

A robotic hand with a black and blue gripper, labeled 'ROBOTIQ', is reaching towards a human hand with red nail polish and a gold ring. The background is a blurred, colorful bokeh of light circles in shades of blue, purple, and orange. The overall scene suggests a human-robot interaction or collaboration.

**HUN  
REN**



**SZTAKI**

## Reference architecture concept on HUN-REN Cloud

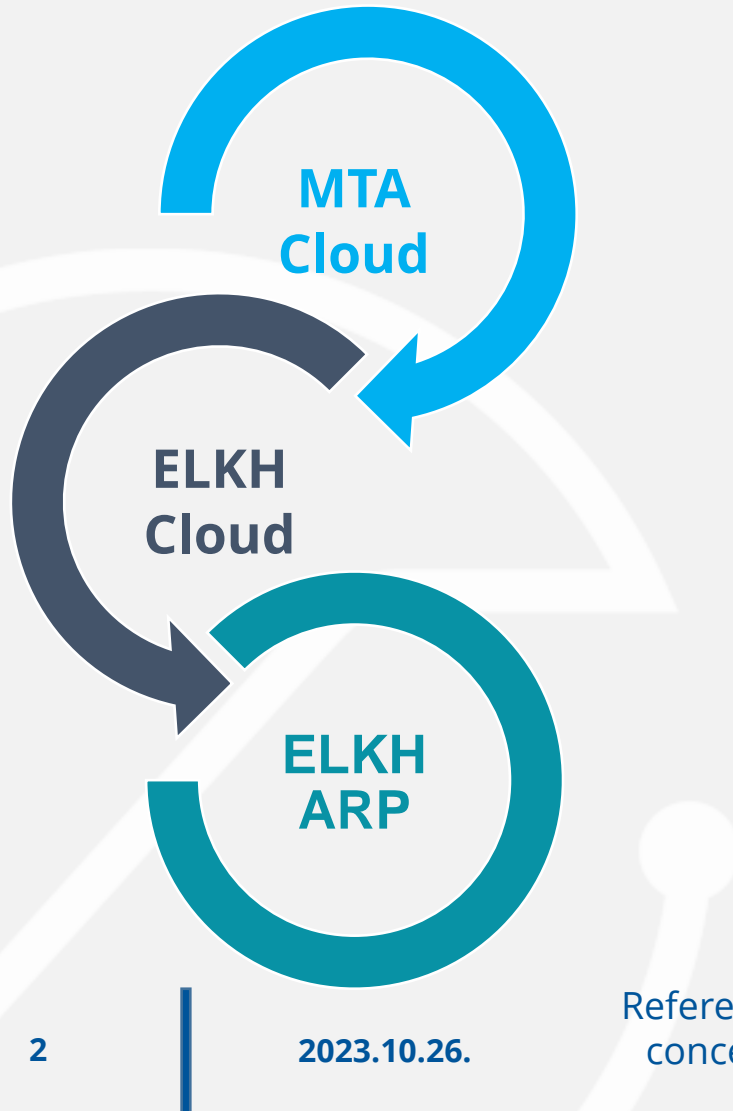
Farkas Attila

HUN-REN SZTAKI LPDS, research associate

**HUN-REN**  
Hungarian Research Network

# Phases of building a research cloud in Hungary

- **MTA Cloud project** (2015 – 2019)
  - Supported by the Hungarian Academy of Sciences (MTA)
  - The same year when the idea of **EOSC initiative** was born
- **ELKH Cloud project** (2020–2022)
  - Supported by the Eötvös Loránd Research Network (ELKH)
  - After the **re-organisation** of the network of research institutes (similarly to the Łukasiewicz Research Network in Poland)
- **ELKH ARP project** (2021–2023)
  - Supported by the Eötvös Loránd Research Network (ELKH)
  - Towards federated **research data** infrastructure



# Capacity comparison and user community statistics

	MTA Cloud	HUN-REN Cloud
<b>vCPU</b>	1356	5900
<b>vGPU</b>	12	512*
<b>RAM</b>	3.25 TB	28 TB
<b>SSD</b>	0 TB	338 TB
<b>HDD</b>	527 TB	1250 TB
<b>Network speed</b>	10 GBPS	100 Gbps

\*theoretical maximum 2060

Aggregated number of supported projects



# Reference architecture concept

- A kind of a PaaS solution
- Reference architectures contain all the necessary building blocks for building a complex software system on cloud-based resources
- Satisfied non-functional requirements:
  - Scalability
  - Availability
  - Configurability
  - Safety
- Provided by well-defined descriptor file for different Infrastructure as Code tools

# Usage of reference architectures

## Necessary steps from the user:

**Step 0 - Preparation** (HUN-REN Cloud project, creation of empty VM if necessary)

**Step 1 - Terraform and Ansible installation**

**Step 2 - Download the descriptor files**

HUN-REN Cloud webpage

**Step 3 - Create firewall rules (optional)**

HUN-REN Cloud OpenStack dashboard or descriptor file

**Step 4 - Personalizing descriptor files**

**Step 5 - Terraform initialization**

```
$ terraform init
```

**Step 6 - Infrastructure deployment**

```
$ terraform apply (--auto-approve)
```

**Step 7 - Infrastructure usage**

**Step 8 - Infrastructure destroy**


```
$ terraform destroy (--auto-approve)
```

Step 0 - 1. Only for the first time.

Step 2 - 4. Only 1 time per reference architecture

Step 5 - 8.

