
Hoverfly Documentation

Release v0.10.0

SpectoLabs

Feb 15, 2017

1	Introduction	3
1.1	Motivation	3
1.2	Download and installation	3
1.3	Getting Started	4
2	Key Concepts	5
2.1	Hoverfly as a proxy server	5
2.2	Hoverfly as a webservice	7
2.3	Hoverfly modes	8
2.4	Simulations	11
2.5	Captured traffic	13
2.6	Templates	14
2.7	Destination filtering	16
2.8	Meta	17
2.9	Delays	17
2.10	Middleware	17
3	Native language bindings	19
3.1	HoverPy	19
3.2	Hoverfly Java	19
4	Tutorials	21
4.1	Creating and exporting a simulation	21
4.2	Importing and using a simulation	22
4.3	Adding delays to a simulation	23
4.4	Adding templates to a simulation	26
4.5	Using middleware to simulate network latency	27
4.6	Using middleware to modify response payload and status code	28
4.7	Simulating HTTPS APIs	29
4.8	Running Hoverfly as a webservice	30
4.9	Capturing or simulating specific URLs	30
4.10	Using Hoverfly behind a proxy	31
5	Reference	33
5.1	hoverctl commands	33
5.2	Hoverfly commands	34
5.3	REST API	36

5.4 Simulation schema 42



Hoverfly is a lightweight, open source API simulation tool. Using Hoverfly, you can create realistic simulations of the APIs your application depends on.

- Replace slow, flaky API dependencies with realistic, re-usable simulations
- Simulate network latency, random failures or rate limits to test edge-cases
- Extend and customize with any programming language
- Export, share, edit and import API simulations
- CLI and native language bindings for Java and Python
- REST API
- Lightweight, high-performance, run anywhere
- Apache 2 license

Hoverfly is developed and maintained by [SpectoLabs](#).

1.1 Motivation

Developing and testing interdependent applications is difficult. Maybe you're working on a mobile application that needs to talk to a legacy API. Or a microservice that relies on two other services that are still in development.

The problem is the same: how do you develop and test against external dependencies which you cannot control?

You could use mocking libraries as substitutes for external dependencies. But mocks are intrusive, and do not allow you to test all the way to the architectural boundary of your application.

Stubbed services are better, but they often involve too much configuration or may not be transparent to your application.

Then there is the problem of managing test data. Often, to write proper tests, you need fine-grained control over the data in your mocks or stubs. Managing test data across large projects with multiple teams introduces bottlenecks that impact delivery times.

Integration testing “over the wire” is problematic too. When stubs or mocks are swapped out for real services (in a continuous integration environment for example) new variables are introduced. Network latency and random outages can cause integration tests to fail unexpectedly.

Hoverfly was designed to provide you with the means to create your own “dependency sandbox”: a simulated development and test environment that you control.

Hoverfly grew out of an effort to build “the smallest service virtualization tool possible”.

1.2 Download and installation

Hoverfly comes with a command line interface called **hoverctl**. Archives containing the Hoverfly and **hoverctl** binaries are available for the major operating systems and architectures.

- MacOS 64bit
- Linux 32bit
- Linux 64bit

- Windows 32bit
- Windows 64bit

Download the correct archive, extract the binaries and place them in a directory on your PATH.

1.2.1 Homebrew (MacOS)

If you have [homebrew](#), you can install Hoverfly using the `brew` command.

```
brew install SpectoLabs/tap/hoverfly
```

1.3 Getting Started

Note

It is recommended that you keep Hoverfly and `hoverctl` in the same directory. However if they are not in the same directory, `hoverctl` will look in the current directory for Hoverfly, then in other directories on the PATH.

Hoverfly is composed of two binaries: **Hoverfly** and **hoverctl**.

`hoverctl` is a command line tool that can be used to configure and control Hoverfly. It allows you to run Hoverfly as a daemon.

Hoverfly is the application that does the bulk of the work. It provides the proxy server or webserver, and the API endpoints.

Once you have extracted both Hoverfly and `hoverctl` into a directory on your PATH, you can run `hoverctl` and Hoverfly.

```
hoverctl --version
hoverfly --version
```

Both of these commands should return a version number. Now you can run an instance of Hoverfly:

```
hoverctl start
```

Check whether Hoverfly is running with the following command:

```
hoverctl logs
```

The logs should contain the string `serving proxy`. This indicates that Hoverfly is running.

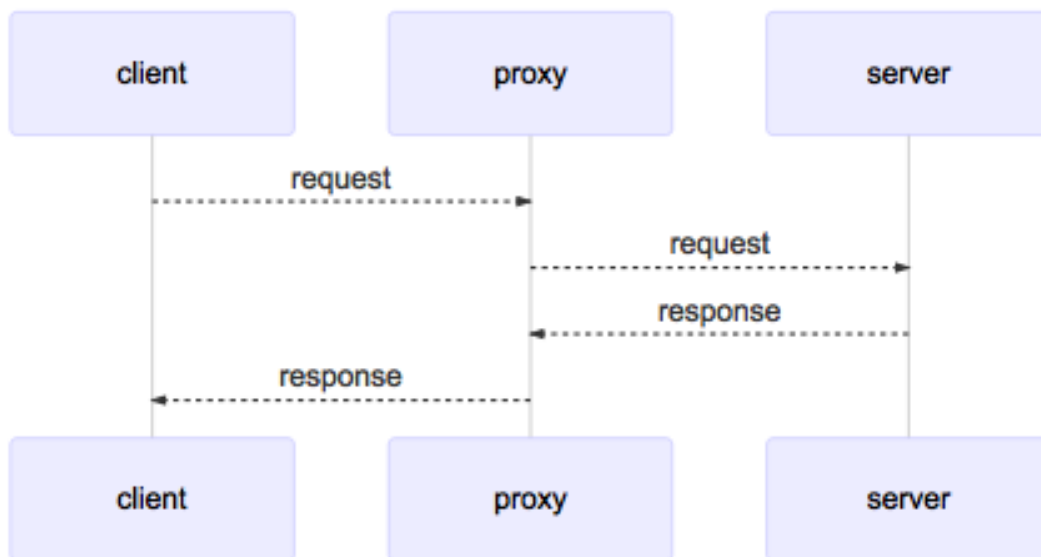
Finally, stop Hoverfly with:

```
hoverctl stop
```


Hoverfly's functionality is quite broad. You are encouraged to take the time to understand these key concepts before jumping into the *Tutorials*.

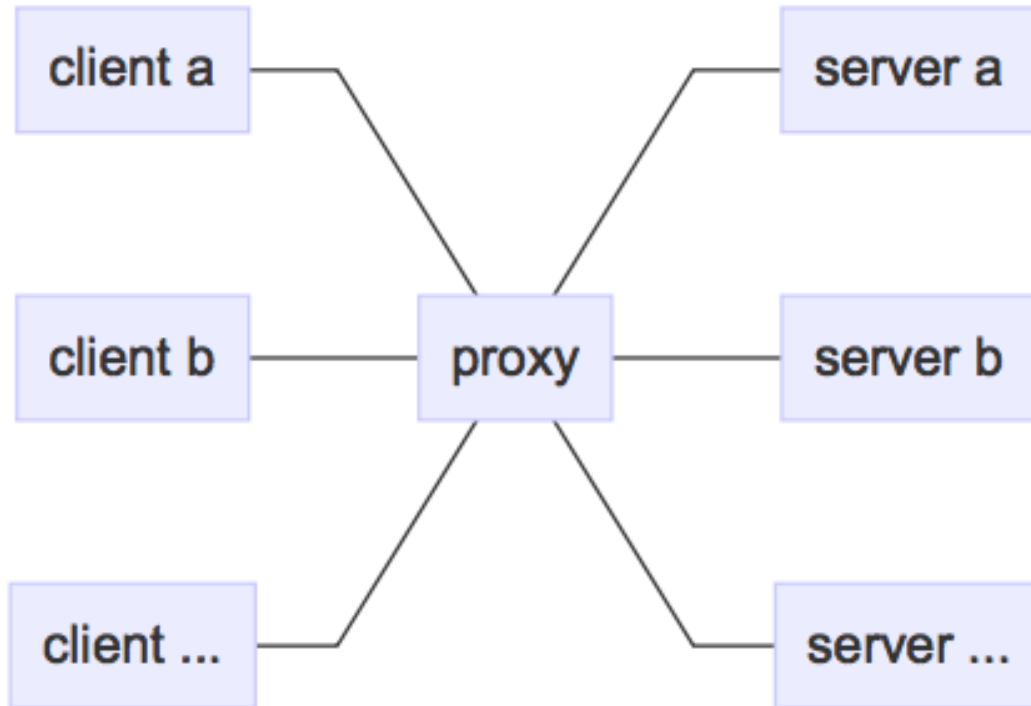
2.1 Hoverfly as a proxy server

A proxy server passes requests between a client and server.



