Fastly Products Archive

Network Services

These articles provide information about Fastly products that focus on performance (speed), availability, and media and that accelerate content delivery with control from an edge cloud platform.

Generated: Thu, 30 Jun 2022 19:46:49 +0000

https://docs.fastly.com/products/product-network-services

<u>Application Programming Interface (API)</u>

Last updated: 2020-05-15

https://docs.fastly.com/products/application-programming-interface ര

The Fastly Application Programming Interface (API) allows you to integrate your applications and services with the Fastly platform.

Our API is presented using a **REST** model. It uses standard HTTP response codes and verbs to allow you to programmatically control all the same features that are available through the Fastly web interface. Our web interface is a client of our public API.

Documentation

The Fastly API provides a variety of endpoints that we document in our API reference documentation. The endpoint documentation for each API call shows the method, path, authentication type, resource, and parameters that must be combined with the base URL to form a request.

Client libraries

To make coding against the Fastly API easier, we maintain a list of <u>client libraries</u> in a variety of languages. Our client list provides more information about each of the client libraries in our list, including third-party tools and integrations in other languages created by the Fastly community.

Base URL and endpoints

With one exception, all API calls referenced in our documentation start with a base URL:

https://api.fastly.com/

The base URL for the <u>real-time analytics API</u> is:

https://rt.fastly.com

Authentication

Nearly all API requests require authentication using an API token. You can limit the capabilities of those tokens using a scope. Our authentication page provides more information about the API tokens you must use to authenticate. Our guide to Using API tokens provides more information about managing these tokens via the web interface depending on the role your organization has assigned to your account.



NOTE

Fastly accounts created before May 15th, 2017 may have used one or more API keys to authenticate API requests. See our note on <u>legacy API keys</u> for more information.

Errors

Fastly uses standardized HTTP response codes to indicate the success or failure of an API request. Codes in the 2xx range indicate success and confirm a request worked as expected. Codes in the 4xx range indicate an error and provide both an error code and a brief explanation.

Rate limiting

Unless otherwise stated in the API reference documentation for individual endpoints, API access is limited to 1,000 non-read requests per hour, per user account. If more than one API token is associated with one user, all that user's tokens share the same limit. Information about rate limit consumption and remaining credit is available in an HTTP response header, examples of which can also be found in our API reference documentation.

TLS version requirement

The Fastly API requires TLS 1.2. Because of the PCI Security Standards Council mandate, TLS versions 1.0 and 1.1 are no longer supported for accessing Fastly's API.

<u>Capacity Reservation</u>

Last updated: 2021-06-03

https://docs.fastly.com/products/capacity-reservation

Capacity Reservation allows you to reserve Fastly traffic capacity for events based on data bandwidth (in gigabits per second), duration, and type of delivery (Media Shield for VOD, Media Shield for Live, Fastly Streaming Delivery, or Fastly Full Site Delivery). You are required to purchase capacity reservations any time you're expecting a utilization spike from your planned events. When purchasing capacity reservations for your events, you are also required to purchase Fastly's Live Event Monitoring service for the duration of the event. To learn more, contact sales@fastly.com.

Billing

Fees for Capacity Reservation are based on the duration, in hours, of the event and the reserved bandwidth, in gigabits per second (Gbps), for Fastly Full Site Delivery and Streaming Delivery traffic (Capacity Reservation - Edge) or Media Shield for VOD and Media Shield for Live traffic (Capacity Reservation - Media Shield). Fees do not include delivery fees for Fastly Full Site Delivery, Fastly Streaming Delivery, Media Shield for VOD, or Media Shield for Live.

Cloud Optimizer

= Last updated: 2020-12-18

https://docs.fastly.com/products/cloud-optimizer

IMPORTANT

This information is part of a limited availability release. For additional details, read our product and feature lifecycle descriptions.

Fastly's Cloud Optimizer product allows customers using one or more non-Fastly content delivery networks (CDNs) to take advantage of Fastly's Full-Site Delivery features without migrating edge delivery traffic to Fastly. Cloud Optimizer works with your existing content delivery infrastructure by designating Fastly as the origin for all of your end-user-serving CDNs. Using Cloud Optimizer provides you with real-time visibility of origin traffic, granular load balancing for your origin infrastructure, and request collapsing to decrease traffic to origin.

To learn more about Fastly's Cloud Optimizer, contact your account manager or email sales@fastly.com for more details.

MOTE

Cloud Optimizer is not available for video streaming activities. Check out Media Shield for Live and Media Shield for VOD instead.

<u>Dedicated IP addresses</u>

Last updated: 2021-05-05

https://docs.fastly.com/products/dedicated-ip-addresses

Fastly's Dedicated Internet Protocol (IP) addresses provide you with a pool of IPv4 and IPv6 addresses, maintained and managed by us, across Fastly's global Edge Cloud. They can be used to support TLS certificate management for non-SNI clients, to support custom cipher suites or IP-to-service pinning, or to help manage <u>zero-rated billing</u> endpoints or security allowlisting.



O NOTE

Purchase of Fastly's Platform TLS product requires you to also have purchased Dedicated IP addresses.

TLS non-SNI client support

Dedicated IP addresses can be used to host customer certificates for non-SNI client support. Fastly can install customer-provided or Fastly-managed certificates at a dedicated set of IP addresses identified via customer-managed DNS records. These DNS records can be set up to use three possible network routing options (sometimes referred to as network maps or domain maps) that allow you to choose which parts of the Fastly network to use. See <u>Fastly's TLS offerings</u> for a more detailed description of the supported TLS options at Fastly.

Custom cipher suites

Fastly supports a number of standard cipher suites. Should you require more personalized control, Fastly supports the creation of custom cipher suites by providing you with dedicated IP addresses that support these custom sets.

IP-to-service pinning

IP-to-service pinning uses dedicated IP addresses to map customer services to specific endpoint IP addresses and direct an end user's request to a specific service based on the requested endpoint IP address.

Zero-rated IP addresses

Zero-rated IP addresses (ZRIPs) allow you to use dedicated IP addresses within Fastly's global Edge Cloud to identify traffic for special treatment. For example, if you need to waive billing charges going to or from specific web pages, ZRIPs can help you to identify traffic for zero billing.

Security allowlisting

Security allowlisting uses dedicated IP addresses to control the set of Fastly global IP addresses seen by third parties. You can incorporate dedicated IPs into access control lists (ACLs) to tighten security between a customer and a third party.

Fastly's Full-Site Delivery

iii Last updated: 2022-06-03

https://docs.fastly.com/products/fastlys-full-site-delivery

Fastly's Full-Site Delivery allows you to speed up websites and mobile applications by pushing content closer to users, providing improved and secure experiences across the world. Full-Site Delivery includes the following features.

Content serving, caching, and control

Full-Site Delivery uses Fastly's global <u>content delivery capabilities</u> to cache and accelerate the delivery of your <u>HTTP-based file</u> <u>content</u> such as video, images, CSS, Javascript files, as well as HTML and API responses. Specifically:

HTTP header controls. Full-Site Delivery obeys standard HTTP caching headers and support forwarding, <u>adding, removing</u>, <u>and modifying the HTTP headers</u> we receive from your origin servers and send to end users, allowing you to send one set of instructions to your Fastly services and another set of instructions to downstream caches, proxies or browsers.

• **Time to Live controls.** Content expiration is controlled via Time to Live (TTL) <u>settings you configure</u> that work as timers on your cached content. You have the option of configuring a global default TTL to control cached content which, when set, will cache objects in a consistent manner even if you have multiple origins or server applications with inconsistent TTL settings.

- Request collapsing. When your content expires, the fetch and refresh process from your origin may take one second or more.
 During that time, your Full-Site Delivery may receive dozens or hundreds of end-user requests for that content. Fastly's request collapsing feature groups those requests and fulfills them together when it receives the refreshed content from your origin. Request collapsing decreases load on your origin servers by keeping your Fastly services from sending duplicate requests for the same expired content to them. Request collapsing is enabled by default.
- Grace mode (Serving stale content). If your origin servers become unavailable for any reason, grace mode can instruct your
 Fastly services to continue to serve stale or expired (but likely still valid) content to end users for a set amount of time. This
 allows you some extra time to return your unavailable servers to normal operations while still serving content instead of error
 messages to end users. Grace mode is not configured by default. To enable it, you must specifically configure your services to
 serve stale content.
- **Purging.** For <u>dynamic or event-based content</u> that doesn't lend itself to predetermined TTL-based content expiration, you can proactively remove or invalidate your content within milliseconds with Fastly's <u>purging features</u>. We limit purging to an average of 100K purges per hour per customer account, inclusive of all services within that account.

Edge logic and advanced content delivery control

Fastly's content delivery capabilities are based on a heavily extended version of the <u>Varnish</u> caching software. Varnish software gives you direct access to content delivery, control and edge logic capabilities, via the expressive HTTP inspection and modification scripting language, <u>Varnish Configuration Language</u> (VCL).

Streaming content delivery

Fastly's Streaming Delivery allows you to stream live and video-on-demand streaming content by leveraging Fastly's native support of common streaming formats. Fastly streaming format support includes HTTP Live Streaming (HLS), HTTP Dynamic Streaming (HDS), Dynamic Adaptive Streaming over HTTP (MPEG-DASH) and HTTP Smooth Streaming.

Precision Path

<u>Precision Path</u> traffic routing proactively identifies network congestion and poorly performing paths and automatically switches your traffic over to better performing alternatives. This improves service availability and resilience. Provisioned at strategic locations across our global fleet, this feature is available to all Fastly customers as part of our platform.

Origin shielding

You can designate a Fastly point of presence (POP) to <u>serve as a shield</u> for your origin servers, thus enabling increased cache hit rates for your Fastly services and potentially protecting your origin servers from unexpected spikes in requests for content. You can optimize this shielding geographically by configuring different shield POPs for different origin server locations. Origin shielding is not enabled by default. To use it, you must specifically <u>enable it</u>.

Load balancing

Services configured with multiple origin servers will <u>automatically distribute requests</u> to those servers evenly. You can modify this default load balancing behavior with a variety of conditions and load balancing rules.

Health checks

The health of your origin servers can be monitored with <u>configurable health checks</u> to help ensure only responsive origin servers are being sent requests.

Fastly web interface

All Fastly accounts have access to <u>Fastly's web interface</u>, allowing it to be <u>managed by multiple users</u> within your organization. You can control each user's role, as well as control the scope of their service access and their specific permission levels. Fastly services can be created, <u>monitored</u>, and managed through the Fastly Web Interface via any standard, modern web browser.

Application programming interface (API)

Fastly provides an <u>application programming interface</u> (API), accessible via HTTPS, through which Fastly services can be created and configured, and customers can access account information and analytics.

Real-time log streaming

To help you tune the performance of your Fastly services, we support <u>real-time log streaming</u> to a variety of locations, including third-party services, for storage and analysis. You can find our supported logging endpoints in our <u>list of streaming log guides</u>. We limit real-time log usage to a monthly average of two log statements per request, per service. If you require a higher volume of logs, Fastly offers <u>High Volume Logging</u>.

Transport Layer Security

Fastly supports a variety of <u>Transport Layer Security (TLS) services</u> that allow websites and applications to serve traffic over HTTP Secure (HTTPS), providing added privacy and data security for your services and end users. All Fastly services have access to our free shared domain option, plus a variety of additional paid TLS services to meet your TLS business and technical needs.

Always-on DDoS mitigation

Fastly's globally distributed network was built to absorb Distributed Denial of Service (DDoS) attacks. As part of Fastly's standard, Full Site Delivery, all customers receive access to a combination of features inherent in Fastly Edge Cloud network capabilities that help protect the availability of your content from DDoS threats.

- Access to origin shielding. Fastly allows you to designate a specific point of presence (POP) to host cached content from
 your origin servers. This POP acts as a <u>shield</u> that protects those servers from every cache miss or pass through the Fastly
 network, reducing the load that directly reaches them.
- Automatic resistance to availability attacks. Before they're even processed by our caching infrastructure, we filter out Layer 3 and 4 attacks (e.g., Ping floods, ICMP floods, UDP abuse) as well as distributed reflection and amplification (DRDoS) attacks that rely on anonymity to abuse internet protocols (e.g., DNS and NTP).
- Access to Fastly cache IP space. Fastly provides an API endpoint to any customer who would like to know which IP
 addresses our caches will use to send traffic from our CDN to your origin servers. We make this data available so you can
 update firewalls at your origin to ensure only our cache traffic can access your resources.
- Custom DDoS filter creation abilities. Using <u>custom VCL</u>, we allow you to craft your own DDoS protection rules to protect
 your network from complex Layer 7 attacks. Once you identify signs of a potential DDoS attack, you can <u>mix and match Fastly</u>
 <u>VCL with custom VCL</u> to construct filter configurations based on a variety of client and request criteria (e.g., headers, cookies,
 request path, client IP, geographic location) that block malicious requests before they hit your origin servers.

In addition to these standard DDoS protection services, Fastly offers a <u>DDoS Protection and Mitigation Service</u>. For more information about this or any of our advanced services, including their subscription costs, contact <u>sales-ddos@fastly.com</u>.

Pricing and billing

Full-Site Delivery <u>prices</u> are based on the volume of content delivered to your end users and the location of the POPs from which that content was served. <u>Fastly billing</u> is done in arrears based on actual usage with month-to-date usage being available via both our web interface and APIs.



NOTE

Fastly maintains partnerships with Google and Microsoft that may provide discounts on outbound data transfer traffic to customers who qualify and configure their Fastly services correctly. See our <u>integrations guides</u> for additional details.

- Fastly's On-the-Fly Packaging service
- **iii** Last updated: 2021-12-09
- https://docs.fastly.com/products/fastlys-onthefly-packaging-service

Fastly offers an "on-the-fly," dynamic, video-on-demand content packaging service. Rather than requiring you to pre-package all protocols of a viewer-requested video, Fastly allows you to dynamically package video content in different HTTP streaming formats in real time, using source files. That video content then becomes immediately available to viewers.



IMPORTANT

Fastly's On-the-Fly Packager (OTFP) for On Demand Streaming service is an add-on service. Our Professional Services team will assist with configuration and testing. To enable OTFP and begin this process, contact your account manager or email sales@fastly.com for more details.

Supported on-the-fly packaging features

Fastly's OTFP service supports the following specific features.

Supported HTTP streaming formats and codecs

- HDS, HLS, and MPEG-DASH packaging. Fastly provides support for version 1 of the Adobe HTTP Dynamic Streaming (HDS) specification and support for the ISO/IEC 23009-1:2014 specification defining Dynamic Adaptive Streaming over HTTP (MPEG-DASH). We support all features included in up to version 3 (draft 6) of the HTTP Live Streaming (HLS) specification and popular features from later versions such as subtitle, trick play and media segmentation in fragmented MPEG-4 (fMP4) format (per ISO/IEC 14996-12:2015 specification).
- Standard codecs. Fastly supports Advanced Video Coding (H.264/AVC/MPEG-4 Part 10) and High Efficiency Video Coding (H.265/HEVC) video codecs. Fastly also supports Advanced Audio Coding (AAC, AAC-LC, HE-AAC), Dolby Digital (AC-3) and MPEG-1 Audio Layer III (MP3) audio codecs.
- Source video container format. Fastly supports the Progressive MP4 specification (specifically the .mp4, unencrypted .mov, and audio-only .m4a extensions) as source container format for packaging into all supported HTTP streaming formats.

Accessibility and user experience

- HLS multi-language subtitles and closed captions. Fastly provides support for both in-band (EIA-608 and CEA-708) and out-of-band (Web Video Text Tracks or WebVTT) subtitle and closed caption delivery.
- HLS trick play. Fastly supports trick play (also called trick mode), a feature that displays video scenes during fast-forwarding and rewinding. The HLS Authoring Specification requires this feature for distributing video on the Apple TV.

Content protection

- Media encryption. Fastly can encrypt videos packaged into HLS (supports both Envelope/AES-128 and SAMPLE-AES methods) and MPEG-DASH (ISO/IEC 23001-7, a common encryption in ISO base media file format file) streaming formats by generating a unique content encryption key for each video, enabling secure video delivery to viewers.
- Multi-DRM. Fastly can support multiple Digital Rights Management (DRM) technologies including <u>Apple FairPlay</u> for HLS and Microsoft PlayReady, Google Widevine and Marlin DRM for MPEG-DASH streaming formats. OTFP is integrated with Multi-DRM service providers that are responsible for content rights management and DRM license delivery.

Dynamic Ad Insertion (DAI) readiness

- HLS timed metadata injection. Fastly supports HLS time-based metadata, which allows you embed custom metadata or ad markers about a stream into video segments at specified time instances in ID3v2 format.
- Content preconditioning. Fastly can segment video at the intended break points, such as for ad markers via HLS and MPEG-DASH protocols. Fastly can also add any third-party service-specific cues or metadata into video manifests at those break points to implement server or client-side ad stitching.

Clip creation

 Clip creation (also known as "timeline trimming"). Fastly supports clip creation features for all supported packaging formats, allowing you to deliver sections of video without segmenting a longer, archived video. Time guery parameters ("start" and "end") allow you to break up videos into discrete sections so users don't have to find the relevant section using the timeline.

Standard content delivery network features

Fastly also provides the following features as part of standard content delivery network services:

- <u>Token-based validation</u> for decreasing response time by placing validation at the edge
- Geolocation and device detection for content targeting
- <u>Dictionaries</u> for real-time business rules and decision making at the edge
- Remote log streaming for data aggregation and viewer diagnostics
- Transport Layer Security (TLS) for secure communications delivery

How the on-the-fly packaging service works

Fastly's OTFP service gets configured between our caching network and your origin storage (e.g., Amazon S3, Google Cloud Storage, or Rackspace Cloud Files).



When users request manifests or video segments, those requests initially come to Fastly caches instead of going to your origin storage. Fastly's edge caches deliver those objects if they are available and valid. If the objects don't already exist in the edge caches, the requests will be passed on to a designated shield cache to be delivered instead as long as the objects are available and valid. If neither the edge caches nor the shield cache can deliver the objects, the requests for those objects will go directly to and be fulfilled by the OTFP service which acts as an origin for Fastly's cache nodes.

The OTFP service will make the necessary request to your origin storage to fulfill the original request from the user. The OTFP service also maintains a small, local, in-memory cache for video metadata indexes. These indexes are created using mp4 moov atom (or movie atom) that provide information about the video file such as its timescale, duration, audio and video codec information, and video resolution (among other characteristics).

For adaptive bitrate playback, the OTFP service will cache indexes of each quality level requested. If a user requests a manifest, OTFP will look for the corresponding indexes and, if it is available and valid, OTFP will generate the manifest and deliver it to the user. Otherwise, OTFP will fetch the moov atom from origin storage to generate the corresponding index. If a user requests video segments, OTFP will look for the corresponding audio and video sample entries in the cached index, download those samples from origin storage, and package them in the format requested.

- Fastly's Streaming Delivery
- Last updated: 2019-08-09
- https://docs.fastly.com/products/fastlys-streaming-delivery

Fastly's Streaming Delivery allows you to scale your streaming content delivery when you will not be using your Fastly services for any of the other HTTP content formats supported by Fastly's Full Site Delivery.



O NOTE

Fastly's Streaming Delivery is a subset of Fastly's Full-Site Delivery and it must be configured in an account separate from other Fastly accounts to allow for separate billing and invoices.

If you have your own video packaging infrastructure, Fastly can act as a globally distributed HTTP streaming network to improve quality of service and increase viewer capacity for both your live and Video On Demand (VOD) content. When a manifest or video segment is requested by an end user's player, your Fastly Streaming Delivery will pull the requested content from your origin media servers and subsequent requests for that stream will be served from Fastly's POPs instead of your origin servers.

Fastly's Streaming Delivery supports the following HTTP-based media streaming protocols:

- HTTP Live Streaming (HLS)
- HTTP Dynamic Streaming (HDS)
- HTTP Smooth Streaming (HSS)
- Dynamic Adaptive Streaming over HTTP (MPEG-DASH)

You can also use Fastly's Full Site Delivery to configure and control live streaming and VOD caching.



NOTE

Fastly maintains partnerships with Google and Microsoft that may provide discounts on outbound data transfer traffic to customers who qualify and configure their Fastly services correctly. See our <u>integrations guides</u> for additional details.

- Image Optimizer
- **±** Last updated: 2020-09-25
- https://docs.fastly.com/products/image-optimizer

The <u>Fastly Image Optimizer (Fastly IO)</u> is a real-time image transformation and optimization service that caches and serves pixel-optimized, bandwidth-efficient images requested from your origin server. Fastly IO specifically supports a variety of <u>input and output image formats</u>.



Image transformation and optimization

When an image is requested from your origin server, Fastly IO can perform <u>transformation tasks</u> before serving and caching the optimized version. Image transformations can be applied programmatically and through dynamic URLs in real-time. You can <u>make images responsive</u> so they automatically adjust to fit the size of the screen viewing the content. As a result, image pre-processing can be offloaded to the edge. Multiple copies of the images, each appropriately sized for different devices, are served from cache instead, which allows you to reduce the number of requests to your origin.

Debugging and troubleshooting

To aid in debugging when serving images, <u>special HTTP headers</u> will be present in a response when an image is requested. The specific header included depends on the response's result. For successful transformations and optimizations, the HTTP header returned provides general information that allows you to compare image dimensions, file sizes, and formats. Additional HTTP headers are included for source image issues that aren't fatal enough to cause an error but could still be problematic, as well as transformations and optimizations that fail outright.

Billing

Billing for Fastly IO is based on the number of monthly image requests that are processed and delivered. When using the animated GIF to video functionality, each image frame delivered as video is counted as an optimized image request.

This article describes a product that may use third-party cloud infrastructure to process or store content or requests for content. For more information, see our <u>cloud infrastructure security and compliance program</u>.

