
gaffer Documentation

Release 0.5.2

Author

May 25, 2018

Contents

1	Gaffer	1
1.1	Features	1
1.2	Contents:	2
2	Indices and tables	33
	Python Module Index	35

Control, Watch and Launch your applications and jobs over HTTP.

Gaffer is a set of Python modules and tools to easily maintain and interact with your applications or jobs launched on different machines over HTTP and websockets.

It promotes distributed and decentralized topologies without single points of failure, enabling fault tolerance and high availability.

1.1 Features

- RESTful HTTP Api
- Websockets and **SOCKJS** support to interact with a gaffer node from any browser or **SOCKJS** client.
- Framework to manage and interact your applications and jobs on different machines
- Server and *Command Line* tools to manage and interact with your processes
- manages topology information. Clients query `gaffer_lookupd` to discover gaffer nodes for a specific job or application.
- Possibility to interact with **STDIO** and **PIPES** to interact with your applications and processes
- Subscribe to process statistics per process or process templates and get them in quasi RT.
- Procfile applications support (see *Gaffer*) but also JSON config support.
- Supervisor-like features.
- Fully evented. Use the libuv event loop using the [pyuv library](#)
- Flapping: handle cases where your processes crash too much
- Easily extensible: add your own endpoint, create your client, embed gaffer in your application, ...
- Compatible with python 2.7x, 3.x

Note: gaffer source code is hosted on [Github](#)

1.2 Contents:

1.2.1 Getting started

This tutorial exposes the usage of gaffer as a tool. For a general overview or how to integrate it in your application you should read the *overview page*.

Introduction

Gaffer allows you to launch OS processes and supervise them. 3 *command line* tools allows you to use it for now:

- *Gafferd* is the process supervisor and should be launched first before to use other tools.
- *Gaffer* is a Procfile application manager and allows you to load your Procfile applications in gafferd and watch their status.
- gafferctl is a more generic tool than gaffer and is more admin oriented. It allows you to setup any process templates and manage your processes. You can also use it to watch the activity in gafferd (process activity or general activity)

A process template is the way you describe the launch of an OS process, how many you want to launch on startup, how many time you want to restart it in case of failures (flapping). . . . A process template can be loaded using any tool or on *gafferd* startup using its configuration file.

Workflow

To use gaffer tools you need to:

1. First launch gafferd
2. Use either gaffer or gafferctl to manage your processes

Launch gafferd

For more informations of gafferd go on its *documentation page* .

To launch gafferd run the following command line:

```
$ gafferd -c /path/to/gaffer.ini
```

If you want to launch **custom plugins** with gafferd you can also set the path to them:

```
$ gafferd -c /path/to/gaffer.ini -p /path/to/plugin
```

Note: default plugin path is relative to the user launching gaffer and is set to *~/.gaffer/plugins*.

Note: To launch it in daemon mode use the `--daemon` option.

Then with the default configuration, you can check if gafferd is alive

The configuration file

The configuration file can be used to set the global configuration of gafferd, setup some processes and webhooks.

Note: Since the configuration is passed to the plugin you can also use this configuration file to setup your plugins.

Here is a simple example of a config to launch the dummy process from the example folder:

```
[process:dummy]
cmd = ./dummy.py
numprocesses = 1
redirect_output = stdout, stderr
```

Note: Process can be grouped. You can then start and stop all processes of a group and see if a process is member of a group using the HTTP api. (sadly this is not yet possible to do it using the command line).

For example if you want dummy be part of the group test, then `[process:dummy]` will become `[process:test:dummy]`. A process template as you can see can only be part of one group.

Groups are useful when you want to manage a configuration for one application or processes / users.

Each process section should be prefixed by `process:`. Possible parameters are:

- **cmd:** the full command line to launch. eg. `./dummy.p`
- **args:** arguments to pass as a string. eg. `-some value --option=a`
- **cwd:** path to working directory
- **uid:** user name or id used to execute the process
- **gid:** group name or id used to execute the process
- **detach:** if you want to completely detach the process from gafferd (gaffer will still continue to supervise it)
- **shell:** The process is executed in a shell (unix only)
- **flapping:** flapping rule. eg. `2, 1., 7., 5` which means `attempts=2, window=1., retry_in=7., max_retry=5`
- **redirect_input:** to allow you to interact with stdin
- **redirect_output:** to watch both stdout & stderr. output names can be whatever you want. For example you. eg. `redirect_output = mystdout, mystderr` stdout will be labelled `mystdout` in this case.
- **graceful_timeout:** time to wait before definitely kill a process. By default 30s. When killing a process, gaffer is first sending a `SIGTERM` signal then after a graceful timeout if the process hasn't stopped by itself send a `SIGKILL` signal. It allows you to handle the way your process will stop.
- **os_env:** true or false, to pass all operating system variables to the process environment.
- **priority:** Integer. Allows you to fix the order in which gafferd will start the processes. 0 is the highest priority. By default all processes have the same order.

Sometimes you also want to pass a custom environment to your process. This is done by creating a special configuration section named `env:processname`. Each environment sections are prefixed by `env:.` For example to pass a special `PORT` environment variable to `dummy`:

```
[env:dummy]
port = 80
```

All environment variables key are passed in uppercase to the process environment.

Manage your Procfile applications

The **gaffer** command line tool is an interface to the *gaffer HTTP api* and include support for loading/unloading Procfile applications, scaling them up and down,

It can also be used as a manager for Procfile-based applications similar to `foreman` but using the *gaffer framework*. It is running your application directly using a Procfile or export it to a `gaffer`d configuration file or simply to a JSON file that you could send to `gaffer`d using the *HTTP api*.

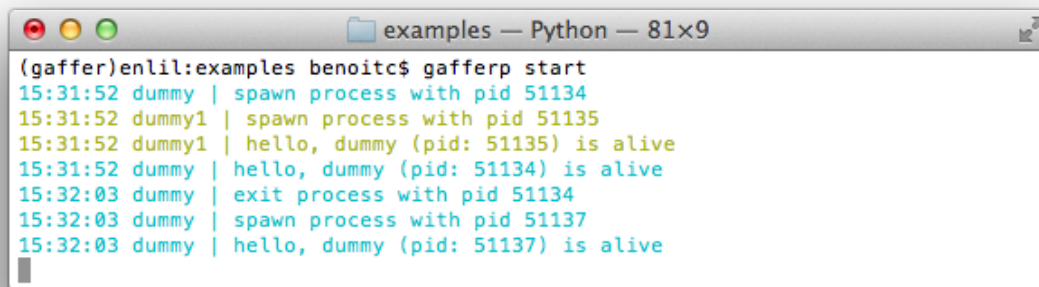
Example of use

For example using the following **Procfile**:

```
dummy: python -u dummy_basic.py
dummy1: python -u dummy_basic.py
```

You can launch all the programs in this procfile using the following command line:

```
$ gaffer start
```

A terminal window titled "examples — Python — 81x9" showing the output of the `gaffer start` command. The output shows two processes, `dummy` and `dummy1`, being spawned and then exiting. The `dummy` process prints "hello, dummy (pid: 51134) is alive" and then "exit process with pid 51134". The `dummy1` process prints "hello, dummy (pid: 51135) is alive".

```
(gaffer)enil:examples benoitc$ gaffer start
15:31:52 dummy | spawn process with pid 51134
15:31:52 dummy1 | spawn process with pid 51135
15:31:52 dummy1 | hello, dummy (pid: 51135) is alive
15:31:52 dummy | hello, dummy (pid: 51134) is alive
15:32:03 dummy | exit process with pid 51134
15:32:03 dummy | spawn process with pid 51137
15:32:03 dummy | hello, dummy (pid: 51137) is alive
```

Or load them on a `gaffer` node:

```
$ gaffer load
```

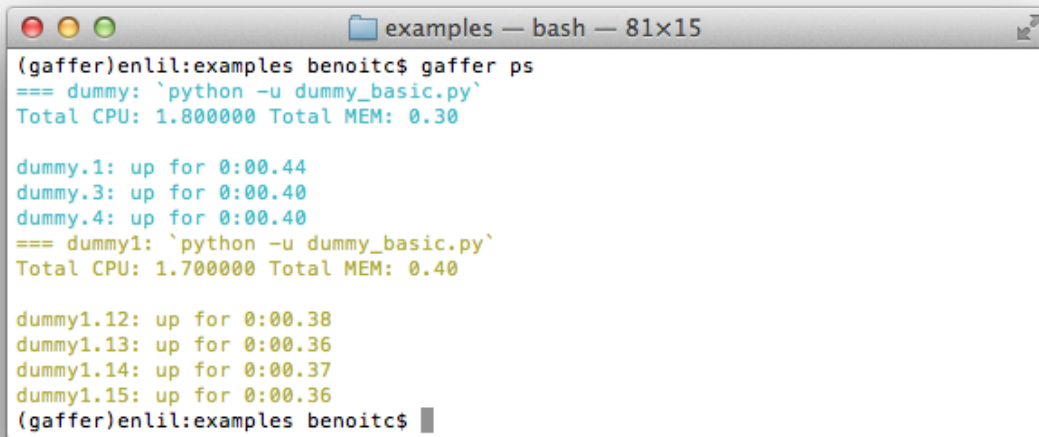
All processes in the Procfile will be then loaded to `gaffer`d and started.

