
Vaurien Documentation

Release 1.8

Mozilla

April 30, 2013

CONTENTS

Ever heard of the [Chaos Monkey](#)?



It's a project at Netflix to enhance the infrastructure tolerance. The Chaos Monkey will randomly shut down some servers or block some network connections, and the system is supposed to survive to these events. It's a way to verify the high availability and tolerance of the system.

Besides a redundant infrastructure, if you think about reliability at the level of your web applications there are many questions that often remain unanswered:

- What happens if the MYSQL server is restarted? Are your connectors able to survive this event and continue to work properly afterwards?
- Is your web application still working in degraded mode when Membase is down?
- Are you sending back the right 503s when postgresql times out ?

Of course you can – and should – try out all these scenarios on stage while your application is getting a realistic load.

But testing these scenarios while you are building your code is also a good practice, and having automated functional tests for this is preferable.

That's where **Vaurien** is useful.

Vaurien is basically a Chaos Monkey for your TCP connections. Vaurien acts as a proxy between your application and any backend.

You can use it in your functional tests or even on a real deployment through the command-line.

INSTALLING VAURIEN

You can install Vaurien directly from PyPI. The best way to do so is via *pip*:

```
$ pip install vaurien
```


DESIGN

Vaurien is a TCP proxy that simply reads data sent to it and pass it to a backend, and vice-versa.

It has built-in **protocols**: TCP, HTTP, Redis & Memcache. The **TCP** protocol is the default one and just sucks data on both sides and pass it along.

Having higher-level protocols is mandatory in some cases, when Vaurien needs to read a specific amount of data in the sockets, or when you need to be aware of the kind of response you're waiting for, and so on.

Vaurien also has **behaviors**. A behavior is a class that's going to be invoked everytime Vaurien proxies a request. That's how you can impact the behavior of the proxy. For instance, adding a delay or degrading the response can be implemented in a behavior.

Both **protocols** and **behaviors** are plugins, allowing you to extend Vaurien by adding new ones.

Last (but not least), Vaurien provides a couple of APIs you can use to change the behavior of the proxy live. That's handy when you are doing functional tests against your server: you can for instance start to add big delays and see how your web application reacts.

USING VAURIEN FROM THE COMMAND-LINE

Vaurien is a command-line tool.

Let's say you want to add a delay for 20% of the HTTP requests made on **google.com**:

```
$ vaurien --protocol http --proxy localhost:8000 --backend google.com:80 \  
    --behavior 20:delay
```

With this set up, Vaurien will stream all the traffic to **google.com** by using the *http* protocol, and will add delays 20% of the time.

You can find a description of all built-in protocols here: [Protocols](#).

You can pass options to the behavior using *-behavior-NAME-OPTION* options:

```
$ vaurien --protocol http --proxy localhost:8000 --backend google.com:80 \  
    --behavior 20:delay \  
    --behavior-delay-sleep 2
```

Passing all options through the command-line can be tedious, so you can also create an *ini* file for this:

```
[vaurien]  
backend = google.com:80  
proxy = localhost:8000  
protocol = http  
behavior = 20:delay
```

```
[behavior:delay]  
sleep = 2
```

You can find a description of all built-in behaviors here: [Behaviors](#).

You can also find some usage examples here: [Examples](#).

CONTROLLING VAURIEN LIVE

Vaurien provides an HTTP server with an API, which can be used to control the proxy and change its behavior on the fly.

To activate it, use the *-http* option:

```
$ vaurien --http
```

By default the server runs on **localhost:8080** but you can change it with the **-http-host** and **-http-port** options.

See [APIs](#) for a full list of APIs.

CONTROLLING VAURIEN FROM YOUR CODE

If you want to run and drive a Vaurien proxy from your code, the project provides a few helpers for this.

For example, if you want to write a test that uses a Vaurien proxy, you can write:

```
import unittest
from vaurien import Client, start_proxy, stop_proxy

class MyTest(unittest.TestCase):

    def setUp(self):
        self.proxy_pid = start_proxy(port=8080)

    def tearDown(self):
        stop_proxy(self.proxy_pid)

    def test_one(self):
        client = Client()
        options = {'inject': True}

        with client.with_behavior('error', **options):
            # do something...
            pass

        # we're back to normal here
```

In this test, the proxy is started and stopped before and after the test, and the Client class will let you drive its behavior.

