

Strategies for common LA installation procedures

A personal view

ALA-LAs Workshop (Madrid, March 2025)

Vicente J. Ruiz Jurado (@vjrj), 2025/03/11

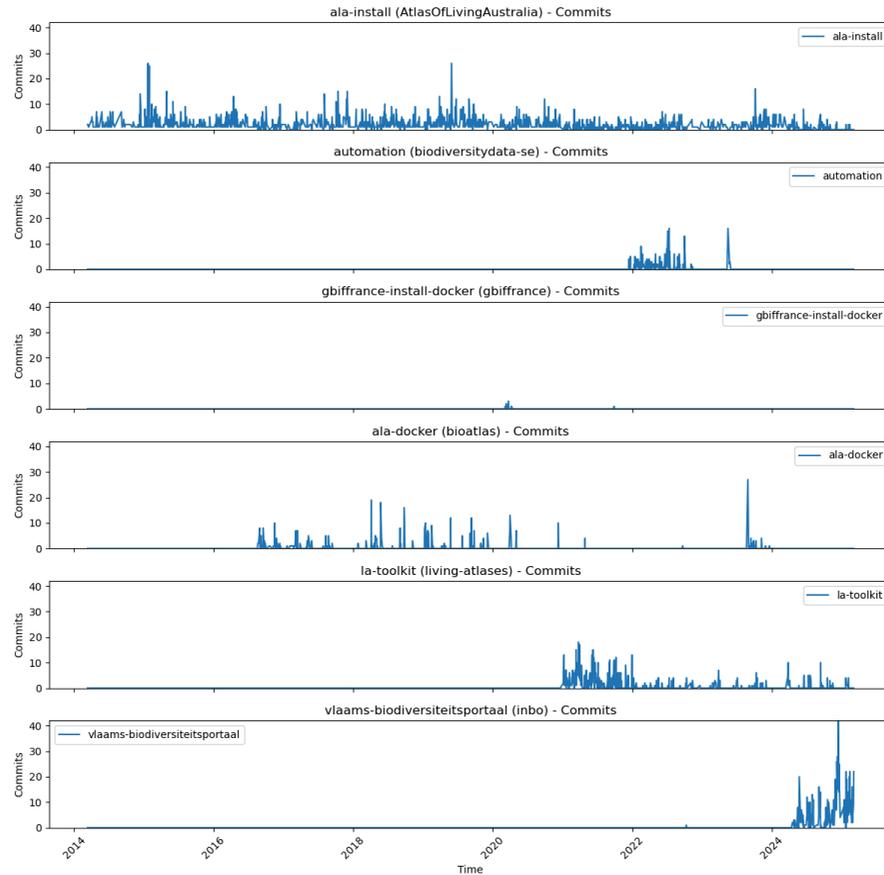
Comparing ALA Deployment Strategies

A look at different install approaches over the years

Past -knew- devops efforts and current situation

From	Code Repository	Based	Active	Out-of-the-box?	SubDomains & /pat
Atlas of Living Australia	ala-install	Ansible	2014-now	Yes	Yes
Atlas of Living Australia	ala-install (Docker)	Ansible + Helm + Docker Swarm	2023-now	Yes	Yes
Atlas of Living Australia	Alerts Userdetails Species Lists	AWS Cloudformation + AWS EKS + Helm	2024-now	No	Yes
LA Community	la-toolkit	Based on ala-install	≈2021-now	Yes	Yes
Sweden (BioAtlas)	ala-docker	Docker Swarm	≈2017-2023	No	? Unknown
Sweden (SBDI)	automation	Docker Swarm + Ansible	2022-2023	No	? Unknown
Sweden (SBDI)	Private repo	Docker Swarm + Ansible	Now	No	? Unknown
Brazil	vertigoala	Docker + K8s + helm + CI/CD with private gitlab	≈2019-now	No	? Unknown
France	gbiffrance-install-docker	Docker-Compose	2020-2021 Last commit 4y ago	No	No
Flanders	vlaams-biodiversiteitsportaal	Docker (public) + Terraform (priv)	2024-now	No	No

Activity of some of these repositories



Summary (IMHO)

- ALA is moving from `ala-install` to `helm` and `AWS` based solutions
- From the LA community we have different container based initiatives (not based in `ala-install`), but:
 - none work out-of-the box for newcomers,
 - so require of months (at least) of IT development for each new deployment (*just to start*)
 - and *not last*, a maintenance of these deployments over time, to keep it compatible with new ALA deployments

My opinion about ALA containers solution

In the past I was waiting for a docker official solution supported by ALA (like `ala-install`).

But now I see that their proposed solution, I think that does not fit the LA community needs:

- it does not work currently out-of-the-box (as `ala-install`),
- because it depends on AWS and similar,
- and `k8s` IMHO it's a solution more complex than what we need (compared with other options like Docker Swarm).

