




**CSCS**

Centro Svizzero di Calcolo Scientifico  
Swiss National Supercomputing Centre

**ETH** zürich



# The Swiss National Supercomputing Centre (CSCS)

Joost VandeVondele

Deputy Director, The Swiss National Supercomputing Centre (CSCS), ETH Zurich

# CSCS: the ALPS research infrastructure (RI)




CSCS develops and operates a high-performance computing and data research infrastructure that supports world-class science in Switzerland.

# A RI connected to experiment, computational science, and the world

PSI




MCH




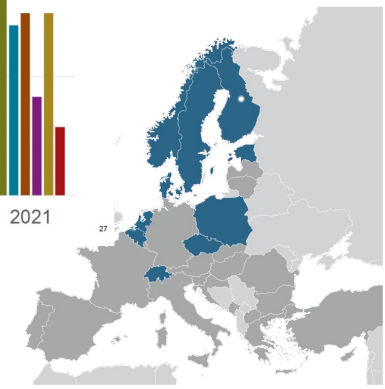

SKA / CTA



CERN




CSCS



# User Program & PASC, and Partnerships

The **User Lab (User Program & PASC)** provides access to resources and knowhow based on a *peer review process* and are funded through the HPCN initiative and ETHZ by the Swiss Government

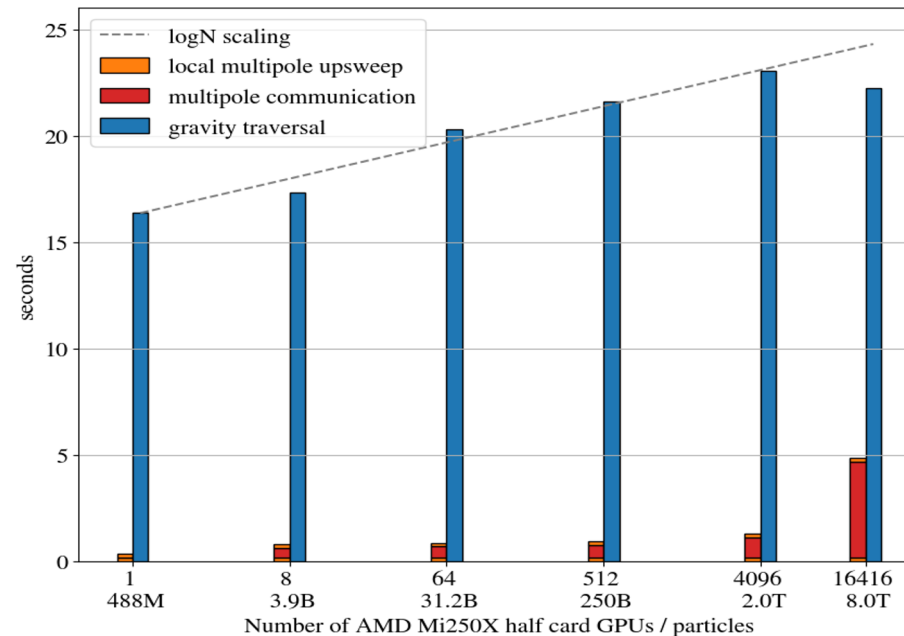


# Educating and moving communities along: The Platform for Advanced Scientific Computing (PASC)

Example: hydrodynamics and gravity  
for astrophysics and cosmology



- PASC project, 2nd round / 8 years
- Essentially code started from scratch, GPU acceleration/portability.
- Full Size runs on LUMI-G (16416 GPUs)
- Won Europe largest (22MGPUh) allocation on LUMI-G

- The broad availability and quality of HPC software developed in Switzerland
- Performance on CSCS infrastructure (GPU-accelerated supercomputing platforms)
- HPC software engineering competence in Swiss academic institutions
- Plays a crucial role in the co-design of CSCS' supercomputing infrastructure
- in its fifth cycle (HP2C: 2009, PASC: 2013, 2017, 2021, 2025)



SPH  
EXA

# CHIPP - Analysis of data from the Large Hadron Collider (LHC)



- Swiss particle physics community in the context of the Large Hadron Collider (LHC) at CERN. The goal of LHC is to help understand the building blocks of our Universe by particle collisions.
- On behalf of CHIPP, CSCS operates a mid-size Tier-2 grid site for three of the four detectors:
  - ATLAS, CMS and LHCb.
  - Currently ~10 PB of storage, ~250 compute nodes
- This grid site is the first in the world to be fully running with HPC resources and have front-end services to be fully kubernetesised
  - (Seamlessly) migrating into Alps

# The Swiss AI initiative



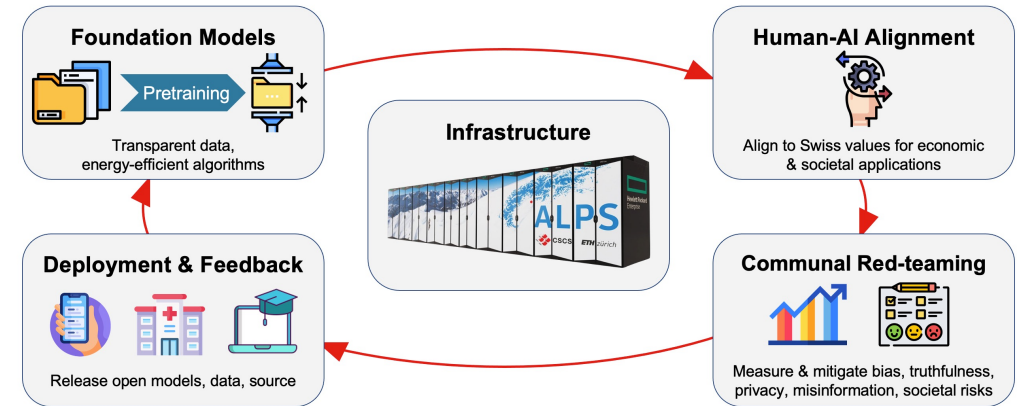
## Joint initiative for trustworthy AI

ETH Zurich and EPFL are launching the "Swiss AI Initiative", whose purpose is to position Switzerland as a leading global hub for the development and implementation of transparent and reliable artificial intelligence (AI). The new Alps supercomputer based at the Swiss National Supercomputing Centre (CSCS) provides the supporting world-class infrastructure.

04.12.2023

Share

## Guiding Principles: Trust, Openness, Collaboration

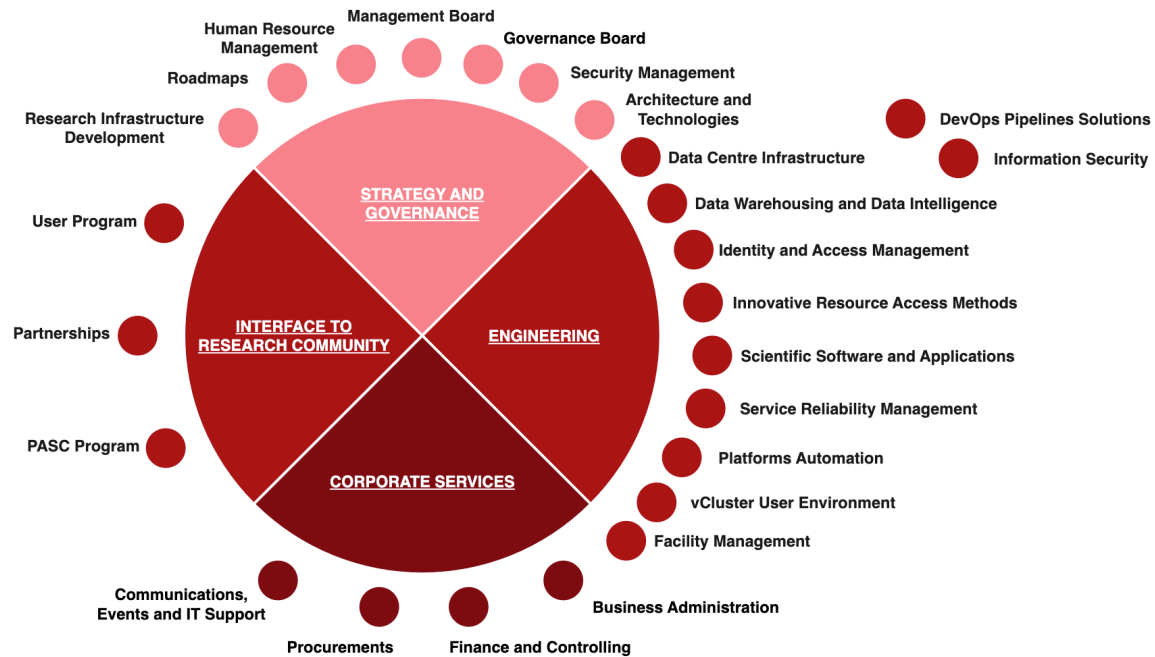


Develop capabilities, know-how & talent to build trustworthy Generative AI aligned with Swiss values

Make these resources available for the benefit of Swiss society and global actors



# CSCS is a world-class agile engineering organization



Introduce CSCS people in this meeting today (Thomas, Pablo, Nur, Mauro, Maxime, Kasia, Maria Grazia, Prashanth), more staff will join in person or online for the panel discussion.



130 Employees, 30 nationalities



# Conclusions

- CSCS mission is to enable world-class science in Switzerland
- With Alps, CSCS has deployed a world-class AI capable supercomputer that offer cloud-native functionality
- Infrastructure is key to modern AI and foundational models, and the Swiss AI initiative unites and will generate top talent in AI
- People are key, employing, educating and enabling.

# Alps Research Infrastructure

Thomas C. Schulthess



## OUR MISSION (since 2015)

“We develop and operate a high-performance computing and data research infrastructure that supports world-class science in Switzerland.”

