
tapes Documentation

Release 0.3

Emils Solmanis

April 06, 2015

| | | |
|----------|-------------------------------------|-----------|
| 1 | Quick start | 3 |
| 1.1 | Dependencies | 3 |
| 1.2 | Single process | 3 |
| 1.3 | Utilities | 4 |
| 1.4 | Reporting | 4 |
| 1.5 | Multi process | 5 |
| 2 | API docs | 7 |
| 3 | Benchmarks | 11 |
| 3.1 | Just the timer | 11 |
| 3.2 | Flask | 12 |
| 3.3 | Tornado, single process | 13 |
| 3.4 | Tornado, forked processes | 15 |
| 4 | Indices and tables | 17 |
| | Python Module Index | 19 |

Contents:

Quick start

1.1 Dependencies

The libraries dependencies are kept to a minimum to avoid the overhead for features you choose not to use, e.g.,

- tornado is not formally a dependency, but is obviously needed should you choose to use the tornado features
- pyzmq is not formally a dependency, but is needed should you choose to use the distributed reporting features

1.2 Single process

If you just need to get off the ground and flying, there's convenience methods in the `tapes` module, so you can do:

```
timer = tapes.timer('my.timer')

@app.route('/widgets')
def moo():
    with timer.time():
        return 'stuff'
```

This will simply use a global `tapes.registry.Registry` object instance under the hood.

You can also explicitly create and maintain a registry (or several, if needed) by just creating some:

```
registry0 = Registry()
registry1 = Registry()

timer0 = registry0.timer('some.name')
timer1 = registry1.timer('some.name')
```

Even though the timers were created with the same name, they're maintained in different registries, they're separate instances.

Repeated calls on the same registry with the same name will return the existing instance:

```
timer0 = tapes.timer('some.timer')
timer1 = tapes.timer('some.timer')
# timer0 and timer1 are actually the same instance
```

1.3 Utilities

It's common that you need to add the same metrics to many similar subclasses but have them all named differently based on the class name. Because you would usually want to make the metric a class level attribute to avoid doing the tedious lookup in particular methods, there is a utility metaclass that adds any necessary metrics to the class.

Example:

```
class MyCommonBaseHandler(tornado.web.RequestHandler):
    __metaclass__ = metered_meta(metrics=[
        ('latency', 'my.http.endpoints.{}.latency', registry.timer),
        ('error_rate', 'my.http.endpoints.{}.error_rate', registry.meter),
    ], base=abc.ABCMeta)
```

The `base` argument wires in any existing metas, most commonly `abc.ABCMeta`. The `metrics` argument takes a list of tuples, in which - the 1st arg is the name of the class attribute - the 2nd arg is a template for the metric name. It will get rendered by calling `template.format(class_name)` - the 3rd arg is a callable taking a single argument - the fully rendered name - and returning a metric instance

The above example would add two attributes to all subclasses of `MyCommonBaseHandler` called `latency` and `error_rate`, which would be a timer and a meter, and correspond to metrics parametrized on the class name.

For more info check the full docs at `tapes.meta.metered_meta()`.

1.4 Reporting

For pull style reporting you need to

- call `tapes.registry.Registry.get_stats()`
- expose the returned dict of values somehow

This should be trivial for most web applications, e.g., for Flask:

```
@app.route('/metrics')
def metrics():
    return jsonify(tapes.get_stats())
```

or Tornado:

```
class MetricsHandler(web.RequestHandler):
    @gen.coroutine
    def get(self, *args, **kwargs):
        self.finish(tapes.get_stats())
```

For push style reporting it gets more complicated because of the GIL. There are a couple of scheduled reporters that report with configurable intervals, but they create a `Thread`. For most Python implementations this essentially means that your application would block for the time it is doing the reporting, and while it is probably fine for most scenarios, it's still something to keep in mind.

There is a scheduled reporter base class for `Tornado` as well, which is non-blocking and uses the `IOLoop` to schedule reporting.

Currently the push reporting is implemented for streams and `StatsD`.

The reporters are all used almost the same, i.e., by creating one and calling `start()` on it:


```
reporter = SomeScheduledReporter(
    interval=timedelta(seconds=10), registry=registry
)
reporter.start()
```

If you want a guaranteed low latency setup, you might want to look into the multi-process options below.

Otherwise, take a stroll through [tapes.reporting](#)

1.5 Multi process

Python web applications are generally run with several forks handling the same socket, either via *uWSGI*, *gunicorn*, or, in the case of Tornado, just natively.

