

Our service guides reflect the latest version of the Fastly web interface. Service guides for the classic web interface can still be accessed at [docs-legacy.fastly.com](https://docs-legacy.fastly.com) (<https://docs-legacy.fastly.com>). For more information see [our announcement \(/changes/2016/11/17/special\)](#).

Generated: Thu, 01 Dec 2016 19:03:49 +0000

# Fastly Help Guides

• [Guides \(/guides/\)](/guides/) > [Account management \(/guides/accounts\)](/guides/accounts) > [User access and control \(/guides/user-access-and-control/\)](/guides/user-access-and-control/)

## § Adding and deleting user accounts (/guides/user-access-and-control/adding-and-deleting-user-accounts)

Fastly allows you to add users to your account and assign them different roles (</guides/user-access-and-control/user-roles-and-how-to-change-them>), delete user accounts when you no longer want them to have access, or add existing users to existing accounts.

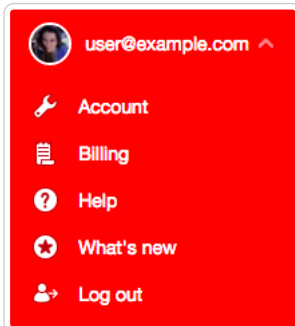
**❗ IMPORTANT:** You must be assigned the [role of superuser \(/guides/user-access-and-control/user-roles-and-how-to-change-them\)](/guides/user-access-and-control/user-roles-and-how-to-change-them) in order to add or delete user accounts.

### Adding new account users

**❗ TIP:** Adding a new user to make them the billing contact for your account? Follow our [billing contact instructions \(/guides/account-types-and-billing/who-receives-your-bill#changing-who-receives-your-bill\)](/guides/account-types-and-billing/who-receives-your-bill#changing-who-receives-your-bill) instead.

To add a user to your account, send them an invitation to join following the steps below:

1. Log in to the Fastly web interface.
2. From the user menu, select **Account**.



Your account information appears.

3. In the **User Invitations** area, click the **Invite** button. The Invite User window appears.

### Invite user

Email

\*

A confirmation code will be sent to this email.


Role

--Select--

INVITE

CANCEL

4. In the **Email** field, type the email address of the user to invite.
5. From the **Role** menu select the role to assign (</guides/user-access-and-control/user-roles-and-how-to-change-them>) once the user accepts the invitation.
6. Click the **Invite** button to send an invitation to the email you specified. The email address of the user you invited appears in the User Invitations area and remains there until the invitation is accepted.



## Deleting existing account users

**TIP:** Deleting the owner of the account? Be sure to [transfer ownership](/guides/user-access-and-control/user-roles-and-how-to-change-them#account-ownership-and-how-to-transfer-it) (</guides/user-access-and-control/user-roles-and-how-to-change-them#account-ownership-and-how-to-transfer-it>) first.

To delete a user from your account, follow the steps below:

1. Log in to the Fastly web interface.
2. From the user menu, select **Account**. Your account information appears.
3. Scroll down to the **Users** area and search for the name of the user to delete.
4. Click the gear icon to the right of the user to be deleted, and then select **Delete** from the menu that appears. A confirmation window appears.
5. Click **Confirm and delete**.

## Adding existing account users

There might be times when you want to invite a user who already has a Fastly account to your account. To add an existing user, follow the steps below:


1. Have the existing user you're inviting cancel their account (</guides/account-types-and-billing/accounts-and-pricing-plans#canceling-your-account>) so you can invite them to yours. We tie a user's email address to the creation of an account. Canceling that account allows the email address to be reused.
2. Log in to the Fastly web interface.
3. Re-invite the user to your account.

**NOTE:** Account cancellation might not be an option in some situations. Contact [support@fastly.com](mailto:support@fastly.com) (<mailto:support@fastly.com>) to discuss how accounts can be combined.

## § Email and password changes (</guides/user-access-and-control/email-and-password-changes>)

### Changing your name or email address

1. Log in to the Fastly web interface.
2. From the user menu, select **Account**.



Your account information appears.

3. In the **Your Profile** area, make changes as follows:

- In the **Name** field, type your name.
- In the **Email (login)** field, type your email address.

4. Click **Update Profile** to save the changes.

5. If you've changed your email address, confirm your password in the window that appears.

## Changing your password

To change the password currently associated with your account:

1. Log in to the Fastly web interface.
2. From the user menu, select **Account**. Your account information appears.
3. In the **Change Password** area, fill out the fields as follows:
  - In the **Current Password** field, type your existing password.
  - In the **New Password** field, type the new password.
  - In the **Confirm Password** field, type the new password a second time.
4. Click **Change Password**.

## Password requirements

When choosing a password keep in mind that it must:

- be at least 7 characters long
- contain at least one letter and one number

In addition, passwords cannot solely contain:

- sequences of letters or numbers (e.g., `12345678`, `abcdefg`)
- repeated characters (e.g., `222222`, `aaaaaa`)
- adjacent key placements on a standard keyboard (e.g., `QWERTY`)

The system will prevent you from choosing a password that:

- matches commonly used passwords (e.g., `password123`, `changeme`)
- uses popular dictionary words in passwords less than 16 characters (e.g., `batterystaple`)
- matches your user name or your email address

## § Merging accounts (/guides/user-access-and-control/merging-accounts)

If several employees at your company independently create testing accounts when learning about Fastly services (/guides/about-fastly-services/), you can have those testing accounts merged into a single account by emailing the Customer Support team at support@fastly.com (mailto:support@fastly.com) with the following information:

- the Customer IDs (/guides/account-management-and-security/finding-and-managing-your-account-info) of the accounts to be merged
- which account should be considered the primary account (any other accounts will be merged into the primary)

After you contact us, we'll reach out to verify the ownership of each account. If we can confirm ownership, we'll initiate a merge.

## § User roles and how to change them (/guides/user-access-and-control/user-roles-and-how-to-change-them)

Your Fastly account can be managed by multiple users. Users can be added to your account by invitation (/guides/user-access-and-control/adding-and-deleting-user-accounts), with their roles assigned and changed by **superusers** at any time via the web interface or via the API (/api/).

### Assignable user roles

Fastly allows you to assign one of four different roles to each user who has access to your account:

	User	Billing	Engineer	Superuser
<b>Analytics &amp; Stats</b>				
Select a service	X	X	X	X
Select a datacenter	X	X	X	X
<b>Configure</b>				
Create and delete services			X	X
Configure services			X	X
Compare service versions			X	X
Deactivate services			X	X
Purge			X	X
View and download generated VCL			X	X
Customize VCL			X	X
<b>Account &amp; Organization</b>				
Update personal profile settings	X	X	X	X
Update company settings				X
Issue user invitations				X
Assign user roles				X
Issue password resets				X
Delete user accounts				X
Enable and disable personal 2FA	X	X	X	X
Enable and disable company-wide 2FA				X
View and generate API keys			X	X
<b>Billing</b>				
View invoices		X		X
View billing history		X		X
Pay bills		X		X
Update credit card info		X		X
Change account type		X		X

By default, the system assigns the role of **superuser** to the person who signs up for your organization's account.


**NOTE:** The role you assign to users does not affect their ability to submit requests to Fastly [Customer Support \(/guides/customer-support/submitting-support-requests\)](/guides/customer-support/submitting-support-requests).

### Changing roles for existing users

Follow these instructions to change the roles for existing users:



1. Log in to the Fastly web interface.
2. From the user menu, select **Account**.



Your account information appears.

3. In the **Users** area, click the gear icon next to a user name and then select **Edit** from the menu that appears. The Edit User window appears for the selected user.

4. From the **Role** menu, select a role for the user.
5. Click **Update**. The user's role will be changed.

## Account ownership and how to transfer it

By default, we assign account "ownership" to the first user who signs up for an account for your organization. Accounts can only be canceled (</guides/account-types-and-billing/accounts-and-pricing-plans#canceling-your-account>) by owners. As the first user of an account, we automatically assign the owner the superuser role, though that role can be changed by another superuser once additional users are added.

An account owner serves as the primary point of contact for billing purposes. Invoices are sent to them, but if a specific billing contact (</guides/account-types-and-billing/who-receives-your-bill>) has been defined, invoices go to that contact instead.

To transfer account ownership to another user, contact [support@fastly.com](mailto:support@fastly.com) (<mailto:support@fastly.com>) for assistance.

• [Guides \(/guides/\)](/guides/) > [Account management \(/guides/accounts\)](/guides/accounts) > [Account management and security \(/guides/account-management-and-security/\)](/guides/account-management-and-security/)

## § Changing your account's company name (</guides/account-management-and-security/changing-your-accounts-company-name>)

Fastly allows you to change your account's company name at any time after it's been created.

**❗ IMPORTANT:** You must be the [account owner \(/guides/user-access-and-control/user-roles-and-how-to-change-them#account-ownership-and-how-to-transfer-it\)](/guides/user-access-and-control/user-roles-and-how-to-change-them#account-ownership-and-how-to-transfer-it) or be assigned the [role of superuser \(/guides/user-access-and-control/user-roles-and-how-to-change-them\)](/guides/user-access-and-control/user-roles-and-how-to-change-them) to change your account name.

To change the company name, follow the steps below.

1. Log in to the Fastly web interface.
2. From the user menu, select **Account**. Your account information appears.
3. In the **Company Name** field of the **Company Settings** area towards the middle of the page, replace the current company name.

## Company Settings

Company Name

Example Company

★ Required

Company ID

ABCDEFGH1234567

Owner

Artur Bergman

Billing Contact

Example Billing Contact

⌵

Phone Number

555 555 5555

★ Required

Address (optional)

IP Whitelist

0.0.0.0/0

★ Required

e.g. 192.168.0.1, 192.168.0.0/32, 192.168.\*.\*

[What is this?](#)

UPDATE COMPANY

CANCEL ACCOUNT

4. Click the **Update Company** button.

## § Enabling an IP whitelist for account logins (/guides/account-management-and-security/enabling-an-ip-whitelist-for-account-logins)

Fastly allows you to define the range of IP addresses authorized to access your Fastly account. This optional IP whitelisting functionality is not enabled by default.

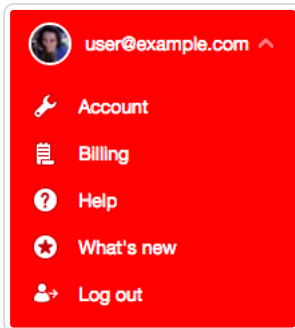
**⚠ WARNING:** *If you decide to use optional IP whitelisting, your account owner must have a valid telephone number on file.* During setup, Fastly checks your current IP address against the list you provide to ensure you don't lock yourself out of your account. If your IP addresses change at a later date (for example, because you move offices) and you forget to update your whitelist configuration, you will be locked out of your account. You will need to contact [support@fastly.com \(mailto:support@fastly.com\)](mailto:support@fastly.com) to request that a Customer Support

representative contact your account's owner via telephone during Fastly's regular business hours. To protect your account's security, **we will not unlock your account (</guides/customer-support/account-lockouts>) based on an email request alone.**

## Enabling an IP whitelist

To restrict access to your Fastly account based on a specific list or range of IP addresses, follow these steps.

1. Log in to the Fastly web interface.
2. From the user menu, select **Account**.



Your account information appears.

3. In the **IP Whitelist** field of the **Company Settings** area, replace `0.0.0.0/0` (the default IP range indicating no whitelisting) with the IP

addresses to which your account access should be restricted.

In the IP Whitelist field you can include single or multiple IP addresses or IP ranges (separated by commas) as follows:

- a single IPv4 address (e.g., replace the default with `192.168.0.1`)
- an IPv4 CIDR range (e.g., replace the default with `192.168.0.0/32`)
- an IPv4 Wildcard range (e.g., replace the default with `192.168.0.*`, `192.168.*.1`, `192.168.*.*`)

4. Click the **Update Company** button.

## Disabling an IP whitelist

To disable IP whitelisting on your Fastly account, follow these steps.

1. Log in to the Fastly web interface.
2. From the user menu, select **Account**. Your account information appears.
3. In the **IP Whitelist** field of the **Company Settings** area, type `0.0.0.0/0` (the default IP range indicating no whitelisting).
4. Click the **Update Company** button.

## § Enabling and disabling two-factor authentication (</guides/account-management-and-security/enabling-and-disabling-two-factor-authentication>)

Fastly supports two-factor authentication, a two-step verification system, for logging in to the web interface. In a two-factor authentication security process, users provide two means of identifying themselves to the system, typically by providing the system with something they know (for example, their login ID and password combination) and something they have (such as an authentication code). Organizations can enable company-wide two-factor authentication to require all users within the organization to use two-factor authentication.

## Before you begin

**You'll need to enter an authentication code regularly.** Once two-factor authentication has been enabled, an authentication code will be requested upon login at least every 14 days for each computer and browser you use to access the Fastly web interface.

**A mobile device is required.** Using this security feature with a Fastly account requires a mobile device capable of scanning a barcode or QR code using a downloadable authenticator application. We recommend the following:

- For Android, iOS, and Blackberry: Google Authenticator (<https://support.google.com/accounts/answer/1066447?hl=en>)
- For Android and iOS: Duo Mobile (<https://guide.duo.com/third-party-accounts>)
- For Android: AWS Virtual MFA (<http://www.amazon.com/Amazon-com-AWS-Virtual-MFA/dp/B0061MU68M>)
- For Windows Phone: Authenticator (<https://www.microsoft.com/en-us/store/apps/authenticator/9nblggh09llf>)

**There are special requirements for using this feature with the API.** If you enable two-factor authentication via the web interface, you will no longer be able to use a username and password when using the API. You must use the organization's API key for authentication.

## Managing two-factor authentication as a user

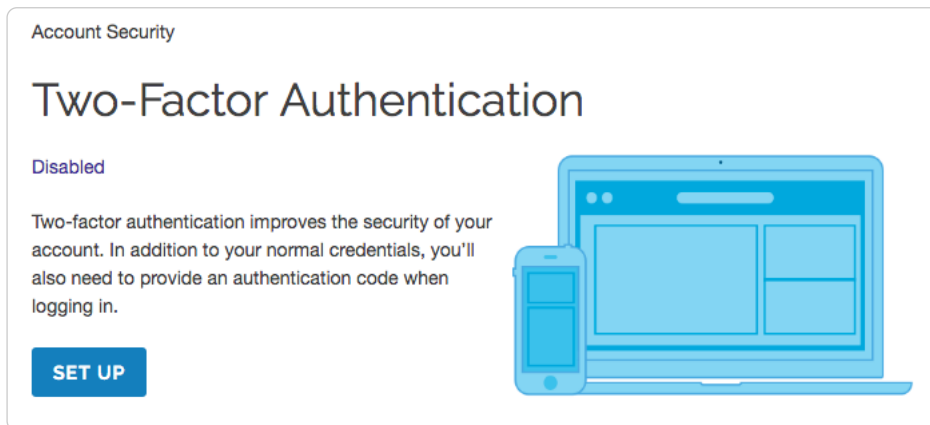
Depending on whether or not your organization has enabled company-wide two-factor authentication, you may be able to enable and disable two-factor authentication for your personal account. We also have instructions for recovering access to your account if you lose your mobile device.

### Enabling two-factor authentication

To enable two-factor authentication for your user account, follow the steps below.

**❗ IMPORTANT:** If your organization has enabled company-wide two-factor authentication, you will be required to set up two-factor authentication when you log in to the Fastly web interface. Skip to step six for instructions.

1. Log in to the Fastly web interface.
2. From the user menu, select **Account**. Your account information appears.
3. Click **Manage your security settings** from the **Account Security** area near the bottom of the page. The Account Security page appears.



4. Click the **Set Up** button. The password verification screen appears.
5. Verify your Fastly password and then click **Continue**. The authentication QR code appears.

## Account Security

## Two-Factor Authentication

We recommend one of the following authenticator apps:



Google Authenticator   

Duo Mobile  

Amazon Virtual MFA 

Authenticator 



**Step 1** Open the authenticator mobile app

**Step 2** Scan the QR code or enter **cvlk bc7x**

**Step 3** Enter the generated code below

**Authentication Code**

 Required

**CONTINUE**

**❗ IMPORTANT:** The QR code above is an example. Scan the one that appears in the Fastly application, not in this guide.

- Launch the authenticator application installed on your mobile device and scan the displayed QR code or manually enter the key displayed in the setup window. A time-based authentication code appears on your mobile device. Depending on your device, however, a browser link may first appear. You need to click this link to save it. When you do, the words `Secret saved` appear briefly.
- In the **Authentication Code** field, type the time-based authentication code displayed on your mobile device.

**❗ "ANDROID:** A common time syncing issue may cause your authenticator codes to fail. You can correct this using [Google's instructions](https://support.google.com/accounts/answer/185834?hl=en#sync) (<https://support.google.com/accounts/answer/185834?hl=en#sync>) for your authenticator application.

- Click **Continue**. The confirmation screen appears along with your recovery codes.

## Account Security

## Two-Factor Authentication

You've enabled two-factor authentication.

To access your account, use your normal login credentials and the generated code from the authenticator application on your mobile device.

**Keep your recovery codes in a safe place!** They're the only alternative way to access your account.

Can't access your mobile device? No worries. You can use one of your [recovery codes to log in](#).

## Recovery codes

1928  
26d9  
2146c  
424d  
4dd7  
5364  
6c87  
8911  
a587  
ed4a

**❗ IMPORTANT:** If you're ever unable to access your mobile device, the displayed recovery codes can be used to log in when your account has two-factor authentication enabled. Each of these recovery codes can only be used once, but you can regenerate a new set of 10 at any time (any unused codes at that time will be invalidated). Store your recovery codes in a safe place.

After you enable two-factor authentication, logging in to your Fastly account will require your email address and password, and then an authentication code generated by the authenticator application you've installed on your mobile device. By default, the system requires you to authenticate your login using an authentication code at least every two weeks for each computer and browser you use to access the Fastly web

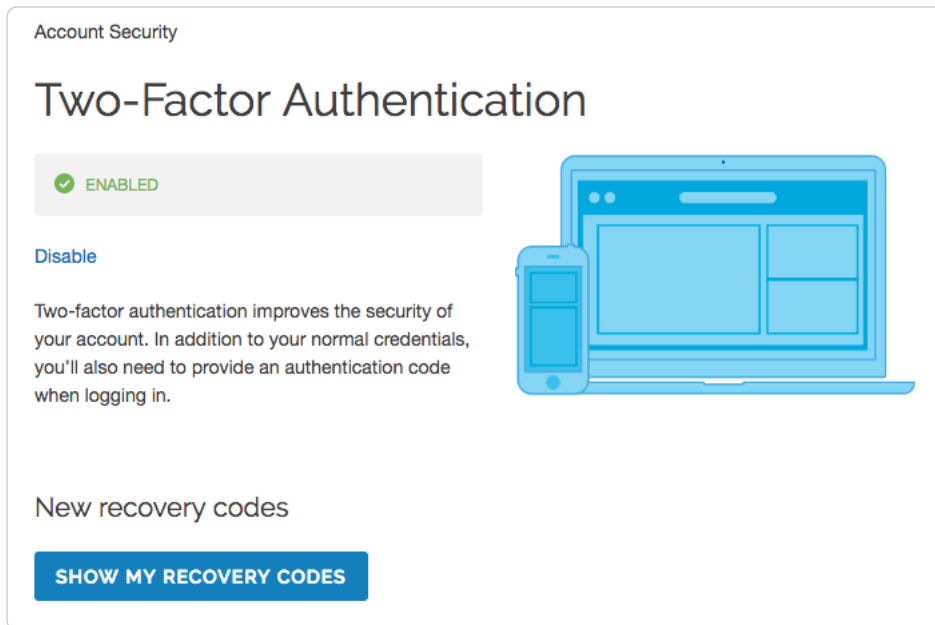
interface.

## Disabling two-factor authentication

Once two-factor authentication is enabled for your account, you can disable it at any time by following the steps below.

**❗ IMPORTANT:** If your organization has enabled [company-wide two-factor authentication](#), you cannot disable two-factor authentication for your account.

1. Log in to the Fastly web interface.
2. From the user menu, select **Account**. Your account information appears.
3. In the **Account Security** area of the settings sidebar, click **Manage your security settings**. The two-factor authentication controls appear.



4. Click **Disable**. The verification screen appears.
5. In the **Authentication Code** field, type the time-based authentication code displayed in the authenticator application on your mobile device, then click **Confirm and Disable**.

**❗ NOTE:** If you have lost your mobile device, you can enter a recovery code in the **Authentication Code** field. For more information, see the section on [what to do if you lose your mobile device](#).

## What to do if you lose your mobile device

If you lose your mobile device after enabling two-factor authentication, use a recovery code to log in to your Fastly account. You can continue to use recovery codes to log in until you get your mobile device back. Recovery codes can only be used once, however, so remember to regenerate a new list of codes to avoid running out before you recover your mobile device.


If you do not believe you will be able to recover your lost mobile device and you still have at least two recovery codes left, you can log in with one recovery code and disable two-factor authentication with a second code. Once two-factor authentication is disabled, you can re-enable it with a new mobile device at a later time and regenerate a new set of codes.

If your organization has enabled company-wide two-factor authentication, you can contact a superuser (</guides/user-access-and-control/user-roles-and-how-to-change-them>) for your organization and ask them to reset your two-factor authentication.

Locked out of your account? See our article on what you can do about it (</guides/customer-support/account-lockouts>).

## Managing two-factor authentication as a superuser

If you are assigned the superuser role (</guides/user-access-and-control/user-roles-and-how-to-change-them>) for your organization, you can view who has two-factor authentication enabled on the account tab in the Users area of your Account settings. Users with this feature enabled have padlocks displayed next to their names.



To disable two-factor authentication for any user within your organization, select **Disable 2FA** from the menu that appears when you click the gear icon next to that user's name.

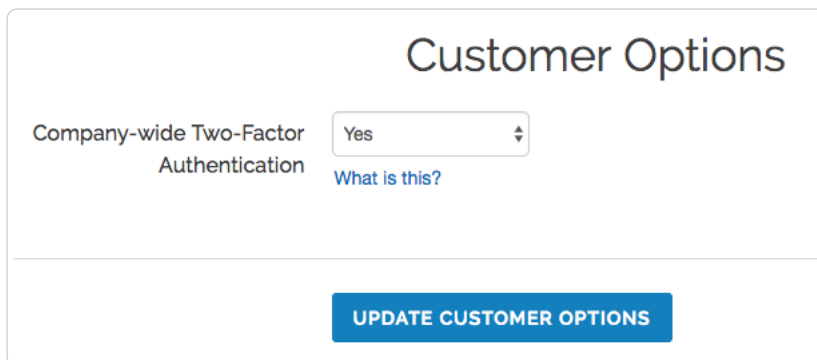
## Managing two-factor authentication as a company

Organizations can enable two-factor authentication for all of their users. When the company-wide two-factor authentication feature is enabled, all users within the organization are required to use two-factor authentication to log in to the Fastly web interface, and they cannot disable two-factor authentication for their accounts.

### Enabling company-wide two-factor authentication

Users assigned the superuser role (/guides/user-access-and-control/user-roles-and-how-to-change-them) can enable this feature on the Account page. To enable company-wide two-factor authentication for all users within your organization, follow the steps below.

1. Log in to the Fastly web interface.
2. From the user menu, select **Account**. Your account information appears.
3. In the **Customer Options** area, select **Yes** from the **Company-wide Two-Factor Authentication** menu.



4. Click **Update Customer Options**. A warning message appears.
5. Click **Continue**. You will be logged out of the Fastly web interface. This completes the setup process for company-wide two-factor authentication.

Users who have not already enabled two-factor authentication for their accounts will be prompted to do so the next time they log in to the Fastly web interface.

### Resetting a user's two-factor authentication

If company-wide two-factor authentication is enabled, and a user within the organization gets locked out of their account or needs to enable a new device, a superuser (/guides/user-access-and-control/user-roles-and-how-to-change-them) can reset the user's two-factor authentication. To reset a user's two-factor authentication, follow the steps below.

1. Log in to the Fastly web interface.
2. From the user menu, select **Account**. Your account information appears.
3. In the **Users** area, click the gear icon next to a user and then select **Reset 2FA**. A warning message appears.
4. Click **Reset**. The user will need to set up two-factor authentication for their account the next time they log in.

### Disabling two-factor authentication for a single user's account

If company-wide two-factor authentication is enabled, a superuser (/guides/user-access-and-control/user-roles-and-how-to-change-them) can disable two-factor authentication for a single user's account. This is typically done for user accounts being used for scripts and session authentication. To disable two-factor authentication for a single user's account, follow the steps below.


1. Log in to the Fastly web interface.
2. From the user menu, select **Account**. Your account information appears.

3. In the **Users** area, click the gear icon next to a user and then select **Ignore 2FA**. A warning message appears.
4. Click **Ignore**. Two-factor authentication will no longer be required for the selected user.

## Disabling company-wide two-factor authentication

A superuser (/guides/user-access-and-control/user-roles-and-how-to-change-them) can disable company-wide two-factor authentication. Once this feature is disabled, existing users within the organization will be able to manage their own two-factor authentication settings, and new users will not be required to set up two-factor authentication to log in to the Fastly web interface. To disable company-wide two-factor authentication, follow the steps below:

1. Log in to the Fastly web interface.
2. From the user menu, select **Account**. Your account information appears.
3. In the **Customer Options** area, select **No** from the **Company-wide Two-Factor Authentication** menu.




4. Click **Update Customer Options**. A warning message appears.
5. Click **Continue**. Company-wide two-factor authentication becomes disabled.

## § Finding and managing your account info (/guides/account-management-and-security/finding-and-managing-your-account-info)

Account information, including your service ID, your customer ID (also called your company ID), and your API key, can be accessed directly from the Fastly web interface (<https://manage.fastly.com/>).

### Finding your service ID

Your **Service ID** appears hidden, immediately below the name of your service on any page. To display the ID, click the **Show service ID** link:



To hide it, click the **Hide service ID** link.

### Finding your customer ID

Your **Customer ID** (also called your **Company ID**) is located in the **Company Settings** area near the middle of your **Account** page:



## Company Settings

Company Name

Example Company

★ Required

Company ID

ABCDEFG1234567

Owner

Artur Bergman

Billing Contact

Example Billing Contact

⬆ ⬇ ⬆

Phone Number

555 555 5555

★ Required

Address (optional)

IP Whitelist

0.0.0.0/0

★ Required

e.g. 192.168.0.1, 192.168.0.0/32, 192.168.\*.\*

[What is this?](#)

UPDATE COMPANY

CANCEL ACCOUNT

## Finding and displaying your API key

To display your API key, follow the steps below:

1. Log in to the Fastly web interface.
2. From the user menu, select **Account**. Your account information appears.
3. In the **API Clients** area near the bottom of the page, click the **Show** button. The password confirmation window appears.

### API Clients

We maintain API client libraries for [Ruby](#), [Perl](#), and [Python](#).

Visit the [Help Guides](#) and [API Documentation](#) for information on how to best use Fastly with your application.

API Key:

SHOW

GENERATE NEW KEY


4. In the **Your password** field, type your password and click the **Show** button. Your API key appears.

★ **TIP:** Your account has a master API key that is scoped to the entire account. To learn more, see our documentation on [API authentication \(/api/auth\)](#).

## Regenerating your API key

To regenerate your API key, follow the steps below:

1. Log in to the Fastly web interface.
2. From the user menu, select **Account**. Your account information appears.
3. In the **API Clients** area near the bottom of the page, click the **Generate New Key** button. A confirmation window appears.



4. In the **Your password** field, type your password and click the **Confirm and regenerate** button. Your key is regenerated.
5. To review your key, click the **Show** button. Your new API key appears.

**❗ IMPORTANT:** Only account users assigned the [engineer or superuser roles \(/guides/user-access-and-control/user-roles-and-how-to-change-them\)](/guides/user-access-and-control/user-roles-and-how-to-change-them) can view and change the API key.

• [Guides \(/guides/\)](/guides/) > [Account management \(/guides/accounts\)](/guides/accounts) > [Account types and billing \(/guides/account-types-and-billing/\)](/guides/account-types-and-billing/)

## § Accounts and pricing plans (/guides/account-types-and-billing/accounts-and-pricing-plans)

### Types of accounts and plans

Fastly offers a variety of accounts and pricing plans.

#### Free developer trials

We offer a development trial that allows you to test our services free of charge. We allow you to test up to \$50 of traffic for free to ensure everything fits your requirements. Simply sign up for a trial (<https://www.fastly.com/signup>) and begin testing.

We don't limit functionality during your trial. Once your testing is complete and you're ready to start pushing production traffic our way, you can switch your account to a Paid account.

#### Paid accounts without contracts

After your trial period ends, you can use Fastly's services on a month-to-month basis without having to sign a contract. Be sure you've provided us with current billing address (</guides/account-types-and-billing/paying-your-bill#changing-your-billing-address>) as well as your credit card information. Then switch to a paid account by accessing the Billing (</guides/about-fastly-services/about-the-web-interface-controls#about-the-billing-page>) page while logged into your account and then selecting "Paid" from the Change pricing plan menu in the Pricing Plan area on the Pricing Plan tab. You'll be prompted to enter your billing address (</guides/account-types-and-billing/paying-your-bill#changing-your-billing-address>) if you haven't already entered that information.

## Pricing Plan

Current Plan Developer

Change Pricing Plan Developer - Free

☐ I agree to [Terms of Service](#).

**UPDATE PLAN**

Once you switch to a paid plan, the development trial option will disappear and we'll bill you automatically (</guides/account-types-and-billing/how-we-calculate-your-bill/>) at the end of every month using the credit card information (</guides/account-types-and-billing/paying-your-bill#changing-your-credit-card-information>) you have on file with us. You can cancel this account at any time.

## Contractual commitments

If you plan to push at least 2TB of data per month and require one of our TLS service options (</guides/securing-communications/ordering-a-paid-tls-option/>), or if you plan to push a minimum of 4TB of data per month, it might be worthwhile to consider a contract with Fastly. Contact us at [sales@fastly.com](mailto:sales@fastly.com) (<mailto:sales@fastly.com>) for more information. We also offer solutions (<https://www.fastly.com/services/web-mobile-performance>) targeted to the needs of specific industries (for example, our Ecommerce package (<https://www.fastly.com/industries/ecommerce>)).

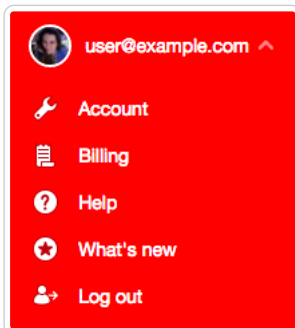
## Free open source developer accounts

We're active open source contributors and support the community (<https://www.fastly.com/open-source>) whenever possible. If you're an open source developer, your bill is on us. Contact us at [community@fastly.com](mailto:community@fastly.com) (<mailto:community@fastly.com>) to get started.

## Canceling your account

To cancel your account, have your account owner (</guides/user-access-and-control/user-roles-and-how-to-change-them#account-ownership-and-how-to-transfer-it>) follow the steps below:

1. Log in to the Fastly web interface.
2. From the user menu, select **Account**.



Your account information appears.

3. In the **Company Settings** area, click the **Cancel Account** button.

## Company Settings

Company Name

Example Company

★ Required

Company ID

ABCDEFGH1234567

Owner

Artur Bergman

Billing Contact

Example Billing Contact

⌵

Phone Number

555 555 5555

★ Required

Address (optional)

IP Whitelist

0.0.0.0/0

★ Required

e.g. 192.168.0.1, 192.168.0.0/32, 192.168.\*.\*

What is this?

UPDATE COMPANY

CANCEL ACCOUNT

A confirmation window appears.

4. In the **Your password** field of the confirmation window, type the password associated with your account and click **Confirm and cancel**.

After your account is canceled, you'll be billed for any outstanding charges (/guides/account-types-and-billing/how-we-calculate-your-bill) accrued through the day you canceled. For questions about your final billing statement, contact our billing team (mailto:billing@fastly.com) for assistance. If you decide at a later date to reactivate your account, contact Customer Support (mailto:support@fastly.com) and request reactivation.

## § How we calculate your bill (/guides/account-types-and-billing/how-we-calculate-your-bill)

### What we do and don't charge you for

We bill you monthly according to that month's use of Fastly's services. Your monthly charges represent the combined total of your bandwidth and request usage, using the rates listed on our pricing page (https://www.fastly.com/pricing). The minimum charge you pay each month is \$50.

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."

Two specific settings related to responses may affect the total charges on your bill. Enabling gzipping (/guides/basic-configuration/enabling-automatic-gzipping) can reduce the size of responses which reduces the bandwidth you use and thus can reduce your total bill. Enabling shielding (/guides/performance-tuning/shielding) may initially result in greater bandwidth use because requests may need to travel between POPs. The reduced load on your origin servers, however, frequently offsets this increased cost and the potential increase in your bill's total.

Charges for TLS service options (/guides/securing-communications/ordering-a-paid-tls-option) you've chosen are applied in addition to the bandwidth and request usage we charge for normal content delivery and streaming. We don't bill you for the amount of space your content takes up on our cache servers.

We don't charge you for Standard Support (/guides/customer-support/support-description-and-sla) service.

## About the monthly minimum charges

We bill a minimum of \$50 per month so we can fully support all of our customers. This is the minimum price you'll pay in any month once you've completed your testing trials (/guides/account-types-and-billing/accounts-and-pricing-plans).

For example, say that you're done testing Fastly's services and you've begun to push production-level traffic through Fastly. If most of your site's traffic for the current month is in North America and Europe and your site uses 10GB of traffic over 10 million requests, the combined bandwidth and request charges would be \$8.70 for the month. Because this amount falls below the \$50 monthly minimum, we would charge you \$50 for that month, not \$8.70.

Asia-Pacific, Australia, New Zealand, and Latin America prices listed are slightly higher. If most of your site's traffic were in these areas instead of North America and Europe, then at the above traffic levels your bandwidth and request usage charges would still fall below the monthly minimum and we would charge you \$50 for that month.

**NOTE:** If you're using Fastly for content delivery via Heroku's cloud development services, see Fastly's Heroku add-ons [pricing plan](https://elements.heroku.com/addons/fastly) (<https://elements.heroku.com/addons/fastly>) for additional details.

## When we charge you for Fastly services

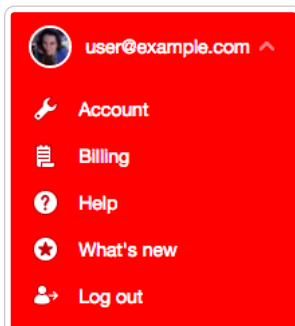
Fastly bills in arrears, not in advance, meaning that we bill you for services after you've used them, not before. For example, if you sign up for and start using Fastly services in January, the bill you receive in February reflects January's charges and services, your March bill reflects February's charges services, and so forth.

## How account cancelation affects your bill

If you ever cancel your account (/guides/account-types-and-billing/accounts-and-pricing-plans#canceling-your-account), you'll be billed for any outstanding charges accrued through the day you canceled, or at least the monthly minimum, whichever amount is greater.

## Reviewing the charges to your account

If you've been assigned a superuser or billing role (/guides/user-access-and-control/user-roles-and-how-to-change-them), you can review your account use and the associated charges via the Billing page in the Fastly web interface. Billing overview and invoice information can be accessed by clicking the Billing link in the user menu at the top right of any page.



By default, the month-to-date information about your account use appears, summarizing the bandwidth you've used and the associated charges, followed by the charges you've incurred to date for requests for the current month. The bottom of the summary displays the grand total dollar amount owed for the month-to-date.

### Billing

Paid in Full

MONTH-TO-DATE

INVOICE HISTORY

PRICING PLAN

TAX ADDRESS

#### September 2016 – Month-to-date

Bandwidth	0 GB	\$0
Incurring		\$0
Developer Plan Minimum		\$0
Grand Total		\$0

You can review similar details (including invoices) for previous months by clicking the **Invoice History** tab and then the link for any month listed in the Date column.

Billing			
2 Unpaid			
MONTH-TO-DATE	INVOICE HISTORY	PRICING PLAN	TAX ADDRESS
DATE	AMOUNT	STATUS	ACTIONS
<a href="#">August 2016</a>	\$50	Outstanding	<a href="#">PRINT</a>

## Regional account use details

To view specific regional use details for your account, click the View details link above the summary info for the current month-to-date. The details that appear include the current month's tiered bandwidth and request use for each geographic region of the world.

United States		
\$0		
Bandwidth		
TIER	PRICE	UNITS
first 10TB	\$0.12 / GB	0
Requests		
TIER	PRICE	UNITS
per 10,000 HTTP requests	\$0.01 / GB	0
per 10,000 HTTPS requests	\$0.01 / GB	0
Europe		
\$0		
Bandwidth		
TIER	PRICE	UNITS
first 10TB	\$0.12 / GB	0
Requests		
TIER	PRICE	UNITS
per 10,000 HTTP requests	\$0.01 / GB	0
per 10,000 HTTPS requests	\$0.01 / GB	0

You can review similar regional, tiered bandwidth and request use details for previous months by clicking on the link for any month under the Date column of the Billing History area below the current month's billing summary.

## Printing account use details

You can print account use details for any month directly from your browser window by viewing the specific summary or regional details and then selecting the Print option from your browser's File menu.

## § Paying your bill (/guides/account-types-and-billing/paying-your-bill)


At the end of each month, your billing contact (/guides/account-types-and-billing/who-receives-your-bill) will be sent an email summarizing your current usage levels (/guides/account-types-and-billing/how-we-calculate-your-bill) and the charges your account incurred for the month. The email contains a link to the related invoice for your account.

You'll need both a valid credit card and current billing address when you switch to a paid, month-to-month account (/guides/account-types-and-billing/accounts-and-pricing-plans#paid-accounts-without-contracts). Once your invoice gets generated, your credit card is automatically charged.

### Changing your credit card information

To change your credit card information, follow the steps below:

1. Log in to the Fastly web interface.
2. From the user menu, select **Billing**.



The Billing page appears.

3. Click the **Settings** tab. The Pricing Plan and Credit Card areas appear.
4. In the **Credit Card** area click **Edit**. Details appear for the credit card you have on file with Fastly.
5. Make any necessary changes to the credit card number, card verification code, and expiration date for the credit card to be used for automatic billing.

### Credit Card

Credit Card Number

\*\*\*\* \* 4242

★ Required

CVC

★ Required

Expiration Date

1 - January

2018

UPDATE

REMOVE

CANCEL EDIT

6. Click **Update** to save your credit card information.

★ **TIP:** Fastly never sees your credit card number. All transactions are handled by our [fully PCI compliant payment gateway](https://stripe.com/docs/security) (<https://stripe.com/docs/security>), Stripe.

### Changing your billing address

To change your billing address, follow the steps below:

1. Log in to the Fastly web interface.
2. From the user menu, select **Billing**. Your account information appears.

3. Click the **Tax Address** tab and type your tax address.

## Billing

Paid in Full

MONTH-TO-DATE

INVOICE HISTORY

PRICING PLAN

TAX ADDRESS

### Tax Address

For tax purposes. Tax address may be different than address on file with your credit card.

Country

United States of America

Street Address 1

★ Required

Street Address 2

City

★ Required

State or Province

--Select--

ZIP Code or Postal Code

★ Required

UPDATE TAX ADDRESS

4. Click the **Update Tax Address** button to save the address.

## § Who receives your bill (/guides/account-types-and-billing/who-receives-your-bill)

By default, your account owner is considered your billing contact and will receive your bill for Fastly services. You can change your billing contact at any time if you've been assigned the superuser role (/guides/user-access-and-control/user-roles-and-how-to-change-them) for an account. If you ever delete your billing contact, billing will automatically revert to the account owner.

**❗ IMPORTANT:** Invoices are only sent to the email addresses of the Account Owner or the Billing Contact. Invoices are not sent every user assigned a [billing role \(/guides/user-access-and-control/user-roles-and-how-to-change-them\)](/guides/user-access-and-control/user-roles-and-how-to-change-them).

## Changing who receives your bill


Follow the steps below to have your billing invoice sent to a person other than the owner of your account.

### For new users

To send the billing invoice to a user who has not yet created an account, follow these steps.


1. Log in to the Fastly web interface.
2. From the user menu, select **Account**.





Your account information appears.

3. In the **User Invitations** area, click the **Invite** button. The Invite User window appears.

A modal window titled 'Invite user' with a close button (X) in the top right corner. The window has a light gray background. It contains two main input fields: 'Email' and 'Role'. The 'Email' field is a text input with a small orange asterisk to its right. Below it, a note says 'A confirmation code will be sent to this email.' The 'Role' field is a dropdown menu with '--Select--' as the current selection. At the bottom of the window, there are two buttons: a blue 'INVITE' button on the left and a gray 'CANCEL' button on the right.

4. In the **Email** field, type the email address of the user you'd like to invite to become a billing contact.

★ **TIP:** To send invoices to multiple people, we recommend setting up a group email address and setting that email address as your Billing Contact user.

5. From the **Role** menu, select the role to assign (/guides/user-access-and-control/user-roles-and-how-to-change-them) once the user accepts the invitation.
6. Click the **Invite** button to send an invitation to the email you specified.
7. Once the user has accepted the invitation, return to the account information in the web interface.
8. In the **Company Settings** area, select the user's name from the **Billing Contact** menu.

## Company Settings

Company Name

Example Company

★ Required

Company ID

ABCDEFGH1234567

Owner

Artur Bergman

Billing Contact

Example Billing Contact

⌵

Phone Number

555 555 5555

★ Required

Address (optional)

IP Whitelist

0.0.0.0/0

★ Required

e.g. 192.168.0.1, 192.168.0.0/32, 192.168.\*.\*

[What is this?](#)

UPDATE COMPANY

CANCEL ACCOUNT

9. Click **Update Company** to set the billing contact.

## For existing users

To send the billing invoice to a user who already has an account, follow these steps.

1. Log in to the Fastly web interface.
2. From the user menu, select **Account**.
3. In the **Company Settings** area, select the user's name from the **Billing Contact** menu. Make sure the user name you select has the correct role (/guides/user-access-and-control/user-roles-and-how-to-change-them) assigned to view and manage billing information.
4. Click **Update Company** to save the billing contact.

• [Guides \(/guides/\)](/guides/) > [Account management \(/guides/accounts\)](/guides/accounts) > [Customer support \(/guides/customer-support/\)](/guides/customer-support/)

## \$ Account lockouts (/guides/customer-support/account-lockouts)


### Why is my account locked?

For security reasons, Fastly limits the number of times someone can try logging in to an account. We don't want to give people unlimited attempts at guessing your password, so we stop them from trying after a limited number of failed attempts to sign in. You can change your password (/guides/user-access-and-control/email-and-password-changes#changing-your-password) at any time when you're logged in to your account.


### I am not using two-factor authentication. How can I access my account?

Once locked, you will not be able to sign in to your account, even with the correct password. To unlock your account because you exceeded the number of guesses you were allowed:

1. Point any standard web browser (/guides/debugging/browser-recommendations-when-using-the-fastly-application) to the Fastly login page (<https://manage.fastly.com>). The login controls appear.

The image shows the Fastly login page. It has a red background with the Fastly logo (a white circle with a dot) at the top. Below the logo, the text "Log in to Fastly" is displayed in white. There are two input fields: one for email (placeholder: "email@example.com") and one for password (placeholder: "password"). To the right of the password field is a link that says "Forgot password". Below the input fields is a blue button with the text "LOG IN" in white. At the bottom, there is a link that says "DON'T HAVE AN ACCOUNT? SIGN UP" in white.

2. Click the **Forgot Password** button underneath the password field. The Reset your password page appears.

The image shows the Fastly "Reset your password" page. It has a red background with the Fastly logo (a white circle with a dot) at the top. Below the logo, the text "Reset your password" is displayed in white. There is one input field for email (placeholder: "email@example.com"). Below the input field is a blue button with the text "RESET" in white. At the bottom, there is a link that says "RETURN TO LOG IN" in white.

3. In the **Email (login)** field, type the email address you normally use to log in to your Fastly account, and then click the **Reset** button. Password reset instructions will be emailed to you.
4. Click on the password reset link in the emailed instructions that the system sends you. The Reset Your Password page appears.
5. Click the **Reset Password** button. The system sends you a temporary password to the email address you supplied.
6. Using the temporary password you receive, log in to your account. The controls to create a new password appear.
7. Fill out the **Reset Password** controls as follows:
  - In the **Current Password** field, type the temporary password that the system emailed to you when you requested a password reset.
  - In the **New Password** field, type a new password to replace the temporary password you were sent.
  - In the **Confirm Password** field, type the new password a second time to confirm it.
8. Click the **Change Password** button. The system changes your password and logs you into your account.

## I am using two-factor authentication. How can I access my account?

### I don't have my mobile device.

If you do not have access to your mobile device, you can complete the login process using one of your recovery codes. These were the recovery codes you saved in a secured location outside of your Fastly account when two-factor authentication was first enabled. You can continue to use your recovery codes until your device is once again accessible. Recovery codes can only be used once, however, so remember to regenerate a new set to avoid running out before you recover your mobile device.

If you don't believe you will be able to recover your lost mobile device and you still have at least two recovery codes left, you can log in with one recovery code and disable two-factor authentication with a second code. Once two-factor authentication is disabled, you can re-enable it with a new mobile device at a later time and regenerate a new set of codes.

## I don't have my mobile device and I don't have my recovery codes.

If you don't have your mobile device and didn't save any recovery codes, have another user at your company with the superuser role (</guides/user-access-and-control/user-roles-and-how-to-change-them>) contact Customer Support at [support@fastly.com](mailto:support@fastly.com) (<mailto:support@fastly.com>). Have them inform Customer Support which user needs assistance with their login. After Customer Support verifies that the request is from a superuser, we will provide them with your recovery code. The superuser will then send you this information and reset your password so that you can access your account.

## I don't have my phone, I didn't save my recovery codes, and I am the only superuser for the account.

Contact Customer Support at [support@fastly.com](mailto:support@fastly.com) (<mailto:support@fastly.com>). We will verify that you are associated with the company by phone. We will use the contact information located on the company website or under the Fastly account tab. Upon verification, we will send you a recovery code and reset your password.

## Was my account compromised?

If a user's account appears to be hacked or phished, we may proactively reset the passwords for the affected accounts to in order to revoke access to the hacker. In these cases, we send an email to the account's real owner (you) with additional information on how to reset the password. If you received one of these emails, follow the instructions in the email.

If you think your account has been hacked or phished, contact Customer Support at [support@fastly.com](mailto:support@fastly.com) (<mailto:support@fastly.com>) immediately.

## How is a locked account different from a blocked account?

Fastly allows you to restrict who can access your Fastly account based on the IP address of the person attempting to log in. This means that even with the correct login name and password, access to your Fastly account may be blocked if the IP doesn't match your company's list of allowed addresses.

If your company enables this optional IP whitelisting (</guides/account-management-and-security/enabling-an-ip-whitelist-for-account-logins>), they must keep the list of restricted IP addresses up to date. Only users with the role of superuser (</guides/user-access-and-control/user-roles-and-how-to-change-them>) can make changes to the IP whitelist settings (your account owner is always a superuser), and your account owner must have a valid telephone number on file to do so.

If your IP addresses change after whitelisting is enabled and you forget to update your whitelist configuration, you will be locked out of your account. You will need to contact [support@fastly.com](mailto:support@fastly.com) (<mailto:support@fastly.com>) to request that a Customer Support representative contact your account's owner via telephone during Fastly's regular business hours. To protect your account's security, we will not unlock your account based on an email request alone.

---

## § Common service and domain errors (</guides/customer-support/common-service-and-domain-errors>)

---

### Exceeding max number of domains

We currently limit the maximum number of services and domains you can configure. Once you reach that limit, error messages may appear that look something like this:

```
{
  "msg": "An error occurred while connecting to the fastly API, please try your request again.",
  "detail": "Exceeding max number of domains: 10"
}
```

If you're receiving a limit message and need to create more services or domains, contact [support@fastly.com](mailto:support@fastly.com) (<mailto:support@fastly.com>) for assistance. Fastly support engineers can not only increase the number of services that you can use, they can suggest other ways to design what you are trying to achieve.

---

## § Compliance Services (</guides/customer-support/compliance-services>)

---

Subscribers who purchase Compliance 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 (/guides/compliance/security-program) and technology compliance (/guides/compliance/technology-compliance) programs, subject to Subscriber's purchase of Professional Services (/guides/customer-support/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 Compliance Services must also purchase Gold or Platinum Support (/guides/customer-support/support-description-and-sla).

## § DDoS Protection and Mitigation Service and SLA (/guides/customer-support/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.
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 (<https://api.fastly.com/public-ip-list>) that allows Subscribers to download an updated list of all Fastly IPs globally and can be used to filter traffic and control communication between Fastly's caches and a Subscriber's origin. Fastly provides the Fastly IP Space to Subscribers in order to ensure known communication between the Fastly cache nodes and a Subscriber's origin datacenter.
- "Fastly Control" means elements entirely under Fastly's control and not a consequence of (a) a Subscriber's hardware or software failures, (b) a Subscriber's or end user's connectivity issues, (c) Subscriber operator errors, (d) Subscriber traffic amounts that exceed a Subscriber's Permitted Utilization as defined in the Service Availability SLA (</guides/customer-support/service-availability-sla>), (e) corrupted Subscriber content, (f) acts of god (any) or war, or earthquakes, or terrorist actions.

## Subscriber responsibilities

As a Subscriber, you:

- must identify and maintain two points of contact to be used during an attack to communicate status, issues, and coordinate with Fastly to successfully protect services.
- must use common best practices for DDoS Attack defense including:
  - using updated white and black lists in the Fastly IP Space at the origin datacenter 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\(/guides/\)](https://docs.fastly.com(/guides/)).

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

## Communications and channels of support

### Support tickets

Create support tickets by sending an email to [support@fastly.com](mailto:support@fastly.com) (<mailto: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](https://slack.com) (<https://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](mailto:billing@fastly.com) (<mailto:billing@fastly.com>) within 30 days of the DDoS Attack.

## DDoS Mitigation response SLA

If, during a DDoS Attack on a Subscriber with 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](mailto:billing@fastly.com) (<mailto: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.

## § Legacy Premium Support and SLA (/guides/customer-support/legacy-premium-support-and-sla)

**NOTE:** Fastly maintains support for its original Premium Support plan. For more information about our current [Gold and Platinum support plans \(/guides/customer-support/support-description-and-sla\)](/guides/customer-support/support-description-and-sla) or for information about our [Professional Services packages \(/guides/customer-support/professional-services\)](/guides/customer-support/professional-services), contact [sales@fastly.com](mailto:sales@fastly.com) (<mailto: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:

Offering	Unpaid Account	Month-to-Month Account	Termed Contract	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

Offering	Unpaid Account	Month-to-Month Account	Termed Contact	Premium Support
Severe Incident Response Email Address	No	No	No	Yes
Support SLA	None	None	Termination Option	Invoice Credits + Termination Option

## Technical support

The following section applies to all Subscribers.

### Definitions

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

### Subscriber responsibilities

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

### Support requests

Subscribers submit support requests by sending email to [support@fastly.com](mailto:support@fastly.com) (<mailto: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 cause(s) of the Incident and resolve them if caused by factors within Fastly's control, or provide information to those who can resolve the factors if the factors are within others' control, as follows:

For a Severity 1 Incident:

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

For a Severity 2 Incident:

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

### Charges for Incident Response

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

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

For all Incidents:

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

### Support Invoice Credits

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

#### Credit Terms:

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

## Legacy Service availability SLA

Support availability and response times vary depending on the type of account (/guides/account-types-and-billing/accounts-and-pricing-plans) you have and the level of support (/guides/customer-support/support-description-and-sla) 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 User Interface or API due to planned maintenance.

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

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

## Termination

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

## 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.

# § Product and feature life cycles (/guides/customer-support/product-and-feature-life-cycles)

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.
- Fastly strongly advises against using production traffic for Beta products due to their dynamic nature.

Beta services are covered by Section 7 of our Terms of Service (<https://www.fastly.com/terms>).

## Limited Availability

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

These guidelines apply to Fastly's Limited Availability program:

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

## General Availability

General Availability products released by Fastly are available for everyone's use. Fastly manages these products in accordance with Fastly's terms and conditions (<https://www.fastly.com/terms>).

## 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 (</guides/third-party-technology/>) 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 (<mailto:support@fastly.com>) or your account team.

---

## § Professional Services (</guides/customer-support/professional-services>)

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

## Implementation Services

### How it works

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

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

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

### Implementation options

Some common implementation options we offer include:

- Initial setup and configuration
- End-to-end encryption setup
- Fine-tuning cache times
- Custom header logic
- 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 and authorization
- ESI (edge side includes)
- Streaming and video packaging
- Site performance analysis
- Managed vendor migration

Fastly offers two Service Implementation packages:

- **Standard Implementation** - Basic implementation for Fastly customers who have simple content configurations.
- **Enterprise Implementation** - Advanced implementation for Fastly customers who have complex, custom configurations.

## Management Services

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

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

## Consulting Engagement Services

For customers who lack 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

## § Service availability SLA (/guides/customer-support/service-availability-sla)

Support availability and response times vary depending on the type of account (/guides/account-types-and-billing/accounts-and-pricing-plans) you have and the level of support (/guides/customer-support/support-description-and-sla) you have purchased.

Agreement Type	Unpaid Account	Month-to-Month Account	Termed Contract	Gold & 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 User Interface or API due to maintenance.

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

**"Fastly Control"** means elements entirely under Fastly's control and not a consequence of (a) Subscriber hardware or software failures, (b) Subscriber or end user connectivity issues, (c) Subscriber operator errors, (d) a Utilization spike (see below), (e) corrupted Subscriber content, (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 Gold or 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%

Availability Percent	Period of Degraded Performance	Monthly Credit Percent
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.

## § Submitting support requests (/guides/customer-support/submitting-support-requests)

Fastly offers standard technical support for all accounts through its support@fastly.com (mailto:support@fastly.com) email address. However, support availability and response times vary depending on the type of account you have and the level of support you have purchased. Our support description and SLA (/guides/customer-support/support-description-and-sla) information discusses these details further.

For all levels of technical support, Fastly's customer support ticketing system automatically generates and assigns each request a unique service ticket number, which is then sent within minutes to the customer at the email address used when submitting the ticket. This automated response also contains a direct link to the service ticket within Fastly's customer support system. Fastly then responds to these requests as appropriate for the level of support purchased by the customer.

## Accounts with Gold or Platinum Support


To submit emergency support tickets, customers must pay for a Gold or Platinum Support plan that includes support for severe incidents. Customers with Gold or Platinum Support plans can submit requests for support any time:

- a service that has been working stops working.
- a production deployment experiences a drastic change in performance that either is customer- or end-user-affecting.

Customers with Gold or Platinum Support accounts submit emergency tickets to an emergency support email address (instead of via the standard Customer Support email address (mailto:support@fastly.com)). Customers with Platinum Support accounts can also submit tickets via a toll-free telephone number.

## Accounts without Gold or Platinum Support

Customers who do not pay for Gold or Platinum Support (/guides/customer-support/support-description-and-sla) levels can submit support requests via email to support@fastly.com (mailto:support@fastly.com) or directly in the Fastly web interface, via the Help link in the user menu, any time they require assistance with self-provisioned tasks (/guides/about-fastly-services/self-provisioned-fastly-services). Fastly then responds to these requests as appropriate for the level of support purchased by the customer.



## Tips on what to include in a support request

To help shorten the time it takes to resolve a submitted support request, here are some tips on key information to include in your request when you reach out to the Fastly support team:

- **Give us your customer ID if you have one.** Your customer ID (</guides/account-management-and-security/finding-and-managing-your-account-info#finding-your-customer-id>), service ID, or the email address associated with your Fastly account help us identify who you are. If those don't apply, any information that would help us verify your identity will help us get started on your issue more quickly.
- **Tell us what's happening.** This helps us better understand your issue. For example, you expected to see something in your service and it's not there, or you're seeing an error message. Be sure to include the same information entered on the subject line within the body of the support request.
- **Describe your problem in detail.** We need information like specific service IDs and relevant URLs that aren't behaving as expected, the exact behavior you observed (perhaps include screenshots of what you saw), and whether or not you recently made changes that might have contributed to the issues (e.g., new DNS entries or new origins). These details help us focus support where it's needed most and makes our response to you faster.
- **Provide the steps you took.** This allows us to quickly reproduce your issue. For example, send the cURL command you ran if it created an issue or if you were trying a new configuration change, indicate whether you were using the API, the web interface, or custom VCL. Also, be sure to remove any passwords and API tokens from the content you provide in the support request.
- **Use Fastly Debug to identify network issues.** Supplying the output of Fastly Debug (<http://www.fastly-debug.com/>) provides us with key diagnostic information. For example, your requests are being routed to a POP you didn't expect, the debugging output helps us understand what's happening between the browser and Fastly POPs.

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

Support Offering	Standard Support	Gold Support	Platinum Support
Phone and chat support	Not available.	Not available.	Toll-free telephone available 24/7/365. Dedicated chat channel available during Fastly business hours.
Emergency Escalation	Not available.	Not available.	Available via email and phone.
Designated Customer Support Engineer	Not available.	Not available.	Available for large accounts on case-by-case basis.
Termination Option	Not available for unpaid and month-to-month customers. Only included for termed contracts.	Available with invoice credits.	Available with invoice credits.

## Technical support

The following section applies to all subscribers.

### Definitions

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

### Subscriber responsibilities

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

### Support requests

Subscribers submit support requests by sending email to [support@fastly.com](mailto:support@fastly.com) (<mailto: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 (</guides/customer-support/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](http://www.slack.com) (<https://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, Gold Support or 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 cause(s) of the Incident and resolve them if caused by factors within Fastly's control, or provide information to those who can resolve the factors if the factors are within others' control, as follows:

For a Severity 1 Incident:

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

For a Severity 2 Incident:

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

### 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.

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

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

**For all incidents:**

- If the Incident-causing factors are within Fastly's control, there will be no hourly charges for Fastly engineering staff time.
- If the factors are within Subscriber's control, Subscriber agrees to pay Fastly its hourly charges for Fastly engineering staff time. If it appears likely the factors are within Subscriber's 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.

## Support invoice credits

In the event a Severity 1 Incident occurs, Subscriber has purchased Gold or 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.

**NOTE:** Fastly maintains support for its [original Premium Support plan \(/guides/customer-support/legacy-premium-support-and-sla\)](/guides/customer-support/legacy-premium-support-and-sla). To convert your account to the current Gold and Platinum Support plans, contact [sales@fastly.com \(mailto:sales@fastly.com\)](mailto:sales@fastly.com).

• [Guides \(/guides/\)](/guides/) > [Compliance and law \(/guides/compliance-law\)](/guides/compliance-law) > [About our terms \(/guides/about-our-terms/\)](/guides/about-our-terms/)

• [Guides \(/guides/\)](/guides/) > [Compliance and law \(/guides/compliance-law\)](/guides/compliance-law) > [Compliance \(/guides/compliance/\)](/guides/compliance/)

## § HIPAA and caching PHI (/guides/compliance/hipaa-and-caching-phi)

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 (</guides/securing-communications/ordering-a-paid-tls-option>) and backend (</guides/securing-communications/connecting-to-origins-over-tls>) TLS to encrypt transmitted data from your origin to your end users.
- Add the `beresp.hipaa` variable (</guides/vcl/miscellaneous-VCL-extensions>) to objects containing PHI to keep that data out of non-volatile disk storage at the edge.

Contact [sales@fastly.com](mailto:sales@fastly.com) (mailto: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 (<https://www.fastly.com/terms>).

**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 \(/guides/compliance/security-program\)](/guides/compliance/security-program) and [technology compliance \(/guides/compliance/technology-compliance\)](/guides/compliance/technology-compliance) guides provide more information about these safeguards.

## § PCI-compliant caching (/guides/compliance/pci-compliant-caching)

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 (</guides/vcl/miscellaneous-VCL-extensions>) 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 (</guides/securing-communications/ordering-a-paid-tls-option>) and backend TLS (</guides/securing-communications/connecting-to-origins-over-tls>), this feature allows you to cache and transmit flagged content through the Fastly network in compliance with our PCI certification.

Contact [sales-ecommerce@fastly.com](mailto:sales-ecommerce@fastly.com) (mailto:sales-ecommerce@fastly.com) for more information on how to enable this feature for your account.

**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 \(/guides/compliance/security-program\)](/guides/compliance/security-program) and [technology compliance \(/guides/compliance/technology-compliance\)](/guides/compliance/technology-compliance) guides provide more information about these safeguards.

## § Security program (/guides/compliance/security-program)

Fastly's security program includes safeguards that help protect your data as it moves through the Fastly service. Information about these safeguards is organized by category. Our technology compliance (</guides/compliance/technology-compliance>) guide describes additional safeguards we maintain.

### Authentication and authorization

**User account assignment.** We assign individual user accounts to personnel who access Fastly systems and devices. These assignments help us monitor and enforce accountability of user activity.

**User-level privileges.** Our systems and devices enforce user roles or similar measures to control the extent of access we grant individual users.

**Multi-factor authentication.** We enforce multi-factor authentication to better secure our computing resources from unauthorized logins.

### Application security

**Secure software development.** We provide annual training to Fastly developers to help identify and prevent common software vulnerabilities, including the OWASP Top 10 ([https://www.owasp.org/index.php/OWASP\\_Top\\_Ten\\_Cheat\\_Sheet](https://www.owasp.org/index.php/OWASP_Top_Ten_Cheat_Sheet)). Developer code undergoes peer review prior to deployment, and internal security engineers and third-party security validators periodically analyze code for software components with higher potential security risk.

**Web application security review.** A third party assesses the security of the Fastly web application annually. We address findings from this assessment according to the risk they pose to the security of the Fastly service.

### Network and infrastructure security

**Network security reviews.** We regularly perform vulnerability scans and third-party penetration tests on the Fastly network. We review and address findings from these activities to help maintain the security of our network.

**Configuration standards.** We document and follow configuration standards to maintain secure systems and network devices. These standards include business justification for used ports, protocols, and services, as well as the removal of insecure default settings.

**Vulnerability and patch management.** To maintain awareness of potential security vulnerabilities, Fastly monitors public and private distribution lists, as well as reports submitted through our responsible disclosure (<https://www.fastly.com/security/report-security-issue>) process. We validate and implement security patches for critical vulnerabilities within 24 hours of discovery. For non-critical vulnerabilities and updates, we schedule and

deploy vendor-provided patches on a regular basis.

## Encryption

**Secure data transmission.** The Fastly service supports TLS configurations (</guides/securing-communications/>) to encrypt connections both externally to end users and backend origin servers, as well as internally within the Fastly network.

**Encryption key management.** We maintain technology and procedures to secure private keys throughout their lifecycle.

**Key storage and access security.** We store private keys in encrypted repositories, and we restrict key storage access to personnel who support our key management processes.

## Datacenter and physical security

**Physical access restrictions.** Our datacenters are fully enclosed with perimeter protection such as fences, gates, and mantraps to prevent unwanted entry. Only authorized people (including datacenter personnel, our employees, and contractors) may enter and move within a datacenter.

**Datacenter access management.** We ensure movement within our datacenters is monitored via onsite safeguards such as security guard assignment, facility access logging and review, and video surveillance. Additionally, we periodically review and adjust the list of personnel who may enter our datacenters.

**Secure asset installation.** We install computer and network hardware in locked cages and racks. Only authorized individuals may physically access this equipment.

**Environmental safeguards.** Our datacenters compensate for environmental disruptions with systems that control backup power, temperature and humidity, and fire suppression.

## Business continuity and operational resilience

**Service failover.** If any of our points of presence (POPs (</guides/about-fastly-services/fastly-pop-locations/>)) experience issues serving content, we can redirect traffic to a neighboring POP without interrupting the delivery of content to end users.

**Internet redundancy.** Our datacenters have connections with multiple Internet service providers. We do not rely on any single carrier to serve content to end users.

**Service monitoring.** We monitor multiple internal and external reporting channels to detect service-related issues. Personnel are available 24x7x365 to confirm and respond to disruptions of the Fastly service.

**Communication and reporting.** We update impacted customers using various communication methods (such as [status.fastly.com](https://status.fastly.com) (<https://status.fastly.com/>)), depending on an incident's scope and severity.

## Security incident management

**Incident response plan.** We maintain a formal incident response plan with established roles and responsibilities, communication protocols, and response procedures. We review and update this plan periodically to adapt it to evolving threats and risks to the Fastly service.

**Incident response team.** Representatives from key departments help address security-related incidents we discover. These personnel coordinate the investigation and resolution of incidents, as well as communication with external contacts as needed.

**Breach notification.** Fastly will notify affected customers within 48 hours of validating an unauthorized disclosure of customer confidential information.

## Logging and monitoring

**Log analysis.** We aggregate and securely store Fastly internal system activity. Monitoring these logs helps us discover and investigate potential security issues.

**Change and configuration monitoring.** We use multiple monitoring and alert mechanisms to enhance the visibility of technology changes and help ensure adherence to our change management process.

**Intrusion detection.** We maintain mechanisms to detect potential intrusions at the network and host level. Our Security department inspects and responds to events these detection measures discover.

## Customer and end user data management

**Cache data and configurations.** Customers manage which content is cached, where, and for how long by setting policies that control that content. See our introduction to caching (</guides/about-fastly-services/how-fastlys-cdn-service-works>) for more information. We may directly access or modify customer accounts or configurations to provide our services, prevent or address service or technical issues, as required by law, or as customers expressly permit. For the same reasons, we may also access or modify equipment, systems, or services that manage customer content.

**Client IP addresses.** As part of our caching network's general interaction with the Internet, Fastly independently collects anonymized and aggregated client IP address information on a limited basis to provide and improve its services. Client IP addresses are retained 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 (such as 503 "Service Unavailable" errors (</guides/debugging/common-503-errors/>)), and are discarded thereafter.

**Subscriber IP addresses.** Fastly independently collects the IP addresses of users who access their services within the Fastly web interface or through the API. We make these IP addresses available to customers through our audit log functionality (<https://www.fastly.com/blog/announcing-audit-logs>). If customers define origin servers or syslog endpoints with IP addresses, we save those IP addresses as part of their configurations. We may retain IP addresses from audit logs or configurations indefinitely. Dynamically-resolved origin IP addresses may be retained for up to two business days, or up to seven days if those addresses are associated with transmission errors (such as 503 "Service Unavailable" errors (</guides/debugging/common-503-errors/>)), and are discarded thereafter.

**IP addresses and security monitoring.** Fastly may retain indefinitely any non-anonymized, non-aggregated client or subscriber IP addresses associated with suspicious activity that may pose a risk to the Fastly network or our customers, or that are associated with administrative connections to the Fastly service.

**Content request data.** Content enters, transits, and departs our network in response to requests. We retain and use data about the operation and reliability of our processing of requests to monitor, maintain, and improve our services, our business operations, and our security and compliance programs. Subject to confidentiality obligations to our customers, we only disclose this data in anonymized and aggregated form.

**Subscriber log streaming.** Subscribers may stream syslog activity (</guides/streaming-logs/>), including end user IP addresses (</guides/basic-configuration/adding-or-modifying-headers-on-http-requests-and-responses#common-sources-of-new-content>), to a remote endpoint for analysis and use. Fastly does not retain subscriber syslog activity, except as described above.

---

## § Technology compliance (</guides/compliance/technology-compliance>)

---

Fastly's technology compliance program includes safeguards that help protect your data as it moves through the Fastly service. Information about these safeguards is organized by category. Our security program (</guides/compliance/security-program>) guide describes additional safeguards we maintain.

---

## Governance

**Information security roles and responsibilities.** We have formally assigned information security duties to Fastly personnel. Our Chief Security Officer and Security organization work with other departments to safeguard sensitive information related to the Fastly service.

**Policies and procedures.** Our policies and procedures help us maintain security in our systems, processes, and employee practices. Fastly's Security organization formally reviews these policies and procedures at least annually.

**Risk management.** We integrate risk assessment activities with various processes to identify and address information security risk to the company and customer data on our network.

**Vendor security oversight.** Fastly performs risk-based evaluations of the security measures of our vendors. We review these security measures before we begin using a vendor, and we ask the vendor to formally acknowledge these measures. We re-evaluate vendor security measures on a recurring basis thereafter.

## Human resources security

**Employee background screening.** We screen new employees as part of the hiring process. Screening activities depend on applicable local regulations and may include criminal background checks and reference checks.

**Confidentiality agreement.** Our employees formally agree to safeguard the sensitive information they may view, process, or transmit as part of their job functions.

**Security awareness training.** We train our people to protect the data and devices they use. Each employee receives security awareness training as part of new hire procedures, and current employees take this training annually.

## Data privacy

**Privacy policy.** Our privacy policy (<https://www.fastly.com/privacy>) describes how we collect, use, and protect the personal information of customer personnel using our websites or configuring services in the Fastly web interface. We certify our privacy policy and practices with TRUSTe (<https://privacy.truste.com/privacy-seal/Fastly/validation?rid=e864daae-e54e-4aba-9ae0-6a76f78e0962>).

**Personal data transfer.** The Fastly services by default do not process personal data. However, our service can be configured or used at the direction of the customer to process personal data. Our guide about our terms (</guides/about-our-terms/#can-i-use-transfer-data-from-the-eu-to-the-us-using-the-fastly-services>) provides additional information about data privacy compliance related to the processing of personal data.

## Technology change management

**Change management process.** We follow a defined set of procedures to develop and deploy technology changes. These changes include updates to software, configurations, and devices that support the Fastly service.

**Testing.** We test technology changes at various stages of development, and we confirm those tests are successful before completing a deployment into the Fastly service.

**Change approval and notification.** As part of our deployment process, we prepare, approve, and communicate change notices to maintain awareness among personnel who manage the Fastly network and systems.

**Post-implementation review.** We confirm the success of changes after their deployment. Should we experience issues during implementation, we also maintain procedures to revert changes.

## Identity and access management

**User requests and approval.** We document and approve requests for user access to the Fastly network. Our security administrators confirm appropriate documentation is in place before granting requested user rights.

**Access modification.** We promptly update or remove an employee's access to the Fastly network to match that employee's current job function or employment status.

**User access review.** We periodically inspect access privileges to make sure our personnel have appropriate access to Fastly systems and data.

• [Guides \(/guides/\)](/guides/) > [Compliance and law \(/guides/compliance-law/\)](/guides/compliance-law/) > [Security measures \(/guides/security-measures/\)](/guides/security-measures/)

• [Guides \(/guides/\)](/guides/) > [Compliance and law \(/guides/compliance-law/\)](/guides/compliance-law/) > [Third-party technology \(/guides/third-party-technology/\)](/guides/third-party-technology/)

• [Guides \(/guides/\)](/guides/) > [Compliance and law \(/guides/compliance-law/\)](/guides/compliance-law/) > [Related offerings \(/guides/related-offerings/\)](/guides/related-offerings/)

• [Guides \(/guides/\)](/guides/) > [Compliance and law \(/guides/compliance-law/\)](/guides/compliance-law/) > [Translations \(/guides/translations/\)](/guides/translations/)

• [Guides \(/guides/\)](/guides/) > [Compliance and law \(/guides/compliance-law/\)](/guides/compliance-law/) > [Archives \(/guides/archives/\)](/guides/archives/)

• [Guides \(/guides/\)](/guides/) > [Basic setup \(/guides/basic\)](/guides/basic/) > [Basic setup \(/guides/basic-setup/\)](/guides/basic-setup/)

## § Adding CNAME records (/guides/basic-setup/adding-cname-records)

To direct traffic from the Internet through Fastly instead of immediately through your site or application, you must create a DNS CNAME record ([https://en.wikipedia.org/wiki/CNAME\\_record](https://en.wikipedia.org/wiki/CNAME_record)) for your domain that points to a Fastly hostname (e.g., `nonssl.global.fastly.net`) instead of directly to your domain (e.g., `www.example.com`). To do this, you must have access privileges that allow you to modify DNS records for your domain.

★ **TIP:** Want to use Fastly on your apex domain (e.g., `example.com` rather than `www.example.com`) where you can't use a CNAME record? See our guide to [using Fastly with apex domains \(/guides/basic-configuration/using-fastly-with-apex-domains/\)](/guides/basic-configuration/using-fastly-with-apex-domains/) for more details.

## Choosing the right Fastly hostname for your CNAME record

The Fastly hostname you use for your DNS CNAME record will differ based on:

- the standard HTTPS (TLS) support requirements for your domain.
- any custom TLS options purchased for your domain.
- whether or not you choose to limit your traffic (</guides/performance-tuning/enabling-global-pops#limiting-pop-use-to-the-united-states-and-european-union>) to the US and EU network or use Fastly's global network (</guides/performance-tuning/enabling-global-pops>).

We've provided recommendations below based on these criteria.

## Non-TLS hostnames and limiting traffic

If you don't require TLS support and only need to accept HTTP (Port 80) connections, use:

- `nonssl.global.fastly.net.` to route traffic through Fastly's entire global network
- `nonssl.us-eu.fastly.net.` to route traffic through Fastly's US and EU POPs only

**❗ IMPORTANT:** Fastly's non-TLS hostnames refuse HTTPS connections (port 443) to prevent TLS certificate mismatch errors.

## TLS-enabled hostnames

If you plan to accept both HTTP (port 80) and HTTPS (port 443) connections and you're using Fastly's free shared TLS wildcard certificate (</guides/securing-communications/setting-up-free-tls>), use:

```
[name].global.ssl.fastly.net.
```

However, if you've purchased a Shared Certificate (</guides/securing-communications/ordering-a-paid-tls-option#shared-certificate>) or Shared Wildcard Certificate (</guides/securing-communications/ordering-a-paid-tls-option#shared-wildcard-certificate-service>) service, use:

- `[letter].ssl.global.fastly.net.` to route traffic through Fastly's entire global network
- `[letter].ssl.us-eu.fastly.net.` to route traffic through Fastly's US and EU POPs only

When you purchase one of these certificate services, Fastly Support will add your domains to a specific TLS Certificate, usually differentiated by a certificate letter (e.g., `a`, `a2`, `b`, `c`). You'll need to add the certificate letter to the beginning of the Fastly hostname noted above for use in your CNAME record. For example, if your domain was added to our `a` certificate the above hostname would become `a.ssl.global.fastly.net.`

**❗ IMPORTANT:** You must use the assigned Fastly TLS hostname provided by Fastly Support. Using the incorrect Fastly hostname will cause a TLS Certificate mismatch error for HTTPS (Port 443) traffic.

## Customer-specific hostnames

If you've purchased our Customer certificate hosting (</guides/securing-communications/ordering-a-paid-tls-option#customer-certificate-hosting>) or SNI customer certificate hosting (</guides/securing-communications/ordering-a-paid-tls-option#sni-customer-certificate-hosting>) option, we'll assign you to a specific map that uses the following format:

```
[name].map.fastly.net.
```

## Updating the CNAME record with your DNS provider

Once you've determined the appropriate Fastly hostname for your domain, create a DNS CNAME record for your domain. You'll need to follow your DNS provider's documentation for creating or updating a DNS CNAME record. The steps you follow will vary depending on your DNS provider's control panel interfaces.

★ **TIP:** If you can't find your provider's CNAME configuration instructions, Google maintains instructions [for most major providers](https://support.google.com/a/topic/1615038?hl=en) (<https://support.google.com/a/topic/1615038?hl=en>). Keep in mind that these instructions are maintained by Google, not Fastly, and are tailored specifically for Google enterprise services.

If you run your own DNS server or are familiar with the format of BIND zone files, the CNAME record would look similar to this:

```
www.example.com.      3600      IN      CNAME    nonssl.global.fastly.net.
```

In the above example, the domain set up on Fastly is `www.example.com.`, with a time-to-live (TTL) of `3600` seconds (1 hour), the Record Type is `CNAME`, and the Fastly hostname is `nonssl.global.fastly.net.` because TLS support isn't required and traffic will be routed through Fastly's entire global network.

★ **TIP:** Test an updated CNAME using `dig` on your computer's terminal (e.g., `dig www.example.com`). If you see the correct Fastly hostname in the CNAME record in the `ANSWER` section of the response, then the CNAME record is created properly.

## Best practices when updating a DNS CNAME record

- Be sure you've added all domains you want served by Fastly to the appropriate service. If you don't and you point your domain to Fastly, an `unknown domain` error will occur.


- Make sure your service is properly configured. You can test a Fastly service on your local machine by using cURL (</guides/debugging/curl-and-other-caching-verification-methods>) and replacing `192.168.1.1` with a Fastly IP address or setting a temporary `/etc/hosts` entry on your local machine.
- If you have multiple hostnames on the same domain (e.g., `api.example.com`, `www.example.com`, `app.example.com`), you can use a DNS wildcard record (`*.example.com`) at your DNS provider so only a single CNAME record is created and maintained. You should also add either a matching `*.example.com` domain or the individual domains to your Fastly service.
- Before changing a CNAME to point to Fastly's hostname, change your service configuration to lower the CNAME's TTL to a small number (we suggest 60 seconds) and wait for the old TTL to expire. Creating a DNS CNAME record for your domain after the TTL expiration ensures you have an easy way to roll back changes if you encounter an issue. Once you confirm everything is working properly using Fastly, you can increase the TTL to its original number.

## § Getting started with Fastly (</guides/basic-setup/getting-started-with-fastly>)

In this article, we explain what Fastly does and how best to use it with your site.

### How Fastly works

Fastly works by storing the content of your website on servers all over the world and quickly delivering that content to your users. We do this using Varnish (<https://www.varnish-cache.org>), an open source web application accelerator.



We track the geo-location of each user and make sure they are connecting to a server that is closest to them. This makes your site faster by reducing the time spent waiting for data to be sent from the server to the user.


We also give you full control over when and how we store content from your servers. You can set a Time To Live (TTL) for any path on your site and instantly invalidate or purge any path on your site using our Purge API (</api/purge>).

By using these tools, you only have to generate pages one time for the site for many millions of page views. This saves time for your users and costs on your server bills.

### Advanced features

Fastly also provides many advanced features that help you monitor how your data is accessed and customize your content delivery.





- **Instant Purging** (</guides/purging/>) allows you to have better control over when and how content is updated. You can update your data when you want and as often as you want, rather than waiting up to 24 hours to change data at the edge.
- **Real Time Analytics** (</guides/about-fastly-services/about-the-web-interface-controls>) provides a top level view of your network and how your site is performing. Every second, we compile the relevant data about all of your traffic into an easy-to-read report.
- **Conditions** (</guides/conditions/>) change how requests are routed, what headers to send, and how content is cached.
- **Health Checks** (</guides/basic-configuration/health-checks-tutorial>) monitors the status of one or many of your back end servers. This way if anything goes wrong with your servers, you immediately know about it.
- **Streaming Logs** (</guides/streaming-logs/setting-up-remote-log-streaming>) are quickly and easily configured to send information from your servers anywhere and in the format you want.
- **Varnish Configuration Language (VCL)** (</guides/vcl/guide-to-vcl>) allows you to modify nearly every aspect of an HTTP request and response. You can upload VCL files with specialized configurations to your account.
- The **Fastly API** (</api/>) can programmatically handle your configuration. This allows you to write scripts to handle basic configuration tasks and create your own administrative views (so they can be directly coupled with your existing admin software).

## Getting started

If you do not have an account, sign up now (<https://www.fastly.com/signup>). Feel free to choose the developer plan so you can test how Fastly works on your site.

If you want to use Fastly but do not know where to begin, check the Basic setup (</guides/basic-setup/>) documentation. You can learn everything you need to set up and configure your first service for your site.

If you want to explore more, check the Basic configuration documentation to learn more about caching, and features such as purging (</guides/purging/>) and shielding (</guides/performance-tuning/shielding>).

If you want information on advanced features, especially related to things like load balancing (</guides/performance-tuning/load-balancing-configuration>) or the Varnish Control Language (</guides/vcl/>) that we support, check the advanced configuration section of our help files or the API Reference (</api/>), which includes a full reference to the Fastly API.

If you are having problems, send us a message at [support@fastly.com](mailto:support@fastly.com) (<mailto:support@fastly.com>).

## § Glossary of terms (</guides/basic-setup/glossary-of-terms>)

These are common Fastly, HTTP, and networking terms you may encounter within our service guides.

### ACL

Access Control List. A list of permissions that can be attached to an object (</guides/vcl/using-access-control-lists>) allowing customers to quickly check a client's IP against a list of known net blocks and then make decisions based on the result.

### Altitude

Fastly's customer summit (<https://www.fastly.com/altitude>).

### Backend

See *origin server*.

### Cache-Control

The specific HTTP header (</guides/tutorials/cache-control-tutorial>) that controls who can cache a response, under which conditions, and for how long. Fastly respects `Cache-Control` headers returned from origin servers as one approach to cache management. See also *surrogate-control*, *max-age*, and *s-maxage*.

### community.fastly.com

Fastly's community discussion forum (<https://community.fastly.com/>).

### Cookie

HTTP headers used to perform certain functions like authenticating login in secure website areas, information tracking, remembering user preferences, and customizing how information is presented.

### cURL

An open-source command line tool (<https://curl.haxx.se/>) for transferring data with URL syntax from or to a server using one of many supported protocols. Fastly users can issue cURL commands to verify requests are caching (</guides/debugging/curl-and-other-caching-verification-methods>) in the Fastly network.

### DNS

Domain Name System. A system for naming computers and network services that translates a domain's numbered IP address into an easy-to-remember alphabetic name.

### Edge Dictionary

A type of container Fastly users can create (</guides/edge-dictionaries/>) to store data as key-value pairs and turn frequently repeated statements into a single function that acts as constant.

### Egress traffic

Bandwidth used when traffic travels from Fastly points of presence (POPs) to the end user.

### ESI

Edge Side Includes. An XML-based markup language (<http://www.w3.org/TR/esi-lang>) that allows content assembly by HTTP surrogates. Allows Fastly users to cache pages that contain both cacheable and uncacheable content (such as user-specific information).

### Gzip

A way of compressing information to make it faster to transmit. Fastly allows users to dynamically gzip content (</guides/basic-configuration/enabling-automatic-gzipping>) based on file extension or content type.

### Header

An HTTP field that precedes the main content of information being sent in a request or response and describes the length of the content, type of content, or other characteristics of the information.

### Host (header)

Information used in addition to the IP address and port number to uniquely identify a domain. See also *Normalizing the host header* (<https://www.fastly.com/blog/varnish-tip-normalize-host-header>).

### Ingress traffic

Bandwidth used when end users make requests that send traffic to Fastly points of presence (POPs).

### Instant Purge

A feature of Fastly's purging functionality (</guides/purging/>) that allows users to actively invalidate content in Fastly caches within milliseconds. See also *Soft Purge*.

### manage.fastly.com

The web interface through which customers access Fastly's CDN services (<https://manage.fastly.com/>).

### max-age

An HTTP Cache-Control directive that specifies how long (in seconds) an object will remain in the cache before Fastly removes the object from storage. See also *surrogate-control*, *cache-control*, and *s-maxage*.

### MTR

A tool that combines traceroute and ping programs in a single network diagnostic tool. Frequently used in debugging network connections (</guides/debugging/debugging-with-mtr>).

### Origin server

The location or address from which Fastly's network requests the content it will serve.

### Origin Shield (Shielding)

A specific Fastly point of presence (POP) designated by users (</guides/performance-tuning/shielding#enabling-shielding>) as the primary source of content through which all content requests from other POPs will be directed in lieu of contacting a customer's origins directly.

## OTFP

On-the-fly packaging. A feature of Fastly's Video on Demand (<https://www.fastly.com/services/video-demand>) media and streaming services that allows customers to dynamically package video for delivery in multiple HTTP streaming formats. Also known as "just in time" video content packaging.

## POP

Point of Presence. Datacenter within which Fastly's globally distributed cache servers (<https://www.fastly.com/network-map>) reside.

## Priority

A setting that allows users to specify the order request and cache settings execute within their subroutines. The Priority can be any whole number and always default to 10. The smaller the assigned priority number, the sooner that condition executes (e.g., 1 executes sooner than 10).

## private

An HTTP Cache-Control directive that allows users to select which objects are not cached. Fastly will not cache responses with a Cache-Control value of `private`.

## Purging

The process of picking out one or more objects from the Fastly cache and discarding it along with its variants. See also *Instant Purge* and *Soft Purge*.

## Redirect

A function that directs requests for information from their originally intended locations to a more desirable destination (</guides/performance-tuning/generating-http-redirects-at-the-edge>).

## Set-Cookie

The header sent by a server in response to an HTTP request and then used to create a cookie on a user's origin. Fastly supports a method for extracting a named value (</guides/vcl/response-cookie-handling>) out of `Set-Cookie` headers no matter how many there are. By default, Fastly will not cache responses that contain a `Set-Cookie` header.

## s-maxage

An HTTP cache control directive similar to *max-age*, but applied only to shared caches. See also *surrogate-control* and *cache-control*.

## Soft Purge

A type of purging (</guides/purging/>) that allows users to easily mark content as outdated (</guides/purging/soft-purges>) (expired) instead of immediately deleting it from Fastly's caches. See also *Instant Purge*.

## status.fastly.com

Fastly's network status (<https://status.fastly.com/>) monitoring site. Allows customers to quickly check whether anomalies they see may be due to a known problem currently being worked on by Fastly or if their issues more likely stem from problems within their own infrastructure.

## support@fastly.com

The main email address of Fastly's Customer Support (<https://www.fastly.com/support>) team through which customers can ask questions and receive assistance.

## Surrogate-Control

An HTTP response header that allows origin servers to use control directives to dictate how intermediate caches, including Fastly, should handle response entities. `Surrogate-Control` will not affect browsers. See also *cache-control*, *max-age*, and *s-maxage*.

## Surrogate Key

A unique identifier that allows customers to group content together for faster processing. Fastly uses surrogate keys (<https://www.fastly.com/blog/surrogate-keys-part-1>) as part of its purging strategy (</guides/purging/>).

## Synthetic response

Custom responses generated within the CDN that users can set if a specific URL is requested or a specific condition, such as a status code, is met. These responses require no origin server interaction.

## TLS (SSL)

A cryptographic protocol Fastly follows (</guides/securing-communications/>) that ensures privacy between communicating applications and their users on the internet.

## URL

Uniform Resource Locator. An address used (</guides/about-fastly-services/domain-names-and-fastly-services>) to find a site or application's objects on the internet.

## Varnish

Caching software (<https://www.fastly.com/blog/benefits-using-varnish>) that helps content-heavy dynamic websites as well as heavily consumed APIs load faster. Fastly's core caching infrastructure is based on a heavily modified version of Varnish.

## VCL

Varnish Configuration Language. A scripting language (</guides/vcl/guide-to-vcl>) used to configure and add logic to Varnish caches. Fastly users can create custom VCL files with specialized configurations.

## [www.fastly-debug.com](https://www.fastly-debug.com)

A network debugging tool designed to provide key info to help a Fastly user troubleshoot issues with Fastly's Customer Support (<https://www.fastly.com/support>) team.

# § Sign up and create your first service (</guides/basic-setup/sign-up-and-create-your-first-service>)

To create a Fastly account and set up your first service, follow the steps below.

## Sign up at Fastly.com

Before you do anything else, you must sign up for a Fastly account.

1. Click on any **Sign Up** button on the Fastly.com website or simply point a browser to the signup form (<https://www.fastly.com/signup>).
2. When the signup form appears, fill in all the fields with your contact information. All the fields are required.

First Name

Last Name

Email

Password ⓘ

Company


Phone

SIGN UP FOR AN ACCOUNT

By clicking on Sign Up you agree to the [Terms of Service](#).

**ⓘ NOTE:** You'll be able to [change your password and email address \(/guides/user-access-and-control/email-and-password-changes\)](/guides/user-access-and-control/email-and-password-changes) any time after signup. Without a valid email we can't send you account verification details during account setup. Without a valid telephone number, we can't assist you with specific kinds of [account lockout issues \(/guides/account-management-and-security/enabling-an-ip-whitelist-for-account-logins\)](/guides/account-management-and-security/enabling-an-ip-whitelist-for-account-logins).

3. Click the **Sign Up For An Account** button. The confirmation screen will appear with instructions on what to do next and you'll be sent an e-mail that contains a verification link.



**Success! You're almost ready for liftoff.**  
**Check your email [email@example.com](mailto:email@example.com) to confirm your account.**

If you haven't received a confirmation email in 10 minutes, please email us [support@fastly.com](mailto:support@fastly.com)

4. Check your inbox and find the confirmation email we sent you.
5. Click the verification link (we need to make sure you're not a spam robot and verify your email). The verification link will immediately take you to the first step of the quick start process so you can create your first service.

## Create your first service

Once you've verified your email, we log you into the application automatically and immediately take you through the quick start process to create your first service.

1. Specify your **domain** and **hostname**. We need this information to properly route requests to your website and to make sure your cache updates properly. If you have many servers you can configure the rest later.



Your website domain for Fastly to use when routing requests.  
For example: [www.example.com](http://www.example.com) or [\\*.example.com](http://*.example.com)


 : 

The hostname (or IP address) and port number for your origin server. [Help with Amazon S3](#) and [Google Cloud Storage](#)

2. Click the **Continue** button.

## Recommendations to enhance your service


All three are included in your Fastly service, and you can add them at any time.



Compress your data for better performance

[+ ENABLE GZIP](#)


[Learn more](#)



Send logs about your Fastly cache activity

[+ ENABLE LOGGING](#)

[Learn more](#)



Monitor the status of your origin server

[+ ENABLE HEALTH CHECK](#)

[Learn more](#)

[CONTINUE >](#)

### 3. Optionally enable the following features:

- Gzip (/guides/basic-configuration/enabling-automatic-gzipping) to dynamically gzip content, resulting in smaller file sizes and faster transfer speeds.
- Logging (/guides/streaming-logs/setting-up-remote-log-streaming) to automatically store logs with a third-party service, resulting in the ability to view cache activity and to troubleshoot.
- Health checks (/guides/basic-configuration/health-checks-tutorial) to monitor the health of your origin server so that Fastly will stop attempting to send requests to it when it is unhealthy.


### 4. Click the **Continue** button. The system configures your service.

5. To test whether or not the configuration was successful, open **yourdomain.global.prod.fastly.net** in a new browser window (replace `yourdomain` with your own website's domain name). Your website should appear, but it may take up to 60 seconds for the new configuration settings to take effect.

## Add CNAME DNS records

You must set the CNAME DNS record for your domain to point to Fastly. For more information, see the instructions in our Adding CNAME records (/guides/basic-setup/adding-cname-records) guide.

### Not quite ready?




Go explore the rest of Fastly. You can go live at any time.

A link to the go live instructions will be in the top navigation.

[EXPLORE >](#)

### Go live!



You have to take one last step: Use your DNS provider to point your CNAME to Fastly.

After that's done,

[VIEW YOUR DASHBOARD >](#)

OR

Once you've completed this step, you should be all ready to go!

If you have any problems, feel free to contact us at [support@fastly.com](mailto:support@fastly.com) (<mailto:support@fastly.com>).

## § Working with services (/guides/basic-setup/working-with-services)

The Fastly web interface allows you to create new services or edit existing ones and then activate your changes once you have things configured. The web interface also allows you to do other things with existing services, like rename them, compare them to each other, deactivate or reactivate them, and delete them.

### Creating a new service

To create a new service, follow the steps below:

1. Log in to the Fastly web interface and click the **Configure** link.
2. Click the **Options** button to open the **Manage service** menu and select **Create**. The Create a new service page appears.

### Create a new service

[Or use the setup guide](#)

Server address and port

:

★ Required

The IP address (or hostname) and port number for your origin server. To [enable TLS](#), use port 443.

[▶ Setting the server address and port](#)

Domain

★ Required

The domain name of your website.

[▶ Setting the domain name](#)

Description

★ Required

The name of your service, such as **My service**.

CREATE

CANCEL

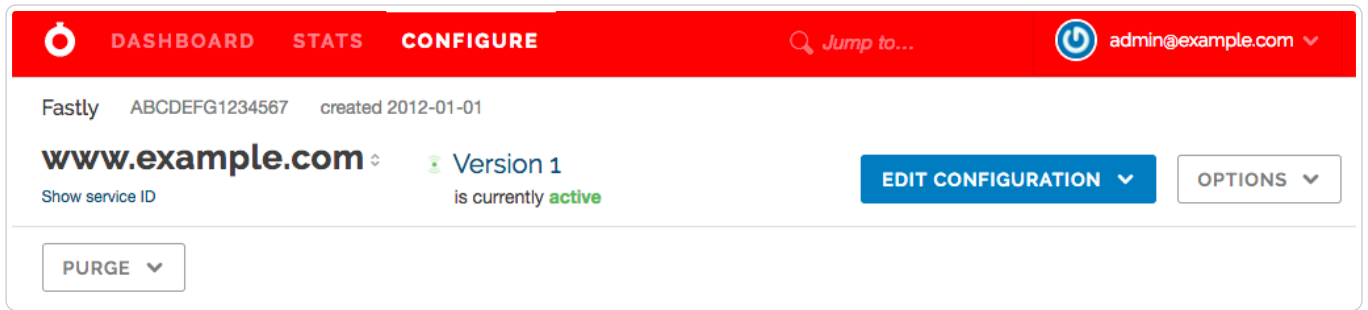
3. Fill out the **Create a new service** page as follows:
  - In the **Server address and port** field, type the IP address (or hostname) and port number (usually `80`) for your website's server.
  - In the **Domain name** field, type the domain name of your website. We need this information to properly route requests to your website.
  - In the **Description** field, type a human-readable description for the new service. You can change this description at any time.
4. Click the **Create** button. The new service appears in the list of services available.

### Editing and activating versions of services

Fastly locks versions of services you've already activated to make rollbacks safer and provide version control. You can duplicate ("clone") any existing service version, active or inactive, and edit that cloned version. You must "activate" new versions of services in order to deploy their configurations. Configuration changes are never automatically activated.