**The research infrastructure (User Laboratory)  
is open to scientists worldwide**

# The issues with digitalisation and data in science

**BIG DATA ANALYTICS**

**FAIR DATA**

**EDI**

**GRID COMPUTING**



**DATA LAKES**

**OPEN RESEARCH DATA**

**EOSC: CLOUD OR COMMONS?**


**EOSC-PORTAL**

**EOSC-NODES**

Important considerations when dealing with digital data:  
(e.g. ask what ChatGPT says about Big Data)


- 1. Velocity
- 2. Volume
- 3. Veracity
- 4. Variety
- 5. Value

But we don't know how big — prepare the infrastructure for/with most ambitious/experienced



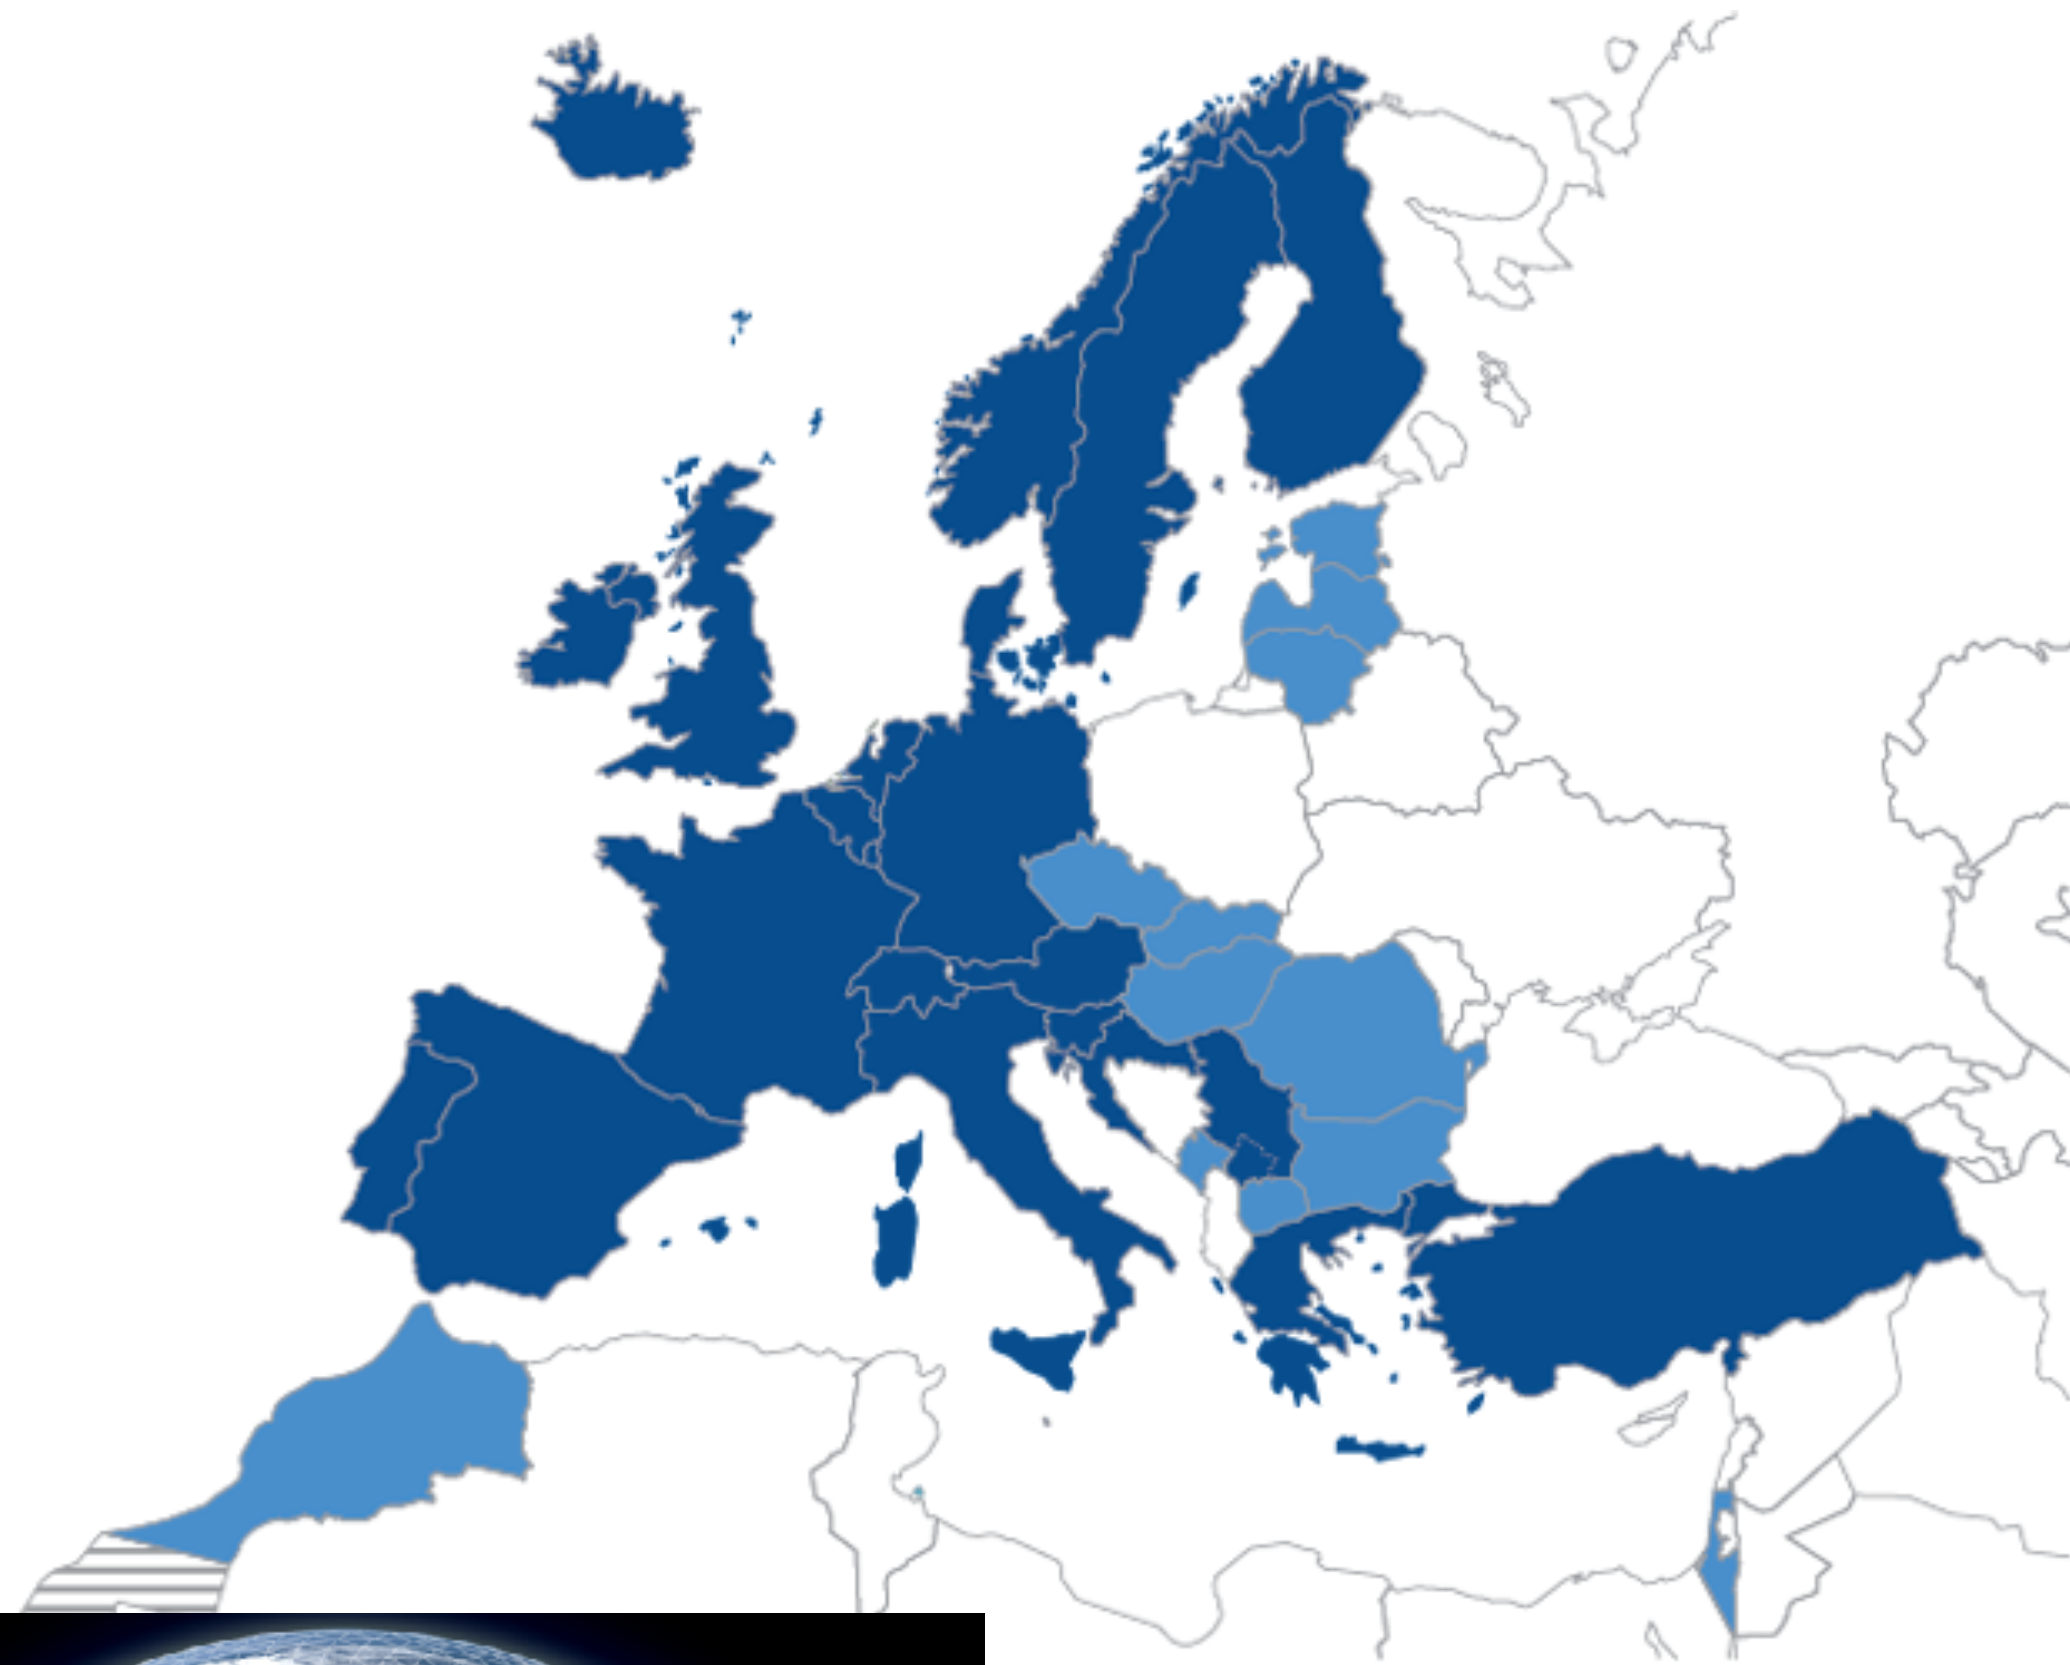
# MeteoSwiss

Federal Department of Home Affairs FDHA  
Federal Office of Meteorology and Climatology MeteoSwiss





# European Center for Medium-Range Weather Forecasts




## ECMWF

An independent intergovernmental organization established in 1975

Switzerland was founding member of ECMWF among 18 countries

Today the worldwide leading numerical weather prediction center

Provides input data for the weather predictions of MeteoSwiss






**CSCS**

Centro Svizzero di Calcolo Scientifico  
Swiss National Supercomputing Centre

**ETH** zürich



M E N U



USER PORTAL

← 1 / 3 →

## World's Most Powerful AI-Capable Supercomputer




**CSCS, Hewlett Packard Enterprise and NVIDIA Announce World's Most...**

12.04.2021


"Alps" system to advance research across climate, physics, life sciences with 7x more powerful AI capabilities than...