Media Shield for Live

± Last updated: 2020-12-18

https://docs.fastly.com/products/media-shield-for-live

Fastly Media Shield for Live offers customers the ability to decrease origin traffic by <u>reducing multiple CDN requests</u> of live video events or live linear channels into a single request back to your origin. Media Shield for Live works with your existing architecture by making Fastly the origin to all of your end-user-serving CDNs. This also allows you to take advantage of Fastly's observability features like <u>real-time analytics</u>, <u>historical stats</u>, and <u>real-time logging</u> in a multi-CDN environment.

To learn more about Fastly's Media Shield for Live, contact your account manager or email sales@fastly.com for details.

Media Shield for VOD

Last updated: 2020-12-18

https://docs.fastly.com/products/media-shield-for-vod

Fastly Media Shield for VOD offers video-on-demand customers the ability to decrease origin traffic by <u>reducing multiple CDN</u> <u>requests</u> into a single request back to your origin. Media Shield for VOD works with your existing architecture by making Fastly the origin to all of your end-user-serving CDNs. Fastly Media Shield for VOD is compatible with Fastly's <u>On-the-Fly-Packaging (OTFP)</u> <u>service</u>.

Media Shield for VOD allows you to take advantage of Fastly's observability features like <u>real-time analytics</u>, <u>historical stats</u>, and <u>real-time logging</u> in a multi-CDN environment.

To learn more about Fastly's Media Shield for VOD, contact your account manager or email sales@fastly.com for details.

Nearline Cache

苗 Last updated: 2021-04-29

https://docs.fastly.com/products/nearline-cache

IMPORTANT

This information is part of a limited availability release. For additional details, read our <u>product and feature lifecycle</u> descriptions.

Fastly's Nearline Cache allows you to manually or automatically populate and store content in third-party cloud storage near a Fastly POP. In the event of a cache miss, content that has previously been written to Nearline Cache is fetched from the storage instead of your origin. Content not present in Nearline Cache will be fetched from your origin and, if configured for automatic migration, will be written to Nearline Cache asynchronously.

There is no charge for data transfer costs between the third-party cloud storage and Fastly. Availability and performance of Fastly's Nearline Cache depends upon the third-party cloud storage provider's services. You are responsible for <u>purging</u> or otherwise rendering inaccessible any stored content that you do not intend to serve to end users. Content in the third-party cloud storage will be programmatically deleted by Fastly after a specified period.

Billing

When purchasing Nearline Cache, you must also purchase **Gold** or **Enterprise** Support.

You will be charged for content in Nearline Cache based on the average amount of data in gigabytes (GB) in Nearline Cache during the billing period and for write operations as specified on your service order.

Origin Connect

a Last updated: 2021-08-05

https://docs.fastly.com/products/origin-connect

Origin Connect provides you with a direct fiber connection between your origin servers and a Fastly shield POP thus reducing the number of organizations (and by association, the number of servers) handling your data.

Prerequisites

To be considered for Origin Connect, you need to:

- have at least one <u>Fastly shield POP</u> configured
- have servers in the same data center as the selected Fastly shield POP (e.g., IAD, AMS, SJC)
- be interviewed by Fastly so we can identify your customer-specific business needs
- have **Enterprise-level support**
- have a publicly routed Autonomous System Number (ASN)

If you are approved for Origin Connect, we'll issue you with a Letter of Authorization and Connecting Facility Assignment (LOA-CFA) that the data center provider will need when you order your cross-network connection (or cross connect). You will need to pay for the cross connect with your facility provider.

For each cross connect, you, as subscriber, will need to provide Fastly with:

- a minimum of a globally unique (non RFC-1918) /31 IPv4 network prefix
- a minimum of a /127 IPv6 network prefix
- a 10G port (we recommend two and will accept up to 4x 10G ports for redundancy)

Both you, as the subscriber, and Fastly will each need to:

- provide the ASN intended for Border Gateway Protocol (BGP) peering use
- provision BGP peering on each interconnect
- provide a BGP prefix filter list
- comply with any other reasonable request to technically provision the Origin Connect product

If the cross connect is not completed within 90 days, the authority granted by the LOA-CFA expires.

In the event of Origin Connect service degradation, congestion, or a failure of one of these interconnects, public internet transit will be used for origin connectivity, and the subscriber will prefer the carrier of Fastly's reasonable request. There is no Service Level Agreement (SLA) available for Origin Connect.

If your origin server is located within a cloud storage provider or your traffic doesn't meet our minimum threshold for Origin Connect, contact us at sales@fastly.com to discuss other options.

- **Subscriber Provided Prefix**
- **t** Last updated: 2021-09-01
- https://docs.fastly.com/products/subscriber-provided-prefix

Fastly's Subscriber Provided Prefix product allows you to have your IP spaces announced, routed, and served by Fastly infrastructure for use with production services. When you purchase this product, you provide your own IP address space to Fastly rather than use Fastly IP addresses. You can then direct traffic to your own IP addresses, which are reachable via HTTP Anycast on Fastly's infrastructure.

We recommend this service for customers who want to control their address space by separating their network layer concerns from their content delivery concerns. By combining Fastly's Subscriber Provided Prefix service with our <u>Origin Connect</u> product and our <u>DDoS Protection and Mitigation service</u>, you can protect your origin servers by directing traffic through Fastly's global network.

For more information about this product, contact sales@fastly.com.

Prerequisites

To purchase Fastly's Subscriber Provided Prefix service you must also purchase Fastly's <u>Enterprise Support</u> package and our IP-to-Service Pinning Setup service.

When you sign up for this product, you'll need to provide Fastly with an executed Letter of Authorization (LOA), on a form we provide, that grants us permission to announce your prefixes. The LOA includes, at a minimum, the IP blocks to announce, the registry and object identifier, as well as the administrative, technical, and abuse contacts for those prefixes.

Using the Subscriber Provided Prefix product requires at least one /24 IPv4 or /48 IPv6 prefix for announcement purposes. Additional prefixes and larger prefixes may also be supported. These prefixes must not be originated from any autonomous system number (ASN) at the time Fastly announces them. They should also be dormant for a period of approximately three months prior to use by Fastly.

How the Subscriber Provided Prefix product works

Fastly will announce the designated prefixes identified in your LOA. Your prefixes will be announced along with existing Fastly prefixes and will be originated from the Fastly Autonomous System (AS) Number AS54113. The Subscriber Provided Prefix product supports HTTP and HTTPS traffic only and your prefixes will be terminated at Fastly for these two protocols. We make routing announcements on a global basis unless you request they be constrained to our defined North America and Europe region.

To enable specific IP addresses within your announced prefix, Fastly combines this Subscriber Provided Prefix product with our IP-to-Service Pinning feature, which must be purchased separately. IP addresses that are not service pinned will not serve your traffic.

After completing all the necessary routing announcements and setup within your CDN services, Fastly needs additional time to complete the setup. In general, you should allow for at least 45 days of lead time for us to set up routing announcements and CDN service. Your service order identifies the specific lead time Fastly needs for full operability.

You may provide Fastly notice at any time to withdraw your prefix announcement by opening a ticket at support@fastly.com. We need at least 45 days of notice to permanently remove routing announcements and CDN service for your designated prefixes. When we receive notice of your request for prefix withdrawal, we will provide you with a withdrawal process timeline. This process starts with us reconfiguring your service within the Fastly network. When that reconfiguration work completes, you must then point your DNS records at Fastly to ensure uninterrupted service. Once your traffic is moved from your prefix to a Fastly prefix, we will withdraw the announcement.

Conditions and limitations

When using Fastly's Subscriber Provided Prefix product you agree to the following limitations:

- Your purchase of the Subscriber Provided Prefix product entitles you to the announcement of the specified IP prefixes identified in your LOA. Any additional prefixes beyond your initial order will require an additional purchase of this product.
- Fastly does not provide termination or proxy services for non-HTTP and non-HTTPS protocols with this product.
- Fastly does not provide general network transit or peering services as part of this product.

When using Fastly's Subscriber Provided Prefix product you agree to the following conditions:

- Your IP addresses are your assets. They belong to you and are not a Fastly service. Fastly has no liability for your assets.
- You will pay additional fees if you withdraw your prefixes for the purpose of replacing or updating them.
- Your provided prefixes will not have any negative IP reputation associated with them as determined by us. Fastly will scan
 your prefixes against common IP reputation databases prior to announcement to ensure your IP reputation remains neutral or
 positive.
- You must maintain transit connectivity to Fastly for origin traffic. Prefixes provided to Fastly for this service must not overlap with IP addressing used by your origin servers.
- Fastly retains exclusive announcement rights for your prefixes. Conflicting announcements will disrupt or prevent traffic delivery.

To specifically mitigate DDoS attacks, you agree that:

- Prefix announcements Fastly makes for you may include regional capacity announcements.
- Fastly may prepend, remove, or blackhole routing announcements in the event of a DDoS attack.
- Fastly may de-aggregate your prefixes at our discretion to improve network reliability.

Fastly may perform these actions even if you have not purchased the <u>Fastly DDoS protection and mitigation service</u>.



1 NOTE

For any IP addresses not pinned to a service but contained within your Subscriber Provided Prefix product, Fastly's Varnish servers will return a TCP reset or an HTTP 500 error response code.

- Varnish Configuration Language (VCL)
- Last updated: 2020-06-10
- https://docs.fastly.com/products/varnish-configuration-language

Fastly Varnish Configuration Language (VCL) is a domain-specific language derived from the Varnish proxy cache, which is part of Fastly's platform architecture.

When you create and enable a VCL service via the <u>web interface</u> or the <u>API</u>, Fastly generates VCL code automatically for the functionality you've specified. You can refine and augment that functionality as needed to create customized configurations by defining your own VCL logic and combining it with the automatically generated VCL. Those configurations are then distributed to all Fastly caches worldwide, loaded, and activated without requiring maintenance windows or service downtimes.

Documentation

Fastly VCL provides a variety of language components, which we document in our VCL reference documentation. The documentation describes each component, including its syntax and structure where necessary. Where applicable, examples of each component's use are included.

Custom VCL

You can create a custom VCL configuration to augment or override the default VCL configuration for your service. To make using custom VCL easier, we maintain a VCL boilerplate that you can use as a starting point when including your own VCL logic. You can use the web interface to <u>upload</u> custom VCL files for your service as well as preview it prior to activation.

VCL Snippets

VCL Snippets are short blocks of VCL logic that can be included directly in your service configurations. They're specifically designed to allow you to add small sections of code when you don't need more complex, specialized configurations.

Security

These articles provide information about Fastly products that deliver web application and API protection.

- https://docs.fastly.com/products/product-security
- DDoS Protection and Mitigation Service and SLA
- Last updated: 2018-04-25 **=**
- https://docs.fastly.com/products/ddos-protection-and-mitigation-service-and-sla ത

Fastly offers DDoS Protection and Mitigation Service to customers with a sustained DDoS threat risk or with short term and seasonal events to protect. While the DDoS Protection and Mitigation Service cannot prevent or eliminate attacks or guarantee the uptime of your origin servers, it offers the following resources to assist you with mitigating the service and financial impacts of DDoS and related attacks.

Fastly's DDoS Protection and Mitigation Service includes:

• Immediate onboarding - We will work directly with you to immediately transition you to Fastly's CDN service if you're not already a customer.

- Emergency configuration and deployment support We will actively work with you to configure your service map and provide an initial filter policy to immediately block an attack.
- Ongoing attack mitigation support We will work directly with you to write custom VCL filters to deal with changing attacks or new attacks. We'll also isolate malicious traffic on your behalf.
- Incident response plan We will create a plan that identifies how communication and escalation will occur between you and your staff and Fastly if an attack occurs. The plan will also describe mitigation and defense details such as any DDoS filters that we can insert into VCL prior to or during an attack.

Using our knowledge of attacks against our network and our customers, we analyze all DDoS Attack vectors using VCL statements, network filters, bulk traffic filtering through regional sinks, or a combination of these techniques.

The following table summarizes what is provided under our DDoS Protection and Mitigation Service:

Support offering	Details
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. Service onboarding beginning within 60 minutes of threat notification.
Overage Insurance	Included.
Access to Fastly IP Space	Included.
Email support	Available.
Phone and chat support	Toll-free telephone available 24/7/365. Dedicated chat channel available during Fastly Business Hours.
Emergency escalation	Available via email and phone support.

Technical support

The following section applies to all Subscribers of the DDoS Protection and Mitigation Service.

Definitions

- "Business Hours" are 8AM-6PM during a Business Day in California, New York, London, or Tokyo.
- "Business Days" are Monday through Friday, excluding any day that is simultaneously a US, UK, and Japanese national or banking holiday.
- A "DDoS Attack" is a Denial of Service (DoS) event (including Distributed Denial of Service (DDoS) or Distributed Reflection Amplification Denial of Service (DRDoS) attacks) that includes both an increase of unwanted traffic beyond two (2) times the average traffic of any Fastly Service for the preceding two (2) month period and a simultaneous increase in error responses from origin sites configured for any Fastly service. Fastly captures and analyzes suspected or actual DDoS Attack traffic to improve and protect its services.
- A "Fastly IP Space" is a <u>published API endpoint</u> that allows Subscribers to download an updated list of all Fastly IPs globally and can be used to filter traffic and control communication between Fastly's caches and a Subscriber's origin. Fastly provides the Fastly IP Space to Subscribers in order to ensure known communication between the Fastly cache nodes and a Subscriber's origin data center.
- "Fastly Control" means elements entirely under Fastly's control and not a consequence of (a) a Subscriber's hardware or software failures, (b) a Subscriber's or end user's connectivity issues, (c) Subscriber operator errors, (d) Subscriber traffic

amounts that exceed a Subscriber's Permitted Utilization as defined in the Terms and Conditions, (e) corrupted Subscriber content, (f) acts of god (any) or war, or earthquakes, or terrorist actions.

Subscriber responsibilities

As a Subscriber, you:

- must identify and maintain two points of contact to be used during an attack to communicate status, issues, and coordinate with Fastly to successfully protect services.
- must use common best practices for DDoS Attack defense including:
 - using updated allow and block lists in the Fastly IP Space at the origin data center to protect against attack traffic bypassing Fastly's infrastructure.
 - limiting or eliminating your origin IP addresses from Domain Name System (DNS) records to avoid these addresses being used as attack targets.
- are responsible for using and configuring services according to the documentation available at https://docs.fastly.com.

Support requests

Subscribers may make support requests by submitting a <u>support ticket</u> which will trigger a system-generated acknowledgement within minutes containing the ticket number and a direct link to the ticket.

DDoS Attack reports should include at least:

- a determination of the severity of the attack.
- the size of the attack threatened or previously observed.
- the type and vector of attack traffic seen or threatened.
- any duration of previous attacks and vector behavior including major source IP addresses.
- attack history for the last 24 months.
- threat specifics including all details of any attacks that the protected services or sites have experienced in the past.

Communications and channels of support

Support tickets

Create support tickets by sending an email to support@fastly.com or calling our dedicated phone line. Filed tickets trigger Fastly's promised response time.

Tickets for communication between Fastly support engineers and a Subscriber's personnel are tracked using a ticketing application, which maintains a time-stamped transcript of communications, and sends emails to Subscriber and Fastly staff as tickets are updated.

Phone support

Subscribers to the DDoS Protection and Mitigation Service receive a dedicated phone number to contact Fastly support engineers. Fastly personnel can also establish audio and video conferencing (free app or browser plug-in required) for real-time voice and video communications.

Chat

To facilitate real-time communication, Subscribers to the DDoS Protection and Mitigation Service receive a dedicated chat channel for real-time communications during Business Hours or as needed by Fastly personnel. Though subject to change, Fastly's current chat provider is Slack (www.slack.com).

Attack traffic

Response time

Fastly commits to responding to DDoS Attack notifications from Subscribers within 15 minutes of notice and, as applicable, will begin on-boarding Subscribers to the DDoS Protection and Mitigation Service within 60 minutes of a DDoS Attack notification.

Related Invoice Credits

Fastly will waive all bandwidth and request charges associated with DDoS Attack traffic and will provide Invoice Credits or adjustments for the same.

Attack traffic credit terms

Subscribers must submit claims for waiver of DDoS Attack-related charges to billing@fastly.com within 30 days of the DDoS Attack.

DDoS Mitigation response SLA

If, during a DDoS Attack on a Subscriber with the DDoS Protection and Mitigation Service, there is a material delay in response time and the cause of the delay is within Fastly's control, a one-time credit of \$500 per incident will be credited to that Subscriber's account.

SLA credit terms

- Requests for Invoice Credits must be made within 30 days of the DDoS Attack that triggered the service credit.
- All requests for Invoice Credits must be made to <u>billing@fastly.com</u>.
- In no event shall Invoice Credits exceed the fee for the DDoS Protection and Mitigation Service payable by a Subscriber for the month in which the Invoice Credits accrued.
- A pending Invoice Credit does not release a Subscriber from the Subscriber's obligation to pay Fastly's submitted invoices in full when due.
- Invoice Credits will be applied to the invoice within the month the credits were incurred.

Termination for SLA

For a Subscriber of the DDoS Protection and Mitigation Service with a <u>Termed Contract</u>, if in any three-month period where three (3) or more support response time objectives are not met and the failure to meet the objectives materially adversely impacted the Subscriber, the Subscriber will have 30 days to terminate the DDoS Protection and Mitigation Service subscription following the third response failure. Subscribers must notify Fastly of their intention to terminate the DDoS Protection and Mitigation Service subscription within 30 days of the triggering event.

- Fastly Next-Gen WAF (powered by Signal Sciences) support description and SLA
- **tast updated: 2022-05-31**
- https://docs.fastly.com/products/fastly-next-gen-waf-support-description-and-sla

Support availability and response times for the <u>Fastly Next-Gen WAF (powered by Signal Sciences)</u> vary depending on the type of account you have and the platform you have purchased.

IMPORTANT

This page was previously named *Signal Sciences WAF support description and SLA*. The Fastly Next-Gen WAF now collectively refers to the products that were previously known as the Signal Sciences Cloud WAF and Signal Sciences Next-Gen WAF. The functionality of those products has not changed as part of the new naming convention. Fastly Next-Gen WAF continues to be powered by Signal Sciences technology.

Support channels and response times

Fastly provides customer support via the support channels listed below.

Support tickets

Fastly Next-Gen WAF (powered by Signal Sciences) includes access to a support portal that allows you to submit requests for support online, update existing support tickets, and track the status of support tickets. As part of submitting a request via the support portal, you may designate a proposed severity level for the issue being reported, but the ultimate classification of a request will be determined by Fastly based on various factors including input you provide.

Email

Fastly's technical support staff can be contacted via <u>support@signalsciences.com</u> during standard business hours. All support tickets generated by email will be designated with a P2 severity level.

Response times

The following table summarizes the response times based on the severity of the reported issue and the platform you have purchased.

Issue severity level	Essential	Professional	Premier
Urgent - Critical Impact (P0)	1 business day with updates every 4 business days	60 minutes or less with updates every 2 hours (24/7/365)	60 minutes or less with updates every 2 hours (24/7/365)
High - Serious Impact (P1)	1 business day with updates every 4 business days	4 business hours with updates every 12 business hours	4 business hours with updates every 12 business hours
Normal - Minor Impact (P2)	1 business day with updates every 4 business days	1 business day with updates every 4 business days	1 business day with updates every 4 business days
Low - Minor Impact (P3)	2 business days with no further updates	2 business days with no further updates	2 business days with no further updates

- Business hours are 4 AM-7 PM Monday through Friday, Pacific Time.
- Business days are Monday through Friday, except standard Fastly holidays.
- **Urgent Critical Impact (P0)** issues have confirmed errors in a production environment that make a solution, its features, or its functionality completely unavailable to users.
- **High Serious Impact (P1)** issues have confirmed errors in a production environment that cause significant loss of functionality for a primary feature of a solution that has notable impacts to customer business.
- **Normal Minor Impact (P2)** issues have confirmed errors in a production environment that cause partial loss of functionality of a non-significant feature or a significant cosmetic issue with the web interface. This severity level also applies to identified errors in a non-production environment.
- Low Minor Impact (P3) issues have confirmed errors that cause minor cosmetic issues with the web interface. This severity level also applies to feature requests and general questions about functionality.
- A **confirmed error** is any failure of the Next-Gen WAF to meet Fastly's specifications outlined in the relevant documentation, found in production uses, and that can reasonably be reproduced by Fastly.



IMPORTANT

Customers purchasing at the Essential platform level are limited to 4 hours of support per month.

SLA

Fastly endeavors to maintain 99.9% availability of the Signal Sciences Hosted Dashboard (<u>Hosted Dashboard</u>), which is used by the Fastly Next-Gen WAF and the hosted infrastructure components of the Edge and Cloud WAF deployment method.

SLA for Hosted Dashboard

Subscribers experiencing unavailability of the Signal Sciences Hosted Dashboard will be entitled to service credits according to the following table.

Monthly Availability of Hosted Dashboard	Service Credit % of Pro-rated Monthly WAF Subscription Fees
<99.9-99.0	5%
<99.0%-98.5%	10%
<98.5%-98.0%	15%
<98.0%	20%

"**Availability**" of the Hosted Dashboard is calculated as follows: ([# of minutes in month]-[# of minutes per month the Hosted Dashboard is Unavailable]) / [# of minutes in month].

"**Unavailable**" with respect to the Hosted Dashboard means the Hosted Dashboard is not available for your access and use through your internet connection, excluding (a) unavailability of the Hosted Dashboard caused by issues not under Fastly Control or (b) unavailability that does not last for a consecutive ten-minute period.

"Fastly Control" means elements entirely under Fastly's control and not a consequence of (a) your hardware or software failures or those of a third party, (b) you or your end users' connectivity issues, (c) operator errors by your employees, (d) alteration, modification, unauthorized or misuse of Fastly products, or use not in accordance with the documentation, (e) your corrupted content, or (f) acts of god (any) or war, or earthquakes, or terrorist actions.

SLA for Cloud WAF Hosted Services

Subscribers experiencing unavailability of the hosted infrastructure component of Cloud WAF Hosted Services will be entitled to service credits according to the following table.

