F5 White Paper

Manageable Application Security

BIG-IP Application Security Manager v10.1 brings the visibility, control, and flexibility necessary to defend applications and comply with regulations without sacrificing manageability.

by Lori MacVittie
Technical Marketing Manager, Application Services

Contents

Introduction

Manageable is rarely a term used in the same sentence as application security, particularly web application security. Aggressive, agile development methodologies combine with the discovery of new—and variations on old—attack vectors, putting nearly constant stress on information security personnel and the solutions they employ to defend web applications from attacks.

Part of the problem is the amount of information that must be collected, collated, combed through, and summarized for management and business stakeholders as well as web application developers. Information security personnel suffer from information overload on a nearly daily basis, but need to sift through myriad articles, reports, blogs, logs, and scans to do their jobs. Information is vital to successful web application security strategies, whether it be information about a new attack, a new twist on an existing attack, or identifying weak points in web applications. Investments in security solutions have to provide a clear value, which equals additional time spent collecting and documenting proof of this value.

The latest version of F5® BIG-IP® Application Security Manager™ (ASM), v10.1, addresses information overload and the need for agility in implementation. This release of BIG-IP ASM includes a variety of new technological advancements in web application security to assist in web application defense; new reporting and configuration options to make collection and summation of information security data more easily accessible; and enhanced integration with F5 and partner solutions designed to improve the overall deployment time as well as the depth and breadth of information available.

Visibility

To enable a successful web application security strategy, administrators need visibility into attacks, policies, and the threat posture of applications these strategies are designed to protect. Understanding what an attack is and how it works, in addition to the mitigation techniques used by web application firewalls can be invaluable, both enabling protection and educating administrators and developers on the attacks.

BIG-IP ASM v10.1 addresses all three areas of visibility with new and enhanced features designed to provide the information necessary to achieve security goals.

Reporting

Reports are a fact of life in IT. In the realm of web application security, reports are even more important because they often contain information vital to keeping web application security strategies agile and to ensuring compliance with regulations like PCI.

In the past, when administrators needed reports from BIG-IP ASM on attacks, violations, audit trails, and vulnerabilities employed a centralized syslog server. Unless the organization had external systems to handle scheduling or compilation of security specific reports, this process could be painstaking and consumed inordinate amounts of time.

BIG-IP ASM v10.1 introduces a completely new reporting system that provides for scheduling as well as customization. This new version combines an external syslog server for reporting on trends and gathering forensic data with a variety of reporting options. Drill-down options in the new GUI provide an easy way for administrators to explore and find the information they need when they need it, without needing to query external systems. In addition, BIG-IP ASM v10.1 includes 17 pre-packaged reports that provide visibility into attacks and usage patterns.

