Defining the Rules of Preemptive Protection: The ISS Intrusion Prevention System

By Christopher J. Rouland
Chief Technology Officer

INTERNET | SECURITY | SYSTEMS®
Ahead of the threat.

Introduction

Intrusion Prevention Systems (IPS) are rapidly becoming an integral part of an effective network defense solution. Unfortunately, finding the truth in today's often over-hyped market of network-based IPS offerings is no easy task. As the technologies behind IPS become increasingly complex, so does determining which IPS solutions can actually deliver preemptive protection, a new standard in security that stops attacks *before* they impact the network.

Before attempting to analyze any vendor's IPS offering, it is important to understand that network security is not an absolute. The network security landscape has become cluttered with buzz-word technologies, snake oil solutions and panaceas all advertising complete protection. Often, vendors making these claims do not account for the dynamic nature of online threats, resulting in solutions that are only effective against a small subset of threats in the wild. It is important to recognize that no singular IPS technique provides adequate protection against all known and unknown network security threats. Like traditional physical security, every unique Internet threat may require a new approach to best detect and neutralize it before it causes damage.

So how can you determine which IPS will deliver accurate, preemptive protection against the next Internet threat? The rules of preemptive protection are clear. Network IPS products that block attacks before impact must offer optimum performance, provide the highest level of protection, and rely on a solid foundation of research covering both threats and vulnerabilities.

Figure 1: Preemptive Protection Requirements

As illustrated in Figure 1, an IPS must have superior characteristics in the following three areas to enable preemptive protection:

- **Performance** – The ability to perform transparently in the network environment while also supporting the other critical areas of preemptive protection.
- **Protection** – The ability to provide a high level of protection requires many protocol identification and analysis techniques to ensure optimum accuracy.
- **Research** – Powerful intrusion prevention is based on up-to-the-second security intelligence that keeps pace with the changing threat landscape. This requires an in-house research team that fully understands network security threats and vulnerabilities, and injects that knowledge into the product as threats adapt and before they impact business.

Now that the three rules of preemptive protection are defined: performance, protection and research, evaluating the efficacy of an IPS offering becomes much easier.

Performance

The first rule of preemptive protection from an Intrusion Prevention System is performance. IPS performance should be ideally matched to the environment being protected. Several sub-categories outlined below contribute to the overall performance of an IPS.

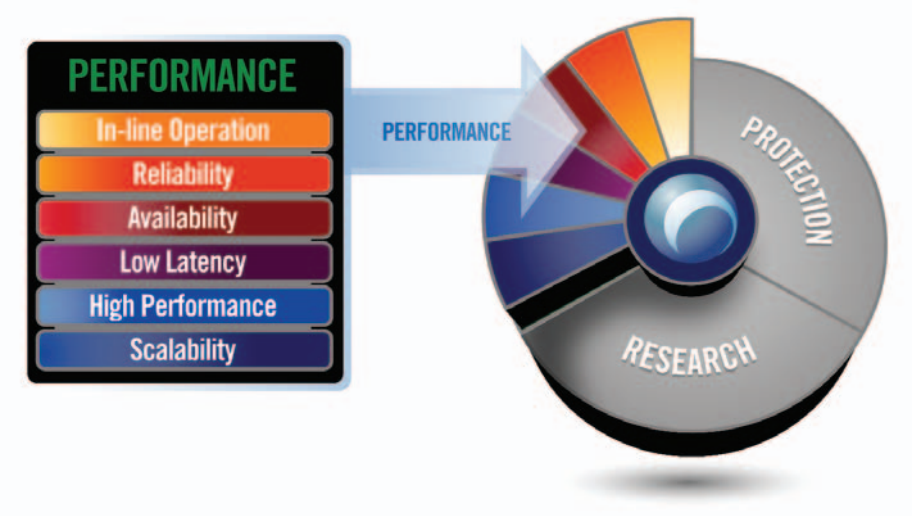

Figure 2: Performance Requirements

In-line Operation
An effective IPS must operate transparently in-line on the network. Transparent in-line operation results in minimal impact to information technology (IT) infrastructure.

