

Technology Dept Tuning Session Q1 FY21-22



WIKIMEDIA
FOUNDATION

Platform Evolution



Overview

Created Execution Plans to understand and document our methodologies and resourcing for achieving our Annual Objectives.

Progress and Challenges

Deprioritized Global Equitable System Performance until Q3 due to staffing changes. Will be planning resourcing and scope in Q2 to understand feasibility or beginning in January.

We are reevaluating our MTP metrics for a secure and sustainable platform. As a department we are beginning adoption of a new framework for our operations which uses 4 KPIs based on industry research (See [Accelerate](#)).

OKRs

Data as a Service	Yellow
Platform Excellence: Global equitable system performance	Red
Product Platform	Yellow

Actions to do in Q2

- Read Accelerate as a department and update metrics
- Create Execution Plan for Global Equitable System Performance



Platform Evolution



MTP Outcomes

We will build tooling for internal and external development and reuse of code and content

MTP Metrics

An 80% increase in structured data used (uptake) across wikis.

Baseline: 87M pages across Wikimedia projects use Wikidata as of April 2020.

An 10% increase in non-text (e.g. Commons) content used across wikis.

Baseline: 31.2M items from Commons are used across Wikimedia projects as of April 2020.

Y3
Goal

Q1
Status

Q2
Status

Q3
Status

Q4
Status

80%
increase

+39.5%
increase in
percentage of
use from 4/20
baseline

10%
increase

+14.5%
increase in
percentage of
use from 4/20
baseline



Platform Evolution Metrics



MTP Outcomes

A secure and sustainable platform that empowers a thriving developer community with the ease of software -as-a-service tooling

MTP Metrics

	Y3 Goal	Q1 Status	Q2 Status	Q3 Status	Q4 Status
30% increase of tool maintainers Baseline: 1880 maintainers in Q2 (FY 19/20)	25% increase from baseline	17.4% (2207)			
10% (4.2 / 5) increase in developer satisfaction Baseline: 2019 developer satisfaction: 3.8 / 5	4% (4.0)	Next survey will be conducted in Q3			
10% decrease in code review time Baseline: 19 days in June 2019	6% (17.6 days)	-72% (31 days)			
20% decrease in outstanding code reviews Baseline: 1032 code reviews in June 2021	4% (991 reviews)	796 (22%)			



MTP Priority OKRs



WIKIMEDIA
FOUNDATION

Data as a Service



Objective:

Wikimedia application data is easily discoverable and well-prepared to enable data-informed decision making, application development, and research by the community and the Foundation.

Wikimedia's new machine learning infrastructure reached its "Hello World" status to host real, production models. This puts the team in a good position to launch an MVP of Lift Wing in the coming months.

Data Management requirements gathering and planning are in progress.

On target to complete hiring to support KR1 & KR2 delivery - Data PM, Data Governance specialist and 3 Data Engineering hires.

Target quarter for completion: Q4 FY21-22



Data as a Service



Key Results

KR1 - Establish organizational data management structure. Build a browseable, shareable data dictionary, and describe 25% of known data elements.

Baseline: Existing informal documentation of misc datasets

KR2 - Enable efficient program evaluation and decision support through three novel use cases.

Baseline: Siloed analytics by team

KR3 - Build machine learning services. Operationalize a ML governance strategy for the Foundation, and create ways to understand, evaluate, and provide feedback on ML models.

Baseline: No MVP, No ML governance strategy in place

	Year Goal	Q1 Status	Q2 Status	Q3 Status	Q4 Status
	Data Discoverability	Resourcing, Info gathering, Planning			
	3 Integrated Analytics towards Self-Service	0 out of 3 On Track			
	Lift Wing MVP ML Governance Strategy	MVP Infra: ~70% MVP Models: ~80% MVP API: ~10% Governance: ~5%			



Global equitable system performance

Objective:

People in emerging communities have the same level performance for their contribution workflows as those in established communities.

Deferred to Q3 due to staff departures in Q1. Planning of resources and reevaluating of scope will happen during Q2. Will provide an update on changes in next tuning session.