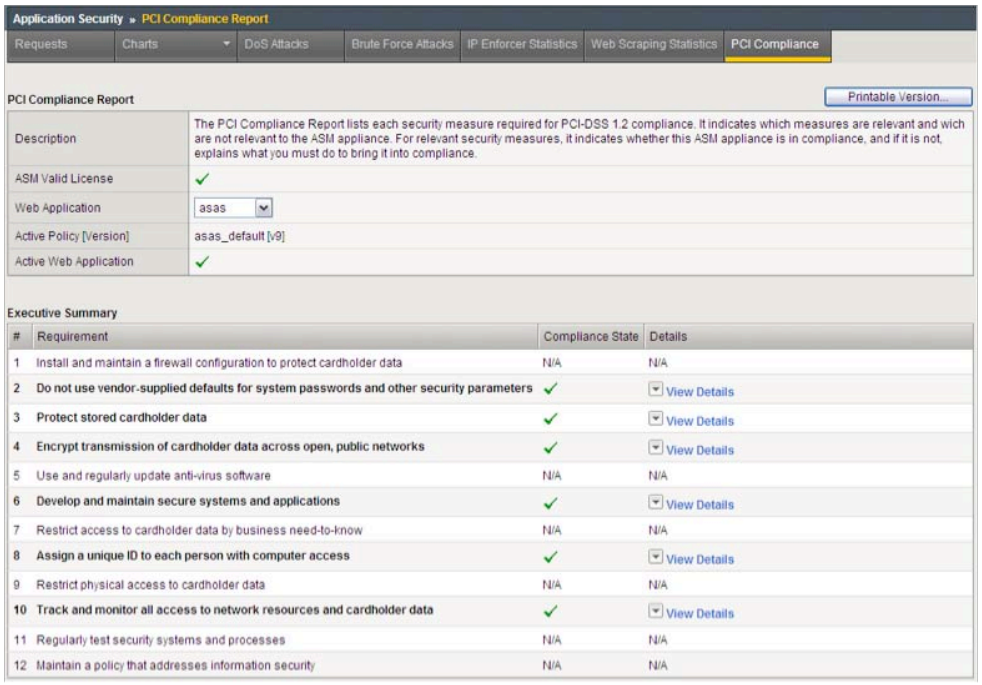

Figure 1: A new PCI compliance report in BIG-IP ASM v10.1 helps organizations understand compliance.

Also, the new reporting system includes highly accurate Geolocation data. These location-based reports provide information about the country from which the threats originated. This combines with the new configurability of BIG-IP ASM to use the X-Forwarded-For HTTP header for reporting and the calculations used to detect brute force attacks and L7 DOS. These capabilities help customers who use services like Akamai where some of the requests are proxied. The ability to configure BIG-IP ASM narrows the identification of violators and, in conjunction with Geolocation, can be an invaluable tool for tightening policies related to the prevention of distributed attacks.

Web Scraping

Web scraping is an old technique that is still used for a variety of reasons. While it was once used primarily as means to "webify" green screens and legacy applications and to provide integration methods for otherwise inaccessible applications, today it is used to harvest e-mail addresses, competitive data, and as a means to find vulnerabilities for later exploitation. Bot activity is often used improve the search ranking of specific terms by automating the process of generating new pages on web platforms. Known as "web spam," this is yet another illegitimate use of web scraping techniques.

Web scraping is, today, more information theft than it is a classical security problem. If a competitor is scraping your content and then publishing it, this can misrepresent your brand or change how the content is presented, which can then negatively impact the business goals for this content. Traditional methods of identifying bots and spiders (such as automated methods) do not accurately detect scraping activity because of the similarities between human users and automated data.

Part visibility enhancement and part protection, BIG-IP ASM's new web scraping detection technology provides valuable insight into these types of attacks. Whether it's spiders or script-driven browsers, BIG-IP ASM provides the means by which such attacks can be detected. While not foolproof, BIG-IP ASM analyzes browsing behavior and characterization to determine when an attack is in progress and can either prevent or simply report on the violation.

Attack Expert System

As threats grow in number and complexity, it becomes increasingly difficult for administrators and information security teams to keep up to date on every attack and countermeasure. BIG-IP ASM v10.1 introduces a new, comprehensive attack expert system that provides an immediate, detailed description of the attack, as

well as enhanced visibility into the mitigation techniques used by BIG-IP ASM to detect and prevent the attack. The expert system also helps network teams—who are often responsible for managing web application firewalls and similar devices—become more familiar with web application security.

Every violation detected by BIG-IP ASM also now includes the risk associated with the violation/check and an example of the attack. Including both risk and example with violation information helps administrators and developers employ a solution based on both the risk and the difficulty of implementation.

