An infrastructure for the reproducible scientific workflows.

Dr. Arman Khalatyan

Researcher at eScience/Supercomputing/IT

Leibniz-Institut für Astrophysik Potsdam (AIP), Germany

17.09.2021

Scientific life

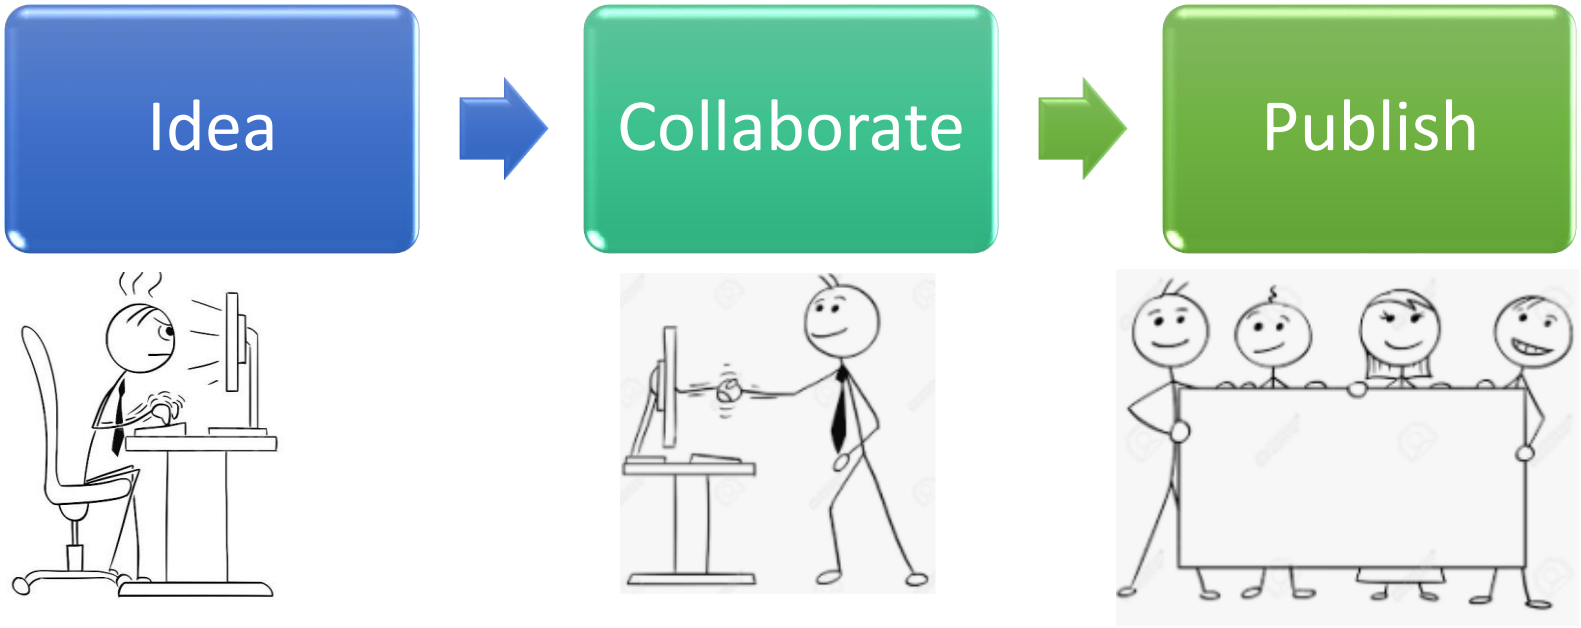

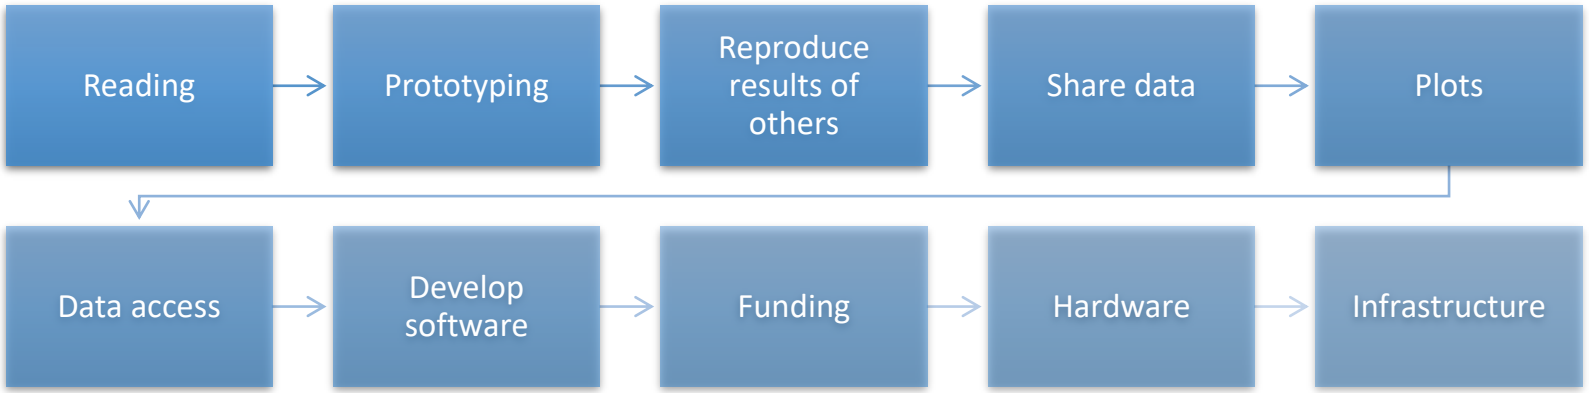

Scientific life (top to down)

