

Category: 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.



Cache Reservation



Last updated: 2025-05-06



</products/cache-reservation>

Cache Reservation provides access to the caching layer at Fastly's edge where you can reserve cache space specifically for your content in Fastly shielding locations. By prioritizing your content's cache storage, Cache Reservation allows that content to stay in cache longer by minimizing content eviction in these multi-tenant environments. Cache Reservation helps optimize your origin's offload from any CDN, including Fastly, reducing your cloud egress costs.

To learn more about Cache Reservation, contact your customer success manager or email sales@fastly.com for additional details.

Prerequisites

To purchase Cache Reservation you must have a paid account for [full-site delivery](#) or [streaming delivery](#) and you must enable [shielding](#). In addition, all cache reservations are subject to pre-qualification, specific to the size of your content and the shielding locations requested. Specifically:

- the ID of the specific service you want to have associated with your cache reservation.
- the shielding location for the traffic managed by the specific service.
- the approximate size of the objects included in this reservation.

- an estimate of how much space (in GB) the actively consumed content for each service will require at the chosen shielding location. We can help you with this estimate. Fastly will use this information in the qualification process to determine if it is possible for your company to use Cache Reservation. To see if your company meets the qualification criteria, contact sales@fastly.com.

Limitations and considerations

Keep in mind the following limitations and considerations:

- Cache reservation is not compatible with Fastly's [Compute platform](#).
- Reservations are restricted to [shield POPs](#) only.
- Each reservation covers a single shielding location.
- Reservations apply to your cache as a whole, not to individual objects, and you cannot see what is in the cache at any given time.
- Reservations minimize cache storage eviction via prioritization up to the specified reservation size, but do not ultimately prevent eviction.

Billing

Cache Reservation charges are billed based on reservation size (in GB) in specific shielding locations. Consider setting a [spend alert](#) for your account to automatically be notified when your month-to-date bill approaches or meets a specific dollar amount.

* * *



Capacity Reservation



Last updated: 2023-04-05



</products/capacity-reservation>

Capacity Reservation allows you to reserve Fastly traffic capacity for events based on data bandwidth (in qigabits 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.

Billing

NOTE

Billing limits for this product may be different depending on your [account type](#), if you've purchased a [packaged offering](#), or are using a [product or feature trial](#).

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. Consider setting a [spend alert](#) for your account to automatically be notified when your month-to-date bill approaches or meets a specific dollar amount.

Some Fastly products can be purchased directly in the web interface on the [Products page](#). For more details about a product, including [pricing information](#) or for help purchasing it, contact your customer success manager or email sales@fastly.com.

* * *



Cloud Optimizer



Last updated: 2020-12-18



</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 customer success manager or email sales@fastly.com for more details.

NOTE

Cloud Optimizer is not available for video streaming activities. Check out [Media Shield for Live](#) and [Media Shield for VOD](#) instead.

Some Fastly products can be purchased directly in the web interface on the [Products page](#). For more details about a product, including [pricing information](#) or for help purchasing it, contact your customer success manager or email sales@fastly.com.

* * *



Dedicated IP addresses



Last updated: 2023-09-28



</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 [zero-rated billing](#)

endpoints or security allowlisting.

Purchase of Fastly's [Platform TLS product](#) requires you to also have purchased Dedicated IP addresses.

TLS non-SNI client support

Fastly-managed certificates require clients to support TLS v1.2 and Server Name Indication (SNI) by default. When there is not an SNI match, a fallback certificate is used. Dedicated IP addresses can be used to host fallback certificates for non-SNI client support. Using the Fastly API, you can place your self-managed or Fastly-managed certificates at a dedicated set of IP addresses in the event there is no SNI match.

Additionally, for self-managed certificates only, you can also indicate a default ECDSA and RSA certificate. When no SNI match is found, Fastly will first check if the client supports ECDSA. If it does, we will send the fallback ECDSA certificate. If there is no SNI match and the client does not support ECDSA, we send the RSA fallback certificate.

Use your DNS records to associate the set of dedicated IP addresses to use to direct incoming traffic to the Fastly edge network. There are three possible network routing options (sometimes referred to as network maps or domain maps) that allow you to choose which sub-regions of the Fastly network to use.

Read [Fastly's TLS offerings](#) 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.

Some Fastly products can be purchased directly in the web interface on the [Products page](#). For more details about a product, including [pricing information](#) or for help purchasing it, contact your customer success manager or email sales@fastly.com.

* * *



Domain Research API



Last updated: 2025-11-04



</products/domain-research-api>

Fastly's [Domain Research API](#) allows you to programmatically retrieve algorithmic domain search results and check domain availability. The Domain Research API might be useful if you want to:

- check if domain names are available for registration or for sale in aftermarkets.
- add a domain discovery experience to your product.
- monitor domain name state changes over time.

How it works

The Domain Research API is part of the Fastly API, which requires a [Fastly API token](#). You

can use the Fastly API to [enable the Domain Research API for your account](#).

The Domain Research API has two endpoints. The Suggest endpoint allows you to retrieve domain search results, and the Status endpoint allows you to check domain availability.

The Status endpoint has two variants. The Precise variant allows you to check the registry-level availability of a domain (e.g., for new registrations), and the Estimated variant allows you to check DNS and aftermarket-level availability of a domain (e.g., not registry-level).

Limitations and considerations

The Domain Research API has the following limitations and considerations:

- The most common use case of the Status endpoint is to determine if a domain is available for registration from a registrar. The `inactive` status means a domain is available for registration.
- By design, the Status endpoint accepts only a single domain per API request. If your use case calls for checking multiple domains, you'll need to check them one at a time.
- Our platform limits Domain Research API requests to a total runtime of 30 seconds. If an upstream provider fails to respond within that timeframe, our API will return a best-effort value, typically `unknown` or `undelegated`.
- The Suggest and Status endpoints are intentionally decoupled. Domain availability information is not present in Suggest endpoint responses, and Status endpoint responses do not include search result suggestions.

Billing

The Domain Research API is disabled by default. Users assigned the superuser role can enable it on the [Products](#) page.

Billing for the Domain Research API is based on Domain Research API requests. There are three types of billable Domain Research API requests:

- **Suggest API requests:** The number of HTTP requests sent to the Suggest endpoint.
- **Status-Precise API requests:** The number of HTTP requests sent to the Status endpoint.

- **Status-Estimated API requests:** The number of HTTP requests sent to the Status endpoint, with the `&scope=estimate` parameter explicitly set. These responses include a subset of the data returned by the Status-Precise API consisting of (i) DNS-level domain delegation (or undelegated) data and (ii) Aftermarket listing metadata (if present for a given domain). These responses do not include domain registry-level availability data.

Consider setting a [spend alert](#) for your account to automatically be notified when your month-to-date bill approaches or meets a specific dollar amount.

* * *



Fastly DNS



Last updated: 2026-07-28



</products/fastly-dns>

[Fastly DNS](#) is our authoritative DNS service that leverages Fastly's global, high-performance edge network to serve DNS query responses. Fastly DNS supports the primary-secondary DNS architecture, which requires you to store and manage your primary “source of truth” DNS records elsewhere. By adding DNS zones, you configure your primary DNS service to add Fastly as a secondary provider. Once delegated at your domain name registrar, Fastly DNS can serve your domains' authoritative responses.

Prerequisites and considerations

Keep in mind the following prerequisites and considerations:

- To purchase DNS, you must have a Service Order for [Fastly's Full-Site Delivery](#) or [Fastly's Streaming Delivery](#) or have purchased one of our [packaged offerings](#).
- Fastly DNS is not a primary DNS service. You must store and manage your primary DNS records outside Fastly. Whichever service you use must support [zone transfers \(AXFR\)](#) in order to sync records with Fastly.
- In line with security best practices, Fastly recommends using TSIG (Transaction Signature) keys for cryptographically securing zone transfers.

signature) keys for cryptographically securing zone transfers.

- DNSSEC and non-standard record types (like ALIAS or CNAME flattening) are not supported.
- Incremental zone transfers (IXFR) are not supported.
- IPv6 is supported for DNS query resolution, but not for zone transfers.

Billing

Billing for Fastly DNS is based on the number of DNS query responses provided by Fastly in units of 1 million query responses. Consider setting a [spend alert](#) for your account to automatically be notified when your month-to-date bill approaches or meets a specific dollar amount.

* * *



Fastly's Full-Site Delivery



Last updated: 2024-10-22



</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 content delivery capabilities to cache and accelerate the delivery of your [HTTP-based file content](#) 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, [adding, removing, and modifying the HTTP headers](#) 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) [settings you configure](#) 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](#).
- **Compression.** To help you speed up information transmission, we allow you to compress static content during transmission thereby making it available to your customers more quickly. To enable static or dynamic content compression, you must either [enable automatic compression](#) or [set up an advanced compression policy](#).
- **Purging.** For [dynamic or event-based content](#) that doesn't lend itself to predetermined TTL-based content expiration, you can proactively remove or invalidate your content within milliseconds with Fastly's [purging features](#). We limit purging to an average of 100K purges per hour per customer account, inclusive of all services within that account and according to any [packaged offering](#) you've purchased.

Edge logic and advanced content delivery control

Fastly's content delivery capabilities are based on a heavily extended version of the

[varnish](#) 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, [Varnish Configuration Language](#) (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

[Precision Path](#) 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.

Request routing

[Request routing](#) allows you to direct requests to a specific Fastly service based on application logic, eliminating the need for custom Compute or VCL logic or service chaining.

Origin shielding

You can designate a Fastly point of presence (POP) to [serve as a shield](#) 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 [enable it](#).

Load balancing

Services configured with multiple origin servers will automatically distribute requests to those servers evenly. You can modify this default load balancing behavior with a variety of conditions and [load balancing rules](#).

or conditions and [load balancing rules](#).

Health checks

The health of your origin servers can be monitored with [configurable health checks](#) to help ensure only responsive origin servers are being sent requests.

Fastly control panel

All Fastly accounts have access to [Fastly's control panel](#), allowing it to be [managed by multiple users](#) 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, [monitored](#), and managed through the Fastly control panel via any standard, modern web browser.

Application programming interface (API)

Fastly provides an [application programming interface](#) (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 [real-time log streaming](#) to a variety of locations, including third-party services, for storage and analysis. You can find our supported logging endpoints in our [list of streaming log guides](#). 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 [High Volume Logging](#).

Transport Layer Security

Fastly supports a variety of [Transport Layer Security \(TLS\) services](#) 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

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 [shield](#) 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 [custom VCL](#), 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 [mix and match Fastly VCL with custom VCL](#) 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 included mitigation capabilities, Fastly offers [Fastly DDoS Protection](#). For more information about this or any of our advanced services, including their subscription costs, contact sales@fastly.com.

Pricing and billing

NOTE

Billing limits for this product may be different depending on your [account type](#), if you've purchased a [packaged offering](#), or are using a [product or feature trial](#).

Full-Site Delivery [charges](#) are based on the volume of content delivered to your end users and the location of the POPs from which that content was served. [Fastly billing](#) is done in arrears based on actual usage with month-to-date usage being available via both our control panel and APIs. Consider setting a [spend alert](#) for your account to automatically

be notified when your month-to-date bill approaches or meets a specific dollar amount.

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 [integrations guides](#) for additional details.

* * *



Fastly's On-the-Fly Packager service



Last updated: 2024-01-16



</products/fastlys-onthefly-packager-service>

Fastly offers an "on-the-fly," dynamic, video-on-demand content packager 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 that requires the purchase of or use of additional professional service hours for first-time implementations. Our Professional Services team will assist with configuration and testing of this product and its features. To enable OTFP and begin this process, contact your customer success manager or email sales@fastly.com for more details.

Supported on-the-fly packager 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), High Efficiency Video Coding (H.265/HEVC), and AOMedia Video 1 (AV1) 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 [Apple FairPlay](#) 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 query 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:

- [Token-based validation](#) for decreasing response time by placing validation at the edge
- [Geolocation](#) and [device detection](#) for content targeting
- [Dictionaries](#) 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 packager 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: 2025-12-19



</products/fastlys-streaming-delivery>

Fastly's Streaming Delivery allows you to scale the delivery of your streaming content independently of any other HTTP content delivery supported by [Fastly's Full Site Delivery](#).

You can also configure and control live and video on demand (VOD) caching using full site delivery, but there are advantages to using streaming delivery, such as more favorable streaming-specific pricing and a lower traffic load on your full site delivery service. Additionally, all of the features available to full site delivery services are available to streaming delivery services.

IMPORTANT

Streaming media delivery refers to the delivery of streaming media content using supported streaming protocols for real-time consumption as the content is received, rather than after it has been downloaded. Streaming media content consists of audio and video segments, and associated content intended to enable or to be consumed with them, such as manifests and files used for captions or subtitles.

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 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 points of presence \(POPs\)](#) instead of your origin servers.

Request collapsing

If many users request the same content at the same time and that content is not cached in Fastly's POPs, your origin will have to serve that content. It doesn't, however, need to know about every individual user request made to Fastly's POPs and it would be inefficient to send the same content out many times. So, Fastly will only request the in-demand content from your origin once, essentially *collapsing* all of the user requests into a single request to your origin. Then we'll respond to each user individually.

Read more about [request collapsing](#).

Streaming miss

When Fastly needs to fetch content from your origin to serve a user request, we minimize the time until that user receives the first response (also called first-byte latency), by sending pieces of your origin's response to the user as soon as Fastly receives them, instead of first buffering the response from your origin, caching the data, and then streaming the data back to the user.

Read more about [Streaming Miss](#).

Origin shielding

You can designate a Fastly POP to serve as a shield for your origin servers, intercepting user requests on behalf of your origins to protect them from spikes in request traffic and also potentially increasing your overall cache hit rates. You can optimize this shielding geographically by configuring different shield POPs for different origin servers.

Origin shielding is not enabled by default. To use it, you must specifically enable it.

Read more about [origin shielding](#).

Real-time log streaming

To help you tune the performance of your Fastly services, we support [real-time log streaming](#) to a variety of locations, including third-party services, for storage and analysis. You can find our supported logging endpoints in our [list of streaming log guides](#). 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 [High Volume Logging](#).

Supported streaming protocols

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

- Dynamic Adaptive Streaming over HTTP (MPEG-DASH)
- High Efficiency Streaming Protocol (HESP)
- HTTP Dynamic Streaming (HDS)
- HTTP Live Streaming (HLS)
- HTTP Smooth Streaming (HSS)
- Low-Latency HTTP Live Streaming (LL-HLS)

Limitations and billing

Fastly's Streaming Delivery is a subset of Fastly's Full Site Delivery for the delivery of streaming media content. Fastly's Streaming Delivery must be configured in a separate

Fastly account, for exclusive use with streaming media content, in order to use separate billing plans and invoices as part of [calculating your bill](#).

Some Fastly products can be purchased directly in the web interface on the [Products page](#). For more details about a product, including [pricing information](#) or for help purchasing it, contact your customer success manager or email sales@fastly.com.

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 [integrations guides](#) for additional details.

* * *



HIPAA-Compliant Caching and Delivery



Last updated: 2018-08-01



</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](#).

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.

* * *



Image Optimizer



Last updated: 2025-11-18



</products/image-optimizer>

The [Fastly Image Optimizer \(Fastly IO\)](#) 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 [input and output image formats](#).



End User



Fastly Edge Cloud



Image Optimizer



Origin



Source Image

Image transformation and optimization

When an image is requested from your origin server, Fastly IO can perform [transformation tasks](#) before serving and caching the optimized version. Image transformations can be applied programmatically and through dynamic URLs in real-time. You can [make images responsive](#) 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, [special HTTP headers](#) 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

NOTE

Billing limits for this product may be different depending on your [account type](#), if you've purchased a [packaged offering](#), or are using a [product or feature trial](#).

Billing for Fastly IO is based on the number of image requests that Image Optimizer processes and delivers for you. A request becomes a billable IO request when your service configuration invokes Image Optimizer. Consider setting a [spend alert](#) for your account to automatically be notified when your month-to-date bill approaches or meets a specific dollar amount.

What triggers a billable IO request

Billing triggers differ depending on whether you are using a VCL service or a Compute service:

- **VCL services:** Any request that sets the `req.http.X-Fastly-Imageonto-Ani`

header is counted as a billable image request. Services using shielding may execute their edge logic twice—once at the edge POP and once at the shield POP. Both executions count as billable image requests.

- **Compute services:** Image Optimizer can be accessed through Compute SDKs as part of Fastly's [Beta program](#). Any request sent to Image Optimizer, including requests made through the SDKs such as `Request::with_image_optimizer`, counts as a billable image request.

How image requests are counted

Once Image Optimizer is invoked, we count image requests as follows:

- Image Optimizer requests are billed whether there is a cache HIT or MISS.
- When using animated GIF-to-video functionality, each frame delivered as video counts as a separate image request. GIF-to-GIF transformations count as one request.
- Non-image content cannot be optimized, but will still be counted as an image request if it is sent through Image Optimizer. Only send image content through IO.

Purchasing

On most accounts, anyone assigned the role of superuser can purchase this product from the Fastly control panel. If you have not been assigned that role, you can use the control panel to request that a superuser purchase it for you.

Some features, such as AVIF and JPEGXL [encoding format](#), may require an account upgrade. Existing customers can contact sales@fastly.com for details about pricing, features, and upgrade options.

This article describes a product that may use third-party cloud infrastructure to process or store content or requests for content. For more information, check out our [cloud infrastructure security and compliance program](#).

* * *



Media Shield for Live



Last updated: 2020-12-18



</products/media-shield-for-live>

Fastly Media Shield for Live offers customers the ability to decrease origin traffic by [reducing multiple CDN requests](#) 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](#) in a multi-CDN environment.

Some Fastly products can be purchased directly in the web interface on the [Products page](#). For more details about a product, including [pricing information](#) or for help purchasing it, contact your customer success manager or email sales@fastly.com.

* * *



Media Shield for VOD



Last updated: 2020-12-18



</products/media-shield-for-vod>

Fastly Media Shield for VOD offers video-on-demand customers the ability to decrease origin traffic by [reducing multiple CDN requests](#) 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 [On-the-Fly-Packaging \(OTFP\) service](#).

Media Shield for VOD allows you to take advantage of Fastly's [Observability features](#) in a multi-CDN environment.

Some Fastly products can be purchased directly in the web interface on the [Products page](#). For more details about a product, including [pricing information](#) or for

[product page](#). For more details about a product, including [pricing information](#) or for help purchasing it, contact your customer success manager or email sales@fastly.com.

* * *



Object Storage



Last updated: 2026-04-14



</products/object-storage>

Fastly Object Storage is an Amazon S3-compatible large object storage solution that works seamlessly with both Deliver and Compute services. Using Fastly Object Storage, you can store larger file sizes with Fastly, improving latency, increasing cache hit ratios, and reducing egress charges. Objects stored in Fastly Object Storage are accessible via an S3 compatible interface.

Fastly Object Storage might be useful if you want to:

- replace an existing storage solution with one that's part of Fastly's network.
- store data in a neutral location to be used across different vendors.
- access data at the edge to reduce overall costs.

Limitations and considerations

Use of Fastly Object Storage is subject to the following limitations:

- Data is stored in containers called buckets, which are limited to 100 per region. There is no limit to the amount of data storage per bucket.
- Object keys are limited to 1024 bytes.
- Objects are limited to 5 TB per object. Object metadata is limited to 1,000 bytes per object.
- No more than 5 GB of data may be included in a single upload, but multipart upload

is supported.

- The minimum size of any part of a multipart upload, except the last one, is 5 MB. The maximum number of parts in a multipart upload is 10,000.
- Uploads must complete within 120 seconds or you'll receive a `408 Request Timeout` error. One way to decrease the likelihood of timeouts is to upload objects in smaller parts, making sure to meet the requirements for multipart upload above.
- Metadata set on objects using `x-amz-meta-` headers is limited to 1000 bytes, including the length of the header name, and must not begin with `x-amz-meta-fst` or `x-amz-meta-fastly`.
- Data accessed from public, Amazon S3-compatible endpoints are subject to a maximum speed of ~150Mbps and limited to a total of 100 requests per second per bucket. Data accessed from within Fastly's network is not subject to these limits. If your needs exceed these limits, reach out to sales@fastly.com.

The following limitations and considerations also apply when using the [S3-compatible API](#) with Object Storage:

- Access keys created through the control panel grant access to Fastly Object Storage at the account-level, not the bucket-level. If you want to create access keys with bucket-level access, use the Fastly API.
- Bucket names cannot start with `fst` or `fastly`.
- Public access (i.e., unauthenticated access to buckets or objects) is not supported.
- Chunked uploads are not supported. Payloads must be transferred in a single chunk per request.
- Fastly Object Storage automatically encrypts uploaded objects. Uploaded objects have ETags that aren't an MD5 digest of their object data, which may require configuration changes to S3-compatible tooling in order to interact with the objects.

Billing

Fastly Object Storage is an add-on and is priced in addition to Fastly services. On most accounts, anyone assigned the role of superuser can purchase this product from the Fastly control panel. If you have not been assigned that role, you can use the control panel to request that a superuser purchase it for you.

