

Budapest DevOps Meetup

November, 18th 2013
Andrey Smirnov (@smira)

smotri.com

skype™




qik
See what happens



Microsoft

Who Am I

DevOps is
Communication

Dev and Ops in separate
orgs  finger-pointing

No Ops is better than
Dev vs. Ops

Communication is better than process

Qik: 08.2010

- No Chef, Puppet
- Everything is set up manually
- Tons of inconsistencies
- Incomplete monitoring (Nagios @ 100% CPU)
- What's running on that box???
- Shared Google Spreadsheet with root passwords

6 months later

- Majority of servers reinstalled from scratch
- Configuration managed by Chef
- Monitoring in Zabbix
- A lot of **pain** and **fun**

Qik at peak

HTTP Requests: **1.3M rps**

Storage requests: **240M rps**

Notifications concurrent
connections: **3M**

Presence updates: **1.8M rps**

Contact lists
unique identities: **600M**

Hardware: **80 servers**

Monitoring items: **30K**

Storages DSL

```
redis 'Dude2ID' do
  base_port 6000
  total_partitions 256
  instance 'cache9.qik.com' do
    base_name :dude2id0
    partitions 0..63
  end
  instance 'cache10.qik.com' do
    base_name :dude2id1
    partitions 64..127
  end
  instance 'cache11.qik.com' do
    base_name :dude2id2
    partitions 128..191
  end
  instance 'cache12.qik.com' do
    base_name :dude2id3
    partitions 192..255
  end
end
```

```
memcachedb 'AuthStorage' do
  base_port 21201
  base_replication_port 31201
  cache_size 3072
  total_partitions 3

  master 'db-qe-master01.qik.com' do
    partition 0
    name :auth0
  end
  slave 'db-qe-slave01.qik.com' do
    partition 0
    name :auth0_s
  end

  master 'db-qe-master01.qik.com' do
    partition 1
    name :auth1
  end
  slave 'db-qe-slave01.qik.com' do
    partition 1
    name :auth1_s
  end
end
```

Storages

- Don't Repeat Yourself:
 - Setup service
 - Setup monitoring
 - Register in application configs