If you want to start a process with a specific environment file you can create a `.env` in he application folder (or use the command line option to tell to `gaffer` which one to use). Each environment variables are passed by lines. Ex:


```
PORT=80
```

and then scale them up and down:

```
$ gaffer scale dummy=3 dummy1+2
Scaling dummy processes... done, now running 3
Scaling dummy1 processes... done, now running 3
```



```
examples — bash — 81x15
(gaffer)enil:examples benoitc$ gaffer ps
=== dummy: `python -u dummy_basic.py`
Total CPU: 1.800000 Total MEM: 0.30

dummy.1: up for 0:00.44
dummy.3: up for 0:00.40
dummy.4: up for 0:00.40
=== dummy1: `python -u dummy_basic.py`
Total CPU: 1.700000 Total MEM: 0.40

dummy1.12: up for 0:00.38
dummy1.13: up for 0:00.36
dummy1.14: up for 0:00.37
dummy1.15: up for 0:00.36
(gaffer)enil:examples benoitc$
```

have a look on the [Gaffer](#) page for more informations about the commands.

Control gafferd with gafferctl

gafferctl can be used to run any command listed below. For example, you can get a list of all processes templates:

```
$ gafferctl processes
```

You can simply add a process using the `load` command:

```
$ gafferctl load_process ../test.json
$ cat ../test.json | gafferctl load_process -
$ gafferctl load_process - < ../test.json
```

test.json can be:

```
{
  "name": "somename",
  "cmd": "cmd to execute",
  "args": [],
  "env": {},
  "uid": int or "",
  "gid": int or "",
  "cwd": "working dir",
  "detach": False,
  "shell": False,
```

(continues on next page)

(continued from previous page)

```
"os_env": False,  
"numprocesses": 1  
}
```

You can also add a process using the `add` command:

```
gafferctl add name inc
```

where `name` is the name of the process to create and `inc` the number of new OS processes to start.

To start a process run the following command:

```
$ gafferctl start name
```

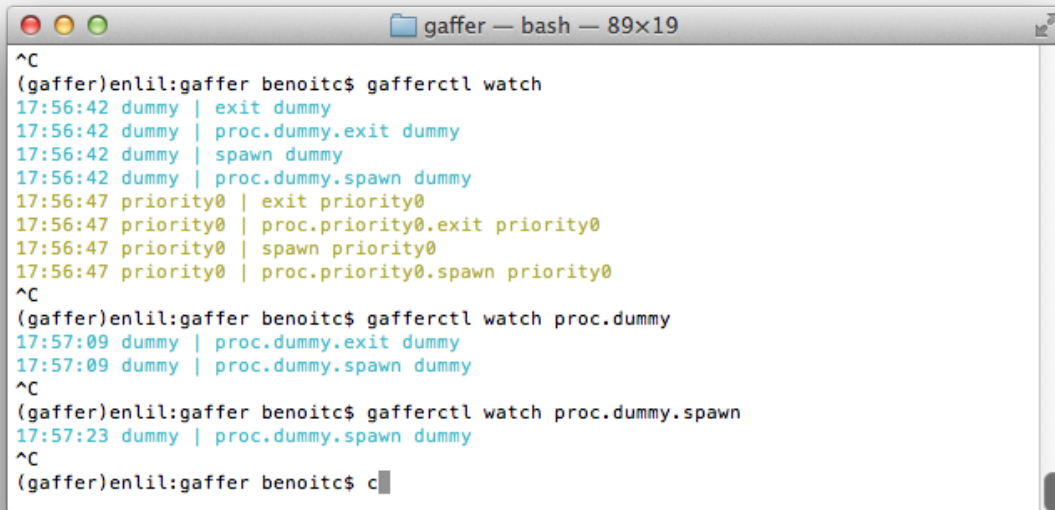
And stop it using the `stop` command.

To scale up a process use the `add` command. For example to increase the number of processes from 3:

```
$ gafferctl add name 3
```

To decrease the number of processes use the command `stop`.

The command `watch` allows you to watch changes in a local or remote gaffer node.



Demo

1.2.2 Overview

Gaffer is a set of Python modules and tools to easily maintain and interact with your processes.

Depending on your needs you can simply use the gaffer tools (eventually extend them) or embed the gaffer possibilities in your own apps.

Design

Gaffer is internally based on an event loop using the `libuv` from Joyent via the `pyuv` binding

All gaffer events are added to the loop and processes asynchronously which make it pretty performant for handling & controlling multiple processes.

At the lowest level you will find the manager. A manager is responsible for maintaining live processes and managing actions on them:

- start/stop processes
- increase/decrease the number of instances of each process, via process templates
- add/remove process templates to manage

A process template describes the way a process will be launched and how many OS processes you want to handle for this template. This number can be changed dynamically.

Current properties for each template:

- **name**: name of the process (eg, 'django-server')
- **cmd**: program command, (eg. '/var/www/CMS/manage.py')
- **args**: the arguments for the command to run. Can be a list or a string. If **args** is a string, it's split using `shlex.split()`. Defaults to None.
- **env**: a mapping containing the environment variables the command will run with. Optional
- **uid**: int or str, user id
- **gid**: int or str, user group id,
- **cwd**: working dir
- **detach**: the process is launched but won't be monitored and won't exit when the manager is stopped.
- **shell**: boolean, run the script in a shell. (UNIX only),
- **os_env**: boolean, pass the os environment to the program
- **numprocesses**: int the number of OS processes to launch for this description
- **flapping**: a `FlappingInfo` instance or, if flapping detection should be used. flapping parameters are:
 - **attempts**: maximum number of attempts before we stop the process and set it to retry later
 - **window**: period in which we are testing the number of retry
 - **retry_in**: seconds, the time after we restart the process and try to spawn them
 - **max_retry**: maximum number of retry before we give up and stop the process.
- **redirect_output**: list of io to redirect (max 2) this is a list of custom labels to use for the redirection. Ex: ["a", "b"] will redirect stdout & stderr and stdout events will be labeled "a"
- **redirect_input**: Boolean (False is the default). Set it if you want to be able to write to stdin.

The manager is also responsible of starting and stopping gaffer applications that you add to the manager to react to different events. A application can fetch information from the manager and interact with it.

Running an application is done like this:

```
# initialize the controller with the default loop loop = pyuv.Loop.default_loop() manager = Manager(loop=loop)
# start the controller manager.start(applications=[HttpHandler()])
```

```
.... # do smth  
manager.stop() # stop the controlller manager.run() # run the event loop
```

The `HttpHandler` application allows you to interact with `gaffer` via HTTP. It is used by the `gaffer` server which is able for now to load process templates via an ini files and maintain an HTTP endpoint which can be configured to be accessible on multiples interfaces and transports (tcp & unix sockets) .

Note: Only applications instances are used by the manager. It allows you to initialize them with your own settings.

Building your own application is easy, basically an application has the following structure:

```
class MyApplication(object):  
  
    def __init__(self):  
        # initialisation  
  
    def start(self, loop, manager):  
        # this method is called by the manager to start the controller  
  
    def stop(self):  
        # method called when the manager stops  
  
    def restart(self):  
        # method called when the manager restarts
```

You can use this structure for anything you want, (even add an app to the loop).

To help you in your work a *pyuv implementation* of tornado is integrated and the powerful *events* module allows you to manage PUB/SUB events (or anything evented) inside your app. An `EventEmitter` is a threadsafe class to manage subscriber and publisher of events. It is internally used to broadcast process and manager events.

