

---

# Tensor Documentation

*Release 1.0*

**Colin Alston**

January 06, 2017



<b>1 Getting started</b>	<b>3</b>
1.1 Installation . . . . .	3
1.2 Creating a configuration file . . . . .	3
1.3 Using outputs . . . . .	4
1.4 Using Elasticsearch instead . . . . .	4
1.5 Using sources . . . . .	4
1.6 Configuration . . . . .	5
1.7 State triggers . . . . .	5
1.8 Routing sources . . . . .	5
1.9 Remote SSH checks . . . . .	6
1.10 Starting Tensor . . . . .	7
<b>2 Sources</b>	<b>9</b>
2.1 Introduction . . . . .	9
2.2 Writing your own sources . . . . .	9
2.3 Using custom sources . . . . .	10
2.4 Handling asynchronous tasks . . . . .	10
2.5 Thinking outside the box . . . . .	11
<b>3 Outputs</b>	<b>13</b>
3.1 Introduction . . . . .	13
3.2 Using TLS with Riemann . . . . .	13
3.3 Writing your own outputs . . . . .	13
<b>4 Example configurations</b>	<b>15</b>
4.1 Replacing Munin . . . . .	15
<b>5 tensor</b>	<b>23</b>
5.1 tensor.aggregators . . . . .	23
5.2 tensor.interfaces . . . . .	23
5.3 tensor.objects . . . . .	23
5.4 tensor.service . . . . .	24
5.5 tensor.utils . . . . .	25
<b>6 tensor.protocol</b>	<b>27</b>
6.1 tensor.protocol.elasticsearch . . . . .	27
6.2 tensor.protocol.icmp . . . . .	27
6.3 tensor.protocol.riemann . . . . .	27
6.4 tensor.protocol.ssh . . . . .	28

6.5	<code>tensor.protocol.sflow</code>	28
<b>7</b>	<b><code>tensor.logs</code></b>	<b>29</b>
7.1	<code>tensor.logs.follower</code>	29
7.2	<code>tensor.logs.parsers</code>	29
<b>8</b>	<b><code>tensor.sources</code></b>	<b>31</b>
8.1	<code>tensor.sources.database.postgresql</code>	31
8.2	<code>tensor.sources.database.elasticsearch</code>	31
8.3	<code>tensor.sources.database.memcache</code>	32
8.4	<code>tensor.sources.docker</code>	32
8.5	<code>tensor.sources.haproxy</code>	33
8.6	<code>tensor.sources.generator</code>	33
8.7	<code>tensor.sources.linux</code>	33
8.8	<code>tensor.sources.media</code>	36
8.9	<code>tensor.sources.munin</code>	36
8.10	<code>tensor.sources.network</code>	36
8.11	<code>tensor.sources.nginx</code>	37
8.12	<code>tensor.sources.python</code>	38
8.13	<code>tensor.sources.rabbitmq</code>	38
8.14	<code>tensor.sources.redis</code>	39
8.15	<code>tensor.sources.riak</code>	39
8.16	<code>tensor.sources.riemann</code>	40
8.17	<code>tensor.sources.sflow</code>	40
8.18	<code>tensor.sources.snmp</code>	40
<b>9</b>	<b><code>tensor.outputs</code></b>	<b>43</b>
9.1	<code>tensor.outputs.riemann</code>	43
9.2	<code>tensor.outputs.elasticsearch</code>	44
<b>10</b>	<b>Indices and tables</b>	<b>47</b>
	<b>Python Module Index</b>	<b>49</b>

Tensor is a modular gateway and event router built using the Twisted framework. It can be used as a monitoring agent and ETL framework for a wide range of applications, and provides a simple yet “powerful” plugin mechanism to expand its capabilities.

Contents:



---

## Getting started

---

### 1.1 Installation

Tensor can be installed from PyPi with pip

```
$ pip install tensor
```

This will also install Twisted, protobuf and PyYAML

Or you can use the .deb package. Let the latest release from <https://github.com/calston/tensor/releases/latest>

```
$ aptitude install python-twisted python-protobuf python-yaml
$ wget https://github.com/calston/tensor/releases/download/0.3.0/tensor_0.3.0_amd64.deb
$ dpkg -i tensor_0.3.0_amd64.deb
```

This also gives you an init script and default config in /etc/tensor/

### 1.2 Creating a configuration file

Tensor uses a simple YAML configuration file

Some first basic requirements are the Riemann server and port (defaults to localhost:5555) and the queue interval:

```
server: localhost
port: 5555
interval: 1.0
proto: udp
```

Tensors checks are Python classes (called sources) which are instantiated with the configuration block which defines them. Rather than being one-shot scripts, a source object remains in memory with its own timer which adds events to a queue. The *interval* defined above is the rate at which these events are rolled up into a message and sent to Riemann.

It is important then that *interval* is set to a value appropriate to how frequently you want to see them in Riemann, as well as the rate at which they collect metrics from the system. All *interval* attributes are floating point in seconds, this means you can check (and send to Riemann) at rates well below 1 second.

Riemann is the default output in this configuration, but there are others.

## 1.3 Using outputs

A better and more explicit means of configuring where events go is to use the *output* framework. Currently there is support for Elasticsearch and Riemann, but this can easily be extended.

You can configure multiple outputs which receive a copy of every message for example

```
outputs:
  - output: tensor.outputs.riemann.RiemannTCP
    server: localhost
    port: 5555
```

If you enable multiple outputs then the global *server*, *port* and *proto* options will go un-used and the default Riemann TCP transport won't start.

You can configure as many outputs as you like, or create your own.

## 1.4 Using Elasticsearch instead

You may wish to output events directly into Elasticsearch in which case your configuration would look like this

```
outputs:
  - output: tensor.outputs.elasticsearch.ElasticSearch
    server: 127.0.0.1
    port: 9200
```

## 1.5 Using sources

To configure the basic CPU usage source add it to the *sources* list in the config file

```
sources:
  - service: cpu
    source: tensor.sources.linux.basic.CPU
    interval: 1.0
    warning: {
      cpu: "> 0.5"
    }
    critical: {
      cpu: "> 0.8"
    }
```

This will measure the CPU from /proc/stat every second, with a warning state if the value exceeds 50%, and a critical state if it exceeds 80%

The *service* attribute can be any string you like, populating the *service* field in Riemann. The logical expression to raise the state of the event is (eg. *critical*) is assigned to a key which matches the service name.

Sources may return more than one metric, in which case it will add a prefix to the service. The state expression must correspond to that as well.

For example, the Ping check returns both latency and packet loss:

```
service: googledns
source: tensor.sources.network.Ping
interval: 60.0
destination: 8.8.8.8
```

```
critical: {
    googledns.latency: "> 100",
    googledns.loss: "> 0"
}
```

This will ping 8.8.8.8 every 60 seconds and raise a critical alert for the latency metric if it exceeds 100ms, and the packet loss metric if there is any at all.

## 1.6 Configuration

Sources can contain any number of configuration attributes which vary between them. All sources take the following options though

service	mandatory	Service name after which extra metric names are appended, dot separated
interval	depends	Clock tick interval, for sources which implement a polling clock
ttl	optional	TTL for metric expiry in Riemann
hostname	optional	Hostname to tag this service with. Defaults to system FQDN but can be overridden.
tags	optional	Comma separated list of tags for metrics

## 1.7 State triggers

*critical* and *warning* matches can also be a regular expression for sources which output keys for different devices and metrics:

```
service: network
source: tensor.sources.linux.basic.Network
...
critical: {
    network.\w+.tx_packets: "> 1000",
}
```

## 1.8 Routing sources

Since multiple outputs can be added, Tensor events can be routed from sources to specific outputs or multiple outputs. By default events are routed to all outputs.

To enable routing, outputs need a unique *name* attribute:

```
outputs:
  - output: tensor.outputs.riemann.RiemannTCP
    name: riemann1
    server: riemann1.acme.com
    port: 5555

  - output: tensor.outputs.riemann.RiemannTCP
    name: riemann2
    server: riemann2.acme.com
    port: 5555

  - output: tensor.outputs.riemann.RiemannUDP
    name: riemannudp
    server: riemann1.acme.com
```

```
port: 5555

sources:
  - service: cpu1
    source: tensor.sources.linux.basic.CPU
    interval: 1.0
    route: riemannudp

  - service: cpu2
    source: tensor.sources.linux.basic.CPU
    interval: 1.0
    route:
      - riemann1
      - riemann2
```

The *route* attribute can also accept a list of output names. The above configuration would route *cpu1* metrics to the UDP output, and the *cpu2* metrics to both *riemann1* and *riemann2* TCP outputs.

## 1.9 Remote SSH checks

A new feature in Tensor is the ability to perform checks on a remote device using SSH. This is currently only supported by certain sources.

To perform a check over SSH we need an *ssh\_host* which defaults to the check hostname, *ssh\_username*, and one of *ssh\_key*, *ssh\_keyfile* or *ssh\_password*. All of these except the *ssh\_host* parameter can be specified globally and/or on a specific source to override the global configuration.

*ssh\_key* allows providing a private key in a YAML text blob. If *ssh\_key* or *ssh\_keyfile* is password encrypted then *ssh\_keypass* can be set to that in plain text - although this isn't really recommendable.

