

---

# **Circus Documentation**

***Release 0.9.2***

**Mozilla Foundation**

July 22, 2013



# CONTENTS





Circus can be used to monitor and control processes and sockets.

Circus can be driven via a command-line interface or programmatically through its python API. Circus is designed using [ZeroMQ](#). See [Design](#) for more details.

If you are curious about what Circus brings compared to other projects, read [Why should I use Circus instead of X ?](#).

**Warning:** By default, Circus doesn't secure its messages when sending information through ZeroMQ. Before running Circus, make sure you read the [Security](#) page.

To install it and try its features check out the [Step-by-step tutorial](#).



# RUNNING CIRCUS

Circus provides a command-line script call **circusd** that can be used to manage one or more *watchers*. Each watcher can have one or more running *processes*.

Circus' command-line tool is configurable using an ini-style configuration file.

Here's a very minimal example:

```
[watcher:myprogram]
cmd = python
args = -u myprogram.py $WID
numprocesses = 5

[watcher:anotherprogram]
cmd = another_program
numprocesses = 2
```

The file is then passed to *circusd*:

```
$ circusd example.ini
```

Besides processes, Circus can also bind sockets. Since every process managed by Circus is a child of the main Circus daemon, that means any program that's controlled by Circus can use those sockets.

Running a socket is as simple as adding a *socket* section in the config file:

```
[socket:mysocket]
host = localhost
port = 8080
```

To learn more about sockets, see *Circus Sockets*.

To understand why it's a killer feature, read *Circus stack v.s. Classical stack*.





# CONTROLLING CIRCUS

Circus provides two command-line tools to manage your running daemon:

- *circusctl*, a management console you can use it to perform actions such as adding or removing *workers*
- *circus-top*, a top-like console you can use to display the memory and cpu usage of your running Circus.

To learn more about these, see [Command-line tools](#)

Circus also offers a small web application that can connect to a running Circus daemon and let you monitor and interact with it.

Running the web application is as simple as adding an **httpd** option in the ini file in the *circus* section:

```
[circus]
httpd = True
```

Or if you want, you can run it as a standalone process with:

```
$ circushttpd
```

By default, **circushttpd** runs on the *8080* port.

To learn more about this feature, see [The Web Console](#)



# DEVELOPING WITH CIRCUS

Circus provides high-level classes and functions that will let you manage processes in your own applications.

For example, if you want to run four processes forever, you could write:

```
from circus import get_arbiter

myprogram = {"cmd": "python myprogram.py", "numprocesses": 4}

arbiter = get_arbiter([myprogram])
try:
    arbiter.start()
finally:
    arbiter.stop()
```

This snippet will run four instances of *myprogram* and watch them for you, restarting them if they die unexpectedly.

To learn more about this, see [\*Circus Library\*](#)



# EXTENDING CIRCUS

It's easy to extend Circus to create a more complex system, by listening to all the **circusd** events via its pub/sub channel, and driving it via commands.

That's how the flapping feature works for instance: it listens to all the processes dying, measures how often it happens, and stops the incriminated watchers after too many restarts attempts.

Circus comes with a plugin system to help you write such extensions, and a few built-in plugins you can reuse. See [Plugins](#).

You can also have a more subtile startup and shutdown behavior by using the **hooks** system that will let you run arbitrary code before and after some processes are started or stopped. See [Hooks](#).



# MORE DOCUMENTATION

## 5.1 Installing Circus

Use pip:

```
$ pip install circus
```

Or download the archive on PyPI, extract and install it manually with:

```
$ python setup.py install
```

If you want to try out Circus, see the *Step-by-step tutorial*.

If you are using debian or any debian based distribution, you also can use the ppa to install circus, it's at <https://launchpad.net/~roman-imankulov/+archive/circus>

### 5.1.1 `zc.buildout`

We provide a `zc.buildout` configuration, you can use it by simply running the bootstrap script, then calling buildout:

```
$ python bootstrap.py  
$ bin/buildout
```

### 5.1.2 More on Requirements

Circus uses:

- Python 2.6, 2.7 (3.x needs to be tested)
- zeromq  $\geq 2.10$  (you need to use the 2.X series)

And on Python side:

- pyzmq 2.2.0.1
- iowait 0.1
- psutil 0.4.1

You can install all the py dependencies with the `pip-requirements.txt` file we provide manually, or just install Circus and have the latest versions of those libraries pulled for you:

```
$ pip install -r pip-requirements.txt
```

If you want to run the Web console you will need to install **circus-web**:

```
$ pip install circus-web
```

## 5.2 Configuration

Circus can be configured using an ini-style configuration file.

Example:

```
[circus]
check_delay = 5
endpoint = tcp://127.0.0.1:5555
pubsub_endpoint = tcp://127.0.0.1:5556
include = *.more.config.ini

[watcher:myprogram]
cmd = python
args = -u myprogram.py $(circus.wid) $(ENV.VAR)
warmup_delay = 0
numprocesses = 5

# hook
hooks.before_start = my.hooks.control_redis

# will push in test.log the stream every 300 ms
stdout_stream.class = FileStream
stdout_stream.filename = test.log
stdout_stream.refresh_time = 0.3

# optionally rotate the log file when it reaches 1 gb
# and save 5 copied of rotated files
stdout_stream.max_bytes = 1073741824
stdout_stream.backup_count = 5

[env:myprogram]
PATH = $PATH:/bin
CAKE = lie

[plugin:statsd]
use = circus.plugins.statsd.StatsdEmitter
host = localhost
port = 8125
sample_rate = 1.0
application_name = example

[socket:web]
host = localhost
port = 8080
```

### 5.2.1 circus - single section

**endpoint** The ZMQ socket used to manage Circus via **circusctl**. (default: *tcp://127.0.0.1:5555*)

**pubsub\_endpoint** The ZMQ PUB/SUB socket receiving publications of events. (default: *tcp://127.0.0.1:5556*)



**statsd** If set to True, Circus runs the circusd-stats daemon. (default: False)

**stats\_endpoint** The ZMQ PUB/SUB socket receiving publications of stats. (default: `tcp://127.0.0.1:5557`)

**statsd\_close\_outputs** If True sends the circusd-stats stdout/stderr to /dev/null. (default: False)

**check\_delay** The polling interval in seconds for the ZMQ socket. (default: 5)

**include** List of config files to include. You can use wildcards (\*) to include particular schemes for your files. The paths are absolute or relative to the config file. (default: None)

**include\_dir** List of config directories. All files matching \*.ini under each directory will be included. The paths are absolute or relative to the config file. (default: None)

**stream\_backend** Defines the type of backend to use for the streaming. Possible values are **thread** or **gevent**. (default: thread)

**warmup\_delay** The interval in seconds between two watchers start. Must be an int. (default: 0)

**httpd** If set to True, Circus runs the circushttpd daemon. (default: False)

**httpd\_host** The host ran by the circushttpd daemon. (default: localhost)

**httpd\_port** The port ran by the circushttpd daemon. (default: 8080)

**httpd\_close\_outputs** If True, sends the circushttpd stdout/stderr to /dev/null. (default: False)

**debug** If set to True, all Circus stout/stderr daemons are redirected to circusd stdout/stderr (default: False)

**pidfile** The file that must be used to keep the daemon pid.

**loglevel** The loglevel that we want to see (default: INFO)

**logoutput** The logoutput file where we want to log (default: stdout)

### 5.2.2 watcher:NAME - as many sections as you want

**NAME** The name of the watcher. This name is used in **circusctl**

**cmd** The executable program to run.

**args** Command-line arguments to pass to the program. You can use the python format syntax here to build the parameters. Environment variables are available, as well as the worker id and the environment variables that you passed, if any, with the “env” parameter. See *Formating the commands and arguments with dynamic variables* for more information on this.

**shell** If True, the processes are run in the shell (default: False)

**working\_dir** The working dir for the processes (default: None)

**uid** The user id or name the command should run with. (The current uid is the default).

**gid** The group id or name the command should run with. (The current gid is the default).

**copy\_env** If set to true, the local environment variables will be copied and passed to the workers when spawning them. (Default: False)

**copy\_path** If set to true, **sys.path** is passed in the subprocess environ using **PYTHONPATH**. **copy\_env** has to be true. (Default: False)

**warmup\_delay** The delay (in seconds) between running processes.

**autostart** If set to true, the watcher will not be started automatically when the arbiter starts. The watcher can be started explicitly (example: `circusctl start myprogram`). (Default: True)

**numprocesses** The number of processes to run for this watcher.

**rlimit\_LIMIT** Set resource limit LIMIT for the watched processes. The config name should match the RLIMIT\_\* constants (not case sensitive) listed in the [Python resource module reference](#). For example, the config line 'rlimit\_nofile = 500' sets the maximum number of open files to 500.

**stderr\_stream.class** A fully qualified Python class name that will be instantiated, and will receive the **stderr** stream of all processes in its `__call__()` method.

Circus provides some stream classes you can use without prefix:

- `FileStream`: writes in a file
- `QueueStream`: write in a memory `Queue`
- `StdoutStream`: writes in the `stdout`
- `FancyStdoutStream`: writes colored output with time prefixes in the `stdout`

**stderr\_stream.\*** All options starting with *stderr\_stream*. other than *class* will be passed the constructor when creating an instance of the class defined in **stderr\_stream.class**.

**stdout\_stream.class** A fully qualified Python class name that will be instantiated, and will receive the **stdout** stream of all processes in its `__call__()` method.

Circus provides soem stream classes you can use without prefix:

- `FileStream`: writes in a file
- `QueueStream`: write in a memory `Queue`
- `StdoutStream`: writes in the `stdout`
- `FancyStdoutStream`: writes colored output with time prefixes in the `stdout`

**stdout\_stream.\*** All options starting with *stdout\_stream*. other than *class* will be passed the constructor when creating an instance of the class defined in **stdout\_stream.class**.

**close\_child\_stdout** If set to True, the sdout stream of each process will be sent to /dev/null after the fork. Defaults to False.

**close\_child\_stderr** If set to True, the sderr stream of each process will be sent to /dev/null after the fork. Defaults to False.

**send\_hup** if True, a process reload will be done by sending the SIGHUP signal. Defaults to False.

**max\_retry** The number of times we attempt to start a process, before we abandon and stop the whole watcher. Defaults to 5. Set to -1 to disable max\_retry and retry indefinitely.

**priority** Integer that defines a priority for the watcher. When the Arbiter do some operations on all watchers, it will sort them with this field, from the bigger number to the smallest. Defaults to 0.

**singleton** If set to True, this watcher will have at the most one process. Defaults to False.

**use\_sockets** If set to True, this watcher will be able to access defined sockets via their file descriptors. If False, all parent fds are closed when the child process is forked. Defaults to False.

**max\_age** If set then the process will be restarted sometime after max\_age seconds. This is useful when processes deal with pool of connectors: restarting processes improves the load balancing. Defaults to being disabled.

**max\_age\_variance** If max\_age is set then the process will live between max\_age and max\_age + random(0, max\_age\_variance) seconds. This avoids restarting all processes for a watcher at once. Defaults to 30 seconds.

**on\_demand** If set to `True`, the processes will be started only after the first connection to one of the configured sockets (see below). If a restart is needed, it will be only triggered at the next socket event.

**hooks.\*** Available hooks: **before\_start**, **before\_spawn**, **after\_start**, **before\_stop**, **after\_stop**

Define callback functions that hook into the watcher startup/shutdown process.

If the hook returns **False** and if the hook is one of **before\_start**, **before\_spawn** or **after\_start**, the startup will be aborted.

Notice that a hook that fails during the stopping process will not abort it.

The callback definition can be followed by a boolean flag separated by a comma. When the flag is set to **true**, any error occurring in the hook will be ignored. If set to **false** (the default), the hook will return **False**.

More on *Hooks*.

**virtualenv** When provided, points to the root of a Virtualenv directory. The watcher will scan the local **site-packages** and loads its content into the execution environment. Must be used with **copy\_env** set to `True`. Defaults to `None`.

**respawn** If set to `False`, the processes handled by a watcher will not be respawned automatically. (default: `True`)

### 5.2.3 socket:NAME - as many sections as you want

**host** The host of the socket. Defaults to `'localhost'`

**port** The port. Defaults to `8080`.

**family** The socket family. Can be `'AF_UNIX'`, `'AF_INET'` or `'AF_INET6'`. Defaults to `'AF_INET'`.

**type** The socket type. Can be `'SOCK_STREAM'`, `'SOCK_DGRAM'`, `'SOCK_RAW'`, `'SOCK_RDM'` or `'SOCK_SEQPACKET'`. Defaults to `'SOCK_STREAM'`.

**path** When provided a path to a file that will be used as a unix socket file. If a path is provided, **family** is forced to `AF_UNIX` and **host** and **port** are ignored.

**umask** When provided, sets the umask that will be used to create an `AF_UNIX` socket. For example, `umask=000` will produce a socket with permission `777`.

Once a socket is created, the `${circus.sockets.NAME}` string can be used in the command (`cmd` or `args`) of a watcher. Circus will replace it by the FD value. The watcher must also have `use_sockets` set to `True` otherwise the socket will have been closed and you will get errors when the watcher tries to use it.

Example:

```
[watcher:webworker]
cmd = chaussette --fd ${circus.sockets.webapp} chaussette.util.bench_app
use_sockets = True
```

```
[socket:webapp]
host = 127.0.0.1
port = 8888
```

### 5.2.4 plugin:NAME - as many sections as you want

**use** The fully qualified name that points to the plugin class.

**anything else** Every other key found in the section is passed to the plugin constructor in the **config** mapping.

Circus comes with a few pre-shipped *plugins* but you can also extend them easily by *developing your own*.

### 5.2.5 env or env[:WATCHERS] - as many sections as you want

**anything** The name of an environment variable to assign value to. bash style environment substitutions are supported. for example, append /bin to *PATH* 'PATH = \$PATH:/bin'

Section responsible for delivering environment variable to run processes.

Example:

```
[watcher:worker1]
cmd = ping 127.0.0.1

[watcher:worker2]
cmd = ping 127.0.0.1

[env]
CAKE = lie
```

The variable *CAKE* will be propagated to all watchers defined in config file.

**WATCHERS** can be a comma separated list of watcher sections to apply this environment to. if multiple env sections match a watcher, they will be combine in the order they appear in the configuration file. later entries will take precedence.

Example:

```
[watcher:worker1]
cmd = ping 127.0.0.1

[watcher:worker2]
cmd = ping 127.0.0.1

[env:worker1,worker2]
PATH = /bin

[env:worker1]
PATH = $PATH

[env:worker2]
CAKE = lie
```

*worker1* will be run with PATH = \$PATH (expanded from the environment circusd was run in) *worker2* will be run with PATH = /bin and CAKE = lie

It's possible to use wildcards as well.

Example:

```
[watcher:worker1]
cmd = ping 127.0.0.1

[watcher:worker2]
cmd = ping 127.0.0.1

[env:worker*]
PATH = /bin
```

Both *worker1* and *worker2* will be run with `PATH = /bin`

## 5.2.6 Using environment variables

When writing your configuration file, you can use environment variables defined in the *env* section or in *os.environ* itself.

You just have to use the *circus.env.* prefix.

Example:

```
[watcher:worker1]
cmd = $(circus.env.shell)

[watcher:worker2]
baz = $(circus.env.user)
bar = $(circus.env.yeah)
sup = $(circus.env.oh)

[env]
yeah = boo

[env:worker2]
oh = ok
```

If a variable is defined in several places, the most specialized value has precedence: a variable defined in *env:XXX* will override a variable defined in *env*, which will override a variable defined in *os.environ*.

