
measure_it Documentation

Release 0.6.0

Pete Fein

Apr 22, 2018

1	Installation	3
1.1	Minimal	3
1.2	Batteries Included	3
2	Basic Usage	5
2.1	Functions	5
2.2	Blocks	6
2.3	Iterables	6
2.4	Generators	7
2.5	Reducers & Producers	7
3	Output	11
3.1	Logging	11
3.2	Comma Separated	11
3.3	Summary Reports	12
3.4	statsd	13
4	API Documentation	15
4.1	instrument	15
4.2	instrument.output	16
4.3	instrument.output.csv	17
4.4	instrument.output.logging	18
4.5	instrument.output.table	18
4.6	instrument.output.plot	18
4.7	instrument.output.statsd	18
4.8	instrument.output._numpy	19
5	Backwards Incompatibilities	21
5.1	0.5 -> 0.6	21
5.2	0.4 -> 0.5	21
5.3	0.3 -> 0.4	21
6	Future Plans	23
6.1	Automagic Instrumentation	23
6.2	More metric backends	23
6.3	Bonus features	23
6.4	Modernization	24

7	Changelog	25
7.1	0.6.0	25
7.2	0.5.1	25
7.3	0.4	25
7.4	0.3	26
7.5	0.2	26
7.6	0.1	26
	Python Module Index	27

`instrument` provides instrumentation primitives for metrics and benchmarking.

Python `instrument` was formerly known as `measure-it`.

author Pete Fein <pete@wearpants.org>

license BSD

versions Python 3.4+

source <https://github.com/wearpants/instrument>

homepage <https://instrument.readthedocs.org/en/latest/>

instrument subscribes to the Python batteries-included philosophy. It ships with support for a number of different *metric backends*, and eventually instrumentation for popular packages. Since most users will want only a subset of this functionality, optional dependencies are not installed by default.

Note that the included `requirements.txt` file includes *all* dependencies (for build/test purposes), and is almost certainly not what you want.

1.1 Minimal

To install the base package (with no dependencies beyond the standard library):

```
pip install instrument
```

This includes: all generic measurement functions, `print_metric()` and `csv` metrics.

1.2 Batteries Included

You should have completed the minimal install already. To install the dependencies for an optional component, specify it in brackets with `--upgrade`:

```
pip install --upgrade instrument[statsd]
```

The following extra targets are supported:

- `statsd`: statsd metric
- `numpy`: statistics metrics, based on numpy
- `plot`: graphs of metrics, based on matplotlib. Because Python packaging is brain damaged, you must install the `numpy` target first. You'll need the `agg backend`.

instrument provides instrumentation primitives for runtime metrics and benchmarking. These helpers report a count of items, elapsed time and an optional name, making it easy to gather information about your code's performance.

```
>>> import instrument
>>> from time import sleep
```

2.1 Functions

The *function()* decorator measures total function execution time:

```
>>> @instrument.function()
... def slow():
...     # you'd do something useful here
...     sleep(.1)
...     return "SLOW"
>>> slow()
__main__.slow: 1 items in 0.10 seconds
'SLOW'
```

This works in classes too. A name will be inferred if you don't provide one:

```
>>> class CrunchCrunch(object):
...     @instrument.function()
...     def slow(self):
...         # you'd do something useful here
...         sleep(.1)
...         return "SLOW"
>>> CrunchCrunch().slow()
__main__.CrunchCrunch.slow: 1 items in 0.10 seconds
'SLOW'
```

2.2 Blocks

To measure the execution time of a block of code, use a `block()` context manager:

```
>>> with instrument.block(name="slowcode") :  
...     # you'd do something useful here  
...     sleep(0.2)  
slowcode: 1 items in 0.20 seconds
```

You can also pass your own value for `count`; this is useful to measure a resource used by a block (the number of bytes in a file, for example):

```
>>> with instrument.block(name="slowcode", count=42):  
...     # you'd do something useful here  
...     sleep(0.2)  
slowcode: 42 items in 0.20 seconds
```

2.3 Iterables

Iterators & generators often encapsulate I/O, number crunching or other operations we want to gather metrics for:

```
>>> def math_is_hard(N):  
...     x = 0  
...     while x < N:  
...         sleep(.1)  
...         yield x * x  
...         x += 1
```

Timing iterators is tricky. `all()`, `each()` and `first()` record metrics for time and item count for iterables.

