

Taming Pythons with ZooKeeper



June 29, 2013

\$ whoami



Jyrki Pulliainen

@nailor

jyrki@spotify.com

CAUTION

Snake Lunges at Face
When Door Opens

ZooKeeper?



Wednesday, 3 July 13

- This talk is about ZooKeeper.
- Preaching Java software in Python conference. I can think of more healthier things to do too.
- On side note, our Product Owner is really damn good at this game

ZooKeeper!



Wednesday, 3 July 13

- Apache project, Yahoo 2007
- Consistency & Partition tolerance
- Filesystem like, actually can be viewed as a trie, store data in directories too
- In memory, limits the dataset you can store. Maximum zookeeper node, znode, data size 1M

The Tao of ZooKeeper



Wednesday, 3 July 13

Orderly, Reliable, Efficient, Timely, Contention free, ambition free

HOW DO I SHOT WEB



One library for the enterprise

Curator, from Netflix

Seven for the pythonistas

gevent-zookeeper

twitter's zookeeper library

zkpython

zoop

zc.zk

txzookeeper

pykeeper

Wednesday, 3 July 13

- gevent-zookeeper -> Spotify
- zkpython segfaults
- Others have somewhat OK implementations, but lack core features

One Library to rule them all*

* unless you are running twisted

Wednesday, 3 July 13

- txzookeeper still valid for twisted

Kazoo

Wednesday, 3 July 13

- Origins from the Nimbus project
- Ben Bangert as the Sauron of ZooKeeper
- Not Frodo, that bastard wanted to destroy the perfectly good ring.
- All Python, including the protocol. No more segfaults!

Let's talk about KaZoo

```
events, unlike the default Zookeeper watches. They will also be
called with a single argument, a
:class:`~kazoo.protocol.states.WatchedEvent` instance.

"""
def __init__(self, hosts='127.0.0.1:2181',
              timeout=10.0, client_id=None, max_retries=None,
              retry_delay=0.1, retry_backoff=2, retry_jitter=0.8,
              retry_max_delay=3600, handler=None, default_acl=None,
              auth_data=None, read_only=None, randomize_hosts=True):
    """Create a :class:`KazooClient` instance. All time arguments
    are in seconds.

    :param hosts: Comma-separated list of hosts to connect to
                  (e.g. 127.0.0.1:2181,127.0.0.1:2182).
    :param timeout: The longest to wait for a Zookeeper connection.
```

CRUD

Wednesday, 3 July 13

Create, Read, Update, Delete

All basic operations available as async too, but we'll focus on the synchronous use

First we need to connect!

```
from kazoo.client import KazooClient

zk = KazooClient(hosts='127.0.0.1:2181')
zk.start()
# ...
zk.stop()
```

Wednesday, 3 July 13

Easy to connect to one host, multiple host, with namespace....
ZooKeeper supports connection namespacing!
Bonus: get notified when the connection state changes

First we need to connect!

```
from kazoo.client import KazooClient
```

```
zk = KazooClient(hosts='127.0.0.1:2181')
```

```
zk.start()
```

```
# ...
```

```
zk.stop()
```

```
zk = KazooClient(hosts='127.0.0.1:2181,127.0.0.2:2181')
```

First we need to connect!

```
from kazoo.client import KazooClient
```

```
zk = KazooClient(hosts='127.0.0.1:2181')
```

```
zk.start()
```

```
# ...
```

```
zk.stop()
```

```
zk = KazooClient(hosts='127.0.0.1:2181,127.0.0.2:2181')
```

```
zk = KazooClient(hosts='127.0.0.1:2181/namespace,127.0.0.2:2181')
```

Create

```
zk.create( '/europython', b'2013' )
```

Wednesday, 3 July 13

Adding nodes easy, helpers exist for recursive creation

Ephemeral ZK feature, session + heartbeats, can't have children!