Duplicated efforts

But we have a lot of duplicated efforts.

For instance, lets think in the container images **registry** or the images build system:

- ALA has a container registry in AWS (private?)
- Sweden has is owns
- same Brazil,
- same France,
- the same Flanders...

do we each have to build and publish the images of all the ALA components ourselves?

I think that we need to **avoid duplicating these efforts**.

More general objective (IMHO)

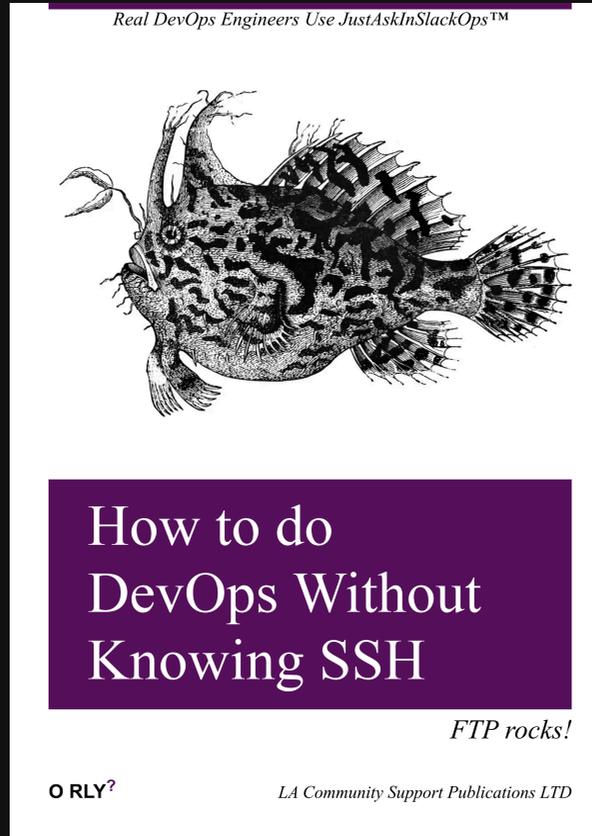
I think that we should join forces to:

- **collaborative maintain** a **common** container based **LA deployment solution** over time,
- that works **out-of-the box** (for newcomers and during updates),
- **without** service provider dependencies,
- that is **FLOSS** (Free/Libre/Open/Source/Software),
- and still we should try to **collaborate** as before **with ALA** in a common ground.

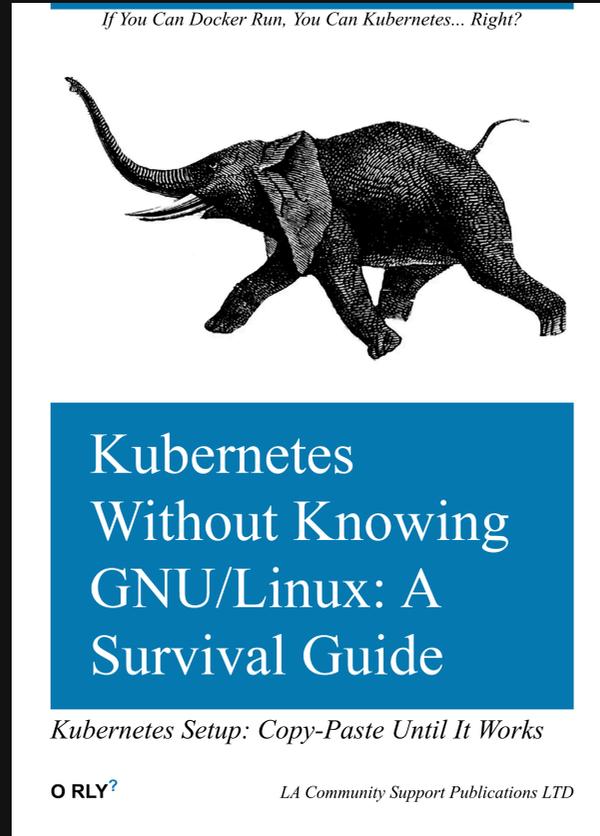
Docker swarm vs k8s / etc.

Let me explain my point of view with a bit of ... fun 🐼.

My bestseller book (6th Edition)!



My new book in pre-print





Major Companies That Use Kubernetes

- Airbnb
- Reddit
- Skyscanner
- Prezi
- Zalando
- Target
- Spotify
- Toyota
- Google (off course)

among others.

Fun Fact

The list above was compiled from [Kubernetes Failure Stories](#). 😊

It's just a compilation of daily use failure stories for people dealing with Kubernetes operations to learn from others.

My Learns

- What I learn is that:
 - we are not so big companies
 - with big IT departments
 - with so big requirements,
 - or money.

