

# TECHNOLOGY DEPT

## Foundational & Sustaining Programs

February 2018 quarterly check-in  
for work done in Q2 FY2017/18



**WIKIMEDIA**  
FOUNDATION

# Program Structure

<b>Sustaining</b>	TP1 Availability, Performance & Maintenance	TP3 Addressing Technical Debt TP8 Multi-datacenter Support
<b>Foundational</b>	TP2 Mediawiki Refresh TP6 Streamlined Service Delivery PP1 Discoverability	X-SPDM Security, Privacy, & Data mgmt X-SDC Structured Data on Commons
<b>Community Support</b>	TP5 Scoring Platform (ORES) TP9 Growing Wikipedia Across Languages [...] TP7 Smart Tools for Better Data	TP11 Citations/Verifiability TP12 Growing Contributor Diversity X-CH Community Health/Anti-harassment
<b>Tech Community Support</b>	TP4 Technical Community Building	TP10 Public Cloud Services & Support

# Program Priorities

	If we don't do this...	Actual FTE (approximate)
<b>Sustaining</b>	The sites go down.	30
<b>Foundational</b>	Performance and data quality decays.	10
<b>Community Support</b>	Become technologically obsolete.	10
<b>Tech Community Support</b>	Lose bots and code contributions.	4

# How we prioritize



## Fundamentals

What is our part in fulfilling the mission?



## Service

What are other people asking from us?



## Improvement

What could we do to improve our offering?



## Maintenance

What will sustain and improve our delivery?

# Agenda:

## **Technology**

Program 1: Availability, performance and maintenance

Program 2: MediaWiki

Program 3: Addressing technical debt

Program 6: Streamlined service delivery

Program 8: Multi-datacenter support

## **Product / Audiences & Cross-Departmental (tech portions only)**

Program 1: Make knowledge more easily discoverable

## **Programs covered in other presentations:**

Program 7: Payment processor investigation and long-term strategy

Program 8: Donor retention

Fundraising Tech / Advancement

Privacy, Security and Data Management

Structured Data on Commons (Database & Search integration)

# Technology Program 1

## Availability, Performance and Maintenance

**Program Structure**

**Sustaining**

TP1 Availability, Performance & Maintenance

We will **maintain** the availability of Wikimedia's sites and services for our global audiences, and ensure they're running **reliably**, **securely** and **performant** while *modernizing* our infrastructure and improving current levels of service when it comes to testing, deployments and maintenance of software and hardware



Lewis Hine Power house mechanic working on steam pump, Public Domain



Feb 2018

**TechOps / Traffic**1.1.3

## New cache edge site *EQSIN*, to better serve Asia:

- Completed purchasing of space and hardware
- Acquired IP address space from APNIC
- Installed all server & network hardware
- Network links and user service on track for Q3





CHECK IN

TEAM/DEPT

PROGRAM

Feb 2018

**TechOps / Core Ops**

1.1.1



Punch cards, Public Domain

Wikidata now has its own dedicated *database section* with dedicated resources: `s8`

After extensive preparation work throughout Q2, the migration was performed seamlessly on January 9th.

# Keeping up...

All Puppet agents in production and Cloud Services have been upgraded to Puppet 4, *enabling new syntax and features*

- The upgrade of PuppetDB (backend) was postponed to Q3

The “*misc*” caching cluster was upgraded to Varnish 5

- all other clusters will follow in Q3



Feb 2018

**Release Engineering**1.5

## Migrate majority of developers to Node.js framework:

We are changing our canonical browser testing framework to being Node.js based (from Ruby).

This past quarter we officially deprecated the Ruby framework and [only 46% of our repositories](#) (“projects”) depend on it at this point (with 7% in progress).



