

Wikimedia TechOps quarterly review

August 2014

Looking back... 3-6 months

- Staffing: new hires
- Off-site meeting in Athens
- Process
- Projects
- Performance & availability metrics, incidents
- Goals

Staffing

- Backfilled 3 Operations Engineers
 - Chase Pettet, Giuseppe Lavagetto and Filippo Giunchedi
- Contracted a Dallas data center engineer
 - Papaul Tshibamba

- This quarter we're missing Faidon :-)

Staffing

We're hiring:

- Technical Project Manager
 - JD will be posted tomorrow
- Operations Security Engineer
 - Suitable candidates have proven hard to find
 - We may need to adjust our requirements & expectations
 - Train someone internally?
- Yuvi Panda joining TechOps (Labs)
 - Coming from Mobile Apps
 - ~October

Off-site meeting in Athens

- Team building! (that means drinks, right)
- Working together, face to face
- Meeting our new team members
- Hacking together (monitoring)
- Firefighting together (Heartbleed)
- Some planning, process changes
- **LOTS** of food



Process: Liaisons

- Zero & Mobile: Brandon
- MediaWiki Core: Faidon, later Giuseppe
- Analytics: Andrew O
- Fundraising: Jeff
- Release Eng/QA: Andrew B
- Services: Faidon

Process: SoS, Phabricator

- Faidon and Andrew O have been attending the weekly Scrum of Scrums meetings
- Greatly improving awareness of cross team blockers & needs
- Plans to migrate to Phabricator (from RT) as much as possible

Projects - Operations driven

- ✓ Puppet 3 upgrade, (Hiera)
- ✓ Labs migration to eqiad (Ashburn)
- ✓ pmtpa (Tampa) reduction
- ✓ Data center RFP completion (CyrusOne)
- ✓ codfw (Dallas) physical buildout
- ✓ ulsfo (San Francisco) floor move
- ✓ Admin accounts Puppet revamp
 - Data retention policy implementation

Projects - Operations driven

- OpenStreetMap databases
- MariaDB 10 upgrade
- ✓ HTTPS PFS
- X Virtualization for miscellaneous services
- X Monitoring revamp
- X HTTPS & IPsec

Projects - externally driven

- HHVM, application servers migration (HAT)
- Phabricator
- Hadoop infrastructure (Refinery)
- ✓ Toolserver migration to (Tool)Labs
 - ElasticSearch support
 - PDF rendering
 - ContentTranslation Handler
 - Mathoid
- ✓ Kafka for Varnish logging