It is sometimes necessary to use a proxy server to reach a network (as a security measure, for example). Because of this, all network-enabled software can be configured to use a proxy server.

The relationship between clients and servers via a proxy server can be one-to-one, one-to-many, many-to-one, or many-to-many.



By default Hoverfly starts as a proxy server.

2.1.1 Using a proxy server

Applications can usually be configured to use a proxy server by setting environment variables:

```
export HTTP_PROXY="http://proxy-address:port "  
export HTTPS_PROXY="https://proxy-address:port "
```

Launching network-enabled software within an environment containing these variables *should* make the application use the specified proxy server. The term *should* is used as not all software respects these environment variables for security reasons.

Alternatively, applications themselves can usually be configured to use a proxy. [Curl](#) can be configured to use a proxy via flags.

```
curl http://hoverfly.io --proxy http://proxy-ip:port
```

Note: The proxy configuration methods described here are intended to help you use the code examples in this documentation. The method of configuring an application or operating system to use a proxy varies depending on the

environment.

- [Windows Proxy Settings Explained](#)
- [Firefox Proxy Settings](#)

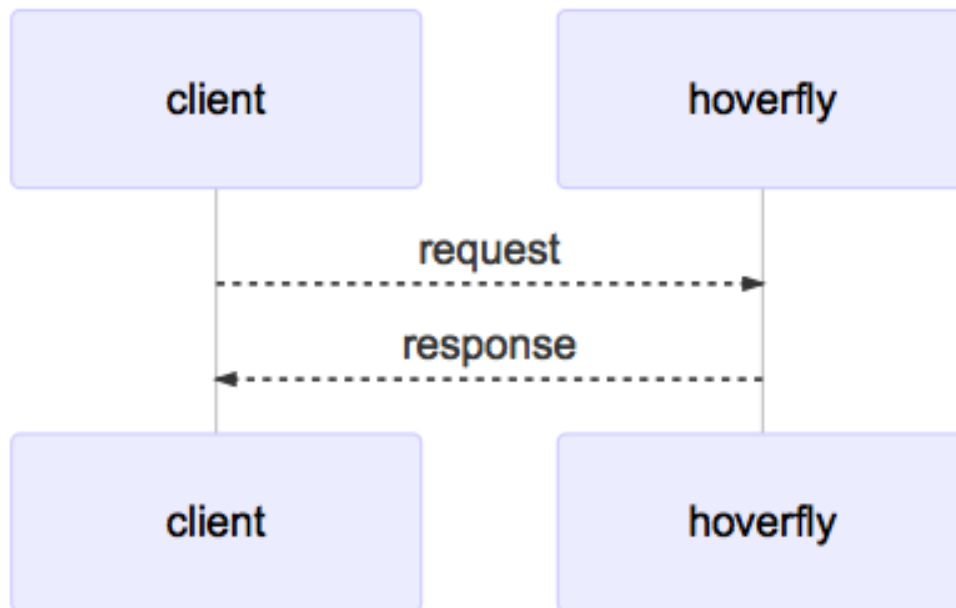
2.1.2 The difference between a proxy server and a webserver

A proxy server is a type of webserver. The main difference is that when a webserver receives a request from a client, it is expected to respond with whatever the intended response is (an HTML page, for example). The data it responds with is generally expected to reside on that server, or within the same network.

A proxy server is expected to pass the incoming request on to another server (the “destination”). It is also expected to set some appropriate headers along the way, such as `X-Forwarded-For`, `X-Real-IP`, `X-Forwarded-Proto` etc. Once the proxy server receives a response from the destination, it is expected to pass it back to the client.

2.2 Hoverfly as a webserver

Sometimes you may not be able to configure your client to use a proxy, or you may want to explicitly point your application at Hoverfly. For this reason, Hoverfly can run as a webserver.



Note: When running as a webserver, Hoverfly cannot capture traffic (see *Capture mode*) - it can only be used to simulate and synthesize APIs (see *Simulate mode*, *Modify mode* and *Synthesize mode*). For this reason, when you use Hoverfly as a webserver, you should have Hoverfly simulations ready to be loaded.

When running as a webserver, Hoverfly strips the domain from the endpoint URL. For example, if you made requests to the following URL while capturing traffic with Hoverfly running as a proxy:

```
http://echo.jsontest.com/key/value
```

And Hoverfly is running in simulate mode as a webserver on:

```
http://localhost:8888
```

Then the URL you would use to retrieve the data from Hoverfly would be:

```
http://localhost:8500/key/value
```

See also:

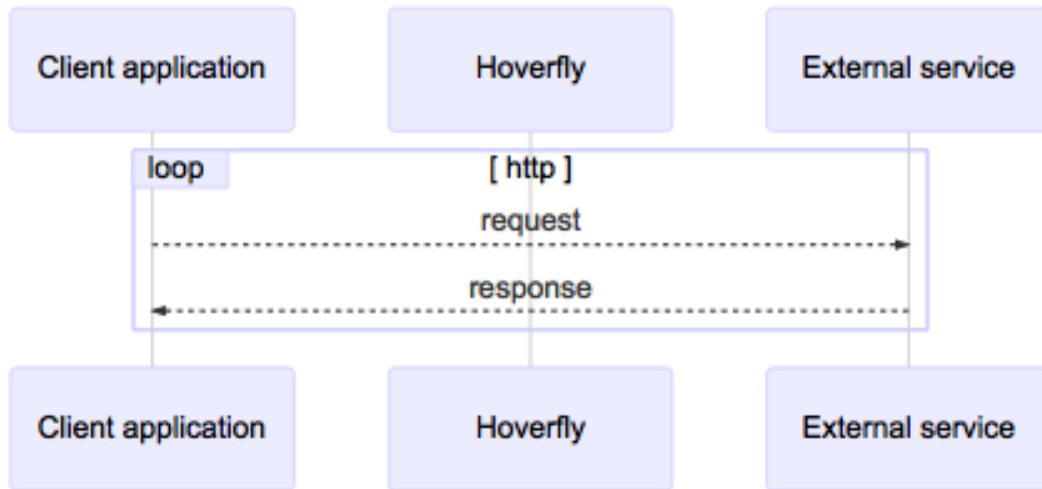
Please refer to the *Running Hoverfly as a webserver* tutorial for a step-by-step example.

2.3 Hoverfly modes

Hoverfly has four different modes. It can only run in one mode at any one time.

2.3.1 Capture mode

Capture mode is used for creating API simulations.



In Capture mode, Hoverfly (running as a proxy server - see *Hoverfly as a proxy server*) intercepts communication between the client application and the external service. It transparently records outgoing requests from the client and the incoming responses from the service API.

