F5 White Paper

Security 101: BIG-IP ASM and IPS Differences Defined

Industry experts have long debated which is a better defense mechanism in defending against Internet based attacks: a web application firewall (WAF) like F5 BIG-IP Application Security Manager (ASM) or an intrusion detection or prevention system (IDS/IPS). But the most ideal scenario would include a WAF and an IDS/IPS. The key to a successful security plan is not to provide the best security at one layer, but the best security at all layers.

by Peter Silva
Technical Marketing Manager

Contents

Introduction

As they endeavor to secure their systems from malicious intrusion attempts, many companies face the same decision: whether to use a web application firewall (WAF) or an intrusion detection or prevention system (IDS/IPS). But this notion that only one or the other is the solution is faulty. Attacks occur at different layers of the OSI model and they often penetrate multiple layers of either the stack or the actual system infrastructure. Attacks are also evolving—what once was only a network layer attack has shifted into a multi-layer network and application attack. For example, malicious intruders may start with a network-based attack, like denial of service (DoS), and once that takes hold, quickly launch another wave of attacks targeted at layer 7 (the application).

Ultimately, this should not be an either/or discussion. Sound security means not only providing the best security at one layer, but at all layers. Otherwise organizations have a closed gate with no fence around it.

Brief History of IDS/IPS

Intrusion detection systems have been around since the mid-80s and are based on behavioral analysis. They grew out of a 1972 paper written for the U.S. Air Force[1] that discussed the need for increased computer security awareness. Initially the idea was that the government could use audited records to identify the misuse of computer systems and potential threats. Early prototypes would track user activity and analyze audit trails, flagging data that was out of the ordinary. Researchers then used models and patterns of behavior to create a statistical baseline of the types of information being sent over the network. With this baseline established, they could audit the traffic against defined patterns to discover deviations that might indicate a threat.

Soon, rather than analyzing audit/system logs, Network System Monitor devices started to analyze network traffic to look for deviations in the baseline patterns. These improvements then lead to the first commercially available IDS. Throughout the 1990s, IDSs quickly improved and matured. The Automated Security Incident Measurement IDS, developed by the U.S. Air Force, incorporated both software- and hardware-based solutions. Toward the end of the 90s, host-based solutions and network routers with built-in detection capabilities had emerged.

[1] Anderson, James P. Computer Security Technology Planning Study, ESD-TR-73-51, ESD/AFSC, Hanscom AFB, Bedford, MA 01731 (Oct. 1972) [NTIS AD-758 206]

The late 90s also saw the development of intrusion prevention systems. Initially, detecting anomalies was a huge task but as the technology matured, it became possible to not only detect, but prevent intrusions in real time. This is the entire concept behind intrusion detection—finding and stopping attacks in real time. Some products could stop or "kill" traffic based on specific patterns and in 1998, Snort, an open source packet sniffer and logger, became available. This enabled IT administrators to test and begin to understand the concept of intrusion detection.

The next generation of IPSs could intercept files and network activity, then make associations with application state and policy rules to make real-time policy decisions based on that activity. IDSs detect and control malicious network traffic akin to packet sniffers; IPSs add to that the ability to prevent malicious attacks like worms, viruses, and Trojans. Most IDSs and IPSs rely on a signature database of known attacks or vulnerabilities, along with malicious "fingerprints" or abnormal packet activity that may signify malicious activity. IDSs and IPSs also monitor protocol deviations in some cases.

IDS/IPS Deployments

Often, IDS and IPS devices are deployed as perimeter defense mechanisms, with an IPS placed in line to monitor network traffic as packets pass through. The IPS tries to match data in the packets to data in a signature database, and it may look for anomalies in the traffic. IPSs can also take action based on what it has detected, for instance by blocking or stopping the traffic. IPSs are designed to block the types of traffic that they identify as threatening, but they do not understand web application protocol logic and cannot decipher if a web application request is normal or malicious. So if the IPS does not have a signature for a new attack type, it could let that attack through without detection or prevention.

With millions of websites and innumerable exploitable vulnerabilities available to attackers, IPSs fail when web application protection is required. They may identify false positives, which can delay response to actual attacks. And actual attacks might also be accepted as normal traffic if they happen frequently enough since an analyst may not be able to review every anomaly.

Brief History of WAFs

The concept of an application layer firewall came about in the early 1990s. In 1990, Bill Cheswick wrote a paper called "The Design of a Secure Internet Gateway," which described a system that kept the corporate network separated from the Internet. This would allow only certain traffic to cross over from the Internet and protect the corporate network from intruders. Marcus J. Ranum followed that in 1992 with "A Network Firewall," which affirmed the need for a secure gateway between the Internet and corporate network "to prevent miscreants and unwelcome visitors from accessing hosts on the private network." At the time, corporate networks had very weak protection, and securing them was extremely difficult.

