

# Quarterly review

## Architecture

### Q2 - 2015/16

Approximate team size during this quarter: 0.5 FTE  
*Time spent: strengthen 50%, focus 50%, experiment 0%*

Key performance indicator

? - we may choose RFC flow, but no metrics gathered yet			
---	--	--	--

## Q2 - Architecture

Objective: Prepare WikiDev '16



Objective	Measure of success	Status
Focus: Prepare WikiDev '16	Demonstrate that we can develop a modern system in an inclusive, consensus-oriented, open manner. Large number of diverse RfCs, with ample discussion prior to WikiDev '16	Preparation happened, summit resulted :-)

**Learning:** Sharpened collective thinking on [good meetings](#) and [how to build consensus](#)

## Q2 - Architecture

Objective: Strengthen: Improve ArchCom utility 

Objective	Measure of success	Status
Strengthen: Improve ArchCom utility	Set up Architecture Committee to succeed. Maintain one public IRC meeting per week	10 IRC RFC review meetings, plus 3 “Agenda bashing” meetings for WikiDev ‘16. Most were at the usual time (14:00 PT) <a href="#">Phab query for meetings</a>

**Learning:** Quote: *”[With] archcom being effective now (Which is great btw, so don't get me wrong) - I have noticed a trend where decisions are made in irc meetings, with little to no public discussion on the mailing list”*. -- [Brian Wolff on T118932, November 30](#)

## Q2 - Architecture

Objective: Strengthen: Improve ArchCom utility



Objective	Measure of success	Status
Experiment: ArchCom naming	Think through our loaded naming scheme, possibly renaming bits (e. g. “Architecture Committee” and “RfC”)	I didn’t get around to proposing anything on this front.

**Learning:** it’s ok that this one fell off. “ArchCom” and “RFC” seem to work well enough for now.

# Quarterly review

## Operations

### Q2 - 2015/16

Approximate team size during this quarter: 17 FTE  
*Time spent: strengthen 40%, focus 40%, experiment 20%*

Key performance indicator

Availability	99.979%	+0.037% from Q1	
--------------	---------	-----------------	--

# Q2 - Technical Operations

Objective: Encrypted PII 

Objective	Measure of success	Status
<p><i>Focus:</i> Encrypt PII on cross-data center links</p> <p><i>Team members involved:</i> 5</p>	<p>Encrypt cross-data center traffic for MySQL/MariaDB, Kafka, Cassandra, (Varnish) HTTP</p> <p>Decommission udp2log (in favor of Kafka)</p>	<p>Cross data center traffic for MariaDB, Kafka, Cassandra and between Varnish cache tiers is now encrypted.</p> <p>Legacy (and unencrypted) logging solution udp2log has been decommissioned.</p>

**Learning:** It has been confirmed by recent security vulnerability announcements ([Juniper](#) and Fortinet VPN backdoors) that we should not rely on transport link encryption by (proprietary) network hardware, as we have been hesitant to do since the start. We will continue to rely on FOSS implementations (TLS, IPsec) for encryption of PII data.

## Q2 - Technical Operations

Objective: User auth security 

Objective	Measure of success	Status
<p><i>Strengthen: Security:</i> Strengthen infrastructure user authentication</p> <p><i>Team members involved: 3</i></p>	<p>Establish org-wide workflow for creation and deactivation of user accounts in Production</p> <p>Implement (but don't enforce) two-factor authentication on SSH bastion hosts</p> <p>Migrate unmaintained OpenDJ LDAP to OpenLDAP</p>	<p>Organization-wide on/offboarding process has been implemented.</p> <p>Two-factor authentication is implemented and under experiment by all root users (using Yubikeys).</p> <p>OpenDJ has been replaced with OpenLDAP.</p>

**Learning:** First results of using Yubikeys for ssh authentication are promising, however some usability and reliability issues remain. Further experimentation with the OTP setup will continue into the next quarter.

## Q2 - Technical Operations

Objective: Incident monitoring



Objective	Measure of success	Status
<p><i>Strengthen:</i> Incident monitoring infrastructure</p> <p><i>Team members involved:</i> 2</p>	<p>Upgrade incident monitoring and expand it to codfw</p> <p>Reduce monitoring check execution lag to 0</p> <p>Create new abstractions for incident monitoring in Puppet</p>	<p>Software options were evaluated and we decided on a migration to Shinken (distributed monitoring). This migration has been prepared but didn't complete in time.</p> <p>New monitoring abstractions have been created in puppet, work is still ongoing.</p>

**Learning:** FQ2 is a challenging quarter every year with reduced time available due to the holiday periods, relatively high amount of staff taking unused vacation days, and deployment freezes. We should take this into account for our goal planning for this quarter.