M O R E

M O R E S C I E N C E







Power: ~8kW

# Alps: 2,688 nodes or 10,752 GH200 Superchips

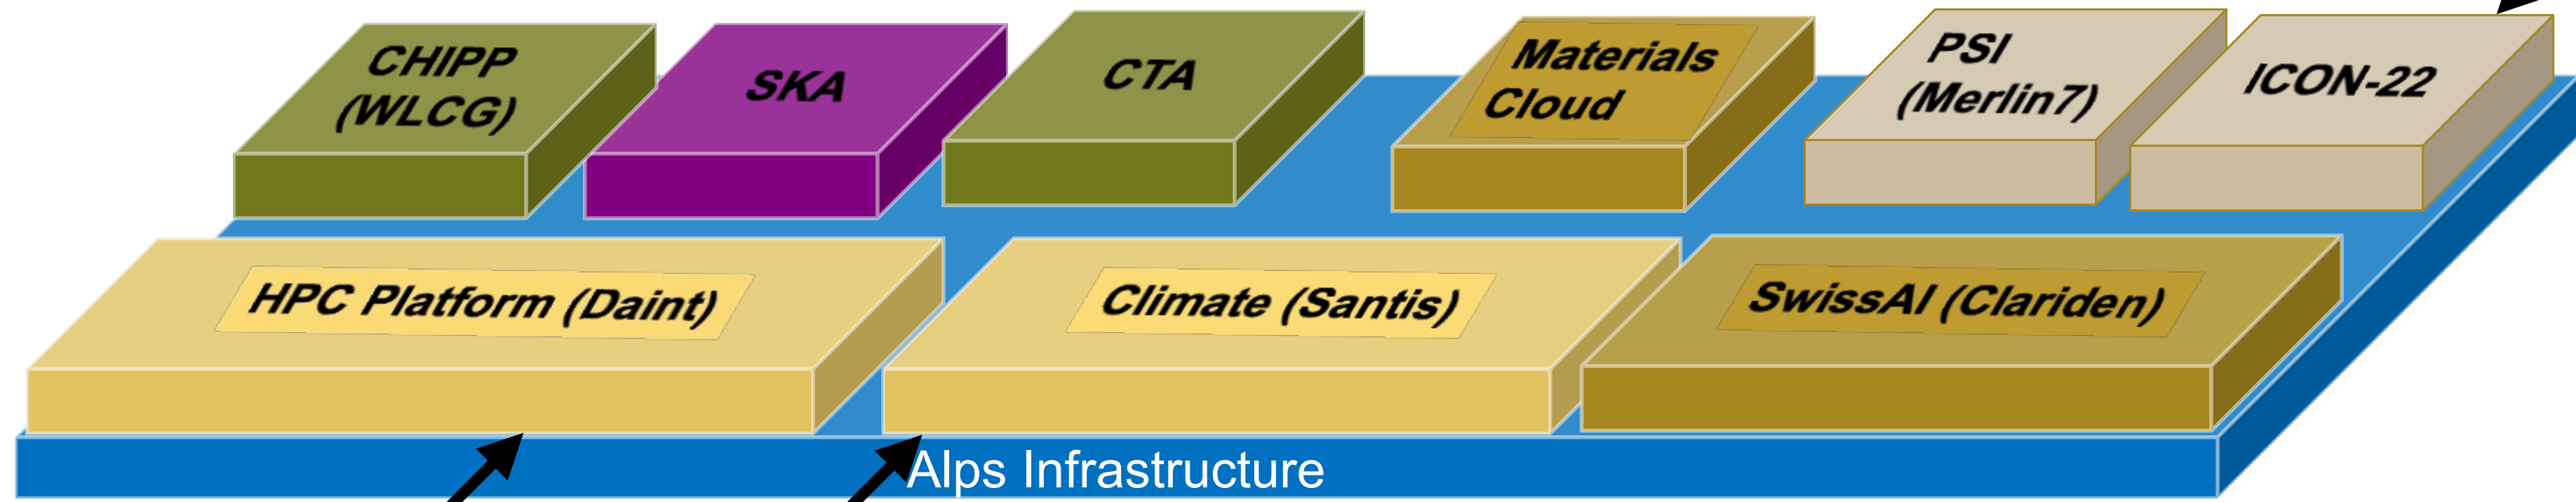


Cray-Shasta slingshot network

Power envelope: 10 MW



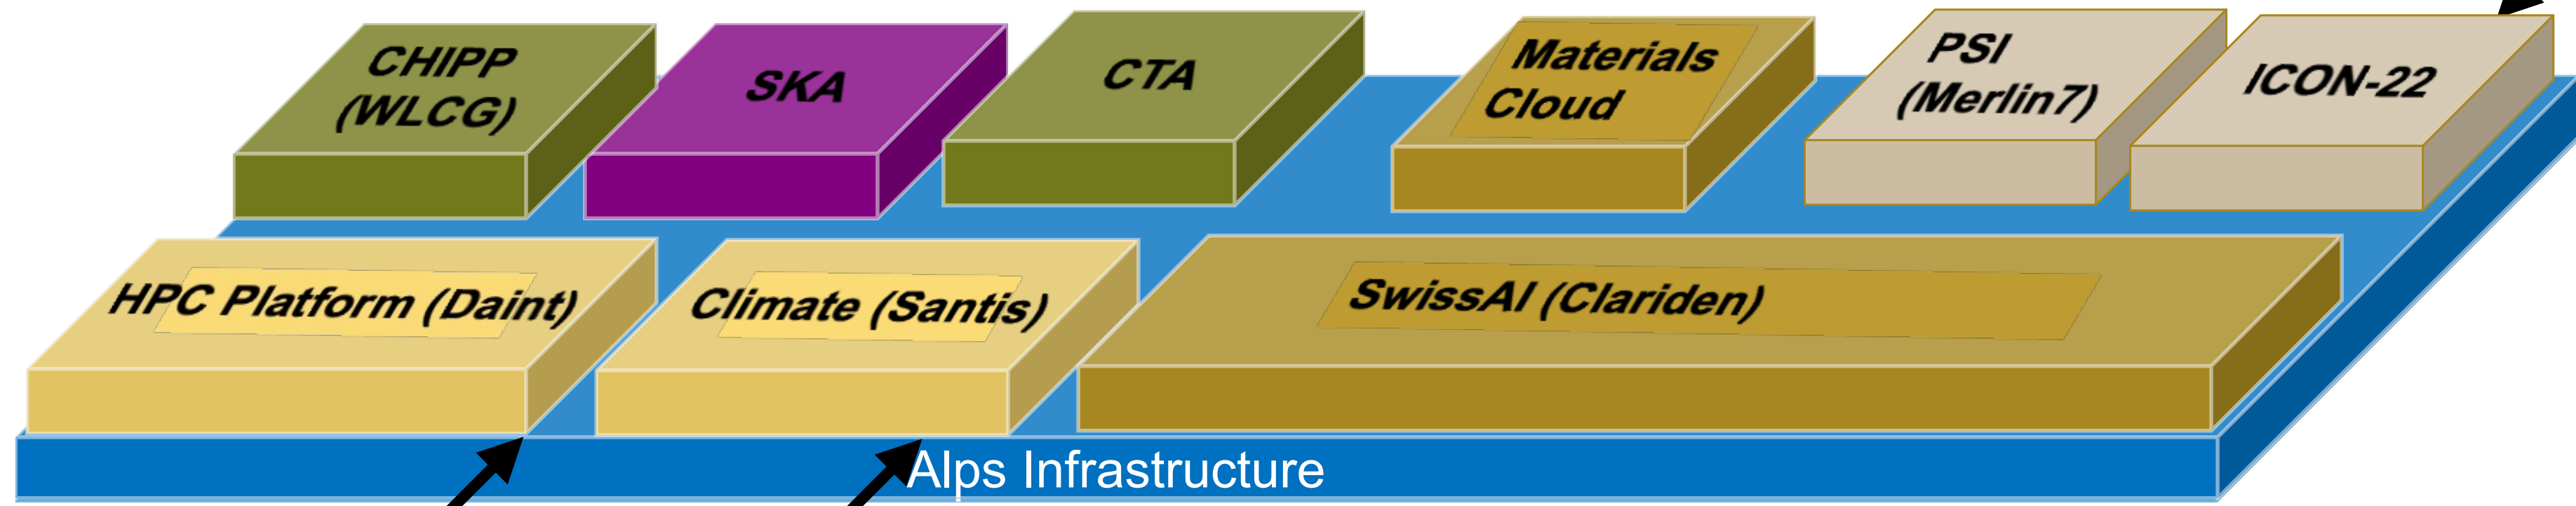
# Alps: beyond traditional supercomputing




User Program (HPCN)



