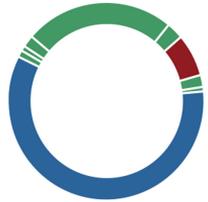


Analytics Quarterly Review

Wikimedia Foundation - Q2 of FY14-15



Analytics Engineering

What we said

Q2 goals (1/3)

Area	Goal	Why
Fill out Vital Signs	Visualize PageViews & PageViews broken down by target site	<i>Make traffic data public (including mobile) in an easy to consume format</i>
	Implement Breakdowns for existing metrics	<i>Start exposing data distinguishing mobile and desktop sites</i>
Prototype	Investigate Open Source BI Tool	<i>Look at free alternatives to visualize data</i>

Q2 goals (2/3)

Area	Goal	Why
Wikimetrics	User can expand uploaded cohorts to include users across all wikis	<i>Grantmaking requested these features for grantees to evaluate the success of their programs</i>
	User can delete a member of a cohort	
	User can tag a cohort with pre-defined tags	
	User can list cohorts from an existing tag	

Q2 goals (3/3)

Area	Goal	Why
EventLogging	ID and direct the purging of raw logs older than 90 days (executed by ops)	<i>Housekeeping & compliance with data retention guidelines</i>
	ID and direct the purging of rows older than 90 days (executed by sean and ops with input from product)	
	Work with product teams who need to keep some data longer than 90 days to comply with the privacy policy	

What we did

Q2 goals (1/3)

Area	Goal	Status
Vital Signs	Visualize PageViews & PageViews broken down by target site	[done]
	Implement Breakdowns for existing metrics	[blocked] cannot get needed performance in labs
Prototype	Investigate Open Source BI Tool	[done] using PV data

Q2 goals (2/3)

Area	Goal	Status
Wikimetrics	User can expand uploaded cohorts to include users across all wikis	[done]
	User can delete a member of a cohort	[done]
	User can tag a cohort with pre-defined tags	[done]
	User can list cohorts from an existing tag	[de-prioritized] implemented reporting instead

Q2 goals (3/3)

Area	Goal	Status
EventLogging	ID and direct the purging of raw logs older than 90 days (executed by ops)	[ongoing] 2 of 3 servers have logs auto-purged
	ID and direct the purging of rows older than 90 days (executed by sean and ops with input from product)	[ongoing] 3 of 110 schemas are auto-purged
	Work with product teams who need to keep some data longer than 90 days to comply with the privacy policy	[ongoing]

Other key accomplishments

Producing Datasets from Hadoop

GLAM, project counts, page counts

Increasing throughput of EventLogging (Batch Inserts)

Inter Team collaboration to support sendBeacon() in EventLogging

Support product teams in the creation of EventLogging Dashboards

Teams: Mobile, Collaboration, Language, Edit

Documented steps to create a dashboard

What we learned

What we learned:

