fastly

Powered by Signal Sciences

PRODUCT DATASHEET

Fastly Next-Gen WAF

Architecture and Deployment Overview

Unified web app and API security for any environment

Fastly offers the most flexibly deployed WAF on the market and can protect your apps and APIs wherever they are—in containers, on-premises, in the cloud, or at the edge—with one integrated solution. Gain comprehensive protection without sacrificing performance or requiring dedicated headcount: the Fastly Next-Gen WAF (powered by Signal Sciences) simply works out of the box and is so effective 90% of our customers run us in full blocking mode.

The Fastly Next-Gen WAF provides the proactive protection modern apps require while integrating into your DevOps and security toolchains for unparalleled visibility. Our flexible architecture can advance your application security strategy by providing developers, operations, and security teams insight into where and how your web applications and APIs are attacked.

This datasheet provides detail into the highly performant, patented architecture of the Fastly Next-Gen WAF, as well as information on the wide array of deployment options available. This document is arranged into the following sections:

- **Architecture overview**
- **Deployment options**
- **DevOps and security toolchain integrations**

Gartner Peer Insights Customers' Choice 2022

Signal Sciences, now part of Fastly, is the only vendor to be named a Gartner Peer Insights Customers' Choice for Web Application and API Protection (WAAP) for four consecutive years and is one of the **highest-rated** WAAP solutions on the market with an overall rating of **4.9/5** as of 31 January 2022[1] based on 267 reviews.

1: Gartner Peer Insights content consists of the opinions of individual end users based on their own experiences with the vendors listed on the platform, should not be construed as statements of fact, nor do they represent the views of Gartner or its affiliates. Gartner does not endorse any vendor, product or service depicted in this content nor makes any warranties, expressed or implied, with respect to this content, about its accuracy or completeness, including any warranties of merchantability or fitness for a particular purpose. GARTNER PEER INSIGHTS Logo is a trademark and service mark of Gartner, Inc. and/or its affiliates and is used herein with permission. All rights reserved.

© 2022 Fastly, Inc. All Rights Reserved

Architecture overview

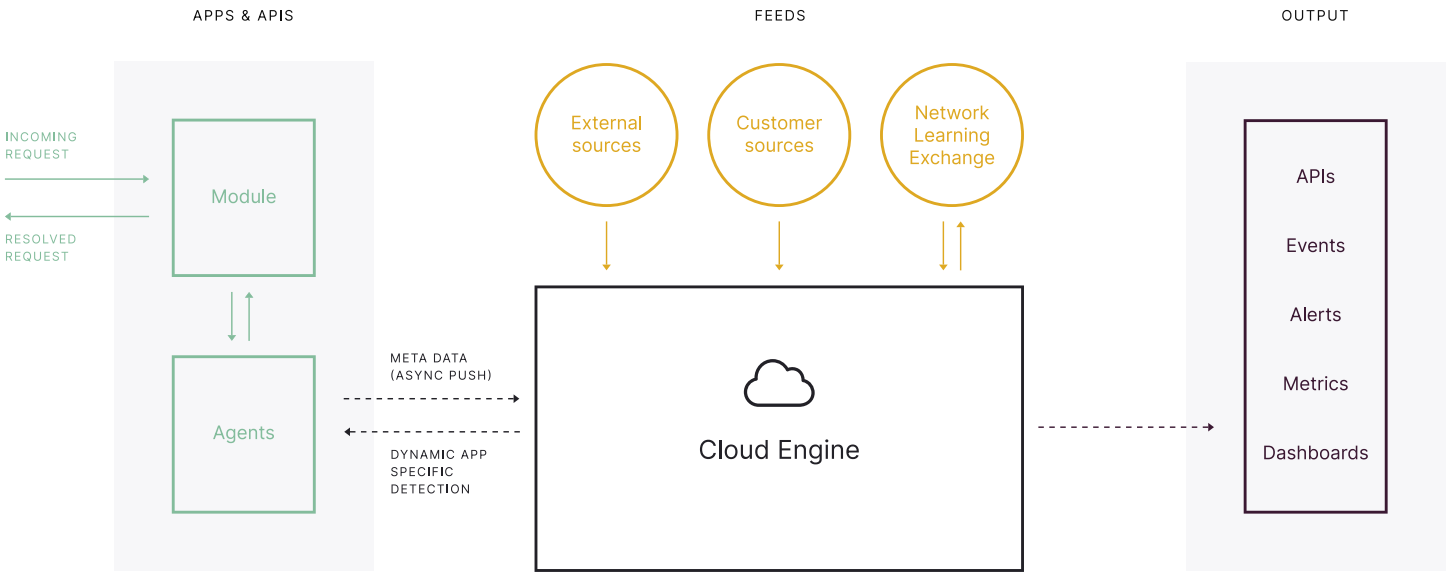

The Fastly Next-Gen WAF is a hybrid software as a service (SaaS) solution with three main components. This patented approach, developed by Signal Sciences, allows us to easily scale and protect even the highest volume applications and APIs without impacting performance.

