WIKIMEDIA FOUNDATION

Finding the Capacity to Grieve Once More

A 15 year tale of celebrity deaths

Alexandros Kosiaris

@akosiaris@hachyderm.io

WIKIMEDIA FOUNDATION

Who's Wikimedia?

Wikipedia
The Free Encyclopedia

WIKIDATA

WIKIQUOTE

Wiktionary
The free dictionary

WIKIMEDIA

wiki
voyage

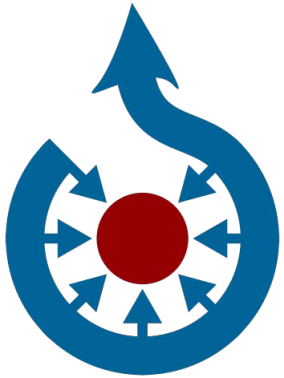

WIKIMEDIA
COMMONS

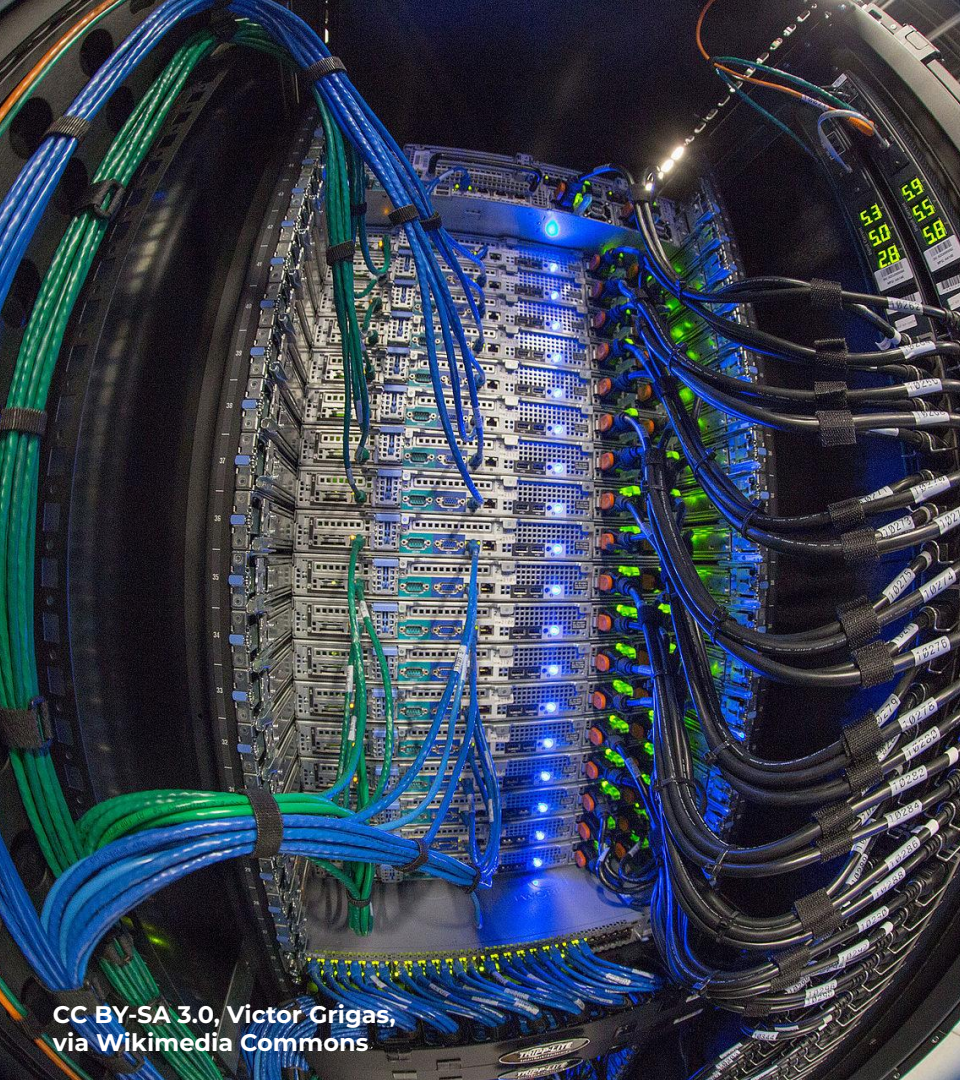

Our infra

- All our own metal

- ~2200 physical servers
- ~270 VMs

- 7 total locations