Watch stats

Stats of a process can be monitored continuously (there is a refresh interval of 0.1s to fetch CPU informations) using the following method:

```
manager.monitor(<nameorid>, <listener>)
```

Where `<nameorid>` is the name of the process template. In this case the statistics of all the the OS processes using this template will be emitted. Stats events are collected in the listener callback.

Callback signature: `callback(evtype, msg)`.

evtype is always "STATS" here and **msg** is a dict:

```
{  
    "mem_info1": int,  
    "mem_info2": int,  
    "cpu": int,  
    "mem": int,  
    "ctime": int,  
    "pid": int,  
    "username": str,  
    "nicce": int,  
    "cmdline": str,
```

(continues on next page)

(continued from previous page)

```
"children": [{ stat dict, ... }]
}
```

To unmonitor the process in your app run:

```
manager.unmonitor(<nameorid>, <listener>)
```

Note: Internally a monitor subscribe you to an EventEmitter. A timer is running until there are subscribers to the process stats events.

Of course you can directly monitor a process using the internal pid:

```
process = manager.running[pid]
process.monitor(<listener>)

...

process.unmonitor(<listener>)
```

IO Events

Subscribe to stdout/stderr process streams

You can subscribe to stdout/stderr process streams and even write to stdin if you want.

To be able to receive the stdout/stderr streams in your application, you need to create a process with the *redirect_output* setting:

```
manager.add_process("nameofprocestemplate", cmd,
    redirect_output["stdout", "stderr"])
```

Note: Name of outputs can be anything, only the order counts. So if you want to name *stdout* as *a* just replace *stdout* by *a* in the declaration.

If you don't want to receive *stderr*, just omit it from the list.

If you want to redirect *stderr* to *stdout* just use the same name.

Then for example, to monitor the stdout output do:

```
process.monitor_io("stdout", somecallback)
```

Callback signature: `callback(evtype, msg)`.

And to unmonitor:

```
process.unmonitor_io("stdout", somecallback)
```

Note: To subscribe to all process streams replace the stream name by `'.'`.

Write to STDIN

Writing to stdin is pretty easy. Just do:

```
process.write("somedata")
```

or to send multiple lines:

```
process.writelines(["line", "line"])
```

You can write lines from multiple publishers and multiple publishers can write at the same time. This method is threadsafe.

HTTP API

See the *HTTP api description* for more informations.

Tools

Gaffer proposes different tools (and more will come soon) to manage your processes without having to code. It can be used like *supervisor*, *god*, *runit* or other tools around. Speaking of *runit* a similar controlling will be available in 0.2 .

See the *Command Line* documentation for more informations.

1.2.3 CHANGES

2013/09/29 - version 0.5.1

- new intermediary release. See <https://github.com/benoitc/gaffer/compare/0.4.4...0.5.0> for a detailed changelog.

2012/12/20 - version 0.4.4

- improve Events dispatching
- add support for multiple channel in a process
- add ping handler for monitoring
- some fixes in the http api
- fix `stop_processes` function

2012/11/02 - version 0.4.3

- process os environment now inherits from the gafferd environment
- fix autorestart feature: now handled asynchronously which allows us to still handle “stop command when a process fails”

2012/11/01 - version 0.4.2

- fix `os_env` option

2012/10/29 - version 0.4.0

- add environment variables support in the gaffer setting file.
- add a plugin system to easily extend *Gaffer* using HTML sites or gaffer applications in Python

2012/10/18 - version 0.3.1

- add environment variables substitution in the process command line and arguments.

2012/10/18 - version 0.3.0

- add the *Gaffer* command line tool: load, unload your procfile applications to gaffer, scale them up and down. Or just use it as a procfile manager just like *foreman* .
- add gafferctl commands/watch command to watch a node activity remotely.
- add priority feature: now processes can be launch in order
- add the possibility to manipulate *groups of processes*
- add the possibility to set the default endpoint in gaffer from the command line
- add `-v` and `--vv` options to gaffer to have a verbose output.
- add an eventsource client in the framework to manipulate gaffer streams.
- add `Manager.start_processes` method. Start all processes.
- add `console_output` application to the framework
- add new global *Gaffer events* to the manager: `spawn`, `reap`, `stop_pid`, `exit`.
- fix shutdown
- fix heartbeat

2012/10/15 - version 0.2.0

- add *Webhooks*: post to an url when a gaffer event is triggered
- add graceful shutdown. kill processes after a graceful time
- add `commands/load_process` command
- code refactoring: make the code simpler

2012/10/12 - version 0.1.0

Initial release

1.2.4 Command Line

Gaffer is a *process management framework* but also a set of command lines tools allowing you to manage on your machine or a cluster. All the command line tools are obviously using the framework.

Gaffer is an interface to the *gaffer HTTP api* and include support for loading/unloadin apps, scaling them up and down, It can also be used as a manager for Procfile-based applications similar to *foreman* but using the *gaffer*

framework. It is running your application directly using a Procfile or export it to a gaffer configuration file or simply to a JSON file that you could send to gaffer using the *HTTP api*.

Gaffer is a server able to launch and manage processes. It can be controlled via the *HTTP api*. It is controlled by gafferctl and can be used to handle many processes.

Gaffer

The **gaffer** command line tool is an interface to the *gaffer HTTP api* and include support for loading/unloading Procfile applications, scaling them up and down,

It can also be used as a manager for Procfile-based applications similar to foreman but using the *gaffer framework*. It is running your application directly using a Procfile or export it to a gaffer configuration file or simply to a JSON file that you could send to gaffer using the *HTTP api*.

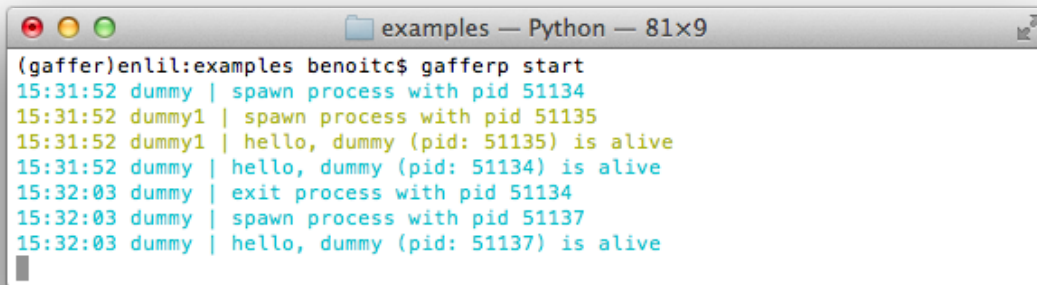
Example of use

For example using the following **Procfile**:

```
dummy: python -u dummy_basic.py
dummy1: python -u dummy_basic.py
```

You can launch all the programs in this procfile using the following command line:

```
$ gaffer start
```

A terminal window titled "examples — Python — 81x9" showing the output of the command "gaffer start". The output shows two processes, "dummy" and "dummy1", being spawned and then exiting and being spawned again. The output is as follows:

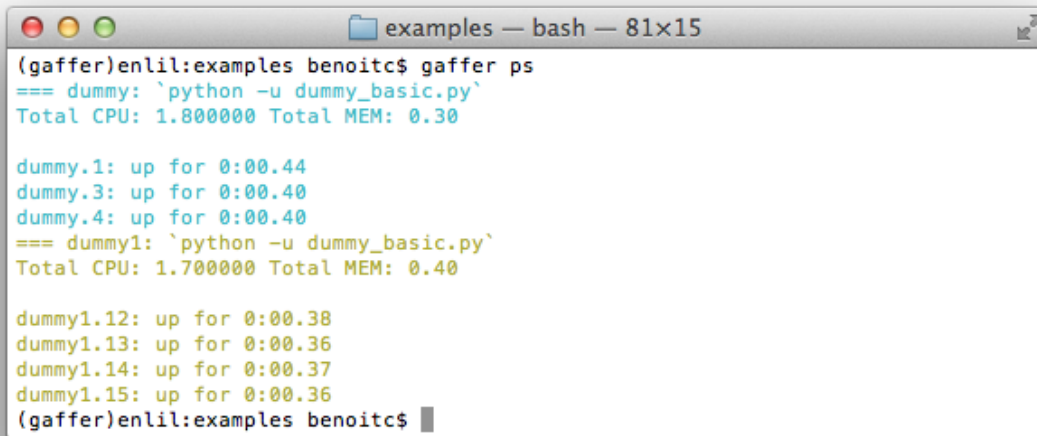
```
(gaffer)enlil:examples benoitc$ gaffer start
15:31:52 dummy | spawn process with pid 51134
15:31:52 dummy1 | spawn process with pid 51135
15:31:52 dummy1 | hello, dummy (pid: 51135) is alive
15:31:52 dummy | hello, dummy (pid: 51134) is alive
15:32:03 dummy | exit process with pid 51134
15:32:03 dummy | spawn process with pid 51137
15:32:03 dummy | hello, dummy (pid: 51137) is alive
```

