
Twiggy Documentation

Release 0.4.7

Peter Fein

Nov 20, 2017

Contents

1	Who, What, When, Where	1
2	Why Twigg Should Be Your New Logger	3
3	Documentation	5
3.1	Logging Messages	5
3.2	Configuring Output	8
3.3	Reference Guide	10
3.4	API Reference	15
3.5	Testing	24
3.6	Glossary	24
3.7	Changelog	25
3.8	Contributors	26
	Python Module Index	27

Who, What, When, Where

Twiggy is a more Pythonic logger. It aims to be easy to setup:

```
>>> from twiggpy import quick_setup
>>> quick_setup()
```

And fun to use!

```
>>> from twiggpy import log
>>> log.name('frank').fields(number=42).info("hello {who}, it's a {0} day", 'sunny',
↳who='world')
INFO:frank:number=42|hello world, it's a sunny day
```

author Peter Fein

email pete@wearpants.org

twitter @petecode

homepage <http://twiggy.readthedocs.io/en/latest/>

hosting <https://github.com/wearpants/twiggy>

IRC <irc://irc.freenode.net/#twiggy>

license BSD

Python 2.6, 2.7

Twiggy was born at Pycon 2010 after I whined about the standard library's logging and Jesse Noller "invited" me to do something about it.

Install straight with distutils from the [Cheeseshop](#) or:

```
pip install Twiggpy
easy_install -U Twiggpy
```

Get the latest version:

```
git clone https://github.com/wearpants/twiggy.git
```

Why Twiggy Should Be Your New Logger

You should use Twiggy because it is awesome. For more information, [see this blog post](#).

Warning: Twiggy works great, but is not rock solid (yet); do not use for nuclear power plants, spaceships or mortgage derivatives trading (not that it'd matter).

3.1 Logging Messages

This part describes how user code can log messages with twiggy.

To get started quickly, use `quick_setup()`:

```
>>> import twiggy
>>> twiggy.quick_setup()
```

See also:

Full details on *Configuring Output*.

3.1.1 The Magic log

The main interface is the the magic `log`.

```
>>> from twiggy import log
>>> log
<twiggy.logger.Logger object at 0x...>
```

It works out of the box, using typical *levels*. Arbitrary levels are *not* supported. Note that when logging, you never need to refer to any level object; just use the methods on the log.

```
>>> log.debug('You may not care')
DEBUG|You may not care
>>> log.error('OMFG! Pants on fire!')
ERROR|OMFG! Pants on fire!
```

The log can handle messages in several styles of *format strings*, defaulting to *new-style*.

```
>>> log.info('I wear {0} on my {where}', 'pants', where='legs')
INFO|I wear pants on my legs
```

You can name your loggers.

```
>>> mylog = log.name('alfredo')
>>> mylog.debug('hello')
DEBUG:alfredo|hello
```

3.1.2 Better output

Twiggy's default output strives to be user-friendly and to avoid pet peeves.

Newlines are suppressed by default; that can be turned off per-message.

```
>>> log.info('user\ninput\nnannoys\nus')
INFO|user\ninput\nnannoys\nus
>>> log.options(suppress_newlines=False).info('we\ndeal')
INFO|we
deal
```

Exceptions are prefixed by TRACE. By default, *tracing* will use the current exception, but you can also pass an `exc_info` tuple.

```
>>> try:
...     1/0
... except:
...     log.trace('error').warning('oh noes')
WARNING|oh noes
TRACE Traceback (most recent call last):
TRACE   File "<doctest better-output[...]>", line 2, in <module>
TRACE     1/0
TRACE ZeroDivisionError: division by zero
```

See also:

How to fold exceptions to a single line

3.1.3 Structured Logging

I like this method chaining style a lot.

```
>>> log.name('benito').info('hi there')
INFO:benito|hi there
```

It makes *structured logging* easy. In the past, fielded data was stuffed in the text of your message:

```
>>> log.info('Going for a walk. path: {0} roads: {1}', "less traveled", 42)
INFO|Going for a walk. path: less traveled roads: 42
```

Instead, you can use *fields()* to add arbitrary key-value pairs. Output is easily parseable.

```
>>> log.fields(path="less traveled", roads=42).info('Going for a walk')
INFO:path=less traveled:roads=42|Going for a walk
```

The *struct()* is a short cut for *only* logging fields. This is great for runtime statistics gathering.

```
>>> log.struct(paths=42, dolphins='thankful')
INFO:dolphins=thankful:paths=42|
```

3.1.4 Partial Binding

Each call to `fields()` or `options()` creates a new, independent log instance that inherits all of the data of the parent. This incremental binding can be useful for webapps.

```
>>> ## an application-level log
... webapp_log = log.name("myblog")
>>> ## a log for the individual request
... current_request_log = webapp_log.fields(request_id='12345')
>>> current_request_log.fields(rows=100, user='frank').info('frobnicating database')
INFO:myblog:request_id=12345:rows=100:user=frank|frobnicating database
>>> current_request_log.fields(bytes=5678).info('sending page over tubes')
INFO:myblog:bytes=5678:request_id=12345|sending page over tubes
>>> ## a log for a different request
... another_log = webapp_log.fields(request_id='67890')
>>> another_log.debug('Client connected')
DEBUG:myblog:request_id=67890|Client connected
```

Chained style is awesome. It allows you to create complex yet parsable log messages in a concise way.

```
>>> log.name('donjuan').fields(pants='sexy').info("hello, {who} want to {what}?", who=
↳ 'ladies', what='dance')
INFO:donjuan:pants=sexy|hello, ladies want to dance?
```