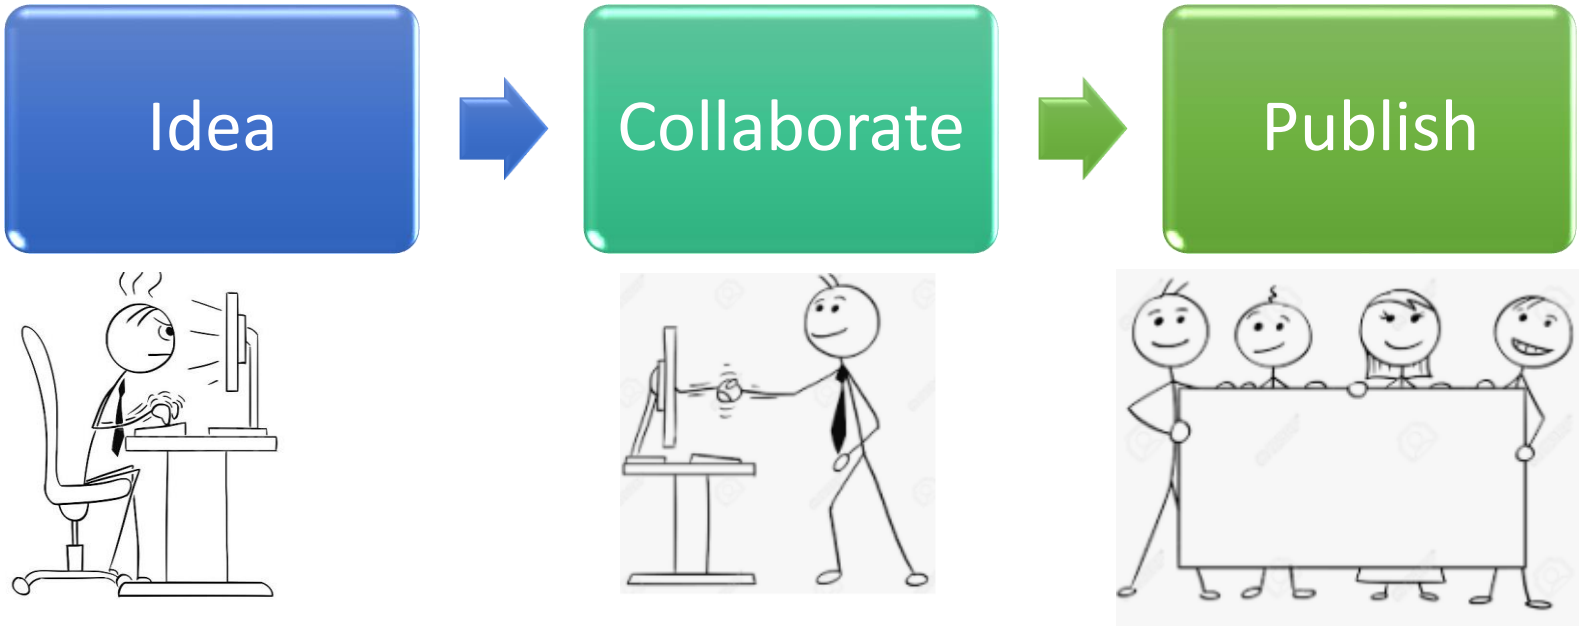

Scientific products

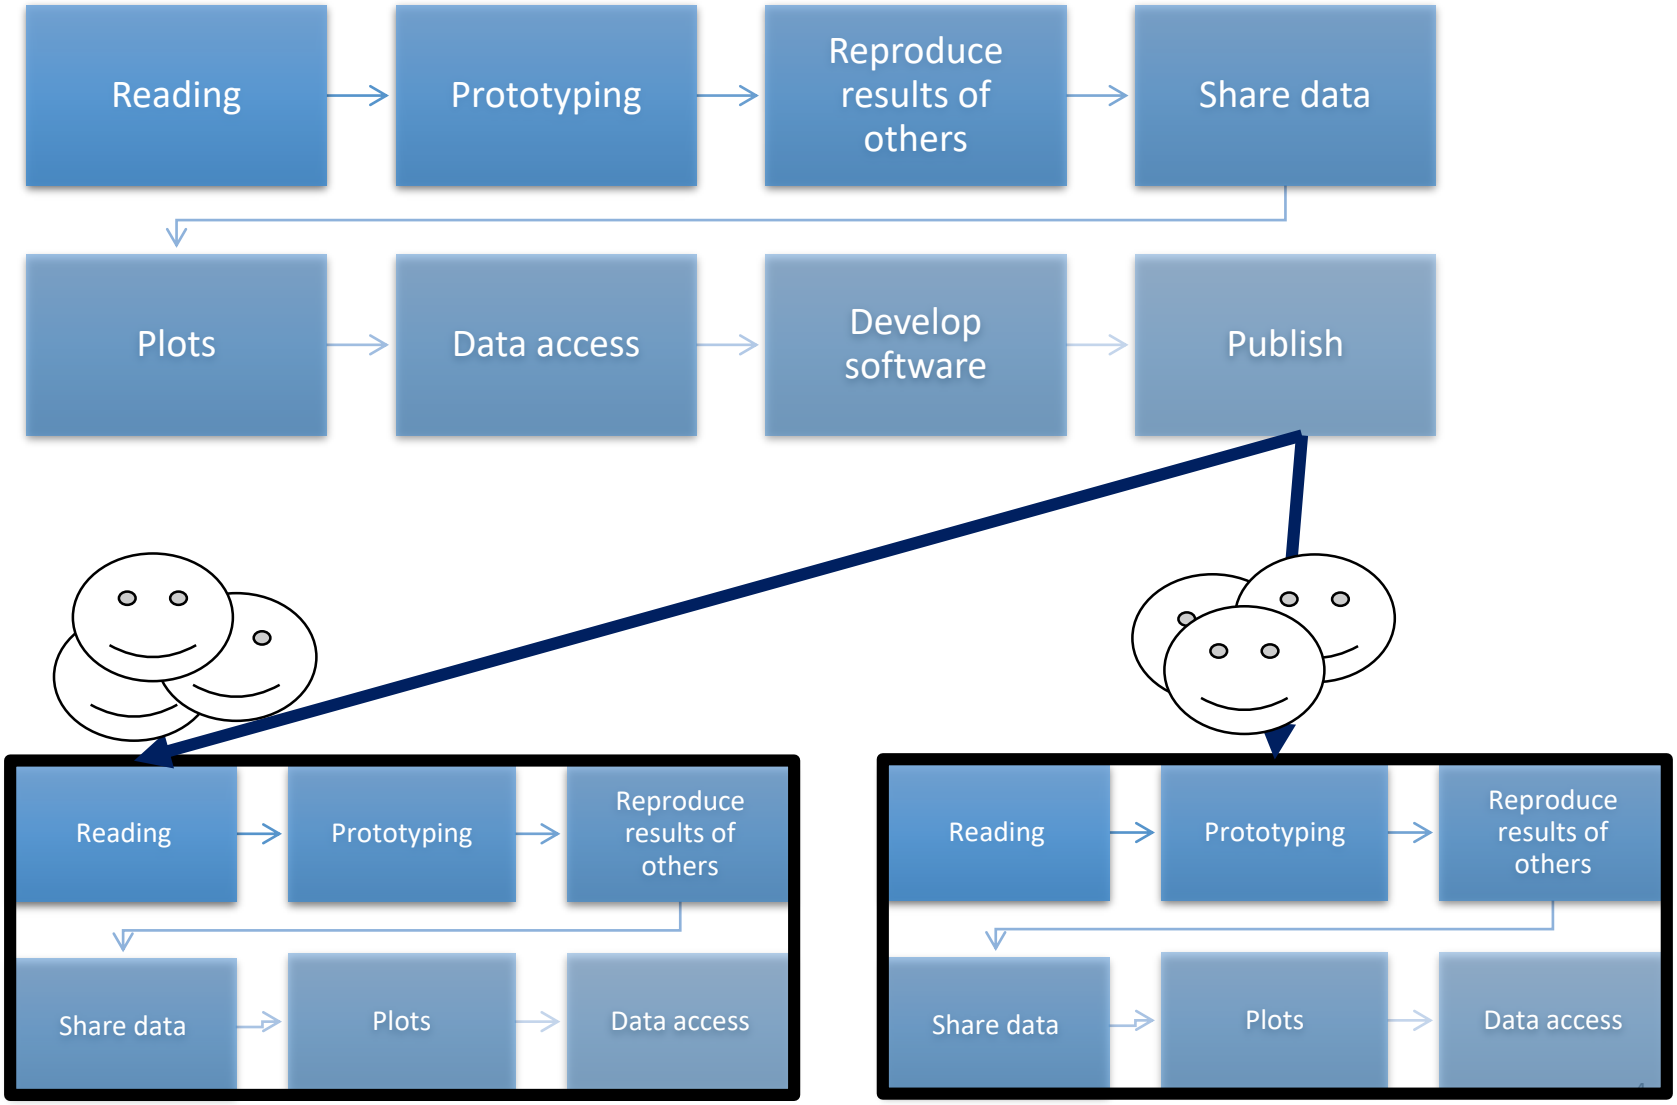

Research technology and infrastructure

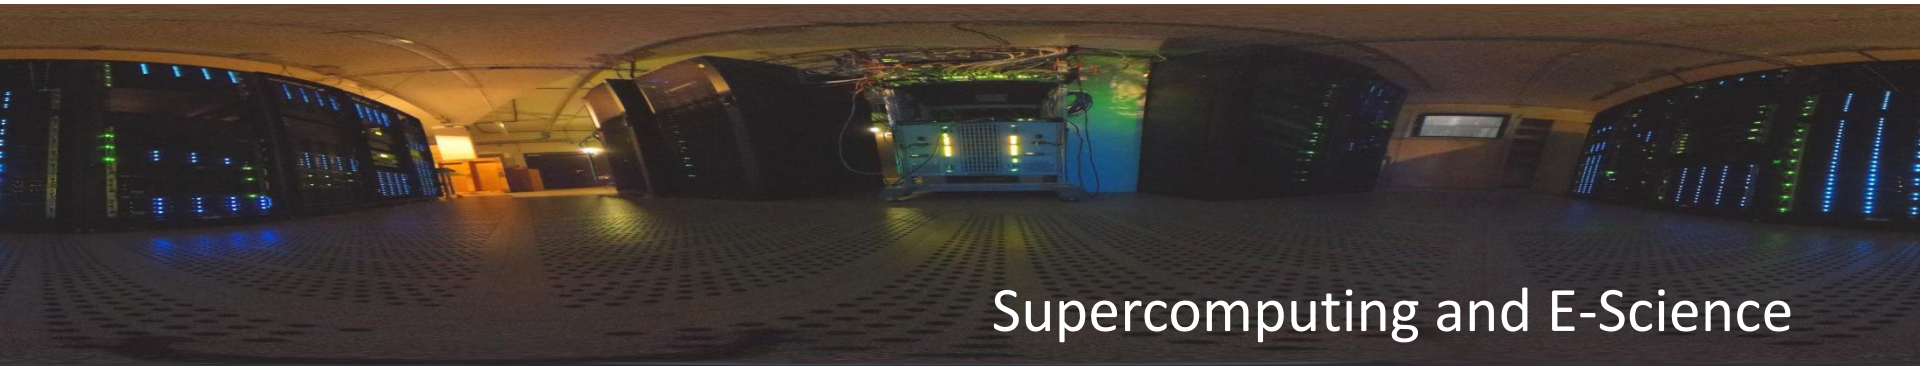

Compute Cluster

Simulations in

Cosmology &

Magneto-Hydro-Dynamic

Leibniz: 1700 Cores

Newton: >3000 Cores

GPUs: >8(NEW)

Storage / Archives / Database

Simulation and observation data

ca. >**5** **P**etabyte Storage,

Virtual Research Environment for CLUES, MUSE

Databases: CosmoSim, RAVE, Gaia,Plates Archives

EUROPÄISCHE UNION
Europäischer Fonds für Regionale Entwicklung

LAND BRANDENBURG

Down to top problem

Hardware

Compute nodes

Storage

Software

Rocky Linux

