

Kubernetes, Resistance is Futile

Adnan Hodzic | Lead SRE: DAP Public Cloud | [@foctrl](#)

April 2023



do your thing

Kubernetes, Resistance is Futile

Adnan Hodzic | Lead SRE: DAP Public Cloud | [@foctrl](https://twitter.com/foctrl)

April 2023



- ★ High level overview
- ★ Talk revolves around my time with MLP
- ★ MLP = Machine Learning Platform



do your thing

Setup and premise



do your thing

- ★ ING has 60000 employees
- ★ 18000 working in Tech
- ★ Offices in 40 countries
- ★ Serving 38 million customers

Setup and premise



do your thing

- ★ ING has 60000 employees
- ★ 18000 working in Tech
- ★ Offices in 40 countries
- ★ Serving 38 million customers



- ★ Not as easy to deploy workloads to Prod
- ★ Even more complicated for a data scientists
- ★ MLP serves as a model hosting platform

Setup and premise



do your thing

- ★ ING has 60000 employees
- ★ 18000 working in Tech
- ★ Offices in 40 countries
- ★ Serving 38 million customers



- ★ Python + Docker > Docker Compose
- ★ Ansible > Docker Compose > VM's
- ★ Kubernetes ☹️

VM setup problems & annoyances

VM setup problems & annoyances



★ VM patching, risk, security operations taking a lot of our time



★ VM's treated like pets, deployment & development too long and error prone



★ Developed "chameleon", mimicking Prod VM's in a local container



★ Risk on horizon if there was a sudden & exponential growth of MLP



★ Large number of unplanned migrations

Making K8s migration a possibility

Making K8s migration a possibility



★ Containerize everything!



★ The more open source/CNCF software you use, the better



★ Spread the principles of SRE among all team members

☆ Building a football/soccer team, not having a single striker or a goalie

☆ Anyone being able to pick up anything, while sticking to their expertise

Making K8s migration a possibility



★ Containerize everything!



★ The more open source/CNCF software you use, the better



★ Spread the principles of SRE among all team members



★ Document everything

"idea is only 5%, the rest of it (95%) is execution"

Resistance to the idea & unwelcoming circumstances



★ Team wasn't onboard with move to Kubernetes



★ Research, POC's & demos of MLP v2 as part of ING "mastery days"

☆ MLP v2 on ArgoCD or Helm charts, Kustomize instead of Ansible?

Resistance to the idea & unwelcoming circumstances



★ Team wasn't onboard with move to Kubernetes



★ Research, POC's & demos of MLP v2 as part of ING "mastery days"



★ Demonstrate how existing problems can be fixed with K8s



★ Attempt to go to public cloud

Resistance to the idea & unwelcoming circumstances



- ★ Attempt to go to public cloud on GKE (Google Kubernetes Engine)
 - ☆ ICHP (ING Container Hosting Platform) as a destination target?
 - ☆ Move to the cloud, MLP GKE POC

Resistance to the idea & unwelcoming circumstances



- ★ Failed attempt to go to public cloud on GKE (Google Kubernetes Engine)
 - ☆ ICHP (ING Container Hosting Platform) as a destination target?
 - ☆ Move to the cloud, MLP GKE POC (success!)
 - ☆ “Schrems 2” law dispute stopped us in our tracks, but things are better now

Why migrate to Kubernetes?

Why migrate to Kubernetes?



★ Eliminate need for patching, other VM maintenance



★ A lot of great out of box features: HPA, cluster autoscaling, service discovery ...



