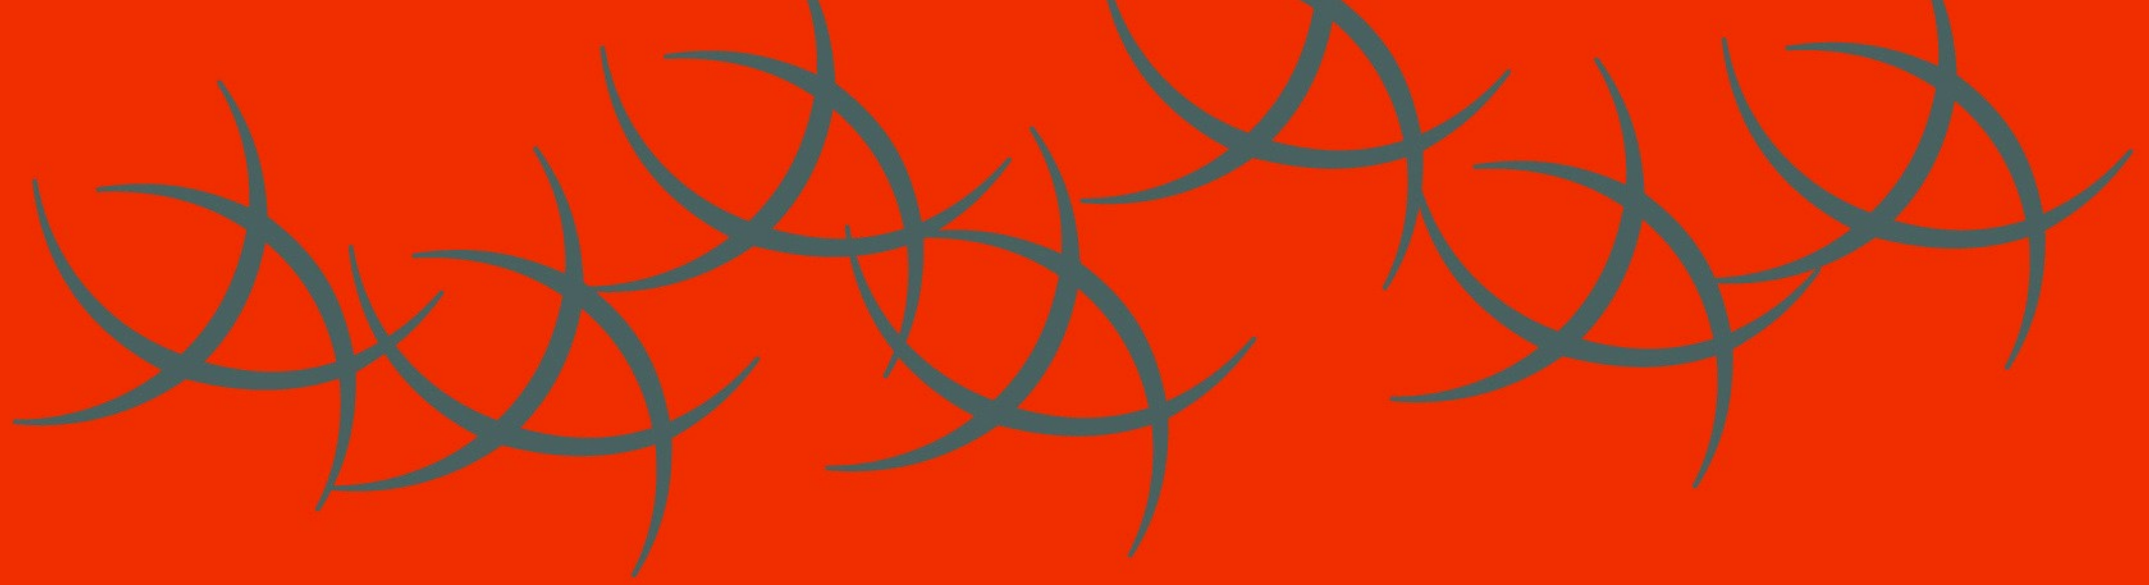

arnes

Running Moodle at ARNES

A year long journey of running Moodle at large scale

Martin Božič, Arnes

Managing and running large scale Moodle system, for IT, GÉANT web conference, 20th April 2021

Introduction

- Backstory
- Spring lockdown
- Fall lockdown
- Architecture
- Operations and monitoring
- Support

