

FAQ



TABLE OF CONTENTS

1	What is a Network Intrusion Prevention System?.....	2
2	Why do I need an IPS if I currently have a Firewall and an Intrusion Detection System (IDS)?	3
3	What is the Return on Investment for a Network Intrusion Prevention System?	4
4	What are the Essential Characteristics of an IPS?	5
5	What Protection Mechanisms Does the Top Layer Solution Use To Block Remote Exploits?	6
6	What Trends does Top Layer see based on Recent Cyber Crime Activity?	9
7	Why is Performance so Important when Considering Network Protection?.....	10
8	Are Zero-Day Attacks Real and Should I really be Worried?.....	11
9	Many Vendors Rely Heavily on Signatures to Identify Exploits; Why is this not the Best Method?	11
10	Why should an IPS be Stateful?	12
11	Why Should I Choose Top Layer's IPS 5500 E-Series IPS Over Other Leading IPS Solutions?	13

1 Question: What is a Network Intrusion Prevention System?

Answer: A Network Intrusion Prevention System (IPS) is an in-line security appliance that inspects network traffic, identifying malicious, harmful, and/or unwanted network activity and blocking it. The inspection performed by an IPS is done in real-time to ensure that good network traffic is able to pass through the IPS without noticeable delay.

There can be some overlap of functionality between network IPS and traditional firewalls, but it is clear that a firewall is not sufficient to protect against today's cyber threats. While each class of devices can block certain types of network transactions, how they affect networking configuration, how they perform traffic inspection, and how they approach system security are fundamentally different.

As a networking component, unlike most firewalls that also act as routers, an IPS is a transparent device on the network that does not have a visible IP address, and requires no network reconfiguration to deploy. While a firewall's basic task is to regulate the type of network "conversations" that are allowed between computer systems of differing trust levels, an IPS's job is to inspect protocol and application content on the network to ensure that it does not contain harmful, malicious, and/or unwanted content. Both firewalls and network IPS are frequently deployed at network perimeters. While both may be used internally in the network, the use of IPS to protect internal data centers and to perform internal network segmentation is far more common than the use of firewalls. Finally, while firewalls allow fine-grained policies to implement their traffic regulation, some IPS solutions are limited in their ability to apply inspection criteria discriminately, and must inspect all network traffic according to a single policy or rule setting.

The limited ability of most IPS solutions to apply firewall-like granularity to their inspection creates inherent limitations in the level of protection that can be realized since inspection rules that are best suited to protect client computers (e.g. desktops) may cause false positive alerts if applied to traffic going to servers, and vice versa. In an effort to reduce the false positive issue, many IPS solutions reduce their recommended signature set to a least-common denominator approach, lessening the very protection the IPS was installed to obtain.

The Top Layer IPS overcomes this limitation since it has an integrated stateful firewall, but remains a transparent network device. Top Layer customers can apply different IPS rules or signature sets to different classes of network traffic, thereby increasing the overall protection realized. This process is as simple as selecting pre-defined recommended or strict rule sets, and applying them to classes of network traffic that are defined in a familiar firewall-like policy.

Another characteristic of IPS products is their suitability for both perimeter and core deployments. Perimeter deployments typically place the IPS behind the firewall, allowing the firewall to apply its access controls first, and then the IPS further inspects

traffic that the firewall allows through. The Top Layer IPS has advanced DDoS protection capabilities which make it well suited to be deployed in front of the firewall to prevent the firewall from becoming a single point of failure in the event of a botnet attack. In fact, the majority of Top Layer perimeter deployments take advantage of this advanced feature. The power of the Top Layer IPS, compared to other IPS products is clearly demonstrated with deployments at the core. Firstly, the Rate-based algorithms protect against traffic floods, the built-in stateful firewall filtering blocks unauthorized access to specific network assets, and finally, with the IPS rule sets and acceptable application use policies, users can define what type of traffic can pass to specific applications.

2 Question: Why do I need an Intrusion Prevention System (IPS) if I currently have a Firewall and an Intrusion Detection System (IDS)?

Answer: Many organizations still rely on firewalls for network access control and Intrusion Detection Systems (IDS) for monitoring and identifying malicious network traffic. With the proliferation of easy-to-obtain-and-use hacking tools and the enormous profits that can be gained from stealing personal data and confidential information, these organizations remain at HIGH RISK of a successful breach.