Billing for Fastly Object Storage is based on a combination of total storage charges and data processing operations for the month. Storage is calculated using GB-month, rounded to the nearest hour, and processing operations are categorized into two groups, each with differing prices:

- **Class A Operations** are write operations and include `CreateMultipartUpload`, `CompleteMultipartUpload`, `PutObject`, and `UploadPart`.
- **Class B Operations** are read operations and include `GetObject` and `HeadObject`.

A full list of supported operations are available in [our documentation](#).

You are responsible for removing your data and you will be charged for storage as long as your data is present in Fastly Object Storage. Consider setting a [spend alert](#) for your account to automatically be notified when your month-to-date bill approaches or meets a specific dollar amount.

For more details about this product, contact your customer success manager or email sales@fastly.com.

* * *



Oblivious HTTP Relay



Last updated: 2023-05-31



</products/oblivious-http-relay>

The Fastly Oblivious HTTP Relay (OHTTP Relay) implements the relay portion of the [Oblivious HTTP specification](#), which allows you to create an OHTTP-compliant service using Fastly. It can be used to build double-blind privacy-enabled Fastly services that transmit requests and responses without direct knowledge of personally identifiable information linked to customers.

How it works

Oblivious HTTP is a protocol for forwarding encrypted messages via HTTP. Specifically, OHTTP facilitates the transmission of an encrypted, encapsulated message to an HTTP

endpoint from a client to a gateway through a trusted relay service, without delivering identifying information about the end user who made the request or other information that is unnecessary for request processing. Fastly's Oblivious HTTP Relay acts as that relay service.

Using Oblivious HTTP, encrypted messages are created by a client and forwarded via HTTPS to a trusted relay, in this case, Fastly's OHTTP Relay. That relay then forwards it via HTTPS to a gateway. The gateway then removes any request encryption and generates an encrypted response to the original request, forwarding it to a target without ever exposing the client originally making the request.

Fastly's OHTTP Relay product serves as the relay portion of the Oblivious HTTP transmission process. Specifically, the Fastly OHTTP Relay does the following:

- **Routes requests and responses.** The OHTTP relay routes encrypted, encapsulated messages and corresponding responses between clients and configured backends (OHTTP Gateways).
- **Performs simple request and response validation.** The OHTTP relay performs simple request and response validation, which you can specify. For example, OHTTP can confirm the message's content type, that the request was received via HTTPS, and that the request was received with a known host and path that maps to a known backend endpoint.
- **Removes non-essential request information.** The OHTTP relay strips all request headers except those that are required for the correct operation of the Fastly service or that must be passed to the OHTTP Gateway. At your request, Fastly can configure specific headers as long as they don't contain personally identifiable information.

Limitations and considerations

To maintain the privacy hygiene of messages and their corresponding responses, OHTTP Relay will not permit the following:

- **You will not be able to use the control panel or API to control your OHTTP-enabled service configuration.** After the OHTTP-enabled service is created, you must contact Fastly to make modifications to the service configuration.
- **You cannot decrypt encapsulated messages.** No visibility or introspection into the nature of the end user request is possible within Fastly's OHTTP Relay. Fastly does not have the keys to decrypt messages.

- **You will not be able to log any personally identifiable information.** No personally identifying data is available for log delivery.

Implementing OHTTP Relay

To implement Fastly's OHTTP Relay, you must contact Fastly at sales@fastly.com to begin the onboarding process. As part of that process, you will be expected to provide Fastly's Professional Services team with a frontend hostname for the relay service and a backend hostname for the gateway service through which headers will pass.

In addition, you can also request the inclusion of additional [HTTP headers](#) beyond `Content-Type`, `Content-Length`, and `Host` that should not be stripped from requests and responses during validation. If you specify additional headers, you must confirm that they will not contain personally identifiable information that can be linked to customers.

Our [Professional Services](#) staff will use this information to guide you through the onboarding process as part of the initial setup and configuration process for your Fastly service.

Once your service configuration settings are confirmed, they will be enabled for you by Fastly. You will have a Fastly account created for you and will be assigned the role of User so that you can view real-time and historical stats about your service. As a standard User, you will not be able to directly control and make changes to your OHTTP-enabled service. Requests for service configuration changes can be submitted directly to Fastly via support@fastly.com.

Billing

We bill you for OHTTP Relay based on a combination of bandwidth (per GB) and requests (per 10,000) for content delivered to clients from Fastly and then for bandwidth for traffic sent from Fastly to your customers' origin. Consider setting a [spend alert](#) for your account to automatically be notified when your month-to-date bill approaches or meets a specific dollar amount.

Some Fastly products can be purchased directly in the web interface on the [Products page](#). For more details about a product, including [pricing information](#) or for help purchasing it, contact your customer success manager or email sales@fastly.com.

* * *



Origin Connect



Last updated: 2024-12-10



</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 [Fastly shield POP](#) 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 100G or 10G port (we recommend two and will accept up to 4× 100/10G 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.

* * *



PCI-Compliant Caching and Delivery



Last updated: 2018-08-01



</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

...more, you can ensure your content is served through the Fastly network in
compliance with our PCI certification.

Contact sales-ecommerce@fastly.com 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.

* * *



Subscriber Provided Prefix



Last updated: 2024-05-06



</products/subscriber-provided-prefix>

Fastly's Subscriber Provided Prefix product allows you to bring your own IP addresses and have them announced, routed, and served by Fastly infrastructure for use with production services. When you purchase this product, you provide your own IP addresses 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 addresses by separating their network layer concerns from their content delivery concerns. By combining Fastly's Subscriber Provided Prefix service with our [Origin Connect](#) product and our [DDoS Protection and Mitigation service](#), you can protect your origin servers by

directing traffic through Fastly's global network.

Prerequisites

To purchase Fastly's Subscriber Provided Prefix service you must also purchase Fastly's [Enterprise Support](#) 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 are also 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 <https://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

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.

Working with the Subscriber Provided Prefix API

After the initial setup is completed, you can use the Subscriber Provided Prefix API to programmatically access the Subscriber Provided Prefix service. You can use the Subscriber Provided Prefix API for tasks like automating the creation and maintenance of your IP addresses or customizing the attributes on IP addresses, such as supported protocols and TLS versions. For more information, contact support@fastly.com.

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 [Fastly DDoS protection and mitigation service](#).

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.

Billing

Fees for Subscriber Provided Prefix include a set up charge and a monthly rate for two anycast prefixes configurations, one IPv4 and one IPv6 IP range. Additional prefixes and larger prefixes can be included for a higher fee. Announcements can be made either globally or for North America and the EU.

Some Fastly products can be purchased directly in the web interface on the [Products page](#). For more details about a product, including [pricing information](#) or for help purchasing it, contact your customer success manager or email sales@fastly.com.



WebSockets



Last updated: 2025-12-16



</products/websockets>

Fastly supports the use of the [WebSocket protocol](#). This protocol allows you to establish long-lived, two-way, interactive communication sessions between an end user's client (such as a web browser) and your servers.

Limitations and considerations

Keep in mind the following limitations and considerations:

- WebSockets is not compatible with [shielding](#) or the [Fastly Next-Gen WAF](#).
- When [adding a host](#) to your Fastly service configuration, only the **Name**, **Address**, **Enable TLS**, and **Override Host** origin server settings are supported with WebSockets.
- When handling a WebSocket request, `vc1_log` will run at the time the request is accepted rather than when the connection ends.
- Client request headers that are added, removed, or modified on your `Request` (or `req.http` in VCL) will be reflected in the WebSocket handoff.
- If you have TLS certificates on your origin server, they must be signed by a public certification authority. Self-signed TLS certificates are not supported.

Billing

NOTE

Billing limits for this product may be different depending on your [account type](#), if you've purchased a [packaged offering](#), or are using a [product or feature trial](#).

Use of the WebSocket protocol is disabled by default. Users assigned the superuser role in eligible accounts can enable it on the [Products](#) page. Once enabled, any user on your account with the [appropriate permissions](#) will be able to use this communications

protocol.

Enabling WebSockets on your account will also result in changes to your monthly bill. We base billing for WebSockets on a combination of bandwidth and connection time. Connection time is measured for each connection in usage minutes (rounded up to whole minutes) and aggregated monthly to millions of minutes. Bandwidth is included as part of your overall delivery bandwidth rate in your monthly billing statement. Consider setting a [spend alert](#) for your account to automatically be notified when your month-to-date bill approaches or meets a specific dollar amount.

Some Fastly products can be purchased directly in the web interface on the [Products page](#). For more details about a product, including [pricing information](#) or for help purchasing it, contact your customer success manager or email sales@fastly.com.

* * *

Category: **Security**

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



API Discovery



Last updated: 2025-09-30



</products/api-discovery>

Fastly's [API Discovery](#) provides a continuously-updating record of incoming application programming interface (API) requests proxied through Fastly's Edge network. Once enabled, API Discovery surfaces API requests passing through your Fastly Compute and Delivery services, providing visibility into your API attack surface. You can view, search, and download records of API traffic to evaluate your API landscape, monitor changes, and inform attack mitigation efforts.

Prerequisites

To use API Discovery, you must purchase a [paid account](#) with a contract for Fastly's services. Once purchased, it can be enabled in the control panel by anyone assigned the role of superuser or engineer.

Limitations and considerations

Keep in mind the following limitations and considerations:

- **API traffic and architecture limitations.** This product will only discover API traffic and will not discover non-API traffic, such as CDN media assets. It is optimized for aggregating REST APIs. While all HTTP-based API traffic will be observed by the product, the product does not explicitly support architectural patterns other than REST.
- **Data processing and normalization.** This product aggregates requests by HTTP components (domains, paths, methods) but does not include URL query string parameters or HTTP request body data in its analysis. APIs with dynamic URL components may not appear immediately as the system needs to observe multiple variations to recognize patterns.
- **Data observation limitations.** This product observes sampled network traffic, not every API call, especially sporadic calls. The interface displays discovered APIs with delays, and timestamp data is estimated based on sampling.
- **Security products note.** No security 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 does 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.

Billing

NOTE

Billing limits for this product may be different depending on your [account type](#), if

you've purchased a [packaged offering](#), or are using a [product or feature trial](#).

Fastly charges for API Discovery based on the volume of requests processed per month. These charges are separate from and do not include charges associated with the Fastly Compute or Delivery services nor with usage of the Fastly Next-Gen WAF. When Fastly DDoS protection is enabled on the service, requests associated with mitigated attacks are excluded from billing. Consider setting a [spend alert](#) for your account to automatically be notified when your month-to-date bill approaches or meets a specific dollar amount.

Some Fastly products can be purchased directly in the web interface on the [Products page](#). For more details about a product, including [pricing information](#) or for help purchasing it, contact your customer success manager or email sales@fastly.com.

* * *



Bot Management



Last updated: 2026-03-31



</products/bot-management>

Fastly's [Bot Management](#) product provides you with visibility into bot traffic, allowing you to identify bots and respond to automated traffic with the nuance your business and security posture requires. Using the knowledge you gain from this detection, you can enforce rulesets and policies to control bots in both the delivery (before cache) or the Next-Gen WAF (after cache) as part of your web asset and application protection measures. Because not all bots are malicious, Bot Management offers controls that can help you decrease unwanted bot activity by allowing you to customize your interactions and automatically decide which bots to allow in your ecosystem.

ContentGuard

ContentGuard provides bot detection and mitigation at the network edge, protecting your traffic before bots can access content in either Fastly cache or your origin. This feature leverages the bot detection engine to perform inspection at the network edge and is the

leverages the bot detection engine to perform inspection at the network edge and is the first line of defense against automated threats. It helps you detect bots that scrape content and may attempt to steal valuable proprietary data like dynamic pricing, real-time inventory, or copyrighted material. The following Bot Management product features are available with ContentGuard:

- Client fingerprinting
- Verified bots
- Block or allow requests based on detection attributes via VCL

Client fingerprinting

Client fingerprinting incorporates [JA3](#) and [JA4](#) fingerprinting and allows you to identify client types as long as that information is available as part of the TLS encrypted communication between a specific client and its server. This feature can help you detect bots designed for malicious activities such as credential stuffing, credential cracking, or IP rotation attacks.

Client challenges

[Client challenges](#) allow you to require users to prove that they are human or that a connection is happening via a full-fledged browser. For each service, you choose whether these challenges are dynamic, interactive, or non-interactive:

- Dynamic challenges allow Fastly to automatically choose the most appropriate client challenge based on the situation, including Private Access Tokens (PATs), non-interactive challenges, and interactive challenges if suspicious activity is detected during the initial check.
- Interactive challenges use configurable CAPTCHA-like challenge-response tests that human users must respond to.
- Non-interactive challenges use JavaScript Proof-of-Work (PoW) to test that the browser supports JavaScript.

To identify when challenges have been initiated and solved, cookies are issued from the customer domain in which the challenges are issued. Specifically:

- the `_fs_ch_st_<RANDOM STRING>` challenge start cookie signals the initiation of the challenge and helps mitigate trivial replay of challenge flows by your service

- the `_fs_ch_cp_<RANDOM_STRING>` challenge complete cookie signals the completion of the challenge and communicates to your service that access to a resource should be permitted

Advanced client-side detections

[Advanced client-side detections](#) allow you to detect bots that leverage headless browsers such as headless Chrome. This feature requires you to modify the HTML code of your website to include a JavaScript snippet. To identify that a browser has run the JavaScript, the `_fs_cd_cp_<RANDOM STRING>` cookie is issued from the customer domain.

Verified bots

Verified bots allow you to add a Next-Gen WAF signal to the logic of your active configuration rules that will help validate self-identified bots and thereby allow or block them as appropriate as requests arrive to the web applications you're protecting.

Organizations wishing to have their bots included in one of the Verified Bot [categories](#) can submit details about the bot using our [submission form](#). Use your business email address to set up your account on community.fastly.com to aid verification. Submissions will be reviewed by Fastly and considered for inclusion in the product.

Private Access Tokens

IMPORTANT

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

The Private Access Token (PAT) verification service allows you to protect access to resources on your origin. When an end user requests information from an origin that Fastly protects, the PATs service specifically requires the requestor to prove that they are human and verify their identity, but does so without directly revealing personal information about the requester or requiring them to solve puzzle-based challenges. It does this based on the settings you specify in your Bot Management implementation and then responds to requests by issuing a validation token granting access or blocking access to those protected resources as appropriate.

Prerequisites

To purchase Bot Management, you must have a [paid account](#) with a contract for Fastly's services and must purchase a Fastly delivery product. For those who are interested in our [Security Subscription packages](#), Bot Management is available in all feature tiers except the Essentials platform.

Limitations and considerations

The following features require Next-Gen WAF to function and can only be done post-cache. They are not available within ContentGuard:

- Client challenges
- Advanced client-side detections

Keep in mind the following limitations and considerations for client challenges:

- The client challenges feature must be [enabled for each individual service](#) using your service ID via an API call.
- Client challenges are issued to fully-qualified domain names (FQDN). If your service includes subdomains that shouldn't receive challenges (e.g., `api.example.com`), be sure to restrict the challenge to the desired subdomains when creating the request rule that adds the challenge.
- Exceptions to client challenges can be used to allow some bots access to your site. These exceptions must be explicitly included in any rule that would otherwise exclude them.

In addition, keep in mind the following limitations and considerations specific to PATs:

- PATs usage is measured based on the number of token redemptions that occur. One token redemption is equal to one request, which affects your billing as described below.
- PATs challenges can only be issued to Apple-supported devices using iOS 16 or higher or macOS Ventura or higher.
- Apple-supported devices are limited to 10 tokens per minute, per device, and only 10 tokens per every 5 minutes are allowed for a single origin server or website. Only 1 token per minute is allowed for a single TLS connection to a server.

Billing

NOTE

Billing limits for this product may be different depending on your [account type](#), if you've purchased a [packaged offering](#), or are using a [product or feature trial](#).

Fastly charges for Bot Management based on the volume of requests (per millions) processed per month. These charges are separate from and do not include charges associated with the Fastly Full Site Delivery service nor with usage of the Fastly Next-Gen WAF. Consider setting a [spend alert](#) for your account to automatically be notified when your month-to-date bill approaches or meets a specific dollar amount.

WARNING

Enabling ContentGuard may increase your Bot Management bill due to the additional requests processed.

Some Fastly products can be purchased directly in the web interface on the [Products page](#). For more details about a product, including [pricing information](#) or for help purchasing it, contact your customer success manager or email sales@fastly.com.

* * *

	Certainly
	Last updated: 2023-08-16
	/products/certainly

Certainly is Fastly's publicly-trusted certification authority (CA) that generates Transport Layer Security (TLS) certificates to provide trusted identification of secured websites. Certainly is only available for use with [Fastly TLS](#).

TLS provides security and legitimacy to your website or application by serving traffic over HTTPS, something users expect to ensure their web activity is secure. For TLS to work, domains need to be validated and approved by a trusted CA that issues certificates to

domains need to be validated and approved by a trusted CA that issues certificates to web servers, thus confirming ownership of a site. In turn, those certificates provide the credentials used by TLS and other security protocols that allow servers to prove their authenticity and to establish encrypted, private sessions with web clients without requiring a pre-existing relationship. Certainly is designed to help simplify, automate, and scale the management of TLS certificates on Fastly.

When using Fastly managed certificates in Fastly TLS, you can generate certificates with Certainly instead of a third-party certification authority. Certainly issues certificates that are valid for 30 days. Fastly will attempt to re-verify your domain and renew your certificate after 20 days as long as your DNS records point to Fastly and your Certification Authority Authorization (CAA), if in use, is set to Certainly. Having certificates re-verified and renewed in a shorter time period provides extra security by forcing the regular re-validation of ownership and rotation of the TLS certificates and asymmetric keys.

NOTE

Customers using Fastly's Platform TLS service are currently not able to use Certainly certificates with that product and must continue to manage their own certificates.

Some Fastly products can be purchased directly in the web interface on the [Products page](#). For more details about a product, including [pricing information](#) or for help purchasing it, contact your customer success manager or email sales@fastly.com.

* * *



Edge Rate Limiting



Last updated: 2025-11-17



</products/edge-rate-limiting>

Fastly's [Edge Rate Limiting](#) provides customers with the ability to count client requests and optionally penalize clients for exceeding set rate limits, thereby controlling the rate of

requests sent to origin servers. Common uses for the Edge Rate Limiting product include mitigating abusive use of a website or service (e.g. by a scraping bot or a denial of service attack) or applying limits on use of an expensive or billable resource (e.g. allowing only up to 1000 requests an hour to an API endpoint). By controlling the rate of requests to your origins, you can help ensure service availability during excessive spikes in traffic.

Limitations and caveats

Edge Rate Limiting is compatible with Fastly's [origin shield](#) feature and both can be used together. If you have shielding enabled, rate limits will be counted twice, once at the edge and once at the origin shield. This has different implications for where protection is occurring and how the client is identified.

Edge Rate Limiting is not intended to compute rates with high precision and may under count by up to 10%. For example, if you have a rate limit of 100 requests per second over a 10 second window, when the real request rate reaches 100 RPS, it may register as low as 90 and therefore may not trigger the limit until the real request rate reaches 110 RPS.

Both rate counters and penalty boxes have a fixed capacity for client entries. Once a rate counter is full, each new entry evicts the entry that was least recently incremented. Once a penalty box is full, each new entry will evict the entry with the shortest remaining time to live (TTL). Penalty box TTLs are enforced by rounding up on the minute, so the effective minimum TTL of an entry in a penalty box is 2 minutes.

Security products note

No security 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 does 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.

Billing

NOTE

Billing limits for this product may be different depending on your [account type](#), if

you've purchased a [packaged offering](#), or are using a [product or feature trial](#).

Fastly charges for Edge Rate Limiting based on the number of rate counters and penalty boxes associated with your account each month. These charges are separate from and do not include charges associated with Fastly Compute, Delivery, or other Fastly services. Consider setting a [spend alert](#) for your account to automatically be notified when your month-to-date bill approaches or meets a specific dollar amount.

Some Fastly products can be purchased directly in the web interface on the [Products page](#). For more details about a product, including [pricing information](#) or for help purchasing it, contact your customer success manager or email sales@fastly.com.

* * *



Fastly Client-Side Protection



Last updated: 2025-04-10



</products/fastly-client-side-protection>

Fastly [Client-Side Protection](#) provides you with the ability to inventory and control the resources (e.g., scripts, images, and fonts) that load on an end user's browser from defined areas of your web applications by building and enforcing [content security policies](#). When a resource violates your content security policy, the end user's browser blocks or logs the resource per your selected protection mode. Based on policy violation reports, you can adjust your content security policies as needed. In addition, you can provide a justification as to why each client-side script is or isn't allowed. These capabilities help you guard against cross-site scripting attacks (e.g., Magecart attack) and enable you to maintain compliance with Payment Card Industry Data Security Standard (PCI DSS) [4.0.1 - Sections 6.4.3 and 11.6.1](#).

Prerequisites

To purchase Fastly Client-Side Protection, you must purchase Fastly's [Next-Gen WAF](#).

