



Network Intrusion Prevention Systems - Why “Always On” Stateful Inspection and Deep Packet Analysis are Essential to Deliver Non-Stop Protection

By Sanjay Raja
Top Layer Networks, Inc.

With the explosion in Internet connectivity and the mainstream use of broadband and mobile technologies, there has been a huge increase in the number of computer systems and storage devices connected to the public network. With an ever-increasing reliance on computing infrastructure, we find that our critical IT assets, confidential data and intellectual property are more susceptible to cyber attack than ever before. In response to the changing threat landscape, Network Intrusion Prevention Systems (IPS) were developed to provide advanced protection beyond that offered by firewalls and Intrusion Detection Systems (IDS). Out of this need appeared many new products claiming to be IPS Solutions, however, the vast majority are unable to handle these new threats despite their data sheet claims. This white paper highlights why “always-on” Stateful Inspection and Deep Packet Analysis are necessary capabilities when selecting an IPS solution.

Table of Contents

1	EXECUTIVE SUMMARY	3
2	INTRODUCTION.....	3
3	WHAT IS STATEFUL INSPECTION?.....	5
4	SOME IPS VENDORS CLAIM THEIR PRODUCTS USE STATEFUL INSPECTION	5
	IS IT AN IDS OR AN IPS?.....	5
	WHAT'S THE REAL PERFORMANCE?	6
5	PROTECTION BUILT ON TOP OF STATEFUL INSPECTION.....	7
6	STEPPING BEYOND STATEFUL INSPECTION: DEEP PACKET ANALYSIS	8
7	TOP LAYER'S ATTACK MITIGATOR IPS 5500	9
8	CONCLUSION.....	12

1 Executive Summary

Since the early days of firewalls, no other technologies have revolutionized security more than Stateful Inspection and Deep Packet Analysis. The ability to understand context information of new incoming packets in relation to previous packets from the same session and inspect its contents for malicious activity has become an expected requirement in offering a network security solution. However, with the advent of a new generation of cyber threats, ranging from hybrid exploits of critical application vulnerabilities, to brute force Distributed Denial of Service (DDoS) attacks carefully designed to overwhelm firewalls and their intended targets, traditional security products are failing to maintain total network and application integrity. Since 2002, a new technology has emerged to address these complex threats - Network Intrusion Prevention Systems (IPS). The diverse spectrum of technologies used in IPS products to perform this advanced protection can be confusing in the eyes of organizations looking to deploy IPS in their network. One thing is certain, only IPS products that use “always-on” Stateful Inspection in conjunction with Deep Packet Analysis should be trusted to provide the most advanced protection against the new threat landscape. Most IPS solutions do not use these technologies or only invoke them under certain conditions. In either case, the risk associated with even a small amount of malicious traffic evading these solutions and arriving at their designated target is substantial, considering the crippling impact to mission critical resources.

2 Introduction

Most organizations are still reliant on security based on 1990s technology to defend against today's complex network and application-based attacks and brute force DDoS attacks. Organizations that rely on firewalls for access control and intrusion detection systems for monitoring network traffic are at high risk of a successful attack from these new threats.

Firewalls have historically been deployed at the network perimeter as the first line of defense against unwanted intruders. Some firewall vendors have tried to address these new attack types by simply augmenting their appliances with software. However, this creates a serious performance bottleneck that at best, increases latency and worse can lead to performance bottlenecks that could allow bad traffic to slip through while dropping good traffic. Today's firewalls are not architected to handle large volumes of DDoS traffic. Most gigabit firewalls that claim DDoS protection are completely crippled by less than 100Mbit/s worth of DDoS traffic. Add to this the implications of attacks being launched from inside by accident or by disgruntled employees or contractors by introducing a Trojan, virus, or worm that was acquired while surfing the Internet at home. A simple SYN Flood attack generated externally using widely available

tools or by a single internally compromised system running a zombie program can overrun a firewall in minutes, resulting in lost Internet connectivity, and therefore lost productivity, a source of great concern for any organization that uses the Internet to conduct business. For organizations where the Internet is their primary source of income and their customers have choices on where to buy products or services, these new threats can put them out of business for days. Organizations are also concerned with the legal repercussions regarding privacy, unavailability, or being the unknowing source of attacks. There have been many examples of businesses under such attacks whose Internet connection was pulled by their Internet Service Provider in order to keep the ISP's other customers connected.