Backstory

ARNES provides ICT services for Higher Ed. and Research institutions as well as primary and middle schools

- belong to Ministry of Education
- the whole country has about 2m people, about 260k pupils in primary and secondary schools and 60k students in higher ed

I'm Martin Božič, system admin/team lead for Moodle at ARNES

- part of internal team for system and application support
  - we wear many hats in our team (at least 2 each)
- I came to ARNES in 2014 to help with migrating the ~160 Moodle sites on their OpenVZ containerized solution to **a single Moodle instance**
- The decision for such solution was made upon the good experience with our Wordpress service (similar migration from Joomlas on containers to Wordpress)

Backstory – moving Moodle from containers to “One to rule them all”

- *BUT! (you might say) Wordpress has a true multi-tennant feature built-in, Moodle wasn't built for that kind of deployment!*
  - *Technically true, but with a couple of additional plugins and top categories separation we've made it work well enough.*

Factors:

- There are little to no tech experience running website/app on internet in primary/secondary schools
- containerized solution from previous project on “auto-pilot” and hard to maintain
- Joomla already migrated to Wordpress

Backstory – October 2019, Migration complete

Migrations complete and the result:

- 366 institutions
- 12000 courses (5000 empty!)
- 58000 registered users
- max 150 “live users” (logged in last 5 min)

Thursday 12th March 2020

The signs

- rumours of state wide lockdown announced for monday
- unusualy high activity of users on our Moodle
- our facilities would be lockdown too (except on-call)
- 2 virtual machines for Moodle and 1 admin
- started preparing before any official announcement