oVirt

docker

RANCHER

HPC Cloud friendly

lustre™

Cloud friendly

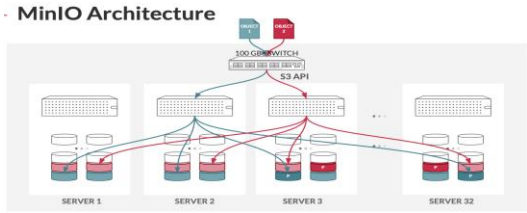

Can we obscure the HW and Infra form the users?

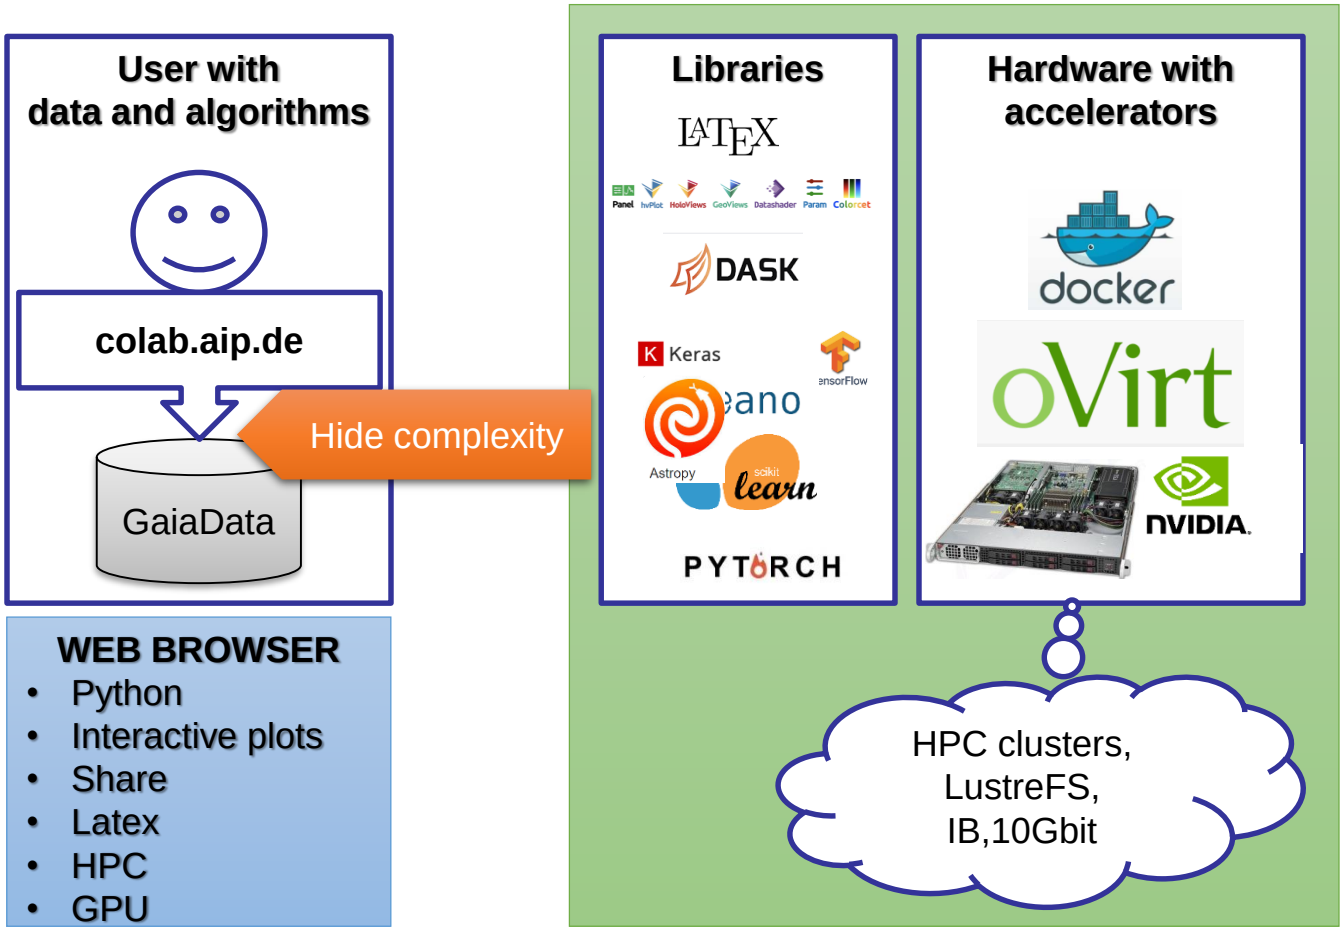

The whole complexity is obscured from the users

Users want the all tools in one place

- Data+LaTex+Code
- Collaborators to share
- Article versions
- Cluster access
- Easy publishing for the demo notebooks

Possible solution:

- dask+Kubernetes
- CoCalc project

CoCalc is a web-based cloud computing and course management platform for computational mathematics. Part of the Sage project, it supports editing of Sage worksheets, LaTeX documents and Jupyter notebooks.

