
ThriftPy Documentation

Release 0.3.9

Lx Yu

December 07, 2016

1	Code Demo	3
2	Features	5
3	Installation	7
4	Usage Notice	9
4.1	Cython Binary Protocol	9
4.2	Better Module	9
5	Benchmarks	11
6	Contribute	13
7	Contributors	15
8	Changelog	17
8.1	Version 0.3.9	17
8.2	Version 0.3.8	17
8.3	Version 0.3.7	17
8.4	Version 0.3.6	17
8.5	Version 0.3.5	18
8.6	Version 0.3.4	18
8.7	Version 0.3.3	18
8.8	Version 0.3.2	18
8.9	Version 0.3.1	19
8.10	Version 0.3.0	19
8.11	Version 0.2.1	19
8.12	Version 0.2.0	19
8.13	Version 0.1.15	20
8.14	Version 0.1.14	20
8.15	Version 0.1.13	20
8.16	Version 0.1.12	21
8.17	Version 0.1.11	21
8.18	Version 0.1.10	21
8.19	Version 0.1.9	21
8.20	Version 0.1.8	21
8.21	Version 0.1.7	22
8.22	Version 0.1.6	22

8.23	Version 0.1.5	22
8.24	Version 0.1.4	22
8.25	Version 0.1.3	23
8.26	Version 0.1.2	23
8.27	Version 0.1.1	23

ThriftPy is a pure python implementation of [Apache Thrift](#) in a pythonic way.

The official thrift python lib is not pythonic at all, it needs a complicated process of installation, and the generated sdk is very ugly. Everytime the thrift file changed you have to re-generate the sdk which causes more pain in development.

ThriftPy helps that, it's compatible with Apache Thrift so you no longer need to install 'thrift' package, it can import thrift file on the fly so you no longer need to re-generate the sdk again and again and again.

Github: <https://github.com/eleme/thriftpy>

Code Demo

ThriftPy make it super easy to write server/client code with thrift. Let's checkout this simple pingpong service demo.

We need a 'pingpong.thrift' file:

```
service PingPong {
    string ping(),
}
```

Then we can make a server:

```
import thriftpy
pingpong_thrift = thriftpy.load("pingpong.thrift", module_name="pingpong_thrift")

from thriftpy.rpc import make_server

class Dispatcher(object):
    def ping(self):
        return "pong"

server = make_server(pingpong_thrift.PingPong, Dispatcher(), '127.0.0.1', 6000)
server.serve()
```

And a client:

```
import thriftpy
pingpong_thrift = thriftpy.load("pingpong.thrift", module_name="pingpong_thrift")

from thriftpy.rpc import make_client

client = make_client(pingpong_thrift.PingPong, '127.0.0.1', 6000)
client.ping()
```

See, it's that easy!

You can refer to 'examples' and 'tests' directory in source code for more usage examples.

Features

Currently ThriftPy have these features (also advantages over the upstream python lib):

- Supports python2.6+, python3.3+, pypy and pypy3.
- Pure python implementation. No longer need to compile & install the ‘thrift’ package. All you need is thriftpy and thrift file.
- Compatible with Apache Thrift. You can use ThriftPy together with the official implementation servers and clients, such as a upstream server with a thriftpy client or the opposite.

Currently implemented protocols and transports:

- binary protocol (python and cython)
 - compact protocol (python and cython)
 - json protocol
 - buffered transport (python & cython)
 - framed transport
 - tornado server and client (with tornado 4.0)
- Can directly load thrift file as module, the sdk code will be generated on the fly.

For example, `pingpong_thrift = thriftpy.load("pingpong.thrift", module_name="pingpong_thrift")` will load ‘pingpong.thrift’ as ‘pingpong_thrift’ module.

Or, when import hook enabled by `thriftpy.install_import_hook()`, you can directly use `import pingpong_thrift` to import the ‘pingpong.thrift’ file as module, you may also use `from pingpong_thrift import PingService` to import specific object from the thrift module.

- Easy RPC server/client setup.

Installation

Install with pip.

```
$ pip install thriftpy
```

You may also install cython first to build cython extension locally.

```
$ pip install cython thriftpy
```

Usage Notice

4.1 Cython Binary Protocol

The cython accelerated binary protocol is enabled by default for CPython if available, and is disabled for Pypy.

To force use pure python version of binary protocol, you must import them from the direct module.

```
from thrift.protocol.binary import TBinaryProtocolFactory
from thrift.transport.buffered import TBufferedTransportFactory
from thrift.transport.framed import TFramedTransportFactory
```

4.2 Better Module

To load thrift file as better module, provide a *module_name* in *load*.