Example

```
ssh_username: tensor
ssh_key: |
  Proc-Type: 4,ENCRYPTED
  DEK-Info: AES-128-CBC,A6588464A721D661311DBCE44C76337E

/bqr0AEIbiWubFiPEcdlNw8WdDrFqELOCzo78ohtDX/2HJhkMCCTauv46is5UCvj
pweYupJQgZZ9g+6rKLhTo6d0VYwaSOuR6OJWEecIr7quyQBgCPOvun2fVGrx6/7U
D9HiXbdBDVc4vcEUviZu5V+E9xLmP9GteD1OR7Tfr1AqFMPzHVvDE9UxrzEacfY4
KP87KP6x+8so5KvZSJkisczc+Jbt+P1ZisDwX9BCHJNmAYYFRm2umY7sCmLNmeoc
Y95E6Tmpze4J1559mLM7nuzOpnnFEii4pA5H7unMUCa9AwkLLYLOV7N8iRETgG0R
snvH5uiVRqEB84ypItCZF+Nk5Y0/WPSWPdq/bhwyQeodEPj1IfiHKzDf9GuuT9E1
Q4dGxA0mLOKMqPDJGGc7mwTTN5icjz94gsLTfI1me1qzTzxdko/BMqsmPSUbknxs
wgkofT+48L00HL9zq0quHkgjoTe1Wud8tI4mG0T19BTFE9P0t1fdJNoEQ1dk9RcR
UkhjmBuN3h8r9w9EVugAvbp/c7SQILXEJ6QZK2NMz001SA5Tv7hmDh1J01cIF1zb
VI+r1xly/riDN6U9w35vOZEzK13qYbAXrnRteo7MEYvc/BahvxBP0ZEGRXeoNfAj
JLvbkhBUVylcH5fGs2SYiwUEKBx5nLL5NeNI1ymRKbsyJ3oTKZU+PQhfarcJD2r
u/dZoDb/AEjxCKaM1EaDG590Bjc6Zxc1ZkF6gSK27iJRP5CCj5XoD7kIpmpZFE+gc
KpVNHH6ef2ptOngkEDUyTmZ7z181VCeC4sBPzrLPDnWB+cie+19/cJDJpRz0n0j
qMkh7MY+FQ8t0AopFAy7r50nV5F1Gt9rG7YaW08j5Lv3TsPPDOxFk5IoB6AtRpr8
tSQCCyCcdHkD3M1wI/PD9bejuELaDG8PaVzOuj5rVyh+saZQeD9GmegsuBkDgb4g
C0tzWOQ1H0ii478rbQAxwsOEMdR5lxEFoo8mC0p4mnWjt12DzJQorQC/fjbRRv7z
vfJamXvfEuHj3NPP9cumrskBtD+kRz/c2zgVJ8vvwRgNPazdfJqGYjmFB0loVVyuu
x+hBHOD5zyMPFrJW9MNDTiTEaQREaje5tUOfNoA1Wa4s2bVLnhHCXdMSWmiDmJQp
HEYAIIZ120JhMe8V431t6dBx+nutApzParWqET5D0DWvlurDWFrHMnazh164RqsGu
E4Dg6ZsRnI+PEJmroia6gYEScUfd5QSUEbxTeLhNzo1Kf5JRBW93NNxhAzn9ZJ90
```

```

2bjvkHOJ1ADnfON5TsPgroXX95P/9V8DWUT+/skelFw43V1pIT+PtraYqrlyvow+
uJMA2q9sRLzXnFb2vg7JdD1XA4f2eUBwzb7q8wSuQexSERWaTx5uAERDnGAWyaN2
3xCY18CTff70xN7j39hG/pI0ghRLGVBmCJ5NRzNZ80SPBE/nzYy/X6pGV+vsjPoZ
S3dBmv1BV/HzB41jsO2pI/VjCJVNZdOWDzy18GQ/jt8/xH8R9Ld6/6tuS0HbiefS
ZefHS5wV1KNZBK+vh08HvX/AY9WBHFH+DEbrpymn/9oAKVmH+f73ADqVOanMPk0
-----END RSA PRIVATE KEY-----
ssh_keypass: testtest

sources:
  - service: load
    use_ssh: True
    ssh_host: myremotebox.acme.net
    source: tensor.sources.linux.basic.LoadAverage
    interval: 2.0

```

Note: Currently Tensor will not perform any host key checking.

## 1.10 Starting Tensor

To start Tensor, simply use twistd to run the service and pass a config file:

```
twistd -n tensor -c tensor.yml
```

If you're using the Debian package then an init script is included.



---

## Sources

---

### 2.1 Introduction

Sources are Python objects which subclass `tensor.objects.Source`. They are constructed with a dictionary parsed from the YAML configuration block which defines them, and as such can read any attributes from that either optional or mandatory.

Since sources are constructed at startup time they can retain any required state, for example the last metric value to report rates of change or for any other purpose. However since a Tensor process might be running many checks a source should not use an excessive amount of memory.

The `source` configuration option is passed a string representing an object in much the same way as you would import it in a python module. The final class name is split from this string. For example specifying:

```
source: tensor.sources.network.Ping
```

is equivalent to:

```
from tensor.sources.network import Ping
```

### 2.2 Writing your own sources

A source class must subclass `tensor.objects.Source` and also implement the interface `tensor.interfaces.ITensorSource`

The source must have a `get` method which returns a `tensor.objects.Event` object. The `Source` parent class provides a helper method `createEvent` which performs the metric level checking (evaluating the simple logical statement in the configuration), sets the correct service name and handles prefixing service names.

A “Hello world” source:

```
from zope.interface import implementer

from tensor.interfaces import ITensorSource
from tensor.objects import Source

@implementer(ITensorSource)
class HelloWorld(Source):

    def get(self):
        return self.createEvent('ok', 'Hello world!', 0)
```

To hold some state, you can re-implement the `__init__` method, as long as the arguments remain the same.

Extending the above example to create a simple flip-flop metric event:

```
from zope.interface import implementer

from tensor.interfaces import ITensorSource
from tensor.objects import Source

@implementer(ITensorSource)
class HelloWorld(Source):
    def __init__(self, *a):
        Source.__init__(self, *a)
        self.bit = False

    def get(self):
        self.bit = not self.bit
        return self.createEvent('ok', 'Hello world!', self.bit and 0.0 or 1.0)
```

You could then place this in a Python module like `hello.py` and as long as it's in the Python path for Tensor it can be used as a source with `hello.HelloWorld`

A list of events can also be returned but be careful of overwhelming the output buffer, and if you need to produce lots of metrics it may be worthwhile to return nothing from `get` and call `self.queueBack` as needed.

## 2.3 Using custom sources

When a source is specified, eg

```
source: tensor.sources.network.Ping
```

Tensor will import and instantiate the `Ping` class from `tensor.sources.network`. Consequently a source can be any installed Python module.

For the sake of convenience, however, Tensor also appends `/var/lib/tensor` to the Python path. This means you can easily create, test and distribute sources in that directory.

For example, create the above `hello.py` file and place it in `/var/lib/tensor` then use the configuration

```
source: hello.HelloWorld
```

You can also always submit Github pull request with sources to have them added to Tensor for others to benefit from!

## 2.4 Handling asynchronous tasks

Since Tensor is written using the Twisted asynchronous framework, sources can (and in most cases *must*) make full use of it to implement network checks, or execute other processes.

The simplest example of a source which executes an external process is the `ProcessCount` check:

```
from zope.interface import implementer

from twisted.internet import defer

from tensor.interfaces import ITensorSource
from tensor.objects import Source
from tensor.utils import fork
```

```

@implementer(ITensorSource)
class ProcessCount(Source):
    @defer.inlineCallbacks
    def get(self):
        out, err, code = yield fork('/bin/ps', args=('-e',))

        count = len(out.strip('\n').split('\n'))

        defer.returnValue(
            self.createEvent('ok', 'Process count %s' % (count), count)
        )

```

For more information please read the Twisted documentation at <https://twistedmatrix.com/trac/wiki/Documentation>  
The `tensor.utils.fork()` method returns a deferred which can timeout after a specified time.

## 2.5 Thinking outside the box

Historically monitoring systems are poorly architected, and terribly inflexible. To demonstrate how Tensor offers a different concept to the boring status quo it's interesting to note that there is nothing preventing you from starting a listening service directly within a source which processes and relays events to Riemann implementing some protocol.

Here is an example of a source which listens for TCP connections to port 8000, accepting any number on a line and passing that to the event queue:

```

from twisted.internet.protocol import Factory
from twisted.protocols.basic import LineReceiver
from twisted.internet import reactor

from zope.interface import implementer

from tensor.interfaces import ITensorSource
from tensor.objects import Source

class Numbers(LineReceiver):
    def __init__(self, source):
        self.source = source

    def lineReceived(self, line):
        """
        Send any numbers received back to the Tensor queue
        """
        print repr(line)
        try:
            num = float(line)
            self.source.queueBack(
                self.source.createEvent('ok', 'Number: %s' % num, num)
            )
        except:
            pass

class NumbersFactory(Factory):
    def __init__(self, source):
        self.source = source

    def buildProtocol(self, addr):

```

```
    return Numbers(self.source)

@implementer(ITensorSource)
class NumberProxy(Source):
    def startTimer(self):
        # Override starting the source timer, we don't need it
        f = NumbersFactory(self)
        reactor.listenTCP(8000, f)

    def get(self):
        # Implement the get method, but we can ignore it
        pass
```

---

## Outputs

---

### 3.1 Introduction

Outputs are Python objects which subclass `tensor.objects.Output`. They are constructed with a dictionary parsed from the YAML configuration block which defines them, and as such can read any attributes from that either optional or mandatory.

Since outputs are constructed at startup time they can retain any required state. A copy of the queue is passed to all `:method:`tensor.objects.Output.eventsReceived`` calls which happen at each queue `interval` config setting as the queue is emptied. This list of `tensor.objects.Event` objects must not be altered by the output.

