CERN openlab and HPC activities

Maria Girone, CERN
Head of CERN openlab

CERN openlab

Outline

01

Introduction: CERN and HPC activities
R&D collaborations with science and industry

02

CoE RAISE

RAISE Center of Excellence

03

interTwin

interTwin

04

Relationship with **EuroHPC** and access to resources

EuroHPC Joint Undertaking

05

Spectrum

01 Introduction: CERN and HPC activities

R&D collaborations with science and industry

Motivation

New Techniques and Technologies

Enables Machine Learning and AI algorithms and processing techniques

Opens the possibility for real-time interactive simulations (Digital Twins)

A path to optimize energy usage

New Challenges and Opportunities

New resources for processing

Encourages stronger engagement with industry, other science communities and the computing community

Requires technology migration and redesigning of applications

Requires strategic planning and communication with the existing distributed computing community

Collaborations and R&D on HPC

New Techniques and Technologies

Enables Machine Learning and AI algorithms and processing techniques

Opens the possibility for real-time interactive simulations (Digital Twins)

A path to optimize energy usage

ODISSEE

Online Data Intensive Solutions for Science in the Exabyte Era

New Challenges and Opportunities

New resources for processing

Requires technology migration and redesigning of applications

Encourages stronger engagement with industry, other science communities and the computing community

Requires strategic planning and communication with the existing distributed computing community

CERN openlab

Since its inception, CERN openlab has fostered the development of big data scientific research through **four primary missions.**

1 | Establishing strategic industry collaborations

2 | Fuelling technological innovation

3 | Exposing technology to scientists

4 | Nurturing knowledge and growth in young STEM researchers

Stakeholders

CERN openlab's primary role is to act as conduit and facilitator for collaboration in computing science and technology between two categories of stakeholders:

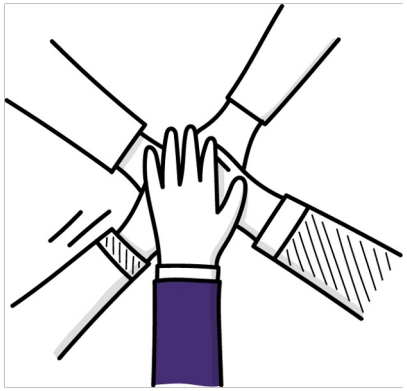

The science communities
(CERN departments and groups; R&D teams at CERN; Research centres)

Technology providers
(Industry)

Examples of Projects in HPC Areas

intel Unified programming models and portability

altera FPGA Acceleration

Micron Expandable Memory and AI on edge devices

E4 COMPUTER ENGINEERING Accelerators and RISC-V

NVIDIA Accelerators, digital twins and AI

SIMONS FOUNDATION AI/HPC Convergence, AI tuning

cerabyte Advanced Archival Storage & Digital Preservation

PURESTORAGE High Performance Scalable Storage

CORNELIS NETWORKS Low latency networking

AMD XILINX FPGA Acceleration


02
CoE RAISE

CoE RAISE

Center of Excellence for Research on AI- and Simulation-based Engineering at Exascale

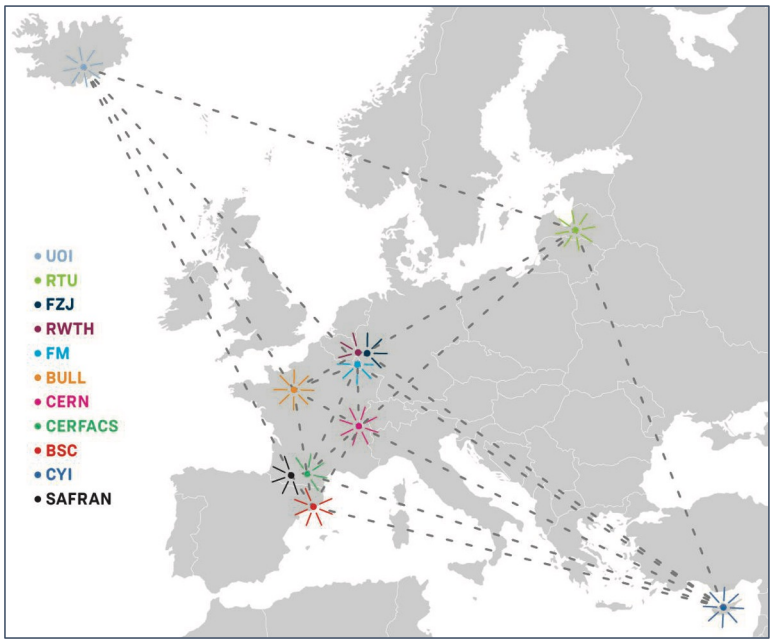

CoE RAISE Partners

https://www.coe-raise.eu/

Develop novel, scalable Artificial Intelligence technologies

Connect

- Hardware infrastructure
- Software infrastructure
- Computer-driven use-cases
- Data-driven use-cases

WP4: Data-Driven Use-Cases towards Exascale
Leader: Dr. Maria Girone

Task 4.1.: Event reconstruction and classification at the CERN HL-LHC → more details later
Leader: Eric Wulff

WP2: AI- and HPC-Cross Methods at Exascale
Leader: Prof. Morris Riedel

Provides expert support on HPC and AI methods to use-cases in WP4

CoE RAISE

Modularity of next generation HPC systems

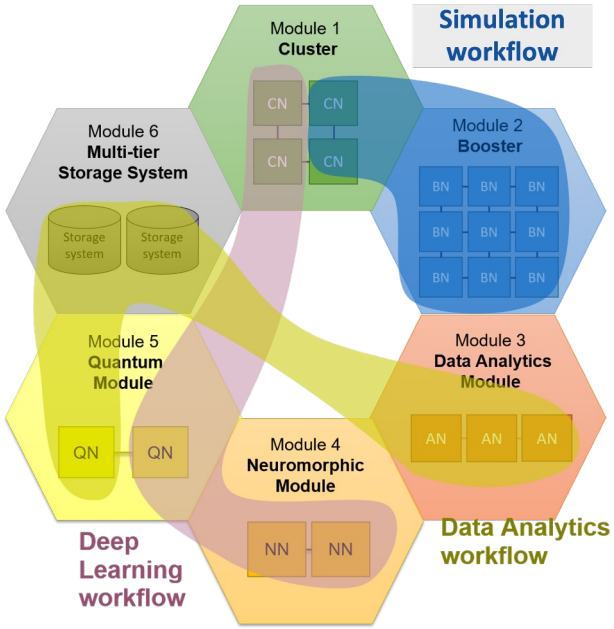

Complex Hardware

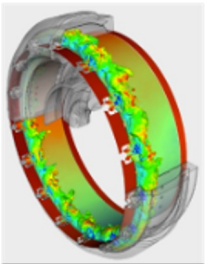

Complex Task

- Find the most suitable hardware for a specific task
- Enable intertwined AI- and HPC workflows
- EuroHPC's roadmap includes integrating Quantum Computers and Quantum simulators in already existing supercomputer centers [1]

[1] https://eurohpc-ju.europa.eu/selection-six-sites-host-first-european-quantum-computers-2022-10-04_en

