CLAIX: One-Cluster-Concept for Convenient and Secure Scientific Computing at RWTH Aachen University

Tim Cramer, Christian Wassermann, Fabian Orland, Christian Terboven, Matthias S. Müller et al.

Hamburg, NUG 2024

NHR4CES NHR for Computational Engineering Science | RWTH AACHEN UNIVERSITY

RWTH Aachen University

A leading university with strong research

- One of the leading Technical Universities in Germany (TU9)
- One of eleven Germen Universities of Excellence
- Ranked among top 10 German universities in THE 2023
- One of the central nodes in the German Initiative for Research Data Management (NFDI)
- Host of many recognized centers: National High Performance Computing Center for Engineering Sciences (NHR4CES), ...

Studies and Teaching

Excellent Teaching, Learning and Assessment

- 47,078 Students
- 14,150 International Students
- 173 courses of study

Employees

- 10,272 Employees

NHR4CES NHR for Computational Engineering Science

IT Center | RWTH AACHEN UNIVERSITY

IT Center @ RWTH Aachen University

Mission

IT-Service Provider for RWTH Aachen University

- From network infrastructure to HPC systems
- E-Learning and SLCM
- Responsible to support Research Data Management at RWTH

National Mission

- **HPC for Computational Engineering Sciences (NHR4CES)**
- **Important node of the NFDI network**
- **Service provider for 42 universities in the state of North Rhine-Westphalia**

Staff

- 360 employees
  (111 scientists, 130 staff, 46 apprentices, 74 students)

Two major platforms to integrate HPC and RDM

NATIONALES HOCHLEISTUNGS RECHNEN

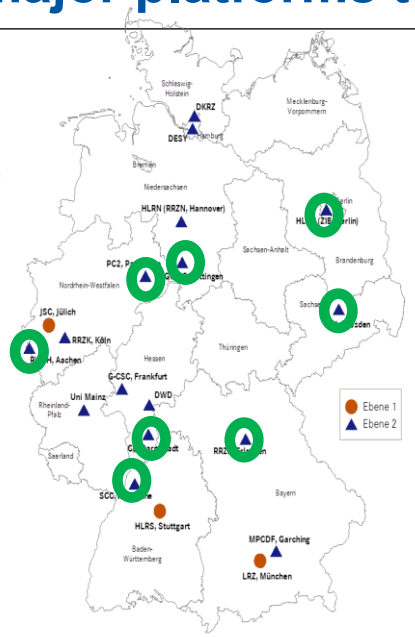

- **An integrated One-Cluster Concept for HPC**
  - Different Segmens for
    - Clasical HPC
    - Maschine Learning
    - Interactive usage, Jupyter Hub
  - Different storage classes for specific needs
    - Home file system for secure, permanent storage
    - Project file system for high bandwidth demand
    - On demand file system for high metadata performance

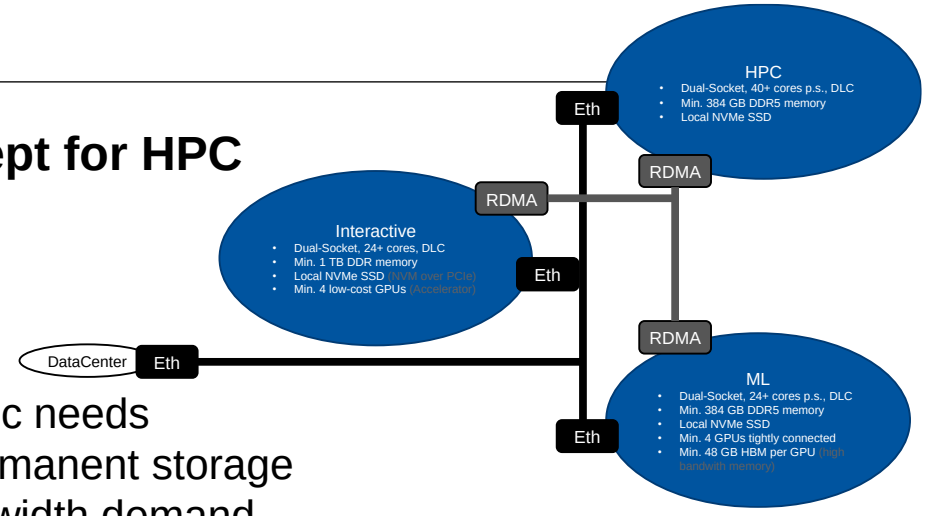

nfdi

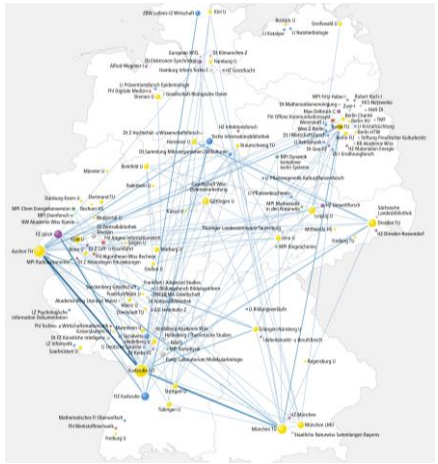

- **Coscine – the data management platform for FAIR data**
  - a data storage and linking environment for arbitrary data sources
  - Datastorage.nrw as a federated storage backend
  - implements FAIR principles based on the FAIR DO concept

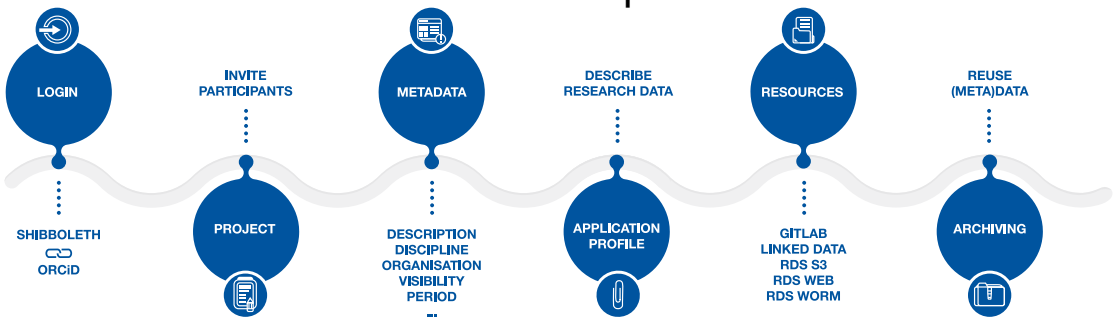

NHR4CES NHR for Computational Engineering Science

IT Center | RWTH AACHEN UNIVERSITY

Key HPC+AI projects @ RWTH

in addition to NHR4CES

Focus here: infrastructure funding sources

Key HPC+AI projects @ RWTH

WestAI

- One of four AI service centers (BMBF funding)
- Research focus: “Large-scale multi-modal transferable learning of complex AI models”

EDIH Rheinland

- European Digital Innovation Hub (EU funding)
- Gives SMEs in the Rhineland and the Euregio access to (HPC) expertise and knowledge.

Tier-3

- HPC+AI infrastructure for the state of NRW and/or institutes + research initiatives at RWTH

HPC systems @ RWTH

CLAIX 23 +25

Site-update: CLAIX-2023

- Pilot phase started January 2024 (production April 2024)
- 632 HPC nodes
  - 2 x Intel Xeon 8468 Sapphire Rapids (48 cores each)
  - 96 cores in total @ 2.1 GHz
  - 256 GB main memory (160 nodes with 512 GB, 2 nodes with 1 TB)
  - Interconnect: Infiniband NDR, 2:1 Blocking
- 52 ML nodes
  - CPU same as in HPC nodes, 512 GB main memory
  - 4 x NVIDIA H100
    - Memory per GPU: 94 GB, 2400 GB/s STREAM bandwidth
  - Interconnect: 2 x Infiniband NDR, 2:1 Blocking
- File systems:
  - High availability file system GPFS $HOME, capacity: 3.8 TB
  - Fast metadata + many files $BEEOND, capacity: 1.4 TB per HPC node, 648 TB per ML node
  - Fast parallel file system Lustre $HPCWORK, capacity: ~26 PB

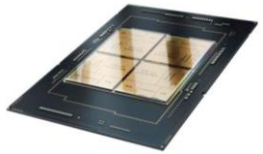

Intel Sapphire Rapids CPU, Source: Intel

NVIDIA H100, Source: NVIDIA

NHR4CES NHR for Computational Engineering Science

WEST AI KI-Servicezentrum

HPC.NRW

One-Cluster-Concept: integration of funding schemes

One-Cluster Concept: Compute and Storage for HPC and AI workflows, CLAIX-2023 and -2025