This causes problems with metrics, namely

- HTTP reporting breaks. The reported metrics are per-fork and you get random forks every time you make a request
- Metrics aren't strictly commutative, so even if you do push-reporting, you end up losing some info in the combiner, e.g., StatsD
- Push-reporting generally means blocking your main application's execution due to the GIL

Hence, with Tapes you have an option to

- run a light-weight proxy registry per fork
- have an aggregating master registry that does all the reporting in a separate fork

The proxy registries communicate with the master registry via OMQ IPC pub-sub. Because OMQ is really, **really** fast, this ends up almost as fast as just computing the metrics locally anyway.

The obvious drawbacks are

- a separate forked process doing the aggregation and reporting
- a slight lag if the traffic volume is low due to batching in OMQ

With that said, it's much simpler to actually use than it sounds. The reporters are the same, so the only changes are the use of an aggregator and proxy registries.

NOTE: due to the way gauge operates, it's unavailable in the distributed mode. If / when I figure out how to combine it, I'll add it.

Tornado example:

```
registry = DistributedRegistry()

class TimedHandler(web.RequestHandler):
    timer = registry.timer('my.timer')

    @gen.coroutine
    def get(self):
        with TimedHandler.timer.time():
            self.write('finished')

RegistryAggregator(HTTPReporter(8889)).start()

server = httpserver.HTTPServer(application)
server.bind(8888)
```

```
server.start(0)

registry.connect()

ioloop.IOLoop.current().start()
```

NOTE

- `tapes.distributed.registry.RegistryAggregator.start()` is called **before** the `fork()` call
- `tapes.distributed.registry.DistributedRegistry.connect()` is called **after** the `fork()` call

Check the API docs for more info – `tapes.distributed.registry`.

class `tapes.registry.Registry`

Factory and storage location for all metrics stuff.

Use producer methods to create metrics. Metrics are hierarchical, the names are split on ‘.’.

counter (*name*)

Creates or gets an existing counter.

Parameters *name* – The name

Returns The created or existing counter for the given name

gauge (*name, producer*)

Creates or gets an existing gauge.

Parameters *name* – The name

Returns The created or existing gauge for the given name

get_stats ()

Retrieves the current values of the metrics associated with this registry, formatted as a dict.

The metrics form a hierarchy, their names are split on ‘.’. The returned dict is an *addict*, so you can use it as either a regular dict or via attributes, e.g.,

```
>>> import tapes
>>> registry = tapes.Registry()
>>> timer = registry.timer('my.timer')
>>> stats = registry.get_stats()
>>> print(stats['my']['timer']['count'])
0
>>> print(stats.my.timer.count)
0
```

Returns The values of the metrics associated with this registry

histogram (*name*)

Creates or gets an existing histogram.

Parameters *name* – The name

Returns The created or existing histogram for the given name

meter (*name*)

Creates or gets an existing meter.

Parameters *name* – The name

Returns The created or existing meter for the given name

timer (*name*)

Creates or gets an existing timer.

Parameters **name** – The name

Returns The created or existing timer for the given name

`tapes.meta.metered_meta` (*metrics*, *base=<type 'type'>*)

Creates a metaclass that will add the specified metrics at a path parametrized on the dynamic class name.

Prime use case is for base classes if all subclasses need separate metrics and / or the metrics need to be used in base class methods, e.g., Tornado's RequestHandler like:

```
import tapes
import tornado
import abc

registry = tapes.Registry()

class MyCommonBaseHandler(tornado.web.RequestHandler):
    __metaclass__ = metered_meta([
        ('latency', 'my.http.endpoints.{}.latency', registry.timer)
    ], base=abc.ABCMeta)

    @tornado.gen.coroutine
    def get(self, *args, **kwargs):
        with self.latency.time():
            yield self.get_impl(*args, **kwargs)

    @abc.abstractmethod
    def get_impl(self, *args, **kwargs):
        pass

class MyImplHandler(MyCommonBaseHandler):
    @tornado.gen.coroutine
    def get_impl(self, *args, **kwargs):
        self.finish({'stuff': 'something'})

class MyOtherImplHandler(MyCommonBaseHandler):
    @tornado.gen.coroutine
    def get_impl(self, *args, **kwargs):
        self.finish({'other stuff': 'more of something'})
```

This would produce two different relevant metrics,

- `my.http.endpoints.MyImplHandler.latency`
- `my.http.endpoints.MyOtherImplHandler.latency`

and, as an unfortunate side effect of adding it in the base class, a `my.http.endpoints.MyCommonBaseHandler.latency` too.