In the late 1990s, as more web and even business applications made their way to the Internet, WAFs were playing an increasing role in enterprises' success. In the rush to "webify," organizations side-stepped many security considerations, exposing applications to myriad vulnerabilities. Since fixing the code was too costly and time-consuming, they turned to WAFs for protection. These early WAFs functioned as HTTP filters that sat in front of the web application, protecting it by refusing to process any artificial character inputs that could seize control of the server. By the late 1990s, the idea to use an application-level security policy, and to enforce every incoming request based on that policy, was introduced. The goal was to secure e-business applications immediately and on the fly since security was not taken into account in the design or implementation of the various applications at the time.

WAF Deployments

WAFs have greatly matured since those early days. They can now create a highly customized security policy for a specific web application. WAFs can not only reference signature databases, but use rules that describe what good traffic should look like with generic attack signatures to give web application firewalls the strongest mitigation possible. WAFs are designed to protect web applications and block the majority of the most common and dangerous web application attacks. They are deployed inline as a proxy, bridge, or a mirror port out of band and can even be deployed on the web server itself, where they can audit traffic to and from the web servers and applications, and analyze web application logic.

They can also manipulate responses and requests and hide the TCP stack of the web server. Instead of matching traffic against a signature or anomaly file, they watch the behavior of the web requests and responses. IPSs and WAFs are similar in that they analyze traffic; but WAFs can protect against web-based threats like SQL injections, session hijacking, XSS, parameter tampering, and other threats identified in the OWASP Top 10. Some WAFs may contain signatures to block well-known attacks, but they also understand the web application logic.

In addition to protecting the web application from known attacks, WAFs can also detect and potentially prevent unknown attacks. For instance, a WAF may observe an unusually large amount of traffic coming from the web application. The WAF can flag it as unusual or unexpected traffic, and can block that data.

Many web applications are still developed with speed, but not security in mind. While the need for secure code has gained the attention of developers, there are still millions of websites corrupt with bugs and vulnerabilities. It is impossible to go back and fix all that code, and traditional network firewalls don't understand application logic. But by deploying a WAF, organizations mitigate the risk of a potentially vulnerable web application, just as they would deploy network firewalls to mitigate network threats.

WAFs Versus IPSs

A signature-based IPS has very little understanding of the underlying application. It cannot protect URLs or parameters. It does not know if an attacker is web-scraping, and it cannot mask sensitive information like credit cards and Social Security numbers. It could protect against specific SQL injections, but it would have to match the signatures perfectly to trigger a response, and it does not normalize or decode obfuscated traffic. One advantage of IPSs is that they protect the most commonly used Internet protocols, such as DNS, SMTP, SSH, Telnet, and FTP.

There are a few additional differences between a WAF like BIG-IP ASM and a traditional IDS or IPS:

- The generic detection capabilities of IPS are falling short in the world of web applications because many of the attacks look like valid HTTP requests. Take for example a recent financial institution incident, where a simple manipulation of a parameter value within the URL allowed an attacker to see account details. The requests looked valid and execution was allowed. Another example is cookie manipulation. Some applications require that cookies not be changed on the client side. IPSs cannot sign cookies against manipulation.

- To provide a high level of protection against well-known web application attacks, a security policy needs to be granular. For example, it must be able to prevent a SQL injection attack. The IPS signatures may only be enforced on certain parameters, which can generate false positives. This level of granularity, essential to a sound security policy, is missing from IPSs.
- IPS signatures are driven by specific vulnerabilities of specific applications, and many have a CERT ID. For example, there are signatures that can mitigate a SQL injection on php-bulletin board login page. These signatures can provide value to php-bulletin board customers; however, within the enterprise, there are many custom, home-grown web applications. IPS vendors will never have access to these home-grown web applications that may contain unclassified vulnerabilities, and therefore they will never have signatures for those vulnerabilities. WAFs have a unique expertise in writing generic attack signatures that, when applied with granular rules, can mitigate attacks for custom and home-grown web applications.
- To prevent web application attacks, organizations must accurately decode and normalize the content before the policy engine analyses it. For example, it is important for the application to know which language to expect to receive from the browser. It is also important to correctly analyze strings and correctly evaluate strings that the back-end parsers are treating as comments. IPSs do not have either of these capabilities.
- Signatures are often simply not enough for a sound security policy—organizations need a positive security model. For example if a given parameter value can only be 12 characters long, it is virtually impossible for a hacker to mount an XSS attack. Simply being able to enforce a different character set per parameter is also important. For example, an administrator might allow the ' character on the username field (so Mr. O'Neil can log in) but then deny it on other parameters. This limits attackers' ability to use those other parameters as a vehicle for SQL injection attacks.

BIG-IP ASM: Total Protection

The best security implementation will likely involve both an IPS and a WAF, but organizations should also consider which attack vectors are getting traction in the malicious hacking community. F5® BIG-IP® Application Security Manager™ (ASM) is a full-featured WAF that delivers comprehensive protection for web applications while maintaining low total cost of ownership. Although it also has some IPS features, BIG-IP ASM's primary focus is protecting web applications at layer 7.