The `output` configuration option is passed a string representing an object the same way as `sources` configurations are. For example this outputs events to Riemann over TCP:

```
outputs:
  - output: tensor.outputs.riemann.RiemannTCP
    server: 127.0.0.1
    port: 5555
```

### 3.2 Using TLS with Riemann

The RiemannTCP output also supports TLS, which can make use of Puppet certs for convenience

```
outputs:
  - output: tensor.outputs.riemann.RiemannTCP
    server: 127.0.0.1
    port: 5554
    tls: true
    cert: /var/lib/puppet/ssl/certs/test.acme.com.pem
    key: /var/lib/puppet/ssl/private_keys/test.acme.com.pem
```

### 3.3 Writing your own outputs

An output class should subclass `tensor.objects.Output`.

The output can implement a `createClient` method which starts the output in whatever way necessary and can be a deferred. The output must also have a `eventsReceived` method which takes a list of `tensor.objects.Event` objects and process them accordingly, it can also be a deferred.

An example logging source:

```
from twisted.internet import reactor, defer
from twisted.python import log

from tensor.objects import Output

class Logger(Output):
    def eventsReceived(self, events):
        log.msg("Events dequeued: %s" % len(events))
```

If you save this as *test.py* the basic configuration you need is simply

```
outputs:
    - output: tensor.outputs.riemann.RiemannUDP
      server: localhost
      port: 5555

    - output: test.Logger
```

You should now see how many events are exiting in the Tensor log file

```
2014-10-24 15:35:27+0200 [-] Starting protocol <tensor.protocol.riemann.RiemannUDP object at 0x7f3b5a
2014-10-24 15:35:28+0200 [-] Events dequeued: 7
2014-10-24 15:35:29+0200 [-] Events dequeued: 2
2014-10-24 15:35:30+0200 [-] Events dequeued: 3
```

Events can be routed in different ways to outputs, see the Getting started guide for more details

---

## Example configurations

---

### 4.1 Replacing Munin

The first step is to create a TRIG stack (Tensor Riemann InfluxDB Grafana).

#### 4.1.1 Step 1: Install Riemann

```
$ wget http://aphyr.com/riemann/riemann_0.2.6_all.deb  
$ aptitude install openjdk-7-jre  
$ dpkg -i riemann_0.2.6_all.deb
```

#### 4.1.2 Step 2: Install InfluxDB

```
$ wget http://s3.amazonaws.com/influxdb/influxdb_latest_amd64.deb  
$ sudo dpkg -i influxdb_latest_amd64.deb
```

Start InfluxDB, then quickly change the root/root default password because it also defaults to listening on all interfaces and apparently this is not important enough for them to fix.

Create a *riemann* and *grafana* database, and some users for them

```
$ curl -X POST 'http://localhost:8086/db?u=root&p=root' \  
-d '{"name": "riemann"}'  
$ curl -X POST 'http://localhost:8086/db?u=root&p=root' \  
-d '{"name": "grafana"}'  
$ curl -X POST 'http://localhost:8086/db/riemann/users?u=root&p=root' \  
-d '{"name": "riemann", "password": "riemann"}'  
$ curl -X POST 'http://localhost:8086/db/grafana/users?u=root&p=root' \  
-d '{"name": "grafana", "password": "grafana"}'
```

NB. InfluxDB is easy to get running but is not production ready or stable so your data can very easily be lost.

#### 4.1.3 Step 3: Install Grafana

```
$ aptitude install nginx  
$ mkdir /var/www  
$ cd /var/www  
$ wget http://grafanarel.s3.amazonaws.com/grafana-1.8.1.tar.gz
```

```
$ tar -zxf grafana-1.8.1.tar.gz
$ mv grafana-1.8.1 grafana
```

Now we must create an nginx configuration in */etc/nginx/sites-enabled*.

You can use something like this

```
server {
    listen 80;
    server_name <your hostname>;
    access_log /var/log/nginx/grafana-access.log;
    error_log /var/log/nginx/grafana-error.log;

    location / {
        alias /var/www/grafana/;
        index index.html;
        try_files $uri $uri/ /index.html;
    }
}
```

Next we need a configuration file for grafana. Open */var/www/grafana/config.js* and use the following configuration

```
define(['settings'],
function (Settings) {
    return new Settings({
        datasources: {
            influxdb: {
                type: 'influxdb',
                url: "http://<your hostname>:8086/db/riemann",
                username: 'riemann',
                password: 'riemann',
            },
            grafana: {
                type: 'influxdb',
                url: "http://<your hostname>:8086/db/grafana",
                username: 'grafana',
                password: 'grafana',
                grafanaDB: true
            },
        },
        search: {
            max_results: 20
        },
        default_route: '/dashboard/file/default.json',
        unsaved_changes_warning: true,
        playlist_timespan: "1m",
        admin: {
            password: ''
        },
        window_title_prefix: 'Grafana - ',
        plugins: {
            panels: [],
            dependencies: []
        }
    });
});
```

#### 4.1.4 Step 4: Glue things together

Lets start by configuring Riemann to talk to InfluxDB. This is the full /etc/riemann/riemann.config file.

```
; -*- mode: clojure; -*-
; vim: filetype=clojure
(require 'capacitor.core)
(require 'capacitor.async)
(require 'clojure.core.async)

(defn make-async-influxdb-client [opts]
  (let [client (capacitor.core/make-client opts)
        events-in (capacitor.async/make-chan)
        resp-out (capacitor.async/make-chan)]
    (capacitor.async/run! events-in resp-out client 100 10000)
    (fn [series payload]
      (let [p (merge payload {
                           :series series
                           :time   (* 1000 (:time payload)) ;; s → ms
                         })]
        (clojure.core.async/put! events-in p)))))

(def influx (make-async-influxdb-client {
  :host      "localhost"
  :port      8086
  :username "riemann"
  :password "riemann"
  :db        "riemann"
}))

(logging/init {:file "/var/log/riemann/riemann.log"})

; Listen on the local interface over TCP (5555), UDP (5555), and websockets
; (5556)
(let [host "0.0.0.0"]
  (tcp-server {:host host})
  (udp-server {:host host})
  (ws-server  {:host host}))

(periodically-expire 60)

(let [index (index)]
  (streams
    index

    (fn [event]
      (let [series (format "%s.%s" (:host event) (:service event))]
        (influx series {
          :time   (:time event)
          :value  (:metric event)
        }))))
```

You're pretty much done at this point, and should see the metrics from the Riemann server process if you open up Grafana and look through the query builder.

#### 4.1.5 Step 5: Using Tensor to retrieve stats from munin-node

First of all, install Tensor

```
$ pip install tensor
```

Next create /etc/tensor and a *tensor.yml* file in that directory.

The *tensor.yml* config file should look like this

```
ttl: 60.0
interval: 1.0

outputs:
  - output: tensor.outputs.riemann.RiemannTCP
    port: 5555
    server: <riemann server>

# Sources
sources:
  - service: mymunin
    source: tensor.sources.munin.MuninNode
    interval: 60.0
    ttl: 120.0
    critical:
      mymunin.system.load.load: "> 2"
}
```

This configures Tensor to connect to the munin-node on the local machine and retrieve all configured plugin values. You can create critical alert levels by setting the dot separated prefix for the service name and munin plugin.

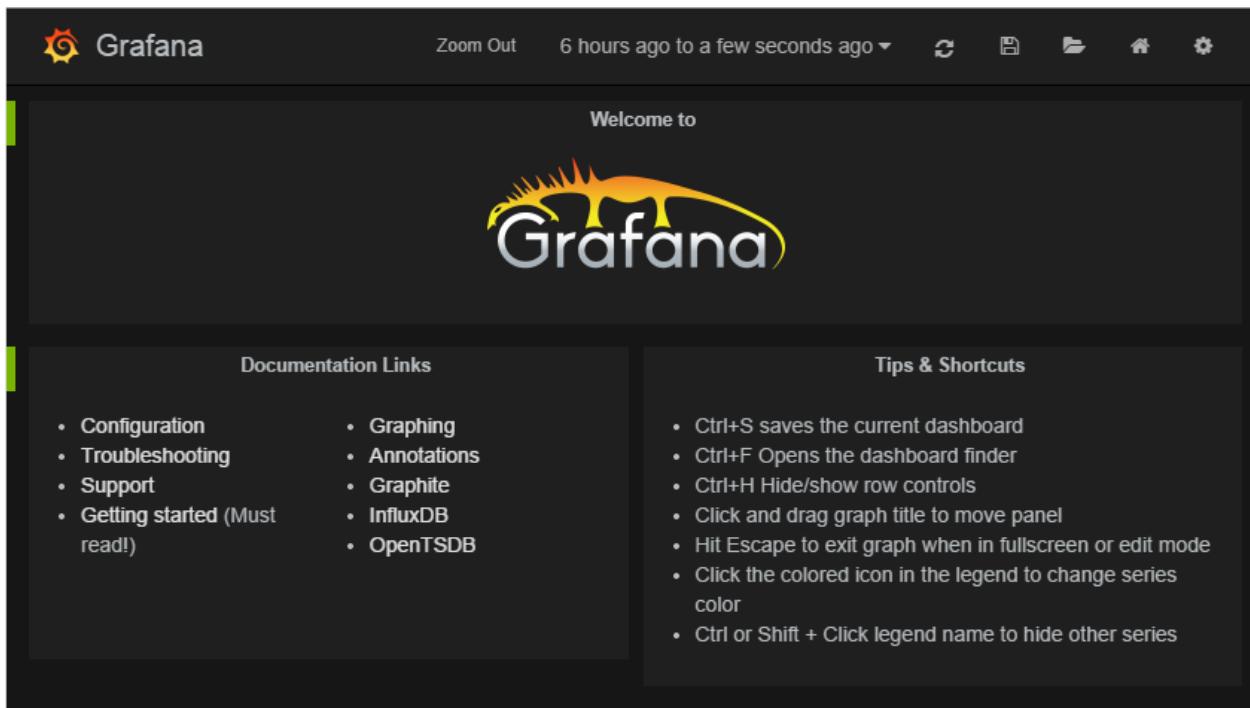
You can now start Tensor

```
$ twistd -n tensor -c /etc/tensor/tensor.yml
2014-10-22 13:30:38+0200 [-] Log opened.
2014-10-22 13:30:38+0200 [-] twistd 14.0.2 (/home/colin/riemann-tensor/ve/bin/python 2.7.6) starting
2014-10-22 13:30:38+0200 [-] reactor class: twisted.internet.epollreactor.EPollReactor.
2014-10-22 13:30:38+0200 [-] Starting factory <tensor.protocol.riemann.RiemannClientFactory instance
```

This pretty much indicates everything is alright, or else we'd see quickly see some errors.