Monthly Availability of Cloud WAF Hosted Services	Service Credit % of Pro-rated Monthly Hosted Services Subscription Fees
<99.9-99.0	5%
<99.0%-98.5%	10%
<98.5%-98.0%	15%
<98.0%	20%

"Availability" of the Hosted Services is calculated as follows: ([# of minutes in month]-[# of minutes per month the Hosted Services is Unavailable]) / [# of minutes in month].

"**Unavailable**" with respect to the Hosted Services means the Hosted Services are not available to process traffic or communicate with Hosted Dashboard, excluding (a) unavailability caused by issues not under Fastly Control or (b) unavailability that does not last for a consecutive ten-minute period.

"Fastly Control" means elements entirely under Fastly's control and not a consequence of (a) your hardware or software failures or those of a third party, (b) you or your end users' connectivity issues, (c) operator errors by your employees, (d) alteration, modification, unauthorized or misuse of Fastly products, or use not in accordance with the documentation, (e) your corrupted content, or (f) acts of god (any) or war, or earthquakes, or terrorist actions.

Additional Terms

Fastly may temporarily limit or disable the inspection and blocking capabilities of the <u>Fastly Next-Gen WAF (Edge)</u> for your service if your traffic threatens to interfere with or disrupt the integrity or performance of Fastly's services. If this is necessary, the edge security service will <u>fail open</u> and your service will continue to serve traffic.

Credit terms

- You must contact us within 15 days of experiencing unavailability to receive a service credit.
- For any given month, the maximum amount of any credit is 20%, regardless of the reason it is owed.
- A pending credit does not release you from your obligation to pay Fastly's submitted invoices in full when due.

• Credits will be applied to the invoice two months following the month an invoice credit was incurred.

Fastly Next-Gen WAF (powered by Signal Sciences)

Last updated: 2022-06-22

https://docs.fastly.com/products/fastly-next-gen-waf

The Fastly Next-Gen WAF (powered by Signal Sciences) is a web application firewall that monitors for suspicious and anomalous web traffic and protects, in real-time, against attacks directed at the applications and origin servers that you specify.

Using default settings created by Fastly and custom settings you specify, the Next-Gen WAF identifies and tracks attacks across all of your deployments and determines whether to flag the originating IP address as potentially problematic, rate limit the IP address, allow the request, tag it with signals, or block it. You can choose to enable or disable these actions at any time. When the Next-Gen WAF determines that an incoming request is anomalous, we collect data from that request and upload it to our cloud engine, allowing us to perform out-of-band analysis of inbound traffic.

The Fastly Next-Gen WAF now collectively refers to the products that were previously known as the Signal Sciences Cloud WAF and Signal Sciences Next-Gen WAF. The functionality of those products has not changed as part of the new naming convention. Fastly Next-Gen WAF continues to be powered by Signal Sciences technology.

Deployment types

The Next-Gen WAF can be deployed in three different ways:

- On Fastly's Edge Cloud platform (Edge). To use the edge deployment method, you must add it to new or existing Fastly services that you create in the Fastly console and update your DNS records to point to Fastly.
- Directly on your web servers within your infrastructure (Core). The Core deployment method consists of two components, the module and the agent. The module can exist as a <u>plugin to your web server</u> or as a <u>language or framework specific</u> <u>implementation</u>. The agent is a small daemon process that provides an interface between your web server and our cloud engine. You can also use the core deployment method without a module by running the agent in <u>reverse proxy mode</u>.
- On Fastly's cloud-hosted infrastructure (Cloud WAF). To use Cloud WAF, you must upload a TLS certificate, add an origin server using the Signal Sciences Hosted Dashboard, and update your DNS records to point to the appropriate servers.

Hosted Dashboard

The <u>Signal Sciences Hosted Dashboard</u> (Hosted Dashboard) is a web interface that you can use to investigate anomalous web traffic and see what actions, if any, Next-Gen WAF performed in response to certain requests. You can also use the Hosted Dashboard to create Workspaces. A *Workspace* (also known as a Site) is a user-defined set of rules and settings for applications and origin servers. The Hosted Dashboard allows you to create multiple Workspaces to differentiate between one or more APIs, microservices, or web applications. For each Workspace, you can use the Hosted Dashboard to add rules for requests, configure site alert thresholds, and add integrations to other systems.

Threat intelligence

As part of Next-Gen WAF, we may <u>aggregate the attack data collected</u> from use of Next-Gen WAF and combine it with data collected from security and other services offered as part of the Fastly platform, including for other subscribers. We use these data insights (threat intelligence) to analyze and detect potential future anomalies or attacks and to improve, secure, provide, and market Fastly services in a manner that does not associate the threat intelligence with or identify any subscriber. For example, you receive the benefits of this threat intelligence via the Network Learning Exchange (NLX) feature that adds a unique signal to information in the Hosted Dashboard and alerts you to potential bad actors that have been identified elsewhere in the subscriber network.

API

The Signal Sciences Application Programming Interface (API) allows you to integrate your applications and services with the Next-Gen WAF. It uses standard HTTP response codes and verbs to allow you to programmatically control all the same features that are available through the Hosted Dashboard. The Signal Sciences API provides a variety of endpoints that we document in our <u>API reference documentation</u>.

Control over data sharing

Next-Gen WAF gives you control over data shared with Fastly. The <u>Hosted Services</u> component of the Cloud WAF deployment does not create copies of or store your data as it passes through. The hosted aspect of the Edge deployment similarly does not create copies of or store your data feed as it passes through.

The security components for all deployment types of Next-Gen WAF do not require transmission or collection of any sensitive or personally identifiable information to function other than IP addresses that are identified as the initiator of anomalous or suspicious requests and related metadata. The Next-Gen WAF is designed to automatically redact certain sensitive or personally identifiable information in fields that are known to commonly contain such information before transmission to the cloud engine component of the Next-Gen WAF. Also, the Next-Gen WAF allows you to manually configure which fields are redacted via the Hosted Dashboard to further limit the sensitive information or other information sent to the cloud engine component of the Next-Gen WAF, other than the limited data required for the functionality of the Next-Gen WAF. If properly configured, for Edge and Cloud WAF deployments, none of your sensitive information other than the IP addresses identified as the initiator of anomalous or suspicious requests will be sent to the cloud engine component of the Next-Gen WAF. For Core deployments of Next-Gen WAF, if properly configured, this means that none of your sensitive information other than the IP addresses identified as the initiator of anomalous or suspicious requests will be shared with Fastly.

DDoS mitigation

Edge and Cloud WAF deployments feature an always-on service integration that examines inbound traffic to detect and mitigate Distributed Denial of Service (DDoS) attacks before they reach the applications and origin servers that you specify.

Edge deployments receive access to a <u>combination of features</u> inherent in the Fastly Edge Cloud network that help protect from DDoS threats. This service requires no additional installation or maintenance.

Cloud WAF deployments use automated mitigation techniques to stop common network protocol-based floods including SYN floods and reflection attacks using UDP, DNS, NTP, and SSDP. This service requires no additional installation or maintenance.

Feature availability

Feature availability depends on the platform you have purchased.

Feature	Essential	Professional	Premier
Default attack signals	Included	Included	Included
Default anomaly signals	Included	Included	Included
Default dashboards	Included	Included	Included
<u>Virtual Patching</u>	Included (BLOCK only)	Included	Included
Custom response codes	Not Included	Included	Included
Custom signals	Not Included	Included	Included
Standard API & ATO signals	Not Included	Included	Included
Advanced Rate Limiting	Not Included	Not Included	Included
Edge Rate Limiting	Not Included	Not Included	Included but requires active <u>full-site delivery</u> account
<u>Deployment Types</u>	Edge Core Cloud WAF	Core Cloud WAF	Core Cloud WAF

Documentation

We provide documentation for the Next-Gen WAF in the <u>Signal Sciences Help Center</u>. Release notes for the agents and modules are also provided in the Signal Sciences Help Center.

Billing

We bill you as specified in your applicable ordering document. We measure months according to Coordinated Universal Time (UTC). All deployments are billed according to the number of Workspaces and the average requests per second (RPS) processed by Next-Gen WAF.

Edge deployments are additionally billed for delivery charges associated with the <u>Full-Site Delivery service</u> on which Edge deployments are hosted. <u>Prices</u> are based on the volume of content delivered to your end users and the location of the POPs from which that content was served. <u>Fastly billing</u> is done in arrears based on actual usage with month-to-date usage being available via both our web interface and APIs.

Cloud WAF deployments are additionally billed for the overall traffic flowing through the <u>Hosted Services</u> in terabytes (TBs) and the number and location of protected origins.

Subscriber responsibilities

From time to time, we may provide error corrections, bug fixes, software updates, and software upgrades to the agent and the module. Notices about updates are included in the <u>documentation</u> and described in the <u>release notes</u>. You can also <u>subscribe to receive emails from us</u> when updates are released or subscribe to our integrations with third-party tools (e.g., <u>Slack</u> or <u>Microsoft Teams</u>). For Core deployments, it is your responsibility to ensure that you are using the most recent version of the Next-Gen WAF components. Agents on Edge and Cloud WAF deployments are kept up to date by Fastly.

As a subscriber, you can identify and maintain up to five points of contact for support communications. All support requests must be initiated from and communicated through the designated points of contact.

Subject to the terms of any open source license applicable to any Fastly software installed in your environment (namely the agents and modules), your subscription for Next-Gen WAF does not include permission to modify the software or create derivative works based upon the software other than as set forth in the Documentation.

Limitations

All WAF products that exist today, including the Next-Gen WAF, have several limitations:

- **False positives.** Any WAF can mistake good traffic for bad. We strongly recommend you monitor your traffic via the Hosted Dashboard for a minimum of two weeks before blocking traffic. You don't want to start blocking traffic with configurations that are generating false positives.
- Custom application vulnerabilities. If attackers discover a vulnerability unique to your application or the technologies you use, and if your WAF configuration does not have a rule to protect against exploits for that particular vulnerability, it will not be able to protect your application in that instance.
- Inspection of HTTP and HTTPS traffic only. A WAF only inspects HTTP or HTTPS requests (layer 7). It will not process any TCP, UDP, or ICMP requests.
- Security products note. No security product, such as a WAF or DDoS mitigation product, including those security services offered by Fastly, will detect or prevent all possible attacks or threats. Subscribers should maintain appropriate security controls on all web applications and origins, and the use of Fastly's security products do not relieve subscribers of this obligation. Subscribers should test and validate the effectiveness of Fastly's security services to the extent possible prior to deploying these services in production, and continuously monitor their performance and adjust these services as appropriate to address changes in the Subscriber's web applications, origin services, and configurations of the other aspects of the Subscriber's Fastly services.

This article describes a product that may use third-party cloud infrastructure to process or store content or requests for content. For more information, see the section on <u>cloud infrastructure</u>, <u>data center</u>, <u>and physical security</u>.

- HIPAA-Compliant Caching and Delivery
- **iii** Last updated: 2018-08-01
- https://docs.fastly.com/products/hipaa-compliant-caching-and-delivery

You can configure the Fastly CDN service to cache and transmit protected health information (PHI) in keeping with Health Information Portability and Accountability Act (HIPAA) security requirements. Use the following features to ensure secure handling of cache data that contains PHI:

• Configure frontend and backend TLS to encrypt transmitted data from your origin to your end users.

• Add the beresp.hipaa variable to objects containing PHI to keep that data out of non-volatile disk storage at the edge.

Contact sales@fastly.com for more information on how to enable the beresp.hipaa feature for your account. For accounts that have this feature enabled, Fastly will enter into a HIPAA business associate agreement (BAA) as an addendum to our terms of service.

IMPORTANT

If you have purchased Fastly's PCI-compliant caching or HIPAA-compliant caching products Fastly will enforce a minimum version of TLS 1.2 or higher for all connections to meet the compliance requirements mandated by the PCI Security Standards Council.

1 NOTE

Fastly's security and technology compliance program includes safeguards for the entire Fastly CDN service, independent of using the beresp.hipaa variable. The Fastly security program and technology compliance content provide more information about these safeguards.

PCI-Compliant Caching and Delivery

Last updated: 2018-08-01 **=**

https://docs.fastly.com/products/pci-compliant-caching-and-delivery

We have designed Fastly's core CDN service with Payment Card Industry Data Security Standard (PCI DSS) compliance in mind. With proper authorization on your account, you can use Fastly's beresp.pci VCL variable to automatically cache content in a manner that satisfies PCI DSS requirements.

Adding the beresp.pci variable to an object prevents writing of that object to non-volatile disk storage on the edge. Combined with frontend and backend TLS, this feature allows you to cache and transmit flagged content through the Fastly network in compliance with our PCI certification.

Contact <u>sales-ecommerce@fastly.com</u> for more information on how to enable this product for your account.

IMPORTANT

If you have purchased Fastly's PCI-compliant caching or HIPAA-compliant caching products Fastly will enforce a minimum version of TLS 1.2 or higher for all connections to meet the compliance requirements mandated by the PCI Security Standards Council.



NOTE

Fastly's security and technology compliance program includes safeguards for the entire Fastly CDN Service, independent of using the beresp.pci variable. The Fastly security program and technology compliance content provide more information about these safeguards.

Platform TLS

Last updated: 2021-10-05

https://docs.fastly.com/products/platform-tls ത

Fastly's Platform TLS product allows you to programmatically manage certificates and keys for Transport Layer Security (TLS) using a web API. This product does not have a web interface component.

Consider this product if:

you need to support thousands of individual X.509 certificates and their associated private keys.

Fastly Products 6/30/22, 2:52 PM

• you own and generate your own certificates and private keys (typically obtained from a third-party certification authority such as Let's Encrypt).

For more information about this product, contact sales@fastly.com.



IMPORTANT

This information is part of a limited availability release. For additional details, read our product and feature lifecycle descriptions.

How Platform TLS works

Platform TLS allows you to programmatically manage certificates and private keys on a special Fastly service provisioned for use with the Platform TLS API. Using the API, you can:

- deploy new X.509 certificates
- retrieve information about deployed certificates
- update and delete existing certificates
- deploy new private keys
- retrieve information about private keys
- delete private keys

You can support your entire certificate lifecycle by replacing expiring certificates with newly generated ones at any time and using the API to rotate your private keys to manage your key management requirements.

Initial setup and configuration

The Platform TLS product will be provisioned by Fastly staff on a dedicated IP address pool (which you purchase separately) in Fastly's infrastructure. We configure your service to skip domain lookups and instead route client requests directly to your service based on the destination IP address that a client is connecting to. Because multiple certificates are served off the same IP address pool, Server Name Indication (SNI) is required for this product to work properly. We then provide you with a custom DNS map to use in your CNAME records and the corresponding Anycast IP addresses (for use with any apex domains you serve through Fastly).

Once setup is complete, certificates you upload using the API will automatically be made available to your dedicated IP address pool. Browser clients initiating a TLS handshake will automatically receive the proper certificate based on the domain indicated in the TLS handshake.

Certificate and key uploads and renewals

Once setup and configuration are complete, you can upload TLS private keys and matching TLS certificates using the Platform TLS API. The Platform TLS product automatically matches certificates to previously uploaded keys. TLS certificates may be procured from the certification authority (CA) of your choice.

When renewing and replacing certificates nearing expiration, you must procure new ones from your CA and then use the Platform TLS API to upload their replacements. You may also rotate your private keys. Any time you decide to swap out your key with a new one, that new key would need to be uploaded first, and then all the certificates associated with the old key would need to be regenerated and uploaded.

Domain configuration

To begin serving traffic through Fastly with the Platform TLS product, you or your customers must modify DNS records for any web properties to point traffic to the IP address pool assigned for your service. Fastly will assign a DNS name for use with your DNS records that can support a CNAME record and the Anycast IPs that can be used with apex domains.

- Using a CNAME record. With this option, a CNAME record gets created with a DNS provider and points to a custom DNS map Fastly provides. This option should be used for subdomains or wildcard domains (e.g., www.example.com or *.example.com).
- Using an A record. With this option, an A record gets created with a DNS provider and points to an Anycast address that Fastly provides. This option should be used for apex domains (e.g., example.com). Map names and Anycast addresses will be provided during initial setup and configuration. To obtain this information again, contact support@fastly.com.



IMPORTANT

For each of your domains, a CNAME or an A record must have been created with a DNS provider and you must have activated a Fastly service for traffic to be properly directed through it.

How TLS is enforced when you have multiple certificates

Fastly will automatically choose the certificate to be delivered for a given request based on the Host requested. The certificate with the most specific matching hostname will be preferred over certificates with less specific hostnames. Fastly's TLS server will always prefer an exact match SAN entry to a wildcard match. For example, on a request for api.example.com, Fastly will serve a certificate with a SAN entry for api.example.com over a different certificate with a SAN entry for *.example.com.

Conditions and limitations

When using Platform TLS, you agree to the following conditions:

- You are responsible for procuring your own certificates from the CA of your choice. Fastly will not procure certificates on your behalf.
- You are responsible for updating certificates prior to expiration. Expired certificates will cause TLS handshake failures that most browsers will display as site errors.

When using Platform TLS, you agree to the following limitations:

- This product requires a dedicated IP address pool on Fastly's infrastructure. If you've previously purchased a dedicated IP address pool from Fastly, Platform TLS may be enabled on it.
- The certificate deployment process is not instantaneous. It takes approximately 20 minutes on average to complete once a certificate is submitted, though the deployment may take as long as one hour.
- If two certificates are uploaded with identical hostnames, the most recently uploaded certificate will be chosen.
- By default certificates uploaded via the Platform API should not exceed one domain per certificate.

As with all API-based activities, standard API rate limits apply.

TLS オプション

Last updated: 2021-02-04

https://docs.fastly.com/products/tls-service-options-ja

Fastly は、Web サイトやアプリケーションがサービスのプラバシーとデーターセキュリティに必要となる HTTPS でトラフィックを配信 できるようにする様々な TLS (Transport Layer Security) サービスを提供しています。Fastly から安全な HTTPS トラフィックを配信する には、Web サイトに一致する秘密鍵を持つ有効な TLS 証明書が必要です。これらの証明書は、お客様が生成してアップロードすること も、Fastly 側で代わりに作成することも可能です。



★ ヒント

Fastly の pricing ページ は、当社の TLS サービスの現在の料金の詳細を記載しています。

重要な検討事項

Let's Encrypt やその他の認証局が提供する証明書は、サードパーティーの技術です。利用者は、自分が正当な登録者であることに責任を 持ち、代わりに手配された証明書に記載されているドメインを管理していることを証明する必要があります。

GlobalSign が提供する証明書は、https://www.globalsign.com/repository に記載された GlobalSign の利用規約に従うものとします。

Fastly の <u>PCI 準拠キャッシング</u>または <u>HIPAA 準拠キャッシング</u> を購入している場合、Fastly は PCI Security Standards Council が定め るコンプライアンス要件を満たすために、すべての接続に TLS 1.2 以上を適用します。

デフォルトでは、Server Name Indication (SNI) 拡張機能を使用します。最新のブラウザはすべて SNI をサポートしています。SNI をサポ ートしていないクライアント(Windows XP や Android 2.x 以前のクライアントなど)は、TLS ハンドシェイクのエラーが発生します。

Fastly は、RSA 公開鍵暗号化のために最小鍵長 2048 ビットの信頼された認証局によって署名された SHA-256 証明書をサポートしています。パフォーマンス上の理由から、それ以上の鍵長が必要ない場合には、2048 ビットの鍵長を使用することを強くお勧めします。

Fastly TLS

Fastly TLS では、 $\underline{\mathbf{6料P}\mathbf{7D}\mathbf{9D}\mathbf{1}}$ のお客様は、Web コントロールパネルまたは API を使用してドメインごとに TLS 証明書を管理することができます($\underline{\mathbf{開発者向けトライアルアカウント}}$ では、Fastly TLS を使用することはできません)。Fastly TLS では、お客様が TLS 証明書と秘密鍵を生成してアップロードする方法と、Fastly がサードパーティの認証局を介して TLS 証明書を代わりに生成して管理する方法があります。

設定・動作

<u>お客様所有の証明書を利用</u>する場合、Fastly の Web コントロールパネルまたは API を使用して TLS ファイルをアップロードすることが できます。証明書をアップロードする前に、対応する秘密鍵を事前にアップロードしておく必要があります。

<u>Fastly が管理する証明書を利用</u>する場合、証明書の発行要求を完了させるには DNS レコードを変更することで、ドメインを管理できることを証明する必要があります。Fastlyはドメインごとに1つの証明書を生成します。

デフォルトでは、Fastly は TLS 証明書を共有 IP アドレスにインストールします。クライアントのリクエストが Fastly に送信されると、TLS の SNI 拡張機能を使用して正しい証明書を選択し、クライアントが TLS ハンドシェイクリクエストでホスト名を提示できるようにします。

● 重要

Fastly TLS の証明書は50までに制限されています。この制限を上げる方法については、<u>sales@fastly.com</u> までお問い合わせください。

課金方法

Fastly TLS は、月末に TLS が有効になっている FQDN (fully qualified domain names) (例: example.com や www.example.com) とワイルドカードドメイン(例: *.example.com)の数に基づいて課金されます。

コンシェルジュTLS

エンタープライズサポートにコンシェルジュ TLS を加える場合には、sales@fastly.com までお問い合わせください。

その他のTLSオプション

Fastly TLS の他にも、共有証明書オプションや商用認証局から調達した証明書を使用するマネージドオプションなど、いくつかの TLS オプションをご用意しています。

Fastlyの共有ドメインを利用した無料 TLS

Fastly は無料の TLS オプションを提供しており、example.global.ssl.fastly.net のような Fastly の共有ドメインで HTTPS トラフィックを配信することができます。

このオプションを使用するには、ガイドの無料 TLS の設定の手順に従い、設定をはじめる前にの項目を確認してください。特定のトラフィックルーティング、ドメイン名や URL の要件がある場合は、Fastly の有料 TLS オプションを利用すると、より柔軟に対応できます。

Dedicated IPアドレス

Fastly は、お客様固有の DNS レコードで指定された <u>Dedicated IP アドレス</u> にお客様証明書をインストールすることができます。これらのDNSレコードは、3つのネットワークルーティングオプション(ネットワークマップまたはドメインマップと呼ばれることもあります)を使用するように設定することができ、Fastly ネットワークのどの部分を使用するかを選択することができます。

お客様がこのオプションの使用基準を満たしているかどうかを確認するには、sales@fastly.comまでお問い合わせください。



Last updated: 2022-04-14

https://docs.fastly.com/products/tls-service-options

Fastly provides a variety of Transport Layer Security (TLS) services that allow websites and applications to serve traffic over HTTPS, offering privacy and data security for services. To serve secure HTTPS traffic from Fastly, your website needs a valid TLS certificate with a matching private key. You can generate and upload these yourself or have Fastly do this automatically on your behalf.