To make changes to a service and activate a new version, follow the steps below:


1. Log in to the Fastly web interface and click the **Configure** link. The Configure page appears.



2. Click the **Edit Configuration** button. You can select **Clone active** to clone the active version of the service for editing, or you can select **Edit latest** to edit the latest draft of the service.



The service version page appears.



3. Click **Activate**. The new version of the service is activated.

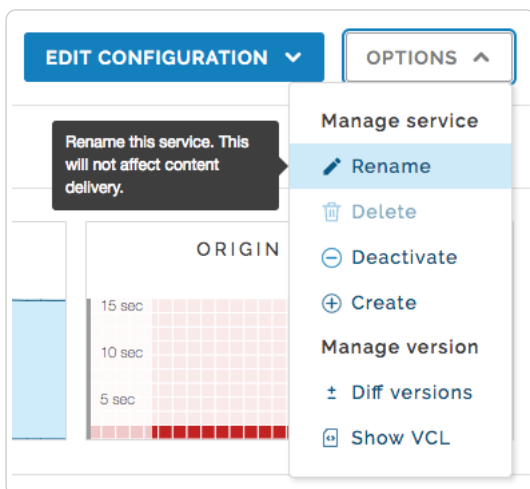
## Other things you can do with services

In addition to creating or editing services, you can rename them, compare versions of them, deactivate or reactivate specific version of them, and delete them.

### Renaming services


To rename your service, follow the steps below:

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Options** button to open the **Manage service** menu and select **Rename**.



The Rename your service pane appears.





Rename your service

★ Required

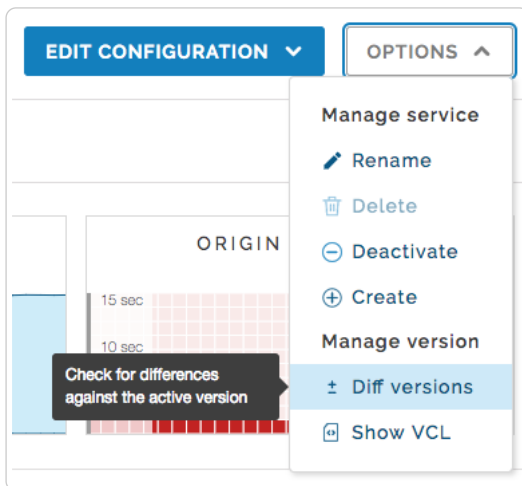
**RENAME**

4. Type the new service name and then click **Rename**. The new service name appears.


## Comparing different service versions

To compare two versions of a service, follow the steps below:

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Options** button to open the **Manage service** menu and select **Diff versions**.



The Diff versions page appears. Removals are highlighted in red with a minus sign at the beginning of the line. Additions are highlighted in green with a plus sign at the beginning of the line.



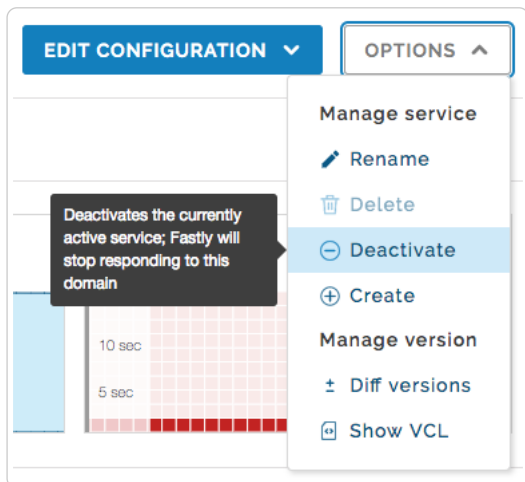
The screenshot shows the Fastly 'Diff' page. The top navigation bar includes 'DASHBOARD', 'STATS', and 'CONFIGURE'. The main header shows 'Fastly ABCDEFG1234567 created 2012-01-01' and a 'Return to service summary' link. The service name 'www.example.com' is displayed. The 'Diff' section shows two versions: 'Version 1' (active) and 'Version 4' (in development). The diff content shows a removal of 'name: 1.1.1.1' (highlighted in red) and an addition of 'name: My origin server' (highlighted in green). The address '1.1.1.1' and 'auto\_loadbalance: false' are also shown.

You can change the compared service versions by selecting a different version number in the selection menus.

## Deactivating a service

To deactivate a service, follow the steps below:

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Options** button to open the **Manage service** menu and select **Deactivate**.



You can also activate or deactivate a service via the API (`/api/config#version`). Did you accidentally delete a service? We can help.

## Reactivating a service

To reactivate a service, follow the steps below:

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.
4. Click **Activate**. The service is reactivated.

## Deleting a service

Fastly allows you to delete any service you create, along with all of its versions. Fastly does not offer a way to delete specific versions of a service, however. Service versions are meant to be an historic log of the changes that were made to a service. To undo changes introduced by a particular service version, you can always go back to a previous version and reactivate it or clone a new service version based on any old version.

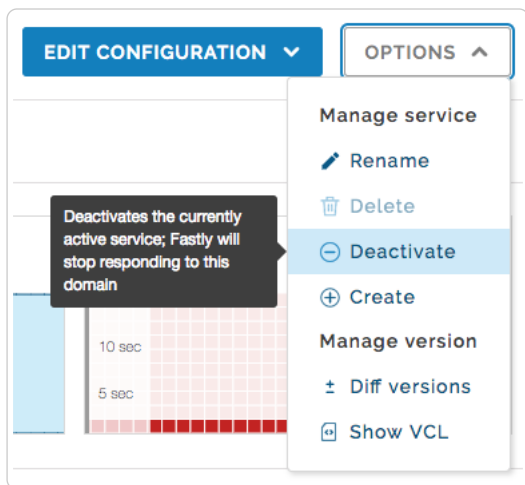
**ⓘ NOTE:** When a service is deleted from your configuration, Fastly maintains a copy in its databases in case the deletion was accidental. Any domains that were associated with that service will remain associated with it even after deletion. Before deleting a service, you should delete the domains within it first and then delete the service itself.

To delete any service along with all of its versions, follow the steps below:

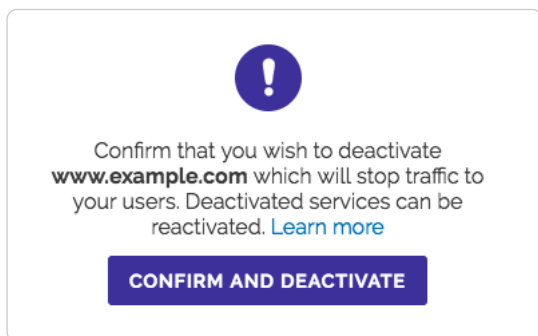
1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Edit configuration** button and then select **Clone active**. The Domains page appears.
4. In the **Domains** area, delete each domain associated with this service (you'll be asked to confirm each deletion).

**ⓘ IMPORTANT:** Be sure to repeat this step until all domains associated with the service have been deleted. If any domains remain associated with a deleted service, the system may not allow you to associate them with a new service at a later date.

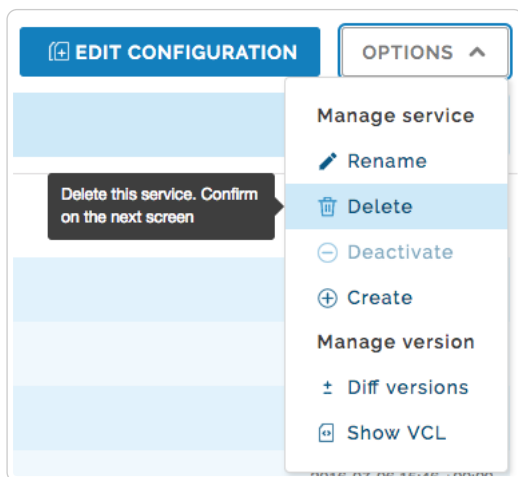
5. Click the **Configure** link at the top of the window.
6. Click the **Options** button to open the **Manage service** menu and select **Deactivate**.



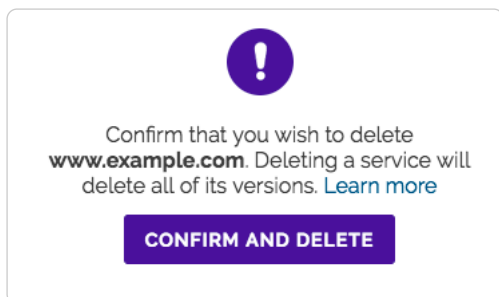
The deactivate service warning appears.



7. Click the **Confirm and deactivate** button to confirm you want to deactivate your service and acknowledge that you no longer want to serve traffic with it.
8. Click the **Options** button again to open the **Manage service** menu and select **Delete**.



The confirm delete window appears.



9. Click the **Confirm and delete** button to confirm that you want to delete the service.

## Getting help with accidental service deletions

Services can be deactivated or deleted. Deactivated services can be reactivated at any time, but once they've been deleted you must contact Customer Support (mailto:support@fastly.com) to have them restored. When sending your request, remember to include:

- your customer ID (/guides/account-management-and-security/finding-and-managing-your-account-info#finding-your-customer-id)
- your company name
- your service ID (/guides/account-management-and-security/finding-and-managing-your-account-info#finding-your-service-id) (the name of the service you want restored)

Customer Support will notify you when your service has been restored.

• Guides (/guides/) > Basic setup (/guides/basic) > Basic configuration (/guides/basic-configuration/)

## § Adding or modifying headers on HTTP requests and responses (/guides/basic-configuration/adding-or-modifying-headers-on-http-requests-and-responses)

HTTP header fields are components of the header section of request and response messages in the Hypertext Transfer Protocol (HTTP). They define the operating parameters of an HTTP transaction. When you create and configure headers, you can determine how you want your content served to your users. The following steps show you how to add and edit headers.

### Create new headers

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.
4. Click the **Content** tab. The Content page appears.

**DOMAINS**  
1 Domain

**ORIGINS**  
1 Host  
1 Health check

**SETTINGS**  
1 Request setting  
0 Cache settings

**CONTENT**  
0 Headers  
0 Gzips  
0 Responses

**LOGGING**  
1 Endpoint

**CUSTOM VCL**  
0 Files

## Headers

*There are no headers.*

**CREATE YOUR FIRST HEADER**

## Gzip

*Gzip is not enabled.*

**SET UP GZIP**

## Responses

*There are no responses.*

**CREATE YOUR FIRST RESPONSE**

### About content

Headers allow you to set, delete, and change request and response headers as your information is handled by Fastly.

Responses allow you to create synthetic responses that exist entirely on the cache servers.

5. Click the **Create header** button. The Create a new header window appears.

## Create a new header

Learn more about this section in our [headers tutorial](#).

CONDITION

This will happen all the time unless you [Attach a condition](#)

Type / Action

Request

Set

The type of header and the action performed on it.

Destination

★ Required

The name of the header that will be affected by the selected action. For example: **http.Content-Type**, **http.Set-Cookie**, **http.Via**, **http.Location**, or **http.Access-Control-Allow-Origin**.

Source

★ Required

New content for the header. Can be a static value (e.g. string or number) or a dynamic value (e.g. existing header or a GeoIP value). Please use quotes for string values.

Ignore if set

No

If switched to **Yes**, the action will not be performed if the header in Destination exists.

Priority

10

★ Required

The order in which the header rules execute within the condition.

Description

★ Required

The name of your header, such as **My header**.

CREATE

CANCEL

6. Fill out the **Create a new header** fields as follows:

- From the **Type** menu, select **Request**, and from the **Action** menu, select **Set**.
- In the **Destination** field, type the name of the header affected by the selected action.
- In the **Source** field, type where the content for the header comes from.
- From the **Ignore if set** menu, select **No** if you want the header in the **Destination** field modified or select **Yes** if you don't want it modified.
- In the **Priority** field, type the order the header rules execute.
- In the **Description** field, type the name of your header rule (for example, `My header`).

The Field description table below provides additional details about each of these controls.

7. Click the **Create** button.

8. Click the **Activate** button to deploy your configuration changes.

## Edit headers

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.
4. Click the **Content** tab. The Content page appears.
5. Click the name of the header you want to edit. The Edit a header page appears.

## Edit a header

Learn more about this section in our [headers tutorial](#).

### CONDITION

This will happen all the time unless you [Attach a condition](#)

Type / Action

Request

Set

The type of header and the action performed on it.

Destination

1.1.1.1

★ Required

The name of the header that will be affected by the selected action. For example: **http.Content-Type**, **http.Set-Cookie**, **http.Via**, **http.Location**, or **http.Access-Control-Allow-Origin**.

Source

string

★ Required

New content for the header. Can be a static value (e.g. string or number) or a dynamic value (e.g. existing header or a GeolIP value). Please use quotes for string values.

Ignore if set

No

If switched to **Yes**, the action will not be performed if the header in **Destination** exists.

Priority

10

★ Required

The order in which the header rules execute within the condition.

Description

My header

★ Required

The name of your header, such as **My header**.

UPDATE

CANCEL

6. Fill out the **Edit a header** fields as follows:

- From the **Type** menu, select **Request**, and from the **Action** menu, select **Set**.
- In the **Destination** field, type the name of the header affected by the selected action.
- In the **Source** field, type where the content for the header comes from.
- From the **Ignore if set** menu, select **No** if you want the header in the **Destination** field modified or select **Yes** if you don't want it modified.
- In the **Priority** field, type the order the header rules execute.
- In the **Description** field, type the name of your header rule (for example, `My header`).

7. Click the **Update** button.

8. Click the **Activate** button to deploy your configuration changes.

## Field description table

This table describes what each field in the Header window means:

Field	Description
Type	The Type menu can be set to <b>Request</b> , <b>Response</b> , or <b>Cache</b> . Selecting <b>Request</b> modifies the request coming from the user, and this will carry through to the request that gets sent to your origin server. Selecting <b>Response</b> affects the HTTP response that is sent back to the user. Selecting <b>Cache</b> affects the HTTP response that your origin server returns before it gets stored on Fastly servers, meaning whatever changes you make there will be remembered on a cache hit.

Field	Description
<b>Action</b>	The Action menu can be set to <b>Set</b> , <b>Append</b> , <b>Delete</b> , <b>Regex</b> , and <b>Regex All</b> . Selecting <b>Set</b> (the default) will write a value into the header (potentially overwriting it, if it already exists). Selecting <b>Append</b> will add a value onto the end of a header or set it if it doesn't exist. Selecting <b>Delete</b> will remove a header. When selected, it hides the Source field in the Header window. Selecting <b>Regex</b> allows you to perform a find and replace on specific text and is based on a regular expression you type in. When selected, the Regex and Substitution controls appear in the Header window. Selecting <b>Regex All</b> allows you to perform the same function as <b>Regex</b> but it performs a find and replace multiple times. When selected, the Regex and Substitution controls appear in the Header window.
<b>Destination</b>	The Destination field determines the name of the header that is going to be affected by our Action. Because header rules can be used to affect more than just HTTP headers, your input to this field should be formatted like this: <code>http.Header-Name</code> .
<b>Source</b>	The Source field is available on Set, Append, Regex, and Regex All actions. This field becomes hidden in the Header window when you select Delete from the Action menu. It determines where the new content for the header comes from. There are a plethora of options for Source. The simplest is a static string such as <code>"My Static String"</code> (including the quotes). Other options include <code>client.ip</code> , <code>req.http.Another-Header</code> , and <code>geoip.city</code> . See the list of Common Sources below for more common sources of new content.
<b>Regex</b>	The Regex field only appears in the Header window when you select Regex or Regex All from the Action menu. It allows you to perform a find and replace on specific text and is based on a regular expression that you type in.
<b>Substitution</b>	The Substitution field only appears in the Header window when you select the Regex and Regex All from the Action menu. It replaces the text that was removed by the regex expression with the text you typed in the Substitution field.
<b>Ignore if set</b>	By default this is set to No, which means that if the header you are modifying already exists, it will be modified.
<b>Priority</b>	The Priority field determines the order in which the header rules execute (e.g., a priority of 1 means the header rule executes first). This can be important if you set headers and then set other headers based on the earlier ones.
<b>Description</b>	The Description field specifies a memorable word or phrase that allows you to recognize and remember a particular Header rule.

## Common sources of new content

Name	Valid Types	Description
<code>req.http.Fastly-Client-IP</code>	Request, Cache, Response	The true IP address of the client.
<code>client.ip</code> and <code>client.identity</code>	Request, Cache, Response	The client IP address. These variables are available, but may not always display the source IP address. For instance, they may show the edge node IP when shielding is enabled. For the true client IP address use <code>req.http.Fastly-Client-IP</code> .  <b>IMPORTANT:</b> In some cases, client IP data may be considered sensitive. Make sure you protect the sensitive IP data you stream or store.
<code>server.identity</code>	Request, Cache, Response	A unique identifier for the Fastly server processing the request.
<code>server.region</code>	Request, Cache, Response	The region in which the Fastly server resides.
<code>server.datacenter</code>	Request, Cache, Response	The datacenter in which the Fastly server resides.
<code>req.url</code>	Request, Cache, Response	The URL of the HTTP Request from the client.
<code>req.http.*</code>	Request, Cache, Response	The headers from the HTTP Request, access as: <code>req.http.HeaderName</code>
<code>beresp.status</code>	Cache	The status returned from the origin server.
<code>beresp.http.*</code>	Cache	The headers from the origin's HTTP Response, access: <code>beresp.http.HeaderName</code>
<code>resp.status</code>	Response	The status that is going to be returned to the client.



Name	Valid Types	Description
<code>resp.http.*</code>	Response	The headers in the HTTP Response to be returned to the client, access: <code>resp.http.HeaderName</code>
<code>geoip.*</code>	Request, Cache, Response	GeoIP values for the client's IP (see our GeoIP article (/guides/vcl/geoip-related-vcl-features) for more information).

## § Creating and customizing a robots.txt file (/guides/basic-configuration/creating-and-customizing-a-robots-file)

Once you've created a CNAME DNS record (/guides/basic-setup/cname-instructions-for-most-providers) to direct your domain's traffic to Fastly, you can set custom responses (/guides/vcl/custom-responses-that-dont-hit-origin-servers) that will be served by the robots.txt file ([https://en.wikipedia.org/wiki/Robots\\_exclusion\\_standard](https://en.wikipedia.org/wiki/Robots_exclusion_standard)) on your website. To create and configure your robots.txt file via the Fastly web interface, follow the steps below:

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.
4. Click the **Content** tab. The Content page appears.

The screenshot shows the Fastly web interface with the 'CONTENT' tab selected. The top navigation bar displays various configuration options: DOMAINS (1 Domain), ORIGINS (1 Host, 1 Health check), SETTINGS (1 Request setting, 0 Cache settings), CONTENT (0 Headers, 0 Gzips, 0 Responses), LOGGING (1 Endpoint), and CUSTOM VCL (0 Files). The 'CONTENT' tab is highlighted. The main content area is divided into three sections: 'Headers' with a 'CREATE YOUR FIRST HEADER' button, 'Gzip' with a 'SET UP GZIP' button, and 'Responses' with a 'CREATE YOUR FIRST RESPONSE' button. An 'About content' sidebar on the right provides additional information: 'Headers allow you to set, delete, and change request and response headers as your information is handled by Fastly.' and 'Responses allow you to create synthetic responses that exist entirely on the cache servers.'

5. Click the **Create response** button. The Create a new synthetic response page appears.

## Create a new synthetic response

More on how synthetic responses work in our [tutorial](#).

### CONDITION

This will happen all the time unless you [Attach a condition](#)

Status

200 OK

The HTTP status code to include in the header of the response.

MIME Type

text/html

The media type to put into the Content-Type header which informs the browser how to display the response. Common MIME types are **application/json**, **text/plain**, **text/html**.

Response

User-agent: \*  
  
Disallow: /tmp/\*  
Disallow: /foo.html

The content to be served when delivering the response.

Description

root.txt

★ Required

The name of your response, such as **My response**.

CREATE

CANCEL

6. Fill out the **Create a new synthetic response** fields as follows:

- Leave the **Status** menu set at its default `200 OK`.
- In the **MIME Type** field, type `text/plain`.
- In the **Response** field, type at least one User-agent string and one Disallow string. For instance, the above example tells all user agents (via the `User-agent: *` string) they are not allowed to crawl anything after `/tmp/` directory or the `/foo.html` file (via the `Disallow: /tmp/*` and `Disallow: /foo.html` strings respectively).
- In the **Description** field, type an appropriate name. For example `robots.txt`.

7. Click the **Create** button. Your new response appears in the list of responses.

8. Click the **Attach a condition** link to the right of the newly created response. The Create a new condition window appears.

### Create a new condition

Learn the basics in our [conditions tutorial](#)

Type

Learn more about the [different types of conditions](#).

Name  \*

Apply if...

The expression to decide whether this is run.

▸ [Examples](#)

[Adjust priority \(default is 10\)](#)

This is an advanced option. For most configurations, the default is best.

**SAVE AND APPLY TO ROBOT.TXT** CANCEL

9. Fill out the **Create a new condition** fields as follows:

- From the **Type** menu, select the desired condition (for example, `Request`).
- In the **Name** field, type a meaningful name for your condition (e.g., `robots.txt`).
- In the **Apply if** field, type the logical expression to execute in VCL to determine if the condition resolves as true or false. In this case, the logical expression would be the location of your robots.txt file (e.g., `req.url ~ "^/robots.txt"`).
- Leave the **Priority** set to 10.

10. Click the **Save and apply to** button.

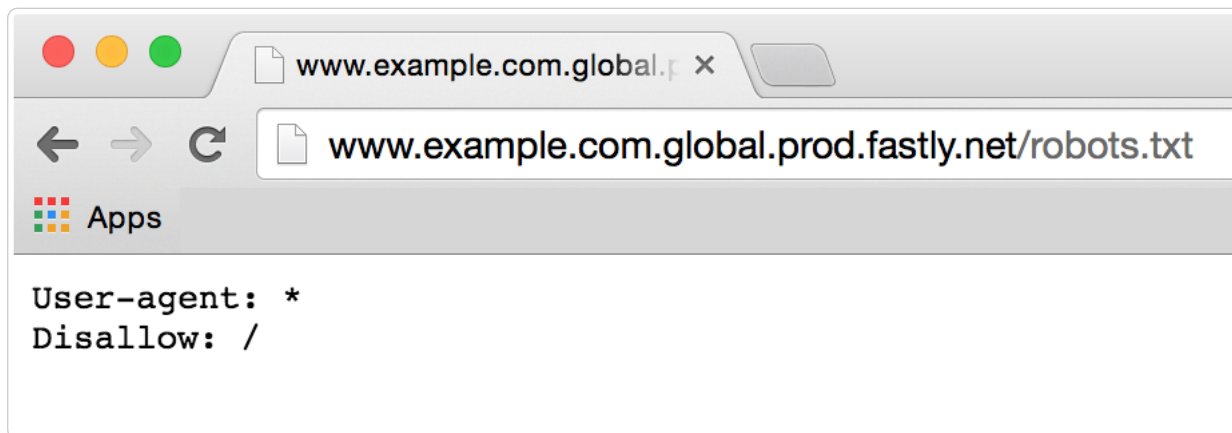
11. Click the **Activate** button to deploy your configuration changes.

**NOTE:** For an in-depth explanation of creating custom responses, check out our [Responses Tutorial \(/guides/basic-configuration/responses-tutorial\)](#).

## Why can't I customize my robots.txt file with global.prod.fastly.net?

Adding the `.global.prod.fastly.net` extension to your domain (for example, `www.example.com.global.prod.fastly.net`) via the browser or in a cURL command can be used to test how your production site will perform using Fastly's services.

To prevent Google from accidentally crawling this test URL, we provide an internal robots.txt file that instructs Google's web crawlers to ignore all pages for all hostnames that end in `.prod.fastly.net`.



This internal robots.txt file cannot be customized via the Fastly web interface until after you have set the CNAME DNS record for your domain to point to `global.prod.fastly.net`.

## § Creating error pages with custom responses (/guides/basic-configuration/creating-error-pages-with-custom-responses)

The default error responses served by Fastly can be jarring for your users, especially when using Fastly for consumer applications. This tutorial shows you how to set up your service configuration to serve a custom page or a synthetic response when we receive an error code from your backend.

★ **TIP:** Fastly can optionally serve stale content when there is a problem with your origin server. For more information, see our guide on [serving stale content \(/guides/performance-tuning/serving-stale-content\)](/guides/performance-tuning/serving-stale-content).

We assume you are already accustomed to editing and deploying configurations using the web interface (<https://manage.fastly.com/>). If you are not familiar with the web interface, see our basic setup (</guides/basic-setup/>) and configuration (</guides/basic-configuration/>) information to learn more before you continue.

### Overview

In this tutorial, we will create a response object that contains the HTML you want to serve for your error page. We provide example HTML, but you can use any HTML you see fit. The response object will require that you use a condition in order for it to be served. If you are not familiar with conditions or responses we suggest you read the following:

- Responses Tutorial (</guides/basic-configuration/responses-tutorial>)
- Conditions guides (</guides/conditions/>)

### Creating the custom response

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.
4. Click the **Content** tab. The Content page appears.

DOMAINS  
1 Domain

ORIGINS  
1 Host  
1 Health check

SETTINGS  
1 Request setting  
0 Cache settings

CONTENT  
0 Headers  
0 Gzips  
0 Responses

LOGGING  
1 Endpoint

CUSTOM VCL  
0 Files

## Headers

*There are no headers.*

CREATE YOUR FIRST HEADER

## Gzip

*Gzip is not enabled.*

SET UP GZIP

## Responses

*There are no responses.*

CREATE YOUR FIRST RESPONSE

### About content

Headers allow you to set, delete, and change request and response headers as your information is handled by Fastly.

Responses allow you to create synthetic responses that exist entirely on the cache servers.

5. Click the **Create response** button. The Create a new synthetic response page appears.

## Create a new synthetic response

More on how synthetic responses work in our [tutorial](#).

### CONDITION

This will happen all the time unless you [Attach a condition](#)

Status

404 Not Found

The HTTP status code to include in the header of the response.

MIME Type

text/html

The media type to put into the Content-Type header which informs the browser how to display the response. Common MIME types are [application/json](#), [text/plain](#), [text/html](#).

Response

```
<html>
<body>
<h1>Custom Handler - Error 404</h1>
</body>
</html>
```

The content to be served when delivering the response.

Description

Custom 404

★ Required

The name of your response, such as **My response**.

CREATE

CANCEL

6. Fill out the **Create a new synthetic response** fields as follows:

- From the **Status** menu, select the appropriate status (e.g., `404 Not Found`).
- In the **MIME Type** field, specify the Content-Type of the response (e.g., `text/html`).
- In the **Response** field, type the content to be served when delivering a response.
- In the **Description** field, type a name for the response you're creating (e.g., `Custom 404`).

7. Click the **Create** button. Your new response appears in the list of responses.

8. Click the **Attach a condition** link to the right of the name of your new response. The Create a new condition window appears.

Create a new condition

Learn the basics in our [conditions tutorial](#)

Type

Cache

Learn more about the [different types of conditions](#).

Name

404 Not found

\*

Apply if...

beresp.status == 404

The expression to decide whether this is run.

Examples

Adjust priority (default is 10)

This is an advanced option. For most configurations, the default is best.

SAVE AND APPLY TO CUSTOM 404

CANCEL

9. Fill out the **Create a new condition** fields as follows:

- From the **Type** menu, select the type of condition you want (e.g., `Cache`).
- In the **Name** field, type a name for the condition you're creating (e.g., `404 Not Found`).
- In the **Apply if** field, type the condition under which the new response occurs in the following format:

```
beresp.status == ###
```

where `###` equals the status condition you're creating the response for. For example, using the value of `beresp.status == 404` in the **Apply if** field here tells Fastly to use this response object whenever origin servers return a 404 status. (See the Conditions guides (</guides/conditions/>) for more detailed information on conditions.)

10. Click the **Save and apply to** button. The condition is created and applied to the custom response object you made earlier.

11. Click the **Activate** button to deploy your configuration changes. Fastly will now serve your custom HTML error page when required.

## § Creating multiple domains at one time (</guides/basic-configuration/creating-multiple-domains-at-one-time>)

It is possible to create multiple domains within in a single service programmatically, using Fastly's API (</api/>). Before you add domains, however, consider that it may be easier to serve multiple domains through custom VCL (</guides/vcl/mixing-and-matching-fastly-vcl-with-custom-vcl>). You'll be able to finely tune requests in this manner, yet remain flexible enough that you'll be able to implement as many processes as needed.

### Are there domain creation limits?

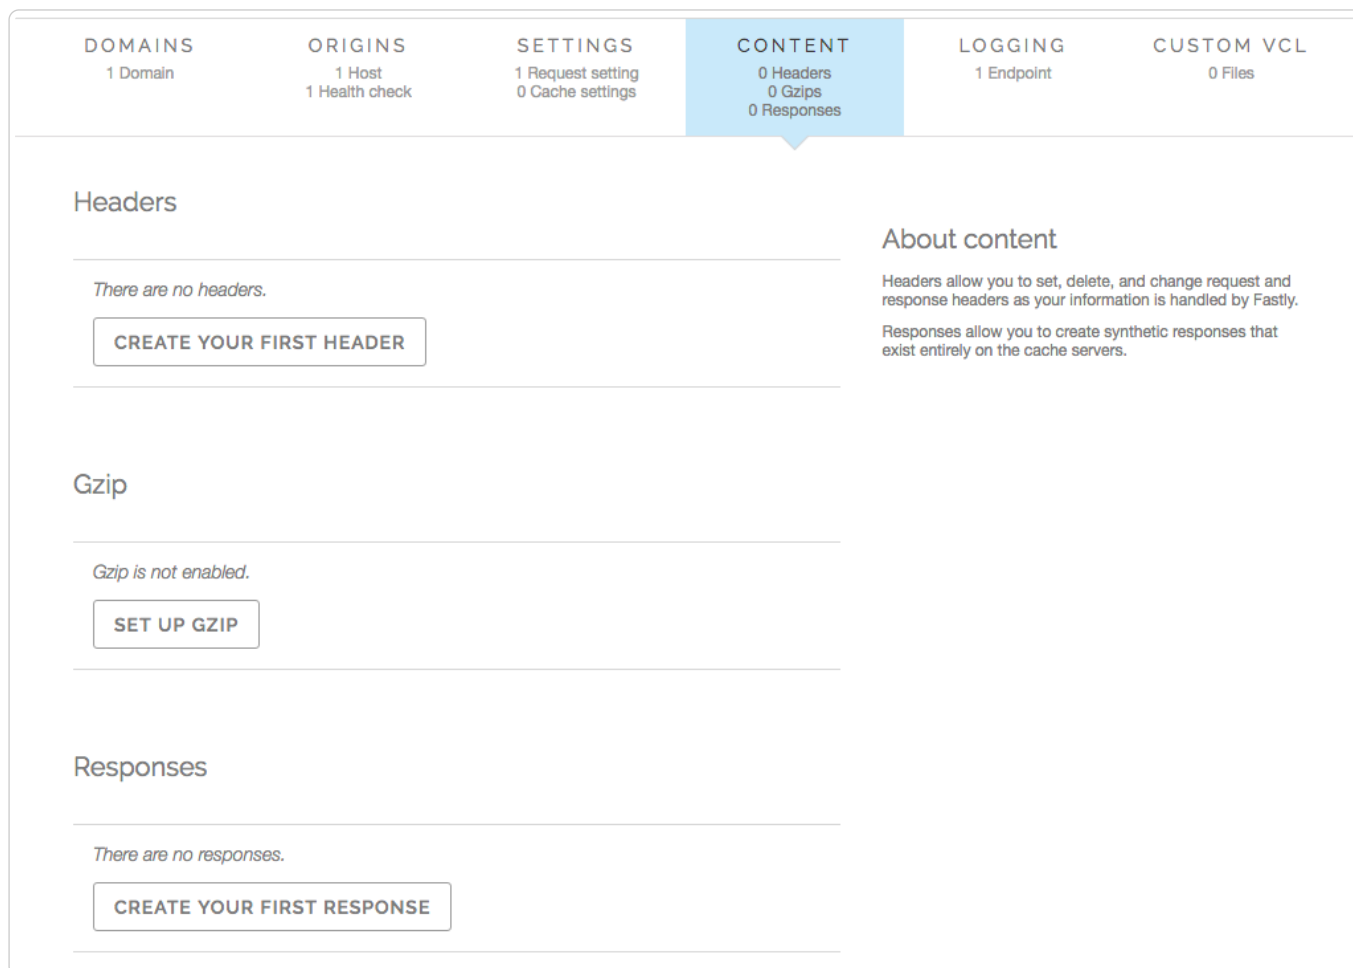
We set a limit (</guides/customer-support/common-service-and-domain-errors>) on the number of domains you can create per service by default. However, if you email [support@fastly.com](mailto:support@fastly.com) (<mailto:support@fastly.com>), we may be able to adjust this number for you by working with you to set up and fine-tune domain handling in your service.

## § Enabling automatic gzipping (</guides/basic-configuration/enabling-automatic-gzipping>)

To dynamically gzip content based on file extension or content-type, follow the steps below.

**⚠ WARNING:** Because dynamic gzipping fetches content from origin, compresses it, and then caches it, this feature doesn't work with our [ESI feature \(/guides/performance-tuning/esi-use\)](/guides/performance-tuning/esi-use). If you enable gzipping, Fastly will stop processing ESI language elements.

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.
4. Click the **Content** tab. The Content page appears.



5. Click the **Create gzip** button. The Create a new gzip page appears.



## Create a new gzip

More on how to enable gzip in our [gzip tutorial](#).

CONDITION

This will happen all the time unless you [Attach a condition](#)

Name

★ Required

The name of your gzip policy, such as **My Gzip Policy**.

By default we gzip css, js, html, eot, ico, otf, ttf, json, and svg.  
[Override these defaults](#)

CREATE

CANCEL

6. Accept the default gzip policy by clicking **Create** button. The Default Gzip policy appears.
7. Click the **Create** button. The new gzip policy appears.
8. Click the **Activate** button to deploy your configuration changes.

## Override the default gzip rule

If you want to override the default gzip rule, follow the steps below:

1. Click the **Create gzip** button. The Create a new gzip page appears.
2. Click the **Override these defaults** link on the page. The additional gzip fields appear.

## Create a new gzip

More on how to enable gzip in our [gzip tutorial](#).

CONDITION

This will happen all the time unless you [Attach a condition](#)

Name

★ Required

The name of your gzip policy, such as **My Gzip Policy**.

Extensions

Enter the file extensions for each file type you wish to have dynamically gzipped, separated by spaces.

Content Types

Enter the content-type for each type of content you wish to have dynamically gzipped, separated by spaces.

CREATE

CANCEL

3. Fill out the **Create a new gzip** fields as follows:

- In the **Name** field, type an arbitrary name for your new gzip rule.
- In the **Extensions** field, type the file extension for each file type to be dynamically gzipped, separated by spaces. Only type the three- or four-letter string representing the file extension. We recommend setting the Extensions to `css js html eot ico otf ttf json`.
- In the **Content Types** field, type the content-type for each type of content you wish to have dynamically gzipped, separated by spaces. Do not use regular expressions. We recommended setting the Content Types to the following:

```
text/html application/x-javascript text/css application/javascript text/javascript application/json application/vnd.ms-fontobject application/x-font-opentype application/x-font-truetype application/x-font-ttf application/xml font/eot font/opentype font/otf image/svg+xml image/vnd.microsoft.icon text/plain text/xml
```

4. Click the **Create** button. The new gzip policy appears.

5. Click the **Activate** button to deploy your configuration changes.

## Automatic exclusion from old browsers

Some web browsers should not receive gzipped content. Fastly's default gzip configuration intentionally excludes old browsers. Specifically, we run the following VCL before any custom VCL:

```
if (req.http.Accept-Encoding) {
  if (req.http.User-Agent ~ "MSIE 6") {
    unset req.http.Accept-Encoding;
  } elseif (req.http.Accept-Encoding ~ "gzip") {
    set req.http.Accept-Encoding = "gzip";
  } elseif (req.http.Accept-Encoding ~ "deflate") {
    set req.http.Accept-Encoding = "deflate";
  } else {
    unset req.http.Accept-Encoding;
  }
}
```

## § Health check frequency (/guides/basic-configuration/health-check-frequency)

Fastly performs health checks (/guides/basic-configuration/health-checks-tutorial) on your origin server based on the Check frequency setting you select in the Create a new health check page. The Check frequency setting you select specifies approximately how many requests per minute Fastly POPs (<https://www.fastly.com/network-map>) are checked to see if they pass. There is roughly one health check per Fastly POP per period. Any checks that pass will be reported as "healthy."

To determine origin health, you can configure health checks with the following frequency options:

- **Low** - Approximately 30 requests/minute, where "healthy" means 1 out of 2 must pass
- **Medium** - Approximately 120 requests/minute, where "healthy" means 3 out of 5 must pass
- **High** - Approximately 900 requests/minute, where "healthy" means 7 out of 10 must pass
- **Custom** - A custom frequency and request health pass levels

To set a custom frequency, select **Custom**. The following controls appear:

- **Threshold & Window** - The number of successes per total number health checks. For example, specifying `1/2` means 1 out of 2 checks must pass to be reported as healthy.
- **Initial** - The number of requests to assume as passing on deploy.
- **Interval & Timeout (ms)** - Interval represents the period of time for the requests to run. Timeout represents the wait time until request is considered failed. Both times are specified in milliseconds.

## § Health checks tutorial (/guides/basic-configuration/health-checks-tutorial)

In this tutorial we show you how to create health checks that periodically contacts your origin servers to make sure they're still working.

We assume that you're already accustomed to editing and deploying configurations using the web interface (<https://manage.fastly.com>). If you're not familiar with basic editing using the web interface, see our Help Guides (/guides/) to learn more before you continue.

# Overview

Health checks, while simple in principle, are more involved than most other configuration objects. Let's run through the options you have when creating a health check so you have a place to start when working through the tutorial.

- **Request** - An HTTP method and path to visit on your origin servers when performing the check. Use a unique path, for example, use `/website-healthcheck.txt`, not `/` or `/healthcheck`.
- **Host header** - The HTTP Host Header to set when making the request (e.g. `example.com`).
- **Description** - A human-readable identifier for the health check (e.g., `West Coast Origin Check`).
- **Expected response** - The HTTP status code the origin servers must respond with for the check to pass (usually `200 OK`).
- **Check frequency** - How often the origin server is checked with additional parameters that determine if the origin server is up or down. For more information about the additional parameters, see our guide on Health check frequency (</guides/basic-configuration/health-check-frequency>).




Fastly will periodically check your origin server based on the options chosen. Pay special attention to the HTTP host header. A common mistake is setting the wrong host. If the origin server does not receive a host it expects, it may issue a 301 or 302 redirect causing the health check to fail. Also, Varnish requires the origin server receiving the health check requests to close the connection for each request. If the origin server does not close the connection, health checks will time out and fail.

If an origin server is marked unhealthy due to health checks, Fastly will stop attempting to send requests to it. If all origin servers are marked unhealthy, Fastly will attempt to serve stale (</guides/performance-tuning/serving-stale-content>). If no stale object is available, a 503 will be returned to the client.

## Creating a health check

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.
4. Click the **Origins** tab. The Origins page appears.

[Return to service summary](#)


**www.eample.com**  Version 4   
Show service ID **in development** [ACTIVATE](#) [OPTIONS](#) 

DOMAINS 1 Domain	ORIGINS 1 Host 0 Health checks	SETTINGS 0 Request settings 0 Cache settings	CONTENT 0 Headers 0 Gzips 0 Responses	LOGGING 1 Endpoint	CUSTOM VCL 0 Files
---------------------	--------------------------------------	--	--	-----------------------	-----------------------

## Hosts

Hosts are used as backends for your site. In addition to the IP address and port, the information is used to uniquely identify a domain. For added security, we recommend [connecting to origins over TLS](#).

[+ CREATE HOST](#)

1.1.1.1 : 80 

Attach a condition TLS Options Delete

www.example.com

Show all details

## Health checks

Health checks [monitor the status of your hosts](#)—you can set Fastly to use a different origin, serve stale content, and more.

There are no health checks.

[CREATE YOUR FIRST HEALTH CHECK](#)

5. Click the **Create health check** button. The Create a new health check page appears.

## Create a new health check

Learn the basics in our [health checks tutorial](#)

Request

An HTTP verb (i.e., **HEAD**, **GET**, or **POST**) and path to visit on your origins when performing the check. Use a unique path. For example, use `/website-healthcheck.txt`, not `/` or `/healthcheck`.

Host header  ★ Required

The Host header to set when making the request (e.g. `example.com`).

▸ [Tips](#)

Description  ★ Required

A comment that describes your health check.

Expected response

The HTTP status code that signifies a healthy state.

Check frequency

Set how often this health check is performed. Checking frequently will result in more origin requests. [Learn more](#).

▸ [Details of check frequency options](#)

CREATE

CANCEL

6. Fill out the **Create a new health check** fields. For more information, review the field descriptions in the Overview section.

7. Click the **Create** button. Your new health check now appears in the list of checks.

## Assigning a health check

Health checks do nothing on their own, but they can be added as a special parameter to an origin server in your configuration.

1. Edit one of your existing origin servers by clicking the origin server's name. The Edit a host page appears.
2. From the **Health checks** menu, select the health check you just created.

## Edit a host

### CONDITION

This will happen all the time unless you [Attach a condition](#)

Address  :  ★ Required

The IP address (or hostname) and port number for your origin server. To [enable TLS](#), use port 443.

Shielding

The shield POP designated to reduce inbound load on your origins by serving cached data. Learn more about [POP request handling](#) and the [caveats of shielding](#).

Health check

[+ Create a new health check](#)

A small request assigned to [test the state of your origin](#).

Auto load balance

Enables Fastly to [load balance between two or more backend hosts](#).

[▶ Tips](#)

Description  ★ Required

The name of your origin, such as [My origin server](#).

3. Click **Update**.

Fastly will now use the health check to monitor the selected origin server.

## § How request settings are applied (/guides/basic-configuration/how-request-settings-are-applied)

Requests settings are applied based on the Action you select in the Create a new request setting page. You can choose any one of the following settings:

- **Do nothing now** - Apply the request setting options, but don't force a lookup or a pass action. The request settings are applied as the system continues through the VCL logic.
- **Lookup** - Immediately search the cache for content. If none is found (a MISS), then send the request to the origin.
- **Pass** - Immediately send the request to the origin each time and ignore additional request configurations. See our info on understanding the different PASS action behaviors ([/guides/vcl/understanding-the-different-pass-action-behaviors](#)) to learn more.

## § Manipulating the X-Forwarded-For header (/guides/basic-configuration/manipulating-the-x-forwarded-for-header)

You can control what happens to the X-Forwarded-For HTTP header via the Create a new request setting page on the Requests settings area of the Settings tab. From the X-Forwarded-For menu, select one of the following behaviors:

- **Append** - Appends the client IP to the X-Forwarded-For header.
- **Append All** - Appends the client IP (and edge-cache IP, in case of shielding) to the X-Forwarded-For header. Creates the header if it does not exist yet.
- **Clear** - Clears the X-Forwarded-For header.

- **Leave** - Leaves the X-Forwarded-For header as is, if it is present.
- **Overwrite** - Overwrites the X-Forwarded-For header with just the client IP.

For more information about requests and responses, see our tutorial (</guides/basic-configuration/adding-or-modifying-headers-on-http-requests-and-responses>).

---

## § Overriding caching defaults based on a backend response (</guides/basic-configuration/overriding-caching-defaults-based-on-backend-responses>)




---


In certain situations you may want to conditionally apply a different caching policy (</guides/performance-tuning/controlling-caching#conditionally-preventing-pages-from-caching>) based on a backend response. In this particular case we have backend that on occasion returns 404 errors (e.g., document not found). We don't want those responses to be cached for full caching period of a day but only for 5 minutes. To override default caching we add a cache object and then create conditions for it.

### Creating the new Cache Object

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.
4. Click the **Settings** tab. The Settings page appears.

[Return to service summary](#)


**www.example.com**   **Version 4**   
Show service ID **in development**

**ACTIVATE** **OPTIONS** 

<b>DOMAINS</b> 1 Domain	<b>ORIGINS</b> 1 Host 0 Health checks	<b>SETTINGS</b> 0 Request settings 0 Cache settings	<b>CONTENT</b> 0 Headers 0 Gzips 0 Responses	<b>LOGGING</b> 1 Endpoint	<b>CUSTOM VCL</b> 0 Files
----------------------------	---	---	---	------------------------------	------------------------------

### Default PCI

This setting allows you to use the **beresp.pci** VCL variable to automatically cache content.

**PCI enabled: No** 

### Override host

Override the **host header** being sent to your origin regardless of the host used in the initial request. Only required if the domain your origin is expecting is different than those that Fastly hosts.

[When should I use an override host?](#)

**SPECIFY AN OVERRIDE HOST**

### Request settings

Request Settings are used to customize Fastly's request handling. When used with **Conditions** the Request Settings allow you to fine tune how specific types of requests are handled.

*There are no request settings.*

**CREATE YOUR FIRST REQUEST SETTING**

### Cache settings


Cache Settings controls how **caching is performed** on Fastly. When used with **Conditions**, the Cache Settings provide you with fine grain control over how long content persists in the cache.

*There are no cache settings.*

**CREATE YOUR FIRST CACHE SETTING**

### Fallback TTL

This setting is used only when there is no **Time to Live (TTL)** set in an object's header.

**TTL (seconds): 3600** 

5. Click the **Create cache setting** button. The Create a new cache setting page appears.



## Create a new cache setting

This will override your cache control headers. In most cases, use with conditions applied.

### CONDITION

This will happen all the time unless you [Attach a condition](#)

TTL (seconds)

The TTL (Time To Live) entered will set the lifespan of the object within our cache nodes.

Action

This setting decides [how the request will be handled](#).

Stale TTL (seconds)

Stale TTL (Time To Live) is used by your application to [serve stale content](#). It determines how long to keep stale data after it has expired. Note there's custom VCL required to [complete this setup](#).

Description

★ Required

The name of your cache setting, such as **My cache setting**.

CREATE

CANCEL

6. Fill out the **Create a new cache setting** fields as follows:

- In the **TTL (seconds)** field, type the amount of time, in seconds, to cache the objects (e.g., ).
- From the **Action** menu, select **Deliver**.
- In the **Stale TTL (seconds)** field, type the amount of time to serve stale or expired responses, in seconds, should the backend become unavailable (e.g., ).
- In the **Description** field, type a descriptive name for the new cache settings.

7. Click the **Create** button.

## Creating an Override Condition for the new Cache Object

Once the object is created, add a condition to it.

1. Click the **Attach a condition** link to the right of the object.
2. Click **Create cache setting** button. The Create a new cache condition window appears.

Create a new cache condition.

Learn the basics in our [conditions tutorial](#)

Type

Cache

Learn more about the [different types of conditions](#).

Name

Override cache default

★ Required

Apply if...

beresp.status == 404 ||  
beresp.status == 403

The expression to decide whether this is run.

[Examples](#)

[Adjust priority \(default is 10\)](#)

This is an advanced option. For most configurations, the default is best.

SAVE AND APPLY TO MY CACHE SETTING

CANCEL

3. Fill out the **Create a new cache setting** fields as follows:

- From the **Type** menu, select the type of condition that is applied.
- In the **Name** field, type a descriptive name for the new condition.
- In the **Apply if** field, type an appropriate backend response header to specify when the condition will be applied. For example, `beresp.status == 404 || beresp.status == 403`.
- Leave the **Priority** field set at its default value.

4. Click the **Save and apply to** button.

5. Click the **Activate** button to deploy your configuration changes.

## Other notes

You can use any backend response header in the **Apply if** field to make decisions on caching.

For example, `beresp.http.Content-Type ~ "^text/html"` can be used to specify different caching rules for HTML documents.

## § Removing headers from backend response (/guides/basic-configuration/removing-headers-from-backend-response)

You can remove headers from any backend response. This may be necessary if your application automatically sets headers. For example, Drupal can set the following Expires and Cache-Control headers to prevent caching:

```
Expires: Sun, 19 Nov 1978 05:00:00 GMT
Last-Modified: Wed, 18 Jul 2012 18:52:16 +0000
Cache-Control: no-cache, must-revalidate, post-check=0, pre-check=0
```

To remove a header from the backend response, add a new header as follows:

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.
4. Click the **Content** tab The Content page appears.

DOMAINS  
1 Domain

ORIGINS  
1 Host  
1 Health check

SETTINGS  
1 Request setting  
0 Cache settings

CONTENT  
0 Headers  
0 Gzips  
0 Responses

LOGGING  
1 Endpoint

CUSTOM VCL  
0 Files

Headers

There are no headers.

CREATE YOUR FIRST HEADER

Gzip

Gzip is not enabled.

SET UP GZIP

Responses

There are no responses.

CREATE YOUR FIRST RESPONSE

About content

Headers allow you to set, delete, and change request and response headers as your information is handled by Fastly.

Responses allow you to create synthetic responses that exist entirely on the cache servers.

5. Click the **Create header** button. The Create a new header window appears.

## Create a new header

Learn more about this section in our [headers tutorial](#).

CONDITION

This will happen all the time unless you [Attach a condition](#)

Type / Action

Cache

Delete

The type of header and the action performed on it.

Destination

http.Expires

★ Required

The name of the header that will be affected by the selected action. For example: **http.Content-Type**, **http.Set-Cookie**, **http.Via**, **http.Location**, or **http.Access-Control-Allow-Origin**.

Ignore if set

No

If switched to **Yes**, the action will not be performed if the header in Destination exists.

Priority

10

★ Required

The order in which the header rules execute within the condition.

Description

Remove Expire Headers

★ Required

The name of your header, such as **My header**.

CREATE

CANCEL

6. Fill out the **Create a new header** fields as follows:

- From the **Type** menu, select **Cache**, and from the **Action** menu, select **Delete**
- In the **Destination** field, type the name of the header (e.g., `http.Expires`).
- From the **Ignore if set** menu, select **No**.
- In the **Priority** field, type `10`.
- In the **Description** field, type a descriptive name for the header rule (e.g., `Remove Expire Headers`).

7. Click the **Create** button.

8. Click the **Activate** button to deploy your configuration changes.

★ **TIP:** You may also be interested in our information on [setting content type based on file extension \(/guides/basic-configuration/setting-content-type-based-on-file-extension\)](/guides/basic-configuration/setting-content-type-based-on-file-extension).

## § Responses tutorial (/guides/basic-configuration/responses-tutorial)

Fastly allows you to create custom HTTP responses that are served directly from the cache without storing the page on a server. Responses are commonly used to serve small static assets that seldom change and maintenance pages that are served when origins are unavailable. This tutorial shows you how to create your own responses.

❗ **NOTE:** We assume that you already know how to edit and deploy configurations using the [web interface \(https://manage.fastly.com/\)](https://manage.fastly.com/). If you are not familiar with basic editing using the application, see [our help guides \(/guides/\)](/guides/) to learn more.

# Overview

A response has three basic attributes:

- Status - An HTTP status code to include in the header of the response
- Response - The content to be served when delivering the response
- Description - A human readable identifier for the response

By setting these three attributes and adding a condition to the response, you can very quickly get one up and running on your service. Let's get started!

## Creating a response

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.
4. Click the **Content** tab. The Content page appears.

The screenshot shows the Fastly Content configuration page. The top navigation bar includes tabs for DOMAINS, ORIGINS, SETTINGS, CONTENT, LOGGING, and CUSTOM VCL. The CONTENT tab is selected, showing 0 Headers, 0 Gzips, and 0 Responses. The main content area is divided into three sections: Headers, Gzip, and Responses. Each section has a message indicating no items are present and a button to create the first one. To the right of the Headers section is an 'About content' sidebar explaining headers and responses.

DOMAINS	ORIGINS	SETTINGS	CONTENT	LOGGING	CUSTOM VCL
1 Domain	1 Host 1 Health check	1 Request setting 0 Cache settings	0 Headers 0 Gzips 0 Responses	1 Endpoint	0 Files

### Headers

There are no headers.

[CREATE YOUR FIRST HEADER](#)

### Gzip

Gzip is not enabled.

[SET UP GZIP](#)

### Responses

There are no responses.

[CREATE YOUR FIRST RESPONSE](#)

### About content

Headers allow you to set, delete, and change request and response headers as your information is handled by Fastly.

Responses allow you to create synthetic responses that exist entirely on the cache servers.

5. Click the **Create response** button. The Create a new synthetic response page appears.

## Create a new synthetic response

More on how synthetic responses work in our [tutorial](#).

### CONDITION

This will happen all the time unless you [Attach a condition](#)

Status

200 OK

The HTTP status code to include in the header of the response.

MIME Type

text/html

The media type to put into the Content-Type header which informs the browser how to display the response. Common MIME types are [application/json](#), [text/plain](#), [text/html](#).

Response

```
<html>
<head><title>Under Construction</title></head>
<body>
<h1>This page is under construction</h1>
</body>
</html>
```

The content to be served when delivering the response.

Description

My first response

★ Required

The name of your response, such as **My response**.

CREATE

CANCEL

6. Fill out the **Create a new synthetic response** fields as follows:

- From the **Status** menu, select the appropriate status (e.g., `200 OK`).
- In the **MIME Type** field, type the content type of the response (e.g., `text/html`).
- In the **Response** field, type the response you want to appear when the conditions are met.
- In the **Description** field, type a human-readable name for the response (e.g., `My First Response`).

7. Click the **Create** button to create your custom response.

Your new response appears in the list of responses.

## Adding conditions

1. Click the **Attach condition** link to the right of the new response. The Create a new condition window appears.

Create a new condition.

Learn the basics in our [conditions tutorial](#)

Type

Request

Learn more about the [different types of conditions](#).

Name

Construction

\*

Apply if...

req.url ~ "^/construction/"

The expression to decide whether this is run.

[Examples](#)

[Adjust priority \(default is 10\)](#)

This is an advanced option. For most configurations, the default is best.

SAVE AND APPLY TO MY FIRST RESPONSE

CANCEL

2. Fill out the **Create a new condition** fields as follows:

- From the **Type** menu, select the type of condition you want to create.
- In the **Name** field, type a human-readable name for the condition so that it can be easily identified in the future.
- In the **Apply if** field, type the condition under which the new response occurs. The condition should take the following format: `|req.url ~ "^/construction/"` equals the request condition you're creating the response for. The Conditions subcategory (</guides/conditions/>) has more detailed information on conditions.
- In the **Priority** field, type a priority if needed. Condition priorities are only needed in "interesting" cases, and can usually be left at the default "10" for all response conditions.

3. Click the **Save and apply to** button.

4. Click the **Activate** button to deploy your configuration changes.

Fastly now serves your custom response page when the condition is met.

## § Setting Content Type based on file extension (/guides/basic-configuration/setting-content-type-based-on-file-extension)

In some situations you may want to override the content type that a backend returns. To do that you will need to create a new header object and an associated condition.

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.
4. Click the **Content** tab. The Content page appears.

DOMAINS  
1 Domain

ORIGINS  
1 Host  
1 Health check

SETTINGS  
1 Request setting  
0 Cache settings

CONTENT  
0 Headers  
0 Gzips  
0 Responses

LOGGING  
1 Endpoint

CUSTOM VCL  
0 Files

Headers

There are no headers.

CREATE YOUR FIRST HEADER

Gzip

Gzip is not enabled.

SET UP GZIP

Responses

There are no responses.

CREATE YOUR FIRST RESPONSE

About content

Headers allow you to set, delete, and change request and response headers as your information is handled by Fastly.

Responses allow you to create synthetic responses that exist entirely on the cache servers.

5. Click the **Create header** button. The Create a new header page appears.



## Create a new header

Learn more about this section in our [headers tutorial](#).

### CONDITION

This will happen all the time unless you [Attach a condition](#)

Type / Action

Cache

Set

The type of header and the action performed on it.

Destination

http.Content-Type

★ Required

The name of the header that will be affected by the selected action. For example: **http.Content-Type**, **http.Set-Cookie**, **http.Via**, **http.Location**, or **http.Access-Control-Allow-Origin**.

Source

"application/javascript; charset=utf-8"

★ Required

New content for the header. Can be a static value (e.g. string or number) or a dynamic value (e.g. existing header or a GeoIP value). Please use quotes for string values.

Ignore if set

No

If switched to **Yes**, the action will not be performed if the header in Destination exists.

Priority

10

★ Required

The order in which the header rules execute within the condition.

Description

Add Content Type

★ Required

The name of your header, such as **My header**.

CREATE

CANCEL

6. Fill out the **Create a new header** fields as follows:

- From the **Type** menu, select **Cache**, and from the **Action** menu, select **Set**.
- In the **Destination** field, type `http.Content-Type`.
- In the **Source** field, type the content type you want to match, such as `"application/javascript; charset=utf-8"`.
- From the **Ignore if set** menu, select **No**.
- In the **Priority** field, type `10`.
- In the **Description** field, type an appropriate name (e.g., `Add Content Type`).

7. Click the **Create** button.

Once you have created the header object, apply a condition (</guides/conditions/>). Otherwise, that particular object is applied to all requests.

1. Click the **Attach a condition** link to the right of the new header name. The Create a new cache condition window appears.

Create a new cache condition.

Learn the basics in our [conditions tutorial](#)

Type

Cache

Learn more about the [different types of conditions](#).

Name

Files ending with .js

\*

Apply if...

req.url.ext == ".js"

The expression to decide whether this is run.

[Examples](#)

[Adjust priority \(default is 10\)](#)

This is an advanced option. For most configurations, the default is best.

SAVE AND APPLY TO THIS ADD CONTENT TYPE

CANCEL

2. Fill out the **Create a new cache condition** fields as follows:

- Leave the **Type** menu set to **Cache**.
- In the **Name** field, type a descriptive name, such as `Files ending with .js`.
- In the **Apply if** field, type the condition that matches your request, such as `req.url.ext == ".js"` (to match the request for files ending in .js).

3. Click the **Save and apply to** button to create the new condition.

4. Click the **Activate** button to deploy your configuration changes.

★ **TIP:** You may also be interested in our guide to [Removing headers from backend response \(/guides/basic-configuration/removing-headers-from-backend-response\)](/guides/basic-configuration/removing-headers-from-backend-response).

## § Specifying an override host (/guides/basic-configuration/specifying-an-override-host)

If you want to rewrite the host header being sent to your origin regardless of the host used in the initial request, specify an override host. Use this if you have multiple domains tied to a service and want them all served by the same origin, or if the domain your origin is expecting is different than one specified in your Fastly service. You most likely won't need to use this feature.

You can override the host header being sent to your origin by specifying the domain name (/guides/about-fastly-services/domain-names-and-fastly-services) of your override host on the **Settings** page for a specific service.




Here are some examples of when to use an override host:


- When using backends such as Amazon S3 (/guides/integrations/amazon-s3), Google Cloud Storage (/guides/integrations/google-cloud-storage#setting-the-default-host-for-your-service-to-your-gcs-bucket), or Heroku (<https://devcenter.heroku.com/articles/fastly>), you want to ensure you use the proper host header so these providers know how to route requests directly to your content. Each provider uses the host header to associate requests with your account's storage location. For example, if you set up your origin using Amazon S3, you send the name of your S3 bucket as your host header. Amazon is set up so that it only accepts host headers that have the same name as the bucket hosting your content. A request to `your-domain.com` must be re-written to `<your-bucket>.s3.amazonaws.com`, or else the request is denied.
- You have a service that contains three sites: `www.abc.com`, `www.myexample.com`, and `www.mysite.com` and you have one origin. You can have the same origin respond to each domain by overriding the host header to one accepted by your origin, for example, `origin.example.com`. The result will be that a request to `www.abc.com`, `www.myexample.com`, or `www.mysite.com` returns content from `origin.example.com`.

## Override a host

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.
4. Click the **Settings** tab. The Settings page appears.

[Return to service summary](#)


**www.example.com**   **Version 4**   
Show service ID **in development**

**ACTIVATE** **OPTIONS** 

<b>DOMAINS</b> 1 Domain	<b>ORIGINS</b> 1 Host 0 Health checks	<b>SETTINGS</b> 0 Request settings 0 Cache settings	<b>CONTENT</b> 0 Headers 0 Gzips 0 Responses	<b>LOGGING</b> 1 Endpoint	<b>CUSTOM VCL</b> 0 Files
----------------------------	---	---	---	------------------------------	------------------------------

### Default PCI

This setting allows you to use the **beresp.pci** VCL variable to automatically cache content.

**PCI enabled: No** 

### Override host

Override the **host header** being sent to your origin regardless of the host used in the initial request. Only required if the domain your origin is expecting is different than those that Fastly hosts.

[When should I use an override host?](#)

**SPECIFY AN OVERRIDE HOST**

### Request settings

Request Settings are used to customize Fastly's request handling. When used with **Conditions** the Request Settings allow you to fine tune how specific types of requests are handled.

*There are no request settings.*

**CREATE YOUR FIRST REQUEST SETTING**

### Cache settings


Cache Settings controls how **caching is performed** on Fastly. When used with **Conditions**, the Cache Settings provide you with fine grain control over how long content persists in the cache.

*There are no cache settings.*

**CREATE YOUR FIRST CACHE SETTING**


### Fallback TTL

This setting is used only when there is no **Time to Live (TTL)** set in an object's header.

**TTL (seconds): 3600** 

5. Click the **Specify an Override Host** button on the lower right side of the page.

6. Type the desired host name. For example, `origin.example.com`.



7. Click the **Save** button. If you are using Amazon S3, set your default host to `<yourbucket>.s3.amazonaws.com`. If you are using Google Cloud Storage, set your default host to `<your bucket name>.storage.googleapis.com`. The override host header appears in the Override host section.
8. Click the **Activate** button to deploy your configuration changes.

## Caveats about using the override host

There are situations when you may not want to use an override host:

- **Using multiple origins.** When you specify a host override, you're specifying what host name is actually sent to your origin. If you have a service with two different origins and each origin requires a different host name, specifying a host override for all requests results in one origin not returning valid responses. If you specify a default host name that matches only one of the origins, then no content is returned from the other origin requests.

**NOTE:** If you want to serve content from multiple backends, you should conditionally route to them. Refer to [Routing assets to different origins \(/guides/performance-tuning/routing-assets-to-different-origins\)](/guides/performance-tuning/routing-assets-to-different-origins) for more information.

- **Shielding is enabled.** If you enable a host override along with shielding and the specified override host doesn't match to a domain within the service, the shield won't route the request properly and an error of 500 is expected. Refer to [Shielding \(/guides/performance-tuning/shielding#performance-caveats\)](/guides/performance-tuning/shielding#performance-caveats) for more information.

## § Testing setup before changing domains (/guides/basic-configuration/testing-setup-before-changing-domains)

After you deploy your service, but before you change your DNS entries to send your domain to our servers, you can check to see how your service is pulled through our network. Testing your domain can help you identify DNS issues or problems with your Fastly configuration.

To test your domain on Fastly before you make a final CNAME change you would:

- find out the IP address of a Fastly pop
- add a domain host entry to your hosts file
- test the domain in browser

## Determine the IP address of a Fastly POP

Run nslookup or a dig command to determine the IP address of a Fastly POP. For example, nslookup on `global.prod.fastly.net` would return:

```
$ nslookup global.prod.fastly.net
Server:      10.100.8.11
Address:     10.100.8.11#53

Non-authoritative answer:
global.prod.fastly.net canonical name = global-ssl.fastly.net.
global-ssl.fastly.net canonical name = fallback.global-ssl.fastly.net.
Name:   fallback.global-ssl.fastly.net
Address: 199.27.79.249
Name:   fallback.global-ssl.fastly.net
Address: 23.235.47.249
```

Find one of the final IP addresses returned at the bottom of the nslookup response. In this example, the ones you're looking for are `199.27.79.249` and `23.235.47.249`.

Alternatively, running a dig command for `global.prod.fastly.net` would return this:

```
$ dig global.prod.fastly.net

; <<>> DiG 9.10.3-P3 <<>> global.prod.fastly.net
;; global options: +cmd
;; Got answer:
;; ->HEADER<- opcode: QUERY, status: NOERROR, id: 658
;; flags: qr rd ra; QUERY: 1, ANSWER: 4, AUTHORITY: 4, ADDITIONAL: 5

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags:; udp: 4096
;; QUESTION SECTION:
;global.prod.fastly.net.      IN  A

;; ANSWER SECTION:
global.prod.fastly.net.  9  IN  CNAME  global-ssl.fastly.net.
global-ssl.fastly.net.  9  IN  CNAME  fallback.global-ssl.fastly.net.
fallback.global-ssl.fastly.net. 9 IN  A      199.27.79.249
fallback.global-ssl.fastly.net. 9 IN  A      23.235.47.249

;; AUTHORITY SECTION:
fastly.net.             12462 IN  NS      ns3.p04.dynect.net.
fastly.net.             12462 IN  NS      ns4.p04.dynect.net.
fastly.net.             12462 IN  NS      ns2.p04.dynect.net.
fastly.net.             12462 IN  NS      ns1.p04.dynect.net.

;; ADDITIONAL SECTION:
ns1.p04.dynect.net. 12415 IN  A      208.78.70.4
ns2.p04.dynect.net. 12415 IN  A      204.13.250.4
ns3.p04.dynect.net. 12415 IN  A      208.78.71.4
ns4.p04.dynect.net. 12415 IN  A      204.13.251.4
```

The final IP addresses (A records) are the last two items in the `ANSWER SECTION` of the dig results: `199.27.79.249` and `23.235.47.249`.

## Modify your hosts file

Temporarily add a static IP and domain host entry to your hosts file on your computer. For example if the domain you are testing is

`www.example.com` and one of the IP addresses returned by nslookup or a dig command is `199.27.79.249` you would add this entry to the file:

```
199.27.79.249 www.example.com
```

and save the changes.

★ **TIP:** On machines running Mac OSX or Linux, your hosts file is `/etc/hosts`. On Windows-based machines, it's

```
C:\Windows\System32\Drivers\etc\hosts.
```

## Test your domain

Test your domain to see how Fastly pulls it through our network by restarting your browser if it's already running, and then typing your domain in the address field. You should now see the updated domain in the address field indicating requests are being sent to the Fastly POP.

Alternatively, you can test the domain using a ping command to verify that your domain is being served by a Fastly POP address. In this case, `ping www.example.com` would display the Fastly POP address `199.27.79.249`.

Be sure to remove the host entry from your hosts file after you make CNAME changes to point your domain to Fastly.

## § Using Fastly with apex domains (/guides/basic-configuration/using-fastly-with-apex-domains)

Some customers use only their second level or apex domain (e.g., `example.com` rather than `www.example.com`) as their canonical domain. Due to limitations in the DNS specification, we don't recommend placing a CNAME record at the apex domain or using the CNAME Flattening (e.g., ALIAS or ANAME) features offered by some DNS providers. Instead, we offer Anycast IP addresses for content that must be hosted on a second-level or apex domain.

## Where problems exist and why

The DNS instructions in RFC1034 (<https://www.ietf.org/rfc/rfc1034.txt>) (section 3.6.2) state that, if a CNAME record is present at a node, no other data should be present. This ensures the data for a canonical name and its aliases cannot be different. Because an apex domain requires NS records and usually other records like MX to make it work, setting a CNAME at the apex would break the "no other data should be present" rule.

In general, the problem with apex domains happens when they fail to redirect to their `www` equivalents (`example.com` points nowhere instead of pointing to `www.example.com`). A couple of workaround options exist:

- Only use Fastly for API or AJAX calls, images, and other static assets (e.g., serve `example.com` yourself and CNAME to Fastly for assets at `assets.example.com`)
- Redirect from the apex domain to the version proxied by Fastly (e.g., redirect any requests for `example.com` to `www.example.com`)

Neither of these workarounds, however, are ideal.

## Anycast option

Fastly can provide Anycast (<https://en.wikipedia.org/wiki/Anycast>) IP addresses for content that must be hosted on a second-level or apex domain. We do not charge extra for this service, however, you must be using one of Fastly's paid plans (</guides/account-types-and-billing/accounts-and-pricing-plans>) (with or without a contract) in order to take advantage of it.

Our Anycast option allows you to add A records that point your apex domain at Fastly. If the Fastly Anycast IP addresses change we will notify you at [status.fastly.com](https://status.fastly.com/) (<https://status.fastly.com/>). Because Anycast doesn't give us as much flexibility in routing your requests, this option may not be as performant as our CNAME-based system (</guides/basic-setup/adding-cname-records>). We recommend you use our CNAME-based system for as much content as possible, particularly for large files or streaming video.

For questions about Anycast and second-level or apex domains, contact [support@fastly.com](mailto:support@fastly.com) (<mailto:support@fastly.com>).

• [Guides \(/guides/\)](/guides/) > [Basic setup \(/guides/basic\)](/guides/basic) > [About Fastly services \(/guides/about-fastly-services/\)](/guides/about-fastly-services/)

## § About Fastly's Application Programming Interface (API) (</guides/about-fastly-services/about-fastlys-application-programming-interface>)

Fastly provides an application programming interface (API) that can be accessed via a number of popular interactive clients. The Fastly API allows you to manage Fastly services via remote procedure calls instead of the web interface (</guides/about-fastly-services/about-the-web-interface-controls>). This currently includes features such as:

- Authentication (</api/auth>)
- Configuration (</api/config>)
- Historical Stats (</api/stats>)
- Purging (</api/purge>)
- Remote Logging (</api/logging>)

The API features do not include customer account setup, which can only occur through the web interface controls. For examples of each API call in action, including full descriptions of the fields used and examples of requests and responses, see Fastly's API Reference (</api/>).

## Available API clients

The API's main entry point is <https://api.fastly.com> (<https://api.fastly.com>). It can be accessed via the following interactive clients:

- a Perl module
- a Ruby gem
- two different Python libraries
- a Node.js client
- a Scala client

Fastly's API Client web page (</api/clients>) contains links to GitHub repositories where these clients can be found. When third-party organizations have supplied these clients, we've noted so on the web page. We also have several integrations (</guides/integrations/>) available.

Fastly makes no warranty on third-party software. We assume no responsibility for errors or omissions in the third-party software or documentation available. Using such software is done entirely at your own discretion and risk.

## Authentication via the API

Nearly all API calls require requests to be authenticated. The Fastly API provides three methods for authenticating API calls: API token, API key, and username-password. Most API calls can be authenticated using an API token or API key alone. Those calls that don't use an API token or API key alone (usually at the account level) typically require using an authenticated session cookie instead. See the API authentication web page (</api/auth>) for additional information.

Authentication Method	Description
API Token	Allows you to create unique authentication identifiers for the users and applications authorized to interact with your service. You can scope the authorization of tokens to a single service, and you can restrict access to purge all ( <a href="/guides/purging/single-purges#purging-all-content">/guides/purging/single-purges#purging-all-content</a> ).
API Key	Allows the API key ( <a href="/guides/account-management-and-security/finding-and-managing-your-account-info#finding-and-regenerating-your-api-key">/guides/account-management-and-security/finding-and-managing-your-account-info#finding-and-regenerating-your-api-key</a> ) located on a customer's account page to be included as a Fastly-Key header. The master API key is scoped to the entire account.
Session	Allows you to create a cookie that can be stashed and used in subsequent requests. You can create the cookie with a POST HTTP command issued with the user and password parameters.

## § About Fastly's Cached Content Control features (</guides/about-fastly-services/about-fastlys-cached-content-control-features>)

Fastly has no set hard limit on how long objects will remain cached (</guides/performance-tuning/controlling-caching#how-long-fastly-caches-content>). Instead, Fastly supports customer-configurable Time to Live (TTL) settings and customer controlled content purging.

### Time to Live support

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 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.

### Instant Purge support

Fastly supports an Instant Purge feature that allows you to actively invalidate content (</guides/purging/>). 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 (</api/purge>). Fastly currently supports Ruby, Python, PHP, and Perl libraries (</api/clients>) 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, customers can eliminate cache-to-origin HTTP traffic that all other CDN services generate to determine if expired objects are still valid.

## § About Fastly's Cloud Accelerator (</guides/about-fastly-services/about-fastlys-cloud-accelerator>)



Fastly's Cloud Accelerator (<https://www.fastly.com/blog/announcing-fastlys-cloud-accelerator-collaboration-with-google-cloud-platform>) provides an integration between Fastly and Google services. Specifically, the integration enables you to connect to Google's Cloud Platform (<https://cloud.google.com>) service via peered network interconnections (direct PNIs) directly to Fastly's content delivery network services, thus speeding up your content delivery and optimizing backend workload.

When you sign up for Fastly services (</guides/basic-setup/sign-up-and-create-your-first-service>) and configure a Google Cloud Platform service as your origin server, you designate a specific point of presence (POP) to serve as an Origin Shield that handles cached content (</guides/about-fastly-services/how-fastlys-cdn-service-works>) from their servers.

Requests from Fastly POPs (</guides/about-fastly-services/fastly-pop-locations>) to the Cloud Accelerator Origin Shields are routed over Fastly's network, which leverages optimized TCP connection handling, quick-start and opened connections to ensure fast response times between POPs and through to the end-user. Fastly ensures that requests go directly to the Origin Shield instead of the origin servers. Only requests that the entire network has never handled will go back to the Google Cloud Platform service.

## § About Fastly's Delivery Acceleration features (</guides/about-fastly-services/about-fastlys-delivery-acceleration-features>)

### 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 (</guides/performance-tuning/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 a customer's origin server. Request collapsing is enabled by default.

### Grace mode

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 instructs Fastly's CDN Service to accelerate content delivery to end users by continuing to serve stale or expired (but likely still valid) content for a set amount of time. This allows customers 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. Customers must specifically configure their service to serve stale content (</guides/performance-tuning/serving-stale-content>).

### HTTP request fulfillment

The Fastly CDN Service responds to HTTP GET requests initiated from end users' using a customer's 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 a customer's 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 customers 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 customers to specify how to forward and transform specific header types.

## § About Fastly's On-the-Fly Packaging service (</guides/about-fastly-services/about-fastlys-onthefly-packaging-service>)

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

**ⓘ IMPORTANT:** Fastly's On-the-Fly Packaging (OTFP) service is an add-on service. Our Professional Services team will assist with configuration and testing. To enable OTFP and begin this process, contact your account manager or email [sales@fastly.com](mailto:sales@fastly.com) (<mailto:sales@fastly.com>) for more details.

### Supported On-the-Fly Packaging features

Fastly's OTFP service supports the following specific features:


- **HDS, HLS, and MPEG-DASH packaging.** Fastly provides support for version 1 of the Adobe HTTPS 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 cherry picked features from later versions to support additional customer use cases, like subtitle support and trick play for HLS.
- **Standard codecs.** Fastly supports H.264 and MPEG-4 Part 10 Advanced Video Coding (AVC). Fastly also supports Advanced Audio Coding (AAC), including AAC-LC (low complexity) and HE-AAC (high efficiency), MPEG-1 or Audio Layer III (MP3), and the Dolby AC3 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.
- **HLS multi-language subtitles and closed captions.** Fastly provides support for both in-band (EIA-608 (<https://en.wikipedia.org/wiki/EIA-608>) and CEA-708 (<https://en.wikipedia.org/wiki/CEA-708>)) and out-of-band (Web Video Text Tracks or WebVTT (<https://w3c.github.io/webvtt/>)) subtitle and closed caption delivery.
- **HLS timed metadata injection.** Fastly supports HLS time-based metadata ([https://developer.apple.com/library/ios/documentation/AudioVideo/Conceptual/HTTP\\_Live\\_Streaming\\_Metadata\\_Spec/Introduction/Introduction.html](https://developer.apple.com/library/ios/documentation/AudioVideo/Conceptual/HTTP_Live_Streaming_Metadata_Spec/Introduction/Introduction.html)) which allows you embed custom metadata or ad markers about a stream into video segments at specified time instances in ID3v2 format.
- **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 ([https://developer.apple.com/library/tvos/documentation/General/Reference/HLSAuthoringSpec/Requirements.html#//apple\\_ref/doc/uid/TP400165CH2-SW1](https://developer.apple.com/library/tvos/documentation/General/Reference/HLSAuthoringSpec/Requirements.html#//apple_ref/doc/uid/TP400165CH2-SW1)) requires this feature for distributing video on the Apple TV.
- **Media encryption.** Fastly can encrypt videos packaged into HLS (supports both Envelope/AES-128 and SAMPLE-AES ([https://developer.apple.com/library/ios/documentation/AudioVideo/Conceptual/HLS\\_Sample\\_Encryption/Intro/Intro.html](https://developer.apple.com/library/ios/documentation/AudioVideo/Conceptual/HLS_Sample_Encryption/Intro/Intro.html)) methods) and MPEG-DASH (ISO/IEC 23001-7: 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.
- **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.

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

- Token-based authentication (</guides/tutorials/enabling-token-authentication/>) for increased response time by placing validation at the edge
- Geo-IP location (</guides/vcl/geoip-related-vcl-features>) and device detection (</guides/vcl/delivering-different-content-to-different-devices>) for content targeting
- Edge dictionaries (</guides/edge-dictionaries/>) for real-time business rules and decision making at the edge
- Remote log streaming (</guides/streaming-logs/>) for data aggregation and viewer diagnostics
- Transport Layer Security (TLS) (</guides/securing-communications/>) for secure communications delivery

## How Fastly's On-the-Fly Packaging service works

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



When users request manifests or video segments, those requests initially come to Fastly caches instead of going to your origin storage. Fastly's edge caches deliver those objects if they are available and valid. If the objects don't already exist in the edge caches, the requests will be passed on to a designated shield cache (</guides/performance-tuning/shielding>) 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 (</guides/on-the-fly-packaging/adaptive-bitrate-playback-url-guidelines>), 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.

---

## § About Fastly's Origin Shielding features (</guides/about-fastly-services/about-fastlys-origin-shielding-features>)

---

When configuring Fastly's CDN Service during the self-provisioning process (</guides/about-fastly-services/self-provisioned-fastly-services>), customers can choose an "origin shield" as a specific point of presence (POP) designated to host cached content from their servers. This server is referred to as a "shield" because it protects a customer's application servers from continuous requests for content.

### Shield POPs

Customers can designate a specific POP to serve as a "shield" (</guides/performance-tuning/shielding>) for their origin servers. If Fastly's caches do not have the content being requested, they fetch it from the shield server instead of the customer's origin servers. Fastly caches fetch content from a customer's origin server only when the shield server does not have the content being requested.

### Load balancing

Customers can designate multiple servers as their origin servers. When two or more application servers are provisioned as origin servers, Fastly's CDN Service will distribute fetch requests for content across those application servers using the round-robin method of distribution. This type of load balancing (</guides/performance-tuning/load-balancing-configuration>) is enabled by default; customers must explicitly disable it.

### Health checks

Customers have the option to configure Fastly's CDN Service to perform health checks (</guides/basic-configuration/health-checks-tutorial>) on their application servers and measure their responsiveness. Health checks are not enabled by default; the customer must specifically enable them. Customers can use health check responsiveness measurements to fine-tune the distribution of fetch requests. Request collapsing is enabled by default.

### Request collapsing

Cached content sometimes must be refreshed when that content becomes "stale" or expires. When multiple end users request cached content that is in the process of being refreshed from origin servers, request collapsing (</guides/performance-tuning/request-collapsing>) groups those requests to be satisfied together. This protects customer application servers by keeping Fastly's CDN Service from sending duplicate requests to the origin information.

### Grace mode

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 customers 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. Customers must specifically configure their service to serve stale content (</guides/performance-tuning/serving-stale-content>).

---

## § About Fastly's Real-Time Log Streaming features (</guides/about-fastly-services/about-fastlys-realtime-log-streaming-features>)

---

Fastly's log streaming features allow logging configuration information to be sent over TLS (Transport Layer Security). This means that logging information can be encrypted, which allows customers to send potentially sensitive information to log files without exposing data.

### Supported protocols and logging providers

Fastly supports a variety of syslog-compatible logging providers, such as Sumo Logic (</guides/streaming-logs/log-streaming-sumologic>), Papertrail (</guides/streaming-logs/log-streaming-papertrail>), and Logentries (</guides/streaming-logs/log-streaming-logentries>). We also support other methods of sending logs besides the syslog (</guides/streaming-logs/log-streaming-syslog>) protocol. We allow pushing of log files to Amazon S3

(/guides/streaming-logs/log-streaming-amazon-s3) buckets as well as any S3-compatible providers (such as DreamHost's DreamObjects). Finally, Fastly supports FTP uploading (/guides/streaming-logs/log-streaming-ftp).

## How Real-Time Log Streaming works

As Varnish Cache software processes transactions, it writes instances to a memory log for engineers' diagnostic use. These memory instances are quickly overwritten, but selected data may be formatted into log lines and those can be continually transmitted or "streamed" in real time to a customer's chosen logging service or any other standard syslogd server for reporting purposes.

## § About Fastly's Video Caching service (/guides/about-fastly-services/about-fastlys-video-caching-service)

For customers with their own video packaging infrastructure, Fastly can act as a globally distributed HTTP cache to improve quality of service and increase viewer capacity. When a manifest or video segment is requested by a customer's player, a Fastly edge or shield POP will pull the requested content from the customer's origin media server. Subsequent requests for that content will be served from Fastly's cache servers instead of the customer's origin (read *How Fastly's CDN Services Work* (/guides/about-fastly-services/how-fastlys-cdn-service-works) for more information).

Fastly can cache and deliver any HTTP based media streaming protocol including:

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

In addition, both live and on-demand streams are supported.

## § About the web interface controls (/guides/about-fastly-services/about-the-web-interface-controls)

In addition to being accessible via Fastly's application programming interface (API) (/guides/about-fastly-services/about-fastlys-application-programming-interface), Fastly services can also be accessed via a web interface for users with the appropriate access permissions.

## Access to Fastly's web interface controls

Access to Fastly's web interface controls requires you to sign up for a Fastly account. Creating an account is free. Once you've created an account, you can navigate to the controls via the Fastly login page at <https://manage.fastly.com> (<https://manage.fastly.com>), either directly using any standard web browser or by clicking the Login link at the top right of almost all pages at the Fastly website (<https://www.fastly.com>).

Once logged in to a Fastly account, the web interface controls appear as appropriate based on the roles and permissions (/guides/user-access-and-control/user-roles-and-how-to-change-them) assigned to you.



The default control groups appear as follows from left to right across the top of the interface:


- the Dashboard page
- the Stats page
- the Configure controls
- the user menu

**NOTE:** Not all Fastly service features are enabled by default (e.g., [custom VCL](/guides/vcl/uploading-custom-vcl) (/guides/vcl/uploading-custom-vcl)). The appearance of the web interface controls may change from the defaults displayed once these services are enabled for your account.

## About the Dashboard page



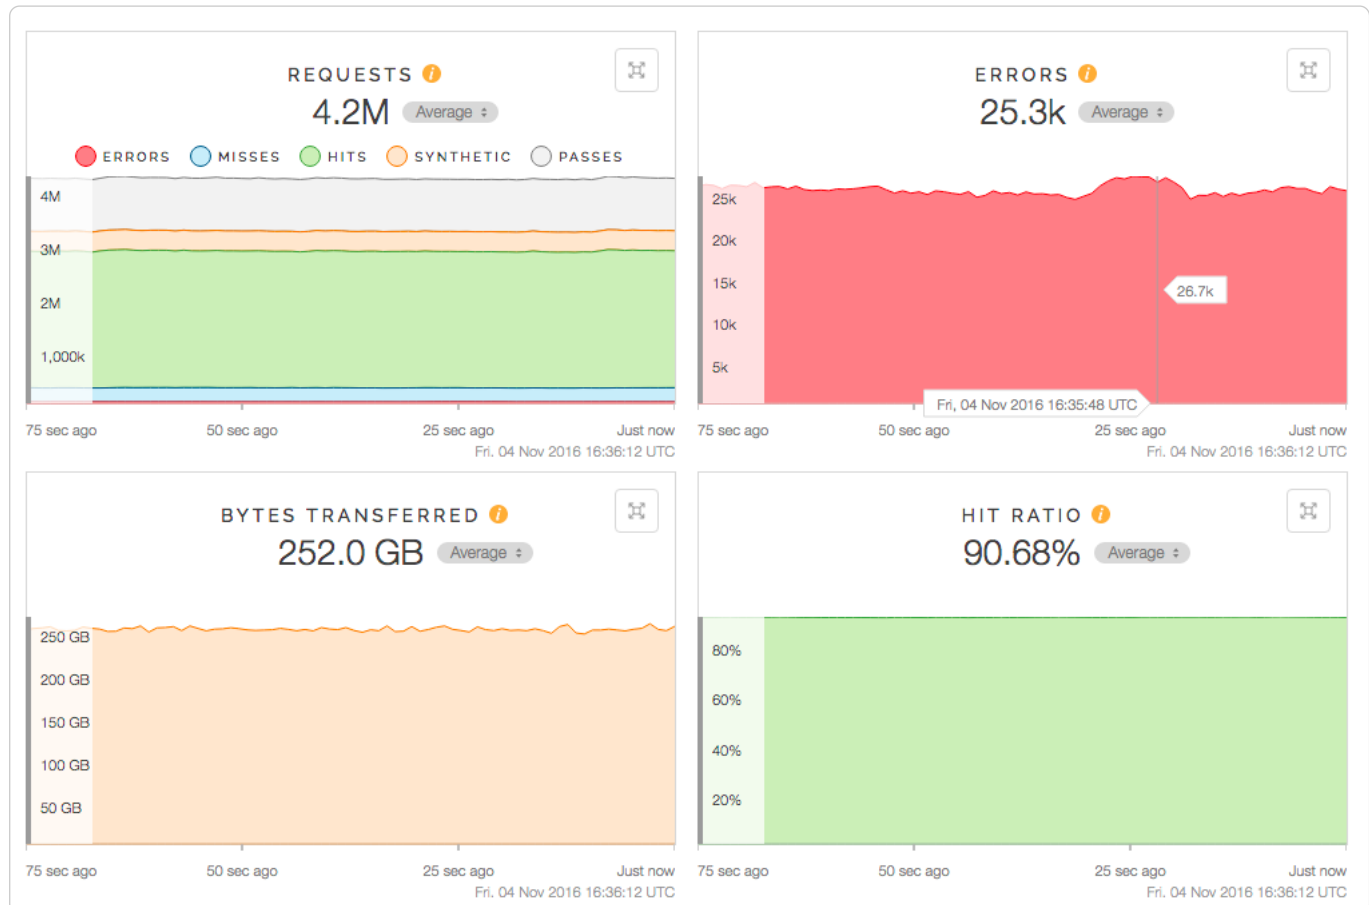
The Dashboard page allows you to separately monitor caching for each of your services in real time, as they operate on a second-by-second basis. The Dashboard link appears automatically for logged in users with the appropriate access permissions (/guides/user-access-and-control/user-roles-and-how-to-change-them); however, data on the page may appear grayed out or blank to some users, with no information displayed in the controls, when a customer's service does not receive enough requests for Fastly to display meaningful information about it in real time.



In addition to a menu allowing you to select the specific datacenter from which to view data (it defaults to data from all datacenters), the top of the dashboard includes the following real-time cache activity:

- **Hits:** the number of times requested data is found in cache
- **Hit Time:** the amount of time spent processing cache hits
- **Misses:** the number of times requested data is not found in cache
- **Hit Ratio:** the percentage of content being accessed that is currently cached by Fastly, defined as the proportion of cache hits to all cacheable content (hits + misses)
- **Miss Time:** the amount of time spent processing cache misses
- **Requests:** the total number of requests received for your site by Fastly
- **Errors:** the number of error requests that occurred

Below the real-time cache activity summary data, several graphs appear:



The graphed cache activity includes:


- **Requests:** a graph of the total number of requests received for your site by Fastly over time.
- **Errors:** a graph of the number of error requests that occurred over time.

- **Bytes Transferred:** a graph of the total number of bytes transferred by Fastly for your service over time.
- **Hit Ratio:** a graph of the percentage of content being accessed that is currently cached by Fastly over time.
- **Origin Latency** a histogram of the average amount of time to first byte (measured in milliseconds) on a cache miss. High origin latency means that your backends are taking longer to process requests.

One minute after Dashboard measurement data in these graphs rolls off the screen, it becomes available for retrieval on the Stats page (/guides/about-fastly-services/about-the-web-interface-controls#about-the-stats-page). You may not see any traffic right away because of the following:

- **Not enough data is going to your site.** If this is the case, visit the site yourself to trigger some traffic.
- **You've made a CNAME change.** If this is the case, it could take from a few minutes to hours for the change to propagate your DNS servers. See how to edit your DNS record (/guides/basic-setup/adding-cname-records#updating-the-cname-record-with-your-dns-provider) to point to Fastly for more information.

You can expand and minimize the view of some of the graphs using the quadruple arrow button in the right-hand corner of the graph. In addition, the Miss Latency graph includes a small gear icon in the upper right corner that allows you to change the interval limit displayed by the graph from the default 15 second interval. Finally, hovering the cursor over any part of a graph displays a timestamp indicator that updates itself as you move the mouse.



## About the Stats page

Fastly's Stats page provides a visual interface to our Stats API (/api/stats) that displays a series of graphs derived from your site's metrics statistical information. The Stats link appears automatically for logged in users with the appropriate access permissions (/guides/user-access-and-control/user-roles-and-how-to-change-them).



The displayed caching and performance metrics help you optimize your website's speed. These metrics include the following:

- **Hit Ratio** metrics tell you how well you are caching content using Fastly. This metric represents the proportion of cache hits versus all cacheable content (hits + misses). Increasing your hit ratio improves the overall performance benefit of using Fastly.
- **Cache Coverage** metrics show how much of your site you are caching with Fastly. This metric represents the ratio of cacheable requests (i.e., non "pass" requests) to total requests. Improving your cache coverage by reducing passes can improve site performance and reduce load on your origin servers.
- **Caching Overview** metrics compare Cache Hits, Cache Misses, Synthetic Responses (in VCL edge responses), and Passes (or requests that cannot be cached according to your configuration).

The traffic metrics analyze your website's traffic as it evolves over time. These metrics include the following:


- **Requests** metrics show you the total number of requests over time that were received for your site by Fastly.
- **Bytes Transferred** metrics show you the total number of bytes transferred by Fastly for your service.
- **Header & Body Bytes Transferred** metrics show you the relative values of bytes transferred when serving the body portion of HTTP requests and the header portion of the requests.
- **Miss Latency:** metrics show the distribution of the miss latency times for your origin.
- **Error Ratio** metrics show you the ratio of error responses (4XX & 5XX status code errors ([https://en.wikipedia.org/wiki/List\\_of\\_HTTP\\_status\\_codes](https://en.wikipedia.org/wiki/List_of_HTTP_status_codes))) compared to the total number of requests for your site. This metric allows you to quickly identify error spikes at given times.
- **HTTP Info, Success, & Redirects** metrics shows the number of HTTP Info (1XX), Success (2XX), and Redirect (3XX) statuses served for your site using Fastly.



- **Status 3XX Details** metrics shows the breakdown between the number of HTTP Status 301s, 302s, 304s, and other 3XX requests.
- **HTTP Client and Server Errors** metrics shows the number of HTTP Client Errors (4XX), and Server Errors (5XX) served for your site by Fastly.

## Controlling the historical information displayed


You can control how you view the historical information.



For all displayed graphs, you can choose:

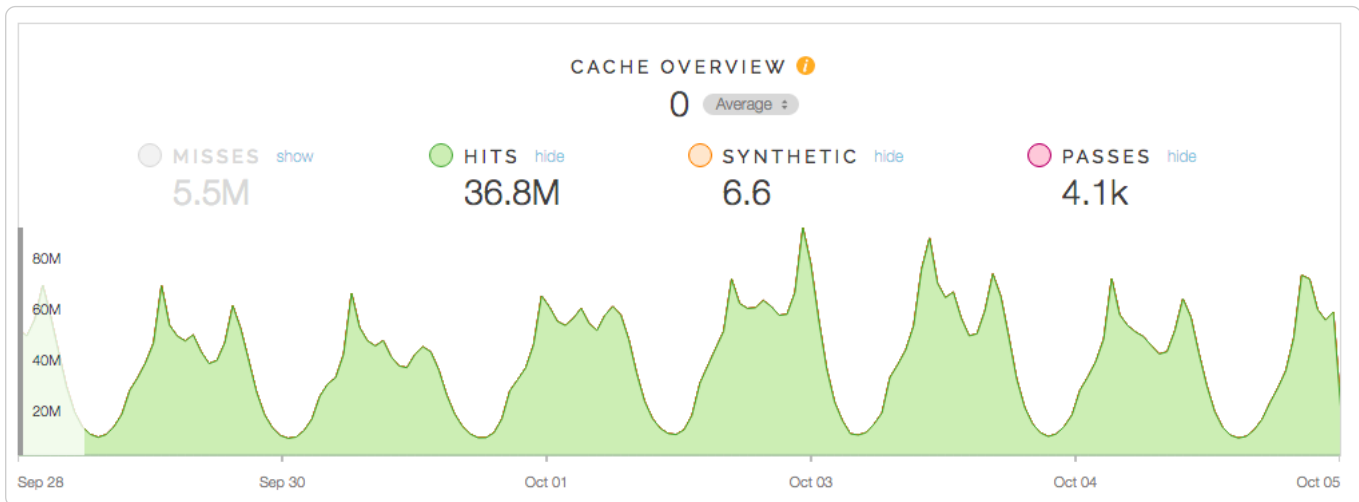
- the exact local date and time range of the graphed data
- how often to sample the data displayed
- whether to view global data for the graphs or only data from a specific region
- how to display the statistical values

You can change the statistics displayed in each graph. For example, notice the Average button on this Requests graph:



Clicking this button on any graph (including those on the Dashboard) lets you change the display of the graph's data to an average, a 95th percentile, a minimum, a maximum, or a total.

