
Statsd Metrics Documentation

Release 0.1.0

Farzad Ghanei

November 16, 2016

1 Metrics	3
1.1 metrics – Metric classes and helper functions	4
2 Client	7
2.1 client – Statsd client	7
2.2 client.tcp – Statsd client sending metrics over TCP	9
3 Introduction	11
3.1 Metric Classes	11
3.2 Clients	11
4 Installation	13
4.1 Dependencies	13
5 License	15
6 Development	17
Python Module Index	19

Contents:

Metrics

Define the data types used in Statsd. Each data type is defined as a class, supported data types are:

- *Counter*
- *Timer*
- *Gauge*
- *Set*
- *GaugeDelta*

Note: The metric classes and helper functions are available from the package directly, but internally they are defined in `metrics` module. So there is no need to import the `metrics` module directly, unless you're trying to access those objects that are not used regularly and hence are not exported, like the `AbstractMetric` class.

Each metric requires a name and a value.

```
from statsdmetrics import Counter, Timer
counter = Counter('event.login', 1)
timer = Timer('db.query.user', 10)
```

An optional sample rate can be specified for the metrics. Sample rate is used by the client and the server to help to reduce network traffic, or reduce the load on the server.

```
>>> from statsdmetrics import Counter
>>> counter = Counter('event.login', 1, 0.2)
>>> counter.name
'event.login'
>>> counter.count
1
>>> counter.sample_rate
0.2
```

All metrics have `name` and `sample_rate` properties, but they store their value in different properties.

Metrics provide `to_request()` method to create the proper value used to send the metric to the server.

```
>>> from statsdmetrics import Counter, Timer, Gauge, Set, GaugeDelta
>>> counter = Counter('event.login', 1, 0.2)
>>> counter.to_request()
'event.login:1|c|@0.2'
>>> timer = Timer('db.query.user', 10, 0.5)
>>> timer.to_request()
```

```
'db.query.user:10|ms@0.5'
>>> gauge = Gauge('memory', 20480)
>>> gauge.to_request()
'memory:20480|g'
>>> set_ = Set('unique.users', 'first')
>>> set_.to_request()
'unique.users:first|s'
>>> delta = GaugeDelta('memory', 128)
>>> delta.to_request()
'memory:+128|g'
>>> delta.delta = -256
>>> delta.to_request()
'memory:-256|g'
```

1.1 metrics – Metric classes and helper functions

1.1.1 Metric Classes

class metrics.AbstractMetric

Abstract class that all metric classes would extend from

name

the name of the metric

sample_rate

the rate of sampling that the client considers when sending metrics

to_request () → str

return the string that is used in the Statsd request to send the metric

class metrics.Counter (name, count[, sample_rate])

A metric to count events

count

current count of events being reported via the metric

class metrics.Timer (name, milliseconds[, sample_rate])

A metric for timing durations, in milliseconds.

milliseconds

number of milliseconds for the duration

class metrics.Gauge (name, value[, sample_rate])

Any arbitrary value, like the memory usage in bytes.

value

the value of the metric

class metrics.Set (name, value[, sample_rate])

A set of unique values counted on the server side for each sampling period. Technically the value could be anything that can be serialized to a string (to be sent on the request).

value

the value of the metric

class metrics.GaugeDelta (name, delta[, sample_rate])

A value change in a gauge, could be a positive or negative numeric value.

delta

the difference in the value of the gauge

1.1.2 Module functions

`metrics.normalize_metric_name(name) → str`

normalize a metric name, removing characters that might not be welcome by common backends.

```
>>> from statsdmetrics import normalize_metric_name
>>> normalize_metric_name("will replace some, and $remove! others*")
'will_replace_some_and_remove_others'
```

If passed argument is not a string, an `TypeError` is raised.

`metrics.parse_metric_from_request(request) → str`

parse a metric object from a request string.

```
>>> from statsdmetrics import parse_metric_from_request
>>> metric = parse_metric_from_request("memory:2048|g")
>>> type(metric)
<class 'statsdmetrics.metrics.Gauge'>
>>> metric.name, metric.value, metric.sample_rate
('memory', 2048.0, 1)
>>> metric = parse_metric_from_request('event.connections:-2|c|@0.6')
>>> type(metric)
<class 'statsdmetrics.metrics.Counter'>
>>> metric.name, metric.count, metric.sample_rate
('event.connections', -2, 0.6)
```

If the request is invalid, a `ValueError` is raised.