Target quarter for completion: Q4 FY21-22



Department: Technology

Global equitable system performance

Key Results

	Year Goal	Q1 Status	Q2 Status	Q3 Status	Q4 Status
<p>KR1 - Plan and resource 3 projects based on our strategy for increasing performance for contributors globally, specifically addressing caching and content composition for logged in users. Baseline: 0 projects planned and resourced</p>	3 projects	0%			
<p>KR2 - Complete 3 performance projects which decrease the time it takes to complete existing contribution workflows within emerging markets (Based on Largest Contentful Paint, an industry metric for perceived load speed) Baseline: 0 projects which result in decreasing the by at least 5% within a laboratory regional test</p>	3 projects	0%			
<p>KR3 - Make 3 product decisions made by using performance data gathered from testing against standard emerging community performance personas. Baseline: 3 product decisions made</p>	3 product decisions	0%			

Product Platform



Objective:

Our platform and processes are ready and able to invite all of the world's population to join us.

Foundational work has begun to enable KR1 and KR3 including new technology to support necessary data platforms (KR1) and initial development of a shared component library to support the move to Vue (KR3).

Parser consolidation work continues with active work to address rendering, performance, and functionality gaps.

Challenges:

The Structured Data team has seen significant staff turnover, slowing velocity and potentially putting our KR at risk for the year.

The Search team is stretched thin between Search & Relevancy, and Query Services, leading to missed deadlines (WDQS Streaming Updater, WCQS production) and inability to address important work (Query Completion). WDQS Blazegraph migration KRs will need to be revised to reflect size of this project (underestimated during annual planning).

Target quarter for completion: Q4 FY21-22



Product Platform



Key Results

KR1: Machines are able to recognize Wikimedia content and suggest relations to other Wikimedia content in at least 2 wikis, enabling experimentation with at least two new strategic features

Baseline: 0 experimental features

KR2: Scale our capability to surface and utilize relationships between content across two different use cases and the current WDQS service level continues to support editing of Wikidata

Baseline: 25 10+ min server events in the past 30 days (11 Sept 2021 to 11 Oct 2021; manual count on [this graph](#))

KR3: Improve the satisfaction and velocity of staff and volunteer developer contributions to frontend components of Wikimedia projects by 10% through migration to the Vue frontend framework

Baseline: 0%

KR4: Reduce platform complexity by fully integrating Parsoid into the rendering pipeline so that by EOFY (a) all talk pages are rendered with Parsoid enabling advanced DiscussionTools features (b) over 95% of Main namespace pages on all non-language-variant wikis have insignificant visual diffs compared to current output (c) Parsoid HTML is used for read views on the main namespace on at least 1 wiki.

Baseline: TBD

Year Goal	Q1 Status	Q2 Status	Q3 Status	Q4 Status
2 new experimental features released in 2 wikis	Requirements for 1 feature developed			
WDQS & WCQS update lag < 10min	Updates to improve lag in progress			
10% increase in satisfaction and velocity	0% increase			
100% talk pages rendered 95% Main namespace pages	Parsoid renders at least ~90% of pages from 20 wikis without visual diffs			

Efficacy & Resilience OKRs



WIKIMEDIA
FOUNDATION

Impact Culture: Business Intelligence

Objective:

Foundation leadership have access to organizational data, insights and dashboards for impactful decision making.

-
- Compiled short list of metrics for measuring ~60% of Annual Plan and MTP activities. This work was scheduled for completion in Q1, but is now slated for completion in early Q2.
 - Met with representatives from Advancement, Tech, T&C, Comms, and F&A to discuss MTP, AP, and business health metric needs for FY 21-22. Alignment with remaining departments planned in Q2.
 - Completed analysis of FY 21-22 Annual Planning Process feedback from APP leads and staff and shared findings with APP team to inform future process and communication improvement.
 - Theory of Change and SROI metric KR is on hold until January 2022 and will be carried forward in collaboration with F&A.

Target quarter for completion: Q4 FY 21-22



Department: Technology

Impact Culture: Business Intelligence

Key Results

KR1: Each quarter, 50% of departments work with the GDI team on the theory of change; develop org-level business metrics, including Social Return on Investment (SROI) measures by Q2

Baseline: No ToC or SROI measures established.

KR2: By end of Jan 2022, design and roll out a data dashboard that can be used by at least 50% of Foundation departments for annual planning.

Baseline: No dashboard.

