Zscaler Intrusion Prevention

zscaler™

While IPS inspection is a key component to a proper defense in depth strategy, understanding how and where to implement it within your organization is just as important. In this paper, we will consider what IDS/IPS devices are designed to do, how their capabilities differ based on their primary function, and where such functionality is best deployed. We will also consider Zscaler's offering, in which IPS capabilities are a major part of the company's Advanced Persistent Threat Protection strategy, in addition to a number of other, complementary techniques.

IDS and IPS Functionality

Intrusion Detection Systems were designed to monitor traffic for unauthorized network activity. They do this by analyzing a packet in its entirety, which includes the header and the payload, and comparing it to the signatures of known malware. An IDS appliance typically sits out of the direct data stream and reports on the bad packets it sees including:

- Malicious code
- Botnets
- Viruses
- Targeted attacks and exploits
- Spyware
- Cross=Site Scripting (XSS) attacks
- SQL injections

...and more

While properly tuned Intrusion Detection System is beneficial, it's important to note that, even at its best, IDS technology will alert an IT Team to an issue, but will not prevent it. An Intrusion Prevention System, on the other hand, is capable of detection bad packets and taking action. An IPS can do this because it is deployed inline, with the network traffic itself. IPS Technology can block malicious traffic by resetting and blocking the connection or by dropping packets. Since packets are forwarded as they are processed by the IPS, detection must be done in real time in order to block attacks before they reach their targeted victims. The IPS can also generate logs and alerts for administrators.

Both an IDS and an IPS typically sit behind a firewall. The firewall analyzes packet headers and enforces policy based on 5-tuple information, including protocol, source/destination address, and sort/destination port. If traffic is allowed based on a firewall scan, it goes to the IDS/IPS, which scans the entire packet and payload.

Advanced Threat Detection Methods

Signature-based detection (IPS)

An IPS primarily detects malicious traffic based on signatures. Signatures are a set of patterns that identify a vulnerability being triggered or an exploit being used. Signatures represent elements of a vulnerability or malware that must be present in an attack seen on the network. Signatures must be specific enough to avoid triggering false positives, in which legitimate traffic is incorrectly identified as malicious. But signatures must also be broad enough to stop variants of a known attack to ensure that a real attack does not get through.

Network traffic is parsed and pre-processed in order to make signature-based detection more efficient and more accurate. With HTTP traffic, for example, signatures can be applied on specific headers, on decoded content, on the request or on the server response. The IPS vendor must monitor and research known vulnerabilities and exploits in order to write new signatures. Typical IPS vendors update their signature database daily.

Anomaly detection (IPS)

As Intrusion Prevention Systems have become more widely used, attackers have found ways to evade signature-based detection. If the attacker can break the traffic pre-processing by generating traffic that will be parsed incorrectly by the IPS, but that will be handled correctly by the target, signatures are applied to the wrong part of the network traffic and may not trigger an action by the IPS. Common evasion techniques include multiple-encoding of the URL, using unusual white spaces to separate HTTP headers, or using unusual encoding techniques (7-bitASCII). IPS technology detects anomalies to prevent such evasion techniques. Anomalous traffic can then be flagged and blocked inline.

Behavior Analysis (Sandboxing)

While an IPS system can provide a strong defense for inbound threats, attackers have found ways to circumvent their detection. By weaponizing a file, and constantly changing the file slightly, a hacker can circumvent both signature and anomaly based IPS detection. Because the file that carries the malicious payload has been slightly changed, the resulting hash of the file changes. The file hash is a mathematical computation against a known file an IPS uses to confirm if the file has been seen before. Behavior analysis, which is common in Sandboxing technology, can perform in-depth analysis on a file's behavior to confirm if the resulting behavior associated with running that file on the target system will be malicious. While outside the scope of this paper, sandboxing is technology you will want to layer into your defense strategy in order to plug existing security gaps.