Limitations and considerations

Keep in mind the following limitations and considerations for Fastly Client-Side Protection:

- Fastly Client-Side Protection is reliant on [browser support](#) for `HTTP Content-Security-Policy` response headers. Older browsers may not support all features.
- The Next-Gen WAF inserts the `HTTP Content-Security-Policy-Report-Only` response header into a sample of responses. This header triggers the inventory process for Fastly Client-Side Protection.
- Depending on whether Client-Side Protection is in blocking or logging mode, the Next-Gen WAF adds either the `HTTP Content-Security-Policy` or the `HTTP Content-Security-Policy-Report-Only` response header to all responses that pass through the WAF. These headers deliver your content security policy.
- If your web application has a broken or insecure connection or certificate, the end user's browser will not forward policy violation reports to Fastly.
- When an object evaluated by the Next-Gen WAF is cached, the content security policy attached to the object is also cached. Both the object and content security policy are served together for as long as the object remains in the cache. If you update the content security policy, cached objects won't reflect the updated content security policy until the object is removed from cache and passes through the WAF again.
- Even with SHA256 hashing, inline scripts can pose a security risk. While compliance with PCI DSS is met, this doesn't guarantee full protection.
- We store policy violation reports and scripts that were observed during the inventory process for 90 days. When we fail to re-inventory a script for 90 days, we stop storing the script.
- Fastly Client-Side Protection can only be accessed using the [Fastly control panel](#).
- Fastly Client-Side Protection uses [Manifest v3](#).

Billing

NOTE

Billing limits for this product may be different depending on your [account type](#) if

Billing limits for this product may be different depending on your [account type](#), if you've purchased a [packaged offering](#), or are using a [product or feature trial](#).

Fastly Client-Side Protection is billed based on the number of [Pages](#) that are activated under the [Websites page](#) for your account each month. As indicated on your service order, your account includes a set number of activated Pages and each additional page incurs a charge.

For example, say you've purchased five Pages as part of your services and that you have five websites on which you wish to use those pages. Each website requires one Page and therefore consumes one Page each from the five you have purchased. During the month, the first five Pages you activate incur no additional charges on your account. Any additional Pages activated throughout the month, however, incur additional month-to-month fees. If an additional Page is activated but deleted before the months' end, it incurs charges only during the month activated and not subsequent months.

Security products note

No security 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 does 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 DDoS Protection



Last updated: 2025-11-20



</products/fastly-ddos-protection>

Fastly [DDoS Protection](#) provides real-time visibility into and defense against common Distributed Denial of Service (DDoS) attacks directed at your applications, APIs, and origin servers. It provides automatic detection and mitigation at Fastly's network edge.

away from your infrastructure.

Prerequisites

To use Fastly DDoS Protection, you must have a [paid account](#) for [Full-Site Delivery](#), [Fastly Streaming Delivery](#), or [Compute](#).

On most accounts, anyone assigned the role of superuser can purchase this product from the Fastly control panel. If you have not been assigned that role, you can use the control panel to request that a superuser purchase it for you.

Once purchased, superusers or engineers on your account can [enable Fastly DDoS Protection](#) in the Fastly control panel to immediately establish a defensive posture against cyber threats.

How it works

Fastly DDoS Protection uses techniques like our proprietary Adaptive Threat Engine (formerly referred to as "Attribute Unmasking") to rapidly fingerprint complex attack traffic and distinguish it from normal, organic traffic to your applications and origin servers. This allows you to continuously evaluate and automatically adapt and respond to DDoS attacks without manual intervention or maintenance.

Once enabled, Fastly's network will dynamically and proactively process, analyze, and diagnose your traffic. Metrics about each service's request traffic are logged, including details about attacks detected and mitigated, and appear in a dedicated dashboard accessible through the Fastly control panel.

Fastly DDoS Protection automatically generates event rules when it detects an attack. Event rules define the conditions and attributes used to filter or mitigate traffic during a DDoS attack. You can modify the behavior of individual event rules to precisely control which traffic is logged or blocked by Fastly DDoS Protection for a specific attack.

Limitations

No security 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 does 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 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.

Billing

NOTE

Billing limits for this product may be different depending on your [account type](#), if you've purchased a [packaged offering](#), or are using a [product or feature trial](#).

Fastly DDoS Protection automatically detects attack traffic and excludes this traffic from billing-related metering. You are never billed for mitigated attacks. We bill you based on non-attack traffic that passes through Fastly's network, which we define as requests that are not identified as malicious by the product. Metered customers are not billed for the first 500,000 non-attack traffic requests each month.

Some Fastly products can be purchased directly in the web interface on the [Products page](#). For more details about a product, including [pricing information](#) or for help purchasing it, contact your customer success manager or email sales@fastly.com.

* * *



Fastly Next-Gen WAF



Last updated: 2026-06-17



</products/fastly-next-gen-waf>

The [Fastly Next-Gen WAF](#) 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, block it, or return a [deceptive response](#). 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.

Feature availability

Feature availability and limitations depend on the Security Core, Security Core Plus, and Security Total [packaged security offering](#) you may have purchased, as well as limitations specifically mentioned on your service order. The Core, Core Plus, and Total columns indicate what's included in each package.

Feature	Core	Core Plus	Total
Deployment Types	Edge On-Prem	Edge On-Prem	Edge On- Prem
API Discovery	Available for purchase	Available for purchase	Included
Bot Management	Available for purchase	Available for purchase	Included
Client-side Protection	Available for purchase	Included	Included
DDoS Protection	Available for purchase	Included	Included
Default attack signals	Included	Included	Included
Default anomaly signals	Included	Included	Included
Signal Sciences Cloud WAF			

Standard API & AIO signals	Included	Included	Included
Custom signals	Included	Included	Included
Lists	Included	Included	Included
Virtual patching	Included (BLOCK only)	Included	Included
Custom rules *	Available for purchase	Included	Included
Signal exclusion rules	Included	Included	Included
Rule action types	Allow and Block only	All	All
Attack thresholds	Included	Included	Included
Edge Rate Limiting **	Available for purchase	Available for purchase	Included
Custom response codes	Not available	Included	Included
Default dashboards	Included	Included	Included

* Custom rules include [request rules](#), [advanced rate limiting rules](#), and [signal thresholds \(site alerts\)](#).

** Requires an active [Full-site Delivery](#) or [Compute](#) account.

Feature availability and limitations depend on the platform and, if applicable, the Security Starter, Security Advantage, or Security Ultimate [packaged security offering](#) you may have purchased. The Essential, Professional, and Premier platform columns indicate what's included in each package.

Feature	Essential	Professional	Premier
Deployment Types	Edge On-Prem Cloud	Edge On-Prem Cloud	Edge On-Prem Cloud
API Discovery	Not available	Available for purchase	Included

Bot Management	Not available	Available for purchase	Available for purchase
Client-side Protection	Not available	Available for purchase	Available for purchase
DDoS Protection	Not available	Available for purchase	Available for purchase
Standard API & ATO signals	Not available	Included	Included
Custom signals	Not available	Included	Included
Lists	Not available	Included	Included
Virtual patching	Included (BLOCK only)	Included	Included
Request rules	Included	Included	Included
Signal exclusion rules	Included	Included	Included
Templated rules	Not available	Included	Included
Rule action types	Allow and Block only	All except Deception	All
Attack thresholds	Included	Included	Included
Site alerts (signal thresholds)	Included (attack signals only)	Included	Included
Edge Rate Limiting	Not available	Included*	Included*
Advanced Rate Limiting	Not available	Not available	Included
Custom response codes	Not available	Included	Included
Default dashboards	Included	Included	Included

* Requires an active [Full-site Delivery](#) or [Compute](#) account.

Documentation

Documentation for the Next-Gen WAF can be found at www.fastly.com/documentation/guides/next-gen-waf. We [announce the most recent changes and updates](#) for the agents and modules in our changelog.

Control panel access

The Next-Gen WAF can be accessed via either the [Next-Gen WAF control panel](#) or the [Fastly control panel](#). Each control panel allows you 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 control panel to create sites (also known as workspaces). A site (workspace) is a user-defined set of rules and settings for applications and origin servers. Each control panel allows you to create multiple sites (workspaces) to differentiate between one or more APIs, microservices, or web applications. For each site (workspace), you can use the control panels to add rules for requests, configure thresholds, and add integrations to other systems.

Deployment types

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

- **On Fastly's Edge platform (Edge WAF).** To use the Edge WAF deployment method with VCL or Compute services, you must add it to new or existing Fastly services that you create in the Fastly control panel and update your DNS records to point to Fastly.
- **Directly on your web servers within your infrastructure (On-Prem WAF).** The On-Prem WAF (formerly known as Core WAF) deployment method consists of two components, the module and the agent. The module can exist as a [plugin to your web server](#) or as a [language or framework-specific implementation](#). The agent is a small process that provides an interface between your web server and our cloud engine. You can also use this deployment method without a module by running the agent in [reverse proxy mode](#).
- **On Fastly's cloud-hosted infrastructure (Cloud WAF).** To use Cloud WAF, you must upload a TLS certificate, add an origin server using the Next-Gen WAF control panel, and update your DNS records to point to the appropriate servers.

The Next-Gen WAF control panel supports all features of all deployment types. The Fastly control panel supports the features of the Edge WAF and On-Prem WAF deployment

control panel supports the features of the Edge WAF and On-Prem WAF deployment types only.

Threat intelligence

As part of Next-Gen WAF, we may [aggregate the attack data collected](#) 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 control panels and alerts you to potential bad actors that have been identified elsewhere in the subscriber network.

API

The [Signal Sciences Application Programming Interface](#) allows you to integrate your applications and services with the Next-Gen WAF via the Next-Gen WAF control panel. The [Fastly Security Application Programming Interface](#) allows you to integrate your applications and services with the Next-Gen WAF via the Fastly control panel. Each uses standard HTTP response codes and verbs to allow you to programmatically control all the same features that are available with the [control panels](#). Each API provides a variety of endpoints that we document in our API reference documentation.

Control over data sharing

Next-Gen WAF gives you control over data shared with Fastly. The hosted Cloud WAF deployment 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 control panel 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 On-Prem WAF 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 WAF deployments receive access to a [combination of features](#) 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.

In addition to these included detection and mitigation capabilities, Fastly offers [Fastly DDoS Protection](#). For more information about this or any of our advanced services, including their subscription costs, contact sales@fastly.com.

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 [documentation](#) and described in the [release notes](#). You can also [subscribe to receive emails from us](#) when updates are released or subscribe to our integrations with third-party tools (e.g., [Slack](#) or [Microsoft Teams](#)). For On-Prem WAF 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 control panel 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.
- **WebSocket traffic inspection.** Next-Gen WAF can only inspect WebSocket traffic when it is deployed using the Core WAF deployment method. Edge WAF and Cloud WAF deployments don't support WebSocket traffic inspection.
- **Security products note.** No security 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 does 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.

This article describes a product that may use third-party cloud infrastructure to process or store content or requests for content. For more information, check out our [cloud infrastructure security and compliance program](#).

Billing

NOTE

Billing limits for this product may be different depending on your [account type](#), if you've purchased a [packaged offering](#), or are using a [product or feature trial](#).

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 sites (workspaces) and either the average requests per second (RPS) or average requests per month (RPM) processed by Next-Gen WAF as appropriate for your packaged entitlements. Consider setting a [spend alert](#) for your account to automatically be notified when your month-to-date bill approaches or meets a specific dollar amount.

Any time you purchase a [deployment option](#) for the first time, your service order will include a one-time purchase of [Implementation Services](#) to assist you with your onboarding experience.

Edge WAF deployments are additionally billed for delivery charges associated with the [Full-Site Delivery service](#) on which those deployments are hosted. [Prices](#) are based on the volume of content delivered to your end users and the location of the POPs from which that content was served. Fastly billing is done in arrears based on [actual usage](#) with month-to-date usage being available via both our control panel and APIs.

Cloud WAF deployments are additionally billed for the overall traffic flowing through the hosted services in terabytes (TBs) and the number and location of protected origins.

Some Fastly products can be purchased directly in the web interface on the [Products page](#). For more details about a product, including [pricing information](#) or for help purchasing it, contact your customer success manager or email sales@fastly.com.

* * *



Platform TLS



Last updated: 2021-10-05



</products/platform-tls>

IMPORTANT

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

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

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](#).

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.

Some Fastly products can be purchased directly in the web interface on the [Products page](#). For more details about a product, including [pricing information](#) or for help purchasing it, contact your customer success manager or email sales@fastly.com.

* * *



TLS service options



Last updated: 2026-06-26



</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.

Fastly TLS allows you to manage TLS certificates on a domain-by-domain or multi-domain basis using our control panel or API. To serve secure HTTPS traffic from Fastly, your website needs a valid TLS certificate with a matching private key. You can either [instruct Fastly to automatically generate and manage TLS keys and certificates](#) on your behalf or generate and [upload your own TLS certificates](#) and private keys.

TIP

Fastly's [pricing page](#) details the current rates for our TLS services.

Important considerations

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. Fastly may revoke certificates if required by the [CA/Browser Forum Baseline Requirements](#) or the certification authority (CA) providing the certificate. We may also revoke certificates if you fail to comply with Fastly's [Acceptable Use Policy](#).

revoke certificates if you fail to comply with Fastly's [Acceptable Use Policy](#).

Certificates provided by Let's Encrypt or GlobalSign are third-party technologies. Certificates provided by GlobalSign are subject to the terms of GlobalSign's Subscriber Agreement, which can be found at <https://www.globalsign.com/en/repository>. An alternative to these third-party technologies is [Certainly](#), Fastly's publicly-trusted CA.

For customers bringing their own certificates, both Fastly TLS and Concierge TLS service support Domain Validated (DV), Organization Validated (OV), Extended Validation (EV), RSA, and ECDSA certificates. If Fastly manages your certificates, use Certainly or Let's Encrypt to issue DV certificates or GlobalSign to issue DV or OV certificates.

If you've 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.

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 Server Name Indication (SNI) extension of TLS that allows clients to present a hostname in the TLS handshake request. 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.

TIP

If you need TLS support for non-SNI clients, see [Dedicated IP addresses](#).

Fastly enforces a match between each HTTP request's Host header and a Subject Alternative Name (SAN) field on the associated TLS certificate to prevent use of a technique known as [domain fronting](#). If you encounter an HTTP 421 error from Fastly, a [mismatch](#) is the likely cause.

Fastly supports SHA-256 certificates signed by publicly trusted certification authorities that have a key size of 2048 bits for RSA public key encryption and key sizes of 256 bits and 384 bits for ECDSA public key encryption. For performance reasons and to help mitigate your security costs, we strongly recommend using an ECDSA certificate.

Fastly TLS

With Fastly TLS, you can either [instruct Fastly to automatically generate and manage TLS certificates](#) on your behalf or generate and [upload your own TLS certificates](#) and private keys.

Fastly-managed TLS subscriptions (managed TLS)

Fastly-managed TLS subscriptions (managed TLS)

Fastly-managed certificates are an option for both paid accounts and free accounts. When Fastly manages your certificates, you use the Fastly control panel 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.

Fastly-managed certificates can be procured from Certainly, a non-profit CA (Let's Encrypt), or a commercial CA (GlobalSign).

TIP

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

Self-managed certificates (Bring Your Own Certificates - BYOC)

If you have a paid account for Fastly's services, you can [bring your own certificates](#) and use the Fastly control panel or API to upload TLS certificates and keys. You must ensure you upload the relevant private key first before uploading the matching certificate.

Billing

NOTE

Billing limits for this product may be different depending on your [account type](#), if you've purchased a [packaged offering](#), or are using a [product or feature trial](#).

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 in an activated state at the end of the month for your account. All domains in an activated state at the end of the month count towards your bill, 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

Specifically, there are pricing differences between Fastly TLS certificates provided by a commercial CA and those provided by a non-profit CA or Certainly. Our [pricing page](#) provides specifics about these differences.

All accounts include the following for free:

- Up to five (5) total Fastly-managed Let's Encrypt subscriptions and/or self-managed certificates (BYOC).
- Up to five (5) Fastly-managed Certainly subscriptions.

Additional subscriptions or certificates, as well as GlobalSign CA certificates, can be purchased for an additional fee. Refer to our [pricing page](#) or contact sales@fastly.com for more information.

Mutual TLS authentication

[Mutual TLS \(mTLS\)](#) is an additional layer of network connection security that is added on top of our existing TLS product. By default, the TLS protocol only requires a server to present a trusted certificate to the client. mTLS requires the client to also present a trusted certificate to the server. Instead of having to rely on traditional authentication methods like passwords or API keys, the server to client connection is secured using TLS certificates.

Billing

All accounts can protect two domains for free with mTLS. Additional domains can be purchased for a flat fee. Refer to our [pricing page](#) or contact sales@fastly.com for more information.

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.

Dedicated IP addresses

Fastly can install customer-provided or Fastly-managed certificates at a [dedicated set of IP addresses](#) 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.

* * *

Category: **Observability**

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



Domain Inspector



Last updated: 2026-02-23



</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 **fully qualified domain name** without requiring you to send log data to a **third-party data collector**.

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 control panel or as a JSON formatted data stream.

IMPORTANT

This product is not available for Next-Gen WAF.

Limitations and considerations

This product has the following limitations and considerations:

We limit you to 250 domains per service with Domain Inspector. Reach out to

- we limit you to 250 domains per service with Domain Inspector. Reach out to support@fastly.com for details on how to increase this limit.
- The data retention period for this product is 45 days if you purchased a [package offering](#).
- To help control data collection and aggregation, [wildcard domains](#) are aggregated against a single `*.example.com` entry.
- Domain Inspector is not supported on services that use [IP-to-service pinning](#).
- Domain Inspector does not support domains that are longer than 64 characters.
- Data related to edge hit ratio and origin offload are not available for Compute services.

Billing

Fees for Domain Inspector are charged based on the number of unique domains sending traffic through Fastly. Consider setting a [spend alert](#) for your account to automatically be notified when your month-to-date bill approaches or meets a specific dollar amount.

Some Fastly products can be purchased directly in the web interface on the [Products page](#). For more details about a product, including [pricing information](#) or for help purchasing it, contact your customer success manager or email sales@fastly.com.

* * *



High Volume Logging



Last updated: 2021-03-09



</products/high-volume-logging>

Fastly's [real-time log streaming](#) 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

service. For customers who need to increase this limit, Fastly offers High Volume Logging.

Billing

High Volume Logging is billed based on the cumulative log statements streamed in excess of the Fastly's Full-Site Delivery [Real-Time Log Streaming](#) 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. Consider setting a [spend alert](#) for your account to automatically be notified when your month-to-date bill approaches or meets a specific dollar amount.

Some Fastly products can be purchased directly in the web interface on the [Products page](#). For more details about a product, including [pricing information](#) or for help purchasing it, contact your customer success manager or email sales@fastly.com.

* * *



Log Explorer & Insights



Last updated: 2025-12-18



</products/log-explorer-and-insights>

The Log Explorer & Insights feature allows you to proactively store, inspect, and monitor your log data on the Fastly Platform. We collect, store, and analyze request and response-related data (e.g., path and status code) to provide aggregated metrics and log data.

The [Insights dashboard](#) provides a variety of views of your logging data that can help you visualize and identify trends. To facilitate troubleshooting, [Log Explorer](#) allows you to view, filter, and analyze request logs using the Fastly control panel and API. By monitoring and inspecting your logging data, you can optimize the performance of your services and troubleshoot and debug related issues.

Prerequisites

Prerequisites

To use Log Explorer & Insights, you must purchase a [paid account](#) and an [Observability package](#). Services using custom VCL are required to have the Fastly macros included as detailed in the [custom VCL documentation](#).

Limitations and considerations

This feature has the following limitations and considerations:

- Data related to hits, misses, cache hit ratio, location data, autonomous system numbers (ASNs), operating system, browser, and device are not available for Compute services.
- The data retention period for this feature is 7 days.
- All data is stored in the United States on Fastly's own infrastructure or via our [sub-processors](#) for a maximum of 7 days.
- Data log formats and their views cannot be customized.
- Alerts cannot be used with Log Explorer & Insights metrics.
- If [segmented caching](#) is enabled on a service, not all requests generated will be logged and visible on the [Insights page](#) or [Log Explorer](#). Only requests where `segmented_caching.is_inner_req` is `false` will be logged.

To support performance at scale, the Log Explorer & Insights feature samples data at a set rate determined by the requests per second (RPS) of your service. The dataset will be large enough to provide statistically significant insights.

Billing

Log Explorer & Insights is included as part of Fastly's Starter, Advantage, or Ultimate [Observability packages](#). The limits on the number of services associated with this feature depends on the Observability package you purchased. For details, refer to the [packaged offering entitlements](#). Once purchased, it can be enabled in the control panel by anyone assigned the role of superuser or engineer.

This article describes a product that may use third-party cloud infrastructure to process or store content or requests for content. For more information, check out our [cloud infrastructure security and compliance program](#).

* * *

Logging Insights Package



Last updated: 2018-06-15



</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 [logging endpoint](#)
- 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 [Sumo Logic App for Fastly](#). You'll need a Sumo Logic account with the appropriate license, and you'll need to enable the [Sumo Logic logging endpoint](#).

Some Fastly products can be purchased directly in the web interface on the [Products page](#). For more details about a product, including [pricing information](#) or for help purchasing it, contact your customer success manager or email sales@fastly.com.

* * *



Observability features



Last updated: 2026-06-18



</products/observability-features>

As part of our [Observability product offerings](#), Fastly provides you with a collection of features that allow you to continuously monitor the status of your website, product, or service using real-time and historical metrics.