Or load them on a gaffer node:

```
$ gaffer load
```

and then scale them up and down:

```
$ gaffer scale dummy=3 dummy1+2
Scaling dummy processes... done, now running 3
Scaling dummy1 processes... done, now running 3
```

```

(gaffer)enlil:examples benoitc$ gaffer ps
=== dummy: `python -u dummy_basic.py`
Total CPU: 1.800000 Total MEM: 0.30

dummy.1: up for 0:00.44
dummy.3: up for 0:00.40
dummy.4: up for 0:00.40
=== dummy1: `python -u dummy_basic.py`
Total CPU: 1.700000 Total MEM: 0.40

dummy1.12: up for 0:00.38
dummy1.13: up for 0:00.36
dummy1.14: up for 0:00.37
dummy1.15: up for 0:00.36
(gaffer)enlil:examples benoitc$

```

OPTIONS

-h **-help** show this help message and exit **-version** show version and exit **-f profile,-profile profile** Specify an alternate Procfile to load **-d root,-directory root** Specify an alternate application root

This defaults to the directory containing the Procfile [default: .]

-e **k=v,-env k=v** Specify one or more .env files to load **-endpoint endpoint** gaffer node URL to connect [default: http://127.0.0.1:5000]

SUBCOMMANDS

export [-c concurrency]-concurrency concurrency] [-format=format] [-out=filename] [<name>]

Export a Procfile

This command export a Procfile to a gaffer process settings format. It can be either a JSON that you could send to gaffer via the JSON API or an ini file that can be included to the gaffer configuration.

<format> ini or json **-out=filename** path of filename where the export will be saved

load [-c concurrency]-concurrency concurrency] [-nostart] [<name>] Load a Procfile application to gaffer

<name> is the name of the application recorded in gaffer. By default it will be the name of your project folder. You can use `.` to specify the current folder.

ps [<appname>] List your processes informations

<appname> he name of the application (session) of process recoreded in gaffer. By default it will be the name of your project folder. You can use `.` to specify the current folder.

run [-c] [<args>]... Run one-off commands using the same environment as your defined processes

-c concurrency Specify the number of each process type to run. The value passed in should be in the format process=num,process=num

--concurrency concurrency same as the -c option.

scale [<appname>] [process=value]... Scaling your process

Procfile applications can scale up or down instantly from the command line or API.

Scaling a process in an application is done using the scale command:

```
$ gaffer scale dummy=3
Scaling dummy processes... done, now running 3
```

Or both at once:

```
$ gaffer scale dummy=3 dummy1+2
Scaling dummy processes... done, now running 3
Scaling dummy1 processes... done, now running 3
```

start [-c concurrency|--concurrency concurrency]

Start a process type or all process types from the Procfile.

-c concurrency Specify the number of each process type to run. The value passed in should be in the format process=num,process=num

--concurrency concurrency same as the -c option.

unload [<name>] Unload a Procfile application from a gaffer node

Gafferd

Gafferd is a server able to launch and manage processes. It can be controlled via the [HTTP api](#).

Usage

```
$ gafferd -h
usage: gafferd [-h] [-c CONFIG_FILE] [-p PLUGINS_DIR] [-v] [-vv] [--daemon]
              [--pidfile PIDFILE] [--bind BIND] [--certfile CERTFILE]
              [--keyfile KEYFILE] [--backlog BACKLOG]
              [config]

Run some watchers.

positional arguments:
  config                configuration file

optional arguments:
  -h, --help            show this help message and exit
  -c CONFIG_FILE, --config CONFIG_FILE
                        configuration file
  -p PLUGINS_DIR, --plugins-dir PLUGINS_DIR
                        default plugin dir
  -v                    verbose mode
  -vv                   like verbose mode but output stream too
  --daemon              Start gaffer in the background
```

(continues on next page)

(continued from previous page)

```

--pidfile PIDFILE
--bind BIND          default HTTP binding
--certfile CERTFILE  SSL certificate file for the default binding
--keyfile KEYFILE    SSL key file for the default binding
--backlog BACKLOG    default backlog

```

Config file example

```

[gaffer]
http_endpoints = public

[endpoint:public]
bind = 127.0.0.1:5000
;certfile=
;keyfile=

[webhooks]
;create = http://some/url
;proc.dummy.spawn = http://some/otherurl

[process:dummy]
cmd = ./dummy.py
;cwd = .
;uid =
;gid =
;detach = false
;shell = false
; flapping format: attempts=2, window=1., retry_in=7., max_retry=5
;flapping = 2, 1., 7., 5
numprocesses = 1
redirect_output = stdout, stderr
; redirect_input = true
; graceful_timeout = 30

[process:echo]
cmd = ./echo.py
numprocesses = 1
redirect_output = stdout, stderr
redirect_input = true

```

Plugins

Plugins are a way to enhance the basic gaffer functionality in a custom manner. Plugins allows you to load any gaffer application and site plugins. You can for example use the plugin system to add a simple UI to administrate gaffer using the HTTP interface.

A plugin has the following structure:

```

/pluginname
  _site/
  plugin/
    __init__.py

```

(continues on next page)

(continued from previous page)

```
...
***.py
```

A plugin can be discovered by adding one or more modules that expose a class inheriting from `gaffer.Plugin`. Every plugin file should have a `__all__` attribute containing the implemented plugin class. Ex:

```
from gaffer import Plugin

__all__ = ['DummyPlugin']

from .app import DummyApp

class DummyPlugin(Plugin):
    name = "dummy"
    version = "1.0"
    description = "test"

    def app(self, cfg):
        return DummyApp()
```

The dummy app here only prints some info when started or stopped:

```
class DummyApp(object):

    def start(self, loop, manager):
        print("start dummy app")

    def stop(self):
        print("stop dummy")

    def restart(self):
        print("restart dummy")
```

See the *Overview* for more info. You can try it in the example folder:

```
$ cd examples
$ gafferd -c gaffer.ini -p plugins/
```

Install plugins

Installing plugins can be done by placing the plugin in the plugin folder. The plugin folder is either set in the setting file using the `plugin_dir` in the `gaffer` section or using the `-p` option of the command line.

The default plugin dir is set to `~/.gafferd/plugins`.

Site plugins

Plugins can have “sites” in them, any plugin that exists under the `plugins` directory with a `_site` directory, its content will be statically served when hitting `/_plugin/[plugin_name]/url`. Those can be added even after the process has started.

Installed plugins that do not contain any Python related content, will automatically be detected as site plugins, and their content will be moved under `_site`.

Mandatory Plugins

If you rely on some plugins, you can define mandatory plugins using the `mandatory` attribute of a the plugin class, for example, here is a sample config:

```
class DummyPlugin(Plugin):
    ...
    mandatory = ['somedep']
```

gaffer_lookupd

Gaffer lookupd server to register gaffer nodes and access to their address.

```
$ gaffer_lookupd -h usage: gaffer_lookupd [--version] [-v] [--daemon] [--pidfile=PIDFILE]
```

```
    [--bind=ADDRESS] [--backlog=BACKLOG] [--certfile=CERTFILE] [--keyfile=KEYFILE] [--cacert=CACERT]
```

Options

```
-h --help show this help message and exit
--version show version and exit
-v verbose mode
--daemon Start gaffer in daemon mode
--pidfile=PIDFILE --backlog=BACKLOG default backlog [default: 128]
--bind=ADDRESS default HTTP binding [default: 0.0.0.0:5010]
--certfile=CERTFILE SSL certificate file
--keyfile=KEYFILE SSL key file
--cacert=CACERT SSL CA certificate
```

1.2.5 HTTP api

an http API provided by the `gaffer.http_handler.HttpHandler`` gaffer application can be used to control gaffer via HTTP. To embed it in your app just initialize your manager with it:

```
manager = Manager(apps=[HttpHandler()])
```

The `HttpHandler` can be configured to accept multiple endpoints and can be extended with new HTTP handlers. Internally we are using Tornado so you can either extend it with rules using pure totornado handlers or wsgi apps.