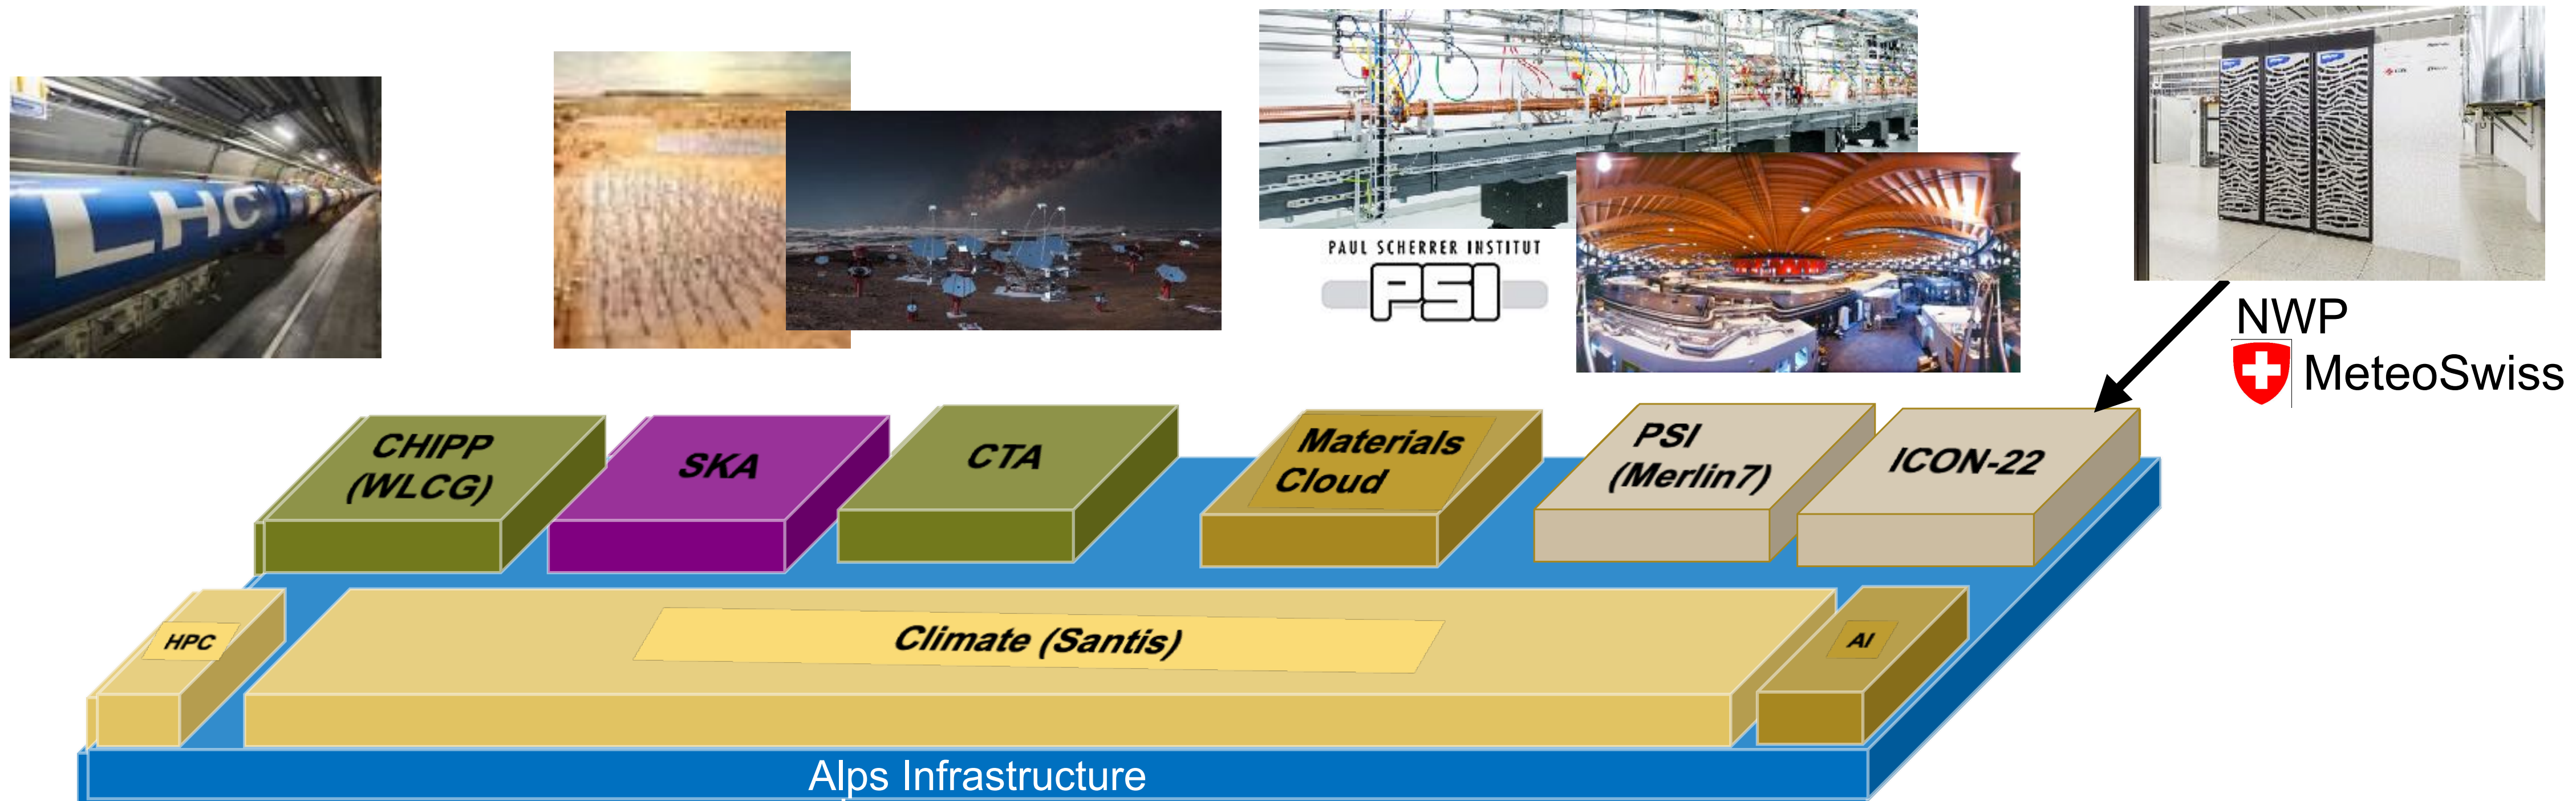
# Alps: elasticity to (e.g.) prioritise AI work



User Program (HPCN)