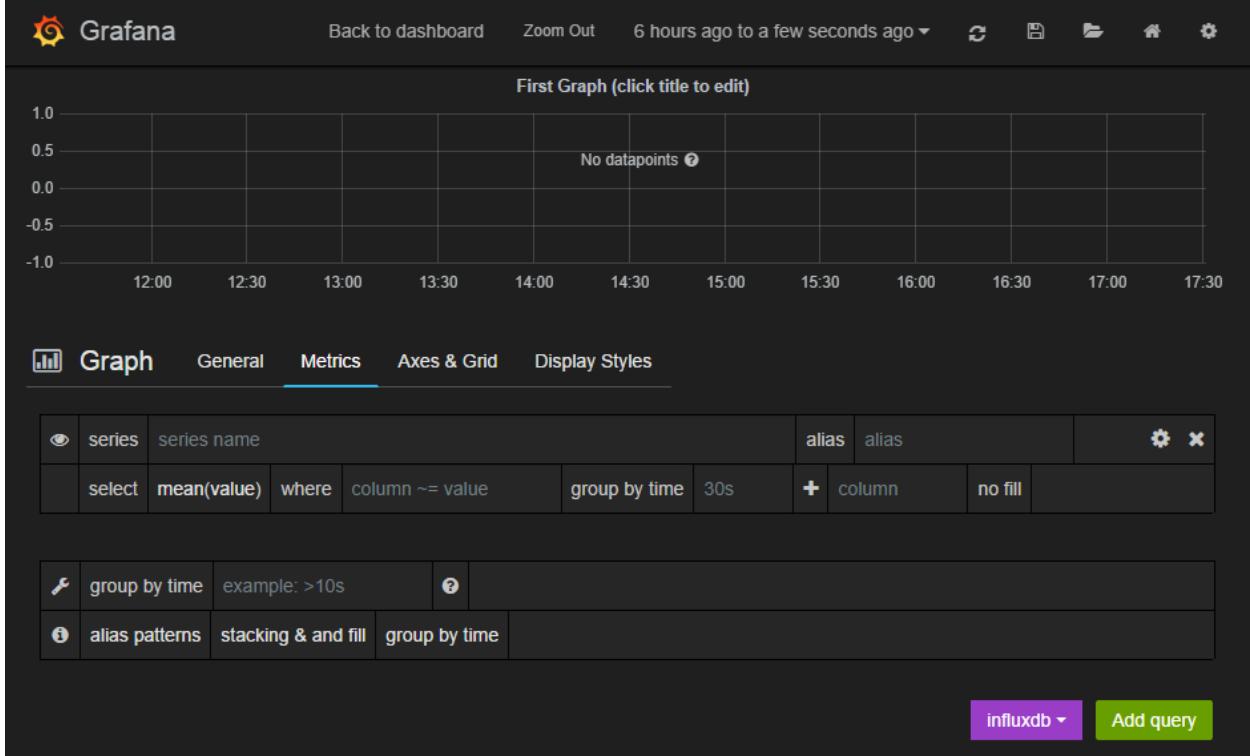
Next we will add some graphs to Grafana

#### 4.1.6 Step 6: Creating graphs in Grafana

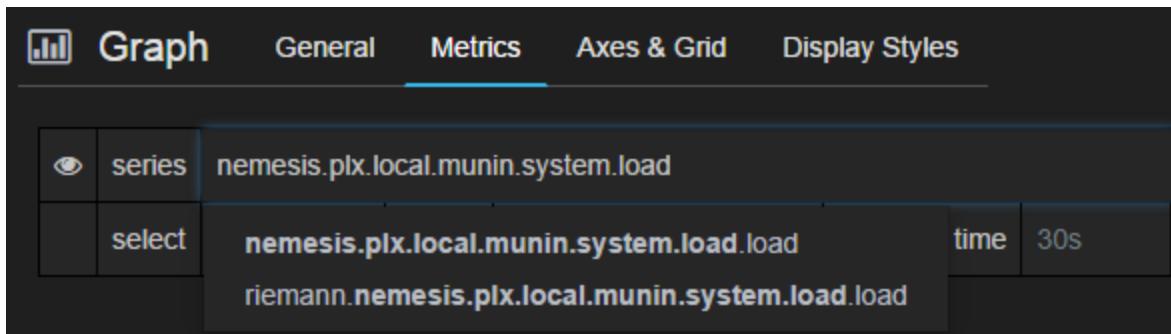


Click on the green row tag on the left, and delete all but the last row. This will leave you with an empty graph.

Click the title of the graph, then click *Edit*.



In the edit screen the Metrics tab will be open already. Now we can add our munin metrics. If you start typing in the *series* field you should see your hosts and metrics autocomplete.



Many Munin metrics are *counter* types which are usually converted to a rate by the RRD aggregation on Munin Graph.

Handily the `tensor.sources.munin.MuninNode` source takes care of this by caching the metric between run intervals when that type is used.

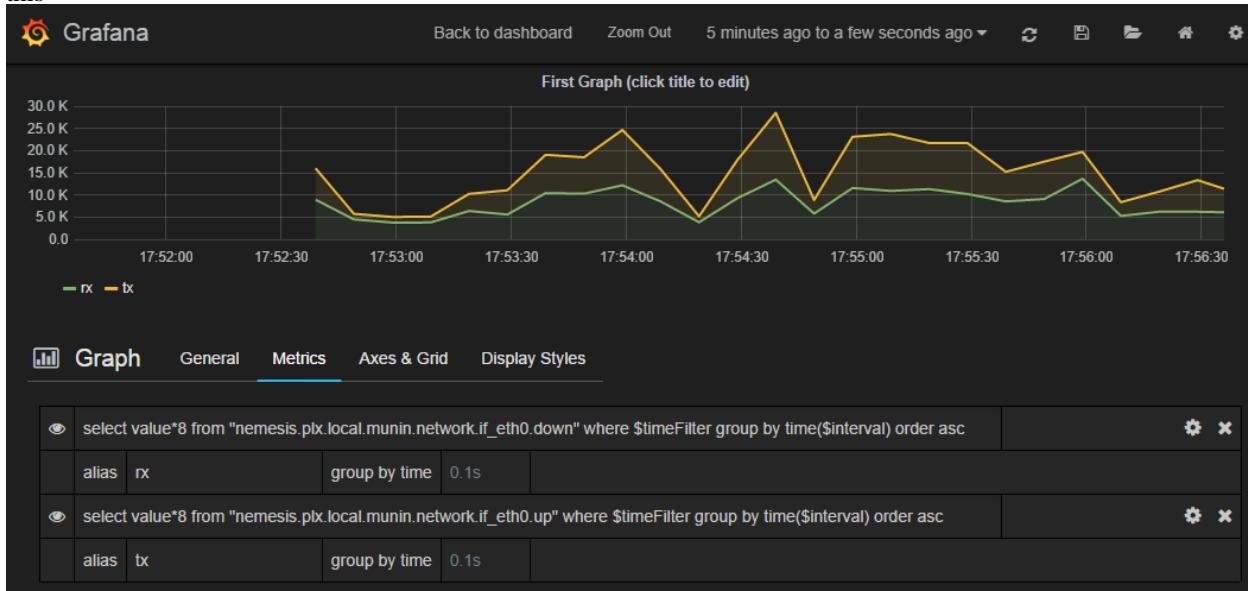
If we wanted to graph our network interface all we need to do is make it a slightly better unit by multiplying the Byte/sec metric by 8, since Grafana provides a bit/sec legend format.

To do this start by clicking the gear icon on the metric query, then select *Raw query mode*.