Reliability
Intrusion prevention is usually applied at critical network infrastructure points. Therefore, IPS failures have the potential to cause system outages. With crucial information and systems on the line, IPS solutions must be highly reliable with a long Mean Time Between Failure (MTBF).

Availability
At a minimum, a network IPS must not interfere with traffic should it malfunction or go into an offline state. To avoid this outcome, network IPS devices should fail open, regardless of network media.

Low Latency
Network-based IPS devices must introduce a minimal amount of latency to network traffic. Low latency is often the most critical performance factor for network intrusion prevention.

Example: Business critical Voice over IP (VoIP) applications begin to degrade noticeably at approximately 1,500 microseconds[1]. An in-line IPS must not introduce significant latency to affect the business continuity of such an application while at the same time providing 100 percent intrusion prevention coverage.

High Performance
A network-based IPS must exhibit many of the performance characteristics of switching and routing equipment, while simultaneously blocking threats to the network and the devices connected to it.

Example: Line-speed refers to the ability of a device to process packets at the maximum speed-rating of the network. The IPS must process traffic at line-speed to avoid "bottlenecks" that might open the door to a denial of service (DOS) condition should the traffic overload the device.

[1] http://www.networkmagazine.com/article/NMG20000710S0012

Scalability

At the network level, IPS devices must scale to a large number of user sessions and transactions without disrupting business continuity.

IPS performance requirements and characteristics differ slightly depending on whether intrusion prevention is deployed on the network or within host-based systems like servers and desktops. In either case, performance remains a key purchase consideration to ensure that the IPS causes no disruption to applications residing on servers and desktops, and application communication within the network.

Protection

An Intrusion Prevention Systems' overall ability to protect against threats comprises the second rule of preemptive protection. Understanding protection technologies starts with understanding their origin. Intrusion Prevention System technologies evolved from Intrusion Detection Systems (IDS). While IDS devices identify threats and send out alerts, Intrusion Prevention Systems go one step further to block attacks from impacting the network.

Security professionals generally agree that Intrusion Detection Systems rely on either signature-based methods or protocol analysis-based methods to accurately identify threats. Within each basic IDS category, more specific technologies exist, ranging in complexity from elementary to highly advanced[2]. The highly evolved IDS analysis techniques prove more accurate than elementary methods and readily apply to Intrusion Prevention Systems, the next generation of security solutions. Since intrusion prevention is designed to block attacks while allowing legitimate traffic, accurate attack detection is critical. Established security vendors with a strong heritage in IDS have optimized the accuracy of their Intrusion Prevention Systems by fine-tuning the underlying analysis techniques over time. Newcomers to the security space now offering IPS haven't had the same opportunity.

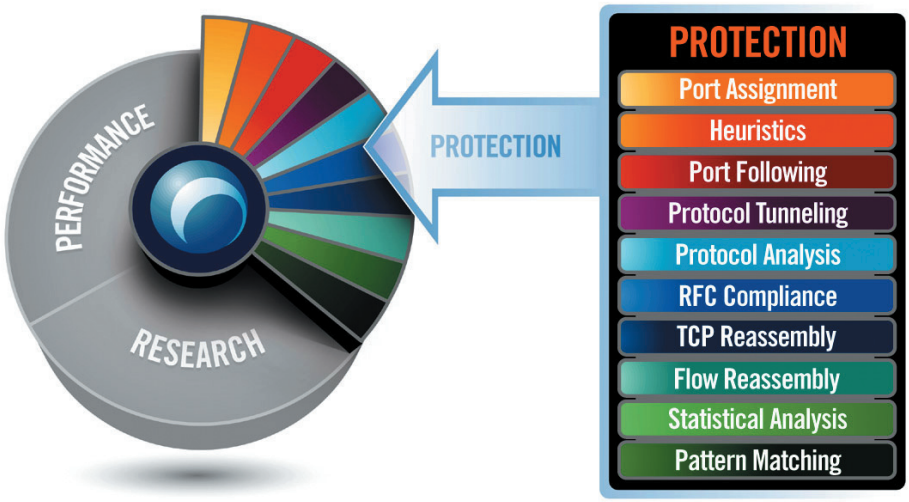

Figure 3: Individual Identification and Analysis Methods that Enable Preemptive Protection.