The direct loaded TObjects can't be pickled.

```
>>> ab = thrift.load("addressbook.thrift")
>>> pickle.dumps(ab.AddressBook())
PicklingError: Can't pickle <class 'addressbook.AddressBook'>
```

TObjects can be pickled when load with *module_name* provided.

```
>>> ab = thrift.load("addressbook.thrift", "addressbook_thrift")
>>> pickle.dumps(ab.AddressBook())
b'\x80\x03addressbook_thrift\nAddressBook\nq\x00}\x81q\x01}q\x02X\x06\x00\x00\x00peopleq\x03Nsb.'
```

You can also use *from ... import ...* style after a standard module load.

```
>>> ab = thrift.load("addressbook.thrift", "addressbook_thrift")
>>> from addressbook_thrift import *
```

Benchmarks

Some benchmark results:

```
# apache thrift py binary
binary protocol struct benchmark for 100000 times:
encode -> 3.74061203003
decode -> 5.02829790115

# apache thrift c binary
accelerated protocol struct benchmark for 100000 times:
encode -> 0.398949146271
decode -> 0.536000013351

# thriftpy & pypy2.3
binary protocol struct benchmark for 100000 times:
encode -> 0.413738965988
decode -> 0.605606079102

# thriftpy & py3.4
binary protocol struct benchmark for 100000 times:
encode -> 3.291545867919922
decode -> 4.337666034698486

# thriftpy & py3.4 + cython
cybin protocol struct benchmark for 100000 times:
encode -> 0.5828649997711182
decode -> 0.8259570598602295
```

Checkout the *benchmark/benchmark.rst* for detailed benchmark scripts and scores.

Contribute

1. Fork the repo and make changes.
2. Write a test which shows a bug was fixed or the feature works as expected.
3. Make sure travis-ci test succeed.
4. Send pull request.

Contributors

<https://github.com/eleme/thriftpy/graphs/contributors>

Changelog

8.1 Version 0.3.9

Released on August 26, 2016.

- add support for timeout and ssl in *make_server / make_client* helper funcs, via #204, #205 and #229.
- add support for *thrift_file* path in http protocol, via #225.
- preserve traceback when re-raise undeclared exception, via #206.
- performance improvement by dynamically compile spec'd *__init__* functions, via #210 and #227.
- performance improvement by refine cython encoding/decoding, via #211 and #212_.
- bugfix for type error in *cast_byte* parser and improve include dirs function, via #214
- bugfix for parse error when field begin with true/false keyword, via #215 and #218.
- bugfix for *is_open* not return false when socket closed after open, via #230.

8.2 Version 0.3.8

Released on May 3, 2016.

- add propagate *decode_response* to nested structs, via #194.
- add support for tornado ssl, via #196.

8.3 Version 0.3.7

Released on Mar 24, 2016.

- bugfix for a possible unicode decode error in *cybin*.
- use a better unhashable implementation for payload.

8.4 Version 0.3.6

Released on Mar 24, 2016.

- add compact protocol support, via #159.
- add option to force return bytes on response, via #190.
- bugfix for ssl socket can't be init without certfile and keyfile, and add additional *capath* argument for SSLContext. via #186.
- bugfix for set_timeout only works before socket open, via #188.

8.5 Version 0.3.5

Released on Feb 16, 2016.

- fix another set_timeout backward compat issue introduced in last version.
- make thrift container struct unhashable, via #184.

8.6 Version 0.3.4

Released on Feb 3, 2016.

- fix backward compat issue introduced in last version, add back *set_timeout* api in socket.

8.7 Version 0.3.3

Released on Jan 21, 2016.

- add support for ssl transport.
- add named loggers, via #169.
- refine socket and serversocket implementation with more configure options.
- bugfix for parser failure on windows under python3.2 caused by samefile method, via #172.

8.8 Version 0.3.2

Released on Oct 12, 2015.

- add *__thrift_meta__* attribute to loaded module, via #138.
- add type validation before write data to transport, via #149 and #150.
- add load_fp api to load thrift from file like object, via #154.
- add support for recursive struct definition, via #155.
- add support for integer boolean constants, via #161.
- simplify the read_i08() bool result cast, via #162.
- performance improvements on payload *init()* func, via #163.
- bugfix for parsing of duplicate field name or id, now will raise error when duplicates detected, via #139.
- bugfix for server side transport not connected error when closing socket, via #143.

- bugfix for a typo error in *default_spec* generation, via #145.
- bugfix for i16 byte swap bug in OS X, via #148.

8.9 Version 0.3.1

Released on May 29, 2015.

- lock down to use pure python only in windows env. (this avoid the cython stuffs on windows totally)
- enable multiple include dirs, via #131.
- bugfix for parsing of constants with separators, via #134.

8.10 Version 0.3.0

Released on April 15, 2015.

Non-Backward Compatible changes:

- migrate multiplexed protocol implementation to the same with upstream, via #117.

8.10.1 0.2.x

8.11 Version 0.2.1

Released on April 15, 2015.

- add an experimental tracking feature in *thriftpy.contrib*, via #96.
- add limitation on thrift reserved keyword for compatible with upstream, via #115.
- bugfix EOF grammar error, via #103.
- bugfix for mis-mach transport in client caused server crash, via #119.
- bugfix for typedef on included thrift files, via #121.