Ongoing Operational responsibilities

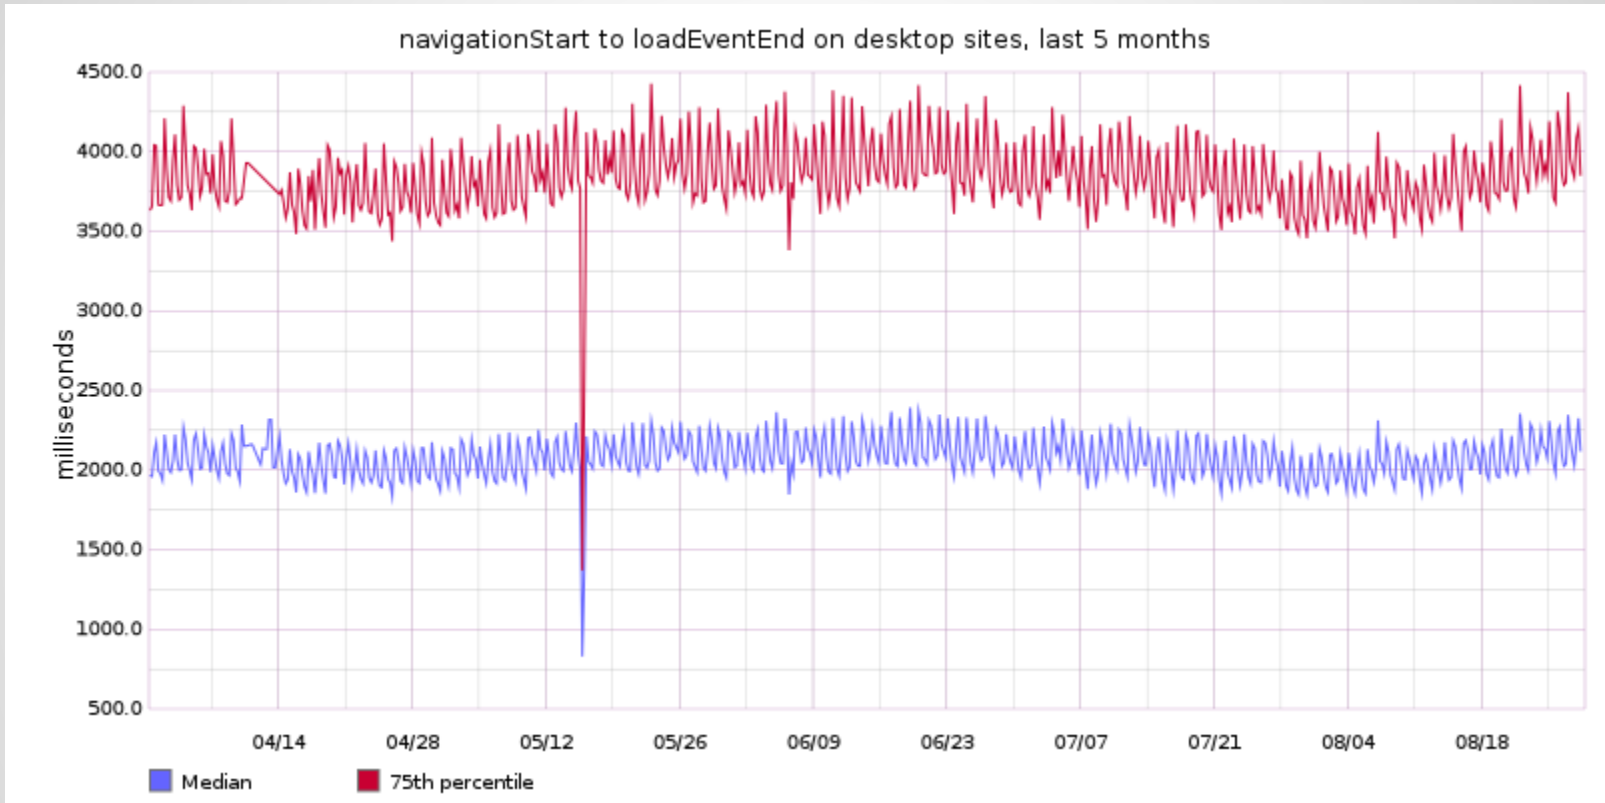
- server & service maintenance
- security upgrades
- incident response
- request tickets
- equipment procurement
- data center management
- capacity management
- Gerrit reviews
- service contract negotiations
- network migrations
- Puppet rework
- Labs user support
- access requests
- saying No :(