AI-based particle flow reconstruction workflow

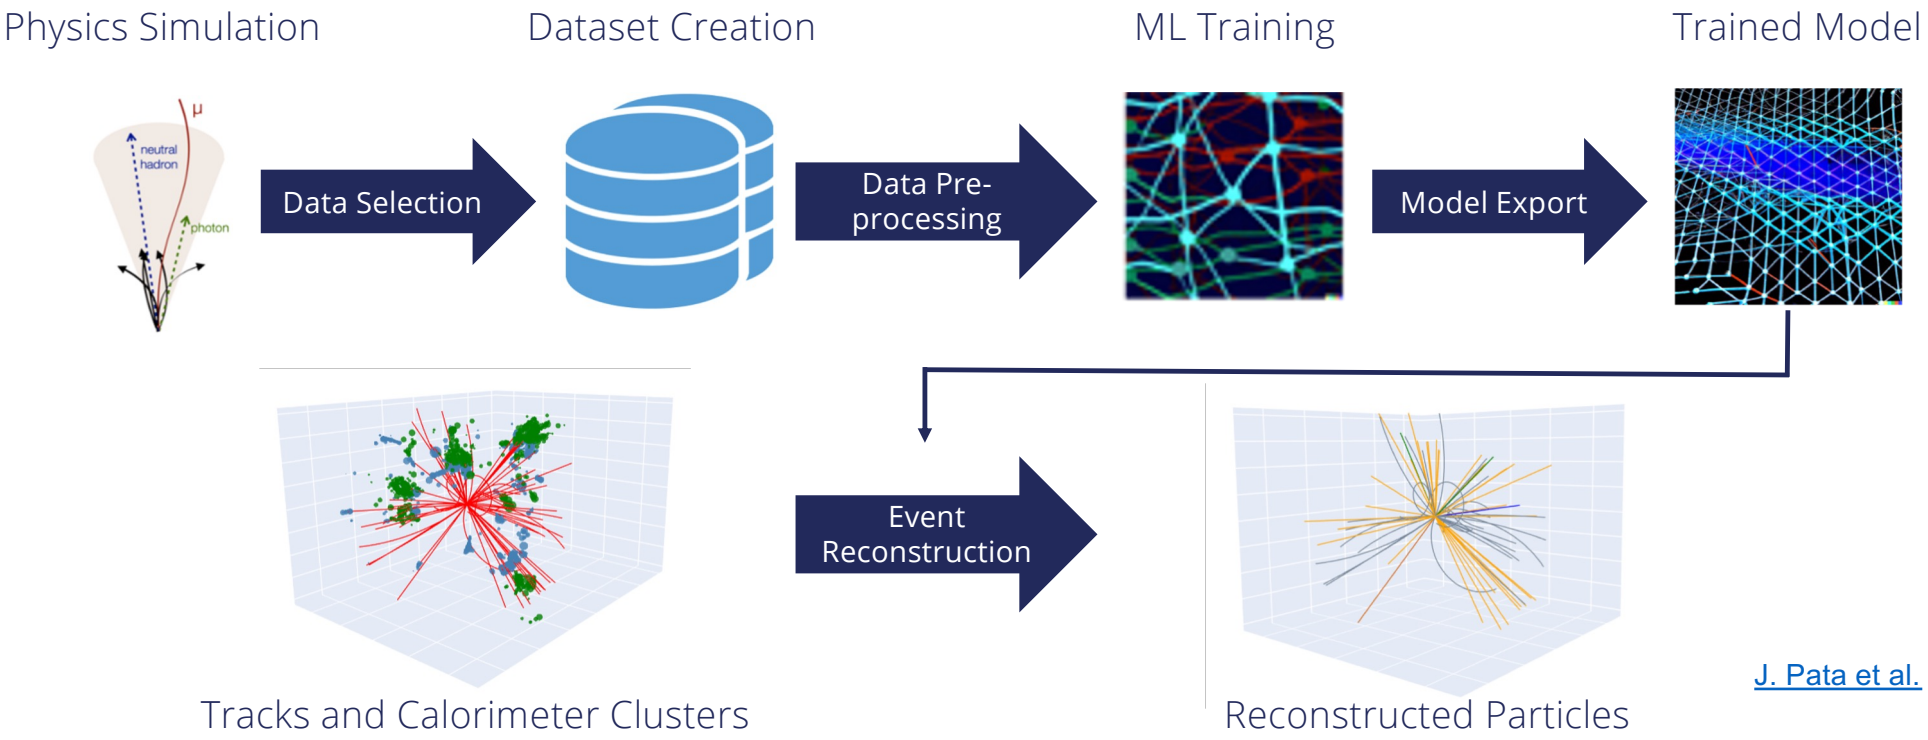

J. Pata et al.

Large-scale distributed hyperparameter optimization

Distributed HPO

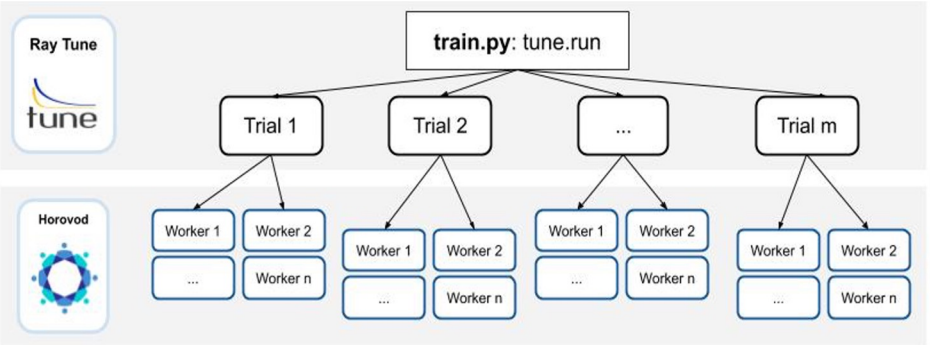

- Using ASHA + Bayesian Optimisation
- Scalable up to hundreds of GPUs
- Final validation loss decreased by giving a significant performance improvement from HPO

Hyperparameter tuning results

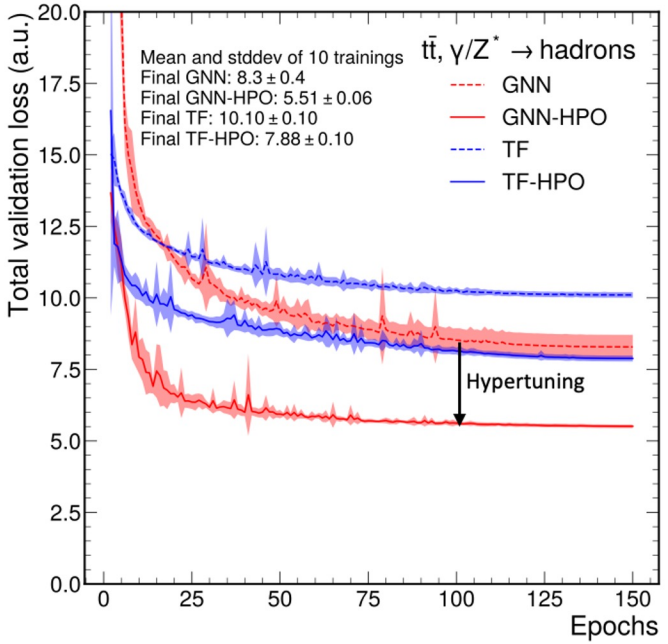

Pata, J., Wulff, E., Mokhtar, F. *et al.* Improved particle-flow event reconstruction with scalable neural networks for current and future particle detectors. *Commun Phys* 7, 124 (2024). https://doi.org/10.1038/s42005-024-01599-5

Scaling of an HPO workflow on HPC

Scaling of a HPO run of MLPF on the JURECA-DC-GPU system at the the Jülich Supercomputer Centre (JSC), 4 NVIDIA A100 and 2× 64 cores AMD EPYC 7742 per node.

- **Superlinear scaling** due to re-loading of models when using fewer nodes
- Using the ASHA algorithm to schedule and terminate trials, in combination with Bayesian optimization

In collaboration with WP2 in CoE RAISE we have also shown scaling of multi-node distributed training of deep learning models on **up to 256 compute nodes (1024 GPUs)**.