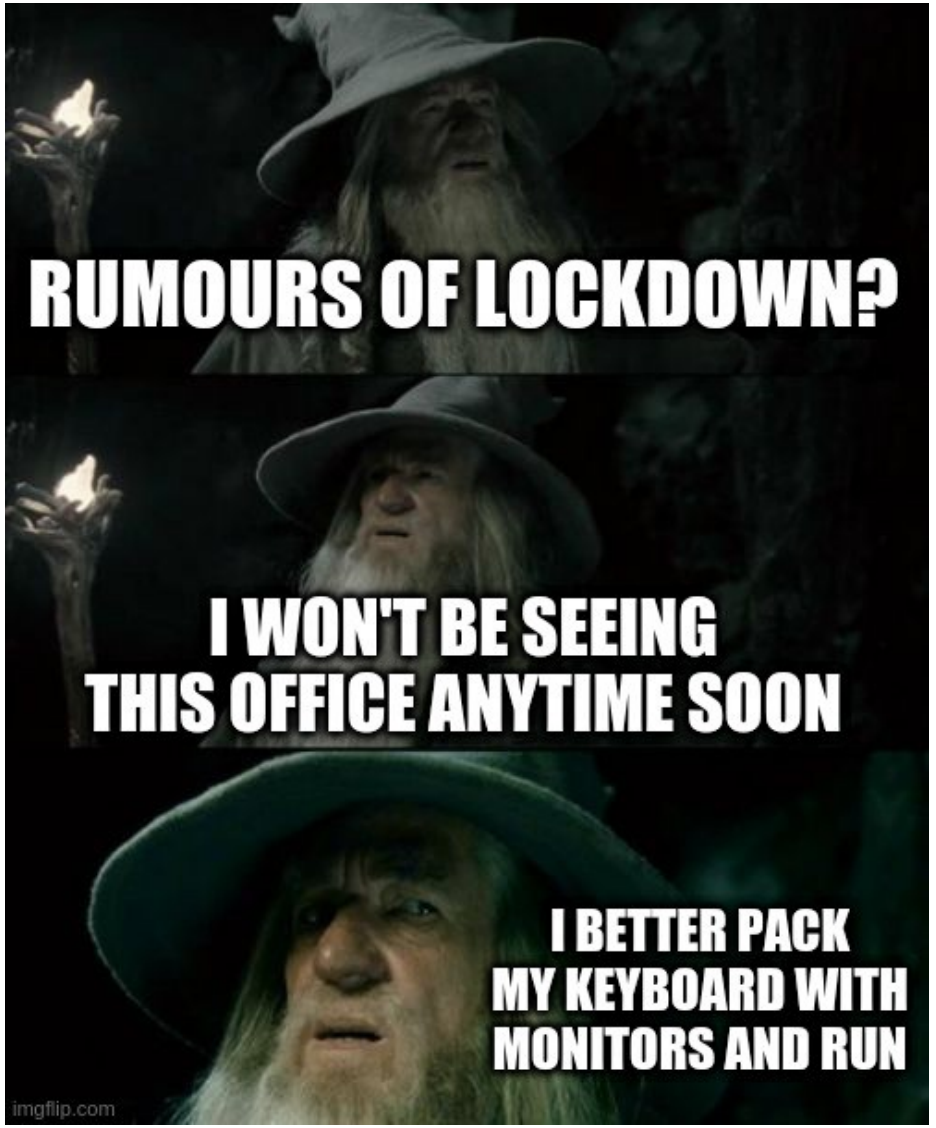

Thursday 12th March 2020

The scramble

- Moodle was the priority – all hands on deck (10 people)
- used what was at hand and were comfortable with – classic cluster setup
- started scale-out with VMs – abandoned quickly
- cannibalized some hardware from HPC and spares we've had in store

The stack

- CentOS 8
- NFS v4/v3
- httpd 2.4
- PHP 7.4 (FPM)
- MariaDB 10.3
- Redis 5
- Shibboleth SP 3.1

