ETH zürich

Using Acceleration as a Service (XaaS): Rethinking HPC Infrastructure in the AI Era

Mauro Bianco, Torsten Hoefler, Mark Klein, Maxime Martinasso, Pim Witlox, Thomas C. Schulthess, and Stefano Schuppli

CSCS
Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich

XaaS: Acceleration as a Service to Enable Productive High-Performance Cloud Computing

JOURNAL ARTICLE · 02 April 2024 · Computing in Science and Engineering

DOI: https://doi.org/10.1109/mcse.2024.3382154 · OSTI ID: 2545755

Hoefler, Torsten [1]; Copik, Marcin [1]; Beckman, Pete [2]; Jones, Andrew [3]; Foster, Ian [2]; Parashar, Manish [4]; Reed, Daniel [4]; Troyer, Matthias [3]; Schulthess, Thomas [5]; Ernst, Daniel [6]; Dongarra, Jack [7]

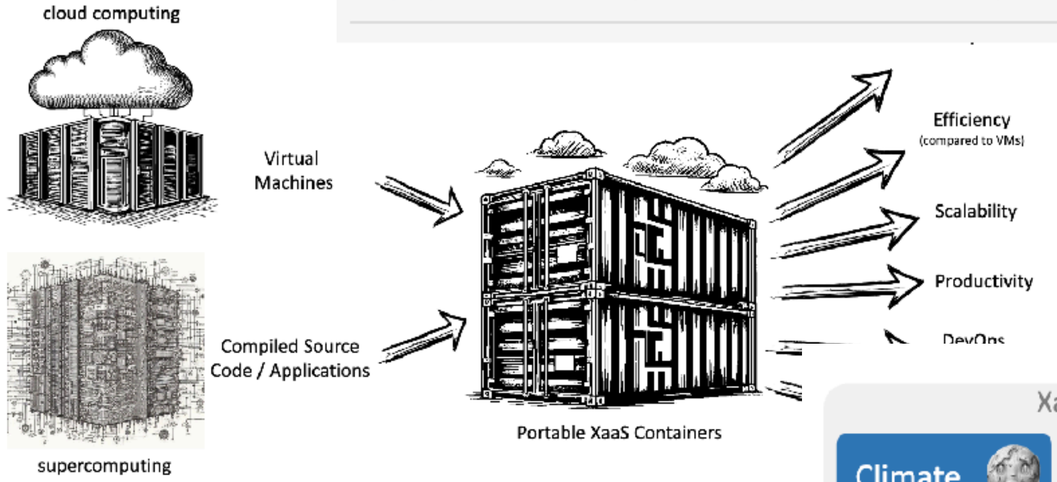

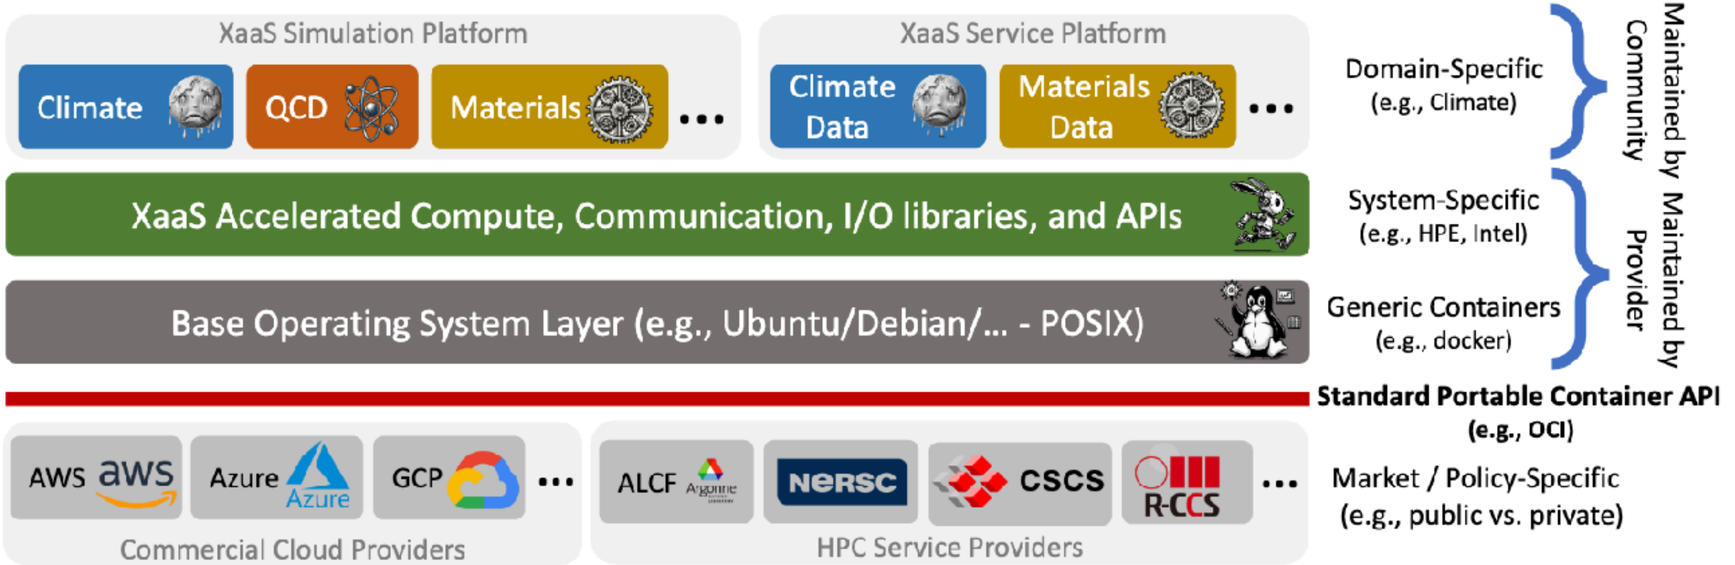

OUR MISSION (since 2015)

"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 RESEACH 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

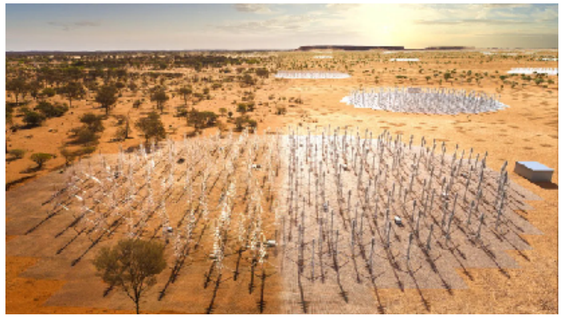

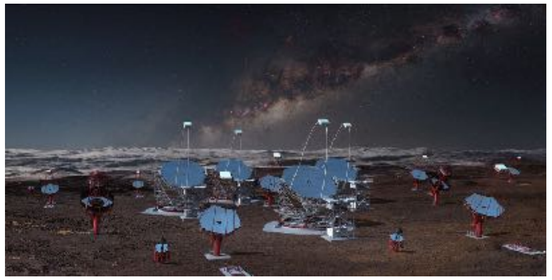

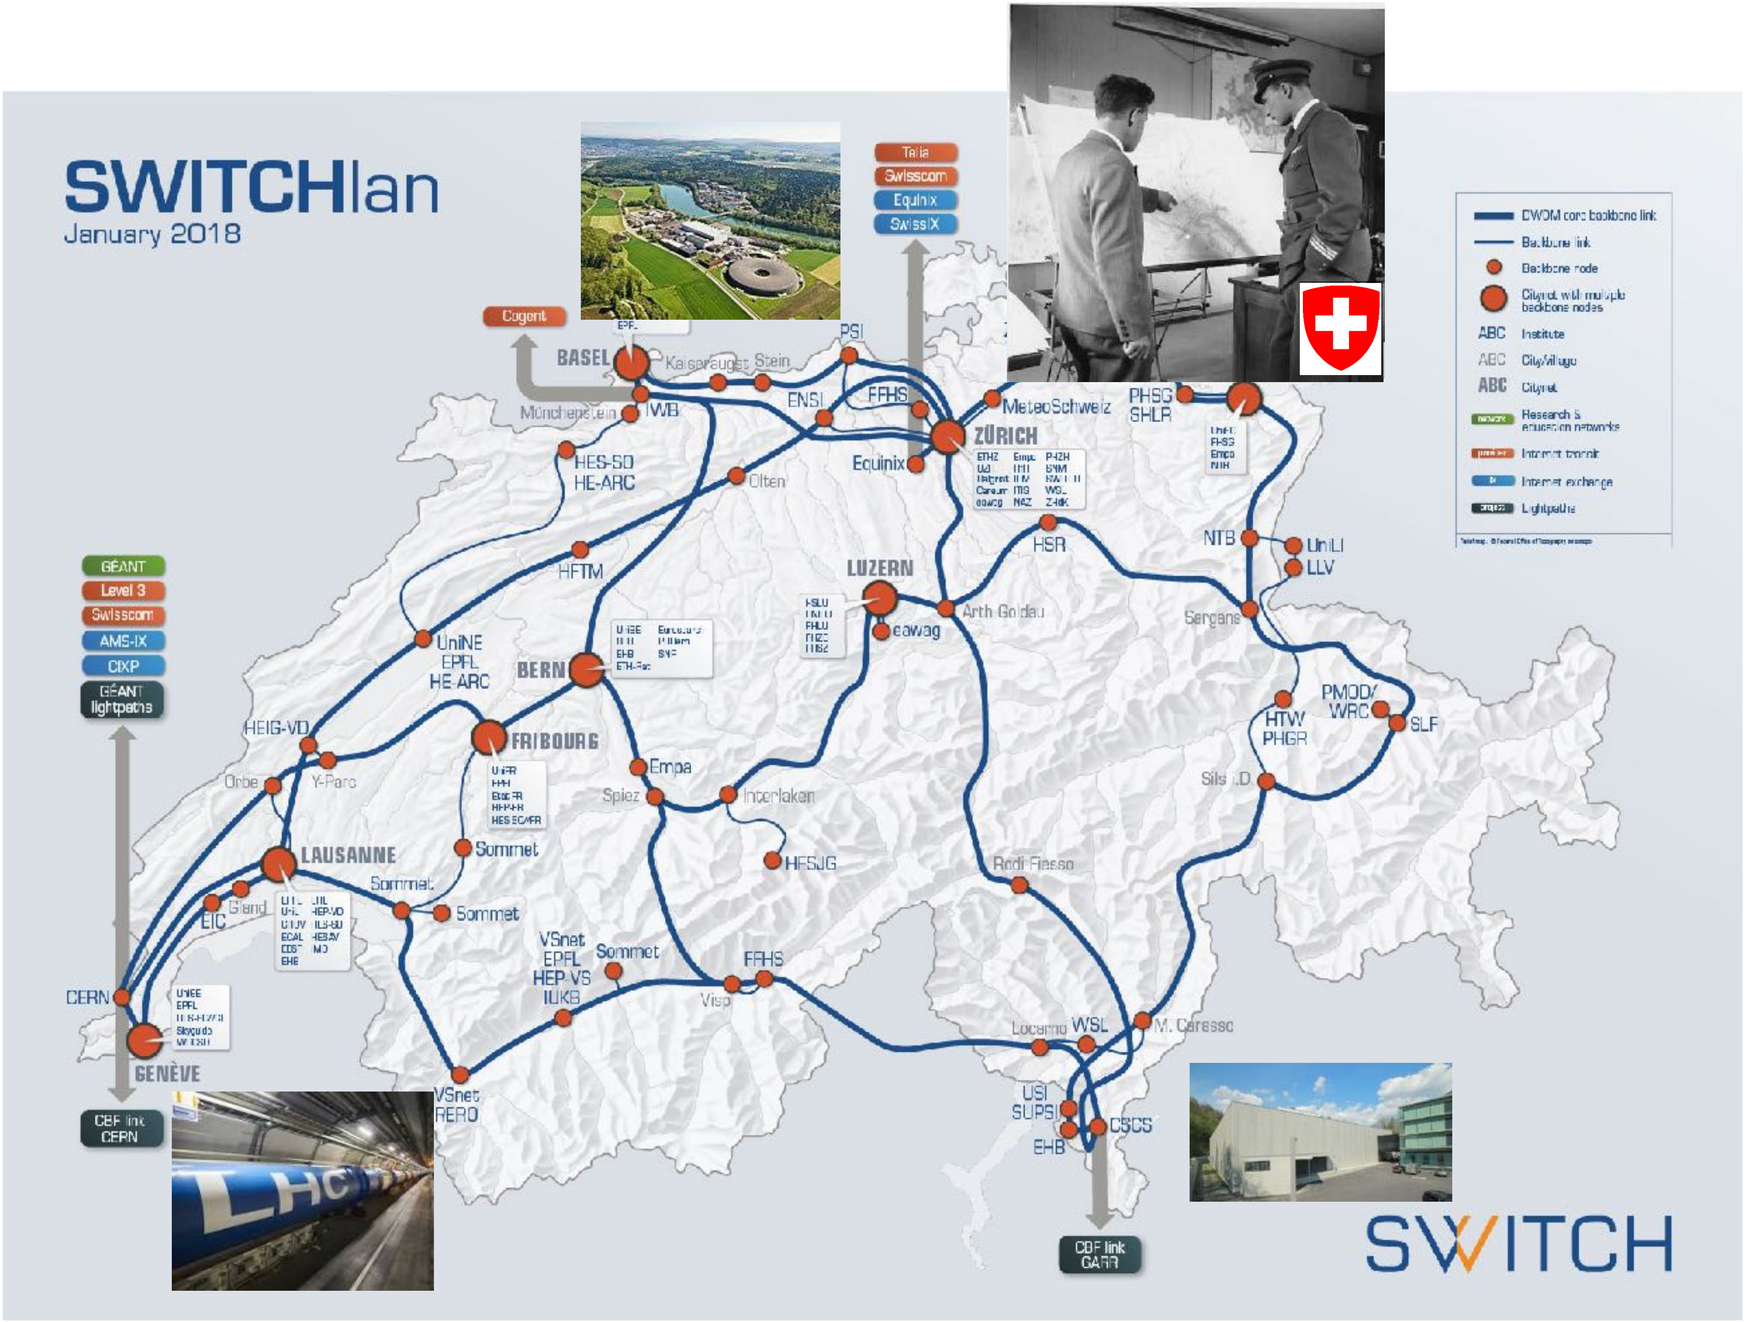
SWITCHlan
January 2018
Cogent
Telia
Swisscom
Equinix
SwissIX
BASEL
Kaiseraugst
Stein
PSI
ENSI
FFHS
MeteoSchweiz
PHSG
SHLR
ZÜRICH
Equinix
Olten
HES-SO
HE-ARC
HSR
NTB
UniLI
LLV
Sargans
LUZERN
eawag
Arth-Goldau
HFTM
UniNE
EPFL
HE-ARC
BERN
FRIBOURG
HEIG-VD
Orbe
Y-Parc
Empa
Spiez
Interlaken
HFSJG
Rodi Fiesso
LAUSANNE
Sommet
Gland
EIC
VSnet
EPFL
HEP-VS
IUKB
Visp
Locarno
WSL
M. Ceresio
USI
SUPSI
EHB
CSCS
PMOD/
WRC
SLF
HTW
PHGR
Sils i.D.
CERN
GENÈVE
VSnet
RERO
GÉANT
Level 3
Swisscom
AMS-IX
CIXP
GÉANT lightpaths
CBF link CERN
CBF link GARR
DWDM core backbone link
Backbone link
Backbone node
City/Village with multiple backbone nodes
Institute
City/Village
City/Village
Research & education networks
Internet transit
Internet exchange
Lightpaths
SWITCH