TIP

Fastly's pricing page details the current rates for our TLS services.

Important considerations

Certificates provided by any certification authority (CA), whether they are non-profit or commercial providers, are third-party technologies. You are responsible for ensuring that you are the legitimate registrant and can demonstrate control of any domain that appears on a certificate procured on your behalf. Certificates provided by GlobalSign are subject to the terms of GlobalSign's Subscriber Agreement, which can be found at https://www.globalsign.com/repository.

For customers bringing their own certificates, both Fastly TLS and Concierge TLS service support Domain Validated (DV), Organization Validated (OV), and Extended Validation (EV) certificates. If Fastly manages your certificates, however, only DV and OV certificates can be used.

If you've purchased Fastly's <u>PCI-compliant caching</u> or <u>HIPAA-compliant caching</u> products, Fastly will enforce a minimum version of TLS 1.2 or higher for all connections to meet the compliance requirements mandated by the PCI Security Standards Council.

By default Fastly uses the Server Name Indication (SNI) extension. All modern browsers support SNI. Clients that do not support SNI (such as those on Windows XP and Android 2.x or earlier) will see a TLS handshake error.

Fastly supports SHA-256 certificates signed by publicly trusted certification authorities that have a minimum key size of 2048 bits for RSA public key encryption. For performance reasons, we strongly recommend using a 2048-bit key size for RSA when larger key sizes are not required for your application.

Fastly TLS

Fastly TLS allows <u>paid account</u> customers to manage TLS certificates on a domain-by-domain or multi-domain basis using our web interface or API (you can't use Fastly TLS with a <u>developer trial</u>). With Fastly TLS, you can either generate and upload your own TLS certificates and private keys or instruct Fastly to automatically generate and manage TLS certificates via a third-party, non-profit or commercial CA on your behalf.

How it works

If you <u>bring your own certificates</u>, you can use the Fastly web interface or API to upload TLS certificates and keys. You must ensure you upload the relevant private key first before uploading the matching certificate.

When Fastly manages your certificates, you use the Fastly web interface or API to select the CA from which Fastly should procure your TLS certificates. Fastly then procures DV certificates from the authority you've chosen. To complete a certificate request, you must prove that you control your domains by modifying DNS records.



TIP

To have Fastly procure organization validated certificates (OV) instead, contact sales@fastly.com.

By default, Fastly installs TLS certificates at a shared set of IP addresses. When client requests get sent to Fastly, we select the correct certificates using the SNI extension of TLS that allows clients to present a hostname in the TLS handshake request.

How we bill for it

Fastly TLS is billed based on the number of fully qualified domain names (e.g., example.com or www.example.com) and wildcard domains (e.g., *.example.com) that are TLS enabled at the end of the month for your account. All domains in an enabled state will be billed at the end of the month, regardless of certificate status (e.g., valid or expired).

Fastly TLS treats all entries on a certificate equally and each entry as its own item. On both certificates you manage and those that Fastly manages for you, an entry can be an apex domain, a subdomain, or a wildcard domain. Charges are based on the combined total of the domains on the certificates you manage as well as certificates that Fastly manages for you.

For Fastly-managed subscriptions, your charges may vary based on the CA you select. Specifically, there are pricing differences between Fastly TLS certificates provided by a commercial CA and those provided by a non-profit CA. Our <u>pricing page</u> provides specifics about these differences.

Concierge TLS

Concierge TLS provides you with TLS-specific advanced configuration support. It is sold as a packaged addition to Fastly's Enterprise Support service option.

To add Concierge TLS to your Enterprise Support option, contact sales@fastly.com.

Other TLS options

In addition to Fastly TLS, we make several other TLS options available including shared certificate options and a managed option that uses a procured certificate from a commercial certification authority.

Free TLS via the shared Fastly domain

Fastly offers a free TLS option that allows you to serve HTTPS traffic using an address like <code>example.global.ssl.fastly.net</code> via a shared Fastly domain.

To use this option, follow the instructions in our guide to <u>setting up free TLS</u> and pay close attention to the <u>noted limitations</u>. If you have specific traffic routing, domain naming, or URL requirements, one of Fastly's paid TLS options will provide you with more flexibility.

Dedicated IP addresses

Fastly can install customer-provided or Fastly-managed certificates at a <u>dedicated set of IP addresses</u> specified via customer-specific DNS records. These DNS records can be set up to use three possible network routing options (sometimes referred to as network maps or domain maps) that allow you to choose which parts of the Fastly network to use.

To see if your company meets the qualification criteria for this option, contact sales@fastly.com.

- **WAF Quick Start Package**
- **tast updated: 2021-11-17**
- https://docs.fastly.com/products/waf-quick-start-package

Fastly's WAF Quick Start Package provides your organization with basic setup and provisioning of your <u>WAF</u> by Fastly. For more information about this package, contact <u>support@fastly.com</u>.

IMPORTANT

As of June 30, 2021, the Fastly WAF (WAF 2020) offering became a legacy product. It will continue to be supported for all existing users. As an alternative, <u>Fastly Next-Gen WAF (powered by Signal Sciences)</u> offers proactive monitoring of and protection against suspicious and anomalous web traffic directed at your applications and origin servers. It can be controlled via the <u>web interface dashboard</u> or <u>application programming interface</u> (API). Contact <u>sales@fastly.com</u> or your Fastly account team to evaluate or move to the Fastly Next-Gen WAF option.

Prerequisites

To use the WAF Quick Start Package, you need to:

- <u>purchase a paid account</u> for Fastly's services
- purchase Fastly's WAF
- have logging enabled for at least one supported <u>logging endpoint</u>



1 NOTE

It's your responsibility to ensure you have a suitable logging service available and properly configured during the onboarding period.

How it works

Fastly Professional Services staff will guide you through the following stages:

- Planning. Professional Services staff will help you configure your WAF based on your security requirements. They will gather protection requirements, select from available WAF rules, and configure the WAF based on known or perceived risks to your application.
- Deployment. Professional Services staff help you configure your Fastly WAF VCL and add it to your Fastly service. They provide best-practice consulting for configuration of your WAF functionality within the Fastly service and will publish the policy to your Fastly service.
- Hand-off. Professional Services staff help you validate that your WAF policy is active and set up in logging only mode.
- WAF Support and SLA
- Last updated: 2019-12-12 ⅲ
- https://docs.fastly.com/products/waf-support-and-sla

IMPORTANT

As of June 30, 2021, the Fastly WAF (WAF 2020) offering became a legacy product. It will continue to be supported for all existing users. As an alternative, Fastly Next-Gen WAF (powered by Signal Sciences) offers proactive monitoring of and protection against suspicious and anomalous web traffic directed at your applications and origin servers. It can be controlled via the web interface dashboard or application programming interface (API). Contact sales@fastly.com or your Fastly account team to evaluate or move to the Fastly Next-Gen WAF option.

IMPORTANT

No security product, such as a WAF or DDoS mitigation product, including those security services offered by Fastly, will detect or prevent all possible attacks or threats. As a subscriber, you should maintain appropriate security controls on all web applications and origins. The use of Fastly's security products do not relieve you of this obligation. As a subscriber, you should test and validate the effectiveness of Fastly's security services to the extent possible prior to deploying these services in production, continuously monitor their performance, and adjust these services as appropriate to address changes in your web applications, origin services, and configurations of the other aspects of your Fastly services.

Fastly WAF Support

Fastly WAF Support offers the following resources to assist you with mitigating the service impacts of unwanted or malicious requests:

- Onboarding We will work with you to enable the initial setup and then do limited monitoring of the designated services for Fastly WAF.
- Initial configuration and deployment support We will actively work with you to select your rules to block Attacks.
- Ongoing Attack mitigation support We will work directly with you to configure and activate existing WAF rule filters to deal with changing Attacks or new Attacks.
- New standard rules We will assist you with the configuration of any new, standard rules introduced in the Fastly WAF.

Definitions

- "Business Hours" are 8AM-6PM during a Business Day in California, New York, and London.
- "Business Days" are Monday through Friday, excluding any day that is a US national or UK banking holiday.
- An "Attack" is a request or requests intended to cause unwanted or error responses from origin sites configured for any Fastly service. Fastly captures and analyzes suspected or actual Attack traffic to improve and protect its services.
- "Fastly Control" means elements entirely under Fastly's control and not a consequence of (a) a Subscriber's hardware or software failures, (b) a Subscriber's or end user's connectivity issues, (c) Subscriber operator errors, (d) Subscriber traffic amounts that exceed a Subscriber's Permitted Utilization as defined in the Terms and Conditions, (e) corrupted Subscriber content, (f) acts of god (any) or war, or earthquakes, or terrorist actions.

Support channels and availability

The following table summarizes support channels and availability for Fastly WAF Support as determined by the support package purchased by a Subscriber:

Support offering	Gold Support	Enterprise Support
Online self-service help	Unlimited access.	Unlimited access.
Availability for general inquiries	Business hours.	24/7/365.
Severity 1 incident report response	2 hours.	15 minutes.
Dedicated chat channel	Not available.	Business hours.
Web and email support	Available.	Available.
Phone support	Not available.	Toll-free telephone available 24/7/365.
Emergency escalation	Available via email.	Available via email and phone support.

Onboarding

As part of onboarding a subscriber service, Fastly support will:

- enable designated services for WAF functionality, providing access to our rule and filter libraries.
- work directly with you to determine the right set of rules and filters for your service.
- publish those rules or filters into your service in logging mode.
- monitor the behavior of those rules for a designated period starting when the rules are published to the service.

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

Subscriber responsibilities

Subscribers must identify and maintain two points of contact to be used during an Attack to communicate status and issues and to coordinate with Fastly to successfully protect services. Subscribers are responsible for using and configuring CDN services according to the documentation available at https://docs.fastly.com.

Support requests

Subscribers may make support requests by submitting a <u>support ticket</u>, which will trigger a system-generated acknowledgement within minutes containing the ticket number and a direct link to the ticket.

In particular, when requesting support related to an Attack, Subscribers should include as much of the following information as available:

- a determination of the severity of the Attack.
- the size of the Attack threatened or previously observed.

- the type and vector of Attack traffic seen or threatened.
- any duration of previous Attacks and vector behavior including major source IP addresses.
- an Attack history for the last 24 months.
- threat specifics including all details of any Attacks that the protected services or sites have experienced in the past.

Communications and channels of support

Support tickets

Create support tickets by sending an email to support@fastly.com. Tickets for communication between Fastly support engineers and a Subscriber's personnel are tracked using a ticketing application, which maintains a time-stamped transcript of communications, and sends emails to Subscriber and Fastly staff as tickets are updated.

Phone support

Subscribers who also purchase <u>Enterprise Support</u> receive a dedicated phone number to contact Fastly support engineers. Fastly personnel can also establish audio and video conferencing (free app or browser plug-in required) for real-time voice and video communications.

Chat

To facilitate real-time communication, Subscribers receive a dedicated chat channel during Onboarding and, for Subscribers that also purchase Enterprise Support, for an Attack for real-time communications about WAF issues during Business Hours or as needed by Fastly personnel. Though subject to change, Fastly's current chat provider is Slack (www.slack.com).

Observational logging

Fastly may from time to time, including as part of initial onboarding and during any period where Subscriber purchases additional Fastly WAF Tuning Package or Fastly WAF Tuning Plus Package, collect and store a copy of logging information from the Fastly WAF (which will include IP addresses) to monitor ruleset behavior, including false positives, by establishing a logging endpoint in your service configuration which will securely collect logging information in a third-party storage provider. Subscriber instructs Fastly to access and use the logs exclusively for providing WAF services, providing support and performance management to Subscriber, monitoring or maintaining Subscriber's Services and the Fastly WAF, threat detection and in accordance with the Documentation. Logged data will be deleted on a rolling basis and in any event retained no longer than thirty (30) days unless otherwise agreed by Subscriber.

- **WAF Tuning Package**
- **iii** Last updated: 2020-07-15
- https://docs.fastly.com/products/waf-tuning-package

IMPORTANT

As of June 30, 2021, the Fastly WAF (WAF 2020) offering became a legacy product. It will continue to be supported for all existing users. As an alternative, <u>Fastly Next-Gen WAF (powered by Signal Sciences)</u> offers proactive monitoring of and protection against suspicious and anomalous web traffic directed at your applications and origin servers. It can be controlled via the <u>web interface dashboard</u> or <u>application programming interface</u> (API). Contact <u>sales@fastly.com</u> or your Fastly account team to evaluate or move to the Fastly Next-Gen WAF option.

Fastly's WAF Tuning Package provides your organization with tuning of your <u>WAF</u> by Fastly. For more information about this package, contact <u>support@fastly.com</u>.

How it works

Fastly Professional Services staff will guide you through the following tuning stages:

• **Planning.** Professional Services staff help you gather protection requirements, define rules or filter policies, and develop policy structure.

- **Deployment.** Professional Services staff help you configure your Fastly WAF VCL and add it to your Fastly service. They provide best-practice consulting for configuration of your WAF functionality within the Fastly service and will publish the policy to your Fastly service.
- **Testing.** Professional Services staff help you validate that your WAF policy is active and set up testing for it.
- Go-Live. Professional Services staff monitor and address issues during final production testing and policy deployment.

Planning, deployment, testing, and go-live may involve some iterative cycles depending on the complexity of your policy.

Policy services

Some common tuning options we offer include:

- Initial setup and configuration
- Validation of policy match to origin systems
- Policy updates and maintenance
- **WAF Tuning Plus Package**
- **tast updated: 2020-07-15**
- https://docs.fastly.com/products/waf-tuning-plus-package

IMPORTANT

As of June 30, 2021, the Fastly WAF (WAF 2020) offering became a legacy product. It will continue to be supported for all existing users. As an alternative, <u>Fastly Next-Gen WAF (powered by Signal Sciences)</u> offers proactive monitoring of and protection against suspicious and anomalous web traffic directed at your applications and origin servers. It can be controlled via the <u>web interface dashboard</u> or <u>application programming interface</u> (API). Contact <u>sales@fastly.com</u> or your Fastly account team to evaluate or move to the Fastly Next-Gen WAF option.

Fastly's WAF Tuning Plus Package provides your organization with enhanced professional maintenance of your <u>WAF</u> by Fastly. The WAF Tuning Plus Package also improves visibility into application layer threats and strengthens your overall security posture. The WAF Tuning Plus Package includes ongoing tuning and configuration services designed to help protect you against critical threats. To protect against WAF bypass attacks, it also includes authenticated TLS to origin.

For more information about the WAF Tuning Plus Package, contact support@fastly.com.

WAF Tuning Plus Package features

Fastly's WAF Tuning Plus Package is a service offering for the term of your contractual agreement. It includes the following features.

Ongoing tuning and configuration

At your request, Fastly will provide you with one report per service protected by the Fastly WAF. Fastly will schedule periodic calls with you to review the reports.

Up to once per quarter, at your request, Fastly will tune previously provisioned WAF services as follows:

- We'll update your original profile, created during your initial WAF tuning, to record any new changes to your application stack or new perceived security risks based on actual or attempted attacks.
- We'll <u>update your WAF rule set</u> to the latest available (if applicable).
- We'll enable, disable, or change new or existing WAF rules based on new traffic patterns or security risks not present in the initial tuning cycle.
- We'll make a set of final recommendations on OWASP thresholds and switch your WAF into blocking mode.

Up to three times per quarter, at your request, Fastly will provide on-demand rule enablement (if available) for critical vulnerabilities.

Proactive notifications

We may notify you of available Fastly rules to help address critical vulnerabilities that we identify.

Authenticated TLS to origin

To mitigate WAF bypass attacks, Fastly will configure client-authenticated connections to your origin server for each service running WAF. This is an additional layer of security on top of network-level ACLs. This service requires a customer-provided TLS certificate, matching private key, and CA certificate or certificate chain.

Fastly will update the certificate on your behalf prior to expiration. Here's how it works:

- Fastly must receive new certificates at least 15 business days prior to expiration.
- Fastly will update the private key on your behalf (with a 15 business day notice) should the key be revoked.
- If you don't have your own key and certificate, Fastly can help you generate the certificates and keys at an additional cost. For
 more information, contact <u>sales@fastly.com</u>.

Observability

These articles provide information about Fastly products that provide visibility and insights into traffic, security, and performance.

- https://docs.fastly.com/products/product-observability
- **Domain Inspector**
- **tast updated: 2022-06-30**
- https://docs.fastly.com/products/domain-inspector

Domain Inspector provides you with a dataset and visualizations that offer real-time visibility and historical reporting of domain-level metrics delivered by a Fastly service. It allows you to monitor traffic for a <u>fully qualified domain name</u> without requiring you to send log data to a <u>third-party data collector</u>.

Specifically, Domain Inspector aggregates response data (for example, requests, bandwidth, edge or origin response codes, and cache hit ratios) received by a specified service and presents that information for each domain either as visualizations in the web interface or as a JSON formatted data stream.

Limitations and considerations

To help control data collection and aggregation, wildcard domains are aggregated against a single *.example.com entry.

Billing

Fees for Domain Inspector are charged based on the number of unique domains sending traffic through Fastly. To discuss pricing in more detail, contact sales@fastly.com.

- High Volume Logging
- **iii** Last updated: 2021-03-09
- https://docs.fastly.com/products/high-volume-logging

Fastly's <u>real-time log streaming</u> features allow you to tune the performance of Fastly services, but are limited to a monthly average of two log statements per request, per service. For customers who need to increase this limit, Fastly offers High Volume Logging. Contact <u>sales@fastly.com</u> to learn more.

Billing

High Volume Logging is billed based on the cumulative log statements streamed in excess of the Fastly's Full-Site Delivery Real-<u>Time Log Streaming</u> limits. Usage is calculated using the average size of all log statements multiplied by the number of statements in excess of the limit. The size is measured in log GBs streamed, pre-compression.

Logging Insights Package

t Last updated: 2018-06-15

https://docs.fastly.com/products/logging-insights-package

Fastly's Logging Insights Package provides you with guidance and customization of dashboard graphs in your third-party logging endpoint. After we've interviewed you to identify your specific business needs, we'll write advanced queries and create customized dashboards for the logs stored in your logging endpoint. You can then analyze and correlate any aspect of HTTP and HTTPS requests and responses to gain visibility into your service, allowing you to make decisions and changes. We'll then answer your questions and incorporate feedback to further customize the dashboards.

Prerequisites

To use the Logging Insights Package, you need to:

- purchase a paid account with a contract for Fastly's services
- have logging enabled for at least one supported <u>logging endpoint</u>
- be interviewed by Fastly so we can identify your customer-specific business needs
- grant Fastly temporary access to your third-party logging endpoint so we can configure your account on your behalf

NOTE

It's your responsibility to grant and revoke Fastly's access to your third-party logging endpoint.

Logging Insights Package features

The Logging Insights Package for Sumo Logic provides you with customization of the following Sumo Logic dashboards:

- The **Overview dashboard** provides you with a high-level overview of your Fastly services, allowing you to identify potential problems within them.
- The Origin Performance dashboard allows you to focus on your origin performance to check for latencies, slow URLs, and error-causing URLs.
- The Quality of Service dashboard allows you to see where your Fastly service's download times, cache performance, and performance by geographic location are below minimum thresholds.
- The Visitors dashboard allows you to see where your traffic is coming from.

The Logging Insights Package supports the <u>Sumo Logic App for Fastly</u>. You'll need a Sumo Logic account with the appropriate license, and you'll need to enable the <u>Sumo Logic logging endpoint</u>. For additional information, contact <u>sales@fastly.com</u>.

Origin Inspector

t Last updated: 2022-06-30

https://docs.fastly.com/products/origin-inspector

Origin Inspector provides you with a dataset and visualizations that offer real-time and historical visibility into responses delivered from your origin servers to the Fastly Edge Cloud. It allows you to monitor origin traffic without requiring you to send log data to a third-party data collector.

Specifically, Origin Inspector aggregates origin response data (for example, egress bandwidth, response status codes, and number of origin responses) received by a specified service and presents that information either as <u>visualizations in the web interface</u> or as a <u>JSON-formatted data stream</u>.

Prerequisites

Origin Inspector can be purchased as part of either Fastly's Essential or Professional Edge Cloud packages.

Customers with a <u>paid account with a contract</u> for Fastly's services can purchase Origin Inspector at the Enterprise level for a separate, additional fee rather than as part of an Edge Cloud package. Contact us at <u>sales@fastly.com</u> to discuss this option.

Data reporting and availability timeframes

Real-time data is reported in one-second periods up to the last complete second for a 120-second window. Historical data aggregation and retention timeframes, however, vary based on the Origin Inspector level you have purchased.

Edge Cloud Package	Aggregation Timeframe	Retention Timeframe
Essential	Per hour	8 days
Professional	Per minute	15 days
Enterprise	Per minute	45 days

Billing

Fees for Origin Inspector are charged based on the number of origins sending traffic through Fastly. Where applicable, an "origin day" is one unique origin per calendar day. To discuss pricing in more detail, contact sales@fastly.com.

Compute

These articles provide information about Fastly's serverless compute environment for building applications and executing at the edge.

- https://docs.fastly.com/products/product-compute
- <u>Compute@Edge</u>
- **iii** Last updated: 2022-05-10
- https://docs.fastly.com/products/compute-at-edge

The Compute@Edge platform helps you compile your custom code to WebAssembly and runs it at the Fastly edge using the WebAssembly System Interface for each compute request. Per-request isolation and lightweight sandboxing create an environment focused on performance and security.