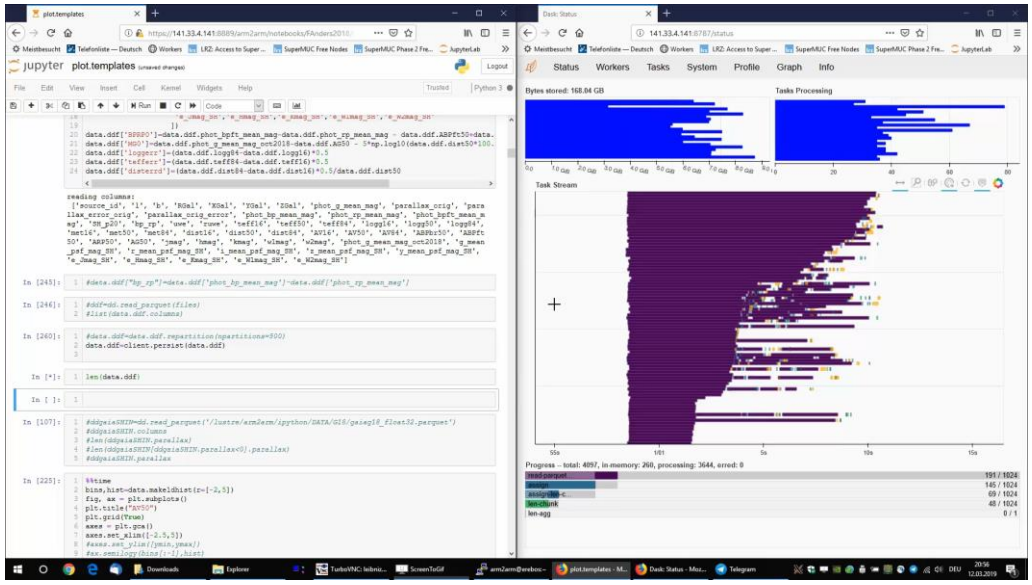

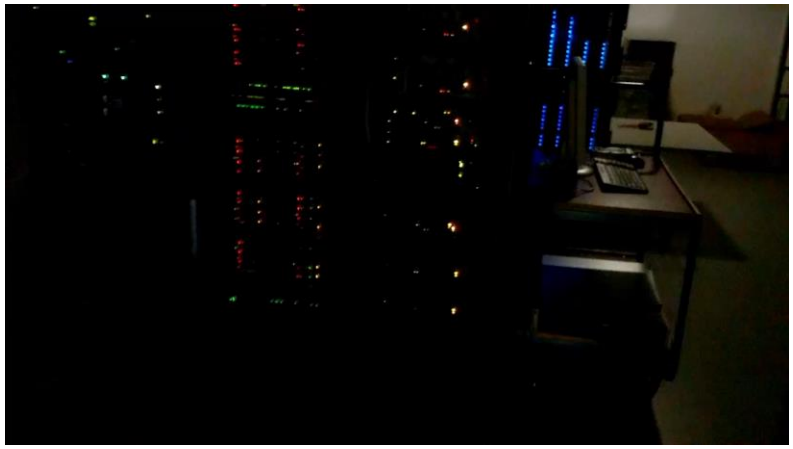

Local workflow:
”..but it works on my laptop”

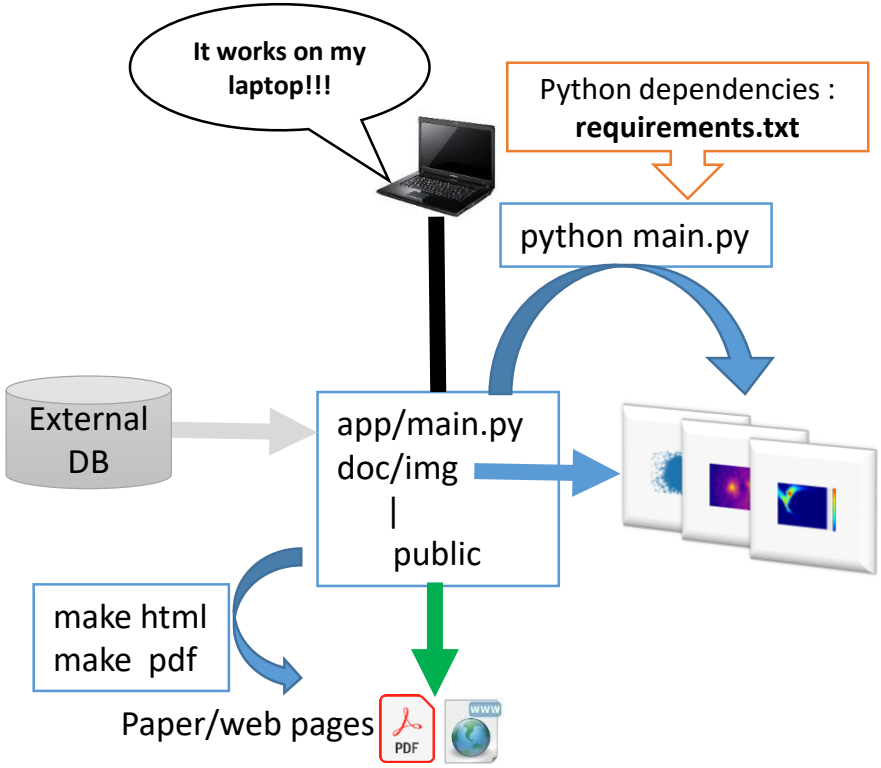

Stages

- RUN
  - edit **main.py**
  - run: **main.py**
  - inspect images
- PAGES
  - edit **doc/index.rst** … etc.
  - compile
  - Inspect pdf/html
  - Publish

Global reproducible workflow: ”…works everywhere”

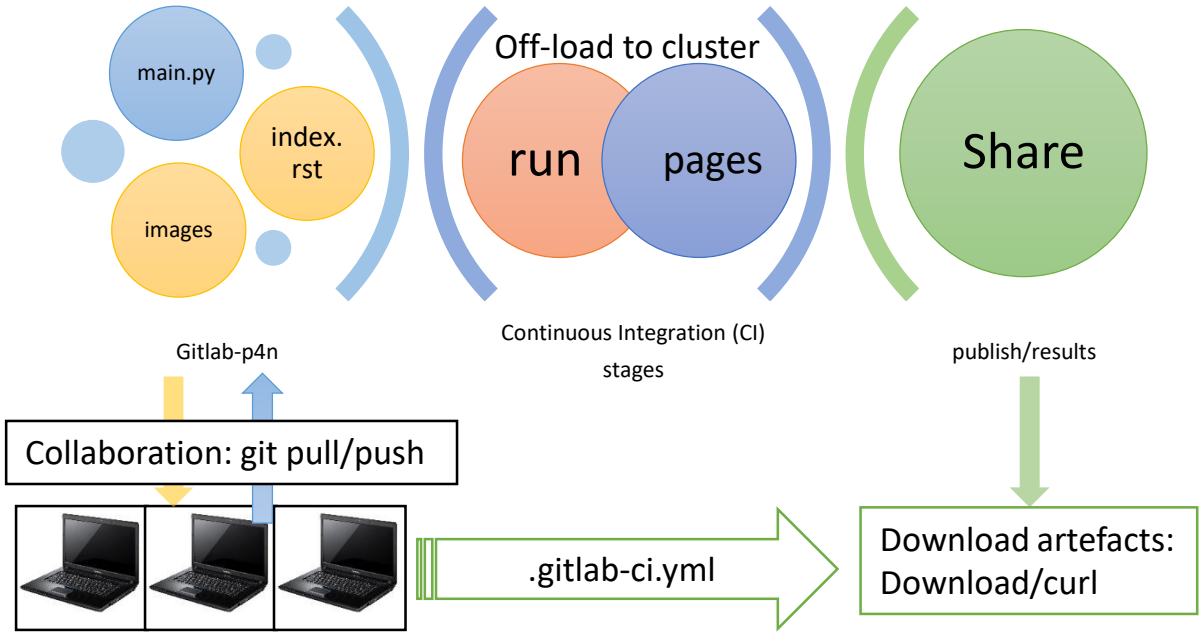

Stages

- RUN
  - [illegible]
  - [illegible]
  - [illegible]
- PAGE
  - [illegible] … etc.
  - compile
  - Inspect pdf/html
  - Publish

Stages

- EDIT
  - **main.py**
  - **index.rst**
- PUBLISH
  - **artefacts**

Global Workflow of StarHorse team

A Bayesian code to estimate the photo-astrometric distances, extinctions, and astrophysical parameters for Gaia DR2 stars
F.Anders et al. (2019)

