Rancher managed Kubernetes for Helmholtz users

Container Orchestration

Tim Wetzel, Michael Schuh, Johannes Reppin, Patrick Fuhrmann, Uwe Jandt
HEPiX workshop 2023, Academia Sinica, Taipei, Taiwan

HELMHOLTZ RESEARCH FOR GRAND CHALLENGES

DESY.

DESY.

What is HIFIS?

Helmholtz federated IT services

- Helmholtz Association with 18 autonomous research centres in Germany
- Incubator platforms for better collaboration between centres
  - Using synergies is key!
- HIFIS is the central IT service federation platform in Helmholtz
- Very good review last year from international experts
- Centres make web services and resources available for all other Helmholtz members
- Central AAI with community attributes and integration
  - Helmholtz, EGI, eduGAIN, ...

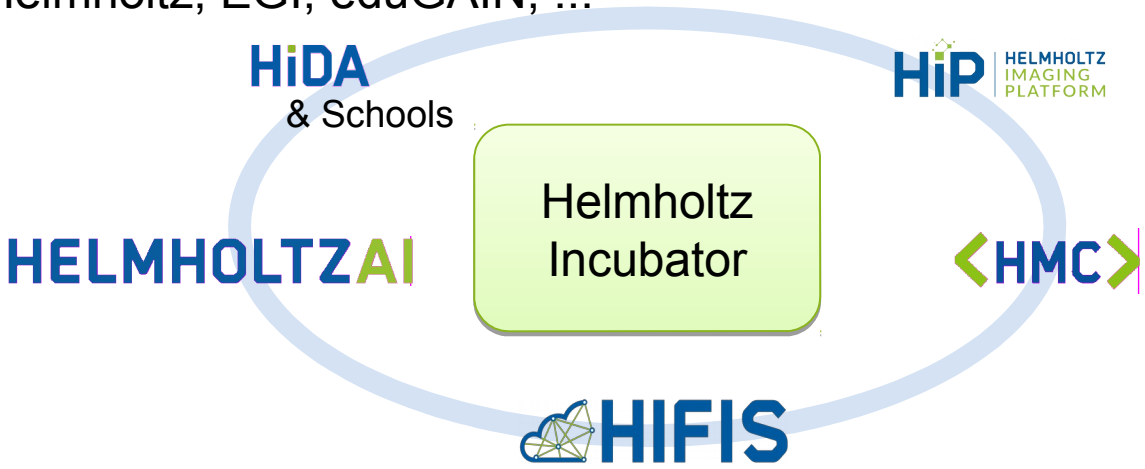

The Cloud Native Landscape

The Cloud Native Trail Map

Container orchestration

1. Containerize applications
2. CI/CD - robust automation
3. Orchestration
4. Monitoring & Analysis

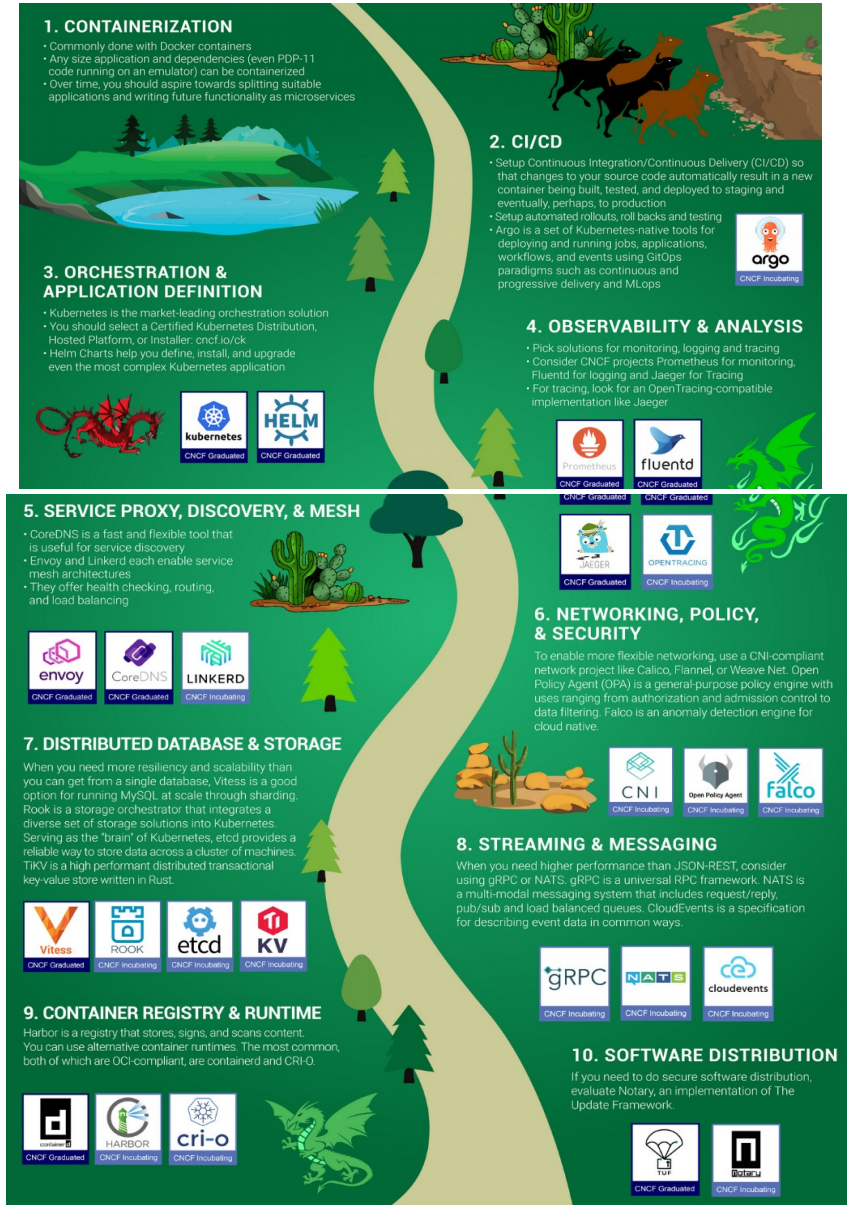

Source: trailmap.cncf.io

Integration with research infrastructure

5. DNS, Load Balancing
6. Network Operations
   - Software-defined networking
   - Firewall
   - Dynamic certificates
7. Scientific data storage
   - dCache
   - High performance storage
   - DBaaS
8. Event streaming platforms
   - Data acquisition streams
   - Function as a service (FaaS)
9. Scale container registry
   - HPC & HTC
   - Docker / Singularity Registry
10. Software Repository
    - CVMFS
    - NIMS