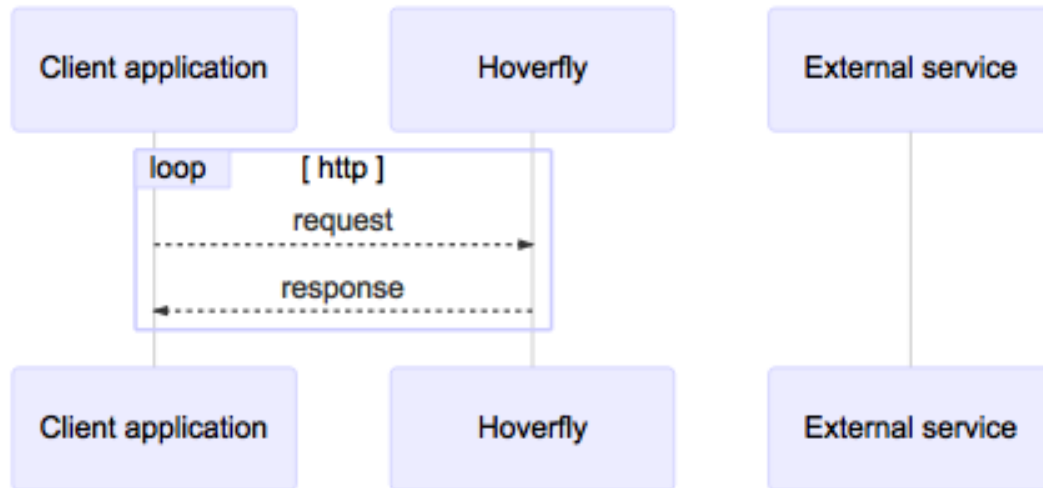
Most commonly, requests to the external service API are triggered by running automated tests against the application that consumes the API. During subsequent test runs, Hoverfly can be set to run in *Simulate mode*, removing the dependency on the real external service API. Alternatively, requests can be generated using a manual process.

Usually, Capture mode is used as the starting point in the process of creating an API simulation. Captured data is then exported and modified before being re-imported into Hoverfly for use as a simulation.

Note: Hoverfly cannot be set to Capture mode when running as a webserver (see *Hoverfly as a webserver*).

2.3.2 Simulate mode

In this mode, Hoverfly uses traffic captured using *Capture mode* (which may also have been manually edited) to mimic external APIs.



Each time Hoverfly receives a request from the application, instead of forwarding it to the intended destination, it looks in the simulation data for a matching response. If it finds a match, it returns the response to the application.


This simple matching strategy won't always be appropriate however. In some cases you may want Hoverfly to return a single response for a number of different possible requests. This can be done by editing the simulation data (see *Templates*).

2.3.3 Synthesize mode

This mode is similar to *Simulate mode*, but instead of looking for a response in stored simulation data, the request is passed directly to a user-supplied executable file. These files are known as *Middleware*.

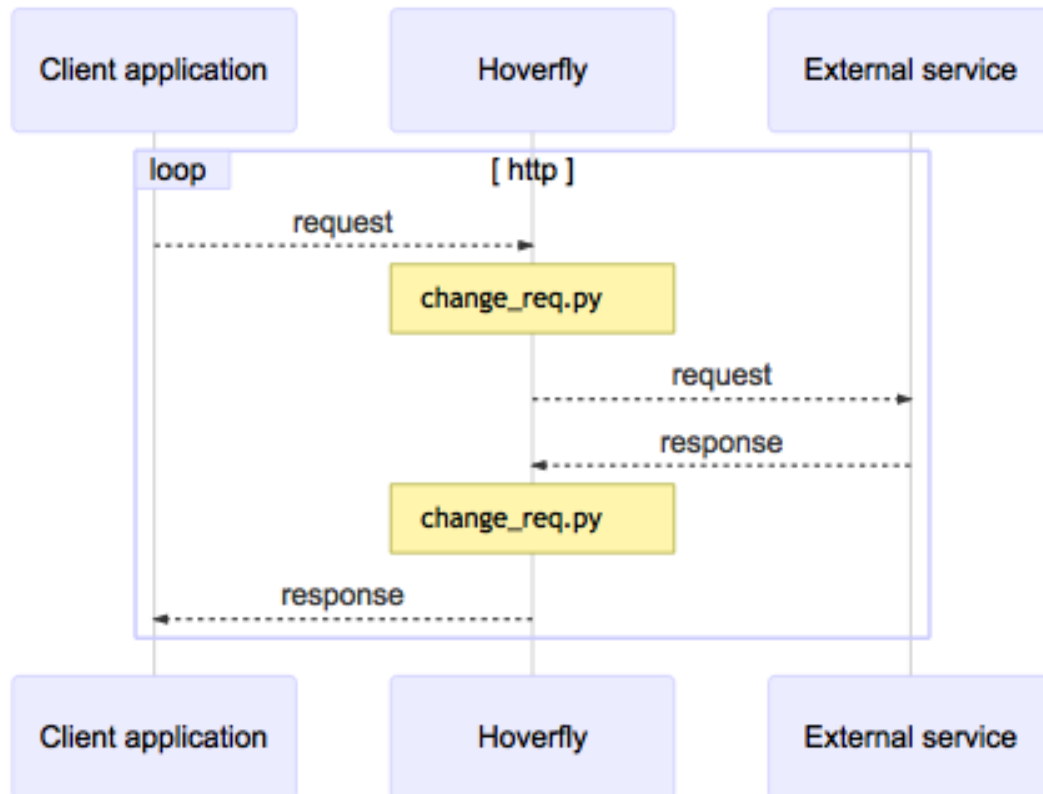
In Synthesize mode, the middleware executable is expected to generate a response to the incoming request “on the fly”. Hoverfly will then return the generated response to the application. For Hoverfly to operate in Synthesize mode, a middleware executable must be specified.

Note: You might use this mode to simulate an API that may be too difficult to record correctly via *Capture mode*. An example would be an API that uses state to change the responses. You could create middleware that manages this state and produces the desired response based on the data in the request.



2.3.4 Modify mode

Modify mode is similar to *Capture mode*, except it **does not save the requests and responses**. In Modify mode, Hoverfly will pass each request to a *Middleware* executable before forwarding it to the destination. Responses will also be passed to middleware before being returned to the client.



You could use this mode to “man in the middle” your own requests and responses. For example, you could change the API key you are using to authenticate against a third-party API.

2.4 Simulations

The core functionality of Hoverfly is to capture HTTP(S) traffic to create API simulations which can be used in testing. Hoverfly stores captured traffic as *simulations*.


Simulations can be written to disk in two different formats: in a JSON, or a BoltDB database format.

2.4.1 <simulation>.json

Simulation JSON can be exported, edited and imported in and out of Hoverfly, and can be shared among Hoverfly users or instances. Simulation JSON files must adhere to the Hoverfly *Simulation schema*.

See also:

For a hands-on tutorial of creating and editing simulations, see *Creating and exporting a simulation*.



2.4.2 requests.db

Hoverfly can also store simulation data on disk in a file called *requests.db*. This file is written to the current working directory. This means simulations do not need to be manually exported and re-imported between Hoverfly invocations.

Warning: Please note although you can persist the Hoverfly data store to disk, Hoverfly does not store all of its internal state in this database currently. The only way to access, modify and share the full state of a running instance of Hoverfly is through the simulation JSON.

This mechanism uses a very high performance Golang database system: [BoltDB](#).

2.5 Captured traffic

Hoverfly's core functionality is to capture requests and responses ("traffic") to create API simulations.

2.5.1 Request response pairs

When you capture traffic using Hoverfly's *Capture mode* and export resulting the simulation to JSON, you will see *request response pairs*:

```
{
  "response": {
    "status": 200,
    "body": "body here",
    "encodedBody": false,
    "headers": {
      "Content-Type": ["text/html; charset=utf-8"]
    }
  },
  "request": {
    "requestType": "recording",
    "path": "/",
    "method": "GET",
    "destination": "myhost.io",
    "scheme": "https",
    "query": "",
    "body": "",
    "headers": {
      "Accept": ["text/html,application/xhtml+xml,application/xml;q=0.9,image/webp,*/*;q=0.8"],
      "User-Agent": ["Mozilla/5.0 (Macintosh; Intel Mac OS X 10_11_6) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/52.0.2743.116 Safari/537.36"]
    }
  }
}
```

Notice the "response" and "request" key values in the document above. They both contain all of the fields from the original request and response. The headers are also stored.

Please also notice the "requestType" key, with its corresponding "recording" value: this denotes that the request response pair was created while using Hoverfly in *Capture mode*, and that in *Simulate mode*, Hoverfly will return this exact response when it receives this request.