3.1.5 Sample Output

Routed to a *file*, the above produces the following:

```
2010-03-28T14:23:34Z:DEBUG:You may not care
2010-03-28T14:23:34Z:ERROR:OMFG! Pants on fire!
2010-03-28T14:23:34Z:INFO:I like bikes
2010-03-28T14:23:34Z:INFO:I wear pants on my legs
2010-03-28T14:23:34Z:DEBUG:alfredo:hello
2010-03-28T14:23:34Z:INFO:user\ninput\nnannoys\nus
2010-03-28T14:23:34Z:INFO:we
deal
2010-03-28T14:23:34Z:WARNING:oh noes
TRACE Traceback (most recent call last):
TRACE   File "<doctest better-output[...]>", line 35, in <module>
TRACE ZeroDivisionError: integer division or modulo by zero
2010-03-28T14:23:34Z:INFO:benito:hi there
2010-03-28T14:23:34Z:INFO:Going for a walk. path: less traveled roads: 42
2010-03-28T14:23:34Z:INFO:path=less traveled:roads=42:Going for a walk
2010-03-28T14:23:34Z:INFO:dolphins=thankful:paths=42:
2010-03-28T14:23:34Z:INFO:myblog:request_id=12345:rows=100:user=frank:frobnicating_
↳ database
2010-03-28T14:23:34Z:INFO:myblog:bytes=5678:request_id=12345:sending page over tubes
2010-03-28T14:23:34Z:INFO:myblog:request_id=67890:Client connected
2010-03-28T14:23:34Z:INFO:donjuan:pants=sexy:hello, ladies want to dance?
2010-03-28T14:23:34Z:INFO:myblog:request_id=12345:rows=100:user=frank:frobnicating_
↳ database
2010-03-28T14:23:34Z:INFO:myblog:bytes=5678:request_id=12345:sending page over tubes
2010-03-28T14:23:34Z:DEBUG:myblog:request_id=67890:Client connected
```

3.2 Configuring Output

This part discusses how to configure twiggy's output of messages. You should do this once, near the start of your application's `__main__`.

3.2.1 Quick Setup

To quickly configure output, use the `quick_setup` function. Quick setup is limited to sending all messages to a file or `sys.stderr`. A timestamp will be prefixed when logging to a file.

```
twiggy.quick_setup(min_level=<LogLevel DEBUG>, file=None, msg_buffer=0)
    Quickly set up emitters.
```

Parameters

- **min_level** (`LogLevel`) – lowest message level to cause output
- **file** (`string`) – filename to log to, or `sys.stdout`, or `sys.stderr`. None means standard error.
- **msg_buffer** (`int`) – number of messages to buffer, see `outputs.AsyncOutput.msg_buffer`

3.2.2 twiggy_setup.py

Twiggy's output side features modern, loosely coupled design.

By convention, your configuration lives in a file in your application called `twiggy_setup.py`, in a function called `twiggy_setup()`. You can of course put your configuration elsewhere, but using a separate module makes integration with configuration management systems easy. You should import and run `twiggy_setup` near the top of your application. It's particularly important to set up Twiggy *before spawning new processes*.

A `twiggy_setup` function should create outputs and use the `add_emitters()` convenience function to link those outputs to the `log`.

```
from twiggy import add_emitters, outputs, levels, filters, formats, emitters # import_
↳* is also ok
def twiggy_setup():
    alice_output = outputs.FileOutput("alice.log", format=formats.line_format)
    bob_output = outputs.FileOutput("bob.log", format=formats.line_format)

    add_emitters(
        # (name, min_level, filter, output),
        ("alice", levels.DEBUG, None, alice_output),
        ("betty", levels.INFO, filters.names("betty"), bob_output),
        ("brian.*", levels.DEBUG, filters.glob_names("brian.*"), bob_output),
    )

# near the top of your __main__
twiggy_setup()
```

`add_emitters()` populates the `emitters` dictionary:

```
>>> sorted(emitters.keys())
['alice', 'betty', 'brian.*']
```

In this example, we create two log destinations: `alice.log` and `bob.log`. `alice` will receive all messages, and `bob` will receive two sets of messages:

- messages with the `name` field equal to `betty` and `level >= INFO`
- messages with the `name` field glob-matching `brian.*`

`Emitters` can be removed by deleting them from this dict. `filter` and `min_level` may be modified during the running of the application, but outputs *cannot* be changed. Instead, remove the emitter and re-add it.

```
>>> # bump level
... emitters['alice'].min_level = levels.WARNING
>>> # change filter
... emitters['alice'].filter = filters.names('alice', 'andy')
>>> # remove entirely
... del emitters['alice']
```

We'll examine the various parts in more detail.

3.2.3 Outputs

Outputs are the destinations to which log messages are written (files, databases, etc.). Several *implementations* are provided. Once created, outputs cannot be modified. Each output has an associated *format*.

Asynchronous Logging

Many outputs can be configured to use a separate, dedicated process to log messages. This is known as *asynchronous logging* and is enabled with the `msg_buffer` argument. Asynchronous mode dramatically reduces the cost of logging, as expensive formatting and writing operations are moved out of the main thread of control.

3.2.4 Formats

Formats transform a log message into a form that can be written by an output. The result of formatting is output dependent; for example, an output that posts to an HTTP server may take a format that provides JSON, whereas an output that writes to a file may produce text.