Software stack container orchestration

Infrastructure as a Service
Cloud Computing
Infrastructure as code

Software stack container orchestration

Container as a Service
Cloud Native CI/CD
Container registry

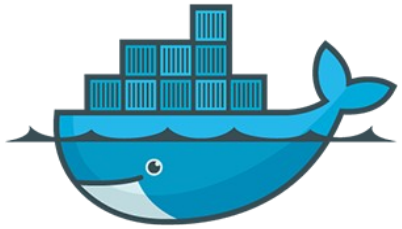

Infrastructure as a Service
Cloud Computing
Infrastructure as code

Software stack container orchestration

Kubernetes as a Service
Container Orchestration
Kubernetes Package Manager

Container as a Service
Cloud Native CI/CD
Container registry

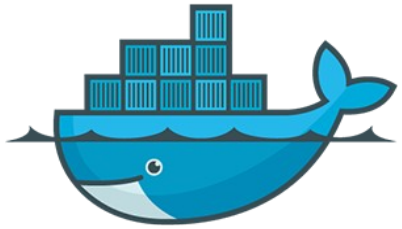

Infrastructure as a Service
Cloud Computing
Infrastructure as code

Software stack container orchestration

Software as a Service
Container-based environments
App deployments as code

Kubernetes as a Service
Container Orchestration
Kubernetes Package Manager

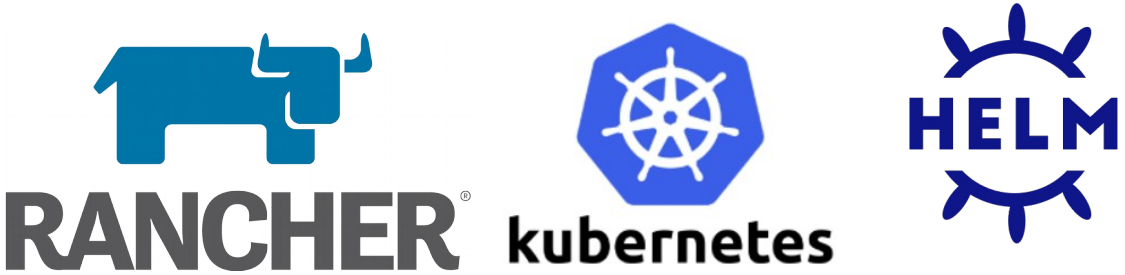

Container as a Service
Cloud Native CI/CD
Container registry

Infrastructure as a Service
Cloud Computing
Infrastructure as code

Architecture

Infrastructure concepts

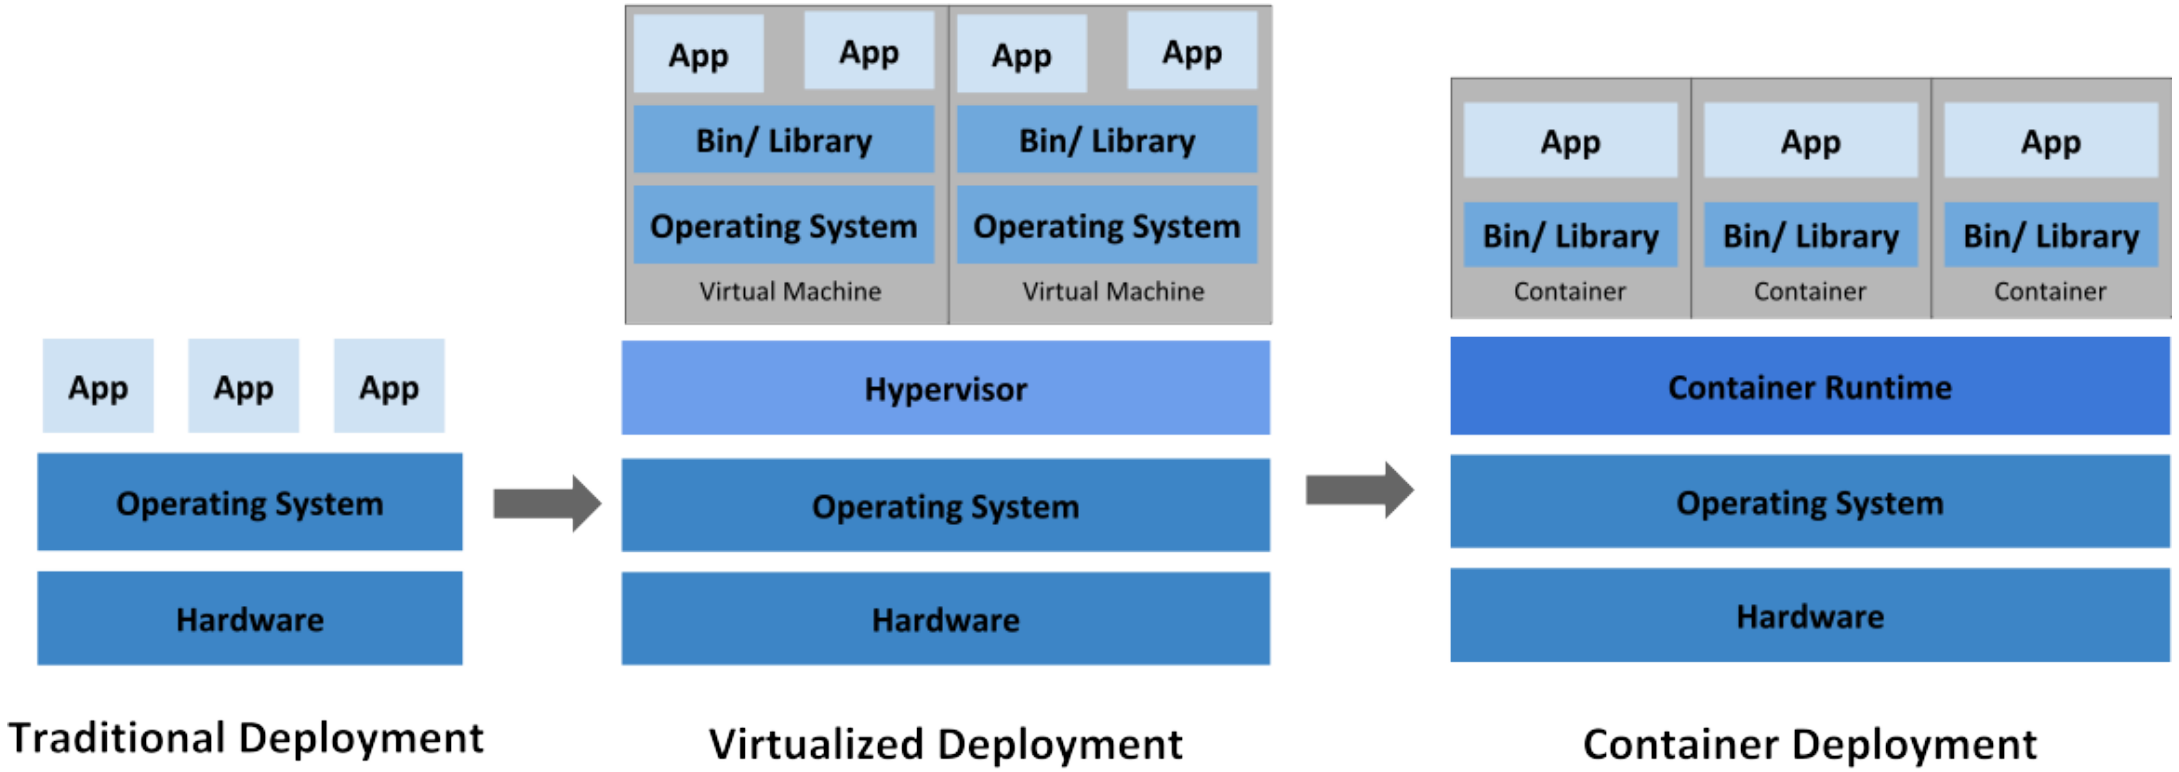

- A container image is a standalone, executable package of software that includes everything needed to run an application: code, application runtime, system tools, system libraries and settings.
- **Container images become containers at run time** and are processes (or group of processes) running in a cgroup, where the Kernel restricts resource usage (CPU, memory, disk I/O, network)

Architecture

Clusters and Schedulers

- Openstack Hypervisors organized as **Host Aggregates**
  - Openstack Nova (compute scheduler) places VMs on Hypervisors (hosts)

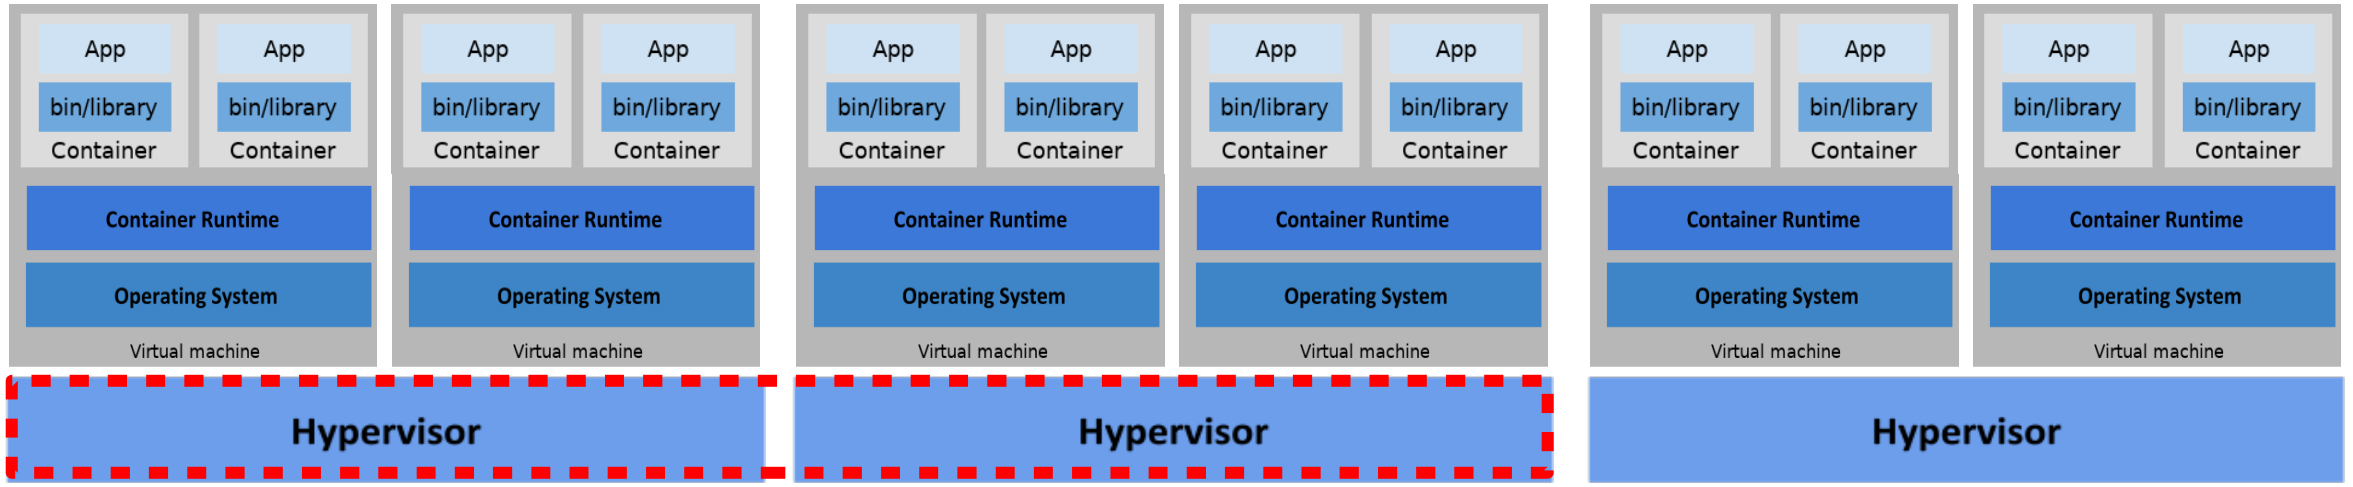

Architecture

Clusters and Schedulers

- VMs configured as Kubernetes nodes form a **Kubernetes Cluster**
  - Dynamically created by Rancher Kubernetes management platform
- Openstack Hypervisors organized as **Host Aggregates**
  - Openstack Nova (compute scheduler) places VMs on Hypervisors

