THE TOLLY GROUP

The authoritative, unbiased source for IT certification, research and testing

WHITE PAPER

May 2006

A white paper commissioned by iPolicy, Karalon, Nortel, Sourcefire and Top Layer

Document #206115

TollyEdge

Benchmarking Strategies for Wired Intrusion Prevention Systems (IPS)

Terms of Usage

USE THIS DOCUMENT ONLY IF YOU AGREE TO THE TERMS LISTED HEREIN.

This document is provided, free-of-charge, to help you understand whether a given product, technology or service merits additional investigation for your particular needs. Any decision to purchase must be based on your own assessment of suitability. This evaluation was focused on illustrating specific features and/or performance of the product(s) and was conducted under controlled, laboratory conditions and certain tests may have been tailored to reflect performance under ideal conditions; performance may vary under real-world conditions. Users should run tests based on their own real-world scenarios to validate performance for their own networks. Commercially reasonable efforts were made to ensure the accuracy of the data contained herein but errors can occur.

The test/audit documented herein may also rely on various test tools the accuracy of which is beyond our control. Furthermore, the document relies on certain representations by the sponsor that are beyond our control to verify. Among these is that the software/hardware tested is production or production track and is, or will be, available in equivalent or better form to commercial customers. The Tolly Group provides a fee-based service to assist users in understanding the applicability of a given test scenario to their specific needs. Contact us for information. When foreign translations exist, the English document is considered authoritative. To assure accuracy, only use documents downloaded directly from The Tolly Group's Web site.

Tolly Group Vendor Service

With more than 17 years experience validating leading-edge Information Technology products and services; The Tolly Group has built a global reputation for producing accurate and unbiased evaluations and analysis. We employ time-proven test methodologies and fair testing principles to benchmark products and services with the highest degree of accuracy.

Our "Up-to-Spec" service provides the custom testing complement to the "standard", granular tests offered in "Tolly Verified". See our "Up-to-Spec" Home Page.

Plus, unlike narrowly focused testing labs, The Tolly Group combines its vast technology knowledge with focused marketing services to help clients better position product benchmarks for maximum exposure.

This document was authored by:

- Charles Bruno,
  Executive Editor
  The Tolly Group

Table of Contents

Benchmarking Strategies for IPS

Overview

With e-Commerce and Internet usage growing at exponential rates each year, data security — defending computer networks against threats and vulnerabilities — has emerged as a top concern for IT professionals and corporate decision makers. Leaving networks exposed results in devastating consequences to the unprotected entity, including:

- The inability to conduct business transactions,
- reduced consumer confidence,
- reduced employee productivity,
- unplanned expenditures,
- unplanned IT resource time allocated to security-related clean up,
- legal liability for organizations and individuals within the organizations,
- theft of proprietary information,
- diminished company reputation, and, ultimately,
- loss of revenue and profits.

Industries that are most vulnerable to attacks include: Banking and Finance, Water Supply Systems, Electric Power Systems, Gas and Oil, Information and Communications, and Emergency Services, to name a few. Legislation has been enacted in the last few years to protect the confidentiality and integrity of private data, holding corporate executives personally accountable if they do not comply with these laws by a mandated date. Therefore, it is extremely important to have a corporate structure that supports network security from the top down.

The latest tools in corporate IT defense strategies are advanced Intrusion Prevention Systems (IPS) that integrate a deep packet inspection-based stateful firewall, intrusion detection and prevention system (IDS/IPS), and content screening technologies, capable not only of identifying unwanted intrusions, but also blocking and preventing them.

The IPS acts much like a security guard at a bank, performing as a deterrent to potential hackers and keeping customers and suppliers confident in their abilities to protect information and investments. What makes IPSs even more effective is their ability to protect all seven layers of the OSI stack, while allowing normal traffic to continue to flow...maintaining normal business operations while thwarting attacks. IPSs with deep packet inspection, stateful firewall and content screening capabilities are critical nowadays to protecting corporate, government and organization networks.

While there are many important considerations with respect to selecting an appropriate IPS solution, it is difficult to argue against "protection" as being the most important element. After all, if an intrusion solution cannot offer the level of protection users need, nothing else matters. So, an IPS system must provide a full array of protection against known and unknown attacks, viruses and worms; offer high accuracy rates with little to no "false negatives" or "false positives" (which we'll explain below); introduce minimal degradation of network performance; ensure low latency rates while providing standard performance; be easy to install and maintain with good reporting tools; and ongoing vendor support with automatic updates/vaccines to the system.

This ***TollyEdge*****: Benchmarking Strategies for Intrusion Prevention Systems** white paper explains what to look for in an IPS — in terms of protection, performance, ease of use and reliability, quantifying measures and benchmarks where possible, to help you decide how much software or other tools are required to adequately protect your network. The IPS market is projected to grow exponentially over the next few years. Don't be left unaware of the risks of being left unprotected.