Note: Since JSON does not support binary data, binary responses are base64 encoded. This is denoted by the `encodedBody` field. Hoverfly automatically encodes and decodes the data during the export and import phases.

2.5.2 Matching

Hoverfly simulates APIs by *matching* responses to incoming requests.

Imagine scanning through a dictionary for a word, and then looking up its definition. Hoverfly does exactly that, but the “word” is the URL that was “recorded” in *Capture mode*, plus all the fields in the `"request"` part of the document above, with the exception of the headers.

Note: The reason headers are not included in the match is because they can vary depending on the client.

This one-to-one matching strategy is extremely fast, but in some cases you may want Hoverfly to return a single response for more than one request. This is possible using *Templates*.

2.6 Templates

Sometimes simple one-to-one matching of responses to requests is not enough.

Request templates are defined in the *Simulation schema* by setting the `"requestType"` property for a request to `"template"` and including only the information in the request that you want Hoverfly to use in the match.

In the example below, Hoverfly will return the same response for any request with the path `/template`:

```
{
  "data": {
    "pairs": [{
      "response": {
        "status": 200,
        "body": "<h1>Matched on template</h1>",
        "encodedBody": false,
        "headers": {
          "Content-Type": ["text/html; charset=utf-8"]
        }
      },
      "request": {
        "requestType": "template",
        "path": "/template"
      }
    }
  ]
}
```

For looser matching on URL paths, it is possible to use a wildcard to substitute characters. This is achieved by using the `*` symbol. This will match any number of characters, and is case sensitive.

In the next example, Hoverfly will return the same response for requests with the path `/api/v1/template` or `/api/v2/template`.

```
{
  "data": {
```

```

    "pairs": [{
      "response": {
        "status": 200,
        "body": "<h1>Matched on template</h1>",
        "encodedBody": false,
        "headers": {
          "Content-Type": ["text/html; charset=utf-8"]
        }
      },
      "request": {
        "requestType": "template",
        "path": "/api/*/template"
      }
    }
  ]
}

```

The *Simulation schema* can contain both request recordings and request templates:

```

{
  "data": {
    "pairs": [{
      "response": {
        "status": 200,
        "body": "<h1>Matched on recording</h1>",
        "encodedBody": false,
        "headers": {
          "Content-Type": [
            "text/html; charset=utf-8"
          ]
        }
      },
      "request": {
        "requestType": "recording",
        "path": "/",
        "method": "GET",
        "destination": "myhost.io",
        "scheme": "https",
        "query": "",
        "body": "",
        "headers": {
          "Accept": [
            "text/html,application/xhtml+xml,application/xml;q=0.9,image/
↪webp,*/*;q=0.8"
          ],
          "Content-Type": [
            "text/plain; charset=utf-8"
          ],
          "User-Agent": [
            "Mozilla/5.0 (Macintosh; Intel Mac OS X 10_11_6) AppleWebKit/
↪537.36 (KHTML, like Gecko) Chrome/52.0.2743.116 Safari/537.36"
          ]
        }
      }
    }
  ], {
    "response": {
      "status": 200,
      "body": "<h1>Matched on template</h1>",

```

```
        "encodedBody": false,
        "headers": {
          "Content-Type": [
            "text/html; charset=utf-8"
          ]
        }
      },
      "request": {
        "requestType": "template",
        "path": "/template",
        "method": null,
        "destination": null,
        "scheme": null,
        "query": null,
        "body": null,
        "headers": null
      }
    }
  ],
  "globalActions": {
    "delays": []
  }
}
```

A standard workflow might be:

1. Capture some traffic
2. Export it to JSON
3. Edit the JSON to set certain requests to templates, removing the properties for these requests that should be excluded from the match or substituting characters in the URL path for wildcards
4. Re-import the JSON to Hoverfly

Note: If the "requestType" property is not defined or not recognized, Hoverfly will treat a request as a "recording".

See also:

Templating is best understood with a practical example, so please refer to [Adding templates to a simulation](#) to get hands on experience with templating.

2.7 Destination filtering

By default, Hoverfly will process every request it receives. However, you may wish to control which URLs Hoverfly processes.

This is done by *filtering* the *destination* URLs using either a string or a regular expression. The *destination* string or regular expression will be compared against the host and the path of a URL.

For example, specifying `hoverfly.io` as the destination value will tell Hoverfly to process only URLs on the `hoverfly.io` host.

```
hoverctl destination "hoverfly.io"
```

Specifying `api` as the *destination* value during *Capture mode* will tell Hoverfly to capture only URLs that contain the string `api`. This would include both `api.hoverfly.io/endpoint` and `hoverfly.io/api/endpoint`.

See also:

This functionality is best understood via a practical example: see *Capturing or simulating specific URLs* in the *Tutorials* section.

Note: The *destination* setting applies to all Hoverfly modes. If a *destination* value is set while Hoverfly is running in *Simulate mode*, requests that are excluded by the *destination* setting will be passed through to the real URLs. This makes it possible to return both real and simulated responses.

2.8 Meta

The last part of the simulation schema is the meta object. Its purpose is to store metadata that is relevant to your simulation. This includes the simulation schema version, the version of Hoverfly used to export the simulation and the date and time at which the simulation was exported.

```
"meta": {
  "schemaVersion": "v1",
  "hoverflyVersion": "v0.9.0",
  "timeExported": "2016-11-11T11:53:52Z"
}
```

2.9 Delays

Once you have created a simulated service by capturing traffic between your application and an external service, you may wish to make the simulation more “realistic” by applying latency to the responses returned by Hoverfly.

Hoverfly can be configured to apply delays to responses based on URL pattern matching or HTTP method. This is done using a regular expression to match against the URL, a delay value in milliseconds, and an optional HTTP method value.

See also:

This functionality is best understood via a practical example: see *Adding delays to a simulation* in the *Tutorials* section.

You can also apply delays to simulations using *Middleware* (see the *Using middleware to simulate network latency* tutorial). Using middleware to apply delays sacrifices performance for flexibility.

2.10 Middleware

Middleware intercepts traffic between the client and the API (whether real or simulated), and allowing you to manipulate it.

You can use middleware to manipulate data in simulated responses, or to inject unpredictable performance characteristics into your simulation.

Middleware works differently depending on the Hoverfly mode.


- Capture mode: middleware affects only **outgoing requests**
- Simulate mode: middleware affects only **incoming responses** (cache contents remain untouched)

- Synthesize mode: middleware **creates responses**
- Modify mode: middleware affects **requests and responses**

You can write middleware in any language as long as it can be executed by the shell on the Hoverfly host. Middleware could be a binary file, or a script.

2.10.1 Middleware Interface

When middleware is called by Hoverfly, it expects to receive and return JSON (see *Simulation schema*). Middleware can be used to modify the values in the JSON but **must not** modify the schema itself.



Hoverfly will send the JSON object to middleware via the standard input stream. Hoverfly will then listen to the standard output stream and wait for the JSON object to be returned.

See also:

Middleware examples are covered in the tutorials section. See *Using middleware to simulate network latency* and *Using middleware to modify response payload and status code*.

Native language bindings

Native language bindings are available for Hoverfly to make it easy to integrate into different environments.

3.1 HoverPy

To get started:

```
sudo pip install hoverpy
python
```

And in Python you can simply get started with:

```
import hoverpy
import requests

# capture mode
with hoverpy HoverPy(capture=True) as hp:
    data = requests.get("http://time.jsontest.com/").json()

# simulation mode
with hoverpy HoverPy() as hp:
    simData = requests.get("http://time.jsontest.com/").json()
    print(simData)

assert(data["milliseconds_since_epoch"] == simData["milliseconds_since_epoch"])
```

For more information, read the [HoverPy documentation](#).

3.2 Hoverfly Java

- Strict or loose HTTP request matching based on URL, method, body and header combinations

- Fluent and expressive DSL for easy generation of simulated APIs
- Automatic marshalling of objects into JSON during request/response body generation
- HTTPS automatically supported, no extra configuration required
- Download via Maven or Gradle

To get started, read the [Hoverfly Java documentation](#).