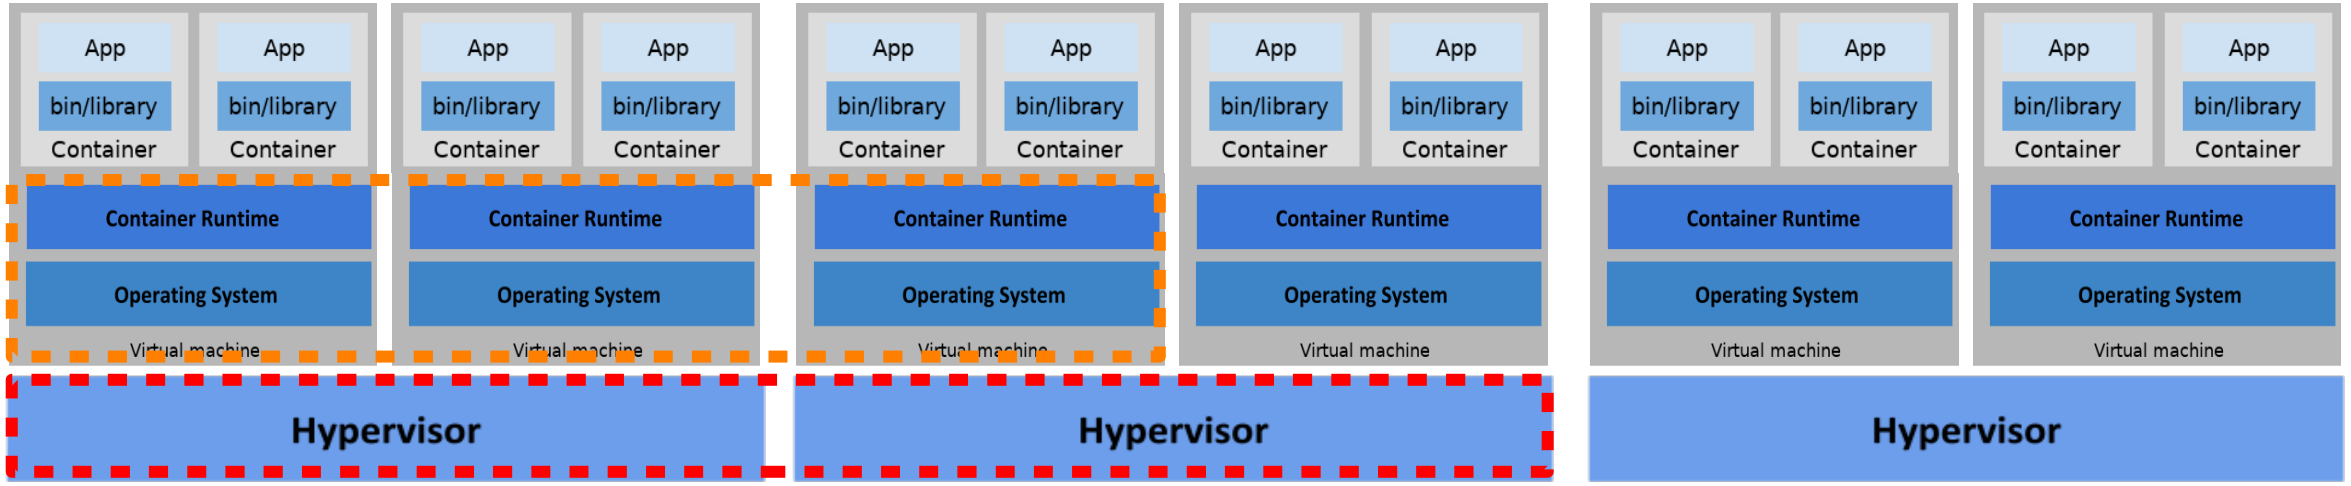

Architecture

Clusters and Schedulers

- Kubernetes scheduler places containers on Kubernetes Nodes
  - Application Deployments form **Container Clusters**
- VMs configured as Kubernetes nodes form a **Kubernetes Cluster**
  - Dynamically created by Rancher Kubernetes Management platform
- Openstack Hypervisors organized as **Host Aggregates**
  - Openstack Nova (compute scheduler) places VMs on Hypervisors

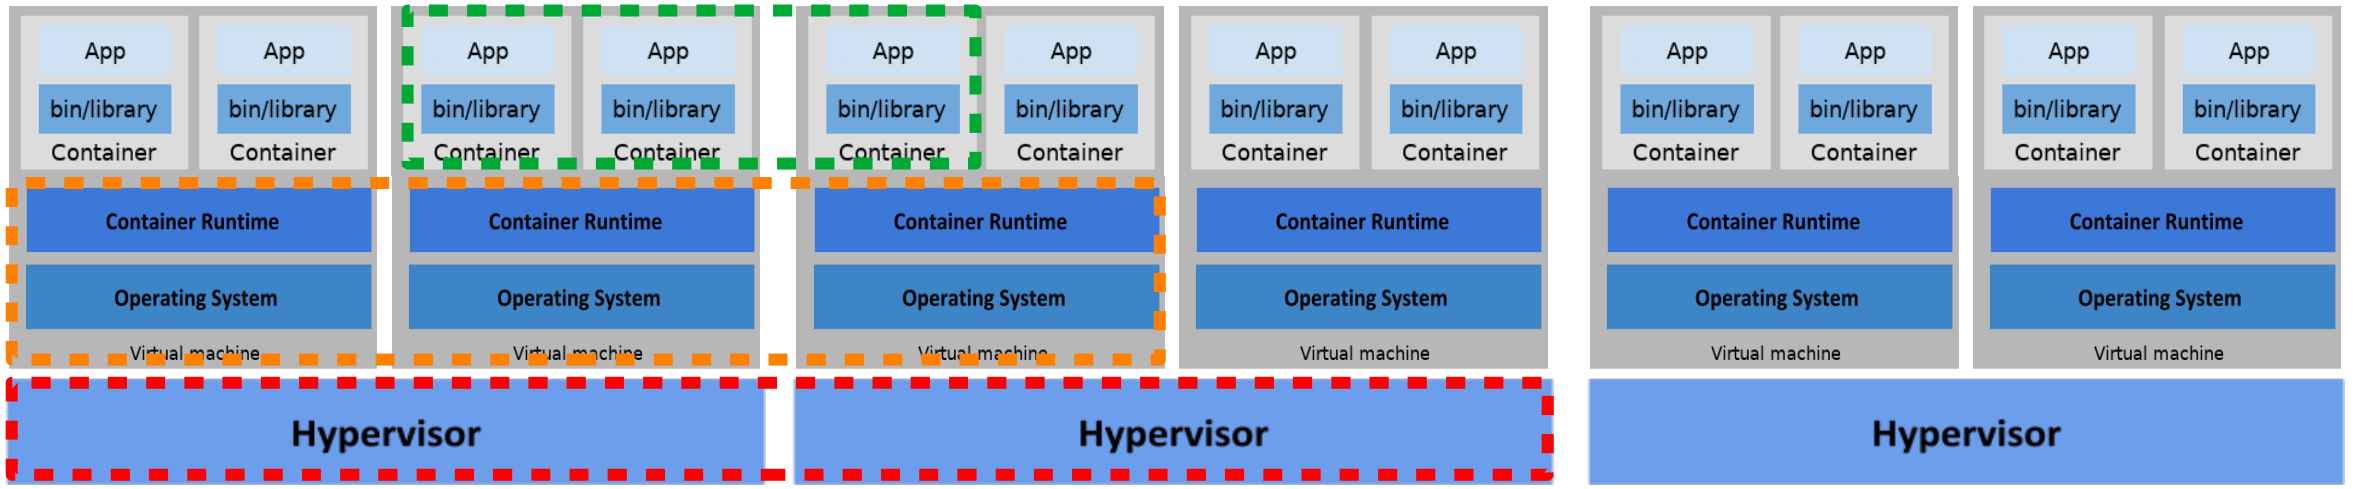

Rancher

K8s cluster management and addtitional RBAC

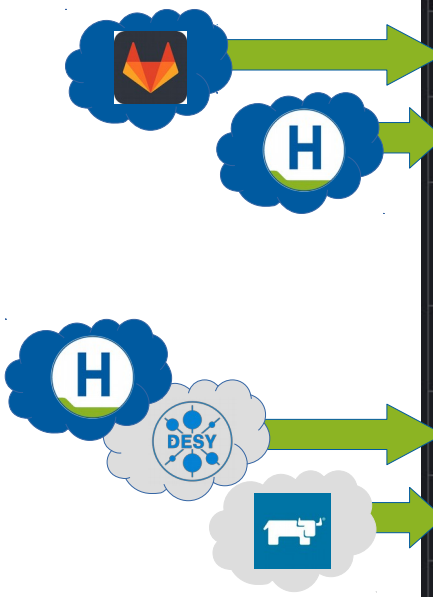

Clusters 9

Import Existing | Create | Filter

State	Name	Provider	Kubernetes Version	CPU	Memory	Pods
Active	gitlab-public-ci	Custom RKE	v1.21.9	72 cores	140 GiB	103/990
Active	guest-cluster2	Custom RKE	v1.22.10	40 cores	78 GiB	170/550
Unavailable	internal-public-ci Cluster agent is not connected	Imported Imported	v1.21.13	24 cores	47 GiB	–
Active	it-paas-k8s-cluster1	OpenStack RKE2	v1.26.2+rke2r1	24 cores	47 GiB	44/330
Active	kube-cluster1	Custom RKE	v1.24.8	176 cores	343 GiB	413/1870
Active	local	Local K3s	v1.21.5	12 cores	23 GiB	27/330
Active	mhamming-test	Custom RKE	v1.24.9	6 cores	11 GiB	28/330
Unavailable	panosc-k8s Cluster agent is not connected	Custom RKE	v1.21.4	12 cores	23 GiB	–
Active	rke-test-cluster	Custom RKE	v1.24.8	12 cores	23 GiB	49/330