You can also exclude certain data entirely. For example, in this Caching Overview graph, hovering the cursor over the word "Hits" next to any of the data values displays a small, clickable **hide** link. Clicking this link will hide that value in the graph's display.



Notice that the actual numbers of the hidden data still appear grayed out in the controls, but the hide link changes to a show link and the graph itself doesn't display the hidden data at all.

## About the Configure page



The Configure page allows you to define exactly how each instance of your cache should behave and deliver content from data sources. The Configure page appears automatically for logged in users with the appropriate access permissions (</guides/user-access-and-control/user-roles-and-how-to-change-them/>). You use the Configure page to create versions of each of your service's configuration settings and then use the controls to activate or deactivate (</guides/basic-setup/working-with-services/>) them.


Specifically, you can configure and manage (</guides/basic-configuration/>):

- the domains used to route requests to a service
- the hosts used as backends for a site and how they should be accessed
- the health checks that monitor backend hosts
- various request and cache settings, headers, and responses that control how Fastly caches and serves content for a service
- how logging (</guides/streaming-logs/setting-up-remote-log-streaming>) should be performed and where server logs should be sent (as specified by an rsyslog endpoint)
- custom Varnish configuration language (</guides/vcl/>) (VCL) files if custom VCL is enabled (</guides/vcl/uploading-custom-vcl>)
- how conditions (</guides/conditions/>) are mapped and used for a service at various times (e.g., during request processing, when Fastly receives a backend response, or just before an object is potentially cached)

With the appropriate permissions, you can activate configuration changes immediately and roll back those changes just as quickly should they not have the intended effect. The Configure page also allows you to compare differences (</guides/basic-setup/working-with-services#comparing-different-service-versions>) between two configuration versions.

## About the user menu

The user menu appears at the far right of the default control group:




Depending on your access permissions (</guides/user-access-and-control/user-roles-and-how-to-change-them>), it provides access to a variety of account-specific and personal settings information. Specifically, it gives you access to Account details, Billing information, Help information, and What's new references. It also provides a way for you log out of the web interface.

## About the Account page

The Account page allows you to manage various user access and control (</guides/user-access-and-control/>) settings including not only the personal settings for your account but also certain roles and access permissions of other users you have invited (</guides/user-access-and-control/adding-and-deleting-user-accounts>) to use the account. In addition, this page allows you to manage account access and security (</guides/account-management-and-security/>) and even cancel your account (</guides/account-types-and-billing/accounts-and-pricing-plans.html#canceling-your-account>) should the need arise.

## About the Billing page

This page shows you at a glance whether or not your account is paid in full.



The Month-To-Date tab on this page provides information about month-to-date account changes for the current billing month, including complete details about the bandwidth and requests consulted when calculating your bill (</guides/account-types-and-billing/how-we-calculate-your-bill>).

The Invoice History tab on this page provides a complete billing history for your account. You can review details for any given month by clicking on the desired link in the Date column.

The Pricing Plan tab on this page provides information about your current pricing plan, and allows you to switch to a paid plan (</guides/account-types-and-billing/accounts-and-pricing-plans.html#paid-accounts-without-contracts>) once your development trial period ends. You can use it to change details about the credit card we'll charge when you pay your bill (</guides/account-types-and-billing/paying-your-bill>).

The Tax Address tab on this page allows you to add or update the tax address.

**NOTE:** Billing controls appear automatically for logged in users assigned the [billing or superuser roles](/guides/user-access-and-control/user-roles-and-how-to-change-them) (</guides/user-access-and-control/user-roles-and-how-to-change-them>).



## About the Help page

The Help page allows you to search Fastly documentation for immediate help with questions you may have about using and accessing Fastly services. If you can't find what you need in the documentation, you can also use this page to submit a request (/guides/customer-support/submitting-support-requests) for help directly from Fastly Customer Support.

## About the What's new panel

The What's new panel provides a brief overview of the changes most recently made to the Fastly web interface.

---

# § Content and its delivery (/guides/about-fastly-services/content-and-its-delivery)

---

## Content types delivered by Fastly

The underlying protocol used by the World Wide Web to define how content is formatted and transmitted is called the Hypertext Transfer Protocol (HTTP). Fastly's CDN Service delivers all HTTP-based file content (e.g., HTML, GIF, JPEG, PNG, JavaScript, CSS) including the following:

- Static content
- Dynamic content
- Video content

Each content type is described below.

### Static content

Static content includes content that remains relatively unchanged. Fastly can control static content in two ways:

- using the time to live (TTL) method, where Fastly's cache re-validates the content after expiration of the TTL, or
- using Fastly's Instant Purge functionality, in which content remains valid until the cache receives a purge request (/guides/purging/) that invalidates the content.

Examples of static content include images, css, and javascript files.

### Dynamic content

Dynamic content includes content that changes at unpredictable intervals, but can still be cached for a fraction of time. We serve this dynamic content by taking advantage of Fastly's Instant Purge functionality. Using this functionality, dynamic content remains valid only until a Fastly cache receives a purge request (/guides/purging/) that invalidates the content. Fastly understands that the rate of those purge requests cannot be predicted. Dynamic content may change frequently as a source application issues purge requests in rapid succession to keep the content up to date. Dynamic content can, however, remain valid for months if there are no changes requested.

Examples of dynamic content include sports scores, weather forecasts, breaking news, user-generated content, and current store item inventory.

### Video content

Video content includes:

- Live video streams
- Video on Demand (VOD) content libraries

Video content can be served using standard HTTP requests. Specifically, Fastly supports HTTP Streaming standards, including HTTP Live Streaming (HLS), HTTP Dynamic Streaming (HDS), HTTP Smooth Streaming (HSS), and MPEG-DASH. For Fastly's CDN Service to deliver video, the video must be packaged.

## Content sources supported by Fastly

Fastly caches deliver various types of content from many different sources. Supported sources include:

- Websites
- Internet APIs
- Internet Applications
- Live and Live Linear Video
- Video on Demand (VOD) Libraries

Regardless of the content source, the content's source server must communicate using HTTP. HTTP defines specific types of "methods" that indicate the desired action to be performed on content. The manner in which those HTTP methods are used (the standard, primary methods being GET, POST, PUT, and DELETE) can be labeled as being RESTful ([https://en.wikipedia.org/wiki/Representational\\_state\\_transfer](https://en.wikipedia.org/wiki/Representational_state_transfer)) or not. Fastly supports RESTful HTTP by default, but also can support the use of non-RESTful HTTP as long as the method used is mapped to its appropriate cache function. Each of the content sources supported by Fastly are described in more detail below.

Websites

Websites are servers that provide content to browser applications (e.g., Google's Chrome, Apple's Safari, Microsoft's Internet Explorer, Opera Software's Opera) when end users request that content. The content contains both the requested data and the formatting or display information the browser needs to present the data visually to the end user.

With no CDN services involved, browsers request data by sending HTTP GET requests that identify the data with a uniform resource locator (URL) address to the origin server that has access to the requested data. The server retrieves the data, then constructs and sends an HTTP response to the requestor. When a CDN Service is used, however, the HTTP requests go to the CDN rather than the origin server because the customer configures it to redirect all requests for data to the CDN instead. Customers do this by adding a CNAME or alias for their origin server that points to Fastly instead.

Internet APIs

Application program interfaces (APIs) serve as a language and message format that defines exactly how a program will interact with the rest of the world. APIs reside on HTTP servers. Unlike the responses from a website, content from APIs contain only requested data and identification information for that data; no formatting or display information is included. Typically the content serves as input to another computing process. If it must be displayed visually to an end user, a device application (such as, an iPad, Android device, or iPhone Weather application) does data display instead.

Legacy internet applications

Legacy Internet applications refer to applications not originally developed for access over the Internet. These applications may use HTTP in a non-RESTful manner. They can be incrementally accelerated without caching, benefiting only from the TCP Stack optimization done between edge Fastly POPs and the Shield POP, and the Shield POP to the origin. Then caching can be enabled incrementally, starting with the exchanges with the greatest user-experienced delay.

Live and live linear video streams & video on demand libraries

Live and live linear video content (for example, broadcast television) is generally delivered as a "stream" of information to users, which they either choose to watch or not during a specific broadcast time. Video on demand (VOD), on the other hand, allows end users to select and watch video content when they choose to, rather than having to watch at a specific broadcast time.

Regardless of which type of video content an end user experiences, a video player can begin playing before its entire contents have been completely transmitted. End users access the video content from a customer's servers via HTTP requests from a video player application that can be embedded as a part of a web browser. Unlike other types of website content, this content does not contain formatting or display information. The video player handles the formatting and display instead.

When the video content is requested, the customer's server sends the content as a series of pre-packaged file chunks along with a manifest file required by the player to properly present the video to the end user. The manifest lists the names of each file chunk. The video player application needs to receive the manifest file first in order to know the names of the video content chunks to request.

"Pre-packaging" in this context refers to the process of receiving the video contents, converting or "transcoding" the stream into segments (chunks) for presentation at a specific dimension and transmission rate, and then packaging it so a video player can identify and request the segments of the live video a user wants to view.

To request video delivery on your account, contact your Fastly Account Representative at [sales@fastly.com](mailto:sales@fastly.com) (<mailto:sales@fastly.com>).

§ Domain names and Fastly's services (/guides/about-fastly-services/domain-names-and-fastly-services)

A domain name is a component of a Uniform Resource Locator ([https://en.wikipedia.org/wiki/Uniform\\_Resource\\_Locator](https://en.wikipedia.org/wiki/Uniform_Resource_Locator)) (URL). A domain name represents an Internet Protocol ([https://en.wikipedia.org/wiki/Internet\\_Protocol](https://en.wikipedia.org/wiki/Internet_Protocol)) (IP) resource and, in this context, refers to the IP address of a server computer hosting a website ([https://en.wikipedia.org/wiki/Web\\_site](https://en.wikipedia.org/wiki/Web_site)), the website itself, or any other service communicated via the Internet. Customers associate their domain names with their origin (content source) when provisioning a Fastly service.

Domain names are made up of components as follows:

Domain Name Component	Example
-----------------------	---------

URL	<a href="http://www.example.com/index.html">http://www.example.com/index.html</a>
Top-Level Domain Name	com
Second-Level Domain Name	example.com
Hostname	www.example.com

Domain names are registered with a domain name registrar. Fastly is not a domain name registrar.

Fastly allows the use of domain names as origins. There is a performance impact, but it allows things like EC2 machines and App Engine instances to be used. In fact, when we detect certain conditions we optimize our routing to be more efficient.

Fastly supports the use of multiple subdomains for the same origin server, and allows the specification of any number of subdomains for each origin. Some customers use only their second level or apex domain (</guides/basic-configuration/using-fastly-with-apex-domains>) (e.g., example.com rather than www.example.com) as their canonical domain. Due to limitations in the DNS specification, we don't recommend placing a CNAME record at the apex domain or using the CNAME Flattening (e.g., ALIAS or ANAME) features offered by some DNS providers. Instead, we offer Anycast IP addresses (</guides/basic-configuration/using-fastly-with-apex-domains#anycast-option>) for content that must be hosted on a second-level or apex domain.

## § Fastly POP locations (</guides/about-fastly-services/fastly-pop-locations>)

Our points of presence (POPs) on the Internet are strategically placed at the center of the highest density Internet Exchange Points around the world. Fastly's Network Map (<https://www.fastly.com/network-map>) shows a detailed view of the current and planned locations of all Fastly POPs. For more information, about how we choose these locations, check out our blog post [How Fastly Chooses POP Locations](https://www.fastly.com/blog/how-fastly-chooses-pop-locations) (<https://www.fastly.com/blog/how-fastly-chooses-pop-locations>).

Once you're signed up for Fastly service (either through a test account (</guides/account-types-and-billing/accounts-and-pricing-plans>) or a paid plan) you can see a live, real-time visual representation (</guides/about-fastly-services/about-the-web-interface-controls#about-the-global-traffic-map>) of the general regions of the world in which Fastly's points of presence (POPs) receive requests for your service.

### Will Fastly ever adjust POP locations or service regions? How will I be notified?

Fastly continues to grow its network footprint, adding new service POPs in the process. At times, expansion may result in the addition of new billable regions to our network. We'll announce new POP locations and new billable regions in advance through our network status page (</guides/debugging/fastlys-network-status>) at [status.fastly.com](https://status.fastly.com) (<https://status.fastly.com/>). Contact [sales@fastly.com](mailto:sales@fastly.com) (<mailto:sales@fastly.com>) with specific contract or billing questions.

## § How caching and CDNs work (</guides/about-fastly-services/how-caching-and-cdns-work>)

Fastly is a Content Delivery Network, or CDN. CDNs work on the principle that once a piece of content has been generated it doesn't need to be generated again for a while so a copy can be kept around in a cache. Cache machines are optimized to serve small files *very very* quickly. CDNs typically have caches placed in datacenters all around the world. When a user requests information from a customer's site they're actually redirected to the set of cache machines closest to them instead of the customer's actual servers. This means that a European user going to an American site gets their content anywhere from 200-500ms faster. CDNs also minimize the effects of a cache miss. A cache miss occurs when a user requests a bit of content and it is not in the cache at that moment (because it's expired, because no-one has asked for it before, or because the cache got too full and old content was thrown out).

### What can be cached?

CDNs are good at managing a cache of small, static resources (for example, static images, CSS files, Javascripts, and animated GIFs). CDNs are also popular for offloading expensive-to-serve files like video and audio media.

At Fastly, our architecture (known as a *reverse proxy*) is designed to enable customers to go a step further and cache entire web pages for even more efficient handling of your traffic.

★ **TIP:** Static files + media objects + web pages = your whole site. With the right service configuration (which we can assist you in setting up) Fastly can reduce your backend traffic by orders of magnitude with no loss in control over the content your users see.

# Managing the Cache

Caching serves as a powerful weapon in your make-the-site-faster arsenal. However, most objects in your cache aren't going to stay there permanently. They'll need to expire so that fresh content can be served. How long that content should stay in the cache might be mere seconds or a number of minutes or even a year or more.

How can you manage which of your content is cached, where, and for how long? By setting policies that control the cached data. Most caching policies are implemented as a set of HTTP headers sent with your content by the web server (as specified in the configuration or the application). These headers were designed with the client (browser) in mind but CDNs like Fastly will also use those headers as a guide on caching policy.

## Expires

The `Expires` header is the original cache-related HTTP header and tells the cache (typically a browser cache) how long to hang onto a piece of content. Thereafter, the browser will re-request the content from its source. The downside is that it's a static date and if you don't update it later, the date will pass and the browser will start requesting that resource from the source every time it sees it.

If none of the following headers are found in the request, Fastly will respect the `Expires` header value.

## Cache-Control

The `Cache-Control` headers (introduced in the HTTP 1.1 specification) cover browser caches and in most cases, intermediate caches as well:

- `Cache-Control: public` - Any cache can store a copy of the content.
- `Cache-Control: private` - Don't store, this is for a single user.
- `Cache-Control: no-cache` - Re-validate before serving this content.
- `Cache-Control: no-store` - Don't ever store this content.
- `Cache-Control: public, max-age=[seconds]` - Caches can store this content for  $n$  seconds.
- `Cache-Control: s-maxage=[seconds]` - Same as max-age but applies specifically to proxy caches.

Only the `max-age`, `s-maxage`, and `private` Cache-Control headers will influence Fastly's caching. All other Cache-Control headers will not, but will be passed through to the browser. For more in-depth information about how Fastly responds to these Cache-Control headers and how these headers interact with Expires and Surrogate-Control, check out our Cache Control Tutorial (/guides/tutorials/cache-control-tutorial).

**NOTE:** For more information on the rest of the Cache-Control headers, see the relevant section in Mark Nottingham's [Caching Tutorial](https://www.mnot.net/cache_docs/#CACHE-CONTROL) ([https://www.mnot.net/cache\\_docs/#CACHE-CONTROL](https://www.mnot.net/cache_docs/#CACHE-CONTROL)).

## Surrogate Headers

`Surrogate` headers are a relatively new addition to the cache management vocabulary (described in this W3C tech note (<http://www.w3.org/TR/edge-arch/>)). These headers provide a specific cache policy for proxy caches in the processing path. `Surrogate-Control` accepts many of the same values as `Cache-Control`, plus some other more esoteric ones (read the tech note for all the options).

One use of this technique is to provide conservative cache interactions to the browser (for example, `Cache-Control: no-cache`). This causes the browser to re-validate with the source on every request for the content. This makes sure that the user is getting the freshest possible content. Simultaneously, a `Surrogate-Control` header can be sent with a longer `max-age` that lets a proxy cache in front of the source handle most of the browser traffic, only passing requests to the source when the proxy's cache expires.

With Fastly, one of the most useful `Surrogate` headers is `Surrogate-Key`. When Fastly processes a request and sees a `Surrogate-Key` header, it uses the space-separated value as a list of tags to associate with the request URL in the cache. Combined with Fastly's Purge API (/api/purge) an entire collection of URLs can be expired from the cache in one API call (and typically happens in around 1ms). `Surrogate-Control` is the most specific.

## Fastly and Cache Control Headers

Fastly looks for caching information in each of these headers as described in our Cache-Control docs (/guides/tutorials/cache-control-tutorial). In order of preference:

- `Surrogate-Control:`
- `Cache-Control: s-maxage`
- `Cache-Control: max-age`
- `Expires:`

## Shielding

When an object or collection of objects in the cache expires, the next time any of those objects are requested, the request is going to get passed through to your application. Generally, with a good caching strategy, this won't break things. However, when a popular object or collection of objects expires from the cache, your backend can be hit with a large influx of traffic as the cache nodes refetch the objects from the source.

In most cases, the object being fetched is not going to differ between requests, so why should every cache node have to get its own copy from the backend? With Shield Nodes, they don't have to. Shielding configured through the Fastly web interface (</guides/performance-tuning/shielding>) allows you to select a specific datacenter (most efficiently, one geographically close to your application) to act as a shield node. When objects in the cache expire, the shield node is the only node to get the content from your source application. All other cache nodes will fetch from the shield node, reducing source traffic dramatically.

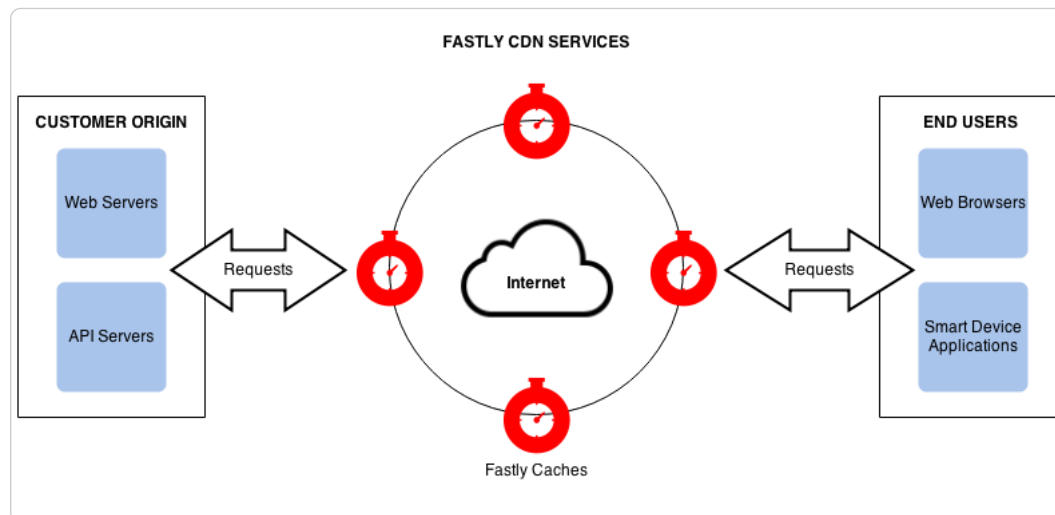
## Resources

- Wikipedia: Reverse Proxy ([https://en.wikipedia.org/wiki/Reverse\\_proxy](https://en.wikipedia.org/wiki/Reverse_proxy))
- Fastly's Cache-Control docs (</guides/tutorials/cache-control-tutorial>)
- Mark Nottingham's Caching Tutorial ([https://www.mnot.net/cache\\_docs/#CACHE-CONTROL](https://www.mnot.net/cache_docs/#CACHE-CONTROL))
- Surrogate header W3C tech note (<http://www.w3.org/TR/edge-arch/>)

## § How Fastly's CDN Service Works (</guides/about-fastly-services/how-fastlys-cdn-service-works>)

Fastly is a content delivery network ([https://en.wikipedia.org/wiki/Content\\_delivery\\_network](https://en.wikipedia.org/wiki/Content_delivery_network)) (CDN). We serve as an Internet intermediary and offer the Fastly CDN Service to make transmission of your content to your end users more efficient.

You can make content available through your websites and Internet-accessible (hosted) application programming interfaces (APIs). You can create content (customer-generated content), as can your end users (user-generated content). Fastly's CDN Service then makes the transmission of that content (which we sometimes refer to as "content objects") more efficient by automatically storing copies at intermediate locations on a temporary basis. The process of storing these copies is known as "caching" and the server locations in which they are stored are referred to as "caches."




Fastly's delivers its CDN service from key access points to the Internet called "points of presence" (POPs). Fastly places POPs (<https://www.fastly.com/blog/how-fastly-chooses-pop-locations>) where their connectivity to the Internet reduces network transit time when delivering content to end-users. Each POP has a cluster of Fastly cache servers. When end users request your content objects, Fastly delivers them from whichever of the cache locations are closest to each end user.


Fastly's caches only receive and process your end user requests for content objects. You decide which objects will be cached, for how long, who can access them, whether they are to be encrypted when transmitted over the Internet, and when the objects will be deleted from the caching service. You make these decisions by specifically configuring Fastly's CDN Service with these requirements. We refer to this configuration process as "provisioning."

To provision Fastly's CDN service (</guides/basic-setup/sign-up-and-create-your-first-service>), you must identify which of your application servers will provide the original content objects for each of your various domains (e.g., `company.com`, `myco.com`). Your application servers can be physical servers in a datacenter or hosting facility, or applications running on cloud services like Amazon, or any combination. Fastly refers to these source servers as "origin" and "backend" servers interchangeably.

The first time each Fastly cache receives a request for a content object, it fetches the object from the appropriate origin server. If multiple origin servers are specified, the cache will distribute the processing load for the fetches across all of them (based on the configuration criteria set by you). After the content object is fetched, the cache stores a copy of it and forwards its response to the end user.



Each time after the first time an end user requests that same content object, the Fastly cache fulfills requests by retrieving the cached copy from storage (or memory) and immediately delivering it to the end user – the fetch step to the original copy is not repeated until the content object either expires or becomes invalidated.



## Can Fastly host my content?

We accelerate your site by caching both static assets and dynamic content by acting as a reverse proxy ([https://en.wikipedia.org/wiki/Reverse\\_proxy](https://en.wikipedia.org/wiki/Reverse_proxy)) to your origin server (also known as "Origin Pull"), but we do not provide services for uploading your content to our servers.

In addition to using your own servers as the source, we also support various "cloud storage" services as your origin, such as Amazon Elastic Compute Cloud (</guides/integrations/amazon-s3>) (EC2), Amazon Simple Storage Service (</guides/integrations/amazon-s3>) (S3), and Google Cloud (</guides/integrations/google-cloud-storage>) as your file origin. Our partnership with Google (<https://www.fastly.com/partner/gcp>) in particular enables us to have direct connectivity to their cloud infrastructure.

## § HTTP status codes cached by default (/guides/about-fastly-services/http-status-codes-cached-by-default)

Fastly caches the following response status codes by default. In addition to these statuses, you can force an object to cache under other states using conditions (</guides/conditions/>) and responses (</guides/basic-configuration/responses-tutorial>).

Code	Message
200	OK
203	Non-Authoritative Information
300	Multiple Choices
301	Moved Permanently
302	Moved Temporarily
404	Not Found
410	Gone

## § Self-provisioned Fastly services (/guides/about-fastly-services/self-provisioned-fastly-services)

You can configure or "provision" Fastly caching and video services personally, independent of Fastly staff, via the Fastly web interface (<https://manage.fastly.com>). Fastly calls this "self-provisioning." Self-provisioning tasks include things like:

- creating and activating services
- adding domains and origin servers
- configuring load balancing

- modifying how services handle HTTP headers
- submitting purge commands

Once provisioned, Fastly services can be activated immediately. If self-provisioned tasks fail to operate in an appropriate or expected manner, you can find answers to a variety of frequently asked questions in Fastly's guides and tutorials (/guides/). You can also receive personalized assistance by submitting requests (/guides/customer-support/submitting-support-requests) directly to Fastly's Customer Support staff.

• Guides (/guides/) > Advanced setup (/guides/advanced) > Securing communications (/guides/securing-communications/)

## § Accessing Fastly's IP ranges (/guides/securing-communications/accessing-fastlys-ip-ranges)

To help you whitelist Fastly's services through your firewall, we provide access to the list of Fastly's assigned IP ranges. You can access the list via URL:

<https://api.fastly.com/public-ip-list> (<https://api.fastly.com/public-ip-list>)

You can then automate the API call (for example, by running a script (<https://github.com/jondade/IP-Whitelist-cron>) as a cron job) to request the list of IPs to detect when the IP ranges change.

To make sure you have plenty of time to stay in sync, we post IP address announcements (<https://status.fastly.com/incidents/sg8850gdhv26>) along with other service announcements to our status page (<https://status.fastly.com/>), which you can subscribe (/guides/debugging/fastlys-network-status#subscribing-to-notifications) to.

## § Allowing only TLS connections to your site (/guides/securing-communications/allowing-only-tls-connections-to-your-site)

If you want to only allow TLS on your site, we have you covered. There is a switch built into the "Request Settings" that will allow you to force unencrypted requests over to TLS. It works by returning a **301 Moved Permanently** response to any unencrypted request, which redirects to the TLS equivalent. For instance, making a request for **http://www.example.com/foo.jpeg** would redirect to **https://www.example.com/foo.jpeg**.

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.
4. Click the **Settings** tab. The Settings page appears.
5. Click the **Create request setting** button. The Create a new request setting page appears.



## Create a new request setting

CONDITION

This will happen all the time unless you [Attach a condition](#)

Action

Do nothing now

Force a VCL action to happen. Options are:

**Do nothing now** - No action.

**Lookup** - Force an immediate lookup in the cache.

**Pass** - Force a request to the origin server.

Learn about [how are request settings applied](#).

Force TLS

Yes

Select **Yes** to force unencrypted requests over to TLS, return a **301 Moved Permanently** response to any unencrypted request, and redirect to the TLS equivalent. See our [TLS encryption guide](#) for more info.

X-Forwarded-For

Leave

Stipulate how the **X-Forwarded-For** header should be treated. Learn about [how to manipulate the X-Forwarded-For header](#).

► [Details for X-Forwarded-For options](#)

Description

Force TLS

★ Required

The name of your request setting, such as **My request setting**.

6. Fill out the **Create a new request setting** fields as follows:

- From the **Force TLS** menu, select **Yes**.
- In the **Description** field, type a human-readable name for the request setting. This name is displayed in the Fastly web interface.

7. Click the **Create** button to save your request setting changes.

8. Click the **Activate** button to deploy your configuration changes.

For more information about TLS-related issues, see our TLS guides (</guides/securing-communications/>) or contact [support@fastly.com](mailto:support@fastly.com) (mailto:support@fastly.com) with questions.

## § Connecting to origins over TLS (</guides/securing-communications/connecting-to-origins-over-tls>)

We provide two ways to use TLS to communicate to your origin servers.

### The simple way

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.
4. Click the **Origins** tab. The Origins page appears.
5. Click the **Create host** button. The Create a new host page appears.



## Create a new host

### CONDITION

This will happen all the time unless you [Attach a condition](#)

#### Address



★ Required

The IP address (or hostname) and port number for your origin server. To [enable TLS](#), use port 443.

#### Shielding

The shield POP designated to reduce inbound load on your origins by serving cached data. [Learn more about POP request handling](#) and the [caveats of shielding](#).

#### Health check

[+ Create a new health check](#)

A small request assigned to [test the state of your origin](#).

#### Auto load balance

Enables Fastly to [load balance between two or more backend hosts](#).

[► Tips](#)

#### Description

★ Required

The name of your origin, such as **My origin server**.

6. Fill out the **Create a new host** fields as follows:

- In the **Address** field, type the address of your secure server (for example, `origin.example.com`), and in the port field type `443`.
- Leave the **Shielding**, **Health check**, and **Auto load balance** controls set to their default values.
- In the **Description** field, type the name of your server (for example, `My Origin Server`). This name is displayed in the Fastly web interface.

7. Click the **Create** button.

8. Click the **Activate** button to deploy your configuration changes.

**❗ IMPORTANT:** For port 443, and only port 443, we will enable TLS automatically in our system. No other configuration is needed and using this port often can help avoid triggering TLS redirect loops. If you find that setting the port specifically to 443 doesn't produce the results you're hoping for or simply isn't working, consider trying the more complicated settings detailed below.

## The detailed way

To use a different port, or for a little more control over the TLS connections, follow these steps.

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.
4. Click the **Origins** tab. The Origins page appears.
5. Click the **Create host** button. The Create a new host page appears.

## Create a new host

**CONDITION** This will happen all the time unless you [Attach a condition](#)

Address  :  ★ Required

The IP address (or hostname) and port number for your origin server. To [enable TLS](#), use port 443.

Shielding

The shield POP designated to reduce inbound load on your origins by serving cached data. Learn more about [POP request handling](#) and the [caveats of shielding](#).

Health check

[+ Create a new health check](#)

A small request assigned to [test the state of your origin](#).

Auto load balance

Enables Fastly to [load balance between two or more backend hosts](#).

[► Tips](#)

Description  ★ Required

The name of your origin, such as [My origin server](#).

6. Fill out the **Create a new host** fields as follows:

- In the **Address** field, type the address of your secure server (for example, `origin.example.com`), and in the port field type a port number (for example, `8443`).
- Leave the **Shielding**, **Health check**, and **Auto load balance** controls set to their default values.
- In the **Description** field, type the name of your server (for example, `My Origin Server`). This name is displayed in the Fastly web interface.

7. Click the **Create** button.

8. Click the **TLS Options** link next to the host you just created. The TLS options page appears.

9. From the **Connect to backend using TLS** menu, select **Yes**.

10. From the **Verify certificate** menu, select **Yes**.

**⚠ WARNING:** Selecting No from the Verify Certificate menu means the certificate will not be verified. Not verifying the certificate has serious security implications, including vulnerability to [man-in-the-middle attack](https://www.owasp.org/index.php/Man-in-the-middle_attack) ([https://www.owasp.org/index.php/Man-in-the-middle\\_attack](https://www.owasp.org/index.php/Man-in-the-middle_attack)). Consider uploading a CA certificate instead of disabling certificate validation.

11. In the **Certificate Hostname** field, type the hostname associated with your TLS certificate. This value is matched against the certificate common name (CN) or a subject alternate name (SAN) depending on the certificate you were issued. For example, if your certificate's CN field is `www.example.com`, type that value for your hostname. If you leave this field blank, the system will use the default host information displayed below this field.

12. If you are specifying a TLS CA certificate, see the section below.

13. Click the **Save** button. Even if you leave the Certificate Hostname field blank to use the default information, you must click the Save button to verify the certificate.

14. Click the **Activate** button to deploy your configuration changes.

And that's all you need to do. Everything else is optional, but just in case you'd like to set them, we've included the information below.

## Setting the TLS hostname

Normally we check the server certificate against the hostname portion of the address for your origin entered in the **Create a new host** window. To have Fastly verify the certificate using a different hostname, specify it via the **SNI Hostname** field under **Advanced Options**.

This information also gets sent to the server in the TLS handshake. If you are using Server Name Indication ([https://en.wikipedia.org/wiki/Server\\_Name\\_Indication](https://en.wikipedia.org/wiki/Server_Name_Indication)) (SNI) to put multiple certificates on your origin, specifying it in the **SNI Hostname** field will select which one is used.

## Specifying a TLS CA certificate

If you're using a certificate that is either self-signed or signed by a Certificate Authority (CA) not commonly recognized by major browsers (and unlikely to be in the Ubuntu bundle that we use), you can provide the certificate in PEM format via the **TLS client certificate** field under **Advanced Options**. The PEM format looks like this:

[illegible]

## Specifying a TLS client certificate and key

To ensure TLS connections to your origin come from Fastly and aren't random, anonymous requests, set your origin to verify the client using a client certificate. Simply paste the certificate and private key in PEM form into the appropriate text boxes on the **TLS options** page.

**❗ IMPORTANT:** The private key must not be encrypted with a passphrase.

Then configure your backend to require client certificates and verify them against the CA cert they were signed with. Here are some ways of doing that:

- Apache ([http://httpd.apache.org/docs/2.4/ssl/ssl\\_howto.html#accesscontrol](http://httpd.apache.org/docs/2.4/ssl/ssl_howto.html#accesscontrol))
- Nginx ([http://nginx.org/en/docs/http/nginx\\_http\\_ssl\\_module.html#ssl\\_client\\_certificate](http://nginx.org/en/docs/http/nginx_http_ssl_module.html#ssl_client_certificate))
- IIS (<https://ondrej.wordpress.com/2010/01/24/iis-7-and-client-certificates/>)

## Specifying acceptable TLS protocol versions

If your origin server is configured with support for modern TLS protocol versions, you can customize the TLS protocols Fastly will use to connect to it by setting a **Minimum TLS Version** and **Maximum TLS Version** under **Advanced Options**. We recommend setting both to the most up-to-date TLS protocol, currently 1.2, if your origin can support it.

Use the `openssl` command to verify your origin supports a given TLS protocol version. For example:

```
openssl s_client -connect origin.example.com:443 -tls1_2
```

Replace `-tls1 2` with `tls1 1` and `tls1 0` to test other protocol versions. Fastly does not support SSLv2 or SSLv3.

## Specifying acceptable TLS cipher suites

Fastly supports configuring the OpenSSL cipher suites used when connecting to your origin server. This allows you to turn specific cipher suites on or off based on security properties and origin server support. The **Ciphersuites** setting under **Advanced Options** accepts an OpenSSL formatted cipher list (<https://www.openssl.org/docs/manmaster/apps/ciphers.html>). We recommend using the strongest cipher suite your origin will support as detailed by the Mozilla SSL Configuration Generator (<https://mozilla.github.io/server-side-tls/ssl-config-generator/>).

Use the `openssl` command to verify your origin supports a given cipher suite. For example:

```
openssl s_client -connect origin.example.com:443 -tls1 2 -cipher ECDHE-RSA-AES128-GCM-SHA256
```

Replace `-cipher ECDHE-RSA-AES128-GCM-SHA256` with the cipher suite to test.

## § Domain validation for TLS certificates (/guides/securing-communications/domain-validation-for-tls-certificates)

---

When you purchase our Shared Certificate (/guides/securing-communications/ordering-a-paid-tls-option#shared-certificate) or Shared Wildcard Certificate (/guides/securing-communications/ordering-a-paid-tls-option#shared-wildcard-certificate-service) TLS options, our partner Certificate Authority (GlobalSign) must verify you control the domains requested and that you authorize us to request a certificate service on your behalf. You can choose:

- DNS text record verification (preferred)
- URL verification
- Email verification

Regardless of the verification method you use, be sure to follow our instructions (/guides/securing-communications/ordering-a-paid-tls-option) to begin the TLS ordering process.

### DNS text record verification

We provide you with a unique DNS TXT record that you need to add for the zone origin ("@"). The text of this entry will change depending on the certificate to which your domain is added, but will generally look something like this (where the `{meta tag}` will change depending on the certificate):

```
@ IN TXT "globalsign-domain-verification={meta tag}"
```

Consult the documentation for your DNS server or hosted DNS provider for more information about how to add the record. This text record must be wholly separate from other text records. A prepended, inserted, or appended record will not work.

### URL verification

We provide you with an HTML meta tag that must be included in the `<head>` section of the web page served at the root of the domain to be added. The meta tag will be formatted similar to the following (where the `{meta tag}` text will change depending on the certificate):

```
<meta name="globalsign-domain-verification" content="{meta tag}" />
```

Note that this text must be served from the actual requested domain or root domain. For example, GlobalSign will specifically query `http://www.example.com`, so the verification tag must be served from whatever resource is returned from that URL. GlobalSign will not follow redirects or request a file on that domain (such as `http://www.example.com/verify.html` or `http://www.example.com/index.html`). It will also work on `http://example.com`.

### Email verification

GlobalSign will give Fastly a list of acceptable email addresses to which they can send a validation email. Generally these email addresses match those that appear on the WHOIS record of the domain requested, plus the following:

- admin@domain.com
- administrator@domain.com
- hostmaster@domain.com
- postmaster@domain.com
- webmaster@domain.com

For entries requested for a subdomain, each of those addresses `@subdomain.domain.com` will also work (e.g., `admin@subdomain.domain.com`).

We will send you the list of acceptable email address. You will need to tell us which email address to use. GlobalSign will then send a verification email to the email address you specify. Once you receive the verification email, you will need to click on a link in that email and follow the instructions to complete the validation.

---

## § Enabling HSTS through Fastly (/guides/securing-communications/enabling-hsts-through-fastly)

---

The HTTP Strict Transport Security (<https://tools.ietf.org/html/rfc6797>) (HSTS) security enhancement specification provides a way to force modern browsers to communicate only via the Transport Layer Security (TLS) protocol. Once enabled, it will force the browser to redirect (typically with a status code 307) to the HTTPS URL.

**NOTE:** HSTS only takes effect *after* a site has been visited on a trusted HTTPS connection. It doesn't replace the need to have redirects from your HTTP site.

## Enabling HSTS

To enable HSTS, add a new header (/guides/basic-configuration/adding-or-modifying-headers-on-http-requests-and-responses) as follows:

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.
4. Click the **Content** tab. The Content page appears.
5. Click the **Create header** button to create a new header. The Create a new header page appears.

### Create a new header

Learn more about this section in our [headers tutorial](#).

**CONDITION** This will happen all the time unless you [Attach a condition](#)

Type / Action Response Set

The type of header and the action performed on it.

Destination http.Strict-Transport-Security ★ Required

The name of the header that will be affected by the selected action. For example: `http.Content-Type`, `http.Set-Cookie`, `http.Via`, `http.Location`, or `http.Access-Control-Allow-Origin`.

Source "max-age=31536000" ★ Required

New content for the header. Can be a static value (e.g. string or number) or a dynamic value (e.g. existing header or a GeolIP value). Please use quotes for string values.

Ignore if set No

If switched to **Yes**, the action will not be performed if the header in Destination exists.

Priority 20 ★ Required

The order in which the header rules execute within the condition.

Description HSTS ★ Required

The name of your header, such as **My header**.

**CREATE** CANCEL

6. Fill out the **Create a new header** fields as follows:

- From the **Type** menu, select **Response**, and from the **Action** menu select **Set**.
- In the **Destination** field, type `http.Strict-Transport-Security`.
- In the **Source** field, type `"max-age=<max age in seconds>"`. For example, `"max-age=31536000"`. As described below, `max-age` is required and two additional HSTS options can be specified.
- Leave the **Ignore if set** menu and the **Priority** field set to their defaults (or set them as appropriate for your service).
- In the **Description** field, type a human-readable name, such as `HSTS`. This name is displayed in the Fastly web interface.

7. Click the **Create** button.
8. Click the **Activate** button to deploy your configuration changes.

## HSTS options

HSTS requires the max-age directive (<https://tools.ietf.org/html/rfc6797#section-6.1.1>) be set in order to function properly. It specifies how long in seconds to remember that the current domain should only be contacted over HTTPS. The example shown above sets `max-age` to one year (31536000 seconds = 1 year). You may want to experiment using a smaller value than what is shown.

Two additional options can be specified with the HSTS response header:

- `includeSubdomains` - This token applies HSTS to all of your site's subdomains. Before you include it, be certain none of your subdomains require functionality on HTTP in a browser. Ensure your TLS certificate is a wildcard or has coverage for all subdomain possibilities.

**❗ IMPORTANT:** All subdomains will be unreachable on HTTP by browsers that have seen the HSTS header once `includeSubdomains` is enabled.

- `preload` - This token allows you to submit your domain for inclusion in a preloaded HSTS list that is built into several major browsers. Although the token is not part of the HSTS specification, including it in the header is a prerequisite for submitting to this preloaded list.

**⚠ WARNING:** Don't request browser preload inclusion unless you're sure that you can support HTTPS for the long term. Inclusion in the HSTS Preload List cannot be undone easily. See <https://hstspreload.appspot.com/> (<https://hstspreload.appspot.com/>) for submission instructions and more information.

Combining all of these options together in the **Source** field would look like this:

```
"Strict-Transport-Security: max-age=<max age in seconds>; includeSubDomains; preload"
```

To disable HSTS for whatever reason, simply set the `max-age` to `0` on an HTTPS connection.

The HSTS Preload List is managed by a third party, not by Fastly. See <https://hstspreload.appspot.com/> (<https://hstspreload.appspot.com/>) for more information.

## Additional reading

- RFC 6797 (<http://tools.ietf.org/html/rfc6797>), which describes the HSTS specification
- the Wikipedia description ([https://en.wikipedia.org/wiki/HTTP\\_Strict\\_Transport\\_Security](https://en.wikipedia.org/wiki/HTTP_Strict_Transport_Security)) of HSTS, including the currently known limitations and a browser support list
- the OWASP.org explanation ([https://www.owasp.org/index.php/HTTP\\_Strict\\_Transport\\_Security](https://www.owasp.org/index.php/HTTP_Strict_Transport_Security)) of HSTS, including descriptions of the threats it addresses
- the Chromium Projects description (<https://www.chromium.org/hsts>) of HSTS and preloading HSTS sites

## § Ordering a paid TLS option (/guides/securing-communications/ordering-a-paid-tls-option)

The Fastly Transport Layer Security (TLS) service provides privacy and data security for your service. You can request different TLS service options for your account. For specific details about pricing for each of Fastly's TLS options, see our pricing page (<https://www.fastly.com/pricing>) or contact [sales@fastly.com](mailto:sales@fastly.com) (<mailto:sales@fastly.com>) for our current rates.

To order TLS service, contact customer support via email at [support@fastly.com](mailto:support@fastly.com) (<mailto:support@fastly.com>). Remember to put "TLS Certificate Request" in the subject of the request, and be sure to include your customer ID (/guides/account-management-and-security/finding-and-managing-your-account-info#finding-your-customer-id) when you let us know which service option you're interested in. We'll walk you through the setup process for the option you choose.

## Shared domain

This free option uses the Fastly SAN certificate's wildcard entry for `global.ssl.fastly.net`. To use this service, add a new domain in the Fastly web interface and set up an origin server for that domain. You can learn more about how to do that in our guide on setting up free TLS (/guides/securing-communications/setting-up-free-tls).

## Shared certificate

This service option uses the Fastly SAN certificate. You can read about SAN certificates here (<https://www.globalsign.com/en/ssl/multi-domain-ssl/>). Specifically:

- You get to use your domain, but Fastly does the certificate administration.
- You provide your domain name list to Fastly and we add those names to the Certificate SAN field.

Our registrar explains the shared (SAN) certificate (<https://www.globalsign.com/en/ssl/multi-domain-ssl/>) as "a way to conserve IP addresses by putting multiple hostnames or domains on one certificate. There are no security implications... Addition of your name to the certificate still needs to be authorized by you."

## Shared wildcard certificate service

This service option uses the Fastly SAN certificate. You can read about SAN certificates here (<https://www.globalsign.com/en/ssl/multi-domain-ssl/>). Specifically:

- You get to use your domain, but Fastly does the certificate administration.
- You provide one or more wildcard domain name entries to Fastly and we add those names to the Certificate SAN field.

## Customer certificate hosting

This service option uses your TLS certificates, which includes Extended Verification (EV) certificates. Specifically:

- You provide the certificates to Fastly.
- Fastly installs those certificates into the caches and allocates IP addresses on each cache.
- Fastly creates a new (customer-specific) DNS Global Domain Map that associates the certificate with the allocated IP addresses.
- Fastly maintains the domain map and the certificate on the caches.

**❗ IMPORTANT:** Fastly supports SHA-256 certificates signed by publicly trusted certificate authorities that have a minimum key size of 2048 bits for RSA or DSA, and 256 bits for ECDSA. For performance reasons, we strongly recommend using a 2048-bit key size for RSA and DSA when larger key sizes are not required for your application.

## SNI customer certificate hosting

Fastly provides, on a limited availability basis, customer certificate hosting that uses Server Name Indication (SNI) for certificate selection. To see if your company meets our SNI certificate hosting qualification criteria, contact [sales@fastly.com](mailto:sales@fastly.com) (<mailto:sales@fastly.com>). For SNI customer certificate hosting:

- You provide the certificates to Fastly.
- Fastly hosts your certificate on a shared domain map and uses SNI to select the correct certificate to serve based on the domain name sent.

**❗ IMPORTANT:** 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 certificate error.

**❗ NOTE:** Any certificates provided by GlobalSign are subject to the terms of GlobalSign's Subscriber Agreement, which can be found at <https://www.globalsign.com/repository/> (<https://www.globalsign.com/en/repository/>).

## § Setting up free TLS (/guides/securing-communications/setting-up-free-tls)

Customers can use our shared domain TLS wildcard certificate for free. Follow the steps below to set up free TLS:

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.
4. Click the **Domains** tab. The Domains page appears.
5. Click the **Create domain** button. The Create a new domain page appears.



## Create a new domain

**Domain Name**

The domain name of your website.

[▶ Setting the domain name](#)  
[▶ What if I'm using apex domains?](#)

**Description**

An optional comment that describes your domain.

★ Required

CREATE
CANCEL

6. Fill out the **Create a new domain** fields as follows:

- In the **Domain Name** field, type `[name].global.ssl.fastly.net`, where `[name]` is a single word. If the name has already been taken, you will need to pick a different one.
- In the **Description** field, type a human-readable name for the domain. This name is displayed in the Fastly web interface.

**❗ IMPORTANT:** In the **Domain Name** field, `[name]` can only be a single word. You cannot use a dot-separated name such as `www.example.org.global.ssl.fastly.net` because nesting in TLS certificates is not supported.

7. Click the **Create** button to save the domain. The new domain appears in the list of domains.

8. Click the **Activate** button to deploy your configuration changes.

You should be able to access the domain via the following URL: `https://<name>.global.ssl.fastly.net/`. Note that you don't need to add CNAME records (</guides/basic-setup/adding-cname-records/>) to use the shared domain certificate.

**❗ IMPORTANT:** If you DNS alias your own domain (`www.example.org`) to that name (`example.global.ssl.fastly.net`) a TLS name mismatch warning will appear in the browser. The only way to fix the mismatch is to add your domain to our TLS certificate. For more information, see our [TLS options guide \(/guides/securing-communications/ordering-a-paid-tls-option/\)](/guides/securing-communications/ordering-a-paid-tls-option/).

## § Support for App Transport Security (</guides/securing-communications/support-for-app-transport-security/>)

Apple uses App Transport Security

([https://developer.apple.com/library/ios/documentation/General/Reference/InfoPlistKeyReference/Articles/CocoaKeys.html#//apple\\_ref/doc/uid/TP40001SW35](https://developer.apple.com/library/ios/documentation/General/Reference/InfoPlistKeyReference/Articles/CocoaKeys.html#//apple_ref/doc/uid/TP40001SW35)) (ATS) to improve the security of connections between web services and applications installed on devices using iOS 9 or later, and OS X 10.11 (El Capitan) and later. Fastly is fully compliant with all ATS requirements. You shouldn't run into any issues supporting iOS or OS X users while using our service.

### Results from the ATS diagnostics tool

We used Apple's ATS diagnostics tool to ensure that Fastly is compliant with all ATS requirements. You can review the output from the diagnostics tool below.



```
$ /usr/bin/nscurl --ats-diagnostics https://www.fastly.com
Starting ATS Diagnostics

Configuring ATS Info.plist keys and displaying the result of HTTPS loads to https://www.fastly.com.
A test will "PASS" if URLSession:task:didCompleteWithError: returns a nil error.

Use '--verbose' to view the ATS dictionaries used and to display the error received in URLSession:task:didCompleteWithError:.
=====

Default ATS Secure Connection
---
ATS Default Connection

Result : PASS
---

=====
Allowing Arbitrary Loads

---
Allow All Loads

Result : PASS
---

=====

Configuring TLS exceptions for www.fastly.com

---
TLSv1.2

Result : PASS
---

---
TLSv1.1

Result : PASS
---

---
TLSv1.0

Result : PASS
---

=====

Configuring PFS exceptions for www.fastly.com

---
Disabling Perfect Forward Secrecy

Result : PASS
---

=====

Configuring PFS exceptions and allowing insecure HTTP for www.fastly.com

---
Disabling Perfect Forward Secrecy and Allowing Insecure HTTP

Result : PASS
---

=====

Configuring TLS exceptions with PFS disabled for www.fastly.com

---
TLSv1.2 with PFS disabled

Result : PASS
---

---
TLSv1.1 with PFS disabled

Result : PASS
```

```

---
---
TLSv1.0 with PFS disabled

Result : PASS
---

=====

Configuring TLS exceptions with PFS disabled and insecure HTTP allowed for www.fastly.com

---
TLSv1.2 with PFS disabled and insecure HTTP allowed

Result : PASS
---

---
TLSv1.1 with PFS disabled and insecure HTTP allowed

Result : PASS
---

---
TLSv1.0 with PFS disabled and insecure HTTP allowed

Result : PASS
---
```

## § TLS termination (/guides/securing-communications/tls-termination)

### Identifying TLS terminated requests

To maintain optimal caching performance, Fastly uses a TLS terminator ([https://en.wikipedia.org/wiki/TLS\\_termination\\_proxy](https://en.wikipedia.org/wiki/TLS_termination_proxy)) separate from the caching engine. This means, however, that the caching engine doesn't know that it was originally a TLS request. As a result, we set the `Fastly-SSL` header when fetching the content from your servers.

Because we set this header, you can check for its presence on your backend by doing something like:

```

if (req.http.Fastly-SSL) {
    set resp.http.X-Is-SSL = "yes";
}
```

and that should tell you if the request was a TLS request or not.

### When using WordPress

If you're using Fastly TLS services with WordPress, you'll want to add a check for the `HTTP_FASTLY_SSL` header so that WordPress can build URLs to your CSS or JS assets correctly. Do this by placing a check in your `wp-config.php` file to override the SSL flag that is checked later:

```

if (!empty( $_SERVER['HTTP_FASTLY_SSL'])) {
    $_SERVER['HTTPS'] = 'on';
}
```

As usual, this must be placed anywhere before the `require_once` line with `wp-settings.php`.

### Finding the original IP when using TLS termination

Because Fastly uses a TLS terminator ([https://en.wikipedia.org/wiki/TLS\\_termination\\_proxy](https://en.wikipedia.org/wiki/TLS_termination_proxy)), separate from the caching engine for performance, the engine overwrites the original IP briefly due to the re-request to your origin servers once decrypted and causes anything that references the original IP to show up as 127.0.0.0/8 IPs. To find the original IP via VCL:

- use `req.http.Fastly-Client-IP` if you're using shielding
- use `client.ip` you're not using shielding or if you're building an ACL

Fastly also sends along the client IP to the origin in a HTTP header, `Fastly-Client-IP`, which can be used by server software to adjust as needed.

For more information about TLS-related issues, see our TLS guides (/guides/securing-communications/) or contact [support@fastly.com](mailto:support@fastly.com) (mailto:support@fastly.com) with questions.

• Guides (/guides/) > Advanced setup (/guides/advanced) > Purging (/guides/purging/)

## § Authenticating URL purge requests via API (/guides/purging/authenticating-api-purge-requests)

Fastly's URL purge (/guides/purging/single-purges#purging-a-url) feature allows you to purge individual URLs on your website. By default, authentication is not required to purge a URL with the Fastly API, but you can enable API key authentication in the Fastly web interface by adding a header or by using custom VCL.

**NOTE:** All purge requests other than URL purges require authentication by default, as indicated in the [API documentation](#) (/api/purge#purge).

### Enabling authentication in the Fastly web interface

You can enable API key authentication for URL purge requests by adding a header and optionally attaching a condition in the Fastly web interface.

#### Adding the header

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.
4. Click the **Content** tab. The Content page appears.

The screenshot shows the Fastly web interface with the 'CONTENT' tab selected. The top navigation bar includes tabs for DOMAINS (1 Domain), ORIGINS (1 Host, 1 Health check), SETTINGS (1 Request setting, 0 Cache settings), CONTENT (0 Headers, 0 Gzips, 0 Responses), LOGGING (1 Endpoint), and CUSTOM VCL (0 Files). The main content area is divided into three sections: Headers, Gzip, and Responses. Each section has a 'There are no [headers/gzips/responses]' message and a button to 'CREATE YOUR FIRST [HEADER/GZIP/RESPONSE]'. To the right of the Headers section is an 'About content' sidebar with text explaining that headers allow setting, deleting, and changing request and response headers, and responses allow creating synthetic responses that exist entirely on the cache servers.

5. Click the **Create header** button. The Create a new header window appears.

## Create a new header

Learn more about this section in our [headers tutorial](#).

CONDITION

This will happen all the time unless you [Attach a condition](#)

Type / Action

Request

Set

The type of header and the action performed on it.

Destination

http.Fastly-Purge-Requires-Auth

★ Required

The name of the header that will be affected by the selected action. For example: `http.Content-Type`, `http.Set-Cookie`, `http.Via`, `http.Location`, or `http.Access-Control-Allow-Origin`.

Source

"1"

★ Required

New content for the header. Can be a static value (e.g. string or number) or a dynamic value (e.g. existing header or a GeolP value). Please use quotes for string values.

Ignore if set

No

If switched to **Yes**, the action will not be performed if the header in Destination exists.

Priority

10

★ Required

The order in which the header rules execute within the condition.

Description

Fastly Purge

★ Required

The name of your header, such as **My header**.

CREATE

CANCEL

6. Fill out the **Create a new header** fields as follows:

- From the **Type** menu, select **Request**, and from the **Action** menu, select **Set**.
- In the **Destination** field, type `http.Fastly-Purge-Requires-Auth`.
- In the **Source** field, type `"1"`.
- From the **Ignore if set** menu, select **No**.
- In the **Priority** field, type `10`.
- In the **Description** field, type the name of your header rule (for example, `Fastly Purge`).

7. Click the **Create** button.

## Attaching a condition

Attaching the following condition is optional. Without the condition, the header you just created will be added to all requests. With the condition, the header will be added to purge requests only.

1. On the Content page, click the **Attach a condition** link to the right of your new header. The Create a new request condition window appears.

Create a new request condition

Learn the basics in our [conditions tutorial](#)

Type

Request

Learn more about the [different types of conditions](#).

Name

Purge

\*

Apply if...

req.request == "FASTLYPURGE"

The expression to decide whether this is run.

[Examples](#)

[Adjust priority \(default is 10\)](#)

This is an advanced option. For most configurations, the default is best.

SAVE AND APPLY TO FASTLY PURGE

CANCEL

2. Fill out the **Create a new request condition** fields as follows:

- From the **Type** menu, select **Request**.
- In the **Name** field, type a descriptive name for the new condition (for example, `Purge`).
- In the **Apply if** field, type `req.request == "FASTLYPURGE"`.

3. Click the **Save and apply to** button.

4. Click the **Activate** button to deploy your configuration changes.

## Enabling authentication with custom VCL

If you'd rather enable API key authentication for URL purge requests using custom VCL (</guides/vcl/uploading-custom-vcl>), add the following to your VCL file:

```
if (req.request == "FASTLYPURGE") {
  set req.http.Fastly-Purge-Requires-Auth = "1";
}
```

**ⓘ IMPORTANT:** The ability to [upload custom VCL code](/guides/vcl/uploading-custom-vcl) (</guides/vcl/uploading-custom-vcl>) is disabled by default when you sign up. To enable this ability for your account, contact [support@fastly.com](mailto:support@fastly.com) (<mailto:support@fastly.com>) and request it.

## Purging URLs with an API key

After you've enabled API key authentication for URL purge requests, you'll need to provide your API key (</guides/account-management-and-security/finding-and-managing-your-account-info#finding-and-regenerating-your-api-key>) in the URL purge API request (`/api/purge#purge_3aa1d66ee81dbfed0b03deed0fa16a9a`):

```
curl -X PURGE -H Fastly-Key:$FASTLY_KEY https://www.example.com/
```

which would return this response:

```
{"status": "ok", "id": "1234567890"}
```

**⚠ WARNING:** If your website is not configured to use HTTPS, we recommend purging using a POST request with a secure Fastly API URL. This will ensure that your API key in the header is encrypted in transit. The request will look like this: `curl -X POST -H Fastly-Key:$FASTLY_KEY`

```
https://api.fastly.com/purge/<your_url_here>.
```

## § Getting started with surrogate keys (/guides/purging/getting-started-with-surrogate-keys)

Efficient cache invalidation is an essential part of keeping your website fast. Purging too much cache using `purge all` (/guides/purging/single-purges#purging-all-content) may increase your website's load time while the cache rebuilds. If you find yourself purging all cache on more than a weekly basis, consider using surrogate keys for more targeted purging.

Surrogate keys allow you to selectively purge related content. Using the `Surrogate-Key` header, you can tag a group of objects with a key and then use it to purge multiple pieces of content at once. This process can occur automatically within your application, making it easier to cache and purge content that changes rapidly and unpredictably.

This guide consists of four parts:

- Understanding surrogate keys
- Looking at a practical example
- Generating and setting surrogate keys
- Limitations

**NOTE:** This guide assumes you're already familiar with [the way content delivery networks \(CDNs\) work](/guides/about-fastly-services/how-caching-and-cdns-work) (/guides/about-fastly-services/how-caching-and-cdns-work) in general, and [the way Fastly's CDN works](/guides/about-fastly-services/how-fastlys-cdn-service-works) (/guides/about-fastly-services/how-fastlys-cdn-service-works) in particular.

### Understanding surrogate keys

After you've signed up for Fastly and added one or more services (/guides/basic-setup/working-with-services#creating-a-new-service), you can start examining how your origin server responds to requests. When your origin server responds to an HTTP request for content, it's because Fastly hasn't yet cached that content or the cache has expired. Your server's response to the request will resemble the example shown below. (Note that you can use the `curl` command (/guides/debugging/curl-and-other-caching-verification-methods) to inspect any of your server's responses.)

```
HTTP/1.1 200 OK
Content-Type: text/html
Connection: keep-alive
...
```

To control how your content is served to users and cached by Fastly, you can add to or modify the headers (/guides/basic-configuration/adding-or-modifying-headers-on-http-requests-and-responses) that are included in your origin server's response. The `Surrogate-Key` header is one of the headers that you can add to the response. It allows you to "tag" an object, such as an image or a blog post, with one or more keys. When the object changes, you can reference the key in a purge request (/guides/purging/single-purges) to remove the object from the cache.

You can add space-delimited strings to the `Surrogate-Key` header, like this:

```
HTTP/1.1 200 OK
Surrogate-Key: key1 key2 key3
Content-Type: text/html
...
```

This response contains three surrogate keys: `key1`, `key2`, and `key3`. When Fastly receives a response like this, we use the surrogate keys to create a mapping from each key to the cached content, then we strip out the `Surrogate-Key` header so it's not included in the response to your readers.

### Creating relationships between keys and objects

One of the major advantages of surrogate keys is that they allow for a many-to-many relationship between keys and objects. An origin server's response can associate multiple keys with the object, and the same key can be provided in different responses. Take a look at these two requests and responses:

```
GET /blog/ HTTP/1.1
Host: www.example.com

HTTP/1.1 200 OK Content-Type: text/html
Content-Length: 1234
Surrogate-Key: mainpage template-a
```

```
GET /blog/article/fastly-rocks HTTP/1.1
Host: www.example.com

HTTP/1.1 200 OK
Content-Type: text/html
Content-Length: 2345
Surrogate-Key: template-a article-fastly-rocks
```

In this example, there are two objects (`/blog` and `/blog/article/fastly-rocks`) with three keys (`mainpage`, `template-a`, and `article-fastly-rocks`). Two of the keys (`mainpage` and `article-fastly-rocks`) are associated with a single object, and a third key (`template-a`) is associated with both objects.

## Purging objects with surrogate keys

By using the `Surrogate-Key` header to associate keys with one or more objects, you can precisely control which objects are removed from cache during a purge. Consider the example presented above. Purging the `mainpage` key would remove only the `/blog` object from the cache. On the other hand, purging the `template-a` key would remove both the `/blog` and `/blog/article/fastly-rocks` objects from the cache.

You can use the Fastly web interface to manually purge objects via key (`/guides/purging/single-purges#purging-with-a-key`), or you can use our Purge API (`/api/purge`) to purge a collection of objects with one API call. If you're using Fastly to cache your API, check out the guide on purging API cache with surrogate keys (`/guides/api-caching/purging-api-cache-with-surrogate-keys`) to learn how surrogate keys can help you purge API cache.

## Looking at a practical example

Let's look at a practical example to learn how surrogate keys work. Imagine you're building a picture-hosting website and you're using Fastly to speed it up. You initially decide to cache picture pages by their URL (e.g., `http://www.example.com/pic/id`). When a picture's information changes, you'll remove the old version by sending a purge request to the Fastly API:

```
PURGE /pic/{id}
Host: example.com
Accept: */*
```

But there's a potential problem with this solution. You display a user's information next to their pictures, so you'll need to purge all of the user's picture pages if they change their information. You could send an individual purge for each picture, but that would take too long. As an alternative, you could cache the pages for a very short amount of time, but that would waste our server resources.

Surrogate keys solve this problem. By adding a surrogate key to all of a user's picture pages (e.g., `/user/542`, `/user/25`), you can purge all of the user's pictures by sending Fastly a purge for the user's surrogate key when they update their information. Now, instead of having to purge each picture individually, you can update them all with just one request:

```
PURGE /service/id/purge/user/542
...
```

## Purging multiple sites at the same time

What if you wanted to build a mobile version of the picture-hosting website to complement the desktop version? You'll need a way to purge both the desktop version and the mobile version at the same time. Surrogate keys can help in this instance. You can tag the different versions of a picture page with the same surrogate key (e.g., `pic/76`, `pic/345`) and purge them all at once. All of the related content on our sites can now be purged with one request.

## Tagging templates with surrogate keys

Surrogate keys come in really handy when making changes to templates. Imagine you have to make a change to the banner of the website. Since you're caching entire page, updating the header template isn't enough. You'll also need to purge all the pages that use the template. You could purge every page on the website, but there's no reason to purge content that doesn't use the header template.

You can make things easy by using surrogate keys. By adding surrogate keys for each template on a page (e.g., `/templates/pic/show`, `/templates/pic/header`, `/templates/pic/comment`), you can check which templates have changed and purge only pages with modified templates.

## Generating and setting surrogate keys

There are two ways to set the `Surrogate-Key` header: by adding the header in the Fastly web interface, or by generating the keys with your own application. We describe how to use the Fastly web interface in our guide to generating `Surrogate-Key` headers based on URLs (`/guides/purging/setting-surrogate-key-headers-based-on-a-url`) (we have a separate guide for Amazon S3 origins (`/guides/purging/setting-surrogate-key-headers-for-amazon-s3-origins`)).

Automatically generating keys with your own application is described below using Fastly's Test Blog (<https://www.fastly.com/blog/introducing-our-open-source-app-training>) as an example. The test blog is a Ruby on Rails application that comes preloaded with example content that can be cached and purged using Fastly via the `fastly-rails` (<https://github.com/fastly/fastly-rails>) Ruby gem. (We also have other API clients (`/api/clients`) that

support surrogate keys.)

**NOTE:** You can install the Fastly Test Blog and recreate this example yourself to practice generating surrogate keys. For the purposes of this example, we've made one deviation from the Fastly Test Blog instructions outlined on the [GitHub page \(https://github.com/fastly/fastly-test-blog\)](https://github.com/fastly/fastly-test-blog). We've added a CNAME record (</guides/basic-setup/adding-cname-records>) for our test blog to create two URLs — <http://origin.example.com>, which is the URL of the origin server not attached to our Fastly service, and <http://fastly.example.com>, which is the URL being cached by Fastly. Having the two URLs available will be useful when we examine the headers.

## Configuring the API client

The `fastly-rails` (<https://github.com/fastly/fastly-rails>) gem provides a `set_surrogate_key_header` method (<https://github.com/fastly/fastly-rails#headers>) which the test blog uses to automatically generate surrogate keys for new articles. You can see how this works by examining the code in `articles_controller.rb` ([https://github.com/fastly/fastly-test-blog/blob/master/app/controllers/articles\\_controller.rb](https://github.com/fastly/fastly-test-blog/blob/master/app/controllers/articles_controller.rb)). A slightly modified excerpt of the code from the controller is shown below.

```
class ArticlesController < ApplicationController
  # include this before_action in controller endpoints that you wish to edge cache
  # This can be used with any custom actions. Set these headers for GETs that you want to cache
  # e.g. before_action :set_cache_control_headers, only: [:index, :show, :my_custom_action]
  before_action :set_cache_control_headers, only: [:index, :show]

  # Returns all Article objects, and sets a table_key of 'articles',
  # and a record_key for each article object: "#{table_key}/#{article_id}"
  def index
    @articles = Article.all
    set_surrogate_key_header 'articles', @articles.map(&:record_key)
  end

  # Sets a surrogate key for the current article.
  #
  # Example:
  #
  # Article[75]
  # Surrogate-Key:articles/75
  def show
    set_surrogate_key_header @article.record_key
  end

  ...
end
```

The `before_action` method creates `Cache-Control` and `Surrogate-Control` HTTP headers with a default TTL of 30 days. This method must be added to any controller action that you want to edge cache. The `set_surrogate_key_header` method sets `Surrogate-Key` headers for objects that you want to be able to purge. In this case, surrogate keys are set for each article and the articles index.

## Examining the headers

Now that you've looked at how the surrogate keys are generated by the test blog, inspect the headers on an article page on the origin server. Here's the partial output from `curl -svo /dev/null origin.example.com/articles/1`:

```
HTTP/1.1 200 OK
X-Frame-Options: SAMEORIGIN
X-Xss-Protection: 1; mode=block
X-Content-Type-Options: nosniff
Cache-Control: public, no-cache
Surrogate-Control: max-age=86400
Surrogate-Key: articles/1
```

Notice how the `Cache-Control`, `Surrogate-Control`, and `Surrogate-Key` headers are present in the response shown above. Thanks to the `set_surrogate_key_header` method, the test blog application automatically generates the unique `articles/1` surrogate key for this article.

Next, inspect the headers on the article index page on the origin server URL. Because this page lists all of the articles, you might expect it to contain the surrogate keys for every article displayed on the page, and that is indeed the case. Here's the partial output from `curl -svo /dev/null origin.example.com`:



```
HTTP/1.1 200 OK
X-Frame-Options: SAMEORIGIN
X-Xss-Protection: 1; mode=block
X-Content-Type-Options: nosniff
Cache-Control: public, no-cache
Surrogate-Control: max-age=86400
Surrogate-Key: articles articles/1 articles/2 articles/3 articles/4 articles/5 articles/6 articles/7 articles/8 articles/9 articles/
10 articles/11 articles/12 articles/15 articles/16 articles/17 articles/18 articles/19 articles/20 articles/27 articles/28 articles/
29 articles/30 articles/31
```

If you purged the `articles/1` surrogate key, both `http://origin.example.com/articles/1` and `http://origin.example.com` would be purged from the cache.

Finally, take a look at the URL that's piped through Fastly: `http://fastly.example.com`. The surrogate keys won't be visible in the headers, but Fastly knows what they are. Recall from understanding surrogate keys that Fastly strips out the `Surrogate-Header` and creates a mapping from each key to the cached content. Here's the partial output from `curl -svo /dev/null fastly.example.com`:

```
HTTP/1.1 200 OK
X-Frame-Options: SAMEORIGIN
X-Xss-Protection: 1; mode=block
X-Content-Type-Options: nosniff
Cache-Control: public, no-cache
Content-Type: text/html; charset=utf-8
Via: 1.1 varnish
Age: 78309
X-Served-By: cache-iad2120-IAD
X-Cache: HIT
X-Cache-Hits: 1
X-Timer: S1449255701.272992,VS0,VE5
```

With the `Surrogate-Key` header present on the index and article pages, you're now able to manually purge blog pages via key (`/guides/purging/single-purges#purging-with-a-key`), or purge a collection of pages with one API call (`/api/purge`).

## Limitations

The surrogate keys sent by your origin server can be as simple or complex as you need, but there are a couple limitations:

- Surrogate keys are limited to a total length of 1,024 bytes each
- Any keys that exceed the limit will be dropped instead of truncated
- Surrogate key headers are limited to a total length of 16,384 bytes
- Any keys past the one that exceeds the limit will be dropped

## § Setting Surrogate-Key headers based on a URL (`/guides/purging/setting-surrogate-key-headers-based-on-a-url`)

You can mark content with a surrogate key (`/guides/purging/getting-started-with-surrogate-keys`) and use it to purge groups of specific URLs (`/guides/purging/single-purges#purging-with-a-key`) at once without purging everything (`/guides/purging/single-purges#purging-all-content`), or purging each URL (`/guides/purging/single-purges#purging-a-url`) singularly.

Follow these instructions to set Surrogate-Key headers based on a URL:

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.
4. Click the **Content** tab. The Content page appears.

DOMAINS  
1 Domain

ORIGINS  
1 Host  
1 Health check

SETTINGS  
1 Request setting  
0 Cache settings

CONTENT  
0 Headers  
0 Gzips  
0 Responses

LOGGING  
1 Endpoint

CUSTOM VCL  
0 Files

Headers

There are no headers.

CREATE YOUR FIRST HEADER

Gzip

Gzip is not enabled.

SET UP GZIP

Responses

There are no responses.

CREATE YOUR FIRST RESPONSE

About content

Headers allow you to set, delete, and change request and response headers as your information is handled by Fastly.

Responses allow you to create synthetic responses that exist entirely on the cache servers.

5. Click the **Create header** button. The Create a new header page appears.

## Create a new header

Learn more about this section in our [headers tutorial](#).

CONDITION

This will happen all the time unless you [Attach a condition](#)

Type / Action

Cache

Set

The type of header and the action performed on it.

Destination

http.Surrogate-Key

★ Required

The name of the header that will be affected by the selected action. For example: `http.Content-Type`, `http.Set-Cookie`, `http.Via`, `http.Location`, or `http.Access-Control-Allow-Origin`.

Source

regex(req.url, "^/(.\*)\\.\\.\\.\\1\$")

★ Required

New content for the header. Can be a static value (e.g. string or number) or a dynamic value (e.g. existing header or a GeoIP value). Please use quotes for string values.

Ignore if set

No

If switched to **Yes**, the action will not be performed if the header in Destination exists.

Priority

10

★ Required

The order in which the header rules execute within the condition.

Description

URL Surrogate Key

★ Required

The name of your header, such as **My header**.

CREATE

CANCEL

6. Fill out the **Create a new header** fields as follows:

- From the **Type** menu, select **Cache**, and from the **Action** menu, select **Set**.
- In the **Destination** field, type `http.Surrogate-Key`.
- In the **Source** field, type `regex(req.url, "^/(.*)\\.\\.\\.\\1$")`. This will accept a URL that looks like `/foo.html` and will create the Surrogate-Key `foo`.
- From the **Ignore if set** menu, select **No**.
- In the **Priority** field, type `10`.
- In the **Description** field, type a human-readable name for the header. This name is displayed in the Fastly web interface.

7. Click the **Create** button to create your header.

8. Click the **Activate** button to deploy your configuration changes.

**NOTE:** There are several limitations to surrogate keys. See the [surrogate key limitations \(/guides/purging/getting-started-with-surrogate-keys#limitations\)](#) section for more information.

## § Setting Surrogate-Key headers for Amazon S3 origins (/guides/purging/setting-surrogate-key-headers-for-amazon-s3-origins)

You can mark content with a surrogate key (/guides/purging/getting-started-with-surrogate-keys) and use it to purge groups of specific URLs (/guides/purging/single-purges#purging-with-a-key) at once without purging everything (/guides/purging/single-purges#purging-all-content), or purging each URL (/guides/purging/single-purges#purging-a-url) singularly. On the Amazon S3 side, you can use the `x-amz-meta-surrogate-key` header to mark your content as you see fit, and then on the Fastly side set up a Header configuration to translate the S3 information into the header we look for.

**❗ IMPORTANT:** Pay close attention to the capitalization. Amazon S3 only accepts all lowercase header names.

Follow these instructions to set Surrogate-Key headers for Amazon S3 origin servers:

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.
4. Click the **Content** tab. The Content page appears.

The screenshot shows the Fastly Content configuration page. The top navigation bar includes tabs for DOMAINS, ORIGINS, SETTINGS, CONTENT, LOGGING, and CUSTOM VCL. The CONTENT tab is selected, showing 0 Headers, 0 Gzips, and 0 Responses. The main content area is divided into three sections: Headers, Gzip, and Responses. Each section has a message indicating that there are no headers, gzip is not enabled, or there are no responses, and a button to create the first header, set up gzip, or create the first response. On the right, there is an 'About content' section explaining that headers allow setting, deleting, and changing request and response headers, and responses allow creating synthetic responses that exist entirely on the cache servers.

5. Click the **Create header** button. The Create a new header page appears.

## Create a new header

Learn more about this section in our [headers tutorial](#).

CONDITION

This will happen all the time unless you [Attach a condition](#)

Type / Action

The type of header and the action performed on it.

Destination  ★ Required

The name of the header that will be affected by the selected action. For example: **http.Content-Type**, **http.Set-Cookie**, **http.Via**, **http.Location**, or **http.Access-Control-Allow-Origin**.

Source  ★ Required

New content for the header. Can be a static value (e.g. string or number) or a dynamic value (e.g. existing header or a GeoIP value). Please use quotes for string values.

Ignore if set

If switched to **Yes**, the action will not be performed if the header in Destination exists.

Priority  ★ Required

The order in which the header rules execute within the condition.

Description  ★ Required

The name of your header, such as **My header**.

CREATE

CANCEL

6. Fill out the **Create a new header** fields as follows:

- From the **Type** menu, select **Cache**, and from the **Action** menu, select **Set**.
- In the **Destination** field, type `http.Surrogate-Key`.
- In the **Source** field, type `beresp.http.x-amz-meta-surrogate-key`.
- From the **Ignore if set** menu, select **No**.
- In the **Priority** field, type `10`.
- In the **Description** field, type a human-readable name for the header. This name is displayed in the Fastly web interface.

7. Click the **Create** button to create your header.

8. Click the **Activate** button to deploy your configuration changes.

**NOTE:** There are several limitations to surrogate keys. See the [surrogate key limitations \(/guides/purging/getting-started-with-surrogate-keys#limitations\)](#) section for more information.

## § Single purges (/guides/purging/single-purges)

Fastly provides several levels of cache purging. You can purge something as small as a single URL via the "Purge URL" command or as large as all content under a service via the "Purge All" command. You can also selectively purge content via key-based purging using the "Purge Key" command. We also provide a purging feature called Soft Purge that allows you to mark content as outdated (stale) instead of permanently deleting it

from Fastly's caches.

**TIP:** To mark content as outdated instead of permanently deleting it, check out our [Soft Purge \(/guides/purging/soft-purges\)](/guides/purging/soft-purges) feature. You may also be interested in our [wildcard purging \(/guides/purging/wildcard-purges\)](/guides/purging/wildcard-purges).

## Purging via the user interface

To purge content using the Fastly web interface, choose one of the purging methods below.

### Purging a URL

To purge a single URL, follow the steps below:

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. From the **Purge** menu, select **Purge URL**. The Purge URL window appears.

**Purge URL**  
Purging tutorial

Domain

Subdomain  ★ Required

Full URL path

▶ [Learn more about purging HTTP](#)

Preview

Password  ★ Required

**PURGE** CANCEL

4. From the **Domain** menu, select the domain (</guides/about-fastly-services/domain-names-and-fastly-services>) on which your content resides. If the domain you select is a wildcard domain (e.g., `*.example.com`) the Subdomain field will appear.
5. If the **Subdomain** field appears, type the subdomain to purge for the wildcard domain you've selected (e.g., `www`).
6. In the **Full URL path** field, type the path to the content you'll be purging (e.g., `/example.jpg`). The Preview field displays the URL that will be purged.
7. In the **Password** field, type your Fastly password.
8. Click the **Purge** button.

### Purging with a key

Key-based purging of cache information allows you to include surrogate key headers (</guides/purging/getting-started-with-surrogate-keys>) with your content, which our caches strip out and index. You can then issue purge commands to selectively flush out anything with a specific surrogate key. For extra usefulness, items can have multiple keys that allow you to set up your own purging hierarchy.

**IMPORTANT:** To enable key purging, your server must also send a special response header along with the content being cached. Have your server send the `Surrogate-Key` header with a list of keys for this URL. If you're using multiple keys, separate them with spaces. For example, to tag a response with both the keys `foo` and `bar` you would send this header: `Surrogate-Key: foo bar`

To purge content tagged with one or more keys, follow the steps below:

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. From the **Purge** menu, select **Purge Key**. The Purge Key window appears.

4. In the **Keys** field, type one or more keys (separate multiple keys with spaces).
5. In the **Password** field, type your Fastly password.
6. Click the **Purge** button.

## Purging all content

**⚠ WARNING:** Do not purge all cached content if you are seeing [503 errors \(/guides/debugging/common-503-errors\)](/guides/debugging/common-503-errors). Purge all overrides [stale-if-error \(/guides/performance-tuning/serving-stale-content\)](/guides/performance-tuning/serving-stale-content) and increases the requests to your origin server, which could result in additional 503 errors.

To instantly purge all content under your service, follow the steps below:

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. From the **Purge** menu, select **Purge All**. The Purge All window appears and displays the exact domain you'll be purging.

4. In the **Password** field, type your Fastly password.
5. Click the **Purge** button.

## Purging via API

The syntax for purging a service through the API can be found in the Purging section (</api/purge>) of the API (</api/>) documentation.

## § Soft purges (/guides/purging/soft-purges)

Fastly provides a Soft Purge feature that allows you to mark content as outdated (stale) instead of permanently purging and thereby deleting it from Fastly's caches. Objects invalidated with Soft Purge will be treated as outdated (stale) while Fastly fetches a new version from origin. You can purge by URL or by surrogate key using Soft Purge.

Before using Soft Purge, we recommend you implement one of the following revalidation methods:

- Configure `stale_while_revalidate` to serve stale content (/guides/performance-tuning/serving-stale-content) and fetch the newest version of the object from origin in the background. If you choose this revalidation method, consider also configuring `stale_if_error` at the same time.
- Set up `ETag` or `Last-Modified` headers for relevant content on your origin servers.

To implement Soft Purge, add a `Fastly-Soft-Purge` request header (such as `Fastly-Soft-Purge: 1`) to any single URL or key-based purge. For example, to purge the URL `www.example.com` with Soft Purge, you would issue the following command:

```
curl -X PURGE -H "Fastly-Soft-Purge:1" http://www.example.com
```

## § Wildcard purges (/guides/purging/wildcard-purges)

Wildcard purging allows you to flush the cache of all pages under a directory branch or URL path; for example, you want to empty the cache of all pages under your `"/service"` path. Having to purge each URL (/guides/purging/single-purges#purging-a-url) one by one using the Fastly API (/api/purge) or via the Fastly app is not very efficient.

Although Fastly does not have a specific wildcard purge function, you can implement the same behavior by making a small configuration change using surrogate keys (<https://www.fastly.com/blog/surrogate-keys-part-1>). Surrogate keys allow you to tag a group of objects with a keyword (key) and then purge multiple pieces of content at once with it via the web interface or via custom VCL.

**❗ IMPORTANT:** Purging will only apply to new objects as they're being put into the cache after you set up configuration changes. It will not apply to objects already in the cache when this configuration is being applied.

To purge content based on wildcard paths, follow the steps below.

### Via the web interface

To purge content based on wildcard paths via the web interface, follow the steps below.

#### Create a default wildcard header

We set a default wildcard so that we have the flexibility to append other surrogate keys to a URL path.

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.
4. Click the **Content** tab. The Content page appears.



DOMAINS  
1 Domain

ORIGINS  
1 Host  
1 Health check

SETTINGS  
1 Request setting  
0 Cache settings

CONTENT  
0 Headers  
0 Gzips  
0 Responses

LOGGING  
1 Endpoint

CUSTOM VCL  
0 Files

Headers

There are no headers.

CREATE YOUR FIRST HEADER

Gzip

Gzip is not enabled.

SET UP GZIP

Responses

There are no responses.

CREATE YOUR FIRST RESPONSE

About content

Headers allow you to set, delete, and change request and response headers as your information is handled by Fastly.

Responses allow you to create synthetic responses that exist entirely on the cache servers.

5. Click the **Create header** button. The Create a new header page appears.

## Create a new header

Learn more about this section in our [headers tutorial](#).

### CONDITION

This will happen all the time unless you [Attach a condition](#)

Type / Action

Cache

Set

The type of header and the action performed on it.

Destination

http.Surrogate-Key

★ Required

The name of the header that will be affected by the selected action. For example: **http.Content-Type**, **http.Set-Cookie**, **http.Via**, **http.Location**, or **http.Access-Control-Allow-Origin**.

Source

""

★ Required

New content for the header. Can be a static value (e.g. string or number) or a dynamic value (e.g. existing header or a GeoIP value). Please use quotes for string values.

Ignore if set

No

If switched to **Yes**, the action will not be performed if the header in Destination exists.

Priority

10

★ Required

The order in which the header rules execute within the condition.

Description

Default Wildcard

★ Required

The name of your header, such as **My header**.

CREATE

CANCEL

6. Fill out the **Create a new header** fields as follows:

- From the **Type** menu, select **Cache** and from the **Action** menu, select **Set**.
- In the **Destination** field, type `http.Surrogate-Key`.
- In the **Source** field, type `""`.
- From the **Ignore if set** menu, select **No**.
- In the **Priority** field, type `10`.
- In the **Description** field, type `Default Wildcard`. This name is displayed in the Fastly web interface.

7. Click the **Create** button. A new header appears in the Headers area of the Content page.

## Create headers for each wildcard path being purged

Next, create a header for each of the wildcard paths you need the ability to purge. For instance, you want to purge the wildcard path `/*/*foo`.

1. Click the **Create header** button to create another new header.

## Create a new header

Learn more about this section in our [headers tutorial](#).

CONDITION

This will happen all the time unless you [Attach a condition](#)

Type / Action

Cache

Append

The type of header and the action performed on it.

Destination

http.Surrogate-Key

★ Required

The name of the header that will be affected by the selected action. For example: **http.Content-Type**, **http.Set-Cookie**, **http.Via**, **http.Location**, or **http.Access-Control-Allow-Origin**.

Source

"\*/foo"

★ Required

New content for the header. Can be a static value (e.g. string or number) or a dynamic value (e.g. existing header or a GeolIP value). Please use quotes for string values.

Ignore if set

No

If switched to **Yes**, the action will not be performed if the header in Destination exists.

Priority

20

★ Required

The order in which the header rules execute within the condition.

Description

/\* /foo Wildcard

★ Required

The name of your header, such as **My header**.

CREATE

CANCEL

2. Fill out the **Create a new header** fields as follows:

- From the **Type** menu, select **Cache**, and from the **Action** menu, select **Append**.
- In the **Destination** field, type `http.Surrogate-Key`.
- In the **Source** field, type `" */foo"`. There is a space before the asterisk in the Source field, which is important when appending multiple surrogate keys to a URL.
- From the **Ignore if set** menu, select **No**.
- In the **Priority** field, type `20`.
- In the **Description** field, type `/* /foo Wildcard`. This name is displayed in the Fastly web interface.

3. Click the **Create** button. A new header appears in the Headers area of the Content page.

Notice the Action is set to Append to add to the default wildcard surrogate key. The Priority is set to 20 so that the Default Wildcard header is executed first and then the wildcard path appends.

## Create conditions for each wildcard path being purged

Finally, create a condition for each of the wildcard paths you need the ability to purge.

- Click the **Attach a condition** link next to the wildcard path header name. The Create a new cache condition window appears.

Create a new cache condition

Learn the basics in our [conditions tutorial](#)

Type

Cache

Learn more about the [different types of conditions](#).

Name

/^ /foo Wildcard Condition

\*

Apply if...

req.url ~ "^[^/]\*foo\$"

The expression to decide whether this is run.

[Examples](#)

[Adjust priority \(default is 10\)](#)

This is an advanced option. For most configurations, the default is best.

SAVE AND APPLY TO /^ /FOO WILDCARD

CANCEL

2. Fill out the **Create a new cache condition** fields as follows:

- In the **Name** field, type `/^ /foo Wildcard Condition`.
- In the **Apply if** field, type `req.url ~ "^[^/]*foo$"`.

3. Click the **Save and apply to** button to create the new condition.

What does the condition mean? In the Apply if field above, the first `"^"` and `"$"` tells Fastly to look for the following pattern:

- Start from the first slash after the request host header.
- There should be one directory.
- It should be followed by the path `/foo` ending the URL.

Some examples would be `/a/foo`, `/bar/foo`, and `/c/foo`. You could also remove the first `"^"` and `">"$"` to allow the condition to be more general so that the pattern can occur in the middle of a URL path.


Some other examples for URL wildcard conditions:

Apply if field	Matched pattern
<code>req.url ~ "[^/]*foo"</code>	<code>/delta/wow/a/foo/neat/cool/img.gif</code>
<code>req.url ~ "^.*/foo\$"</code>	<code>/a/b/c/d/e/f/foo</code>

## Purge the wildcard

Ready to purge that wildcard? You can do this through the UI using the steps below.


1. Log in to the Fastly web interface and click the **Configure** link.
2. From the **Purge** menu, select **Purge Key**.



PURGE ▾	
Purge URL	
Purge Key	
Purge All	
Hits	391,508
Misses	93,908
Errors	90

The Purge Key window appears.

3. In the **Keys** field, type the surrogate key you want to purge. Continuing with our example, you would type `*/foo` without the quotes that were entered in the Source field of the New Header window above.



Purge Key [Purging tutorial](#)

Keys  ★ Required

Enter the keys (tags). Separate multiple with spaces.

Password  ★ Required

PURGE CANCEL

4. In the **Password** field, type your Fastly password.
5. Click the **Purge** button.

## Via custom VCL

To purge content based on wildcard paths via custom VCL, follow the steps below.

1. Add the following code to the VCL template:

```

sub construct_skey {
  if (req.url ~ "^(^/?)*([/?]*)(^/?)*([/?]*)(^/?)*" && !std.strstr(beresp.http.Surrogate-Key,re.group.1)) {
    if (!re.group.1) {
      return;
    }
    set req.http.skey = re.group.1;
    set beresp.http.Surrogate-Key = if(beresp.http.Surrogate-Key, beresp.http.Surrogate-Key " ", "") req.http.skey;

    if (!re.group.2) {
      return;
    }
    set req.http.skey = req.http.skey re.group.2;
    set beresp.http.Surrogate-Key = beresp.http.Surrogate-Key " " req.http.skey;

    if (!re.group.3) {
      return;
    }
    set req.http.skey = req.http.skey re.group.3;
    set beresp.http.Surrogate-Key = beresp.http.Surrogate-Key " " req.http.skey;

    if (!re.group.4) {
      return;
    }
    set req.http.skey = req.http.skey re.group.4;
    set beresp.http.Surrogate-Key = beresp.http.Surrogate-Key " " req.http.skey;

    if (!re.group.5) {
      return;
    }
    set req.http.skey = req.http.skey re.group.5;
    set beresp.http.Surrogate-Key = beresp.http.Surrogate-Key " " req.http.skey;
  }
}

```

1. Call the subroutine in `vcl_fetch`:

```

sub vcl_fetch {
  call construct_skey;
}

```

1. Check your success by curling an object not already in cache with the `Fastly-Debug:1` header to expose the surrogate keys. For example:

```

$ curl -svo /dev/null http://www.example.com/test/test2/file3.txt -HFastly-Debug:1
* Trying 192.0.2.0...
* Connected to www.example.com (192.0.2.0) port 80 (#0)
> Host: www.example.com
> User-Agent: curl/7.43.0
> Accept: */*
> Fastly-Debug:1
>
< HTTP/1.1 200 OK
< Server: Apache
< Content-Type: text/plain
< Surrogate-Key: /test /test/test2 /test/test2/file3.txt
< Via: 1.1 varnish
< X-Backend-IP: 203.0.113.0
< Cache-Control: max-age=31536000, stale-while-revalidate=31536000, stale-if-error=31536000
< Content-Length: 19
< Accept-Ranges: bytes
< Date: Fri, 29 Jan 2016 21:30:08 GMT
< Via: 1.1 varnish
< Age: 1035
< Connection: keep-alive
< Fastly-Debug-Path: (D cache-sjc3123-SJC 1454103008) (F cache-sjc3134-SJC 1454101973) (D cache-den6026-DEN 1454101973) (F cache-den6027-DEN 1454101973)
< Fastly-Debug-TTL: (H cache-sjc3123-SJC - - 1035) (M cache-den6026-DEN - - 0)
< Fastly-Debug-Digest: b43bd38cf940e1669c2927c8662660e5170758053dda42e772ce3fc34ee57fc1
< X-Served-By: cache-den6026-DEN, cache-sjc3123-SJC
< X-Cache: MISS, HIT
< X-Cache-Hits: 0, 1
< Vary: Accept-Encoding
<
{ [19 bytes data]
* Connection #0 to host www.example.com left intact

```

In the above example, the `< Surrogate-Key: /test /test/test2 /test/test2/file3.txt` headers show the addition of the three surrogate keys.

## Via the API

You can also use our key-based purging via the API to perform wildcard purging using an HTTP request:

```
POST /service/<Fastly Service ID>/purge/*/foo
Fastly-Key: <Fastly Key>
```

This will purge any content that was associated with the `"*/foo"` surrogate key according to the setup in your header rules. Additional syntax for purging a service through the API can be found in the Purging section (</api/purge/>) of the API (</api/>) documentation.

• [Guides \(/guides/\)](/guides/) > [Developer's tools \(/guides/devtools/\)](/guides/devtools/) > [Access Control Lists \(/guides/access-control-lists/\)](/guides/access-control-lists/)

## § About Edge ACLs (/guides/access-control-lists/about-edge-acls)

Edge ACLs allow you to attach an access control list (ACL) to a service with versionless ACL entries that are stored separately from your VCL configuration. You can use the Fastly API to programmatically add, remove, and update ACLs and their entries.

### How Edge ACLs work

Like Edge Dictionaries (</guides/edge-dictionaries/about-edge-dictionaries/>), Edge ACLs have two major components: The ACL itself (</guides/access-control-lists/creating-and-using-edge-acls/>), and the ACL entries (</guides/access-control-lists/creating-and-manipulating-edge-acl-entries/>) within it. ACLs are containers for ACL entries, which store IP addresses that can be used to whitelist or blacklist access to resources. Every ACL is attached to a version of a service, but ACL entries are versionless and any changes will become effective immediately.

### When ACLs might be useful

- E-commerce companies preventing scraping from certain IP ranges
- Offices restricting access to their administrative portals
- Advertising technology companies blocking bad-actors at edge
- Mobile applications accepting only calls from specific proxies or IP ranges
- System administrators restricting access to groups of backends from an office IP address or subnet range

### Getting started

To create an ACL and use it within your service, you'll need to perform the following steps:

1. Create an empty ACL (</guides/access-control-lists/creating-and-using-edge-acls/>) in a working version of a service that's unlocked and not yet activated.
2. Activate the new version of the service (</guides/basic-setup/working-with-services#service-versions-and-their-activation/>) you associated with the empty ACL.
3. Add ACL entries (</guides/access-control-lists/creating-and-manipulating-edge-acl-entries/>) to the newly created ACL.

Once the ACL is created, properly associated, and filled with ACL entries, it can be called in your service.

### Putting Edge ACLs to work

After you've used the Fastly API to create an ACL and add ACL entries, the VCL for the ACLs and ACL entries will be automatically generated, as shown below. In this example, the name of the ACL is `office_ip_ranges`.

```
# This VCL is automatically generated when you add an Edge ACL and entries with the Fastly API
# In this example, the ACL name is "office_ip_ranges"
acl office_ip_ranges {
  "10.1.1.0"/16;                # internal office
  "8.19.222.194";              # remote VPN office
  "FE80:0000:0000:0000:0202:B3FF:FE1E:8329"; # ipv6 address remote
}
```

You can upload custom VCL (</guides/vcl/uploading-custom-vcl/>) to add logic to interact with ACLs. In this example, the `office_ip_ranges` ACL is used as a whitelist. All access to `/admin` is denied by default, but IP addresses in the `office_ip_ranges` ACL are allowed to access `/admin` without restriction.

```
sub vcl_recv {
  # block all requests to Admin pages from IP addresses not in office_ip_ranges
  if (req.url ~ "^/admin" && ! (client.ip ~ office_ip_ranges)) {
    error 403 "Forbidden";
  }
}
```

ACL entries have a boolean option for negation which allows you to specify whether or not the IP address is whitelisted or blacklisted. You can set the option to true (1) to blacklist (deny), or false (0) to whitelist (allow). We *do not* recommend mixing both negated and non-negated entries in the same ACL. Doing so might have the opposite effect depending on the condition you specify in the VCL.