## 5.2.7 Formating the commands and arguments with dynamic variables

As you may have seen, it is possible to pass some information that are computed dynamically when running the processes. Among other things, you can get the worker id (WID) and all the options that are passed to the `Process`. Additionally, it is possible to access the options passed to the `Watcher` which instanciated the process.

---

**Note:** The worker id is different from the process id. It's a unique value, starting at 1, which is only unique for the watcher.

---

For instance, if you want to access some variables that are contained in the environment, you would need to do it with a setting like this:

```
cmd = "make-me-a-coffee --sugar $(CIRCUS.ENV.SUGAR_AMOUNT) "
```

This works with both *cmd* and *args*.

**Important:**

- All variables are prefixed with *circus*.
- The replacement is case insensitive.

## 5.3 Plugins

Circus comes with a few pre-shipped plugins you can use easily. The configuration of them is as follows:

### 5.3.1 Statsd

**use** set to 'circus.plugins.statsd.StatsdEmitter'

**application\_name** the name used to identify the bucket prefix to emit the stats to (it will be prefixed with "circus." and suffixed with ".watcher")

**host** the host to post the statsd data to

**port** the port the statsd daemon listens on

**sample\_rate** if you prefer a different sample rate than 1, you can set it here

### 5.3.2 FullStats

An extension on the Statsd plugin that is also publishing the process stats. As such it has the same configuration options as Statsd and the following.

**use** set to 'circus.plugins.statsd.FullStats'

**loop\_rate** the frequency the plugin should ask for the stats in seconds. Default: 60.

### 5.3.3 RedisObserver

This services observers a redis process for you, publishes the information to statsd and offers to restart the watcher when it doesn't react in a given timeout. This plugin requires [redis-py](#) to run.

It has the same configuration as statsd and adds the following:

**use** set to 'circus.plugins.redis\_observer.RedisObserver'

**loop\_rate** the frequency the plugin should ask for the stats in seconds. Default: 60.

**redis\_url** the database to check for as a redis url. Default: "redis://localhost:6379/0"

**timeout** the timeout in seconds the request can take before it is considered down. Defaults to 5.

**restart\_on\_timeout** the name of the process to restart when the request timed out. No restart triggered when not given. Default: None.

### 5.3.4 HttpObserver

This services observers a http process for you by pinging a certain website regularly. Similar to the redis observer it offers to restart the watcher on an error. It requires [tornado](#) to run.

It has the same configuration as statsd and adds the following:

**use** set to 'circus.plugins.http\_observer.HttpObserver'

**loop\_rate** the frequency the plugin should ask for the stats in seconds. Default: 60.

**check\_url** the url to check for. Default: "http://localhost/"

**timeout** the timeout in seconds the request can take before it is considered down. Defaults to 10.

**restart\_on\_error** the name of the process to restart when the request timed out or returned any other kind of error. No restart triggered when not given. Default: None.

### 5.3.5 ResourceWatcher

This services watches the resources of the given process and triggers a restart when they exceed certain limitations too often in a row.

It has the same configuration as statsd and adds the following:

**use** set to 'circus.plugins.resource\_watcher.ResourceWatcher'

**loop\_rate** the frequency the plugin should ask for the stats in seconds. Default: 60.

**watcher** the watcher this resource watcher should be looking after. (previously called `service` but `service` is now deprecated)

**max\_cpu** The maximum cpu one process is allowed to consume (in %). Default: 90

**min\_cpu** The minimum cpu one process should consume (in %). Default: None (no minimum) You can set the `min_cpu` to 0 (zero), in this case if one process consume exactly 0% cpu, it will trigger an exceeded limit.

**max\_mem** The amount of memory one process of this watcher is allowed to consume (in %). Default: 90

**min\_mem** The minimum memory one process should consume (in %). Default: None (no minimum) You can set the `min_mem` to 0 (zero), in this case if one process consume exactly 0% memory, it will trigger an exceeded limit.

**health\_threshold** The health is the average of cpu and memory (in %) the watchers processes are allowed to consume (in %). Default: 75

**max\_count** How often these limits (each one is counted separately) are allowed to be exceeded before a restart will be triggered. Default: 3

Example:

```
[circus]
...

[watcher:program]
cmd = sleep 120

[plugin:myplugin]
use = circus.plugins.resource_watcher.ResourceWatcher
watcher = program
min_cpu = 10
max_cpu = 70
min_mem = 0
max_mem = 20
```

### 5.3.6 Watchdog

Plugin that binds an udp socket and wait for watchdog messages. For “watchdoged” processes, the watchdog will kill them if they don’t send a heartbeat in a certain period of time materialized by `loop_rate` \* `max_count`. (circus will automatically restart the missing processes in the watcher)

Each monitored process should send udp message at least at the `loop_rate`. The udp message format is a line of text, decoded using **msg\_regex** parameter. The heartbeat message **MUST** at least contain the pid of the process sending the message.

The list of monitored watchers are determined by the parameter **watchers\_regex** in the configuration.

Configuration parameters:

**loop\_rate** watchdog loop rate in seconds. At each loop, WatchDog will look for “dead” processes.

**watchers\_regex** regex for matching watcher names that should be monitored by the watchdog (default: “.\*” all watchers are monitored)

**msg\_regex** regex for decoding the received heartbeat message in udp (default: “^(?P<pid>.\*)((?P<timestamp>.\*)\$”) the default format is a simple text message: “pid;timestamp”

**max\_count** max number of passed loop without receiving any heartbeat before restarting process (default: 3)

**ip** ip the watchdog will bind on (default: 127.0.0.1)

**port** port the watchdog will bind on (default: 1664)

## 5.4 Hooks

Circus provides four hooks that can be used to trigger actions when a watcher is starting or stopping.

A typical use case is to control that all the conditions are met for a process to start.

Let’s say you have a watcher that runs *Redis* and a watcher that runs a Python script that works with *Redis*.

With Circus you can order the startup by using the **priority** option:

```
[watcher:queue-worker]
cmd = python -u worker.py
priority = 2
```

```
[watcher:redis]
cmd = redis-server
priority = 1
```

With this setup, Circus will start **Redis** then the queue worker.

But Circus does not really control that *Redis* is up and running. It just starts the process it was asked to start.

What we miss here is a way to control that *Redis* is started, and fully functional. A function that controls this could be:

```
import redis
import time

def check_redis(*args, **kw):
    time.sleep(.5) # give it a chance to start
    r = redis.StrictRedis(host='localhost', port=6379, db=0)
    r.set('foo', 'bar')
    return r.get('foo') == 'bar'
```

This function can be plugged into Circus as a *after\_start* hook:

```
[watcher:queue-worker]
cmd = python -u worker.py
hooks.before_start = mycoolapp.myplugins.check_redis
priority = 2
```

```
[watcher:redis]
cmd = redis-server
priority = 1
```



Once Circus has started the **redis** watcher, it will start the **queue-worker** watcher, since it follows the **priority** ordering.

Just before starting the second watcher, it will run the **check\_redis** function, and in case it returns **False** will abort the watcher starting process.

Available hooks are:

- **before\_start**: called before the watcher is started. If the hook returns **False** the startup is aborted.
- **after\_start**: called after the watcher is started. If the hook returns **False** the watcher is immediately stopped and the startup is aborted.
- **before\_spawn**: called before the watcher spawns a new process. is stopped. If the hook returns **False** the watcher is immediately stopped and the startup is aborted.
- **before\_stop**: called before the watcher is stopped. The hook result is ignored.
- **after\_stop**: called after the watcher is stopped. The hook result is ignored.

### 5.4.1 Hook signature

A hook must follow this signature:

```
def hook(watcher, arbiter, hook_name):
    ...
```

Where **watcher** is the **Watcher** class instance, **arbiter** the **Arbiter** one, and **hook\_name** the hook name.

You can ignore those but being able to use the watcher and/or arbiter data and methods can be useful in some hooks.

### 5.4.2 Hook events

Everytime a hook is run, its result is notified as an event in Circus.

There are two events related to hooks:

- **hook\_success**:: a hook was successfully called. The event keys are **name** the name of the event, and **time**: the date of the events.
- **hook\_failure**:: a hook has failed. The event keys are **name** the name of the event, **time**: the date of the events and **error**: the exception that occurred in the event, if any.

## 5.5 Command-line tools

### 5.5.1 circus-top

*circus-top* is a top-like console you can run to watch live your running Circus system. It will display the CPU, Memory usage and socket hits if you have some.

Example of output:

```
-----
circusd-stats
PID                CPU (%)           MEMORY (%)
14252              0.8              0.4
                  0.8 (avg)       0.4 (sum)
```

dummy		
PID	CPU (%)	MEMORY (%)
14257	78.6	0.1
14256	76.6	0.1
14258	74.3	0.1
14260	71.4	0.1
14259	70.7	0.1
	74.32 (avg)	0.5 (sum)

-----

*circus-top* is a read-only console. If you want to interact with the system, use *circusctl*.

### 5.5.2 circusctl

*circusctl* can be used to run any command listed in *Commands* . For example, you can get a list of all the watchers, you can do

```
$ circusctl list
```

Besides supporting a handful of options you can also specify the endpoint *circusctl* should use using the `CIRCUSCTL_ENDPOINT` environment variable.

## 5.6 Commands

At the epicenter of circus lives the command systems. *circusctl* is just a zeromq client, and if needed you can drive programmatically the Circus system by writing your own zmq client.

All messages are Json mappings.

For each command below, we provide a usage example with *circusctl* but also the input / output zmq messages.

### 5.6.1 circusctl commands

- **add:** *Add a watcher*
- **decr:** *Decrement the number of processes in a watcher*
- **dstats:** *Get circusd stats*
- **get:** *Get the value of a watcher option*
- **globaloptions:** *Get the arbiter options*
- **incr:** *Increment the number of processes in a watcher*
- **list:** *Get list of watchers or processes in a watcher*
- **listen:** *Subscribe to a watcher event*
- **listsockets:** *Get the list of sockets*
- **numprocesses:** *Get the number of processes*
- **numwatchers:** *Get the number of watchers*
- **options:** *Get the value of a watcher option*
- **quit:** *Quit the arbiter immediately*

- **reload:** *Reload the arbiter or a watcher*
- **reloadconfig:** *Reload the configuration file*
- **restart:** *Restart the arbiter or a watcher*
- **rm:** *Remove a watcher*
- **set:** *Set a watcher option*
- **signal:** *Send a signal*
- **start:** *Start the arbiter or a watcher*
- **stats:** *Get process infos*
- **status:** *Get the status of a watcher or all watchers*
- **stop:** *Stop the arbiter or a watcher*

## Add a watcher

This command add a watcher dynamically to a arbiter.

### ZMQ Message

```
{
  "command": "add",
  "properties": {
    "cmd": "/path/to/commandline --option"
    "name": "nameofwatcher"
    "args": [],
    "options": {},
    "start": false
  }
}
```

A message contains 2 properties:

- **cmd:** Full command line to execute in a process
- **args:** array, arguments passed to the command (optional)
- **name:** name of watcher
- **options:** options of a watcher
- **start:** start the watcher after the creation

The response return a status “ok”.

### Command line

```
$ circusctl add [--start] <name> <cmd>
```

### Options

- <name>: name of the watcher to create
- <cmd>: full command line to execute in a process
- -start: start the watcher immediately

### Decrement the number of processes in a watcher

This command decrements the number of processes in a watcher by -1.

### ZMQ Message

```
{
  "command": "decr",
  "properties": {
    "name": "<watchername>",
    "nb": <nbprocess>
  }
}
```

The response returns the number of processes in the 'numprocesses' property:

```
{ "status": "ok", "numprocesses": <n>, "time", "timestamp" }
```

### Command line

```
$ circusctl descr <name> [<nb>]
```

### Options

- <name>: name of the watcher
- <nb>: the number of processes to remove.

### Get circusd stats

You can get at any time some statistics about circusd with the dstat command.

### ZMQ Message

To get the circusd stats, simply run:

```
{
  "command": "dstats"
}
```

The response returns a mapping the property "infos" containing some process informations:

```
{
  "info": {
    "children": [],
    "cmdline": "python",
    "cpu": 0.1,
    "ctime": "0:00.41",
    "mem": 0.1,
    "mem_info1": "3M",
    "mem_info2": "2G",
    "nice": 0,
    "pid": 47864,
    "username": "root"
  },
  "status": "ok",
  "time": 1332265655.897085
}
```

## Command Line

```
$ circusctl dstats
```

## Get the value of a watcher option

This command return the watchers options values asked.

## ZMQ Message

```
{
  "command": "get",
  "properties": {
    "keys": ["key1", "key2"]
    "name": "nameofwatcher"
  }
}
```

A response contains 2 properties:

- keys: list, The option keys for which you want to get the values
- name: name of watcher

The response return an object with a property “options” containing the list of key/value returned by circus.

eg:

```
{
  "status": "ok",
  "options": {
    "flapping_window": 1,
    "times": 2
  },
  'time': 1332202594.754644
}
```

See Options for for a description of options enabled?

### Command line

```
$ circusctl get <name> <key> <value> <key1> <value1>
```

### Get the arbiter options

This command return the arbiter options

### ZMQ Message

```
{
  "command": "globaloptions",
  "properties": {
    "key1": "val1",
    ..
  }
}
```

A message contains 2 properties:

- keys: list, The option keys for which you want to get the values

The response return an object with a property “options” containing the list of key/value returned by circus.

eg:

```
{
  "status": "ok",
  "options": {
    "check_delay": 1,
    ...
  },
  'time': 1332202594.754644
}
```

### Command line

```
$ circusctl globaloptions
```

### Options

Options Keys are:

- endpoint: the controller ZMQ endpoint
- pubsub\_endpoint: the pubsub endpoint
- check\_delay: the delay between two controller points
- multicast\_endpoint: the multicast endpoint for circusd cluster auto-discovery

### Increment the number of processes in a watcher

This comment increment the number of processes in a watcher by +1.

### ZMQ Message

```
{
  "command": "incr",
  "properties": {
    "name": "<watchername>",
    "nb": <nbprocess>
  }
}
```

The response return the number of processes in the ‘numprocesses’ property:

```
{ "status": "ok", "numprocesses": <n>, "time", "timestamp" }
```

### Command line

```
$ circusctl incr <name> [<nb>]
```

### Options

- <name>: name of the watcher.
- <nb>: the number of processes to add.

### Get list of watchers or processes in a watcher

#### ZMQ Message

To get the list of all the watchers:

```
{
  "command": "list",
}
```

To get the list of active processes in a watcher:

```
{
  "command": "list",
  "properties": {
    "name": "nameofwatcher",
  }
}
```

The response return the list asked. the mapping returned can either be ‘watchers’ or ‘pids’ depending the request.

### Command line

```
$ circusctl list [<name>]
```

### Subscribe to a watcher event

#### ZMQ

At any moment you can subscribe to a circus event. Circus provides a PUB/SUB feed on which any clients can subscribe. The subscriber endpoint URI is set in the circus.ini configuration file.

Events are pubsub topics:

- *watcher.<watchername>.reap*: when a process is reaped
- *watcher.<watchername>.spawn*: when a process is spawned
- *watcher.<watchername>.kill*: when a process is killed
- *watcher.<watchername>.updated*: when watcher configuration is updated
- *watcher.<watchername>.stop*: when a watcher is stopped
- *watcher.<watchername>.start*: when a watcher is started

All events messages are in a json struct.