Mitigating the OWASP Top 10: BIG-IP ASM Versus IPSs

According to WhiteHat Security's winter 2011 report,[2] the top ten website vulnerability classes of 2010 were:

Vulnerability	Prevalence (chance a given website has this vulnerability)	BIG-IP ASM protection	IPS protection
Information leakage	64%	DataGuard and streams	None
Cross-site scripting (XSS)	64%	Signatures, metacharacter enforcement, parameter protection	Signatures
Content spoofing	43%	URL and parameter protection	None
Cross-site request forgery	24%	Token injection, flows, referrer header, dynamic parameters	None
Brute force	17%	Brute force thresholds per URL	None
Insufficient authorization	15%	Authentication, ACA, iRules®, virtual patching	None
Predictable resource location	14%	Signatures, parameter and URL protection	Signatures
SQL injection	14%	Signatures, metacharacter enforcement, parameter protection	Signatures
Session fixation	14%	Signatures, iRules, cookie signing/encryption	Signatures
Abuse of functionality	10%	Virtual patching, iRules	None

These are just a few of the many attacks BIG-IP ASM protects against at layer 7. An IDS or IPS has only one solution to those problems: signatures. Signatures alone can't protect against zero-day attacks for example; proactive URLs, parameters, allowed methods, and deep application knowledge are essential to this task. And if a zero-day attack does occur, an IPS's signatures can't offer any protection. However if a zero-day attack occurs that BIG-IP ASM doesn't detect, it can still be virtually patched using F5's iRules until a there's a permanent fix.

[2] WhiteHat Website Security Statistic Report. Winter 2011, 11th Edition – Measuring Website Security: Windows of Exposure. 2011

A security conversation should be about how to provide the best layered defense. BIG-IP ASM protects traffic at multiple levels, using several techniques and mechanisms. IPS just reads the stream of data, hoping that traffic matches its one technique: signatures.

Is Secure Coding the Answer?

While secure coding is good practice, it also has a few challenges. It can be difficult to implement since developers are not trained to write secure code. Developers' focus is functionality and usability, and while secure code is ideal, it's seldom a realistic expectation of developers.

Even if an organization is heavily invested in secure coding practices, nothing will ever negate 100 percent of human error. In the software world, these mistakes are known as bugs. Any bug with a security angle is a vulnerability, and it's generally acknowledged that all products are shipped with bugs.

Some problems simply cannot be fixed with secure coding; for example, sometimes the vulnerably exists in the underlying platform, like the web server, or the OS. Sometimes it's within a third-party library in use within the application, or else the application was developed by a third party. It could be a legacy application for which no one even knows the code.

Once a vulnerability is discovered, fixing it can take a lot of time and effort. This requires development team cooperation, which can be hard to come by if they are working on an upcoming release, and it requires test team cooperation to make sure the fix doesn't break any functionality. Overall, this approach is time-consuming and can be very expensive.

Security officers often find themselves responsible for website security, but without any ability to control that security. They know where the vulnerabilities are, but they do not control the development team that would be able fix the issues.

Additionally, without a WAF, no one knows if a given application is being hacked. No one logs the full transactions, no one can look into the SSL traffic, and no one knows if an attacker is attempting a breach on the website. Only a WAF like BIG-IP ASM can look into the SSL traffic and log the full requests that are needed for a good audit trail and forensics and to detect sophisticated application attacks.

A WAF can help both development and IT by giving the development team time to fix the code, and giving web application access control back to the IT/security group.

Conclusion

Web application firewalls like BIG-IP ASM are unique in that they can detect and prevent attacks against a web application. They provide an in-depth inspection of web traffic and can protect against many of the same vulnerabilities that IPSs look for. They are not designed, however, to purely inspect network traffic like an IPS.

If an organization already has an IPS as part of the infrastructure, the ideal secure infrastructure would include a WAF to enhance the capabilities offered with an IPS. This is a best practice of layered defenses. The WAF provides yet another layer of protection within an organization's infrastructure and can protect against many attacks that would sail through an IPS. If an organization has neither, the WAF would provide the best application protection overall.

BIG-IP ASM gives organizations the fastest, most comprehensive, and scalable web application firewall and protects companies from the most serious security threats that cyber attacks pose. It can help organizations quickly pass a security audit without requiring changes to the application code, and it ensures application availability by delivering comprehensive, flexible protection for web applications.

F5 Networks, Inc. 401 Elliott Avenue West, Seattle, WA 98119 888-882-4447 www.f5.com

F5 Networks, Inc.
Corporate Headquarters
info@f5.com

F5 Networks
Asia-Pacific
apacinfo@f5.com

F5 Networks Ltd.
Europe/Middle-East/Africa
emeainfo@f5.com

F5 Networks
Japan K.K.
f5j-info@f5.com