fastly

Fastly WAF

CUSTOMER FAQ

This document provides answers to some of the most frequently asked questions about Fastly's Web Application Firewall (WAF) service which is currently in Limited Availability (LA).

General Questions

How does Fastly WAF work?

Fastly WAF is integrated into our global nodes and runs as part of the cache logic within our CDN service. In each cache node, a policy is published consisting of rules and mitigations for common web application attacks and exploits. Fastly WAF rules can be configured in real-time via our API. Our WAF can run either in active blocking mode or passive logging mode only.

How does your WAF interact with my current Fastly service?

Fastly WAF is built into our Varnish platform. Our engine consumes trusted third party rulesets from OWASP and TrustWave, in addition to open source and Fastly-generated rules. These rules are converted into VCL which can then be deployed into your Fastly service. The Fastly WAF ruleset operates as an edge dictionary attached to your Fastly service. Any requests that miss or pass cache are automatically inspected by our WAF ruleset.

Does your WAF service comply with PCI DSS 6.6 requirements?

Compliance with the PCI Standard is based on fully audited assessments of standards compliance done by 3rd parties. Fastly WAF is designed to adhere to the PCI DSS 6.6 requirement. However you will need to validate appropriate use of our WAF while working with your auditors for PCI DSS 3.0 certification.

What rules do you use?

Fastly utilizes the OWASP Top 10 Threat Protection ruleset which covers common exploits to web services. We have also partnered with TrustWave SpiderLabs to leverage their ModSec rule support. These rules protect common content management systems as well as web server and application back-ends, like SQL databases. Additionally, Fastly's Security Research team has added rules which detect commonly known attacks: IPs, bad user agents, and known botnet command and control nodes.

How do you update your rulesets?

As part of our WAF service, we manage rule updates from commercial third parties, Fastly research, and open sources. We update published rules into a policy as needed, or when changes to the rules are available from their respective sources. New rules which match the published classes of rules are also inserted into the WAF instance of any service once it is enabled. This helps ensure immediate coverage for new or evolving exploits.

What types of security threats do you protect against?

Fastly WAF secures web applications from key application-layer attacks, such as injection attacks and malicious inputs, cross site scripting, data exfiltration, HTTP protocol violations and other OWASP Top 10 threats.

How do you protect against these kinds of attacks?

- **SQL injection attacks:** Both the OWASP ModSecurity Core Rule Set and the TrustWave commercial ruleset contain specific filters for SQL injection attacks and its variants.
- **Cross-site injection attacks:** The OWASP ruleset protects against cross-site injection attacks. Fastly leverages a scoring mechanism for each request looking for cross-site injection and other threats to the origin. We score every request against the entire core ruleset and validate that the request score is below a configurable threshold in order for it to pass.
- **Brute force attacks:** Brute force attacks are covered in the OWASP ruleset which we use. Fastly also blocks brute force activity by leveraging VCL to recognize specific sources, requests, or attempts to brute force or overwhelm security controls prior to any traffic reaching the origin datacenter.
- **Sustained attacks:** These can be mitigated through Fastly WAF or by creating specific rules and filters within VCL. Our DDoS mitigation capabilities also help ensure the availability and integrity of origin service content. Fastly has protected its clients from sustained attacks against high profile targets over long periods of time, including active attack vector evolution and multiple attack vectors simultaneously.
- **Burst attacks:** Burst attacks designed to overwhelm available resources are dealt with seamlessly within the Fastly infrastructure. Fastly builds in necessary bandwidth and transactional capacity to ensure that burst attacks do not affect customers through the exhaustion of resources needed to deliver legitimate traffic to end users.
- **Network attacks:** Network attacks, or attacks targeting network infrastructure are managed automatically by Fastly. We do not pass DNS to origin, and traffic that does not match a narrow HTTP, HTTPS or DNS profile is discarded at the edge of the network. Attacks targeting control protocols are defended against through authentication of endpoints throughout the network. Additionally network protocols used within the Fastly network are hardened to ensure that they cannot be leveraged as a means of amplification or reflection. Customers are responsible for protecting against attacks that bypass the Fastly network by leveraging the Fastly Cache IP address space, published to our customers as a component of our CDN service. It's recommended that origin IP address space not be published in public DNS to ensure bypass attacks cannot use these addresses as targets.
- **JavaScript injection:** JavaScript injection attacks are protected through the WAF and enablement of existing rules to protect against malicious code being inserted into client communications with web services. Common exploit patterns or scores are filtered through the WAF to ensure the integrity of the origin service.

How do your DDoS and WAF solutions differ in terms of Layer 7 protection?