  - 2x "core" + CDN
  - 5x just CDN

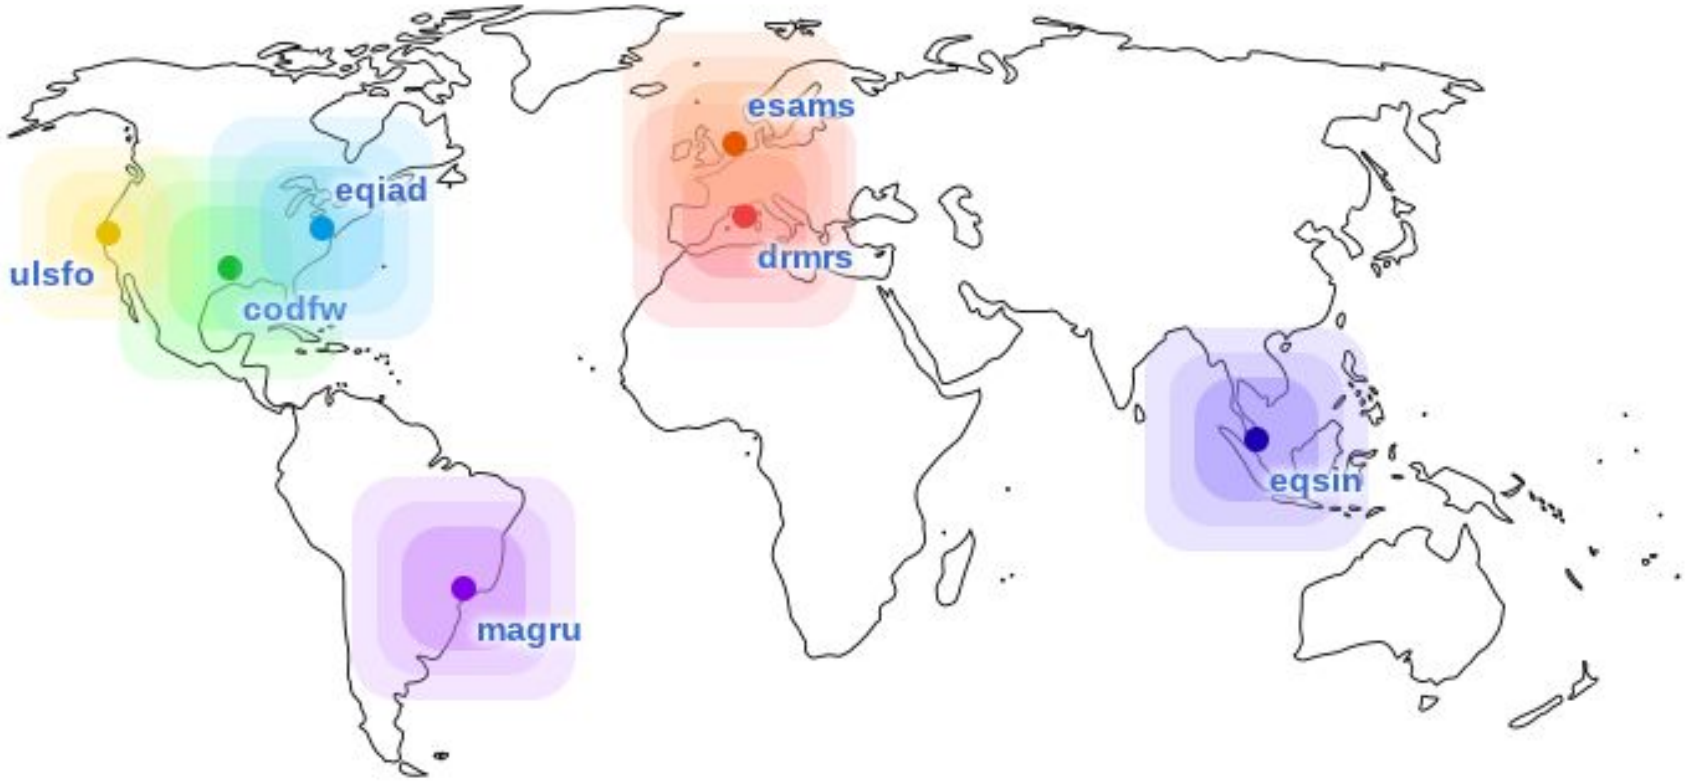

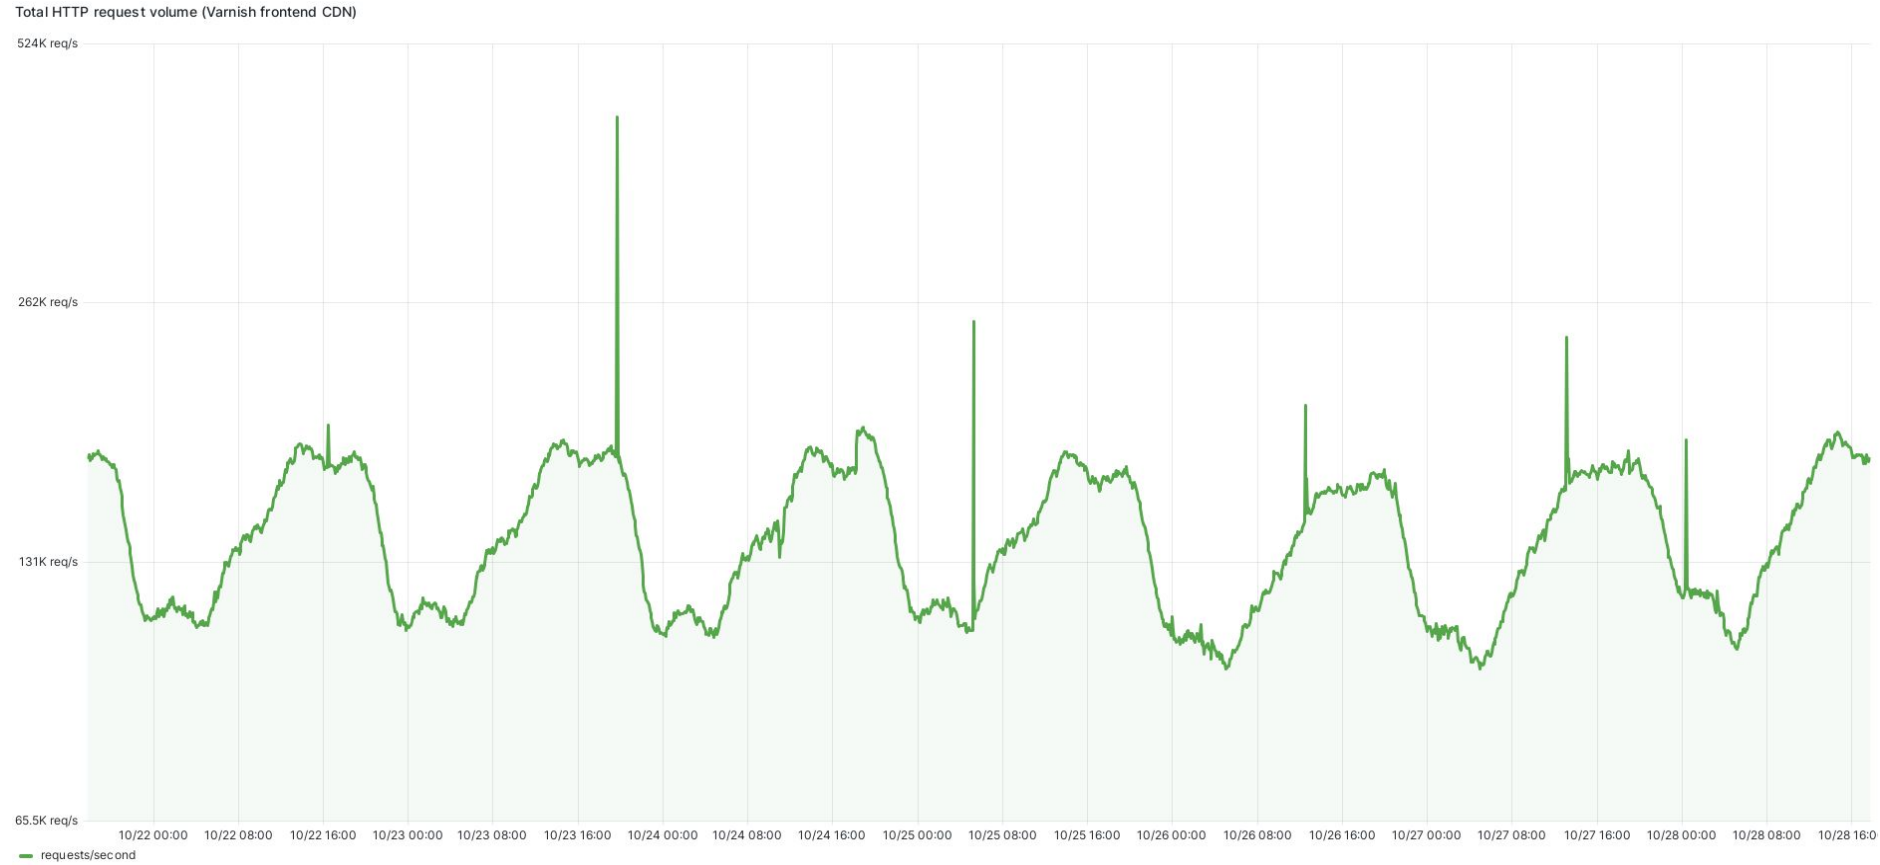

— https://grafana.wikimedia.org

WIKIMEDIA FOUNDATION

A short story

How we fared from 2009 to 2020

25 June 2009

- Wikipedia was 8 years old
- Michael Jackson departs from this world

WIKIMEDIA FOUNDATION

The Michael Jackson Effect

WIKIMEDIA FOUNDATION

Michael Jackson effect

Page Discussion Read Edit Edit source View history Tools

The **Michael Jackson effect** (also **Michael Jackson problem**[1]) is a technical term used in the Wikimedia movement to refer to a cache stampede. A cache stampede is the system failure that results when there is high demand for a computed object that is presently uncomputed.

Jackson performing in June 1988.

Event [edit | edit source]

The term was coined *[when?][by whom?]* after the death of Michael Jackson on 25 June 2009, which resulted in an unprecedented amount of page views and combined edit traffic.