Request Format and Responses

Gaffer supports **GET, POST, PUT, DELETE, OPTIONS** HTTP verbs.

All messages (except some streams) are JSON encoded. All messages sent to gaffers should be json encoded.

Gaffer supports cross-origin resource sharing (aka CORS).

HTTP endpoints

Main http endpoints are described in the description of the `gafferctl` commands in `gafferctl`:

Gafferctl is using extensively this HTTP api.

Output streams

The output streams can be fetched by doing:

```
GET /streams/<pid>/<nameoffeed>
```

It accepts the following query parameters:

- **feed** : continuous, longpoll, eventsource
- **heartbeat**: true or seconds, send an empty line each sec (if true 60)

ex:

```
$ curl localhost:5000/streams/1/stderr?feed=continuous
STDERR 12
STDERR 13
STDERR 14
STDERR 15
STDERR 16
STDERR 17
STDERR 18
STDERR 19
STDERR 20
STDERR 21
STDERR 22
STDERR 23
STDERR 24
STDERR 25
STDERR 26
STDERR 27
STDERR 28
STDERR 29
STDERR 30
STDERR 31

$ curl localhost:5000/streams/1/stderr?feed=longpoll
STDERR 215

$ curl localhost:5000/streams/1/stderr?feed=eventsource
event: stderr
data: STDERR 20

event: stderr
data: STDERR 21

event: stderr
data: STDERR 22

$ curl localhost:5000/streams/1/stdout?feed=longpoll
STDOUTi 14
```

Write to STDIN

It is now possible to write to stdin via the HTTP api by sending:

```
POST to /streams/<pid>/ttin
```

Where <pid> is an internal process ide that you can retrieve by calling *GET /processses/<name>/_pids*

ex:

```
$ curl -XPOST -d '$ECHO\n' localhost:5000/streams/2/stdin
{"ok": true}

$ curl localhost:5000/streams/2/stdout?feed=longpoll
ECHO
```

Websocket stream for STDIN/STDOUT

It is now possible to get stdin/stdout via a websocket. Writing to `ws://HOST:PORT/wstreams/<pid>` will send the data to stdin any information written on stdout will be then sent back to the websocket.

See the echo client/server example in the example folder:

```
$ python echo_client.py
Sent
Receiving...
Received 'ECHO

'

(test)enlil:examples benoitc$ python echo_client.py
Sent
Receiving...
Received 'ECHO
```

Note: unfortunately the echo_client script can only be launched with python 2.7 :/

Note: to redirect stderr to stdout just use the same name when you setting the redirect_output property on process creation.

1.2.6 Webhooks

Webhooks allow to register an url to a specific event (or alls) and the event will be posted on this URL. Each events can trigger a post on a given url.

for example to listen all create events on <http://echohttp.com/echo> you can add this line in the webhooks sections of the gaffer setting file:

```
[webhooks]
create = http://echohttp.com/echo you
```

Or programatically:

```
from gaffer.manager import Manager
from gaffer.webhooks import WebHooks
hooks = [("create", "http://echohttp.com/echo you ")
webhooks = WebHooks(hooks=hooks)

manager = Manager()
manager.start(apps=[webhooks])
```

This gaffer application is started like other applications in the manager. All *Gaffer events* are supported.

The webhooks Module

```
class gaffer.webhooks.WebHooks (hooks=[])
    Bases: object

    webhook app

    close ()

    decref ()

    incref ()

    maybe_start_monitor ()

    maybe_stop_monitor ()

    refcount

    register_hook (event, url)
        associate an url to an event

    restart ()

    start (loop, manager)
        start the webhook app

    stop ()
        stop the webhook app, stop monitoring to events

    unregister_hook (event, url)
        unregister an url for this event
```

1.2.7 Core gaffer framework

manager Module

The manager module is a core component of gaffer. A Manager is responsible of maintaining processes and allows you to interact with them.

Classes

```
class gaffer.manager.Manager (loop=None)
    Bases: object
```

Manager - maintain process alive

A manager is responsible of maintaining process alive and manage actions on them:

- increase/decrease the number of processes / process template
- start/stop processes
- add/remove process templates to manage

The design is pretty simple. The manager is running on the default event loop and listening on events. Events are sent when a process exit or from any method call. The control of a manager can be extended by adding apps on startup. For example gaffer provides an application allowing you to control processes via HTTP.

Running an application is done like this:


```
# initialize the application with the default loop
loop = pyuv.Loop.default_loop()
m = Manager(loop=loop)

# start the application
m.start(apps=[HttpHandler])

.... # do smth

m.stop() # stop the controller
m.run() # run the event loop
```

Note: The loop can be omitted if the first thing you do is launching a manager. The run function is here for convenience. You can of course just run `loop.run()` instead

Warning: The manager should be stopped the last one to prevent any lock in your application.

active

commit (*name*, *graceful_timeout=0*, *env=None*)

Like “scale(1) but the process won’t be kept alived at the end. It is also not handled uring scaling or reaping.

get (*name*)

get a job config

get_process (*pid*)

get an OS process by ID. A process is a `gaffer.Process` instance attached to a process state that you can use.

get_process_id ()

generate a process id

info (*name*)

get job’ infos

jobs (*sessionid=None*)

jobs_walk (*callback*, *sessionid=None*)

kill (*pid*, *sig*)

send a signal to a process

killall (*name*, *sig*)

send a signal to all processes of a job

list (*name=None*)

load (*config*, *sessionid=None*, *env=None*, *start=True*)

load a process config object.

Args:

- **config**: a `process.ProcessConfig` instance
- **sessionid**: Some processes only make sense in certain contexts. this flag instructs gaffer to maintain this process in the sessionid context. A context can be for example an application. If no session is specified the config will be attached to the `default` session.

- **env**: dict, None by default, if specified the config env variable will be updated with the env values.

manage (*name*)

monitor (*listener, name*)

get stats changes on a process template or id

pids (*name=None*)

reload (*name, sessionid=None*)

reload a process config. The number of processes is reseted to the one in settings and all current processes are killed

restart (*callback=None*)

restart all processes in the manager. This function is threadsafe

run ()

Convenience function to use in place of *loop.run()* If the manager is not started it raises a *RuntimeError*.

Note: if you want to use separately the default loop for this thread then just use the start function and run the loop somewhere else.

scale (*name, n*)

Scale the number of processes in for a job. By using this function you can increase, decrease or set the number of processes in a template. Change is handled once the event loop is idling

n can be a positive or negative integer. It can also be a string containing the opetation to do. For example:

```
m.scale("sometemplate", 1) # increase of 1
m.scale("sometemplate", -1) # decrease of 1
m.scale("sometemplate", "+1") # increase of 1
m.scale("sometemplate", "-1") # decrease of 1
m.scale("sometemplate", "=1") # set the number of processess to 1
```

send (*pid, lines, stream=None*)

send some data to the process

sessions

start (*apps=[]*)

start the manager.

start_job (*name*)

Start a job from which the config have been previously loaded

stats (*name*)

return job stats

stop (*callback=None*)

stop the manager. This function is threadsafe

stop_job (*name*)

stop a jon. All processes of this job are stopped and won't be restarted by the manager

stop_process (*pid*)

stop a process

stopall (*name*)

stop all processes of a job. Processes are just exiting and will be restarted by the manager.

subscribe (*topic*)

unload (*name_or_process, sessionid=None*)

unload a process config.

unmonitor (*listener, name*)
 get stats changes on a process template or id

unsubscribe (*topic, channel*)

update (*config, sessionid=None, env=None, start=False*)
 update a process config. All processes are killed

wakeup ()

walk (*callback, name=None*)

process Module

The process module wrap a process and IO redirection

```
class gaffer.process.Process (loop, pid, name, cmd, args=None, env=None, uid=None,
                                gid=None, cwd=None, detach=False, shell=False, redi-
                                rect_output=[], redirect_input=False, custom_streams=[],
                                custom_channels=[], on_exit_cb=None)
```