Intrusion Detection Systems (IDS) may be effective at detecting suspicious activity, but do not provide adequate protection against attacks. Worm attacks like Slammer and Blaster spread so fast that by the time an alert was generated and the IT administrator has time to peruse reams of alerts and reports, the damage was done.

The need for advanced protection systems has caused many technology vendors to change their marketing to make their products, traditionally IDS, but sometimes Firewalls, look and feel like IPS solutions. It appears that every IDS or security vendor has developed an IPS solution, seemingly overnight. However, there are key architectural features necessary for an IPS solution to be capable of offering proper protection under different network environments and user traffic conditions. Software modifications or minor hardware modifications to traditionally offline IDS systems or existing firewall architectures cannot provide the same capabilities as a purpose-built IPS Solution. In selecting the right solution, customers are left with lots of marketing claims and no easy way to evaluate these offerings until painfully struck by an attack.

The purpose of this white paper is to explain why "Always On" Stateful Inspection is an essential element for any IPS to provide advanced security protection against these new hybrid and brute force attacks. We will also help you understand what an IPS vendor really means when they claim their product performs Stateful inspection and why you need to be wary of half truths that could lead you down a path where you will not be as sufficiently protected as you had been led to believe.

In addition, this paper explains why an IPS product must be able to perform high speed Deep Packet Inspection in order to provide the most robust security against these new threats.

Finally, this paper describes how Top Layer's approach to network intrusion prevention utilizes the power of always on Stateful inspection and combines it with deep packet inspection to provide

the most reliable, high performance non-stop IPS product that you can safely rely on to protect your organization.

3 What is Stateful Inspection?

Every operating system implementation has security leaks that are known to hackers throughout the world. In the 1990's, Stateful Inspection became the industry standard for Firewalls to address protect against undesired access and other malicious behavior including protection against low-level DoS attacks. As well as examining header information, Stateful Inspection can examine the contents of a packet (up through the application layer) to determine more context about the packet beyond its source and destination information. In addition, Stateful Inspection monitors the state of a connection and compiles historic information in a state table. As a result, dynamic filtering decisions can be expanded beyond administrator-defined rules that simply block known IP addresses or TCP ports (as in static packet filtering) to take into account the context of a packet that has been established by packets that passed through the firewall earlier.

For the same reason that most firewalls use Stateful Inspection as a mechanism to monitor traffic anomalies, an IPS device must utilize Stateful Inspection to perform advanced protection against new types of attack as well as defend against the growing frequency and scale of DDoS attacks. Always-on Stateful Inspection introduces the most advanced application security by incorporating full-time session state awareness and extra security context information, which is stored and updated dynamically in order to more quickly mitigate threats. This form of Stateful Inspection provides cumulative data against which subsequent communication attempts can be evaluated and acted upon in real time. It also delivers the ability to create virtual session information for tracking connectionless protocols (e.g. Microsoft RPC and UDP-based applications).

4 Some IPS Vendors Claim Their Products Use Stateful Inspection

Is it an IDS or an IPS?