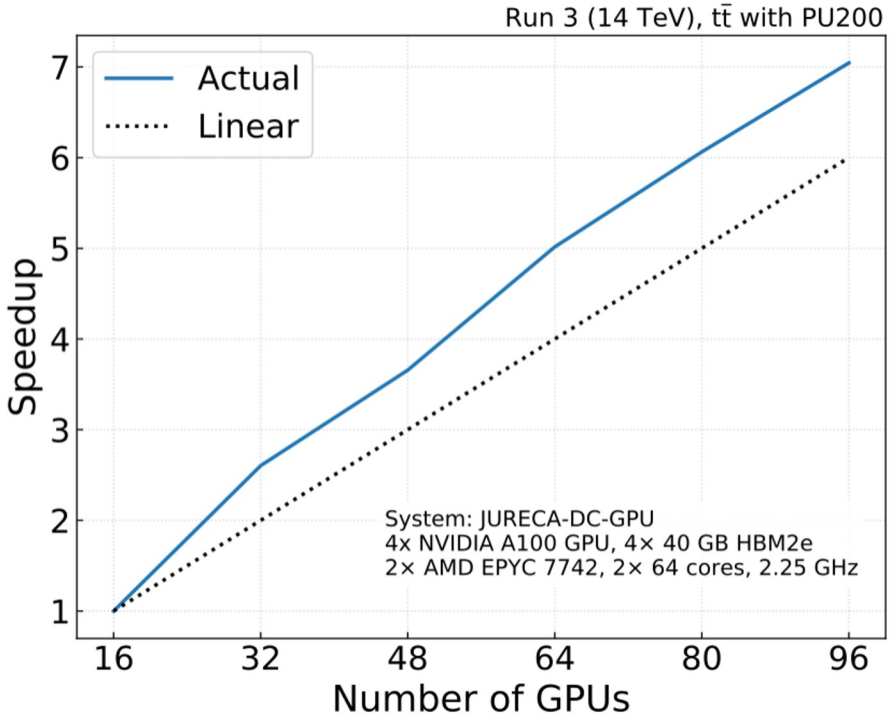

Data used: *Simulated particle-level events of ttbar and QCD with PU200 using Pythia8+Delphes3 for machine learned particle flow (MLPF),* https://doi.org/10.5281/zenodo.4559324

03 | interTwin

interTwin

A Digital Twin Engine for Science

- Co-design and implement the prototype of an **interdisciplinary Digital Twin Engine**
- **Open-source** platform based on **open standards**
- Large spectrum of **diverse use cases** from **physics** and **earth observation sciences**

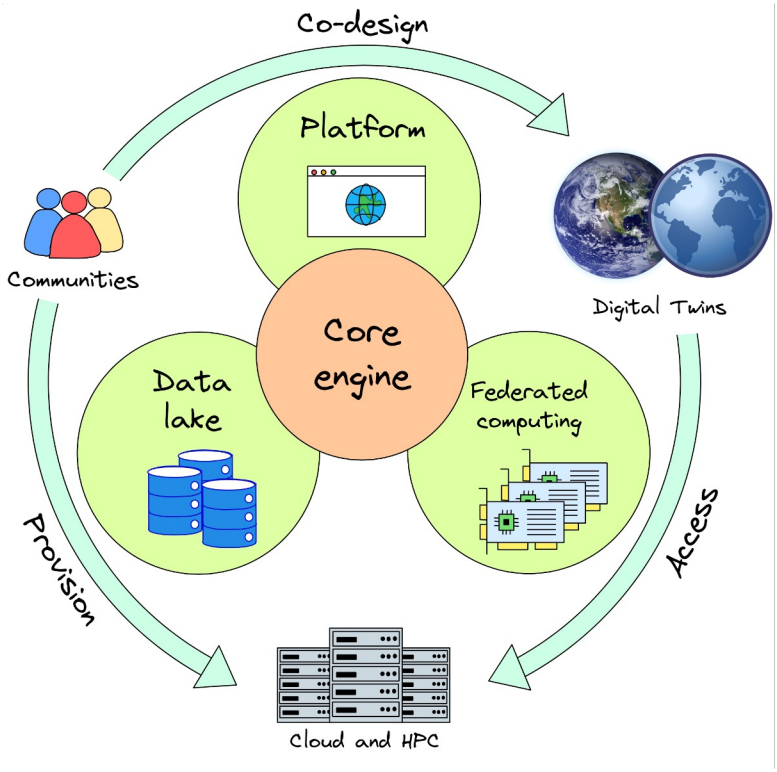

https://www.intertwin.eu/

interTwin

A Digital Twin Engine for Science

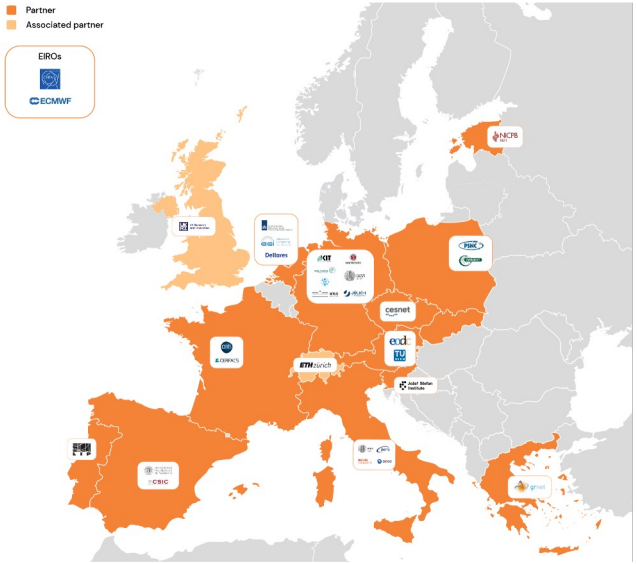

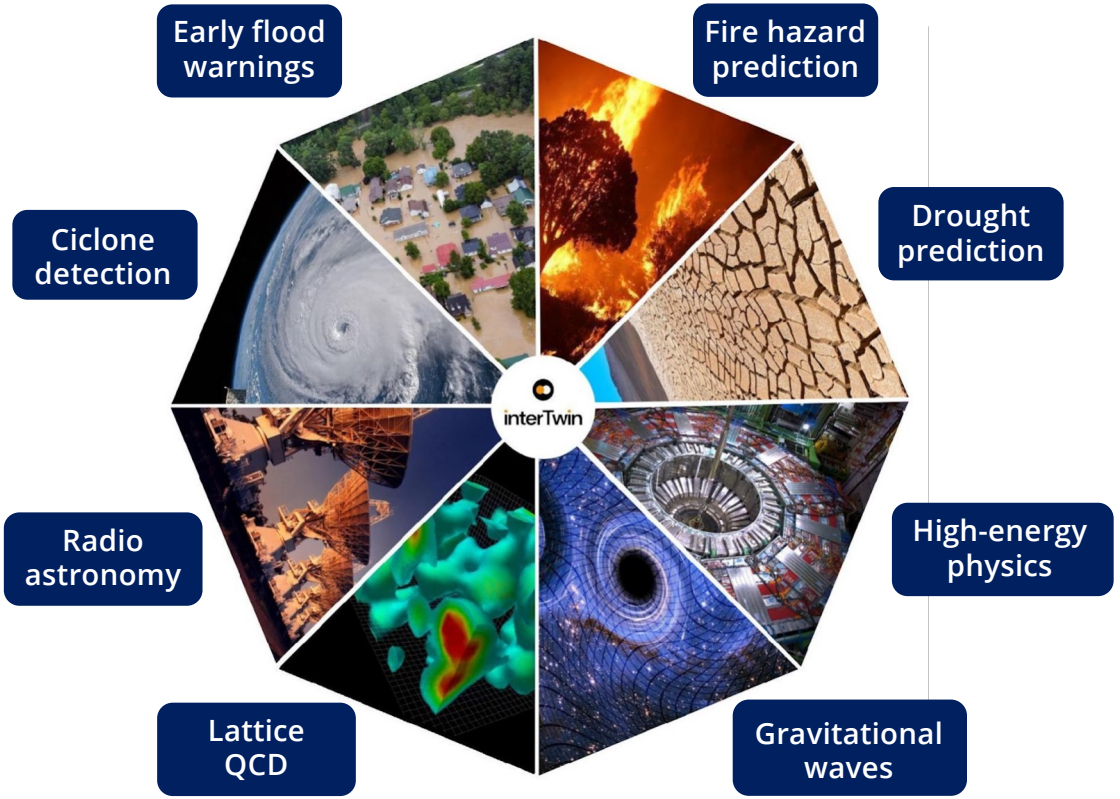

29 participants, including sciences, technology and resource providers.

