

# Using an Intrusion Prevention System as Part of a Layered Security Approach

By Shawn Conaway

## Securing the Enterprise

Intrusion prevention is the art of keeping your network free from attack. It is a holistic approach to security that includes hardening computer systems, deploying utility servers like email gateways and antivirus servers, and, of course, deploying intrusion prevention systems (IPS).

An IPS is, at its most basic, a firewall that has the ability to respond to perceived attacks based on criteria and behaviors identified by the IPS vendor and honed by the IPS administrator. When an attack is identified, the IPS has the ability to drop, reset, redirect traffic, or simply log the traffic and pass it through.

**Intrusion prevention is the art of keeping your network free from attack. It is a holistic approach to security that includes hardening computer systems, deploying utility servers like email gateways and antivirus servers, and, of course, deploying intrusion prevention systems (IPS).**

Employing an IPS is a financial investment in the security of your network and your networked systems. The cost of a disrupted network is far more costly than the implementation costs for an IPS in all but the smallest companies. The IPS prevents hackers from easily compromising internal resources in addition to preventing the overload of a company network from unnecessary traffic.

The bulk of the work that an IPS does is preventing virus and worm attacks. An IPS also prevents directed common attacks from amateur hackers running tool kits and designed by more skilled hackers. These 'script kiddies' run attacks with tool kits and signatures that will be prevented by a good IPS. An IPS is less capable of blocking unknown 'zero-day' attacks because the signature of the attack is unknown, although behavioral analysis may be able to thwart the attack.

## Analyzing the Traffic

An adequate IPS has to be robust enough to handle the volume of traffic passing through it, normalize the packets, statefully analyze the packets, and then act upon the information to identify an attack. Packet scrubbing and packet analysis has to occur at line speeds so the rest of the network is not disrupted.

Packet normalization, or packet scrubbing, is the way data packets are rearranged so they cannot be misinterpreted. When packets are normalized, they are reassembled so that the packets contain complete datagrams with a standardized length and configuration. Hackers purposely send malformed packets with illegal TCP flag combinations in an attempt to exploit weaknesses in routers and systems whose operating systems have not been kept up to date. Since packet fragments are reassembled when they reach the intended target, malevolent packets that are not normalized and filtered at the IPS level are still able to exploit vulnerabilities on the target of the attack.

Other harmful traffic that can be stopped at the IPS includes duplicate overlapping packet fragments and packets that have a Time to Live (TTL) stamp that is too short. Packets with overlapping datagrams can be misinterpreted by the target if duplicates are received that are out of order. An exceedingly eloquent attack is if these overlapping packets also have a TTL stamp that is configured to be too short for the good set of packets to reach its destination and the right

length for the malicious packets. These seemingly legitimate paired packet fragments appear like benign data. The result is that only the bad packets reach their destination.

Stateful pattern matching involves analyzing the chronological order of the packets in a TCP stream to identify known attack patterns. Once the packets are assembled, the IPS can analyze the packets to see if they have any pattern matching signatures of known exploits. An IPS that does not perform stateful pattern matching will not be able to catch attacks where a known signature is split across two packets. Stateful pattern matching also avoids the high number of false positives that would normally be generated by an IPS that just does pattern matching.

In addition to stateful pattern matching and normalization, an IPS can perform other forms of detective work, searching for aberrations like protocol anomalies, traffic anomalies, and layer two attacks.

**Protocol anomalies** are detected when a protocol does not do what it normally should do, such as an HTTP response without an HTTP request, or a protocol that has a non-standard value.

**Traffic anomalies** are attacks that span a number of separate connections for a set amount of time with a predefined number of packets. Attacks start and stop repeatedly in an attempt to avoid triggering the IPS by exceeding the default threshold. An example of a traffic anomaly

is a recurring ICMP flood. An ICMP flood is an ICMP ping attack that causes a system to respond to so many ICMP echo requests that it can no longer effectively handle replying to any other requests. The default threshold on many IPSes is set at 1,000 ICMP packets per second.

**Layer 2 attacks**—A man-in-the-middle attack is a type of layer 2 attack where the hacker is able to intercept, read and insert unwanted information into the stream. An example is if a third party is listening in on an exchange of public keys. When one party requests a public key from a second party, the third party can respond to the request with his own key. From then on, the third party is able to control the supposedly secured connection and receive any secured data.

## **How Can an IPS Identify an Attack?**

There are four common methods that an IPS can use to identify attacks. They are rated-based, signature matching, policy matching, and behavioral.

