

Routing with EdgeJS

The `@layer0/core` package provides a JavaScript API for controlling routing and caching from your code base rather than a CDN web portal. Using this *EdgeJS* approach allows this vital routing logic to be properly tested, reviewed, and version controlled, just like the rest of your application code.

Using the Router, you can:

- Proxy requests to upstream sites
- Send redirects from the network edge
- Render responses on the server using Next.js, Nuxt.js, Angular, or any other framework that supports server side rendering.
- Alter request and response headers
- Send synthetic responses
- Configure multiple destinations for split testing

Configuration

You define routes for Edgio using the `routes.js` file.

Before continuing, if you have not already initialized your project with Edgio, do so using the instructions in [Web CDN](#).

The `routes.js` file should export an instance of `@layer0/core/router/Router`:

JavaScript ./routes.js

Copy

```
1const { Router } = require('@layer0/core/router')
2
3module.exports = new Router()
```

Declare Routes

Declare routes using the method corresponding to the HTTP method you want to match.

JavaScript ./routes.js

Copy

```
1const { Router } = require('@layer0/core/router')
2
3module.exports = new Router().get('/some-path', ({ cache, proxy }) => {
4 // handle the request here
5})
```

All HTTP methods are available:

- get
- put
- post

- patch
- delete
- head

To match all methods, use match:

```
JavaScript ./routes.js
```

```
Copy
```

```
1const { Router } = require('@layer0/core/router')
2
3module.exports = new Router().match('/some-path', ({ cache, proxy }) => {
4  // handle the request here
5})
```

Route Execution

When Edgio receives a request, it executes **each route that matches the request** in the order in which they are declared until one sends a response. The following methods return a response:

- [appShell](#)
- [compute](#)
- [jwtVerification](#)
- [proxy](#)
- [redirect](#)
- [send](#)
- [serveStatic](#)
- [serviceWorker](#)
- [stream](#)
- [use](#)

Multiple routes can therefore be executed for a given request. A common pattern is to add caching with one route and render the response with a later one using middleware. In the following example we cache then render a response with Next.js:

```
JavaScript
```

```
Copy
```

```
1const { Router } = require('@layer0/core/router')
2const { nextRoutes } = require('@layer0/next')
3
4// In this example a request to /products/1 will be cached by the first
5// route, then served by the `nextRoutes` middleware
6new Router()
7  .get('/products/:id', ({ cache }) => {
8    cache({
9      edge: { maxAgeSeconds: 60 * 60, staleWhileRevalidateSeconds: 60 *
10     }
11  })
```

```

9     })
10    })
11    .use(nextRoutes)

```

Alter Requests and Responses

Edgio offers APIs to manipulate request and response headers and cookies. The APIs are:

Operation	Request	Upstream Response	Response sent to B
Set header	<code>setRequestHeader</code>	<code>setUpstreamResponseHeader</code>	<code>setResponseHeader</code>
Add cookie	*	<code>addUpstreamResponseCookie</code>	<code>addResponseCookie</code>
Update header	<code>updateRequestHeader</code>	<code>updateUpstreamResponseHeader</code>	<code>updateResponseHeader</code>
Update cookie	*	<code>updateUpstreamResponseCookie</code>	<code>updateResponseCookie</code>
Remove header	<code>removeRequestHeader</code>	<code>removeUpstreamResponseHeader</code>	<code>removeResponseHeader</code>
Remove cookie	*	<code>removeUpstreamResponseCookie</code>	<code>removeResponseCookie</code>

* Adding, updating, or removing a request cookie can be achieved with `updateRequestHeader` applied to cookie header.

You can find detailed descriptions of these APIs in the @layer0/core [documentation](#).

Embedded Values

You can inject values from the request or response into headers or cookies as template literals using the `${value}` format. For example: `setResponseHeader('original-request-path', '${path}')` would add an `original-request-path` response header whose value is the request path.

Value	Embedded value	Description
HTTP method	<code>\${method}</code>	The value of the HTTP method used for the request (e.g. GET)
URL	<code>\${url}</code>	The complete URL path including any query strings (e.g. <code>/search?query=docs</code>). Protocol, hostname, and port are not included.
Path	<code>\${path}</code>	The URL path excluding any query strings (e.g. <code>/search</code>)
Query string	<code>\${query:<name>}</code>	The value of the <code><name></code> query string or empty if not available.
Request header	<code>\${req:<name>}</code>	The value of the <code><name></code> request header or empty if not available.
Request cookie	<code>\${req:cookie:<name>}</code>	The value of the <code><name></code> cookie in <code>cookie</code> request header or empty if not available.
Request named parameter	<code>\${req:param:<name>}</code>	The value of the <code><name></code> param defined in the route or empty if not available.
Response header	<code>\${res:<name>}</code>	The value of the <code><name></code> response header or empty if not available.

Route Pattern Syntax

The syntax for route paths is provided by [path-to-regexp](#), which is the same library used by [Express](#).

Named Parameters

Named parameters are defined by prefixing a colon to the parameter name (`:foo`).

JavaScript

Copy

```
1 new Router().get('/:foo/:bar', res => {
2   /* ... */
3 })
```

Please note: Parameter names must use “word characters” ([A-Za-z0-9_]).

Custom Matching Parameters

Parameters can have a custom regexp, which overrides the default match ([^/]+). For example, you can match digits or names in a path:

JavaScript

Copy

```
1new Router().get('/icon-:foo(\\d+).png', res => {
2  /* ... */
3})
```

Tip: Backslashes need to be escaped with another backslash in JavaScript strings.

Custom Prefix and Suffix

Parameters can be wrapped in { } to create custom prefixes or suffixes for your segment:

JavaScript

Copy

```
1new Router().get('/:attr1?{-:attr2}?{-:attr3}?') , res => {
2  /* ... */
3})
```

Unnamed Parameters

It is possible to write an unnamed parameter that only consists of a regexp. It works the same the named parameter, except it will be numerically indexed:

JavaScript

Copy

```
1new Router().get('/:foo/(.*)', res => {
2  /* ... */
3})
```

Modifiers

Modifiers must be placed after the parameter (e.g. /:foo?, /(test)?, /:foo(test)?, or {-:foo(test)}?).

Optional

Parameters can be suffixed with a question mark (?) to make the parameter optional.

JavaScript

Copy

```
1new Router().get('/:foo/:bar?', res => {
2  /* ... */
3})
```

Tip: The prefix is also optional, escape the prefix \/ to make it required.

Zero or More

Parameters can be suffixed with an asterisk (*) to denote zero or more parameter matches.

JavaScript

Copy

```
1new Router().get('/:foo*', res => {
```

```
2 /* res.params.foo will be an array */
3})
```

The captured parameter value will be provided as an array.