Agents

Lightweight agents you deploy on your existing infrastructure to perform detection and decisioning against requests quickly and accurately.

Modules

Optional but powerful component that pairs with our agents to enforce high performance and reliability.

Cloud Engine

Cloud-hosted analytics backend that enriches the agent asynchronously with intelligence gathered from external and proprietary sources to make dynamic, application-specific detections.

© 2022 Fastly, Inc. All Rights Reserved

Agents

Agents consist of a small daemon process and are designed to handle extremely heavy loads while making highly-performant and accurate detections and decisions locally. The agent also collects metadata about the malicious requests it has processed and shares that metadata with the Cloud Engine. We protect some of the highest volume sites on the Internet, where tens of thousands of agents collectively process trillions of production requests without impacting app or API performance. Agents block attacks before they hit applications or APIs and provide visibility into not only requests that come in but also server responses and anomalies that show how the application is behaving.

Modules

Modules run on virtually any web server (NGINX, Apache, IIS, and more) or application language (.NET, Java, Python, PHP, .nodeJS, and more). The module is just a few hundred lines of code to ensure both reliability and extreme performance. Its sole job is to pass requests through to the agent and receive and enforce decisions from the agent to allow the request through to the application or log/block it (depending on the mode set in the console).

Cloud Engine

The Cloud Engine collects and analyzes anonymized attack data and telemetry from the many thousands of software agents across our customer base. The output from the Cloud Engine is used by the agent locally to perform better detection and make more aggressive blocking decisions. The agent decisioning is enhanced by our Network Learning Exchange (NLX) which shares confirmed malicious IP sources within the management console, alerting you to suspicious actors before they are a threat to your applications and APIs. Other feeds include external lists of malicious IPs and customers' custom IP lists, all of which provide additional request context that enriches the agent decisioning. This visibility and context is shared via our API and native integrations with the DevOps tools your team already uses, including Slack, PagerDuty, Jira and more as well as security tools like Splunk, Elastic, and Palo Alto Networks Cortex XSOAR. Metrics and event reporting for your entire application footprint are also readily available via dashboards in a unified management console.

© 2022 Fastly, Inc. All Rights Reserved

Deployment options

Native deployment options for data center, cloud, containers, and serverless

Deployment option 1: Cloud and container-native

Agent-module pair installs at your web server, API gateway, or at the app level within minutes. Our agent is infrastructure-agnostic which provides you the flexibility to deploy where you need it, without worrying about dependencies on underlying languages or frameworks.

Deploying in Kubernetes and service mesh

New application tools and frameworks, such as Kubernetes, are quickly moving companies into a DevOps-focused world. Companies now release code faster than ever before and Fastly offers flexible deployment options to fit within your container strategy with three “layers” where you can install our WAF in Kubernetes and four methods for how you deploy. Additionally, our native integrations with Envoy Proxy and Istio service meshes mean Fastly provides visibility into both north-south (client-server) and east-west (service to service) requests.

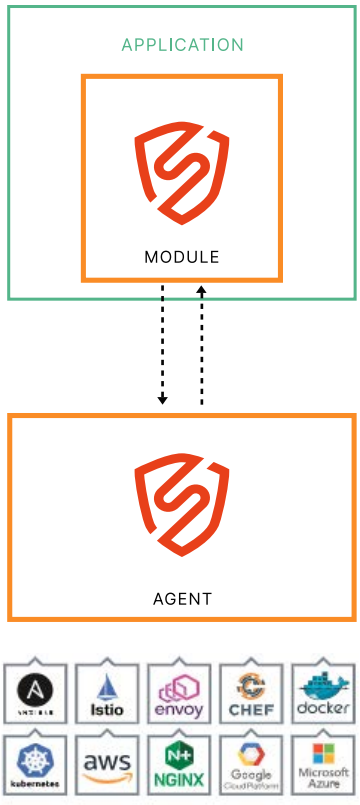

Install Method	Layer 1: Ingress Controller	Layer 2: Mid-Tier Service	Layer 3: App Tier
Agent + module in same app container	✓	✓	✓
Agent + module in different containers	✓	✓	✓
Agent in reverse proxy mode in same container as app	✓	✓	✓
Agent in reverse proxy in sidecar container	✓	✓	✓

