

Wikipedia: Open Infrastructure for Open Knowledge

Jaime Crespo (@jynus)
esLibre 2023 Zaragoza - 12 May

Introduction

How it all started?

- **Jan 2001:** UseModWiki (Perl + text files)
- **Jan 2002:** "The PHP script" (PHP + MySQL 3.23). Hosted on a Bomis server
- **July 2002:** "Phase III" (MediaWiki)
 - Apache 1.3 + PHP 4.2 + MySQL 3.23. Wikipedia's first own server
- **May 2003:** Separate web server for enwiki + database server/web server for other wikis
- **Sept 2003:** Two front-end web servers, each with its own cache (file-based). About 8-11 page req/s.
- **Jan 2004:** Tampa datacenter. Primary and replica dbs + dumps, en + foundation and other webservers, misc server (mail list). Files locally.
- **Feb 2004:** 3 squid servers over 2 DCs, file cache turned off. 5 Apaches using NFS
- **2005:** 40 servers
- **2006:** 120 servers. Tampa + Amsterdam + Seoul
- **2009:** 350 servers
- **2013:** New primary DC in Ashburn, Virginia
- **2014:** New secondary DCs in Dallas, Texas and San Francisco (4 in total, including Amsterdam)
- **2016:** ~1200 bare metal servers. First datacenter failover test.
- **2018:** New facility in Singapore, 5 colocation facilities
- **2020:** ~1300 bare metal servers
- **2022:** Multi-dc app servers




Monthly overview

Reading

Total page views

24B

April ↓ -9.05% month over month



281B ↑ 7.57% year over year

Last 12 Months (May 2022 - Apr 2023)

Page views by country

Countries with the most views for April

3B United States of A

918M Japan

863M Germany


863M United Kingdom

Contributing

Edits

47M

April ↓ -3.86% month over month




560M ↑ 2.00% year over year

Last 12 Months (May 2022 - Apr 2023)

New registered users

258K


April ↓ -10.47% month over month



3M ↑ 4.95% year over year

Last 12 Months (May 2022 - Apr 2023)

Some stats



100-170K [HTTP\(S\) requests/s](#)



92 million [multimedia files](#)




970 [wikis](#)



~320 [languages](#)

2057 volumes

14 stacks



IT Philosophy

- Only using free software (Debian)
- Self-hosted (for privacy)
- Managed by volunteers and staff
- Transparency:
 - Open code ([Gerrit](#)/[GitLab](#))
 - Open infrastructure ([configuration management](#) and orchestration)
 - Open bug tracker ([Phabricator](#))
 - Open metrics ([Grafana](#))
 - Open discussions ([mail lists](#) & [IRC](#))
 - Open docs ([mw.org](#) & [WikiTech](#))
 - Open data ([dumps.wm.org](#), [Wikireplicas](#))

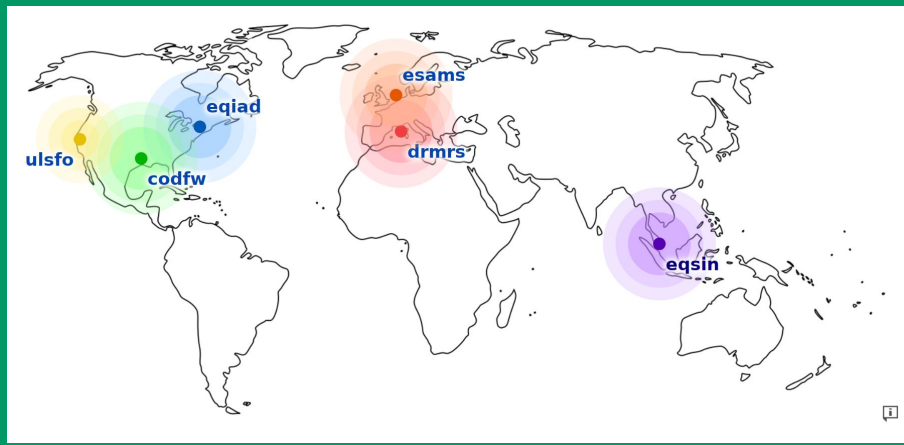


“Imagine a world in which every single person on the planet is given free access to the sum of all human knowledge.
That's what we're doing.”