Distributed Machine Learning training

Built on top of the expertise acquired in RAISE, developing a unified abstraction layer for distributed training with PyTorch, seamlessly integrating Torch Distributed (DDP), Microsoft Deepspeed, and Horovod.

Scaling test demonstrated on up to **512 GPUs** JSC.

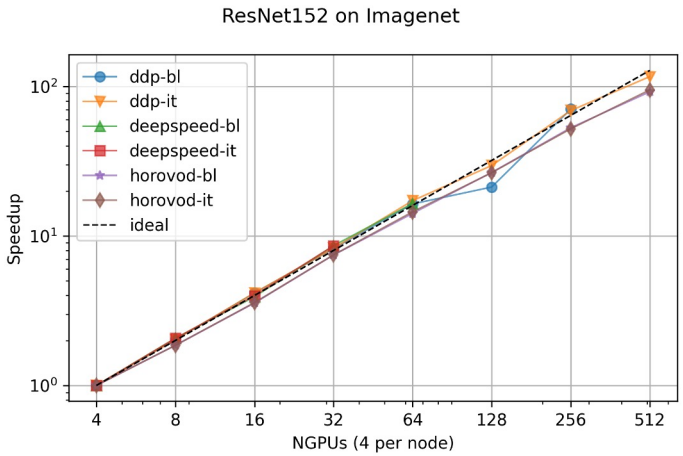

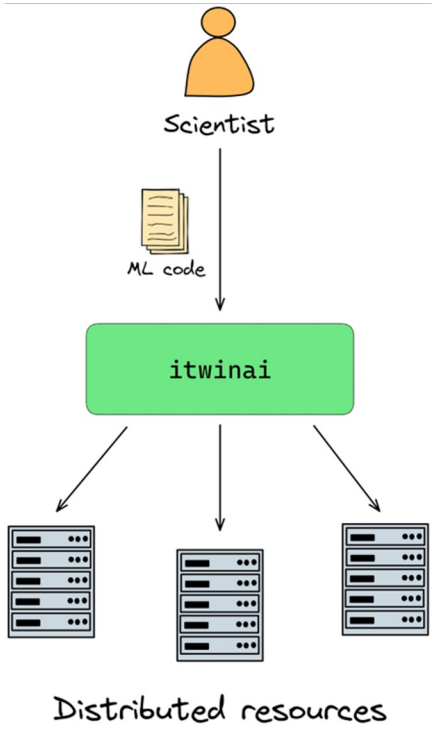

Machine learning on cloud and HPC

InterLink: leverages k8s' Virtual Kubelet and microservice architecture of the DTE

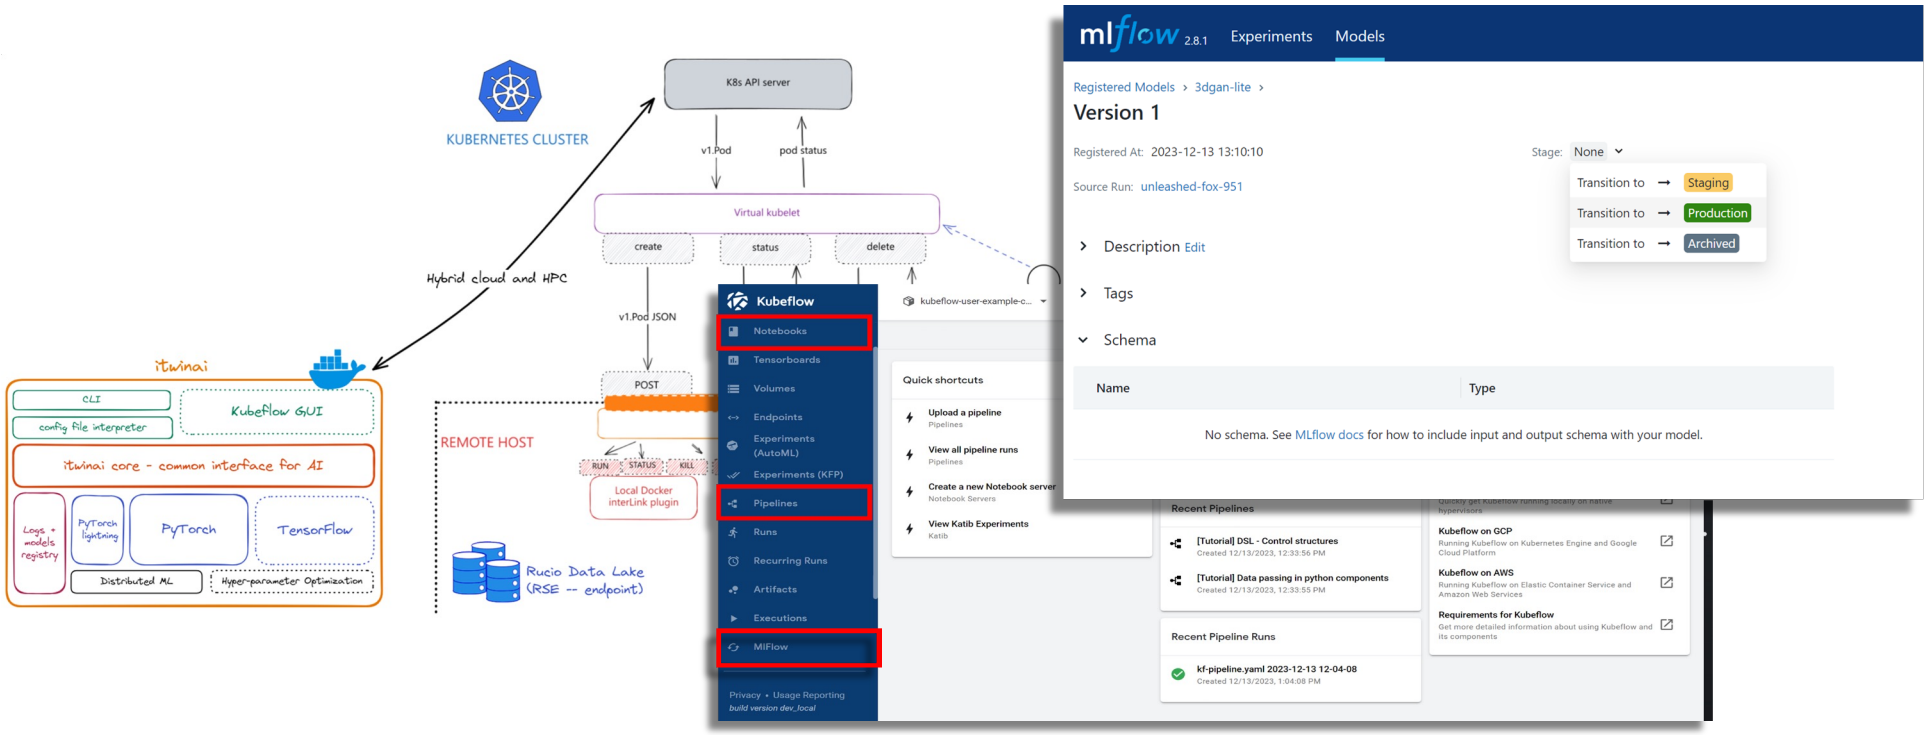

04 Relationship with **EuroHPC** and access to resources

EuroHPC

EuroHPC JU is the main organization enabling HPC access and development in EU

Budget EUR 7 billion for the period 2021-2027

Made up of EU, Member states, and several private partners (consortiums)

