Guide

WAF Buyer’s Guide

This checklist will guide you in asking the right questions and evaluating options to choose a WAF that meets your needs now and scales with you in the future.

fastly®

Introduction

Applications and APIs are the backbone of modern businesses. They power customer interactions, process sensitive data, and drive revenue. However, the growing reliance on these digital assets makes them attractive targets for cyberattacks. Web application firewalls (WAFs) are built to secure applications and APIs, but because they've been around for decades, there are massive discrepancies between vendors. To make matters worse, buyers are bombarded with marketing claims and technical jargon, making it difficult to truly differentiate between them.

The fear of making a poor investment or missing crucial features is a significant concern, and this buyer's guide aims to bridge the information gap and empower you to make a confident decision.

WAF buyer’s guide checklist

Table stakes: ‘must have’ features

- □ Protects against OWASP Top 10 Threats
- □ Supports IP/CIDRs, GEO & ASN allow and block lists
- □ Granular and hierarchical policy enforcement

Read the details →

Additional WAF capabilities

- □ Bot mitigation to protect against unwanted automated traffic
- □ DDoS protection to safeguard against denial-of-service attacks
- □ API security controls that support your API formats
- □ Threat intelligence for proactive security measures

Read the details →

Deployment flexibility and speed

- □ Flexible deployment models including cloud, containers, edge, on-premises, and hybrid
- □ Proven track record of rapid deployments for faster time-to-value
- □ Automated deployment through Infrastructure as Code (IaC)

Read the details →

Usability that improves productivity

- □ User-friendly with easy configuration and management
- □ Visibility, insights, and faster decisioning - no more “black boxes”
- □ Simpler rule building
- □ Minimal need for continuous configuration, tuning, and maintenance
- □ Pre-built integrations for DevOps and security toolchains

Read the details →

Scalability and performance

- □ Scalable to accommodate traffic spikes and application growth
- □ Low latency and high throughput
- □ Global network presence for geographically distributed users
- □ Advanced security use cases built on edge computing

Read the details →

Accuracy

- □ Low false positive rate to not break apps
- □ Thresholding as a method to more confidently move to blocking mode

Read the details →

Additional considerations

- □ Single platform for vendor consolidation, unified management, and cost savings
- □ Strong customer support
- □ Clarity around professional services engagements (no surprises!)
- □ Team-augmenting managed services options
- □ Competitive and flexible pricing options
- □ Purchasable using existing credits in cloud marketplaces like AWS and GCP

Read the details →

Learn more about why it’s important

Table stakes: ‘must have’ features

A few starting requirements are necessary for a WAF solution to even be considered. The OWASP Top 10 outlines the top ten most critical risks for web applications. If a WAF can’t identify and block the OWASP Top 10, it leaves you vulnerable to the most dangerous threats.

Other essential WAF features include support for IP/CIDRs, GEO, and ASN allow/block lists. These features let you take broad strokes against malicious traffic by allowing or blocking traffic based on IP addresses, geographic locations, or Autonomous System Numbers (ASNs). This reduces the burden on your security team by minimizing the need for a multitude of complex, granular rules. However, granular policy enforcement remains crucial. A WAF should allow you to define rules with varying levels of detail, applicable globally, for groups of domains, or for individual domains. This ensures a balance between efficiency and customization for different applications or regions.

Additional WAF capabilities

A comprehensive solution offers a layered defense against various web applications and API security threats. Some vendors have adopted the term web application and API protection (WAAP), first coined by Gartner, to refer to their solution. Most WAAP platforms begin with a WAF and other capabilities are either included or provided as add-on components. Depending on your organization’s needs, you may need all of these capabilities or a subset. For the sake of this buyer’s guide, we will consider all components below as part of the overall WAF solution:

Bot mitigation

Protects against automated bots that can scrape data, launch denial-of-service attacks, or engage in credential stuffing while allowing good bots and human traffic. Bot mitigation employs various techniques like CAPTCHA and JavaScript challenges, client fingerprinting, and IP reputation checks to identify and block automated bot traffic. Bot mitigation solutions should have the granularity to differentiate between good bots (e.g., search engine bots) and bad bots (e.g., scraper bots).