The article received a record-breaking 5.9 million visits on a single day (26 June), of which one million were during a single hour.[2]

The article received 1.2 million visits on the day of Jackson's death (25 June), which with the combined edit traffic caused several server overloads that made Wikipedia intermittently unavailable to the public.[2]

Technical impact [edit | edit source]

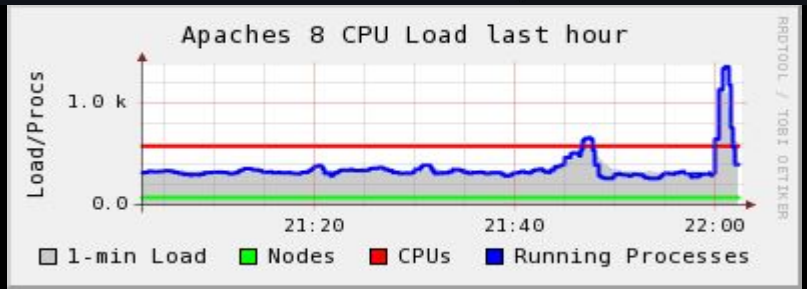

Cache Stampede

- A wikipedia article has many revisions.
- During the edit save process, wikitext is parsed.
- There are pages, even today, that take > 30 seconds to parse.
- When done, the page's entry in cache (ParserCache) is overwritten
- And the CDN cache is purged
- Edit rates skyrocketed
- View rates skyrocketed even more
- Race conditions amongst servers trying to parse MJ's page
    - At times the exact same revision

Solution: Poolcounter

- A small C daemon and the accompanying MediaWiki core code
- A distributed mutex that allows to request a lock on a given page
- Run a couple of instances, at most a couple of servers will be parsing the pages wikitext.
- Protocol is simple
  - Commands: ACQ4ANY, ACQ4ME, RELEASE
  - Replies: LOCKED, DONE, QUEUE_FULL, TIMEOUT, RELEASED

A (then) recently created article got its first major entry

[Wikipedia:Article_traffic_jumps]

WIKIMEDIA FOUNDATION

CC BY-SA 2.0, Zoran Veselinovic

Michael Jackson
2009-06-26, 1.4M views, 2009-06-27, 5.9M views

And life went on

And people visited and edited Wikipedia

WIKIMEDIA
FOUNDATION

Amy Winehouse
2011-07-23: 4.2M views. 2011-07-24: 2.3M views

And staff and volunteers recounted the tale

WIKIMEDIA
FOUNDATION

Steve Jobs
2011-10-06: 7.4M views, 2011-10-07: 1.6M views

Whitney Houston
2012-02-12: 5.9M views

Paul Walker
2013-12-01: 4.3M views

And the tale was told
sometimes in awe,
sometimes in wonder,
sometimes in fear.

