

12th CENTRAL & EASTERN EUROPEAN
SOFTWARE ENGINEERING CONFERENCE IN RUSSIA

October 28 - 29, Moscow



Using Hadoop stack to build a cloud VAT declarations revising service

Alex Chistyakov

ITMO University, ISST Lab

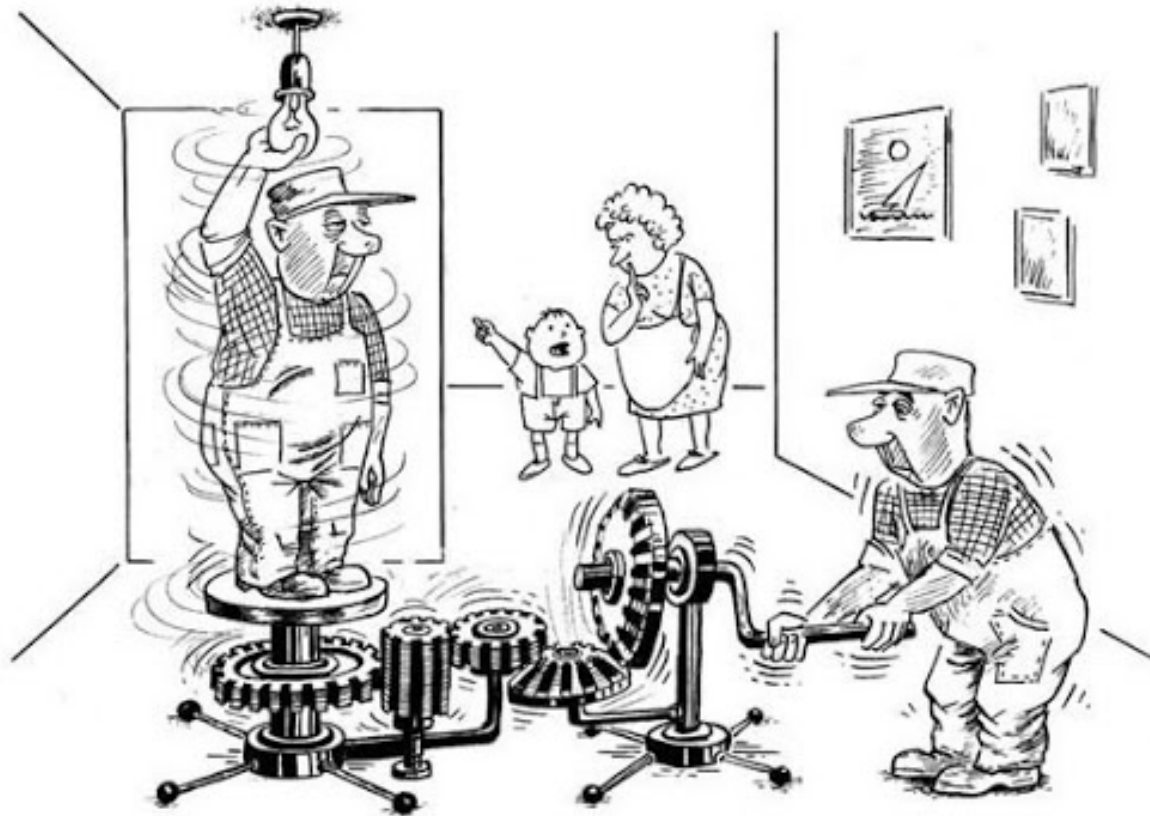
Hello, my name is Alex

Currently a researcher at ISST Lab (ITMO University, SPb., Russia)

This is a slightly better version of me =>



How many engineers do we need?



Science is tough

Well, at least there are students to help us researchers to do math



VAT revising service is...