Spring lockdown

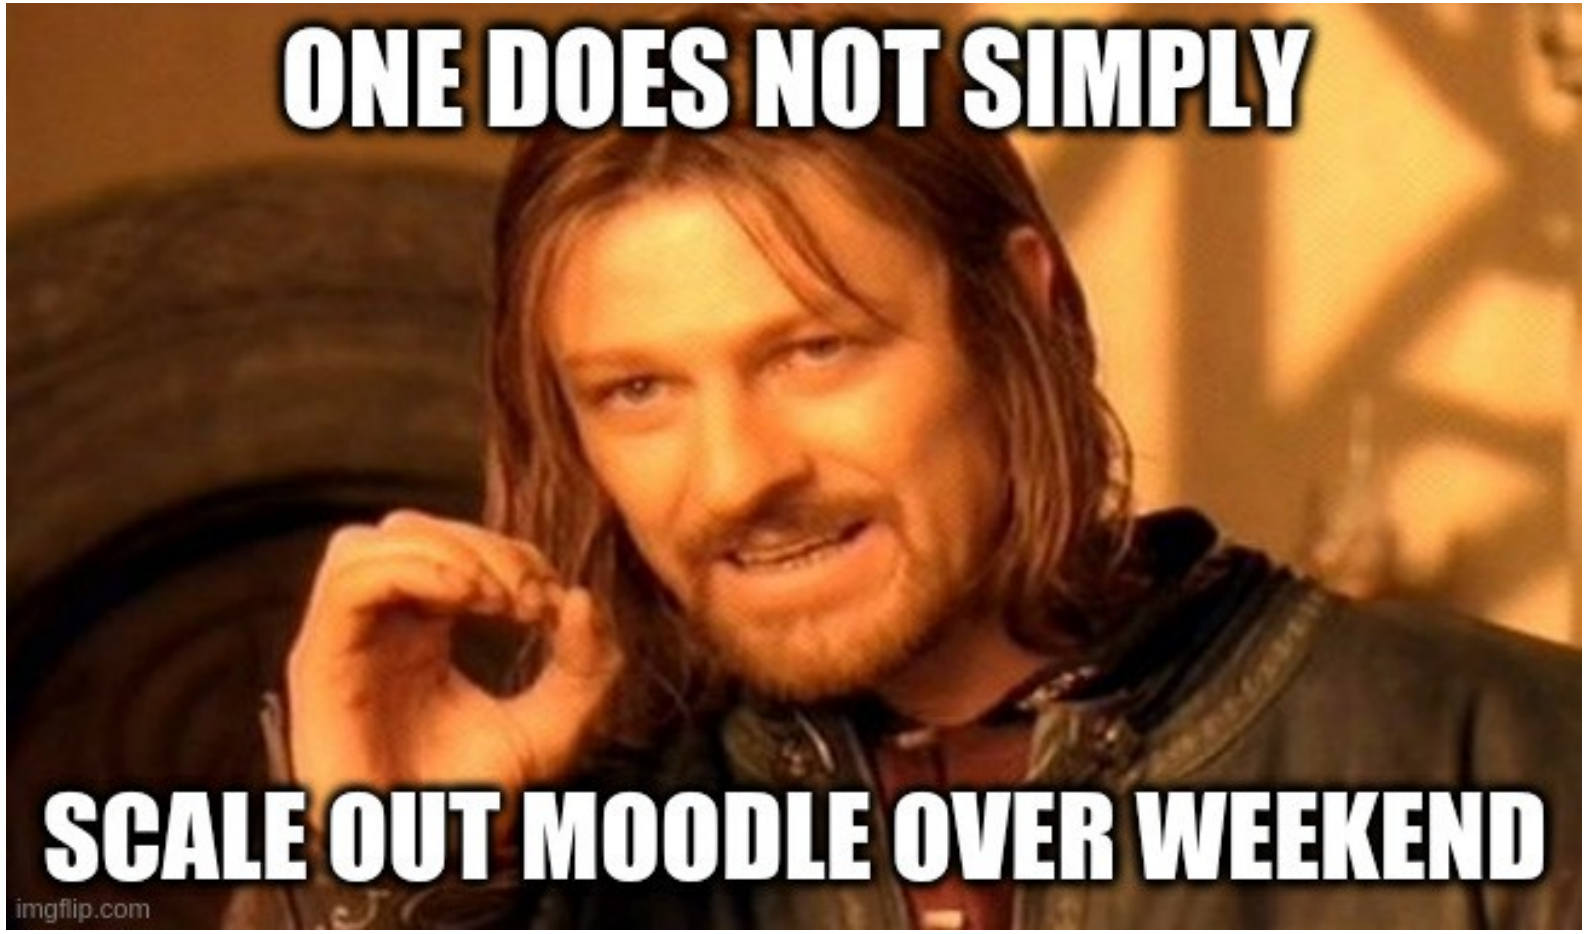

Spring Lockdown – the numbers

Lockdown lasted from 16th March – 14th May 2020

Mid april 2020 stats

- peak at 7k live users (150 peak in Oct 2019)
- 90k users per day
- 102k registered users
- 23k courses (12k in Oct 2019)
- 0.5 seconds average response time

The infrastructure

- 3 HW web nodes
- 1 SQL HW master, 1 SQL VM slave
- 1 NFS HW server

Preparations for the 2nd Wave

- analyzed the data and estimated possible bottlenecks in the second wave and subsequent lockdown
- had to tear down and develop a new staging environment for the Moodle upgrade
- upgraded Moodle from 3.7 to 3.9 to use the new read-only SQL slave feature
- developed new plugin for users and classrooms provisioning from an identity service
- ordered new and better hardware to replace the current makeshift setup

Education of teachers

- a special team at ARNES prepared online courses for teaching users which contributed to the usage increase in 2nd wave beyond our expectations

The Fall/Winter Lockdown – Hell Week of October 2020

The Fall/Winter Lockdown – Hell Week of October 2020

- how do you know you now run a Moodle at large scale?
- you run into the same bugs as other big sites (MDL-51111)

The load testing

- had little experience (we did load testing on Moodle years ago)
- had little time because there were so many things to do just to make things work
- opted for custom scenarios with JMeter
- needed as realistic scenario as possible
  - testing mostly login, dashboard view, course visit, a couple of resource lookups and logoff
    - login including the whole federated SAML shebang for better realism
- did it on production
  - enrolling 10000 test accounts
  - test late at night
- slowly cranking up the load until anomalies started to appear

The Fall/Winter Lockdown – Hell Week of October 2020

- some hard hitting DDoS attacks while we were resolving the bottlenecks
- couldn't really afford downtime anymore, so we went with Cloudflare Business as DoS protection and load balancer.
- decided to let in mostly neighbouring states, EU and some specific exceptions

The Fall/Winter Lockdown

- 3 months of complete schools lockdown
- growth 400 MB/day uploads – 2TB / week
- *moodledata/filedir* from 0.7 TB in March 2020 to 35 TB in March 2021