The Joint Undertaking manages the Union's access time for HPC (from 35% up to 50% of their total capacity)

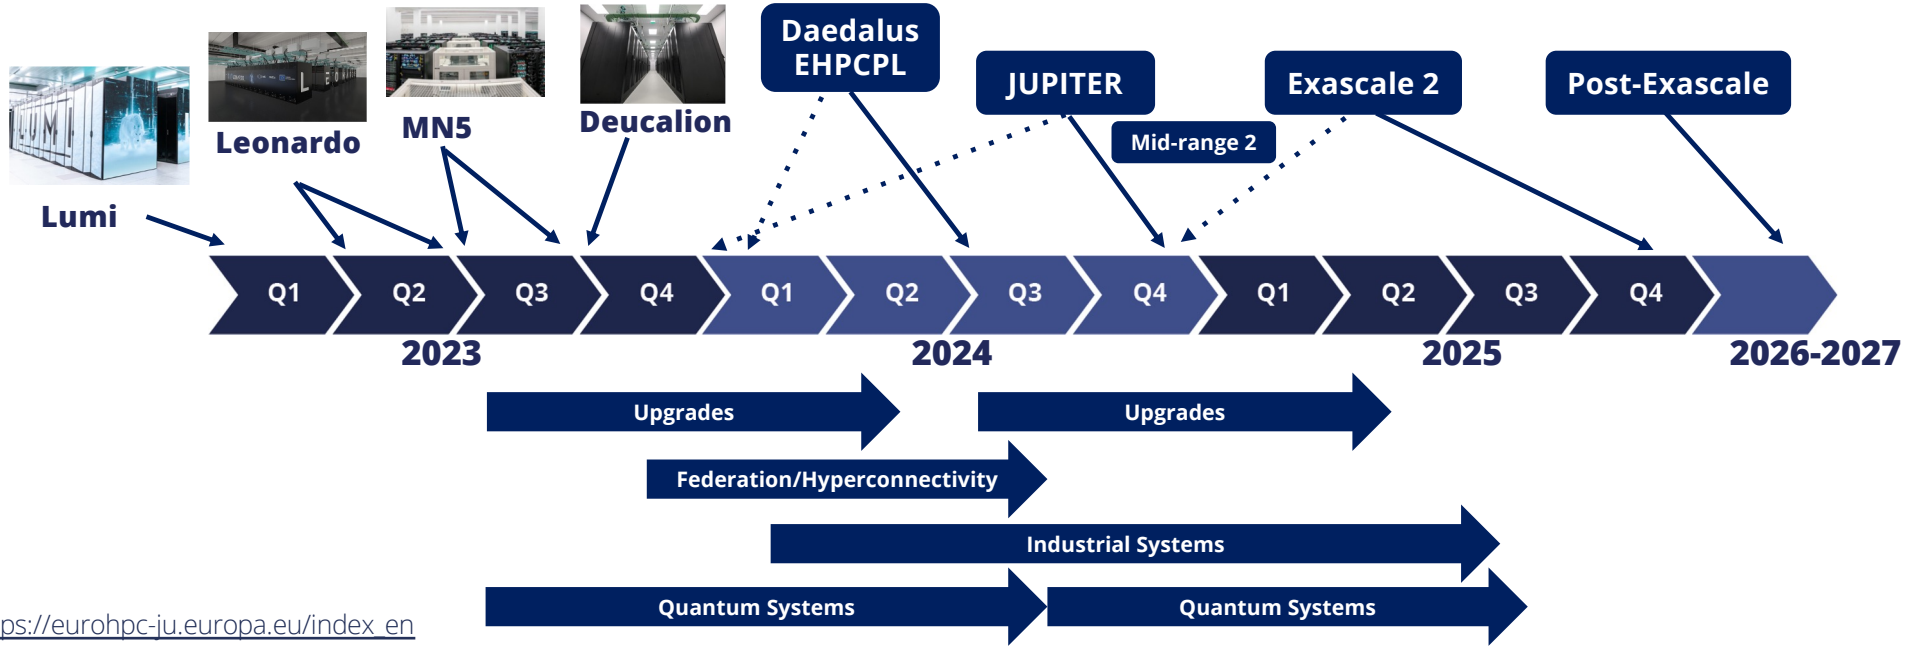

https://eurohpc-ju.europa.eu/index_en

EuroHPC

CERN participates in EuroHPC user related working groups, including requirements and is developing close relationships and a roadmap for seamless integration

EuroHPC is the primary portal for research access, via regular **Open Calls:**

- Benchmarking (3 months)
- Development (1 year, low capacity) **(extensively used at CERN by developers)**
- Regular (extended use)
- Extreme (high impact/use)

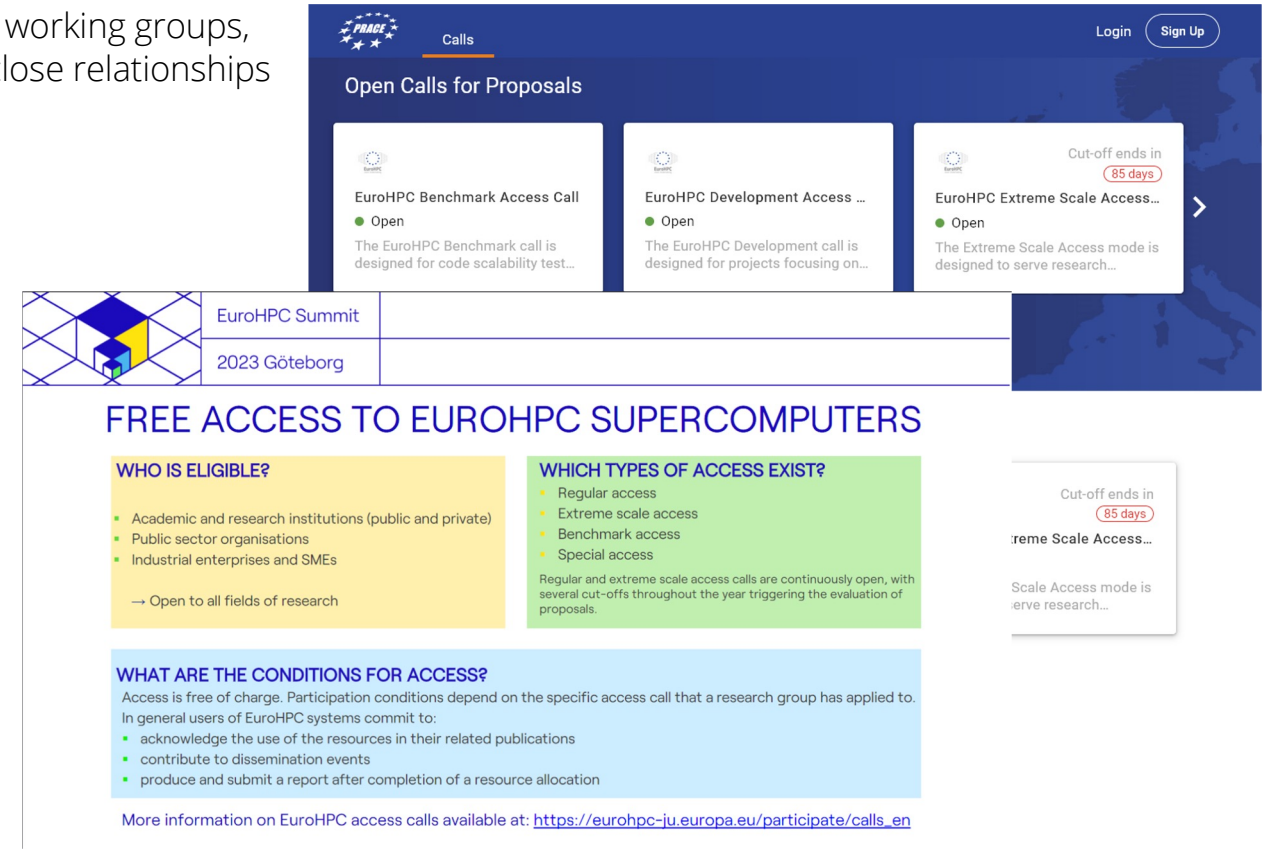

https://pracecalls.eu

EuroHPC
Quantum Computing Infrastructure (2024+)