In these examples, we will use the `hoverctl` (CLI tool for Hoverfly) to interact with Hoverfly.

Hoverfly can also be controlled via its *REST API*, or via *Native language bindings*.

4.1 Creating and exporting a simulation

Note: If you are running Hoverfly on a machine that accesses the internet via a proxy (for example if you are on a corporate network), please follow the *Using Hoverfly behind a proxy* tutorial before proceeding.

Start Hoverfly and set it to Capture mode

```
hoverctl start
hoverctl mode capture
```

Make a request with cURL, using Hoverfly as a proxy server:

```
curl --proxy http://localhost:8500 http://time.jsontest.com
```

View the Hoverfly logs

```
hoverctl logs
```

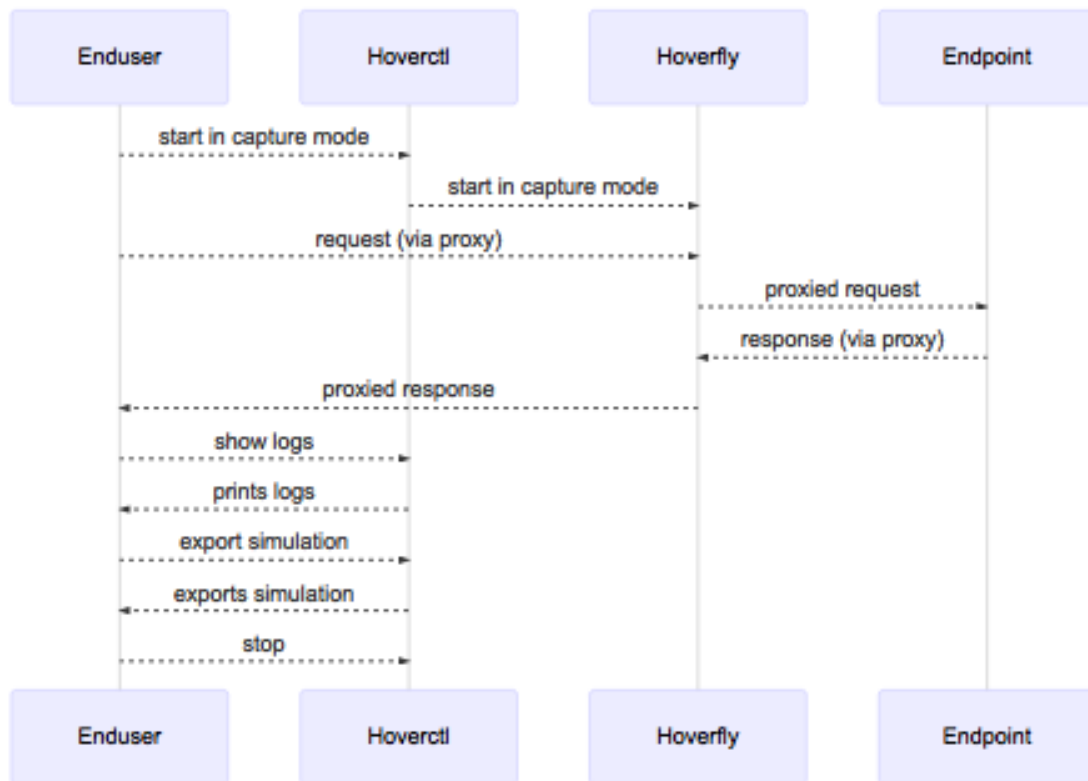
Export the simulation to a JSON file

```
hoverctl export simulation.json
```

Stop hoverfly

```
hoverctl stop
```

You'll now see a `simulation.json` file in your current working directory, which contains all your simulation data. In case you are curious, the sequence diagram for this process looks like this:



4.2 Importing and using a simulation

In this tutorial we are going to import the simulation we created in the previous tutorial.

```
hoverctl start
hoverctl import simulation.json
```

Hoverfly can also import simulation data that is stored on a remote host via HTTP:

```
hoverctl import https://example.com/example.json
```

Make a request with cURL, using Hoverfly as a proxy.

```
curl --proxy localhost:8500 http://time.jsontest.com
```

This outputs the time at the time the request was captured.

```
{
  "time": "02:07:28 PM",
  "milliseconds_since_epoch": 1482242848562,
  "date": "12-20-2016"
}
```

Stop Hoverfly:

```
hoverctl stop
```

4.3 Adding delays to a simulation

Simulating API latency during development allows you to write code that will deal with it gracefully.

In Hoverfly, this is done by applying “delays” to responses in a simulation.

Delays are applied by editing the Hoverfly simulation JSON file. Delays can be applied selectively according to request URL pattern and/or HTTP method.

4.3.1 Applying a delay to all responses

Let’s apply a 2 second delay to all responses. First, we need to create and export a simulation.

```
hoverctl start
hoverctl mode capture
curl --proxy localhost:8500 http://time.jsontest.com
curl --proxy localhost:8500 http://date.jsontest.com
hoverctl export simulation.json
hoverctl stop
```

Take a look at the "globalActions" property within the simulation.json file you exported. It should look like this:

```
"globalActions": {
  "delays": []
}
```

Edit the file so the "globalActions" property looks like this:

```
"globalActions": {
  "delays": [
    {
      "urlPattern": ".",
      "delay": 2000
    }
  ]
}
```

Hoverfly will apply a delay of 2000ms to all URLs that match the "urlPattern" value. We want the delay to be applied to **all URLs**, so we set the "urlPattern" value to the regular expression ".".

Now import the edited simulation.json file, switch Hoverfly to Simulate mode and make the requests again.

```
hoverctl start
hoverctl import simulation.json
curl --proxy localhost:8500 http://time.jsontest.com
curl --proxy localhost:8500 http://date.jsontest.com
hoverctl stop
```

The responses to both requests are delayed by 2 seconds.

4.3.2 Applying different delays based on host

Now let's apply a delay of 1 second on responses from `time.jsontest.com` and a delay of 2 seconds on responses from `date.jsontest.com`.

Run the following to create and export a simulation.

```
hoverctl start
hoverctl mode capture
curl --proxy localhost:8500 http://time.jsontest.com
curl --proxy localhost:8500 http://date.jsontest.com
hoverctl export simulation.json
hoverctl stop
```

Edit the `simulation.json` file so that the "globalActions" property looks like this:

```
    "globalActions": {
      "delays": [
        {
          "urlPattern": "time\\.jsontest\\.com",
          "delay": 1000
        },
        {
          "urlPattern": "date\\.jsontest\\.com",
          "delay": 2000
        }
      ]
    }
```

Now run the following to import the edited `simulation.json` file and run the simulation:

```
hoverctl start
hoverctl import simulation.json
curl --proxy localhost:8500 http://time.jsontest.com
curl --proxy localhost:8500 http://date.jsontest.com
hoverctl stop
```

You should notice a 1 second delay on responses from `time.jsontest.com`, and a 2 second delay on responses from `date.jsontest.com`.

Note: You can easily get into a situation where your request URL has multiple matches. In this case, the first successful match wins.

4.3.3 Applying different delays based on URI

Now let's apply different delays based on location. Run the following to create and export a simulation.

```
hoverctl start
hoverctl mode capture
curl --proxy localhost:8500 http://echo.jsontest.com/a/b
curl --proxy localhost:8500 http://echo.jsontest.com/b/c
curl --proxy localhost:8500 http://echo.jsontest.com/c/d
hoverctl export simulation.json
hoverctl stop
```

Edit the `simulation.json` file so that the `"globalActions"` property looks like this:

```

    "globalActions": {
      "delays": [
        {
          "urlPattern": "echo\\.jsontest\\.com\\/a\\/b",
          "delay": 2000
        },
        {
          "urlPattern": "echo\\.jsontest\\.com\\/b\\/c",
          "delay": 2000
        },
        {
          "urlPattern": "echo\\.jsontest\\.com\\/c\\/d",
          "delay": 3000
        }
      ]
    }

```

Now run the following to import the edited `simulation.json` file and run the simulation:

```

hoverctl start
hoverctl import simulation.json
hoverctl mode simulate
curl --proxy localhost:8500 http://echo.jsontest.com/a/b
curl --proxy localhost:8500 http://echo.jsontest.com/b/c
curl --proxy localhost:8500 http://echo.jsontest.com/c/d
hoverctl stop

```

You should notice a 2 second delay on responses from `echo.jsontest.com/a/b` and `echo.jsontest.com/b/c`, and a 3 second delay on the response from `echo.jsontest.com/c/d`.