ETH zürich

MeteoSwiss

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

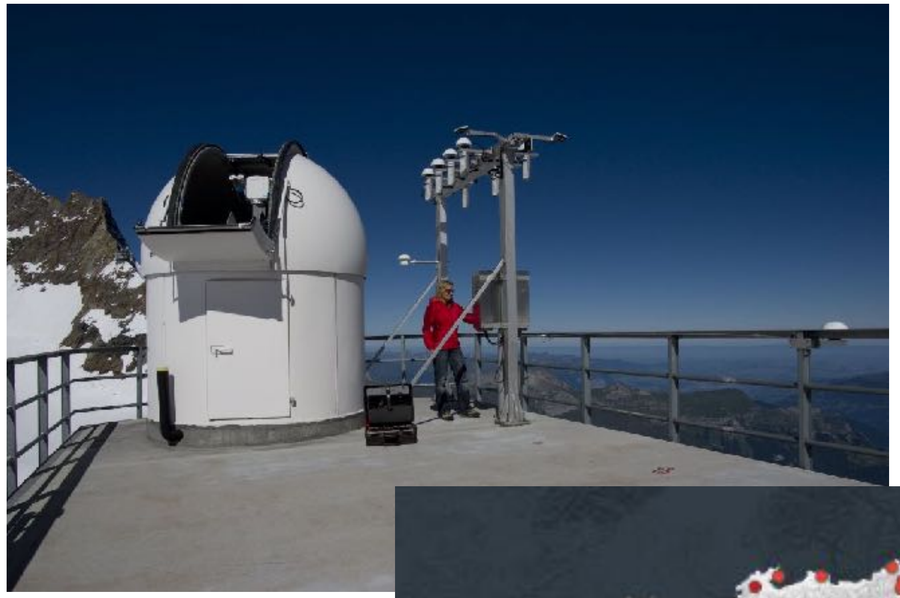

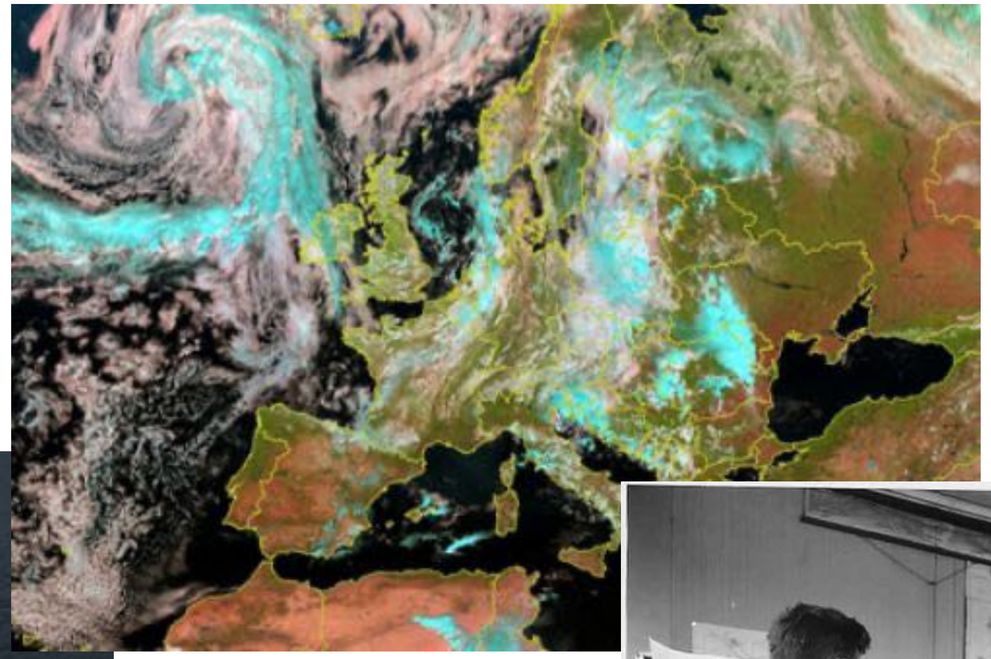

Source: meteoswiss.admin.ch

SwissFEL, SLS, SinQ, and SμS at PSI

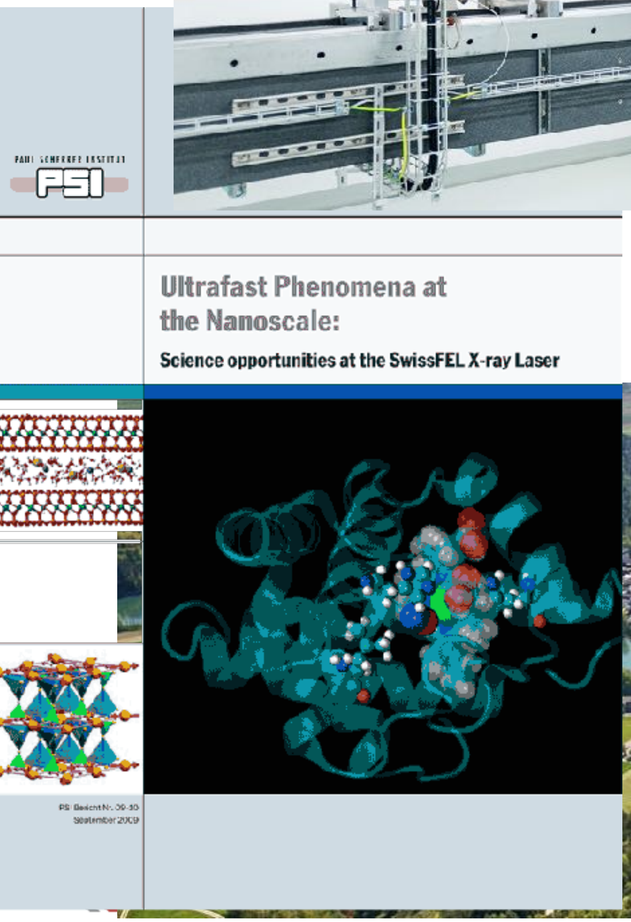

4 User Laboratories

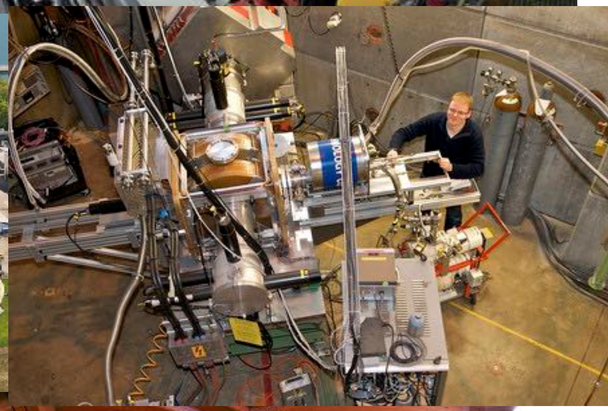

Specifications (co-designe use cases) for “Alps”

1. Climate and weather community for “big data” problem — most complete in terms for 5 Vs
2. Strong multi-tenancy: PSI has to run their own user program on the infrastructure

Idea: creating specialised clusters

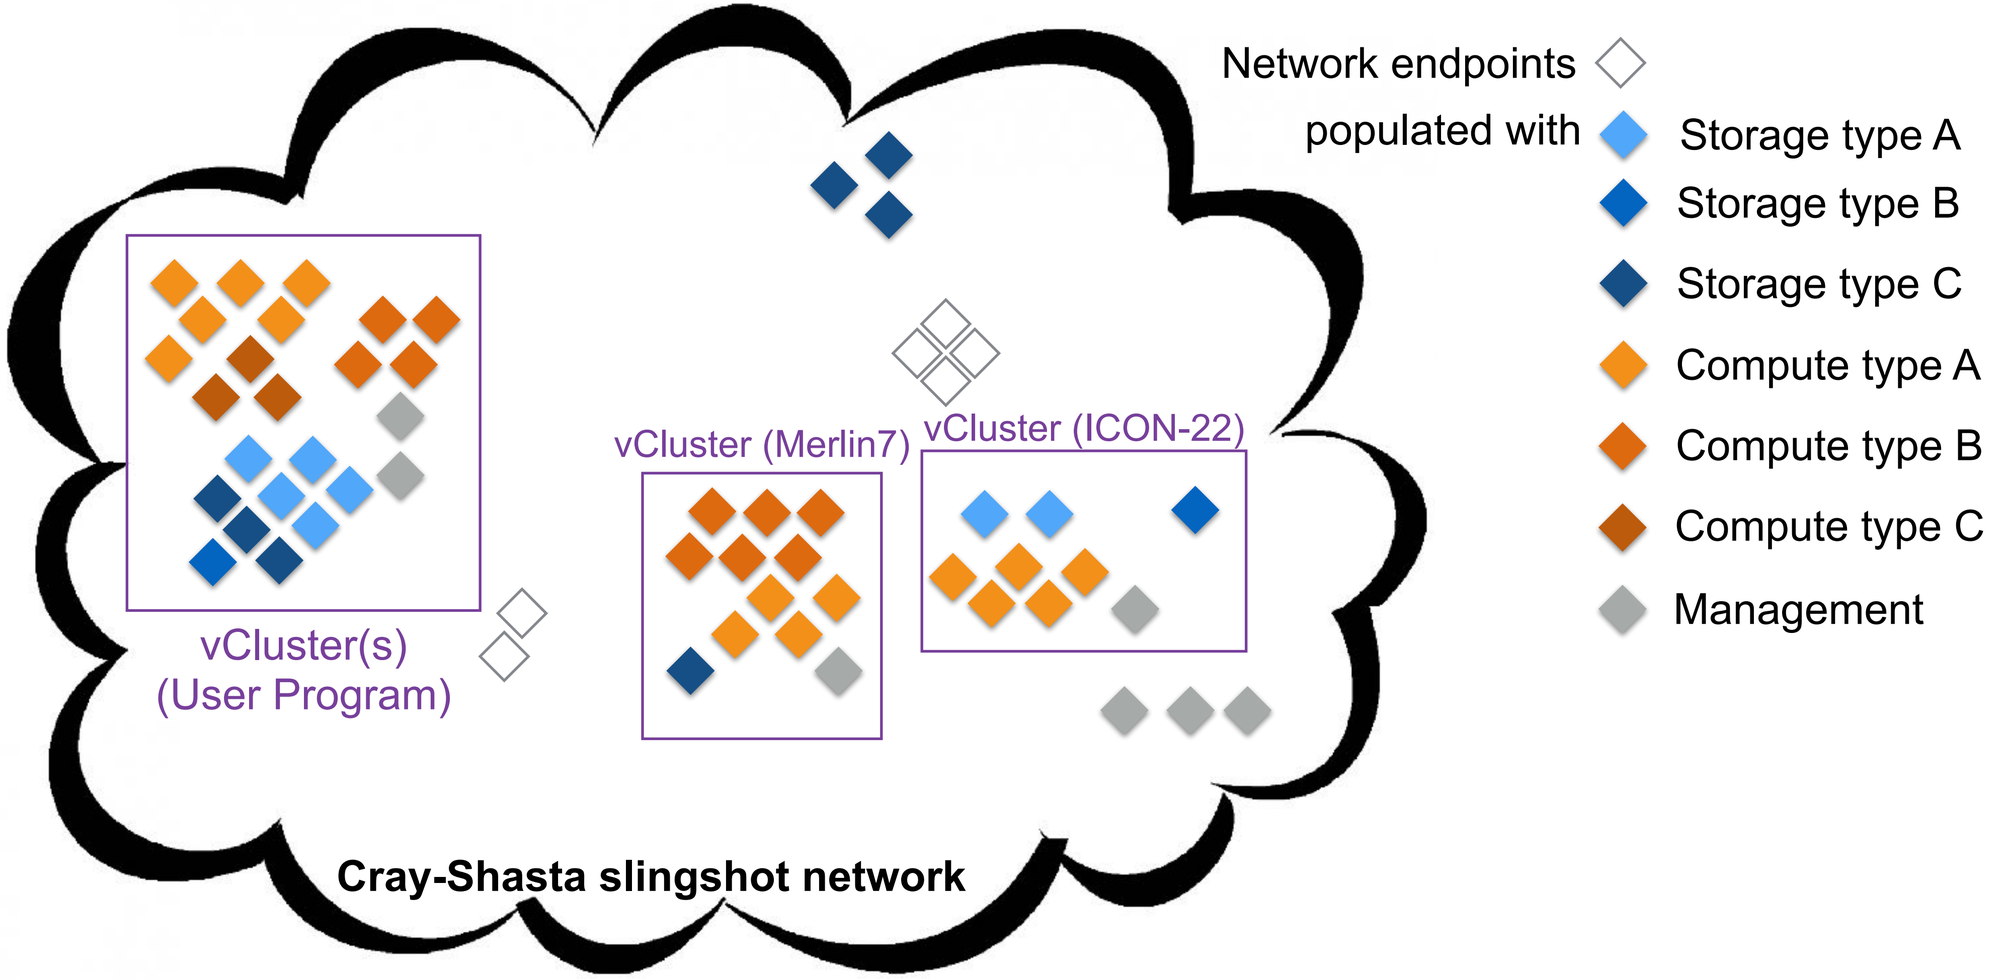

(*) currently based on HPE's Cray-Shasta system architecture

CSCS
Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich

MENU

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...