It is well known that many "IDS-based" IPS systems are capable of some Stateful inspection while operating in an offline IDS mode. IDS-based IPS solutions were spawned from the IDS vendors that had their roots firmly planted in their ability to alert, report, and correlate attacks. Despite poor reviews and experience with IDS systems in generating huge volumes of alerts and false positives, the concept of taking these offline devices and putting them inline and allowing them to block attacks based, primarily on signature or pattern matching techniques is the approach many vendors commonly use. In fact a portion of these vendors utilize a lesser form of Stateful Inspection to complete simple pattern matching (a.k.a. signature matching) on packets to establish whether the packet contains a known exploit. As a result, these IPS vendors will claim

their products have Stateful inspection capabilities. However, as soon as these IPS products are deployed in-line to do proactive blocking rather than simple off line detection, many of these devices lose their Stateful inspection capabilities and simply inspect packets coming in, without maintaining full context across the session. Typically, if these devices try to maintain “always on” state, the performance and latency fall away dramatically. This highlights the fact that the performance requirements for an offline IDS device are dramatically different from an inline IPS. It is well known that Stateful Inspection capabilities allow an IPS device to maintain more information on packet flows and are therefore able to make more intelligent decisions in terms of detecting malicious traffic.

What’s the Real Performance?

In some cases, an IPS device may turn on Stateful inspection as soon as it detects an attack so that the device can more closely monitor packet flows and relevant context on future transmissions. This is typically a short-term burst of increased protection, which, after a while, reverts back to the stateless mode. The advantage this provides to those IPS vendors is that they are able to quote much higher performance numbers in their data sheets based on passing legitimate traffic through the device without performing Stateful inspection. As previously stated, the moment these devices go into Stateful mode, their performance drops off dramatically and there is a high risk that legitimate packets will be dropped and that the IPS device becomes a performance bottleneck in the network.

Having an IPS that offers part-time Stateful inspection creates a real challenge to network security managers. For instance, new hybrid attacks that split the malicious code across multiple packets are more likely to be missed by this type of IPS. Another problem is with asymmetrical network topologies where packets can come and go out on different network segments in order to improve redundancy and reliability of the network. In an asymmetric network, almost all IPS are incapable of maintaining any state for all transactions. In that case, the IPS is most likely doing simple packet filtering based on address information only causing most attacks to not be identified. This is protection is less than the capability of any stateful firewall.

To get around the challenge of performance bottlenecks with always-on Stateful inspection, an IPS vendor must invest heavily in developing special-purpose, yet programmable hardware that is seamlessly integrated together to minimize latency (ideally switch-like microsecond latency) concerns while passing traffic under load or attack. Only the most advanced hardware architecture allows for excellent protection at all times with minimal degradation in performance.

5 Protection Built On Top of Stateful Inspection

It is important to understand why Stateful inspection is important in terms of providing relevant context of packets, which in turn allows more intelligent decisions to be made on attack traffic versus legitimate traffic. For example, many rate-based attacks utilize tens of thousands of computers to send legitimate traffic to their target. There are no signatures to be matched and there are no protocol anomalies that warrant blocking those attacks. The only sure way to block that type of attack is to keep state, and therefore context, on a large number of IP addresses. Here in lies another consideration in understanding what a vendor means when they say their device is Stateful. In this example, the larger the address table, the better context can be maintained over more IP addresses and therefore better protection can be provided while allowing legitimate transactions to continue flowing.

The first generation firewall used simple packet filtering to allow or deny the passage of traffic between networks based on the header information in each data packet. In this scenario, traffic is filtered based on rules that are established to allow or deny the flow of network traffic, primarily based on the source, destination, and port (service or application).

The limitations of packet filtering were seen when dealing with more complex applications, like "NetMeeting", which uses H.323 as its underlying protocol. In this case, the user should leave a large number of ports open, since NetMeeting opens/closes them dynamically on demand for audio/video conferencing. For a simple packet filtering hardware firewall this is a major problem, since no context or awareness exists for the application on the intranet side, which needs to access these ports. Therefore, a huge hole exists in the firewall since the administrator must open all these ports beforehand. The situation is not the same with a Stateful inspection firewall, which knows which application listens to which port.