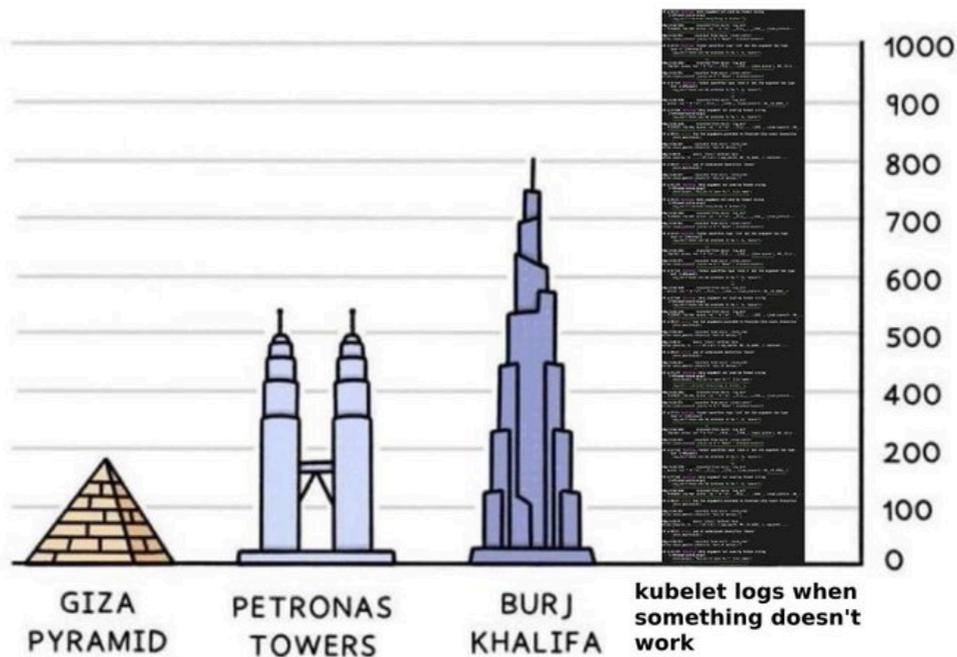
Kubernetes engineers when they don't have a kitchen towel



Some more k8s memes... (sorry!)

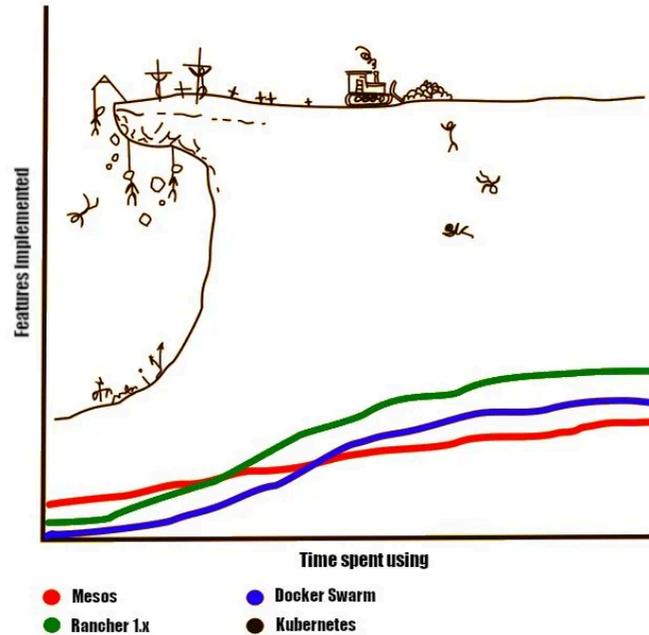


Debugging Kubernetes



k8s Learning curve (I)

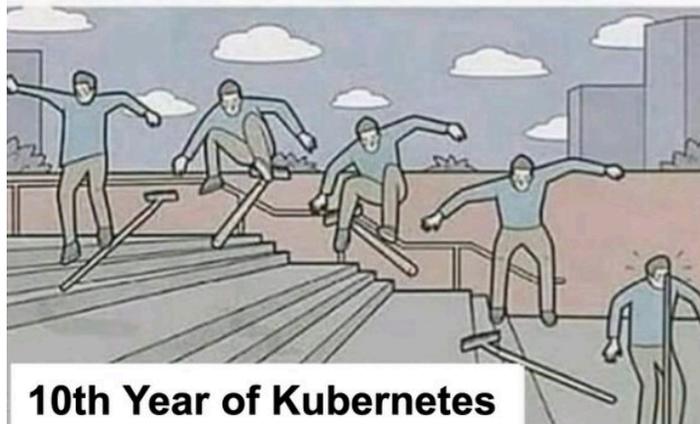
Learning curves of some Container Orchestration Engines



k8s Learning curve (II)



1st Year of Kubernetes



10th Year of Kubernetes

My (impossible) goal during these years

To setting up an LA portal with just one simple command (`apt install la` or `docker compose up`)

At some point, I really believed that this could be possible. (Yes, a bit utopian 🦄)



Reality Check

- Once upon a time, installing a LA demo was quick and easy.
- Now, dependencies keep growing, requirements are piling up, and what was once simple has turned into a complex labyrinth.
- That utopia has been moving further away over time.