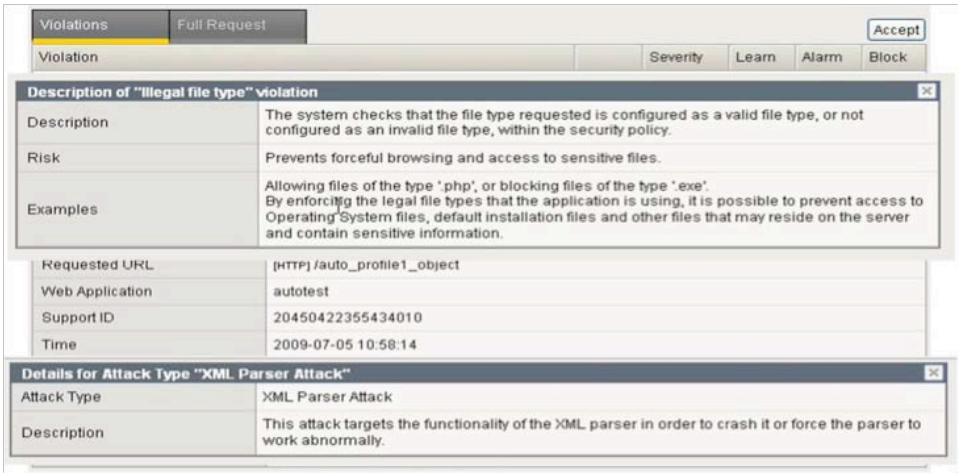

Figure 2: New attack expert system provides detailed information on attacks, risk, and mitigation techniques.

Control

One of the most difficult aspects of managing a web application firewall is dealing with nearly constant policy changes due to frequent website modifications. When policies are automatically adjusted for application changes, they must then be double-checked for accuracy. This process is just as difficult as determining the best way to restrict access to clients. Both tasks require a delicate balance between ensuring the strictest security controls possible without compromising legitimate user access. This means suggested changes must be scrutinized and tested to ensure that violators can be penalized at the most granular level possible without impacting legitimate users.

BIG-IP ASM v10.1 includes additional object types in staged policies, integration with F5 iRules™ for better customization and flexibility, and IP-based penalties for repeat violators.

Staging

The staging of policies is not a new concept in web application security or in BIG-IP ASM. However, BIG-IP ASM v10.1 includes file types, URLs, and parameters. Policies can contain frequently changing objects in a web application. Object-only policies can then be staged for transparent testing in production environments, leaving all other policy entities in blocking mode. Staging enables testing of updated policies until the policy is deemed ready for promotion to a full, blocking application security policy, without reducing current protection levels.

iRules

BIG-IP ASM v10.1 now integrates with iRules, F5's network-side scripting implementation platform. New iRules events that are specific to BIG-IP ASM violations enable administrators more flexibility in how they respond to malicious data and attacks. For example, custom response pages could be returned based on specific violations, providing more information on the attack. Suspicious user input could be manipulated to mitigate the attack, which is useful if the incoming data has been inadvertently tainted or manipulated without user consent. Or perhaps the administrator prefers to force the user to logout upon violation.

Different attack responses vary by organization and specific violations, thus every conceivable response could not possibly be covered by a turn-key web application security solution. Because of this, iRule integration is included in BIG-IP ASM v10 to provide the control over web application security policy behaviors, better addressing the needs of the IT and business organizations.

IP Penalties

With BIG-IP ASM v10.1 it is possible to restrict access from a single IP that generates too many violations over a period of time. This granular violation control provides the ability to restrict or completely block access from the IP addresses with continuous policy violations, potentially indicating an ongoing attack.

Combined with the use of Geolocation and the X-Forwarded-For header, this new feature enables administrators to granularly control client requests and prevent overwhelming internal applications and infrastructure from repeated attempts at unauthorized access or content.

Flexibility

Flexibility is key in web application security because the frequent changes to both the environment and attack methods. A static, inflexible web application security solution provides excellent protection for web applications at a single point in time, but does not afford organizations the ability to adapt to new attacks, new content, and new user requirements.

The ability to integrate with third-party solutions providing complementary services is paramount to success. Integration with the rest of the infrastructure is the hallmark of a dynamic infrastructure and thus enables a more dynamic, adaptable overall security solution. BIG-IP ASM v10.1 enables this type of integration and, in this release, improves its integration with WhiteHat Security's Sentinel for a better overall integration experience.

Integration is not limited to just third-party solutions. When a web application security solution such as BIG-IP ASM is deployed on a Unified Application Delivery and Data Services platform like BIG-IP® Local Traffic Manager™ (LTM), it is expected to integrate well with other solutions on the same platform. BIG-IP ASM v10.1 integrates seamlessly with additional BIG-IP LTM features and modules, supporting a high-performance, unified application delivery infrastructure.