- EuroHPC **currently exposes 9 HPC sites** for use
- **6 quantum computers** will be hosted by EuroHPC facilities
- Free access for Research, Development and Innovation
- Integration of QC with HPC and development of hybrid HPC-QC applications and workflows
- European Quantum Excellence Centres in quantum applications, for science and industry

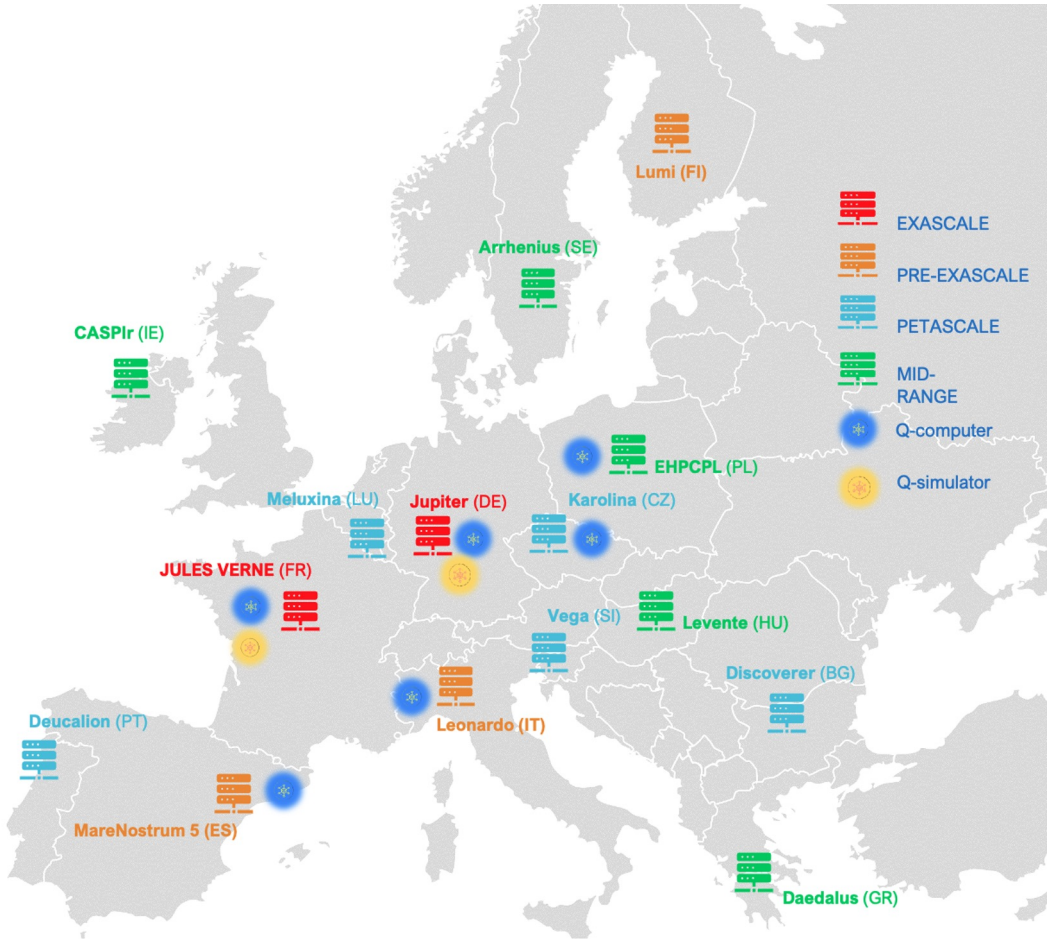

Image by Gustav Kalbe, Opening Plenary Session at EuroHPC Summit 2024.

05 Spectrum

What is SPECTRUM?

A project granted under the call HORIZON-INFRA-2023-DEV-01-05, which **aims to the preparation of a Computing Strategy for Data-intensive Science Infrastructures in Europe for the High Energy Physics (HEP) and Radio Astronomy (RA) domains.**

Who is SPECTRUM?

SPECTRUM **gathers selected stakeholders in the HEP and RA research domains, and at the same time experts from the e-Infrastructures** (HPCs, Clouds, Quantum Computing). The former group brings directions and future needs, the latter expectations for new e-Infrastructures about technical and policy aspects.

Coordinator

- EGI Foundation

Research Infrastructure representatives

- **LHC:** CERN, INFN
- **SKA:** CNRS/OCA
- **LOFAR:** NWO-I through ASTRON

e-Infrastructure representatives

- FZJ (HPC Exascale and quantum computing)
- CINECA (HPC & Quantum)
- SURF (HTC, HPC, Cloud)
- (also EGI Foundation, INFN)

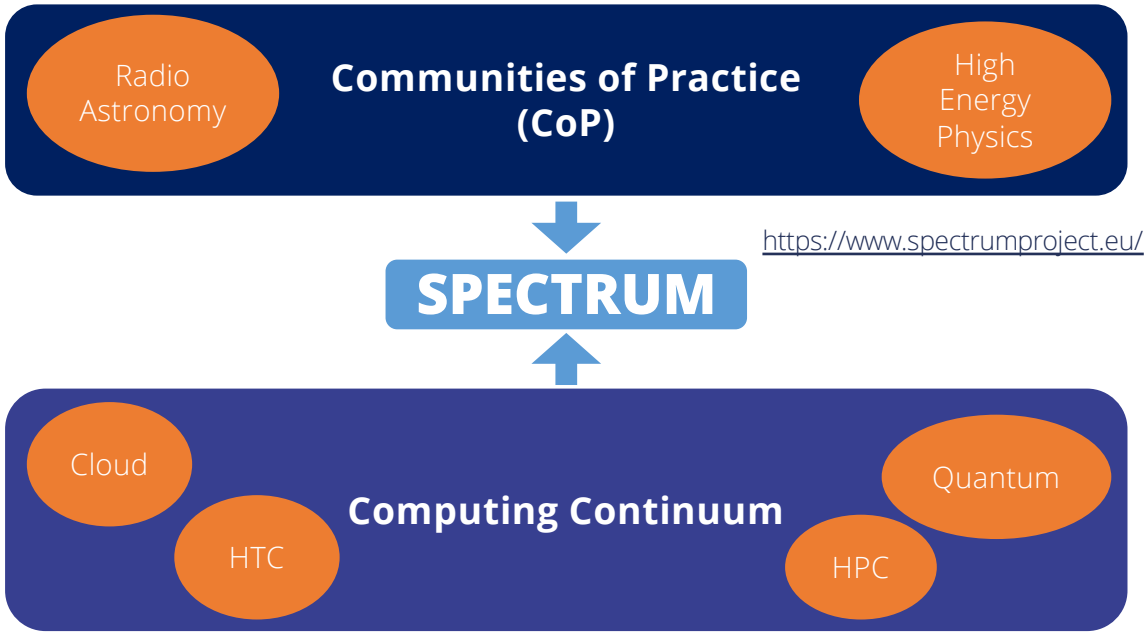

https://www.spectrumproject.eu/

What are the **expected outcomes** of SPECTRUM?