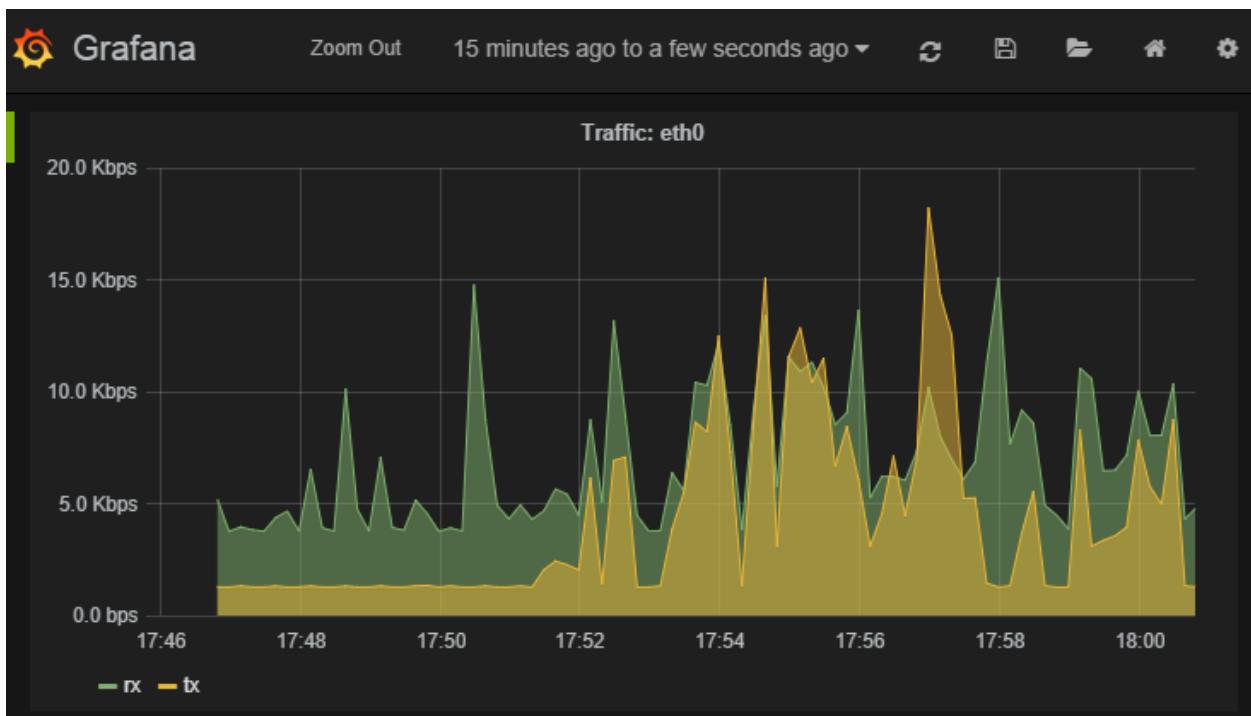
Use the following query

```
select value * 8 from "<your hostname>.munin.network.if_eth0.down" where $timeFilter group by time($interval) order asc
```

And chose an alias of “RX”. Do the same for if\_eth0.up and alias that “TX”. You should end up with something like this



Click on *General* to edit the title, and then on *Axes & Grid* change the Format to *bps*. Under *Display Styles* you can stack the data or play around with the look of the graph. Click *Back to dashboard* and you should end up with something as follows



API Documentation:



**tensor**

---

## 5.1 tensor.aggregators

```
tensor.aggregators.Counter(a, b, delta)
    Counter derivative

tensor.aggregators.Counter32(a, b, delta)
    32bit counter aggregator with wrapping

tensor.aggregators.Counter64(a, b, delta)
    64bit counter aggregator with wrapping
```

## 5.2 tensor.interfaces

### 5.3 tensor.objects

```
class tensor.objects.Event(state, service, description, metric, ttl, tags=None, hostname=None, aggregation=None, evtime=None, attributes=None, type='riemann')
    Bases: object
    Tensor Event object

All sources pass these to the queue, which form a proxy object to create protobuf Event objects
```

#### Parameters

- **state** – Some sort of string < 255 chars describing the state
- **service** – The service name for this event
- **description** – A description for the event, ie. “My house is on fire!”
- **metric** – int or float metric for this event
- **ttl** – TTL (time-to-live) for this event
- **tags** – List of tag strings
- **hostname** – Hostname for the event (defaults to system fqdn)
- **aggregation** – Aggregation function
- **attributes** – A dictionary of key/value attributes for this event
- **evtime** – Event timestamp override

```
class tensor.objects.Output (config, tensor)
Bases: object

Output parent class

Outputs can inherit this object which provides a construct for a working output

Parameters

- config – Dictionary config for this queue (usually read from the yaml configuration)
- tensor – A TensorService object for interacting with the queue manager

createClient ()
Deferred which sets up the output

eventsReceived ()
Receives a list of events and processes them

Arguments: events – list of tensor.objects.Event

stop ()
Called when the service shuts down

class tensor.objects.Source (config, queueBack, tensor)
Bases: object

Source parent class

Sources can inherit this object which provides a number of utility methods.

Parameters

- config – Dictionary config for this queue (usually read from the yaml configuration)
- queueBack – A callback method to receive a list of Event objects
- tensor – A TensorService object for interacting with the queue manager

createEvent (state, description, metric, prefix=None, hostname=None, aggregation=None, ev-time=None)
Creates an Event object from the Source configuration

createLog (type, data, evtime=None, hostname=None)
Creates an Event object from the Source configuration

startTimer ()
Starts the timer for this source

stopTimer ()
Stops the timer for this source

tick (*args, **kwargs)
Called for every timer tick. Calls self.get which can be a deferred and passes that result back to the queueBack method

Returns a deferred
```

## 5.4 tensor.service

```
class tensor.service.TensorService (config)
Bases: twisted.application.service.Service

Tensor service
```

Runs timers, configures sources and manages the queue

**sendEvent** (*source, events*)

Callback that all event sources call when they have a new event or list of events

**setupOutputs** (*config*)

Setup output processors

**setupSources** (*config*)

Sets up source objects from the given config

**sourceWatchdog** ()

Watchdog timer function.

Recreates sources which have not generated events in 10\*interval if they have watchdog set to true in their configuration

## 5.5 tensor.utils

**class** tensor.utils.**BodyReceiver** (*finished*)

Bases: twisted.internet.protocol.Protocol

Simple buffering consumer for body objects

**class** tensor.utils.**PersistentCache** (*location='var/lib/tensor/cache'*)

Bases: object

A very basic dictionary cache abstraction. Not to be used for large amounts of data or high concurrency

**contains** (*k*)

Return True if key *k* exists

**delete** (*k*)

Remove key *k* from the cache

**expire** (*age*)

Expire any items in the cache older than *age* seconds

**get** (*k*)

Returns key contents, and modify time

**set** (*k, v*)

Set a key *k* to value *v*

**class** tensor.utils.**ProcessProtocol** (*deferred, timeout*)

Bases: twisted.internet.protocol.ProcessProtocol

ProcessProtocol which supports timeouts

**class** tensor.utils.**Resolver**

Bases: object

Helper class for DNS resolution

**class** tensor.utils.**StringProducer** (*body*)

Bases: object

String producer for writing to HTTP requests

**exception** tensor.utils.**Timeout**

Bases: exceptions.Exception

Raised to notify that an operation exceeded its timeout.

`tensor.utils.fork(executable, args=(), env={}, path=None, timeout=3600)`

Provides a deferred wrapper function with a timeout function

#### Parameters

- **executable** (*str.*) – Executable
- **args** (*tupple.*) – Tuppble of arguments
- **env** (*dict.*) – Environment dictionary
- **timeout** (*int.*) – Kill the child process if timeout is exceeded

---

**tensor.protocol**

---

## 6.1 tensor.protocol.elasticsearch

```
class tensor.protocol.elasticsearch.ElasticSearch(url='http://localhost:9200',
                                                user=None,           password=None,
                                                index='tensor-%Y.%m.%d')
Bases: object
Twisted ElasticSearch API
```

## 6.2 tensor.protocol.icmp

```
class tensor.protocol.icmp.EchoPacket(seq=0, id=None, data=None, packet=None)
Bases: object
ICMP Echo packet encoder and decoder

class tensor.protocol.icmp.ICMPPing(d, dst, count, inter=0.2, maxwait=1000, size=64)
Bases: twisted.internet.protocol.DatagramProtocol
ICMP Ping implementation

class tensor.protocol.icmp.ICMPPort(port, proto, interface='', maxPacketSize=8192, reactor=None)
Bases: twisted.internet.udp.Port
Raw socket listener for ICMP

class tensor.protocol.icmp.IP(packet)
Bases: object
IP header decoder

tensor.protocol.icmp.ping(dst, count, inter=0.2, maxwait=1000, size=64)
Sends ICMP echo requests to destination dst count times. Returns a deferred which fires when responses are finished.
```

## 6.3 tensor.protocol.riemann

```
class tensor.protocol.riemann.RiemannClientFactory(hosts, failover=False)
Bases: twisted.internet.protocol.ReconnectingClientFactory
```

---

A reconnecting client factory which creates RiemannProtocol instances

**class** tensor.protocol.riemann.**RiemannProtocol**

Bases: twisted.protocols.basic.Int32StringReceiver, tensor.protocol.riemann.RiemannProtocol

Riemann protobuf protocol

**class** tensor.protocol.riemann.**RiemannUDP** (*host, port*)

Bases: twisted.internet.protocol.DatagramProtocol, tensor.protocol.riemann.RiemannProtocol

UDP datagram protocol for Riemann

## 6.4 tensor.protocol.ssh

## 6.5 tensor.protocol.sflow

### 6.5.1 tensor.protocol.sflow.server

**class** tensor.protocol.sflow.server.**DatagramReceiver**

Bases: twisted.internet.protocol.DatagramProtocol

DatagramReceiver for sFlow packets

### 6.5.2 tensor.protocol.sflow.protocol

---

## tensor.logs

---

### 7.1 tensor.logs.follower

```
class tensor.logs.follower.LogFollower(logfile, parser=None, tmp_path='/var/lib/tensor/', history=False)
```

Bases: object

Provides a class for following log files between runs

#### Parameters

- **logfile** (*str*) – Full path to logfile
- **parser** (*str*) – Optional parser method for log lines

**get** (*max\_lines=None*)

Returns a big list of all log lines since the last run

**get\_fn** (*fn, max\_lines=None*)

Passes each parsed log line to *fn* This is a better idea than storing a giant log file in memory

### 7.2 tensor.logs.parsers

```
class tensor.logs.parsers.ApacheLogParser(format)
```

Parses Apache log format

Adapted from <http://code.google.com/p/apachelog>

**Parameters** **format** (*str*) – Apache log format definition eg r'%h %l %u %t "%r" %>s %b  
“%{Referer}i” “%{User-Agent}i” or one of ‘common’, ‘vhcommon’ or ‘combined’

**names** ()

Returns the field names the parser extracted from the input format (a list)

**parse** (*line*)

Parses a single line from the log file and returns a dictionary of it’s contents.