Users love Pentaho

Labs is not a stable environment for

Generating metrics for Vital Signs

Wikimetrics

Resources

EventLogging maintenance takes significant time

We don't have the DBA resources to build a Data Warehouse quickly

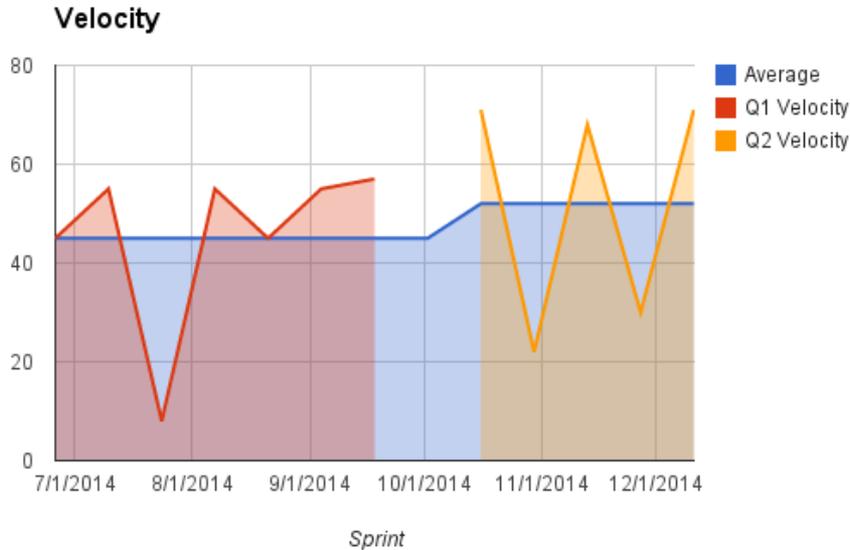
What's next

Q3 goals

Area	Goal	Why
VisualEditor	<p>Collect and continually visualize key metrics including but not limited to:</p> <ul style="list-style-type: none">● initialization time and save time (perceived and actual)● edit funnel from start of an edit action to successful save<ul style="list-style-type: none">○ save rates, abort rates, error rates○ across experiences: mobile, desktop, wikitext, VE.● Prioritization of metrics will be driven by VisualEditor success criteria.● Provide ad-hoc support as needed to VE team	<p><i>Support the Editing team so they can be ready to roll out VE to larger wikis</i></p>
Unique Clients	<p>Develop a prototype report and visualization on Unique Clients per project per day and month.</p> <p>Goal is dependent on community engagement, engineering design and Mediawiki work.</p>	<p><i>Unique Clients is a much needed metric</i></p>