The realization of a **Community of Practice to gather and inform about future directions and needs in data-intensive research** on the one side, about future e-Infrastructures on the other.

A **Strategic Research, Innovation and Deployment Agenda (SRIDA) and a Technical Blueprint about agreed processing models and solutions**, to provide feedback on investment to funding agencies and policy makers.

Why is SPECTRUM different from previous attempts?

Previous interactions between the research and the e-Infrastructures communities have been a-posteriori, attempting to adapt scientific workflows to already operational facilities. This has been only partially successful due technical (non-compliant system architectures, …) and policy (user access, …) incompatibilities.

SPECTRUM **wants to move the handshaking process a-priori**, before e-Infrastructures are designed and deployed.

SPECTRUM: Plan of Work

6 Work Packages Identified

- **WP1:** Project Management and Communications (first period)
- **WP2:** Project Management and Communications (second period)
- **WP3:** Community of Practice Creation and First Consultation
- **WP4:** Community of Practice Second Consultation and Long-term Sustainability
- **WP5:** Landscape, use-cases, challenges and gaps
- **WP6:** Strategy Research, Innovation and Deployment Agenda

Project length: 30 months (Jan24 - Jun26)

- **First period:** January 2024 - March 2025
- **Second period:** April 2025 - June 2026

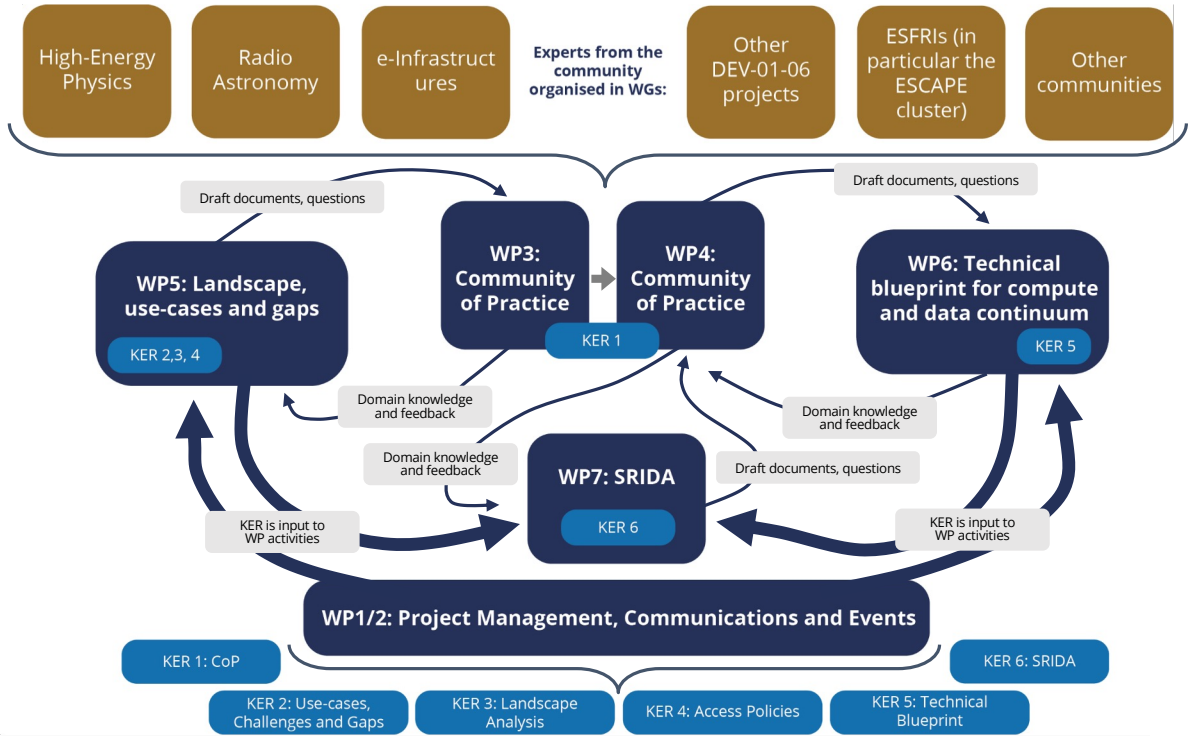

Relationship between WPs and and their related KERS

SPECTRUM CoP
Working Groups

WG1: Data Management and Access

Data Management - including data location,
movement and registration

Data Access protocols

Data Archiving

Security

WG2: Workflow Management and organisation

Resource Discovery and Workflows Submission

Resource Allocation

Complex Workflows

WG3: Compute Environment

Expected Tools and Services - Describes the tools and services expected at centers, including virtualization and CVMFS (CernVM File System).

Facility Expectations - Outlines expectations for facilities, including workflow execution times ("be able to execute workflows xx hours long), memory requirements ("nodes with at least YY GB RAM"), networking capabilities ("networking in/out at least at YY Mbps per core") and virtualization.

Edge Services - Discusses edge services needed to establish environment(s).

Library Provisioning - Discuss the minimal set of libraries and how to provide them, including virtualization, Spack, and module.

SPECTRUM CoP
Working Groups

WG4: SW tools

Machine Learning Frameworks

Multithreading Frameworks

Multi-Node Tools - Including MPI (Message Passing Interface) and similar tools

Compilers, toolchains, …

Quantum Computing Tools and Frameworks

Code Management Practices

WG5: Scientific Use-Cases

Typical Use-Cases - Describes typical use-cases with quantitative descriptions

Requirements and Needs

Best Practices Collection - Gathers best practices and existing documents

Data Fluxes and Paths

WG6: Facilities

HPC Centers - Discusses drivers and future directions of HPC centers

Access to Quantum Computing Hardware

Access to Commercial and Public Clouds

Sustainability - Examines economic and carbon footprint aspects

Security - Including access and translation from global SSO to local credentials

The SPECTRUM project and the JENA Computing Initiative are conducting a **survey** to gather insights on current best practice and expected future evolutions in the domain of large scale / data intensive scientific computing

Outlook

We have already begun to see the benefits from engaging with the HPC community

- Opportunistic access to resources, technical progress in AI/ML applications, and real-time interactive simulation

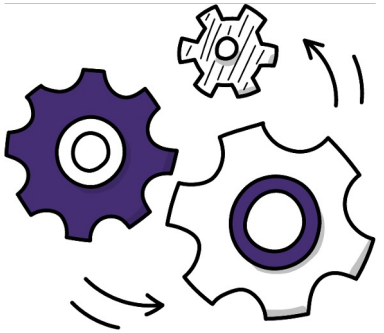

We are continuing with community engagement and strategic planning

- Activities in SPECTRUM, and soon ODISSEE, will bring the communities closer together and more seamlessly integrate the computing environments

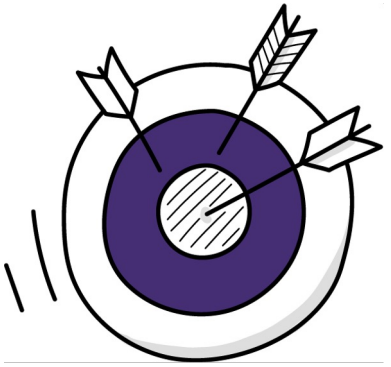

Computing is experiencing a rapid evolution of technology and techniques with large improvements in accelerators and the applications

- We are trying to stay on the forefront while stably supporting programmes that run for decades

Thank you

Maria Girone (CERN)
Head of CERN openlab

CERN openlab

Get in touch!