The Winter/Fall Lockdown

Lasted from 19th October 2020 - 15th February 2021

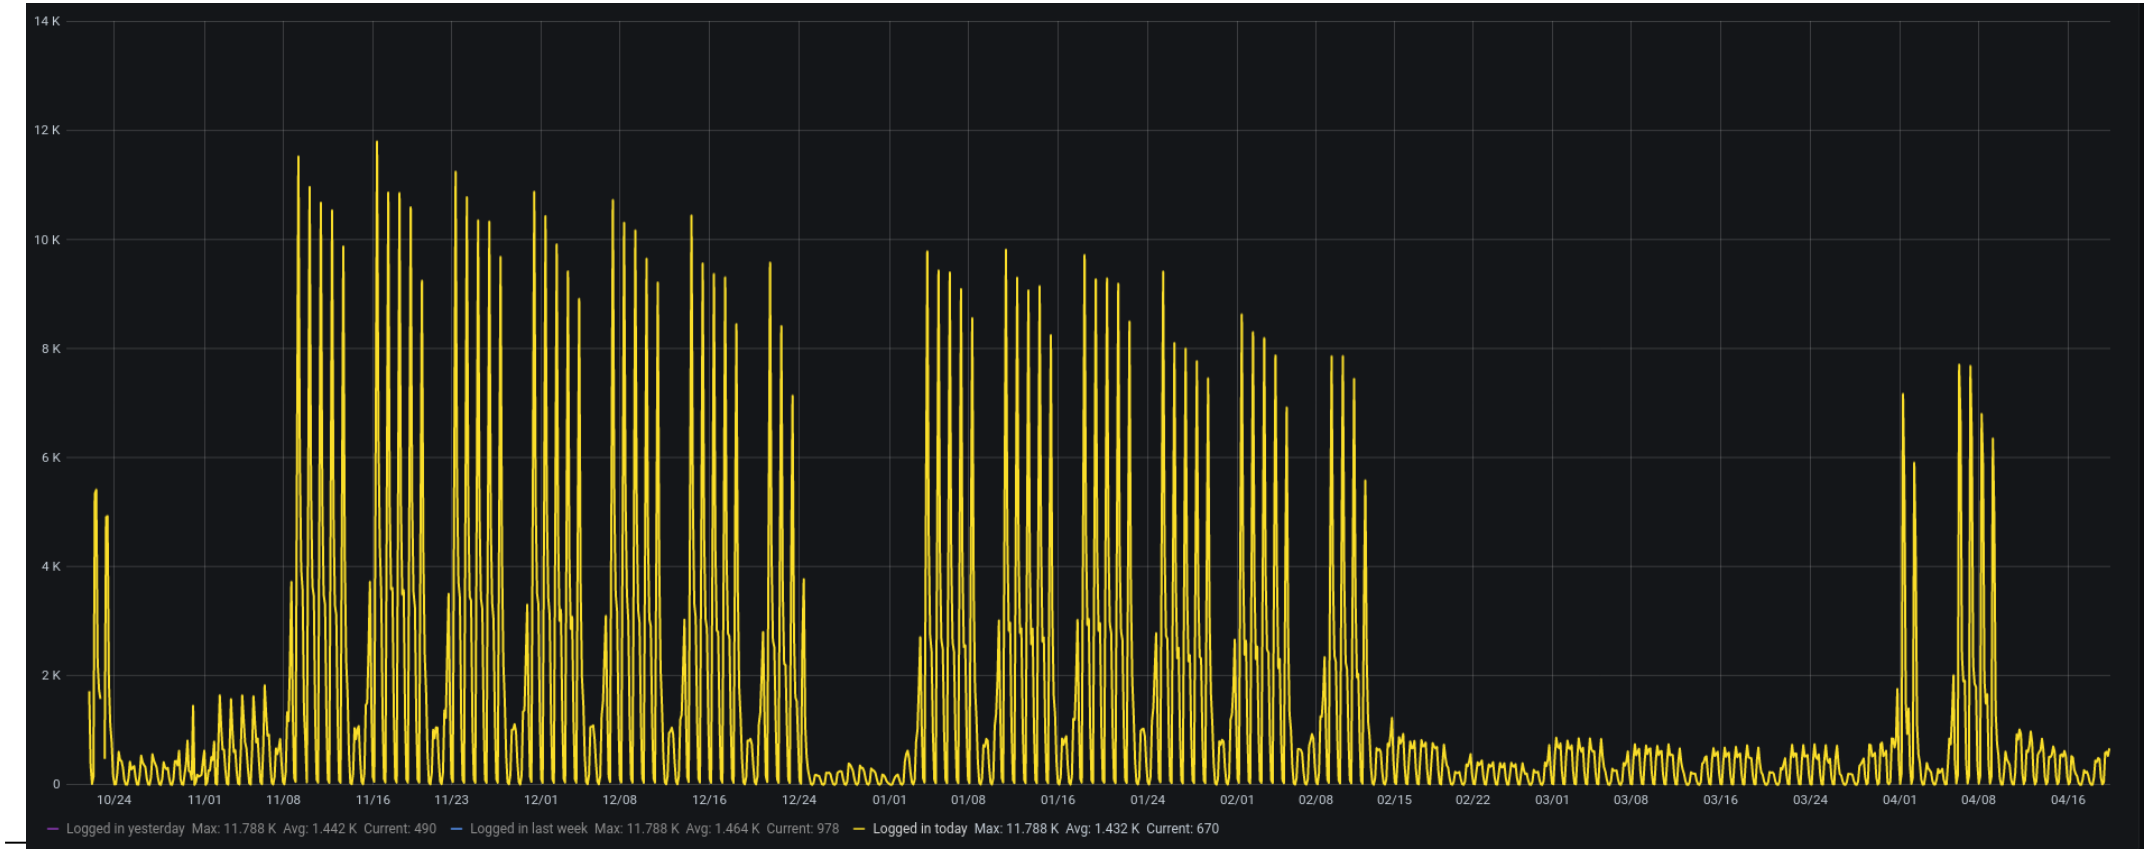

Easter Lockdown – the numbers

April 2021 stats

- lockdown for a week after Easter
- peak at 12k live users (150 peak in Oct 2019)
- 136k active users per day
- 203k registered users
- 57k courses (12k in Oct 2019)
- 1.5 seconds average response time

New bottlenecks discovered from the accumulated data:

- too many events trashing *mdl_event* table
- too much H5P content uploaded
- but too short of a lockdown to really reproduce some scenarios
- left too long *logstore_standard_log* (300m lines for 6 months)

Architecture

Production

- 9 bare-metal web frontends behind Cloudflare
- 2 bare-metal SQL servers in master-slave setup; 320G database + 350G separate logging DB
- 1 NetApp NFS server (for other services too) ~45 TB used
- 1 VM for cronjob server
- 4 VM web frontends repurposed for testing, long running queries and additional ad-hoc task servers
- 1 VM for OnlyOffice document converter