	Year Goal	Q1 Status	Q2 Status	Q3 Status	Q4 Status
<p>KR1: Each quarter, 50% of departments work with the GDI team on the theory of change; develop org-level business metrics, including Social Return on Investment (SROI) measures by Q2</p> <p>Baseline: No ToC or SROI measures established.</p>	<p>50% of departments work on Theory of Change</p> <p>SROI metrics defined</p>	<p>On hold until January</p>			
<p>KR2: By end of Jan 2022, design and roll out a data dashboard that can be used by at least 50% of Foundation departments for annual planning.</p> <p>Baseline: No dashboard.</p>	<p>Dashboard launched by Jan. 2022</p> <p>50% of departments use dashboard for annual planning</p>	<p>60% of required metrics identified</p>			



Technical Community Building



Objective:

Our technical community is thriving and has a clear, consistent means to discover, build, and deploy applications that support community workflows, invent new forms of content creation and consumption, and leverage Wikimedia's APIs and data beyond the core wiki experience.


Discoverable Docs (KR 1): Content strategy, requirements for MVP of the Developer Portal, and acceptance criteria for accuracy have been defined in Q1. In Q2 we will start the work on improving key documents and technical implementation of the Developer Portal.

Curated Pathways (KR 2): We conducted an initial audit of PAWS and Quarry tasks and implemented two new features into Quarry. Work on automated workflows for enabling or deleting tools from Toolforge is almost completed. A call for participation in the Small Wiki Toolkits initiative at the end of Q1 resulted in interest of communities from West Africa and Malta. Exploring how to scale the Small Wiki Toolkits initiative further will be a focus of our work in Q2.

Toolhub (KR 3): A series of blockers delayed the launch of Toolhub into early Q2. In parallel to the engineering work, we started with the research on “quality signals” through a number of input sessions with users. The outcome as well as the feedback we’re expecting after the launch will inform the next feature set of Toolhub. We’re still on track with our planned deliverables for FY 21/22.



Target quarter for completion: Q4 FY 21/22



[Toolhub](#)



[Small Wiki Toolkits project logo](#) by Sethiasta, cc-by sa 4.0

Technical Community Building



Key Results

KR1: A single entry point into our technical documentation, combined with accurate landing pages, enables technologists of all levels to discover Wikimedia's technical areas and engage in communities of practice (acceptance criteria for accuracy defined by Q1, 100% of landing pages meet criteria by end of Q3, launch of developer portal by Q4)

Baseline: Content draft and research was done in FY 20/21

KR2: At least 10 underserved wiki communities have a set of technical skills and capabilities enabled through globally localized partnerships and curated pathways for learning and contribution, in support of scaling the ecosystem of tool developers (10%)

Baseline: 2162 Tool maintainers; communities: TBD in Q2

KR3: Users can assess the reliability of tools for adoption, contribution or research based on a system of quality signals (co-maintainers, docs, recent editing usage, published source code, endorsements, etc) within Toolhub (research by Q2, code by Q3, adoption by Q4).

Baseline: 0

	Year Goal	Q1 Status	Q2 Status	Q3 Status	Q4 Status
	Developer portal is launched 100% of landing pages meet accuracy criteria	Q1 goals completed			
	10% increase in tool maintainers 10 communities	2.4% (2207) increase in tool maintainers			
	Release of Toolhub 1.1; adoption	On track			



Platform Excellence: Resilience



Objective:

Our services, infrastructure and data are resilient to and/or quick to recover from unexpected malicious or non-malicious events

The [global industry-wide chip shortage](#) continued to affect lead times for hardware, creating a significant risk of (further) delays to our two data center projects (**KR1**). We have **responded** to this emerging challenge by spending significant time in complex vendor negotiations, evaluation of alternative solutions, and by drafting various stopgap solutions. We **remain on-track** for the EMEA region expansion in **Marseille**—meeting the previously revised target of Q2—as well as the **expansion** of our **main data center** in Virginia.

For our **incident response** initiative (**KR2**), we **remain on-track**, having succeeded in the foundational goals set out for Q1, namely revamping and restructuring a cross-team working group responsible for this effort, performing a gap analysis on on existing People, Processes and Tooling (PPT), and defining an incident engagement scorecard to measure our future progress against. In Q2, we expect to set a baseline score for incidents, and formulate a revamped roadmap.

Finally, in Q1 we set baselines for our current **security services** (**KR3**) to help us identify where our time is spent and which services are good candidates for enhancing our capabilities. In Q2, we expect to **further refine** our service portfolio, explore what can we **checklist**, what can we **automate** and where can we start using **risk assessments** to further accelerate our service delivery model.