Wrap an iterator in `all()` to time how long it takes to consume entirely:

```
>>> underlying = math_is_hard(5)  
>>> underlying  
<generator object math_is_hard at ...>  
>>> _ = instrument.all(underlying)  
>>> squares = list(_)  
5 items in 0.50 seconds
```

The wrapped iterator yields the same results as the underlying iterable:

```
>>> squares  
[0, 1, 4, 9, 16]
```

The `each()` wrapper measures the time taken to produce each item:

```
>>> _ = instrument.each(math_is_hard(5))  
>>> list(_)  
1 items in 0.10 seconds  
[0, 1, 4, 9, 16]
```

The `first()` wrapper measures the time taken to produce the first item:

```
>>> _ = instrument.first(math_is_hard(5))
>>> list(_)
1 items in 0.10 seconds
[0, 1, 4, 9, 16]
```

You can provide a custom name for the metric:

```
>>> _ = instrument.all(math_is_hard(5), name="bogomips")
>>> list(_)
bogomips: 5 items in 0.50 seconds
[0, 1, 4, 9, 16]
```

2.4 Generators

You can use `all()`, `each()` and `first()` as decorators to wrap a generator function (one that uses `yield`):

```
>>> @instrument.each()
... def slow(N):
...     for i in range(N):
...         sleep(.1)
...         yield i
>>> list(slow(3))
__main__.slow: 1 items in 0.10 seconds
__main__.slow: 1 items in 0.10 seconds
__main__.slow: 1 items in 0.10 seconds
[0, 1, 2]
```

Decorators work inside classes too. If you don't provide a name, a decent one will be inferred:

```
>>> class Database(object):
...     @instrument.all()
...     def batch_get(self, ids):
...         # you'd actually hit database here
...         for i in ids:
...             sleep(.1)
...             yield {"id":i, "square": i*i}
>>> database = Database()
>>> _ = database.batch_get([1, 2, 3, 9000])
>>> list(_)
__main__.Database.batch_get: 4 items in 0.40 seconds
[{'id': 1, 'square': 1}, {'id': 2, 'square': 4}, {'id': 3, 'square': 9}, {'id': 9000,
↪ 'square': 81000000}]
```

2.5 Reducers & Producers

`reducer()` and `producer()` are decorators for functions, *not* iterables.

The `reducer()` decorator measures functions that consume many items. Examples include aggregators or a `batch_save()`:

```
>>> @instrument.reducer()
... def sum_squares(L):
...     total = 0
...     for i in L:
...         sleep(.1)
...         total += i*i
...     return total
...
>>> sum_squares(range(5))
__main__.sum_squares: 5 items in 0.50 seconds
30
```

This works with functions taking a variable number of `*args` too:

```
>>> @instrument.reducer()
... def sum_squares2(*args):
...     total = 0
...     for i in args:
...         sleep(.1)
...         total += i*i
...     return total
...
>>> sum_squares2(*range(5))
__main__.sum_squares2: 5 items in 0.50 seconds
30
```

The `producer()` decorator measures a function that produces many items. This is similar to using `all()` as a decorator, but for functions that return lists instead of iterators. It works with any returned object supporting `len()`:

```
>>> @instrument.producer()
... def list_squares(N):
...     sleep(0.1 * N)
...     return [i * i for i in range(N)]
>>> list_squares(5)
__main__.list_squares: 5 items in 0.50 seconds
[0, 1, 4, 9, 16]
```