[One or More](#)

Parameters can be suffixed with a plus sign (+) to denote one or more parameter matches.

JavaScript

Copy

```
1new Router().get('/:foo+', res => {
2 /* res.params.foo will be an array */
3})
```

The captured parameter value will be provided as an array.

[Matching Method, Query Parameters, Cookies, and Headers](#)

Match can either take a URL path, or an object which allows you to match based on method, query parameters, cookies, or request headers:

JavaScript

Copy

```
1router.match(
2 {
3   path: '/some-path', // value is route-pattern syntax
4   method: /GET|POST/i, // value is a regular expression
5   cookies: { currency: /^(usd)$/i }, // keys are cookie names, values
are regular expressions
6   headers: { 'x-moov-device': /^desktop$/i }, // keys are header names,
values are regular expressions
7   query: { page: /^(1|2|3)$/ }, // keys are query parameter names,
values are regular expressions
8 },
9 () => {},
10)
```

[Body Matching for POST requests](#)

You can also match HTTP POST requests based on their request body content as in the following example:

JavaScript

Copy

```
1router.match(
2 {
3   body: { parse: 'json', criteria: { operationName: 'GetProducts' } },
// the body content will be parsed as JSON and the parsed JSON matched
against the presence of the criteria properties (in this case a GraphQL
operation named 'GetProducts')
4 },
5 () => {},
6)
```

Currently the only body content supported is JSON. Body content is parsed as JSON and is matched against the presence of the fields specified in the `criteria` field. The [POST Body Matching Criteria](#) section below contains examples of using the `criteria` field.

Body matching can be combined with other match parameters such as headers and cookies. For example,

JavaScript

Copy

```
1router.match(  
2  {  
3    // Only matches GetProducts operations to the /graphql endpoint  
4    // for logged in users  
5    path: '/graphql',  
6    cookies: { loginStatus: /^(loggedIn)$/i }, // loggedin users  
7    body: { parse: 'json', criteria: { operationName: 'GetProducts' } },  
8  },  
9  () => {},  
10)
```

Caching & POST Body Matching

When body matching is combined with `cache` in a route, **the HTTP request body will automatically be used as the cache key**. For example, the code below will cache GraphQL `GetProducts` queries using the entire request body as the cache key:

JavaScript

Copy

```
1router.match(  
2  {  
3    body: { parse: 'json', criteria: { operationName: 'GetProducts' } },  
4  },  
5  ({ cache }) => {  
6    cache({  
7      edge: {  
8        maxAgeSeconds: 60 * 60,  
9        staleWhileRevalidateSeconds: 60 * 60 * 24, // this way stale  
items can still be prefetched  
10      },  
11    })  
12  },  
13)
```

You can still add additional parameters to the cache key using the normal EdgeJS `key` property. For example, the code below will cache GraphQL `GetProducts` queries separately for each user based on their `userID` cookie *and* the HTTP body of the request.

JavaScript

Copy

```

1router.match(
2  {
3    body: { parse: 'json', criteria: { operationName: 'GetProducts' } },
4  },
5  ({ cache }) => {
6    cache({
7      edge: {
8        maxAgeSeconds: 60 * 60,
9        staleWhileRevalidateSeconds: 60 * 60 * 24, // this way stale
items can still be prefetched
10     },
11     key: new CustomCacheKey().addCookie('userID'), // Split cache by
userID
12   })
13 },
14)

```

POST Body Matching Criteria

The `criteria` property can be a string or regular expression.

For example, the router below,

JavaScript

Copy

```

1router.match(
2  {
3    body: { parse: 'json', criteria: { foo: 'bar' } },
4  },
5  () => {},
6)

```

would match an HTTP POST request body containing:

JavaScript

Copy

```

1{
2  "foo": "bar",
3  "bar": "foo"
4}

```

Regular Expression Criteria

Regular expressions can also be used as `criteria`. For example,

JavaScript

Copy

```

1router.match(
2  {
3    body: { parse: 'json', criteria: { operationName: /^Get/ } },
4  },
5  () => {},
6)

```


would match an HTTP POST body containing:

JavaScript

Copy

```
1{
2  "operationName": "GetShops",
3  "query": "...",
4  "variables": {}
5}
```

Nested JSON Criteria

You can also use a nested object to match a field at a specific location in the JSON. For example,

JavaScript

Copy

```
1router.match(
2  {
3    body: {
4      parse: 'json',
5      criteria: {
6        operation: {
7          name: 'GetShops',
8        },
9      },
10   },
11 },
12 () => {},
13)
```

would match an HTTP POST body containing:

JavaScript

Copy

```
1{
2  "operation": {
3    "name": "GetShops",
4    "query": "...",
5  }
6}
```

GraphQL Queries

The EdgeJS router provides a `graphqlOperation` method for matching GraphQL.

JavaScript

Copy

```
1router.graphqlOperation('GetProducts', res => {
2  /* Handle the POST for the GetProducts query specifically */
3})
```

By default, the `graphqlOperation` assumes your GraphQL endpoint is at `/graphql`. You can alter this behavior by using the `path` property as shown below:

JavaScript

Copy

```
1router.graphqlOperation({ path: '/api/graphql', name: 'GetProducts' },  
res => {  
2  /* Handle the POST for the GetProducts query specifically */  
3})
```

Note that when the `graphqlOperation` function is used, the HTTP request body will automatically be included in the cache key.

The `graphqlOperation` function is provided to simplify matching of common GraphQL scenarios. For complex GraphQL matching (such as authenticated data), you can use the generic *Body Matching for POST requests* feature.

See the guide on [Implementing GraphQL Routing](#) in your project.

Request Handling

The second argument to routes is a function that receives a `ResponseWriter` and uses it to send a response. Using `ResponseWriter` you can:

- Proxy a backend configured in `layer0.config.js`
- Serve a static file
- Send a redirect
- Send a synthetic response
- Cache the response at edge and in the browser
- Manipulate request and response headers

[See the API Docs for Response Writer](#)

Blocking Search Engine Crawlers

If you need to block all search engine bot traffic to specific environments (such as your default or staging environment), the easiest way is to include the `x-robots-tag` header with the same directives you would otherwise set in a `meta` tag.

To block search engine traffic for Edgio edge links and permalinks, you can use the built-in `.noIndexPermalink()` call on the router:

JavaScript

Copy

```
1router.noIndexPermalink()
```

This will match requests with the `host` header matching `/layer0.link|layer0-perma.link/` and set a response header of `x-robots-tag: noindex`.