How It All works?



Datacenters

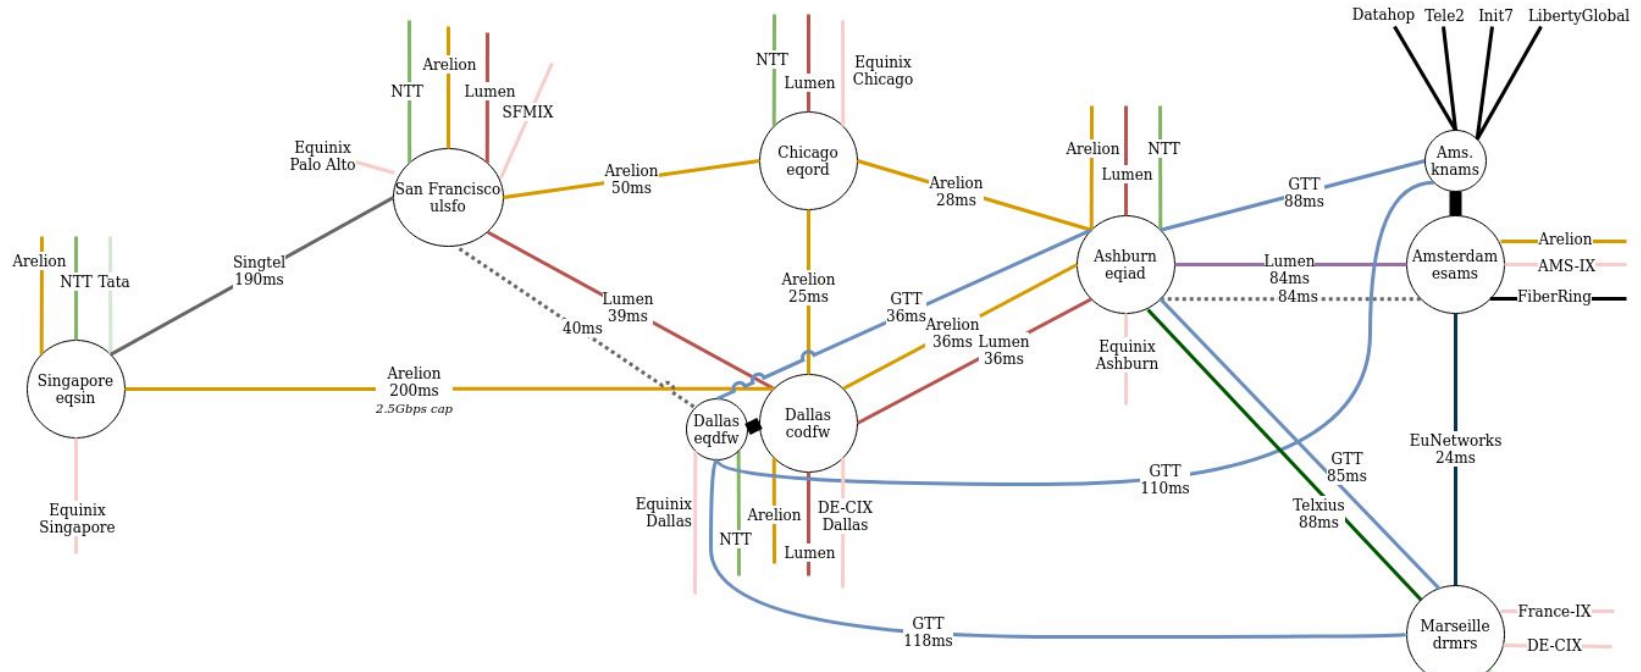


- 6 colocation facilities around the world
 - Want to expand to Africa, South America
- ~1800 bare metal servers
 - No public cloud, ext. CDNs
- gdnssd: GeoDNS to redirect you to our closest edge



Network

- **AS14907** operated with per-site and global prefixes
- IPv4 and IPv6 connectivity
- Internet transit and IXP peering at each POP
- Private wavelength services interconnect POPs