Raises and exception if it couldn’t parse the line

**pattern** ()

Returns the compound regular expression the parser extracted from the input format (a string)



---

**tensor.sources**

---

## 8.1 tensor.sources.database.postgresql

```
class tensor.sources.database.postgresql.PostgreSQL(*a, **kw)
Bases: tensor.objects.Source
```

Reads PostgreSQL metrics

**Configuration arguments:**

**Parameters**

- **host** (*str.*) – Database host
- **port** (*int.*) – Database port
- **user** (*str.*) – Username
- **password** (*str.*) – Password

**Metrics:**

**(service name).(database name).(metrics)** Metrics from pg\_stat\_database

## 8.2 tensor.sources.database.elasticsearch

```
class tensor.sources.database.elasticsearch.ElasticSearch(*a, **kw)
Bases: tensor.objects.Source
```

Reads elasticsearch metrics

**Configuration arguments:**

**Parameters**

- **url** (*str.*) – Elasticsearch base URL (default: <http://localhost:9200>)
- **user** (*str.*) – Basic auth username
- **password** (*str.*) – Password

**Metrics:**

**(service name).cluster.status** Cluster status (Red=0, Yellow=1, Green=2)

**(service name).cluster.nodes** Cluster node count

(service name).indices Total indices in cluster  
(service name).shards.total Total number of shards  
(service name).shards.primary Number of primary shards  
(service name).documents.total Total documents  
(service name).documents.rate Documents per second  
(service name).documents.size Size of document store in bytes

## 8.3 tensor.sources.database.memcache

**class** `tensor.sources.database.memcache(*a, **kw)`

Bases: `tensor.objects.Source`

Reads memcache metrics

**Configuration arguments:**

**Parameters**

- **host** (`str.`) – Database host (default localhost)
- **port** (`int.`) – Database port (default 11211)

**Metrics:**

(service name).(metrics) Metrics from memcached

## 8.4 tensor.sources.docker

**class** `tensor.sources.docker.ContainerStats(*a, **kw)`

Bases: `tensor.objects.Source`

Returns stats for Docker containers on this host

**Configuration arguments:**

**Parameters** **url** (`str.`) – Docker stats URL

**Metrics:**

(service name).(container name).mem\_limit Maximum memory for container  
(service name).(container name).mem\_used Memory used by container  
(service name).(container name).cpu Percentage of system CPU in use  
(service name).(container name).io\_read IO reads per second  
(service name).(container name).io\_write IO writes per second  
(service name).(container name).io\_sync IO synchronous op/s  
(service name).(container name).io\_async IO asynchronous op/s  
(service name).(container name).io\_total Total IOPS

Note. If a MARATHON\_APP\_ID environment variable exists on the container then *container name* will be used instead of that.

## 8.5 tensor.sources.haproxy

```
class tensor.sources.haproxy.HAProxy(*a, **kw)
Bases: tensor.objects.Source
```

Reads Nginx stub\_status

**Configuration arguments:**

**Parameters**

- **url** (*str.*) – URL to fetch stats from
- **user** (*str.*) – Username
- **password** (*str.*) – Password

**Metrics:**

(**service name**).(**backend|frontend|nodes**).(**stats**) Various statistics

## 8.6 tensor.sources.generator

```
class tensor.sources.generator.Function(config, queueBack, tensor)
Bases: tensor.objects.Source
```

Produces an arbitrary function

Functions can contain the functions sin, cos, sinh, cosh, tan, tanh, asin, acos, atan, atan, asinh, acosh, atanh, log(n, [base]), abs

Or the constants e, pi, and variable x

**Configuration arguments:**

**Parameters**

- **dx** (*float.*) – Resolution with time (steps of x)
- **function** (*string.*) – Function to produce

## 8.7 tensor.sources.linux

### 8.7.1 tensor.sources.linux.basic

```
class tensor.sources.linux.basic.CPU(*a)
Bases: tensor.objects.Source
```

Reports system CPU utilisation as a percentage/100

**Metrics:**

(**service name**) Percentage CPU utilisation

(**service name**).(**type**) Percentage CPU utilisation by type

```
class tensor.sources.linux.basic.DiskFree(config, queueBack, tensor)
Bases: tensor.objects.Source
```

Returns the free space for all mounted filesystems

**Configuration arguments:**

**Parameters** **disks** (*list.*) – List of devices to check (optional)

**Metrics:**

(**service name**).(**device**).**used** Used space (%)  
(**service name**).(**device**).**bytes** Used space (kbytes)  
(**service name**).(**device**).**free** Free space (kbytes)

**class** `tensor.sources.linux.basic.DiskIO(*a, **kw)`

Bases: `tensor.objects.Source`

Reports disk IO statistics per device

**Configuration arguments:**

**Parameters** **devices** (*list.*) – List of devices to check (optional)

**Metrics:**

(**service name**).(**device name**).**reads** Number of completed reads  
(**service name**).(**device name**).**read\_bytes** Bytes read per second  
(**service name**).(**device name**).**read\_latency** Disk read latency  
(**service name**).(**device name**).**writes** Number of completed writes  
(**service name**).(**device name**).**write\_bytes** Bytes written per second  
(**service name**).(**device name**).**write\_latency** Disk write latency

**class** `tensor.sources.linux.basic.LoadAverage(config, queueBack, tensor)`

Bases: `tensor.objects.Source`

Reports system load average for the current host

**Metrics:**

(**service name**) Load average

**class** `tensor.sources.linux.basic.Memory(config, queueBack, tensor)`

Bases: `tensor.objects.Source`

Reports system memory utilisation as a percentage/100

**Metrics:**

(**service name**) Percentage memory utilisation

**class** `tensor.sources.linux.basic.Network(config, queueBack, tensor)`

Bases: `tensor.objects.Source`

Returns all network interface statistics

**Configuration arguments:**

**Parameters** **interfaces** (*list.*) – List of interfaces to check (optional)

**Metrics:**

(**service name**).(**device**).**tx\_bytes** Bytes transmitted  
(**service name**).(**device**).**tx\_packets** Packets transmitted  
(**service name**).(**device**).**tx\_errors** Errors

**(service name).(device).rx\_bytes** Bytes received  
**(service name).(device).rx\_packets** Packets received  
**(service name).(device).rx\_errors** Errors

## 8.7.2 tensor.sources.linux.process

**class** `tensor.sources.linux.process.ProcessCount` (*config, queueBack, tensor*)  
Bases: `tensor.objects.Source`

Returns the ps count on the system

**Metrics:**

**(service name)** Number of processes

**class** `tensor.sources.linux.process.ProcessStats` (*config, queueBack, tensor*)  
Bases: `tensor.objects.Source`

Returns memory used by each active parent process

**Metrics:**

**(service name).proc.(process name).cpu** Per process CPU usage

**(service name).proc.(process name).memory** Per process memory use

**(service name).proc.(process name).age** Per process age

**(service name).user.(user name).cpu** Per user CPU usage

**(service name).user.(user name).memory** Per user memory use

## 8.7.3 tensor.sources.linux.sensors

**class** `tensor.sources.linux.sensors.SMART` (\**a*, \*\**kw*)  
Bases: `tensor.objects.Source`

Returns SMART output for all disks

**Metrics:**

**(service name).(disk).(sensor)** Sensor value

**class** `tensor.sources.linux.sensors.Sensors` (*config, queueBack, tensor*)  
Bases: `tensor.objects.Source`

Returns lm-sensors output

NB. This is very untested on different configurations and versions. Please report any issues with the output of the `sensors` command to help improve it.

**Metrics:**

**(service name).(adapter).(sensor)** Sensor value

## 8.8 tensor.sources.media

### 8.8.1 tensor.sources.media.libav

```
class tensor.sources.media.libav.DarwinRTSP (config, queueBack, tensor)
Bases: tensor.objects.Source
```

Makes avprobe requests of a Darwin RTSP sample stream (sample\_100kbit.mp4)

#### Configuration arguments:

**Parameters** **destination** – Host name or IP address to check

**Metrics:** :(service name): Time to complete request

You can also override the *hostname* argument to make it match metrics from that host.

## 8.9 tensor.sources.munin

```
class tensor.sources.munin.MuninNode (config, queueBack, tensor)
Bases: tensor.objects.Source
```

Connects to munin-node and retrieves all metrics

#### Configuration arguments:

**Parameters**

- **host** (*str.*) – munin-node hostname (probably localhost)
- **port** (*int.*) – munin-node port (probably 4949)

#### Metrics:

(**service name**).(**plugin name**).(**keys...**) A dot separated tree of munin plugin keys

```
class tensor.sources.munin.MuninProtocol
```

Bases: twisted.protocols.basic.LineReceiver

MuninProtocol - provides a line receiver protocol for making requests to munin-node

Requests must be made sequentially

## 8.10 tensor.sources.network

```
class tensor.sources.network.HTTP (config, queueBack, tensor)
Bases: tensor.objects.Source
```

Performs an HTTP request

#### Configuration arguments:

**Parameters**

- **url** (*str.*) – HTTP URL
- **method** (*str.*) – HTTP request method to use (default GET)
- **match** (*str.*) – A text string to match in the document when it is correct
- **useragent** (*str.*) – User-Agent header to use

- **timeout** (*int.*) – Timeout for connection (default 60s)

**Metrics:**

**(service name).latency** Time to complete request

**class** `tensor.sources.network.Ping` (*config, queueBack, tensor*)

Bases: `tensor.objects.Source`

Performs an Ping checks against a destination

**Configuration arguments:**

**Parameters** `destination` (*str.*) – Host name or IP address to ping

**Metrics:**

**(service name).latency** Ping latency

**(service name).loss** Packet loss

You can also override the *hostname* argument to make it match metrics from that host.