Project: Tampa/Dallas dc migration

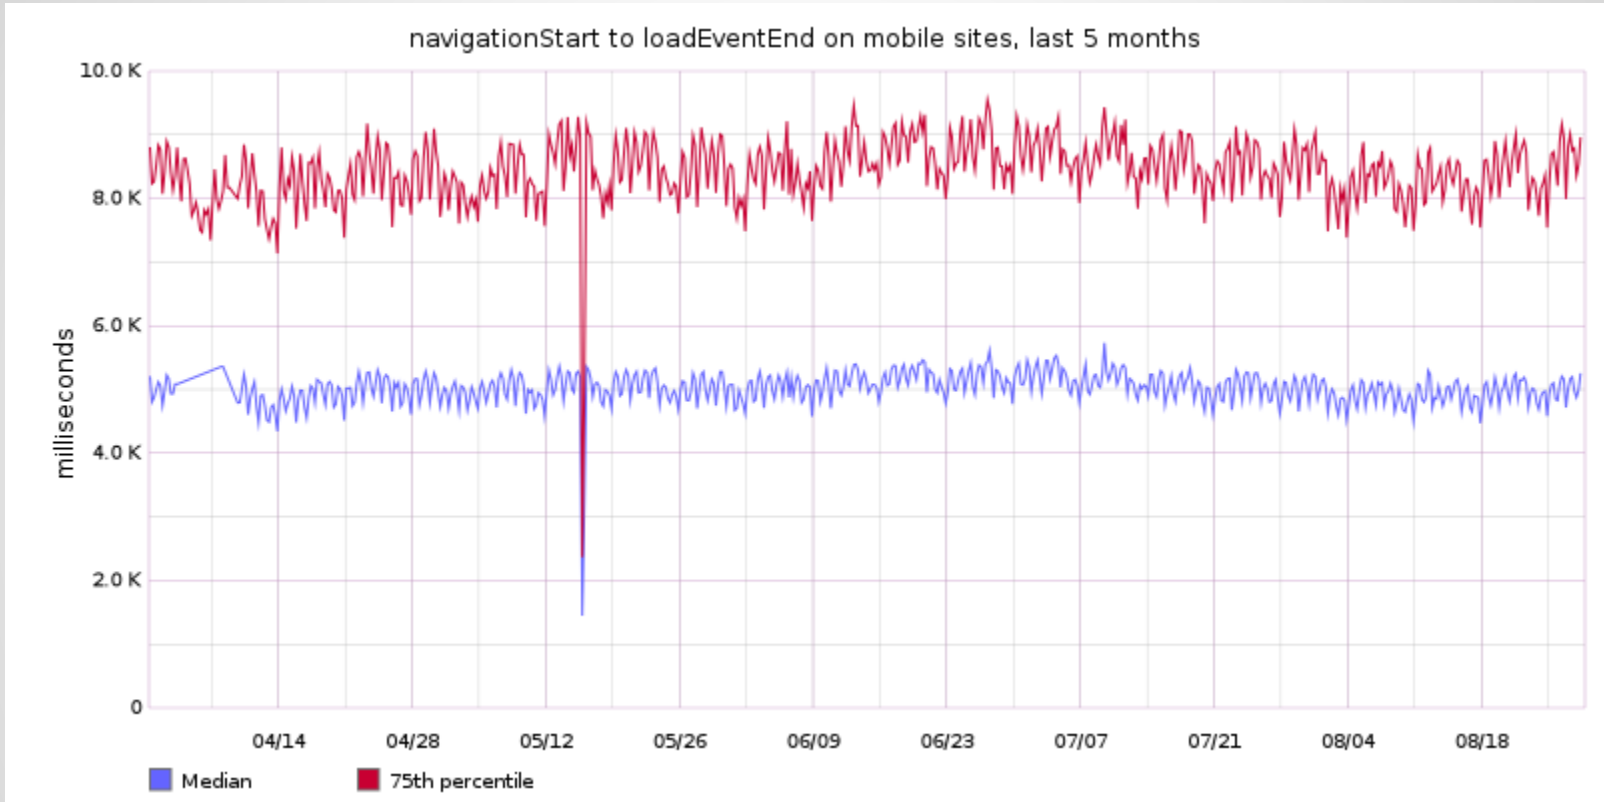
- First connectivity available since last week
- First server installs this week
 - Essential infrastructure coming online now
- Data migration starting next week
 - Database replication, Swift images, backups, etc...
- Plan to turn off all lights in Tampa end of September

