
MmStats Documentation

Release 0.6.2

Michael Schurter

October 02, 2012

CONTENTS

[Documentation](#) | [Package](#) | [Code](#)

ABOUT

Mmstats is a way to expose and read diagnostic values and metrics for applications.

Think of mmstats as /proc for your application and the readers as procps utilities.

This project is a Python implementation, but compatible implementations can be made in any language (see Goals).

Discuss at <https://groups.google.com/group/python-introspection>

1.1 Goals

- Separate publishing/writing from consuming/reading tools
- Platform/language independent (a Java writer can be read by a Python tool)
- Predictable performance impact for writers via:
 - No locks (1 writer per thread)
 - No syscalls (after instantiation)
 - All in userspace
 - Reading has no impact on writers
- Optional persistent (writer can sync anytime)
- 1-way (Publish/consume only; mmstats are not management extensions)

USAGE

2.1 Requirements

Cython 2.6 or 2.7 (Windows is untested)

PyPy (only tested in 1.7, should be faster in 1.8)

2.2 Using

1. `python setup.py install`
2. `import mmstats`
3. Create a subclass of `mmstats.MmStats` like:

```
class WebStats(mmstats.MmStats):  
    status2xx = mmstats.CounterField(label='status.2XX')  
    status3xx = mmstats.CounterField(label='status.3XX')  
    status4xx = mmstats.CounterField(label='status.4XX')  
    status5xx = mmstats.CounterField(label='status.5XX')  
    last_hit = mmstats.DoubleField(label='timers.last_hit')
```

4. Instantiate it once per thread/process:

```
webstats = WebStats(label_prefix='web.stats.')
```

5. Record some data:

```
if response.status_code == 200:  
    webstats.status2xx.inc()
```

```
webstats.last_hit = time.time()
```

6. Run `slurpstats` to read it

7. Run `mmash` to create a web interface for stats

8. Run `pollstats -p web.stats.status 2XX,3XX,4XX,5XX /tmp/mmstats-*` for a vmstat/dstat like view.

9. Did a process die unexpectedly and leave around a stale mmstat file? `cleanstats /path/to/mmstat/files` will check to see which files are stale and remove them.

2.3 Development

It's easiest to develop mmstats within a virtualenv:

```
$ git clone git://github.com/schmichael/mmstats.git
$ cd mmstats
$ virtualenv .
$ source bin/activate
$ python setup.py develop
$ ./run_flask_example # This starts up a sample web app
$ curl http://localhost:5001/
$ curl http://localhost:5001/500
$ curl http://localhost:5001/status
$ # If you have ab installed:
$ ab -n 50 -c 10 http://localhost:5001/
```

Now to view the stats run the following in a new terminal:

```
$ # To get a raw view of the data:
$ slurpstats mmstats-*
$ # Or start up the web interface:
$ mmash
$ # Run pollstats while ab is running:
$ pollstats -p flask.example. ok,bad,working mmstats-*
```

To cleanup stray mmstats files: rm mmstats-flask-*

The web interface will automatically reload when you change source files.

Put static files into static/ and template files into templates/

TODO: Factor mmash out into its own project (with a meaningful name?)

2.4 Testing

Feel free to use your favorite test runner like [nose](#) or [pytest](#) or just run:

```
$ python setup.py test
```

FURTHER TOPICS

3.1 Development

3.1.1 mmap Format

Structure of version 1 mmstat's mmaps:

version number	fields...
byte = 01	...

Fields

There are two types of field structures so far in mmstats:

1. buffered
2. unbuffered

Buffered fields use multiple buffers for handling values which cannot be written atomically.

Unbuffered structures have ff in the write buffer field.

Buffered

label size	label	type size	type	write buffer	buffer 1	buffer 2
ushort	char[]	ushort	char[]	byte	varies	varies

The buffers field length = sizeof(type) * buffers.

The current write buffer is referenced by: write_buffer * sizeof(type)

TODO: field for total number of buffers?

