Key Considerations in Choosing a Web Application Firewall

Today, enterprises are extending their businesses by using more web-based and cloud-hosted applications, so a robust and agile web application firewall (WAF) isn't a luxury—it's a requirement. As these web- and cloud-based applications become more popular, attacks become increasingly sophisticated and frequent, threatening enterprise data.

White Paper

Introduction

Today, enterprises are extending their businesses by using more web-based and cloud-hosted applications, so a robust and agile web application firewall (WAF) isn't a luxury—it's a requirement. As these web- and cloud-based applications become more popular, attacks become increasingly sophisticated and frequent, threatening enterprise data. This makes it far more difficult for administrators and security teams to keep up to date on the latest attacks and protection measures. At the same time, they must meet the stringent compliance requirements for online commerce and data sharing across traditional and cloud environments.

Whether you need a WAF to comply with Payment Card Industry Data Security Standard (PCI DSS) regulations or simply to protect your business-critical web applications from common attacks—such as SQL injection, DDoS attacks, and complex, multi-faceted zero-day attacks—choosing the right product can be complicated. Here's a breakdown of some of the key factors you should consider when selecting a WAF to protect your business.

Deployment Models

Traditionally, WAFs were deployed as hardware appliances on premises in enterprise data centers. But with applications migrating to cloud-based Infrastructure-as-a-Service (IaaS) environments and organizations leveraging cloud Software-as-a-Service (SaaS) apps, security teams are challenged to protect applications beyond the data center—without compromising performance, scalability, and manageability.

In moving toward this hybrid application deployment model, organizations struggle to maintain sufficient control across new infrastructures that provide limited security options for critical web applications residing outside the controlled environment. You can either trust the built-in security of your cloud provider or leverage a cloud-based service to protect all your web applications.

Many organizations will continue to use a hardware WAF appliance to protect critical applications maintained in a traditional data center, but they can also meet their application security requirements using other WAF deployment options.

One is to deploy a WAF as a software-based virtual edition (VE)—a cost-effective option for small-to-medium-size businesses or those wanting to deploy protections closer to the app. Virtual WAFs can also be effective for enterprises that build their own software-defined data center (SDDC) environments or leverage IaaS solutions. Finally, many organizations deploy cloud-based WAF (WAF-as-a-Service) to intercept web traffic before it enters the network or reaches the server in the cloud.

No matter which deployment model is right for your organization, the solution you select should have some essential capabilities.

Basic Considerations When Selecting a WAF

Keep the following in mind when selecting a WAF to protect your web-facing applications—whether they reside in a traditional data center or in the cloud.

Network Architecture and Application Infrastructure

Web application firewalls are designed to watch and respond to HTTP/S traffic. They are most often deployed as appliances in the line of traffic between the requester and the application server, inspecting requests and responses before forwarding them. Inline deployments tend to be most effective in actively blocking malicious traffic based on policies and rules that must be applied judiciously to avoid dropping legitimate traffic.

In this inline model, there are three specific methods that can be used to pass traffic: reverse-proxy mode, router mode, and bridge mode.

- Reverse proxy is the most common mode of operation, terminating all incoming traffic and interacting with the server on behalf of a requestor. As the most feature-rich mode, reverse-proxy is required for most application security capabilities.
- Router mode (a.k.a. transparent mode) is similar to reverse proxy mode, but does not terminate requests intended for the server and offers few services. Frequently, transparent mode is used for traffic logging and reporting.
- In bridge mode, the WAF functions as a layer 2 switch with very limited firewall services.

The mode of operation will be determined by how your application is set up on the network. Other considerations are whether the architecture uses a reverse proxy, if SSL termination is required, and how many WAFs will be deployed.

A WAF can also be deployed "out of band," which allows the WAF to observe traffic from a monitoring port. This non-intrusive "passive" deployment option is ideal for testing the WAF without impacting traffic, yet still enabling the WAF to block malicious requests.

Before selecting a WAF, consider which deployment option best suits your network infrastructure and network environment, and understand the scope of services you will need to use. Also, check what modes of operation the WAF supports, because not all firewalls support every mode mentioned.

Security Effectiveness and Detection Techniques

Today's leading WAFs employ a combination of techniques to ensure accurate detection coverage that does not block legitimate traffic. Traditionally, the most widely used WAF configuration has been a negative security model, which allows all transactions except those that contain a threat/attack. Negative security utilizes signatures and rules designed to detect known threats and attacks. The signature rules database should be quite substantial, as attack knowledge has built up over the years. This is a great model for out-of-the-box protection, blocking commonly known threats, including Web injections, OWASP top 10, XSS and beyond.

