White Paper

Vulnerability Assessment Plus Web Application Firewall (VA+WAF)

June 2008

by F5 Networks® and WhiteHat Security Solution

Contents

Introduction

Inside an enterprise lives an IT security professional responsible for website security. He takes his job seriously, because if his employer's websites get hacked, he gets the late night call from the boss upstairs. A big part of the job requires educating developers on the importance of secure coding and informing the business owners of web security risks. He does this because no amount of patching or firewalling will fend off an attacker with a web browser. While doing everything within his power, there's still a total lack of control in protecting the websites he's responsible for. He can't find the vulnerabilities with a traditional network scanner, nor can he fix them in website(s) when they're found without developer involvement. But, this is all about to change.

New technology developed by WhiteHat Security and F5 Networks enables organizations to immediately mitigate discovered vulnerabilities using a web application firewall. WhiteHat's Sentinel service provides continuous assessments of web applications for vulnerabilities. Once detected and validated by WhiteHat, detailed information about these vulnerabilities are passed to the F5 BIG-IP® Application Security Manager™ (ASM) web application firewall, where they are implemented as blocking rules and prohibit exploiting the detected vulnerability. IT Security professionals are able to get timely and accurate application security assessments and immediately block exploitation of vulnerabilities. The best part is that it all happens almost immediately and under the control of the security group with no dependence upon developers for patches. The end result is responsive and manageable web security.

The Current State of Vulnerability Management

This situation has become all too familiar with today's e-business enabled enterprises that are at risk. Recent studies say 9 in 10 websites contain serious security issues[1, 2, 3] are now the #1 target for malicious hackers. The problem is: when website vulnerabilities are identified by a pen-tester, developer, outsider, or whomever, there is always a certain amount of time required to determine the appropriate solution. Resolution could take the form of a software update, configuration change, web application firewall rule, etc. In any case, the time to fix should be swift because hackers will exploit the websites' vulnerabilities when no immediate remedy is implemented. Published reports state that nearly 80 percent of website hosting malware are legitimate and have been hacked[4].

While the source code is being fixed or system configuration updated, the an organization has three options:

1. Take the website down
2. Revert to an older version of the website/code (if it's secure)
3. Stay up while exposed

The cold reality is vulnerabilities happen despite the most regimented software development lifecycle. Historically option #1 (taking down the website) is employed when an incident has occurred; option #2 (rolling back the code) is preferable when a hot fix is not back-ported to development and is later overwritten. Practically speaking, the vast majority of website owners default to option #3 (do nothing), essentially assuming the risk rather than halt business.

Why do so many companies choose not to act? While organizations and their security teams have good intentions, the challenges associated with remediation vulnerabilities in web applications are daunting. For most, this involves the time consuming process of allocating the proper personnel, prioritization of tasks, QA regression testing the fix, and finally scheduling a production release. Figure 1 illustrates just how long it takes for the average organization to fix some of the most pervasive and widely exploited vulnerabilities.

Clearly, organizations must become more efficient at identifying web application security problems, remediate more quickly, and adapt better to new attack techniques. When speaking with IT security personnel, the issues they voice speak to the disconnect between them and the software development groups. IT security possesses little control over the security of the website in comparison to their control of the network or hosts. Patches cannot be applied to resolve custom web application vulnerabilities. So, they must coordinate with development, which typically does not report to them, to get a code fix in place. Also, IT security has a difficult time explaining the details and associated risk of a vulnerability to this less security savvy audience.