Target quarter for completion: Q4 FY21/22



Department: Technology

Platform Excellence: Resilience



Key Results

KR1: Wikimedia's infrastructure is scaled to address known compute, storage and traffic capacity risks, by adding a new data center in EMEA (by end of Q2), expanding our main data center by at least 20% (by end of Q2), and by documenting two new capacity plans (by end of Q4)

Baseline: One EMEA data center; 32 [racks](#) of space in our main data center; zero documented capacity plans

KR2: Service and security operational issues are detected, escalated, remediated and communicated to stakeholders and the movement, as measured by a 20% incident score improvement

Baseline: To be defined in Q1/Q2

KR3: Security and privacy services are enterprise wide, centrally coordinated, scalable and resilient in a way that empowers all users to make good security and privacy decisions, measured by a 10% increase in consumption of consultation services and a 30% decrease in operational services

Baseline: To be defined in Q1

Year Goal	Q1 Status	Q2 Status	Q3 Status	Q4 Status
One new data center in EMEA Expand main data center by >20% Two new capacity plans	Contracts signed and orders placed Awaiting equipment			
20% incident engagement improvement	Scorecard defined			
10% increase in consultation services 30% decrease in operational services	Set baselines for services			

Culture, Equity & Team Practices



Objective:

Culture, Equity and Team Practices: Modern and Inclusive Practices. Modern software practices and inclusive communication practices will remove or significantly reduce barriers to entry and collaboration, and streamline deployments for all technical contributors in the wiki ecosystem, enabling a faster path to innovation.

We developed and executed on, an extensive [Discovery to Implementation](#) phase for our Effective and Responsible Communications Project involving all teams and individuals in Technology. We defined [accountabilities and responsibilities](#) for every member of the Technology department, we [created alignment](#) within the department leadership team. Additionally, we held team level discussions across the 20 teams in Technology to receive nominations for the top 3 [actionable recommendations](#).

We discussed and selected “user saves wiki edit” as an important, impactful user workflow to focus on this year in defining SLOs accurately representing the user experience.

Target quarter for completion: Q4 FY21/22



Culture, Equity & Team Practices



Key Results

	Year Goal	Q1 Status	Q2 Status	Q3 Status	Q4 Status
<p>KR1: At least one user-journey, and all its supporting infrastructure services, are covered by SLOs, and those SLOs are used to drive engineering decisions including when to roll back deployments. For all supporting services within this slice developed at the Foundation, including MediaWiki, we measure and track lead-time as an indicator of the objective of maximizing the speed of innovation.</p> <p>Baseline: 0 user-journey SLOs, 3 infrastructure service SLOs</p>	1 user journey N infra services (with SLOs defined)	1 user journey selected 4 infra svc SLOs defined			
<p>KR2: 100% of non-inclusive language is removed from our documentation and code repositories according to https://www.mediawiki.org/wiki/Inclusive_language_list.</p> <p>Baseline: 0</p>	100% of non-inclusive language removed	0%			
<p>KR3: Deliver 4 3 milestones towards one recommendation resulting from the Effective and Responsible Communication discovery phase (by Q4)</p> <p>Baseline: 0</p>	4 3 (see speaker notes for this change request)	0			
<p>KR4: 100% of development teams across Technology and Product departments build using the DEI framework.</p> <p>Baseline: 0%</p>	100% of teams	0%			

Service Level Objectives

Reporting period:
June 1 - August 31, 2021

Service	SLO		Performance	SLO met?
etcd	Availability	99.9% of requests are successful	99.999%	✓
	Latency	99.8% of requests under 32 ms	99.99%	✓
API Gateway	Availability	99.9% of requests proxied successfully	99.94%	✓
	Read latency	99% of read requests under 2000 ms	442 ms	✓
	Write latency	99% of write requests under 1500 ms	2328 ms	✗ <i>see notes</i>
Docker Registry	Availability	99% of requests are successful	100%	✓
	Manifest read ltcy.	95% of manifest reads under 2000 ms	497 ms	✓
	Tag read latency	95% of tag reads under 2000 ms	49 ms	✓
	Write latency	95% of manifest writes under 3000 ms	1302 ms	✓

Appendix



WIKIMEDIA
FOUNDATION