FORTINET

How to Choose a Next-Generation Web Application Firewall

Table of Contents

Executive Overview

The more companies rely on web applications and application programming interfaces (APIs) to support basic business processes, the more crucial web application firewalls (WAFs) become for protecting corporate data and operations against the latest advanced threats. Not just any WAF will suffice, but rather organizations need to ensure it also protects APIs and includes advanced capabilities.

As a starting point, a WAF must include core features such as antivirus and malware protection, a signature engine, IT-reputation checks, and protocol validation. And while application learning for behavioral threat detection is also a critical element, it creates serious security challenges in the form of false-positive rates that impact overburdened security teams. To minimize these false positives and the staff time they require, a WAF must also leverage artificial intelligence (AI) and machine learning (ML) to increase the accuracy of threat detection. In addition to the above, a WAF should integrate into the broader security architecture.

Web Application Security Challenges

Companies rely on web applications—on-premises, in the cloud, or both—for all sorts of functions. When these applications fail, the disruption not only affects corporate operations but also ripples up and down the supply chain. Web applications also access and process critical data such as customer and financial data. Without vigilant protection, that data may be compromised.

Greater dependence on web applications increases the need for a WAF to protect those applications (and their associated API interfaces) against external and internal threats. This is due to the fact that the corporate attack surface is rapidly expanding, while at the same time the volume and variety of cyberattacks grows each day. Web application attacks are involved in 26% of all breaches, making the second most common attack pattern.[1] An organization that only depends on firewall and intrusion prevention system (IPS) security is inadequately prepared to detect and repel the latest advanced attacks—including those that are polymorphic and/or that simultaneously employ multiple attack vectors.

Companies need a WAF that effectively identifies and protects against both known and unknown exploits while minimizing false positives, which may dilute the resources available for threat response. Some WAF technologies struggle to meet this objective.

Chief information security officers (CISOs) looking to improve protection for their organization's web applications need to find a WAF that offers:

- Excellence in core WAF capabilities
- Sophisticated behavioral threat detection with minimal demand on management resources
- Scalability that does not reduce throughput

The global average cost of a data breach is currently $4.45 million.[2]

The web application firewall market was valued at roughly $6 billion in 2021 and is expected to reach over $14 billion by 2027, at a compound annual growth rate (CAGR) of 15.91%.[3]

Cloud Security Lacks API Integration

To provide effective cloud protection for using Software-as-a-Service (SaaS) applications and B2B communications tools, organizations need tools that natively integrate into the cloud in order to run in the same elastic and distributed way that cloud applications run. Traditional security tools cannot offer this kind of functionality when operating as a cloud overlay solution. This is due in large part to the fact that these solutions lack the ability to integrate with the cloud API.

An API is the interface that connects an organization to an application or service. The security and availability of cloud services depend on the integrity of its APIs. These interfaces should be designed with security at the forefront, to protect against malicious attacks as well as accidental insider threats.

There is growing awareness of API security. It is now commonly listed as a top priority CISOs are preparing to, or currently trying to, address. But despite this desire, successful API attacks continue to occur, as each API brings its own unique attack vector and basic application security solutions fail to meet the challenge.

87% of enterprises have built up to a multi-cloud strategy.[4]

Whether you have a single public cloud or multi-cloud scenario, proper integration is a must. In order to perform vital functions, including proper configuration and monitoring of cloud resources, security analytics and management tools need to have the ability to integrate into public cloud APIs. Without API integration, cloud application security inherently lacks consistent visibility of serious problems, such as misconfigurations that expose the organization to potential attacks and also possible regulatory compliance violations. A WAF solution's API security capabilities should include:

- Visibility that analyzes the organization's mobile and web applications to discover, categorize, and secure all APIs
- Centralized security and API management controls that help prioritize protection according to quantified risks
- Distributed enforcement that can protect APIs across the entire architecture, not just at the perimeter
- Mobile device authenticity verification to also protect mobile APIs from malicious attacks

API exploitation by attackers is accelerating. According to a recent Forrester report, malicious API traffic has risen by 117%.[5]

Insist on Best-in-Class Core Capabilities

Because attacks on web applications and APIs are so varied, security needs to employ a combination of protection approaches. These approaches not only must be varied to combat the diverse attack vectors but they also need to be correlated. And the WAF should take appropriate action to protect web applications and APIs whenever it detects threats. Following are core capabilities that security architects need to seek in a WAF solution:

Antivirus and malware protection

The most fundamental of security products, an antivirus and malware-detection engine is a crucial building block for any successful WAF. The engine needs to scan all web application traffic for threats that could potentially infect servers and other devices on the corporate network.