4.3.4 Applying different delays based on HTTP method

Let's apply a delay of 2 seconds on responses to **GET requests only** made to `echo.jsontest.com/b/c`.

Run the following to create and export a simulation.

```

hoverctl start
hoverctl mode capture
curl --proxy localhost:8500 http://echo.jsontest.com/b/c
curl --proxy localhost:8500 -X POST http://echo.jsontest.com/b/c
hoverctl export simulation.json
hoverctl stop

```

Edit the `simulation.json` file so that the `"globalActions"` property looks like this:

```

    "globalActions": {
      "delays": [
        {
          "urlPattern": "echo\\.jsontest\\.com\\/b\\/c",
          "delay": 2000,
          "httpMethod": "GET"
        }
      ]
    }

```

Now run the following to import the edited `simulation.json` file and run the simulation:

```
hoverctl start
hoverctl import simulation.json
curl --proxy localhost:8500 http://echo.jsontest.com/b/c
curl -X POST --proxy localhost:8500 http://echo.jsontest.com/b/c
hoverctl stop
```

You should notice a 2 second delay on the response to the GET request and no delay on the response to the POST request.

4.4 Adding templates to a simulation

See also:

Please carefully read through *Templates* alongside this tutorial to gain a high-level understanding of what we are about to cover.

In this tutorial, we are going to go through the steps required to generate and use a matching template.

Let's begin by capturing some traffic and exporting a simulation.

```
hoverctl start
hoverctl mode capture
curl --proxy http://localhost:8500 http://echo.jsontest.com/foo/baz/bar/spam
hoverctl export simulation.json
hoverctl stop
```

Which gives us this output:

```
time="2016-12-22T08:56:24Z" level=info msg="Hoverfly is now running" admin-
↪port=8888 proxy-port=8500
time="2016-12-22T08:56:24Z" level=info msg="Hoverfly has been set to capture ↪
↪mode"
  % Total      % Received % Xferd  Average Speed   Time    Time       Time  ↪
↪Current
                               Dload  Upload  Total  Spent    Left  Speed
  0     0    0     0     0     0     0     0  --:--:--  --:--:--  --:--:--  ↪
↪0
  0     0    0     0     0     0     0     0  --:--:--  --:--:--  --:--:--  ↪
↪0
100   38 100   38     0     0   167     0  --:--:--  --:--:--  --:--:--  ↪
↪168
{
  "foo": "baz",
  "bar": "spam"
}
time="2016-12-22T08:56:25Z" level=info msg="Successfully exported to ↪
↪simulation.json"
time="2016-12-22T08:56:25Z" level=info msg="Hoverfly has been stopped"
```

If you take a look at your `simulation.json` you should notice these lines in your request.

```
"request": {
  "requestType": "recording",
```

```
"path": "/foo/baz/bar/spam",  
"method": "GET",
```

Modify them to:

```
"request": {  
  "requestType": "template",  
  "path": "/foo/*/bar/spam",  
  "method": "GET",
```

Save the file as `simulationimport.json` and run the following command to import it and cURL the simulated endpoint:

```
hoverctl start  
hoverctl mode simulate  
hoverctl import simulationimport.json  
curl --proxy http://localhost:8500 http://echo.jsonstest.com/foo/QUX/bar/spam  
hoverctl stop
```

The same response is returned, even though we created our simulation with a request to `http://echo.jsonstest.com/foo/baz/bar/spam` in Capture mode and then sent a request to `http://echo.jsonstest.com/foo/QUX/bar/spam` in Simulate mode.

As you can see, templating allows us to match URLs using [globbing](#).

Note: Key points:

- To do templating, capture a simulation, export it, edit it
- While editing, change "requestTypes" value to "template"
- Substitute strings in URLs with the wildcard * to return one response for more than one request
- Re-import the simulation

4.5 Using middleware to simulate network latency

See also:

Please carefully read through [Middleware](#) alongside these tutorials to gain a high-level understanding of what we are about to cover.

We will use a Python script to apply a random delay of less than one second to every response in a simulation.

Before you proceed, please ensure that you have Python installed.

Let's begin by writing our middleware. Save the following as `middleware.py`:

```
#!/usr/bin/env python  
import sys  
import logging  
import random  
from time import sleep  
  
logging.basicConfig(filename='random_delay_middleware.log', level=logging.DEBUG)  
logging.debug('Random delay middleware is called')
```

```
# set delay to random value less than one second

SLEEP_SECS = random.random()

def main():

    data = sys.stdin.readlines()
    # this is a json string in one line so we are interested in that one line
    payload = data[0]
    logging.debug("sleeping for %s seconds" % SLEEP_SECS)
    sleep(SLEEP_SECS)

    # do not modifying payload, returning same one
    print(payload)

if __name__ == "__main__":
    main()
```

The middleware script delays each response by a random value of less than one second.

```
hoverctl start
hoverctl mode capture
curl --proxy http://localhost:8500 http://time.jsontest.com
hoverctl mode simulate
hoverctl middleware --binary python --script middleware.py
curl --proxy http://localhost:8500 http://time.jsontest.com
hoverctl stop
```

Middleware gives you control over the behaviour of a simulation, as well as the data.

Note: Middleware gives you flexibility when simulating network latency - allowing you to randomize the delay value for example - but a new process is spawned every time the middleware script is executed. This can impact Hoverfly's performance under load.

If you need to simulate latency during a load test, it is recommended that you use Hoverfly's native *Delays* functionality to simulate network latency (see *Adding delays to a simulation*) instead of writing middleware. The delays functionality sacrifices flexibility for performance.

4.6 Using middleware to modify response payload and status code

See also:

Please carefully read through *Middleware* alongside these tutorials to gain a high-level understanding of what we are about to cover.

We will use a python script to modify the body of a response and randomly change the status code.

Let's begin by writing our middleware. Save the following as `middleware.py`:

```
#!/usr/bin/env python

import sys
import json
import logging
```



```

import random

logging.basicConfig(filename='middleware.log', level=logging.DEBUG)
logging.debug('Middleware "modify_request" called')

def main():
    payload = sys.stdin.readlines()[0]

    logging.debug(payload)

    payload_dict = json.loads(payload)
    payload_dict['response']['status'] = random.choice([200, 201])

    if "response" in payload_dict and "body" in payload_dict["response"]:
        payload_dict["response"]["body"] = '{"foo': 'baz'}\n"

    print(json.dumps(payload_dict))

if __name__ == "__main__":
    main()

```

The middleware script randomly toggles the status code between 200 and 201, and changes the response body to a dictionary containing `{ 'foo': 'baz' }`.

```

hoverctl start
hoverctl mode capture
curl --proxy http://localhost:8500 http://time.jsontest.com
hoverctl mode simulate
hoverctl middleware --binary python --script middleware.py
curl --proxy http://localhost:8500 http://time.jsontest.com
hoverctl stop

```

As you can see, middleware allows you to completely modify the content of a simulated HTTP response.

4.7 Simulating HTTPS APIs

To capture HTTPS traffic, you need to use Hoverfly's SSL certificate.

First, download the certificate:

```
wget https://raw.githubusercontent.com/SpectoLabs/hoverfly/master/core/cert.pem
```

We can now run Hoverfly with the standard capture then simulate workflow.

```

hoverctl start
hoverctl mode capture
curl --proxy https://localhost:8500 https://example.com --cacert cert.pem
hoverctl mode simulate
curl --proxy https://localhost:8500 https://example.com --cacert cert.pem
hoverctl stop

```

Curl was able to make the HTTPS request using an HTTPS proxy because we provided it with Hoverfly's SSL certificate.