Line-oriented formatting

`LineFormat` formats messages for text-oriented outputs such as a file or standard error. It uses a `ConversionTable` to stringify the arbitrary fields in a message. To customize, copy the default `line_format` and modify:

```
# in your twiggy_setup
import copy
my_format = copy.copy(formats.line_format)
my_format.conversion.add(key = 'address', # name of the field
                        convert_value = hex, # gets original value
                        convert_item = "{0}={1}".format, # gets called with: key,
                        ↪ converted_value
                        required = True)

# output messages with name 'memory' to stderr
add_emitters(('memory', levels.DEBUG, filters.names('memory'), outputs.
            ↪ StreamOutput(format = my_format)))
```

3.2.5 Filtering Output

The messages output by an emitter are determined by its `min_level` and `filter` (a *function* which take a *Message* and returns bool). These attributes may be changed while the application is running. The `filter` attribute of emitters is *intelligent*; you may assign strings, bools or functions and it will magically do the right thing. Assigning a list indicates that *all* of the filters must pass for the message to be output.

```
e = emitters['memory']
e.min_level = levels.WARNING
# True allows all messages through (None works as well)
e.filter = True
# False blocks all messages
e.filter = False
# Strings are interpreted as regexes (regex objects ok too)
e.filter = "^mem.*y$"
# functions are passed the message; return True to emit
e.filter = lambda msg: msg.fields['address'] > 0xDECAF
# lists are all()'d
e.filter = ["^mem.y$", lambda msg: msg.fields['address'] > 0xDECAF]
```

See also:

Available *filters*

3.3 Reference Guide

3.3.1 Dynamic Logging

Any functions in message args/fields are called and the value substituted.

```
>>> import os
>>> from twiggy.lib import thread_name
>>> thread_name()
'MainThread'
>>> log.fields(pid=os.getpid).info("I'm in thread {0}", thread_name)
INFO:pid=...|I'm in thread MainThread
```

This can be useful with partially-bound loggers, which lets us do some cool stuff. Here's a proxy class that logs which thread accesses attributes.

```
class ThreadTracker(object):
    """a proxy that logs attribute access"""
    def __init__(self, obj):
        self.__obj = obj
        # a partially bound logger
        self.__log = log.name("tracker").fields(obj_id=id(obj), thread=thread_name)
        self.__log.debug("started tracking")
    def __getattr__(self, attr):
        self.__log.debug("accessed {0}", attr)
        return getattr(self.__obj, attr)

class Bunch(object):
    pass
```

Let's see it in action.

```

>>> foo = Bunch()
>>> foo.bar = 42
>>> tracked = ThreadTracker(foo)
DEBUG:tracker:obj_id=...:thread=MainThread|started tracking
>>> tracked.bar
DEBUG:tracker:obj_id=...:thread=MainThread|accessed bar
42
>>> import threading
>>> t=threading.Thread(target = lambda: tracked.bar * 2, name = "TheDoubler")
>>> t.start(); t.join()
DEBUG:tracker:obj_id=...:thread=TheDoubler|accessed bar

```

If you really want to log a callable, `repr()` it or wrap it in `lambda`.

See also:

procinfo feature

3.3.2 Features!

Features are optional additions of logging functionality to the *log*. They encapsulate common logging patterns. Code can be written using a feature, enhancing what information is logged. The feature can be disabled at *runtime* if desired.

Warning: Features are currently deprecated, pending a reimplementation in version 0.5

```

>>> from twigg.features import socket as socket_feature
>>> log.addFeature(socket_feature.socket)
>>> import socket
>>> s = socket.socket(socket.AF_INET, socket.SOCK_STREAM)
>>> s.connect(('www.python.org', 80))
>>> log.socket(s).debug("connected")
DEBUG:host=...:ip_addr=...:port=80:service=http|connected
>>> # turn off the feature - the name is still available
... log.disableFeature('socket')
>>> log.socket(s).debug("connected")
DEBUG|connected
>>> # use a different implementation
... log.addFeature(socket_feature.socket_minimal, 'socket')
>>> log.socket(s).debug("connected")
DEBUG:ip_addr=...:port=80|connected

```

3.3.3 Stays Out of Your Way

Twigg tries to stay out of your way. Specifically, an error in logging should **never** propagate outside the logging subsystem and cause your main application to crash. Instead, errors are trapped and reported by the *internal_log*.

Instances of *InternalLogger* only have a single *Output* - they do not use emitters. By default, these messages are sent to standard error. You may assign an alternate output (such as a file) to `twigg.internal_log.output` if desired, with the following conditions:

- the output should be failsafe - any errors that occur during internal logging will be dumped to standard error, and suppressed, causing the original message to be discarded.

- accordingly, networked or asynchronous outputs are not recommended.
- make sure someone is reading these log messages!

3.3.4 Concurrency

Locking in twiggy is as fine-grained as possible. Each individual output has its own lock (if necessary), and only holds that lock when writing. Using redundant outputs (ie, pointing to the same file) is not supported and will cause logfile corruption.

Asynchronous loggers never lock.

3.3.5 Use by Libraries

Libraries require special care to be polite and usable by application code. The library should have a single bound in its top-level package that's used by modules. Library logging should generally be silent by default.

```
# in mylib/__init__.py
log = twiggy.log.name('mylib')
log.min_level = twiggy.levels.DISABLED

# in mylib/some_module.py
from . import log
log.debug("hi there")
```

This allows application code to enable/disable all of library's logging as needed.

```
# in twiggy_setup
import mylib
mylib.log.min_level = twiggy.levels.INFO
```

In addition to *min_level*, loggers also have a *filter*. This filter operates *only on the format string*, and is intended to allow users to selectively disable individual messages in a poorly-written library.

```
# in mylib:
for i in xrange(1000000):
    log.warning("blah blah {0}", 42)

# in twiggy_setup: turn off stupidity
mylib.log.filter = lambda format_spec: format_spec != "blah blah {0}"
```

Note that using a filter this way is an optimization - in general, application code should use *emitters* instead.