Frequent signature updates

Another capability included in every WAF is signature detection, which compares the contents of incoming packets against the signatures of known web attacks. These can include botnets, advanced threats, and distributed denial-of-service (DDoS) attacks.

To provide effective signature detection, a WAF requires:

1. Resources of a large and reputable threat research organization.

2. The ability to incorporate threat research insights into its signature-detection database. Ideally, these updates will flow into the WAF in real time.

3. The faculty to either redirect potentially malicious packets to a sandboxing tool or else block them from the corporate network.

IP-reputation verification

Like signature detection, IP-reputation checks compare incoming traffic against known threats. The difference is that they look not at the content of the incoming packets but at their source. The WAF maintains a blacklist of IP addresses known to be

associated with delivery of botnets and other types of attack. The WAF compares traffic hitting protected web applications against its malicious source blacklist, and when it detects a match, it prevents the associated traffic from entering the network.

As with signature detection, excellence in IP-reputation checks requires a WAF to tie into a threat-intelligence service that provides frequent updates to the blacklist. The more sophisticated WAFs are also able to identify and blacklist the sources of malicious packets tagged by their signature-detection engine.

Protocol validation

A WAF must also be able to root out improper HTTP code. When applications that use different communications protocols interact, they create a vulnerability. An attack might be able to bypass strong security measures in each application by mimicking the other application's protocol. Such an exploit could bypass even strong security measures by feigning translation errors. All web-based applications should comply with HTTP RFC specifications.

To nip prospective protocol exploits in the bud, a WAF needs to validate the protocol of any code that protected web applications try to execute.

Integration of capabilities

Each of these WAF capabilities alone is key to protecting web applications from one or more common types of exploits. To optimize protection, the WAF should integrate these functions in two ways:

- **Data correlation.** Data on application-layer signatures, malicious bots, suspicious IP addresses, and emerging viruses should be correlated so that threat intelligence is shared across capabilities. For example, when the WAF identifies a botnet, it can add the botnet's originating IP address to its IP-reputation blacklist, automatically flagging any future traffic coming from that address.
- **Intelligence sharing.** The WAF should also fit seamlessly into the organization's broader security architecture. Many cyberattackers employ polymorphic malware and simultaneously take multivector attack approaches. Combating such threats requires real-time intelligence sharing across all network components. For example, an attack might probe vulnerabilities across multiple vectors (for example, endpoints, email, cloud services) and employ ML to hone the exploits based on information learned. In these instances, the WAF needs to support real-time, two-way sharing of threat intelligence with each of the corresponding security elements in order to successfully thwart these kinds of advanced threats.

Flexibility of deployment

Deployment flexibility should include the ability to protect applications wherever they occur across an organization's hybrid environment. This may require a variety of WAF form factors—including physical appliances, virtual machines (VMs), public cloud instances, and SaaS.

In addition, the ongoing cybersecurity skills shortage has yielded over-burdened and under-resourced security teams worldwide. The option of delivering WAF-as-a-Service (WAFaaS) provides a model that can reduce deployment time from hours or even days down to minutes. WAFaaS not only alleviates initial set-up and configuration of WAFs, it also supports elastic scalability for growing and/or distributed businesses. WAFaaS can also help reduce operational churn for ongoing maintenance and management to further reduce demands on limited staff resources.

There are almost 3.5 million unfilled security positions worldwide today.[6]

Harnessing Artificial Intelligence for Threat Detection

Blacklisting and whitelisting security technologies may catch a significant proportion of threats, but they are only as good as the lists they rely upon—namely, they can identify only previously recognized exploits. With security providers unable to create signatures for unknown threats, traditional security approaches are unable to prevent and detect emerging and zero-day attacks.

In addition to list-based monitoring capabilities, most WAFs incorporate behavior-based threat detection, an approach to web application security that compares the actions of users or applications against expected behaviors to recognize and flag anomalies. Solutions that incorporate application-learning technologies monitor responses to certain inputs over time and extrapolate—or "learn"—what responses they should expect to receive in the future.