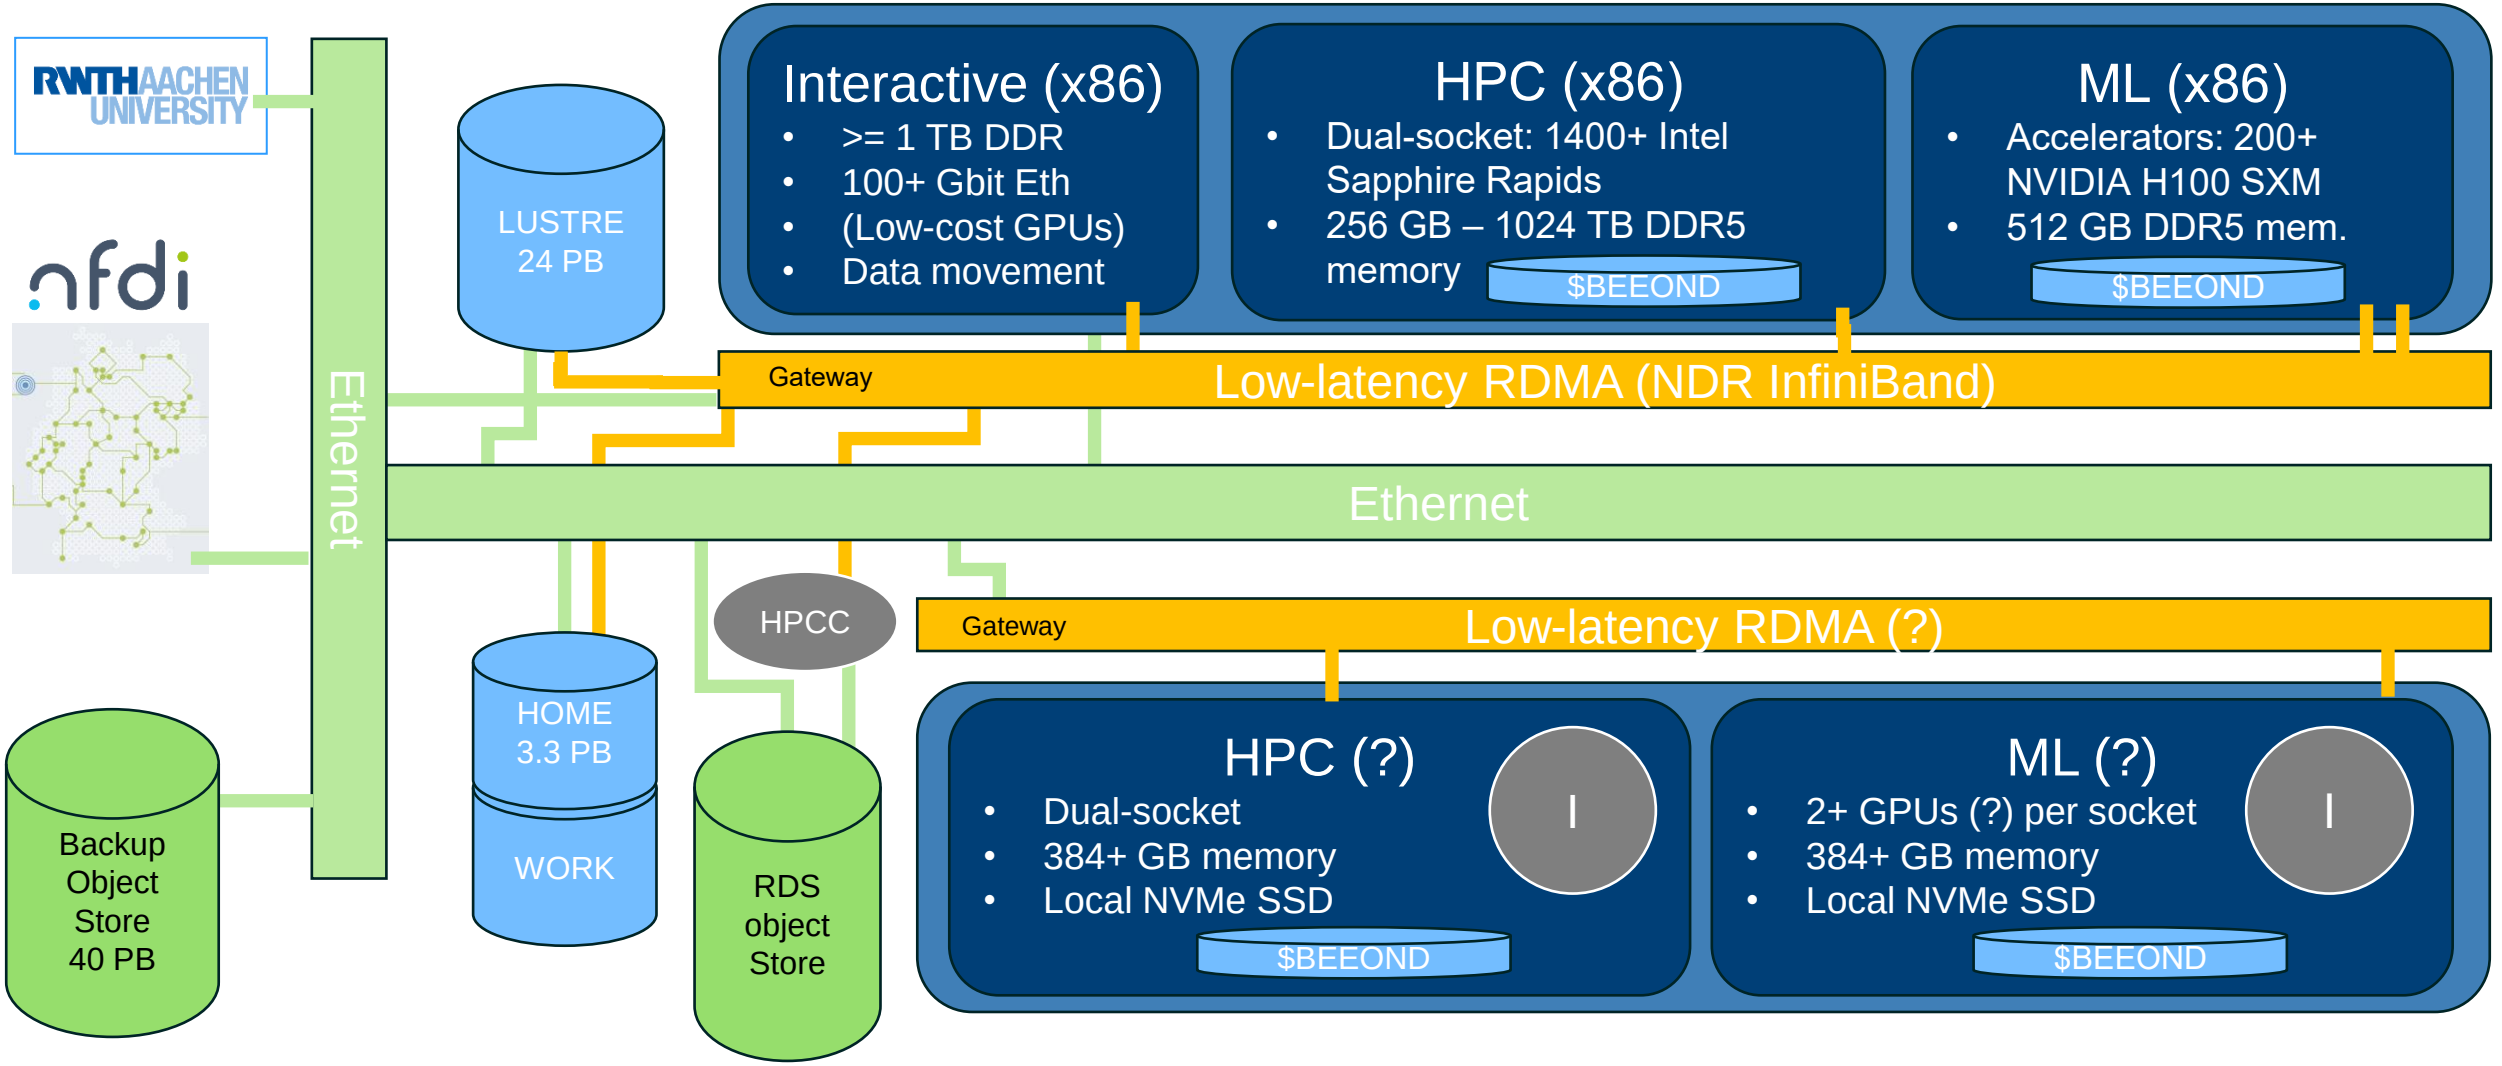

One-Cluster Concept: Compute and Storage for HPC and AI workflows, CLAIX-2023 and -2025