IMPORTANT

This feature is not permitted for use in cryptocurrency mining.

Serverless isolation technology

Compute@Edge runs <u>WebAssembly</u> (Wasm). When a Compute request is received by Fastly, an instance is created and the serverless function is run, allowing developers to apply custom business logic on demand.

Global deployment

Deploying to a Compute@Edge service leverages Fastly's software-defined network and globally distributed points of presence. A single deploy action makes customer logic available across the Fastly network.

Available programming languages

By running Wasm on the Fastly network, Compute@Edge creates a serverless environment suitable for multiple programming languages. Fastly collaborates with the **Bytecode Alliance** and other open source communities to actively grow the number of supported languages. Support level per language varies. Resources per language are available on developer.fastly.com.

Logging endpoint compatibility

Compute@Edge supports sending user-specified logs to a variety of logging endpoints. These connections can be created and managed via manage.fastly.com and by using a supported language.

Continuous integration and deployment

Deployment to the Compute@Edge platform can be accomplished via the Fastly web interface, the Fastly API, and via Fastly's Terraform provider plugin (Limited Availability). The Fastly CLI (Limited Availability) also provides a local toolchain with features for creating, debugging, and deploying to Wasm services, including Log Tailing and Local Testing.



O NOTE

Some Compute@Edge features available through the Fastly CLI are disabled by default. To learn more about them, contact your account manager or email sales@fastly.com for details.

Log Tailing (Limited Availability)

<u>Log Tailing</u> allows you to stream custom log messages from your Compute@Edge application so you can respond quickly when debugging the application without setting up a third-party logging tool.

Local Testing Limited Availability

Local Testing allows you to run your work-in-progress applications locally on your laptop, server, or Cl system, so you can test your Compute@Edge applications without hosting them on public staging or production environments. Local environments support a subset of Compute@Edge features.

Billing

Fastly charges for Compute@Edge based on the total number of compute requests you make plus their compute duration and CPU time. Each compute request uses the Fastly delivery architecture and the associated Compute@Edge charges apply in addition to any already applicable delivery charges.

- Compute requests. Compute requests represent a count of the number of times you invoke a function instance. Each incoming request creates one function instance.
- Compute duration. Compute duration represents the total memory allocation over time required to process a compute request. We measure memory allocation in GB-seconds and calculate it based on the time it takes for a function instance to execute, multiplied by the memory allocated to that function. We measure function instance time in wall clock time from the start of a function to its completion or termination, rounded up to the nearest 50ms. The minimum function size for billing purposes is 128MB, though actual memory allocation may differ.
- **CPU time.** CPU time represents the sum of the total time spent by the CPU running your code.

For specific pricing information, contact your Fastly account team.

Resource limits

Fastly services and individual instances are allowed a specific range of resources per service and per execution as described in our developer learning resources. These limits are lower for free trial accounts.

Keep these limitations and constraints in mind especially when <u>testing and debugging</u> on Compute@Edge, when <u>sending Fastly</u> <u>logs</u> to third party logging providers, and when using <u>Log Tailing</u>.

Enable

These articles provide information about Fastly services and support solutions.

https://docs.fastly.com/products/product-enable

Assurance Services

tast updated: 2018-03-30

https://docs.fastly.com/products/assurance-services

Subscribers who purchase Assurance Services will:

- have access to a library of third-party audit reports and certification attestations (most recent 12 months).
- have access to executive summary reports for penetration tests and network scans (most recent 12 months).
- have access to a library of security-related policies and procedures.
- have access to a library of executive summaries of annual risk assessments (most recent 12 months).
- have access to a library of historical Fastly Service Advisory (FSA) documents (most recent 12 months).
- be able to perform unlimited audits of Fastly's <u>security</u> and <u>technology compliance</u> programs, subject to Subscriber's purchase of <u>Professional Services</u>. Audits require advance notice of at least 10 business days and shall be performed by Subscriber (or a mutually acceptable third party) according to standard audit practices.
- have the ability to be added as an Additional Insured on Fastly's General Commercial Liability Insurance for an additional fee.

Subscribers who wish to purchase Assurance Services must also purchase Gold or Enterprise Support.

- Designated Technical Specialist
- **t** Last updated: 2021-02-19
- https://docs.fastly.com/products/designated-technical-specialist

Fastly offers the ability to purchase the support of a Cloud Engineer to serve as a Designated Technical Specialist for your organization. These specialists help you optimize your use of Fastly's products and features by providing proactive check-ins and regular reviews to help you analyze your account's service configurations and their performance. Designated Technical Specialists also provide enhanced troubleshooting coordination with Fastly's support and professional services organizations.

Fastly offers three Designated Technical Specialist packages: Essentials, Premier, and Enterprise. A Designated Technical Specialist's available hours of service each month to your organization depend on the package you purchase and could include the following summarized activities:

Support Offering	Essentials	Premier	Enterprise
Total available hours	Up to 15 hours/month	Up to 80 hours/month	Up to 160 hours/month
VCL code configuration reviews	1 review/month	4 reviews/month	8 reviews/month
Architecture and configuration recommendations	Included	Included	Included
White glove escalation support	Included	Included	Included

Support Offering	Essentials	Premier	Enterprise
Email support	Available	Available	Available
Private chat support	Available	Available	Available
Emergency phone support	Available	Available	Available
Availability for general inquiries	Business hours	Business hours	Business hours
Initial response time	Next business day	Next business day	Next business day
Proactive account management	Account Manager	Account Manager	Account Manager
Support coordination	Included	Included	Included
Scheduled technical check-ins	Monthly	2x monthly	Weekly (as requested)
Account reporting	Monthly	Weekly	Weekly
Business reviews	Annually	Quarterly	Monthly
On-site travel for business reviews	Not included	Up to 2x annually (as requested)	Quarterly (as requested)
Custom reporting	Not included	Not included	By request
Early beta program access	Not included	Not included	Included

Definitions

- "Business Hours" are 8AM-6PM during a Business Day in California, New York, London, or Tokyo.
- "Business Days" are Monday through Friday, excluding any day that is simultaneously a US, UK, and Japanese national or banking holiday.



IMPORTANT

Designated Technical Specialists provide support during Fastly business hours to facilitate non-urgent discussions. They are not a 24×7 resource. Always rely on <u>normal support communications channels</u> for urgent issues and escalations.

Designated Technical Specialist packages

Each Designated Technical Specialist package includes **Enterprise Support** as well as the following core features:

- Email and private chat channel support during business hours between you and Fastly to facilitate guick guestions and answers for general inquiries and communication.
- Regular guidance on topics like configuration analysis, account performance, infrastructure and company plans, and roadmap requests.
- Advice on best practices for implementing and using Fastly with your infrastructure.
- Priority engagement and coordination with appropriate support resources as necessary during normal Fastly business hours.
- Comprehensive documentation of your implementation of Fastly's services and requirements to enable better support by Fastly teams.
- Basic reports on utilization and performance of Fastly services.

For all Designated Technical Specialist packages, keep in mind that other than regularly scheduled business reviews on site (as applicable for your package) additional reviews or on site travel must be agreed upon in advance via a statement of work. Also, any unused hours or other scheduled availability does not carry forward to future months. You won't be entitled to any refunds or credits for unused hours or other scheduled availability for any one month.



English is the primary language used by Designated Technical Specialists.

Essentials Designated Technical Specialist

In addition to the core features noted above, included hours could be used for:

- Monthly account reports with an overview of services, traffic, and tickets.
- An annual business review.

Premier Designated Technical Specialist

In addition to the core features noted above, included hours could be used for:

- Weekly account reports with an overview of services, traffic, and tickets.
- Scheduled check-ins via phone, 2x monthly.
- A quarterly business review (QBR), with onsite option, twice annually.

Enterprise Designated Technical Specialist

In addition to the core features noted above, included hours could be used for:

- Weekly account reports with an overview of services, traffic, and tickets.
- Weekly scheduled check-ins via phone (as requested).
- A monthly business review (QBR), with onsite option, quarterly as requested.
- Custom reporting upon reasonable request.
- Early access to beta programs.
- Fastly Next-Gen WAF (powered by Signal Sciences) professional services
- **t** Last updated: 2022-03-17
- https://docs.fastly.com/products/fastly-next-gen-waf-professional-services

IMPORTANT

This page was previously named "Signal Sciences professional services" and described services associated with the products previously offered as the Signal Sciences Cloud WAF and Next-Gen WAF, which are now referred to collectively as the <u>Fastly Next-Gen WAF (powered by Signal Sciences)</u>.

Fastly Next-Gen WAF (powered by Signal Sciences) (Next-Gen WAF) professional services provide your organization with training, implementation, and maintenance services for the Next-Gen WAF. Depending on the service offerings you select, our team will provide training and work with you to plan, test, deploy, and maintain a solution to protect your applications and origin servers. All Fastly Next-Gen WAF professional services are designed to be delivered remotely and can be purchased a la carte or in bundles.

A la carte services

The following services can be purchased a la carte.

Implementation services

Fastly will help you implement a solution using the Next-Gen WAF. Implementation services include personalized meetings to help you plan and deploy a customized solution for your applications and origin servers. Fastly will help you test your configuration prior to deployment. Implementation services are required for all new customers.

https://docs.fastly.com/products/aio

Implementation services provide a tried-and-tested process to help you choose the right approach for your environment and use case. Your implementation consultant will work with you to ensure you have the training and support required to unlock the potential of our security products.

Training services

Fastly can provide two different types of training sessions. A free introductory virtual training session available on Fastly Academy teaches skills for using Fastly Next-Gen WAF products and provides real-world examples. You can purchase advanced training sessions with an instructor to learn skills for using the Signal Sciences Application Programming Interface (API) and troubleshooting the Next-Gen WAF Agents and Modules.

Managed rules

Managed rules are rules created and managed by Fastly for your organization. Fastly will create and deploy managed rules for your organization after speaking with you about your organization's requirements. Managed rules are sold in packs of five.

Health checks

Fastly can perform a health check of your deployment of Next-Gen WAF to ensure that your deployment is in a "good" state and being fully utilized. At your request, our team will speak with you to understand how you currently use the Next-Gen WAF and then provide you with an assessment of your deployment with suggestions for improvement.

General services

General services is an hourly service offering (an eight hour minimum) that provides you with access to the Solutions Engineering team.

Essentials subscriptions

Essentials subscriptions are bundles of implementation services, training services, managed rules, and health checks.

Continuity Essentials

Continuity Essentials is an annual service offering that provides introductory training and an onboarding call with our team. At your request, Fastly will provide up to four quarterly reviews of your implementation and an assessment of your deployment.

TAM Essentials

TAM Essentials is an annual service offering that provides you with access to Shared TAMs. Your organization will receive up to two training sessions and an onboarding call with our team. At your request, Fastly will provide a quarterly review of your implementation and an assessment of your deployment. Our team will also provide a roadmap session and review all of your organization's open support cases, bugs, and feature requests.

Strategic Essentials

Strategic TAM Essentials is an annual service offering that provides access to a designated TAM with access to the TAM enterprise team for additional support as needed. Your organization will receive up to four training sessions, including self-paced and instructor-led sessions, an onboarding call with our team, and a monthly check-in call. At your request, Fastly will provide a quarterly review of your implementation and an assessment of your deployment. Our team will also provide a roadmap session and review all of your organization's open support cases, bugs, and feature requests.

- Live Event Monitoring Service
- **iii** Last updated: 2020-06-25
- <u>https://docs.fastly.com/products/live-event-monitoring-service</u>

With Fastly's Live Event Monitoring Service, our Customer Support engineers will monitor your scheduled event's performance and help troubleshoot issues with your Fastly service. We will also alert you as we detect issues with internet congestion and with upstream or downstream providers. We do this in real time throughout your event using a dedicated chat channel. This allows you

https://docs.fastly.com/products/aio

to receive alerts and notifications as well as ask questions without losing time spent contacting support and recounting what the issue is. Fastly's Live Event Monitoring Service is performed from Fastly's offices and does not include support on-site at your facilities.

For additional information about this service, contact sales@fastly.com.



IMPORTANT

This information is part of a limited availability release. For additional details, read our <u>product and feature lifecycle</u> descriptions.

Prerequisites

To use the Live Event Monitoring Service, you must purchase a paid account with a contract for Fastly's services.

You must schedule the start and end times of your event. These times will appear on your service order.

Event Monitoring service features

For the duration of your scheduled event, the Live Event Monitoring service reserves Fastly support staff who will perform the following:

Monitoring:

- Drops or spikes in bandwidth and request levels
- 5xx and 4xx errors
- Cache hit ratio
- Origin latency
- · Upstream issues with origin
- Internet congestion events

Alerting and real-time communication:

- Kick-off call to define alerting thresholds
- Real-time notifications via instant messaging

Troubleshooting:

- Rapid response from personnel who know your configuration and have been monitoring the scheduled event
- Accelerated escalation to senior support teams

Observational logging

In the course of performing Live Event Monitoring services, Fastly may collect and store a copy of logging information from Fastly Services by establishing a logging endpoint in your service configuration that will securely collect logging information in a third-party storage provider. When you purchase Live Event Monitoring services, you allow Fastly to access and use the logs exclusively for providing performance management, monitoring, and troubleshooting of your Fastly services during the event and for analysis after the event. Logged data will be deleted on a rolling basis and in any event retained no longer than thirty (30) days unless otherwise agreed to by you.

- Performance Optimization Package
- **iii** Last updated: 2020-04-01
- https://docs.fastly.com/products/performance-optimization-package

https://docs.fastly.com/products/aio

Fastly's Performance Optimization Package allows you to take advantage of configuration expertise to analyze and tune the performance of your Fastly services. Fastly's Professional Services team can help you use real-time analytics to identify potential improvements for your site's performance.

Prerequisites

To use the Performance Optimization Package, you need to:

- purchase a <u>paid account with a contract</u> for Fastly's services
- provide Fastly with a batch of representative site URLs that Fastly will analyze for various performance-related factors and use to suggest changes to increase performance

Performance Optimization Package features

The Fastly Performance Optimization Package specifically includes the following analyses and recommendations by Fastly Professional Services staff:

- Cache Hit Ratio, Shielding, and Clustering. We'll review your existing configuration and service settings and recommend incremental performance improvements you can make to ensure you're taking advantage of Fastly's network architecture.
- **Gzip and Brotli (origin based) compression.** We'll suggest configuration changes needed to ensure requested objects have the proper compression for each content type.
- HTTP/2 readiness. We'll assess your site and suggest network protocol changes to support HTTP/2, and provide recommendations on how to optimize for it.
- **TCP/IP protocols.** We'll analyze how your Fastly services send data via TCP/IP to end users and suggest the configuration changes needed to maximize request throughput while reducing last mile latency.

As part of this package, we'll provide you with a written assessment of our recommendations. Implementation of those recommendations by Fastly's <u>Professional Services team</u> can be purchased at an additional cost. For more information, contact <u>sales@fastly.com</u>.

Professional Services

tast updated: 2018-10-05

https://docs.fastly.com/products/professional-services

Fastly offers a range of Professional Services to help you begin using Fastly services. Choose between <u>Service Implementation</u>, <u>Service Management</u>, or <u>Consulting Engagement Services</u>, depending on your needs. For more information about any of our Professional Services packages, contact <u>sales@fastly.com</u>.

Service Implementation

How it works

Fastly Professional Services staff will personally guide you through the following stages:

- **Planning:** Professional Services staff help you with requirements gathering, solution design, documentation and resource allocation.
- **Implementation:** Professional Services staff help you with configuration of Fastly services and custom VCL development. They provide best-practice consulting for configuration of your origins.
- Testing: Professional Services staff help you validate configurations and set up testing.
- **Go-Live:** Professional Services staff monitor and address issues during final production testing and deployment.

Implementation, Testing, and Go-Live may involve some iterative cycles depending on the complexity of your configuration.

Implementation options

https://docs.fastly.com/products/aio 40/70

Some common implementation options we offer include:

- · Initial setup and configuration
- End-to-end encryption setup
- Fine-tuning cache times
- <u>Custom header</u> logic
- Dynamic content delivery optimization
- Multi-tiered caching setup
- Lightweight web page hosting
- Custom <u>purging</u> and event-driven content management
- Geographic or localization detection
- Edge logic and device detection
- <u>Stale content</u> configuration and origin outage handling
- Edge authentication and authorization
- ESI (<u>edge side includes</u>
- Streaming and video packaging
- Site performance analysis
- Managed vendor migration

Fastly offers two Service Implementation packages:

- Standard: Basic implementation for Fastly customers with simple content configurations.
- Enterprise: Advanced implementation for Fastly customers with complex, custom configurations.

Service Management

For customers who require ongoing configuration and technical assistance, Fastly offers Service Management that provide professional services to you and your staff on an as-needed basis. These hours may be used to supplement your existing Support Plan or Service Implementation.

Some common activities you may need assistance with:

- Site performance analysis
- Varnish and VCL training
- <u>Service</u> configuration
- End-to-end encryption setup
- Cache time fine-tuning
- Custom header logic creation
- Dynamic content delivery optimization
- Multi-tiered caching setup
- Lightweight web page hosting
- Custom purging and event-driven content management
- Geographic or localization detection
- Edge logic and device detection
- Stale content configuration and origin outage handling

https://docs.fastly.com/products/aio 41/70

- · Edge authentication
- ESI (edge side includes) configuration
- Streaming and video packaging

Consulting Engagement Services

For customers who require in-house expertise or dedicated resources, Fastly's Support Engineers are available to provide a range of more technical professional services, including:

- Technical advisory services
- Translating configurations to VCL
- Optimization of website performance
- On-site Varnish and VCL training
- Non-Fastly related performance tuning
- · Adapting Fastly features to a particular customer use case
- Related offerings
- **iii** Last updated: 2018-04-24
- https://docs.fastly.com/products/related-offerings

Fastly offers service level agreements to customers based on the nature of their agreement with Fastly and the Fastly products they have purchased. These service level agreements offered by Fastly provide information to customers based on the nature of their agreement with Fastly and the Fastly products they have purchased.

- Service availability SLA
- Support description and SLA

We understand that some customers may require more support from Fastly to meet their additional security and compliance needs. Customers with these needs may subscribe to packages that include our <u>Assurance Services</u> offering.

- Response Security Service
- **iii** Last updated: 2022-02-03
- https://docs.fastly.com/products/response-security-service

Fastly offers <u>Fastly Next-Gen WAF (powered by Signal Sciences)</u> customers a Response Security Service (RSS) that provides your organization with enhanced access to our Customer Security Operations Center (CSOC) team and periodic consultation with a Designated Security Specialist for strategic security solutions reviews and planning. Together, Fastly's CSOC team and your Designated Security Specialist support the design, implementation, and maintenance of your security solutions by assisting with initial configuration, requested maintenance, and attack support.

The Fastly Next-Gen WAF now collectively refers to the products that were previously known as the Signal Sciences Cloud WAF and Signal Sciences Next-Gen WAF. The functionality of those products has not changed as part of the new naming convention. Fastly Next-Gen WAF continues to be powered by Signal Sciences technology.

The following table summarizes what RSS provides:

Support offering	Details
CSOC availability for general requests and inquiries	24/7/365.

https://docs.fastly.com/products/aio 42/70

Support offering	Details
CSOC availability for critical security incidents	24/7/365.
General inquiries response times	Within 24 hours.
Critical security incident response times	Within 15 minutes of notice with active troubleshooting starting within 30 minutes of acknowledgement of incident severity.
Emergency phone number	Included.
Emergency email address	Included.
Dedicated chat channel	Included.
Online self-service help	Unlimited access.
Virtual, self-paced training	Included.
Quarterly reviews	Upon request, once per quarter, during US business hours.

Definitions

- "Business hours" are 8AM-6PM during a business day in California or New York.
- "Business days" are Monday through Friday, excluding any day that is a US national holiday.
- A "critical security incident" is an event that creates significant business impact or loss of availability for your production environments, or that threatens the integrity or confidentiality of your proprietary information.
- "Fastly control" means elements entirely under Fastly's control and within its products' capabilities, and not a consequence of (a) your hardware or software failures or those of a third party, (b) you or your end user's connectivity issues, (c) operator errors by your employees, (d) alteration, modification, unauthorized or misuse of Fastly products, or use not in accordance with the documentation, (e) your corrupted content, and (f) acts of god (any) or war, or earthquakes, or terrorist actions.

Prerequisites

To purchase and use RSS, you must also purchase a Premier Platform subscription for Fastly Next-Gen WAF.

To ensure accurate response to requests and incident reports, you must ensure your account contact information remains up-to-date. CSOC can help you verify this information at any time.

Response Security Service features

Fastly offers RSS for the term of your contractual agreement. It includes the following features.

Priority CSOC access

By purchasing RSS, you will be entitled to 24/7 access to Fastly's CSOC for assistance with incidents, configuration changes, and general inquiries. To receive this assistance, you may initiate contact via:

- **Phone number.** You will receive a dedicated, toll-free phone number to initiate contact with Fastly's CSOC and to report critical security incidents. Fastly personnel can also establish audio and video conferencing (free app or browser plug-in required) for real-time voice and video communications.
- **Email address.** You will receive an email address to initiate contact with Fastly's CSOC for general support questions as well as an emergency email address for reporting of critical security incidents.
- **Chat channel.** You will receive a dedicated security chat channel for real-time communications to discuss general security product support and questions during business hours or as needed by Fastly personnel. Though subject to change, Fastly's current chat provider is Slack (www.slack.com).

Each of these contact methods will be provided to you during your onboarding period.

https://docs.fastly.com/products/aio 43/70

Online help and virtual training

In addition to unlimited access to online self-service documentation at <u>docs.fastly.com</u> and <u>developer.fastly.com</u>, you will have access to virtual, self-paced Fastly Next-Gen WAF application training scenarios.