Additionally, you can customize this to block traffic to development or staging websites based on the `host` header of the request:

JavaScript

Copy

```
1router
2  .noIndexPermalink()
3  .get(
4    {
5      headers: {
6        // Regex to catch multiple hostnames
7        host: /dev.example.com|staging.example.com/,
8      },
9    },
10   ({ setResponseHeader }) => {
11     setResponseHeader('x-robots-tag', 'noindex')
12   },
13  )
```

Full Example

This example shows typical usage of @layer0/core, including serving a service worker, next.js routes (vanity and conventional routes), and falling back to a legacy backend.

JavaScript ./routes.js

Copy

```
1const { Router } = require('@layer0/core/router')
2
3module.exports = new Router()
4 // adds `x-robots-tag: noindex` response header to Edgio
5 // edge links and permalinks to prevent bot indexing
6 .noIndexPermalink()
7 .get('/service-worker.js', ({ serviceWorker }) => {
8   // serve the service worker built by webpack
9   serviceWorker('dist/service-worker.js')
10 })
11 .get('/p/:productId', ({ cache }) => {
12   // cache products for one hour at edge and using the service worker
13   cache({
14     edge: {
15       maxAgeSeconds: 60 * 60,
16       staleWhileRevalidateSeconds: 60 * 60,
17     },
18     browser: {
19       maxAgeSeconds: 0,
20       serviceWorkerSeconds: 60 * 60,
21     },
22   })
23   proxy('origin')
24 })
25 .fallback(({ proxy }) => {
```

```

26 // serve all unmatched URLs from the origin backend configured in
layer0.config.js
27 proxy('origin')
28 })

```

Errors Handling

You can use the router's `catch` method to return specific content when the request results in an error status (For example, a status code of 537). Using `catch`, you can also alter the `statusCode` and `response` on the edge before issuing a response to the user.

JavaScript

Copy

```
1router.catch(RegExp | string | number, (routeHandler: Function))
```

Examples

For example, to issue a custom error page when the origin returns any 5xx status code:

JavaScript

routes.js

```

Copy
1const { Router } = require('@layer0/core/router')
2
3module.exports = new Router()
4 // Example route that returns with a 5xx error status code
5 .get('/failing-route', ({ proxy }) => {
6   proxy('broken-origin')
7 })
8 // So let's assume that the route above returns 5xx, so instead of
rendering
9 // the broken-origin response we can alter that by specifying .catch
10 .catch(/5[0-9][0-9]/, ({ serveStatic }) => {
11   // The file below is present at the root of the directory
12   serveStatic('customized-error-page.html', { statusCode: 502 })
13 })

```

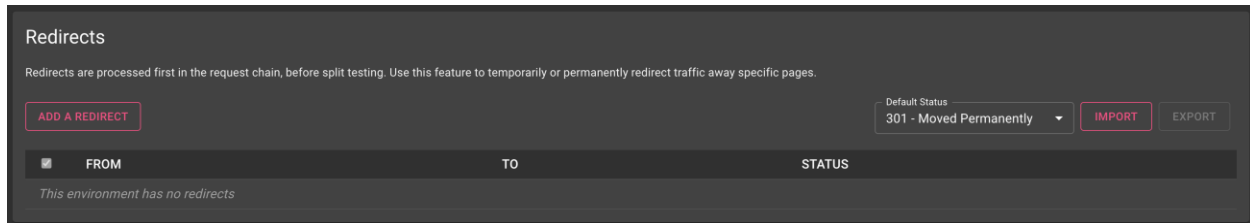
The `.catch` method allows the edge router to render a response based on the result preceding routes. So in the example above whenever we receive a 5xx, we respond with `customized-error-page.html` from the application's root directory, and change the status code to 502.

- Your catch callback is provided a [ResponseWriter](#) instance. You can use any ResponseWriter method except `proxy` inside `.catch`.
- We highly recommend keeping catch routes simple. Serve responses using `serveStatic` instead of `send` to minimize the size of the edge bundle.

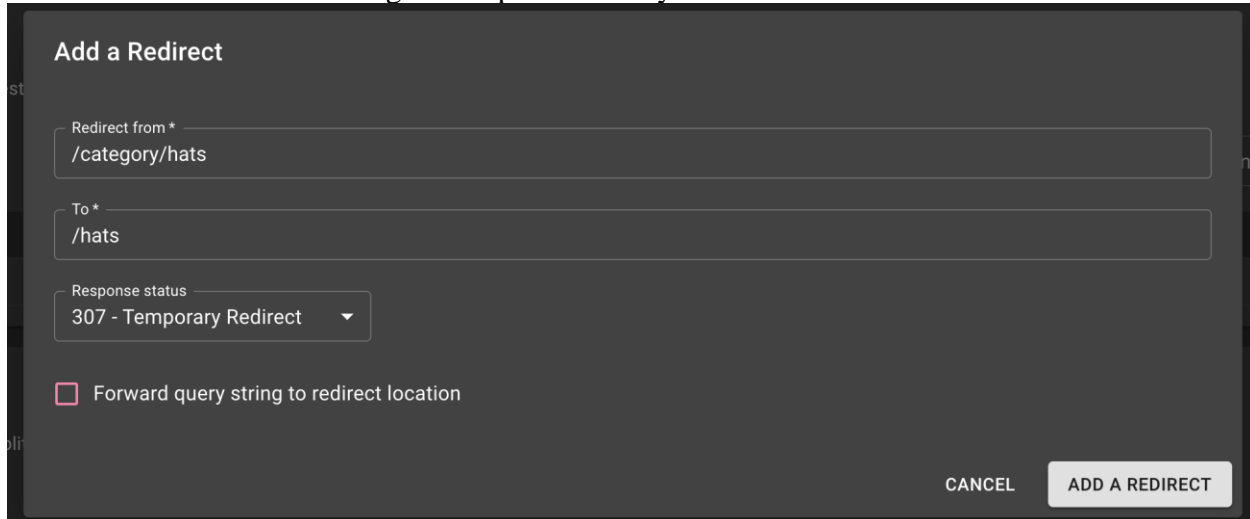
Environment Edge Redirects

In addition to sending redirects at the edge within the router configuration, this can also be configured at the environment level within the Edgio Developer Console.

Under *<Your Environment>* → *Configuration*, click *Edit* to draft a new configuration. Scroll down to the *Redirects* section:



Click *Add A Redirect* to configure the path or host you wish to redirect to:



Note: you will need to activate and redeploy your site for this change to take effect.

[Verifying JWTs on the Edge](#)

JWTs can be verified on the edge either in cookies or in a header (`Authorization: Bearer . . .`), with either symmetric or asymmetric encryption with signature lengths 256, 384 or 512 bits. Redirects or 403 Forbidden responses can be given individually on expired (`iat` claim, or premature `nfb` claim) or invalid (invalid signature, wrong algorithm, inadequate claims.)

JWTs; the 403 Forbidden can be intercepted and clients can be redirected to another page, which can optionally pass the original URL which intercepted the invalid/expired JWT.