GitLab Integration

Flux CD and Helm Releases

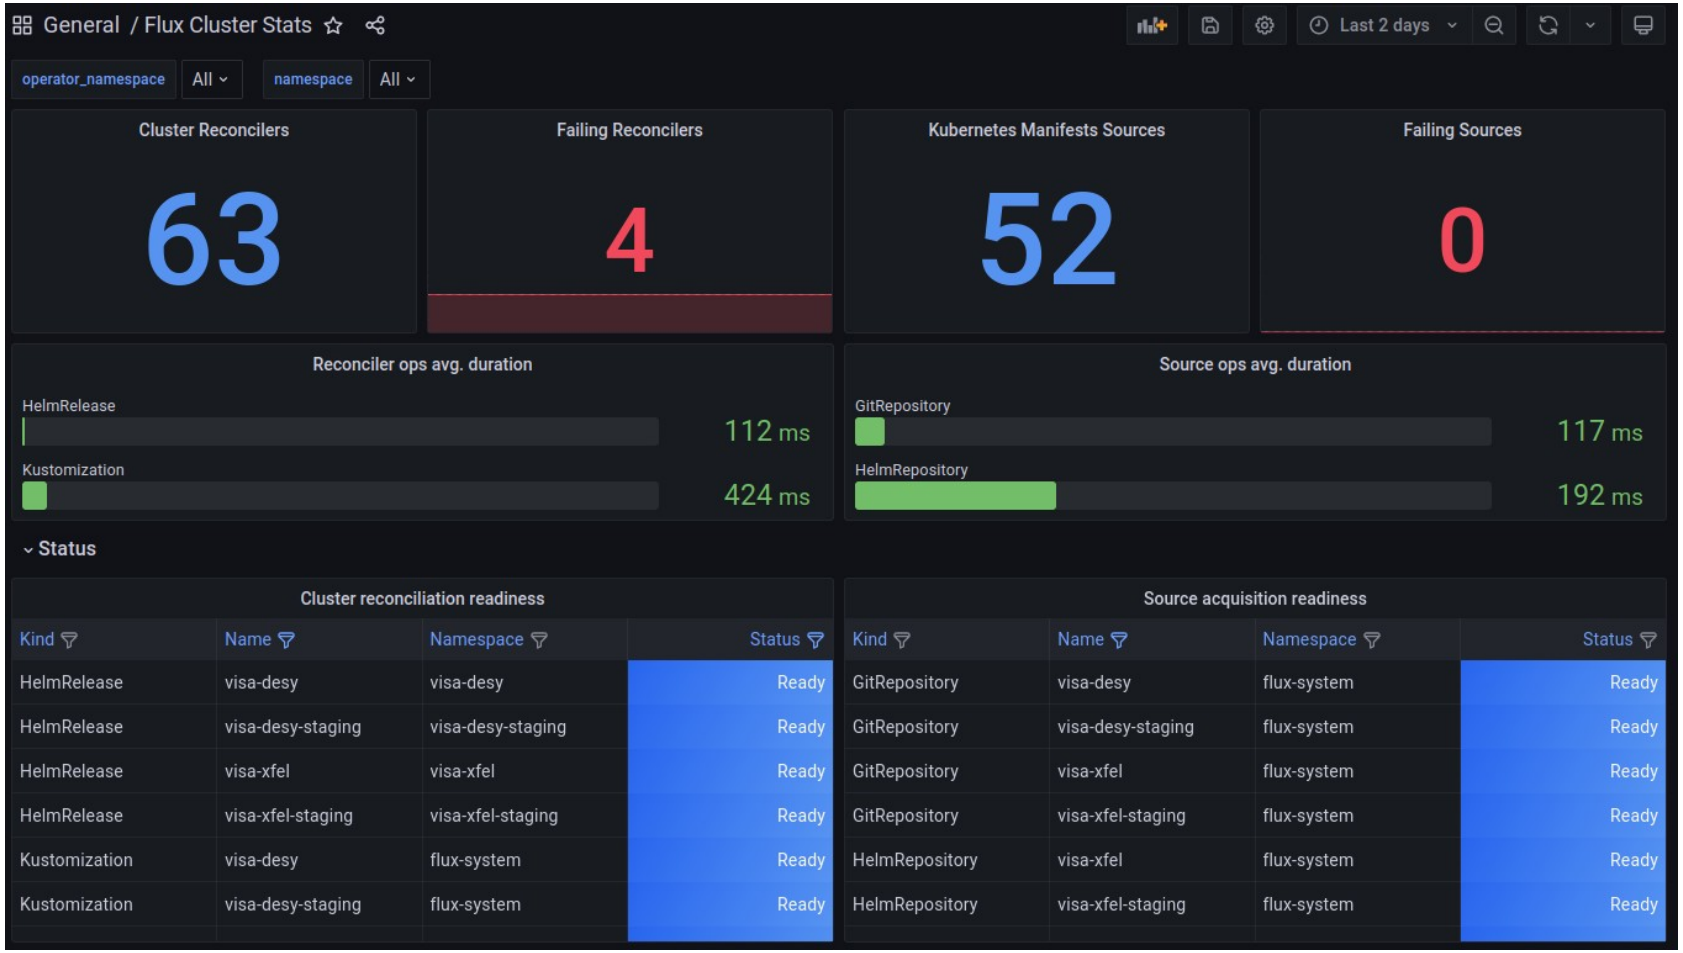

GitLab HELM

Sources
- Helm Charts
- Git Repositories
- Sealed Secrets

flux

Targets
- Clusters
- Namespaces

kubernetes

Gitlab Runners for gitlab.desy.de

Johannes Reppin, Michael Schuh – DESY IT / HIFIS, PaNOSC, ExPaNDS

Helm Deployment

- Gitlab CI/CD pipelines / jobs
- Autoscaling service
- Shared runner pilot for gitlab.desy.de
- At ~ 50k pipelines per month