## Limitations

When creating Edge ACLs (/guides/access-control-lists/creating-and-using-edge-acls), keep the following limitations in mind as you develop your service configurations:

- **ACLs created with custom VCL cannot be manipulated using the API.** If you create an ACL using the API, you can use the API to make changes to it, and use custom VCL to interact with it. If you create an ACL using custom VCL (/guides/access-control-lists/manually-creating-access-control-lists), that ACL must always be manipulated via custom VCL.
- **Custom VCL manipulations are always versioned.** This is true for both ACLs created using custom VCL and for any logic created to interact with ACLs.
- **ACLs are limited to 1000 ACL entries.** If you find your ACLs approaching this limit, you may want to review this with your account manager. We may be able to help you figure out an even more efficient way to do things with your ACLs.
- **Audit logs don't exist for ACL changes.** If you use the API to add, update, or remove an ACL entry, there will be no record of it. The only record of a change will exist when you compare service versions (<https://www.fastly.com/blog/introducing-version-diff>) to view the point at which the ACL was associated with the service version in the first place.
- **When you delete an ACL, you'll only delete it from the service version you're editing.** ACLs are tied to versions and can be cloned and reverted. When using ACLs, we want you to be able to do things like delete an ACL from a current version of your service in order to roll back your configuration to a previous version using as few steps as possible.
- **When you delete an ACL, we remove the ACL entries inside it.** When you delete an ACL we remove the entries that live inside that container, but they are not deleted from the ACL in earlier versions of your configuration. You can roll back to a previous version of a service to restore the ACL and the ACL entries.
- **ACL entry deletions are permanent.** If you delete an ACL entry, the entry is permanently removed from all service versions and cannot be recovered.

## § Creating and manipulating Edge ACL entries (/guides/access-control-lists/creating-and-manipulating-edge-acl-entries)

Edge ACL (/guides/access-control-lists/about-edge-acls) entries contain IP addresses that can be used to whitelist or blacklist access to resources. ACL entries are versionless, and any changes will become effective immediately, regardless of version.

## Parameters

Edge ACL entries have the following parameters:

- **service\_id:** The ID of the Fastly service the ACL is associated with.
- **acl\_id:** The ID of the ACL.
- **id:** The ID of the ACL entry.
- **ip:** The IP address contained within the ACL entry.
- **subnet:** Optional. The range of IP addresses within a single ACL entry.
- **negated:** Defaults to false. A number indicating true (1) to blacklist (deny), or false (0) to whitelist (allow). Note that we *do not* recommend mixing both negated and non-negated entries in the same ACL.
- **comment:** Optional. A descriptive comment indicating why you created the ACL entry.

## Creating an ACL entry

To add an entry to an existing ACL, make the following API call in a terminal application:



```
curl -H "Fastly-Key: FASTLY_KEY" -X POST https://api.fastly.com/service/<service_id>/acl/<acl_id>/entry -d 'ip=127.0.0.1&subnet=16&negated=1&comment=do not allow'
```

The response will look like this:

```
{
  "acl_id": "<acl_id>",
  "comment": "do not allow",
  "created_at": "2016-04-22T19:14:02+00:00",
  "deleted_at": null,
  "id": "<acl_entry_id>",
  "ip": "127.0.0.1",
  "negated": "1",
  "service_id": "<service_id>",
  "subnet": 16,
  "updated_at": "2016-04-22T19:14:02+00:00"
}
```

## Viewing ACL entries

To see information related to a single ACL entry, make the following API call in a terminal application:

```
curl -H "Fastly-Key: FASTLY_KEY" -H 'Content-Type: application/vnd.api+json' https://api.fastly.com/service/<service_id>/acl/<acl_id>/entry/<acl_entry_id>
```

The response will look like this:

```
{
  "acl_id": "<acl_id>",
  "comment": "do not allow",
  "created_at": "2016-04-22T19:18:42+00:00",
  "deleted_at": null,
  "id": "<acl_entry_id>",
  "ip": "127.0.0.5",
  "negated": "1",
  "service_id": "<service_id>",
  "subnet": 16,
  "updated_at": "2016-04-22T19:18:42+00:00"
}
```

To view a list of all ACL entries attached to a particular ACL, make the following API call in a terminal application:

```
curl -H "Fastly-Key: FASTLY_KEY" https://api.fastly.com/service/<service_id>/acl/<acl_id>/entries
```

The response will look like this:

```
[
  {
    "acl_id": "<acl_id>",
    "comment": "",
    "created_at": "2016-04-22T19:13:03+00:00",
    "deleted_at": null,
    "id": "<acl_entry_1_id>",
    "ip": "127.0.0.1",
    "negated": "0",
    "service_id": "<service_id>",
    "subnet": 16,
    "updated_at": "2016-04-22T19:13:03+00:00"
  },
  {
    "acl_id": "<acl_id>",
    "comment": "",
    "created_at": "2016-04-22T19:14:02+00:00",
    "deleted_at": null,
    "id": "<acl_entry_2_id>",
    "ip": "127.0.0.2",
    "negated": "1",
    "service_id": "<service_id>",
    "subnet": 16,
    "updated_at": "2016-04-22T19:14:02+00:00"
  }
]
```

## Updating ACL entries

There are two ways to update ACL entries: you can update a single ACL entry, or you can update multiple ACL entries at the same time.

## Updating a single ACL entry

To update an existing ACL entry, make the following API call in a terminal application:

```
curl -H "Fastly-Key: FASTLY_KEY" -H "Content-type: application/vnd.api+json" -X PATCH https://api.fastly.com/service/<service_id>/acl/<acl_id>/entry/<acl_entry_id> -d 'ip=127.0.0.2&subnet=32&negated=0&comment=allow'
```

The response will look like this:

```
{
  "acl_id": "<acl_id>",
  "comment": "allow",
  "created_at": "2016-04-22T19:18:42+00:00",
  "deleted_at": null,
  "id": "<acl_entry_id>",
  "ip": "127.0.0.2",
  "negated": "0",
  "service_id": "<service_id>",
  "subnet": 32,
  "updated_at": "2016-04-22T19:18:42+00:00"
}
```

## Updating multiple ACL entries

You can also update multiple ACL entries at the same time. Include an `entries` array of changes in the API call and pass an operation (`op`) parameter for every change. Possible `op` values are `create`, `update`, and `delete`.

To update multiple ACL entries at the same time, make the following API call in a terminal application:

```
curl -H "Fastly-Key: FASTLY_KEY" -H "Content-type: application/json" -X PATCH https://api.fastly.com/service/<service_id>/acl/<acl_id>/entries -d '{"entries":[{"op": "create", "ip": "192.168.0.1", "subnet": "8"}, {"op": "update", "id": "<acl_entry_id>", "ip": "192.168.0.2", "subnet": "16"}, {"op": "delete", "id": "<acl_entry_id>"}]}'
```

The response will look like this:

```
{
  "status": "ok"
}
```

## Deleting an ACL entry

**⚠ WARNING:** ACL entry deletions are permanent. If you delete an ACL entry, the entry is permanently removed from all service versions and cannot be recovered.

To permanently delete an ACL entry, make the following API call in a terminal application:

```
curl -H "Fastly-Key: FASTLY_KEY" -X DELETE https://api.fastly.com/service/<service_id>/acl/<acl_id>/entry/<acl_entry_id>
```

The response will look like this:

```
{
  "status": "ok"
}
```

# § Creating and using Edge ACLs (/guides/access-control-lists/creating-and-using-edge-acls)

Edge ACLs (/guides/access-control-lists/about-edge-acls) are lists of permissions that Varnish uses to grant or restrict access to URLs within your services. You can use the Fastly API to programmatically add, remove, and update Edge ACLs and their entries. To create an Edge ACL and use it within your service, you'll need to perform the following steps:

1. Create an empty ACL (/guides/access-control-lists/creating-and-using-edge-acls) in a working version of a service that's unlocked and not yet activated.
2. Activate the new version of the service (/guides/basic-setup/working-with-services#service-versions-and-their-activation) you associated with the empty ACL.
3. Add ACL entries (/guides/access-control-lists/creating-and-manipulating-edge-acl-entries) to the newly created ACL.

Once the ACL is created, properly associated, and filled with ACL entries, it can be called in your service.

## Attributes

Edge ACLs have the following attributes:

- **Service ID:** The ID of the Fastly service the ACL is associated with.
- **VCL Version Number:** The VCL version number the ACL is associated with. Note that the ACL will continue to reside within subsequently cloned counterparts.
- **ACL Name:** The name of the ACL.
- **ACL ID:** The unique identifier of the ACL.

## Creating an ACL

To start using an ACL, you'll need to create an empty one within a version of a service that's unlocked and not yet activated. Make the following API call in a terminal application:

```
curl -H "Fastly-Key: FASTLY_KEY" -X POST https://api.fastly.com/service/<service_id>/version/<vcl_version_number>/acl -d name=my_acl
```

The response will look like this:

```
{
  "id": "<vcl_version_number>",
  "name": "my_acl",
  "service_id": "<service_id>",
  "version": "1",
  "created_at": "2016-04-14 21:23:21",
  "updated_at": "2016-04-14 21:23:21"
}
```

Be sure to activate the new version of the service you associated with the empty ACL.

## Viewing ACLs

To see information related to a single ACL (in this example, `my_acl`) attached to a particular version of a service, make the following API call in a terminal application:

```
curl -H "Fastly-Key: FASTLY_KEY" https://api.fastly.com/service/<service_id>/version/<vcl_version_number>/acl/my_acl
```

The response will look like this:

```
{
  "id": "<acl_id>",
  "name": "my_acl",
  "service_id": "<service_id>",
  "version": "<vcl_version_number>",
  "created_at": "2016-04-14 21:23:21",
  "updated_at": "2016-04-14 21:23:21"
}
```

To view a list of all ACLs attached to a particular version of a service, make the following API call in a terminal application:

```
curl -H "Fastly-Key: FASTLY_KEY" https://api.fastly.com/service/<service_id>/version/<vcl_version_number>/acl
```

The response will look like this:

```
[
  {
    "id": "<acl_1_id>",
    "name": "my_new_acl",
    "service_id": "<service_id>",
    "version": "<vcl_version_number>",
    "created_at": "2016-04-14 21:23:21",
    "updated_at": "2016-04-15 17:23:09"
  },
  {
    "id": "<acl_2_id>",
    "name": "my_other_acl",
    "service_id": "<service_id>",
    "version": "<vcl_version_number>",
    "created_at": "2016-04-14 21:23:21",
    "updated_at": "2016-04-15 17:23:09"
  }
]
```

## Deleting an ACL

Deleting an ACL deletes the ACL and all of its associated entries. To delete an ACL (in this example, `my_new_acl`), make the following API call in a terminal application:

```
curl -H "Fastly-Key: FASTLY_KEY" -X DELETE https://api.fastly.com/service/<service_id>/version/<vcl_version_number>/acl/my_new_acl
```

The response will look like this:

```
{
  "status": "ok"
}
```

## § Manually creating access control lists (/guides/access-control-lists/manually-creating-access-control-lists)

Varnish (/guides/vcl/guide-to-vcl) allows you to use Access Control Lists (ACLs), a feature that enables fast matching of a client's IP address against a list of defined IP addresses. An ACL looks like this:

```
# Who is allowed access ...
acl local {
  "localhost";
  "192.168.1.0"/24; /* and everyone on the local network */
  ! "192.168.1.23"; /* except for the dial-in router */
}
```

## Defining an ACL

Using ACLs requires you to create and add custom VCL to Fastly's boilerplate VCL. To define an ACL in your Fastly configuration:

1. Ask support (mailto:support@fastly.com) to enable custom VCL uploading (/guides/vcl/uploading-custom-vcl) for your account.
2. Read about how to mix and match custom VCL (/guides/vcl/mixing-and-matching-fastly-vcl-with-custom-vcl) with Fastly's VCL.
3. Create a custom VCL file with your ACL definitions included in the appropriate location. Use the example shown below as a guide. You can reference the ACL in your configuration (`vcl_recv`) using a match operation which can be located above or below `#FASTLY_recv` - the placement only matters for the order of operations within Varnish's execution of your configuration.

```
# If you are using the "include" keyword
include "myACL1.vcl";

# And/or if you are using an actual ACL block
acl local {
    "localhost";
    "192.168.1.0"/24; /* and everyone on the local network */
    ! "192.168.1.23"; /* except for the dial-in router */
}

sub vcl_recv {
    # block any requests to Admin pages not from local IPs
    if (req.url ~ "^/admin" && client.ip !~ local) {
        error 403 "Forbidden";
    }
}
```

4. Upload the file (/guides/vcl/uploading-custom-vcl) in the Varnish Configuration area of your service.

## Shielding Caveats

Be aware that if you've enabled shielding (/guides/performance-tuning/shielding#enabling-shielding), you need to ensure the client IP check is only executed at the edge. For example:

```
sub vcl_recv {
    # block any requests to Admin pages not from local IPs
    if (req.url ~ "^/admin" && client.ip !~ local && !req.http.Fastly-FF){
        error 403 "Forbidden";
    }
}
```

The `client.ip` provides the source address connecting to Fastly. In the case of Origin Shielding (/guides/about-fastly-services/about-fastlys-origin-shielding-features), that address gets overwritten as one Fastly POP (<https://www.fastly.com/network-map>) connects to another. The `Fastly-FF` header in the above example ensures this code is not executed on the shield server because it gets added by the edge node.

You can create different behaviors based on any other attributes from a request as well, such as location and cookie presence.

• Guides (/guides/) > Developer's tools (/guides/devtools) > API Caching (/guides/api-caching/)

## § Enabling API caching (/guides/api-caching/enabling-api-caching)

Application Programming Interfaces (APIs) allow you to retrieve data from a variety of web services. Fastly makes it possible for you to cache your API so you can accelerate the performance of your service-oriented architecture. It optimizes your API's performance by efficiently handling traffic bursts and reducing latency.

This guide consists of three parts:

- Determining which API endpoints to cache
- Configuring PURGE requests
- Setting up Fastly

## An example

Let's look at an example to learn how API caching works. Imagine we're an online magazine with articles on which users can make comments. Each article can have many comments, and each comment is authored by exactly one user.

Relational Schema



©fastly 2013

We'll design a RESTful API specification and use it to manipulate and retrieve comments:

- `GET /comment` - Returns a list of all comments
- `GET /comment/:id` - Returns a comment with the given ID
- `POST /comment` - Creates a new comment
- `PUT /comment/:id` - Updates a comment with the given ID
- `DELETE /comment/:id` - Deletes a comment with the given ID

The create, read, update, and delete (CRUD) methods ensure the API can perform its basic operations, but they don't expose the relational aspect of the data. To do so, you would add a couple of relational endpoints:

- `GET /articles/:article_id/comments` - Get a list of comments for a given article
- `GET /user/:user_id/comments` - Get all comments for a given user

Endpoints like these allow programmers to get the information they need to do things like render the HTML page for an article, or display comments on a user's profile page. While there are many other possible endpoints we could construct, this set should suffice for the purposes of this guide. Let's assume that the API has been programmed to use an Object-Relational Mapper (ORM), such as ActiveRecord, when interacting with the database.

## Determining which API endpoints to cache

Start by identifying the URLs you want to cache. We recommend splitting the specification endpoints into two groups.

The first group, called "accessors," retrieves or accesses the comment data. These are the endpoints you want to cache using Fastly. Using the example, four endpoints match this description:

- `GET /comment`
- `GET /comment/:id`
- `GET /article/:article_id/comments`
- `GET /user/:user_id/comments`

The second group, called "mutators," changes or mutates the comment data. These endpoints are always dynamic, and are therefore uncacheable. Using the example, three endpoints match this description:

- `POST /comment`
- `PUT /comment/:id`
- `DELETE /comment/:id`

You should see a pattern emerging. Because the example API is RESTful, we can use a simple rule to identify the accessor and mutator endpoints: GET endpoints can be cached, but PUT, POST, and DELETE endpoints cannot.

Once you've gathered this information, you're ready to program the API to configure PURGE requests.

## Configuring PURGE requests

Don't be tempted to point at the PUT, POST, and DELETE endpoints as the place where data is modified. In most modern APIs, these endpoints represent an interface to the actual model code responsible for handling the database modifications.

In the example, we assumed that we'd be using an ORM to perform the actual database work. Most ORMs allow programmers to set special "callbacks" on models that will fire when certain actions have been performed (e.g., before or after validation, or after creating a new record).

For purging, we are interested in whether a model has saved information to the database — whether it's a new record, an update to an existing record, or the deleting of a record. At this point, we'd add a callback that tells the API to send a PURGE request to Fastly for each of the cacheable endpoints.

For an ActiveRecord comments model, you could do something like this:

```
require 'fastly'

class Comment < ActiveRecord::Base
  fastly = Fastly.new(api_key: 'MY_API_KEY')

  after_save do
    fastly.purge "/comment"
    fastly.purge "/comment/#{self.id}"
    fastly.purge "/article/#{self.article_id}/comments"
    fastly.purge "/user/#{self.user_id}/comments"
  end
end
```

Keep two things in mind when creating the callback:

- The purge code should be triggered *after* the information has been saved to the database, otherwise a race condition could be created where Fastly fetches the data from the origin server before the data has been saved to the database. This would cache the old data instead of the new data.
- These URLs are being purged because they have content that changes when a comment is changed.

With the model code in place, the API is now ready to be cached.

## Setting up Fastly

The final step is setting up Fastly. You'll need to perform the following steps:

- Create a new service
- Add the domain for the API
- Add the origin server that powers the API

In addition, you can optionally create rules that tell Fastly how to work with the specific elements that are exclusive to your API.

**NOTE:** By default, Fastly will not cache PUT, POST, and DELETE requests. For more information, see our [article on default caching behavior of HTTP methods \(/guides/debugging/using-get-instead-of-head-for-command-line-caching-tests\)](/guides/debugging/using-get-instead-of-head-for-command-line-caching-tests).

## Creating a new service

Follow these instructions to add a new service to your Fastly account:

1. Log in to the Fastly web interface and click the **Configure** link.
2. Click the **Options** button to open the **Manage service** menu and select **Create**. The Create a new service page appears.

### Create a new service

[Or use the setup guide](#)

Server address and port

:

★ Required

The IP address (or hostname) and port number for your origin server. To [enable TLS](#), use port 443.

▸ [Setting the server address and port](#)

Domain

★ Required

The domain name of your website.

▸ [Setting the domain name](#)

Description

★ Required

The name of your service, such as **My service**.

CREATE

CANCEL

3. Fill out the **Create a new service** fields as follows:
  - In the **Server address and port** fields, type the IP address or hostname of your API server and the port number.
  - In the **Domain** field, type the domain name associated with your API (e.g., `api.example.com`).
  - In the **Description** field, type a name for this service (e.g., `My API Service`).
4. Click the **Create** button.

Your service now appears in the service menu.

## Adding the domain


Follow these instructions to add the API's domain name to your Fastly service:

1. On the **Configure** page, click the **Edit configuration** button and then select **Clone active**. The Domains page appears.
2. Click the **Create domain** button. The Create a new domain page appears.

## Create a new domain

Domain Name

api.example.com

 Required

The domain name of your website.

[▶ Setting the domain name](#)

[▶ What if I'm using apex domains?](#)

Description

My API domain name

An optional comment that describes your domain.

CREATE

CANCEL

3. Fill out the **Create a new domain** fields as follows:
  - In the **Domain Name** field, type the domain name for the API.
  - In the **Description** field, type a description for the domain name.
4. Click **Create**. Your API's domain name appears in the list of domains.

## Adding the origin server

Follow these instructions to add the origin server that runs the API software to your Fastly service:

1. Click the **Origins** tab. The Origins page appears.
2. Click the **Create host** button. The Create a new host page appears.



## Create a new host

CONDITION

This will happen all the time unless you [Attach a condition](#)

Address

: 80

★ Required

The IP address (or hostname) and port number for your origin server. To [enable TLS](#), use port 443.

Shielding

The shield POP designated to reduce inbound load on your origins by serving cached data. Learn more about [POP request handling](#) and the [caveats of shielding](#).

Health check

[+ Create a new health check](#)

A small request assigned to [test the state of your origin](#).

Auto load balance

Enables Fastly to [load balance between two or more backend hosts](#).

[► Tips](#)

Description

★ Required

The name of your origin, such as [My origin server](#).

3. Fill out the **Create a new host** fields as follows:

- In the **Address** field, type the IP address or hostname of the origin server, and in the port field, type the port number of the origin server.
- From the **Shielding** menu, optionally select a POP to enable the shielding feature (</guides/performance-tuning/shielding>)
- From the **Health check** menu, optionally select a health check for this origin server. For more information, see our health checks tutorial (</guides/basic-configuration/health-checks-tutorial>).
- From the **Auto load balance** menu, select an option to enable or disable load balancing for this origin server. For more information, see our guide on load balancing (</guides/performance-tuning/load-balancing-configuration>).
- If you enabled load balancing, type a weight in the **Weight** field.
- In the **Description** field, type a description for the origin server.

4. Click **Create**. The origin server appears in the list of origin servers.

## § Implementing API cache control (</guides/api-caching/implementing-api-cache-control>)

This guide explains how to implement API cache control. Once you've enabled API caching (</guides/api-caching/enabling-api-caching>), and ensured purging works properly with your cached data, you can set up specific headers like cache-control and surrogate-control to change when data is cached.

This guide consists of three parts:

- Understanding cache control headers
- Updating the example to use cache-control
- Implementing cache-control

## Understanding cache control headers

In general, we assume that GET requests are cached and PUT, POST, and DELETE requests are not. For an ideal REST API, this rule works well. Unfortunately, most APIs are far from ideal and require additional caching rules for some requests.

For these reasons, it's a good idea to set cache-control headers when migrating APIs to Fastly. Cache-control, as defined by RFC 7234 (<https://tools.ietf.org/html/rfc7234#section-5.2>) (the HTTP specification), includes many different options for appropriate handling of cached data. Specifically, cache-control headers tell user agents (e.g., web browsers) how to handle the caching of server responses. For example:

- `Cache-Control: private`
- `Cache-Control: max-age=86400`

In the first example, `private` tells the user agent the information is specific to a single user and should not be cached for other users. In the second example, `max-age=86400` tells the user agent the response can be cached, but that it expires in exactly 86,400 seconds (one day).

Fastly respects cache-control headers by default, but you can also use another proxy-specific header: surrogate-control. Surrogate-control headers are similar to cache-control headers, but provide instructions to reverse proxy caches like Fastly. You can use cache-control and surrogate-control headers together. For more information about cache-control and surrogate-control headers, see our cache control tutorial (</guides/tutorials/cache-control-tutorial>).

## An updated example

Let's take a look at how the cache-control headers could be used in our original example, the comments API (</guides/api-caching/enabling-api-caching#an-example>). Recall the API endpoint that provided a list of comments for a given article:

```
GET /article/:article_id/comments
```

When a user submits a comment for a given article, the response from this endpoint will be purged from the Fastly cache by the comment model (</guides/api-caching/enabling-api-caching#configuring-purge-requests>). It's hard to predict when content will change. Therefore, we'd like to ensure the following:

1. If the content doesn't change, it should stay cached in Fastly for a reasonable amount of time.
2. If the content does change, it should not be cached by the client longer than it needs to be.

The goal is to ensure that API responses will reach clients in a timely manner, but we also want to ensure that clients always have the most up-to-date information. The first constraint can be solved by using the surrogate-control header, and the second constraint can be solved by using the cache-control header:

```
Surrogate-Control: max-age=86400
Cache-Control: max-age=60
```

These headers tell Fastly that it is allowed to cache the content for up to one day. In addition, the headers tell the client that it is allowed to cache the content for 60 seconds, and that it should go back to its source of truth (in this case, the Fastly cache) after 60 seconds.

## Implementing cache control

Migrating APIs isn't easy, even for experienced teams. When migrating an API to Fastly, we recommend separating the task into three strategic endpoint migrations to make the process more manageable while still maintaining the validity of the API as a whole.

### Preparing the API

To ensure that the API bypasses the cache during the piecewise migration, we must have every API endpoint return a specific control header:

```
Cache-Control: private
```

This header tells Fastly that a request to any endpoint on the API should bypass the cache and be sent directly to the origin. This will allow us to serve the API via Fastly and have it work as expected.

**NOTE:** Modern web frameworks allow for blanket rules to be overridden by specific endpoints (for example, by the use of middlewares). Depending on how the API has been implemented, this step might be as simple as adding a single line of code.

### Serving traffic with Fastly

The next step is configuring a Fastly service (</guides/api-caching/enabling-api-caching#setting-up-fastly>) to serve the API's traffic. After you save the configuration, there will be an immediate speed improvement. This happens because Fastly's cache servers keep long-lasting connections to the API's origin servers, which reduces the latency overhead of establishing multiple TCP connections.

### Migrating endpoints

Now we can implement instant purge caching for each cacheable API endpoint, one at a time. The order in which this is done depends on the API, but by targeting the slowest endpoints first, you can achieve dramatic improvements for endpoints that need them the most. Because each endpoint can be worked on independently, the engineering process is easier to manage.

## Excluding endpoints

The last step is deciding which API endpoints you don't want Fastly to cache. To disable caching for endpoints, you'll need to add new conditions for the endpoints (</guides/performance-tuning/controlling-caching#conditionally-preventing-pages-from-caching>). As you learned in [Preparing the API](#), using the `Cache-Control: private` header is another option for disabling caching.

## § Purging API cache with surrogate keys (</guides/api-caching/purging-api-cache-with-surrogate-keys>)

Fastly makes it possible for you to cache your API so you can accelerate the performance of your service-oriented architecture. Of course, caching your API is one thing - efficiently invalidating the API cache is another matter entirely. If you've already enabled API caching (</guides/api-caching/enabling-api-caching>) and implemented API cache control (</guides/api-caching/implementing-api-cache-control>), you've probably run into this problem, which was aptly described by Phil Karlton (<http://martinfowler.com/bliki/TwoHardThings.html>):

*There are only two hard things in computer science: cache invalidation and naming things.*

This guide explains how to use the Fastly API to purge your API cache with surrogate keys (</guides/purging/getting-started-with-surrogate-keys>). Surrogate keys allow you to reduce the complexity of caching an API by combining multiple cache purges into a single key-based purge.

## What's a surrogate key?

Surrogate keys allow you to selectively purge related content. Using the `Surrogate-Key` header, you can "tag" an object, such as an image or a blog post, with one or more keys. When Fastly fetches an object from your origin server, we check to see if you've specified a `Surrogate-Key` header. If you have, we add the response to a list we've set aside for each of the keys.

When you want to purge all of the responses associated with a key, issue a key purge (</guides/purging/single-purges#purging-with-a-key>) request and all of the objects associated with that key will be purged. This makes it possible to combine many purges into a single request. Ultimately, it makes it easier to manage categorically related data.

To learn more about surrogate keys and to see how you can integrate them into your application, see our [guide on getting started with surrogate keys](/guides/purging/getting-started-with-surrogate-keys).

## Example: Purging categories

To see how surrogate keys work in conjunction with an API endpoint, imagine you have an online store and an API endpoint that returns the details of a product. When a user wants to get information about a specific product, like a keyboard, the request might look like this:

```
GET /product/12345
```

If your API is using Fastly and the response is not already cached, Fastly will make a request to your API's origin server and receive a response like this:

```
HTTP/1.1 200 OK
Content-Type: text/json
Cache-Control: private
Surrogate-Control: max-age=86400
Surrogate-Key: peripherals keyboards

{"id: 12345, name: "Uber Keyboard", price: "$124.99"}
```

You knew that entire product categories would occasionally need to be purged, so you thoughtfully included the `peripherals` and `keyboards` product categories as keys in the `Surrogate-Key` header. When Fastly receives a response like this, we add it to an internal map, strip out the `Surrogate-Key` header, cache the response, and then deliver it to the end user.

Now imagine that your company decides to apply a 10% discount to all peripherals. You could issue the following key purge to invalidate all objects tagged with the `peripherals` surrogate key:

```
PURGE /service/:service_id/peripherals
```

When Fastly receives this request, we reference the list of content associated with the `peripherals` surrogate key and systematically purge every piece of content in the list.

## Relational dependencies

Your API can use surrogate keys to group large numbers of items that may eventually need to be purged at the same time. Consider the example presented above. The API for your online store could have surrogate keys for product types, specific sales, or manufacturers.

From this perspective, the `Surrogate-Key` header provides Fastly with information about relations and possible dependencies between different API endpoints. Wherever there's a relation between two different types of resources in an API, there might be a good reason to keep them categorized by using a surrogate key.

## Example: Purging product reviews and action shots

To learn how surrogate keys can help with relational dependencies, imagine that your online store wants to allow buyers to post product reviews and "action shots" depicting the products in use. To support these new features, you'll need to change your API. First, you'll need to create a new `review` endpoint:

```
GET /review/:id
POST /review
```

Next, you'll need to create a new `action_shot` endpoint:

```
POST /product/:id/action_shot
GET /product/:id/action_shot/:shot_id
```

Since both of the new endpoints refer to specific products, they'll need to be purged when relevant product information changes. Surrogate keys are a perfect fit for this use case. You can implement them by modifying the `review` and `action_shot` to return the following header:

```
Surrogate-Key: product/:id
```

This relates each of the endpoints to a specific product in the cache (where `:id` is the product's unique identifier). When the product information changes, your API issues the following a key purge:

```
PURGE /service/:service_id/product/:id
```

When Fastly receives this request, we purge each of the related endpoints at the same time.

## Variations on a theme

You'll also want to consider using surrogate keys if your API has many different endpoints that all derive from a single source. Any time the source data changes, each of the endpoints associated with it will need to be purged from the cache. By associating each of the endpoints with a surrogate key, a single purge can be issued to purge them from the cache when the source changes.

## Example: Purging product images

To understand how this works, imagine that your online store has an API endpoint for retrieving product images in various sizes:

```
GET /product/:id/image/:size
```

This endpoint returns an image of the appropriate `:size` (e.g., `small`, `medium`, `large`) for the product of the given `:id`. To save disk space, you opt to have the API generate each specifically sized image from a single source image using an imaging library like ImageMagick (<http://www.imagemagick.org/>). Since the sales and marketing team uses the API to upload new product images, you set up the endpoint to include a surrogate key:

```
Surrogate-Key: product/:id/image
```

When the PUT endpoint for uploading a product image is called, the API sends the following purge request:

```
PURGE /service/:service_id/product/:id/image
```

When Fastly receives this request, we purge all size variations of the product image.

• [Guides \(/guides/\)](/guides/) > [Developer's tools \(/guides/devtools/\)](/guides/devtools/) > [Conditions \(/guides/conditions/\)](/guides/conditions/)

## § About conditions (/guides/conditions/about-conditions)

Use conditions to specify how to edit and activate configurations. This allows you to add logic to any basic configuration object in a service. When you assign a condition to an object in your service, you are wrapping the object in a VCL if-statement. If the condition is met, then the object is used when processing an end-user request.

## Before you begin

Because of their complexity and power, be sure you're comfortable with the following before you begin using conditions (</guides/conditions/using-conditions/>):

- **Know how to edit and activate configurations using the Fastly web interface.** If you are not familiar with basic editing using the web interface, see our documentation (</guides/basic-setup/working-with-services/>) to learn more before moving forward.
- **Understand basic C-style logical expression syntax** (e.g., basic logic, operators such as `&&`, and precedence). If you do not have experience with these types of expressions, refer to a basic programming guide (such as the Tizag Perl tutorial (<http://www.tizag.com/perlT/>)) that deals with "IF" style expressions in either C or Perl. (These languages are not directly applicable to the tutorial (</guides/conditions/using-conditions#an-example-of-adding-conditions>) because VCL is neither C nor Perl. However, the syntax is similar.)

## Condition basics

Conditions are saved in your service like any other configuration object, but they can only be accessed, assigned, and manipulated (</guides/conditions/using-conditions#an-example-of-adding-conditions>) via the Attach a condition link next to that object. The link launches an Add a condition window that allows you to select an existing condition or create a new one and then apply it in multiple places across your configuration.

Properly configured conditions require only a few parameters, making them appear deceptively simple:

- a **Type** - The type of condition being added, usually automatically selected by the system. For example, if you create a cache header, you can only apply a cache condition. Note that with a response setting, you can select either a request condition (based on what the user is requesting) or a cache condition (based on what the backend is returning).
- a **Name** - A human-readable identifier for the condition.
- an **Apply if** statement - The logical expression to execute in VCL to determine if condition resolves as True or False.
- the **Priority** - A number used to determine the order in which multiple conditions execute. Lower numbers execute first. Defaults to 10.

Most problems occur in the Apply if parameter because it uses actual VCL code to specify when a condition should be applied.

## Working with conditions

A basic condition used by Fastly might look like this:

```
req.url ~ "^/api/"
```

This condition tells the cache server to pay attention if a request it receives begins with the string `"/api/"`. The `~` operator performs a Perl-style regular expression match with the `"^/api/"` regex and returns true if `req.url` matches. You might use this condition when you want to specifically route certain requests for your service. In this example, you might be routing requests to an API backend that exists on your website.

## Matching conditions on logical expressions

Conditions can perform matches on more complicated logical expressions. For example:

- `client.ip == "127.0.0.1"`: The client requesting a resource on your service has the IP `"127.0.0.1"`
- `req.http.host ~ "example.com"`: The host header of the incoming request is `"example.com"`
- `req.request ~ "POST|PUT" && req.url ~ "^/api/articles/"`: The request is a POST or a PUT and the URL begins with `"/api/articles/"`

The `client.ip`, `req.http.host`, `req.request`, and `req.url` conditions shown above all represent configuration variables in VCL.

## Operators for complex conditions

You could also get creative and create a more complex condition used by Fastly that might look like this:

```
req.http.host = "www.example.com" && (req.url !~ "^/foo" || req.url !~ "^/bar/" || req.url !~ "^/baz/")
```

This condition tells the cache server that the URL should equal `www.example.com` and the url cannot match `www.example.com/foo`, or `www.example.com/bar`, or `www.example.com/baz`. You might use this type of condition when you have multiple variables or options and want to fine-tune your results. In this example, you are indicating that you don't want URLs that contain `foo`, `bar`, or `baz` by using the following operators:

- `()` groups your expressions and restricts alteration to part of the regex
- `||` performs an alternation where each variable is checked until it finds a variable that is true
- `!~` excludes any urls that include the specified variables

## Where to go for more information

The Varnish Cache documentation (<https://www.varnish-cache.org/docs/2.1/reference/vcl.html#variables>) on variables provides a complete list of variables you can use to craft conditions (</guides/conditions/using-conditions/>). Keep in mind, however, some of the variables Varnish allows may not be available or may have no meaning in the context of Fastly services. For more information, see our Guide to VCL (</guides/vcl/guide-to-vcl/>).

## § Troubleshooting conditions (/guides/conditions/troubleshooting-conditions)

If you are having problems using conditions, here are some common mistakes to look for:


- **Don't put the `if ()` statement in the Apply if field of the condition window.** The actual if statement is implied. You only need to type an evaluated expression. For example: `req.url ~ "/special/"`.
- **Don't use the `!~` (inverse regex match) to build regular expressions that exclude particular URLs.** For example, if you want to apply something to all URLs except those that start with `/admin`, the condition for this is `req.url !~ "^/admin"` and is entered in the Apply If field. An alternative to this expression would be `!(req.url !~ "^/admin")`.

## Editing conditions

You can edit your condition that are listed on the same page as your configurations. Depending on the condition, you can edit the cache, request, responses, and header conditions.

For example, if you have a cache condition that was being applied to a response, you would edit that condition in your service's configuration as follows:


1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.
4. Click the desired tab that has the condition that you want to edit. For example, click the **Settings** tab and find the desired condition.



The screenshot shows a condition entry with the following details:

- IF Article page**
- `req.url ~ "/articles"`
- THEN**
- Article page**
- [Show details](#)
- [Detach from condition](#)
- [Delete](#)

5. Click the **IF** statement link of the condition (for example, click **IF Article pages**). For this example, the Edit this cache condition window appears.



The screenshot shows the 'Edit this request condition' window with the following details:

- Edit this request condition.**
- [Learn the basics in our conditions tutorial](#)
- Type** Request
- Name** Special API Request
- Apply if...** `req.url ~ "/articles"`
- The expression to decide whether this is run.
- [Examples](#)
- [Adjust priority \(default is 10\)](#)
- This is an advanced option. For most configurations, the default is best.
- SAVE**
- CANCEL**

6. Make your edits and click **Save**.

## Helpful links

The following links provide additional information useful when working with and troubleshooting conditions:

- Fastly Help Guides (/guides/) provides basic editing information when using the web interface.
- Tizag Perl tutorial (<http://www.tizag.com/perlT/perlif.php>) explains basic C-style logical expression syntax (e.g., basic logic, operators such as &&, precedence) that users need to be familiar with when creating conditions.
- Varnish-Cache.org documentation (<https://www.varnish-cache.org/docs/2.1/reference/vcl.html#variables>) explains different set of VCL variables that a condition can use to determine if an object applies.
- Varnish Cache documentation on variables (<https://www.varnish-cache.org/docs/2.1/reference/vcl.html#variables>) describes variables that are used to craft conditions.

## § Using conditions (/guides/conditions/using-conditions)

Conditions use the Varnish Control Language (VCL) (/guides/vcl/guide-to-vcl#about-varnish-and-why-fastly-uses-it) to define when an object should be applied while processing requests to a cache server. Once you understand some basics about conditions (/guides/conditions/about-conditions), use this guide to learn about when to use conditions and how to create them using the Fastly web interface.

### About request, response, and cache conditions

Conditions can occur in various stages of caching: when a request is being processed, when a response to that request is being processed, and when a response from your origin server arrives but before that information is (potentially) cached. Each stage of caching has access to a different set of VCL variables (as described in the Varnish-Cache.org documentation (<https://www.varnish-cache.org/docs/2.1/reference/vcl.html#variables>)) that can be used to create conditions.

To help distinguish conditions during these stages we group them into three types: request, response, and cache. Configuration objects also only apply in various stages of the pipeline and can have different types of conditions associated with them.

The following list of the condition types describes the objects to which conditions can apply and the VCL variable types that are available to them:

- **Request:** Conditions applied to objects while processing a request.
  - Objects: Request Settings, Backends, Headers, and Responses
  - VCL Variables: `client.*`, `server.*`, `req.*`
  - Example: `req.url ~ "^/foo/bar$"`
- **Response:** Conditions applied to objects while processing a response to a request.
  - Objects: Headers, Syslog
  - VCL Variables: `client.*`, `server.*`, `req.*`, `resp.*`
  - Example: `resp.status == 404`
- **Cache:** Conditions applied when Fastly receives a response from your origin, just before that response is (potentially) cached.
  - Objects: Cache Settings, Headers, Gzip, Responses
  - VCL Variables: `client.*`, `server.*`, `req.*`, `beresp.*`
  - Example: `!beresp.cachable && req.url ~ "^/articles/"`

These conditions allow you to control your configuration with minimal programming.

### An example of adding conditions

The scenario: You want to add a new origin server that handles a specific portion of your API requests. Some requests to this API must be cached differently than other requests to your API, so you want to set special headers for specific types of requests. Specifically, you don't want your new origin server to cache PUT, POST, or DELETE requests because they're special for this particular API and send back extra, time dependent, meta-information in each response. And finally, you want to track the effectiveness of doing this. To accomplish all of this using conditions via the Fastly web interface, you would:

1. Create a new origin server to handle the special API traffic.
2. Create a new condition that tells the cache how to route requests to that origin server.
3. Create a new cache setting object and then a new condition that applies that object at the appropriate time.
4. Create a new header to track the specific type of API requests.

5. Create a new cache condition to make sure that the header is only set on specific type of request.
6. Check your work.

## Create a new origin server

To create a new origin server that will handle the special API traffic, follow the steps below.

1. Log in to the Fastly web interface and click the **Configure** link. The Configure page appears.
2. Select the appropriate service from the service menu.
3. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.
4. Click the **Origins** tab. The Origins page appears.
5. Click the **Create host** button to create a new backend for your special API traffic. The Create a new host page appears.

## Create a new host

**CONDITION** This will happen all the time unless you [Attach a condition](#)

**Address**  :  ★ Required

The IP address (or hostname) and port number for your origin server. To [enable TLS](#), use port 443.

**Shielding**

The shield POP designated to reduce inbound load on your origins by serving cached data. [Learn more about POP request handling](#) and the [caveats of shielding](#).

**Health check**

[+ Create a new health check](#)

A small request assigned to [test the state of your origin](#).

**Auto load balance**

Enables Fastly to [load balance between two or more backend hosts](#).

[► Tips](#)

**Description**  ★ Required

The name of your origin, such as [My origin server](#).

6. Fill out the **Create a new host** fields as follows:

- In the **Address** field, type the address of your API server (for example, `867.53.0.9`) and in the **Port** field, type `80`.
- Leave the **Shielding** and **Health check** controls set to their defaults.
- From the **Auto load balance** menu, select **No** to exclude this server from the normal, automatic load balancing pool.
- In the **Description** field, type the name of your API server (for example, `Special API Set Header`).

7. Click the **Create** button. Your new host appears in the Host area.

## Create a request condition

Once you've created a new origin serve to handle the special API traffic, tell the cache how to route requests to this origin server by creating a request condition.

1. In the **Hosts** area, click the **Attach a condition** link next to the name of the origin server you just created. The Add a condition to window appears.
2. You can either select an available condition or you can click the **Create a new request condition** button. The Create a new request condition window appears.



Create a new request condition.

Learn the basics in our [conditions tutorial](#)

Type

Request

Learn more about the [different types of conditions](#).

Name

Special API Request

★ Required

Apply if...

req.url ~ "^/special/"

The expression to decide whether this is run.

[Examples](#)

[Adjust priority \(default is 10\)](#)

This is an advanced option. For most configurations, the default is best.

SAVE AND APPLY TO SPECIAL API SET HEADER

CANCEL

3. Fill out the **Create a new request condition** fields as follows:

- In the **Type** field, the type of condition that you are creating appears.
- In the **Name** field, type a descriptive name for the new condition (for example `Special API Request`).
- In the **Apply if** field, type the appropriate request condition that will be applied (for example, `req.url ~ "^/special/"` to address all requests related to the special API server).
- Leave the **Priority** field set to its default value.

4. Click the **Save and apply to** button to create the new condition for the host.

## Create a cache settings object and its condition

The requests now are being properly routed to the new origin, you need to make sure that it doesn't cache any responses from PUT, POST, or DELETE requests because they are special for this particular API and send back extra, time dependent, meta-information in each response. To do this, create a cache settings object in the configuration.

1. Click the **Settings** tab. The Settings page appears.
2. In the **Cache Settings** area, click the **Create cache setting** button. The Create a new cache setting page appears.

## Create a new cache setting

This will override your cache control headers. In most cases, use with conditions applied.

### CONDITION

This will happen all the time unless you [Attach a condition](#)

TTL (seconds)

The TTL (Time To Live) entered will set the lifespan of the object within our cache nodes.

Action

This setting decides [how the request will be handled](#).

Stale TTL (seconds)

Stale TTL (Time To Live) is used by your application to [serve stale content](#). It determines how long to keep stale data after it has expired. Note there's custom VCL required to [complete this setup](#).

Description

★ Required

The name of your cache setting, such as **My cache setting**.

CREATE

CANCEL

3. Fill out the **Create a new cache setting** fields as follow:

- Leave the **TTL (seconds)** field set to its default value.
- From the **Action** menu, select **Pass**.
- Leave the **Stale TTL (seconds)** field set to its default value.
- In the **Description** field, type a descriptive name for the new cache settings.

4. Click the **Create** button.

Then, create a new condition that specifies when the cache settings object should be applied.

1. In the **Cache Settings** area, click the **Attach a condition** link next to the name of the cache setting you just created. The Add a condition to window appears.
2. Click **Create a new cache condition** button. The Create a new condition window appears.

Create a new cache condition.

Learn the basics in our [conditions tutorial](#)

Type

Cache

Learn more about the [different types of conditions](#).

Name

Special API Set Header

★ Required

Apply if...

req.request ~ "PUT|POST|DELETE" && beresp.status == 200

The expression to decide whether this is run.

[Examples](#)

[Adjust priority \(default is 10\)](#)

This is an advanced option. For most configurations, the default is best.

SAVE AND APPLY TO THIS CACHE SETTING

CANCEL

3. Fill out the **Create a new condition** fields as follows:

- In the **Type** field, the type of condition you are creating appears.
- In the **Name** field, type a descriptive name for the new condition (for example, `Special API Set Header`).
- In the **Apply if** field, type the appropriate request condition that will be applied (for example, `req.request ~ "PUT|POST|DELETE" && beresp.status == 200`).
- Leave the **Priority** field set to its default value.

4. Click the **Save and apply to** button to create the new condition for the cache setting.

## Create a new header and its condition

To make sure you can track the effectiveness the new API, create a new header so you can use it to gather information about the special API requests as they happen.

1. Click the **Content** tab. The Content page appears.
2. In the **Headers** area, click the **Create header** button to create a new header. The Create a new header page appears.

## Create a new header

Learn more about this section in our [headers tutorial](#).

### CONDITION

This will happen all the time unless you [Attach a condition](#)

Type / Action

The type of header and the action performed on it.

Destination  ★ Required

The name of the header that will be affected by the selected action. For example: `http.Content-Type`, `http.Set-Cookie`, `http.Via`, `http.Location`, or `http.Access-Control-Allow-Origin`.

Source

Create a new response condition.

Learn the basics in our [conditions tutorial](#)

Type

Response

Learn more about the [different types of conditions](#).

Name

Special API Set Header

★ Required

Apply if...

req.url ~ "^/special" && resp.status == 200

The expression to decide whether this is run.

[Examples](#)

[Adjust priority \(default is 10\)](#)

This is an advanced option. For most configurations, the default is best.

SAVE AND APPLY TO SPECIAL API SET HEADER

CANCEL

6. Fill out the **Create a new request condition** fields as follows:

- In the **Type** field, the type of condition you are creating appears. Be aware that the type for an associated condition is inherited from the Type field in the header you create. Also, when you create a new condition for a header, you can change the type of condition applied (for example, if the condition type is `Response`, you can change it to `Request` or `Cache`).
- In the **Name** field, type a descriptive name for the new condition (for example, `Special API Response Condition`).
- In the **Apply if** field, type the appropriate request condition that will be applied (for example, `req.url ~ "^/special" && resp.status == 200`).
- Leave the **Priority** field set to its default value.

7. Click the **Save and apply to** button to create the new condition for the header.

## Check your work

Before activating the application, review the generated VCL (</guides/vcl/previewing-and-testing-vcl-before-activating-it>) to see how Fastly converted the objects and conditions into actual VCL. For the example show above, the VCL for the request condition appears as:

```
# Condition: Special API Request Prio: 10
if (req.url ~ "^/special/") {
    set req.backend = F_Special_API_Server;
}
#end condition
```

The cache settings and condition VCL appears as:

```
if (req.request ~ "PUT|POST|DELETE" && beresp.status == 200) {
    set beresp.ttl = 0s;
    set beresp.grace = 0s;
    return(pass);
}
```

And the new header response condition VCL appears as:

```
# Condition Special API Response Condition Prio: 10
if (req.url ~ "^/special" && resp.status == 200) {

    # Header rewrite Special API Set Header: 10
    set resp.http.super = "Thanks for asking!";
}
```

As you become more familiar with the VCL syntax and programming, look at the generated VCL to see if the configuration is doing what you think it is doing (most VCL is pretty simple once you know what the variables are referring to).

• [Guides \(/guides/\)](/guides/) > [Developer's tools \(/guides/devtools/\)](/guides/devtools/) > [Dynamic Servers \(/guides/dynamic-servers/\)](/guides/dynamic-servers/)

## § About Dynamic Servers (/guides/dynamic-servers/about-dynamic-servers)

Fastly's Dynamic Servers feature distributes and directs incoming requests to a pool of servers that you dynamically manage. The benefits to using Dynamic Servers include:

- support for any cloud infrastructure provider and any type of server instances
- high availability of web applications when using health checks
- compatibility with TLS termination, HTTP/2, and IPv6
- server stickiness without requiring a cookie-based approach

To set up Dynamic Servers, you attach a pool to a service (</guides/dynamic-servers/creating-and-using-pools-with-dynamic-servers/>), then add versionless server entries that are stored separately from your VCL configuration. You can use the Fastly API (</api/dynamicservers/>) to programmatically add, remove, and update pools and server entries.

**ⓘ IMPORTANT:** This feature is part of a limited availability release. For more information, see our [product and feature life cycle \(/guides/customer-support/product-and-feature-life-cycles#limited-availability\)](/guides/customer-support/product-and-feature-life-cycles#limited-availability) descriptions.

## How Dynamic Servers work

Like Edge Dictionaries (</guides/edge-dictionaries/about-edge-dictionaries/>) and Edge ACLs (</guides/access-control-lists/about-edge-acls/>), Dynamic Servers have two major components: the pool and the server entries within it. Pools act as containers for server entries that store the hostnames or IP addresses of server to which incoming requests can be directed. Each pool is attached to a version of a service, but server entries are versionless and any changes will become effective immediately.

## When Dynamic Servers might be useful

Dynamic Servers might be useful for organizations that need to load balance requests among servers in a single region, such as AWS EC2 instances in the US East region. They're also useful for evaluating new server instance types and new software deployments. Dynamic Servers are a great way to scale micro services.

## Getting started

You'll need to follow these steps:

1. Create an empty pool (</guides/dynamic-servers/creating-and-using-pools-with-dynamic-servers/>) in a working version of a service that's unlocked and not yet activated.
2. Activate (</guides/basic-setup/working-with-services#editing-and-activating-versions-of-services/>) the new version of the service you associated with the empty pool.
3. Add server entries (</guides/dynamic-servers/creating-and-using-pools-with-dynamic-servers/>) to the newly created pool.

Once the pool is created, properly associated, and filled with server entries, it can be called in your service.

## Limitations

Keep the following limitations in mind as you use Dynamic Servers:

- **There are limits to the number of pools you can attach and server entries you can add.** You attach a maximum of two pools per service, and add a maximum of three servers per pool.
- **Pools created with custom VCL cannot be manipulated using the API.** If you create a pool using the API, you can use the API to make changes to it, and use custom VCL to interact with it. If you create a pool using custom VCL, that pool's server entries must always be manipulated via custom VCL.
- **Custom VCL manipulations are always versioned.** This is true for pools and server entries created using custom VCL and for any logic created to interact with them.

- **Pools need at least one server entry.** Server entries cannot be deleted when a pool only has one entry left.
- **Server entry deletions are permanent.** If you delete a server entry, the entry is permanently removed from all service versions and cannot be recovered.
- **When you delete a pool, you'll only delete it from the service version you're editing.** Pools are tied to versions and can be cloned and reverted. When using pools, we want you to be able to do things like delete a pool from a current version of your service in order to roll back your configuration to a previous version using as few steps as possible.
- **Audit logs don't exist for server entry changes.** If you use the API to add, update, or remove a server entry, there will be no record of it. The only record of a change will exist when you compare service versions to view the point at which the pool was associated with the service version in the first place.

## § Creating and using pools with Dynamic Servers (/guides/dynamic-servers/creating-and-using-pools-with-dynamic-servers)

Fastly's Dynamic Servers (/guides/dynamic-servers/about-dynamic-servers) feature distributes and directs incoming requests to a pool of servers (/guides/dynamic-servers/creating-and-using-server-entries-with-dynamic-servers) that you dynamically manage. A pool is responsible for balancing requests among a group of servers. In addition to balancing, pools can be configured to attempt retrying failed requests. Pools have a quorum setting which can be used to determine when the pool as a whole is considered `up` in order to prevent problems following an outage as servers come back up.

**ⓘ IMPORTANT:** This feature is part of a limited availability release. For more information, see our [product and feature life cycle \(/guides/customer-support/product-and-feature-life-cycles#limited-availability\)](/guides/customer-support/product-and-feature-life-cycles#limited-availability) descriptions.

### Creating a pool

To start using Dynamic Servers, you'll need to create an empty pool within a version of a service that's unlocked and not yet activated. Make the following API call in a terminal application:

```
curl -vs -H "Fastly-Key: $FASTLY_KEY" -X POST https://api.fastly.com/service/<service_id>/version/<service_version>/pool -d 'name=<pool_name>&comment=<comment>'
```

The response will look like this:

```
{
  "between_bytes_timeout": "10000",
  "comment": "<comment>",
  "connect_timeout": "1000",
  "created_at": "2016-08-01T14:43:22+00:00",
  "deleted_at": null,
  "first_byte_timeout": "15000",
  "healthcheck": null,
  "id": "2IpWU5CGzPpbpGsABSDops",
  "max_conn_default": "200",
  "max_tls_version": null,
  "min_tls_version": null,
  "name": "<pool_name>",
  "quorum": "75",
  "request_condition": null,
  "service_id": "<service_id>",
  "shield": null,
  "tls_ca_cert": null,
  "tls_cert_hostname": null,
  "tls_check_cert": 1,
  "tls_ciphers": null,
  "tls_client_cert": null,
  "tls_client_key": null,
  "tls_sni_hostname": null,
  "type": "random",
  "updated_at": "2016-08-01T14:43:22+00:00",
  "use_tls": 0,
  "version": "<service_version>"
}
```

Be sure to activate the new version of the service you associated with the empty ACL.

**NOTE:** You can attach a maximum of two pools per service.

## Viewing pools

To see information related to a single pool (in this example, `SP_Prod_Pool_1`) attached to a particular version of a service, make the following API call in a terminal application:

```
curl -vs -H "Fastly-Key: $FASTLY_KEY" -X POST https://api.fastly.com/service/<service_id>/version/<service_version>/pool/<pool_id>
```

The response will look like this:

```
{
  "between_bytes_timeout": "10000",
  "comment": "just my first pool",
  "connect_timeout": "1000",
  "created_at": "2016-08-01T14:43:22+00:00",
  "deleted_at": null,
  "first_byte_timeout": "15000",
  "healthcheck": null,
  "id": "<pool_id>",
  "max_conn_default": "200",
  "max_tls_version": null,
  "min_tls_version": null,
  "name": "SP_Prod_Pool_1",
  "quorum": "75",
  "request_condition": null,
  "service_id": "<service_id>",
  "shield": null,
  "tls_ca_cert": null,
  "tls_cert_hostname": null,
  "tls_check_cert": 1,
  "tls_ciphers": null,
  "tls_client_cert": null,
  "tls_client_key": null,
  "tls_sni_hostname": null,
  "type": "random",
  "updated_at": "2016-08-01T14:43:22+00:00",
  "use_tls": 0,
  "version": "<service_version>"
}
```

To view a list of all pools attached to a particular version of a service, make the following API call in a terminal application:

```
curl -vs -H "Fastly-Key: $MY_KEY" -X POST https://api.fastly.com/service/<service_id>/version/<service_version>/pool
```

The response will look like this:



```
[
{
  "between_bytes_timeout": "10000",
  "comment": "just my first pool",
  "connect_timeout": "1000",
  "created_at": "2016-08-01T14:43:22+00:00",
  "deleted_at": null,
  "first_byte_timeout": "15000",
  "healthcheck": null,
  "id": "2IpWU5CGzPpbGgsABSDops",
  "max_conn_default": "200",
  "max_tls_version": null,
  "min_tls_version": null,
  "name": "SP_Prod_Pool_1",
  "quorum": "75",
  "request_condition": null,
  "service_id": "<service_id>",
  "shield": null,
  "tls_ca_cert": null,
  "tls_cert_hostname": null,
  "tls_check_cert": 1,
  "tls_ciphers": null,
  "tls_client_cert": null,
  "tls_client_key": null,
  "tls_sni_hostname": null,
  "type": "random",
  "updated_at": "2016-08-01T14:43:22+00:00",
  "use_tls": 0,
  "version": "<service_version>"
}
]
```

## Deleting a pool

Deleting a pool deletes the pool and all of its associated server entries. To delete a pool (in this example, `SP_Prod_Pool_1`), make the following API call in a terminal application:

```
curl -vs -H "Fastly-Key: $MY_KEY" -X DELETE https://api.fastly.com/service/<service_id>/version/<service_version>/pool/<pool_id>
```

The response will look like this:

```
{
  "status": "ok"
}
```

## § Creating and using server entries with Dynamic Servers (/guides/dynamic-servers/creating-and-using-server-entries-with-dynamic-servers)

Fastly's Dynamic Servers (/guides/dynamic-servers/about-dynamic-servers) feature distributes and directs incoming requests to a pool (/guides/dynamic-servers/creating-and-using-pools-with-dynamic-servers) of servers that you dynamically manage. A server entry is an address (IP address or hostname) of a server to which the Dynamic Servers feature can forward requests. You can add up to three server entries to a pool. Fastly can then select any one of the server entries based on a selection policy defined for the pool.

**❗ IMPORTANT:** This feature is part of a limited availability release. For more information, see our [product and feature life cycle \(/guides/customer-support/product-and-feature-life-cycles#limited-availability\)](/guides/customer-support/product-and-feature-life-cycles#limited-availability) descriptions.

## Creating a server entry

To add a server entry to the pool, make the following API call in a terminal application:

```
curl -vs -H "Fastly-Key: $MY_KEY" -X POST https://api.fastly.com/service/<service_id>/pool/<pool_id>/server -d 'address=<hostname_or_ip_address>'
```

The response will look like this:

```
{
  "id": "6kEuoknxiaDBCLiAjKqyXq",
  "service_id": "<service_id>",
  "pool_id": "<pool_id>",
  "weight": "100",
  "max_conn": "200",
  "port": "80",
  "address": "<hostname_or_ip_address>",
  "comment": "",
  "created_at": "2016-06-20T08:20:36+00:00",
  "updated_at": "2016-06-20T08:20:36+00:00",
  "deleted_at": null
}
```

**NOTE:** You can add a maximum of three server entries per pool.

## Viewing server entries

To see information related to a single server entry, make the following API call in a terminal application:

```
curl -vs -H "Fastly-Key: $MY_KEY" https://api.fastly.com/service/<service_id>/pool/<pool_id>/server/<hostname_or_ip_address>
```

The response will look like this:

```
{
  "id": "6kEuoknxiaDBCLiAjKqyXq",
  "service_id": "<service_id>",
  "pool_id": "<pool_id>",
  "weight": "100",
  "max_conn": "200",
  "port": "80",
  "address": "<hostname_or_ip_address>",
  "comment": "",
  "created_at": "2016-06-20T08:20:36+00:00",
  "updated_at": "2016-06-20T08:20:36+00:00",
  "deleted_at": null
}
```

To view a list of all server entries attached a particular pool, make the following API call in a terminal application:

```
curl -vs -H "Fastly-Key: $MY_KEY" https://api.fastly.com/service/<service_id>/pool/<pool_id>/servers
```

The response will look like this:

```
[
  {
    "id": "6kEuoknxiaDBCLiAjKqyXq",
    "service_id": "<service_id>",
    "pool_id": "<pool_id>",
    "weight": "100",
    "max_conn": "200",
    "port": "80",
    "address": "<hostname_or_ip_address>",
    "comment": "",
    "created_at": "2016-06-20T08:20:36+00:00",
    "updated_at": "2016-06-20T08:20:36+00:00",
    "deleted_at": null
  }
]
```

## Deleting a server entry

To permanently delete a server entry, make the following API call in a terminal application:

```
curl -vs -H "Fastly-Key: $MY_KEY" -X DELETE https://api.fastly.com/service/<service_id>/pool/<pool_id>/server/<server_id>
```

The response will look like this:

```
{
  "status": "ok"
}
```

**NOTE:** Pools must have at least one server entry. The API won't allow you to delete the last server entry.

• [Guides \(/guides/\)](/guides/) > [Developer's tools \(/guides/devtools/\)](/guides/devtools/) > [Edge Dictionaries \(/guides/edge-dictionaries/\)](/guides/edge-dictionaries/)

## § About Edge Dictionaries (/guides/edge-dictionaries/about-edge-dictionaries)

Fastly offers instantly updatable, global, Edge Dictionaries for use with your services (</guides/basic-setup/working-with-services/>).

### Edge Dictionaries simplify services

Edge Dictionaries are made up of dictionary containers and the dictionary items within them. In combination, containers and items allow you to store data as key-value pairs and turn frequently repeated statements like this:

```
if (something == "value1") {
  set other = "result1";
} else if (something == "value2") {
  set other = "result2";
}
```

into a single function that acts as constant:

```
table <ID> {
  "KEY_STRING": "VALUE_STRING",
  "KEY_STRING2": "VALUE_STRING2",
  ...
}
```

Once you attach a dictionary container to a version of your service with the API and that service is activated, the data in it becomes "versionless." This means you can add to and update (</guides/edge-dictionaries/creating-and-manipulating-dictionary-items/>) the data an Edge Dictionary contains using a single API call at any time after it is created, without ever incrementing a service's version. You can perform lookups on the dictionary items in an Edge Dictionary using functions like `table.lookup(<ID>, "KEY_STRING")` OR `table.lookup(<ID>, "KEY_STRING", "DEFAULT_VALUE_STRING")` in your configuration.

### When Edge Dictionaries might be useful

- Content sharing and social media outlets updating large referer blacklists
- Mobile advertisers validating a key to prevent cache-bust guessing
- Customers authenticating valid user keys at the edge
- Global publishers redirecting users to a specific country site based on geo-location
- Image providers performing token checks for certain objects
- Advertising technology companies blocking bad actors at edge
- Customers deploying user interface versions with simple value change via API

### Limitations and considerations

When creating Edge Dictionaries (</guides/edge-dictionaries/creating-and-using-dictionaries/>), keep the following things in mind as you develop your service configurations:

- **Edge Dictionaries created with custom VCL cannot be manipulated using the API.** If you create a dictionary container using the API (`/api/config#dictionary`), you can use the API to make changes to it, and use custom VCL to interact with it. If you create a dictionary container using custom VCL, that dictionary must always be manipulated via custom VCL.
- **Dictionary containers are limited to 1000 dictionary items.** If you find your dictionary containers approaching this item limit, contact us (<mailto:support@fastly.com>). We may be able to help you figure out an even more efficient way to do things with your Edge Dictionaries.
- **Dictionary item keys are limited 256 characters and values are limited to 8000 characters.** Dictionary items are made up of key-value pairs with character limits. Be sure to take this into account when designing your Edge Dictionaries.
- **Dictionary item keys are case sensitive.** Dictionary item names are case sensitive. Be sure to take this into account when designing your Edge Dictionaries.
- **Audit logs don't exist for Edge Dictionary changes.** If you add, update, or remove a dictionary item, there will be no record of it. The only record of a change will exist when you compare service versions (<https://www.fastly.com/blog/introducing-version-diff>) to view the point at

which the dictionary container was associated with the service version in the first place.

- **When you delete a dictionary container, you'll only delete it from the service version you're editing.** Dictionary containers are tied to versions and can be cloned and reverted. When using Edge Dictionaries, we want you to be able to do things like delete a dictionary container from a current version of your service in order to roll back your configuration to a previous version using as few steps as possible.
- **When you delete a dictionary container, we don't delete the dictionary items inside it.** The dictionary items in a dictionary container are versionless. When you change service versions, we want you to still be able to access the data.
- **Dictionary item deletions are permanent.** Because we don't store data, if you delete a dictionary item, the entry is gone forever from all service versions.

## § Creating and manipulating dictionary items (/guides/edge-dictionaries/creating-and-manipulating-dictionary-items)

A dictionary item is a key-value pair that makes up an entry in a dictionary container in an Edge Dictionary. Once you create an Edge Dictionary (/guides/edge-dictionaries/creating-and-using-dictionaries) and associate the dictionary container with a service, any dictionary items created will appear in your generated VCL.

For example, if you were using Edge Dictionaries to control geoIP redirects, the table would appear similar to this:

```
table geoip_redirect {  
  "GB" : "www.example.co.uk",  
  "IE" : "www.example.co.uk",  
  "IT" : "www.example.com.it",  
  "AU" : "www.example.com.au",  
}
```

If you already have a dictionary container associated with an active version of your service, you can easily add, update, or delete the items in it as long as you know the `dictionary_id`.

In our geoIP example, you would find your `dictionary_id` using the following API call:

```
curl -H 'Fastly-Key: <API_key>'   
https://api.fastly.com/service/<service_id>/version/<version_number>/dictionary/geoip_redirect
```

which would return this response:

```
{  
  "version": <version_number>,  
  "name": "geoip_redirect",  
  "id": "<dictionary_id>",  
  "service_id": "<service_id>"  
}
```

## Adding new items to dictionary

You can add new dictionary items without having to increment your service version number. For example, this API call to a geoIP table to add a new dictionary item:

```
curl -X POST -H 'Fastly-Key: <API_key>' -d 'item_key=NZ&item_value=www.example.com.au' "https://api.fastly.com/service/<service_id>/dictionary/<dictionary_id>/item"
```

returns this response:

```
{  
  "dictionary_id": "<dictionary_id>",  
  "service_id": "<service_id>",  
  "item_key": "NZ",  
  "item_value": "www.example.com.au"  
}
```

The table in the generated VCL would then be updated with the new dictionary item and look like this:

```
table geoip_redirect {  
  "GB" : "www.example.co.uk",  
  "IE" : "www.example.co.uk",  
  "IT" : "www.example.com.it",  
  "AU" : "www.example.com.au",  
  "NZ" : "www.example.com.au",  
}
```

## Upserting dictionary items

You can create and update dictionary items regardless of whether or not they exist. For example, the following API call to the geoIP table to update an existing dictionary item or create it if it doesn't exist:

```
curl -X PUT -H 'Fastly-Key: <API_key>' -d 'item_value=www.example.co.aq' "https://api.fastly.com/service/<service_id>/dictionary/<dictionary_id>/item/AQ"
```

returns this response:

```
{
  "dictionary_id": "<dictionary_id>",
  "item_key": "AQ",
  "item_value": "www.example.co.aq",
  "service_id": "<service_id>"
}
```

The table in the generated VCL would then be updated with the new dictionary item and look like this:

```
table geoip_redirect {
  "GB" : "www.example.co.uk",
  "IE" : "www.example.co.uk",
  "IT" : "www.example.com.it",
  "AU" : "www.example.com.au",
  "NZ" : "www.example.com.au",
  "AQ" : "www.example.co.aq",
}
```

## Updating dictionary items

You can also update any dictionary item without having to increment your service version number. For example, the following API call to the geoIP table to update an existing dictionary item:

```
curl -X PATCH -H 'Fastly-Key: <API_key>' -d 'item_value=www.example.co.uk' "https://api.fastly.com/service/<service_id>/dictionary/<dictionary_id>/item/NZ"
```

returns this response:

```
{
  "dictionary_id": "<dictionary_id>",
  "item_key": "NZ",
  "item_value": "www.example.co.uk",
  "service_id": "<service_id>"
}
```

The table in the generated VCL would then be updated with the new dictionary item and look like this:

```
table geoip_redirect {
  "GB" : "www.example.co.uk",
  "IE" : "www.example.co.uk",
  "IT" : "www.example.com.it",
  "AU" : "www.example.com.au",
  "NZ" : "www.example.co.uk",
  "AQ" : "www.example.co.aq",
}
```

## Batch updating dictionary items

You can update up to 1,000 dictionary items with a single API call. The following actions are available within a batch update:

- **Upsert** - Creates an item if it doesn't exist, otherwise modifies the existing one.
- **Create** - Creates a new item, but will not update an existing one.
- **Update** - Updates an existing item, but will not create a new one if it doesn't exist.
- **Delete** - Permanently deletes the item from the dictionary.

For example, to batch update existing dictionary items in the geoIP table, create a new file called `batch.json` that contains the following JSON-encoded data:

```
{
  "items": [
    {
      "op": "create",
      "item_key": "JP",
      "item_value": "www.example.co.jp"
    },
    {
      "op": "update",
      "item_key": "GB",
      "item_value": "www.example.co.uk"
    },
    {
      "op": "delete",
      "item_key": "IT"
    }
  ]
}
```

Now you can make the following API call:

```
curl -X PATCH -H 'Content-Type: application/json' -H 'Fastly-Key: <API_key>' -d @batch.json "https://api.fastly.com/service/<service_id>/dictionary/<dictionary_id>/items"
```

See the API documentation ([/api/config#dictionary\\_item\\_dc826ce1255a7c42bc48eb204eed8f7f](https://api.fastly.com/api/config#dictionary_item_dc826ce1255a7c42bc48eb204eed8f7f)) for more information.

## Deleting a dictionary item

**⚠ WARNING:** Dictionary item deletions are permanent. Fastly does not store data. If you delete a dictionary item, the entry is gone forever from all versions of your service.

To remove an item from your table, use this API call:

```
curl -X DELETE -H 'Fastly-Key: <API_key>' https://api.fastly.com/service/<service_id>/dictionary/<dictionary_id>/item/NZ
```

Unlike creation and update of dictionary items, the API call returns no response.

## § Creating and using Edge Dictionaries (/guides/edge-dictionaries/creating-and-using-dictionaries)

Edge Dictionaries (</guides/edge-dictionaries/about-edge-dictionaries>) allow you to create logic that doesn't need to be attached to a configuration service version. Edge Dictionaries are made up of dictionary containers and dictionary items. You can use dictionary items to create and store key-value pairs. Attaching dictionary containers to a service version allows you to turn frequently repeated statements into single function statements that act as a constant.

To create an Edge Dictionary and use it within your service you need to:

1. Create an empty dictionary container in a working version of a service that's unlocked and not yet activated.
2. Activate the new version of the service you associated with the empty dictionary container.
3. Add dictionary items to the newly created dictionary container.

Once the dictionary container is created, properly associated, and filled with dictionary items, it can be called in your service.

For example, say you have a referer blacklist that changes frequently and you want to associate it with a service. Any time that service's configuration changes, especially if the configuration rolls back to a previous version, you would want the blacklisted referer domains to continue to remain with the service configuration instead of being removed. Edge Dictionaries would help you do this.

## Create an empty dictionary container within a service

In order to use a dictionary container, start by creating an empty one within an unlocked version of a service.

Before an Edge Dictionary can be manipulated, its dictionary container must be associated with at least one service version that is not locked and not active so that the service becomes aware of the dictionary's existence.

For example, if you were creating a referer blacklist via the API, you would make an API call by running this command:

```
curl -X POST -H 'Fastly-Key: <API_key>' -d 'name=referer_blacklist' https://api.fastly.com/service/<service_id>/version/<version_number>/dictionary
```

which would return:

```
{
  "name": "referrer_blacklist",
  "service_id": "<service_id>",
  "version": <version_number>,
  "id": "<dictionary_id>"
}
```

## Activate the service associated with the dictionary container

In order for an Edge Dictionary to appear in generated VCL so it can be referred to later, the version associated with the dictionary container must be activated.

In our referer blacklist example, you would make this API call to activate service version associated with the empty dictionary container you created:

```
curl -X PUT -H 'Fastly-Key: <API_key>' https://api.fastly.com/service/<service_id>/version/<version_number>/activate
```

The response would be this:

```
{
  "number": <version_number>,
  "active": true,
  "service_id": "<service_id>"
}
```

## Add dictionary items

Once the dictionary container becomes associated with the configuration of a service, you can begin populating it with dictionary items.

For example, you would use the following API call for each URL you want to add a URLs to your referer blacklist:

```
curl -X POST -H 'Fastly-Key: <API_key>' -d 'item_key=example-referer.org&item_value=true' "https://api.fastly.com/service/<service_id>/dictionary/<dictionary_id>/item"
```

The response for each URL added would look similar to this:

```
{
  "dictionary_id": "<dictionary_id>",
  "service_id": "<service_id>",
  "item_key": "example-referer.org",
  "item_value": "true"
}
```

Once the blacklisted URLs are added as items in your dictionary container, you can find them in your generated VCL by looking for a table similar to this:

```
table referer_blacklist {
  "example-referer.org": "true",
  "another-referer.net": "true",
  "sample-referer.com": "true",
}
```

## Using a service to call Edge Dictionaries

When you create Edge Dictionaries via API calls, the dictionary contents aren't tied to any single version of your service.

The logic needed to *interact* with the table of information the Edge Dictionary creates, however, is always tied to a service version.

For example, adding a new referer to your blacklist requires that you specifically interact with the Edge Dictionary at some point after you create it. You could do this via API calls because its data would not require a service version activation. The dictionary was created via API calls not via custom VCL.

Specifically, you would set the host of the referer to a header by including custom VCL like this:

```
// versioned vcl
sub vcl_recv {

    # capture host of referer into a header
    set req.http.Referer-Host = regsub(req.http.Referer, "^https?:/?([^\s/]+).*$", "\1");

    # check if referer host is in blacklisted table
    if (table.lookup(referer_blacklist, req.http.Referer-Host)) {
        # responseObject: forbidden-referrer
        error 900 "Fastly Internal";
    }
    #end condition
}

sub vcl_error {

    if (obj.status == 900) {
        set obj.http.Content-Type = "";
        synthetic {" "};
        return(deliver);
    }
}
```

• [Guides \(/guides/\)](/guides/) > [Developer's tools \(/guides/devtools/\)](/guides/devtools/) > [VCL \(/guides/vcl/\)](/guides/vcl/)

# \$ Accept-Language header VCL features (/guides/vcl/accept-language%20header-vcl-features)

Fastly provides functions in VCL to parse and normalize the `Accept-Language` header.

## Language lookup

An implementation of the Lookup functionality as defined by RFC 4647, section 3.4 (<http://tools.ietf.org/html/rfc4647#section-3.4>).

### Syntax

```
accept.language_lookup(<available languages>, <default>, <priority list>)
```

Argument	Explanation
available languages	A colon-separated list of languages to choose from. Typically the languages that the origin can provide. For example: <code>en:de:fr:pt:es:zh-CN</code>
default	The default language to return if none from the <code>priority list</code> match. For example: <code>en</code>
priority list	The <code>Accept-Language</code> header. A comma-separated list of languages, optionally accompanied by weights (q-values). For example: <code>pt-BR,pt;q=0.8,en-US;q=0.6,en;q=0.4</code>

### Return values

The best matching language (as per the RFC) is returned. If none are found, the default language is returned, unless a weight of zero (`q=0`) was indicated by the priority list, in which case `NULL` is returned.

### Examples

```
set req.http.Normalized-Language =
    accept.language_lookup("en:de:fr:pt:es:zh-CN", "en", req.http.Accept-Language);
```

The above would result in `Normalized-Language: pt` given an `Accept-Language: pt-BR,pt;q=0.8,en-US;q=0.6,en;q=0.4` header.

`accept.language_lookup("en", "nl", "en-GB")` results in `en`, as each subtag is removed and the match retried when a tag does not match.

`accept.language_lookup("en:nl", "nl", "en-GB,nl;q=0.5")` results in `en` still, even if `nl` is a more exact match, because the q-value of `nl` is lower, and that has precedence. Exactness just does not come into the equation.

`accept.language_lookup("en-US:nl", "nl", "en-GB,nl;q=0.5")` results in `nl`, because subtags are not removed from the available languages, only from language tags on the priority list.



`accept.language_lookup("en-US:nl", "nl", "en-us,nl;q=0.5")` results in `en-US`, as the lookup is case insensitive.

`accept.language_lookup("en-US:nl", "nl", "en-GB,nl;q=0")` results in `NULL` (the value, not a string) since `en-GB` and `en` do not match, and `nl` (the default) is listed as unacceptable.

If `q=0` for the default language is to be ignored, the following VCL can be used:

```
set req.http.Normalized-Language =
  accept.language_lookup("en-US:nl", "nl", req.http.Accept-Language);
if (!req.http.Normalized-Language) {
  # User will get Dutch even if he doesn't want it!
  # (Because none of our languages were acceptable)
  set req.http.Normalized-Language = "nl";
}
```

## Language filter (Basic)

An implementation of the Basic Filtering functionality as defined by RFC 4647, section 3.3.1 (<http://tools.ietf.org/html/rfc4647#section-3.3.1>).

The implementation is not exact when the wildcard tag (`*`) is used. If a wildcard is encountered and no matches have been found yet, the default is returned. If there are matches, those are returned and the remainder of the priority list is ignored.

(There is no implementation of Extended Filtering, but if you are in need you could always file a feature request with Support. :))

## Syntax

`accept.language_filter_basic(<available languages>, <default>, <priority list>, <limit>)`

Argument	Explanation
<code>available languages</code>	A colon-separated list of languages choose from. Typically the languages that the origin can provide. For example: <code>en:de:fr:pt:es:zh-CN</code>
<code>default</code>	The default language to return if none from the <code>priority list</code> match. For example: <code>en</code>
<code>priority list</code>	The <code>Accept-Language</code> header. A comma-separated list of languages, optionally accompanied by weights (q-values). For example: <code>pt-BR,pt;q=0.8,en-US;q=0.6,en;q=0.4</code>
<code>limit</code>	The maximum amount of languages returned.

## Return values

The best matching language (as per the RFC) is returned. If none are found, the default language is returned, unless a weight of zero (`q=0`) was indicated by the priority list, in which case `NULL` is returned.

## Examples

```
set req.http.Filtered-Language =
  accept.language_filter_basic("en:de:fr:pt:es:zh-CN", "en", req.http.Accept-Language, 2);
```

The above would result in `Filtered-Language: pt,en` given an `Accept-Language: pt-BR,pt;q=0.8,en-US;q=0.6,en;q=0.4` header.

`accept.language_filter_basic("en", "nl", "en-GB", 2)` results in `en`, as each subtag is removed and the match retried when a tag does not match.

`accept.language_filter_basic("en:nl", "nl", "en-GB,nl;q=0.5", 2)` results in `en,nl`, even if `nl` is a more exact match, because the q-value of `nl` is lower, and that has precedence. Exactness just does not come into the equation.

`accept.language_filter_basic("en-US:nl", "nl", "en-GB,nl;q=0.5", 2)` results in `nl`, because subtags are not removed from the available languages during the search.

`accept.language_filter_basic("en-US:nl", "nl", "en-us,nl;q=0.5", 2)` results in `en-US,nl`, as the lookup is case insensitive.

`accept.language_filter_basic("en-US:nl", "nl", "en-GB,nl;q=0", 2)` results in `NULL` (the value, not a string) since `en-GB` and `en` do not match, and `nl` (the default) is listed as unacceptable.

If `q=0` for the default language is to be ignored, the following VCL can be used:

```
set req.http.Filtered-Language =  
  accept.language_filter_basic("en-US:nl", "nl", req.http.Accept-Language, 2);  
if (!req.http.Filtered-Language) {  
  # User will get Dutch even if he doesn't want it!  
  # (Because none of our languages were acceptable)  
  set req.http.Filtered-Language = "nl";  
}
```

`accept.language_filter_basic("en:nl:de:fr", "nl", "en-GB,*;q=0.5", 2)` results in `en` and `accept.language_filter_basic("en:nl:de:fr", "nl", "*", 2)` results in `nl`.

---

## § Authenticating before returning a request (/guides/vcl/authenticating-before-returning-a-request)

---

Performing authentication before returning a request is possible if your authentication is completely header-based and you do something like the following using custom VCL (</guides/vcl/uploading-custom-vcl>):

```

sub vcl_recv {

  /* unset state tracking header to avoid client sending it */
  if (req.restarts == 0) {
    unset req.http.X-Authed;
  }

  if (!req.http.X-Authed) {
    /* stash the original URL and Host for later */
    set req.http.X-Orig-URL = req.url;

    /* set the URL to what the auth backend expects */
    set req.url = "/authenticate";

    /* Auth requests won't be cached, so pass */
    return(pass);
  }

  if (req.http.X-Authed == "true") {
    /* were authed, so proceed with the request */
    /* reset the URL */
    set req.url = req.http.X-Orig-URL;
  } else {
    /* the auth backend refused the request, so 403 the client */
    error 403;
  }

#FASTLY recv

  ...etc...
}

sub vcl_deliver {

  /* if we are in the auth phase */
  if (!req.http.X-Authed) {

    /* if we got a 5XX from the auth backend, we should fail open */
    if (resp.status >= 500 && resp.status < 600) {
      set req.http.X-Authed = "true";
    }

    if (resp.status == 200) {

      /* the auth backend responded with 200, allow the request and restart */
      set req.http.X-Authed = "true";
    } else if (resp.status == 401) {

      return(deliver);
    } else {

      /* the auth backend responded with non-200, deny the request and restart */
      set req.http.X-Authed = "false";
    }

    restart;
  }

#FASTLY deliver

  ...etc...
}

```

**NOTE:** Be sure to change `/authenticate` to whatever your authentication endpoint is.

If you feel like you can cache the authentication, then add the appropriate headers to the hash in `vcl_hash` and `return(lookup)` instead of `(pass)`.

**IMPORTANT:** The ability to [upload custom VCL code](/guides/vcl/uploading-custom-vcl) (</guides/vcl/uploading-custom-vcl>) is disabled by default when you sign up. To enable this ability for your account, contact [support@fastly.com](mailto:support@fastly.com) (<mailto:support@fastly.com>) and request it.

## § Conditionally changing a URL (/guides/vcl/conditionally-changing-a-url)

To conditionally change a URL based on the domain, include VCL that looks something like this:

```
if (req.http.host ~ "^restricted") {  
    set req.url = "/sanitized" req.url;  
}
```

If you have shielding enabled, however, add the following code instead to avoid rewriting the URL twice:

```
if (req.http.host ~ "^restricted" && req.url !~ "^/sanitized") {  
    set req.url = "/sanitized" req.url;  
}
```

In Fastly's web interface, this VCL would be the equivalent of creating a new Header (/guides/basic-configuration/adding-or-modifying-headers-on-http-requests-and-responses):

### Create a new header

Learn more about this section in our [headers tutorial](#).

**CONDITION** This will happen all the time unless you [Attach a condition](#)

Type / Action	<div>Request</div>	<div>Set</div>
The type of header and the action performed on it.		
Destination	<div>url</div>	★ Required
The name of the header that will be affected by the selected action. For example: <b>http.Content-Type</b> , <b>http.Set-Cookie</b> , <b>http.Via</b> , <b>http.Location</b> , or <b>http.Access-Control-Allow-Origin</b> .		
Source	<div>"/sanitized" req.url</div>	★ Required
New content for the header. Can be a static value (e.g. string or number) or a dynamic value (e.g. existing header or a GeolIP value). Please use quotes for string values.		
Ignore if set	<div>No</div>	
If switched to <b>Yes</b> , the action will not be performed if the header in Destination exists.		
Priority	<div>10</div>	★ Required
The order in which the header rules execute within the condition.		
Description	<div>Restricted URL</div>	★ Required
The name of your header, such as <b>My header</b> .		

CREATE

CANCEL

and then creating a request condition that restricts connections to that host:

Create a new request condition.

Learn the basics in our [conditions tutorial](#)

Type

Request

Learn more about the [different types of conditions](#).

Name

Restricted Host

★

 Required

Apply if...

req.http.host ~ "^restricted"

The expression to decide whether this is run.

▶

 Examples

[Adjust priority \(default is 10\)](#)

This is an advanced option. For most configurations, the default is best.

SAVE AND APPLY TO RESTRICTED URL

CANCEL

## § Creating location-based tagging (/guides/vcl/creating-location-based-tagging)

You can set custom HTTP headers in your varnish configuration (VCL) based on the variables we expose. Use the GeoIP features we have built into Varnish to create location-based tagging. We provide a list of geographic information based on a client's IP address. For a complete list of available GeoIP variables, read about which GeoIP features are accessible to VCL (/guides/vcl/geoip-related-vcl-features).

In the example below, an HTTP header Fastly-GeoIP-CountryCode is created with the two letter country code of the client's IP address.

```
sub vcl_recv {
  ...
  set req.http.fastly-GeoIP-CountryCode = geoip.country_code;
  ...
}
```

🔔

**IMPORTANT:** The ability to [upload custom VCL code](#) (/guides/vcl/uploading-custom-vcl) is disabled by default when you sign up. To enable this ability for your account, contact [support@fastly.com](mailto:support@fastly.com) (mailto:support@fastly.com) and request it.

## § Cryptographic- and hashing-related VCL functions (/guides/vcl/cryptographic-and-hashing-related-vcl-functions)

Fastly provides several functions in VCL (/guides/vcl/) for cryptographic- and hashing-related purposes. It is based very heavily on Kristian Lyngstøl's digest vmod (<https://github.com/varnish/libvmod-digest>) for Varnish 3 (which means you can also refer to that documentation for more detail).

### Functions

Function	Description
<code>digest.hmac_md5(&lt;key&gt;,&lt;message&gt;)</code>	Hash-based message authentication code ( <a href="https://en.wikipedia.org/wiki/Hash-based_message_authentication_code">https://en.wikipedia.org/wiki/Hash-based_message_authentication_code</a> ) using MD5. Returns a hex-encoded string prepended with 0x.

Function	Description
<code>digest.hmac_md5_base64(&lt;key&gt;, &lt;message&gt;)</code>	Hash-based message authentication code ( <a href="https://en.wikipedia.org/wiki/Hash-based_message_authentication_code">https://en.wikipedia.org/wiki/Hash-based_message_authentication_code</a> ) using MD5. Returns a base64-encoded ( <a href="https://en.wikipedia.org/wiki/Base64">https://en.wikipedia.org/wiki/Base64</a> ) string.
<code>digest.hmac_sha1(&lt;key&gt;, &lt;message&gt;)</code>	Hash-based message authentication code ( <a href="https://en.wikipedia.org/wiki/Hash-based_message_authentication_code">https://en.wikipedia.org/wiki/Hash-based_message_authentication_code</a> ) using SHA1 ( <a href="https://en.wikipedia.org/wiki/Secure_Hash_Algorithm">https://en.wikipedia.org/wiki/Secure_Hash_Algorithm</a> ). Returns a hex-encoded string prepended with 0x.
<code>digest.hmac_sha1_base64(&lt;key&gt;, &lt;message&gt;)</code>	Hash-based message authentication code ( <a href="https://en.wikipedia.org/wiki/Hash-based_message_authentication_code">https://en.wikipedia.org/wiki/Hash-based_message_authentication_code</a> ) using SHA1 ( <a href="https://en.wikipedia.org/wiki/Secure_Hash_Algorithm">https://en.wikipedia.org/wiki/Secure_Hash_Algorithm</a> ). Returns a base64-encoded ( <a href="https://en.wikipedia.org/wiki/Base64">https://en.wikipedia.org/wiki/Base64</a> ) string.
<code>digest.hmac_sha256(&lt;key&gt;, &lt;message&gt;)</code>	Hash-based message authentication code ( <a href="https://en.wikipedia.org/wiki/Hash-based_message_authentication_code">https://en.wikipedia.org/wiki/Hash-based_message_authentication_code</a> ) using SHA256 ( <a href="https://en.wikipedia.org/wiki/Secure_Hash_Algorithm">https://en.wikipedia.org/wiki/Secure_Hash_Algorithm</a> ). Returns a hex-encoded string prepended with 0x.
<code>digest.hmac_sha256_base64(&lt;key&gt;, &lt;message&gt;)</code>	Hash-based message authentication code ( <a href="https://en.wikipedia.org/wiki/Hash-based_message_authentication_code">https://en.wikipedia.org/wiki/Hash-based_message_authentication_code</a> ) using SHA256 ( <a href="https://en.wikipedia.org/wiki/Secure_Hash_Algorithm">https://en.wikipedia.org/wiki/Secure_Hash_Algorithm</a> ). Returns a base64-encoded ( <a href="https://en.wikipedia.org/wiki/Base64">https://en.wikipedia.org/wiki/Base64</a> ) string.
<code>digest.base64(&lt;string&gt;)</code>	Base64 ( <a href="https://en.wikipedia.org/wiki/Base64">https://en.wikipedia.org/wiki/Base64</a> ) encoding. Returns the base64-encoded version of the input-string.
<code>digest.base64url(&lt;string&gt;)</code>	Base64 ( <a href="https://en.wikipedia.org/wiki/Base64#URL_applications">https://en.wikipedia.org/wiki/Base64#URL_applications</a> ) encoding. Returns the base64-encoded version of the input-string. Replaces +/ with -_ for url safety.
<code>digest.base64url_nopad(&lt;string&gt;)</code>	Base64 ( <a href="https://en.wikipedia.org/wiki/Base64#URL_applications">https://en.wikipedia.org/wiki/Base64#URL_applications</a> ) encoding. Returns the base64-encoded version of the input-string. Replaces +/ with -_ for url safety. Has no length padding ( <a href="https://en.wikipedia.org/wiki/Base64#Padding">https://en.wikipedia.org/wiki/Base64#Padding</a> ).
<code>digest.base64_decode(&lt;string&gt;)</code>	Decode Base64. Returns a string.
<code>digest.base64url_decode(&lt;string&gt;)</code>	Decode Base64 with url safe characters in. Returns a string.
<code>digest.base64url_nopad_decode(&lt;string&gt;)</code>	Decode Base64 with url safe characters. Returns a string. Identical to <code>base64_url_decode</code> .
<code>digest.hash_sha1(&lt;string&gt;)</code>	Use the SHA1 ( <a href="https://en.wikipedia.org/wiki/Secure_Hash_Algorithm">https://en.wikipedia.org/wiki/Secure_Hash_Algorithm</a> ) hash. Returns a hex-encoded string.
<code>digest.hash_sha224(&lt;string&gt;)</code>	Use the SHA224 ( <a href="https://en.wikipedia.org/wiki/Secure_Hash_Algorithm">https://en.wikipedia.org/wiki/Secure_Hash_Algorithm</a> ) hash. Returns a hex-encoded string.
<code>digest.hash_sha256(&lt;string&gt;)</code>	Use the SHA256 ( <a href="https://en.wikipedia.org/wiki/Secure_Hash_Algorithm">https://en.wikipedia.org/wiki/Secure_Hash_Algorithm</a> ) hash. Returns a hex-encoded string.
<code>digest.hash_sha384(&lt;string&gt;)</code>	Use the SHA384 ( <a href="https://en.wikipedia.org/wiki/Secure_Hash_Algorithm">https://en.wikipedia.org/wiki/Secure_Hash_Algorithm</a> ) hash. Returns a hex-encoded string.
<code>digest.hash_sha512(&lt;string&gt;)</code>	Use the SHA512 ( <a href="https://en.wikipedia.org/wiki/Secure_Hash_Algorithm">https://en.wikipedia.org/wiki/Secure_Hash_Algorithm</a> ) hash. Returns a hex-encoded string.
<code>digest.hash_gost(&lt;string&gt;)</code>	Use the Gost ( <a href="https://en.wikipedia.org/wiki/GOST_(hash_function)">https://en.wikipedia.org/wiki/GOST_(hash_function)</a> ) hash function. Returns a hex-encoded string.
<code>digest.hash_md2(&lt;string&gt;)</code>	Use the MD2 ( <a href="https://en.wikipedia.org/wiki/MD2_(cryptography)">https://en.wikipedia.org/wiki/MD2_(cryptography)</a> ) hash. Returns a hex-encoded string.
<code>digest.hash_md4(&lt;string&gt;)</code>	Use the MD4 ( <a href="https://en.wikipedia.org/wiki/MD4">https://en.wikipedia.org/wiki/MD4</a> ) hash. Returns a hex-encoded string.
<code>digest.hash_md5(&lt;string&gt;)</code>	Use the MD5 ( <a href="https://en.wikipedia.org/wiki/MD5">https://en.wikipedia.org/wiki/MD5</a> ) hash. Returns a hex-encoded string.
<code>digest.hash_crc32(&lt;string&gt;)</code>	Use a 32 bit Cyclic Redundancy Checksum ( <a href="https://en.wikipedia.org/wiki/CRC32">https://en.wikipedia.org/wiki/CRC32</a> ). Returns a hex-encoded string.
<code>digest.hash_crc32b(&lt;string&gt;)</code>	A reversed CRC32 (for compatibility with some PHP applications ( <a href="http://php.net/manual/en/function.hash-file.php#104836">http://php.net/manual/en/function.hash-file.php#104836</a> )). Returns a hex-encoded string.