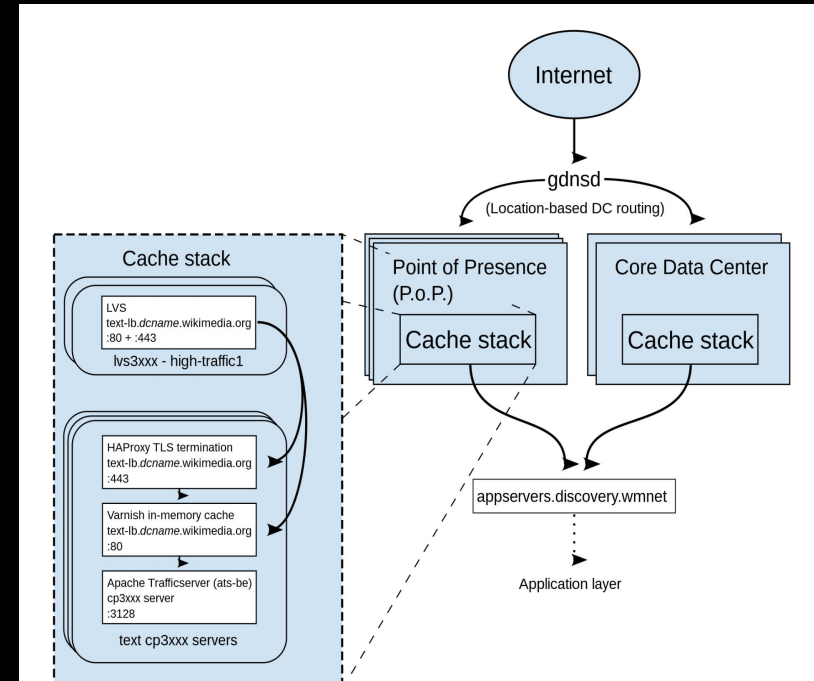
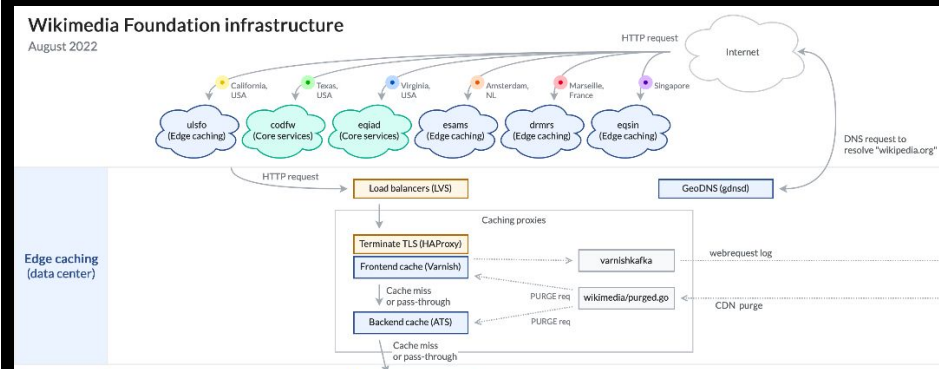


Edges & Cache

- **Linux Virtual Server** for load balancing all layers below
- **HAProxy** for TLS termination and HTTP 2
- **Varnish** for the in-memory cache ("frontend"), and
- **Apache Traffic Server** is responsible for on-disk persistent caching ("backend").