Web application security professionals and management need to be able to examine and audit the security posture of the entire infrastructure, including web application security policies. As these policies can also be integral to compliance efforts, it is important that a web application firewall be able to provide detailed, understandable information to those who may not be familiar with the system. To assist in these goals, BIG-IP ASM v10.1 exports policies in human readable format.

Enabling Offsite Audits

BIG-IP ASM policies can now be exported into a flat, readable XML file. Anyone can read and understand the policies governing web application security without needing physical access to the console or GUI and without requiring extensive education on interpreting configuration options.

This is particularly useful when enabling PCI audit processes with an offsite auditor or without access to BIG-IP ASM. Auditors can view the policies offsite, which frees up IT resources that would otherwise used in assisting them. Allowing auditors to periodically keep themselves up to date on policy changes via exported files also

enables a smoother compliance process by catching potential problems when they are in the staging process rather than after they are in production. This feature also enables external changes to be imported back into BIG-IP ASM.

```
<?xml version="1.0" encoding="utf-8"?>
<policy name="phpauction_default" xmlns="http://www.f5.com/ASM/CP/policy">
  <encoding>utf-8</encoding>
  <web_application>phpauction</web_application>
  <maximum_http_length>8192</maximum_http_length>
  <maximum_cookie_length>8192</maximum_cookie_length>
  <description></description>
  <blocking>
    <enforcement_mode>transparent</enforcement_mode>
    <violation id="EVASION_DETECTED" name="Evasion technique detected">
      <alarm>true</alarm>
      <block>true</block>
      <learn>true</learn>
    </violation>
    <violation id="REQUEST_TOO_LONG" name="Request length exceeds defined
buffer size">
      <alarm>true</alarm>
      <block>true</block>
      <learn>true</learn>
    </violation>
```

Figure 3: Example of a human, readable exported BIG-IP ASM policy in XML format.

Integration Options

Improving integration with other F5 solutions as well as third-party solutions is paramount to enforcing security aspects seamlessly and without negatively impacting application performance.

RamCache integration

RamCache is often used on BIG-IP platforms to improve performance by offloading frequently accessed content to a memory-based cache. Caching of any kind can be problematic for web application security solutions because of the lack of control over the cache and the potential to deliver content in a variety of unauthorized situations.

Deeper integration enables BIG-IP ASM v10.1 to take advantage of RamCache for performance improvements without sacrificing security. BIG-IP ASM v10.1 automatically disables the use of the cache on blocked responses and during the

use of the policy builder in real-time. Additionally, the cache is cleared when a new policy is applied, ensuring that the most current security policies affect requests and responses immediately.

Improved coverage with WhiteHat integration

WhiteHat Security's Sentinel service is used by organizations to provide ongoing vulnerability assessment of web applications. F5's integration with WhiteHat Security enables Sentinel to create BIG-IP ASM policies that prevent the vulnerabilities discovered by a WhiteHat Sentinel scan, giving IT and developers the time necessary to address, deploy, and test fixes to those vulnerabilities.

BIG-IP ASM v10.1 improves upon the existing integration and now provides the ability for WhiteHat Sentinel to create and deploy BIG-IP ASM policies that prevent additional attack methods including:

- Command injection
- XPath injection
- Path traversal
- HTTP response splitting

Conclusion

Web application security is a stressful undertaking with constantly changing and increasingly difficult pressures coming from both inside and out of the organization. Compliance, management, reporting, and the ever-revolving door that is attack vectors puts stress on infrastructure as well as the administrators who must implement security policies to prevent unauthorized access, stop overwhelming attacks, and constantly adjust to new applications and users.

The combined complexity results in web application security that is, at worst, unmanageable and at best, difficult without considerable investment in training and experience with specific products.

F5 BIG-IP ASM v10.1 introduces new features designed to enhance flexibility while improving the control and visibility required by security professionals and administrators. With this new version, BIG-IP ASM brings manageability to web application security.

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
info.asia@f5.com

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

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

CS10084 1209