MORE

MORE SCIENCE

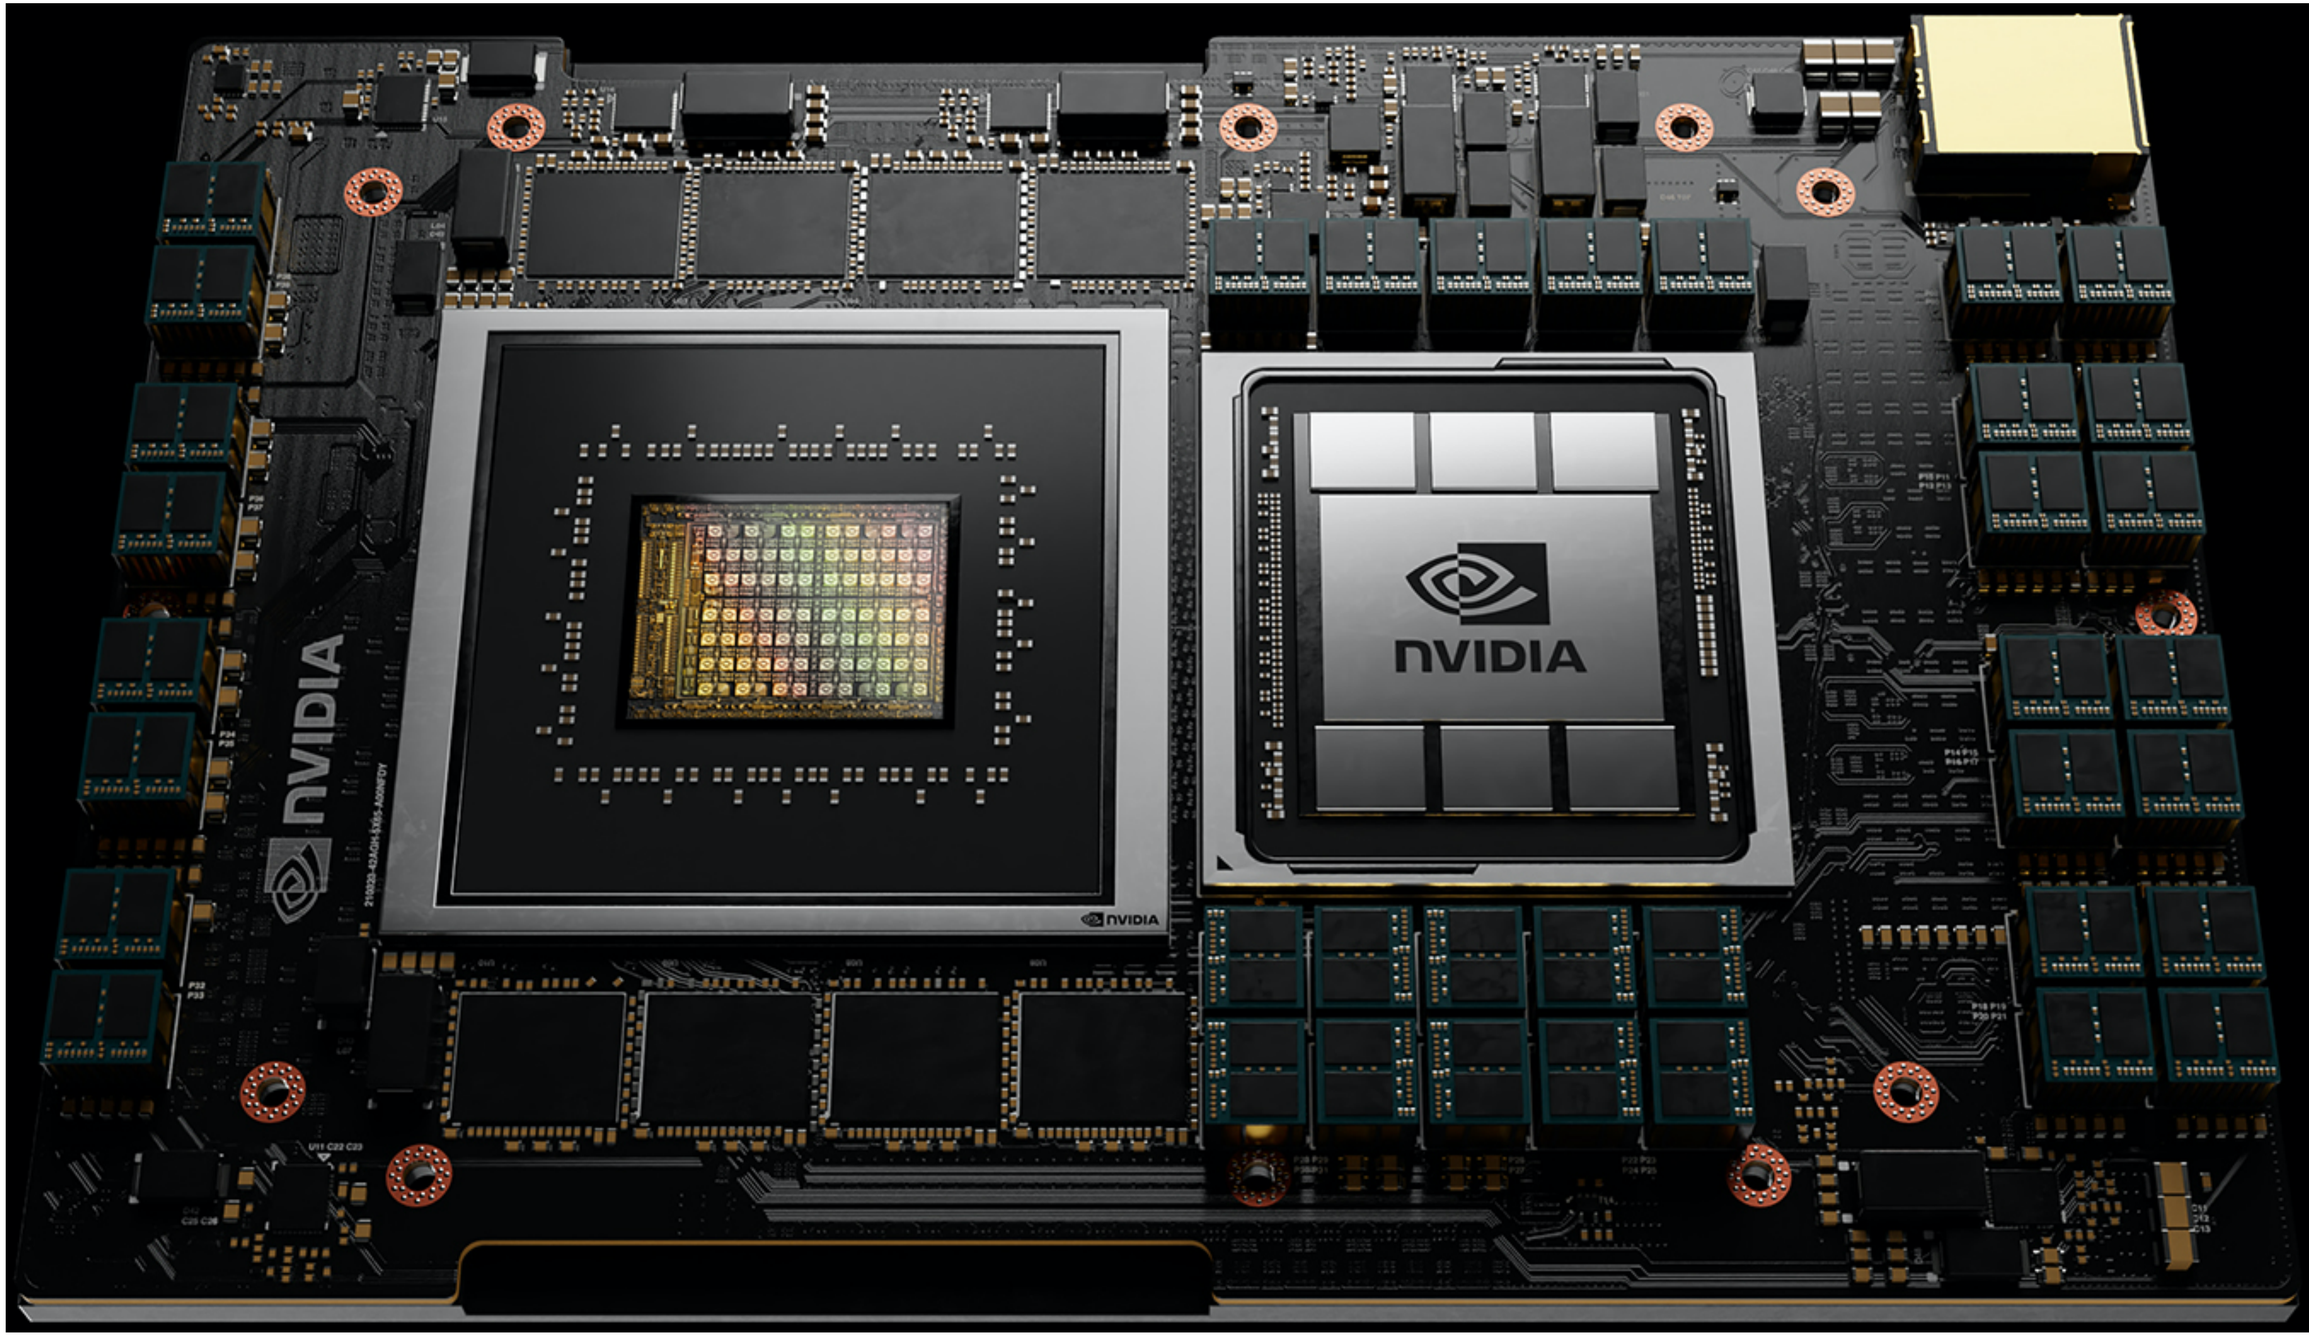
NVIDIA
NVIDIA
NVIDIA

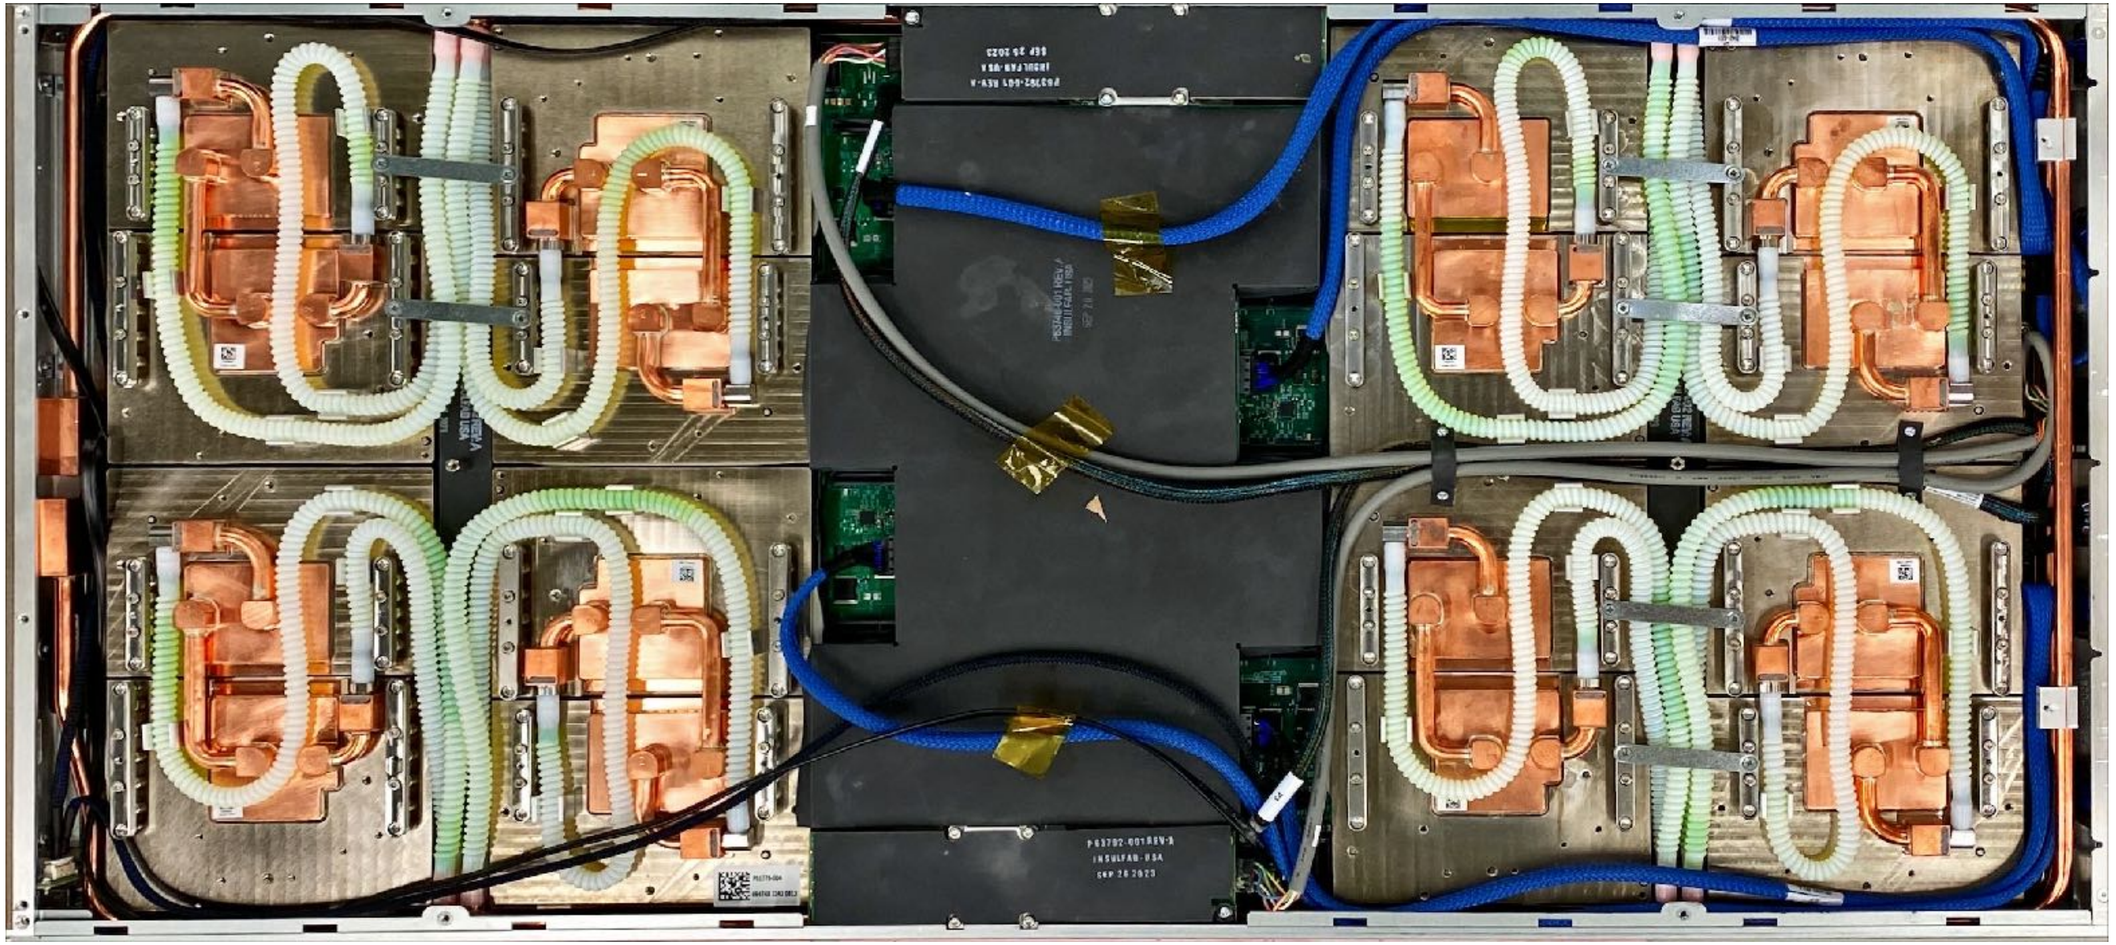

Power: ~8kW

First 64 Blanca Peak blades arrive on 01/24/2024

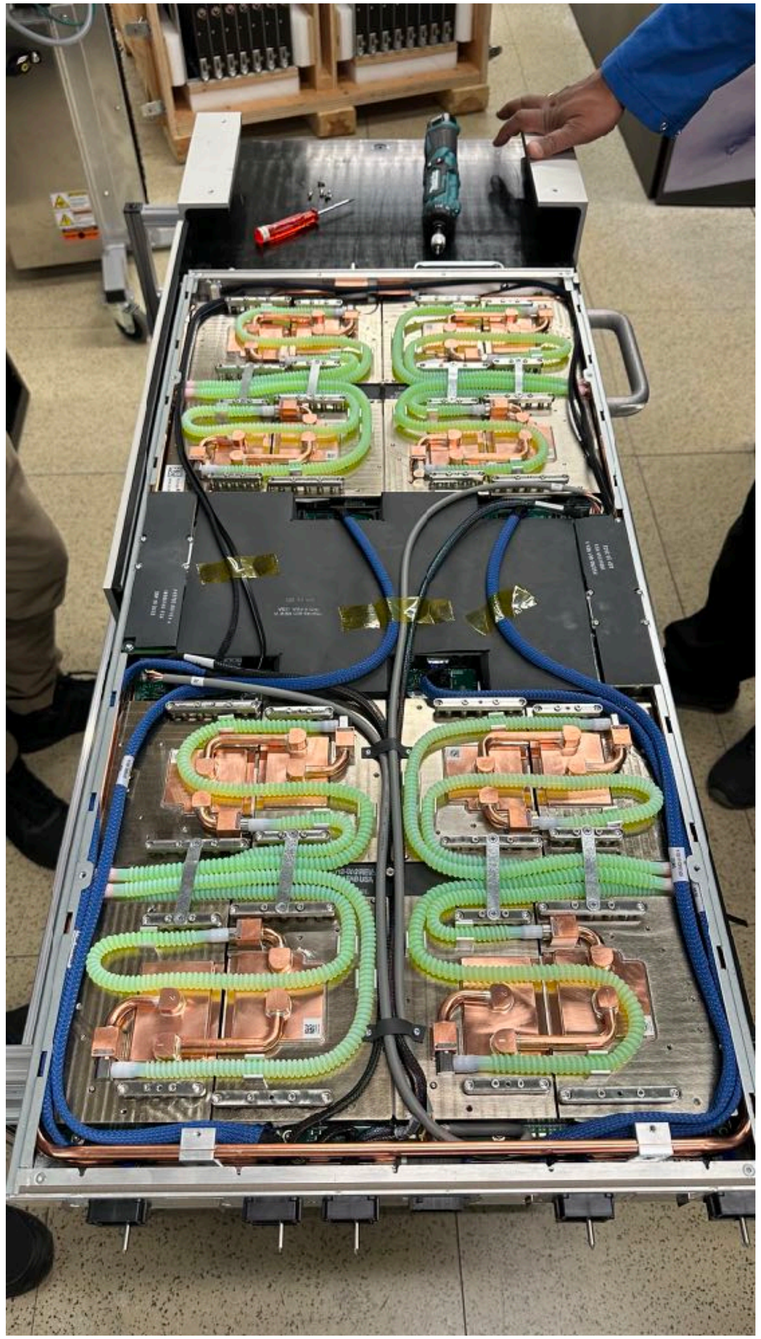