Fastly fully supports deployments for: kubernetes docker envoy Istio

© 2022 Fastly, Inc. All Rights Reserved

Deployment Option #2: Data center and legacy applications

Customers who need protection for legacy applications or those deployed in data centers typically choose one of two deployment options: install the Fastly Next-Gen WAF to inspect traffic prior to web requests reaching the app or API endpoint, or install our agent in reverse proxy mode. For example, our module can be installed at the load balancer (HAProxy, NGINX) or at the API gateway (Ambassador, Kong, Cloudentity). For customers with requirements that don't allow for installation at the load balancer or API gateway, our agent can be deployed in reverse proxy mode. Either deployment option provides the same level of visibility and actionable insights and alerts as our other deployment options with full feature parity.

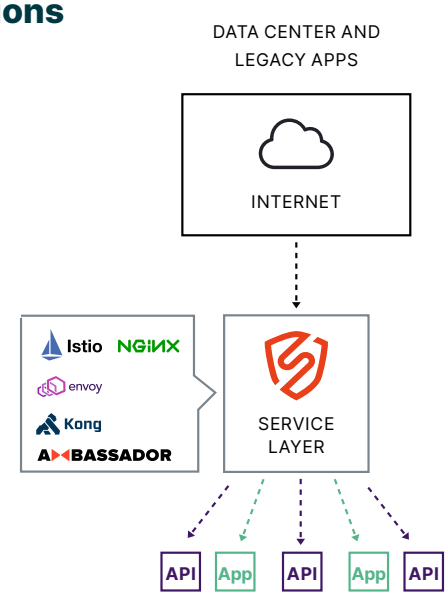

Deployment Option #3: At the edge

The Fastly Next-Gen WAF is available on the Fastly Edge Cloud Network, allowing customers to enforce security controls as part of Fastly delivery services. The edge cloud deployment option is seamlessly integrated with Fastly's caching layer, Varnish.

This provides protection and acceleration closer to users and shields origin systems from abusive attack traffic while delivering world-class performance. Our edge deployment is ideal for customers who are unable to install software on existing infrastructure and for those who want to take advantage of the performance benefits of Fastly's global content delivery network (CDN). This deployment option also offers additional features including Layer 3 and 4 always-on DDoS protection and TLS management.