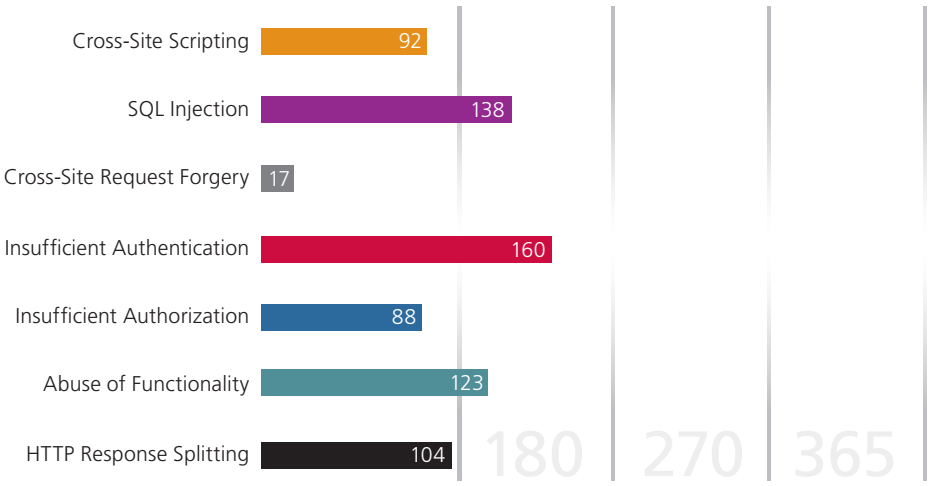

Figure 1: *Average Time to Fix by Class of Attack Measured in Days*

Overcoming these challenges requires a cutting-edge yet pragmatic approach: leveraging the tight integration of precise, comprehensive vulnerability assessments with web application firewall technology. Such a solution,

1. Measurably improves security;
2. Drastically reduces the time-to-fix from months or years to days or hours;
3. Assists organizations meet industry and governmental regulations, such as PCI-DSS 6.6 compliance;
4. Enables vulnerability assessment results to be immediately actionable;
5. Eases WAF configuration and management, demonstrates due care. And, most importantly provides a revolutionary fourth remediation option to those discussed above.

"Virtual patching"—enabling IT security professionals to regain control over website security by eliminating vulnerabilities as they are detected without developer intervention.

Vulnerability Assessment

Sophisticated cyber-criminals supported by organized crime or by nation states are exploiting website vulnerabilities at an alarming rate. Hardened networks have caused the attackers to focus on more vulnerable web application targets. Beyond just stealing identity information (social security and credit card numbers), cyber-criminals commonly penetrate one of any number of weak spots in an organization's website and silently lace the web pages with malicious code. When users visit the organization's website, their web browser is automatically exploited and their machine loaded with Trojan horses designed to steal passwords, send spam, attack other computers, and more. In April 2008, a single massive hack infected hundreds

of thousands of web pages with malicious code using a sophisticated form of blind SQL injection. With 9 out of 10 websites possessing serious vulnerabilities, it's best to know what issues exist before they can be exploited.

Every effective vulnerability assessment program requires a cohesive combination of people, process, and technology. Qualified people are necessary to carry out day-to-day tasks, manage the technology, and interpret the results to make them meaningful to the business. Process is required for coordinated efforts between executive management, IT security, and software development groups to share information, prioritize vulnerability fixes, and enable organizational improvements. The right technology is essential for consistency, efficiency, and comprehensiveness. Whether an organization chooses to perform vulnerability assessments with internal resources, a consultancy, or a software-as-a-services vendor, the overall vulnerability program must always account for people, process, and technology. If not, the effort will cost more in time and dollars than it should. Or worse, simply not work.

Still, no matter how perfect any vulnerability assessment product or service, the challenge remains: Any identified custom web application issues must be resolved by the organization—a task inevitably falling to the software developers and not IT security. This is a problem that cannot be solved automatically with a vendor-supplied software patch or new network firewall rule. This is where web application firewalls play a powerful role.

Web Application Firewalls

Web application firewalls (WAFs) are hardware or software devices positioned to monitor website traffic, with the ability to enforce policy on browser/server transactions. WAFs are similar, though not identical to, network firewalls where policies are typically applied to IP addresses, ports, and protocols. WAFs are specifically designed to inspect HTTP(s) traffic and regulate data contained within headers, URL parameters, and web content. Another similarity: network firewalls are used to protect insecure hosts from remote exploitation. WAFs do the same for insecure websites. With a WAF in place, malicious hackers may target insecure websites, but attacks are intercepted and denied before reaching the custom web application code. Beyond offering defense-in-depth, WAFs are now considered a fundamental part of any website security program, especially in light of the upcoming PCI 6.6 compliance deadline.