- **Alerts.** The [Alerts feature](#) provides visibility into service health conditions via automated notifications for service-related performance metrics. You can set up

alerts via the web interface or the API to receive notifications to various integrations when a metric you define either goes above or remains below a threshold.

- **Domain Inspector.** Our [Domain Inspector interface](#) feature (available for Delivery only) offers a dashboard designed to provide you with real-time and historic visibility into response data for traffic from your domains and subdomains to a Fastly Deliver service. It allows you to monitor traffic for a [fully qualified domain name](#) without requiring you to send log data to a [third-party data collector](#).
- **Edge Observer.** Our [Edge Observer interface](#) serves as the entry point of all our Observability features. It provides you with customizable views that give insights into key metrics that can inform your decisions and goals.
- **Fastly Mobile App (Beta).** Available as an iOS and Android app providing remote access to real-time platform diagnostics, service monitoring, and operational alerting. It enables users to track account-wide metrics, access Observability tools (including Domain and Origin Inspectors, Custom Dashboards, and Live Logs), view global Fastly network status, audit deployed services and account activity, and configure customized push notifications for active incidents and maintenance events.
- **Log Explorer & Insights.** The [Log Explorer & Insights feature](#) offers a dashboard designed to provide you with visibility into edge request performance. It allows you to monitor and analyze request traffic, detect error patterns, and help identify performance issues with your web properties.
- **Origin Inspector.** Our [Origin Inspector interface](#) offers a dashboard designed to provide you with 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.

Entitlements and feature availability may be different if you've purchased a [package offering](#) or are using a [product or feature trial](#).

Some Fastly products can be purchased directly in the web interface on the [Products page](#). For more details about a product, including [pricing information](#) or for help purchasing it, contact your customer success manager or email sales@fastly.com.



Origin Inspector



Last updated: 2024-01-18



</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 [visualizations in the control panel](#) or as a [JSON-formatted data stream](#).

Prerequisites

Origin Inspector can be purchased as part of either Fastly's Essential or Professional [Edge Cloud packages](#).

Customers with a [paid account with a contract](#) 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 sales@fastly.com to discuss this option.

Limitations and considerations

We limit you to 250 origins per service with Origin Inspector. Reach out to support@fastly.com for details on how to increase this limit.

In addition, 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

The data retention period for this product is 45 days if you purchased a [package offering](#).

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 per service. Consider setting a [spend alert](#) for your account to automatically be notified when your month-to-date bill approaches or meets a specific dollar amount. To discuss pricing in more detail, contact sales@fastly.com.

* * *

Category: Compute

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



AI Accelerator



Last updated: 2024-12-10



</products/ai-accelerator>

[AI Accelerator](#) is a semantic caching solution for large language model (LLM) APIs used in generative artificial intelligence (AI) applications.

AI accelerator caches LLM queries based on semantic similarity, allowing for flexible caching of query results and faster response times to end users. AI Accelerator is compatible with multiple LLM products performing pass-through caching. Your Fastly credentials allow you to call supported LLM APIs, caching queries automatically. AI Accelerator passes through any responses, errors, or rate limits associated with your LLM provider.

Fastly AI Accelerator can help:

- Reduce user wait times between query and response
- Reduce fees by making fewer direct queries to LLM services

Prerequisites

To use AI Accelerator, you must have a contract for Fastly services or a valid credit card on file.

On most accounts, anyone assigned the role of superuser can purchase this product from the Fastly control panel. If you have not been assigned that role, you can use the control panel to request that a superuser purchase it for you.

Limitations and considerations

AI Accelerator has been designed to be vendor neutral, so it works with many different LLM providers, but you must have valid credentials for at least one supported LLM service. [A list of supported LLM providers is available here](#). LLM providers are third-party technologies subject to their own terms and conditions for which you are responsible.

Billing

AI Accelerator is an add-on and is priced in addition to Fastly services. Billing for Fastly AI Accelerator is based on the total number of incoming monthly requests to AI Accelerator. Consider setting a [spend alert](#) for your account to automatically be notified when your month-to-date bill approaches or meets a specific dollar amount.

You are responsible for any fees associated with the LLMs you choose. These fees will not be a part of your Fastly bill.

Some Fastly products can be purchased directly in the web interface on the [Products page](#). For more details about a product, including [pricing information](#) or for help purchasing it, contact your customer success manager or email sales@fastly.com.



Cache APIs



Last updated: 2025-02-05



</products/cache-apis>

The following APIs enable developers to build apps on Compute using Fastly's global cache network.

Core Cache API

The [Core Cache API](#) uses languages like [Rust](#), [JavaScript](#), and [Go](#) to expose the API primitives used to implement cache applications.

HTTP Cache API

The [HTTP Cache API](#) gives developers the ability to control caching within the context of an HTTP request through actions like `Cache-Control` parameter customization and request and response header modification.

Simple Cache API

Limited Availability

The [Simple Cache API](#) exposes a key-value cache interface developers can use to cache small, arbitrary data to help reduce repetitive compute costs.

Limitations and Considerations

Compute trials and paid accounts include the use of Cache APIs. A monthly average of up to 5 cache operations per Compute request are included. A *cache operation* is a read from cache or an operation where a read is followed by a write. There is no separate cap on the maximum number of cache operations per individual Compute request. Contact support@fastly.com if you need to raise the included number of monthly average cache operations.

* * *



Compute



Last updated: 2026-03-31



</products/compute>

Fastly's Compute platform, formerly known as Compute@Edge, 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 runs [WebAssembly](#) (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 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 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 in our [documentation](#).

Logging endpoint compatibility

Compute 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 platform can be accomplished via [the Fastly control panel](#), the [Fastly API](#), and via Fastly's [Terraform provider plugin](#). The [Fastly CLI](#) also provides a local toolchain with features for creating, debugging, and deploying to Wasm services, including Log Tailing and Local Testing.

NOTE

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

Compute features

As part of our Compute product offering, Fastly provides you access to the following collection of features.

NOTE

If you're using a [free account](#), you may need to contact [support](#) to access some features.

Cache APIs

[Cache APIs](#) enable developers to build apps on Compute using Fastly's global cache network.

Dynamic Backends

[Dynamic Backends](#) extends the Fastly API and allows you to decide which origin to connect to at runtime instead of having to define it ahead of time in your configuration. With Dynamic Backends, you can dynamically adjust your origin definitions, dispatch to new hosts based on computed values, and connect to a wider variety of origin servers.

new hosts based on computed values, and connect to a wider variety of origin servers.

Edge Data Storage

Use one of our [Edge Data Storage](#) solutions to store data at the edge as key-value pairs in versionless containers.

Log Tailing

[Log Tailing](#) allows you to stream custom log messages from your Compute application so you can respond quickly when debugging the application without setting up a third-party logging tool.

Local Testing

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

Resource limits

Fastly services and individual instances are allowed a specific range of resources per service and per execution as described in our [Compute resource limits](#). These limits change based on whether or not you've purchased one of our packaged offerings or you're using a [free account](#).

Keep these limitations and constraints in mind especially when [testing and debugging](#) on Compute, when [sending Fastly logs](#) to third party logging providers, and when using [Log Tailing](#).

Billing

NOTE

Billing limits for this product may be different depending on your [account type](#), if you've purchased a [packaged offering](#), or are using a [product or feature trial](#).

Users assigned the superuser role can purchase Compute by directly enabling it, or by enabling a product that runs on Compute, on the [Products page](#) as necessary. Once enabled, any user on your account with the appropriate permissions will be able to use this product.

this product.

Fastly charges for **Compute** 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 charges apply in addition to any already applicable **delivery charges**. Consider setting a **spend alert** for your account to automatically be notified when your month-to-date bill approaches or meets a specific dollar amount.

- **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 vCPU.** Compute vCPU represents the amount of time spent by a single thread of a CPU to process a Compute request. This measures work time (the time spent to execute code) only, not time the process spends waiting on responses or other idle time. Note that Compute vCPU can have minor variations based on typical processor execution variables, so this metric is reported in aggregate and as an average per service.
- **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.

IMPORTANT

Charges based on Compute duration only apply to customers who purchased Compute on or before March 17, 2026 where the Compute duration metric remains on their service order.

* * *



Edge Data Storage



Last updated: 2026-03-31



</products/edge-data-storage>

Compute gives you the option of storing the data you need at the edge as key-value pairs in versionless containers. You can attach these containers to an active service and update the data at any time after it's created, without ever incrementing a service's version. Additionally, these containers can be shared by multiple [Compute services](#) in your account.

Prerequisites

Edge data storage options are only available for Fastly's Compute services, not for VCL-based services.

Important considerations

- Compute is required to use any of the Edge Data Storage options, and separate Compute charges apply in addition to Edge Data Storage billing. For more information on how Compute is billed, refer to the [Compute product documentation](#).
- There is a limit of 100 edge data stores total per account, including [ACLs on Compute](#). To raise this limit, contact your customer success manager.
- Personal information, secrets, or sensitive data should not be included in edge data stores or incorporated into edge logic other than secrets included in Secret Store.
- We do not maintain version histories of your stores. Our [Compliance and Law FAQ](#) describes in detail how Fastly handles personal data privacy.

Config Store

With fast and secure read performance, [Config Store](#) is useful for moving simple functions to the edge. You can store environment variables, redirect lists, and more in Config Store, where they can be shared across services and referenced in your edge logic. A config store might be useful if you:

- have essential data that you don't want to expire.
- have a small dataset.
- need to make frequent changes to your data.

Billing

NOTE

Billing limits for this product may be different depending on your [account type](#), if you've purchased a [packaged offering](#), or are using a [product or feature trial](#).

Config Store is included with your Compute service with the following limitations:

- All accounts include up to five (5) config stores (each having a maximum of 500 entries). You can add additional stores for a monthly fee.
- Write operations to Config Store are limited to 1000 writes per hour. Write limits are shared with all other Fastly API calls. Exceeding the limit will result in an API limit error. Refer to [API rate limiting](#) for more info.
- Config store keys are limited to 255 characters and their values are limited to 8,000 characters.

KV Store

[KV Store](#) is a key-value store that provides high performance reads and writes across Fastly's network to enable powerful edge applications. A KV store might be useful if you:

- require global access to your data from hundreds of POPs.
- require frequent and intensive read operations.
- write or update infrequently (one operation per second, per key).

Limitations and considerations

KV Store supports 1024 byte UTF-8 files with a maximum size of 25MB but that limit can be raised to 100MB upon request.

Billing

NOTE

Billing limits for this product may be different depending on your [account type](#), if you've purchased a [packaged offering](#), or are using a [product or feature trial](#).

Billing for KV Store is based on a combination of total storage charges and data

processing operations for the month. Storage is calculated using GB-months rounded to the nearest hour for each object and processing operations are categorized into three groups, each with differing prices:

- **Class A Operations (writes)** include operations like `SetKey`, `CreateStore`, and `UpdateStore`. Free accounts are limited to 250,000 operations per month.
- **Class B Operations (reads)** include operations like `GetKey` and `GetStore`. Free accounts are limited to 5 million operations per month.
- **Free Operations** include operations like `DeleteKey` and `DeleteStore`.

Consider setting a [spend alert](#) for your account to automatically be notified when your month-to-date bill approaches or meets a specific dollar amount.

All accounts include two KV stores with the following limitations:

- Storage must not exceed 1GB between both stores.
- Class A Operations are limited to 250,000 per month.
- Class B Operations are limited to 5 million per month.

Use of KV Store is disabled by default for new accounts. Users assigned the superuser role can purchase KV Store by contacting sales@fastly.com. Once enabled, any user on your account with the [appropriate permissions](#) will be able to use this product. Additional KV stores and higher storage and operations limits can be purchased as part of an add-on priced in addition to Compute services.

IMPORTANT

Compute is a prerequisite for KV Store, and separate Compute charges may apply once KV Store is enabled. For more information on how Compute is billed, refer to the [Compute product documentation](#).

Secret Store

Fastly Secret Store is a secrets management service that helps you extend protected access to your origins, applications, and other resources on the Fastly edge. This service allows you to store, manage, and access credentials and tokens from your Fastly Compute applications.

A secret store might be useful if you need to securely store:

- Database credentials
- API keys
- Access tokens

Limitations and considerations

Read operations to Secret Store are limited to 5 reads per Compute Request.

Billing

NOTE

Billing limits for this product may be different depending on your [account type](#), if you've purchased a [packaged offering](#), or are using a [product or feature trial](#).

Users assigned the superuser role can purchase Secret Store by enabling it on the [Products page](#). Once enabled, any user on your account with the [appropriate permissions](#) will be able to use this product.

IMPORTANT

Compute is a prerequisite for Secret Store, and separate Compute charges may apply once Secret Store is enabled. For more information on how Compute is billed, refer to the [Compute product documentation](#).

All accounts include ten (10) total secrets across any number of stores. You can add additional secrets for a monthly fee.

Security products notice

No security product, including products storing secrets and sensitive information, 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 and stores does 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 them 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.

* * *

Fanout



Last updated: 2026-03-31



</products/fanout>

Fastly Fanout is a publish-subscribe message broker operating at the Fastly edge. It provides customers with the ability to push messages via direct connections that remain open indefinitely rather than requiring clients to poll for data using short-lived HTTPS requests.

Prerequisites

To use Fanout you must purchase a [paid account](#) with a contract for Fastly's services. Once purchased, it can be [enabled via API](#).

Limitations and considerations

Keep in mind the following limitations and considerations:

- Fanout is only available for Compute services, not Delivery services. Consider using service chaining as an alternative for Delivery services.
- [Service chaining](#) can be used with Fanout. During configuration, the service that initially handles requests from an end user (the first service) must be a Compute service and also must be the Compute service that hands off the request to Fanout.
- Fanout is not compatible with [backend Directors](#), [shielding](#), or the [Fastly Next-Gen WAF](#).
- Only the **Name**, **Address**, **Enable TLS**, and **Override Host** origin server settings are supported.
- Self-signed TLS certificates are not supported. TLS certificates must be signed by a public certification authority.

Billing

NOTE

Billing limits for this product may be different depending on your [account type](#), if you've purchased a [packaged offering](#), or are using a [product or feature trial](#).

Users assigned the superuser role can purchase Fanout by enabling it on the [Products page](#). Once enabled, any user on your account with the [appropriate permissions](#) will be able to use this product.

IMPORTANT

Compute is a prerequisite for Fanout, and separate Compute charges may apply once Fanout is enabled. For more information on how Compute is billed, refer to the [Compute product documentation](#).

We bill for Fanout based on a combination of bandwidth, connection time, and number of messages. Bandwidth is included as part of your overall delivery bandwidth rate in your monthly billing statement. Connection time is measured for each connection in usage minutes (rounded up to whole minutes) and aggregated monthly to millions of minutes. Number of messages is the sum of the total published messages received from the publish API and the total published messages sent to end users. Consider setting a [spend alert](#) for your account to automatically be notified when your month-to-date bill approaches or meets a specific dollar amount.

For more details about a product, including [pricing information](#) or for help purchasing it, contact your customer success manager or email sales@fastly.com.

* * *

Category: **Support**

These articles provide information about Fastly services and support solutions.



Assurance Services



Last updated: 2018-03-30

</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 [security](#) and [technology compliance](#) programs, subject to Subscriber's purchase of [Professional Services](#). 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](#).

Some Fastly products can be purchased directly in the web interface on the [Products page](#). For more details about a product, including [pricing information](#) or for help purchasing it, contact your customer success manager or email sales@fastly.com.

* * *

DDoS Protection and Mitigation



Service and SLA



Last updated: 2023-05-26



</products/ddos-protection-and-mitigation-service-and-sla>

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.

help	
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 [published API endpoint](#) that allows you 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 your origin. Fastly provides the Fastly IP Space to you in order to ensure known communication between the Fastly cache nodes and your origin's data center.
- **"Fastly Control"** means elements entirely under Fastly's control and not a consequence of (a) your hardware or software failures, (b) you or your end user's connectivity issues, (c) your operator errors, (d) traffic amounts that exceed your Permitted Utilization as defined in the Terms and Conditions, (e) your corrupted 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 [support ticket](#) 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 size of the attack threatened or previously experienced.

- 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 visiting <https://support.fastly.com/>, sending 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 billing@fastly.com.
- 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 [Termed Contract](#), 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.

Some Fastly products can be purchased directly in the web interface on the [Products page](#). For more details about a product, including [pricing information](#) or for help purchasing it, contact your customer success manager or email sales@fastly.com.

* * *



Fastly Next-Gen WAF professional services



Last updated: 2024-04-01



</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 [Fastly Next-Gen WAF](#).

Fastly Next-Gen WAF (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.

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.

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.

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.

For purchases of more than 200 hours in a year, hours are evenly allocated throughout the contract term on a quarterly interval. If the hours allocated for the quarter are not used, they will expire at the calendar quarter end.

Some Fastly products can be purchased directly in the web interface on the [Products page](#). For more details about a product, including [pricing information](#) or for help purchasing it, contact your customer success manager or email sales@fastly.com.

* * *



Live Event Monitoring Service



Last updated: 2022-09-01



</products/live-event-monitoring-service>

IMPORTANT

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

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 to receive alerts and notifications as well as ask questions without losing time spent contacting support and recounting what the issue is.

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.

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

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. Fastly will disable logging at the conclusion of troubleshooting, and collected log data will be retained for no more than 30 days after logging is disabled, unless otherwise instructed by you, your company, or organization.

Some Fastly products can be purchased directly in the web interface on the [Products page](#). For more details about a product, including [pricing information](#) or for help purchasing it, contact your customer success manager or email sales@fastly.com.

* * *



Managed Security Enterprise



Last updated: 2025-10-28



</products/managed-security-enterprise>

Managed Security Enterprise (formerly referred to as "Fastly Managed Security Service") provides your organization with continuous monitoring of your included Fastly security products, proactive action by Fastly in the event of an identified security incident, enhanced access to our Customer Security Operations Center (CSOC), 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 consulting on initial configuration, requested maintenance, monitoring, and attack support.

The following table summarizes what Managed Security Enterprise provides:

Support offering	Details
CSOC proactive monitoring of Fastly security products	24/7/365.
CSOC availability for general requests and inquiries	24/7/365.
CSOC availability for critical security incidents	24/7/365.
General inquiries response times	Within 24 hours.
Customer-identified critical security incident response times	Within 15 minutes of notice with active troubleshooting starting within 30 minutes of acknowledgement of incident severity.
Fastly-identified critical security incident notification times	Within 30 minutes of notice with active troubleshooting starting within 30 minutes of acknowledgement of incident severity.
Emergency phone number	Included.
Emergency email address	Included.
Dedicated security chat channel	Included.
Post-event report	Upon request or at Fastly's discretion.
Monthly security report	Once per month.
Monthly reviews	Upon request, up to one per month, during business days and business hours, remote-only, limited to a maximum

	duration of 2 hours, and must be requested 10 business days in advance.
Readiness drill	Up to once every six months, at your request, Fastly will partner with you to execute a readiness drill.

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 not a consequence of (a) your hardware or software failures, (b) you or your end user's connectivity issues, (c) your operator errors, (d) traffic amounts that exceed your Permitted Utilization as defined in the Terms and Conditions, (e) your corrupted content, (f) acts of god (any) or war, or earthquakes, or terrorist actions.
- **"Full-Site Delivery Services"** means the configuration for a website, app, API, or anything else to be served through Fastly's Full-Site Delivery product.

Prerequisites

To purchase and use Managed Security Enterprise, you must also purchase [Fastly Full-Site Delivery](#) and delivery [Enterprise Support](#), along with either [Next-Gen WAF](#) or [Edge Rate Limiting](#).

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

Features

Fastly offers Managed Security Enterprise for the term as indicated in your Service Order. It includes the following features.

Continuous security product monitoring

Continuous security product monitoring

Fastly will continuously monitor your included [Fastly Security products](#) and [Fastly Full-Site Delivery](#) for early detection of security events and take proactive action in the event we identify a security event resulting in a security incident. The set of security events we monitor may change over time. You can request a current listing of the security incidents we monitor by submitting a general request to the CSOC. We will follow the contact procedures defined in your runbook in the event that we need to contact you prior to taking an action. It is your responsibility to notify us if any of your contact methods or permitted actions need to change.

The proactive actions Fastly may take on your behalf are explicitly permitted by you and documented in your customer runbook. During onboarding, we'll agree on which actions we may take and you can update these actions by contacting the CSOC. Actions you may permit us to take can include, but are not limited to contacting you if we identify a security event requiring mitigation action and taking proactive action to mitigate the event.

Threat hunting

Fastly's threat hunting proactively analyzes customer data up to once per month, searching for security risks and weaknesses. This approach strengthens your security posture and helps identify and mitigate potential threats.

Post-event reports

At Fastly's discretion or at your request, Fastly will provide post-event reports for security incidents. These reports will document what Fastly observed and what actions were taken.

Monthly security report

Once a month, Fastly will send you a report documenting events observed and actions taken over the past month, recommendations for configuration changes and maintenance, results of threat hunting (when available), threat intelligence updates, and product updates.

Readiness drill

Up to once every six months, at your request, Fastly will partner with you to execute a readiness drill. This readiness drill simulates multiple phases of a security event with the objective of improving incident response. The scope of the readiness drill is at Fastly's

discretion, but will typically include how we will engage and escalate during an attack scenario. You must schedule a readiness drill at least two weeks in advance by making a request by contacting the CSOC. You won't be entitled to any refunds or credits for unused Readiness drills availability.

