

---

# **uwsgiconf Documentation**

*Release 0.20.0*

**Igor 'idle sign' Starikov**

**Jan 06, 2020**



---

# Contents

---

<b>1</b>	<b>Description</b>	<b>3</b>
<b>2</b>	<b>Requirements</b>	<b>5</b>
<b>3</b>	<b>Table of Contents</b>	<b>7</b>
3.1	Quickstart . . . . .	7
3.2	Presets . . . . .	9
3.3	Command-line interface (CLI) . . . . .	12
3.4	Python uwsgi module stub . . . . .	13
3.5	Contrib . . . . .	34
3.6	FAQ . . . . .	36
3.7	Configuration [Dynamic/Runtime] . . . . .	37
3.8	Configuration [Static] . . . . .	57
<b>4</b>	<b>Get involved into uwsgiconf</b>	<b>159</b>
	<b>Python Module Index</b>	<b>161</b>
	<b>Index</b>	<b>163</b>



<https://github.com/idlesign/uwsgiconf>



### *Configure uWSGI from your Python code*

If you think you know uWSGI you're probably wrong. It is always more than you think it is. There are so many subsystems and [options](#) (800+) it is difficult to even try to wrap your mind around.

**uwsgi**conf allowing to define uWSGI configurations in Python tries to improve things the following ways:

- It structures options for various subsystems using classes and methods;
- It uses docstrings and sane naming to facilitate navigation;
- It ships some useful presets to reduce boilerplate code;
- It encourages configuration reuse;
- It comes with CLI to facilitate configuration;
- It features easy to use and documented **uwsgi stub** Python module;
- It offers **runtime** package, similar to **uwsgidecorators**, but with more abstractions;
- It features integration with Django Framework;
- It is able to generate configuration files for Systemd, Upstart.
- It can use `pyuwsgi`.

*Consider using IDE with autocompletion and docstrings support to be more productive with uwsgi*conf.

By that time you already know that **uwsgi**conf is just another configuration method. [Why?](#)





## CHAPTER 2

---

### Requirements

---

1. Python 2.7+, 3.5+
2. `click` package (optional, for CLI)
3. `uWSGI` (`uwsgi` or `pyuwsgi`)



## 3.1 Quickstart

### 3.1.1 Install

You can get and install **uwsgi**conf from PyPI using `pip`:

```
$ pip install uwsgi
```

#### CLI

**uwsgi**conf expects `click` package available for CLI but won't install this dependency by default.

Use the following command to install **uwsgi**conf with `click`:

```
$ pip install uwsgi[cli]
```

### 3.1.2 Using a preset to run Python web application

Let's make `uwsgi`conf.py. There we configure uWSGI using nice `PythonSection` preset to run our web app.

```
from uwsgi.conf import configure_uwsgi
from uwsgi.presets.nice import PythonSection

def get_configurations():
    """This should return one or more Section or Configuration objects.
    In such a way you can configure more than one uWSGI instance in the same place.

    """
    my_app_dir = '/home/idle/myapp/'
```

(continues on next page)

(continued from previous page)

```

section = PythonSection(
    # Reload uWSGI when this file is updated.
    touch_reload=__file__,

    params_python=dict(
        # Let's add something into Python path.
        search_path='/opt/apps_shared/',
    ),

    wsgi_module=my_app_dir + 'wsgi.py',

    # We'll redirect logs into a file.
    log_into=my_app_dir + 'app.log',

    # If your uWSGI has no basic plugins embedded
    # (i.e. not from PyPI) you can give uwsgi.conf a hint:
    # embedded_plugins=False,

).networking.register_socket(
    # Make app available at http://127.0.0.1:8000
    PythonSection.networking.sockets.http('127.0.0.1:8000'))

return section

# Almost done. One more thing:
configure_uwsgi(get_configurations)

```

Now we are ready to use this configuration:

```

$ uwsgi.conf compile > myconf.ini
$ uwsgi myconf.ini

; or instead just
$ uwsgi.conf run

```

### 3.1.3 Configuration with multiple sections

Let's configure uWSGI to use Emperor Broodlord mode as described [here](#) using Broodlord preset.

```

from uwsgi.conf import Section, Configuration
from uwsgi.presets.empire import Broodlord

emperor, zerg = Broodlord(

    zerg_socket='/tmp/broodlord.sock',
    zerg_count=40,
    zerg_die_on_idle=30,

    vassals_home='/etc/vassals',
    vassal_queue_items_sos=10,

    # We'll use the same basic params both for Broodlord Emperor and his zergs.
    section_emperor=(Section()).

```

(continues on next page)

(continued from previous page)

```

# NOTE. Here we use a shortcut for ``set_basic_params`` methods:
# E.g.: instead of `master_process.set_basic_params(enable=True)`
# you say `master_process(enable=True)`.
# But in that case you won't get any arg hints from you IDE.
master_process(enable=True).
workers(count=1).
logging(no_requests=True).
python.set_wsgi_params(module='werkzeug.testapp:test_app')
),
).configure()

# Bind Emperor to socket.
emperor.networking.register_socket(Section.networking.sockets.default(':3031'))

# Put Emperor and zerg sections into configuration.
multisection_config = Configuration([emperor, zerg])

```

## 3.2 Presets

Presets are means to reduce boilerplate code.

Use them as shortcuts to spend less time on configuring.

### 3.2.1 Preset: Empire

This preset offers configuration helpers related to Emperor and his vassals.

```

class uwsgi.conf.presets.empire.Broodlord(zerg_socket, zerg_die_on_idle=None, vas-
sals_home=None, zerg_count=None,
vassal_overload_sos_interval=None,
vassal_queue_items_sos=None, section_
tion_emperor=None, section_zerg=None)

```

This mode is a way for a vassal to ask for reinforcements to the Emperor.

Reinforcements are new vassals spawned on demand generally bound on the same socket.

**Warning:** If you are looking for a way to dynamically adapt the number of workers of an instance, check the Cheaper subsystem - adaptive process spawning mode.

*Broodlord mode is for spawning totally new instances.*

- <http://uwsgi-docs.readthedocs.io/en/latest/Broodlord.html>

#### Parameters

- **zerg\_socket** (*str|unicode*) – Unix socket to bind server to.
- **zerg\_die\_on\_idle** (*int*) – A number of seconds after which an idle zerg will be destroyed.
- **vassals\_home** (*str|unicode|list[str|unicode]*) – Set vassals home.
- **zerg\_count** (*int*) – Maximum number of zergs to spawn.

- **vassal\_overload\_sos\_interval** (*int*) – Ask emperor for reinforcement when overloaded. Accepts the number of seconds to wait between asking for a new reinforcements.
- **vassal\_queue\_items\_sos** (*int*) – Ask emperor for sos if backlog queue has more items than the value specified

**configure** ()

Configures broodlord mode and returns emperor and zerg sections.

**Return type** tuple

### 3.2.2 Preset: Nice

This preset offers nice configuration defaults.

```
class uwsgi.conf.presets.nice.Section (name=None, touch_reload=None, workers=None,  
threads=None, mules=None, owner=None,  
log_into=None, log_dedicated=None, process_prefix=None,  
ignore_write_errors=None,  
**kwargs)
```

Basic nice configuration.

#### Parameters

- **name** (*str/unicode*) – Section name.
- **touch\_reload** (*str/list*) – Reload uWSGI if the specified file or directory is modified/touched.
- **workers** (*int*) – Spawn the specified number of workers (processes). Default: workers number equals to CPU count.
- **threads** (*int/bool*) – Number of threads per worker or `True` to enable user-made threads support.
- **mules** (*int*) – Number of mules to spawn.
- **owner** (*str/unicode*) – Set process owner user and group.
- **log\_into** (*str/unicode*) – Filepath or UDP address to send logs into.
- **log\_dedicated** (*bool*) – If `True` all logging will be handled with a separate thread in master process.
- **process\_prefix** (*str/unicode*) – Add prefix to process names.
- **ignore\_write\_errors** (*bool*) – If `True` no annoying SIGPIPE/write/writew errors will be logged, and no related exceptions will be raised.

---

**Note:** Usually such errors could be seen on client connection cancelling and are safe to ignore.

---

- **kwargs** –

**get\_log\_format\_default** ()

Returns default log message format.

---

**Note:** Some params may be missing.

---

**classmethod** `get_bundled_static_path` (*filename*)

Returns a full path to a static HTML page file bundled with uwsgi.conf.

**Parameters** `filename` (*str*) – File name to return path to.

**Examples:**

- 403.html
- 404.html
- 500.html
- 503.html

**Return type** `str`

**configure\_maintenance\_mode** (*trigger, response*)

Allows maintenance mode when a certain response is given for every request if a trigger is set.

**Parameters**

- **trigger** (*str*) – This triggers maintenance mode responses. Should be a path to a file: if file exists, maintenance mode is on.
- **response** (*str*) – Response to to give in maintenance mode.

**Supported:**

- File path - this file will be served in response.
- URLs starting with `http` - requests will be redirected there using 302. This is often discouraged, because it may have search ranking implications.

**configure\_owner** (*owner='www-data'*)

Shortcut to set process owner data.

**Parameters** `owner` (*str/unicode*) – Sets user and group. Default: `www-data`.

**configure\_certbot\_https** (*domain, webroot, address=None, allow\_shared\_sockets=None*)

Enables HTTPS using certificates from Certbot <https://certbot.eff.org>.

---

**Note:** This relies on `webroot` mechanism of Certbot - <https://certbot.eff.org/docs/using.html#webroot>

Sample command to get a certificate: `certbot certonly --webroot -w /webroot/path/ -d mydomain.org`

---

**Parameters**

- **domain** (*str*) – Domain name certificates issued for (the same as in `-d` option in the above command).
- **webroot** (*str*) – Directory to store challenge files to get and renew the certificate (the same as in `-w` option in the above command).
- **address** (*str*) – Address to bind socket to.
- **allow\_shared\_sockets** (*bool*) – Allows using shared sockets to bind to privileged ports. If not provided automatic mode is enabled: shared are allowed if current user is not root.

```
class uwsgiconf.presets.nice.PythonSection (name=None,          params_python=None,
                                             wsgi_module=None,    wsgi_callable=None,
                                             embedded_plugins=True, require_app=True,
                                             threads=True, **kwargs)
```

Basic nice configuration using Python plugin.

#### Parameters

- **name** (*str/unicode*) – Section name.
- **params\_python** (*dict*) – See Python plugin basic params.
- **wsgi\_module** (*str/unicode*) – WSGI application module path or filepath.  
**Example:** `mypackage.my_wsgi_module` – read from `application` attr of `mypackage/my_wsgi_module.py` `mypackage.my_wsgi_module:my_app` – read from `my_app` attr of `mypackage/my_wsgi_module.py`
- **wsgi\_callable** (*str/unicode/callable*) – WSGI application callable name. Default: `application`.
- **embedded\_plugins** (*bool/None*) – This indicates whether plugins were embedded into uWSGI, which is the case if you have uWSGI from PyPI.
- **require\_app** (*bool*) – Exit if no app can be loaded.
- **threads** (*int/bool*) – Number of threads per worker or `True` to enable user-made threads support.
- **kwargs** –

## 3.3 Command-line interface (CLI)

`uwsgiconf` comes with CLI (`click` package required) to facilitate configuration.

```
; To install uwsgiconf with click:
$ pip install uwsgiconf[cli]
```

### 3.3.1 Compile

Compiles classic uWSGI configuration from a given `uwsgiconf` configuration module (or from the default one - `uwsgiconf.py`).

---

**Note:** Be sure that your configuration module defines `configuration` attribute. It must hold one or more Configuration or Section (those will be automatically casted to configurations) objects. Callable as attribute value is supported.

---

```
; This compiles uwsgiconf.py from the current directory
; into .ini and prints that out:
$ uwsgiconf compile

; This compiles there/thisfile.py file:
$ uwsgiconf compile there/thisfile.py

; Add "> target_file.ini" to redirect output (configuration) into a file.
```



### 3.3.2 Run

Runs uWSGI using configuration from a given *uwsgi<sup>conf</sup>* configuration module (or from the default one - *uwsgi<sup>conf</sup>.py*).

**Note:** uWSGI process will replace uwsgi<sup>conf</sup> process.

```
; This runs uWSGI using uwsgiconf.py from the current directory.
$ uwsgiconf run

; This runs uWSGI using configuration from there/thisfile.py:
$ uwsgiconf run there/thisfile.py
```

### 3.3.3 Probe plugins

Shows available uWSGI plugins.

```
$ uwsgiconf probe_plugins
```

### 3.3.4 Systemd and other configs

You can generate configuration files to launch uwsgi<sup>conf</sup> automatically using system facilities.

Config contents in sent to stdout and could be redirected into a file.

```
$ uwsgiconf sysinit systemd
$ uwsgiconf sysinit upstart
```

Usage example for Systemd:

```
; Generate and save config into `myapp.service` file
$ uwsgiconf sysinit --project myapp > myapp.service

; Copy config into standard location
$ sudo cp myapp.service /etc/systemd/system/

; Reload available configs information and run service
$ sudo sh -c "systemctl daemon-reload; systemctl start myapp.service"

; Watch application log realtime (if syslog is used)
$ journalctl -fu myapp.service
```

## 3.4 Python uwsgi module stub

**uwsgi<sup>conf</sup>** comes with documented **uwsgi** module that you can import instead of `import uwsgi`.

```
# Instead of
import uwsgi

# you can do.
from uwsgiconf import uwsgi
```

That way **uwsgi** will be available runtime as usual, besides you will get autocompletion and hints in IDE, and won't get `ImportError` when run without **uwsgi**.

This also will facilitate your testing a bit, for those simple cases when you won't expect any result from **uwsgi** function.

**Warning:** This is a stub module, so it doesn't really implement functions defined in it. Use it for documentation purposes.

`uwsgi.conf.uwsgi_stub.is_stub = True`

Indicates whether stub is used instead of real *uwsgi* module.

`uwsgi.conf.uwsgi_stub.SPOOL_IGNORE = 0`

Spooler function result.

Ignore this task, if multiple languages are loaded in the instance all of them will fight for managing the task. This return values allows you to skip a task in specific languages.

- <http://uwsgi-docs.readthedocs.io/en/latest/Spooler.html#setting-the-spooler-function-callable>

`uwsgi.conf.uwsgi_stub.SPOOL_OK = -2`

Spooler function result.

The task has been completed, the spool file will be removed.

- <http://uwsgi-docs.readthedocs.io/en/latest/Spooler.html#setting-the-spooler-function-callable>

`uwsgi.conf.uwsgi_stub.SPOOL_RETRY = -1`

Spooler function result.

Something is temporarily wrong, the task will be retried at the next spooler iteration.

- <http://uwsgi-docs.readthedocs.io/en/latest/Spooler.html#setting-the-spooler-function-callable>

`uwsgi.conf.uwsgi_stub.SymbolsImporter = None`

SymbolsImporter type.

`uwsgi.conf.uwsgi_stub.SymbolsZipImporter = None`

SymbolsZipImporter type.

`uwsgi.conf.uwsgi_stub.ZipImporter = None`

ZipImporter type.

`uwsgi.conf.uwsgi_stub.applications = None`

Applications dictionary mapping mountpoints to application callables.

---

**Note:** Can be None.

---

- <http://uwsgi.readthedocs.io/en/latest/Python.html#application-dictionary>

`uwsgi.conf.uwsgi_stub.buffer_size = 0`

The current configured buffer size in bytes.

`uwsgi.conf.uwsgi_stub.cores = 0`

Detected number of processor cores.

`uwsgi.conf.uwsgi_stub.env = {}`

Request environment dictionary.

`uwsgi.conf.uwsgi_stub.has_threads = False`

Flag indicating whether thread support is enabled.

`uwsgi.conf.uwsgi_stub.hostname = ''`  
Current host name.

`uwsgi.conf.uwsgi_stub.magic_table = {}`  
Current mapping of configuration file “magic” variables.

- <http://uwsgi.readthedocs.io/en/latest/Configuration.html#magic-variables>

`uwsgi.conf.uwsgi_stub.numproc = 0`  
Number of workers (processes) currently running.

`uwsgi.conf.uwsgi_stub.opt = {}`  
The current configuration options, including any custom placeholders.

`uwsgi.conf.uwsgi_stub.post_fork_hook = <bound method type.run of <class 'uwsgi.conf.runtime.p`  
Function to be called after process fork (spawning a new worker/mule).

`uwsgi.conf.uwsgi_stub.spooler = <bound method type._process_message_raw of <class 'uwsgi.conf`  
Function to be called for spooler messages processing.

`uwsgi.conf.uwsgi_stub.sockets = []`  
Current list of file descriptors for registered sockets.

`uwsgi.conf.uwsgi_stub.start_response = None`  
Callable spitting UWSGI response.

`uwsgi.conf.uwsgi_stub.started_on = 0`  
uWSGI’s startup Unix timestamp.

`uwsgi.conf.uwsgi_stub.unbit = False`  
Unbit internal flag.

`uwsgi.conf.uwsgi_stub.version = '0.0.0'`  
The uWSGI version string.

<b>Warning:</b> Bytes are returned for Python 3.
--

**Return type** byteslstr

`uwsgi.conf.uwsgi_stub.version_info = (0, 0, 0, 0, '')`  
Five-elements version number tuple.

**Return type** tuple[int, int, int, int, byteslstr]

`uwsgi.conf.uwsgi_stub.mule_msg_hook (message)`  
Registers a function to be called for each mule message.

`uwsgi.conf.uwsgi_stub.accepting ()`  
Called to notify the master that the worker is accepting requests, this is required for `touch_chain_reload` to work.

**Return type** None

`uwsgi.conf.uwsgi_stub.add_cron (signal, minute, hour, day, month, weekday)`  
Adds cron. The interface to the uWSGI signal cron facility. The syntax is

---

**Note:** The last 5 arguments work similarly to a standard crontab, but instead of “\*”, use -1, and instead of “/2”, “/3”, etc. use -2 and -3, etc.

---

### Parameters

- **signal** (*int*) – Signal to raise.
- **minute** (*int*) – Minute 0-59. Defaults to *each*.
- **hour** (*int*) – Hour 0-23. Defaults to *each*.
- **day** (*int*) – Day of the month number 1-31. Defaults to *each*.
- **month** (*int*) – Month number 1-12. Defaults to *each*.
- **weekday** (*int*) – Day of a the week number. Defaults to *each*. 0 - Sunday 1 - Monday 2 - Tuesday 3 - Wednesday 4 - Thursday 5 - Friday 6 - Saturday

**Return type** bool

**Raises** **ValueError** – If unable to add cron rule.

`uwsgi.conf.uwsgi_stub.add_file_monitor(signal, filename)`

Maps a specific file/directory modification event to a signal.

### Parameters

- **signal** (*int*) – Signal to raise.
- **filename** (*str/unicode*) – File or a directory to watch for its modification.

**Return type** None

**Raises** **ValueError** – If unable to register monitor.

`uwsgi.conf.uwsgi_stub.add_ms_timer(signal, period)`

Add a millisecond resolution timer.

### Parameters

- **signal** (*int*) – Signal to raise.
- **period** (*int*) – The interval (milliseconds) at which to raise the signal.

**Return type** None

**Raises** **ValueError** – If unable to add timer.

`uwsgi.conf.uwsgi_stub.add_rb_timer(signal, period, repeat=0)`

Add a red-black timer.

### Parameters

- **signal** (*int*) – Signal to raise.
- **period** (*int*) – The interval (seconds) at which the signal is raised.
- **repeat** (*int*) – How many times to send signal. Will stop after ther number is reached. Default: 0 - infinitely.

**Return type** None

**Raises** **ValueError** – If unable to add timer.

`uwsgi.conf.uwsgi_stub.add_timer(signal, period)`

Add timer.

### Parameters

- **signal** (*int*) – Signal to raise.
- **period** (*int*) – The interval (seconds) at which to raise the signal.

**Return type** None

**Raises** **ValueError** – If unable to add timer.

`uwsgiconf.uwsgi_stub.add_var(name, value)`

Registers custom request variable.

Can be used for better integration with the internal routing subsystem.

**Parameters**

- **name** (*str/unicode*) –
- **value** (*str/unicode*) –

**Return type** bool

**Raises** **ValueError** – If buffer size is not enough.

`uwsgiconf.uwsgi_stub.alarm(name, message)`

Issues the given alarm with the given message.

**Parameters**

- **name** (*str/unicode*) –
- **message** (*str/unicode*) – Message to pass to alarm.

**Return type** None

`uwsgiconf.uwsgi_stub.async_connect(socket)`

Issues socket connection. And returns a file descriptor.

- <http://uwsgi.readthedocs.io/en/latest/Async.html>

**Parameters** **socket** (*str/unicode*) –

**Return type** int

`uwsgiconf.uwsgi_stub.async_sleep(seconds)`

Suspends handling the current request and passes control to the next async core.

- <http://uwsgi.readthedocs.io/en/latest/Async.html>

**Parameters** **seconds** – Sleep time, in seconds.

**Rtype** None

`uwsgiconf.uwsgi_stub.cache_clear(cache)`

Clears cache with the given name.

**Parameters** **cache** (*str/unicode*) – Cache name with optional address (if @-syntax is used).

**Return type** None

`uwsgiconf.uwsgi_stub.cache_dec(key, value=1, expires=None, cache=None)`

Decrements the specified key value by the specified value.

- <http://uwsgi.readthedocs.io/en/latest/Changelog-1.9.9.html#math-for-cache>

**Parameters**

- **key** (*str/unicode*) –
- **value** (*int*) –

- **expires** (*int*) – Expire timeout (seconds).
- **cache** (*str/unicode*) – Cache name with optional address (if @-syntax is used).

**Return type** bool

`uwsgi.conf.uwsgi_stub.cache_del` (*key, cache=None*)

Deletes the given cached key from the cache.

**Parameters**

- **key** (*str/unicode*) – The cache key to delete.
- **cache** (*str/unicode*) – Cache name with optional address (if @-syntax is used).

**Return type** None

`uwsgi.conf.uwsgi_stub.cache_div` (*key, value=2, expires=None, cache=None*)

Divides the specified key value by the specified value.

- <http://uwsgi.readthedocs.io/en/latest/Changelog-1.9.9.html#math-for-cache>

**Parameters**

- **key** (*str/unicode*) –
- **value** (*int*) –
- **expires** (*int*) – Expire timeout (seconds).
- **cache** (*str/unicode*) – Cache name with optional address (if @-syntax is used).

**Return type** bool

`uwsgi.conf.uwsgi_stub.cache_exists` (*key, cache=None*)

Checks whether there is a value in the cache associated with the given key.

**Parameters**

- **key** (*str/unicode*) – The cache key to check.
- **cache** (*str/unicode*) – Cache name with optional address (if @-syntax is used).

**Return type** bool

`uwsgi.conf.uwsgi_stub.cache_get` (*key, cache=None*)

Gets a value from the cache.

**Warning:** Bytes are returned for Python 3.

**Parameters**

- **key** (*str/unicode*) – The cache key to get value for.
- **cache** (*str/unicode*) – Cache name with optional address (if @-syntax is used).

**Return type** bytes|str|None

`uwsgi.conf.uwsgi_stub.cache_inc` (*key, value=1, expires=None, cache=None*)

Increments the specified key value by the specified value.

- <http://uwsgi.readthedocs.io/en/latest/Changelog-1.9.9.html#math-for-cache>

**Parameters**

- **key** (*str/unicode*) –
- **value** (*int*) –
- **expires** (*int*) – Expire timeout (seconds).
- **cache** (*str/unicode*) – Cache name with optional address (if @-syntax is used).

**Return type** booluwsgi.conf.uwsgi\_stub.**cache\_keys** (*cache=None*)

Returns a list of keys available in cache.

**Parameters** **cache** (*str/unicode*) – Cache name with optional address (if @-syntax is used).**Return type** list**Raises** **ValueError** – If cache is unavailable.uwsgi.conf.uwsgi\_stub.**cache\_mul** (*key, value=2, expires=None, cache=None*)

Multiplies the specified key value by the specified value.

- <http://uwsgi.readthedocs.io/en/latest/Changelog-1.9.9.html#math-for-cache>

**Parameters**

- **key** (*str/unicode*) –
- **value** (*int*) –
- **expires** (*int*) – Expire timeout (seconds).
- **cache** (*str/unicode*) – Cache name with optional address (if @-syntax is used).

**Return type** booluwsgi.conf.uwsgi\_stub.**cache\_num** (*key, cache=None*)

Gets the 64bit number from the specified item.

- <http://uwsgi.readthedocs.io/en/latest/Changelog-1.9.9.html#math-for-cache>

**Parameters**

- **key** (*str/unicode*) – The cache key to get value for.
- **cache** (*str/unicode*) – Cache name with optional address (if @-syntax is used).

**Return type** int|long|Noneuwsgi.conf.uwsgi\_stub.**cache\_set** (*key, value, expires=None, cache=None*)

Sets the specified key value.

**Parameters**

- **key** (*str/unicode*) –
- **value** (*str/unicode*) –
- **expires** (*int*) – Expire timeout (seconds).
- **cache** (*str/unicode*) – Cache name with optional address (if @-syntax is used).

**Return type** bool

`uwsgi.conf.uwsgi_stub.cache_update` (*key*, *value*, *expires=None*, *cache=None*)  
Updates the specified key value.

**Parameters**

- **key** (*str/unicode*) –
- **value** (*int*) –
- **expires** (*int*) – Expire timeout (seconds).
- **cache** (*str/unicode*) – Cache name with optional address (if @-syntax is used).

**Return type** bool

`uwsgi.conf.uwsgi_stub.call` (*func\_name*, \**args*)  
Performs an [RPC] function call with the given arguments.

**Warning:** Bytes are returned for Python 3.

**Parameters**

- **func\_name** (*str/unicode*) – Function name to call with optional address (if @-syntax is used).
- **args** (*list[str/bytes]*) –

**Return type** byteslstr

`uwsgi.conf.uwsgi_stub.chunked_read` (*timeout*)  
Reads chunked input.

**Warning:** Bytes are returned for Python 3.

- <http://uwsgi.readthedocs.io/en/latest/Chunked.html>
- <http://uwsgi.readthedocs.io/en/latest/Changelog-1.9.13.html#chunked-input-api>

**Parameters** **timeout** (*int*) – Wait timeout (seconds).

**Return type** byteslstr

**Raises** **IOError** – If unable to receive chunked part.

`uwsgi.conf.uwsgi_stub.chunked_read_nb` ()  
Reads chunked input without blocking.

**Warning:** Bytes are returned for Python 3.

- <http://uwsgi.readthedocs.io/en/latest/Chunked.html>
- <http://uwsgi.readthedocs.io/en/latest/Changelog-1.9.13.html#chunked-input-api>

**Return type** byteslstr

**Raises** **IOError** – If unable to receive chunked part.



`uwsgi.conf.uwsgi_stub.cl()`  
Returns current post content length.

**Return type** `int`

`uwsgi.conf.uwsgi_stub.close(fd)`  
Closes the given file descriptor.

**Parameters** `fd (int)` – File descriptor.

**Return type** `None`

`uwsgi.conf.uwsgi_stub.connect(socket, timeout=0)`  
Connects to the socket.

**Parameters**

- **socket** (`str/unicode`) – Socket name.
- **timeout** (`int`) – Timeout (seconds).

**Return type** `int`

`uwsgi.conf.uwsgi_stub.connection_fd()`  
Returns current request file descriptor.

**Return type** `int`

`uwsgi.conf.uwsgi_stub.disconnect()`  
Drops current connection.

**Return type** `None`

`uwsgi.conf.uwsgi_stub.embedded_data(symbol_name)`  
Reads a symbol from the uWSGI binary image.

- <http://uwsgi.readthedocs.io/en/latest/Embed.html>

**Warning:** Bytes are returned for Python 3.

**Parameters** `symbol_name (str/unicode)` – The symbol name to extract.

**Return type** `bytes`

**Raises** `ValueError` – If symbol is unavailable.

`uwsgi.conf.uwsgi_stub.extract(fname)`  
Extracts file contents.

**Warning:** Bytes are returned for Python 3.

**Parameters** `fname (str/unicode)` –

**Return type** `bytes`

`uwsgi.conf.uwsgi_stub.farm_get_msg()`  
Reads a mule farm message.

- <http://uwsgi.readthedocs.io/en/latest/Embed.html>

**Warning:** Bytes are returned for Python 3.

**Return type** bytes|str|None

**Raises** **ValueError** – If not in a mule

`uwsgi.conf.uwsgi_stub.farm_msg` (*farm, message*)

Sends a message to the given farm.

**Parameters**

- **farm** (*str/unicode*) – Farm name to send message to.
- **message** (*str/unicode*) –

**Return type** None

`uwsgi.conf.uwsgi_stub.get_logvar` (*name*)

Return user-defined log variable contents.

- <http://uwsgi.readthedocs.io/en/latest/LogFormat.html#user-defined-logvars>

**Warning:** Bytes are returned for Python 3.

**Parameters** **name** (*str/unicode*) –

**Return type** bytes|str

`uwsgi.conf.uwsgi_stub.green_schedule` ()

Switches to another green thread.

---

**Note:** Alias for `suspend`.

---

- <http://uwsgi.readthedocs.io/en/latest/Async.html#suspend-resume>

**Type** bool

`uwsgi.conf.uwsgi_stub.i_am_the_spooler` ()

Returns flag indicating whether you are the Spooler.

**Return type** bool

`uwsgi.conf.uwsgi_stub.in_farm` (*name*)

Returns flag indicating whether you are (mule) belong to the given farm.

**Parameters** **name** (*str/unicode*) – Farm name.

**Return type** bool

`uwsgi.conf.uwsgi_stub.is_a_reload` ()

Returns flag indicating whether reloading mechanics is used.

**Return type** bool

`uwsgi.conf.uwsgi_stub.is_connected` (*fd*)

Checks the given file descriptor.

**Parameters** `fd (int)` – File descriptor

**Return type** bool

`uwsgi.conf.uwsgi_stub.is_locked (lock_num=0)`  
Checks for the given lock.

---

**Note:** Lock 0 is always available.

---

**Parameters** `lock_num (int)` – Lock number.

**Return type** bool

**Raises** `ValueError` – For Spooler or invalid lock number

`uwsgi.conf.uwsgi_stub.listen_queue (socket_num=0)`  
Returns listen queue (backlog size) of the given socket.

**Parameters** `socket_num (int)` – Socket number.

**Return type** int

**Raises** `ValueError` – If socket is not found

`uwsgi.conf.uwsgi_stub.lock (lock_num=0)`  
Sets the given lock.

---

**Note:** Lock 0 is always available.

---

**Parameters** `lock_num (int)` – Lock number.

**Return type** None

**Raises** `ValueError` – For Spooler or invalid lock number

`uwsgi.conf.uwsgi_stub.log (message)`  
Logs a message.

**Parameters** `message (str/unicode)` –

**Return type** bool

`uwsgi.conf.uwsgi_stub.log_this_request ()`  
Instructs uWSGI to log current request data.

**Return type** None

`uwsgi.conf.uwsgi_stub.logsize ()`  
Returns current log size.

**Return type** long

`uwsgi.conf.uwsgi_stub.loop ()`  
Returns current event loop name or None if loop is not set.

**Return type** `stlunicode|None`

`uwsgi.conf.uwsgi_stub.masterpid ()`  
Return the process identifier (PID) of the uWSGI master process.

**Return type** int

`uwsgi.conf.uwsgi_stub.mem()`

Returns memory usage tuple of ints: (rss, vsz).

**Return type** tuple

`uwsgi.conf.uwsgi_stub.metric_dec(key, value=1)`

Decrements the specified metric key value by the specified value.

**Parameters**

- **key** (*str/unicode*) –
- **value** (*int*) –

**Return type** bool

`uwsgi.conf.uwsgi_stub.metric_div(key, value=1)`

Divides the specified metric key value by the specified value.

**Parameters**

- **key** (*str/unicode*) –
- **value** (*int*) –

**Return type** bool

`uwsgi.conf.uwsgi_stub.metric_get(key)`

Returns metric value by key.

**Parameters** **key** (*str/unicode*) –

**Return type** int/long

`uwsgi.conf.uwsgi_stub.metric_inc(key, value=1)`

Increments the specified metric key value by the specified value.

**Parameters**

- **key** (*str/unicode*) –
- **value** (*int*) –

**Return type** bool

`uwsgi.conf.uwsgi_stub.metric_mul(key, value=1)`

Multiplies the specified metric key value by the specified value.

**Parameters**

- **key** (*str/unicode*) –
- **value** (*int*) –

**Return type** bool

`uwsgi.conf.uwsgi_stub.metric_set(key, value)`

Sets metric value.

**Parameters**

- **key** (*str/unicode*) –
- **value** (*int/long*) –

**Return type** bool

`uwsgi.conf.uwsgi_stub.metric_set_max(key, value)`

Sets metric value if it is greater than the current one.

**Parameters**

- **key** (*str/unicode*) –
- **value** (*int/long*) –

**Return type** bool

`uwsgiconf.uwsgi_stub.metric_set_min(key, value)`  
Sets metric value if it is less that the current one.

**Parameters**

- **key** (*str/unicode*) –
- **value** (*int/long*) –

**Return type** bool

`uwsgiconf.uwsgi_stub.micros()`  
Returns uWSGI clock microseconds.

**Return type** long

`uwsgiconf.uwsgi_stub.mule_get_msg(signals=None, farms=None, buffer_size=65536, timeout=-1)`  
Block until a mule message is received and return it.

This can be called from multiple threads in the same programmed mule.

**Warning:** Bytes are returned for Python 3.

**Parameters**

- **signals** (*bool*) – Whether to manage signals.
- **farms** (*bool*) – Whether to manage farms.
- **buffer\_size** (*int*) –
- **timeout** (*int*) – Seconds.

**Return type** bytes/str**Raises** **ValueError** – If not in a mule.

`uwsgiconf.uwsgi_stub.mule_id()`  
Returns current mule ID. 0 if not a mule (e.g. worker).

**Return type** int

`uwsgiconf.uwsgi_stub.mule_msg(message, mule_farm=None)`  
Sends a message to a mule(s)/farm.

**Parameters**

- **message** (*str/unicode*) –
- **mule\_farm** – Mule ID, or farm name.

**Return type** bool**Raises** **ValueError** – If no mules, or mule ID or farm name is not recognized.

`uwsgi.conf.uwsgi_stub.offload(filename)`  
Offloads a file.

**Warning:** Currently not implemented.

**Parameters** `filename` –

**Return type** `byteslstr`

**Raises** `ValueError` – If unable to offload.

`uwsgi.conf.uwsgi_stub.parsefile(fpath)`  
Parses the given file.

Currently implemented only Spooler file parsing.

**Parameters** `fpath` (`str/unicode`) –

**Return type** `None`

`uwsgi.conf.uwsgi_stub.ready()`  
Returns flag indicating whether we are ready to handle requests.

**Return type** `bool`

`uwsgi.conf.uwsgi_stub.ready_fd()`  
Returns flag indicating whether file description related to request is ready.

**Return type** `bool`

`uwsgi.conf.uwsgi_stub.recv(fd, maxsize=4096)`  
Reads data from the given file descriptor.

**Warning:** Bytes are returned for Python 3.

**Parameters**

- `fd` (`int`) –
- `maxsize` (`int`) – Chunk size (bytes).

**Return type** `byteslstr`

`uwsgi.conf.uwsgi_stub.register_rpc(name, func)`  
Registers RPC function.

- <http://uwsgi.readthedocs.io/en/latest/RPC.html>

**Parameters**

- `name` (`str/unicode`) –
- `func` (`callable`) –

**Return type** `bool`

**Raises** `ValueError` – If unable to register function

`uwsgi.conf.uwsgi_stub.register_signal(number, target, func)`  
Registers a signal handler.

**Parameters**

- **number** (*int*) – Signal number.
- **target** (*str/unicode*) –
  - `workers` - run the signal handler on all the workers
  - `workerN` - run the signal handler only on worker N
  - `worker/worker0` - run the signal handler on the first available worker
  - `active-workers` - run the signal handlers on all the active [non-cheaped] workers
  - `mules` - run the signal handler on all of the mules
  - `muleN` - run the signal handler on mule N
  - `mule/mule0` - run the signal handler on the first available mule
  - `spooler` - run the signal on the first available spooler
  - `farmN/farm_XXX` - run the signal handler in the mule farm N or named XXX
  - <http://uwsgi.readthedocs.io/en/latest/Signals.html#signals-targets>
- **func** (*callable*) –

**Return type** None**Raises** **ValueError** – If unable to register

`uwsgiconf.uwsgi_stub.reload()`  
Gracefully reloads uWSGI.

- <http://uwsgi.readthedocs.io/en/latest/Management.html#reloading-the-server>

**Return type** bool

`uwsgiconf.uwsgi_stub.request_id()`  
Returns current request number (handled by worker on core).

**Return type** int

`uwsgiconf.uwsgi_stub.route(name, args_str)`  
Registers a named route for internal routing subsystem.

**Parameters**

- **name** (*str*) – Route name
- **args\_str** – Comma-separated arguments string.

**Return type** int

`uwsgiconf.uwsgi_stub.rpc(address, func_name, *args)`  
Performs an RPC function call with the given arguments.

- <http://uwsgi.readthedocs.io/en/latest/RPC.html>

**Warning:** Bytes are returned for Python 3.

**Parameters**

- **address** (*str/unicode*) –

- **func\_name** (*str/unicode*) – Function name to call.
- **args** (*list[str/bytes]*) –

**Return type** bytes|str

**Raises** **ValueError** – If unable to call RPC function.

`uwsgi.conf.uwsgi_stub.rpc_list()`  
Returns registered RPC functions names.

**Return type** tuple

`uwsgi.conf.uwsgi_stub.send(fd_or_data, data=None)`  
Puts data into file descriptor.

- One argument. Data to write into request file descriptor.
- Two arguments. 1. File descriptor; 2. Data to write.

#### Parameters

- **df\_or\_data** –
- **data** –

**Return type** bool

`uwsgi.conf.uwsgi_stub.sendfile(fd_or_name, chunk_size=0, start_pos=0, filesize=0)`  
Runs a sendfile.

#### Parameters

- **fd\_or\_name** (*int|str/unicode*) – File path or descriptor number.
- **chunk\_size** (*int*) – Not used.
- **start\_pos** (*int*) –
- **filesize** (*int*) – Filesize. If 0 will be determined automatically.

**Return type** bool|None

`uwsgi.conf.uwsgi_stub.send_to_spooler(message=None, **kwargs)`  
Send data to the The uWSGI Spooler. Also known as `spool()`.

**Warning:** Either *message* argument should contain a dictionary this message dictionary will be constructed from *kwargs*.

#### Parameters

- **bytes] message** (*dict[bytes,)* – The message to spool. Keys and values are bytes.
- **kwargs** (*dict*) – Possible kwargs (these are also reserved *message* argument dictionary keys):
  - **spooler: The spooler (id or directory) to use.** Specify the ABSOLUTE path of the spooler that has to manage this task
  - **priority:** Number. The priority of the message. Larger - less important.



**Warning:** This works only if you enable `order_tasks` option in `spooler.set_basic_params()`.

This will be the subdirectory in the spooler directory in which the task will be placed, you can use that trick to give a good-enough prioritization to tasks.

---

**Note:** This is for systems with few resources. For better approach use multiple spoolers.

---

- **at: Unix time at which the task must be executed.** The task will not be run until the ‘at’ time is passed.
- **body: A binary body to add to the message,** in addition to the message dictionary itself. Use this key for objects bigger than 64k, the blob will be appended to the serialized uwsgi packet and passed back to the spooler function as the ‘body’ argument.

**Return type** str|None

`uwsgi.conf.uwsgi_stub.set_logvar` (*name*, *value*)  
Sets log variable.

**Parameters**

- **name** (*str|unicode*) –
- **value** (*str|unicode*) –

**Return type** None

`uwsgi.conf.uwsgi_stub.set_spooler_frequency` (*seconds*)  
Sets how often the spooler runs.

**Parameters** **seconds** (*int*) –

**Return type** bool

`uwsgi.conf.uwsgi_stub.set_user_harakiri` (*timeout=0*)  
Sets user level harakiri.

**Parameters** **timeout** (*int*) – Seconds. 0 disable timer.

**Return type** None

`uwsgi.conf.uwsgi_stub.set_warning_message` (*message*)  
Sets a warning. This will be reported by pingers.

**Parameters** **message** (*str|unicode*) –

**Return type** bool

`uwsgi.conf.uwsgi_stub.setprocname` (*name*)  
Sets current process name.

**Parameters** **name** (*str|unicode*) –

**Return type** bool

`uwsgi.conf.uwsgi_stub.signal` (*num*, *remote=""*)  
Sends the signal to master or remote.

**Parameters**

- **num** – Signal number.
- **remote** (*str/unicode*) – Remote address.

**Return type** None

**Raises**

- **ValueError** – If remote rejected the signal.
- **IOError** – If unable to deliver to remote.

`uwsgi.conf.uwsgi_stub.signal_received()`

Get the number of the last signal received.

Used in conjunction with `signal_wait`.

- <http://uwsgi-docs.readthedocs.io/en/latest/Signals.html#signal-wait-and-signal-received>

**Return type** int

`uwsgi.conf.uwsgi_stub.signal_registered(num)`

Verifies the given signal has been registered.

**Parameters** `num` (*int*) –

**Return type** bool|None

`uwsgi.conf.uwsgi_stub.signal_wait(num=None)`

Waits for the given of any signal.

Block the process/thread/async core until a signal is received. Use `signal_received` to get the number of the signal received. If a registered handler handles a signal, `signal_wait` will be interrupted and the actual handler will handle the signal.

- <http://uwsgi-docs.readthedocs.io/en/latest/Signals.html#signal-wait-and-signal-received>

**Parameters** `num` (*int*) –

**Return type** str|unicode

**Raises** **SystemError** – If something went wrong.

`uwsgi.conf.uwsgi_stub.spool(message=None, **kwargs)`

Send data to the The uWSGI Spooler. Also known as `spool()`.

<p><b>Warning:</b> Either <code>message</code> argument should contain a dictionary this message dictionary will be constructed from <code>kwargs</code>.</p>
---

**Parameters**

- **bytes** `message` (*dict [bytes,)* – The message to spool. Keys and values are bytes.
- **kwargs** (*dict*) – Possible kwargs (these are also reserved `message` argument dictionary keys):
  - **spooler:** The spooler (id or directory) to use. Specify the ABSOLUTE path of the spooler that has to manage this task
  - **priority:** Number. The priority of the message. Larger - less important.

**Warning:** This works only if you enable `order_tasks` option in `spooler.set_basic_params()`.

This will be the subdirectory in the spooler directory in which the task will be placed, you can use that trick to give a good-enough prioritization to tasks.

---

**Note:** This is for systems with few resources. For better approach use multiple spoolers.

---

- **at: Unix time at which the task must be executed.** The task will not be run until the ‘at’ time is passed.
- **body: A binary body to add to the message,** in addition to the message dictionary itself. Use this key for objects bigger than 64k, the blob will be appended to the serialized uwsgi packet and passed back to the spooler function as the ‘body’ argument.

**Return type** str|None

`uwsgi.conf.uwsgi_stub.spooler_get_task(path)`

Returns a spooler task information.

**Parameters** `path (str)` – The relative or absolute path to the task to read.

**Return type** dict|None

`uwsgi.conf.uwsgi_stub.spooler_jobs()`

Returns a list of spooler jobs (filenames in spooler directory).

**Return type** list[str]

`uwsgi.conf.uwsgi_stub.spooler_pid()`

Returns first spooler process ID

**Return type** int

`uwsgi.conf.uwsgi_stub.spooler_pids()`

Returns a list of all spooler processes IDs.

**Return type** list[int]

`uwsgi.conf.uwsgi_stub.stop()`

Stops uWSGI.

**Return type** bool|None

`uwsgi.conf.uwsgi_stub.suspend()`

Suspends handling of current coroutine/green thread and passes control to the next async core.

- <http://uwsgi.readthedocs.io/en/latest/Async.html#suspend-resume>

**Return type** bool

`uwsgi.conf.uwsgi_stub.total_requests()`

Returns the total number of requests managed so far by the pool of uWSGI workers.

**Return type** int

`uwsgi.conf.uwsgi_stub.unlock(lock_num=0)`  
Unlocks the given lock.

---

**Note:** Lock 0 is always available.

---

**Parameters** `lock_num` (*int*) – Lock number.

**Return type** None

**Raises** `ValueError` – For Spooler or invalid lock number

`uwsgi.conf.uwsgi_stub.wait_fd_read(fd, timeout=None)`

Suspends handling of the current request until there is something to be read on file descriptor.

May be called several times before yielding/suspending to add more file descriptors to the set to be watched.

- <http://uwsgi-docs.readthedocs.io/en/latest/Async.html#waiting-for-i-o>

**Parameters**

- `fd` (*int*) – File descriptor number.
- `timeout` (*int*) – Timeout. Default: infinite.

**Return type** byteslstr

**Raises** `IOError` – If unable to read.

`uwsgi.conf.uwsgi_stub.wait_fd_write(fd, timeout=None)`

Suspends handling of the current request until there is nothing more to be written on file descriptor.

May be called several times to add more file descriptors to the set to be watched.

- <http://uwsgi-docs.readthedocs.io/en/latest/Async.html#waiting-for-i-o>

**Parameters**

- `fd` (*int*) – File descriptor number.
- `timeout` (*int*) – Timeout. Default: infinite.

**Return type** byteslstr

**Raises** `IOError` – If unable to read.

`uwsgi.conf.uwsgi_stub.websocket_handshake(security_key=None, origin=None, proto=None)`

Waits for websocket handshake.

**Parameters**

- `security_key` (*str/unicode*) – Websocket security key to use.
- `origin` (*str/unicode*) – Override `Sec-WebSocket-Origin`.
- `proto` (*str/unicode*) – Override `Sec-WebSocket-Protocol`.

**Return type** None

**Raises** `IOError` – If unable to complete handshake.

`uwsgi.conf.uwsgi_stub.websocket_recv(request_context=None)`

Receives data from websocket.

**Parameters** `request_context` –

**Return type** `byteslstr`

**Raises** `IOError` – If unable to receive a message.

`uwsgiconf.uwsgi_stub.websocket_recv_nb` (*request\_context=None*)

Receives data from websocket (non-blocking variant).

**Parameters** `request_context` –

**Return type** `byteslstr`

**Raises** `IOError` – If unable to receive a message.

`uwsgiconf.uwsgi_stub.websocket_send` (*message, request\_context=None*)

Sends a message to websocket.

**Parameters**

- `message` (*str*) – data to send
- `request_context` –

**Raises** `IOError` – If unable to send a message.

`uwsgiconf.uwsgi_stub.websocket_send_binary` (*message, request\_context=None*)

Sends binary message to websocket.

**Parameters**

- `message` (*str*) – data to send
- `request_context` –

**Raises** `IOError` – If unable to send a message.

`uwsgiconf.uwsgi_stub.worker_id` ()

Returns current worker ID. 0 if not a worker (e.g. mule).

**Return type** `int`

`uwsgiconf.uwsgi_stub.workers` ()

Gets statistics for all the workers for the current server.

Returns tuple of dicts.

**Return type** `tuple[dict]`

`uwsgiconf.uwsgi_stub.i_am_the_lord` (*legion\_name*)

Returns flag indicating whether you are the lord of the given legion.

- <http://uwsgi.readthedocs.io/en/latest/Legion.html#legion-api>

**Parameters** `legion_name` (*str|unicode*) –

**Return type** `bool`

`uwsgiconf.uwsgi_stub.lord_scroll` (*legion\_name*)

Returns a Lord scroll for the Legion.

- <http://uwsgi.readthedocs.io/en/latest/Legion.html#lord-scroll-coming-soon>

**Parameters** `legion_name` (*str|unicode*) –

**Return type** `bool`

`uwsgiconf.uwsgi_stub.scrolls` (*legion<sup>\_</sup>name*)

Returns a list of Legion scrolls defined on cluster.

**Parameters** `legion_name` (*str/unicode*) –

**Return type** list

## 3.5 Contrib

Additional integrations with third parties.

### 3.5.1 Django uwsgify

`uwsgify` adds integration with Django Framework.

First add `uwsgify` into `INSTALLED_APPS`.

```
INSTALLED_APPS = [  
    ...  
    'uwsgiconf.contrib.django.uwsgify',  
    ...  
]
```

### uwsgi<sup>\_</sup>run

`uwsgi_run` management command runs uWSGI to serve your Django-based project.

```
$ ./manage.py uwsgi_run  
  
; Options are available, use --help switch to get help.  
$ ./manage.py uwsgi_run --help
```

Now your project is up and running on `http://127.0.0.1:8000`.