WAFs at their core are designed to separate safe web traffic from malicious traffic before it's received by the website. And, if an attack does find a way to sneak past a WAF, it still has the ability to prevent sensitive information from leaving the trusted network. To get a better understanding of how the technology works, it's helpful to view a WAF's functionality as three discrete components—policies, policy generation, and policy enforcement. Depending on the particular WAF in use, they may go about implementing each component in a number of different ways. No one particular way has proven to be the right way, as each has its pros and cons.

Every website is different, as is every business. Accordingly, WAF policies are all unique and customized for each website except for basic, universal security requirements like valid HTTP protocol enforcement. WAF policies are created to programmatically describe what a website should or should not do. They'll take the form of either a white list, black list, or a combination of both. A white list policy might allow only explicitly defined web pages (URLs) to be served from the website, with all other requests rejected. A black list policy might deny any requests containing "' SELECT * FROM", which indicates a possible SQL injection attack. The challenge is that websites are diverse, complex, and constantly changing, requiring policies with hundreds if not thousands of clear and precise rules. One can see how difficult this would be to manage with a typical security team. This makes policy generation an extremely important function to streamline the process.

WAF policy generation may take place in three ways: "learning," via vulnerability assessments, or manually. "Learning" is a process whereby website traffic is passively monitored for what types of requests are normally received, then dropping the rest. This approach often requires weeks or months to complete initially, plus the subsequent correction of any mistakes. While very helpful at eliminating much of the initial grunt work, this method is often made complex when rogue attack traffic is mingled with good traffic in a production environment, confusing the learning engine. The site is also left vulnerable while the process plays out.

Another option is leveraging the results from the vulnerability assessment process, assuming it is an ongoing, methodical process, which we'll discuss in greater detail. Vulnerability assessment generated policies provide a WAF with intelligence regarding where and what type of vulnerabilities exist on a website to narrow the WAF's focus and provide more clearly defined policies.

Lastly, and least common, is a purely manual approach in which each URL, parameter value and overall aspect of a website is defined by hand. While theoretically the most accurate approach, it's also the most time consuming and difficult to manage and therefore impractical for most organizations in the Web 2.0 world.

The remaining piece, policy enforcement, is a challenge because it must be extremely accurate, otherwise attacks might slip by or perhaps worse for many, the WAF may block legitimate traffic from reaching a website. When a request to the website is received, the WAF has three enforcement options: allow, block, or alert. These options are straightforward and should speak for themselves. The following quote should provide context for how most WAFs are typically configured in the field:

"When you know nothing, permit-all is the only option. When you know something, default-permit is what you can and should do. When you know everything, default-deny becomes possible, and only then."

Economics and Strategies of Data Security

The lesson is: no matter how proficient "learning" technology becomes or how diligent a person manually enters policies, the process is unlikely to be either perfect or trusted enough for the business to rely on it solely. This is precisely why only a razor thin percentage of deployed WAFs are configured in default-deny mode, as opposed to their network firewall cousins. Most WAFs are deployed in either permit-all mode—where they only alert on suspicious traffic—or default-permit mode, blocking only a strict few of the most common and easy to identify web attacks. The rest easily pass through.

What's key to realize is that our opponents are intelligent, much more so than any computer or learning algorithm. Our security solutions must be able to keep pace with their capabilities. If we know the weak spots in a website, we should focus our time and energy on attacks targeting the areas where we are vulnerable instead of where we're not. Instead of taking a heavily restrictive policy up front and then relaxing it to allow the website to function, we can create a solid minimum security baseline, and then ratchet up security based upon the vulnerabilities the website actually has. Vulnerability assessment and web application firewalls have a way of complementing each other where the whole is greater than the sum of its parts.

