Security Management: Access Control.
<http://www.cccure.org/Documents/HISM/039-041.html>