3.3.6 Tips And Tricks

Alternate Styles

In addition to the default new-style (braces) format specs, twiggy also supports old-style (percent, aka printf) and templates (dollar). The aliases {}, % and \$ are also supported.

```
>>> log.options(style='percent').info('I like %s', "bikes")
INFO|I like bikes
>>> log.options(style='dollar').info('$what kill', what='Cars')
INFO|Cars kill
```


Use Fields

Use `fields()` to include key-value data in a message instead of embedding it the human-readable string.

```
# do this:
log.fields(key1='a', key2='b').info("stuff happened")

# not this:
log.info("stuff happened. key1: {0} key2: {1}", 'a', 'b')
```

3.3.7 Technical Details

Independence of logger instances

Each log instance created by partial binding is independent from each other. In particular, a logger's `name()` has no relation to the object; it's just for human use.

```
>>> log.name('bob') is log.name('bob')
False
```

Optimizations

Twiggy has been written to be fast, minimizing the performance impact on the main execution path. In particular, messages that will cause no output are handled as quickly as possible. Users are therefore encouraged to add lots of logging for development/debugging purposes and then turn them off in production.

The emit methods can be hidden behind an appropriate `assert`. Python will eliminate the statement entirely when run with bytecode optimization (`python -O`).

```
assert log.debug("This goes away with python -O") is None
assert not log.debug("So does this")
```

Note: The author doesn't particularly care for code written like this, but likes making his users happy more.

3.3.8 Extending Twiggy

When developing extensions to twiggy, use the `devel_log`. An `InternalLogger`, the `devel_log` is completely separate from the main `log`. By default, messages logged to the `devel_log` are discarded; assigning an appropriate `Output` to its `output` attribute before using.

Writing Features

Warning: Features are currently deprecated, pending a reimplementaion in version 0.5

Features are used to encapsulate common logging patterns. They are implemented as methods added to the `Logger` class. They receive an instance as the first argument (ie, `self`). *Enable the feature* before using.

Features come in two flavors: those that add information to a message's fields or set options, and those that cause output.

Features which only add fields/set options should simply call the appropriate method on `self` and return the resultant object.:

```
def dimensions(self, shape):
    return self.fields(height=shape.height, width=shape.width)
```

Features can also emit messages as usual. Do not return from these methods.:

```
def sayhi(self, lang):
    if lang == 'en':
        self.info("Hello world")
    elif lang == 'fr':
        self.info("Bonjour tout le monde")
```

If the feature should add fields *and* emit in the same step (like `struct()`), use the `emit()` decorators. Here's a prototype feature that dumps information about a WSGI environ.:

```
from twiggy.logger import emit

@emit.info
def dump_wsgi(self, wsgi_environ):
    keys = ['SERVER_PROTOCOL', 'SERVER_PORT', 'SERVER_NAME', 'CONTENT_LENGTH',
    ↪ 'CONTENT_TYPE', 'QUERY_STRING', 'PATH_INFO', 'SCRIPT_NAME', 'REQUEST_METHOD']
    d = {}
    for k in keys:
        d[k] = wsgi_environ.get(k, '')

    for k, v in wsgi_environ.iteritems():
        if k.startswith('HTTP_'):
            k = k[5:].title().replace('_', '-')
            d[k] = v

    # if called on an unnamed logger, add a name
    if name not in self._fields:
        self = self.name('dumpwsgi')

    return self.fields_dict(d)
```

Writing Outputs and Formats

Outputs do the work of writing a message to an external resource (file, socket, etc.). User-defined outputs should inherit from `Output` or `AsyncOutput` if they wish to support *asynchronous logging* (preferred).

An Output subclass's `__init__` should take a *format* and any parameters needed to acquire resources (filename, hostname, etc.), but *not the resources themselves*. These are created in `__open()`. Implementations supporting asynchronous logging should also take a `msg_buffer` argument.

Outputs should define the following:

`Output.__open()`
Acquire any resources needed for writing (files, sockets, etc.)

`Output.__close()`
Release any resources acquired in `__open`

Output.**_write**(*x*)

Do the work of writing

Parameters *x* – an implementation-dependent object to be written.

If the output requires locking to be thread-safe, set the class attribute `use_locks` to True (the default). Turning off may give slightly higher throughput.

The `format` callable is Output-specific; it should take a `Message` and return an appropriate object (string, database row, etc.) to be written. **Do not modify** the received message - it is shared by all outputs. `ConversionTables` are particularly useful for formatting fields. They are commonly used with `LineFormat` to format messages for text-oriented output.

```
from twiggy.lib.converter import ConversionTable
conversion = ConversionTable()

fields = {'shape': 'square',
          'height': 10,
          'width': 5,
          'color': 'blue'}

# hide shape field name
# uppercase value
# make mandatory
conversion.add(key = 'shape',
              convert_value = str.upper,
              convert_item = '{1}'.format, # stringify 2nd item (value)
              required = True)

# format height value with two decimal places
# show as "<key> is <value>"
conversion.add('height', '{0:.2f}'.format, "{0} is {1}".format)

# separate fields in final output by colons
conversion.aggregate = ':'.join

# unknown items are sorted by key

# unknown values are stringified
conversion.generic_value = str

# show unknown items as "<key>=<value>"
conversion.generic_item = "{0}={1}".format

# convert!
print(conversion.convert(fields))
```

```
SQUARE:height is 10.00:color=blue:width=5
```

3.4 API Reference

3.4.1 Global Objects

`twiggy.log`

the magic log object

`twiggy.internal_log`

InternalLogger for reporting errors within Twiggy itself

`twiggy.devel_log`

InternalLogger for use by developers writing extensions to Twiggy

`twiggy.emitters`

the global *emitters* dictionary, tied to the *log*

`twiggy.add_emitters(*tuples)`

Add multiple emitters

tuples should be (name_of_emitter, min_level, filter, output). The last three are passed to *Emitter*.