The Solution: Total Website Security

WhiteHat Sentinel + F5 BIG-IP Application Security Manager (ASM)

Combining vulnerability assessment results with firewalls has been used successfully in the past, just never in web application security. Kavado (defunct) tried in 2002-2003; and other vendors tried again in 2003-2004 using an open standard called AVDL. Ultimately, all proved unsuccessful due to a few obstacles. Chief among them is that commercial scanning products would dump hundreds or even thousands of unvalidated results loaded with false positives and duplicate vulnerabilities into WAFs. The implementations not only heavily slowed WAF performance, but blocked other business critical traffic from accessing a website. Practically speaking, no fully automated solution is capable of the level of accuracy required to create WAF policies safely deployable in default-deny mode.

Only recently have web application vulnerability assessment solutions, specifically through WhiteHat Sentinel, matured to the point where the combined solution has become truly viable. With people, process, and technology we've overcome the hurdle of obtaining highly accurate vulnerability data that can be made actionable. At the same time, today's WAF products, specifically the F5 BIG-IP ASM, are significantly more technologically advanced than in years past. WAFs have become easier to set-up, manage, and also integrate with. The best part is when WAFs are armed with timely and precise vulnerability intelligence; they can actually be deployed reliably in block mode.

WhiteHat Security's flagship offering, WhiteHat Sentinel, is the only website vulnerability management solution that enables organizations to address all website vulnerability issues with accuracy and confidence. As a web-based subscription service, WhiteHat Sentinel combines advanced proprietary precision scanning technology with expert analysis, enabling customers to identify, prioritize, manage, and remediate website vulnerabilities as they occur. This comprehensive, laser-focused approach gives all parties a clear view into the organization's website security posture in an easy-to-manage, cost-effective manner.

WhiteHat Sentinel's Service Oriented Architecture (SOA) was built to assess hundreds, even thousands of the largest and most complex websites simultaneously. This scalability of both the methodology and the technology enables WhiteHat to streamline the process of website security and also enables

rapid and highly accurate identification of new threats, so, you stay on top of the latest attack vectors. WhiteHat Sentinel is also the only solution that assesses the 24 classes of vulnerabilities identified by the web Application Security Consortium (including Cross-Site Scripting and SQL injection).

Organizations can immediately focus on remediation of vulnerabilities as false positives are virtually eliminated by the WhiteHat operations team. Security and development teams are free to focus on real problems instead of validating long lists of possible vulnerabilities. Sentinel's web interface allows authorized users to schedule scans or access reports at any time from any location.

WhiteHat Sentinel provides comprehensive and unlimited assessments to keep up with rapidly changing websites, prioritization recommendations are delivered based on threat and severity levels, and "one-click" vulnerability retesting to confirm that remediations are successful.

The BIG-IP ASM, available as a software module for the BIG-IP Application Delivery Networking system or a stand alone appliance, is an extremely powerful and flexible web application firewall. BIG-IP ASM provides proactive network and application-layer protection from generalized and targeted attacks by understanding the user interaction with the application. Through the F5 iControl® API, WhiteHat Sentinel is able to directly configure precise policies ("virtual patch") on BIG-IP ASM to protect against vulnerability exploits (e.g., Cross-site Scripting, SQL injection) found during the scanning process (Figure 2). From the customer perspective the process is simple and straightforward.

1. Sentinel notifies the customer (via email) to newly identified vulnerabilities reports available
2. Customer may choose to issue a virtual patch on their BIG-IP ASM in either transparent or block mode.
3. When the virtual patch is placed in block mode, the Sentinel re-test button will confirm that the issue is mitigated and if so, close out the issue.