**Rate-based attacks** are attacks that flood the attack target with more traffic than it can handle. In especially effective attacks, a rate-based attack can also overwhelm the network. Denial of service (DoS) and distributed denial of service (DDoS) attacks are rate-based IPS attacks that send overwhelming amounts of data to the attack target. Smurf attacks and Fraggle attacks are common types of DoS attacks.

In a smurf attack, an ICMP echo request (a.k.a. a 'ping') is sent to IP broadcast addresses. The packet is spoofed so that the source IP address of the packet is the address of the attack target. When an attacker sends out the spoofed packet to unsuspecting compromised zombie systems, also known collectively as a botnet, they respond by echoing packets back at the IP address of the victim. The zombies are systems that have been compromised without their knowledge by opening malicious spam mail or by running Trojans with intriguing names like BritneySpearsNaked.exe. In this way, a hacker or script-kiddie is able to amplify the attack because each packet the attacker sends is answered by hundreds of packets each time a single attack packet is sent out by the zombies. The Kournikova virus propagated in a similar way, tempting email recipient for naked pictures of the tennis star Anna Kournikova.

A SYN Attack is another rate-based attack whereby an attacker sends a flood of connection requests, or SYN requests, to an intended victim faster than they can be processed. Each SYN request packet contains a randomized spoofed IP address. The victim replies with a SYN ACK, an acknowledgement of the SYN request, then waits for the system that generated the SYN request to reply with an ACK, or an acknowledgement that the session is being created. However, the ACK never comes in this attack. The result is that the victim's connection table eventually fills, preventing the victim from being able to reply to any session request.

SYN attacks are sometimes able to crash servers by overflowing their memory buffers. Commonly, up-to-date systems drop the incomplete session requests from their buffers after a predetermined amount of time.

**Signature matching attacks** are attacks that can be identified by the data coming across the network. For example, the I Blaster virus has the string "Billy Gates why do you make this possible? Stop

making money and fix your software!!" which converts into a specific binary signature with a specific length and arrangement. IPS block the worm when it tries to traverse the network by identifying the signature of the worm.

**Policy matching** is when the IPS prevents traffic based on what is acceptable use. For instance, an organization may decide that running peer-to-peer application like Napster or Kazaa is a security threat and that peer-to-peer connections are not allowed. Also, the company might decide to disallow MIME traffic that is not going through approved email gateway. Commonly, instant messaging is blocked to prevent vulnerabilities such as the ability for file sharing and remote desktop in AOL's AIM Pro instant messenger.

**Behavioral patterns**—An IPS is able to stop unusual behavioral patterns by responding to anomalous traffic and traffic patterns. These anomalies are based upon information detailing how specific protocols work as well as on historical information that baselines traffic patterns on a company's network. For instance, if a baseline shows that a specific network has relatively high TCP traffic and low UDP traffic, the IPS will be triggered if UDP traffic spikes. In another case, an IPS may see an unusual increase in an incoming SYN request, but not enough traffic to trigger a response. The IPS will trigger if it is able to correlate the increased SYN attack traffic with a subsequent console session directed outwards to the attacker.

This ability to react to the behavior of an attack allows the IPS to stop the attack without knowing the attack signature. The drawback is that there are many false positives. Another problem is that the IPS has is that it can block normal traffic that does not meet the baseline parameters, such as when an unusually large amount of traffic is generated by employees during the yearly open enrollment for health benefits.

## **Types of IPS**

There are many flavors of IPS, but the basic function of each is to detect and stop attacks by either dropping sessions, resetting sessions, blocking packets, or proxying traffic. An IPS can be either hardware, software, or a combination hardware/software solution. The five main

**There are many flavors of IPS, but the basic function of each is to detect and stop attacks by either dropping sessions, resetting sessions, blocking packets, or proxying traffic. An IPS can be either hardware, software, or a combination hardware/software solution.**

types of IPSes are in-line detection, layer seven switches, deceptive systems, application firewalls, and hybrid switches.

An **in-line detection system** is a direct barrier between your network and the rest of the world. It is commonly placed at the outer edge of the network perimeter in front of the firewall. An in-line IPS can also be installed on the interior network for finer tuned prevention.

An in-line IPS works like a combination firewall and layer-two bridge. When good data comes across the line, the in-line IPS passes the packets through to the rest of the network. When the packets contain any known vulnerabilities, the firewall blocks the packet or drops the connection.

**Layer seven switches** are placed in front of the firewall, often acting as a load-balancer for web-based applications. The switch inspects HTTP, SMTP and DNS requests to identify where to direct the traffic. The switch is also able to read the intended URL of the HTTP request to route the traffic appropriately.