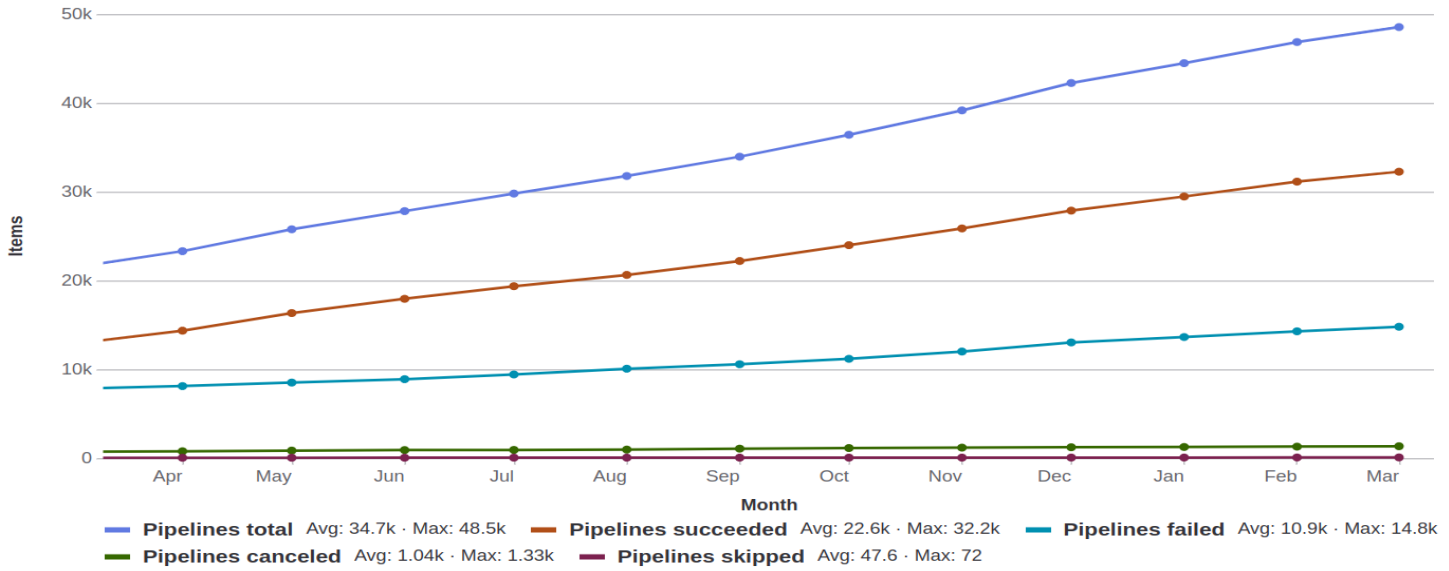

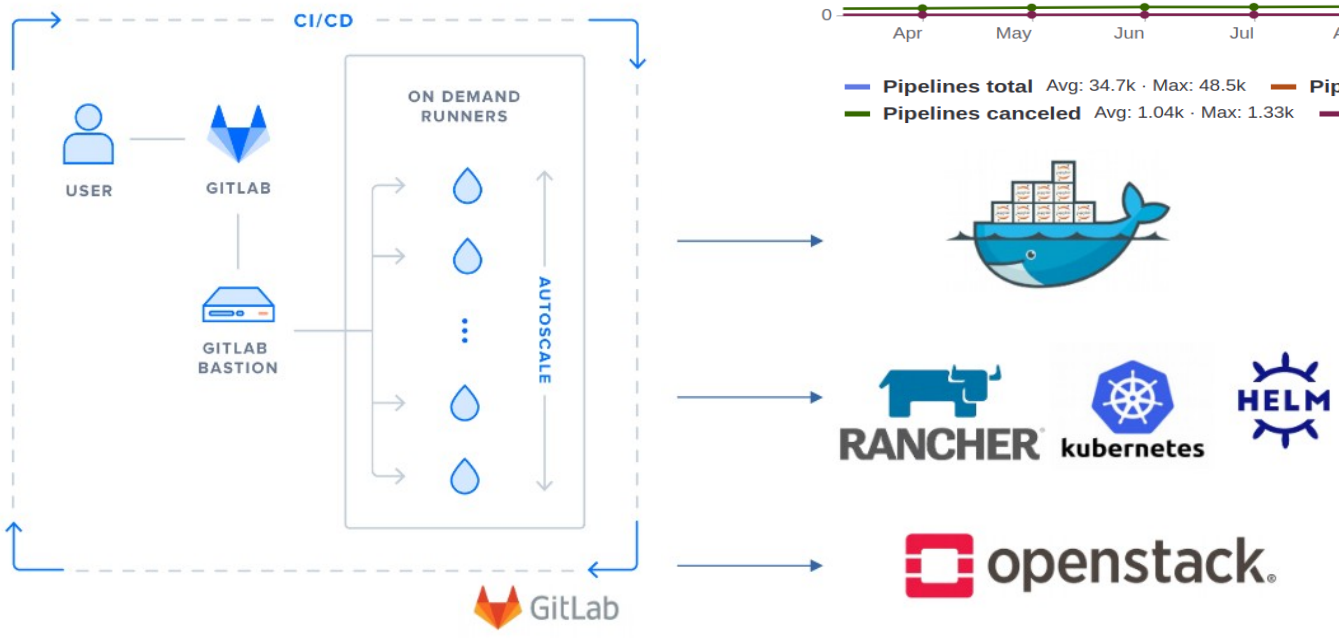

JupyterHub - jupyter.desy.de

Johannes Reppin, Tim Wetzel – DESY IT / HIFIS

Helm Deployment with Flux CD

- Merge group memberships from Helmholtz AAI in Keycloak
- Run Jupyter Servers, set UID/GIDs
- Integrate with dcache-demo.desy.de

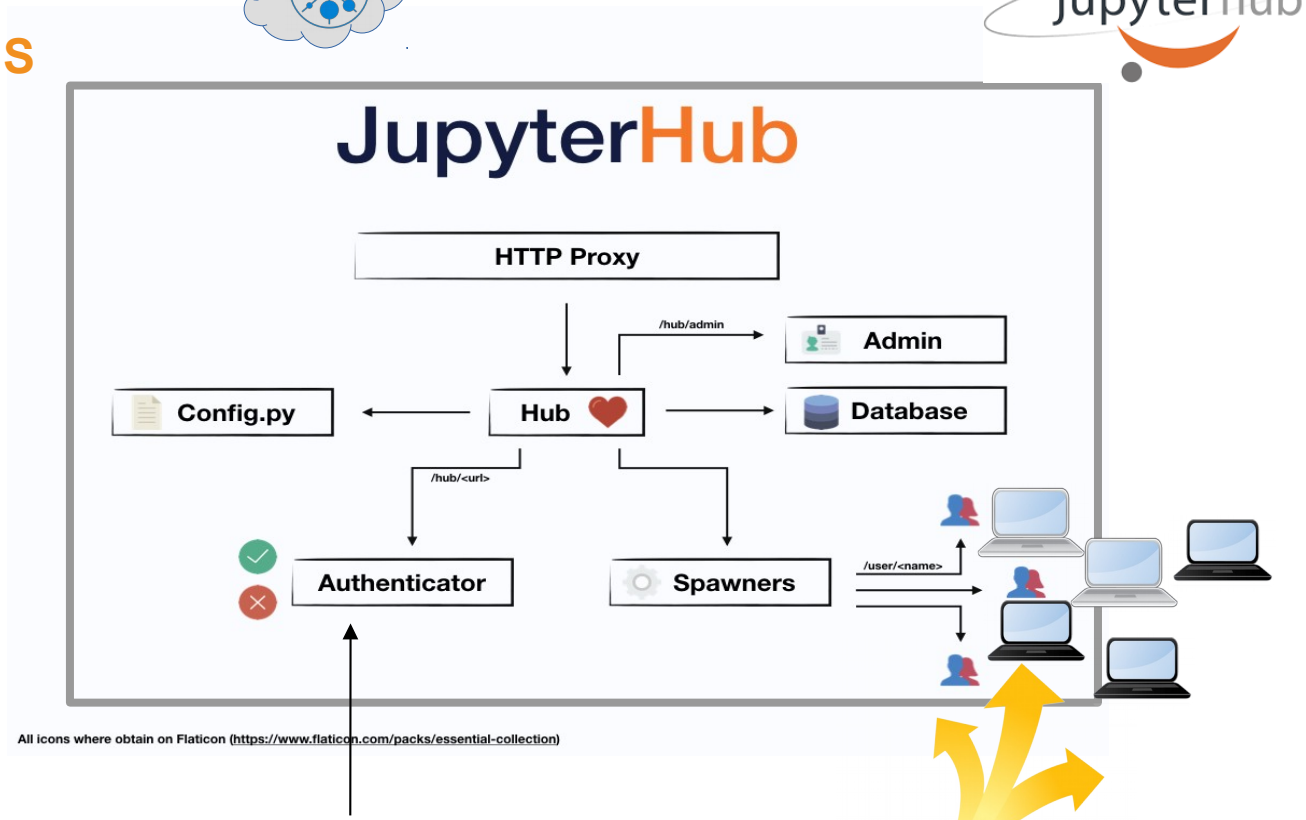

User Federation
- LDAP (DESY)
- OIDC (EGI Check-in)
- OIDC (Helmholtz AAI)

PetaByte Storage
- NFS Mounts

Collaborative document editing - notes.desy.de

Johannes Reppin – DESY IT / HIFIS

Helm Deployment with Flux CD

- HedgeDoc instance
- Collaborative note taking for Helmholtz
- Storage and backups via S3 on CEPH