Account and event reviews

At your request, Fastly will schedule 2-hour account and event reviews during US business hours, up to once per month during the term of your subscription, with a designated security specialist who will review recent security events and actions taken, 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 by contacting the CSOC. You won't be entitled to any refunds or credits for unused reviews.

Priority CSOC access

By purchasing Managed Security Enterprise, 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 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 security event notifications, general security product support and questions. The chat channel will be monitored 24/7 by Fastly's CSOC. Inquiries regarding critical security incidents should be communicated using the notification mechanism that will be described during onboarding. Though subject to change, Fastly's current chat provider is Slack (www.slack.com).

Each of these contact methods will be provided to you (as applicable) during your onboarding period.

Support requests and response times

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 Managed Security Enterprise customers or by submitting a ticket via the control panel 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

Although Managed Security Enterprise includes continuous monitoring and proactive response to security incidents, there may be times where you need to notify us of a critical security incident requiring support. Support for critical security incidents that you identify can only be initiated via the emergency email address provided to Managed Security Enterprise customers (not chat) or by selecting the Urgent priority when submitting a ticket via the control panel. The ultimate classification of a request submitted by either of these methods will be determined solely by Fastly based on various factors including input and information you provide.

Fastly will acknowledge your critical security incident outreach within fifteen (15) minutes of its receipt. Alternatively, Fastly will notify you of critical incidents that we detect within thirty (30) minutes of detection. If classified as a critical security incident, we will begin actively troubleshooting these incidents within thirty (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 four (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.

Response SLA and credit terms

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

- Requests for invoice credits must be made within 30 calendar days of the critical

security incident that triggered the service credit.

- All requests for invoice credits must be made to billing@fastly.com.
- A pending invoice credit does not release you from your obligation to pay all 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 consecutive three-month period where three (3) or more support response or notification time objectives are not met and the failure to meet the objectives materially adversely impacted you, you will have 30 days to terminate the Managed Security Enterprise subscription following the third response failure. You must notify Fastly of your intention to terminate the Managed Security Enterprise subscription, or the Managed Security Enterprise portion of any bundled subscription, within thirty (30) days of the triggering event. No other remedy or refund will be available other than your ability to terminate your subscription to Managed Security Enterprise.

Observational logging

Fastly will access and use your logs as part of Managed Security Enterprise. Logs will be used by Fastly to provide you with support, to monitor and maintain your Fastly security products, and as a means of threat detection.

Fastly will collect, store, and use a sampled subset of logging information generated by the Fastly content delivery network and security products (including IP addresses) for purposes including, but not limited to, monitoring product behavior, managing false positives, making configuration adjustments, producing periodic customer reports, making improvements to our products and services, improving our detection capabilities, and detecting potential security incidents. Fastly will do this by establishing a logging endpoint in your service configuration that will securely collect logging information in a third-party storage provider. Fastly will derive aggregated, anonymized data from the logs collected. This data will be used to improve security products and services for all subscribers, and includes statistical analyses as well as the development of security research and threat intelligence products.

By subscribing to Managed Security Enterprise, you instruct Fastly to access and use the logs for providing the above purpose in accordance with the Fastly Documentation. Sampled 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. Aggregated data will be deleted on a rolling basis and in any event retained no longer than ninety (90) days unless

on a rolling basis and in any event retained no longer than ninety (90) days unless otherwise agreed to by you.

Limitations

Managed Security Enterprise has the following limitations:

- **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.
- **Services monitored.** We will monitor up to ten Full-Site Delivery Services. You may request additional monitoring by submitting a general request to the CSOC.

No security 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 does 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.

Some Fastly products can be purchased directly in the web interface on the [Products page](#). For more details about a product, including [pricing information](#) or for help purchasing it, contact your customer success manager or email sales@fastly.com.



Managed Security Professional



Last updated: 2025-10-28



</products/managed-security-professional>

Managed Security Professional provides your organization with continuous monitoring of your included Fastly security products, proactive action by Fastly in the event of an identified security incident, enhanced access to our Customer Security Operations Center (CSOC), 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 consulting on initial configuration, requested maintenance, monitoring, and attack support.

The following table summarizes what Managed Security Professional provides:

Support offering	Details
CSOC proactive monitoring of Fastly security products	24/7/365.
CSOC availability for general requests and inquiries	24/7/365.
CSOC availability for critical security incidents	24/7/365.
General inquiries response times	Within 24 hours.
Customer-identified 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 security chat channel	Included.
Post-event report	Upon request or at Fastly's discretion.
Security report	Once per quarter.
Reviews	Upon request, up to one per quarter, during business days and business hours, remote-only, limited to a maximum duration of 2 hours, and must be requested 10 business days in advance.

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 not a consequence of (a) your hardware or software failures, (b) you or your end user's connectivity issues, (c) your operator errors, (d) traffic amounts that exceed your Permitted Utilization as defined in the Terms and Conditions, (e) your corrupted content, (f) acts of god (any) or war, or earthquakes, or terrorist actions.
- **"Full-Site Delivery Services"** means the configuration for a website, app, API, or anything else to be served through Fastly's Full-Site Delivery product.

Prerequisites

To purchase and use Managed Security Professional, you must also purchase Fastly [Full-Site Delivery](#) and Network Services [Enterprise Support](#), along with either [Next-Gen WAF](#)

or [Edge Rate Limiting](#).

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

Features

Fastly offers Managed Security Professional for the term as indicated in your Service Order. It includes the following features.

Fastly will continuously monitor your included [Fastly Security products](#) and [Fastly Full-Site Delivery](#) for early detection of security events and take proactive action in the event we identify a security event resulting in a security incident. The set of security events we monitor may change over time. You can request a current listing of the security incidents we monitor by submitting a general request to the CSOC. We will follow the contact procedures defined in your runbook in the event that we need to contact you prior to taking an action. It is your responsibility to notify us if any of your contact methods or permitted actions need to change.

The proactive actions Fastly may take on your behalf are explicitly permitted by you and documented in your customer runbook. During onboarding, we'll agree on which actions we may take and you can update these actions by contacting the CSOC. Actions you may permit us to take can include, but are not limited to contacting you if we identify a security event requiring mitigation action and taking proactive action to mitigate the event.

Post-event reports

At Fastly's discretion or at your request, Fastly will provide post-event reports for security incidents. These reports will document what Fastly observed and what actions were taken.

Quarterly security report

Once a quarter, Fastly will send you a report documenting events observed and actions taken over the past month, recommendations for configuration changes and maintenance, and product updates.

Account and event reviews

At your request, Fastly will schedule 2-hour account and event reviews during US

At your request, Fastly will schedule a team account and event review during our business hours, up to once per quarter during the term of your subscription, with a designated security specialist who will review recent security events and actions taken, 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 by contacting the CSOC. You won't be entitled to any refunds or credits for unused reviews.

Priority CSOC access

By purchasing Managed Security Professional, 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 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 security event notifications, general security product support and questions. The chat channel will be monitored 24/7 by Fastly's CSOC. Inquiries regarding critical security incidents should be communicated using the notification mechanism that will be described during onboarding. Though subject to change, Fastly's current chat provider is Slack (www.slack.com).

Each of these contact methods will be provided to you (as applicable) during your onboarding period.

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 Managed Security Professional customers or by

submitting a ticket via the Next-Gen WAF control panel 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

Although Managed Security Professional includes continuous monitoring and proactive response to security incidents, there may be times where you need to notify us of a critical security incident requiring support. Support for critical security incidents that you identify can only be initiated via the emergency email address provided to Managed Security Professional customers (not chat) or by selecting the Urgent priority when submitting a ticket via the Next-Gen WAF control panel. The ultimate classification of a request submitted by either of these methods will be determined solely by Fastly based on various factors including input and information you provide.

Fastly will acknowledge your critical security incident outreach within fifteen (15) minutes of its receipt. If classified as a critical security incident, we will begin actively troubleshooting these incidents within thirty (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 four (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.

Response SLA and credit terms

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

- Requests for invoice credits must be made within 30 calendar days of the critical security incident that triggered the service credit.
- All requests for invoice credits must be made to billing@fastly.com.
- A pending invoice credit does not release you from your obligation to pay all 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 consecutive three-month period where three (3) or more support response or notification time objectives are not met and the failure to meet the objectives materially adversely impacted you, you will have 30 days to terminate the Managed Security Professional subscription following the third response failure. You must notify Fastly of your intention to terminate the Managed Security Professional subscription, or the Managed Security Professional portion of any bundled subscription, within thirty (30) days of the triggering event. No other remedy or refund will be available other than your ability to terminate your subscription to Managed Security Professional.

Observational logging

Fastly will access and use your logs as part of Managed Security Professional. Logs will be used by Fastly to provide you with support, to monitor and maintain your Fastly security products, and as a means of threat detection.

Fastly will collect, store, and use a sampled subset of logging information generated by the Fastly content delivery network and security products (including IP addresses) for purposes including, but not limited to, monitoring product behavior, managing false positives, making configuration adjustments, producing periodic customer reports, making improvements to our products and services, improving our detection capabilities, and detecting potential security incidents. Fastly will do this by establishing a logging endpoint in your service configuration that will securely collect logging information in a third-party storage provider. Fastly will derive aggregated, anonymized data from the logs collected. This data will be used to improve security products and services for all subscribers, and includes statistical analyses as well as the development of security research and threat intelligence products.

By subscribing to Managed Security Professional, you instruct Fastly to access and use the logs for providing the above purpose in accordance with the Fastly Documentation. Sampled 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. Aggregated data will be deleted on a rolling basis and in any event retained no longer than ninety (90) days unless otherwise agreed to by you.

Limitations

Managed Security Professional has the following limitations:

- **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.
- **Services monitored.** We will monitor up to two Full-Site Delivery Services. Monitoring for additional services is available via add-on purchase.

No security 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 does 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.

Some Fastly products can be purchased directly in the web interface on the [Products page](#). For more details about a product, including [pricing information](#) or for help purchasing it, contact your customer success manager or email sales@fastly.com.

* * *



Network Services service availability SLA



Last updated: 2024-09-04



</products/network-services-service-availability-sla>

Network Services service level agreements depend on the [type of account](#) you have as summarized in the following table:

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 you or Fastly, (3) requiring you 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 you or Fastly.
- **"Fastly Control"** means elements entirely under Fastly's control and not a consequence of (a) your hardware or software failures, (b) you or your end user's connectivity issues, (c) your operator errors, (d) traffic amounts that exceed your Permitted Utilization as defined in the Terms and Conditions, (e) your corrupted 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 (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%
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.

* * *



Network Services Support description and SLA



Last updated: 2025-05-08



</products/network-services-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, including those in any [packaged offering](#). 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	Business hours, including weekends &	24/7/365.	24/7/365.

Reports	weekends & holidays.		
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.
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.
Technical Account Manager	Not available.	Not available.	Available with the purchase of a Technical Account Manager add-on package.
Discounted Professional Services	Not available.	Not available.	30% discount on Professional Services packages. Does not apply to Fastly Next-Gen WAF service packages.
PCI and HIPAA configuration	Not available.	Not available.	Available via email, phone, and chat

services		support.	
Compute 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 partners@fastly.com 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 on-demand 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.

- 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) your hardware or software failures, (b) you or your end user's connectivity issues, (c) your operator errors, (d) traffic amounts that exceed your Permitted Utilization as defined in the Terms and Conditions, (e) your corrupted 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 visiting <https://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 [Service Availability SLA](#)),

- 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

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.

Enterprise 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 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 [Premium Support](#) and [Platinum Support](#) plans. To convert your account to the current Gold and Enterprise Support plans, contact sales@fastly.com. If you have an agreement that requires the purchase of Platinum support, converting to Enterprise support satisfies that requirement.

* * *



Performance Optimization Package



Last updated: 2020-04-01



</products/performance-optimization-package>

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 [paid account with a contract](#) 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 [Professional Services team](#) can be purchased at an additional cost.

Some Fastly products can be purchased directly in the web interface on the [Products page](#). For more details about a product, including [pricing information](#) or for help purchasing it, contact your customer success manager or email sales@fastly.com.

* * *



Professional Services



Last updated: 2025-08-15



</products/professional-services>

Fastly allows customers to purchase assistance from Solutions Engineering staff to help you onboard and optimize your use of Fastly's services. Our team will help you through planning, implementing, testing, and launching your project. Choose between one of our [Onboarding Packages](#) and [Professional Service Hours](#), depending on your needs.

Prerequisites

To purchase a Professional Services offering, you must have a [paid account with a contract](#) for Fastly's services.

Onboarding Packages

First time Fastly Network Services customers must purchase one of our Onboarding Packages to help you onboard to Fastly in a way suited to your needs. When you purchase an Onboarding Package, you'll get access to a member of our Solutions Engineering staff to help you with tasks related to onboarding your Delivery or Compute services, such as service implementation and initial setup and configuration, along with regular check-ins during your initial use of Fastly. Contact your customer success manager or email sales@fastly.com for more information.

Professional Service Hours

If you're unsure of your project needs, our Professional Service Hours offering allows you to purchase a number of hours with our Solutions Engineering team to use over a set period of time.

Professional Service Hours can be used for things like:

- Feature and product configuration and integration
- Performance and caching optimization
- Compute code development and code review
- Consultation with guidance and advice for best practices

When you're ready to get started on a project that requires Professional Services assistance, reach out to your account team or request a consultation by contacting sales@fastly.com. Fastly Solutions Engineering staff will set up an initial meeting to understand your requirements and success criteria. They'll help you determine the type of assistance you'll need and, when the project doesn't have a predefined scope, help you estimate the number of Professional Service hours the project might need. The number of hours and the term start and end date will appear on your service order.

For purchases of more than 200 hours in a year, hours are evenly allocated throughout the contract term on a quarterly interval. If the hours allocated for the quarter are not used, they will expire at the calendar quarter end.

How it works

Depending on the Onboarding Package or Professional Service Hours you have purchased, our Solutions Engineering staff will personally guide you through one or more of the following stages:

- **Planning.** They'll help you with things like requirements gathering, solution design, documentation, and resource allocation.
- **Training.** They'll demo the Fastly control panel, walk through the request and response flow through Fastly, and share best practices, caveats, and support methods.

- **Implementation.** They'll help you with things like configuration of Fastly services and custom VCL or Compute development.
- **Testing.** They'll help you validate configurations and set up testing.
- **Go-live.** They'll supervise your traffic moving to Fastly, post go-live monitoring, and help address issues during final production testing and deployment.

It typically takes about 10 business days for Fastly's Solutions Engineering staff to scope your problem and initiate a solution. The implementation, testing, and go-live stages may involve some iterative cycles depending on the complexity of your configuration.

* * *



Response Security



Last updated: 2025-10-28



</products/response-security>

Fastly offers [Fastly Next-Gen WAF](#) customers a Response Security service (formerly referred to as a "Response Security Service") 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 following table summarizes what Response Security provides:

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

General inquiries response times	Within 24 hours.
Customer-identified 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 not a consequence of (a) your hardware or software failures, (b) you or your end user's connectivity issues, (c) your operator errors, (d) traffic amounts that exceed your Permitted Utilization as defined in the Terms and Conditions, (e) your corrupted content, (f) acts of god (any) or war, or earthquakes, or terrorist actions.

Prerequisites

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 features

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

Priority CSOC access

By purchasing Response Security, 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.

Online help and virtual training

In addition to unlimited access to online self-service documentation at docs.fastly.com 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 Response Security

at least two weeks in advance by making a request via the provided Response Security general support question email address. You won't be entitled to any refunds or credits for unused scheduled availability.

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 Response Security ticket at <https://support.fastly.com> 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 Response Security customers (not chat) or by selecting the Urgent priority when submitting a ticket at <https://support.fastly.com>. 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.

Response SLA and credit terms

If you have purchased Response Security 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 billing@fastly.com.
- In no event shall invoice credits exceed the fee for Response Security 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 Response Security subscription following the third response failure. You must notify Fastly of your intention to terminate the Response Security subscription within 30 days of the triggering event.

Limitations

Response Security has the following limitations:

- **Product applicability.** This service only applies to subscribed and fully provisioned Fastly security products. 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.
- **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, 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 does 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.

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.

Some Fastly products can be purchased directly in the web interface on the [Products page](#). For more details about a product, including [pricing information](#) or for help purchasing it, contact your customer success manager or email sales@fastly.com.

* * *



Security service SLA



Last updated: 2023-04-05



</products/security-service-sla>

Fastly endeavors to maintain 99.9% availability of the Signal Sciences Hosted Dashboard ([Hosted Dashboard](#)), 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: $([\text{\# of minutes in month}] - [\text{\# of minutes per month the Hosted Dashboard is Unavailable}]) / [\text{\# of minutes in month}]$.

"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, (b) you or your end user's connectivity issues, (c) your operator errors, (d) traffic amounts that exceed your Permitted Utilization as defined in the Terms and Conditions, (e) your corrupted content, (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: $([\text{\# of minutes in month}] - [\text{\#$

Availability of the Hosted Services is calculated as follows: $(\text{[# of minutes in month]} - \text{[# of minutes per month the Hosted Services is Unavailable]}) / \text{[# of minutes in month]}$.

"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, (b) you or your end user's connectivity issues, (c) your operator errors, (d) traffic amounts that exceed your Permitted Utilization as defined in the Terms and Conditions, (e) your corrupted content, (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 [Fastly Next-Gen WAF \(Edge\)](#) 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 [fail open](#) 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.

* * *



Security support description and SLA



Last updated: 2023-04-05

</products/security-support-description-and-sla>

Security support availability and response times vary depending on the [type of account](#) you have and the level of support you have purchased, including those in any [packaged offering](#). 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 for Urgent - Critical Impact (P0) issues	1 business day with updates every 4 business days	60 minutes or less with updates every 2 hours (24/7/365)	15 minutes or less with updates every 2 hours (24/7/365)
Email support	Available.	Available, with priority over Standard Support.	Available, with priority over Standard and Gold Support.
Phone support	Not available.	Not available.	Toll-free telephone available 24/7/365.
Chat channel support	Not available.	Not available.	Dedicated chat channel available during Fastly business hours.
Security Technical			Available for

Account Manager	Not available.	Not available.	purchase.
-----------------	----------------	----------------	-----------

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 your services is adversely impacted.
- **"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.

Communications

The following section applies to all subscribers. To ensure accurate responses, you must keep your account contact information up-to-date. Subscribers can submit support requests by visiting <https://support.fastly.com/> or they use one of the additional options described below.

Tickets

TICKETS

Fastly Next-Gen WAF 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 support@fastly.com during standard business hours. All support tickets generated by email will be designated with a P2 severity level.

Chat

Subscribers to Enterprise Security Support receive a dedicated chat channel for real-time communications to discuss general security product support topics 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).

Phone support

Subscribers to Enterprise Security 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.

Incident reporting and response times

The following section applies to all subscribers.

Response times

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

Issue severity level	Standard Support	Gold Support	Enterprise Support
P1	1 business day	2 hours	15 minutes

Urgent - Critical Impact (P0)	with updates every 4 business days	60 minutes or less with updates every 2 hours (24/7/365)	15 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

Additional fees

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 your control, you agree to pay Fastly its hourly charges for Fastly engineering staff time. If it appears likely the factors are within your control, you may tell Fastly staff to stop working on troubleshooting the incident (thereby stopping the hourly charges from being incurred). You agree 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.

Security support invoice credits

In the event a P0 incident occurs, you have purchased Gold or Enterprise Security Support, the cause of the incident is within Fastly's control, and the communication or response timeframes are materially not met, a one-time credit of \$500 per incident will be credited to your account.

Credit Terms:

- Requests for Invoice Credits must be made within 30 days of the incident which

triggered the service credit.

- All requests for invoice credits must be made to billing@fastly.com.
- To the extent you purchase Enterprise Security Support a la carte and not as part of a packaged offering, in no event shall Invoice Credits exceed the invoice value for Enterprise Security Support in the month in which they are accrued.
- A pending credit does not release you from its 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 Gold or Enterprise Security Support subscription following the third response failure. You must notify Fastly of your intention to terminate the Enterprise Security Support subscription within 30 days of the triggering event. Please note, the termination right discussed in this paragraph does not apply when Gold or Enterprise Security Support are purchased as part of a packaged offering.

* * *



Security Technical Account Manager



Last updated: 2024-02-20



</products/security-technical-account-manager>

Fastly offers customers the ability to purchase the support of a Security Technical Account Manager for your organization. These specialists help you optimize your use of Fastly's security products and features by providing proactive check-ins and regular reviews to help you analyze your account's security service configurations and their performance. Security Technical Account Managers also provide enhanced troubleshooting coordination with Fastly's support and professional services organizations.

TIP

Need help with products or features that aren't security related? Check out Fastly's [Technical Account Manager](#) offering instead.