# Technology Program 2

## MediaWiki

**Program Structure**

**Foundational**

TP2 Mediawiki Refresh  
TP6 Streamlined Service Delivery  
PP1 Discoverability

We will strive for a  
refreshed and  
highly performant  
core platform by  
bringing renewed  
focus on  
MediaWiki







**Outcome 1:** Stakeholders in MediaWiki development will **have a sense of progress and direction** in MediaWiki.

**Objective 1:** **develop** a MediaWiki **roadmap**

**Milestone 1:** Hire a product manager for MediaWiki by 2017-08-31 ✓

**Outcome 2:** MediaWiki **code quality** will be improved.

**Objective 1:** Increase measured unit test coverage

**Objective 2:** Break up large classes and source files

**Outcome 3:** MediaWiki **security and stability** will be improved.

**Objective 1:** Address the backlog of action items that arise from security and downtime post-mortems





# How can we reduce complexity?

2 parsers

2 HTML  
storage  
systems

3 UIs

2 APIs

Co-leading the cross-departmental  
Audiences Technology Working Group

Developing a cross-departmental  
program:

- annual plan goals
- 3-5 year plan

**Status: Planning in progress**

# Namespacization

	A	B	C
815	RevDelRevisionItem	MediaWiki\RevisionDelete\RevDelRevisionItem	
816	RevDelRevisionList	MediaWiki\RevisionDelete\RevDelRevisionList	
817	RevisionDeleteUser	MediaWiki\RevisionDelete\RevisionDeleteUser	
818	RevisionDeleter	MediaWiki\RevisionDelete\RevisionDeleter	
819	RevisionItem	MediaWiki\RevisionItem	
820	RevisionItemBase	MediaWiki\RevisionItemBase	
821	RevisionListBase	MediaWiki\RevisionListBase	
822	RunningStat	MediaWiki\RunningStat	
823	ScopedCallback	MediaWiki\ScopedCallback	
824	AugmentPageProps	MediaWiki\Search\AugmentPageProps	
825	DummySearchIndexFieldDefinition	MediaWiki\Search\DummySearchIndexFieldDefinition	
826	NullIndexField	MediaWiki\Search\NullIndexField	
827	PerRowAugmentor	MediaWiki\Search\PerRowAugmentor	
828	ResultAugmentor	MediaWiki\Search\ResultAugmentor	
829	ResultSetAugmentor	MediaWiki\Search\ResultSetAugmentor	
830	SearchDatabase	MediaWiki\Search\SearchDatabase	
831	SearchEngine	MediaWiki\Search\SearchEngine	
832	SearchEngineConfig	MediaWiki\Search\SearchEngineConfig	
833	SearchEngineDummy	MediaWiki\Search\SearchEngineDummy	
834	SearchEngineFactory	MediaWiki\Search\SearchEngineFactory	
835	SearchExactMatchRescorer	MediaWiki\Search\SearchExactMatchRescorer	
836	SearchHighlighter	MediaWiki\Search\SearchHighlighter	
837	SearchIndexField	MediaWiki\Search\SearchIndexField	
838	SearchIndexFieldDefinition	MediaWiki\Search\SearchIndexFieldDefinition	
839	SearchMssql	MediaWiki\Search\SearchMssql	
840	SearchMySQL	MediaWiki\Search\SearchMySQL	
841	SearchNearMatchResultSet	MediaWiki\Search\SearchNearMatchResultSet	
842	SearchNearMatcher	MediaWiki\Search\SearchNearMatcher	
843	SearchOracle	MediaWiki\Search\SearchOracle	
844	SearchPostgres	MediaWiki\Search\SearchPostgres	
845	SearchResult	MediaWiki\Search\SearchResult	
846	SearchResultSet	MediaWiki\Search\SearchResultSet	
847	SearchSqlite	MediaWiki\Search\SearchSqlite	
848	SearchSuggestion	MediaWiki\Search\SearchSuggestion	
849	SearchSuggestionSet	MediaWiki\Search\SearchSuggestionSet	
850	SqlSearchResultSet	MediaWiki\Search\SqlSearchResultSet	
851	SiteConfiguration	MediaWiki\SiteConfiguration	
852	SiteStats	MediaWiki\SiteStats	
853	SiteStatsInit	MediaWiki\SiteStatsInit	
854	CachingSiteStore	MediaWiki\Site\CachingSiteStore	
855	DBSiteStore	MediaWiki\Site\DBSiteStore	
856	FileBasedSiteLookup	MediaWiki\Site\FileBasedSiteLookup	
857	HashSiteStore	MediaWiki\Site\HashSiteStore	
858	MediaWikiSite	MediaWiki\Site\MediaWikiSite	
859	Site	MediaWiki\Site\Site	

## Namespaces:

- Aid in readability
- Reduce ambiguity
- Avoid naming collisions
- Promote structuring of code

Developing a tool to rename almost all of MediaWiki core's 1,700 classes.