The firewall is generally an organization's first line of defense. A firewall's basic task is to regulate the type of network "conversations" allowed between computer systems of differing trust levels. Typically, they block unauthorized access while permitting authorized types of communications. They are frequently used to prevent unauthorized Internet users from accessing private networks connected to the Internet. Modern firewalls can filter traffic based on many packet attributes like source IP address, IP source, source port, destination IP address or TCP/UDP ports. Since they are not designed to inspect application content, an attack from an allowed IP address will simply pass straight through the firewall. This is particularly a problem when it comes to handling services that must be open to the general Internet (web service, DNS, etc.). Also, consider an employee, third party contractor, or visitor who logs on to the corporate network inside the perimeter with an infected laptop computer. In this case any firewall security is circumvented altogether.

A network IDS is software and/or hardware that is deployed as a network monitoring tool and is designed to detect unwanted attempts at accessing, manipulating, and/or disabling of computer systems, mainly through a network, such as the Internet. IDS products were not designed to operate in-line in the network since they would become unacceptable choke points on the network. Although IDS might be effective at detecting suspicious activity, it does not provide any protection against attacks since it does not block the malicious packets or terminate the connection.

An IPS typically does not replace a firewall, but instead is used in conjunction with the firewall to provide a robust security infrastructure. Most IDS products have a very large database of known attack signatures which can provide valuable forensic information after an attack has occurred. However, since an IPS is deployed in-line, it would not be practical to have all possible rules or signatures enabled for automatic blocking as that could cause an unacceptably high level of false positives (blocking of legitimate network

Top Layer Security Intrusion Prevention System (IPS) FAQ

traffic). Since all networks are not alike, an IPS may require tuning during the deployment process. An IPS with flexible protection policies can provide an excellent balance of automatic blocking (without false positives) of harmful, malicious, and/or unwanted network transactions, and detection of other suspicious behavior that is less deterministic. In summary, an IPS will provide more value than just simply using IDS.

The Top Layer IPS has advanced detection capabilities with its Protection Processor inspection engine. Unlike other IPS approaches, the Top Layer IPS uses a state-of-the-art, multi-tiered "Protection Processor Architecture" that couples industry-proven protocol validation modules with data validation modules that inspect file content regardless of the protocol over which the files are being transported. This approach requires fewer rules or signatures and dramatically reduces the risk of false positives compared to other IPS technologies.

3 Question: What is the Return on Investment for a Network Intrusion Prevention System?

Answer: Most Top Layer IPS customers tell us that the payback from their IPS investment is rapid. Customers cite several reasons for this, including:

- Dangerous cyber attacks are stopped in real-time, eliminating the probable loss of data and almost certain remediation costs.
- Security team has adequate time to properly test operating system patches.
- Mission critical server downtime is eliminated, thereby maximizing revenue and maintaining high user/customer satisfaction.
- Unwanted network traffic is eliminated and available bandwidth and perceived network performance is increased.
- Operating expenses incurred by maintaining and running older, ineffective security solutions are reduced.
- Regulatory compliance standards are met reducing expensive remediation costs for audit findings.

In addition to ROI, when Top Layer customers evaluated IPS solutions, they were concerned about the relative total cost of ownership (TCO) of each solution. Ranking highly in the top five reasons for choosing the Top Layer IPS solution was the fact that the TCO was approximately 50% of other leading IPS solutions over a three year period.

4 What are the Essential Characteristics of an IPS?

Answer: For an IPS to provide effective non-stop protection against network and application-level attacks, the following aspects of a complete solution must be addressed:

Protection Capabilities

- Accurately block known and unknown (including zero-day) attacks.
- Security research team provides fast protection for newly discovered vulnerabilities and exploits.
- Not reliant on signatures as the primary form of defense (a method adopted by IPS products that spawned from IDS technologies that are susceptible to false positives).
- Always allow the good traffic to flow even when under attack.

Usability & Reporting

- Easy and fast out-of-the-box configuration providing immediate protection with minimal ongoing operational maintenance.
- Real-time security event reporting and alerting.
- A centralized management solution that has configurable reporting capabilities.
- Report relevant data for incident response and forensic analysis.

Resilience & Scalability

- Since it operates in-line, it must be a resilient hardware solution that will not be a single point of network failure.
- Not add any discernable latency under extreme load or attack, since this will negatively impact business users.
- As network capacity and performance increases over time, the IPS solution must be saleable in line with those requirements.
- Provide protection in complex network topologies such as asymmetrical networks.

In addition, an IPS should be validated using Common Criteria methodology. Common Criteria is an international standard for validating Information Assurance (IA) software and appliances, with a comprehensive range of evaluation criteria for government-use installations and corporate security products. While the products of 26 technology vendors have received validation through Common Criteria for their IPS/IDS capabilities, Top Layer is the only one to have received product validation at EAL4 with a specialized focus on IPS.