Designs of the Times

IPS systems, traditionally, utilize a scanning engine to inspect incoming data against a database of virus and malware "signatures" to properly identify and thwart threats before they become security breaches.

As packets pass through the traditional IPS, they are fully inspected to determine whether they are legitimate or malicious. Each specialized security appliance or each security application if executed on an individual hardware platform opens the packet, inspects it, analyzes it and makes a decision to let the data through or discard it.

Such products predicate their packet inspections on an extensive database of exploit and malware signatures. Some vendors offer products with comprehensive signature databases that are updated regularly to reflect the current attacks and threats. Some products inspect at Layers 3/4, others extend up to Layer 7 to protect application level threats.

Yet another advanced IPS design goes beyond this deep packet analysis. Several of the emerging threats require multiple security defense mechanisms in addition to just the single function intrusion prevention. Some vendors perform multiple security processes based on a one-time Layer 2 to Layer 7 deep inspection of the packets. This then is followed by a shared multifunction security analysis and the result could be a correlated multilevel security response — e.g. across intrusion prevention, content filtering and stateful firewall.

The benefit of such an approach, proponents say, is a higher level of security and ability to thwart both current and emerging threats while maintaining high performance and lower latency. Low latency is important in streaming applications such as Voice over IP (VoIP). In this approach, for example, the security system can protect against a worm attack by blocking the worm through signature-based intrusion detection and prevention, block the worm communication through a firewall action and disallow the worm-infected machines to communicate with the command site through URL filtering, all with a single packet inspection application thereby not degrading performance or latency.

Increasingly, there is another new IPS design that is popping up that claims to be more efficient than signature-based IPS solutions. So-called behavioral-based IPS solutions do not rely upon virus signatures to analyze incoming traffic. Instead, they rely upon an extensive rules base to examine known vulnerabilities of a system and look for and respond to network activity that tries to take advantage of network vulnerabilities. The essence of such designs is that rules-based systems deal with the root of the security vulnerability rather than just the exploitation.

The behavior-based systems often focus on three areas: real-time protection, rules-based behavior scanning, and real-time discovery of network resources and the conversations between those devices to identify and deal with threats targeting devices and applications.

Behavior-based solutions use a rules-based decision engine that can be configured to detect both signature-based events for known exploits and anomalous behavior for yet unknown threats. Rules are used to examine packets at both the IP protocol level and at the application level and can be set to look for specific occurrences of attacks against a protocol or set to look for the conditions of an attack.

The idea with behavior-based systems is that users can adjust rules to fine tune deep-packet inspection; specify source and destination addresses, and offset the inspection point to improve performance and catch threats that are hidden deep within the data payload. Such a system also makes it easy for security administrators to modify or disable rules that are triggering threat events on legitimate traffic.

Another key distinction between traditional IPS solutions and behavior-based solutions is the concept of network discovery. Traditional IPS solutions, by their nature, focus on identifying and stopping incoming malware threats by comparing income data streams to signatures. Such IPSs are unaware of the underlying network and the devices they protect.

Behavior-based IPS solutions, by contrast, build a map of the network and the devices on the network. They track the nature of the conversations those devices are having by employing endpoint intelligence to produce an impact flag, or a signal that an attack could be underway. IPS rules then can be fine-tuned to respond to various network conditions either automatically, without human intervention, or fire off a message to a network administrator with notification of the event.

The degree of risk assigned to the various rules in the product is a fundamental differentiator when compared to traditional signature-based IPS tools. The behavior-based IPS solutions are more effective to detect and block zero-day attacks than signature-based alternatives. In other words, the behavior-based IPS solutions more dynamically cope with new types of attacks than signature-based tools. The pure signature-based IPS solutions cannot detect and block without the appropriate signature and there is high possibility they do not have an appropriate signature available for zero-day attacks.

That said, signature-based IPS solutions generally process packets faster than behavior-based solutions because the signature-based solutions mainly look for the particular contents of the packets but the behavior-based solutions look for many different aspects of the connections and have heavier algorithms running for the same reason.

Designers of behavior-based IPSs say such products offer a degree of automation and flexibility unavailable from signature-based systems. Such systems offer a continuous effort to refine the product in a way that allows network managers to minimize false positives.

Despite their architectural differences, signature-based IPS and behavior-based IPS solutions still perform similar tasks — they detect and stop network threats before they infiltrate and impact network applications or network performance. There are even hybrid solutions that embrace both approaches.

Deployment Modes

There are several modes in which to deploy an IPS that can improve performance, or adapt it to a particular network configuration.

Inline Transparent

In this mode, the IPS is deployed much like a Layer 2 bridge. It sits, inline in the network, but does not have an IP address and does not handle any routing.

The benefit of this stealth mode is that since there is no assigned IP address for each interface, there is no need to alter the network configuration and the IPS system is not visible to the attackers. Some vendors offer full stateful firewall with NAT and VLAN functionality even in transparent mode, allowing seamless deployment in existing VLAN networks without any reconfiguration.