## Q2 - Technical Operations

Objective: Bare metal Labs 

Objective	Measure of success	Status
<i>Experiment:</i> Support bare-metal servers in Labs  <i>Team members involved:</i> 2	Allow provisioning of physical servers within the Labs network, governed by Labs Puppetmaster and LDAP user accounts	Labs infrastructure and Puppet manifests now support bare metal servers in the Labs context. One bare metal server has been provisioned.

**Learning:** A separate small new test cluster was built for testing infrastructure changes to Labs itself, which proved very helpful for this goal and many projects in the future. OpenStack Ironic was also evaluated for this goal, but our current reliance on Nova-Network would imply a significant effort and create more technical debt until our migration to OpenStack Neutron completes.

## Q2 - Technical Operations

Greatly simplified, consolidated and optimized cache clusters and configuration

Migrated to a more reliable SMS paging service

Migrated video scalers to HHVM

Prepared OTRS upgrade to version 5 and fixed authorization security bug

Expanded Labs infrastructure with new virtualization nodes

Acquired two new IP transit port donations

Extended PyBal with etcd integration and alerting [with Performance]

Upgraded LibreNMS to a recent revision

Cleaned up and upgraded HHVM packages to 3.11

Deployed ORES [with Research]

## Other successes and misses

Implemented multiple instances per server for RESTbase [with Services]

Expanded esams network with a new core router

Handled a Mailman & Google Apps security incident [with OIT, Legal]

Implemented request rate limiting in Varnish

Discussed strategy for dealing with (hundreds of) domain names and SSL certificates [with Legal]

Migrated Procurement process from RT to Phabricator

Implemented canonical X-Client-IP header on all cache clusters

Implemented out-of-band network management for all server data center sites

EventBus consultation and deployment [with Analytics/Services]

# Q2 - Technical Operations

External monitoring **availability** metrics based on Catchpoint data.

Services audience	Availability cumulative average over 1 day interval	
	July 1 - September 30	Oct 1 - Dec 31 2015
(R)eaders	99.942%	<b>99.979%</b>
(C)ontent Contributors	99.935%	<b>99.974%</b>
(M)ovement Partners (volunteers, developers, etc.)	99.943%	<b>99.972%</b>
(E)xternal Partners	99.970%	<b>99.966%</b>
(D)onors	99.975%	<b>99.992%</b>

# Quarterly review

## Release Engineering

### Q2 - 2015/16

Approximate team size during this quarter: ... 6  
*Time spent: strengthen 9%, focus 28%, experiment 0%,  
Maintenance/Other 63%*

Key performance indicator

Time to merge in MW Core	~6 mins	-45% from Q4	<i>(still analyzing/making sense)</i>
--------------------------	---------	--------------	---------------------------------------

# Q2 - Release Engineering

**Objective:** Reduce CI wait time



Objective	Measure of success	Status
<b>STRENGTHEN</b> Reduce CI wait time	<ul style="list-style-type: none"><li>• CI cluster responds to spike in queued builds by starting and registering additional jenkins slaves (+)</li><li>• Migrate majority of CI jobs to Nodepool (-)</li></ul>	Nodepool is up and working.  However, the majority of CI jobs are not yet migrated (waiting on a solution for caching dependencies).


**Learning:** More collaboration is needed to address some of the more annoying problems (see: caching solution for dependencies).

# Q2 - Release Engineering

**Objective:** Reduce CI wait time



MediaWiki changes resident time in gate-and-submit



**Time for a MediaWiki core change to be merged by CI. Releng KPI.**

**Down from 11 minutes to average ~ 6 minutes**

**Spikes are MediaWiki releases**

# Q2 - Release Engineering

**Objective:** Reduce number of deploy tools



Objective	Measure of success	Status
<p><b>FOCUS</b></p> <p>Reduce number of deploy tools from 3 to 2</p>	<ul style="list-style-type: none"><li>• Migrate all Service team owned services and MW deploys to scap3</li></ul>	<p>Not completed. We are behind schedule in migration of services.</p> <p>The Dev Summit discussion, however, was very fruitful.</p> <p>Migrated some services on Beta Cluster to scap3 pre Holidays.</p>

**Learning:** After the slow beginnings to this project (slow due to other priorities taking time) we didn't allocate enough time in Q2 for "the last 10%" of functionality needed by Services.