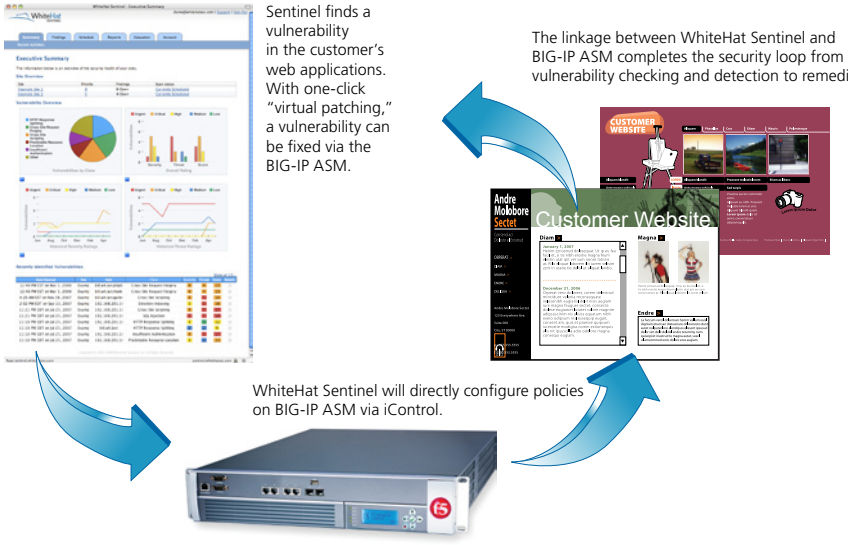

Figure 2: *WhiteHat Sentinel + BIG-IP ASM Integration Flow*

The linkage between WhiteHat Sentinel and BIG-IP ASM completes the circle from vulnerability checking and detection to remediation of specific vulnerabilities. This makes it easier for the user—find the problem, then fix the problem with one click. This integration makes "virtual patching" a reality. The combination of the WhiteHat Sentinel Service with BIG-IP ASM allows for a more sophisticated and precise vulnerability detection and resolution (website security) service. The end result is total website security:

- Highly targeted vulnerability remediation (virtual patching); keeps bad traffic out, allows good traffic in
- Closes the loop from vulnerability detection to remediation
- Find and fix website vulnerabilities with one turnkey solution
- Automated remediation reduces the time, complexity, and cost of protecting customers' web applications without the need for manual intervention
- Reducing false positives while still allowing increasingly high-levels of security

For companies concerned about the upcoming PCI 6.6 deadline, the combined solution fully meets, and exceeds, this requirement of the PCI compliance standards developed by VISA, MasterCard, and other major credit card companies. According to the standard, an organization must do at least one of the following:

- Undergo application scanning and code review by an application security specialist; or,
- Install a web application firewall in front of the web-facing applications.

With the F5/WhiteHat solution, companies can obtain compliance and feel more confident about their website security.

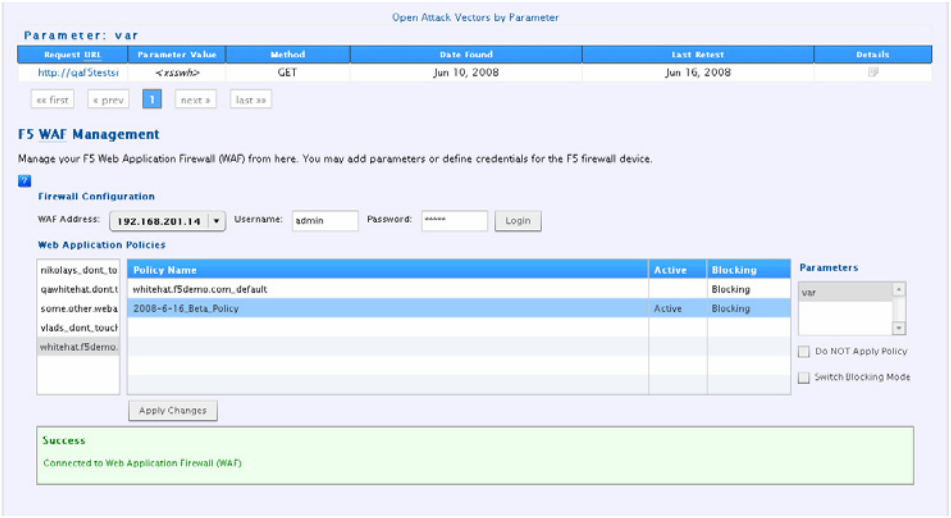

Figure 3: *WhiteHat Sentinel displaying open attack vectors by parameter. 2008-6-16_Beta_Policy is being selected and the policy will be applied to BIG-IP ASM.*