`twiggy.quick_setup(min_level=<LogLevel DEBUG>, file=None, msg_buffer=0)`

Quickly set up *emitters*.

Parameters

- **min_level** (*LogLevel*) – lowest message level to cause output
- **file** (*string*) – filename to log to, or `sys.stdout`, or `sys.stderr`. None means standard error.
- **msg_buffer** (*int*) – number of messages to buffer, see `outputs.AsyncOutput.msg_buffer`

3.4.2 Features

Optional additions of logging functionality

procinfo

Logging feature to add information about process, etc.

`twiggy.features.procinfo.procinfo(self)`

Adds the following fields:

Hostname current hostname

Pid current process id

Thread current thread name

socket

Logging feature to add information about a socket

`twiggy.features.socket.socket(self, s)`

Adds the following fields:

ip_addr numeric IP address

port port number

host peer hostname, as returned by `getnameinfo()`

service the human readable name of the service on `port`

Parameters **s** (*socket*) – the socket to extract information from

twiggy.features.socket.**socket_minimal** (*self*, *s*)
 Like *socket*, but only log *ip_addr* and *port*

3.4.3 Filters

twiggy.filters.**filter** (*msg* : *Message*) → bool
 A *filter* is any function that takes a *Message* and returns True if it should be *emitted*.

twiggy.filters.**msg_filter** (*x*) → filter
 create a *filter* intelligently

You may pass:

- None, True** the filter will always return True
- False** the filter will always return False
- string** compiled into a regex
- regex** `match()` against the message text
- callable** returned as is
- list** apply *msg_filter* to each element, and `all()` the results

Return type *filter* function

twiggy.filters.**names** (**names*) → filter
 create a *filter*, which gives True if the message's name equals any of those provided

names will be stored as an attribute on the filter.

Parameters **names** (*strings*) – names to match

Return type *filter* function

twiggy.filters.**glob_names** (**names*) → filter
 create a *filter*, which gives True if the message's name globs those provided.

names will be stored as an attribute on the filter.

This is probably quite a bit slower than *names()*.

Parameters **names** (*strings*) – glob patterns.

Return type *filter* function

class twiggy.filters.**Emitter**
 Hold and manage an *Output* and associated *filter()*

min_level
 only emit if greater than this *LogLevel*

filter
 arbitrary *filter()* on message contents. Assigning to this attribute is *intelligent*.

_output
Output to emit messages to. Do not modify.

3.4.4 Formats

Formats are single-argument callables that take a *Message* and return an object appropriate for the *Output* they are assigned to.

```
class twiggy.formats.LineFormat (separator=':', traceback_prefix='\nTRACE', conversion=line_conversion)
```

separator

string to separate line parts. Defaults to `:`.

traceback_prefix

string to prepend to traceback lines. Defaults to `\nTRACE`. Set to `'\n'` (double backslash n) to roll up tracebacks to a single line.

conversion

ConversionTable used to format *fields*. Defaults to *line_conversion*

format_text (*msg*)

format the text part of a message

format_fields (*msg*)

format the fields of a message

format_traceback (*msg*)

format the traceback part of a message

`twiggy.formats.line_conversion`

a default line-oriented *ConversionTable*. Produces a nice-looking string from *fields*.

Fields are separated by a colon (`:`). Resultant string includes:

time in iso8601 format (required)

level message level (required)

name logger name

Remaining fields are sorted alphabetically and formatted as `key=value`

`twiggy.formats.line_format`

a default *LineFormat* for output to a file. *Sample output*.

Fields are formatted using *line_conversion* and separated from the message *text* by a colon (`:`). Traceback lines are prefixed by `TRACE`.

`twiggy.formats.shell_conversion`

a default line-oriented *ConversionTable* for use in the shell. Returns the same string as *line_conversion* but drops the `time` field.

`twiggy.formats.shell_format`

a default *LineFormat* for use in the shell. Same as *line_format* but uses *shell_conversion* for *fields*.

3.4.5 Levels

Levels include (increasing severity): `DEBUG`, `INFO`, `NOTICE`, `WARNING`, `ERROR`, `CRITICAL`, `DISABLED`

class `twiggy.levels.LogLevel` (*name*, *value*)

A log level. Users should *not* create new instances.

Levels are opaque; they may be compared to each other, but nothing else.

`twiggy.levels.name2level` (*name*)
 return a *LogLevel* from a case-insensitive string

3.4.6 Library

`twiggy.lib.iso8601time` (*gmtime=None*)
 convert time to ISO 8601 format - it sucks less!

Parameters `gmtime` (*time.struct_time*) – time tuple. If None, use `time.gmtime()` (UTC)

XXX timezone is not supported

`twiggy.lib.thread_name` ()
 return the name of the current thread

Converter

class `twiggy.lib.converter.Converter` (*key, convert_value, convert_item, required=False*)
 Holder for *ConversionTable* items

Variables

- **key** – the key to apply the conversion to
- **convert_value** (*function*) – one-argument function to convert the value
- **convert_item** (*function*) – two-argument function converting the key and converted value
- **required** (*bool*) – is the item required to present. Items are optional by default.