Function	Description
<code>digest.hash_adler32(&lt;string&gt;)</code>	Use the Adler32 ( <a href="https://en.wikipedia.org/wiki/Adler32">https://en.wikipedia.org/wiki/Adler32</a> ) checksum algorithm. Returns a hex-encoded string.
<code>digest.hash_haval128(&lt;string&gt;)</code>	Use the 128 bit Haval ( <a href="https://en.wikipedia.org/wiki/HAVAL">https://en.wikipedia.org/wiki/HAVAL</a> ) hash function. Returns a hex-encoded string.
<code>digest.hash_haval160(&lt;string&gt;)</code>	Use the 160 bit Haval ( <a href="https://en.wikipedia.org/wiki/HAVAL">https://en.wikipedia.org/wiki/HAVAL</a> ) hash function. Returns a hex-encoded string.
<code>digest.hash_haval192(&lt;string&gt;)</code>	Use the 192 bit Haval ( <a href="https://en.wikipedia.org/wiki/HAVAL">https://en.wikipedia.org/wiki/HAVAL</a> ) hash function. Returns a hex-encoded string.
<code>digest.hash_haval224(&lt;string&gt;)</code>	Use the 224 bit Haval ( <a href="https://en.wikipedia.org/wiki/HAVAL">https://en.wikipedia.org/wiki/HAVAL</a> ) hash function. Returns a hex-encoded string.
<code>digest.hash_haval256(&lt;string&gt;)</code>	Use the 256 bit Haval ( <a href="https://en.wikipedia.org/wiki/HAVAL">https://en.wikipedia.org/wiki/HAVAL</a> ) hash function. Returns a hex-encoded string.
<code>digest.hash_ripemd128(&lt;string&gt;)</code>	Use the 128 bit RIPEMD ( <a href="https://en.wikipedia.org/wiki/RIPEMD">https://en.wikipedia.org/wiki/RIPEMD</a> ) hash function. Returns a hex-encoded string.
<code>digest.hash_ripemd160(&lt;string&gt;)</code>	Use the 160 bit RIPEMD ( <a href="https://en.wikipedia.org/wiki/RIPEMD">https://en.wikipedia.org/wiki/RIPEMD</a> ) hash function. Returns a hex-encoded string.
<code>digest.hash_ripemd256(&lt;string&gt;)</code>	Use the 256 bit RIPEMD ( <a href="https://en.wikipedia.org/wiki/RIPEMD">https://en.wikipedia.org/wiki/RIPEMD</a> ) hash function. Returns a hex-encoded string.
<code>digest.hash_ripemd320(&lt;string&gt;)</code>	Use the 320 bit RIPEMD ( <a href="https://en.wikipedia.org/wiki/RIPEMD">https://en.wikipedia.org/wiki/RIPEMD</a> ) hash function. Returns a hex-encoded string.
<code>digest.hash_tiger(&lt;string&gt;)</code>	Use the standard, 192 bits Tiger ( <a href="https://en.wikipedia.org/wiki/Tiger_(cryptography)">https://en.wikipedia.org/wiki/Tiger_(cryptography)</a> ) hash function. Returns a hex-encoded string.
<code>digest.hash_tiger128(&lt;string&gt;)</code>	Use the truncated 128 bit Tiger ( <a href="https://en.wikipedia.org/wiki/Tiger_(cryptography)">https://en.wikipedia.org/wiki/Tiger_(cryptography)</a> ) hash function. Returns a hex-encoded string.
<code>digest.hash_tiger160(&lt;string&gt;)</code>	Use the truncated 160 bit Tiger ( <a href="https://en.wikipedia.org/wiki/Tiger_(cryptography)">https://en.wikipedia.org/wiki/Tiger_(cryptography)</a> ) hash function. Returns a hex-encoded string.
<code>digest.hash_snefru128(&lt;string&gt;)</code>	Use the 256 bit Snefru ( <a href="https://en.wikipedia.org/wiki/Snefru">https://en.wikipedia.org/wiki/Snefru</a> ) hash function. Returns a hex-encoded string.
<code>digest.hash_snefru256(&lt;string&gt;)</code>	Use the 128 bit Snefru ( <a href="https://en.wikipedia.org/wiki/Snefru">https://en.wikipedia.org/wiki/Snefru</a> ) hash function. Returns a hex-encoded string.
<code>digest.hash_whirlpool(&lt;string&gt;)</code>	Use the Whirlpool ( <a href="https://en.wikipedia.org/wiki/Whirlpool_(cryptography)">https://en.wikipedia.org/wiki/Whirlpool_(cryptography)</a> ) hash function. Returns a hex-encoded string.
<code>digest.rsa_verify(ID hash_method, STRING_LIST public_key, STRING_LIST payload, STRING_LIST digest [, ID base64_method ])</code>	A boolean function that returns true if the RSA digest using <code>public_key</code> of <code>payload</code> matches <code>digest</code> . The <code>hash_method</code> parameter selects the digest function to use. It can be <code>sha256</code> , <code>sha384</code> , <code>sha512</code> , or <code>default</code> ( <code>default</code> is equivalent to <code>sha256</code> ). The <code>base64_method</code> parameter is optional. It can be <code>standard</code> , <code>url</code> , <code>url_nopad</code> , or <code>default</code> ( <code>default</code> is equivalent to <code>url_nopad</code> ).
<code>digest.secure_is_equal(STRING_LIST s1, STRING_LIST s2)</code>	A boolean function that returns true if s1 and s2 are equal. The comparison is done in constant time to defend against timing attacks.
<code>digest.time_hmac_md5(&lt;base64 encoded key&gt;, &lt;interval&gt;, &lt;offset&gt;)</code>	Time based One Time Password using MD5. Returns base64 encoded output.
<code>digest.time_hmac_sha1(&lt;base64 encoded key&gt;, &lt;interval&gt;, &lt;offset&gt;)</code>	Time based One Time Password using SHA1. Returns base64 encoded output.
<code>digest.time_hmac_sha256(&lt;base64 encoded key&gt;, &lt;interval&gt;, &lt;offset&gt;)</code>	Time based One Time Password using SHA256. Returns base64 encoded output.

## Notes

In base64 decoding, the output theoretically could be in binary but is interpreted as a string. So if the binary output contains '\0' then it could be truncated.

The time based One-Time Password algorithm initializes the HMAC using the key and appropriate hash type. Then it hashes the message

```
(<time now in seconds since UNIX epoch> / <interval>) + <offset>
```

as a 64bit unsigned integer (little endian) and base64 encodes the result.

## Examples

### One-Time Password Validation (Token Authentication)

Use this to validate tokens with a URL format like the following:

```
http://cname-to-fastly/video.mp4?6h2YU1lCB4C50SbkZ0E6U3dZGjh+84dz3+Zope2UhiK=
```

Example implementations for token generation in various languages can be found in GitHub (<https://github.com/fastly/token-functions>).

#### Example VCL

```
sub vcl_recv {
  #FASTLY recv

  set req.http.token = regsub(req.url, ".*\?(.*)$", "\1");
  if (req.http.token != digest.time_hmac_sha256("RmFzdGx5IFRva2VuIFRlc3Q=", 60, 0) &&
      req.http.token != digest.time_hmac_sha256("RmFzdGx5IFRva2VuIFRlc3Q=", 60, -1)) {
    error 403;
  }
}
```

### Signature

```
set resp.http.x-data-sig = digest.hmac_sha256("secretkey",resp.http.x-data);
```

### Base64 decoding

A snippet like this in `vcl_error` would set the response body to the value of the request header field named `x-parrot` after base64-decoding the value:

```
synthetic digest.base64_decode(req.http.x-parrot);
```

However, if the base64-decoded string contains a NUL byte (0x00), then that byte and any bytes following it will not be included in the response. Keep that in mind if you intend to send a synthetic response that contains binary data. There is currently no way to send a synthetic response containing a NUL byte.

## § Custom responses that don't hit origin servers (/guides/vcl/custom-responses-that-dont-hit-origin-servers)

Fastly can send custom responses for certain requests that you don't want to hit your origin servers. For example, if you wanted to restrict caching to a URL subtree that contains images and scripts, you could configure Fastly to return an `HTTP 404 Not Found` response to requests for anything other than `/Content/*` or `/Scripts/*`. To illustrate how to implement this example, we'll show you how to create a response and corresponding request condition.

### Creating the response

Follow these instructions to create a new response for your service:

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.
4. Click the **Content** tab. The Content page appears.



DOMAINS  
1 Domain

ORIGINS  
1 Host  
1 Health check

SETTINGS  
1 Request setting  
0 Cache settings

CONTENT  
0 Headers  
0 Gzips  
0 Responses

LOGGING  
1 Endpoint

CUSTOM VCL  
0 Files

## Headers

*There are no headers.*

CREATE YOUR FIRST HEADER

## Gzip

*Gzip is not enabled.*

SET UP GZIP

## Responses

*There are no responses.*

CREATE YOUR FIRST RESPONSE

### About content

Headers allow you to set, delete, and change request and response headers as your information is handled by Fastly.

Responses allow you to create synthetic responses that exist entirely on the cache servers.

5. Click the **Create response** button. The Create a new synthetic response page appears.

## Create a new synthetic response

More on how synthetic responses work in our [tutorial](#).

CONDITION

This will happen all the time unless you [Attach a condition](#)

Status

200 OK

The HTTP status code to include in the header of the response.

MIME Type

The media type to put into the Content-Type header which informs the browser how to display the response. Common MIME types are **application/json**, **text/plain**, **text/html**.

Response

The content to be served when delivering the response.

Description

★ Required

The name of your response, such as **My response**.

CREATE

CANCEL

6. Fill out the **Create a new synthetic response** fields as follows:

- From the **Status** menu, select an HTTP code to return to the client. For example, `404 Not Found`.
- In the **MIME Type** field, type the MIME type of the response. For example, `text/html`.
- In the **Response** field, type the plaintext or HTML content to return to the client. For example `Page not found`.
- In the **Description** field, type a human-readable name for the response. For example `Return Not Found`.

7. Click the **Create** button to create the response. The new response appears in the Responses area of the Content tab.

Responses

CREATE RESPONSE

Return Not Found


Attach a condition Delete

Status: 404 (text/html)

## Creating the request condition


Follow these instructions to attach a request condition to the response you just created:

1. Click the **Attach a condition** link next to the response that you just created. The Add a condition window appears.



The dialog box has a title bar with a close button (X). The main heading is "Add a condition to Return Not Found" with a subtitle "Learn the basics in our [conditions tutorial](#)". Below this is a search bar with the placeholder text "Select a condition (9 total)". In the center, there is a blue button labeled "CREATE A NEW CONDITION". At the bottom right, there is a "CANCEL" button.

2. Click the **Create a new condition** button. The Create a new condition window appears.



The dialog box has a title bar with a close button (X). The main heading is "Create a new condition" with a subtitle "Learn the basics in our [conditions tutorial](#)". Below this are three input fields: "Type" with a dropdown menu showing "Request", "Name" with a text box containing "Return Not Found" and an orange asterisk icon, and "Apply if..." with a text box containing "! ( req.url ~ \"^/(Content|Scripts)/\" )". Below the "Apply if..." field is a note: "The expression to decide whether this is run." followed by a link "Examples". Below that is a link "Adjust priority (default is 10)" and a note: "This is an advanced option. For most configurations, the default is best." At the bottom, there is a blue button labeled "SAVE AND APPLY TO RETURN NOT FOUND" and a "CANCEL" button.

3. Fill out the **Create a new condition** fields as follows:

- From the **Type** menu, select **Request**.
- In the **Name** field, type a human-readable name for the condition. For example, `Return Not Found`.
- In the **Apply if** field, type the request condition you want inserted into a VCL if statement. For example, `! ( req.url ~ \"^/(Content|Scripts)/\" )`. See below for more examples of request conditions.

4. Click the **Save and apply to** button. The Responses area now displays the condition that must be met in order for your response to begin being used.

Responses

CREATE RESPONSE

IF Return Not Found

! ( req.url ~ "^/(Content|Scripts)/" )

THEN

Return Not Found

Detach from conditionDelete

Status: 404 (text/html)

5. Click the **Activate** button to deploy your configuration changes.

## Example request conditions

Respond only if URLs don't match a certain mask, in this case `/Content/*` or `/Scripts/*`:

```
! (req.url ~ "^/(Content|Scripts)/")
```

Respond only if URLs match `/secret/*` or are Microsoft Word or Excel documents (`*.doc` and `*.xls` file extensions):

```
! (req.url ~ "^/secret/" || req.url ~ "\.(xls|doc)$")
```

Ignore POST and PUT HTTP requests:

```
req.request == "POST" || req.request == "PUT"
```

Deny a spider or crawler using `user-agent` `"annoying_robot"`:

```
req.http.user-agent ~ "annoying_robot"
```

Prevent a specific IP from connecting, in this case the IP `225.0.0.1`:

```
client.ip == "225.0.0.1"
```

Use geographic variables (`/guides/vcl/geoip-related-vcl-features`) to block traffic from a specific location (e.g., China):

```
geoip.country_code == "CN"
```

Match the `client.ip` against a CIDR range, such as `240.24.0.0/16` (this requires first creating an ACL object in VCL (`/guides/vcl/using-access-control-lists`)):

```
client.ip ~ ipRangeObject
```

## § Date- and time-related VCL features (`/guides/vcl/date-and-time-related-vcl-features`)

By default VCL includes the `now` variable, which provides the current time (for example, `Wed, 17 Sep 2025 23:19:06 GMT`). Fastly adds several new Varnish variables and functions (`/guides/vcl/`) that allow more flexibility when dealing with dates and times.

### Variables

The following variables have been added:

Name	Location	Description
<code>now</code>	all	The current time in RFC 1123 format ( <a href="https://www.ietf.org/rfc/rfc1123.txt">https://www.ietf.org/rfc/rfc1123.txt</a> ) format.
<code>now.sec</code>	all	Like the <code>now</code> variable, but in seconds since the UNIX Epoch ( <a href="https://en.wikipedia.org/wiki/Unix_time">https://en.wikipedia.org/wiki/Unix_time</a> ).
<code>time.start</code>	all	The time the request started, after TLS termination, using RFC 1123 format ( <a href="https://www.ietf.org/rfc/rfc1123.txt">https://www.ietf.org/rfc/rfc1123.txt</a> ).

Name	Location	Description
<code>time.start.sec</code>	all	The time the request started in seconds since the UNIX Epoch ( <a href="https://en.wikipedia.org/wiki/Unix_time">https://en.wikipedia.org/wiki/Unix_time</a> ), after TLS termination.
<code>time.start.msec</code>	all	The time the request started in milliseconds since the UNIX Epoch ( <a href="https://en.wikipedia.org/wiki/Unix_time">https://en.wikipedia.org/wiki/Unix_time</a> ), after TLS termination.
<code>time.start.msec_frac</code>	all	The time the request started in milliseconds since the last whole second, after TLS termination.
<code>time.start.usec</code>	all	The time the request started in microseconds since the UNIX Epoch ( <a href="https://en.wikipedia.org/wiki/Unix_time">https://en.wikipedia.org/wiki/Unix_time</a> ), after TLS termination.
<code>time.start.usec_frac</code>	all	The time the request started in microseconds since the last whole second, after TLS termination.
<code>time.end</code>	deliver, log	The time the request ended, using RFC 1123 format ( <a href="https://www.ietf.org/rfc/rfc1123.txt">https://www.ietf.org/rfc/rfc1123.txt</a> ). Also useful for <code>strftime</code> .
<code>time.end.sec</code>	deliver, log	The time the request ended in seconds since the UNIX Epoch ( <a href="https://en.wikipedia.org/wiki/Unix_time">https://en.wikipedia.org/wiki/Unix_time</a> ).
<code>time.end.msec</code>	deliver, log	The time the request ended in milliseconds since the UNIX Epoch ( <a href="https://en.wikipedia.org/wiki/Unix_time">https://en.wikipedia.org/wiki/Unix_time</a> ).
<code>time.end.msec_frac</code>	deliver, log	The time the request started in milliseconds since the last whole second.
<code>time.end.usec</code>	deliver, log	The time the request ended in microseconds since the UNIX Epoch ( <a href="https://en.wikipedia.org/wiki/Unix_time">https://en.wikipedia.org/wiki/Unix_time</a> ).
<code>time.end.usec_frac</code>	deliver, log	The time the request started in microseconds since the last whole second.
<code>time.elapsed</code>	deliver, log	The time since the request start, using RFC 1123 format ( <a href="https://www.ietf.org/rfc/rfc1123.txt">https://www.ietf.org/rfc/rfc1123.txt</a> ). Also useful for <code>strftime</code> .
<code>time.elapsed.sec</code>	deliver, log	The time since the request start in seconds.
<code>time.elapsed.msec</code>	deliver, log	The time since the request start in milliseconds.
<code>time.elapsed.msec_frac</code>	deliver, log	The time the request started in milliseconds since the last whole second.
<code>time.elapsed.usec</code>	deliver, log	The time since the request start in microseconds.
<code>time.elapsed.usec_frac</code>	deliver, log	The time the request started in microseconds since the last whole second
<code>time.to_first_byte</code>	deliver, log	The time interval since the request started up to the point before the <code>vcl_deliver</code> function ran. When used in a string context, an RTIME variable like this one will be formatted as a number in seconds with 3 decimal digits of precision. In <code>vcl_deliver</code> this interval will be very close to <code>time.elapsed</code> . In <code>vcl_log</code> , the difference between <code>time.elapsed</code> and <code>time.to_first_byte</code> will be the time that it took to send the response body.

## Functions

The following functions have been added:

Name	Description
<code>time.hex_to_time(&lt;divider&gt;, &lt;hextime&gt;)</code>	Takes a hexadecimal string value, divides by <code>&lt;divider&gt;</code> and interprets the result as seconds since UNIX Epoch ( <a href="https://en.wikipedia.org/wiki/Unix_time">https://en.wikipedia.org/wiki/Unix_time</a> ).
<code>time.add(&lt;time&gt;, &lt;offset&gt;)</code>	Adds <code>offset</code> to <code>time</code> .
<code>time.sub(&lt;time&gt;, &lt;offset&gt;)</code>	Subtracts <code>offset</code> from <code>time</code> .
<code>time.is_after(&lt;time1&gt;, &lt;time2&gt;)</code>	Returns <code>TRUE</code> if <code>time1</code> is after <code>time2</code> . (Normal timeflow and causality required.)
<code>strftime(&lt;format&gt;, &lt;time&gt;)</code>	Formats a time to a string. This uses standard POSIX <code>strftime</code> formats ( <a href="http://www.unix.com/man-page/FreeBSD/3/strftime/">http://www.unix.com/man-page/FreeBSD/3/strftime/</a> ).

Name	Description
<code>std.time(&lt;string_to_parse&gt;, &lt;fallback_value&gt;)</code>	Converts a string to a time variable. The following string formats are supported: <code>Sun, 06 Nov 1994 08:49:37 GMT</code> , <code>Sunday, 06-Nov-94 08:49:37 GMT</code> , <code>Sun Nov 6 08:49:37 1994</code> , <code>784111777.00</code> , <code>784111777</code> . Useful because time variables are needed as arguments for functions like <code>time.add</code> and <code>strftime</code> .
<code>std.integer2time(&lt;seconds_since_epoch&gt;)</code>	Converts an integer to a time variable. To use a string, use <code>std.atoi</code> .

## Examples

The following examples illustrate how to use the variables and functions.

★ **TIP:** Regular strings ("short strings") in VCL use `%xx` escapes (percent encoding) for special characters, which would conflict with the `%` used in the `strftime` format. For the `strftime` examples, we use VCL "long strings" `{ "... }`, which do not use the `%xx` escapes. Alternatively, you could use `%25` for each `%`.

### time.hex\_to\_time, time.add, and time.is\_after

```
if (time.is_after(time.add(now, 10m), time.hex_to_time(1, "d0542d8"))) {
  ...
}
```

### strftime

```
set resp.http.Now = strftime({"%Y-%m-%d %H:%M"}, now)

set resp.http.Start = strftime({"%a, %d %b %Y %T %z"}, time.start)
```

### std.time and std.integer2time

```
set resp.http.X-Seconds-Since-Modified = strftime({"%s"}, time.sub(now, std.time(resp.http.Last-Modified, now)));

std.integer2time(std.atoi("1445445162"));
```

Comparison operators like `>` `<` `>=` `<=` `==` `!=` do not work with `std.time` or `std.integer2time`. Instead, you can compare two times using something similar to this:

```
if (time.is_after(now, std.integer2time(std.atoi("1445445162")))) {
  # do something
}
```

## § Delivering different content to different devices (/guides/vcl/delivering-different-content-to-different-devices)

The easiest way to deliver different content based on the device being used is to rewrite the URL of the request based on what the user agent is. We've written an article that describes how to change the URL based on conditions (</guides/vcl/conditionally-changing-a-url>) using our user interface but in pure VCL it would look something like this:

```
sub vcl_recv {
  if (req.http.User-Agent ~ "(?i)ip(hone|od)") {
    set req.url = "/mobile" req.url;
  } elseif (req.http.User-Agent ~ "(?i)ipad") {
    set req.url = "/tablet" req.url;
  }
  #FASTLY recv
}
```

Obviously the code fragment above doesn't contain a comprehensive list of mobile and tablet devices. Google has an official blog post (<https://webmasters.googleblog.com/2011/03/mo-better-to-also-detect-mobile-user.html>) on detecting Android mobile versus tablet and this VCL fragment (<https://github.com/varnishcache/varnish-devicedetect/blob/master/devicedetect.vcl>) from Varnish Software can detect several different types of devices quite reliably, although it doesn't include Windows mobile and tablet, BlackBerry Playbook, and the Kindle user agents.

The most comprehensive device detection routine we've seen so far is this one:

```
# based on https://github.com/varnish/varnish-devicedetect/blob/master/devicedetect.vcl
sub detect_device {
    unset req.http.X-UA-Device;
    unset req.http.X-UA-Vendor;

    set req.http.X-UA-Device = "desktop";
    set req.http.X-UA-Vendor = "generic";

    # Handle that a cookie or url param may override the detection altogether
    if (req.url ~ "[&|?]device_force=([^\s]+)") {
        set req.http.X-UA-Device = regsub(req.url, ".*[&|?]device_force=([^\s]+).*", "\1");
    } elseif (req.http.Cookie ~ "(?i)X-UA-Device-force") {
        # ;?? means zero or one ;, non-greedy to match the first
        set req.http.X-UA-Device = regsub(req.http.Cookie, "(?i).*X-UA-Device-force=([^;]+)??.*", "\1");
        # Clean up our mess in the cookie header
        set req.http.Cookie = regsuball(req.http.Cookie, "(^|; )X-UA-Device-force=^[^;]+? ", "\1");
        # If the cookie header is now empty, or just whitespace, unset it
        if (req.http.Cookie ~ "^ *$") { unset req.http.Cookie; } # "$ # stupid syntax highlighter
    } else {
        if (req.http.User-Agent ~ "(?i)(ads|google|bing|msn|yandex|baidu|ro|career|)bot" ||
            req.http.User-Agent ~ "(?i)(baidu|jike|symantec)spider" ||
            req.http.User-Agent ~ "(?i)scanner" ||
            req.http.User-Agent ~ "(?i)(web)crawler") {
            set req.http.X-UA-Device = "bot";
        } elseif (req.http.User-Agent ~ "(?i)ipad") {
            set req.http.X-UA-Device = "tablet";
            set req.http.X-UA-Vendor = "apple";
        } elseif (req.http.User-Agent ~ "(?i)ip(hone|od)") {
            set req.http.X-UA-Device = "smartphone";
            set req.http.X-UA-Vendor = "apple";
        }
        # how do we differ between an android phone and an android tablet?
        # http://stackoverflow.com/questions/5341637/how-do-detect-android-tablets-in-general-useragent
        # http://googlewebmastercentral.blogspot.com/2011/03/mo-better-to-also-detect-mobile-user.html
        } elseif (req.http.User-Agent ~ "(?i)android.*(mobile|mini)") {
            set req.http.X-UA-Device = "smartphone";
            set req.http.X-UA-Vendor = "android";
        }
        # android 3/honeycomb was just about tablet-only, and any phones will probably handle a bigger page layout
        } elseif (req.http.User-Agent ~ "(?i)android") {
            set req.http.X-UA-Device = "tablet";
            set req.http.X-UA-Vendor = "android";
        }
        # see http://my.opera.com/community/openweb/idopera/
        } elseif (req.http.User-Agent ~ "Opera Mobi") {
            set req.http.X-UA-Device = "smartphone";
            set req.http.X-UA-Vendor = "android";
        }
        } elseif (req.http.User-Agent ~ "PlayBook; U; RIM Tablet") {
            set req.http.X-UA-Device = "tablet";
            set req.http.X-UA-Vendor = "blackberry";
        }
        } elseif (req.http.User-Agent ~ "hp-tablet.*TouchPad") {
            set req.http.X-UA-Device = "tablet";
            set req.http.X-UA-Vendor = "hp";
        }
        } elseif (req.http.User-Agent ~ "Kindle/3") {
            set req.http.X-UA-Device = "tablet";
            set req.http.X-UA-Vendor = "kindle";
        }
        } elseif (req.http.User-Agent ~ "Mobile.+Firefox") {
            set req.http.X-UA-Device = "mobile";
            set req.http.X-UA-Vendor = "firefoxos";
        }
        } elseif (req.http.User-Agent ~ "^HTC") {
            set req.http.X-UA-Device = "smartphone";
            set req.http.X-UA-Vendor = "htc";
        }
        } elseif (req.http.User-Agent ~ "Fennec") {
            set req.http.X-UA-Device = "smartphone";
            set req.http.X-UA-Vendor = "fennec";
        }
        } elseif (req.http.User-Agent ~ "IEMobile") {
            set req.http.X-UA-Device = "smartphone";
            set req.http.X-UA-Vendor = "microsoft";
        }
        } elseif (req.http.User-Agent ~ "BlackBerry" || req.http.User-Agent ~ "BB10.*Mobile") {
            set req.http.X-UA-Device = "smartphone";
            set req.http.X-UA-Vendor = "blackberry";
        }
        } elseif (req.http.User-Agent ~ "GT-.*Build/GINGERBREAD") {
            set req.http.X-UA-Device = "smartphone";
            set req.http.X-UA-Vendor = "android";
        }
        } elseif (req.http.User-Agent ~ "SymbianOS.*AppleWebKit") {
            set req.http.X-UA-Device = "smartphone";
            set req.http.X-UA-Vendor = "symbian";
        }
        } elseif (req.http.User-Agent ~ "(?i)symbian" ||
            req.http.User-Agent ~ "(?i)^sonyericsson" ||
            req.http.User-Agent ~ "(?i)^nokia" ||
            req.http.User-Agent ~ "(?i)^samsung" ||
            req.http.User-Agent ~ "(?i)^lg" ||
            req.http.User-Agent ~ "(?i)bada" ||
            req.http.User-Agent ~ "(?i)blazer" ||
```

```

req.http.User-Agent ~ "(?i)cellphone" ||
req.http.User-Agent ~ "(?i)iemobile" ||
req.http.User-Agent ~ "(?i)midp-2.0" ||
req.http.User-Agent ~ "(?i)u990" ||
req.http.User-Agent ~ "(?i)netfront" ||
req.http.User-Agent ~ "(?i)opera mini" ||
req.http.User-Agent ~ "(?i)palm" ||
req.http.User-Agent ~ "(?i)nintendo wii" ||
req.http.User-Agent ~ "(?i)playstation portable" ||
req.http.User-Agent ~ "(?i)portalmmm" ||
req.http.User-Agent ~ "(?i)proxinet" ||
req.http.User-Agent ~ "(?i)sonyericsson" ||
req.http.User-Agent ~ "(?i)symbian" ||
req.http.User-Agent ~ "(?i)windows\ ?ce" ||
req.http.User-Agent ~ "(?i)winwap" ||
req.http.User-Agent ~ "(?i)eudoraweb" ||
req.http.User-Agent ~ "(?i)htc" ||
req.http.User-Agent ~ "(?i)240x320" ||
req.http.User-Agent ~ "(?i)avantgo" {
  set req.http.X-UA-Device = "mobile";
}
}
}

```

## § GeoIP-related VCL features (/guides/vcl/geoip-related-vcl-features)

Fastly exposes a number of geographic variables for you to take advantage of inside VCL (/guides/vcl/guide-to-vcl). These are based on the MaxMind IP Geolocation database and appear as follows:

Variable	Description
<code>geoip.latitude</code>	The latitude associated with the IP address.
<code>geoip.longitude</code>	The longitude associated with the IP address.
<code>geoip.city</code>	The city or town name associated with the IP address. See MaxMind's list of cities ( <a href="http://www.maxmind.com/GeoIPCity-534-Location.csv">http://www.maxmind.com/GeoIPCity-534-Location.csv</a> ) for a list of all the possible return values. Their list is updated on a regular basis. The returned string uses the character encoding ISO-8859-1 (latin1).
<code>geoip.city.latin1</code>	An alias of <code>geoip.city</code> . The returned string uses the character encoding ISO-8859-1 (latin1).
<code>geoip.city.utf8</code>	The string from <code>geoip.city</code> encoded using the UTF-8 character encoding.
<code>geoip.city.ascii</code>	The string from <code>geoip.city</code> encoded using the ASCII character encoding (an ASCII approximation of the original string with diacritics removed).
<code>geoip.continent_code</code>	A two-character code representing the continent associated with the IP address. Possible codes are: <b>AF</b> - Africa, <b>AS</b> - Asia, <b>EU</b> - Europe, <b>NA</b> - North America, <b>OC</b> - Oceania, <b>SA</b> - South America.
<code>geoip.country_code</code>	A two-character ISO 3166-1 ( <a href="https://en.wikipedia.org/wiki/ISO_3166-1">https://en.wikipedia.org/wiki/ISO_3166-1</a> ) country code for the country associated with the IP address. In addition to the standard codes, we may also return one of the following: <b>A1</b> - an anonymous proxy ( <a href="http://dev.maxmind.com/faq/category/geoip-general/#anonproxy">http://dev.maxmind.com/faq/category/geoip-general/#anonproxy</a> ), <b>A2</b> - a satellite provider ( <a href="http://dev.maxmind.com/faq/category/geoip-general/#satellite">http://dev.maxmind.com/faq/category/geoip-general/#satellite</a> ), <b>EU</b> - an IP in a block used by multiple European ( <a href="http://dev.maxmind.com/faq/category/geoip-general/#euapcodes">http://dev.maxmind.com/faq/category/geoip-general/#euapcodes</a> ) countries, <b>AP</b> - an IP in a block used by multiple Asia/Pacific region ( <a href="http://dev.maxmind.com/faq/category/geoip-general/#euapcodes">http://dev.maxmind.com/faq/category/geoip-general/#euapcodes</a> ) countries. Note: The <b>US</b> country code is returned for IP addresses associated with overseas US military bases.
<code>geoip.country_code3</code>	The same as <code>country_code</code> ; however, it returns a ISO 3166-1 alpha-3 ( <a href="https://en.wikipedia.org/wiki/ISO_3166-1_alpha-3">https://en.wikipedia.org/wiki/ISO_3166-1_alpha-3</a> ) three-character code.
<code>geoip.country_name</code>	The country name associated with the IP address. The returned string uses the character encoding ISO-8859-1 (latin1).
<code>geoip.country_name.ascii</code>	The string from <code>geoip.country_name</code> encoded using the ASCII character encoding (an ASCII approximation of the original string with diacritics removed).
<code>geoip.country_name.latin1</code>	An alias of <code>geoip.country_name</code> . The returned string uses the character encoding ISO-8859-1 (latin1).
<code>geoip.country_name.utf8</code>	The string from <code>geoip.country_name</code> encoded using the UTF-8 character encoding.



Variable	Description
<code>geoip.postal_code</code>	The postal code associated with the IP address. These are available for some IP addresses in Australia, Canada, France, Germany, Italy, Spain, Switzerland, United Kingdom, and the US. We return the first 3 characters for Canadian postal codes. We return the first 2-4 characters (outward code) for postal codes in the United Kingdom.
<code>geoip.region</code>	The region name associated with the IP address. See MaxMind's FAQ ( <a href="https://dev.maxmind.com/faq/how-do-i-convert-region-codes-to-names/">https://dev.maxmind.com/faq/how-do-i-convert-region-codes-to-names/</a> ) about converting region codes to names. The returned string uses the character encoding ISO-8859-1 (latin1).
<code>geoip.region.latin1</code>	An alias of <code>geoip.region</code> . The returned string uses the character encoding ISO-8859-1 (latin1).
<code>geoip.region.utf8</code>	The string from <code>geoip.region</code> encoded using the UTF-8 character encoding.
<code>geoip.region.ascii</code>	The string from <code>geoip.region</code> encoded using the ASCII character encoding (an ASCII approximation of the original string with diacritics removed).
<code>geoip.area_code</code>	The telephone area code associated with the IP address. These are only available for IP addresses in the US.
<code>geoip.metro_code</code>	The metro code associated with the IP address. These are only available for IP addresses in the US.

**NOTE:** GeoIP information, including data streamed by our [log streaming service \(/guides/streaming-logs/\)](/guides/streaming-logs/), is intended to be used only in connection with your use of Fastly services. Use of GeoIP data for other purposes may require permission of a GeoIP vendor, such as [MaxMind \(https://www.maxmind.com/en/license\\_agreement\)](https://www.maxmind.com/en/license_agreement).

Fastly also exposes codes that describe the location of the datacenter the request came through as follows:

Variable	Description
<code>server.region</code>	A code representing the general region of the world in which the POP location resides. One of: APAC, Asia, EU-Central, EU-East, EU-West, Middle-East, North-America, SA-East, South-America, US-Central, US-East, US-West.
<code>server.datacenter</code>	A code representing one of Fastly's POP locations ( <a href="/guides/about-fastly-services/fastly-pop-locations/">/guides/about-fastly-services/fastly-pop-locations/</a> ).

**NOTE:** All strings are returned using [ISO-8859-1 character encoding \(https://en.wikipedia.org/wiki/ISO/IEC\\_8859-1\)](https://en.wikipedia.org/wiki/ISO/IEC_8859-1) (also known as latin1).

## § Guide to VCL (/guides/vcl/guide-to-vcl)

### About Varnish and why Fastly uses it

Varnish is the open source software (<https://www.fastly.com/open-source>) Fastly commercialized with performance and capacity (among other) enhancements. Fastly's Varnish is based on Varnish 2.1 and our Varnish syntax is specifically compatible with Varnish 2.1.5. The principal configuration mechanism of Varnish software is the Varnish Configuration Language (VCL), the scripting language used to configure and add logic to Varnish caches. Varnish allows Fastly to apply changes to the cache software as it is executing. Specifically, VCL is generated, compiled, transmitted to all Fastly caches, loaded into the operating software, and activated immediately, with no waiting for maintenance windows and no server downtime.

Fastly generates VCL automatically per your specifications via the web interface (</guides/about-fastly-services/>). We allow you to create your own VCL files with specialized configurations. Your custom VCL files can be uploaded (</guides/vcl/uploading-custom-vcl/>) into Fastly caches and activated. You can also mix and match (</guides/vcl/mixing-and-matching-fastly-vcl-with-custom-vcl/>) custom VCL and Fastly VCL, using them together at the same time. You will never lose the options on the Fastly user interface when you use custom VCL, but keep in mind that custom VCL always takes precedence over any VCL generated by the user interface. Be mindful of where your custom VCL sits in the default VCL.

**IMPORTANT:** The ability to [upload custom VCL code \(/guides/vcl/uploading-custom-vcl/\)](/guides/vcl/uploading-custom-vcl/) is disabled by default when you sign up. To enable this ability for your account, contact [support@fastly.com \(mailto:support@fastly.com\)](mailto:support@fastly.com) and request it.

### Fastly's VCL Extensions

In addition, Fastly has included a number of extensions to VCL that won't be covered by any other documentation. Specifically:

Extension	Description
cryptographic and hashing functions ( <a href="/guides/vcl/cryptographic-and-hashing-related-vcl-functions/">/guides/vcl/cryptographic-and-hashing-related-vcl-functions/</a> )	Supports Hash-based Message Authentication Code (HMAC), a message authentication code that uses a cryptographic key in conjunction with a hash function.

Extension	Description
date- and time-related features (/guides/vcl/date-and-time-related-vcl-features)	Supports the default VCL "now" variable that provides the current time as an RFC 850 formatted date (e.g., Tuesday, 29-Apr-14 08:41:55), as well as several new functions that allow you to have more flexibility when dealing with dates and times.
GeoIP features (/guides/vcl/geoip-related-vcl-features)	Provides the ability to search a database from MaxMind.com for a given host or IP address, and return information about the country, city or Internet Service Provider (ISP) for that IP address.
randomness features (/guides/vcl/randomness-related-vcl-features)	Supports the insertion of random strings, content cookies, and decisions into requests.
size-related variables (/guides/vcl/size-related-vcl-variables)	Supports reporting variables that offer insight into what happened in a request.
TLS and HTTP/2 variables (/guides/vcl/tls-and-http2-vcl-variables)	Supports the use of variables related to TLS and HTTP/2.
miscellaneous features and variables (/guides/vcl/miscellaneous-VCL-extensions)	Provides miscellaneous VCL extensions not easily grouped into other categories.

## Embedding inline C code in VCL

Currently, we don't provide embedded C access to our users. Fastly is a shared infrastructure. By allowing the use of inline C code, we could potentially give a single user the power to read, write to, or write from everything. As a result, our varnish process (i.e., files on disk, memory of the varnish user's processes) would become unprotected because inline C code opens the potential for users to do things like crash servers, steal data, or run a botnet.

We appreciate feedback from our customers. If you are interested in a feature that requires C code, contact [support@fastly.com](mailto:support@fastly.com) (mailto:support@fastly.com). Our engineering team looks forward to these kinds of challenges.

## Where to learn more about Varnish and VCL

The official Varnish documentation (<https://www.varnish-cache.org/docs/2.1/tutorial/vcl.html>) is a good place to start when looking for online information. In addition, Varnish Software, who provides commercial support for Varnish, has written a free online book (<http://info.varnish-software.com/the-varnish-book>).

Roberto Moutinho's book *Instant Varnish Cache* (<http://www.amazon.com/Instant-Varnish-Cache-Roberto-Moutinho/dp/178216040X>) also provides information.

## § Isolating header values without regular expressions (/guides/vcl/isolating-header-values-without-regular-expressions)

Fastly supports the ability to extract header subfield values without regular expressions in a human-readable way.

"Headers subfields" refer to headers with a body syntax style similar to `value1=123value123; testValue=asdf_true; loggedInTest=true;` Or `max-age=0, surrogate-control=3600`. These headers include Cookie, Set-Cookie, Cache-Control, or a custom header. Fastly allows you to isolate these key values with the following syntax:

```
req.http.Header-Name:key-name
```

In cases where a `Set-Cookie` response from origin is `value1=123value123; testValue=asdf_true; loggedInTest=true;`, the code for isolating the `loggedInTest` value would be:

```
beresp.http.Set-Cookie:loggedInTest
```

This logic can be used in uploaded custom VCL, as well as throughout the web interface. For example, using VCL this logic would execute based on the value of `staff_user` within `req.http.Cookie`.

```
# in vcl_recv
if (req.http.Cookie:staff_user ~ "true") {
    # some logic goes here
    return(pass);
}
```

For example, to isolate the value of `ab_test_value` from `Cookie` to the header `req.http.AB-Test-Value`, create a custom header (</guides/basic-configuration/adding-or-modifying-headers-on-http-requests-and-responses>) with the following settings:

## Create a new header

Learn more about this section in our [headers tutorial](#).

**CONDITION** This will happen all the time unless you [Attach a condition](#)

Type / Action Request Set

The type of header and the action performed on it.

Destination http.AB-Test-Value ★ Required

The name of the header that will be affected by the selected action. For example: `http.Content-Type`, `http.Set-Cookie`, `http.Via`, `http.Location`, or `http.Access-Control-Allow-Origin`.

Source req.http.Cookie:ab\_test\_value ★ Required

New content for the header. Can be a static value (e.g. string or number) or a dynamic value (e.g. existing header or a GeoIP value). Please use quotes for string values.

Ignore if set No

If switched to **Yes**, the action will not be performed if the header in Destination exists.

Priority 10 ★ Required

The order in which the header rules execute within the condition.

Description AB Test Header ★ Required

The name of your header, such as **My header**.

CREATE
CANCEL

Fill out the **Create a new header** fields as follows:

- From the **Type** menu, select **Request**, and from the **Action** menu select **Set**.
- In the **Destination** field, type `http.AB-Test-Value`.
- In the **Source** field, type `req.http.Cookie:ab_test_value`.
- From the **Ignore if set** menu, select **No**.
- In the **Priority** field, type `10`.
- In the **Description** field, type `AB Test Header`.

This will send the `AB-Test-Value` header in every inbound request.

## § Manipulating the cache key (</guides/vcl/manipulating-the-cache-key>)

### Redefining the cache key

#### ⚠ WARNING:

By default, Fastly uses the URL and the host of a request (plus a special, internal Fastly variable for [purging](/guides/purging/single-purges) (</guides/purging/single-purges>))

purposes) to create unique HTTP objects. Although Fastly allows you to explicitly set the cache key to define this more precisely, changing the default behavior risks the following:

1. If you add too much information to the cache key, you can significantly reduce your hit ratio.
2. If you make a mistake when explicitly setting the cache key, you can cause all requests to get the same object.
3. If you add anything to the hash, you will need to send a purge for each combination of the URL and value you add in order to purge that specific information from the cache.

To avoid these dangers, consider [using the Vary header \(https://www.fastly.com/blog/best-practices-for-using-the-vary-header\)](https://www.fastly.com/blog/best-practices-for-using-the-vary-header) instead of following the instructions below.

## Explicitly setting the cache key

You can set the cache key explicitly (including attaching conditions) by adding a request setting via the Settings tab on the Configure page of your service and including a comma-separated list of cache keys. The values of the cache keys listed are combined to make a single hash, and each unique hash is considered a unique object.

For example, if you don't want the query string to be part of the cache key, but you don't want to change `req.url` so that the query string still ends up in your logs, you could use the following text for the hash keys:

```
regsub(req.url, "\?.*$", ""), req.http.host
```

In the user interface, the text would appear in the Cache keys field:

ADVANCED OPTIONS

Override host

The Host header to be sent to your origins, regardless of the Host header in the initial end user request. For example, if your origin is Amazon, you would send the S3-bucket-name as your host header.

Timer support

Not selected

Whether or not to include the response time on MISSES from the origin server. Select **No** to hide this. [Learn more about understanding the X-Timer header.](#)

Maximum stale age (seconds)

The maximum time in seconds during which stale object content remains in Fastly's cache nodes. [Learn more about serving stale content.](#)

Force miss

Not selected

Select **Yes** to ignore a cached object and retrieve a new version from your origin, which will remain in cache after the request is complete. To refresh the object, use [the API's Purge function](#) instead. To bypass the object entirely, select **Pass** in the Action field instead.

Request collapsing

Enabled

By default, request collapsing is turned on. Before using this advanced feature, read about [request collapsing](#). Select to disable using `req.hash_ignore_busy`.

Cache keys

The hash keys or criteria used to define caching. You must send the same criteria when making a purge request.  
**Warning:** Changing Fastly's default cache behavior is risky. [Learn about manipulating the cache key.](#)

CREATE

CANCEL

As a general rule, you should always have `req.url` as one of your cache keys or as part of one. The example above includes `req.url` inside the `regsub()` therefore passing this requirement.

## Purging adjustments when making additions to cache keys

Because purging works on individual hashes, additions to cache keys can complicate purging URLs. However, it can also be simplified.

For example, if you were to change your cache key to just `req.url` and not the default `req.url, req.http.host`, then purging `http://foo.example.com/file.html` would also purge `http://bar.example.com/file.html`. Keep in mind this is because they're actually the same object in the cache!

On the other hand, if you were to change your cache key `req.url, req.http.host, req.http.Fastly-SSL`, you would have to purge `http://example.com/` and `https://example.com/` individually.

In the latter case, if you were to use the Vary header instead of changing the cache key, you could still have different content on the two URLs, yet purge them with a single purge. In this case you would add a new Cache Header (/guides/basic-configuration/adding-or-modifying-headers-on-http-requests-and-responses), use `http.Vary` as the Destination, and use the following as the Source:

```
if(beresp.http.Vary, beresp.http.Vary ",", "") "Fastly-SSL"
```

## Using a POST request body as a cache key

As long as the body of a POST request is less than 2K in size and the content type is `application/x-www-form-urlencoded`, then we allow you to use it as part of the cache key. Your VCL (/guides/vcl/) should look something like:

```
sub vcl_hash {
    set req.hash += req.url;
    set req.hash += req.http.host;
    if (req.request == "POST" && req.postbody) {
        set req.hash += req.postbody;
    }
    return(hash);
}
```

You'll also need to force a cache lookup, but only for requests that can be cached, by doing something like this in `vcl_recv`:

```
if (req.request == "POST" && req.postbody - "(^|&)action=list(&|$)") {
    return(lookup);
}
```

★ **TIP:** To refine this, you could add only the important parts of `req.postbody` to the hash using `regsub()`.

## Using a cookie as a cache key

You can use a cookie as a cache key or just check for the presence of a cookie set to a specific value by controlling its request conditions (/guides/conditions/). Both methods are simple and shown in the steps below.

### To use a cookie as a cache key

Using a cookie as a cache key looks complicated but it's actually quite simple. Let's say your cookie is called "MyCookie" and it looks like

```
mycookie=
```

#### Create new headers

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.
4. Click the **Content** tab. The Content page appears.

DOMAINS  
1 Domain

ORIGINS  
1 Host  
1 Health check

SETTINGS  
1 Request setting  
0 Cache settings

CONTENT  
0 Headers  
0 Gzips  
0 Responses

LOGGING  
1 Endpoint

CUSTOM VCL  
0 Files

## Headers

*There are no headers.*

CREATE YOUR FIRST HEADER

## Gzip

*Gzip is not enabled.*

SET UP GZIP

## Responses

*There are no responses.*

CREATE YOUR FIRST RESPONSE

### About content

Headers allow you to set, delete, and change request and response headers as your information is handled by Fastly.

Responses allow you to create synthetic responses that exist entirely on the cache servers.

5. Click the **Create header** button. The Create a new header page appears.

## Create a new header

Learn more about this section in our [headers tutorial](#).

### CONDITION

This will happen all the time unless you [Attach a condition](#)

Type / Action

Request

Set

The type of header and the action performed on it.

Destination

★ Required

The name of the header that will be affected by the selected action. For example: **http.Content-Type**, **http.Set-Cookie**, **http.Via**, **http.Location**, or **http.Access-Control-Allow-Origin**.

Source

★ Required

New content for the header. Can be a static value (e.g. string or number) or a dynamic value (e.g. existing header or a GeolIP value). Please use quotes for string values.

Ignore if set

No

If switched to **Yes**, the action will not be performed if the header in Destination exists.

Priority

10

★ Required

The order in which the header rules execute within the condition.

Description

★ Required

The name of your header, such as **My header**.

CREATE

CANCEL

6. Fill out the **Create a new header** fields as follows:


- From the **Type** menu, select **Request**, and from the **Action** menu select **Set**.
- In the **Destination** field, type `http.X-MyCookie`.
- In the **Source** field, type `"0"` (with quotes).
- Leave the **Ignore if set** menu set to the default, **No**.
- In the **Priority** field, type a number representing the order in which the header rule should execute. The default is set to `10` for new headers.
- In the **Description** field, type `Set MyCookie Header Default`.

7. Click the **Create** button. The new header appears in the Headers area of the Content tab.

8. Click the **Create header** button again and create a second new header by filling out the fields as follows:

- From the **Type** menu, select **Request**, and from the **Action** menu select **Set**.
- In the **Destination** field, type `http.X-MyCookie`.
- In the **Source** field, type `regsub(req.http.cookie, ".mycookie =([^;]+);.*", "\1")`.
- Leave the **Ignore if set** menu set to the default, **No**.
- In the **Priority** field, type a larger number than the priority of previous header you just created. For example, if you left the default priority set to `10`, type `20`.
- In the **Description** field, type `Set MyCookie Header from Cookie`.

- Click the **Create** button. The second header appears in the Headers area of the Content tab.




Headers

**Set MyCookie Header Default** [Attach a condition](#) [Delete](#)  
*Request / Set*  
[Show details](#)

**Set MyCookie Header from Cookie** [Attach a condition](#) [Delete](#)  
*Request / Set*  
[Show details](#)

### Attach conditions to the new headers

- Click the **Attach a condition** link next to the `Set MyCookie Header from Cookie` header. The add a condition window appears.
- Click the **Create a new request condition** button. The Create a new request condition window appears.



Create a new request condition. ×  
Learn the basics in our [conditions tutorial](#)

Type  ⌵  
Learn more about the [different types of conditions](#).

Name  ★ Required

Apply if...   
The expression to decide whether this is run.  
[▶ Examples](#)

[Adjust priority \(default is 10\)](#)  
This is an advanced option. For most configurations, the default is best.

**SAVE AND APPLY TO SET MYCOOKIE HEADER FROM COOKIE** **CANCEL**

- Fill out the fields of the **Create a new request condition** window as follows:
  - Leave the **Type** menu set to **Request**.
  - In the **Name** field, type `Has MyCookie cookie`.
  - In the **Apply if** field, type `req.http.cookie ~ "mycookie="`.
- Click the **Save and apply to** button. The Headers area now displays the condition that must be met in order for your header to begin being used.



5. Click the **Settings** tab. The Settings page appears.
6. Click the **Create request setting** button. The Create a new request setting page appears.
7. Fill out the **Create a new request setting** fields as follows:
  - In the **Description** field type `Set Hash from Cookie`.
  - In the **Cache keys** field, type `req.url, req.http.host, req.http.X-MyCookie`
8. Click the **Create** button. The new request appears in the Request settings area.
9. Click the **Attach a condition** link next to the new request. The Add a condition window appears.

10. From the **Select a condition** menu, select `Has MyCookie cookie`. The Request settings area now displays the condition that must be met in order for your request to begin being used.

11. Click the **Activate** button to deploy your configuration changes.

## To check for the presence of a cookie set to a specific value

An alternative way if you're just checking for the presence of the cookie set to some specific value (e.g., 1):

1. Add a new Request setting where the Cache key field is set to `req.url, req.http.host, "Has mycookie"`.
2. Add a condition to that Request setting where the Apply if field contains `req.http.cookie ~ ".*mycookie =1;.*"`.

## § Miscellaneous VCL features (/guides/vcl/miscellaneous-VCL-extensions)

Fastly has added several miscellaneous features (/guides/vcl/guide-to-vcl) to Varnish that don't easily fit into specific categories (/guides/vcl/).

Feature	Description
<code>bereq.url.basename</code>	Same as <code>req.url.basename</code> , except for use between Fastly and your origin servers.
<code>bereq.url.dirname</code>	Same as <code>req.url.dirname</code> , except for use between Fastly and your origin servers.
<code>beresp.backend.ip</code>	The IP of the backend this response was fetched from (backported from Varnish 3).
<code>beresp.backend.name</code>	The name of the backend this response was fetched from (backported from Varnish 3).
<code>beresp.backend.port</code>	The port of the backend this response was fetched from (backported from Varnish 3).
<code>beresp.grace</code>	Defines how long an object can remain overdue and still have Varnish consider it for grace mode. Fastly has implemented <code>stale-if-error</code> (/guides/performance-tuning/serving-stale-content#usage) as a parallel implementation of <code>beresp.grace</code> .
<code>beresp.pci</code>	Specifies that content be cached in a manner that satisfies PCI DSS requirements. See our PCI compliance description (/guides/compliance/pci-compliant-caching) for instructions on enabling this feature for your account.
<code>beresp.hipaa</code>	Specifies that content not be cached in non-volatile memory to help customers meet HIPAA security requirements. See our guide on HIPAA and caching PHI (/guides/compliance/hipaa-and-caching-phi) for instructions on enabling this feature for your account. An alias of <code>beresp.pci</code> .
<code>boltsort.sort</code>	Sorts URL parameters. For example, <code>boltsort.sort("/foo?b=1&amp;a=2&amp;c=3");</code> returns <code>"/foo?a=2&amp;b=1&amp;c=3"</code> .
<code>client.port</code>	Returns the remote client port. This could be useful as a seed that returns the same value both in an ESI and a top level request. For example, you could hash <code>client.ip</code> and <code>client.port</code> to get a value used both in ESI and the top level request.
<code>goto</code>	Performs a one-way transfer of control to another line of code. See the example for more information.
<code>http_status_matches</code>	Determines whether or not an HTTP status code matches a pattern. The arguments are an integer (usually <code>beresp.status</code> or <code>resp.status</code> ) and a comma-separated list of status codes, optionally prefixed by a <code>!</code> to negate the match. It returns <code>TRUE</code> or <code>FALSE</code> . For example, <code>if (http_status_matches(beresp.status, "!200,404")) {}</code> .
<code>if()</code>	Implements a ternary operator for strings; if the expression is true, it returns <code>TRUE</code> ; if the expression is false, it returns <code>FALSE</code> . For example, you have an <code>if(x, true-expression, false-expression)</code> ; if this argument is true, the <code>true-expression</code> is returned. Otherwise, the <code>false-expression</code> is returned. See the example for more information.
<code>req.grace</code>	Defines how long an object can remain overdue and still have Varnish consider it for grace mode.
<code>req.http.host</code>	The full host name, without the path or query parameters. For example, in the request <code>www.example.com/index.html?a=1&amp;b=2</code> , <code>req.http.host</code> will return <code>www.example.com</code> .
<code>req.restarts</code>	Counts the number of times the VCL has been restarted.
<code>req.topurl</code>	In an ESI subrequest, returns the URL of the top-level request.
<code>req.url.basename</code>	The file name specified in a URL. For example, in the request <code>www.example.com/1/hello.gif?foo=bar</code> , <code>req.url.basename</code> will return <code>hello.gif</code> .
<code>req.url.dirname</code>	The directories specified in a URL. For example, in the request <code>www.example.com/1/hello.gif?foo=bar</code> , <code>req.url.dirname</code> will return <code>/1</code> . In the request <code>www.example.com/5/inner/hello.gif?foo=bar</code> , <code>req.url.dirname</code> will return <code>/5/inner</code> .
<code>req.url.ext</code>	The file extension specified in a URL. For example, in the request <code>www.example.com/1/hello.gif?foo=bar</code> , <code>req.url.ext</code> will return <code>gif</code> .
<code>req.url.path</code>	The full path, without any query parameters. For example, in the request <code>www.example.com/index.html?a=1&amp;b=2</code> , <code>req.url.path</code> will return <code>/index.html</code> .
<code>req.url</code>	The full path, including query parameters. For example, in the request <code>www.example.com/index.html?a=1&amp;b=2</code> , <code>req.url</code> will return <code>/index.html?a=1&amp;b=2</code> .

Feature	Description
<code>return</code>	Returns (with no return value) from a custom subroutine to exit early. See the example for more information.
<code>stale.exists</code>	Specifies if a given object has stale content ( <a href="/guides/performance-tuning/serving-stale-content">/guides/performance-tuning/serving-stale-content</a> ) in cache. Returns <code>TRUE</code> or <code>FALSE</code> .
<code>std.atoi</code>	Takes a string (which represents an integer) as an argument and returns its value.
<code>std.strstr</code>	Finds the first occurrence of a byte string and returns its value.
<code>std.tolower</code>	Changes the case of a string to lower case. For example, <code>std.tolower("HELLO");</code> will return <code>"hello"</code> .
<code>std.toupper</code>	Changes the case of a string to upper case. For example, <code>std.toupper("hello");</code> will return <code>"HELLO"</code> .
<code>urldecode</code>	Decodes a percent-encoded string. For example, <code>urldecode("hello%20world!");</code> will return <code>"hello world !"</code> .
<code>urlencode</code>	Encodes a string for use in a URL. This is also known as percent-encoding ( <a href="https://en.wikipedia.org/wiki/Percent-encoding">https://en.wikipedia.org/wiki/Percent-encoding</a> ). For example, <code>urlencode("hello world");</code> will return <code>"hello%20world"</code> .

## Examples

Use the following examples to learn how to implement the features.

### Goto

In the C programming language, `goto` is used to jump to a common error handling block before returning from a function. It can be used to do something similar in VCL, such as perform logic or set headers before returning lookup, error, or pass. It also works in custom subroutines. You can only use it to jump forward in the code, not backward.

```
sub vcl_recv {
    if (req.http.Foo) {
        goto foo;
    }

    foo:
        set req.http.Foo = "1";
    }
```

### Return

You can use `return` to exit early from a custom subroutine.

```
sub custom_subroutine {
    if (req.http.Cookie:user_id) {
        return;
    }

    # do a bunch of other stuff
}
```

### If

You can use `if()` as a construct to make simple conditional expressions more concise.

```
set req.http.foo-status = if(req.http.foo, "present", "absent");
```

## § Mixing and matching Fastly VCL with custom VCL (</guides/vcl/mixing-and-matching-fastly-vcl-with-custom-vcl>)

**ⓘ IMPORTANT:** The ability to [upload custom VCL code](/guides/vcl/uploading-custom-vcl) (</guides/vcl/uploading-custom-vcl>) is disabled by default when you sign up. To enable this ability for your account, contact [support@fastly.com](mailto:support@fastly.com) (<mailto:support@fastly.com>) and request it.

Fastly Varnish syntax is specifically compatible with Varnish 2.1.5 (<https://www.varnish-cache.org/docs/2.1>). We run a custom version with added functionality and our VCL parser has its own pre-processor. To mix and match Fastly VCL with your custom VCL successfully, remember the following:

- **You can only restart Varnish requests three times.** This limit exists to prevent infinite loops.

- **VCL doesn't take kindly to Windows newlines (line breaks).** It's best to avoid them entirely.
- **It's best to use `curl -X PURGE` to initiate purges via API (/api/purge).** To restrict access to purging, check for the `FASTLYPURGE` method not the `PURGE` method. When you send a request to Varnish to initiate a purge, the HTTP method that you use is "PURGE", but it has already been changed to "FASTLYPURGE" by the time your VCL runs that request.
- **If you override TTLs with custom VCL, your default TTL set in the configuration (/guides/performance-tuning/serving-stale-content) will not be honored** and the expected behavior may change.

## Inserting custom VCL in Fastly's VCL boilerplate

**⚠ DANGER:** Include all of the Fastly VCL boilerplate as a template in your custom VCL file, especially the VCL macro lines (they start with `#FASTLY`). VCL macros expand the code into generated VCL. Add your custom code *in between* the different sections as shown in the example unless you specifically intend to override the VCL at that point.

### Custom VCL placement example

```
sub vcl_miss {  
  # my custom code  
  if (req.http.User-Agent ~ "Googlebot") {  
    set req.backend = F_special_google_backend;  
  }  
  #FASTLY miss  
  return(fetch);  
}
```

### Fastly's VCL boilerplate

```
sub vcl_recv {
#FASTLY recv

    if (req.request != "HEAD" && req.request != "GET" && req.request != "FASTLYPURGE") {
        return(pass);
    }

    return(lookup);
}

sub vcl_fetch {
#FASTLY fetch

    if ((beresp.status == 500 || beresp.status == 503) && req.restarts < 1 && (req.request == "GET" || req.request == "HEAD")) {
        restart;
    }

    if (req.restarts > 0) {
        set beresp.http.Fastly-Restarts = req.restarts;
    }

    if (beresp.http.Set-Cookie) {
        set req.http.Fastly-Cachetype = "SETCOOKIE";
        return(pass);
    }

    if (beresp.http.Cache-Control ~ "private") {
        set req.http.Fastly-Cachetype = "PRIVATE";
        return(pass);
    }

    if (beresp.status == 500 || beresp.status == 503) {
        set req.http.Fastly-Cachetype = "ERROR";
        set beresp.ttl = 1s;
        set beresp.grace = 5s;
        return(deliver);
    }

    if (beresp.http.Expires || beresp.http.Surrogate-Control ~ "max-age" || beresp.http.Cache-Control ~ "(s-maxage|max-age)") {
        # keep the ttl here
    } else {
        # apply the default ttl
        set beresp.ttl = 3600s;
    }

    return(deliver);
}

sub vcl_hit {
#FASTLY hit

    if (!obj.cacheable) {
        return(pass);
    }
    return(deliver);
}

sub vcl_miss {
#FASTLY miss
    return(fetch);
}

sub vcl_deliver {
#FASTLY deliver
    return(deliver);
}

sub vcl_error {
#FASTLY error
}

sub vcl_pass {
#FASTLY pass
}

sub vcl_log {
#FASTLY log
}
```

## § Overriding which IP address the GeoIP features use (/guides/vcl/overriding-which-ip-address-the-geoip-features-use)

By default the GeoIP lookup is based on the IP address of the user. In some cases, such as with traffic through proxies, this type of lookup doesn't work properly. In particular, users of Opera Mini always browse through a proxy and the true IP address appears in the X-Forwarded-For header (<https://dev.opera.com/articles/opera-mini-request-headers/#x-forwarded-for>). Similarly, the Amazon Silk browser (<https://developer.amazon.com/appsandservices/community/post/TxOMW3RNF3FYRK/Amazon-Silk-Tips-for-Site-Owners>) can optionally come through a proxy, indicated via the User Agent string.

To work around this, you can set this variable to account for both the Opera Mini and Amazon Silk browsers:

```
set geoip.use_x_forwarded_for = true;
```

To do so, you would use code like this in `vcl_recv`:

```
if (req.http.X-OperaMini-Features || req.http.User-Agent ~ " Silk-Accelerated=true$") {  
    set geoip.use_x_forwarded_for = true;  
}
```

which tells Fastly to use the X-Forwarded-For header for the IP address. If it is not available, then the code will fall back to using the IP address of the client.

Finally, just in case there's some scenario or browser we haven't anticipated, you can also override based on an arbitrary header:

```
set geoip.ip_override = req.http.Custom-IP-Override;
```

Setting either of these variables:

```
set geoip.use_x_forwarded_for = true;
```

```
set geoip.ip_override = req.http.Custom-IP-Override;
```

will force the GeoIP information to be reloaded.

**⚠ IMPORTANT:** The ability to [upload custom VCL code \(/guides/vcl/uploading-custom-vcl\)](/guides/vcl/uploading-custom-vcl) is disabled by default when you sign up. To enable this ability for your account, contact [support@fastly.com](mailto:support@fastly.com) (<mailto:support@fastly.com>) and request it.

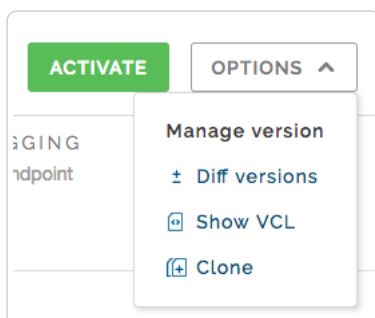
## § Previewing and testing VCL before activating it (/guides/vcl/previewing-and-testing-vcl-before-activating-it)

Any time you upload VCL files (</guides/vcl/uploading-custom-vcl>) you can preview and test the VCL prior to activating a new version of your service.

### Previewing VCL before activation

To preview VCL prior to activating a service version.

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.
4. Click the **Options** button to open the **Manage version** menu and select **Show VCL**.



The VCL preview page appears.

## Testing VCL configurations

You don't need a second account to test your VCL configurations. We recommend adding a new service within your existing account that's specifically designed for testing. A name like "QA" or "testing" or "staging" makes distinguishing between services much easier.

Once created, simply point your testing service to your testing or QA environment. Edit your Fastly configurations for the testing service as if you were creating them for production. Preview your VCL, test things out, and tweak them to get them perfect.

When your testing is complete, make the same changes in your production service that you made to your testing service. If you are using custom VCL, upload the VCL file (</guides/vcl/uploading-custom-vcl>) to the production service you'll be using.

**❗ IMPORTANT:** The ability to [upload custom VCL code](/guides/vcl/uploading-custom-vcl) (</guides/vcl/uploading-custom-vcl>) is disabled by default when you sign up. To enable this ability for your account, contact [support@fastly.com](mailto:support@fastly.com) (<mailto:support@fastly.com>) and request it.

## § Query string manipulation VCL features (</guides/vcl/query-string-manipulation-vcl-features>)

Fastly provides several functions in VCL for query-string manipulation based on Dridi Boukelmoune's `vmod-querystring` for Varnish. Refer to the original documentation (<https://github.com/Dridi/libvmod-querystring>) for more information.

### Functions

Function	Description	Example
<code>querystring.clean(&lt;string&gt;)</code>	Returns the given URL without empty parameters. The query-string is removed if empty (either before or after the removal of empty parameters). Note that a parameter with an empty value does not constitute an empty parameter, so a query string "?something" would not be cleaned.	<pre>set req.url = querystring.clean(req.url);</pre>
<code>querystring.filter(&lt;string&gt;, &lt;string_list&gt;)</code>	Returns the given URL without the listed parameters.	<pre>set req.url = querystring.filter(req.url, "utm_source" + querystring.filtersep() + "utm_medium" + querystring.filtersep() + "utm_campaign");</pre>
<code>querystring.filter_except(&lt;string&gt;, &lt;string_list&gt;)</code>	Returns the given URL but only keeps the listed parameters.	<pre>set req.url = querystring.filter_except(req.url, "q" + querystring.filtersep() + "p");</pre>
<code>querystring.filtersep()</code>	Returns the separator needed by the <code>filter</code> and <code>filter_except</code> functions.	
<code>querystring.globfilter(&lt;string&gt;, &lt;string&gt;)</code>	Returns the given URL without the parameters matching a glob.	<pre>set req.url = querystring.globfilter(req.url, "utm_*");</pre>
<code>querystring.globfilter_except(&lt;string&gt;, &lt;string&gt;)</code>	Returns the given URL but only keeps the parameters matching a glob.	<pre>set req.url = querystring.globfilter_except(req.url, "sess*");</pre>
<code>querystring.regfilter(&lt;string&gt;, &lt;string&gt;)</code>	Returns the given URL without the parameters matching a regular expression.	<pre>set req.url = querystring.regfilter(req.url, "utm_.*");</pre>
<code>querystring.regfilter_except(&lt;string&gt;, &lt;string&gt;)</code>	Returns the given URL but only keeps the parameters matching a regular expression.	<pre>set req.url = querystring.regfilter_except(req.url, "^(q p)\$");</pre>
<code>querystring.remove(&lt;string&gt;)</code>	Returns the given URL with its query-string removed.	<pre>set req.url = querystring.remove(req.url);</pre>

Function	Description	Example
<code>querystring.sort(&lt;string&gt;)</code>	Returns the given URL with its query-string sorted.	<pre>set req.url = querystring.sort(req.url);</pre>

## Examples

In your VCL, you could use `querystring.sort` as follows:

```
import querystring;

sub vcl_recv {
  # return this URL with only the parameters that match this regular expression
  set req.url = querystring.regfilter_except(req.url, "^(q|p)$");
}
```

You can use `querystring.regfilter` to specify a list of arguments that must not be removed (everything else will be) with a negative look-ahead expression:

```
set req.url = querystring.regfilter(req.url, "^(?!param1|param2)");
```

## § Randomness-related VCL features (/guides/vcl/randomness-related-vcl-features)

Fastly provides several functions in VCL (/guides/vcl/) that control randomness-related activities.

**⚠ WARNING:** We use BSD random number functions from the [GNU C Library \(http://www.gnu.org/software/libc/manual/html\\_node/BSD-Random.html\)](http://www.gnu.org/software/libc/manual/html_node/BSD-Random.html), not true randomizing sources. These VCL functions should not be used for [cryptographic \(/guides/vcl/cryptographic-and-hashing-related-vcl-functions\)](/guides/vcl/cryptographic-and-hashing-related-vcl-functions) or security purposes.

## Random strings

Use the function `randomstr(length [, characters])`. When characters aren't provided, the default will be the 64 characters of `A-Za-z0-9_-`.

```
sub vcl_deliver {
  set resp.http.Foo = "randomstuff=" randomstr(10);
  set resp.http.Bar = "morsecode=" randomstr(50, "-."); # 50 dots and dashes
}
```

## Random content cookies in pure VCL

```
sub vcl_deliver {
  add resp.http.Set-Cookie = "somerandomstuff=" randomstr(10) "; expires=" now + 180d "; path=/;";
}
```

This adds a cookie named "somerandomstuff" with 10 random characters as value, expiring 180 days from now.

## Random decisions

Use the function `randombool( _numerator_, _denominator_ )`, which has a numerator/denominator chance of returning true.

```
sub vcl_recv {
  if (randombool(1, 4)) {
    set req.http.X-AB = "A";
  } else {
    set req.http.X-AB = "B";
  }
}
```

This will add a X-AB header to the request, with a 25% (1 out of 4) chance of having the value "A", and 75% chance of having the value "B".

The `randombool()` function accepts INT function return values, so you could do something this:

```
if (randombool(std.atoi(req.http.Some-Header), 100)) {
  # do something
}
```



Another function, `randombool_seeded()`, takes an additional seed argument. Results for a given seed will always be the same. For instance, in this example the value of the response header will always be `no`:

```
if (randombool_seeded(50, 100, 12345)) {  
  set resp.http.Seeded-Value = "yes";  
} else {  
  set resp.http.Seeded-Value = "no";  
}
```

This could be useful for stickiness. For example, if you based the seed off of something that identified a user, you could perform A/B testing without setting a special cookie.

**⚠ WARNING:** The `randombool` and `randombool_seeded` functions do not use secure random numbers and should not be used for cryptographic purposes.

## § Response Cookie handling (/guides/vcl/response-cookie-handling)

The traditional way to read response cookies in VCL is to inspect either the `beresp.http.Set-Cookie` or the `resp.http.Set-Cookie` variables and then extract values using regular expressions. However this is not ideal since attempting to parse potentially complicated or quoted strings with regular expressions is brittle and prone to being tripped up by edge cases. It also doesn't allow for reading multiple headers with the same name such as when an origin sends multiple `Set-Cookie` headers. Because of these two reasons Fastly supports a method for extracting a named value out of `Set-Cookie` headers no matter how many there are.

To access a named value simply use the function with either `beresp` or `resp` depending on what part of the request you're in - so either

```
setcookie.get_value_by_name(beresp, "name")
```

or

```
setcookie.get_value_by_name(resp, "name")
```

as appropriate, replacing `"name"` with whatever the name of the value is. So for example, given this HTTP response from an origin

```
HTTP/1.1 200 OK  
Cache-Control: max-age=60  
Content-Type: text/html; charset=utf-8  
Content-Length: 80806  
Accept-Ranges: bytes  
Date: Tue, 11 Aug 2015 19:00:04 GMT  
Age: 123  
Connection: keep-alive  
Set-Cookie: one=a; httponly; secure  
Set-Cookie: two=b or not to b; httponly
```

then using the function like this

```
set resp.http.X-One = setcookie.get_value_by_name(resp, "one");  
set resp.http.X-Two = setcookie.get_value_by_name(resp, "two");
```

will set `resp.http.X-One` to be "a" and `resp.http.X-Two` to "b or not to b".

This logic can be used in uploaded custom VCL (/guides/vcl/uploading-custom-vcl), as well as throughout the web interface. For example:

## Create a new header

Learn more about this section in our [headers tutorial](#).

**CONDITION** This will happen all the time unless you [Attach a condition](#)

Type / Action	<input type="text" value="Request"/>	<input type="text" value="Set"/>
The type of header and the action performed on it.		
Destination	<input type="text" value="http.x-One"/>	★ Required
The name of the header that will be affected by the selected action. For example: <code>http.Content-Type</code> , <code>http.Set-Cookie</code> , <code>http.Via</code> , <code>http.Location</code> , or <code>http.Access-Control-Allow-Origin</code> .		
Source	<input name")"="" type="text" value="setcookie.get_value_by_name(resp, "/>	★ Required
New content for the header. Can be a static value (e.g. string or number) or a dynamic value (e.g. existing header or a GeoIP value). Please use quotes for string values.		
Ignore if set	<input type="text" value="No"/>	
If switched to <b>Yes</b> , the action will not be performed if the header in Destination exists.		
Priority	<input type="text" value="10"/>	★ Required
The order in which the header rules execute within the condition.		
Description	<input type="text" value="x-One header"/>	★ Required
The name of your header, such as <b>My header</b> .		

CREATE

CANCEL

## § Setting up redundant origin servers (/guides/vcl/setting-up-redundant-origin-servers)

Sometimes you want to set up two different origin servers, one as a primary and one as a backup in case the primary becomes unavailable. You can do this via the web interface or using custom VCL.

### Using the user interface

Set up redundant origins via the user interface using these steps.

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.
4. Click the **Origins** tab. The Origins page appears.
5. In the **Health Checks** area, define a health check (/guides/basic-configuration/health-checks-tutorial) and assign it to the primary origin server.
6. In the **Hosts** area, find your secondary origin server and click the **Attach a condition** link.

Hosts

CREATE HOST

192.168.1.1 : 80

Attach a condition

TLS Options

Delete

Secondary Origin Server

Show all details

The Add a condition window appears.

Add a condition to Secondary Origin Server

Learn the basics in our [conditions tutorial](#)

Select a condition (7 total)

or

CREATE A NEW REQUEST CONDITION

CANCEL

7. Click **Create a new request condition**. The Create a new request condition window appears.

Create a new request condition.

Learn the basics in our [conditions tutorial](#)

Type

Request

Learn more about the [different types of conditions](#).

Name

Primary Unhealthy

★ Required

Apply if...

!req.backend.healthy

The expression to decide whether this is run.

▶ [Examples](#)

[Adjust priority \(default is 10\)](#)

This is an advanced option. For most configurations, the default is best.

SAVE AND APPLY TO SECONDARY ORIGIN SERVER

CANCEL

8. Fill out the **Create a new request condition** fields as follows:

- From the **Type** menu, select **Request**.
- In the **Name** field, type the name of your request condition (for example, `Primary Unhealthy`).
- In the **Apply if** field, type `!req.backend.healthy`.

9. Click the **Save and apply to** button. The Hosts area now displays the condition that must be met (Primary Unhealthy) in order for your secondary origin server to begin being used.

Hosts

CREATE HOST

IF **Primary Unhealthy**  
!req.backend.healthy

THEN  
**192.168.1.1 : 80**  
Secondary Origin Server  
Show all details

Detach from condition

TLS Options

Delete

Once you've added the condition to your secondary origin server, the VCL generated by Fastly will reflect the new condition.

10. Preview the VCL (/guides/vcl/previewing-and-testing-vcl-before-activating-it), and confirm the following snippets appear in `vcl_recv`:

```
# default conditions
set req.backend = F_primary;

# Request Condition: primary unhealthy Prio: 10
if (!req.backend.healthy) {
  set req.backend = F_secondary;
}
#end condition
```

Using custom VCL

**ⓘ IMPORTANT:** The ability to upload custom VCL code (/guides/vcl/uploading-custom-vcl) is disabled by default when you sign up. To enable this ability for your account, contact [support@fastly.com](mailto:support@fastly.com) (mailto:support@fastly.com) and request it.

Set up redundant origins with custom VCL using these steps.

- 1. In the Fastly web interface, define a health check (/guides/basic-configuration/health-checks-tutorial) and assign it to the primary origin server.
- 2. Copy the boilerplate VCL from our guide on mixing Fastly VCL with custom VCL (/guides/vcl/mixing-and-matching-fastly-vcl-with-custom-vcl), and paste it into a new file.
- 3. Replace the `vcl_recv` sub with:

```
sub vcl_recv {
  #FASTLY recv
  set req.backend = F_<primary_origin>;
  if (!req.backend.healthy) {
    set req.backend = F_<secondary_origin>;
  }
  if (req.request != "HEAD" && req.request != "GET" && req.request != "FASTLYPURGE") {
    return(pass);
  }
  return(lookup);
}
```

To find the exact backend names, view the generated VCL (/guides/vcl/previewing-and-testing-vcl-before-activating-it).

- 4. Upload (/guides/vcl/uploading-custom-vcl) your VCL file.

§ Size-related VCL variables (/guides/vcl/size-related-vcl-variables)

To allow better reporting, Fastly has added several variables (/guides/vcl/guide-to-vcl#fastly-39-s-vcl-extensions) to VCL to give more insight into what happened in a request.

Name	Where	Description
<code>req.bytes_read</code>	deliver log	How big a request was in total bytes.
<code>req.header_bytes_read</code>	all	How big the header of a request was in total bytes.

Name	Where	Description
<code>req.body_bytes_read</code>	deliver log	How big the body of a request was in total bytes.
<code>resp.bytes_written</code>	log	How many bytes in total were sent as a response.
<code>resp.header_bytes_written</code>	log	How many bytes were written for the header of a response.
<code>resp.body_bytes_written</code>	log	How many bytes were written for body of a response.
<code>resp.completed</code>	log	Whether the response completed successfully or not.

## § Support for the Edge-Control header (/guides/vcl/support-for-the-edge-control-header)

VCL provides the building blocks to access information inside the Edge-Control response header field from the origin. We support this by honoring `cache-maxage` from Edge-Control as the time to live (TTL) of the object on the Fastly edge, and honoring `downstream-ttl` from Edge-Control as the TTL to be sent down from the Fastly edge to the end user's browser.

In order to incorporate this Edge-Control header support, include the following snippet in your `vcl_fetch` function via custom VCL (/guides/vcl/uploading-custom-vcl):

```
if (parse_time_delta(subfield(beresp.http.Edge-Control, "downstream-ttl")) >= 0) {
    set beresp.http.Cache-Control = "max-age=" parse_time_delta(subfield(beresp.http.Edge-Control, "downstream-ttl"));
}

if (parse_time_delta(subfield(beresp.http.Edge-Control, "cache-maxage")) >= 0) {
    set beresp.ttl = parse_time_delta(subfield(beresp.http.Edge-Control, "cache-maxage"));
}
```

The `subfield` function parses the Edge-Control field for subfields, and the `parse_time_delta` function converts time values like "7m" into a number of seconds. You can then use that number of seconds to populate `beresp.ttl` (the TTL of the object on the Fastly edge) or you can use it to construct a Cache-Control header field for downstream. The `parse_time_delta` function will return -1 if the subfield is not well-formed as a time value, or if it is entirely absent. The above snippet honors `cache-maxage` and `downstream-ttl` from Edge-Control if present and usable.

**⚠ IMPORTANT:** The ability to [upload custom VCL code \(/guides/vcl/uploading-custom-vcl\)](/guides/vcl/uploading-custom-vcl) is disabled by default when you sign up. To enable this ability for your account, contact [support@fastly.com](mailto:support@fastly.com) (<mailto:support@fastly.com>) and request it.

## § TLS and HTTP/2 VCL variables (/guides/vcl/tls-and-http2-vcl-variables)

Fastly has added several variables (/guides/vcl/guide-to-vcl#fastly-39-s-vcl-extensions) that expose information about the TLS and HTTP/2 attributes of a request.

**⚠ IMPORTANT:** These VCL variables are currently only allowed to appear within the VCL hooks `vcl_recv`, `vcl_hash`, `vcl_deliver` and `vcl_log`.

Feature	Description
<code>fastly_info.is_h2</code>	Whether or not the request was made using HTTP/2. Returns a boolean value.
<code>fastly_info.h2.is_push</code>	If the request was made over HTTP/2, whether or not the request is part of a push request. Returns a boolean value.
<code>fastly_info.h2.stream_id</code>	If the request was made over HTTP/2, the underlying HTTP/2 stream id.
<code>tls.client.protocol</code>	The TLS protocol version that this connection is speaking over. Example: <code>"TLSv1.2"</code>
<code>tls.client.servername</code>	The Server Name Indication (SNI) that the client sent in the <code>ClientHello</code> TLS record. Returns <code>""</code> if the client did not send SNI. Returns <code>NULL</code> (the undefined string) if the request is not a TLS request.
<code>tls.client.cipher</code>	The cipher suite used to secure the client TLS connection. Example: <code>"ECDHE-RSA-AES128-GCM-SHA256"</code>

Feature	Description
<code>tls.client.ciphers_sha</code>	A SHA1 of the two-byte cipher suite identifiers sent from the client as part of the TLS handshake, represented in base64.
<code>tls.client.tlsexts_sha</code>	A SHA1 of the TLS extension identifiers sent from the client as part of the TLS handshake, represented in base64.

## § Understanding the different PASS action behaviors (/guides/vcl/understanding-the-different-pass-action-behaviors)

Passing with a request setting and with a cache setting triggers very different behavior in Varnish (/guides/vcl/guide-to-vcl). Within VCL, passing with a request setting is the same as `return(pass)` in `vcl_recv`. Passing with a cache setting is the same as `return(pass)` in `vcl_fetch`. If you are familiar with Varnish 3+, passing with a cache setting is equivalent to `return(hit_for_pass)`.

### Using a request setting

### Create a new request setting

CONDITION

This will happen all the time unless you [Attach a condition](#)

Action

Pass (do not cache)

Force a VCL action to happen. Options are:  
**Do nothing now** - No action.  
**Lookup** - Force an immediate lookup in the cache.  
**Pass** - Force a request to the origin server.  
[Learn about how are request settings applied.](#)

Force TLS

Not selected

Select **Yes** to force unencrypted requests over to TLS, return a **301 Moved Permanently** response to any unencrypted request, and redirect to the TLS equivalent. See our [TLS encryption guide](#) for more info.

X-Forwarded-For

Leave

Stipulate how the **X-Forwarded-For** header should be treated. [Learn about how to manipulate the X-Forwarded-For header.](#)  
[Details for X-Forwarded-For options](#)

Description

★ Required

The name of your request setting, such as **My request setting**.

Passing with a request setting translates within your generated VCL to `return(pass)` in `vcl_recv`. Varnish will not perform a lookup to see if an object is in cache and the response from the origin will not be cached.

Passing in this manner disables request collapsing. Normally simultaneous requests for the same object that result in cache misses will be collapsed down to a single request to the origin. While the first request is sent to the origin, the other requests for that object are queued until a response is received. When requests are passed in `vcl_recv`, they will all go to the origin separately without being collapsed.

### Using a cache setting

## Create a new cache setting

This will override your cache control headers. In most cases, use with conditions applied.

CONDITION

This will happen all the time unless you [Attach a condition](#)

TTL (seconds)

The TTL (Time To Live) entered will set the lifespan of the object within our cache nodes.

Action

This setting decides [how the request will be handled](#).

Stale TTL (seconds)

Stale TTL (Time To Live) is used by your application to [serve stale content](#). It determines how long to keep stale data after it has expired. Note there's custom VCL required to [complete this setup](#).

Description

★ Required

The name of your cache setting, such as **My cache setting**.

CREATE

CANCEL

Passing with a cache setting translates within your generated VCL to `return(pass)` in `vcl_fetch`. At this point in the flow of a request, Varnish has performed a lookup and determined that the object is not in cache. A request to the origin has been made; however, in `vcl_fetch` we have determined that the response is not cacheable. In Fastly's default VCL, this can happen based on the presence of a `Set-Cookie` response header from the origin.

Passing in `vcl_fetch` is often not desirable because request collapsing is *not* disabled. This makes sense since Varnish is not aware in `vcl_recv` that the object is uncacheable. On the first request for an object that will be later passed in `vcl_fetch`, all other simultaneous cache misses will be queued. Once the response from the origin is received and Varnish has realized that the request should be passed, the queued requests are sent to the origin.

This creates a scenario where two users request an object at the same time, and one user must wait for the other before being served. If these requests were passed in `vcl_recv`, neither user would need to wait.

To get around this disadvantage, when a request is passed in `vcl_fetch`, Varnish creates what is called a hit-for-pass object. These objects have their own TTLs and while they exist, Varnish will pass any requests for them as if the pass had been triggered in `vcl_recv`. For this reason, it is important to set a TTL that makes sense for your case when you pass in `vcl_fetch`. All future requests for the object will be passed until the hit-for-pass object expires. Hit-for-pass objects can also be purged like any other object.

Even with this feature, there will be cases where simultaneous requests will be queued and users will wait. Whenever there is not a hit-for-pass object in cache, these requests will be treated as if they are normal cache misses and request collapsing will be enabled. Whenever possible it is best avoid relying on passing in `vcl_fetch`.

## Using `req.hash_always_miss` and `req.hash_ignore_busy`

Setting `req.hash_always_miss` forces a request to miss whether it is in cache or not. It is different when passing in `vcl_recv` in that the response will be cached and request collapsing will not be disabled. Later on the request can still be passed in `vcl_fetch` if desired.

A second relevant variable is `req.hash_ignore_busy`. Setting this to true disables request collapsing so that each request is sent separately to origin. When `req.hash_ignore_busy` is enabled all responses will be cached and each response received from the origin will overwrite the last. Future requests for the object that are served from cache will receive the copy of the object from the last cache miss to complete. `req.hash_ignore_busy` is used mostly for avoiding deadlocks in complex multi-Varnish setups.

Setting both these variables can be useful to force requests to be sent separately to the origin while still caching the responses.

## § Uploading custom VCL (/guides/vcl/uploading-custom-vcl)

Fastly allows you create your own Varnish Configuration Language (VCL) files with specialized configurations. By uploading custom VCL files, you can use custom VCL and Fastly VCL together at the same time (</guides/vcl/mixing-and-matching-fastly-vcl-with-custom-vcl/>). Keep in mind that your custom VCL always takes precedence over VCL generated by Fastly.

**❗ IMPORTANT:** The ability to [upload custom VCL code](/guides/vcl/uploading-custom-vcl/) (</guides/vcl/uploading-custom-vcl/>) is disabled by default when you sign up. To enable this ability for your account, contact [support@fastly.com](mailto:support@fastly.com) (<mailto:support@fastly.com>) and request it.

## Uploading a VCL file

After Fastly support has enabled the ability to upload custom VCL to your account, follow these instructions to upload a custom VCL file:

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.
4. Click the **Custom VCL** tab. The Custom VCL page appears.
5. Click the **Upload a new VCL file** button. The Upload a new VCL file page appears.

6. In the **Name** field, enter the name of the VCL file. For included files, this name must match the include statement in the main VCL file. See how to include additional VCL configurations for more information.
7. Click **Upload file** and select a file to upload. The name of the uploaded file appears next to the button.

**❗ IMPORTANT:** Don't upload generated VCL that you've downloaded from the Fastly web interface. Instead, edit and then upload a copy of Fastly's [VCL boilerplate](/guides/vcl/mixing-and-matching-fastly-vcl-with-custom-vcl/#fastly-39-s-vcl-boilerplate) (</guides/vcl/mixing-and-matching-fastly-vcl-with-custom-vcl/#fastly-39-s-vcl-boilerplate>) to avoid errors.

8. Click the **Create** button. The VCL file appears in the Varnish Configurations area.

9. Click the **Activate** button to deploy your configuration changes.

## Editing a VCL file

To edit an existing VCL file, follow these instructions:

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.



4. Click the **Custom VCL** tab. The Custom VCL page appears.
5. In the **Varnish Configurations** area, click the VCL file you want to edit. The Edit an existing VCL file page appears.

## Edit an existing VCL file

Our [VCL tutorial](#) will help you get started with creating VCL files.

Name
★ Required

For included files, this name must exactly match the include statement in the main VCL file.

Existing file
[View Source](#)
[Download](#)

Config file

6. In the **Name** field, optionally enter a new name of the VCL file.
7. Click the **Download** link to download the appropriate file.
8. Make the necessary changes to your file and save them.
9. Click the **Replace file** button and select the file you updated. The selected file replaces the current VCL file and the file name appears next to the button.
10. Click the **Update** button to update the VCL file in the Fastly application.
11. Click the **Activate** button to deploy your configuration changes.

## Including additional VCL configurations

You can apply additional VCL files along with your main VCL by including their file names in the main VCL file using the syntax `include "VCL Name"` where `VCL Name` is the name of an included VCL object you've created.

For example, if you've created an included VCL object called "ACL" (to use an access control list (</guides/vcl/using-access-control-lists>) for code manageability) and the file is named `acl.vcl`, your main VCL configuration file would need to contain this line:

```
include "ACL"
```

## § VCL regular expression cheat sheet (</guides/vcl/vcl-regular-expression-cheat-sheet>)

❗ **IMPORTANT:** [Varnish \(/guides/vcl/\)](/guides/vcl/) regular expressions are case sensitive and forward slashes don't need to be escaped.

### Basic matching

```
req.url == "phrase"
```

Matches only if `req.url` is exactly `phrase`.

```
req.url ~ "phrase"
```

Matches `phrase` anywhere.

### Matching at the beginning or end of a string

```
req.http.host ~ "^www"
```

Matches if `req.http.host` starts with `www`.

```
req.url ~ "\.jpg$"
```

Matches if `req.url` ends with `.jpg`.

## Multiple matches

```
req.url ~ "\.(png|jpg|css|js)$"
```

Matches if `req.url` ends with either `.png`, `.jpg`, `.css`, or `.js`.

```
req.url ~ "\.php(?:.*)?$"
```

Matches if `req.url` ends with `.php`, `.php?foo=bar` or `.php?`, but not `.phpa`.

**NOTE:** You can also use `req.url.ext` to find the file extension specified in a URL. For example, in the request `www.example.com/1/hello.gif?foo=bar`, `req.url.ext` will return `gif`. See our [Miscellaneous VCL features \(/guides/vcl/miscellaneous-VCL-extensions\)](/guides/vcl/miscellaneous-VCL-extensions) guide for more information.

```
req.url ~ "\.[abc]server$"
```

Matches if `req.url` ends with `.aserver`, `.bserver` or `.cserver`.

## Matching wildcards

```
req.url ~ "jp.g$"
```

Matches if `req.url` ends with `jpeg`, `jpg`, and `jpg`, but doesn't match if `req.url` ends with `jpg`. It also matches if any other character is between the `jp` and the `g`.

```
req.url ~ "jp.*g$"
```

Matches `jp` followed by 0 or more random characters ending with the letter `g` (`jpeg`, `jpg`, and `jpeeeeg` all match).

## Capturing matches

```
set req.http.Foo = "abbbccccc";
if (req.http.Foo ~ "^(a+)(b+)(c+)") {
  set resp.http.match0 = re.group.0; # now equals 'abbbccccc'
  set resp.http.match1 = re.group.1; # now equals 'a'
  set resp.http.match2 = re.group.2; # now equals 'bbb'
  set resp.http.match3 = re.group.3; # now equals 'ccccc'
}
```

The `re.group.[0-9]` objects allow you to capture matches. The `re.group.0` object evaluates to the entire matched string even if no capture groups have been supplied. You can use these objects to replace this example:

```
if (req.url ~ "(?i)\?.*some_query_arg=([^&]*)") {
  set req.http.Thing-I-Want = regsub(req.url, "(?i)\?.*some_query_arg=([^&]*).*", "\1");
}
```

You can use `re.group` to greatly simplify the previous example:

```
if (req.url ~ "(?i)\?.*some_query_arg=([^&]*)") {
  set req.http.Thing-I-Want = re.group.1;
}
```

You could even get really fancy and do something like this:

```
set req.http.Thing-I-Want = if(req.url ~ "(?i)\?.*some_query_arg=([^&]*)", re.group.1, "");
```

## Replacing content

```
set req.http.host = regsub(req.http.host, "^www\.","");
```

Removes a leading `www.` from the host, if present.

```
set req.http.api-test = regsub(req.http.host, "^www\.","api.");
```

Sets the `api-test` header to contain the host-header, but replaces a leading `www.` with `api:`

```
Host: www.example.com ->
Host: www.example.com
api-test: api.example.com
Host: example.com ->
Host: example.com
api-test: example.com
```

• [Guides \(/guides/\)](#) > [Developer's tools \(/guides/devtools\)](#) > [VCL Snippets \(/guides/vcl-snippets/\)](#)

## § About VCL Snippets (/guides/vcl-snippets/about-vcl-snippets)

VCL Snippets allow you to insert small chunks of VCL code into your Fastly service configurations. You can use these snippets without having to follow the more complex method of mixing and matching custom VCL ([/guides/vcl/mixing-and-matching-fastly-vcl-with-custom-vcl](#)) with regular VCL settings. Should you choose to, however, both types of VCL Snippets can be included in customer VCL.

**❗ IMPORTANT:** This feature is part of a limited availability release. For more information, see our [product and feature life cycle \(/guides/customer-support/product-and-feature-life-cycles#limited-availability\)](#) descriptions.

### Types of VCL Snippets

Fastly supports two types of VCL Snippets:

- **Regular VCL Snippets ([/guides/vcl-snippets/using-regular-vcl-snippets](#))** get created as you create versions of your Fastly configurations. In effect, they belong to a specific service and any modifications you make to the snippet are locked and deployed when you deploy the a new service configuration. You can treat regular snippets like any other Fastly objects, which continue to be deployed with a service until you specifically delete them.
- **Dynamic VCL Snippets ([/guides/vcl-snippets/using-dynamic-vcl-snippets](#))** can be modified and deployed any time they're changed. We treat them as versionless objects, much like Edge Dictionaries ([/guides/edge-dictionaries/](#)) or Edge ACLs ([/guides/access-control-lists/](#)). Versionless objects like dynamic snippets can be modified independently from service changes, allowing you to modify code rapidly without deploying a service version that may not be ready for production.

### Parameters

These are the parameters to specify when creating or updating VCL Snippets:

Parameter	Description
<code>name</code>	Describes your snippet.
<code>dynamic</code>	Sets the snippet version to regular ( <code>0</code> ) or dynamic ( <code>1</code> ).
<code>type</code>	Specifies the location in which to place the snippet. Valid values include <code>init</code> (above all subroutines), <code>recv</code> (within the <code>vcl_recv</code> subroutine but below the boilerplate VCL ( <a href="#">/guides/vcl/mixing-and-matching-fastly-vcl-with-custom-vcl#fastly-39-s-vcl-boilerplate</a> ) and above any objects), <code>hit</code> , <code>miss</code> , <code>pass</code> , <code>fetch</code> , <code>error</code> , <code>deliver</code> , <code>log</code> , and <code>none</code> (see the info on includes below).
<code>content</code>	Specifies exactly what the snippet does, via the format <code>if ( any valid piece of VCL for the type specified ){ command to execute }</code> .

### Including snippets in custom VCL

By specifying a location of `none` for the `type` parameter, snippets will not be rendered in VCL. This allows you to include snippets in custom VCL using the following syntax for both regular and dynamic snippets:

```
include "snippet::<snippet name>"
```

The same VCL Snippet can be included in custom VCL in as many places as needed.

### Example use cases

The follow examples detail how each type of VCL Snippet might be used.

#### Using regular VCL Snippets for location-based redirection

Say that you work at a large content publisher and you want to redirect users to different editions of your publication depending on which country their request comes from. Say also that you want the ability to override the edition you deliver to them based on a cookie.

Using regular VCL snippets (</guides/vcl-snippets/using-regular-vcl-snippets/>), you could add a new object with the relevant VCL as follows:

```
if (req.http.Cookie:edition == "US" || geoip.country_code == "US" || ) {
  set req.http.Edition = "US";
  set req.backend = F_US;
} elseif (req.http.Cookie:edition == "Europe" || server.region - "^EU-" ) {
  set req.http.Edition = "EU";
  set req.backend = F_European;
} else {
  set req.http.Edition = "INT";
  set req.backend = F_International;
}
```

This creates an Edition header in VCL but allows you to override it by setting a condition. You would add the Edition header into Vary (<https://www.fastly.com/blog/best-practices-for-using-the-vary-header>) and then add a false condition (</guides/conditions/about-conditions#operators-for-complex-conditions>) (e.g., `[!req.url]`) to your other backends to ensure the correct edition of your publication gets delivered. (Remember: VCL Snippets get added to VCL before backends are set.)

## Using dynamic VCL Snippets to block site scrapers

Say you wanted to implement some pattern matching against incoming requests to block someone trying to scrape your site. Say also that you've developed a system that looks at all incoming requests and generates a set of rules that can identify scrapers using a combination of the incoming IP address, the browser, and the URL they're trying to fetch. Finally, say that the system updates the rules every 20 minutes.

