Technology Tuning Session Q1 FY20-21

WIKIMEDIA
FOUNDATION

MTP Priority Update

WIKIMEDIA FOUNDATION

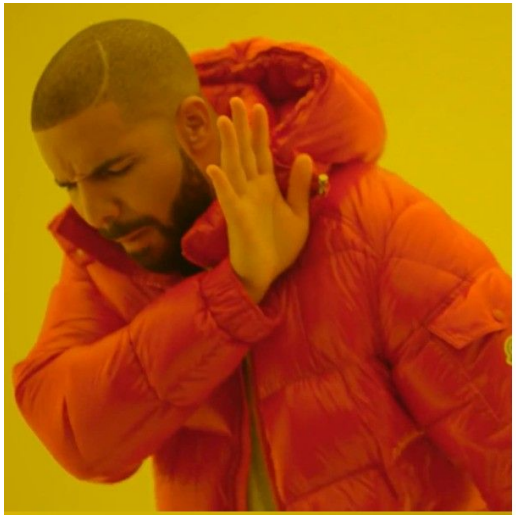

Q4: Cache

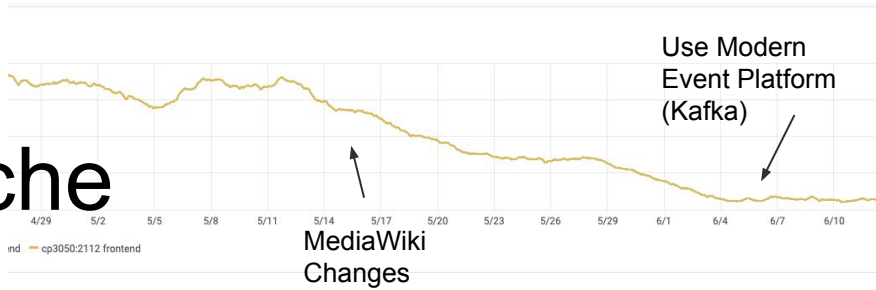

Q1: Cash

Our Data Center Operations team:

- Reduced OpEx costs by 11%, saving ~$21k/mo
- Saved $100k in CapEx

Department: Technology

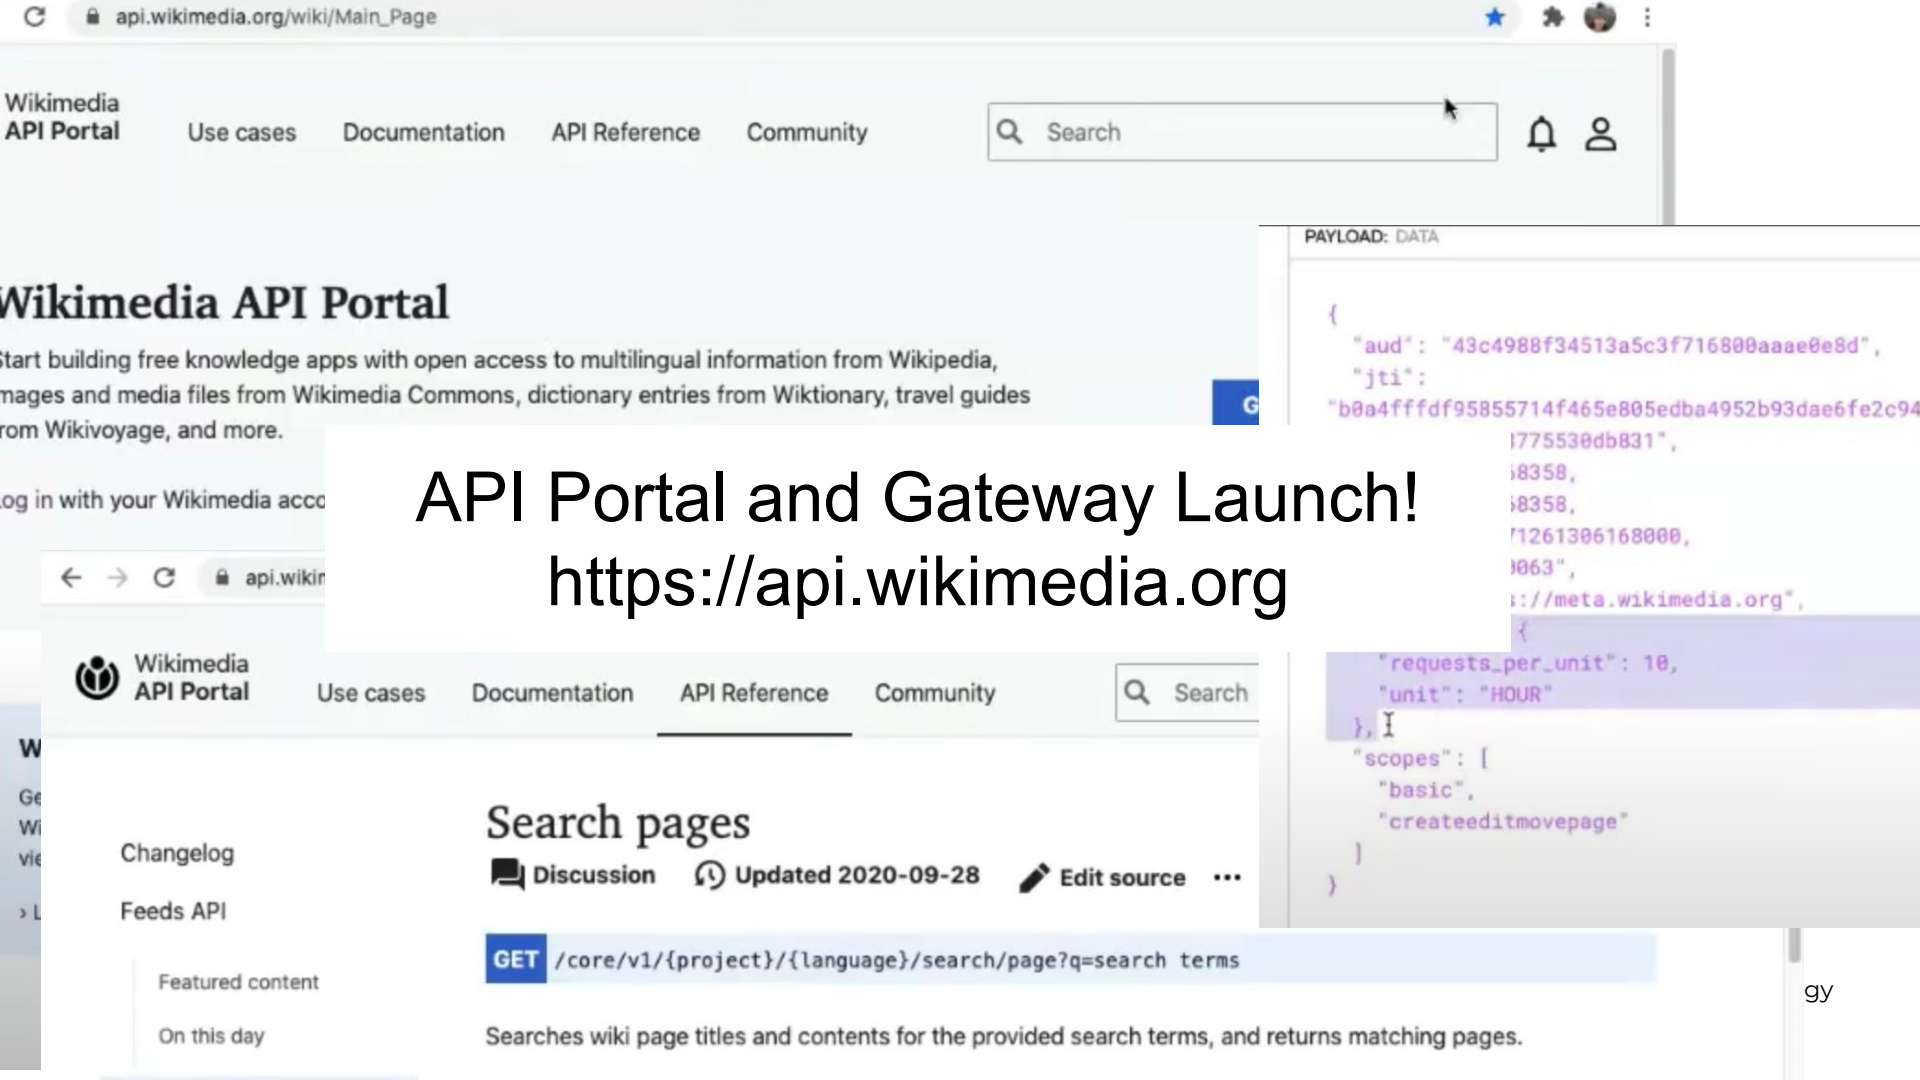
api.wikimedia.org/wiki/Main_Page
Wikimedia API Portal
Use cases
Documentation
API Reference
Community
Search
Wikimedia API Portal
Start building free knowledge apps with open access to multilingual information from Wikipedia,
mages and media files from Wikimedia Commons, dictionary entries from Wiktionary, travel guides
rom Wikivoyage, and more.
Log in with your Wikimedia acco
PAYLOAD: DATA
{
"aud": "43c4988f34513a5c3f716800aaae0e8d",
"jti":
"b0a4fffdf95855714f465e805edba4952b93dae6fe2c94
775530db831",
8358,
8358,
1261306168000,
0063",
://meta.wikimedia.org",
{
"requests_per_unit": 10,
"unit": "HOUR"
},
"scopes": [
"basic",
"createeditmovepage"
]
}
API Portal and Gateway Launch!
https://api.wikimedia.org
Wikimedia API Portal
Use cases
Documentation
API Reference
Community
Search
Changelog
Feeds API
Search pages
Discussion
Updated 2020-09-28
Edit source
GET /core/v1/{project}/{language}/search/page?q=search terms
Featured content
On this day
Searches wiki page titles and contents for the provided search terms, and returns matching pages.

Platform Evolution

Overview

We are continuing to evolve our architecture by introducing more structured data and leveraging machine learning — this is unlocking new capabilities in our products to support our communities. By improving our code quality, automation and developer tooling, we are increasing our capacity to innovate, experiment, learn, and deliver. Supporting our diverse technical communities is resulting in growth and enables solutions that engage with our content and data, and also leverages Wikimedia data beyond the core wiki experience. In support of Content Integrity, we continue to build and expand our tools and knowledge.

Progress and Challenges

The new technical decision making model is making good progress and will have a pilot project soon. We've also defined growth rates for our metrics on structured data and media usage and the creation of models for identifying misinformation is well underway.