The Protection Toolkit

Every skilled craftsman uses a collection of tools to perform a job. Unfortunately, hackers are no different and often use collections of purpose-built tools to make breaking into systems easier. IPS devices, therefore, must rely on their own set of tools to combat attacks. The individual tools in a robust IPS toolkit fall into two high-level categories: identification and analysis (see Figure 3 for a full list of IPS techniques). The identification category consists of tools that help the IPS accurately identify the protocol encountered within the network traffic. In the analysis category, tools analyze identified protocol traffic for malicious behavior, indicating what should be blocked or allowed.

[2] http://www.giac.org/practical/Paul_Barry_GSEC.doc

For accurate, preemptive protection, IPS solutions should use multiple identification and analysis methods (listed in Figure 3). On their own, each method has inherent strengths and weaknesses, accounting for why no single intrusion prevention technique offers acceptable levels of protection. *False positives* occur when an analysis technique triggers an alert and/or response action when there is no actual threat. When an IPS produces false positives, it can disrupt normal business operations by interfering with legitimate traffic. The opposite error, *false negatives*, occurs when the IPS does not accurately detect an actual security threat, which usually results in more damage. In the case of false negatives, the IPS provides no warning for successful attacks—making the outcome particularly dangerous. Intrusion Prevention Systems that employ a combination of identification and analysis techniques coupled with rich, up-to-date security content, effectively reduce false positives and false negatives.

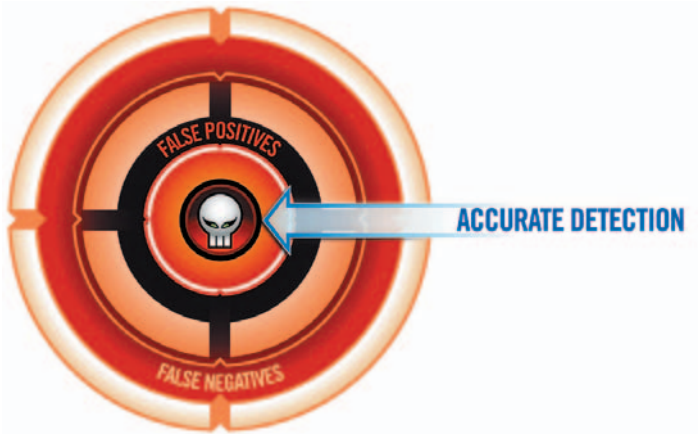

Figure 4: Accurate Detection

Protocol Recognition and Identification Techniques

Before analysis of protocol traffic can begin, the traffic must be accurately identified. All remaining steps of traffic inspection hinge on the accuracy of this initial process. Traffic parsed incorrectly will render false positives at best and false negatives at worst. Using multiple techniques, protocols can be accurately identified with a high degree of confidence. The following is an overview of some of the more popular recognition and identification techniques used in modern Intrusion Prevention Systems.

Port Assignment

Port assignment is the most elementary method of identifying application protocol types. The technique of port assignment assumes the application protocol type based upon the Transmission Control Protocol/Internet Protocol (TCP/IP) port being used for the connection.

Example: Hypertext Transfer Protocol (HTTP), an application protocol used primarily for World Wide Web (WWW) communication, is most commonly found on Transmission Control Protocol (TCP) ports 80 or 8080. Since it is such a common communication protocol, traffic on these ports is almost always HTTP.

Port assignment is valuable as a preliminary protocol identification technique. However, because protocols are not bound absolutely to particular ports, using port assignment alone poses problems. An IPS that assumes protocols are always bound to particular ports provides intruders with a very elementary way to evade the system, possibly resulting in a successful attack. In order to reduce false negatives, traffic identified by port assignment should always be double-checked with another recognition technique to ensure that attacks are blocked.