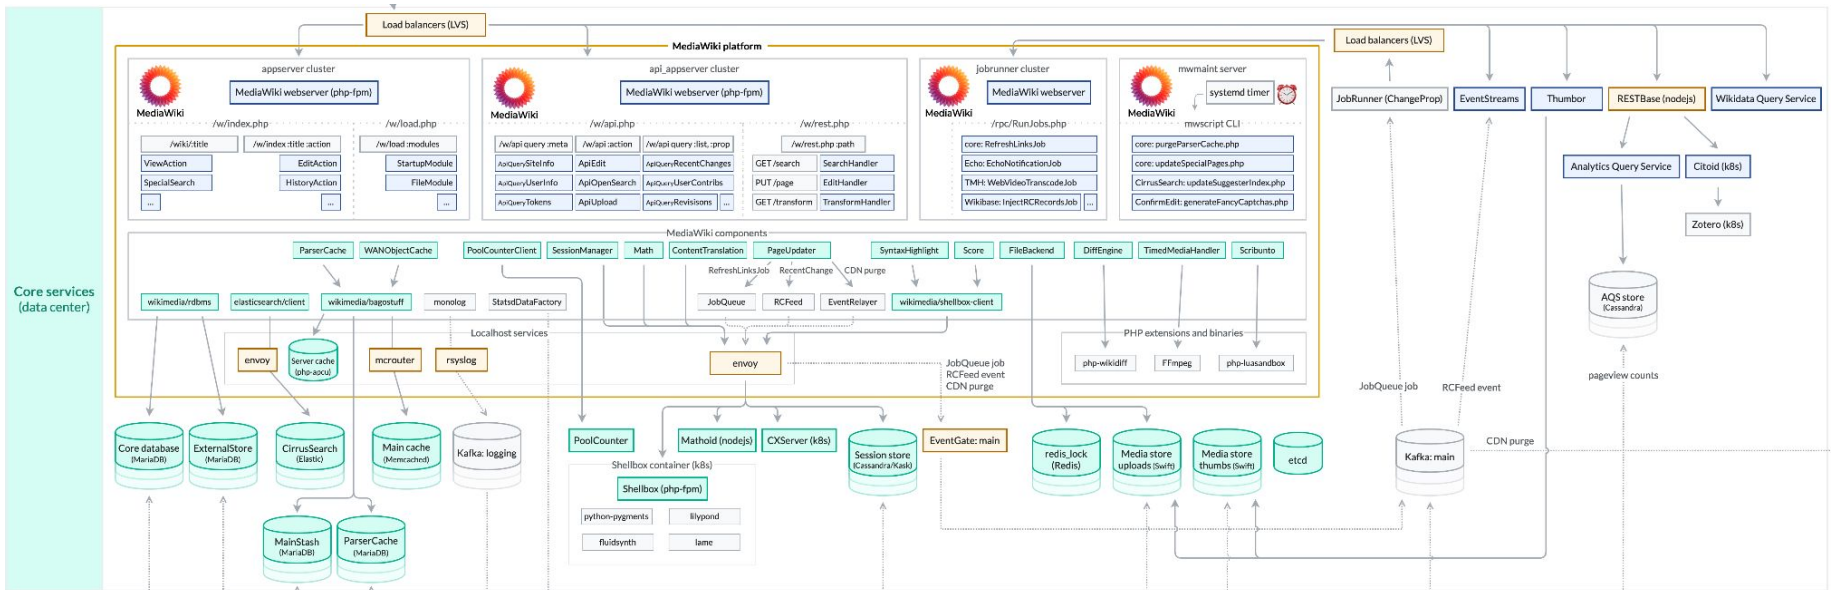
upload (media files) and text (HTML, js, CSS) separate clusters

- ~90% of traffic ends at the cache layers



App servers

- MediaWiki monolith (Apache + PHP 7.4)
- 4 (+ 1) clusters: app_servers, api_servers, parsoid, jobrunners/videoscalers (+ mwdebug)
- In migration to Kubernetes
- Multiple levels of caching (code, raw wikitext, parsed HTML, internal state) on memcache, app cache, databases



Services

- Handled in [Kubernetes](#):
 - Several media rendering services: Thumbor, PDF generation, timeline, mathoid
 - Mobile apps APIs
 - Services depending on external requests: Citoid, Translation
 - Events and streams
 - Machine learning



```
128 mw:
129 # The datacenter. To be overridden in deployments.
130 datacenter: local
131 domain_suffix: "local"
132 egress:
133   database_networks: "10.0.0.0/8"
134   etcd_servers: [] # If present, they need to have the form "{ip: <>, port: <>}"
135 httpd:
136   image_tag: latest
137   # Add here any apache configuration you might want to test/preload.
138   # It will be loaded before all the rest of conf-available and after modules have loaded.
139   # See https://github.com/wikimedia/operations-docker-images-production-images/blob/master/ima
140   additional_config: false
141 localmemcached:
142   # Enable to start memcache daemons in the pod, one for each of the
143   # ports listed. These can serve as placeholders for
144   # mcrouter to allow for image testing in non-production
145   # environments. If this section is enabled then the mcrouter
146   # sections must be disabled.
147   enabled: false
```