#### Command line

The client has been updated to provide a simple way to listen on the events:

```
circusctl listen [<topic>, ...]
```

#### Example of result:

```
$ circusctl listen tcp://127.0.0.1:5556
watcher.refuge.spawn: {u'process_id': 6, u'process_pid': 72976,
                        u'time': 1331681080.985104}
watcher.refuge.spawn: {u'process_id': 7, u'process_pid': 72995,
                        u'time': 1331681086.208542}
watcher.refuge.spawn: {u'process_id': 8, u'process_pid': 73014,
                        u'time': 1331681091.427005}
```

### Get the list of sockets

#### ZMQ Message

To get the list of sockets:

```
{
  "command": "listsockets",
}
```

The response return a list of json mappings with keys for fd, name, host and port.

#### Command line

```
$ circusctl listsockets
```



## Get the number of processes

Get the number of processes in a watcher or in a arbiter

### ZMQ Message

```
{
  "command": "numprocesses",
  "properties": {
    "name": "<watchername>"
  }
}
```

The response return the number of processes in the ‘numprocesses’ property:

```
{ "status": "ok", "numprocesses": <n>, "time", "timestamp" }
```

If the property name isn’t specified, the sum of all processes managed is returned.

### Command line

```
$ circusctl numprocesses [<name>]
```

### Options

- <name>: name of the watcher

## Get the number of watchers

Get the number of watchers in a arbiter

### ZMQ Message

```
{
  "command": "numwatchers",
}
```

The response return the number of watchers in the ‘numwatchers’ property:

```
{ "status": "ok", "numwatchers": <n>, "time", "timestamp" }
```

### Command line

```
$ circusctl numwatchers
```

## Get the value of a watcher option

This command return the watchers options values asked.

## ZMQ Message

```
{
  "command": "options",
  "properties": {
    "name": "nameofwatcher",
  }
}
```

A message contains 1 property:

- name: name of watcher

The response return an object with a property “options” containing a dictionary of key/value returned by circus.

eg:

```
{
  "status": "ok",
  "options": {
    "flapping_window": 1,
    "flapping_attempts": 2,
    ...
  },
  'time': 1332202594.754644
}
```

## Command line

```
$ circusctl options <name>
```

## Options

- <name>: name of the watcher

Options Keys are:

- numprocesses: integer, number of processes
- warmup\_delay: integer or number, delay to wait between process spawning in seconds
- working\_dir: string, directory where the process will be executed
- uid: string or integer, user ID used to launch the process
- gid: string or integer, group ID used to launch the process
- send\_hup: boolean, if TRU the signal HUP will be used on reload
- shell: boolean, will run the command in the shell environment if true
- cmd: string, The command line used to launch the process
- env: object, define the environnement in which the process will be launch
- flapping\_attempts: integer, number of times we try to relaunch a process in the flapping\_window time before we stop the watcher during the retry\_in time.
- flapping\_window: integer or number, times in seconds in which we test the number of process restart.

- `retry_in`: integer or number, time in seconds we wait before we retry to launch the process if the maximum number of attempts has been reach.
- `max_retry`: integer, The maximum of retries loops
- `graceful_timeout`: integer or number, time we wait before we definitely kill a process.
- `priority`: used to sort watchers in the arbiter
- `singleton`: if True, a singleton watcher.
- `max_age`: time a process can live before being restarted
- `max_age_variance`: variable additional time to live, avoids stampeding herd.

### Quit the arbiter immediately

When the arbiter receive this command, the arbiter exit.

#### ZMQ Message

```
{  
  "command": "quit"  
}
```

The response return the status “ok”.

#### Command line

```
$ circusctl quit
```

### Reload the arbiter or a watcher

This command reload all the process in a watcher or all watchers. If a the option `send_hup` is set to true in a watcher then the HUP signal will be sent to the process. A graceful reload follow the following process:

1. Send a SIGQUIT signal to a process
2. Wait until graceful timeout
3. Send a SIGKILL signal to the process to make sure it is finally killed.

#### ZMQ Message

```
{  
  "command": "reload",  
  "propeties": {  
    "name": '<name>',  
    "graceful": true  
  }  
}
```

The response return the status “ok”. If the property `graceful` is set to true the processes will be exited gracefully.

If the property `name` is present, then the reload will be applied to the watcher.

### Command line

```
$ circusctl reload [<name>] [--terminate]
```

### Options

- <name>: name of the watcher
- --terminate; quit the node immediately

### Reload the configuration file

This command reloads the configuration file, so changes in the configuration file will be reflected in the configuration of circus.

### ZMQ Message

```
{  
  "command": "reloadconfig",  
}
```

The response return the status “ok”. If the property graceful is set to true the processes will be exited gracefully.

### Command line

```
$ circusctl reloadconfig
```

### Restart the arbiter or a watcher

This command restart all the process in a watcher or all watchers. This function simply stop a watcher then restart it.

### ZMQ Message

```
{  
  "command": "restart",  
  "properties": {  
    "name": '<name>'  
  }  
}
```

The response return the status “ok”.

If the property name is present, then the reload will be applied to the watcher.

### Command line

```
$ circusctl restart [<name>] [--terminate]
```

### Options

- <name>: name of the watcher
- `-terminate`; quit the node immediately

### Remove a watcher

This command remove a watcher dynamically from the arbiter. The watchers are gracefully stopped.

#### ZMQ Message

```
{
  "command": "rm",
  "properties": {
    "name": "<nameofwatcher>",
  }
}
```

A message contains 1 property:

- name: name of watcher

The response return a status “ok”.

### Command line

```
$ circusctl rm <name>
```

### Options

- <name>: name of the watcher to remove

### Set a watcher option

#### ZMQ Message

```
{
  "command": "set",
  "properties": {
    "name": "nameofwatcher",
    "options": {
      "key1": "val1",
    }
  }
}
```

The response return the status “ok”. See the command Options for a list of key to set.

## Command line

```
$ circusctl set <name> <key1> <value1> <key2> <value2>
```

## Send a signal

This command allows you to send a signal to all processes in a watcher, a specific process in a watcher or its children.

### ZMQ Message

To send a signal to all the processes for a watcher:

```
{
  "command": "signal",
  "property": {
    "name": <name>,
    "signal": <signal>
  }
}
```

To send a signal to a process:

```
{
  "command": "signal",
  "property": {
    "name": <name>,
    "pid": <processid>,
    "signal": <signal>
  }
}
```

An optional property “children” can be used to send the signal to all the children rather than the process itself:

```
{
  "command": "signal",
  "property": {
    "name": <name>,
    "pid": <processid>,
    "signal": <signal>,
    "children": True
  }
}
```

To send a signal to a process child:

```
{
  "command": "signal",
  "property": {
    "name": <name>,
    "pid": <processid>,
    "signal": <signal>,
    "child_pid": <childpid>,
  }
}
```

It is also possible to send a signal to all the processes of the watcher and its childs:

```
{
  "command": "signal",
  "property": {
    "name": <name>,
  }
}
```

```
    "signum": <signum>,
    "children": True
}
```

Last, you can send a signal to the process *and* its children, with the *recursive* option:

```
{
  "command": "signal",
  "property": {
    "name": <name>,
    "signum": <signum>,
    "recursive": True
  }
}
```

### Command line

```
$ circusctl signal <name> [<process>] [<pid>] [--children]
    [recursive] <signum>
```

#### Options:

- <name>: the name of the watcher
- <pid>: integer, the process id.
- <signum>: the signal number to send.
- <childpid>: the pid of a child, if any
- <children>: boolean, send the signal to all the children
- <recursive>: boolean, send the signal to the process and its children

Allowed signals are:

- 3: QUIT
- 15: TERM
- 9: KILL
- 1: HUP
- 21: TTIN
- 22: TTOU
- 30: USR1
- 31: USR2

### Start the arbiter or a watcher

This command starts all the processes in a watcher or all watchers.

### ZMQ Message

```
{
  "command": "start",
  "properties": {
    "name": '<name>',
  }
}
```

The response return the status “ok”.

If the property name is present, the watcher will be started.

### Command line

```
$ circusctl start [<name>]
```

### Options

- <name>: name of the watcher

### Get process infos

You can get at any time some statistics about your processes with the stat command.

### ZMQ Message

To get stats for all watchers:

```
{
  "command": "stats"
}
```

To get stats for a watcher:

```
{
  "command": "stats",
  "properties": {
    "name": <name>
  }
}
```

To get stats for a process:

```
{
  "command": "stats",
  "properties": {
    "name": <name>,
    "process": <processid>
  }
}
```

The response return an object per process with the property “info” containing some process informations:



```
{
  "info": {
    "children": [],
    "cmdline": "python",
    "cpu": 0.1,
    "ctime": "0:00.41",
    "mem": 0.1,
    "mem_info1": "3M",
    "mem_info2": "2G",
    "nice": 0,
    "pid": 47864,
    "username": "root"
  },
  "process": 5,
  "status": "ok",
  "time": 1332265655.897085
}
```

### Command Line

```
$ circusctl stats [<watchername>] [<processid>]
```

### Get the status of a watcher or all watchers

This command start get the status of a watcher or all watchers.

### ZMQ Message

```
{
  "command": "status",
  "properties": {
    "name": '<name>',
  }
}
```

The response return the status “active” or “stopped” or the status / watchers.

### Command line

```
$ circusctl status [<name>]
```

### Options

- <name>: name of the watcher

### Example

```
$ circusctl status dummy
active
$ circusctl status
dummy: active
```

```
dummy2: active
refuge: active
```

## Stop the arbiter or a watcher

This command stop all the process in a watcher or all watchers.

### ZMQ Message

```
{
  "command": "stop",
  "properties": {
    "name": '<name>',
  }
}
```

The response return the status “ok”.

If the property name is present, then the stop will be applied to the watcher.

### Command line

```
$ circusctl stop [<name>]
```

### Options

- <name>: name of the watcher

## 5.7 The Web Console

Circus comes with a Web Console that can be used to manage the system.

The Web Console lets you:

- Connect to any running Circus system
- Watch the processes CPU and Memory usage in real-time
- Add or kill processes
- Add new watchers

---

**Note:** The real-time CPU & Memory usage feature uses the stats socket. If you want to activate it, make sure the Circus system you’ll connect to has the stats endpoint enabled in its configuration:

```
[circus]
statsd = True
```

By default, this option is not activated.

---

The web console is its own package, you need to install:

```
$ pip install circus-web
```

To enable the console, add a few options in the Circus ini file:

```
[circus]
httpd = True
httpd_host = localhost
httpd_port = 8080
```

*httpd\_host* and *httpd\_port* are optional, and default to *localhost* and *8080*.

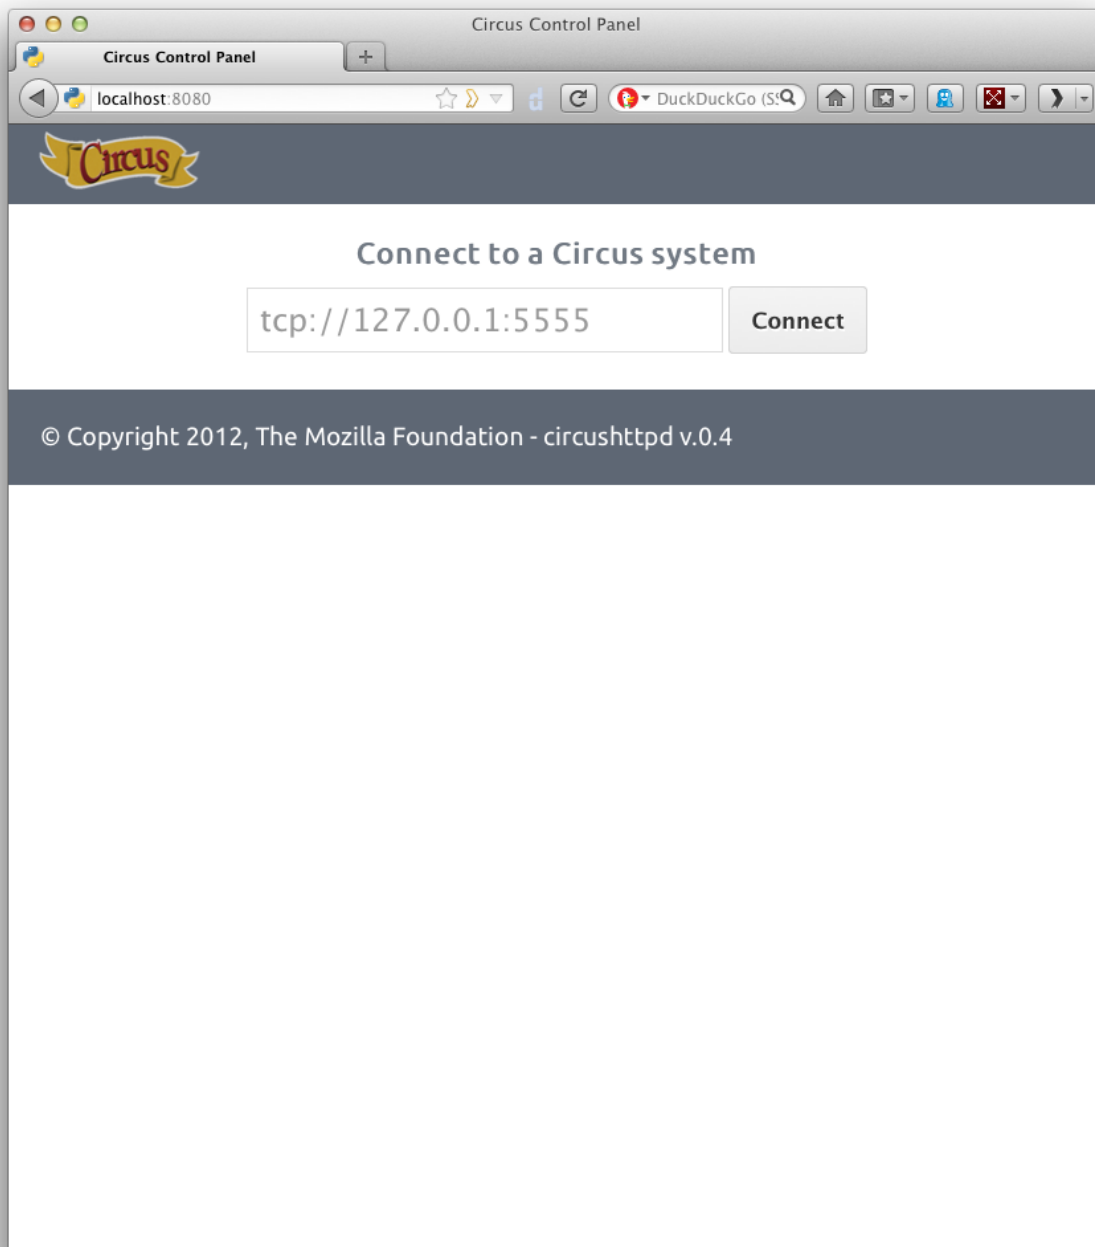
If you want to run the web app on its own, just run the **circushttpd** script:

```
$ circushttpd
Bottle server starting up...
Listening on http://localhost:8080/
Hit Ctrl-C to quit.
```

By default the script will run the Web Console on port 8080, but the `-port` option can be used to change it.

### 5.7.1 Using the console


Once the script is running, you can open a browser and visit *http://localhost:8080*. You should get this screen:



The Web Console is ready to be connected to a Circus system, given its **endpoint**. By default the endpoint is *tcp://127.0.0.1:5555*.