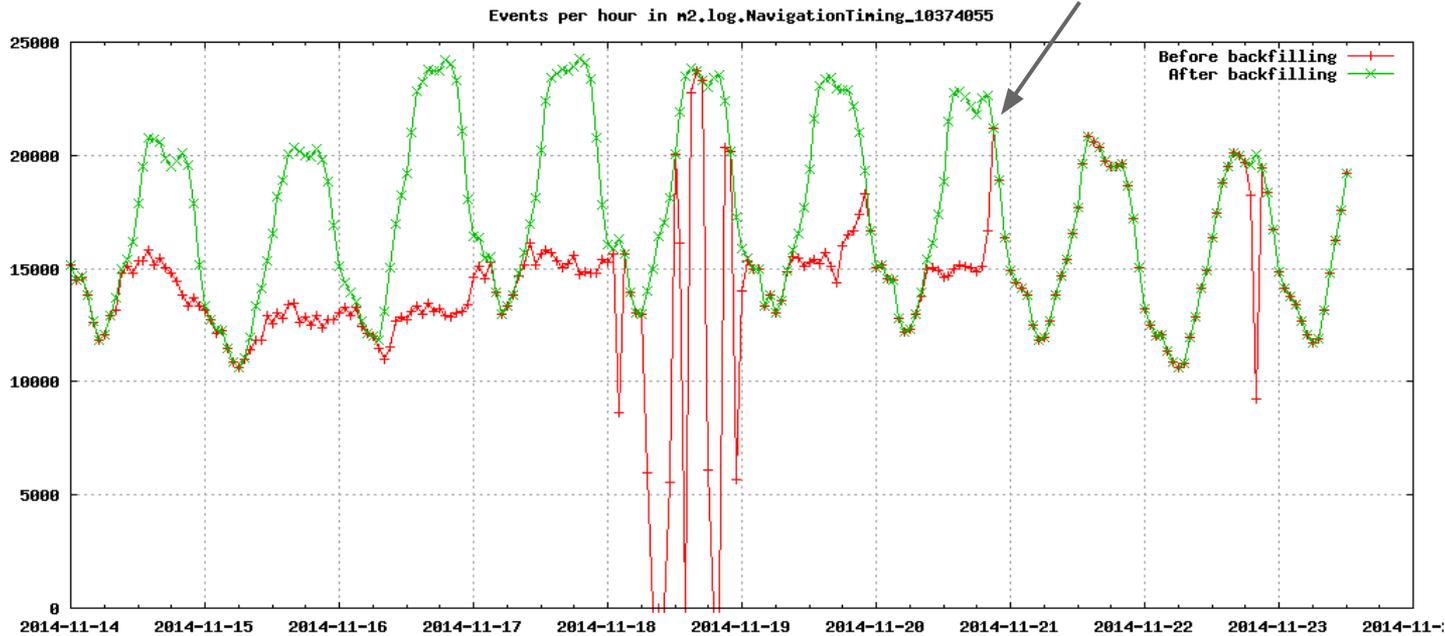
Metrics and callouts

Metrics and callouts: Velocity

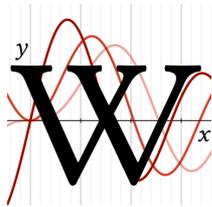


Velocity increased this quarter with addition of Marcel and Grace to the team.

Metrics and callouts: EventLogging



Source: <https://lists.wikimedia.org/pipermail/analytics/2014-November/002808.html>



Research and Data

What we said

Background: Focus areas in Q2

Readership

Mobile contributions

Performance

Fundraising

Expand
team capacity

Q2 goals (1/3)

Area	Goal	Why
Readership	Finalize the new pageview definition , incorporating feedback from the first draft	<i>Legacy PV data cannot answer many questions we have about traffic / consumption</i>
	Deliver a unique client definition and implementation proposal	<i>Overcome the limitations of PVs and low-quality external data (comScore)</i>
	Generate reports on readership trends from historical data, based on the new pageview definitions	<i>Build organizational alignment around traffic trends</i>
	Generate recurrent reports of session-based metrics	<i>Help Mobile App team drive reader engagement</i>

Q2 goals (2/3)

Area	Goal	Why
Mobile contributions	Instrumentation, experimental design and evaluation strategy for mobile microcontributions	<i>Design a sound experimental strategy to measure the impact of new ways of contributing via mobile interfaces</i>
	Analysis of data collected from the first rollout of microcontributions	
	Analysis of statements suitable for microcontributions	
	Provide breakdown definitions (mobile/desktop/apps) for Vital Signs metrics	<i>Standardize engagement metrics across platforms</i>
Performance	Run a test on the impact of HHVM on newcomer engagement	<i>Determine if HHVM affects engagement beyond sheer site performance</i>
	Analyze the performance of editing interfaces	<i>Provide baselines for editing interfaces</i>
Fundraising	Ad hoc analytics to support the English campaign	<i>Pay our salaries better and faster</i>

Q2 goals (3/3)

Area	Goal	Why
Expand team capacity	Start formal collaborations with external researchers under NDAs	<i>Leverage research community to produce datasets and research we have no internal capacity for</i>
	Hire a traffic researcher	<i>Increase team capacity to analyze readership</i>
	Provide ad hoc support to other teams in the organization	<i>Provide data and research support where needed</i>

What we did

Background: Q2 achievements

Q2 was a great quarter, we achieved a lot

Deviations from the plan were mostly caused by external dependencies

Other achievements beyond stated goals

Q2 goals (1/3)

Area	Goal	Status
Readership	Finalize the new pageview definition , incorporating feedback from the first draft	[done] PV definition and implementation prototyped and evaluated against the sampled logs, ported to a UDF for productization in hadoop
	Deliver a unique client definition and implementation proposal	[in progress] provided definition in coordination with the devs, pending public consultation and implementation details (opt out)
	Generate reports on readership trends from historical data, based on the new pageview definitions	[done] used data from the new PV definition to create reports for the Board, the FR team and a public Monthly Metrics presentation.
	Generate recurrent reports of session-based metrics	[done] generated ad-hoc reports for the mobile team, now handed off to the Devs for productization.