## 8.11 `tensor.sources.nginx`

**class** `tensor.sources.nginx.Nginx` (*config, queueBack, tensor*)

Bases: `tensor.objects.Source`

Reads Nginx stub\_status

**Configuration arguments:**

**Parameters** `stats_url` (*str.*) – URL to fetch stub\_status from

**Metrics:**

**(service name).active** Active connections at this time

**(service name).accepts** Accepted connections

**(service name).handled** Handled connections

**(service name).requests** Total client requests

**(service name).reading** Reading requests

**(service name).writing** Writing responses

**(service name).waiting** Waiting connections

**class** `tensor.sources.nginx.NginxLog` (*\*a*)

Bases: `tensor.objects.Source`

Tails Nginx log files, parses them and returns log events for outputs which support them.

**Configuration arguments:****Parameters**

- **log\_format** (*str.*) – Log format passed to parser, same as the config definition (default: combined)
- **file** (*str.*) – Log file
- **max\_lines** (*int.*) – Maximum number of log lines to read per interval to prevent overwhelming Tensor when reading large logs (default 2000)

```
class tensor.sources.nginx.NginxLogMetrics (*a)
Bases: tensor.objects.Source

Tails Nginx log files, parses them and returns metrics for data usage and requests against other fields.

Configuration arguments:

Parameters

- log_format (str.) – Log format passed to parser, same as the config definition
- file (str.) – Log file
- max_lines (int.) – Maximum number of log lines to read per interval to prevent overwhelming Tensor when reading large logs (default 2000)
- resolution (int.) – Aggregate bucket resolution in seconds (default 10)
- history (bool.) – Read the entire file from scratch if we've never seen it (default false)

```

**Metrics:**

```
(service name).total_rbytes Bytes total for all requests
(service name).total_requests Total request count
(service name).stats.(code).(requests|rbytes) Metrics by status code
(service name).user-agent.(agent).(requests|rbytes) Metrics by user agent
(service name).client.(ip).(requests|rbytes) Metrics by client IP
(service name).request.(request path).(requests|rbytes) Metrics by request path
```

## 8.12 tensor.sources.python

### 8.12.1 tensor.sources.python.uwsgi

```
class tensor.sources.python.uwsgi.Emperor (config, queueBack, tensor)
Bases: tensor.objects.Source

Connects to UWSGI Emperor stats and creates useful metrics

Configuration arguments:
```

**Parameters**

- **host** (*str.*) – Hostname (default localhost)
- **port** (*int.*) – Port

```
class tensor.sources.python.uwsgi.JSONProtocol
Bases: twisted.internet.protocol.Protocol

JSON line protocol
```

## 8.13 tensor.sources.rabbitmq

```
class tensor.sources.rabbitmq.Queues (*a, **kw)
Bases: tensor.objects.Source

Returns Queue information for a particular vhost
```

**Configuration arguments:**

**Parameters** **vhost** (*str.*) – Vhost name

**Metrics:**

**(service\_name).(queue).ready** Ready messages for queue

**(service\_name).(queue).unack** Unacknowledged messages for queue

**(service\_name).(queue).ready\_rate** Ready rate of change per second

**(service\_name).(queue).unack\_rate** Unacknowledge rate of change per second

## 8.14 tensor.sources.redis

```
class tensor.sources.redis.Queues (*a, **kw)
Bases: tensor.objects.Source
```

Query llen from redis-cli

**Configuration arguments:****Parameters**

- **queue** (*str.*) – Queue name (defaults to ‘celery’, just because)
- **db** (*int.*) – DB number
- **clipath** (*str.*) – Path to redis-cli (default: /usr/bin/redis-cli)

**Metrics:**

**(service\_name)** Queue length

**(service\_name)** Queue rate

## 8.15 tensor.sources.riak

```
class tensor.sources.riak.RiakStats (config, queueBack, tensor)
Bases: tensor.objects.Source
```

Returns GET/PUT rates for a Riak node

**Configuration arguments:****Parameters**

- **url** (*str.*) – Riak stats URL
- **useragent** (*str.*) – User-Agent header to use

**Metrics:**

**(service name).latency** Time to complete request

## 8.16 tensor.sources.riemann

```
class tensor.sources.riemann.RiemannTCP (config, queueBack, tensor)
```

Bases: *tensor.objects.Source*

Provides a listening server which accepts Riemann metrics and proxies them to our queue.

**Configuration arguments:**

**Parameters** **port** (*int.*) – Port to listen on (default 5555)

**startTimer()**

Creates a Riemann TCP server instead of a timer

```
class tensor.sources.riemann.RiemannTCPServer (source)
```

Bases: *tensor.protocol.riemann.RiemannProtocol*

Server implementation of the Riemann protocol

## 8.17 tensor.sources.sflow

```
class tensor.sources.sflow.sFlow (config, queueBack, tensor)
```

Bases: *tensor.objects.Source*

Provides an sFlow server Source

**Configuration arguments:**

**Parameters**

- **port** (*int.*) – UDP port to listen on

- **dnslookup** (*bool.*) – Enable reverse DNS lookup for device IPs (default: True)

**Metrics:**

Metrics are published using the key patterns (device).(service name).(interface).(inlout)Octets (device).(service name).(interface).ip (device).(service name).(interface).port

**startTimer()**

Creates a sFlow datagram server

```
class tensor.sources.sflow.sFlowReceiver (source)
```

Bases: *tensor.protocol.sflow.server.DatagramReceiver*

sFlow datagram protocol

## 8.18 tensor.sources.snmp

```
class tensor.sources.snmp.SNMP (*a, **kw)
```

Bases: *tensor.objects.Source*

Connects to an SNMP agent and retrieves OIDs

**Configuration arguments:**

**Parameters**

- **ip** (*str.*) – SNMP agent host (default: 127.0.0.1)

- **port** (*int.*) – SNMP port (default: 161)
- **community** (*str.*) – SNMP read community

```
class tensor.sources.snmp.SNMPCisco837(*a, **kw)
Bases: tensor.sources.snmp.SNMP
```

Connects to a Cisco 837 and makes metrics

#### Configuration arguments:

##### Parameters

- **ip** (*str.*) – SNMP agent host (default: 127.0.0.1)
- **port** (*int.*) – SNMP port (default: 161)
- **community** (*str.*) – SNMP read community

```
class tensor.sources.snmp.SNMPConnection(host, port, community)
```

Bases: *object*

A wrapper class for PySNMP functions

##### Parameters

- **host** (*str.*) – SNMP agent host
- **port** (*int.*) – SNMP port
- **community** (*str.*) – SNMP read community

(This is not a source and you shouldn't try to use it as one)



**tensor.outputs**

## 9.1 tensor.outputs.riemann

**class** `tensor.outputs.riemann.RiemannTCP(*a)`  
 Bases: `tensor.objects.Output`

Riemann TCP output

**Configuration arguments:**

### Parameters

- **server** (`str.`) – Riemann server hostname (default: localhost)
- **port** (`int.`) – Riemann server port (default: 5555)
- **failover** (`bool.`) – Enable server failover, in which case `server` may be a list
- **maxrate** (`int.`) – Maximum de-queue rate (0 is no limit)
- **maxsize** (`int.`) – Maximum queue size (0 is no limit, default is 250000)
- **interval** (`float.`) – De-queue interval in seconds (default: 1.0)
- **pressure** (`int.`) – Maximum backpressure (-1 is no limit)
- **tls** (`bool.`) – Use TLS (default false)
- **cert** (`str.`) – Host certificate path
- **key** (`str.`) – Host private key path
- **allow\_nan** (`bool.`) – Send events with None metric value (default true)

### `createClient()`

Create a TCP connection to Riemann with automatic reconnection

### `emptyQueue()`

Remove all or self.queueDepth events from the queue

### `eventsReceived(events)`

Receives a list of events and transmits them to Riemann

Arguments: `events` – list of `tensor.objects.Event`

### `stop()`

Stop this client.

```
    tick()
        Clock tick called every self.inter

class tensor.outputs.riemann.RiemannUDP(*a)
    Bases: tensor.objects.Output

    Riemann UDP output (spray-and-pray mode)

Configuration arguments:

    Parameters
        • server (str.) – Riemann server IP address (default: 127.0.0.1)
        • port (int.) – Riemann server port (default: 5555)

    createClient()
        Create a UDP connection to Riemann

    eventsReceived(events)
        Receives a list of events and transmits them to Riemann

        Arguments: events – list of tensor.objects.Event
```

## 9.2 tensor.outputs.elasticsearch

```
    class tensor.outputs.elasticsearch.ElasticSearch(*a)
        Bases: tensor.objects.Output

        ElasticSearch HTTP API output

    Configuration arguments:

    Parameters
        • url (str) – Elasticsearch URL (default: http://localhost:9200)
        • maxsize (int) – Maximum queue backlog size (default: 250000, 0 disables)
        • maxrate (int) – Maximum rate of documents added to index (default: 100)
        • interval (int) – Queue check interval in seconds (default: 1.0)
        • user (str) – Optional basic auth username
        • password (str) – Optional basic auth password
        • index (str) – Index name format to store documents in Elastic (default: tensor-%Y.%m.%d)

    createClient()
        Sets up HTTP connector and starts queue timer

    eventsReceived(events)
        Receives a list of events and queues them

        Arguments: events – list of tensor.objects.Event

    stop()
        Stop this client.

    tick(*args, **kwargs)
        Clock tick called every self.inter
```