Performance & uptime metrics

April 2014 - now



Navigation Timing desktop sites, <http://gdash.wikimedia.org>

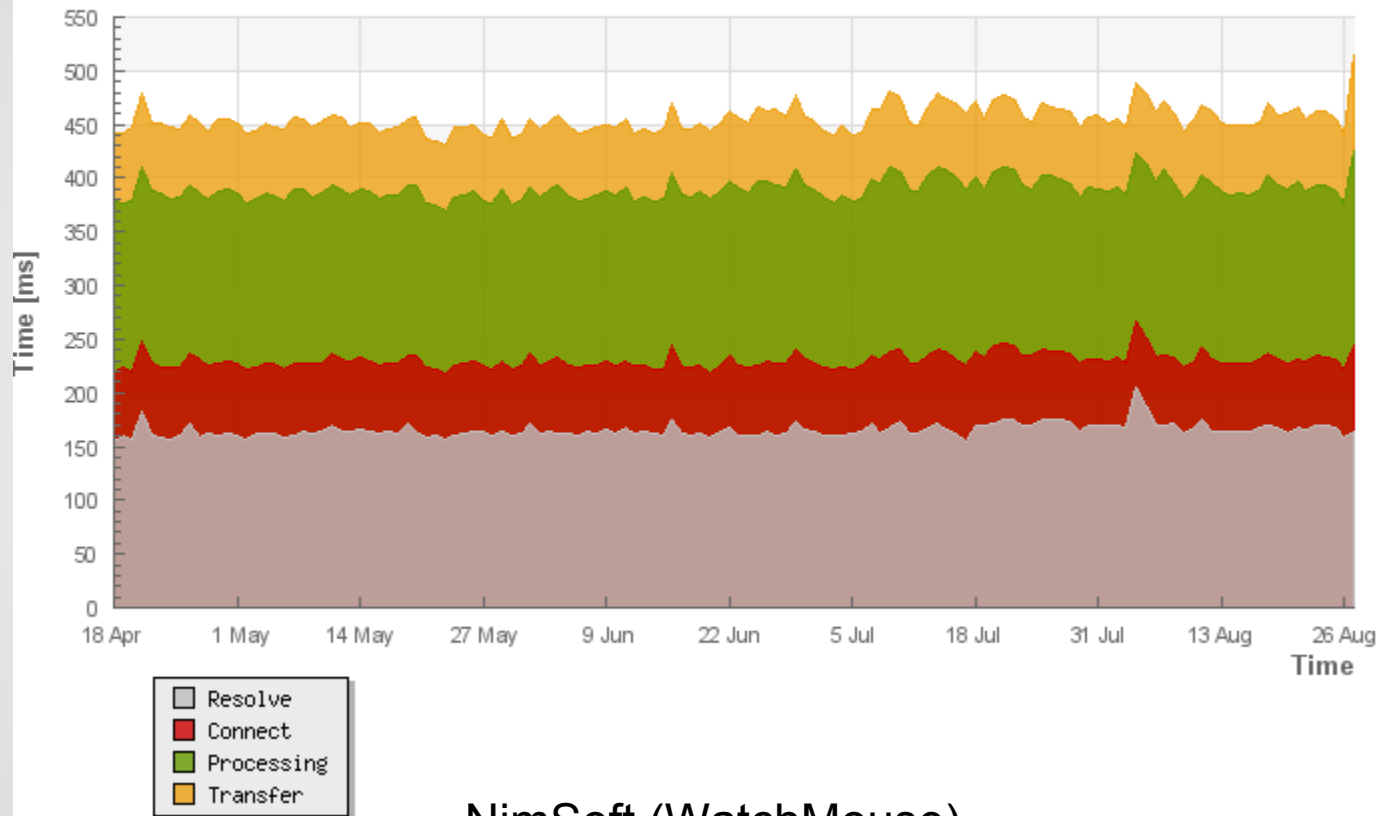


Navigation Timing mobile sites, <http://gdash.wikimedia.org>

Performance

1 April to 28 August 2014

Monitors: DNS, GeolP lookup, https services - commons, and 32 more...




NimSoft (WatchMouse)

Performance

1 April to 28 August 2014

Monitors: Wiki platform [[w:en:Main Page]] (s1)




NimSoft (WatchMouse)