Heuristics

Heuristics, in the context of protocol identification and recognition, involves developing algorithms used to positively identify traffic. The algorithms are based on sets of rules that uniquely identify the protocol's behavior. For example, Instant Messaging (IM) applications often purposely avoid using specific ports so that they can take advantage of whatever ports remain accessible via the firewall. Heuristics is often the only method of correctly identifying certain protocols.

Example: The Remote Procedure Call (RPC) protocol can be found on virtually any port, making port assignment an ineffective tool for accurate identification. In certain cases, RPC traffic can even be found on TCP port 80, which is almost always HTTP traffic. To accurately identify RPC protocol traffic so that analysis can be performed, traffic must be closely monitored for key fields inside valid RPC messages for indications that are unique to RPC protocol traffic. If all conditions are met, regardless of port, the IPS engine should parse and analyze the traffic as the RPC protocol.

Heuristic techniques assume that unique identifiers in the traffic always exist. But due to some protocol designs, unique traffic identifiers are not always present. The next technique, port following, is sometimes used as an additional method for accurately identifying protocols and their traffic flows.

Port Following

The *port following* technique monitors previously identified communication sessions for additional connections on random ports. Some application protocols use an initial port to control a connection, but then negotiate and open a random port to transfer data between the client and the server endpoints of the connection.

Example: The File Transfer Protocol (FTP) commonly uses TCP port 21 to initiate and control an FTP session. When data transfer is initiated, however, the client and server usually negotiate an additional random port number to transfer the raw data. This traffic does not exhibit any characteristics of FTP traffic because it is the actual binary data with no control structures inline. The port following protocol identification technique proves useful in this situation because it will group the negotiated file transfer port with its associated connection and the traffic will be analyzed appropriately.

Protocol Tunneling Recognition

Protocol tunneling is the practice of "embedding" one application protocol within another—a common occurrence in modern network communication. In some cases hackers will use protocol tunneling to disguise their attacks—so the ability to recognize this evasion technique is critical to preemptive protection.

Example: The AOL Instant Messenger (AIM) protocol is designed to work in the most restrictive environments. By default the application attempts to use TCP port 5190 for AIM protocol connections to the Instant Messenger (IM) server. If this port is blocked by firewalls, the application can "tunnel" the AIM protocol within the HTTP protocol on TCP port 80, which is almost universally allowed through a firewall, to connect to the AIM server. In this instance, protocol tunneling ensures that the AIM application will work in most environments. An IPS that recognizes port tunneling will still be able to block an attack embedded in the AIM protocol.

Traffic Analysis Techniques

Traffic analysis takes place after traffic has been correctly identified. Further analysis beyond basic identification helps the IPS determine the intent of the traffic and take appropriate steps to block malicious traffic. As with identification techniques, no single method is effective enough on its own. Therefore, an IPS with multiple analysis techniques working in tandem provides additional protection. Below are some examples of analysis techniques that IPS solutions should employ.

Protocol Analysis

Protocol analysis is a popular technique used by IPS devices to stop known and unknown threats. Known threats consist of attacks and exploit code already released into the wild, while unknown threats are yet to be released and also have the potential to target known and unknown vulnerabilities. Protocol analysis can be performed on protocols down to level 2 of the Open Systems Interface (OSI) Model layer[3]. Using protocol analysis techniques, the IPS double-checks a connection's communication against the generally accepted behavior for the protocol. If a network transaction does not follow the accepted behavior, the traffic is blocked or an alert is generated, depending on the configuration of the IPS engine.

[3] http://www.webopedia.com/quick_ref/OSI_Layers.asp

Example: The Local Security Authority Subsystem Service (LSASS) vulnerability that was released on April 13, 2004 in conjunction with Microsoft Security Bulletin MS04-011[4] provides a good example of how protocol analysis can be used for preemptive protection against an unknown attack. This particular LSASS vulnerability allows a remote attacker to execute arbitrary code with system privileges. After studying the LSASS vulnerability and its underlying protocol, ISS released a security content update to its IPS in tandem with the vulnerability disclosure (April 13, 2004) to protect against potential attempts to exploit this vulnerability.