5 Question: What Protection Mechanisms Does the Top Layer Solution Use To Block Remote Exploits?

Answer: It is estimated that over 6,000 web pages are infected with malware every day, the equivalent of one every 14 seconds. What is more staggering is the fact that over 80% of these web pages actually belongs to innocent organizations. The point is that vulnerabilities are widespread in almost every conceivable piece of software ever written. The problem is that over 90% of vulnerabilities discovered in 2008 were remotely exploitable, tools and frameworks to remotely exploit vulnerable systems are widely available, and the profits from doing so are astronomical. Within this landscape you have the makings of the perfect business. In fact during 2008, the total number of malware variants documented surpassed the numbers from the previous four years combined.

In 2008 it was reported that there were over 400,000 variants of malware, ranging from Trojans, worms, viruses, downloaders, dialers, key loggers, rootkits and Spyware to name a few. Most of this malware is quite harmless against patched systems and a significant proportion of them are so old it is unlikely that an enterprise would even run systems and applications that could be compromised. There certainly is a benefit to an IPS having a good library of rules and signatures that cover the more important malware, but the focus should be on the capabilities surrounding updates for newly discovered exploits and vulnerabilities. The Top Layer IPS has a multitude of threat detection engines with specialized hardware to maximize performance and minimize latency.

Protocol Anomaly Detection

All protocols should adhere to standards such as Request for Comments (RFCs). An IPS must be able to determine whether the packets violate those standards which may be indicative of malware being present. In addition to determining whether the packets violate the standards, it must also be able to determine whether the data within the protocol adheres to expected usage. This expected usage could be industry-wide or at the enterprise level. For example, if peer-to-peer (P2P) applications were disallowed by an enterprise by policy, legitimate P2P traffic would traverse the firewall but should be blocked by the IPS. In contrast, a corporate policy may allow P2P, but disallow file sharing or other attachments. In this case the IPS must be able to identify any attachments associated with the protocol and strip out the attachments to be discarded. Since the Top Layer IPS applies stateful protocol inspection, it makes more intelligent decisions than those that rely primarily on signatures.

Data File Inspection

A significant proportion of attacks seen today results from malware contained in data that are used by applications, even though the transport protocol may adhere to the appropriate RFCs. For instance many attackers take advantage of vulnerabilities in Microsoft Office applications to launch their attack once the application runs the data with the embedded malware. Therefore an IPS must have the ability to inspect the data files. Here lies an interesting architectural conundrum. Theoretically you could create a signature for the exploit or vulnerability but you would need to be able to apply that signature to the applicable data file regardless of the transport protocol. For instance, it is quite common to send the malformed data packets across P2P, email, HTTP, and any number of other protocols. If there is no easy way to separately identify the data file

Top Layer Security Intrusion Prevention System (IPS) FAQ

format and therefore apply the relevant signature table, you would be forced to create a custom signature for every transport protocol. Not only would this become a time-consuming exercise, it may also lead to a greater incidence of false positives.

The Top Layer IPS uses a state-of-the-art, multi-tiered "Protection Processor Architecture" that couples industry-proven protocol validation modules with a new set of data validation modules that inspect file- content regardless of the protocol over which the files are being transported. This approach requires fewer rules or signatures than alternative solutions, which dramatically reduces the incidence of false positives compared to other IPS technologies.

Acceptable Application Usage

It is important that an IPS can restrict what an application is able to process thereby preventing unauthorized operations. The ability to combine access control and approved usage checks on application layer traffic is important. For example, a web server is able to process far more commands than a typical user would use in practice. By only permitting traffic to the web server that utilizes the allowed commands you would eliminate complete classes of potential attacks. When applied by the IPS, this type of protection can be effective at blocking zero-day exploits.

Signature Matching

Signatures are a dangerous term in the world of IPS. In the early days, IPS vendors touted the number of signatures they had as an indication of how good their products were. With a little probing, it was quite easy to see that there was a considerable difference between the numbers of signatures that could be applied to real-time inline blocking of attacks vs. those that could only be used for detection purposes only - in some cases the block to detect ratio was 1: 10!

There are several techniques that have been created over the years for applying signatures to network traffic to determine whether the packets contain malware. The earliest and most simple version was referred to as simple pattern matching. If the malware was buried deep within the packet payload, this technique may require inspecting a tremendous amount of data until the malware was discovered, causing an unacceptable performance degradation of the IPS. A more efficient form of pattern matching referred to as regular expression defines complex search patterns that increase the accuracy of malware detection. In order to minimize latency, a significant amount of hardware acceleration needs to be built in to the IPS device. It also makes sense that a signature that targets a vulnerability is more effective than one that targets a single exploit for the simple reason that there may be hundreds of variants of an exploit for a single vulnerability and having a signature for each variant has a greater potential.