Client

To send the metrics to Statsd server, client classes are available in the `client` package and `client.tcp` module.

2.1 `client` – Statsd client

```
class client.Client(host[, port=8125][, prefix=''])
    Default Statsd client that sends each metric in a separate UDP request

    host
        the host name (or IP address) of Statsd server. This property is readonly.

    port
        the port number of Statsd server. This property is readonly.

    prefix
        the default prefix for all metric names sent from the client

    remote_address
        tuple of resolved server address (host, port). This property is readonly.

    increment(name[, count=1][, rate=1])
        Increase a Counter metric by count with an integer value. An optional sample rate can be specified.

    decrement(name[, count=1][, rate=1])
        Decrease a Counter metric by count with an integer value. An optional sample rate can be specified.

    timing(name, milliseconds[, rate=1])
        Send a Timer metric for the duration of a task in milliseconds. The milliseconds should be a non-negative numeric value. An optional sample rate can be specified.

    gauge(name, value[, rate=1])
        Send a Gauge metric with the specified value. The value should be a none-negative numeric value. An optional sample rate can be specified.

    set(name, value[, rate=1])
        Send a Set metric with the specified value. The server will count the number of unique values during each sampling period. The value could be any value that can be converted to a string. An optional sample rate can be specified.

    gauge_delta(name, delta[, rate=1])
        Send a GaugeDelta metric with the specified delta. The delta should be a numeric value. An optional sample rate can be specified.
```

batch_client([size=512])

Create a `BatchClient` object, using the same configurations of current client. This batch client could be used as a context manager in a `with` statement. After the `with` block when the context manager exits, all the metrics are flushed to the server in batch requests.

Note: Most Statsd servers do not apply the sample rate on timing metrics calculated results (mean, percentile, max, min), gauge or set metrics, but they take the rate into account for the number of received samples. Some statsd servers totally ignore the sample rate for metrics other than counters.

2.1.1 Examples

```
from statsdmetrics.client import Client
client = Client("stats.example.org")
client.increment("login")
client.timing("db.search.username", 3500)
client.prefix = "other"
client.gauge_delta("memory", -256)
client.decrement(name="connections", count=2)
```

```
from statsdmetrics.client import Client

client = Client("stats.example.org")
with client.batch_client() as batch_client:
    batch_client.increment("login")
    batch_client.decrement(name="connections", count=2)
    batch_client.timing("db.search.username", 3500)
# now all metrics are flushed automatically in batch requests
```

class client.BatchClient(host[, port=8125][, prefix=''][, batch_size=512])

Statsd client that buffers all metrics and sends them in batch requests over UDP when instructed to flush the metrics explicitly.

Each UDP request might contain multiple metrics, but limited to a certain batch size to avoid UDP fragmentation.

The size of batch requests is not the fixed size of the requests, since metrics can not be broken into multiple requests. So if adding a new metric overflows this size, then that metric will be sent in a new batch request.

batch_size

Size of each batch request. This property is **readonly**.

clear()

Clear buffered metrics

flush()

Send the buffered metrics in batch requests.

unit_client()

Create a `Client` object, using the same configurations of current batch client to send the metrics on each request. The client uses the same resources as the batch client.

```
from statsdmetrics.client import BatchClient

client = BatchClient("stats.example.org")
client.set("unique.ip_address", "10.10.10.1")
client.gauge("memory", 20480)
client.flush() # sends one UDP packet to remote server, carrying both metrics
```

2.2 client.tcp – Statsd client sending metrics over TCP

```
class client.tcp.TCPClient(host[, port=8125][, prefix=''])
    Statsd client that sends each metric in separate requests over TCP.

    Provides the same interface as Client.
```