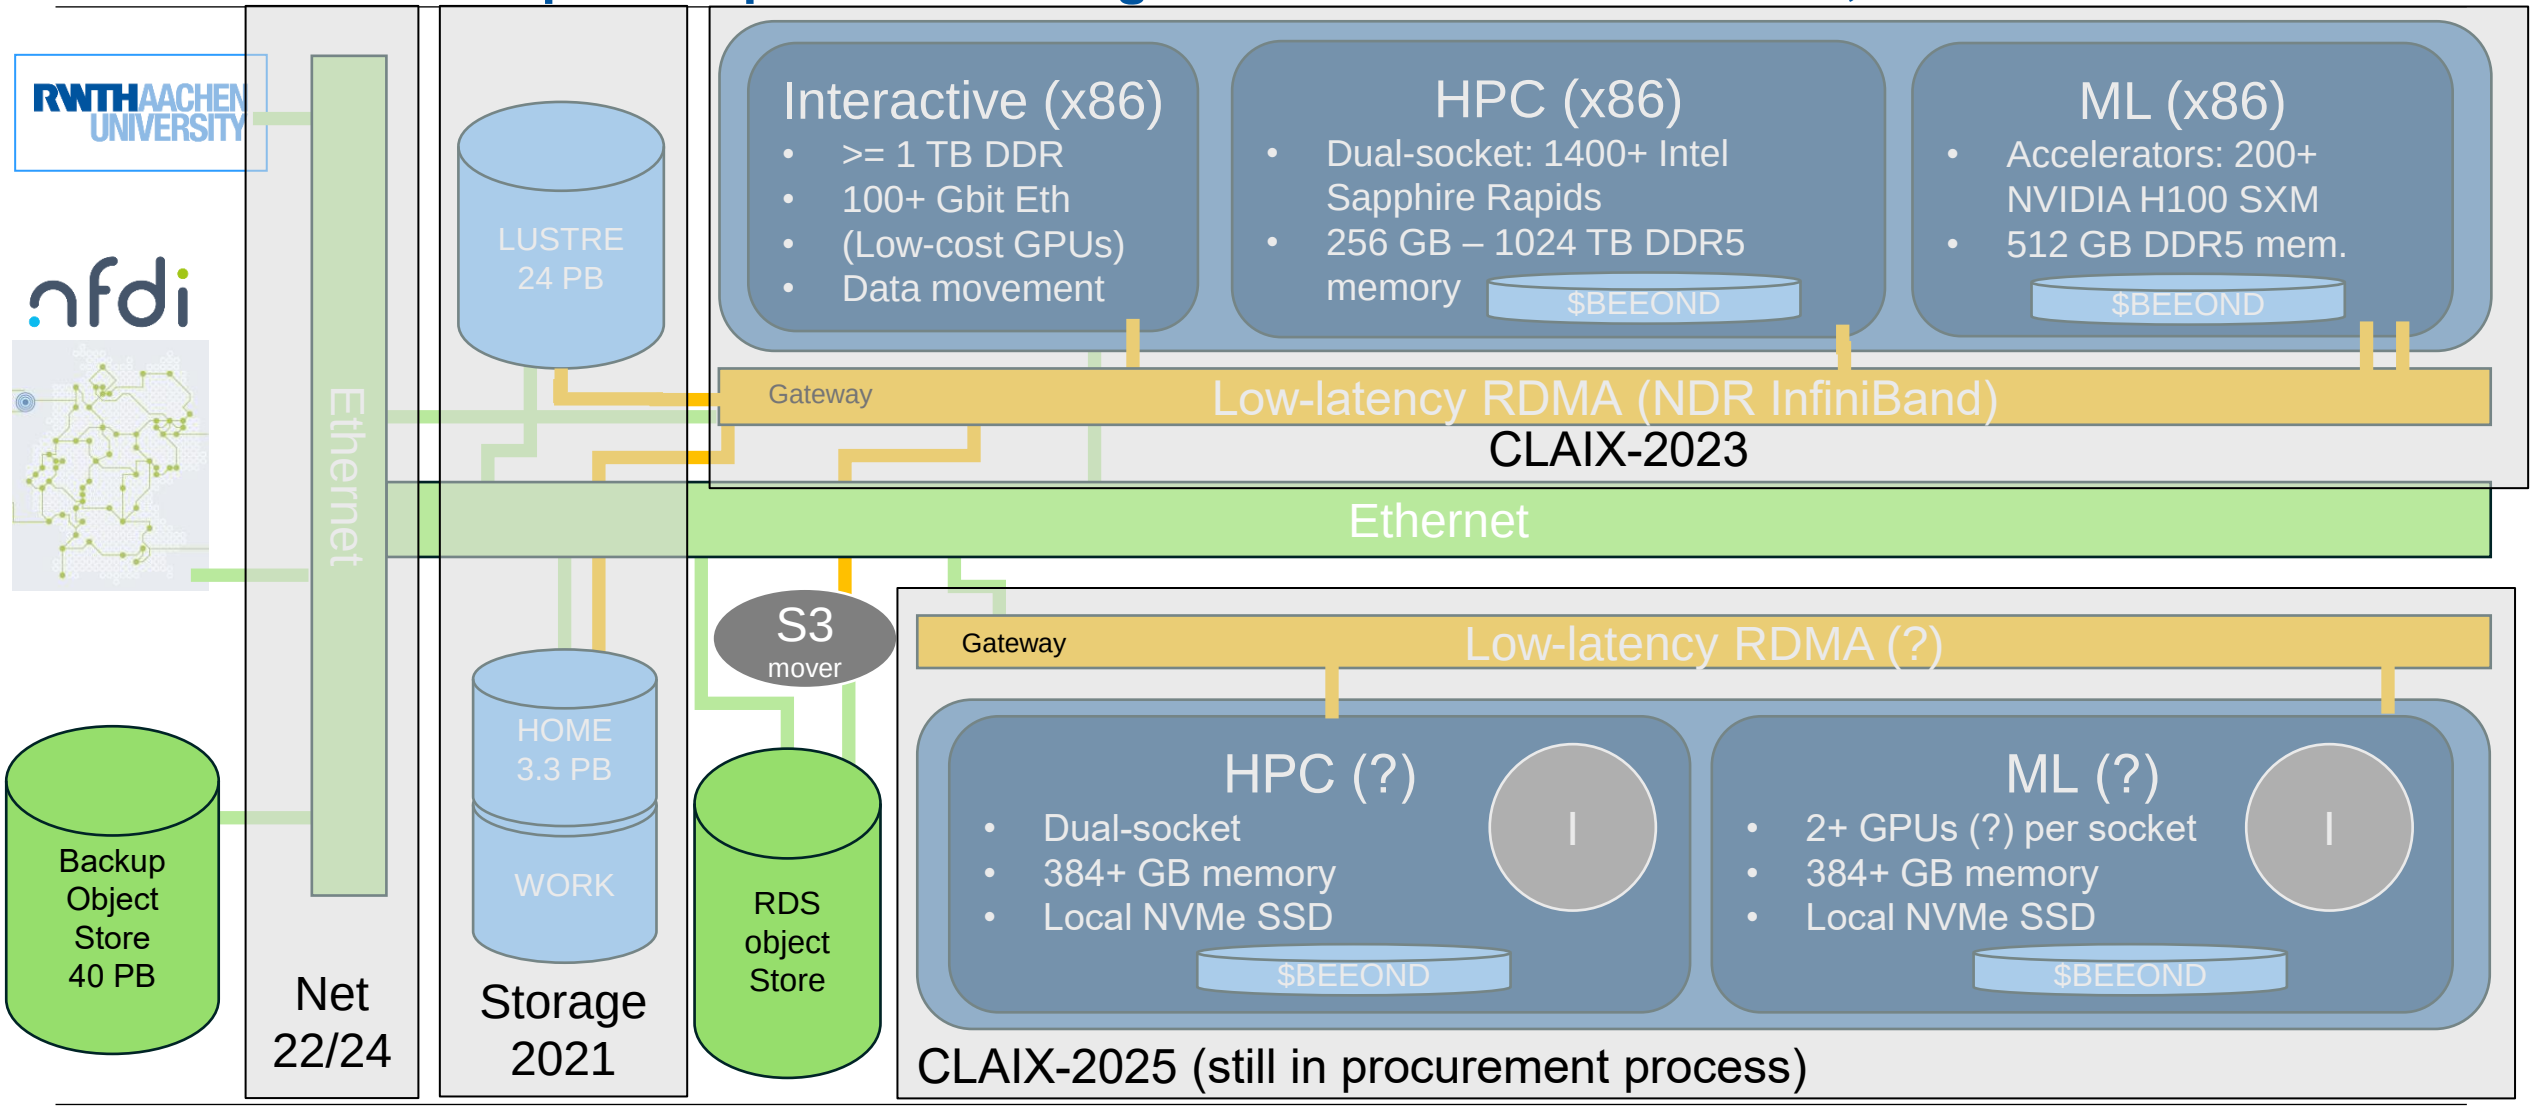

Cluster Performance Monitoring

Cluster Performance Monitoring @ RWTH

Cluster Performance Monitoring @ RWTH

Cluster Performance Monitoring @ RWTH

Cluster Performance Monitoring @ RWTH

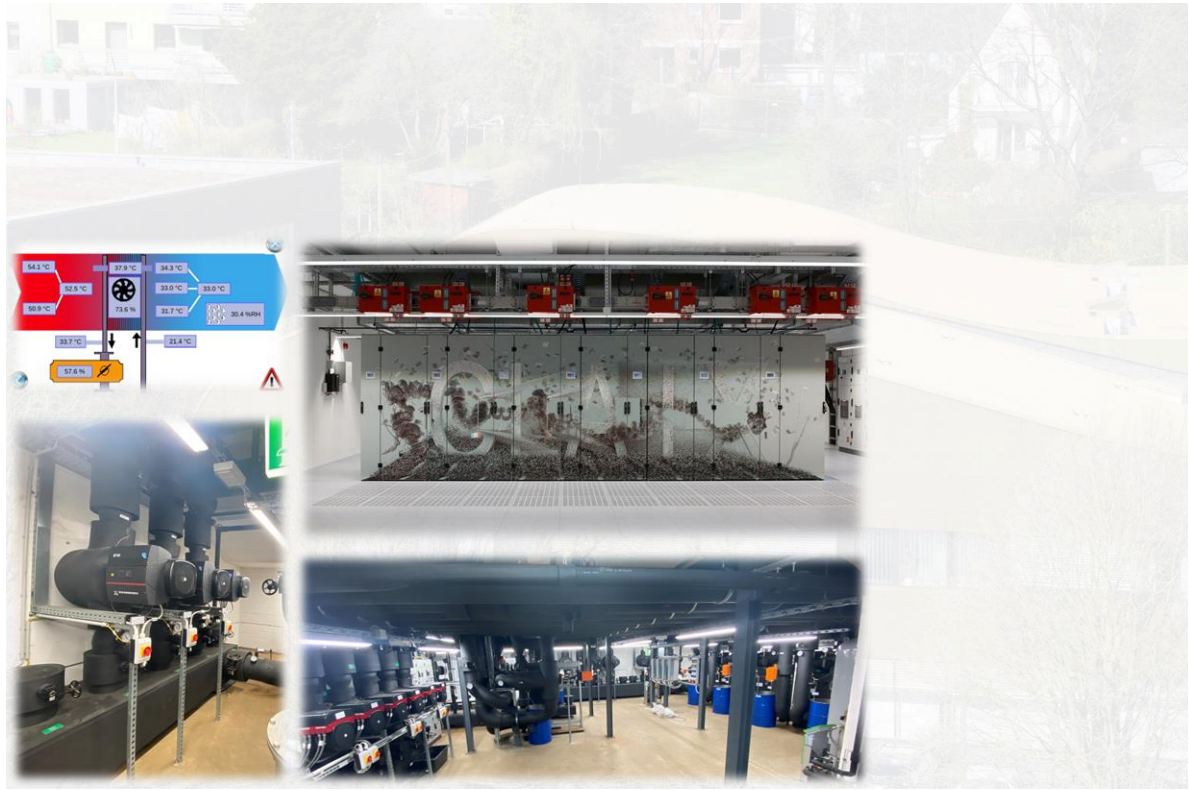

Cluster Performance Monitoring @ RWTH

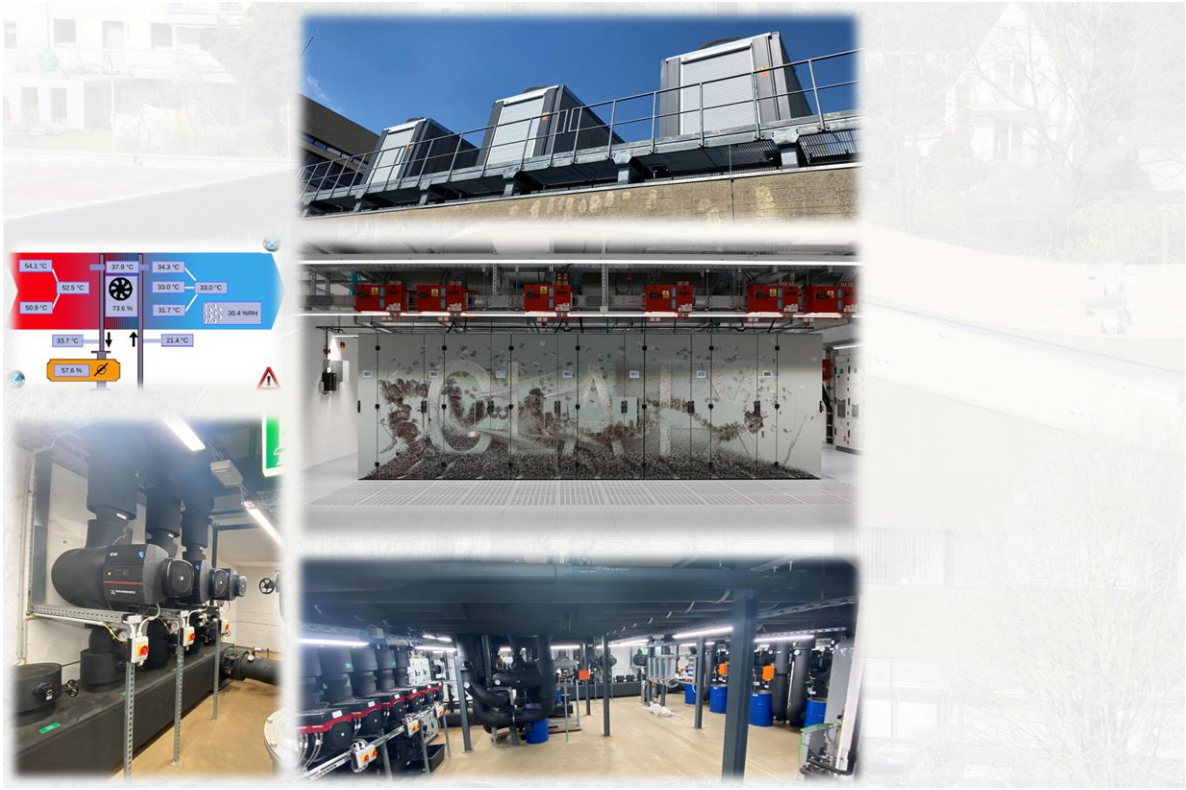

Cluster Performance Monitoring @ RWTH

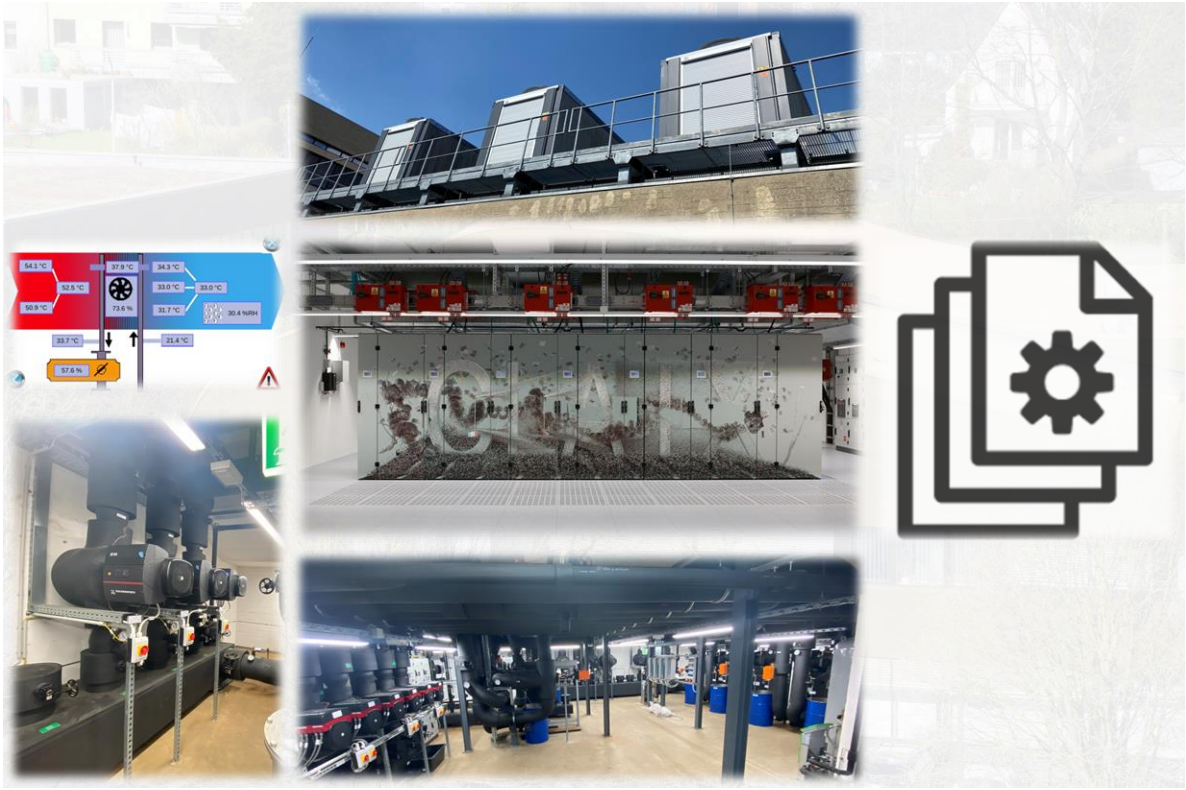

Cluster Performance Monitoring @ RWTH

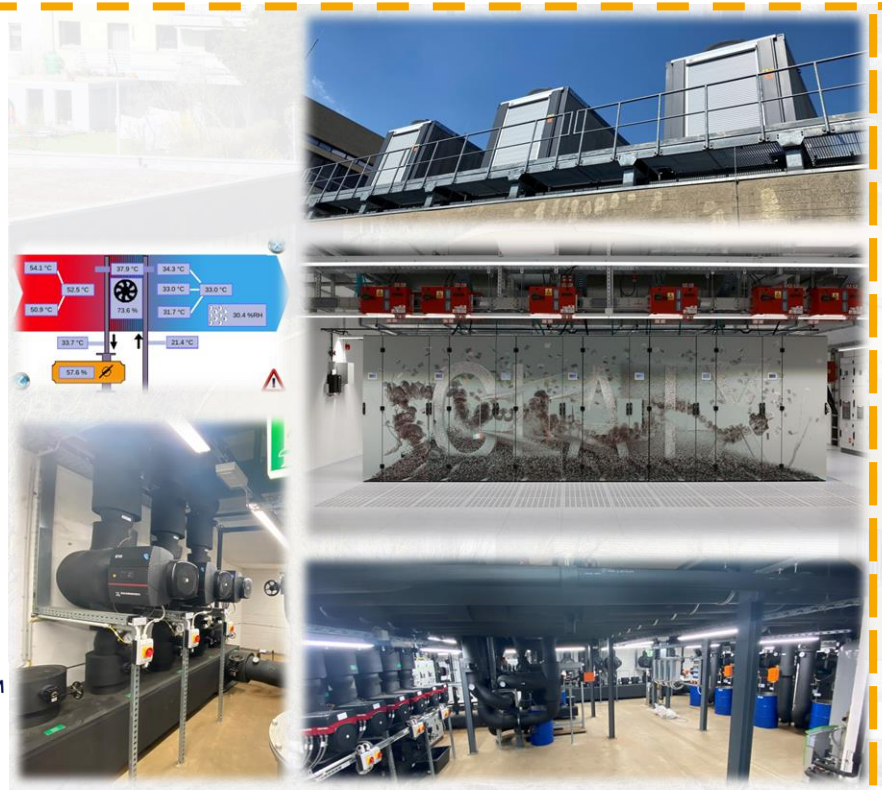

influxdb™

Time series database for cluster and infrastructure data

MariaDB

Relational database for job data

Digital Twin – Focus: Data integration

Holistic view of the data center including cooling, power supply and job load

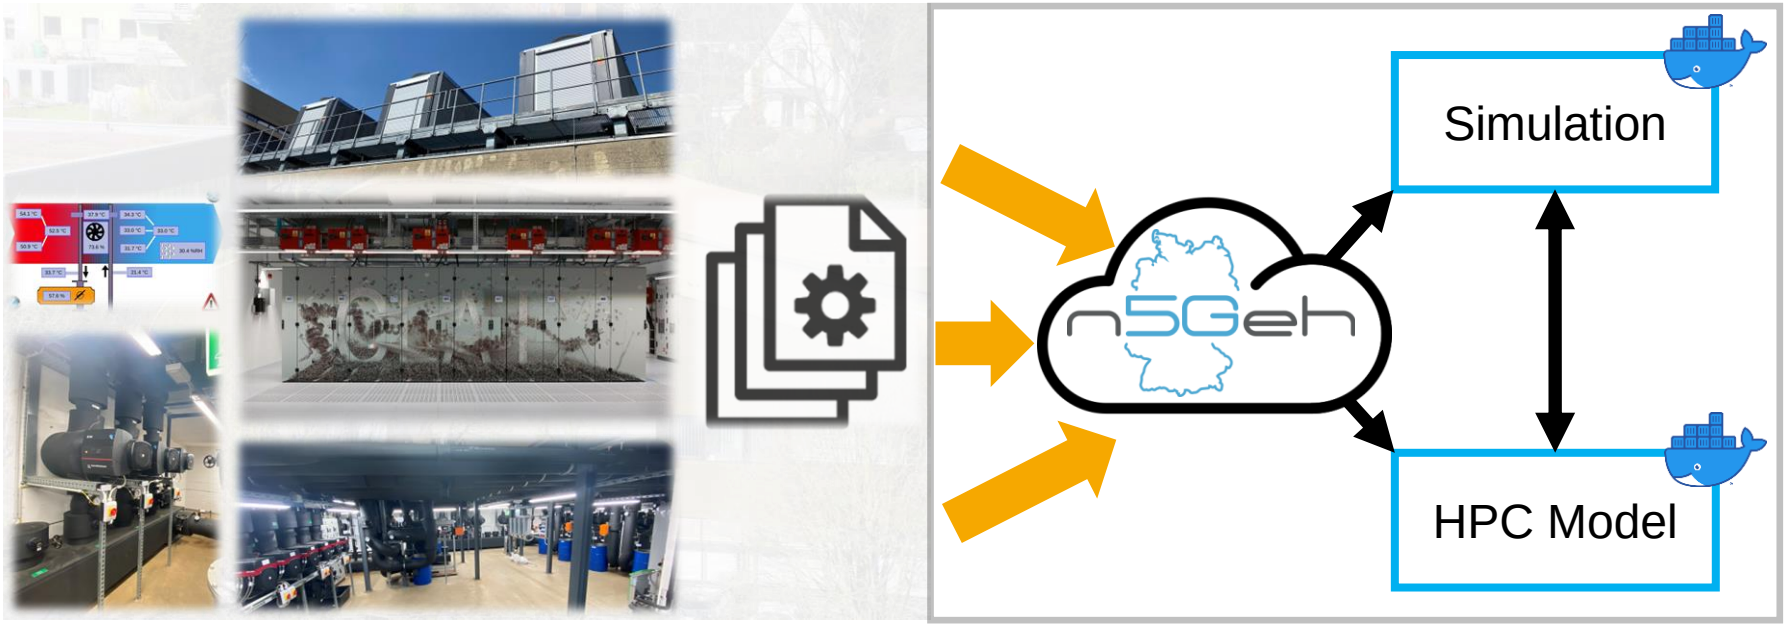

Physical data center **Digital Twin**

Accessing the Monitoring System – Joblist Overview

- History of Finished & Running Jobs

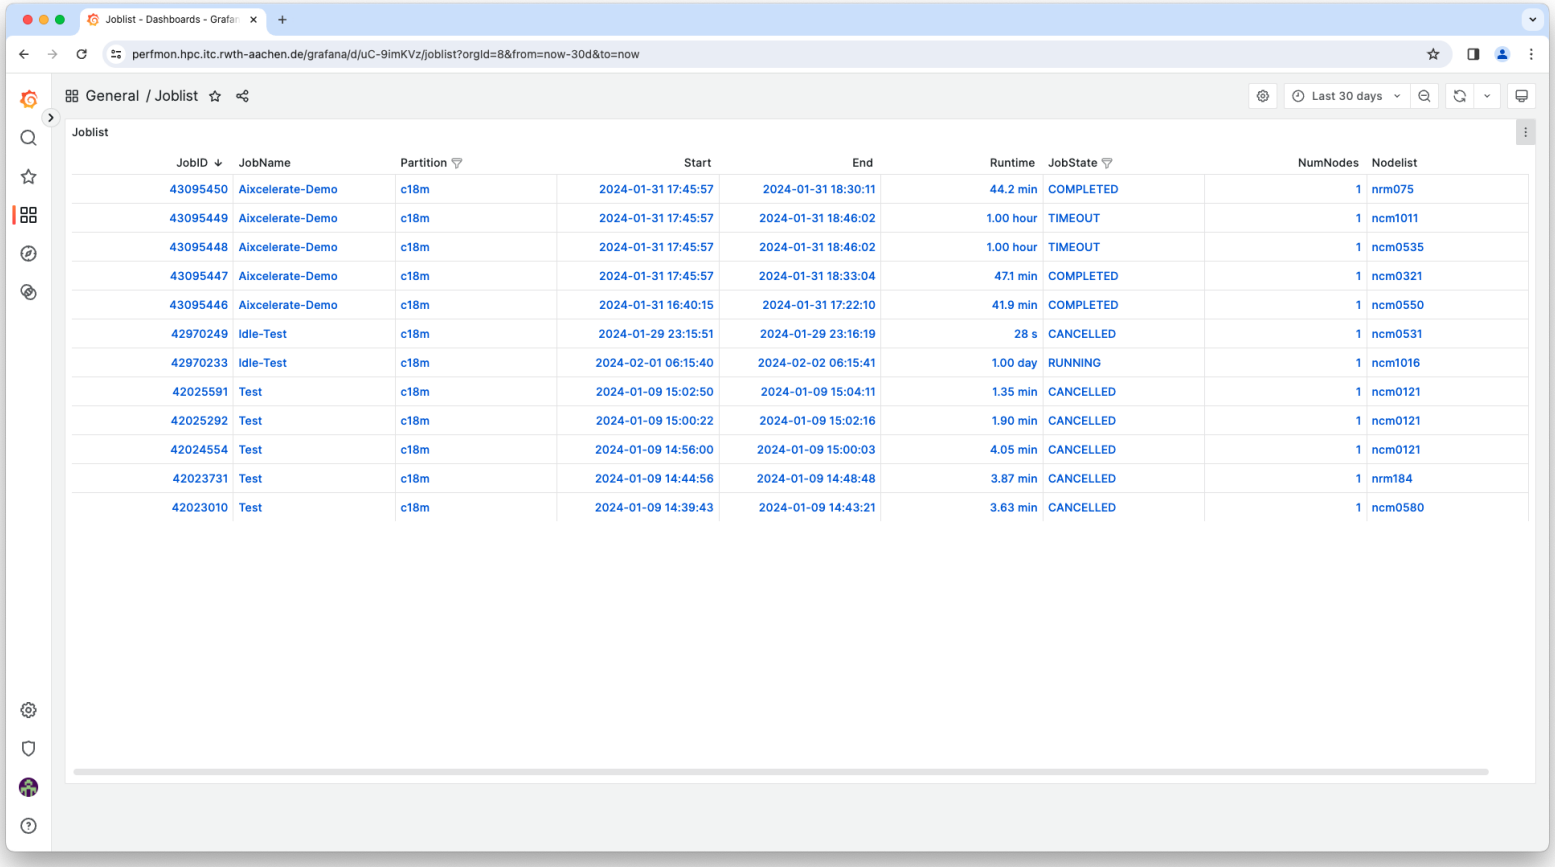

JobID	JobName	Partition	Start	End	Runtime	JobState	NumNodes	Nodelist
43095450	Aixcelerate-Demo	c18m	2024-01-31 17:45:57	2024-01-31 18:30:11	44.2 min	COMPLETED	1	nrm075
43095449	Aixcelerate-Demo	c18m	2024-01-31 17:45:57	2024-01-31 18:46:02	1.00 hour	TIMEOUT	1	ncm1011
43095448	Aixcelerate-Demo	c18m	2024-01-31 17:45:57	2024-01-31 18:46:02	1.00 hour	TIMEOUT	1	ncm0535
43095447	Aixcelerate-Demo	c18m	2024-01-31 17:45:57	2024-01-31 18:33:04	47.1 min	COMPLETED	1	ncm0321
43095446	Aixcelerate-Demo	c18m	2024-01-31 16:40:15	2024-01-31 17:22:10	41.9 min	COMPLETED	1	ncm0550
42970249	Idle-Test	c18m	2024-01-29 23:15:51	2024-01-29 23:16:19	28 s	CANCELLED	1	ncm0531
42970233	Idle-Test	c18m	2024-02-01 06:15:40	2024-02-02 06:15:41	1.00 day	RUNNING	1	ncm1016
42025591	Test	c18m	2024-01-09 15:02:50	2024-01-09 15:04:11	1.35 min	CANCELLED	1	ncm0121