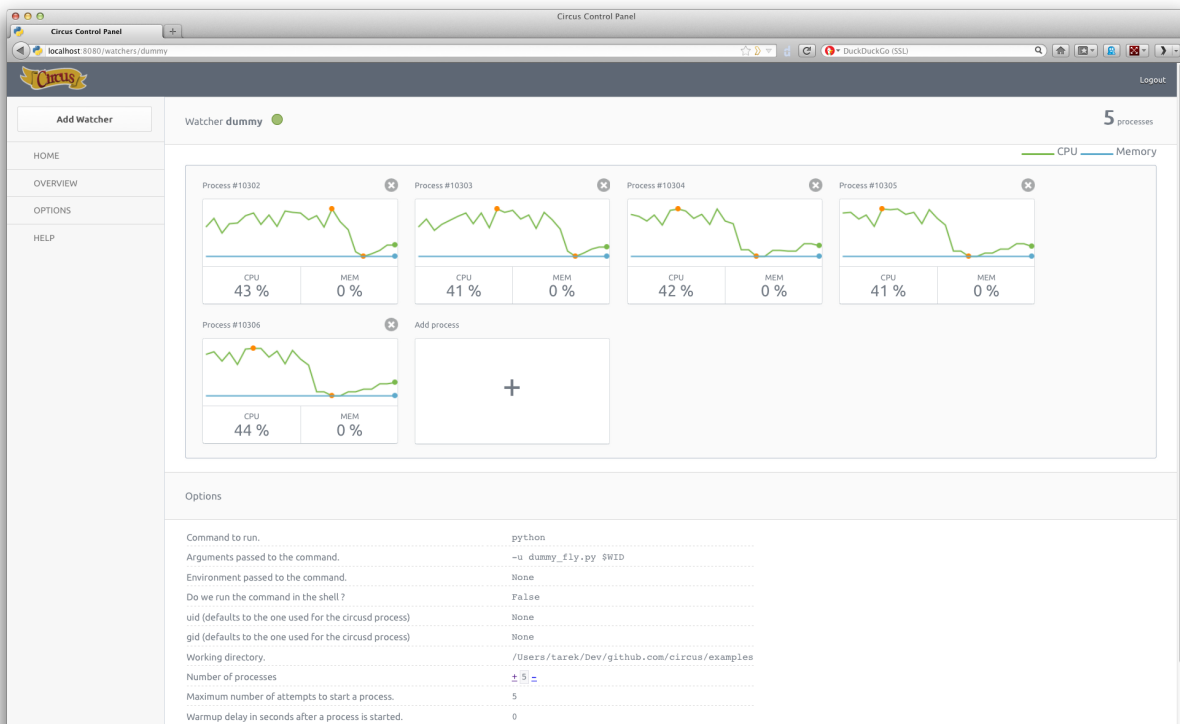
Once you hit *Connect*, the web application will connect to the Circus system.

With the Web Console logged in, you should get a list of watchers, and a real-time status of the two Circus processes (circusd and circusd-stats).



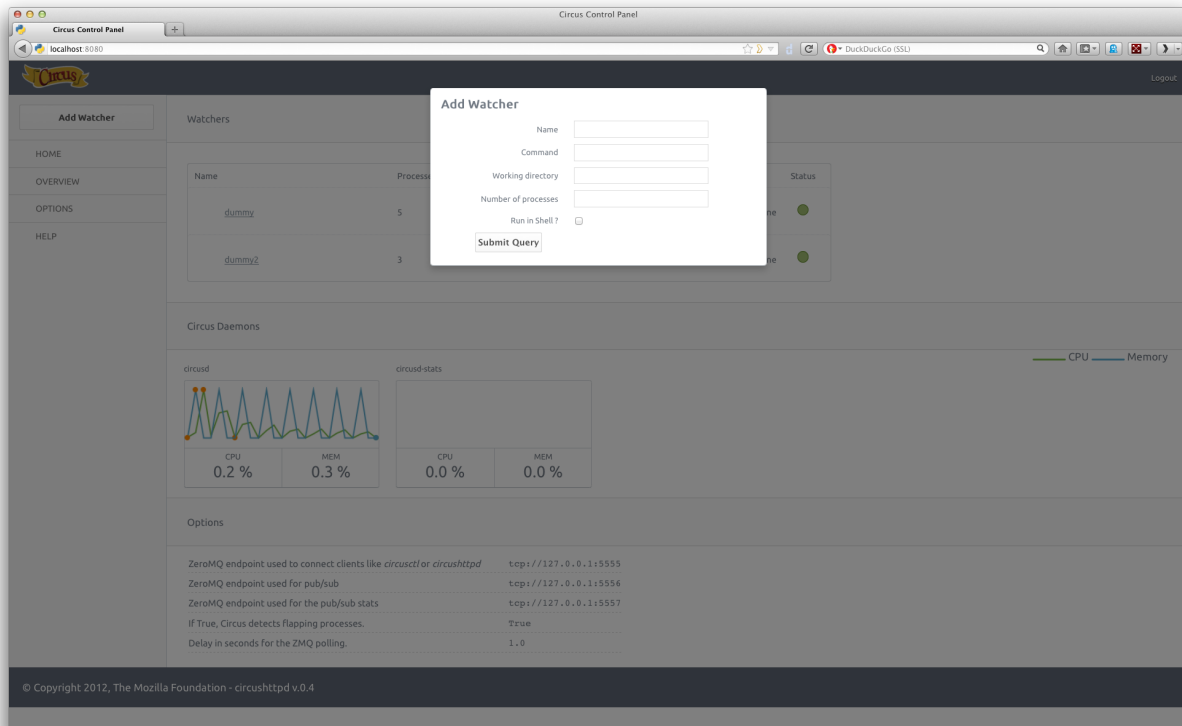
You can click on the status of each watcher to toggle it from **Active** (green) to **Inactive** (red). This change is effective immediately and let you start & stop watchers.

If you click on the watcher name, you will get a web page for that particular watcher, with its processes:



On this screen, you can add or remove processes, and kill existing ones.

Last but not least, you can add a brand new watcher by clicking on the *Add Watcher* link in the left menu:



### 5.7.2 Running behind Nginx and Varnish

Nginx can act as a proxy in front of Circus. It can also deal with security.

To hook Nginx, you define a *location* directive that proxies the calls to Circus.

Example:

```
location ~/media/*(.jpg|.css|.js)$ {
    alias /path/to/circus/web/;
}

location / {
    proxy_set_header X-Forwarded-For $proxy_add_x_forwarded_for;
    proxy_set_header Host $http_host;
    proxy_redirect off;
    proxy_pass http://127.0.0.1:8080;
}
```

If you want more configuration options, see <http://wiki.nginx.org/HttpProxyModule>.

Websockets in Nginx (v1.2.5) is currently unsupported, although it will be implemented in 1.3. To receive real-time statuses and graphs in the web console, you need to use a websocket-compatible proxy like Varnish or HAProxy. In Varnish, two backends can be defined: one for serving the web console and one for the handling the socket connections.

Example:

```

backend default {
    .host = "127.0.0.1";
    .port = "8001";
}

backend socket {
    .host = "127.0.0.1";
    .port = "8080";
    .connect_timeout = 1s;
    .first_byte_timeout = 2s;
    .between_bytes_timeout = 60s;
}

sub vcl_pipe {
    if (req.http.upgrade) {
        set bereq.http.upgrade = req.http.upgrade;
    }
}

sub vcl_recv {
    if (req.http.Upgrade ~ "(?i)websocket") {
        set req.backend = socket;
        return (pipe);
    }
}

```

Here, web console requests are bound to port 8001, and Nginx should be configured to listen on that port. Websocket connections are upgraded and piped directly to the circushttpd process listening on port 8080.

### 5.7.3 Running behind Nginx >= 1.3.13

As of [Nginx>=1.3.13](#) websockets are supported by the web server. With Nginx>=1.3.13 there is no longer a need to reroute websocket traffic via Varnish or HAProxy.

On Ubuntu you can install Nginx>=1.3.13 from Chris Lea's development branch [PPA](#), as so:

```

sudo apt-get install python-software-properties
sudo add-apt-repository ppa:chris-lea/nginx-devel
sudo apt-get update
sudo apt-get install nginx
nginx -v

```

An example Nginx config with websocket support:

```

# /etc/nginx/sites-enabled/default

upstream circusweb_server {
    server localhost:8080;
}

server {
    listen 80;
    server_name _;

    location / {
        proxy_pass http://circusweb_server;
        proxy_http_version 1.1;
        proxy_set_header Upgrade $http_upgrade;
    }
}

```

```
proxy_set_header Connection "upgrade";
proxy_set_header Host $http_host;
proxy_set_header X-Forwarded-Host $http_host;
proxy_set_header X-Real-IP $remote_addr;
proxy_set_header X-Forwarded-For $proxy_add_x_forwarded_for;
proxy_set_header X-Forwarded-Proto $scheme;
proxy_redirect off;
}

location ~/media/\*(.png|.jpg|.css|.js|.ico)$ {
    alias /path_to_site-packages/circusweb/media/;
}
}
```

### 5.7.4 Password-protect circushttpd

As explained in the [Security](#) page, running *circushttpd* is pretty unsafe. We don't provide any security in Circus itself, but you can protect your console at the NGinx level, by using <http://wiki.nginx.org/HttpAuthBasicModule>

Example:

```
location / {
    proxy_set_header X-Forwarded-For $proxy_add_x_forwarded_for;
    proxy_set_header Host $http_host;
    proxy_set_header X-Forwarded-Host: $http_host;
    proxy_set_header X-Forwarded-Proto: $scheme;
    proxy_redirect off;
    proxy_pass http://127.0.0.1:8080;
    auth_basic "Restricted";
    auth_basic_user_file /path/to/htpasswd;
}
```

The **htpasswd** file contains users and their passwords, and a password prompt will pop when you access the console.

You can use Apache's `htpasswd` script to edit it, or the Python script they provide at: <http://trac.edgewall.org/browser/trunk/contrib/htpasswd.py>

However, there's no native support for the combined use of HTTP Authentication and WebSockets (the server will throw HTTP 401 error codes). A workaround is to disable such authentication for the socket.io server.

Example (needs to be added before the previous rule):

```
location /socket.io {
    proxy_set_header X-Forwarded-For $proxy_add_x_forwarded_for;
    proxy_set_header Host $http_host;
    proxy_set_header X-Forwarded-Host: $http_host;
    proxy_set_header X-Forwarded-Proto: $scheme;
    proxy_redirect off;
    proxy_pass http://127.0.0.1:8080;
}
```

Of course that's just one way to protect your web console, you could use many other techniques.

### 5.7.5 Extending the web console

We picked *bottle* to build the webconsole, mainly because it's a really tiny framework that doesn't do much. By having a look at the code of the web console, you'll eventually find out that it's really simple to understand.



Here is how it's split:

- The *circushttpd.py* file contains the “views” definitions and some code to handle the socket connection (via *socketio*).
- the *controller.py* contains a single class which is in charge of doing the communication with the circus controller. It allows to have a nicer high level API when defining the web server.

If you want to add a feature in the web console you can reuse the code that's existing. A few tools are at your disposal to ease the process:

- There is a *render\_template* function, which takes the named arguments you pass to it and pass them to the template renderer and return the resulting HTML. It also passes some additional variables, such as the session, the circus version and the client if defined.
- If you want to run commands and do a redirection depending the result of it, you can use the *run\_command* function, which takes a callable as a first argument, a message in case of success and a redirection url.

The *StatsNamespace* class is responsible for managing the websocket communication on the server side. Its documentation should help you to understand what it does.

## 5.8 Circus Sockets

Circus can bind network sockets and manage them as it does for processes.

The main idea is that a child process that's created by Circus to run one of the watcher's command can inherit from all the opened file descriptors.

That's how Apache or Unicorn works, and many other tools out there.

### 5.8.1 Goal

The goal of having sockets managed by Circus is to be able to manage network applications in Circus exactly like other applications.

For example, if you use Circus with *Chaussette* – a WSGI server, you can get a very fast web server running and manage “*Web Workers*” in Circus as you would do for any other process.

Splitting the socket management from the network application itself offers a lot of opportunities to scale and manage your stack.

### 5.8.2 Design

The gist of the feature is done by binding the socket and start listening to it in **circusd**:

```
import socket
```

```
sock = socket.socket(FAMILY, TYPE)
sock.bind((HOST, PORT))
sock.listen(BACKLOG)
fd = sock.fileno()
```

Circus then keeps track of all the opened fds, and let the processes it runs as children have access to them if they want.

If you create a small Python network script that you intend to run in Circus, it could look like this:

```
import socket
import sys

fd = int(sys.argv[1])    # getting the FD from circus
sock = socket.fromfd(fd, FAMILY, TYPE)

# dealing with one request at a time
while True:
    conn, addr = sock.accept()
    request = conn.recv(1024)
    .. do something ..
    conn.sendall(response)
    conn.close()
```

Then Circus could run like this:

```
[circus]
check_delay = 5
endpoint = tcp://127.0.0.1:5555
pubsub_endpoint = tcp://127.0.0.1:5556
stats_endpoint = tcp://127.0.0.1:5557

[watcher:dummy]
cmd = mycoolscript $(circus.sockets.foo)
use_sockets = True
warmup_delay = 0
numprocesses = 5

[socket:foo]
host = 127.0.0.1
port = 8888
```

*\$(circus.sockets.foo)* will be replaced by the FD value once the socket is created and bound on the 8888 *port*.

---

**Note:** Starting at Circus 0.8 there's an alternate syntax to avoid some conflicts with some config parsers. You can write:

```
*((circus.sockets.foo))*
```

---

### 5.8.3 Real-world example

**Chaussette** is the perfect Circus companion if you want to run your WSGI application.

Once it's installed, running 5 **meinheld** workers can be done by creating a socket and calling the **chaussette** command in a worker, like this:

```
[circus]
endpoint = tcp://127.0.0.1:5555
pubsub_endpoint = tcp://127.0.0.1:5556
stats_endpoint = tcp://127.0.0.1:5557

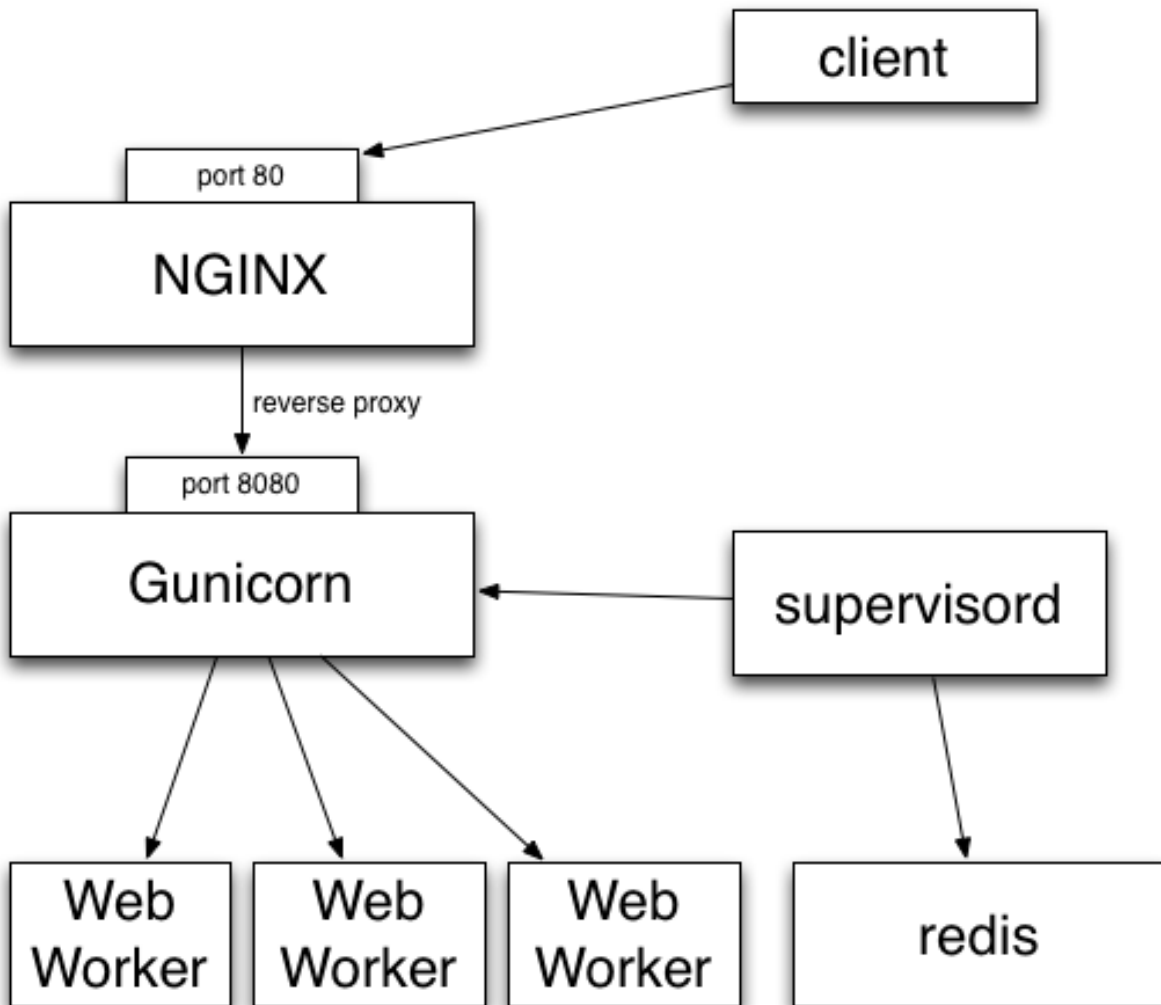
[watcher:web]
cmd = chaussette --fd $(circus.sockets.web) --backend meinheld mycool.app
use_sockets = True
numprocesses = 5
```

```
[socket:web]
host = 0.0.0.0
port = 8000
```

We did not publish benchmarks yet, but a Web cluster managed by Circus with a Gevent or Meinheld backend is as fast as any pre-fork WSGI server out there.