2.2.1 Examples

```
from statsdmetrics.client.tcp import TCPClient
client = TCPClient("stats.example.org")
client.increment("login")
client.timing("db.search.username", 3500)
client.prefix = "other"
client.gauge_delta("memory", -256)
client.decrement(name="connections", count=2)
```

```
from statsdmetrics.client.tcp import TCPClient

client = TCPClient("stats.example.org")
with client.batch_client() as batch_client:
    batch_client.increment("login")
    batch_client.decrement(name="connections", count=2)
    batch_client.timing("db.search.username", 3500)
# now all metrics are flushed automatically in batch requests
```

```
class client.tcp.TCPBatchClient(host[, port=8125][, prefix=''][, batch_size=512])
    Statsd client that buffers all metrics and sends them in batch requests over TCP when instructed to flush the metrics explicitly.
```

Provides the same interface as *BatchClient*.

```
from statsdmetrics.client.tcp import TCPBatchClient

client = TCPBatchClient("stats.example.org")
client.set("unique.ip_address", "10.10.10.1")
client.gauge("memory", 20480)
client.flush() # sends one TCP packet to remote server, carrying both metrics
```

Introduction

Statsd metrics is an API to create, parse or send metrics to a Statsd server.

3.1 Metric Classes

Metric classes are used to define the data type and values for each metric, and to create the contents of the request that will be sent to the Statsd server.

Available metrics:

- *Counter*
- *Timer*
- *Gauge*
- *Set*
- *GaugeDelta*

The *metrics* module also provides helper functions to normalize metric names, and a parse a Statsd request and return the corresponding metric object. This could be used on the server side to parse the received requests.

3.2 Clients

- *Client*: Default client, sends request on each call using UDP
- *BatchClient*: Buffers metrics and flushes them in batch requests using UDP
- *TCPClient*: Sends request on each call using TCP
- *TCPBatchClient*: Buffers metrics and flushes them in batch requests using TCP

Installation

```
pip install statsdmetrics
```

4.1 Dependencies

There are no specific dependencies, it runs on Python 2.7+ (CPython 2.7, 3.2, 3.3 3.4 and 3.5, PyPy 2.6 and PyPy3 2.4, and Jython 2.7 are tested)

However on development (and test) environment `mock` is required, and `distutilazy` (or `setuptools` as a fallback) is used to run the tests.

```
# on dev/test env
pip install -r requirements-dev.txt
```


License

Statsd metrics is released under the terms of the [MIT license](#).

Development

- Code is on [GitHub](#)
- Documentations are on [Read The Docs](#)

C

`client`, [7](#)
`client.tcp`, [9](#)

m

`metrics`, [4](#)

A

AbstractMetric (class in metrics), [4](#)
AbstractMetric.name (in module metrics), [4](#)
AbstractMetric.sample_rate (in module metrics), [4](#)

B

batch_client() (client.Client method), [7](#)
BatchClient (class in client), [8](#)
BatchClient.batch_size (in module client), [8](#)

C

clear() (client.BatchClient method), [8](#)
Client (class in client), [7](#)
client (module), [7](#)
Client.host (in module client), [7](#)
Client.port (in module client), [7](#)
Client.prefix (in module client), [7](#)
Client.remote_address (in module client), [7](#)
client.tcp (module), [9](#)
Counter (class in metrics), [4](#)
Counter.count (in module metrics), [4](#)

D

decrement() (client.Client method), [7](#)

F

flush() (client.BatchClient method), [8](#)

G

Gauge (class in metrics), [4](#)
gauge() (client.Client method), [7](#)
Gauge.value (in module metrics), [4](#)
gauge_delta() (client.Client method), [7](#)
GaugeDelta (class in metrics), [4](#)
GaugeDelta.delta (in module metrics), [4](#)

I

increment() (client.Client method), [7](#)

M

metrics (module), [4](#)

N

normalize_metric_name() (in module metrics), [5](#)

P

parse_metric_from_request() (in module metrics), [5](#)

S

Set (class in metrics), [4](#)
set() (client.Client method), [7](#)
Set.value (in module metrics), [4](#)

T

TCPBatchClient (class in client.tcp), [9](#)
TCPClient (class in client.tcp), [9](#)
Timer (class in metrics), [4](#)
Timer.milliseconds (in module metrics), [4](#)
timing() (client.Client method), [7](#)
to_request() (metrics.AbstractMetric method), [4](#)

U

unit_client() (client.BatchClient method), [8](#)