DDoS protection

Safeguards your applications and APIs from Distributed Denial-of-Service (DDoS) attacks that overwhelm systems with traffic, causing outages. WAF solutions can mitigate DDoS attacks by filtering malicious traffic, absorbing attack traffic volume, and maintaining application availability. Attacks can come at different layers in your network, so a solution that offers Layer 3/4 and Layer 7 DDoS protection will provide greater protection than just a Layer 7 solution.

API security

Many modern applications rely heavily on APIs to connect and exchange data. WAF solutions should offer specific security controls to protect APIs, including authentication, authorization, and API traffic monitoring. With API security, you'll want to ensure the WAF supports your API formats (REST, GraphQL, gRPC, etc).

Threat Intelligence for proactive security measures

While not a core component, threat intelligence is a valuable addition to a WAF solution. It provides real-time insights into evolving cyber threats and attack methods with the goal of allowing security teams to be proactive in their defense.

First-party IP reputation intelligence feeds, updated daily, offer more accurate and stronger security. This prevents yesterday's malicious activities from affecting today's legitimate traffic, especially from shared IPs. When evaluating IP reputation intelligence, consider whether it's presented as a binary (attacker/legitimate traffic) or a risk score. While both are additive, risk scores present questions of what arbitrary expression was used to create it, how often it's updated, and how to treat a risk score of 59 vs. 65 without inadvertently impacting legitimate users and creating a spike in false positives. A system with high accuracy can bypass that complexity and simply let you block attackers, and let legitimate traffic through without false positives.

Usability

Security shouldn't come at the cost of complexity, especially in the application security realm, where tools may be managed by security teams and/or DevOps teams. Cumbersome interfaces and workflows hinder user adoption and overall effectiveness. An ideal WAF provides a user-friendly experience in visibility, rule building and maintenance, and integrations.

Visibility, insights, and faster decisioning

Visibility has historically been a problem for WAF users. Many WAFs cannot provide real-time, granular, and useful visibility into their decisioning activities. Also, many organizations deploy multiple WAFs across their footprint because they haven't found one that covers all their environments. This results in poor and disjointed visibility as their WAF data is partitioned across different tools and dashboards, and teams lose time when switching between different consoles. Getting all your WAF data in a "single pane of glass" without having to tie different platforms together or constantly do data exports and merges is invaluable.

The WAF should provide an "at-a-glance" utility with intuitive, customizable dashboards and reports that offer real-time insights into ongoing attacks and potential security incidents across all deployments. "Black boxes" and a lack of visibility are common complaints with WAF solutions - users struggle to analyze and mitigate threats because data is not presented in a straightforward way. What criteria were used to make a given decision? How are different types of traffic being handled? Where are the outliers, and is anything behaving unexpectedly?

These kinds of questions should be easy to answer at a glance, and in one location across all of your organization's WAF deployments. These insights can drive measurable improvements in security, revenue, and performance, and they should be at your fingertips. This kind of visibility also helps show leadership that the investment is paying for itself (and then some), and that the budget is doing more to help the bottom line by effectively blocking malicious traffic, not just checking off a compliance requirement. Bottom line on visibility: It should be easy for your WAF to show you how it's making your team more efficient and more effective, and clearly report on the ROI it's delivering.

Simpler rule building, tuning, and maintenance

Traditional WAFs rely heavily on manual regex rule creation. While effective for known threats, signature-based detection struggles with constantly evolving attack methods and zero-day exploits. It is also error-prone. Overly broad signatures can mistakenly block legitimate traffic, causing disruptions and frustration. Additionally, security teams are constantly under pressure to update and maintain libraries of thousands of regex rules, a time-consuming and error-prone process. The burden of maintenance and the risk of false positives is so high for some organizations that they disable the WAF altogether or never turn on blocking mode, leaving their applications vulnerable to attack.