Please see the [API documentation](#) for this feature for more detailed examples.

Common Routing Patterns

This guide gives examples of common routing patterns using Edgio.

[Proxying an Origin Same Path](#)

To forward a request to the same path on one of the backends listed in `layer0.config.js`, use the `proxy` method of `ResponseWriter`:

```
JavaScript
Copy
1router.get('/some-path', ({ proxy }) => {
2  proxy('origin')
3})
```

The first argument corresponds to the name of a backend in `layer0.config.js`. For example:

```
JavaScript
Copy
1module.exports = {
2  backends: {
3    origin: {
4      domainOrIp: 'my-shop.example.com',
5      hostHeader: 'my-shop.example.com',
6    },
7  },
8}
```

[Different Path](#)

To forward the request to a different path, use the `path` option of the `ProxyOptions` interface:

```
JavaScript
Copy
1router.get('/products/:productId', ({ proxy }) => {
2  proxy('origin', { path: '/p/:productId' })
3})
```

[Adding Caching](#)

To cache proxied requests at the edge, use the `cache` method.

```
JavaScript
Copy
1router.get('/products/:productId', ({ cache, proxy }) => {
2  cache({
3    edge: {
4      maxAgeSeconds: 60 * 60 * 24 // keep entries in the cache
for 24 hours
5      staleWhileRevalidateSeconds: 60 * 60 // when a cached page is
older than 24 hours, serve it one more time
6 // for up to 60 minutes while
fetching a new version from the origin
7    }
8  })
9})
```

```
8  })
9  proxy('origin')
10 }
```

Altering the Request

You can alter request headers when forwarding a request to a backend:

JavaScript

Copy

```
1 router.get(
2   '/products/:productId',
3   ({ setRequestHeader, updateRequestHeader, removeRequestHeader, proxy })
=> {
4     setRequestHeader('header-name', 'header-value')
5     updateRequestHeader('header-name', /some-.*-part/gi, 'some-
replacement')
6     removeRequestHeader('header-name')
7     proxy('origin')
8   },
9 )
```

The above example makes use of `setRequestHeader`, `updateRequestHeader`, and `removeRequestHeader` API calls.

Altering the Response

You can also alter the response before and after the cache:

JavaScript

Copy

```
1 router.get(
2   '/products/:productId',
3   ({
4     setUpstreamResponseHeader,
5     setResponseHeader,
6     removeResponseHeader,
7     removeUpstreamResponseHeader,
8     updateResponseHeader
9     updateUpstreamResponseHeader
10    proxy,
11  }) => {
12    proxy('origin')
13
14    // applied before the cache
15    setUpstreamResponseHeader('header-name', 'header-value')
16    updateUpstreamResponseHeader('header-name', /some-.*-part/gi, 'some-
replacement')
17    removeUpstreamResponseHeader('header-name')
18  }
```

```

19 // applied after the cache
20 setResponseHeader('header-name', 'header-value')
21 updateResponseHeader('header-name', /some-.*-part/gi, 'some-
replacement')
22 removeResponseHeader('header-name')
23 },
24)

```

[Altering All Responses](#)

You can also write catch-all routes that will alter all responses. One example where this is useful is injecting [Content Security Policy](#) headers.

Another example is adding response headers for debugging, which is often useful if [Edgio is behind another CDN](#) or if you are troubleshooting your router rules. For example, you could respond with the value of request `x-forwarded-for` into `x-debug-xff` to see the value that Edgio is receiving from the CDN:

```

JavaScript
Copy
1router.match(
2  {
3    path: '/*',
4    query: {
5      my_site_debug: 'true',
6    },
7  },
8  ({ setResponseHeader }) => {
9    setResponseHeader('x-debug-xff', `${req.x-forwarded-for}`)
10  },
11)
12// The rest of your router...

```

The rules for interpolating the values of request and response objects can be found in the [routing](#) guide. Note that catch-all routes that alter headers, cookies, or caching can be placed at the start of your router while allowing subsequent routes to run because they alter the request or the response without actually sending a response. See [route execution](#) for more information on route execution order and sending responses.

[Manipulating Cookies](#)

You can manipulate cookies before they are sent to the browser using cookie response API calls like `addResponseCookie`:

```

JavaScript
Copy
1router.get('/some/path', ({

```



```

2  addUpstreamResponseCookie,
3  addResponseCookie,
4  removeResponseCookie,
5  removeUpstreamResponseCookie,
6  updateResponseCookie
7  updateUpstreamResponseCookie,
8  proxy
9}) => {
10 proxy('origin')
11
12 // applied before the cache
13 addUpstreamResponseCookie('cookie-to-add', 'cookie-value')
14 removeUpstreamResponseCookie('cookie-to-remove')
15 updateUpstreamResponseCookie('cookie-to-alter', /Domain=.+;/,
'Domain=mydomain.com;')
16
17 // applied after the cache
18 addResponseCookie('cookie-to-add', 'cookie-value')
19 removeResponseCookie('cookie-to-remove')
20 updateResponseCookie('cookie-to-alter', /Domain=.+;/,
'Domain=mydomain.com;')
21})

```

[Adding Options to Cookies](#)

In addition to the name and value of the cookie, you can also add attributes to each cookie. For information on possible cookie attributes, please refer to <https://developer.mozilla.org/en-US/docs/Web/HTTP/Headers/Set-Cookie>

```

JavaScript
Copy
1router.get('/some/path', ({ addUpstreamResponseCookie, addResponseCookie,
proxy }) => {
2  proxy('origin')
3
4  addUpstreamResponseCookie('cookie-to-add', 'cookie-value', {
5    domain: 'test.com',
6  })
7
8  addResponseCookie('cookie-to-add', 'cookie-value', { 'max-age': 50000
9  })
9})

```

[Proxying to Different Backends Based on Different Host Names](#)

To proxy to different backends by matching the `host` header (e.g. different backends for different international sites):

```

JavaScript

```

Copy

```
1router
2  .match(
3    {
4      path: '/*:path*',
5      headers: {
6        host: 'yoursite.c1',
7      },
8    },
9    ({ proxy }) => {
10     proxy('country1-backend')
11   },
12 )
13 .match(
14   {
15     path: '/*:path*',
16     headers: {
17       host: 'yoursite.c2',
18     },
19   },
20   ({ proxy }) => {
21     proxy('country2-backend')
22   },
23 )
24 .match(
25   {
26     path: '/*:path*',
27   },
28   ({ proxy }) => {
29     proxy('everybody-else-backend')
30   },
31 )
```

Serving a Static File

To serve a specific file use the [serveStatic](#) API:

JavaScript

Copy

```
1router.get('/favicon.ico', ({ serveStatic, cache }) => {
2  cache({
3    edge: {
4      maxAgeSeconds: 60 * 60 * 24, // cache at the edge for 24 hours
5    },
6    browser: {
7      maxAgeSeconds: 60 * 60 * 24, // cache for 24 hours
8    },
9  })
```

```
10  serveStatic('assets/favicon.ico') // path is relative to the root of
    your project
11})
```

Serving Static Files From a Directory

Here's an example that serves all requests by sending the corresponding file in the `public` directory

```
JavaScript
Copy
1router.get('/:path*', ({ serveStatic, cache }) => {
2  cache({
3    edge: {
4      maxAgeSeconds: 60 * 60 * 24, // cache at the edge for 24 hours
5    },
6    browser: false, // prevent caching of stale html in the browser
7  })
8  serveStatic('public/:path*')
9})
```

Routing to Serverless

If your request needs to be run on the serverless tier, you can use the `renderWithApp` handler to render your result using your application. Use this method to respond with an SSR or API result from your application.