Fastly WAF operates at Layer 7, providing protection against known *integrity threats*. It can be used to mitigate application-layer attacks which work by abusing web applications. Fastly's DDoS solution also operates at Layer 7, protecting against *availability threats*. It addresses both application layer attacks and network threats such as SYN floods and bandwidth exhaustion attacks.

Does your WAF solution slow down performance?

No, Fastly WAF is actually integrated into our global Varnish environment. This means we can save you valuable milliseconds that are lost with other WAF solutions which have to send traffic through an external WAF platform for scrubbing. Also, since our service places heavy focus on caching as a strategy, less traffic needs to be sent to origin, thereby improving performance. On average, we are currently seeing approx. 20 ms of induced latency per process request.

How do I deploy your WAF into my Fastly service?

Enabling Fastly WAF doesn't require any modifications to your web application or origin servers. Our Customer Support team will work closely with you to ensure a smooth deployment. This involves setting up a logging endpoint, selecting rulesets and optionally customizing the response. Once our WAF is set up, you can begin monitoring logs to determine which requests to your origin are legitimate and which you should consider blocking to protect your origin. For more details, check out our **documentation**. If you require fully managed threat analysis and WAF management please contact **sales@fastly.com**.

Additional Features & Functionality

Do you offer automated blocking capabilities?

We automatically block traffic matching any published rule for a service as long as those rules are set to blocking and logging modes within our WAF. Automatic blocking of network level threats (such as DDoS attacks to IP addresses, services or infrastructure) happens seamlessly within our network for all non HTTP (port 80) or HTTPS (port 443) traffic. Fastly's Varnish Configuration Language (VCL) can also be altered so that known malicious clients, source IP addresses, request types, HTTP methods, or specific requests can be blocked at the cache prior to transiting to the origin data center.

How do you protect against direct-IP attacks?

Fastly **cache IPs** are available at no extra charge, for any customer to use to validate requests that have been processed through the Fastly edge. Fastly's Origin Cloaking feature prevents these kinds of attacks by hiding your origin from attackers. Using private network interconnections, we connect directly with your origin server, hiding the IP address from the public internet. This forces all attack traffic to go through our network, where we apply DDoS mitigation rules.

Is your WAF aware of protected applications (e.g. WordPress platform)?

Yes. We can configure policies for the application stack of your website and apply pre-defined templates using Trustwave SpiderLabs ruleset classes. We cover well known content management systems like WordPress, and Drupal, common technologies like XML and HTML and popular coding languages like Python, Perl and Ruby. We can also apply exceptions to eliminate the need for fine tuning (e.g. you can do conditional inspections in VCL for our WAF).

Will your WAF support "blacklists" and "whitelists"?

Yes, white / black listing is inherent in our service, and can actually be customized with VCL settings / conditions. In addition, our WAF rulesets are enabled by default with known Access Control Lists (ACLs) to check against. Fastly also provides IP ACLs for known abusive space as well as legitimate space for customers to leverage in filtering, as needed.

Are custom rules supported? What kind of firewall rules can be created?

Because our WAF service was built on our Varnish-based platform, we have complete flexibility when it comes to creating firewall rules. This allows customers to create custom WAF rules, written in VCL, in addition to the provided rulesets. Any aspect of the HTTP/HTTPS request can be inspected and any rule can be created to match any aspect of that request. These rules can be created by our Customer Engineering team, or by your in-house support team

Does your WAF support virtual patching?

Yes, our WAF supports virtual patching through publication of immediate vulnerability patches to the policy globally, in under a second. We can apply virtual patches per platform, specific to different versions of platforms. We can also patch dynamically — e.g. we can identify an exploit by string or strict match with Regex, and create a rule for it which we then import into our WAF. To create customized virtual patches we recommend you contact support@fastly.com.

How do you provide real-time information about an attack?

Fastly provides real-time access to WAF logs, events, and notifications through our logging functionality. We log 100% of all WAF activity within 1-2 seconds of it happening at the edge. This allows you to manage events, notifications, and activity immediately with first hand visibility. All WAF logs will follow Mod_Security formats for compatibility with event management systems out of the box.

Can you inspect TLS / SSL traffic?

Fastly WAF is part of our CDN, so we can terminate TLS traffic on the edge, prior to WAF inspection. This ensures consistent inspection of all traffic, both TLS and non-TLS terminated sessions. If a connection is encrypted, we decrypt it, run it through all Varnish logic (including WAF), reencrypt it, and send it to origin. All of this happens in memory, so we do not write anything to disk or retain any of the content.

Does your WAF provide advanced bot protection?

Known malicious bots and TOR nodes are detected through WAF rulebase rules that inspect for HTTP header and IP address data. The rules identify and log all activity seen from these endpoints and logs can be streamed to your log collection system. Fastly has also partnered with PerimeterX for advanced client-side bot detection. We can recognize and enforce critical business logic to protect against affiliate fraud, loyalty and gift cards as well as attempts to brute force elements of the site associated with account hijacking activity. For more information on the Fastly and PerimeterX offering contact PerimeterX@fastly.com.