Modern WAF solutions move beyond signature-based detection, aiming to more effectively block both known and unknown threats with near-zero tuning. This minimizes disruptions to legitimate traffic, reduces toil, and lessens reliance on manual error-prone processes. They often employ advanced techniques like machine learning (ML) or context-aware detection. ML can analyze vast amounts of traffic data to identify anomalies indicative of malicious activity, even if it's a novel attack. Contextual analysis can allow solutions to go beyond just looking at the "what" of an attack (specific patterns) and delve into the "how" (attacker behavior). This allows them to detect suspicious behavior patterns even if the attack is new. Context-aware solutions consider the context of a request, such as the request parameters, user location, request frequency, and historical behavior. This context helps distinguish legitimate requests from malicious attempts.

As every environment is unique, there will be a need for custom rules and logic. That's why it's also important to look for pre-configured security policies and an intuitive interface for creating and editing your custom rules.

Pre-built integrations for DevOps and security toolchains

Your WAF should fit into how your organization works, rather than forcing your teams to adapt to something new and difficult. Thus, there should also be a robust set of pre-built integrations with DevOps and security toolchains. You want to instantly enable your team to take advantage of the new and better data from your WAF within the tools and CI/CD workflows that they already use to keep daily operations efficient and make scaling easier without security bottlenecks. Examples include getting real-time alerts in Slack so your team can respond quickly, sending logs to your SIEM solution for further analysis and correlation, and automating rule updates using Infrastructure as Code (IaC) to minimize manual work.

Ultimately, a user-friendly WAF translates to faster implementation, reduced operational costs, improved team efficiency, and enhanced security visibility for your organization.

Deployment flexibility and speed

When evaluating WAF vendors, consider not just the ease and speed of deployment but also the coverage of their deployment options. As mentioned earlier, visibility suffers when all your applications and APIs aren't covered by the same WAF. It's important to consider other deployment options even if you only need one type right now. Think about a future where you might adopt, acquire, or build new infrastructure differently. Choosing a WAF with various deployment options future-proofs the simplicity and cost-effectiveness of your security posture. Stay flexible and maintain your ability to adopt new DevOps trends or make decisions based on what's best for your organization, rather than being limited by a vendor's constraints. Make sure you choose a WAF that covers all the environments listed here, and for more on the benefits of selecting a single WAF that provides comprehensive coverage, see the section on vendor consolidation below.

Protect your applications across your entire footprint

1. **Cloud and container-native:** Application tools and frameworks like Kubernetes, Envoy Proxy, and Istio, are all moving further into a DevOps-focused worldview. Not all WAF solutions are flexible enough to deploy in these types of environments, which can be a limitation for your current architecture, but could also create pressure to avoid architecture decisions you would like to employ in the future.
2. **Edge:** The addition of a CDN can provide exceptional Layer 3 and 4 DDoS protection and resiliency. The deployment can also offer greater flexibility around deception techniques and cutting-edge serverless computing use cases.
3. **On-premises:** This deployment is typically best for protecting legacy applications or those deployed in data centers, and you need a WAF that can install and inspect traffic before web requests reach the app or API endpoint. For example, at the load balancer or the API gateway. For customers with requirements that don't allow for installation at the load balancer or API gateway, look for a WAF that can be deployed as a reverse proxy.
4. **Hybrid:** Hybrid deployments provide the greatest level of flexibility as your infrastructure evolves and shifts between cloud, on-premises, containers, and more. It's ideal for phased cloud migrations or securing applications in diverse locations. Hybrid deployments can offer the benefits of cloud-based WAF solutions but may be countered by the drawbacks of on-premises. Choose a security vendor that secures across different deployments seamlessly, with centralized management for usability and a unified security posture across the entire footprint. You don't want to manage multiple security vendors across your different deployment types, and you don't want to use a vendor where the offerings are so disconnected that it still feels like using multiple vendors.