If, during system updates, your colleagues are also making changes to the rest of your Fastly configuration, you probably don't want the system to automatically deploy the latest version of the service since it might be untested. Instead you could generate the rules as a Dynamic VCL Snippet (</guides/vcl-snippets/using-dynamic-vcl-snippets/>). Whenever the snippet is updated, all other logic remains the same as the currently deployed version and only your rules are modified.

## Limitations of VCL Snippets

- Snippets can only be 64kb long.
- Snippets do not currently support conditions, though if statements can be used instead.
- Snippets cannot currently be shared between services.

## § Using dynamic VCL Snippets (</guides/vcl-snippets/using-dynamic-vcl-snippets/>)

Fastly's VCL Snippets (</guides/vcl-snippets/about-vcl-snippets/>) allow you to create complex configuration objects at the edge without having to mix and match custom VCL (</guides/vcl/mixing-and-matching-fastly-vcl-with-custom-vcl/>) with regular VCL settings in the UI.

**❗ IMPORTANT:** This feature is part of a limited availability release. For more information, see our [product and feature life cycle](/guides/customer-support/product-and-feature-life-cycles#limited-availability) (</guides/customer-support/product-and-feature-life-cycles#limited-availability>) descriptions.

## Creating and using a dynamic VCL Snippet

Using the cURL command line tool, make the following API call in a terminal application:

```
curl -X POST -s https://manage.fastly.com/service/<Service ID>/version/<Editable Version #>/snippet -H "Fastly-Key:<API TOKEN>" -H 'Content-Type: application/x-www-form-urlencoded' --data $'name=my_dynamic_snippet&type=recv&dynamic=1&content=if ( req.url ) {\n set req.http.my-snippet-test-header = "true";\n}';
```

Fastly returns a JSON response that looks like this:

```
{
  "service_id": "<Service Id>",
  "version": "<Editable Version>",
  "name": "my_dynamic_snippet",
  "type": "recv",
  "priority": 100,
  "dynamic": 1,
  "content": null,
  "id": "decafbad12345",
  "created_at": "2016-09-09T20:34:51+00:00",
  "updated_at": "2016-09-09T20:34:51+00:00",
  "deleted_at": null
}
```

**NOTE:** The returned JSON includes `"content": null`. This happens because the content is stored in a separate, unversioned object.

## Listing all dynamic VCL Snippets

To list all dynamic VCL Snippets attached to a service, make the following API call in a terminal application:

```
curl -X GET -s https://manage.fastly.com/service/<Service ID>/version/<Editable Version #>/snippet/ -H "Fastly-Key:<API TOKEN>"
```

## Fetching an individual dynamic VCL Snippet

To fetch an individual snippet, make the following API call in a terminal application:

```
curl -X GET -s https://manage.fastly.com/service/<Service ID>/snippet/<Snippet ID e.g decafbad12345> -H "Fastly-Key:<API TOKEN>"
```

Unlike fetching regular VCL Snippets (`/guides/vcl-snippets/using-regular-vcl-snippets#fetching-an-individual-regular-vcl-snippet`), you do not include the version in the URL and you must use the ID returned when the snippet was created, not the name.

## Updating an existing dynamic VCL Snippet

To update an individual snippet, make the following API call in a terminal application:

```
curl -X PUT -s https://manage.fastly.com/service/<Service ID>/snippet/<Snippet ID e.g decafbad12345> -H "Fastly-Key:<API TOKEN>" -H 'Content-Type: application/x-www-form-urlencoded' --data $'content=if ( req.url ) {\n set req.http.my-snippet-test-header = \"affirmative\";\n}';
```

## Deleting an existing regular VCL Snippet

To delete an individual snippet, make the following API call in a terminal application:

```
curl -X DELETE -s https://manage.fastly.com/service/<Service ID>/version/<Editable Version #>/snippet/<Snippet Name e.g my_dynamic_snippet> -H "Fastly-Key:<API TOKEN>"
```

## § Using regular VCL Snippets (`/guides/vcl-snippets/using-regular-vcl-snippets`)

Fastly's VCL Snippets (`/guides/vcl-snippets/about-vcl-snippets`) allow you to create complex configuration objects at the edge without having to mix and match custom VCL (`/guides/vcl/mixing-and-matching-fastly-vcl-with-custom-vcl`) with regular VCL settings in the web interface.

**IMPORTANT:** This feature is part of a limited availability release. For more information, see our [product and feature life cycle](/guides/customer-support/product-and-feature-life-cycles#limited-availability) (`/guides/customer-support/product-and-feature-life-cycles#limited-availability`) descriptions.

## Creating and using a regular VCL Snippet

Using the cURL command line tool, make the following API call in a terminal application:

```
curl -X POST -s https://manage.fastly.com/service/<Service ID>/version/<Editable Version #>/snippet -H "Fastly-Key:<API TOKEN>" -H 'fastly-cookie' -H 'Content-Type: application/x-www-form-urlencoded' --data $'name=my_regular_snippet&type=recv&dynamic=0&content=if ( req.url ) {\n set req.http.my-snippet-test-header = \"true\";\n}';
```

Fastly returns a JSON response that looks like this:

```
{
  "service_id": "<Service Id>",
  "version": "<Editable Version>",
  "name": "my_regular_snippet",
  "type": "recv",
  "content": "if ( req.url ) {\n set req.http.my-snippet-test-header = \"true\";\n}",
  "priority": 100,
  "dynamic": 0,
  "id": "56789exampleid",
  "created_at": "2016-09-09T20:34:51+00:00",
  "updated_at": "2016-09-09T20:34:51+00:00",
  "deleted_at": null
}
```

**NOTE:** When regular VCL snippets get created, an `id` field will be returned that isn't used. It only applies to [dynamic VCL Snippets](/guides/vcl-snippets/using-dynamic-vcl-snippets) (</guides/vcl-snippets/using-dynamic-vcl-snippets>). In addition, the returned JSON includes a populated `"content"` field. This happens because the content is stored in a versioned object.

## Listing all regular VCL Snippets

To list all regular VCL Snippets attached to a service, make the following API call in a terminal application:

```
curl -X GET -s https://manage.fastly.com/service/<Service ID>/version/<Editable Version #>/snippet/ -H "Fastly-Key:<API TOKEN>"
```

## Fetching an individual regular VCL Snippet

To fetch an individual snippet, make the following API call in a terminal application:

```
curl -X GET -s https://manage.fastly.com/service/<Service ID>/version/<Editable Version #>/snippet/<Snippet Name e.g my_regular_snippet> -H "Fastly-Key:<API TOKEN>"
```

Unlike fetching dynamic VCL Snippets (</guides/vcl-snippets/using-dynamic-vcl-snippets#fetching-n-individual-dynamic-regular-vcl-snippet>) you include the version in the URL and you must use the name of the snippet, not the ID.

## Updating an existing regular VCL Snippet

To update an individual snippet, make the following API call in a terminal application:

```
curl -X PUT -s https://manage.fastly.com/service/<Service ID>/version/<Editable Version #>/snippet/<Snippet Name e.g my_regular_snippet> -H "Fastly-Key:<API TOKEN>" -H 'Content-Type: application/x-www-form-urlencoded' --data $'content=if ( req.url ) {\n set req.ht\n tp.my-snippet-test-header = \"affirmative\";\n}';
```

## Deleting an existing regular VCL Snippet

To delete an individual snippet, make the following API call in a terminal application:

```
curl -X DELETE -s https://manage.fastly.com/service/<Service ID>/version/<Editable Version #>/snippet/<Snippet Name e.g my_regular_snippet> -H "Fastly-Key:<API TOKEN>"
```

• [Guides \(/guides/\)](/guides/) > [Developer's tools \(/guides/devtools/\)](/guides/devtools/) > [Tutorials \(/guides/tutorials/\)](/guides/tutorials/)

## § Cache control tutorial (</guides/tutorials/cache-control-tutorial>)

You are in full control of how Fastly caches your resources. The most preferred way of instructing Fastly is to use backend HTTP headers. The other way is to use the Varnish Configuration Language (</guides/vcl/guide-to-vcl>) (VCL).

**IMPORTANT:** The ability to [upload custom VCL code \(/guides/vcl/uploading-custom-vcl\)](/guides/vcl/uploading-custom-vcl) is disabled by default when you sign up. To enable this ability for your account, contact [support@fastly.com](mailto:support@fastly.com) (<mailto:support@fastly.com>) and request it.

## Backend HTTP headers

You can set four different types of HTTP headers which will have different effects on our caches and on web browsers. If you use more than one type, they are prioritized in the order listed below:

## Surrogate-Control

Format:

```
Surrogate-Control: max-age=(time in seconds)
```

Example:

```
Surrogate-Control: max-age=3600
```

will cache something on our caches for one hour. This header gets stripped and is only visible to Fastly caches.

## Cache-Control: s-maxage

Format:

```
Cache-Control: s-maxage=(time in seconds)
```

This is the same as `Surrogate-Control`, except the header is not stripped and will be respected by Fastly caches and any caches between Fastly and the browser, but not the browser itself.

## Cache-Control: max-age

Format:

```
Cache-Control: max-age=(time in seconds)
```

This header will be respected by Fastly caches, any caches between Fastly and the browser, and the browser itself.

## Expires

This header caches content until it expires as specified. It will be respected by Fastly caches, any caches between Fastly and the browser and the browser itself. See section 14.21 of RFC2616 (<http://www.w3.org/Protocols/rfc2616/rfc2616-sec14.html#sec14.21>) for an explanation of the format.

## Do not cache

If you don't want certain resources cached, set the header as follows:

```
Cache-Control: private
```

If you just set `max-age=0` or an `Expires` in the past, there is still no guarantee that the content won't be used to satisfy multiple outstanding requests at the same time. Also, an expired (stale) object might be used in case of errors or while a request to your backend for the same object is already in progress.

## When Cache-Control and Surrogate-Control headers co-exist

Say that you want Fastly to cache your resources forever but send headers to browsers so that they don't cache it at all (so that every browser miss hits Fastly but isn't a cache miss from your service).

The best way to do this would be to send Fastly both the `Cache-Control` header as you want it to go to the browsers, and use `Surrogate-Control` to tell us how long to cache for. For example:

```
Cache-Control: no-cache, no-store, must-revalidate
Surrogate-Control: max-age=3600
Pragma: no-cache
Expires: 0
```

Except for when the `Cache-Control` header is set to `private`, the `Surrogate-Control` header takes priority over `Cache-Control`, but unlike `Cache-Control` it is stripped so the browsers don't see it.

## Example backend configs

### Apache Config

If you are using Apache, the easiest way to add headers is to use the `mod_expires` module ([http://httpd.apache.org/docs/current/mod/mod\\_expires.html](http://httpd.apache.org/docs/current/mod/mod_expires.html)). For example, to cache GIF images for 75 minutes after the image was last accessed by the cache server, you would add a directive like this under the `VirtualHost` (or globally). For example:

```
ExpiresActive On
ExpiresByType image/gif "access plus 1 hours 15 minutes"
```

You can also cache whole URL subtrees. For example:

```
<Location "/css">
  ExpiresActive On
  ExpiresDefault "access plus 1 year"
</Location>

<Location "/static/">
  ExpiresActive On
  ExpiresDefault "access plus 1 day"
</Location>
```

## NGINX Configuration

To configure NGINX, add the expires directive. For example:

```
location ~* \.(js|css|png|jpg|jpeg|gif|ico)$ {
    expires 1h;
}
```

Alternatively, if you need more flexibility in modifying headers you can try the `HttpHeadersMore` Module ([http://nginx.org/en/docs/http/nginx\\_headers\\_module.html](http://nginx.org/en/docs/http/nginx_headers_module.html)).

## Amazon S3 configuration

By default, S3 doesn't have a facility for setting Cache-Control headers across multiple objects, so you will have to do this file-by-file using the S3Explorer, or in an automated fashion by using a cloud library like boto. Remember that you can combine long cache time with instant purges to enhance your performance.

★ **TIP:** While it's difficult to get S3 to set `Surrogate-Control`, you can set `x-amz-meta-surrogate-control` (</guides/purging/setting-surrogate-key-headers-for-amazon-s3-origins>) instead on origin and Fastly will honor that.

```
from boto.s3.connection import S3Connection

connection = S3Connection('aws access key', 'aws secret key')

buckets = connection.get_all_buckets()

for bucket in buckets:
    for key in bucket.list():
        print('%s' % key)

        if key.name.endswith('.jpg'):
            contentType = 'image/jpeg'
        elif key.name.endswith('.png'):
            contentType = 'image/png'
        else:
            continue

        key.metadata.update({
            'Content-Type': contentType,
            'Cache-Control': 'max-age=864000'
        })
        key.copy(
            key.bucket.name,
            key.name,
            key.metadata,
            preserve_acl=True
        )
```

❗ **IMPORTANT:** The above example provides an S3 configuration option for customers with small- to medium-sized buckets. However, it iterates over every object in those buckets. If you have millions of objects this may not be the right approach. For millions of objects, we recommend using [VCL](/guides/vcl/guide-to-vcl) (</guides/vcl/guide-to-vcl>). Be sure to [contact us \(mailto:support@fastly.com\)](mailto:support@fastly.com) for assistance.

## Custom Headers in Programming Languages and Frameworks

### PHP

More information: <http://php.net/manual/en/function.header.php> (<http://php.net/manual/en/function.header.php>)

Example: add this to your PHP code before you send any output to cache certain HTML for an hour



```
header('Cache-Control: max-age=3600');
```

## Django

More information: <https://docs.djangoproject.com/en/dev/ref/request-response/#setting-headers>  
(<https://docs.djangoproject.com/en/dev/ref/request-response/#setting-headers>)

Example:

```
response = HttpResponse()
response['Cache-Control'] = 'max-age=3600'
```

## Sinatra

More information: <http://sinatra.rubyforge.org/doc/> (<http://sinatra.rubyforge.org/doc/>)

Example:

```
get '/' do
  headers['Cache-Control'] = 'max-age=3600'
end
```

★ **TIP:** Expiration times in these examples are provided for guidance only. You can use longer expirations coupled with our [purging API](#) ([/api/purge](#)) to make your site faster and your backend less loaded.

# § Enabling URL token validation (/guides/tutorials/enabling-url-token-validation)

Token validation allows you to create URLs that expire. Tokens are generated within your web application and appended to URLs in a query string. Requests are authenticated at Fastly's edge instead of your origin server. When Fastly receives a request for the URL, the token is validated before serving the content. After a configurable period of time, the token expires.

## Adding custom VCL

To enable token validation, you'll need to create a Varnish configuration ([/guides/vcl/uploading-custom-vcl](#)) named `vcl_recv` and add the following example code to it.

❗ **IMPORTANT:** The ability to [upload custom VCL code](#) ([/guides/vcl/uploading-custom-vcl](#)) is disabled by default when you sign up. To enable this ability for your account, contact [support@fastly.com](mailto:support@fastly.com) (<mailto:support@fastly.com>) and request it.

```
/* make sure there is a token */
if (req.url != ".+\.?\.?token=(\d{10,11})_([^\&]+)") {
  error 403;
}

/* extract token expiration and signature */
set req.http.X-Exp = re.group.1;
set req.http.X-Sig = re.group.2;

/* validate signature */
if (req.http.X-Sig == regsub(digest.hmac_shal(digest.base64_decode("YOUR%SECRET%KEY%IN%BASE64%HERE"),
req.url.path req.http.X-Exp), "^0x", "")) {

  /* check that expiration time has not elapsed */
  if (time.is_after(now, std.integer2time(std.atol(req.http.X-Exp)))) {
    error 410;
  }

} else {
  error 403;
}

/* cleanup variables */
unset req.http.X-Sig;
unset req.http.X-Exp;
```

**⚠ WARNING:** You must replace `YOUR%SECRET%KEY%IN%BASE64%HERE` in the example VCL with your own randomly generated secret key. The example key will intentionally cause an error if you use it. Anyone who learns this key can bypass your token validation, so it's critical that you keep this key secret.

The key found in `digest.hmac_sha1` can be any string. The one in this example was generated with the command `openssl rand -base64 32`.

A token is expected in the `?token=` GET parameter. Tokens take the format `[expiration]_[signature]` and look like this:

```
1441307151_4492f25946a2e8e1414a8bb53dab8a6ba1cf4615
```

The full request URL looks like this:

```
http://www.example.com/foo/bar.html?token=1441307151_4492f25946a2e8e1414a8bb53dab8a6ba1cf4615
```

## What the custom VCL does

The custom VCL code checks for two things:

- Is the current time greater than the expiration time specified in the token?
- Does our signature match the signature of the token?

If the signature is invalid, Varnish returns a 403 response. If the signature is valid but the expiration time has elapsed, Varnish returns a 410 response. The different response codes are helpful for debugging.

## Configuring your application

You'll need to write custom code in your application to generate tokens and authenticate with Varnish. We provide examples in our token functions (<https://github.com/fastly/token-functions>) repository on GitHub. Review the examples in the repository to learn how to generate custom tokens within your application.

## Testing

To test your configuration, append a token generated by your application to a URL in a query string. For example:

```
http://www.example.com/foo/bar.html?token=1441307151_4492f25946a2e8e1414a8bb53dab8a6ba1cf4615
```

If the token is valid, you will receive a normal response. If it is invalid, you will receive a 403 response.

## Troubleshooting NUL bytes

You should verify that your secret key is devoid of NUL bytes. If the base64-decoded string contains a NUL byte (0x00), then that byte and any bytes following it will not be included in the response. See [base64 decoding \(/guides/vcl/cryptographic-and-hashing-related-vcl-functions#base64-decoding\)](/guides/vcl/cryptographic-and-hashing-related-vcl-functions#base64-decoding) for more information.

- [Guides \(/guides/\)](/guides/) > [Diagnostics and performance \(/guides/diagnostics\)](/guides/diagnostics/) > [Streaming logs \(/guides/streaming-logs/\)](/guides/streaming-logs/)

## § Changing log line formats (/guides/streaming-logs/changing-log-line-formats)

Fastly's Real-Time Log Streaming (</guides/about-fastly-services/about-fastlys-realtime-log-streaming-features/>) feature allows you to change the format that your log messages are delivered in. By default, we send log messages out in standard syslog format. The prefix for this format (as defined in RFC 3164 (<https://tools.ietf.org/html/rfc3164>)) appears as follows:

```
<134>2016-07-04T22:37:26Z cache-sjc3128 LogTest[62959]: <your log message>
```

The prefix begins with the message priority (always `<134>`, which means `Facility=Local0, Severity=Informational`), followed the date and time the log was sent (`2016-07-04T22:37:26Z`), the cache node it came from (in this case `cache-sjc3128`), the name of your log (`LogTest`) and the ID of the process sending it (`62959`).

## Available message formats

Although the default message prefix works for most logging services and processors, we allow you to choose one of several formats:

- `classic` is the default prefix format. A standard syslog prefix as defined by RFC 3164 (<https://tools.ietf.org/html/rfc3164>).

- `loggly` is a structured syslog prefix format based on RFC 5424 (<https://tools.ietf.org/html/rfc5424>).
- `logplex` is a Heroku-style length prefixed syslog format ([https://github.com/heroku/logplex/blob/master/doc/README.http\\_drains.md](https://github.com/heroku/logplex/blob/master/doc/README.http_drains.md)).
- `blank` means no prefix. Just your log message. Useful when writing to JSON and CSV files.

## Updating endpoints to use a different format

The following logging endpoints can be updated to use a message format other than the default:

- Amazon S3 (</guides/streaming-logs/log-streaming-amazon-s3>)
- Cloud Files (</guides/streaming-logs/log-streaming-cloudfiles>)
- FTP (</guides/streaming-logs/log-streaming-ftp>)
- Google Cloud Storage (</guides/streaming-logs/log-streaming-google-cloud-storage>)
- OpenStack (</guides/streaming-logs/log-streaming-openstack>)
- SFTP (</guides/streaming-logs/log-streaming-sftp>)
- Sumo Logic (</guides/streaming-logs/log-streaming-sumologic>)
- Syslog (</guides/streaming-logs/log-streaming-syslog>)

To update one of the above logging endpoints, run the following command:

```
curl -X PUT -H 'Fastly-Key: <your Fastly API key>' -H 'Content-Type: application/json' 'https://api.fastly.com/service/<your Fastly service ID>/version/<version number>/logging/<log type>/<log name>' --data-binary '{"message_type": "<type>"}'
```

where `log type` is one of the eligible endpoint types:

- `s3`
- `cloudfiles`
- `ftp`
- `gcs`
- `openstack`
- `sftp`
- `sumologic`
- `syslog`

Keep in mind that the `message_type` field is a per-object field. Updating it on one logging object *will not* change it on any other objects. For example, to update a Google Cloud Storage endpoint to the `blank` message type the cURL command would look something like this if the endpoint was named "GCS Test":

```
curl -X PUT -H 'Fastly-Key: d3cafb4dde4dbeef' -H 'Content-Type: application/json' 'https://api.fastly.com/service/SU1Z0isxPaozGVKXdv0eY/version/1/logging/gcs/GCS%20Test' --data-binary '{"message_type": "blank"}'
```

**NOTE:** The log name `GCS Test` needed to be URL encoded (<https://en.wikipedia.org/wiki/Percent-encoding>), which turned the space into `%20`.

## § Custom log formats (/guides/streaming-logs/custom-log-formats)

Fastly provides two versions of custom log formats. All logs use the version 1 custom log format by default, but you can upgrade to the version 2 custom log format and then make version 2 look like version 1 for the sake of continuity. We've described the key advantages of the version 2 custom log format below.

**WARNING:** Logging objects using version 2 formatting cannot be downgraded to version 1 because version 1 only provides a limited subset of logging directives and is not strictly compatible with Apache log format.

See instructions below for how to upgrade logging object to version 2 custom log formats. For help upgrading contact [support@fastly.com](mailto:support@fastly.com) (<mailto:support@fastly.com>).

## Version 1 log format

This table details version 1 of Fastly's custom log formats.

String	Description
<code>b</code>	The content size of the response, calculated using the <code>Content-Length</code> header rather than actually checking the length of the response (and may therefore be wrong).
<code>h</code>	The remote IP address.
<code>l</code>	The remote log name. Always returns the hardcoded value <code>"-"</code> .
<code>r</code>	The HTTP verb and request path (e.g., <code>GET /index.html</code> ). Unlike Apache and version 2 log formats, the protocol version is not included.
<code>&gt;s</code>	The status of the last request.
<code>t</code>	The time the request was received, in Unix <code>ctime</code> format (e.g., <code>Thu, 01 Jan 1970 00:00:00 GMT</code> ) rather than Apache's Standard English format (e.g., <code>01/Jan/1970:00:00:00 -0700</code> ).
<code>u</code>	The remote user. Always returns the hardcoded value <code>"-"</code> .

## Version 2 log format

This table details version 2 of Fastly's custom log formats.

String	Description
<code>%</code>	The percent sign.
<code>a</code>	The client IP address of the request.
<code>A</code>	The local IP address.
<code>B</code>	The size of response in bytes, excluding HTTP headers.
<code>b</code>	The size of response in bytes, excluding HTTP headers. In Common Log Format ( <a href="https://httpd.apache.org/docs/trunk/logs.html#common">https://httpd.apache.org/docs/trunk/logs.html#common</a> ) (CLF), that means a <code>"-"</code> rather than a <code>0</code> when no bytes are sent.
<code>{Foobar}C</code>	The contents of cookie <code>Foobar</code> in the request sent to the server.
<code>D</code>	The time taken to serve the request, in microseconds.
<code>{FOOBAR}e</code>	The contents of the environment variable <code>FOOBAR</code> . Always returns <code>NIL</code> .
<code>f</code>	The filename.
<code>h</code>	The remote hostname.
<code>H</code>	The request protocol.
<code>{Foobar}i</code>	The contents of <code>Foobar:</code> header lines in the request sent to the server.
<code>I</code>	Bytes received, including request and headers. Cannot be zero.
<code>k</code>	The number of keepalive requests handled on this connection. Always returns <code>0</code> .
<code>l</code>	The remote logname (from identd, if supplied). This will return a dash unless <code>mod_ident</code> is present and <code>IdentityCheck</code> is set On. Always returns <code>"-"</code> .
<code>m</code>	The request method.
<code>{Foobar}n</code>	The contents of note <code>Foobar</code> from another module. Always returns <code>NIL</code> .
<code>{Foobar}o</code>	The contents of <code>Foobar:</code> header lines in the reply.
<code>O</code>	Bytes sent, including headers. Cannot be zero.
<code>p</code>	The canonical port of the server serving the request. Always returns <code>80</code> .
<code>{format}p</code>	The canonical port of the server serving the request. Valid formats are <code>canonical</code> , <code>local</code> , or <code>remote</code> . Returns <code>80</code> for HTTP requests and <code>443</code> for HTTPS requests.
<code>P</code>	The process ID of the child that serviced the request. Always returns <code>NIL</code> .
<code>{format}P</code>	The process ID or thread ID of the child that serviced the request. Valid formats are <code>pid</code> , <code>tid</code> , and <code>hextid</code> . Always returns <code>NIL</code> .
<code>q</code>	The query string (prepended with a <code>?</code> if a query string exists, otherwise an empty string).

String	Description
<code>r</code>	The first line of the request.
<code>R</code>	The handler generating the response (if any).
<code>s</code>	The status. For requests that got internally redirected, this is the status of the <i>original</i> request. Use <code>%&gt;s</code> for the final status.
<code>t</code>	The time the request was received, in Standard English format (e.g., <code>01/Jan/1970:00:00:00 -0700</code> ). The last number indicates the timezone offset from GMT.
<code>{format}t</code>	The time, in the form given by <code>format</code> , which should be in <code>strftime(3)</code> format (potentially localized). If the format starts with <code>begin:</code> (the default) the time is taken at the beginning of the request processing. If it starts with <code>end:</code> it is the time when the log entry gets written, close to the end of the request processing. In addition to the formats supported by <code>strftime(3)</code> , the following format tokens are supported: <code>sec</code> (number of seconds since the Epoch), <code>msec</code> (number of milliseconds since the Epoch), <code>usec</code> (number of microseconds since the Epoch), <code>msec_frac</code> (millisecond fraction), and <code>usec_frac</code> (microsecond fraction).
<code>T</code>	The time taken to serve the request, in seconds.
<code>u</code>	The remote user if the request was authenticated. May be bogus if return status ( <code>%s</code> ) is 401 (unauthorized). Always returns <code>"-"</code> .
<code>U</code>	The URL path requested, not including any query string.
<code>v</code>	The canonical <code>ServerName</code> of the server serving the request.
<code>V</code>	The server name according to the <code>UseCanonicalName</code> setting.
<code>{vcl}v</code>	The literal VCL to include without quoting (a Fastly extension).
<code>x</code>	The connection status when response is completed. Statuses include <code>x</code> (connection aborted before the response completed), <code>+</code> (connection may be kept alive after the response is sent), and <code>-</code> (connection will be closed after the response is sent).

## Upgrading endpoints to use version 2 log format

To upgrade a logging endpoints, run the following command:

```
curl -X PUT -H 'Fastly-Key: <your Fastly API key>' -H 'Content-Type: application/json' 'https://api.fastly.com/service/<your Fastly service ID>/version/<version number>/logging/<log type>/<log name>' --data-binary '{"format_version":"2"}'
```

Where `log_type` is type of the endpoint you want to upgrade i.e one of

- `ftp`
- `heroku`
- `logentries`
- `loggly`
- `logshuttle`
- `papertrail`
- `s3`
- `scalyr`
- `sftp`
- `sumologic`
- `syslog`

Keep in mind that the `format_version` field is a per-object field. Updating it on one logging object will *not* change it on any other objects. For example, to upgrade a Google Cloud Storage endpoint the cURL command would look something like this if the endpoint was named "GCS Test":

```
curl -X PUT -H 'Fastly-Key: d3cafb4dde4dbeef' -H 'Content-Type: application/json' 'https://api.fastly.com/service/SU1Z0isxPaozGVKXdv0eY/version/1/logging/gcs/GCS%20Test' --data-binary '{"format_version":"2"}'
```

**NOTE:** The log name `GCS Test` needed to be URL encoded (<https://en.wikipedia.org/wiki/Percent-encoding>), which turned the space into `%20`.

## Determining which logging version is being used

To determine which logging version your service currently uses, issue the following cURL command:

```
curl -X GET -H 'Fastly-Key: <your Fastly API key>' 'https://api.fastly.com/service/<your Fastly service ID>/version/<version number>/logging/<log type>/<log name>'
```

where `<log type>` is `ftp`, `heroku`, `logentries`, `loggly`, `logshuttle`, `papertrail`, `s3`, `scalyr`, `sftp`, `sumologic`, or `syslog`. The `cURL` command will produce JSON output detailing the configuration of your service version. For example:

```
{
  "address": "logs.papertrailapp.com",
  "created_at": "2016-04-01T15:37:30+00:00",
  "deleted_at": null,
  "format": "time.start.msec time.to_first_byte time.elapsed req.body_bytes_read req.bytes_read resp.http.content-length server.region client.ip %>s \"req.request req.url req.proto\" \"req.http.referer\" \"req.http.user-agent\"",
  "format_version": "2",
  "hostname": "logs.papertrailapp.com",
  "name": "fastly",
  "port": "11111",
  "public_key": null,
  "response_condition": "LOG /",
  "service_id": "1a2b3c4d5e6f7g8h9j0k",
  "updated_at": "2016-04-01T19:47:47+00:00",
  "version": "123"
}
```

The `format_version` field displays either a `1` or a `2` as appropriate for the custom log format being used.

## Advantages of using the version 2 custom log format

The key advantages of using the version 2 custom log format include the following:

- Log lines are generated in `vcl_log` instead of `vcl_deliver` to allow us to accurately set the various size variables because `vcl_log` is run after the object has been delivered to the browser.
- The `%t` time directive is compatible with Apache log format. In version 1, we used a non-standard time format.
- The `%r` "first line of request" directive is compatible with Apache log format. In version 1, we incorrectly left off the protocol.
- When using the `%b` directive, which represents the size of a response in bytes, excluding HTTP headers is more accurate. In version 1, we used the reported Content-Length from the origin, which could be inaccurate (especially with ESI (</guides/performance-tuning/using-edge-side-includes>)).
- We've added all Apache logging directives that make sense. In version 1, we used a smaller subset.

## Making version 2 logs look like version 1

The default logging format for version 1 is as follows:

```
%h %l %u %t %r %>s
```

Most of the directives in version 2 are exactly the same - only `%t` and `%r` are different. After you upgrade to version 2 log formats, you can recreate the appearance of version 1 logs using the new `%{...}V` directive, which allows you to specifically include VCL in logging directives:

```
%h %l %u %{now}V %{req.request}V %{req.url}V %>s
```

In addition, if you are using the `%b` directive in version 1, then you can use this directive instead:

```
%{resp.http.Content-Length}V
```

## § Log streaming: Amazon S3 (</guides/streaming-logs/log-streaming-amazon-s3>)

Fastly's Real-Time Log Streaming (</guides/about-fastly-services/about-fastlys-realtime-log-streaming-features>) feature can send log files to Amazon Simple Storage Service (<https://aws.amazon.com/s3/>) (Amazon S3). Amazon S3 is a static file storage service used by developers and IT teams. You can also use the instructions in this guide to configure log streaming to another S3-compatible service.

## Prerequisites

Before adding Amazon S3 as a logging endpoint for Fastly services, we recommend creating an Identity and Access Management (IAM) user in Amazon S3 specifically for Fastly. Grant the user `ListBucket`, `GetObject`, and `PutObject` permissions for the directory in which you want to store logs. For more information, see Amazon's [Getting Your Access Key ID and Secret Access Key](http://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSGettingStartedGuide/AWSCredentials.html) (<http://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSGettingStartedGuide/AWSCredentials.html>) page.

# Adding Amazon S3 as a logging endpoint

After you've registered for an Amazon S3 account and created an IAM user in Amazon S3, follow these instructions to add Amazon S3 as a logging endpoint:

1. Review the information in our Setting Up Remote Log Streaming (/guides/streaming-logs/setting-up-remote-log-streaming) guide.
2. Click the Amazon Web Services logo. The Create a new S3 endpoint page appears.

**S3 ACCOUNT REQUIRED**

Before continuing, you will need the **Access key** and **Secret key** from your IAM console. If you don't have an S3 account, now is the time to set one up - [Sign up for S3.](#)

## Create a new S3 endpoint

Learn the basics in our [S3 logging endpoint documentation.](#)

CONDITION

This will happen all the time unless you [Attach a condition](#)

Description

My S3 endpoint

★ Required

The name of your endpoint, such as **My Endpoint**.

Log format

%h %l %u %t %r %>s

An Apache-style string or VCL variables to use for log formatting (the Apache Common Log format string appears by default). See [Fastly's log files docs](#), [Varnish's descriptions of VCL variables](#), and [Fastly's available VCL variables](#) for more info.

Bucket Name

★ Required

The name of the bucket in which to store the logs.

Access Key

★ Required

The Access Key associated with the target S3 bucket. See [for more info](#).

Secret Key

★ Required

The Secret Key associated with the target S3 bucket. See [Amazon's AWS Getting Started Guide](#) for more info.

Period

★ Required

The frequency for log files to be rotated. **3600** is recommended (default unit is seconds). Number only. No need to enter unit.

3. Fill out the **Create a new S3 endpoint** fields as follows:

- In the **Description** field, type a human-readable name for the endpoint.
- In the **Log Format** field, optionally type an Apache-style string or VCL variables to use for log formatting. The Apache Common Log format string appears in this field by default.
- In the **Bucket Name** field, type the name of the Amazon S3 bucket in which to store the logs.
- In the **Access Key** field, type the access key associated with the Amazon S3 bucket. See Amazon's S3 Getting Started Guide (<http://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSGettingStartedGuide/AWSCredentials.html>) for more information.
- In the **Secret Key** field, type the secret key associated with the Amazon S3 bucket. See Amazon's S3 Getting Started Guide (<http://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSGettingStartedGuide/AWSCredentials.html>) for more information.
- In the **Period** field, optionally type an interval (in seconds) to control how frequently your log files are rotated. This value defaults to 3600 seconds.

4. Review the **Advanced Options** of the **Create a new S3 endpoint** page and decide which of the optional fields to change, if any.

ADVANCED OPTIONS

Path

The path within the bucket in which to place the files. If empty, files will be placed in its root. Ends with trailing slash. Not required.

Domain

The region-specific endpoint for your domain. If your S3 bucket was not created in the US Standard region, set as per [Amazon's Web Services](#).

Gzip Level

0

The level of gzip compression, if any. Ranges from 0 (fastest and not compressed) to 9 (slowest and most compressed).

Redundancy Level

Standard

The S3 redundancy level to store logs with. Defaults to **Standard**. See [Amazon's docs](#) on using reduced redundancy storage for more information.

5. Fill out the **Advanced Options** of the **Create a new S3 endpoint** page as follows:

- In the **Path** field, optionally type the path within the bucket to store the files. The path ends with a trailing slash. If this field is left empty, the files will be saved in the bucket's root path.
- In the **Domain** field, optionally type the domain of the Amazon S3 endpoint. If your Amazon S3 bucket was not created in the US Standard region, you must set the domain to match the appropriate endpoint URL. Use the table in the S3 section of the Regions and Endpoints ([http://docs.aws.amazon.com/general/latest/gr/rande.html#s3\\_region](http://docs.aws.amazon.com/general/latest/gr/rande.html#s3_region)) Amazon S3 documentation page. If you want to use an S3-compatible storage system (such as Dreamhost's DreamObjects (<https://www.dreamhost.com/cloud/storage/>)), set the domain to match the domain name for that service (for example, in the case of DreamObjects, the domain name would be `objects.dreamhost.com`).
- In the **Gzip Level** field, optionally type the level of gzip compression you want applied to the log files. You can specify any whole number from 1 (fastest and least compressed) to 9 (slowest and most compressed). This value defaults to 0 (no compression).
- From the **Redundancy Level** menu, select a setting. This value defaults to **Standard**. Amazon's Using Reduced Redundancy Storage Guide (<http://docs.aws.amazon.com/AmazonS3/latest/dev/UsingRRS.html>) provides more information on using reduced redundancy storage. Our discussion of format strings (</guides/streaming-logs/custom-log-formats>) also provides more information.

6. Click the **Create** button to create the new logging endpoint.

7. Click the **Activate** button to deploy your configuration changes.

**NOTE:** Although Fastly continuously streams logs into Amazon S3, the Amazon S3 website and API do not make files available for access until after their upload is complete.

## § Log streaming: Cloud Files (/guides/streaming-logs/log-streaming-cloudfiles)

Fastly's Real-Time Log Streaming (</guides/about-fastly-services/about-fastlys-realtime-log-streaming-features>) feature can send log file to Cloud Files (<https://www.rackspace.com/cloud/files>). Operated by Rackspace, Cloud Files is a file storage service used by developers and IT teams.

**NOTE:** This logging endpoint is disabled by default. To enable this endpoint for your account, contact [support@fastly.com](mailto:support@fastly.com) (<mailto:support@fastly.com>) and request it.

## Prerequisites

If you don't already have a Rackspace Cloud account, you'll need to register (<https://cart.rackspace.com/cloud>) for one. Follow the instructions on Rackspace's website (<https://cart.rackspace.com/cloud>).

## Creating a Cloud Files user and container



Start by creating a Cloud Files user with restricted permissions via Rackspace's cloud control panel (<https://mycloud.rackspace.com/>).

1. Log in to Rackspace's cloud control panel (<https://mycloud.rackspace.com/>).
2. From the user account menu, select **User Management**.
3. Click **Create User** and fill in all appropriate details.
4. In the **Product Access** section, set **User Role** to **Custom**.
5. Review the **Product Access** list. For all items in the **Product** column, set **Role** to **No Access** except the **Files** item.
6. Set the **Files** item **Role** to **Admin**. This will allow you to create the files to store the logs in, but not access any other services.

Next, find the API key for your Cloud Files account. You'll use this later to authenticate using the Cloud Files API.

1. From the user account menu, select **Account Settings**.
2. Show the API key in the **Login details** and make a note of it.

Now that you've created the Cloud Files user and found the API key, you can set up a Cloud Files container.

1. From the **Storage** menu, select **Files**.
2. Click **Create Container**.
3. Assign the container a meaningful name like `Fastly logs - my service`.
4. Choose a region to keep the files in and make sure the container is private.
5. Click **Create Container**.

## Adding a Cloud Files logging endpoint

Once you have created the Cloud Files user and container, follow these instructions to add Cloud Files as a logging endpoint:

1. Review the information in our Setting Up Remote Log Streaming (</guides/streaming-logs/setting-up-remote-log-streaming>) guide.
2. Click the Cloud Files logo. The Create a new Cloud Files endpoint page appears.



## CLOUDFILES ACCOUNT REQUIRED

Before continuing, you will need the **Container Name** and **Access Key** associated with your Rackspace Cloud Files account. If you don't have a Cloud Files account, now is the time to set one up—[Sign up for Cloud Files](#).

## Create a new Cloud Files endpoint

Learn the basics in our Rackspace Cloud Files logging endpoint [documentation](#).

## CONDITION

This will happen all the time unless you [Attach a condition](#)

Description

★ Required

The name of your endpoint, such as **My Endpoint**.

Log format

An Apache-style string or VCL variables to use for log formatting (the Apache Common Log format string appears by default). See [Fastly's log files docs](#), [Varnish's descriptions of VCL variables](#), and [Fastly's available VCL variables](#) for more info.

Bucket name

★ Required

The name of the bucket in which to store the logs.

User

★ Required

The username for your Cloud Files account.

Access key

★ Required

The API key for your Cloud Files account.

Period

★ Required

This manages how frequently in seconds to rotate your log files. Use numbers only in this field.

3. Fill out the **Create a new Cloud Files endpoint** fields as follows:

- In the **Description** field, type a human-readable name for the endpoint.
- In the **Log format** field, optionally type an Apache-style string or VCL variables to use for log formatting. The Apache Common Log format string appears in this field by default. See our guidance on format strings ([/guides/streaming-logs/custom-log-formats](#)) for more information.
- In the **Bucket name** field, type the name of the Cloud Files container in which to store the logs.
- In the **User** field, type the username of the Cloud Files user you created above.
- In the **Access key** field, type the API key of your Cloud Files account.
- In the **Period** field, type an interval (in seconds) to manage how frequently in seconds to rotate your log files. This value defaults to  seconds.

4. Review the **Advanced Options** of the **Create a new Cloud Files endpoint** page and decide which of the optional fields to change, if any.

ADVANCED OPTIONS

Path

The path within the bucket for placing files. It defaults to /, which means files will be placed in its root.

Gzip level

The level of gzip compression, if any, to apply to log files.  
The setting defaults to 0 (no compression). You can specify any whole number from 1 (fastest and least compressed) to 9 (slowest and most compressed).

5. Fill out the **Advanced Options** of the **Create a new Cloud Files endpoint** page as follows

- In the **Path** field, optionally type the path within the container to store the files. The path ends with a trailing slash. If this field is left empty, the files will be saved in the container's root path.
- In the **Gzip level** field, optionally type the level of gzip compression you want applied to the log files. You can specify any whole number from  (fastest and least compressed) to  (slowest and most compressed). This value defaults to  (no compression).

6. Click the **Create** button to create the new logging endpoint.

7. Click the **Activate** button to deploy your configuration changes.

---

## § Log streaming: FTP (/guides/streaming-logs/log-streaming-ftp)

---

Fastly's Real-Time Log Streaming (/guides/about-fastly-services/about-fastlys-realtime-log-streaming-features) feature can send log files to password-protected and anonymous FTP servers.

### Adding FTP as a logging endpoint

After Fastly support has enabled the FTP endpoint for your account, follow these instructions to add FTP as a logging endpoint:

1. Review the information in our Setting Up Remote Log Streaming (/guides/streaming-logs/setting-up-remote-log-streaming) guide.
2. Click the FTP logo. The Create a new File Transfer Protocol (FTP) endpoint page appears.

# Create a new File Transfer Protocol (FTP) endpoint

Learn the basics in our FTP logging endpoint [documentation](#).

## CONDITION

This will happen all the time unless you [Attach a condition](#)

**Description**  ★ Required  
 The name of your endpoint, such as **My Endpoint**.

**Log format**   
 An Apache-style string or VCL variables to use for log formatting (the Apache Common Log format string appears by default). See [Fastly's log files docs](#), [Varnish's descriptions of VCL variables](#), and [Fastly's available VCL variables](#) for more info.

**Address**  :  ★ Required  
 The hostname (or IP address) and port of the FTP server to deliver logs to. For example, **logging.example.com**.

**Path**  ★ Required  
 The path to the directory where logs are to be delivered. For example, **/fastly-logs**.

**User**  ★ Required  
 The username for the server. For anonymous access, use the username **anonymous**.

**Password**  ★ Required  
 The password for the server. For anonymous access, use the **email address** as the password.

**Period**  ★ Required  
 This manages how frequently in seconds to rotate your log files.  
 Use numbers only in this field.

3. Fill out the **Create a new File Transfer Protocol (FTP) endpoint** fields as follows:

- In the **Description** field, type a human-readable name for the endpoint.
- In the **Log format** field, optionally type an Apache-style string or VCL variables to use for log formatting. The Apache Common Log format string appears in this field by default.
- In the **Address** field, type the hostname or IP address of the FTP server. In the port field, type the port number you're using for FTP (the default is ).
- In the **Path** field, optionally type the path to store the files. The path ends with a trailing slash. If this field is left empty, the files will be saved in the root path.
- In the **User** field, type the username used to authenticate to the FTP server. For anonymous access, use the username .
- In the **Password** field, type the password used to authenticate to the FTP server. For anonymous access, use an email address as the password.
- In the **Period** field, type an interval (in seconds) to control how frequently your log files are rotated. This value defaults to  seconds

4. In the **Gzip level** field of the **Advanced Options**, optionally type the level of gzip compression you want applied to the log files. You can specify any whole number from  (fastest and least compressed) to  (slowest and most compressed). This value defaults to  (no compression).

5. Click the **Create** button to create the new logging endpoint.

6. Click the **Activate** button to deploy your configuration changes.

---

## § Log streaming: Google Cloud Storage (/guides/streaming-logs/log-streaming-google-cloud-storage)

---

Fastly's Real-Time Log Streaming feature (/guides/about-fastly-services/about-fastlys-realtime-log-streaming-features) can send log files to Google Cloud Storage (<https://cloud.google.com/storage/>) (GCS). GCS is an online file storage service used for storing and accessing data on Google's infrastructure. One advantage of using GCS is that you can use Google BigQuery (<https://cloud.google.com/bigquery/>) to analyze the log files.

### Prerequisites

Before adding GCS as a logging endpoint for Fastly services, you will need to register for a GCS account, create a bucket and service account on Google's website, and obtain the `private_key` and `client_email` from the JSON file associated with the service account.

### Creating a GCS bucket

You can create a new GCS bucket to hold the logs, or you can use an existing bucket. Be sure to note the name of the bucket as you will need it later. To learn how to create a GCS bucket, see Google's guide on creating a bucket ([https://cloud.google.com/storage/docs/getting-started-console#create\\_a\\_bucket](https://cloud.google.com/storage/docs/getting-started-console#create_a_bucket)).

### Creating a service account

GCS uses service accounts for third-party application authentication. You will need to create a new service account on Google's website. To learn how to create a service account, see Google's guide on generating a service account credential (<https://cloud.google.com/storage/docs/authentication#generating-a-private-key>). When you create the service account, be sure to set the **Key Type** to `JSON`.

### Obtaining the private key and client email

After you create the service account, a JSON file will be downloaded to your computer. This file contains the credentials for the GCS service account you just created. Open the file with a text editor and make a note of the `private_key` and `client_email`.

## Adding GCS as a logging endpoint

After Fastly support has enabled the GCS endpoint for your account, follow these instructions to add GCS as a logging endpoint:

1. Review the information in our Setting Up Remote Log Streaming (/guides/streaming-logs/setting-up-remote-log-streaming) guide.
2. Click the Google Cloud Services logo. The Create a new Google Cloud Storage (GCS) endpoint page appears.



## GCS ACCOUNT REQUIRED

Before continuing, you will need the **private\_key** and **client\_email** associated with a GCS service account. If you don't have an account, now is the time to set one up—[Sign up for Google Cloud Storage](#).

## Create a new Google Cloud Storage (GCS) endpoint

Learn the basics in our Google Cloud Storage logging endpoint [documentation](#).

## CONDITION

This will happen all the time unless you [Attach a condition](#)

Description

★ Required

The name of your endpoint, such as **My Endpoint**.

Log format

An Apache-style string or VCL variables to use for log formatting (the Apache Common Log format string appears by default). See [Fastly's log files docs](#), [Varnish's descriptions of VCL variables](#), and [Fastly's available VCL variables](#) for more info.

Email

★ Required

The [Email address](#) associated with the target GCS bucket on your account.

Bucket name

★ Required

The name of the bucket in which to store the logs.

Secret key

★ Required

The [Secret Key](#) associated with the target GCS bucket on your account.

Period

★ Required

This manages how frequently in seconds to rotate your log files.

Use numbers only in this field.

### 3. Fill out the **Create a new Google Cloud Storage (GCS) endpoint** fields as follows:

- In the **Description** field, type a human-readable name for the endpoint.
- In the **Log Format** field, optionally type an Apache-style string or VCL variables to use for log formatting. The Apache Common Log format string appears in this field by default. See our guidance on format strings ([/guides/streaming-logs/custom-log-formats](#)) for more information.
- In the **Email** field, type the `client_email` address listed in the JSON file associated with the service account you created on Google's website.
- In the **Bucket Name** field, type the name of the GCS bucket in which to store the logs.
- In the **Secret Key** field, type the `private_key` value listed in the JSON file associated with the service account you created on Google's website. We strip out the JSON newline escape characters for you so don't worry about removing them.
- In the **Period** field, optionally type an interval (in seconds) to control how frequently your log files are rotated. This value defaults to `3600` seconds.

### 4. Review the **Advanced Options** of the **Create a new Google Cloud Storage (GCS) endpoint** page and decide which of the optional fields to change, if any.

ADVANCED OPTIONS

Path

The path within the bucket for placing files. It defaults to /, which means files will be placed in its root.

Gzip level

The level of gzip compression, if any, to apply to log files.  
The setting defaults to 0 (no compression). You can specify any whole number from 1 (fastest and least compressed) to 9 (slowest and most compressed).

5. Fill out the **Advanced Options** of the **Create a new Google Cloud Storage (GCS) endpoint** page as follows:

- In the **Path** field, optionally type the path within the bucket to store the files. Specify a directory by ending the path with a trailing slash (/). Leaving this field empty saves the files in the bucket's root path.
- In the **Gzip Level** field, optionally type the level of gzip compression you want applied to the log files. You can specify any whole number from 1 (fastest and least compressed) to 9 (slowest and most compressed). This value defaults to 0 (no compression).

6. Click **Create** to create the new logging endpoint.

7. Click the **Activate** button to deploy your configuration changes.

## § Log streaming: Log Shuttle (/guides/streaming-logs/log-streaming-log-shuttle)

Fastly's Real-Time Log Streaming (/guides/about-fastly-services/about-fastlys-realtime-log-streaming-features) feature can send log files to Log Shuttle (<https://github.com/heroku/log-shuttle>). Log Shuttle is an open source application designed to provide simpler encrypted and authenticated log delivery.

**NOTE:** This logging endpoint is disabled by default. To enable this endpoint for your account, contact [support@fastly.com](mailto:support@fastly.com) (<mailto:support@fastly.com>) and request it.

## Adding Log Shuttle as a logging endpoint

After Fastly support has enabled the Log Shuttle endpoint for your account, follow these instructions to add Log Shuttle as a logging endpoint:

1. Review the information in our Setting Up Remote Log Streaming (/guides/streaming-logs/setting-up-remote-log-streaming) guide.
2. Click the Log Shuttle logo. The Create a new Log Shuttle endpoint page appears.

## Create a new Log Shuttle endpoint

Learn the basics in our Log Shuttle logging endpoint [documentation](#).

**CONDITION** This will happen all the time unless you [Attach a condition](#)

**Description**  ★ Required

The name of your endpoint, such as **My Endpoint**.

**Log format**

An Apache-style string or VCL variables to use for log formatting (the Apache Common Log format string appears by default). See [Fastly's log files docs](#), [Varnish's descriptions of VCL variables](#), and [Fastly's available VCL variables](#) for more info.

**Token**  ★ Required

The authentication token to send in front of each log line.

**URL**  ★ Required

The URL to send log data to. For example, <https://logs.example.com>.

**CREATE**

**CANCEL**

3. Fill out the **Create a new Log Shuttle endpoint** fields as follows:

- In the **Description** field, type a human-readable name for the endpoint.
- In the **Log format** field, optionally type an Apache-style string or VCL variables to use for log formatting. The Apache Common Log format string appears in this field by default.
- In the **Token** field, type the data authentication token. This is required for some endpoints like Heroku's Log Integration.
- In the **URL** field, type the URL to which log data will be sent (e.g., <https://logs.example.com/>).

4. Click the **Create** button to create the new logging endpoint.

5. Click the **Activate** button to deploy your configuration changes.

## § Log streaming: Logentries (/guides/streaming-logs/log-streaming-logentries)

Fastly's Real-Time Log Streaming (/guides/about-fastly-services/about-fastlys-realtime-log-streaming-features) feature can send log files to Logentries (<https://logentries.com/>). Logentries is a real-time log management and analytics system that you can use to monitor your Fastly logs.

### Prerequisites

1. Register for a Logentries account.
2. Create a new log in the Logentries application by following the instructions on the Logentries website (<https://logentries.com/doc/fastly/>).
3. During new log creation, select **Manual Configuration** and **Token TCP**.
4. Make a note of the token provided in the Logentries configuration panel. We recommend you use this token when you create the Logentries logging endpoint for Fastly services.

## Adding Logentries as a logging endpoint

After you've created a new log in Logentries and found the token, follow these instructions to add Logentries as a logging endpoint for Fastly services:



1. Review the information in our Setting Up Remote Log Streaming (/guides/streaming-logs/setting-up-remote-log-streaming) guide.
2. Click the Logentries logo. The Create a new Logentries endpoint page appears.

## Create a new Logentries endpoint

Learn the basics in our Logentries logging endpoint [documentation](#).

**CONDITION** This will happen all the time unless you [Attach a condition](#)

**Description**  ★ Required  
 The name of your endpoint, such as **My Endpoint**.

**Log format**   
 An Apache-style string or VCL variables to use for log formatting (the Apache Common Log format string appears by default). See [Fastly's log files docs](#), [Varnish's descriptions of VCL variables](#), and [Fastly's available VCL variables](#) for more info.

**Port**  ★ Required  
 You will need to change your port if you use TLS or a Token.

**TLS**  ⌵  
 Use Transport Layer Security (TLS) for secure logging.

**Token**   
 The [authentication token](#) to send in front of each log line. It's shown when you create or edit a log set.

**CREATE** **CANCEL**

3. Fill out the **Create a new Logentries endpoint** fields as follows:

- In the **Description** field, type a human-readable name for the endpoint.
- In the **Log format** field, optionally type an Apache-style string or VCL variables to use for log formatting. The Apache Common Log format string appears in this field by default. Our discussion of format strings (/guides/streaming-logs/custom-log-formats) also provides more information.
- In the **Port** field, type `20000`. Though we recommend this specific setting when adding your endpoint, there are additional options to consider when deciding on the port, TLS, and token settings.
- From the **TLS** menu, select **Yes**.
- In the **Token** field, type the token provided in the Logentries configuration panel.

4. Click the **Create** button to create the new logging endpoint.

5. Click the **Activate** button to deploy your configuration changes.

## Additional selections for ports, TLS, and tokens

The Port, TLS, and Token selections shown above are what we recommend when you add a Logentries endpoint. However, you do have other choices when you add your endpoint.

**Using your port number.** You can add a Logentries endpoint by using your secret Logentries port number (e.g., `56789`). However, we strongly recommend sending your logs via TLS. To do this, set TLS to `Yes` and add `10000` to your secret port number (e.g., `66789`). See the Logentries guide Plain TCP/UDP (<https://docs.logentries.com/docs/input-plaintextudp>) for more information.

**Using your token.** You can add a Logentries endpoint by using your secret account token. To use your token, set your port to `10000`. However, we strongly recommend sending your logs via TLS. To do this, set TLS to `Yes` and the port number to `20000`. See the Logentries guide Token TCP (<https://docs.logentries.com/docs/input-token>) for more information.

## Next steps

Logentries maintains the Fastly Community Pack (<https://blog.logentries.com/2015/01/fastly-community-pack/>) that leverages custom VCL to provide advanced User-Agent statistics, regional statistics, error tracking, and more.

**❗ IMPORTANT:** The ability to [upload custom VCL code \(/guides/vcl/uploading-custom-vcl\)](/guides/vcl/uploading-custom-vcl) is disabled by default when you sign up. To enable this ability for your account, contact [support@fastly.com \(mailto:support@fastly.com\)](mailto:support@fastly.com) and request it.

## § Log streaming: Loggly (/guides/streaming-logs/log-streaming-loggly)

Fastly's Real-Time Log Streaming (/guides/about-fastly-services/about-fastlys-realtime-log-streaming-features) feature can send log files to Loggly (<https://www.loggly.com/>). Loggly is an agent-less log collection and management tool.

**❗ IMPORTANT:** This feature is part of a limited availability release. For more information, see our [product and feature life cycle \(/guides/customer-support/product-and-feature-life-cycles#limited-availability\)](/guides/customer-support/product-and-feature-life-cycles#limited-availability) descriptions.

### Prerequisites

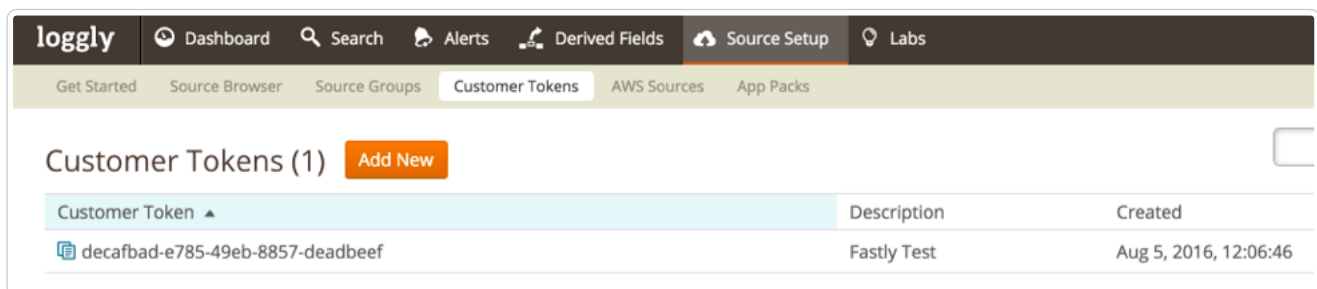
If you don't already have a Loggly account, you'll need to register for one. Follow the signup instructions (<https://www.loggly.com/signup/>) on the Loggly website.

You should also have your Fastly API key (/guides/account-management-and-security/finding-and-managing-your-account-info.html#finding-and-displaying-your-api-key) and the service ID (/guides/account-management-and-security/finding-and-managing-your-account-info#finding-your-service-id) and version number of the Fastly service for which you'll be enabling Loggly logging.

### Adding Loggly as a logging endpoint

To use Loggly as a logging endpoint follow the steps below:

1. Navigate to the **Customer Tokens** (<https://www.loggly.com/docs/customer-token-authentication-token/>) area in the **Source Setup** on your Loggly dashboard.



2. Make note of your Loggly customer token. Loggly uses this to associate data you send them with your account.
3. Run the following cURL command:

```
curl -s -X POST -H 'Fastly-Key: <your Fastly API key>' -H 'Content-Type: application/json' 'https://api.fastly.com/service/<your Fastly service ID>/version/<version number>/logging/loggly' --data-binary '{"name": "<log name>", "format": "%h %l %u %t %r %>s", "token": "<your Loggly token>"}'
```

## § Log streaming: Heroku's Logplex (/guides/streaming-logs/log-streaming-logplex)

As part of our Real-Time Log Streaming (/guides/about-fastly-services/about-fastlys-realtime-log-streaming-features) feature, if you use our Heroku add-on (<https://elements.heroku.com/addons/fastly>), you can send log files directly to Heroku's Logplex (<https://devcenter.heroku.com/articles/logplex>) system. Logplex is Heroku's distributed syslog router that collates and distributes log entries from a variety of sources into a single channel.

**❗ IMPORTANT:** This feature is part of a limited availability release. For more information, see our [product and feature life cycle \(/guides/customer-support/product-and-feature-life-cycles#limited-availability\)](/guides/customer-support/product-and-feature-life-cycles#limited-availability) descriptions.

To enable this feature for your account, contact [support@fastly.com \(mailto:support@fastly.com\)](mailto:support@fastly.com) and request it.

Once enabled, your Fastly logs will be available in exactly the same way (<https://devcenter.heroku.com/articles/logging>) as your regular app and hosted service logs. You can view them using the Heroku command line log viewer or send them to a logging add-on (<https://elements.heroku.com/addons/#logging>).

## § Log streaming: OpenStack (/guides/streaming-logs/log-streaming-openstack)

Fastly's Real-Time Log Streaming (/guides/about-fastly-services/about-fastlys-realtime-log-streaming-features) feature can send log files to OpenStack (<http://www.openstack.org/>). OpenStack is an open-source platform for cloud-computing that many companies deploy as an infrastructure-as-a-service.

**NOTE:** This logging endpoint is disabled by default. To enable this endpoint for your account, contact [support@fastly.com](mailto:support@fastly.com) (<mailto:support@fastly.com>) and request it.

### Adding OpenStack as a logging endpoint

After Fastly support has enabled the OpenStack endpoint for your account, follow these instructions to add OpenStack as a logging endpoint:

1. Review the information in our Setting Up Remote Log Streaming (/guides/streaming-logs/setting-up-remote-log-streaming) guide.
2. Click the OpenStack logo. The Create a new OpenStack endpoint page appears.

### Create a new OpenStack endpoint

Learn the basics in our OpenStack logging endpoint [documentation](#).

**CONDITION** This will happen all the time unless you [Attach a condition](#)

Description

★ Required

The name of your endpoint, such as **My Endpoint**.

Log format

An Apache-style string or VCL variables to use for log formatting (the Apache Common Log format string appears by default). See [Fastly's log files docs](#), [Varnish's descriptions of VCL variables](#), and [Fastly's available VCL variables](#) for more info.

Auth URL

★ Required

The URL used for OpenStack authentication. For example, <https://auth.api.rackspacecloud.com/v1.0>, <https://auth.storage.memset.com/v1.0>.

Bucket name

★ Required

The name of the bucket in which to store the logs.

User

★ Required

Your OpenStack username.

Access key

★ Required

Your OpenStack access key.

Period

★ Required

Use numbers only in this field. This manages how frequently in seconds to rotate your log files.

3. Fill out the **Create a new OpenStack endpoint** fields as follows:

- In the **Description** field, type a human-readable name for the endpoint.
  - In the **Log format** field, optionally type an Apache-style string or VCL variables to use for log formatting. The Apache Common Log format string appears in this field by default. See our guidance on format strings (</guides/streaming-logs/custom-log-formats>) for more information.
  - In the **Auth URL** field, type the URL used for OpenStack authentication (e.g., `https://auth.api.rackspacecloud.com/v1.0`).
  - In the **Bucket Name** field, type the name of the OpenStack bucket in which to store the logs.
  - In the **Username** field, type your OpenStack username.
  - In the **Access Key** field, type your OpenStack access key.
  - In the **Period** field, type an interval (in seconds) to control how frequently to rotate your log files. This value defaults to `3600` seconds.
4. Review the **Advanced Options** of the **Create a new OpenStack endpoint** page and decide which of the optional fields to change, if any.

ADVANCED OPTIONS

Path

The path within the bucket for placing files. It defaults to `/`, which means files will be placed in its root.

Gzip level

The level of gzip compression, if any, to apply to log files.

The setting defaults to `0` (no compression). You can specify any whole number from `1` (fastest and least compressed) to `9` (slowest and most compressed).

5. Fill out the **Advanced Options** of the **Create a new OpenStack endpoint** page as follows:
- In the **Path** field, optionally type the path within the bucket to store the files. The path ends with a trailing slash. If this field is left empty, the files will be saved in the bucket's root path.
  - In the **Gzip Level** field, optionally type the level of gzip compression you want applied to the log files. You can specify any whole number from `1` (fastest and least compressed) to `9` (slowest and most compressed). This value defaults to `0` (no compression).
6. Click the **Create** button to create the new logging endpoint.
7. Click the **Activate** button to deploy your configuration changes.

## § Log streaming: Papertrail (</guides/streaming-logs/log-streaming-papertrail>)

Fastly's Real-Time Log Streaming (</guides/about-fastly-services/about-fastlys-realtime-log-streaming-features>) feature can send log files to Papertrail (<https://papertrailapp.com>). Papertrail is a web-based log aggregation application used by developers and IT teams. Instructions for setting up remote log streaming via Papertrail are detailed in the Papertrail setup and configuration documentation (<http://help.papertrailapp.com/kb/hosting-services/fastly/>).

## § Log streaming: Scalyr (</guides/streaming-logs/log-streaming-scalyr>)

Fastly's Real-Time Log Streaming (</guides/about-fastly-services/about-fastlys-realtime-log-streaming-features>) feature can send log files to Scalyr (<https://www.scalyr.com/>). Scalyr pulls all your server logs and metrics into a centralized, searchable system in real-time.

### Prerequisites

If you don't already have a Scalyr account, you'll need to register for one. Follow the signup instructions (<https://www.scalyr.com/signup>) on the Scalyr website.

Once you've signed up, navigate to the **API Keys** area in the **Settings** on your Scalyr dashboard and make note of your Scalyr Write Token. Scalyr uses this to associate data you send them with your account. You'll need this token when you set up your endpoint with Fastly.

If you're adding the Scalyr endpoint via the command line, instead of the web interface, you should also have your Fastly API key (</guides/account-management-and-security/finding-and-managing-your-account-info.html#finding-and-displaying-your-api-key>) and the service ID (</guides/account-management-and-security/finding-and-managing-your-account-info.html#finding-your-service-id>) and version number of the Fastly service for which you'll be enabling Scalyr logging.

## Adding Scalyr as a logging endpoint via the web interface

Follow these instructions to add Scalyr as a logging endpoint:

1. Review the information in our [Setting Up Remote Log Streaming](/guides/streaming-logs/setting-up-remote-log-streaming) guide.
2. Click the Scalyr logo. The Create a new Scalyr endpoint page appears.

### Create a new Scalyr endpoint

Learn the basics in our Scalyr logging endpoint [documentation](#).

**CONDITION** This will happen all the time unless you [Attach a condition](#)

**Description**  ★ Required  
 The name of your endpoint, such as **My Endpoint**.

**Log format**   
 An Apache-style string or VCL variables to use for log formatting (the Apache Common Log format string appears by default). See [Fastly's log files docs](#), [Varnish's descriptions of VCL variables](#), and [Fastly's available VCL variables](#) for more info.

**Token**  ★ Required  
 The [authentication token](#) to send in front of each log line.

**CREATE** **CANCEL**

3. Fill out the **Create a new Scalyr endpoint** fields as follows:
  - In the **Description** field, type a human-readable name for the endpoint.
  - In the **Log format** field, optionally type an Apache-style string or VCL variables to use for log formatting. The Apache Common Log format string appears in this field by default. See our guidance on format strings (</guides/streaming-logs/custom-log-formats>) for more information.
  - In the **Token** field, type the Scalyr Write Token provided in the Scalyr dashboard.
4. Click the **Create** button to create the new logging endpoint.
5. Click the **Activate** button to deploy your configuration changes.

## Adding Scalyr as a logging endpoint via command line

Run the following cURL command:

```

```plaintext
curl -s -X POST -H 'Fastly-Key: <your Fastly API key>' -H 'Content-Type: application/json' 'https://api.fastly.com/service/<your Fastly service ID>/version/<version number>/logging/scalyr' --data-binary '{"name": "<log name>", "format": "%h %l %u %t %r %>s", "token": "<your Scalyr token>"}'
```

```

Be sure to replace `<your Fastly API key>`, `<your Fastly service ID>`, `<version number>`, `<log name>`, and `<your Scalyr token>` the appropriate details specific to your account.

## § Log streaming: SFTP (/guides/streaming-logs/log-streaming-sftp)

Fastly's Real-Time Log Streaming (/guides/about-fastly-services/about-fastlys-realtime-log-streaming-features) feature can send log files to SFTP, a secure file transfer subsystem for the Secure Shell (SSH) protocol. Our SFTP endpoint supports both password-based authentication and SSH public-key authentication, with SSH public-key authentication being preferred. To learn more about SSH public-key authentication, or to learn how to generate public and private key pairs, see this guide (<https://git-scm.com/book/en/v2/Git-on-the-Server-Generating-Your-SSH-Public-Key>).

❗ **IMPORTANT:** This feature is part of a limited availability release. For more information, see our [product and feature life cycle \(/guides/customer-support/product-and-feature-life-cycles#limited-availability\)](/guides/customer-support/product-and-feature-life-cycles#limited-availability) descriptions.

## Adding SFTP as a logging endpoint

To create an SFTP endpoint, run the following cURL command:

```
curl -s -X POST -H 'Fastly-Key: <your Fastly API key>' -H 'Content-Type: application/json' 'https://api.fastly.com/service/<your Fastly service ID>/version/<version number>/logging/sftp' --data-binary '{"name": "<log name>", "address": "<your hostname>", "port": "<port number>", "format": "%h %l %u %t %r %s", "user": "<your username>", "secret_key": "<your secret key>"}'
```

Secret keys must be passed in "PEM" format. A PEM file is a base64 (ASCII) encoded block of data encapsulated between two delimiters, e.g. something that looks like:

```
-----BEGIN RSA PRIVATE KEY-----
....
-----END RSA PRIVATE KEY-----
```

To pass your secret key on the command line, you'll need to convert all the newlines to escaped `\n` characters, as shown in the example below.

```
curl -s -X POST -H 'Fastly-Key: <your Fastly API key>' -H 'Content-Type: application/json' 'https://api.fastly.com/service/<your Fastly service ID>/version/<version number>/logging/sftp' --data-binary '{"name": "<log name>", "address": "<your hostname>", "port": "22", "format": "%h %l %u %t %r %s", "user": "<your username>", "secret_key": "-----BEGIN PRIVATE KEY-----\nMIIEwAIBADANBgkqhkiG9w0BAQEF\nAASCBKowggSmAgEAAoIBAQCWAY2o6Rxh2NdT\nu509AvEek9mWV7LVje/tFnlMzf9t8nRhiLrL83N+u0NnZpnpPsI4lRT88zi9fAj4\nniMeSFJcF7dUdN259ZiubF/aWfDSZ\nUaLOd8Jw1z+L5Vgpc2sERN/N+UA/Sn6NEk2M\nn6DliRuTYhwxaJjpWIG206Noa66tyJAbbsuMRDA0B9b6avd2xQ4UkUWjFwNtwX5b3\nnWhWoxTL/eVs7JPX8Ju0jfsGa13x8\nUcnMhojwzY7wDU3WURO/kp5LgeVY/r1p0fVZ\nnRaDEcakA1ZMMqCr3DcEzFsmi5h/7ncn+r1lRmIxExBHawZ3MPVhWAEQg0Ceo5TIY\nn8xK9CEBrAgMBAAECggEBAIAA1lMt\nU4JQYz2bimwJjs2Y5dWBww7wjRsd4tLcuPmz\nnp16uj0OrQj33mpf0mJbFhkzGQ9XxvkX8+uUERLX3yqA3oUiE67kTLpGpa9WLOgm2\nnGiyYxTnEDzmmWczTKS4wbQE+b5UZ\nH9NfomerTlyASlKsG+lZHLFIYct+nUyYgDFa\nn6ARFUcLgVBwOjGlnvYi2emLunomsLIkvPFdMdJzFw8hP4YyZ7GHoKpXb0CHGINTS\nng5IAFnUNvNlIjjz7qk+7W6w6VPt5\nYb9rcFA6Q+90hr3tEaK9erEi0w+xJqC8M+O1\nnzvwxOnMUI0thLHK4HG7f1KWMEUDjjBLj8xLRq5EDfEzhRoppuT30ToWw5NAb65S\nngWS6EQswlEpAN4ulBZitilS+GWj\nN/cYkukzo2NrEp+ECgYEAzRKU/h6tLH5DsLWM\nnNO/Fldbj3gimW60sPe08S1FsAu7c4DuHRhpYvm+9v0xPKm5IZPCa7P1sJWFY4WIq\nnFJNq+zsDJSNkTbNHTcLfNR5A/OY\nZTrDqr2OUMVJfzB67Is89BUIqYKh32NV4c36E\nnzcf94JMaKHLxV1V87YUeX1V63ECgYEAu0Ii/ypOwL7TSuOh65VdmQH2d9xtPmFS\nnouAx10d6FBsqreElretqQEYtpAkG\ni6RX8JYJSn7gfsczm087F+GBHY1Hx80mpOMGA\nnjgWSijRWY5sCgYEAstTUMvvyqUBrmXAZRbHOQGtTOLneovLNCEKRDh4Z12eT0JId\nnAtVumsxUnA9E6RC1cRmsSEW0gcb\n3+Jb0jKbK2tMLphIpRgSvwlD213clyuczsj\nnz2LhiBYnZrztDrq3SEcu55RTT9dhWv5F7bZ9i\nJuax08vVCexoSRwa/kJBECgYEA/nmS0zhXis2T2iJytnYlcuViSF5ew\n4rInD3n/37u8erF3Anwf9dGre3wsj3cr2S6ZB\nnOMmPHrmlq3KJjrjItJD1JoOic1ikFuw7W878qWcd\nDN1J3z4/Er2XrnWiXefxfay\nnmvUF7zoQYsNjzEhHvwDCBwodVve\nPENo+AiJOE97QujScXw6cakN0T8xcYIGddzP\nnQypNTw8MSGU/ZeKyji9h5jZ62bc33FZTSBAlEwQULCsCgYEA\nmKYqSAUWTKn2QrPq\nnj9RrhfjSeekLok2u4Vuy1fH/vqc\nn1WLthUodL9PtGD6hUOqaXP6INl3w2yIKjAla\nnuG+3m9U7MRqkH1ecobtvGN/q1r8=\n\n-----END PRIVATE KEY-----"}'
```

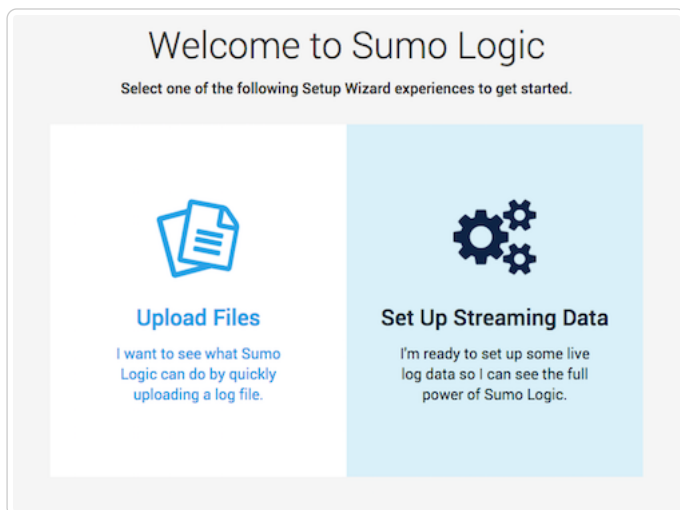
## § Log streaming: Sumo Logic (/guides/streaming-logs/log-streaming-sumologic)

Fastly's Real-Time Log Streaming (/guides/about-fastly-services/about-fastlys-realtime-log-streaming-features) feature can send log files to Sumo Logic (<https://www.sumologic.com/>). Sumo Logic is a web-based log analytics platform used by developers and IT teams.

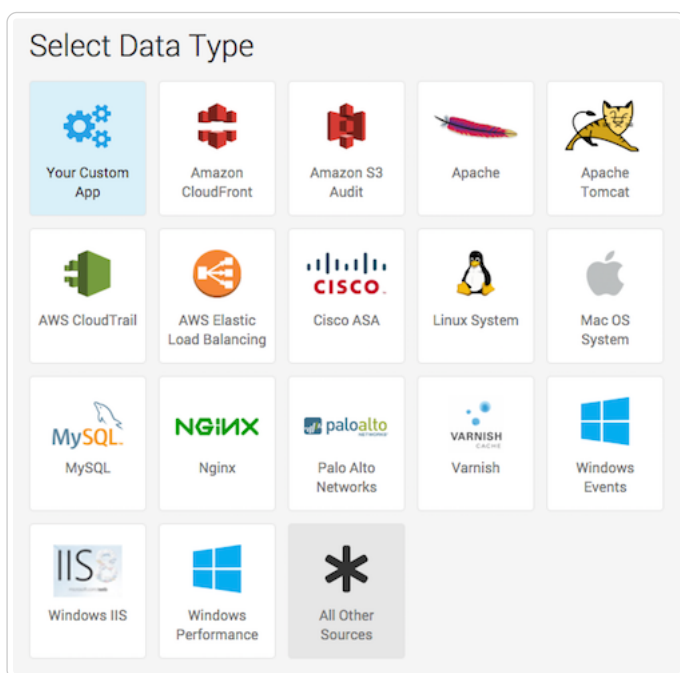
## Setting up Sumo Logic

To use Sumo Logic as a logging endpoint, you'll need to create a Sumo Logic account, add a new source, and save the HTTP Source URL. Follow these instructions to add a new source in the Sumo Logic website:

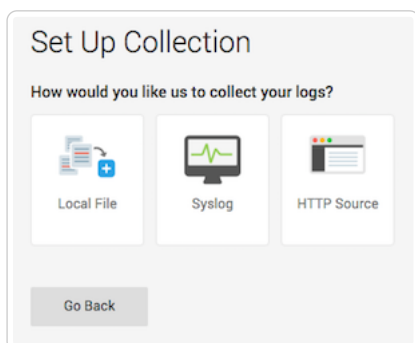
1. The process starts with the Sumo Logic Setup Wizard, which appears immediately after you create your Sumo Logic account. If you already have an account, you can access the wizard by selecting **Setup Wizard** from the **Manage** menu at the top of the Sumo Logic application.



2. Click **Set Up Streaming Data**. The Select Data Type window appears.



3. Click **All Other Sources**. The Set Up Collection window appears.



4. Click **HTTP Source**. The Configure Source: HTTP Source window appears.

Configure Source: HTTP Source

1. Enter a Source Category that will help you search your logs later.

Source Category ⓘ

custom\_http\_logs

2. Select a time zone for your log file.

☒ Use time zone from log file. If none present use:

(UTC) Etc/UTC

☐ Ignore time zone from log file and instead use:

(UTC) Etc/UTC

Go Back Continue

5. In the **Source Category** field, type a human-readable name for the category (e.g., `fastly_cdn`) and select a time zone for your log file.
6. Click **Continue**. The HTTP Source URL appears.

Configure Source: HTTP Source

1. An HTTP Source has been automatically configured for you. Copy the following URL.

https://endpoint1.collection.us2.sumologic.com/receive/r/v1/http/ZaVnC4dhaV1\_3H4t1DeNhc0Sq-AFYok\_X\_kEx9wEYsmV0Xb7X1Evail0RU0r1XPr5tuc13w4FbjEoW5ZZSpKTWL\_vWnhzChI3Z1v4bY8XW8qd1ofVD74Kw== Select All

2. Use the URL as the target for your Source. [Learn more](#)  
> [View cURL example](#)

Go Back Continue

7. Copy the HTTP Source URL. You will enter this value in the Fastly web interface.
8. Click **Continue**. Sumo Logic will add the new source.

## Adding Sumo Logic as a logging endpoint

After you've created a Sumo Logic account and obtained the HTTP Source URL, follow these instructions to add Sumo Logic as a logging endpoint for Fastly services:

1. Review the information in our Setting Up Remote Log Streaming (/guides/streaming-logs/setting-up-remote-log-streaming) guide.
2. Click the Sumo Logic logo. The Create a new Sumo Logic endpoint page appears.





## SUMO LOGIC ACCOUNT REQUIRED

Before continuing, you will need the **HTTP Source URL** from your Sumo Logic account. If you don't have a Sumo Logic account, now is the time to set one up—[Sign up for Sumo Logic](#).

## Create a new Sumo Logic endpoint

Learn the basics in our Sumo Logic logging endpoint [documentation](#).

## CONDITION

This will happen all the time unless you [Attach a condition](#)

Description

★ Required

The name of your endpoint, such as **My Endpoint**.

Log format

An Apache-style string or VCL variables to use for log formatting (the Apache Common Log format string appears by default). See [Fastly's log files docs](#), [Varnish's descriptions of VCL variables](#), and [Fastly's available VCL variables](#) for more info.

Collector URL

★ Required

The address of the HTTP Source Collector in Sumo Logic.

CREATE

CANCEL

3. Fill out the **Create a new Sumo Logic endpoint** fields as follows:

- In the **Description** field, type a human-readable name for the endpoint.
- In the **Log format** field, optionally type an Apache-style string or VCL variables to use for log formatting. The Apache Common Log format string appears in this field by default. Our discussion of format strings ([/guides/streaming-logs/custom-log-formats](#)) provides more information.
- In the **Collector URL** field, type the address of the HTTP Source URL you found in the Sumo Logic website.

4. Click the **Create** button to create the new logging endpoint.

5. Click the **Activate** button to deploy your configuration changes.

## Troubleshooting

The Sumo Logic logging endpoint is designed for services with sustained levels of traffic. If you aren't seeing any logs in Sumo Logic, try waiting a bit.

## § Log streaming: Syslog ([/guides/streaming-logs/log-streaming-syslog](#))

Fastly's Real-Time Log Streaming ([/guides/about-fastly-services/about-fastlys-realtime-log-streaming-features](#)) feature can send log files to syslog-based logging software. Syslog is a widely used standard for message logging.

**NOTE:** [Splunk Storm \(http://www.splunk.com/\)](http://www.splunk.com/) was end-of-lifed on April 1st, 2015. If you are using a Splunk product, you should configure it to [receive log streams \(/guides/streaming-logs/setting-up-remote-log-streaming\)](#) over syslog.

## Adding syslog as a logging endpoint

Follow these instructions to add syslog as a logging endpoint:

1. Review the information in our Setting Up Remote Log Streaming (/guides/streaming-logs/setting-up-remote-log-streaming) guide.
2. Click the syslog icon. The Create a new Syslog endpoint page appears.

## Create a new Syslog endpoint

Learn the basics in our Syslog logging endpoint [documentation](#).

**CONDITION** This will happen all the time unless you [Attach a condition](#)

**Description**

★ Required

The name of your endpoint, such as **My Endpoint**.

**Log format**

An Apache-style string or VCL variables to use for log formatting (the Apache Common Log format string appears by default). See [Fastly's log files docs](#), [Varnish's descriptions of VCL variables](#), and [Fastly's available VCL variables](#) for more info.

**Syslog address**

 : 
★ Required

The domain name (or IP address) and port to which log should be sent. Logs are streamed using TCP so be sure this port can receive incoming TCP traffic.

**Token**

Optional authentication token to send in front of each log line.

**TLS**

Use Transport Layer Security (TLS) for secure logging.

CREATE

CANCEL

3. Fill out the **Create a new Syslog endpoint** fields as follows:

- In the **Description** field, type a human-readable name for the endpoint.
- In the **Log format** field, optionally type an Apache-style string or VCL variables to use for log formatting. The Apache Common Log format string appears in this field by default. See our guidance on format strings (/guides/streaming-logs/custom-log-formats) for more information.
- In the **Syslog address** field, type the domain name or IP address and port to which logs should be sent. Be sure this port can receive incoming TCP traffic from Fastly. See the firewall considerations section for more information.
- In the **Token** field, optionally type a string prefix (line prefix) to send in front of each log line.
- From the **TLS** menu, select **No** to disable encryption for the syslog endpoint, or **Yes** to enable it. When you select Yes, the TLS Hostname and TLS CA Certificate fields both appear.
- In the **TLS Hostname** field, optionally type the hostname used to verify the syslog server's certificate. This can be either the Common Name (CN) or Subject Alternate Name (SAN). This field only appears when you select Yes from the Use TLS menu.
- In the **TLS CA Certificate** field, optionally copy and paste the Certificate Authority (CA) certificate used to verify that the origin server's certificate is valid. The certificate you upload must be in PEM format. Consider uploading the certificate if it's not signed by a well-known certificate authority. This value is not required if your TLS certificate is signed by a well-known authority. This field only appears when you select Yes from the Use TLS menu.

4. Click the **Create** button to create the new logging endpoint.
5. Click the **Activate** button to deploy your configuration changes.

## Adding separators or static strings

To insert a separator or other arbitrary string into the syslog endpoint format:

1. Create a new header (/guides/basic-configuration/adding-or-modifying-headers-on-http-requests-and-responses) with the following fields:
  - From the **Type** menu, select **Request**, and from the **Action** menu, select **Set**.
  - In the **Destination** field, type any suitable header name (for example, `http.X-Separator`).
  - In the **Source** field, type any special character or string you want (for example, `"|"`).
2. Reference the new header variable in the log format box for your specific provider (for example, `req.http.X-Separator`).

## Syslog facility and severity

The syslog output includes the following facility and severity values:

```
facility: local0
severity: info
```

## Firewall considerations

Syslog has limited security features. For this reason, it's best to create a firewall for your syslog server and only accept TCP traffic on your configured port from our address blocks. Our list of address blocks is dynamic, so we recommend programmatically obtaining the list from our JSON feed (<https://api.fastly.com/public-ip-list>) whenever possible.

## § Setting up remote log streaming (/guides/streaming-logs/setting-up-remote-log-streaming)

Fastly's Real-Time Log Streaming feature (/guides/about-fastly-services/about-fastlys-realtime-log-streaming-features) allows you to automatically save logs to a third-party service for storage and analysis. Logs provide an important resource for troubleshooting connectivity problems, pinpointing configuration areas that could use performance tuning (/guides/performance-tuning/), and identifying the causes of service disruptions. We recommend setting up remote log streaming when you start using Fastly services.

**NOTE:** Fastly does not provide direct support for third-party services. See [Fastly's Terms of Service \(https://www.fastly.com/terms\)](https://www.fastly.com/terms) for more information.

## Configuring logging endpoints

You can configure one or more logging endpoints for Fastly services. Follow these instructions to access the logging settings:




1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.
4. Click the **Logging** tab. The logging endpoints page appears.




| DOMAINS  | ORIGINS                   | SETTINGS                               | CONTENT                             | LOGGING     | CUSTOM VCL |
|----------|---------------------------|--|-------------------------------------|-------------|------------|
| 1 Domain | 1 Host<br>0 Health checks | 0 Request settings<br>0 Cache settings | 2 Headers<br>0 Gzips<br>0 Responses | 0 Endpoints | 0 Files    |



RECOMMENDED

## Choose your logging endpoint

For more information, read our [logging documentation](#).

**NOTE:** Some third-party services are disabled by default, so you won't see them in the Fastly web interface user interface until they've specifically been enabled for your account. To enable these services, contact [support@fastly.com](mailto:support@fastly.com) (<mailto:support@fastly.com>).

5. Follow the instructions in one of our guides for third-party services to complete the set up process and deploy your changes:

- Amazon S3 (</guides/streaming-logs/log-streaming-amazon-s3>)
- Cloud Files (</guides/streaming-logs/log-streaming-cloudfiles>)
- FTP (</guides/streaming-logs/log-streaming-ftp>)
- Google Cloud Storage (</guides/streaming-logs/log-streaming-google-cloud-storage>)
- Log Shuttle (</guides/streaming-logs/log-streaming-log-shuttle>)
- Logentries (</guides/streaming-logs/log-streaming-logentries>)
- Loggly (</guides/streaming-logs/log-streaming-loggly>)
- Heroku's Logplex (</guides/streaming-logs/log-streaming-logplex>)
- Papertrail (</guides/streaming-logs/log-streaming-papertrail>)
- OpenStack (</guides/streaming-logs/log-streaming-openstack>)
- Scalyr (</guides/streaming-logs/log-streaming-scalyr>)
- SFTP (</guides/streaming-logs/log-streaming-sftp>)
- Sumo Logic (</guides/streaming-logs/log-streaming-sumologic>)
- Syslog (</guides/streaming-logs/log-streaming-syslog>)

Once you've clicked Activate to deploy your changes, events will begin being logged immediately. The logs may take a few moments to appear on your log server.

## How, when, and where logs are streamed

To control log streaming, Fastly provides two versions of custom log formats (/guides/streaming-logs/custom-log-formats), each of which uses Apache-style logging directives ([http://httpd.apache.org/docs/2.4/mod/mod\\_log\\_config.html](http://httpd.apache.org/docs/2.4/mod/mod_log_config.html)). The logging format strings in each of these versions are based on the Common Log Format (<https://httpd.apache.org/docs/trunk/logs.html#common>) (CLF).

Logs are streamed over TCP, not UDP, optionally using TLS for security with supported endpoints. Additionally, if you are using custom VCL (/guides/vcl/uploading-custom-vcl) be sure to include the `#FASTLY deliver` (/guides/vcl/mixing-and-matching-fastly-vcl-with-custom-vcl) macro in your `vcl_deliver` handler.

By default, logs are placed in your root directory every hour using the file naming format `YYYY-mm-ddThh:mm:ss-<server id>`. You can change both the frequency and path of these files. Fastly uses several different log-server aggregation points and each will send logs files, none of which contain duplicate entries. These log files are created as soon as streaming starts and they're written to over the entire time period you specify (or the default). Once that time has passed, the files aren't touched any more and the logging process creates a new batch of files.

## Examples of other useful things you can log

Other useful things you can log include the following:

- `resp.http.X-Cache` - hits and misses
- `fastly_info.state` - misses, passes, hits, hit-passes, errors, stales, synthetics, and whether a request waited for request collapsing or whether it clustered
- `obj.hits` - the number of hits for an asset from a cache node—`server.identity` can then provide the specific cache node that received the hit
- `obj.lastuse` - approximate time elapsed since the object was last accessed, specific to each cache node
- `time.elapsed` - the time between the start of the request when it hits Fastly and when we deliver the first byte in the response
- `req.http.LiterallyAnyHeaderYouWant` - any request header from the client
- `req.http.Accept-Language` - users' Accept Language

You can also log any Varnish variable (<https://www.varnish-cache.org/docs/2.1/reference/vcl.html#variables>). Consider taking advantage of some of Fastly's extensions to VCL (/guides/vcl/guide-to-vcl#fastly-39-s-vcl-extensions) as well.

## Escaping characters in logs

Logs respond to VCL (/guides/vcl/) like any other object. For example, the following code can escape quotes from User-Agent your log stream:

```
log {"syslog serviceid endpointname :: "} {""} regsuball(req.http.user-agent, {""}, {"\\"}) {""};
```

## Preventing duplicate log entries when using custom VCL

If you use custom VCL (/guides/vcl/uploading-custom-vcl) commands for logging, you may notice duplicate entries in your logs. This happens because logs are being generated by both Fastly and the custom VCL logging commands. You can eliminate the duplicate entries by adding a condition that prevents Fastly from generating log entries. Follow these instructions to add the condition:

1. On the Logging endpoints page, click the **Attach a condition** link next to the appropriate logging service. The Add a condition window appears.
2. Click **Create a new response condition**. The Create a new response condition window appears.
3. Fill out the **Create a new response condition** window as follows:
  - Leave the **Type** menu set to **Response**.
  - In the **Name** field, type a human-readable name for the condition.
  - In the **Apply if** field, type `!req.url`. That condition will never be met, and is what prevents Fastly from generating log entries.
  - Leave the default value set in the **Priority** field.
4. Click **Save and apply to**.
5. Click the **Activate** button to deploy your configuration changes.

Fastly will stop generating log entries, and your logs will only contain entries generated by the custom VCL logging commands.

• [Guides \(/guides/\)](/guides/) > [Diagnostics and performance \(/guides/diagnostics\)](/guides/diagnostics/) > [Debugging \(/guides/debugging/\)](/guides/debugging/)

## § Browser recommendations when using the Fastly web interface (/guides/debugging/browser-recommendations-when-using-the-fastly-web-interface)

---

We support the following browsers:

- Google Chrome (<http://www.google.com/chrome>)
- Firefox (<https://www.mozilla.org/en-US/firefox/products/>)
- Safari (<http://www.apple.com/safari/download/>)

If you aren't using one of these browsers, then some visual styling may not be correct (<https://docs-next.fastly.com/next-gen-app-known-issues>) when using the Fastly web interface (<https://manage.fastly.com/>).

We strongly recommend updating your browser before beginning any debugging (/guides/debugging/) of Fastly services and before reporting problems to Fastly Customer Support (<mailto:support@fastly.com>). You can find the latest, downloadable versions of all major browsers online. The list at Browse Happy (<http://browsehappy.com/>) may help you.

---

## § Changing connection timeouts to your origin (/guides/debugging/changing-connection-timeouts-to-your-origin)

---

Connection timeouts to your origin server control how long Fastly will wait for a response from your origin server before exiting with an error. Changing the connection timeout is a good way to start troubleshooting 503 backend read errors (/guides/debugging/common-503-errors#503-backend-read-error). Follow the steps below to change the connection timeouts to your origin server:

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.
4. Click the **Origins** tab.
5. Click the link of the host that you want to edit. The Edit a host page appears.

| TIMEOUTS              |  |
|-----------------------|--|
| Connection timeout    | <input type="text" value="1000"/><br>How long to wait for a timeout in milliseconds.       |
| First byte timeout    | <input type="text" value="15000"/><br>How long to wait for the first byte in milliseconds. |
| Between bytes timeout | <input type="text" value="10000"/><br>How long to wait between bytes in milliseconds.      |

6. Type the new timeout in the appropriate field of the **Timeouts** section.
7. Click the **Update** button.

★ **TIP:** Additional techniques that help you gain insights into your service configurations can be found in our [Debugging \(/guides/debugging/\)](/guides/debugging/) guides.

## § Common 503 errors (/guides/debugging/common-503-errors)

---

Varnish (</guides/vcl/guide-to-vcl>), the software that powers the Fastly CDN, will sometimes return standardized 503 responses due to various issues that can occur when attempting to fetch data from your origin servers. The generic status text associated with a 503 error is "Service Unavailable." It can mean a wide variety of things. The most common reasons this generic text appears include:

1. The origin server generated a 503 error and Fastly passed it through as is.
2. The origin returned a 503 error without a response header, so Fastly used the default response.
3. The status line of the HTTP response from the origin was not parseable.
4. VCL code was run that used the "error" statement without an appropriate response status (e.g., `error 503` instead of `error 503 " _broken thing_ "`).

The following list provides the most common non-generic, standardized 503 responses and basic explanations for each.

**⚠ WARNING:** If you are seeing 503 errors, do not [purge all](/guides/purging/single-purges#purging-all-content) (</guides/purging/single-purges#purging-all-content>) cached content. Purge all overrides [stale-if-error](/guides/performance-tuning/serving-stale-content) (</guides/performance-tuning/serving-stale-content>) and increases the requests to your origin server, which could result in additional 503 errors.

## 503 Backend Read Error

This error typically appears if a timeout error occurs when Fastly cache servers attempt to fetch content from your origins. It can also be due to a variety of transient network issues that might cause a read to fail (e.g., router failovers, packet loss) or an origin overload.

**Benchmarking your backend response times.** Many outside factors cause backends response times to vary. Repeated, consistent Backend Read Errors frequently can be prevented by changing your backend timeout settings in the Fastly web interface. Start by running the following command to estimate response time for benchmarking purposes:

```
curl -s -w "%{time_total}\n" -o /dev/null http://example.com/path/to/file
```

**Increasing your backend timeout settings.** After benchmarking some of the slower paths in your application, you should have an idea of your ideal backend response time. Adjust the backend timeout values on the Edit a host page in the Advanced Options area. For more information, see our guide on changing connection timeouts to your origin (</guides/debugging/changing-connection-timeouts-to-your-origin>).