Real-time Shunning

The Top Layer IPS has an effective protection capability called shunning that can quickly block traffic from IP addresses, temporarily or permanently, that are suspected of originating or being related to an attack. The advanced protection capabilities from shunning can be summarized as follows:

- **Attack Source Identification** - The "Security Event Viewer" enables users to identify a set of attacker IP addresses associated with blocked and detected attacks.

Top Layer Security Intrusion Prevention System (IPS) FAQ

- **Malicious IP Address Shunning** - isolate "events of interest" and automatically shun all IP addresses associated with a particular attack event. Users can set time periods for how long each address should be shunned, as well as manually unshun addresses that are determined safe.
- **Attack Defense Dashboards** – The user interface allows Security Operations Center personnel to switch between "quiet-time" monitoring and "under siege" incident response.
- **Additional Router Protection** - Administrators can export a list of IP addresses being shunned so that they can be imported into a router for blocking by the router.

6 Question: What Trends does Top Layer see based on Recent Cyber Crime Activity?

Answer: Valuable data loss (impacting brand reputation and attracting large monetary fines) together with regulations continue to be top-of-mind issues for CSO's and CIO's when it comes to network and application security. Much of this concern stems from the growing sophistication of cyber attacks and the multitude of ways they are being launched. There are so many entry points on today's network, whether it's email, FTP, Web services or wireless, security defenses need to be more comprehensive than ever before. What has been discovered from investigating the large number of recent high-profile data losses was that in many instances, systems became infected with remote exploits that could have been prevented in the first place with an IPS.

The heavy focus on operational and tactical issues by CSO's and CIO's comes amid a growing realization for the need of security managers to take a more strategic focus, in other words, you cannot separate the operation issues from the business issues. Maintaining secure business operations means the security manager needs to proactively address the key network and application threats for an organization before they happen. One thing is certain, at some point, every organization will be the target of an attack. Only those organizations that address the threat now will be ready to tackle it when it occurs. The difficulty is that these threats take many forms:

System Penetration	DDoS Attacks	Insider Abuse
Spoofing	Data/Network Sabotage	Unauthorized Insider Access
Worms and Trojans	Viruses	Hijacking IT Resources
Zero-Day Attacks	Compliance with Legislation	Loss of Intellectual Property
Unprotected Remote Links	Lack of Redundancy	Rolling Out New Applications

Amid this growing number of potential pain points is the fact that the attacks may span the range from simple brute force attacks to highly sophisticated and targeted ones. Over-provisioning with more servers and more bandwidth is not enough to defend against today's attacks. Current network intrusion prevention solutions provide the answer for enterprises to defend against known and unknown attacks while allowing legitimate business transactions to continue to flow to their destination.

7 Question: Why is Performance so Important when Considering Network Protection?

Answer: Performance is critical for an in-line IPS. The key performance aspects for an in-line IPS are latency, throughput, DDoS rejection rates, operation load, and scalability. The Top Layer IPS delivers industry-leading performance across all of these key attributes.

- **Lowest Latency Of Any IPS Device Ever Tested** – The Top Layer IPS was the first IPS to seamlessly integrate multiple protection mechanisms on a distributed ASIC platform. NSS Labs has tested over 25 IPS appliances and the results showed that Top Layer's IPS has the lowest latency of them all – generally measuring below 50 microseconds.
- **Saleable Performance and Capacity** - The Top Layer IPS ProtectionCluster™ technology is a proprietary load sharing technology built-in to each IPS. For example, two Top Layer IPS devices can be connected directly to each other, and up to eight Top Layer IPS devices can be interconnected using standard switches. In addition to increased throughput, attack defense capability, and session capacity, this configuration provides a transparent solution in standard and asymmetric redundant network configurations.
- **Industry Leading DDoS Rejection Rates** - Today, botnet attacks can be launched simultaneously from botnet armies of tens of thousands of compromised machines, delivering seemingly harmless legitimate traffic at multi-gigabit per second rates. Today, attackers target e-commerce sites, email servers, DNS servers, and VoIP providers to prevent legitimate transactions or data from reaching the desired target. Only the most advanced DDoS capabilities, designed in hardware, can stop these attacks while allowing legitimate traffic to continue to flow to the intended destination. Top Layer has been at the leading edge of stopping high volume DDoS attacks for many years. The Top Layer IPS incorporates this technology in all of its IPS products and allows customers to combine traditional IPS protection features with full DDoS protection.
- **Performance When Under Load** - This is the one performance metric missing from most vendors datasheets. As a result of the tight integration of the protection mechanisms with the hardware architecture, datasheet performance for the Top Layer IPS is what you can expect when deployed in live networks (with small packets), even while under attack.