Within the **with** block, the proxy will error out any call by using the *errors* behavior, so you can verify that your application is behaving as expected when it happens.

EXTENDING VAURIEN

Vaurien comes with a handful of useful *Behaviors* and *Protocols*, but you can create your own ones and plug them in a configuration file.

In fact, that's the best way to create realistic issues: imagine you have a very specific type of error on your LDAP server everytime your infrastructure is under heavy load. You can reproduce this issue in your behavior and make sure your web application behaves as it should.

Creating new behaviors and protocols is done by implementing classes with specific signatures.

For example if you want to create a “*super*” behavior, you just have to write a class with two special methods: **on_before_handle** and **on_after_handle**.

Once the class is ready, you can register it with **Behavior.register**:

```
from vaurien.behaviors import Behavior

class MySuperBehavior(object):

    name = 'super'
    options = {}

    def on_before_handle(self, protocol, source, dest, to_backend):
        # do something here
        return True

    def on_after_handle(self, protocol, source, dest, to_backend):
        # do something else
        return True

Behavior.register(MySuperBehavior)
```

You will find a full tutorial in *Extending Vaurien*.

CONTRIBUTE

The code repository & bug tracker are located at <https://github.com/mozilla-services/vaurien>
Don't hesitate to send us pull requests or open issues!

MORE DOCUMENTATION

And there is more! Have a look at the other sections of the documentation:

8.1 Behaviors

Vaurien provides a collections of behaviors, all of them are listed on this page. You can also write your own behaviors if you need. Have a look at [Extending Vaurien](#) to learn more.

8.1.1 blackout

Reads the packets that have been sent then hangs.

Acts like a `pdb.set_trace()` you'd forgot in your code ;)

8.1.2 delay

Adds a delay before or after the backend is called.

The delay can happen *after* or *before* the backend is called.

Options:

- **before**: If True adds before the backend is called. Otherwise after (bool, default: True)
- **sleep**: Delay in seconds (int, default: 1)

8.1.3 dummy

Transparent behavior. Nothing's done.

8.1.4 error

Reads the packets that have been sent then send back "errors".

Used in conjunction with the HTTP Protocol, it will randomly send back a 501, 502 or 503.

For other protocols, it returns random data.

The *inject* option can be used to inject data within valid data received from the backend. The Warmup option can be used to deactivate the random data injection for a number of calls. This is useful if you need the communication to settle in some specific protocols before the random data is injected.

The *inject* option is deactivated when the *http* protocol is used.

Options:

- **inject**: Inject errors inside valid data (bool, default: False)
- **warmup**: Number of calls before erroring out (int, default: 0)

8.1.5 hang

Reads the packets that have been sent then hangs.

Acts like a `pdb.set_trace()` you'd forgot in your code ;)

8.2 Protocols

Vaurien provides a collections of protocols, which are all listed on this page. You can also write your own protocols if you need. Have a look at [Extending Vaurien](#) to learn more.

8.2.1 http

HTTP protocol.

Options:

- **buffer**: Buffer size (int, default: 8124)
- **keep_alive**: Keep the connection alive (bool, default: False)
- **overwrite_host_header**: If True, the HTTP Host header will be rewritten with backend address. (bool, default: True)
- **reuse_socket**: If True, the socket is reused. (bool, default: False)

8.2.2 memcache

Memcache protocol.

Options:

- **buffer**: Buffer size (int, default: 8124)
- **keep_alive**: Keep the connection alive (bool, default: False)
- **reuse_socket**: If True, the socket is reused. (bool, default: False)

8.2.3 mysql

No documentation. Boooo!

Options:

- **buffer**: Buffer size (int, default: 8124)

- **reuse_socket**: If True, the socket is reused. (bool, default: False)

8.2.4 redis

Redis protocol.

Options:

- **buffer**: Buffer size (int, default: 8124)
- **keep_alive**: Keep the connection alive (bool, default: False)
- **reuse_socket**: If True, the socket is reused. (bool, default: False)

8.2.5 smtp

SMTP Protocol.

Options:

- **buffer**: Buffer size (int, default: 8124)
- **reuse_socket**: If True, the socket is reused. (bool, default: False)

8.2.6 tcp

TCP handler.