Send us a message

Email | openlab-communications@cern.ch

Website | openlab.cern

Follow us!

Phase VIII Brochure

CERN openlab LinkedIn

Thank you!

Maria Girone
Head of CERN openlab

CERN openlab

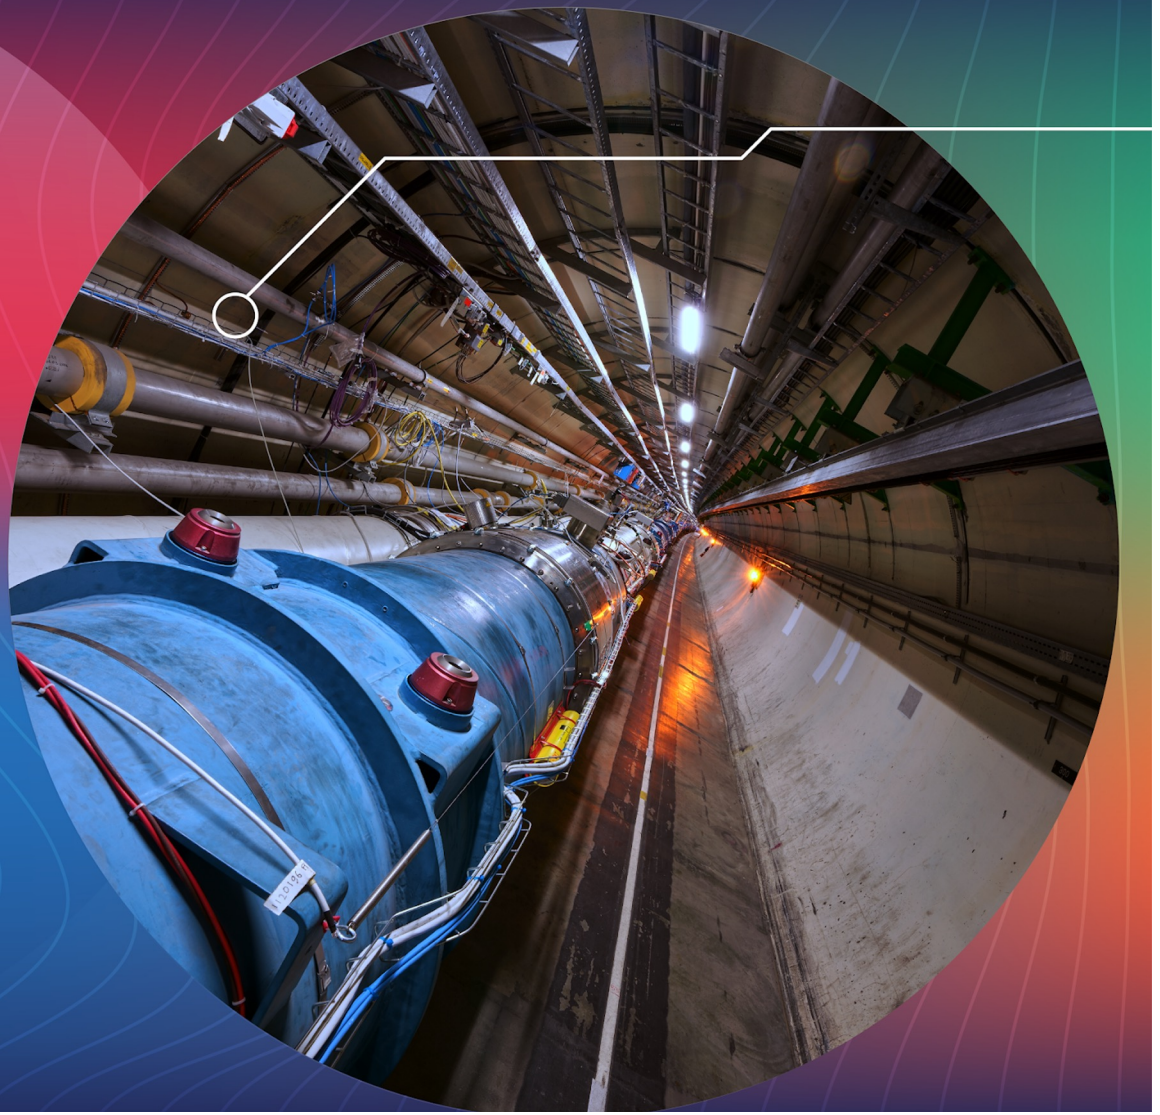

interTwin

A Digital Twin Engine for Science

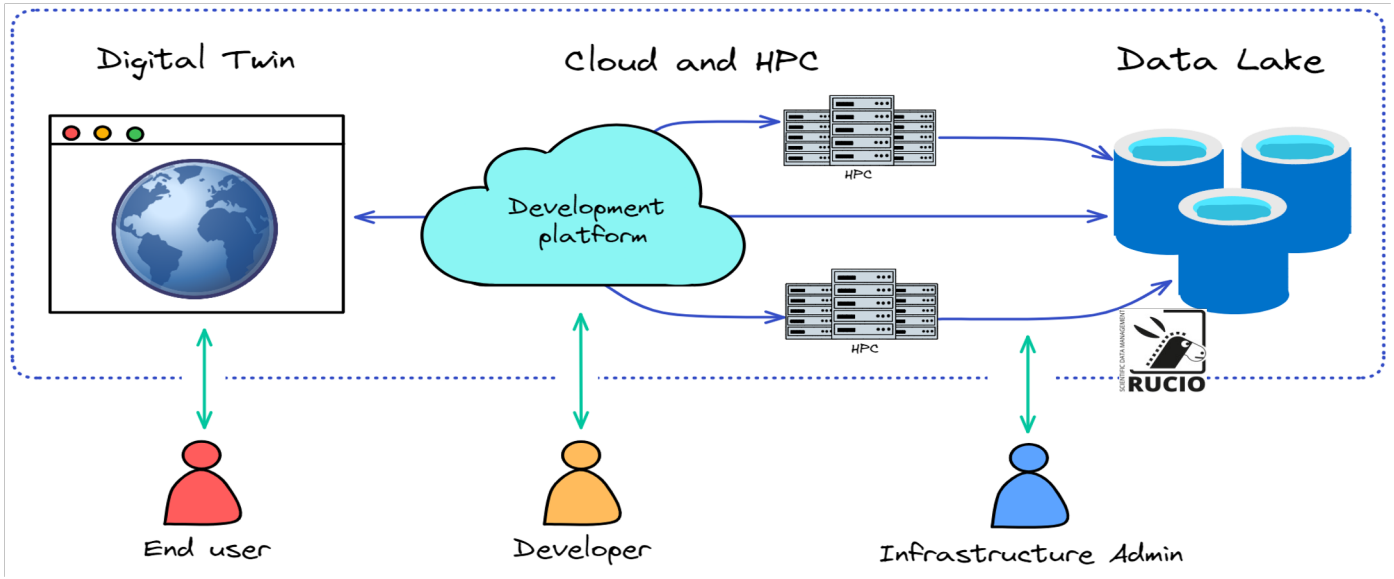

SPECTRUM:
Goals (up to now)

Specific scientific (sub)domains working as silos:

- Very intense R&D, well performing computing models, but very few cross domain communications
  - Have HEP and RA share experience, tools, solutions, visions
- Roadmap for next 5-10 years often available, but not compared with other subdomains needs
  - Conflicting models?

Many attempts to extend the computing to large infrastructures (HPC, Clouds, ..)

- But mostly a-posteriori: try to use an already deployed system, facing with limitations (technical and political)
  - No participation to the planning and deployment phases; no participation to the definition of target use cases