Dev and testing

- 1 VM for QA
- Docker images in GitLab CI

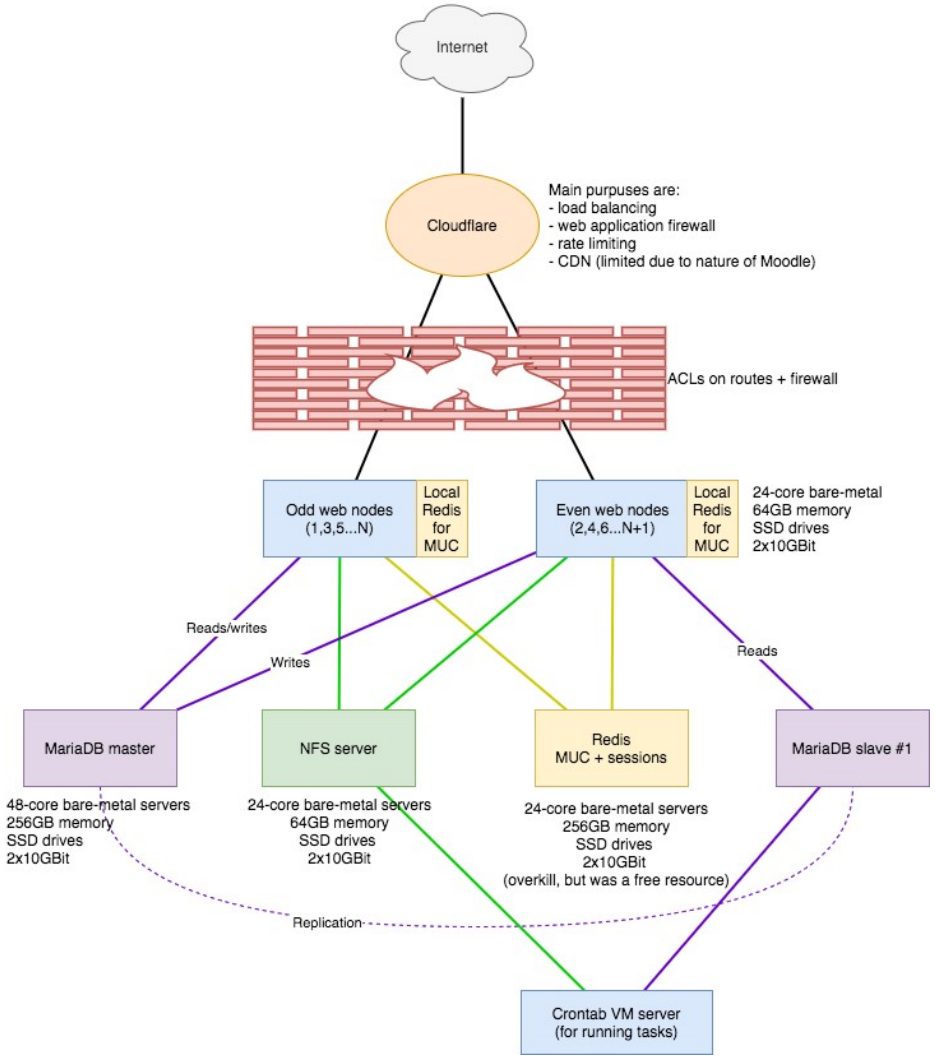

Architecture

Production

- web nodes 24-core AMD EPYC, 64G, 2x10 Gbit
- SQL server 48-core CPU, 256G, 2x10 Gbit
- the old NFS server was the same as web nodes
- cronjob VM 4 vCPU, 8G RAM
- QA VM (LAMP), 4 vCPU, 16 G

Development

- git repos, centralized build system in GitLab

Deployment

- development: local, Review Apps from GitLab