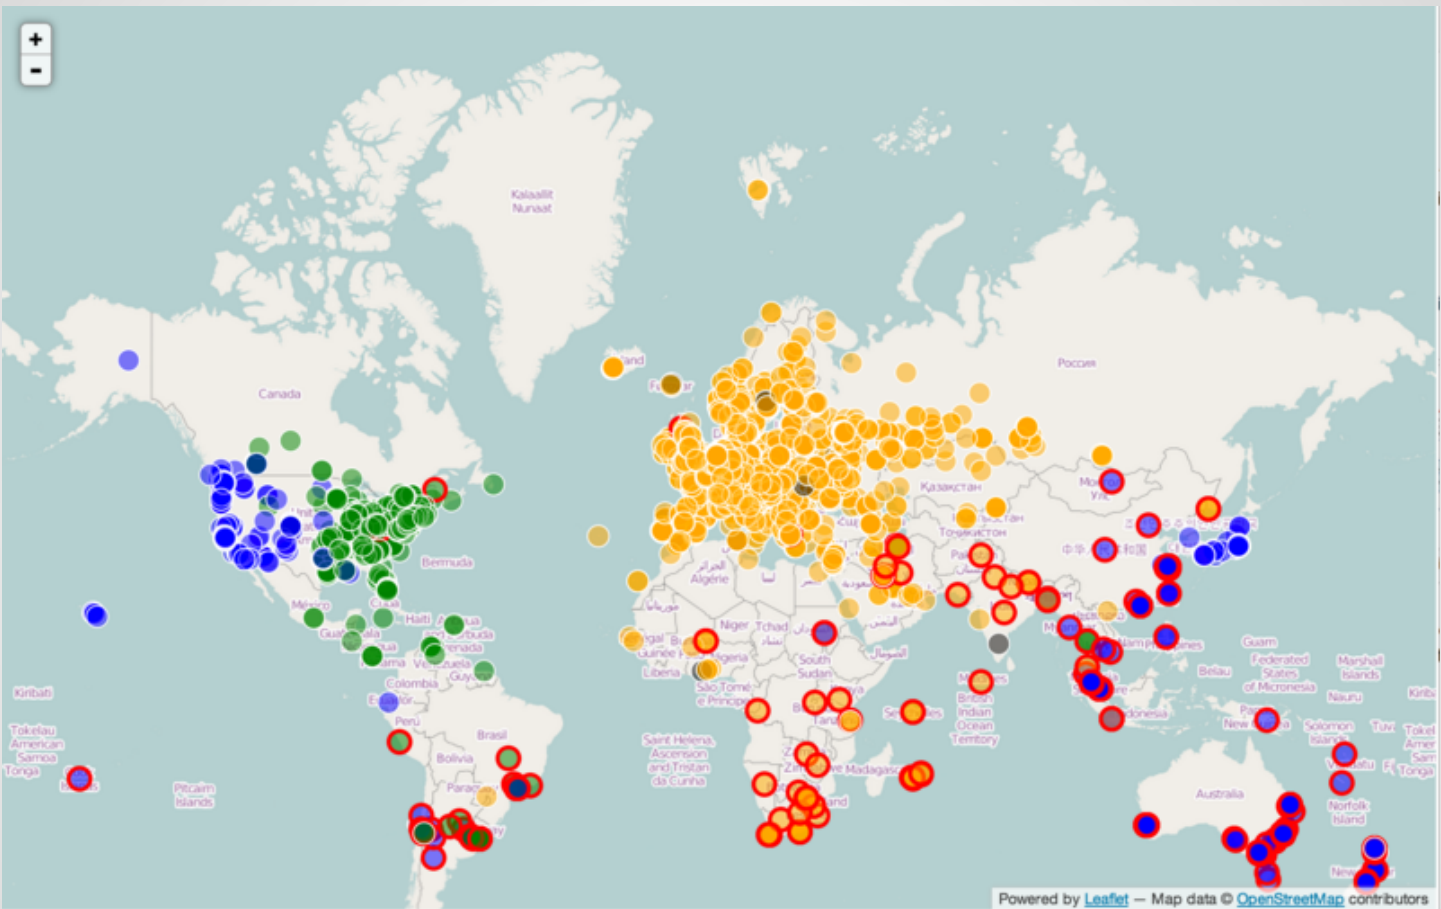
Performance

1 April to 28 August 2014

Monitors: Wiki platform [[w:en:Main Page]] (s1) - UNCACHED



NimSoft (WatchMouse)