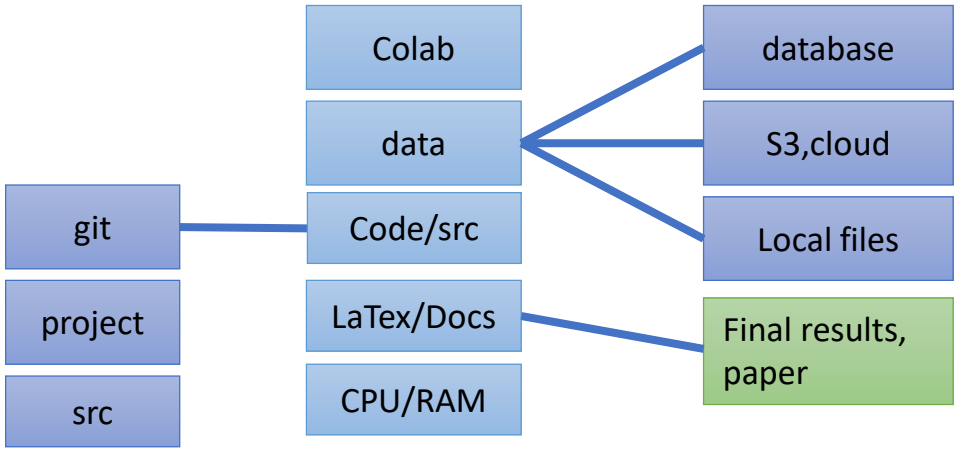

Getting the data

Get the list of the files: `wget --no-check-certificate http://data.aip.de/data/starhorse/fits/list-fits.txt`

Download the data: wget --no-check-certificate -i list-fits.txt

- Access examples: starhorse_db
- cmd_from_db: launch binder Launch on Google Colab
- cmd_from_db_chunking: launch binder Launch on Google Colab

https://data.aip.de/projects/starhorse2019.html

An interactive StarHorse data access example (github+gaia.aip.de+binder or googlecloud)

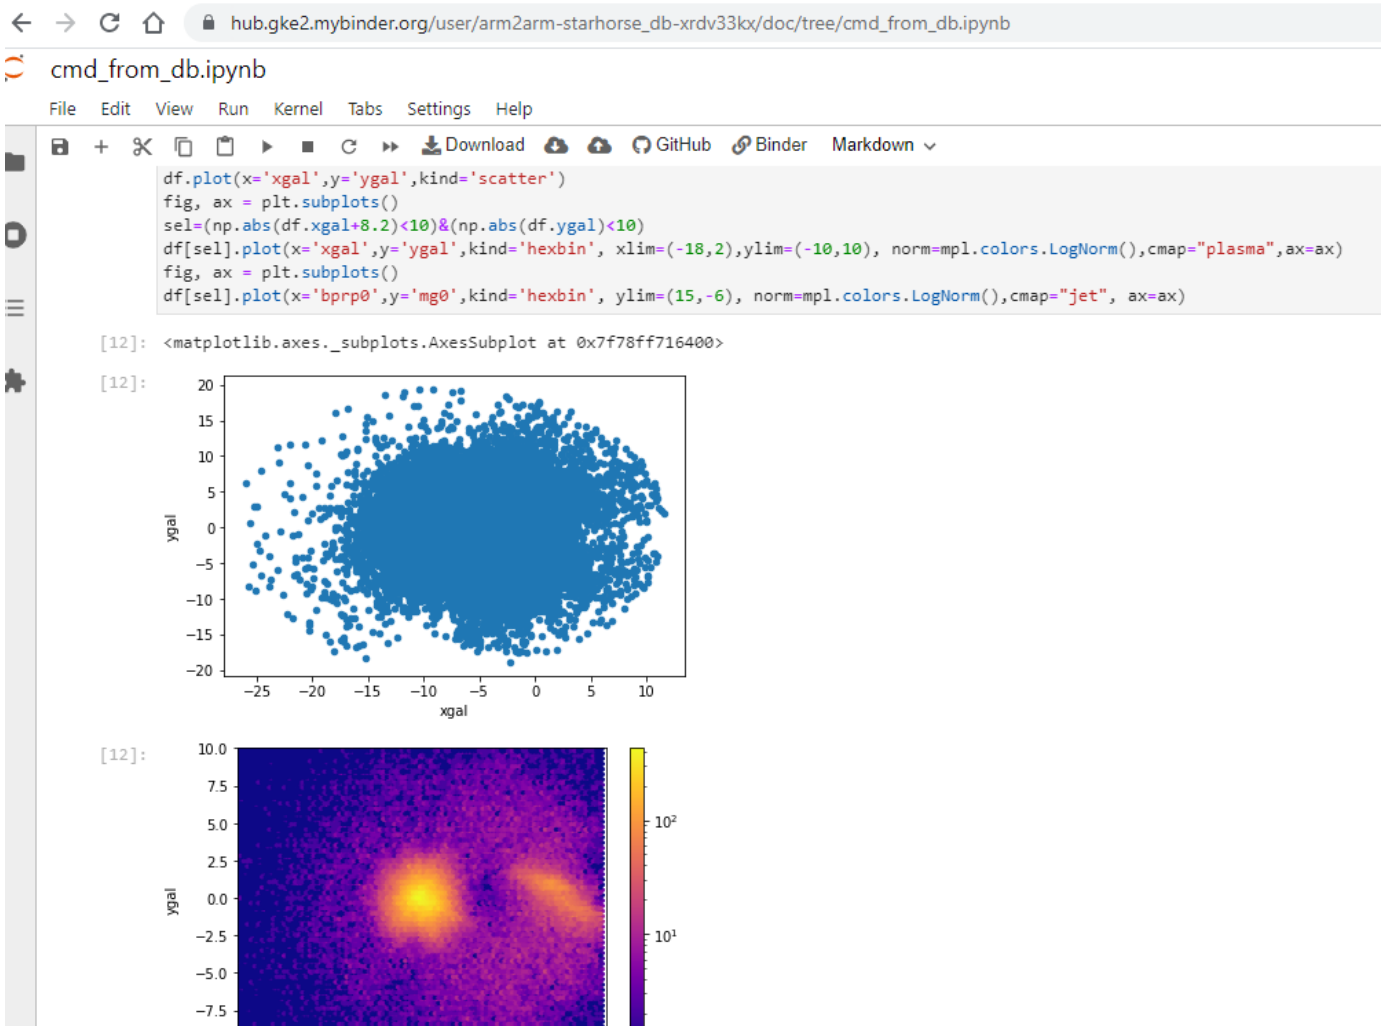

What is **docker** as an astronomer? A streamlit example app.