Example using the `renderWithApp` handler:

```
JavaScript
Copy
1router.get('/some/:path*', ({ renderWithApp, cache }) => {
2  cache(CACHE_PAGES)
3  renderWithApp()
4})
```

Falling Back to Server-side Rendering

If you render some but not all paths for a given route at build time, you can fall back to server side rendering using the `onNotFound` option. Add the `loadingPage` option to display a loading page while server-side rendering is in progress.

```
JavaScript
Copy
1router.get('/products/:id', ({ serveStatic, cache, renderWithApp }) => {
2  cache({
3    edge: {
4      maxAgeSeconds: 60 * 60 * 24, // cache at the edge for 24 hours
5    },
```

```

6  })
7  serveStatic('dist/products/:id.html', {
8    onNotFound: () => renderWithApp(),
9    loadingPage: 'dist/products/loading.html',
10 })
11})

```

This hybrid of static and dynamic rendering was first introduced in Next.js as [Incremental Static Generation \(ISG\)](#). In Next.js apps, developers enable this behavior by returning `fallback: true` from `getStaticPaths()`.

The `@layer0/next` package automatically configures the routes for ISG pages to use `onNotFound` and `loadingPage`.

[Returning a Custom 404 Page](#)

When a request matches a route with `serveStatic`, but no matching static asset exists, you can serve a custom 404 page using the `onNotFound` option.

```

JavaScript
Copy
1router.get('/products/:id', ({ serveStatic, cache }) => {
2  cache({
3    edge: {
4      maxAgeSeconds: 60 * 60 * 24, // cache at the edge for 24 hours
5    },
6  })
7  serveStatic('dist/products/:id.html', {
8    onNotFound: async () => {
9      await serveStatic('/products/not-found.html', {
10       statusCode: 404,
11       statusMessage: 'Not Found',
12     })
13   },
14 })
15})

```

[Responding with a String Response Body](#)

To respond with a simple, constant string as the response body use the `send` method:

```

JavaScript
Copy
1router.get('/some-path', ({ cache, setResponseHeader, send }) => {
2  cache({
3    edge: {
4      maxAgeSeconds: 60 * 60 * 24, // cache for 24 hours
5    },
6  })

```

```

7  setResponseHeader('Content-Type', 'text/html')
8  send(`
9    <!doctype html>
10   <html>
11     <body>Hello World</body>
12   </html>
13 `)
14})

```

To compute a dynamic response use the [compute](#) method:

```

JavaScript
Copy
1router.get('/hello/:name', ({ cache, setResponseHeader, compute, send })
=> {
2  cache({
3    edge: {
4      maxAgeSeconds: 60 * 60 * 24, // cache for 24 hours
5    },
6  })
7  setResponseHeader('Content-Type', 'text/html')
8  compute((request, response) => {
9    send(`
10     <!doctype html>
11     <html>
12       <body>Hello ${request.params.name}</body>
13     </html>
14     `)
15  })
16})

```

[Redirecting](#)

To redirect the browser to a different URL, use the [redirect](#) API:

```

JavaScript
Copy
1router.get('/p/:productId', ({ redirect }) => {
2  return redirect('/products/:productId', 301) // overrides the default
status of 302 (Temporary Redirect)
3})

```

If you need to compute the destination with sophisticated logic:

```

JavaScript
Copy
1router.get('/p/:productId', ({ redirect, compute, cache }) => {
2  cache({
3    edge: {
4      maxAgeSeconds: 60 * 60 * 24, // cache for 24 hours

```

```

5     },
6   })
7   compute(async request => {
8     const destination = await
getDestinationFromMyAPI(request.params.productId)
9     redirect(destination)
10  })
11})

```

[Redirecting All Traffic to a Different Domain](#)

JavaScript

Copy

```

1// Redirect all traffic except those with host header starting with www.
to www.mydomain.com
2router.match({ headers: { host: /^(?!www\.).*$/ } }, ({ redirect }) => {
3  redirect('https://www.mydomain.com${path}')
4})
5
6// Redirect all traffic from www.domain.com to domain.com
7router.match({ headers: { host: /^(www\.).*$/ } }, ({ redirect }) => {
8  redirect('https://domain.com${path}')
9})

```

[Blocking Unwanted Traffic](#)

[Blocking traffic from specific countries](#)

If you need to block all traffic from a specific country or set of countries, you can do so by matching requests by the [country code](#) geolocation header:

JavaScript

Copy

```

1router.get(
2  {
3    headers: {
4      'x-0-geo-country-code': /XX|XY|XZ/, // Regex matching two-letter
country codes of the countries you want to block
5    },
6  },
7  ({ send }) => {
8    send('Blocked', 403)
9  },
10)

```

You can find more about geolocation headers [here](#).

[Allowing Specific IPs](#)

If you need to block all traffic except requests that originate from specific IP addresses, you can do so by matching requests by the [x-0-client-ip](#) header:

```
JavaScript
Copy
1router.get(
2  {
3    headers: {
4      // Regex that will do a negative lookahead for IPs you want to
allow.
5      // In this example 172.16.16.0/24 and 10.10.10.3/32 will be allowed
and everything else will receive a 403
6      'x-0-client-ip':
/\b((?!172\.16\.16)(?!10.10.10.3)\d{1,3}\.\d{1,3}\.\d{1,3}\.\d{1,3})\b/,
7    },
8  },
9  ({ send }) => {
10   send('Blocked', 403)
11 },
12)
```

[Blocking Search Engine Crawlers](#)

If you need to block all search engine bot traffic to specific environments (such as your default or staging environment), the easiest way is to include the `x-robots-tag` header with the same directives you would otherwise set in a `meta` tag.

To block search engine traffic for Edgio edge links and permalinks, you can use the built-in `.noIndexPermalink()` call on the router:

```
JavaScript
Copy
1router.noIndexPermalink()
```

This will match requests with the `host` header matching `/layer0.link|layer0-perma.link/` and set a response header of `x-robots-tag: noindex`.