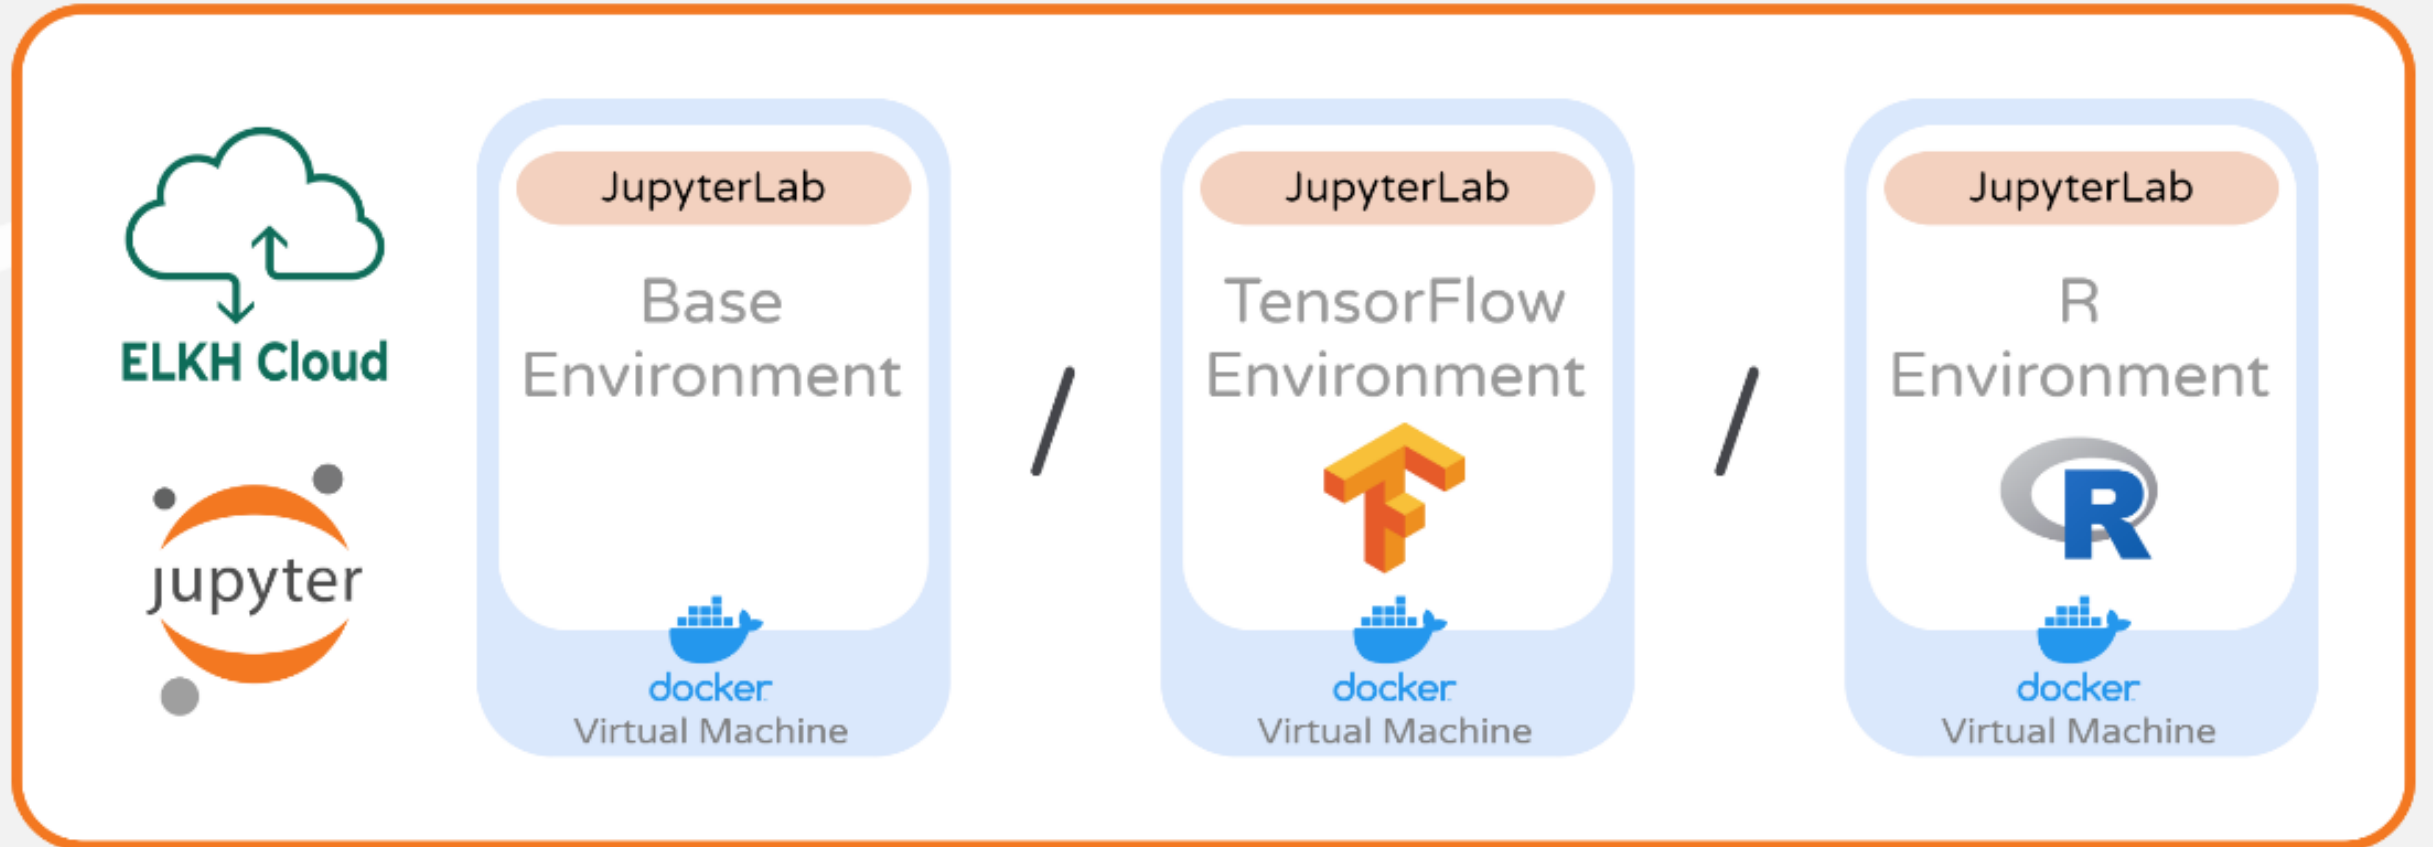
1 command



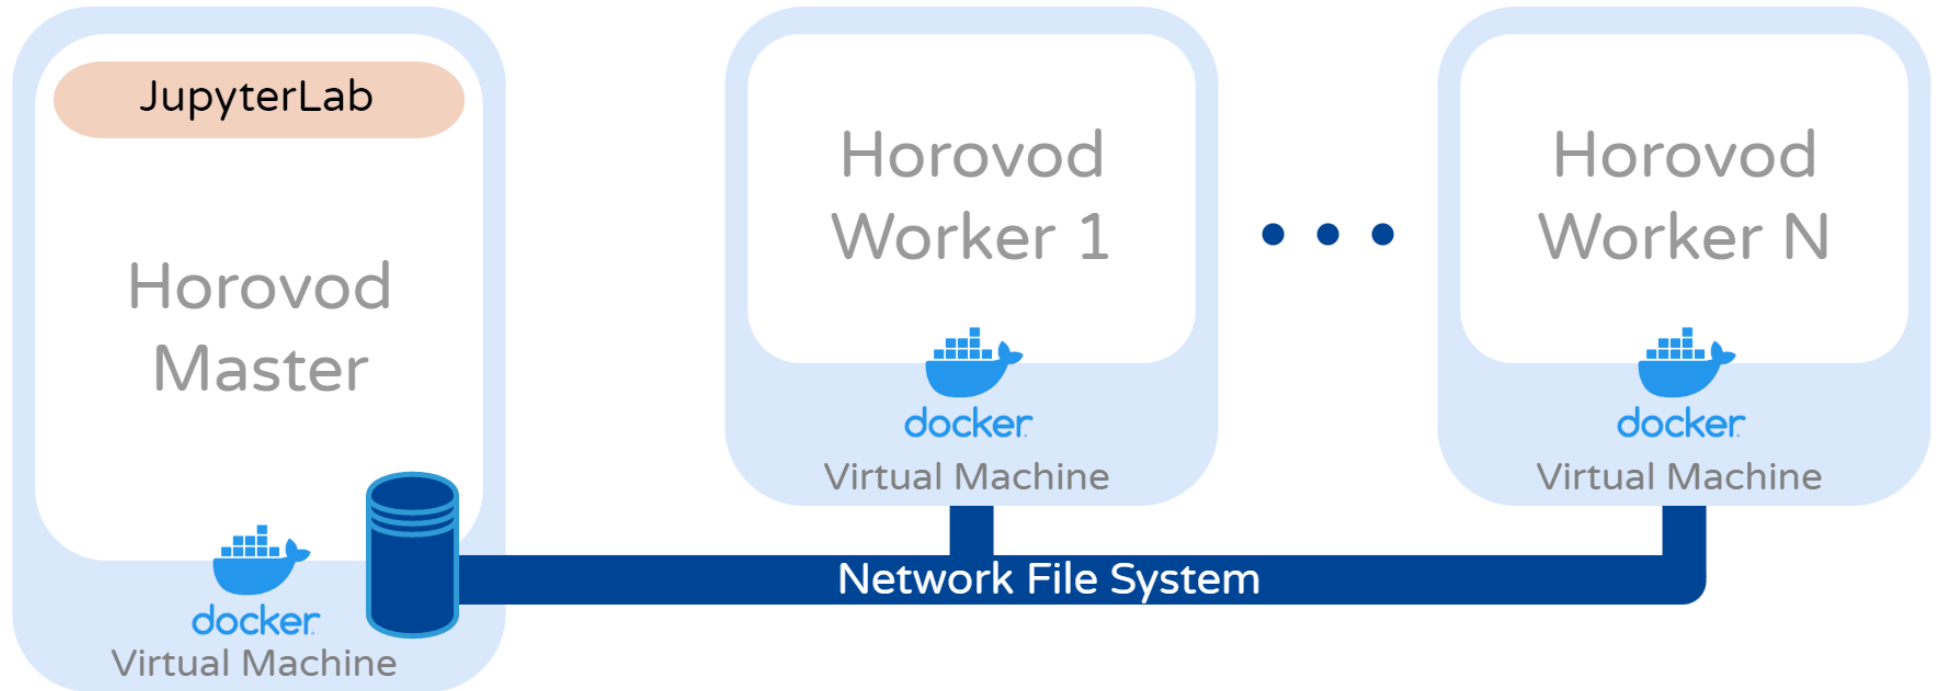
# Reference architecture types

- Development environment
- Machine learning support
- Big data and IoT platforms
- Container platforms
- Workload management
- Quantum resource support