By default the command runs your project using some defaults, but you can configure it to your needs with the help of `uwsgicfg.py` (constructed in a usual for **uwsgi<sup>conf</sup>** manner) placed near your `manage.py`.

```
from uwsgiconf.config import configure_uwsgi  
  
def get_configurations():  
  
    from os.path import dirname, abspath, join  
    from uwsgiconf.presets.nice import PythonSection  
  
    section = PythonSection.bootstrap(  
        'http://127.0.0.1:8000',  
        wsgi_module=join(dirname(abspath(__file__)), 'wsgi.py')  
    )  
  
    ...  
  
    return section
```

(continues on next page)

(continued from previous page)

```
configure_uwsgi(get_configurations)
```

**Note: Embedding.** if you're using **pyuwsgi** having uWSGI and your entire project compiled into a single binary, and your *manage.py* is the entrypoint, use **-embedded** option: `myproject uwsgi_run --embedded`.

### uwsgi\_reload

`uwsgi_reload` management command reloads uWSGI master process, workers.

```
$ ./manage.py uwsgi_reload  
  
; Options are available, use --help switch to get help.  
$ ./manage.py uwsgi_reload --help
```

### uwsgi\_stop

`uwsgi_stop` management command allows you to shutdown uWSGI instance.

```
$ ./manage.py uwsgi_stop  
  
; Options are available, use --help switch to get help.  
$ ./manage.py uwsgi_stop --help
```

### uwsgi\_stats

`uwsgi_stats` management command allows you to dump uWSGI configuration and current stats into the log.

```
$ ./manage.py uwsgi_stats
```

### uwsgi\_log

`uwsgi_log` management command allows you to manage uWSGI log related stuff.

```
$ ./manage.py uwsgi_log --rotate  
  
; Options are available, use --help switch to get help.  
$ ./manage.py uwsgi_log --help
```

### uwsgi\_sysinit

`uwsgi_sysinit` management command allows you to generate system service configs (e.g. `systemd`) to start your Django project on system start.

```
; Dump config to file.
$ ./manage.py uwsgi_sysinit > myapp.service

; Wire up the service config into system directory and start service
$ sudo systemctl enable --now myapp.service

; Watch application log realtime
$ sudo journalctl -fu myapp.service
```

## 3.6 FAQ

### 3.6.1 How to get through

*There are some many options, how to start?*

Start with a preset configuration. For example, if you have a Python web-app try out `uwsgiconf.presets.nice.PythonSection`. Basic things to do are: define `wsgi_module` and do `.networking.register_socket`.

This should already give a more or less decent decent configuration.

After that you can skim through option groups (such as `.networking`, `.main_process`, `.workers` etc.) and deep into uWSGI abilities.

### 3.6.2 Use from virtualenv

*I have a virtualenv in venv/ directory (where I have uWSGI and uwsgi`conf`) and configuration module outside of it, how do I run it?*

You can try the following trick (from directory containing `venv/` and `uwsgiconf.py`):

```
$ venv/bin/uwsgiconf run
```

### 3.6.3 Unknown config directive

*I use PythonSection for configuration and get [strict-mode] unknown config directive: wsgi-file on start. What's that?*

**uwsgi`conf`** enables configuration options check (aka `strict-mode`) by default.

If uWSGI plugin which provides some options is not available, you'll get the message. That's because `PythonSection` by default won't instruct uWSGI to load Python plugin (since if you get uWSGI from PyPI you already have Python and a bunch of other plugins embedded, so there's no need to load them).

If you get that message most probably uWSGI is provided by your OS distribution (e.g. on Debian you'll need to install plugin packages separately from uWSGI itself).

In that case you can try to set `embedded_plugins=False` for `PythonSection` (see Quickstart example).

Another option is to quickly fire up uWSGI to check what plugins are embedded (the same can be achieved with `$ uwsgiconf probe_plugins` command).

**uwsgi`conf`** can also do it for you automatically on configuration stage:

```
Section(embedded_plugins=Section.embedded_plugins_presets.PROBE)
```



Using the above, `embedded_plugins` will be inhabited by plugins actually available in uWSGI.

## 3.7 Configuration [Dynamic/Runtime]

**uwsgiconf** comes with `runtime` package which is similar to **uwsgidecorators** but offers different abstractions to provide useful shortcuts and defaults.

Various modules from that package can be imported and used runtime to configure different aspects of **uWSGI**, such as *caching*, *locks*, *signals*, *spooler*, *rpc*, etc.

### 3.7.1 Alarms

`uwsgiconf.runtime.alarms.issue_alarm` (*name*, *message*)

Issues the given alarm with the given message.

#### Parameters

- **name** (*str/unicode*) –
- **message** (*str/unicode*) – Message to pass to alarm.

**Return type** None

### 3.7.2 Asynced

`uwsgiconf.runtime.asynced.connect` (*socket*)

Issues socket connection. And returns a file descriptor.

- <http://uwsgi.readthedocs.io/en/latest/Async.html>

**Parameters** **socket** (*str/unicode*) –

**Return type** int

`uwsgiconf.runtime.asynced.sleep` (*seconds*)

Suspends handling the current request and passes control to the next async core.

- <http://uwsgi.readthedocs.io/en/latest/Async.html>

**Parameters** **seconds** – Sleep time, in seconds.

**Rtype** None

`uwsgiconf.runtime.asynced.suspend` ()

Suspends handling of current coroutine/green thread and passes control to the next async core.

- <http://uwsgi.readthedocs.io/en/latest/Async.html#suspend-resume>

**Return type** bool

`uwsgiconf.runtime.asynced.get_loop_name` ()

Returns current event loop name or None if loop is not set.

**Return type** `stlunicode|None`

`uwsgi.conf.runtime.asynced.wait_for_fd_read(fd, timeout=None)`

Suspends handling of the current request until there is something to be read on file descriptor.

May be called several times before yielding/suspending to add more file descriptors to the set to be watched.

- <http://uwsgi-docs.readthedocs.io/en/latest/Async.html#waiting-for-i-o>

**Parameters**

- **fd** (*int*) – File descriptor number.
- **timeout** (*int*) – Timeout. Default: infinite.

**Return type** byteslstr

**Raises** **IOError** – If unable to read.

`uwsgi.conf.runtime.asynced.wait_for_fd_write(fd, timeout=None)`

Suspends handling of the current request until there is nothing more to be written on file descriptor.

May be called several times to add more file descriptors to the set to be watched.

- <http://uwsgi-docs.readthedocs.io/en/latest/Async.html#waiting-for-i-o>

**Parameters**

- **fd** (*int*) – File descriptor number.
- **timeout** (*int*) – Timeout. Default: infinite.

**Return type** byteslstr

**Raises** **IOError** – If unable to read.

### 3.7.3 Caching

```
from uwsgi.conf.runtime.caching import Cache

# We'll access preconfigured cache named `mycache`.
cache = Cache('mycache')

key_exists = 'mykey' in cache

def my_setter(key):
    if key == 'anotherkey':
        return 'yes'
    return 'no'

# Getting cached value and populating it if required in one pass:
yes_or_no = cache.get('anotherkey', setter=my_setter)
```

**class** `uwsgi.conf.runtime.caching.Cache` (*name=None, timeout=None*)

Interface for uWSGI Caching subsystem.

**Warning:** To use this helper one needs to configure cache(s) in uWSGI config beforehand.

**Parameters**

- **name** (*str/unicode*) – Cache name with optional address (if @-syntax is used).

- **timeout** (*int*) – Expire timeout (seconds). Default: 300 (5 minutes). Use 0 to not to set a timeout (not to expire).

---

**Note:** This value is ignore if cache is configured not to expire.

---

### **keys**

Returns a list of keys available in cache.

**Return type** list

**Raises** **ValueError** – If cache is unavailable.

### **clear** ()

Clears cache the cache.

### **get** (*key*, *default=None*, *as\_int=False*, *setter=None*)

Gets a value from the cache.

#### **Parameters**

- **key** (*str/unicode*) – The cache key to get value for.
- **default** – Value to return if none found in cache.
- **as\_int** (*bool*) – Return 64bit number instead of str.
- **setter** (*callable*) – Setter callable to automatically set cache value if not already cached. Required to accept a key and return a value that will be cached.

**Return type** str/unicode/int

### **set** (*key*, *value*, *timeout=None*)

Sets the specified key value.

#### **Parameters**

- **key** (*str/unicode*) –
- **value** (*str/unicode*) –
- **timeout** (*int*) – 0 to not to expire. Object default is used if not set.

**Return type** bool

### **delete** (*key*)

Deletes the given cached key from the cache.

**Parameters** **key** (*str/unicode*) – The cache key to delete.

**Return type** None

### **incr** (*key*, *delta=1*)

Increments the specified key value by the specified value.

#### **Parameters**

- **key** (*str/unicode*) –
- **delta** (*int*) –

**Return type** bool

### **decr** (*key*, *delta=1*)

Decrements the specified key value by the specified value.

**Parameters**

- **key** (*str/unicode*) –
- **delta** (*int*) –

**Return type** bool**mul** (*key, value=2*)

Multiplies the specified key value by the specified value.

**Parameters**

- **key** (*str/unicode*) –
- **value** (*int*) –

**Return type** bool**div** (*key, value=2*)

Divides the specified key value by the specified value.

**Parameters**

- **key** (*str/unicode*) –
- **value** (*int*) –

**Return type** bool

### 3.7.4 Control

```
from uwsgi.conf.runtime.control import harakiri_imposed, reload

@harakiri_imposed(1)
def doomed():
    """Master process will kill this function after 1 sec."""

# or

with harakiri_imposed(30):
    # Master will kill worker if code under that manager won't finish in 30 sec.

# We'll reload uWSGI.
reload()
```

**uwsgi.conf.runtime.control.stop()**

Stops uWSGI.

**Return type** bool|None**uwsgi.conf.runtime.control.reload()**

Gracefully reloads uWSGI.

- <http://uwsgi.readthedocs.io/en/latest/Management.html#reloading-the-server>

**Return type** bool**uwsgi.conf.runtime.control.disconnect()**

Drops current connection.

**Return type** None

`uwsgi.conf.runtime.control.set_process_name(name)`  
Sets current process name.

**Parameters** `name` (*str/unicode*) –

**Return type** `bool`

**class** `uwsgi.conf.runtime.control.harakiri_imposed(timeout)`  
Decorator and context manager. Allows temporarily setting harakiri timeout for a function or a code block.

Examples:

```
@harakiri_imposed(1)
def doomed():
    do()
```

```
with harakiri_imposed(10):
    do()
```

**Parameters** `timeout` (*int*) – Timeout (seconds) before harakiri.

### 3.7.5 Locking

```
from uwsgi.conf.runtime.locking import lock

@lock()
def locked():
    """This function will be locked with default (0) lock."""
    ...

# or

with lock(2):
    # Code under this context manager will be locked with lock 2.
    ...
```

**class** `uwsgi.conf.runtime.locking.Lock(num=0)`  
Locks related stuff.

Lock number 0 is always available. More locks need to be registered with `.config.locking.set_basic_params(count=X)` where X is the number of locks.

---

**Note:** The same lock should be released before next acquiring.

---

Can be used as context manager:

```
with Lock():
    do()
```

Can be used as a decorator:

```
@Lock()
def do():
    pass
```

**Parameters** `num` (*int*) – Lock number (0-64). 0 is always available and is used as default.

**is\_set**

“Checks whether the lock is active.

**Return type** bool

**Raises** **ValueError** – For Spooler or invalid lock number

**acquire()**

Sets the lock.

**Return type** None

**Raises** **ValueError** – For Spooler or invalid lock number

**release()**

Unlocks the lock.

**Return type** None

**Raises** **ValueError** – For Spooler or invalid lock number

`uwsgi.conf.runtime.locking.lock`

Convenience alias for `Lock`.

alias of `uwsgi.conf.runtime.locking.Lock`

### 3.7.6 Logging

`uwsgi.conf.runtime.logging.variable_set(name, value)`

Sets log variable.

**Parameters**

- **name** (*str/unicode*) –
- **value** (*str/unicode*) –

**Return type** None

`uwsgi.conf.runtime.logging.variable_get(name)`

Return user-defined log variable contents.

- <http://uwsgi.readthedocs.io/en/latest/LogFormat.html#user-defined-logvars>

<b>Warning:</b> Bytes are returned for Python 3.
--

**Parameters** **name** (*str/unicode*) –

**Return type** bytes/str

`uwsgi.conf.runtime.logging.log_message(message)`

Logs a message.

**Parameters** **message** (*str/unicode*) –

**Return type** bool

`uwsgi.conf.runtime.logging.get_current_log_size()`

Returns current log size.

**Return type** long

### 3.7.7 Monitoring

`uwsgi.conf.runtime.monitoring.register_file_monitor(filename, target=None)`

Maps a specific file/directory modification event to a signal.

#### Parameters

- **filename** (*str/unicode*) – File or a directory to watch for its modification.
- **target** (*int/Signal/str/unicode*) – Existing signal to raise or Signal Target to register signal implicitly.

Available targets:

- `workers` - run the signal handler on all the workers
- `workerN` - run the signal handler only on worker N
- `worker/worker0` - run the signal handler on the first available worker
- `active-workers` - run the signal handlers on all the active [non-cheaped] workers
- `mules` - run the signal handler on all of the mules
- `muleN` - run the signal handler on mule N
- `mule/mule0` - run the signal handler on the first available mule
- `spooler` - run the signal on the first available spooler
- `farmN/farm_XXX` - run the signal handler in the mule farm N or named XXX

**Raises ValueError** – If unable to register monitor.

**class** `uwsgi.conf.runtime.monitoring.Metric(name)`

User metric related stuff.

**Parameters** **name** (*str/unicode*) – Metric name.

**value**

Current metric value.

**Return type** `intllong`

**set** (*value, mode=None*)

Sets metric value.

**Parameters**

- **value** (*int/long*) – New value.
- **mode** (*str/unicode*) – Update mode.
  - `None` - Unconditional update.
  - `max` - Sets metric value if it is greater that the current one.
  - `min` - Sets metric value if it is less that the current one.

**Return type** `bool`

**incr** (*delta=1*)

Increments the specified metric key value by the specified value.

**Parameters** **delta** (*int*) –

**Return type** `bool`

**decr** (*delta=1*)

Decrements the specified metric key value by the specified value.

**Parameters** **delta** (*int*) –

**Return type** bool

**mul** (*value=1*)

Multiplies the specified metric key value by the specified value.

**Parameters** **value** (*int*) –

**Return type** bool

**div** (*value=1*)

Divides the specified metric key value by the specified value.

**Parameters** **value** (*int*) –

**Return type** bool

### 3.7.8 Mules

```
from uwsgi.conf.runtime.mules import Mule, Farm

first_mule = Mule(1)

@first_mule.offload()
def for_mule(*args, **kwargs):
    # This function will be offloaded to and handled by mule 1.
    ...

farm_two = Farm('two')

@farm_two.offload()
def for_farm(*args, **kwargs):
    # And this one will be offloaded to farm `two` and handled by any mule from that_
    ↪ farm.
    ...
```

`uwsgi.conf.runtime.mules.mule_offload` (*mule\_or\_farm=None*)

Decorator. Use to offload function execution to a mule or a farm.

**Parameters** **mule\_or\_farm** (*int|str|Mule|Farm*) –

**class** `uwsgi.conf.runtime.mules.Mule` (*id*)

Represents uWSGI Mule.

**Parameters** **id** (*int*) – Mule ID

**offload** ()

Decorator. Allows to offload function execution on this mule.

```
first_mule = Mule(1)

@first_mule.offload()
def for_mule(*args, **kwargs):
    # This function will be offloaded to and handled by mule 1.
    ...
```



**classmethod** `get_current_id()`  
Returns current mule ID. Returns 0 if not a mule.

**Return type** `int`

**classmethod** `get_current()`  
Returns current mule object or None if not a mule.

**Return type** `Optional[Mule]`

**classmethod** `get_message(signals=True, farms=False, buffer_size=65536, timeout=-1)`  
Block until a mule message is received and return it.

This can be called from multiple threads in the same programmed mule.

**Parameters**

- **signals** (*bool*) – Whether to manage signals.
- **farms** (*bool*) – Whether to manage farms.
- **buffer\_size** (*int*) –
- **timeout** (*int*) – Seconds.

**Return type** `strlunicode`

**Raises** `ValueError` – If not in a mule.

**send** (*message*)  
Sends a message to a mule(s)/farm.

**Parameters** **message** (*str/unicode*) –

**Return type** `bool`

**Raises** `ValueError` – If no mules, or mule ID or farm name is not recognized.

**class** `uwsgiconf.runtime.mules.Farm(name, mules=None)`  
Represents uWSGI Mule Farm.

**Parameters**

- **name** (*str/unicode*) – Mule farm name.
- **mules** (*list[int]*) – Attached mules.

**classmethod** `get_farms()`  
Returns a list of registered farm objects.

```
farms = Farm.get_farms()
first_farm = farms[0]
first_farm_first_mule = first_farm.mules[0]
```

**Return type** `list[Farm]`

**offload** ()  
Decorator. Allows to offload function execution on mules of this farm.

```
first_mule = Farm('myfarm')

@first_mule.offload()
def for_mule(*args, **kwargs):
```

(continues on next page)

(continued from previous page)

```
# This function will be offloaded to farm `myfarm` and handled by any
↳mule from that farm.
...
```

**is\_mine**

Returns flag indicating whether current mule belongs to this farm.

**Parameters** `name` (*str/unicode*) – Farm name.**Return type** bool**classmethod** `get_message()`

Reads a mule farm message.

- <http://uwsgi.readthedocs.io/en/latest/Embed.html>

**Warning:** Bytes are returned for Python 3.

**Return type** `str/unicode/None`**Raises** `ValueError` – If not in a mule**send** (*message*)

Sends a message to the given farm.

**Parameters** `message` (*str/unicode*) –

### 3.7.9 Platform

Platform object is available in `uwsgi` module attribute:

```
from uwsgiconf.runtime.platform import uwsgi

rss, vsz = uwsgi.memory

print(uwsgi.config)

@uwsgi.postfork_hooks.add()
def db_close_connections():
    """This will be called after fork()."""
    print('Forked!')
```

**class** `uwsgiconf.runtime.platform._Platform`**request**alias of `uwsgiconf.runtime.request._Request`**postfork\_hooks**uWSGI is a preforking server, so you might need to execute a fixup tasks (hooks) after each `fork()`. Each hook will be executed in sequence on each process (worker/mule).

**Note:** The `fork()` happen before app loading, so there's no hooks for dynamic apps. But one can still move `postfork` hooks in a `.py` file and import it on server startup with `python.import_module()`.

alias of `_PostForkHooks`

**workers\_count = 0**

Number of workers (processes) currently running.

**cores\_count = 0**

Detected number of processor cores.

**buffer\_size = 0**

The current configured buffer size in bytes.

**threads\_enabled = False**

Flag indicating whether thread support is enabled.

**started\_on = 0**

uWSGI's startup Unix timestamp.

**apps\_map = None**

Applications dictionary mapping mountpoints to application callables.

**hostname**

Current host name.

**Return type** str

**config**

The current configuration options, including any custom placeholders.

**Return type** dict

**config\_variables**

Current mapping of configuration file "magic" variables.

**Return type** dict

**worker\_id**

Returns current worker ID. 0 if not a worker (e.g. mule).

**Return type** int

**workers\_info**

Gets statistics for all the workers for the current server.

Returns tuple of dicts.

**Return type** tuple[dict]

**ready\_for\_requests**

Returns flag indicating whether we are ready to handle requests.

**Return type** bool

**master\_pid**

Return the process identifier (PID) of the uWSGI master process.

**Return type** int

**memory**

Returns memory usage tuple of ints: (rss, vsz).

**Return type** tuple[int, int]

**clock**

Returns uWSGI clock microseconds.

**Return type** long

**get\_listen\_queue** (*socket\_num=0*)

Returns listen queue (backlog size) of the given socket.

**Parameters** **socket\_num** (*int*) – Socket number.

**Return type** int

**Raises** **ValueError** – If socket is not found

**get\_version** (*as\_tuple=False*)

Returns uWSGI version string or tuple.

**Parameters** **as\_tuple** (*bool*) –

**Return type** str/tuple

**class** uwsgi.conf.runtime.request.\_Request

Current request information.

**env**

Request environment dictionary.

**id**

Returns current request number (handled by worker on core).

**Return type** int

**total\_count**

Returns the total number of requests managed so far by the pool of uWSGI workers.

**Return type** int

**fd**

Returns current request file descriptor.

**Return type** int

**content\_length**

Returns current post content length.

**Return type** int/long

**log** ()

Instructs uWSGI to log current request data.

**Return type** None

**add\_var** (*name, value*)

Registers custom request variable.

Can be used for better integration with the internal routing subsystem.

**Parameters**

- **name** (*str/unicode*) –
- **value** (*str/unicode*) –

**Return type** bool

**Raises** **ValueError** – If buffer size is not enough.

### 3.7.10 RPC

```

from uwsgiconf.runtime.rpc import register_rpc, make_rpc_call, get_rpc_list

@register_rpc()
def expose_me(arg1, arg2=15):
    print('RPC called %s' % arg1)

make_rpc_call('expose_me', ['value1'])

all_rpc = get_rpc_list() # Registered RPC items list.

```

`uwsgiconf.runtime.rpc.register_rpc` (*name=None*)  
 Decorator. Allows registering a function for RPC.

- <http://uwsgi.readthedocs.io/en/latest/RPC.html>

```

@register_rpc()
def expose_me(arg1, arg2=15):
    print('RPC called %s' % arg1)

make_rpc_call('expose_me', ['value1'])

```

**Parameters** `name` (*str/unicode*) – RPC function name to associate with decorated function.

**Return type** callable

`uwsgiconf.runtime.rpc.make_rpc_call` (*func\_name, args=None, remote=None*)  
 Performs an RPC function call (local or remote) with the given arguments.

**Parameters**

- `func_name` (*str/unicode*) – RPC function name to call.
- `args` (*Iterable*) – Function arguments.

**Warning:** Strings are expected.

- `remote` (*str/unicode*) –

**Return type** bytes|str

**Raises** `ValueError` – If unable to call RPC function.

`uwsgiconf.runtime.rpc.get_rpc_list` ()  
 Returns registered RPC functions names.

**Return type** tuple

### 3.7.11 Scheduling

```

from uwsgiconf.runtime.scheduling import register_timer_rb, register_cron

@register_timer_rb(10, repeat=2)
def repeat_twice():
    """This function will be called twice with 10 seconds interval

```

(continues on next page)

(continued from previous page)

```

        (by default in in first available mule) using red-black tree based timer.

        """

@register_cron(day=-3, hour='10-18/2')
def do_something():
    """This will be run every 3rd day, from 10 till 18 every 2 hours."""

```

uwsgi.conf.runtime.scheduling.**register\_timer** (*period*, *target=None*)

Add timer.

Can be used as a decorator:

```

@register_timer(3)
def repeat():
    do()

```

### Parameters

- **period** (*int*) – The interval (seconds) at which to raise the signal.
- **target** (*int*/*Signal*/*str*/*unicode*) – Existing signal to raise or Signal Target to register signal implicitly.

Available targets:

- workers - run the signal handler on all the workers
- workerN - run the signal handler only on worker N
- worker/worker0 - run the signal handler on the first available worker
- active-workers - run the signal handlers on all the active [non-cheaped] workers
- mules - run the signal handler on all of the mules
- muleN - run the signal handler on mule N
- mule/mule0 - run the signal handler on the first available mule
- spooler - run the signal on the first available spooler
- farmN/farm\_XXX - run the signal handler in the mule farm N or named XXX

**Return type** boolcallable

**Raises** **ValueError** – If unable to add timer.

uwsgi.conf.runtime.scheduling.**register\_timer\_rb** (*period*, *repeat=None*, *target=None*)

Add a red-black timer (based on black-red tree).

```

@register_timer_rb(3)
def repeat():
    do()

```

### Parameters

- **period** (*int*) – The interval (seconds) at which the signal is raised.
- **repeat** (*int*) – How many times to send signal. Will stop after ther number is reached. Default: None - infinitely.

- **target** (*int/Signal/str/unicode*) – Existing signal to raise or Signal Target to register signal implicitly.

Available targets:

- `workers` - run the signal handler on all the workers
- `workerN` - run the signal handler only on worker N
- `worker/worker0` - run the signal handler on the first available worker
- `active-workers` - run the signal handlers on all the active [non-cheaped] workers
- `mules` - run the signal handler on all of the mules
- `muleN` - run the signal handler on mule N
- `mule/mule0` - run the signal handler on the first available mule
- `spooler` - run the signal on the first available spooler
- `farmN/farm_XXX` - run the signal handler in the mule farm N or named XXX

**Return type** boolcallable

**Raises** `ValueError` – If unable to add timer.

`uwsgi.conf.runtime.scheduling.register_timer_ms` (*period, target=None*)

Add a millisecond resolution timer.

```
@register_timer_ms(300)
def repeat():
    do()
```

#### Parameters

- **period** (*int*) – The interval (milliseconds) at which the signal is raised.
- **target** (*int/Signal/str/unicode*) – Existing signal to raise or Signal Target to register signal implicitly.

Available targets:

- `workers` - run the signal handler on all the workers
- `workerN` - run the signal handler only on worker N
- `worker/worker0` - run the signal handler on the first available worker
- `active-workers` - run the signal handlers on all the active [non-cheaped] workers
- `mules` - run the signal handler on all of the mules
- `muleN` - run the signal handler on mule N
- `mule/mule0` - run the signal handler on the first available mule
- `spooler` - run the signal on the first available spooler
- `farmN/farm_XXX` - run the signal handler in the mule farm N or named XXX

**Return type** boolcallable

**Raises** `ValueError` – If unable to add timer.

uwsgi.conf.runtime.scheduling.**register\_cron** (*weekday=None, month=None, day=None, hour=None, minute=None, target=None*)

Adds cron. The interface to the uWSGI signal cron facility.

```
@register_cron(hour=-3) # Every 3 hours.
def repeat():
    do()
```

---

**Note:** Arguments work similarly to a standard crontab, but instead of “\*”, use -1, and instead of “/2”, “/3”, etc. use -2 and -3, etc.

---

**Note:** Periods - rules like hour='10-18/2' (from 10 till 18 every 2 hours) - are allowed, but they are emulated by uwsgi.conf. Use strings to define periods.

Keep in mind, that your actual function will be wrapped into another one, which will check whether it is time to call your function.

---

### Parameters

- **weekday** (*int|str|unicode*) – Day of a the week number. Defaults to *each*. 0 - Sunday 1 - Monday 2 - Tuesday 3 - Wednesday 4 - Thursday 5 - Friday 6 - Saturday
- **month** (*int|str|unicode*) – Month number 1-12. Defaults to *each*.
- **day** (*int|str|unicode*) – Day of the month number 1-31. Defaults to *each*.
- **hour** (*int|str|unicode*) – Hour 0-23. Defaults to *each*.
- **minute** (*int|str|unicode*) – Minute 0-59. Defaults to *each*.
- **target** (*int|Signal|str|unicode*) – Existing signal to raise or Signal Target to register signal implicitly.

Available targets:

- *workers* - run the signal handler on all the workers
- *workerN* - run the signal handler only on worker N
- *worker/worker0* - run the signal handler on the first available worker
- *active-workers* - run the signal handlers on all the active [non-cheaped] workers
- *mules* - run the signal handler on all of the mules
- *muleN* - run the signal handler on mule N
- *mule/mule0* - run the signal handler on the first available mule
- *spooler* - run the signal on the first available spooler
- *farmN/farm\_XXX* - run the signal handler in the mule farm N or named XXX

**Return type** boolcallable

**Raises** **ValueError** – If unable to add cron rule.



### 3.7.12 Signals

`uwsgi.conf.runtime.signals.registry_signals = []`  
Registered signals.

**class** `uwsgi.conf.runtime.signals.SignalDescription` (*num, target, func*)  
Registered signal information.

**func**  
Alias for field number 2

**num**  
Alias for field number 0

**target**  
Alias for field number 1

`uwsgi.conf.runtime.signals.get_available_num()`  
Returns first available signal number.

**Return type** `int`

**Raises** `UwsgiConfException` – If no signal is available.

`uwsgi.conf.runtime.signals.get_last_received()`  
Get the last signal received.

**Return type** `Signal`

**class** `uwsgi.conf.runtime.signals.Signal` (*num=None*)  
Represents uWSGI signal.

**Warning:** If you define a new function in worker1 and register it as a signal handler, only worker1 can run it. The best way to register signals is defining them in the master (`.runtime.uwsgi.postfork_hooks.add`), so all workers see them.

```
signal = Signal()

@signal.register_handler()
def somefunc():
    pass

# or the same:

@signal
def somefunc():
    pass
```

**Parameters** `num` (*int*) – Signal number (0-255).

---

**Note:** If not set it will be chosen automatically.

---

**registered**  
Whether the signal is registered.

**Return type** `bool`

**register\_handler** (*target=None*)

Decorator for a function to be used as a signal handler.

```
signal = Signal()

@signal.register_handler()
def somefunc():
    pass
```

**Parameters** **target** (*str/unicode*) – Where this signal will be delivered to. Default: worker.

- workers - run the signal handler on all the workers
- workerN - run the signal handler only on worker N
- worker/worker0 - run the signal handler on the first available worker
- active-workers - run the signal handlers on all the active [non-cheaped] workers
- mules - run the signal handler on all of the mules
- muleN - run the signal handler on mule N
- mule/mule0 - run the signal handler on the first available mule
- spooler - run the signal on the first available spooler
- farmN/farm\_XXX - run the signal handler in the mule farm N or named XXX
- <http://uwsgi.readthedocs.io/en/latest/Signals.html#signals-targets>

**send** (*remote=None*)

Sends the signal to master or remote.

When you send a signal, it is copied into the master's queue. The master will then check the signal table and dispatch the messages.

**Parameters** **remote** (*str/unicode/None*) – Remote address.

**Return type** None

**Raises**

- **ValueError** – If remote rejected the signal.
- **IOError** – If unable to deliver to remote.

**wait** ()

Waits for the given of any signal.

Block the process/thread/async core until a signal is received. Use `signal_received` to get the number of the signal received. If a registered handler handles a signal, `signal_wait` will be interrupted and the actual handler will handle the signal.

- <http://uwsgi-docs.readthedocs.io/en/latest/Signals.html#signal-wait-and-signal-received>

**Raises** **SystemError** – If something went wrong.

### 3.7.13 Spooler

```

my_spooler = Spooler.get_by_basename('myspooler')

# @Spooler.task() to run on first available or to run on `my_spooler`:
@my_spooler.task(postpone=timedelta(seconds=1))
def run_me(a, b='c'):
    # We do:
    # * return True if task processed
    # * return None if task was ignored
    # * raise an exception to force task retry
    return True

# Now call this function as usual and it'll run in a spooler.
...
run_me('some', b='other')
...

```

**class** uwsgiconf.runtime.spooler.**Spooler** (*name*)  
 Gives an access to uWSGI Spooler related functions.

**Warning:** To use this helper one needs to configure spooler(s) in uWSGI config beforehand.

```

my_spooler = Spooler.get_by_basename('myspooler')

# @Spooler.task() to run on first available or to run on `my_spooler`:
@my_spooler.task(postpone=timedelta(seconds=1))
def run_me(a, b='c'):
    ...

# Now call this function as usual and it'll run in a spooler.
...
run_me('some', b='other')
...

```

**task** (*postpone=None*)  
 Decorator. Used to register a function which should be run in Spooler.

```

my_spooler = Spooler.get_by_basename('myspooler')

# @Spooler.task() to run on first available or to run on `my_spooler`:
@my_spooler.task(postpone=timedelta(seconds=1))
def run_me(a, b='c'):
    ...

```

**classmethod** **send\_message\_raw** (*message*, *spooler=None*, *priority=None*, *postpone=None*,  
*payload=None*)

Sends a message to a spooler.

#### Parameters

- **msg** (*str*) – Message to pass using spooler.
- **spooler** (*str/Spooler*) – The spooler (id or directory) to use. Specify the ABSOLUTE path of the spooler that has to manage this task
- **priority** (*int*) – Number. The priority of the message. Larger - less important.

**Warning:** This works only if you enable `order_tasks` option in `spooler.set_basic_params()`.

- **postpone** (*datetime/timedelta*) – Postpone message processing till.
- **payload** – Object to pickle and pass within message.

**Return type** str

**classmethod** `get_spoolers()`  
Returns a list of registered spoolers.

**Return type** list[*Spooler*]

**classmethod** `get_by_basename(name)`  
Returns spooler object for a given directory name.

If there is more than one spooler with the same directory base name, the first one is returned.

If not found *None* is returned.

**Parameters** `name` (*str*) – Directory base name. E.g.: ‘mydir’ to get spooler for ‘/some-where/here/is/mydir’

**Return type** Spooler|None

**classmethod** `get_pids()`  
Returns a list of all spooler processes IDs.

**Return type** list

**classmethod** `set_period(seconds)`  
Sets how often the spooler runs.

**Parameters** `seconds` (*int*) –

**Return type** bool

**classmethod** `get_tasks()`  
Returns a list of spooler jobs (filenames in spooler directory).

**Return type** list[str]

**classmethod** `read_task_file(path)`  
Returns a spooler task information.

**Parameters** `path` (*str*) – The relative or absolute path to the task to read.

**Return type** dict

**class** `uwsgiconf.runtime.spooler.TaskResult(result=None, exception=None)`  
Represents a task processing result.

**class** `uwsgiconf.runtime.spooler.ResultProcessed(result=None, exception=None)`  
Treat task as processed.

**class** `uwsgiconf.runtime.spooler.ResultSkipped(result=None, exception=None)`  
Treat task as skipped (ignored).

**class** `uwsgiconf.runtime.spooler.ResultRescheduled(result=None, exception=None)`  
Treat task as rescheduled (being due to retry).

**class** `uwsgiconf.runtime.spooler.SpoolerTask(name, message, payload)`  
Consolidates information for a spooler task.

**mark\_processed**  
alias of *ResultProcessed*

**mark\_skipped**  
alias of *ResultSkipped*

**mark\_rescheduled**  
alias of *ResultRescheduled*

**process ()**  
Processes the task.

**Supported results:**

- *None* - mark as ignored (skipped)
- *TaskResult* - result type logic
- exception - mark to retry
- other - mark as processed

**Return type** TaskResult|None|bool

**class** uwsgi.conf.runtime.spooler.**SpoolerFunctionCallTask** (*name, message, payload*)  
Function call type. Allows delegating function calls to spoolers.

**process ()**  
Processes the task.

**Supported results:**

- *None* - mark as ignored (skipped)
- *TaskResult* - result type logic
- exception - mark to retry
- other - mark as processed

**Return type** TaskResult|None|bool

uwsgi.conf.runtime.spooler.**spooler\_task\_types** = {'fcall': <class 'uwsgi.conf.runtime.spooler...>  
Known task types handlers will store here runtime.

## 3.8 Configuration [Static]

### 3.8.1 Configuration and Section

Configuration and Section are two types you'll have to mainly deal with.

**class** uwsgi.conf.config.**Section** (*name=None, runtime\_dir=None, project\_name=None, strict\_config=None, style\_prints=False, embedded\_plugins=None, \*\*kwargs*)

Configuration section.

Options within configuration section are gathered into groups:

- alarms
- caching

- `master_process`
- `workers`
- etc.

Next to all public methods of groups are for setting configuration parameters. Such methods return section object to allow chaining.

You can pass options group basic parameters into (the following are all the same):

- `set_basic_params()` as in `section.workers.set_basic_params(count=3)`
- `__call__` as in `section.workers(count=3)`
- section initializer using *params\_* prefixed group name:

```
Section(  
    params_workers=dict(count=3),  
)
```

### Parameters

- **`strict_config`** (*bool*) – Enable strict configuration parsing. If any unknown option is encountered in a configuration file, an error is shown and uWSGI quits.

To use placeholder variables when using strict mode, use the `set_placeholder` option.

- **`name`** (*str/unicode*) – Configuration section name.
- **`runtime_dir`** (*str/unicode*) – Directory to store runtime files. See `.replace_placeholders()`

---

**Note:** This can be used to store PID files, sockets, master FIFO, etc.

---

- **`project_name`** (*str/unicode*) – Project name (alias) to be used to differentiate projects. See `.replace_placeholders()`.
- **`style_prints`** (*bool*) – Enables styling (e.g. colouring) for `print_` family methods. Could be nice for console and distracting in logs.
- **`embedded_plugins`** (*list/callable*) – List of embedded plugins. Plugins from that list will be considered already loaded so uwsgiconf won't instruct uWSGI to load it if required.

See `.embedded_plugins_presets` for shortcuts.

---

#### Note:

- If you installed uWSGI using PyPI package there should already be basic plugins embedded.
  - If using Ubuntu distribution you have to install plugins as separate packages.
- 

- <http://uwsgi-docs.readthedocs.io/en/latest/BuildSystem.html#plugins-and-uwsgiplugin-py>

**alarms**

alias of `uwsgi.conf.options.alarms.Alarms`

**applications**

alias of `uwsgi.conf.options.applications.Applications`

**caching**

alias of `uwsgi.conf.options.caching.Caching`

**cheapening**

alias of `uwsgi.conf.options.workers_cheapening.Cheapening`

**empire**

alias of `uwsgi.conf.options.empire.Empire`

**locks**

alias of `uwsgi.conf.options.locks.Lock`

**logging**

alias of `uwsgi.conf.options.logging.Logging`

**main\_process**

alias of `uwsgi.conf.options.main_process.MainProcess`

**master\_process**

alias of `uwsgi.conf.options.master_process.MasterProcess`

**monitoring**

alias of `uwsgi.conf.options.monitoring.Monitoring`

**networking**

alias of `uwsgi.conf.options.networking.Networking`

**queue**

alias of `uwsgi.conf.options.queue.Queue`

**routing**

alias of `uwsgi.conf.options.routing.Routing`

**spooler**

alias of `uwsgi.conf.options.spooler.Spooler`

**statics**

alias of `uwsgi.conf.options.statics.Statics`

**subscriptions**

alias of `uwsgi.conf.options.subscriptions.Subscriptions`

**workers**

alias of `uwsgi.conf.options.workers.Workers`

**python**

alias of `uwsgi.conf.options.python.Python`

**class embedded\_plugins\_presets**

These are plugin presets that can be used as `embedded_plugins` values.

**BASIC** = ['ping', 'cache', 'nagios', 'rrdtool', 'carbon', 'rpc', 'corerouter', 'fas']  
Basic set of embedded plugins. This set is used in uWSGI package from PyPI.

**static PROBE** (`uwsgi_binary=None`)

This preset allows probing real uWSGI to get actual embedded plugin list.

**replace\_placeholders** (*value*)

Replaces placeholders that can be used e.g. in filepaths.

**Supported placeholders:**

- {project\_runtime\_dir}
- {project\_name}
- {runtime\_dir}

**Parameters** **value** (*str/unicode/list[str/unicode]/None*) –

**Return type** `None|str|unicodellist[str|unicode]`

**project\_name**

Project name (alias) to be used to differentiate projects. See `.replace_placeholders()`.

**Return type** `str|unicode`

**get\_runtime\_dir** (*default=True*)

Directory to store runtime files. See `.replace_placeholders()`

---

**Note:** This can be used to store PID files, sockets, master FIFO, etc.

---

**Parameters** **default** (*bool*) – Whether to return [system] default if not set.

**Return type** `str|unicode`

**set\_runtime\_dir** (*value*)

Sets user-defined runtime directory value.

**Parameters** **value** (*str/unicode*) –

**set\_basic\_params** (*strict\_config=None, \*\*kwargs*)

**Return type** *Section*

**as\_configuration** (*\*\*kwargs*)

Returns configuration object including only one (this very) section.

**Parameters** **kwargs** – Configuration objects initializer arguments.

**Return type** *Configuration*

**print\_plugins** ()

Print out enabled plugins.

**print\_stamp** ()

Prints out a stamp containing useful information, such as what and when has generated this configuration.

**print\_out** (*value, indent=None, format\_options=None, asap=False*)

Prints out the given value.

**Parameters**

- **value** –
- **indent** (*str/unicode*) –
- **format\_options** (*dict|str/unicode*) – text color
- **asap** (*bool*) – Print as soon as possible.



**print\_variables** ()

Prints out magic variables available in config files alongside with their values and descriptions. May be useful for debugging.

<http://uwsgi-docs.readthedocs.io/en/latest/Configuration.html#magic-variables>

**set\_plugins\_params** (*plugins=None, search\_dirs=None, autoload=None, required=False*)

Sets plugin-related parameters.

**Parameters**

- **plugins** (*list|str|unicode|OptionsGroup|list[OptionsGroup]*) – uWSGI plugins to load
- **search\_dirs** (*list|str|unicode*) – Directories to search for uWSGI plugins.
- **autoload** (*bool*) – Try to automatically load plugins when unknown options are found.
- **required** (*bool*) – Load uWSGI plugins and exit on error.

**set\_fallback** (*target*)

Sets a fallback configuration for section.

Re-exec uWSGI with the specified config when exit code is 1.

**Parameters** **target** (*str|unicode|Section*) – File path or Section to include.

**set\_placeholder** (*key, value*)

Placeholders are custom magic variables defined during configuration time.

---

**Note:** These are accessible, like any uWSGI option, in your application code via `.runtime.platform.uwsgi.config`.

---

**Parameters**

- **key** (*str|unicode*) –
- **value** (*str|unicode*) –

**env** (*key, value=None, unset=False, asap=False*)

Processes (sets/unsets) environment variable.

If is not given in *set* mode value will be taken from current env.

**Parameters**

- **key** (*str|unicode*) –
- **value** –
- **unset** (*bool*) – Whether to unset this variable.
- **asap** (*bool*) – If True env variable will be set as soon as possible.

**include** (*target*)

Includes target contents into config.

**Parameters** **target** (*str|unicode|Section|list*) – File path or Section to include.

**classmethod derive\_from** (*section, name=None*)

Creates a new section based on the given.

**Parameters**

- **section** (*Section*) – Section to derive from,
- **name** (*str/unicode*) – New section name.

**Return type** *Section***class vars**

The following variables also known as magic variables could be used as option values where appropriate.

- <http://uwsgi-docs.readthedocs.io/en/latest/Configuration.html#magic-variables>

**VERSION** = '%V'

uWSGI version number

**FORMAT\_ESCAPE** = '%['

ANSI escape 033. useful for printing colors

**CONF\_CURRENT\_SECTION** = '%x'

The current section identifier, eg. conf.ini:section.

**CONF\_NAME\_ORIGINAL** = '%o'

The original conf filename, as specified on the command line

**TIMESTAMP\_STARTUP\_S** = '%t'

Unix time s, gathered at instance startup.

**TIMESTAMP\_STARTUP\_MS** = '%T'

Unix time ms, gathered at instance startup

**DIR\_VASSALS** = '%v'

Vassals directory - pwd.

**HOST\_NAME** = '%h'

Host name.

**CPU\_CORES** = '%k'

Detected CPU count.

**USER\_ID** = '%u'

User ID.

**USER\_NAME** = '%U'

User name.

**GROUP\_ID** = '%g'

Use group ID.

**GROUP\_NAME** = '%G'

Use group name.

**classmethod** `get_descriptions()`

Returns variable to description mapping.

**Return type** dict**classmethod** `bootstrap(dsn, allow_shared_sockets=None, **init_kwargs)`

Constructs a section object performing it's basic (default) configuration.

**Parameters**

- **dsn** (*str/unicode/list[str/unicode]*) – Data source name, e.g: \* <http://127.0.0.1:8000> \* <https://127.0.0.1:443?cert=/here/there.crt&key=/that/my.key>

---

**Note:** Some schemas: fastcgi, http, https, raw, scgi, shared, udp, uwsgi, suwsgi, zeromq

---

- **allow\_shared\_sockets** (*bool*) – Allows using shared sockets to bind to privileged ports. If not provided automatic mode is enabled: shared are allowed if current user is not root.
- **init\_kwargs** – Additional initialization keyword arguments accepted by section type.

**Return type** *Section*

**class** uwsgi.conf.config.**Configuration** (*sections=None*, *autoinclude\_sections=False*, *alias=None*)

Configuration is comprised from one or more Sections and could be represented in format natively supported by uWSGI.

**Parameters**

- **sections** (*list [Section]*) – If not provided, empty section will be automatically generated.
- **autoinclude\_sections** (*bool*) – Whether to include in the first sections all subsequent sections.
- **alias** (*str/unicode*) – Configuration alias. This will be used in `tofile` as file name.

**format** (*do\_print=False*, *stamp=True*, *formatter='ini'*)

Applies formatting to configuration.

**Parameters**

- **do\_print** (*bool*) – Whether to print out formatted config.
- **stamp** (*bool*) – Whether to add stamp data to the first configuration section.
- **formatter** (*str/unicode*) – Formatter alias to format options. Default: ini.

**Return type** *strlunicodellist*

**print\_ini** ()

Print out this configuration as .ini.

**Return type** *strlunicode*

**tofile** (*filepath=None*)

Saves configuration into a file and returns its path.

Convenience method.

**Parameters** **filepath** (*str/unicode*) – Filepath to save configuration into. If not provided a temporary file will be automatically generated.

**Return type** *strlunicode*

**uwsgi.conf.config.configure\_uwsgi** (*configurator\_func*)

Allows configuring uWSGI using Configuration objects returned by the given configuration function.

**Parameters** **configurator\_func** (*callable*) – Function which return a list on configurations.

**Return type** *list|None*

**Returns** A list with detected configurations or `None` if called from within uWSGI (e.g. when trying to load WSGI application).

**Raises** `ConfigurationError` –

## 3.8.2 Alarms

### Alarm Types

**class** `uwsgi.conf.options.alarms.AlarmType` (*alias*, \*args, \*\*kwargs)

**class** `uwsgi.conf.options.alarms.AlarmCommand` (*alias*, *command*)  
Run a shell command, passing info into its stdin.

**class** `uwsgi.conf.options.alarms.AlarmSignal` (*alias*, *sig\_num*)  
Raise an uWSGI signal.

**class** `uwsgi.conf.options.alarms.AlarmLog` (*alias*)  
Print line into log.

**class** `uwsgi.conf.options.alarms.AlarmMule` (*alias*, *mule\_id*)  
Send info to a mule waiting for messages.

**class** `uwsgi.conf.options.alarms.AlarmCurl` (*alias*, *url*, *method=None*, *ssl=None*,  
*ssl\_insecure=None*, *auth\_user=None*,  
*auth\_pass=None*, *timeout=None*,  
*conn\_timeout=None*, *mail\_from=None*,  
*mail\_to=None*, *subject=None*)  
Send info to a cURL-able URL.

**class** `uwsgi.conf.options.alarms.AlarmXmpp` (*alias*, *jid*, *password*, *recipients*)  
Send info via XMPP/jabber.

**class** `uwsgi.conf.options.alarms.Alarms` (\*args, \*\*kwargs)  
Alarms.

This subsystem allows the developer/sysadmin to “announce” special conditions of an app via various channels.

- <http://uwsgi-docs.readthedocs.io/en/latest/AlarmSubsystem.html>

#### **class** `alarm_types`

Alarm types available for `.register_alarm()`.

##### **command**

alias of `uwsgi.conf.options.alarms.AlarmCommand`

##### **curl**

alias of `uwsgi.conf.options.alarms.AlarmCurl`

##### **log**

alias of `uwsgi.conf.options.alarms.AlarmLog`

##### **mule**

alias of `uwsgi.conf.options.alarms.AlarmMule`

##### **signal**

alias of `uwsgi.conf.options.alarms.AlarmSignal`

##### **xmpp**

alias of `uwsgi.conf.options.alarms.AlarmXmpp`

**set\_basic\_params** (*msg\_size=None*, *cheap=None*, *anti\_loop\_timeout=None*)

**Parameters**

- **msg\_size** (*int*) – Set the max size of an alarm message in bytes. Default: 8192.
- **cheap** (*bool*) – Use main alarm thread rather than create dedicated threads for curl-based alarms
- **anti\_loop\_timeout** (*int*) – Tune the anti-loop alarm system. Default: 3 seconds.

**print\_alarms** ()

Print out enabled alarms.

**register\_alarm** (*alarm*)

Register (create) an alarm.

**Parameters** **alarm** (*AlarmType* | *list* [*AlarmType*]) – Alarm.

**alarm\_on\_log** (*alarm*, *matcher*, *skip=False*)

Raise (or skip) the specified alarm when a log line matches the specified regexp.

**Parameters**

- **alarm** (*AlarmType* | *list* [*AlarmType*]) – Alarm.
- **matcher** (*str* | *unicode*) – Regular expression to match log line.
- **skip** (*bool*) –

**alarm\_on\_fd\_ready** (*alarm*, *fd*, *message*, *byte\_count=None*)

Triggers the alarm when the specified file descriptor is ready for read.