Stateful inspection uses information that utilizes layers 3-7 of the OSI model (network layer and upwards), to obtain knowledge such as allowing traffic sessions to pass over specific ports *dynamically*. By combining information from various layers (transport, session, and network), an IPS is able to better understand the protocol that it is inspecting and make more intelligent decisions on whether packets are legitimate or not. The following looks at the interaction of always on Stateful inspection in the context of an IPS and protocol types:

Advanced Internet Protocols

User-level applications such as FTP and the Web can create complex patterns of network traffic, which requires that the IPS analyze groups of network connection "states". A central cache within the IPS keeps track of the state information associated with each network connections. Each packet received by the IPS device is analyzed against the state table to assess whether or not it will be allowed to pass through to its destination or be rejected.

TCP Connections

The first packet of any new connection has its SYN flag set and its ACK flag cleared, referred to as "initiation" packets. All packets that do not have this flag structure are called "subsequent" packets, since they represent data, which occurs later in the TCP stream. These subsequent packets are bi-directional and Stateful Inspection provides the capability of following all the flows in this bi-directional conversation to make sure they are legitimate.

UDP/ICMP

User Datagram Protocol ("UDP") and ICMP do not themselves contain any connection information (such as sequence numbers). However they both contain an IP address pair and UDP contains port pairs while ICMP packets have type and code information. All of this data can be analyzed in order to build "virtual connections" in the cache. For instance, a cache entry will be created by any UDP packet, which originates on the LAN. Its IP address and port pairs will be stored. For a short period of time, UDP packets from the WAN which have matching IP and UDP information will be allowed back in through the firewall.

An analogous situation exists for ICMP, except with some additional restrictions. Specifically, only outgoing echoes will allow incoming echo replies, outgoing address mask requests will allow incoming address mask replies, and outgoing timestamp requests will allow incoming timestamp replies. No other ICMP packets are allowed in through the firewall, simply because they are too dangerous and contain too little tracking information. For instance, ICMP redirect packets are never allowed in, since they could be used to reroute traffic through attacking machines.

Application Protocols

Some higher layer protocols (such as FTP and RealAudio) utilize multiple simultaneous network connections. In general terms, they usually have a "control connection", which is used for sending commands between endpoints, and "data connections", which are used for transmitting bulk information. Consider the FTP protocol. A user on the LAN opens a control connection to a server on the Internet and requests a file. At this point, the remote server will open a data connection from the Internet. For FTP to work properly, this connection must be allowed to pass through even though a connection from the Internet would normally be rejected. In order to achieve this, a device inspects the application-level FTP data. Specifically, it searches for outgoing "PORT" commands, and subsequently adds a cache entry for the anticipated data connection. This can be done safely, since the PORT command contains address and port information, which can be used to uniquely identify the connection. Any protocol, which operates in this way, must be supported on a case-by-case basis.

6 Stepping Beyond Stateful Inspection: Deep Packet Analysis

The term "Deep Packet Inspection" describes a variety of features that enable the IPS to scour individual data packets or streams of packets to spot malicious code or other anomalies that might be part of an attack. Deep Packet Inspection directs, persists, filters and logs IP-based applications, including Web traffic, based on content encapsulated in a packet's header or payload.

As already discussed, always-on Stateful inspection features enabled IPS devices to move beyond just filtering traffic based on the information contained in data packet headers to monitor active connections. Deep packet inspection allows the IPS to dig even deeper into traffic flows to spot hidden attacks on targets like Web, e-mail, and DNS servers. By way of example, Deep Packet Inspection lets the IPS device look deep into the content of a TCP or UDP flow for a complete view. This is accomplished by reassembling IP datagrams, TCP datastreams and UDP packets as they flow through the device to view the entire application content and then act on it according to user-defined policies.

To perform effective Deep Packet Inspection, it is vital that the protocol traffic be reordered in the form it was originally transmitted. Here in lies the interaction of Stateful inspection and Deep Packet Inspection. To accurately reorder the transmission, the received packets must be compared to the state table to understand where they are in the transmission flow. To bridge the gap between receiving packets out of order and/or in fragments and the requirement to perform Deep Packet Inspection, the IPS device requires a “reorder engine” to take all of the packets and put them in original transmission order. Once put in the proper sequence, the IPS can fully inspect the content of each transmission to confirm whether it is legitimate or not.