Faster time to value through rapid deployment

Everyone wants fast time-to-value – if you pay for a 12-month contract and it takes you a month to deploy in production, that's a huge amount to overpay compared to the benefits provided. In addition to deployment flexibility, look for options that can deploy on a scale of minutes or days, not weeks or months. Look for customer references who have had quick deployment successes, even in emergency situations, like while they were under attack.

Automated deployment through Infrastructure as Code (IaC)

While not a deployment type itself, automating deployments with infrastructure as code (IaC), like Terraform, helps reduce lead time for provisioning and security changes while allowing for more trust in the application developers by empowering them with deployment automation.

Scalability and performance

Application traffic can fluctuate significantly in today's dynamic online environment. Whether it's a viral marketing campaign or the latest malicious attack, the sudden spike in traffic can overwhelm traditional WAFs, leading to slowdowns, outages, and lost revenue. That's why beyond feature and capability comparisons, it's also important to look at the underlying network, architecture, and platform a WAF is built on. You want to be sure that you're investing in a solution that scales alongside your business and traffic growth and evolves to meet the needs of new and upcoming threats.

Look for a solution that automatically scales resources to meet demand, offers a globally distributed network with low latency, and delivers high-throughput processing without sacrificing security effectiveness. This ensures a smooth user experience, robust security, and improved business continuity for your critical applications.

A globally distributed edge network will handle protection closer to your end users no matter where they are, making their experiences as fast as possible. An architecture like this will also make rule propagation much faster throughout the network, ensuring newly added security measures take effect across the entire network immediately.

Edge computing is gaining traction for security use cases where custom code or edge storage can enable creative security solutions. For example, edge computing customers can go beyond the feature set of the WAF to serve up customized honeypots to malicious requests. Customers concerned about price scraping can serve fake pricing data to scraping bots using a combination of WAF and edge computing capabilities. Both edge computing and WAF can run on the same network, which accelerates performance and further reduces load on your origin servers.

Accuracy

If a WAF fits all your other requirements but isn't accurate in identifying and blocking malicious actors, it will likely end up as shelfware, stuck forever in monitoring or logging mode only. All WAFs will claim great accuracy, but there are a few ways to dig deeper into the truth. What percentage of the WAF's customer base has it deployed in blocking mode? If a high rate of customers has the WAF in logging mode most of the time, it indicates low confidence in its accuracy and high concern that turning it to blocking mode would harm legitimate traffic. You should also check that success doesn't require continuous configuration and tuning because you should not have to constantly babysit a WAF for it to be effective.

Thresholding is a useful way to lower false positives and provides more confidence and flexibility when setting up a new WAF or creating new rules. With thresholding, attacks are not blocked immediately but only once they hit a certain threshold in the number of malicious requests. This works because, in the real world, you'll often see hundreds or thousands of malicious requests coming through rapidly, not just one or two. Once you've seen how the rule responds to production traffic, you can lower the threshold or move to instant blocking.

Additional considerations

There are a few other things that round out your experience with a security vendor to ensure your engagement with them delivers as much value as possible, with as few surprises and costs as possible. This can include understanding the terms for professional services engagements so you're not hit with surprise costs every quarter for services you thought were included, or just making sure they have extended services or flexible pricing available in case your needs change in the future. Here are a few more things to consider to make the best long-term selection for your organization.

Single platform for vendor consolidation

Vendor consolidation is a big topic as everyone wants to cut costs, reduce duplicate efforts, and simplify both their technology stack and procurement efforts. Inferior WAFs can't protect your entire technology footprint, and they can't protect against the full breadth of threats that your applications, APIs, and microservices are facing. For example, DDoS protection and bot management may be handled by separate products and not included as part of a WAF's coverage. When split between vendors there's *usually* inefficiency in pricing, but there's *always* inefficiency for your visibility into what's happening, your management of your security posture, and your time to resolution for new attacks and issues that appear. Keeping all your security data and tool management under a single pane of glass delivers faster and more useful insights. You will have better control when your security policies are all in agreement, and if something goes wrong, it's much easier to identify where the problem is.