==> v1/Deployment

NAME	DESIRED	CURRENT	UP-TO-DATE	AVAILABLE	AGE
mathoid-staging	1	1	1	1	148d


==> v1/NetworkPolicy

NAME	POD-SELECTOR	AGE
mathoid-staging	app=mathoid,release=staging	148d

==> v1/Pod(related)

NAME	READY	STATUS	RESTARTS	AGE
mathoid-staging-6b84b94b47-wtp5m	2/2	Running	0	34d

NOTES:



Q great| Search

Great
Topics referred to by the same term

Great **Depression**
Worldwide economic depression (1929–1939)

Great **Britain**
Island northwest of continental Europe

Great **Pyramid of Giza**
Largest pyramid in the Giza Necropolis, Egypt

Great **Famine (Ireland)**

edia,
/one can edit.
glish

In the new

Supporting clusters

- Sessions
- Search
- Analytics
- Cloud Services

37.02%
Percentage of total edits made to Wikimedia wikis by Wikimedia Cloud Services hosted bots and tools.
April-June 2019

2.6 Billion
Number of API requests made to Wikimedia wikis by Wikimedia Cloud Services hosted bots and tools.
January-March 2019



177 Projects
699 Instances
2019-06-07

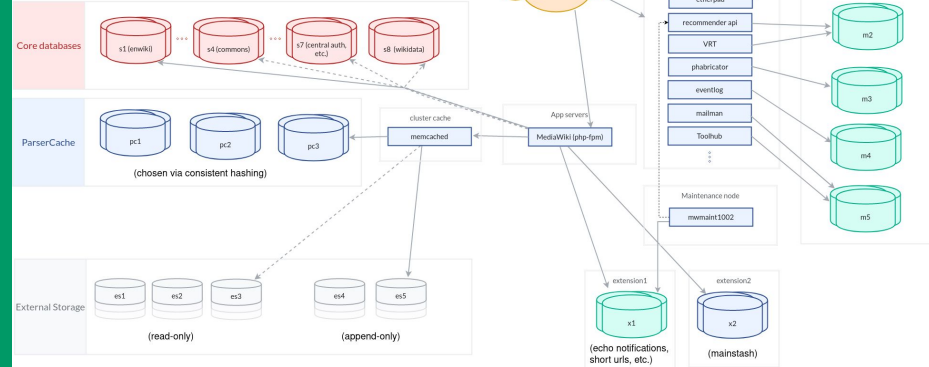
2,367 Tools
1,818 Maintainers
2019-06-07

Databases & Storage

- **MariaDB** (10.4 -> 10.6) main structured data storage (metadata and wikitext revisions)
 - Also used by many internal services (bug tracker, backups, etc)
- Openstack **Swift** for multimedia and thumbnails/transcoded video
- **Cassandra** for sessions, additional storage
- Other specialized data stores*: Hadoop, Druid, Redis, PostgresQL, Memcache, MiniIO, Etc, Kafka, Bacula



*In some cases, for a generous definition of "data stores"