Form Follows Function

While the basics of IDS/IPS functionality are widely used, the manner in which that functionality is enabled, the place where it is deployed, and the way the information is used can be very different. There are still "pure play" IDS/IPS vendors, which are typically deployed in data centers to protect servers or to protect aggregated user traffic going out to the internet. In many cases, however, IDS/IPS technology has been absorbed into other products.

Unified Threat Management Appliances

One of the first product categories to incorporate IDS/IPS functionality was Unified Threat Management (UTM), which combines firewalling, IDS/IPS functionality, and gateway antivirus into a single appliance. Gartner defines the UTM market as multifunctional network security products used by small or midsize businesses[1];

When enterprises consider how to protect remote or branch offices, UTMs are often their first thought, since a single appliance appears to be fairly affordable. Unfortunately, looks can be deceiving, and the cost of purchasing UTM devices across several branch offices can be dauntingly high. When you add in the cost to install and deploy the appliances, ensure that policies interact with devices up and downstream, ascertain that policies are consistent across branches, handle updates and maintenance, and then correlate logs for a complete, company-wide picture, even the lowest-priced single UTM becomes exorbitant.

"Next-Generation" Appliances

IDS/IPS functionality is often discussed as a component of a Next Gen Firewall (NGFW). While the term remains somewhat nebulous, most agree that an NGFW is a device that enforces policy unilaterally and its inspection encompasses more than just the network packet header information of "traditional" firewalls. NGFWs can be set up in front of company servers to protect against illegitimate access to company assets. They are particularly useful in the case of the data center, where inbound or internal attacks are possible. NGFW appliances can also be set up inside the LAN to protect clients and servers against internal attacks, and can be placed at egress points to protect users accessing the Internet. But because NGFW appliances must cover all ports and protocols, their deployment in branch offices is generally unnecessary and definitely too costly, since the vast majority of traffic at the branch is HTTP/HTTPS.

[1] Gartner Magic Quadrant for Unified Threat Management Devices

Zscaler Cloud-Based Advanced Threat Protection

While the data center remains a central location that must be well defended, other locations, including remote or branch offices, can present challenges. In most cases, the majority of traffic in remote/branch offices is web traffic, and the only apparently cost-effective defense is typically a UTM system. While a single device might be affordable, duplication of these devices across an entire branch deployment is not. As a result, many branch offices are left with substandard or inconsistent intrusion detection and prevention. Also, remote user traffic is generally not examined at all, unless it is run through the data center IPS via latency-inducing VPN tunnels or costly MPLS links. Hackers are well aware of this reality.

There will always be groups targeting the largest organizations, and they know that they will find the most robust defenses there. But the largest organizations have branches and remote offices, with less robust defenses, and once a user in one of these locations has clicked on a phishing email or unwittingly downloaded malware on their device via a zero-frame exploit, they become the hacker's way into even the largest enterprise.

Zscaler is the answer. Based in the cloud, Zscaler's Advanced Threat Protection engine combines the best in IPS protection with a number of other capabilities, including antivirus/antimalware, blacklisting, and sandboxing. Branch offices and remote users simply direct their outbound HTTP and HTTPS traffic at a Zscaler Enforcement Node (ZEN), and Zscaler will geolocate each connection to the closest ZEN in one of the hundreds of data centers in which Zscaler is located. Zscaler then examines every byte of traffic as well as every response. There is no hardware to buy, software to update, or versions to maintain or coordinate. Zscaler was built from the ground up to eliminate issues around processing, performance, and scale that are inherent in perimeter-based appliances.

Challenges of IPS Appliances

Delivering IPS functionality in a perimeter-based appliance creates a number of issues that go far beyond the cost of deploying and maintaining a hardware or software device. These challenges are common to perimeter-based security devices, particularly those outside the data center perimeter. Issues include:

Limited Technology

Performance of an IPS is linked to the amount of protocol decoding the appliance has to do and the amount of pattern matching it has to perform. IPS technology is not designed to handle large blacklists of URLs or IP addresses. It also lacks support for antivirus and more advanced file analysis technologies such as file sandboxing mentioned previously. Today, IPS is integrated into more complete UTM systems to offer broader security, but as previously discussed, UTM systems are not a logical option for remote users or branch offices.