# Q2 - Release Engineering

**Objective:** Retire Gerrit and Gitblit



Objective	Measure of success	Status
<p><b>FOCUS</b></p> <p>Retire Gerrit and Gitblit in favor of Phabricator</p>	<ul style="list-style-type: none"><li>• Decommission Gitblit (-)</li><li>• Code review RFC: creation, publication, discussion, feedback etc (+)</li></ul>	<p>Not completed.</p> <p>Gitblit is still live.</p> <p>Differential Migration RFC is published, has been discussed, and feedback being addressed (mostly from the Dev Summit).</p>

**Learning: Happy surprise-** The discussion at the Dev Summit about the migration was predominantly positive and there were no new blockers identified. It was a very encouraging conversation.



# Q2 - Release Engineering

**Objective:** Release MediaWiki 1.26



Objective	Measure of success	Status
<b>FOCUS</b> Release MediaWiki 1.26	<ul style="list-style-type: none"><li>• A quality MW 1.26 successfully released</li></ul>	Released.

**Learning:** Less of a learning and more of a perennial known: The releases we make for third-parties are still sub-par due to a number of reasons, one being internal priority setting (it is far less than 1 FTE per year).

## Q2 - Release Engineering

## Other successes and misses

Success: We are exploring non-Ruby browser test automation frameworks based on fruitful conversations with WMF engineering teams. Just in the exploratory stages so far.

Success: Positive working relationship with upstream Phabricator; some of our custom code is being subsumed by upstream features (great!).

Miss: WMF Log Errors graph (see appendix) interpretation is sub-optimal due to issues with tracking and config changes (to logging).

Miss: The silo-ification of the CI Scaling project is still too high.

# Q2 - Release Engineering

Category	Workflow	Comments	Type
Continuous Integration / Test maintenance	CI Config changes	<p><a href="#">Changes to our CI configuration</a> to enable/disable tests in a repository, etc</p> <ul style="list-style-type: none"><li>• 215 config change requests closed by RelEng members in Q2</li><li>• 610 total</li></ul> <p>Migrating repos to CI test entry point thanks to volunteers help (eg: Paladox)</p>	R
	MediaWiki Selenium (automated browser tests library)	<p>3 releases in Q2</p> <ul style="list-style-type: none"><li>• <a href="https://rubygems.org/gems/mediawiki_selenium/versions">https://rubygems.org/gems/mediawiki_selenium/versions</a></li></ul>	M
	Browser Test maintenance	Updating libraries, removing unused tests, etc	M

# Q2 - Release Engineering

Category	Workflow	Comments	Type
Service maintenance	Phabricator maintenance	Phabricator upgrades (bi-weekly, as needed), and user account help	M
	MediaWiki + SWAT deployments	Daily SWAT deployments Weekly (“train”) production deploys of MW+Extensions MW train deploys now a rolling responsibility within the team	M

# Quarterly review

## Services

### Q2 - 2015/16

Approximate team size during this quarter: 4 FTE  
*Time spent: strengthen 40%, focus 40%, experiment 20%*

Key performance indicators

REST API usage	248 req/s mean	15% up from Q4	∞% YoY
REST API HTML uptime (catchpoint)	99.97%	99.82%	∞% YoY

## Q2 - Services

Objective: API build-out



Objective	Measure of success	Status
Strengthen: Build out the REST API	Provide cached & purged REST API entry points for at least two high-traffic API use cases, and work with users to speed up page loads.	New entry points: page summary, wiktory definition, math SVG & MathML renders. Math used in production web; apps and hovercards (web) in process of switching over to summary & definition end points.

### Learning:

Need a well-defined thumbnail API. Started discussion & proposed strawman.

Broad support for maximizing caching with more REST end points at dev summit.

## Q2 - Services

### Objective: Event bus & change propagation



Objective	Measure of success	Status
Focus: Set up an event bus and change propagation solution	Key change events like edits, renames, deletes available in a publish-subscribe event bus. Basic change propagation / pre-generation functionality for REST services operational.	EventBus operational, MediaWiki event production just deployed. Change propagation service deployment delayed by dependency on event bus.