# Alps: supporting Gordon Bell Prize winner SC24

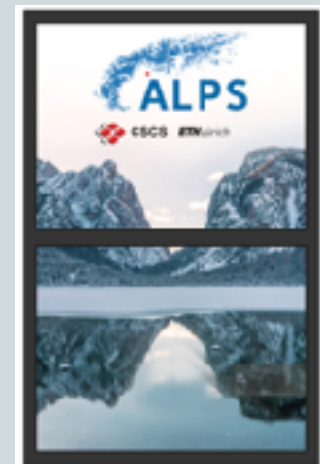


0.375 EFlop/s on 3,072 nodes (18,432 NVIDIA V100 GPUs) of Summit

0.243 EFlop/s on 1,024 nodes (4,096 NVIDIA A100 GPUs) of Leonardo

0.739 EFlop/s on 1,936 nodes (7,744 NVIDIA GH200 GPU superchips) of Alps

0.976 EFlop/s on 9,025 nodes (36,100 AMD MI250X GPUs) of Frontier

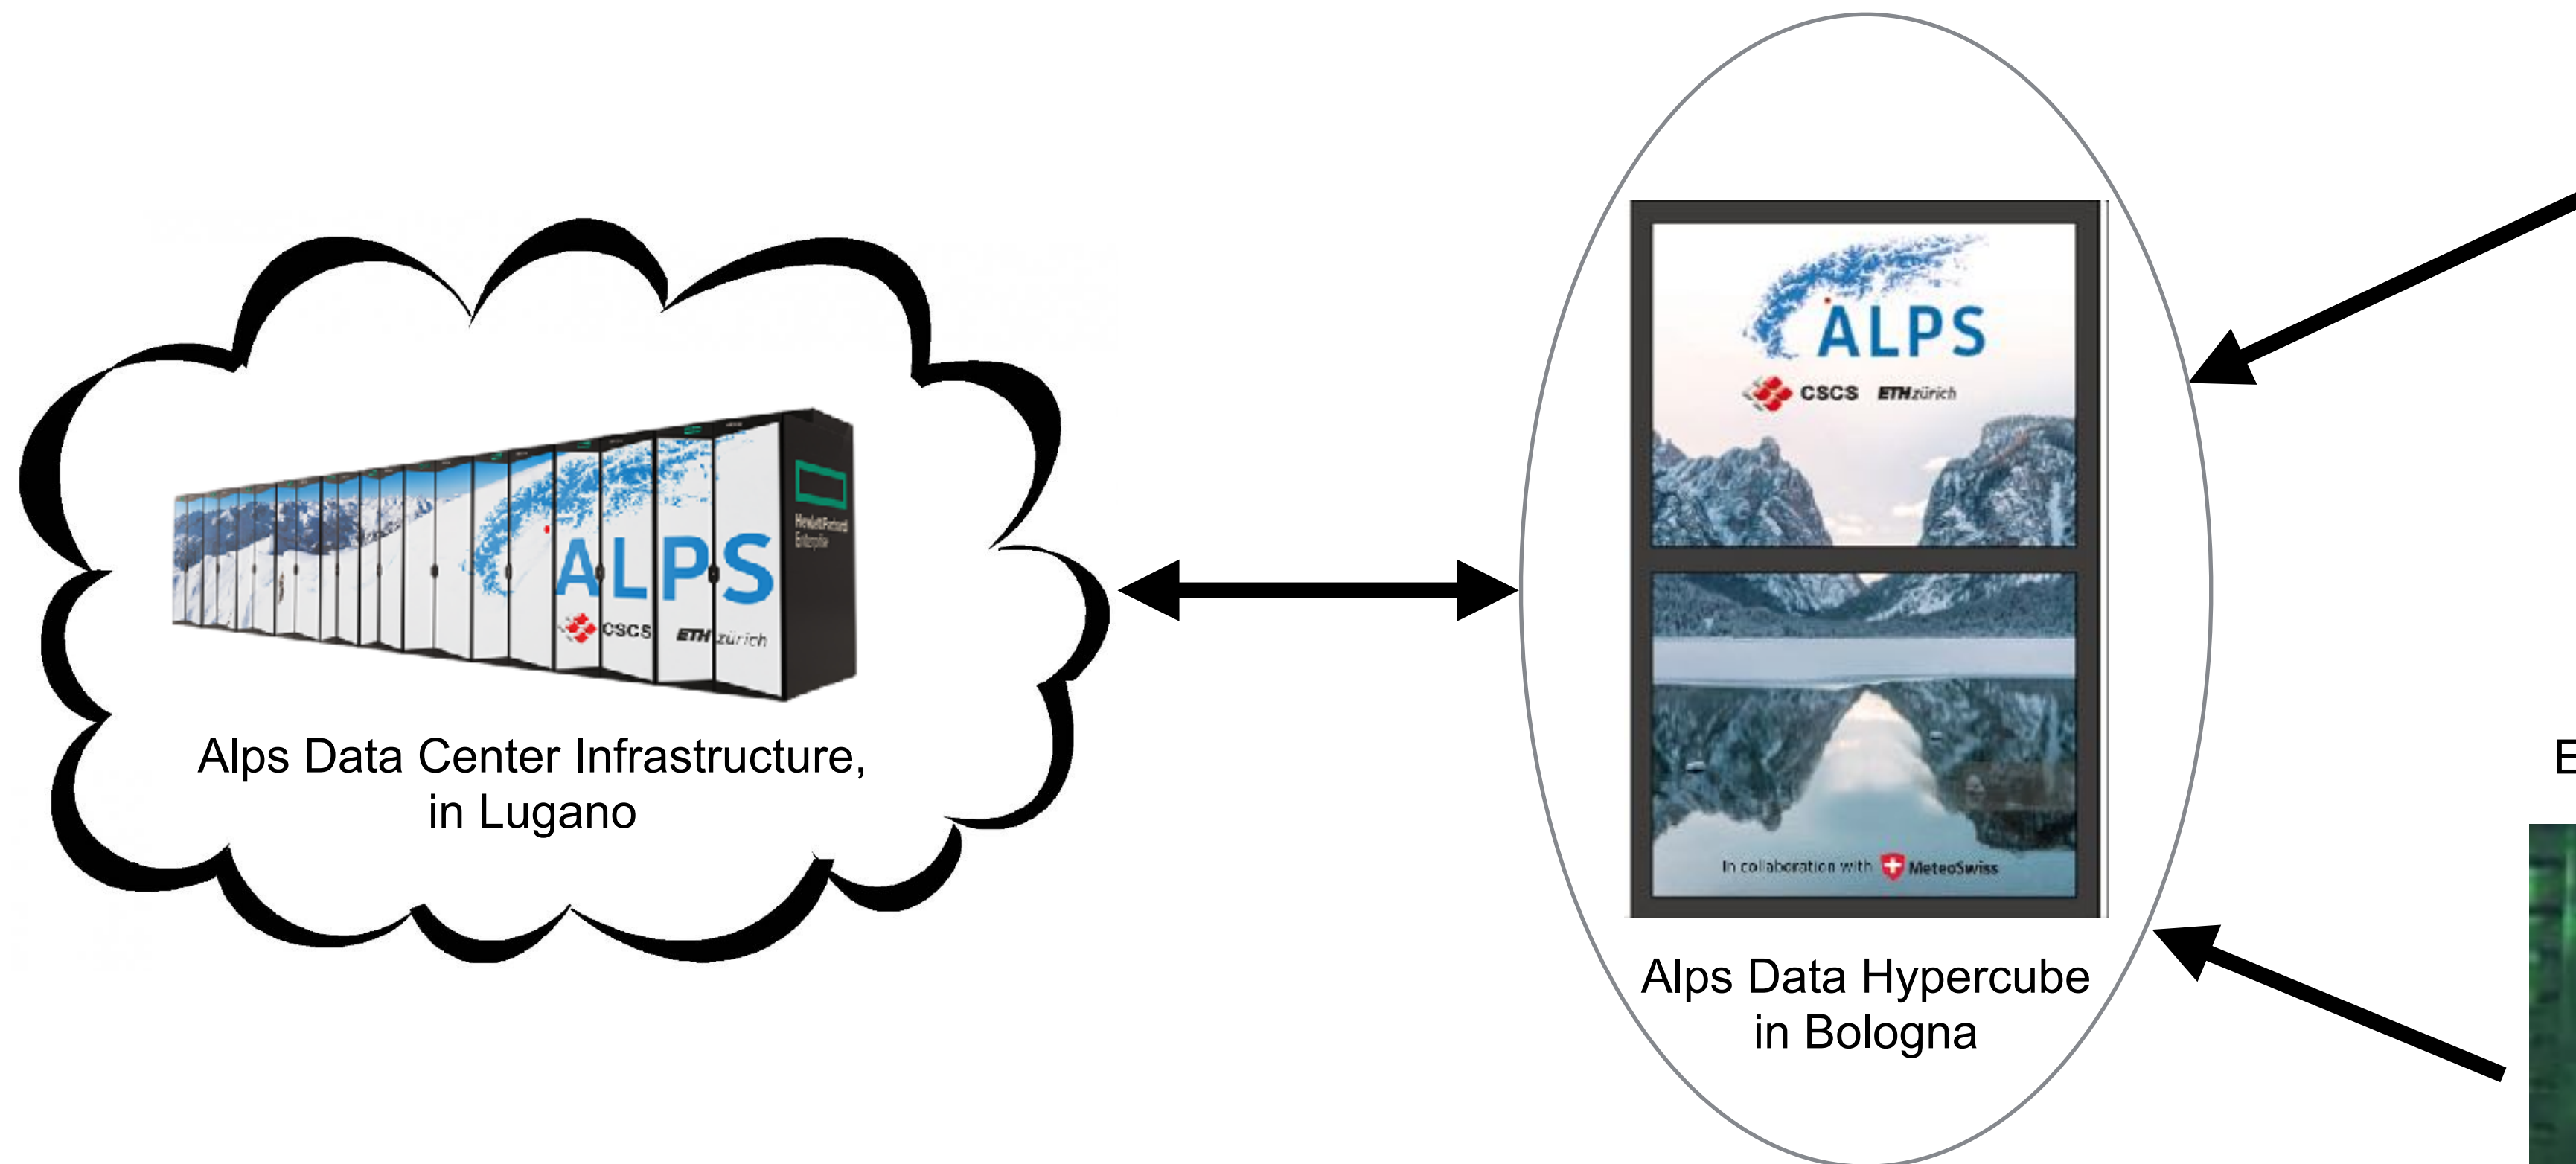


# Alps data hyper-cube at ECMWF in Bologna

IFS running @ ECMF in Bologna




ECMWF MARS archive in Bologna



# LUMI is a consortium that has deployed the first EuroHPC pre-exascale supercomputer



Scaling from 10s to 100s of MW