Quarterly reviews

At your request, Fastly will schedule 2-hour account check-ins during US business hours, up to one per quarter during the term of your subscription, with a Designated Security Specialist who will help you review configurations, consult with you on rule creation, review security product roadmaps with you, and discuss your overall security health. Because some review discussions require advance preparation, you must schedule them at least two weeks in advance by making a request via the provided RSS general support question email address. You won't be entitled to any refunds or credits for unused scheduled availability.

RSS support requests and response times

Fastly's response times and status updates vary based on request and incident severity.

General requests and inquiries

You may initiate general requests and inquiries by creating a ticket via the general support email address provided to RSS customers or by submitting a ticket via the <u>Signal Sciences console</u> and we will acknowledge your general outreach within two hours of its receipt. We will begin addressing your ticket within 24 hours of acknowledging its receipt and will provide status updates to you once daily on each subsequent day until the incident is resolved or is believed to be outside of Fastly's control.

Critical security incidents

Support for critical security incidents can only be initiated via the emergency email address provided to RSS customers (not chat) or by selecting the Urgent priority when submitting a ticket via the Signal Sciences console. The ultimate classification of a request submitted by either of these methods will be determined by Fastly based on various factors including input you provide.

Fastly will acknowledge your critical security incident outreach within 15 minutes of its receipt. If classified as a critical security incident, we will begin actively troubleshooting these incidents within 30 minutes of acknowledging your ticket and will provide an initial status update within an hour of acknowledging your ticket, with subsequent updates at least every 4 hours thereafter unless an alternative update cadence has been agreed upon. Fastly will continue to work until the incident impact has been mitigated or is believed to be outside of Fastly's control.

RSS response SLA and credit terms

If you have purchased RSS and, during a critical security incident, there is a material delay in response time and the cause of the delay is within Fastly's control, a one-time credit of \$500 per incident will be credited to your account. Specifically:

- Requests for invoice credits must be made within 30 days of the critical security incident that triggered the service credit.
- All requests for invoice credits must be made to <u>billing@fastly.com</u>.
- In no event shall invoice credits exceed the fee for RSS payable by you for the month in which the invoice credits accrued.
- A pending invoice credit does not release you from your obligation to pay Fastly's submitted invoices in full when due.
- Invoice credits will be applied to the invoice generated two months following the month in which the credits were incurred.

If in any three-month period where three (3) or more support response time objectives are not met and the failure to meet the objectives materially adversely impacted you, you will have 30 days to terminate the RSS subscription following the third response failure. You must notify Fastly of your intention to terminate the RSS subscription within 30 days of the triggering event.

Limitations

RSS has the following limitations:

- Product applicability. This service only applies to the <u>Fastly Next-Gen WAF</u> product. No other products are included in this service.
- **Service monitoring.** This is a reactive service, not a pro-active one. You must initiate all requests for action. Fastly does not monitor your services for security events or suspected attacks.

https://docs.fastly.com/products/aio 44/70

• Origin administration and access. Fastly will not directly access or administer your origin systems at any time.

- Third-party product administration. Fastly will not administer third-party products or services.
- Identity Verification. For contacts via telephone, we encourage you to establish authentication methods to verify that individuals reporting issues via telephone are authorized to make inquiries or request changes to account configurations on your behalf. Authentication methods may include use of an account authorization passphrase, Slack challenge process, or email verification. If an individual reporting an issue via telephone is not able to have their identity verified, they may report issues but not receive any account information or initiate account changes and your account's administrators will be notified of requests or inquiries.

No security product, such as a WAF or DDoS mitigation product, including those security services offered by Fastly, will detect or prevent all possible attacks or threats. You should maintain appropriate security controls on all web applications and origins, and the use of Fastly's security products do not relieve you of this obligation. You should test and validate the effectiveness of Fastly's security services to the extent possible prior to deploying them in production, and you should continuously monitor their performance and adjust the services as appropriate to address changes in your web applications, origin services, and configurations of the other aspects of your Fastly services.

Service availability SLA

tast updated: 2019-06-07

https://docs.fastly.com/products/service-availability-sla

Support availability and response times vary depending on the <u>type of account</u> you have and the <u>level of support</u> you have purchased.

Agreement Type	Unpaid Account	Month-to-Month Account	Termed Contract	Gold & Enterprise Support
Service Level Agreement	None	None	Termination Option	Invoice Credits + Termination Option

Definitions

"Degraded Performance" means the Services are experiencing Error Conditions that are (1) caused by issues under Fastly Control, (2) observable or reproducible by Subscriber or Fastly, (3) requiring Subscriber to redirect traffic off the Services. Degraded Performance does not include any reduction on availability of the Application Web Interface or API due to maintenance.

"Error Condition" means the Services are (1) not responding to end user requests, (2) incorrectly sending end users error condition messages or (3) sending incorrect partial content to end users and these conditions are observable or reproducible by Subscriber or Fastly.

"Fastly Control" means elements entirely under Fastly's control and not a consequence of (a) Subscriber hardware or software failures, (b) Subscriber or end user connectivity issues, (c) Subscriber operator errors, (d) a Utilization spike (see below), (e) corrupted Subscriber content, or (f) acts of god (any) or war, or earthquakes, or terrorist actions.

Termination

Any Subscriber that has a contract with a term and a minimum commitment shall have thirty (30) days to terminate their subscription agreement following (1) a period of Degraded Performance longer than 7.2 hours in any one month, or (2) three contiguous months that have periods of Degraded performance longer than 43.8 minutes each.

Availability invoice credits

Subscribers who purchase Gold or Enterprise Support shall be entitled to Invoice Credits according to the following table.

Availability Percent	Period of Degraded Performance	Monthly Credit Percent
Below 100% - 99.99%	Up to 4.32 minutes	1%
Below 99.99% – 99.9%	Up to 43.8 minutes	5%

https://docs.fastly.com/products/aio 45/70

Availability Percent	Period of Degraded Performance	Monthly Credit Percent
Below 99.9% – 99.0%	Up to 7.2 hours	10%
Below 99.0% - 98.0%	Up to 14.4 hours	25%
Below 98.0%	Greater than 864 minutes	50%

Invoice Credits for unavailability will accrue on a monthly basis. The Credit Amount for a month is equal to the monthly usage charge multiplied by Monthly Credit Percent.

Credit terms

- Requests for Invoice Credits for Degraded Performance must be made within 30 days of the period of Degraded Performance.
- The maximum amount of any credit is the Invoice Amount for the month the Degraded Performance occurred.
- A pending credit does not release Subscriber from its obligation to pay Fastly's submitted invoices in full when due.
- Credits will be applied to the Invoice two months following the month an invoice credit was incurred.

Utilization Spikes

Subscriber's bandwidth utilization, measured in megabits per second, will be sampled every five (5) minutes on a region-by-region basis each month (the "Samples"). Subscriber's "Average Utilization" for a region in a month will be the average of the Samples. Subscriber's "Peak Utilization" for a region in a month will be calculated by the 95th percentile method, according to which the Samples will then be ordered from highest to lowest, and the highest five percent (5%) of Samples will be discarded and the remaining highest Sample will be Subscriber's Peak Utilization for the region in that month. Subscriber's "Permitted Utilization" in a month for a region will be five (5) times Subscriber's Average Utilization in that month for that region. A "Utilization Spike" will occur if Subscriber's Peak Utilization exceeds its Permitted Utilization in a region. Utilization Spikes may interfere with or disrupt the integrity or performance of the Services. Subscribers should contact Support in advance of any planned utilization spike and respond immediately to any communications from Fastly regarding an actual or suspected Utilization Spike.

Support description and SLA

 Last updated: 2022-02-02

 https://docs.fastly.com/products/support-description-and-sla

Support availability and response times vary depending on the <u>type of account</u> you have and the level of support you have purchased. The following table summarizes those offerings:

Support Offering	Standard Support	Gold Support	Enterprise Support
Online Self- Service Help	Unlimited access.	Unlimited access.	Unlimited access.
Availability for General Inquiries	Business hours.	Business hours.	24/7/365.
Availability for Incident Reports	Business hours, including weekends & holidays.	24/7/365.	24/7/365.
Initial Response Times	By the next business day.	Severity 1 Incidents within 2 hours. Severity 2 Incidents within same day. All other Incidents by the next business day.	Severity 1 Incidents within 15 minutes. Severity 2 Incidents within 2 hours. All other Incidents by the next business day.
Email support	Available.	Available, with priority over Standard Support.	Available, with priority over Standard and Gold Support.

https://docs.fastly.com/products/aio 46/70

22, 2:52 PM		Fastly Products	
Support Offering	Standard Support	Gold Support	Enterprise Support
Phone and chat support	Not available.	Not available.	Toll-free telephone available 24/7/365. Dedicated chat channel available during Fastly business hours.
Emergency Escalation	Not available.	Not available.	Available via email and phone.
Designated Technical Resource	Not available.	Not available.	Available with <u>designated technical</u> <u>resource add-on</u> package.
Discounted Professional Services	Not available.	Not available.	30% discount on <u>Service Management</u> packages. Does not apply to Fastly Next-Gen WAF (powered by Signal Sciences) service packages.
PCI and HIPAA configuration services	Not available.	Not available.	Available via email, phone, and chat support.
Enhanced compliance support (including GDPR)	Not available.	Not available.	Available via email, phone, and chat support.
Compute@Edge code configuration support	Interoperability and configuration guidance and troubleshooting.	Interoperability and configuration guidance and troubleshooting.	Interoperability and configuration guidance and troubleshooting.
Termination Option	Not available for unpaid and month-to-month customers. Only included for termed contracts.	Available with invoice credits.	Available with invoice credits.

Partner Support Services

For Fastly customers approved as Partners, additional partner support products become available. To be eligible as a Partner, customers must be classified and approved as such. Contact <u>partners@fastly.com</u> for details.

Partners will not be entitled to Standard Support that customers receive automatically on the platform. All Partners will be required to purchase either Partner Gold or Partner Enterprise support. The corresponding support availability and response times vary depending on the purchased support level.

In addition to the Gold and Enterprise support offerings, all Partners purchasing Partner Support receive access to a library of ondemand online training modules.

Technical support

The following section applies to all subscribers.

Definitions

- "Business Hours" are 8AM-6PM during a Business Day in California, New York, London, or Tokyo.
- "Business Days" are Monday through Friday, excluding any day that is simultaneously a US, UK, and Japanese national or banking holiday.
- An "Incident" is an occurrence during which end users' use of Subscriber's services is adversely impacted.

https://docs.fastly.com/products/aio 47/70

 A "Severity 1 Incident" is an incident resulting in a major service outage requiring Subscriber to redirect all traffic from Fastly to another CDN.

- A "Severity 2 Incident" is an incident resulting in minor or intermittent outage not requiring Subscriber to redirect traffic to another CDN.
- "Fastly Control" means elements entirely under Fastly's control and not a consequence of (a) a Subscriber's hardware or software failures, (b) a Subscriber's or end user's connectivity issues, (c) Subscriber operator errors, (d) Subscriber traffic amounts that exceed a Subscriber's Permitted Utilization as defined in the Terms and Conditions, (e) corrupted Subscriber content, (f) acts of god (any) or war, or earthquakes, or terrorist actions.

Subscriber responsibilities

Subscriber is responsible for using and configuring services according to the Documentation available at https://docs.fastly.com.

Support requests

Subscribers submit support requests by sending email to support@fastly.com. Subscribers receive a system-generated response within minutes containing the ticket number and a direct link to the ticket.

Reasons to contact us for incidents include:

- Services are not responding to end user requests.
- Services incorrectly send end users error condition messages.
- Services send incorrect or partial content to end users.

Incident reports should include all relevant information such as:

- Subscriber's determination of the Severity Level of the incident,
- · Subscriber hardware failures,
- Subscriber operator errors,
- · Services configuration errors made by Subscriber employees,
- A potential Utilization Spike (see the <u>Service Availability SLA</u>),
- · Corrupted Subscriber content,
- DDOS attacks, and
- Relevant force majeure acts such as extreme weather, earthquakes, strikes or terrorist actions.

Communications

Tickets

Communications between Fastly support engineers and Subscriber personnel are conducted using a ticketing application that maintains a time-stamped transcript of communications and sends emails to Subscriber and Fastly staff as tickets are updated.

Chat

Subscribers to Enterprise Support receive a dedicated chat channel for real-time communications during Business Hours. Though subject to change, Fastly's current chat provider is Slack (www.slack.com).

Phone support

Subscribers to Enterprise Support receive a dedicated, toll-free phone number to contact Fastly support engineers. Fastly personnel can also establish audio and video conferencing (free app or browser plug-in required) for real-time voice and video communications.

Response time

https://docs.fastly.com/products/aio 48/70

Fastly shall use best efforts to respond in a timely fashion.

Termed contracts

The following applies to any subscriber that has a contract with a term and a minimum commitment.

Response times

Fastly commits to acknowledging receipt of a support ticket within the next Business Day following submission of a support request by a Subscriber with a Termed Contract.

Termination

In any three-month period where three (3) or more support Response Time objectives are not met and the failure to meet the objectives materially adversely impacted Subscriber, Subscribers with a Termed Contract, Gold Support, or Enterprise Support shall have thirty (30) days to terminate their subscription agreement following the third failure.

Incident response times

Incident reporting

Severity 1 Incidents: Fastly will provide Subscriber an Incident Support Email address for Subscriber to report Incidents. Subscriber should report Incidents promptly using the Incident Support email.

Severity 2 Incidents: Subscriber should report Severity 2 Incidents by submitting a Support Request.

Incident reporting and additional fees

For Severity 1 Incidents caused by factors within Subscriber's control, a flat fee of \$1500 will be assessed, and any time spent beyond three (3) hours will be invoiced at Subscriber's undiscounted Professional Services rates. For Severity 2 Incidents caused by factors within Subscriber's control, Subscriber will be invoiced at Subscriber's undiscounted Professional Services Rates.

For all incidents:

- If the Incident-causing factors are within Fastly's control, there will be no hourly charges for Fastly engineering staff time.
- If the factors are within Subscriber's control, Subscriber agrees to pay Fastly its hourly charges for Fastly engineering staff time. If it appears likely the factors are within Subscriber's control, Subscriber may tell Fastly staff to stop working on troubleshooting the Incident (thereby stopping the hourly charges from being incurred). Subscriber agrees to tell Fastly to stop working on an Incident via an email sent to Fastly's Incident Support email address. The timestamp on the email will be the time charges cease to be incurred.

Gold Support

Fastly will respond to the report of an Incident by troubleshooting the causes of the Incident and resolve them if caused by factors within Fastly's control, or provide information to those who can resolve the factors if the factors are within others' control, as follows:

For a Severity 1 Incident:

- Fastly support staff will acknowledge receipt of the email within two hours.
- Fastly engineers will begin actively troubleshooting within the same day, will work on the Incident during the same day, and will provide status updates to Subscriber daily on each subsequent day.
- Fastly staff will work until (a) the incident is resolved or (b) the incident is believed to be outside of Fastly's control.

For a Severity 2 Incident:

- Fastly support staff will acknowledge receipt of the email within the same day.
- Fastly engineers will begin actively troubleshooting within the same day, will work on the Incident during the same day or next day, and will provide status updates to Subscriber daily on each subsequent day.

https://docs.fastly.com/products/aio 49/70

Enterprise Support

Fastly will respond to the report of an Incident by troubleshooting the cause(s) of the Incident and resolve them if caused by factors within Fastly's control, or provide information to those who can resolve the factors if the factors are within others' control, as follows.

For a Severity 1 Incident:

- Fastly support staff will acknowledge receipt of the email within 15 minutes.
- Fastly will start actively troubleshooting within 30 minutes of receipt of the email.
- Fastly will perform its tasks on a 24/7 basis.
- Fastly and Subscriber will immediately communicate upon learning new information that may be useful in troubleshooting the incident, and status updates between Fastly and Subscriber staff will take place no less frequently than every 30 minutes for the first two hours, and no less frequently than every hour thereafter.
- Fastly staff will work until (a) the incident is resolved or (b) the incident is believed to be outside of Fastly's control.

For a Severity 2 Incident:

- Fastly support staff will acknowledge receipt of the email within two hours.
- Fastly engineers will begin actively troubleshooting within the same day, will work on the Incident during the same day, and will provide status updates to Subscriber daily on each subsequent day.

Support invoice credits

In the event a Severity 1 Incident occurs, Subscriber has purchased Gold or Enterprise Support, the cause of the Incident is within Fastly's control, and any of the communication or response timeframes are materially not met, a one-time credit of \$500 per incident will be credited to Subscriber's account.

Credit Terms:

- Requests for Invoice Credits must be made within 30 days of the incident which triggered the service credit.
- In no event shall Invoice Credits exceed the invoice value of the month in which they are accrued.
- A pending credit does not release Subscriber from its obligation to pay Fastly's submitted invoices in full when due.
- Credits will be applied to the invoice two months following the month an invoice credit was incurred.

NOTE

Fastly maintains support for its original <u>Premium Support</u> and <u>Platinum Support</u> plans. To convert your account to the current Gold and Enterprise Support plans, contact <u>sales@fastly.com</u>. If you have an agreement that requires the purchase of Platinum support, converting to Enterprise support satisfies that requirement.

Resources

These articles provide information about Fastly's product lifecycle and the products within it.

- https://docs.fastly.com/products/product-resources
- Fastly product lifecycle
- https://docs.fastly.com/products/fastly-product-lifecycle

Fastly releases or retires its products and features as detailed below.

https://docs.fastly.com/products/aio 50/70

Product and feature releases

We release our products and features in the following stages.

Beta

Beta products are initial releases of potential future products or features. We provide customers who participate in our Beta program the opportunity to test, validate, and provide feedback on future functionality. Feedback gathered during this phase helps us to determine which features and functionality provide the most value to our customers and helps us focus our efforts accordingly.

These guidelines apply to Fastly's Beta program:

- Customers can choose or elect to participate in a Beta program.
- Fastly does not make any promises on the features, functionality, or performance of our Beta products.
- We reserve the right to change the scope of or discontinue a Beta product or feature at any point in time.
- We do not charge our subscribers for using our Beta products or features.
- Beta products or features are not included in any existing support contracts or obligations.
- Fastly does not provide Beta customers with discounts on future purchases of any products or services.
- Fastly strongly advises against using production traffic for Beta products due to their dynamic nature.

Beta services are covered by Section 6 of our Terms of Service.

Limited Availability

Limited Availability products are ready to be released to the world, pending some fine tuning. Limited Availability allows us to test out a product or service with a limited number of customers, so we can closely monitor it and make any necessary adjustments before rolling it out more broadly. Our goal is to make it easy for customers to set up our products with their services and take advantage of the features that come along with them.

These guidelines apply to Fastly's Limited Availability program:

- Fastly may charge its Limited Availability customers and pricing may vary depending on features.
- Fastly does not make any promises on the features, functionality, or performance of our Limited Availability products.
- Fastly does not provide its Limited Availability customers with discounts on future purchases of any products or services.
- Fastly does provide limited product and customer engineering support and documentation for Limited Availability products.

General Availability

General Availability products released by Fastly are available for everyone's use. Fastly manages these products in accordance with <u>Fastly's terms and conditions</u>.

Product or feature retirement

The decision to retire or deprecate Limited Availability or General Availability features always follows a rigorous process including understanding the demand, use, impact of feature retirement and, most importantly, customers' feedback. Our goal is to always invest resources in areas that will add the most value for customers. When low-value functionality or less successful features compete for resources or create confusion, we may decide that retirement or deprecation is the best solution. In the most difficult of scenarios, feature retirement may cause temporary challenges for some customers. Focusing on the highest priorities of the greatest number of customers, however, allows us to continue to deliver a superior solution with the most benefit.

Fastly is committed to transparency in everything we do, particularly when that activity has implications on the functionality of our features or platform. In the interest of building trust and clarifying change, we have established a number of guidelines around communication of feature retirement, end-of-life, and deprecation.

When a decision to retire a feature is reached Fastly will strive to provide:

• **Advance notice:** We will provide notification proportional to feature criticality. For minor changes with improved functionality we will notify customers no less than three (3) months prior to deprecation, and for major changes we will notify no less than

https://docs.fastly.com/products/aio 51/70

six (6) months in advance.

• Alternative functionality: We will include guidance and direction on new features in our services which replace retired or deprecated functionality. New features and functionality will always be provided in advance to ensure customers have time to understand and transition to new functionality prior to the retirement of previous functionality. In some cases this may be with partners or other approved third-party services.

- Continuous support: Fastly commits to providing continuous support for all features until the retirement date.
- Considerate scheduling: When planning significant changes, including feature or product retirement, we will align retirement as close to major updates or releases as possible to limit the scope of impact on your services.

In some extreme cases Fastly may need to accelerate the retirement of functionality timeline:

- Essential changes that are necessary or appropriate to protect the integrity of our service may occasionally be required. In these cases, it is important that those changes occur as quickly as possible. We will communicate with customers transparently with as much advance notice as possible in these situations.
- <u>Integrated third-party software or services</u> may need to be retired due to the third-party decision to change or retire their solution. In these situations, the pace of the retirement will be out of our control, although we remain committed to transparency and will strive to provide as much notice as possible.

For more information, contact customer support or your account team.



Fastly defines each of its products as follows. For more information about any of our products, contact sales@fastly.com.

Application Programming Interface

Fastly provides an <u>application programming interface (API)</u> that can be accessed via a number of popular <u>interactive clients</u> and allows you to manage Fastly services via remote procedure calls. These services include features such as <u>authentication</u>, <u>configuration</u>, <u>historical stats</u>, <u>purging</u>, and <u>remote logging</u>. In addition to being accessible via Fastly's API, Fastly services can also be accessed via a web interface for users with the appropriate <u>access permissions</u>; however, API features do not include customer account setup, which can only occur through the <u>web interface controls</u>.

Assurance Services

<u>Assurance Services</u> offers access to third-party audit reports, certification attestations, and unlimited audits of Fastly's security and technology compliance programs. In addition, it provides access to libraries with summary reports of penetration tests, risk assessments, and security policies, as well as an historical archive of security advisories.