Additionally, you can customize this to block traffic to development or staging websites based on the `host` header of the request:

```
JavaScript
Copy
1router
2  .noIndexPermalink()
3  .get(
4    {
5      headers: {
6        // Regex to catch multiple hostnames
```

```
7     host: /dev.example.com|staging.example.com/,
8   },
9 },
10 ({ setResponseHeader }) => {
11   setResponseHeader('x-robots-tag', 'noindex')
12 },
13 )
```

Custom Domains & SSL

This guide covers the steps you need to take your site live on Edgio with a secure, custom domain.

Creating custom domains is always done in the context of creating or updating an environment.

Configuration Overview


1. If needed, create an environment using instructions in [Environments](#).
2. Create the [Custom Domain](#).
3. Do [Network Configuration](#) (DNS and the IP allow list) for the domain.
4. Configure [TLS/SSL](#) for the domain.

Custom Domains

Before going live, you must create a production environment and configure your domains.

To configure your custom domains:

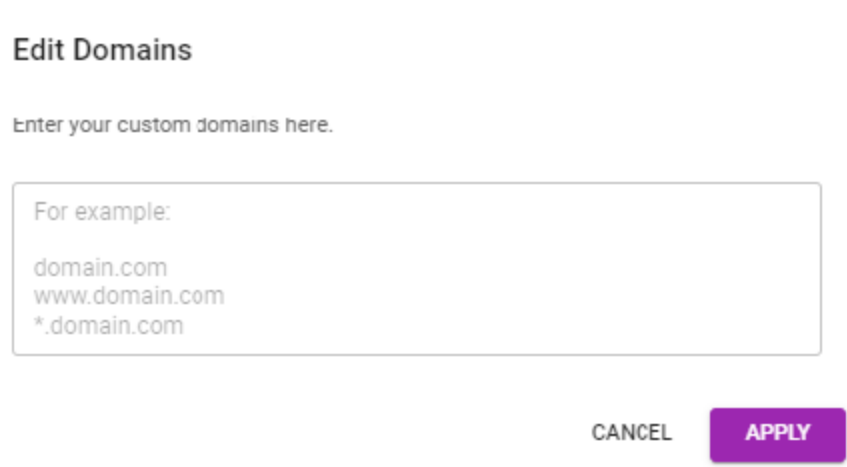
1. Navigate to a site, then open an existing environment or create a new environment. (To create an environment, use instructions in [Environments](#).)
 - For an existing environment, select the *ENVIRONMENTS* tab header, then click an environment name in the list of environments. Continue with the numbered steps below.
 - For a new environment, the *DEPLOYMENTS* tab is displayed. Continue with the following steps.
2. Select the *CONFIGURATION* tab header.




3. Create a new draft version of the environment by clicking *EDIT* at the top of the page.
4. In the *Domains* section, click *EDIT DOMAINS*.



5. Enter a name in the *Edit Domains* dialog, then click *APPLY*.



6. Click *ACTIVATE* at the top of the page to enable the updated environment.



[Migrating from Fastly](#)

If you're migrating to Edgio from [Fastly](#), you will need to do the following before adding your domains to your Edgio environment:

- Contact [Fastly support](#) and request that control of your domains be transferred to Edgio. Be sure to explicitly list each domain that needs to be transferred and ask Fastly to contact support(at)layer0.co if they need Edgio to confirm the transfer.
- Before going live with Edgio, you will need to ensure that you've removed your domains from all active Fastly services. To remove domains from a service, clone the service, remove the domains, then activate the new version of the service. Once the new service version is activated you can add the domains to your Edgio environment and activate it.

Network Configuration

You can find the DNS and allowed IP configurations in the *Networking* tab for your environment.

The screenshot shows the 'Networking' tab in the Edgio dashboard. It contains two main sections: 'DNS Configuration' and 'IP Whitelist'. The 'DNS Configuration' section explains that DNS records are needed to route traffic to Layer0. It provides instructions for using a sub-domain (adding a CNAME record) and an apex domain (creating multiple A records). A code block shows a CNAME record value: `1032280f-3ce1-4521-bcb0-801135d91948.layer0.link`. Another code block shows four A record values: `151.101.1.79`, `151.101.65.79`, `151.101.129.79`, and `151.101.193.79`. The 'IP Whitelist' section explains that origin servers should not block requests from Layer0. A code block lists 25 IP address ranges, such as `23.235.32.0/20`, `43.249.72.0/22`, `103.244.50.0/24`, `103.245.222.0/23`, `103.245.224.0/24`, `104.156.80.0/20`, `146.75.0.0/16`, `151.101.0.0/16`, `157.52.64.0/18`, `167.82.0.0/17`, `167.82.128.0/20`, `167.82.160.0/20`, `167.82.224.0/20`, `172.111.64.0/18`, `185.31.16.0/22`, `199.27.72.0/21`, `199.232.0.0/16`, `35.171.218.41`, `107.20.40.25`, `3.18.205.112`, and `3.23.32.142`.

DNS

In order to configure your DNS provider to direct traffic for a particular set of domains to Edgio, you must create DNS records for your website. If you are launching a new site, then you can create the records whenever you feel ready. For sites that are already live, the DNS update is the last step. Once you have updated your DNS you are committed to launching.

Using a Sub-domain (e.g. [www.mywebsite.xyz](#))

To host your site on a subdomain, add a CNAME record with the value shown under *DNS Configuration* (see above).

```
1# To verify your DNS entry, run the following command
2dig <your-sub-domain>
3
4# Example
5dig www.mywebsite.xyz
```

```
6
7# Result
8www.mywebsite.xyz. 599 IN CNAME d12ea738-71b3-25e8-c771-
6fdd3f6bd8ba.layer0-limelight.link.
```

[Using an Apex Domain \(e.g. mywebsite.xyz\)](#)

To host your site on the apex domain, create multiple A records on your apex domain, with the following Anycast IP address values: 208.69.180.11, 208.69.180.12, 208.69.180.13, 208.69.180.14

```
1# To verify your DNS entry, run the following command
2dig <your-apex-domain>
3
4# Example
5dig mywebsite.xyz
6
7# Result
8mywebsite.xyz. 599 IN A 208.69.180.11
9mywebsite.xyz. 599 IN A 208.69.180.12
10mywebsite.xyz. 599 IN A 208.69.180.13
11mywebsite.xyz. 599 IN A 208.69.180.14
```

[Using Both an Apex Domain and a Sub-domain \(e.g. mywebsite.xyz and www.mywebsite.xyz\)](#)

- Create the multiple A records with the IPs, on your apex domain (see above).
- Create a CNAME record for your sub-domain, with the value of your apex domain.

```
1# To verify your DNS entries, run the following command
2dig <your-sub-domain>
3# Example
4dig www.mywebsite.xyz
5# Result
6www.mywebsite.xyz. 599 IN CNAME. mywebsite.xyz.
7mywebsite.xyz. 599 IN A 208.69.180.11
8mywebsite.xyz. 599 IN A 208.69.180.12
9mywebsite.xyz. 599 IN A 208.69.180.13
10mywebsite.xyz. 599 IN A 208.69.180.14
```

[Allowing Edgio IP Addresses](#)

Before going live, ensure that all Edgio IP addresses are allowed in the security layer in front of your origin and/or API servers. The IP addresses you need to allow can be found on the *Allowlisting* section of the *Networking* tab of the *Environment* page. Note that each team may have their own set of IPs so these values cannot be copied from one team to another.