### Learning:

EventBus collaboration between Services & Analytics initially suffered from an unclear division of responsibilities. Clarifying responsibilities and dependencies helped to make progress.

## Q2 - Services

Objective: API-driven web front-end



Objective	Measure of success	Status
Experiment: Prototype an API-driven web front-end	Gather information to inform our longer-term front-end architecture by prototyping a service-worker / node.js based front-end.	Prototyped a service-worker front-end and established performance characteristics. Worked with Reading team & held a productive session on the topic at the dev summit.

### Learning:

Session at the dev summit showed broad support for the direction & a preference for progressive enhancement.



## Q2 - Services

## Other successes and misses

- Built API & storage behind VisualEditor's HTML / wikitext switch feature.
- Helped Analytics to build & publish the pageview API using RESTBase and Cassandra.
- Created a docker-based mediawiki-containers distribution solution.
- Worked with RelEng on Scap3, began to start being used (EventBus HTTP proxy service).
- Initiated a Content working group with several productive discussions on structured content, separation of content from presentation & an API-driven front-end architecture. Many of these ideas featured at the dev summit.

# Q2 - Services

Category	Workflow	Comments	Type
Service development & operations	Mentor new service development	<ul style="list-style-type: none"><li>- AQS (pageview) service deployment and refinement, including schema improvements and Cassandra tuning.</li></ul> <p>Approx effort: 0.5 FTE</p>	N
	Maintenance of existing third-party services	<ul style="list-style-type: none"><li>- Mentored Citoid development (Marielle Volz): backend for VisualEditor citation feature. Transitioning ownership to Parsing team.</li><li>- Worked with volunteer Moritz Schubotz towards a cleaner Math render mode.</li><li>- Worked with the Reading team to expand mobileapps / content service functionality.</li></ul> <p>Approx effort: 0.5 FTE</p>	M
	Cassandra / service operations	<p>Expanded Cassandra capacity, upgraded to new versions &amp; refined configuration.</p> <p>Prepared service deployments, investigated service issues &amp; improved shared operational service infrastructure.</p> <p>Approx effort: 1 FTE</p>	R

# Quarterly review

## Security

### Q2 - 2015/16

Approximate team size during this quarter: 2.1 FTE (+Moritz)  
*Time spent (strategic): strengthen 25%, focus 0%, experiment 75%*

Key performance indicator

<a href="#">Critical and High Priority Security Bugs</a>	REDACTED	REDACTED from Q4	n/a
--	----------	------------------	-----

## Q2 - Security

### Objective: Static Security Testing



Objective	Measure of success	Status
Automated security static analysis of MediaWiki	Automatically scan core and one extension weekly.	Security team chose Veracode for scanning, and are submitting core and two extensions for each deployment branch.

#### Learning:

The quality of static security testing tools for PHP is still immature. The WMF could significantly contribute in this area, but the Security Team lacks the resources to pursue this.

## Q2 - Security

Objective: Expand Training 

Objective	Measure of success	Status
Expand developer training	Present secure SDL training for community and staff.  (stretch) Develop and present security training materials for DevOps and Mobile	The team documented the optimal process for other teams to integrate security into their development process, but training has not been given.

### Learning:

The team adapted to account for a security incident. Our goals should be more flexible to handle this in the future.

## Q2 - Security

Objective: Expand Metrics 

Objective	Measure of success	Status
(stretch) Document and report initial metrics for security reviews	Metrics reported	This stretch goal was not started.

**Learning:**

## Q2 - Security

## Other successes and misses

Learning: This quarter we asked teams to schedule anticipated reviews at the start of the quarter, to allow the security team to plan for an increasing number of reviews. This allowed the Security Team to better schedule capacity, but required teams with unscheduled work to either delay or reprioritize reviews they had previously requested. While this is probably healthy for teams, it may discourage personal/skunkworks projects, and the Security Team should ideally have capacity to handle some of those as well.

Success: The Security Team supported or drove two community RFC's (~~one passed, one anticipated to pass~~ both have passed) related to password policies for accounts with privileges. When implemented, these policies should improve account security for privileged users, and overall safety for the users of our sites.

Success: Regular vulnerability scanning (for PCI compliance) established, and could be provided to other teams as well.