# JupyterLab reference architecture

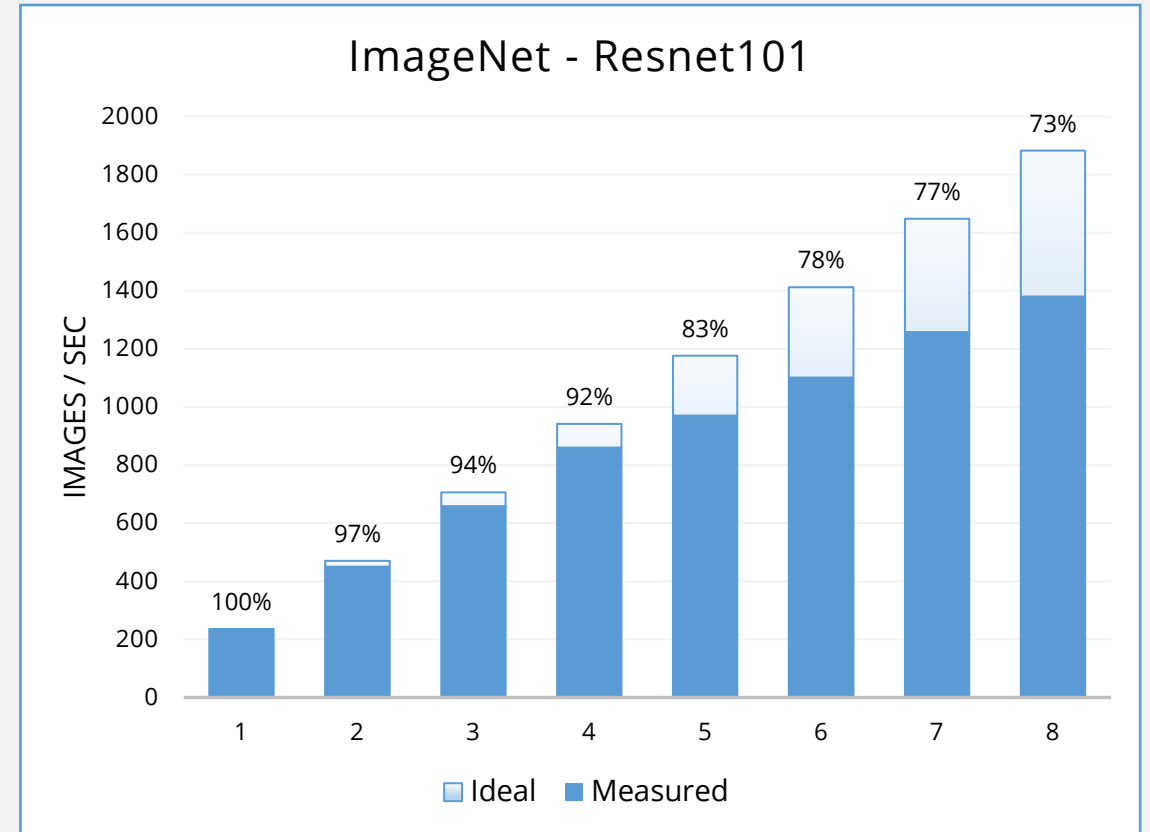
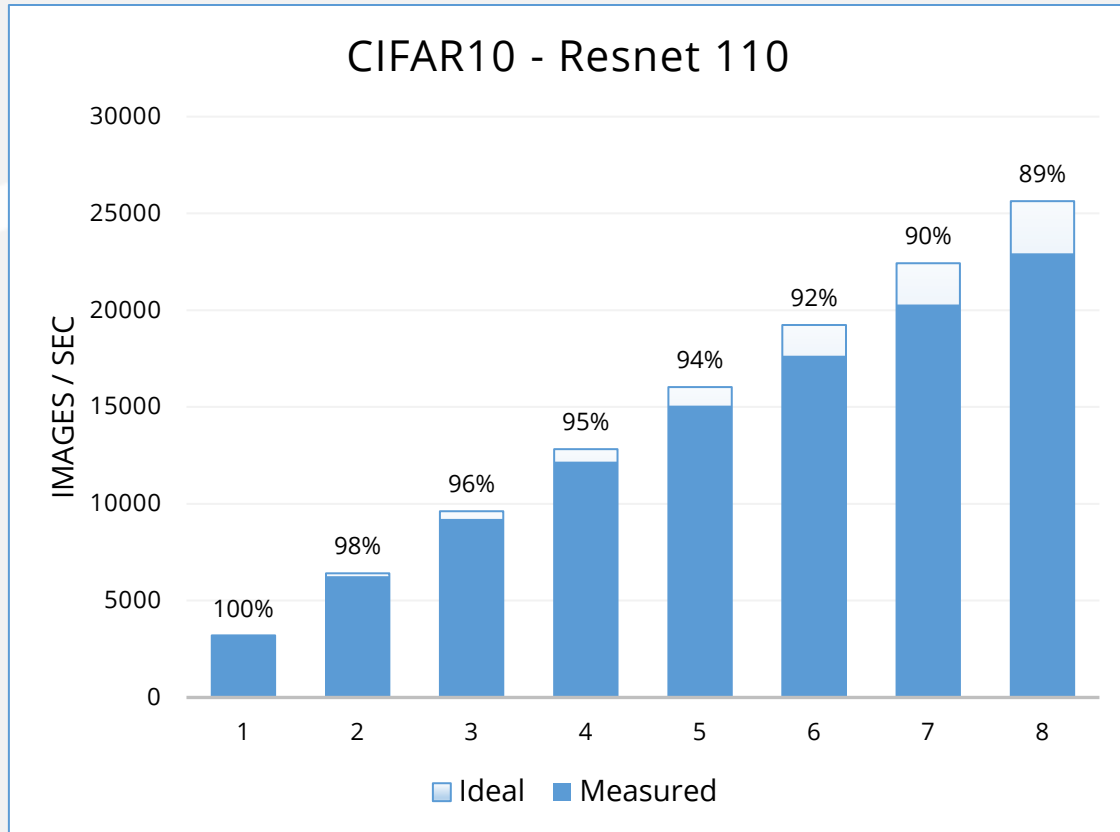