In recent years, positive security models have become popular. This approach blocks all traffic, allowing only those transactions that are known to be valid and safe. The positive approach is based on strict content validation and statistical analysis, which can be more effective in preventing zero-day threats and vulnerability manipulation.

To be truly effective, a positive security approach requires deep knowledge of the application and its expected uses. Rules and signatures should cross-reference to the specific OS, application types, and version numbers. When properly configured, this model can prove to be more efficient, as it can use fewer rules than a negative security model.

However, not all manufacturers provide an efficient means to configure and update policies and rules. Users of some WAF offerings may find it difficult to provision, configure, and maintain rules and policies to support a positive security model. Security administrators should look for automation tools that streamline updates and the configuration process, without requiring re-learning or additional training.

Positive and negative models are both capable of achieving the delicate balance between "security" and "functionality." The difference between these models stems from where each begins and where they collide. This can be as simple as the number of rules required to meet the end goal, or even the choice of whether to err on the side of functionality or security if something is missed in the policy. However, neither a positive nor negative security model alone can deliver the most economical solution in every situation or environment. When merged with business requirements, an integrated positive and negative approach can enable organizations to realize the greatest ROI from any security policy implementation.

Overall, the effectiveness of any WAF solution is difficult to discern without testing. NSS Labs is a good resource for gaining insight into the effectiveness of leading WAFs on the market. Consider reviewing the Security Value Map™ (SVM) and Comparative Analysis Report™ (CAR) series by NSS Labs, which evaluates leading WAF products on their ability to prevent intrusions, and detect and mitigate threats.

Performance, High Availability, and Reliability

Web application firewalls play an essential role in maximizing throughput and ensuring the high availability of the application(s) they protect. WAF capabilities should include features that address these factors directly:

- Caching copies of regularly requested web content reduces repeated requests to back-end servers.
- Automatic content compression provides for more efficient network transport.
- Hardware-based SSL acceleration speeds SSL processing and reduces the burden on back-end web servers.
- Load balancing web requests across multiple back-end web servers optimizes performance.
- Connection pooling reduces back-end server TCP overhead by allowing multiple requests to use the same back-end connection.

Typically a reverse-proxy deployment will support each of these features. However, not all WAF solutions have these functionalities, and many are dependent upon the deployment mode chosen, so check vendor specifications.

Virtual Patching and Scanner Integration

Web application vulnerabilities are among the most common causes of data breaches. Vulnerabilities—unique to each application—leave companies' web infrastructures exposed to attacks such as cross-site scripting, SQL injections, cookie poisoning, and others. And although developers apply best practices in secure coding, and perform adequate security testing of applications, all applications are prone to vulnerabilities. Given this, and the time it takes to release new software updates, additional tools are needed to detect, validate, and patch software exposures until new application code is made available.

When defects are found in software code, organizations with a WAF can rapidly apply fixes (virtual patches) to prevent exploitation by an attacker. Virtual patching requires no immediate changes to the software, and it allows organizations to secure applications immediately—and in some cases, automatically—upon dynamic application testing. Virtual patches are a key component of a strong WAF, often requiring integration with a vulnerability scanner. The level of WAF–scanner integration varies with WAF vendor, and not all cloud-based WAFs integrate seamlessly. Make sure to pick a WAF solution that seamlessly integrates with leading scanner technologies; otherwise, you'll need to manually configure policies in place to address any vulnerabilities, which is a time- and labor-intensive process.

PCI DSS Compliance

Malicious attacks designed to steal sensitive credit card information are increasing, with more and more security breaches and data thefts occurring daily. The PCI DSS requirements have been revised in an attempt to prevent these types of attacks and keep customer data secure. If your organization works with, processes, or stores sensitive credit card information, you must comply with PCI DSS requirements. You must strengthen your security posture by protecting your critical web applications, which are often easy pathways for malicious attackers to gain access to sensitive cardholder data.

While you can adhere to PCI DSS standards by deploying a vulnerability scanner or a WAF, the most effective solution is to integrate the data from scanning technology with the attack-mitigation power of a WAF. The best WAFs can identify, isolate, and block sophisticated attacks without impacting legitimate application transactions. In addition, some WAFs offer PCI reporting, which determines if compliance regulations are being met, and if they are not, details the steps required to become compliant.