## 503 Connection Timed Out

This error occurs if the request times out while waiting for Fastly to establish a TCP connection to your origin or waiting for your origin to respond to the request. Similar to Backend Read Errors, connection timeouts can be caused by transient network issues, long trips to origin, and origin latency. Two common ways to alleviate these timeout errors include increasing the connection timeout values set for the Fastly host (</guides/debugging/changing-connection-timeouts-to-your-origin>) or setting up an origin shield (</guides/performance-tuning/shielding>).

## 503 Backend Write Error

This error is similar to Backend Read Error but occurs when Fastly sends information in the form of a POST request to the backend.

## 503 Client Read Error

This error generally occurs because of a network issue between the client and Fastly. It is similar to the Backend Read Error but occurs when reading information from a client.

## 503 Connection Refused

This error occurs when Fastly attempts to make a connection to your origin over a specific port and the server refuses the connection. It typically appears when the wrong port is specified for the host in the Fastly web interface.

## 503 Network Unreachable

This error appears when Fastly cannot find a route to the given IP range. This generally occurs because of misconfigured or non-operational routers.

## 503 Backend Is Unhealthy

This error appears when custom health checks report a backend as down. It typically occurs when a Fastly edge server receives a client request and must make a request to your origin, but because the backend is considered unhealthy, Fastly doesn't try to send the request at all. This error may mistakenly appear instead of the Backend.max\_conn Reached error the first time Fastly encounters the maximum number of connections to your backend.

## 503 Backend.max\_conn Reached

This error occurs when Varnish makes a request to a backend in your Fastly service that has reached its defined maximum number of connections. By default, Fastly limits you to 200 origin connections from a single cache node to protect the origins from overload. To correct this error, increase the maximum connections limit to your origin. This error may also appear as "Maximum Threads for Service Reached."

## 503 Maximum Threads for Service Reached

See `Backend.max_conn` Reached.

## 503 Illegal Vary Header From Backend

This error occurs when a backend returns a malformed vary header with its response. A well-formed vary header (<https://www.fastly.com/blog/best-practices-for-using-the-vary-header>) tells Fastly to serve a different version of an object based on the value of the request header included within it.

## 503 No Stale Object Available

This error occurs when you configure Fastly to serve stale objects (</guides/performance-tuning/serving-stale-content>) in the event of a backend failure but the stale object has expired and your backend is still failing for some reason (thus, no stale object is available).

## 503 Quorum Wait Not Reached

This error occurs when a Director (`/api/config#director`) used for balancing requests among a group of backends (only available via the Fastly API) cannot serve traffic based on its configuration because it does not have enough available backends in its group.

## 503 No Healthy Backends

This error occurs when a Director (`/api/config#director`) used for balancing requests among a group of backends (only available via the Fastly API) cannot cache the specified content because there are no healthy backends available in its group.

## 503 All Backends Failed or Unhealthy

This error occurs when a Director (`/api/config#director`) used for balancing requests among a group of backends (only available via the Fastly API) fails because all the backends are unhealthy or multiple backends from which the Director tried to fetch information failed with the same error.

## 503 SSL Handshake error

This error occurs when TLS negotiation between Fastly and your origin fails. To fix this error, review and correct your host's TLS configurations (</guides/securing-communications/connecting-to-origins-over-tls>).

---

# § Curl and other caching verification methods

## (</guides/debugging/curl-and-other-caching-verification-methods>)

---

The easiest way to tell if your request is caching in the Fastly network is to use a cURL (</guides/basic-setup/glossary-of-terms#curl>) on the command line. We recommend using one of two cURL commands for debugging purposes:

- a simple cURL command that displays the request and response headers for a given object
- a slightly more complex cURL command that uses the `Fastly-Debug` header to expose information normally stripped by the simple cURL

## Using the simple cURL command

The following cURL command displays the request and response headers for a given object:

```
curl -svo /dev/null www.example.com/index.html
```

where `www.example.com/index.html` is replaced with the full object path of the object you're testing.

For example, using `curl -svo /dev/null www.example.com` produces something like the following section of output:

```
[...]  
< Age: 142  
< X-Served-By: cache-jfk1041-JFK, cache-ord1720-ORD  
< X-Cache: HIT, HIT  
< X-Cache-Hits: 1, 7  
[...]
```



This output tells us the current age of the object in cache. It also shows shielding is enabled because two cache nodes display in `X-Served-By`. However, we're most interested in the output of the `X-Cache` header. A properly caching object displays a value of `X-Cache: HIT`, `X-Cache: HIT, HIT`, `X-Cache: HIT, MISS`, or `X-Cache: MISS, HIT`.

## Using a Fastly-Debug header with cURL

The `Fastly-Debug` header provides additional information for debugging by exposing specific information that is normally stripped when using a simple cURL command:

```
curl -svo /dev/null -H "Fastly-Debug:1" www.example.com/index.html
```

where `www.example.com/index.html` is replaced with the full object path of the object you're testing.

For example, with optional shielding being used and a TTL set to 86400 (24 hours) using `Surrogate-Control`, the command `curl -svo /dev/null -H "Fastly-Debug:1" www.example.com` produces something like the following section of output:

```
[...]
< Surrogate-Control: max-age=86400
< Surrogate-Key: articles articles/1 articles/2

[...]
< Age: 403
< Fastly-Debug-Path: (D cache-ord1722-ORD 1470672957) (F cache-ord1743-ORD 1470672629) (D cache-jfk1041-JFK 1470672629) (F cache-jfk1030-JFK 1470672554)
< Fastly-Debug-TTL: (H cache-ord1722-ORD 85997.246 0.000 403) (H cache-jfk1041-JFK - - 75)
< X-Served-By: cache-jfk1041-JFK, cache-ord1722-ORD
< X-Cache: HIT, HIT
< X-Cache-Hits: 1, 6

[...]
```

Because surrogate keys are present, the `Fastly-Debug` header exposes them. As with the simple cURL command, this section of output tells us the current age of the object in cache. In addition, `Fastly-Debug` exposes specific header details to help with debugging as noted below.

## Information exposed by the Fastly-Debug header

**Fastly-Debug Path** contains information about which cache server handles fetching and delivery of an object and object TTL information. The edge POP appears first in the sequence and the shield POP appears second.

- `D` represents which cache by name in the edge or shield ran `vcl_deliver`
- `F` represents which cache by name in the edge or shield ran `vcl_fetch`
- the number following each specific server name is a timestamp in seconds

With shielding enabled, you should generally see four cache servers listed in this header. In rare cases where a cache server exists as both an edge and a shield within the cluster for that object, you may see two or three caches listed.

**Fastly-Debug-TTL** provides information on HIT and MISS timings.

- `H` represents a HIT, meaning the object was found in the cache
- `M` represents a MISS, meaning the object was not cached at the time of the query

For each of these timings:

- the first number specifies the TTL remaining for the object
- the second number specifies the grace period
- the third number specifies the current age of the object in cache

Because the cluster node determines these values, it may take a few requests to see these numbers populate as expected.

**X-Served-By** indicates the shield and edge servers that were queried for the request. The shield POP appears first in the sequence and the edge POP appears second.

**X-Cache:** indicates whether the request was a HIT or a MISS for the datacenter.

**NOTE:** Our guide to [understanding cache HIT and MISS headers \(/guides/performance-tuning/understanding-cache-hit-and-miss-headers-with-shielded-services\)](/guides/performance-tuning/understanding-cache-hit-and-miss-headers-with-shielded-services) provides in depth details key to understanding the `X-Served-By`, `X-Cache`, and `X-Cache-Hits` headers with shielded services.

## § Debugging with mtr (/guides/debugging/debugging-with-mtr)

For diagnostics and debugging in the Fastly network, we think the mtr (<http://www.bitwizard.nl/mtr/>) tool offers a great way to test network speed, evaluate performance, and perform connection diagnostics. The program's source and installation instructions live in GitHub (<https://github.com/traviscross/mtr>).

While mtr provides a number of practical uses for network engineering needs, the following command works well:

```
mtr -c 20 -w -r www.example.com
```

Be sure to replace `www.example.com` with the hostname of the domain you're working with. The command will generate the network hops to the destination you specify, any packet loss experienced, and aggregate connection statistics.

For example, if we wanted to test the network connection from Fastly's San Francisco office to the CDN, we would use the above command for `www.fastly.com`. The following output would appear:

```
~ mtr -c 20 -w -r www.fastly.com
Start: Mon Feb  2 15:27:20 2015
HOST: test-local-machine.local
Loss%  Snt  Last  Avg  Best  Wrst  StDev
1. |-- 10.100.20.2                0.0%   20   2.1   2.2   1.6   4.3   0.5
2. |-- ge-4-3-4.mpr4.sfo7.us.zip.zayo.com 0.0%   20   2.3   2.4   1.8   5.2   0.6
3. |-- ae5.cr2.sjc2.us.zip.zayo.com       0.0%   20   4.6   6.5   2.9  35.3   7.7
4. |-- ae10.mpr4.sjc7.us.zip.zayo.com     0.0%   20   4.7   4.8   3.6  14.5   2.3
5. |-- be6461.ccr21.sjc03.atlas.cogentco.com 5.0%   20   5.1   5.9   4.2  15.3   2.6
6. |-- fastly-inc.edge2.sanjose3.level13.net 0.0%   20   5.0   4.7   4.2   8.2   0.8
7. |-- ???                        100.0   20   0.0   0.0   0.0   0.0   0.0
8. |-- 23.235.47.184              0.0%   20   4.7  14.3   3.8  74.6  20.3
```

## § Debugging with WebPageTest (/guides/debugging/debugging-with-webpagetest)

It's important to establish habits of testing and performance before, during, and after migrating to Fastly. This allows you to clearly measure the impact of tests and changes to your infrastructure.

One tool that Fastly recommends for this purpose is WebPageTest.org (<http://www.WebPageTest.org>). WebPageTest provides a free and open source testing tool for deep performance analysis. It is built on browser technology to accurately replicate what your end users encounter when visiting a website.

We recommend using the WebPageTest defaults for basic testing, but keep a few rules in mind:

- On the Test Settings tab under Advanced Settings, Connection should always be set to Native Connection during initial benchmarks.
- Two to three test runs may be required before a site is properly caching in Fastly.
- Using WebPageTest's "Visual Comparison" (<http://www.webpagetest.org/video/>) feature offers an ideal way to A/B test potential changes.

## § Error 1000 with CloudFlare DNS (/guides/debugging/error-1000-with-cloudflare-dns)

Using CloudFlare for DNS and other CDNs can cause CloudFlare to show an Error 1000 indicating that your DNS points to prohibited IP addresses. This occurs when the hostnames are CNAME'd to Fastly (</guides/basic-setup/adding-cname-records>) and an origin server is configured as a fully qualified domain name (FQDN) within Fastly:

**old-server.example.com : 80**    Attach a condition    TLS Options    Delete  
*My old server*  
[Show all details](#)

To solve this error, direct Fastly to use the IP address as the host for any backend origin servers. This removes the need to resolve the hostname for traffic to the servers:

**192.0.1.2 : 80**    Attach a condition    TLS Options    Delete  
*AWS EC2 instance*  
[Show all details](#)

You can also change this by modifying the VCL configuration files directly. For example, this VCL:

```
backend F_Hosting_server_Example_Backend {  
    ...  
    .port = "80";  
    .host = "exampleserver.exampledomain.tld";  
}
```

would become:

```
backend F_Hosting_server_Example_Backend {  
    ...  
    .port = "80";  
    .host = "12.34.56.78";  
}
```

---

## § Fastly's network status (/guides/debugging/fastlys-network-status)

---

Fastly continuously monitors the status of our global network and all related services. In the event of a service interruption, an update will be posted on the Fastly status page at [status.fastly.com](https://status.fastly.com) (<https://status.fastly.com>). If you are experiencing problems and do not see a notice posted, email [support@fastly.com](mailto:support@fastly.com) (<mailto:support@fastly.com>) for assistance.

**fastly** | Service Status [Support](#) [Documentation](#) [Blog](#) [Login](#) [SUBSCRIBE TO UPDATES](#)

Fastly's network has built-in redundancies and automatic failover routing to ensure optimal performance and uptime. But when a network issue does arise, we think our customers deserve clear, transparent communication so they can maintain trust in our service and our team. Notices will be posted here when we re-route traffic, upgrade hardware, or in the extremely rare case our network isn't serving traffic. If you are experiencing issues and do not see a notice posted, please email [support@fastly.com](mailto:support@fastly.com) for assistance.

**All Systems Operational** Refreshed 2 minutes ago

|                         |             |
|-------------------------|-------------|
| API                     | Operational |
| Stats <a href="#">?</a> | Operational |
| Configuration panel     | Operational |
| + North America         | Operational |
| + South America         | Operational |
| + Europe                | Operational |
| + Asia/Pacific          | Operational |

## Scheduled Maintenance

**New Datacenter: Dubai** Scheduled for Jul 15, 00:00-00:15 UTC

As part of Fastly's global network expansion, we're excited to let you know that we will be adding a Dubai, United Arab Emirates (FJR) point of presence (POP) to Fastly's network next month.

## Overall system status

The current system status appears at the top of the Fastly status page and includes the last time the status was refreshed so that you know how current the information is.

**All Systems Operational** Refreshed less than one minute ago

## Individual component statuses

The status of the Fastly API ([/api/](https://api/)), the Fastly web interface (<https://manage.fastly.com/>), statistics collection and delivery, and each Fastly point of presence (<https://www.fastly.com/network-map>) (POP) appears immediately below the overall status. POPs are grouped by region. You can see the status of all POPs in a region by clicking the + icon next to the region's name.

|                      |             |
|----------------------|-------------|
| API                  | Operational |
| Stats <span>?</span> | Operational |
| Configuration panel  | Operational |
| ⊕ North America      | Operational |
| ⊕ South America      | Operational |
| ⊕ Europe             | Operational |
| ⊕ Asia/Pacific       | Operational |

## Past incident statuses

Fastly keeps track of past incidents. Past incidents, if any, for approximately the past two weeks appear immediately below the individual component statuses.

### Past Incidents

Jun 30, 2016

---

No incidents reported today.

In addition to the textual description, each incident status appears in a color that indicates the level of service impact. The color indicators are as follows:

- Black - no component marked specifically as out of service or degraded
- Yellow - minor degradation or disruption
- Orange - significant degradation or traffic rerouting
- Red - component offline

We also keep track of all past incidents in an incident history page (<https://status.fastly.com/history>).

## Subscribing to notifications

Fastly allows you to subscribe to status notifications via email or SMS text messaging. Simply click the **Subscribe to Updates** button in the upper right corner of the status page screen. Once subscribed, we'll email you any time we create or update an incident.

The image shows two side-by-side screenshots of the Fastly status page. Both screenshots have a 'Login' button and a red 'SUBSCRIBE TO UPDATES' button in the top right corner. Below the 'SUBSCRIBE TO UPDATES' button is a row of icons: an envelope (email), a phone (SMS), a code symbol (API), a speech bubble (chat), a signal tower (RSS), and an 'X' (unsubscribe).

The left screenshot shows the email subscription modal. It contains the text: 'Get email notifications whenever Fastly creates or updates an incident.' Below this is a text input field labeled 'Email Address' and a red button labeled 'SUBSCRIBE VIA EMAIL'. At the bottom of the modal, it says 'Refreshed 15 minutes ago'.

The right screenshot shows the SMS subscription modal. It contains the text: 'Get SMS notifications whenever Fastly creates or resolves an incident.' Below this is a text input field with the placeholder 'ex. 6505551234'. Below the input field is a red button labeled 'SUBSCRIBE VIA SMS'. At the bottom of the modal, it says 'Message and data rates may apply. Go to [www.statuspage.io/terms](http://www.statuspage.io/terms) for terms and conditions.'

To subscribe to email notifications, click the letter icon, type your email address in the displayed field, and click **Subscribe Via Email**. You can unsubscribe at any time by clicking the unsubscribe link that appears at the bottom of every status email.

To subscribe via SMS text messaging, click the telephone icon, type your telephone number in the displayed field, and click **Subscribe Via SMS**. Unsubscribe from SMS text messaging at any time by replying STOP to any status message you receive.

## § Google Pagespeed module errors (/guides/debugging/google-pagespeed-module-errors)

If you are using the Google Pagespeed module and notice constant MISSES for HTML pages, check the Cache-Control settings in the module's .htaccess file.

By default, Google Pagespeed serves all HTML with `Cache-Control: no-cache, max-age=0`. This setting conflicts with Fastly's default configuration. If your origin sends the headers `Cache-Control: private` or `Cache-Control: max-age=0`, Fastly will pass requests straight to the origin.

To change the Google Pagespeed directive and leave the original HTML caching headers, update your origin's .htaccess file with:

```
ModPagespeedModifyCachingHeaders off
```

More details about the Pagespeed Module (<https://developers.google.com/speed/pagespeed/module/install>) can be found within Google Developers directory. For additional information about controlling how long Fastly caches your resources, start with our Cache Control Tutorial (/guides/tutorials/cache-control-tutorial)

## § Googlebot crawl stats (/guides/debugging/googlebot-crawl-stats)

Any time you notice any major changes in your SEO stats, indexing, or crawler behavior, start troubleshooting by asking these questions:

- Did you read the Google FAQs (<https://sites.google.com/site/webmasterhelpforum/en/faq--crawling--indexing---ranking>) for indexing, crawling, and ranking?
- Is your robots.txt file still accessible and were there any changes to it?
- Is your sitemap testing without errors ([https://support.google.com/webmasters/answer/183669?hl=en&ref\\_topic=6080662&rd=2](https://support.google.com/webmasters/answer/183669?hl=en&ref_topic=6080662&rd=2))?
- Did you adjust your Googlebot crawl rate (<https://support.google.com/webmasters/answer/48620>)?
- Have you had Google's "Fetch as Google" (<https://support.google.com/webmasters/answer/6066468>) tool re-crawl the URLs?

We recommend exploring Google's Webmaster Tools (<https://www.google.com/webmasters/>) if you're experiencing issues. Their "Fetch as Google" tool article and their article on troubleshooting sitmemap errors offer specific help for help debugging Googlebot crawl stats in this situation. Google also includes an entire section in their tools documentation on getting additional support (<https://support.google.com/webmasters/answer/1249981?hl=en>) if you're experiencing trouble.

★ **TIP:** Our [debugging articles \(/guides/debugging/\)](/guides/debugging/) contain a variety of troubleshooting tips.

## § Loop detection (/guides/debugging/loop-detection)

Fastly automatically detects loops resulting from service configuration errors. When a loop is detected, Fastly blocks the requests and generates an error message. Loops can occur when the same hostname is configured as both the domain and the origin server, and the CNAME record for the domain is pointed at Fastly. For example, loop detection will be triggered if you set `www.example.com` as the domain and the origin server (/guides/basic-setup/sign-up-and-create-your-first-service#create-your-first-service-and-test-it) in your Fastly service and you add a CNAME DNS record (/guides/basic-setup/adding-cname-records) for `www.example.com` that points at Fastly.

### How to avoid triggering loop detection

To avoid triggering loop detection, you should verify the hostname of your origin server is not the same as the domain using one of the following two options:

- Create a DNS hostname (`origin.example.com`) with the appropriate A and AAAA DNS records for your origin server, and use that origin DNS hostname in your Fastly service configuration. This ensures the origin (`origin.example.com`) is different than the domain (`www.example.com`) on your service. We recommend this option. If you make changes to the DNS records for `origin.example.com` in the future, Fastly will automatically detect and use those changes.

- Use an IPv4 address instead of a DNS hostname for your origin's address within your Fastly service's configuration. If the origin server's IP address changes in the future, you'll need to update and activate a new version of your Fastly service configuration.

## Example error message

When Fastly detects a loop, an error message similar to the one displayed below will appear in the headers.

```
HTTP/1.1 503 Loop detected
Error-Reason: loop detected
Connection: close
Content-Type: text/plain
Fastly-Host: <hostname>
Fastly-FF: <hostname>
Server: Varnish
```

## § Temporarily disabling caching (/guides/debugging/temporarily-disabling-caching)

Caching can be disabled:

- at the individual URL level,
- at the browser level, and
- at the site level.

### Disabling caching at the individual URL level

To disable caching at the individual URL level:

1. Create a request setting that always forces a pass (/guides/basic-configuration/how-request-settings-are-applied).
2. Add a condition to the request setting (/guides/vcl/understanding-the-different-pass-action-behaviors) that looks for specific URLs.
3. Activate the new version of your service to enable the setting.

### Disabling caching at the browser level

Theoretically, all browsers should follow the stated rules of the HTTP standard. In practice, however, some browsers don't strictly follow these rules. The following combination of headers seems to force absolutely no caching with every browser we've tested.

```
Cache-Control: no-cache, no-store, private, must-revalidate, max-age=0, max-stale=0, post-check=0, pre-check=0
Pragma: no-cache
Expires: 0
```

In addition, IE8 has some odd behavior to do with the back button. Adding `Vary: *` to the headers seems to fix the problem.

**❗ IMPORTANT:** If you want your content cached in Fastly but not cached on the browser, you must not add these headers on your origin server. Instead, add these as new Headers on the Content page and be sure the Type is set to Response (/guides/basic-configuration/responses-tutorial).

### Disabling caching at the site level

You can disable caching at the site level via the Fastly UI or via custom VCL. Via the UI, create a request setting that always forces a pass (/guides/basic-configuration/how-request-settings-are-applied) and then activate the new version of your service to enable the setting. To disable caching at the site level via custom VCL (/guides/vcl/mixing-and-matching-fastly-vcl-with-custom-vcl), add this to your Fastly VCL:

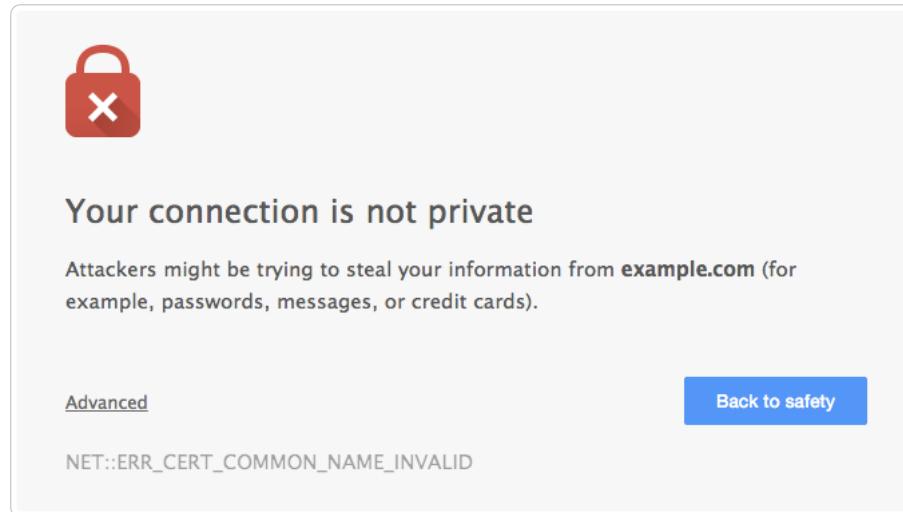
```
sub vcl_recv {
    return(pass);
}
```

**❗ IMPORTANT:** The ability to upload custom VCL code (/guides/vcl/uploading-custom-vcl) is disabled by default when you sign up. To enable this ability for your account, contact [support@fastly.com](mailto:support@fastly.com) (mailto:support@fastly.com) and request it.

## § TLS certificate errors (/guides/debugging/tls-certificate-errors)

## "Your connection is not private"

If you've recently started testing Fastly services, you may see errors like the following:



These errors appear because your domain has not been provisioned with TLS across the Fastly network. We offer a number TLS options (/guides/securing-communications/ordering-a-paid-tls-option) that may work for you. Contact support@fastly.com (mailto:support@fastly.com) to begin the provisioning process.

If you don't want to use TLS for your site, set the CNAME DNS record for your domain to point to `global-nossl.fastly.net`. This network endpoint only accepts requests over port 80, and will not expose your users to these certificate errors.

## Errors when using Wget

When connecting to a Fastly service using Wget, you may see errors along the lines of

```
ERROR: Certificate verification error for mysite.example.com: unable to get local issuer certificate
ERROR: certificate common name `*.a.ssl.fastly.net' doesn't match requested host name `mysite.example.com'.
To connect to mysite.example.com insecurely, use `--no-check-certificate'.
Unable to establish TLS connection.
```

Checking with a browser or cURL will show that there really is no problem, however. The errors appear because a previous version of Wget (wget-1.12-2.fc13) that shipped with some versions of Red Hat Enterprise Linux (RHEL) was buggy and failed to check Subject Alternative Names (SAN) properly.

Upgrading Wget will correct this problem and eliminate the errors. For more information you can read this Red Hat bug report (https://bugzilla.redhat.com/show\_bug.cgi?id=674186) or this Debian one (https://bugs.debian.org/cgi-bin/bugreport.cgi?bug=409938). For more information about TLS-related issues, see our TLS guides (/guides/securing-communications/) or contact support@fastly.com (mailto:support@fastly.com) with questions.

## § TLS origin configuration messages (/guides/debugging/tls-origin-configuration-messages)

When you are connecting to origins over TLS (/guides/securing-communications/connecting-to-origins-over-tls), you may have errors.

### Hostname mismatches

- `Error: Hostname mismatch`

### Why the error appears

Your origin server is serving a TLS certificate with a Common Name (CN) or list of Subject Alternate Names (SAN) that does not match the origin host or the origin's SSL hostname setting.

### How to fix it

You can fix this by telling Fastly what to match against in the CN or SAN field in your origin's certificate.

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.



3. Click the **Edit configuration** button and then select **Clone active**.
4. Click the **Origins** tab. The Origins page appears.
5. Click the **TLS Options** link next to the affected host. The TLS Options page appears.
6. In the **Certificate Hostname** field, type the hostname associated with your TLS certificate. This value is matched against the certificate common name (CN) or a subject alternate name (SAN) depending on the certificate you were issued. For example, if your certificate's CN field is `www.example.com`, type that value for your hostname. If you leave this field blank, the system will use the default host information displayed below this field.
7. Click the **Save** button. Even if you leave the Certificate Hostname field blank to use the default information, you must click the Save button to verify the certificate.
8. Activate a new version of the service to complete the configuration changes.

When using custom VCL (/guides/vcl/uploading-custom-vcl), you can specify the hostname to match against the certificate by using the

`.ssl_cert_hostname` field of your origin's definition. For example: `.ssl_cert_hostname = www.example.com;`.

## Certificate chain mismatches

- `Error: unable to verify the first certificate`
- `Error: self signed certificate`
- `Error: unable to get local issuer certificate`
- `Error: self signed certificate in certificate chain`
- `Error: unable to get issuer certificate`

## Why the errors appear

Your origin server is serving a certificate chain that can not be validated using any of the Certificate Authorities (CAs) that Fastly knows. This can happen for two reasons:

- Your certificate is self-signed or self-issued and you did not provide your generated CA certificate to Fastly for us to use for verification.
- Your certificate is issued by a CA that isn't in Fastly's CA certificates bundle.

## How to fix them

In both cases, you can fix your configuration by adding the CA certificate that Fastly should use to verify the certificate to your service configuration:

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Edit configuration** button and then select **Clone active**.
4. Click the **Origins** tab. The Origins page appears.
5. Click the **TLS Options** link next to the affected host. The TLS Options page appears.
6. In the **TLS CA certificate** field, copy and paste a PEM-formatted CA certificate.
7. Click the **Save** button.
8. Activate a new version of the service to complete the configuration changes.

If you are using custom VCL, you can specify the CA for Fastly to use by setting the `.ssl_ca_cert` backend parameter to a PEM encoded CA certificate.

Alternatively, you can get a new certificate issued by a CA in Fastly's CA certificate bundle (e.g., Globalsign).

## Connection failures

- `Error: Gethostbyname`
- `Error: Connection timeout`
- `Error: Connection refused`

## Why each error appears and how to fix it

For `Gethostbyname` failures, the configured backend Host domain is returning NXDOMAIN. Double check that the DNS settings for your backend are correct.

For `Connection time out` failures, the connection to your server is timing out. Double check that your backend is accessible and responding in a timely fashion.

For `Connection refused` failures, the connection to your server is being refused, potentially by a firewall or network ACL. Double check that you have whitelisted the Fastly IP addresses (/guides/securing-communications/accessing-fastlys-ip-ranges) and that your backend is accessible from our network.

## Certificate expirations

`Error: Certificate has expired`

The certificate your backend server is presenting Fastly has expired and needs to be reissued with an updated validity period.

If this is a self-signed certificate you can perform this update on your own by issuing a new CSR with your private key, creating the corresponding certificate, and installing it on the server.

If this is a CA signed certificate you will need to issue a new CSR with your private key, submit it to your CA, and install the signed certificate they provide you.

## SSL and old TLS protocol errors

- `Error: Unknown protocol`
- `Error: SSL handshake failure`
- `Error: TLSv1 alert internal error`

### Why the errors appear

Either your origin server is not configured to use TLS or it only supports older, outdated versions of the protocol. We do not support SSLv2 or SSLv3.

### How to fix them

If the origin server is configured to use TLS, make sure you are using the latest version of both the server and the TLS library (e.g., OpenSSL). You may have to explicitly enable a newer protocol version. Fastly supports TLSv1, TLSv1.1 and TLSv1.2.

If the origin server is not configured to use TLS, change your service configuration to disable TLS and communicate with it on port 80 instead of port 443:

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Edit configuration** button and then select **Clone active**.
4. Click the **Origins** tab. The Origins page appears.
5. Click the **TLS Options** link next to the affected host. The TLS Options window appears.
6. From the **Connect to backend using TLS** menu, select **No**.
7. Click the **Save** button.
8. Activate a new version of the service to complete the configuration changes.

## RC4 cipher error

- `Error: Using RC4 Cipher`

### Why the error appears

When Fastly connects to your origin server using TLS, the only cipher suite your server supports for establishing a connection is the RC4 cipher. This cipher is considered to be unsafe for general use and should be deprecated.

### How to fix it

You can fix this on your origin by using the latest version of both the server and the TLS library (e.g., OpenSSL) and ensuring the cipher suites offered are tuned to best practices. You may need to explicitly blacklist the RC4 cipher.

## § Using GET instead of HEAD for command line caching tests (/guides/debugging/using-get-instead-of-head-for-command-

# line-caching-tests)

If you're testing on the command line to determine an object's caching status, then use GET instead of HEAD. For example:

```
curl -svo /dev/null www.example.com
```

## Default caching behavior of HTTP verbs

By default, the results of GET requests are cached. HEAD requests are not proxied as is, but are handled locally if an object is in cache or a GET is done to the backend to get the object into the cache. Anything other than HEAD or GET requests are proxied and not cached by default. However, POSTs *can* be cached if required, and caching of that can be keyed off of the contents of the POST request body if that is below 2K. We've posted an example of how to do this (</guides/vcl/manipulating-the-cache-key#using-a-post-request-body-as-a-cache-key>).

• Guides (</guides/>) > Diagnostics and performance (</guides/diagnostics>) > Performance tuning (</guides/performance-tuning/>)

## § Changing origins based on user location (</guides/performance-tuning/changing-origins-based-on-user-location>)

Fastly allows you to change origins based on user location. The VCL looks something like this:

```
# default conditions
set req.backend = F_global;

# Use restricted content if the user is in Asia, France or Germany
if (geoip.continent_code == "AS" || geoip.country_code == "fr" || geoip.country_code == "de") {
    set req.backend = F_restricted_content;
}
```

So, how does this translate to the Fastly UI? First, create a new header (</guides/basic-configuration/adding-or-modifying-headers-on-http-requests-and-responses>) on the Content page of the Configure page (</guides/about-fastly-services/about-the-web-interface-controls#about-the-configure-page>). The Type is Request and the Action is Set. Note that the priority is 10. The backend (here it's F\_global) should be the name of whatever your global origin server is. You can see the name if you look at your VCL.

## Create a new header

Learn more about this section in our [headers tutorial](#).

**CONDITION** This will happen all the time unless you [Attach a condition](#)

|   |   |                                  |
|---|---|----------------------------------|
| Type / Action   | <input type="text" value="Request"/>            | <input type="text" value="Set"/> |
| The type of header and the action performed on it.  |   |                                  |
| Destination   | <input type="text" value="backend"/>            | <span>★ Required</span>          |
| The name of the header that will be affected by the selected action. For example: <b>http.Content-Type</b> , <b>http.Set-Cookie</b> , <b>http.Via</b> , <b>http.Location</b> , or <b>http.Access-Control-Allow-Origin</b> . |   |                                  |
| Source  | <input type="text" value="F_global"/>           | <span>★ Required</span>          |
| New content for the header. Can be a static value (e.g. string or number) or a dynamic value (e.g. existing header or a GeoIP value). Please use quotes for string values.  |   |                                  |
| Ignore if set   | <input type="text" value="No"/>                 |                                  |
| If switched to <b>Yes</b> , the action will not be performed if the header in Destination exists.   |   |                                  |
| Priority  | <input type="text" value="10"/>                 | <span>★ Required</span>          |
| The order in which the header rules execute within the condition.   |   |                                  |
| Description   | <input type="text" value="Set default origin"/> | <span>★ Required</span>          |
| The name of your header, such as <b>My header</b> .   |   |                                  |

[CREATE](#)[CANCEL](#)

Now create another header almost exactly the same with two differences: it sets the backend to be your restricted origin server and the priority is higher than the previous header.

## Create a new header

Learn more about this section in our [headers tutorial](#).

**CONDITION** This will happen all the time unless you [Attach a condition](#)

|   |  |                                  |
|---|--|----------------------------------|
| Type / Action   | <input type="text" value="Request"/>               | <input type="text" value="Set"/> |
| The type of header and the action performed on it.  |  |                                  |
| Destination   | <input type="text" value="backend"/>               | Required                         |
| The name of the header that will be affected by the selected action. For example: <code>http.Content-Type</code> , <code>http.Set-Cookie</code> , <code>http.Via</code> , <code>http.Location</code> , or <code>http.Access-Control-Allow-Origin</code> . |  |                                  |
| Source  | <input type="text" value="F_restricted_content"/>  | Required                         |
| New content for the header. Can be a static value (e.g. string or number) or a dynamic value (e.g. existing header or a GeoIP value). Please use quotes for string values.  |  |                                  |
| Ignore if set   | <input type="text" value="No"/>                    |                                  |
| If switched to <b>Yes</b> , the action will not be performed if the header in Destination exists.   |  |                                  |
| Priority  | <input type="text" value="11"/>                    | Required                         |
| The order in which the header rules execute within the condition.   |  |                                  |
| Description   | <input type="text" value="Set restricted origin"/> | Required                         |
| The name of your header, such as <b>My header</b> .   |  |                                  |

[CREATE](#)[CANCEL](#)

Finally, attach a condition (</guides/conditions/>) to the restricted origin header which checks the GeoIP header (</guides/vcl/geoip-related-vcl-features/>):

Create a new request condition.  
Learn the basics in our [conditions tutorial](#)

Type

Request

Learn more about the [different types of conditions](#).

Name

From Restricted Location

\*

Apply if...

geoip.content\_code == "AS" || geoip.content\_code == "EU"

The expression to decide whether this is run.

[Examples](#)

[Adjust priority \(default is 10\)](#)

This is an advanced option. For most configurations, the default is best.

SAVE AND APPLY TO FORCE TLS

CANCEL

## § Checking multiple backends for a single request (/guides/performance-tuning/checking-multiple-backends-for-a-single-request)

Using a restart is a good option to check multiple backends for a single request. This can be created using a cache setting rule (/guides/tutorials/cache-control-tutorial) and request headers.

### Create a new cache setting rule

Follow these steps to create a cache restart within `vcl_fetch`.

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.
4. Click the **Settings** tab. The Settings page appears.
5. Click the **Create cache setting** button. The Create a new cache setting page appears.

## Create a new cache setting

This will override your cache control headers. In most cases, use with conditions applied.

**CONDITION** This will happen all the time unless you [Attach a condition](#)

TTL (seconds)

The TTL (Time To Live) entered will set the lifespan of the object within our cache nodes.

Action

This setting decides [how the request will be handled](#).

Stale TTL (seconds)

Stale TTL (Time To Live) is used by your application to [serve stale content](#). It determines how long to keep stale data after it has expired. Note there's custom VCL required to [complete this setup](#).

Description

★ Required

The name of your cache setting, such as **My cache setting**.

CREATE

CANCEL

6. Fill out the **Create a new cache setting** fields as follows:

- In the **TTL (seconds)** field, type .
- From the **Action** menu, select **Restart processing**.
- In the **Stale TTL (seconds)** field, type .
- In the **Description** field, type  (or any meaningful, preferred name).

7. Click the **Create** button. The new cache setting appears on the Settings page.

8. On the **Settings** page, click the **Attach a condition** link next to the cache setting you just created. The Create a new cache condition window appears.

Create a new cache condition

Learn the basics in our [conditions tutorial](#)

Type

Cache

Learn more about the [different types of conditions](#).

Name

Request Restart

\*

Apply if...

beresp.status != 200 && beresp.status != 304

The expression to decide whether this is run.

[Examples](#)

[Adjust priority \(default is 10\)](#)

This is an advanced option. For most configurations, the default is best.

SAVE AND APPLY TO RETURN RESTART

CANCEL

9. Fill out the **Create a new cache condition** fields as follows:

- Leave the **Type** menu set to **Cache**.
- In the **Name** field, type `Restart Request` (or any meaningful, preferred name).
- In the **Apply if** field, type `beresp.status != 200 && beresp.status != 304`.

10. Click the **Save and apply to** button to create the condition.

## Create new request headers

Follow these steps to create a request header within `vcl_recv`.

1. Click the **Content** tab. The Content page appears.
2. Click the **Create header** button. The Create a new header page appears.



## Create a new header

Learn more about this section in our [headers tutorial](#).

### CONDITION

This will happen all the time unless you [Attach a condition](#)

Type / Action

Request

Set

The type of header and the action performed on it.

Destination

http.Fastly-Force-Shield

★ Required

The name of the header that will be affected by the selected action. For example: **http.Content-Type**, **http.Set-Cookie**, **http.Via**, **http.Location**, or **http.Access-Control-Allow-Origin**.

Source

"yes"

★ Required

New content for the header. Can be a static value (e.g. string or number) or a dynamic value (e.g. existing header or a GeolP value). Please use quotes for string values.

Ignore if set

No

If switched to **Yes**, the action will not be performed if the header in Destination exists.

Priority

10

★ Required

The order in which the header rules execute within the condition.

Description

Fastly Internal Shielding

★ Required

The name of your header, such as **My header**.

CREATE

CANCEL

3. Fill out the **Create a new header** fields as follows:

- From the **Type** menu, select **Request**, and from the **Action** menu, select **Set**.
- In the **Destination** field, type `http.Fastly-Force-Shield`.
- In the **Source** field, type `"yes"`.
- From the **Ignore if set** menu, select **No**.
- In the **Priority** field, type `10`.
- In the **Description** field, type `Fastly Internal Shielding` (or any meaningful, preferred name).

4. Click the **Create** button. The new header appears on the Content page.

5. Click the **Create header** button to create another header to switch to the next backend. The Create a new header page appears.

## Create a new header

Learn more about this section in our [headers tutorial](#).

CONDITION

This will happen all the time unless you [Attach a condition](#)

Type / Action

Request

Set

The type of header and the action performed on it.

Destination

backend

★ Required

The name of the header that will be affected by the selected action. For example: **http.Content-Type**, **http.Set-Cookie**, **http.Via**, **http.Location**, or **http.Access-Control-Allow-Origin**.

Source

Second\_Backend

★ Required

New content for the header. Can be a static value (e.g. string or number) or a dynamic value (e.g. existing header or a GeolIP value). Please use quotes for string values.

Ignore if set

No

If switched to **Yes**, the action will not be performed if the header in Destination exists.

Priority

11

★ Required

The order in which the header rules execute within the condition.

Description

Second Backend

★ Required

The name of your header, such as **My header**.

CREATE

CANCEL

6. Fill out the **Create a new header** fields as follows:

- From the **Type** menu, select **Request**, and from the **Action** menu, select **Set**.
- In the **Destination** field, type `backend`.
- In the **Source** field, type `Second_Backend` (this should match the name of your other backend).
- From the **Ignore if set** menu, select **No**.
- In the **Priority** field, type `11`.
- In the **Description** field, type `Second Backend` (or any meaningful, preferred name).

7. Click the **Create** button. The new header appears on the Content page.

## Create new header conditions

Follow these steps to create conditions for the headers.

1. On the **Content** page, click the **Attach a condition** link next to one of the headers you just created. The Create a new request condition window appears.

Create a new request condition

Learn the basics in our [conditions tutorial](#)

Type

Request

Learn more about the [different types of conditions](#).

Name

Req.restart

\*

Apply if...

req.restarts == 1

The expression to decide whether this is run.

[Examples](#)

[Adjust priority \(default is 10\)](#)

This is an advanced option. For most configurations, the default is best.

SAVE AND APPLY TO FASTLY INTERNAL SHIELDING

CANCEL

2. Fill out the **Create a new request condition** fields as follows:

- Leave the **Type** menu set to **Request**.
- In the **Name** field, type `Req.request` (or any meaningful, preferred name).
- In the **Apply if** field, type `req.restarts == 1`.

3. Click **Save and apply to**. The condition appears on the Content page.

4. Repeat steps 1-3 for the other header.

5. Click the **Activate** button to deploy your configuration changes.

## § Controlling caching (/guides/performance-tuning/controlling-caching)

### How long Fastly caches content

The maximum amount of time we cache content depends on a number of factors including the TTL (Time To Live) and Grace Period, how often an object gets accessed, and how busy other customers are. Setting TTL and Grace Period to a week, possibly even two weeks should be absolutely fine. For more information about controlling how long Fastly caches your resources, start with our Cache Control Tutorial (/guides/tutorials/cache-control-tutorial).

In general, we will honor any cache-control headers (/guides/about-fastly-services/how-fastlys-cdn-service-works) you send to us from your origin; however, if you do not send any cache-control headers, you will most likely reach the Default TTL limit for your service. You can change this limit on the Configuration page (/guides/about-fastly-services/about-the-web-interface-controls.html#about-the-configure-page).

### Changing caching times for different users

You can change the caching times for different users through Surrogate-Control headers defined by the W3C (<http://www.w3.org/TR/edge-arch/>). If, for example, you wanted Fastly to cache something for a month (clearing with API purges, if necessary) but you also wanted to set a maximum age of a single day for users viewing that object in a browser, then you could return the HTTP header:

```
Surrogate-Control: max-age=2629744
Cache-Control: max-age=86400
```

The Surrogate Control header in this example tells Fastly to cache the object for a maximum of 2629744 seconds (one month). The Cache Control header in this example tells the browser to cache the object for a maximum of 86400 seconds (1 day).

For more information about controlling caching, see our [Cache Control Tutorial \(/guides/tutorials/cache-control-tutorial/\)](/guides/tutorials/cache-control-tutorial/).

## Conditionally preventing pages from caching

To conditionally prevent pages from caching, follow the steps below.

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.
4. Click the **Settings** tab. The Settings page appears.
5. Click the **Create cache setting** button to create a new cache setting. The Create a new cache setting page appears.

### Create a new cache setting

This will override your cache control headers. In most cases, use with conditions applied.

CONDITION

This will happen all the time unless you [Attach a condition](#)

TTL (seconds)

The TTL (Time To Live) entered will set the lifespan of the object within our cache nodes.

Action

Pass (do not cache)

This setting decides [how the request will be handled](#).

Stale TTL (seconds)

Stale TTL (Time To Live) is used by your application to [serve stale content](#). It determines how long to keep stale data after it has expired. Note there's custom VCL required to [complete this setup](#).

Description

★ Required

The name of your cache setting, such as **My cache setting**.

CREATE

CANCEL

6. Create a new cache setting and then click the **Create** button. The new cache setting you created appears on the Settings page.
7. Click the **Attach a condition** link to the right of the newly created cache setting. The Create a new cache condition window appears.

Create a new cache condition

Learn the basics in our [conditions tutorial](#)

Type

Cache

Learn more about the [different types of conditions](#).

Name

Cacheable URLs

\*

Apply if...

req.url !~ "^(cacheable|images|assets)"

The expression to decide whether this is run.

[Examples](#)

[Adjust priority \(default is 10\)](#)

This is an advanced option. For most configurations, the default is best.

SAVE AND APPLY TO FORCE PASS

CANCEL

8. Create a condition that matches the URLs you want and then click the **Save and apply to** button. In this example, we set the condition to look for URLs containing `/cacheable`, `/images`, or `/assets`. If the condition finds them, the URLs should be cached. If the condition doesn't find them, the URLs are explicitly not cached by the apply if statement shown above.

9. Click the **Activate** button to deploy your configuration changes.

★ **TIP:** You can use these steps to [override default caching based on a backend response](#) (</guides/basic-configuration/overriding-caching-defaults-based-on-backend-responses>).

## Caching action descriptions

You can use actions to tell Fastly what to do with cached objects and what to do with additional cache configurations as a result. The following actions are available:

- **Do nothing now** - Only set the TTL or stale TTL.
- **Deliver** - Deliver the object to the client. Usually returned from `vcl_fetch`.
- **Pass** - Pass the request and subsequent response to and from the origin server without caching the object. Usually returned from `vcl_recv`.
- **Restart** - Restart the request processing for the object. You can restart the processing of the whole transaction. Changes to the `req` object are retained.

## § Enabling cross-origin resource sharing (CORS)

### (</guides/performance-tuning/enabling-cross-origin-resource-sharing>)

We recommend enabling CORS (Cross-Origin Resource Sharing (<http://docs.aws.amazon.com/AmazonS3/latest/dev/cors.html>)) when using Amazon S3 (</guides/integrations/amazon-s3>) as your backend server. To enable CORS, set up a custom HTTP header for your service by following the steps below.

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.

3. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.
4. Click the **Content** tab. The Content page appears.

The screenshot displays the Fastly Content management interface. At the top, a navigation bar contains six tabs: DOMAINS (1 Domain), ORIGINS (1 Host, 1 Health check), SETTINGS (1 Request setting, 0 Cache settings), CONTENT (0 Headers, 0 Gzips, 0 Responses), LOGGING (1 Endpoint), and CUSTOM VCL (0 Files). The CONTENT tab is selected and highlighted in blue. Below the navigation bar, the main content area is divided into three sections: Headers, Gzip, and Responses. Each section has a heading, a message indicating no items are present, and a button to create the first item. On the right side, there is an 'About content' section with explanatory text.

| DOMAINS  | ORIGINS                  | SETTINGS                              | CONTENT                             | LOGGING    | CUSTOM VCL |
|----------|--------------------------|---------------------------------------|-------------------------------------|------------|------------|
| 1 Domain | 1 Host<br>1 Health check | 1 Request setting<br>0 Cache settings | 0 Headers<br>0 Gzips<br>0 Responses | 1 Endpoint | 0 Files    |

### Headers

*There are no headers.*

**CREATE YOUR FIRST HEADER**

### Gzip

*Gzip is not enabled.*

**SET UP GZIP**

### Responses

*There are no responses.*

**CREATE YOUR FIRST RESPONSE**

### About content

Headers allow you to set, delete, and change request and response headers as your information is handled by Fastly.

Responses allow you to create synthetic responses that exist entirely on the cache servers.

5. Click the **Create header** button. The Create a new header page appears.

## Create a new header

Learn more about this section in our [headers tutorial](#).

### CONDITION

This will happen all the time unless you [Attach a condition](#)

Type / Action

The type of header and the action performed on it.

Destination  ★ Required

The name of the header that will be affected by the selected action. For example: **http.Content-Type**, **http.Set-Cookie**, **http.Via**, **http.Location**, or **http.Access-Control-Allow-Origin**.

Source  ★ Required

New content for the header. Can be a static value (e.g. string or number) or a dynamic value (e.g. existing header or a GeolIP value). Please use quotes for string values.

Ignore if set

If switched to **Yes**, the action will not be performed if the header in Destination exists.

Priority  ★ Required

The order in which the header rules execute within the condition.

Description  ★ Required

The name of your header, such as **My header**.

CREATE

CANCEL

6. Fill out the **Create a new header** fields as follows:

- From the **Type** menu, select **Cache**, and from the **Action** menu, select **Set**.
- In the **Destination** field, type `http.Access-Control-Allow-Origin`.
- In the **Source** field, type `***`.
- Leave the **Ignore if set** menu and the **Priority** field set to their default values.
- In the **Description** field, type a descriptive name for the new header (e.g., `CORS S3 Allow`). This name is displayed in the Fastly web interface.

7. Click the **Create** button. The new header appears on the Content page.

8. Click the **Activate** button to deploy your configuration changes.

**❗ IMPORTANT:** Objects already cached won't have this header applied until you [purge them \(/guides/purging/\)](/guides/purging/).

## Test it out

Running the command `curl -I your-hostname.com/path/to/resource` should include similar information to the following in your header:

```
Access-Control-Allow-Origin: http://your-hostname.tld
Access-Control-Allow-Methods: GET
Access-Control-Expose-Headers: Content-Length, Connection, Date...
```

## § Enabling global POPs (/guides/performance-tuning/enabling-global-pops)

The sun never sets on the Fastly empire, but how can you take full advantage? Simply set your CNAME record (/guides/basic-setup/adding-cname-records) to `nonssl.global.fastly.net` for non TLS traffic. You'll now have access to all of our worldwide POPs (<https://www.fastly.com/network-map>) as they come online. We don't restrict POP access. Instead, you control it.

### How to check if your CNAME is set to nonssl.global.fastly.net

Run the following command in your terminal:

```
$ dig www.example.com +short
```

Your output should appear similar to the following:

```
nonssl.global.fastly.net.  
151.101.117.57
```

If you don't see `nonssl.global.fastly.net.` in your output, then your CNAME isn't properly set. We link to instructions for setting your CNAME (/guides/basic-setup/adding-cname-records) for a number of popular providers.

Instead of using the above command in your terminal, you can also use various online DNS checking tools, such as the OpenDNS Cache Check (<https://cachecheck.opendns.com/>).

### Limiting POP use to the United States and European Union

If you're not using TLS, simply set your CNAME record (/guides/basic-setup/adding-cname-records) to `nossl.us-eu.fastly.net.` instead of `nonssl.global.fastly.net.`. If you're using TLS, see our guide on CNAME records (/guides/basic-setup/adding-cname-records) to find the appropriate entry.

## § Failover configuration (/guides/performance-tuning/failover-configuration)

This guide describes how to configure a failover origin server. A failover (backup) server ensures you can maintain availability of your content if your primary server is not available.

### Before you begin

Before you configure failover, keep in mind the following:

- To configure a failover origin server you must make sure you have health checks (/guides/basic-configuration/health-checks-tutorial) configured for your primary server. If you configure your failover server but don't configure health checks (/guides/basic-configuration/health-checks-tutorial) on the primary server, the failover won't work properly if your primary server stops responding.
- Many customers configure load balancing at the same time they configure failover functionality. Our guide on configuring load balancing (/guides/performance-tuning/load-balancing-configuration) can show you how.

### Configuring a failover origin server

Once you've confirmed health checks are configured, you must:

1. Turn off automatic load balancing on both the primary origin server and the server that will become your failover.
2. Create headers that configure both the primary and failover origin servers.
3. Create a header condition that specifies exactly when to use the failover server.

### Turn off automatic load balancing

To configure a failover origin server you must turn off automatic load balancing for both the server that will act as your primary origin server and the server that will become your failover origin server.

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.



- Click the **Origins** tab. The Origins page appears.
- Click the name of the origin server you want to configure. The Edit a host page appears.

## Edit a host

CONDITION

This will happen all the time unless you [Attach a condition](#)

Address

:

★ Required

The IP address (or hostname) and port number for your origin server. To [enable TLS](#), use port 443.

Shielding

The shield POP designated to reduce inbound load on your origins by serving cached data. [Learn more about POP request handling](#) and the [caveats of shielding](#).

Health check

[+ Create a new health check](#)

A small request assigned to [test the state of your origin](#).

Auto load balance

Enables Fastly to [load balance between two or more backend hosts](#).

[▶ Tips](#)

Description

★ Required

The name of your origin, such as [My origin server](#).

- From the **Auto load balance** menu, select **No**.
- Click the **Update** button to apply the changes.

## Configure the primary and failover origin servers

Once you've turned off automatic load balancing, create two new request headers, one each for your primary and failover servers.

- Click the **Content** tab. The Content page appears.
- Click the **Create header** button to create the first request header. The Create a new header window appears.

## Create a new header

Learn more about this section in our [headers tutorial](#).

### CONDITION

This will happen all the time unless you [Attach a condition](#)

Type / Action

Request

Set

The type of header and the action performed on it.

Destination

backend

★ Required

The name of the header that will be affected by the selected action. For example: **http.Content-Type**, **http.Set-Cookie**, **http.Via**, **http.Location**, or **http.Access-Control-Allow-Origin**.

Source

F\_global

★ Required

New content for the header. Can be a static value (e.g. string or number) or a dynamic value (e.g. existing header or a GeoIP value). Please use quotes for string values.

Ignore if set

No

If switched to **Yes**, the action will not be performed if the header in Destination exists.

Priority

10

★ Required

The order in which the header rules execute within the condition.

Description

Set default origin

★ Required

The name of your header, such as **My header**.

CREATE

CANCEL

3. Fill out the **Create a new header** fields as follows:

- From the **Type** menu, select **Request**, and from the **Action** menu, select **Set**.
- In the **Destination** field, type the name of the header that will be affected by the selected action.
- In the **Source** field, type where the new content for the header comes from.
- Leave the **Ignore if set** and **Priority** controls at their default settings.
- In the **Description** field, type a descriptive name for the header. This name is displayed in the Fastly web interface.

4. Click the **Create** button to create the first header. A new header appears on the Content page.

5. Click the **Create header** button to create a second request header. The Create a new header window appears.

## Create a new header

Learn more about this section in our [headers tutorial](#).

CONDITION

This will happen all the time unless you [Attach a condition](#)

Type / Action

Request

Set

The type of header and the action performed on it.

Destination

backend

★ Required

The name of the header that will be affected by the selected action. For example: **http.Content-Type**, **http.Set-Cookie**, **http.Via**, **http.Location**, or **http.Access-Control-Allow-Origin**.

Source

F\_mirror

★ Required

New content for the header. Can be a static value (e.g. string or number) or a dynamic value (e.g. existing header or a GeoIP value). Please use quotes for string values.

Ignore if set

No

If switched to **Yes**, the action will not be performed if the header in Destination exists.

Priority

11

★ Required

The order in which the header rules execute within the condition.

Description

set failover origin

★ Required

The name of your header, such as **My header**.

CREATE

CANCEL

6. Fill out the **Create a new header** fields as follows:

- From the **Type** menu, select **Request**, and from the **Action** menu, select **Set**.
- In the **Destination** field, type the name of the header that will be affected by the selected action.
- In the **Source** field, type where the new content for the header comes from.
- Leave the **Ignore if set** control at the default setting.
- In the **Priority** field, type a number at least one higher than the priority you set on the primary server's request header. For example, if you left the first header's priority set to the default, **10**, you would set the second header's priority to **11** or higher.
- In the **Description** field, type a descriptive name for the header. This name is displayed in the Fastly web interface.

7. Click the **Create** button to create the second header. A new header appears on the Content page.

## Specify when to use the failover server

Once you've configured your primary and failover servers, create an associated header condition that specifies exactly when the failover server should be used.

- On the **Content** page, click the **Attach a condition** link next to the new header you just created for the failover origin server. The Create a new request condition window appears.

Create a new request condition

Learn the basics in our [conditions tutorial](#)

Type

Request

Learn more about the [different types of conditions](#).

Name

Primary Down

\*

Apply if...

req.restarts > 0 || !req.backend.healthy

The expression to decide whether this is run.

▸ Examples

Priority

11

Determines ordering for multiple conditions. Lower numbers execute first.

[Reset to default \(priority 10\)](#)

SAVE AND APPLY TO SET FAILOVER ORIGIN

CANCEL

2. Fill out the **Create a new request condition** fields as follows:

- From the **Type** menu, select **Request**.
- In the **Name** field, type a descriptive name for the new condition (for example, `Primary Down`).
- In the **Apply if** field, type the appropriate request condition that will be applied. For example, `req.restarts > 0 || !req.backend.healthy` would tell the system only to use the failover server if the number of restarts is more than 0 or the origin is unhealthy.
- In the **Priority** field, type `11`.

3. Click the **Save and apply to** button to create the new condition for the header.

4. Click the **Activate** button to deploy your configuration changes.

## § Generating HTTP redirects at the edge (/guides/performance-tuning/generating-http-redirects-at-the-edge)

When users request information from your origin servers, you may want to redirect them for various reasons. For example, you may want to redirect them to either a secure or non-secure version of a web page, or to pages that have been moved or updated since the last time they were requested.

You can send these redirects from the edge rather than having to go to origin by creating a synthetic response with the appropriate redirect status code and then creating a content rule with the proper Location header.

### Create a new response and condition

To generate redirects at the edge, start by creating a new response with the appropriate status code and a new condition describing when the response can be applied.

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.

4. Click the **Content** tab. The Content page appears.
5. Click the **Create response** button. The Create a new synthetic response page appears.

## Create a new synthetic response

More on how synthetic responses work in our [tutorial](#).

CONDITION

This will happen all the time unless you [Attach a condition](#)

Status

301 Moved Permanently

The HTTP status code to include in the header of the response.

MIME Type

The media type to put into the Content-Type header which informs the browser how to display the response. Common MIME types are `application/json`, `text/plain`, `text/html`.

Response

The content to be served when delivering the response.

Description

Redirect to blog

★ Required

The name of your response, such as `My response`.

CREATE

CANCEL

6. Fill out the **Create a new synthetic response** fields as follows:

- From the **Status** menu, select the HTTP status code that should be included in the header of the response (e.g., `301 Moved Permanently` or `302 Moved Temporarily` for redirections).
- Leave the **MIME Type** field blank.
- In the **Description** field, type a meaningful name for your response (e.g., `Redirect to blog`). This name is displayed in the Fastly web interface.

7. Click the **Create** button to create the new response.
8. On the **Content** page, click the **Attach a condition** link to the right of the new response you just created. The Create a new request condition window appears.

Create a new condition

Learn the basics in our [conditions tutorial](#)

Type

Request

Learn more about the [different types of conditions](#).

Name

URL is /wordpress

\*

Apply if...

req.url ~ "^/wordpress"

The expression to decide whether this is run.

[Examples](#)

[Adjust priority \(default is 10\)](#)

This is an advanced option. For most configurations, the default is best.

SAVE AND APPLY TO REDIRECT TO BLOG

CANCEL

9. Fill out the **Create a new request condition** fields as follows:

- From the **Type** menu, select **Request**.
- In the **Name** field, type a meaningful name for your condition (e.g., `URL is /wordpress`). This name is displayed in the Fastly web interface.
- In the **Apply if** field, type the logical expression to execute in VCL to determine if the condition resolves as True or False (e.g., `req.url ~ "^/wordpress"`).
- Leave the **Priority** set to 10.

10. Click the **Save and apply to** button to create the new request condition.

## Create a new header and condition

Complete the creation of a synthetic redirect by creating a new header and condition that modifies that response by adding the location header based on the status code and the matching URL. This ensures the redirect only applies when both of those are true.

1. On the **Content** page, click the **Create header** button. The Create a new header page appears.

## Create a new header

Learn more about this section in our [headers tutorial](#).

**CONDITION** This will happen all the time unless you [Attach a condition](#)

|   |   |            |
|---|---|------------|
| Type / Action   | Response                                      | Set        |
| The type of header and the action performed on it.  |   |            |
| Destination   | http.Location                                 | ★ Required |
| The name of the header that will be affected by the selected action. For example: <code>http.Content-Type</code> , <code>http.Set-Cookie</code> , <code>http.Via</code> , <code>http.Location</code> , or <code>http.Access-Control-Allow-Origin</code> . |   |            |
| Source  | "http://www.example.com/new-location/of/item" | ★ Required |
| New content for the header. Can be a static value (e.g. string or number) or a dynamic value (e.g. existing header or a GeoIP value). Please use quotes for string values.  |   |            |
| Ignore if set   | No  |            |
| If switched to <b>Yes</b> , the action will not be performed if the header in Destination exists.   |   |            |
| Priority  | 10  | ★ Required |
| The order in which the header rules execute within the condition.   |   |            |
| Description   | Location for WordPress redirect               | ★ Required |
| The name of your header, such as <b>My header</b> .   |   |            |

CREATE

CANCEL

2. Fill out the **Create a new header** fields as follows:

- From the **Type** menu, select **Response**, and from the **Action** menu, select **Set**.
- In the **Destination** field, type `http.Location`.
- In the **Source** field, type the source location of the new content (e.g., `"http://www.example.com/new-location/of/item"`).
- Leave the **Ignore if set** and **Priority** fields at their default settings.
- In the **Description** field, type a meaningful name for your header (e.g., `Location for WordPress redirect`).

3. Click the **Create** button to create the new header.

4. On the **Content** page, click the **Attach a condition** link to the right of the new header you just created. The Create a new request condition window appears.

Create a new condition.

Learn the basics in our [conditions tutorial](#)

Type

Response

Learn more about the [different types of conditions](#).

Name

Set location for blog redirect

\*

Apply if...

req.url ~ "^/wordpress" && resp.status == 301

The expression to decide whether this is run.

[Examples](#)

[Adjust priority \(default is 10\)](#)

This is an advanced option. For most configurations, the default is best.

SAVE AND APPLY TO THIS HEADER

CANCEL

5. Fill out the **Create a new request condition** fields as follows:

- From the **Type** menu, select **Response**.
- In the **Name** field, type a meaningful name for your condition (e.g., `Set location for blog redirect`).
- In the **Apply if** field, type the logical expression to execute in VCL to determine if the condition resolves as true or false (e.g., `req.url ~ "^/wordpress" && resp.status == 301`). The `resp.status` needs to match the response code generated in the response above.
- Leave the **Priority** set to 10.

6. Click the **Save and apply to** button to create the new condition.

7. Click the **Activate** button to deploy your configuration changes.

**NOTE:** These responses use a custom status number >500. They will appear as errors on the [Dashboard page \(/guides/about-fastly-services/about-the-web-interface-controls.html#about-the-dashboard-page\)](#) even though they are desired behavior.

## § HTTP/2 server push (/guides/performance-tuning/http2-server-push)

Fastly recognizes `link` headers with the preload keyword (<https://w3c.github.io/preload/>) sent by an origin server and pushes the designated resource to a client. For example, this `link` response header triggers an HTTP/2 push:

```
link: </assets/jquery.js>; rel=preload
```

We support multiple `link` headers and multiple assets in one `link` header:

```
link: </assets/jquery.js>; rel=preload, </assets/base.css>; rel=preload
```

## Link headers and Amazon S3 buckets

If you're using an Amazon Simple Storage Service (S3) (</guides/integrations/amazon-s3>) bucket as your origin server, you can still use `link` headers by applying a cache setting condition (</guides/performance-tuning/controlling-caching>) like this one:

```
set beresp.http.Link = beresp.http.x-amz-meta-Link
```



# § Improving caching performance with large files (/guides/performance-tuning/improving-caching-performance-with-large-files)

Fastly provides two features to enhance performance specifically for large files up to 5GB: Streaming Miss and Large File Support.

## Streaming Miss

When fetching an object from the origin, Streaming Miss ensures the response is streamed back to the client immediately and is written to cache only after the whole object has been fetched. This reduces the first-byte latency, which is the time that the client must wait before it starts receiving the response body. The larger the response body, the more pronounced the benefit of streaming.

### NOTE:

If you enable Streaming Miss, be aware that if an error occurs while transferring the response body, Fastly cannot send an error because the headers are already sent to the client. All we can do is close the connection, truncating the response.

## Configuration

Configuration is simple. In VCL, simply set `beresp.do_stream` to true in `vcl_fetch`:

```
sub vcl_fetch {  
  ...  
  set beresp.do_stream = true;  
  ...  
  return(deliver);  
}
```

The same can be achieved by creating a new header of type `Cache`, action `Set`, destination `do_stream` and source `true` (this can, of course, be controlled with conditions (/guides/conditions/)).

## Create a new header

Learn more about this section in our [headers tutorial](#).

**CONDITION** This will happen all the time unless you [Attach a condition](#)

|   |                       |            |
|---|-----------------------|------------|
| Type / Action   | Cache                 | Set        |
| The type of header and the action performed on it.  |                       |            |
| Destination   | do_stream             | ★ Required |
| The name of the header that will be affected by the selected action. For example: <code>http.Content-Type</code> , <code>http.Set-Cookie</code> , <code>http.Via</code> , <code>http.Location</code> , or <code>http.Access-Control-Allow-Origin</code> . |                       |            |
| Source  | true                  | ★ Required |
| New content for the header. Can be a static value (e.g. string or number) or a dynamic value (e.g. existing header or a GeoIP value). Please use quotes for string values.  |                       |            |
| Ignore if set   | No                    |            |
| If switched to <b>Yes</b> , the action will not be performed if the header in Destination exists.   |                       |            |
| Priority  | 10                    | ★ Required |
| The order in which the header rules execute within the condition.   |                       |            |
| Description   | Enable Streaming Miss | ★ Required |
| The name of your header, such as <b>My header</b> .   |                       |            |

CREATE

CANCEL

## Limitations

There are several limitations to using Streaming Miss.

### Origins cannot use TLS

Streaming Miss does not currently support HTTPS (TLS) origin servers. The content can be served to the client over HTTPS, but it cannot be fetched with Streaming Miss over HTTPS. Objects fetched from HTTPS origins are therefore limited to the non-Streaming Miss size of 2GB.

### Streaming Miss is not available to HTTP/1.0 clients

If an HTTP/1.0 request triggers a fetch, and the response header from the origin does not contain a Content-Length field, then Streaming Miss will be disabled for the fetch and the fetched object will be subject to the non-streaming-miss object size limit.

If an HTTP/1.0 request is received while a Streaming Miss for an object is in progress, the HTTP/1.0 request will wait for the response body to be downloaded before it will receive the response header and the response body, as if the object was being fetched without Streaming Miss.

Cache hits are not affected. An HTTP/1.0 client can receive a large object served from cache, just like an HTTP/1.1 client.

### Streaming Miss is not compatible with on-the-fly gzip compressing of the fetched object

Streaming Miss can handle large files whether or not they are compressed. However, on-the-fly compression of objects that are not already compressed is not compatible with Streaming Miss. If the VCL sets `beresp.gzip` to true, Streaming Miss will be disabled.

### Streaming Miss is not compatible with ESI (Edge-Side Includes)

Responses that are processed through ESI cannot be streamed. Responses that are included from an ESI template cannot be streamed. When ESI is enabled for the response or when the response is fetched using `<esi:include>`, then Streaming Miss will be disabled and the fetched object will be subject to the non-streaming-miss object size limit of 2GB.

# Large File Support

Large File Support is automatically enabled for all clients — there's no need to manually configure anything. You should, however, be aware that there are maximum file sizes and several failure modes.

## Maximum file size

If Streaming Miss is enabled, then the maximum size is slightly below 5GB (specifically 5,368,578,048 bytes). With Streaming Miss disabled, the maximum size is still higher than the previous maximum, but is limited to a little under 2GB (specifically 2,147,352,576 bytes).

## Failure modes

There are several failure modes you may encounter while using Large File Support.

### What happens when the maximum object size limit is exceeded?

If the response from the origin has a Content-Length header field which exceeds the maximum object size, Fastly will immediately generate a 503 response to the client unless specific VCL is put in place to act on the error.

If no Content-Length header field is returned, Fastly will start to fetch the response body. If while fetching the response body we determine that the object exceeds maximum object size, we will generate a status 503 response to the client (again, unless specific VCL is in place to act on the error).

If no Content-Length header field is present and Streaming Miss is in effect, Fastly will stream the content back to the client. However, if while streaming the response body Fastly determines that the object exceeds the maximum object size, it will terminate the client connection abruptly. The client will detect a protocol violation, because it will see its connection close without a properly terminating 0-length chunk.

### What happens when an origin read fails?

A failure to read the response header from the origin, regardless of Streaming Miss, causes a 503 response (which can be acted on in VCL).

If reading the response body from the origin fails or times out, the problem will be reported differently depending on whether Streaming Miss is in effect for the fetch. Without Streaming Miss, a 503 response will be generated. With Streaming Miss, however, it is already too late to send an error response since the header will already have been sent. In this case, Fastly will again abruptly terminate the client connection and the client will detect a protocol violation. If the response was chunked, it will see its connection close without a properly terminating 0-length chunk. If Content-Length was known, it will see the connection close before the number of bytes given.

Incidentally, this is the reason why HTTP/1.0 clients cannot be supported by Streaming Miss in the cases when the Content-Length is not yet known or available. Without the client receiving a Content-Length and without support for chunking, the client cannot distinguish the proper end of the download from an abrupt connection breakage anywhere upstream from it.

---

## § Load-balancing configuration (/guides/performance-tuning/load-balancing-configuration)

This guide describes how to automatically load balance between two or more origin servers. Load balancing distributes requests across multiple servers to optimize resource use and avoid overloading any single resource.

## Before you begin

Before you configure load balancing, keep in mind the following:

- To prevent errors when shielding is enabled (/guides/performance-tuning/shielding#enabling-shielding), all backends in the automatic load balancing group must use the same shielding location.
- Conditions on your origin server can directly change how automatic load balancing behaves. Be sure to review conditions behavior to ensure automatic load balancing works properly.
- Many customers configure failover at the same time they configure load balancing functionality. Our guide on configuring failover (/guides/performance-tuning/failover-configuration) can show you how.

## Setting up load balancing between origins

If you want to randomly load balance between two or more origin servers, enable the Fastly automatic load balancing feature.

## Enabling load balancing

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.
4. Click the **Origins** tab. The Origins page appears.

5. Click the name of the host you want to edit. The Edit a host page appears.

## Edit a host

CONDITION

This will happen all the time unless you [Attach a condition](#)

Address

:

★ Required

The IP address (or hostname) and port number for your origin server. To [enable TLS](#), use port 443.

Shielding

The shield POP designated to reduce inbound load on your origins by serving cached data. Learn more about [POP request handling](#) and the [caveats of shielding](#).

Health check

[+ Create a new health check](#)

A small request assigned to [test the state of your origin](#).

Auto load balance

Enables Fastly to [load balance between two or more backend hosts](#).

[▶ Tips](#)

6. From the **Auto load balance** menu, select **Yes**.
7. Click the **Update** button.
8. Repeat steps 5, 6, and 7 for each origin server you want to include in the automatic load balancing group.
9. Click the **Activate** button to deploy your configuration changes.

## How Weight affects load balancing

The Weight setting on the Create a new host or Edit a host page allows you to specify the percentage of the total traffic that is sent to specific origin servers. This setting only appears on the pages when you set Auto load balance to Yes.

## Create a new host

**CONDITION** This will happen all the time unless you [Attach a condition](#)

Address  :  ★ Required

The IP address (or hostname) and port number for your origin server. To [enable TLS](#), use port 443.

Shielding

The shield POP designated to reduce inbound load on your origins by serving cached data. Learn more about [POP request handling](#) and the [caveats of shielding](#).

Health check

[+ Create a new health check](#)

A small request assigned to [test the state of your origin](#).

Auto load balance

Enables Fastly to [load balance](#) between two or more backend hosts.

[► Tips](#)

Weight

A portion of traffic to send to a specific origin. Each origin receives (weight ÷ total weight across origins) of the traffic.

When you specify a whole number in the Weight field, you specify the percentage of the total traffic to send to a specific origin server. Each origin server receives the percentage ( $\frac{\text{weight}}{\text{total}}$ ) of the total traffic equal to the number you specify.

For example, if you have two origin servers, A and B, setting the Weight field to 50 on both splits the traffic between them equally. Each origin server receives 50 percent of your total traffic.

If you increase the Weight setting on origin server A to 55 and decrease the Weight setting on origin server B to 45, the percentage of traffic changes to 55 percent and 45 percent respectively.

### How conditions affect load balancing

When you set conditions (</guides/conditions/>) on origin servers you can potentially change how automatic load balancing works. The load balancing autodirector groups servers together based on like conditions, giving you the flexibility to effectively create subsets of the autodirector by assigning a condition to one group of origins and another condition to another set of origins. If each group of origin servers has a different condition that affects load balancing, the auto load function will not randomly load balance between the different servers.

Conditions can also be assigned to a server through a header. For example, you have three servers called F\_Fastly, F\_Second\_backend, and F\_Third\_backend and want all URLs with a certain prefix to default to the second server.

## Edit a header

Learn more about this section in our [headers tutorial](#).

### CONDITION

This will happen all the time unless you [Attach a condition](#)

Type / Action

Request

Set

The type of header and the action performed on it.

Destination

backend

★ Required

The name of the header that will be affected by the selected action. For example: `http.Content-Type`, `http.Set-Cookie`, `http.Via`, `http.Location`, or `http.Access-Control-Allow-Origin`.

Source

F\_Second\_backend

★ Required

New content for the header. Can be a static value (e.g. string or number) or a dynamic value (e.g. existing header or a GeoIP value). Please use quotes for string values.

Ignore if set

No

If switched to **Yes**, the action will not be performed if the header in Destination exists.

Priority

10

★ Required

The order in which the header rules execute within the condition.

Description

Media Requests

★ Required

The name of your header, such as **My header**.

UPDATE

CANCEL

After the header is created, click the **Attach a condition** link to the right of the header name, then add a new condition to apply if the URL matches the desired prefix.

Create a new request condition

Learn the basics in our [conditions tutorial](#)

Type

Request

Learn more about the [different types of conditions](#).

Name

Media Files

Apply if...

req.url-"/media"

The expression to decide whether this is run.

[Examples](#)

[Adjust priority \(default is 10\)](#)

This is an advanced option. For most configurations, the default is best.

SAVE AND APPLY TO MEDIA REQUESTS

CANCEL

The generated VCL below illustrates the autodirector set for all three servers. Within the section **sub vcl\_recv**, the default origin server is set to the autodirector and, if the media condition is met, requests are forwarded to the second server.

```
director autodirector_ random {
{
.backend = F_Second_backend;
.weight = 100;
}{
.backend = F_Third_backend;
.weight = 100;
}{
.backend = F_Fastly;
.weight = 100;
}
}

sub vcl_recv {
#--FASTLY RECV CODE START
if (req.restarts == 0) {
if (!req.http.X-Timer) {
set req.http.X-Timer = "S" time.start.sec "." time.start.usec_frac;
}
set req.http.X-Timer = req.http.X-Timer ",VS0";
}

# default conditions
set req.backend = autodirector_;

# end default conditions

# Request Condition: Media files Prio: 10
if (req.url ~ "/media") {

# Header rewrite Media Requests : 10
set req.backend = F_Second_backend;
}
#end condition

#--FASTLY RECV CODE END
}
```

## § Maintaining separate HTTP and HTTPS requests to origin servers (/guides/performance-tuning/maintaining-separate-http-and-https-requests-to-backend-servers)

It is common to use the same origin web application to serve both HTTP and HTTPS requests and let the application determine which actions to take to secure communications (/guides/securing-communications/) depending on the incoming protocol. Fastly allows users to set this up to preserve this functionality within their servers. To set Fastly up to send HTTP requests to the non-secure service and HTTPS requests to the secure service, configure two origins, one each for the secure and non-secure ports, then set up the conditions under which requests will be sent there.

### Create multiple origins

Begin by configuring the same origin address with a different port as a separate origin by following the steps below.

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.
4. Click the **Origins** tab. The Origins page appears.
5. Click the **Create host** button to create a new backend for your non-secure server. The Create a new host page appears.

## Create a new host

**CONDITION** This will happen all the time unless you [Attach a condition](#)

**Address**  :  ★ Required

The IP address (or hostname) and port number for your origin server. To [enable TLS](#), use port 443.

**Shielding** (none) ⌵

The shield POP designated to reduce inbound load on your origins by serving cached data. Learn more about [POP request handling](#) and the [caveats of shielding](#).

**Health check** (none) ⌵

+ [Create a new health check](#)

A small request assigned to [test the state of your origin](#).

**Auto load balance** No ⌵

Enables Fastly to [load balance between two or more backend hosts](#).

▶ [Tips](#)

**Description**  ★ Required

The name of your origin, such as [My origin server](#).

6. Fill out the **Create a new host** fields as follows:
  - In the **Address** field, type the address of your non-secure server (for example, `a-server.my.com`), and in the port field, type `80`.
  - Leave the **Shielding**, **Health check**, and **Auto load balance** controls set to their default values.
  - In the **Description** field, type the name of your non-secure server (for example, `Server Name (plain)`).
7. Click the **Create** button. The non-secure HTTP server appears on the Origins page.
8. Click the **Create host** button a second time to create a second backend, this time for your secure server.



## Create a new host

CONDITION

This will happen all the time unless you [Attach a condition](#)

Address

: 443

★ Required

The IP address (or hostname) and port number for your origin server. To [enable TLS](#), use port 443.

Shielding

The shield POP designated to reduce inbound load on your origins by serving cached data. Learn more about [POP request handling](#) and the [caveats of shielding](#).

Health check

[+ Create a new health check](#)

A small request assigned to [test the state of your origin](#).

Auto load balance

Enables Fastly to [load balance between two or more backend hosts](#).

[▶ Tips](#)

Description

★ Required

The name of your origin, such as [My origin server](#).

9. Fill out the **Create a new host** fields as follows:

- In the **Address** field, type the address of your secure server (for example, `a-server.my.com`), and in the port field, type `443`.
- Leave the **Shielding**, **Health check**, and **Auto load balance** controls set to their default values.
- In the **Description** field, type the name of your secure server (for example, `Server Name (secure)`).

10. Click the **Create** button. The secure HTTP server appears in the Backends area.

## Conditionally send traffic to origins

To conditionally determine which server receives secure and non-secure requests, Fastly relies on the presence or absence of a specific header when the backend is selected. When an incoming connection is received over TLS, Fastly sets the `req.http.fastly-ssl` header to determine which server to use.

Set a condition for this header on each origin by following the steps below.

1. On the **Origins** page, click the **Attach a condition** link next to the name of the non-secure server. The Create a new request condition window appears.

### Create a new request condition

Learn the basics in our [conditions tutorial](#)

Type

Learn more about the [different types of conditions](#).

Name  \*

Apply if...

The expression to decide whether this is run.

▸ [Examples](#)

[Adjust priority \(default is 10\)](#)

This is an advanced option. For most configurations, the default is best.

**SAVE AND APPLY TO 1.1.1.1** CANCEL

2. Fill out the **Create a new request** fields as follows:

- Leave the **Type** menu set to **Request**.
- In the **Name** field, type the name of the condition specifying use of the non-secure server (for example, `Use non-secure`).
- In the **Apply if** field, type `!req.http.fastly-ssl`.
- Leave the priority set to its default value.

3. Click the **Save and apply to** button to create the new condition.

4. On the **Origins** page, click the **Attach a condition** link next to the name of the secure server. The Create a new request condition window appears.

Create a new request condition

Learn the basics in our [conditions tutorial](#)

Type

Request

Learn more about the [different types of conditions](#).

Name

Use secure

\*

Apply if...

req.http.fastly-ssl

The expression to decide whether this is run.

[Examples](#)

[Adjust priority \(default is 10\)](#)

This is an advanced option. For most configurations, the default is best.

SAVE AND APPLY TO 1.1.1.1

CANCEL

5. Fill out the **Create a new request condition** window as follows:

- Leave the **Type** menu set to **Request**.
- In the **Name** field, type the name of the condition specifying use of the secure server (for example, `Use secure`).
- In the **Apply if** field, type `req.http.fastly-ssl`.
- Leave the priority set to its default value.

6. Click the **Save and apply to** button to create the new condition.

7. Click the **Activate** button to deploy your configuration changes.

## § Making query strings agnostic (/guides/performance-tuning/making-query-strings-agnostic)

Under normal circumstances, Fastly would consider these URLs different objects that are cached separately:

- `http://example.com`
- `http://example.com?asdf=asdf`
- `http://example.com?asdf=zxcv`

It is possible, however, to have them all ignore the query string and return the same cached file.

1. Log in the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.
4. Click the **Content** tab. The Content page appears.
5. Click the **Create header** button. The Create a new header window appears.

## Create a new header

Learn more about this section in our [headers tutorial](#).

### CONDITION

This will happen all the time unless you [Attach a condition](#)

Type / Action

The type of header and the action performed on it.

Destination  ★ Required

The name of the header that will be affected by the selected action. For example: **http.Content-Type**, **http.Set-Cookie**, **http.Via**, **http.Location**, or **http.Access-Control-Allow-Origin**.

Source  ★ Required

New content for the header. Can be a static value (e.g. string or number) or a dynamic value (e.g. existing header or a GeolP value). Please use quotes for string values.

Ignore if set

If switched to **Yes**, the action will not be performed if the header in Destination exists.

Priority  ★ Required

The order in which the header rules execute within the condition.

Description  ★ Required

The name of your header, such as **My header**.

CREATE

CANCEL

6. Fill out the **Create a new header** fields as follows:

- From the **Type** menu, select **Request**, and from **Action** menu, select **Set**.
- In the **Destination** field, type `url`.
- In the **Source** field, type `req.url.path`.
- From the **Ignore if set** menu, select **No**.
- Set the **Priority** field to whatever priority you want.
- In the **Description** field, type a description for the header (e.g., `New query string name`).

7. Click the **Create** button to create the new header. The new header you created appears on the Content page.

8. Click the **Activate** button to deploy your configuration changes.

The request will be sent to the origin as a URL without the query string.

For more information about controlling caching, see our [Cache Control Tutorial \(/guides/tutorials/cache-control-tutorial\)](#).

## § Request collapsing (/guides/performance-tuning/request-collapsing)

This guide describes Fastly's Request Collapsing feature, frequently used when creating advanced service configurations.

**NOTE:** This guide requires advanced knowledge of [Varnish and the VCL language \(/guides/vcl/guide-to-vcl\)](/guides/vcl/guide-to-vcl).

## The basics

Request Collapsing causes simultaneous cache misses within a single Fastly datacenter to be "collapsed" into a single request to an origin server. While the single request is being processed by the origin, the other requests wait in a queue for it to complete. Two types of Request Collapsing exist:

1. Collapsing on a single cache server
2. Collapsing within the datacenter between cache servers

Each cache server will automatically queue duplicate requests for the same hash and only allow one request to origin. You can disable this behavior by setting `req.hash_ignore_busy` to `true` in `vcl_recv`.

Within a datacenter, not every cache stores every object. Only two servers in each datacenter will store an object: one as a primary and one as a backup. Only those two servers will fetch the object from origin.

## How it works

In Fastly's version of Varnish, VCL subroutines often run on different caches during a request. For a particular request, both an edge cache and a shield cache will exist (though a single cache can, in some cases, fulfill both of these roles). The edge cache receives the HTTP request from the client and determines via a hash which server in the datacenter is the shield cache. If this cache determines it is the shield cache and has the object in cache, it fulfills both the edge cache and the shield cache roles.

Certain VCL subroutines run on the edge cache and some on the shield cache:

- Edge Cache: `vcl_recv`, `vcl_deliver`, `vcl_log`, `vcl_error`
- Shield Cache: `vcl_miss`, `vcl_hit`, `vcl_pass`, `vcl_fetch`, `vcl_error`

## Determining if a cache is an edge or a shield

The `fastly_info.is_cluster_edge` VCL variable will be true if the cache currently running the VCL is the edge cache and false if it is the shield cache.

## Caveats

Keep in mind the following limitations when using the Request Collapsing feature:

1. Any `req.http.*` headers are not transferred from the shield cache back to the edge cache. Remember this when writing advanced configurations that use headers to keep track of state. If you set a `req.http.*` header in any of the subroutines that run on the shield cache, expect that the change will not persist on the edge cache.
2. A single, slow request to origin can sometimes cause a great many other requests for the same object to hang and fail. Because many requests for a single object are being collapsed down to one, they all succeed or fail based on the request that reaches the origin.

## § Routing assets to different origins (/guides/performance-tuning/routing-assets-to-different-origins)

Some customers have assets stored on multiple origin servers and want to route various requests to specific, different servers based on criteria they supply (e.g., asset type, file directory, host header). Fastly offers customers the ability to set conditions on their origins, which simply adds an if statement block to your VCL.

### Basic setup: Create conditions for each origin

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.
4. Click the **Origins** tab. The Origins page appears.
5. Click the **Attach a condition** link to the right of the name of an origin server. The Create a new request condition window appears.

Create a new request condition

Learn the basics in our [conditions tutorial](#)

Type

Request

Learn more about the [different types of conditions](#).

Name

Hosts

\*

Apply if...

req.http.host ~ "www.example.com"

The expression to decide whether this is run.

Examples

Adjust priority (default is 10)

This is an advanced option. For most configurations, the default is best.

SAVE AND APPLY TO 1.1.1.1

CANCEL

6. Fill out the **Create a new request condition** fields as follows:

- Leave the **Type** menu set to **Request**.
- In the **Name** field, type a human-readable name for the condition.
- In the **Apply if** field, type the conditions that you want to apply to your origin server. For example, for hosts, you could type `req.http.host ~ "www.example.com"`. Or, for content-type / URL, you could type `req.url ~ "\.(jpg|png|gif)(\|\?)"`.

7. Click the **Save and apply to** button. The new condition appears on the Origins page.

8. Click the **Activate** button to deploy your configuration changes.

## Backup setup: Create a header

What if you have a condition already assigned to your origin? Although you can group request conditions on the origin with an 'and' or 'or' clause, there can only ever be one condition rule attached to that origin. If you want to separate your request conditions instead of grouping them, you can use header rules to route assets to different origins instead.

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.
4. Click the **Content** tab. The Content page appears.
5. Click the **Create header** button. The Create a new header page appears.

## Create a new header

Learn more about this section in our [headers tutorial](#).

### CONDITION

This will happen all the time unless you [Attach a condition](#)

Type / Action

Request

Set

The type of header and the action performed on it.

Destination

backend

★ Required

The name of the header that will be affected by the selected action. For example: **http.Content-Type**, **http.Set-Cookie**, **http.Via**, **http.Location**, or **http.Access-Control-Allow-Origin**.

Source

Image\_Backend

★ Required

New content for the header. Can be a static value (e.g. string or number) or a dynamic value (e.g. existing header or a GeoIP value). Please use quotes for string values.

Ignore if set

No

If switched to **Yes**, the action will not be performed if the header in Destination exists.

Priority

10

★ Required

The order in which the header rules execute within the condition.

Description

Image Backend

★ Required

The name of your header, such as **My header**.

CREATE

CANCEL

6. Fill out the **Create a new header** fields as follows:

- From the **Type** menu, select **Request**, and from the **Action** menu, select **Set**.
- In the **Destination** field, type `backend`.
- In the **Source** field, type `Image_Backend`. (This should match the name of your global origin server. You can see the exact name if you look at your VCL. Click on the **VCL** button at the top of the page.)
- From the **Ignore if set** menu, select **No**.
- In the **Priority** field, type `10`.
- In the **Description** field, type `Image Backend` (or any meaningful, preferred name).

7. Click the **Create** button. The new header appears on the Content page.

8. On the **Content** page, click the **Attach a condition** link next to the header you just created. The Create a new request condition window appears.

Create a new request condition

Learn the basics in our [conditions tutorial](#)

Type

Request

Learn more about the [different types of conditions](#).

Name

Redirect Images

\*

Apply if...

req.url ~ "\.(jpg|png|gif)(\$|\?)"

The expression to decide whether this is run.

Examples

Adjust priority (default is 10)

This is an advanced option. For most configurations, the default is best.

SAVE AND APPLY TO IMAGE BACKEND

CANCEL

9. Fill out the **Create a new request condition** fields as follows:

- Leave the **Type** menu set to **Request**.
- In the **Name** field, type `Redirect Images` (or any meaningful, preferred name).
- In the **Apply if** field, type `req.url ~ "\.(jpg|png|gif)($|\?)"`.

10. Click the **Save and apply to** button. The condition appears on the Content page.

11. Click the **Activate** button to deploy your configuration changes.

★ **TIP:** Learn more about conditions in our [Conditions subcategory \(/guides/conditions/\)](/guides/conditions/).

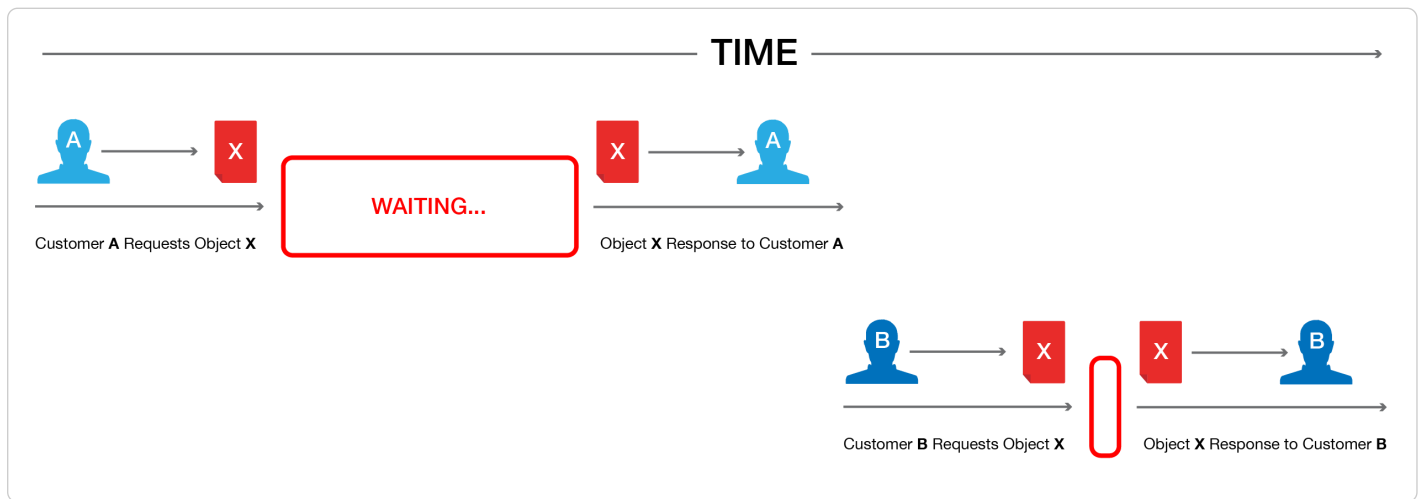
## § Serving stale content (/guides/performance-tuning/serving-stale-content)

Fastly can optionally serve stale content when there is a problem with your origin server, or if new content is taking a long time to fetch from your origin server. For example, if Fastly can't contact your origin server, our POPs (/guides/about-fastly-services/fastly-pop-locations) will continue to serve cached content when users request it. These features are not enabled by default.

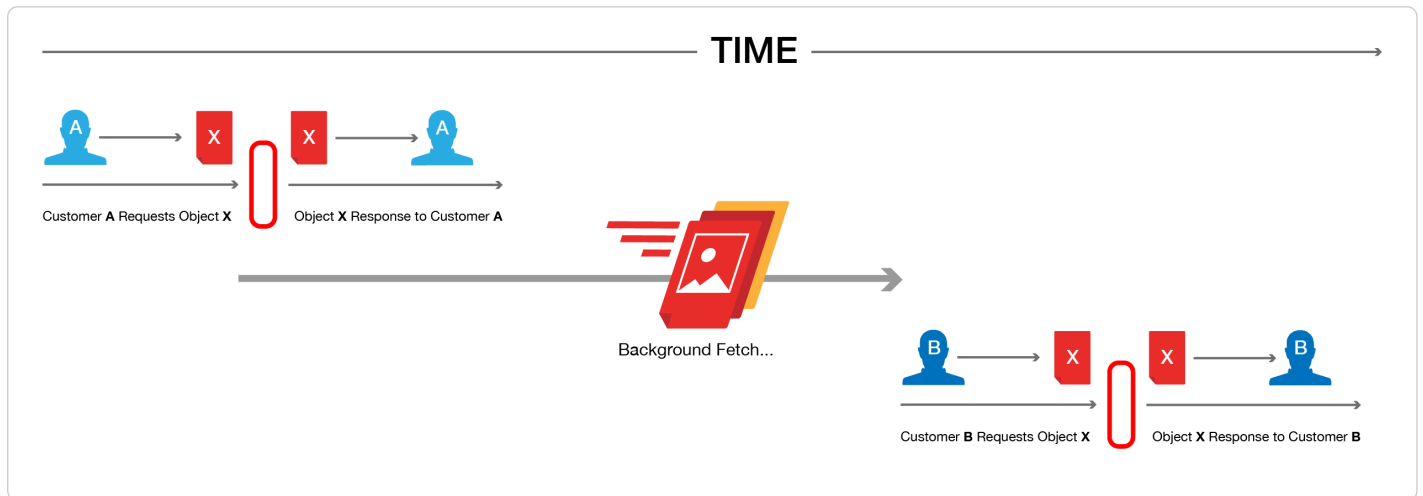
### Serving old content while fetching new content

Certain pieces of content can take a long time to generate. Once the content is cached it will be served quickly, but the first user to try and access it will pay a penalty.





This is unavoidable if the cache is completely cold, but if this is happening when the object is in cache and its TTL is expired, then Fastly can be configured to show the stale content while the new content is fetched in the background.