Does you WAF protect against malware?

Malware protection is included within our WAF for known malware types and their variants, to the extent that the available ruleset provides. Fastly also runs active validation of content within the cache in a generalized way to ensure that traffic is not coming from or going to known malware distribution locations. We run frequent testing against our platform to ensure no malware is resident or being inserted into the infrastructure for infection or distribution.

Does your WAF offer DLP functionality?

Fastly WAF does not currently offer Data Loss Prevention (DLP) or any out-bound content inspection. However, outbound content inspection is on our WAF roadmap. We will look to be able to inspect web application output and respond (block, allow, mask and / or alert) based on policies or rules.

Does your WAF allow for creation of rules based on usage?

Fastly provides request by request processing — we do not keep state across a client over time. As a result, we do not support the creation of rules based on usage.

Does Fastly have a security sandbox?

Sandboxing is not supported today for file types outside of the normal WAF functionality. If the integrity of file types is a concern, you can leverage Fastly's conditional VCL capabilities. We support multiple backends to integrate a service like file type inspection, ensuring that content is not weaponized, negatively impacting the origin.

Does Fastly have a Security Operations Center?

Fastly does not maintain a physical Security Operations Center (SOC). Instead we maintain a security operations process that allows us to engage the right resources to respond to incidents in a real-time manner. For major incidents, Fastly maintains an Incident command (IC) structure which engages during events with customers that have experienced noticeable impact. We also offer 24/7/365 follow-the-sun support.

Do you offer behavioral analysis to uncover attacks that could otherwise go undetected?

Fastly provides a stream of real-time logs which collection systems can use to detect repeated attempts against a service or endpoints. We can stream log data on every request. All records are streamed in Modsec format. Logs can be streamed to common log collection systems such as Sumo Logic, Splunk, Logentries, Papertrail and Datadog.

Support

What service support is available for WAF Limited Availability?

The onboarding process for our WAF service will be manual for Limited Availability (LA). Our Customer Support team will build a default policy in logging only mode, log this information into your facility, validate it and run positive tests. We will then hand it off to you, at which point you can choose to leave the service in logging mode, or switch to blocking mode. While we can help with that switch, you will need to manage whichever mode you chose. You will need to open a ticket request for any rule changes during the LA period. If you are using our fully managed WAF we will work with you on the above steps. Please see the below table for more details on our WAF support offering:

Support offering	Details
Note: All Fastly WAF customers must be on our Gold or Platinum Support Plan	
Online self-service help	Unlimited access
Availability for general inquiries	24/7
Availability for incident reports	24/7
Initial response times	Attack notification response within 15 minutes
Web and email support	Available
Phone and chat support	Toll-free telephone available 24/7/365. **Platinum Support:** Dedicated chat channel available during Fastly Business Hours. **Gold Support:** Temporary Slack channel for setup / testing period (approx. 30 days). Channel archived unless any issues arise.
Emergency escalation	Available via email and phone support

What is included in your WAF onboarding service?

The following will be included as part of our onboarding service for Limited Availability:

- Enable the designated service(s) for WAF functionality and provide access to our rule and filter libraries
- Work directly with you to determine the right set of rules and filters for your service
- Support deploying default ruleset — OWASP, Trustwave SpiderLabs, basic threat filters
- Publish those rules or filters into your service in logging mode
- Monitor the behavior of those rules for a one week period, starting when the rules are published to the service

What support do you provide for rule setup and integration?

We provide rigorous validation of rule setup and integration. This includes the following steps:

- Validate via positive testing that the ruleset is deployed in the service with the classes of rules specified (will require live testing, either ongoing or periodically)
- Support for triage, analysis and improvements to rules creating false positives
- Integration support of WAF functionality with other VCL and Fastly platform functionality

What ongoing support do you provide?

For Fastly WAF customers on either Gold or Platinum Support Plans, the following ongoing support will be provided:

- Break / fix of WAF functionality
- False positives with rules
- Updating off-the-shelf rules
- Support for integration of other Fastly functionality with WAF (conditional responses)
- Product evolution

How do you handle false positive triage?

False positive triage will resolve instances where legitimate requests have triggered a WAF rule or filter. We will either remove the rule from the policy or, where possible, modify the rule or policy to address the legitimate request properly.

Does Fastly offer a fully managed threat analysis service?

If your organization requires real-time threat analysis please contact sales@fastly.com to discuss requirements.

More questions?

If you have any questions that are not answered in this FAQ please contact support@fastly.com.