Status: In progress, target end of Q3

# Code Quality



ILLUSTRATED BY SEGUE TECHNOLOGIES

## Improvements to:

- continuous integration
- code coverage analysis
- extension management

## Core and extension bug fixes

## Code review

**Status: On-going**



# Firejail

(a sandbox for external  
commands run by  
MediaWiki core)

## Worked to:

- integrate Firejail into MediaWiki core
- contribute new features to Firejail upstream maintainers

**Status: Development complete,  
deployment in Q3**



# Remove use of PHP serialization in revision storage

**Addresses RFC** *“Deprecate using php serialization inside MediaWiki”*

**Status:** Code written and in review,  
target completion end of Q3

# Technology Program 3

## Addressing technical debt


**Program Structure**

**Sustaining**

TP3 Addressing Technical Debt  
TP8 Multi-datacenter Support



We will lead an organization-wide effort to reduce technical debt



Feb 2018

## Release Engineering

3 - Tech Debt

### Defined what it means to be a code steward

- pending a conversation with Toby and Victoria for wider dissemination

### Defined the “Technical Debt Project Manager” role

- regularly communicates with all Foundation engineering teams regarding their technical debt

### In collaboration with the Sunsetting Working Group

- developed a process (“Code Stewardship Reviews”) for dealing with orphaned code; the first items are going through now

### Continuing conversations

- with teams in support of finding stewards for orphaned code

# Technology Program 6

## Streamlined service delivery

**Program Structure**

**Foundational**

TP2 Mediawiki Refresh  
TP6 Streamlined Service Delivery  
PP1 Discoverability

We will **streamline** and *integrate* the delivery of services, by building a **new production platform** for integrated development, testing, deployment and hosting of applications




**TechOps** (with Release Engineering) **extended the design of our new service delivery platform** with a design for *pod-level monitoring and service-level alerting* - monitoring standard metrics of CPU & memory usage, IO operations, etc...

**Created scaffolding** (of Helm charts) to setup services for either personal development or production deployment.




## Operation Latencies

### Kubelet operation latencies eqiad




- kubernetes1001.eqiad.wmnet
- kubernetes1002.eqiad.wmnet
- kubernetes1003.eqiad.wmnet
- kubernetes1004.eqiad.wmnet

### Kubelet operation latencies codfw




- kubernetes2001.codfw.wmnet
- kubernetes2002.codfw.wmnet
- kubernetes2003.codfw.wmnet
- kubernetes2004.codfw.wmnet

### Kubelet operation latencies eqiad



- container\_status
- create\_container
- exec
- image\_status
- list containers


### Kubelet operation latencies codfw



- run\_podsandbox
- start\_container
- status
- stop\_podsandbox
- version

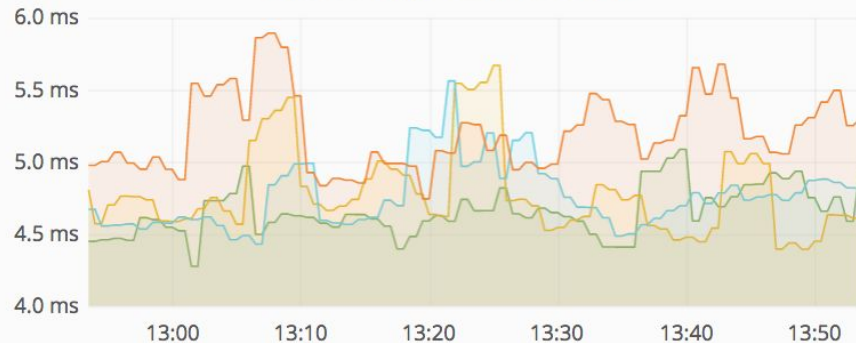
## Response times

### Kubelet response times eqiad



- kubernetes1001.eqiad.wmnet
- kubernetes1002.eqiad.wmnet
- kubernetes1003.eqiad.wmnet
- kubernetes1004.eqiad.wmnet