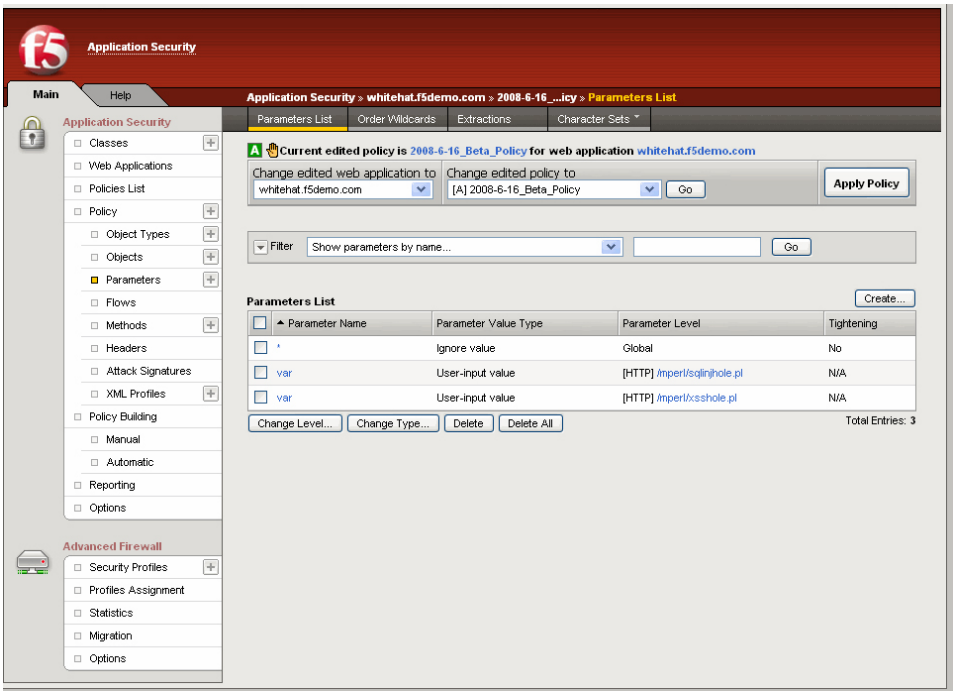

Figure 4: *F5 interface showing the 2008-6-16_Beta_Policy parameters have been applied and are blocking attack vectors.*

Conclusion

Neither software nor developers will ever be perfect or even close to it, leaving the concept of Secure Software as a lofty but unattainable goal. As important as SDLC processes are, they can't always take into consideration unknown attack techniques, current techniques we don't fully appreciate and ignore, or the massive amounts of old insecure code we depend upon already in circulation.

New attack techniques are being published all the time with the existing ones often becoming ever more powerful. To compensate, new, highly accurate solutions are demanded to help organizations identify emerging threats, react faster to them, and adapt better to a constantly changing landscape. Now when a vulnerability or new attack class appears, IT security has a fourth option for the business to consider giving the developers time to fix the code:

1. Take the website off-line
2. Revert to older code (known to be secure)
3. Leave the known vulnerable code online
4. Virtual Patch

---

[1] WhiteHat Security website Security Statistics Report (March)
http://www.whitehatsec.com/home/resource/stats.html

[2] Facing up to the threat of cyber-crime
http://www.continuitycentral.com/feature0555.htm

[3] 70% of websites at immediate risk of being hacked!
http://www.acunetix.com/news/security-audit-results.htm

[4] Sophos: One web page infected every five seconds
http://news.zdnet.com/2424-1009_22-198647.html

F5 Networks, Inc.
Corporate Headquarters
401 Elliott Avenue West
Seattle, WA 98119
+1-206-272-5555 Phone
(888) 88BIGIP Toll-free
+1-206-272-5556 Fax
www.f5.com
info@f5.com

F5 Networks
Asia-Pacific
+65-6533-6103 Phone
+65-6533-6106 Fax
info.asia@f5.com

F5 Networks Ltd.
Europe/Middle-East/Africa
+44 (0) 1932 582 000 Phone
+44 (0) 1932 582 001 Fax
emeainfo@f5.com

F5 Networks
Japan K.K.
+81-3-5114-3200 Phone
+81-3-5114-3201 Fax
info@f5networks.co.jp

WP-WHITEHAT 08/08