Capacity Reservation

<u>Capacity Reservation</u> allows you to reserve Fastly traffic capacity for events.

Cloud Optimizer Limited Availability

Fastly's <u>Cloud Optimizer</u> product allows you to use Fastly's Full-Site Delivery features without migrating edge delivery traffic to Fastly by designating Fastly as the origin to one or more existing non-Fastly content delivery networks currently serving your infrastructure.

Compute@Edge

Fastly's <u>Compute@Edge</u> is a serverless compute environment that allows you to develop, deploy, and operate serverless applications using Fastly's edge cloud platform. You can use Compute@Edge to write new applications, move critical logic closer to end users, and bring enhanced security and performance to current workflows.

https://docs.fastly.com/products/aio 52/70

Concierge TLS

<u>Concierge TLS</u> is a packaged addition to Fastly's <u>Enterprise Support</u> service option that includes one hundred (100) TLS enabled domains, as well as advanced TLS support and configuration options for Enterprises.

Consulting Engagement Services

Fastly <u>Consulting Engagement Services</u> provide high levels of expert support and implementation assistance for customers who require in-house expertise or dedicated resources from our Professional Services and Senior Engineering teams.

Customer Support Services

Fastly Customer Support Services provide answers to questions about features of Fastly products and services. Each member of the Fastly support team provides technical support to resolve questions about account configuration, operation, and management. <u>Support availability and response times</u> vary depending on the level of support you have purchased.

Dedicated IP addresses

Fastly's <u>Dedicated Internet Protocol (IP) addresses</u> provide you with a pool of IP addresses across Fastly's global Edge Cloud.

Designated Technical Specialist

Fastly offers the ability to purchase the support of a Cloud Engineer to serve as a <u>Designated Technical Specialist</u> for your organization. These specialists help you optimize your use of Fastly's products and features and provide enhanced troubleshooting coordination with Fastly's support and professional services organizations.

DDoS Protection and Mitigation Service

Fastly's <u>DDoS Protection and Mitigation Service</u> helps protect against volumetric and targeted distributed denial of service attacks against origin servers. It provides overage insurance for unplanned or unexpected traffic patterns, immediate onboarding assistance, emergency configuration and deployment support, ongoing attack mitigation support, and an incident response plan.

Domain Inspector

<u>Domain Inspector</u> provides you with a dataset and visualizations that offer real-time visibility and historical reporting of domain-level metrics delivered by a Fastly service. It allows you to monitor traffic for a <u>fully qualified domain name</u> without requiring you to send log data to a third-party data collector.

Fastly TLS

<u>Fastly TLS</u> provides TLS management using either certificates that customers upload themselves or Fastly-managed certificates generated by a third-party Certification Authority.

Full-Site Delivery

Fastly's <u>Full-Site Delivery</u> uses Fastly's global content delivery capabilities to cache and accelerate the delivery of static, dynamic, and streaming <u>HTTP-based file content</u>. Full-Site Delivery allows you to tailor delivery of content using features like <u>HTTP header</u> manipulation, <u>time-to-live</u> (TTL) settings, <u>purging</u>, <u>origin shielding</u>, and the advanced edge logic capabilities provided via scripting with the <u>Varnish Configuration Language (VCL)</u>. <u>Transport Layer Security (TLS)</u> and <u>Always-on DDoS mitigation</u> provide security for Fastly services, with real-time monitoring via the <u>Fastly web interface</u>. <u>Log streaming</u> to a variety of third-party endpoints provides observability. Comprehensive <u>APIs</u> power Fastly's web interface and provide programmatic access to Fastly's Full-Site Delivery features.

High Volume Logging

Fastly's <u>High Volume Logging</u> allows you to increase your real-time log streaming log limit beyond the monthly average of two log statements per request, per service.

HIPAA-Compliant Caching and Delivery

https://docs.fastly.com/products/aio 53/70

Fastly offers a <u>HIPAA-Compliant</u> Caching and Delivery product that allows you to transmit protected information like protected health information through Fastly's network.

Image Optimizer

Fastly's <u>Image Optimizer</u> product provides real-time image transformation that caches optimized images requested from your origin server. This product may use third-party cloud infrastructure to process or store content or requests for content.

Live Event Monitoring Limited Availability

Fastly's <u>Live Event Monitoring service</u> offers customers the ability to reserve Fastly customer support resources during their scheduled event's specified hours to proactively monitor key availability and performance metrics. It also offers a dedicated chat channel to communicate with Fastly customer support engineers in real-time.

Logging Insights Package

Fastly's <u>Logging Insights Package</u> helps you analyze and interpret your streaming log data. This professional services offering includes a guided customization of preconfigured third-party logging endpoint dashboards tailored to your specific business needs. Fastly assists with advanced queries, customizations, and best practices.

Media Shield for Live

Fastly's <u>Media Shield for Live</u> product offers the ability to decrease origin traffic of live video events or live linear channels by reducing multiple CDN requests into a single request per shield point of presence (POP) back to your origin. Media Shield for Live works with your existing architecture by making Fastly the origin to all of your end-user-serving CDNs.

Media Shield for VOD

Fastly's <u>Media Shield for VOD</u> product offers the ability to decrease origin traffic of video on demand by reducing multiple CDN requests into a single request per shield point of presence (POP) back to your origin. Media Shield for VOD works with your existing architecture by making Fastly the origin to all of your end-user-serving CDNs.

Nearline Cache Limited Availability

Fastly's <u>Nearline Cache</u> allows you to manually or automatically populate and store content in third-party cloud storage near a Fastly POP.

Next-Gen WAF

The <u>Fastly Next-Gen WAF (powered by Signal Sciences)</u> is a web application firewall that monitors for suspicious and anomalous web traffic and protects, in real-time, against attacks directed at the applications and origin servers that you specify.

Origin Connect

Fastly's Origin Connect product offers a private network interconnect between your origin servers and your Fastly shield POP.

Origin Inspector

<u>Origin Inspector</u> provides you with a dataset and visualizations that offer real-time and historical visibility into responses delivered from your origin servers to the Fastly Edge Cloud. It allows you to monitor origin traffic without requiring you to send log data to a third-party data collector.

Partner Support Services

Fastly partners are required to purchase <u>Partner Support Services</u>. These services provide Fastly platform support benefits as well as a library of on-demand training modules.

PCI-Compliant Caching and Delivery

https://docs.fastly.com/products/aio 54/70

Fastly offers a <u>PCI-Compliant Caching and Delivery</u> product that allows you to transmit protected information like cardholder data through Fastly's network.

Performance Optimization Package

Fastly's <u>Performance Optimization Package</u> provides configuration expertise for analysis and tuning of Fastly services using real-time analytics to identify potential improvements for site performance. This professional services offering includes an assessment, followed by specific recommendations and implementation work.

Platform TLS Limited Availability

Fastly's <u>Platform TLS</u> product allows you to programmatically manage certificates and keys for Transport Layer Security (TLS) using a web API. Use this service if you need to upload thousands or hundreds of thousands of individual X.509 certificates and their associated private keys to Fastly.

Response Security Service

Fastly offers <u>Fastly Next-Gen WAF (powered by Signal Sciences)</u> customers a <u>Response Security Service (RSS)</u> that provides your organization with enhanced access to our Customer Security Operations Center (CSOC) team and periodic consultation with a Designated Security Specialist for strategic security solutions reviews and planning.

Service Implementation

Fastly <u>Service Implementation</u> offers remote planning, customized configurations, testing, and go-live assistance from our Professional Services team for your initial Fastly service implementation and implementation of new Fastly products and services.

Service Management

Fastly <u>Service Management</u> offers ongoing configuration and advanced technical assistance from our Professional Services team on an as-needed basis.

Streaming Delivery

Fastly's <u>Streaming Delivery</u> allows you to use Fastly as a globally distributed HTTP streaming network to improve quality of service and increase viewer capacity for both live and Video On Demand (VOD) content. Streaming Delivery provides all the capabilities of Fastly's <u>Full-Site Delivery</u>, but only for HTTP-based media streaming protocols including HTTP Live Streaming (HLS), HTTP Dynamic Streaming (HDS), HTTP Smooth Streaming (HSS), and Dynamic Adaptive Streaming over HTTP (MPEG-DASH). Fastly's Streaming Delivery must be configured in an account separate from other Fastly accounts.

Subscriber Provided Prefix

Fastly's <u>Subscriber Provided Prefix</u> product allows you to have your IP spaces announced, routed, and served by Fastly infrastructure for use with production services. When you purchase this product, you provide your own IP address space to Fastly rather than use Fastly IP addresses. You can then direct traffic to your own IP addresses, which are reachable via HTTP Anycast on Fastly's infrastructure.

Web Application Firewall (WAF) Quick Start Package

Fastly's <u>WAF Quick Start Package</u> provides you with assistance configuring the initial setup of the Fastly WAF. This professional services offering helps you set up a default policy and configure your WAF in logging mode.

Web Application Firewall (WAF) Tuning Package

Fastly's <u>WAF Tuning Package</u> provide tuning assistance with the configuration of the Fastly WAF. This professional services offering helps you plan your WAF policies and the configuration of the WAF VCL for your Fastly service.

Web Application Firewall (WAF) Tuning Plus Package

https://docs.fastly.com/products/aio 55/70

Fastly's WAF Tuning Plus Package provides ongoing enhanced professional maintenance of your WAF by Fastly. For each service running WAF, the WAF Tuning Plus Package includes ongoing tuning and configuration services as well as authenticated TLS to origin to help protect you against critical security threats. To purchase the WAF Tuning Plus Package, you must have already purchased and provisioned our WAF product. Once purchased, these professional services continue for the term of your WAF contract.

Legacy offerings

These articles provide information about Fastly's legacy offerings.

- https://docs.fastly.com/products/legacy-offerings
- Fastly's Legacy Full-site Delivery services
- Last updated: 2019-08-09
- https://docs.fastly.com/products/fastlys-legacy-full-site-delivery-services

Fastly offers full-site delivery that allows you to speed up websites and mobile apps by pushing content closer to users, providing improved and secure experiences across the world.

IMPORTANT

These terms apply to Subscribers who purchased Full-site Delivery on or before October 8, 2018. For more information about our current <u>Full-Site Delivery</u>, contact <u>sales@fastly.com</u>.

HTTP request fulfillment

The Fastly CDN Service responds to HTTP GET requests initiated from end users' using your website, or from a program making calls to an internet-hosted API.

Header support

Fastly's CDN Service supports forwarding HTTP headers to end users when they are received from your origin server. Alternatively, headers can be added, removed, or modified using our edge scripting language either before or after caching a response from the origin. This includes the Cache-Control and Expires headers as well as the Surrogate-Control header. HTTP header support allows you to send one set of instructions to the Fastly cache servers and another set of instructions to downstream caches, such as proxies or browsers. In particular, the Surrogate-Control header allows you to specify how to forward and transform specific header types.

Time to Live support

Fastly has no set hard limit on how long objects will remain cached. Instead, Fastly supports the expiration of content via Time to Live (TTL) settings that you configure. TTL settings work as timers on your cached content. When content has resided in the cache for the entire TTL interval, that content is given the status of "expired." Before Fastly delivers requested content that is expired, the cache checks to see if the content is still valid by checking with your application server first.

If the application server says the content remains unchanged, the cache sets the content's status to "valid" and resets its TTL value. If the object has been changed, it is declared "invalid" because the content has expired. The application server delivers updated content. Fastly's CDN Service caches the updated content with the status of "valid," and its TTL timer begins to run.

The fetch and refresh process may take a second or more, and during that time, a Fastly cache may receive dozens or hundreds of end-user requests for that content. Fastly's request collapsing feature groups these requests and fulfills them at once when the application server response is received.

Fastly offers you the option of setting a global, default TTL for cached content control. When set, Fastly's CDN service caches objects in a consistent manner even when applications are inconsistent in doing so.

56/70 https://docs.fastly.com/products/aio

Origin shielding

When configuring Fastly's CDN Service, you can designate a specific point of presence (POP) to serve as a shield for your origin servers. This server is referred to as a "shield" because it protects your application servers from continuous requests for content. By default, no origin shield is enabled for you. You must specifically enable shielding to use it.

If Fastly's caches do not have the content being requested, they fetch it from the shield server instead of your origin servers. Fastly caches fetch content from your origin server only when the shield server does not have the content being requested.

Load balancing

You can designate multiple servers as your origin servers. When two or more application servers are provisioned as origin servers, Fastly's CDN Service will distribute requests to fetch content across those application servers. This type of <u>load balancing</u> is enabled by default. You must explicitly disable it if you don't want to use it.

Request collapsing

Cached content sometimes must be refreshed when that content becomes "stale" or expires. When multiple end users request content that is in the process of being refreshed, <u>request collapsing</u> groups those requests to be satisfied together, as soon as the content is received. This accelerates content delivery by keeping Fastly's CDN Service from repeating duplicate requests to your origin server. Request collapsing is enabled by default.

Instant Purge support

Fastly supports an Instant Purge feature that allows you to <u>actively invalidate content</u>. Rather than requiring your network operations and application staff to guess how frequently each bit of content may change, Fastly allows you to generate an HTTP Purge method that is sent to the CDN Service whenever an application changes or deletes data in its database. The Fastly CDN Service invalidates the associated content throughout the service's cache network, causing a new version of that content to be retrieved from the application server the next time it is requested.

Fastly allows URL-based and key-based purging, as well as purging of all content at once via specific, <u>configurable purging</u> <u>commands</u>. Fastly currently supports <u>Ruby, Python, PHP, and Perl libraries</u> for instant purging.

When purging by URL or surrogate key, Fastly's CDN Service can process thousands of changes per second. The invalidation process takes less than 300 milliseconds, making it possible to deliver dynamic content that changes rapidly and unpredictably. Using Instant Purge, you can eliminate cache-to-origin HTTP traffic that all other CDN services generate to determine if expired objects are still valid.

Health checks

You have the option to configure Fastly's CDN Service to <u>perform health checks</u> on your application servers and measure their responsiveness. You can use health check responsiveness measurements to fine-tune the distribution of fetch requests. Health checks are not enabled by default. You must specifically enable them.


Grace mode support

When an application server becomes unavailable for any reason, end users will normally receive error messages indicating the content they've requested cannot be retrieved. When enabled, grace mode shields application servers by instructing Fastly's CDN Service to continue to serve stale or expired (but likely still valid) content to end users for a set amount of time. This allows you to return otherwise unavailable application servers to normal operations and still serve content rather than error messages to end users. By default, grace mode is not configured. You must specifically configure you service to serve stale content to use grace mode.

- Fastly's Legacy Media Shield
- **t** Last updated: 2019-09-12
- https://docs.fastly.com/products/fastlys-legacy-media-shield

IMPORTANT

These terms apply to Subscribers who purchased Media Shield on or before September 12, 2019. For more information about our current Media Shield product, contact sales@fastly.com.



Fastly Media Shield offers customers the ability to decrease origin traffic by <u>reducing multiple CDN requests</u> into a single request back to your origin. Media Shield works with your existing architecture by making Fastly the origin to all of your end-user-serving CDNs.

To learn more about Fastly's Media Shield, contact your account manager or email sales@fastly.com for more details.

IMPORTANT

This information is part of a limited availability release. For additional details, read our <u>product and feature lifecycle</u> descriptions.

- Legacy Customer-Provided TLS Certificate Hosting Service
- **t** Last updated: 2020-04-01
- <u>https://docs.fastly.com/products/legacy-customer-provided-tls-certificate-hosting-service</u>

O NOTE

Fastly maintains support for its original Customer-Provided TLS Certificate Hosting Service. For more information about our current <u>TLS service options</u>, contact <u>sales@fastly.com</u>.

Fastly offers a Customer-Provided TLS Certificate Hosting Service where you provide TLS certificates and private keys which we then install at a shared set of IP addresses. Each are selected using the SNI extension of TLS that allows clients to present a hostname in the TLS handshake request. Choose this option if you have requirements that prevent you from using the Fastly TLS interface to upload your TLS certificates and private keys personally.

To purchase this option, contact sales@fastly.com.

Legacy Platinum Support and SLA

https://docs.fastly.com/products/aio 58/70

tast updated: 2018-03-30

Q

https://docs.fastly.com/products/legacy-platinum-support-and-sla



NOTE

Fastly maintains support for its original Platinum Support plan. For more information about our current <u>Gold and Enterprise</u> <u>Support plans</u> or for information about our <u>Professional Services packages</u>, contact <u>sales@fastly.com</u>.

Legacy Platinum Support description and SLA

Support availability and response times vary depending on the type of account you have and the level of support you have purchased. The following table summarizes those offerings:

Support Offering	Platinum Support
Online Self-Service Help	Unlimited access.
Availability for General Inquiries	24/7/365.
Availability for Incident Reports	24/7/365.
Initial Response Times	Severity 1 Incidents within 15 minutes. Severity 2 Incidents within 2 hours. All other Incidents by the next business day.
Email support	Available, with priority over Standard and Gold Support.
Phone and chat support	Toll-free telephone available 24/7/365. Dedicated chat channel available during Fastly business hours.
Emergency Escalation	Available via email and phone.
Designated Customer Support Engineer	Available for large accounts on case-by-case basis.
Termination Option	Available with invoice credits.

Technical support

The following section applies to all subscribers.

Definitions

- "Business Hours" are 8AM-6PM during a Business Day in California, New York, London, or Tokyo.
- "Business Days" are Monday through Friday, excluding any day that is simultaneously a US, UK, and Japanese national or banking holiday.
- An "Incident" is an occurrence during which end users' use of Subscriber's services is adversely impacted.
- A "Severity 1 Incident" is an incident resulting in a major service outage requiring Subscriber to redirect all traffic from Fastly to another CDN.
- A "Severity 2 Incident" is an incident resulting in minor or intermittent outage not requiring Subscriber to redirect traffic to another CDN.
- "Fastly Control" means elements entirely under Fastly's control and not a consequence of (a) a Subscriber's hardware or software failures, (b) a Subscriber's or end user's connectivity issues, (c) Subscriber operator errors, (d) Subscriber traffic amounts that exceed a Subscriber's Permitted Utilization as defined in the Terms and Conditions, (e) corrupted Subscriber content, (f) acts of god (any) or war, or earthquakes, or terrorist actions.

Subscriber responsibilities

Subscriber is responsible using and configuring services according to the Documentation available at https://docs.fastly.com.

https://docs.fastly.com/products/aio 59/70

Support requests

Subscribers submit support requests by sending email to support@fastly.com. Subscribers receive a system-generated response within minutes containing the ticket number and a direct link to the ticket.

Incident reports should include at the least the following:

- Services are not responding to end user requests.
- Services incorrectly send end users error condition messages.
- Services send incorrect or partial content to end users.

Incident reports should include all relevant information such as:

- Subscriber's determination of the Severity Level of the incident,
- Subscriber hardware failures,
- · Subscriber operator errors,
- · Services configuration errors made by Subscriber employees,
- A potential Utilization Spike (see the <u>Service Availability SLA</u>),
- Corrupted Subscriber content,
- DDOS attacks, and
- Relevant force majeure acts such as extreme weather, earthquakes, strikes or terrorist actions.

Communications

Tickets

Communications between Fastly support engineers and Subscriber personnel are conducted using the ticketing application, which maintains a time-stamped transcript of communications, and sends emails to Subscriber and Fastly staff as tickets are updated.

Chat

Subscribers to Platinum Support receive a dedicated chat channel for real-time communications during Business Hours. Though subject to change, Fastly's current chat provider is Slack (www.slack.com).

Phone support

Subscribers to Platinum Support receive a dedicated phone number to contact Fastly support engineers. Fastly personnel can also establish audio and video conferencing (free app or browser plug-in required) for real-time voice and video communications.

Response time

Fastly shall use best efforts to respond in a timely fashion.

Termed contracts

The following applies to any subscriber that has a contract with a term and a minimum commitment.

Response times

Fastly commits to acknowledging receipt of a support ticket within the next Business Day following submission of a support request by a Subscriber with a Termed Contract.

Termination

In any three-month period where three (3) or more support Response Time objectives are not met and the failure to meet the objectives materially adversely impacted Subscriber, Subscribers with a Termed Contract, Platinum Support shall have thirty (30) days to terminate their subscription agreement following the third failure.

Incident response times

Incident reporting

https://docs.fastly.com/products/aio 60/70

Severity 1 Incidents: Fastly will provide Subscriber an Incident Support Email address for Subscriber to report Incidents. Subscriber should report Incidents promptly using the Incident Support email.

Severity 2 Incidents: Subscriber should report Severity 2 Incidents by submitting a Support Request.

Platinum Support

Fastly will respond to the report of an Incident by troubleshooting the cause(s) of the Incident and resolve them if caused by factors within Fastly's control, or provide information to those who can resolve the factors if the factors are within others' control, as follows:

For a Severity 1 Incident:

- Fastly support staff will acknowledge receipt of the email within 15 minutes.
- Fastly will start actively troubleshooting within 30 minutes of receipt of the email.
- Fastly will perform its tasks on a 24/7 basis.
- Fastly and Subscriber will immediately communicate upon learning new information that may be useful in troubleshooting the incident, and status updates between Fastly and Subscriber staff will take place no less frequently than every 30 minutes for the first two hours, and no less frequently than every hour thereafter.
- Fastly staff will work until (a) the incident is resolved or (b) the incident is believed to be outside of Fastly's control.

For a Severity 2 Incident:

- Fastly support staff will acknowledge receipt of the email within two hours.
- Fastly engineers will begin actively troubleshooting within the same day, will work on the Incident during the same day, and will provide status updates to Subscriber daily on each subsequent day.

Support invoice credits

In the event a Severity 1 Incident occurs, Subscriber has purchased Platinum Support, the cause of the Incident is within Fastly's control, and any of the communication or response timeframes are materially not met, a one-time credit of \$500 per incident will be credited to Subscriber's account.

Credit Terms:

- Requests for Invoice Credits must be made within 30 days of the incident which triggered the service credit.
- In no event shall Invoice Credits exceed the invoice value of the month in which they are accrued.
- A pending credit does not release Subscriber from its obligation to pay Fastly's submitted invoices in full when due.
- Credits will be applied to the invoice two months following the month an invoice credit was incurred.