OKRs

Content Integrity	
Evolutionary Architecture	

Actions

- Pilot new technical decision making process
- Hire the Director of Disinformation (Position opened at the start of Q2) and Research Scientist roles.
- Build the Machine Learning Platform

Department: Technology

Platform Evolution

MTP Outcomes

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

MTP Metrics	Y2 Goal	Q1 Status	Q2 Status	Q3 Status	Q4 Status
A 80% increase in structured data used (uptake) across wikis. **Baseline:** 87M of pages across Wikimedia projects use Wikidata as of April 2020.	80% increase	16% increase in percentage of use from 4/20 baseline	-	-	-
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.	10% increase	0.7% increase in percentage of use from 4/20 baseline	-	-	-

Department: Technology

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	Y2 Goal	Q1 Status	Q2 Status	Q3 Status	Q4 Status
30% increase of tool maintainers **Baseline:** 1880 maintainers in Q2 (FY 19/20)	15% increase from baseline	8% (2033) increase from baseline			
10% (4.2 / 5) increase in developer satisfaction **Baseline:** 2019 developer satisfaction: 3.8 / 5	4% (4.0)	Next survey will be conducted in early Q3			
10% decrease in code review time **Baseline:** 19 days in June 2019	4% (18 days)	-5% (20 days)			
20% decrease in outstanding code reviews **Baseline:** 1088 code reviews in June 2019	4% (1043 reviews)	15% (918 reviews)			

MTP Priority OKRs

WIKIMEDIA
FOUNDATION

Content Integrity

Objective:

The Wikimedia movement has the tools and knowledge to identify and respond to abuse, misinformation, and disinformation campaigns, in order to more efficiently and effectively improve the quality of content, defend the projects and retain the public's trust.

We build the **technologies** that empower the editor and patroller communities enforce content policies such as *verifiability*.

We build and sustain **relationships** with academic and industry partners to more effectively address disinformation.

We coordinate and report on existing and future **disinformation projects** within the Wikimedia Foundation.

Target quarter for completion: Q4 FY20-21

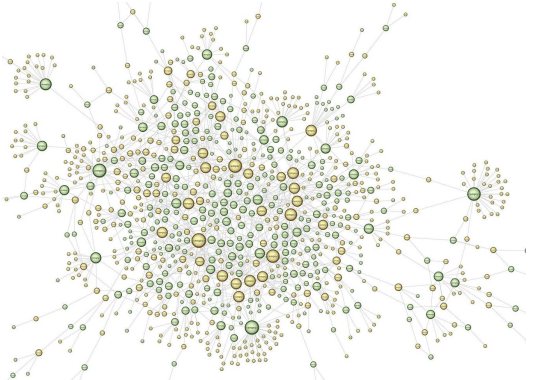

The OrangeMoody network of socks (ongoing)

Department: Technology

OKR Change: Content Integrity

Original

KR2: Complete 4 product integrations (e.g. link recommendations) of internally built machine learning models and services for improving content integrity and deploy 7 new community-driven models that extend our existing content integrity approach.

Proposed

KR2: Complete 4 product integrations (e.g. link recommendations) of internally built machine learning models and services for improving content integrity.

Department:
Advancement

Content Integrity

Key Results	Year Goal	Q1 Status	Q2 Status	Q3 Status	Q4 Status
KR1: Develop 1 model to identify misinformation (by Q3) and deliver 4 milestones towards 1 model to understand the diffusion of content on the Wikimedia projects (by Q4). **Baseline:** 0 milestones	4	1	-	-	-
KR2: Complete 4 product integrations (e.g. link recommendations) of internally built machine learning models and services for improving content integrity **Baseline:** 0 models in the new ML platform	4	0	-	-	-
KR3: Coordinate and report quarterly on the disinformation activities across teams, establish a community support forum (by Q3), and participate in 10 external relationships (by Q4) including academia and industry partners. **Baseline:** 0	10	5	-	-	-

Department: Technology

Evolutionary Architecture

Objective:

A new Wikimedia Knowledge Platform is defined, building upon our key technologies, enabling our development teams to deliver value independently, and empowering our communities to share the world's knowledge in the spaces and formats which reflect their values and cultures.

Work on predictably structuring content and adding semantic meaning (metadata) to that content

- Completed initial development of canonical data model
- Completed initial import of 1 wiki into Structured Content Store

Maturing our architecture practice

- Draft of information architecture created
- Architecture decision making process aligned with (new) Technical Decision Making Process
- Modeling with Abstract Wikipedia Team

Completion of API Gateway

- This is crucial foundational technology that will enable streamlined decoupling of services.

Target quarter for completion: Q4 FY20-21

Department: Technology

Evolutionary Architecture

Key Results	Goal	Q1 Status	Q2 Status	Q3 Status	Q4 Status
KR1: Reached level 3 of the Architecture Maturity Model (AMM) for at least 85% of capabilities by Q4. **Baseline:** 4 capabilities at Level 1 (50%), 4 capabilities at Level 2 (50%)	move 85% of capabilities to Level 3	37.5% at Level 1 62.5% at Level 2	-	-	-
KR2: Engineers have an improved understanding of current architecture and target architecture, enabled by the outputs of reaching level 3 for the AMM Architecture Development capability, "Gap analysis, modeling and iterative strategy for reaching a target architecture are completed", measured quantitatively through surveys in Q1, Q2 and Q4. **Baseline:** TBD in Q2	Baseline TBD	Baseline survey pushed to 3rd week of October	-	-	-
KR3: 4 new service components in production delivering new capabilities, enabling and adopting industry standard best practices for architecture, engineering and deployment, allowing for the mitigation of risks for both development teams and operational stakeholders, building trust in our development processes. 1 service in be end of Q2, 2 by end of Q3, and 4 by end of Q4 **Baseline:** Zero services	4 services decoupled	50% (2 services)	-	-	-

Department: Technology

Drill Down: Evolutionary Architecture

The situation

A baseline survey of developer understanding of current and future architecture has not been performed yet.

The impact

Baseline data is delayed.

The recommendation

Collect baseline data at the beginning of Q2 and then resurvey in Q3 and Q4.