Performance

The performance of IPS appliances depends on the number of signatures enabled and the throughput of content to inspect. For HTTP, where the response is typically much larger than the request, enabling Server-to-Client HTTP signatures dramatically slows IPS appliance performance; in fact, most IPS appliances turn off response scanning by default at the highest throughput claimed.

Unfortunately, the result is that security drops dramatically. A full scan of the HTTP response is required to detect many types of attacks including browser exploits, exploit kits, and more. In addition, the majority of botnet traffic is dynamic and the IP addresses and domains involved are constantly changing. In this case, the malicious traffic can be more efficiently blocked by inspecting the response rather than the request as it enables visualization of elements including the configuration download, list of ISP or domains to connect to, and more.

Lack of SSL Decryption

As transparent devices, many IPSs struggle to perform man-in-the-middle (MiTM) SSL decryption. HTTPS decryption is processor-intensive and severely limits the performance throughput IPS hardware appliances can deliver. This problem becomes even more significant as the amount of SSL traffic grows. Accord Google Transparency report, over 90% of traffic traveling through Google is now encrypted[2]. In addition, free SSL certificate sites like LetsEncrypt have enabled hackers can now even enable SSL delivery from malicious websites. As more threats and hackers migrated to SSL, a strong ability to inspect all SSL within a company's defense strategy is needed.

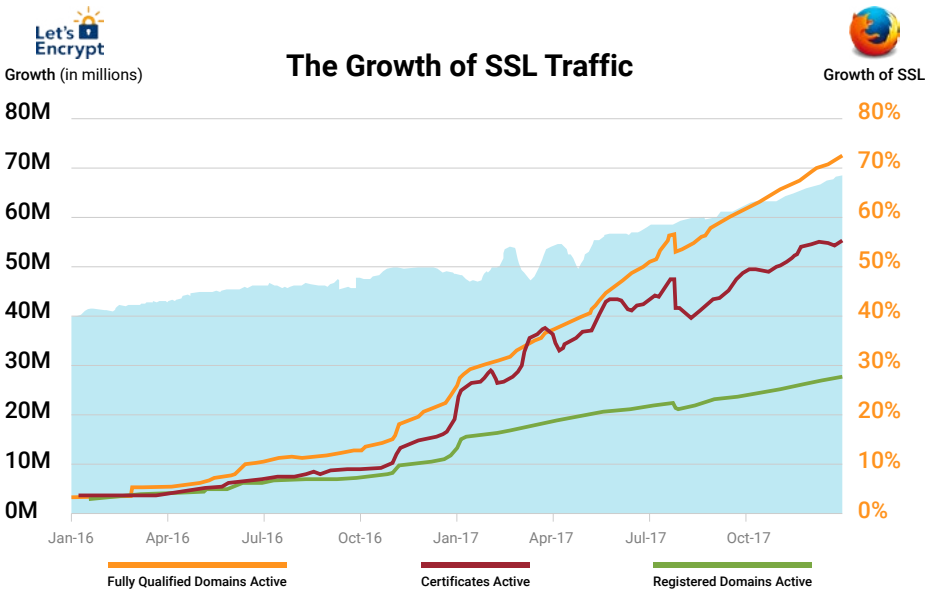

[2] Google Transparency Report: https://transparencyreport.google.com/https/overview?hl=en

Asynchronous Traffic

In many networks, traffic from a user can often egress and ingress from different gateway locations within the network. Called asynchronous traffic, this routing situation often causes IPS and other inspection devices to miss relevant threats. For proper threat detection, an IPS or other inspection appliance often needs to correlate both the client and server side of the user communication pattern together to properly perform signature detection. Because it is often impossible to place a single IPS appliance across multiple gateways locations, this separated inbound and outbound user traffic can't be properly tracked and inspected by the appliance, which causes missed attacks.