Bases: object

class wrapping a process

Args:

- **loop**: main application loop (a pyuv Loop instance)
- **name**: name of the process
- **cmd**: program command, string)
- **args**: the arguments for the command to run. Can be a list or a string. If **args** is a string, it's splitted using `shlex.split()`. Defaults to None.
- **env**: a mapping containing the environment variables the command will run with. Optional
- **uid**: int or str, user id
- **gid**: int or st, user group id,
- **cwd**: working dir
- **detach**: the process is launched but won't be monitored and won't exit when the manager is stopped.
- **shell**: boolean, run the script in a shell. (UNIX only)
- **redirect_output**: list of io to redict (max 2) this is a list of custom labels to use for the redirection. Ex: ["a", "b"] will redirect stdoutt & stderr and stdout events will be labeled "a"
- **redirect_input**: Boolean (False is the default). Set it if you want to be able to write to stdin.
- **custom_streams**: list of additional streams that should be created and passed to process. This is a list of streams labels. They become available through `streams` attribute.
- **custom_channels**: list of additional channels that should be passed to process.

active

close ()

closed

info

return the process info. If the process is monitored it return the last informations stored asynchronously by the watcher

kill (*signum*)
send a signal to the process

monitor (*listener=None*)
start to monitor the process

Listener can be any callable and receive (“*stat*”, *process_info*)

monitor_io (*io_label, listener*)
subscribe to registered IO events

os_pid
return the process pid

spawn (*once=False, graceful_timeout=None, env=None*)
spawn the process

stats

status
return the process status

stop ()
stop the process

unmonitor (*listener*)
stop monitoring this process.

listener is the callback passed to the monitor function previously.

unmonitor_io (*io_label, listener*)
unsubscribe to the IO event

write (*data*)
send data to the process via stdin

writelines (*data*)
send data to the process via stdin

class `gaffer.process.ProcessConfig` (*name, cmd, **settings*)

Bases: `object`

object to maintain a process config

DEFAULT_PARAMS = {'args': [], 'custom_channels': [], 'custom_streams': [], 'cwd': ''}

classmethod `from_dict` (*config*)

get (*key, default=None*)

make_process (*loop, pid, label, env=None, on_exit=None*)

create a Process object from the configuration

Args:

- **loop**: main pyuv loop instance that will maintain the process
- **pid**: process id, generally given by the manager
- **label**: the job label. Usually the process type. context. A context can be for example an application.
- **on_exit**: callback called when the process exited.

to_dict ()

```

class gaffer.process.ProcessWatcher (loop, process)
    Bases: object

    object to retrieve process stats

    active

    refresh (interval=0)

    stop (all_events=False)

    subscribe (listener)

    subscribe_once (listener)

    unsubscribe (listener)

class gaffer.process.RedirectIO (loop, process, stdio=[])
    Bases: object

    pipes_count = 2

    start ()

    stdio

    stop (all_events=False)

    subscribe (label, listener)

    unsubscribe (label, listener)

class gaffer.process.RedirectStdin (loop, process)
    Bases: object

    redirect stdin allows multiple sender to write to same pipe

    start ()

    stop (all_events=False)

    write (data)

    writelines (data)

class gaffer.process.Stream (loop, process, id)
    Bases: gaffer.process.RedirectStdin

    create custom stdio

    start ()

    subscribe (listener)

    unsubscribe (listener)

gaffer.process.get_process_stats (process=None, interval=0)
    Return information about a process. (can be an pid or a Process object)

    If process is None, will return the information about the current process.

```

Gaffer events

Many events happend in gaffer.

Manager events

Manager events have the following format:

```
{
  "event": "<nameofevent>>",
  "name": "<templatename>"
}
```

- **create**: a process template is created
- **start**: a process template start to launch OS processes
- **stop**: all OS processes of a process template are stopped
- **restart**: all processes of a process template are restarted
- **update**: a process template is updated
- **delete**: a process template is deleted
- **spawn**: a new process is spawned
- **reap**: a process is reaped
- **exit**: a process exited
- **stop_pid**: a process has been stopped

Processes events

All processes' events are prefixed by `proc.<name>` to make the pattern matching easier, where `<name>` is the name of the process template

Events are:

- **proc.<name>.start** : the template `<name>` start to spawn processes
- **proc.<name>.spawn** : one OS process using the process `<name>` template is spawned. Message is:

```
{
  "event": "proc.<name>.spawn">>,
  "name": "<name>",
  "detach": false,
  "pid": int
}
```

Note: pid is the internal pid

- **proc.<name>.exit**: one OS process of the `<name>` template has exited. Message is:

```
{
  "event": "proc.<name>.exit">>,
  "name": "<name>",
  "pid": int,
  "exit_code": int,
  "term_signal": int
}
```

- **proc.<name>.stop**: all OS processes in the template <name> are stopped.
- **proc.<name>.stop_pid**: One OS process of the template <name> is stopped. Message is:

```
{
  "event": "proc.<name>.stop_pid">>,
  "name": "<name>",
  "pid": int
}
```

- **proc.<name>.stop_pid**: One OS process of the template <name> is reaped. Message is:

```
{
  "event": "proc.<name>.reap">>,
  "name": "<name>",
  "pid": int
}
```

The events Module

This module offers a common way to subscribe and emit events. All events in gaffer are using.

Example of usage

```
event = EventEmitter()

# subscribe to all events with the pattern a.*
event.subscribe("a", subscriber)

# subscribe to all events "a.b"
event.subscribe("a.b", subscriber2)

# subscribe to all events (wildcard)
event.subscribe(".", subscriber3)

# publish an event
event.publish("a.b", arg, namedarg=val)
```

In this example all subscribers will be notified of the event. A subscriber is just a callable (*event*, **args*, ***kwargs*)

Classes

class gaffer.events.**EventEmitter** (*loop*, *max_size=10000*)

Bases: object

Many events happen in gaffer. For example a process will emit the events “start”, “stop”, “exit”.

This object offers a common interface to all events emitters

close ()

close the event

This function clears the list of listeners and stops all idle callbacks

publish (*evtype*, **args*, ***kwargs*)
emit an event **evtype**

The event will be emitted asynchronously so we don't block here

subscribe (*evtype*, *listener*, *once=False*)
subscribe to an event

subscribe_once (*evtype*, *listener*)
subscribe to event once. Once the event is triggered we remove ourselves from the list of listeners

unsubscribe (*evtype*, *listener*, *once=False*)
unsubscribe from an event

unsubscribe_all (*events=[]*)
unsubscribe all listeners from a list of events

unsubscribe_once (*evtype*, *listener*)

procfile Module

module to parse and manage a Procfile

class `gaffer.procfile.Procfile` (*procfile*, *root=None*, *envs=None*)
Bases: `object`

Procfile object to parse a procfile and a list of given environment files.

as_configparser (*concurrency_settings=None*)
return a ConfigParser object. It can be used to generate a gaffer setting file or a configuration file that can be included.

as_dict (*name*, *concurrency_settings=None*)
return a procfile line as a JSON object usable with the command `gafferctl load`.

get_appname ()

parse (*procfile*)
main function to parse a procfile. It returns a dict

parse_cmd (*v*)

processes ()
iterator over the configuration

`gaffer.procfile.get_env` (*envs=[]*)

pidfile Module

class `gaffer.pidfile.Pidfile` (*fname*)
Bases: `object`

Manage a PID file. If a specific name is provided it and “`”%s.oldpid” % name’` will be used. Otherwise we create a temp file using `os.mkstemp`.

create (*pid*)

rename (*path*)

unlink ()
delete pidfile

validate()
 Validate pidfile and make it stale if needed

util Module

`gaffer.util.bind_sockets` (*addr*, *backlog=128*, *allows_unix_socket=False*)

`gaffer.util.bytes2human` (*n*)
 Translates bytes into a human repr.

`gaffer.util.bytestring` (*s*)

`gaffer.util.check_gid` (*val*)
 Return a gid, given a group value
 If the group value is unknown, raises a ValueError.

`gaffer.util.check_uid` (*val*)
 Return an uid, given a user value. If the value is an integer, make sure it's an existing uid.
 If the user value is unknown, raises a ValueError.

`gaffer.util.daemonize` ()
 Standard daemonization of a process.