42025292	Test	c18m	2024-01-09 15:00:22	2024-01-09 15:02:16	1.90 min	CANCELLED	1	ncm0121
42024554	Test	c18m	2024-01-09 14:56:00	2024-01-09 15:00:03	4.05 min	CANCELLED	1	ncm0121
42023731	Test	c18m	2024-01-09 14:44:56	2024-01-09 14:48:48	3.87 min	CANCELLED	1	nrm184
42023010	Test	c18m	2024-01-09 14:39:43	2024-01-09 14:43:21	3.63 min	CANCELLED	1	ncm0580

Accessing the Monitoring System – Joblist Overview

- History of Finished & Running Jobs

Increase Time Range

Filter by Cluster Partition

Sort by Runtime

Access Job Details

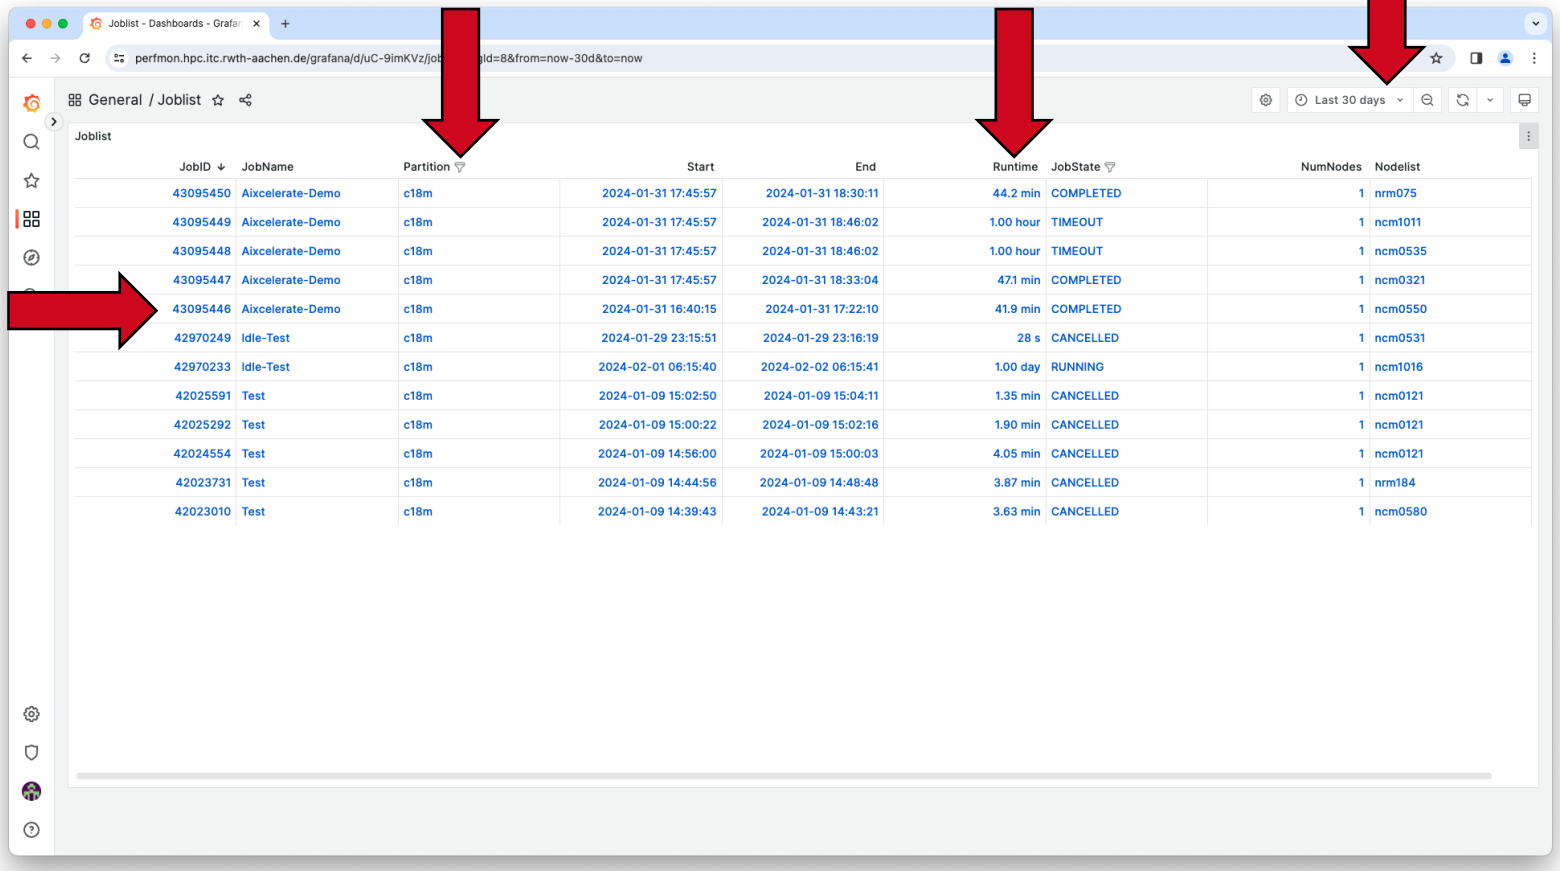

JobID	JobName	Partition	Start	End	Runtime	JobState	NumNodes	Nodelist
43095450	Aixcelerate-Demo	c18m	2024-01-31 17:45:57	2024-01-31 18:30:11	44.2 min	COMPLETED	1	nrm075
43095449	Aixcelerate-Demo	c18m	2024-01-31 17:45:57	2024-01-31 18:46:02	1.00 hour	TIMEOUT	1	ncm1011
43095448	Aixcelerate-Demo	c18m	2024-01-31 17:45:57	2024-01-31 18:46:02	1.00 hour	TIMEOUT	1	ncm0535
43095447	Aixcelerate-Demo	c18m	2024-01-31 17:45:57	2024-01-31 18:33:04	47.1 min	COMPLETED	1	ncm0321
43095446	Aixcelerate-Demo	c18m	2024-01-31 16:40:15	2024-01-31 17:22:10	41.9 min	COMPLETED	1	ncm0550
42970249	Idle-Test	c18m	2024-01-29 23:15:51	2024-01-29 23:16:19	28 s	CANCELLED	1	ncm0531
42970233	Idle-Test	c18m	2024-02-01 06:15:40	2024-02-02 06:15:41	1.00 day	RUNNING	1	ncm1016
42025591	Test	c18m	2024-01-09 15:02:50	2024-01-09 15:04:11	1.35 min	CANCELLED	1	ncm0121
42025292	Test	c18m	2024-01-09 15:00:22	2024-01-09 15:02:16	1.90 min	CANCELLED	1	ncm0121
42024554	Test	c18m	2024-01-09 14:56:00	2024-01-09 15:00:03	4.05 min	CANCELLED	1	ncm0121
42023731	Test	c18m	2024-01-09 14:44:56	2024-01-09 14:48:48	3.87 min	CANCELLED	1	nrm184
42023010	Test	c18m	2024-01-09 14:39:43	2024-01-09 14:43:21	3.63 min	CANCELLED	1	ncm0580