There are two major attacks that employ fragmentation: retransmitted fragments with invalid or spoofed data, and overlapping or adjacent fragments that can exploit systems. Some vendor implementations of the TCP/IP IP fragmentation re-assembly code do not properly handle overlapping or retransmitted IP fragments. Since exploits can be broken up in order to evade detection, it is critical to perform fragmentation reassembly prior to inspection. However, fragment attacks are much more difficult to protect against due to the intense processing required in coalescing and inspecting all fragments as part of a packet stream. The IPS 5500 is the only solution that can perform this function with minimal latency based on dedicated processing for handling fragments that maintains state on the fragment number and offset. In addition, security solutions that do not have the proper processing capability are susceptible to DoS attacks using the "teardrop," "teardrop2," "boink," and "bonk" attacks. This attack fills up the buffer space of any device claiming to be 'fragment smart', but lacking the proper processing capability. Only the IPS 5500 is architected properly to process and perform full stateful inspection on high-volumes of fragments, offering non-stop protection from malicious content and DOS attacks. A side benefit of the IPS 5500's ability to coalesce high volumes of fragmented traffic is that it can also improve the performance of downstream network devices by offering already sequenced packet streams.

The ability to perform deep packet inspection means that an IPS device is able to provide higher-level logic functions based upon the content of an individual data packet, or a stream of fragmented data. These higher level functions would include the ability correlate the data content against a series of rules. Through this process an IPS can better protect against content that may be associated directly with threats such as worms, Trojans, or other exploits and threats.

7 Top Layer's Attack Mitigator IPS 5500

The IPS 5500 is the only product that has seamlessly integrated “always on” Stateful inspection and Deep Packet Inspection to provide a superior defense against application attacks. Top Layer's Attack Mitigator IPS 5500 is the worlds first Non-Stop Intrusion Prevention System that delivers dynamic, real-time proactive defense for network and application-based attacks.

The Attack Mitigator stops intrusions, automatically blocking internal and external cyber threats and other bad traffic before they can degrade and/or damage critical IT infrastructures. The Attack Mitigator is the only IPS product to have been awarded the NSS “Approved” status for both content-based and rate-based attacks (visit http://www.toplayer.com/content/news/nss_approved.jsp for more details on this stringent test).

Protection For Content-Based Attacks

Based on integrated “always on” Stateful inspection and Deep Packet Inspection, the IPS 5500 provides the strongest defense against content-based attacks. Most application attacks, by their very nature, contain malicious code that would cause the attack packet stream to violate acceptable protocol usage. Simply put, since it is clearly known whether or not a packet stream is compliant with acceptable protocol usage, any data stream that violates that acceptable usage or user-defined usage can be deemed malicious and can be consequently blocked.

Example: Blaster and Lovsan were caused by sending a malformed DCOM activation packet with a very large “servername” entry in the UNC path parameter of a remote request over the Microsoft RPC protocol. An attack signature might work for known variants, but small changes to the attack would pass through the IPS and infect the network and assets. By using the Attack Mitigator’s Protocol Validation Filters, all variants of the exploit were stopped because the packets were identified as not conforming to an acceptable protocol use.

The Protocol Validation Filters in the Attack Mitigator operate by invoking several simultaneous detection mechanisms that detect whether a packet stream should be allowed to pass through or not. Using these mechanisms, prior to performing signature-based detection, provides the most comprehensive IPS protection while at the same time eliminating the concerns regarding false positives, the bane of the network administrator.