Inline Gateway

The IPS is deployed much like a Layer 3 router, with IP addresses assigned to each port. IP subnets exist on each side of the IPS.

Network managers who deploy an IPS in inline gateway mode may have to alter the IP addresses of surrounding devices to make sure the adjacent network interfaces support the same subnet of the IPS ports.

The benefit of inline gateway mode, though, is that network managers gain better control over traffic, and can support IP subnets and virtual LANs (VLANs) by segregating incoming traffic.

High Availability: Active-Active

With network security becoming critical to the survivability of key applications, users may opt to deploy tandems of IPS appliances to ensure availability of the intrusion prevention service in the event a primary device fails.

Some IPS solutions vendors offer what is referred to as "active-active" mode IPS. Here, two identical IPS devices are deployed in the user network, sharing the incoming traffic between them. Operating software manages a "heartbeat" between the devices; in the event of an outage on one device, the sibling IPS takes full control of the incoming traffic when it fails to detect the heartbeat of the other IPS.

High Availability: Active-Standby

Think of the IPS active-active scenario, but this time the sibling IPS does not share the active traffic, but instead sits in hot standby reserve. In the event of an outage, the standby detects the lack of a heartbeat from the primary IPS and cuts over to handle traffic scanning.

The downside to this approach is that companies pay for two IPS appliances but utilize just one, while the standby sits idle.

Segmentation with VLAN

Some IPS solutions come with multiple ports, some of which are reserved for local traffic and some of which are set aside for WAN connections.

Users are able to dedicate and pair ports for a given VLAN to segregate traffic and ensure that only traffic destined for a given VLAN will ride across th dedicated ports. Some vendors support VLAN-tagged traffic in transparent mode and use the VLAN tags to apply different security policies to each segment or group of segments, thereby allowing advanced security virtualization.

Segmentation without VLANs

Users may also dedicate existing IPS ports to specific IP subnets. While all ports are typically shared in inline transparent mode, users who deploy an IPS in inline gateway mode can segment traffic by IP subnets.

The Architecture of the Threat

The architectural choices behind an IPS deployment are worth checking into, but eventually all IPS solutions sooner or later have to deal with the various security threats that threaten the network.

Bots, viruses, spyware, malware, worms — the list of malware goes on and on. While the threats seem endless, it is possible to characterize threats based on some rather simple elements. By doing this one can better understand the nature of new threats and how likely it is that your chosen IPS can deal with the threat.

Source of an Attack

Understanding the origin of an attack can help you better understand the nature of the attack, as well as how a given IPS architecture might defend against it. While we normally think of attacks coming just from the outside, attacks can originate from both sides of your IPS:

External

Someone trying to steal your data or infect corporate resources. All IPS solutions deal with this type.

Internal

Less often considered but also important are threats that originate from internal network hosts including those that use trojan horse techniques to plant malicious code inside the organization. It is the job of this code either to compromise corporate information by sending it outside the organization and/or using the internal corporate computers as platforms for launching distributed attacks on yet other targets. Thus, IPS solutions must also be called upon to block not only "inbound" attacks, but "outbound" malicious traffic, as well.

Note that some of these "launched from inside" attacks may not be as the result of the failure of the IPS to "catch" them on the way in. Rather, some attacks could end up on corporate notebooks that became infected while outside the control of the corporate IPS.

Some vendors offer advanced virtualization and segmentation in both transparent and gateway mode that enables further control against "inside-to-inside" attacks and malware propagation across internal network segments.

Carriers of an Attack

Understanding the carrier of the attack — that is, the legitimate method that an attacker uses to insert the attack into a network — is an important factor. Here are the main carriers that transport malware in your network:

E-mail

Malware, viruses, bots, malicious scripts and programs are frequently delivered as attachments to E-mail messages. It is not uncommon for the perpetrator to make the attack appear to be a harmless file, like an Adobe PDF, when in fact it is an executable program that can cause harm. From a technical perspective, these attacks are often referred to by the E-mail protocols they ride on (SMTP, POP3, IMAP). Stateful IPSs can help thwart these types of attacks.

Web Browsing

More frequently, the simple act of browsing a Web page can trigger an attack. Malicious scripts embedded in innocuous-looking Web pages can result in "nuisance" attacks, like changing your home page or adding a bookmark, but could also result in more serious

Attack Coverage Capabilities

IPS solutions may offer a range of attack coverage capabilities, including:

- Bidirectional Intrusion detection and Prevention
- Number of signatures and attack categories
- Number of IPS signatures (signatures for which IPS action is supported)
- Signature customization & support for custom signatures
- Protection against latest threats, e.g. worms, spyware, bots
- Protocol and traffic anomaly support
- DoS/DDoS attack protection
- Vulnerability scanning and exploit coverage (IPS as a "Pre-Patch" solution)
- Integration of additional security functions for better intrusion prevention — e.g. URL filtering and Layer 7 firewall for worm and spyware protection

attacks, as well. From a protocol point of view, these are referred to as HTTP attacks. URL filtering with IPS can help protect against these types of threats.