### 5.8.4 Circus stack v.s. Classical stack

In a classical WSGI stack, you have a server like Gunicorn that serves on a port or an unix socket and is usually deployed behind a web server like Nginx:



Clients call Nginx, which reverse proxies all the calls to Gunicorn.

If you want to make sure the Gunicorn process stays up and running, you have to use a program like Supervisord or upstart.

Gunicorn in turn watches for its processes (“workers”).

In other words you are using two levels of process management. One that you manage and control (supervisord), and a second one that you have to manage in a different UI, with a different philosophy and less control over what’s going

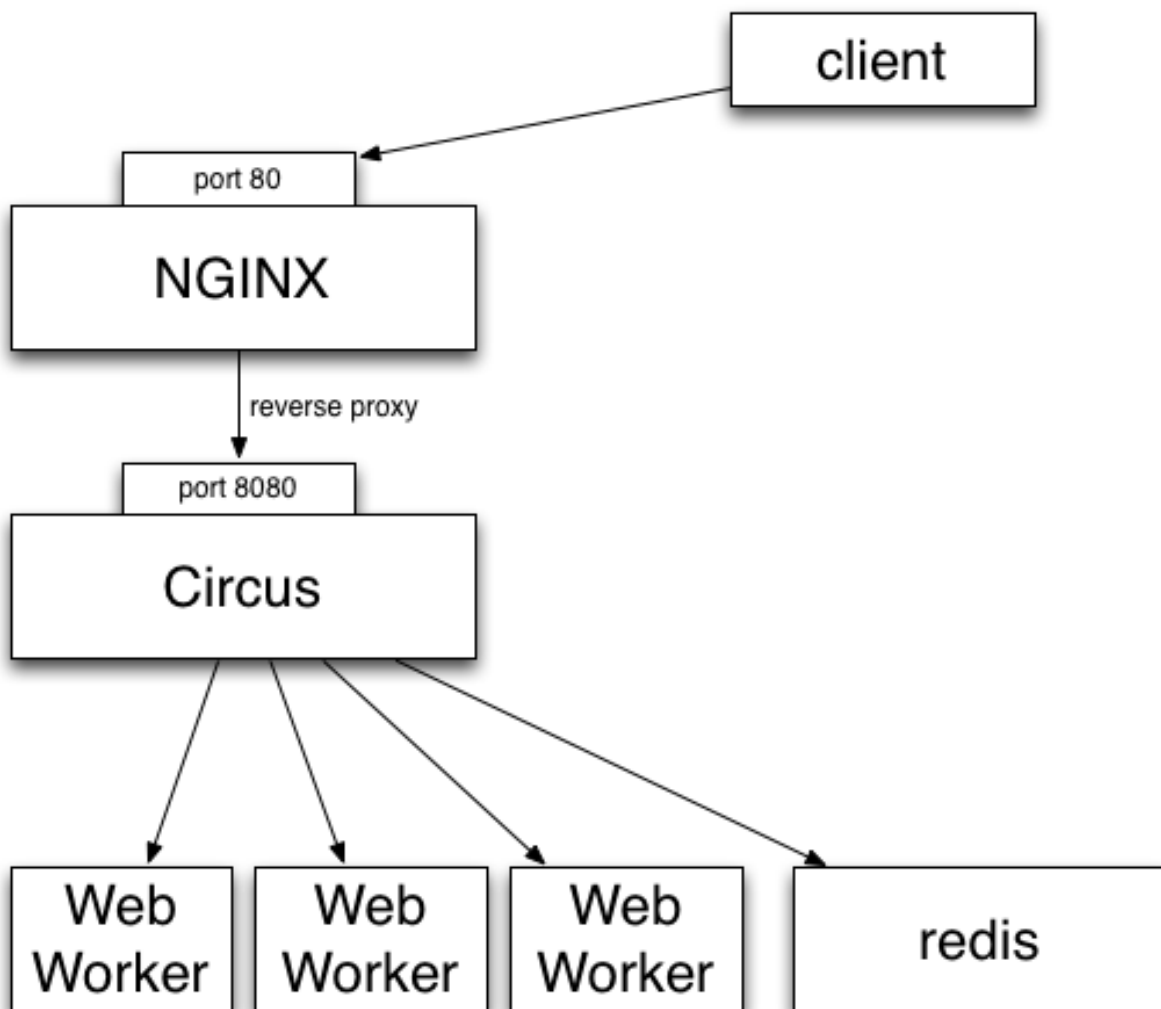
on (the wsgi server's one)

This is true for Gunicorn and most multi-processes WSGI servers out there I know about. uWsgi is a bit different as it offers plethoras of options.

But if you want to add a Redis server in your stack, you *will* end up with managing your stack processes in two different places.

Circus' approach on this is to manage processes *and* sockets.

A Circus stack can look like this:



So, like Gunicorn, Circus is able to bind a socket that will be proxied by Nginx. Circus don't deal with the requests but simply binds the socket. It's then up to a web worker process to accept connections on the socket and do the work.

It provides equivalent features than Supervisor but will also let you manage all processes at the same level, whether they are web workers or Redis or whatever. Adding a new web worker is done exactly like adding a new Redis process.

## Benchches

We did a few benches to compare Circus & Chaussette with Gunicorn. To summarize, Circus is not adding any overhead and you can pick up many different backends for your web workers.

See:

- <http://blog.ziade.org/2012/06/28/wsgi-web-servers-bench>
- <http://blog.ziade.org/2012/07/03/wsgi-web-servers-bench-part-2>

## 5.9 Circus Library

The Circus package is composed of a high-level `get_arbiter()` function and many classes. In most cases, using the high-level function should be enough, as it creates everything that is needed for Circus to run.

You can subclass Circus' classes if you need more granularity than what is offered by the configuration.

### 5.9.1 The `get_arbiter` function

`get_arbiter()` is just a convenience on top of the various circus classes. It creates a *arbiter* (class `Arbiter`) instance with the provided options, which in turn runs a single `Watcher` with a single `Process`.

```
circus.get_arbiter()
```

Example:

```
from circus import get_arbiter

arbiter = get_arbiter({"cmd": "myprogram", "numprocesses": 3})
try:
    arbiter.start()
finally:
    arbiter.stop()
```

### 5.9.2 The classes collection

Circus provides a series of classes you can use to implement your own process manager:

- `Process`: wraps a running process and provides a few helpers on top of it.
- `Watcher`: run several instances of `Process` against the same command. Manage the death and life of processes.
- `Arbiter`: manages several `Watcher`.

```
class circus.process.Process(wid, cmd, args=None, working_dir=None, shell=False,
                             uid=None, gid=None, env=None, rlimits=None, executable=None, use_fds=False, watcher=None, spawn=True,
                             pipe_stdout=True, pipe_stderr=True, close_child_stdout=False,
                             close_child_stderr=False)
```

Wraps a process.

Options:

- **wid**: the process unique identifier. This value will be used to replace the `$WID` string in the command line if present.
- **cmd**: the command to run. May contain any of the variables available that are being passed to this class. They will be replaced using the python format syntax.
- **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`.

- executable**: When executable is given, the first item in the args sequence obtained from **cmd** is still treated by most programs as the command name, which can then be different from the actual executable name. It becomes the display name for the executing program in utilities such as **ps**.
- working\_dir**: the working directory to run the command in. If not provided, will default to the current working directory.
- shell**: if *True*, will run the command in the shell environment. *False* by default. **warning: this is a security hazard.**
- uid**: if given, is the user id or name the command should run with. The current uid is the default.
- gid**: if given, is the group id or name the command should run with. The current gid is the default.
- env**: a mapping containing the environment variables the command will run with. Optional.
- rlimits**: a mapping containing rlimit names and values that will be set before the command runs.
- use\_fds**: if *True*, will not close the fds in the subprocess. default: *False*.
- pipe\_stdout**: if *True*, will open a PIPE on stdout. default: *True*.
- pipe\_stderr**: if *True*, will open a PIPE on stderr. default: *True*.
- close\_child\_stdout**: If *True*, redirects the child process' stdout to /dev/null after the fork. default: *False*.
- close\_child\_stderr**: If *True*, redirects the child process' stdout to /dev/null after the fork. default: *False*.

**pid**

Return the *pid*

**stdout**

Return the *stdout* stream

**stderr**

Return the *stdout* stream

**send\_signal** (\*args, \*\*kw)

Sends a signal **sig** to the process.

**stop** (\*args, \*\*kw)

Terminate the process.

**age** ()

Return the age of the process in seconds.

**info** ()

Return process info.

The info returned is a mapping with these keys:

- mem\_info1**: Resident Set Size Memory in bytes (RSS)
- mem\_info2**: Virtual Memory Size in bytes (VMS).
- cpu**: % of cpu usage.
- mem**: % of memory usage.
- ctime**: process CPU (user + system) time in seconds.
- pid**: process id.
- username**: user name that owns the process.
- nice**: process niceness (between -20 and 20)
- cmdline**: the command line the process was run with.

**children()**  
Return a list of children pids.

**is\_child(pid)**  
Return True if the given *pid* is a child of that process.

**send\_signal\_child(\*args, \*\*kw)**  
Send signal *signum* to child *pid*.

**send\_signal\_children(\*args, \*\*kw)**  
Send signal *signum* to all children.

**status**  
Return the process status as a constant

- RUNNING
- DEAD\_OR\_ZOMBIE
- UNEXISTING
- OTHER

Example:

```
>>> from circus.process import Process
>>> process = Process('Top', 'top', shell=True)
>>> process.age()
3.0107998847961426
>>> process.info()
'Top: 6812 N/A tarek Zombie N/A N/A N/A N/A N/A'
>>> process.status
1
>>> process.stop()
>>> process.status
2
>>> process.info()
'No such process (stopped?)'
```

**class circus.watcher.Watcher**(*name, cmd, args=None, numprocesses=1, warmup\_delay=0.0, working\_dir=None, shell=False, uid=None, max\_retry=5, gid=None, send\_hup=False, env=None, stopped=True, graceful\_timeout=30.0, preruleload\_fn=None, rlimits=None, executable=None, stdout\_stream=None, stderr\_stream=None, priority=0, loop=None, singleton=False, use\_sockets=False, copy\_env=False, copy\_path=False, max\_age=0, max\_age\_variance=30, hooks=None, respawn=True, autostart=True, on\_demand=False, virtualenv=None, close\_child\_stdout=False, close\_child\_stderr=False, \*\*options*)

Class managing a list of processes for a given command.

Options:

- name**: name given to the watcher. Used to uniquely identify it.
- cmd**: the command to run. May contain *\$WID*, which will be replaced by **wid**.
- 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`.
- numprocesses**: Number of processes to run.
- working\_dir**: the working directory to run the command in. If not provided, will default to the current working directory.

- shell**: if *True*, will run the command in the shell environment. *False* by default. **warning: this is a security hazard.**
- uid**: if given, is the user id or name the command should run with. The current uid is the default.
- gid**: if given, is the group id or name the command should run with. The current gid is the default.
- send\_hup**: if *True*, a process reload will be done by sending the SIGHUP signal. Defaults to *False*.
- env**: a mapping containing the environment variables the command will run with. Optional.
- rlimits**: a mapping containing rlimit names and values that will be set before the command runs.
- stdout\_stream**: a mapping that defines the stream for the process stdout. Defaults to *None*.

Optional. When provided, *stdout\_stream* is a mapping containing up to three keys:

- class**: the stream class. Defaults to *circus.stream.FileStream*
- filename**: the filename, if using a *FileStream*
- refresh\_time**: the delay between two stream checks. Defaults to 0.3 seconds.
- max\_bytes**: maximum file size, after which a new output file is opened. defaults to 0 which means no maximum size.
- backup\_count**: how many backups to retain when rotating files according to the *max\_bytes* parameter. defaults to 0 which means no backups are made.

This mapping will be used to create a stream callable of the specified class. Each entry received by the callable is a mapping containing:

- pid** - the process pid
- name** - the stream name (*stderr* or *stdout*)
- data** - the data

- stderr\_stream**: a mapping that defines the stream for the process stderr. Defaults to *None*.

Optional. When provided, *stderr\_stream* is a mapping containing up to three keys: - **class**: the stream class. Defaults to *circus.stream.FileStream* - **filename**: the filename, if using a *FileStream* - **refresh\_time**: the delay between two stream checks. Defaults

to 0.3 seconds.

- max\_bytes**: maximum file size, after which a new output file is opened. defaults to 0 which means no maximum size.
- backup\_count**: how many backups to retain when rotating files according to the *max\_bytes* parameter. defaults to 0 which means no backups are made.

This mapping will be used to create a stream callable of the specified class.