Parameters

- **metrics** – list of (*attr_name*, *metrics_path_template*, *metrics_factory*)
- **base** – optional meta base if other than *type*

Returns a metaclass that populates the class with the needed metrics at paths based on the dynamic class name

class `tapes.reporting.ScheduledReporter` (*interval*, *registry=None*)

Super class for scheduled reporters. Handles scheduling via a Thread.

report ()

Override in subclasses.

A Python Thread is used for scheduling, so whatever this ends up doing, it should be pretty fast.

class `tapes.reporting.http.HTTPReporter` (*port*, *registry=None*)

Exposes metrics via HTTP.

For web applications, you should almost certainly just use your existing framework's capabilities. This is for applications that don't have HTTP easily available.

class `tapes.reporting.statsd.StatsdReporter` (*interval*, *host='localhost'*, *port=8125*, *prefix=None*, *registry=None*)

Reporter for StatsD.

class `tapes.reporting.stream.ThreadedStreamReporter` (*interval*, *stream=<open file '<stdout>', mode 'w'>*, *registry=None*)

Dumps JSON serialized metrics to a stream with an interval

class `tapes.reporting.tornado.TornadoScheduledReporter` (*interval*, *registry=None*, *io_loop=None*)

Scheduled reporter that uses a tornado IOLoop for scheduling

class `tapes.reporting.tornado.statsd.TornadoStatsdReporter` (*interval*, *host='localhost'*, *port=8125*, *prefix=None*, *registry=None*)

Reports to StatsD using an IOLoop for scheduling

class `tapes.reporting.tornado.stream.TornadoStreamReporter` (*interval*, *stream=<open file '<stdout>', mode 'w'>*, *registry=None*, *io_loop=None*)

Writes JSON serialized metrics to a stream using an IOLoop for scheduling

class `tapes.distributed.registry.DistributedRegistry` (*socket_addr='ipc://tapes_metrics.ipc'*)

A registry proxy that pushes metrics data to a RegistryAggregator.

connect ()

Connects to the 0MQ socket and starts publishing.

class `tapes.distributed.registry.RegistryAggregator` (*reporter*, *socket_addr='ipc://tapes_metrics.ipc'*)

Aggregates multiple registry proxies and reports on the unified metrics.

start (*fork=True*)

Starts the registry aggregator.

Parameters **fork** – whether to fork a process; if False, blocks and stays in the existing process

stop ()

Terminates the forked process.

Only valid if started as a fork, because... well you wouldn't get here otherwise. :return:

Benchmarks

The benchmarking scripts are in the source code and have `tox` environments defined for Python 2.7, 3.4 and PyPy. All you need to do to run them yourself is check out the code and run:

```
tox -e benchmark-py27,benchmark-py34,benchmark-pypy
```

The results shown here are from an old-ish i5 laptop and are by no means 100% conclusive. With that said...

3.1 Just the timer

This just runs some timer metric calls via Python's `timeit`. Try it yourself by checking out the code and doing:


```
python -m benchmark.util.quick
```

The times (lower is better) are


| Version | Tapes | Scales |
|---------|--------------|---------------|
| Py 2.7 | 3.0038728714 | 12.1747989655 |
| Py 3.4 | 3.3859353300 | 15.8657573729 |
| PyPy | 0.4108750820 | 11.163599968 |