Unfriendly Packet Drops

IPS blocks traffic at Layer 4, but applications such as Web browsers work at Layer 5 and above. Some IPS platforms do not provide a graceful TCP reset when packets are dropped. Many applications will therefore keep retrying connections that get reset, resulting in an increase of traffic on the network.

Zscaler's Cloud IPS

Zscaler integrates IPS technology as part of Advanced Threats detection. IPS is one of many types of inspection done along with blacklisting, heuristics (Page Risk), antivirus, file sandboxing, and more. Because Zscaler protects users, not servers, the IPS does not look for server attacks such as SQL injections, denial of service, or remote code execution, which are the types of threats that you would expect to see in the data center. Instead, Zscaler's IPS focuses on detecting threats to users over HTTP and HTTPS, making it the ideal solution for branch and remote offices, as well as mobile users. Because if its performance, scalability and integrated security services, Zscaler can also easily add an extra layer of security to any deployment anywhere.

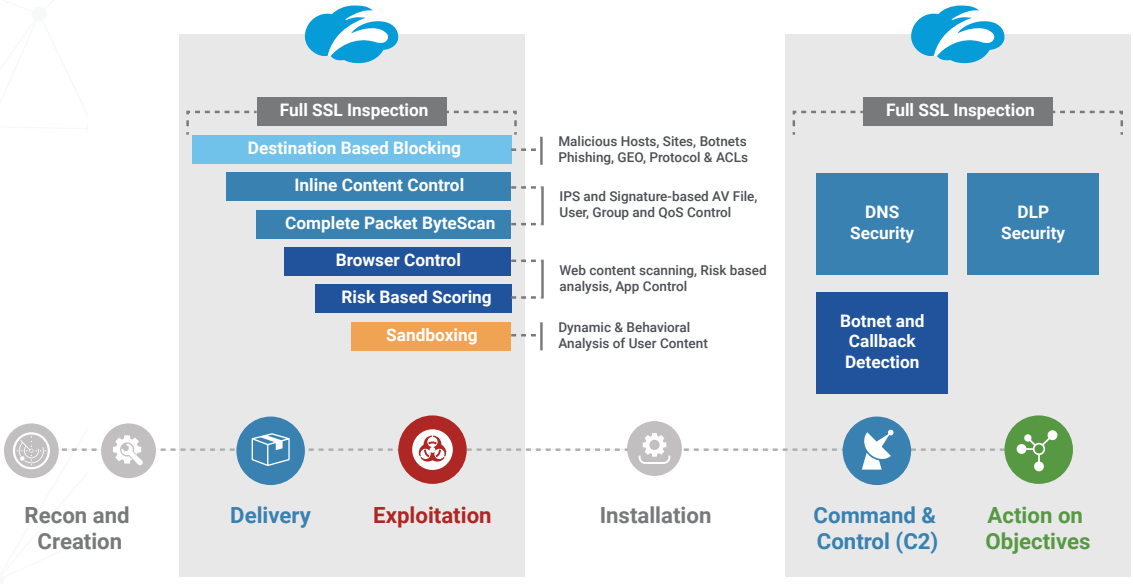

The Zscaler Cloud Platform integrates IPS protection across a complete security stack. You get full protection across all threats, even in SSL, without the hardware performance limitations or complex integration efforts.