# Horovod reference architecture

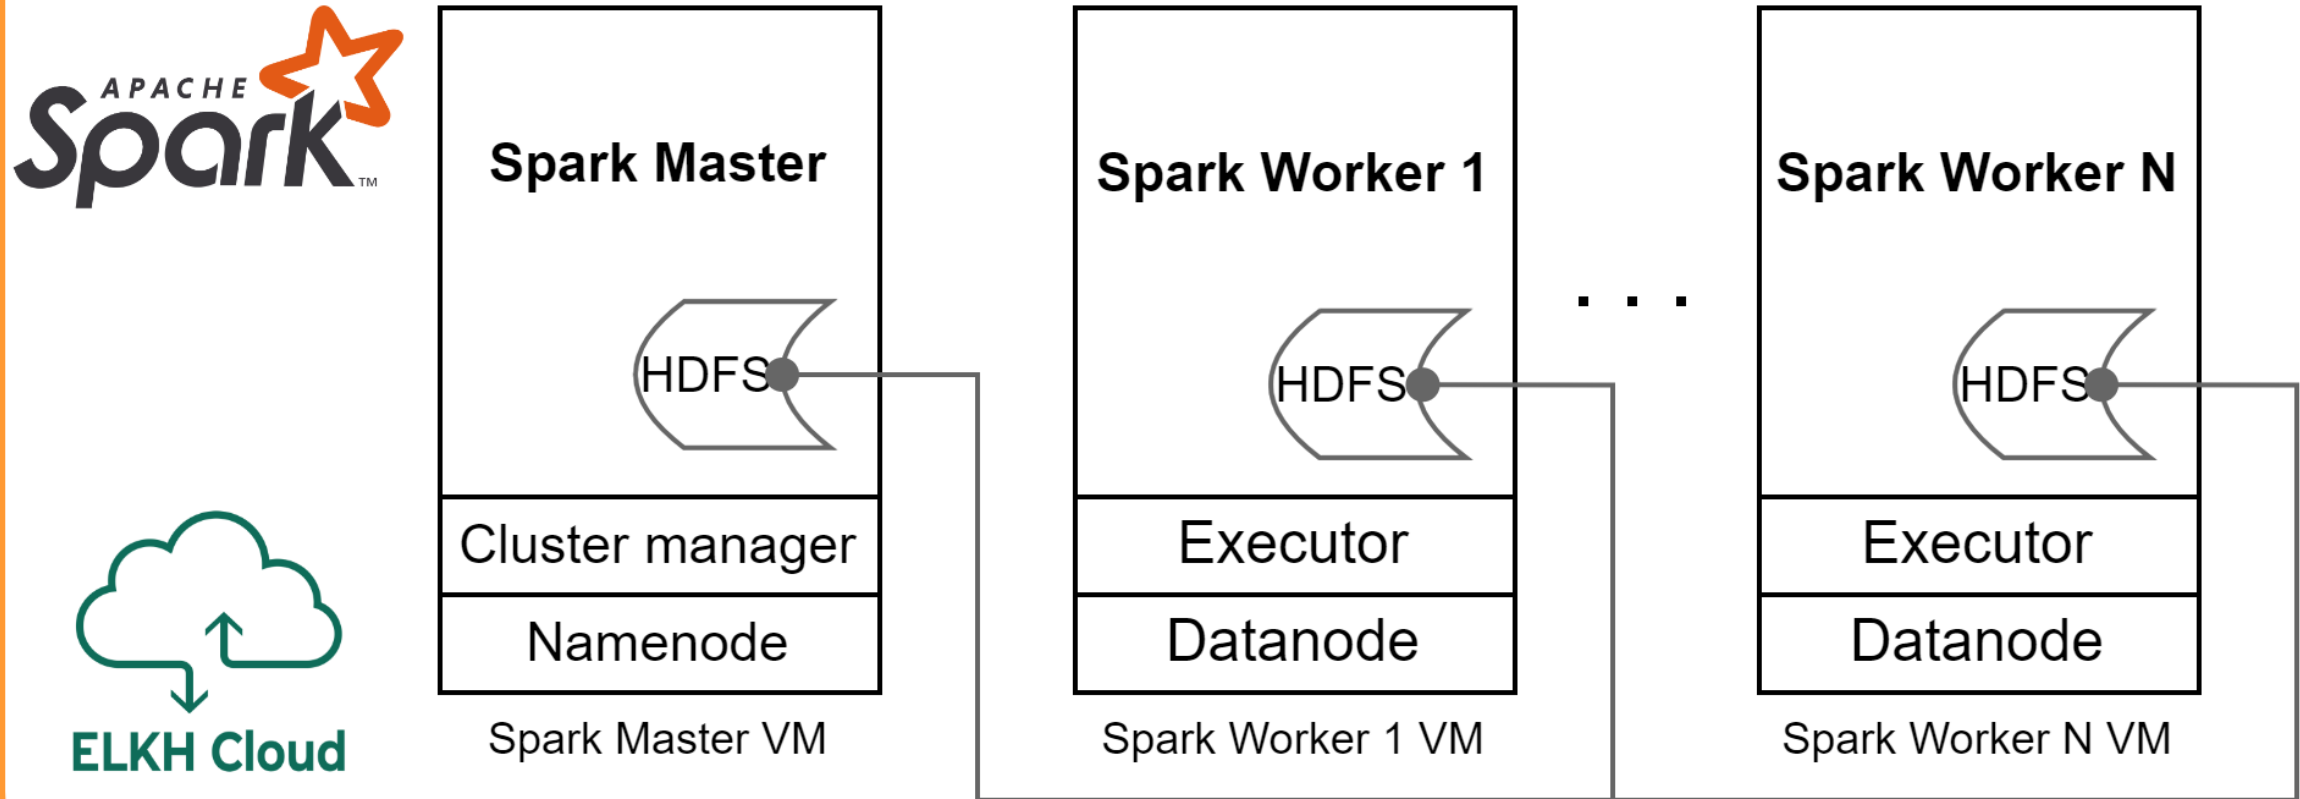




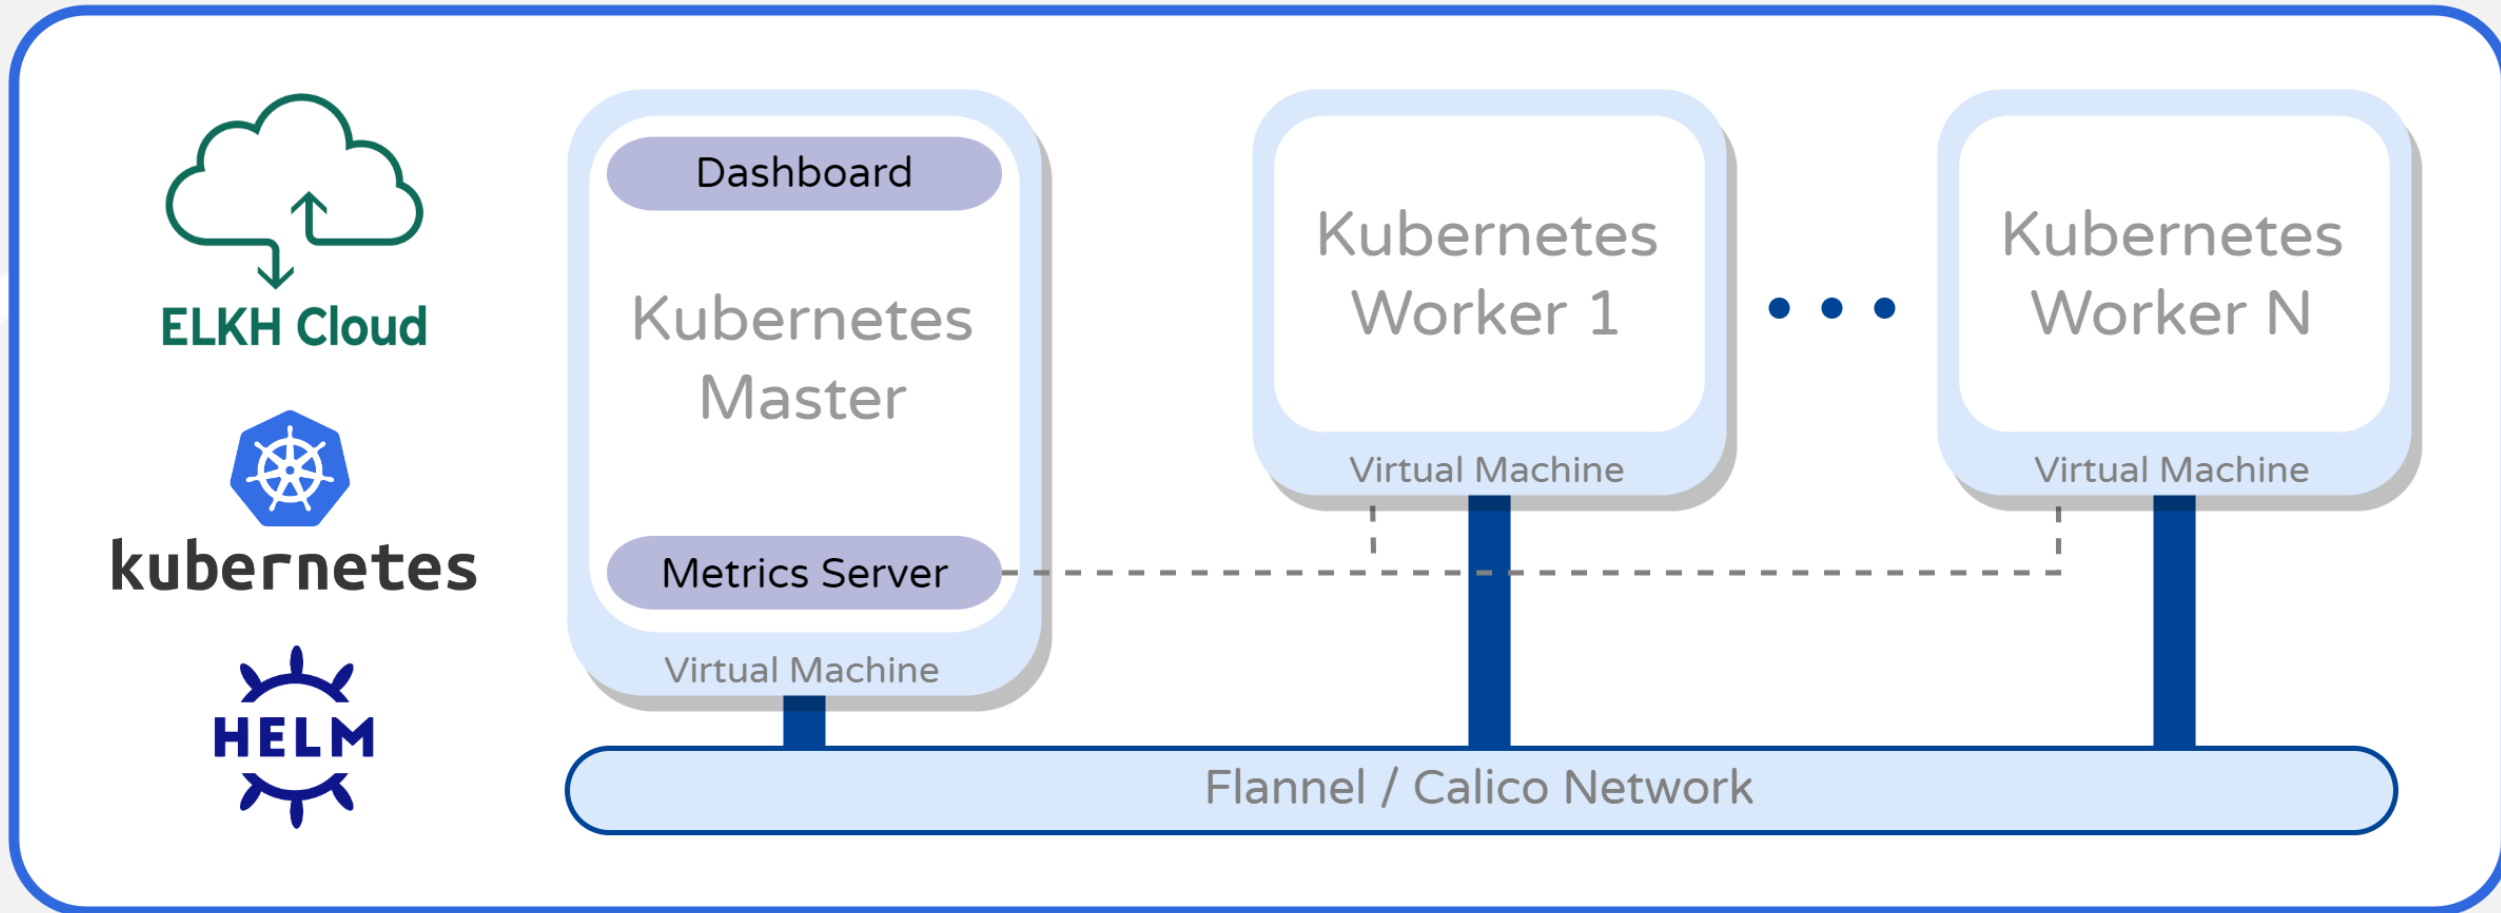
# Horovod Performance Evaluation



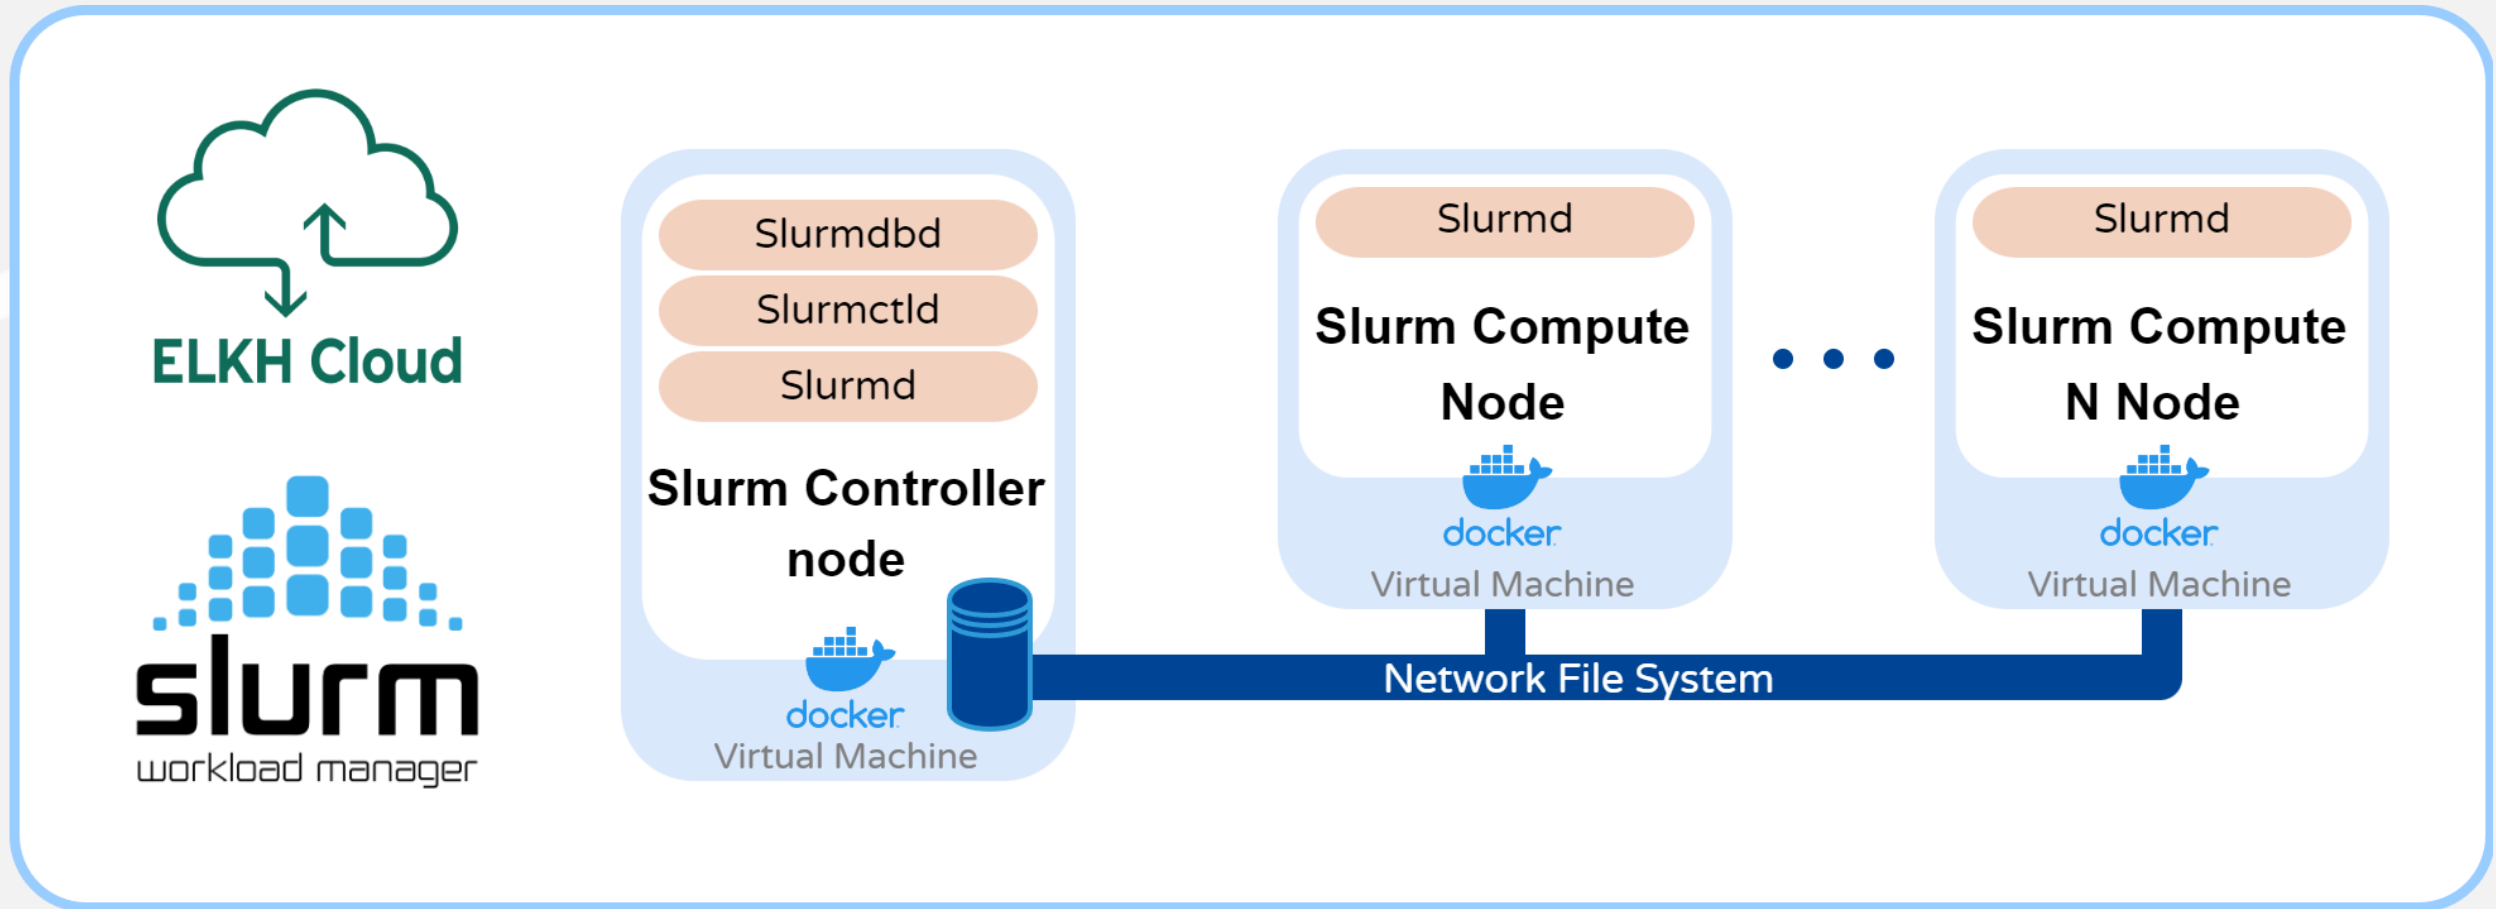
# Apache Spark reference architecture



# Kubernetes reference architecture

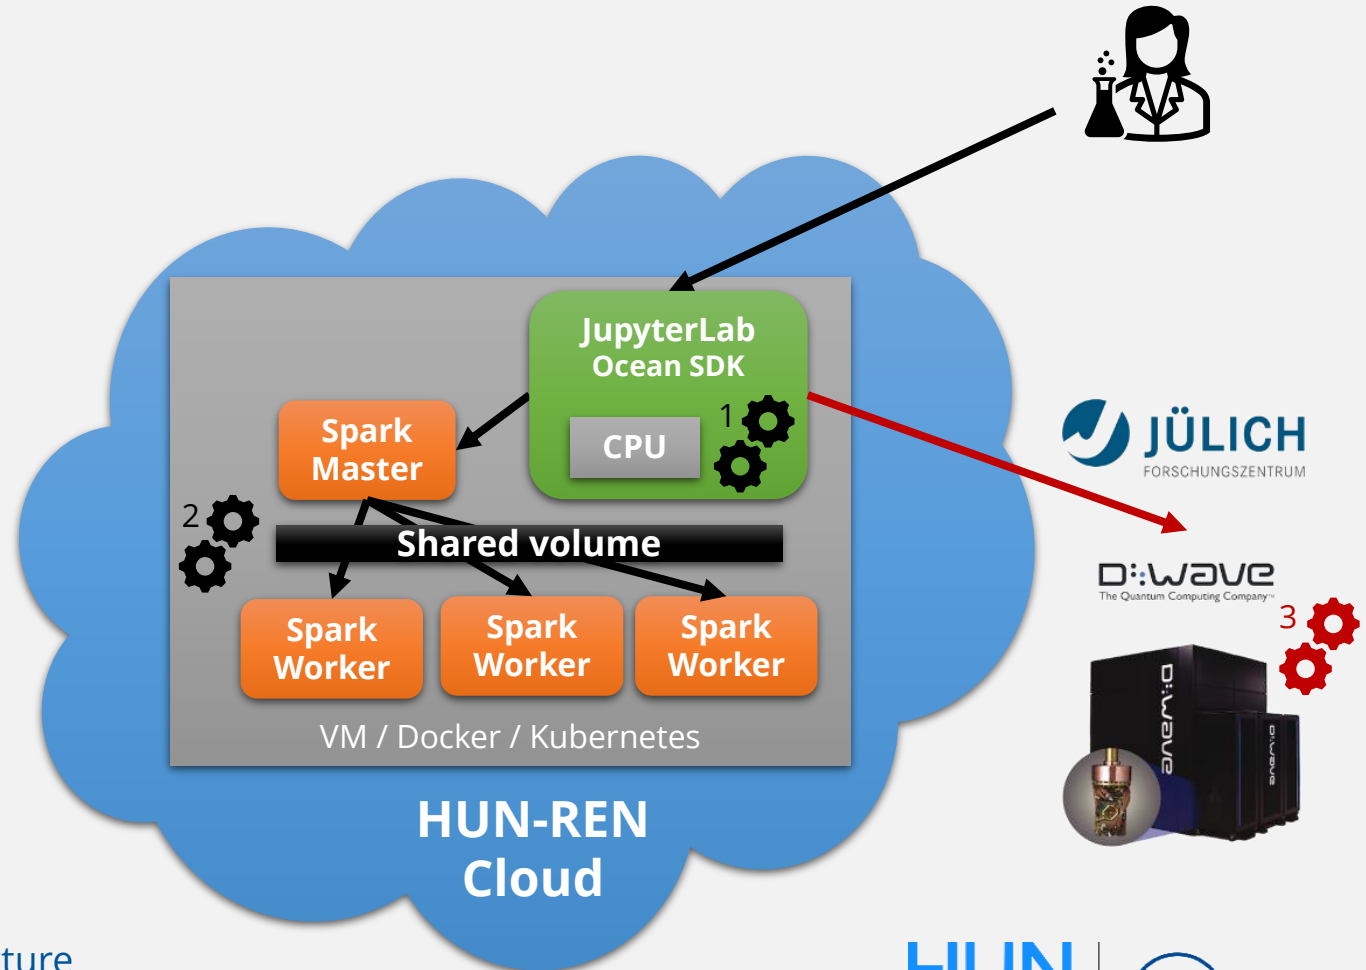


# SLURM reference architecture

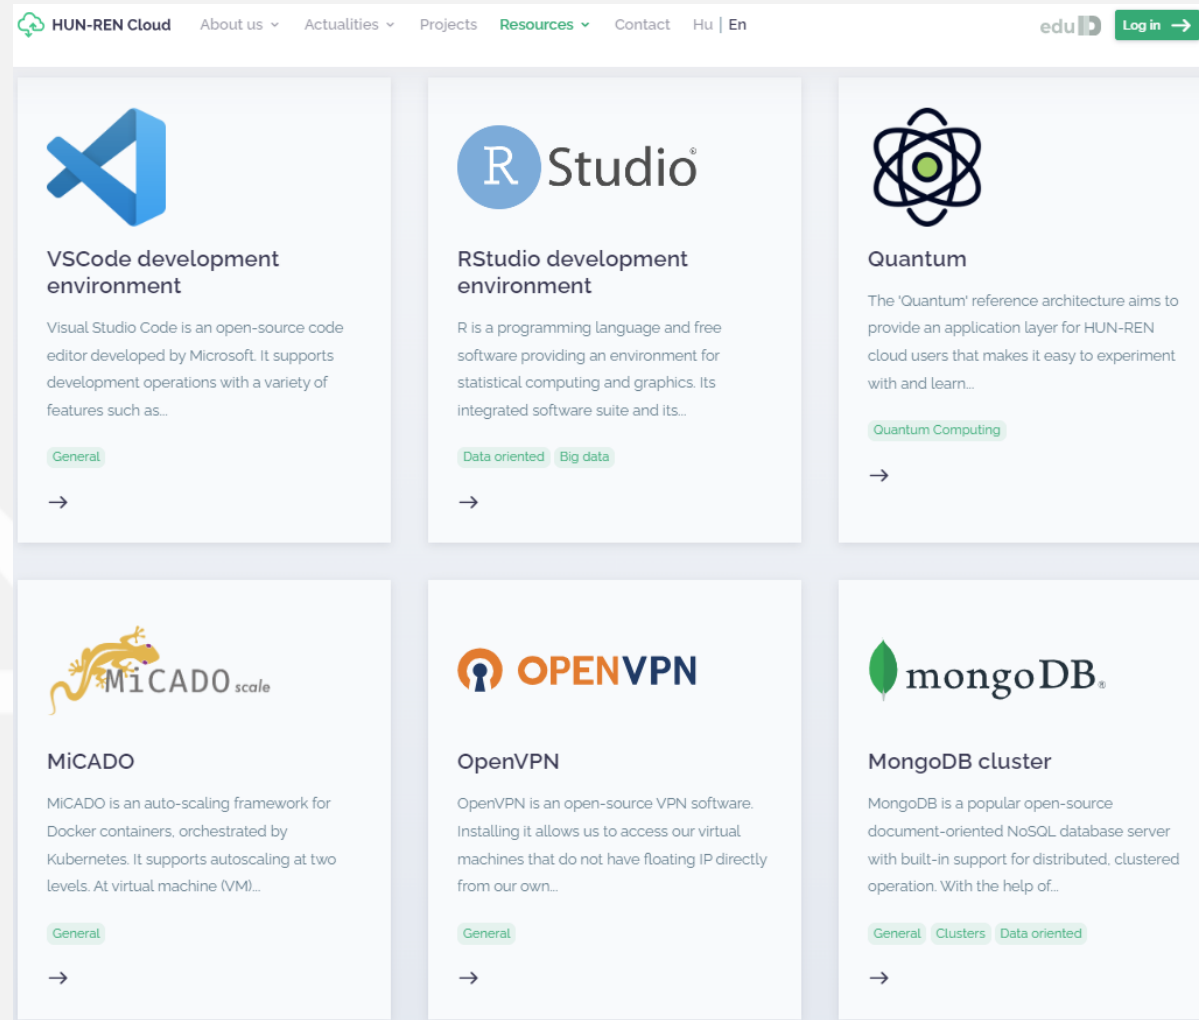


# Quantum reference architecture

- An **increasing number** of experimental academic and commercial services provide computing capacities based on **quantum** principles
- but their underlying technologies, availability and interfaces differ significantly
- **New reference architecture** significantly lowers the entry barrier for quantum programming for future users



