
Tensor Documentation

Release 0.0.6

Colin Alston

February 07, 2015

1	Getting started	3
1.1	Installation	3
1.2	Creating a configuration file	3
1.3	Starting Tensor	4
2	Sources	5
2.1	Introduction	5
2.2	Writing your own sources	5
2.3	Handling asynchronous tasks	6
2.4	Thinking outside the box	7
3	Outputs	9
3.1	Introduction	9
3.2	Writing your own outputs	9
4	Example configurations	11
4.1	Replacing Munin	11
5	tensor	19
5.1	tensor.interfaces	19
5.2	tensor.objects	19
5.3	tensor.service	20
5.4	tensor.utils	20
6	tensor.protocol	23
6.1	tensor.protocol.riemann	23
6.2	tensor.protocol.sflow	23
7	tensor.sources	25
7.1	tensor.sources.network	25
7.2	tensor.sources.nginx	25
7.3	tensor.sources.munin	26
7.4	tensor.sources.riak	27
7.5	tensor.sources.rabbitmq	27
7.6	tensor.sources.linux	27
7.7	tensor.sources.media	29
7.8	tensor.sources.sflow	29
7.9	tensor.sources.snmp	30

8	tensor.outputs	31
8.1	tensor.outputs.riemann	31
9	Indices and tables	33
	Python Module Index	35

Tensor is a modular gateway and event router for Riemann, built using the Twisted framework.

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/0.2.0/tensor_0.2.0_amd64.deb
$ dpkg -i tensor_0.2.0_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

The 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.

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

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

If you enable multiple outputs then the *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.

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.

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",
}
```

1.3 Starting Tensor

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

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


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 implements

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

class HelloWorld(Source):
    implements(ITensorSource)

    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 implements

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

class HelloWorld(Source):
    implements(ITensorSource)

    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 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 implements

from twisted.internet import defer

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

class ProcessCount(Source):
    implements(ITensorSource)

    @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.

2.4 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 implements

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)

class NumberProxy(Source):
    implements(ITensorSource)

    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

```
output: tensor.sources.network.Ping
```

3.2 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 0x7f3b5
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
```

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"}'
```

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 -zxvf 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 $timeFilter group by time($
```

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:

5.1 tensor.interfaces

5.2 tensor.objects

class `tensor.objects.Event` (*state, service, description, metric, ttl, tags=[]*, *hostname=None, aggregation=None, evtime=None*)

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
- **tags** – List of tag strings
- **hostname** – Hostname for the event (defaults to system fqdn)
- **aggregation** – Aggregation function
- **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

startTimer()

Starts 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

5.3 tensor.service

class *tensor.service.TensorService*(*config*)

Bases: *twisted.application.service.Service*

Tensor service

Runs timers, configures sources and manages the queue

sendEvent(*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

5.4 tensor.utils

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

Bases: *twisted.internet.protocol.Protocol*

Simple buffering consumer for body objects

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

`tensor.utils.fork` (*executable, args=(), env={}, path=None, timeout=3600*)
Provides a deferred wrapper function with a timeout function

Parameters

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

tensor.protocol

6.1 tensor.protocol.riemann

class `tensor.protocol.riemann.RiemannClientFactory`

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.RiemannProtobuf`

Riemann protobuf protocol

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

Bases: `twisted.internet.protocol.DatagramProtocol`, `tensor.protocol.riemann.RiemannProtobuf`

UDP datagram protocol for Riemann

6.2 tensor.protocol.sflow

6.2.1 tensor.protocol.sflow.server

class `tensor.protocol.sflow.server.DatagramReceiver`

Bases: `twisted.internet.protocol.DatagramProtocol`

DatagramReceiver for sFlow packets

6.2.2 tensor.protocol.sflow.protocol

tensor.sources

7.1 tensor.sources.network

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

Bases: `tensor.objects.Source`

Performs an HTTP request

Configuration arguments:

Parameters

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

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

This is a horrible implementation which forks to *ping*

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.

7.2 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.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)

Metrics:

`(service name).total_bytes` Bytes total for all requests

`(service name).total_requests` Total request count

`(service name).stats.(code).(requests|bytes)` Metrics by status code

`(service name).user-agent.(agent).(requests|bytes)` Metrics by user agent

`(service name).client.(ip).(requests|bytes)` Metrics by client IP

`(service name).request.(request path).(requests|bytes)` Metrics by request path

7.3 tensor.sources.munin

`class tensor.sources.munin.MuninNode (*a, **kw)`

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

7.4 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

7.5 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

7.6 tensor.sources.linux

7.6.1 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

Metrics:

(service name).(device) Used space (%)

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

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

7.6.2 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

7.7 tensor.sources.media

7.7.1 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.