Strong vendor support and services

Customer success, support, and satisfaction are a good way to investigate whether all the other claims a WAF vendor makes can be trusted. Third-party review sites, analyst reviews, and peer feedback will give you a pulse on overall customer satisfaction and how their customer support is viewed. Do their customers think their support is fast, helpful, and comprehensive? Do they feel like they have a working partnership? If you see tons of complaints about the service or lack thereof, that's a red flag you should pay attention to. When in doubt, ask if they maintain a customer satisfaction score (CSAT), and it will likely offer the insight you seek.

You should have a clear understanding from the start about what will fall under professional services engagements and what will not. While you may not intend to use the vendor's services, it's also worth noting what service level agreements (SLAs) they commit themselves to. Should you need to leverage services, stronger SLAs guarantee the quality if the time comes.

There should be a wealth of high-quality and up-to-date resources for customer onboarding, migration, and product documentation. You should also look for an active customer community where you can ask questions to other customers along with experts from the vendor, and gain insights from what your peers are doing.

Team-augmenting managed services

No matter how easy a WAF is to use, you may find it makes more sense for your organization to contract expert AppSec help rather than try to staff up a full-time team. Even if you don't think you want to engage in managed services at the start, you should see if it's offered in case it becomes a need in the future. Scaling a security team is difficult – SecOps and AppSec experts are expensive and in high demand, and maintaining the option to scale the impact of your existing team with cost-effective help from the vendor side could be a useful option.

The WAF you select should have an option for managed and professional services staffed by true AppSec experts who can help you solve problems and achieve the outcomes you want. The pricing should be straightforward without any loopholes for surprise billing, and other customers should hold it in high regard. Make sure that response times are fast and customer satisfaction scores are high. Check for different tiers of support or types of engagement designed to work for different needs or sizes of organizations. In short, make sure there's an option that would work for you, even if you don't need it right away.

Competitive pricing and licensing options

Pricing should be clear, with no hidden charges, and no overages or surprises after the fact. Know what's included in the base package and what features (bot management, API security, DDoS, logging, etc) are add-ons. Also, know at what level you'll be moved up into the next tier. We've heard from plenty of organizations that got in on a low-cost tier with a vendor but were then surprised at renewal time with the news of having to move up to a much more expensive pricing tier.

A vendor should be dedicated to helping you find the right plan for the service you need. Professional services can be useful for some projects you want to take on with help from your vendor, but it shouldn't be something that is required for simple things that a better WAF would let you update on your own, or handle through a better customer support offering.

If you already spend a lot of your budget with AWS or GCP you should look for a WAF that you can purchase through their marketplaces. This can simplify your billing and allow you to use budget or credits that are already locked into those systems rather than freeing dollars from elsewhere in your organization.

Conclusion

Selecting a WAF that is simpler to use, simpler to understand, AND more effective is the only way to go. But WAF buyers often aren't aware that this is an option because complexity and poor performance has been normalized in the market. You can't afford to settle for the old way of doing things when security best practices continue to evolve at such a rapid pace. With the right vendor, you can reduce risk to your applications and toil for your teams while increasing their effectiveness, and seeing ROI quickly.

Get in touch to learn more →

Additional resources

Learn more about a few of the topics mentioned in the buyer guide:

1. Navigating the OWASP Top 10
2. Escaping the Black Box of Security Visibility with Signals
3. Why regex isn't the best for security rules
4. Fastly's expert and managed security solutions
5. Fastly's flexible WAF deployment options
6. Fastly's integration partners
7. Faster time-to-value with Fastly's Next-Gen WAF
8. Why the Fastly WAF's blocking mode is so powerful
9. Preemptive blocking with Fastly's Network Learning Exchange