`reducer()` and `producer()` can be used inside classes:

```
>>> class Database(object):
...     @instrument.reducer()
...     def batch_save(self, rows):
...         for r in rows:
...             # you'd actually commit to your database here
...             sleep(0.1)
...
...     @instrument.reducer()
...     def batch_save2(self, *rows):
...         for r in rows:
...             # you'd actually commit to your database here
...             sleep(0.1)
...
...     @instrument.producer()
...     def dumb_query(self, x):
...         # you'd actually talk to your database here
...         sleep(0.1 * x)
...         return [{"id":i, "square": i*i} for i in range(x)]
...
... 
```

```
>>> rows = [{'id':i} for i in range(5)]
>>> database = Database()
>>> database.batch_save(rows)
__main__.Database.batch_save: 5 items in 0.50 seconds
>>> database.batch_save2(*rows)
__main__.Database.batch_save2: 5 items in 0.50 seconds
>>> database.dumb_query(3)
__main__.Database.dumb_query: 3 items in 0.30 seconds
[{'id': 0, 'square': 0}, {'id': 1, 'square': 1}, {'id': 2, 'square': 4}]
```


By default, metrics are *printed* to standard output. You can provide your own metric recording function. It should take three arguments: count of items, elapsed time in seconds, and *name*, which can be None:

```
>>> def my_metric(name, count, elapsed):
...     print("Iterable %s produced %d items in %d milliseconds"%(name, count,
↳int(round(elapsed*1000))))
...
>>> _ = instrument.all(math_is_hard(5), metric=my_metric, name="bogomips")
>>> list(_)
Iterable bogomips produced 5 items in 5000 milliseconds
[0, 1, 4, 9, 16]
```

Unless individually specified, metrics are reported using the global `instrument.default_metric()`. To change the active default, simply assign another metric function to this attribute. In general, you should configure your metric functions at program startup, **before** recording any metrics. `make_multi_metric()` creates a single metric function that records to several outputs.

3.1 Logging

`logging` writes metrics to a standard library logger, using the metric's name.

```
>>> from instrument.output.logging import log_metric
>>> _ = instrument.all(math_is_hard(5), metric=log_metric, name="bogomips")
>>> list(_)
INFO:instrument.bogomips:5 items in 5.00 seconds
[0, 1, 4, 9, 16]
```

3.2 Comma Separated

`csv` saves raw metrics as comma separated text files. This is useful for conducting external analysis. `csv` is thread-

safe; use under multiprocessing requires some care.

CSVFileMetric saves all metrics to a single file with three columns: metric name, item count & elapsed time. Create an instance of this class and pass its *CSVFileMetric.metric()* method to measurement functions. The *outfile* parameter controls where to write data; an existing file will be overwritten.

```
>>> from instrument.output.csv import CSVFileMetric
>>> csvfm = CSVFileMetric("/tmp/my_metrics_file.csv")
>>> _ = instrument.all(math_is_hard(5), metric=csvfm.metric, name="bogomips")
>>> list(_)
[0, 1, 4, 9, 16]
```

CSVDirMetric saves metrics to multiple files, named after each metric, with two columns: item count & elapsed time. This class is global to your program; do not manually create instances. Instead, use the classmethod *CSVDirMetric.metric()*. Set the class variable *outdir* to a directory in which to store files. The contents of this directory will be deleted on startup.

```
>>> from instrument.output.csv import CSVDirMetric
>>> CSVDirMetric.outdir = "/tmp/my_metrics_dir"
>>> _ = instrument.all(math_is_hard(5), metric=CSVDirMetric.metric, name="bogomips")
>>> list(_)
[0, 1, 4, 9, 16]
```

Both classes support a *dump_atexit* flag, which will register a handler to write data when the interpreter finishes execution. Set to false to manage yourself.

3.3 Summary Reports

table reports aggregate statistics and *plot* generates plots (graphs). These are useful for benchmarking or batch jobs; for live systems, *statsd* is a better choice. *table* and *plot* are threadsafe; use under multiprocessing requires some care.

TableMetric and *PlotMetric* are global to your program; do not manually create instances. Instead, use the classmethod *metric()*. The *dump_atexit* flag will register a handler to write data when the interpreter finishes execution. Set to false to manage yourself.