Each entry received by the callable is a mapping containing:

- pid** - the process pid
- name** - the stream name (*stderr* or *stdout*)
- data** - the data

- priority** – integer that defines a priority for the watcher. When the Arbiter do some operations on all watchers, it will sort them with this field, from the bigger number to the smallest. (default: 0)
- singleton** – If *True*, this watcher has a single process. (default: *False*)



- use\_sockets** – If True, the processes will inherit the file descriptors, thus can reuse the sockets opened by circusd. (default: False)
- on\_demand** – If True, the processes will be started only at the first connection to the socket (default: False)
- copy\_env** – If True, the environment in which circus is running run will be reproduced for the workers. (default: False)
- copy\_path** – If True, circusd *sys.path* is sent to the process through *PYTHONPATH*. You must activate **copy\_env** for **copy\_path** to work. (default: False)
- max\_age**: If set after around *max\_age* seconds, the process is replaced with a new one. (default: 0, Disabled)
- max\_age\_variance**: The maximum number of seconds that can be added to *max\_age*. This extra value is to avoid restarting all processes at the same time. A process will live between *max\_age* and *max\_age* + *max\_age\_variance* seconds.
- hooks**: callback functions for hooking into the watcher startup and shutdown process. **hooks** is a dict where each key is the hook name and each value is a 2-tuple with the name of the callable or the callable itself and a boolean flag indicating if an exception occurring in the hook should not be ignored. Possible values for the hook name: *before\_start*, *after\_start*, *before\_spawn*, *before\_stop*, *after\_stop*.
- options** – extra options for the worker. All options found in the configuration file for instance, are passed in this mapping – this can be used by plugins for watcher-specific options.
- respawn** – If set to False, the processes handled by a watcher will not be respawned automatically. (default: True)
- virtualenv** – The root directory of a virtualenv. If provided, the watcher will load the environment for its execution. (default: None)
- close\_child\_stdout**: If True, closes the stdout after the fork. default: False.
- close\_child\_stderr**: If True, closes the stderr after the fork. default: False.

**notify\_event** (*topic, msg*)

Publish a message on the event publisher channel

**reap\_processes** (*\*args, \*\*kw*)

Reap all the processes for this watcher.

**manage\_processes** (*\*args, \*\*kw*)

Manage processes.

**reap\_and\_manage\_processes** (*\*args, \*\*kw*)

Reap & manage processes.

**spawn\_processes** (*\*args, \*\*kw*)

Spawn processes.

**spawn\_process** ()

Spawn process.

**kill\_process** (*process, sig=15*)

Kill process.

**kill\_processes** (*\*args, \*\*kw*)

Kill all the processes of this watcher.

**send\_signal\_child** (*\*args, \*\*kw*)

Send signal to a child.

**stop** (\*args, \*\*kw)

Stop.

**start** (\*args, \*\*kw)

Start.

**restart** (\*args, \*\*kw)

Restart.

**reload** (\*args, \*\*kw)

```
class circus.arbiter.Arbiter(watchers, endpoint, pubsub_endpoint, check_delay=0.5, pre-
                             reload_fn=None, context=None, loop=None, statsd=False,
                             stats_endpoint=None, statsd_close_outputs=False, mul-
                             ticast_endpoint=None, plugins=None, sockets=None,
                             warmup_delay=0, httpd=False, httpd_host='localhost',
                             httpd_port=8080, httpd_close_outputs=False, debug=False,
                             ssh_server=None, proc_name='circusd', pidfile=None,
                             loglevel=None, logoutput=None, fqdn_prefix=None)
```

Class used to control a list of watchers.

Options:

- **watchers** – a list of `Watcher` objects
- **endpoint** – the controller ZMQ endpoint
- **pubsub\_endpoint** – the pubsub endpoint
- **statsd** – If True, a `circusd-stats` process is run (default: False)
- **stats\_endpoint** – the stats endpoint.
- **statsd\_close\_outputs** – if True sends the `circusd-stats` stdout/stderr to `/dev/null` (default: False)
- **multicast\_endpoint** – the multicast endpoint for `circusd` cluster auto-discovery (default: `udp://237.219.251.97:12027`) Multicast addr should be between `224.0.0.0` to `239.255.255.255` and the same for the all cluster.
- **check\_delay** – the delay between two controller points (default: 1 s)
- **prereload\_fn** – callable that will be executed on each reload (default: None)
- **context** – if provided, the `zmq` context to reuse. (default: None)
- **loop**: if provided, a `zmq.eventloop.ioloop.IOLoop` instance to reuse. (default: None)
- **plugins** – a list of plugins. Each item is a mapping with:
  - **use** – Fully qualified name that points to the plugin class
  - every other value is passed to the plugin in the **config** option
- **sockets** – a mapping of sockets. Each key is the socket name, and each value a `CircusSocket` class. (default: None)
- **warmup\_delay** – a delay in seconds between two watchers startup. (default: 0)
- **httpd** – If True, a `circushttpd` process is run (default: False)
- **httpd\_host** – the `circushttpd` host (default: `localhost`)
- **httpd\_port** – the `circushttpd` port (default: 8080)
- **httpd\_close\_outputs** – if True, sends `circushttpd` stdout/stderr to `/dev/null`. (default: False)
- **debug** – if True, adds a lot of debug info in the stdout (default: False)

- proc\_name** – the arbiter process name

- fqdn\_prefix** – a prefix for the unique identifier of the circus instance on the cluster.

**start** (\*args, \*\*kw)

Starts all the watchers.

The start command is an infinite loop that waits for any command from a client and that watches all the processes and restarts them if needed.

**reload** (\*args, \*\*kw)

Reloads everything.

Run the `prereload_fn()` callable if any, then gracefully reload all watchers.

**numprocesses** ()

Return the number of processes running across all watchers.

**numwatchers** ()

Return the number of watchers.

**get\_watcher** (name)

Return the watcher *name*.

**add\_watcher** (name, cmd, \*\*kw)

Adds a watcher.

Options:

- name**: name of the watcher to add

- cmd**: command to run.

- all other options defined in the `Watcher` constructor.

## 5.10 Deployment

Although the Circus daemon can be managed with the `circusd` command, it's easier to have it start on boot. If your system supports Upstart, you can create this Upstart script in `/etc/init/circus.conf`.

```
start on filesystem and net-device-up IFACE=lo

stop on shutdown

respawn
exec /usr/local/bin/circusd --log-output /var/log/circus.log \
                                --pidfile /var/run/circusd.pid \
                                /etc/circus.ini
```

This assumes that `circus.ini` is located at `/etc/circus.ini`. After rebooting, you can control `circusd` with the service command:

```
$ service circus start/stop/restart
```

### 5.10.1 Recipes

This section will contain recipes to deploy Circus. Until then you can look at Pete's [Puppet recipe](#) or at Remy's [Chef recipe](#)

## 5.11 The Plugin System

Circus comes with a plugin system which let you interact with **circusd**.

---

**Note:** We might add circusd-stats support to plugins later on

---

A Plugin is composed of two parts:

- a ZMQ subscriber to all events published by **circusd**
- a ZMQ client to send commands to **circusd**

Each plugin is run as a separate process under a custom watcher.

A few examples of some plugins you could create with this system:

- a notification system that sends e-mail alerts when a watcher is flapping
- a logger
- a tool that add or remove processes depending on the load
- etc.

Circus itself provides a few plugins:

- a statsd plugin, that sends to statsd all events emitted by circusd
- the flapping feature which avoid to re-launch processes infinitely when they die too quickly.
- many more to come !

### 5.11.1 The CircusPlugin class

Circus provides a base class to help you implement plugins: `circus.plugins.CircusPlugin`

```
class circus.plugins.CircusPlugin(endpoint, pubsub_endpoint, check_delay, ssh_server=None,  
                                **config)
```

Base class to write plugins.

Options:

- **context** – the ZMQ context to use
- **endpoint** – the circusd ZMQ endpoint
- **pubsub\_endpoint** – the circusd ZMQ pub/sub endpoint
- **check\_delay** – the configured check delay
- **config** – free config mapping

```
call (command, **props)
```

Sends to **circusd** the command.

Options:

- **command** – the command to call
- **props** – keywords argument to add to the call

Returns the JSON mapping sent back by **circusd**

**cast** (*command*, *\*\*props*)

Fire-and-forget a command to **circusd**

Options:

- command** – the command to call
- props** – keywords argument to add to the call

**handle\_recv** (*data*)

Receives every event published by **circusd**

Options:

- data** – a tuple containing the topic and the message.

**handle\_stop** ()

Called right before the plugin is stopped by Circus.

**handle\_init** ()

Called right before a plugin is started - in the thread context.

When initialized by Circus, this class creates its own event loop that receives all **circusd** events and pass them to `handle_recv()`. The data received is a tuple containing the topic and the data itself.

`handle_recv()` **must** be implemented by the plugin.

The `call()` and `cast()` methods can be used to interact with **circusd** if you are building a Plugin that actively interacts with the daemon.

`handle_init()` and `handle_stop()` are just convenience methods you can use to initialize and clean up your code. `handle_init()` is called within the thread that just started. `handle_stop()` is called in the main thread just before the thread is stopped and joined.

### 5.11.2 Writing a plugin

Let's write a plugin that logs in a file every event happening in **circusd**. It takes one argument which is the filename.

The plugin could look like this:

```
from circus.plugins import CircusPlugin

class Logger(CircusPlugin):

    name = 'logger'

    def __init__(self, filename, **kwargs):
        super(Logger, self).__init__(**kwargs)
        self.filename = filename
        self.file = None

    def handle_init(self):
        self.file = open(self.filename, 'a+')

    def handle_stop(self):
        self.file.close()

    def handle_recv(self, data):
        topic, msg = data
        self.file.write('%s::%s' % (topic, msg))
```

That's it ! This class can be saved in any package/module, as long as it can be seen by Python.

For example, `Logger` could be found in a *plugins* module in a *myproject* package.

### Async requests

In case you want to make any asynchronous operations (like a Tornado call or using `periodicCall`) make sure you are using the right loop. The loop you always want to be using is `self.loop` as it gets set up by the base class. The default loop often isn't the same and therefore code might not get executed as expected.

### 5.11.3 Trying a plugin

You can run a plugin through the command line with the **circus-plugin** command, by specifying the plugin fully qualified name:

```
$ circus-plugin --endpoint tcp://127.0.0.1:5555 --pubsub tcp://127.0.0.1:5556 myproject.plugins.Logger
[INFO] Loading the plugin...
[INFO] Endpoint: 'tcp://127.0.0.1:5555'
[INFO] Pub/sub: 'tcp://127.0.0.1:5556'
[INFO] Starting
```

Another way to run a plugin is to let Circus handle its initialization. This is done by adding a **[plugin:NAME]** section in the configuration file, where *NAME* is a unique name for your plugin:

```
[plugin:logger]
use = myproject.plugins.Logger
filename = /var/myproject/circus.log
```

**use** is mandatory and points to the fully qualified name of the plugin.

When Circus starts, it creates a watcher with one process that runs the pointed class, and pass any other variable contained in the section to the plugin constructor via the **config** mapping.

You can also programmatically add plugins when you create a `circus.arbiter.Arbiter` class or use `circus.get_arbiter()`, see *Circus Library*.

### 5.11.4 Performances

Since every plugin is loaded in its own process, it should not impact the overall performances of the system as long as the work done by the plugin is not doing too many calls to the **circusd** process.

## 5.12 Security

Circus is built on the top of the ZeroMQ library and comes with no security at all in its protocols. However, you can run a Circus system on a server and set up an SSH tunnel to access it from another machine.

This section explains what Circus does on your system when you run it, and ends up describing how to use an SSH tunnel.

You can also read <http://www.zeromq.org/area:faq#toc5>

### 5.12.1 TCP ports

By default, Circus opens the following TCP ports on the local host:

- **5555** – the port used to control circus via **circusctl**
- **5556** – the port used for the Publisher/Subscriber channel.
- **5557** – the port used for the statistics channel – if activated.
- **8080** – the port used by the Web UI – if activated.

These ports allow client apps to interact with your Circus system, and depending on how your infrastructure is organized, you may want to protect these ports via firewalls **or** configure Circus to run using **IPC** ports.

Here's an example of running Circus using only IPC entry points:

```
[circus]
check_delay = 5
endpoint = ipc:///var/circus/endpoint
pubsub_endpoint = ipc:///var/circus/pubsub
stats_endpoint = ipc:///var/circus/stats
```

When Configured using IPC, the commands must be run from the same box, but no one can access them from outside, unlike using TCP.

Of course, if you activate the Web UI, the **8080** port will still be open.

### 5.12.2 circushttd

When you run **circushttd** manually, or when you use the **httd** option in the ini file like this:

```
[circus]
check_delay = 5
endpoint = ipc:///var/circus/endpoint
pubsub_endpoint = ipc:///var/circus/pubsub
stats_endpoint = ipc:///var/circus/stats
httd = 1
```

The web application will run on port **8080** and will let anyone accessing the web page manage the **circusd** daemon.

That includes creating new watchers that can run any command on your system !

#### Do not make it publicly available

If you want to protect the access to the web panel, you can serve it behind Nginx or Apache or any proxy-capable web server, than can take care of the security.

### 5.12.3 User and Group Permissions

By default, all processes started with Circus will be running with the same user and group than **circusd**. Depending on the privileges the user has on the system, you may not have access to all the features Circus provides.

For instance, some statistics features on a running processes require extended privileges. Typically, if the CPU usage numbers you get using the **stats** command are *N/A*, it means your user can't access the proc files. This will be the case by default under Mac OS X.

You may run **circusd** as root to fix this, and set the **uid** and **gid** values for each watcher to get all the features.

But beware that running **circusd** as root exposes you to potential privilege escalation bugs. While we're doing our best to avoid any bugs, running as root and facing a bug that performs unwanted actions on your system may be dangerous.

The best way to prevent this is to make sure that the system running Circus is completely isolated (like a VM) **or** to run the whole system under a controlled user.

### 5.12.4 SSH tunneling

Clients can connect to a **circusd** instance by creating an SSH tunnel. To do so, pass the command line option **-ssh** followed by **user@address**, where **user** is the user on the remote server and **address** is the server's address as seen by the client. The SSH protocol will require credentials to complete the login.