Just two weeks later on April 29, the "Sasser" worm emerged targeting the LSASS vulnerability. ISS customers were already protected against the Sasser worm and any other potential attacks using this vector (in this case the LSASS vulnerability) as a point of entry into a host. By studying the underlying LSASS vulnerability and using protocol analysis to block any attempt to exploit it, ISS was able to preempt the Sasser attack along with any Sasser variants. Many other IPS solutions had to wait for the specific Sasser exploit logic to emerge before they were able to protect against the attack.

RFC Compliance Checking

Request for Comments (RFC) compliance checking, also commonly called *protocol validation* or *protocol anomaly* detection, triggers when network traffic does not conform to the RFC standard[5]. This technique produces a high rate of false positives because developers are not required to adhere to the application protocol's RFC. RFC compliance checking also tends to produce a lot of false negatives because most attacks are considered "legal" according to the application protocol's RFC standard. Therefore, RFC compliance checking should rarely be used by itself and is most effective when combined with another technique.

Example: VoIP messaging protocol H.323 contains several vulnerabilities and provides a good example of network traffic which must adhere to the protocol closely in order to work properly. Checking RFC compliance is very useful in this situation to protect against known and unknown attempts to exploit devices which support the protocol.

TCP Reassembly

Packet fragmentation, the splitting of one original packet of information on the network into two or more packets, is a normal networking operation due to varying transport protocols. Hackers also employ fragmentation as a method for evading elementary detection systems. Tools such as Fragroute make it easy to break malicious attack packets into smaller fragments before sending them across the network. To handle the normal conditions that exist in a network environment, as well as abnormal attempts to obfuscate an attack, IPS devices must be able to reconnect pieces of traffic that belong together. This preprocessing is called *TCP reassembly* and should always be used to analyze traffic for hidden signs of malicious intent.

Flow Assembly/Simulation

Flow assembly or *simulation* is similar to TCP reassembly, but requires that the IPS keep up with a connection in its entirety (as opposed to a packet or a portion of the data flow). Flow assembly must analyze the connection as a whole rather than inspect individual portions of the traffic as they are encountered. A variety of modern threats use fragmentation techniques to avoid detection by security devices. By reconstructing the traffic flow of the connection, the IPS can identify threats that would have evaded the system otherwise.

Statistical Threshold Analysis

Statistical threshold analysis is based upon detection and blocking of network anomalies. This technique is also sometimes called *statistical anomaly* or *threshold analysis,* and usually involves monitoring the network for a period of time to create a "baseline" of what normal traffic patterns look like. Once the baseline is established, patterns that exceed the threshold of the baseline are suppressed. Establishing baselines and using other statistical anomaly techniques can effectively stop threats that generate obvious deviations from normal traffic. Other, more subtle threats may slip under the radar of IPS devices relying solely on statistical analysis. Many vendors claim that statistical analysis stops all unknown threats, but their theory is flawed because of the dynamic nature of most modern computer networks. In a dynamic environment establishing baselines is difficult, cost-prohibitive and limited in scalability as a stand-alone component of an IPS.

[4] http://www.microsoft.com/technet/security/bulletin/MS04-011.mspx

[5] http://dict.die.net/rfc/

Example: Port scanning is one type of traffic pattern that statistical analysis techniques identify very efficiently. Port scanning involves connecting to TCP/IP ports in an attempt to discover available services on a host. Port scanning is considered a form of reconnaissance, and port scanning results are usually used later to attempt to break into open ports[6].

Pattern Matching

Pattern matching, the most popular method of analyzing threats, also maintains the worst reputation. Pattern matching is also called *regular expression (regex) matching*. In lieu of using regular expressions, some security vendors have implemented custom pattern matching language that simulates the effect of using regular expressions. Pattern matching involves scanning network traffic as it passes through the IPS for patterns that have been predefined to signal malicious behavior.