# International Computing and AI Network (ICAIN)

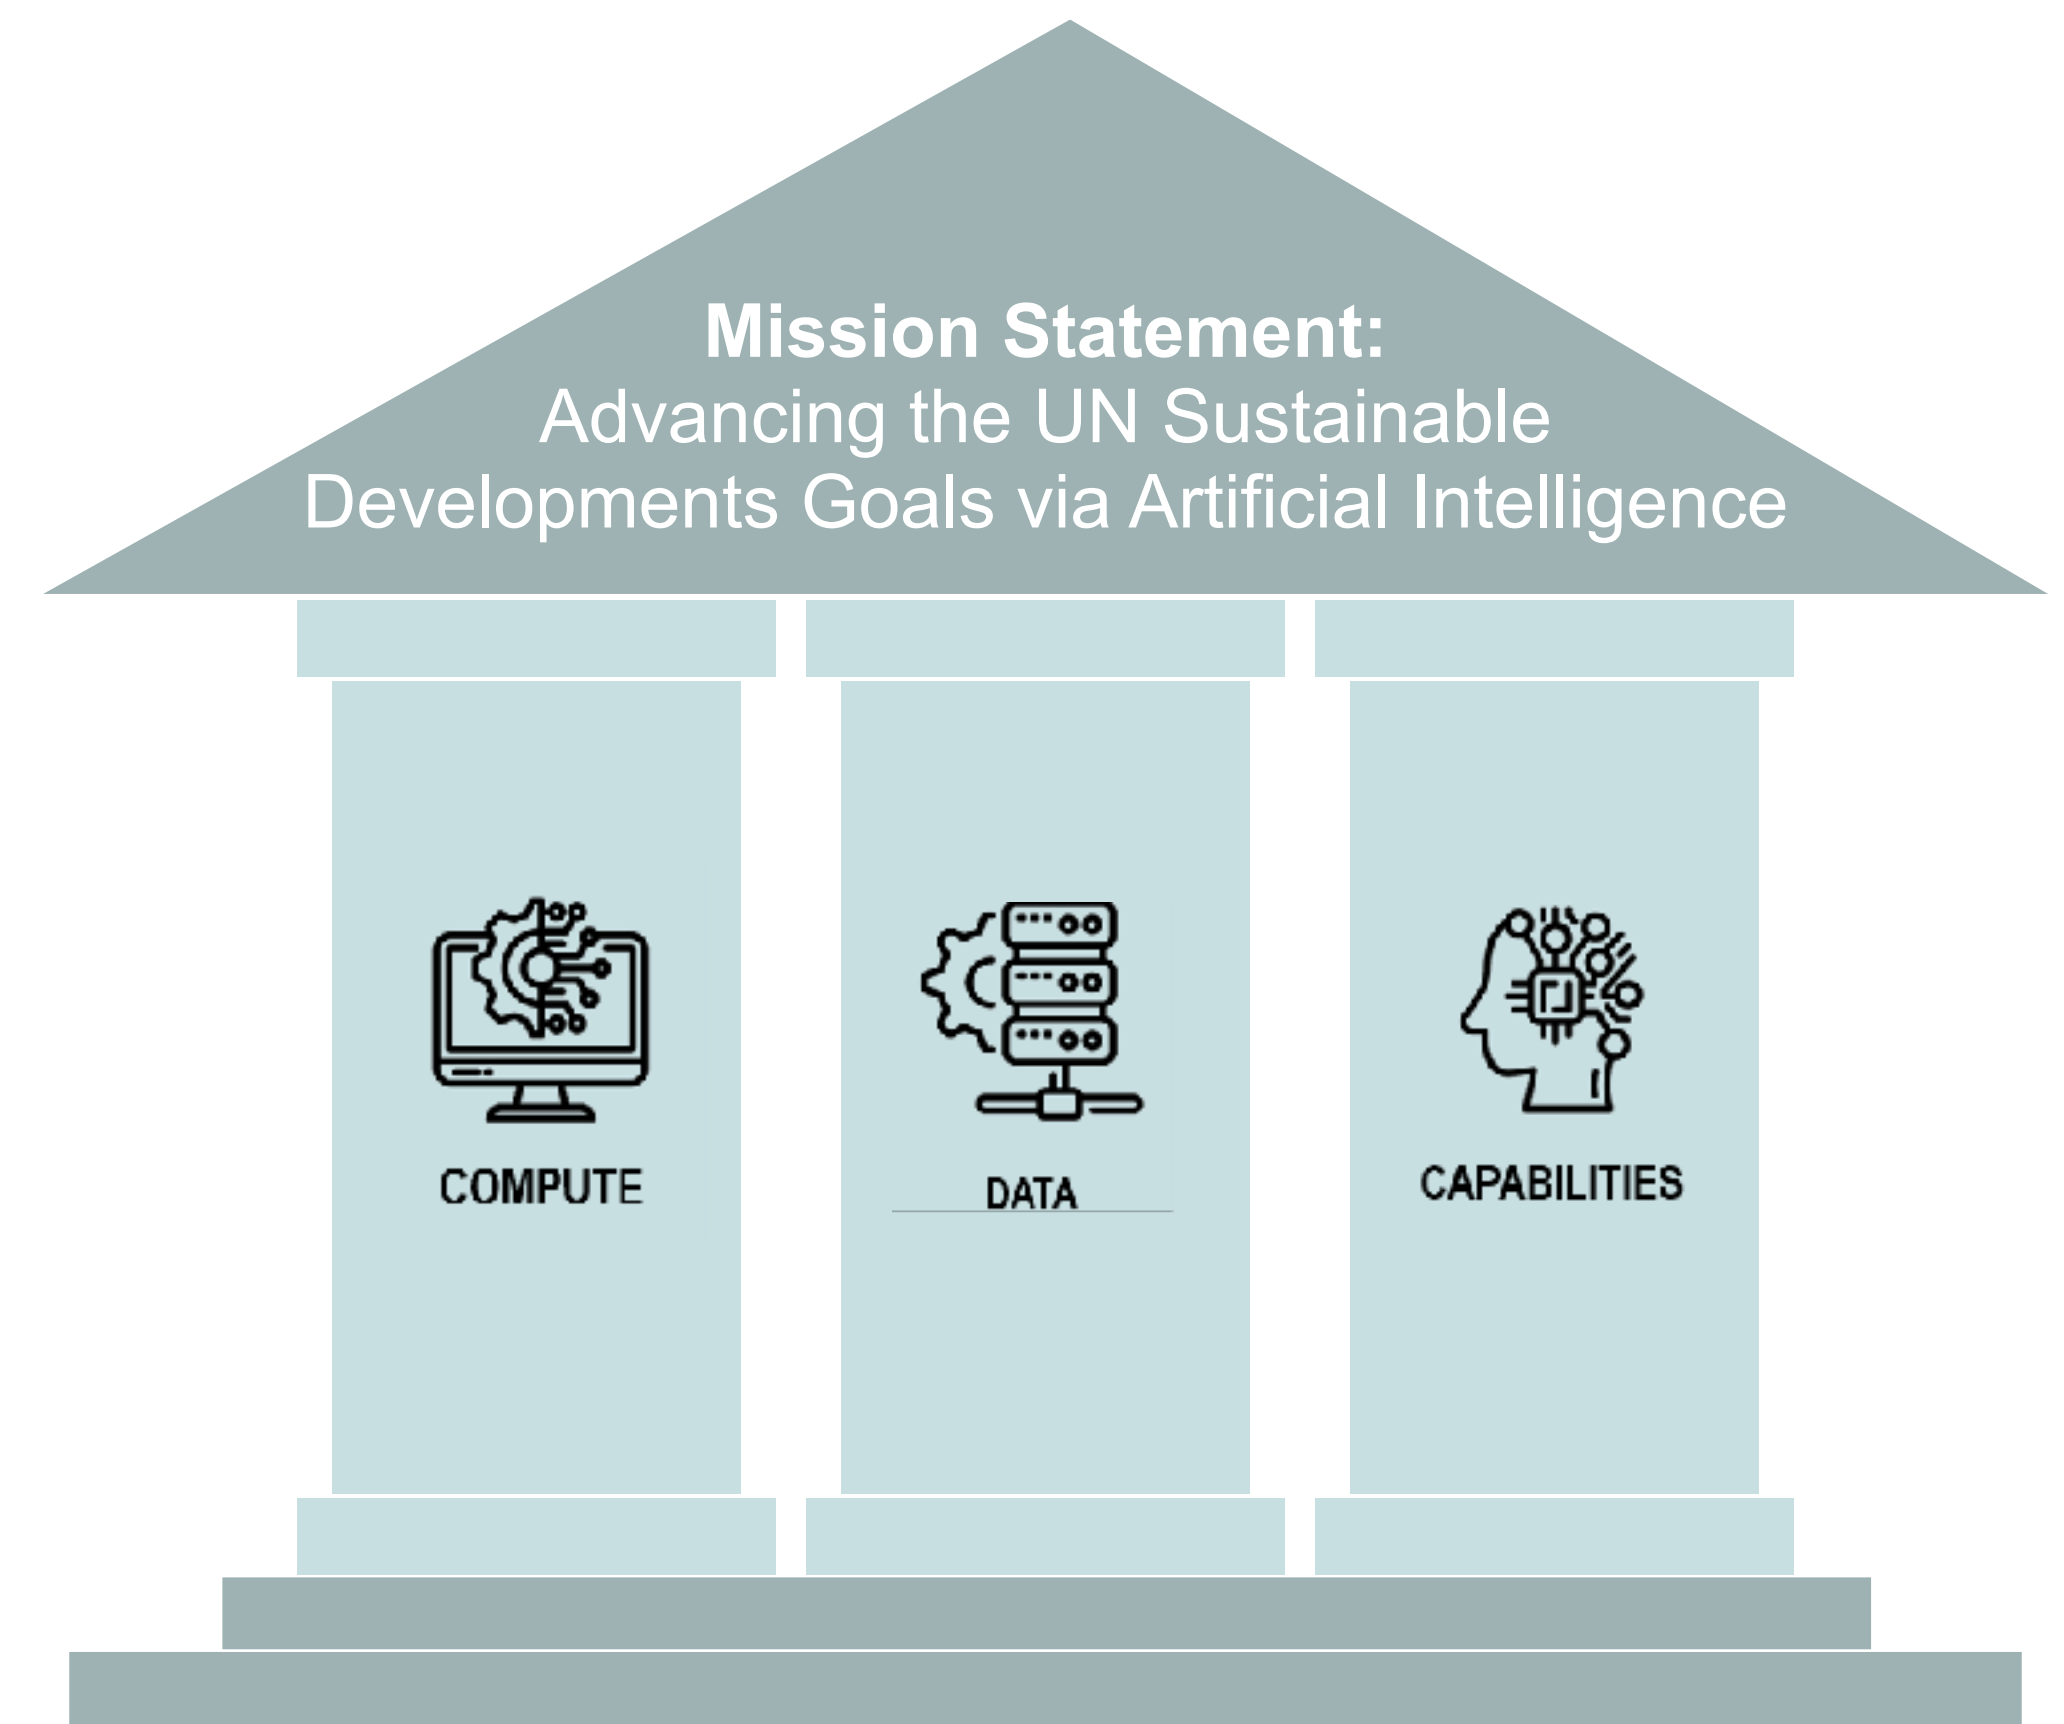
## The Vision of ICAIN

We envision a world where cutting-edge AI capabilities are harnessed to their fullest potential in pursuit of the United Nations Sustainable Development Goals, such as:


- Eradicating poverty in all its forms
- Achieving global food security and sustainable agriculture
- Combating climate change and its far-reaching impacts

## The Values of ICAIN

- **Equity:** Democratizing global AI technology access
- **Sustainability:** Harnessing global supercomputing for sustainability
- **Safety:** Anticipating and mitigating potential risks



