




Past, Present, and Future of WMCS

Cloud VPS, Toolforge, PAWS, Quarry and friends



Athens, May 2023



Arturo Borrero González

Senior SRE @ WMF




Bryan Davis

Principal Software Engineer @ WMF



Past



How we got to the Wikimedia Cloud Services we have today

Ancient times: 2001-2010



- 2001 [Wikipedia started](#)
- 2001 [Bots started doing things](#)
- [...]
- 2005 [Toolserver created by WMDE](#)
- [...]



Past: 2011 - 2014

- 2011 [Wikimedia Labs project started](#)
- 2012 [Wikimedia Labs has 120+ projects and 600+ users](#)
- 2013 ["Tool Labs" project started with intent to replace Toolserver](#)
- 2014 Quarry introduced
- 2014 [Toolserver was shutdown](#)




Past: 2015 - 2017

- 2015 PAWS introduced
- 2015 [Initial kubernetes support](#)
- 2015 [“Labs labs labs” naming problem](#)
- 2017 [Wikimedia Cloud Services team created](#)
- 2017 [Labs and Tool Labs rebranded as Cloud VPS and Toolforge](#)
- 2017 [Wiki-Replicas redesigned](#)



Past: 2018 - 2020

- 2018 [Openstack Neutron network deployed](#)
- 2019 [HTTPS enforced for Toolforge](#)
- 2019 [Ubuntu replaced with Debian in Toolforge](#)
- 2020 [Toolforge Kubernetes cluster rebuilt](#)
- 2020 [New toolforge.org domain introduced](#)
- 2020 [Wiki-Replicas redesigned](#)





Celebrate !

Almost 20 years of technical community efforts.

Thanks to you, the movement.



Present



What “just” happened, is happening “now” or is about to happen “soon”


Some numbers

- ~30.5% of total wiki edits come from WMCS [\[April 2023\]](#)
- ~ 3200 Toolforge tools
- ~ 2500 Toolforge maintainers
- ~ 1000 Cloud VPS virtual machines
- ~ 200 Cloud VPS projects
- 4000 vCPU, 27 TB memory, 140 TB storage




Recent changes


- [Toolhub](#)
- New domain names were introduced:
 - [wmcloud.org](#)
 - [wikimediacloud.org](#)
 - [wikimedia.cloud](#)
- [HTTPS enforced at front proxy layer](#)
- [Openstack Trove](#)
- [Terraform support](#)
 - publicly reachable openstack API endpoints
- Ceph distributed network storage
- [Cinder volumes \(attachable disk space\)](#)
- [Toolforge jobs framework](#)



Now and soon

- ToolsDB improvements
- [Toolforge builds service](#)
- Openstack Magnum (k8s-as-a-Service)
- Cloud VPS object storage (swift)





Future

Let's talk about the next 10 years




Toolforge

- Toolforge as a platform
- Toolforge push-to-deploy
- No ssh interaction required
- Central log system
- NFS ?
- ToolsDB ?



Cloud VPS

- Multiple DC support
- Tenant Networks
- Kubernetes as undercloud
- IPv6
- Beyond virtual machines ?
- GPUs ?



Data Services

- Wiki-Replicas improvements ?
- Superset as a replacement for Quarry ?
 - preview: superset.wmcloud.org
- Database-as-a-Service beyond Trove ?
- Data pipelines ?
- Analytics / Data lake ?

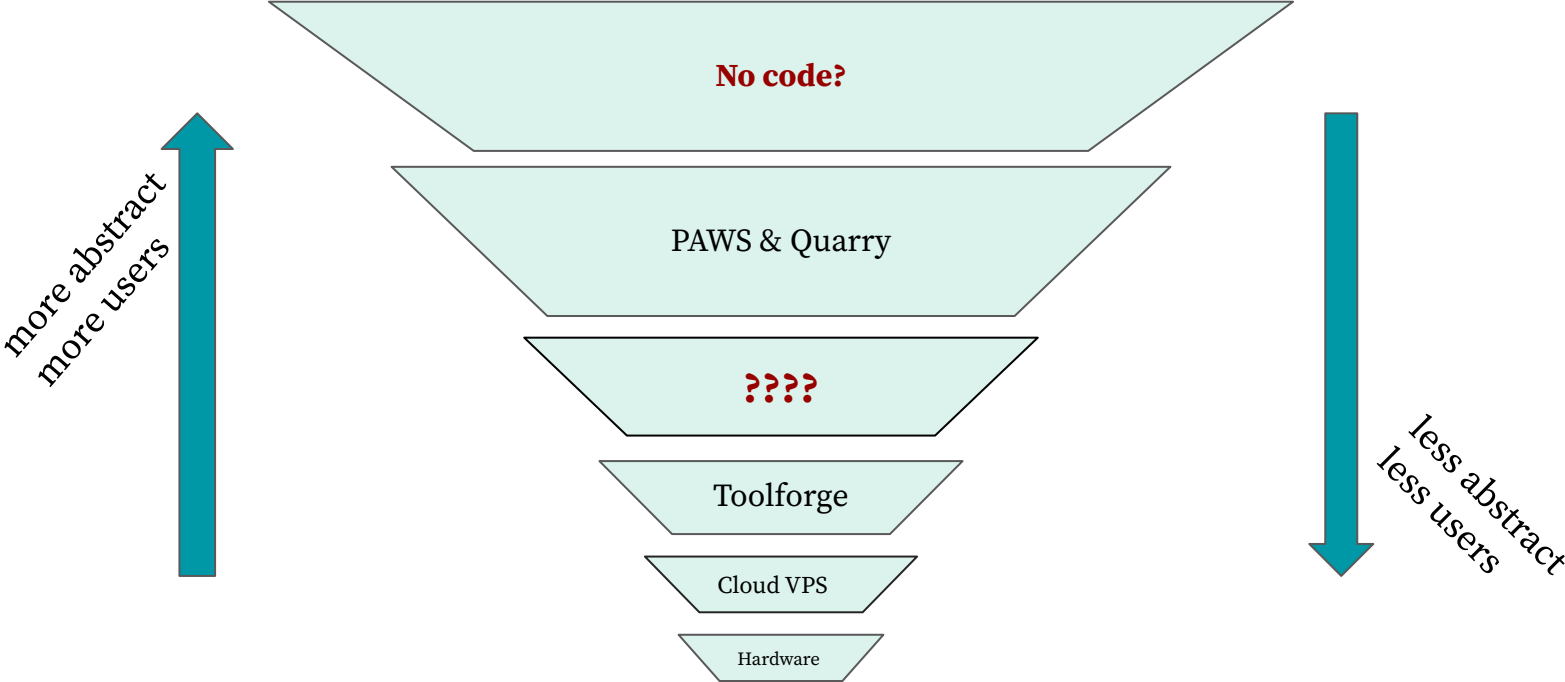


New computing abstractions?


- Helm charts as deployment unit ?
- What does the Wiki community need ?



New computing abstractions?



Open debate




- Cloud VPS ?
- Toolforge ?
- Data Services ?
- PAWS & Quarry ?
- New computing abstractions ?





Past, Present, and Future of WMCS

Cloud VPS, Toolforge, PAWS, Quarry and friends



Athens, May 2023