The last delivery with 128 of 1,344 blades arrives on 06/05/2024 (4+ months later)

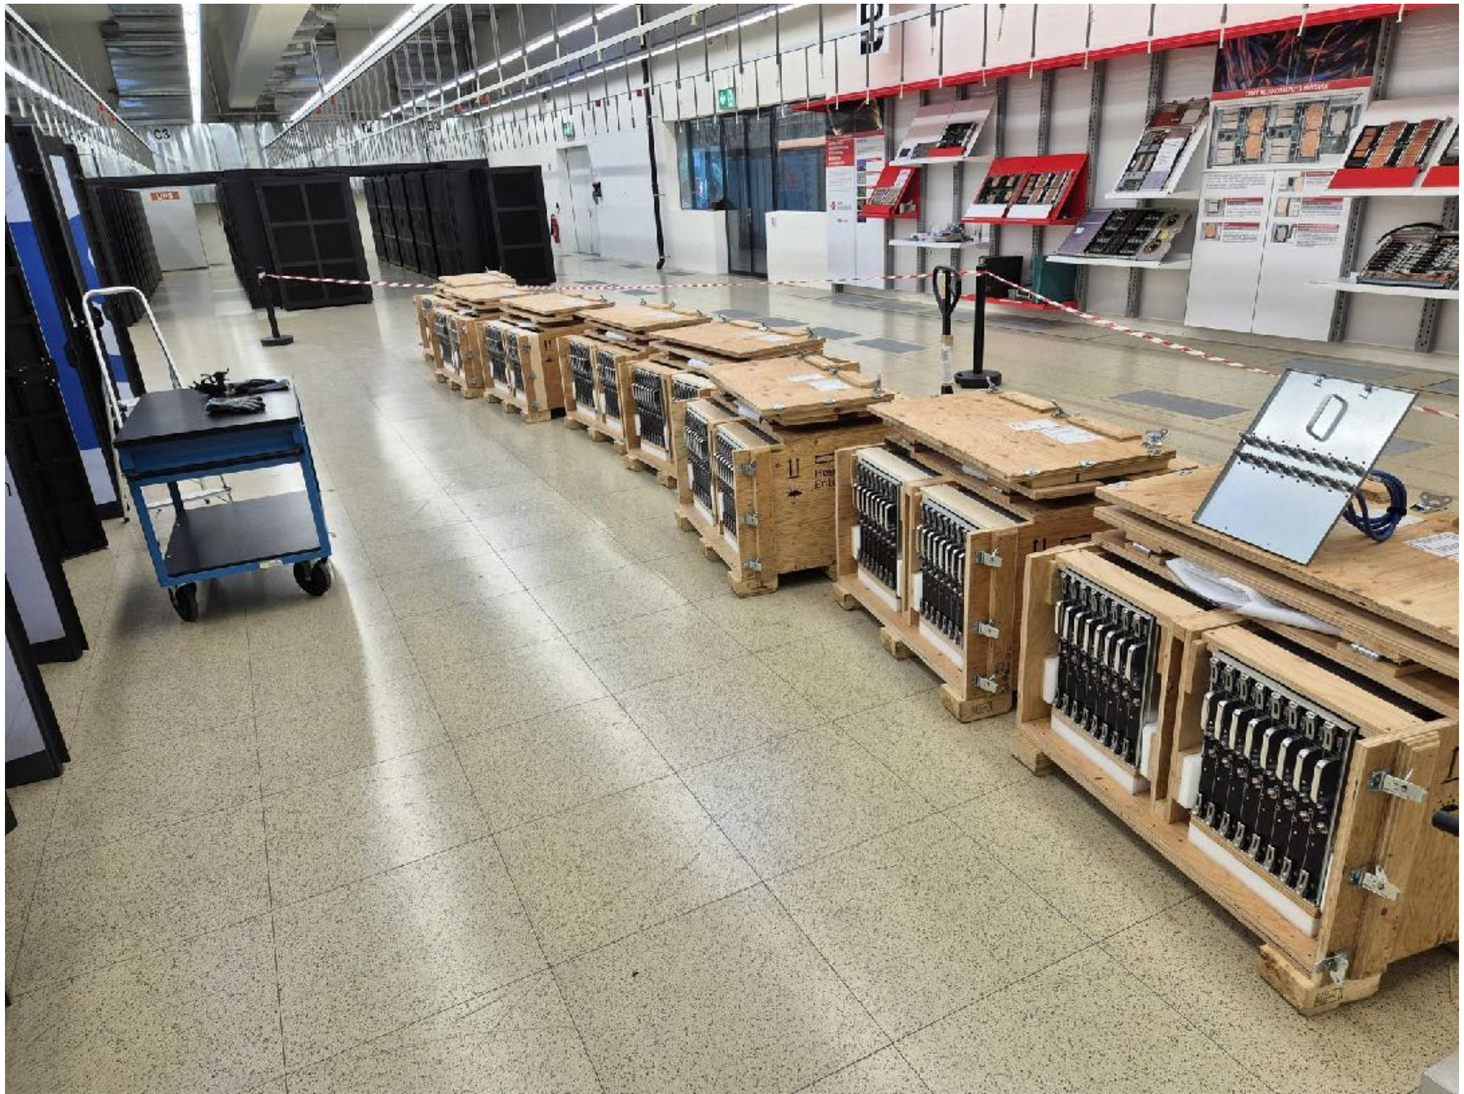

Alps with 2688 CG4 nodes @ 896 GB

or 10,752 GH200 Superchips @ 224 GB (each with H100 @ 96GB)

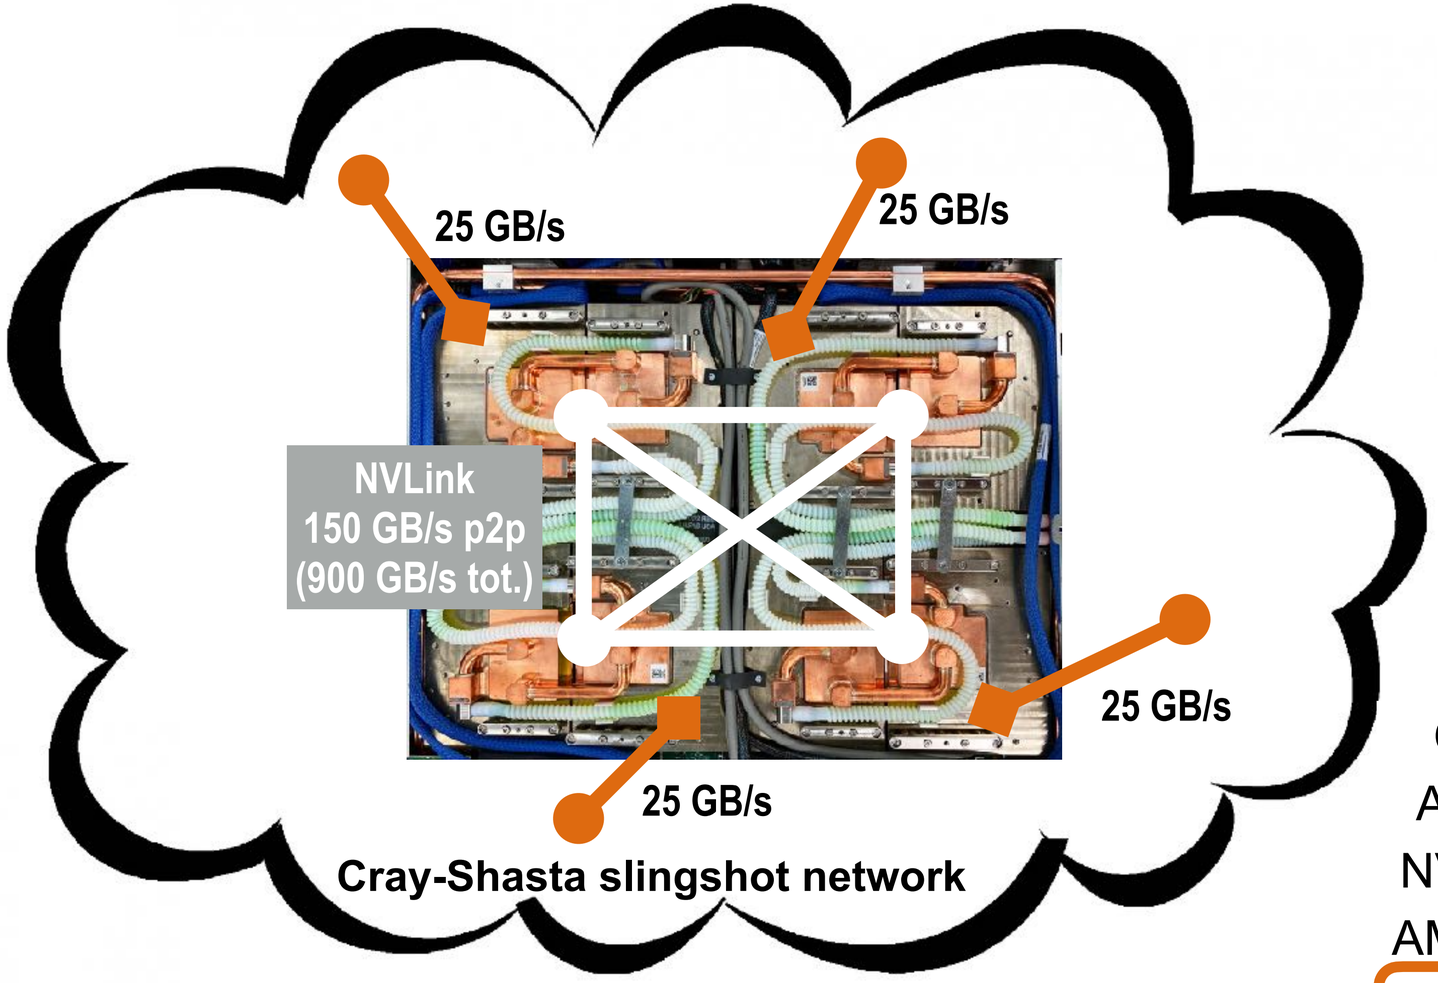

Power envelope: 10 MW

Compute node types:

AMD Rome (since 2020)

NVIDIA A100 (since 2021)

AMD Mi 250X (since 2022)

NVDIA GH200 (since 2024)

AMD Mi 300A (H2 2024)

Versatile Software Defined Cluster (vCluster)

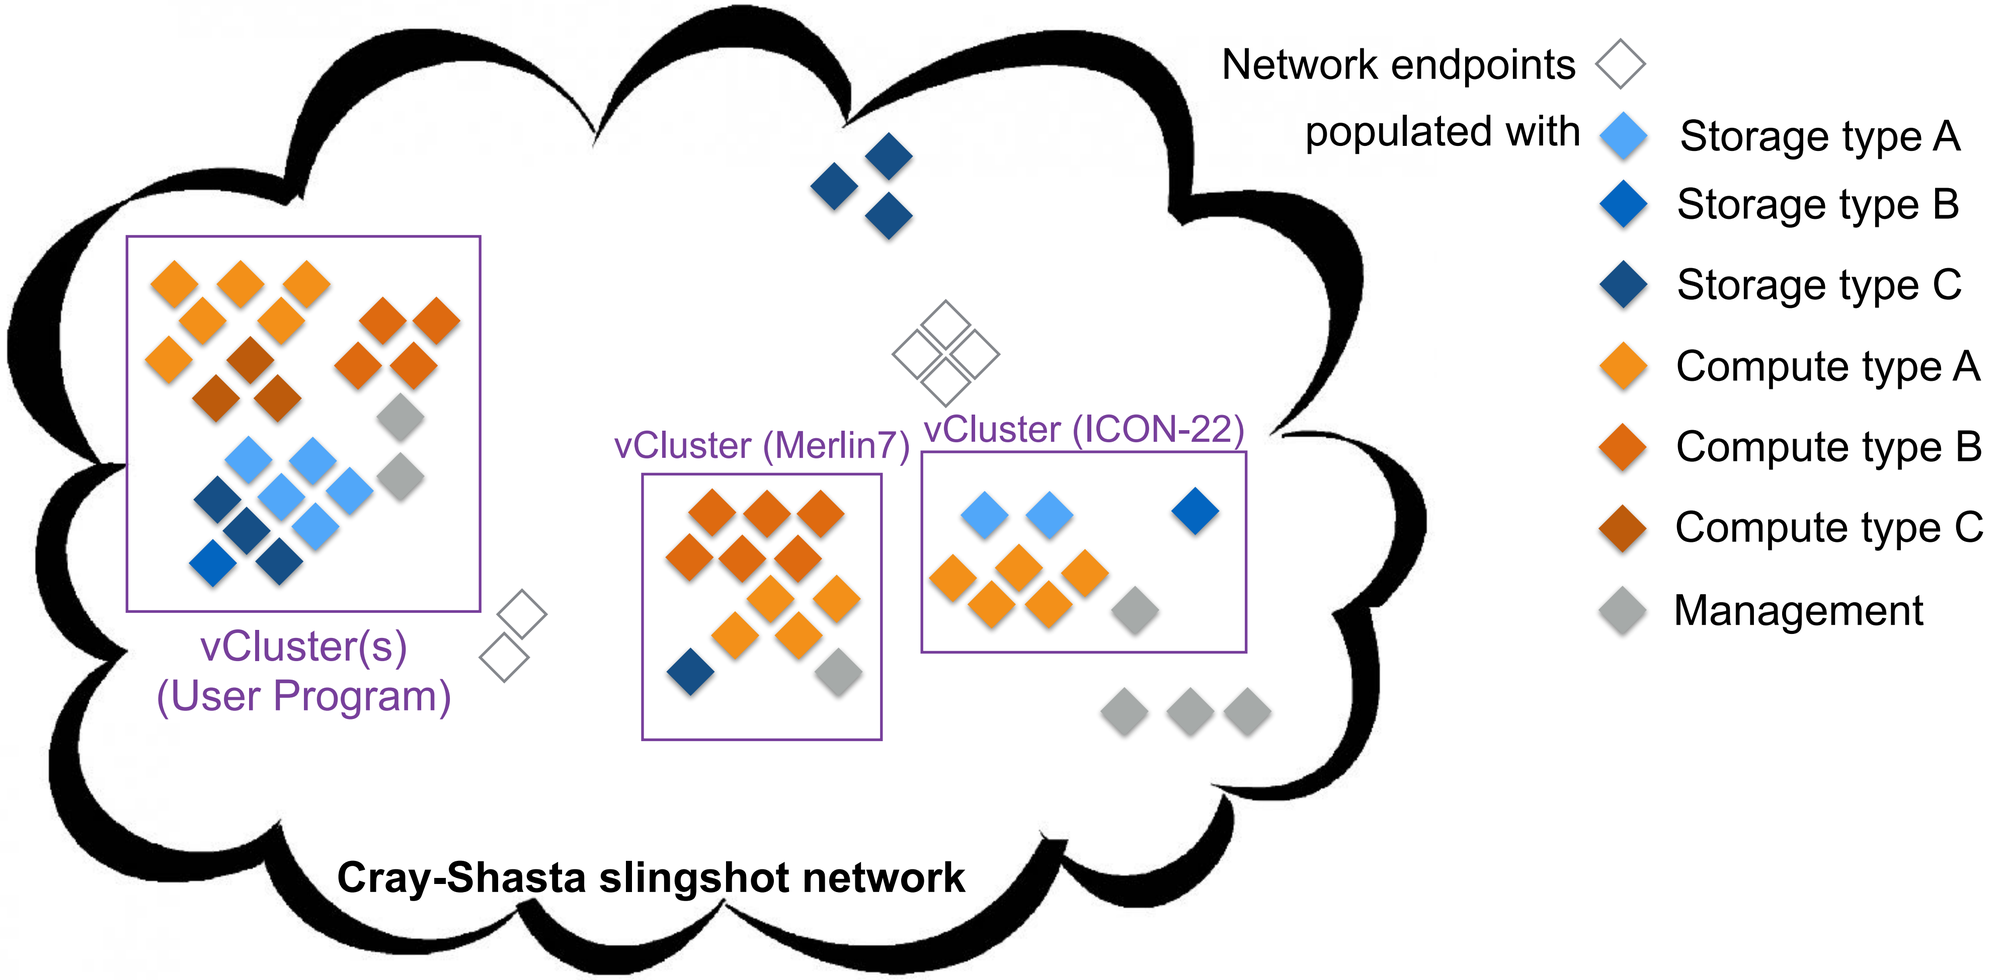

(*) currently based on HPE's Cray-Shasta system architecture

HPC & Cloud convergence

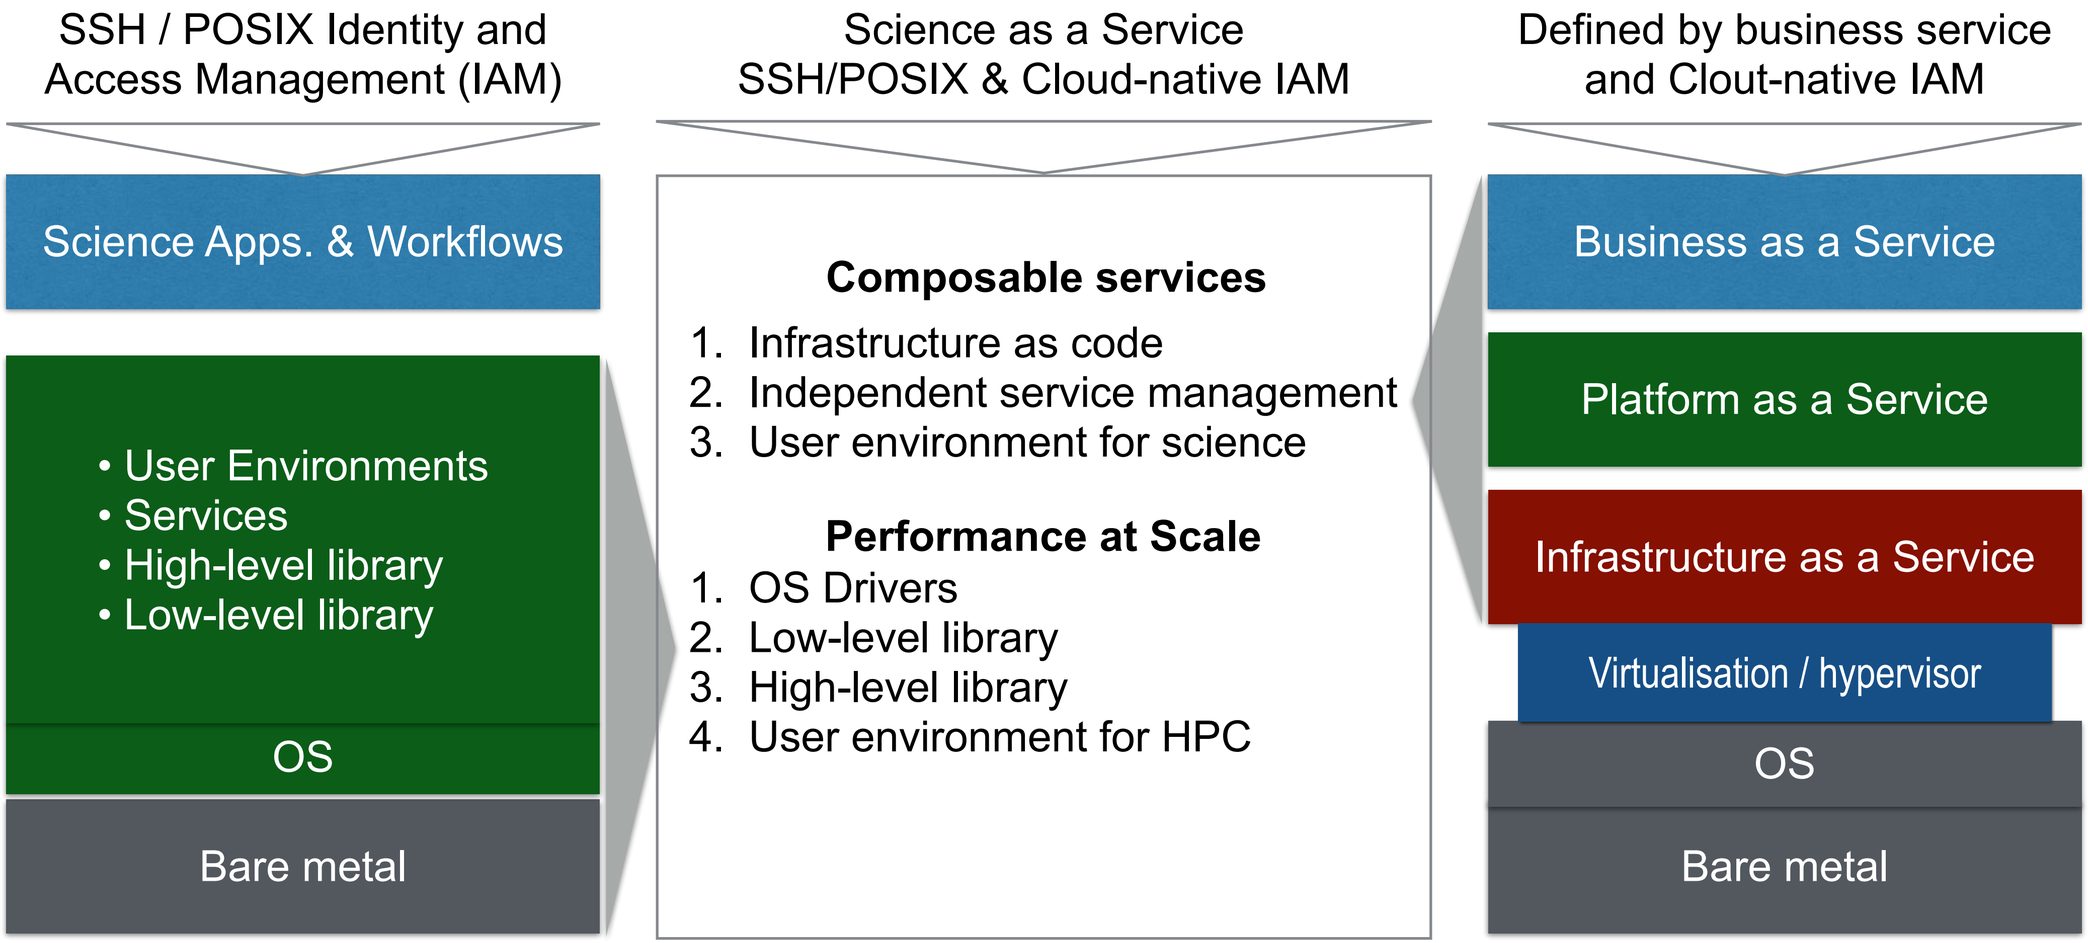

Separation of concerns

User environment

Platform engineer (typically familiar with domain)