- **What if I w**
  environme
  collaborato
- **docker run**

```
[root@nnewl1 ~]# podman run arm2
✔ docker.io/arm2arm/astro-aiko:l
Trying to pull docker.io/arm2arm
Getting image source signatures
Copying blob ee776f0e36af done
Copying blob 11a88e764313 done
Copying blob 989e6b19a265 done
Copying blob e9afc4f90ab0 done
Copying blob af14b6c2f878 done
Copying blob 5573c4b30949 done
Copying blob 0f6c875a5369 done
Copying blob a2726e1e2770 done
Copying blob 029629755bc9 done
Copying blob b7ed565dbb37 done
Copying blob d42ad7cae96b done
Copying blob 2296d57805b5 done
Copying blob cff49450c17f done
Copying blob 6ee6bc466e7c [-----
Copying blob cd13f0cc7f40 done
Copying blob 7f3da3d220c6 done
Copying blob b41173953f00 done
```

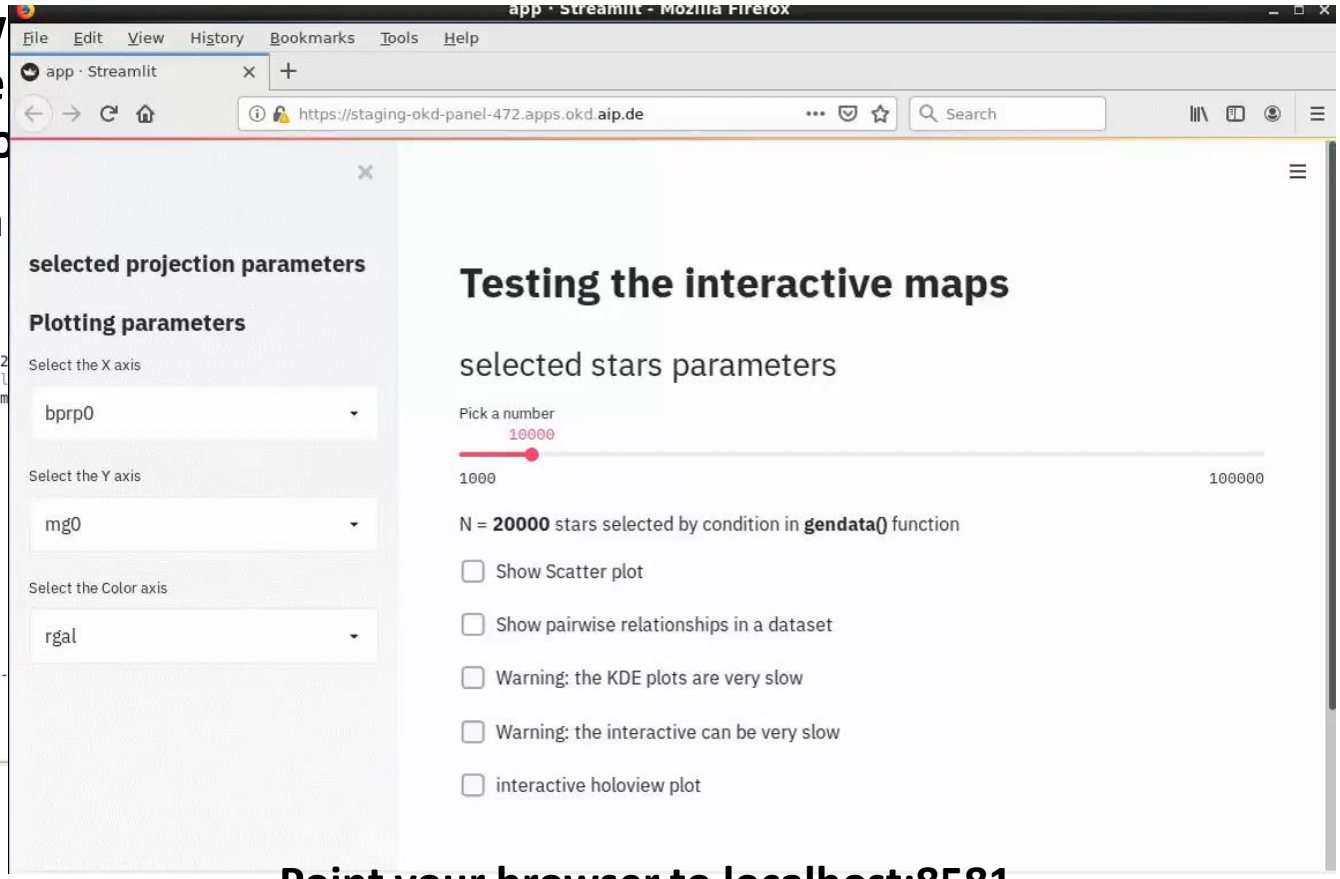

Point your browser to localhost:8581
You have an interactive plots

Docker

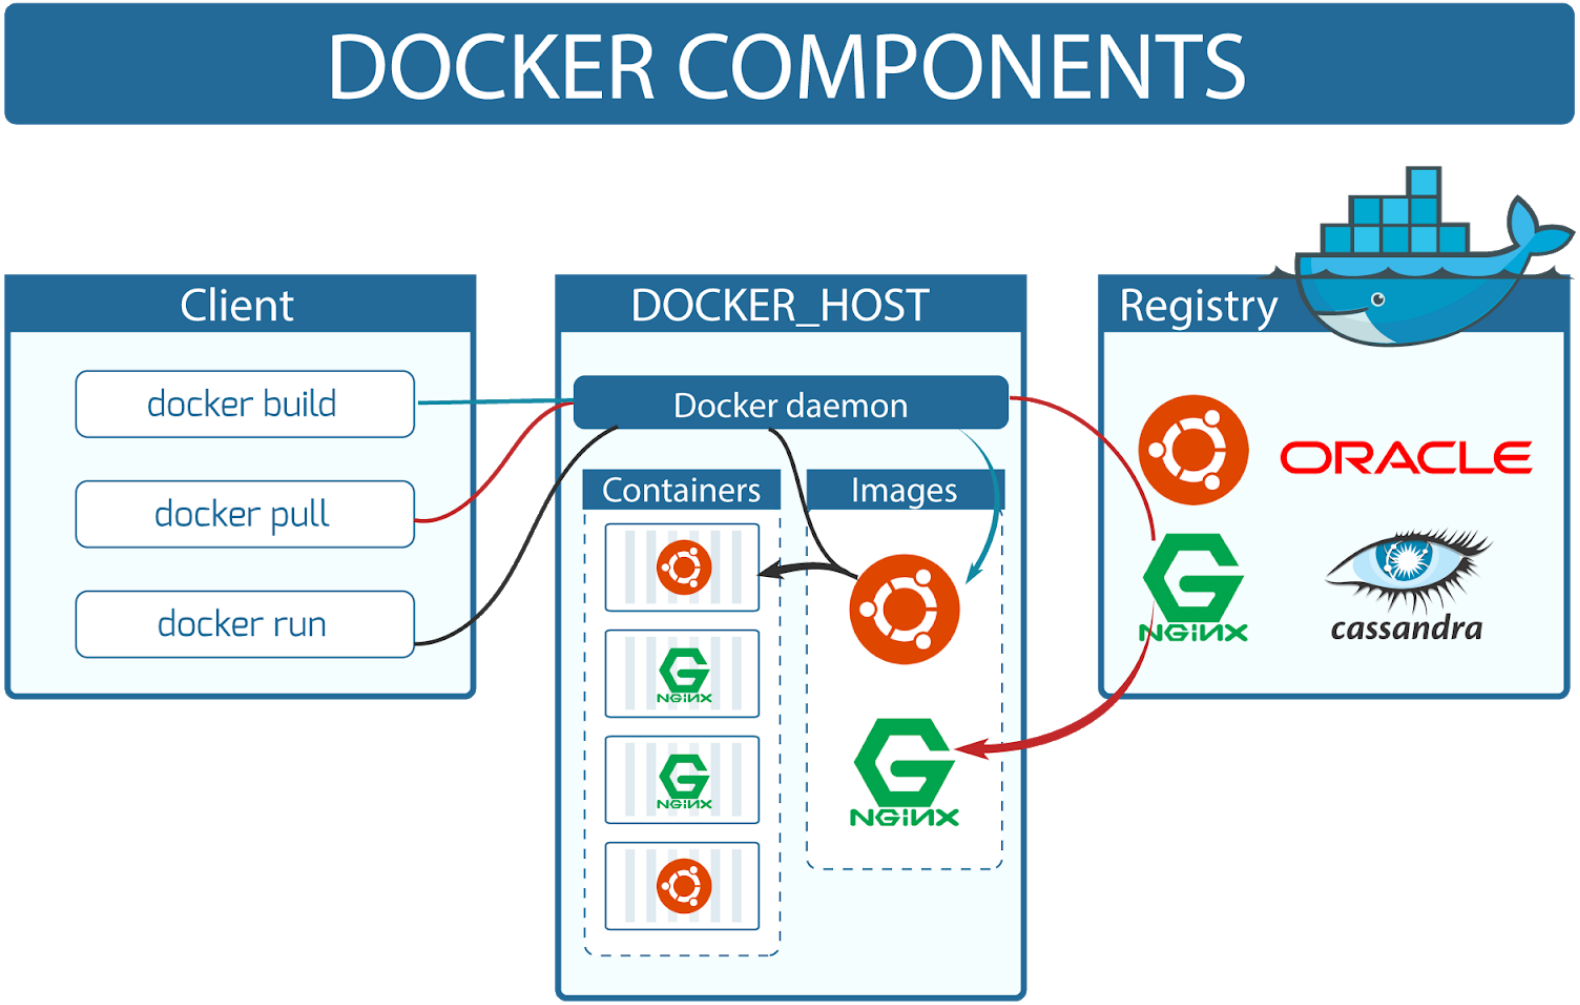
DOCKER COMPONENTS
Client
docker build
docker pull
docker run
DOCKER_HOST
Docker daemon
Containers
Images
Registry
ORACLE
NGINX
cassandra

What about kubernetes?