This is really useful for integration with the Linux eventfd() facility. Pretty low-level and the basis of most of the alarm plugins.

- <http://uwsgi-docs.readthedocs.io/en/latest/Changelog-1.9.7.html#alarm-fd>

**Parameters**

- **alarm** (*AlarmType* | *list* [*AlarmType*]) – Alarm.
- **fd** (*str* | *unicode*) – File descriptor.
- **message** (*str* | *unicode*) – Message to send.
- **byte\_count** (*int*) – Files to read. Default: 1 byte.

---

**Note:** For `eventfd` set 8.

---

**alarm\_on\_queue\_full** (*alarm*)

Raise the specified alarm when the socket backlog queue is full.

**Parameters** **alarm** (*AlarmType* | *list* [*AlarmType*]) – Alarm.

**alarm\_on\_segfault** (*alarm*)

Raise the specified alarm when the segmentation fault handler is executed.

Sends a backtrace.

**Parameters** **alarm** (*AlarmType* | *list* [*AlarmType*]) – Alarm.

### 3.8.3 Applications

**class** uwsgi.conf.options.applications.**Applications** (\*args, \*\*kwargs)  
Applications.

**set\_basic\_params** (exit\_if\_none=None, max\_per\_worker=None, single\_interpreter=None, no\_default=None, manage\_script\_name=None)

#### Parameters

- **exit\_if\_none** (*bool*) – Exit if no app can be loaded.
- **max\_per\_worker** (*int*) – Set the maximum number of per-worker applications.
- **single\_interpreter** (*bool*) – Do not use multiple interpreters (where available). Some of the supported languages (such as Python) have the concept of “multiple interpreters”. By default every app is loaded in a new python interpreter (that means a pretty-well isolated namespace for each app). If you want all of the app to be loaded in the same python vm, use the this option.
- **no\_default** (*bool*) – Do not automatically fallback to default app. By default, the first loaded app is mounted as the “default one”. That app will be served when no mountpoint matches.
- **manage\_script\_name** (*bool*) – You can to instruct uWSGI to map specific apps in the so called “mountpoint” and rewrite SCRIPT\_NAME and PATH\_INFO automatically. See `.mount()`. The WSGI standard dictates that SCRIPT\_NAME is the variable used to select a specific application.

**mount** (mountpoint, app, into\_worker=False)  
Load application under mountpoint.

#### Example:

- `.mount('', 'app0.py')` – Root URL part
- `.mount('/app1', 'app1.py')` – URL part
- `.mount('/pinax/here', '/var/www/pinax/deploy/pinax.wsgi')`
- `.mount('the_app3', 'app3.py')` – Variable value: application alias (can be set by UWSGI\_APPID)
- `.mount('example.com', 'app2.py')` – Variable value: Hostname (variable set in nginx)
- <http://uwsgi-docs.readthedocs.io/en/latest/Nginx.html#hosting-multiple-apps-in-the-same-process-aka-managing-sc>

#### Parameters

- **mountpoint** (*str/unicode*) – URL part, or variable value.

---

**Note:** In case of URL part you may also want to set `manage_script_name` basic param to `True`.

---

**Warning:** In case of URL part a trailing slash may case problems in some cases (e.g. with Django based projects).

- **app** (*str/unicode*) – App module/file.

- **into\_worker** (*bool*) – Load application under mountpoint in the specified worker or after workers spawn.

**switch\_into\_lazy\_mode** (*affect\_master=None*)

Load apps in workers instead of master.

This option may have memory usage implications as Copy-on-Write semantics can not be used.

---

**Note:** Consider using `touch_chain_reload` option in `workers` basic params for lazy apps reloading.

---

**Parameters** **affect\_master** (*bool*) – If **True** only workers will be reloaded by uWSGI’s reload signals; the master will remain alive.

**Warning:** uWSGI configuration changes are not picked up on reload by the master.

### 3.8.4 Caching

**class** uwsgi.conf.options.caching.**Caching** (*\*args, \*\*kwargs*)

Caching.

uWSGI includes a very fast, all-in-memory, zero-IPC, SMP-safe, constantly-optimizing, highly-tunable, key-value store simply called “the caching framework”.

A single uWSGI instance can create an unlimited number of “caches” each one with different setup and purpose.

- <http://uwsgi-docs.readthedocs.io/en/latest/Caching.html>
- <http://uwsgi-docs.readthedocs.io/en/latest/tutorials/CachingCookbook.html>

**set\_basic\_params** (*no\_expire=None, expire\_scan\_interval=None, report\_freed=None*)

**Parameters**

- **no\_expire** (*bool*) – Disable auto sweep of expired items. Since uWSGI 1.2, cache item expiration is managed by a thread in the master process, to reduce the risk of deadlock. This thread can be disabled (making item expiry a no-op) with the this option.
- **expire\_scan\_interval** (*int*) – Set the frequency (in seconds) of cache sweeper scans. Default: 3.
- **report\_freed** (*bool*) – Constantly report the cache item freed by the sweeper.

**Warning:** Use only for debug.

**add\_item** (*key, value, cache\_name=None*)

Add an item into the given cache.

This is a commodity option (mainly useful for testing) allowing you to store an item in a uWSGI cache during startup.

**Parameters**

- **key** (*str/unicode*) –

- **value** –
- **cache\_name** (*str/unicode*) – If not set, default will be used.

**add\_file** (*filepath*, *gzip=False*, *cache\_name=None*)

Load a static file in the cache.

---

**Note:** Items are stored with the filepath as is (relative or absolute) as the key.

---

### Parameters

- **filepath** (*str/unicode*) –
- **gzip** (*bool*) – Use gzip compression.
- **cache\_name** (*str/unicode*) – If not set, default will be used.

**add\_cache** (*name*, *max\_items*, *no\_expire=None*, *store=None*, *store\_sync\_interval=None*,  
*store\_delete=None*, *hash\_algo=None*, *hash\_size=None*, *key\_size=None*,  
*udp\_clients=None*, *udp\_servers=None*, *block\_size=None*, *block\_count=None*,  
*sync\_from=None*, *mode\_bitmap=None*, *use\_lastmod=None*, *full\_silent=None*,  
*full\_purge\_lru=None*)

Creates cache. Default mode: single block.

---

**Note:** This uses new generation `cache2` option available since uWSGI 1.9.

---

---

**Note:** When at least one cache is configured without `full_purge_lru` and the master is enabled a thread named “the cache sweeper” is started. Its main purpose is deleting expired keys from the cache. If you want auto-expiring you need to enable the master.

---

### Parameters

- **name** (*str/unicode*) – Set the name of the cache. Must be unique in an instance.
- **max\_items** (*int*) – Set the maximum number of cache items.

---

**Note:** Effective number of items is **max\_items - 1** - the first item of the cache is always internally used as “NULL/None/undef”.

---

- **no\_expire** (*bool*) – If `True` cache items won’t expire even if instructed to do so by cache set method.
- **store** (*str/unicode*) – Set the filename for the persistent storage. If it doesn’t exist, the system assumes an empty cache and the file will be created.
- **store\_sync\_interval** (*int*) – Set the number of seconds after which `msync()` is called to flush memory cache on disk when in persistent mode. By default it is disabled leaving the decision-making to the kernel.
- **store\_delete** (*bool*) – uWSGI, by default, will not start if a cache file exists and the store file does not match the configured items/blocksize. Setting this option will make uWSGI delete the existing file upon mismatch and create a new one.



- **hash\_algo** (*str/unicode*) – Set the hash algorithm used in the hash table. Current options are:
  - djb33x (default)
  - murmur2
- **hash\_size** (*int*) – This is the size of the hash table in bytes. Generally 65536 (the default) is a good value.

---

**Note:** Change it only if you know what you are doing or if you have a lot of collisions in your cache.

---

- **key\_size** (*int*) – Set the maximum size of a key, in bytes. Default: 2048.
- **udp\_clients** (*str/unicode/list*) – List of UDP servers which will receive UDP cache updates.
- **|list udp\_servers** (*str/unicode*) – List of UDP addresses on which to bind the cache to wait for UDP updates.
- **block\_size** (*int*) – Set the size (in bytes) of a single block.

---

**Note:** It's a good idea to use a multiple of 4096 (common memory page size).

---

- **block\_count** (*int*) – Set the number of blocks in the cache. Useful only in bitmap mode, otherwise the number of blocks is equal to the maximum number of items.
- **sync\_from** (*str/unicode/list*) – List of uWSGI addresses which the cache subsystem will connect to for getting a full dump of the cache. It can be used for initial cache synchronization. The first node sending a valid dump will stop the procedure.
- **mode\_bitmap** (*bool*) – Enable (more versatile but relatively slower) bitmap mode.

<http://uwsgi-docs.readthedocs.io/en/latest/Caching.html#single-block-faster-vs-bitmaps-slower>

**Warning:** Considered production ready only from uWSGI 2.0.2.

- **use\_lastmod** (*bool*) – Enabling will update `last_modified_at` timestamp of each cache on every cache item modification. Enable it if you want to track this value or if other features depend on it. This value will then be accessible via the stats socket.
- **full\_silent** (*bool*) – By default uWSGI will print warning message on every cache set operation if the cache is full. To disable this warning set this option.

---

**Note:** Available since 2.0.4.

---

- **full\_purge\_lru** (*bool*) – Allows the caching framework to evict Least Recently Used (LRU) item when you try to add new item to cache storage that is full.

---

**Note:** `no_expire` argument will be ignored.

---

### 3.8.5 Empire

**class** uwsgi.conf.options.empire.**Empire** (\*args, \*\*kwargs)

Emperor and his vassals.

If you need to deploy a big number of apps on a single server, or a group of servers, the Emperor mode is just the ticket.

- <http://uwsgi-docs.readthedocs.io/en/latest/Emperor.html>

**set\_emperor\_params** (vassals\_home=None, name=None, scan\_interval=None, pid\_file=None, spawn\_asap=None, stats\_address=None, trigger\_socket=None, links\_no\_follow=None)

---

**Note:** The emperor should generally not be run with master, unless master features like advanced logging are specifically needed.

---

---

**Note:** The emperor should generally be started at server boot time and left alone, not reloaded/restarted except for uWSGI upgrades; emperor reloads are a bit drastic, reloading all vassals at once. Instead vassals should be reloaded individually when needed, in the manner of the imperial monitor in use.

---

#### Parameters

- **vassals\_home** (*str|unicode|list[str|unicode]*) – Set vassals home and enable Emperor mode.
- **name** (*str|unicode*) – Set the Emperor process name.
- **scan\_interval** (*int*) – Set the Emperor scan frequency. Default: 3 seconds.
- **pid\_file** (*str|unicode*) – Write the Emperor pid in the specified file.
- **spawn\_asap** (*bool*) – Spawn the Emperor as soon as possible.
- **stats\_address** (*str|unicode*) – Run the Emperor stats server on specified address.
- **trigger\_socket** (*str|unicode*) – Enable the Emperor trigger socket.
- **links\_no\_follow** (*bool*) – Do not follow symlinks when checking for mtime.

**print\_monitors** ()

Print out enabled imperial monitors.

**set\_emperor\_command\_params** (command\_socket=None, wait\_for\_command=None, wait\_for\_command\_exclude=None)

Emperor commands related parameters.

- <http://uwsgi-docs.readthedocs.io/en/latest/tutorials/EmperorSubscriptions.html>

#### Parameters

- **command\_socket** (*str|unicode*) – Enable the Emperor command socket. It is a channel allowing external process to govern vassals.
- **wait\_for\_command** (*bool*) – Always wait for a ‘spawn’ Emperor command before starting a vassal.

- **wait\_for\_command\_exclude** (*str/unicode/list[str/unicode]*) – Vassals that will ignore `wait_for_command`.

**set\_vassals\_wrapper\_params** (*wrapper=None, overrides=None, fallbacks=None*)

Binary wrapper for vassals parameters.

#### Parameters

- **wrapper** (*str/unicode*) – Set a binary wrapper for vassals.
- **overrides** (*str/unicode/list[str/unicode]*) – Set a binary wrapper for vassals to try before the default one
- **fallbacks** (*str/unicode/list[str/unicode]*) – Set a binary wrapper for vassals to try as a last resort. Allows you to specify an alternative binary to execute when running a vassal and the default `binary_path` is not found (or returns an error).

**set\_throttle\_params** (*level=None, level\_max=None*)

Throttling options.

- <http://uwsgi-docs.readthedocs.io/en/latest/Emperor.html#throttling>
- <http://uwsgi-docs.readthedocs.io/en/latest/Emperor.html#loyalty>

#### Parameters

- **level** (*int*) – Set throttling level (in milliseconds) for bad behaving vassals. Default: 1000.
- **level\_max** (*int*) – Set maximum throttling level (in milliseconds) for bad behaving vassals. Default: 3 minutes.

**set\_tolerance\_params** (*for\_heartbeat=None, for\_cursed\_vassals=None*)

Various tolerance options.

#### Parameters

- **for\_heartbeat** (*int*) – Set the Emperor tolerance about heartbeats.
  - <http://uwsgi-docs.readthedocs.io/en/latest/Emperor.html#heartbeat-system>
- **for\_cursed\_vassals** (*int*) – Set the Emperor tolerance about cursed vassals.
  - <http://uwsgi-docs.readthedocs.io/en/latest/Emperor.html#blacklist-system>

**set\_mode\_tyrant\_params** (*enable=None, links\_no\_follow=None, use\_initgroups=None*)

Tyrant mode (secure multi-user hosting).

In Tyrant mode the Emperor will run the vassal using the UID/GID of the vassal configuration file.

- <http://uwsgi-docs.readthedocs.io/en/latest/Emperor.html#tyrant-mode-secure-multi-user-hosting>

#### Parameters

- **enable** – Puts the Emperor in Tyrant mode.
- **links\_no\_follow** (*bool*) – Do not follow symlinks when checking for uid/gid in Tyrant mode.
- **use\_initgroups** (*bool*) – Add additional groups set via `initgroups()` in Tyrant mode.

**set\_mode\_broodlord\_params** (*zerg\_count=None, vassal\_overload\_sos\_interval=None, vassal\_queue\_items\_sos=None*)

This mode is a way for a vassal to ask for reinforcements to the Emperor.

Reinforcements are new vassals spawned on demand generally bound on the same socket.

**Warning:** If you are looking for a way to dynamically adapt the number of workers of an instance, check the Cheaper subsystem - adaptive process spawning mode.

*Broodlord mode is for spawning totally new instances.*

#### Parameters

- **zerg\_count** (*int*) – Maximum number of zergs to spawn.
- **vassal\_overload\_sos\_interval** (*int*) – Ask emperor for reinforcement when overloaded. Accepts the number of seconds to wait between asking for a new reinforcements.
- **vassal\_queue\_items\_sos** (*int*) – Ask emperor for sos if listen queue (backlog) has more items than the value specified

### 3.8.6 Locks

**class** uwsgi.conf.options.locks.**Locks** (*\*args, \*\*kwargs*)

Locks.

- <http://uwsgi.readthedocs.io/en/latest/Locks.html>

**set\_basic\_params** (*count=None, thunder\_lock=None, lock\_engine=None*)

#### Parameters

- **count** (*int*) – Create the specified number of shared locks.
- **thunder\_lock** (*bool*) – Serialize accept() usage (if possible) Could improve performance on Linux with robust pthread mutexes.

<http://uwsgi.readthedocs.io/en/latest/articles/SerializingAccept.html>

- **lock\_engine** (*str|unicode*) – Set the lock engine.

#### Example:

– ipcsem

**set\_ipcsem\_params** (*ftok=None, persistent=None*)

Sets ipcsem lock engine params.

#### Parameters

- **ftok** (*str|unicode*) – Set the ipcsem key via ftok() for avoiding duplicates.
- **persistent** (*bool*) – Do not remove ipcsem's on shutdown.

**lock\_file** (*fpath, after\_setup=False, wait=False*)

Locks the specified file.

#### Parameters

- **fpath** (*str|unicode*) – File path.
- **after\_setup** (*bool*) – True - after logging/daemon setup False - before starting

- **wait** (*bool*) – True - wait if locked False - exit if locked

### 3.8.7 Logging

#### Loggers

**class** uwsgiconf.options.logging\_loggers.**Logger** (*alias, \*args*)

**class** uwsgiconf.options.logging\_loggers.**LoggerFile** (*filepath, alias=None*)  
Allows logging into files.

##### Parameters

- **filepath** (*str/unicode*) – File path.
- **alias** (*str/unicode*) – Logger alias.

**class** uwsgiconf.options.logging\_loggers.**LoggerSocket** (*addr\_or\_path, alias=None*)  
Allows logging into UNIX and UDP sockets.

##### Parameters

- **addr\_or\_path** (*str/unicode*) – Remote address or filepath.

##### Examples:

- /tmp/uwsgi.logsock
- 192.168.173.19:5050

- **alias** (*str/unicode*) – Logger alias.

**class** uwsgiconf.options.logging\_loggers.**LoggerSyslog** (*app\_name=None, facility=None, alias=None*)  
Allows logging into Unix standard syslog.

##### Parameters

- **app\_name** (*str/unicode*) –
- **facility** (*str/unicode*) –  
– <https://en.wikipedia.org/wiki/Syslog#Facility>
- **alias** (*str/unicode*) – Logger alias.

**class** uwsgiconf.options.logging\_loggers.**LoggerRsyslog** (*app\_name=None, host=None, facility=None, split=None, packet\_size=None, alias=None*)  
Allows logging into Unix standard syslog or a remote syslog.

##### Parameters

- **app\_name** (*str/unicode*) –
- **host** (*str/unicode*) – Address (host and port) or UNIX socket path.
- **facility** (*str/unicode*) –  
– <https://en.wikipedia.org/wiki/Syslog#Facility>
- **split** (*bool*) – Split big messages into multiple chunks if they are bigger than allowed packet size. Default: False.

- **packet\_size** (*int*) – Set maximum packet size for syslog messages. Default: 1024.

**Warning:** using packets > 1024 breaks RFC 3164 (#4.1)

- **alias** (*str/unicode*) – Logger alias.

**class** uwsgi.conf.options.logging\_loggers.**LoggerRedis** (*host=None, command=None, prefix=None, alias=None*)

Allows logging into Redis.

---

**Note:** Consider using `dedicate_thread` param.

---

#### Parameters

- **host** (*str/unicode*) – Default: 127.0.0.1:6379
- **command** (*str/unicode*) – Command to be used. Default: publish uwsgi

#### Examples:

- publish foobar
- rpush foo
- **prefix** (*str/unicode*) – Default: <empty>
- **alias** (*str/unicode*) – Logger alias.

**class** uwsgi.conf.options.logging\_loggers.**LoggerMongo** (*host=None, collection=None, node=None, alias=None*)

Allows logging into Mongo DB.

---

**Note:** Consider using `dedicate_thread` param.

---

#### Parameters

- **host** (*str/unicode*) – Default: 127.0.0.1:27017
- **collection** (*str/unicode*) – Command to be used. Default: uwsgi.logs
- **node** (*str/unicode*) – An identification string for the instance sending logs Default: <server hostname>
- **alias** (*str/unicode*) – Logger alias.

**class** uwsgi.conf.options.logging\_loggers.**LoggerZeroMq** (*connection\_str, alias=None*)

Allows logging into ZeroMQ sockets.

#### Parameters

- **connection\_str** (*str/unicode*) –

#### Examples:

- tcp://192.168.173.18:9191
- **alias** (*str/unicode*) – Logger alias.

## Encoders

```

class uwsgi.conf.options.logging_encoders.Encoder (*args)
class uwsgi.conf.options.logging_encoders.EncoderPrefix (value)
    Add a raw prefix to each log msg.

    Parameters value (str/unicode) – Value to be used as affix
class uwsgi.conf.options.logging_encoders.EncoderSuffix (value)
    Add a raw suffix to each log msg

    Parameters value (str/unicode) – Value to be used as affix
class uwsgi.conf.options.logging_encoders.EncoderNewline (*args)
    Add a newline char to each log msg.
class uwsgi.conf.options.logging_encoders.EncoderGzip (*args)
    Compress each msg with gzip (requires zlib).
class uwsgi.conf.options.logging_encoders.EncoderCompress (*args)
    Compress each msg with zlib compress (requires zlib).
class uwsgi.conf.options.logging_encoders.EncoderFormat (template)
    Apply the specified format to each log msg.

    Parameters template (str/unicode) – Template string. Available variables are listed in
        FormatEncoder.Vars.

class vars
    Variables available to use.

    MESSAGE = '${msg}'
        Raw log message (newline stripped).

    MESSAGE_NEWLINE = '${msgnl}'
        Raw log message (with newline).

    TIME = '${unix}'
        Current unix time.

    TIME_MS = '${micros}'
        Current unix time in microseconds.
class uwsgi.conf.options.logging_encoders.EncoderJson (template)
    Apply the specified format to each log msg with each variable json escaped.

    Parameters template (str/unicode) – Template string. Available variables are listed in
        FormatEncoder.Vars.
class uwsgi.conf.options.logging.Var (name)
class uwsgi.conf.options.logging.VarMetric (name)
class uwsgi.conf.options.logging.VarRequestVar (name)
class uwsgi.conf.options.logging.Logging (*args, **kwargs)
    Logging.

    • http://uwsgi.readthedocs.io/en/latest/Logging.html
    • http://uwsgi-docs.readthedocs.io/en/latest/LogFormat.html

class loggers
    Loggers available for add_logger().

```

**file**

alias of `uwsgi.conf.options.logging_loggers.LoggerFile`

**mongo**

alias of `uwsgi.conf.options.logging_loggers.LoggerMongo`

**redis**

alias of `uwsgi.conf.options.logging_loggers.LoggerRedis`

**socket**

alias of `uwsgi.conf.options.logging_loggers.LoggerSocket`

**syslog**

alias of `uwsgi.conf.options.logging_loggers.LoggerSyslog`

**rsyslog**

alias of `uwsgi.conf.options.logging_loggers.LoggerRsyslog`

**zeromq**

alias of `uwsgi.conf.options.logging_loggers.LoggerZeroMq`

**class encoders**

Loggers available for `add_logger_encoder()`.

**compress**

alias of `uwsgi.conf.options.logging_encoders.EncoderCompress`

**format**

alias of `uwsgi.conf.options.logging_encoders.EncoderFormat`

**gzip**

alias of `uwsgi.conf.options.logging_encoders.EncoderGzip`

**json**

alias of `uwsgi.conf.options.logging_encoders.EncoderJson`

**newline**

alias of `uwsgi.conf.options.logging_encoders.EncoderNewline`

**prefix**

alias of `uwsgi.conf.options.logging_encoders.EncoderPrefix`

**suffix**

alias of `uwsgi.conf.options.logging_encoders.EncoderSuffix`

**set\_basic\_params** (*no\_requests=None, template=None, memory\_report=None, prefix=None, prefix\_date=None, apply\_strftime=None, response\_ms=None, ip\_x\_forwarded=None*)

**Parameters**

- **no\_requests** (*bool*) – Disable requests logging - only uWSGI internal messages and errors will be logged.
- **template** (*str/unicode*) – Set advanced format for request logging. This template string can use variables from `Logging.Vars`.
- **prefix** (*str/unicode*) – Prefix log items with a string.
- **prefix\_date** (*str/unicode/bool*) – Prefix log items with date string.

---

**Note:** This can be `True` or contain formatting placeholders (e.g. `%Y-%m-%d %H:%M:%S`) if used with `apply_strftime`.

---



- **memory\_report** (*int*) – Enable memory report. \* 1 - basic (default); \* 2 - uss/pss (Linux only)
- **apply\_strftime** (*bool*) – Apply strftime to dates in log entries. E.g. `prefix_date` can contain format placeholders. See also `vars.REQ_START_FORMATTED`.
- **response\_ms** (*bool*) – Report response time in microseconds instead of milliseconds.
- **ip\_x\_forwarded** (*bool*) – Use the IP from X-Forwarded-For header instead of `REMOTE_ADDR`. Used when uWSGI is run behind multiple proxies.

**log\_into** (*target, before\_priv\_drop=True*)  
Simple file or UDP logging.

---

**Note:** This doesn't require any Logger plugin and can be used if no log routing is required.

---

#### Parameters

- **target** (*str/unicode*) – Filepath or UDP address.
- **before\_priv\_drop** (*bool*) – Whether to log data before or after privileges drop.

**set\_file\_params** (*reopen\_on\_reload=None, truncate\_on\_startup=None, max\_size=None, rotation\_fname=None, touch\_reopen=None, touch\_rotate=None, owner=None, mode=None*)  
Set various parameters related to file logging.

#### Parameters

- **reopen\_on\_reload** (*bool*) – Reopen log after reload.
- **truncate\_on\_startup** (*bool*) – Truncate log on startup.
- **max\_size** (*int*) – Set maximum logfile size in bytes after which log should be rotated.
- **rotation\_fname** (*str/unicode*) – Set log file name after rotation.
- **touch\_reopen** (*str/unicode/list*) – Trigger log reopen if the specified file is modified/touched.

---

**Note:** This can be set to a file touched by `postrotate` script of `logrotate` to implement rotation.

---

- **touch\_rotate** (*str/unicode/list*) – Trigger log rotation if the specified file is modified/touched.
- **owner** (*str/unicode*) – Set owner `chown()` for logs.
- **mode** (*str/unicode*) – Set mode `chmod()` for logs.

**set\_filters** (*include=None, exclude=None, write\_errors=None, write\_errors\_tolerance=None, sigpipe=None*)  
Set various log data filters.

#### Parameters

- **include** (*str/unicode/list*) – Show only log lines matching the specified regexp.

---

**Note:** Requires enabled PCRE support.

---

- **exclude** (*str/unicode/list*) – Do not show log lines matching the specified regexp.

---

**Note:** Requires enabled PCRE support.

---

- **write\_errors** (*bool*) – Log (annoying) write()/writev() errors. Default: True.

---

**Note:** If both this and `sigpipe` set to False, it's the same as setting `write-errors-exception-only` uWSGI option.

---

- **write\_errors\_tolerance** (*int*) – Set the maximum number of allowed write errors before exception is raised. Default: no tolerance.

---

**Note:** Available for Python, Perl, PHP.

---

- **sigpipe** (*bool*) – Log (annoying) SIGPIPE. Default: True.

---

**Note:** If both this and `write_errors` set to False, it's the same as setting `write-errors-exception-only` uWSGI option.

---

**set\_requests\_filters** (*slower=None, bigger=None, status\_4xx=None, status\_5xx=None, no\_body=None, sendfile=None, io\_errors=None*)

Set various log data filters.

#### Parameters

- **slower** (*int*) – Log requests slower than the specified number of milliseconds.
- **bigger** (*int*) – Log requests bigger than the specified size in bytes.
- **status\_4xx** – Log requests with a 4xx response.
- **status\_5xx** – Log requests with a 5xx response.
- **no\_body** (*bool*) – Log responses without body.
- **sendfile** (*bool*) – Log sendfile requests.
- **io\_errors** (*bool*) – Log requests with io errors.

**set\_master\_logging\_params** (*enable=None, dedicate\_thread=None, buffer=None, sock\_stream=None, sock\_stream\_requests\_only=None*)

Sets logging params for delegating logging to master process.

#### Parameters

- **enable** (*bool*) – Delegate logging to master process. Delegate the write of the logs to the master process (this will put all of the logging I/O to a single process). Useful for system with advanced I/O schedulers/elevators.

- **dedicate\_thread** (*bool*) – Delegate log writing to a thread.  
As error situations could cause the master to block while writing a log line to a remote server, it may be a good idea to use this option and delegate writes to a secondary thread.
- **buffer** (*int*) – Set the buffer size for the master logger in bytes. Bigger log messages will be truncated.
- **sock\_stream** (*bool/tuple*) – Create the master logpipe as SOCK\_STREAM.
- **sock\_stream\_requests\_only** (*bool/tuple*) – Create the master requests logpipe as SOCK\_STREAM.

**print\_loggers** ()

Print out available (built) loggers.

**add\_logger** (*logger, requests\_only=False, for\_single\_worker=False*)

Set/add a common logger or a request requests only.

#### Parameters

- **logger** (*str/unicode/list/Logger/list [Logger]*) –
- **requests\_only** (*bool*) – Logger used only for requests information messages.
- **for\_single\_worker** (*bool*) – Logger to be used in single-worker setup.

**add\_logger\_route** (*logger, matcher, requests\_only=False*)

Log to the specified named logger if regexp applied on log item matches.

#### Parameters

- **logger** (*str/unicode/list/Logger/list [Logger]*) – Logger to associate route with.
- **matcher** (*str/unicode*) – Regular expression to apply to log item.
- **requests\_only** (*bool*) – Matching should be used only for requests information messages.

**add\_logger\_encoder** (*encoder, logger=None, requests\_only=False, for\_single\_worker=False*)

Add an item in the log encoder or request encoder chain.

- <http://uwsgi-docs.readthedocs.io/en/latest/LogEncoders.html>

---

**Note:** Encoders automatically enable master log handling (see `.set_master_logging_params()`).

---



---

**Note:** For best performance consider allocating a thread for log sending with `dedicate_thread`.

---

#### Parameters

- **encoder** (*str/unicode/list/Encoder*) – Encoder (or a list) to add into processing.
- **logger** (*str/unicode/Logger*) – Logger apply associate encoders to.
- **requests\_only** (*bool*) – Encoder to be used only for requests information messages.

- `for_single_worker` (*bool*) – Encoder to be used in single-worker setup.

**class vars**

Variables available for custom log formatting.

**REQ\_URI** = '%(uri)'

REQUEST\_URI from `wsgi_request` of the current request.

**REQ\_METHOD** = '%(method)'

REQUEST\_METHOD from `wsgi_request` of the current request.

**REQ\_REMOTE\_USER** = '%(user)'

REMOTE\_USER from `wsgi_request` of the current request.

**REQ\_REMOTE\_ADDR** = '%(addr)'

REMOTE\_ADDR from `wsgi_request` of the current request.

**REQ\_HTTP\_HOST** = '%(host)'

HTTP\_HOST from `wsgi_request` of the current request.

**REQ\_SERVER\_PROTOCOL** = '%(proto)'

SERVER\_PROTOCOL from `wsgi_request` of the current request.

**REQ\_USER\_AGENT** = '%(uagent)'

HTTP\_USER\_AGENT from `wsgi_request` of the current request.

**REQ\_REFERER** = '%(referer)'

HTTP\_REFERER from `wsgi_request` of the current request.

**REQ\_START\_TS** = '%(time)'

Timestamp of the start of the request. E.g.: 1512623650

**REQ\_START\_CTIME** = '%(ctime)'

Ctime of the start of the request. E.g.: Thu Dec 7 08:05:35 2017

**REQ\_START\_UNIX\_US** = '%(tmsecs)'

Timestamp of the start of the request in milliseconds since the epoch.

---

**Note:** since 1.9.21

---

**REQ\_START\_UNIX\_MS** = '%(tmicros)'

Timestamp of the start of the request in microseconds since the epoch.

---

**Note:** since 1.9.21

---

**REQ\_START\_HUMAN** = '%(ltime)'

Human-formatted (Apache style) request time.

**REQ\_START\_FORMATTED** = '%(ftime)'

Request time formatted with `apply_strftime`.

---

**Note:** Use `apply_strftime` and placeholders.

---

**REQ\_SIZE\_BODY** = '%(cl)'

Request content body size.

**REQ\_COUNT\_VARS\_CGI = '%(vars)'**  
Number of CGI vars in the request.

**REQ\_COUNT\_ERR\_READ = '%(rerr)'**  
Number of read errors for the request.

---

**Note:** since 1.9.21

---

**REQ\_COUNT\_ERR\_WRITE = '%(werr)'**  
Number of write errors for the request.

---

**Note:** since 1.9.21

---

**REQ\_COUNT\_ERR = '%(ioerr)'**  
Number of write and read errors for the request.

---

**Note:** since 1.9.21

---

**RESP\_STATUS = '%(status)'**  
HTTP response status code.

**RESP\_TIME\_US = '%(micros)'**  
Response time in microseconds. E.g.: 1512623650704

**RESP\_TIME\_MS = '%(msecs)'**  
Response time in milliseconds. E.g.: 1512623650704413

**RESP\_SIZE = '%(size)'**  
Response body size + response headers size.

**RESP\_SIZE\_HEADERS = '%(hsize)'**  
Response headers size.

**RESP\_SIZE\_BODY = '%(rsize)'**  
Response body size.

**RESP\_COUNT\_HEADERS = '%(headers)'**  
Number of generated response headers.

**TIME\_UNIX = '%(epoch)'**  
The current time in Unix format.

**WORKER\_PID = '%(pid)'**  
pid of the worker handling the request.

**WORKER\_ID = '%(wid)'**  
id of the worker handling the request.

**ASYNC\_SWITCHES = '%(switches)'**  
Number of async switches.

**CORE = '%(core)'**  
The core running the request.

**MEM\_VSZ = '%(vsz)'**  
Address space/virtual memory usage (in bytes).

**MEM\_RSS** = `'%(rss)'`  
RSS memory usage (in bytes).

**MEM\_VSZ\_MB** = `'%(vszM)'`  
Address space/virtual memory usage (in megabytes).

**MEM\_RSS\_MV** = `'%(rssM)'`  
RSS memory usage (in megabytes).

**SIZE\_PACKET\_UWSGI** = `'%(pktsize)'`  
Size of the internal request uwsgi packet.

**MOD1** = `'%(modifier1)'`  
modifier1 of the request. See `.routing.modifiers`.

**MOD2** = `'%(modifier2)'`  
modifier2 of the request. See `.routing.modifiers`.

**metric**  
Metric value (see The Metrics subsystem).  
alias of `VarMetric`

**request\_var**  
Request variable value.  
alias of `VarRequestVar`

## 3.8.8 Main process

### Actions

**class** uwsgi.conf.options.main\_process\_actions.**HookAction** (\*args)

**class** uwsgi.conf.options.main\_process\_actions.**ActionMount** (mountpoint, fs=None, src=None, flags=None)

Mount or unmount filesystems.

#### Examples:

- Mount: `proc none /proc`
- Unmount: `/proc`

#### Parameters

- **mountpoint** (*str/unicode*) –
- **fs** (*str/unicode*) – Filesystem. Presence indicates mounting.
- **src** (*str/unicode*) – Presence indicates mounting.
- **flags** (*str/unicode/list*) – Flags available for the operating system. As an example on Linux you will options like: `bind`, `recursive`, `readonly`, `rec`, `detach` etc.

**class** uwsgi.conf.options.main\_process\_actions.**ActionExecute** (command)

Run the shell command.

Command run under `/bin/sh`. If for some reason you do not want to use `/bin/sh`, use `binsh` option,

#### Examples:

- `cat /proc/self/mounts`

**class** uwsgiConf.options.main\_process\_actions.**ActionCall** (*target*, *honour\_exit\_status=False*, *arg\_int=False*)

Call functions in the current process address space.

#### Parameters

- **target** (*str/unicode*) – Symbol and args.
- **honour\_exit\_status** (*bool*) – Expect an int return. Anything != 0 means failure.
- **arg\_int** (*bool*) – Parse the argument as an int.

**class** uwsgiConf.options.main\_process\_actions.**ActionDirChange** (*target\_dir*)  
Changes a directory.

Convenience action, same as `call:chdir <directory>`.

**class** uwsgiConf.options.main\_process\_actions.**ActionDirCreate** (*target\_dir*)  
Creates a directory with 0777.

**class** uwsgiConf.options.main\_process\_actions.**ActionFileCreate** (*fpath*)  
Creates a directory with 0666.

**class** uwsgiConf.options.main\_process\_actions.**ActionExit** (*status\_code=None*)  
Exits.

Convenience action, same as `callint:exit [num]`.

**class** uwsgiConf.options.main\_process\_actions.**ActionPrintout** (*text=None*)  
Prints.

Convenience action, same as calling the `uwsgi_log` symbol.

**class** uwsgiConf.options.main\_process\_actions.**ActionSetHostName** (*name*)  
Sets a host name.

**class** uwsgiConf.options.main\_process\_actions.**ActionAlarm** (*alarm, message*)  
Issues an alarm. See `.alarms` options group.

**class** uwsgiConf.options.main\_process\_actions.**ActionFileWrite** (*target, text, append=False, newline=False*)

Writes a string to the specified file.

If file doesn't exist it will be created.

---

**Note:** Since 1.9.21

---

#### Parameters

- **target** (*str/unicode*) – File to write to.
- **text** (*str/unicode*) – Text to write into file.
- **append** (*bool*) – Append text instead of rewrite.
- **newline** (*bool*) – Add a newline at the end.

**class** uwsgiConf.options.main\_process\_actions.**ActionFifoWrite** (*target, text, wait=False*)

Writes a string to the specified FIFO (see `fifo_file` from `master_process` params).

**Parameters** **wait** (*bool*) – Wait until FIFO is available.

**class** uwsgi.conf.options.main\_process\_actions.**ActionUnlink** (*target*)  
Unlink the specified file.

---

**Note:** Since 1.9.21

---

**class** uwsgi.conf.options.main\_process.**MainProcess** (*\*args, \*\*kwargs*)  
Main process is the uWSGI process.

**Warning:** Do not run uWSGI instances as root. You can start your uWSGIs as root, but be sure to drop privileges with the `uid` and `gid` options from `set_owner_params`.

**class actions**

Actions available for `.set_hook()`.

**alarm**

alias of `uwsgi.conf.options.main_process_actions.ActionAlarm`

**call**

alias of `uwsgi.conf.options.main_process_actions.ActionCall`

**dir\_change**

alias of `uwsgi.conf.options.main_process_actions.ActionDirChange`

**dir\_create**

alias of `uwsgi.conf.options.main_process_actions.ActionDirCreate`

**execute**

alias of `uwsgi.conf.options.main_process_actions.ActionExecute`

**exit**

alias of `uwsgi.conf.options.main_process_actions.ActionExit`

**fifo\_write**

alias of `uwsgi.conf.options.main_process_actions.ActionFifoWrite`

**file\_create**

alias of `uwsgi.conf.options.main_process_actions.ActionFileCreate`

**file\_write**

alias of `uwsgi.conf.options.main_process_actions.ActionFileWrite`

**mount**

alias of `uwsgi.conf.options.main_process_actions.ActionMount`

**printout**

alias of `uwsgi.conf.options.main_process_actions.ActionPrintout`

**set\_host\_name**

alias of `uwsgi.conf.options.main_process_actions.ActionSetHostName`

**unlink**

alias of `uwsgi.conf.options.main_process_actions.ActionUnlink`

**class phases**

Phases available for hooking using `.set_hook()`.

Some of them may be **fatal** - a failing hook for them will mean failing of the whole uWSGI instance (generally calling `exit(1)`).



**ASAP = 'asap'**  
As soon as possible. **Fatal**  
Run directly after configuration file has been parsed, before anything else is done.

**JAIL\_PRE = 'pre-jail'**  
Before jailing. **Fatal**  
Run before any attempt to drop privileges or put the process in some form of jail.

**JAIL\_IN = 'in-jail'**  
In jail after initialization. **Fatal**  
Run soon after jayling, but after post-jail. If jailing requires fork(), the children run this phase.

**JAIL\_POST = 'post-jail'**  
After jailing. **Fatal**  
Run soon after any jailing, but before privileges drop. If jailing requires fork(), the parent process run this phase.

**PRIV\_DROP\_PRE = 'as-root'**  
Before privileges drop. **Fatal**  
Last chance to run something as root.

**PRIV\_DROP\_POST = 'as-user'**  
After privileges drop. **Fatal**

**MASTER\_START = 'master-start'**  
When Master starts.

**EMPEROR\_START = 'emperor-start'**  
When Emperor starts.

**EMPEROR\_STOP = 'emperor-stop'**  
When Emperor sent a stop message.

**EMPEROR\_RELOAD = 'emperor-reload'**  
When Emperor sent a reload message.

**EMPEROR\_LOST = 'emperor-lost'**  
When Emperor connection is lost.

**EXIT = 'as-user-atexit'**  
Before app exit and reload.

**APP\_LOAD\_PRE = 'pre-app'**  
Before app loading. **Fatal**

**APP\_LOAD\_POST = 'post-app'**  
After app loading. **Fatal**

**VASSAL\_ON\_DEMAND\_IN = 'as-on-demand-vassal'**  
Whenever a vassal enters on-demand mode.

**VASSAL\_CONFIG\_CHANGE\_POST = 'as-on-config-vassal'**  
Whenever the emperor detects a config change for an on-demand vassal.

**VASSAL\_START\_PRE = 'as-emperor-before-vassal'**  
Before the new vassal is spawned.

**VASSAL\_PRIV\_DRP\_PRE = 'as-vassal-before-drop'**  
In vassal, before dropping its privileges.

**VASSAL\_SET\_NAMESPACE = 'as-emperor-setns'**

In the emperor entering vassal namespace.

**VASSAL\_START\_IN = 'as-vassal'**

In the vassal before executing the uwsgi binary. **Fatal**

In vassal on start just before calling exec() directly in the new namespace.

**VASSAL\_START\_POST = 'as-emperor'**

In the emperor soon after a vassal has been spawn.

**Setting 4 env vars:**

- UWSGI\_VASSAL\_CONFIG
- UWSGI\_VASSAL\_PID
- UWSGI\_VASSAL\_UID
- UWSGI\_VASSAL\_GID

**GATEWAY\_START\_IN\_EACH = 'as-gateway'**

In each gateway on start.

**MULE\_START\_IN\_EACH = 'as-mule'**

In each mule on start.

**WORKER\_ACCEPTING\_PRE\_EACH = 'accepting'**

Before the each worker starts accepting requests.

---

**Note:** Since 1.9.21

---

**WORKER\_ACCEPTING\_PRE\_FIRST = 'accepting1'**

Before the first worker starts accepting requests.

---

**Note:** Since 1.9.21

---

**WORKER\_ACCEPTING\_PRE\_EACH\_ONCE = 'accepting-once'**

Before the each worker starts accepting requests, one time per instance.

---

**Note:** Since 1.9.21

---

**WORKER\_ACCEPTING\_PRE\_FIRST\_ONCE = 'accepting1-once'**

Before the first worker starts accepting requests, one time per instance.

---

**Note:** Since 1.9.21

---

**set\_basic\_params** (*touch\_reload=None, priority=None, vacuum=None, binary\_path=None, honour\_stdin=None*)

**Parameters**

- **touch\_reload** (*str/list*) – Reload uWSGI if the specified file or directory is modified/touched.
- **priority** (*int*) – Set processes/threads priority (*nice*) value.

- **vacuum** (*bool*) – Try to remove all of the generated files/sockets (UNIX sockets and pidfiles) upon exit.
- **binary\_path** (*str/unicode*) – Force uWSGI binary path. If you do not have uWSGI in the system path you can force its path with this option to permit the reloading system and the Emperor to easily find the binary to execute.
- **honour\_stdin** (*bool*) – Do not remap stdin to `/dev/null`. By default, `stdin` is remapped to `/dev/null` on uWSGI startup. If you need a valid stdin (for debugging, piping and so on) use this option.

**set\_memory\_params** (*ksm\_interval=None, no\_swap=None*)

Set memory related parameters.

#### Parameters

- **ksm\_interval** (*int*) – Kernel Samepage Merging frequency option, that can reduce memory usage. Accepts a number of requests (or master process cycles) to run page scanner after.

---

**Note:** Linux only.

---

– <http://uwsgi.readthedocs.io/en/latest/KSM.html>

- **no\_swap** (*bool*) – Lock all memory pages avoiding swapping.

**daemonize** (*log\_into, after\_app\_loading=False*)

Daemonize uWSGI.

#### Parameters

- **log\_into** (*str/unicode*) – Logging destination:
  - File: `/tmp/mylog.log`
  - UDP: `192.168.1.2:1717`

---

**Note:** This will require an UDP server to manage log messages. Use `networking.register_socket('192.168.1.2:1717, type=networking.SOCK_UDP)` to start uWSGI UDP server.

---

- **bool after\_app\_loading** (*str/unicode*) – Whether to daemonize after or before applications loading.

**change\_dir** (*to, after\_app\_loading=False*)

Chdir to specified directory before or after apps loading.

#### Parameters

- **to** (*str/unicode*) – Target directory.
- **after\_app\_loading** (*bool*) – *True* - after load *False* - before load

**set\_owner\_params** (*uid=None, gid=None, add\_gids=None, set\_asap=False*)

Set process owner params - user, group.

#### Parameters

- **uid** (*str/unicode/int*) – Set uid to the specified username or uid.

- **gid** (*str/unicode/int*) – Set gid to the specified groupname or gid.
- **add\_gids** (*list/str/unicode/int*) – Add the specified group id to the process credentials. This options allows you to add additional group ids to the current process. You can specify it multiple times.
- **set\_asap** (*bool*) – Set as soon as possible. Setting them on top of your vassal file will force the instance to `setuid()/setgid()` as soon as possible and without the (theoretical) possibility to override them.

**get\_owner** (*default=True*)

Return (User ID, Group ID) tuple

**Parameters** **default** (*bool*) – Whether to return default if not set.

**Return type** tuple[int, int]

**set\_hook** (*phase, action*)

Allows setting hooks (attaching actions) for various uWSGI phases.

**Parameters**

- **phase** (*str/unicode*) – See constants in `.phases`.
- **action** (*str/unicode/list/HookAction/list [HookAction]*) –

**set\_hook\_touch** (*fpath, action*)

Allows running certain action when the specified file is touched.

**Parameters**

- **fpath** (*str/unicode*) – File path.
- **action** (*str/unicode/list/HookAction/list [HookAction]*) –

**set\_hook\_after\_request** (*func*)

Run the specified function/symbol (C level) after each request.

**Parameters** **func** (*str/unicode*) –

**set\_on\_exit\_params** (*skip\_hooks=None, skip\_teardown=None*)

Set params related to process exit procedure.

**Parameters**

- **skip\_hooks** (*bool*) – Skip EXIT phase hook.

---

**Note:** Ignored by the master.

---

- **skip\_teardown** (*bool*) – Allows skipping teardown (finalization) processes for some plugins.

---

**Note:** Ignored by the master.

---

**Supported by:**

- Perl
- Python

**run\_command\_on\_event** (*command*, *phase='asap'*)

Run the given command on a given phase.

**Parameters**

- **command** (*str/unicode*) –
- **phase** (*str/unicode*) – See constants in `Phases` class.

**run\_command\_on\_touch** (*command*, *target*)

Run command when the specified file is modified/touched.

**Parameters**

- **command** (*str/unicode*) –
- **target** (*str/unicode*) – File path.

**set\_pid\_file** (*fpath*, *before\_priv\_drop=True*, *safe=False*)

Creates pidfile before or after privileges drop.

**Parameters**

- **fpath** (*str/unicode*) – File path.
- **before\_priv\_drop** (*bool*) – Whether to create pidfile before privileges are dropped.

---

**Note:** Vacuum is made after privileges drop, so it may not be able to delete PID file if it was created before dropping.

---

- **safe** (*bool*) – The safe-pidfile works similar to pidfile but performs the write a little later in the loading process. This avoids overwriting the value when app loading fails, with the consequent loss of a valid PID number.

**set\_naming\_params** (*autonaming=None*, *prefix=None*, *suffix=None*, *name=None*)

Setups processes naming parameters.

**Parameters**

- **autonaming** (*bool*) – Automatically set process name to something meaningful. Generated process names may be ‘uWSGI Master’, ‘uWSGI Worker #’, etc.
- **prefix** (*str/unicode*) – Add prefix to process names.
- **suffix** (*str/unicode*) – Append string to process names.
- **name** (*str/unicode*) – Set process names to given static value.

### 3.8.9 Master process

**class** uwsgi.conf.options.master\_process.**MasterProcess** (*\*args*, *\*\*kwargs*)

Master process is a separate process offering mentoring capabilities for other processes. Only one master process per uWSGI instance.

uWSGI’s built-in `prefork+threading` multi-worker management mode, activated by flicking the master switch on.

---

**Note:** For all practical serving deployments it’s not really a good idea not to use master mode.

---

**set\_basic\_params** (*enable=None, name=None, no\_orphans=None, as\_root=None, subproc\_check\_interval=None, fifo\_file=None*)

**Parameters**

- **enable** (*bool*) – Enable uWSGI master process.
- **master\_enabled** (*str/unicode*) – Set master process name to given value.
- **no\_orphans** (*bool*) – Automatically kill workers if master dies (can be dangerous for availability).
- **as\_root** (*bool*) – Leave master process running as root.
- **subproc\_check\_interval** (*int*) – Set the interval (in seconds) of master checks. Default: 1 The master process makes a scan of subprocesses, etc. every N seconds.