Note: This example uses Curl. If you are using Hoverfly in another environment, you will need to add the certificate

to your trust store. This is done automatically by the Hoverfly Java library (see *Hoverfly Java*).

4.8 Running Hoverfly as a webserver

See also:

Please carefully read through *Hoverfly as a webserver* alongside this tutorial to gain a high-level understanding of what we are about to cover.

Below is a complete example how to capture data with Hoverfly running as a proxy, and how to save it in a simulation file.

```
hoverctl start
hoverctl mode capture
curl --proxy http://localhost:8500 http://echo.jsonstest.com/a/b
hoverctl export simulation.json
hoverctl stop
```

Now we can use Hoverfly as a **webserver** in Simulate mode.

```
hoverctl start webserver
hoverctl import simulation.json
curl http://localhost:8500/a/b
hoverctl stop
```

Hoverfly returned a response to our request while running as a webserver, not as a proxy.

Notice that we issued a cURL command to `http://localhost:8500/a/b` instead of `http://echo.jsonstest.com/a/b`. This is because when running as a webserver, Hoverfly strips the domain from the endpoint URL in the simulation.

This is explained in more detail in the *Hoverfly as a webserver* section.

Note: Hoverfly starts in Simulate mode by default.

4.9 Capturing or simulating specific URLs

We can use the `hoverctl destination` command to specify which URLs to capture or simulate. The destination setting can be tested using the `--dry-run` flag. This makes it easy to check whether the destination setting will filter out URLs as required.

```
hoverctl start
hoverctl destination "ip" --dry-run http://ip.jsonstest.com
hoverctl destination "ip" --dry-run http://time.jsonstest.com
hoverctl stop
```

This tells us that setting the destination to `ip` will allow the URL `http://ip.jsonstest.com` to be captured or simulated, while the URL `http://time.jsonstest.com` will be ignored.

Now we have checked the destination setting, we can apply it to filter out the URLs we don't want to capture.

```
hoverctl start
hoverctl destination "ip"
hoverctl mode capture
curl --proxy http://localhost:8500 http://ip.jsontest.com
curl --proxy http://localhost:8500 http://time.jsontest.com
hoverctl logs
hoverctl export simulation.json
hoverctl stop
```

If we examine the logs and the `simulation.json` file, we can see that *only* a request response pair to the `http://ip.jsontest.com` URL has been captured.

The destination setting can be either a string or a regular expression.

```
hoverctl start
hoverctl destination "^.*api.*com" --dry-run https://api.github.com
hoverctl destination "^.*api.*com" --dry-run https://api.slack.com
hoverctl destination "^.*api.*com" --dry-run https://github.com
hoverctl stop
```

Here, we can see that setting the destination to `^.*api.*com` will allow the `https://api.github.com` and `https://api.slack.com` URLs to be captured or simulated, while the `https://github.com` URL will be ignored.

4.10 Using Hoverfly behind a proxy

In some environments, you may only be able to access the internet via a proxy. For example, your organization may route all traffic through a proxy for security reasons.

If this is the case, you will need to configure Hoverfly to work with the ‘upstream’ proxy.

This configuration value can be easily set when starting an instance of Hoverfly.

```
hoverctl start --upstream-proxy http://localhost:8080
```

Currently, Hoverfly will only work with unauthenticated proxies.

These reference documents contain information regarding invoking the `hoverctl` command, `hoverfly` command, and interacting with the APIs.

5.1 `hoverctl` commands

This page contains the output of:

```
hoverctl --help
```

The command's help content has been placed here for convenience.

```
usage: hoverctl [<flags>] <command> [<args> ...]

Flags:
  --help                Show context-sensitive help (also try
                        --help-long and --help-man).
  -v, --verbose         Verbose mode.
  --host=HOST           Set the host of Hoverfly
  --admin-port=ADMIN-PORT Set the admin port of Hoverfly
  --proxy-port=PROXY-PORT Set the proxy port of Hoverfly
  --certificate=CERTIFICATE Supply path for custom certificate
  --key=KEY             Supply path for custom key
  --disable-tls        Disable TLS verification
  --version            Show application version.

Commands:
  help [<command>...]
    Show help.

  mode [<name>]
    Get Hoverfly's current mode

  destination [<flags>] [<name>]
```

```

    Get Hoverfly's current destination

middleware [<path>]
    Get Hoverfly's middleware

start [<server type>]
    Start a local instance of Hoverfly

stop
    Stop a local instance of Hoverfly

export <name>
    Exports data out of Hoverfly

import [<flags>] <name>
    Imports data into Hoverfly

delete [<resource>]
    Delete test data from Hoverfly

delays [<path>]
    Get per-host response delay config currently loaded in Hoverfly

logs [<flags>]
    Get the logs from Hoverfly

templates [<path>]
    Get set of request templates currently loaded in Hoverfly

config
    Get the config being used by hoverctl and Hoverfly

```

5.2 Hoverfly commands

This page contains the output of:

```
hoverfly --help
```

The command's help content has been placed here for convenience.

```

Usage of hoverfly:
  -add
        add new user '-add -username hfdadmin -password hfpass'
  -admin
        supply '-admin false' to make this non admin user (defaults to 'true') ↵
↪(default true)
  -ap string
        admin port - run admin interface on another port (i.e. '-ap 1234' to run ↵
↪admin UI on port 1234)
  -auth
        enable authentication, currently it is disabled by default
  -capture
        start Hoverfly in capture mode - transparently intercepts and saves ↵
↪requests/response
  -cert string

```

```

    CA certificate used to sign MITM certificates
-cert-name string
    cert name (default "hoverfly.proxy")
-cert-org string
    organisation name for new cert (default "Hoverfly Authority")
-db string
    Persistence storage to use - 'boltdb' or 'memory' which will not write
↳ anything to disk (default "boltdb")
-db-path string
    database location - supply it to provide specific database location (will
↳ be created there if it doesn't exist)
-dest value
    specify which hosts to process (i.e. '-dest fooservice.org -dest
↳ barservice.org -dest catservice.org') - other hosts will be ignored will passthrough
↳
-destination string
    destination URI to catch (default ".")
-dev
    supply -dev flag to serve directly from ./static/dist instead from statik
↳
binary
-generate-ca-cert
    generate CA certificate and private key for MITM
-httpptest.serve string
    if non-empty, httpptest.NewServer serves on this address and blocks
-import value
    import from file or from URL (i.e. '-import my_service.json' or '-import
↳ http://mypage.com/service_x.json')
-key string
    private key of the CA used to sign MITM certificates
-metrics
    supply -metrics flag to enable metrics logging to stdout
-middleware string
    should proxy use middleware
-modify
    start Hoverfly in modify mode - applies middleware (required) to both
↳ outgoing and incoming HTTP traffic
-password string
    password for new user
-pp string
    proxy port - run proxy on another port (i.e. '-pp 9999' to run proxy on
↳ port 9999)
-synthesize
    start Hoverfly in synthesize mode (middleware is required)
-test.bench string
    regular expression per path component to select benchmarks to run
-test.benchmem
    print memory allocations for benchmarks
-test.benchmark duration
    approximate run time for each benchmark (default 1s)
-test.blockprofile string
    write a goroutine blocking profile to the named file after execution
-test.blockprofilerate int
    if >= 0, calls runtime.SetBlockProfileRate() (default 1)
-test.count n
    run tests and benchmarks n times (default 1)
-test.coverprofile string
    write a coverage profile to the named file after execution
-test.cpu string

```

```

        comma-separated list of number of CPUs to use for each test
-test.cpubprofile string
        write a cpu profile to the named file during execution
-test.memprofile string
        write a memory profile to the named file after execution
-test.memprofilerate int
        if >=0, sets runtime.MemProfileRate
-test.outputdir string
        directory in which to write profiles
-test.parallel int
        maximum test parallelism (default 8)
-test.run string
        regular expression to select tests and examples to run
-test.short
        run smaller test suite to save time
-test.timeout duration
        if positive, sets an aggregate time limit for all tests
-test.trace string
        write an execution trace to the named file after execution
-test.v
        verbose: print additional output
-tls-verification
        turn on/off tls verification for outgoing requests (will not try to
↪verify certificates) - defaults to true (default true)
-username string
        username for new user
-v
        should every proxy request be logged to stdout
-version
        get the version of hoverfly
-webserver
        start Hoverfly in webserver mode (simulate mode)

```

5.3 REST API

5.3.1 GET /api/v2/simulation

Gets all simulation data. The simulation JSON contains all the information Hoverfly can hold; this includes recordings, templates, delays and metadata.