3.3.1 Tables

TableMetric prints pretty tables of aggregate population statistics. Set the class variable *outfile* to a file-like object (defaults to *stderr*):

```
>>> from instrument.output.table import TableMetric
>>> _ = instrument.all(math_is_hard(5), metric=TableMetric.metric, name="bogomips")
>>> list(_)
[0, 1, 4, 9, 16]
```

You'll get a nice table for output:




Name	Count Mean	Count Stddev	Elapsed Mean	Elapsed Stddev
alice	47.96	28.44	310.85	291.16
bob	50.08	28.84	333.98	297.11
charles	51.79	29.22	353.58	300.82

3.3.2 Plots

PlotMetric generates plots using matplotlib. Plots are saved to multiple files, named after each metric. Set the class variable `outdir` to a directory in which to store files. The contents of this directory will be deleted on startup.

3.4 statsd

For monitoring production systems, the `statsd_metric()` function can be used to record metrics to `statsd` and `graphite`. Each metric will generate two buckets: a count and a timing.



4.1 instrument

you are not expected to understand this implementation. that's why it has tests. the above-mentioned 'you' includes the author. :-}

`instrument.all` (*iterable=None*, *, *name=None*, *metric=<function call_default>*)
Measure total time and item count for consuming an iterable

Parameters

- **iterable** – any iterable
- **metric** (*function*) – f(name, count, total_time)
- **name** (*str*) – name for the metric

`instrument.block` (*, *name=None*, *metric=<function call_default>*, *count=1*)
Context manager to measure execution time of a block

Parameters

- **metric** (*function*) – f(name, 1, time)
- **name** (*str*) – name for the metric
- **count** (*int*) – user-supplied number of items, defaults to 1

`instrument.call_default` (*name*, *count*, *elapsed*)
call the `default_metric()` global in this module

Parameters

- **name** (*str*) – name of the metric
- **count** (*int*) – number of items
- **elapsed** (*float*) – time in seconds

class `instrument.counted_iterable` (*iterable*)

helper class that wraps an iterable and counts items

`instrument.each` (*iterable=None*, *, *name=None*, *metric=<function call_default>*)

Measure time elapsed to produce each item of an iterable

Parameters

- **iterable** – any iterable
- **metric** (*function*) – *f*(name, 1, time)
- **name** (*str*) – name for the metric

`instrument.first` (*iterable=None*, *, *name=None*, *metric=<function call_default>*)

Measure time elapsed to produce first item of an iterable

Parameters

- **iterable** – any iterable
- **metric** (*function*) – *f*(name, 1, time)
- **name** (*str*) – name for the metric

`instrument.function` (*, *name=None*, *metric=<function call_default>*)

Decorator to measure function execution time.

Parameters

- **metric** (*function*) – *f*(name, 1, total_time)
- **name** (*str*) – name for the metric

`instrument.producer` (*, *name=None*, *metric=<function call_default>*)

Decorator to measure a function that produces many items.

The function should return an object that supports `__len__` (ie, a list). If the function returns an iterator, use `iter()` instead.

Parameters

- **metric** (*function*) – *f*(name, count, total_time)
- **name** (*str*) – name for the metric

`instrument.reducer` (*, *name=None*, *metric=<function call_default>*)

Decorator to measure a function that consumes many items.

The wrapped `func` should take either a single `iterable` argument or `*args` (plus keyword arguments).

Parameters

- **metric** (*function*) – *f*(name, count, total_time)
- **name** (*str*) – name for the metric

4.2 instrument.output

`instrument.output.make_multi_metric` (**metrics*)

Make a new metric function that calls the supplied metrics

Parameters *metrics* (*functions*) – metric functions

Return type *function*

`instrument.output.print_metric` (*name*, *count*, *elapsed*)

A metric function that prints to standard output

Parameters

- **name** (*str*) – name of the metric
- **count** (*int*) – number of items
- **elapsed** (*float*) – time in seconds

`instrument.output.stderr_metric` (*name*, *count*, *elapsed*)

A metric function that prints to standard error

Parameters

- **name** (*str*) – name of the metric
- **count** (*int*) – number of items
- **elapsed** (*float*) – time in seconds

4.3 instrument.output.csv

write metrics to csv files

class `instrument.output.csv.CSVDirMetric` (*name*)

Write metrics to multiple CSV files

Do not create instances of this class directly. Simply pass the classmethod `metric()` to a measurement function. Output using `dump()`. These are the only public methods.