**Warning:** You can increase this time if you need to, but it's DISCOURAGED.

- **fifo\_file** (*str/unicode*) – Enables the master FIFO.

---

**Note:** Placeholders can be used to build paths, e.g.: {project\_runtime\_dir}.fifo See [Section.project\\_name](#) and [Section.runtime\\_dir](#).

---

Instead of signals, you can tell the master to create a UNIX named pipe (FIFO) that you may use to issue commands to the master.

Up to 10 different FIFO files supported. By default the first specified is bound (mapped as '0').

– <http://uwsgi.readthedocs.io/en/latest/MasterFIFO.html#the-master-fifo>

---

**Note:** Since 1.9.17

---

**set\_exit\_events** (*no\_workers=None, idle=None, reload=None, sig\_term=None*)

Do exit on certain events

**Parameters**

- **no\_workers** (*bool*) – Shutdown uWSGI when no workers are running.
- **idle** (*bool*) – Shutdown uWSGI when idle.
- **reload** (*bool*) – Force exit even if a reload is requested.
- **sig\_term** (*bool*) – Exit on SIGTERM instead of brutal workers reload.

---

**Note:** Before 2.1 SIGTERM reloaded the stack while SIGINT/SIGQUIT shut it down.

---

**set\_exception\_handling\_params** (*handler=None, catch=None, no\_write\_exception=None*)

Exception handling related params.

**Parameters**

- **handler** (*str/unicode/list[str/unicode]*) – Register one or more exception handling C-functions.
- **catch** (*bool*) – Catch exceptions and report them as http output (including stack trace and env params).

**Warning:** Use only for testing purposes.

- **no\_write\_exception** (*bool*) – Disable exception generation on write()/writev().

---

**Note:** This can be combined with logging.  
`set_filters(write_errors=False, sigpipe=False)`.

---

**set\_idle\_params** (*timeout=None, exit=None*)

Activate idle mode - put uWSGI in cheap mode after inactivity timeout.

**Parameters**

- **timeout** (*int*) – Inactivity timeout in seconds.
- **exit** (*bool*) – Shutdown uWSGI when idle.

**set\_reload\_params** (*mercy=None, exit=None*)

Set reload related params.

**Parameters**

- **mercy** (*int*) – Set the maximum time (in seconds) we wait for workers and other processes to die during reload/shutdown.
- **exit** (*bool*) – Force exit even if a reload is requested.

**add\_cron\_task** (*command, weekday=None, month=None, day=None, hour=None, minute=None, legion=None, unique=None, harakiri=None*)

Adds a cron task running the given command on the given schedule. <http://uwsgi.readthedocs.io/en/latest/Cron.html>

**HINTS:**

- **Use negative values to say every:** hour=-3 stands for *every 3 hours*
- **Use - (minus) to make interval:** minute='13-18' stands for *from minute 13 to 18*

---

**Note:** We use cron2 option available since 1.9.11.

---

**Parameters**

- **command** (*str/unicode*) – Command to execute on schedule (with or without path).
- **weekday** (*int/str/unicode*) – Day of a the week number. Defaults to *each*. 0 - Sunday 1 - Monday 2 - Tuesday 3 - Wednesday 4 - Thursday 5 - Friday 6 - Saturday
- **month** (*int/str/unicode*) – Month number 1-12. Defaults to *each*.
- **day** (*int/str/unicode*) – Day of the month number 1-31. Defaults to *each*.
- **hour** (*int/str/unicode*) – Hour 0-23. Defaults to *each*.

- **minute** (*int|str|unicode*) – Minute 0-59. Defaults to *each*.
- **legion** (*str|unicode*) – Set legion (cluster) name to use this cron command against. Such commands are only executed by legion lord node.
- **unique** (*bool*) – Marks command as unique. Default to not unique. Some commands can take a long time to finish or just hang doing their thing. Sometimes this is okay, but there are also cases when running multiple instances of the same command can be dangerous.
- **harakiri** (*int*) – Enforce a time limit (in seconds) on executed commands. If a command is taking longer it will be killed.

**attach\_process\_classic** (*command\_or\_pid\_path*, *background*, *control=False*,  
*for\_legion=False*)

Attaches a command/daemon to the master process optionally managed by a pidfile.

This will allow the uWSGI master to control/monitor/respawn this process.

---

**Note:** This uses old classic uWSGI means of process attaching To have more control use `.attach_process()` method (requires uWSGI 2.0+)

---

<http://uwsgi-docs.readthedocs.io/en/latest/AttachingDaemons.html>

#### Parameters

- **command\_or\_pid\_path** (*str|unicode*) –
- **background** (*bool*) – Must indicate whether process is in background.
- **control** (*bool*) – Consider this process a control: when the daemon dies, the master exits.

---

**Note:** pidfile managed processed not supported.

---

- **for\_legion** (*bool*) – Legion daemons will be executed only on the legion lord node, so there will always be a single daemon instance running in each legion. Once the lord dies a daemon will be spawned on another node.

---

**Note:** uWSGI 1.9.9+ required.

---

**attach\_process** (*command*, *for\_legion=False*, *broken\_counter=None*, *pidfile=None*, *control=None*, *daemonize=None*, *touch\_reload=None*, *signal\_stop=None*, *signal\_reload=None*, *honour\_stdin=None*, *uid=None*, *gid=None*, *new\_pid\_ns=None*, *change\_dir=None*)

Attaches a command/daemon to the master process.

This will allow the uWSGI master to control/monitor/respawn this process.

<http://uwsgi-docs.readthedocs.io/en/latest/AttachingDaemons.html>

#### Parameters

- **command** (*str|unicode*) – The command line to execute.
- **for\_legion** (*bool*) – Legion daemons will be executed only on the legion lord node, so there will always be a single daemon instance running in each legion. Once the lord dies a daemon will be spawned on another node.



- **broken\_counter** (*int*) – Maximum attempts before considering a daemon “broken”.
- **pidfile** (*str/unicode*) – The pidfile path to check (enable smart mode).
- **control** (*bool*) – If True, the daemon becomes a *control* one: if it dies the whole uWSGI instance dies.
- **daemonize** (*bool*) – Daemonize the process (enable smart2 mode).
- **touch\_reload** (*list/str/unicode*) – List of files to check: whenever they are ‘touched’, the daemon is restarted
- **signal\_stop** (*int*) – The signal number to send to the daemon when uWSGI is stopped.
- **signal\_reload** (*int*) – The signal number to send to the daemon when uWSGI is reloaded.
- **honour\_stdin** (*bool*) – The signal number to send to the daemon when uWSGI is reloaded.
- **uid** (*str/unicode/int*) – Drop privileges to the specified uid.

---

**Note:** Requires master running as root.

---

- **gid** (*str/unicode/int*) – Drop privileges to the specified gid.

---

**Note:** Requires master running as root.

---

- **new\_pid\_ns** (*bool*) – Spawn the process in a new pid namespace.

---

**Note:** Requires master running as root.

---



---

**Note:** Linux only.

---

- **change\_dir** (*str/unicode*) – Use chdir() to the specified directory before running the command.

### 3.8.10 Monitoring

#### Metric Types

```
class uwsgi.conf.options.monitoring_metric_types.Metric(name, oid=None,
                                                         alias_for=None, collector=None,
                                                         initial_value=None, collect_interval=None,
                                                         reset_after_push=None)
```

#### Parameters

- **name** (*str/unicode*) – Metric name.

---

**Note:** Only numbers, letters, underscores, dashes and dots.

---

- **alias\_for** (*str/unicode*) – If set metric will be a simple alias for the specified one.
- **oid** (*str/unicode*) – Metric OID.  
Required for SNMP.
  - <http://uwsgi-docs.readthedocs.io/en/latest/Metrics.html#oid-assignment-for-plugins>
- **collector** (*Collector*) – Collector to be used. If not set it is considered that the value must be updated manually from applications using the metrics API.
  - <http://uwsgi-docs.readthedocs.io/en/latest/Metrics.html#api>
- **initial\_value** (*int*) – Set the metric to a specific value on startup.
- **collect\_interval** (*int*) – How often the metric should be gathered. In seconds.
- **reset\_after\_push** (*bool*) – Reset the metric to zero (or the configured initial\_value) after it's been pushed.

```
class uwsgi.conf.options.monitoring_metric_types.MetricTypeCounter(name,
                                                                    oid=None,
                                                                    alias_for=None,
                                                                    collector=None,
                                                                    initial_value=None,
                                                                    collect_interval=None,
                                                                    reset_after_push=None)
```

A generally-growing up number.

**Example:**

- number of requests

**Parameters**

- **name** (*str/unicode*) – Metric name.

---

**Note:** Only numbers, letters, underscores, dashes and dots.

---

- **alias\_for** (*str/unicode*) – If set metric will be a simple alias for the specified one.
- **oid** (*str/unicode*) – Metric OID.  
Required for SNMP.
  - <http://uwsgi-docs.readthedocs.io/en/latest/Metrics.html#oid-assignment-for-plugins>
- **collector** (*Collector*) – Collector to be used. If not set it is considered that the value must be updated manually from applications using the metrics API.
  - <http://uwsgi-docs.readthedocs.io/en/latest/Metrics.html#api>

- **initial\_value** (*int*) – Set the metric to a specific value on startup.
- **collect\_interval** (*int*) – How often the metric should be gathered. In seconds.
- **reset\_after\_push** (*bool*) – Reset the metric to zero (or the configured initial\_value) after it's been pushed.

```
class uwsgi.conf.options.monitoring_metric_types.MetricTypeGauge (name,
                                                                oid=None,
                                                                alias_for=None,
                                                                collec-
                                                                tor=None, ini-
                                                                tial_value=None,
                                                                col-
                                                                lect_interval=None,
                                                                re-
                                                                set_after_push=None)
```

A number that can increase or decrease dynamically.

**Example:**

- memory used by a worker
- CPU load

**Parameters**

- **name** (*str/unicode*) – Metric name.

---

**Note:** Only numbers, letters, underscores, dashes and dots.

---

- **alias\_for** (*str/unicode*) – If set metric will be a simple alias for the specified one.
- **oid** (*str/unicode*) – Metric OID.  
Required for SNMP.
  - <http://uwsgi-docs.readthedocs.io/en/latest/Metrics.html#oid-assignment-for-plugins>
- **collector** (*Collector*) – Collector to be used. If not set it is considered that the value must be updated manually from applications using the metrics API.
  - <http://uwsgi-docs.readthedocs.io/en/latest/Metrics.html#api>
- **initial\_value** (*int*) – Set the metric to a specific value on startup.
- **collect\_interval** (*int*) – How often the metric should be gathered. In seconds.
- **reset\_after\_push** (*bool*) – Reset the metric to zero (or the configured initial\_value) after it's been pushed.

```
class uwsgi.conf.options.monitoring_metric_types.MetricTypeAbsolute (name,
                                                                    oid=None,
                                                                    alias_for=None,
                                                                    collector=None,
                                                                    initial_value=None,
                                                                    collect_interval=None,
                                                                    reset_after_push=None)
```

An absolute number.

**Example:**

- memory of the whole server
- size of the hard disk.

**Parameters**

- **name** (*str/unicode*) – Metric name.

---

**Note:** Only numbers, letters, underscores, dashes and dots.

---

- **alias\_for** (*str/unicode*) – If set metric will be a simple alias for the specified one.
- **oid** (*str/unicode*) – Metric OID.  
Required for SNMP.
  - <http://uwsgi-docs.readthedocs.io/en/latest/Metrics.html#oid-assignment-for-plugins>
- **collector** (*Collector*) – Collector to be used. If not set it is considered that the value must be updated manually from applications using the metrics API.
  - <http://uwsgi-docs.readthedocs.io/en/latest/Metrics.html#api>
- **initial\_value** (*int*) – Set the metric to a specific value on startup.
- **collect\_interval** (*int*) – How often the metric should be gathered. In seconds.
- **reset\_after\_push** (*bool*) – Reset the metric to zero (or the configured initial\_value) after it's been pushed.

```
class uwsgi.conf.options.monitoring_metric_types.MetricTypeAlias (name,
                                                                    oid=None,
                                                                    alias_for=None,
                                                                    collector=None,
                                                                    initial_value=None,
                                                                    collect_interval=None,
                                                                    reset_after_push=None)
```

This is a virtual metric pointing to another one .

You can use it to give different names to already existing metrics.

### Parameters

- **name** (*str/unicode*) – Metric name.

---

**Note:** Only numbers, letters, underscores, dashes and dots.

---

- **alias\_for** (*str/unicode*) – If set metric will be a simple alias for the specified one.
- **oid** (*str/unicode*) – Metric OID.  
Required for SNMP.
  - <http://uwsgi-docs.readthedocs.io/en/latest/Metrics.html#oid-assignment-for-plugins>
- **collector** (*Collector*) – Collector to be used. If not set it is considered that the value must be updated manually from applications using the metrics API.
  - <http://uwsgi-docs.readthedocs.io/en/latest/Metrics.html#api>
- **initial\_value** (*int*) – Set the metric to a specific value on startup.
- **collect\_interval** (*int*) – How often the metric should be gathered. In seconds.
- **reset\_after\_push** (*bool*) – Reset the metric to zero (or the configured initial\_value) after it's been pushed.

### Pushers

**class** uwsgi.conf.options.monitoring\_pushers.**Pusher** (\*args)

**class** uwsgi.conf.options.monitoring\_pushers.**PusherSocket** (address, prefix=None)  
Push metrics to a UDP server.

Uses the following format: <metric> <type> <value> <type> - is in the numeric form of metric type.

#### Parameters

- **address** (*str/unicode*) –
- **prefix** (*str/unicode*) – Arbitrary prefix to differentiate sender.

**class** uwsgi.conf.options.monitoring\_pushers.**PusherRrdtool** (target\_dir, library=None, push\_interval=None)

This will store an rrd file for each metric in the specified directory.

Each rrd file has a single data source named “metric”.

#### Parameters

- **target\_dir** (*str/unicode*) – Directory to store rrd files into.
- **library** (*str/unicode*) – Set the name of rrd library. Default: librrd.so.
- **push\_interval** (*int*) – Set push frequency.

**class** uwsgi.conf.options.monitoring\_pushers.**PusherStatsd** (address, prefix=None, no\_workers=None, all\_gauges=None)

Push metrics to a statsd server.

#### Parameters

- **address** (*str/unicode*) –

- **prefix** (*str/unicode*) – Arbitrary prefix to differentiate sender.
- **no\_workers** (*bool*) – Disable generation of single worker metrics.
- **all\_gauges** (*bool*) – Push all metrics to statsd as gauges.

```
class uwsgi.conf.options.monitoring_pushers.PusherCarbon (address,  
                                                    node_realm=None,  
                                                    node_root=None,  
                                                    push_interval=None,  
                                                    idle_avg_source=None,  
                                                    use_metrics=None,  
                                                    no_workers=None, time-  
                                                    out=None, retries=None,  
                                                    retries_delay=None, host-  
                                                    name_dots_replacer=None)
```

Push metrics to a Carbon server of Graphite.

Metric node format: `<node_root>.hostname.<node_realm>.metrics_data`.

- <http://uwsgi.readthedocs.io/en/latest/Carbon.html>
- <http://uwsgi.readthedocs.io/en/latest/tutorials/GraphiteAndMetrics.html>

#### Parameters

- **address** (*str/unicode/list[str/unicode]*) – Host and port. Example: 127.0.0.1:2004
- **node\_realm** (*str/unicode*) – Set carbon metrics realm node.
- **node\_root** (*str/unicode*) – Set carbon metrics root node. Default: uwsgi.
- **push\_interval** (*int*) – Set carbon push frequency in seconds. Default: 60.
- **no\_workers** (*bool*) – Disable generation of single worker metrics.
- **idle\_avg\_source** (*str/unicode*) – Average values source during idle period (no requests).

#### Variants:

- last (default)
- zero
- none
- **use\_metrics** (*bool*) – Don't compute all statistics, use metrics subsystem data instead.

**Warning:** Key names of built-in stats are different from those of metrics system.

- **timeout** (*int*) – Set carbon connection timeout in seconds. Default: 3.
- **retries** (*int*) – Set maximum number of retries in case of connection errors. Default: 1.
- **retries\_delay** (*int*) – Set connection retry delay in seconds. Default: 7.
- **hostname\_dots\_replacer** (*str/unicode*) – Set char to use as a replacement for dots in hostname in `<node_root>.hostname.<node_realm>.metrics_data'`

This affects Graphite aggregation mechanics.

---

**Note:** Dots are not replaced by default.

---

**class** uwsgi.conf.options.monitoring\_pushers.**PusherZabbix** (*address, prefix=None, template=None*)

Push metrics to a zabbix server.

#### Parameters

- **address** (*str/unicode*) –
- **prefix** (*str/unicode*) – Arbitrary prefix to differentiate sender.
- **template** (*str/unicode*) – Print (or store to a file) the zabbix template for the current metrics setup.

**class** uwsgi.conf.options.monitoring\_pushers.**PusherMongo** (*address=None, collection=None, push\_interval=None*)

Push statistics (as JSON) the the specified MongoDB database.

#### Parameters

- **address** (*str/unicode*) – Default: 127.0.0.1:27017
- **collection** (*str/unicode*) – MongoDB collection to write into. Default: uwsgi.statistics
- **push\_interval** (*int*) – Write interval in seconds.

**class** uwsgi.conf.options.monitoring\_pushers.**PusherFile** (*fpath=None, separator=None, push\_interval=None*)

Stores stats JSON into a file.

---

**Note:** Mainly for demonstration purposes.

---

#### Parameters

- **fpath** (*str/unicode*) – File path. Default: uwsgi.stats
- **separator** (*str/unicode*) – New entry separator. Default:
- **push\_interval** (*int*) – Write interval in seconds.

## Collectors

**class** uwsgi.conf.options.monitoring\_collectors.**Collector** (*\*args, \*\*kwargs*)

**class** uwsgi.conf.options.monitoring\_collectors.**CollectorPointer** (*\*args, \*\*kwargs*)

The value is collected from memory pointer.

**class** uwsgi.conf.options.monitoring\_collectors.**CollectorFile** (*fpath, get\_slot=None*)

The value is collected from a file.

#### Parameters

- **fpath** (*str/unicode*) – File path.
- **get\_slot** (*int*) – Get value from the given slot number. Slots: the content is split (using n, t, spaces, r and zero as separator) and the item (the returned array is zero-based) used as the return value.

**class** uwsgi.conf.options.monitoring\_collectors.**CollectorFunction** (*func*)

The value is computed calling a specific C function every time.

---

**Note:**

- The argument it takes is a `uwsgi_metric` pointer. You generally do not need to parse the metric, so just casting to void will avoid headaches.
  - The function must return an `int64_t` value.
- 

**Parameters** **func** (*str/unicode*) – Function to call.

**class** uwsgi.conf.options.monitoring\_collectors.**CollectorSum** (*what*)

The value is the sum of other metrics.

**class** uwsgi.conf.options.monitoring\_collectors.**CollectorAvg** (*what*)

The value is the algebraic average of the children.

---

**Note:** Since 1.9.20

---

**class** uwsgi.conf.options.monitoring\_collectors.**CollectorAccumulator** (*what*)

Always add the sum of children to the final value.

Example:

- Round 1: child1 = 22, child2 = 17 -> metric\_value = 39
- Round 2: child1 = 26, child2 = 30 -> metric\_value += 56

**class** uwsgi.conf.options.monitoring\_collectors.**CollectorAdder** (*what, value*)

Add the specified argument (`arg1n`) to the sum of children.

**Parameters** **value** (*int*) – Value to add (multiply if it is `CollectorMultiplier`).

**class** uwsgi.conf.options.monitoring\_collectors.**CollectorMultiplier** (*what, value*)

Multiply the sum of children by the specified argument.

Example:

- child1 = 22, child2 = 17, arg1n = 3 -> metric\_value = (22+17)\*3

**Parameters** **value** (*int*) – Value to add (multiply if it is `CollectorMultiplier`).

**class** uwsgi.conf.options.monitoring.**Monitoring** (*\*args, \*\*kwargs*)

Monitoring facilities.

- **SNMP** - <http://uwsgi.readthedocs.io/en/latest/SNMP.html>
- **Stats** - <http://uwsgi.readthedocs.io/en/latest/StatsServer.html> Set of metrics gathered from uWSGI internals.
- **Metrics** - <http://uwsgi.readthedocs.io/en/latest/Metrics.html> Basic set of metrics gathered from uWSGI internals + user defined metrics.



**class metric\_types**

Various metric types to represent data of various nature.

User metrics must inherit from one of those.

**absolute**

alias of `uwsgi.conf.options.monitoring_metric_types.MetricTypeAbsolute`

**alias**

alias of `uwsgi.conf.options.monitoring_metric_types.MetricTypeAlias`

**counter**

alias of `uwsgi.conf.options.monitoring_metric_types.MetricTypeCounter`

**gauge**

alias of `uwsgi.conf.options.monitoring_metric_types.MetricTypeGauge`

**class collectors**

Metric collection and accumulation means.

**accumulator**

alias of `uwsgi.conf.options.monitoring_collectors.CollectorAccumulator`

**adder**

alias of `uwsgi.conf.options.monitoring_collectors.CollectorAdder`

**avg**

alias of `uwsgi.conf.options.monitoring_collectors.CollectorAvg`

**file**

alias of `uwsgi.conf.options.monitoring_collectors.CollectorFile`

**function**

alias of `uwsgi.conf.options.monitoring_collectors.CollectorFunction`

**multiplier**

alias of `uwsgi.conf.options.monitoring_collectors.CollectorMultiplier`

**pointer**

alias of `uwsgi.conf.options.monitoring_collectors.CollectorPointer`

**sum**

alias of `uwsgi.conf.options.monitoring_collectors.CollectorSum`

**class pushers**

Means to deliver metrics to various remotes or locals.

These are available for `.register_stats_pusher()`.

**carbon**

alias of `uwsgi.conf.options.monitoring_pushers.PusherCarbon`

**file**

alias of `uwsgi.conf.options.monitoring_pushers.PusherFile`

**mongo**

alias of `uwsgi.conf.options.monitoring_pushers.PusherMongo`

**rrdtool**

alias of `uwsgi.conf.options.monitoring_pushers.PusherRrdtool`

**socket**

alias of `uwsgi.conf.options.monitoring_pushers.PusherSocket`

**statsd**

alias of `uwsgi.conf.options.monitoring_pushers.PusherStatsd`

**zabbix**

alias of `uwsgi.conf.options.monitoring_pushers.PusherZabbix`

**register\_metric** (*metric*)

Officially Registered Metrics:

- **worker 3 - exports information about workers.** Example: `worker.1.requests [or 3.1.1]` reports the number of requests served by worker 1.
- **plugin 4 - namespace for metrics automatically added by plugins.** Example: `plugins.foo.bar`
- `core 5` - namespace for general instance information.
- **router 6 - namespace for corerouters.** Example: `router.http.active_sessions`
- **socket 7 - namespace for sockets.** Example: `socket.0.listen_queue`
- **mule 8 - namespace for mules.** Example: `mule.1.signals`
- **spooler 9 - namespace for spoolers.** Example: `spooler.1.signals`
- `system 10` - namespace for system metrics, like `loadavg` or free memory.

**Parameters** *metric* (*Metric*/*list* [*Metric*]) – Metric object.

**set\_metrics\_params** (*enable=None, store\_dir=None, restore=None, no\_cores=None*)

Sets basic Metrics subsystem params.

uWSGI metrics subsystem allows you to manage “numbers” from your apps.

When enabled, the subsystem configures a vast amount of metrics (like requests per-core, memory usage, etc) but, in addition to this, you can configure your own metrics, such as the number of active users or, say, hits of a particular URL, as well as the memory consumption of your app or the whole server.

- <http://uwsgi.readthedocs.io/en/latest/Metrics.html>
- SNMP Integration - <http://uwsgi.readthedocs.io/en/latest/Metrics.html#snmp-integration>

**Parameters**

- **enable** (*bool*) – Enables the subsystem.
- **store\_dir** (*str/unicode*) – Directory to store metrics. The metrics subsystem can expose all of its metrics in the form of text files in a directory. The content of each file is the value of the metric (updated in real time).

---

**Note:** Placeholders can be used to build paths, e.g.: `{project_runtime_dir}/metrics/`  
See `Section.project_name` and `Section.runtime_dir`.

---

- **restore** (*bool*) – Restore previous metrics from `store_dir`. When you restart a uWSGI instance, all of its metrics are reset. Use the option to force the metric subsystem to read-back the values from the metric directory before starting to collect values.
- **no\_cores** (*bool*) – Disable generation of cores-related metrics.

**set\_metrics\_threshold** (*name, value, check\_interval=None, reset\_to=None, alarm=None, alarm\_message=None*)

Sets metric threshold parameters.

**Parameters**

- **name** (*str/unicode*) – Metric name.
- **value** (*int*) – Threshold value.
- **reset\_to** (*int*) – Reset value to when threshold is reached.
- **check\_interval** (*int*) – Threshold check interval in seconds.
- **alarm** (*str/unicode/AlarmType*) – Alarm to trigger when threshold is reached.
- **alarm\_message** (*str/unicode*) – Message to pass to alarm. If not set metrics name is passed.

**set\_stats\_params** (*address=None, enable\_http=None, minify=None, no\_cores=None, no\_metrics=None, push\_interval=None*)  
Enables stats server on the specified address.

- <http://uwsgi.readthedocs.io/en/latest/StatsServer.html>

**Parameters**

- **address** (*str/unicode*) – Address/socket to make stats available on.

**Examples:**

- 127.0.0.1:1717
- /tmp/statsock
- :5050
- **enable\_http** (*bool*) – Server stats over HTTP. Prefixes stats server json output with http headers.
- **minify** (*bool*) – Minify statistics json output.
- **no\_cores** (*bool*) – Disable generation of cores-related stats.
- **no\_metrics** (*bool*) – Do not include metrics in stats output.
- **push\_interval** (*int*) – Set the default frequency of stats pushers in seconds/

**register\_stats\_pusher** (*pusher*)

Registers a pusher to be used for pushing statistics to various remotes/locals.

**Parameters** *pusher* (*Pusher/list [Pusher]*) –

**enable\_snmp** (*address, community\_string*)

Enables SNMP.

uWSGI server embeds a tiny SNMP server that you can use to integrate your web apps with your monitoring infrastructure.

- <http://uwsgi.readthedocs.io/en/latest/SNMP.html>

---

**Note:** SNMP server is started in the master process after dropping the privileges. If you want it to listen on a privileged port, you can either use Capabilities on Linux, or use the `as-root` option to run the master process as root.

---

**Parameters**

- **address** (*str/unicode*) – UDP address to bind to.

Examples:

– 192.168.1.1:2222

- **community\_string** (*str/unicode*) – SNMP instance identifier to address it.

## 3.8.11 Networking

### Sockets

**class** uwsgi.conf.options.networking\_sockets.**Socket** (*address*, *bound\_workers=None*, *modifier=None*)

#### Parameters

- **address** (*str/unicode/SocketShared*) – Address ([host]:port or socket file) to bind socket to.
- **bound\_workers** (*str/unicode/int/list*) – Map socket to specific workers. As you can bind a uWSGI instance to multiple sockets, you can use this option to map specific workers to specific sockets to implement a sort of in-process Quality of Service scheme. If you host multiple apps in the same uWSGI instance, you can easily dedicate resources to each of them.
- **modifier** (*Modifier*) – Socket routing modifier.

**class** uwsgi.conf.options.networking\_sockets.**SocketDefault** (*address*, *bound\_workers=None*, *modifier=None*)

Bind using default protocol. See `default_socket_type` option.

#### Parameters

- **address** (*str/unicode/SocketShared*) – Address ([host]:port or socket file) to bind socket to.
- **bound\_workers** (*str/unicode/int/list*) – Map socket to specific workers. As you can bind a uWSGI instance to multiple sockets, you can use this option to map specific workers to specific sockets to implement a sort of in-process Quality of Service scheme. If you host multiple apps in the same uWSGI instance, you can easily dedicate resources to each of them.
- **modifier** (*Modifier*) – Socket routing modifier.

**class** uwsgi.conf.options.networking\_sockets.**SocketHttp** (*address*, *http11=False*, *bound\_workers=None*, *modifier=None*)

Bind to the specified socket using HTTP

#### Parameters

- **address** (*str/unicode/SocketShared*) – Address ([host]:port or socket file) to bind socket to.
- **http11** (*bool*) – Keep-Alive support. If set the server will try to maintain the connection opened if a bunch of rules are respected.

This is not a smart http 1.1 parser (to avoid parsing the whole response) but assumes the developer is generating the right headers.

This has been added to support RTSP protocol for video streaming.

- **bound\_workers** (*str|unicode|int|list*) – Map socket to specific workers. As you can bind a uWSGI instance to multiple sockets, you can use this option to map specific workers to specific sockets to implement a sort of in-process Quality of Service scheme. If you host multiple apps in the same uWSGI instance, you can easily dedicate resources to each of them.
- **modifier** (*Modifier*) – Socket routing modifier.

```
class uwsgi.conf.options.networking_sockets.SocketHttps (address, cert, key,
                                                    ciphers=None,
                                                    client_ca=None,
                                                    bound_workers=None,
                                                    modifier=None)
```

Bind to the specified socket using HTTPS

#### Parameters

- **address** (*str|unicode|SocketShared*) – Address ([host]:port or socket file) to bind socket to.
- **cert** (*str|unicode*) – Certificate file.
- **key** (*str|unicode*) – Private key file.
- **ciphers** (*str|unicode*) – Ciphers [alias] string.

#### Example:

- DEFAULT
- HIGH
- DHE, EDH

– <https://www.openssl.org/docs/man1.1.0/apps/ciphers.html>

- **client\_ca** (*str|unicode*) – Client CA file for client-based auth.
- **bound\_workers** (*str|unicode|int|list*) – Map socket to specific workers. As you can bind a uWSGI instance to multiple sockets, you can use this option to map specific workers to specific sockets to implement a sort of in-process Quality of Service scheme. If you host multiple apps in the same uWSGI instance, you can easily dedicate resources to each of them.
- **modifier** (*Modifier*) – Socket routing modifier.

```
classmethod get_certbot_paths (domain)
```

Returns a tuple of paths for files (certificates\_chain, private\_key) from Certbot <https://certbot.eff.org>

Those paths can be used to pass into Socket initializer.

---

**Note:** If files not found empty strings are returned.

---

**Parameters** **domain** (*str*) – Domain name to get filepaths for.

**Return type** tuple[str, str]

```
class uwsgi.conf.options.networking_sockets.SocketUwsgi (address, persistent=False,  
                                                    bound_workers=None,  
                                                    modifier=None)
```

uwsgi specific socket using uwsgi protocol.

#### Parameters

- **address** (*str/unicode/SocketShared*) – Address ([host]:port or socket file) to bind socket to.
- **persistent** (*bool*) – Use persistent uwsgi protocol (puwsgi).
- **bound\_workers** (*str/unicode/int/list*) – Map socket to specific workers. As you can bind a uWSGI instance to multiple sockets, you can use this option to map specific workers to specific sockets to implement a sort of in-process Quality of Service scheme. If you host multiple apps in the same uWSGI instance, you can easily dedicate resources to each of them.
- **modifier** (*Modifier*) – Socket routing modifier.

```
class uwsgi.conf.options.networking_sockets.SocketUwsgis (address, cert, key,  
                                                    ciphers=None,  
                                                    client_ca=None,  
                                                    bound_workers=None,  
                                                    modifier=None)
```

uwsgi specific socket using uwsgi protocol over SSL.

#### Parameters

- **address** (*str/unicode/SocketShared*) – Address ([host]:port or socket file) to bind socket to.
- **cert** (*str/unicode*) – Certificate file.
- **key** (*str/unicode*) – Private key file.
- **ciphers** (*str/unicode*) – Ciphers [alias] string.

#### Example:

- DEFAULT
- HIGH
- DHE, EDH

– <https://www.openssl.org/docs/man1.1.0/apps/ciphers.html>

- **client\_ca** (*str/unicode*) – Client CA file for client-based auth.
- **bound\_workers** (*str/unicode/int/list*) – Map socket to specific workers. As you can bind a uWSGI instance to multiple sockets, you can use this option to map specific workers to specific sockets to implement a sort of in-process Quality of Service scheme. If you host multiple apps in the same uWSGI instance, you can easily dedicate resources to each of them.
- **modifier** (*Modifier*) – Socket routing modifier.

```
class uwsgi.conf.options.networking_sockets.SocketUdp (address,  
                                                    bound_workers=None, modi-  
                                                    fier=None)
```

Run the udp server on the specified address.

---

**Note:** Mainly useful for SNMP or shared UDP logging.

---

### Parameters

- **address** (*str|unicode|SocketShared*) – Address ([host]:port or socket file) to bind socket to.
- **bound\_workers** (*str|unicode|int|list*) – Map socket to specific workers. As you can bind a uWSGI instance to multiple sockets, you can use this option to map specific workers to specific sockets to implement a sort of in-process Quality of Service scheme. If you host multiple apps in the same uWSGI instance, you can easily dedicate resources to each of them.
- **modifier** (*Modifier*) – Socket routing modifier.

```
class uwsgi.conf.options.networking_sockets.SocketFastcgi (address,      nph=False,
                                                         bound_workers=None,
                                                         modifier=None)
```

Bind to the specified socket using FastCGI.

### Parameters

- **address** (*str|unicode|SocketShared*) – Address ([host]:port or socket file) to bind socket to.
- **nph** (*bool*) – Use NPH mode (“no-parsed-header” - bypass the server completely by sending the complete HTTP header directly to the browser).
- **bound\_workers** (*str|unicode|int|list*) – Map socket to specific workers. As you can bind a uWSGI instance to multiple sockets, you can use this option to map specific workers to specific sockets to implement a sort of in-process Quality of Service scheme. If you host multiple apps in the same uWSGI instance, you can easily dedicate resources to each of them.
- **modifier** (*Modifier*) – Socket routing modifier.

```
class uwsgi.conf.options.networking_sockets.SocketScgi (address,      nph=False,
                                                         bound_workers=None,
                                                         modifier=None)
```

Bind to the specified UNIX/TCP socket using SCGI protocol.

### Parameters

- **address** (*str|unicode|SocketShared*) – Address ([host]:port or socket file) to bind socket to.
- **nph** (*bool*) – Use NPH mode (“no-parsed-header” - bypass the server completely by sending the complete HTTP header directly to the browser).
- **bound\_workers** (*str|unicode|int|list*) – Map socket to specific workers. As you can bind a uWSGI instance to multiple sockets, you can use this option to map specific workers to specific sockets to implement a sort of in-process Quality of Service scheme. If you host multiple apps in the same uWSGI instance, you can easily dedicate resources to each of them.
- **modifier** (*Modifier*) – Socket routing modifier.

```
class uwsgi.conf.options.networking_sockets.SocketRaw(address,  
                                                    bound_workers=None, modi-  
                                                    fier=None)
```

Bind to the specified UNIX/TCP socket using RAW protocol.

Raw mode allows you to directly parse the request in your application callable. Instead of getting a list of CGI vars/headers in your callable you only get the file descriptor soon after accept().

You can then read()/write() to that file descriptor in full freedom.

---

**Note:** Raw mode disables request logging.

---

**Warning:** Use it as a low-level socket wrapper.

#### Parameters

- **address** (*str/unicode/SocketShared*) – Address ([host]:port or socket file) to bind socket to.
- **bound\_workers** (*str/unicode/int/list*) – Map socket to specific workers. As you can bind a uWSGI instance to multiple sockets, you can use this option to map specific workers to specific sockets to implement a sort of in-process Quality of Service scheme. If you host multiple apps in the same uWSGI instance, you can easily dedicate resources to each of them.
- **modifier** (*Modifier*) – Socket routing modifier.

```
class uwsgi.conf.options.networking_sockets.SocketShared(address,           unde-  
                                                    ferred=False,  
                                                    bound_workers=None,  
                                                    modifier=None)
```

Create a shared socket for advanced jailing or IPC purposes.

Allows you to create a socket early in the server's startup and use it after privileges drop or jailing. This can be used to bind to privileged (<1024) ports.

Shared sockets are a way to share sockets among various uWSGI components: you can use that to share a socket between the fastrouter and uWSGI instance.

#### Parameters

- **address** (*str/unicode*) – Address ([host]:port or socket file) to bind socket to.
- **undeferred** (*bool*) – Use shared socket undeferred mode.
- **bound\_workers** (*str/unicode/int/list*) – Map socket to specific workers. As you can bind a uWSGI instance to multiple sockets, you can use this option to map specific workers to specific sockets to implement a sort of in-process Quality of Service scheme. If you host multiple apps in the same uWSGI instance, you can easily dedicate resources to each of them.
- **modifier** (*Modifier*) – Socket routing modifier.

```
class uwsgi.conf.options.networking_sockets.SocketZeromq(address,  
                                                         bound_workers=None,  
                                                         modifier=None)
```

Introduce zeromq pub/sub pair.



### Parameters

- **address** (*str/unicode/SocketShared*) – Address ([host]:port or socket file) to bind socket to.
- **bound\_workers** (*str/unicode/int/list*) – Map socket to specific workers. As you can bind a uWSGI instance to multiple sockets, you can use this option to map specific workers to specific sockets to implement a sort of in-process Quality of Service scheme. If you host multiple apps in the same uWSGI instance, you can easily dedicate resources to each of them.
- **modifier** (*Modifier*) – Socket routing modifier.

**class** uwsgi.conf.options.networking.**Networking** (\*args, \*\*kwargs)  
Networking related stuff. Socket definition, binding and tuning.

#### **class sockets**

Available socket types to use with `.register_socket()`.

#### **default**

alias of `uwsgi.conf.options.networking_sockets.SocketDefault`

#### **fastcgi**

alias of `uwsgi.conf.options.networking_sockets.SocketFastcgi`

#### **http**

alias of `uwsgi.conf.options.networking_sockets.SocketHttp`

#### **https**

alias of `uwsgi.conf.options.networking_sockets.SocketHttps`

#### **raw**

alias of `uwsgi.conf.options.networking_sockets.SocketRaw`

#### **scgi**

alias of `uwsgi.conf.options.networking_sockets.SocketScgi`

#### **shared**

alias of `uwsgi.conf.options.networking_sockets.SocketShared`

#### **udp**

alias of `uwsgi.conf.options.networking_sockets.SocketUdp`

#### **uwsgi**

alias of `uwsgi.conf.options.networking_sockets.SocketUwsgi`

#### **uwsgis**

alias of `uwsgi.conf.options.networking_sockets.SocketUwsgis`

#### **zeromq**

alias of `uwsgi.conf.options.networking_sockets.SocketZeromq`

**classmethod from\_dsn** (*dsn, allow\_shared\_sockets=None*)

Constructs socket configuration object from DSN.

---

**Note:** This will also automatically use shared sockets to bind to privileged ports when non root.

---

### Parameters

- **dsn** (*str/unicode*) – Data source name, e.g: \* `http://127.0.0.1:8000` \* `https://127.0.0.1:443?cert=/here/there.crt&key=/that/my.key`

---

**Note:** Some schemas: fastcgi, http, https, raw, scgi, shared, udp, uwsgi, suwsgi, zeromq

---

- **allow\_shared\_sockets** (*bool*) – Allows using shared sockets to bind to privileged ports. If not provided automatic mode is enabled: shared are allowed if current user is not root.

:rtype Socket

**set\_basic\_params** (*queue\_size=None, freebind=None, default\_socket\_type=None*)

#### Parameters

- **queue\_size** (*int*) – Also known as a backlog. Every socket has an associated queue where request will be put waiting for a process to become ready to accept them. When this queue is full, requests will be rejected.

Default: 100 (an average value chosen by the maximum value allowed by default by your kernel).

---

**Note:** The maximum value is system/kernel dependent. Before increasing it you may need to increase your kernel limit too.

---

- **freebind** (*bool*) – Put socket in freebind mode. Allows binding to non-existent network addresses.

---

**Note:** Linux only.

---

- **default\_socket\_type** (*str|unicode*) – Force the socket type as default. See `.socket_types`.

**set\_socket\_params** (*send\_timeout=None, keep\_alive=None, no\_defer\_accept=None, buffer\_send=None, buffer\_receive=None*)

Sets common socket params.

#### Parameters

- **send\_timeout** (*int*) – Send (write) timeout in seconds.
- **keep\_alive** (*bool*) – Enable TCP KEEPALIVES.
- **no\_defer\_accept** (*bool*) – Disable deferred `accept()` on sockets by default (where available) uWSGI will defer the `accept()` of requests until some data is sent by the client (this is a security/performance measure). If you want to disable this feature for some reason, specify this option.
- **buffer\_send** (*int*) – Set `SO_SNDBUF` (bytes).
- **buffer\_receive** (*int*) – Set `SO_RCVBUF` (bytes).

**set\_unix\_socket\_params** (*abstract=None, permissions=None, owner=None, umask=None*)

Sets Unix-socket related params.

#### Parameters

- **abstract** (*bool*) – Force UNIX socket into abstract mode (Linux only).

- **permissions** (*str*) – UNIX sockets are filesystem objects that obey UNIX permissions like any other filesystem object.

You can set the UNIX sockets' permissions with this option if your webserver would otherwise have no access to the uWSGI socket. When used without a parameter, the permissions will be set to 666. Otherwise the specified chmod value will be used.

- **owner** (*str/unicode*) – Chown UNIX sockets.
- **umask** (*str/unicode*) – Set UNIX socket umask.

**set\_bsd\_socket\_params** (*port\_reuse=None*)

Sets BSD-sockets related params.

**Parameters** **port\_reuse** (*bool*) – Enable REUSE\_PORT flag on socket to allow multiple instances binding on the same address (BSD only).

**register\_socket** (*socket*)

Registers the given socket(s) for further use.

**Parameters** **socket** (*Socket|list[Socket]*) – Socket type object. See `.sockets`.

**set\_ssl\_params** (*verbose\_errors=None, sessions\_cache=None, sessions\_timeout=None, session\_context=None, raw\_options=None, dir\_tmp=None, client\_cert\_var=None*)

#### Parameters

- **verbose\_errors** (*bool*) – Be verbose about SSL errors.
- **sessions\_cache** (*str/unicode/bool*) – Use uWSGI cache for ssl sessions storage.

Accepts either bool or cache name string.

– <http://uwsgi.readthedocs.io/en/latest/SSLScaling.html>

**Warning:** Please be sure to configure cache before setting this.

- **sessions\_timeout** (*int*) – Set SSL sessions timeout in seconds. Default: 300.
- **session\_context** (*str/unicode*) – Session context identifying string. Can be set to static shared value to avoid session rejection.

Default: a value built from the HTTP server address.

– <http://uwsgi.readthedocs.io/en/latest/SSLScaling.html#setup-2-synchronize-caches-of-different-https-routers>

- **raw\_options** (*int|list[int]*) – Set a raw ssl option by its numeric value.
- **dir\_tmp** (*str/unicode*) – Store ssl-related temp files (e.g. pem data) in the specified directory.
- **client\_cert\_var** (*str/unicode*) – Export uWSGI variable HTTPS\_CC containing the raw client certificate.

**set\_sni\_params** (*name, cert, key, ciphers=None, client\_ca=None, wildcard=False*)

Allows setting Server Name Identification (virtual hosting for SSL nodes) params.

- <http://uwsgi.readthedocs.io/en/latest/SNI.html>

#### Parameters

- **name** (*str/unicode*) – Node/server/host name.

- **cert** (*str/unicode*) – Certificate file.
- **key** (*str/unicode*) – Private key file.
- **ciphers** (*str/unicode*) – Ciphers [alias] string.

**Example:**

- DEFAULT
- HIGH
- DHE, EDH
  
- <https://www.openssl.org/docs/man1.1.0/apps/ciphers.html>

- **client\_ca** (*str/unicode*) – Client CA file for client-based auth.
- **wildcard** (*bool*) – Allow regular expressions in *name* (used for wildcard certificates).

**set\_sni\_dir\_params** (*dir, ciphers=None*)

Enable checking for cert/key/client\_ca file in the specified directory and create a sni/ssl context on demand.

**Expected filenames:**

- <sni-name>.cert
- <sni-name>.key
- <sni-name>.ca - this file is optional
  
- <http://uwsgi.readthedocs.io/en/latest/SNI.html#massive-sni-hosting>

**Parameters**

- **dir** (*str/unicode*) –
- **ciphers** (*str/unicode*) – Ciphers [alias] string.

**Example:**

- DEFAULT
- HIGH
- DHE, EDH
  
- <https://www.openssl.org/docs/man1.1.0/apps/ciphers.html>

### 3.8.12 Queue

**class** uwsgi.conf.options.queue.**Queue** (*\*args, \*\*kwargs*)  
Queue.

At the low level it is a simple block-based shared array, with two optional counters, one for stack-style, LIFO usage, the other one for FIFO.

<http://uwsgi-docs.readthedocs.io/en/latest/Queue.html>

**enable** (*size, block\_size=None, store=None, store\_sync\_interval=None*)  
Enables shared queue of the given size.

**Parameters**

- **size** (*int*) – Queue size.
- **block\_size** (*int*) – Block size in bytes. Default: 8 KiB.
- **store** (*str/unicode*) – Persist the queue into file.
- **store\_sync\_interval** (*int*) – Store sync interval in master cycles (usually seconds).

### 3.8.13 Routing

**Actions**

**class** uwsgiconf.options.routing\_actions.**RouteAction** (*\*args*)

**class** uwsgiconf.options.routing\_actions.**ActionFlush**

Send the current contents of the transformation buffer to the client (without clearing the buffer).

- <http://uwsgi.readthedocs.io/en/latest/Transformations.html#flushing-magic>

**class** uwsgiconf.options.routing\_actions.**ActionGzip**

Encodes the response buffer to gzip.

**class** uwsgiconf.options.routing\_actions.**ActionToFile** (*filename, mode=None*)

Used for caching a response buffer into a static file.

**class** uwsgiconf.options.routing\_actions.**ActionUpper**

Transforms each character in uppercase.

Mainly as an example of transformation plugin.

**class** uwsgiconf.options.routing\_actions.**ActionChunked**

Encodes the output in HTTP chunked.

**class** uwsgiconf.options.routing\_actions.**ActionTemplate**

Allows using a template file to expose everything from internal routing system into it.

- <http://uwsgi.readthedocs.io/en/latest/Changelog-1.9.19.html#the-template-transformation>

**class** uwsgiconf.options.routing\_actions.**ActionFixContentLen** (*add\_header=False*)

Fixes Content-length header.

**Parameters** **add\_header** (*bool*) – Force header add instead of plain fix of existing header.

**class** uwsgiconf.options.routing\_actions.**ActionDoContinue**

Stop scanning the internal routing table and continue to the selected request handler.

**class** uwsgiconf.options.routing\_actions.**ActionDoBreak** (*code, return\_body=False*)

Stop scanning the internal routing table and close the request.

**Parameters**

- **code** (*int*) – HTTP code
- **return\_body** – Uses uWSGI's built-in status code and returns both status code and message body.

**class** uwsgiconf.options.routing\_actions.**ActionLog** (*message*)

Print the specified message in the logs or do not log a request is message is None.

**Parameters** **message** (*str/unicode/None*) – Message to add into log. If None logging will be disabled for this request.

**class** uwsgi.conf.options.routing\_actions.**ActionOffloadOff**  
Do not use offloading.

**class** uwsgi.conf.options.routing\_actions.**ActionAddVarLog** (*name*, *val*)  
Add the specified logvar.

**Parameters**

- **name** (*str/unicode*) – Variable name.
- **val** – Variable value.

**class** uwsgi.conf.options.routing\_actions.**ActionDoGoto** (*where*)  
Make a forward jump to the specified label or rule position.

**Parameters** **where** (*str/unicode/int*) – Rule number of label to go to.

**class** uwsgi.conf.options.routing\_actions.**ActionAddVarCgi** (*name*, *val*)  
Add the specified CGI (environment) variable to the request.

**Parameters**

- **name** (*str/unicode*) – Variable name.
- **val** – Variable value.

**class** uwsgi.conf.options.routing\_actions.**ActionHeaderAdd** (*name*, *val*)  
Add the specified HTTP header to the response.

**Parameters**

- **name** (*str/unicode*) – Header name.
- **val** – Header value.

**class** uwsgi.conf.options.routing\_actions.**ActionHeaderRemove** (*name*)  
Remove the specified HTTP header from the response.

**Parameters** **name** (*str/unicode*) – Header name.

**class** uwsgi.conf.options.routing\_actions.**ActionHeadersOff**  
Disable headers.

**class** uwsgi.conf.options.routing\_actions.**ActionHeadersReset** (*code*)  
Clear the response headers, setting a new HTTP status code, useful for resetting a response.