RIPE latency data to our data centers

<http://blog.wikimedia.org/2014/07/09/how-ripe-atlas-helped-wikipedia-users/>

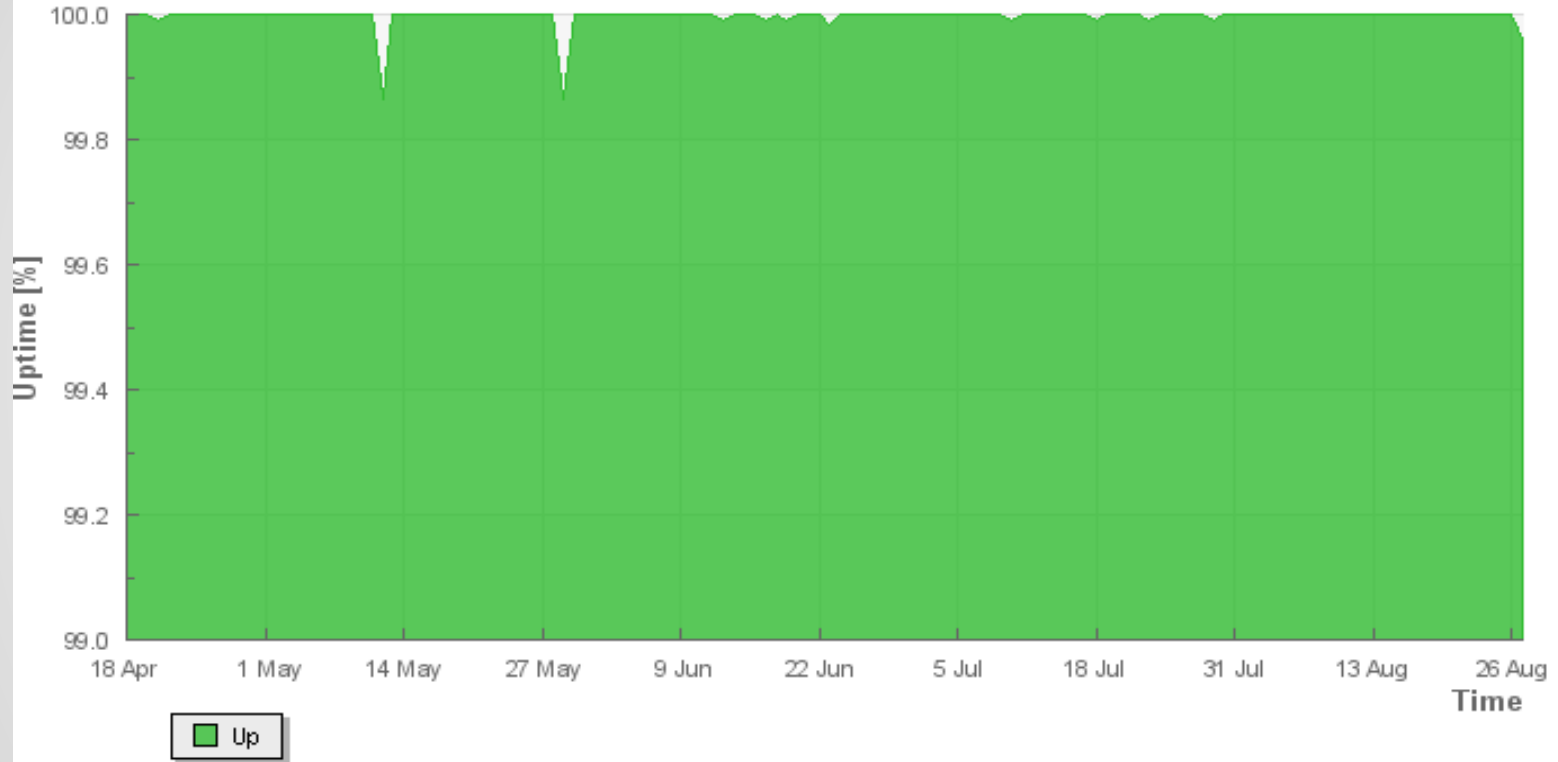
Incidents & uptime

- From March 2014 - June 2014:
 - 17 incident reports
 - Discussed in a quarterly incident review meeting
- From July 2014 - now
 - 7 incident reports

Uptime

1 April to 28 August 2014

Monitors: DNS, GeolP lookup, https services - commons, and 32 more...



NimSoft (WatchMouse)

Core site services aggregated (wikis & supportive infrastructure)

This quarter: Site Operations goals

- X Virtualization for miscellaneous servers at codfw
 - Core data backups/replication at codfw
- X IPsec on cross-data centre links for private data
 - Migrate application servers to Ubuntu 14.04 Trusty, and assist with HHVM

This quarter: Labs goals

- ✓ Reorganization/improvement of new-user & project documentation
 - Storage capacity & redundancy expansion
 - Process for user backups

Next quarter: Site Operations goals

- Core site services at codfw
- Test deployment of HTTP 2.0 / SPDY
- SSL with PFS (“done”)
 - We’ll work on IPsec and SSL infrastructure
- Varnish 4 in limited production
- codfw fundraising infrastructure
 - fully functional in time for the fundraiser

Next quarter: Labs goals

- Labs available in the new data centre (with Neutron/IPv6)
- Distributing tools, deployment-prep to both data centers (availability/redundancy)

Project: new data center (codfw)

- Phase 1
 - Essential infrastructure (now)
 - All data replicated & backed up (~end of Q1)
- Phase 2 (~Q2)
 - Ability to serve wikis from codfw
 - Labs availability zone
 - Fundraising replica
- Phase 3 (Q3, Q4)
 - Migration/duplication of remaining (misc) services

Involvement of the majority of the team

Q2 dependency: Platform

Project: metrics

- Clear need for better performance metrics
 - User perceived performance
 - Data driven decisions, A/B testing
- Metrics need to be consistently well defined across the infrastructure
- Requires robust & scalable metric infrastructure
 - Scalability work on Graphite has already started
 - Tie-in with Analytics infrastructure?
- Plus better third party monitoring
 - RIPE Anchors, NimSoft replacement

Dependencies: Platform, Analytics

Project: monitoring

- Graphite scaling
 - Graphite in Labs
 - Yuvi Panda working on this
- Incident monitoring
 - Our Icinga (Nagios) install is legacy
 - Configuration management not up to standards
 - Not a scalable setup
 - Research to possibly replace with Shinken (compatible)
 - Yuvi also looking into this (for Labs)