Flanders

The recent work of the Flanders Teams is remarkable and quite close to that goal, (kudos 🍻 Stefan Van Dyck et al.)

Things I like:

- Modern versions (including `pipelines`)
- test and prod based on the same repo
- single `docker-compose`
- autobuild of images in the `Dockerfiles`
- ...

Flanders (II)

Things I miss:

- **subdomains** support,
- **internationalization** for any language,
- **optional** ALA modules,
- **docker swarm** support,
- **configurations** in sync with ALA ones,
- **easy setup for others** LA portals,
- better software **versions** and **tags** selection,
- public Docker **registry**,
- missing `ala-install` tasks so requires extra **manual** configurations

Can we continue their work in order to fill that gaps?

I can do the same checklist with other deployment initiatives.

Flanders vs our TDWG 2023 POC

Our 2023 POC had other gaps but not that ones:

- **subdomains** support, ✓
- **internationalization** for any language, ✓
- **optional** ALA modules, ✓
- **docker swarm** support, ✓
- **configurations** in sync with ALA ones, ✓
- **easy setup for others** LA portals, ✓
- better software **versions** and **tags** selection, ✓
- public Docker **registry**, ✓
- missing `ala-install` tasks so requires extra **manual** configurations ✓

One Problem Multiplied by 20 Is Still 20 Problems

No problem should ever have to be solved twice. **Eric S. Raymond**

Every time I solve a problem, and think about the effort it did,

- I also think about the rest of the community and how to prevent others from having to go through the same struggle.
- That's why I try to document the solution, submit a patch, or automate part of the installation process.

It makes no sense for each of us to repeat the same effort that, for instance, **Stefan, ALA, or others before** already did to deploy a portal with containers—and more importantly, to maintain it over the years.



Containers: Solution or More Complexity?

- In theory, containerization should make things easier.
- In practice, it sometimes feels like just another layer of problems.
- Different container solutions might improve our LA deployments... or just make it worse.
- I still think that Docker Swarm or similar fits better for the LA Community than k8s solutions.
- (...)

KUBERNETES “SEMANTIC” VERSIONING EXPLAINED





Containers: Solution or More Complexity? (cont)

K8s makes sense for large-scale applications, but does it really fit our LA community's needs?

I think we require a simpler, community-driven alternative instead of adopting a tool designed for hyperscalers.

Final Thoughts (I)

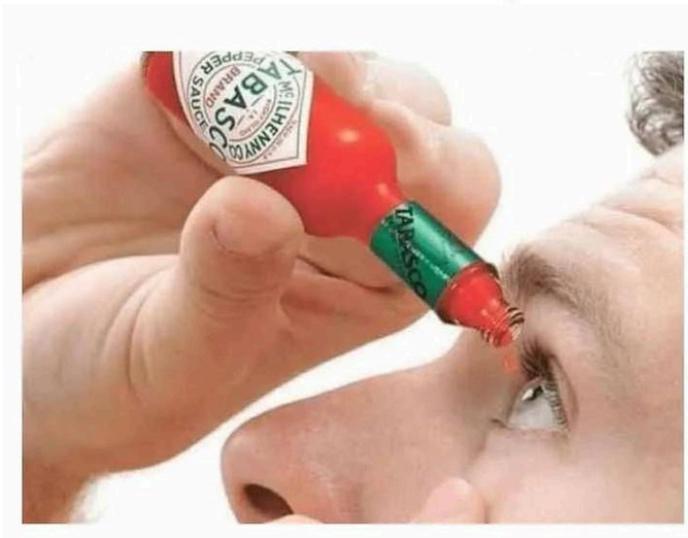
Let me repeat what we think we should do:

-  **Collaborate** to maintain a **common container-based LA deployment**
-  Ensure it **works out-of-the-box** for newcomers
-  **Avoid vendor lock-in** (optional AWS or other cloud dependencies)
-  Keep it **FLOSS (Free/Libre/Open Source Software)**
-  Keep it **Cheap**
-  **Stay compatible** with ALA while keeping our collaboration

🚩 Final Thoughts (II)

The dream of a single-command deploy for ALA might be a fantasy... but at least we have memes to cope with the pain.

Looking at the average Helm chart template





Thank You!

Questions? Comments?

Let's discuss!