WIKIMEDIA
FOUNDATION

Nelson Mandela
2013-12-06: 2.7M views

And sometimes something (small) happened

WIKIMEDIA FOUNDATION

Robin Williams
2014-08-12: 6.5M views, 2014-08-13: 1.4M views

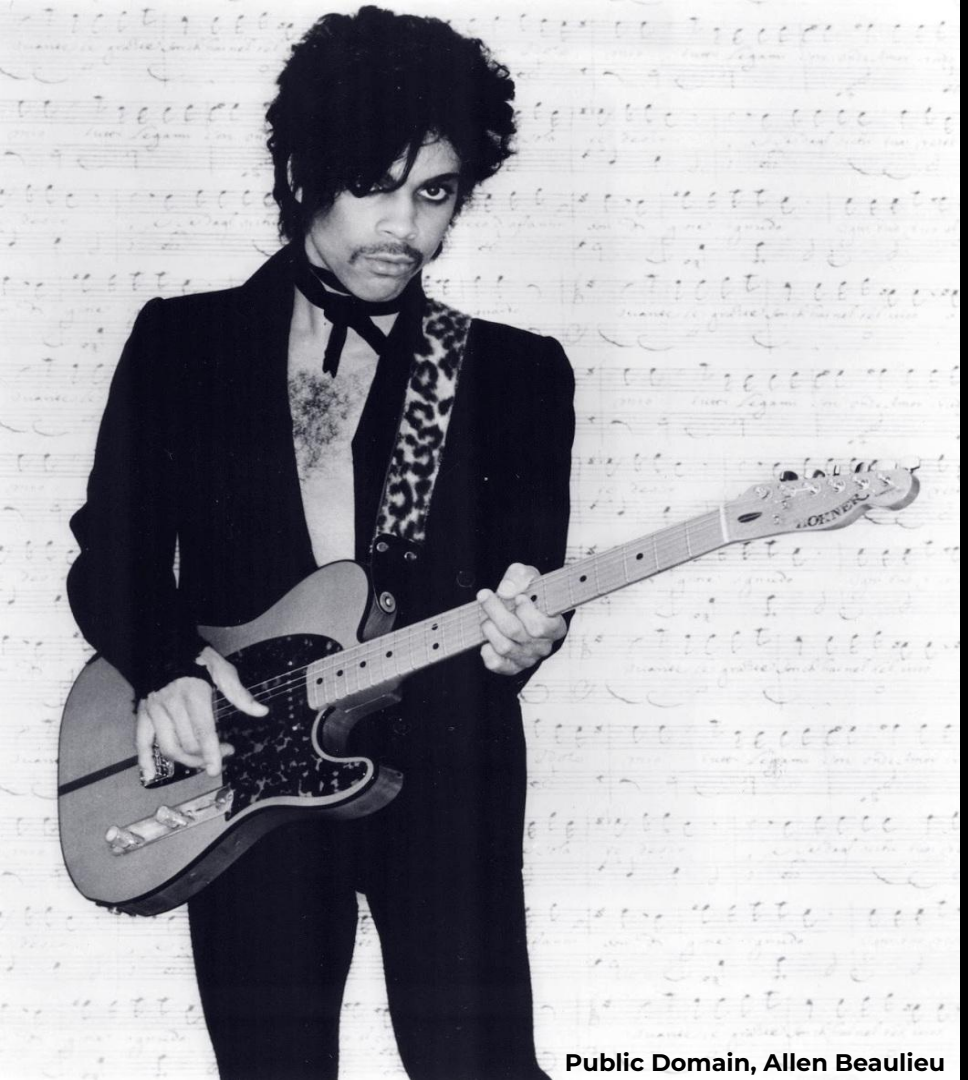

Public Domain, Allen Beaulieu

And then

The composer and singer of “When Doves Cry” passes away

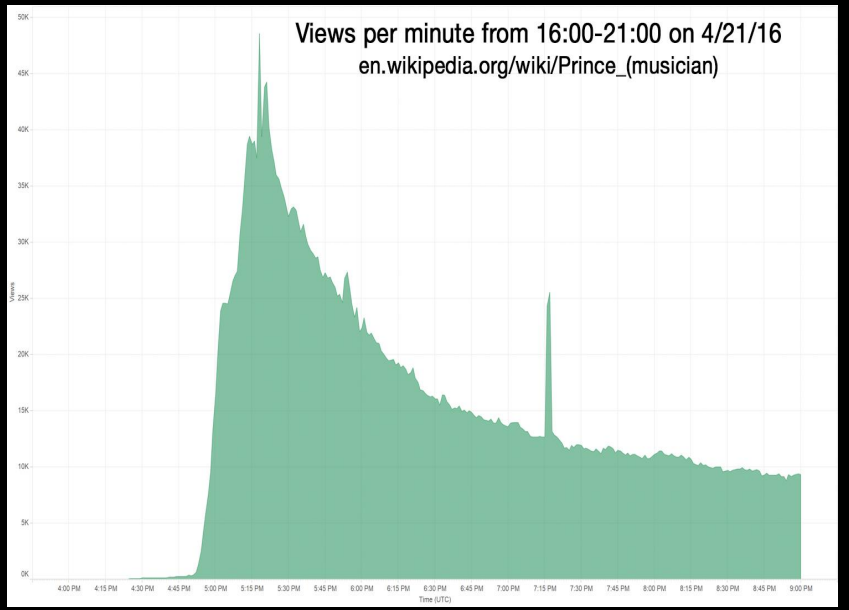

2016-04-21: 5.8M, 2016-04-22: 5.5M,
2016-04-23: 2.1M, 2016-04-24: 1.4M

~3.5 years later

- 26 January 2020
- Covid was about to hit hard (but we didn’t know)
- Helicopter crash, Kobe Bryant passes away in a terrible accident
- During the last All Hands event
- After a day of DDoS attacks
- Wikipedia collapses