These WAFs automatically build a profile of the structure of a protected application, as well as how the application is used in the organization. Then, they associate rules for threat response to characteristics of these profiles. Behaviors that trigger an alert might cause incoming web application traffic to be blocked entirely or to be routed to a corporate sandboxing tool.

Challenges of application learning

WAFs with these capabilities incorporate application "learning," but that does not mean they are very intelligent. They develop the profile for a protected application by observing data entries and other facets of user behavior as it relates to each parameter of the application, including value ranges for form fields, HTTP methods, cookies, etc. And they do continue to update these profiles over time, as they gather more and more data on user behavior.

The problem is that any behavior which does not fit into a WAF's specified profile—in other words, any behavior that the WAF has not previously observed—triggers an alert. This creates an exorbitantly high rate of false positives in the WAF's threat detection. Anytime a new data trend in user behavior emerges, application traffic may be blocked

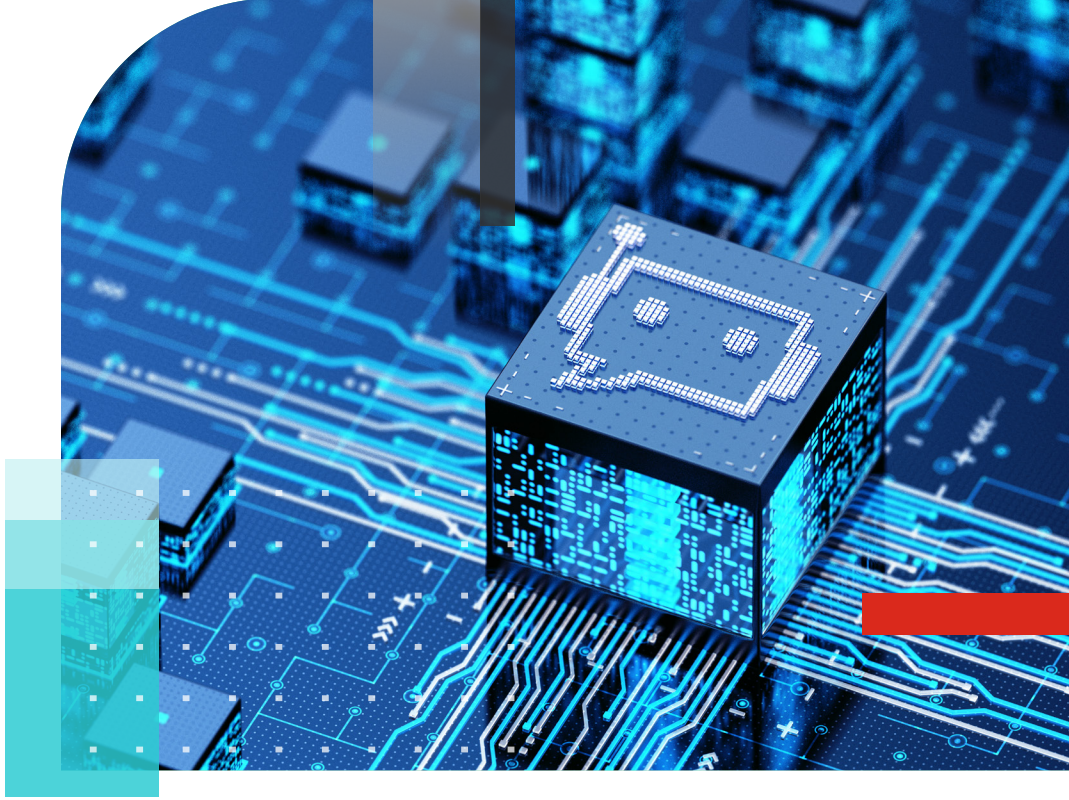

until a human can review and decide that it is not actually a threat. Over time, the new behavior will become expected, but many actions that present no threat to the organization get flagged and require manual follow-up processes.

Benefits of AI-based ML in a WAF

Companies with limited security staff should look for ways to reduce the often-high resource requirements entailed in managing a WAF's application-learning capabilities. One area where WAF providers can turn is true AI-based ML.

ML can enable a WAF to automatically detect malicious web traffic. It can view deviations from normal user or application behavior, not as an immediate cause for alarm but as a context for considering security concerns. Potential alerts are not evaluated from the perspective of a simple “yes” or “no” rule violation; they instead inform the WAF’s automated calculation of the probability that a user or application behavior represents a threat requiring a security response. The ML process should apply a two-layer approach of first identifying whether a request is an anomaly, and if so, whether the anomaly is then an attack.