[TLS/SSL](#)

All data transmitted to and from your Edgio site must be secured with TLS (Transport Layer Security). TLS, also known as SSL (Secure Sockets Layer), is a cryptographic protocol to communicate securely over the Internet. TLS provides end-to-end data encryption and data integrity for all web requests.

Edgio provides a wildcard TLS certificate that covers the auto-generated domains that it assigns to your site (e.g {team}-{site}-{branch}-{version}.layer0-limelight.link). You need to provide your own certificate for your site's custom domains.



If you already have an existing certificate, you can use it by skipping ahead to [Uploading your Certificate](#). Many customers who have existing certificates still choose to obtain a new one when adopting Edgio so as not to reuse the same private key with more than one vendor/system.

Obtaining a Certificate Automatically

Edgio can generate SSL Certificates on your behalf using [Let's Encrypt](#). Certificates are free, valid for 3 months, and automatically renewed as long as the technical requirements, shown below, remain met:

1. Make sure each environment is configured with the custom domains on which it will receive traffic. For more information on configuring custom domains, see [Custom Domains](#) above.
2. Using your DNS provider, verify and possibly add a CAA record to allow *Let's Encrypt* to generate certificates for your domains.
 - The CAA DNS entries of a domain behave like an allow list to indicate whether **any** or only **certain** Certificate Authorities are allowed to generate certificates for that domain.
 - If there are no CAA records, it means that **any** Certificate Authority is allowed to generate certificates for that domain.
 - If there are CAA records, it means that only **certain** Certificate Authorities are allowed to generate certificates for that domain.
 - So in order for *Let's Encrypt* to be able to generate a certificate for your domains, you must either not have defined any CAA records, or *Let's Encrypt's* CAA entry must be among those defined in the list of CAA records.

You can verify the value of the CAA records for your domain from the command line using the command below.

```
Bash
Copy
1# Run the following command
2dig caa +short <your-apex-domain>
3
4# Example
5dig caa +short mywebsite.xyz
```

Example of a CAA query showing that only **certain** Certificate Authorities are allowed to generate certificates for that domain:

```
Bash
Copy
10 issue "amazon.com"
20 issue "digicert.com"
30 issue "globalsign.com"
40 issue "letsencrypt.org"
```

If the result of the CAA DNS query is empty, it means that **any** Certificate Authority is allowed to generate certificates on that domain. If so, you can directly go to the next step.

If there are already some CAA DNS entries defined on your domain, and if *Let's Encrypt's* CAA entry is not among those, you will have to add an additional CCA entry for *Let's Encrypt*.

To do so, log into your DNS provider, and add a CAA type DNS record with the following values:

- Type : CAA
- Name : empty (or @, depending on the DNS provider)
- Flags: 0
- Tag: issue
- Value: letsencrypt.org (or "letsencrypt.org")

Example with GoDaddy:

Type *	Name *	Flags *
CAA	@	0
Tag *	Value *	TTL *
issue	letsencrypt.org	1 Hour
		<input type="button" value="Save"/> <input type="button" value="Cancel"/>

Example with Gandi:

TTL *

Unit *

The minimum TTL value for Gandi's LiveDNS is 300 seconds.

Name

Flags *

Tag *

Hostname *


Value *


You can use the following links to see how to configure the CAA record on commonly used DNS providers:


- [How to add a CAA record on Gandi](#)
- [How to add a CAA record on Godaddy](#)
- [How to add a CAA record on AWS](#)
- [How to add a CAA record on NameCheap](#)

Once the DNS entry has been added, you can verify the CAA record using one of the following:

- [CAA Test](#)
- [Entrust CAA Lookup](#)

 Many DNS providers have already added this CAA DNS record by default

 Some DNS providers does not allow the creation of CAA DNS records and therefore allow any Certificate Authority to generate certificates

 You can learn more about CAA DNS records on [Let's Encrypt website](#), on [Wikipedia](#), on [Gandi](#) and on [eff.org](#)

3. Add an `_acme-challenge`. CNAME DNS entry to allow Edgio to issue a certificate request on your behalf.
Log into your DNS provider and add one CNAME type DNS entry with the value `_acme-challenge.<your-domain-here>` for each domain you use on your Edgio website. For example, if your domain is `mywebsite.xyz`, the DNS entry should have a value of `_acme-challenge.mywebsite.xyz`. This record should point to `_acme-challenge.xdn-validation.com`. Repeat the operation of each domain associated with your Edgio website.

Example with Godaddy:



The screenshot shows a DNS configuration form with the following fields:

- Type ***: CNAME
- Host ***: `_acme-challenge.www`
- Points to ***: `_acme-challenge.xdn-validation.com`
- TTL ***: 1 Hour

Buttons: Save, Cancel

Example with Gandi:

Type *

CNAME

TTL * 1800 Unit * seconds

The minimum TTL value for Gandi's LiveDNS is 300 seconds.

Name

_acme-challenge.www .mywebsite.xyz

Hostname *

_acme-challenge.xdn-validation.com.

Once the DNS entries have been added, you can use one of the following to verify that they are correctly configured:

- [MX ToolBox DNS Lookup](#)
- [Nslookup DNS Lookup](#)

You can also verify the CNAME records using the command line:

```
Bash
Copy
```

```
1# Run the following 'dig' command to verify the presence of the
'_acme-challenge.' CNAME :
```

```
2dig +short cname _acme-challenge.<your-domain>
```

```
3
```

```
4# For example:
```

```
5dig +short cname _acme-challenge.mywebsite.xyz
```

Expected result for the DNS query:

```
1_acme-challenge.xdn-validation.com.
```

If you use multiple domains for your website,

like `mywebsite.xyz` and `www.mywebsite.xyz`, you will have to make sure that the `_acme-challenge` DNS record has been added for both domains:

1 `_acme-challenge.mywebsite.xyz -> _acme-challenge.xdn-validation.com.`

2 `_acme-challenge.www.mywebsite.xyz -> _acme-challenge.xdn-validation.com.`

If you have been previously using *Let's Encrypt* to generate certificates for this domain, please verify that there are no remaining TXT records named `_acme-challenge.mywebsite.xyz`.