2020-01-26: 9.5M views, 2020-01-27: 8M views

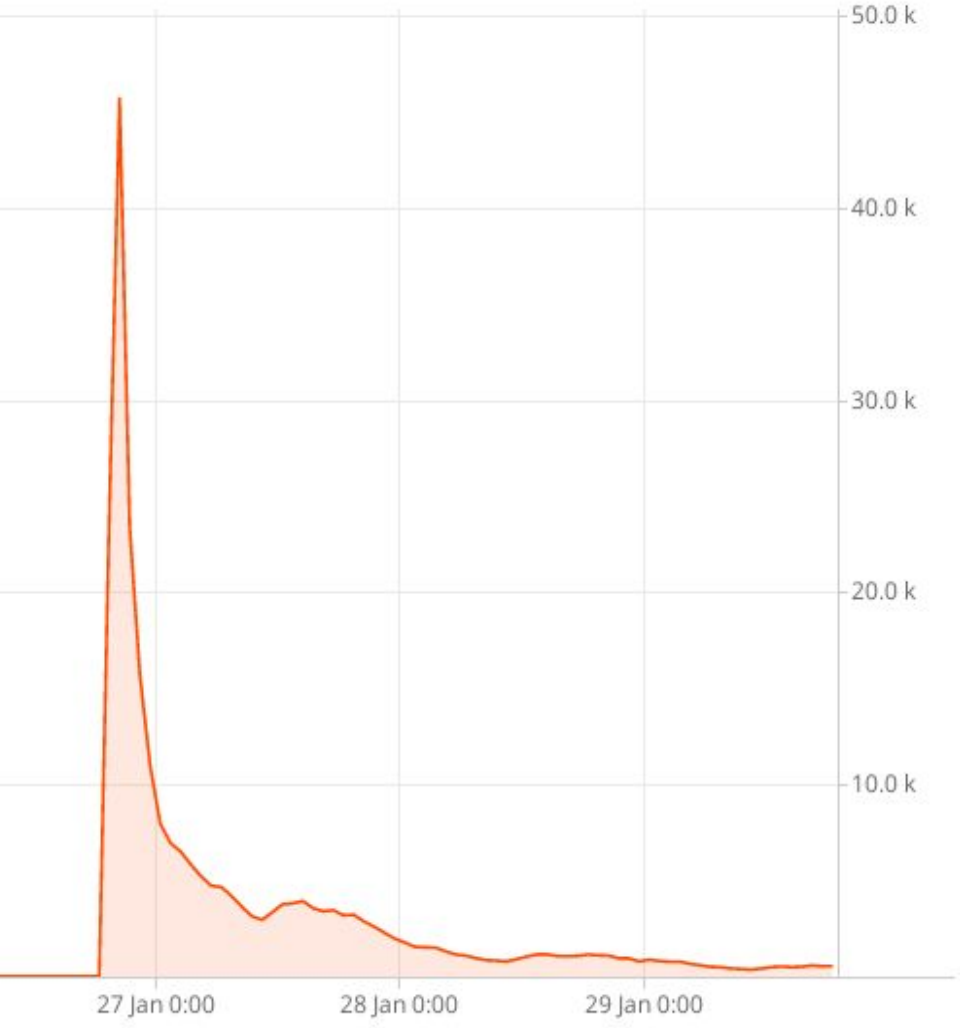

Why?

- People able to respond were either jet-lagged, tired, sleeping or in planes
- DDoS attack before the event
- The event was sudden
- The previous one was years ago
- Wikidata, started in 2012, had by 2020 become a major data provider for Wikipedias. A regression had been inserted
- Wikipedia usage had increased considerably since 2009, but our investment hadn't.

WIKIMEDIA FOUNDATION

What did we do?

Multi pronged approach

- Ran a post mortem
- Beefed up DDoS protection
- Altered the Wikipedia <-> Wikidata communication paths
- Increase memcached capacity
- Utilize the secondary DC (Multi-DC)
- Improve incident response culture and processes
- Improve onboarding
- Never travel again
  - Just joking, but Covid has indeed forced reduced travel

WIKIMEDIA FOUNDATION

Anti-DDoS

Procure and integrate better

- We had already some DDoS scrubbing, recently procured
- Worked with suppliers
  - Reported bugs
  - Request functionality
  - Integrated their API into our automation cookbooks
- Wrote a playbook
- Trained ourselves to use it

WIKIMEDIA FOUNDATION

Wikidata changes

Wikidata + Babel

- Open knowledge base of structured linked data
- Powers a lot of infoboxes in Wikipedia, including user language proficiency (Babel)
- Babel uses memcached but falls back to (expensive) HTTPS API calls to MediaWiki to fetch the primary data.
- Cue in memcached link saturation and...

Solution

- A service mesh
  - Envoy based
  - Persistent HTTPS connection
  - Retries
  - Circuit breaking
- Memcached
  - Invest into a local memcached instance per MediaWiki server
  - Invest into a centralized memached “gutter pool”
- Make Babel not talk to the MediaWiki API
  - It is part of MediaWiki anyway, it can fetch the relevant data “cheaply” from the database