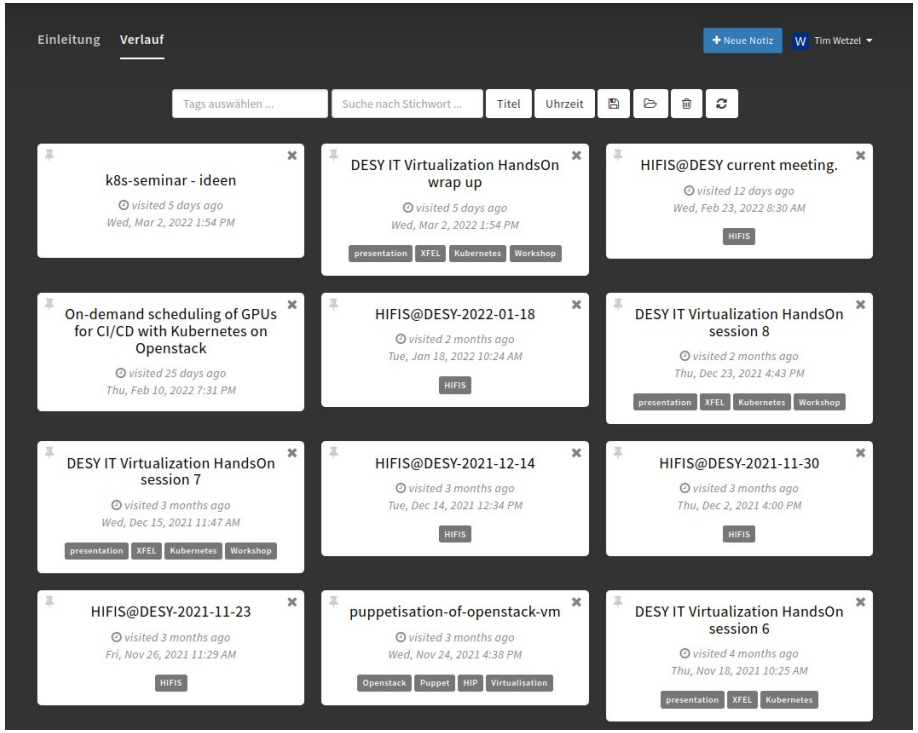

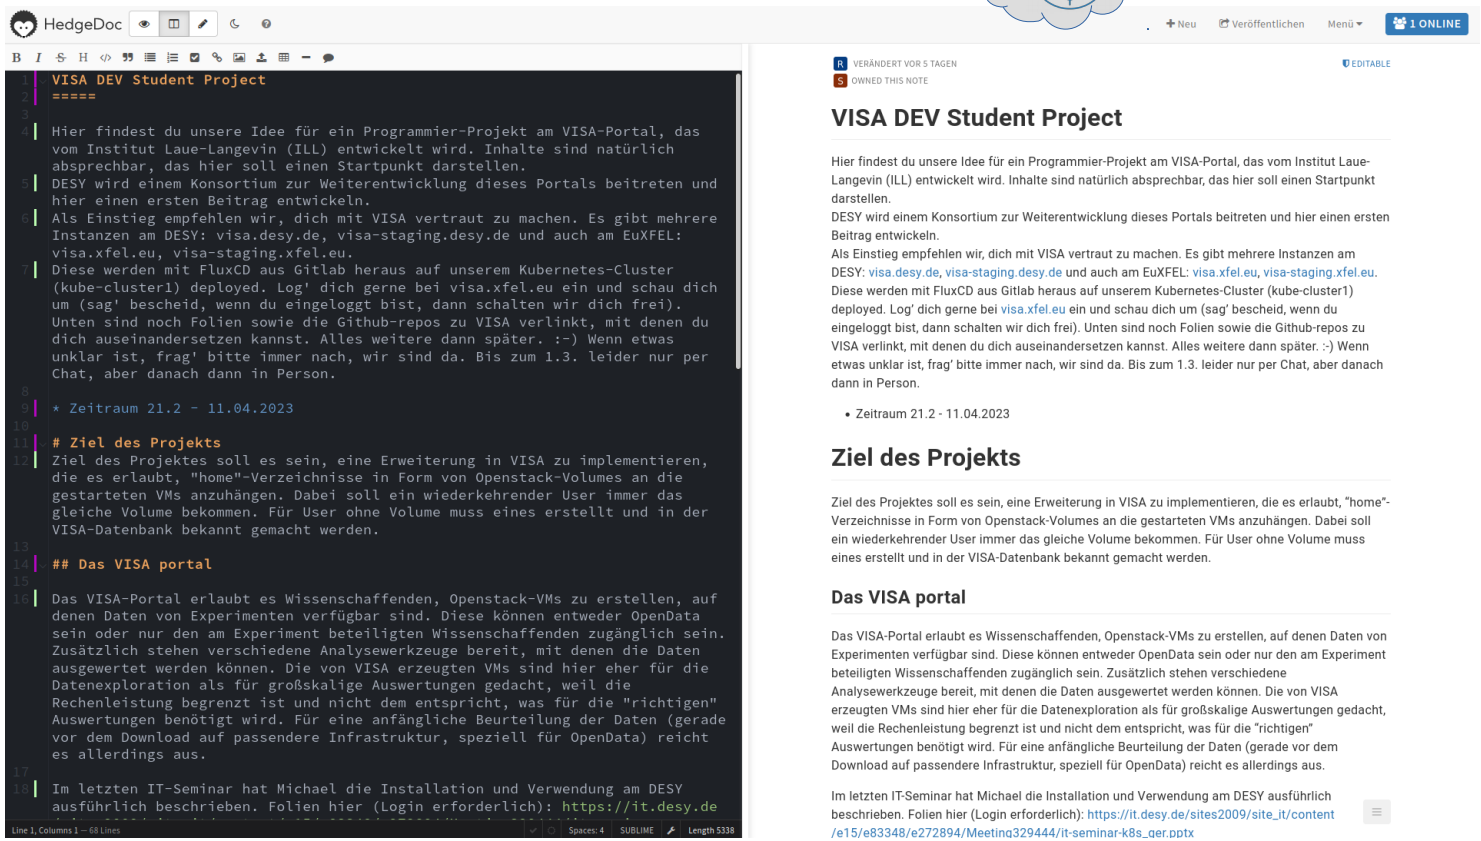

Collaborative note taking, presentation generation, LaTeX-like mathematical notation and more. More than 1.5k AAI-registered users from all Helmholtz and beyond.

Virtual Infrastructure for Scientific Applications (VISA portal)

Michael Schuh, Tim Wetzel – DESY IT

- Developed at ILL, France
- Deployed at 6+ institutes in Europe
- Scientists can spawn VMs with specific datasets mounted for data exploration after experiments
- Applications and workflows integrated into the infrastructure according to user needs

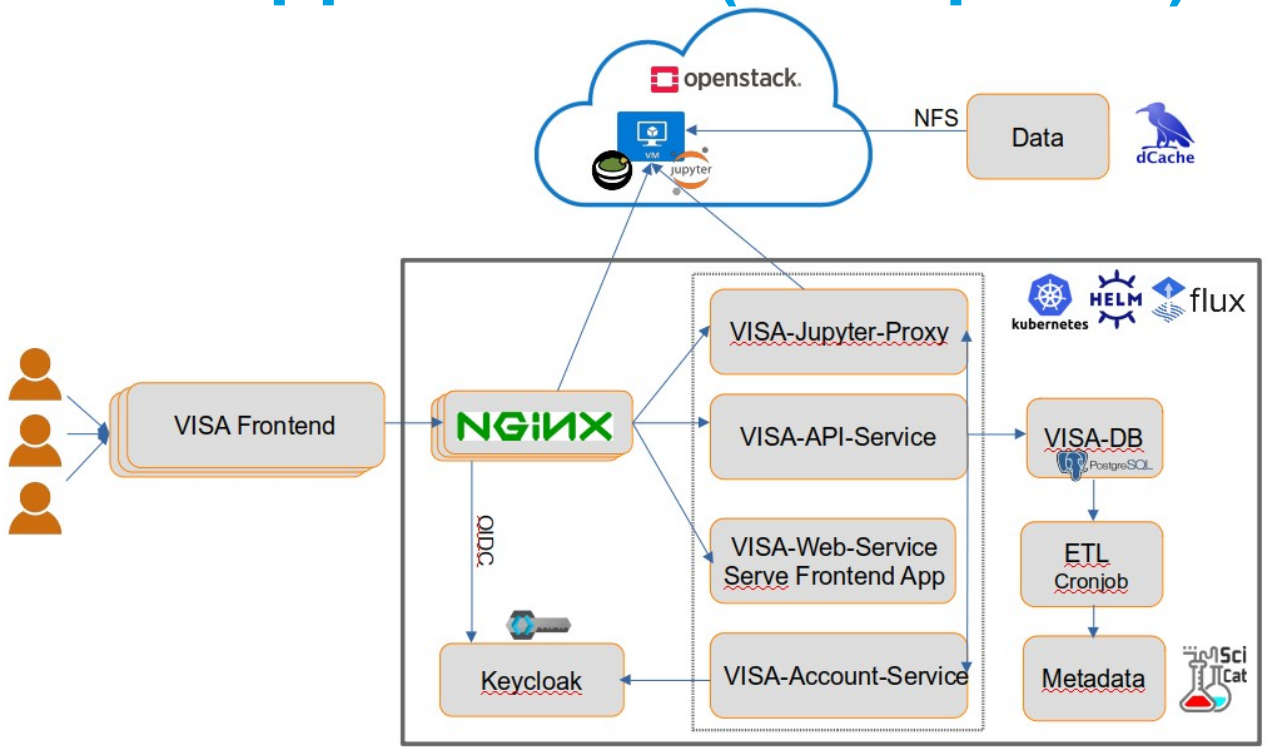

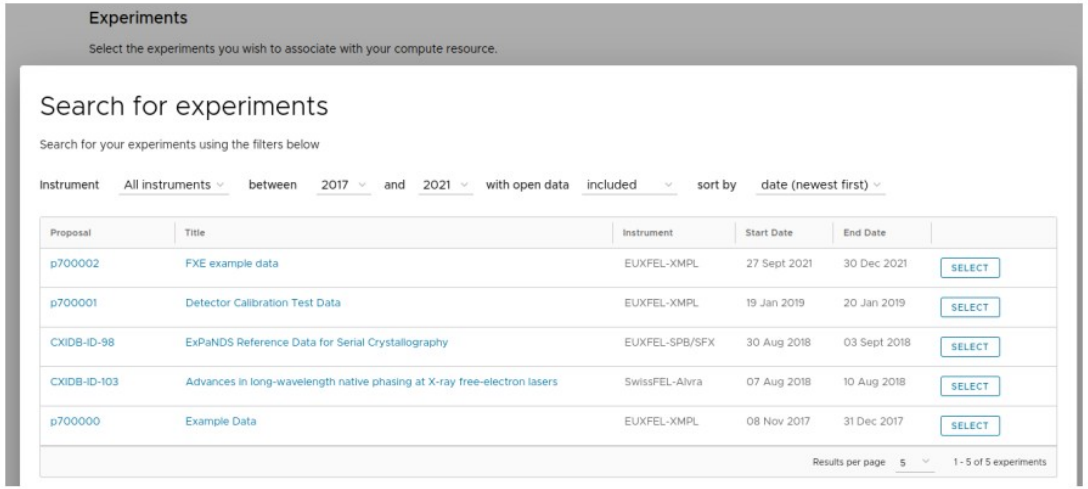

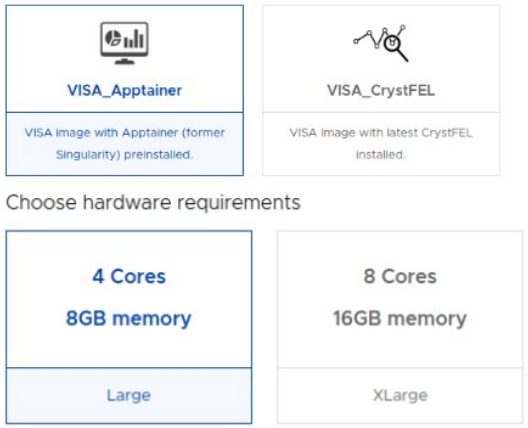

Helmholtz Cloud Portal - cloud.helmholtz.de

Thomas Beermann, Sebastian Wagner - DESY IT / HIFIS

- Service discovery portal for Helmholtz federated web service
- Allows users to find, understand and access services in Helmholtz
- Service quotas, access, monitoring to be integrated into portal
- Communication channel between services and portal in development
- Developed and hosted at DESY
- Code available on Helmholtz codebase
- Deployed on k8s via CI/CD pipelines
- Testing, monitoring, dependency management & error detection integrated into pipelines