Each metric consumes one open file and 32K of memory while running.

Variables

- **dump_atexit** (*bool*) – automatically call `dump()` when the interpreter exits. Defaults to True.
- **outdir** (*str*) – directory to save CSV files in. Defaults to `./instrument_csv`.

classmethod `dump()`

Output all recorded metrics

classmethod `metric` (*name*, *count*, *elapsed*)

A metric function that writes multiple CSV files

Parameters

- **name** (*str*) – name of the metric
- **count** (*int*) – number of items
- **elapsed** (*float*) – time in seconds

class `instrument.output.csv.CSVFileMetric` (*outfile*=`'instrument.csv'`, *dump_atexit*=`True`)

Write metrics to a single CSV file

Pass the method `metric()` to a measurement function. Output using `dump()`. These are the only public methods.

Variables

- **outfile** – file to save to. Defaults to `./instrument.csv`.

- **dump_atexit** – automatically call `dump()` when the interpreter exits. Defaults to True.

dump()

Output all recorded metrics

metric (*name, count, elapsed*)

A metric function that writes a single CSV file

Parameters

- **name** (*str*) – name of the metric
- **count** (*int*) – number of items
- **elapsed** (*float*) – time in seconds

4.4 instrument.output.logging

save metrics to standard library logging

`instrument.output.logging.make_log_metric` (*level=20, msg='%d items in %.2f seconds'*)

Make a new metric function that logs at the given level

Parameters

- **level** (*int*) – logging level, defaults to `logging.INFO`
- **msg** (*string*) – logging message format string, taking count and elapsed

Return type *function*

4.5 instrument.output.table

print pretty tables of statistics

class `instrument.output.table.TableMetric` (*name*)

Print a table of statistics. See *NumpyMetric* for usage.

Variables `outfile` – output file. Defaults to `sys.stderr`.

4.6 instrument.output.plot

plot metrics with matplotlib

class `instrument.output.plot.PlotMetric` (*name*)

Plot graphs of metrics. See *NumpyMetric* for usage.

Variables `outdir` – directory to save plots in. Defaults to `./instrument_plots`.

4.7 instrument.output.statsd

save metrics to statsd

`instrument.output.statsd.statsd_metric` (*name, count, elapsed*)

Metric that records to statsd & graphite

4.8 instrument.output._numpy

numpy-based metrics

class `instrument.output._numpy.NumpyMetric` (*name*)

Base class for numpy-based metrics

Do not create instances of this class directly. Simply pass the classmethod `metric()` to a measurement function. Output using `dump()`. These are the only public methods. This is an abstract base class; you should use one of the concrete subclasses in this module instead.

Each metric consumes one open file and 32K of memory while running. Output requires enough memory to load all data points for each metric.

Variables `dump_atexit` (*bool*) – automatically call `dump()` when the interpreter exits. Defaults to True.

classmethod `dump` ()

Output all recorded metrics

classmethod `metric` (*name*, *count*, *elapsed*)

A metric function that buffers through numpy

Parameters

- **name** (*str*) – name of the metric
- **count** (*int*) – number of items
- **elapsed** (*float*) – time in seconds

Backwards Incompatibilities

5.1 0.5 -> 0.6

- drop support for Python 2.7
- *name* and *metric* are keyword only arguments
- various renames; see *Changelog* for details

5.2 0.4 -> 0.5

- main package renamed from `measure_it` to `instrument`
- prefixed `measure_iter`, etc. functions no longer available; use `instrument.iter` instead

5.3 0.3 -> 0.4

- remove deprecated `measure()`; use `measure_iter()` instead
- convert to package; `statsd_metric()` moved to its own module
- swap order of name & metric arguments

Some rough thoughts.