Example response body:

```

{
  "data": {
    "pairs": [
      {
        "response": {
          "status": 200,
          "body": "<h1>Matched on recording</h1>",
          "encodedBody": false,
          "headers": {
            "Content-Type": [
              "text/html; charset=utf-8"
            ]
          }
        }
      },

```



```

    "request": {
      "requestType": "recording",
      "path": "/",
      "method": "GET",
      "destination": "myhost.io",
      "scheme": "https",
      "query": "",
      "body": "",
      "headers": {
        "Accept": [
          ↪q=0.8"
          "text/html,application/xhtml+xml,application/xml;q=0.9,image/webp,*/*;
          ],
          "Content-Type": [
            "text/plain; charset=utf-8"
          ],
          "User-Agent": [
            ↪(KHTML, like Gecko) Chrome/52.0.2743.116 Safari/537.36"
            "Mozilla/5.0 (Macintosh; Intel Mac OS X 10_11_6) AppleWebKit/537.36_
          ]
        ]
      }
    },
    {
      "response": {
        "status": 200,
        "body": "<h1>Matched on template</h1>",
        "encodedBody": false,
        "headers": {
          "Content-Type": [
            "text/html; charset=utf-8"
          ]
        }
      }
    },
    "request": {
      "requestType": "template",
      "path": "/template",
      "method": null,
      "destination": null,
      "scheme": null,
      "query": null,
      "body": null,
      "headers": null
    }
  ],
  "globalActions": {
    "delays": []
  }
},
"meta": {
  "schemaVersion": "v1",
  "hoverflyVersion": "v0.9.0",
  "timeExported": "2016-11-11T11:53:52Z"
}

```

5.3.2 PUT /api/v2/simulation

This puts the supplied simulation JSON into Hoverfly, overwriting any existing simulation data.

Example request body:

```
{
  "data": {
    "pairs": [
      {
        "response": {
          "status": 200,
          "body": "<h1>Matched on recording</h1>",
          "encodedBody": false,
          "headers": {
            "Content-Type": [
              "text/html; charset=utf-8"
            ]
          }
        },
        "request": {
          "requestType": "recording",
          "path": "/",
          "method": "GET",
          "destination": "myhost.io",
          "scheme": "https",
          "query": "",
          "body": "",
          "headers": {
            "Accept": [
              "text/html,application/xhtml+xml,application/xml;q=0.9,image/webp,*/*;
↪q=0.8"
            ],
            "Content-Type": [
              "text/plain; charset=utf-8"
            ],
            "User-Agent": [
              "Mozilla/5.0 (Macintosh; Intel Mac OS X 10_11_6) AppleWebKit/537.36_
↪(KHTML, like Gecko) Chrome/52.0.2743.116 Safari/537.36"
            ]
          }
        }
      }
    ],
    {
      "response": {
        "status": 200,
        "body": "<h1>Matched on template</h1>",
        "encodedBody": false,
        "headers": {
          "Content-Type": [
            "text/html; charset=utf-8"
          ]
        }
      },
      "request": {
        "requestType": "template",
        "path": "/template",
        "method": null,
        "destination": null,

```

```
    "scheme": null,
    "query": null,
    "body": null,
    "headers": null
  }
},
"globalActions": {
  "delays": []
}
},
"meta": {
  "schemaVersion": "v1",
  "hoverflyVersion": "v0.9.0",
  "timeExported": "2016-11-11T11:53:52Z"
}
```

5.3.3 GET /api/v2/hoverfly

Gets configuration information from the running instance of Hoverfly.

Example response body:

```
{
  "destination": ".",
  "middleware": {
    "binary": "python",
    "script": "# a python script would go here",
    "remote": ""
  },
  "mode": "simulate",
  "usage": {
    "counters": {
      "capture": 0,
      "modify": 0,
      "simulate": 0,
      "synthesize": 0
    }
  }
}
```

5.3.4 GET /api/v2/hoverfly/destination

Gets the current destination setting for the running instance of Hoverfly.

Example response body:

```
{
  destination: "."
}
```

5.3.5 PUT /api/v2/hoverfly/destination

Sets a new destination for the running instance of Hoverfly, overwriting the existing destination setting.

Example request body:

```
{
  destination: "new-destination"
}
```

5.3.6 GET /api/v2/hoverfly/middleware

Gets the middleware settings for the running instance of Hoverfly. This could be either an executable binary, a script that can be executed with a binary or a URL to remote middleware.

Example response body:

```
{
  "binary": "python",
  "script": "#python code goes here",
  "remote": ""
}
```

5.3.7 PUT /api/v2/hoverfly/middleware

Sets new middleware, overwriting the existing middleware for the running instance of Hoverfly. The middleware being set can be either an executable binary located on the host, a script and the binary to execute it or the URL to a remote middleware.

Example request body:

```
{
  "binary": "python",
  "script": "#python code goes here",
  "remote": ""
}
```

5.3.8 GET /api/v2/hoverfly/mode

Gets the mode for the running instance of Hoverfly.

Example response body:

```
{
  mode: "simulate"
}
```

5.3.9 PUT /api/v2/hoverfly/mode

Changes the mode of the running instance of Hoverfly.

Example request body:

```
{
  mode: "simulate"
}
```

5.3.10 GET /api/v2/hoverfly/usage

Gets metrics information for the running instance of Hoverfly.

Example response body:

```
{
  "metrics": {
    "counters": {
      "capture": 0,
      "modify": 0,
      "simulate": 0,
      "synthesize": 0
    }
  }
}
```

5.3.11 GET /api/v2/hoverfly/version

Gets the version of Hoverfly.

Example response body:

```
{
  "version": "v0.10.1"
}
```

5.3.12 GET /api/v2/hoverfly/upstream-proxy

Gets the upstream proxy configured for Hoverfly.

Example response body:

```
{
  "upstream-proxy": "proxy.corp.big-it-company.org:8080"
}
```

5.4 Simulation schema

```

{
  "data": {
    "pairs": [
      {
        "response": {
          "status": 200,
          "body": "<h1>Matched on recording</h1>",
          "encodedBody": false,
          "headers": {
            "Content-Type": [
              "text/html; charset=utf-8"
            ]
          }
        },
        "request": {
          "requestType": "recording",
          "path": "/",
          "method": "GET",
          "destination": "myhost.io",
          "scheme": "https",
          "query": "",
          "body": "",
          "headers": {
            "Accept": [
              "text/html,application/xhtml+xml,application/xml;q=0.9,image/webp,*/*;
↪q=0.8"
            ],
            "Content-Type": [
              "text/plain; charset=utf-8"
            ],
            "User-Agent": [
              "Mozilla/5.0 (Macintosh; Intel Mac OS X 10_11_6) AppleWebKit/537.36_
↪(KHTML, like Gecko) Chrome/52.0.2743.116 Safari/537.36"
            ]
          }
        }
      },
      {
        "response": {
          "status": 200,
          "body": "<h1>Matched on template</h1>",
          "encodedBody": false,
          "headers": {
            "Content-Type": [
              "text/html; charset=utf-8"
            ]
          }
        },
        "request": {
          "requestType": "template",
          "path": "/template",
          "method": null,
          "destination": null,
          "scheme": null,
          "query": null,
          "body": null,

```

```
        "headers": null
      }
    ],
    "globalActions": {
      "delays": []
    }
  },
  "meta": {
    "schemaVersion": "v1",
    "hoverflyVersion": "v0.9.0",
    "timeExported": "2016-11-11T11:53:52Z"
  }
}
```