Fastly offers Professional and Premier Security Technical Account Manager packages. Available hours of service each month to your organization depend on the package you purchase and could include the following summarized activities:

Support Offering	Professional	Premier	Enterprise	Global Add-on
Total available hours	Up to 4 hours/month	Up to 20 hours/month	Up to 40 hours/month	Up to 20 additional hours/month
Scheduled technical check-ins	Quarterly	1x month	4x monthly (as requested)	-
Named point of contact	Included	Included	Included	-
Hours of operation/availability	Business hours in single time zone	Business hours in single time zone	Business hours in single time zone	Business hours in 1 additional time zone (max 2 regions)
Regional coverage	1 region	1 region	1 region	1 additional region
Architecture and configuration recommendations	Included	Included	Included	-
Health check reporting	Quarterly	Quarterly	Quarterly	-
Escalation support				

and coordination	Included	Included	Included	-
Email support	Available	Available	Available	-
Private chat 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	12 business hours
Virtual, self-paced training	Included	Included	Included	-
Live training	Up to 4 sessions	Up to 8 sessions	Up to 8 sessions	Up to 8 sessions
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.
- **"Regions"** are areas of the world where Security TAM coverage is available and are restricted to North America, Europe, and Japan, plus Australia and New Zealand.

IMPORTANT

Security Technical Account Managers provide support during Fastly business hours and facilitate non-urgent discussions. They are not a 24×7 resource. Always rely on [normal support communications channels](#) for urgent issues and escalations. To engage the services of our Customer Security Operations Center (CSOC) team, check out our [Response Security](#) offering.

Prerequisites

To purchase and use a Security Technical Account Manager package, you must also purchase either [Enterprise Security Support \(ESS\)](#), [Response Security](#), or [Managed Security Enterprise](#).

To ensure accurate responses to requests, you must keep your account contact information up-to-date.

Security Technical Account Manager packages

Each Security Technical Account Manager package includes the following core features:

- Priority engagement and coordination with the support resources as necessary during Fastly business hours.
- Technical guidance on topics like security-related configuration, service performance, and roadmap requests.
- Advice on security best practices when implementing and using Fastly with your infrastructure.
- Quarterly health check reporting.

For all Security Technical Account Manager packages, 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.

NOTE

English is the primary language used by our Security Technical Account Managers.

Professional Security Technical Account Manager

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

- Architectural and configuration recommendations, as requested.
- Scheduled technical check-ins, via phone or video conference, 1x quarterly.
- Additional, live training, up to 4 sessions.

Premier Security Technical Account Manager

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

- Architectural and configuration recommendations, as requested.
- More frequent scheduled technical check-ins, 1x monthly.
- Additional live training, up to 8 sessions.

Enterprise Security Technical Account Manager

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

- More frequent scheduled technical check-ins, 4x monthly upon request.
- Additional live training, up to 8 sessions.

With this level of support, you also receive early beta program access.

Global Add-On Security Technical Account Manager

When combined with the Enterprise Security Technical Account Manager package, included hours could be used for:

- Enterprise Security Support in an additional time zone (max 2 regions).
- Additional live training, up to 16 sessions.

With this level of support, you also receive early beta program access and increased initial response time for inquiries.

Some Fastly products can be purchased directly in the web interface on the [Products page](#). For more details about a product, including [pricing information](#) or for help purchasing it, contact your customer success manager or email sales@fastly.com.

* * *



Technical Account Manager



Last updated: 2024-02-21

 Last updated: 2024-02-21

 </products/technical-account-manager>

Fastly offers the ability to purchase the support of a Fastly Engineer to serve as a Technical Account Manager (formerly referred to as a "Designated Technical Specialist") for your organization. These specialists act as an extension of your team. They are product experts who 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. Technical Account Managers also provide enhanced troubleshooting coordination with Fastly's support and professional services organizations.

TIP

Looking to optimize a security product or feature? Check out Fastly's [Security Technical Account Manager](#) offering instead.

Fastly offers Professional, Premier, and Enterprise Technical Account Manager packages, and also offers a Global Add-On for customers purchasing the Enterprise Technical Account Manager package. A Technical Account Manager'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	Professional	Premier	Enterprise	Global Add-on
Total available hours	Up to 40 hours/month	Up to 80 hours/month	Up to 160 hours/month	Up to 50 additional hours/month
Scheduled technical check-ins	Monthly	2x monthly	4x monthly (as requested)	-
Named point of contact	Included	Included	Included	-
	Business hours	Business hours	Business hours	Business hours in 4 time zones

Hours of operation/ availability	Business hours in single time zone	Business hours in single time zone	Business hours in single time zone	hours in 1 additional time zone (max 2 regions)
Regional coverage	1 region	1 region	1 region	1 additional region
Architecture and configuration recommendations	Included	Included	Included	-
Escalation support and coordination	Included	Included	Included	-
Email support	Available	Available	Available	-
Private chat 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	12 business hours
Custom reporting	Not included	Not included	By request	-

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.
- **"Regions"** are areas of the world where Designed Technical Specialist coverage is available and are restricted to North America, Europe, Japan, and Australia.

IMPORTANT

Technical Account Managers provide support during Fastly business hours to facilitate *non-urgent* discussions. They are not a 24×7 resource. Always rely on [normal support communications channels](#) for urgent issues and escalations.

Prerequisites

To purchase and use a Technical Account Manager package, you must have an [Enterprise Support plan](#).

To ensure accurate responses to requests, you must keep your account contact information up-to-date.

Technical Account Manager packages

Each Technical Account Manager package includes the following core features:

- 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 Technical Account Manager packages, 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.

NOTE

English is the primary language used by Technical Account Managers.

Professional Technical Account Manager

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

- architectural and configuration recommendations, as requested.

- scheduled technical check-ins, via phone or video conference, 1x monthly.

Premier Technical Account Manager

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

- account reports with an overview of services, traffic, and tickets.
- more frequent scheduled technical check-ins, 2x monthly.
- a quarterly service review for optimizing your Fastly strategy.

Enterprise Technical Account Manager

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).
- custom reporting upon reasonable request.

Global Technical Account Manager Add-On

When combined with the Enterprise Designed Technical Specialist offering, you receive:

- an additional Technical Account Manager assigned to your account.
- Technical Account Manager support in one additional region or time zone.
- increased initial response time for inquiries.

The Global Technical Account Manager Add-On can only be added to an existing Enterprise Technical Account Manager package.

Some Fastly products can be purchased directly in the web interface on the [Products page](#). For more details about a product, including [pricing information](#) or for help purchasing it, contact your customer success manager or email sales@fastly.com.

Category: Resources

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



Compute resource limits



Last updated: 2026-06-16



</products/compute-resource-limits>

This guide details Fastly resource limits for [Compute](#) services. These limits change based on whether you're using a free account or have purchased one of our [packaged offerings](#). These limits may be further defined on your service order.

NOTE

In addition to the limits described here, Compute services can also be limited by the products or platform services that you access through your application code. Check out our [network services resource limits](#) for additional details.

Port limits

By default, Compute services only accept client connections on port 443. To learn more about allowing client connections on port 80, contact your account manager or sales@fastly.com.

NOTE

By default, if your Compute service receives a request on port 80, Fastly automatically returns a 308 Permanent Redirect response status with the `Location` header indicating the HTTPS version of the same URL.

Default limits

Item	Limit	Scope
------	-------	-------

Item	Limit	Scope
Maximum compiled package size	100MB	per service
Maximum cached object size	100MB	per object
Maximum CPU time available to a single request instance	50ms	per execution
Maximum runtime for a single request instance	2 min (60s for trial accounts)	per execution
Maximum memory consumption	1M bytes stack, 128MB heap	per execution
Maximum number of dictionary lookups	32	per execution
Maximum number of backend requests	32 (10 for trial accounts)	per execution
Maximum number of dynamic backends	200	per service
Vary rules	8	per cache object
Maximum length of a request header	128KB	per header
Maximum length of a response header	128KB	per header
Maximum length of a method name	8192 bytes	per request
Maximum length of a URL	8192 bytes	per request

An *execution* refers to a single instance of a Compute program being executed, normally in response to a client HTTP request. Separate limits also apply to the use of [log tailing](#) with Compute services.

Service chain limits

Item	Limit	Implications
Service chains and hops	20 hops, 6 unique services	Read our developer information about loop detection for more info.

Edge data storage limitations

When working with edge data stores, there is a limit of 100 shared resources total per account (Config Store, KV Store, Secret Store, and ACLs on Compute combined). Limits for each of our [edge data storage](#) options are outlined below, but may differ if you've purchased a [package offering](#) or are using a [product or feature trial](#). Some limits may be raised; contact your customer success manager for details.

Edge data storage naming restrictions

When naming a store, whether through the [web interface](#) or [API](#), be aware of the following restrictions:

- Store names must be unique.
- Store names must begin with an ASCII alphabetical character and contain only ASCII alphanumeric and underscore characters.
- Store names cannot be an empty string.

Config Store

Config Store is included with your Compute service with the following limitations:

Item	Limit	Implications
Config stores	5	Paid accounts can purchase additional stores for a monthly fee.
Config store entries (per store)	500	An error informs you that you've exceeded the limit.
Config store key length	255 characters	An error informs you that you've exceeded the limit.
Config store value length	8,000 characters	An error informs you that you've exceeded the limit.
Config store write operations	1000 writes per hour	Write limits are shared with all other Fastly API calls. Exceeding the limit will result in an API limit error. Refer to API rate limiting for more info

[for more info.](#)

KV Store

KV Store is included with your Compute service with the following limitations:

Item	Limit	Implications
KV stores	5	Paid accounts can purchase additional stores for a monthly fee.
File size limit	25 MB	This limit can be raised to as much as 100 MB upon request. Contact your customer success manager or sales@fastly.com for more information.
Metadata storage	2000 bytes	Must be composed of Unicode characters encoded using UTF-8.
Class A operations (writes)	250,000 per month	Paid accounts can raise this limit for an additional fee.
Class B operations (reads)	5 million per month	Paid accounts can raise this limit for an additional fee.
Writes per second (per store)	1000	This is a soft limit for the durable store and it auto-scales as needed.
Reads per second (per store)	5000	This is a soft limit for the durable store and it auto-scales as needed.
Rate limit	1 write per second per item	An error informs you that you've exceeded the limit.
Batch requests	100,000 items per request	Exceeding this limit will result in long execution time before eventual failure.

Key limitations

- Keys cannot start with `.well-known/acme-challenge/`, `.`, or `..`.
- Key names can contain any sequence of valid Unicode characters of length 1-1024 bytes when UTF-8 encoded. The following characters can be used in key names but can't be searched in a prefix query string: `!`, `"`, `$`, `%`, `&`, `(`, `)`, `*`, `+`, `,`, `/`, `\`, `:`, `<`, `=`, `>`, `@`, `[`, `]`, ```, `{`, `}`.
- Keys cannot use [Unicode characters](#) 0 through 32, 65534, and 65535 as single-character key names.
- The following characters cannot be used in key names:
 - `#`, `;`, `?`, `^`, and `|`
 - line feed (LF) characters
 - carriage-return (CR) characters

Secret Store

Secret Store is included with your Compute service with the following limitations:

Item	Limit	Implications
Secrets (total across all stores)	10	Paid accounts can purchase additional secrets for a monthly fee.
Maximum secret size	64 KB	An error informs you that you've exceeded the limit.
Rate operations	5 reads per Compute request	An error informs you that you've exceeded the limit.

* * *



Fastly product lifecycle

 Last updated: 2022-08-09

 </products/fastly-product-lifecycle>

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

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.

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 get up our

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 [Fastly's terms and conditions](#).

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 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.
- [Integrated third-party software or services](#) 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.

* * *



How we calculate your delivery bill



Last updated: 2023-12-12



</products/how-we-calculate-your-delivery-bill>

Your bill for content delivery and streaming delivery services is affected by a combination of things including the actual traffic Fastly has served on your behalf, the products you've purchased, the features you've enabled, and the specific configuration settings you've chosen (like enabling [compression](#)).

NOTE

Billing limits for this product may be different depending on your [account type](#), if you've purchased a [packaged offering](#), or are using a [product or feature trial](#).

We charge for egress traffic from our POPs, including traffic served to end users and, if [shielding](#) is enabled, traffic served from the shield POP to other POPs. Specifically, we charge for each response and for the size of the response (which includes the header and body). Each response is billed as a single request, and the response size in bytes is billed as bandwidth. We charge for bandwidth and requests for content delivered to clients from the CDN and for bandwidth for traffic sent from the CDN to our customers' origins.

NOTE

If you're using [Anycast IP addresses](#), these IPs use our global network and will route a request to the nearest [POP](#) located in a billing region that may charge a higher rate. Our billing regions can be found on the [Fastly Pricing](#) page. We announce new billing regions regularly via our [service status page](#).

Charges for any options you've chosen are applied in addition to the bandwidth and request usage we charge for normal content delivery and streaming and may be adjusted according to any [packaged offering](#) you've purchased. In accordance with local laws, Fastly may also collect sales tax based on your shipping or billing address on file.

About the measurements and calculations we use

We measure months according to Coordinated Universal Time (UTC). For usage-based charges, bandwidth is recorded in bytes and presented in gigabytes (GB), and requests are recorded individually and presented in units of 10,000.

Fastly uses [The International System of Units](#) (SI Units) to measure bandwidth. In our calculations, 1 gigabyte (GB) = 10^9 (1,000,000,000) bytes, 1 terabyte (TB) = 10^{12} bytes (or 1,000 GB), and 1 petabyte (PB) = 10^{15} bytes (or 1,000 TB). Your [invoice](#) shows your usage

and that matches the usage shown on the [Observability](#) page.

* * *



Network services resource limits



Last updated: 2026-01-13



</products/network-services-resource-limits>

This guide details Fastly resource limits for [network services](#) products and features based on your account type. These limits change based on whether you're using a free account or have purchased one of our [packaged offerings](#). These limits may be further defined on your service order.

Cache limits

The cache limits for your account depend on when you became a Fastly customer.

Account created on or after June 17, 2020

If you created your account on or after June 17, 2020, the following cache limits apply.

Item	Limit	Implications
Cache file size (with Segmented Caching enabled)	unlimited	None
Cache file size (without Segmented Caching enabled)	20MB	Exceeding this limit when trying to cache a file results in a <code>503 Response object too large</code> error.

Account created prior to June 17, 2020

If you created your account prior to June 17, 2020, the following cache limits apply.

Item	Limit	Implications
Cache file size (with Segmented Caching enabled)	unlimited	None
Cache file size (with streaming miss and without Segmented Caching enabled)	5GB	Exceeding this limit when trying to cache a file results in a <code>503 Response object too large</code> error unless Segmented Caching is enabled.
Cache file size (without streaming miss and without Segmented Caching enabled)	2GB	Exceeding this limit when trying to cache a file results in a <code>503 Response object too large</code> error unless Segmented Caching is enabled.

Rate and time limits

Item	Limit	Implications
API rate limit	1000 requests/hour	Exceeding this limit results in a <code>Too many requests</code> error. The limit is applied to the authenticated user making the request. See API rate limiting for more info.
TLS connections limit	10 minutes	Exceeding this limit results in a <code>502 gateway timeout</code> error.

Request and response limits

Item	Limit	Implications
URL size	8KB	Exceeding the limit results in a <code>414 URI Too Long</code> error.
Cookie size	32KB	Exceeding the limit results in Fastly stripping the cookie and setting <code>req.http.Fastly-Cookie-Overflow = "1"</code> .
Maximum		

maximum response size for tarpitting	4KB	Exceeding the limit results in no error. Tarpitting isn't applied to responses that exceed the limit.
Request header size	128KB	Depending on the circumstances, exceeding the limit can result in Fastly closing the client connection abruptly, or it can result in the client either receiving a <code>502 Gateway Error</code> response with an <code>I/O error</code> in the body, or receiving a <code>503 Service Unavailable</code> response with a <code>Header overflow</code> error in the body.
Response header size	128KB	Exceeding the limit results in a <code>503 backend read error</code> . Read our developer information about errors for more info.
Request header count	255	Exceeding the limit results in a <code>Header overflow</code> error. A small portion of this limit is reserved for internal Fastly use, making the practical limit closer to 200.
Response header count	96	Exceeding the limit results in a <code>Header overflow</code> error. A small portion of this limit is reserved for internal Fastly use, making the practical limit closer to 85.
<code>req.body</code> size	8KB	Exceeding the limit results in the <code>req.body</code> variable being blank. Request body payload is available in <code>req.body</code> only for payloads smaller than 8KB. <code>req.postbody</code> is an alias for <code>req.body</code> .
Surrogate key size	1KB	Exceeding the limit results in purging API failures stating "surrogate key too long, must be less than 1024 bytes." Any keys that exceed the limit will be dropped instead of truncated.
Surrogate key header size	16KB	Exceeding the limit results in no error and any keys past the one that exceeds the limit will be dropped.
Idle client reuse time	10 min	Client connections cannot be used after this period of idle and must be reestablished.

Service, domain, and origin limits

Item	Limit	Implications
------	-------	--------------

Services total per account	10	Exceeding this limit results in an <code>Exceeding max_total_services</code> error. Contact support to discuss raising this limit.
Origins per service	5	Exceeding this limit results in an <code>Exceeding max_backends</code> error. Contact support to discuss raising this limit.
Classic domains per service	20	Exceeding this limit results in an <code>Exceeding max number of domains</code> error. Contact support to discuss raising this limit.
Connections per service	200	Exceeding this limit results in an <code>Error 503 backend.max_conn reached</code> error. You can increase this limit as high as 1000 by updating the backend connection setting to limit the connections a single Fastly cache server will make to a specific origin server.

VCL and configuration limits

Item	Limit	Implications
Custom VCL file size	1MB	Exceeding the limit results in a <code>Content too long</code> error.
Maximum VCL file size	3MB	Exceeding the limit results in a <code>VCL is too long</code> error.
Varnish restart limit	3 restarts	Exceeding the limit results in a <code>Service Unavailable</code> error. This limit exists to prevent infinite loops.
ACL container entries count	1000	Exceeding the limit results in an <code>Exceeding max ACL entries</code> error. Contact support to discuss raising this limit.
Dictionary items count	1000	Exceeding the limit results in an <code>Exceeding max dictionary items</code> error. Contact support to discuss raising this limit.

Dictionary item key length	255 characters	Exceeding the limit results in an <code>Item key is too long</code> error.
Dictionary item value length	8000 characters	Exceeding the limit results in an <code>Item value cannot be greater than</code> error.
Log line size	16KB Deliver, 64KB Compute	Exceeding the limit results in truncated logs that could result in delivery errors.
Service chains and hops	20 hops, 6 unique services	Exceeding the limit results in a <code>Service unavailable</code> error with the <code>Loop detected</code> response reason text. Read our developer information about loop detection for more info.
Synthetic response characters	No character limit	Synthetic responses have no character limit, but large responses may trigger an error for the custom VCL file size limit.
Vary objects count	50 soft, 70 hard	Exceeding the soft limit results in no error. Newer variants displace the oldest. Active fetches from backends are limited to 70 variants. Exceeding this hard limit results in a <code>Too many variants</code> response. Once fetches complete, objects will be removed until the soft limit is reached. Image Optimizer transformations count towards the variant limit for the original image.

Varnish workspace limits

A Varnish workspace is an allotted amount of memory that's dedicated to each request process. It's used as temporary storage for objects like headers and local variables, as well as to execute some VCL functions.

Each request process has 256K of workspace per cache node to work with. Subroutines that run on the delivery node (e.g., `recv`, `hash`, `deliver`, `log`, and `error`) have 256K of workspace to use for each individual request process, while the subroutines that run on the fetch node (e.g., `miss`, `hit`, `pass`, and `fetch`) also have 256K of workspace to

use.

You can check how much workspace is being used for each request by using the `workspace.bytes_total` and `workspace.bytes_free` variables in any subroutine.

If a request process exceeds the workspace limit, Varnish will return a `503 Service Unavailable` or `503 Header Overflow` error.

* * *



Next-Gen WAF resource limits



Last updated: 2026-01-13



</products/next-gen-waf-resource-limits>

This guide details the Next-Gen WAF resource limits we set based on your platform. These limits change based on whether you're using a free account or have purchased one of our [packaged offerings](#). These limits may be further defined on your service order.

NOTE

In addition to the limits described here, you may also encounter resource limits related to the [Network services](#) or [Compute](#) products and features on your account.

Agent response codes

Item	Limit	Implication
Unique custom agent response codes	5 per site (also known as workspace)	An error informs you that you've reached the limit.
Rules that use custom agent response codes	1000 per site (workspace)	An error informs you that you've reached the limit.

Lists

Item	Limit	Implication
Lists	25 per corp (also known as account) plus 25 per site (workspace)	An error informs you that you've reached the limit.
Items per list	5000 items	An error informs you that you've reached the limit.

Redactions

Item	Limit	Implication
Custom redactions	100 per site (workspace)	An error informs you that you've reached the limit.

Rules

Item	Limit	Implication
Request rules	1000 per corp (account) plus 1000 per site (workspace)	An error informs you that you've reached the limit.
Signal exclusion rules	1000 per corp (account) plus 1000 per site (workspace). Signal exclusion rules count against the request rules limits.	An error informs you that you've reached the limit.
Advanced rate limit rules	15 per site (workspace)	An error informs you that you've reached the limit.

Signals

Item	Limit	Implication
Custom signals	200 per corp (account) plus 200 per site (workspace)	An error informs you that you've reached the limit.

Thresholds

Item	Limit	Implication
Site alerts (signal thresholds)	50 per site (workspace)	An error informs you that you've reached the limit.

* * *



Summary product definitions



Last updated: 2025-11-04



</products/summary-product-definitions>

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

AI Accelerator

[AI Accelerator](#) is a semantic caching solution for large language model (LLM) APIs used in generative artificial intelligence (AI) applications.