WIKIMEDIA FOUNDATION

Multi DC

Wait, you said 2 “core” DCs

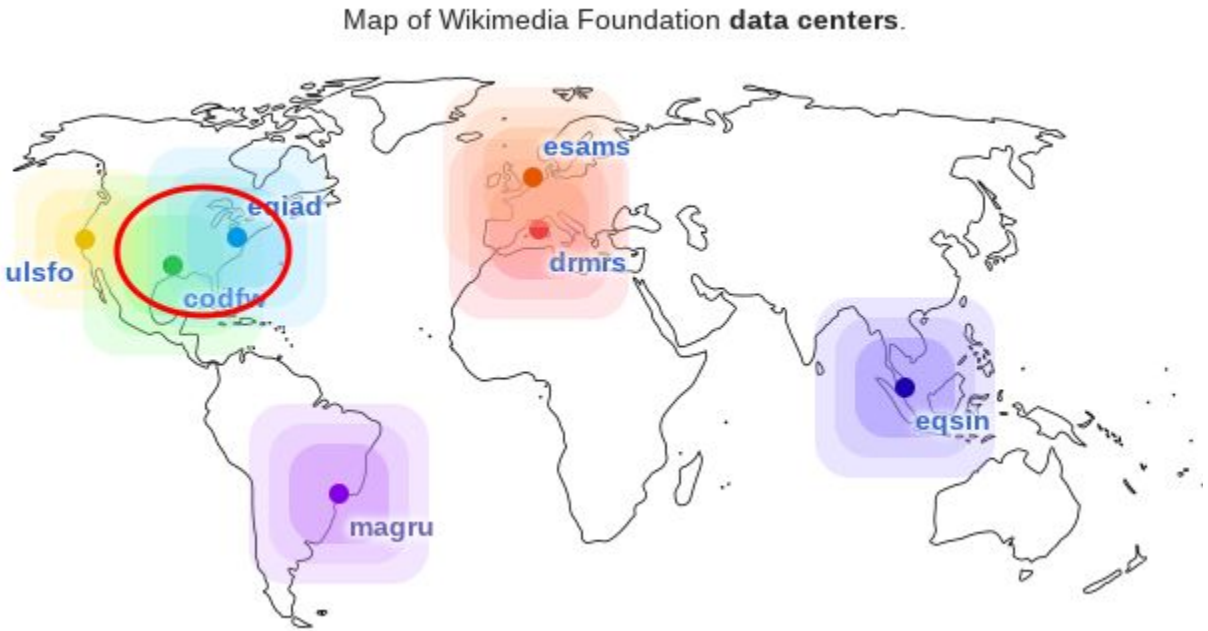

- We haven’t been using both DCs simultaneously until 2022.
- Active/Passive with some asynchronous jobs being executed on the secondary
- We ‘ve been switching occasionally (~1 year) to the secondary
- But the CAPACITY was mostly unused

But always wanted to

- Multi-DC strategy RFC in 2015
- Implementation happened in “stages”
- By 2020, ready to go for the last few parts
  - CDN improvements
  - Session Store/MainStash
  - Cross-DC MariaDB secure writes
- The 2 DCs are 40ms apart, 1 is always in read only mode

CDN improvements

Which DC should a request land to?

- HTTP POST => RW DC
- ?UseDC=master => RW DC
- ?cpPosIndex= => RW DC
- ?action=rollback => RW DC
- Various SingleSignOn system workflows => RW DC
- Everything else => local DC, aka closest

Cache invalidation

- Used to be UDP multicast and packets got …
- Switch to a kafka-based local daemon per CDN node (purged)

MainStash/Session Store

- Sessions
  - From Redis -> Cassandra (via a Golang proxy named Kask)
  - Sessions now are in both DCs
  - Happily peaks at 1.5k rps
- MainStash
  - Non critical, non derivative data, strong persistence.
  - From Redis -> MariaDB
  - `MediaWikiServices->getMainObjectStash()`

Cross-DC MariaDB Secure writes

- Edge case handling where for $reasons an action that would result to a DB write ends up in the read only DC.
- The write needs to go to the RW DC, but encryption was deemed a MUST
- Various approaches evaluated:
  - ProxySQL
  - “dumb” TCP tunnel
  - Envoy as TCP tunnel
  - HAProxy in TCP mode.
- What did we end up doing?
  - Directly from MediaWiki without additional proxies or tunnels.

WIKIMEDIA FOUNDATION

Culture

Blameless Postmortems

- We had an incident response process setup, including incident status doc but… no structured postmortem meetings
- Introduced blameless postmortems meetings
  - Review incident docs, discuss and figure out causes
  - Share the experience and difficulties met
  - Create specific action items and address them
- Overall, it’s been a success

Onboarding

- We all had our interesting onboarding tales, we were remote heavy pre-Covid anyway
- Recognized the deficiency, invested in a proper onboarding framework
- Onboarding buddy
- Checklist of things to know and by when
- Incident response is part of the onboarding process

WIKIMEDIA FOUNDATION

Did it work?