Evolutionary Architecture

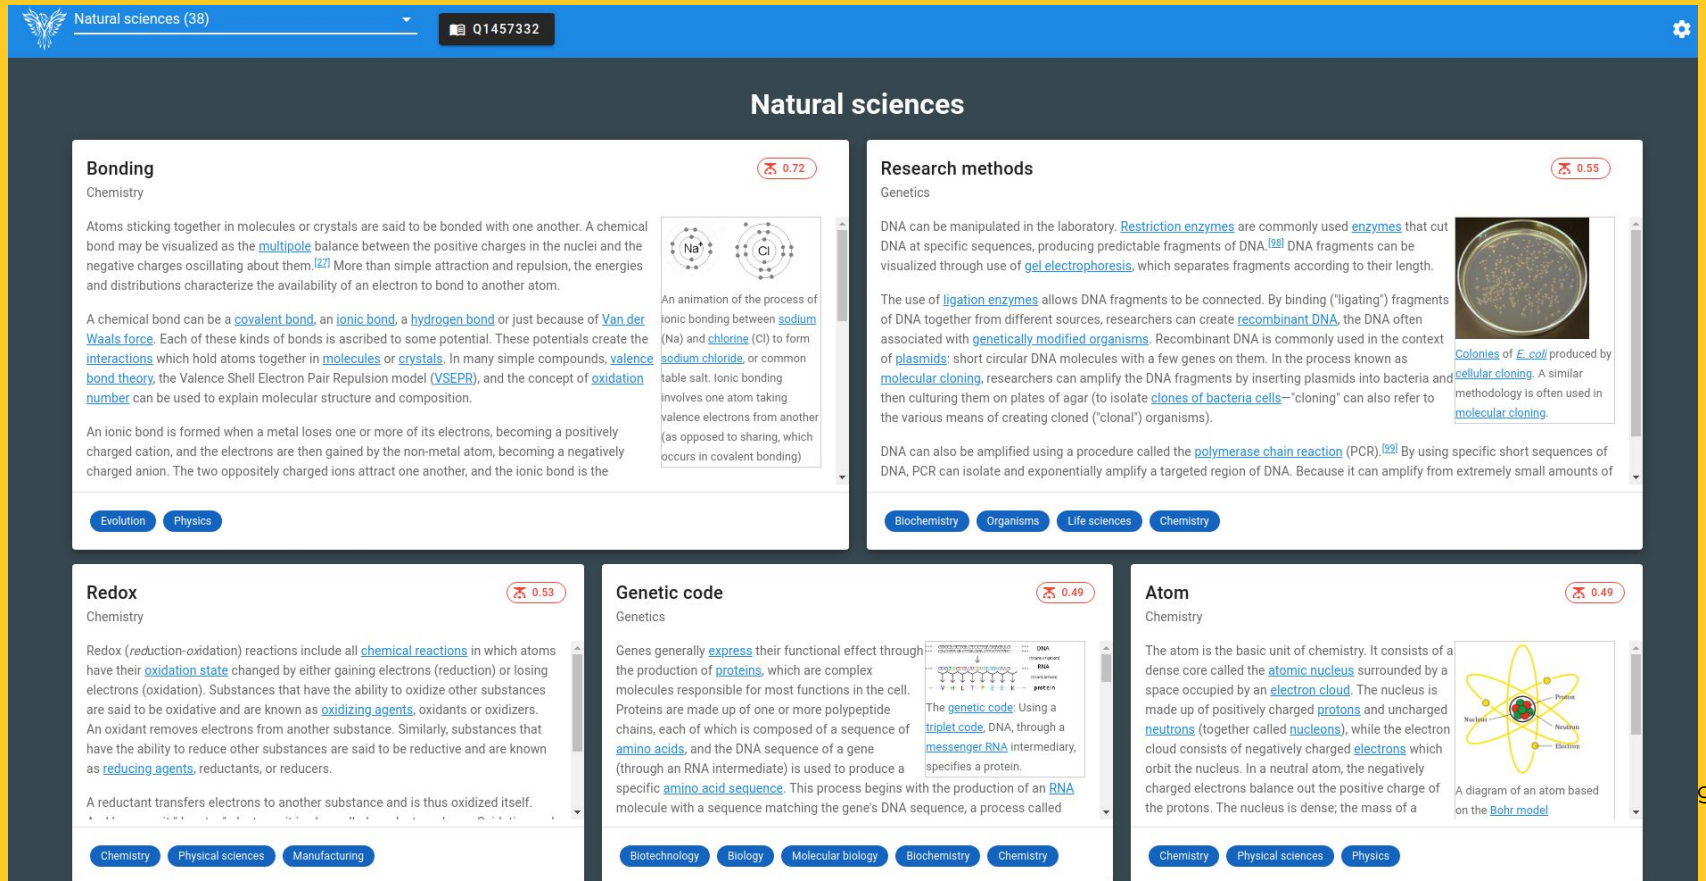

Efficacy & Resilience OKRs

Front Line Defenses

Objective:

Our infrastructure and data are staffed, secured and provisioned appropriately in each area to successfully prevent or handle malicious attacks, the unavailability of one system component, or the unavailability of a staff member.

We are actively interviewing candidates for one of the Single Point of Failure (SPoF) positions (in SRE Service Ops), and are about to kick off two additional hiring processes for two other SPoF positions (in Data Persistence and in Traffic) shortly. The unexpected setback of one of our completed SRE team manager hirings falling through has caused us to reorder our original hiring schedule, to better optimize our available hiring and onboarding bandwidth and support and maximize our ability to hire successfully and on time.

We will start early with the initial stages of planning our second EU data center (KR2), ahead of schedule, in Q2.

We are maturing our ability to identify threats and incorporate the concept of risk into our cybersecurity practices.