- in Astrophysics infrastructur
  still in the same stage as "**D**
  Inc." was in 2014.

Why?

- It was complex
- Rapid development in the Ir
- No LTS

Situation is matured in 2

- Because of https://www.cnc

CLOUD NATIVE COMPUTING FOUNDATION

- We are ready to adopt some
  from industry into to scienti

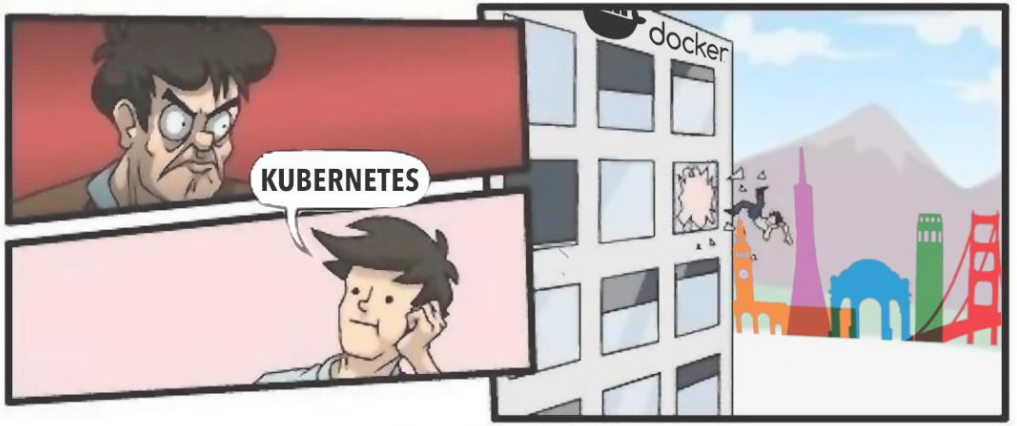

Kubernetes

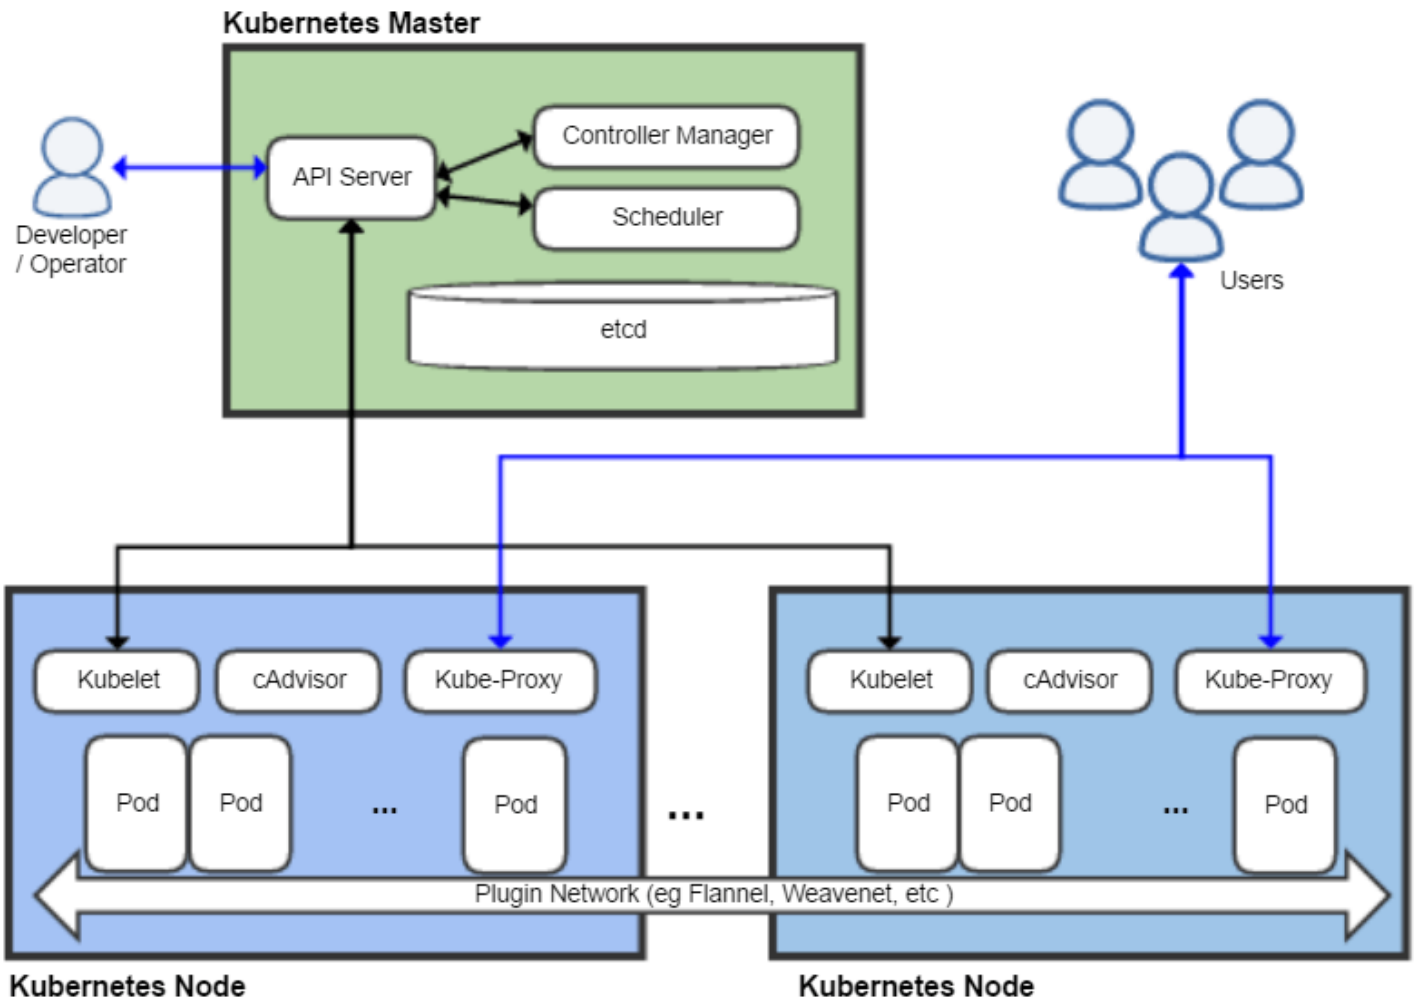
Kubernetes Master
Developer
/ Operator
API Server
Controller Manager
Scheduler
etcd
Users
Kubelet
cAdvisor
Kube-Proxy
Pod
Pod
...
Pod
...
Kubelet
cAdvisor
Kube-Proxy
Pod
Pod
...
Pod
Plugin Network (eg Flannel, Weavenet, etc)
Kubernetes Node
Kubernetes Node