`twiggy.lib.converter.same_value` (*v*)
 return the value unchanged

`twiggy.lib.converter.same_item` (*k, v*)
 return the item unchanged

`twiggy.lib.converter.drop` (*k, v*)
 return None, indicating the item should be dropped

New in version 0.5.0: Add *same_value*, *same_item*, *drop*.

class `twiggy.lib.converter.ConversionTable` (*seq*)
 Convert data dictionaries using *Converters*

For each item in the dictionary to be converted:

1. Find one or more corresponding converters *c* by matching key.
2. Build a list of converted items by calling `c.convertItem(item_key, c.convertValue(item_value))`. The list will have items in the same order as converters were supplied.
3. Dict items for which no converter was found are sorted by key and passed to *generic_value / generic_item*. These items are appended to the list from step 2.
4. If any required items are missing, `ValueError` is raised.
5. The resulting list of converted items is passed to *aggregate*. The value it returns is the result of the conversion.

Users may override *generic_value/generic_item/aggregate* by subclassing or assigning a new function on a `ConversionTable` instance.

Really, it's *pretty intuitive*.

__init__ (*seq=None*)

Parameters *seq* – a sequence of `Converters`

You may also pass 3-or-4 item arg tuples or kwarg dicts (which will be used to create `Converters`)

convert (*d*)

do the conversion

Parameters *d* (*dict*) – the data to convert. Keys should be strings.

generic_value (*value*)

convert values for which no specific `Converter` is supplied

generic_item (*key, value*)

convert items for which no specific `Converter` is supplied

aggregate (*converted*)

aggregate list of converted items. The return value of *convert*

copy ()

make an independent copy of this `ConversionTable`

get (*key*)

return the *first* converter for key

get_all (*key*)

return a list of all converters for key

add (**args, **kwargs*)

Append a `Converter`.

args & *kwargs* will be passed through to its constructor

delete (*key*)

delete the *all* of the converters for key

3.4.7 Logger

Loggers should not be created directly by users; use the global *log* instead.

class `twiggy.logger.BaseLogger` (*fields=None, options=None, min_level=None*)

Base class for loggers

_fields

dictionary of bound fields for *structured logging*. By default, contains a single field `time` with value `time.gmtime()`. This function will be called for each message emitted, populating the field with the current `time.struct_time`.

_options

dictionary of bound *options*.

min_level

minimum `LogLevel` for which to emit. For optimization purposes only.

fields (***kwargs*) → bound `Logger`

bind fields for *structured logging*. *kwargs* are interpreted as names/values of fields.

fields_dict (*d*) → bound Logger
 bind fields for structured logging. Use this instead of *fields* if you have keys which are not valid Python identifiers.

Parameters *d* (*dict*) – dictionary of fields. Keys should be strings.

options (***kwargs*) → bound Logger
 bind *options* for message creation.

trace (*trace='error'*) → bound Logger
 convenience method to enable *traceback logging*

name (*name*) → bound Logger
 convenience method to bind *name* field

struct (***kwargs*) → bound Logger
 convenience method for *structured logging*. Calls *fields()* and emits at *info*

struct_dict (*d*) → bound Logger
 convenience method for *structured logging*. Use instead of *struct* if you have keys which are not valid Python identifiers.

Parameters *d* (*dict*) – dictionary of fields. Keys should be strings.

The following methods cause messages to be emitted. *format_spec* is a template string into which *args* and *kwargs* will be substituted.

debug (*format_spec='', *args, **kwargs*)
 Emit at DEBUG level

info (*format_spec='', *args, **kwargs*)
 Emit at INFO level

notice (*format_spec='', *args, **kwargs*)
 Emit at NOTICE level

warning (*format_spec='', *args, **kwargs*)
 Emit at WARNING level

error (*format_spec='', *args, **kwargs*)
 Emit at ERROR level

critical (*format_spec='', *args, **kwargs*)
 Emit at CRITICAL level

class `twiggy.logger.Logger` (*fields=None, options=None, min_level=None*)
 Logger for end-users. The type of the magic *log*

filter

Filter on *format_spec*. For optimization purposes only. Should have the following signature:

func (*format_spec : string*) → bool
 Should the message be emitted.

classmethod **addFeature** (*func, name=None*)
 add a feature to the class

Parameters

- **func** – the function to add
- **name** (*string*) – the name to add it under. If None, use the function’s name.

classmethod disableFeature (*name*)

disable a feature.

A method will still exist by this name, but it won't do anything.

Parameters *name* (*string*) – the name of the feature to disable.

classmethod delFeature (*name*)

delete a feature entirely

Parameters *name* (*string*) – the name of the feature to remove

class `twiggy.logger.InternalLogger` (*output*, *fields=None*, *options=None*, *min_level=None*)

Special-purpose logger for internal uses

Sends messages directly to output, bypassing *emitters*.

Variables *output* (*Output*) – an output to write to

`twiggy.logger.emit` (*level*)

a decorator that emits at *level* after calling the method. The method should return a *Logger* instance.

For convenience, decorators for the various levels are available as `emit.debug`, `emit.info`, etc..

3.4.8 Message

class `twiggy.message.Message` (*level*, *format_spec*, *fields*, *options*, *args*, *kwargs*)

A logging message. Users never create these directly.

Changed in version 0.4.1: Pass args/kwags as list/dict instead of via **/*** expansion. The constructor takes a dict of *options* to control message creation. In addition to *suppress_newlines*, the following options are recognized:

trace control traceback inclusion. Either a traceback tuple, or one of the strings `always`, `error`, in which case a traceback will be extracted from the current stack frame.

style the style of template used for *format_spec*. One of `braces`, `percent`, `dollar`. The aliases `{}`, `%` and `$` are also supported.

Any callables passed in *fields*, *args* or *kwargs* will be called and the returned value used instead. See *dynamic messages*.