Target quarter for completion: Q1 FY21-22

Front Line Defenses

Key Results	Goal	Q1 Status	Q2 Status	Q3 Status	Q4 Status
KR1: We will plan (by end of Q3), build (Q4) a second EU data center, to better serve EMEA users (by end of FY21-22 Q1). **Baseline:** one EU data center	Two EU data centers	(planning not started yet)	-	-	-
KR2: The number of staff members solely responsible for a system or piece of software in production is reduced by 50%, from 8 to 4, by end of Q4. **Baseline:** 0 SPoF positions filled	Hire for 4 SPoF positions	0% filled 1 hiring process in final stages	-	-	-
KR3: Implement a threat identification and risk treatment program (Q1) and deploy effective countermeasures for the top 2 threats and risks per quarter thereafter. **Baseline:** Partial risk management policy and 0 threat intelligence feeds.	6 threat and risk countermeasures deployed	Completed Q1 items			

Department: Technology

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.

We have started with our work on the Toolhub catalogue project and are exploring new ways to foster collaboration, including setting up an advisory council of staff and community members, weekly progress reports, and a decision record.

PAWS Kubernetes rebuild: We made significant improvements to our Jupyter notebooks deployment PAWS. PAWS enables users with little programming skills to perform automated one time tasks on the Wikimedia projects, or to analyse and visualise data.

PAWS

PAWS logo by Barquill0, cc-by sa 3.0

Building upon lessons learned from the first year of Small Wiki Toolkits, we are planning (virtual) regular regional technical workshops, starting with the South Asian region in collaboration with Indic TechCom.

Target quarter for completion: Q4 FY 21/22

Department: Technology

Technical Community Building

Key Results	Year Goal	Q1 Status	Q2 Status	Q3 Status	Q4 Status
KR1: Communities find the tools that they need through the new Toolhub catalog system (technical plan Q1, working prototype Q2, MVP Q4), and developers create more diverse solutions using a user-extensible, container-based default deployment process (working prototype Q3, running at least 1 workflow by Q4) in Toolforge. **Baseline:** n/a	Toolhub MVP, 1 workflow in new deployment process	Q1 goal complete	-	-	-
KR2: A 10% increase in number of tool maintainers (4% by end of Q2) speaks to a vibrant ecosystem of technical contributors, supported by an iterative model and practice of community and capacity building (Q1 draft, Q4 MVP) which has been refined through 3 initiatives in focused outreach to technical communities (Q1-Q4). **Baseline:** 1974 tool maintainers	10% increase in number of tool maintainers, MVP model of community and capacity building	3% (2033), draft concept	-	-	-
KR3: An organization strategy for key technical documents informs a continuing roadmap based on a process of drafts (Q1) refined by consultation with staff (Q2) and community input (Q3) and a prototype of a single entry point to lower barriers finding existing documentation (Q4). **Baseline:** n/a	Prototype for a single entry point for technical documentation	Q1 goals complete	-	-	-

Department: Technology

Production Quality

Objective:

Engineering teams at WMF have a shared understanding of development velocity, production health, and code quality, and they develop commitments and workflows for improving overall code health.

The purpose of this Objective is to streamline and strengthen our development, deployment, and hosting practices to improve the daily lives of our technical community.

We're proud of the significant progress made by several teams this past quarter — to organize SLO *(service level objective)* and Error Budget definition;, organize a community-wide consultation, and build the foundation for MediaWiki on Kubernetes — there are many challenges ahead.

Target quarter for completion: Q4 FY20-21

Department: Technology

OKR Change: Production Quality

Original

KR3: Evolve and replace our Continuous Integration and Review systems with 1 project moved by end of Q2 and 5 additional projects moved by Q3. In parallel, 70% of Wikimedia developed application layer production service traffic is served by the Deployment Pipeline in Q3 and 95% by Q4.)

Proposed

KR3: Evaluate and shepherd a consultation and best practices recommendations on a potential move to GitLab for code review by the end of Q2. In parallel, 70% of Wikimedia developed application layer production service traffic is served by the Deployment Pipeline in Q3 and 95% by Q4.

Department:
Advancement

Production Quality

Key Results	Goal	Q1 Status	Q2 Status	Q3 Status	Q4 Status
KR1: Evangelize, implement tooling for (by end of Q1) and define and implement Service Level Objectives (SLO) and Error Budgets for our top 10 services and systems (Q2: 2 services; Q3: 4 services; Q4: 4 services) and report and iterate on them with product owners on a quarterly basis to optimize the balance of speed of innovation and reliability. **Baseline:** 0 services with SLOs & Error Budgets defined	10 services with SLOs and error budgets	0 services Developed SLO definition worksheet, Technical SLO baseline work completed	-	-	-
KR2: Educate teams on and define Code Health Objectives for our production-deployed code repositories (20% in Q1, 40% in Q2, 70% in Q3, 100% in Q4) and report and iterate on them with code owners on a quarterly basis. **Baseline:** 0% coverage of CHOs *(Code Health Objectives)*	100% coverage	17% of deployed repos have CHOs identified	-	-	-
KR3: Evaluate and shepherd a consultation and best practices recommendations on a potential move to GitLab for code review by the end of Q2. In parallel, 70% of Wikimedia developed application layer production service traffic is served by the Deployment Pipeline in Q3 and 95% by Q4 **Baseline:** 27k request per second (rqs) *(baseline was measured in FY1920Q4 - eqiad usage)*	Code Review decision, 95% of traffic	GitLab consultation in progress and on-track; 54% of traffic			

Drill Down: Production Quality

The situation

The MediaWiki on Kubernetes work is still in the very early stages and we are discovering new challenges to address.

The impact

The cross-team collaboration is working hard to find those challenges together early.

The recommendation

We need continued time and attention from all teams involved. Having a unified team is under discussion to help provide that momentum.