Instant Messaging

Similar to worms spreading through networks via E-mails, Instant Messaging (IM) programs are the latest vehicles being used to launch attacks across networks. IM networks not only allow the exchange of text messages, they also allow the transmission of files. Plus, hackers use IM to gain backdoor access to networks, bypassing listening ports and firewall filters. In today's world there are already dozens of IM-specific worms, such as the W32choke.worm that affects MSN Messenger. IPSs combined with stateful deep inspection firewalls can help prevent IM threats.

File Transfer

Files that are brought passively into the environment may also contain attacks that can cause harm when the file is opened. These need to be stopped before they get in. This transport is referred to as FTP for the protocol used.

Network Management

Even legitimate traffic, like requesting a DNS name resolution or testing for connectivity using a TCP Ping command, can be used maliciously and be turned into an attack.

Threats by Category

While most users are familiar with the various threats and vulnerabilities, we've highlighted some of the major categories and detailed their correlating business exposures. While some of these dominate the current security landscape, new ones are emerging and enterprises should be on the look out for them.

Viruses and Worms

Definition: A virus is a portion of programming code inserted into other programming to cause some unexpected and usually undesirable event. A worm is the network variant of the host-based viruses. Unlike viruses, these are not transmitted by physical media like floppies. They exploit the vulnerabilities in the software packages installed on user platforms to propagate and take over the networks. A worm is often different pieces of memory-resident and disk-resident code stitched together. The worm or virus might lie dormant until circumstances cause its code to be executed by the computer. Some viruses and worms are playful in intent and effect and some can be quite harmful, erasing or stealing private data or causing your hard disk to require reformatting.

Exposure: Results in loss of privacy, lost revenue, inability to conduct business transactions, reduced employee productivity and irreparable brand damage for the attacked.

Botnets

Definition: A botnet is a group of machines which have been compromised and made available to the external world under a common command and control infrastructure. Botnet hackers like to hide their presence as far as possible and use these resources for either storing malicious files or conducting network-wide activity which requires CPU/network resources. The compromised machine generally belongs to a user who got lured with a malicious E-mail or is not properly patched allowing other machines to exploit the intrinsic vulnerabilities to take over the system.

Exposure: Results in legal issues as the hijacked infrastructure could be used to launch Denial of Service attacks on other businesses and external networks, reducing IT productivity as the infrastructure is being utilized for alternate activities.

Denial of Service (DoS) and Rate-Based Attacks

Definition: A user or program takes up all the system or network resources by launching a multitude of requests, intentionally overloading the network, leaving no resources and thereby "denying" service to legitimate traffic. Typically, Denial-of-Service (DoS) attacks are aimed at consuming most, or all, of a network's bandwidth or a device's processing capacity thus denying that resource to legitimate users.

Exposure: Results in lost revenue, inability to conduct business transactions, reduced employee productivity and irreparable brand damage for the attacked.

Hybrid Exploits

Definition: A hybrid threat is one composed of multiple malware components and at times multiple carriers. This is one of the most serious threats which is facing the industry today as no single patch is sufficient to tackle them. They possess knowledge of multiple vulnerabilities and adapt themselves to exploit a set of vulnerabilties based on the surrounding environment.

Exposure: Results in lost revenue, inability to conduct business transactions, reduced employee productivity and irreparable brand damage for the victim business enterprise.

Remote Exploits

Definition: A means of gaining access to a computer system from a remote access point, typically through a known bug in a program or operating system. Many networks connected to the Internet that are not up-to-date with IPS platforms and security patches are vulnerable to exploits, and the effects of these exploits are seen when malicious worms run rampant and spread to unprotected systems.

Exposure: Results in lost revenue, inability to conduct business transactions, reduced employee productivity and irreparable brand damage for the attacked.

Malicious Content

Definition: Malicious content is any code added, changed, or removed from a software system, application or network in order to intentionally cause harm or subvert the intended function of the system. Traditional examples of malicious code include viruses, worms, and attack scripts, while more modern examples include Java attack applets and dangerous ActiveX controls.

Exposure: Results in costly and lengthy computer downtime and loss of stored data.

Spyware

Definition: A broad category of malicious software designed to intercept or control a computer's operation without the informed consent of that machine's owner or legitimate user. While the term had its genesis in network-habits monitoring software to improve the browsing experience for users, it has come to refer more broadly to software that subverts the computer's operation for the benefit of a third party. A typical example of spyware includes adware which would redirect the user to third-party sites based on what the user has been browsing while many of the malicious versions are involved in stealing user names, passwords, etc. These software either come packaged with shareware P2P software or exploit the weaknesses in browsers to infect the system.

Exposure: Results in loss of privacy, eats up system resources and makes it vulnerable to future threats.