Accessing the Monitoring System – Job Details

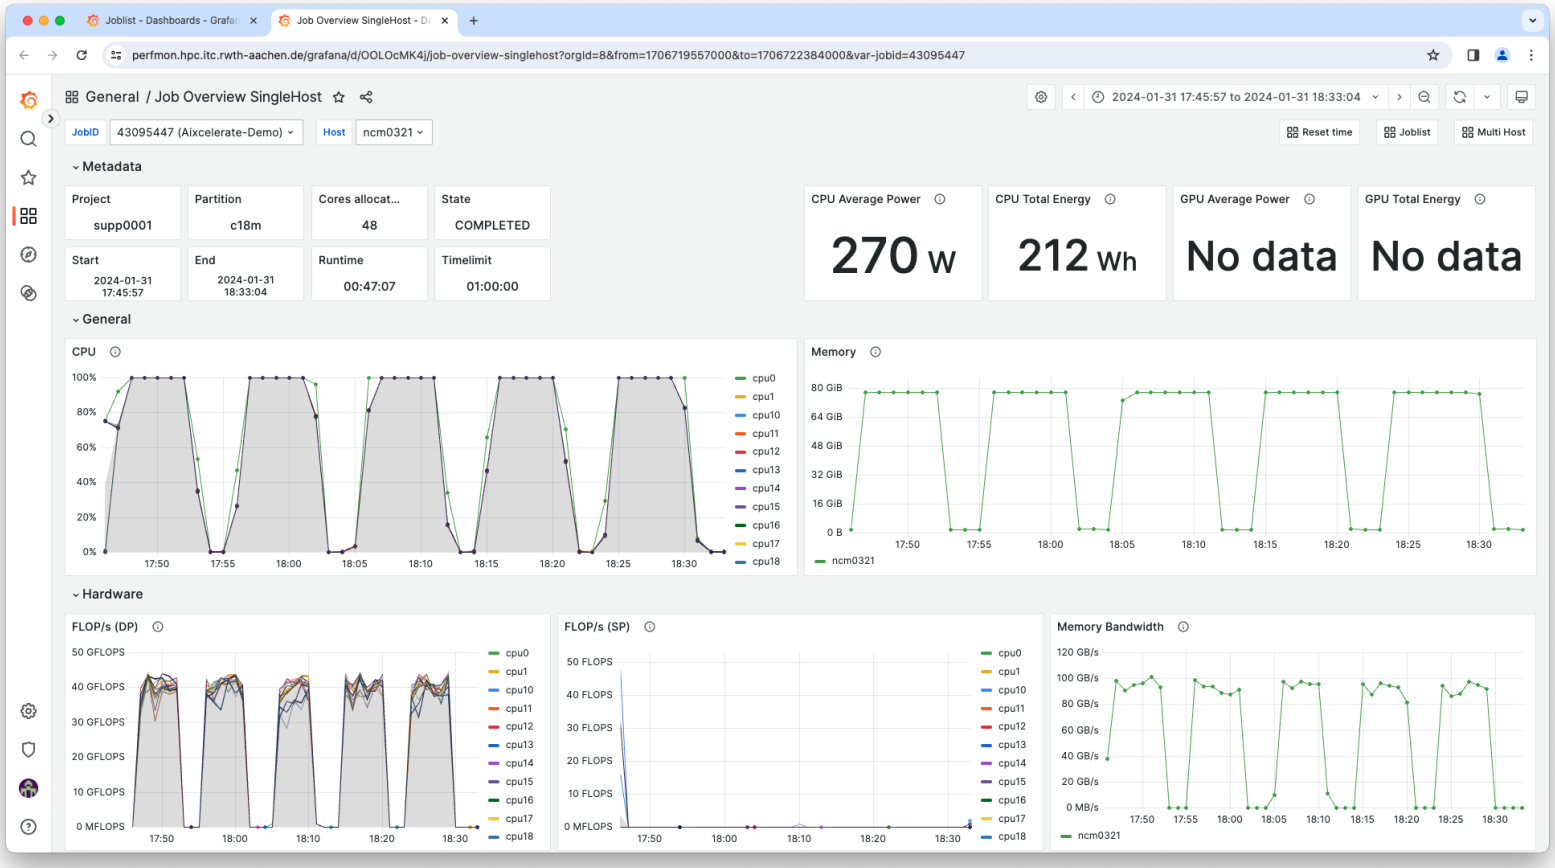

Accessing the Monitoring System – Job Details

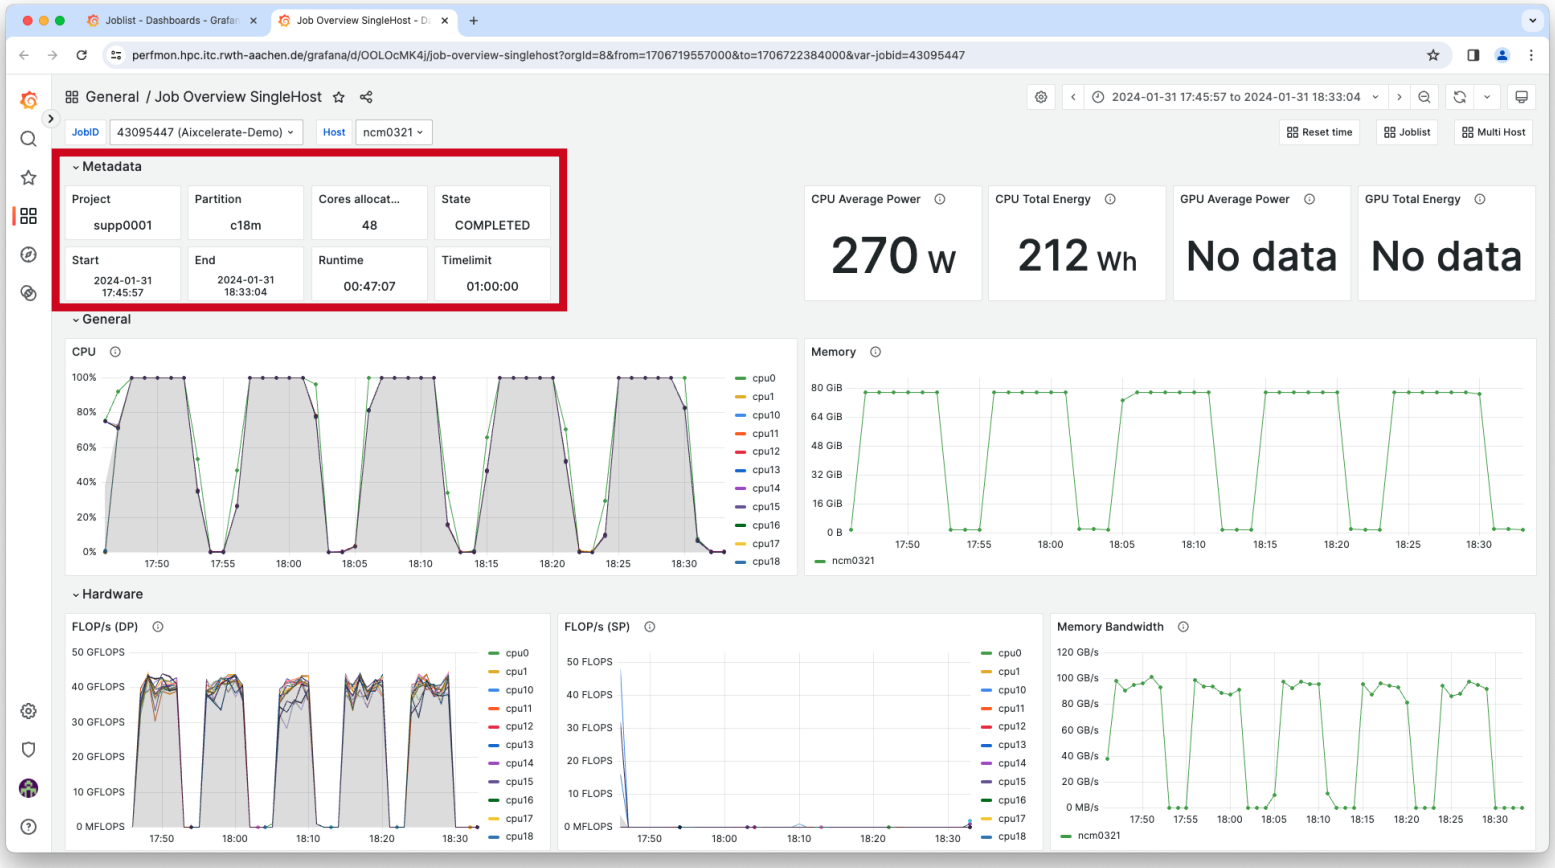

Accessing the Monitoring System – Job Details

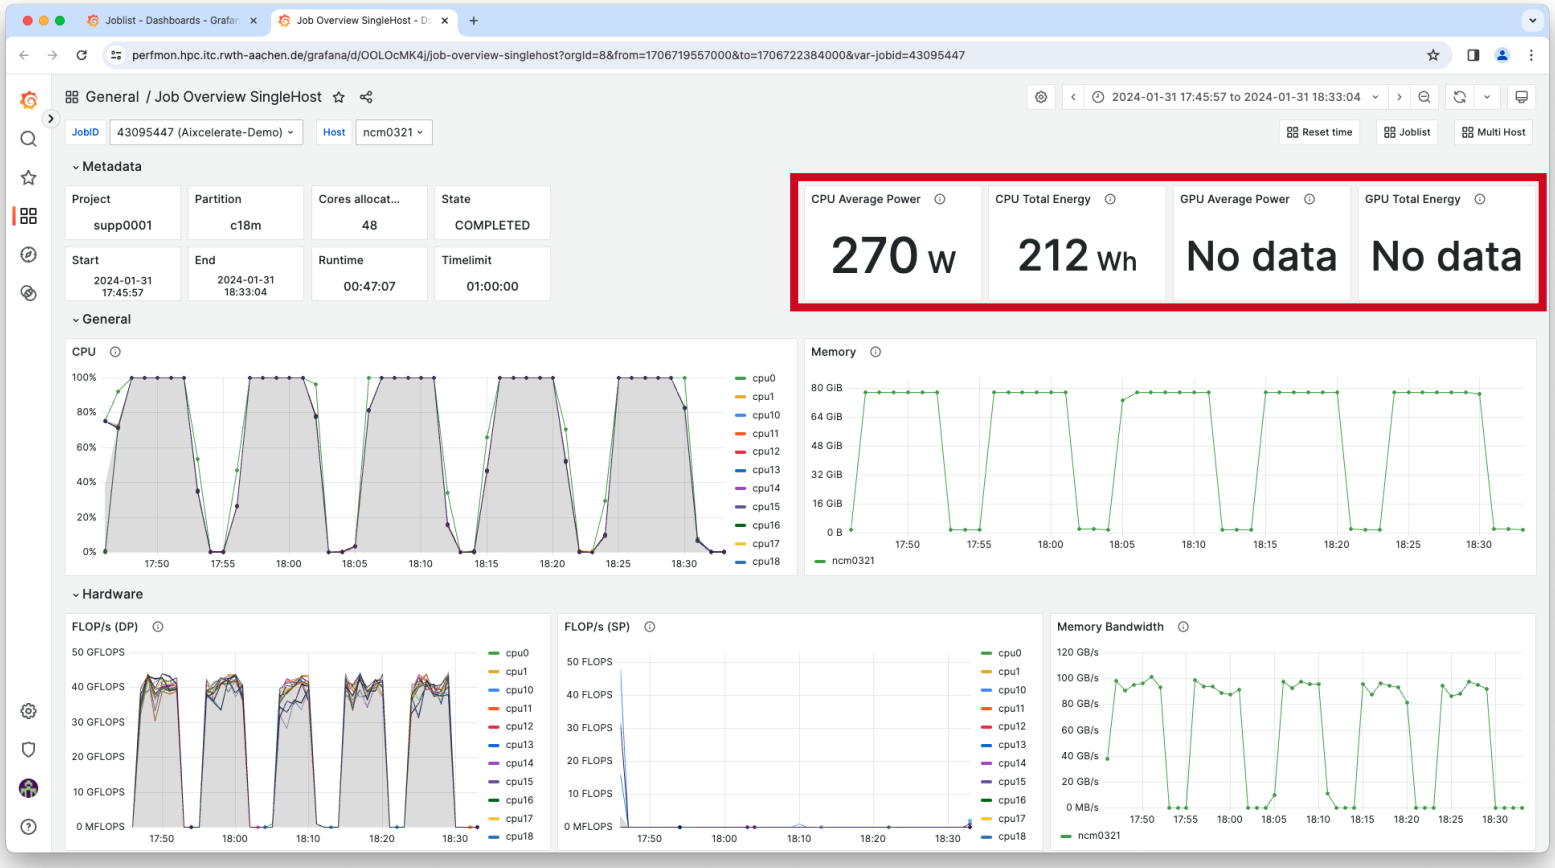

Accessing the Monitoring System – Job Details

Select Different Host

Switch Single-/Multi-Host View

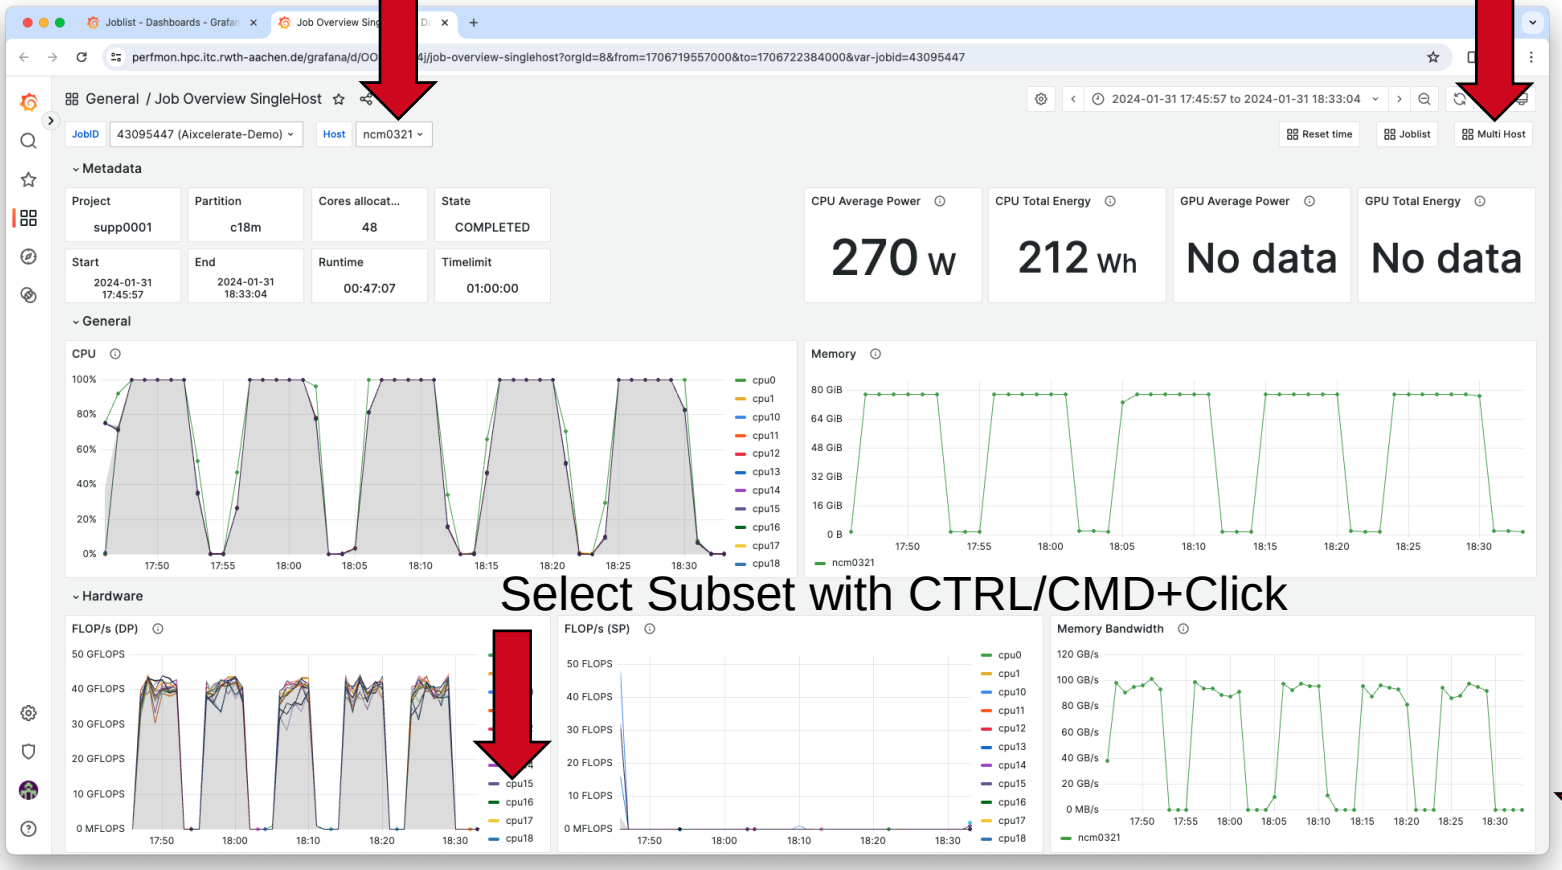

Select Subset with CTRL/CMD+Click

Scroll Down for More Details

Detecting Misconfiguration – Single Core Binding

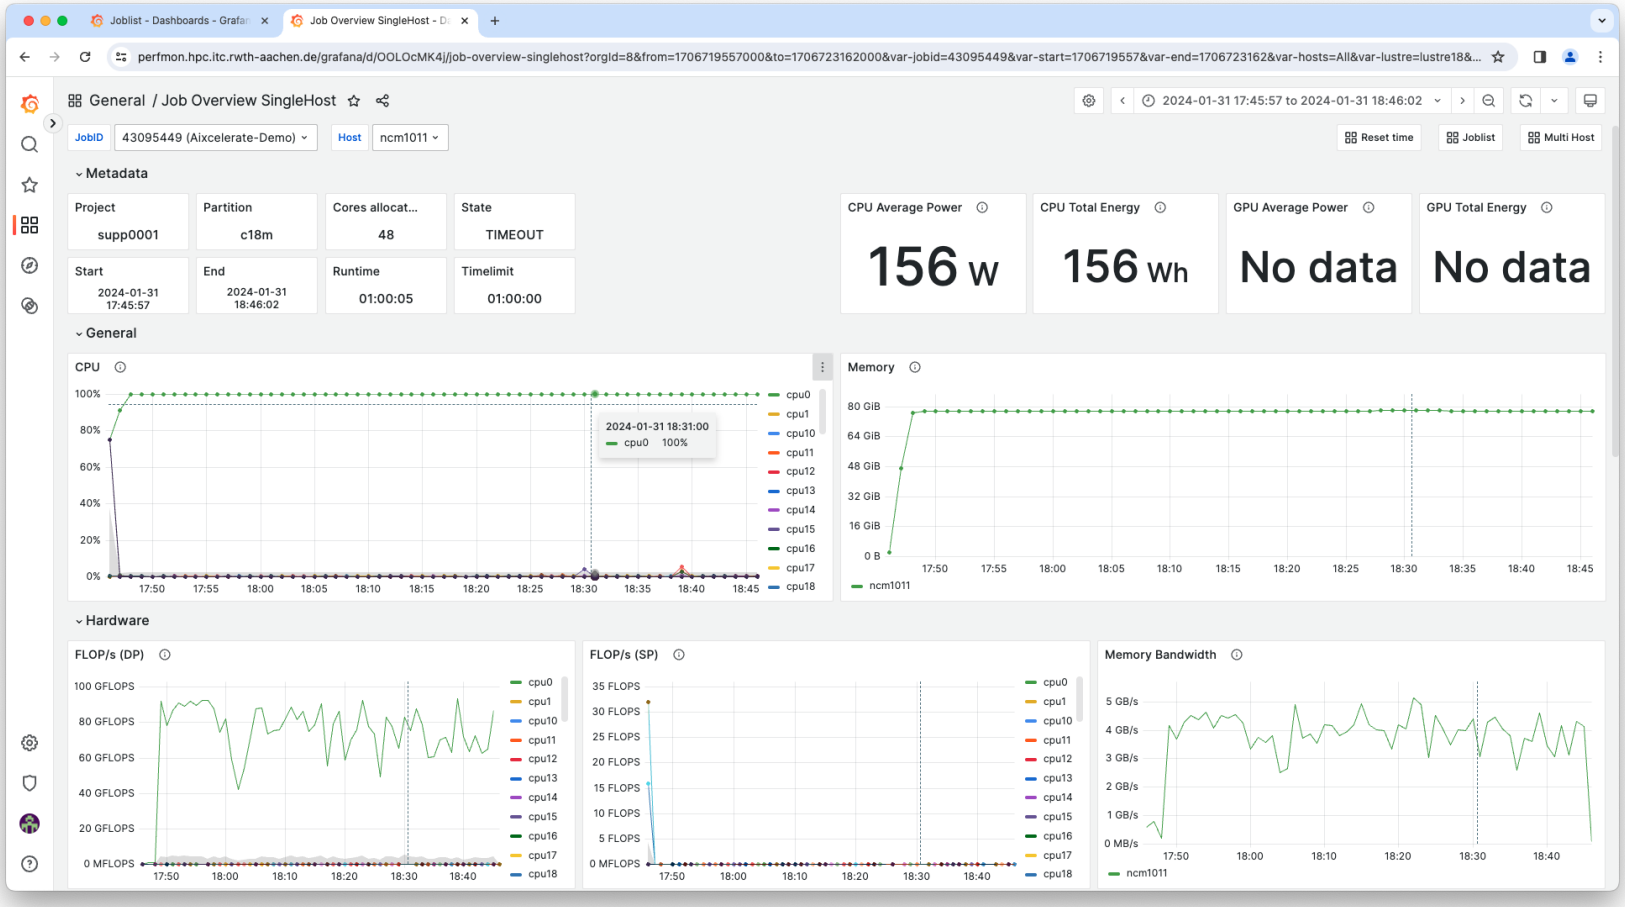

Detecting Misconfiguration – Single Core Binding

✅ 48 Cores Allocated

⚠ Only a Single Core Used

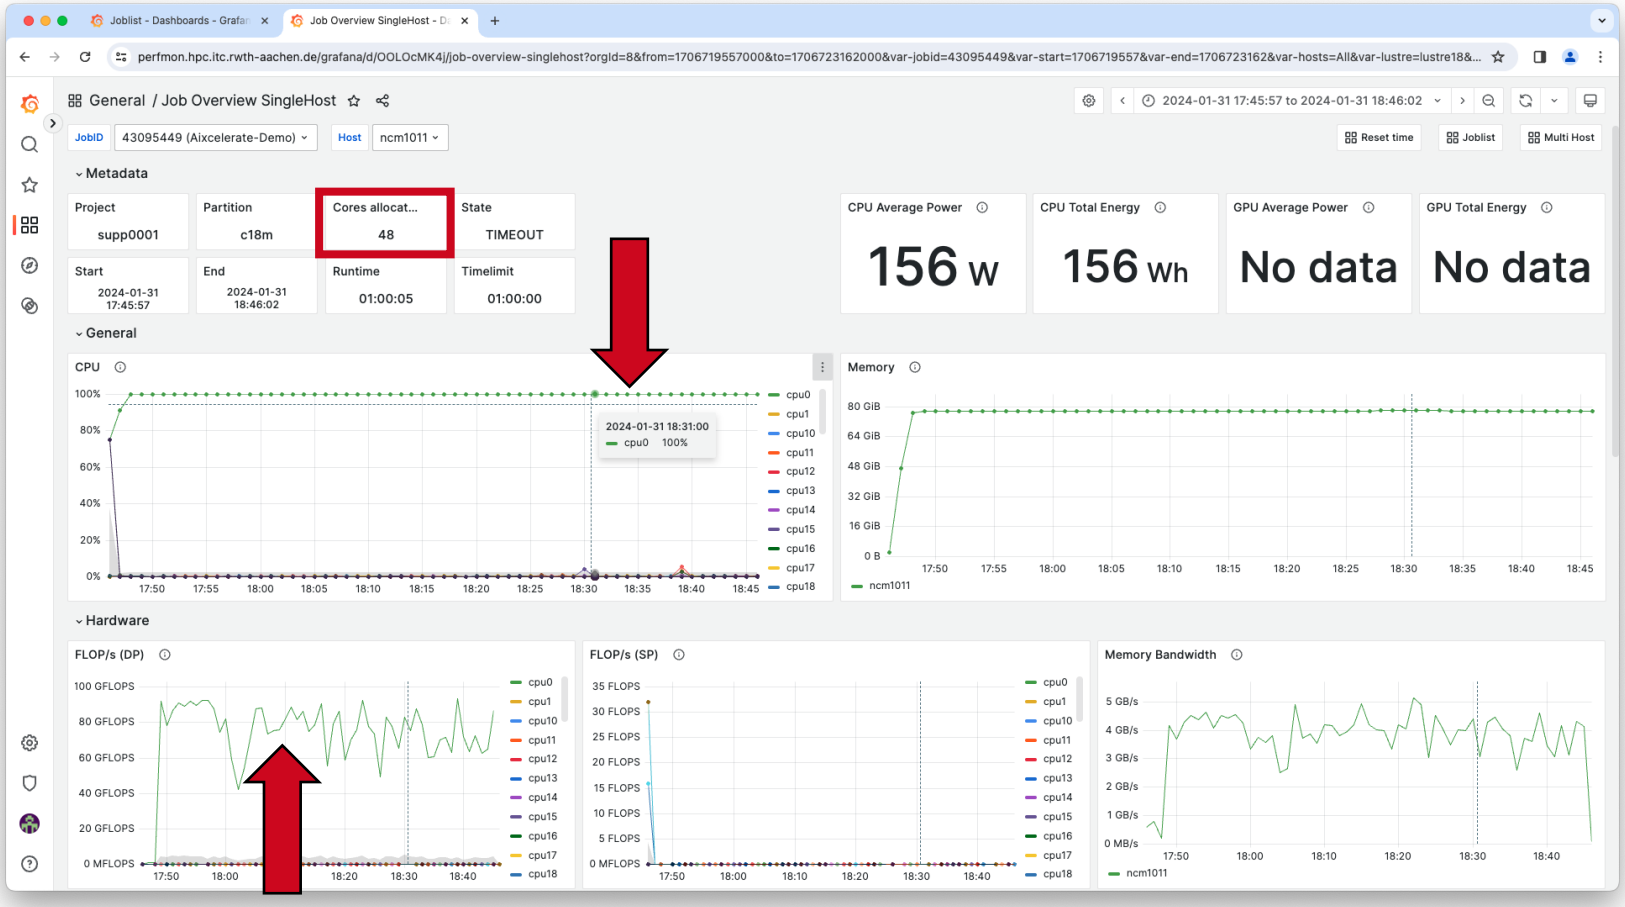