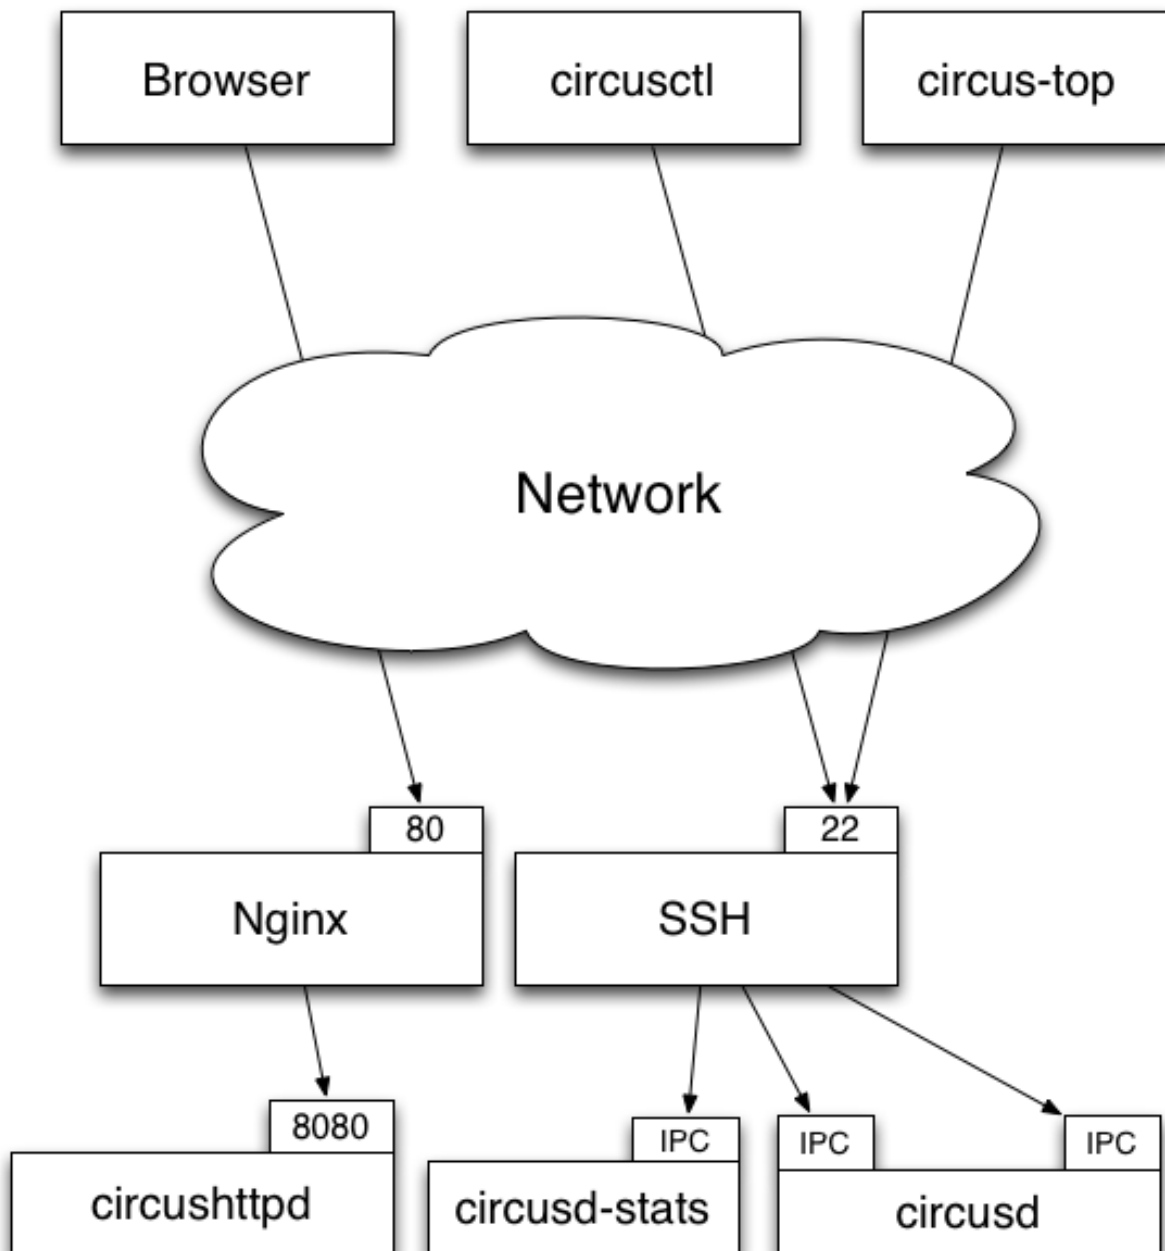
If **circusd** as seen by the SSH server is not at the default endpoint address **localhost:5555** then specify the **circusd** address using the option **-endpoint**

### 5.12.5 Secured setup example

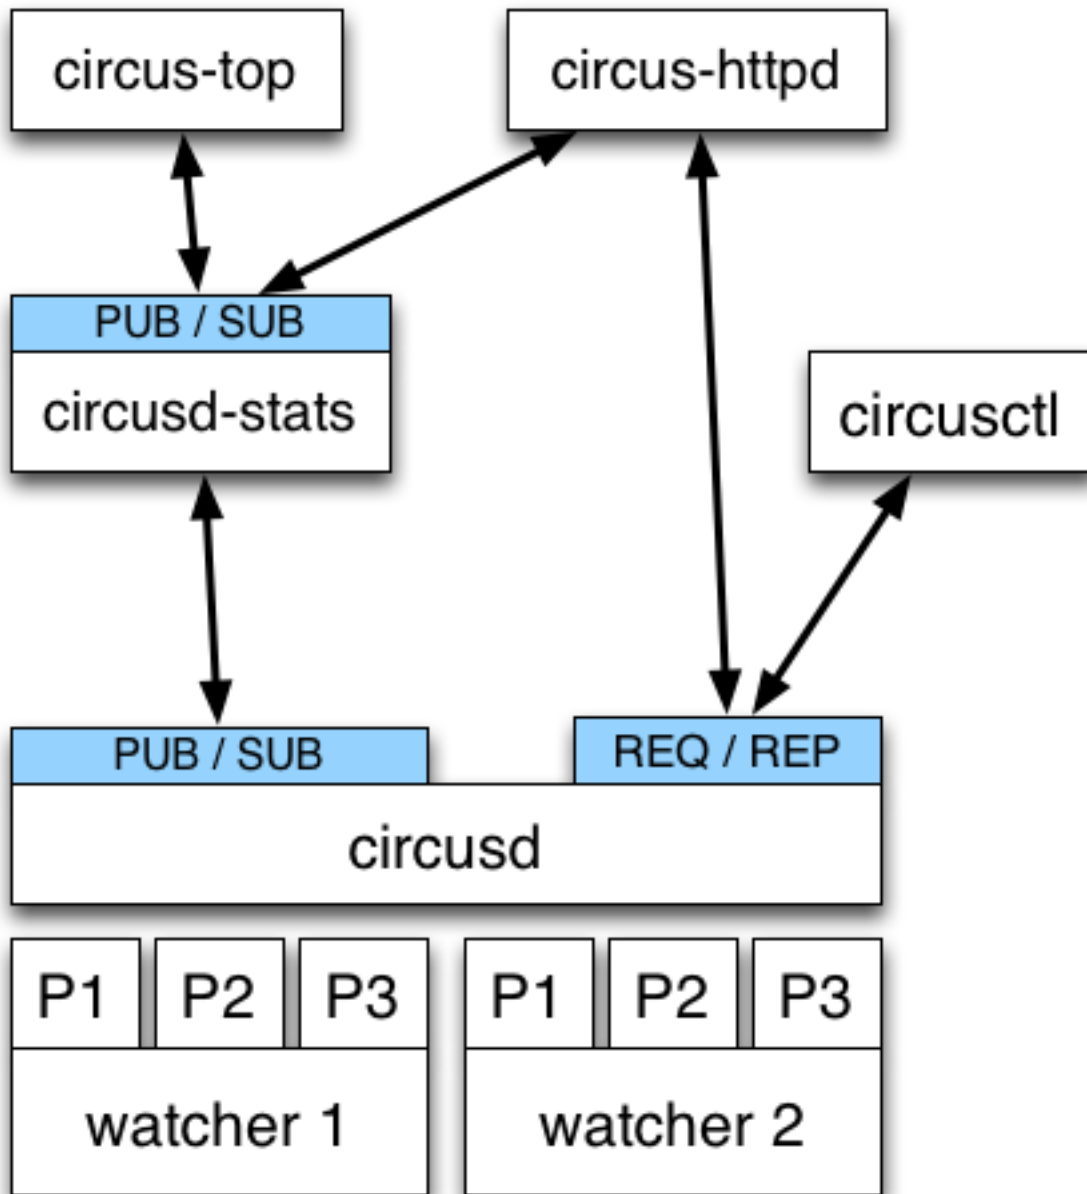
Setting up a secured Circus server can be done by:

- Running an SSH Server
- Running Apache or Nginx on the *80* port, and doing a reverse-proxy on the *8080* port.
- Blocking the *8080* port from outside access.
- Running all ZMQ Circusd ports using IPC files instead of TCP ports, and tunneling all calls via SSH.





## 5.13 Design



Circus is composed of a main process called **circusd** which takes care of running all the processes. Each process managed by Circus is a child process of **circusd**.

Processes are organized in groups called **watchers**. A **watcher** is basically a command **circusd** runs on your system, and for each command you can configure how many processes you want to run.

The concept of *watcher* is useful when you want to manage all the processes running the same command – like restart them, etc.

**circusd** binds two ZeroMQ sockets:

- **REQ/REP** – a socket used to control **circusd** using json-based *commands*.

- **PUB/SUB** – a socket where **circusd** publishes events, like when a process is started or stopped.

---

**Note:** Despite its name, ZeroMQ is not a queue management system. Think of it as an inter-process communication (IPC) library.

---

Another process called **circusd-stats** is run by **circusd** when the option is activated. **circusd-stats**'s job is to publish CPU/Memory usage statistics in a dedicated **PUB/SUB** channel.

This specialized channel is used by **circus-top** and **circus-httpd** to display a live stream of the activity.

**circus-top** is a console script that mimics **top** to display all the CPU and Memory usage of the processes managed by Circus.

**circus-httpd** is the web management interface that will let you interact with Circus. It displays a live stream using web sockets and the **circusd-stats** channel, but also let you interact with **circusd** via its **REQ/REP** channel.

Last but not least, **circusctl** is a command-line tool that let you drive **circusd** via its **REQ/REP** channel.

You can also have plugins that subscribe to **circusd**'s **PUB/SUB** channel and let you send commands to the **REQ/REP** channel like **circusctl** would.

## 5.14 Why should I use Circus instead of X ?

### 1. Circus simplifies your web stack process management

Circus knows how to manage processes *and* sockets, so you don't have to delegate web workers management to a WSGI server.

See *Circus stack v.s. Classical stack*

### 2. Circus provides pub/sub and poll notifications via ZeroMQ

Circus has a *pub/sub* channel you can subscribe to. This channel receives all events happening in Circus. For example, you can be notified when a process is *flapping*, or build a client that triggers a warning when some processes are eating all the CPU or RAM.

These events are sent via a ZeroMQ channel, which makes it different from the stdin stream Supervisor uses:

- Circus sends events in a fire-and-forget fashion, so there's no need to manually loop through *all* listeners and maintain their states.
- Subscribers can be located on a remote host.

Circus also provides ways to get status updates via one-time polls on a req/rep channel. This means you can get your information without having to subscribe to a stream. The *Command-line tools* command provided by Circus uses this channel.

See *Step-by-step tutorial*.

### 3. Circus is (Python) developer friendly

While Circus can be driven entirely by a config file and the *circusctl* / *circusd* commands, it is easy to reuse all or part of the system to build your own custom process watcher in Python.

Every layer of the system is isolated, so you can reuse independently:

- the process wrapper (*Process*)
- the processes manager (*Watcher*)
- the global manager that runs several processes managers (*Arbiter*)

- and so on...

#### 4. Circus scales

One of the use cases of Circus is to manage thousands of processes without adding overhead – we’re dedicated to focus on this.

## 5.15 Step-by-step tutorial

The [examples directory](#) in the Circus repository contains many examples to get you started, but here’s a full tutorial that gives you an overview of the features.

We’re going to supervise a WSGI application.

### 5.15.1 Installation

Circus is tested under Mac OS X and Linux, on the latest Python 2.6 and 2.7. To run a full Circus, you will also need **libzmq**, **libevent** & **virtualenv**.

Under Debuntu:

```
$ sudo apt-get install libzmq-dev libevent-dev python-dev python-virtualenv
```

Create a virtualenv and install *circus*, *circus-web* and *chaussette* in it

```
$ virtualenv /tmp/circus
$ cd /tmp/circus
$ bin/pip install circus
$ bin/pip install circus-web
$ bin/pip install chaussette
```

Once this is done, you’ll find a plethora of commands in the local bin dir.

### 5.15.2 Usage

*Chaussette* comes with a default Hello world app, try to run it:

```
$ bin/chaussette
```

You should be able to visit <http://localhost:8080> and see *hello world*.

Stop Chaussette and add a *circus.ini* file in the directory containing:

```
[circus]
statsd = 1
httpd = 1

[watcher:webapp]
cmd = bin/chaussette --fd $(circus.sockets.web)
numprocesses = 3
use_sockets = True

[socket:web]
host = 127.0.0.1
port = 9999
```

This config file tells Circus to bind a socket on port 9999 and run 3 chaussettes workers against it. It also activates the Circus web dashboard and the statistics module.

Save it & run it using **circusd**:

```
$ bin/circusd --daemon circus.ini
```

Now visit <http://127.0.0.1:9999>, you should see the hello world app. The difference now is that the socket is managed by Circus and there are several web workers that are accepting connections against it.

---

**Note:** The load balancing is operated by the operating system so you're getting the same speed as any other pre-fork web server like Apache or NGinx. Circus does not interfere with the data that goes through.

---

You can also visit <http://localhost:8080/> and enjoy the Circus web dashboard.

### 5.15.3 Interaction

Let's use the **circusctl** shell while the system is running:

```
$ bin/circusctl
circusctl 0.7.1
circusd-stats: active
circushttd: active
webapp: active
(circusctl)
```

You get into an interactive shell. Type **help** to get all commands:

```
(circusctl) help

Documented commands (type help <topic>):
=====
add      get      list      numprocesses  quit      rm        start      stop
decr     globaloptions listen    numwatchers  reload    set       stats
dstats   incr     listsockets options        restart   signal    status

Undocumented commands:
=====
EOF  help
```

Let's try basic things. Let's list the web workers processes and add a new one:

```
(circusctl) list webapp
13712,13713,13714
(circusctl) incr webapp
4
(circusctl) list webapp
13712,13713,13714,13973
```

Congrats, you've interacted with your Circus! Get off the shell with Ctrl+D and now run **circus-top**:

```
$ bin/circus-top
```

This is a **top**-like command to watch all your processes' memory and CPU usage in real time.

Hit Ctrl+C and now let's quit Circus completely via **circus-ctl**:

```
$ bin/circusctl quit
ok
```

### 5.15.4 Next steps

You can plug your own WSGI application instead of Chaussette's hello world simply by pointing the application callable.

Chaussette also comes with many backends like Gevent or Meinheld.

Read <https://chaussette.readthedocs.org/> for all options.

## 5.16 Circus Use Cases

This chapter presents a few use cases, to give you an idea on how to use Circus in your environment.

### 5.16.1 Running a WSGI application

Running a WSGI application with Circus is quite interesting because you can watch & manage your *web workers* using *circus-top*, *circusctl* or the Web interface.

This is made possible by using Circus sockets. See *Circus stack v.s. Classical stack*.

Let's take an example with a minimal Pyramid application:

```
from pyramid.config import Configurator
from pyramid.response import Response

def hello_world(request):
    return Response('Hello %(name)s!' % request.matchdict)

config = Configurator()
config.add_route('hello', '/hello/{name}')
config.add_view(hello_world, route_name='hello')
application = config.make_wsgi_app()
```

Save this script into an **app.py** file, then install those projects:

```
$ pip install Pyramid
$ pip install chaussette
```

Next, make sure you can run your Pyramid application using the **chaussette** console script:

```
$ chaussette app.application
Application is <pyramid.router.Router object at 0x10a4d4bd0>
Serving on localhost:8080
Using <class 'chaussette.backend._waitress.Server'> as a backend
```

And check that you can reach it by visiting **<http://localhost:8080/hello/tarek>**

Now that your application is up and running, let's create a Circus configuration file:

```
[circus]
check_delay = 5
endpoint = tcp://127.0.0.1:5555
pubsub_endpoint = tcp://127.0.0.1:5556
```

```
stats_endpoint = tcp://127.0.0.1:5557

[watcher:webworker]
cmd = chaussette --fd $(circus.sockets.webapp) app.application
use_sockets = True
numprocesses = 3

[socket:webapp]
host = 127.0.0.1
port = 8080
```

This file tells Circus to bind a socket on port *8080* and run *chaussette* workers on that socket – by passing its fd.

Save it to *server.ini* and try to run it using **circusd**

```
$ circusd server.ini
[INFO] Starting master on pid 8971
[INFO] sockets started
[INFO] circusd-stats started
[INFO] webapp started
[INFO] Arbiter now waiting for commands
```

Make sure you still get the app on **<http://localhost:8080/hello/tarek>**.

Congrats ! you have a WSGI application running 3 workers.