Signature based categories covered by Zscaler Cloud IPS

Botnet Protection

- Command and Control Server
- Command and Control Traffic

Malicious Active Content Protection

- Malicious Content & Sites
- Vulnerable ActiveX Controls
- Browser Exploits
- File Format Vulnerabilities
- Blocked Malicious URLs

Fraud Protection

- Known Phishing Sites
- Suspected Phishing Sites
- Spyware Callback
- Web Spam

Unauthorized Communication Protection

- IRC Tunneling
- SSH Tunneling
- Anonymizers

Cross-Site Scripting (XSS) Protection

- Cookie Stealing
- Potentially Malicious Requests

Suspicious Destinations Protection

P2P File Sharing Protection

P2P Anonymizer Protection

P2P VOiP Protection

Crypto Mining

Zscaler use its proprietary ByteScan technology to scan all traffic, both client-to-server and server-to-client. The request and response are both parsed to match vulnerability and exploit signatures on all web traffic. Zscaler is able to consider both the request and the much larger response together, giving you a full picture of threats. All signatures are applied in real time on the request and response for every transaction. All web traffic goes through the Zscaler IPS, whether it originates from a web browser or any application running on the client device.

Signature-Based Detection

All inbound and outbound HTTP and HTTPS traffic is parsed to extract URL, headers, POST data, response body, and more. The Zscaler Security Team releases over 2,000 new signatures every year that cover browser and application vulnerabilities, including those from the Microsoft Active Protections Program (MAPP) as well as disclosure programs from other vendors. Other signatures include those from exploit kits, Command & Control traffic, cross site scripting, and more.

Perhaps more important, with Zscaler, signatures are updated and added transparently several times a day throughout the entire Zscaler cloud. This is in stark contrast to appliance based IPS functionality, in which signature updates may happen once a day at best. Still another benefit is that once Zscaler detects any threat, it is blocked for all users. This means that every user is protected the first time any threat is detected.

Anomaly Detection - Proxy IPS

Zscaler Enforcement Nodes are proxies and are not transparent devices. ZENs receive a request from the clients and create a new request to the destination server. In order to perform that function, ZENs must be able to correctly parse the entire request. Evasion techniques that trick security devices into misunderstanding the requests do not work against a proxy architecture such as Zscaler's. If a ZEN cannot parse a request, the request is not forwarded to the destination. Attempts to evade Zscaler's protocol decoders prevents malicious requests from being forwarded.

Prevention

Zscaler Enforcement Nodes send user-friendly HTML notification pages to the users and applications when malicious traffic is blocked. This means that the user understands that a request was blocked and why, and can then take action. All attacks are blocked, logged, and reported in real time. Administrators can easily correlate attacks or review the user activity before the attacks, making forensics much simpler to perform.

SSL Inspection

As a proxy, Zscaler can decrypt SSL and fully scan the HTTPS traffic in order to match signatures. All web traffic, encrypted or not, goes through the same level of security protection. This eliminates the significant issue of SSL "blind spot," and Zscaler enables it without the performance impact felt by appliances. SSL inspection be turned on and off based on URL categories or Cloud applications and locations.

High Performance, Low Latency across all traffic

Zscaler has developed its own IPS engine for scale. Signature-based detection on both request and response is always turned on, and there is no additional latency when the Advanced Threat policy is set to block threats. In addition, the challenges of inspecting asynchronous traffic is no longer an issue, as the Zscaler cloud easy covers all connection routes a user can use accessing the internet. The results is Zscaler customers do not have to choose between speed and coverage; they get great speed with great coverage all the time.

Conclusion

IPS has a role in protecting users against certain types of threats. Zscaler leverages ByteScan and signature-based detection, along with a fully integrated security stack that includes Sandboxing and SSL inspection in order to block the most exploits and malicious traffic targeting users. By enabling SSL decryption, all traffic can get the same level of security inspection. Without the limitations of hardware, and the ability to elastically scale to an organizations traffic demands, the Zscaler suite of Advanced Threat Protection techhniques can far surpass the level of protection delivered by other IPS offerings in the market.

To learn more about **Zscaler Cloud IPS, Zscaler Cloud Sandboxing,** or the complete suite of **Zscaler Advanced Threat Protection** offering, visit our website, or reach out to us **for a demo.**

Zscaler, Inc.
110 Rose Orchard Way
San Jose, CA 95134
+1 408.533.0288
www.zscaler.com