Department Update

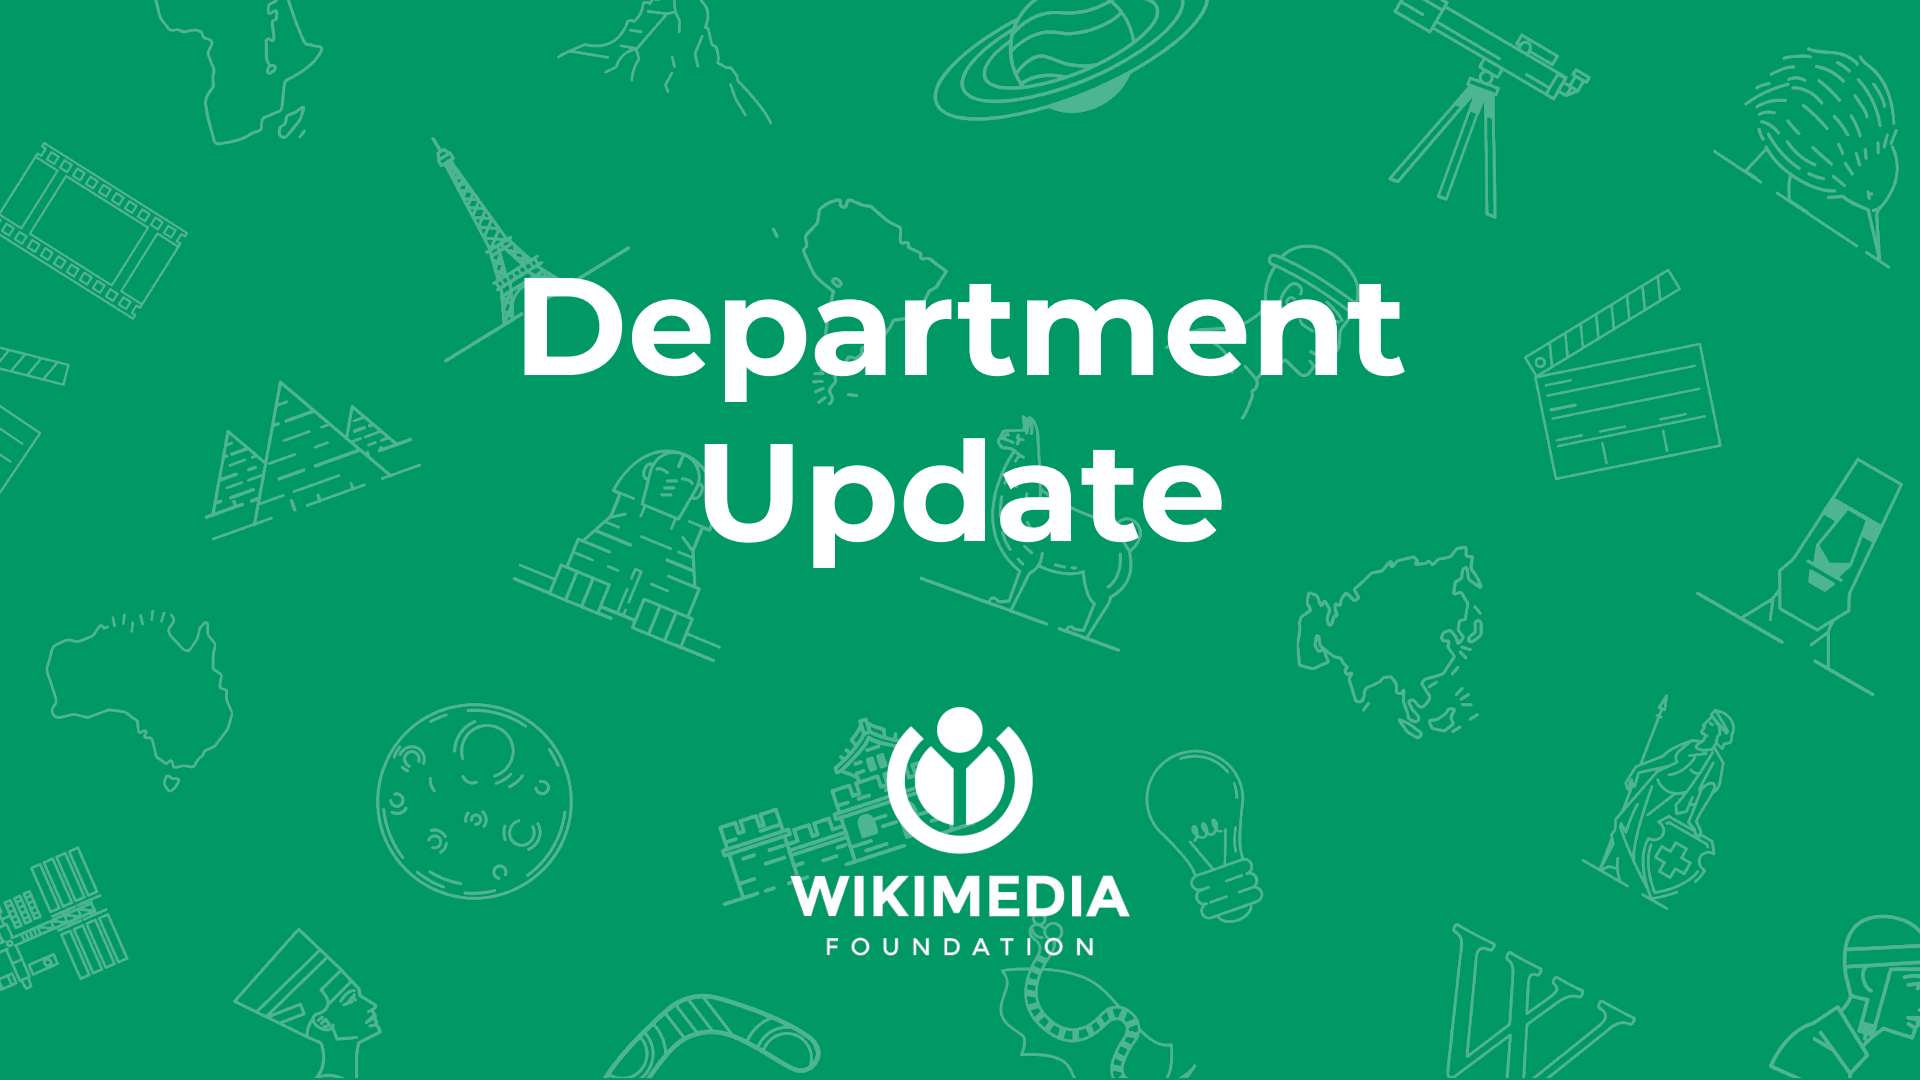

Technology Department

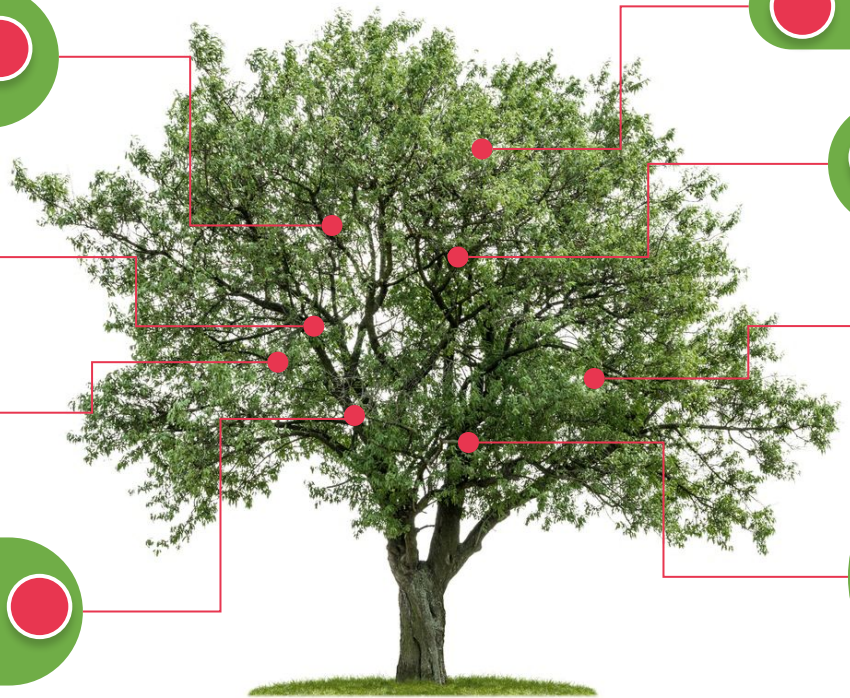

Engineering Productivity
Performance
Quality & Test Engineering
Release Engineering

Technical Engagement
Developer Advocacy
Wikimedia Cloud Services

Fundraising Tech

Data Science & Engineering
Analytics Engineering
Search Platform
Machine Learning Platform

Security

Platform
Architecture
Platform Engineering

Research

SRE
Data Center Operations
Traffic
Observability
Data Persistence
Infrastructure Foundations
Service Operations

New Hires

SRE

Lukasz Sobanski

Analytics Engineering

Razzi Abuissa

Scoring

Tobias Klausmann

Architecture

Diana Montation
Moriel Schottlender *(internal transfer)*

Anniversaries *(July - September)*

14 years
Mark Bergsma

9 years
Chris Johnson
Aaron Schulz
Jeff Green

8 years
Dan Andreescu
Dan Zahn
Andrew Bogott

7 years
Bryan Davis

6 years
Marcel Ruiz Forns

5 years
Peter Hedenskog
Trey Jones

4 years
Manuel Arostegui
Riccardo Coccioli

3 years
Miriam Redi
Diego Saez-Trumper
Cindy Cicalese

2 years
Daniel Kinzler
Scott Bassett
Effie Mouseli
Cole White
Bill Pirkle

1 year
Martin Gerlach
Kevin Bazira
Dallas Wisehaupt
Alex Paskulin
Hieu Pham
Sukhbir Singh
Grant Ingersoll