Protection Against Application Attacks

With the continued growth of multi-layered attacks, IT managers need a strong web application firewall solution. A good WAF ensures application security and availability by providing comprehensive geolocation attack protection from layer 7 DDoS, SQL injection, OWASP Top Ten application security risks, cross-site scripting, and zero-day web application attacks. It also can prevent execution of fraudulent transactions, stop in-browser session hijacking, and secure AJAX applications and JSON payloads, but not all WAFs provide complete coverage in these areas. Note, however, that some WAFs do not defend against content and cookie modification, brute force login attempts and HTTP Parameter Pollution attacks. When evaluating a WAF, make sure you understand the full scope of protections it offers to ensure that your business receives the best coverage.

Data Classification of Protected Applications

If you are thinking about deploying a WAF, you should also consider whether the information that flows to and from your applications will be encrypted. More and more attackers are encrypting their attacks—and traditional security appliances like intrusion prevention systems (IPS) cannot see inside the encryption.

Your WAF solution needs to be able to understand the application and the data that it is protecting. If that data is encrypted, your WAF must be able to decrypt the information and then classify the data within the apps in order to provide additional protection. A strong WAF can terminate SSL traffic, expose what is inside it, and make security decisions based on the encrypted data.

Visibility and Reporting

Protecting your web applications and mitigating threats are two of the essential requirements of a WAF; a third is that the solution gives your organization the ability to collect and analyze the data so that you have a better understanding of the current threat landscape—and how secure your applications are.

Reports provide visibility into attack and traffic trends, long-term data aggregation for forensics, acceleration of incident response, and identification of unanticipated threats before exposure occurs. Many WAFs also integrate with database security products to give administrators a real-time view into the operation of their websites, and provide reports on web-based attempts to gain access to sensitive data, subvert the database, or execute DoS attacks against the database.

Advanced Considerations When Selecting a WAF

In addition to the basic functionality every WAF should offer, there are other capabilities that can influence your decision about which solution to deploy.

Automatic Attack Detection

Inbound automated attack or botnet traffic such as DDoS and malware activity can penetrate security layers and consume valuable processing power. When a system is under attack, bot detection investigates whether a web client source is human, an automated browser script, or even a headless browser.

A strong WAF extends bot-defense capabilities to deliver always-on protection—preventing automated layer 7 DDoS attacks, web scraping, and brute force attacks from ever materializing. This proactive approach to detection identifies more evasive bot sequences that may escape traditional detection methods, and identifies unauthorized, automated attacks upon the first attempt to access an application.

Again, these capabilities are not available from all manufacturers. At a minimum, a WAF today should be able to detect attacks designed to run JavaScript, respond to challenges, and mimic human and browser capabilities without overburdening the applications it protects.

Device ID and Fingerprinting

Browser fingerprinting captures browser attributes in order to identify a client. This is a great way to identify or re-identify a visiting user, user agent, or device. This persistent identification of a client is important in that it allows tracking across sites. Attributes can be very revealing, enabling you to draw inferences about visitors, track users across origins, and share information, all to identify repeat offenders.

Browser fingerprinting techniques vary, but they commonly include passive, active, and cookie-like approaches. However, the effectiveness of the technique lies within the quality and quantity of attributes collected, as well as the ability to protect against tampering. Additionally, any fingerprinting mechanism should include a secure channel for data transmission and a method for hiding any associated security keys and algorithms.

Fingerprinting-based identification is not always reliable and may not work with all device or browser types. Check with your WAF vendor for a list of supported devices/browsers, specific features supported, a list of attributes collected, and what information is reported (e.g., the number of cookies deleted, unique data found).

SSL Offload

SSL processing can put a strain on application resources. Offloading SSL computation to other network resources allows applications to dedicate important CPU resources to other processing tasks, which can improve performance. WAFs that support SSL offloading maximize the utilization of the applications they protect, eliminate the need to buy additional hardware, and increase the value of the WAF itself. Make sure that the WAF you’re considering can offload that processing work to keep everything running smoothly.

Behavioral Analysis

Some WAFs can analyze and understand volumetric traffic patterns and then scan for anomalous behavior based on a set of related rules. A good WAF also assesses average server response time, transactions per second, and sessions that request too much traffic to use as a baseline for determining whether an attack has commenced. Such anomalous patterns uncover attacks that can go undetected by some WAF technologies.