API Discovery

Fastly's [API Discovery](#) continuously finds, organizes, and updates the record of APIs passing through the Fastly edge network to your applications. It helps you monitor API changes, understand your API attack surface area, and prioritize mitigation efforts.

Application Programming Interface

Fastly provides an [application programming interface \(API\)](#) that can be accessed via a number of popular [interactive clients](#) and allows you to manage Fastly services via remote procedure calls. These services include features such as [authentication](#), [configuration](#), [historical stats](#), [purging](#), and [remote logging](#). In addition to being accessible via Fastly's API, Fastly services can also be accessed via a control panel for

users with the appropriate [access permissions](#); however, API features do not include customer account setup, which can only occur through the [web interface controls](#).

Assurance Services

[Assurance Services](#) 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.

Bot Management

Fastly's [Bot Management](#) product provides you with visibility into bot traffic, allowing you to identify bots and automations directly at Fastly's network edge, closer to where requests arrive and further away from your application layer.

Cache Reservation

[Cache Reservation](#) provides access to the caching layer at Fastly's edge where you can reserve cache space specifically for your content in Fastly shielding locations.

Capacity Reservation

[Capacity Reservation](#) allows you to reserve Fastly traffic capacity for events.

Certainly

[Certainly](#) is Fastly's publicly-trusted certification authority (CA) that generates Transport Layer Security (TLS) certificates to provide trusted identification of secured websites. Certainly is only available for use with Fastly TLS.

Cloud Optimizer

Limited Availability

Fastly's [Cloud Optimizer](#) 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

Fastly's [Compute](#) platform, formerly known as Compute@Edge, 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 to write new applications, move critical logic closer to end users, and bring enhanced security and performance to current workflows.

Concierge TLS

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

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. [Support availability and response times](#) vary depending on the level of support you have purchased.

Dedicated IP addresses

Fastly's [Dedicated Internet Protocol \(IP\) addresses](#) provide you with a pool of IP addresses across Fastly's global Edge Cloud.

DDoS Protection and Mitigation Service

Fastly's [DDoS Protection and Mitigation Service](#) helps protect against volumetric and targeted distributed denial of service attacks against origin servers. It provides coverage 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 Research API

Fastly's [Domain Research API](#) allows you to programmatically retrieve algorithmic domain

search results and check domain availability.

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 [fully qualified domain name](#) without requiring you to send log data to a third-party data collector.

Edge Data Storage

Compute gives you the option of [storing the data you need at the edge](#) as key-value pairs in versionless containers. You can attach these containers to an active service and update the data at any time after it's created, without ever incrementing a service's version. Additionally, these containers can be shared by multiple Compute services in your account.

Edge Rate Limiting

Fastly's [Edge Rate Limiting](#) product helps you control the rate of requests sent to your origin servers.

Fanout

Fastly [Fanout](#) is a publish-subscribe message broker operating at the Fastly edge. It provides customers with the ability to push messages via direct connections that remain open indefinitely rather than requiring clients to poll for data using short-lived HTTPS requests.

Fastly Client-Side Protection

Fastly [Client-Side Protection](#) provides you with the ability to inventory and control the resources (e.g., scripts, images, and fonts) that load on an end user's browser from defined areas of your web applications by building and enforcing [content security policies](#). When a resource violates your content security policy, the end user's browser blocks or logs the resource per your selected protection mode. Based on policy violation reports, you can adjust your content security policies as needed. In addition, you can provide a justification as to why each client-side script is or isn't allowed. These capabilities help you guard against cross-site scripting attacks (e.g., Magecart attack)

and enable you to maintain compliance with Payment Card Industry Data Security Standard (PCI DSS) [4.0.1 - Sections 6.4.3 and 11.6.1](#).

Fastly DDoS Protection

[Fastly DDoS Protection](#) provides real-time visibility into and defense against common Distributed Denial of Service (DDoS) attacks directed at your applications, APIs, and origin servers. It provides automatic detection and mitigation at Fastly's network edge, away from your infrastructure.

Fastly DNS

[Fastly DNS](#) is our authoritative DNS service that leverages Fastly's global, high-performance edge network to serve DNS query responses.

Fastly TLS

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

High Volume Logging

Fastly's [High Volume Logging](#) 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

Fastly offers a [HIPAA-Compliant](#) Caching and Delivery product that allows you to transmit protected information like protected health information through Fastly's network.

Image Optimizer

Fastly's [Image Optimizer](#) 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 [Live Event Monitoring service](#) 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.

Log Explorer & Insights

The [Log Explorer & Insights](#) feature allows you to proactively store, inspect, and monitor your log data on the Fastly Platform. We collect, store, and analyze request and response-related data to provide aggregated metrics and log data.

Logging Insights Package

Fastly's [Logging Insights Package](#) 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.

Managed Security Enterprise

Fastly offers [Fastly Next-Gen WAF](#) customers a [Managed Security Enterprise](#) service that provides your organization with continuous security product monitoring, proactive response to identified security events, 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.

Managed Security Professional

[Managed Security Professional](#) provides your organization with continuous monitoring of your included Fastly security products, proactive action by Fastly in the event of an identified security incident, enhanced access to our Customer Security Operations Center (CSOC), and periodic consultation with a Designated Security Specialist for strategic security solutions reviews and planning.

Media Shield for Live

Fastly's [Media Shield for Live](#) 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 [Media Shield for VOD](#) 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.

Next-Gen WAF

The [Fastly Next-Gen WAF](#) 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.

Object Storage

[Fastly Object Storage](#) is an S3-compatible large object storage solution that works seamlessly with Fastly services. Using Fastly Object Storage, you can store larger file sizes with Fastly, improving latency, increasing cache hit ratios, and reducing egress charges.

Oblivious HTTP Relay

The Fastly [Oblivious HTTP Relay](#) (OHTTP Relay) implements the relay portion of the [Oblivious HTTP specification](#), which allows you to create an OHTTP-compliant service using Fastly.

Origin Connect

Fastly's [Origin Connect](#) product offers a private network interconnect between your origin servers and your Fastly shield POP.

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.

Partner Support Services

Fastly partners are required to purchase [Partner Support Services](#). These services provide Fastly platform support benefits as well as a library of on-demand training modules.

PCI-Compliant Caching and Delivery

Fastly offers a [PCI-Compliant Caching and Delivery](#) product that allows you to transmit protected information like cardholder data through Fastly's network.

Performance Optimization Package

Fastly's [Performance Optimization Package](#) 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 [Platform TLS](#) 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

Fastly offers [Fastly Next-Gen WAF](#) customers a [Response Security](#) service 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.

Security Technical Account Manager

Fastly offers customers the ability to purchase the support of a [Security Technical Account Manager](#) for your organization. These specialists help you optimize your use of Fastly's security products and features by providing proactive check-ins and regular reviews to help you analyze your account's security service configurations and their performance. Security Technical Account Managers also provide enhanced troubleshooting coordination with Fastly's support and professional services organizations.

Staging

Fastly's [Staging](#) feature allows customers to test changes to Delivery and Compute service configurations on a staging network before deploying them to the production network.

Streaming Delivery

Fastly's [Streaming Delivery](#) 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 [Full-Site Delivery](#), 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 [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.

Technical Account Manager

Fastly offers the ability to purchase the support of a Fastly Engineer to serve as a [Technical Account Manager](#) 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.

WebSockets

Fastly supports the use of the [WebSocket protocol](#). This protocol allows you to establish long-lived, two-way, interactive communication sessions between clients and servers.

* * *

Category: [Platform features](#)

These articles provide information about Fastly's platform features



Staging



Last updated: 2026-07-07



</products/staging>

Fastly's [Staging feature](#) allows customers to test changes to Delivery and Compute service configurations on a staging environment before deploying them to the production network.

You can stage a service version using the Fastly control panel, API, or CLI. You must modify your computer's `hosts` file or create a new DNS record to access a staged version of your service.

Limitations and considerations

Fastly's Staging feature has the following limitations:

- To stage a service, you must deploy at least one production version of it.
- The Fastly control panel and API don't indicate whether or not changes have been made to a staged service version.
- Changes to domains and TLS configurations can't be staged.
- There's no shared cache between the staging and production environments.
- Changes made to versionless objects and settings, like VCL snippets and Fastly Image Optimizer, will be reflected in both the staging and production environments.
- Fastly's [Subscriber Provided Prefix](#) has unique requirements that must be met before working with the Staging feature. [Contact support](#) to see if your organization meets those requirements.
- Staging environments are assigned IPv4 addresses. IPv6 addresses are not supported.
- The IP address for your staging environment could change if you modify your TLS certificates. Always refer to the Fastly control panel to find the current IP address for your staging environment.
- [Service chaining](#) is not supported.
- If your service is configured to use Next-Gen WAF, traffic to your staging environment will be routed through the production version of your Next-Gen WAF (e.g., requests can be logged and blocked). Changes to Next-Gen WAF configuration settings cannot be staged because Next-Gen WAF is *versionless*. This means that any changes you make to your Next-Gen WAF configuration settings in the staging service become effective immediately for all service versions, including the production version.

You can't stage services that use the following products and features:

- [Fastly's Managed CDN \(mCDN\)](#)
- [Fastly's On-the-Fly Packager \(OTFP\)](#)

Stats and logging

Traffic to your staging environment can be logged separately from your production service. Stats and metrics—such as hits, misses, errors and response time—will not count towards the stats of your production version.

Billing

Traffic to staging services is not distinguished from traffic for production services for billing purposes and will be included with production traffic on your invoice. Consider setting a [spend alert](#) for your account to automatically be notified when your month-to-date bill approaches or meets a specific dollar amount. Refer to the pricing and billing information for any products used with the Staging feature.

* * *

Category: [Legacy offerings](#)

These articles provide information about Fastly's legacy offerings.



Next-Gen WAF legacy support description and SLA



Last updated: 2023-04-05



</products/fastly-next-gen-waf-legacy-support-description-and-sla>

IMPORTANT

These terms apply to customers who purchased [Fastly's Next-Gen WAF](#) (previously referred to as the Signal Sciences Cloud WAF and Signal Sciences

Next-Gen WAF) on or before April 5, 2023. These subscriptions will continue to be supported for all existing users until the end of their current contracts. For more information about our [current support description and service level agreements for Next-Gen WAF](#), contact sales@fastly.com.

Support channels and response times

Fastly provides customer support via the support channels listed below.

Support tickets

Fastly Next-Gen WAF 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 support@fastly.com 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 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.
- **"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 control panel. 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 control panel. 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 ([Hosted Dashboard](#)), 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: $([\text{\# of minutes in month}] - [\text{\# of minutes per month the Hosted Dashboard is Unavailable}]) / [\text{\# of minutes in month}]$.

"**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, (b) you or your end user's connectivity issues, (c) your operator errors, (d) traffic amounts that exceed your Permitted Utilization as defined in the Terms and Conditions, (e) your corrupted content, (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
---	---

Hosted Services	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: $([\text{\# of minutes in month}] - [\text{\# of minutes per month the Hosted Services is Unavailable}]) / [\text{\# of minutes in month}]$.

"**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, (b) you or your end user's connectivity issues, (c) your operator errors, (d) traffic amounts that exceed your Permitted Utilization as defined in the Terms and Conditions, (e) your corrupted content, (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 [Fastly Next-Gen WAF \(Edge\)](#) 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 [fail open](#) 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's Legacy Full-site Delivery services



Last updated: 2019-08-09



</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 [Full-Site Delivery](#), contact sales@fastly.com.

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.

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 [load balancing](#) 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, [request collapsing](#) 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 [actively invalidate content](#). 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, [configurable purging commands](#). Fastly currently supports [Ruby](#), [Python](#), [PHP](#), and [Perl libraries](#) 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 [perform health checks](#) 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




Last updated: 2019-09-12




</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 [reducing multiple CDN requests](#) 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 customer success manager or email sales@fastly.com for more details.

IMPORTANT

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

* * *

Legacy Customer-Provided TLS Certificate Hosting Service



Last updated: 2020-04-01



</products/legacy-customer-provided-tls-certificate-hosting-service>

NOTE

Fastly maintains support for its original Customer-Provided TLS Certificate Hosting Service. For more information about our current [TLS service options](#), contact sales@fastly.com.

Fastly offers a Customer-Provided TLS Certificate Hosting Service where you provide

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.

* * *



Legacy Platinum Support and SLA



Last updated: 2018-03-30



</products/legacy-platinum-support-and-sla>

NOTE

Fastly maintains support for its original Platinum Support plan. For more information about our current [Gold and Enterprise Support plans](#) or for information about our [Professional Services packages](#), contact sales@fastly.com.

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) your hardware or software failures, (b) you or your end user's connectivity issues, (c) your operator errors, (d) traffic amounts that exceed your Permitted Utilization as defined in the Terms and Conditions, (e) your corrupted 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>.

Support requests

Subscribers submit support requests by visiting <https://support.fastly.com/> or 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 [Service Availability SLA](#)),
- 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

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 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 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 [type of account](#) you have and the [level of support](#) 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.
- **"Fastly Control"** means elements entirely under Fastly's control and not a consequence of (a) your hardware or software failures, (b) you or your end user's connectivity issues, (c) your operator errors, (d) traffic amounts that exceed your Permitted Utilization as defined in the Terms and Conditions, (e) your corrupted 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



</products/legacy-premium-support-and-sla>

NOTE

Fastly maintains support for its original Premium Support plan. For more information about our current [Gold and Enterprise support plans](#) or for information about our [Professional Services packages](#), contact sales@fastly.com.

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:

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) your hardware or software failures, (b) you or your end user's connectivity issues, (c) your operator errors, (d) traffic amounts that exceed your Permitted Utilization as defined in the Terms and Conditions, (e) your corrupted 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 visiting <https://support.fastly.com/> or 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),
- 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 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 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

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

- 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 [type of account](#) you have and the [level of support](#) 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) your hardware or software failures, (b) you or your end user's connectivity issues, (c) your operator errors, (d) traffic amounts that exceed your Permitted Utilization as defined in the Terms and Conditions, (e) your corrupted 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.

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.

* * *



Legacy Professional Services



Last updated: 2023-10-30



</products/legacy-professional-services>

NOTE

Fastly maintains support for its original Service Implementation, Service Management, and Consulting Engagement Professional Services offerings. For more information about our current [Professional Services](#) offering, contact sales@fastly.com.

Legacy Service Implementation

For projects with a predefined scope, you can purchase our Service Implementation

For projects with a predefined scope, you can purchase our Service Implementation offering. You'll work with one of our experts to help you with tasks related to your service implementation and initial setup and configuration, like onboarding for Delivery or Compute services.

Legacy Service Management

For customers who require ongoing configuration and technical assistance, Fastly offers Service Management that provides 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
- [Service](#) 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
- Edge authentication
- ESI (edge side includes) configuration
- Streaming and video packaging

Legacy 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

* * *



Legacy shared TLS domain



Last updated: 2024-04-18



</products/legacy-shared-tls>

NOTE

This page describes a legacy offering. To secure your domains with Fastly TLS, we recommend exploring one of our other [TLS service options](#).

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

This TLS option is subject to the following limitations:

- This option uses a shared domain name and does not support use with your own domain name (`www.example.com`). Customers typically use this TLS option when linking directly to assets (e.g., linking to `https://`

`example.global.ssl.fastly.net/example.jpg`) or for testing purposes.

- You cannot DNS alias your own domain to the shared domain. If you do, a TLS name mismatch warning will appear in the browser.
- When using this TLS option, all traffic is routed through Fastly's entire global network.

If you have specific traffic routing, domain naming, or URL requirements, one of Fastly's other [TLS service options](#) can provide you with more flexibility.

* * *



Legacy TAM and Strategic Essentials



Last updated: 2023-01-09



</products/legacy-tam-and-strategic-essentials>

NOTE

As of January 9, 2023, Fastly's TAM Essentials and Strategic Essentials subscriptions within our Fastly Next-Gen WAF professional services offering became legacy products. These subscriptions will continue to be supported for all existing users until the end of their current contracts. As an alternative, our [Security Technical Account Manager](#) offering provides similar services. Contact sales@fastly.com or your Fastly account team to evaluate or move to this Fastly Next-Gen WAF professional services option.

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.

* * *

Category: **Third-party information**

These articles provide information about third-party technology and services incorporated into Fastly's products and services.



Open source software in downloadable components



Last updated: 2026-04-13



</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.

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
------------	---------	---------

<code>apimachinery</code>	v0.35.3	Apache-2.0
<code>archiver</code>	v4.0.0-alpha.9	MIT
<code>AWS SDK for Go</code>	v1.55.6	Apache-2.0
<code>AWS SDK for Go v2</code>	v1.32.13	Apache-2.0
<code>client-go</code>	v0.35.3	Apache-2.0
<code>Cobra</code>	v1.10.2	Apache-2.0
<code>control-plane</code>	v1.36.0	Apache-2.0
<code>Datadog Go</code>	v2.2.0+incompatible	MIT
<code>Elastic Load Balancing</code>	v1.33.23	Apache-2.0
<code>exp</code>	v0.0.0-20230905200255-921286631fa9	NewBSD
<code>gabs</code>	v2.7.0	MIT
<code>glob.go</code>	v0.2.3	MIT
<code>Go App Engine</code>	v1.6.8	Apache-2.0
<code>Go JOSE</code>	v4.1.3	Apache-2.0
<code>Go Networking</code>	v0.49.0	NewBSD
<code>Go support for Protocol Buffers</code>	v1.36.10	NewBSD
<code>Go support for Protocol Buffers</code>	v1.5.4	NewBSD
<code>Go Sync</code>	v0.19.0	NewBSD
<code>Go Text</code>	v0.35.0	NewBSD
<code>go-codec</code>	v1.2.7	MIT
<code>go-diff</code>	v1.0.1-0.20180205163309-da645544ed44	MIT

<code>go-license</code>	<code>v0.0.0-20180405065157-c69f41c2c8d6</code>	MIT
<code>go-reuseport</code>	<code>v0.0.1</code>	ISC
<code>gohistogram - Histograms in Go</code>	<code>v1.0.0</code>	MIT
<code>Google APIs Client Library for Go</code>	<code>v0.273.0</code>	NewBSD
<code>Google Compute Engine API</code>	<code>v1.57.0</code>	Apache-2.0
<code>Gorilla WebSocket</code>	<code>v1.5.3</code>	FreeBSD
<code>gorilla/mux</code>	<code>v1.8.1</code>	NewBSD
<code>GoVersionInfo</code>	<code>v0.0.0-20190209210621-63e6d1acd3dd</code>	MIT
<code>gRPC Proxy</code>	<code>v0.0.0-20181017164139-0f1106ef9c76</code>	Apache-2.0
<code>gRPC-Go</code>	<code>v1.79.3</code>	Apache-2.0
<code>kin-openapi</code>	<code>v0.131.0</code>	MIT
<code>Kubernetes API</code>	<code>v0.35.3</code>	Apache-2.0
<code>Kubernetes controller-runtime</code>	<code>v0.16.3</code>	Apache-2.0
<code>Logrus</code>	<code>v1.9.3</code>	MIT
<code>mapstructure</code>	<code>v1.1.2</code>	MIT
<code>MessagePack Code Generator</code>	<code>v1.2.4</code>	MIT
<code>mongo-driver</code>	<code>v1.9.1</code>	Apache-2.0
<code>Oxy</code>	<code>v1.4.2</code>	Apache-2.0
<code>Package for equality of Go values</code>	<code>v0.7.0</code>	NewBSD

<code>pflag</code>	v1.0.10	NewBSD
<code>profile</code>	v1.7.0	FreeBSD
<code>psutil for Go</code>	v4.25.8	NewBSD
<code>PTerm</code>	v0.12.83	MIT
<code>snappy</code>	v0.0.1	NewBSD
<code>supports-color</code>	v1.2.0	MIT
<code>sys</code>	v0.42.0	NewBSD
<code>Testify</code>	v1.11.1	MIT
<code>Versioning Library for Go</code>	v1.2.1	MPL-2.0
<code>Viper</code>	v1.7.0	MIT
<code>YAML support for the Go language</code>	v3.0.1	MIT

* * *



Sub-processors



Last updated: 2026-05-05



</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 types, 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, read our [data processing terms](#). You 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](#).

Fastly affiliates

These sub-processors serve as processor affiliates for Fastly. For more information about these affiliates, such as privacy concerns, email gc@fastly.com for assistance.