- User env. tailored to domain (uenv)
- “Bring your own stack” (containers)
- Programmable resource access (API)

Service management

Platform tenant engineer/administrator

- Orchestration of platform services
- Self-healing / vetting of compute nodes
- Execution environments

Infrastructure provisioning

Infrastructure engineer/tenant administrator

- Interface to management infrastructure
- Resource provisioning
- Tenant separation (network segregation)

Cloud-native systems architecture

To a particular community, a platform will look like a dedicated supercomputer

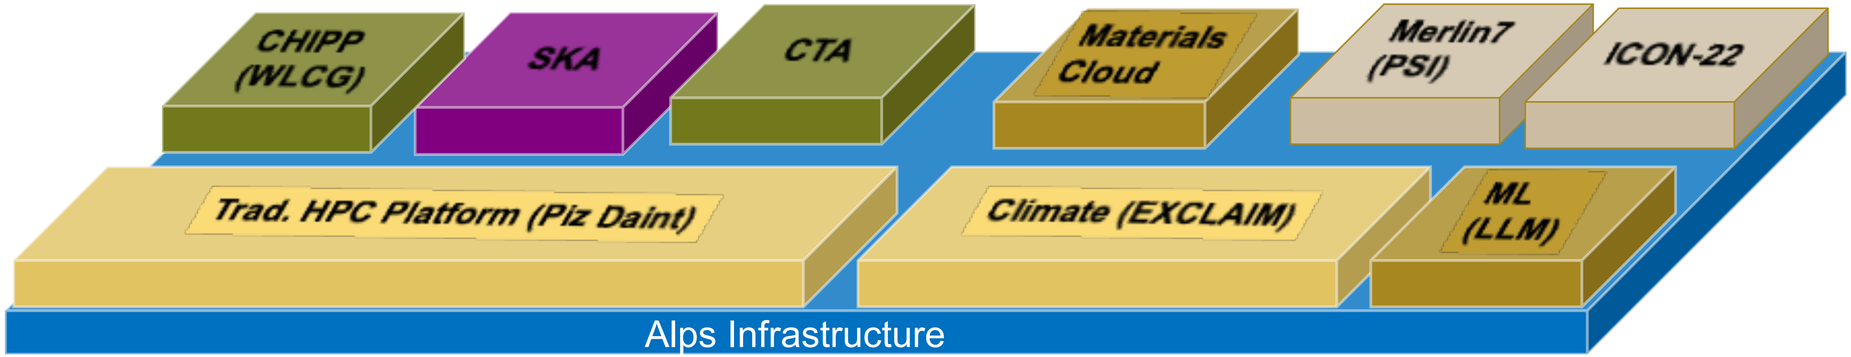

vClusters

FIRECREST

SARUS

Slingshot network

Cray System Management software (μ-service arch.)

Infrastructure partitioning

- API to control plane
- Labelling of resources
- Immutability with workflows
- Deploy a base image