In addition, a WAF should be able to detect an anomaly when either too many sessions are opened from an IP address or when the number of sessions exceeds a set threshold. Strong behavioral analysis capabilities can make it easier for your organization to predict, identify, and respond to attacks.

Security Operations Center

Having a strong connection with your WAF vendor's security team is essential, especially if you're using a cloud-based WAF. A knowledgeable and responsive security team includes experts who analyze threats and malware, and who reverse engineer code to uncover how attacks work—and how they can best be mitigated. Some WAF vendors now provide 24x7x365 support that helps you to stop online fraud by identifying the source of threats and taking down sites operated by cyber criminals. Up-to-date intelligence comes from analysis of malware and assessment of financial fraud information from a variety of sources to provide you with detailed reports of attacks on your business.

Your WAF vendor should also understand the global threat landscape as it relates to the relationship between their product and your environment. They should work with you as a consultant to not only mitigate threats as they arise, but also to enhance your organization's own security practices.

Anti-Fraud Capabilities

Evolving and sophisticated fraud threats that target businesses and their customers have made effective online fraud protection an essential part of enterprise security. Many stand-alone solutions that protect against financial malware, Man-in-the-Browser attacks, and account takeover can require changes to already existing applications and firewalls—and often do not provide adequate visibility into violations.

More advanced WAF solutions integrate with web fraud detection services to simplify deployment, streamline reporting, and strengthen the overall application security posture by thwarting requests from validated fraudsters. These integrated services should enable organizations to rapidly respond to threats at the network and application level.

WAFs should efficiently and accurately correlate application attacks—including web scraping, and DDoS, brute force attempts—with client-side attacks targeting end users. Moreover, a good WAF should allow you to easily understand the full scope of the fraud threat across the network, application, and user.

Ease of Management

Deploying a WAF used to be a difficult and time-consuming exercise in configuring and implementing manual rules. The strongest solutions now simplify policy creation so you can deploy your WAF with security policies that immediately address common attacks on web applications, including HTTP(S) attacks.

Some cloud-based WAFs can deliver a simplified approach to deploying policies across a growing WAF infrastructure that includes traditional and cloud environments. With the automatic learning offered by some WAFs, as well as the centralized deployment and management of WAF policies, you can reduce IT overhead, minimize configuration errors, and ensure the overall effectiveness of each policy. This protects web applications no matter where they reside in your network or across hybrid cloud. In addition, centralized policy management allows you to compare policies and evaluate their effectiveness across different firewalls, thus strengthening your overall security posture.

Scalability and Performance

Organizations need to ensure application availability, even when under attack. The best WAFs can help you dynamically boost performance with application optimization and acceleration technologies like fast caching, compression, SSL offloading, and TCP optimization. An enterprise-grade WAF, with robust appliances and through centralized management, can easily scale to handle large volumes of traffic. In addition, cloud-based WAFs can be deployed on demand to achieve seamless and limitless scalability, resulting in better performance, faster response times, and cost efficiencies.

Vendor Release Cycle

You should also ask your WAF vendor about their release cycle. With the threat landscape changing so quickly and dynamically, vendors that offer more frequent release (quarterly vs. annually, for example) can help decrease your window of exposure and reduce the risk of your applications becoming compromised by a new or emerging threat. In addition, a good WAF can reduce the labor involved in updating signatures by providing automatic signature updates in addition to the manual or scheduled option.

Conclusion

Organizations that deliver today’s rich and complex Internet content to users without having adequate security incur significant risk and are exposed to a variety of potentially malicious attacks from rapidly changing IP addresses. By deploying a strong web application firewall, you can secure your critical web applications wherever they reside—within a virtual software-defined data center (SDDC), managed cloud service environment, public cloud, or traditional data center.

A powerful WAF solution enables organizations to protect against OWASP top 10 threats, application vulnerabilities, and zero-day attacks. With strong Layer 7 DDoS defenses, detection and mitigation techniques, virtual patching, and granular attack visibility thwart even the most sophisticated threats before they reach your servers. A good WAF also enables compliance with key regulatory standards like HIPAA and PCI DSS.

For more information on how F5 Networks can help your organization protect your web applications—and your business—visit f5.com.

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

Americas
info@f5.com

Asia-Pacific
apacinfo@f5.com

Europe/Middle-East/Africa
emeainfo@f5.com

Japan
f5j-info@f5.com