8.12 Version 0.2.0

Released on March 3, 2015.

- support for default enum values that reference the original enum, via #69.
- support for *require* keyword, via #72.
- support for allow use and definition of types in the same file, via #77.
- support for multiplexing for services, via #88.
- support for cython accelerated memory transport and framed transport, via #93
- bugfix for transport clean in read_struct in cybin, via #70.
- bugfix for large reading size in framed transport, via #73.
- bugfix for cython build failed in older CentOS, via #92.

- bugfix for thrift file version mis-match caused message corrupt in `read_struct`, via #95.

Non-Backward Compatible changes:

- refined new parser, the parser now behaves very similar to Apache Thrift, and supports a lot more features than the old one, via #80. Refer to the pull request for more detailed changes.
- refined transport, all transports have cython accelerated version. The cython version of protocol and transport are enabled by default now.

8.12.1 0.1.x

8.13 Version 0.1.15

Released on December 12, 2014.

- add MIT *LICENSE* file as requested.
- tests refines with tox and pytest fixtures.
- support for a mostly cythonized version of framed transport, via #66.
- bugfix for unix socket param in rpc.
- bugfix for receiving 0-length strings & framed transport, via #63.
- bugfix for json protocol unicode decode error, via #65.
- bugfix for operator `__ne__` implementation error, via #68.

8.14 Version 0.1.14

Released on November 8, 2014.

- support for python2.6.
- support for testing by tox.
- support for oneway keyword, via #49.
- bugfix for wrong type args, via #48.
- bugfix for thrift file include keyword, via #53.
- bugfix for skip method not found in protocol, via #55.
- bugfix for set type support, via #59.
- bugfix for 'api' arg name collision in client.

8.15 Version 0.1.13

Released on September 24, 2014.

- bugfix for TPayload not able to be hashed in py3, via #44.
- bugfix for cython buffered transport read issue, via #46.

8.16 Version 0.1.12

Released on September 18, 2014.

- bugfix for lack of *skip* func in cython binary protocol, via #43.

8.17 Version 0.1.11

Released on September 16, 2014.

- bugfix for init func generator for TStruct.
- bugfix for set constants in parser, via #39.
- add support for “includes” and service “extends”, via #37.
- add close() to servers, via #38.
- implement non-strict mode for binary protocol, via #40.
- removed cython ext in pypy, and add pypy3 support.
- some args updates: * add *trans_factory* arg to *make_server* * rename *rbuf_size* in buffered transport to *buf_size*. * rename *isOpen* to *is_open*, *readFrame* to *read_frame*.

8.18 Version 0.1.10

Released on September 4, 2014.

- bugfix for memory free in cython buffered transport, via #35.
- new thrift parser by PLY, removed cache since the performance is much more faster now, via #36.

8.19 Version 0.1.9

Released on September 1, 2014.

- refine cython binary protocol, add cython buffered transport, via #32.
- param name change, rename *transport_factory* to *trans_factory* in rpc.

8.20 Version 0.1.8

Released on August 28, 2014.

- faster thrift file parse speed, via #30.
- bugfix for cybin buffer read, via #31.

8.21 Version 0.1.7

Released on August 19, 2014.

- use args instead of kwargs in api calling to match upstream behavior.
- cython binary protocol auto grow buffer size, via #29.
- bugfix for void api exception handling in processor.
- bugfix for cybin protocol buffer overflow and memcpy, via #27 and #28.

8.22 Version 0.1.6

Released on August 14, 2014.

- json protocol, via #21.
- more standard module for loaded sdk, now generated TPayload objects can be pickled when module_name provided, via #22.
- unicorn_thrift integration pingpong example, via #24.
- token cache now only checks python's major and minor version.
- bugfix for exception handling in void api in RPC request.
- bugfix for negative number value not recognized.
- bugfix for cybin protocol to allow None value in struct.
- bugfix for double free or corruption in cybin protocol, via #26.

8.23 Version 0.1.5

Released on July 25, 2014.

- tornado client, server and framed transport support with tornado 4.0, via #15.
- immediately read from TMemoryBuffer after writing to it, via #20.
- cache *load* function to avoid duplicate module generation.
- support client with socket timeout
- enum struct now has VALUES_TO_NAMES and NAMES_TO_VALUES.

8.24 Version 0.1.4

Released on July 17, 2014.

- parser token cache, speed boost for thrift file parsing, via #12.
- new cython binary protocol with speed very close to c ext, via #16.

8.25 Version 0.1.3

Released on June 19, 2014.

- support for union, binary fields, support for empty structs, support for Apache Storm thrift file, via #14.
- bugfix for import hook
- bugfix for skip function in binary protocols

8.26 Version 0.1.2

Released on June 7, 2014.

- disabled the magic import hook by default. and add install/remove function to switch the hook on and off.
- reworked benchmark suit and add benchmark results.
- new `__init__` function code generator. get a noticable speed boost.
- bug fixes

8.27 Version 0.1.1

First public release.