Incremental: guaranteed ever increasing 10 digit number in node name

Create

```
zk.create( '/europython', b'2013' )
```

```
zk.create( '/europython/jyrki', ephemeral=True )
```

Wednesday, 3 July 13

Adding nodes easy, helpers exist for recursive creation

Ephemeral ZK feature, session + heartbeats, can't have children!

Incremental: guaranteed ever increasing 10 digit number in node name

Create

```
zk.create( '/europython', b'2013' )
```

```
zk.create( '/europython/jyrki', ephemeral=True )
```

```
zk.create( '/europython/sequential', sequence=True )
```

Wednesday, 3 July 13

Adding nodes easy, helpers exist for recursive creation

Ephemeral ZK feature, session + heartbeats, can't have children!

Incremental: guaranteed ever increasing 10 digit number in node name

Read

```
zk.get( '/europython' )
```

Read

```
zk.get( '/europython' )
```

```
zk.exists( '/europython/jyrki' )
```

Read

```
zk.get( '/europython' )
```

```
zk.exists( '/europython/jyrki' )
```

```
zk.get_children( '/europython' )
```


Update & Delete


```
zk.set_data( '/europython/jyrki', b'nervous' )
```

Update & Delete

```
zk.set_data( '/europython/jyrki', b'nervous' )
```

```
zk.delete( '/europython/jyrki' )
```

**SO YEAH, LIBRARY IS NICE AND ALL BUT IF
YOU COULD TELL ME WHAT TO DO WITH IT**



THAT WOULD BE GREAT

quickmeme.com

Distributed locks?

Barrier?

Semaphores?

Counters?

Elections?

```
from kazoo.recipe import <your-favourite-thing>
```

Want to know when things change?

watchers

```
zk.exists('/europython/wine', watch=callback)
```

```
zk.get('/europython/wine', watch=callback)
```

```
zk.get_children('/europython/dinners', watch=callback)
```

```
from kazoo.recipe.watchers import DataWatch, ChildWatch
```

```
@DataWatch( '/path/to/node' )  
def data_callback(data, stat):  
    # ...  
    do_something
```

```
@ChildWatch( '/path/to/node' )  
def child_callback(children):  
    # ...  
    do_something
```


TESTS

What if something goes

WRONG?



`sys.exit()`

Stand back!

It's time for real life example

Spotify Premium

◀▶

🔍 Search

MAIN

Discover

Follow

Inbox

Play Queue

Devices

APPS

App Finder

Top Lists

Radio

Classify

The Guardian

Last.fm

Moodagent

Pitchfork

Songkick Concerts

Soundrop


Top Spots

PINK FLOYD

Wish You Were Here

PINK FLOYD

Wish You Were Here




Wish You Were Here - 2011 - R

Pink Floyd

New Release

You listened to Telepopmusik. Like to hear the new album?




Try Me Anyway / Fever
Telepopmusik

Save as a Playlist

Artist

Like Beach House and Junip?
Check out Toro Y Moi.



Toro Y Moi

+ Follow Artist


Chaz Bundick started making bedroom recordings under the name Toro y Moi in his native Columbia, South Carolina.

Album

You listened to The Rapture.
Here's an album you might like.

Activity


Who to Follow




Nick Barkas
63 followers

+ Follow


Followed by Felix Bruns




Pontus Alexander listened to
Imagine Reality
Uppermost




Sanna Putila listened to
Há Uma Música Do Povo
Mariza




Nick Ingall listened to
Pompeii
Bastille



Pontus Alexander listened to
Fearless
Uppermost



Roger Schildmeijer listened to
Höj volymen
Lilla Namo



Roger Schildmeijer listened to
Vissa tycker
Mange Schmidt, Timbu...

⏮⏪⏩⏭

0:02

5:35

🔁🔄

Wednesday, 3 July 13

SEARCHED FOR THE NEWEST JUSTIN BIEBER



Encryption keys were
not available to play it.

SUMMARY

Guess what, we're hiring

Thank You!

jyrki@spotify.com

@nailor