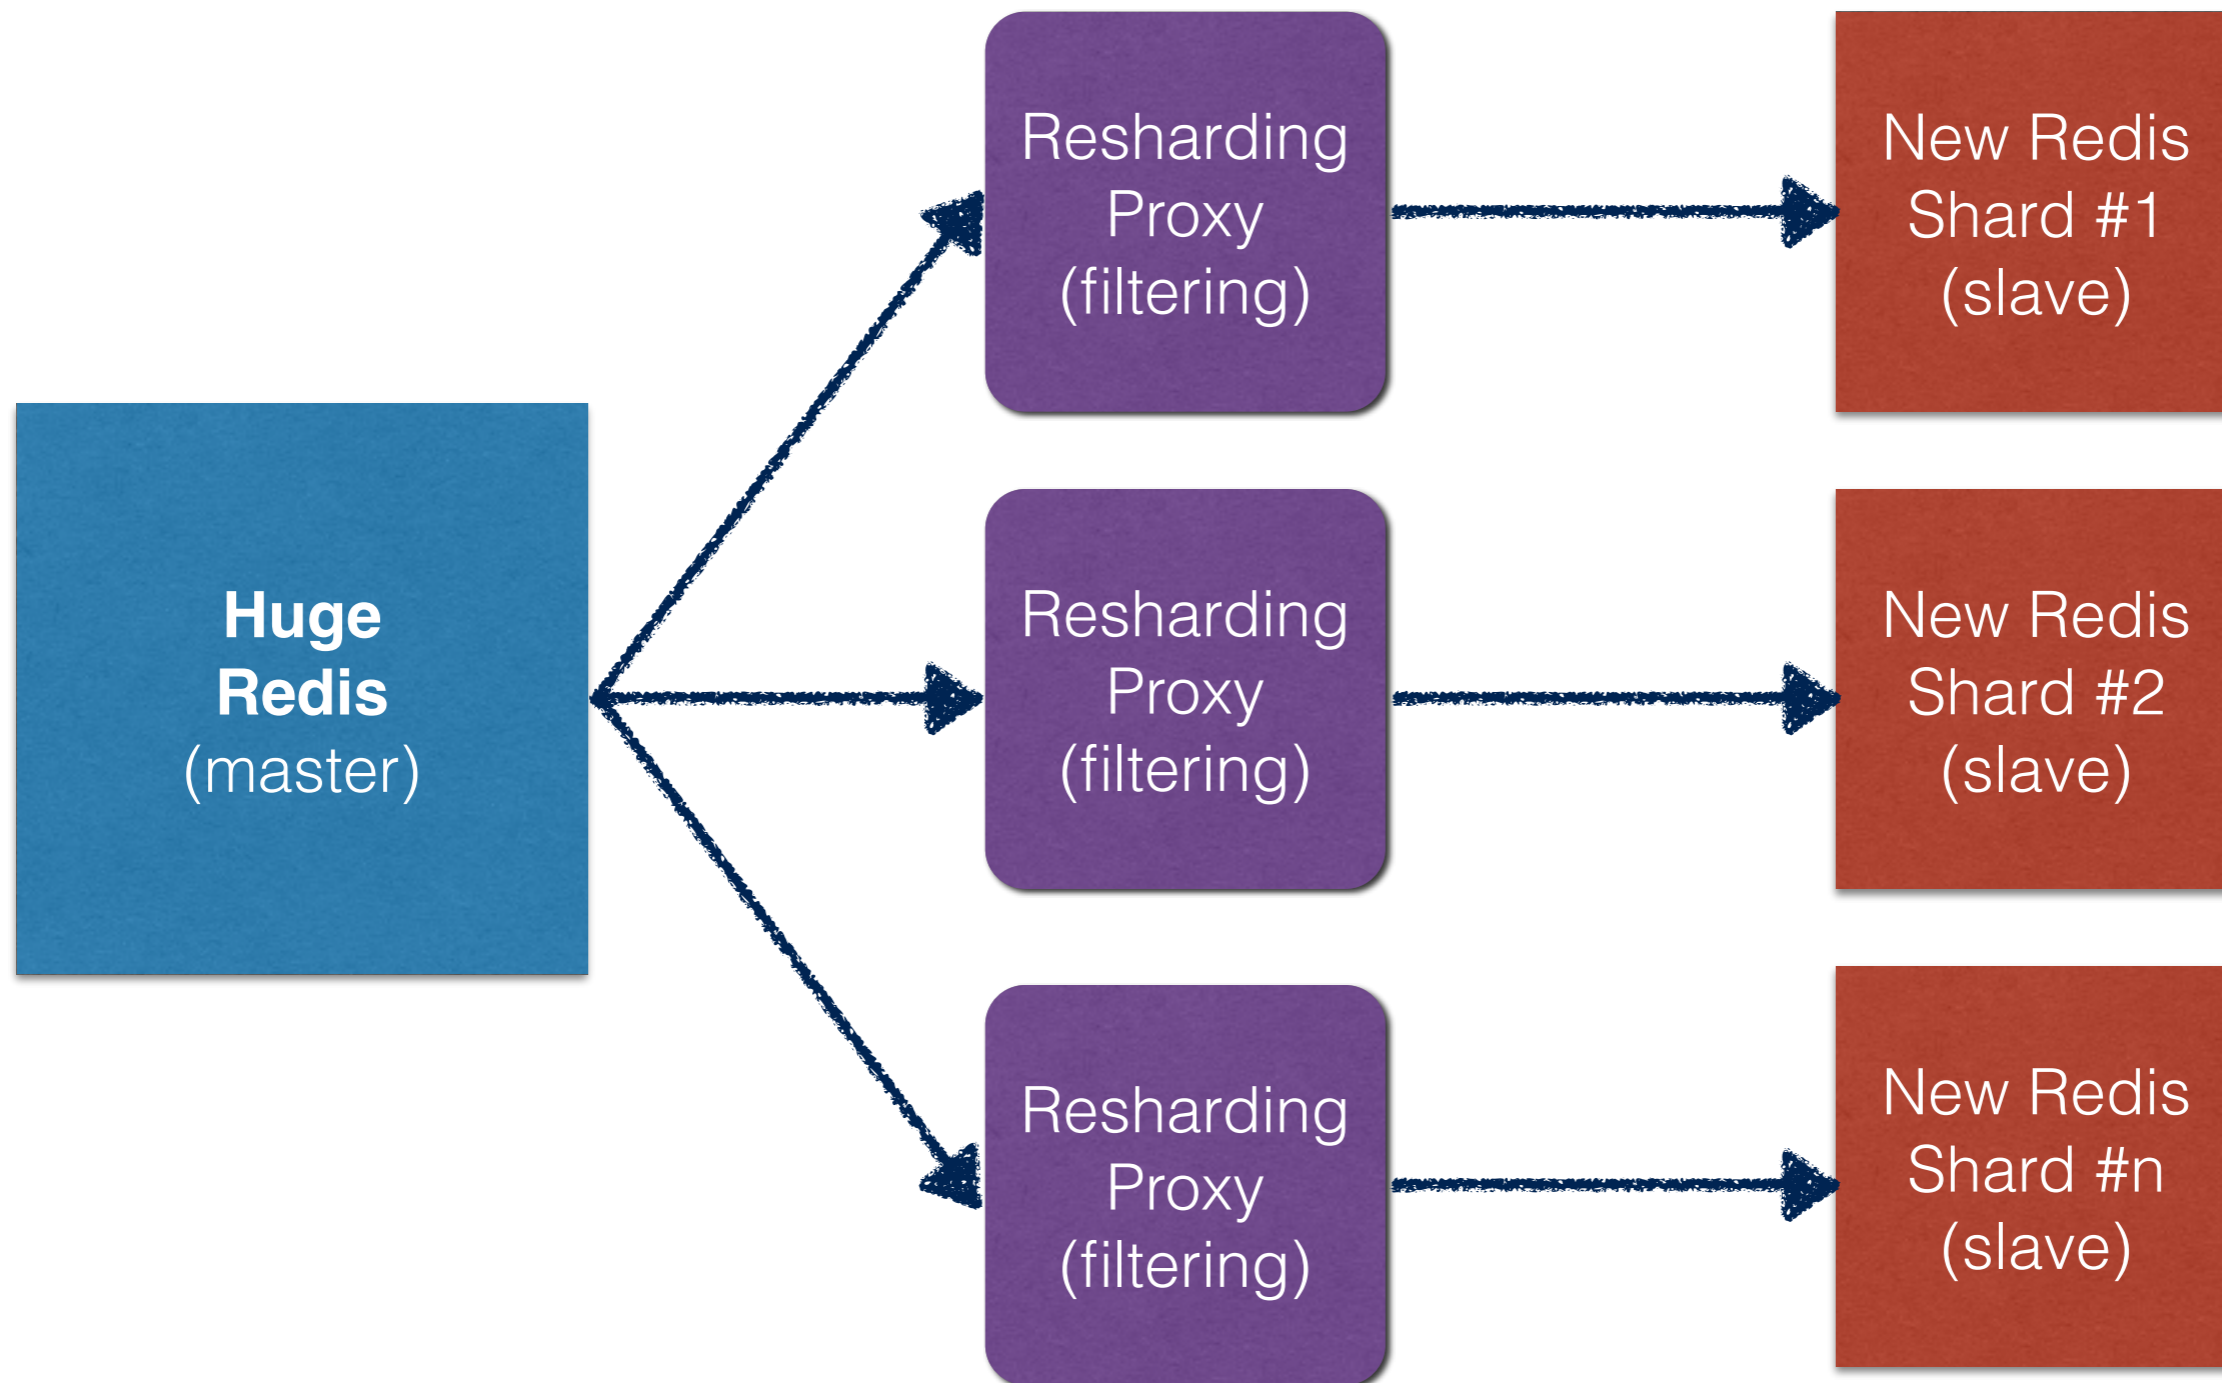
Grizmo

- Standard checks require launching process for each metric
- Asynchronous, fast checks
- Sets up Zabbix automatically based on description

Qik Logs

- Kicked off as a fork of Graylog2
- Accepting logs over UDP using syslog protocol
- Main storage in time-partitioned PostgreSQL tables
- Redis for exception aggregation (a-la Airbrake)

Redis Resharding



Redis Resharding

- RDB filtering + command filtering by key
- <https://github.com/smira/redis-resharding-proxy>
Expect to finish this week

What I'm busy with

- Chef is set of LEGO™ blocks, everybody builds his own city
- Health, monitoring and configuration are highly interconnected
- Integration of components and handling of transitions is open problem

Köszönöm!

me@smira.ru

@smira

<https://github.com/smira>

Andrey Smirnov