Yet another cloud SaaS



Architecture from 50000ft

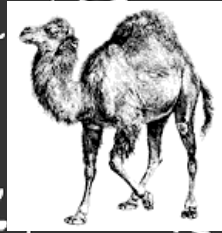


Architecture explained

Solr



APACHE
HBASE



CPAN

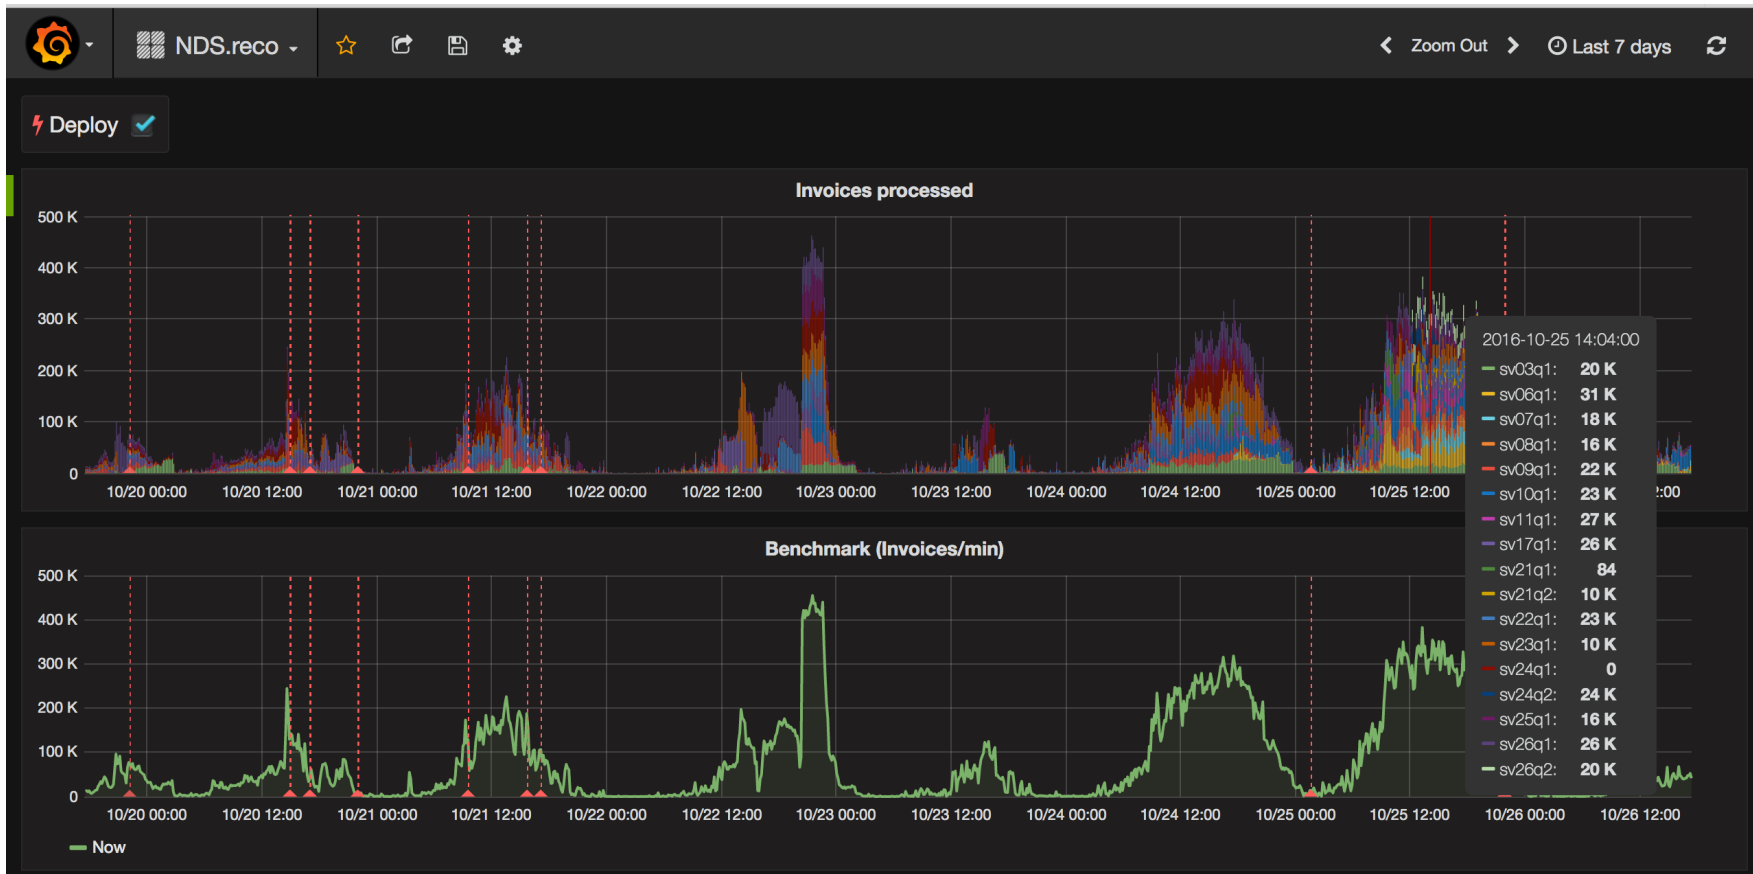


Grafana.net

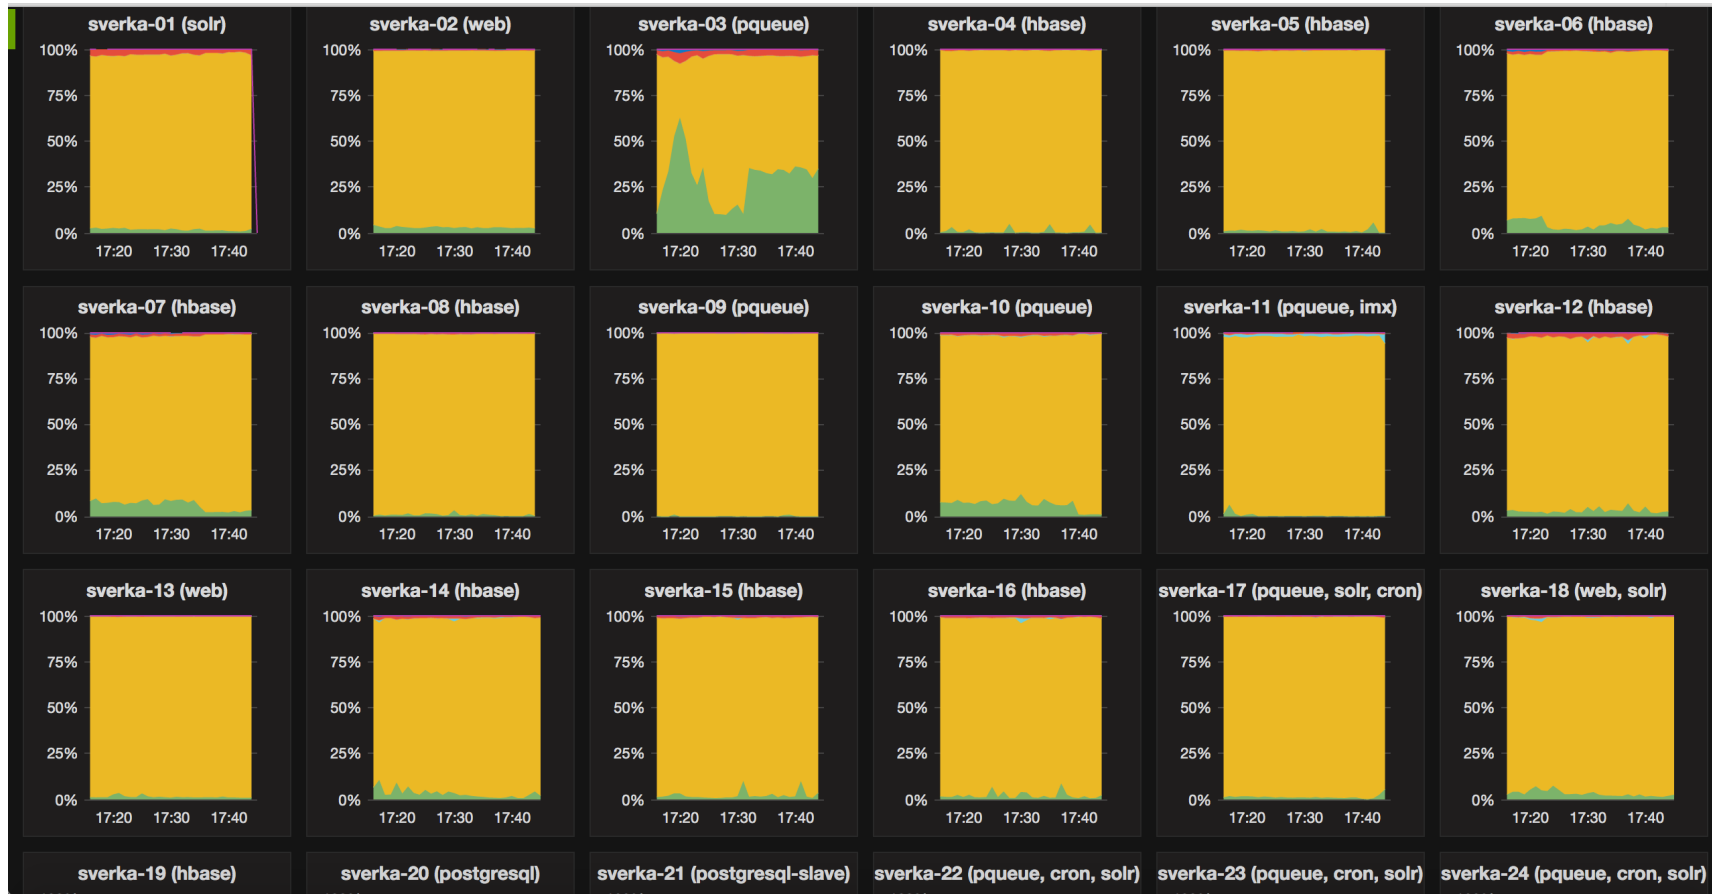