All attributes are read-only.

fields

dictionary of *structured logging* fields. Keys are string, values are arbitrary. A *level* item is required.

suppress_newlines

should newlines be escaped in output. Boolean.

traceback

a stringified traceback, or `None`.

text

the human-readable message. Constructed by substituting *args*/*kwargs* into *format_spec*. String.

__init__ (*level*, *format_spec*, *fields*, *options*, *args*, *kwargs*)

Parameters

- **level** (*LogLevel*) – the level of the message
- **format_spec** (*string*) – the human-readable message template. Should match the style in options.

- **fields** (*dict*) – dictionary of fields for *structured logging*
- **args** (*tuple*) – substitution arguments for `format_spec`.
- **kwargs** (*dict*) – substitution keyword arguments for `format_spec`.
- **options** (*dict*) – a dictionary of *options* to control message creation.

3.4.9 Outputs

class `twigg.outputs.Output` (*format=None, close_atexit=True*)

`__format__`

a *callable* taking a *Message* and formatting it for output. None means return the message unchanged.

`use_locks`

Class variable, indicating that locks should be used when running in a synchronous, multithreaded environment. Threadsafe subclasses may disable locking for higher throughput. Defaults to True.

`__init__` (*format=None, close_atexit=True*)

Parameters

- **format** (*format*) – the format to use. If None, return the message unchanged.
- **close_atexit** (*bool*) – should `close()` be registered with `atexit`. If False, the user is responsible for closing the output.

New in version 0.4.1: Add the `close_atexit` parameter.

`close()`

Finalize the output.

The following methods should be implemented by subclasses.

`__open()`

Acquire any resources needed for writing (files, sockets, etc.)

`__close()`

Release any resources acquired in `__open`

`__write(x)`

Do the work of writing

Parameters **x** – an implementation-dependent object to be written.

class `twigg.outputs.AsyncOutput` (*msg_buffer=0*)

An *Output* with support for *asynchronous logging*.

Inheriting from this class transparently adds support for asynchronous logging using the multiprocessing module. This is off by default, as it can cause log messages to be dropped.

Parameters **msg_buffer** (*int*) – number of messages to buffer in memory when using asynchronous logging. 0 turns asynchronous output off, a negative integer means an unlimited buffer, a positive integer is the size of the buffer.

class `twigg.outputs.FileOutput` (*name, format, mode='a', buffering=1, msg_buffer=0, close_atexit=True*)

Output messages to a file

`name, mode, buffering` are passed to `open()`

class `twiggy.outputs.StreamOutput` (*format, stream=sys.stderr*)
Output to an externally-managed stream.

The stream will be written to, but otherwise left alone (i.e., it will *not* be closed).

class `twiggy.outputs.NullOutput` (*format=None, close_atexit=True*)
An output that just discards its messages

class `twiggy.outputs.ListOutput` (*format=None, close_atexit=True*)
an output that stuffs messages in a list

Useful for unittesting.

Variables `messages` (*list*) – messages that have been emitted

Changed in version 0.4.1: Replace `DequeOutput` with more useful `ListOutput`.

3.5 Testing

This part discusses how to test Twiggy to ensure that Twiggy is built and installed correctly.

3.5.1 Requirements

The following need to be installed prior to testing:

- Python 2.7.1 or greater.
- The `coverage` module.
- `sphinx` 1.0.8 or greater. You'll need to get and build the `sphinx` source.
- `Twiggy` source.

3.5.2 Running Tests

Note: Tests **must** be run from the Twiggy root directory to work.

To run all tests (unittest and Sphinx doctests):

```
./scripts/run-twiggy-tests.sh
```

To run coverage tests, run:

```
./scripts/cover-twiggy-tests.sh discover -b
```

To run coverage tests on a specific module, run:

```
./scripts/cover-twiggy-tests.sh tests.test_levels
```

3.6 Glossary

asynchronous logging performance enhancement that moves formatting and writing messages to a separate process.
See *Asynchronous Logging*.

structured logging logging information in easy-to-parse key-value pairs, instead of embedded in a human-readable message. See an *example*

3.7 Changelog

3.7.1 0.5.0

XXX Unreleased

- add a NOTICE level between INFO and WARNING
- add sameValue, sameItem, drop helper functions to lib.converter
- support {}, %, \$ as style aliases.
- PEP8 name compliance
- add logging_compat module for compatibility with stdlib's logging

3.7.2 0.4.7

03/09/2015 - add missing classifiers to setup.py

3.7.3 0.4.6

03/09/2015 - also suppress newlines in fields output - Python 3 support

3.7.4 0.4.5

03/18/2013 - documentation update, move to Github

3.7.5 0.4.4

07/12/2011 - support Python 2.6

3.7.6 0.4.3

12/20/2010 - add check for Python ≥ 2.7 to setup.py, to reduce invalid bug reports.

3.7.7 0.4.2

11/11/2010 - fix broken installer

3.7.8 0.4.1

11/8/2010

- full test coverage; numerous bug fixes
- add `close_atexit` parameter to `Outputs`
- replace `DequeOutput` with `ListOutput`
- deprecate features, pending a rewrite in 0.5
- minor internal API changes

3.7.9 0.4.0

10/18/2010

First serious public release

3.8 Contributors

Twiggy would not be possible without the support of the following people. You have our thanks.