Legacy Service availability SLA

Support availability and response times vary depending on the <u>type of account</u> you have and the <u>level of support</u> you have purchased.

Agreement Type	Unpaid Account	Month-to-Month Account	Termed Contract	Platinum Support
Service Level Agreement	None	None	Termination Option	Invoice Credits + Termination Option

Definitions

"Degraded Performance" means the Services are experiencing Error Conditions that are (1) caused by issues under Fastly Control, (2) observable or reproducible by Subscriber or Fastly, (3) requiring Subscriber to redirect traffic off the Services. Degraded Performance does not include any reduction on availability of the Application Web Interface or API due to maintenance.

"Error Condition" means the Services are (1) not responding to end user requests, (2) incorrectly sending end users error condition messages or (3) sending incorrect partial content to end users and these conditions are observable or reproducible by Subscriber or Fastly.

https://docs.fastly.com/products/aio 61/70

"Fastly Control" means elements entirely under Fastly's control and not a consequence of (a) Subscriber hardware or software failures, (b) Subscriber or end user connectivity issues, (c) Subscriber operator errors, (d) a Utilization spike (see below), (e) corrupted Subscriber content, (f) acts of god (any) or war, or earthquakes, or terrorist actions.

Termination

Any Subscriber that has a contract with a term and a minimum commitment shall have thirty (30) days to terminate their subscription agreement following (1) a period of Degraded Performance longer than 7.2 hours in any one month, or (b) three contiguous months that have periods of Degraded performance longer than 43.8 minutes each.

Availability invoice credits

Subscribers who purchase Platinum Support shall be entitled to Invoice Credits according to the following table.

Availability Percent	Period of Degraded Performance	Monthly Credit Percent
Below 100% - 99.99%	Up to 4.32 minutes	1%
99.99% – 99.9%	Up to 43.8 minutes	5%
99.89% – 99.0%	Up to 7.2 hours	10%
98.99% - 98.0%	Up to 14.4 hours	25%
Below 97.99%	Greater than 864 minutes	50%

Invoice Credits for unavailability will accrue on a monthly basis. The Credit Amount for a month is equal to the monthly usage charge multiplied by Monthly Credit Percent.

Credit terms

- Requests for Invoice Credits for Degraded Performance must be made within 30 days of the period of Degraded Performance.
- The maximum amount of any credit is the Invoice Amount for the month the Degraded Performance occurred.
- A pending credit does not release Subscriber from its obligation to pay Fastly's submitted invoices in full when due.
- Credits will be applied to the Invoice two months following the month an invoice credit was incurred.

Utilization Spikes

Subscriber's bandwidth utilization, measured in megabits per second, will be sampled every five (5) minutes on a region-by-region basis each month (the "Samples"). Subscriber's "Average Utilization" for a region in a month will be the average of the Samples. Subscriber's "Peak Utilization" for a region in a month will be calculated by the 95th percentile method, according to which the Samples will then be ordered from highest to lowest, and the highest five percent (5%) of Samples will be discarded and the remaining highest Sample will be Subscriber's Peak Utilization for the region in that month. Subscriber's "Permitted Utilization" in a month for a region will be five (5) times Subscriber's Average Utilization in that month for that region. A "Utilization Spike" will occur if Subscriber's Peak Utilization exceeds its Permitted Utilization in a region. Utilization Spikes may interfere with or disrupt the integrity or performance of the Services. Subscribers should contact Support in advance of any planned utilization spike and respond immediately to any communications from Fastly regarding an actual or suspected Utilization Spike.

- Legacy Premium Support and SLA
 ★ Last updated: 2018-03-30
 ★ https://docs.fastly.com/products/legacy-premium-support-and-sla
 - NOTE

Fastly maintains support for its original Premium Support plan. For more information about our current <u>Gold and Enterprise</u> <u>support plans</u> or for information about our <u>Professional Services packages</u>, contact <u>sales@fastly.com</u>.

https://docs.fastly.com/products/aio 62/70

Legacy Premium Support description and SLA

Support availability and response times vary depending on the type of account you have and the level of support you have purchased. The following table summarizes those offerings:

Offering	Unpaid Account	Month-to- Month Account	Termed Contact	Premium Support
Online Forums	Yes	Yes	Yes	Yes
Email Support Response Time Commitment	Best Effort	Best Effort	Next Business Day	Severity 1 Incidents: 15 minutes*. All Others: Next Business Day
Severe Incident Response Email Address	No	No	No	Yes
Support SLA	None	None	Termination Option	Invoice Credits + Termination Option

Technical support

The following section applies to all Subscribers.

Definitions

- "Business Hours" are 8AM-6PM Monday through Friday, Pacific Time.
- "Business Days" are Monday through Friday excluding US and UK national and banking holidays.
- An "Incident" is an occurrence during which an end user's use of Subscriber's services is adversely impacted.
- A "Severity 1 Incident" is an Incident resulting in a major service outage requiring Subscriber to redirect all traffic from Fastly to another CDN.
- A "Severity 2 Incident" is an Incident resulting in minor or intermittent outage not requiring Subscriber to redirect traffic to another CDN.
- "Fastly Control" means elements entirely under Fastly's control and not a consequence of (a) Subscriber hardware or software failures, (b) Subscriber or end user connectivity issues, (c) Subscriber operator errors, (d) Subscriber traffic amounts that exceed Subscriber's Permitted Utilization as defined in the Terms and Conditions, (e) corrupted Subscriber content, (f) acts of god (any) or war, or earthquakes, or terrorist actions.

Subscriber responsibilities

Subscriber is responsible for using and configuring services according to the Documentation available at https://docs.fastly.com.

Support requests

Subscribers submit support requests by sending email to support@fastly.com. Subscribers receive a system-generated response within minutes containing the ticket number and a direct link to the ticket.

Incident reports should include at the least the following:

- Services not responding to end user requests.
- Services incorrectly sending end users error condition messages.
- Services sending incorrect or partial content to end-users.

Incident reports should include all relevant information, such as:

- Subscriber's determination of the Severity Level of the Incident,
- Subscriber hardware failures,
- Subscriber operator errors,
- Services configuration errors made by Subscriber employees,
- Potential Excess Utilization (as defined in the Terms of Use or master services agreement),

https://docs.fastly.com/products/aio 63/70

- · Corrupted Subscriber content,
- · DDOS attacks, and
- Relevant force majeure acts such as extreme weather, earthquakes, strikes or terrorist actions.

Communications

Communications between Fastly support engineers and Subscriber staff are conducted using the ticketing application, which maintains a time-stamped transcript of all communications, and sends emails to Subscriber and Fastly staff as tickets are updated.

Response time

Fastly shall use best efforts to respond in a timely fashion.

Termed contracts

The following applies to any Subscriber that has a contract with a term and a minimum commitment.

Response times

Fastly commits to acknowledging receipt of a support ticket within the next business day following submission of a support request.

Termination

In any three-month period where three (3) or more support Response Time objectives are not met and the failure to meet the objectives materially adversely impacted Subscriber, Subscriber shall have thirty (30) days to terminate their subscription agreement following the third failure.

Premium Support

The following applies to Subscribers who have purchased Premium Support.

Incident reporting

Severity 1 Incidents: Fastly will provide Subscriber an Incident Support Email address for Subscriber to report Incidents. Subscriber should report Incidents promptly using the Incident Support email.

Severity 2 Incidents: Subscriber should report Severity 2 Incidents by submitting a Support Request.

Response times

Fastly will respond to the report of an Incident by troubleshooting the cause(s) of the Incident and resolve them if caused by factors within Fastly's control, or provide information to those who can resolve the factors if the factors are within others' control, as follows:

For a Severity 1 Incident:

- Fastly support staff will acknowledge receipt of the email within 15 minutes.
- Fastly will start actively troubleshooting within 30 minutes of receipt of the email.
- Fastly will perform its tasks on a 24/7 basis.
- Fastly and Subscriber will immediately communicate upon learning new information that may be useful in troubleshooting the Incident, and status updates between Fastly and Subscriber staff will take place no less frequently than every 30 minutes for the first two hours, and no less frequently than every hour thereafter.
- Fastly staff will work until (a) the Incident is resolved or (b) the Incident is believed to be outside of Fastly's control.

For a Severity 2 Incident:

- During Business Hours, Fastly support staff will acknowledge receipt of the email within two hours or within two hours of the start of the next business day if the Incident does not come in during a Business Day.
- Fastly engineers will begin actively troubleshooting within one business day, will work on the Incident during Business Hours, and will provide status updates to Subscriber daily on each subsequent Business Day.

Charges for Incident Response

https://docs.fastly.com/products/aio 64/70

For Severity 1 Incidents caused by factors within Subscriber's control, a flat fee of \$1500 will be assessed, and any time spent beyond three hours will be invoiced at Subscriber's undiscounted Professional Services rates.

For Severity 2 Incidents caused by factors within Subscriber's control, Subscriber will be invoiced at Subscriber's undiscounted Professional Services Rates.

For all Incidents:

- If the Incident-causing factors are within Fastly's control, there will be no hourly charges for Fastly engineering staff time.
- If the factors are within Subscriber's control, Subscriber agrees to pay Fastly its hourly charges for Fastly engineering staff time. If it appears likely the factors are within Subscriber's, Subscriber may tell Fastly staff to stop working on troubleshooting the Incident (thereby stopping the hourly charges from being incurred). Subscriber agrees to tell Fastly to stop working on an Incident via an email sent to Fastly's Incident Support email address. The timestamp on the email will be the time charges cease to be incurred.

Support Invoice Credits

In the event a Severity 1 Incident occurs, Subscriber has purchased Premium Support, the cause of the Incident is within Fastly's control, and any of the communication or response timeframes are materially not met, a one-time credit of \$500 per Incident will be credited to Subscriber's account.

Credit Terms:

- Requests for Invoice Credits must be made within 30 days of the Incident which triggered the service credit.
- In no event shall Invoice Credits exceed the invoice value of the month in which they are accrued.
- A pending credit does not release Subscriber from its obligation to pay Fastly's submitted invoices in full when due.
- Credits will be applied to the invoice two months following the month an invoice credit was incurred.

Legacy Service availability SLA

Support availability and response times vary depending on the <u>type of account</u> you have and the <u>level of support</u> you have purchased.

Agreement Type	Unpaid Account	Month-to-Month Account	Termed Contract	Premium Support
Service Level Agreement	None	None	Termination Option	Invoice Credits + Termination Option

Definitions

"Degraded Performance" for the Services means the Services are experiencing Error Conditions that are (1) caused by issues under Fastly Control, (2) observable or reproducible by Subscriber or Fastly, (3) requiring Subscriber to redirect traffic off the Services. Degraded Performance does not include any reduction on availability of the Application Web Interface or API due to planned maintenance.

"Error Condition" means the Services are (1) not responding to end user requests, (2) incorrectly sending end users error condition messages or (3) sending incorrect partial content to end users and these conditions are observable or reproducible by Subscriber or Fastly.

"Fastly Control" means elements entirely under Fastly's control and not a consequence of (a) Subscriber hardware or software failures, (b) Subscriber or end user connectivity issues, (c) Subscriber operator errors, (d) Subscriber traffic amounts that exceed Subscriber's Permitted Utilization, (e) corrupted Subscriber content, (f) acts of god (any) or war, or earthquakes, or terrorist actions.

Termination

Any Subscriber that has a contract with a term and a minimum commitment shall have thirty (30) days to terminate their subscription agreement if the Services experience Degraded Performance (a) for longer than 7.2 hours in any one month, or (b) for longer than 43.8 minutes each month in any three contiguous months. Subscriber shall have thirty (30) days to terminate their contract following the third failure.

https://docs.fastly.com/products/aio 65/70

Availability of invoice credits

Subscribers who purchase Premium Support shall be entitled to Invoice Credits according to the following table.

Availability Percent	Period of Degraded Performance	Monthly Credit Percent
Below 100% - 99.99%	Up to 4.32 minutes	1%
99.99% – 99.9%	Up to 43.8 minutes	5%
99.89% – 99.0%	Up to 7.2 hours	10%
98.99% - 98.0%	Up to 14.4 hours	25%
Below 97.99%	Greater than 864 minutes	50%

Invoice Credits for unavailability will accrue on a monthly basis. The Credit Amount for a month is equal to the monthly usage charge multiplied by Monthly Credit Percent.

Credit terms

- Requests for Invoice Credits for Degraded Performance must be made within 30 days of the period of Degraded Performance.
- The maximum amount of any credit is the Invoice Amount for the month the Degraded Performance occurred.
- A pending credit does not release Subscriber from its obligation to pay Fastly's submitted invoices in full when due.
- Credits will be applied to the Invoice two months following the month an invoice credit was incurred.
- Web Application Firewall (WAF) (2020)
- **tast updated: 2020-07-14**
- https://docs.fastly.com/products/web-application-firewall-2020

IMPORTANT

As of June 30, 2021, the Fastly WAF (WAF 2020) offering became a legacy product. It will continue to be supported for all existing users. As an alternative, <u>Fastly Next-Gen WAF (powered by Signal Sciences)</u> offers proactive monitoring of and protection against suspicious and anomalous web traffic directed at your applications and origin servers. It can be controlled via the <u>web interface dashboard</u> or <u>application programming interface</u> (API). Contact <u>sales@fastly.com</u> or your Fastly account team to evaluate or move to the Fastly Next-Gen WAF option.

The Fastly WAF is a <u>Web Application Firewall (WAF)</u> security product that detects malicious request traffic sent over HTTP and HTTPS. Once properly configured and enabled for a service, the Fastly WAF can help protect against application-layer (layer 7) attacks such as SQL injection, cross-site scripting, and HTTP protocol violations.

Enabling Fastly WAF doesn't require modifications to your web application or origin servers. Contact <u>our sales team</u> to get started. Once you purchase the Fastly WAF, our <u>customer support team</u> will enable it with a default WAF policy for any service you've provided a service ID for. They will then work closely with you on additional configuration refinements. Once configured, you can then begin monitoring logs to determine which requests to your origin are legitimate and which you should consider blocking.

Limitations

All WAF products that exist today, including the Fastly WAF, have several limitations:

- **False positives.** Any WAF can mistake good traffic for bad. We strongly recommend you <u>monitor your logs</u> for a minimum of two weeks before blocking traffic. You don't want to start blocking traffic with rules that are generating false positives.
- **DNS configuration.** A WAF only works when traffic is directed through it. It cannot protect against malicious requests that are sent to domain names or IP addresses that are not specified in your WAF configuration.
- **Effective rules.** A WAF is only as effective as the provisioned and tuned rules. You can add, remove, or modify rule modes using rule management web interface or the API.

https://docs.fastly.com/products/aio 66/70

• Custom application vulnerabilities. If attackers discover a vulnerability unique to your application or the technologies you use, and if your WAF configuration does not have a rule to protect against exploits for that particular vulnerability, it will not be able to protect your application in that instance. You can add additional rules to help protect against these types of attacks. If you need more protection than the selected rules provide, customer support can work with you to create custom VCL to help block malicious requests.

• Inspection of HTTP and HTTPS traffic only. A WAF only inspects HTTP or HTTPS requests (layer 7). It will not process any TCP, UDP, or ICMP requests.

Security products note



IMPORTANT

To ensure your web application only receives traffic from your WAF-enabled Fastly service, we strongly recommend you configure <u>TLS client authentication</u> for that service and allowlist <u>Fastly's assigned IP ranges</u>.

No security product, such as a WAF or DDoS mitigation product, including those security services offered by Fastly, will detect or prevent all possible attacks or threats. As a subscriber, you should maintain appropriate security controls on all web applications and origins. The use of Fastly's security products do not relieve you of this obligation. As a subscriber, you should test and validate the effectiveness of Fastly's security services to the extent possible prior to deploying these services in production, continuously monitor their performance, and adjust these services as appropriate to address changes in your web applications, origin services, and configurations of the other aspects of your Fastly services.

Third-party information

These articles provide information about third-party technology and services incorporated into the Fastly CDN service.

https://docs.fastly.com/products/third-party-information

Cloud-hosted products

iii Last updated: 2022-02-02

https://docs.fastly.com/products/cloud-hosted-products

The following Fastly products may use third-party cloud infrastructure to process or store content or requests for content according to our <u>cloud infrastructure security program</u>:

- Fastly Image Optimizer
- Fastly WAF Support and SLA
- <u>Live Event Monitoring</u>

The following Fastly products may use third-party cloud infrastructure to process or store content or requests for content according to this <u>cloud infrastructure data center and physical security program</u>:

- Fastly Next-Gen WAF (powered by Signal Sciences)
- Open source software in downloadable components
- **iii** Last updated: 2022-01-26
- https://docs.fastly.com/products/open-source-software-in-downloadable-components

The following table provides information about open source technology incorporated into the downloaded components for Next-Gen WAF.

https://docs.fastly.com/products/aio 67/70



O NOTE

This page will be updated to identify changes to open source third-party technology included in the current version of the downloadable software components.

Technology	Version	License
Datadog Go	2.2.0+incompatible	MIT
glob	0.2.3	MIT
Go App Engine packages	1.6.1	Apache 2.0
go-codec	1.1.7	MIT
go-diff	1.0.1-0.20180205163309-da645544ed44	MIT
gogoprotobuf	1.3.1	NewBSD
gohistogram	1.0.0	MIT
<u>go-license</u>	0.0.0-20180405065157-c69f41c2c8d6	MIT
Go Networking	0.0.0-20211112202133-69e39bad7dc2	NewBSD
Google APIs generated by gogoprotobuf	1.3.0	Apache 2.0
gopsutil	3.21.8+incompatible	NewBSD
<u>go-reuseport</u>	0.0.1	ISC
gorilla/mux	1.7.3	NewBSD
Gorilla Websocket	1.4.2	FreeBSD
<u>go-version</u>	1.2.1	MPL-2.0
GoVersionInfo	0.0.0-20190209210621-63e6d1acd3dd	MIT
gRPC-Go	1.27.0	Apache 2.0
gRPC Go Middleware	1.2.0	Apache 2.0
gRPC Go Proxy	0.0.0-20181017164139-0f1106ef9c76	Apache 2.0
jaeger-client-go	2.22.1+incompatible	Apache 2.0
jose	1.1.2	Apache 2.0
Logrus	1.4.2	MIT
<u>mapstructure</u>	1.1.2	MIT
<u>mgo</u>	0.0.0-20181015135952-eeefdecb41b8	FreeBSD
msgp	1.1.2	MIT
OpenTracing API for Go	1.1.0	Apache 2.0
<u>Oxy</u>	1.1.0	Apache 2.0
<u>pflag</u>	1.0.5	NewBSD
pkg/errors	0.9.1	FreeBSD
<pre>pkg/profile</pre>	1.4.0	FreeBSD
protobuf	1.4.2	NewBSD
<u>protoc-gen-validate</u>	0.1.0	Apache 2.0
reopen	1.0.0	MIT

https://docs.fastly.com/products/aio 68/70

Technology	Version	License
snappy	0.0.1	NewBSD
sys	0.0.0-20220114195835-da31bd327af9	NewBSD
viper	1.7.0	MIT

■ Sub-processors
 ■ Last updated: 2022-05-23
 ② https://docs.fastly.com/products/sub-processors

Fastly engages certain sub-processors in connection with the provision of the Fastly Services. A sub-processor is a Fastly affiliate engaged in the processing of personal data (each a "Fastly Affiliate") or a third-party service provider engaged by Fastly, Inc. or a Fastly Affiliate to process personal data on behalf of Fastly's Subscribers.

Fastly maintains a list of the names, entity type and locations of all sub-processors of personal data contained in Subscriber Data and caused to be submitted to Fastly via the Services according to Subscriber's configuration of the Services, which is set forth below. For more information on Fastly's data processing obligations, please see our <u>data processing terms</u>.

Name of Sub-Processor	Entity Type	Entity Location
Fastly Australia Pty Ltd	Fastly Affiliate	Australia
Fastly Canada Inc.	Fastly Affiliate	Canada
Fastly Cloud Iberica, S.L.	Fastly Affiliate	Spain
Fastly GmbH	Fastly Affiliate	Germany
Fastly India Private Limited	Fastly Affiliate	India
Fastly International (Holdings) Limited	Fastly Affiliate	United Kingdom
Fastly Kabushiki Kaisha	Fastly Affiliate	Japan
Fastly Limited	Fastly Affiliate	United Kingdom
Fastly Netherlands B.V.	Fastly Affiliate	The Netherlands
Fastly SG Pte. Ltd.	Fastly Affiliate	Singapore
Fastly Sweden AB	Fastly Affiliate	Sweden
Signal Sciences, LLC	Fastly Affiliate	United States
Google LLC	Third-party sub-processor	United States
Amazon Web Services, Inc.	Third-party sub-processor	United States
Microsoft Corporation	Third-party sub-processor	United States
MongoDB Atlas	Third-party sub-processor	United States

Subscribers may subscribe to receive email notifications of sub-processor changes at https://docs.fastly.com/changes. Notices of updates to this page will be posted on our changelog.

Third-party technology

Last updated: 2022-02-02

https://docs.fastly.com/products/third-party-technology

These articles provide information about third-party technology and services incorporated into the Fastly CDN service:

https://docs.fastly.com/products/aio 69/70

- Geolocation VCL features
- <u>Device detection variables</u>
- Nearline Cache
- TLS service options

In addition, these articles describe third-party services that interoperate with Fastly CDN services:

- Integrations with Non-Fastly Services
- Streaming logs

The Fastly Next-Gen WAF (powered by Signal Sciences) uses the MaxMind GeoLite2 databases.

• <u>Fastly status</u> <u>www.fastly.com</u>

<u>Sitemap</u> | <u>Changelog</u> | <u>Translations</u> | <u>Archives</u>

Copyright © 2022 Fastly Inc. All Rights Reserved.

Policy FAQ | Acceptable Use | Terms of Service | Privacy

https://docs.fastly.com/products/aio 70/70