# Meditron: LLM for clinical applications




**As we speak: Swiss AI Initiative is training 70B parameter LLM base on fully open data**



# Nudging imperfect foundation with (moove) Massive Open Online Validation and Evaluation



Nudging through use



Source: Annie Hartley

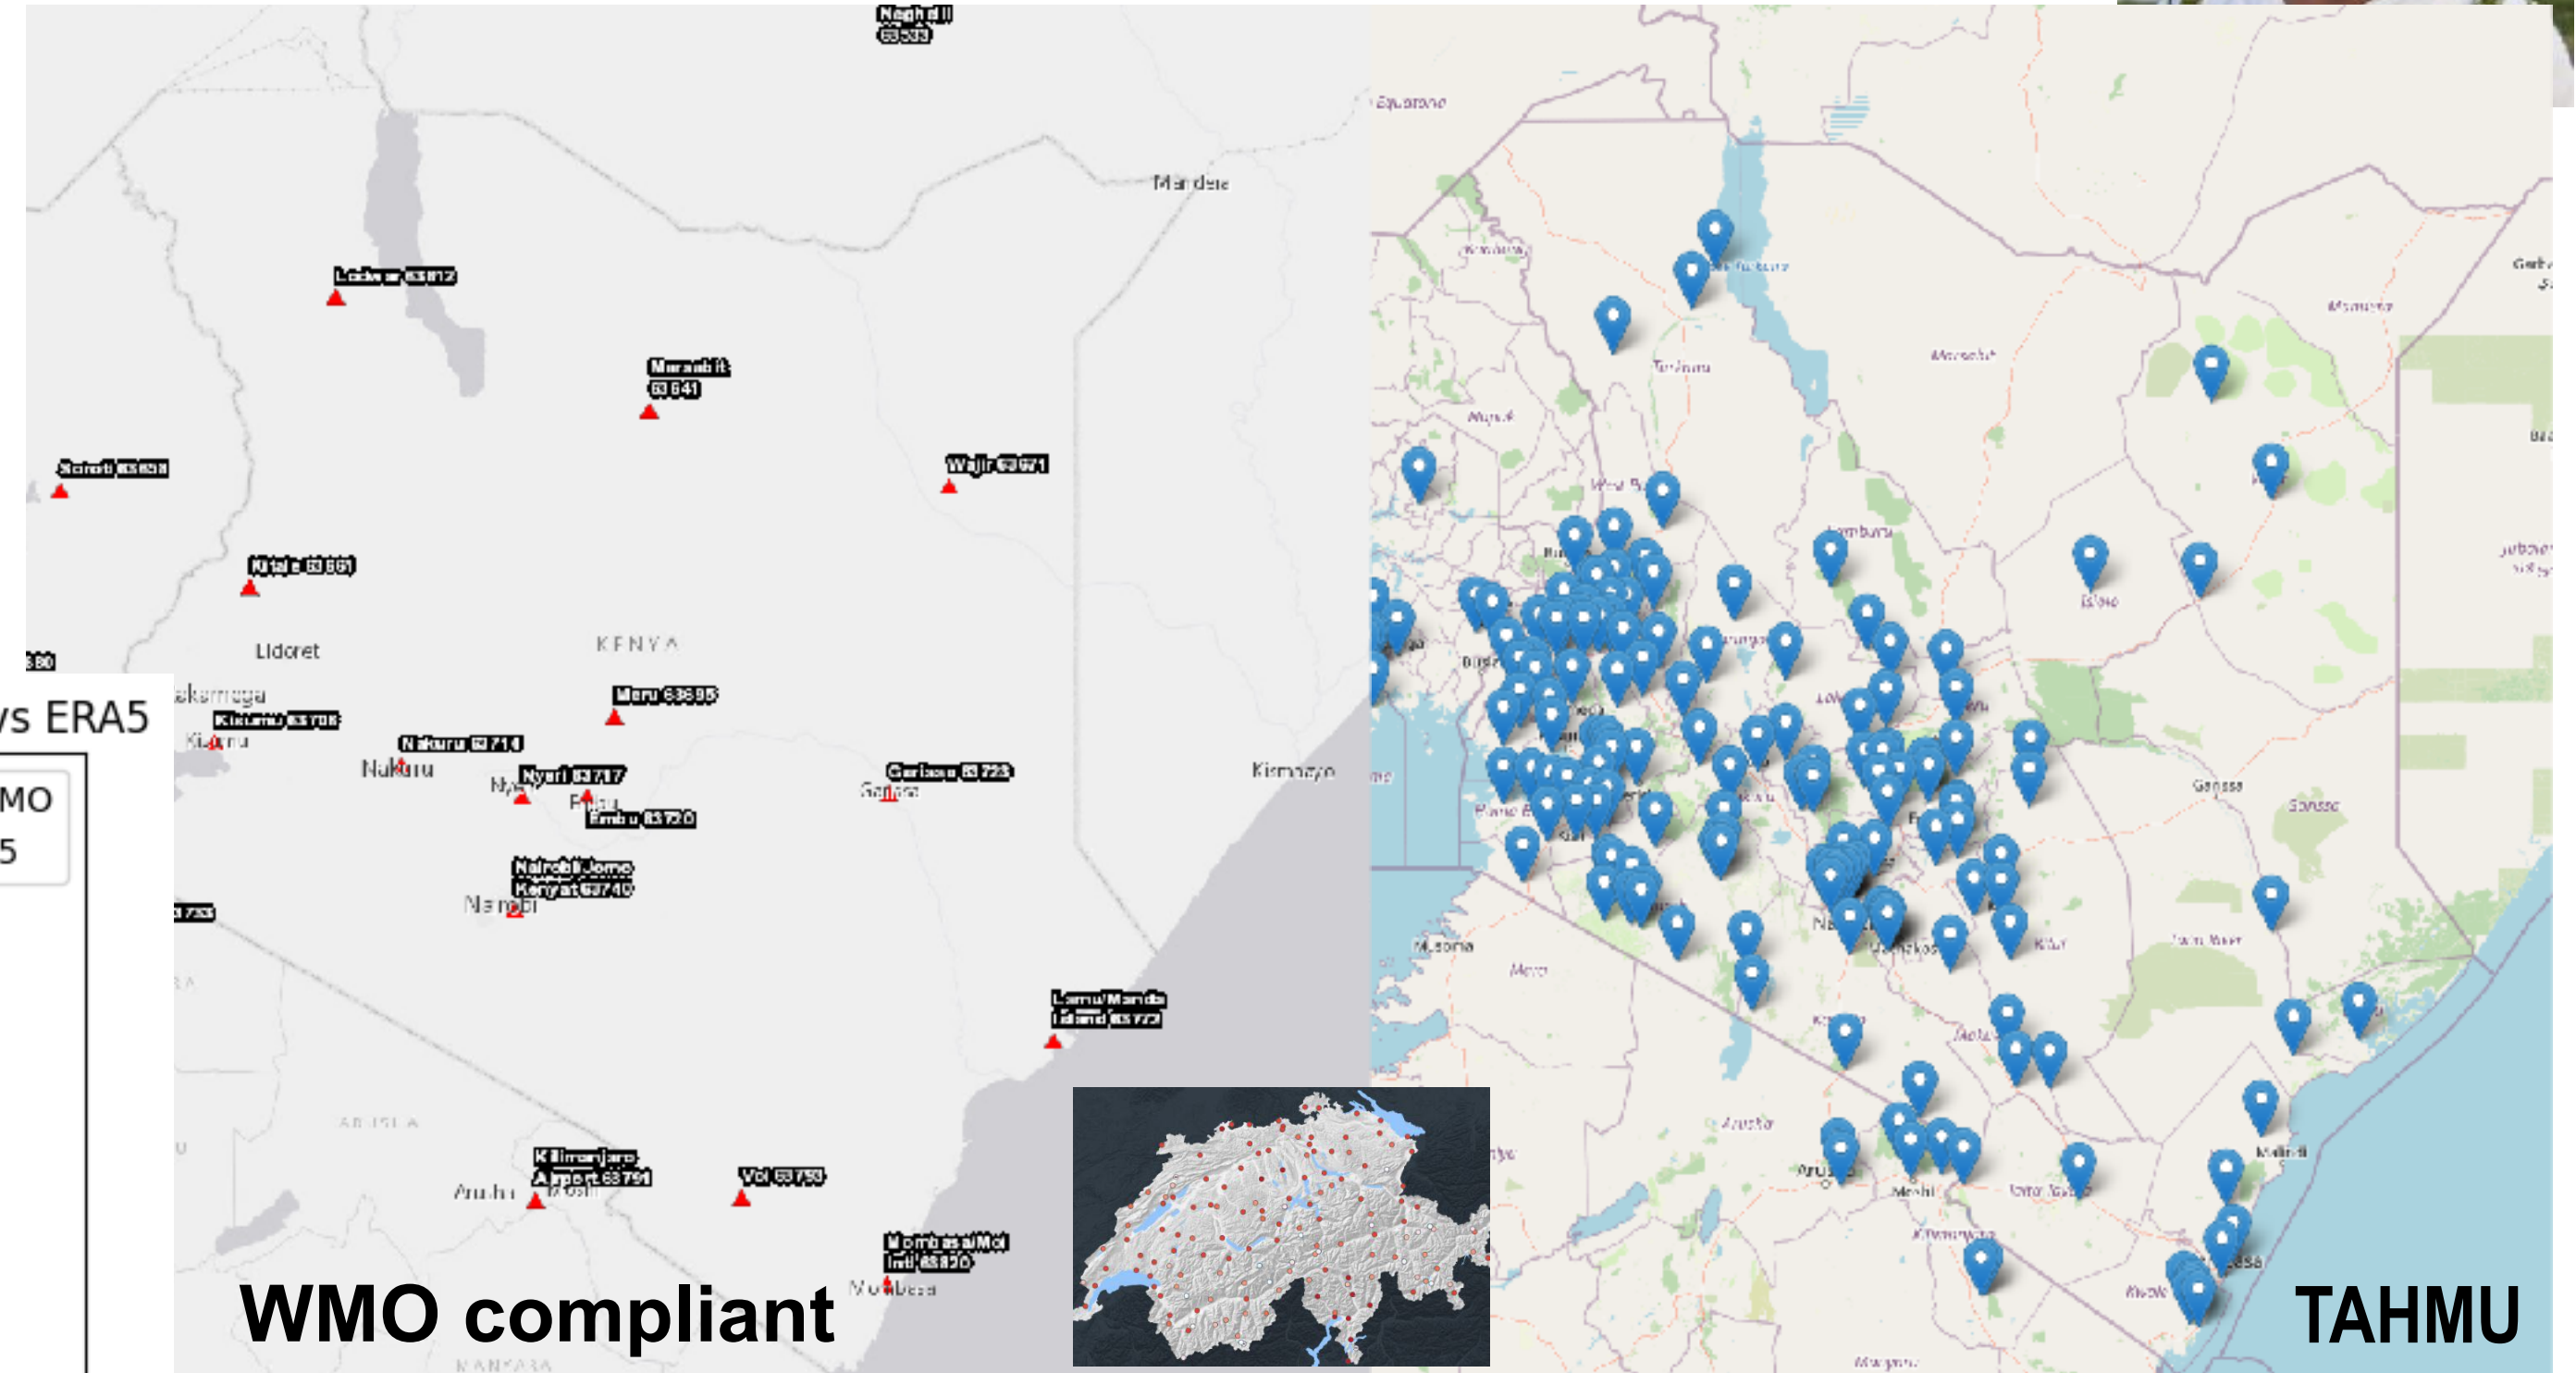
How?

- DPO
- RLHF
- RAG
- Finetuning




Early diagnosis of plant diseases through spectroscopy

# Climate and Agriculture



TA00453 Probability distribution comparison of 2m temperature (degrees celcius) - TAHMO vs ERA5



AI models (e.g. Graphcast) trained with ERA5 data  
 Comparing (ERA5) reanalysis data with TAHMU data