Detecting Misconfiguration – Single Core Binding

✅ 48 Cores Allocated

⚠ Only a Single Core Used

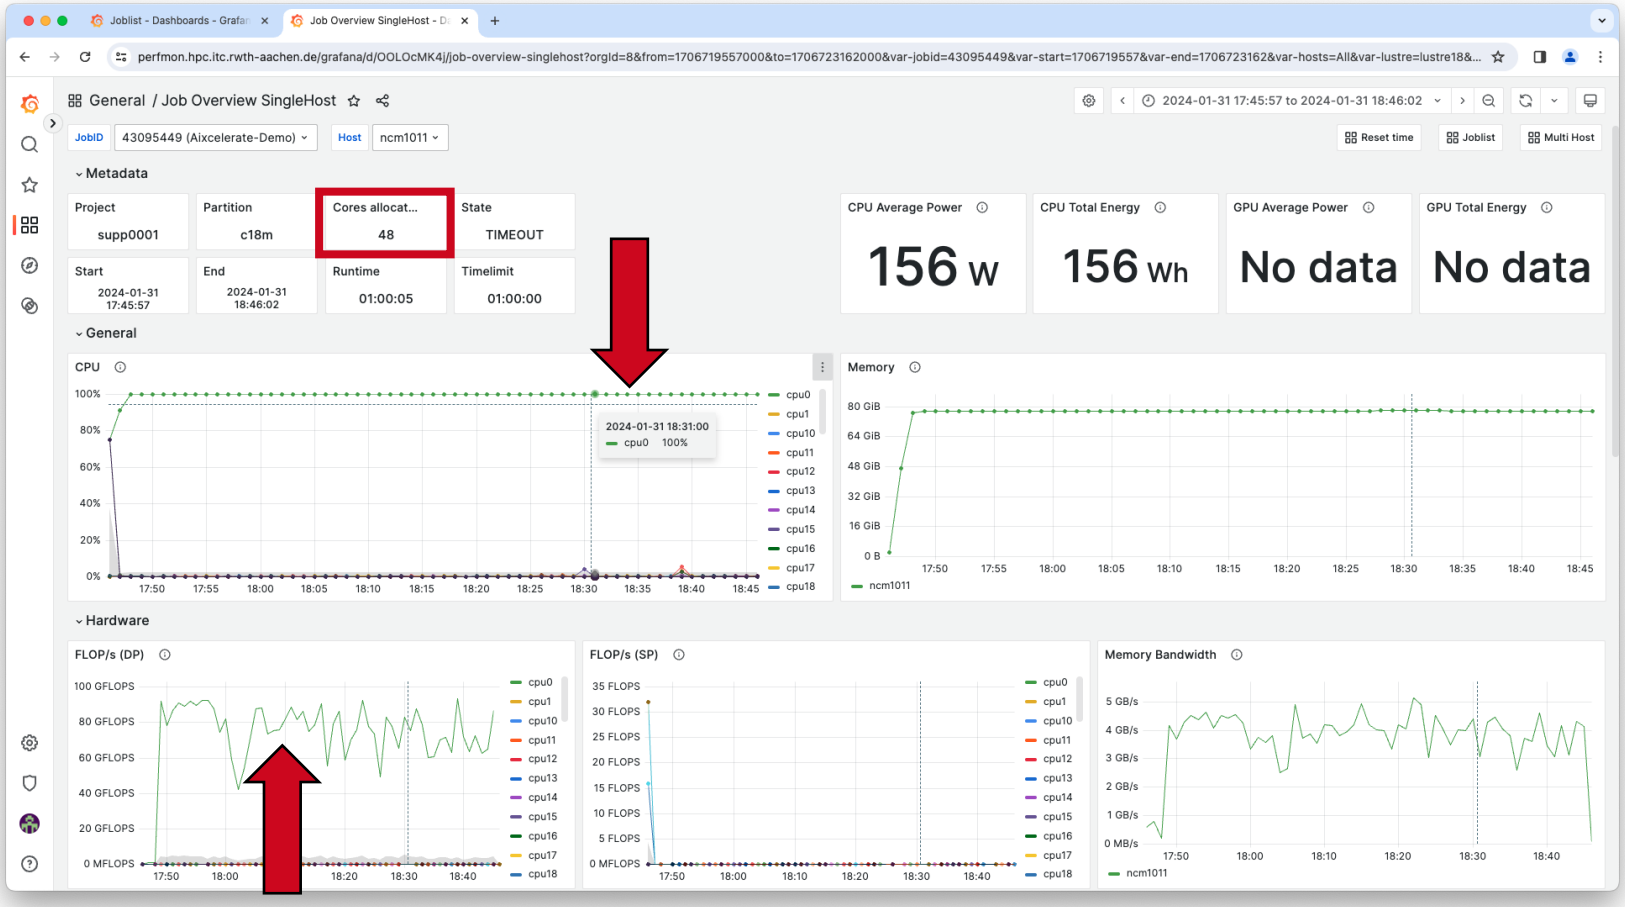

- Lessons Learned:
  - Check thread / process placement
- Consult documentation for
  - `OMP_PLACES`
  - `mpiexec / srun`

Thanks to Christian Wassermann, PhD student @ RWTH

Security Concept & Federal Identity Management

Security Concept for CLAIX

Motivation

- In May 2020 a dozen of Europe's supercomputers were taken offline simultaneously due to an attack
- Hacked user accounts used in combination with a root exploit
- Password less ssh keys used to jump from one system to the others
- Many other examples for compromised IT systems