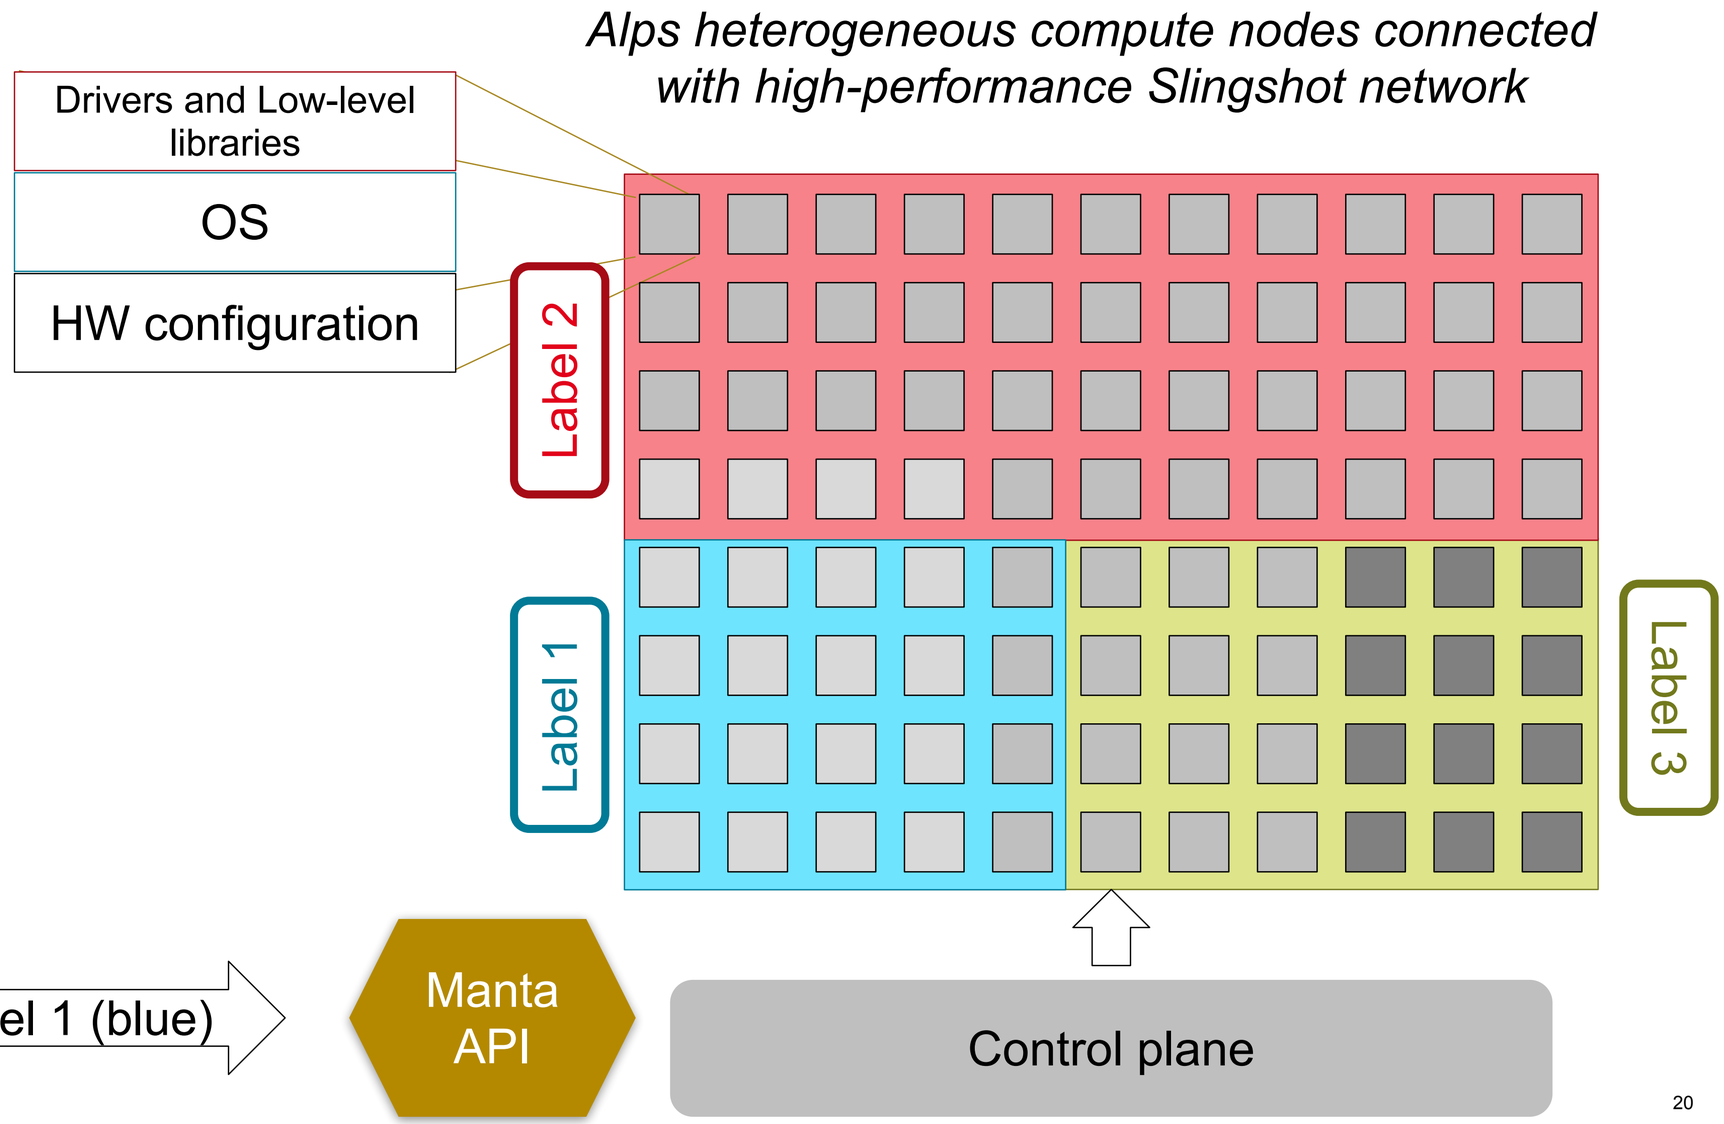

Service Management

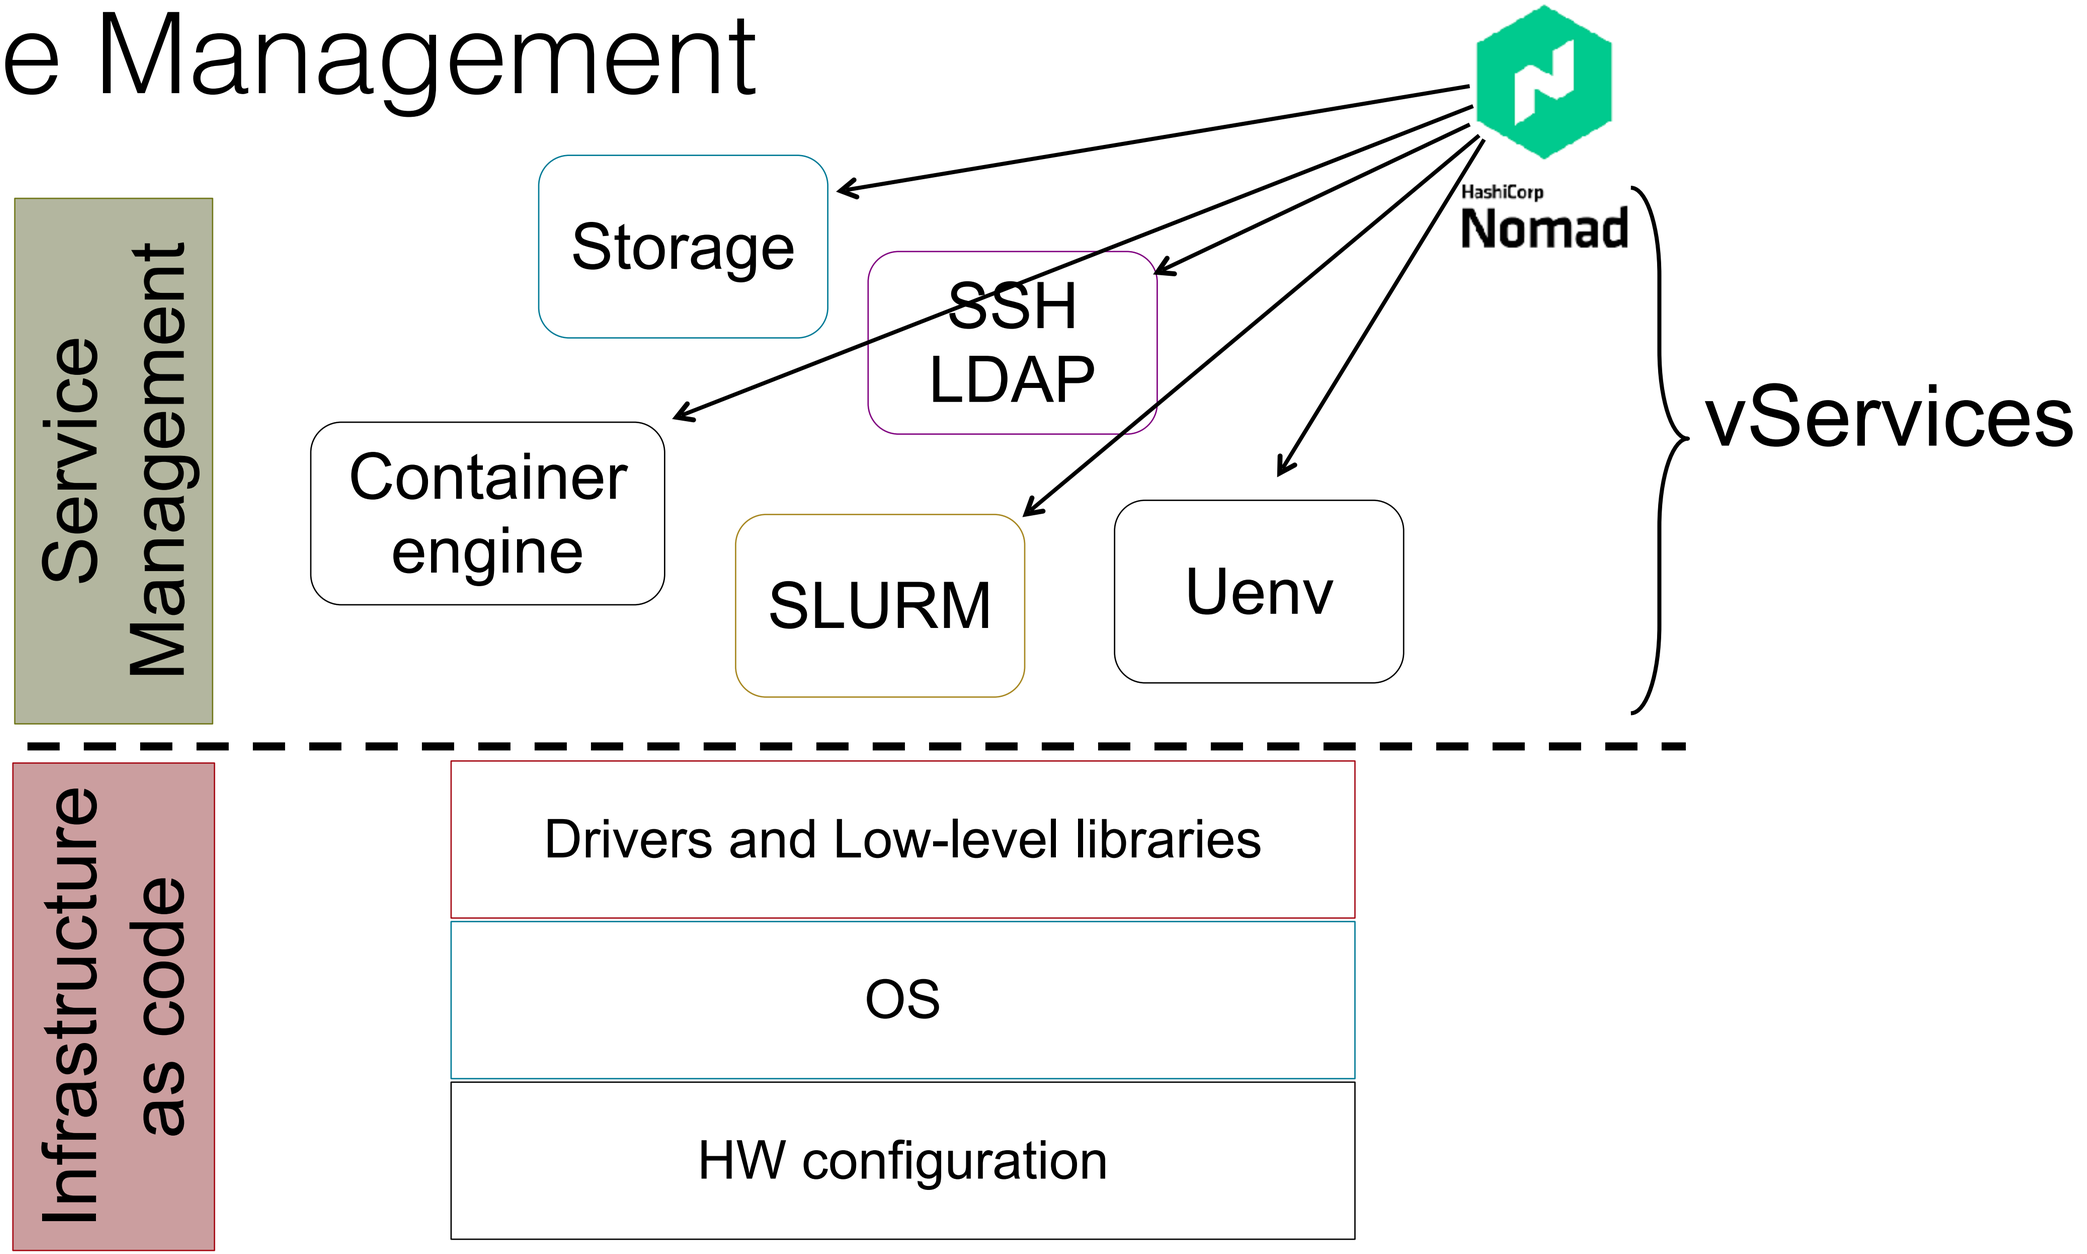

Management of services

vCluster manifest

- List of vServices
- ~10 vServices / vCluster
- vServices
  - SSH/LDAP
  - Storage
  - Uenv
  - Container engine
  - Slurm
  - Node health
  - …

vService definition in the manifest

Resource label to identify compute nodes

```
34 module "cscs-config" {
35   source       = "git@git.cscs.ch:alps-platforms/vservices/vs-cscs-config.git?ref=v1.0.1"
36   deploy       = true
37   vcluster     = module.vcluster
38   hsm_groups   = "daint"
39   ansible_vars = file("platform/config/cscs-config/vars.yml")
40   ccm_version  = "cscs-24.8.0"
41   playbook     = "site.yml"
42 }
43
44 module "storage" {
45   source            = "git@git.cscs.ch:alps-platforms/vservices/vs-storage.git?ref=v1.0.6"
46   deploy            = true
47   vcluster          = module.vcluster
48   ansible_vars      = file("platform/config/storage/vars.yml")
49   node_dependencies = ["cscs-config"]
50
51   capstor_scratch_cscs_state  = "mounted"
52   capstor_store_cscs_state    = "mounted"
53   capstor_users_cscs_state    = "mounted"
54   iopsstor_scratch_cscs_state = "mounted"
55   iopsstor_store_cscs_state   = "mounted"
56   vast_users_cscs_state       = "mounted"
57 }
```

vService code & config.

vService version

User Environment

User environment

Build your own stack

Communities to manage their own PE and applications

CI external from GitHub

Service Management

User Environment

User environment

Bring your own stack

Enable external stack with access to HPC hardware with OCI hooks

Service Management

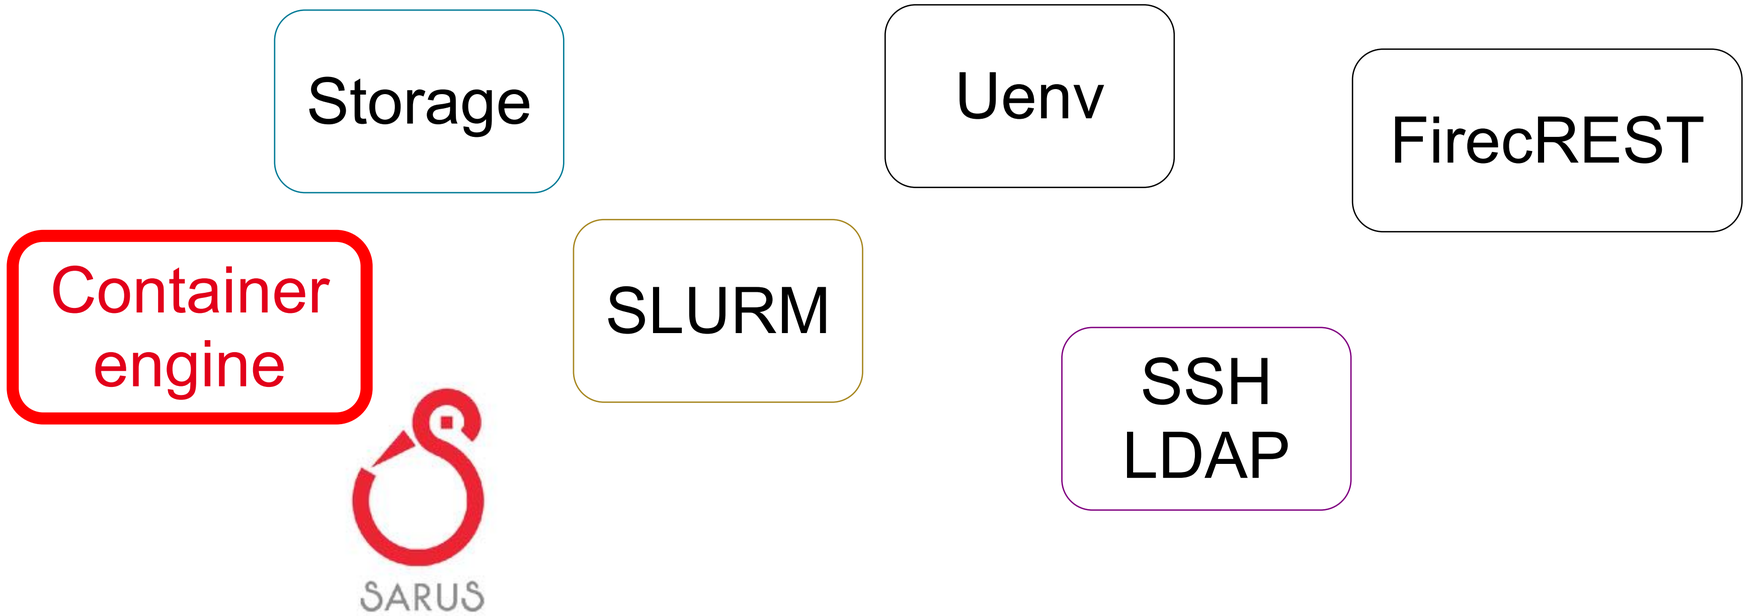

User Environment

User environment

Programmatic access

Web-facing RESTful API to move data and submit jobs

Connect to: Jupyter, AiiDA, reframe, Airflow, OOD, WebUI, CI/CD, python wrapper, **build your own**!

Service Management

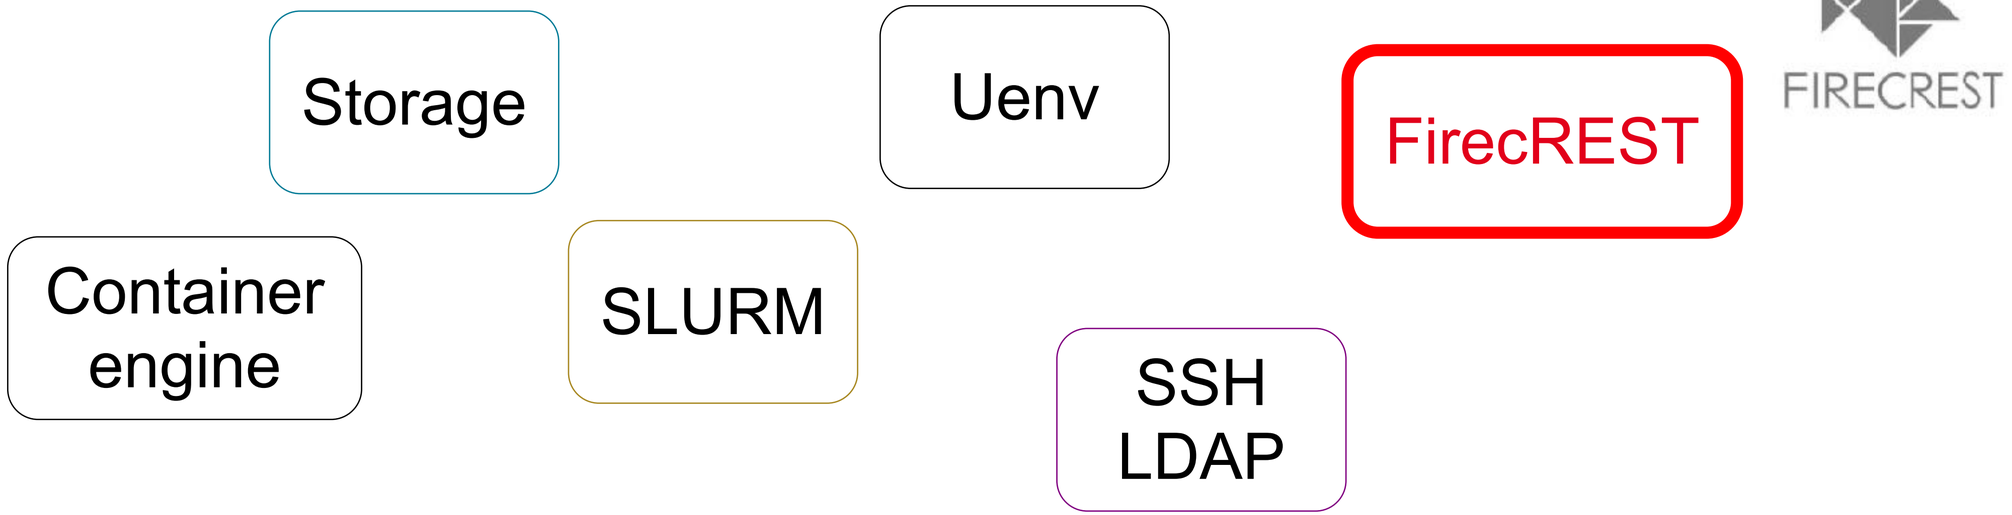

Alps: beyond traditional supercomputing

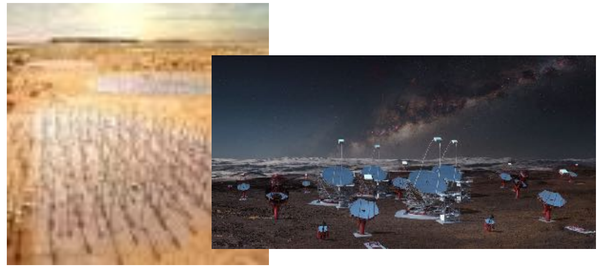

NWP
MeteoSwiss

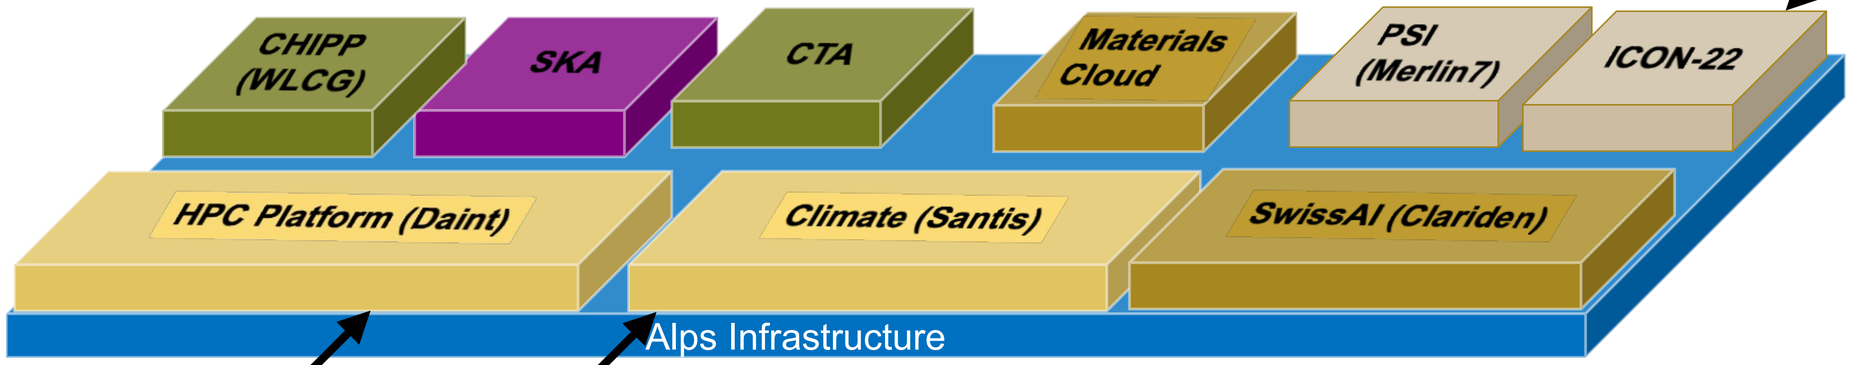

User Program (HPCN)