User	Permissions	Logging	Tags
<ul style="list-style-type: none"> user user_id INT user_name BINARY(255) user_real_name BINARY(255) user_password BLOB(255) user_newpassword BLOB(255) user_newpasswd_expire MWTIMESTAMP user_email TEXT(255) user_touched MWTIMESTAMP user_tokens BINARY(32) user_email_authenticated MWTIMESTAMP user_email_token BINARY(32) user_registration MWTIMESTAMP user_endcount INT user_password_expires MWTIMESTAMP 	<ul style="list-style-type: none"> user_properties user_id INT user_property BINARY(255) user_value BLOB user_newtalk user_id INT user_ip BINARY(40) user_last_timestamp MWTIMESTAMP actor actor_id BIGINT actor_user INT actor_name BINARY(255) bot_passwords bp_user INT bp_key_id BINARY(32) bp_password BLOB(255) bp_token BINARY(32) bp_restrictions BLOB bp_grants BLOB 	<ul style="list-style-type: none"> logging log_id INT log_type BINARY(32) log_action BINARY(128) log_timestamp MWTIMESTAMP log_actor BIGINT log_namespace INT log_title BINARY(255) log_page INT log_params BLOB log_deleted TINYINT log_search ls_field BINARY(32) ls_value STRING(255) ls_log_id INT comment comment_id BIGINT comment_text INT comment_text BLOB comment_data BLOB 	<ul style="list-style-type: none"> change_tag ct_id INT ct_rc_id INT ct_log_id INT ct_rev_id INT ct_params BLOB ct_tag_id INT change_tag_def ctd_id INT ctd_name BINARY(255) ctd_user_defined TINYINT(1) ctd_count BIGINT
Recent changes	Pages	archive	redirect
<ul style="list-style-type: none"> recentchanges rc_id INT rc_timestamp MWTIMESTAMP rc_actor BIGINT rc_namespace INT rc_title BINARY(255) rc_comment_id BIGINT rc_minor TINYINT rc_bot TINYINT rc_new TINYINT rc_cur_id INT rc_old_id INT rc_last_old_id INT rc_type TINYINT rc_namespace BINARY(4) 	<ul style="list-style-type: none"> page page_id INT page_namespace INT page_title BINARY(255) page_restrictions BLOB(255) page_is_new TINYINT page_random FLOAT page_touched MWTIMESTAMP page_links_updated MWTIMESTAMP page_labels INT page_len INT page_content_model BINARY(32) page_lang BINARY(35) 	<ul style="list-style-type: none"> archive ar_id INT ar_namespace INT ar_title BINARY(255) ar_comment_id BIGINT ar_actor BIGINT ar_timestamp MWTIMESTAMP ar_minor_edit TINYINT ar_rev_id INT ar_deleted TINYINT ar_len INT ar_pages_id INT ar_parent_id INT ar_sha1 BINARY(32) 	<ul style="list-style-type: none"> redirect rd_id INT rd_namespace INT rd_title BINARY(255) rd_interwiki STRING(32) rd_fragment BINARY(255) category cat_id INT cat_name BINARY(255) cat_subcats INT cat_files INT


Total HTTP request volume (Varnish frontend CDN)



HTTP 5xx error responses (CDN)



MediaWiki backend response time



Monitoring


- ELK (OpenSearch) for logs
- Prometheus + Thanos + [Grafana](#) for metrics (publicly available)
- Icinga/Alertmanager for alerting
- [Server Admin Log](#) registers deploys and changes on production (wiki, IRC, twitter)
- A few specialized dashboard apps
- Statuspage: *

<https://www.wikimediastatus.net/>

*Exception to self-hosted policy

User-rep [14:11] <jinxer-wm> (KubernetesAPILatency) firing: High Kubernetes API latency (POST pods) on k8s@eqiad - <https://wikitech.wikimedia.org/wiki/Kubernetes> - <https://grafana.wikimedia.org/d/000000435?var-site=eqiad&var-cluster=k8s> - <https://alerts.wikimedia.org/?q=alertname%3DKubernetesAPILatency>

[14:16] <jinxer-wm> (KubernetesAPILatency) resolved: High Kubernetes API latency (POST pods) on k8s@eqiad - <https://wikitech.wikimedia.org/wiki/Kubernetes> - <https://grafana.wikimedia.org/d/000000435?var-site=eqiad&var-cluster=k8s> - <https://alerts.wikimedia.org/?q=alertname%3DKubernetesAPILatency>



NEL country map (if nothing loads, play with the zoom)

NELs by country code

http_request_headers.x-group-country: Descending	Count
RU	10,302
RU	4,598
US	4,219
NL	3,708
JP	2,955
DE	2,204
GB	2,170
FR	1,837
CN	1,573
IT	1,338