### Kubelet response times codfw



- kubernetes2001.codfw.wmnet
- kubernetes2002.codfw.wmnet
- kubernetes2003.codfw.wmnet
- kubernetes2004.codfw.wmnet



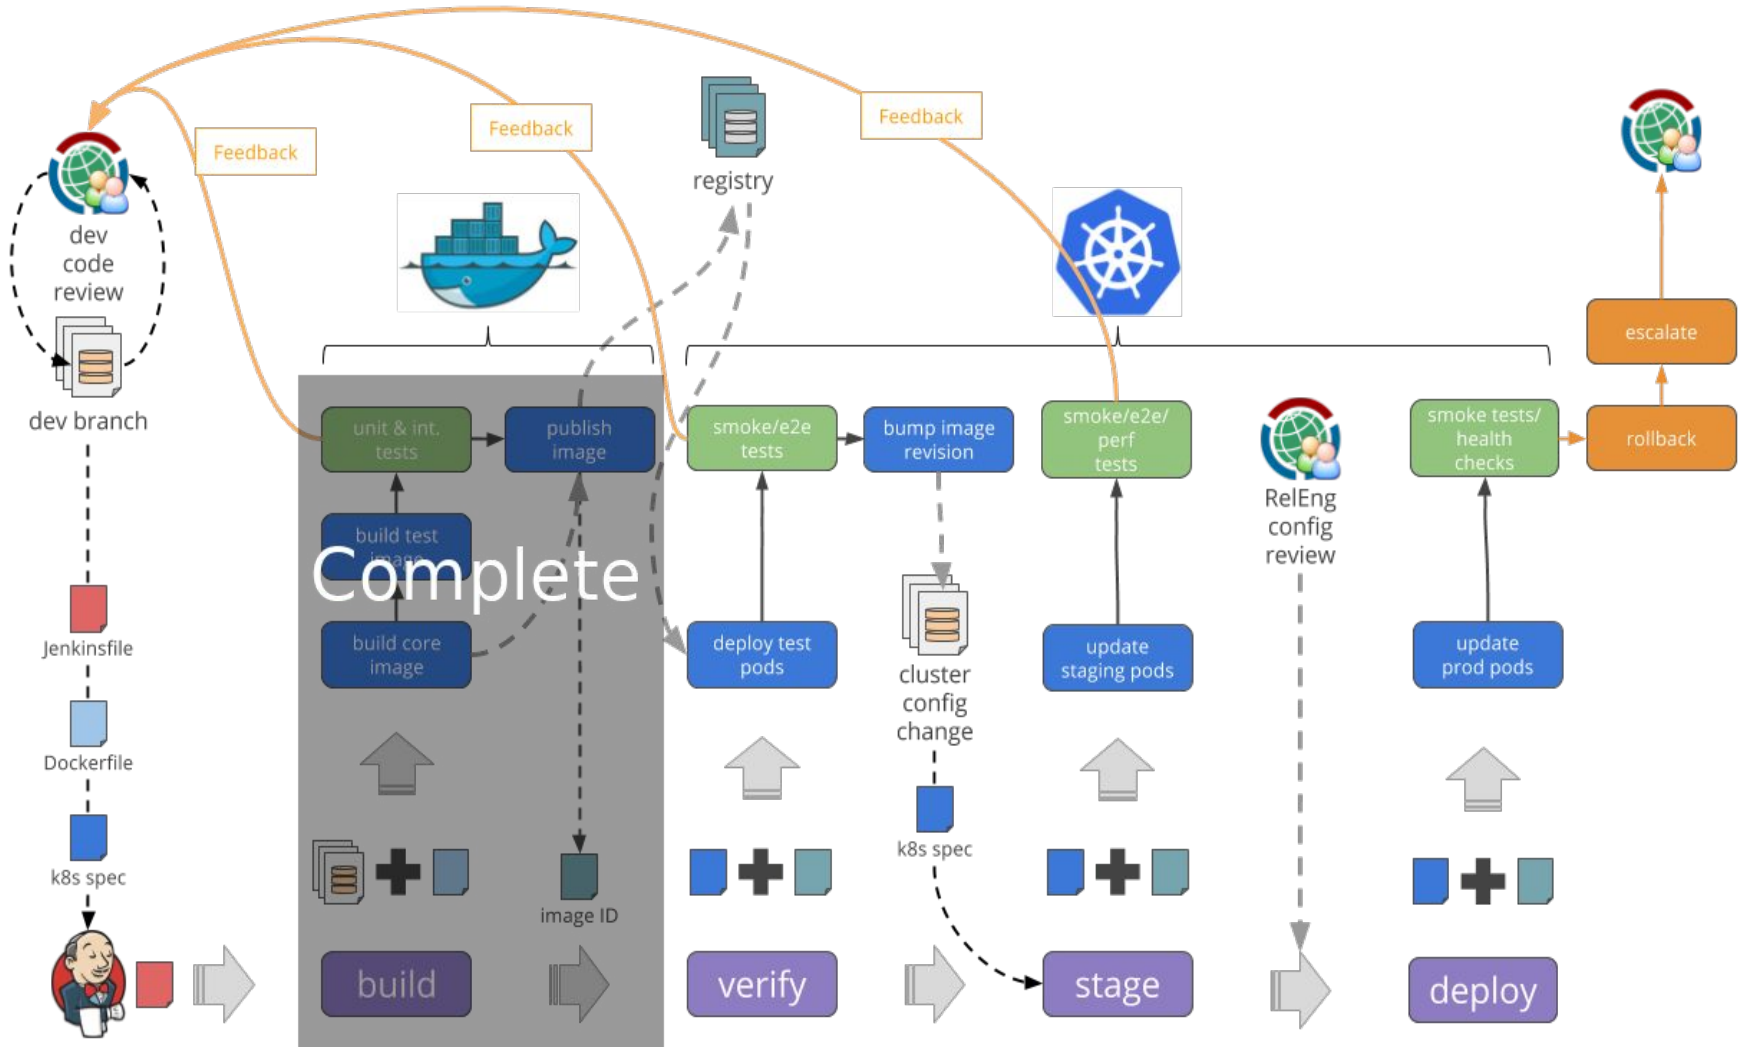
Feb 2018

# Release Engineering

6.1.2

End-to-end smoke tests are included to verify the functional integrity of production images before promotion to a staging environment.

Release Engineering also completed the push phase of the production variant of build and release pipeline.



# Technology Program 8

## Multi-datacenter support

**Program Structure**

**Sustaining**

TP3 Addressing Technical Debt  
TP8 Multi-datacenter Support

## **Program 8** Multi-datacenter support

- 1. Faster primary datacenter switchover**
- 2. Dual primary data centers**

Feb 2018

**Performance**

## mcrouter deployment

- A Facebook product which provides memcached replication
- Goal: install and use in deployment-prep (Beta Cluster)
- Puppet module deployed, mostly complete
- Initial testing done
- **Status: in progress**