6.1 Automagic Instrumentation

Support for automagic instrumentation of popular 3rd-party packages:

- django, using introspection logic from [django-statsd](#)
- generic WSGI middleware. Possibly flask.
- any dbapi-compatible database, with names derived by parsing SQL for table/query type
- HTTP clients: [requests](#) and [urllib](#)
- storage engines: MongoDB, memcached, redis, Elastic Search. Possibly sqlalchemy

6.2 More metric backends

- lightweight running stats, based on forthcoming [stdlib statistics](#) module. May include support for periodic stats output, as a low-budget alternative to statsd.
- Prometheus, Datadog, etc.

6.3 Bonus features

- sampling & filtering for metric functions
- integration of nice Jupyter notebook for analysis

6.4 Modernization

- pypy test & support

See also: *Backwards Incompatibilities*.

7.1 0.6.0

- drop Python 2.7 support
- *all*, *each* and *first* can now be used directly as decorators
- add *logging_metric* & *stderr_metric* outputs
- move outputs to *instrument.output* subpackage
- split *instrument.output.numpy* to *plot* & *table* modules
- instrumenter renames: *reducer*, *producer*, *function*, *all*, *first*, *each*
- move *print_metric* and *make_multi_metric* to *instrument.output*

7.2 0.5.1

- rename project to `instrument` from `measure_it`
- update to modern tooling: `pytest`, `pipenv`, etc..
- improved testing: `tox`, `travis`
- fix to work with newer versions of `statsd`, including its `django` support

7.3 0.4

- add `default_metric`

- add `make_multi_metric`
- add `TableMetric` and `PlotMetric`, based on `numpy`
- add `CSVDirMetric` and `CSVFileMetric`
- `measure_block` supports user-supplied count

7.4 0.3

- add `measure_first`, `measure_produce`, `measure_func`, `measure_block`
- rename `measure` to `measure_iter` and deprecate old name

7.5 0.2

- add `measure_reduce`

7.6 0.1

Initial release

i

instrument, 15
instrument.output, 11
instrument.output._numpy, 19
instrument.output.csv, 17
instrument.output.logging, 18
instrument.output.plot, 18
instrument.output.statsd, 18
instrument.output.table, 18

A

all() (in module instrument), 15

B

block() (in module instrument), 15

C

call_default() (in module instrument), 15

counted_iterable (class in instrument), 15

CSVDirMetric (class in instrument.output.csv), 17

CSVFileMetric (class in instrument.output.csv), 17

D

dump() (instrument.output._numpy.NumpyMetric class method), 19

dump() (instrument.output.csv.CSVDirMetric class method), 17

dump() (instrument.output.csv.CSVFileMetric method), 18

E

each() (in module instrument), 16

F

first() (in module instrument), 16

function() (in module instrument), 16

I

instrument (module), 15

instrument.output (module), 11, 16

instrument.output._numpy (module), 19

instrument.output.csv (module), 17

instrument.output.logging (module), 18

instrument.output.plot (module), 18

instrument.output.statsd (module), 18

instrument.output.table (module), 18

M

make_log_metric() (in module instrument.output.logging), 18

make_multi_metric() (in module instrument.output), 16

metric() (instrument.output._numpy.NumpyMetric class method), 19

metric() (instrument.output.csv.CSVDirMetric class method), 17

metric() (instrument.output.csv.CSVFileMetric method), 18

N

NumpyMetric (class in instrument.output._numpy), 19

P

PlotMetric (class in instrument.output.plot), 18

print_metric() (in module instrument.output), 16

producer() (in module instrument), 16

R

reducer() (in module instrument), 16

S

statsd_metric() (in module instrument.output.statsd), 18

stderr_metric() (in module instrument.output), 17

T

TableMetric (class in instrument.output.table), 18