- staging: fully automated
- production: manual, partialy automated and heavily documented procedure

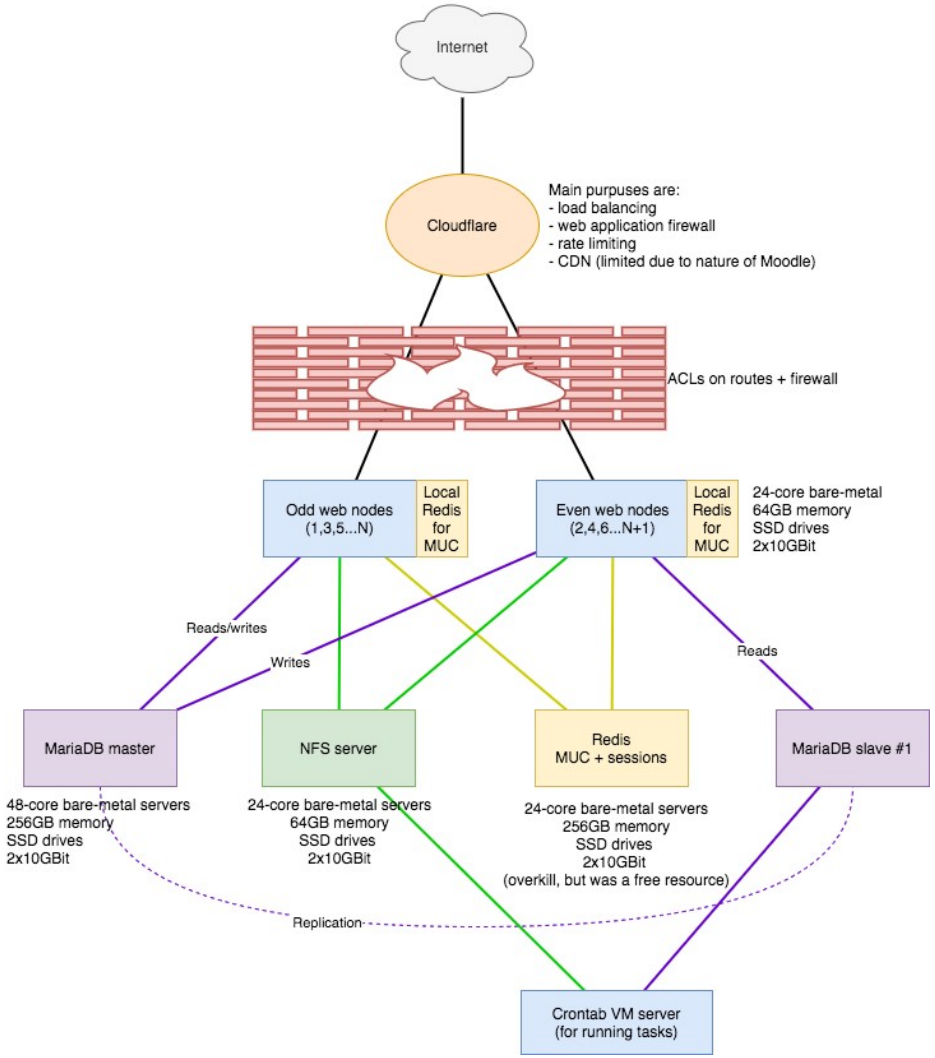

Monitoring

- Nagios – our legacy alerting system, does it's job fine
- ELK – our go-to logging and analytics Swiss Army Knife
- Grafana – day to day use
- Prometheus – slowly dipping our toes into it
- New Relic – very useful for performance tracing, only on one host