★ No need to spend time creating your own tooling, like service discovery

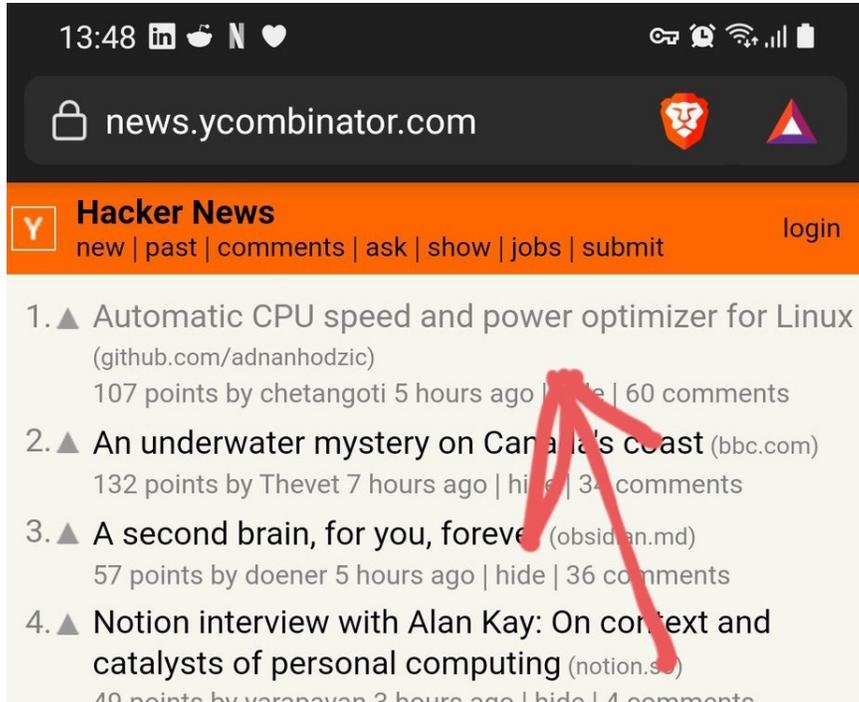


★ Fits ING's ambition to go green



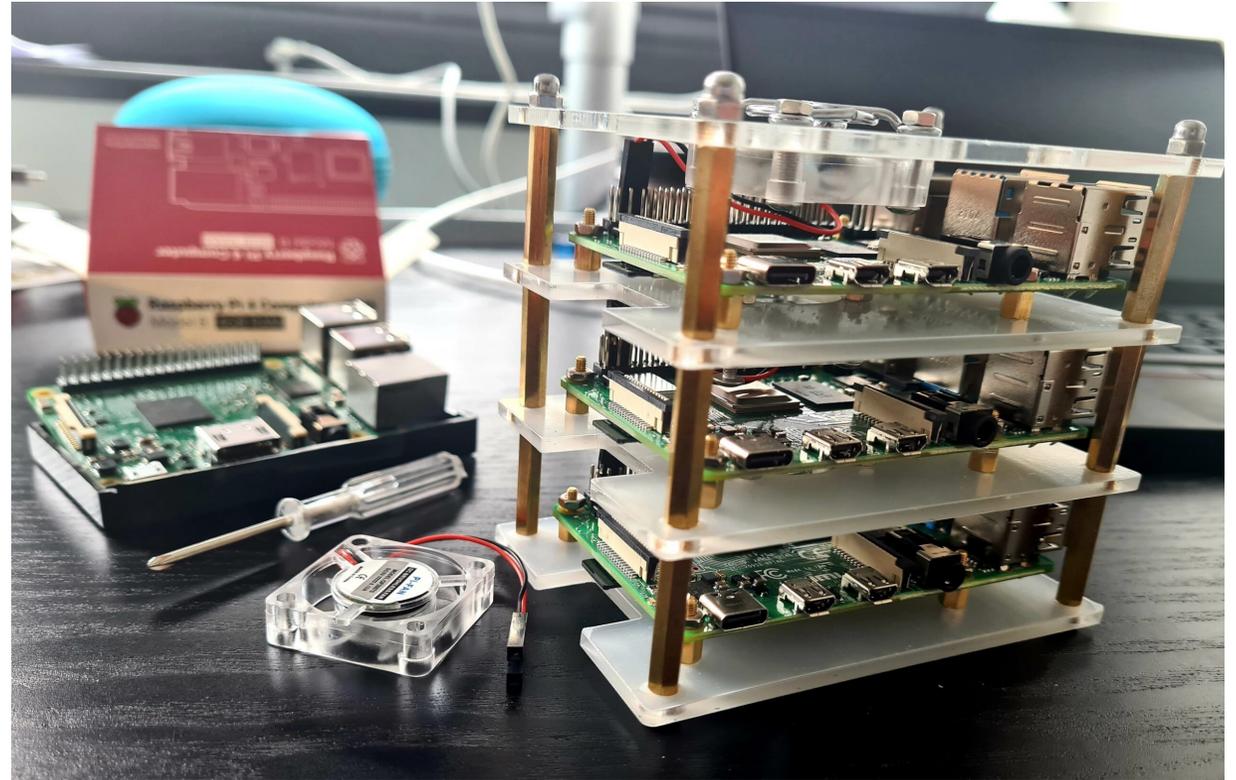
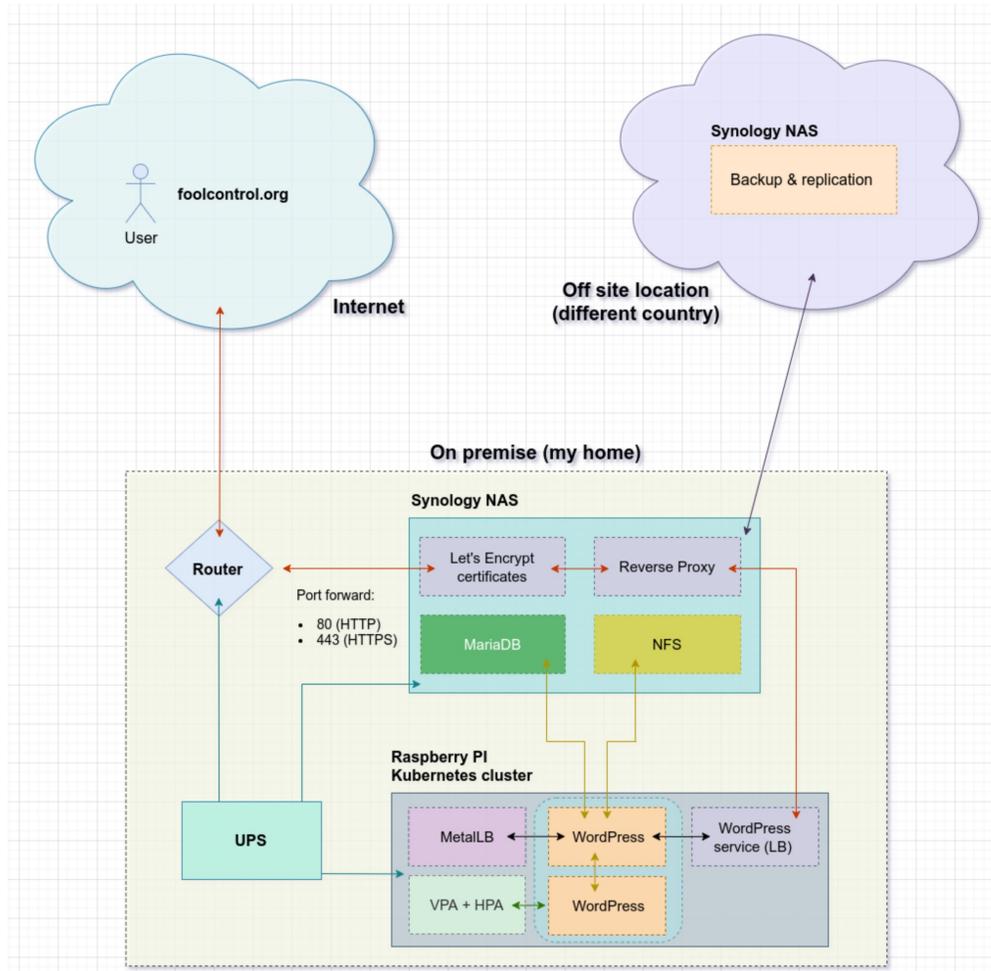
★ Once you're on K8s, you can seamlessly move anywhere