Q2 goals (2/3)

Mobile contributions	Instrumentation, experimental design and evaluation strategy for mobile microcontributions	[done] completed instrumentation and experimental design for 3 tests
	Analysis of data collected from the first rollout of microcontributions	[done] performed data QA and analysis for the 2 tests deployed in the quarter
	Analysis of statements suitable for microcontributions	[done] produced reports for the Mobile team to design WikiGrok campaigns
	Provide breakdown definitions (mobile/desktop/apps) for Vital Signs metrics	[on hold] stability issues on Labs
Performance	Run a test on the impact of HHVM on newcomer engagement	[done] test found no significant evidence of HHVM's impact on user engagement
	Analyze the performance of editing interfaces	[done] designed instrumentation and passed off to Analytics Dev for reporting in Q3
Fundraising	Support the team with the English campaign	[done] fully redesigned FR analytics tooling

Q2 goals (3/3)

Area	Goal	Status
Expand team capacity	Start formal collaborations with external researchers under NDAs	[in progress] started a collaboration with Stanford. Made substantial progress on collaboration with Los Alamos, blocked on Legal issues.
	Hire a traffic researcher	[on hold] posted the job, interviewed candidates but froze the opening after decision to re-scope it
	Provide ad hoc support to other teams in the organization	[done] supported other teams, including: Design, Zero, Language Engineering, Legal.

Other key accomplishments

Supported **strategy team**, particularly on partnerships/syndication track

Measuring “value added”: completed infrastructure work and test results for English Wikipedia importance classes.

Revision scoring: High accuracy “revert” classifier for Enwiki and Ptwiki

What's next

Background Focus areas in Q3

Services

Reporting pipeline

Mobile

Fundraising

Formal collaborations
continued

Streamline process

Q3 goals (1/3)

Area	Goal	Why
Services	Prototype revscoring service	<i>Provide quality measurement as a service</i>
	Prototype services to support mobile products (stretch)	<i>Improve reader engagement (related articles), content discoverability (trending topics, relevant collections)</i>
Reporting pipeline	Finalize UC implementation requirements	<i>Clear all research dependencies for implementing UCs</i>
	Support Dev team with PV productization	<i>Help Devs replace legacy PV data in official dashboards</i>
	Continued responsibility for ad-hoc traffic analysis and reports	<i>Need a deeper understanding of reader behavior, support customers where automated reporting is not available</i>

Q3 goals (2/3)

Area	Goal	Why
Mobile	Continue research support to WikiGrok	<i>Help Mobile optimize the feature through further testing</i>
	Research support for Collections	<i>Help Mobile design /assess first prototype</i>
Fundraising	Develop method for A/B testing with sampled data (stretch)	<i>Identify winning banners more accurately</i>
	Develop proposal to run FR as a minimization problem	<i>Minimize impact of campaigns on site usability while meeting the budget</i>

Q3 goals (3/3)

Area	Goal	Why
Formal collaborations	Wrap up first research project with Stanford (improving linking structure)	<i>Improve readers' access to content</i>
	Kick off second research project with Stanford (knowledge gaps)	<i>Identify gaps in Wikipedia content across languages and work towards closing the gaps</i>
Streamline process	Continue to provide ad hoc support to other teams , pending prioritization	<i>Provide research and data support where it's needed</i>
	Continue work with Team Practices on team process and iterations	<i>Make process more transparent to customers, prepare for Phab migration</i>

Asks

Asks

- **Capacity management** - so we are not reacting to bottlenecks in Hadoop, Kafka, EL and the stat boxen
- **DBA support** - It is on the critical path to improving EventLogging, Wikimetrics and Vital Signs
- **Event Logging** - we need more resources to monitor and troubleshoot and maintain this system as more teams rely on it
- **Labs support**: teams rely on Limn, Vital Signs, Wikimetrics and fires start when Labs goes down or changes are made
- **Hadoop/Kafka** - Investigate outside hosting for crunching public data to alleviate capacity issues on our cluster
- **Dumps** - we rely on the dumps to populate much of stats.wikimedia.org. Delays in dump processing impact our ability to deliver
- Better support for **data QA** - most costly and painful part of the data analysis workflow
- Access to Yuvi and Gabriel's time - continuous support is critical for **service-related work**