## Q2 - Security

Category	Workflow	Comments	Type
MediaWiki Security	Security Bugs	REDACTED. MediaWiki 1.26.1 released, fixing 6 core issues.	R
	Security Reviews	19 reviews opened, performed 13 reviews	R
Privacy and Security Consulting	Architecture and design consultations	Reading Infrastructure - AuthManager / SessionManager Privacy - Data Access Policies, Privacy by Design, Pen tests	N
DFIR	Incident Response	REDACTED	R



## Q2 - Team name

Category	Workflow	Comments	Type
Training	MediaWiki Security	Secure Coding for MediaWiki: 14 attendees	M
	Operational Security for Staff	34 attendees	N
Compliance	Monthly vulnerability scan for PCI	OpenVAS packaged and run monthly	N



	CI Scaling	Scap3 Migratio	Diff Migration	Maintenance	Nonsense stuff	Total/Total
Antoine	36.92%	3.08%	2.31%	36.54%	8.85%	87.69%
Chad	0.00%	7.50%	11.67%	42.92%	24.17%	86.25%
Dan	0.00%	38.00%	2.50%	31.00%	9.44%	80.94%
Mukunda	0.00%	32.50%	15.83%	27.50%	24.17%	100.00%
Tyler	0.00%	40.00%	1.54%	43.85%	3.85%	89.23%
Zeljko	9.23%	0.00%	0.00%	37.69%	53.08%	100.00%
Avg	46.15%	109.23%	31.15%	206.92%	116.92%	510.38%
Percent of Cap.	7.69%	18.21%	5.19%	34.49%	19.49%	85.06%
Percent of total/total	9%	21%	6%	41%	23%	100%
	Strengthen	9%				
	Focus	28%				
	Maint/Nonsense	63%				
	Experiment	0%				

Warning: Hand wavey self reported data.

# Q2 - Release Engineering

2016-01 [ [edit](#) | [edit source](#) ]

On a scale of 1-5 where 1 means "no idea" and 5 means "expert", rate your knowledge of the below skills.

	<u>Continuous Integration</u>			<u>Integration Environments</u>	<u>Deploying software</u>					<u>Development environment</u>	<u>Developer Tools Support</u>		<u>MediaWiki Releases</u>	
	Jenkins maint	unit tests	integration tests	Beta Cluster	Deploying new MW branches	backports SWAT deploys	Developing scap	Dev'ing Trebuchet	Debugging or Reporting log errors	MediaWiki-Vagrant	Gerrit maint	Phab maint	Doing major releases	Doing security releases
<b>Antoine</b>	★★★★★	★★★★★	★★★★★	★★★★★	★★★★★ <sup>(-1)</sup>	★★★★★ <sup>(-1)</sup>	★★★★★ <sup>(-2)</sup>	★★★★★ <sup>(-1)</sup>	★★★★★	★★★★★	★★★★★ <sup>(+1)</sup>	★★★★★	★★★★★	★★★★★
<b>Chad</b>	★★★★★	★★★★★	★★★★★	★★★★★	★★★★★	★★★★★	★★★★★	★★★★★	★★★★★	★★★★★	★★★★★	★★★★★ <sup>(+1)</sup>	★★★★★ <sup>(+2)</sup>	★★★★★
<b>Dan</b>	★★★★★ <sup>(-1)</sup>	★★★★★	★★★★★	★★★★★	★★★★★	★★★★★	★★★★★	★★★★★	★★★★★	★★★★★	★★★★★	★★★★★ <sup>(+1)</sup>	★★★★★	★★★★★
<b>Mukunda</b>	★★★★★ <sup>(+1)</sup>	★★★★★	★★★★★	★★★★★	★★★★★	★★★★★	★★★★★	★★★★★	★★★★★	★★★★★ <sup>(-1)</sup>	★★★★★ <sup>(-1)</sup>	★★★★★	★★★★★	★★★★★
<b>Tyler</b>	★★★★★	★★★★★	★★★★★	★★★★★ <sup>(-1)</sup>	★★★★★ <sup>(+1)</sup>	★★★★★	★★★★★	★★★★★	★★★★★ <sup>(+1)</sup>	★★★★★	★★★★★	★★★★★	★★★★★	★★★★★
<b>Zeljko</b>	★★★★★	★★★★★	★★★★★	★★★★★	★★★★★	★★★★★	★★★★★	★★★★★	★★★★★	★★★★★	★★★★★	★★★★★	★★★★★	★★★★★

Warning: Hand wavy self reported data.

REDACTED