**Deception systems and honeypots**—Honeypots are systems that are online for the sole purpose of being attacked. They are good detectors for both virus attacks and hackers. Honeypots are used to deceive attackers into thinking that they are attacking a valid internal system. When an attacker finds a honeypot through port scanning and then tries to attack the system, the deception system returns a packet marked with misleading data. When the attacker attempts to use that data in his next attack, subsequent traffic is blocked.

An **application firewall** is software that gets installed on each server being monitored. The software monitors the application(s) on a server, watching the API calls the applications make, the memory utilization, and the way software interacts with the system. The application firewall has to have a behavioral profile built for each system it is installed on so that the firewall knows how to differentiate between proper use and malicious use. Application firewalls are distinct from application layer firewalls or web application firewalls, which are also known as layer seven firewalls.

A variant of the application firewall is the personal application firewall, of which Windows Firewall and ZoneAlarm are the most widely known. They block access into and out of a client system and block certain types of application behavior, such as instant messaging applications. There is no profiling done to configure the application, although the personal firewall may notify the user the first time it blocks a specific type of traffic so that the user can make the appropriate changes.

A **hybrid switch** is a combination of an application firewall and a layer seven switch. The switch is a hardware front-end for the protected server. Unlike a layer seven switch that is configured to generally protect the entire network, the hybrid switch is configured to protect one or more similarly configured servers. The switch is set up initially in learning mode so that a policy can be created.

## The Rest of the Story

---

While focusing on a robust Intrusion Prevention System is crucial to network security, it is only one facet of a defense-in-depth security posture. Securing the network from external attack while leaving internal systems unpatched and unsecured is like locking all the windows to a house and leaving the front door open. A company that uses a high-end IPS without having any internal security controls is vulnerable to inadvertent virus attacks by any employee connecting from home over VPN or by a consultant or with an infected laptop.

The most basic approach to preventing an intrusion is to employ systems hardening techniques like disabling unneeded services, applying regular security patches, using antivirus servers or appliances, and keeping passwords secure by enforcing complexity standards and frequency of changing passwords. Hardening techniques can be found in the NIST Security Configuration Checklists Repository at [www.nist.gov](http://www.nist.gov).

**System hardening** is mostly passive because it is normally done when the server is built or when a new application is installed. Applying application and operating system patches is the active component

of system hardening. A system can be considered as hardened one day, then vulnerable the next day simply by someone identifying a vulnerability.

An **in-cloud email gateway** is a way to prevent spam, virus attacks and phishing exploits from reaching the company network by stopping the traffic at the Internet service provider. Using the gateway offers better protection than an internal solution because attacks never make it to the company. Also, the company saves on bandwidth charges and storage space because the traffic is blocked.

**Antivirus and spam filtering** services can be provided by filtering spam and viruses before they reach the internal company network. If a known worm attack gets past the ISP, it won't be able to bring down the company network if systems have up-to-date virus definitions.

Reviewing **event logs** is probably the most neglected intrusion prevention and intrusion detection task because of the brain-numbing repetitive nature of the task. However, reviewing event logs may be the most likely way to identify when a system is being attacked. Companies can automate review of the logs using Microsoft Operations Manager to forward unknown events to Operations or a monitoring group. Effort is still involved in tuning the monitoring application so that the application does not raise alerts for too many false positives.

## The Cost of Inaction

---

Implementing an IPS can be risky because it has the potential to slow down network traffic or to set up a self-imposed denial of service attack by blocking legitimate traffic. Not implementing an IPS is even riskier.

**All war is deception.**  
-Sun Tzu

According to the 2006 CSI/FBI Computer Crime and Security Survey, the total losses for security incidents per respondent was \$167,713. The Federal trade commission estimated higher, stating that in 2003 there was \$48 billion dollars in institutional losses and \$5 billion in losses to individuals by security incidents.

In marketing material generated by McAfee, Inc. the company calculated that companies implementing an intrusion prevention system realized a 19-to-1 return on investment. They based their ROI on an average investment of \$200,000 for an intrusion prevention system compared with avoiding the cost of an outbreak. They estimated that the slammer outbreak alone cost the average company \$240,000.

Implementing an IPS is a wise investment when considering that, if the average loss is around \$167,713 and an IPS system costs \$200,000, the breakeven would be just over one year. An IPS appears to be an expense that companies can't afford to avoid. 🐼

---

NaSPA member Shawn Conaway is a Systems Administrator for a Fortune 100 retailer. He currently holds the Microsoft Certified Systems Engineer, Citrix Certified Administrator, and A+ certifications. Send questions or comments to [s.conaway@naspa.com](mailto:s.conaway@naspa.com).