Options:

- **buffer**: Buffer size (int, default: 8124)
- **keep_alive**: Keep the connection alive (bool, default: False)
- **reuse_socket**: If True, the socket is reused. (bool, default: False)

8.3 APIs

You can control vaurien from its APIs. There is a REST API and a command-line API

8.3.1 The REST API

GET /behavior

Returns the current behavior in use, as a json object.

Example:

```
$ curl -XGET http://localhost:8080/behavior
{
  "behavior": "dummy"
}
```

PUT /behavior

Set the behavior. The behavior must be provided in a JSON object, in the body of the request, with a **name** key for the behavior name, and any option to pass to the behavior class.

Note: Don't forget to set the "application/json" Content-Type header when doing your calls.

Example:

```
$ curl -XPUT -d '{"sleep": 2, "name": "delay"}' http://localhost:8080/behavior \
-H "Content-Type: application/json"
{
  "status": "ok"
}
```

GET /behaviors

Returns a list of behaviors that are possible to use

Example:

```
$ curl -XGET http://localhost:8080/behaviors
{
  "behaviors": [
    "blackout",
    "delay",
    "dummy",
    "error",
    "hang"
  ]
}
```

If you want to control vaurien from the command-line, you can do so by using [vaurienclient](#). `vaurienctl -help` will provide you some help.

8.4 Extending Vaurien

You can extend Vaurien by writing new **protocols** or new **behaviors**.

8.4.1 Writing Protocols

Writing a new protocol is done by creating a class that inherits from the `vaurien.protocols.base.BaseProtocol` class.

The class needs to provide three elements:

- a **name** class attribute, the protocol will be known under that name.
- an optional **options** class attribute - a mapping containing options for the protocol. Each option value is composed of a description, a type and a default value. The mapping is wired in the command-line when you run `vaurien` - and is also used to generate the protocol documentation.
- a **_handle** method, that will be called everytime some data is ready to be read on the proxy socket or on the backend socket.

The `vaurien.protocols.base.BaseProtocol` class also provides a few helpers to work with the sockets:

- **_get_data**: a method to read data in a socket. Catches *EWOULDBLOCK* and *EAGAIN* errors and loops until they happen.

- *option*: a method to get the value of an option

Example:

```
class TCP(BaseProtocol):
    name = 'tcp'
    options = {'reuse_socket': ("If True, the socket is reused.",
                               bool, False),
              'buffer': ("Buffer size", int, 8124),
              'keep_alive': ("Keep the connection alive", bool, False)}

    def _handle(self, source, dest, to_backend):
        # default TCP behavior
        data = self._get_data(source)
        if data:
            dest.sendall(data)
            if not self.option('keep_alive'):
                data = ''
                while True:
                    data = self._get_data(dest)
                    if data == '':
                        break
                    source.sendall(data)

            if not self.option('reuse_socket'):
                dest.close()
                dest._closed = True
            return False
        return data != ''
```

Once the protocol class is ready, it can be registered via the Protocol class:

```
from vaurien.protocols import Protocol
Protocol.register(TCP)
```

8.4.2 Writing Behaviors

Creating new behaviors is very similar to creating protocols.

XXX

8.4.3 Using your protocols and behaviors

XXX

8.5 Examples

Proxying on an HTTP backend and sending back 50x errors 20% of the time:

```
$ vaurien --protocol http --proxy 0.0.0.0:8888 --backend blog.ziade.org:80 \
  --behavior 20:error
```

An SSL SMTP proxy with a 5% error rate and 10% delays:

```
$ vaurien --proxy 0.0.0.0:6565 --backend mail.example.com:465 \  
    --protocol smtp --behavior 5:error,10:delay
```

An SSL SMTP Proxy that starts to error out after 12 calls (so in the middle of the transaction):

```
$ vaurien --proxy 0.0.0.0:6565 --backend mail.example.com:465 \  
    --protocol smtp --behavior 100:error --behavior-error-warmup 12
```

Adding a 1 second delay on **every** call to a MySQL server:

```
$ vaurien --proxy 0.0.0.0:3307 --backend 0.0.0.0:3306 --stay-connected --behavior 100:delay \  
    --behavior-delay-sleep 1
```

A quick'n'dirty SSH tunnel from your box to another box:

```
$ vaurien --stay-connected --proxy 0.0.0.0:8887 --backend 192.168.1.276:22 \  
    --protocol-tcp-keep-alive
```