Source: Freepik

Concept @ RWTH

- Limited IP range only (basically DFN + selected German universities)
- Deployment of critical patches ASAP
- Open communication / cooperation within HPC.NRW, NHR and DFN Cert
- 2FA for
  - RWTH VPN connections (in production)
  - The ssh connection to the cluster (in production)
- One FTE dedicated for the cluster security

CLAIX Identity Management

IDM Infrastructure

- Introduction of RegApp [1] @RWTH: Janurary 2023
- Provisioning and deprovisioning of user identities
- Transparentconnection of SAML federations
- SSO protocol proxy OIDC ↔ SAML
- Infrastructure proxy for LDAP services
- Solution for RWTH and NHR

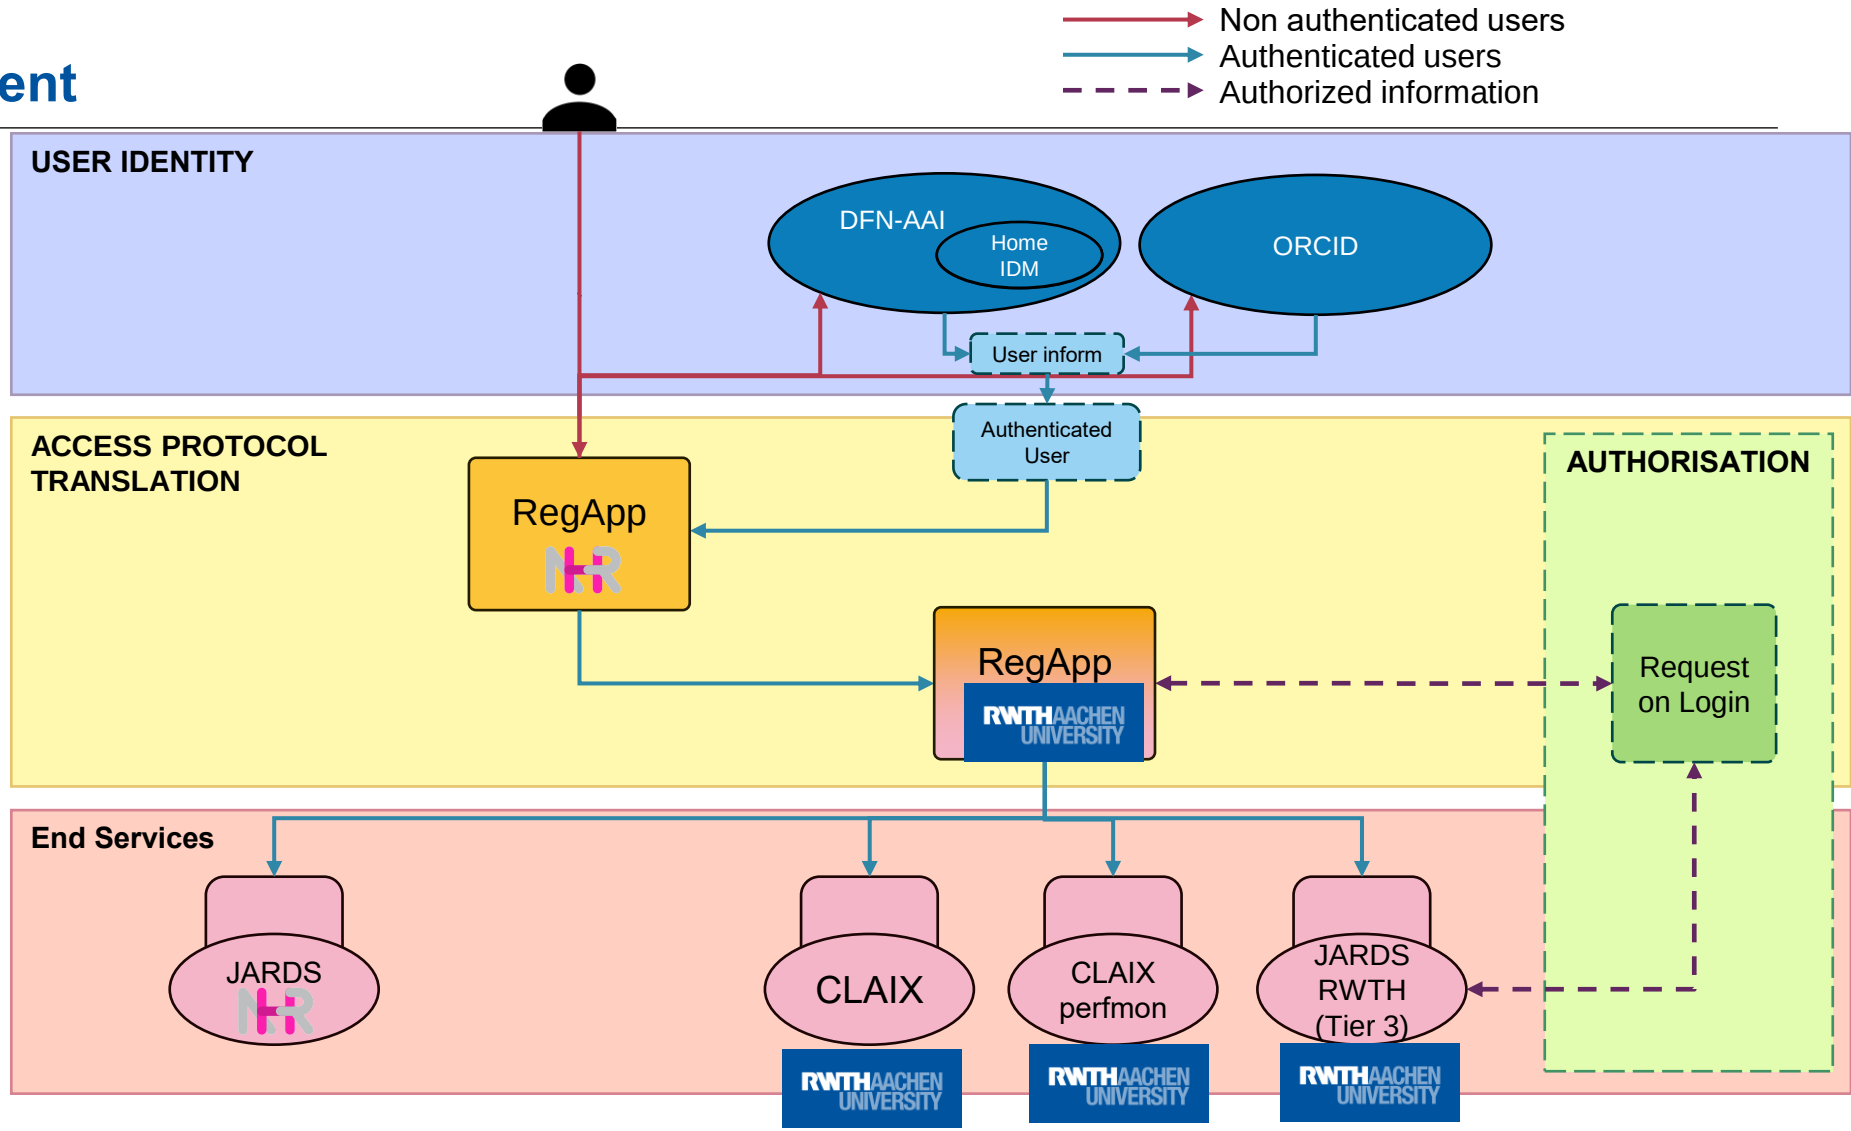

[1] Developed by Karlsruhe Institute for Technology (KIT) https://www.scc.kit.edu/en/services/regapp.php

Introduction 2FA

Experiences

- Migration of > 8000 accounts migrated
- Pilot Phase > one year
  - Technical challenges
  - Preparation of users
- Close cooperation with users
  - Documenation / step-by-step manuals
  - Many online courses

→ Only few support tickets after go live

HPC & AI

Efficient Coupling of Highly Parallel Computational Fluid Dynamics Simulations on Heterogeneous Architectures

Efficient Coupling of Highly Parallel Computational Fluid Dynamics Simulations on Heterogeneous Architectures

Motivation

- Convergence HPC & AI: Deep Learning (DL) as important tool to complement traditional numerical simulations
- Coupling of HPC codes and DL model is challenging
- Two use cases of deploying DL models into HPC simulation codes from the field of simulating turbulent reactive flows
  1. Modeling chemical reaction kinetics [1]

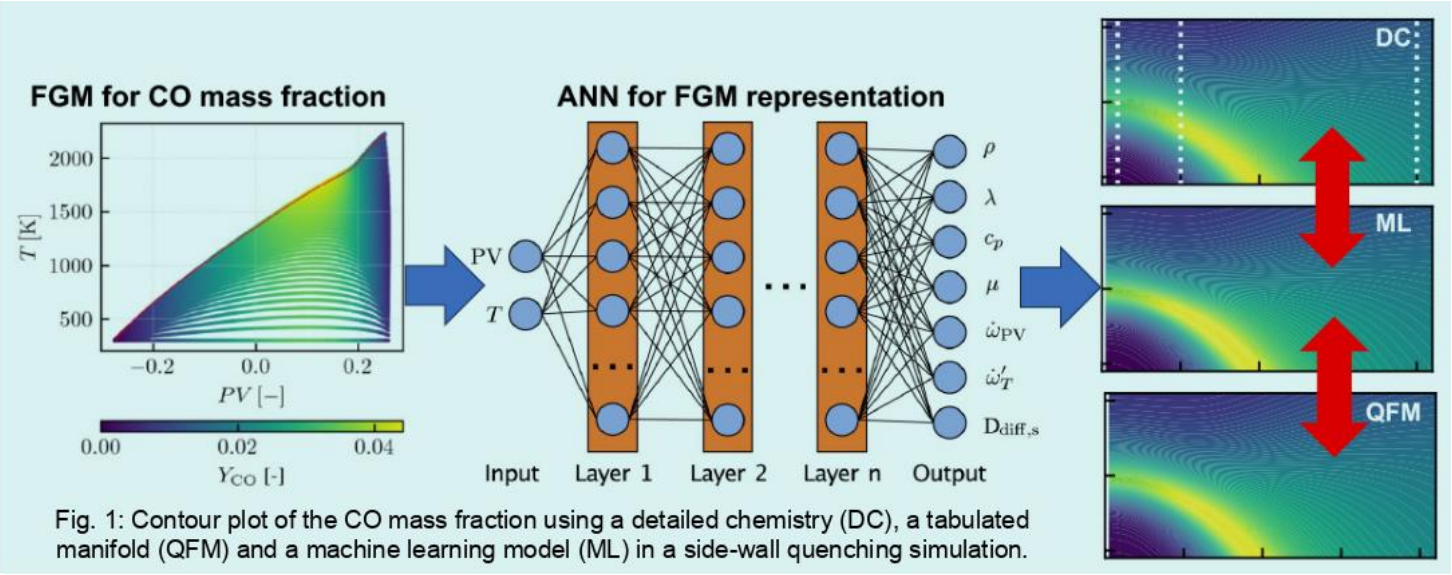

Fig. 1: Contour plot of the CO mass fraction using a detailed chemistry (DC), a tabulated manifold (QFM) and a machine learning model (ML) in a side-wall quenching simulation.

[1] J. Bissantz, J. Karpowski, M. Steinhausen, Y. Luo, F. Ferraro, A. Scholtissek, C. Hass, L. Vervisch, „Application of dense neural networks for manifold-based modeling of flame-wall interactions“, Applications in Energy and Combustion Science, Volume 13, March 2023

Efficient Coupling of Highly Parallel Computational Fluid Dynamics Simulations on Heterogeneous Architectures

Motivation

- Convergence HPC & AI: Deep Learning (DL) as important tool to complement traditional numerical simulations
- Coupling of HPC codes and DL model is challenging
- Two use cases of deploying DL models into HPC simulation codes from the field of simulating turbulent reactive flows
  2. Data-driven super-resolution turbulence modeling [2]

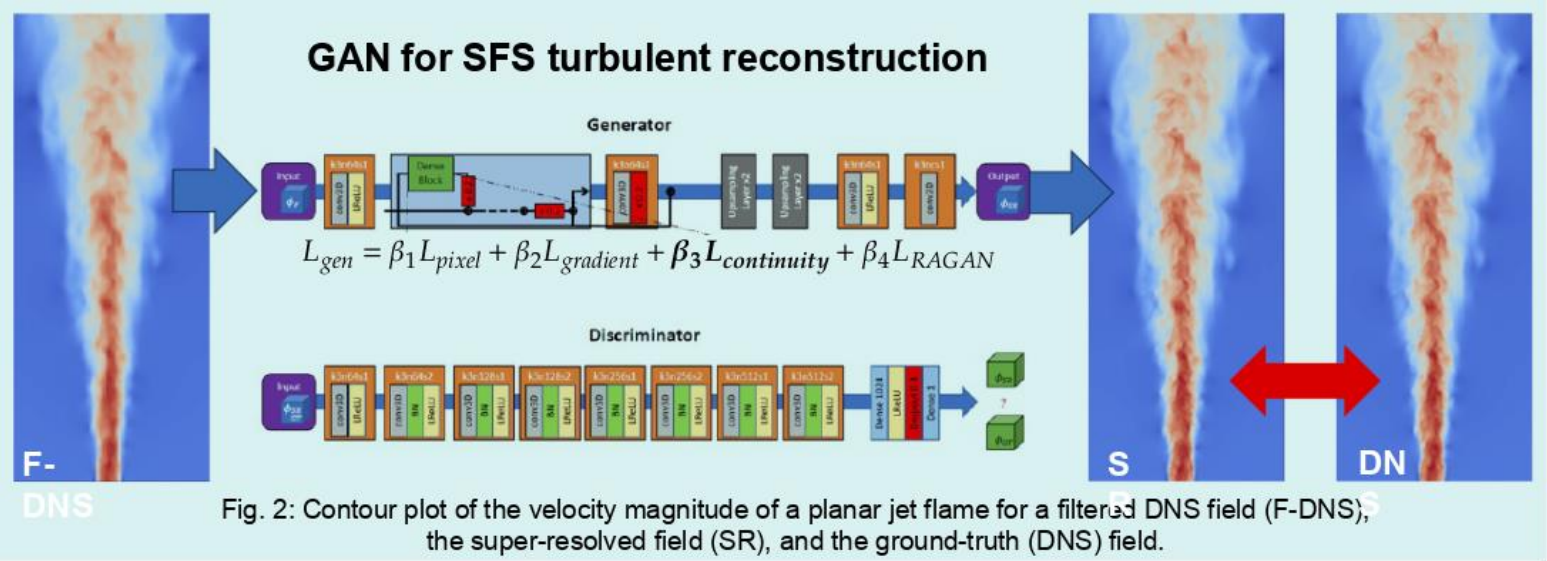

Fig. 2: Contour plot of the velocity magnitude of a planar jet flame for a filtered DNS field (F-DNS), the super-resolved field (SR), and the ground-truth (DNS) field.