`tensor.outputs.elasticsearch.ElasticSearchLog`  
alias of *ElasticSearch*



## **Indices and tables**

---

- genindex
- modindex
- search



**d**

`docker (Any)`, 32

**e**

`elasticsearch (Unix)`, 31

**g**

`generator (Any)`, 33

**h**

`haproxy (Unix)`, 33

**m**

`memcache (Unix)`, 32

`munin (Any)`, 36

**n**

`network (Unix)`, 36

`nginx (Unix)`, 37

**p**

`postgresql (Unix)`, 31

**r**

`riak (Any)`, 39

`riemann (Unix)`, 40

**s**

`sflow (Unix)`, 40

`snmp (Unix)`, 40

**t**

`tensor.aggregators`, 23

`tensor.interfaces`, 23

`tensor.logs.follower`, 29

`tensor.logs.parsers`, 29

`tensor.objects`, 23

`tensor.outputs.elasticsearch`, 44

`tensor.outputs.riemann`, 43

`tensor.protocol.elasticsearch`, 27

`tensor.protocol.icmp`, 27  
`tensor.protocol.riemann`, 27  
`tensor.protocol.sflow.protocol`, 28  
`tensor.protocol.sflow.server`, 28  
`tensor.protocol.ssh`, 28  
`tensor.service`, 24  
`tensor.sources.database.elasticsearch`,  
    31  
`tensor.sources.database.memcache`, 32  
`tensor.sources.database.postgresql`, 31  
`tensor.sources.docker`, 32  
`tensor.sources.generator`, 33  
`tensor.sources.haproxy`, 33  
`tensor.sources.linux.basic`, 33  
`tensor.sources.linux.process`, 35  
`tensor.sources.linux.sensors`, 35  
`tensor.sources.media.libav`, 36  
`tensor.sources.munin`, 36  
`tensor.sources.network`, 36  
`tensor.sources.nginx`, 37  
`tensor.sources.python.uwsgi`, 38  
`tensor.sources.rabbitmq`, 38  
`tensor.sources.redis`, 39  
`tensor.sources.riak`, 39  
`tensor.sources.riemann`, 40  
`tensor.sources.sflow`, 40  
`tensor.sources.snmp`, 40  
`tensor.utils`, 25

**u**

`uwsgi (Any)`, 38



**A**

ApacheLogParser (class in tensor.logs.parsers), 29

**B**

BodyReceiver (class in tensor.utils), 25

**C**

ContainerStats (class in tensor.sources.docker), 32

contains() (tensor.utils.PersistentCache method), 25

Counter() (in module tensor.aggregators), 23

Counter32() (in module tensor.aggregators), 23

Counter64() (in module tensor.aggregators), 23

CPU (class in tensor.sources.linux.basic), 33

createClient() (tensor.objects.Output method), 24

createClient() (tensor.outputs.elasticsearch.ElasticSearch method), 44

createClient() (tensor.outputs.riemann.RiemannTCP method), 43

createClient() (tensor.outputs.riemann.RiemannUDP method), 44

createEvent() (tensor.objects.Source method), 24

createLog() (tensor.objects.Source method), 24

**D**

DarwinRTSP (class in tensor.sources.media.libav), 36

DatagramReceiver (class in tensor.protocol.sflow.server), 28

delete() (tensor.utils.PersistentCache method), 25

DiskFree (class in tensor.sources.linux.basic), 33

DiskIO (class in tensor.sources.linux.basic), 34

docker (module), 32

**E**

EchoPacket (class in tensor.protocol.icmp), 27

ElasticSearch (class in tensor.outputs.elasticsearch), 44

ElasticSearch (class in tensor.protocol.elasticsearch), 27

ElasticSearch (class in tensor.sources.database.elasticsearch), 31

elasticsearch (module), 31

ElasticSearchLog (in module tensor.outputs.elasticsearch), 44

Emperor (class in tensor.sources.python.uwsgi), 38

emptyQueue() (tensor.outputs.riemann.RiemannTCP method), 43

Event (class in tensor.objects), 23

eventsReceived() (tensor.objects.Output method), 24

eventsReceived() (tensor.outputs.elasticsearch.ElasticSearch method), 44

eventsReceived() (tensor.outputs.riemann.RiemannTCP method), 43

eventsReceived() (tensor.outputs.riemann.RiemannUDP method), 44

expire() (tensor.utils.PersistentCache method), 25

**F**

fork() (in module tensor.utils), 25

Function (class in tensor.sources.generator), 33

**G**

generator (module), 33

get() (tensor.logs.follower.LogFollower method), 29

get() (tensor.utils.PersistentCache method), 25

get\_fn() (tensor.logs.follower.LogFollower method), 29

**H**

HAProxy (class in tensor.sources.haproxy), 33

haproxy (module), 33

HTTP (class in tensor.sources.network), 36

**I**

ICMPPing (class in tensor.protocol.icmp), 27

ICMPPPort (class in tensor.protocol.icmp), 27

IP (class in tensor.protocol.icmp), 27

**J**

JSONProtocol (class in tensor.sources.python.uwsgi), 38

**L**

LoadAverage (class in tensor.sources.linux.basic), 34

LogFollower (class in tensor.logs.follower), 29

## M

Memcache (class in tensor.sources.database.memcache), 32

memcache (module), 32

Memory (class in tensor.sources.linux.basic), 34

munin (module), 36

MuninNode (class in tensor.sources.munin), 36

MuninProtocol (class in tensor.sources.munin), 36

## N

names() (tensor.logs.parsers.ApacheLogParser method), 29

Network (class in tensor.sources.linux.basic), 34

network (module), 36

Nginx (class in tensor.sources.nginx), 37

nginx (module), 37

NginxLog (class in tensor.sources.nginx), 37

NginxLogMetrics (class in tensor.sources.nginx), 37

## O

Output (class in tensor.objects), 24

## P

parse() (tensor.logs.parsers.ApacheLogParser method), 29

pattern() (tensor.logs.parsers.ApacheLogParser method), 29

PersistentCache (class in tensor.utils), 25

Ping (class in tensor.sources.network), 37

ping() (in module tensor.protocol.icmp), 27

PostgreSQL (class in tensor.sources.database.postgresql), 31

postgresql (module), 31

ProcessCount (class in tensor.sources.linux.process), 35

ProcessProtocol (class in tensor.utils), 25

ProcessStats (class in tensor.sources.linux.process), 35

## Q

Queues (class in tensor.sources.rabbitmq), 38

Queues (class in tensor.sources.redis), 39

## R

Resolver (class in tensor.utils), 25

riak (module), 39

RiakStats (class in tensor.sources.riak), 39

riemann (module), 40

RiemannClientFactory (class in tensor.protocol.riemann), 27

RiemannProtocol (class in tensor.protocol.riemann), 28

RiemannTCP (class in tensor.outputs.riemann), 43

RiemannTCP (class in tensor.sources.riemann), 40

RiemannTCPServer (class in tensor.sources.riemann), 40

RiemannUDP (class in tensor.outputs.riemann), 44

RiemannUDP (class in tensor.protocol.riemann), 28

## S

sendEvent() (tensor.service.TensorService method), 25

Sensors (class in tensor.sources.linux.sensors), 35

set() (tensor.utils.PersistentCache method), 25

setupOutputs() (tensor.service.TensorService method), 25

setupSources() (tensor.service.TensorService method), 25

sFlow (class in tensor.sources.sflow), 40

sflow (module), 40

sFlowReceiver (class in tensor.sources.sflow), 40

SMART (class in tensor.sources.linux.sensors), 35

SNMP (class in tensor.sources.snmp), 40

snmp (module), 40

SNMPCisco837 (class in tensor.sources.snmp), 41

SNMPConnection (class in tensor.sources.snmp), 41

Source (class in tensor.objects), 24

sourceWatchdog() (tensor.service.TensorService method), 25

startTimer() (tensor.objects.Source method), 24

startTimer() (tensor.sources.riemann.RiemannTCP method), 40

startTimer() (tensor.sources.sflow.sFlow method), 40

stop() (tensor.objects.Output method), 24

stop() (tensor.outputs.elasticsearch.ElasticSearch method), 44

stop() (tensor.outputs.riemann.RiemannTCP method), 43

stopTimer() (tensor.objects.Source method), 24

StringProducer (class in tensor.utils), 25

## T

tensor.aggregators (module), 23

tensor.interfaces (module), 23

tensor.logs.follower (module), 29

tensor.logs.parsers (module), 29

tensor.objects (module), 23

tensor.outputs.elasticsearch (module), 44

tensor.outputs.riemann (module), 43

tensor.protocol.elasticsearch (module), 27

tensor.protocol.icmp (module), 27

tensor.protocol.riemann (module), 27

tensor.protocol.sflow.protocol (module), 28

tensor.protocol.sflow.server (module), 28

tensor.protocol.ssh (module), 28

tensor.service (module), 24

tensor.sources.database.elasticsearch (module), 31

tensor.sources.database.memcache (module), 32

tensor.sources.database.postgresql (module), 31

tensor.sources.docker (module), 32

tensor.sources.generator (module), 33

tensor.sources.haproxy (module), 33

tensor.sources.linux.basic (module), 33

tensor.sources.linux.process (module), 35  
tensor.sources.linux.sensors (module), 35  
tensor.sources.media.libav (module), 36  
tensor.sources.munin (module), 36  
tensor.sources.network (module), 36  
tensor.sources.nginx (module), 37  
tensor.sources.python.uwsgi (module), 38  
tensor.sources.rabbitmq (module), 38  
tensor.sources.redis (module), 39  
tensor.sources.riak (module), 39  
tensor.sources.riemann (module), 40  
tensor.sources.sflow (module), 40  
tensor.sources.snmp (module), 40  
tensor.utils (module), 25  
TensorService (class in tensor.service), 24  
tick() (tensor.objects.Source method), 24  
    tick() (tensor.outputs.elasticsearch.ElasticSearch  
        method), 44  
    tick() (tensor.outputs.riemann.RiemannTCP method), 43  
Timeout, 25

## U

uwsgi (module), 38