**Parameters** **code** (*int*) – HTTP code.

**class** uwsgi.conf.options.routing\_actions.**ActionSignal** (*num*)  
Raise the specified uwsgi signal.

**Parameters** **num** (*int*) – Signal number.

**class** uwsgi.conf.options.routing\_actions.**ActionSend** (*data*, *crnl=False*)  
Extremely advanced (and dangerous) function allowing you to add raw data to the response.

**Parameters**

- **data** – Data to add to response.
- **crnl** (*bool*) – Add carriage return and new line.

**class** uwsgi.conf.options.routing\_actions.**ActionRedirect** (*url*, *permanent=False*)  
Return a HTTP 301/302 Redirect to the specified URL.

**Parameters**

- **unicode url** (*str/*) – URL to redirect to.

- **permanent** (*bool*) – If True use 301, otherwise 302.

**class** uwsgi.conf.options.routing\_actions.**ActionRewrite** (*rule, do\_continue=False*)  
A rewriting engine inspired by Apache mod\_rewrite.

Rebuild PATH\_INFO and QUERY\_STRING according to the specified rules before the request is dispatched to the request handler.

#### Parameters

- **rule** (*str/unicode*) – A rewrite rule.
- **do\_continue** (*bool*) – Stop request processing and continue to the selected request handler.

**class** uwsgi.conf.options.routing\_actions.**ActionRouteUwsgi** (*external\_address=”, modifier=”, app=”*)

Rewrite the modifier1, modifier2 and optionally UWSGI\_APPID values of a request or route the request to an external uwsgi server.

#### Parameters

- **external\_address** (*str/unicode*) – External uWSGI server address (host:port).
- **modifier** (*Modifier*) – Set request modifier.
- **app** (*str/unicode*) – Set UWSGI\_APPID.

**class** uwsgi.conf.options.routing\_actions.**ActionRouteExternal** (*address, host\_header=None*)

Route the request to an external HTTP server.

#### Parameters

- **address** (*str/unicode*) – External HTTP address (host:port)
- **host\_header** (*str/unicode*) – HOST header value.

**class** uwsgi.conf.options.routing\_actions.**ActionAlarm** (*name, message*)  
Triggers an alarm.

- <http://uwsgi.readthedocs.io/en/latest/Changelog-1.9.6.html#the-alarm-routing-action>

#### Parameters

- **name** (*str/unicode*) – Alarm name
- **message** (*str/unicode*) – Message to pass into alarm.

**class** uwsgi.conf.options.routing\_actions.**ActionServeStatic** (*fpath*)  
Serve a static file from the specified physical path.

**Parameters** **fpath** (*str/unicode*) – Static file path.

**class** uwsgi.conf.options.routing\_actions.**ActionAuthBasic** (*realm, user=None, password=None, do\_next=False*)

Use Basic HTTP Auth.

#### Parameters

- **realm** (*str/unicode*) –
- **user** (*str/unicode*) –
- **password** (*str/unicode*) – Password or htpasswd-like file.

- **do\_next** (*bool*) – Allow next rule.

```
class uwsgi.conf.options.routing_actions.AuthLdap (realm, address,  
                                                base_dn=None, bind_dn=None,  
                                                bind_password=None, filter=None,  
                                                login_attr=None, log_level=None,  
                                                do_next=False)
```

Use Basic HTTP Auth.

#### Parameters

- **realm** (*str/unicode*) –
- **address** (*str/unicode*) – LDAP server URI
- **base\_dn** (*str/unicode*) – Base DN used when searching for users.
- **bind\_dn** (*str/unicode*) – DN used for binding. Required if the LDAP server does not allow anonymous searches.
- **bind\_password** (*str/unicode*) – Password for the `bind_dn` user.
- **filter** (*str/unicode*) – Filter used when searching for users. Default: (objectClass=\*)
- **login\_attr** (*str/unicode*) – LDAP attribute that holds user login. Default: uid.
- **log\_level** (*str/unicode*) – Log level.

#### Supported values:

- 0 - don't log any binds
- 1 - log authentication errors,
- 2 - log both successful and failed binds

- **do\_next** (*bool*) – Allow next rule.

```
class uwsgi.conf.options.routing_actions.ActionSetHarakiri (timeout)  
Set harakiri timeout for the current request.
```

**Parameters** **timeout** (*int*) –

```
class uwsgi.conf.options.routing_actions.ActionDirChange (dir)  
Changes a directory.
```

**Parameters** **dir** (*str/unicode*) – Directory to change into.

```
class uwsgi.conf.options.routing_actions.ActionSetVarUwsgiAppid (app)  
Set UWSGI_APPID.
```

Bypass `SCRIPT_NAME` and `VirtualHosting` to let the user choose the mountpoint without limitations (or headaches).

The concept is very generic: `UWSGI_APPID` is the identifier of an application. If it is not found in the internal list of apps, it will be loaded.

**Parameters** **app** (*str/unicode*) – Application ID.

```
class uwsgi.conf.options.routing_actions.ActionSetVarRemoteUser (user)  
Set REMOTE_USER
```

**Parameters** **user** (*str/unicode*) – Username.

```
class uwsgi.conf.options.routing_actions.ActionSetVarUwsgiHome (dir)  
Set UWSGI_HOME
```



**Parameters** `dir` (*str/unicode*) – Directory to make a new home.

**class** `uwsgi.conf.options.routing_actions.ActionSetVarUwsgiScheme` (*value*)  
Set UWSGI\_SCHEME.

Set the URL scheme when it cannot be reliably determined. This may be used to force HTTPS (with the value `https`), for instance.

- <http://uwsgi.readthedocs.io/en/latest/Changelog-1.9.6.html#configuring-dynamic-apps-with-internal-routing>

**Parameters** `value` (*str/unicode*) –

**class** `uwsgi.conf.options.routing_actions.ActionSetVarScriptName` (*name*)  
Set SCRIPT\_NAME

**Parameters** `name` (*str/unicode*) – Script name

**class** `uwsgi.conf.options.routing_actions.ActionSetVarRequestMethod` (*name*)  
Set REQUEST\_METHOD

**Parameters** `name` (*str/unicode*) – Method name.

**class** `uwsgi.conf.options.routing_actions.ActionSetVarRequestUri` (*value*)  
Set REQUEST\_URI

**Parameters** `value` (*str/unicode*) – URI

**class** `uwsgi.conf.options.routing_actions.ActionSetVarRemoteAddr` (*value*)  
Set REMOTE\_ADDR

**Parameters** `value` (*str/unicode*) – Address.

**class** `uwsgi.conf.options.routing_actions.ActionSetVarPathInfo` (*value*)  
Set PATH\_INFO

**Parameters** `value` (*str/unicode*) – New info.

**class** `uwsgi.conf.options.routing_actions.ActionSetVarDocumentRoot` (*value*)  
Set DOCUMENT\_ROOT

**Parameters** `value` (*str/unicode*) –

**class** `uwsgi.conf.options.routing_actions.ActionSetUwsgiProcessName` (*name*)  
Set uWSGI process name.

**Parameters** `name` (*str/unicode*) – New process name.

**class** `uwsgi.conf.options.routing_actions.ActionFixVarPathInfo`  
Fixes PATH\_INFO taking into account script name.

This action allows you to set SCRIPT\_NAME in nginx without bothering to rewrite the PATH\_INFO (something nginx cannot afford).

- <http://uwsgi.readthedocs.io/en/latest/Changelog-2.0.11.html#fixpathinfo-routing-action>

**class** `uwsgi.conf.options.routing_actions.ActionSetScriptFile` (*fpath*)  
Set script file.

- <http://uwsgi.readthedocs.io/en/latest/Changelog-1.9.6.html#configuring-dynamic-apps-with-internal-routing>

**Parameters** `fpath` (*str/unicode*) – File path.

## Dedicated routers

**class** uwsgi.conf.options.routing\_routers.**RouterBase** (*on=None*)

**Parameters** **on** (*SocketShared|str|unicode*) – Activates the router on the given address.

**class** uwsgi.conf.options.routing\_routers.**Forwarder** (*\*args*)

**class** uwsgi.conf.options.routing\_routers.**ForwarderPath** (*sockets\_dir*)

Use the specified base (allows %s pattern) for mapping requests to UNIX sockets.

### Examples:

- /tmp/sockets/
- /tmp/sockets/%s/uwsgi.sock
  
- <http://uwsgi.readthedocs.io/en/latest/Fastrouter.html#way-1-fastrouter-use-base>
- <http://uwsgi.readthedocs.io/en/latest/Fastrouter.html#way-2-fastrouter-use-pattern>

**Parameters** **sockets\_dir** (*str|unicode*) – UNIX sockets directory. Allows %s to denote key (domain).

**class** uwsgi.conf.options.routing\_routers.**ForwarderCode** (*script, func, modifier=None*)

Forwards requests to nodes returned by a function.

This allows using user defined functions to calculate. Function must accept key (domain).

- <http://uwsgi.readthedocs.io/en/latest/Fastrouter.html#way-5-fastrouter-use-code-string>

**Warning:** Remember to not put blocking code in your functions. The router is totally non-blocking, do not ruin it!

### Parameters

- **script** (*str|unicode*) – Script (module for Python) name to get function from.
- **func** (*str|unicode*) – Function name.
- **modifier** (*Modifier*) – Routing modifier.

**class** uwsgi.conf.options.routing\_routers.**ForwarderCache** (*cache\_name=None*)

Uses uWSGI cache to get target nodes from.

- <http://uwsgi.readthedocs.io/en/latest/Fastrouter.html#way-3-fastrouter-use-cache>

**Parameters** **cache\_name** (*str|unicode*) – Cache name to use.

**class** uwsgi.conf.options.routing\_routers.**ForwarderSocket** (*socket*)

Forwards request to the specified uwsgi socket.

**Parameters** **socket** (*str|unicode*) – Socket filepath.

**class** uwsgi.conf.options.routing\_routers.**ForwarderSubscriptionServer** (*address*)

Forwards requests to nodes returned by the subscription server.

Subscriptions are simple UDP packets that instruct the router which domain maps to which instance or instances.

To subscribe to such a subscription server use `.subscriptions.subscribe()`.

- <http://uwsgi.readthedocs.io/en/latest/Fastrouter.html#way-4-fastrouter-subscription-server>

**Parameters** `address` (*str/unicode*) – Address (including port) to run the subscription server on.

**class** `uwsgi.conf.options.routing_routers.RouterHttp` (*on=None, forward\_to=None*)  
uWSGI includes an HTTP router/proxy/load-balancer that can forward requests to uWSGI workers.

The server can be used in two ways:

- embedded - automatically spawn workers and setup the communication socket
- standalone - you have to specify the address of a uwsgi socket to connect to

See `subscribe_to` argument to `.set_basic_params()`

---

**Note:** If you want to go massive (virtualhosting and zero-conf scaling) combine the HTTP router with the uWSGI Subscription Server.

---

Activates the router on the given address.

#### Parameters

- **on** (*SocketShared/str/unicode*) – Activates the router on the given address.
- **forward\_to** (*Forwarder/str/unicode/list[str/unicode]*) – Where to forward requests. Expects a forwarder instance or one or more node names.

**set\_basic\_params** (*workers=None, zerg\_server=None, fallback\_node=None, concurrent\_events=None, cheap\_mode=None, stats\_server=None, quiet=None, buffer\_size=None, keepalive=None, resubscribe\_addresses=None*)

#### Parameters

- **workers** (*int*) – Number of worker processes to spawn.
- **zerg\_server** (*str/unicode*) – Attach the router to a zerg server.
- **fallback\_node** (*str/unicode*) – Fallback to the specified node in case of error.
- **concurrent\_events** (*int*) – Set the maximum number of concurrent events router can manage.

Default: system dependent.

- **cheap\_mode** (*bool*) – Enables cheap mode. When the router is in cheap mode, it will not respond to requests until a node is available. This means that when there are no nodes subscribed, only your local app (if any) will respond. When all of the nodes go down, the router will return in cheap mode.
- **stats\_server** (*str/unicode*) – Router stats server address to run at.
- **quiet** (*bool*) – Do not report failed connections to instances.
- **buffer\_size** (*int*) – Set internal buffer size in bytes. Default: page size.
- **keepalive** (*int*) – Allows holding the connection open even if the request has a body.  
– <http://uwsgi.readthedocs.io/en/latest/HTTP.html#http-keep-alive>

---

**Note:** See http11 socket type for an alternative.

---

- **resubscribe\_addresses** (*str|unicode|list[str|unicode]*) – Forward subscriptions to the specified subscription server.

**set\_connections\_params** (*harakiri=None, timeout\_socket=None, retry\_delay=None, timeout\_headers=None, timeout\_backend=None*)  
Sets connection-related parameters.

#### Parameters

- **harakiri** (*int*) – Set gateway harakiri timeout (seconds).
- **timeout\_socket** (*int*) – Node socket timeout (seconds). Used to set the SPDY timeout. This is the maximum amount of inactivity after the SPDY connection is closed.  
Default: 60.
- **retry\_delay** (*int*) – Retry connections to dead static nodes after the specified amount of seconds. Default: 30.
- **timeout\_headers** (*int*) – Defines the timeout (seconds) while waiting for http headers.  
Default: *socket\_timeout*.
- **timeout\_backend** (*int*) – Defines the timeout (seconds) when connecting to backend instances.  
Default: *socket\_timeout*.

**set\_manage\_params** (*chunked\_input=None, chunked\_output=None, gzip=None, websockets=None, source\_method=None, rtsp=None, proxy\_protocol=None*)  
Allows enabling various automatic management mechanics.

- <http://uwsgi.readthedocs.io/en/latest/Changelog-1.9.html#http-router-keepalive-auto-chunking-auto-gzip-and-transp>

#### Parameters

- **chunked\_input** (*bool*) – Automatically detect chunked input requests and put the session in raw mode.
- **chunked\_output** (*bool*) – Automatically transform output to chunked encoding during HTTP 1.1 keepalive (if needed).
- **gzip** (*bool*) – Automatically gzip content if uWSGI-Encoding header is set to gzip, but content size (Content-Length/Transfer-Encoding) and Content-Encoding are not specified.
- **websockets** (*bool*) – Automatically detect websockets connections and put the session in raw mode.
- **source\_method** (*bool*) – Automatically put the session in raw mode for *SOURCE* HTTP method.  
– <http://uwsgi.readthedocs.io/en/latest/Changelog-2.0.5.html#icecast2-protocol-helpers>
- **rtsp** (*bool*) – Allow the HTTP router to detect RTSP and chunked requests automatically.
- **proxy\_protocol** (*bool*) – Allows the HTTP router to manage PROXY1 protocol requests, such as those made by Haproxy or Amazon Elastic Load Balancer (ELB).

**set\_owner\_params** (*uid=None, gid=None*)

Drop http router privileges to specified user and group.

#### Parameters

- **uid** (*str/unicode/int*) – Set uid to the specified username or uid.
- **gid** (*str/unicode/int*) – Set gid to the specified groupname or gid.

**class** uwsgi.conf.options.routing\_routers.**RouterHttps** (*on, cert, key, ciphers=None, client\_ca=None, session\_context=None, use\_spdy=None, export\_cert\_var=None*)

uWSGI includes an HTTPS router/proxy/load-balancer that can forward requests to uWSGI workers.

The server can be used in two ways:

- embedded - automatically spawn workers and setup the communication socket
- standalone - you have to specify the address of a uwsgi socket to connect to

See *subscribe\_to* argument to *.set\_basic\_params()*

---

**Note:** If you want to go massive (virtualhosting and zero-conf scaling) combine the HTTP router with the uWSGI Subscription Server.

---

Binds https router to run on the given address.

#### Parameters

- **on** (*SocketShared/str/unicode*) – Activates the router on the given address.
- **cert** (*str/unicode*) – Certificate file.
- **key** (*str/unicode*) – Private key file.
- **ciphers** (*str/unicode*) – Ciphers [alias] string.

#### Example:

- DEFAULT
- HIGH
- DHE, EDH

– <https://www.openssl.org/docs/man1.1.0/apps/ciphers.html>

- **client\_ca** (*str/unicode*) – Client CA file for client-based auth.
- **session\_context** (*str/unicode*) – Session context identifying string. Can be set to static shared value to avoid session rejection.

Default: a value built from the HTTP server address.

– <http://uwsgi.readthedocs.io/en/latest/SSLScaling.html#setup-2-synchronize-caches-of-different-https-routers>

- **use\_spdy** (*bool*) – Use SPDY.
- **export\_cert\_var** (*bool*) – Export uwsgi variable *HTTPS\_CC* containing the raw client certificate.

**class** uwsgi.conf.options.routing\_routers.**RouterSsl** (*on, cert, key, forward\_to=None, ciphers=None, client\_ca=None, session\_context=None, use\_sni=None*)

Works in the same way as the RouterRaw, but will terminate ssl connections.

Supports SNI for implementing virtual hosting.

Activates the router on the given address.

#### Parameters

- **on** (*SocketShared|str|unicode*) – Activates the router on the given address.
- **cert** (*str|unicode*) – Certificate file.
- **key** (*str|unicode*) – Private key file.
- **forward\_to** (*Forwarder|str|unicode|list[str|unicode]*) – Where to forward requests. Expects a forwarder instance or one or more node names.
- **ciphers** (*str|unicode*) – Ciphers [alias] string.

#### Example:

- DEFAULT
- HIGH
- DHE, EDH

– <https://www.openssl.org/docs/man1.1.0/apps/ciphers.html>

- **client\_ca** (*str|unicode*) – Client CA file for client-based auth.
- **session\_context** (*str|unicode*) – Session context identifying string. Can be set to static shared value to avoid session rejection.

Default: a value built from the HTTP server address.

– <http://uwsgi.readthedocs.io/en/latest/SSLScaling.html#setup-2-synchronize-caches-of-different-https-routers>

- **use\_sni** (*bool*) – Use SNI to route requests.

**set\_connections\_params** (*harakiri=None, timeout\_socket=None, retry\_delay=None, retry\_max=None*)

Sets connection-related parameters.

#### Parameters

- **harakiri** (*int*) – Set gateway harakiri timeout (seconds).
- **timeout\_socket** (*int*) – Node socket timeout (seconds). Default: 60.
- **retry\_delay** (*int*) – Retry connections to dead static nodes after the specified amount of seconds. Default: 30.
- **retry\_max** (*int*) – Maximum number of retries/fallbacks to other nodes. Default: 3.

**class** uwsgi.conf.options.routing\_routers.**RouterFast** (*on=None, forward\_to=None*)

A proxy/load-balancer/router speaking the uwsgi protocol.

You can put it between your webserver and real uWSGI instances to have more control over the routing of HTTP requests to your application servers.

Activates the router on the given address.

#### Parameters

- **on** (*SocketShared|str|unicode*) – Activates the router on the given address.
- **forward\_to** (*Forwarder|str|unicode|list[str|unicode]*) – Where to forward requests. Expects a forwarder instance or one or more node names.

**set\_basic\_params** (*workers=None, zerg\_server=None, fallback\_node=None, concurrent\_events=None, cheap\_mode=None, stats\_server=None, quiet=None, buffer\_size=None, fallback\_nokey=None, subscription\_key=None, emperor\_command\_socket=None*)

#### Parameters

- **workers** (*int*) – Number of worker processes to spawn.
- **zerg\_server** (*str|unicode*) – Attach the router to a zerg server.
- **fallback\_node** (*str|unicode*) – Fallback to the specified node in case of error.
- **concurrent\_events** (*int*) – Set the maximum number of concurrent events router can manage.

Default: system dependent.

- **cheap\_mode** (*bool*) – Enables cheap mode. When the router is in cheap mode, it will not respond to requests until a node is available. This means that when there are no nodes subscribed, only your local app (if any) will respond. When all of the nodes go down, the router will return in cheap mode.
- **stats\_server** (*str|unicode*) – Router stats server address to run at.
- **quiet** (*bool*) – Do not report failed connections to instances.
- **buffer\_size** (*int*) – Set internal buffer size in bytes. Default: page size.
- **fallback\_nokey** (*bool*) – Move to fallback node even if a subscription key is not found.
- **subscription\_key** (*str|unicode*) – Skip uwsgi parsing and directly set a key.
- **emperor\_command\_socket** (*str|unicode*) – Set the emperor command socket that will receive spawn commands.

See `.empire.set_emperor_command_params()`.

**set\_resubscription\_params** (*addresses=None, bind\_to=None*)

You can specify a dgram address (udp or unix) on which all of the subscriptions request will be forwarded to (obviously changing the node address to the router one).

The system could be useful to build ‘federated’ setup.

- <http://uwsgi.readthedocs.io/en/latest/Changelog-2.0.1.html#resubscriptions>

#### Parameters

- **addresses** (*str|unicode|list[str|unicode]*) – Forward subscriptions to the specified subscription server.
- **bind\_to** (*str|unicode|list[str|unicode]*) – Bind to the specified address when re-subscribing.

**set\_connections\_params** (*harakiri=None, timeout\_socket=None, retry\_delay=None, retry\_max=None, defer=None*)

Sets connection-related parameters.

#### Parameters

- **harakiri** (*int*) – Set gateway harakiri timeout (seconds).
- **timeout\_socket** (*int*) – Node socket timeout (seconds). Default: 60.
- **retry\_delay** (*int*) – Retry connections to dead static nodes after the specified amount of seconds. Default: 30.
- **retry\_max** (*int*) – Maximum number of retries/fallbacks to other nodes. Default: 3
- **defer** (*int*) – Defer connection delay, seconds. Default: 5.

**set\_postbuffering\_params** (*size=None, store\_dir=None*)

Sets buffering params.

Web-proxies like nginx are “buffered”, so they wait til the whole request (and its body) has been read, and then it sends it to the backends.

#### Parameters

- **size** (*int*) – The size (in bytes) of the request body after which the body will be stored to disk (as a temporary file) instead of memory.
- **store\_dir** (*str|unicode*) – Put buffered files to the specified directory. Default: TMPDIR, /tmp/

**set\_owner\_params** (*uid=None, gid=None*)

Drop http router privileges to specified user and group.

#### Parameters

- **uid** (*str|unicode|int*) – Set uid to the specified username or uid.
- **gid** (*str|unicode|int*) – Set gid to the specified groupname or gid.

**class** uwsgi.conf.options.routing\_routers.**RouterRaw** (*on=None, forward\_to=None*)

A pure-TCP load balancer.

Can be used to load balance between the various HTTPS routers.

Activates the router on the given address.

#### Parameters

- **on** (*SocketShared|str|unicode*) – Activates the router on the given address.
- **forward\_to** (*Forwarder|str|unicode|list[str|unicode]*) – Where to forward requests. Expects a forwarder instance or one or more node names.

**set\_connections\_params** (*harakiri=None, timeout\_socket=None, retry\_delay=None, retry\_max=None, use\_xclient=None*)

Sets connection-related parameters.

#### Parameters

- **harakiri** (*int*) – Set gateway harakiri timeout (seconds).
- **timeout\_socket** (*int*) – Node socket timeout (seconds). Default: 60.
- **retry\_delay** (*int*) – Retry connections to dead static nodes after the specified amount of seconds. Default: 30.



- **retry\_max** (*int*) – Maximum number of retries/fallbacks to other nodes. Default: 3.
- **use\_xclient** (*bool*) – Use the xclient protocol to pass the client address.

**class** uwsgi.conf.options.routing\_routers.RouterForkPty (*on=None, undeferred=False*)  
Allows allocation of pseudoterminals in jails.

Dealing with containers is now a common deployment pattern. One of the most annoying tasks when dealing with jails/namespaces is ‘attaching’ to already running instances. The forkpty router aims at simplifying the process giving a pseudoterminal server to your uWSGI instances. A client connect to the socket exposed by the forkpty router and get a new pseudoterminal connected to a process (generally a shell, but can be whatever you want).

---

**Note:** To be used in cooperation with *pty* plugin.

---

Binds router to run on the given address.

#### Parameters

- **on** (*SocketShared|str|unicode*) – Activates the router on the given address.
- **undeferred** (*bool*) – Run router in undeferred mode.

**set\_basic\_params** (*workers=None, zerg\_server=None, fallback\_node=None, concurrent\_events=None, cheap\_mode=None, stats\_server=None, run\_command=None*)

#### Parameters

- **workers** (*int*) – Number of worker processes to spawn.
- **zerg\_server** (*str|unicode*) – Attach the router to a zerg server.
- **fallback\_node** (*str|unicode*) – Fallback to the specified node in case of error.
- **concurrent\_events** (*int*) – Set the maximum number of concurrent events router can manage.  
Default: system dependent.
- **cheap\_mode** (*bool*) – Enables cheap mode. When the router is in cheap mode, it will not respond to requests until a node is available. This means that when there are no nodes subscribed, only your local app (if any) will respond. When all of the nodes go down, the router will return in cheap mode.
- **stats\_server** (*str|unicode*) – Router stats server address to run at.
- **run\_command** (*str|unicode*) – Run the specified command on every connection. Default: /bin/sh.

**set\_connections\_params** (*harakiri=None, timeout\_socket=None*)  
Sets connection-related parameters.

#### Parameters

- **harakiri** (*int*) – Set gateway harakiri timeout (seconds).
- **timeout\_socket** (*int*) – Node socket timeout (seconds). Default: 60.

**set\_window\_params** (*cols=None, rows=None*)  
Sets pty window params.

**Parameters**

- **cols** (*int*) –
- **rows** (*int*) –

**class** uwsgi.conf.options.routing\_routers.**RouterTunTap** (*on=None, device=None, stats\_server=None, gateway=None*)

The tuntap router is a non-blocking highly optimized ip router translating from tuntap device to socket streams.

Allows full user-space networking in jails.

It is meant as a replacement for the currently available networking namespaces approaches. Compared to *veth* or *macvlan* it is really simple and allows total control over the routing subsystem (in addition to a simple customizable firewalling engine).

Generally you spawn the tuntap router in the Emperor instance. Vassals will run in new namespaces in which they create a tuntap device attached to the tuntap router. UNIX sockets are the only way to connect to the tuntap router after jailing.

Vassals should connect to tuntap device.

Passing params will create a router device.

**Parameters**

- **on** (*str|unicode*) – Socket file.
- **device** (*str|unicode*) – Device name.
- **stats\_server** (*str|unicode*) – Router stats server address to run at.
- **gateway** (*str|unicode*) – Gateway address.

**set\_basic\_params** (*use\_credentials=None, stats\_server=None*)

**Parameters**

- **use\_credentials** (*str|unicode*) – Enable check of SCM\_CREDENTIALS for tuntap client/server.
- **stats\_server** (*str|unicode*) – Router stats server address to run at.

**register\_route** (*src, dst, gateway*)

Adds a routing rule to the tuntap router.

**Parameters**

- **src** (*str|unicode*) – Source/mask.
- **dst** (*str|unicode*) – Destination/mask.
- **gateway** (*str|unicode*) – Gateway address.

**device\_connect** (*socket, device\_name*)

Add a tuntap device to the instance.

To be used in a vassal.

**Parameters**

- **socket** (*str|unicode*) – Router socket.  
Example: */run/tuntap\_router.socket*.
- **device\_name** (*str|unicode*) – Device.  
Example: *uwsgi0*.

**device\_add\_rule** (*direction, action, src, dst, target=None*)

Adds a tuntap device rule.

To be used in a vassal.

#### Parameters

- **direction** (*str/unicode*) – Direction:
  - in
  - out.
- **action** (*str/unicode*) – Action:
  - allow
  - deny
  - route
  - gateway.
- **src** (*str/unicode*) – Source/mask.
- **dst** (*str/unicode*) – Destination/mask.
- **target** (*str/unicode*) – Depends on action.
  - Route / Gateway: Accept addr:port

**add\_firewall\_rule** (*direction, action, src=None, dst=None*)

Adds a firewall rule to the router.

The TunTap router includes a very simple firewall for governing vassal’s traffic. The first matching rule stops the chain, if no rule applies, the policy is “allow”.

#### Parameters

- **direction** (*str/unicode*) – Direction:
  - in
  - out
- **action** (*str/unicode*) – Action:
  - allow
  - deny
- **src** (*str/unicode*) – Source/mask.
- **dst** (*str/unicode*) – Destination/mask

## Modifiers

**class** uwsgi.conf.options.routing\_modifiers.**ModifierWsgi** (*submod=None*)

Standard WSGI request followed by the HTTP request body.

**class** uwsgi.conf.options.routing\_modifiers.**ModifierPsgi** (*submod=None*)

Standard PSGI request followed by the HTTP request body.

**class** uwsgi.conf.options.routing\_modifiers.**ModifierLua** (*submod=None*)

Standard LUA/WSAPI request followed by the HTTP request body.

**class** uwsgi.conf.options.routing\_modifiers.**ModifierRack** (*submod=None*)  
Standard RACK request followed by the HTTP request body.

**class** uwsgi.conf.options.routing\_modifiers.**ModifierJvm** (*submod=None*)  
Standard JVM request for The JWSGI interface and The Clojure/Ring JVM request handler followed by the HTTP request body.

**SUB\_RING = 1**  
Use Clojure/Ring JVM request handler.

**class** uwsgi.conf.options.routing\_modifiers.**ModifierCgi** (*submod=None*)  
Standard Running CGI scripts on uWSGI request followed by the HTTP request body.

**class** uwsgi.conf.options.routing\_modifiers.**ModifierManage** (*submod=None*)  
Management interface request: setup flag specified by `modifier2`.  
  
For a list of management flag look at `ManagementFlag`.

**class** uwsgi.conf.options.routing\_modifiers.**ModifierPhp** (*submod=None*)  
Standard Running PHP scripts in uWSGI request followed by the HTTP request body.

**class** uwsgi.conf.options.routing\_modifiers.**ModifierMono** (*submod=None*)  
Standard The Mono ASP.NET plugin request followed by the HTTP request body.

**class** uwsgi.conf.options.routing\_modifiers.**ModifierSpooler** (*submod=None*)  
The uWSGI Spooler request, the block vars is converted to a dictionary/hash/table and passed to the spooler callable.

**class** uwsgi.conf.options.routing\_modifiers.**ModifierSymcall** (*submod=None*)  
Direct call to C-like symbols.

**class** uwsgi.conf.options.routing\_modifiers.**ModifierEval** (*submod=None*)  
Raw Code evaluation. The interpreter is chosen by the `modifier2`.  
  
..note:: It does not return a valid uwsgi response, but a raw string (that may be an HTTP response).

**class** uwsgi.conf.options.routing\_modifiers.**ModifierXslt** (*submod=None*)  
Invoke the The XSLT plugin.

**class** uwsgi.conf.options.routing\_modifiers.**ModifierV8** (*submod=None*)  
Invoke the uWSGI V8 support.

**class** uwsgi.conf.options.routing\_modifiers.**ModifierGridfs** (*submod=None*)  
Invoke the The GridFS plugin.

**class** uwsgi.conf.options.routing\_modifiers.**ModifierFastfunc** (*submod=None*)  
Call the FastFuncs specified by the `modifier2` field.

**class** uwsgi.conf.options.routing\_modifiers.**ModifierGlusterfs** (*submod=None*)  
Invoke the The GlusterFS plugin.

**class** uwsgi.conf.options.routing\_modifiers.**ModifierRados** (*submod=None*)  
Invoke the The RADOS plugin.

**class** uwsgi.conf.options.routing\_modifiers.**ModifierManagePathInfo** (*submod=None*)  
Standard WSGI request followed by the HTTP request body.  
  
The `PATH_INFO` is automatically modified, removing the `SCRIPT_NAME` from it.

**class** uwsgi.conf.options.routing\_modifiers.**ModifierMessage** (*submod=None*)  
Generic message passing (reserved).

**class** uwsgi.conf.options.routing\_modifiers.**ModifierMessageArray** (*submod=None*)  
Array of char passing (reserved).

---

```

class uwsgi.conf.options.routing_modifiers.ModifierMessageMarshal (submod=None)
    Marshalled/serialized object passing (reserved).

class uwsgi.conf.options.routing_modifiers.ModifierSnmp (submod=None)
    Identify a SNMP request/response (mainly via UDP).

class uwsgi.conf.options.routing_modifiers.ModifierRaw (submod=None)
    Corresponds to the HTTP string and signals that this is a raw HTTP response.

class uwsgi.conf.options.routing_modifiers.ModifierMulticastAnnounce (submod=None)
    Announce message.

class uwsgi.conf.options.routing_modifiers.ModifierMulticast (submod=None)
    Array of chars; a custom multicast message managed by uwsgi.

class uwsgi.conf.options.routing_modifiers.ModifierClusterNode (submod=None)
    Add/remove/enable/disable node from a cluster.

    Add action requires a dict of at least 3 keys:

        • hostname
        • address
        • workers

class uwsgi.conf.options.routing_modifiers.ModifierRemoteLogging (submod=None)
    Remote logging (clustering/multicast/unicast).

class uwsgi.conf.options.routing_modifiers.ModifierReload (submod=None)
    Graceful reload request.

class uwsgi.conf.options.routing_modifiers.ModifierReloadBrutal (submod=None)
    Brutal reload request.

class uwsgi.conf.options.routing_modifiers.ModifierConfigFromNode (submod=None)
    Request configuration data from a uwsgi node (even via multicast).

class uwsgi.conf.options.routing_modifiers.ModifierPing (submod=None)
    PING-PONG. Useful for cluster health check.

    SUB_PING = 0
        Request.

    SUB_PONG = 1
        Response.

class uwsgi.conf.options.routing_modifiers.ModifierEcho (submod=None)
    ECHO service.

class uwsgi.conf.options.routing_modifiers.ModifierLegionMsg (submod=None)
    Legion msg (UDP, the body is encrypted).

class uwsgi.conf.options.routing_modifiers.ModifierSignal (submod=None)
    uwsgi_signal framework (payload is optional).

```

---

**Note:** modifier2 is the signal num.

---

```

class uwsgi.conf.options.routing_modifiers.ModifierCache (submod=None)
    Cache operations.

    SUB_GET = 0
        Simple cache get for values not bigger than 64k.

```

**SUB\_SET = 1**  
Simple cache set for values not bigger than 64k.

**SUB\_DELETE = 2**  
Simple cache del.

**SUB\_DICT\_BASED = 3**  
Simple dict based get command.

**SUB\_STREAM = 5**  
Get and stream.

**SUB\_DUMP = 6**  
Dump the whole cache.

**SUB\_MAGIC = 17**  
Magic interface for plugins remote access.

**class** uwsgi.conf.options.routing\_modifiers.**ModifierCorerouterSignal** (*submod=None*)  
Special modifier for signaling corerouters about special conditions.

**class** uwsgi.conf.options.routing\_modifiers.**ModifierRpc** (*submod=None*)  
RPC. The packet is an uwsgi array where

- the first item - the name of the function
- the following - the args

**SUB\_DEFAULT = 0**  
Return uwsgi header + rpc response.

**SUB\_RAW = 1**  
Return raw rpc response, uwsgi header included, if available.

**SUB\_USE\_PATH\_INFO = 2**  
Split PATH\_INFO to get func name and args and return as HTTP response with content\_type as application/binary or Accept request header (if different from \*).

**SUB\_XMLRPC = 3**  
Set xmlrpc wrapper (requires libxml2).

**SUB\_JSONRPC = 4**  
Set jsonrpc wrapper (requires libjansson).

**SUB\_DICT = 5**  
Used in uwsgi response to signal the response is a uwsgi dictionary followed by the body (the dictionary must contains a CONTENT\_LENGTH key).

**class** uwsgi.conf.options.routing\_modifiers.**ModifierPersistentClose** (*submod=None*)  
Close mark for persistent connections.

**class** uwsgi.conf.options.routing\_modifiers.**ModifierSubscription** (*submod=None*)  
Subscription packet. See subscriptions.

**class** uwsgi.conf.options.routing\_modifiers.**ModifierExample** (*submod=None*)  
Modifier used in dummy example plugin.

**class** uwsgi.conf.options.routing\_modifiers.**ModifierResponse** (*submod=None*)  
Generic response. Request dependent.

Example: a spooler response set 0 for a failed spool or 1 for a successful one.

## Subjects

**class** uwsgiconf.options.routing\_subjects.**SubjectCustom** (*subject*, *negate=False*)

Represents a routing subject that supports various check.

### Parameters

- **subject** (*Var/str/unicode*) – Handwritten subject or a Var heir representing it.
- **negate** (*bool*) – Use to negate subject for rule. .. note:: You can also use tilde (~) instead of this argument for negation.

**exists** ()

Check if the subject exists in the filesystem.

**isfile** ()

Check if the subject is a file.

**isdir** ()

Check if the subject is a directory.

**islink** ()

Check if the subject is a link.

**isexec** ()

Check if the subject is an executable file.

**islord** ()

Check if the subject is a Legion Lord.

**contains\_ipv4** ()

Check if the subject is ip v4.

**contains\_ipv6** ()

Check if the subject is ip v6.

**eq** (*val*)

Check if the subject is equal to the specified pattern.

**ge** (*val*)

Check if the subject is greater than or equal to the specified pattern.

**le** (*val*)

Check if the subject is less than or equal to the specified pattern.

**gt** (*val*)

Check if the subject is greater than the specified pattern.

**lt** (*val*)

Check if the subject is less than the specified pattern.

**startswith** (*val*)

Check if the subject starts with the specified pattern.

**endswith** (*val*)

Check if the subject ends with the specified pattern.

**matches** (*regexp*)

Check if the subject matches the specified regexp.

**isempty** ()

Check if the subject is empty.

**contains** (*val*)

Check if the subject contains the specified pattern.

**class** uwsgi.conf.options.routing\_subjects.**SubjectPathInfo** (*regexp*)

Default subject, maps to PATH\_INFO.

**class** uwsgi.conf.options.routing\_subjects.**SubjectRequestUri** (*regexp*)

Checks REQUEST\_URI for a value.

**class** uwsgi.conf.options.routing\_subjects.**SubjectQueryString** (*regexp*)

Checks QUERY\_STRING for a value.

**class** uwsgi.conf.options.routing\_subjects.**SubjectRemoteAddr** (*regexp*)

Checks REMOTE\_ADDR for a value.

**class** uwsgi.conf.options.routing\_subjects.**SubjectRemoteUser** (*regexp*)

Checks REMOTE\_USER for a value.

**class** uwsgi.conf.options.routing\_subjects.**SubjectHttpHost** (*regexp*)

Checks HTTP\_HOST for a value.

**class** uwsgi.conf.options.routing\_subjects.**SubjectHttpReferer** (*regexp*)

Checks HTTP\_REFERER for a value.

**class** uwsgi.conf.options.routing\_subjects.**SubjectHttpUserAgent** (*regexp*)

Checks HTTP\_USER\_AGENT for a value.

**class** uwsgi.conf.options.routing\_subjects.**SubjectStatus** (*regexp*)

Checks HTTP response status code.

**Warning:** Not available in the request chain.

## Variables and Functions

**class** uwsgi.conf.options.routing\_vars.**Var** (*name*)

**class** uwsgi.conf.options.routing\_vars.**Func** (*name*)

**class** uwsgi.conf.options.routing\_vars.**VarGeoip** (*name*)

Returns Geoip data.

<http://uwsgi.readthedocs.io/en/latest/GeoIP.html>

**vars\_country** = ['country\_code', 'country\_code3', 'country\_name']

Keys available for country database.

**vars\_city** = ['continent', 'country\_code', 'country\_code3', 'country\_name', 'region', '']

Keys available for city database.

**class** uwsgi.conf.options.routing\_vars.**VarRequest** (*name*)

Returns request variable. Examples: PATH\_INFO, SCRIPT\_NAME, REQUEST\_METHOD.

**class** uwsgi.conf.options.routing\_vars.**VarMetric** (*name*)

Returns metric (see monitoring) variable.

**class** uwsgi.conf.options.routing\_vars.**VarCookie** (*name*)

Returns cookie variable

**class** uwsgi.conf.options.routing\_vars.**VarQuery** (*name*)

Returns query string variable.



**class** uwsgi.conf.options.routing\_vars.**VarUwsgi** (*name*)  
Returns internal uWSGI information.

**Supported variables:**

- wid
- pid
- uuid
- status

**class** uwsgi.conf.options.routing\_vars.**VarTime** (*name*)  
Returns time/date in various forms.

**Supported variables:**

- unix

**class** uwsgi.conf.options.routing\_vars.**VarHttpptime** (*name*)  
Returns http date adding the numeric argument (if specified) to the current time (use empty arg for current server time).

**class** uwsgi.conf.options.routing\_vars.**FuncMime** (*name*)  
Returns mime type of a variable.

**class** uwsgi.conf.options.routing\_vars.**FuncMath** (*name*)  
Perform a math operation. Example: CONTENT\_LENGTH+1  
Supported operations: + - \* /

**Warning:** Requires matheval support.

**class** uwsgi.conf.options.routing\_vars.**FuncBase64** (*name*)  
Encodes the specified var in base64

**class** uwsgi.conf.options.routing\_vars.**FuncHex** (*name*)  
Encodes the specified var in hex.

**class** uwsgi.conf.options.routing\_vars.**FuncUpper** (*name*)  
Uppercase the specified var.

**class** uwsgi.conf.options.routing\_vars.**FuncLower** (*name*)  
Lowercase the specified var.

**class** uwsgi.conf.options.routing.**RouteRule** (*action*, *subject=None*, *stage=""*)  
Represents a routing rule.

**Parameters**

- **action** (*RouteAction*) – Action (or transformation) to perform. See `.actions` and `.transforms`.
- **subject** (*SubjectCustom/SubjectBuiltin/str*) – Subject to verify before action is performed. See `.subjects`.
  - String values are automatically transformed into `subjects.path_info`.
  - If `None` action is performed always w/o subject check.
- **stage** (*str/unicode*) – Stage on which the action needs to be performed. See `.stages`.

**class vars**

Routing variables.

**cookie**

alias of *uwsgiconf.options.routing\_vars.VarCookie*

**geoip**

alias of *uwsgiconf.options.routing\_vars.VarGeoip*

**httptime**

alias of *uwsgiconf.options.routing\_vars.VarHttpptime*

**metric**

alias of *uwsgiconf.options.routing\_vars.VarMetric*

**query**

alias of *uwsgiconf.options.routing\_vars.VarQuery*

**request**

alias of *uwsgiconf.options.routing\_vars.VarRequest*

**time**

alias of *uwsgiconf.options.routing\_vars.VarTime*

**uwsgi**

alias of *uwsgiconf.options.routing\_vars.VarUwsgi*

**class var\_functions**

Functions that can be applied to variables.

**base64**

alias of *uwsgiconf.options.routing\_vars.FuncBase64*

**hex**

alias of *uwsgiconf.options.routing\_vars.FuncHex*

**lower**

alias of *uwsgiconf.options.routing\_vars.FuncLower*

**math**

alias of *uwsgiconf.options.routing\_vars.FuncMath*

**mime**

alias of *uwsgiconf.options.routing\_vars.FuncMime*

**upper**

alias of *uwsgiconf.options.routing\_vars.FuncUpper*

**class stages**

During the request cycle, various stages (aka chains) are processed.

Chains can be “recursive”. A recursive chain can be called multiple times in a request cycle.

**REQUEST = ''**

Applied before the request is passed to the plugin.

**ERROR = 'error'**

Applied as soon as an HTTP status code is generate. **Recursive chain.**

**RESPONSE = 'response'**

Applied after the last response header has been generated (just before sending the body).

**FINAL = 'final'**

Applied after the response has been sent to the client.

**class subjects**

Routing subjects. These can be request's variables or other entities.

---

**Note:** Non-custom subjects can be pre-optimized (during startup) and should be used for performance reasons.

---

**custom**

alias of `uwsgi.conf.options.routing_subjects.SubjectCustom`

**http\_host**

alias of `uwsgi.conf.options.routing_subjects.SubjectHttpHost`

**http\_referer**

alias of `uwsgi.conf.options.routing_subjects.SubjectHttpReferer`

**http\_user\_agent**

alias of `uwsgi.conf.options.routing_subjects.SubjectHttpUserAgent`

**path\_info**

alias of `uwsgi.conf.options.routing_subjects.SubjectPathInfo`

**query\_string**

alias of `uwsgi.conf.options.routing_subjects.SubjectQueryString`

**remote\_addr**

alias of `uwsgi.conf.options.routing_subjects.SubjectRemoteAddr`

**remote\_user**

alias of `uwsgi.conf.options.routing_subjects.SubjectRemoteUser`

**request\_uri**

alias of `uwsgi.conf.options.routing_subjects.SubjectRequestUri`

**status**

alias of `uwsgi.conf.options.routing_subjects.SubjectStatus`

**class transforms**

A transformation is like a filter applied to the response generated by your application.

Transformations can be chained (the output of a transformation will be the input of the following one) and can completely overwrite response headers.

- <http://uwsgi.readthedocs.io/en/latest/Transformations.html>

**chunked**

alias of `uwsgi.conf.options.routing_actions.ActionChunked`

**fix\_content\_len**

alias of `uwsgi.conf.options.routing_actions.ActionFixContentLen`

**flush**

alias of `uwsgi.conf.options.routing_actions.ActionFlush`

**gzip**

alias of `uwsgi.conf.options.routing_actions.ActionGzip`

**template**

alias of `uwsgi.conf.options.routing_actions.ActionTemplate`

**to\_file**

alias of `uwsgi.conf.options.routing_actions.ActionToFile`

**upper**

alias of *uwsgi.conf.options.routing\_actions.ActionUpper*

**class actions**

Actions available for routing rules.

Values returned by actions:

- NEXT - continue to the next rule
- CONTINUE - stop scanning the internal routing table and run the request
- BREAK - stop scanning the internal routing table and close the request
- GOTO x - go to rule x

**add\_var CGI**

alias of *uwsgi.conf.options.routing\_actions.ActionAddVarCgi*

**add\_var\_log**

alias of *uwsgi.conf.options.routing\_actions.ActionAddVarLog*

**alarm**

alias of *uwsgi.conf.options.routing\_actions.ActionAlarm*

**auth\_basic**

alias of *uwsgi.conf.options.routing\_actions.ActionAuthBasic*

**auth\_ldap**

alias of *uwsgi.conf.options.routing\_actions.AuthLdap*

**dir\_change**

alias of *uwsgi.conf.options.routing\_actions.ActionDirChange*

**do\_break**

alias of *uwsgi.conf.options.routing\_actions.ActionDoBreak*

**do\_continue**

alias of *uwsgi.conf.options.routing\_actions.ActionDoContinue*

**do\_goto**

alias of *uwsgi.conf.options.routing\_actions.ActionDoGoto*

**fix\_var\_path\_info**

alias of *uwsgi.conf.options.routing\_actions.ActionFixVarPathInfo*

**header\_add**

alias of *uwsgi.conf.options.routing\_actions.ActionHeaderAdd*

**header\_remove**

alias of *uwsgi.conf.options.routing\_actions.ActionHeaderRemove*

**headers\_off**

alias of *uwsgi.conf.options.routing\_actions.ActionHeadersOff*

**headers\_reset**

alias of *uwsgi.conf.options.routing\_actions.ActionHeadersReset*

**log**

alias of *uwsgi.conf.options.routing\_actions.ActionLog*

**offload\_off**

alias of *uwsgi.conf.options.routing\_actions.ActionOffloadOff*

**redirect**  
alias of `uwsgi.conf.options.routing_actions.ActionRedirect`

**rewrite**  
alias of `uwsgi.conf.options.routing_actions.ActionRewrite`

**route\_external**  
alias of `uwsgi.conf.options.routing_actions.ActionRouteExternal`

**route\_uwsgi**  
alias of `uwsgi.conf.options.routing_actions.ActionRouteUwsgi`

**send**  
alias of `uwsgi.conf.options.routing_actions.ActionSend`

**serve\_static**  
alias of `uwsgi.conf.options.routing_actions.ActionServeStatic`

**set\_harakiri**  
alias of `uwsgi.conf.options.routing_actions.ActionSetHarakiri`

**set\_script\_file**  
alias of `uwsgi.conf.options.routing_actions.ActionSetScriptFile`

**set\_uwsgi\_process\_name**  
alias of `uwsgi.conf.options.routing_actions.ActionSetUwsgiProcessName`

**set\_var\_document\_root**  
alias of `uwsgi.conf.options.routing_actions.ActionSetVarDocumentRoot`

**set\_var\_path\_info**  
alias of `uwsgi.conf.options.routing_actions.ActionSetVarPathInfo`

**set\_var\_remote\_addr**  
alias of `uwsgi.conf.options.routing_actions.ActionSetVarRemoteAddr`

**set\_var\_remote\_user**  
alias of `uwsgi.conf.options.routing_actions.ActionSetVarRemoteUser`

**set\_var\_request\_method**  
alias of `uwsgi.conf.options.routing_actions.ActionSetVarRequestMethod`

**set\_var\_request\_uri**  
alias of `uwsgi.conf.options.routing_actions.ActionSetVarRequestUri`

**set\_var\_script\_name**  
alias of `uwsgi.conf.options.routing_actions.ActionSetVarScriptName`

**set\_var\_uwsgi\_appid**  
alias of `uwsgi.conf.options.routing_actions.ActionSetVarUwsgiAppid`

**set\_var\_uwsgi\_home**  
alias of `uwsgi.conf.options.routing_actions.ActionSetVarUwsgiHome`

**set\_var\_uwsgi\_scheme**  
alias of `uwsgi.conf.options.routing_actions.ActionSetVarUwsgiScheme`

**signal**  
alias of `uwsgi.conf.options.routing_actions.ActionSignal`

**class** `uwsgi.conf.options.routing.Routing(*args, **kwargs)`  
Routing subsystem.

You can use the internal routing subsystem to dynamically alter the way requests are handled.