Pattern matching remains a very useful tool in the detection of security threats. The truth is that *all* IPS vendors use pattern matching to some degree in their traffic analysis. Pattern matching's lowly reputation as a weak IPS technology results from its history as the first method of detecting threats. In its infancy, pattern matching was very elementary, effectively triggering on any traffic that matched the pattern of bad behavior. This basic technique is commonly referred to as *packet-grepping* or *blind pattern matching*. The packet-grepping name was derived from the popular "grep" tool for UNIX-based systems which is a utility that finds patterns in strings. Initially, pattern matching triggered a high volume of false positives resulting in a higher cost of ownership for those using IDS.

The pattern matching analysis technique has evolved, however, and current solutions use algorithms that trigger a match only if the pattern matches in a portion of the traffic that could actually result in successful vulnerability exploitation. This technique is sometimes called *stateful pattern matching*. As the name implies, the IPS signals a match only if the attack appears in the particular portion of the traffic where an attack would actually exist.

An effective Intrusion Prevention System must employ a combination of the techniques discussed above to provide accurate, preemptive protection against known attacks like Sasser and MS Blaster, and new attacks not yet released into the wild.

Research

The final rule of preemptive protection hinges on research. Up-to-the-second security research must be incorporated into the IPS as rich security content, often in the form of logic or algorithms. Not all IPS vendors collect the same caliber of research, resulting in security content that varies in effectiveness. As with the performance and protection components of an IPS, no single research methodology is adequate. For preemptive IPS solutions, research must encompass both proactive and reactive methods covering both threats and vulnerabilities; global event monitoring; and information-sharing with other research organizations, industry consortiums and government entities.

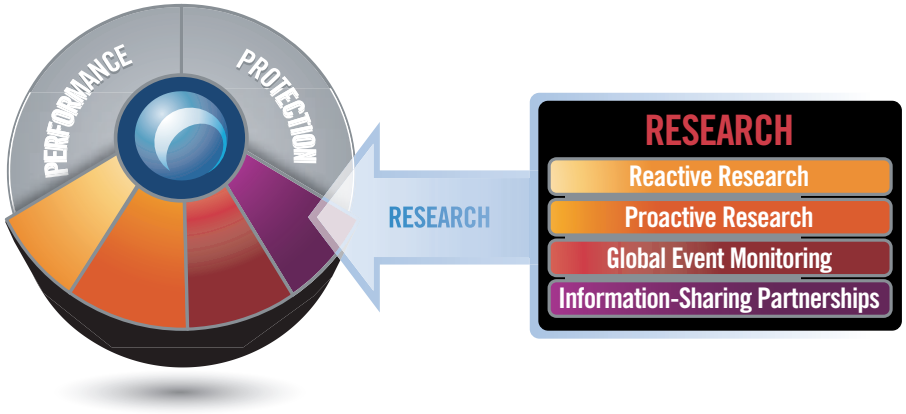

Figure 5: Research Requirements

[6] http://www.iss.net/security_center/advice/Underground/Hacking/Methods/Technical/Port_Scan/

Reactive Research

The fact remains that hackers today still manage to bypass many security devices—primarily due to the security industry's over-dependence on reactive threat coverage. The reactive nature of most of today's IPS solutions stems from reactive research methods. Many security vendors who do not possess an internal security research team are forced to rely solely on security intelligence available in the public domain, including attack exploit code posted on public security and hacker Web sites and vulnerability announcements released by software vendors. These IPS vendors are caught in a holding pattern, waiting for attack code to be publicized before they can update the protection in their IPS.

Collecting research on exploit tools commonly used by hackers and studying different vulnerabilities after they're announced does provide insight into the nature of attacks, but it should not be the sole form of security research supporting an IPS. Hackers generally don't post their exploit code until they've already used it to break into a system and by then it may be too late for the reactive security vendor and its customers. Plus, no security vendor should rely entirely on the hacker community for education about stopping attacks. Preemptive protection against Internet attacks requires proactive research above and beyond what is made public after an exploit or vulnerability is released.