Few WAFs have incorporated this type of ML. Those that have done so are able to respond to behavioral anomalies according to predefined rules depending on the threat likelihood determined across multiple parameters. This can virtually eliminate the false-positive problem created by application learning. Thus, unlike organizations that rely on WAF application learning, those that use ML can avoid allocating valuable staff resources to resolve false positives.

Moreover, ML enables the WAF to classify files and data sources much more accurately. Combined with core WAF capabilities, ML can detect almost all legitimate threats. This helps protect the network against scanners, crawlers, scrapers, credential stuffing, and a host of unknown attacks.

Bot mitigation

ML capabilities that utilized real-time threat intelligence are especially effective tools for bot mitigation. Bots are applications that run automated tasks over the internet. “Good” bots may include search engines, virtual assistants, and chatbots.

"Bad" bots perform invasive tasks, including web scraping, competitive data mining, harvesting of personal and financial data, account takeover, spam, and transaction fraud. Bots' sources can be very difficult to identify and trace—which has accelerated their proliferation as a tool for harm. Therefore, filtering out automated bad bot threats is now a critical need for effective cybersecurity. In this regard, a WAF solution's ML capabilities should include:

- **Biometrics-based detection** to verify whether a client is a bot by monitoring for things like mouse movement, keyboard, screen touch, and scroll
- **Threshold-based detection** that helps to define specific elements of a suspicious behavior, such as the time of day when it occurred
- **Bot deception tools** that insert a hidden link into response pages that trick the bot into revealing itself
- **Mobile application identification capabilities** that automatically verify that a request is legitimate by verifying the token a mobile application carries when it accesses a web server

Keeping up with DevOps

ML capabilities should also allow WAF security to keep up with the speed of development operations (DevOps) to provide advanced protection without impacting availability or performance. Many organizations end up forgoing or disabling DevOps security due to performance bottlenecks created in order to ensure protection. But this leaves the organization exposed to any number of sophisticated attacks.

In a recent finding, bad bots were responsible for over 95% of the malicious web traffic detected in a test of almost 40,000 unique IPs.[7]

Performance and Operational Considerations

Clearly, a WAF's ability to protect the network's web applications and APIs is a key factor in the CISO's solution research. But these are not the only considerations.

- **Throughput.** As business-critical as security may be, few organizations can afford to have traffic slowed down when their WAF conducts comparisons against blacklists, whitelists, and behavioral profiles. Further, security architects evaluating WAFs need to understand not only typical throughput for their different options but also the characteristics of their network and security architecture that might reduce each device's throughput in their unique environment.
- **Scalability.** Related to throughput concerns are the struggles some WAFs have in supporting a large volume of web application traffic. Most companies will continue to see their data volumes grow rapidly for the foreseeable future. Thus, their WAF needs to be scalable enough to support not only the organization's current traffic volume at its desired level of throughput but also its anticipated future web application traffic.
- **Administrative resources.** In addition to the massive amount of staff time that false positives in application learning can consume, WAF buyers should consider each solution's ease of use, as well as how much effort the security team must dedicate to configuring and fine-tuning threat-response rules.
- **Reporting and compliance.** The reporting provided by a WAF needs to comply with all the appropriate regulatory requirements, such as National Institute of Standards and Technology (NIST) 800 security controls, the Payment Card Industry Data Security Standard (PCI DSS), etc.

Evaluating WAF options against all these criteria takes time, but the huge potential to effectively and efficiently protect web applications, APIs, and data makes the process well worth the effort.

[1] 2023 Data Breach Investigate Report, Verizon, 2023.[2]

[2] "2023 Cost of a Data Breach Report," IBM Security/Ponemon Institute, July 2023.

[3] Global Web Application Firewall Market Scope and Overview Report 2023-2030, Market Reports World, August 2023

[4] 2023 State of the Cloud Report, Flexera, 2023.

[5] The State Of Application Security, 2023, Forrester, June 2023.

[6] Cybersecurity Jobs Report: 3.5 Million Unfilled Positions In 2025, Cybercrime Magazine, April 2023.

[7] Botnets responsible for nearly all malicious web traffic, Techradar, July 2023.

FORTINET

www.fortinet.com

August 25, 2023 11:39 PM

486237-B-0-EN