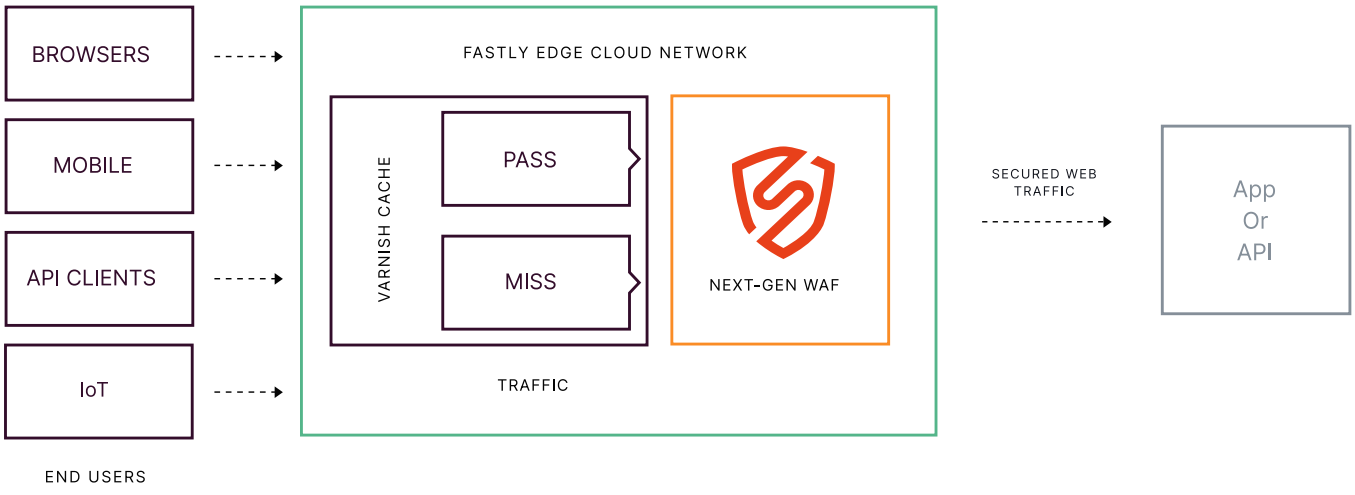

© 2022 Fastly, Inc. All Rights Reserved

Deployment Option #4: Cloud WAF

Cloud WAF empowers you to quickly and easily protect web applications, APIs, microservices and serverless applications—without installing software on your infrastructure. Once deployed, a simple DNS change to point application traffic to Cloud WAF is all that's needed to enable the visibility and protection of the Fastly Next-Gen WAF for your applications. All web requests are redirected to our cloud enforcement layer where bad requests are detected and blocked. All good, legitimate traffic is then forwarded to your application origin server. Cloud WAF is ideal for customers wanting to add an easy-to-manage WAF without making upstream changes to their CDN layer.

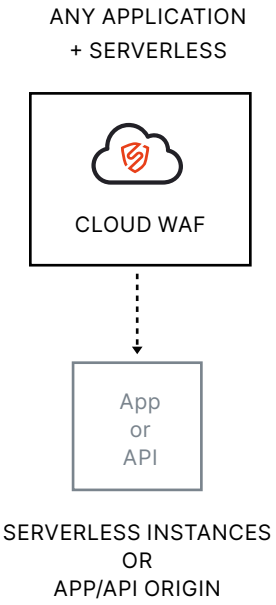

Protection that's committed to data privacy

Many leading financial services firms, healthcare companies, and others with strict data privacy requirements all utilize Fastly's next-gen WAF because of our strong architecture built for data privacy. All sensitive data is handled entirely within the customer environment and only sanitized and redacted portions of requests that are marked as attacks or anomalies are then sent to the Fastly Cloud Engine.

Once the agent identifies a potential attack or anomaly in a request, a set of fully customizable redactions are applied locally and then the agent sends only the redacted individual parameter of the request which contains the attack payload, as well as a few other non-sensitive or benign portions of the request, such as client IP, user agent, URI, etc. Our backend only collects the response's metadata e.g. response codes, sizes, and times. We provide customers the ability to fully customize redaction policies and fields as needed. For additional protection, Fastly automatically enforces redaction of common sensitive data types—such as passwords, keys, GUIDs, and any type of PII or PHI—before the request is sent to our backend.

Betterment

"It works straight out of the box, scales automatically, and does a great job at providing visibility while securing the application."

Anson Gomes
Lead Security Engineer, Betterment

© 2022 Fastly, Inc. All Rights Reserved

DevOps and security toolchain integrations

The best path to success for effective application and API protection is to provide the same baseline of security data to development, operations, and security teams in the tools they're already using. Fastly works with the industry's best tools and platforms to provide real-time alerting into your DevOps and security toolchains and to ensure it's easy for your teams to leverage our production security telemetry within your organization's current tools and processes for further investigation and analysis.

Out-of-the-box technology integrations help teams make or continue their transition to modern development models and architectures. Our single-click integrations include the most common development and operations alerting engines, chat-ops, project management, and incident tracking systems.

Technology and platform integrations

Platform integrations & partners

Run the Fastly Next-Gen WAF anywhere

WEB SERVERS

- NGINX
- Microsoft IIS
- Apache
- Apache Tomcat

IAAS

- Google Cloud
- Microsoft Azure
- aws

PAAS

- Microsoft Azure
- HEROKU
- Red Hat
- IBM Cloud
- VMware Tanzu

CONTAINERS

- kubernetes
- docker
- OPENSHIFT
- Istio
- envoy
- MESOS

CONFIG MANAGEMENT

- CHEF
- Red Hat Ansible
- puppet
- SALTSTACK
- Terraform

Feed integrations & partners

Send and receive data from the Fastly Next-Gen WAF

DEVOPS TOOLCHAIN

- Jira Software
- DATADOG
- PivotalTracker
- VictorOps
- slack
- PagerDuty
- Opsgenie
- Microsoft Teams

SOC/SIEM

- splunk>
- elastic
- ELK Stack
- Radar
- sumo logic
- McAfee
- NETWITNESS
- ArcSight
- LogRhythm
- CISCO
- CORTEX XSOAR

Getting started

Unlock highly effective security without impacting performance.

To learn more about our security solutions, visit our website or contact us at sales@fastly.com.

© 2022 Fastly, Inc. All Rights Reserved