- **Fragment Reorder Engine** – Dedicated ASIC to receive and reorder packets that make up an entire session which are then forwarded to the Deep Packet Inspection Engine where key checks are performed. This engine avoids common hacker evasion techniques and improves IPS performance and downstream network operation.
- **Deep Packet Inspection Engine** - Dedicated FPGA that performs checks against acceptable protocol use. If the entire transmission violates acceptable use, the transmission is blocked, reported on and sent to a discard port for forensics.
- **Session Aware Application Inspection Engine** – To prevent attackers from using evasion techniques such as slow attacks, the Session Aware Application Inspection Engine maintains real-time intelligence on over two million IP addresses. If malicious transmissions from any of these IP addresses are seen, any future traffic originating from them will automatically be closely scrutinized.

- **Security Profiles** – Security policies developed for application and network behavior that define expected and compliant operation. These profiles are developed via a white-list approach – defining what is “good” and stopping everything else. This protects critical resources from design vulnerabilities and implementation vulnerabilities.

Because the Protocol Validation Filters do their work before signature-matching mechanisms are employed against specific exploits, they effectively act as a virtual software patch for unprotected vulnerabilities, including those susceptible to buffer overflow attacks. This is important in the context that we continue to see a narrowing of the time between discovering a vulnerability and the time an exploit is launched, thus reinforcing the literal meaning of the term “Zero-Day vulnerability”. IPS products that do not perform deep packet inspection on entire sessions (with the packets reordered to their original form) do provide the same level of protection from next generation attacks such as hybrid and zero-day attacks that are now commonly seen.

Protocol Validation Filters only block packet streams that breach acceptable application protocol usage rules, therefore legitimate traffic is never blocked, a key differentiator to those IPS solutions that use attack signatures as a first line of defense. The following example shows how the Attack Mitigator is able to distinguish between a legitimate SMTP transmission and one that contains a malicious payload.

Consider an SMTP connection between a mail client and a server. The client opens the connection with the typical TCP three-way handshake. For most IPS devices, the establishment of the connection and the monitoring of it for when connection is terminated are sufficient. However, most IPS devices do not look further up the protocol stack for events that may be considered "out-of-bounds" in the application. With an IPS that is capable of Deep Packet Inspection the IPS can look at the SMTP protocol and monitor it for any attacks. When a client establishes the SMTP connection by following the RFC defined protocol steps of issuing a HELO, waiting for the response by the mail server. The client may then issue a variety of commands include sending e-mail by specifying the SMTP command **MAIL FROM:** If a client tries to issue a **VERFY** command, the IPS monitoring the communication between the client and the mail server can be configured to respond detection of the **VERFY** command by disallowing it. The client may also try to exploit the sendmail address token overflow (discussed in the CERT bulletin CA-2003-12) in order to gain shell access to the server. The Attack Mitigator, because it is capable of Deep Packet Inspection, is able to identify the exploit attempt and deny the connection. Additionally, it may deny the connection from the client altogether.

Protection For Rate-Based Protection

Intelligent Rate-Based Filters allow for connection and application rate limiting, which in conjunction with Patented DDoS Algorithms, provide industry leading DDoS protection capabilities. The IP address table that is used to keep state on IP address can store up to five million addresses when the Attack Mitigator is defending against a DDoS attack. During an attack, the Attack Mitigator will continue to pass good traffic with the lowest latency of any other inline IPS solution.

8 Conclusion

Only Intrusion Prevention System products that utilize both “always-on” Stateful Inspection and Deep Packet Analysis are able to provide the necessary level of protection required to defend against new advanced cyber threats. Top Layer’s Attack Mitigator has taken Stateful Inspection to the next new level by providing the only multi-gigabit solution that offers “always-on” Stateful Inspection and Deep Packet Analysis. The IPS 5500 implements the largest state table available in any IPS product to protect against threats offering the most solid protection available in the market today. The integration of the various protection mechanisms with a powerful, but flexible hardware architecture means the Attack Mitigator is the only product that scales to multi-gigabit networks that will pass 100% of legitimate transactions, even while under a sustained attack.

Before selecting any other IPS product, make sure you ask the vendor how they propose to protect your network from sophisticated new attacks without consistently performing “always-on” Stateful Inspection and Deep Packet Analysis.