7.8 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).(in/out)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

7.9 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

8.1 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)
- **maxrate** (*int.*) – Maximum de-queue rate (0 is no limit)
- **interval** (*float.*) – De-queue interval in seconds (default: 1.0)
- **pressure** (*int.*) – Maximum backpressure (-1 is no limit)

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*

Indices and tables

- *genindex*
- *modindex*
- *search*

m

`munin` (*Any*), 26

n

`network` (*Unix*), 25

`nginx` (*Unix*), 25

r

`riak` (*Any*), 27

s

`sflow` (*Unix*), 29

`snmp` (*Unix*), 30

t

`tensor.interfaces`, 19

`tensor.objects`, 19

`tensor.outputs.riemann`, 31

`tensor.protocol.riemann`, 23

`tensor.protocol.sflow.protocol`, 23

`tensor.protocol.sflow.server`, 23

`tensor.service`, 20

`tensor.sources.linux.basic`, 27

`tensor.sources.linux.process`, 28

`tensor.sources.media.libav`, 29

`tensor.sources.munin`, 26

`tensor.sources.network`, 25

`tensor.sources.nginx`, 25

`tensor.sources.rabbitmq`, 27

`tensor.sources.riak`, 27

`tensor.sources.sflow`, 29

`tensor.sources.snmp`, 30

`tensor.utils`, 20

B

BodyReceiver (class in tensor.utils), 20

C

CPU (class in tensor.sources.linux.basic), 27

createClient() (tensor.objects.Output method), 19

createClient() (tensor.outputs.riemann.RiemannTCP method), 31

createClient() (tensor.outputs.riemann.RiemannUDP method), 32

createEvent() (tensor.objects.Source method), 20

D

DarwinRTSP (class in tensor.sources.media.libav), 29

DatagramReceiver (class in tensor.protocol.sflow.server), 23

DiskFree (class in tensor.sources.linux.basic), 28

E

emptyQueue() (tensor.outputs.riemann.RiemannTCP method), 31

Event (class in tensor.objects), 19

eventsReceived() (tensor.objects.Output method), 19

eventsReceived() (tensor.outputs.riemann.RiemannTCP method), 31

eventsReceived() (tensor.outputs.riemann.RiemannUDP method), 32

F

fork() (in module tensor.utils), 21

H

HTTP (class in tensor.sources.network), 25

L

LoadAverage (class in tensor.sources.linux.basic), 28

M

Memory (class in tensor.sources.linux.basic), 28

munin (module), 26

MuninNode (class in tensor.sources.munin), 26

MuninProtocol (class in tensor.sources.munin), 27

N

Network (class in tensor.sources.linux.basic), 28

network (module), 25

Nginx (class in tensor.sources.nginx), 25

nginx (module), 25

NginxLogMetrics (class in tensor.sources.nginx), 26

O

Output (class in tensor.objects), 19

P

Ping (class in tensor.sources.network), 25

ProcessCount (class in tensor.sources.linux.process), 28

ProcessProtocol (class in tensor.utils), 20

ProcessStats (class in tensor.sources.linux.process), 28

Q

Queues (class in tensor.sources.rabbitmq), 27

R

Resolver (class in tensor.utils), 21

riak (module), 27

RiakStats (class in tensor.sources.riak), 27

RiemannClientFactory (class in tensor.protocol.riemann), 23

RiemannProtocol (class in tensor.protocol.riemann), 23

RiemannTCP (class in tensor.outputs.riemann), 31

RiemannUDP (class in tensor.outputs.riemann), 31

RiemannUDP (class in tensor.protocol.riemann), 23

S

sendEvent() (tensor.service.TensorService method), 20

setupOutputs() (tensor.service.TensorService method), 20

setupSources() (tensor.service.TensorService method), 20

sFlow (class in tensor.sources.sflow), 29

sflow (module), 29

sFlowReceiver (class in tensor.sources.sflow), 29
SNMP (class in tensor.sources.snmp), 30
snmp (module), 30
SNMPCisco837 (class in tensor.sources.snmp), 30
SNMPConnection (class in tensor.sources.snmp), 30
Source (class in tensor.objects), 20
startTimer() (tensor.objects.Source method), 20
startTimer() (tensor.sources.sflow.sFlow method), 29
stop() (tensor.objects.Output method), 20
stop() (tensor.outputs.riemann.RiemannTCP method), 31

T

tensor.interfaces (module), 19
tensor.objects (module), 19
tensor.outputs.riemann (module), 31
tensor.protocol.riemann (module), 23
tensor.protocol.sflow.protocol (module), 23
tensor.protocol.sflow.server (module), 23
tensor.service (module), 20
tensor.sources.linux.basic (module), 27
tensor.sources.linux.process (module), 28
tensor.sources.media.libav (module), 29
tensor.sources.munin (module), 26
tensor.sources.network (module), 25
tensor.sources.nginx (module), 25
tensor.sources.rabbitmq (module), 27
tensor.sources.riak (module), 27
tensor.sources.sflow (module), 29
tensor.sources.snmp (module), 30
tensor.utils (module), 20
TensorService (class in tensor.service), 20
tick() (tensor.objects.Source method), 20
tick() (tensor.outputs.riemann.RiemannTCP method), 31