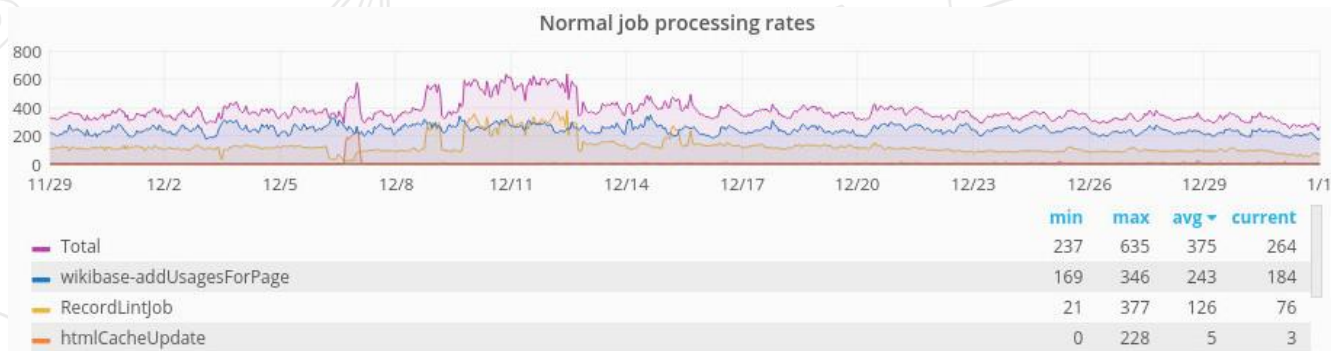
Feb 2018

# Service Platform

8

## Modernize the JobQueue transport

- Migrated 3.5 high-traffic and 1 low-traffic job to EventBus
  - Q2 avg: 400 / 1100 jobs
- Identified the needs for the migration of the remaining jobs
  - refreshLinks (deduplication), CirrusSearch jobs (delayed execution)
- Documented metrics and common operational tasks
  - [https://wikitech.wikimedia.org/wiki/Kafka\\_Job\\_Queue](https://wikitech.wikimedia.org/wiki/Kafka_Job_Queue)
  - <https://grafana.wikimedia.org/dashboard/db/jobqueue-eventbus>





# Product Program 1

Make knowledge more easily discoverable

**Program Structure**


**Foundational**

TP2 Mediawiki Refresh  
TP6 Streamlined Service Delivery  
PP1 Discoverability

Through incremental search technology improvements, we will help readers more effectively discover and search for content.



## Continued to improve **Machine Learning to Rank**



- Implemented **more automation** of the training pipeline
- Added machine learning to rank for **8 additional languages**
- **Upgraded Elasticsearch to v5.5**

Follow the Epic: [T161632](#)

# Improving Search Across Languages

- Investigated seven open source language analysis software components for feasibility as Elasticsearch plugins  
([T171652](#))
- Investigated usage of fallback languages  
([T147959](#))
- Up Next for Q3: Phonetic Matching  
([T182708](#))




Photo by Man vyi

[https://commons.wikimedia.org/wiki/File:Welcome\\_multilingual\\_Guernsey\\_tourism.jpg](https://commons.wikimedia.org/wiki/File:Welcome_multilingual_Guernsey_tourism.jpg)


Preparing for

## Structured Data on Commons

- **Category Search** via dumps for select wikis using SPARQL on WDQS
- Wikidata Query Service **Full-Text Search**
- Ready for next steps SDoC...



Questions?