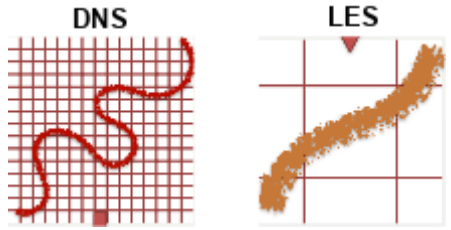

[2] L. Nista, C.D.K. Schumann, T. Grenga, A. Attili, H. Pitsch, „Investigation of the generalization capability of a GAN for LES“, in Proc. Combust. Inst., 2023

Efficient Coupling of Highly Parallel Computational Fluid Dynamics Simulations on Heterogeneous Architectures

Concept

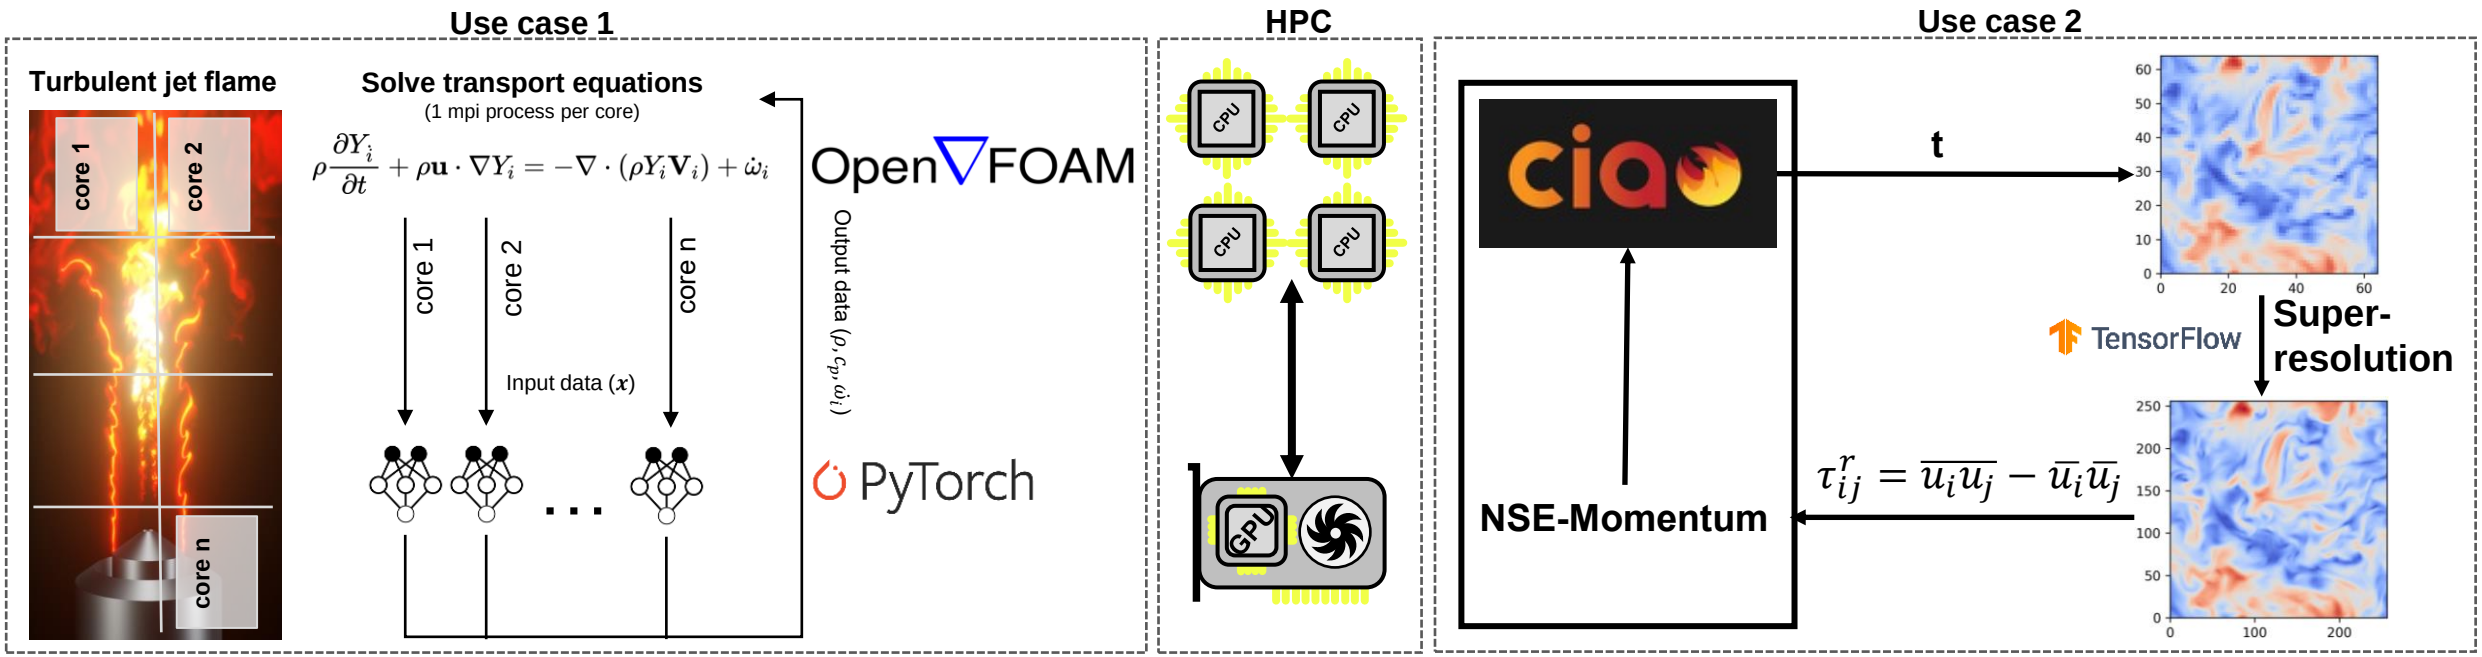

Efficient Coupling of Highly Parallel Computational Fluid Dynamics Simulations on Heterogeneous Architectures

AIxeleratorService Library [3, 4]

- Library Development: Ease the coupling of traditional HPC simulation code (e.g. CFD simulations) with DL models
- User-friendly interface abstraction DL frameworks (PyTorch, Tensorflow, SOL4VE)
- Hides data communication
- Hetergeneous Slurm Jobs

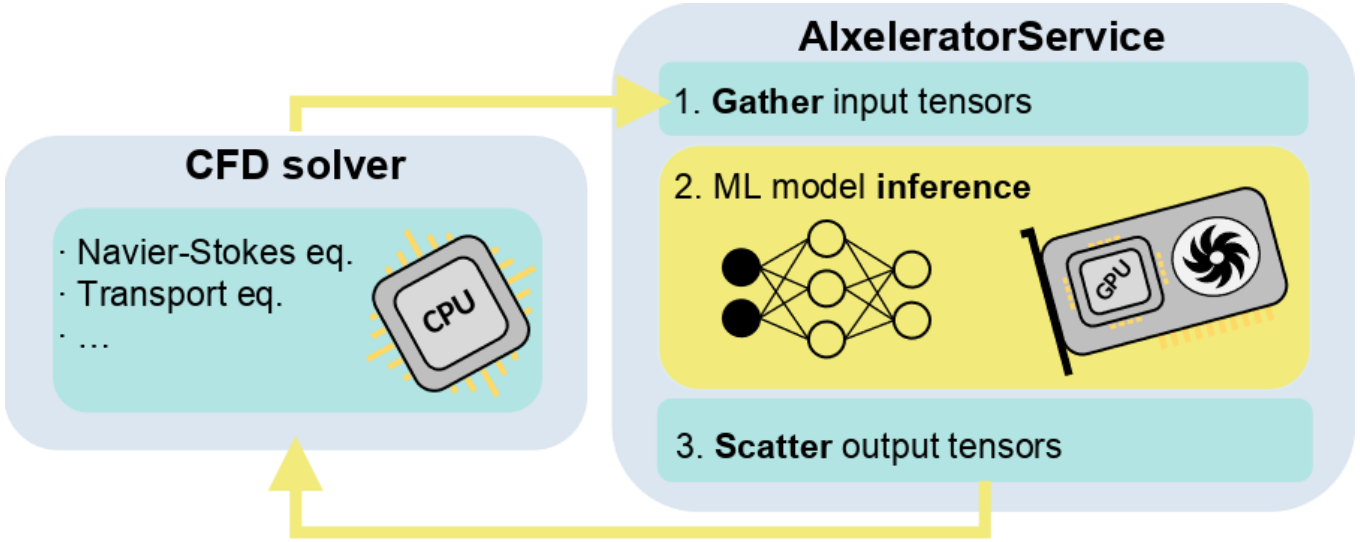

Fig. 3: Simulation loop of a CFD solver coupled with ML inference.

https://github.com/RWTH-HPC/AIxeleratorService

[3] F. Orland, K. S. Brose, J. Bissantz, F. Ferraro, C. Terboven, C. Hasse, „A Case Study on Coupling OpenFOAM with Different Machine Learning Frameworks“, Proceedings of AI4S: Artificial Intelligence and Machine Learning for Scientific Applications, 2022

[4] F. Orland, N. Kocher, L. Nista, H. Pitsch, C. Terboven, „AIxelerating Deep Learning Inference on Heterogeneous Architectures in Turbulent Flow Simulations“, unpublished

Efficient Coupling of Highly Parallel Computational Fluid Dynamics Simulations on Heterogeneous Architectures

Results Use Case 1 (OpenFOAM): Inference

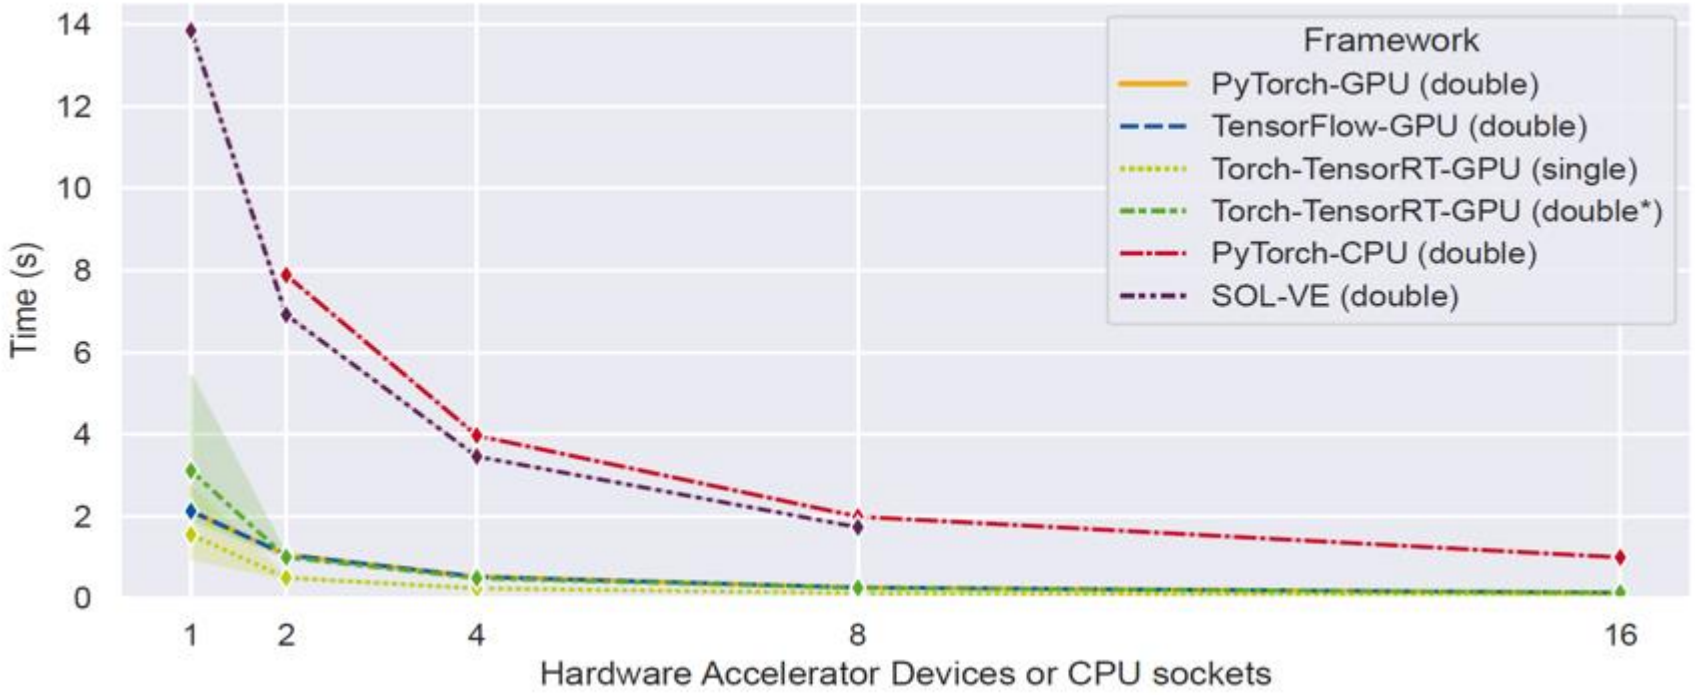

Executed on CLAIX-2018

- CPU: 48 Skylake Platinum 8160 cores @ 2.1 GHz
- GPU: Nvidia V100
- VE: Type 10B (8 cores @ 1.4 GHz)

Thanks to Fabian Orland [3], PhD student @ RWTH

One-Cluster-Concept in CLAIX

- All segements are integrated in one cluster
- Support for classical HPC + new maschine learning
- Different storage classes for specific needs, including research data managment
- Performance Monitoring
  - User perspective
  - Holistic view of the data center including cooling, power supply and job load
- Security concept includes 2FA for the cluster

Thank you for your attention.