---

**Note:** Since 1.9

---

- <http://uwsgi.readthedocs.io/en/latest/InternalRouting.html>
- <http://uwsgi.readthedocs.io/en/latest/Transformations.html>

**route\_rule**

alias of *RouteRule*

**class routers**

Dedicated routers, which can be used with *register\_router()*.

**http**

alias of *uwsgi.conf.options.routing\_routers.RouterHttp*

**https**

alias of *uwsgi.conf.options.routing\_routers.RouterHttps*

**ssl**

alias of *uwsgi.conf.options.routing\_routers.RouterSsl*

**fast**

alias of *uwsgi.conf.options.routing\_routers.RouterFast*

**raw**

alias of *uwsgi.conf.options.routing\_routers.RouterRaw*

**forkpty**

alias of *uwsgi.conf.options.routing\_routers.RouterForkPty*

**tuntap**

alias of *uwsgi.conf.options.routing\_routers.RouterTunTap*

**class modifiers**

Routing modifiers.

- <http://uwsgi.readthedocs.io/en/latest/Protocol.html>

**cache**

alias of *uwsgi.conf.options.routing\_modifiers.ModifierCache*

**cgi**

alias of *uwsgi.conf.options.routing\_modifiers.ModifierCgi*

**cluster\_node**

alias of *uwsgi.conf.options.routing\_modifiers.ModifierClusterNode*

**config\_from\_node**

alias of *uwsgi.conf.options.routing\_modifiers.ModifierConfigFromNode*

**corerouter\_signal**

alias of *uwsgi.conf.options.routing\_modifiers.ModifierCorerouterSignal*

**echo**

alias of *uwsgi.conf.options.routing\_modifiers.ModifierEcho*

**eval**

alias of *uwsgi.conf.options.routing\_modifiers.ModifierEval*

**example**

alias of *uwsgi.conf.options.routing\_modifiers.ModifierExample*

**fastfunc**  
alias of `uwsgi.conf.options.routing_modifiers.ModifierFastfunc`

**glusterfs**  
alias of `uwsgi.conf.options.routing_modifiers.ModifierGlusterfs`

**gridfs**  
alias of `uwsgi.conf.options.routing_modifiers.ModifierGridfs`

**jvm**  
alias of `uwsgi.conf.options.routing_modifiers.ModifierJvm`

**legion\_msg**  
alias of `uwsgi.conf.options.routing_modifiers.ModifierLegionMsg`

**lua**  
alias of `uwsgi.conf.options.routing_modifiers.ModifierLua`

**manage**  
alias of `uwsgi.conf.options.routing_modifiers.ModifierManage`

**manage\_path\_info**  
alias of `uwsgi.conf.options.routing_modifiers.ModifierManagePathInfo`

**message**  
alias of `uwsgi.conf.options.routing_modifiers.ModifierMessage`

**message\_array**  
alias of `uwsgi.conf.options.routing_modifiers.ModifierMessageArray`

**message\_marshall**  
alias of `uwsgi.conf.options.routing_modifiers.ModifierMessageMarshal`

**mono**  
alias of `uwsgi.conf.options.routing_modifiers.ModifierMono`

**multicast**  
alias of `uwsgi.conf.options.routing_modifiers.ModifierMulticast`

**multicast\_announce**  
alias of `uwsgi.conf.options.routing_modifiers.ModifierMulticastAnnounce`

**persistent\_close**  
alias of `uwsgi.conf.options.routing_modifiers.ModifierPersistentClose`

**php**  
alias of `uwsgi.conf.options.routing_modifiers.ModifierPhp`

**ping**  
alias of `uwsgi.conf.options.routing_modifiers.ModifierPing`

**psgi**  
alias of `uwsgi.conf.options.routing_modifiers.ModifierPsgi`

**rack**  
alias of `uwsgi.conf.options.routing_modifiers.ModifierRack`

**rados**  
alias of `uwsgi.conf.options.routing_modifiers.ModifierRados`

**raw**  
alias of `uwsgi.conf.options.routing_modifiers.ModifierRaw`

**reload**

alias of `uwsgi.conf.options.routing_modifiers.ModifierReload`

**reload\_brutal**

alias of `uwsgi.conf.options.routing_modifiers.ModifierReloadBrutal`

**remote\_logging**

alias of `uwsgi.conf.options.routing_modifiers.ModifierRemoteLogging`

**response**

alias of `uwsgi.conf.options.routing_modifiers.ModifierResponse`

**rpc**

alias of `uwsgi.conf.options.routing_modifiers.ModifierRpc`

**signal**

alias of `uwsgi.conf.options.routing_modifiers.ModifierSignal`

**snmp**

alias of `uwsgi.conf.options.routing_modifiers.ModifierSnmp`

**spooler**

alias of `uwsgi.conf.options.routing_modifiers.ModifierSpooler`

**subscription**

alias of `uwsgi.conf.options.routing_modifiers.ModifierSubscription`

**symcall**

alias of `uwsgi.conf.options.routing_modifiers.ModifierSymcall`

**v8**

alias of `uwsgi.conf.options.routing_modifiers.ModifierV8`

**wsgi**

alias of `uwsgi.conf.options.routing_modifiers.ModifierWsgi`

**xslt**

alias of `uwsgi.conf.options.routing_modifiers.ModifierXslt`

**use\_router** (*router*, *force=None*)

**Parameters**

- **router** (*RouterBase*) – Dedicated router object. See `.routers`.
- **force** (*bool*) – All of the gateways (routers) has to be run under the master process, supplying this you can try to bypass this limit.

**register\_route** (*route\_rules*, *label=None*)

Registers a routing rule.

**Parameters**

- **route\_rules** (*RouteRule* / *list* [*RouteRule*]) –
- **label** (*str* / *unicode*) – Label to mark the given set of rules. This can be used in conjunction with `do_goto` rule action.
  - <http://uwsgi.readthedocs.io/en/latest/InternalRouting.html#goto>

**print\_routing\_rules** ()

Print out supported routing rules (actions, transforms, etc.).

**set\_error\_page** (*status*, *html\_fpath*)

Add an error page (html) for managed 403, 404, 500 response.



**Parameters**

- **status** (*int*) – HTTP status code.
- **html\_fpath** (*str/unicode*) – HTML page file path.

**set\_error\_pages** (*codes\_map=None, common\_prefix=None*)

Add an error pages for managed 403, 404, 500 responses.

Shortcut for `.set_error_page()`.

**Parameters**

- **codes\_map** (*dict*) – Status code mapped into an html filepath or just a filename if `common_prefix` is used.

If not set, filename containing status code is presumed: 400.html, 500.html, etc.

- **common\_prefix** (*str/unicode*) – Common path (prefix) for all files.

**set\_geoiip\_params** (*db\_country=None, db\_city=None*)

Sets GeoIP parameters.

- <http://uwsgi.readthedocs.io/en/latest/GeoIP.html>

**Parameters**

- **db\_country** (*str/unicode*) – Country database file path.
- **db\_city** (*str/unicode*) – City database file path. Example: `GeoLiteCity.dat`.

**header\_add** (*name, value*)

Automatically add HTTP headers to response.

**Parameters**

- **name** (*str/unicode*) –
- **value** (*str/unicode*) –

**header\_remove** (*value*)

Automatically remove specified HTTP header from the response.

**Parameters** **value** (*str/unicode*) –

**header\_collect** (*name, target\_var, pull=False*)

Store the specified response header in a request var (optionally removing it from the response).

**Parameters**

- **name** (*str/unicode*) –
- **target\_var** (*str/unicode*) –
- **pull** (*bool*) – Whether to remove header from response.

### 3.8.14 Spooler

**class** uwsgi.conf.options.spooler.**Spooler** (*\*args, \*\*kwargs*)

Spooler.

---

**Note:** Supported on: Perl, Python, Ruby.

---

---

**Note:** Be sure the `spooler` plugin is loaded in your instance, but generally it is built in by default.

---

The Spooler is a queue manager built into uWSGI that works like a printing/mail system. You can enqueue massive sending of emails, image processing, video encoding, etc. and let the spooler do the hard work in background while your users get their requests served by normal workers.

<http://uwsgi-docs.readthedocs.io/en/latest/Spooler.html>

**set\_basic\_params** (*touch\_reload=None, quiet=None, process\_count=None, max\_tasks=None, order\_tasks=None, harakiri=None, change\_dir=None, poll\_interval=None, signal\_as\_task=None, cheap=None, base\_dir=None*)

#### Parameters

- **touch\_reload** (*str|unicode|list*) – reload spoolers if the specified file is modified/touched
- **quiet** (*bool*) – Do not log spooler related messages.
- **process\_count** (*int*) – Set the number of processes for spoolers.
- **max\_tasks** (*int*) – Set the maximum number of tasks to run before recycling a spooler (to help alleviate memory leaks).
- **order\_tasks** (*int*) – Try to order the execution of spooler tasks (uses `scandir` instead of `readdir`).
- **harakiri** (*int*) – Set harakiri timeout for spooler tasks.
- **change\_dir** (*str|unicode*) – `chdir()` to specified directory before each spooler task.
- **poll\_interval** (*int*) – Spooler poll frequency in seconds. Default: 30.
- **signal\_as\_task** (*bool*) – Treat signal events as tasks in spooler. To be used with `spooler-max-tasks`. If enabled spooler will treat signal events as task. Run signal handler will also increase the spooler task count.
- **cheap** (*bool*) – Use spooler cheap mode.
- **base\_dir** (*str*) – Base directory to prepend to `work_dir` argument of `.add()`.

**add** (*work\_dir, external=False*)

Run a spooler on the specified directory.

#### Parameters

- **work\_dir** (*str|unicode|list[str|unicode]*) – Spooler working directory path or it's name if `base_dir` argument of `spooler.set_basic_params()` is set.

---

**Note:** Placeholders can be used to build paths, e.g.: `{project_runtime_dir}/spool/`  
See `Section.project_name` and `Section.runtime_dir`.

---

- **external** (*bool*) – map spoolers requests to a spooler directory managed by an external instance

### 3.8.15 Statics

**class** uwsgi.conf.options.statics.**Statics** (\*args, \*\*kwargs)  
Statics.

Unfortunately you cannot live without serving static files via some protocol (HTTP, SPDY or something else). Fortunately uWSGI has a wide series of options and micro-optimizations for serving static files.

---

**Note:** This subsystem automatically honours the `If-Modified-Since` HTTP request header.

---

- <http://uwsgi.readthedocs.io/en/latest/StaticFiles.html>

**DIR\_DOCUMENT\_ROOT** = 'docroot'

Used to check for static files in the requested DOCUMENT\_ROOT. Pass into `static_dir`.

**class expiration\_criteria**

Expiration criteria (subjects) to use with `.add_expiration_rule()`.

**FILENAME** = 'filename'

Allows setting the Expires header for the specified file name pattern.

**MIME\_TYPE** = 'type'

Allows setting the Expires header for the specified MIME type.

**PATH\_INFO** = 'path-info'

Allows setting the Expires header for the specified PATH\_INFO pattern.

**REQUEST\_URI** = 'uri'

Allows setting the Expires header for the specified REQUEST\_URI pattern.

**class transfer\_modes**

File transfer (serving) modes.

With this, uWSGI will only generate response headers and the web server will be delegated to transferring the physical file.

- <http://uwsgi.readthedocs.io/en/latest/StaticFiles.html#transfer-modes>

**SENDFILE** = 'x-sendfile'

Use X-Sendfile mode. Apache.

**ACCEL\_REDIRECT** = 'x-accel-redirect'

Use X-Accel-Redirect mode. Nginx.

**set\_basic\_params** (*static\_dir=None, index\_file=None, mime\_file=None, skip\_ext=None, transfer\_mode=None*)

#### Parameters

- **static\_dir** (*str|unicode|list[str|unicode]*) – Check for static files in the specified directory.

---

**Note:** Use `DIR_DOCUMENT_ROOT` constant to serve files under `DOCUMENT_ROOT`.

---

- **index\_file** (*str|unicode|list[str|unicode]*) – Search for specified file if a directory is requested.

Example: `index.html`

- **mime\_file** (*str/unicode/list[str/unicode]*) – Set mime types file path to extend uWSGI builtin list.

Default: `/etc/mime.types` or `/etc/apache2/mime.types`.

- **skip\_ext** (*str/unicode/list[str/unicode]*) – Skip specified extension from static file checks.

Example: add `.php` to not serve it as static.

- **transfer\_mode** (*str/unicode*) – Set static file serving (transfer) mode.

See `.transfer_modes`.

---

**Note:** Another option is to specify `count_offload` in `.workers.set_thread_params()`.

---

**register\_static\_map** (*mountpoint, target, retain\_resource\_path=False, safe\_target=False*)

Allows mapping mountpoint to a static directory (or file).

- <http://uwsgi.readthedocs.io/en/latest/StaticFiles.html#mode-3-using-static-file-mount-points>

#### Parameters

- **mountpoint** (*str/unicode*) –
- **target** (*str/unicode*) –
- **retain\_resource\_path** (*bool*) – Append the requested resource to the doc-root.

Example: if `/images` maps to `/var/www/img` requested `/images/logo.png` will be served from:

– True: `/var/www/img/images/logo.png`

– False: `/var/www/img/logo.png`

- **safe\_target** (*bool*) – Skip security checks if the file is under the specified path. Whether to consider resolved (real) target a safe one to serve from.
  - <http://uwsgi.readthedocs.io/en/latest/StaticFiles.html#security>

**add\_expiration\_rule** (*criterion, value, timeout, use\_mod\_time=False*)

Adds statics expiration rule based on a criterion.

#### Parameters

- **criterion** (*str/unicode*) – Criterion (subject) to base expiration on. See `.expiration_criteria`.
- **value** (*str/unicode/list[str/unicode]*) – Value to test criteria upon.

---

**Note:** Usually a regular expression.

---

- **timeout** (*int*) – Number of seconds to expire after.
- **use\_mod\_time** (*bool*) – Base on file modification time instead of the current time.

**set\_paths\_caching\_params** (*timeout=None, cache\_name=None*)

Use the uWSGI caching subsystem to store mappings from URI to filesystem paths.

- <http://uwsgi.readthedocs.io/en/latest/StaticFiles.html# caching-paths-mappings-resolutions>

#### Parameters

- **timeout** (*int*) – Amount of seconds to put resolved paths in the uWSGI cache.
- **cache\_name** (*str/unicode*) – Cache name to use for static paths.

## 3.8.16 Subscriptions

### Balancing Algorithms

**class** uwsgi.conf.options.subscriptions\_algos.**BalancingAlgorithm** (*\*args*)

**class** uwsgi.conf.options.subscriptions\_algos.**BalancingAlgorithmWithBackup** (*backup\_level=None*)

**class** uwsgi.conf.options.subscriptions\_algos.**WeightedRoundRobin** (*backup\_level=None*)  
Weighted round robin algorithm with backup support. The default algorithm.

**class** uwsgi.conf.options.subscriptions\_algos.**LeastReferenceCount** (*backup\_level=None*)  
Least reference count algorithm with backup support.

**class** uwsgi.conf.options.subscriptions\_algos.**WeightedLeastReferenceCount** (*backup\_level=None*)  
Weighted least reference count algorithm with backup support.

**class** uwsgi.conf.options.subscriptions\_algos.**IpHash** (*backup\_level=None*)  
IP hash algorithm with backup support.

**class** uwsgi.conf.options.subscriptions.**Subscriptions** (*\*args, \*\*kwargs*)

This allows some uWSGI instances to announce their presence to subscriptions managing server, which in its turn can address those nodes (e.g. delegate request processing to them) and automatically remove dead nodes from the pool.

Some routers provide subscription server functionality. See *.routing.routers*.

---

**Note:** Subscription system in many ways relies on Master Process.

---

**Warning:** The subscription system is meant for “trusted” networks. All of the nodes in your network can potentially make a total mess with it.

- <http://uwsgi.readthedocs.io/en/latest/SubscriptionServer.html>

**class algorithms**

Balancing algorithms available to use with subscribe.

**ip\_hash**

alias of *uwsgi.conf.options.subscriptions\_algos.IpHash*

**least\_reference\_count**

alias of *uwsgi.conf.options.subscriptions\_algos.LeastReferenceCount*

**weighted\_least\_reference\_count**  
alias of `uwsgi.conf.options.subscriptions_algos.WeightedLeastReferenceCount`

**weighted\_round\_robin**  
alias of `uwsgi.conf.options.subscriptions_algos.WeightedRoundRobin`

**set\_server\_params** (*client\_notify\_address=None, mountpoints\_depth=None, re-quire\_vassal=None, tolerance=None, tolerance\_inactive=None, key\_dot\_split=None*)

Sets subscription server related params.

#### Parameters

- **client\_notify\_address** (*str|unicode*) – Set the notification socket for subscriptions. When you subscribe to a server, you can ask it to “acknowledge” the acceptance of your request. pointing address (Unix socket or UDP), on which your instance will bind and the subscription server will send acknowledgements to.
- **mountpoints\_depth** (*int*) – Enable support of mountpoints of certain depth for subscription system.
  - <http://uwsgi-docs.readthedocs.io/en/latest/SubscriptionServer.html#mountpoints-uwsgi-2-1>
- **require\_vassal** (*bool*) – Require a vassal field (see `subscribe`) from each subscription.
- **tolerance** (*int*) – Subscription reclaim tolerance (seconds).
- **tolerance\_inactive** (*int*) – Subscription inactivity tolerance (seconds).
- **key\_dot\_split** (*bool*) – Try to fallback to the next part in (dot based) subscription key. Used, for example, in SNI.

**set\_server\_verification\_params** (*digest\_algo=None, dir\_cert=None, tolerance=None, no\_check\_uid=None, dir\_credentials=None, pass\_unix\_credentials=None*)

Sets peer verification params for subscription server.

These are for secured subscriptions.

#### Parameters

- **digest\_algo** (*str|unicode*) – Digest algorithm. Example: SHA1

---

**Note:** Also requires `dir_cert` to be set.

---
- **dir\_cert** (*str|unicode*) – Certificate directory.

---

**Note:** Also requires `digest_algo` to be set.

---
- **tolerance** (*int*) – Maximum tolerance (in seconds) of clock skew for secured subscription system. Default: 24h.
- **no\_check\_uid** (*str|unicode|int|list[str|unicode|int]*) – Skip signature check for the specified uids when using unix sockets credentials.
- **dir\_credentials** (*str|unicode|list[str|unicode]*) – Directories to search for subscriptions key credentials.

- **pass\_unix\_credentials** (*bool*) – Enable management of SCM\_CREDENTIALS in subscriptions UNIX sockets.

**set\_client\_params** (*start\_unsubscribed=None, clear\_on\_exit=None, unsubscribe\_on\_reload=None, announce\_interval=None*)  
Sets subscribers related params.

#### Parameters

- **start\_unsubscribed** (*bool*) – Configure subscriptions but do not send them. .. note:: Useful with master FIFO.
- **clear\_on\_exit** (*bool*) – Force clear instead of unsubscribe during shutdown.
- **unsubscribe\_on\_reload** (*bool*) – Force unsubscribe request even during graceful reload.
- **announce\_interval** (*int*) – Send subscription announce at the specified interval. Default: 10 master cycles.

**subscribe** (*server=None, key=None, address=None, address\_vassal=None, balancing\_weight=None, balancing\_algo=None, modifier=None, signing=None, check\_file=None, protocol=None, sni\_cert=None, sni\_key=None, sni\_client\_ca=None*)  
Registers a subscription intent.

#### Parameters

- **server** (*str/unicode*) – Subscription server address (UDP or UNIX socket).

#### Examples:

– 127.0.0.1:7171

- **key** (*str/unicode*) – Key to subscribe. Generally the domain name (+ optional ‘/< mountpoint>’). Examples:
  - mydomain.it/foo
  - mydomain.it/foo/bar (requires mountpoints\_depth=2)
  - mydomain.it
  - ubuntu64.local:9090
- **address** (*str/unicode*) – Address to subscribe (the value for the key) or zero-based internal socket number (integer).
- **address** – Vassal node address.
- **balancing\_weight** (*int*) – Load balancing value. Default: 1.
- **balancing\_algo** – Load balancing algorithm to use. See `balancing_algorithms` .. note:: Since 2.1
- **modifier** (*Modifier*) – Routing modifier object. See `.routing.modifiers`
- **signing** (*list/tuple*) – Signing basics, expects two elements list/tuple: (signing\_algorithm, key).

#### Examples:

– SHA1:idlessh001

- **check\_file** (*str/unicode*) – If this file exists the subscription packet is sent, otherwise it is skipped.

- **protocol** (*str/unicode*) – the protocol to use, by default it is uwsgi. See `.networking.socket_types`.

---

**Note:** Since 2.1

---

- **sni\_cert** (*str/unicode*) – Certificate file to use for SNI proxy management. \* <http://uwsgi.readthedocs.io/en/latest/SNI.html#subscription-system-and-sni>
- **sni\_key** (*str/unicode*) – sni\_key Key file to use for SNI proxy management. \* <http://uwsgi.readthedocs.io/en/latest/SNI.html#subscription-system-and-sni>
- **sni\_client\_ca** (*str/unicode*) – Ca file to use for SNI proxy management. \* <http://uwsgi.readthedocs.io/en/latest/SNI.html#subscription-system-and-sni>

## 3.8.17 Workers

### Cheapening

**class** uwsgi.conf.options.workers\_cheapening.**Algo** (\*args, \*\*kwargs)

**class** uwsgi.conf.options.workers\_cheapening.**AlgoSpare** (\*args, \*\*kwargs)  
The default algorithm.

If all workers are busy for a certain amount of time seconds then uWSGI will spawn new workers. When the load is gone it will begin stopping processes one at a time.

- <http://uwsgi.readthedocs.io/en/latest/Cheaper.html#spare-cheaper-algorithm>

**set\_basic\_params** (*check\_interval\_overload=None*)

**Parameters** **check\_interval\_overload** (*int*) – Interval (sec) to wait after all workers are busy before new worker spawn.

**class** uwsgi.conf.options.workers\_cheapening.**AlgoSpare2** (\*args, \*\*kwargs)

This algorithm is similar to `spare`, but suitable for large scale by increase workers faster (before overload) and decrease them slower.

- <http://uwsgi.readthedocs.io/en/latest/Cheaper.html#spare2-cheaper-algorithm>

**set\_basic\_params** (*check\_interval\_idle=None*)

**Parameters** **check\_interval\_idle** (*int*) – Decrease workers after specified idle. Default: 10.

**class** uwsgi.conf.options.workers\_cheapening.**AlgoQueue** (\*args, \*\*kwargs)

If the socket's listen queue has more than `cheaper_overload` requests waiting to be processed, uWSGI will spawn new workers.

If the backlog is lower it will begin killing processes one at a time.

- <http://uwsgi.readthedocs.io/en/latest/Cheaper.html#backlog-cheaper-algorithm>

**set\_basic\_params** (*check\_num\_overload=None*)

**Parameters** **check\_num\_overload** (*int*) – Number of backlog items in queue.

**class** uwsgi.conf.options.workers\_cheapening.**AlgoBusyness** (\*args, \*\*kwargs)

Algorithm adds or removes workers based on average utilization for a given time period. It's goal is to keep more workers than the minimum needed available at any given time, so the app will always have capacity for new requests.



- <http://uwsgi.readthedocs.io/en/latest/Cheaper.html#busyness-cheaper-algorithm>

---

**Note:** Requires `cheaper_busyness` plugin.

---

**set\_basic\_params** (*check\_interval\_busy=None, busy\_max=None, busy\_min=None, idle\_cycles\_max=None, idle\_cycles\_penalty=None, verbose=None*)

#### Parameters

- **check\_interval\_busy** (*int*) – Interval (sec) to check worker busyness.
- **busy\_max** (*int*) – Maximum busyness (percents). Every time the calculated busyness is higher than this value, uWSGI will spawn new workers. Default: 50.
- **busy\_min** (*int*) – Minimum busyness (percents). If busyness is below this value, the app is considered in an “idle cycle” and uWSGI will start counting them. Once we reach needed number of idle cycles uWSGI will kill one worker. Default: 25.
- **idle\_cycles\_max** (*int*) – This option tells uWSGI how many idle cycles are allowed before stopping a worker.
- **idle\_cycles\_penalty** (*int*) – Number of idle cycles to add to `idle_cycles_max` in case worker spawned too early. Default is 1.
- **verbose** (*bool*) – Enables debug logs for this algo.

**set\_emergency\_params** (*workers\_step=None, idle\_cycles\_max=None, queue\_size=None, queue\_nonzero\_delay=None*)

Sets busyness algorithm emergency workers related params.

Emergency workers could be spawned depending upon uWSGI backlog state.

---

**Note:** These options are Linux only.

---

#### Parameters

- **workers\_step** (*int*) – Number of emergency workers to spawn. Default: 1.
- **idle\_cycles\_max** (*int*) – Idle cycles to reach before stopping an emergency worker. Default: 3.
- **queue\_size** (*int*) – Listen queue (backlog) max size to spawn an emergency worker. Default: 33.
- **queue\_nonzero\_delay** (*int*) – If the request listen queue is > 0 for more than given amount of seconds new emergency workers will be spawned. Default: 60.

**class** `uwsgi.conf.options.workers_cheapening.AlgoManual` (*\*args, \*\*kwargs*)

Algorithm allows to adjust number of workers using Master FIFO commands.

- <http://uwsgi.readthedocs.io/en/latest/MasterFIFO.html#available-commands>

**class** `uwsgi.conf.options.workers_cheapening.Cheapening` (*\*args, \*\*kwargs*)

uWSGI provides the ability to dynamically scale the number of running workers (adaptive process spawning) via pluggable algorithms.

---

**Note:** This uses master process.

---

**class algorithms**

Algorithms available to use with `cheaper_algorithm`.

**busyness**

alias of `AlgoBusyness`

**manual**

alias of `AlgoManual`

**queue**

alias of `AlgoQueue`

**spare**

alias of `AlgoSpare`

**spare2**

alias of `AlgoSpare2`

**set\_basic\_params** (*spawn\_on\_request=None, cheaper\_algo=None, workers\_min=None, workers\_startup=None, workers\_step=None*)

**Parameters**

- **spawn\_on\_request** (*bool*) – Spawn workers only after the first request.
- **cheaper\_algo** (*Algo*) – The algorithm object to be used used for adaptive process spawning. Default: `spare`. See `.algorithms`.
- **workers\_min** (*int*) – Minimal workers count. Enables cheaper mode (adaptive process spawning).

---

**Note:** Must be lower than max workers count.

---

- **workers\_startup** (*int*) – The number of workers to be started when starting the application. After the app is started the algorithm can stop or start workers if needed.
- **workers\_step** (*int*) – Number of additional processes to spawn at a time if they are needed,

**set\_memory\_limits** (*rss\_soft=None, rss\_hard=None*)

Sets worker memory limits for cheapening.

**Parameters**

- **rss\_soft** (*int*) – Don't spawn new workers if total resident memory usage of all workers is higher than this limit in bytes.

**Warning:** This option expects memory reporting enabled: `.logging.set_basic_params(memory_report=1)`

- **rss\_hard** (*int*) – Try to stop workers if total workers resident memory usage is higher than this limit in bytes.

**print\_algorithms** ()

Print out enabled cheaper algorithms.

**class uwsgiconf.options.workers.MuleFarm** (*name, mule\_numbers*)

Represents a mule farm.

**Parameters**

- **name** (*str/unicode*) – Farm alias.
- **mule\_numbers** (*int/list[int]*) – Total mules on farm count, or a list of mule numbers.

**class** uwsgi.conf.options.workers.**Workers** (\*args, \*\*kwargs)  
Workers aka [working] processes.

**mule\_farm**  
alias of *MuleFarm*

**set\_basic\_params** (*count=None, touch\_reload=None, touch\_chain\_reload=None, zombie\_reaper=None, limit\_addr\_space=None, limit\_count=None, cpu\_affinity=None*)

### Parameters

- **count** (*int*) – Spawn the specified number of workers (processes). Set the number of workers for preforking mode. This is the base for easy and safe concurrency in your app. More workers you add, more concurrent requests you can manage.  
  
Each worker corresponds to a system process, so it consumes memory, choose carefully the right number. You can easily drop your system to its knees by setting a too high value.  
  
Setting `workers` to a ridiculously high number will **not** magically make your application web scale - quite the contrary.
- **touch\_reload** (*str/list*) – Trigger reload of (and only) workers if the specified file is modified/touched.
- **touch\_chain\_reload** (*str/list*) – Trigger chain workers reload on file touch. When in `lazy/lazy_apps` mode, you can simply destroy a worker to force it to reload the application code. A new reloading system named “chain reload”, allows you to reload one worker at time (opposed to the standard way where all of the workers are destroyed in bulk)  
  
– <http://uwsgi-docs.readthedocs.io/en/latest/articles/TheArtOfGracefulReloading.html#chain-reloading-lazy-apps>
- **zombie\_reaper** (*bool*) – Call `waitpid(-1, ...)` after each request to get rid of zombies. Enables reaper mode. After each request the server will call `waitpid(-1)` to get rid of zombie processes. If you spawn subprocesses in your app and you happen to end up with zombie processes all over the place you can enable this option. (It really would be better if you could fix your application’s process spawning usage though.)
- **limit\_addr\_space** (*int*) – Limit process address space (`vsz`) (in megabytes) Limits the address space usage of each uWSGI (worker) process using `POSIX/UNIX setrlimit()`. For example, `limit-as 256` will disallow uWSGI processes to grow over 256MB of address space. Address space is the virtual memory a process has access to. It does *not* correspond to physical memory. Read and understand this page before enabling this option: [http://en.wikipedia.org/wiki/Virtual\\_memory](http://en.wikipedia.org/wiki/Virtual_memory)
- **limit\_count** (*int*) – Limit the number of spawnable processes.
- **cpu\_affinity** (*int*) – number of cores for each worker (Linux only) Set the number of cores (CPUs) to allocate to each worker process.  
  
– **4 workers, 4 CPUs, affinity is 1**, each worker is allocated one CPU.  
  
– **4 workers, 2 CPUs, affinity is 1**, workers get one CPU each (0; 1; 0; 1).

- **4 workers, 4 CPUs, affinity is 2**, workers get two CPUs each in a round-robin fashion (0, 1; 2, 3; 0, 1; 2; 3).
- **8 workers, 4 CPUs, affinity is 3**, workers get three CPUs each in a round-robin fashion (0, 1, 2; 3, 0, 1; 2, 3, 0; 1, 2, 3; 0, 1, 2; 3, 0, 1; 2, 3, 0; 1, 2, 3).

**run\_command\_as\_worker** (*command*, *after\_post\_fork\_hook=False*)

Run the specified command as worker.

**Parameters**

- **command** (*str/unicode*) –
- **after\_post\_fork\_hook** (*bool*) – Whether to run it after *post\_fork* hook.

**set\_count\_auto** (*count=None*)

Sets workers count.

By default sets it to detected number of available cores

**Parameters count** (*int*) –

**set\_thread\_params** (*enable=None*, *count=None*, *count\_offload=None*, *stack\_size=None*, *no\_wait=None*)

Sets threads related params.

**Parameters**

- **enable** (*bool*) – Enable threads in the embedded languages. This will allow to spawn threads in your app.

**Warning:** Threads will simply *not work* if this option is not enabled. There will likely be no error, just no execution of your thread code.

- **count** (*int*) – Run each worker in prethreaded mode with the specified number of threads per worker.

**Warning:** Do not use with `gevent`.

---

**Note:** Enables threads automatically.

---

- **count\_offload** (*int*) – Set the number of threads (per-worker) to spawn for offloading. Default: 0.

These threads run such tasks in a non-blocking/evented way allowing for a huge amount of concurrency. Various components of the uWSGI stack are offload-friendly.

---

**Note:** Try to set it to the number of CPU cores to take advantage of SMP.

---

– <http://uwsgi-docs.readthedocs.io/en/latest/OffloadSubsystem.html>

- **stack\_size** (*int*) – Set threads stacksize.
- **no\_wait** (*bool*) – Do not wait for threads cancellation on quit/reload.

**set\_mules\_params** (*mules=None, touch\_reload=None, harakiri\_timeout=None, farms=None, reload\_mercy=None, msg\_buffer=None, msg\_buffer\_recv=None*)  
Sets mules related params.

<http://uwsgi.readthedocs.io/en/latest/Mules.html>

Mules are worker processes living in the uWSGI stack but not reachable via socket connections, that can be used as a generic subsystem to offload tasks.

#### Parameters

- **mules** (*int/list*) – Add the specified mules or number of mules.
- **touch\_reload** (*str/list*) – Reload mules if the specified file is modified/touched.
- **harakiri\_timeout** (*int*) – Set harakiri timeout for mule tasks.
- **farms** (*list [MuleFarm]*) – Mule farms list.

#### Examples:

- `cls_mule_farm('first', 2)`
- `cls_mule_farm('first', [4, 5])`
- **reload\_mercy** (*int*) – Set the maximum time (in seconds) a mule can take to reload/shutdown. Default: 60.
- **msg\_buffer** (*int*) – Set mule message buffer size (bytes) given for mule message queue.
- **msg\_buffer** – Set mule message recv buffer size (bytes).

**set\_reload\_params** (*min\_lifetime=None, max\_lifetime=None, max\_requests=None, max\_requests\_delta=None, max\_addr\_space=None, max\_rss=None, max\_uss=None, max\_pss=None, max\_addr\_space\_forced=None, max\_rss\_forced=None, watch\_interval\_forced=None, mercy=None*)  
Sets workers reload parameters.

#### Parameters

- **min\_lifetime** (*int*) – A worker cannot be destroyed/reloaded unless it has been alive for N seconds (default 60). This is an anti-fork-bomb measure. Since 1.9
- **max\_lifetime** (*int*) – Reload workers after this many seconds. Disabled by default. Since 1.9
- **max\_requests** (*int*) – Reload workers after the specified amount of managed requests (avoid memory leaks). When a worker reaches this number of requests it will get recycled (killed and restarted). You can use this option to “dumb fight” memory leaks.

Also take a look at the `reload-on-as` and `reload-on-rss` options as they are more useful for memory leaks.

**Warning:** The default min-worker-lifetime 60 seconds takes priority over *max-requests*.

Do not use with benchmarking as you'll get stalls such as *worker respawning too fast !!! i have to sleep a bit (2 seconds)*...

- **max\_requests\_delta** (*int*) – Add (*worker\_id* \* *delta*) to the *max\_requests* value of each worker.
- **max\_addr\_space** (*int*) – Reload a worker if its address space usage is higher than the specified value in megabytes.
- **max\_rss** (*int*) – Reload a worker if its physical unshared memory (resident set size) is higher than the specified value (in megabytes).
- **max\_uss** (*int*) – Reload a worker if Unique Set Size is higher than the specified value in megabytes.

---

**Note:** Linux only.

---

- **max\_pss** (*int*) – Reload a worker if Proportional Set Size is higher than the specified value in megabytes.

---

**Note:** Linux only.

---

- **max\_addr\_space\_forced** (*int*) – Force the master to reload a worker if its address space is higher than specified megabytes (in megabytes).
- **max\_rss\_forced** (*int*) – Force the master to reload a worker if its resident set size memory is higher than specified in megabytes.
- **watch\_interval\_forced** (*int*) – The memory collector [per-worker] thread memory watch interval (seconds) used for forced reloads. Default: 3.
- **mercy** (*int*) – Set the maximum time (in seconds) a worker can take before reload/shutdown. Default: 60.

**set\_reload\_on\_exception\_params** (*do\_reload=None*, *etype=None*, *value=None*, *erepr=None*)

Sets workers reload on exceptions parameters.

#### Parameters

- **do\_reload** (*bool*) – Reload a worker when an exception is raised.
- **etype** (*str*) – Reload a worker when a specific exception type is raised.
- **value** (*str*) – Reload a worker when a specific exception value is raised.
- **erepr** (*str*) – Reload a worker when a specific exception type+value (language-specific) is raised.

**set\_harakiri\_params** (*timeout=None*, *verbose=None*, *disable\_for\_arh=None*)

Sets workers harakiri parameters.

#### Parameters

- **timeout** (*int*) – Harakiri timeout in seconds. Every request that will take longer than the seconds specified in the harakiri timeout will be dropped and the corresponding worker is thereafter recycled.
- **verbose** (*bool*) – Harakiri verbose mode. When a request is killed by Harakiri you will get a message in the uWSGI log. Enabling this option will print additional info (for example, the current syscall will be reported on Linux platforms).
- **disable\_for\_arh** (*bool*) – Disallow Harakiri killings during after-request hook methods.

**set\_zerg\_server\_params** (*socket, clients\_socket\_pool=None*)  
Zerg mode. Zerg server params.

When your site load is variable, it would be nice to be able to add workers dynamically. Enabling Zerg mode you can allow zerg clients to attach to your already running server and help it in the work.

- <http://uwsgi-docs.readthedocs.io/en/latest/Zerg.html>

#### Parameters

- **socket** (*str/unicode*) – Unix socket to bind server to.

#### Examples:

- unix socket - `/var/run/mutalisk`
- Linux abstract namespace - `@nydus`

- **clients\_socket\_pool** (*str/unicode/list[str/unicode]*) – This enables Zerg Pools.

---

**Note:** Expects master process.

---

Accepts sockets that will be mapped to Zerg socket.

- <http://uwsgi-docs.readthedocs.io/en/latest/Zerg.html#zerg-pools>

**set\_zerg\_client\_params** (*server\_sockets, use\_fallback\_socket=None*)  
Zerg mode. Zergs params.

#### Parameters

- **server\_sockets** (*str/unicode/list[str/unicode]*) – Attaches zerg to a zerg server.
- **use\_fallback\_socket** (*bool*) – Fallback to normal sockets if the zerg server is not available

### 3.8.18 Python

**class** uwsgiconf.options.python.**Python** (*\*args, \*\*kwargs*)  
Python plugin options.

---

**Note:** By default the plugin does not initialize the GIL. This means your app-generated threads will not run. If you need threads, remember to enable them with `enable_threads`.

---

**set\_basic\_params** (*version=(1, ), python\_home=None, enable\_threads=None, search\_path=None, python\_binary=None, traceback\_path=None, plugin\_dir=None, os\_env\_reload=None, optimization\_level=None*)

#### Parameters

- **version** (*str/unicode/int*) – Python version plugin supports.

#### Example:

- 3 - version 3
- <empty> - version 2

– <default> - version deduced by uwsgi.conf

- **python\_home** (*str/unicode*) – Set python executable directory - PYTHON-HOME/virtualenv.
- **search\_path** (*str/unicode*) – Add directory (or an .egg or a glob) to the Python search path.

---

**Note:** This can be specified up to 64 times.

---

- **python\_binary** (*str/unicode*) – Set python program name.
- **tracebacker\_path** (*str/unicode*) – Enable the uWSGI Python tracebacker. <http://uwsgi-docs.readthedocs.io/en/latest/Tracebacker.html>
- **plugin\_dir** (*str/unicode*) – Directory to search for plugin.
- **enable\_threads** (*bool*) – Enable threads in the embedded languages. This will allow to spawn threads in your app.

**Warning:** Threads will simply *not work* if this option is not enabled. There will likely be no error, just no execution of your thread code.

- **os\_env\_reload** (*bool*) – Force `os.environ` reloading for every request. Used to allow setting of `UWSGI_SETENV` for Python applications.
- **optimization\_level** (*int*) – Python optimization level (see `-O` argument). .. warning:: This may be dangerous for some apps.

**set\_app\_args** (*\*args*)

Sets `sys.argv` for python apps.

**Examples:**

- `pyargv="one two three"` will set `sys.argv` to `('one', 'two', 'three')`.

**Parameters args** –

**set\_wsgi\_params** (*module=None, callable\_name=None, env\_strategy=None*)

Set wsgi related parameters.

**Parameters**

- **module** (*str/unicode*) –
  - load .wsgi file as the Python application
  - load a WSGI module as the application.

---

**Note:** The module (sans .py) must be importable, ie. be in `PYTHONPATH`.

---

**Examples:**

- `mypackage.my_wsgi_module` – read from `application` attr of `mypackage/my_wsgi_module.py`
- `mypackage.my_wsgi_module:my_app` – read from `my_app` attr of `mypackage/my_wsgi_module.py`



- **callable\_name** (*str/unicode*) – Set WSGI callable name. Default: application.
- **env\_strategy** (*str/unicode*) – Strategy for allocating/deallocating the WSGI env, can be:
  - **cheat** - preallocates the env dictionary on uWSGI startup and clears it after each request. Default behaviour for uWSGI <= 2.0.x
  - **holy** - creates and destroys the environ dictionary at each request. Default behaviour for uWSGI >= 2.1

**eval\_wsgi\_entrypoint** (*code*)

Evaluates Python code as WSGI entry point.

**Parameters** **code** (*str/unicode*) –

**set\_autoreload\_params** (*scan\_interval=None, ignore\_modules=None*)

Sets autoreload related parameters.

**Parameters**

- **scan\_interval** (*int*) – Seconds. Monitor Python modules' modification times to trigger reload.

**Warning:** Use only in development.

- **ignore\_modules** (*list/str/unicode*) – Ignore the specified module during auto-reload scan.

**register\_module\_alias** (*alias, module\_path, after\_init=False*)

Adds an alias for a module.

<http://uwsgi-docs.readthedocs.io/en/latest/PythonModuleAlias.html>

**Parameters**

- **alias** (*str/unicode*) –
- **module\_path** (*str/unicode*) –
- **after\_init** (*bool*) – add a python module alias after uwsgi module initialization

**import\_module** (*modules, shared=False, into\_spooler=False*)

Imports a python module.

**Parameters**

- **modules** (*list/str/unicode*) –
- **shared** (*bool*) – If shared import is done once in master process. Otherwise import a python module in all of the processes. This is done after fork but before request processing.
- **into\_spooler** (*bool*) – Import a python module in the spooler. <http://uwsgi-docs.readthedocs.io/en/latest/Spooler.html>

**run\_module** (*module*)

Runs a Python script in the uWSGI environment.

**Parameters** **module** (*str/unicode*) –

### 3.8.19 Config Formatters

Here belong tools for configuration formatting purposes.

`uwsgi.conf.formatters.format_print_text` (*text*, *color\_fg=None*, *color\_bg=None*)

Format given text using ANSI formatting escape sequences.

Could be useful gfor print command.

#### Parameters

- **text** (*str/unicode*) –
- **color\_fg** (*str/unicode*) – text (foreground) color
- **color\_bg** (*str/unicode*) – text (background) color

**Return type** `strlunicode`

**class** `uwsgi.conf.formatters.FormatterBase` (*sections*)

Base class for configuration formatters.

**iter\_options** ()

Iterates configuration sections groups options.

**class** `uwsgi.conf.formatters.IniFormatter` (*sections*)

Translates a configuration as INI file.

**class** `uwsgi.conf.formatters.ArgsFormatter` (*sections*)

Translates a configuration to command line arguments.

`uwsgi.conf.formatters.FORMATTERS = {'args': <class 'uwsgi.conf.formatters.ArgsFormatter'>,`

Available formatters by alias.

---

### Get involved into uwsgiconf

---

**Submit issues.** If you spotted something weird in application behavior or want to propose a feature you can do that at <https://github.com/idlesign/uwsgiconf/issues>

**Write code.** If you are eager to participate in application development, fork it at <https://github.com/idlesign/uwsgiconf>, write your code, whether it should be a bugfix or a feature implementation, and make a pull request right from the forked project page.

**Spread the word.** If you have some tips and tricks or any other words in mind that you think might be of interest for the others — publish it.