Unbuffered

label size	label	type size	type	write buffer	value
ushort	char[]	ushort	char[]	byte = ff	varies

The value field length = sizeof(type).

3.2 History

3.2.1 0.6.2 “Graphtastic” released 2012-03-23

- Added live graphing of numeric metrics thanks to @haard’s work at PyCon
- Documentation improvements

3.2.2 0.6.1 “MANIFEST.out” released 2012-03-08

- Fix packaging issue

3.2.3 0.6.0 “PyCon 2012” released 2012-03-08

- [API CHANGE] - MovingAverageField’s kwarg changed from window_size => size
- Refactored __init__.py into fields, models, and default (and imported public bits into __init__)
- Added TimerField (MovingAverageField + context manager)
- Added docs (don’t get too excited, just a start)

3.2.4 0.5.0 “100% More Average” released 2012-02-25

- [API CHANGE] - RunningAverage field is now AverageField
- Added MovingAverageField with window_size=100 parameter
- Tests can now be run via “python setup.py test”

3.2.5 0.4.1 “Derpstats” released 2012-01-31

- Fixed pollstats
- Updated README slightly

3.2.6 0.4.0 “On the Road to Pycon” released 2012-01-17

- Added clean module and cleanstats script to clean stale mmstat files
- Added path kwarg to MmStats class to allow easy path overriding
- Added StringField for UTF-8 encoded strings
- Added StaticFloatField & StaticDoubleField
- Added created UNIX timestamp (sys.created) to default MmStats class
- Moved all modules into mmstats package
- Fixed mmash template packaging
- Fixed test mmstat file cleanup
- Refactored reading code into mmstats.reader module

3.2.7 0.3.12 “Meow” released 2011-11-29

- Use ctypes.get_errno() instead of Linux specific wrapper

3.2.8 0.3.11 “Rawr” released 2011-11-29

- Fix libc loading on OSX

3.2.9 0.3.10 “ π^2 ” released 2011-11-28

- PyPy support (switched from ctypes._CData.from_buffer to .from_address)
- Multiple calls to MmStats().remove() no longer error (makes testing easier)

3.2.10 0.3.9 “MLIT” released 1970-01-01

- Mistag of 0.3.8

3.2.11 0.3.8 “Hapiness” released 2011-11-20

- Allow filename templating with %PID% and %TID% placeholders
- Allow setting filename template via MMSTATS_FILES environment variable
- Improved docs slightly
- Fixed Ctrl-Cing run_flask_example script
- Fixed 64 bit integer fields on 32 bit platforms
- Fixed StaticInt64Field (was actually a uint64 field before)
- Fixed slurpstats debug output (always showed first 40 bytes of file)
- Strip newlines from org.python.version

3.2.12 0.3.7 “Depressive Realism is for Winners” released 2011-11-17

- Add pollstats utility (similar to dstat/vmstat)
- Cleanup development/testing section of the README
- Slight improvements to basic flask integration example

3.2.13 0.3.6 “The M is for Mongo” released 2011-11-09

- Allow setting the value of CounterFields

3.2.14 0.3.5 “Ornery Orangutan” released 2011-10-20

- Added a running average field
- Made mmash more configurable and added a console entry point
- Updated TODO

3.3 Contributors

Copyright 2012 Urban Airship and contributors

3.3.1 Developers

- Dan Colish <dcolish@gmail.com>
- Fredrik Håård - <https://github.com/haard>
- Adam Lowry <adam@therobots.org>
- Michael Schurter <m@schmichael.com> (creator/lead)

3.3.2 Reviewers

- Niall Kelly

API REFERENCE

4.1 API

4.1.1 Models

```
class mmstats.models.BaseMmStats(path='/tmp',      filename='mmstats-%PID%-%TID%',      la-  
bel_prefix=None)
```

Stats models should inherit from this

```
flush(async=False)
```

Flush mmaped file to disk

```
remove()
```

Close and remove mmap file - No further stats updates will work

```
class mmstats.models.FieldState(field)
```