Unwanted Access

Definition: Unwanted Access includes any attempts to access, modify, delete or retrieve proprietary information, intellectual property or consumer identity/credit information from files, applications, etc.

Exposure: Results in lost revenue, theft/loss of proprietary information, loss of consumer confidence, legal liability, and irreparable brand damage for the attacked.

IPS Benchmarking Model

Benchmarking Protection

While protection and performance are inextricably linked, it is complicated and often counterproductive to benchmark two things at once. Thus, while some testing certainly involves "threat detection under load" and "load handling" as separate topics and then combine the two as needed in your evaluation process.

Ultimately, one can best test for threat protection by taking the approach that ethical hackers take — but in a controlled laboratory environment. That is, one can only really determine if an IPS can stop an attack by attempting to mount that attack.

In recent years, several providers of benchmark tools have responded to this need. Vendors like Spirent Communications and Ixia have updated their load-generation tools to emulate specific threat signatures to the data payload. In other cases, vendors like Karalon and Blade Software have developed products that can generate hundreds of distinct threats to allow for a very thorough test of protection capability, albeit not with high loading, unless that load is provided by another test tool.

A number of freeware and shareware utilities that generate attacks for testing can easily by found on the Internet.

Finally, certain IPS devices may require simulation of a real application flow with appropriate responses from the server to work. In specific cases, a test tool's "artificial" traffic might simply be blocked because of its very nature. In these cases, users might need to build an actual environment to test a vulnerability.

Protection Response Mechanisms:

Every IPS offers a variety of mechanisms that can be employed to respond to threats or analogous network activity.

Basic:

- Silent packet drops
- Session (TCP) reset
- Alerts an packet logs

Advanced:

- Firewall hardening
- Throttling controls for preventing outbound DoS and DDoS attacks
- Network segmentation based on VLANs and IP subnets in both transparent and gateway mode
- Host/network quarantine
- Multi-function integrated defense mechanisms against hybrid attacks e.g. layer 7 firewall and URL filtering actions

Commercial Attack Sets

When evaluating IPS products, IT security professionals must be able to test for the detection and prevention of DoS attacks to the network, such as a Juno attack, while allowing legitimate traffic to process quickly and efficiently. During testing, all filters should be enabled for maximum security performance.

In this group, we find test tools that originated as Layer 2-3 or Layer 4-7 load generators that were enhanced to carry some specific threat signatures in their test payloads. While they typically can offer a high loading rate of multiple Gigabits per second or higher, the variety of attacks that they can be configured to carry is typically limited. In essence, no matter what type of traffic the test tools generate, that traffic must appear real to the IPS devices under test. Further, even though the same test tool may generate test traffic, that traffic may be considered as real by one IPS and subsequently handled as legitimate traffic or attack traffic, while another device may not categorize the traffic as real and let the packets pass.

Extensions to Load Testing Tools

Vendor: Ixia

IxLoad: This is a highly scalable solution for accurately assessing the performance of content-aware devices and networks by creating real-world traffic scenarios at the TCP/UDP (Layer 4) and Application (Layer 7) layers. It emulates clients and servers for Web (HTTP and SSL), FTP, E-mail (SMTP, POP3 and IMAP), Streaming (RTP and RTSP) and services such as DNS, DHCP, LDAP and Telnet and also provides emulation of DDoS attacks inline with the above protocols.

For more info on Ixia's IxLoad, visit *http://www.ixiacom.com*.

Vendor: Spirent Communications

SmartBits: Used for high port density testing of 10/100/Gigabit and 10 Gigabit Ethernet, ATM, POS, Fibre Channel and Frame Relay networks and network devices, featuring test applications for xDSL, cable modem, IP QoS, VoIP, MPLS, IP Multicast, TCP/IP, IPv6, MPLS, routing, SANs, IDS/IPS and VPNs.

Avalanche: Challenges any computing infrastructure's ability to stand up to the load and complexity of the real world by generating in excess of 45,000 HTTP requests per second, and by supporting HTTP1.0/1.1, HTTPS, RTSP/RTP (Apple QuickTime™), RealSystem™ streaming, Microsoft™ Windows Media™ 8 and 9 Series, SMTP, POP3, DNS, Telnet, DDoS and FTP. Through clustering, it scales to the needs of even the largest infrastruc-

For more info on Spirent's products, visit *http://www.spirentcom.com*.

tures. Real-world conditions are accurately replicated by simulating error conditions, realistic user behavior, and maintaining over one million open connections from distinct IP addresses.

Reflector: Simulates the behavior of a large cluster of Web servers, responding to user requests generated by Avalanche. It is the only commercially available product that can withstand the traffic levels generated by Avalanche. As a result, any system placed in the middle of these two products can be assessed for their impact on system performance.

Purpose-Built Vulnerability/Recognition Testers

Vendor: Karalon Software

For more info on Traffic IQ Pro, visit http://www.karalon.com.