3.2 Flask


3.2.1 Python 2.7



3.2.2 Python 3.4

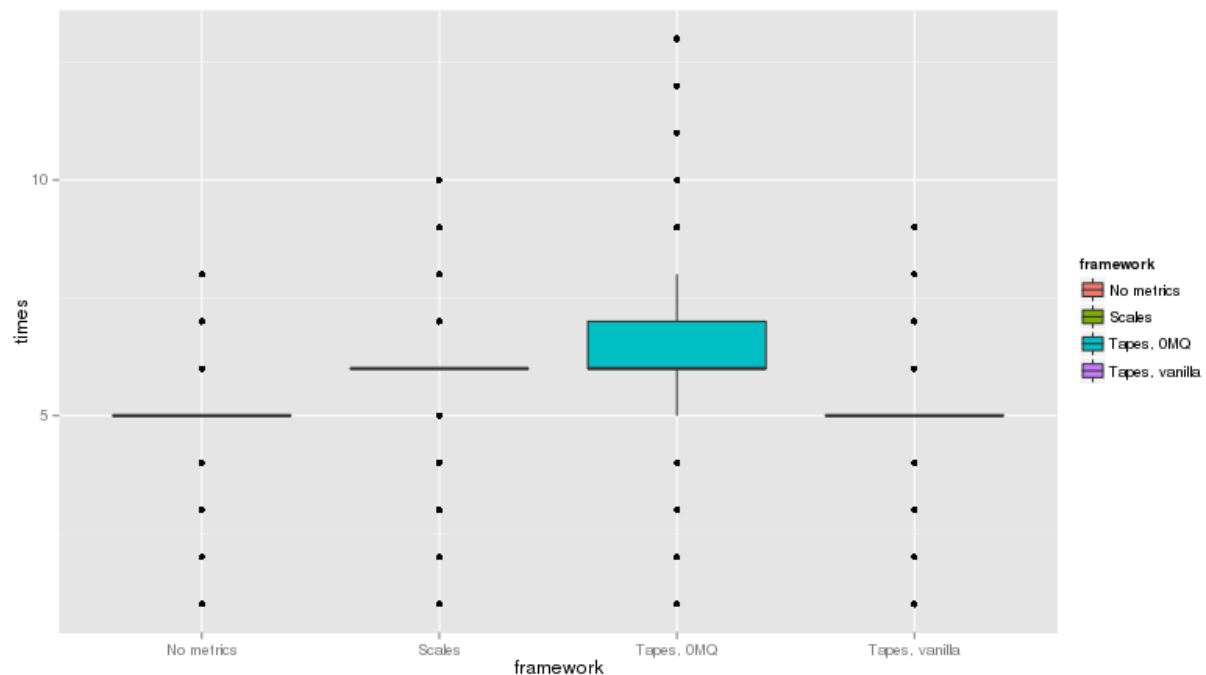


3.2.3 PyPy 2.4




3.3 Tornado, single process


3.3.1 Python 2.7



3.3.2 Python 3.4




3.3.3 PyPy 2.4




3.4 Tornado, forked processes


3.4.1 Python 2.7



3.4.2 Python 3.4



3.4.3 PyPy 2.4



This is a native Python metrics library implementation. It currently supports Python 2.7, 3.4 and PyPy, but most other versions should probably work.

Compared to other libraries, tapes

- doesn't use a separate thread to decay the moving average's weights
- doesn't use "clever tricks" with stack frames, which means it also doesn't break in weird ways (e.g., when used within a method decorated with `@staticmethod`)
- has clean code, that passes `flake8` validation
- is also performance-oriented

You can check the [Benchmarks](#) yourself!

Indices and tables

- `genindex`
- `modindex`
- `search`

t

- tapes.distributed.registry, 9
- tapes.meta, 8
- tapes.registry, 7
- tapes.reporting, 9
 - tapes.reporting.http, 9
 - tapes.reporting.statsd, 9
 - tapes.reporting.stream, 9
 - tapes.reporting.tornado, 9
 - tapes.reporting.tornado.statsd, 9
 - tapes.reporting.tornado.stream, 9

C

`connect()` (tapes.distributed.registry.DistributedRegistry method), 9
`counter()` (tapes.registry.Registry method), 7

D

`DistributedRegistry` (class in tapes.distributed.registry), 9

G

`gauge()` (tapes.registry.Registry method), 7
`get_stats()` (tapes.registry.Registry method), 7

H

`histogram()` (tapes.registry.Registry method), 7
`HTTPReporter` (class in tapes.reporting.http), 9

M

`meter()` (tapes.registry.Registry method), 7
`metered_meta()` (in module tapes.meta), 8

R

`Registry` (class in tapes.registry), 7
`RegistryAggregator` (class in tapes.distributed.registry), 9
`report()` (tapes.reporting.ScheduledReporter method), 9

S

`ScheduledReporter` (class in tapes.reporting), 9
`start()` (tapes.distributed.registry.RegistryAggregator method), 9
`StatsdReporter` (class in tapes.reporting.statsd), 9
`stop()` (tapes.distributed.registry.RegistryAggregator method), 9

T

tapes.distributed.registry (module), 9
tapes.meta (module), 8
tapes.registry (module), 7
tapes.reporting (module), 9
tapes.reporting.http (module), 9

tapes.reporting.statsd (module), 9
tapes.reporting.stream (module), 9
tapes.reporting.tornado (module), 9
tapes.reporting.tornado.statsd (module), 9
tapes.reporting.tornado.stream (module), 9
`ThreadedStreamReporter` (class in tapes.reporting.stream), 9
`timer()` (tapes.registry.Registry method), 8
`TornadoScheduledReporter` (class in tapes.reporting.tornado), 9
`TornadoStatsdReporter` (class in tapes.reporting.tornado.statsd), 9
`TornadoStreamReporter` (class in tapes.reporting.tornado.stream), 9