Public Domain, Donald McKague

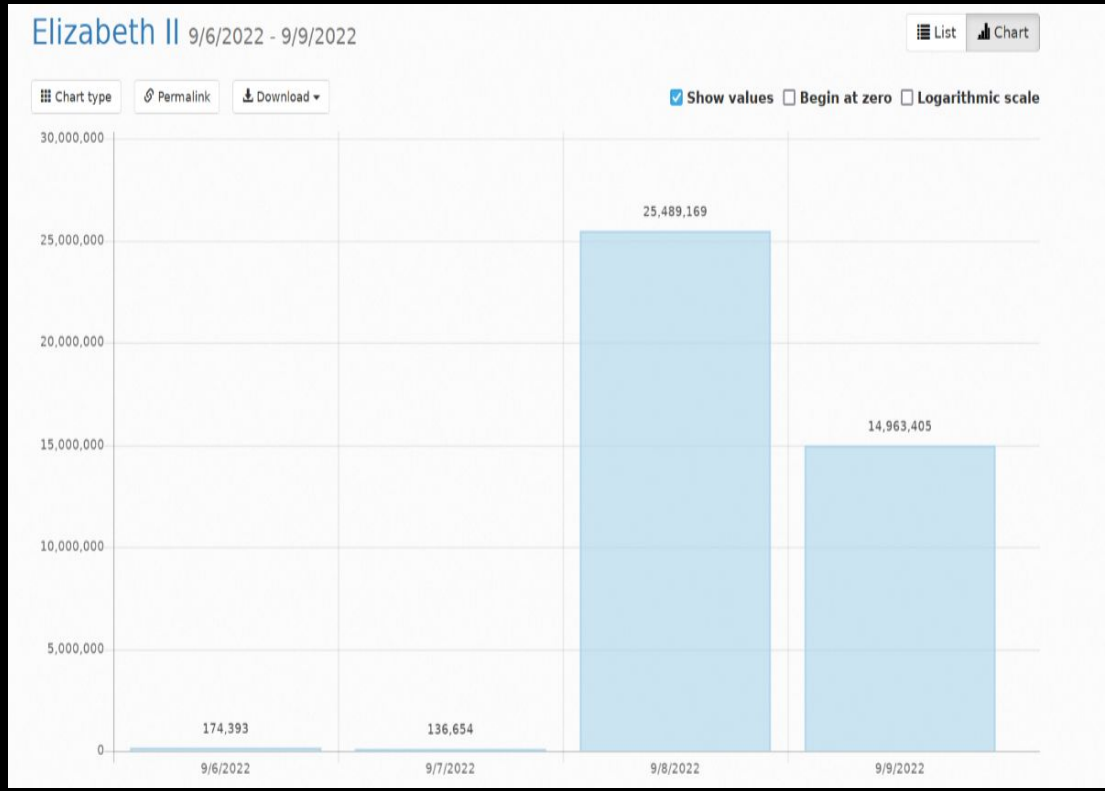

Elizabeth II
2022-09-08: 25.4M views
2022-09-09: 14.9M views

Public Domain, Office of National Drug Control Policy

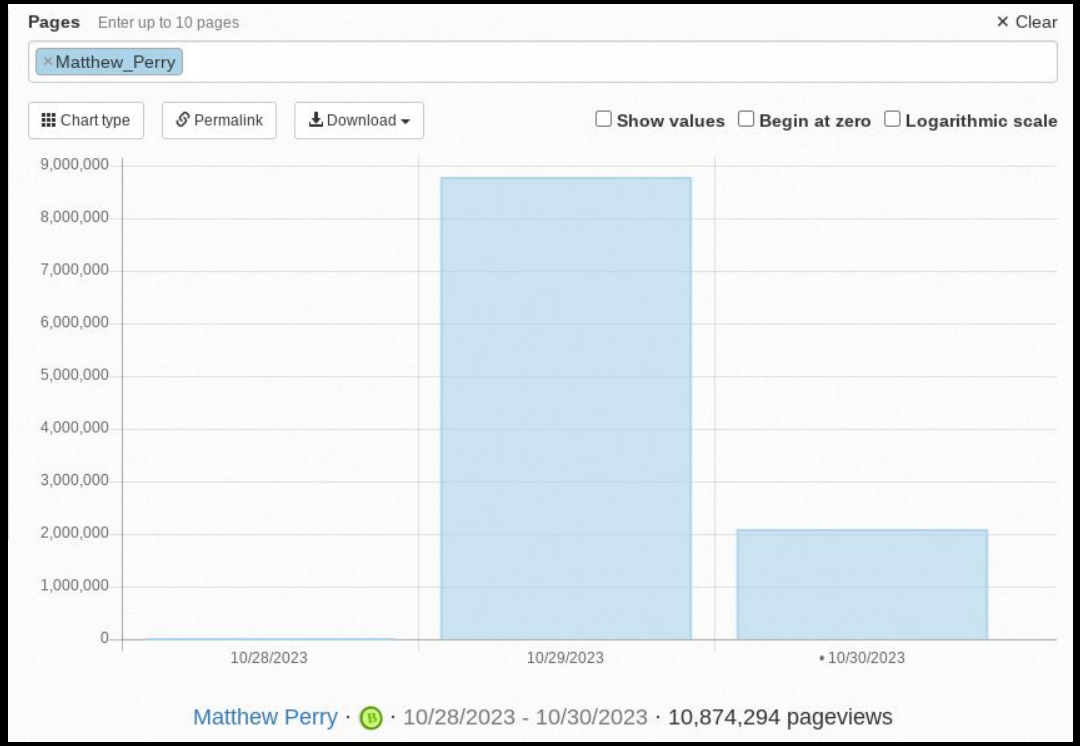

Matthew Perry
2023-10-29: 9M views
2023-10-30: 2M views

Thanks

- Multi-DC (**Aaron Schulz, Tim Starling, Timo Tijhof**)
- Kask (**Eric Evans**)
- Feedback (**Service Ops SRE team @ WMF**)
- Wikibase changes (**Wikimedia Deutschland**)
- Slide deck template (**Chris Danis**)
- The idea of how to talk about this (**Giuseppe Lavagetto**)

Be brave, be curious, be determined, overcome the odds. It can be done