`gaffer.util.from_nanotime` (*n*)
 convert from nanotime to seconds

`gaffer.util.getcwd` ()
 Returns current path, try to use PWD env first

`gaffer.util.hostname` ()

`gaffer.util.is_ipv6` (*addr*)

`gaffer.util.is_ssl` (*url*)

`gaffer.util.nanotime` (*s=None*)
 convert seconds to nanoseconds. If *s* is None, current time is returned

`gaffer.util.ord_` (*c*)

`gaffer.util.parse_address` (*netloc*, *default_port=8000*)

`gaffer.util.parse_job_name` (*name*, *default='default'*)

`gaffer.util.parse_signal_value` (*sig*)

`gaffer.util.parse_ssl_options` (*client_options*)

`gaffer.util.setproctitle_` (*title*)

`gaffer.util.substitute_env` (*s*, *env*)

tornado_pyuv Module

class `gaffer.tornado_pyuv.IOLoop` (*impl=None*, *_loop=None*)
 Bases: object

ERROR = 24

NONE = 0

READ = 1

WRITE = 4

add_callback (*callback*)

add_handler (*fd, handler, events*)

add_timeout (*deadline, callback*)

close ()

handle_callback_exception (*callback*)

This method is called whenever a callback run by the IOloop throws an exception.

By default simply logs the exception as an error. Subclasses may override this method to customize reporting of exceptions.

The exception itself is not passed explicitly, but is available in `sys.exc_info`.

static initialized ()

Returns true if the singleton instance has been created.

install ()

Installs this IOloop object as the singleton instance.

This is normally not necessary as `instance()` will create an IOloop on demand, but you may want to call `install` to use a custom subclass of IOloop.

static instance ()

log_stack (*signal, frame*)

remove_handler (*fd*)

remove_timeout (*timeout*)

running ()

Returns true if this IOloop is currently running.

set_blocking_log_threshold (*seconds*)

set_blocking_signal_threshold (*seconds, action*)

start ()

stop ()

update_handler (*fd, events*)

class `gaffer.tornado_pyuv.PeriodicCallback` (*callback, callback_time, io_loop=None*)

Bases: `object`

start ()

stop ()

class `gaffer.tornado_pyuv.Waker` (*loop*)

Bases: `object`

wake ()

`gaffer.tornado_pyuv.install` ()

1.2.8 httpclient Module

Gaffer provides you a simple Client to control a gaffer node via HTTP.

Example of usage:

```

import pyuv

from gaffer.httpclient import Server

# initialize a loop
loop = pyuv.Loop.default_loop()

s = Server("http://localhost:5000", loop=loop)

# add a process without starting it
process = s.add_process("dummy", "/some/path/to/dummy/script", start=False)

# start a process
process.start()

# increase the number of process by 2 (so 3 will run)
process.add(2)

# stop all processes
process.stop()

loop.run()

```

1.2.9 Gaffer applications

Gaffer applications are applications that are started by the manager. A gaffer application can be used to interact with the manager or listening on events.

An application is a class with the following structure:

```

class Myapplication(object):

    def __init__(self):
        # do inti

    def start(self, loop, manager):
        # this method is call by the manager to start the
        application

    def stop(self):
        # method called when the manager stop

    def restart(self):
        # methhod called when the manager restart

```

Following applications are provided by gaffer:

http_handler Module

console_output Module

module to return all streams from the managed processes to the console. This application is subscribing to the manager to know when a process is created or killed and display the information. When an OS process is spawned it then subscribe to its streams if any are redirected and print the output on the console. This module is used by *Gaffer* .

Note: if `colorize` is set to true, each templates will have a different colour

class `gaffer.console_output.Color`

Bases: `object`

wrapper around `colorama` to ease the output creation. Don't use it directly, instead, use the `colored(name_of_color, lines)` to return the colored output.

Colors are: `cyan`, `yellow`, `green`, `magenta`, `red`, `blue`, `intense_cyan`, `intense_yellow`, `intense_green`, `intense_magenta`, `intense_red`, `intense_blue`.

`lines` can be a list or a string.

output (*color_name, lines*)

class `gaffer.console_output.ConsoleOutput` (*colorize=True, output_streams=True, actions=None*)

Bases: `object`

The application that need to be added to the gaffer manager

DEFAULT_ACTIONS = ['start', 'stop', 'spawn', 'reap', 'exit', 'stop_pid']

restart ()

start (*loop, manager*)

stop ()

`gaffer.console_output.status_bar` (*s*)

sig_handler Module

class `gaffer.sig_handler.BaseSigHandler`

Bases: `object`

A simple gaffer application to handle signals

QUIT_SIGNALS = (3, 15, 2)

handle_quit (*handle, signum*)

handle_reload (*handle, signum*)

restart ()

start (*loop*)

stop ()

class `gaffer.sig_handler.SigHandler`

Bases: `gaffer.sig_handler.BaseSigHandler`

A simple gaffer application to handle signals

handle_quit (*handle, *args*)

handle_reload (*handle, *args*)

start (*loop, manager*)

CHAPTER 2

Indices and tables

- [genindex](#)
- [modindex](#)
- [search](#)

g

`gaffer.console_output`, 31
`gaffer.events`, 25
`gaffer.httpclient`, 30
`gaffer.manager`, 20
`gaffer.pidfile`, 28
`gaffer.process`, 23
`gaffer.procfile`, 28
`gaffer.sig_handler`, 32
`gaffer.tornado_pyuv`, 29
`gaffer.util`, 29
`gaffer.webhooks`, 19

A

active (gaffer.manager.Manager attribute), 21
 active (gaffer.process.Process attribute), 23
 active (gaffer.process.ProcessWatcher attribute), 25
 add_callback() (gaffer.tornado_pyuv.IOLoop method), 30
 add_handler() (gaffer.tornado_pyuv.IOLoop method), 30
 add_timeout() (gaffer.tornado_pyuv.IOLoop method), 30
 as_configparser() (gaffer.procfile.Procfile method), 28
 as_dict() (gaffer.procfile.Procfile method), 28

B

BaseSigHandler (class in gaffer.sig_handler), 32
 bind_sockets() (in module gaffer.util), 29
 bytes2human() (in module gaffer.util), 29
 bytestring() (in module gaffer.util), 29

C

check_gid() (in module gaffer.util), 29
 check_uid() (in module gaffer.util), 29
 close() (gaffer.events.EventEmitter method), 27
 close() (gaffer.process.Process method), 23
 close() (gaffer.tornado_pyuv.IOLoop method), 30
 close() (gaffer.webhooks.WebHooks method), 20
 closed (gaffer.process.Process attribute), 23
 Color (class in gaffer.console_output), 32
 commit() (gaffer.manager.Manager method), 21
 ConsoleOutput (class in gaffer.console_output), 32
 create() (gaffer.pidfile.Pidfile method), 28

D

daemonize() (in module gaffer.util), 29
 decref() (gaffer.webhooks.WebHooks method), 20
 DEFAULT_ACTIONS (gaffer.console_output.ConsoleOutput attribute), 32
 DEFAULT_PARAMS (gaffer.process.ProcessConfig attribute), 24

E

ERROR (gaffer.tornado_pyuv.IOLoop attribute), 29

EventEmitter (class in gaffer.events), 27

F

from_dict() (gaffer.process.ProcessConfig class method), 24
 from_nanotime() (in module gaffer.util), 29

G

gaffer.console_output (module), 31
 gaffer.events (module), 25
 gaffer.httpclient (module), 30
 gaffer.manager (module), 20
 gaffer.pidfile (module), 28
 gaffer.process (module), 23
 gaffer.procfile (module), 28
 gaffer.sig_handler (module), 32
 gaffer.tornado_pyuv (module), 29
 gaffer.util (module), 29
 gaffer.webhooks (module), 19
 get() (gaffer.manager.Manager method), 21
 get() (gaffer.process.ProcessConfig method), 24
 get_appname() (gaffer.procfile.Procfile method), 28
 get_env() (in module gaffer.procfile), 28
 get_process() (gaffer.manager.Manager method), 21
 get_process_id() (gaffer.manager.Manager method), 21
 get_process_stats() (in module gaffer.process), 25
 getcwd() (in module gaffer.util), 29

H

handle_callback_exception()
 (gaffer.tornado_pyuv.IOLoop method), 30
 handle_quit() (gaffer.sig_handler.BaseSigHandler method), 32
 handle_quit() (gaffer.sig_handler.SigHandler method), 32
 handle_reload() (gaffer.sig_handler.BaseSigHandler method), 32
 handle_reload() (gaffer.sig_handler.SigHandler method), 32
 hostname() (in module gaffer.util), 29

