

10th Central and Eastern European Software
Engineering Conference in Russia - CEE-SECR
2014

October 23 - 25, Moscow



Tarantool: a Russian open source story

Konstantin Osipov,
kostja@tarantool.org
<http://try.tarantool.org>

whoami


- MSU CMS alumni
- software engineer with MySQL AB 2003-2010: took part in creating MySQL 4.1, 5.0, 5.1, 5.5
- joined Mail.Ru in 2010, Tarantool/Box key/value store
- run a team of tarantoolers



tarantool, the database (1)

... or in the beginning...

- **unique** (among open source databases):
 - in-memory
 - transactional
 - master-master replication**= low latency + high availability**



tarantool, the database (2)



- **modern**
 - MsgPack data model
 - L1/L2 cache- optimized indexing
 - lock-free architecture

tarantool, the database (3)



- **open & extensible**

- pluggable index API
- pluggable storage engine API
- user-defined triggers on database events
- pluggable authentication

why MsgPack?

... it's like JSON, but fast and small ...

- value-based encoding
- extensible and schema-less
- up to 2x more compact than Redis

JSON 27 bytes

```
{ "compact": true, "schema": 0 }
```

MessagePack 18 bytes



in-memory + persistency



- atomic transactions (no locks)
 - begin, commit, rollback
- write ahead log for durability
- log compaction with MVCC-enhanced snapshots

high availability



- multi-master eventual consistency (async)
- vector clock to denote node position
- synchronous replication is coming!

extensions & plug-ins



- **Sophia** engine: diskstore
 - user-defined triggers:
 - on replace,
 - on commit, on rollback,
 - on connect, on disconnect
- ... wait, there is more with Lua*

tarantool, the app server (1)



- **freedom** to ship code to data
 - background fibers
 - introspection
 - connection-local functions
 - scripted startup



tarantool, the app server (2)



- **simple and small:**
 - libcoro
 - libeio
 - libevent
 - all bound to LuaJIT 2.0 via FFI

tarantool, the app server (3)



- **a toolkit**, not a cookbook
 - `require('fiber')`, `require('ipc')`, `require('socket')`,
 - `json`, `yaml`, `msgpack`, `digest (md5, sha, etc)` on board
 - `net.tarantool`, `net.mysql`, `net.postgresql`

Lua all the way down



- `#!/usr/bin/tarantool`
- non-blocking I/O = Node.js + Redis in one tool
- Lua packages, practices and community = mutual endorsement

next slide in the deck

```
1 #!/usr/bin/env tarantool
2
3 box.cfg{}
4 hosts = box.space.hosts
5 if not hosts then
6     hosts = box.schema.create_space('hosts')
7     hosts:create_index('primary', { parts = {1, 'STR'} })
8 end
9
10 local function handler(self)
11     local host = self.req.peer.host
12     local response = {
13         host = host;
14         counter = hosts:inc(host);
15     }
16     self:render{ json = response }
17 end
18
19 httpd = require('http.server')
20 server = httpd.new('127.0.0.1', 8080)
21 server:route({ path = '/' }, handler)
22 server:start()
```

community (1)

- free as in freedom
 - BSD license
 - no closed source, no paid add-ons
 - <http://github.com/tarantool>

... but there is no such thing as free lunch: the team's employed by Mail.Ru



community (2)

- hundreds of servers within Mail.Ru:
 - replaces memcache, redis, queueing**= smart stateful first tier of the cloud**
- Badoo
 - session store
- Sberbank Digital Ventures
 - fraud detection
- growing number of deployments in online games, RTB, etc

community (3)

- #tarantool at irc.freenode.net
- tarantool@googlegroups.com
- try it now at <http://try.tarantool.org>

Questions?