You can run the *The Web Console* or the *Command-line tools*, and enjoy Circus management.

## 5.16.2 Running a Django application

Running a Django application is done exactly like running a WSGI application. Use the *PYTHONPATH* to import the directory the project is in, the directory that contains the directory that has settings.py in it (with Django 1.4+ this directory has manage.py in it)

```
[socket:dwebapp]
host = 127.0.0.1
port = 8080

[watcher:dwebworker]
cmd = chaussette --fd $(circus.sockets.dwebapp) dproject.wsgi.application
use_sockets = True
numprocesses = 2

[env:dwebworker]
PYTHONPATH = /path/to/parent-of-dproject
```

If you need to pass the *DJANGO\_SETTINGS\_MODULE* for a backend worker for example, you can pass that also though the *env* configuration option:

```
[watcher:dbackend]
cmd = /path/to/script.py
numprocesses=3

[env:dbackend]
PYTHONPATH = /path/to/parent-of-dproject
DJANGO_SETTINGS_MODULE=dproject.settings
```

See <http://chaussette.readthedocs.org> for more about chaussette.

## 5.17 Troubleshooting

By default, *circusd* keeps its logging to *stdout* rather sparse. This lack of output can make things hard to troubleshoot when processes seem to be having trouble starting.

To increase the logging *circusd* provides, try increasing the log level. To see the available log levels just use the *-help* flag.

```
$ circus --log-level debug test.ini
```

One word of warning. If a process is flapping and the debug log level is turned on, you will see messages for each start attempt. It might be helpful to configure the app that is flapping to use a *warmup\_delay* to slow down the messages to a manageable pace.

```
[watcher:webapp]
cmd = python -m myapp.wsgi
warmup_delay = 5
```

By default, *stdout* and *stderr* are captured by the *circusd* process. If you are testing your config and want to see the output in line with the *circusd* output, you can configure your watcher to use the *StdoutStream* class.

```
[watcher:webapp]
cmd = python -m myapp.wsgi
stdout_stream.class = StdoutStream
stderr_stream.class = StdoutStream
```

If your application is producing a traceback or error when it is trying to start up you should be able to see it in the output.

## 5.18 Glossary

- arbiter** The *arbiter* is responsible for managing all the watchers within circus, ensuring all processes run correctly.
- controller** A *controller* contains the set of actions that can be performed on the arbiter.
- flapping** The *flapping detection* subscribes to events and detects when some processes are constantly restarting.
- process, worker, workers, processes** A *process* is an independent OS process instance of your program. A single watcher can run one or more processes. We also call them workers.
- pub/sub** Circus has a *pubsub* that receives events from the watchers and dispatches them to all subscribers.
- remote controller** The *remote controller* allows you to communicate with the controller via ZMQ to control Circus.
- watchers, watcher** A *watcher* is the program you tell Circus to run. A single Circus instance can run one or more watchers.

## 5.19 Contributing to Circus

Circus has been started at Mozilla but its goal is not to stay only there. We're trying to build a tool that's useful for others, and easily extensible.

We really are open to any contributions, in the form of code, documentation, discussions, feature proposal etc.

You can start a topic in our mailing list : <http://tech.groups.yahoo.com/group/circus-dev/>

Or add an issue in our [bug tracker](#)



### 5.19.1 Fixing typos and enhancing the documentation

It's totally possible that your eyes are bleeding while reading this half-english half-french documentation, don't hesitate to contribute any rephrasing / enhancement on the form in the documentation. You probably don't even need to understand how Circus works under the hood to do that.

### 5.19.2 Adding new features

New features are of course very much appreciated. If you have the need and the time to work on new features, adding them to Circus shouldn't be that complicated. We tried very hard to have a clean and understandable API, hope it serves the purpose.

You will need to add documentation and tests alongside with the code of the new feature. Otherwise we'll not be able to accept the patch.

### 5.19.3 How to submit your changes

We're using git as a DVCS. The best way to propose changes is to create a branch on your side (via `git checkout -b branchname`) and commit your changes there. Once you have something ready for prime-time, issue a pull request against this branch.

We are following this model to allow to have low coupling between the features you are proposing. For instance, we can accept one pull request while still being in discussion for another one.

Before proposing your changes, double check that they are not breaking anything! You can use the `tox` command to ensure this, it will run the testsuite under the different supported python versions.

Please use : <http://issue2pr.herokuapp.com/> to reference a commit to an existing circus issue, if any.

### 5.19.4 Avoiding merge commits

Avoiding merge commits allows to have a clean and readable history. To do so, instead of doing "git pull" and letting git handling the merges for you, using `git pull --rebase` will put your changes after the changes that are committed in the branch, or when working on master.

That is, for us core developers, it's not possible anymore to use the handy github green button on pull requests if developers didn't rebased their work themselves or if we wait too much time between the request and the actual merge. Instead, the flow looks like this:

```
git remote add name repo-url
git fetch name
git checkout feature-branch
git rebase master

# check that everything is working properly and then merge on master
git checkout master
git merge feature-branch
```

### 5.19.5 Discussing

If you find yourself in need of any help while looking at the code of Circus, you can go and find us on irc at #mozilla-circus on irc.freenode.org (or if you don't have any IRC client, use [the webchat](#))

You can also start a thread in our mailing list - <http://tech.groups.yahoo.com/group/circus-dev>

## 5.20 How to add new commands in circus

If you want to add a new command, we tried to make this as simple as possible. You need to do three main things:

1. create a “your\_command.py” file under *circus/commands/*.
2. Implement a single class in there, with predefined methods
3. Add the new command in *circus/commands/\_\_init\_\_.py*.

Let’s say we want to add a command which returns the number of watchers actually in use, we would do something like this (extensively commented to allow you to follow more easily):

```
class NumWatchers(Command):
    """It is a good practice to describe what the class does here.

    Have a look at other commands to see how we are used to format this
    text. It will be used to automatically appear in the documentation of
    circus, so don't be affraid of being exhaustive, that's what it is made
    for.
    """
    # all the commands need to inherit from `circus.commands.base.Command`

    name = "numwatchers"
    # you need to specify a name so we find back the command somehow

    options = [(' ', 'optname', default_value, 'description')]
    # XXX describe options

    properties = ['foo', 'bar']
    # properties list the command arguments that are mandatory. If they are
    # not provided, then an error will be thrown

    def execute(self, arbiter, props):
        # the execute method is the core of the command: put here all the
        # logic of the command and return a dict containing the values you
        # want to return, if any
        return {"numwatchers": arbiter.numwatchers()}

    def console_msg(self, msg):
        # msg is what is returned by the execute method.
        # this method is used to format the response for a console (it is
        # used for instance by circusctl to print its messages)
        return "a string that will be displayed"

    def validate(self, props):
        # this method is used to validate that the arguments passed to the
        # command are correct. An ArgumentError should be thrown in case
        # there is an error in the passed arguments (for instance if they
        # do not match together.
        # In case there is a problem wrt their content, a MessageError
        # should be thrown. This method can modify the content of the props
        # dict, it will be passed to execute afterwards.
```

## 5.21 History

### 5.21.1 0.9.1 - 2013-07-17

- When a PYTHONPATH is defined in a config file, it's loaded in sys.path so hooks can be located there - #477, #481
- Use a single argument for add\_callback so it works with PyZMQ < 13.1.x - see #478

### 5.21.2 0.9 - 2013-07-16

- added [env] sections wildcards
- added global [env] section
- fixed hidden exception when circus-web is not installed - #424
- make sure incr/decr commands really use the nb option - #421
- Fix watcher virtualenv site-packages not in PYTHONPATH
- make sure we don't try to remove more processes than 0 - #429
- updated bootstrap.py - #436
- fixed multiplatform separator in pythonpath virtualenv watcher
- refactored socket close function
- Ensure env sections are applied to all watchers - #437
- added the reloadconfig command
- added circus.green and removed gevent from the core - #441, #452
- silenced spurious stdout & warnings in the tests - #438
- \$(circus.env.\*) can be used for all options in the config now
- added a before\_spawn hook
- correct the path of circusd in systemd service file - #450
- make sure we can change hooks and set streams via CLI - #455
- improved doc
- added a spawn\_count stat in watcher
- added min\_cpu and min\_mem parameters in ResourceWatcher plugin
- added the FQDN information to the arbiter.

### 5.21.3 0.8.1 - 2013-05-28

- circusd-stats was choking on unix sockets - #415
- circusd-stats & circushttd child processes stdout/stderr are now left open by default. Python <= 2.7.5 would choke in the logging module in case the 2/3 fds were closed - #415
- Now redirecting to /dev/null in the child process instead of closing. #417

#### 5.21.4 0.8 - 2013-05-24

- Integrated log handlers into zmq io loop.
- Make redirector restartable and subsequently more robust.
- Uses zmq.green.eventloop when gevent is detected
- Added support for CIRCUSCTL\_ENDPOINT environment variable to circusctl - #396
- util: fix bug in to\_uid function - #397
- Remove handler on ioloop error - #398.
- Improved test coverage
- Deprecated the 'service' option for the ResourceWatcher plugin - #404
- removed psutil.error usage
- Added UDP discovery in circusd - #407
- Now allowing globs at arbitrary directory levels - #388
- Added the 'statd' configuration option - #408
- Add pidfile, logoutput and loglevel option to circus configuration file - #379
- Added a tutorial in the docs.
- make sure we're merging all sections when using include - #414
- added pipe\_stdout, pipe\_stderr, close\_child\_stderr & close\_child\_stdout options to the Process class
- added close\_child\_stderr & close\_child\_stdout options to the watcher

#### 5.21.5 0.7.1 - 2013-05-02

- Fixed the respawn option - #382
- Make sure we use an int for the timeout - #380
- display the unix sockets as well - #381
- Make sure it works with the latest pyzmq
- introduced a second syntax for the fd notation

#### 5.21.6 0.7 - 2013-04-08

- Fix get\_arbiter example to use a dict for the watchers argument. #304
- Add some troubleshooting documentation #323
- Add python buildout support
- Removed the gevent and the thread redirectors. now using the ioloop - fixes #346. Relates #340
- circus.web is now its own project
- removed the pyzmq patching
- Allow the watcher to be configured but not started #283
- Add an option to load a virtualenv site dir

- added on\_demand watchers
- added doc about nginx+websockets #371
- now properly parsing the options list of each command #369
- Fixed circusd-stats events handling #372
- fixed the overflow issue in circus-top #378
- many more things...

### 5.21.7 0.6 - 2012-12-18

- Patching protocols name for sockets - #248
- Don't autoscale graphs. #240
- circusctl: add per command help, from docstrings #217
- Added workers hooks
- Added Debian package - #227
- Added Redis, HTTP Observer, Full stats & Resource plugins
- Now processes can have titles
- Added autocompletion
- Added process/watcher age in the webui
- Added SSH tunnel support
- Now using pyzmq.green
- Added upstart script & Varnish doc
- Added environment variables & sections
- Added unix sockets support
- Added the *respawn* option to have single-run watchers
- Now using tox in the tests
- Allow socket substitution in args
- New doc theme
- New rotation options for streams: max\_bytes/backup\_count

### 5.21.8 0.5.2 - 2012-07-26

- now patching the thread module from the stdlib to avoid some Python bugs - #203
- better looking circusctl help screen
- uses pustil get\_nice() when available (nice was deprecated) - #208
- added max\_age support - #221
- only call listen() on SOCK\_STREAM or SOCK\_SEQPACKET sockets
- make sure the controller empties the plugins list in update\_watchers() - #220
- added -log-level and -log-output to circushttd

- fix the process killing via the web UI - #219
- now circus is `zc.buildout` compatible for scripts.
- cleanup the websocket when the client disconnect - #225
- fixed the default value for the endpoint - #199
- splitted `circushttd` in logical modules

### **5.21.9 0.5.1 - 2012-07-11**

- Fixed a bunch of typos in the documentation
- Added the debug option
- Package `web-requirements.txt` properly
- Added a `errno` error code in the messages - fixes #111

### **5.21.10 0.5 - 2012-07-06**

- added socket support
- added a `listsocket` command
- sockets have stats too !
- fixed a lot of small bugs
- removed the `wid` - now using `pid` everywhere
- faster tests
- changed the variables syntax
- use `pyzmq`'s `ioloop` in more places
- now using `iowait` for all `select()` calls
- `incr/decr` commands now have an `nbprocess` parameter
- Add a `reproduce_env` option to watchers
- Add a new `UNEXISTING` status to the processes
- Added the global `httd` option to run `circushttd` as a watcher

### **5.21.11 0.4 - 2012-06-12**

- Added a plugin system
- Added a “singleton” option for watchers
- Fixed circus-top screen flickering
- Removed threads from `circus.stats` in favor of `zmq` periodic callbacks
- Enhanced the documentation
- Circus client now have a `send_message` api
- The flapping feature is now a plugin

- Every command line tool have a `--version` option
- Added a statsd plugin (sends the events from circus to statsd)
- The web UI now uses websockets (via socketio) to get the stats
- The web UI now uses sessions for “flash messages” in the web ui

#### 5.21.12 0.3.4 - 2012-05-30

- Fixed a race condition that prevented the controller to cleanly reap finished processes.
- Now `check_flapping` can be controlled in the configuration. And activated/deactivated per watcher.

#### 5.21.13 0.3.3 - 2012-05-29

- Fixed the regression on the uid handling

#### 5.21.14 0.3.2 - 2012-05-24

- allows optional `args` property to `add_watcher` command.
- added `circushttd`, `circus-top` and `circusd-stats`
- allowing `Arbiter.add_watcher()` to set all `Watcher` option
- make sure the redirectors are re-created on restarts

#### 5.21.15 0.3.1 - 2012-04-18

- fix: make sure watcher’ defaults aren’t overridden
- added a `StdoutStream` class.

#### 5.21.16 0.3 - 2012-04-18

- added the streaming feature
- now displaying coverage in the Sphinx doc
- fixed the way the processes are killed (no more `SIGQUIT`)
- the configuration has been factored out
- `setproctitle` support

#### 5.21.17 0.2 - 2012-04-04

- Removed the *show* name. replaced by *watcher*.
- Added support for setting process **r**limit.
- Added support for include dirs in the config file.
- Fixed a couple of leaking file descriptors.
- Fixed a core dump in the flapping

- Doc improvements
- Make sure circusd errors properly when another circusd is running on the same socket.
- get\_arbiter now accepts several watchers.
- Fixed the cmd vs args vs executable in the process init.
- Fixed `--start` on `circusctl add`

### 5.21.18 0.1 - 2012-03-20

- initial release

## 5.22 Copyright

Circus was initiated by Tarek Ziade and is licenced under APLv2

Benoit Chesneau was an early contributor and did many things, like most of the `circus.commands` work.

### 5.22.1 Licence

Copyright 2012 - Mozilla Foundation  
Copyright 2012 - Benoit Chesneau

Licensed under the Apache License, Version 2.0 (the "License");  
you may not use this file except in compliance with the License.  
You may obtain a copy of the License at

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software  
distributed under the License is distributed on an "AS IS" BASIS,  
WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.  
See the License for the specific language governing permissions and  
limitations under the License.

### 5.22.2 Contributors

See the full list at <https://github.com/mozilla-services/circus/blob/master/NOTICE>



# CONTRIBUTIONS AND FEEDBACK

More on contribution: *Contributing to Circus*.

Useful Links:

- There's a mailing list for any feedback or question: <http://tech.groups.yahoo.com/group/circus-dev/>
- The repository and issue tracker is at GitHub : <https://github.com/mozilla-services/circus>
- Join us on the IRC : Freenode, channel **#mozilla-circus**