I

incref() (gaffer.webhooks.WebHooks method), 20
info (gaffer.process.Process attribute), 23
info() (gaffer.manager.Manager method), 21
initialized() (gaffer.tornado_pyuv.IOLoop static method), 30
install() (gaffer.tornado_pyuv.IOLoop method), 30
install() (in module gaffer.tornado_pyuv), 30
instance() (gaffer.tornado_pyuv.IOLoop static method), 30
IOLoop (class in gaffer.tornado_pyuv), 29
is_ipv6() (in module gaffer.util), 29
is_ssl() (in module gaffer.util), 29

J

jobs() (gaffer.manager.Manager method), 21
jobs_walk() (gaffer.manager.Manager method), 21

K

kill() (gaffer.manager.Manager method), 21
kill() (gaffer.process.Process method), 23
killall() (gaffer.manager.Manager method), 21

L

list() (gaffer.manager.Manager method), 21
load() (gaffer.manager.Manager method), 21
log_stack() (gaffer.tornado_pyuv.IOLoop method), 30

M

make_process() (gaffer.process.ProcessConfig method), 24
manage() (gaffer.manager.Manager method), 22
Manager (class in gaffer.manager), 20
maybe_start_monitor() (gaffer.webhooks.WebHooks method), 20
maybe_stop_monitor() (gaffer.webhooks.WebHooks method), 20
monitor() (gaffer.manager.Manager method), 22
monitor() (gaffer.process.Process method), 24
monitor_io() (gaffer.process.Process method), 24

N

nanotime() (in module gaffer.util), 29
NONE (gaffer.tornado_pyuv.IOLoop attribute), 29

O

ord_() (in module gaffer.util), 29
os_pid (gaffer.process.Process attribute), 24
output() (gaffer.console_output.Color method), 32

P

parse() (gaffer.procfile.Procfile method), 28
parse_address() (in module gaffer.util), 29

parse_cmd() (gaffer.procfile.Procfile method), 28
parse_job_name() (in module gaffer.util), 29
parse_signal_value() (in module gaffer.util), 29
parse_ssl_options() (in module gaffer.util), 29
PeriodicCallback (class in gaffer.tornado_pyuv), 30
Pidfile (class in gaffer.pidfile), 28
pids() (gaffer.manager.Manager method), 22
pipes_count (gaffer.process.RedirectIO attribute), 25
Process (class in gaffer.process), 23
ProcessConfig (class in gaffer.process), 24
processes() (gaffer.procfile.Procfile method), 28
ProcessWatcher (class in gaffer.process), 24
Procfile (class in gaffer.procfile), 28
publish() (gaffer.events.EventEmitter method), 27

Q

QUIT_SIGNALS (gaffer.sig_handler.BaseSigHandler attribute), 32

R

READ (gaffer.tornado_pyuv.IOLoop attribute), 29
RedirectIO (class in gaffer.process), 25
RedirectStdin (class in gaffer.process), 25
refcount (gaffer.webhooks.WebHooks attribute), 20
refresh() (gaffer.process.ProcessWatcher method), 25
register_hook() (gaffer.webhooks.WebHooks method), 20
reload() (gaffer.manager.Manager method), 22
remove_handler() (gaffer.tornado_pyuv.IOLoop method), 30
remove_timeout() (gaffer.tornado_pyuv.IOLoop method), 30
rename() (gaffer.pidfile.Pidfile method), 28
restart() (gaffer.console_output.ConsoleOutput method), 32
restart() (gaffer.manager.Manager method), 22
restart() (gaffer.sig_handler.BaseSigHandler method), 32
restart() (gaffer.webhooks.WebHooks method), 20
run() (gaffer.manager.Manager method), 22
running() (gaffer.tornado_pyuv.IOLoop method), 30

S

scale() (gaffer.manager.Manager method), 22
send() (gaffer.manager.Manager method), 22
sessions (gaffer.manager.Manager attribute), 22
set_blocking_log_threshold() (gaffer.tornado_pyuv.IOLoop method), 30
set_blocking_signal_threshold() (gaffer.tornado_pyuv.IOLoop method), 30
setproctitle_() (in module gaffer.util), 29
SigHandler (class in gaffer.sig_handler), 32
spawn() (gaffer.process.Process method), 24
start() (gaffer.console_output.ConsoleOutput method), 32
start() (gaffer.manager.Manager method), 22
start() (gaffer.process.RedirectIO method), 25

start() (gaffer.process.RedirectStdin method), 25
 start() (gaffer.process.Stream method), 25
 start() (gaffer.sig_handler.BaseSigHandler method), 32
 start() (gaffer.sig_handler.SigHandler method), 32
 start() (gaffer.tornado_pyuv.IOLoop method), 30
 start() (gaffer.tornado_pyuv.PeriodicCallback method), 30
 start() (gaffer.webhooks.WebHooks method), 20
 start_job() (gaffer.manager.Manager method), 22
 stats (gaffer.process.Process attribute), 24
 stats() (gaffer.manager.Manager method), 22
 status (gaffer.process.Process attribute), 24
 status_bar() (in module gaffer.console_output), 32
 stdio (gaffer.process.RedirectIO attribute), 25
 stop() (gaffer.console_output.ConsoleOutput method), 32
 stop() (gaffer.manager.Manager method), 22
 stop() (gaffer.process.Process method), 24
 stop() (gaffer.process.ProcessWatcher method), 25
 stop() (gaffer.process.RedirectIO method), 25
 stop() (gaffer.process.RedirectStdin method), 25
 stop() (gaffer.sig_handler.BaseSigHandler method), 32
 stop() (gaffer.tornado_pyuv.IOLoop method), 30
 stop() (gaffer.tornado_pyuv.PeriodicCallback method), 30
 stop() (gaffer.webhooks.WebHooks method), 20
 stop_job() (gaffer.manager.Manager method), 22
 stop_process() (gaffer.manager.Manager method), 22
 stopall() (gaffer.manager.Manager method), 22
 Stream (class in gaffer.process), 25
 subscribe() (gaffer.events.EventEmitter method), 28
 subscribe() (gaffer.manager.Manager method), 22
 subscribe() (gaffer.process.ProcessWatcher method), 25
 subscribe() (gaffer.process.RedirectIO method), 25
 subscribe() (gaffer.process.Stream method), 25
 subscribe_once() (gaffer.events.EventEmitter method), 28
 subscribe_once() (gaffer.process.ProcessWatcher method), 25
 substitute_env() (in module gaffer.util), 29

T

to_dict() (gaffer.process.ProcessConfig method), 24

U

unlink() (gaffer.pidfile.Pidfile method), 28
 unload() (gaffer.manager.Manager method), 22
 unmonitor() (gaffer.manager.Manager method), 22
 unmonitor() (gaffer.process.Process method), 24
 unmonitor_io() (gaffer.process.Process method), 24
 unregister_hook() (gaffer.webhooks.WebHooks method), 20
 unsubscribe() (gaffer.events.EventEmitter method), 28
 unsubscribe() (gaffer.manager.Manager method), 23
 unsubscribe() (gaffer.process.ProcessWatcher method), 25

unsubscribe() (gaffer.process.RedirectIO method), 25
 unsubscribe() (gaffer.process.Stream method), 25
 unsubscribe_all() (gaffer.events.EventEmitter method), 28
 unsubscribe_once() (gaffer.events.EventEmitter method), 28
 update() (gaffer.manager.Manager method), 23
 update_handler() (gaffer.tornado_pyuv.IOLoop method), 30

V

validate() (gaffer.pidfile.Pidfile method), 28

W

wake() (gaffer.tornado_pyuv.Waker method), 30
 Waker (class in gaffer.tornado_pyuv), 30
 wakeup() (gaffer.manager.Manager method), 23
 walk() (gaffer.manager.Manager method), 23
 WebHooks (class in gaffer.webhooks), 20
 WRITE (gaffer.tornado_pyuv.IOLoop attribute), 29
 write() (gaffer.process.Process method), 24
 write() (gaffer.process.RedirectStdin method), 25
 writelines() (gaffer.process.Process method), 24
 writelines() (gaffer.process.RedirectStdin method), 25