Traffic IQ Pro: Tests the effectiveness of any in-line device including intrusion prevention (IPS) or intrusion detection (IDS) systems, in lab or production environments by generating source and destination virtual machines and transmitting standard or malicious traffic statefully between them. This software tool includes an extensive library of over 700 captured protocol and attack files. Users can pick any number of traffic files from the library or create their own traffic files with the built it conversion process. This tool can test the breadth of detection and prevention for the IPS devices. There is no need for advanced knowledge of scripts or multiple platforms.

The Internet offers some decent programs designed to aid with IPS testing. While such program maybe found by Web searching, The Tolly Group cannot attest to their effectiveness.

Shareware Attack Sets

The Internet offers some decent programs designed to aid with IPS testing. While such programs may be found by Web searching, The Tolly Group cannot attest to their effectiveness.

Home-grown Attack Sets

By definition, these are attack scenarios that, for whatever reason, testers cannot find in commercially available offerings. In most cases, this involves setting up a "client" to attempt the attack and a "server" as the intended victim on opposite sides.

To conclude, the IPS should effectively protect your network from undesired access with stateful firewall filters, from network layer attacks, from DoS and DDoS attacks, from known remote exploits, from unknown and zero-day exploits, while ensuring positive traffic management capabilities that protect against resource consumption and inappropriate traffic rates. However, in order for the IPS to protect your network properly, outstanding performance is required.

Benchmarking Protection Actions

The set of protection actions which an IPS vendor has to offer is also very critical in evaluating the efficiency of the application. These go hand in hand with the detection mechanisms and often the choice of the protection mechanism is dictated by the kind of threat which has been detected.

Drop/Reset Action

The most common method of protection is to drop the packet or reset the TCP session. These actions are often set up to drop the malicious exploit packets. The downside is that for protocols like TCP, they could often leave the sessions hanging in a retransmission mode or in some cases could be used to launch a DoS attack by causing the IPS to generate a lot of reset packets. Also, a reset action can reveal the location of an IPS even if it is in transparent 'stealth' mode.

Firewall Hardening Action

This is when an attacker IP address is blacklisted for a certain duration. A very effective technique when a network is under a repeated attack from a particular network address.

Session/Bandwidth Limiting Action

Session limiting is when the system does not allow more than a fixed number of sessions to a particular server. In bandwidth limiting, the total allocated bandwidth to the server is restricted. DoS/DDoS attacks are often effectively blocked by defeating the purpose for which they have been initiated. By effectively limiting the number of sessions or the bandwidth usage to a critical resource, this action allows critical servers to remain online while effectively keeping away attackers. Advanced functionality offered by some IPS vendors includes selection of these limiting actions on a attack-definition basis.

Network Segmentation

Segmenting the network and preventing attacks from infiltrating from one network segment to another offers an effective protection strategy against worms and internal DoS attacks. Segmentation capability combined with the firewall hardening and session/bandwidth throttling can thwart potential packet flooding and zero-day internal DoS events.

Benchmarking Performance

Performance is equally crucial to protection in measuring the effectiveness of an IPS, which needs to run seamlessly throughout the network — always allowing genuine traffic while always blocking malicious access. The IPS must simultaneously protect your assets against threats and allow legitimate business transactions. The IPS should have high packet processing rates, high network throughput, low network latency, high session capacity, the ability to switchover to a redundant device in an application-aware fashion and high transaction inspection rates. Measuring performance also involves the ongoing support of simultaneous sessions, while introducing large quantities of new sessions per second. In other words, the IPS should not be the performance bottleneck in the network and should have headroom in performance for the preparation of any future network upgrade.

Essential Metrics

Throughput

The first area to consider is throughput speed, which is measured in gigabits or megabits per second which can be sustained with zero drops. Your intrusion prevention solution must be able to inspect traffic for attacks in real-time and in a ***bidirectional*** basis without performance degradation. The IPS must also perform scans and process data at multi-Gigabit speeds and while handling a variety of packet sizes and types: UDP and TCP packets, both large and small, even an Internet mix.

Throughput is critical to an inline security device such as an IPS. Since the device is processing every packet coming into the network, network managers want to make sure it is not a bottleneck. Companies don't want to pay for Gigabit connectivity on the local network, only to find that it is cinched down to 100 Mbps because the IPS has become a bottleneck.

Throughput measurements assure network managers what packet size or traffic mix that passes through the IPS. It represents, on average, how many bits per second can be processed by the IPS but, doesn't tell you how long it takes for any packet to go through the processing chain — that's where latency comes in. Note that different IPS devices can have different processing chains.

Latency

The next area is system latency — the time difference introduced in the overall end-to-end packet delivery by deployment of the IPS system. It is typically measured in milliseconds or microseconds — for example, less than 225 microseconds. The resulting figure must remain low, regardless of traffic mix or number of attack filters installed, in order to avoid becoming the performance bottleneck. A poorly designed IPS may result in introducing intolerable levels of latency into the network.