8 Question: Are Zero-Day Attacks Real and Should I really be Worried?

Answer: Yes, and yes. Zero-day exploits occur when an exploit for vulnerability is created before, or on the same day that a vulnerability becomes known to the world at large. IT organizations are constantly fighting to keep their systems patched and updated, but the reality is it takes time to adequately test a patch against all applications running on the servers. This leaves organizations exposed to the narrowing of the time between discovering a vulnerability and the time an exploit is launched. As such, an attacker can effectively compromise unprotected servers at will.

In mid-2009 Microsoft took the rare step of announcing two critical vulnerabilities for which it did not have a fix, and for which remote exploits were being used in the wild. In each case an attacker could take control of vulnerable systems by enticing victims to simply open web pages that contain malware. In both cases, Top Layer's Security Team issued same-day protection to customers.

9 Question: Many IPS Vendors Rely Heavily on Signatures to Identify and Block Exploits; Why is this not the Best Method?

Answer: Signatures, or pattern matching is one of a number of methods that are used in an IPS to detect and block exploits of vulnerabilities. However, if used as the primary protection mechanism, you will face limitations in what will be successfully blocked. Signatures are notorious for generating false positives, which means that on their own, legitimate traffic will be blocked. In addition, attackers have found ways around pattern matching methods by making relatively small changes to the attack code that renders the detection useless; and therefore, not successfully blocked by the IPS. Another trick has a packet reorder engine and is fully stateful, the attack will never be recognized and will simply pass through to the ultimate target. It is therefore important to have multiple protection mechanisms all working simultaneously.

In the case of the Top Layer IPS, the IPS inspects 100% of the packets and integrates many protection mechanisms, including its Deep Packet Inspection and stateful Analysis Engines to understand an application's behavior and usage across the entire session. The reordered packets that comprise a transmission are inspected to establish whether it is legitimate or malicious. If deemed malicious, the entire packet stream is discarded before reaching its intended target.

10 Question: Why should an IPS be Stateful?

Answer: 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 network security solutions to address malicious attacker behavior. An IPS should also incorporate "always on" stateful inspection to allow continuous monitoring of packets. As well as examining header information, stateful inspection allows the entire packet content (up through the application layer) to be examined 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 or session 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 previously passed through the 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's were spawned from the IDS vendors that had their roots firmly planted in their ability to alert, report, and correlate attacks. The concept of taking these offline devices and putting them in-line and allowing them to block attacks based primarily on signature or pattern-matching techniques was quite logical. In fact, most of these vendors utilize a form of stateful inspection to complete simple pattern matching (also known as, signature matching) on packets to establish whether the packet contains a known exploit. As a result, these IPS vendors will claim that their products have stateful inspection capabilities. However, as soon as these IPS products are deployed in-line to perform proactive blocking rather than simple offline 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 an "always on" state, the performance and latency decline dramatically.

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 that, 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 then, the IPS device becomes a performance bottleneck in the network.

Having an IPS that is sometimes stateful and sometimes not creates a real challenge to network security managers. For instance, 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 in and go out on different network segments. If the IPS is not maintaining state for all transactions, it is again highly likely that attacks will not be identified and will be able to continue on their way to deliver their payload to their destination.

To get around the challenge of performance bottlenecks when stateful inspection is enabled at all times, an IPS vendor must invest heavily in developing purpose-built network processors that are seamlessly integrated together to reduce latency concerns while passing good traffic under load or attack. Only the most advanced hardware architecture allows for excellent protection at all times with no degradation in performance.

11 Question: Why Should I Choose Top Layer's IPS 5500 E-Series IPS Over Other Leading IPS Solutions?

Answer: The best way to summarize this is to articulate what customers regularly tell us:

- "Better protection with fewer false positives"
- "With the free IPS appliance program the total cost of ownership over three years is half of that using other leading IPS products I was considering"
- "No discernable latency added to the network with all the recommended protection features turned on"
- "The appliances are exceptionally well built"
- "Helpful and knowledgeable customer support team"
- "Impressive speed at which new protection is provided"
- "Only solution we've tested that really works in an asymmetric network"
- "The DDoS features alone are reason to buy this product at this price"
- "Common Criteria EAL4 validation was a necessity for selecting an IPS"