Monitoring

- servers performance
- services performance
- user behaviour (current live users count, logins, current visits on resources and activities)
- cronjobs buildup

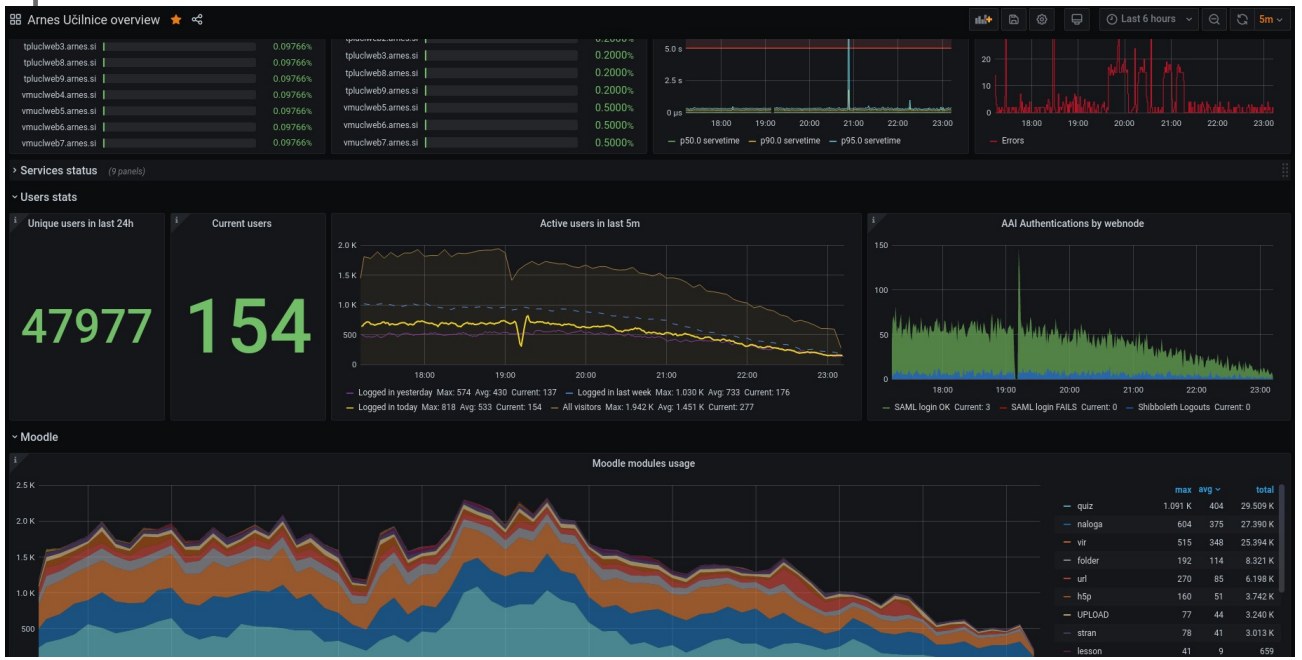

Support

- helpdesk bears the brunt of support
- our team acted as 3rd level of support
- copy pasted Word documents again and again
- one-Moodle-for-all specifics
- Cloudflare limitations (timeouts on reports, upload limits)
- our misconfigurations
- failing cronjobs (mails not sent)
- real bugs beyond our current ability to resolve

The Epilogue

Achieved

- managed to get through the whole ordeal
- scaled the service beyond our wildest imagination
- learned a ton of new stuff
- appeared in national media, but not in a good way :-/
- became 2nd most googled term in 2020 in Slovenia :-)
- connecting with other NRENs running Moodle

Lessons learned

- ask for help sooner
- you'll fail and then some more
- but you'll learn a LOT!
- load testing is hard but inevitable!

The Epilogue

What next?

- course reset cycle for next year
- course templates
- forks of core and plugins dilemma
- problems of running behind Cloudflare
- Moodle documentation for clusters
- better cronjob monitoring and troubleshooting


THANKS FOR SHARING
THE BURDEN
imgflip.com