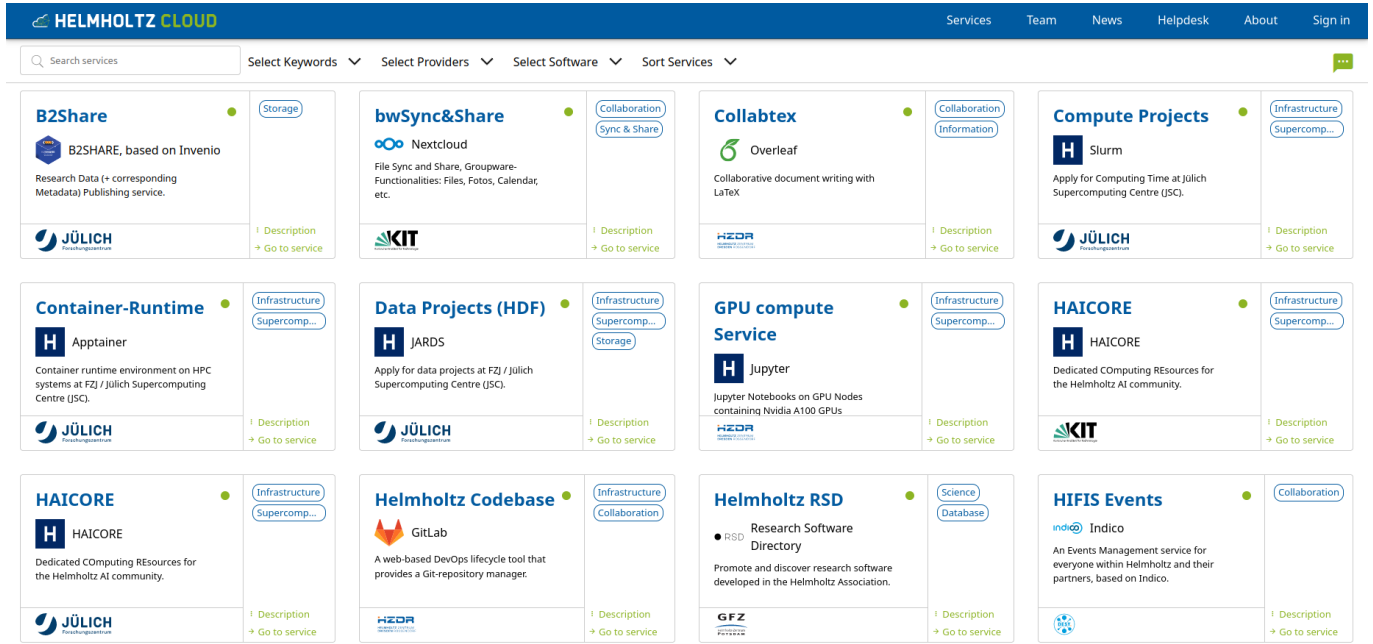

PIA eResearch system (https://info-pia.de)

Castell, Heise, Klett-Tammen, Barohn, Jemric, Steguweit – Helmholtz Centre for Infection Research (HZI); Xapling and Inqode

- System for conducting epidemiologic studies and repetitive surveys
- Available as web and mobile application for participants
- Currently built on docker-compose (see figure)
- Development for k8s-ready version facilitated with HIFIS resources

- Onboarding and introduction to the infrastructure via video call between project lead, developers and DESY team
- Access to Rancher / k8s cluster given live
- Ticket queue and chat for support by DESY
- Colleagues at HZI satisfied by onboarding and platform stability

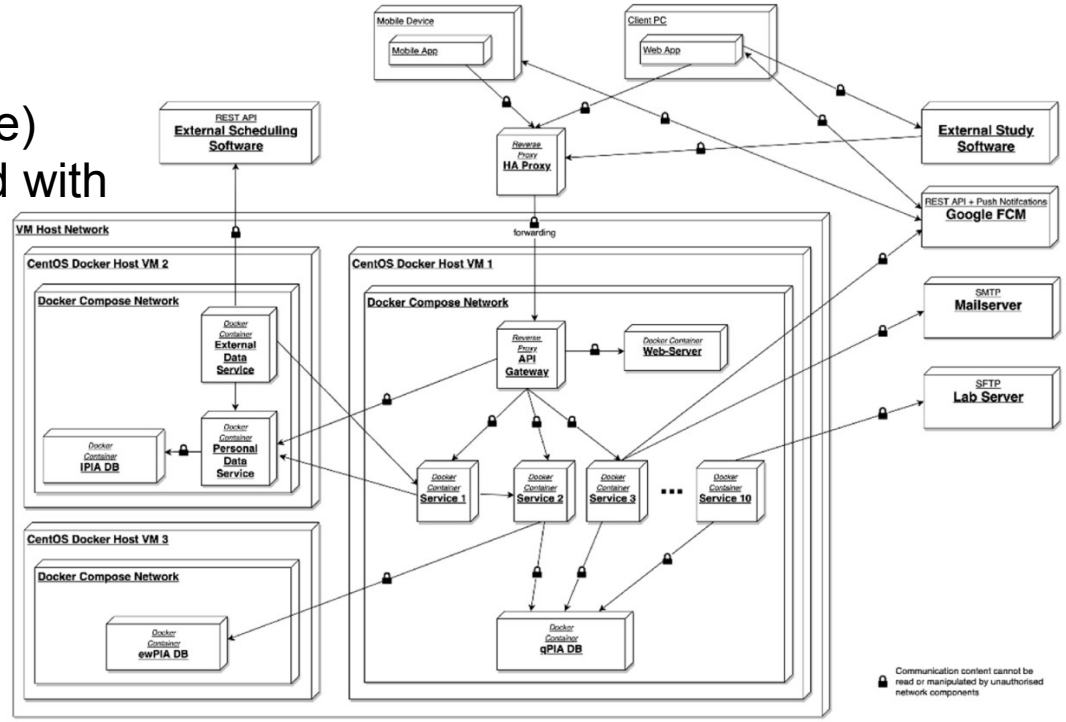

Fig. 1. Formalisation of network environment of PIA (basic version without connection to SORMAS).

Heise et.al. „Putting digital epidemiology into practice: PIA- Prospective Monitoring and Management Application“, 2022 Informatics in Medicine Unlocked 30

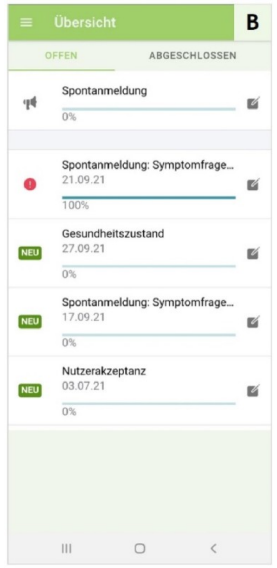

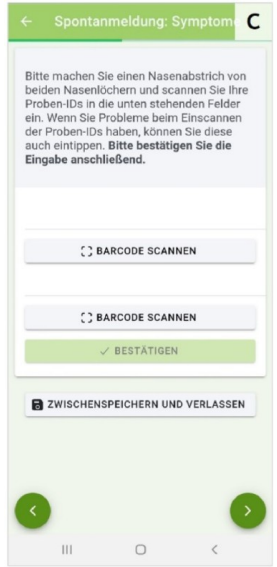

Sensor Management System – Controlled Vocabulary

Norman Ziegner, Daniel Sielaff – Helmholtz Centre for Environmental Research (UFZ)