NGINX

The Church of Metrics



Everything is counted!



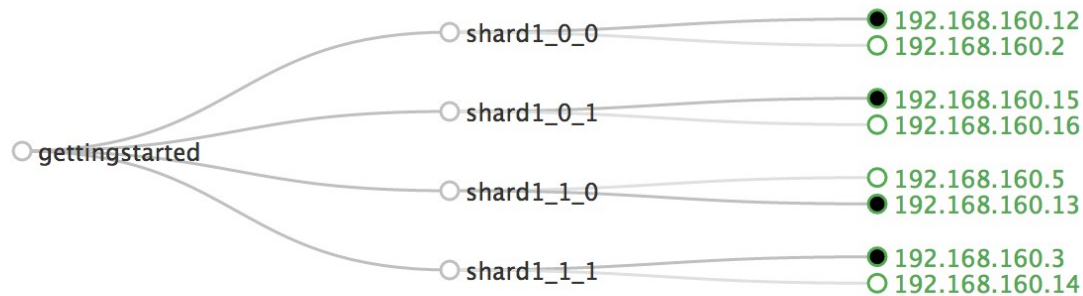
Party time!



Challenges and opportunities



Cloud and proud



- Leader
- Active
- Recovering
- Down
- Recovery Failed
- Gone

Questions

- Perl is great, after all (well, at least it was)
- A custom queue in PostgreSQL is not our biggest problem, believe me
- Maybe, you have a question too

Conclusions

- We must not fear the Big Data
- Big Data is a process
- Know your tools