# Reference architectures



The screenshot displays the 'Resources' section of the HUN-REN Cloud website. It features a grid of six reference architecture cards, each with a logo, title, description, and category tags. The cards are: VSCode development environment (General), RStudio development environment (Data oriented, Big data), Quantum (Quantum Computing), MiCADO (General), OpenVPN (General), and MongoDB cluster (General, Clusters, Data oriented). Each card includes a right-pointing arrow at the bottom.

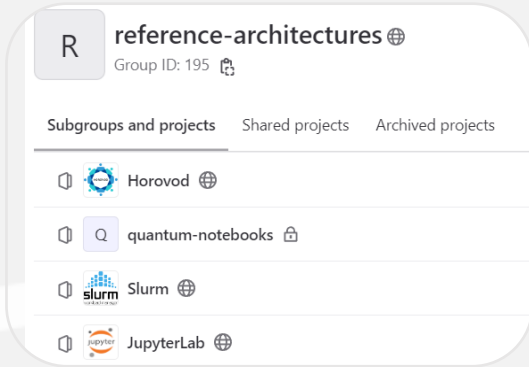
**HUN-REN Cloud** About us Actualities Projects Resources Contact Hu | En edu ID Log in →

- VSCode development environment**  
Visual Studio Code is an open-source code editor developed by Microsoft. It supports development operations with a variety of features such as...  
General →
- RStudio development environment**  