Latency measurements are important indicators of the intrinsic packet-processing delay introduced by the IPS — or lack thereof. An IPS with a high throughput but with a one-second delay in processing the traffic would not be good because it would impact latency-sensitive applications, like voice and video, adversely affecting quality or rendering the applications useless.

Session Rate

This is another important performance metric when the tested device deals with the real-world stateful traffic — how many connections can it handle per second? This looks at the maximum connections per second.

Transaction rate measurements are helpful in determining how many transactions a device can handle — which ultimately helps users determine if the IPS can keep up with the traffic hitting the device.

Session Scalability

The last area for consideration is the number of simultaneous sessions supported when the IPS is running in operating mode, along with the number of new sessions introduced, measured on a per-second basis. It's best to test in the Inline mode to gauge packet and session-loss rates (striving for zero). It is also recommended to test with real-sessions: e.g., complete TCP handshake (UDP only, SYNs only) to measure DoS protection.

Essential Variables

Load

You must test the IPS for its rated capacity — e.g. a 4 Gbps system should be tested for 4 Gbps throughput.

Testing different percentages of "clean" traffic vs. attack traffic is also essential in determining the effectiveness of the performance. We recommend first testing with 100% clean traffic, then follow it up with a slow ramped up induction of attack traffic regulated by an arbitrary percentage of increment. The enforcement platform should demonstrate the ability to drop the malicious frames and the monitoring software should accurately follow the alerts generated by the malicious frame rate, thus providing the requisite data for post event analysis.

Signature Base

The majority of IPS solutions offer different levels of signature matching. You need to ask the question, "How many signatures are being searched in the database? How many are relevant to your network scenario? What kind of security are the enabled signatures providing?"

Not only is the bottom line number important, the IPS should also allow for signature customization.

IPS-specific Configuration Modes

Often, the IPS will offer different modes of operation that trade off function for speed. These modes might radically change the processing that takes place with each packet/connection and thus result in radically different results. So, test with the IPS turned off and on, and at varying levels of protection to measure how it impacts performance.

Benchmarking Usage Model

Before you select the IPS that's right for your organization, you must consider the system's ease of use. You must ask, "How easily will the IPS fit into our network?" Be sure to take into consideration how it fits physically, topologically (its continuity and connectedness), and logically within the network architecture and design. Since you may have to convince the NetOps team to install the device, we're providing you with some specific areas to address, which will help alleviate any concerns.

Tuning & Customization

Definition

Some IPS solutions allow you to customize the IPS and/or modify the base rule set and the IPS action associated with individual rules to accurately represent the threat model to suit your environment.

Benefits

With some architectures, the size of the rule set or signature file impacts performance. In such cases, reducing the size by eliminating signatures that are not relevant to your implementation can improve performance.

Sometimes valid traffic "looks" like an attack and "false positives" are generated. False positives occur when legal traffic flowing in your network is viewed incorrectly as a threat and blocked, resulting in a direct negative business impact. Some devices will let you tune the set to minimize or eliminate false positives and false negatives (when a virus enters your system undetected by the IPS).

Evaluation strategy

Identify specific capabilities of tuning and customizing the IPS for your network.

IPS Management

Definition

Your organization must effectively utilize the IPS to realize its full value. Accept the detailed training offered by the vendors on what your equipment offers. The vendor should offer automatic weekly or on-demand updates with patches to address the continuous onslaught of newer and more sophisticated attacks and zero-day threats. Not enough emphasis can be placed on staying on top of all advisories provided by your vendors, as they closely scrutinize all newly released vulnerabilities and maintain a staff of highly skilled IPS professionals on hand to address new and emerging threats.

Some vendors will even offer a subscription service for virus signature updates, and the IPS refreshes with the vendors' servers on a predetermined basis — even every two hours in some cases. The upshot here is that timely signature updates mean that IPS appliances are well positioned to defend against fast-replicating new threats.

Benefits

You know you are cutting down the window of exposure to vulnerabilities and achieving the full potential out of your IPS. Enhances the effectiveness of the IPS device.

Evaluation Strategy

Measure the ease of use and the learnability of the user interface. Ask these questions:

- Does the IPS allow for event suppression-vulnerabilities that you can't dismiss, but don't have the time at present to review? Do not classify your threats as simply black and white. The IPS should allow for a grey area to be reviewed later.
- Does it group events by severity and threat type?
- Is there a network-wide consolidated reporting ability to identify warning signs of change in the health of the whole network?
- Is there real-time consolidated reporting?
- Does it prioritize risks and vulnerabilites — automate low and medium risk functions as much as possible, while reserving human interaction for the highest risks?
- Does it help to meet government regulations?
- Does it help to meet compliance standards?

Benchmarking Reliability

When considering the reliability of the IPS solution, you must judge both the hardware and software provided. With respect to hardware, being an inline device, the IPS should display the same level of reliability as your switches and routers, and preferably have a high MTBF rating with redundancy features.