What we've done and what we will do

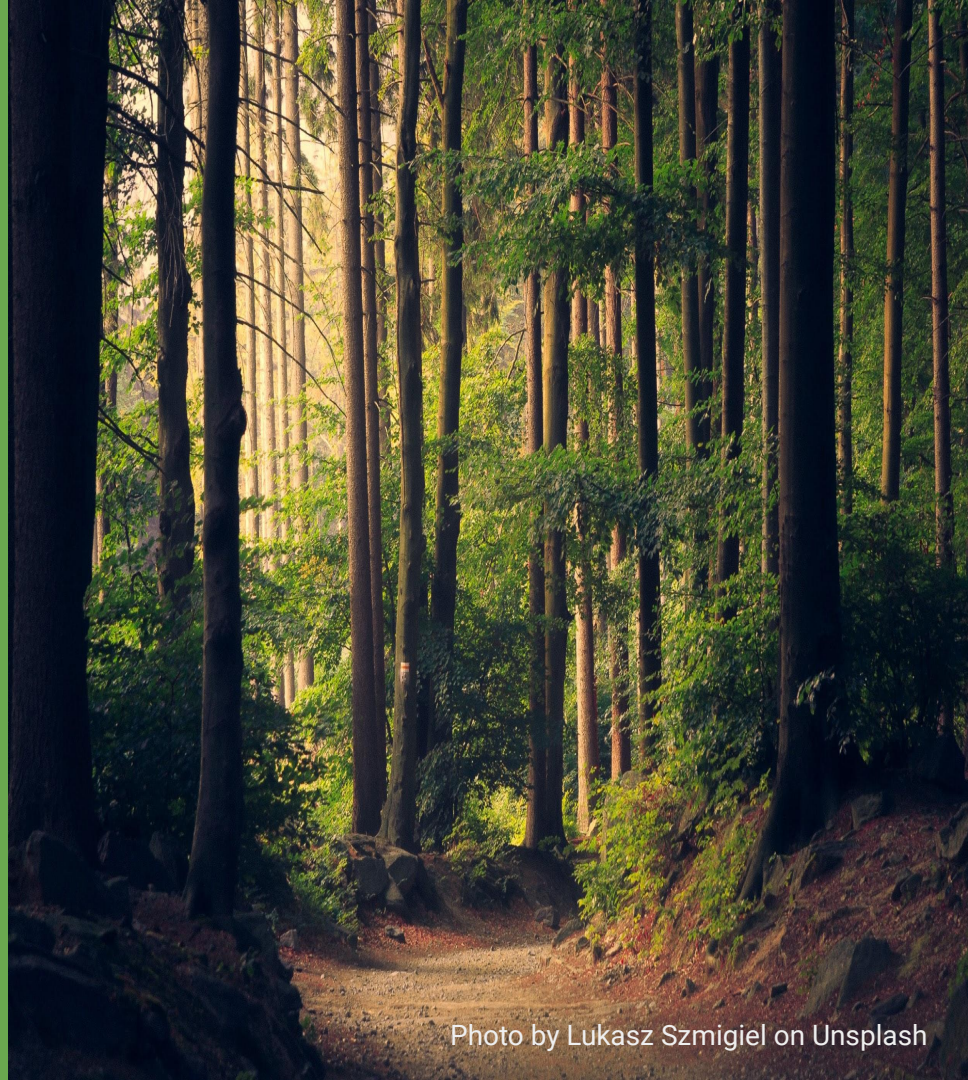

Photo by Lukasz Szmigiel on Unsplash

Department OKR Status *Q1*

Objectives Grouped by Risk

Progress Over Time for Objectives in this View

Categories

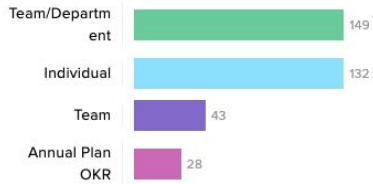

Objectives Contributing to Top Company Objectives

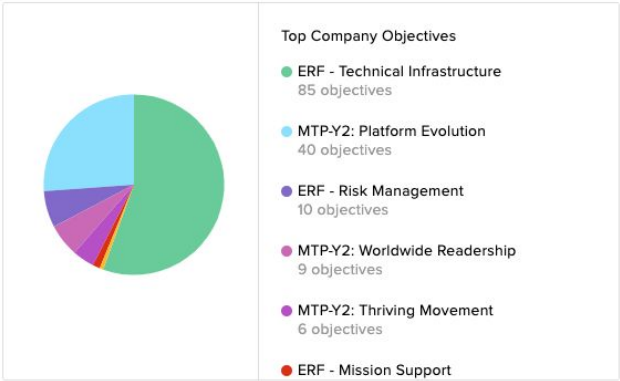

Supporting wins

Service Ops: Data center switchover was uneventful and quick

Data Persistence: new db testing environment for Puppet

Data Center Ops: Reduced OpEx costs of 11% (~$21k/month) by locking in vendor renewals

Platform Engineering: API Portal MVP launched, API Gateway & rate limiter in production

Security: Using CAPTCHA without decreasing security or privacy

Performance: Building consensus on active-active MediaWiki project *("Multi-DC")*

RelEng: Gitlab community consultation to ensure a more equitable review system

FR-Tech: Supported the most complex campaign *ever* while exceeding goals

Analytics: Enabled even more complex queries on Wikistats; added Canary Events

Search: WDQS streaming updater updates, working collaboratively to do task estimations

Architecture: First full wiki import into Structured Content Store; became a team!

Traffic: Eliminated 40+ backported Varnish patches while increasing privacy & security

Infra Foundations: Automation & single sources of truth == less manual toil, fewer mistakes

WMCS: Ceph rolled out - migrated over 150 virtual machines *(VM's)*

DevAdvocacy: Collaboration and communication pays off

Challenges *(aka our new normal)*

- Delays in shipping hardware due to supply chain challenges requires increased planning cycles, injuries to data center personnel
- Capacity for daily work still stretched thin on some teams due to global pandemic, fires, 2020 in general
- Software always has bugs 😑
- Update our thinking on how we do technical decision making
- Mobile device lab: trial phase vs production caused various breakages with vendor (Kobiton)
- Deployment of machine learning models that require one-off solutions

WIKIMEDIA FOUNDATION

Upcoming

- Cross-team planning for Mediawiki migration to Kubernetes *(RelEng, PET, Service Ops)*
- Eliminating blockers and expanding services by paying down technical debt
- Get better at dealing with ad hoc requests and not planning too much
- Announce Gitlab consultation decision
- Figuring out how to implement CAPTCHA while maintaining security and privacy
- Toolhub working prototype up and running
- Centralizing data collection and analysis *(sandbox for Data³)*
- Finalize artifacts for structured content store
- Updating Hadoop to be fully open sourced *(Big Top)*
- Get new Flink-based WDQS streaming updater productionized
- Defining our Department Vision and Values
- Keeping on blogging!
- Build a modern machine learning training, management and deployment infrastructure

Questions