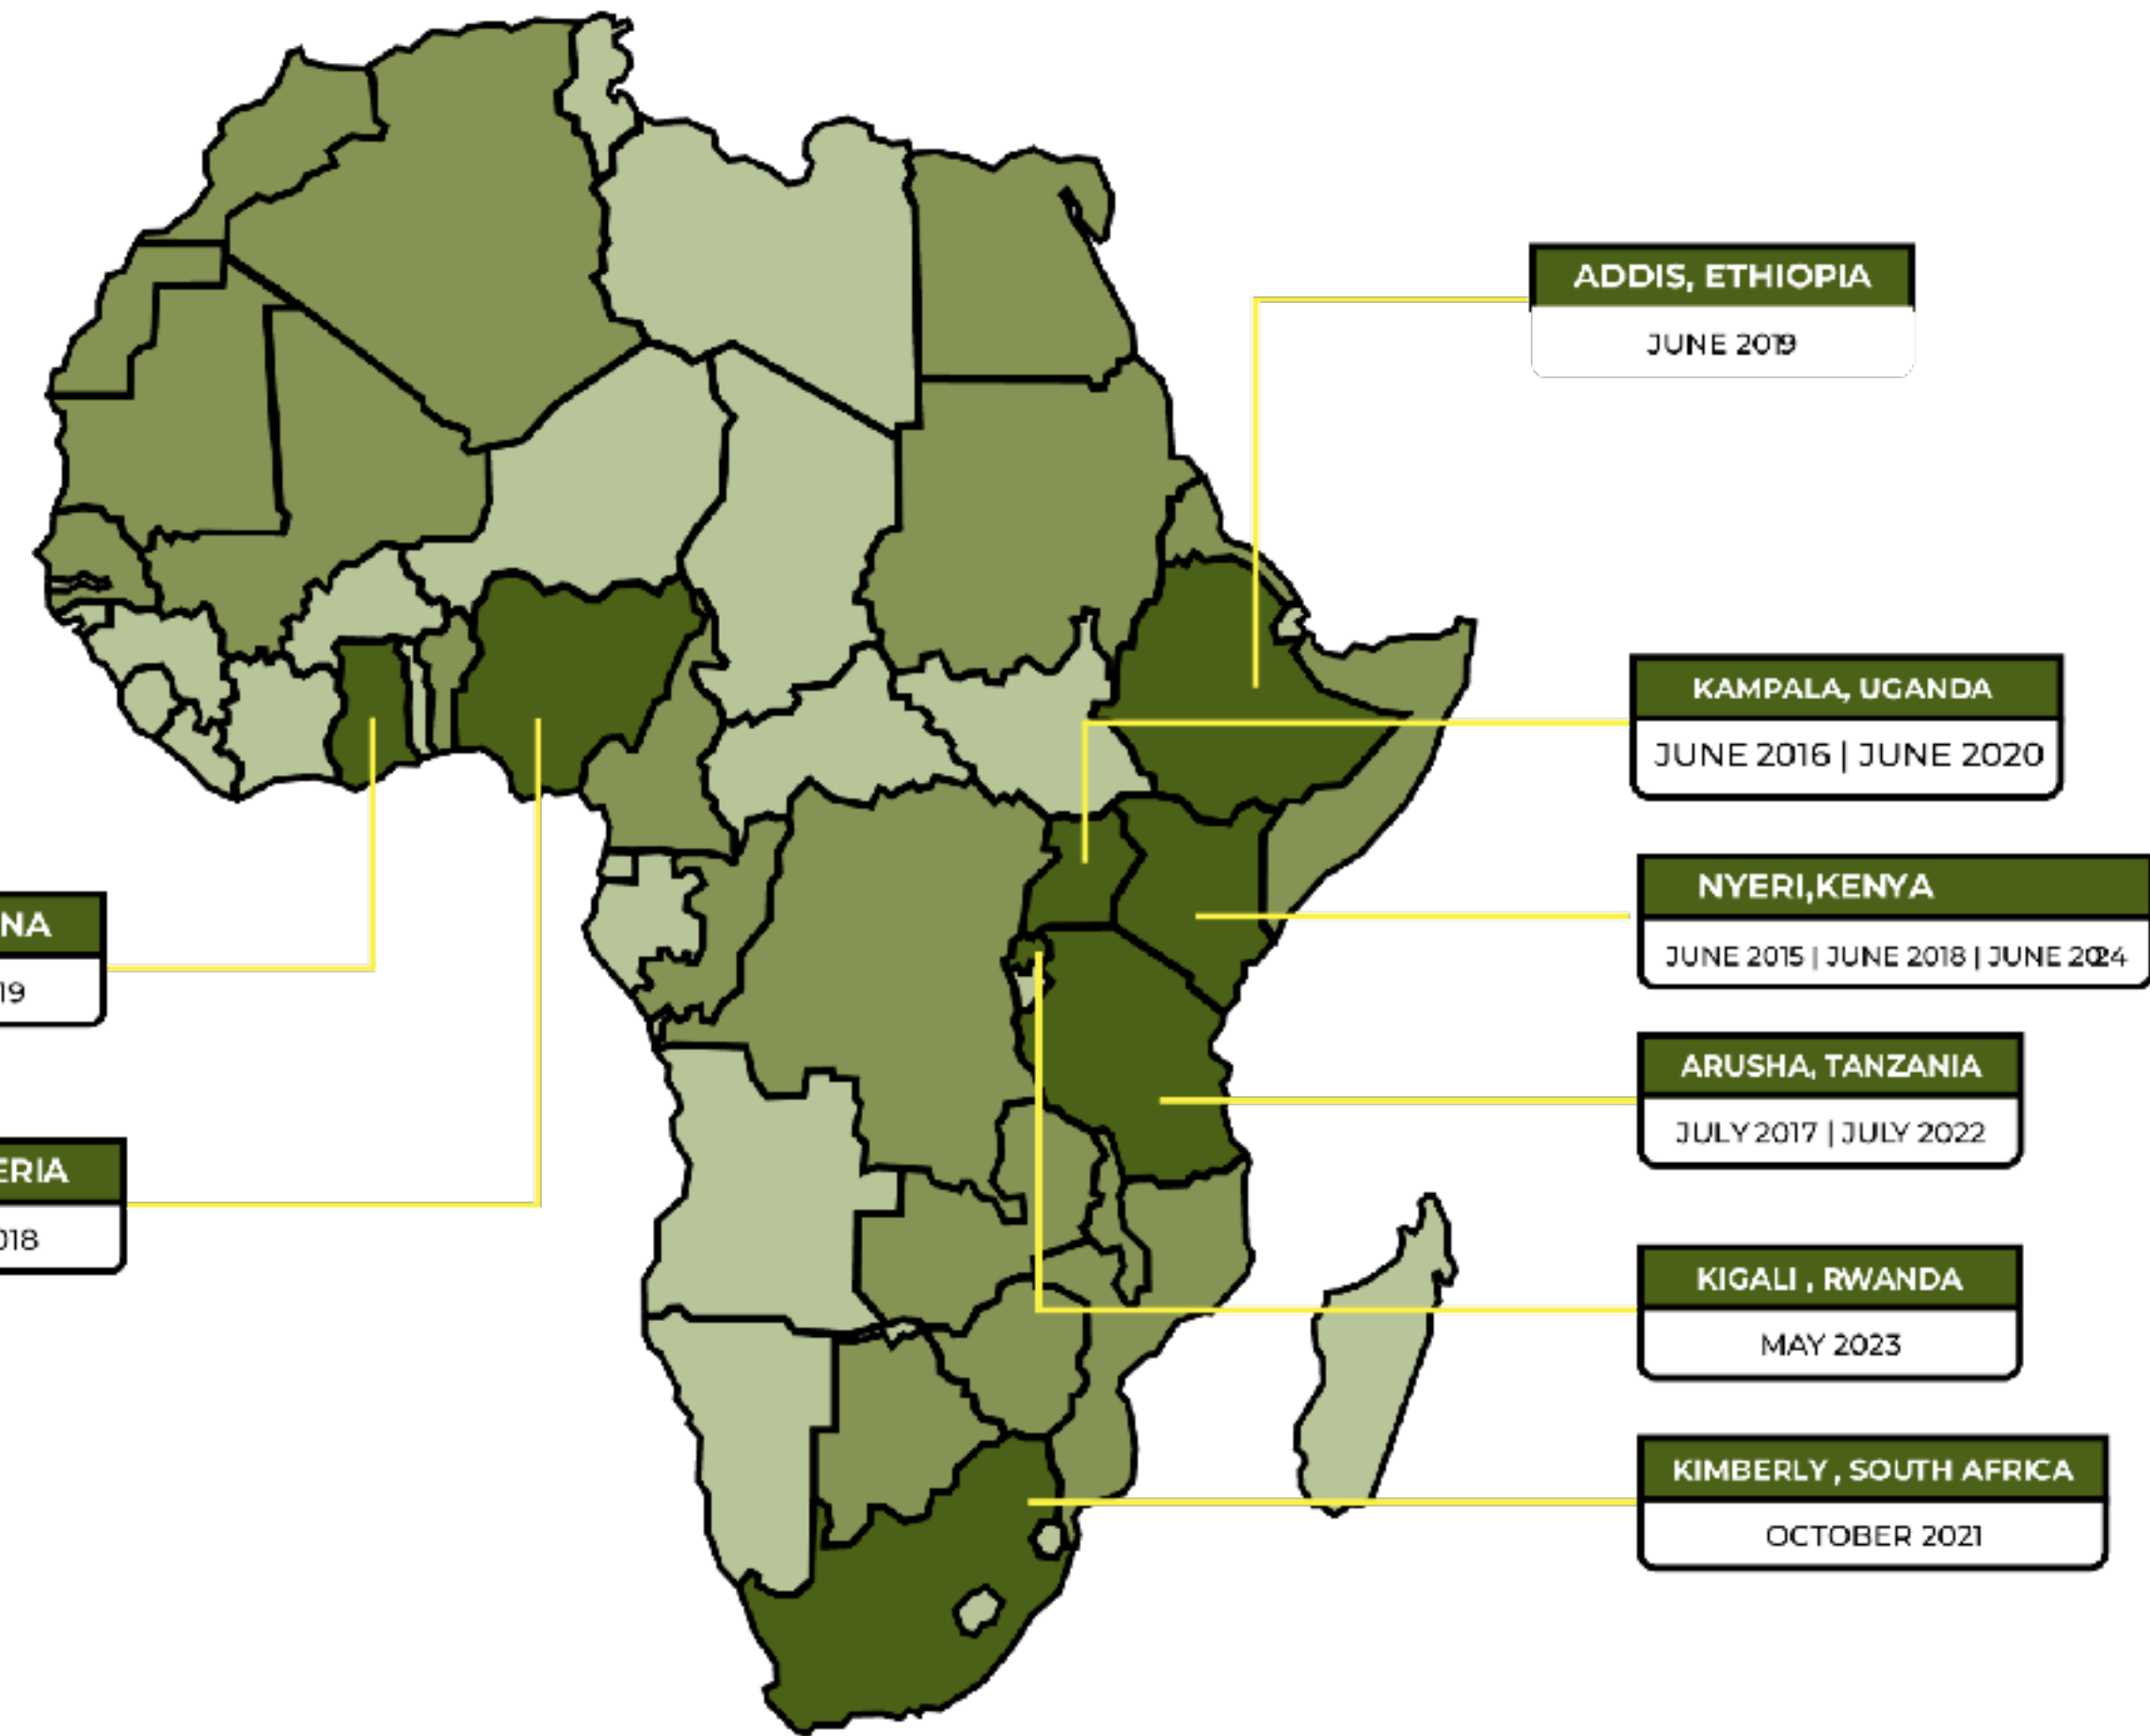
# Data Science Africa

Grassroots capability building organisation

- By Africa
- Student focus
- About solutions
- Sustainable and inclusive
- Agility




June 2024, Nyeri Kenya



# The opportunity has implications ...

- Budget to fund cost of electricity and future investments
  - 90 GiWh in 2025 -> CHF 20.7 Mio.
  - Alps would cost CHF 400 Mio. if we hadn't placed the order in 2020
- Legislative and societal handling of artificial intelligence
  - Copyright and data protection
  - Bridging academic research and commercial applications of AI



Source: DALL-E

Thank you to the teams at CSCS, HPE and  
NVIDIA, as well as our partners,  
and thank for your interest