Proactive Research

Sophisticated modern attacks move across the network at a rapid pace. If an IPS merely reacts to new threats after they appear, business systems will likely suffer negative impacts like corruption and downtime. Proactive security research is a pivotal requirement for preemptive protection. To conduct proactive security research, vendors must maintain a highly trusted and qualified team of security professionals who conduct primary research on the nature of vulnerabilities and attacks. Proactive research also requires extensive capital resources to acquire thousands of different types of hardware and software that can be studied for vulnerabilities and tested when new attacks appear. An IPS updated with proactive, vulnerability-based research focuses on the weak spots that are targets of attack, rather than the actual attack payload. Vulnerability-based protection is preemptive because it blocks any attack targeting the known weakness in the system—whether that attack has been seen before or represents a new variant or unknown threat.

IPS solutions powered by proactive research can also provide a viable alternative to the current patching crisis — offering what is called a "virtual patch." As software vulnerabilities continue to increase, so will the number of patches to install. Today, many organizations remain in an ongoing "triage" mode trying to determine which critical patches to apply first. If an IPS has the benefit of proactive security content updates focused on vulnerabilities, the system will protect those vulnerabilities during the window of exposure between vulnerability announcement and patch application—thus the term, virtual patch.

Example: ISS discovered the SSL PCT1[7] vulnerability which was publicly disclosed on April 13, 2004 in the Microsoft Security Bulletin MS04-011. Before Microsoft issued a patch and before the vulnerability was announced publicly, ISS silently incorporated protection for the vulnerability into its protection products in September of 2003. ISS' focus on primary, vulnerability-based research resulted in products that provided a virtual patch, or buffer of protection before the vulnerability was officially announced and a vendor-supplied patch issued. This level of protection is extremely difficult to provide without a dedicated group of in-house security experts.

ISS' X-Force® security research team conducts original, primary research on vulnerabilities and threats, which is applied to the ISS Intrusion Prevention System in the form of security content updates. The X-Force is credited with discovering and mitigating more major software vulnerabilities since 1998 than all other commercial security research organizations combined[8] (see Figure 6, next page).

[7] http://xforce.iss.net/xforce/alerts/id/168
[8] Source: Frost & Sullivan 2003, Internet Security Systems and public Web sites.

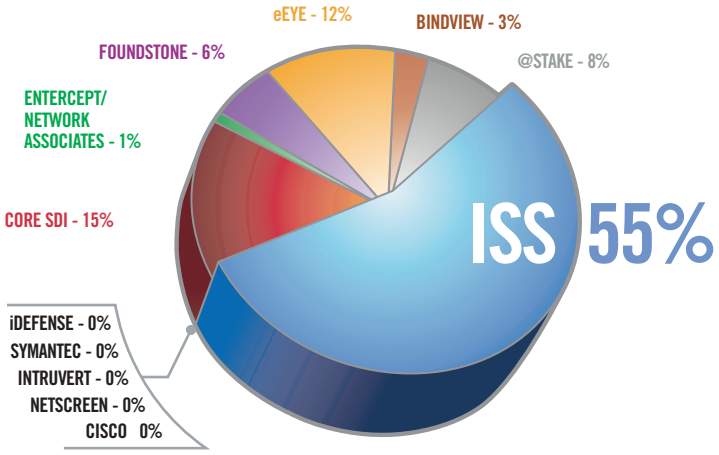

Figure 6: Preemptive Research (vulnerabilities discovered by research firms between 1998 and 2003)

Global Event Monitoring

Studying the ongoing threat landscape also contributes to security research. Some security vendors perform global event monitoring to understand where new security threats appear and how they spread. Centralized attack monitoring is typically conducted at security operation centers (SOCs), large data centers that monitor computer systems for signs of attack. Some SOCs resemble NASA-like mission control rooms with large screens showing the spread of Internet attacks on oversized world maps. The information gathered from security intelligence centers helps IPS vendors educate their customers about the nature of online threats. For global event monitoring to be a truly effective form of research, vendors must continuously poll thousands of computer systems for an adequate sampling of Internet traffic. Security vendors should also establish and maintain SOCs around the work so that data isn't skewed according to particular regions. Finally, SOCs need to be manned 24/7 by experienced security professionals that can analyze exactly what is happening, where it is happening and what it means for customers. When abnormal traffic is observed, global tracking and analysis centers can provide the most up-to-date threat information to alert and protect customers.