Microservices: Reproducible science

Use Cases at AIP

- **Colab.aip.de**
  - Quotas
  - Project isolation
- **Data analysis pipelines with versioning**
  - Reproducible science
  - Pipeline versioning
    - GaiaDR1,2,3
    - RAVEDR1-6
    - StarHorse-18,19,20
- **Publish papers with interactive plots**
  - like binder
  - Example: distill.pub by google
- **Dynamically Scalable webpages**
- **gitlabs @ aip:** CI integration

Pros:

- ***Direct GIT integration***
- ***Scalability***
- Modularity
- Distributed development
- Integration
- Save resources/power

Concerns:

- Complexity
- Design
- ***Testing, debugging***
- ***Inter-service call latency***

Upcoming colab.aip.de vision

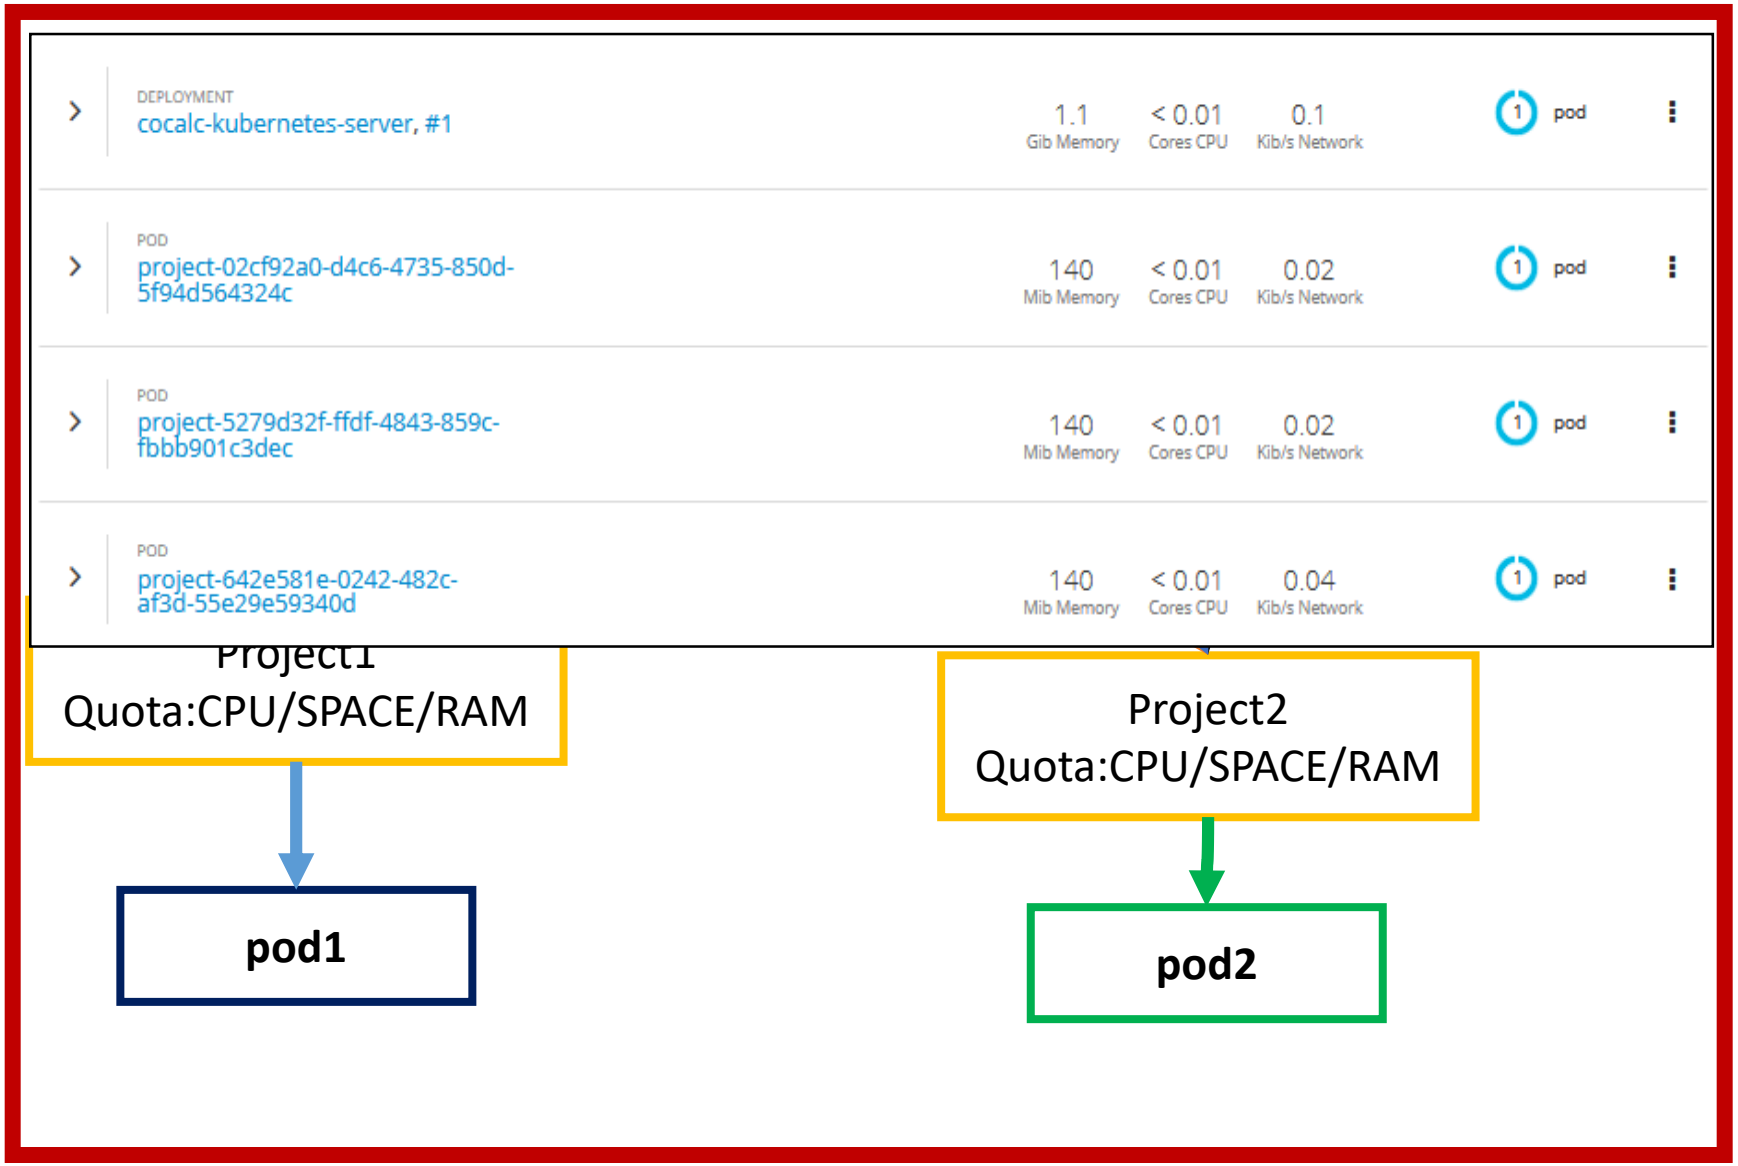

What's next?

- Large Kubernetes clusters deployments on demand via PUNCH4NFDI
- Parallel disk IO to deal with TBs of data
  - Connection with Lustrefs, S3
- Easy to maintain
- Clone service on demand
- Reuse/reshare unused resources
- Snapshots for versioning of the data-pipelines
- 24/7 HA
- Special hardware share: GPU, SSD, NVME
- **LTS Container Registry for each publication**
- **Share docker containers with publications**