The main concern for the software end is to look for robust software with well-defined dependencies of the third-party software. Third-party software components are often the target of hackers and it is essential to have the process to track these known vulnerabilities and analyze their impact on the core IPS software.

The second area of review should be the average lifecycle of the software versions. Like other network devices, this software is going to become an integral part of your network and an administrator cannot afford frequent upgrades to the software to support the next generation of threats.

Similar to software upgrades, a non-disruptive signature update process is also very critical to allow the system to provide protection from the threats as and when they emerge without requiring the network to be brought down.

Some vendors believe Open Source provides the best foundation for any IPS device because any "commonly available" OS is inherently susceptible to infiltration. Some IPS vendors have embraced "SNORT," an Open Source intrusion detection and prevention system.

Proponents of the Open Source approach point out that an Open Source product provides buyers visibility into the IPS protection layer, which is fairly unique in the intrusion market. Peer review also comes into play, proponents of Open Source say, because there is a community of interest using SNORT and when a problem is identified and a fix developed, it is disseminated quickly throughout the SNORT community.

Security buyers, Open Source advocates say, want and need unlimited access to how the IPS makes decisions. Such peer review provides users the opportunity to comment on the product to make it better.

Either way, be sure to identify the underlying Open Source approach being used by vendors and investigate the strengths and vulnerabilities of the underlying Open Source approach.

Hardware: Purpose-built Appliance

Definition

System or key elements built specifically for the security application and used instead of "general purpose" hardware components.

Evaluation Strategy

It is becoming increasingly difficult to differentiate between "purpose built" and "general purpose" implementations without dissecting the devices. (You would often need not only to "look under the hood" but to understand the chips and their functions — not something an Enterprise implementer has time to do.) Here is a punchlist of preferred high availability and stateful network redundancy requirements for your IPS:

- High MTBF hardware with redundancy features and support (ability to synchronize large number of active sessions across redundant systems to minimize disruption)
- Separation of control plane and data plane architecture
- Dual power supply on high-end platforms
- Fail-open as an option
- Switch-like performance — many vendors offer a range of per-second attack filtering levels, depending on your network capacity.

Summing It Up

Sizing up an IPS is no small task. There are a multitude of variables to consider.

From the start, users need to consider the underlying architecture of the IPS. Do the traditional signature-based IPS solutions serve your particular needs? Or would your enterprise be better off with a newer IPS design that relies on the behavior of network and application activity to make threat assessments? Or, should you invest in a multi-security integrated solution that includes integrated IPS with stateful firewall, URL filtering, etc. for adaptive multi-threat prevention.

Beyond the architecture, network managers need to mull over the myriad of deployment options available. Inline transparent mode deployment of an IPS may be expedient, but if you have a highly segmented network or one that relies heavily on VLANs make sure the selected product supports VLANs in transparent mode. Or is an inline gateway more suitable since it delivers greater control over traffic?

Don't forget about reliability and survivability. When and where does it make sense to deploy a pair of IPS appliances, and if you do, how exactly should they be deployed?

When it comes to benchmarking performance of the IPS, throughput, latency, transaction rate and session scalability reveal critical information about the capabilities of the device. Network managers also need to make sure if the IPS device has enough performance headroom for future upgrades.

The Tolly Group recommends that users measure these capabilities in a test bed that closely maps to the real-world network situation. Accurate representations of traffic mixes, including variable packet sizes, should be tested. Wide-area connections should be simulated to parallel what is used in the enterprise. And network gear representing the makes and models used in the live network should be deployed in the test bed.

When it comes to test tools, vendors have developed a reliable stable of tools to simulate real-world network traffic flows, libraries of attack signatures and client/server interactions. The Tolly Group recommends using any of the commercially available test tools to simulate real-world network traffic conditions or introduce threat signatures onto the test network. Such tools will provide the most cost efficient way to replicate actual network conditions.

Along with network and security performance, the reporting scheme is a must-to-have feature for the IPS and one of the most important aspects for the purchasing decision.

Lastly, examine the tuning and customization, management facilities and reliability aspects of the IPS solutions under review. Each of those areas will provide a window into the flexibility and ease-of-use offered by the IPS.

Using these IPS benchmarking strategies will help ensure optimal performance from the IPS solution that's right for your business or organization.

# # #

Information technology is an area of rapid growth and constant change. The Tolly Group conducts engineering-caliber testing in an effort to provide the internetworking industry with valuable information on current products and technology. While great care is taken to assure utmost accuracy, mistakes can occur. In no event shall The Tolly Group be liable for damages of any kind including direct, indirect, special, incidental, and consequential damages which may result from the use of information contained in this document. All trademarks are the property of their respective owners.

TOLLY 17 YEARS 1989-2006

THE TOLLY GROUP

The Tolly Group, Inc.
3701 FAU Blvd. Suite 100
Boca Raton, FL 33431
Phone: 561.391.5610
Fax: 561.391.5810
http://www.tolly.com
info@tolly.com

TOLLY Up to Spec CERTIFIED