Information-Sharing Partnerships

Sharing security information with other entities will enhance the level of research collected. Security vendors that collaborate with law enforcement agencies, government entities and industry consortiums maintain an advantage in terms of protection for customers. Several different vertical industry groups have established Information Sharing and Analysis Centers (ISACs) to formalize the exchange of security information. ISACs serve as clearinghouses for breaking security news and events across corporate and industry boundaries. As with global event monitoring, information-sharing partnerships extend the overall awareness of emerging Internet threats, what they target and how they spread. This research only serves to improve protection from IPS solutions and can provide potentially critical advance warning to customers.

Example: In December 2002, the information technology (IT) sector chose ISS to maintain a forum for sharing information about network vulnerabilities and effective protection methods—resulting in the creation of the IT-ISAC. Through the IT-ISAC, the federal government in 2003 informed ISS that a new strand of the WebDAV virus had appeared. The government did not have any additional information at the time so ISS' X-Force located the threat, tracked it and within hours, provided the first analysis of the exploit to the entire IT sector through the IT-ISAC. Because ISS followed the evolution of the variant across the Internet, it was able to determine that its own customers were not at risk from this particular threat. In this case, information-sharing worked to inform ISS of the potential attack. In turn, ISS supplied the IT industry with an initial assessment of the threat.

Security research determines how well an IPS protects against Internet threats. As with performance and protection, no single research method is adequate. IPS vendors must couple reactive research with the proactive study of vulnerabilities and threats. Global event monitoring and information-sharing partnerships round out the research category, providing a macro view of the threat landscape.

Summary

The rules of preemptive protection are clear. An effective IPS solution relies on a solid foundation of performance, protection and research. Organizations evaluating IPS should look beyond point products or even so-called best-of-breed vendors that offer only a handful of protection, performance or research components. Without all three, the IPS can not stop threats *before* they impact business. High-performing IPS products keep pace with network traffic and avoid disrupting legitimate operations. With multiple protection methods for identifying and analyzing traffic, the IPS maintains accuracy. Finally, a combination of research, including a focus on vulnerabilities as well as threats, must be infused into the IPS to enable true preemptive protection.

About the Author

Christopher Rouland is the Chief Technology Officer (CTO) for Internet Security Systems, Inc. Rouland is responsible for guiding the company's overall technology strategy with a commitment to developing products and services that preemptively protect organizations from cyber threats.

Prior to his appointment to Chief Technology Officer, Rouland served as the vice president of X-Force R&D and was instrumental in building and growing the X-Force organization. The X-Force is a group of security experts dedicated to understanding, documenting and coding new vulnerability checks and tests, attack signatures and solutions to global security issues. The X-Force also maintains the industry's most comprehensive online knowledge base for rapid look up of information on thousands of risks and threats. Additionally, the X-Force team supports the U.S. Department of Homeland Security with daily briefings to update and advise the U.S. government on the current health of the Internet, as well as new types and sources of attacks.

Since joining Internet Security Systems in 1998, Rouland is a frequent spokesperson for national media outlets such as CNN, Fox News, and the Associated Press, as well as most technical print publications. Rouland was credited with the discovery and naming of the SQL Slammer worm, and initiated the White House press conference to alert the world media to the secondary damaging impacts of the Code Red worm.

Rouland has 14 years of valuable experience in information technology. Prior to joining Internet Security Systems, Rouland began his career deploying most of the first Internet connections for UUNet in the late 80's and early 90's, including sites like 'senate.gov'. Rouland has contracted his security expertise to several three letter government agencies. He has also held positions as a software developer, network architect, and vice president of Distributed Technology for Lehman Brothers, Inc.

Special thanks to the following contributors:

Dan Ingevaldson, Mike Lynn, Freddy Mangum, Paul Palmer and Chris Simmons

INTERNET | SECURITY | SYSTEMS®
Ahead of the threat.

6303 BARFIELD ROAD | ATLANTA, GA 30328 | 800.776.2362 | FAX 404.236.2626