R is a programming language and free software providing an environment for statistical computing and graphics. Its integrated software suite and its...  
Data oriented Big data →
- Quantum**  
The 'Quantum' reference architecture aims to provide an application layer for HUN-REN cloud users that makes it easy to experiment with and learn...  
Quantum Computing →
- MiCADO**  
MiCADO is an auto-scaling framework for Docker containers, orchestrated by Kubernetes. It supports autoscaling at two levels. At virtual machine (VM)...  
General →
- OpenVPN**  
OpenVPN is an open-source VPN software. Installing it allows us to access our virtual machines that do not have floating IP directly from our own...  
General →
- MongoDB cluster**  
MongoDB is a popular open-source document-oriented NoSQL database server with built-in support for distributed, clustered operation. With the help of...  
General Clusters Data oriented →

<https://science-cloud.hu/en/reference-architectures>

Reference architecture  
concept on HUN-REN  
Cloud

# Reference architectures: Aiming towards uniformity

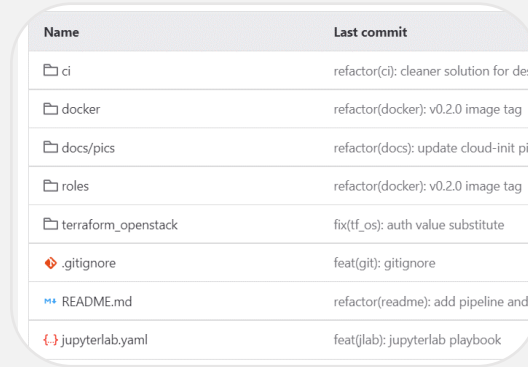


reference-architectures  
Group ID: 195

Subgroups and projects Shared projects Archived projects

- Horovod
- quantum-notebooks
- Slurm
- JupyterLab

Repository



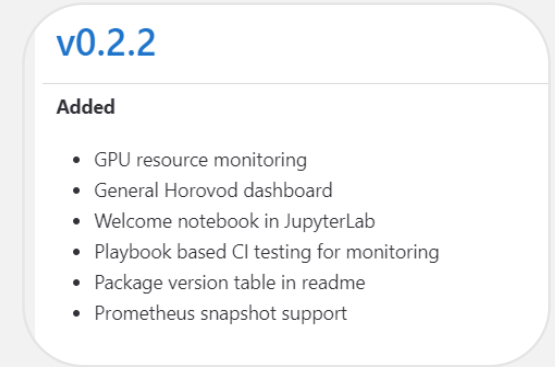
Name	Last commit
ci	refactor(ci): cleaner solution for dest
docker	refactor(docker): v0.2.0 image tag
docs/pics	refactor(docs): update cloud-init pic
roles	refactor(docker): v0.2.0 image tag
terraform_openstack	fix(tf_os): auth value substitute
gitignore	feat(git): gitignore
README.md	refactor(readme): add pipeline and r
jupyterlab.yaml	feat(jlab): jupyterlab playbook