NELs by country over time

Deployment cycle

- Users report bugs and feature requests in [Phabricator](#)
- Developers (both WMF staff and community members) create and review patches in [Gerrit/GitLab](#)
- Testing envs for CI/CD
- Integration testing in [beta](#)
- MediaWiki branch cut weekly from HEAD (e.g. 1.41.0-wmf.6)
 - Canary [deployments](#)
 - Deployed in 3 phases: production test, small wikis & large wikis
 - Scap tool (pre-k8s)



Home

Zuul status

Jenkins

Test coverage

Zuul Status

Real-time status monitor of Zuul, the pipeline manager between Gerrit and Jenkins. [more info](#) »

Filter:

e.g. 1234 or mediawiki... [/]

Expand

by default

test

Pipeline for jobs on patch submission that will apply the "Verified" score in Gerrit. To retry these jobs, submit the comment "recheck".


Queue: mediawiki



Queue: mediawiki



Queue: mediawiki



Queue: mediawiki



Incident handling

- 2 SREs are on call and normally one becomes **Incident Coordinator**
- IC coordinates response, escalates it
- Workbooks are written for the most common debugging and fixing procedures
- Incidents inform gaps in procedures and team knowledge
- Postmortems published publicly at https://wikitech.wikimedia.org/wiki/Incident_status
- Status page: <https://www.wikimediastatus.net/> (exception to not self-hosted/open source solution)



Main page
Recent changes
Server admin log: Prod
Admin log: RelEng
Incident status
Deployments [curr]
SRE Team Help

Cloud VPS & Toolforge
Cloud VPS portal
Toolforge portal
Request VPS project
Admin log: Cloud VPS

Tools
What links here
Related changes
Special pages
Permanent link
Page information
Cite this page

Print/export
Create a book
Download as PDF
Printable version

Incidents/2023-02-22 read only

< Incidents

document status: final

Summary

Incident metadata (see [Incident Scorecard](#))

Incident ID	2023-02-22 read only	Start	2023-02-22 11:03:25 (major impact starts at 2023-02-22 12:16:21)
Task	T330300	End	2023-02-22 12:18:48
People paged	0	Responder count	~7
Coordinators	Jcrespo	Affected metrics/SLOs	?
Impact	For approximately 2 minutes, editing was disabled site-wide. For approximately 54 minutes, editing failed for some users in the codfw datacenter (around 1-2% of all edits)		

While performing a [live switchover test](#) in advance of the [2023 WMF datacenter switchover](#), an existing logical bug on the switchover test script accidentally set the secondary datacenter in read-only mode. While this didn't disrupt most users, mobile editing for people geolocated to codfw app servers (mostly, people in the Americas, and part of Asia and Oceania) had the editing interface disabled (while desktop users were redirected to edit through eqiad). While trying to fix this issue, a tooling interface issue caused all datacenters to be set in read-only mode, disabling editing for all users. This was quickly reverted for both datacenters and editing was restored.

Timeline

All times in UTC.

Contents [hide]

- 1 Summary
- 2 Timeline


Conclusion

Takeaways?




Community is key

Our community is faster at reporting ongoing issues, finding bugs and sending patches (although not every job is suitable to be left for volunteers)



Openness in everything, not only code

Not only free software is necessary, also open infrastructure, open data and open procedures.



It is not always easy!

The technological challenges are huge, and the resources are limited- it is not always easy to succeed.

We need your help!

https://www.mediawiki.org/wiki/How_to_contribute



Questions?

A large, dimly lit crowd of people is shown from a low angle, looking towards the front of a room. Many people have their hands raised in the air, suggesting an interactive session or a Q&A period. The lighting is dark, with some highlights on the people's faces and hands. The overall atmosphere is one of active participation and engagement.

Wikipedia: Open Infrastructure for Open Knowledge

Thank you for attending!

You can find me as @jynus on most services or at jcrespo@wikimedia.org

Credits:

Presentation: Jaime Crespo - License: CC-BY-SA-4.0 Wikimedia Foundation

Thanks to everyone that contributed: marostegui, godog, akosiaris, kwakuofori, _joe_, moritz, topbanks, simon, bblack

To Know More:

<https://wikitech.wikimedia.org>

<https://mediawiki.org/>

<https://meta.wikimedia.org>