U

uwsgi.conf.config, 57  
uwsgi.conf.formatters, 158  
uwsgi.conf.options.alarm\_types, 64  
uwsgi.conf.options.alarms, 64  
uwsgi.conf.options.applications, 66  
uwsgi.conf.options.caching, 67  
uwsgi.conf.options.empire, 70  
uwsgi.conf.options.locks, 72  
uwsgi.conf.options.logging, 75  
uwsgi.conf.options.logging\_encoders, 75  
uwsgi.conf.options.logging\_loggers, 73  
uwsgi.conf.options.main\_process, 84  
uwsgi.conf.options.main\_process\_actions, 82  
uwsgi.conf.options.master\_process, 89  
uwsgi.conf.options.monitoring, 100  
uwsgi.conf.options.monitoring\_collectors, 99  
uwsgi.conf.options.monitoring\_metric\_types, 93  
uwsgi.conf.options.monitoring\_pushers, 97  
uwsgi.conf.options.networking, 109  
uwsgi.conf.options.networking\_sockets, 104  
uwsgi.conf.options.python, 155  
uwsgi.conf.options.queue, 112  
uwsgi.conf.options.routing, 133  
uwsgi.conf.options.routing\_actions, 113  
uwsgi.conf.options.routing\_modifiers, 127  
uwsgi.conf.options.routing\_routers, 118  
uwsgi.conf.options.routing\_subjects, 131  
uwsgi.conf.options.routing\_vars, 132  
uwsgi.conf.options.spooler, 141  
uwsgi.conf.options.statics, 143  
uwsgi.conf.options.subscriptions, 145  
uwsgi.conf.options.subscriptions\_algos, 145  
uwsgi.conf.options.workers, 150  
uwsgi.conf.options.workers\_cheapening, 148  
uwsgi.conf.presets.empire, 9  
uwsgi.conf.presets.nice, 10  
uwsgi.conf.runtime.alarms, 37  
uwsgi.conf.runtime.asynced, 37  
uwsgi.conf.runtime.caching, 38  
uwsgi.conf.runtime.control, 40  
uwsgi.conf.runtime.locking, 41  
uwsgi.conf.runtime.logging, 42  
uwsgi.conf.runtime.monitoring, 43  
uwsgi.conf.runtime.mules, 44  
uwsgi.conf.runtime.rpc, 49  
uwsgi.conf.runtime.scheduling, 50  
uwsgi.conf.runtime.signals, 53  
uwsgi.conf.runtime.spooler, 55  
uwsgi.conf.uwsgi\_stub, 14



## Symbols

`_Platform` (class in `uwsgi.conf.runtime.platform`), 46  
`_Request` (class in `uwsgi.conf.runtime.request`), 48

## A

`absolute` (`uwsgi.conf.options.monitoring.Monitoring.metric_types` attribute), 101  
`ACCEL_REDIRECT` (`uwsgi.conf.options.statics.Statics.transfer_modes` attribute), 143  
`accepting()` (in module `uwsgi.conf.uwsgi_stub`), 15  
`accumulator` (`uwsgi.conf.options.monitoring.Monitoring.collectors` attribute), 101  
`acquire()` (`uwsgi.conf.runtime.locking.Lock` method), 42  
`ActionAddVarCgi` (class in `uwsgi.conf.options.routing_actions`), 114  
`ActionAddVarLog` (class in `uwsgi.conf.options.routing_actions`), 114  
`ActionAlarm` (class in `uwsgi.conf.options.main_process_actions`), 83  
`ActionAlarm` (class in `uwsgi.conf.options.routing_actions`), 115  
`ActionAuthBasic` (class in `uwsgi.conf.options.routing_actions`), 115  
`ActionCall` (class in `uwsgi.conf.options.main_process_actions`), 82  
`ActionChunked` (class in `uwsgi.conf.options.routing_actions`), 113  
`ActionDirChange` (class in `uwsgi.conf.options.main_process_actions`), 83  
`ActionDirChange` (class in `uwsgi.conf.options.routing_actions`), 116  
`ActionDirCreate` (class in `uwsgi.conf.options.main_process_actions`), 83  
`ActionDoBreak` (class in `uwsgi.conf.options.routing_actions`), 113  
`ActionDoContinue` (class in `uwsgi.conf.options.routing_actions`), 113  
`ActionDoGoto` (class in `uwsgi.conf.options.routing_actions`), 114  
`ActionExecute` (class in `uwsgi.conf.options.main_process_actions`), 82  
`ActionExit` (class in `uwsgi.conf.options.main_process_actions`), 83  
`ActionFifoWrite` (class in `uwsgi.conf.options.main_process_actions`), 83  
`ActionFileCreate` (class in `uwsgi.conf.options.main_process_actions`), 83  
`ActionFileWrite` (class in `uwsgi.conf.options.main_process_actions`), 83  
`ActionFixContentLen` (class in `uwsgi.conf.options.routing_actions`), 113  
`ActionFixVarPathInfo` (class in `uwsgi.conf.options.routing_actions`), 117  
`ActionFlush` (class in `uwsgi.conf.options.routing_actions`), 113  
`ActionGzip` (class in `uwsgi.conf.options.routing_actions`), 113  
`ActionHeaderAdd` (class in `uwsgi.conf.options.routing_actions`), 114  
`ActionHeaderRemove` (class in `uwsgi.conf.options.routing_actions`), 114  
`ActionHeadersOff` (class in `uwsgi.conf.options.routing_actions`), 114  
`ActionHeadersReset` (class in `uwsgi.conf.options.routing_actions`), 114  
`ActionLog` (class in `uwsgi.conf.options.routing_actions`), 113  
`ActionMount` (class in `uwsgi.conf.options.main_process_actions`), 82  
`ActionOffloadOff` (class in `uwsgi.conf.options.routing_actions`), 113  
`ActionPrintout` (class in `uwsgi.conf.options.main_process_actions`), 83  
`ActionRedirect` (class in `uwsgi.conf.options.routing_actions`), 114  
`ActionRewrite` (class in `uwsgi-`

- `conf.options.routing_actions`), 115
- ActionRouteExternal (class in `uwsgi-conf.options.routing_actions`), 115
- ActionRouteUwsgi (class in `uwsgi-conf.options.routing_actions`), 115
- ActionSend (class in `uwsgi-conf.options.routing_actions`), 114
- ActionServeStatic (class in `uwsgi-conf.options.routing_actions`), 115
- ActionSetHarakiri (class in `uwsgi-conf.options.routing_actions`), 116
- ActionSetHostName (class in `uwsgi-conf.options.main_process_actions`), 83
- ActionSetScriptFile (class in `uwsgi-conf.options.routing_actions`), 117
- ActionSetUwsgiProcessName (class in `uwsgi-conf.options.routing_actions`), 117
- ActionSetVarDocumentRoot (class in `uwsgi-conf.options.routing_actions`), 117
- ActionSetVarPathInfo (class in `uwsgi-conf.options.routing_actions`), 117
- ActionSetVarRemoteAddr (class in `uwsgi-conf.options.routing_actions`), 117
- ActionSetVarRemoteUser (class in `uwsgi-conf.options.routing_actions`), 116
- ActionSetVarRequestMethod (class in `uwsgi-conf.options.routing_actions`), 117
- ActionSetVarRequestUri (class in `uwsgi-conf.options.routing_actions`), 117
- ActionSetVarScriptName (class in `uwsgi-conf.options.routing_actions`), 117
- ActionSetVarUwsgiAppid (class in `uwsgi-conf.options.routing_actions`), 116
- ActionSetVarUwsgiHome (class in `uwsgi-conf.options.routing_actions`), 116
- ActionSetVarUwsgiScheme (class in `uwsgi-conf.options.routing_actions`), 117
- ActionSignal (class in `uwsgi-conf.options.routing_actions`), 114
- ActionTemplate (class in `uwsgi-conf.options.routing_actions`), 113
- ActionToFile (class in `uwsgi-conf.options.routing_actions`), 113
- ActionUnlink (class in `uwsgi-conf.options.main_process_actions`), 84
- ActionUpper (class in `uwsgi-conf.options.routing_actions`), 113
- `add()` (`uwsgiconf.options.spooler.Spooler` method), 142
- `add_cache()` (`uwsgiconf.options.caching.Caching` method), 68
- `add_cron()` (in module `uwsgiconf.uwsgi_stub`), 15
- `add_cron_task()` (`uwsgi-conf.options.master_process.MasterProcess` method), 91
- `add_expiration_rule()` (`uwsgi-conf.options.statics.Static` method), 144
- `add_file()` (`uwsgiconf.options.caching.Caching` method), 68
- `add_file_monitor()` (in module `uwsgi-conf.uwsgi_stub`), 16
- `add_firewall_rule()` (`uwsgi-conf.options.routing_routers.RouterTunTap` method), 127
- `add_item()` (`uwsgiconf.options.caching.Caching` method), 67
- `add_logger()` (`uwsgiconf.options.logging.Logging` method), 79
- `add_logger_encoder()` (`uwsgi-conf.options.logging.Logging` method), 79
- `add_logger_route()` (`uwsgi-conf.options.logging.Logging` method), 79
- `add_ms_timer()` (in module `uwsgiconf.uwsgi_stub`), 16
- `add_rb_timer()` (in module `uwsgiconf.uwsgi_stub`), 16
- `add_timer()` (in module `uwsgiconf.uwsgi_stub`), 16
- `add_var()` (in module `uwsgiconf.uwsgi_stub`), 17
- `add_var()` (`uwsgiconf.runtime.request._Request` method), 48
- `add_var_cgi` (`uwsgi-conf.options.routing.RouteRule.actions` attribute), 136
- `add_var_log` (`uwsgi-conf.options.routing.RouteRule.actions` attribute), 136
- `adder` (`uwsgiconf.options.monitoring.Monitoring.collectors` attribute), 101
- `alarm` (`uwsgiconf.options.main_process.MainProcess.actions` attribute), 84
- `alarm` (`uwsgiconf.options.routing.RouteRule.actions` attribute), 136
- `alarm()` (in module `uwsgiconf.uwsgi_stub`), 17
- `alarm_on_fd_ready()` (`uwsgi-conf.options.alarms.Alarms` method), 65
- `alarm_on_log()` (`uwsgiconf.options.alarms.Alarms` method), 65
- `alarm_on_queue_full()` (`uwsgi-conf.options.alarms.Alarms` method), 65
- `alarm_on_segfault()` (`uwsgi-conf.options.alarms.Alarms` method), 65
- AlarmCommand (class in `uwsgi-conf.options.alarm_types`), 64
- AlarmCurl (class in `uwsgiconf.options.alarm_types`), 64
- AlarmLog (class in `uwsgiconf.options.alarm_types`), 64
- AlarmMule (class in `uwsgiconf.options.alarm_types`), 64
- Alarms (class in `uwsgiconf.options.alarms`), 64



- alarms (*uwsgi<sup>2</sup>conf.config.Section* attribute), 58
- Alarms.alarm\_types (class in *uwsgi<sup>2</sup>conf.options.alarms*), 64
- AlarmSignal (class in *uwsgi<sup>2</sup>conf.options.alarm\_types*), 64
- AlarmType (class in *uwsgi<sup>2</sup>conf.options.alarm\_types*), 64
- AlarmXmpp (class in *uwsgi<sup>2</sup>conf.options.alarm\_types*), 64
- Algo (class in *uwsgi<sup>2</sup>conf.options.workers\_cheapening*), 148
- AlgoBusyness (class in *uwsgi<sup>2</sup>conf.options.workers\_cheapening*), 148
- AlgoManual (class in *uwsgi<sup>2</sup>conf.options.workers\_cheapening*), 149
- AlgoQueue (class in *uwsgi<sup>2</sup>conf.options.workers\_cheapening*), 148
- AlgoSpare (class in *uwsgi<sup>2</sup>conf.options.workers\_cheapening*), 148
- AlgoSpare2 (class in *uwsgi<sup>2</sup>conf.options.workers\_cheapening*), 148
- alias (*uwsgi<sup>2</sup>conf.options.monitoring.Monitoring.metric\_types* attribute), 101
- APP\_LOAD\_POST (*uwsgi<sup>2</sup>conf.options.main\_process.MainProcess.phases* attribute), 85
- APP\_LOAD\_PRE (*uwsgi<sup>2</sup>conf.options.main\_process.MainProcess.phases* attribute), 85
- Applications (class in *uwsgi<sup>2</sup>conf.options.applications*), 66
- applications (in module *uwsgi<sup>2</sup>conf.uwsgi\_stub*), 14
- applications (*uwsgi<sup>2</sup>conf.config.Section* attribute), 59
- apps\_map (*uwsgi<sup>2</sup>conf.runtime.platform.\_Platform* attribute), 47
- ArgsFormatter (class in *uwsgi<sup>2</sup>conf.formatters*), 158
- as\_configuration() (*uwsgi<sup>2</sup>conf.config.Section* method), 60
- ASAP (*uwsgi<sup>2</sup>conf.options.main\_process.MainProcess.phases* attribute), 84
- async\_connect() (in module *uwsgi<sup>2</sup>conf.uwsgi\_stub*), 17
- async\_sleep() (in module *uwsgi<sup>2</sup>conf.uwsgi\_stub*), 17
- ASYNC\_SWITCHES (*uwsgi<sup>2</sup>conf.options.logging.Logging.vars* attribute), 81
- attach\_process() (*uwsgi<sup>2</sup>conf.options.master\_process.MasterProcess* method), 92
- attach\_process\_classic() (*uwsgi<sup>2</sup>conf.options.master\_process.MasterProcess* method), 92
- auth\_basic (*uwsgi<sup>2</sup>conf.options.routing.RouteRule.actions* attribute), 136
- auth\_ldap (*uwsgi<sup>2</sup>conf.options.routing.RouteRule.actions* attribute), 136
- AuthLdap (class in *uwsgi<sup>2</sup>conf.options.routing\_actions*), 116
- avg (*uwsgi<sup>2</sup>conf.options.monitoring.Monitoring.collectors* attribute), 101
- ## B
- BalancingAlgorithm (class in *uwsgi<sup>2</sup>conf.options.subscriptions\_algos*), 145
- BalancingAlgorithmWithBackup (class in *uwsgi<sup>2</sup>conf.options.subscriptions\_algos*), 145
- base64 (*uwsgi<sup>2</sup>conf.options.routing.RouteRule.var\_functions* attribute), 134
- BASIC (*uwsgi<sup>2</sup>conf.config.Section.embedded\_plugins\_presets* attribute), 59
- bootstrap() (*uwsgi<sup>2</sup>conf.config.Section* class method), 62
- Broodlord (class in *uwsgi<sup>2</sup>conf.presets.empire*), 9
- buffer\_size (in module *uwsgi<sup>2</sup>conf.uwsgi\_stub*), 14
- buffer\_size (*uwsgi<sup>2</sup>conf.runtime.platform.\_Platform* attribute), 47
- busyness (*uwsgi<sup>2</sup>conf.options.workers\_cheapening.Cheapening.algorithm* attribute), 150
- ## C
- Cache (class in *uwsgi<sup>2</sup>conf.runtime.caching*), 38
- cache (*uwsgi<sup>2</sup>conf.options.routing.Routing.modifiers* attribute), 138
- cache\_clear() (in module *uwsgi<sup>2</sup>conf.uwsgi\_stub*), 17
- cache\_dec() (in module *uwsgi<sup>2</sup>conf.uwsgi\_stub*), 17
- cache\_del() (in module *uwsgi<sup>2</sup>conf.uwsgi\_stub*), 18
- cache\_div() (in module *uwsgi<sup>2</sup>conf.uwsgi\_stub*), 18
- cache\_exists() (in module *uwsgi<sup>2</sup>conf.uwsgi\_stub*), 18
- cache\_get() (in module *uwsgi<sup>2</sup>conf.uwsgi\_stub*), 18
- cache\_inc() (in module *uwsgi<sup>2</sup>conf.uwsgi\_stub*), 18
- cache\_keys() (in module *uwsgi<sup>2</sup>conf.uwsgi\_stub*), 19
- cache\_mul() (in module *uwsgi<sup>2</sup>conf.uwsgi\_stub*), 19
- cache\_num() (in module *uwsgi<sup>2</sup>conf.uwsgi\_stub*), 19
- cache\_set() (in module *uwsgi<sup>2</sup>conf.uwsgi\_stub*), 19
- cache\_update() (in module *uwsgi<sup>2</sup>conf.uwsgi\_stub*), 19
- Caching (class in *uwsgi<sup>2</sup>conf.options.caching*), 67
- caching (*uwsgi<sup>2</sup>conf.config.Section* attribute), 59
- call (*uwsgi<sup>2</sup>conf.options.main\_process.MainProcess.actions* attribute), 84
- call() (in module *uwsgi<sup>2</sup>conf.uwsgi\_stub*), 20
- carbon (*uwsgi<sup>2</sup>conf.options.monitoring.Monitoring.pushers* attribute), 101

- cgi (uwsgi.conf.options.routing.Routing.modifiers attribute), 138  
 change\_dir() (uwsgi.conf.options.main\_process.MainProcess method), 87  
 Cheapening (class in uwsgi.conf.options.workers\_cheapening), 149  
 cheapening (uwsgi.conf.config.Section attribute), 59  
 Cheapening.algorithms (class in uwsgi.conf.options.workers\_cheapening), 149  
 chunked (uwsgi.conf.options.routing.RouteRule.transforms attribute), 135  
 chunked\_read() (in module uwsgi.conf.uwsgi\_stub), 20  
 chunked\_read\_nb() (in module uwsgi.conf.uwsgi\_stub), 20  
 cl() (in module uwsgi.conf.uwsgi\_stub), 20  
 clear() (uwsgi.conf.runtime.caching.Cache method), 39  
 clock (uwsgi.conf.runtime.platform.\_Platform attribute), 47  
 close() (in module uwsgi.conf.uwsgi\_stub), 21  
 cluster\_node (uwsgi.conf.options.routing.Routing.modifiers attribute), 138  
 Collector (class in uwsgi.conf.options.monitoring\_collectors), 99  
 CollectorAccumulator (class in uwsgi.conf.options.monitoring\_collectors), 100  
 CollectorAdder (class in uwsgi.conf.options.monitoring\_collectors), 100  
 CollectorAvg (class in uwsgi.conf.options.monitoring\_collectors), 100  
 CollectorFile (class in uwsgi.conf.options.monitoring\_collectors), 99  
 CollectorFunction (class in uwsgi.conf.options.monitoring\_collectors), 100  
 CollectorMultiplier (class in uwsgi.conf.options.monitoring\_collectors), 100  
 CollectorPointer (class in uwsgi.conf.options.monitoring\_collectors), 99  
 CollectorSum (class in uwsgi.conf.options.monitoring\_collectors), 100  
 command (uwsgi.conf.options.alarms.Alarms.alarm\_types attribute), 64  
 compress (uwsgi.conf.options.logging.Logging.encoders attribute), 76  
 CONF\_CURRENT\_SECTION (uwsgi.conf.config.Section.vars attribute), 62  
 CONF\_NAME\_ORIGINAL (uwsgi.conf.config.Section.vars attribute), 62  
 config (uwsgi.conf.runtime.platform.\_Platform attribute), 47  
 config\_from\_node (uwsgi.conf.options.routing.Routing.modifiers attribute), 138  
 config\_variables (uwsgi.conf.runtime.platform.\_Platform attribute), 47  
 Configuration (class in uwsgi.conf.config), 63  
 configure() (uwsgi.conf.presets.empire.Broodlord method), 10  
 configure\_certbot\_https() (uwsgi.conf.presets.nice.Section method), 11  
 configure\_maintenance\_mode() (uwsgi.conf.presets.nice.Section method), 11  
 configure\_owner() (uwsgi.conf.presets.nice.Section method), 11  
 configure\_uwsgi() (in module uwsgi.conf.config), 63  
 connect() (in module uwsgi.conf.runtime.asynced), 37  
 connect() (in module uwsgi.conf.uwsgi\_stub), 21  
 connection\_fd() (in module uwsgi.conf.uwsgi\_stub), 21  
 contains() (uwsgi.conf.options.routing\_subjects.SubjectCustom method), 131  
 contains\_ipv4() (uwsgi.conf.options.routing\_subjects.SubjectCustom method), 131  
 contains\_ipv6() (uwsgi.conf.options.routing\_subjects.SubjectCustom method), 131  
 content\_length (uwsgi.conf.runtime.request.\_Request attribute), 48  
 cookie (uwsgi.conf.options.routing.RouteRule.vars attribute), 134  
 CORE (uwsgi.conf.options.logging.Logging.vars attribute), 81  
 corerouter\_signal (uwsgi.conf.options.routing.Routing.modifiers attribute), 138  
 cores (in module uwsgi.conf.uwsgi\_stub), 14  
 cores\_count (uwsgi.conf.runtime.platform.\_Platform attribute), 47  
 counter (uwsgi.conf.options.monitoring.Monitoring.metric\_types attribute), 101  
 CPU\_CORES (uwsgi.conf.config.Section.vars attribute), 62  
 curl (uwsgi.conf.options.alarms.Alarms.alarm\_types attribute), 64  
 custom (uwsgi.conf.options.routing.RouteRule.subjects attribute), 135  
**D**  
 daemonize() (uwsgi.conf.options.main\_process.MainProcess method), 87

- `decr()` (*uwsgi<sup>2</sup>conf.runtime.caching.Cache method*), 39  
`decr()` (*uwsgi<sup>2</sup>conf.runtime.monitoring.Metric method*), 43  
`default` (*uwsgi<sup>2</sup>conf.options.networking.Networking.socket attribute*), 109  
`delete()` (*uwsgi<sup>2</sup>conf.runtime.caching.Cache method*), 39  
`derive_from()` (*uwsgi<sup>2</sup>conf.config.Section class method*), 61  
`device_add_rule()` (*uwsgi<sup>2</sup>conf.options.routing\_routers.RouterTunTap method*), 127  
`device_connect()` (*uwsgi<sup>2</sup>conf.options.routing\_routers.RouterTunTap method*), 126  
`dir_change` (*uwsgi<sup>2</sup>conf.options.main\_process.MainProcess.actions attribute*), 84  
`dir_change` (*uwsgi<sup>2</sup>conf.options.routing.RouteRule.actions attribute*), 136  
`dir_create` (*uwsgi<sup>2</sup>conf.options.main\_process.MainProcess.actions attribute*), 84  
`DIR_DOCUMENT_ROOT` (*uwsgi<sup>2</sup>conf.options.statics.Statics attribute*), 143  
`DIR_VASSALS` (*uwsgi<sup>2</sup>conf.config.Section.vars attribute*), 62  
`disconnect()` (*in module uwsgi<sup>2</sup>conf.runtime.control*), 40  
`disconnect()` (*in module uwsgi<sup>2</sup>conf.uwsgi\_stub*), 21  
`div()` (*uwsgi<sup>2</sup>conf.runtime.caching.Cache method*), 40  
`div()` (*uwsgi<sup>2</sup>conf.runtime.monitoring.Metric method*), 44  
`do_break` (*uwsgi<sup>2</sup>conf.options.routing.RouteRule.actions attribute*), 136  
`do_continue` (*uwsgi<sup>2</sup>conf.options.routing.RouteRule.actions attribute*), 136  
`do_goto` (*uwsgi<sup>2</sup>conf.options.routing.RouteRule.actions attribute*), 136
- ## E
- `echo` (*uwsgi<sup>2</sup>conf.options.routing.Routing.modifiers attribute*), 138  
`embedded_data()` (*in module uwsgi<sup>2</sup>conf.uwsgi\_stub*), 21  
`EMPEROR_LOST` (*uwsgi<sup>2</sup>conf.options.main\_process.MainProcess.phases attribute*), 85  
`EMPEROR_RELOAD` (*uwsgi<sup>2</sup>conf.options.main\_process.MainProcess.phases attribute*), 85  
`EMPEROR_START` (*uwsgi<sup>2</sup>conf.options.main\_process.MainProcess.phases attribute*), 85  
`EMPEROR_STOP` (*uwsgi<sup>2</sup>conf.options.main\_process.MainProcess.phases attribute*), 85  
`Empire` (*class in uwsgi<sup>2</sup>conf.options.empire*), 70  
`empire` (*uwsgi<sup>2</sup>conf.config.Section attribute*), 59  
`enable()` (*uwsgi<sup>2</sup>conf.options.queue.Queue method*), 112  
`enable_snmp()` (*uwsgi<sup>2</sup>conf.options.monitoring.Monitoring method*), 103  
`Encoder` (*class in uwsgi<sup>2</sup>conf.options.logging\_encoders*), 75  
`EncoderCompress` (*class in uwsgi<sup>2</sup>conf.options.logging\_encoders*), 75  
`EncoderFormat` (*class in uwsgi<sup>2</sup>conf.options.logging\_encoders*), 75  
`EncoderFormat.vars` (*class in uwsgi<sup>2</sup>conf.options.logging\_encoders*), 75  
`EncoderGzip` (*class in uwsgi<sup>2</sup>conf.options.logging\_encoders*), 75  
`EncoderJson` (*class in uwsgi<sup>2</sup>conf.options.logging\_encoders*), 75  
`EncoderNewline` (*class in uwsgi<sup>2</sup>conf.options.logging\_encoders*), 75  
`EncoderPrefix` (*class in uwsgi<sup>2</sup>conf.options.logging\_encoders*), 75  
`EncoderSuffix` (*class in uwsgi<sup>2</sup>conf.options.logging\_encoders*), 75  
`endswith()` (*uwsgi<sup>2</sup>conf.options.routing\_subjects.SubjectCustom method*), 131  
`env` (*in module uwsgi<sup>2</sup>conf.uwsgi\_stub*), 14  
`env` (*uwsgi<sup>2</sup>conf.runtime.request.\_Request attribute*), 48  
`env()` (*uwsgi<sup>2</sup>conf.config.Section method*), 61  
`eq()` (*uwsgi<sup>2</sup>conf.options.routing\_subjects.SubjectCustom method*), 131  
`ERROR` (*uwsgi<sup>2</sup>conf.options.routing.RouteRule.stages attribute*), 134  
`eval` (*uwsgi<sup>2</sup>conf.options.routing.Routing.modifiers attribute*), 138  
`eval_wsgi_entrypoint()` (*uwsgi<sup>2</sup>conf.options.python.Python method*), 157  
`example` (*uwsgi<sup>2</sup>conf.options.routing.Routing.modifiers attribute*), 138  
`execute` (*uwsgi<sup>2</sup>conf.options.main\_process.MainProcess.actions attribute*), 84  
`exists()` (*uwsgi<sup>2</sup>conf.options.routing\_subjects.SubjectCustom method*), 131  
`exit` (*uwsgi<sup>2</sup>conf.options.main\_process.MainProcess.actions attribute*), 84  
`EXIT` (*uwsgi<sup>2</sup>conf.options.main\_process.MainProcess.phases attribute*), 85  
`extract()` (*in module uwsgi<sup>2</sup>conf.uwsgi\_stub*), 21

## F

- Farm (class in `uwsgi2conf.runtime.mules`), 45
- farm\_get\_msg() (in module `uwsgi2conf.uwsgi_stub`), 21
- farm\_msg() (in module `uwsgi2conf.uwsgi_stub`), 22
- fast (`uwsgi2conf.options.routing.Routing.routers` attribute), 138
- fastcgi (`uwsgi2conf.options.networking.Networking.sockets` attribute), 109
- fastfunc (`uwsgi2conf.options.routing.Routing.modifiers` attribute), 138
- fd (`uwsgi2conf.runtime.request._Request` attribute), 48
- fifo\_write (`uwsgi2conf.options.main_process.MainProcess.actions` attribute), 84
- file (`uwsgi2conf.options.logging.Logging.loggers` attribute), 75
- file (`uwsgi2conf.options.monitoring.Monitoring.collectors` attribute), 101
- file (`uwsgi2conf.options.monitoring.Monitoring.pushers` attribute), 101
- file\_create (`uwsgi2conf.options.main_process.MainProcess.actions` attribute), 84
- file\_write (`uwsgi2conf.options.main_process.MainProcess.actions` attribute), 84
- FILENAME (`uwsgi2conf.options.statics.Statics.expiration_criteria` attribute), 143
- FINAL (`uwsgi2conf.options.routing.RouteRule.stages` attribute), 134
- fix\_content\_len (`uwsgi2conf.options.routing.RouteRule.transforms` attribute), 135
- fix\_var\_path\_info (`uwsgi2conf.options.routing.RouteRule.actions` attribute), 136
- flush (`uwsgi2conf.options.routing.RouteRule.transforms` attribute), 135
- forkpty (`uwsgi2conf.options.routing.Routing.routers` attribute), 138
- format (`uwsgi2conf.options.logging.Logging.encoders` attribute), 76
- format() (`uwsgi2conf.config.Configuration` method), 63
- FORMAT\_ESCAPE (`uwsgi2conf.config.Section.vars` attribute), 62
- format\_print\_text() (in module `uwsgi2conf.formatters`), 158
- FormatterBase (class in `uwsgi2conf.formatters`), 158
- FORMATTERS (in module `uwsgi2conf.formatters`), 158
- Forwarder (class in `uwsgi2conf.options.routing_routers`), 118
- ForwarderCache (class in `uwsgi2conf.options.routing_routers`), 118
- ForwarderCode (class in `uwsgi2conf.options.routing_routers`), 118
- ForwarderPath (class in `uwsgi2conf.options.routing_routers`), 118
- ForwarderSocket (class in `uwsgi2conf.options.routing_routers`), 118
- ForwarderSubscriptionServer (class in `uwsgi2conf.options.routing_routers`), 118
- from\_dsn() (`uwsgi2conf.options.networking.Networking.sockets` class method), 109
- Func (class in `uwsgi2conf.options.routing_vars`), 132
- func (`uwsgi2conf.runtime.signals.SignalDescription` attribute), 53
- FuncBase64 (class in `uwsgi2conf.options.routing_vars`), 133
- FuncHex (class in `uwsgi2conf.options.routing_vars`), 133
- FuncLower (class in `uwsgi2conf.options.routing_vars`), 133
- FuncMath (class in `uwsgi2conf.options.routing_vars`), 133
- FuncMime (class in `uwsgi2conf.options.routing_vars`), 133
- function (`uwsgi2conf.options.monitoring.Monitoring.collectors` attribute), 101
- FuncUpper (class in `uwsgi2conf.options.routing_vars`), 133

## G

- GATEWAY\_START\_IN\_EACH (`uwsgi2conf.options.main_process.MainProcess.phases` attribute), 86
- gauge (`uwsgi2conf.options.monitoring.Monitoring.metric_types` attribute), 101
- ge() (`uwsgi2conf.options.routing_subjects.SubjectCustom` method), 131
- geoup (`uwsgi2conf.options.routing.RouteRule.vars` attribute), 134
- get() (`uwsgi2conf.runtime.caching.Cache` method), 39
- get\_available\_num() (in module `uwsgi2conf.runtime.signals`), 53
- get\_bundled\_static\_path() (`uwsgi2conf.presets.nice.Section` class method), 11
- get\_by\_basename() (`uwsgi2conf.runtime.spooler.Spooler` class method), 56
- get\_certbot\_paths() (`uwsgi2conf.options.networking_sockets.SocketHttps` class method), 105
- get\_current() (`uwsgi2conf.runtime.mules.Mule` class method), 45
- get\_current\_id() (`uwsgi2conf.runtime.mules.Mule` class method), 44
- get\_current\_log\_size() (in module `uwsgi2conf.runtime.logging`), 42

- get\_descriptions() (uwsgi<sup>2</sup>conf.config.Section.vars class method), 62  
 get\_farms() (uwsgi<sup>2</sup>conf.runtime.mules.Farm class method), 45  
 get\_last\_received() (in module uwsgi<sup>2</sup>conf.runtime.signals), 53  
 get\_listen\_queue() (uwsgi<sup>2</sup>conf.runtime.platform.\_Platform method), 47  
 get\_log\_format\_default() (uwsgi<sup>2</sup>conf.presets.nice.Section method), 10  
 get\_logvar() (in module uwsgi<sup>2</sup>conf.uwsgi\_stub), 22  
 get\_loop\_name() (in module uwsgi<sup>2</sup>conf.runtime.asynced), 37  
 get\_message() (uwsgi<sup>2</sup>conf.runtime.mules.Farm class method), 46  
 get\_message() (uwsgi<sup>2</sup>conf.runtime.mules.Mule class method), 45  
 get\_owner() (uwsgi<sup>2</sup>conf.options.main\_process.MainProcess method), 88  
 get\_pids() (uwsgi<sup>2</sup>conf.runtime.spooler.Spooler class method), 56  
 get\_rpc\_list() (in module uwsgi<sup>2</sup>conf.runtime.rpc), 49  
 get\_runtime\_dir() (uwsgi<sup>2</sup>conf.config.Section method), 60  
 get\_spoolers() (uwsgi<sup>2</sup>conf.runtime.spooler.Spooler class method), 56  
 get\_tasks() (uwsgi<sup>2</sup>conf.runtime.spooler.Spooler class method), 56  
 get\_version() (uwsgi<sup>2</sup>conf.runtime.platform.\_Platform method), 48  
 glusterfs (uwsgi<sup>2</sup>conf.options.routing.Routing.modifiers attribute), 139  
 green\_schedule() (in module uwsgi<sup>2</sup>conf.uwsgi\_stub), 22  
 gridfs (uwsgi<sup>2</sup>conf.options.routing.Routing.modifiers attribute), 139  
 GROUP\_ID (uwsgi<sup>2</sup>conf.config.Section.vars attribute), 62  
 GROUP\_NAME (uwsgi<sup>2</sup>conf.config.Section.vars attribute), 62  
 gt() (uwsgi<sup>2</sup>conf.options.routing\_subjects.SubjectCustom method), 131  
 gzip (uwsgi<sup>2</sup>conf.options.logging.Logging.encoders attribute), 76  
 gzip (uwsgi<sup>2</sup>conf.options.routing.RouteRule.transforms attribute), 135
- ## H
- harakiri\_imposed (class in uwsgi<sup>2</sup>conf.runtime.control), 41  
 has\_threads (in module uwsgi<sup>2</sup>conf.uwsgi\_stub), 14  
 header\_add (uwsgi<sup>2</sup>conf.options.routing.RouteRule.actions attribute), 136  
 header\_add() (uwsgi<sup>2</sup>conf.options.routing.Routing method), 141  
 header\_collect() (uwsgi<sup>2</sup>conf.options.routing.Routing method), 141  
 header\_remove (uwsgi<sup>2</sup>conf.options.routing.RouteRule.actions attribute), 136  
 header\_remove() (uwsgi<sup>2</sup>conf.options.routing.Routing method), 141  
 headers\_off (uwsgi<sup>2</sup>conf.options.routing.RouteRule.actions attribute), 136  
 headers\_reset (uwsgi<sup>2</sup>conf.options.routing.RouteRule.actions attribute), 136  
 hex (uwsgi<sup>2</sup>conf.options.routing.RouteRule.var\_functions attribute), 134  
 HookAction (class in uwsgi<sup>2</sup>conf.options.main\_process\_actions), 82  
 HOST\_NAME (uwsgi<sup>2</sup>conf.config.Section.vars attribute), 62  
 hostname (in module uwsgi<sup>2</sup>conf.uwsgi\_stub), 15  
 hostname (uwsgi<sup>2</sup>conf.runtime.platform.\_Platform attribute), 47  
 http (uwsgi<sup>2</sup>conf.options.networking.Networking.sockets attribute), 109  
 http (uwsgi<sup>2</sup>conf.options.routing.Routing.routers attribute), 138  
 http\_host (uwsgi<sup>2</sup>conf.options.routing.RouteRule.subjects attribute), 135  
 http\_referer (uwsgi<sup>2</sup>conf.options.routing.RouteRule.subjects attribute), 135  
 http\_user\_agent (uwsgi<sup>2</sup>conf.options.routing.RouteRule.subjects attribute), 135  
 https (uwsgi<sup>2</sup>conf.options.networking.Networking.sockets attribute), 109  
 https (uwsgi<sup>2</sup>conf.options.routing.Routing.routers attribute), 138  
 httptime (uwsgi<sup>2</sup>conf.options.routing.RouteRule.vars attribute), 134
- ## I
- i\_am\_the\_lord() (in module uwsgi<sup>2</sup>conf.uwsgi\_stub), 33  
 i\_am\_the\_spooler() (in module uwsgi<sup>2</sup>conf.uwsgi\_stub), 22  
 id (uwsgi<sup>2</sup>conf.runtime.request.\_Request attribute), 48  
 import\_module() (uwsgi<sup>2</sup>conf.options.python.Python method), 157  
 in\_farm() (in module uwsgi<sup>2</sup>conf.uwsgi\_stub), 22

- include() (*uwsgiconf.config.Section* method), 61
- incr() (*uwsgiconf.runtime.caching.Cache* method), 39
- incr() (*uwsgiconf.runtime.monitoring.Metric* method), 43
- IniFormatter (class in *uwsgiconf.formatters*), 158
- ip\_hash (*uwsgiconf.options.subscriptions.Subscriptions.algorithms* attribute), 145
- IpHash (class in *uwsgiconf.options.subscriptions\_algos*), 145
- is\_a\_reload() (in module *uwsgiconf.uwsgi\_stub*), 22
- is\_connected() (in module *uwsgiconf.uwsgi\_stub*), 22
- is\_locked() (in module *uwsgiconf.uwsgi\_stub*), 23
- is\_mine (*uwsgiconf.runtime.mules.Farm* attribute), 46
- is\_set (*uwsgiconf.runtime.locking.Lock* attribute), 42
- is\_stub (in module *uwsgiconf.uwsgi\_stub*), 14
- isdir() (*uwsgiconf.options.routing\_subjects.SubjectCustom* method), 131
- isempty() (*uwsgiconf.options.routing\_subjects.SubjectCustom* method), 131
- isexec() (*uwsgiconf.options.routing\_subjects.SubjectCustom* method), 131
- isfile() (*uwsgiconf.options.routing\_subjects.SubjectCustom* method), 131
- islink() (*uwsgiconf.options.routing\_subjects.SubjectCustom* method), 131
- islord() (*uwsgiconf.options.routing\_subjects.SubjectCustom* method), 131
- issue\_alarm() (in module *uwsgiconf.runtime.alarms*), 37
- iter\_options() (*uwsgiconf.formatters.FormatterBase* method), 158
- ## J
- JAIL\_IN (*uwsgiconf.options.main\_process.MainProcess.phase* attribute), 85
- JAIL\_POST (*uwsgiconf.options.main\_process.MainProcess.phase* attribute), 85
- JAIL\_PRE (*uwsgiconf.options.main\_process.MainProcess.phase* attribute), 85
- json (*uwsgiconf.options.logging.Logging.encoders* attribute), 76
- jvm (*uwsgiconf.options.routing.Routing.modifiers* attribute), 139
- ## K
- keys (*uwsgiconf.runtime.caching.Cache* attribute), 39
- ## L
- le() (*uwsgiconf.options.routing\_subjects.SubjectCustom* method), 131
- least\_reference\_count (*uwsgiconf.options.subscriptions.Subscriptions.algorithms* attribute), 145
- LeastReferenceCount (class in *uwsgiconf.options.subscriptions\_algos*), 145
- least\_refs\_msg (*uwsgiconf.options.routing.Routing.modifiers* attribute), 139
- listen\_queue() (in module *uwsgiconf.uwsgi\_stub*), 23
- Lock (class in *uwsgiconf.runtime.locking*), 41
- lock (in module *uwsgiconf.runtime.locking*), 42
- lock() (in module *uwsgiconf.uwsgi\_stub*), 23
- lock\_file() (*uwsgiconf.options.locks.Lock* method), 72
- Locks (class in *uwsgiconf.options.locks*), 72
- locks (*uwsgiconf.config.Section* attribute), 59
- log (*uwsgiconf.options.alarms.Alarms.alarm\_types* attribute), 64
- log (*uwsgiconf.options.routing.RouteRule.actions* attribute), 136
- log() (in module *uwsgiconf.uwsgi\_stub*), 23
- log() (*uwsgiconf.runtime.request.\_Request* method), 48
- log\_into() (*uwsgiconf.options.logging.Logging* method), 77
- log\_message() (in module *uwsgiconf.runtime.logging*), 42
- log\_this\_request() (in module *uwsgiconf.uwsgi\_stub*), 23
- Logger (class in *uwsgiconf.options.logging\_loggers*), 73
- LoggerFile (class in *uwsgiconf.options.logging\_loggers*), 73
- LoggerMongo (class in *uwsgiconf.options.logging\_loggers*), 74
- LoggerRedis (class in *uwsgiconf.options.logging\_loggers*), 74
- LoggerRsyslog (class in *uwsgiconf.options.logging\_loggers*), 73
- LoggerSocket (class in *uwsgiconf.options.logging\_loggers*), 73
- LoggerSyslog (class in *uwsgiconf.options.logging\_loggers*), 73
- LoggerZeroMq (class in *uwsgiconf.options.logging\_loggers*), 74
- Logging (class in *uwsgiconf.options.logging*), 75
- logging (*uwsgiconf.config.Section* attribute), 59
- Logging.encoders (class in *uwsgiconf.options.logging*), 76
- Logging.loggers (class in *uwsgiconf.options.logging*), 75
- Logging.vars (class in *uwsgiconf.options.logging*), 80
- logsize() (in module *uwsgiconf.uwsgi\_stub*), 23
- loop() (in module *uwsgiconf.uwsgi\_stub*), 23

- lord\_scroll() (in module uwsgiconf.uwsgi\_stub), 33
- lower (uwsgiconf.options.routing.RouteRule.var\_functions attribute), 134
- lt() (uwsgiconf.options.routing\_subjects.SubjectCustom method), 131
- lua (uwsgiconf.options.routing.Routing.modifiers attribute), 139
- ## M
- magic\_table (in module uwsgiconf.uwsgi\_stub), 15
- main\_process (uwsgiconf.config.Section attribute), 59
- MainProcess (class in uwsgiconf.options.main\_process), 84
- MainProcess.actions (class in uwsgiconf.options.main\_process), 84
- MainProcess.phases (class in uwsgiconf.options.main\_process), 84
- make\_rpc\_call() (in module uwsgiconf.runtime.rpc), 49
- manage (uwsgiconf.options.routing.Routing.modifiers attribute), 139
- manage\_path\_info (uwsgiconf.options.routing.Routing.modifiers attribute), 139
- manual (uwsgiconf.options.workers\_cheapening.Cheapening attribute), 150
- mark\_processed (uwsgiconf.runtime.spooler.SpoolerTask attribute), 56
- mark\_rescheduled (uwsgiconf.runtime.spooler.SpoolerTask attribute), 57
- mark\_skipped (uwsgiconf.runtime.spooler.SpoolerTask attribute), 57
- master\_pid (uwsgiconf.runtime.platform.\_Platform attribute), 47
- master\_process (uwsgiconf.config.Section attribute), 59
- MASTER\_START (uwsgiconf.options.main\_process.MainProcess.phases attribute), 85
- masterpid() (in module uwsgiconf.uwsgi\_stub), 23
- MasterProcess (class in uwsgiconf.options.master\_process), 89
- matches() (uwsgiconf.options.routing\_subjects.SubjectCustom method), 131
- math (uwsgiconf.options.routing.RouteRule.var\_functions attribute), 134
- mem() (in module uwsgiconf.uwsgi\_stub), 23
- MEM\_RSS (uwsgiconf.options.logging.Logging.vars attribute), 81
- MEM\_RSS\_MV (uwsgiconf.options.logging.Logging.vars attribute), 82
- MEM\_VSZ (uwsgiconf.options.logging.Logging.vars attribute), 81
- MEM\_VSZ\_MB (uwsgiconf.options.logging.Logging.vars attribute), 82
- memory (uwsgiconf.runtime.platform.\_Platform attribute), 47
- MESSAGE (uwsgiconf.options.logging\_encoders.EncoderFormat.vars attribute), 75
- message (uwsgiconf.options.routing.Routing.modifiers attribute), 139
- message\_array (uwsgiconf.options.routing.Routing.modifiers attribute), 139
- message\_marshall (uwsgiconf.options.routing.Routing.modifiers attribute), 139
- MESSAGE\_NEWLINE (uwsgiconf.options.logging\_encoders.EncoderFormat.vars attribute), 75
- Metric (class in uwsgiconf.options.monitoring\_metric\_types), 93
- Metric (class in uwsgiconf.runtime.monitoring), 43
- metric (uwsgiconf.options.logging.Logging.vars attribute), 82
- metric\_algorithm (uwsgiconf.options.routing.RouteRule.vars attribute), 134
- metric\_dec() (in module uwsgiconf.uwsgi\_stub), 24
- metric\_div() (in module uwsgiconf.uwsgi\_stub), 24
- metric\_get() (in module uwsgiconf.uwsgi\_stub), 24
- metric\_inc() (in module uwsgiconf.uwsgi\_stub), 24
- metric\_mul() (in module uwsgiconf.uwsgi\_stub), 24
- metric\_set() (in module uwsgiconf.uwsgi\_stub), 24
- metric\_set\_max() (in module uwsgiconf.uwsgi\_stub), 24
- metric\_set\_min() (in module uwsgiconf.uwsgi\_stub), 25
- MetricTypeAbsolute (class in uwsgiconf.options.monitoring\_metric\_types), 95
- MetricTypeAlias (class in uwsgiconf.options.monitoring\_metric\_types), 96
- MetricTypeCounter (class in uwsgiconf.options.monitoring\_metric\_types), 94
- MetricTypeGauge (class in uwsgiconf.options.monitoring\_metric\_types), 95
- micros() (in module uwsgiconf.uwsgi\_stub), 25
- mime (uwsgiconf.options.routing.RouteRule.var\_functions attribute), 134
- MIME\_TYPE (uwsgiconf.options.statics.Statics.expiration\_criteria attribute), 143
- MOD1 (uwsgiconf.options.logging.Logging.vars attribute), 82
- MOD2 (uwsgiconf.options.logging.Logging.vars attribute), 82

- tribute), 82
- ModifierCache (class in uwsgi-  
conf.options.routing\_modifiers), 129
- ModifierCgi (class in uwsgi-  
conf.options.routing\_modifiers), 128
- ModifierClusterNode (class in uwsgi-  
conf.options.routing\_modifiers), 129
- ModifierConfigFromNode (class in uwsgi-  
conf.options.routing\_modifiers), 129
- ModifierCorerouterSignal (class in uwsgi-  
conf.options.routing\_modifiers), 130
- ModifierEcho (class in uwsgi-  
conf.options.routing\_modifiers), 129
- ModifierEval (class in uwsgi-  
conf.options.routing\_modifiers), 128
- ModifierExample (class in uwsgi-  
conf.options.routing\_modifiers), 130
- ModifierFastfunc (class in uwsgi-  
conf.options.routing\_modifiers), 128
- ModifierGlusterfs (class in uwsgi-  
conf.options.routing\_modifiers), 128
- ModifierGridfs (class in uwsgi-  
conf.options.routing\_modifiers), 128
- ModifierJvm (class in uwsgi-  
conf.options.routing\_modifiers), 128
- ModifierLegionMsg (class in uwsgi-  
conf.options.routing\_modifiers), 129
- ModifierLua (class in uwsgi-  
conf.options.routing\_modifiers), 127
- ModifierManage (class in uwsgi-  
conf.options.routing\_modifiers), 128
- ModifierManagePathInfo (class in uwsgi-  
conf.options.routing\_modifiers), 128
- ModifierMessage (class in uwsgi-  
conf.options.routing\_modifiers), 128
- ModifierMessageArray (class in uwsgi-  
conf.options.routing\_modifiers), 128
- ModifierMessageMarshal (class in uwsgi-  
conf.options.routing\_modifiers), 128
- ModifierMono (class in uwsgi-  
conf.options.routing\_modifiers), 128
- ModifierMulticast (class in uwsgi-  
conf.options.routing\_modifiers), 129
- ModifierMulticastAnnounce (class in uwsgi-  
conf.options.routing\_modifiers), 129
- ModifierPersistentClose (class in uwsgi-  
conf.options.routing\_modifiers), 130
- ModifierPhp (class in uwsgi-  
conf.options.routing\_modifiers), 128
- ModifierPing (class in uwsgi-  
conf.options.routing\_modifiers), 129
- ModifierPsgi (class in uwsgi-  
conf.options.routing\_modifiers), 127
- ModifierRack (class in uwsgi-  
conf.options.routing\_modifiers), 127
- ModifierRados (class in uwsgi-  
conf.options.routing\_modifiers), 128
- ModifierRaw (class in uwsgi-  
conf.options.routing\_modifiers), 129
- ModifierReload (class in uwsgi-  
conf.options.routing\_modifiers), 129
- ModifierReloadBrutal (class in uwsgi-  
conf.options.routing\_modifiers), 129
- ModifierRemoteLogging (class in uwsgi-  
conf.options.routing\_modifiers), 129
- ModifierResponse (class in uwsgi-  
conf.options.routing\_modifiers), 130
- ModifierRpc (class in uwsgi-  
conf.options.routing\_modifiers), 130
- ModifierSignal (class in uwsgi-  
conf.options.routing\_modifiers), 129
- ModifierSnmp (class in uwsgi-  
conf.options.routing\_modifiers), 129
- ModifierSpooler (class in uwsgi-  
conf.options.routing\_modifiers), 128
- ModifierSubscription (class in uwsgi-  
conf.options.routing\_modifiers), 130
- ModifierSymcall (class in uwsgi-  
conf.options.routing\_modifiers), 128
- ModifierV8 (class in uwsgi-  
conf.options.routing\_modifiers), 128
- ModifierWsgi (class in uwsgi-  
conf.options.routing\_modifiers), 127
- ModifierXslt (class in uwsgi-  
conf.options.routing\_modifiers), 128
- mongo (uwsgi<sup>2</sup>conf.options.logging.Logging.loggers attribute), 76
- mongo (uwsgi<sup>2</sup>conf.options.monitoring.Monitoring.pushers attribute), 101
- Monitoring (class in uwsgi<sup>2</sup>conf.options.monitoring), 100
- monitoring (uwsgi<sup>2</sup>conf.config.Section attribute), 59
- Monitoring.collectors (class in uwsgi-  
conf.options.monitoring), 101
- Monitoring.metric\_types (class in uwsgi-  
conf.options.monitoring), 100
- Monitoring.pushers (class in uwsgi-  
conf.options.monitoring), 101
- mono (uwsgi<sup>2</sup>conf.options.routing.Routing.modifiers attribute), 139
- mount (uwsgi<sup>2</sup>conf.options.main\_process.MainProcess.actions attribute), 84
- mount () (uwsgi<sup>2</sup>conf.options.applications.Applications method), 66
- mul () (uwsgi<sup>2</sup>conf.runtime.caching.Cache method), 40
- mul () (uwsgi<sup>2</sup>conf.runtime.monitoring.Metric method), 44
- Mule (class in uwsgi<sup>2</sup>conf.runtime.mules), 44