You can read more about the `acme-challenge.` process by visiting [Let's Encrypt Website](#)


4. Once the requirements above are met, you can generate the certificate using the [Edgio Developer Console](#):
 1. Select your site and navigate to *Settings > SSL Certificate*
 2. Verify the state of your certificate (you should see that there's no certificate provided yet for your website):

SSL Certificate

Use XDN to generate and upload SSL certificate for provided domains

Certificate not yet provided 

The certificate generation and uploading may take up to 15 minutes.

NAME	DNS	SSL	CREATED AT	EXPIRATION DATE
www.mywebsite.xyz		<input type="radio"/>	4 minutes ago	-

[GENERATE SSL](#)

- Click on the *Generate SSL Certificate* button:

SSL Certificate

Use XDN to generate and upload SSL certificate for provided domains

Certificate not yet provided

The certificate generation and uploading may take up to 15 minutes.

NAME	DNS	SSL	CREATED AT	EXPIRATION DATE
www.mywebsite.xyz	✓	●	4 minutes ago	-

[GENERATE SSL](#) ←

4. After a couple of minutes, you should see that your website has received a valid certificate:

SSL Certificate

Use XDN to generate and upload SSL certificate for provided domains

Certificate is active →

The certificate generation and uploading may take up to 15 minutes.

NAME	DNS	SSL	CREATED AT	EXPIRATION DATE
www.mywebsite.xyz	✓	●	8 minutes ago	2021-05-19T12:51:11Z

[GENERATE SSL](#) ↑

[Creating a Certificate Manually](#)

TLS certificates are issued by Certificate Authorities (CA) based on Certificate Signing Request (CSR) that they receive from you. Alongside the CSR the same process creates the certificate's private key. You only need to share your CSR with CA, not the private key which you should store securely.

The following steps describe the creation of the CSR and private key with OpenSSL. OpenSSL is an open-source toolkit for the TLS protocol. We recommend using OpenSSL because it ensures that your private key will only be stored locally on your infrastructure.

Your CA may have more customized guides or an entirely customized certification process.

To create CSR and private key do the following:

1. Open your terminal window and make sure that you have OpenSSL installed:
 - On MacOS you can install it by using [brew](#) package manager (e.g. `brew install openssl`)
 - On Windows you can install it by using [Chocolatey](#) package manager (e.g. `choco install openssl`)
 - On Linux/Unix you can install it by running the built-in OS package manager (e.g. `apt-get install openssl`, `apk add openssl` and so on)
2. Go to the directory of your choice and create a configuration

file `layer0.conf` based on this template:

```
1[req]
2default_bits=2048
3distinguished_name = req_distinguished_name
4req_extensions = v3_req
5
6[req_distinguished_name]
7countryName=Country Name (2 letter code)
8countryName_default=US
9stateOrProvinceName=State or Province Name (full name)
10stateOrProvinceName_default=California
11localityName=Locality Name (e.g. city)
12localityName_default=San Francisco
13organizationName=Organization Name (e.g. company)
14organizationName_default=YourCompanyName
15commonName=Fully Qualified Domain Name (FQDN) e.g. www.your-
company-name.com
16commonName_default=www.your-company-domain.com
17
18[ v3_req ]
19subjectAltName=@alt_names
20
21[alt_names] # Other domains: apex domain, wildcard domain for
staging and dev, and so on
22DNS.1=*.your-main-domain.com
23DNS.2=*.your-dev-domain.com
24DNS.3=your-apex-domain.com
25# And so on
```

Replace the country, state/province, locality, organization name and, most importantly Common Name (CN), for the cert which must be the fully qualified domain name for your domain (e.g. for Edgio that is `www.layer0.co`)

You will want to add all the additional domains into the `alt_names` section. There you should add your development, staging and other domains although Edgio strongly encourages the use of wildcard certs.

3. Run `openssl req -out layer0.csr -newkey rsa:2048 -nodes -keyout layer0.key -config layer0.conf -batch`. This should generate your CSR in `layer0.csr` and private key in `layer0.key`. If you want OpenSSL to ask you for each different input, remove the `-batch` option and re-run the command.
4. Verify your CSR contains the expected domains by running `openssl req -in layer0.csr -noout -text | grep DNS`
5. Read the CSR (e.g. `cat layer0.csr`) or copy to your clipboard (on OSX `cat layer0.csr | pbcopy`) and send it to your CA for certification.

Uploading Your Certificate

Prerequisites

To upload a certificate, you must have the **Admin** role on your team, and your team must be upgraded to Edgio Enterprise.

Edgio needs the following to correctly host your certificate:

- Certificate issued by CA
- Intermediate certificates (IC) used by CA, including CA's signing certificate
- Private key that was generated at the time of the CSR.

Uploading the certificate

To upload your SSL certificate, do the following:

1. Navigate to the *Settings* tab on your site:



2. Scroll to *TLS Certificate*.

TLS Certificate

Not yet generated

Automatically create an TLS certificate for my custom domains.

In order to configure TLS for your domains, you need to generate a single certificate for all custom domains. Provide the certificate, intermediate certificates, and private key using the form below.

Primary Certificate

Paste certificate here, e.g:

```
-----BEGIN CERTIFICATE-----  
ABCDEFghijklMnOPq.....  
-----END CERTIFICATE-----
```

Intermediate Certificate

Paste certificate(s), e.g:

```
-----BEGIN CERTIFICATE-----  
ABCDEFghijklMnOPq.....  
-----END CERTIFICATE-----
```

Private Key

Paste private key here, e.g:

```
-----BEGIN PRIVATE KEY-----  
ABCDEFghijklMnOPq.....  
-----END PRIVATE KEY-----
```

Your private key will not be shown again once saved.

CHANGES SAVED

3. Toggle *Automatically create an TLS certificate for my custom domains* to the *on* position.
4. Copy the certificate, intermediate certificates, and the private key into the corresponding edit boxes.



The private key is non-public data and must not be shared with parties other than Edgio. Edgio stores your private key securely at rest. It is never shown in the developer console and only used to provision parts of the infrastructure that are used to terminate TLS connections.

5. Click *CHANGES SAVED*.

The certificate's status becomes *Activating*:

TLS Certificate

Activating Your certificate has been generated and is propagating through the Layer0 edge network

After the certificate is activated, its status becomes *Active*:

TLS Certificate

Active (expires , automatically renewed) [DETAILS](#) ▼

Note: Certificate activation should take just a few minutes. If the status does not become Active within an hour, please contact [support](#).