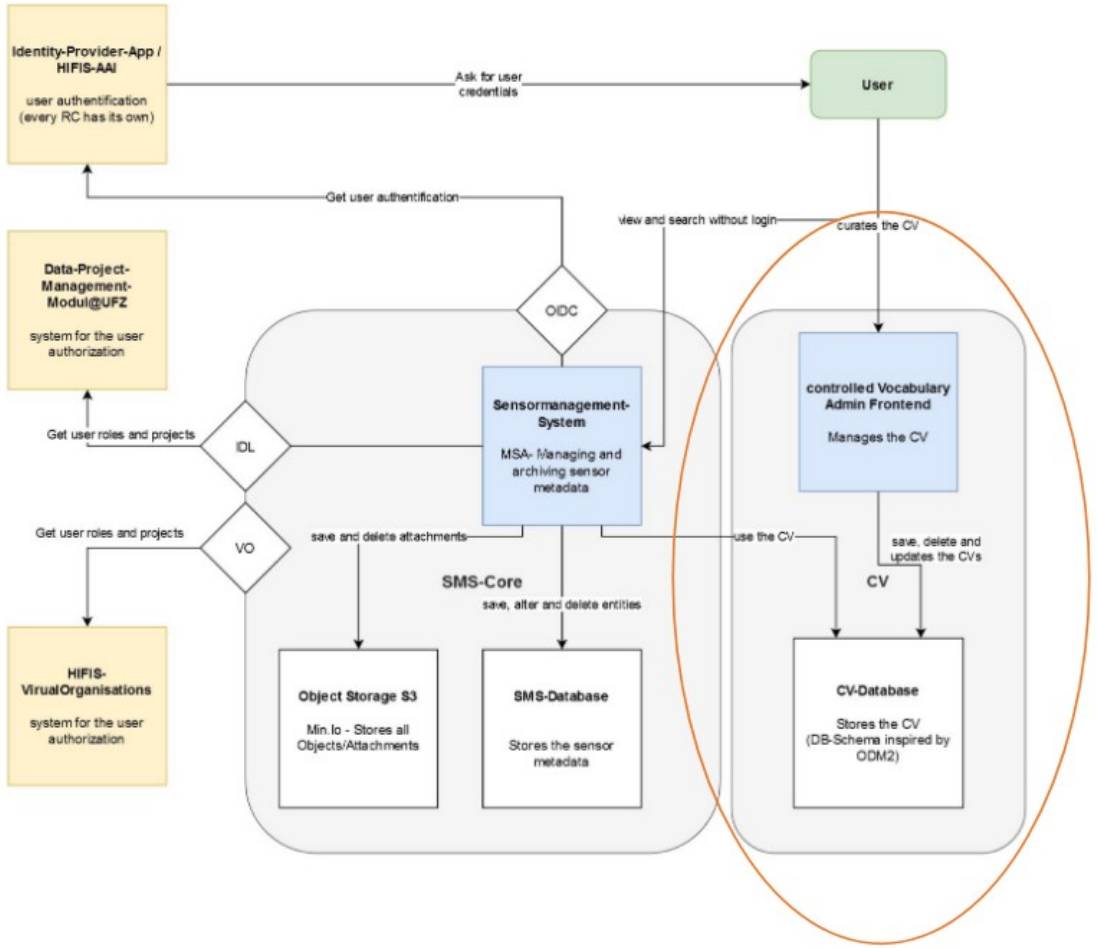

- Sensor Management for environmental research
- Distributed systems for measurement campaigns
- Sensor metadata administrated in SMS
- Metadata is shaped by controlled vocabulary (CV)
  - Ensures FAIRness
  - Instance hosted in Kubernetes@DESY

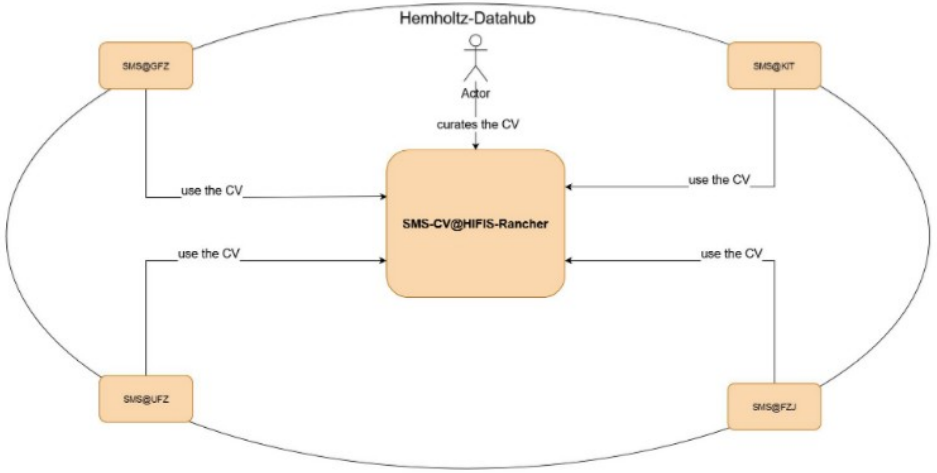

Remote Component Environment (RCE)

Robert Mischke – German Aerospace Center (DLR)

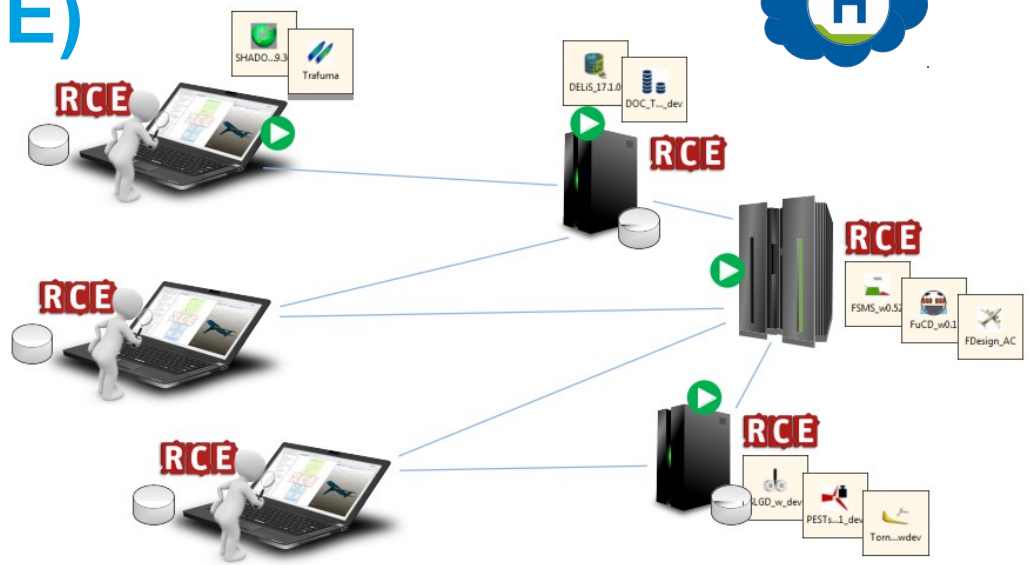

- Open Source integration environment for distributed workflows (https://rcenvironment.de/)
- Comes with GUI, batch mode, external tool integration, …
- Used for simulation and design of aircraft, ships and their impact on climate and environment

- Automated creation and disposal of test networks of distributed software (like RCE) for automated system testing on Kubernetes@DESY
- Automated deployments and chaos testing for resiliency optimization

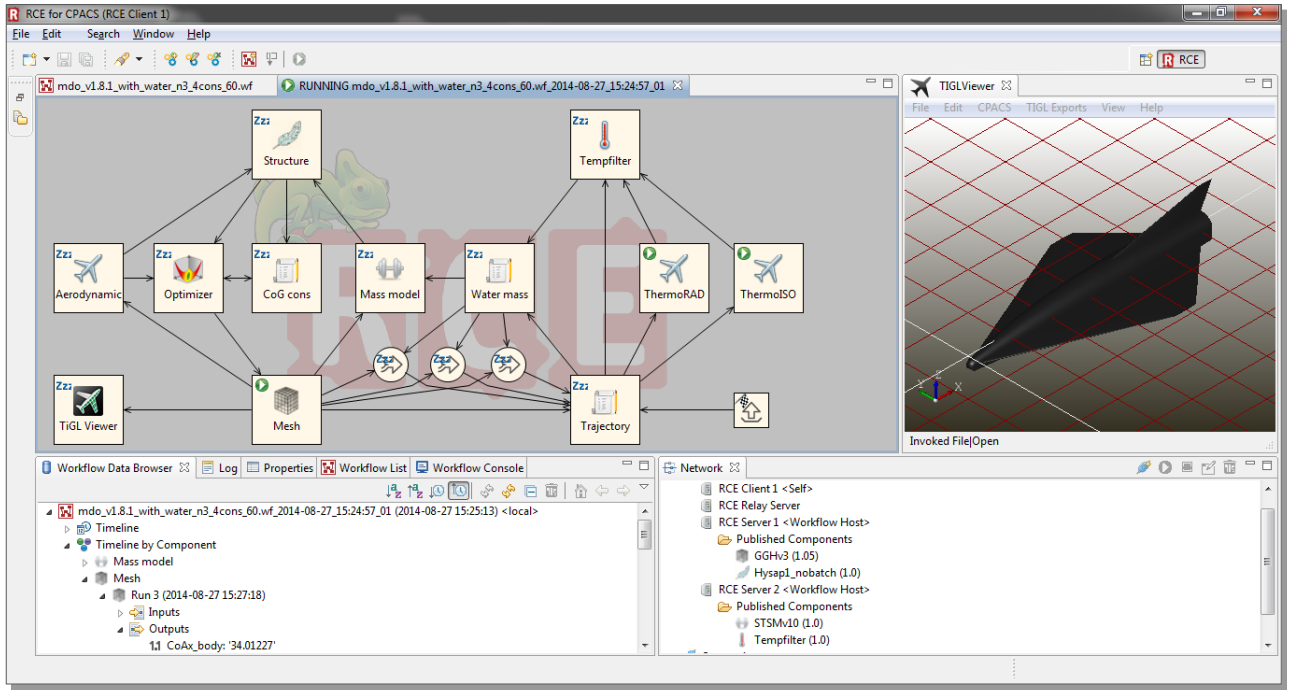

Summary

- Kubernetes clusters for use within DESY as well as for Helmholtz users via HIFIS
- Managed by Rancher instance for dynamical cluster deployments and additional RBAC
- Many internal & external use cases already hosted on clusters
  - Jupyter, Keycloak, HedgeDoc, Gitlab runners, …
  - Helmholtz Cloud Portal, VISA Portal, SciCat, …
  - Use cases from 5 other Helmholtz Centres

Outlook

- Exploring feasibility of k8s deployments on bare metal by DESY colleagues
- Better integration with Gitlab-agent (currently only for cluster admins)
- Optimization of cluster deployments with Terraform for easy base cluster spawning
- Strategy development for individual service backups
- Monitoring solution for external users‘ services
- ...

Thank you

Contact

DESY.

Michael Schuh
IT-RIC
michael.schuh@desy.de

Tim Wetzel
IT-RIC
tim.wetzel@desy.de

Johannes Reppin
IT-INFA
johannes.reppin@desy.de

Deutsches Elektronen Synchrotron
www.desy.de