- mule (*uwsgi<sup>2</sup>conf.options.alarms.Alarms.alarms\_types* attribute), 64
- mule\_farm (*uwsgi<sup>2</sup>conf.options.workers.Workers* attribute), 151
- mule\_get\_msg() (in module *uwsgi<sup>2</sup>conf.uwsgi\_stub*), 25
- mule\_id() (in module *uwsgi<sup>2</sup>conf.uwsgi\_stub*), 25
- mule\_msg() (in module *uwsgi<sup>2</sup>conf.uwsgi\_stub*), 25
- mule\_msg\_hook() (in module *uwsgi<sup>2</sup>conf.uwsgi\_stub*), 15
- mule\_offload() (in module *uwsgi<sup>2</sup>conf.runtime.mules*), 44
- MULE\_START\_IN\_EACH (*uwsgi<sup>2</sup>conf.options.main\_process.MainProcess.phases* attribute), 86
- MuleFarm (class in *uwsgi<sup>2</sup>conf.options.workers*), 150
- multicast (*uwsgi<sup>2</sup>conf.options.routing.Routing.modifiers* attribute), 139
- multicast\_announce (*uwsgi<sup>2</sup>conf.options.routing.Routing.modifiers* attribute), 139
- multiplier (*uwsgi<sup>2</sup>conf.options.monitoring.Monitoring.collectors* attribute), 101
- ## N
- Networking (class in *uwsgi<sup>2</sup>conf.options.networking*), 109
- networking (*uwsgi<sup>2</sup>conf.config.Section* attribute), 59
- Networking.sockets (class in *uwsgi<sup>2</sup>conf.options.networking*), 109
- newline (*uwsgi<sup>2</sup>conf.options.logging.Logging.encoders* attribute), 76
- num (*uwsgi<sup>2</sup>conf.runtime.signals.SignalDescription* attribute), 53
- numproc (in module *uwsgi<sup>2</sup>conf.uwsgi\_stub*), 15
- ## O
- offload() (in module *uwsgi<sup>2</sup>conf.uwsgi\_stub*), 25
- offload() (*uwsgi<sup>2</sup>conf.runtime.mules.Farm* method), 45
- offload() (*uwsgi<sup>2</sup>conf.runtime.mules.Mule* method), 44
- offload\_off (*uwsgi<sup>2</sup>conf.options.routing.RouteRule.actions* attribute), 136
- opt (in module *uwsgi<sup>2</sup>conf.uwsgi\_stub*), 15
- ## P
- parsefile() (in module *uwsgi<sup>2</sup>conf.uwsgi\_stub*), 26
- path\_info (*uwsgi<sup>2</sup>conf.options.routing.RouteRule.subjects* attribute), 135
- PATH\_INFO (*uwsgi<sup>2</sup>conf.options.statics.Statics.expiration\_criteria* attribute), 143
- persistent\_close (*uwsgi<sup>2</sup>conf.options.routing.Routing.modifiers* attribute), 139
- php (*uwsgi<sup>2</sup>conf.options.routing.Routing.modifiers* attribute), 139
- ping (*uwsgi<sup>2</sup>conf.options.routing.Routing.modifiers* attribute), 139
- pointer (*uwsgi<sup>2</sup>conf.options.monitoring.Monitoring.collectors* attribute), 101
- post\_fork\_hook (in module *uwsgi<sup>2</sup>conf.uwsgi\_stub*), 15
- postfork\_hooks (*uwsgi<sup>2</sup>conf.runtime.platform.\_Platform* attribute), 46
- prefix (*uwsgi<sup>2</sup>conf.options.logging.Logging.encoders* attribute), 76
- print\_alarms() (*uwsgi<sup>2</sup>conf.options.alarms.Alarms* method), 65
- print\_algorithms() (*uwsgi<sup>2</sup>conf.options.workers\_cheapening.Cheapening* method), 150
- print\_ini() (*uwsgi<sup>2</sup>conf.config.Configuration* method), 63
- print\_loggers() (*uwsgi<sup>2</sup>conf.options.logging.Logging* method), 79
- print\_monitors() (*uwsgi<sup>2</sup>conf.options.empire.Empire* method), 70
- print\_out() (*uwsgi<sup>2</sup>conf.config.Section* method), 60
- print\_plugins() (*uwsgi<sup>2</sup>conf.config.Section* method), 60
- print\_routing\_rules() (*uwsgi<sup>2</sup>conf.options.routing.Routing* method), 140
- print\_stamp() (*uwsgi<sup>2</sup>conf.config.Section* method), 60
- print\_variables() (*uwsgi<sup>2</sup>conf.config.Section* method), 60
- printout (*uwsgi<sup>2</sup>conf.options.main\_process.MainProcess.actions* attribute), 84
- PRIV\_DROP\_POST (*uwsgi<sup>2</sup>conf.options.main\_process.MainProcess.phases* attribute), 85
- PRIV\_DROP\_PRE (*uwsgi<sup>2</sup>conf.options.main\_process.MainProcess.phases* attribute), 85
- PROBE() (*uwsgi<sup>2</sup>conf.config.Section.embedded\_plugins\_presets* static method), 59
- process() (*uwsgi<sup>2</sup>conf.runtime.spooler.SpoolerFunctionCallTask* method), 57
- process() (*uwsgi<sup>2</sup>conf.runtime.spooler.SpoolerTask* method), 57
- project\_name (*uwsgi<sup>2</sup>conf.config.Section* attribute), 60
- psgi (*uwsgi<sup>2</sup>conf.options.routing.Routing.modifiers* attribute), 139

- Pusher (class in `uwsgi-conf.options.monitoring_pushers`), 97  
 PusherCarbon (class in `uwsgi-conf.options.monitoring_pushers`), 98  
 PusherFile (class in `uwsgi-conf.options.monitoring_pushers`), 99  
 PusherMongo (class in `uwsgi-conf.options.monitoring_pushers`), 99  
 PusherRrdtool (class in `uwsgi-conf.options.monitoring_pushers`), 97  
 PusherSocket (class in `uwsgi-conf.options.monitoring_pushers`), 97  
 PusherStatsd (class in `uwsgi-conf.options.monitoring_pushers`), 97  
 PusherZabbix (class in `uwsgi-conf.options.monitoring_pushers`), 99  
 Python (class in `uwsgi-conf.options.python`), 155  
 python (`uwsgi-conf.config.Section` attribute), 59  
 PythonSection (class in `uwsgi-conf.presets.nice`), 12
- ## Q
- query (`uwsgi-conf.options.routing.RouteRule.vars` attribute), 134  
 query\_string (`uwsgi-conf.options.routing.RouteRule.subjects` attribute), 135  
 Queue (class in `uwsgi-conf.options.queue`), 112  
 queue (`uwsgi-conf.config.Section` attribute), 59  
 queue (`uwsgi-conf.options.workers_cheapening.Cheapening.algorithms` attribute), 150
- ## R
- rack (`uwsgi-conf.options.routing.Routing.modifiers` attribute), 139  
 rados (`uwsgi-conf.options.routing.Routing.modifiers` attribute), 139  
 raw (`uwsgi-conf.options.networking.Networking.sockets` attribute), 109  
 raw (`uwsgi-conf.options.routing.Routing.modifiers` attribute), 139  
 raw (`uwsgi-conf.options.routing.Routing.routers` attribute), 138  
 read\_task\_file() (`uwsgi-conf.runtime.spooler.Spooler` class method), 56  
 ready() (in module `uwsgi-conf.uwsgi_stub`), 26  
 ready\_fd() (in module `uwsgi-conf.uwsgi_stub`), 26  
 ready\_for\_requests (`uwsgi-conf.runtime.platform._Platform` attribute), 47  
 recv() (in module `uwsgi-conf.uwsgi_stub`), 26  
 redirect (`uwsgi-conf.options.routing.RouteRule.actions` attribute), 136  
 redis (`uwsgi-conf.options.logging.Logging.loggers` attribute), 76  
 register\_alarm() (`uwsgi-conf.options.alarms.Alarms` method), 65  
 register\_cron() (in module `uwsgi-conf.runtime.scheduling`), 51  
 register\_file\_monitor() (in module `uwsgi-conf.runtime.monitoring`), 43  
 register\_handler() (`uwsgi-conf.runtime.signals.Signal` method), 53  
 register\_metric() (`uwsgi-conf.options.monitoring.Monitoring` method), 102  
 register\_module\_alias() (`uwsgi-conf.options.python.Python` method), 157  
 register\_route() (`uwsgi-conf.options.routing.Routing` method), 140  
 register\_route() (`uwsgi-conf.options.routing_routers.RouterTunTap` method), 126  
 register\_rpc() (in module `uwsgi-conf.runtime.rpc`), 49  
 register\_rpc() (in module `uwsgi-conf.uwsgi_stub`), 26  
 register\_signal() (in module `uwsgi-conf.uwsgi_stub`), 26  
 register\_socket() (`uwsgi-conf.options.networking.Networking` method), 111  
 register\_static\_map() (`uwsgi-conf.options.statics.Statics` method), 144  
 register\_stats\_pusher() (`uwsgi-conf.options.monitoring.Monitoring` method), 103  
 register\_timer() (in module `uwsgi-conf.runtime.scheduling`), 50  
 register\_timer\_ms() (in module `uwsgi-conf.runtime.scheduling`), 51  
 register\_timer\_rb() (in module `uwsgi-conf.runtime.scheduling`), 50  
 registered (`uwsgi-conf.runtime.signals.Signal` attribute), 53  
 registry\_signals (in module `uwsgi-conf.runtime.signals`), 53  
 release() (`uwsgi-conf.runtime.locking.Lock` method), 42  
 reload (`uwsgi-conf.options.routing.Routing.modifiers` attribute), 139  
 reload() (in module `uwsgi-conf.runtime.control`), 40  
 reload() (in module `uwsgi-conf.uwsgi_stub`), 27  
 reload\_brutal (`uwsgi-conf.options.routing.Routing.modifiers` attribute), 140  
 remote\_addr (`uwsgi-conf.uwsgi_stub` attribute), 27

<code>conf.options.routing.RouteRule.subjects</code> attribute), 135	<code>remote_logging</code> (uwsgi- <code>conf.options.routing.Routing.modifiers</code> attribute), 140	<code>remote_user</code> (uwsgi- <code>conf.options.routing.RouteRule.subjects</code> attribute), 135	<code>replace_placeholders()</code> (uwsgi- <code>conf.config.Section</code> method), 59	<code>REQ_COUNT_ERR</code> (uwsgi- <code>conf.options.logging.Logging.vars</code> attribute), 81	<code>REQ_COUNT_ERR_READ</code> (uwsgi- <code>conf.options.logging.Logging.vars</code> attribute), 81	<code>REQ_COUNT_ERR_WRITE</code> (uwsgi- <code>conf.options.logging.Logging.vars</code> attribute), 81	<code>REQ_COUNT_VARS_CGI</code> (uwsgi- <code>conf.options.logging.Logging.vars</code> attribute), 80	<code>REQ_HTTP_HOST</code> (uwsgi- <code>conf.options.logging.Logging.vars</code> attribute), 80	<code>REQ_METHOD</code> (uwsgi <sup>conf</sup> .options.logging.Logging.vars attribute), 80	<code>REQ_REFERER</code> (uwsgi- <code>conf.options.logging.Logging.vars</code> attribute), 80	<code>REQ_REMOTE_ADDR</code> (uwsgi- <code>conf.options.logging.Logging.vars</code> attribute), 80	<code>REQ_REMOTE_USER</code> (uwsgi- <code>conf.options.logging.Logging.vars</code> attribute), 80	<code>REQ_SERVER_PROTOCOL</code> (uwsgi- <code>conf.options.logging.Logging.vars</code> attribute), 80	<code>REQ_SIZE_BODY</code> (uwsgi- <code>conf.options.logging.Logging.vars</code> attribute), 80	<code>REQ_START_CTIME</code> (uwsgi- <code>conf.options.logging.Logging.vars</code> attribute), 80	<code>REQ_START_FORMATTED</code> (uwsgi- <code>conf.options.logging.Logging.vars</code> attribute), 80	<code>REQ_START_HUMAN</code> (uwsgi- <code>conf.options.logging.Logging.vars</code> attribute), 80	<code>REQ_START_TS</code> (uwsgi- <code>conf.options.logging.Logging.vars</code> attribute), 80	<code>REQ_START_UNIX_MS</code> (uwsgi- <code>conf.options.logging.Logging.vars</code> attribute), 80	<code>REQ_START_UNIX_US</code> (uwsgi- <code>conf.options.logging.Logging.vars</code> attribute), 80	<code>REQ_URI</code> (uwsgi <sup>conf</sup> .options.logging.Logging.vars attribute), 80	<code>REQ_USER_AGENT</code> (uwsgi- <code>conf.options.logging.Logging.vars</code> attribute), 80	<code>REQUEST</code> (uwsgi <sup>conf</sup> .options.routing.RouteRule.stages attribute), 134	<code>request</code> (uwsgi <sup>conf</sup> .options.routing.RouteRule.vars attribute), 134	<code>request</code> (uwsgi <sup>conf</sup> .runtime.platform._Platform attribute), 46	<code>request_id()</code> (in module uwsgi <sup>conf</sup> .uwsgi_stub), 27	<code>request_uri</code> (uwsgi- <code>conf.options.routing.RouteRule.subjects</code> attribute), 135	<code>REQUEST_URI</code> (uwsgi- <code>conf.options.statics.Static.expiration_criteria</code> attribute), 143	<code>request_var</code> (uwsgi- <code>conf.options.logging.Logging.vars</code> attribute), 82	<code>RESP_COUNT_HEADERS</code> (uwsgi- <code>conf.options.logging.Logging.vars</code> attribute), 81	<code>RESP_SIZE</code> (uwsgi <sup>conf</sup> .options.logging.Logging.vars attribute), 81	<code>RESP_SIZE_BODY</code> (uwsgi- <code>conf.options.logging.Logging.vars</code> attribute), 81	<code>RESP_SIZE_HEADERS</code> (uwsgi- <code>conf.options.logging.Logging.vars</code> attribute), 81	<code>RESP_STATUS</code> (uwsgi- <code>conf.options.logging.Logging.vars</code> attribute), 81	<code>RESP_TIME_MS</code> (uwsgi- <code>conf.options.logging.Logging.vars</code> attribute), 81	<code>RESP_TIME_US</code> (uwsgi- <code>conf.options.logging.Logging.vars</code> attribute), 81	<code>RESPONSE</code> (uwsgi <sup>conf</sup> .options.routing.RouteRule.stages attribute), 134	<code>response</code> (uwsgi <sup>conf</sup> .options.routing.Routing.modifiers attribute), 140	<code>ResultProcessed</code> (class in uwsgi- <code>conf.runtime.spooler</code> ), 56	<code>ResultRescheduled</code> (class in uwsgi-
--	--	--	---	---	--	---	--	---	---	---	---	---	---	---	---	---	---	--	---	---	--	--	---	---	--	---	--	--	---	--	--	--	---	---	--	--	--	---	--	---

- `conf.runtime.spooler`), 56
- `ResultSkipped` (class in `uwsgi2conf.runtime.spooler`), 56
- `rewrite` (`uwsgi2conf.options.routing.RouteRule.actions` attribute), 137
- `route()` (in module `uwsgi2conf.uwsgi_stub`), 27
- `route_external` (`uwsgi2conf.options.routing.RouteRule.actions` attribute), 137
- `route_rule` (`uwsgi2conf.options.routing.Routing` attribute), 138
- `route_uwsgi` (`uwsgi2conf.options.routing.RouteRule.actions` attribute), 137
- `RouteAction` (class in `uwsgi2conf.options.routing_actions`), 113
- `RouterBase` (class in `uwsgi2conf.options.routing_routers`), 118
- `RouterFast` (class in `uwsgi2conf.options.routing_routers`), 122
- `RouterForkPty` (class in `uwsgi2conf.options.routing_routers`), 125
- `RouterHttp` (class in `uwsgi2conf.options.routing_routers`), 119
- `RouterHttps` (class in `uwsgi2conf.options.routing_routers`), 121
- `RouterRaw` (class in `uwsgi2conf.options.routing_routers`), 124
- `RouterSsl` (class in `uwsgi2conf.options.routing_routers`), 121
- `RouterTunTap` (class in `uwsgi2conf.options.routing_routers`), 126
- `RouteRule` (class in `uwsgi2conf.options.routing`), 133
- `RouteRule.actions` (class in `uwsgi2conf.options.routing`), 136
- `RouteRule.stages` (class in `uwsgi2conf.options.routing`), 134
- `RouteRule.subjects` (class in `uwsgi2conf.options.routing`), 134
- `RouteRule.transforms` (class in `uwsgi2conf.options.routing`), 135
- `RouteRule.var_functions` (class in `uwsgi2conf.options.routing`), 134
- `RouteRule.vars` (class in `uwsgi2conf.options.routing`), 133
- `Routing` (class in `uwsgi2conf.options.routing`), 137
- `routing` (`uwsgi2conf.config.Section` attribute), 59
- `Routing.modifiers` (class in `uwsgi2conf.options.routing`), 138
- `Routing.routers` (class in `uwsgi2conf.options.routing`), 138
- `rpc` (`uwsgi2conf.options.routing.Routing.modifiers` attribute), 140
- `rpc()` (in module `uwsgi2conf.uwsgi_stub`), 27
- `rpc_list()` (in module `uwsgi2conf.uwsgi_stub`), 28
- `rrdtool` (`uwsgi2conf.options.monitoring.Monitoring.pushers` attribute), 101
- `rsyslog` (`uwsgi2conf.options.logging.Logging.loggers` attribute), 76
- `run_command_as_worker()` (`uwsgi2conf.options.workers.Workers` method), 152
- `run_command_on_event()` (`uwsgi2conf.options.main_process.MainProcess` method), 88
- `run_command_on_touch()` (`uwsgi2conf.options.main_process.MainProcess` method), 89
- `run_module()` (`uwsgi2conf.options.python.Python` method), 157
- ## S
- `scgi` (`uwsgi2conf.options.networking.Networking.sockets` attribute), 109
- `scrolls()` (in module `uwsgi2conf.uwsgi_stub`), 33
- `Section` (class in `uwsgi2conf.config`), 57
- `Section` (class in `uwsgi2conf.presets.nice`), 10
- `Section.embedded_plugins_presets` (class in `uwsgi2conf.config`), 59
- `Section.vars` (class in `uwsgi2conf.config`), 62
- `send` (`uwsgi2conf.options.routing.RouteRule.actions` attribute), 137
- `send()` (in module `uwsgi2conf.uwsgi_stub`), 28
- `send()` (`uwsgi2conf.runtime.mules.Farm` method), 46
- `send()` (`uwsgi2conf.runtime.mules.Mule` method), 45
- `send()` (`uwsgi2conf.runtime.signals.Signal` method), 54
- `send_message_raw()` (`uwsgi2conf.runtime.spooler.Spooler` class method), 55
- `send_to_spooler()` (in module `uwsgi2conf.uwsgi_stub`), 28
- `SENDFILE` (`uwsgi2conf.options.statics.Statics.transfer_modes` attribute), 143
- `sendfile()` (in module `uwsgi2conf.uwsgi_stub`), 28
- `serve_static` (`uwsgi2conf.options.routing.RouteRule.actions` attribute), 137
- `set()` (`uwsgi2conf.runtime.caching.Cache` method), 39
- `set()` (`uwsgi2conf.runtime.monitoring.Metric` method), 43
- `set_app_args()` (`uwsgi2conf.options.python.Python` method), 156
- `set_autoreload_params()` (`uwsgi2conf.options.python.Python` method), 157
- `set_basic_params()` (`uwsgi2conf.config.Section` method), 60
- `set_basic_params()` (`uwsgi2conf.options.alarms.Alarms` method), 64

`set_basic_params()` (*uwsgi-conf.options.applications.Applications method*), 66  
`set_basic_params()` (*uwsgi-conf.options.caching.Caching method*), 67  
`set_basic_params()` (*uwsgi-conf.options.locks.Lock method*), 72  
`set_basic_params()` (*uwsgi-conf.options.logging.Logging method*), 76  
`set_basic_params()` (*uwsgi-conf.options.main\_process.MainProcess method*), 86  
`set_basic_params()` (*uwsgi-conf.options.master\_process.MasterProcess method*), 89  
`set_basic_params()` (*uwsgi-conf.options.networking.Networking method*), 110  
`set_basic_params()` (*uwsgi-conf.options.python.Python method*), 155  
`set_basic_params()` (*uwsgi-conf.options.routing\_routers.RouterFast method*), 123  
`set_basic_params()` (*uwsgi-conf.options.routing\_routers.RouterForkPty method*), 125  
`set_basic_params()` (*uwsgi-conf.options.routing\_routers.RouterHttp method*), 119  
`set_basic_params()` (*uwsgi-conf.options.routing\_routers.RouterTunTap method*), 126  
`set_basic_params()` (*uwsgi-conf.options.spooler.Spooler method*), 142  
`set_basic_params()` (*uwsgi-conf.options.statics.Statics method*), 143  
`set_basic_params()` (*uwsgi-conf.options.workers.Workers method*), 151  
`set_basic_params()` (*uwsgi-conf.options.workers\_cheapening.AlgoBusyness method*), 149  
`set_basic_params()` (*uwsgi-conf.options.workers\_cheapening.AlgoQueue method*), 148  
`set_basic_params()` (*uwsgi-conf.options.workers\_cheapening.AlgoSpare method*), 148  
`set_basic_params()` (*uwsgi-conf.options.workers\_cheapening.AlgoSpare2 method*), 148  
`set_basic_params()` (*uwsgi-conf.options.workers\_cheapening.Cheapening method*), 150  
`set_bsd_socket_params()` (*uwsgi-conf.options.networking.Networking method*), 111  
`set_client_params()` (*uwsgi-conf.options.subscriptions.Subscriptions method*), 147  
`set_connections_params()` (*uwsgi-conf.options.routing\_routers.RouterFast method*), 123  
`set_connections_params()` (*uwsgi-conf.options.routing\_routers.RouterForkPty method*), 125  
`set_connections_params()` (*uwsgi-conf.options.routing\_routers.RouterHttp method*), 120  
`set_connections_params()` (*uwsgi-conf.options.routing\_routers.RouterRaw method*), 124  
`set_connections_params()` (*uwsgi-conf.options.routing\_routers.RouterSsl method*), 122  
`set_count_auto()` (*uwsgi-conf.options.workers.Workers method*), 152  
`set_emergency_params()` (*uwsgi-conf.options.workers\_cheapening.AlgoBusyness method*), 149  
`set_emperor_command_params()` (*uwsgi-conf.options.empire.Empire method*), 70  
`set_emperor_params()` (*uwsgi-conf.options.empire.Empire method*), 70  
`set_error_page()` (*uwsgi-conf.options.routing.Routing method*), 140  
`set_error_pages()` (*uwsgi-conf.options.routing.Routing method*), 141  
`set_exception_handling_params()` (*uwsgi-conf.options.master\_process.MasterProcess method*), 90  
`set_exit_events()` (*uwsgi-conf.options.master\_process.MasterProcess method*), 90  
`set_fallback()` (*uwsgi-conf.config.Section method*), 61  
`set_file_params()` (*uwsgi-conf.options.logging.Logging method*), 77  
`set_filters()` (*uwsgi-conf.options.logging.Logging method*), 77  
`set_geoip_params()` (*uwsgi-conf.options.routing.Routing method*), 141  
`set_harakiri` (*uwsgi-conf.options.routing.RouteRule.actions attribute*), 137  
`set_harakiri_params()` (*uwsgi-conf.options.workers.Workers method*), 154  
`set_hook()` (*uwsgi-conf.options.main\_process.MainProcess method*), 88

`set_hook_after_request()` (*uwsgi-conf.options.main\_process.MainProcess method*), 88  
`set_hook_touch()` (*uwsgi-conf.options.main\_process.MainProcess method*), 88  
`set_host_name` (*uwsgi-conf.options.main\_process.MainProcess.actions attribute*), 84  
`set_idle_params()` (*uwsgi-conf.options.master\_process.MasterProcess method*), 91  
`set_ipcsem_params()` (*uwsgi-conf.options.locks.Lock method*), 72  
`set_logvar()` (*in module uwsgi<sup>2</sup>conf.uwsgi\_stub*), 29  
`set_manage_params()` (*uwsgi-conf.options.routing\_routers.RouterHttp method*), 120  
`set_master_logging_params()` (*uwsgi-conf.options.logging.Logging method*), 78  
`set_memory_limits()` (*uwsgi-conf.options.workers\_cheapening.Cheapening method*), 150  
`set_memory_params()` (*uwsgi-conf.options.main\_process.MainProcess method*), 87  
`set_metrics_params()` (*uwsgi-conf.options.monitoring.Monitoring method*), 102  
`set_metrics_threshold()` (*uwsgi-conf.options.monitoring.Monitoring method*), 102  
`set_mode_broodlord_params()` (*uwsgi-conf.options.empire.Empire method*), 71  
`set_mode_tyrant_params()` (*uwsgi-conf.options.empire.Empire method*), 71  
`set_mules_params()` (*uwsgi-conf.options.workers.Workers method*), 152  
`set_naming_params()` (*uwsgi-conf.options.main\_process.MainProcess method*), 89  
`set_on_exit_params()` (*uwsgi-conf.options.main\_process.MainProcess method*), 88  
`set_owner_params()` (*uwsgi-conf.options.main\_process.MainProcess method*), 87  
`set_owner_params()` (*uwsgi-conf.options.routing\_routers.RouterFast method*), 124  
`set_owner_params()` (*uwsgi-conf.options.routing\_routers.RouterHttp method*), 120  
`set_paths_caching_params()` (*uwsgi-conf.options.statics.Statics method*), 144  
`set_period()` (*uwsgi<sup>2</sup>conf.runtime.spooler.Spooler class method*), 56  
`set_pid_file()` (*uwsgi-conf.options.main\_process.MainProcess method*), 89  
`set_placeholder()` (*uwsgi<sup>2</sup>conf.config.Section method*), 61  
`set_plugins_params()` (*uwsgi<sup>2</sup>conf.config.Section method*), 61  
`set_postbuffering_params()` (*uwsgi-conf.options.routing\_routers.RouterFast method*), 124  
`set_process_name()` (*in module uwsgi<sup>2</sup>conf.runtime.control*), 40  
`set_reload_on_exception_params()` (*uwsgi-conf.options.workers.Workers method*), 154  
`set_reload_params()` (*uwsgi-conf.options.master\_process.MasterProcess method*), 91  
`set_reload_params()` (*uwsgi-conf.options.workers.Workers method*), 153  
`set_requests_filters()` (*uwsgi-conf.options.logging.Logging method*), 78  
`set_resubscription_params()` (*uwsgi-conf.options.routing\_routers.RouterFast method*), 123  
`set_runtime_dir()` (*uwsgi<sup>2</sup>conf.config.Section method*), 60  
`set_script_file` (*uwsgi-conf.options.routing.RouteRule.actions attribute*), 137  
`set_server_params()` (*uwsgi-conf.options.subscriptions.Subscriptions method*), 146  
`set_server_verification_params()` (*uwsgi<sup>2</sup>conf.options.subscriptions.Subscriptions method*), 146  
`set_sni_dir_params()` (*uwsgi-conf.options.networking.Networking method*), 112  
`set_sni_params()` (*uwsgi-conf.options.networking.Networking method*), 111  
`set_socket_params()` (*uwsgi-conf.options.networking.Networking method*), 110  
`set_spooler_frequency()` (*in module uwsgi<sup>2</sup>conf.uwsgi\_stub*), 29  
`set_ssl_params()` (*uwsgi-conf.options.networking.Networking method*), 111  
`set_stats_params()` (*uwsgi-conf.options.monitoring.Monitoring method*),

- 103
- set\_thread\_params() (uwsgi-conf.options.workers.Workers method), 152
- set\_throttle\_params() (uwsgi-conf.options.empire.Empire method), 71
- set\_tolerance\_params() (uwsgi-conf.options.empire.Empire method), 71
- set\_unix\_socket\_params() (uwsgi-conf.options.networking.Networking method), 110
- set\_user\_harakiri() (in module uwsgi-conf.uwsgi\_stub), 29
- set\_uwsgi\_process\_name (uwsgi-conf.options.routing.RouteRule.actions attribute), 137
- set\_var\_document\_root (uwsgi-conf.options.routing.RouteRule.actions attribute), 137
- set\_var\_path\_info (uwsgi-conf.options.routing.RouteRule.actions attribute), 137
- set\_var\_remote\_addr (uwsgi-conf.options.routing.RouteRule.actions attribute), 137
- set\_var\_remote\_user (uwsgi-conf.options.routing.RouteRule.actions attribute), 137
- set\_var\_request\_method (uwsgi-conf.options.routing.RouteRule.actions attribute), 137
- set\_var\_request\_uri (uwsgi-conf.options.routing.RouteRule.actions attribute), 137
- set\_var\_script\_name (uwsgi-conf.options.routing.RouteRule.actions attribute), 137
- set\_var\_uwsgi\_appid (uwsgi-conf.options.routing.RouteRule.actions attribute), 137
- set\_var\_uwsgi\_home (uwsgi-conf.options.routing.RouteRule.actions attribute), 137
- set\_var\_uwsgi\_scheme (uwsgi-conf.options.routing.RouteRule.actions attribute), 137
- set\_vassals\_wrapper\_params() (uwsgi-conf.options.empire.Empire method), 71
- set\_warning\_message() (in module uwsgi-conf.uwsgi\_stub), 29
- set\_window\_params() (uwsgi-conf.options.routing\_routers.RouterForkPty method), 125
- set\_wsgi\_params() (uwsgi-conf.options.python.Python method), 156
- set\_zerg\_client\_params() (uwsgi-conf.options.workers.Workers method), 155
- set\_zerg\_server\_params() (uwsgi-conf.options.workers.Workers method), 154
- setprocname() (in module uwsgi-conf.uwsgi\_stub), 29
- shared (uwsgi-conf.options.networking.Networking.sockets attribute), 109
- Signal (class in uwsgi-conf.runtime.signals), 53
- signal (uwsgi-conf.options.alarms.Alarms.alarm\_types attribute), 64
- signal (uwsgi-conf.options.routing.RouteRule.actions attribute), 137
- signal (uwsgi-conf.options.routing.Routing.modifiers attribute), 140
- signal() (in module uwsgi-conf.uwsgi\_stub), 29
- signal\_received() (in module uwsgi-conf.uwsgi\_stub), 30
- signal\_registered() (in module uwsgi-conf.uwsgi\_stub), 30
- signal\_wait() (in module uwsgi-conf.uwsgi\_stub), 30
- SignalDescription (class in uwsgi-conf.runtime.signals), 53
- SIZE\_PACKET\_UWSGI (uwsgi-conf.options.logging.Logging.vars attribute), 82
- sleep() (in module uwsgi-conf.runtime.asynced), 37
- snmp (uwsgi-conf.options.routing.Routing.modifiers attribute), 140
- Socket (class in uwsgi-conf.options.networking\_sockets), 104
- socket (uwsgi-conf.options.logging.Logging.loggers attribute), 76
- socket (uwsgi-conf.options.monitoring.Monitoring.pushers attribute), 101
- SocketDefault (class in uwsgi-conf.options.networking\_sockets), 104
- SocketFastcgi (class in uwsgi-conf.options.networking\_sockets), 107
- SocketHttp (class in uwsgi-conf.options.networking\_sockets), 104
- SocketHttps (class in uwsgi-conf.options.networking\_sockets), 105
- SocketRaw (class in uwsgi-conf.options.networking\_sockets), 107
- sockets (in module uwsgi-conf.uwsgi\_stub), 15
- SocketScgi (class in uwsgi-conf.options.networking\_sockets), 107
- SocketShared (class in uwsgi-conf.options.networking\_sockets), 108
- SocketUdp (class in uwsgi-conf.options.networking\_sockets), 106
- SocketUwsgi (class in uwsgi-

- `conf.options.networking_sockets`), 105
- SocketUwsgis (class in uwsgi-  
`conf.options.networking_sockets`), 106
- SocketZeromq (class in uwsgi-  
`conf.options.networking_sockets`), 108
- spare (`uwsgiconf.options.workers_cheapening.CheapeningSubPools`  
attribute), 150
- spare2 (`uwsgiconf.options.workers_cheapening.CheapeningSubPools`  
attribute), 150
- spool () (in module `uwsgiconf.uwsgi_stub`), 30
- SPOOL\_IGNORE (in module `uwsgiconf.uwsgi_stub`), 14
- SPOOL\_OK (in module `uwsgiconf.uwsgi_stub`), 14
- SPOOL\_RETRY (in module `uwsgiconf.uwsgi_stub`), 14
- Spooler (class in `uwsgiconf.options.spooler`), 141
- Spooler (class in `uwsgiconf.runtime.spooler`), 55
- spooler (in module `uwsgiconf.uwsgi_stub`), 15
- spooler (`uwsgiconf.config.Section` attribute), 59
- spooler (`uwsgiconf.options.routing.Routing.modifiers`  
attribute), 140
- spooler\_get\_task () (in module `uwsgi-  
conf.uwsgi_stub`), 31
- spooler\_jobs () (in module `uwsgiconf.uwsgi_stub`),  
31
- spooler\_pid () (in module `uwsgiconf.uwsgi_stub`),  
31
- spooler\_pids () (in module `uwsgiconf.uwsgi_stub`),  
31
- spooler\_task\_types (in module `uwsgi-  
conf.runtime.spooler`), 57
- SpoolerFunctionCallTask (class in `uwsgi-  
conf.runtime.spooler`), 57
- SpoolerTask (class in `uwsgiconf.runtime.spooler`), 56
- ssl (`uwsgiconf.options.routing.Routing.routers` at-  
tribute), 138
- start\_response (in module `uwsgiconf.uwsgi_stub`),  
15
- started\_on (in module `uwsgiconf.uwsgi_stub`), 15
- started\_on (`uwsgiconf.runtime.platform._Platform`  
attribute), 47
- startswith () (uwsgi-  
`conf.options.routing_subjects.SubjectCustom`  
method), 131
- Statics (class in `uwsgiconf.options.statics`), 143
- statics (`uwsgiconf.config.Section` attribute), 59
- Statics.expiration\_criteria (class in `uwsgi-  
conf.options.statics`), 143
- Statics.transfer\_modes (class in `uwsgi-  
conf.options.statics`), 143
- statsd (`uwsgiconf.options.monitoring.Monitoring.pushers`  
attribute), 101
- status (`uwsgiconf.options.routing.RouteRule.subjects`  
attribute), 135
- stop () (in module `uwsgiconf.runtime.control`), 40
- stop () (in module `uwsgiconf.uwsgi_stub`), 31
- SUB\_DEFAULT (uwsgi-  
`conf.options.routing_modifiers.ModifierRpc`  
attribute), 130
- SUB\_DELETE (`uwsgiconf.options.routing_modifiers.ModifierCache`  
attribute), 130
- SUB\_DELETE\_CONDITIONS (`uwsgiconf.options.routing_modifiers.ModifierRpc`  
attribute), 130
- SUB\_DEFAULTS\_BASED (uwsgi-  
`conf.options.routing_modifiers.ModifierCache`  
attribute), 130
- SUB\_DUMP (`uwsgiconf.options.routing_modifiers.ModifierCache`  
attribute), 130
- SUB\_GET (`uwsgiconf.options.routing_modifiers.ModifierCache`  
attribute), 129
- SUB\_JSONRPC (uwsgi-  
`conf.options.routing_modifiers.ModifierRpc`  
attribute), 130
- SUB\_MAGIC (`uwsgiconf.options.routing_modifiers.ModifierCache`  
attribute), 130
- SUB\_PING (`uwsgiconf.options.routing_modifiers.ModifierPing`  
attribute), 129
- SUB\_PONG (`uwsgiconf.options.routing_modifiers.ModifierPing`  
attribute), 129
- SUB\_RAW (`uwsgiconf.options.routing_modifiers.ModifierRpc`  
attribute), 130
- SUB\_RING (`uwsgiconf.options.routing_modifiers.ModifierJvm`  
attribute), 128
- SUB\_SET (`uwsgiconf.options.routing_modifiers.ModifierCache`  
attribute), 129
- SUB\_STREAM (`uwsgiconf.options.routing_modifiers.ModifierCache`  
attribute), 130
- SUB\_USE\_PATH\_INFO (uwsgi-  
`conf.options.routing_modifiers.ModifierRpc`  
attribute), 130
- SUB\_XMLRPC (`uwsgiconf.options.routing_modifiers.ModifierRpc`  
attribute), 130
- SubjectCustom (class in `uwsgi-  
conf.options.routing_subjects`), 131
- SubjectHttpHost (class in `uwsgi-  
conf.options.routing_subjects`), 132
- SubjectHttpReferer (class in `uwsgi-  
conf.options.routing_subjects`), 132
- SubjectHttpUserAgent (class in `uwsgi-  
conf.options.routing_subjects`), 132
- SubjectPathInfo (class in `uwsgi-  
conf.options.routing_subjects`), 132
- SubjectQueryString (class in `uwsgi-  
conf.options.routing_subjects`), 132
- SubjectRemoteAddr (class in `uwsgi-  
conf.options.routing_subjects`), 132
- SubjectRemoteUser (class in `uwsgi-  
conf.options.routing_subjects`), 132
- SubjectRequestUri (class in `uwsgi-  
conf.options.routing_subjects`), 132



SubjectStatus (class in `uwsgi-conf.options.routing_subjects`), 132  
 subscribe() (`uwsgi-conf.options.subscriptions.Subscriptions` method), 147  
 subscription (`uwsgi-conf.options.routing.Routing.modifiers` attribute), 140  
 Subscriptions (class in `uwsgi-conf.options.subscriptions`), 145  
 subscriptions (`uwsgi-conf.config.Section` attribute), 59  
 Subscriptions.algorithms (class in `uwsgi-conf.options.subscriptions`), 145  
 suffix (`uwsgi-conf.options.logging.Logging.encoders` attribute), 76  
 sum (`uwsgi-conf.options.monitoring.Monitoring.collectors` attribute), 101  
 suspend() (in module `uwsgi-conf.runtime.asynced`), 37  
 suspend() (in module `uwsgi-conf.uwsgi_stub`), 31  
 switch\_into\_lazy\_mode() (`uwsgi-conf.options.applications.Applications` method), 67  
 SymbolsImporter (in module `uwsgi-conf.uwsgi_stub`), 14  
 SymbolsZipImporter (in module `uwsgi-conf.uwsgi_stub`), 14  
 symcall (`uwsgi-conf.options.routing.Routing.modifiers` attribute), 140  
 syslog (`uwsgi-conf.options.logging.Logging.loggers` attribute), 76

## T

target (`uwsgi-conf.runtime.signals.SignalDescription` attribute), 53  
 task() (`uwsgi-conf.runtime.spooler.Spooler` method), 55  
 TaskResult (class in `uwsgi-conf.runtime.spooler`), 56  
 template (`uwsgi-conf.options.routing.RouteRule.transforms` attribute), 135  
 threads\_enabled (`uwsgi-conf.runtime.platform._Platform` attribute), 47  
 TIME (`uwsgi-conf.options.logging_encoders.EncoderFormatvars` attribute), 75  
 time (`uwsgi-conf.options.routing.RouteRule.vars` attribute), 134  
 TIME\_MS (`uwsgi-conf.options.logging_encoders.EncoderFormatvars` attribute), 75  
 TIME\_UNIX (`uwsgi-conf.options.logging.Logging.vars` attribute), 81  
 TIMESTAMP\_STARTUP\_MS (`uwsgi-conf.config.Section.vars` attribute), 62  
 TIMESTAMP\_STARTUP\_S (`uwsgi-conf.config.Section.vars` attribute), 62  
 to\_file (`uwsgi-conf.options.routing.RouteRule.transforms` attribute), 135  
 tofile() (`uwsgi-conf.config.Configuration` method), 63  
 total\_count (`uwsgi-conf.runtime.request._Request` attribute), 48  
 total\_requests() (in module `uwsgi-conf.uwsgi_stub`), 31  
 tuntap (`uwsgi-conf.options.routing.Routing.routers` attribute), 138

## U

udp (`uwsgi-conf.options.networking.Networking.sockets` attribute), 109  
 unbit (in module `uwsgi-conf.uwsgi_stub`), 15  
 unlink (`uwsgi-conf.options.main_process.MainProcess.actions` attribute), 84  
 unlock() (in module `uwsgi-conf.uwsgi_stub`), 31  
 upper (`uwsgi-conf.options.routing.RouteRule.transforms` attribute), 135  
 upper (`uwsgi-conf.options.routing.RouteRule.var_functions` attribute), 134  
 use\_router() (`uwsgi-conf.options.routing.Routing` method), 140  
 USER\_ID (`uwsgi-conf.config.Section.vars` attribute), 62  
 USER\_NAME (`uwsgi-conf.config.Section.vars` attribute), 62  
 uwsgi (`uwsgi-conf.options.networking.Networking.sockets` attribute), 109  
 uwsgi (`uwsgi-conf.options.routing.RouteRule.vars` attribute), 134  
 uwsgi-conf.config (module), 57  
 uwsgi-conf.formatters (module), 158  
 uwsgi-conf.options.alarms\_types (module), 64  
 uwsgi-conf.options.alarms (module), 64  
 uwsgi-conf.options.applications (module), 66  
 uwsgi-conf.options.caching (module), 67  
 uwsgi-conf.options.empire (module), 70  
 uwsgi-conf.options.locks (module), 72  
 uwsgi-conf.options.logging (module), 75  
 uwsgi-conf.options.logging\_encoders (module), 75  
 uwsgi-conf.options.logging\_loggers (module), 73  
 uwsgi-conf.options.main\_process (module), 84  
 uwsgi-conf.options.main\_process\_actions (module), 82  
 uwsgi-conf.options.master\_process (module), 89  
 uwsgi-conf.options.monitoring (module), 100  
 uwsgi-conf.options.monitoring\_collectors (module), 99

- uwsgi<sup>2</sup>conf.options.monitoring\_metric\_type (module), 93
  - uwsgi<sup>2</sup>conf.options.monitoring\_pushers (module), 97
  - uwsgi<sup>2</sup>conf.options.networking (module), 109
  - uwsgi<sup>2</sup>conf.options.networking\_sockets (module), 104
  - uwsgi<sup>2</sup>conf.options.python (module), 155
  - uwsgi<sup>2</sup>conf.options.queue (module), 112
  - uwsgi<sup>2</sup>conf.options.routing (module), 133
  - uwsgi<sup>2</sup>conf.options.routing\_actions (module), 113
  - uwsgi<sup>2</sup>conf.options.routing\_modifiers (module), 127
  - uwsgi<sup>2</sup>conf.options.routing\_routers (module), 118
  - uwsgi<sup>2</sup>conf.options.routing\_subjects (module), 131
  - uwsgi<sup>2</sup>conf.options.routing\_vars (module), 132
  - uwsgi<sup>2</sup>conf.options.spooler (module), 141
  - uwsgi<sup>2</sup>conf.options.statics (module), 143
  - uwsgi<sup>2</sup>conf.options.subscriptions (module), 145
  - uwsgi<sup>2</sup>conf.options.subscriptions\_algos (module), 145
  - uwsgi<sup>2</sup>conf.options.workers (module), 150
  - uwsgi<sup>2</sup>conf.options.workers\_cheapening (module), 148
  - uwsgi<sup>2</sup>conf.presets.empire (module), 9
  - uwsgi<sup>2</sup>conf.presets.nice (module), 10
  - uwsgi<sup>2</sup>conf.runtime.alarms (module), 37
  - uwsgi<sup>2</sup>conf.runtime.asynced (module), 37
  - uwsgi<sup>2</sup>conf.runtime.caching (module), 38
  - uwsgi<sup>2</sup>conf.runtime.control (module), 40
  - uwsgi<sup>2</sup>conf.runtime.locking (module), 41
  - uwsgi<sup>2</sup>conf.runtime.logging (module), 42
  - uwsgi<sup>2</sup>conf.runtime.monitoring (module), 43
  - uwsgi<sup>2</sup>conf.runtime.mules (module), 44
  - uwsgi<sup>2</sup>conf.runtime.rpc (module), 49
  - uwsgi<sup>2</sup>conf.runtime.scheduling (module), 50
  - uwsgi<sup>2</sup>conf.runtime.signals (module), 53
  - uwsgi<sup>2</sup>conf.runtime.spooler (module), 55
  - uwsgi<sup>2</sup>conf.uwsgi\_stub (module), 14
  - uwsgis (uwsgi<sup>2</sup>conf.options.networking.Networking.sockets attribute), 109
- V**
- v8 (uwsgi<sup>2</sup>conf.options.routing.Routing.modifiers attribute), 140
  - value (uwsgi<sup>2</sup>conf.runtime.monitoring.Metric attribute), 43
  - Var (class in uwsgi<sup>2</sup>conf.options.logging), 75
  - Var (class in uwsgi<sup>2</sup>conf.options.routing\_vars), 132
  - VarCookie (class in uwsgi<sup>2</sup>conf.options.routing\_vars), 132
  - VarGeoip (class in uwsgi<sup>2</sup>conf.options.routing\_vars), 132
  - VarHttpptime (class in uwsgi<sup>2</sup>conf.options.routing\_vars), 133
  - variable\_get () (in module uwsgi<sup>2</sup>conf.runtime.logging), 42
  - variable\_set () (in module uwsgi<sup>2</sup>conf.runtime.logging), 42
  - VarMetric (class in uwsgi<sup>2</sup>conf.options.logging), 75
  - VarMetric (class in uwsgi<sup>2</sup>conf.options.routing\_vars), 132
  - VarQuery (class in uwsgi<sup>2</sup>conf.options.routing\_vars), 132
  - VarRequest (class in uwsgi<sup>2</sup>conf.options.routing\_vars), 132
  - VarRequestVar (class in uwsgi<sup>2</sup>conf.options.logging), 75
  - vars\_city (uwsgi<sup>2</sup>conf.options.routing\_vars.VarGeoip attribute), 132
  - vars\_country (uwsgi<sup>2</sup>conf.options.routing\_vars.VarGeoip attribute), 132
  - VarTime (class in uwsgi<sup>2</sup>conf.options.routing\_vars), 133
  - VarUwsgi (class in uwsgi<sup>2</sup>conf.options.routing\_vars), 132
  - VASSAL\_CONFIG\_CHANGE\_POST (uwsgi<sup>2</sup>conf.options.main\_process.MainProcess.phases attribute), 85
  - VASSAL\_ON\_DEMAND\_IN (uwsgi<sup>2</sup>conf.options.main\_process.MainProcess.phases attribute), 85
  - VASSAL\_PRIV\_DRP\_PRE (uwsgi<sup>2</sup>conf.options.main\_process.MainProcess.phases attribute), 85
  - VASSAL\_SET\_NAMESPACE (uwsgi<sup>2</sup>conf.options.main\_process.MainProcess.phases attribute), 85
  - VASSAL\_START\_IN (uwsgi<sup>2</sup>conf.options.main\_process.MainProcess.phases attribute), 86
  - VASSAL\_START\_POST (uwsgi<sup>2</sup>conf.options.main\_process.MainProcess.phases attribute), 86
  - VASSAL\_START\_PRE (uwsgi<sup>2</sup>conf.options.main\_process.MainProcess.phases attribute), 85
  - version (in module uwsgi<sup>2</sup>conf.uwsgi\_stub), 15
  - VERSION (uwsgi<sup>2</sup>conf.config.Section.vars attribute), 62
  - version\_info (in module uwsgi<sup>2</sup>conf.uwsgi\_stub), 15
- W**
- wait () (uwsgi<sup>2</sup>conf.runtime.signals.Signal method), 54

`wait_fd_read()` (in module `uwsgi2conf.uwsgi_stub`), 32  
`wait_fd_write()` (in module `uwsgi2conf.uwsgi_stub`), 32  
`wait_for_fd_read()` (in module `uwsgi2conf.runtime.asynced`), 37  
`wait_for_fd_write()` (in module `uwsgi2conf.runtime.asynced`), 38  
`websocket_handshake()` (in module `uwsgi2conf.uwsgi_stub`), 32  
`websocket_recv()` (in module `uwsgi2conf.uwsgi_stub`), 32  
`websocket_recv_nb()` (in module `uwsgi2conf.uwsgi_stub`), 33  
`websocket_send()` (in module `uwsgi2conf.uwsgi_stub`), 33  
`websocket_send_binary()` (in module `uwsgi2conf.uwsgi_stub`), 33  
`weighted_least_reference_count` (`uwsgi2conf.options.subscriptions.Subscriptions.algorithms` attribute), 145  
`weighted_round_robin` (`uwsgi2conf.options.subscriptions.Subscriptions.algorithms` attribute), 146  
`WeightedLeastReferenceCount` (class in `uwsgi2conf.options.subscriptions_algos`), 145  
`WeightedRoundRobin` (class in `uwsgi2conf.options.subscriptions_algos`), 145  
`WORKER_ACCEPTING_PRE_EACH` (`uwsgi2conf.options.main_process.MainProcess.phases` attribute), 86  
`WORKER_ACCEPTING_PRE_EACH_ONCE` (`uwsgi2conf.options.main_process.MainProcess.phases` attribute), 86  
`WORKER_ACCEPTING_PRE_FIRST` (`uwsgi2conf.options.main_process.MainProcess.phases` attribute), 86  
`WORKER_ACCEPTING_PRE_FIRST_ONCE` (`uwsgi2conf.options.main_process.MainProcess.phases` attribute), 86  
`WORKER_ID` (`uwsgi2conf.options.logging.Logging.vars` attribute), 81  
`worker_id` (`uwsgi2conf.runtime.platform._Platform` attribute), 47  
`worker_id()` (in module `uwsgi2conf.uwsgi_stub`), 33  
`WORKER_PID` (`uwsgi2conf.options.logging.Logging.vars` attribute), 81  
`Workers` (class in `uwsgi2conf.options.workers`), 151  
`workers` (`uwsgi2conf.config.Section` attribute), 59  
`workers()` (in module `uwsgi2conf.uwsgi_stub`), 33  
`workers_count` (`uwsgi2conf.runtime.platform._Platform` attribute), 47  
`workers_info` (`uwsgi2conf.runtime.platform._Platform` attribute), 47  
`wsgi` (`uwsgi2conf.options.routing.Routing.modifiers` attribute), 140  
**X**  
`xmpp` (`uwsgi2conf.options.alarms.Alarms.alarm_types` attribute), 64  
`xslt` (`uwsgi2conf.options.routing.Routing.modifiers` attribute), 140  
**Z**  
`zabbix` (`uwsgi2conf.options.monitoring.Monitoring.pushers` attribute), 102  
`zeromq` (`uwsgi2conf.options.logging.Logging.loggers` attribute), 76  
`zeromq` (`uwsgi2conf.options.networking.Networking.sockets` attribute), 109  
`ZipImporter` (in module `uwsgi2conf.uwsgi_stub`), 14