# Additional work:

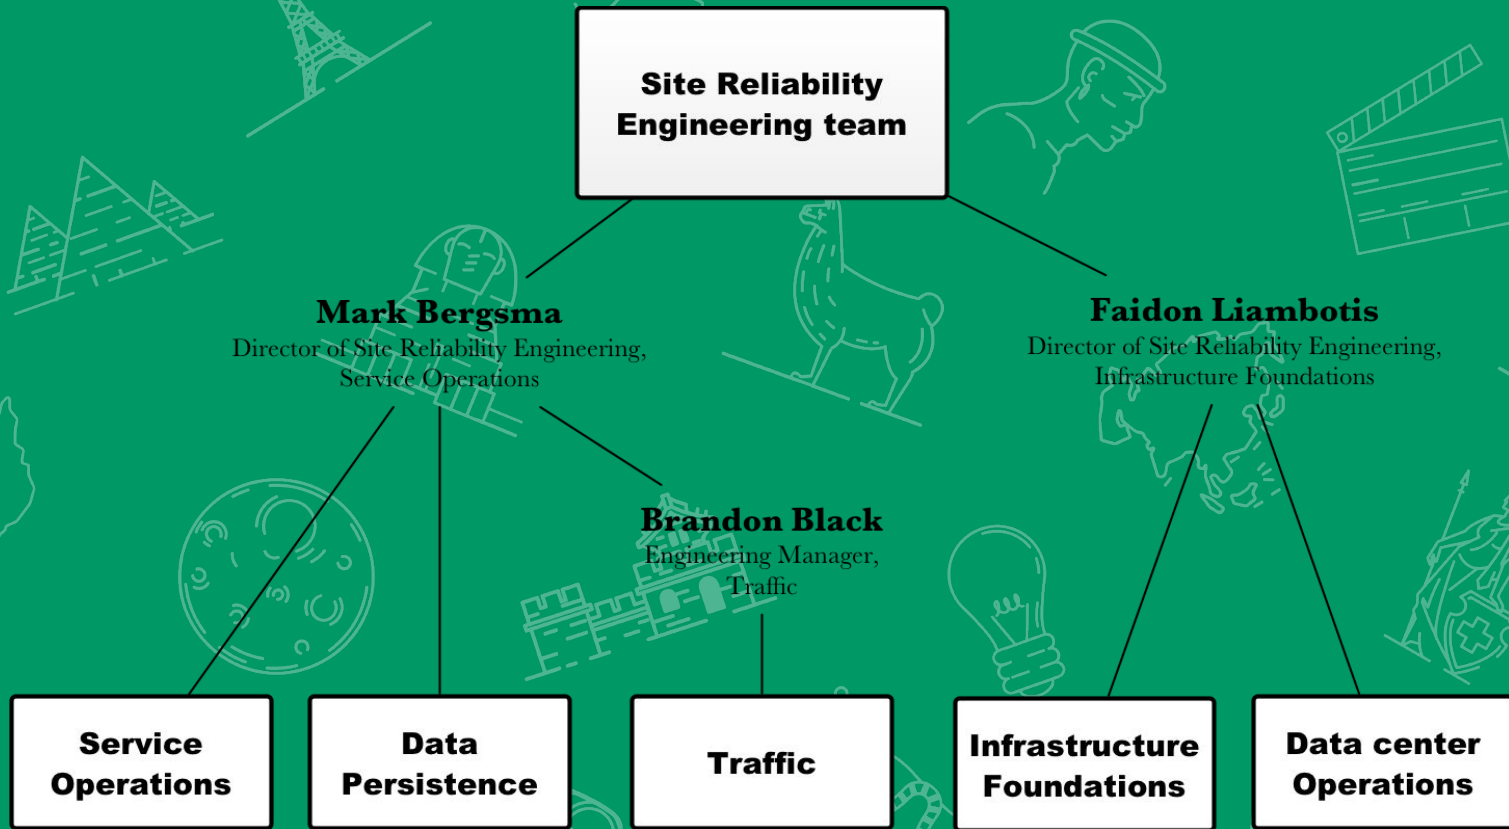
Reactive, one-offs and notable achievements not mentioned previously



**s/Tech Ops/SRE/**



# Site Reliability Engineering



# Other achievements

- Deployed new MediaWiki appservers in eqiad
- Implemented BGP RPKI
- Setup server PSU redundancy monitoring/alerting
- Ramped up encryption of syslog traffic
- Upgraded Grafana to 4.6
- Upgraded HHVM to 3.18.5 across the fleet, fixing crash
- Removed duplicate db indexes on pagelinks and templatelinks
- Assisted with MCR schema design and testing
- Released 2 Pybal releases
- Finished 3DES TLS cipher removal process [w/Community Liaisons]
- Implemented Pybal multiple BGP sessions, per-svc MED support
- Ported Puppet cache manifests to role/profile
- Implemented nginx prometheus monitoring
- Migrated European ASN 43821 to ASN 14907
- Kept Jenkins updated to latest LTS releases [w/RelEng]
- Upgraded PDU firmware and hardened configuration
- Renewed network support contracts
- Upgraded OTRS to 5.0.23, then .24, .25 and .26
- Decommissioned Ganglia
- Migrated mail servers (MXes) certificates to Let's Encrypt