Holds field state for each Field instance

```
class mmstats.models.MmStats(path='/tmp',      filename='mmstats-%PID%-%TID%',      la-  
bel_prefix=None)
```

Mmstats default model base class

Just subclass, add your own fields, and instantiate:

```
>>> from mmstats.models import MmStats  
>>> from mmstats.fields import CounterField  
>>> class MyStats(MmStats):  
...     errors = CounterField()  
...  
>>> stats = MyStats()  
>>> stats.errors.inc()  
>>> stats.errors.value  
1L
```

4.1.2 Model Fields

```
mmstats.fields.AverageField
```

Average field supporting an add() method and value attribute

```
mmstats.fields.BoolField
```

Boolean Field

```
mmstats.fields.BufferedDescriptorField
```

Base class for double buffered descriptor fields

```
mmstats.fields.BufferedDescriptorMixin
    Mixin to add double buffered descriptor methods

    Always read/write as double buffering doesn't make sense for readonly fields

mmstats.fields.ByteField
    8bit Signed Integer Field

mmstats.fields.ComplexDoubleBufferedField
    Base Class for fields with complex internal state like Counters

    Set InternalClass in your subclass

mmstats.fields.CounterField
    Counter field supporting an inc() method and value attribute

mmstats.fields.DataDescriptorMixin
    Mixin to add single buffered __set__ method

class mmstats.fields.DoubleBufferedField (label=None)
    Base class for double buffered writable fields

mmstats.fields.DoubleField
    64bit Double Precision Float Field

exception mmstats.fields.DuplicateFieldName
    Cannot add 2 fields with the same name to MmStat instances

mmstats.fields.FloatField
    32bit Float Field

mmstats.fields.IntegerField
    32bit Double Buffered Signed Integer field

mmstats.fields.NonDataDescriptorMixin
    Mixin to add single buffered __get__ method

mmstats.fields.ReadWriteField
    Base class for simple writable fields

mmstats.fields.ShortField
    16bit Double Buffered Signed Integer field

mmstats.fields.StaticDoubleField
    Unbuffered read-only 64bit Float field

mmstats.fields.StaticFloatField
    Unbuffered read-only 32bit Float field

mmstats.fields.StaticInt64Field
    Unbuffered read-only 64bit Signed Integer field

mmstats.fields.StaticTextField
    Unbuffered read-only UTF-8 encoded String field

mmstats.fields.StaticUInt64Field
    Unbuffered read-only 64bit Unsigned Integer field

mmstats.fields.StaticUIntField
    Unbuffered read-only 32bit Unsigned Integer field

mmstats.fields.StringField
    UTF-8 String Field
```

mmstats.fields.TimerField

Moving average field that provides a context manager for easy timings

As a context manager: >>> class T(MmStats): ... timer = TimerField() >>> t = T() >>> with t.timer as ctx: ... assert ctx.elapsed > 0.0 >>> assert t.timer.value > 0.0 >>> assert t.timer.last > 0.0

mmstats.fields.UInt64Field

Unbuffered read-only 64bit Unsigned Integer field

mmstats.fields.UIntField

32bit Double Buffered Unsigned Integer field

mmstats.fields.UShortField

16bit Double Buffered Unsigned Integer field

4.1.3 Reader API

mmstats reader implementation

exception mmstats.reader.InvalidMmStatsVersion

Unsupported mmstats version

class mmstats.reader.Stat

Stat(label, value)

label

Alias for field number 0

value

Alias for field number 1

4.1.4 Internal Defaults

4.1.5 mmap Compatibility Wrapper

mmstats._mmap.init_mmap(path='/tmp', filename='mmstats-%PID%-%TID%', size=4096)

Given path, filename => filename, size, mmap

In *filename* “%PID%” and “%TID%” will be replaced with pid and thread id

INDICES AND TABLES

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

PYTHON MODULE INDEX

m

mmstats._mmap, ??
mmstats.defaults, ??
mmstats.fields, ??
mmstats.models, ??
mmstats.reader, ??