Meanwhile, in my private life ...



- ★ [auto-cpufreq](#) is trending as #1 on Hacker News
- ★ Under the load, my blog crashes
- ★ Inability to scale in current setup on VM's
- ★ Work in progress: [wp-k8s](#) project (scalable WordPress on Kubernetes)
- ★ Fix the scalability issue in minutes by moving to K8s
- ★ Birth of wp-k8s project for GKE: foolcontrol.org/?p=3754

Meanwhile, in my private life ...



★ wp-k8s: WordPress on privately hosted Kubernetes cluster (RPI 4 + Synology): foolcontrol.org/?p=4004

★ Potentially useful project for the attendees: atuf.app (Amsterdam Toilet & Urinal Finder)

Hitting the wall



★ Shared my private life K8s event & learnings with the team



★ More and more time needed to keep the lights on in current setup



★ Influx of a lot of big models which are requesting a lot of resources



★ Inability to grow at scale in current setup

Hitting the wall



★ Shared my private life K8s event & learnings with the team



★ More and more time needed to keep the lights on in current setup



★ Influx of a lot of big models which are requesting a lot of resources



★ Inability to grow at scale in current setup

☆ 1. Hire more people to keep scaling in current setup?

☆ 2. Move to Kubernetes?

Decision has been made to migrate to K8s



★ Create a design page, with acceptance criteria and scope of work



★ Conclude research and choose your target destination

☆ **ICHP** (ING Container Hosting Platform) vs **DAP** (Data Analytics Platform) Kubernetes cluster

☆ Option 1: **DAP** Kubernetes cluster:

☆ Vanilla Kubernetes, dedicated nVidia GPU cores, part of our area, match made in heaven?

☆ Option 2: **ICHP**:

☆ Stable, running in multiple zones, handles huge portion of risk & maintenance

Decision has been made to migrate to K8s



★ Create a design page, with acceptance criteria and scope of work



★ Conclude research and choose your target destination



★ At this point, it was a “just do it”. Step consisting of putting all the pieces together

Takeaways

Takeaways



★ During the migration, best approach is run 2 setup in parallel



★ Good monitoring & alerting setup will go a long way



★ Be brutally pragmatic, stick to the vision & design



★ Whole journey will take time



★ Perseverance and not giving up, re-evaluate the landscape

Takeaways

★ keep in mind that everything changes, certain ideas along with certain circumstances will match the changes in the landscape, and it might lead to a perfect storm as it did in the MLP's case

Thank you / Questions



do your thing

★ We're hiring: ing.jobs/tech

★ Slides available on:
foolcontrol.org

★ Questions & comments about the talk:
adnan.hodzic@ing.com

★ Connect with me on LinkedIn:
<https://linkedin.com/in/adnanhodzic>

★ Visit us at ING booth (S75):

☆ Meet today's ING Speakers, ask any remaining questions to Adnan, Diana or Thijs (Wed: 18:30 – 19:30)

☆ Meet the ING Speaker: "Kubernetes, Resistance is Futile" Adnan Hodzic (Thu: 15:30 – 16:30)



Please scan the QR above to leave feedback on this session



ing.com



@ING_News



LinkedIn.com/company/ING



SlideShare.net/ING



Medium.com/ing-blog



ingwb.com



Facebook.com/ING



YouTube.com/ING



Flickr.com/INGGroup