Fastly builds on the behavior proposed in RFC 5861 (<http://tools.ietf.org/html/rfc5861>) "HTTP Cache-Control Extensions for Stale Content" by Mark Nottingham, which is under consideration for inclusion in Google's Chrome browser ([https://www.mnot.net/blog/2014/06/01/chrome\\_and\\_stale-while-revalidate](https://www.mnot.net/blog/2014/06/01/chrome_and_stale-while-revalidate)).

## Usage

To activate this behavior simply add a `stale-while-revalidate` or `stale-if-error` statement to either the Cache-Control or Surrogate-Control headers in the response from your origin server. For example:

```
Cache-Control: max-age=600, stale-while-revalidate=30
```

will cache some content for 10 minutes and, at the end of that 10 minutes, will serve stale content for up to 30 seconds while new content is being fetched.

Similarly, this statement:

```
Surrogate-Control: max-age=3600, stale-if-error=86400
```

instructs the cache to update the content every hour (3600 seconds) but if the origin is down then show stale content for a day (86400 seconds).

Alternatively, these behaviors can be controlled from within VCL by setting the following variables in `vcl_fetch`:

```
set beresp.stale_while_revalidate = 30s;
```

```
set beresp.stale_if_error = 86400s;
```

## Interaction with grace

Stale-if-error works exactly the same as Varnish's `grace` variable such that these two statements are equivalent:

```
set beresp.grace = 86400s;
```

```
set beresp.stale_if_error = 86400s;
```

However, if a grace statement is present in VCL it will override any `stale-while-revalidate` or `stale-if-error` statements in any Cache-Control or Surrogate-Control response headers.

Setting `beresp.stale_if_error` either via header or via VCL does nothing on its own. In order to serve stale, follow the instructions below.

## Serving stale content on errors

In certain situations where your origin server becomes unavailable, you may want to serve stale content. These instructions provide an advanced configuration that allows all three possible origin failure cases to be handled using VCL. These instructions require the ability to upload custom VCL.

**❗ IMPORTANT:** The ability to [upload custom VCL code](/guides/vcl/uploading-custom-vcl) (</guides/vcl/uploading-custom-vcl>) is disabled by default when you sign up. To enable this ability for your account, contact [support@fastly.com](mailto:support@fastly.com) (<mailto:support@fastly.com>) and request it.

In the context of Varnish, there are three ways an origin can fail:

- The origin can be marked as unhealthy by failing health checks.
- If Varnish cannot contact the origin for any reason, a 503 error will be generated.
- The origin returns a valid HTTP response, but that response is not one we wish to serve to users (for instance, a 503).

The custom VCL shown below handles all three cases. If the origin is unhealthy, the default serve stale behavior is triggered by `stale-if-error`. In between the origin failing and being marked unhealthy, Varnish would normally return 503s. The custom VCL allows us to instead either serve stale if we have a stale copy, or to return a synthetic error page. The error page can be customized. The third case is handled by intercepting all 5XX errors in `vcl_fetch` and either serving stale or serving the synthetic error page.

**⚠ WARNING:** Do not [purge all](/guides/purging/single-purges#purging-all-content) (</guides/purging/single-purges#purging-all-content>) cached content if you are seeing [503 errors](/guides/debugging/common-503-errors) (</guides/debugging/common-503-errors>). Purge all overrides `stale-if-error` and increases the requests to your origin server, which could result in additional 503 errors.

Although not strictly necessary, health checks (</guides/basic-configuration/health-checks-tutorial>) should be enabled in conjunction with this VCL. Without health checks enabled, all of the functionality will still work, but serving stale or synthetic responses will take much longer while waiting for an origin to timeout. With health checks enabled, this problem is averted by the origin being marked as unhealthy.

The custom VCL shown below includes the Fastly standard boilerplate. Before uploading this to your service, be sure to customize or remove the following values to suit your specific needs:

- `if (beresp.status >= 500 && beresp.status < 600)` should be changed to include any HTTP response codes you wish to serve stale/synthetic for.
- `set beresp.stale_if_error = 86400s;` controls how long content will be eligible to be served stale and should be set to a meaningful amount for your configuration. If you're sending `stale-if-error` in Surrogate-Control or Cache-Control from origin, remove this entire line.
- `set beresp.stale_while_revalidate = 60s;` controls how long the `stale_while_revalidate` feature will be enabled for an object and should be set to a meaningful amount for your configuration. This feature causes Varnish to serve stale on a cache miss and fetch the newest version of the object from origin in the background. This can result in large performance gains on objects with short TTLs, and in general on any cache miss. Note that `stale_while_revalidate` overrides `stale-if-error`. That is, as long as the object is eligible to be served stale while revalidating, `stale-if-error` will have no effect. If you're sending `stale_while_revalidate` in Surrogate-Control or Cache-Control from origin, remove this entire line.
- `synthetic {"<!DOCTYPE html>Your HTML!</html>"};` is the synthetic response Varnish will return if no stale version of an object is available and should be set appropriately for your configuration. You can embed your HTML, CSS, or JS here. Use caution when referencing external CSS and JS documents. If your origin is offline they may be unavailable as well.

```

sub vcl_recv {
    if (req.http.Fastly-FF) {
        set req.max_stale_while_revalidate = 0s;
    }

#FASTLY recv

    if (req.request != "HEAD" && req.request != "GET" && req.request != "FASTLYPURGE") {
        return(pass);
    }

    return(lookup);
}

sub vcl_fetch {
    /* handle 5XX (or any other unwanted status code) */
    if (beresp.status >= 500 && beresp.status < 600) {

        /* deliver stale if the object is available */
        if (stale.exists) {
            return(deliver_stale);
        }

        if (req.restarts < 1 && (req.request == "GET" || req.request == "HEAD")) {
            restart;
        }

        /* else go to vcl_error to deliver a synthetic */
        error 503;
    }

    /* set stale_if_error and stale_while_revalidate (customize these values) */
    set beresp.stale_if_error = 86400s;
    set beresp.stale_while_revalidate = 60s;

#FASTLY fetch

    if ((beresp.status == 500 || beresp.status == 503) && req.restarts < 1 && (req.request == "GET" || req.request == "HEAD")) {
        restart;
    }

    if (req.restarts > 0) {
        set beresp.http.Fastly-Restarts = req.restarts;
    }

    if (beresp.http.Set-Cookie) {
        set req.http.Fastly-Cachetype = "SETCOOKIE";
        return(pass);
    }

    if (beresp.http.Cache-Control ~ "private") {
        set req.http.Fastly-Cachetype = "PRIVATE";
        return(pass);
    }

    /* this code will never be run, commented out for clarity */
    /* if (beresp.status == 500 || beresp.status == 503) {
        set req.http.Fastly-Cachetype = "ERROR";
        set beresp.ttl = 1s;
        set beresp.grace = 5s;
        return(deliver);
    } */

    if (beresp.http.Expires || beresp.http.Surrogate-Control ~ "max-age" || beresp.http.Cache-Control ~ "(s-maxage|max-age)") {
        # keep the ttl here
    } else {
        # apply the default ttl
        set beresp.ttl = 3600s;
    }

    return(deliver);
}

sub vcl_hit {
#FASTLY hit

    if (!obj.cacheable) {
        return(pass);
    }

    return(deliver);
}

```

```

sub vcl_miss {
#FASTLY miss
    return(fetch);
}

sub vcl_deliver {
    if (resp.status >= 500 && resp.status < 600) {
        /* restart if the stale object is available */
        if (stale.exists) {
            restart;
        }
    }

#FASTLY deliver
    return(deliver);
}

sub vcl_error {
#FASTLY error

    /* handle 503s */
    if (obj.status >= 500 && obj.status < 600) {

        /* deliver stale object if it is available */
        if (stale.exists) {
            return(deliver_stale);
        }

        /* otherwise, return a synthetic */

        /* include your HTML response here */
        synthetic {"<!DOCTYPE html><html>Replace this text with the error page you would like to serve to clients if your origin is offline.</html>"};
        return(deliver);
    }
}

sub vcl_pass {
#FASTLY pass
}

sub vcl_log {
#FASTLY log
}

```

## Why serving stale content may not work as expected

Here are some things to consider if Fastly isn't serving stale content:

- **Cache:** Stale objects are only available for cacheable content.
- **Shielding:** If you don't have shielding (</guides/performance-tuning/shielding>) enabled, a POP (</guides/about-fastly-services/fastly-pop-locations>) can only serve stale if a request for that cacheable object was made through that POP before. We recommend enabling shielding to increase the probability that stale exists. Shielding is also a good way to quickly refill the cache after a performing a purge all (</guides/purging/single-purges#purging-all-content>).
- **Requests:** As traffic to your site increases, you're more likely to see stale objects available (even if shielding is disabled). It's reasonable to assume that popular assets will be cached at multiple POPs.
- **Least Recently Used (LRU):** Fastly has an LRU list, so objects are not necessarily guaranteed to stay in cache for the entirety of their TTL (time to live). But eviction is dependent on many factors, including the object's request frequency, its TTL, the POP from which it's being served. For instance, objects with a TTL of longer than 3700s get written to disk, whereas objects with shorter TTLs end up in transient, in-memory-only storage. We recommend setting your TTL to more than 3700s when possible.
- **Purges:** Whenever possible, you should purge content using our soft purge feature (</guides/purging/soft-purges>). Soft purge allows you to easily mark content as outdated (stale) instead of permanently deleting it from Fastly's caches. If you can't use soft purge, we recommend purging by URL (</guides/purging/single-purges#purging-a-url>) or using surrogate keys (</guides/purging/getting-started-with-surrogate-keys>) instead of performing a purge all (</guides/purging/single-purges#purging-all-content>).

## Additional reading

- RFC 5861 (<http://tools.ietf.org/html/rfc5861>)
- Mark Nottingham's blog post ([https://www.mnot.net/blog/2014/06/01/chrome\\_and\\_stale-while-revalidate](https://www.mnot.net/blog/2014/06/01/chrome_and_stale-while-revalidate))
- Chrome discussion (<https://groups.google.com/a/chromium.org/forum/#!msg/chromium-dev/zchogDvIYrY/ZqWSdt3LJdMJ>)

## § Shielding (/guides/performance-tuning/shielding)


Fastly's shielding service feature (/guides/about-fastly-services/about-fastlys-origin-shielding-features) allows you to designate a specific Point of Presence (POP) as a shield node to your origins. Once designated, all requests to your origin will go through that datacenter, increasing the chances of getting a HIT (/guides/performance-tuning/understanding-cache-hit-and-miss-headers-with-shielded-services) for a given resource. If a different POP doesn't have a specific object, it will query the shield (provided it's not down for maintenance) instead of your origins.

### How shielding works


When a user requests brand new content from a customer's server and that content has never been cached by any Fastly POP (/guides/about-fastly-services/fastly-pop-locations), this is what happens to their request when shielding is and is not enabled.

#### Without shielding enabled

Without shielding enabled, when the first request for content arrives at POP A, the POP does not have the content cached. It passes the request along to a customer's origin server to get the content. Once the content is retrieved, POP A caches it and sends it on to the user.




When a second request for that same content arrives at POP A, the content is already cached. No request goes to the customer's origin server. It's merely sent from the cached copy.




If the second request were to arrive at POP B instead of POP A, however, the request would once again be passed along to the customer's origin server. It would then be cached and passed back to the end user, just like POP A did when that info was first requested.

#### With shielding enabled

With shielding enabled (/guides/performance-tuning/shielding), when the first request for content arrives at the POP A, that POP does not have the content cached. It passes the request along to the shield POP, which also doesn't have the content cached. The shield POP passes the request along to the customer's origin server. It then caches the content that's retrieved and passes it along to POP A. POP A then passes the content along to the user.



When a second request for that same content arrives at POP A, the content is already cached, so no request goes to the shield POP or the customer's origin server.



If the second request were to arrive at POP B instead of POP A, however, the request would be passed along to the shield POP. That shield POP already has a cached copy from the first request to POP A. No future requests for the content would be passed along to the customer's origin server until the shield POP's cached copy of it expires.

## Enabling shielding

**❗ IMPORTANT:** If you are using Google Cloud Storage as your origin, you need to follow the steps in our [GCS setup guide \(/guides/integrations/google-cloud-storage\)](https://docs.fastly.com/guides/integrations/google-cloud-storage) instead of the steps below.

Enable shielding with these steps:

1. Read the caveats of shielding information below for details about the implications of and potential pitfalls involved with enabling shielding for your organization.
2. Log in to the Fastly web interface and click the **Configure** link.
3. From the service menu, select the appropriate service.
4. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.
5. Click the **Origins** tab. The Origins page appears.
6. Click the name of the host you want to edit. The Edit a host page appears.

## Edit a host

### CONDITION

This will happen all the time unless you [Attach a condition](#)

Address  :  ★ Required

The IP address (or hostname) and port number for your origin server. To [enable TLS](#), use port 443.

Shielding

The shield POP designated to reduce inbound load on your origins by serving cached data. Learn more about [POP request handling](#) and the [caveats of shielding](#).

Health check

[+ Create a new health check](#)

A small request assigned to [test the state of your origin](#).

Auto load balance

Enables Fastly to [load balance between two or more backend hosts](#).

[► Tips](#)

Weight

A portion of traffic to send to a specific origin. Each origin receives (weight ÷ total weight across origins) of the traffic.

Description  ★ Required

The name of your origin, such as [My origin server](#).

- From the **Shielding** menu, select the datacenter to use as your shield.

**★ TIP:** Generally, we recommend selecting a datacenter close to your backend. Doing this allows faster content delivery because we optimize requests between the edge POP (the one close to your user) and the shield POP (the one close to your server).

**❗ NOTE:** With multiple backends, each backend will have its own shield defined. This allows flexibility if your company has backends selected geographically and different shield POPs are desired.

- Click **Update** to save your changes.
- If you have changed the default host or have added a header to change the host, add the modified hostname to your list of domains. Do this by clicking the **Domains** tab and checking to make sure the host in question appears on the page. If it isn't included, add it by clicking the **Create domain** button.

| DOMAINS<br>2 Domains | ORIGINS<br>2 Hosts<br>0 Health checks | SETTINGS<br>0 Request settings<br>1 Cache setting | CONTENT<br>1 Header<br>0 Gzips<br>0 Responses | LOGGING<br>0 Endpoints | CUSTOM VCL<br>0 Files |
|----------------------|---------------------------------------|---|---|------------------------|-----------------------|
|----------------------|---------------------------------------|---|---|------------------------|-----------------------|

### Domains

Domains are used to route requests to your service. Customers associate their domain names with their origin (content source) when provisioning a Fastly service.

Test the domain to see how your service is being pulled through our network. This will help identify if a problem is a DNS issue, or a Fastly configuration problem.

+ CREATE DOMAIN

|                 |             |        |
|-----------------|-------------|--------|
| www.example.org | Test domain | Delete |
| www.test.co     | Test domain | Delete |

With shielding enabled, queries from other POPs appear as incoming requests to the shield. If the shield doesn't know about the modified hostname, it doesn't know which service to match the request to. Including the origin's hostname in the domain list eliminates this concern.

10. Click the **Activate** button to deploy your configuration changes.

## Caveats of shielding

Shielding not only impacts traffic and hit ratios, it affects configuration and performance. When you configure shielding, be aware of the following caveats.

### Traffic and hit ratio caveats

**Inbound traffic to a shield will be billed as regular traffic, including requests to populate remote POPs.** Enabling shielding will incur some additional Fastly bandwidth charges, but will be offset by savings of your origin bandwidth (and origin server load). Pass-through requests will not go directly to the origin, they will go through the shield first.

**Global HIT ratio calculation may seem lower than the actual numbers.** Shielding is not taken into account when calculating the global hit ratio. If an edge node doesn't have an object in its cache, it reports a miss. Local MISS/Shield HIT gets reported as a miss and a hit in the statistics, even though there is no call to the backend. It will also result in one request from the edge node to the shield. Local MISS/Shield MISS will result in two requests, because we will subsequently fetch the resource from your origin. For more information about caching with shielding see our article [Understanding Cache HIT and MISS with Shielding Services \(/guides/performance-tuning/understanding-cache-hit-and-miss-headers-with-shielded-services\)](/guides/performance-tuning/understanding-cache-hit-and-miss-headers-with-shielded-services).

### Configuration caveats

**You will be unable to manually define backends using VCL.** Shielding at this level is completely dependent on backends being defined as actual objects through the web interface or API. Other custom VCL (</guides/vcl/uploading-custom-vcl>) will work just fine.

**You can only set one shield total if automatic load balancing is selected via the web interface.** If you've selected auto load balancing, you can only select one shield total. You must use custom VCL to use multiple shields when auto load balancing is set.

**Enabling sticky load balancing and shielding at the same time requires custom VCL.** Sticky load balancers use `client.identity` to choose where to send the session. The `client.identity` defaults to the IP request header. That's fine under normal circumstances, but if you enable shielding, the IP will be the original POP's IP, not the client's IP. Thus, to enable shielding and a custom sticky load balancer, you want to use the following:

```
if (req.http.fastly-ff) {
  set client.identity = req.http.Fastly-Client-IP;
}
```

**You'll need to use caution when changing the host header before it reaches the shield.** Fastly matches a request with a host header. If the host header doesn't match to a domain within the service an error of 500 is expected. Also, purging conflicts can occur if the host header is changed to a domain that exists in a different service.



For example, say Service A has hostname `a.example.com` and Service B has hostname `b.example.com`. If Service B changes the host header to `a.example.com`, then the edge will think the request is for Service B but the shield will think the request is for Service A.

When you purge an object from Service B and not from Service A, the shield will serve the old object that you wanted to purge to the edge, since the purge went out to Service B and not Service A. You will want to purge the object from both Service A and Service B. However, this opens the door for confusion and error.

**VCL gets executed twice: once on the shield node and again on the edge node.** Changes to `beresp` and `resp` can affect the caching of a URL on the shield and edge. Consider the following examples.

Say you want Fastly to cache an object for one hour (3600 seconds) and then ten seconds on the browser. The origin sends `Cache-Control: max-age=3600`. You unset `beresp.http.Cache-Control` and then reset `Cache-Control` to `max-age=10`. With shielding enabled, however, the result will not be what you expect. The object will have `max-age=3600` on the shield and reach the edge with `max-age=10`.

A better option in this instance would be to use `Surrogate-Control` and `Cache-Control` response headers. `Surrogate-Control` overrides `Cache-Control` and is stripped after the edge node. The `max-age` from `Cache-Control` will then communicate with the browser. The origin response headers would look like this:

```
Surrogate-Control: max-age=3600
Cache-Control: max-age=10
```

Another common pitfall involves sending the wrong `vary` header to an edge POP. For example, there's VCL that takes a specific value from a cookie, puts it in a header, and that header is then added to the `vary` header. To maximize compatibility with any caches outside of your control (such as with shared proxies as commonly seen in large enterprises), the `vary` header is updated in `vcl_deliver`, replacing the custom header with `Cookie`. The code might look like this:

```
vcl_recv {
    # Set the custom header
    if (req.http.Cookie ~ "ABtesting=B") {
        set req.http.X-ABtesting = "B";
    } else {
        set req.http.X-ABtesting = "A";
    }
    ...
}

...

sub vcl_fetch {
    # Vary on the custom header
    if (beresp.http.Vary) {
        set beresp.http.Vary = beresp.http.Vary ", X-ABtesting";
    } else {
        set beresp.http.Vary = "X-ABtesting";
    }
    ...
}

...

sub vcl_deliver {
    # Hide the existence of the header from downstream
    if (resp.http.Vary) {
        set resp.http.Vary = regsub(resp.http.Vary, "X-ABtesting", "Cookie");
    }
}
```

When combined with shielding, however, the effect of the above code will be that edge POPs will have `Cookie` in the `vary` header, and thus will have a terrible hit rate. To work around this, amend the above VCL so that `vary` is only updated with `Cookie` when the request is not coming from another Fastly cache. The `Fastly-FF` header is a good way to tell. The code would look something like this (including the same `vcl_recv` from the above example):

```
# Same vcl_recv from above code example

sub vcl_fetch {
  # Vary on the custom header, don't add if shield POP already added
  if (beresp.http.Vary != "X-ABtesting") {
    if (beresp.http.Vary) {
      set beresp.http.Vary = beresp.http.Vary ", X-ABtesting";
    } else {
      set beresp.http.Vary = "X-ABtesting";
    }
  }
  ...
}

...

sub vcl_deliver {
  # Hide the existence of the header from downstream
  if (resp.http.Vary && !req.http.Fastly-FF) {
    set resp.http.Vary = regsub(resp.http.Vary, "X-ABtesting", "Cookie");
  }
}
```

## § Tracking your origin's name, IP, and port (/guides/performance-tuning/tracking-your-origins-name-ip-and-port)

Fastly provides three values captured in `vcl_fetch` that allow you to see and track origin information:

- `beresp.backend.name`
- `beresp.backend.port`
- `beresp.backend.ip`

While these three values are immensely useful, you may want to use this information within `vcl_deliver` for things like response information or remote log streaming. You can do this by:

1. Creating cache headers that capture the origin information.
2. Adding a response header to the log format to capture the response output.

## Capturing the Origin Information

To track your origin's name, IP, and port, you need to create two separate headers: one that captures the origin name and another that captures the origin's IP and port (e.g., 80, 443).

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.
4. Click the **Content** tab. The Content page appears.
5. Click the **Create header** button to create a header that captures the name of your origin. The Create a new header page appears.

## Create a new header

Learn more about this section in our [headers tutorial](#).

### CONDITION

This will happen all the time unless you [Attach a condition](#)

Type / Action

Cache

Set

The type of header and the action performed on it.

Destination

http.Backend-Name

★ Required

The name of the header that will be affected by the selected action. For example: **http.Content-Type**, **http.Set-Cookie**, **http.Via**, **http.Location**, or **http.Access-Control-Allow-Origin**.

Source

beresp.backend.name

★ Required

New content for the header. Can be a static value (e.g. string or number) or a dynamic value (e.g. existing header or a GeoIP value). Please use quotes for string values.

Ignore if set

Yes

If switched to **Yes**, the action will not be performed if the header in Destination exists.

Priority

10

★ Required

The order in which the header rules execute within the condition.

Description

Backend Name

★ Required

The name of your header, such as **My header**.

CREATE

CANCEL

6. Fill out the **Create a new header** fields as follows:

- From the **Type** menu, select **Cache**, and from the **Action** menu, select **Set**.
- In the **Destination** field, type `http.Backend-Name` (or preferred header variable).
- In the **Source** field, type `beresp.backend.name`.
- From the **Ignore if set** menu, select **Yes**.
- In the **Priority** field, type `10`.
- In the **Description** field, type `Backend Name` (or preferred name). This name is displayed in the Fastly web interface.

7. Click the **Create** button. The new header appears on the **Content** page.

8. Click the **Create header** button to create a header that captures the IP and port of your origin. The Create a new header page appears.

## Create a new header

Learn more about this section in our [headers tutorial](#).

CONDITION

This will happen all the time unless you [Attach a condition](#)

Type / Action

Cache

Set

The type of header and the action performed on it.

Destination

http.Backend-IP-Port

★ Required

The name of the header that will be affected by the selected action. For example: **http.Content-Type**, **http.Set-Cookie**, **http.Via**, **http.Location**, or **http.Access-Control-Allow-Origin**.

Source

beresp.backend.ip ", " beresp.backend.port

★ Required

New content for the header. Can be a static value (e.g. string or number) or a dynamic value (e.g. existing header or a GeoIP value). Please use quotes for string values.

Ignore if set

Yes

If switched to **Yes**, the action will not be performed if the header in Destination exists.

Priority

10

★ Required

The order in which the header rules execute within the condition.

Description

Backend IP and Port

★ Required

The name of your header, such as **My header**.

CREATE

CANCEL

9. Fill out the **Create a new header** fields as follows:

- From the **Type** menu, select **Cache**, and from the **Action** menu, select **Set**.
- In the **Destination** field, type `http.Backend-IP-Port` (or preferred header variable).
- In the **Source** field, type `beresp.backend.ip ", " beresp.backend.port`.
- From the **Ignore if set** menu, select **Yes**.
- In the **Priority** field, type `10`.
- In the **Description** field, type `Backend IP and Port` (or preferred name). This name is displayed in the Fastly interface.

10. Click the **Create** button. The new header appears on the **Content** page.

11. Click the **Activate** button to deploy your configuration changes.

## Adding a Response Header to the Log Format

The values captured in a header within `vcl_fetch` will flow to `vcl_deliver`. For example, there will exist a `resp.http.Backend-Name` header in `vcl_deliver` that corresponds to `beresp.http.Backend-Name` in `vcl_fetch`. By default, the response header will be included in the response output.

Unfortunately, with remote log streaming, you cannot add the `vcl_fetch` header, `beresp.http.Header-Name`, to the log format. However, you can add its cousin in `vcl_deliver`, `resp.http.Header-Name`.

Add `resp.http.Header-Name` to the log format that you configured by following the instructions in the Remote Log Streaming guide (/guides/streaming-logs/setting-up-remote-log-streaming). Using the example above, you would add `resp.http.Backend-Name` and `resp.http.Backend-IP-Port`.

## Important Notes

### Regarding Shielding

Notice within the example the field **Ignore if set** is set to **Yes**. With shielding, the VCL is executed twice, once on the shield and again on the edge node. This setting will display the original information from the origin without overriding it on the edge node.

You can also set the field **Ignore if set** to **No** with shielding enabled. In this scenario, the edge node captures the shield's information within `beresp.backend.name`, `beresp.backend.ip`, and `beresp.backend.port`.

If remote log streaming is configured, remember it is executed twice. Thus the first log (from the shield node) will have the origin's information and the second log (from the edge node) will have the shield's information.

### Regarding Security

For security purposes, you may want to track the information in logging but not display all or some of it in the response. This is possible but requires custom VCL to strip the information after sending the log line from the edge node.

**❗ IMPORTANT:** The ability to [upload custom VCL code](/guides/vcl/uploading-custom-vcl) (/guides/vcl/uploading-custom-vcl) is disabled by default when you sign up. To enable this ability for your account, contact [support@fastly.com](mailto:support@fastly.com) (mailto:support@fastly.com) and request it.

Don't forget to read our guide to using custom VCL (/guides/vcl/mixing-and-matching-fastly-vcl-with-custom-vcl) before you begin. Remember to include the entire boilerplate if you do not intend to override the Fastly default settings.

Then add the following snippet within `vcl_deliver` with the headers you want to strip. Using our example above, we continue to send the value of the origin's name within the response. We strip the origin's IP and port from the response information.

```
sub vcl_deliver {
#FASTLY deliver

    if (!req.http.Fastly-FF) {
        unset resp.http.Backend-IP-Port;
    }

    return(deliver);
}
```

## § Understanding cache HIT and MISS headers with shielded services (/guides/performance-tuning/understanding-cache-hit-and-miss-headers-with-shielded-services)

Here's some help deciphering cache hit and miss headers when you have shielding enabled (/guides/performance-tuning/shielding). Let's look at the following requests for the same object if you had run a cURL command in your terminal (for example, `curl -svo /dev/null www.example.com`) to return the Fastly headers.

The first request for an object using the above cURL command might produce output something like this:

```
X-Served-By: cache-iad2120-IAD, cache-sjc3120-SJC
X-Cache: MISS, MISS
X-Cache-Hits: 0, 0
```

For this first request, the two cache-nodes in `X-Served-By` show that shielding is turned on, with `cache-iad2120-IAD` serving as the cache node and `cache-sjc3120-SJC` serving as the local delivering node. The `X-Cache: MISS, MISS` indicates that the requested object was neither in the shield cache (a MISS) nor the local delivering node (also a MISS). The `X-Cache-Hits` reflects that same MISS information because it displays `0, 0`.

The second request for an object using the above cURL command might produce output something like this:

```
X-Served-By: cache-iad2120-IAD, cache-sjc3120-SJC
X-Cache: MISS, HIT
X-Cache-Hits: 0, 1
```

This second time, we hit the same local cache-node ( `cache-sjc3120-SJC` ) and got a HIT. The MISS listed for `cache-iad2120-IAD` reflects the state of that node the last time it queried that node for that object (not its current state), which at that time was a MISS, but in fact, the object resides in that cache now because it was requested from origin.

Waiting a minute or two and requesting the same object a third time using the above cURL command might produce output something like this:

```
X-Served-By: cache-iad2120-IAD, cache-sjc3120-SJC
X-Cache: HIT, MISS
X-Cache-Hits: 1, 0
```

This third request shows a new cache ( `cache-sjc3120-SJC` ) being selected, MISSING, and retrieving the object from the shield node ( `cache-iad2120-IAD` ), which had the object (a HIT). In total, there should have only been a single request for this object to the origin.

Keep in mind that if the closest delivering cache node exists in the shield datacenter, you will only see a single server and HIT data such as:

```
X-Served-By: cache-iad2120-IAD
X-Cache: HIT
X-Cache-Hits: 1
```

After a purge of the object, requesting the object again via the above cURL command will produce results similar to the first request scenario. For example:

```
X-Served-By: cache-iad2120-IAD, cache-sjc3120-SJC
X-Cache: MISS, MISS
X-Cache-Hits: 0, 0
```

## § Understanding the X-Timer header (/guides/performance-tuning/understanding-the-xtimer-header)

If you look at the raw headers returned with a response from a Fastly cached asset, you will notice some extra headers tacked on. One in particular is X-Timer. This header provides timing information about the journey of a request from end to end.

Here are two examples of X-Timer headers:

- `S1392947468.641059399,VS0,VE0` (a cache HIT)
- `S1392951663.217806578,VS0,VE31` (a cache MISS)

Let's break these headers down into their parts, separated by commas, and examine what each part means.

The diagram illustrates the structure of the X-Timer header. It shows two examples of headers: `S1392947468.641059399,VS0,VE0` and `S1392951663.217806578,VS0,VE31`. Brackets below each header group the characters into three sections:
 

- Unix time stamp:** The first section, starting with 'S', contains a timestamp in seconds and microseconds (e.g., `S1392947468.641059399`).
- Varnish Start:** The second section, starting with 'VS', represents the start of the varnish part of the request's journey (e.g., `VS0`).
- Varnish End:** The third section, starting with 'VE', represents the sum of the length of the trip (e.g., `VE0` for a hit and `VE31` for a miss).

The first section of the header, starting with `S`, represents a Unix timestamp ([https://en.wikipedia.org/wiki/Unix\\_time](https://en.wikipedia.org/wiki/Unix_time)) of the start of the request on our edges.

The next section, `VS` or "varnish start," represents the start of the varnish part of the request's journey. This should always be 0 (we've got to start counting somewhere).

And the last section, `VE` or "varnish end," represents the sum of the length of the trip. For cache HITs, the length of the trip will nearly always be 0 (not actually zero, but less than a millisecond is rounded down). For cache MISSs, the number represents the number of milliseconds it took to retrieve the data from your origin server and send the response back to the requester. In the example above, it took 31ms to retrieve the data.

## § Using edge side includes (ESI) (/guides/performance-tuning/using-edge-side-includes)

Fastly supports the following edge side includes (ESI) language (<http://www.w3.org/TR/esi-lang>) elements:

- include
- comment
- remove

We don't support the following ESI language elements:

- inline
- choose | when | otherwise
- try | attempt | except
- vars
- ESI Variables

• [Guides \(/guides/\)](/guides/) > [Diagnostics and performance \(/guides/diagnostics\)](/guides/diagnostics/) > [Security \(/guides/security/\)](/guides/security/)

## § Monitoring account activity with audit logs (/guides/security/monitoring-account-activity-with-audit-logs)

Audit logs keep track of events related to your services, account, and users. You can use audit logs to determine which changes were made and by whom. Audit logs are available for retrieval via the Fastly API. A limited subset of your services' audit logs are also available within the web interface. Audit log data is currently retained indefinitely.

### Types of events that are logged

Audit logs record many types of events that users can perform:

- **Version management:** version activation, version deactivation, and version lock
- **User behavior:** log in, log out, failed login, account locked, account unlocked, and password updated
- **Security:** two-factor authentication enabled, two-factor authentication disabled, API token created, API token deleted, and API key created
- **Account management:** new invitation, invitation accepted, user updated, user deleted, new company addition, and company settings updated
- **Billing:** pricing plan change and billing contact change
- **Purging:** purge all
- **Configuration settings (creation, deletion, and updates):** domain, backend, header, health check, cache setting, request setting, logging endpoint, VCL file, gzip setting, and synthetic response

### When audit logs might be useful

- Seeing who activated the most recent version of your service
- Reviewing who logged in to your account via the web interface
- Learning which users have two-factor authentication enabled
- Viewing recent service configuration setting changes

### Accessing audit logs

There are two ways to access audit logs: by using the Fastly API or by viewing the latest events in the web interface. The events displayed in the web interface are a subset of those available via the API.

#### Using the API

You can use the `/events` API endpoint to access your audit logs. Events can be filtered by `user_id`, `service_id`, `customer_id`, and `event_type`. For example, you could make the following API call in a terminal application to view all recent events:

```
curl -H "Fastly-Key: FASTLY_KEY" https://api.fastly.com/events?filter[customer_id]=x4xCwxxJxGCx123Rx5xTx&page[number]=1&page[size]=1
```

The response will look like this:

```
{
  "data": [
    {
      "attributes": {
        "admin": false,
        "created_at": "2016-06-06T20:05:10Z",
        "customer_id": "x4xCwxxJxGCx123Rx5xTx",
        "description": "Version 2 was activated",
        "event_type": "version.activate",
        "ip": "127.0.0.0",
        "metadata": {
          "version_number": 2
        },
        "service_id": "SU1Z0isxPaozGVKXdv0eY",
        "user_id": "4Pp0BW3UkBEJhG3N0kovLP"
      },
      "id": "5IH1QmNSV1Qi7jXc4oIZ1Z",
      "type": "event"
    }
  ],
  "links": {
    "last": "https://api.fastly.com/events?filter[customer_id]=x4xCwxxJxGCx123Rx5xTx&page[number]=1&page[size]=1"
  }
}
```

See the API documentation ([/api/account#events](https://api.fastly.com/events)) for more information.

## Using the web interface

The web interface displays the last 20 service-related events for the selected service. Events related to users and accounts are not displayed in the web interface.

Follow these instructions to access the audit logs for a service:

1. Log in to the Fastly web interface and click the **Configure** link.
2. The most recent service-related events are displayed near the bottom of the page, in the Event Log area.



Fastly
ABCDEFG1234567
created 2012-01-01

www.example.com
Version 1
is currently active
EDIT CONFIGURATION
OPTIONS

Show service ID

PURGE

HIT RATIO  
90.5%

HITS  
2,627,015

MISSES  
274,618

ERRORS  
28,607

REQUESTS

ORIGIN LATENCY

Recent Non-Active Versions

|           |                               |                         |
|-----------|-------------------------------|-------------------------|
| Version 5 | ± Compare with active version | 2016-11-15 21:35 +00:00 |
| Version 4 | ± Compare with active version | 2016-11-15 21:34 +00:00 |
| Version 3 | ± Compare with active version | 2016-11-11 23:51 +00:00 |
| Version 2 | ± Compare with active version | 2016-07-06 15:46 +00:00 |

Event Log

|  |                         |
|--|-------------------------|
| Version 1 was activated<br>by Test Admin | 2016-11-15 21:56 +00:00 |
| Version 4 was cloned<br>by Test Admin    | 2016-11-15 21:35 +00:00 |

## § Penetration testing your service behind Fastly (/guides/security/penetration-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](https://www.fastly.com/acceptable-use) (<https://www.fastly.com/acceptable-use>). 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, create a Customer Support ticket by contacting Fastly via email at [support@fastly.com](mailto:support@fastly.com) (<mailto:support@fastly.com>) at least two (2) business days before you begin any security testing. In your ticket, include these details:

- the IDs of the services (/guides/account-management-and-security/finding-and-managing-your-account-info#finding-your-service-id) that will be tested
- the source IP address 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 (/guides/compliance/security-program) and technology compliance (/guides/compliance/technology-compliance). To perform an independent audit of these programs, contact sales@fastly.com (mailto:sales@fastly.com) to discuss purchase of Compliance Services (/guides/customer-support/compliance-services).

★ **TIP:** We welcome security professionals researching potential vulnerabilities in our network under our guidelines for [reporting a security issue](https://www.fastly.com/security/report-security-issue) (<https://www.fastly.com/security/report-security-issue>).

• Guides (/guides/) > Migrations and integrations (/guides/mig-int) > Migrations (/guides/migrations/)

## § WordPress plugins: Migrating to Purgely from Fastly's original plugin (/guides/migrations/migrating-to-purgely-from-fastlys-original-plugin)

Purgely (<https://wordpress.org/plugins/purgely/>) is an open source WordPress plugin written by Condé Nast that integrates Fastly services into the WordPress platform. We recommend customers use Purgely over Fastly's original WordPress plugin (<https://wordpress.org/plugins/fastly/>) released in 2011, because Purgely takes advantage of newer WordPress and Fastly features, including Fastly's surrogate key (/guides/purging/getting-started-with-surrogate-keys) purging. Purgely also takes fewer steps to configure and delivers content to users faster than the original plugin.

❗ **IMPORTANT:** You can continue to use [Fastly's original WordPress plugin](https://wordpress.org/plugins/fastly/) (<https://wordpress.org/plugins/fastly/>) without service interruption. However, we do not plan to continue its development.

## Recommended migration path

To migrate from Fastly's WordPress plugin to the Purgely plugin, follow these steps.

1. Find your Fastly API key and Service ID in one of the following ways:
  - Look in the Fastly web interface (/guides/account-management-and-security/finding-and-managing-your-account-info).
  - Look in the Fastly WordPress plugin settings in the Plugins menu ([https://codex.wordpress.org/Plugins\\_Screen](https://codex.wordpress.org/Plugins_Screen)) of your WordPress instance.
2. Uninstall Fastly's WordPress plugin from your WordPress instance by following the Uninstalling Plugins ([https://codex.wordpress.org/Managing\\_Plugins#Uninstalling\\_Plugins](https://codex.wordpress.org/Managing_Plugins#Uninstalling_Plugins)) instructions.
3. Install the Purgely plugin using one of the following the methods:
  - **GUI installation (recommended).** Follow the Automatic Plugin Installation ([https://codex.wordpress.org/Managing\\_Plugins#Automatic\\_Plugin\\_Installation](https://codex.wordpress.org/Managing_Plugins#Automatic_Plugin_Installation)) instructions to add it as a new plugin by searching for it in the listed plugins and then clicking **Install Now** on the plugin's **Details** screen.
  - **Command line installation.** Install the WordPress command line tool (<https://codex.wordpress.org/wp-cli>) and then run the install command by typing `wp plugin install --activate purgely`.
  - **Manual installation.** Follow the Manual Plugin Installation ([https://codex.wordpress.org/Managing\\_Plugins#Manual\\_Plugin\\_Installation](https://codex.wordpress.org/Managing_Plugins#Manual_Plugin_Installation)) instructions to upload Purgely to your `/wp-content/plugins/` directory and then activate it from the Plugins menu ([https://codex.wordpress.org/Plugins\\_Screen](https://codex.wordpress.org/Plugins_Screen)) in the WordPress interface.
4. From the WordPress **Plugins** menu, click the **Settings** link in the navigation bar.
5. Click **Purgely**. The Purgely settings window appears.

**Purgely Settings**

**Fastly settings**

Please enter details related to your Fastly account. A Fastly API key and service ID are required for some operations (e.g., surrogate key and full cache purges).

**API Key**  *Required for surrogate key and full cache purges*  
API key for the Fastly account associated with this site. Please see Fastly's documentation for [more information on finding your API key](#).

**Service ID**  *Required for surrogate key and full cache purges*  
Fastly service ID for this site. Please see Fastly's documentation for [more information on finding your service ID](#).

**API Endpoint**   
API endpoint for this service.

**General settings**

This section allows you to configure general cache settings. Note that changes to these settings can cause destabilization to your site if misconfigured. The default setting should be sufficient for most sites.

**Cache TTL (in Seconds)**   
This setting controls the "surrogate-control" header's "max-age" value. It defines the cache duration for all pages on the site.

**Default Purge Type**  Soft  
The purge type setting controls the manner in which the cache is purged. Instant purging causes the cached object(s) to be purged immediately. Soft purging causes the origin to revalidate the cache and Fastly will serve stale content until revalidation is completed. For more information, please see [Fastly's documentation for more information on soft purging](#).

6. Configure the Purgely plugin by doing the following:

- In the **API Key** field, type your Fastly API key.
- In the **Service ID** field, type your Fastly service ID.

7. Click the **Save** button.

★ **TIP:** For other Purgely configuration settings, see the configuration instructions on the Purgely plugin [installation tab](https://wordpress.org/plugins/purgely/installation/) (<https://wordpress.org/plugins/purgely/installation/>).

• Guides (/guides/) > Migrations and integrations (/guides/mig-int) > Integrations (/guides/integrations/)

## § Acquia Cloud (/guides/integrations/acquia-cloud)

To use Acquia Cloud as an origin, you must sign up for both an Acquia Cloud subscription and Fastly services and connect the two.

### Sign up for an Acquia Cloud subscription

1. Using a web browser, navigate to the Acquia Cloud signup page (<https://www.acquia.com/free>).
2. Select the Acquia Cloud Free option. The account subscription form appears.
3. Fill out the form and click **Create** to sign up for a subscription and start Acquia's automated site creation process.

The automated portion of the Acquia subscription process can take three to five minutes to complete. You'll know the entire process ends successfully when you see the checkmark appear next to the word "Done."

### Check for domain alias conflicts

Ensure that your new site domain does not have a conflicting alias by running the `host` command on a command line. For example, the `host` command for the Test-Example-2 domain would be:

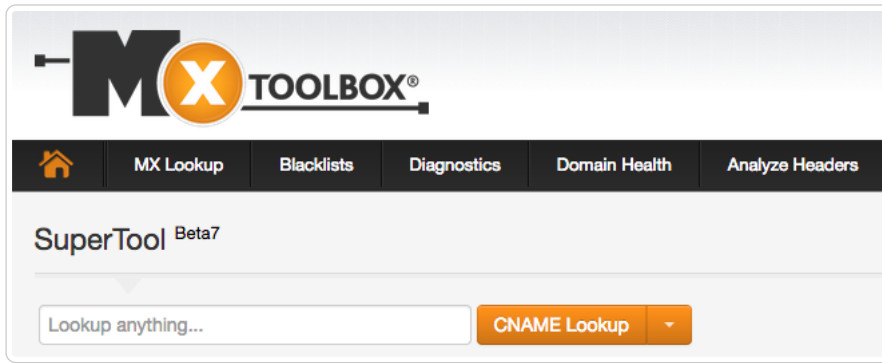
```
host test-example-2.devcloud.acquia-sites.com.
```

and would produce the following sample output:

```
test-example-2.devcloud.acquia-sites.com has address 127.0.0.1
```

### Determine your CNAME

1. Using a web browser, navigate to the MX Toolbox SuperTool (<http://mxtoolbox.com/SuperTool.aspx>).



2. In the **Lookup anything** field, type the name of your website domain.
3. From the menu to the right of the field, select **CNAME Lookup**. The domain name and canonical name appear below the field.
4. Save this information to use when you sign up for Fastly services.

## Sign up for Fastly CDN services

1. Using a web browser, navigate to the Fastly signup page (<https://www.fastly.com/signup>) and sign up for a Fastly account. The system sends a confirmation email to the address you specified during signup.
2. Verify your new Fastly account by clicking on the verification link sent to the email address used to sign up for Fastly service.
3. Log in to the Fastly web interface and complete your account configuration (</guides/basic-setup/sign-up-and-create-your-first-service>).

## Complete the integration

Once you have completed the signup and configuration steps, send an email to [acquia@fastly.com](mailto:acquia@fastly.com) (<mailto:acquia@fastly.com>) to complete the integration. Fastly will need to know:

- the email address associated with your new Fastly account and
- whether or not you require TLS (</guides/securing-communications/ordering-a-paid-tls-option>) for your customer-facing domain.

Keep in mind that Fastly and Acquia both run a Varnish Cache (<https://varnish-cache.org>). In order to properly configure your service, Fastly needs to make a few modifications during the final setup process to ensure compatibility between the two. In this final setup process, Fastly runs a script (<https://gist.github.com/vvuksan/66cc45e09812fbf90808>) to automatically provision and configure your Fastly account. This script will:

- create an Acquia-specific service within your Fastly account,
- add your Acquia origin to the new service,
- add your end-user-facing domain to the new service, and
- add caching configurations to your service to optimize content delivery.

As soon as the configuration on Fastly's side is complete, you will receive email notification that you can change your CNAME records (</guides/basic-setup/cname-instructions-for-most-providers>) to point at the appropriate corresponding Fastly endpoints. As soon as the CNAME process is complete, you'll be ready to start using Acquia Cloud as an origin for Fastly services.

## § Amazon S3 (</guides/integrations/amazon-s3>)

Amazon S3 (<http://docs.aws.amazon.com/AmazonS3/latest/gsg/GetStartedWithS3.html>) can be used as an origin or as a private bucket.

## Using Amazon S3 as an origin

To make your S3 data buckets available through Fastly, follow the steps below.

1. Log in to the Fastly web interface and click the **Configure** link.
2. Click the **Options** button to open the **Manage service** menu and select **Create**. The Create a new service page appears.

## Create a new service

[Or use the setup guide](#)

Server address and port

:

★ Required

The IP address (or hostname) and port number for your origin server. To [enable TLS](#), use port 443.

▶ [Setting the server address and port](#)

Domain

★ Required

The domain name of your website.

▶ [Setting the domain name](#)

Description

★ Required

The name of your service, such as **My service**.

CREATE

CANCEL

3. Fill out the **Create a new service** fields as follows:




- In the **Server address and port** field, type `s3.amazonaws.com`. If you are using a non-standard S3 region (anything other than us-east), you must include that region in the server address field (e.g., `s3.us-west-2.amazonaws.com`). In the port field, you can type either `80` for HTTP or `443` for HTTPS.
- In the **Domain name** field, type the hostname you want to use as the URL (e.g., `cdn.example.com`).
- In the **Description** field, type any descriptive name for your service.


4. Click the **Create** button. A new service appears in the list of services available.

Now that the service is created, you will need to set the default host to `<yourbucket>.s3.amazonaws.com` by following the steps below:

1. From the service menu, select the new service you just created.
2. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.
3. Click the **Settings** tab. The Settings page appears.

[Return to service summary](#)

**www.example.com**   **Version 4**   
Show service ID **in development**


**ACTIVATE** **OPTIONS** 

|                            |   |   |   |                              |                              |
|----------------------------|---|---|---|------------------------------|------------------------------|
| <b>DOMAINS</b><br>1 Domain | <b>ORIGINS</b><br>1 Host<br>0 Health checks | <b>SETTINGS</b><br>0 Request settings<br>0 Cache settings | <b>CONTENT</b><br>0 Headers<br>0 Gzips<br>0 Responses | <b>LOGGING</b><br>1 Endpoint | <b>CUSTOM VCL</b><br>0 Files |
|----------------------------|---|---|---|------------------------------|------------------------------|

### Default PCI

This setting allows you to use the **beresp.pci** VCL variable to automatically cache content.

---

**PCI enabled: No** 

---

### Override host

Override the **host header** being sent to your origin regardless of the host used in the initial request. Only required if the domain your origin is expecting is different than those that Fastly hosts.

[When should I use an override host?](#)

---

**SPECIFY AN OVERRIDE HOST**

---

### Request settings

Request Settings are used to customize Fastly's request handling. When used with **Conditions** the Request Settings allow you to fine tune how specific types of requests are handled.

---

*There are no request settings.*

**CREATE YOUR FIRST REQUEST SETTING**

---

### Cache settings

Cache Settings controls how **caching is performed** on Fastly. When used with **Conditions**, the Cache Settings provide you with fine grain control over how long content persists in the cache.

---

*There are no cache settings.*


**CREATE YOUR FIRST CACHE SETTING**

---

### Fallback TTL

This setting is used only when there is no **Time to Live (TTL)** set in an object's header.

---

**TTL (seconds): 3600** 


---

4. In the **Override host** area, click the **Specify an override host** button. The Add an override host header window appears.

<https://docs.fastly.com/guides/aio>

322/345

5. Type the hostname of your S3 bucket. For example, `<yourbucket>.s3.amazonaws.com`.



6. Click the **Save** button.
7. Click the **Activate** button to deploy your configuration changes.

## Testing your results

By default, we create DNS mapping called **yourdomain.global.prod.fastly.net**. In the example above, it would be `cdn.example.com.global.prod.fastly.net`. Create a DNS alias for the domain name you specified (e.g., CNAME `cdn.example.com` to `global-nossl.fastly.net`).

Fastly will cache any content without an explicit `Cache-Control` header for 1 hour. You can verify whether you are sending any cache headers using `cURL`. For example:

```
$ curl -I opcode-full-stack.s3.amazonaws.com

HTTP/1.1 200 OK
x-amz-id-2: ZpzRp7Iwc6MJ8NtDEFGH12QBdk2CM1+RzVongQbhMp2f2ZyalkFsZd4qPaLMkS1h
x-amz-request-id: ABV5032583242618
Date: Fri, 18 Mar 2012 17:15:38 GMT
Content-Type: application/xml
Transfer-Encoding: chunked
Server: AmazonS3
```

In this example, no cache control headers are set so default TTL will be applied.

## Enhanced cache control

If you need more control over how different types of assets are cached (e.g., Javascript files, images) media check out our Amazon S3 configuration (</guides/tutorials/cache-control-tutorial>) in our Cache Control tutorial.

## Using an Amazon S3 private bucket

To use an Amazon S3 private bucket with Fastly, follow the instructions below.

### Before you begin

Be sure you've already made your S3 data buckets available to Fastly by pointing to the right S3 bucket and setting your origin to port 443. This needs to be done before authenticating.

Be sure you've got the AWS access key ID, AWS secret key ID, and AWS Bucket name on hand. The Amazon S3 Authorization header takes the following form:

```
Authorization: AWS &_AWSAccessKeyId_`:_Signature_`
```

From your developer Amazon account you will need the following information:

1. The **AWS access key ID** and **AWS secret access key**. The AWS secret access key is issued when you register. If you do not have or remember your AWS secret access key, create a new AWS access key ID. The AWS secret access key will be displayed before disappearing

again.

2. Your **AWS Bucket** name.

## Setting up Fastly to use an Amazon S3 private bucket

In order to use an Amazon S3 private bucket with Fastly, create two headers (/guides/basic-configuration/adding-or-modifying-headers-on-http-requests-and-responses), a Date header (for use with the authorization Signature) and an Authorization header.

### Create a Date header

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.
4. Click the **Content** tab. The Content page appears.
5. Click the **Create header** button. The Create a new header page appears.

## Create a new header

Learn more about this section in our [headers tutorial](#).

CONDITION

This will happen all the time unless you [Attach a condition](#)

Type / Action

Request

Set

The type of header and the action performed on it.

Destination

http.Date

★ Required

The name of the header that will be affected by the selected action. For example: **http.Content-Type**, **http.Set-Cookie**, **http.Via**, **http.Location**, or **http.Access-Control-Allow-Origin**.

Source

no

★ Required

New content for the header. Can be a static value (e.g. string or number) or a dynamic value (e.g. existing header or a GeoIP value). Please use quotes for string values.

Ignore if set

No

If switched to **Yes**, the action will not be performed if the header in Destination exists.

Priority

10

★ Required

The order in which the header rules execute within the condition.

Description

Date

★ Required

The name of your header, such as **My header**.

CREATE

CANCEL

6. Fill out the **Create a new header** fields as follows:

- From the **Type** menu, select **Request**, and from the **Action** menu, select **Set**.
- In the **Destination** field, type `http.Date`.
- In the **Source** field, type `now`.
- From the **Ignore if set** menu, select **No**.
- In the **Priority** field, type `10`.



- In the **Description** field, type `Date`.

7. Click the **Create** button. A new Date header appears on the Content page. You will use this later within the Signature of the Authorization header.

### Create an Authorization header

Next create the Authorization header with the specifications listed below.

1. Click the **Create header** button again to create another new header. The Create a new header page appears.

## Create a new header

Learn more about this section in our [headers tutorial](#).

**CONDITION** This will happen all the time unless you [Attach a condition](#)

**Type / Action** Request Set

The type of header and the action performed on it.

**Destination** http.Authorization ★ Required

The name of the header that will be affected by the selected action. For example: `http.Content-Type`, `http.Set-Cookie`, `http.Via`, `http.Location`, or `http.Access-Control-Allow-Origin`.

**Source** "AWS <AWS access key ID>:" digest.hmac\_sha1\_base64("<AWS secret key ★ Required

New content for the header. Can be a static value (e.g. string or number) or a dynamic value (e.g. existing header or a GeoIP value). Please use quotes for string values.

**Ignore if set** No

If switched to **Yes**, the action will not be performed if the header in Destination exists.

**Priority** 20 ★ Required

The order in which the header rules execute within the condition.

**Description** S3 Authorization ★ Required

The name of your header, such as **My header**.

CREATE
CANCEL

2. Fill out the **Create a new header** fields as follows:

- From the **Type** menu, select **Request**, and from the **Action** menu, select **Set**.
- In the **Destination** field, type `http.Authorization`.
- From the **Ignore if set** menu, select **No**.
- In the **Priority** field, type `20`.
- In the **Description** field, type `S3 Authorization`.

3. In the **Source** field, type the header authorization information using the following format:

```
"AWS <AWS access key ID>:" digest.hmac_sha1_base64("<AWS secret key ID>", if(req.request == "HEAD", "GET", req.request) LF LF LF
F req.http.Date LF "/"<AWS Bucket name>" req.url.path)
```

replacing `<AWS access key ID>`, `<AWS secret key ID>`, and `<AWS Bucket name>` with the information you gathered before you began. For example:

```
"AWS JKCAUEFV2ONFFOFMSSLA:" digest.hmac_sha1_base64("P2WPSu68BfI89j72vT+bXYZB7SjI0whT4whqt27", if(req.request == "HEAD", "GET", req.request) LF LF LF req.http.Date LF "/test123" req.url.path)
```

4. Click the **Create** button. The new Authorization header appears on the Content page.

### A detailed look at the Source field

So what's going on in the Source field of the Authorization header? Here's the basic format:

```
AWS<Access Key><Signature Function><key><message>
```

It tells us the following:

| Element                         | Description  |
|---------------------------------|--|
| <code>AWS</code>                | A constant placed before the access key. It's always AWS.  |
| <code>access key</code>         | The access key ID from your Amazon developer's account. We used <code>JKCAUEFV2ONFFOFMSSLA</code> in this example.   |
| <code>signature function</code> | The algorithm used to validate the key and message of the signature. We used <code>digest.hmac_sha1_base64(&lt;key&gt;, &lt;message&gt;)</code> in this example. |
| <code>key</code>                | The secret key ID from your Amazon developer's account. We used <code>P2WPSu68BfI89j72vT+bXYZB7SjI0whT4whqt27</code> in this example.                            |
| <code>message</code>            | The UTF-8 encoding of the StringToSign. See the table below for a break down of each portion of the message.   |

The message that's part of the Source field in the Authorization header takes on this basic format:

```
<HTTP-verb></n><Content-MD5></n><Content-Type></n><Date></n><CanonicalizedAmzHeader></n><CanonicalizedResource>
```

It tells us the following:

| Element                             | Description  |
|-------------------------------------|--|
| <code>HTTP-verb</code>              | The REST verb. We use <code>req.request</code> in this example. We rewrite HEAD to GET because Varnish does this internally before sending requests to origin. |
| <code>/n</code>                     | A newline indicator constant. It's always /n.  |
| <code>Content-MD5</code>            | The content-md5 header value, used as a message integrity check. It's often left blank. We use <code>LF</code> (line feed) in this example.                    |
| <code>Content-Type</code>           | The content-type header value, used to specify the MIME-type. It's often left blank. We use <code>LF</code> in this example.                                   |
| <code>Date</code>                   | The date and time stamp. We use <code>req.http.Date</code> (which we created first as a separate header in the steps above).                                   |
| <code>CanonicalizedAmzHeader</code> | The x-amz headers, which customize your S3 implementation. It's often left blank. We use <code>LF</code> in this example.                                      |
| <code>CanonicalizedResource</code>  | Your Amazon private bucket name. We use <code>"/test123"</code> in this example.   |

## Following redirects to S3 objects and caching S3 responses

With custom VCL, Fastly can follow redirects to S3 objects and cache the s3 response as well as the 301 or 302 response separately.

**❗ IMPORTANT:** The ability to [upload custom VCL code \(/guides/vcl/uploading-custom-vcl\)](/guides/vcl/uploading-custom-vcl) is disabled by default when you sign up. To enable this ability for your account, contact [support@fastly.com \(mailto:support@fastly.com\)](mailto:support@fastly.com) and request it.

Once the ability to upload custom VCL has been enabled, be sure to read our "How do I mix and match Fastly VCL with custom VCL? (/guides/vcl/mixing-and-matching-fastly-vcl-with-custom-vcl)" instructions. It's important to include the entire VCL boilerplate if you do not intend to override the Fastly default settings.

To configure Fastly to follow redirects to S3 objects, insert the following VCL snippets in your custom VCL:

### Within `vcl_recv`

```

sub vcl_recv {

    if (req.http.redir != "true") {
        set req.backend = Main-Origin;
    } else {
        set req.backend = s3_backend;
        set req.http.host = "s3.amazonaws.com";
    }

    #FASTLY recv

    if (req.request != "HEAD" && req.request != "GET" && req.request != "FASTLYPURGE") {
        return(pass);
    }

    return(lookup);
}

```

#### Within vcl\_deliver

```

sub vcl_deliver {

    if (resp.status == 302 || resp.status == 301) {
        set req.http.redir = "true";
        set req.url = regsub(resp.http.Location, "http://s3.amazonaws.com/(.*)$", "/\1");
        set req.http.Fastly-Force-Shield = "yes";
        restart;
    }

    #FASTLY deliver

    return(deliver);
}

```

Be sure to set the `Main-Origin` and `s3_backend` to the actual name of your backends in the service to which you're applying these redirects. You can find the exact names by reviewing your VCL; simply click on the VCL button at the top of the page while viewing the service.

Once you added these VCL snippets to your custom VCL, upload the VCL file and then activate the new version of your service to apply the changes.

## § Google Cloud Storage (/guides/integrations/google-cloud-storage)

Google Cloud Storage (<https://cloud.google.com/storage/>) (GCS) can be used as an origin server with your Fastly services once you set up and configure your GCS account and link it to a Fastly service. It can also be configured to use private content. This speeds up your content delivery and reduces your origin's workload and response times with the dedicated links between Google and Fastly's POPs.

★ **TIP:** Fastly offers a [Cloud Accelerator \(/guides/about-fastly-services/about-fastlys-cloud-accelerator\)](/guides/about-fastly-services/about-fastlys-cloud-accelerator) integration discount to any customer who uses GCS. Take advantage of this discount by emailing [salesgcp@fastly.com](mailto:salesgcp@fastly.com) (<mailto:salesgcp@fastly.com>) once you've completed the steps below. If you already have GCS configured as your origin and would like to take advantage of this discount, [contact support](mailto:support@fastly.com) (<mailto:support@fastly.com>).

## Using GCS as an origin server

To make your GCS data available through Fastly, follow the steps below.

### Setting up and configuring your GCS account

1. Sign up for Google Cloud Storage (<https://cloud.google.com/storage/docs/signup>) and start the basic setup.
2. Create a bucket (<https://cloud.google.com/storage/docs/getting-started-console>) to store your origin's data and remember the name. You'll need the bucket name to connect your GCS account to your Fastly service.

### Create a bucket

**Name** ?  
The bucket name must be unique across Cloud Storage.

**Storage class** ?  
Standard

**Location** ?  
United States

Privacy: Do not include sensitive information in the bucket name. Users cannot access your data without permission, but they can still try to access or create buckets to find out if the name exists.

Create Cancel

3. Add a file to the bucket and then make the file public.

Browser UPLOAD FILES CREATE FOLDER REFRESH

Buckets / test123 Filter by prefix...

| <input type="checkbox"/> | Name               | Size     | Type       | Last modified | Share publicly                                  |   |
|--------------------------|--------------------|----------|------------|---------------|---|---|
| <input type="checkbox"/> | horseshoe-bend.jpg | 688.8 KB | image/jpeg | 2 minutes ago | <input checked="" type="checkbox"/> Public link | ⋮ |
| <input type="checkbox"/> | kirkpano2.jpg      | 359.3 KB | image/jpeg | 2 minutes ago | <input type="checkbox"/>                        | ⋮ |

## Creating a new origin in your Fastly service for your GCS account

Link your GCS account to a Fastly service following the steps below.

1. Log in to the Fastly web interface and click the **Configure** link.
2. Create a new service (/guides/basic-setup/working-with-services) if you don't already have one set up.
3. From the service menu, select the appropriate service.
4. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.
5. Click the **Origins** tab. The Origins page appears.
6. Click the **Create host** button. The Create a new host page appears.

## Create a new host

**CONDITION** This will happen all the time unless you [Attach a condition](#)

**Address**  :  ★ Required

The IP address (or hostname) and port number for your origin server. To [enable TLS](#), use port 443.

**Shielding**

The shield POP designated to reduce inbound load on your origins by serving cached data. Learn more about [POP request handling](#) and the [caveats of shielding](#).

**Health check**

[+ Create a new health check](#)

A small request assigned to [test the state of your origin](#).

**Auto load balance**

Enables Fastly to [load balance between two or more backend hosts](#).

[► Tips](#)

**Description**  ★ Required

The name of your origin, such as **My origin server**.

7. Fill out the **Create a new host** fields as follows:

- In the **Address** field, type the address of your secure server (for example, `origin.example.com`), and in the port field, type `443`.
- From the **Shielding** menu, select an available interconnect location from the list of shielding locations. See our information on interconnect locations below for more details.
- Leave the **Health check** and **Auto load balance** menus set to their default values.
- In the **Description** field, type the name of your server (for example, `Google Cloud Storage`).

8. Click the **Create** button. A new host appears on the Origins page.

## Interconnect locations

The following interconnects allow you to establish direct links with Google's edge network when you choose your shielding location. By selecting one of the locations listed below, you will be eligible to receive discounted pricing (<https://cloud.google.com/interconnect/cdn-interconnect#pricing>) from Google CDN Interconnect for traffic traveling from Google Cloud Platform to Fastly's network. Most customers select interconnects closest to their origin.

Interconnects exist in the following locations within North America:

- Ashburn
- Atlanta
- Chicago
- Dallas
- Los Angeles
- New York City
- San Jose
- Seattle

Interconnects outside of North America exist in:

- Amsterdam

- Frankfurt
- Hong Kong
- London
- Singapore
- Tokyo

Review our caveats of shielding (/guides/performance-tuning/shielding#caveats-of-shielding) and select an interconnect accordingly.

## Setting the Cache-Control header for your GCS bucket

GCS performs its own caching, which may complicate efforts to purge cache. To avoid potential problems, we recommend using the gsutil (https://cloud.google.com/storage/docs/gsutil\_install) command line utility to set the Cache-Control header for one or more files in your GCS bucket:




```
gsutil setmeta -h "Cache-Control: max-age=0, s-maxage=86400" gs://bucket/*.html
```


Replace `bucket` in the example above with your GCS bucket's name. Note that `max-age` should instruct GCS to cache your content for zero seconds, and Fastly to cache your content for one day. See Google's setmeta docs (https://cloud.google.com/storage/docs/gsutil/commands/setmeta) for more information.

## Setting the default host for your service to your GCS bucket

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.
4. Click the **Settings** tab. The Settings page appears.

[Return to service summary](#)

**www.example.com**   **Version 4**   
Show service ID **in development**


**ACTIVATE** **OPTIONS** 

|                            |   |   |   |                              |                              |
|----------------------------|---|---|---|------------------------------|------------------------------|
| <b>DOMAINS</b><br>1 Domain | <b>ORIGINS</b><br>1 Host<br>0 Health checks | <b>SETTINGS</b><br>0 Request settings<br>0 Cache settings | <b>CONTENT</b><br>0 Headers<br>0 Gzips<br>0 Responses | <b>LOGGING</b><br>1 Endpoint | <b>CUSTOM VCL</b><br>0 Files |
|----------------------------|---|---|---|------------------------------|------------------------------|

### Default PCI

This setting allows you to use the **beresp.pci** VCL variable to automatically cache content.

---

**PCI enabled: No** 

---

### Override host

Override the **host header** being sent to your origin regardless of the host used in the initial request. Only required if the domain your origin is expecting is different than those that Fastly hosts.

[When should I use an override host?](#)

---

**SPECIFY AN OVERRIDE HOST**

---

### Request settings

Request Settings are used to customize Fastly's request handling. When used with **Conditions** the Request Settings allow you to fine tune how specific types of requests are handled.

---

*There are no request settings.*

**CREATE YOUR FIRST REQUEST SETTING**

---

### Cache settings

Cache Settings controls how **caching is performed** on Fastly. When used with **Conditions**, the Cache Settings provide you with fine grain control over how long content persists in the cache.

---

*There are no cache settings.*


**CREATE YOUR FIRST CACHE SETTING**

---

### Fallback TTL


This setting is used only when there is no **Time to Live (TTL)** set in an object's header.

---

**TTL (seconds): 3600** 

---

5. In the **Override host** area, click the **Specify an override host** button. The Add an override host header window appears.



6. Type the name of the override host for this service. The name you type should match the name of the bucket you created in your GCS account and will take the format `<your bucket name>.storage.googleapis.com`. For example, if your bucket name is `test123`, your override host name would be `test123.storage.googleapis.com`.
7. Click the **Save** button. A new override host appears on the Settings page.


## Changing the default TTL for your GCS bucket

If you want to change the default TTL for your GCS bucket, if at all, keeping the following in mind:

- Your GCS account controls the default TTL for your GCS content. GCS currently sets the default TTL to 3600 seconds. Changing the default TTL will not override the default setting in your GCS account.
- To override the default TTL set by GCS from within the Fastly web interface, create a new cache setting (/guides/performance-tuning/controlling-caching) and enter the TTL there.
- To override the default TTL in GCS, download the gsutil tool ([https://cloud.google.com/storage/docs/gsutil\\_install](https://cloud.google.com/storage/docs/gsutil_install)) and then change the cache-control headers (<https://cloud.google.com/storage/docs/gsutil/addlhelp/WorkingWithObjectMetadata#cache-control>) to delete the default TTL or change it to an appropriate setting.

## Creating new domains for GCS to respond to

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.
4. Click the **Domains** tab. The Domains page appears.



5. Click the **Create domain** button. The Create a new domain page appears.



## Create a new domain

**Domain Name**

The domain name of your website.

▶ [Setting the domain name](#)

▶ [What if I'm using apex domains?](#)

**Description**

An optional comment that describes your domain.

★ Required

CREATE

CANCEL

6. In the **Domain Name** field, type the name users will type in their browsers to access your site.
7. Click the **Create** button. A new domain appears on the Domains page.
8. Because GCS responds to different hostnames than your Fastly service, click the **Create domain** button to create a second domain. The Domains page appears. Follow the domain creation steps immediately above.
9. In the **Domain Name** field of the second domain you create, type the same value as the default host you created earlier (e.g., `<your bucket name>.storage.googleapis.com`) and click the **Create** button. A new domain appears on the Domains page. Shielding POPs need this additional domain so they can route requests correctly. (See Caveats of Shielding (/guides/performance-tuning/shielding#caveats-of-shielding) for more information.)
10. Click the **Activate** button to deploy your configuration changes.

Once you deploy, you can use `http://<domain>.global.prod.fastly.net/<filename>` to access the files you uploaded.

## Using GCS with private objects

To use Fastly with GCS private objects, be sure you've already made your GCS data available to Fastly by pointing to the right GCS bucket, then follow the steps below.

### Setting up interoperable access

By default, GCS authenticates requests using OAuth2, which Fastly does not support. To access private objects on GCS, your project must have HMAC authentication (<https://cloud.google.com/storage/docs/gsutil/addlhelp/CredentialTypesSupportingVariousUseCases#supported-credential-types>) enabled and interoperable storage access keys (an "Access Key" and "Secret" pair) created. Do this by following the steps below.



1. Open the Google Cloud Platform console and select the appropriate project.
2. Click **Settings**. The Settings appear with the Project Access controls highlighted.
3. Click the **Interoperability** tab. The Interoperability API access controls appear.
4. If you have not set up interoperability before, click **Enable interoperability access**.
5. Click **Make** `<PROJECT-ID>` **your default project** for interoperable access. If that project already serves as the default project, that information appears instead.

Project Access
On-Demand I/O
Interoperability

The Interoperability API lets you use HMAC authentication and lets Cloud Storage interoperate with tools written for other cloud storage systems. Turn on this API only if you require interoperable access for the current user. This API is enabled per project member, not per project. Each member can set a default project and maintain their own access keys. [Learn more](#)

**Default project for interoperable access**

The default project is used with the Interoperability API for all create bucket and list bucket requests.




is your default project for interoperable access

**Interoperable storage access keys**

Use an access key to authenticate yourself when making requests to Cloud Storage. The key is linked to your Google user account. [Learn more](#)

Create a new key

- Click **Create a new key**. An access key and secret code appear.

| Access Key           | Secret                                   |   |
|----------------------|--|---|
| GOOG2E7G4GU46REOYYCV | ROGiAZq5zW4x/hPriINtAo9d9Giuq+fcgSgOvW9y |  |

- Save the access key and secret code that appear. You'll need these later when you're creating an authorization header.

## Setting up Fastly to use GCS private content

To use GCS private content with Fastly, create two headers (/guides/basic-configuration/adding-or-modifying-headers-on-http-requests-and-responses), a Date header (required Authorization Signature) and an Authorization header.

### Creating a Date header

- Log in to the Fastly web interface and click the **Configure** link.
- From the service menu, select the appropriate service.
- Click the **Edit configuration** button and then select **Clone active**. The service version page appears.
- Click the **Content** tab. The Content page appears.
- Click the **Create header** button. The Create a new header page appears.

## Create a new header

Learn more about this section in our [headers tutorial](#).

### CONDITION

This will happen all the time unless you [Attach a condition](#)

Type / Action

Request

Set

The type of header and the action performed on it.

Destination

http.Date

★ Required

The name of the header that will be affected by the selected action. For example: **http.Content-Type**, **http.Set-Cookie**, **http.Via**, **http.Location**, or **http.Access-Control-Allow-Origin**.

Source

no

★ Required

New content for the header. Can be a static value (e.g. string or number) or a dynamic value (e.g. existing header or a GeoIP value). Please use quotes for string values.

Ignore if set

No

If switched to **Yes**, the action will not be performed if the header in Destination exists.

Priority

10

★ Required

The order in which the header rules execute within the condition.

Description

Date

★ Required

The name of your header, such as **My header**.

CREATE

CANCEL

6. Fill out the **Create a new header** fields as follows:

- From the **Type** menu, select **Request**, and from the **Action** menu, select **Set**.
- In the **Destination** field, type `http.Date`.
- In the **Source** field, type `now`.
- From the **Ignore if set** menu, select **No**.
- In the **Priority** field, type `10`.
- In the **Description** field, type `Date`.

7. Click the **Create** button. A new Date header appears on the Content page. You will use this later within the Signature of the Authorization header.

### Creating an Authorization header

1. Click the **Create header** button again to create another new header. The Create a new header page appears.

## Create a new header

Learn more about this section in our [headers tutorial](#).

### CONDITION

This will happen all the time unless you [Attach a condition](#)

Type / Action

The type of header and the action performed on it.

Destination  ★ Required

The name of the header that will be affected by the selected action. For example: **http.Content-Type**, **http.Set-Cookie**, **http.Via**, **http.Location**, or **http.Access-Control-Allow-Origin**.

Source

```
AWS<access key><signature function><key><message>
```

It tells us the following:

| Element                         | Description  |
|---------------------------------|--|
| <code>AWS</code>                | A constant placed before the access key. It's always AWS.  |
| <code>access key</code>         | The access key ID from your GCS developer's account. We used <code>GOOGQORE5WOJJHLXH6OD</code> in this example.  |
| <code>signature function</code> | The algorithm used to validate the key and message of the signature. We used <code>digest.hmac_sha1_base64(&lt;key&gt;, &lt;message&gt;)</code> in this example. |
| <code>key</code>                | The secret key ID from your GCS developer's account. We used <code>oQb0hdmaxF0c5UmC6F833Cde0+ghR8gsr7CCnX62</code> in this example.                              |
| <code>message</code>            | The UTF-8 encoding of the StringToSign. See the table below for a break down of each portion of the message.   |

The message that's part of the Source field in the Authorization header takes on this basic format:

```
<HTTP-verb><\n><Content-MD5>\n<Content-Type><\n><Date><\n><CanonicalExtensionHeaders><\n><CanonicalizedResource>
```

It tells us the following:

| Element                                | Description  |
|--|--|
| <code>HTTP-verb</code>                 | The REST verb. We use <code>req.request</code> in this example.  |
| <code>\n</code>                        | A newline indicator constant. It's always \n.  |
| <code>Content-MD5</code>               | The content-md5 header value, used as a message integrity check. It's often left blank. We use <code>LF</code> (line feed) in this example.          |
| <code>Content-Type</code>              | The content-type header value, used to specify the MIME-type. It's often left blank. We use <code>LF</code> in this example.                         |
| <code>Date</code>                      | The date and time stamp. We use <code>req.http.Date</code> (which we created first as a separate header in the steps above).                         |
| <code>CanonicalExtensionHeaders</code> | The x-amz- or x-goog- headers, which customize your GCS implementation. It's often left blank. We use <code>LF</code> in this example.               |
| <code>CanonicalizedResource</code>     | Your GCS resource path name. We're concatenating GCS bucket name <code>"/test123"</code> with object path <code>req.url.path</code> in this example. |

## § Google Compute Engine (/guides/integrations/google-compute-engine)

Google Compute Engine (GCE) lets you create and run a virtual machine (VM) on the Google infrastructure. The VM can be used as an origin server with your Fastly service once you set up and configure your VM instance and link your instance to a Fastly service.

★ **TIP:** Fastly offers a [Cloud Accelerator \(/guides/about-fastly-services/about-fastlys-cloud-accelerator\)](/guides/about-fastly-services/about-fastlys-cloud-accelerator) integration discount to any customer who uses GCE. Take advantage of this discount by emailing [salesgcp@fastly.com](mailto:salesgcp@fastly.com) (<mailto:salesgcp@fastly.com>) once you've completed the steps below. If you already have GCE configured as your origin and would like to take advantage of this discount, [contact support](mailto:support@fastly.com) (<mailto:support@fastly.com>).

## Creating and setting up your GCE instance

1. Sign up for Google Compute Engine (<https://cloud.google.com/compute/docs/>) and start the basic set up. If you are already signed up and at your dashboard, click the **Get started** link in the Try Compute Engine area.
2. Create or select a project to hold your origin's data.

Select or create a project

The Google Developers Console uses projects to manage resources. To use Google Compute Engine, select an existing project or create a new one.

Select a project

Create a new project

Project name ?  
My test project 123

Your project ID will be my-test-project-123-1218 [Edit](#)

[Hide advanced options...](#)

App Engine location ?  
us-central

Create

3. Click **Create instance** to set up your VM. You can set up your instance using either Windows or Linux.

Compute Engine  
VM instances

Compute Engine lets you use virtual machines that run on Google's infrastructure. You can choose from micro-VMs to large instances running Debian, Windows, or other standard images. Create your first VM instance or try the quickstart to build a sample app.

Create instance or Take the quickstart

4. Fill in the necessary fields and click **Create**. When making your firewall selection, select either **Allow HTTPS traffic** (port 443) or **Allow HTTP traffic** (port 80); you will use one of those ports when you create your new origin in your Fastly service. If you select HTTPS traffic, you need to configure the VM to respond on port 443 with a valid TLS certificate.

VM instances [CREATE INSTANCE](#) [CREATE INSTANCE GROUP](#) [RESET](#)

Filter by label or name

CPU utilization 1 hour 6 hours

CPU

% CPU

100  
80  
60  
40  
20

Feb 26, 11:45 AM Feb 26, 12:00 PM Feb 26, 12:15 PM

☒ CPU: 0

| <input type="checkbox"/> Name ^                | Zone          | Disk       | Network | In use by | External IP | Connect |
|--|---------------|------------|---------|-----------|-------------|---------|
| <input checked="" type="checkbox"/> instance-1 | us-central1-b | instance-1 | default |           | 867.53.0.9  | SSH     |

5. Make note of the following information for when you create your new origin in your Fastly service:
- The instance's IP address (located in the External IP column at the bottom of the page). You'll use this in the **Address** field when you create your new origin.
  - The zone you are using (located in the Zone column at the bottom of the page). You'll use this to guide your selection of an appropriate shielding location for your origin.

# Creating a new origin in your Fastly service for your GCE account

Link your GCE account to a Fastly service following the steps below.

1. Log in to the Fastly web interface and click the **Configure** link.
2. Create a new service (/guides/basic-setup/working-with-services) if you don't already have one set up.
3. From the service menu, select the appropriate service.
4. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.
5. Click the **Origins** tab. The Origins page appears.
6. Click the **Create host** button. The Create a new host page appears.

## Create a new host

CONDITION

This will happen all the time unless you [Attach a condition](#)

Address

: 
★ Required

The IP address (or hostname) and port number for your origin server. To [enable TLS](#), use port 443.

Shielding

The shield POP designated to reduce inbound load on your origins by serving cached data. [Learn more about POP request handling](#) and the [caveats of shielding](#).

Health check

[+ Create a new health check](#)

A small request assigned to [test the state of your origin](#).

Auto load balance

Enables Fastly to [load balance between two or more backend hosts](#).

[► Tips](#)

Description

★ Required

The name of your origin, such as [My origin server](#).

7. Fill out the **Create a new host** fields as follows:

- In the **Address** field, type the IP address of your server, and in the port field, type  or . This should match the port that you selected in the GCE interface.
- From the **Shielding** menu, select an available interconnect location from the list of shielding locations. See our information on interconnect locations below for more details.
- Leave the **Health check** and **Auto load balance** menus set to their default values.
- In the **Description** field, type the name of your server (for example, .

8. Click the **Create** button. A new host appears on the Origins page.

In this example, no cache control headers are set so the default TTL will be applied.

## Interconnect locations

The following interconnects allow you to establish direct links with Google's edge network when you choose your shielding location. By selecting one of the locations listed below, you will be eligible to receive discounted pricing (<https://cloud.google.com/interconnect/cdn-interconnect#pricing>) from Google CDN Interconnect for traffic traveling from Google Cloud Platform to Fastly's network. Most customers select interconnects closest to their origin.

Interconnects exist in the following locations within North America:

- Ashburn
- Atlanta
- Chicago
- Dallas
- Los Angeles
- New York City
- San Jose
- Seattle

Interconnects outside of North America exist in:

- Amsterdam
- Frankfurt
- Hong Kong
- London
- Singapore
- Tokyo

Review our caveats of shielding (/guides/performance-tuning/shielding#caveats-of-shielding) and select an interconnect accordingly.

## Creating new domains for GCE to respond to

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.
4. Click the **Domains** tab The Domains page appears.



5. Click the **Create domain** button. The Create a new domain page appears.



## Create a new domain

Domain Name

The domain name of your website.

[▶ Setting the domain name](#)  
[▶ What if I'm using apex domains?](#)

Description

An optional comment that describes your domain.

★ Required

CREATE

CANCEL

6. In the **Domain Name** field, type the name that users will type in their browsers to access your site.
7. Click the **Create** button. The new domain appears on the Domains page.
8. Click the **Activate** button to deploy your configuration changes.

## Creating a CNAME record

You can now test your configuration (</guides/basic-configuration/testing-setup-before-changing-domains>). In the example above, your domain would appear as `www.example.com.global-ssl.fastly.net`. After you test and you're satisfied with the results, create a CNAME record (</guides/basic-setup/adding-cname-records>) for your domain (e.g., `www.example.com`) pointing to `global-ssl.fastly.net`.

## § PerimeterX Bot Defender (</guides/integrations/perimeterx-bot-defender>)

Fastly provides direct integration between PerimeterX Bot Defender (<https://www.perimeterx.com/bot-defender>) and Fastly edge servers. By placing a snippet of JavaScript (or HTML5) on your site and custom VCL (</guides/vcl/guide-to-vcl>) directly into your Fastly service configuration, this integration allows you to gather behavioral data and statistics that may help you do things like detect invalid traffic and mitigate automated web attacks.

❗ **IMPORTANT:** This feature is part of a limited availability release. For more information, see our [product and feature life cycle](/guides/customer-support/product-and-feature-life-cycles#limited-availability) (</guides/customer-support/product-and-feature-life-cycles#limited-availability>) descriptions.

## How to get started

Integration with PerimeterX Bot Defender requires an account with PerimeterX. Once you have this account set up, contact your Fastly account manager or email [sales@fastly.com](mailto:sales@fastly.com) (<mailto:sales@fastly.com>) to begin the integration process with Fastly. We'll work with you to configure your service to include the required code to enforce bot mitigation policies.

This article describes an integration with a service provided by a third party. Please see our note on integrations (<https://docs.fastly.com/guides/integrations/>).

## § Zencoder (</guides/integrations/zencoder>)

Zencoder provides HLS output that Fastly can cache and deliver. To start, you'll need to create a Zencoder job that outputs live HLS to S3. Zencoder provides documentation on this setup process (<https://app.zencoder.com/docs/guides/getting-started/live-s3-quickstart>).

Next, configure a new Fastly service using the S3 bucket that you are using with Zencoder as an origin. You can refer to our documentation on setting up S3 as an origin (</guides/integrations/amazon-s3>).

Once this is complete, you can serve your HLS streams through Fastly as you would any other content.

• [Guides \(/guides/\)](/guides/) > [Online video streaming \(/guides/online-video\)](/guides/online-video/) > [Live streaming \(/guides/live-streaming/\)](/guides/live-streaming/)

## § Configuration guidelines for live streaming (/guides/live-streaming/configuration-guidelines-for-live-streaming)

The Fastly network can deliver live streams for any HTTP streaming technology (/guides/about-fastly-services/about-fastlys-video-caching-service), archived or recorded, on any public or private cloud storage service. When configuring VCL (/guides/vcl/uploading-custom-vcl) to deliver live streams, we recommend following these guidelines, which Customer Support (mailto:support@fastly.com) can help you with.

### Configure shielding

Configure shielding (/guides/performance-tuning/shielding) by designating a specific shield POP for your origin to ensure live streams remain highly available within the Fastly network. If your setup includes primary and alternate origins (e.g., for high profile live streams), be sure to select a shield POP, close to each origin, one for each origin you define.

### Set different caching TTLs

Set manifest file TTLs to less than half of the video segment duration, typically 1-2 seconds for 5-second video segments. For long DVRs and live-to-VID transitions, set segment TTLs longer on shields and shorter on edge POPs such that they are served from memory (that is, less than 3600s).

### Configure Streaming Miss

Configure Streaming Miss (/guides/performance-tuning/improving-caching-performance-with-large-files#streaming-miss) to reduce the time clients (players) must wait to begin downloading streams when Fastly's edge servers must fetch content from your origin.

### Enable TCP optimizations

Enable TCP tuning optimizations (e.g., CWND size) between cache servers and clients. For example, consider setting the CWND value higher, between 40-50. Enabling TCP optimizations allows clients to better estimate network bandwidth and thereby select the ideal rendition and bitrate (/guides/on-the-fly-packaging/adaptive-bitrate-playback-url-guidelines) for the best playback quality.

The following VCL sample may help you implement this:

```
sub vcl_deliver {
  # increase init cwnd
  if (client.requests == 1) {
    set client.socket.cwnd = 45;
  }

  #FASTLY deliver
  return(deliver);
}
```

### Configure origin timeouts

Set appropriate origin timeouts (/guides/debugging/changing-connection-timeouts-to-your-origin) to ensure new live video chunks and segments are fetched in a timely manner. Start by setting the First Byte (ms) and Between Bytes (ms) timeout values to less than half of the segment duration and then adjust the values accordingly to optimize them for your origin. For example, for a live stream configured with a 5s segment duration, you could start by setting both timeouts to 2000ms.

### Consider setting up failover (fallback) origins

Consider configuring your VCL to allow your origins to failover (/guides/performance-tuning/failover-configuration) from high-profile primary streams to alternate streams in case of encoder failures or other issues (e.g., high resource utilization).

### Configure real-time log streaming

For troubleshooting and debugging any end-to-end live streaming latency issues (also known as "hand-waving latency" or "glass latency"), configure real-time log streaming (/guides/streaming-logs/) and include TCP connection, caching, and different time-related metrics (/guides/vcl/date-and-time-related-vcl-features) in `vcl_log`. For example:

- `fastly_info.state` (cache hits or missess)
- `client.socket.tcpi_rtt` (client round trip time)
- `time.to_first_byte` (time from client request to the first byte being received)

- `time.elapsed` (time since the request started)

These metrics can help you analyze throughput and may help you determine reasons a video player might switch quality levels during ABR playback (/guides/on-the-fly-packaging/adaptive-bitrate-playback-url-guidelines).

## Take advantage of surrogate key purging

All video segments and the manifest for a live stream can be purged using a single API call by using Fastly’s surrogate key feature (/guides/purging/getting-started-with-surrogate-keys).

## Manage live-to-VOD smoothly

Most encoders generate a separate video manifest when making the same live stream available for VOD. If your VOD manifest has the same URL as the live one, purge the live stream video manifest or wait for the caches to invalidate (as they will be set with low TTLs). If your setup archives the live stream as progressive mp4s, consider delivering them using Fastly’s OTFP service (/guides/about-fastly-services/about-fastlys-onthefly-packaging-service).

**NOTE: Wowza integrations.** When configuring your Wowza origin server, be sure to select the [Live HTTP Origin \(https://www.wowza.com/forums/content.php?2227-How-to-configure-a-live-stream-repeater\)](https://www.wowza.com/forums/content.php?2227-How-to-configure-a-live-stream-repeater) application type. If you select Live Edge, Wowza will always return a unique URL for manifest requests, resulting in extremely low cache hit.

- Guides (/guides/) > Online video streaming (/guides/online-video) > On-the-fly packaging (/guides/on-the-fly-packaging/)

# § Adaptive bitrate playback URL guidelines (/guides/on-the-fly-packaging/adaptive-bitrate-playback-url-guidelines)

Fastly's On-the-Fly Packaging (OTFP) service (/guides/about-fastly-services/about-fastlys-onthefly-packaging-service) supports any directory structure you might use to store different quality levels of a video. To construct adaptive bitrate (ABR) playback URLs for a video, make directory paths to that video unique. Ensure all the files associated with a particular video (e.g., quality levels, subtitles) exist under a single directory.

For example, say you had a video called Example Video. Assuming you had multiple quality levels and associated files for Example Video, the following directory structure would provide the best start to constructing ABR playback URLs:

| Directory path example                 | Description  |
|--|--|
| /foo/bar/example-video/                | Base folder unique to this video                                 |
| /foo/bar/example-video/480p_30fps.mp4  | Quality level 480p with 30 frames per sec with audio             |
| /foo/bar/example-video/720p_30fps.mp4  | Quality level 720p with audio with 30 frames per sec with audio  |
| /foo/bar/example-video/720p_60fps.mp4  | Quality level 720p with audio with 60 frames per sec with audio  |
| /foo/bar/example-video/1080p_30fps.mp4 | Quality level 1080p with audio with 30 frames per sec with audio |
| /foo/bar/example-video/1080p_60fps.mp4 | Quality level 1080p with audio with 60 frames per sec with audio |
| /foo/bar/example-video/4k_30fps.mp4    | Quality level 4k with audio with 30 frames per sec with audio    |

With this directory structure, the ABR playback URL for all videos in the base directory would follow this template:

http://example.com/path/to/dir/<video\_id>/<quality\_file1\_name\_wo\_ext>,<quality\_file2\_name\_wo\_ext>,...,<quality\_fileN\_name\_wo\_ext>/master.<f4m|m3u8|mpd>

For example, the ABR playback URLs for Example Video in every format would be:

| Format    | Example URL  |
|-----------|--|
| HDS       | http://example.com/foo/bar/example-video/480p_30fps,720p_30fps,720p_60fps,1080p_30fps,1080p_60fps,4k_30fps/master.f4m  |
| HLS       | http://example.com/foo/bar/example-video/480p_30fps,720p_30fps,720p_60fps,1080p_30fps,1080p_60fps,4k_30fps/master.m3u8 |
| MPEG-DASH | http://example.com/foo/bar/example-video/480p_30fps,720p_30fps,720p_60fps,1080p_30fps,1080p_60fps,4k_30fps/master.mpd  |

You can reduce the duplication in ABR playback URLs separating out the repeated prefix and suffix info as follows:

<filename\_prefix><filename\_variable><filename\_suffix\_wo\_ext>.mp4

and the template would change to one of the following:

```
http://example.com/path/to/dir/<video_id>/<filename_prefix><quality_file1_variable_name_wo_ext>,<quality_file2_variable_name_wo_ext>,...,<quality_fileN_variable_name_wo_ext>,<filename_suffix_wo_ext>/master.<f4m|m3u8|mpd>
```

```
http://example.com/path/to/dir/<video_id>/<filename_prefix><quality_file1_variable_name_wo_ext>,<quality_file2_variable_name_wo_ext>,...,<quality_fileN_variable_name_wo_ext>/master.<f4m|m3u8|mpd>
```

```
http://example.com/path/to/dir/<video_id>/<quality_file1_variable_name>,<quality_file2_variable_name>,...,<quality_fileN_variable_name>,<filename_suffix_wo_ext>/master.<f4m|m3u8|mpd>
```

**⚠ IMPORTANT:** To use [token authentication \(/guides/tutorials/enabling-token-authentication\)](/guides/tutorials/enabling-token-authentication) with ABR manifest URLs, special modifications must be made using [custom VCL \(/guides/vcl/uploading-custom-vcl\)](/guides/vcl/uploading-custom-vcl). Contact [support@fastly.com \(mailto:support@fastly.com\)](mailto:support@fastly.com) for assistance.

## § Collecting OTFP metrics (</guides/on-the-fly-packaging/collecting-otfp-metrics>)

Fastly allows you to collect and process On-the-Fly Packaging (OTFP) service (</guides/about-fastly-services/about-fastlys-onthefly-packaging-service>) metrics for analysis using a combination of custom VCL updates and specific log streaming settings. Once you've set up OTFP metrics collection through remote log streaming you can use any of a number of third-party and open source software options to aggregate your logging data for visualization and further analysis.

### Upload custom VCL

1. If you have not already done so, request the ability to upload custom VCL (</guides/vcl/uploading-custom-vcl>) code be enabled for your account.
2. Review the caveats of mixing and matching Fastly VCL with custom VCL (</guides/vcl/mixing-and-matching-fastly-vcl-with-custom-vcl>).
3. Add the following custom VCL to your Fastly VCL:

```
sub vcl_deliver {
  # Identify Request type
  if (req.url.ext ~ "m3u8|ts|aac|webvtt") {
    set resp.http.Otfp-Format = "HLS";
  } else if (req.url.ext ~ "mpd|m4s") {
    set resp.http.Otfp-Format = "DASH";
  } else {
    set resp.http.Otfp-Format = "OTHER";
  }

  # Extract name-value pairs Otfp Info header
  if (resp.http.X-Fastly-Otfp-Info) {
    set resp.http.Otfp-SS = regsub(resp.http.X-Fastly-Otfp-Info, ".*ss=(\S+).*", "\1");
    set resp.http.Otfp-SL = regsub(resp.http.X-Fastly-Otfp-Info, ".*sl=(\S+).*", "\1");
    set resp.http.Otfp-VL = regsub(resp.http.X-Fastly-Otfp-Info, ".*vl=(\S+).*", "\1");

    # Resolution (rs name-value) not available for audio-only segments
    if (resp.http.X-Fastly-Otfp-Info ~ ".*rs=(\S+).*") {
      set resp.http.Otfp-RS = re.group.1;
    } else {
      set resp.http.Otfp-RS = "-";
    }
  }
}

#FASTLY deliver

return(deliver);
}
```

### Create a logging endpoint

Follow the instructions to set up remote log streaming (</guides/streaming-logs/setting-up-remote-log-streaming>) for your account and when creating your specific logging endpoint, set the **Format String** field to the following:

```
%h now.sec %r %s %b resp.http.Otfp-Format resp.http.Otfp-SS resp.http.Otfp-SL resp.http.Otfp-VL resp.http.Otfp-RS
```

### Control log file timing with a logging endpoint condition

To avoid excess log files, consider attaching a condition to the logging endpoint so logs are only sent when video segments are requested so that logging specifically exclude those files sent from Fastly's Origin Shield.

1. Log in to the Fastly web interface and click the **Configure** link.
2. From the service menu, select the appropriate service.
3. Click the **Edit configuration** button and then select **Clone active**. The service version page appears.
4. Click the **Logging** tab. The logging page appears.
5. In the list of logging endpoints, find the endpoint you enabled when setting up remote log streaming (/guides/streaming-logs/setting-up-remote-log-streaming), then click **Attach a condition**. The Create a new condition window appears.
6. Fill out the **Create a new condition** window as follows:
  - From the **Type** menu, select **Response**.
  - In the **Name** field, type a human-readable name for the condition.
  - In the **Apply if** field, type `resp.http.X-Fastly-Otfp-Info && !req.http.Fastly-FF`.
7. Click **Save and apply to**.

## Analyze logging data

In addition to any Varnish variable (<https://www.varnish-cache.org/docs/2.1/reference/vcl.html#variables>), and a variety of Fastly's extensions to VCL (/guides/vcl/guide-to-vcl#fastly-39-s-vcl-extensions), log files include the following video-specific fields:

- `ss` - video segment start presentation time in seconds
- `sl` - video segment duration in seconds
- `vl` - video duration in seconds
- `rs` - video track display resolution in pixels

You can use these fields to run queries for analysis and use what you discover to refine your video delivery settings.