Name	Entity Location	Processing Location
Fastly Australia Pty Ltd	Tricor Chew and Dormers Pty Ltd Level 3 1049 Victoria Road West Ryde, NSW 2114	Australia
Fastly Canada Inc.	2900 - 733 Seymour Street, PO Box 1, Vancouver, BC V6B OS6	Canada
Fastly France SAS	31-35 Rue De La Fédération, 75015 Paris, France	France
Fastly GmbH	c/o Rueter & Partner, Prielmayerstraße 3, 80335 München, Germany	Germany
Fastly India Private Limited	56/57 Mittal Chambers 228 Barrister Rajani Patel Marg, Nariman Point Mumbai 400-021	India
Fastly Kabushiki Kaisha	Tsao Hibiya 8F, 3-3-13 Shinbashi, Minato-ku, Tokyo 105-0004 Japan	Japan
Fastly Netherlands B.V.	Markt 14, 6211CJ Maastricht, Netherlands	Netherlands
Fastly International (Holdings) Limited	Birchin Court, 5th Floor, 19-25 Birchin Lane, London, England, EC3V 9DU	United Kingdom
Fastly Limited	Birchin Court, 5th Floor, 19-25 Birchin Lane, London, England, EC3V 9DU	United Kingdom
Fastly SG Pte. Ltd.	101 Thomson Road #14-02/03 United	Singapore

Square 307591 Singapore		
Fastly Cloud Iberica, S.L.	Calle Marques De Urquijo 26, Planta 3, Puerta IZ, 28008 Madrid, Spain	Spain
Fastly Sweden AB	c/o Svalner Skatt & Transaktion Box 16115, 103 23 Stockholm, Sweden	Sweden

Third-party sub-processors

These sub-processors serve as third-party service providers engaged by Fastly.

Name and Entity Location	Processing Locations	Purpose	Sub-processor Resources
Google LLC 1600 Amphitheatre Parkway, Mountain View CA 94043	Australia, Belgium, Canada, Germany, Japan, Netherlands, United Kingdom, United States	Fastly uses Google's services for backend data hosting, data insights, and observability. Products that use Google's services include Image Optimizer, AI Accelerator, KV Store, Secret Store, Observability features, and event monitoring services.	Google Resources
Amazon Web Services, Inc. 410 Terry Avenue North, Seattle, WA 98109-5210	Australia, Brazil, Canada, France, Germany, India, Ireland, Japan, Singapore, South Africa, South Korea, Sweden, United Kingdom, United States	Fastly uses AWS's services for backend data hosting and observability. Products that use AWS's services are the Next-Gen WAF, Bot Management and the Glitch Platform. AWS is also used for customer dashboards and audit logs, and in the	AWS Resources

		provisioning and nosting of the control plane.	
MongoDB Atlas	United States	Fastly uses MongoDB Atlas' services for backend data hosting. Products that use MongoDB Atlas' services include Next-Gen WAF and its associated services.	MongoDB Atlas Resources
1633 Broadway, 38th Floor New York NY 10019			
Backblaze, Inc.	Netherlands, United States	Fastly uses Backblaze's services for data hosting. The product that uses Backblaze's services is Fastly Object Storage.	Backblaze Resources
201 Baldwin Ave, San Mateo, California, 94401			
Wasabi Technologies, LLC	Australia, France, Italy, Japan, United Kingdom, United States	Fastly uses Wasabi Technologies services for data hosting. The product that uses Wasabi Technologies' services is Fastly Object Storage.	Wasabi Resources
75 Arlington Street, Suite 810, Boston, MA 02116			

* * *



Third-party technology



Last updated: 2026-03-19



</products/third-party-technology>

These articles provide information about third-party technology and services incorporated into Fastly's products and services:

- [Geolocation VCL features](#)
- [Device detection variables](#)
- [TLS service options](#)
- [Next-Gen WAF geolocation data](#)

In addition, these articles describe third-party services that interoperate with Fastly's products and services:

- [About the AI Accelerator page](#)
- [Integrations with Non-Fastly Services](#)
- [Streaming logs](#)

* * *

Category: **Compliance**

These articles provide information about the administrative, physical, and technical safeguards that protect Fastly's products and services.



Compliance processes and safeguards



Last updated: 2025-12-09



</products/compliance-processes-and-safeguards>

These articles provide information about the processes and safeguards we follow to protect your data:

- [EU Data Act Addendum](#)
- [HIPAA-Compliant Caching and Delivery](#)

PCI-Compliant Caching and Delivery

- [PCI-Compliant Caching and Delivery](#)
- [Security program](#)

* * *



Security Measures



Last updated: 2024-12-11



</products/security-measures>

These articles provide information about the administrative, physical, and technical safeguards that protect Fastly's products and services.

- [Data Management](#)
- [Security program](#)

Reporting discovered security issues

We take the security of our network seriously and support the disclosure of security issues related to our service. If you believe you have found a vulnerability, we encourage you to [report your discovery](#) to our Support team so we can investigate further. If you plan to do security testing of your service behind Fastly, notify Fastly at least two (2) business days prior to the test. See our [security testing guidelines](#) for more information.

Related features

- [Access control lists](#)
- [Configuring user roles and permissions](#)
- [Cryptographic VCL features](#)
- [Managing two-factor authentication](#)

... " ... VCL ...

- [Miscellaneous VCL features](#)
- [Monitoring account activity with the audit log](#)
- [Monitoring account activity with the event log](#)
- [HUMAN Bot Defender](#)
- [Streaming logs](#)
- [TLS](#)

* * *

Category: [Security measures](#)

These articles provide information about the administrative, physical, and technical safeguards that protect Fastly's product and services.



Data management



Last updated: 2026-07-08



</products/data-management>

Fastly maintains a “privacy and protection by design” approach as part of its data governance program. Fastly is intentional about data processing and implements safeguards designed to limit and control access to customer and Fastly data. Fastly limits data collection to what is needed to provide, maintain, and improve its services and considers legal, compliance, regulatory, and commercial obligations when working with data.

Fastly production data

In addition to the Fastly [security program](#), Fastly implements its data governance program with practices designed to manage data in its production environments as follows:

Network data

- *Service management:* Fastly monitors, collects, and processes data related to the functional performance of Fastly's network and services, anomalous activity, and suspicious behavior detected by the services. Fastly retains and uses this data to monitor, maintain, and improve its services, business operations, and security and compliance programs.
- *Confidentiality:* Fastly only discloses this data in an anonymized and aggregated form and subject to its confidentiality obligations to customers.

IP addresses

- *Administrative connections:* Fastly may retain any non-anonymized, non-aggregated client or customer IP addresses associated with administrative connections to Fastly's services indefinitely.
- *Fastly application:* Fastly independently collects the IP addresses of users who access services within the Fastly control panel or through the Fastly API.
- *Endpoints:* If a customer defines origin servers or syslog endpoints with IP addresses, Fastly will save those IP addresses as part of the customer's configurations.
- *Client IPs:* Fastly retains client IP addresses in a non-anonymized, non-aggregated fashion for up to two business days, or up to seven days if those addresses are associated with transmission errors.
- *Origin IPs:* Fastly may retain dynamically-resolved origin IP addresses for up to two business days, or up to seven days if associated with transmission errors.
- *Suspicious activity:* Internal systems logs, including access logs and associated client and customer IP addresses, related to events triggered by anomalous activity or suspicious behavior are retained for at least one year. Fastly may retain IP addresses from Fastly event logs or configurations indefinitely.
- *Security events:* Fastly may retain non-anonymized, non-aggregated client or customer IP addresses associated with security-related incidents indefinitely.

Customer data management

- *Customer content:* Customer content enters, transits, and departs Fastly's network

in response to requests. Generally, customers manage which content is processed, where, and for how long by setting policies that control that content.

- *Customer configurations:* Customer configurations may be stored indefinitely and can be deleted upon request. Fastly may directly access or modify customer accounts or configurations as necessary to provide and improve the services, to prevent or address service or technical issues, as required by law, or as customers expressly permit. Fastly retains encrypted backups of customer configurations, including VCL, and customer provided packages for business continuity purposes.
- *Cached and hosted content:* Cached and hosted content is retained per customer configuration and use of purge functionality. Customers may control length and type of retention through configuration options to meet requirements for regulatory reasons such as [HIPAA](#) or [PCI DSS](#). Fastly deletes cached and hosted content according to a customer's use of the purge or deletion functionality as described in the Documentation.
- *Customer data processed by Next-Gen WAF:* Customer data processed via the security components of Next-Gen WAF and sent to Fastly may be stored for a maximum of 30 days (or such shorter period as may be specified in the Documentation). Customer data processed in local environments is under customer control and there is no remote access by Fastly.
- *Customer packages deployed to Compute:* Customer provided compiled code may be stored indefinitely and can be deleted upon request. Fastly may directly access or modify customer packages as necessary to provide and improve services, to prevent or address service or technical issues, as required by law, or as customers expressly permit.

Customer request logs and observability features

- *Content request logs:* Customers may stream their content request logs, which may include request headers, including client IP addresses, to a customer-owned and managed or Fastly-owned and managed endpoint for analysis and use as described in the Documentation.
- *Request logs retention:* Fastly does not retain customers' request logs except where explicitly stated in the Documentation and related to the functional performance of the services
- *Observability features:* Customers may enable observability features to capture and store data as described in the Documentation.

Note regarding privacy law

Fastly services are delivered through a global network. Points of presence (PoPs) can be found on our [network map](#). Fastly employees provide support and technical services from various locations in the United States, Europe, Japan and other countries and regions. For more information regarding Fastly's compliance with global privacy laws and regulations, refer to the [Fastly data processing terms](#), the [list of sub-processors](#), Fastly's [privacy policy](#), and additional resources on the Fastly [Trust page](#).

* * *



Security program



Last updated: 2026-07-08



</products/security-program>

Fastly operates a comprehensive information security program that includes administrative, physical, and technical safeguards to protect its infrastructure, data, services, and customers.

Foundation

Fastly's security program is based on the [NIST Cybersecurity Framework](#) and includes annually reviewed security policies, designated roles and responsibilities, and formal procedures focused on risk.

Security policies

Fastly institutes information security policies that are published internally and reviewed annually. The policies contain principles and point to standards that cover controls and procedures designed to protect Fastly and its customers.

Roles and responsibilities

Fastly maintains designated roles and responsibilities within its security program,

including a Chief Information Security Officer responsible for comprehensive oversight and a team of security professionals responsible for implementing the program.

Risk-based approach

Fastly maintains formal procedures for the identification, assessment, and treatment of information security and availability risks, threats, and vulnerabilities to its services.

Fastly implements its risk-based approach with controls designed to do the following:

- *Annual risk assessment:* Conduct an annual risk assessment to measure the state of security risk across the company. The results of this assessment are shared with the executive leadership team to ensure appropriate visibility and treatment.
- *Risk analysis and mitigation plan:* Evaluate each identified enterprise security risk and associated controls and implement mitigation plans commensurate with the risk to manage the risk to acceptable levels.
- *Risk register:* Maintain documentation of identified risks, threats, and vulnerabilities related to Fastly services. Assigned personnel help assess and remediate identified items in accordance with established risk and vulnerability management procedures.

Defense-in-depth

Fastly's security program is designed with multiple layers of protection. Fastly's processes, technology, and physical security controls are designed specifically to provide a defense-in-depth approach and can be categorized as follows:

Identity and access management

Fastly manages access to its production systems with controls designed to do the following:

- *Authentication:* Require unique user accounts and multi-factor authentication for remote access to production systems by all personnel.
- *Authorization:* Limit access using role-based access controls. Access is granted based on the principle of least privilege and manager approval and is reviewed on a periodic basis. Systems are in place to remove access when no longer needed, including upon personnel separation.
- *Audit:* Maintain and monitor logs of access attempts (both success and failure).

Data security

Fastly manages data security using controls designed to do the following:

- *Customer credentials management:*
 - Secure customer-provided private keys and credentials throughout their lifecycle.
 - Encrypt customer-provided private keys at rest.
 - Re-encrypt customer-provided private keys on a regular interval.
 - Store key encryption keys in a secrets management system.
 - Decrypt private keys in memory at the edge when requested.
 - Remove private keys from memory after a short period of time.
 - Salt and hash customer passwords at rest.
 - Enable encryption for customer account passwords in transit.
 - Limit access to private keys to only those individuals whose role requires it.
- *Data in transit:* Implement encryption technologies based on specific use cases and customer configurations for data in transit. Customers are responsible for configuring encryption of their data in transit to end-users. Product-specific default encryption settings are described in the Documentation.
- *Data governance program:* Maintain a "privacy and protection-by-design" approach manifested in a data governance program and documented separately in the [data management documentation](#).

Fastly may access or modify equipment, systems, or services that manage customer data as necessary to provide and improve the services, prevent or address service or technical issues, as required by law, or as customers expressly permit.

Application security

Fastly manages application security using controls designed to do the following:

- *Secure development practices:* Train Fastly engineers annually on secure coding concepts, including the OWASP Top 10 and CWE Top 25, and implement testing and deployment controls. Code is peer-reviewed and run through automated testing

before deployment to production systems. After review and testing, code is initially deployed to a limited number of locations in the Fastly network for further monitoring. If no problems are encountered, code is gradually deployed across the Fastly network.

- *Application security analysis*: Conduct periodic analysis and regular penetration testing of Fastly-written code with Fastly security engineers and third-party validators.
- *Automated code analysis*: Deploy technology to automatically identify and report on identified vulnerable dependencies.

System and network security

Fastly manages system and network security in its production systems using controls designed to do the following:

- *Asset management*: Maintain an inventory of hardware and services deployed within the Fastly network.
- *Configuration standards*: Maintain secure configuration standards, including restricted ports, protocols, and services, and remove insecure default settings.
- *Patch management*: Patch systems on a regular basis and apply out-of-band patches for newly-identified risks.
- *System management*: Manage systems by verifying the controls addressed in Fastly's security program are in place, including logging and monitoring, host-based firewalls, and session management.
- *Audit and monitoring*: Log relevant security-related events, including authentication successes or failures and the use of certain commands. Events triggered by anomalous activity or suspicious behavior are investigated and alerts are monitored for effectiveness.
- *Documentation*: Maintain accurate network diagrams and internal documentation of Fastly production systems and services.
- *Access Control List (ACL) review*: On at least a semi-annual basis, review ACLs of host-based firewall and router rulesets.
- *Intrusion detection*: Maintain mechanisms designed to detect potential intrusions at the network and host level. Fastly inspects and responds to detected events, as applicable, to address threats.

applicable, to address threats.

Physical security

Fastly production systems reside in a combination of Fastly-managed data centers and cloud and internet service provider environments. Regardless of the physical location of the infrastructure or its operator, minimum, mandatory physical security controls are evaluated and applied to all environments in which Fastly production systems reside. Fastly manages physical security using controls designed to do the following:

- **Physical access to provider environments:** Only use providers that maintain physical and environmental protections, including perimeter protection, access logging and review, and other surveillance systems.
- **Physical access to production systems:** Only grant physical access to approved personnel. Requests for access are evaluated by authorized personnel based on proof of proper credentials and access is limited to specified areas.
- **Environmental security safeguards:** Require Fastly-managed data centers to protect their systems with controls including power redundancy, fire suppression, and other environmental controls.
- **Secure hardware destruction:** Use secure destruction of production hardware prior to disposal.
- **Disk encryption and secure boot:** Maintain additional technical security controls, such as disk encryption and secure boot, where deemed appropriate.

Data centers used for Fastly's PCI and HIPAA-Compliant Caching and Delivery are Payment Card Industry Data Security Standard (PCI DSS) certified locations.

Human security

Fastly manages human security using controls designed to do the following:

- **Employee background screening:** Conduct background screenings on each employee upon hire and maintain a policy requiring employees to report any criminal convictions during the course of employment, each as permitted by applicable local regulations.
- **Confidentiality agreements:** Require all employees to enter into confidentiality agreements with Fastly to help safeguard sensitive information that employees may view, process, or transmit as part of their job functions.

- **Awareness training:** Provide security training upon hire and annually thereafter for all employees. Mandatory annual training includes security awareness that covers application of best security and privacy practices in day-to-day work to help ensure each employee understands how to identify sensitive information and comply with applicable regulations.

Continuous monitoring and improvement

Fastly uses strict processes and testing procedures designed to continuously monitor, improve, and apply its security program. These are categorized as follows:

Change management process

Fastly follows a defined set of procedures to develop and deploy technology changes. These changes include updates to software, configurations, and devices that support Fastly's services. Fastly manages its change management process using controls designed to do the following:

- **Testing:** Require changes to be tested at various stages of development to confirm the changes operate as expected in a non-production environment before completing a deployment into its production services.
- **Change approval and notification:** Prepare, approve, and communicate change notices to maintain awareness among employees who manage the Fastly network and systems. Fastly maintains rollback procedures to address deployment issues if they arise.
- **Post-implementation review:** Confirm the success of changes after deployment.
- **Change monitoring:** Use multiple monitoring and alert mechanisms to enhance the visibility of technical changes and help ensure adherence to change management processes.

Vulnerability management

Fastly monitors for vulnerabilities in its production systems using controls designed to do the following:

- **Internal and external vulnerability scanning:** On a regular basis, automatically analyze production systems for vulnerabilities.
- **Vulnerability mitigation:** Assess the risk of identified or reported vulnerabilities, and

mitigate vulnerabilities in a timely manner based on the risk assessment. Mitigations for vulnerabilities deemed highest severity are implemented within twenty-four (24) hours of validation.

- **Distribution lists and vendor notification:** Monitor publicly disclosed vulnerabilities and confidential vulnerability distribution lists and notifications from software vendors.

Penetration testing

On an annual basis, Fastly engages a third party to conduct a penetration test of Fastly production systems. Identified issues are prioritized and handled based upon the severity of the evaluated risk they pose.

Compliance and audits

Fastly maintains recurring [audits and assessments](#) to confirm its security program meets various industry standards and regulatory requirements.

Fastly vendor management

Fastly uses third-party vendors and service providers to support its services. Fastly evaluates its vendors for security controls and risk to Fastly, its customers, and its services prior to using vendor services and regularly thereafter based on vendor risk.

When something goes wrong

Fastly aims to provide a reliable and secure platform and monitors for threats and system disruptions to detect, respond to, and recover from incidents in a timely manner. Fastly's implements an incident response process and business continuity controls as follows:

Incident management plan

Fastly maintains a formal incident response plan to address security-related incidents. The plan contains established roles and responsibilities, communication protocols, and response procedures. Fastly reviews and updates the plan periodically to adapt it to evolving threats and risks to its services. Representatives from key departments are assigned to address security-related incidents. These personnel coordinate the full lifecycle of incidents, from detection, through response, and recovery. These processes include communication with external parties, as needed.

Incident notification

Incident notification

Fastly notifies affected customers within forty-eight (48) hours of validating any unauthorized disclosure of customer data. Following any security-related incident, Fastly investigates and takes corrective action in a timely manner according to the incident management plan and provides affected customers with periodic updates.

Business continuity

Fastly manages business continuity using controls designed to do the following:

- **Service failover:** Prepare production systems for service failover. Production systems are deployed on infrastructure in multiple regions or zones to provide redundancy in the event of degraded performance or operational issues with a third party provider. If failure of a service occurs within a single region or zone, Fastly will automatically attempt to use infrastructure in another region or, if possible, another infrastructure provider.
- **Internet redundancy:** Only use data centers and cloud infrastructure providers that have connections with multiple internet service providers.
- **Service monitoring:** Monitor reporting channels to detect service-related issues. Personnel are available 24×7×365 to confirm and respond to disruptions of Fastly services.
- **Communication and reporting:** Provide service interruption updates to customers using various communication methods (including fastlystatus.com), depending on an incident's scope and severity.
- **Business continuity plan and testing:** Review, approve, update, and test a business continuity plan for Fastly production systems annually.
- **Data backups:** Conduct regular backups of data, excluding cached customer data, to support the recovery and availability of Fastly services. Data backups are tested on a periodic basis to validate backup recovery procedures.

* * *



Security testing your service behind Fastly



Last updated: 2025-01-06



</products/security-testing-your-service-behind-fastly>

We understand the need for our customers to validate the security of their service behind Fastly.

IMPORTANT

Penetration tests that interfere with or disrupt the integrity or performance of Fastly services violate our [acceptable use policy](#). You must respond immediately to any communication from Fastly regarding your test to help ensure your testing does not adversely affect other customers or the Fastly network.

To perform security testing of your Fastly service configurations, including penetration testing and DDoS testing, [contact support](#) at least two (2) business days before you begin any security testing. In your ticket, include these details:

- the [IDs of the services](#) that will be tested
- the source IP addresses of the test
- the date of the test
- the start and end time of the test, including the time zone
- the contact information for the individual or third party performing the test, including a phone number and e-mail address
- whether or not the security test is likely to lead to significantly increased traffic volume

The following requirements apply to any security testing you perform:

- Only test Fastly services you own or are authorized by the owner to test. You may not perform tests against other customers without explicit permission or against Fastly-owned resources.
- Do not begin testing until after Fastly has responded affirmatively to your ticket and authorized your request.

- Update the ticket if either the scope or timeframe of your testing changes.
- If you discover vulnerabilities in the Fastly platform during your test, update the ticket with your findings as soon as possible so we can address them.

Fastly maintains programs for [security](#) and [technology compliance](#). To perform an independent audit of these programs, contact sales@fastly.com to discuss purchase of [Assurance Services](#).

TIP

We welcome security professionals researching potential vulnerabilities in our network under our guidelines for [reporting a security issue](#).

* * *

Products

Edge Cloud Platform

DDoS Protection

Next-Gen WAF

Pricing

Try Fastly Free

Network Map

Services

Professional Services

Managed CDN

Support Plans

Talk to an Expert

Learn

Documentation

Developers

Resource Library

Blog

Events

Support

Support Center

Network Status

Contact Us

Company

.. . .

ABOUT US

Careers

Customer Stories

Partners

News

Investor Relations

Trust

© Fastly 2026

Archives

Changelog

Translations

Terms of Service

Policy FAQ