# Other achievements

- Addressed various HTTP caching invalidation issues (incl. Commons)
- Increased MTU on all transit links to support Jumbo frames
- Puppetized and deployed Netbox experimental instance
- Decommissioned OCG
- Reduced paging false positives
- Fought seemingly unbounded growth of the JobQueue
- Implemented monitoring of SMART (HD/SSD health) metrics
- Performed MediaWiki schema changes for "refactored comments" [w/MediaWiki]
- Optimized Varnish timeouts, reducing HTTP errors
- Upgraded LibreNMS to latest release
- Redesigned apt.wikimedia.org repository structure
- Released new version of Cumin with additional features
- Sending "slow Varnish requests" to Logstash
- Implemented anycast for recursive DNS
- Implemented Icinga dashboard URLs to Grafana
- Migrating XML dumps to new (NFS) cluster [w/WMCS]
- Assisted with migration to Cassandra 3 [w/Services]

# Other achievements

- Introduced database multi-instance hosts in core shards
- Standardized declaration and generation of certificates and keys for TLS/SSL
- Assisted Scoring platform team with ORES outages
- Assisted Scoring platform with ORES stress tests
- Assisted with scap upgrade to 3.7.4 and then 3.7.5
- Provided Performance with dedicate infra for web performance
- Moved package builder infrastructure to stretch and added support for buster
- Implemented monitoring for PuppetDB
- Restricted Varnish execution environment capabilities
- Improved script that automatically reimage servers
- Rollout of new wikidiff 1.5.1 [w/WMDE] (supports moved paragraphs)
- Revised repository structure for apt.wikimedia.org
- Video scalers migrated to stretch
- Hired a Traffic Security Engineer
- Assisted Services in the ongoing migration of the Jobqueue
- Began renewal of our network services, saving \$45k/yr so far

# Release Engineering






# Other Highlights

- Released MediaWiki 1.30, with help from Cindy/MW Platform Team
- Released MediaWiki 1.28.1/1.27.2/1.23.16 security releases
- Conducted 10 MediaWiki weekly train rollouts
- Reduced production logspam issues by closing a net of 19 tasks (50 created, 69 closed), see graphs on next slide.
- Conducted 2 incident post-mortems
- Gave a Tech Talk on Selenium tests in Node.js
- Enabled real-time pop-up notifications in Phabricator

# Logspam fighting



In essence: We reduced the rate of logging in the WARNING level, while ERROR (more severe) level stayed consistent.

# Performance



# Other achievements

- Onboarded new Performance Team manager
- Established stable, repeatable synthetic testing infrastructure
- Wrote and published the first comprehensive guide to performance testing Mediawiki features
- Replaced thumbnail generation infrastructure with Thumbor for all public wikis, and published a detailed set of blog posts on the process and rationale for the migration
- Worked with Mozilla to identify performance impacts of Firefox 57 (“Quantum”), and filed a number of bugs related to performance regressions

# Cloud Services




# Other achievements

- Migrated all Wiki Replica traffic to new cluster and decommissioned old servers
- Onboarded Arturo Borrero González as Operations Engineer
- Worked with Administration and CTO to make Sarah Rodlund available for 20 hours of Technical Writing per week starting in Q3
- Team offsite in Austin in conjunction with KubeCon/CloudNativeCon
- Introduced Arturo to smoked brisket
- Hired Brooke Storm as Operations Engineer
- Filled Tech Support contract role
- Filled Product Manager contract role



# MediaWiki



# Other achievements

## Major architectural work for Structured Data on Commons supporting Multi-Content Revisions:

- Comment table development complete, schema changes in progress
- Actor table development nearly complete, in code review



## Other work to improve the MediaWiki platform:

- Helped Wikidata move to the normal deployment/branching system instead of a manual build
- Added a PSR-4 autoloader to MediaWiki core so developers aren't maintaining manual class and file listings
- Created "MinusX", a build tool to identify files with an executable bit set, that shouldn't
- Prettified test coverage listing on doc.wikimedia.org (see screenshot), in preparation for more work relating to test coverage
- Refactored edit conflict handling to support two column conflict (WMDE) Development