Code Structure



Terraform  
ANSIBLE  
docker

Toolset

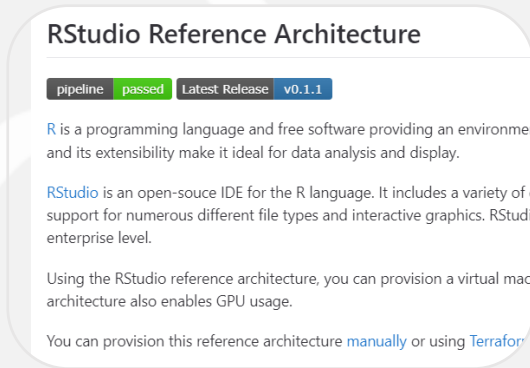


v0.2.2

Added

- GPU resource monitoring
- General Horovod dashboard
- Welcome notebook in JupyterLab
- Playbook based CI testing for monitoring
- Package version table in readme
- Prometheus snapshot support

Release Notes



RStudio Reference Architecture

pipeline passed Latest Release v0.1.1

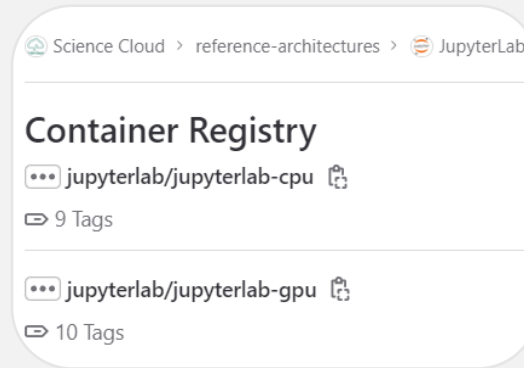
R is a programming language and free software providing an environment for statistical computing and its extensibility make it ideal for data analysis and display.

RStudio is an open-source IDE for the R language. It includes a variety of support for numerous different file types and interactive graphics. RStudio enterprise level.

Using the RStudio reference architecture, you can provision a virtual machine architecture also enables GPU usage.

You can provision this reference architecture manually or using Terraform

Documentation

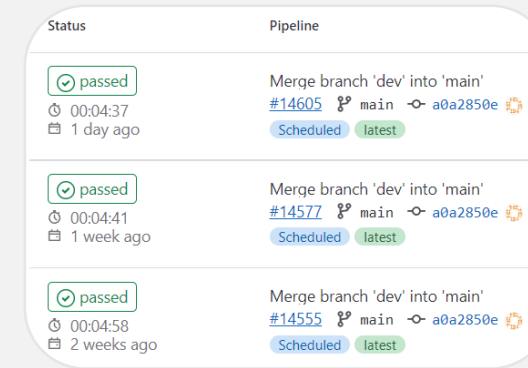


Science Cloud > reference-architectures > JupyterLab

Container Registry

- jupyterlab/jupyterlab-cpu  
9 Tags
- jupyterlab/jupyterlab-gpu  
10 Tags

Image store



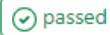








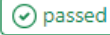








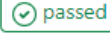








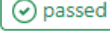








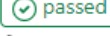








Status Pipeline

- passed 00:04:37 1 day ago  
Merge branch 'dev' into 'main' #14605 main -> a0a2850e  
Scheduled latest
- passed 00:04:41 1 week ago  
Merge branch 'dev' into 'main' #14577 main -> a0a2850e  
Scheduled latest
- passed 00:04:58 2 weeks ago  
Merge branch 'dev' into 'main' #14555 main -> a0a2850e  
Scheduled latest

Testing

# Reference architectures: Automatic testing

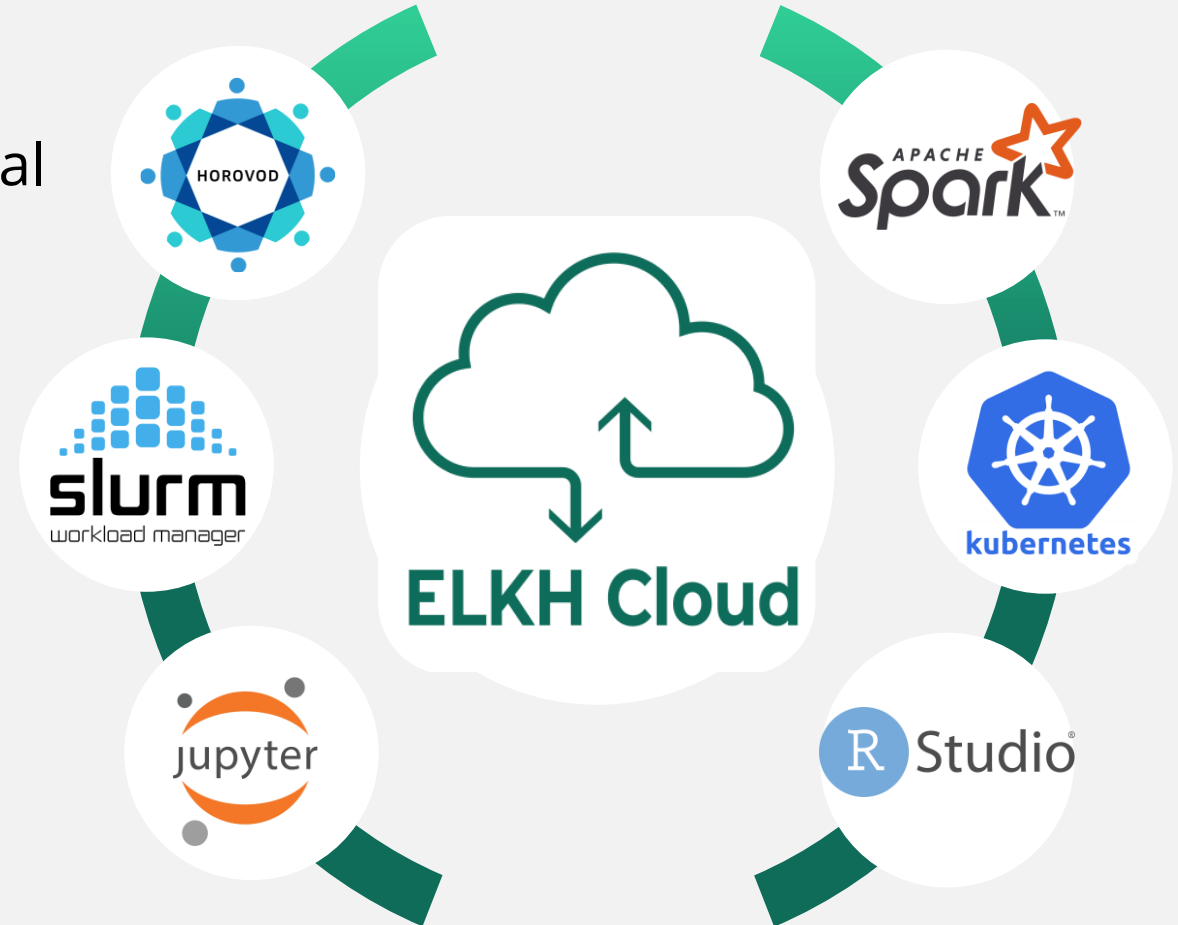
- Prevent software rot
- Dedicated project
- Weekly scheduled pipelines
- Email notifications
- Failsafe cleanup after each pipeline
- Modular functional tests
- Easily expandable



Status	Pipeline	Triggerer	Stages
 00:04:37 2 days ago	Merge branch 'dev' into 'main' <a href="#">#14605</a>  main  a0a2850e  Scheduled latest		   
 00:04:41 1 week ago	Merge branch 'dev' into 'main' <a href="#">#14577</a>  main  a0a2850e  Scheduled latest		   
 00:04:58 2 weeks ago	Merge branch 'dev' into 'main' <a href="#">#14555</a>  main  a0a2850e  Scheduled latest		   
 00:04:31 3 weeks ago	Merge branch 'dev' into 'main' <a href="#">#14515</a>  main  a0a2850e  Scheduled latest		   
 00:04:26 4 weeks ago	Merge branch 'dev' into 'main' <a href="#">#14484</a>  main  a0a2850e  Scheduled latest		   



# HUN-REN Cloud Reference Architectures Summary

- Quickly and easily deployable digital research infrastructures
- Cover a variety of common use-cases
- Open-source and up-to-date
- Infrastructure as Code
- Based on widespread tools
- Long-term support
- Satisfied non-functional requirements





Thank you for your attention!