- Peter Fein `pfein@pobox.com`
- Ian Foote `ianfoote@f2s.com`
- Kyle Rickey `kwkard@gmail.com`
- Lin O'Driscoll `nzlinus@gmail.com`

t

- twiggy, 15
- twiggy.features, 16
- twiggy.features.procinfo, 16
- twiggy.features.socket, 16
- twiggy.filters, 17
- twiggy.formats, 18
- twiggy.levels, 18
- twiggy.lib, 19
- twiggy.lib.converter, 19
- twiggy.logger, 20
- twiggy.message, 22
- twiggy.outputs, 23

Symbols

__init__() (twiggy.lib.converter.ConversionTable method), 20
 __init__() (twiggy.message.Message method), 22
 __init__() (twiggy.outputs.Output method), 23
 _close() (twiggy.outputs.Output method), 23
 _fields (twiggy.logger.BaseLogger attribute), 20
 _format (twiggy.outputs.Output attribute), 23
 _open() (twiggy.outputs.Output method), 23
 _options (twiggy.logger.BaseLogger attribute), 20
 _output (twiggy.filters.Emitter attribute), 17
 _write() (twiggy.outputs.Output method), 23

A

add() (twiggy.lib.converter.ConversionTable method), 20
 add_emitters() (in module twiggy), 16
 addFeature() (twiggy.logger.Logger class method), 21
 aggregate() (twiggy.lib.converter.ConversionTable method), 20
 asynchronous logging, 24
 AsyncOutput (class in twiggy.outputs), 23

B

BaseLogger (class in twiggy.logger), 20

C

close() (twiggy.outputs.Output method), 23
 conversion (twiggy.formats.LineFormat attribute), 18
 ConversionTable (class in twiggy.lib.converter), 19
 convert() (twiggy.lib.converter.ConversionTable method), 20
 Converter (class in twiggy.lib.converter), 19
 copy() (twiggy.lib.converter.ConversionTable method), 20
 critical() (twiggy.logger.BaseLogger method), 21

D

debug() (twiggy.logger.BaseLogger method), 21

delete() (twiggy.lib.converter.ConversionTable method), 20
 delFeature() (twiggy.logger.Logger class method), 22
 devel_log (in module twiggy), 16
 disableFeature() (twiggy.logger.Logger class method), 21
 drop() (in module twiggy.lib.converter), 19

E

emit() (in module twiggy.logger), 22
 Emitter (class in twiggy.filters), 17
 emitters (in module twiggy), 16
 error() (twiggy.logger.BaseLogger method), 21

F

fields (twiggy.message.Message attribute), 22
 fields() (twiggy.logger.BaseLogger method), 20
 fields_dict() (twiggy.logger.BaseLogger method), 20
 FileOutput (class in twiggy.outputs), 23
 filter (twiggy.filters.Emitter attribute), 17
 filter (twiggy.logger.Logger attribute), 21
 filter() (in module twiggy.filters), 17
 format_fields() (twiggy.formats.LineFormat method), 18
 format_text() (twiggy.formats.LineFormat method), 18
 format_traceback() (twiggy.formats.LineFormat method), 18

G

generic_item() (twiggy.lib.converter.ConversionTable method), 20
 generic_value() (twiggy.lib.converter.ConversionTable method), 20
 get() (twiggy.lib.converter.ConversionTable method), 20
 get_all() (twiggy.lib.converter.ConversionTable method), 20
 glob_names() (in module twiggy.filters), 17

I

info() (twiggy.logger.BaseLogger method), 21
 internal_log (in module twiggy), 15

InternalLogger (class in twiggy.logger), 22
iso8601time() (in module twiggy.lib), 19

L

line_conversion (in module twiggy.formats), 18
line_format (in module twiggy.formats), 18
LineFormat (class in twiggy.formats), 18
ListOutput (class in twiggy.outputs), 24
log (in module twiggy), 15
Logger (class in twiggy.logger), 21
LogLevel (class in twiggy.levels), 18

M

Message (class in twiggy.message), 22
min_level (twiggy.filters.Emitter attribute), 17
min_level (twiggy.logger.BaseLogger attribute), 20
msg_filter() (in module twiggy.filters), 17

N

name() (twiggy.logger.BaseLogger method), 21
name2level() (in module twiggy.levels), 18
names() (in module twiggy.filters), 17
notice() (twiggy.logger.BaseLogger method), 21
NullOutput (class in twiggy.outputs), 24

O

options() (twiggy.logger.BaseLogger method), 21
Output (class in twiggy.outputs), 23

P

procinfo() (in module twiggy.features.procinfo), 16

Q

quick_setup() (in module twiggy), 16

S

same_item() (in module twiggy.lib.converter), 19
same_value() (in module twiggy.lib.converter), 19
separator (twiggy.formats.LineFormat attribute), 18
shell_conversion (in module twiggy.formats), 18
shell_format (in module twiggy.formats), 18
socket() (in module twiggy.features.socket), 16
socket_minimal() (in module twiggy.features.socket), 17
StreamOutput (class in twiggy.outputs), 23
struct() (twiggy.logger.BaseLogger method), 21
struct_dict() (twiggy.logger.BaseLogger method), 21
structured logging, 25
suppress_newlines (twiggy.message.Message attribute), 22

T

text (twiggy.message.Message attribute), 22
thread_name() (in module twiggy.lib), 19

trace() (twiggy.logger.BaseLogger method), 21
traceback (twiggy.message.Message attribute), 22
traceback_prefix (twiggy.formats.LineFormat attribute), 18

twiggy (module), 15
twiggy.features (module), 16
twiggy.features.procinfo (module), 16
twiggy.features.socket (module), 16
twiggy.filters (module), 17
twiggy.formats (module), 18
twiggy.levels (module), 18
twiggy.lib (module), 19
twiggy.lib.converter (module), 19
twiggy.logger (module), 20
twiggy.message (module), 22
twiggy.outputs (module), 23

U

use_locks (twiggy.outputs.Output attribute), 23

W

warning() (twiggy.logger.BaseLogger method), 21