XCLAIM

C2SM
Center for Climate Systems Modeling

EPFL AI Center

ETH AI CENTER

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

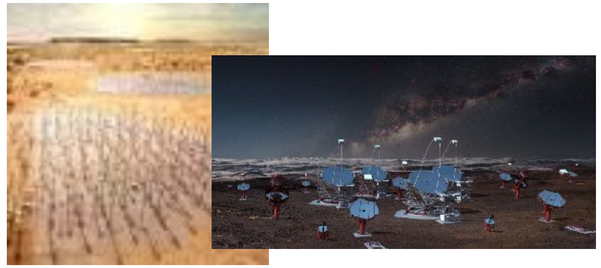

NWP
MeteoSwiss

CHIPP (WLCG)
SKA
CTA
Materials Cloud
PSI (Merlin7)
ICON-22
HPC Platform (Daint)
Climate (Santis)
SwissAI (Clariden)
Alps Infrastructure

User Program (HPCN)

XCLAIM

C2SM
Center for Climate Systems Modeling

EPFL AI Center

ETH AI CENTER

Alps: supporting Gordon Bell Prize winner SC24

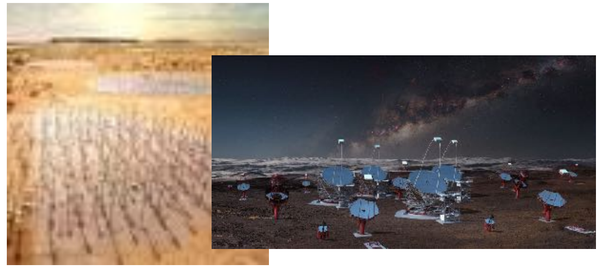

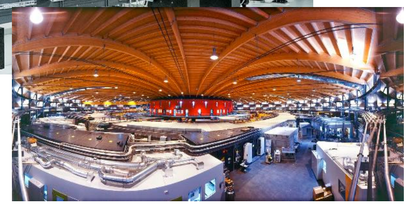

NWP
MeteoSwiss

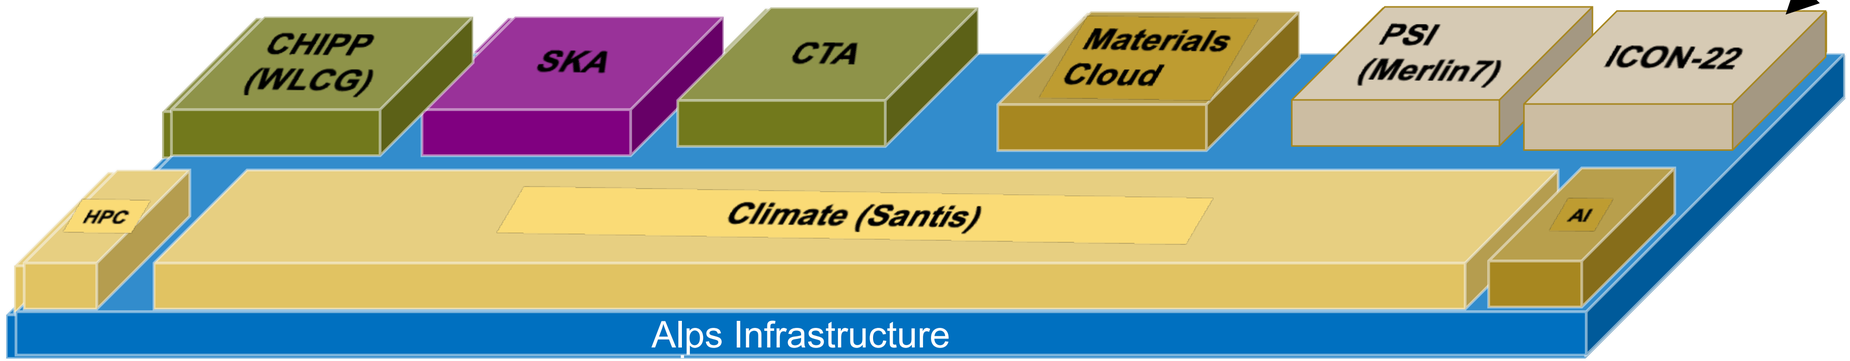

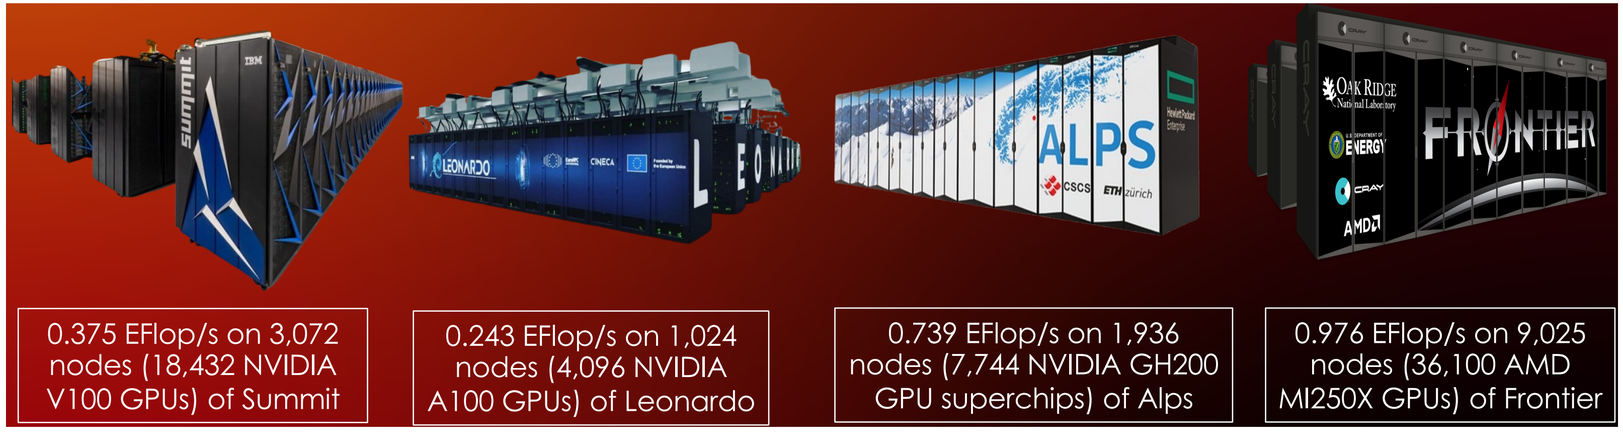

0.375 EFlop/s on 3,072 nodes (18,432 NVIDIA V100 GPUs) of Summit

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

Courtesy of David Keyes, KAUST

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

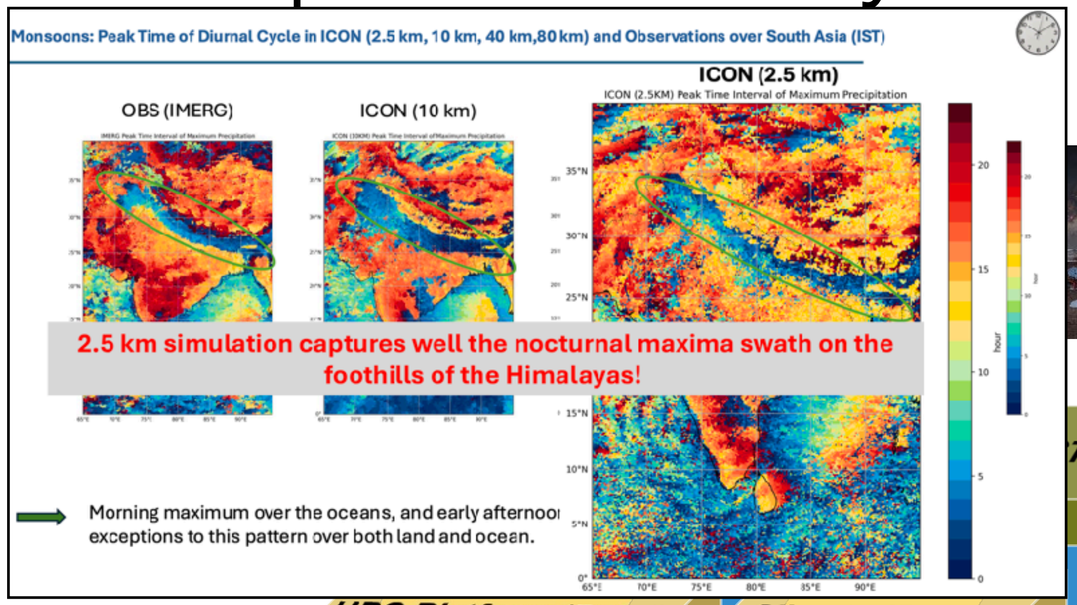

eteoSwiss

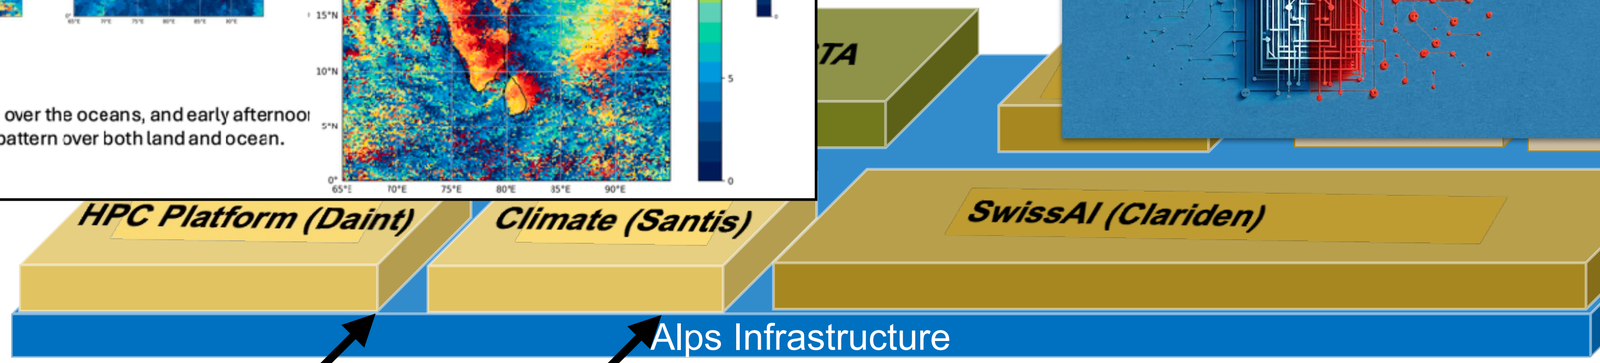

User Program (HPCN)

EXCLAIM

C2SM
Center for Climate
Systems Modeling

EPFL
AI Center

ETH AI CENTER

Apertus LLM, 70B parameters

On target for a 15T release this summer, a demonstration of capability.

Meta: Llama 70B, 1856 H100, Meta, ckpt/25 iters

Fig 1: 1856 training run with high frequency checkpointing. The first checkpoint (drop down in tps) does not have a cached save plan, and the background processing takes far longer than the rest where the cached plan is used.

https://pytorch.org/blog/6x-faster-async-checkpointing/

CSCS: Apertus 70B, 2048 GH, ckpt/250 iters

Checkpoint
(capstor: HDD)

Python GC

https://wandb.ai/ischlag/main-runs-v1/panel/vtozve3zv?nw=nwuserischlag

AI PLATFORM ON ALPS PERFORMS LIKE A DEDICATED AI SYSTEMS WITH COMPARABLE HARDWARE

Source: Joost VandeVondele, CSCS Webinar

ETH zürich

EPFL AI Center

ETH AI CENTER

Goal of Apertus LLM for science

Produce one of the most advanced open and transparent LLMs to date, keeping science at the forefront of AI research

You can consider this an AI infrastructure for science

Specific outcomes:

- A **70B LLM** trained on **15T tokens** (currently training on 4,096 GPUs)
- A **8B LLM** trained on **15T tokens** (trains on 512 GPUs)

CSCS
Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

APERTVS

EPFL ETH zürich CSCS

Inside Apertus: compliant and responsible

- Public & versioned data only, primarily sourced from CommonCrawl
- Present model memorisation of data, specifically copyrighted content
- Respect AI opt-out preferences of content authors
  - Retroactively remove content as per Jan. 2025 robots.txt crawl
  - ~10% English and ~5% non-English data removed from training sets
- Quality: Remove personal identifiable information, toxic & harmful content, and low-quality data (e.g. spam, non-human text, etc.)

Inside Apertus: multilingual design

- Over 1000 languages in training data
- Reflects the natural distribution of languages on the web globally:
  ~60% English web data
  ~40% non-English web data
- Includes code, math, and a large number of low-resourced languages, such as as Romance and Swiss-German dialects

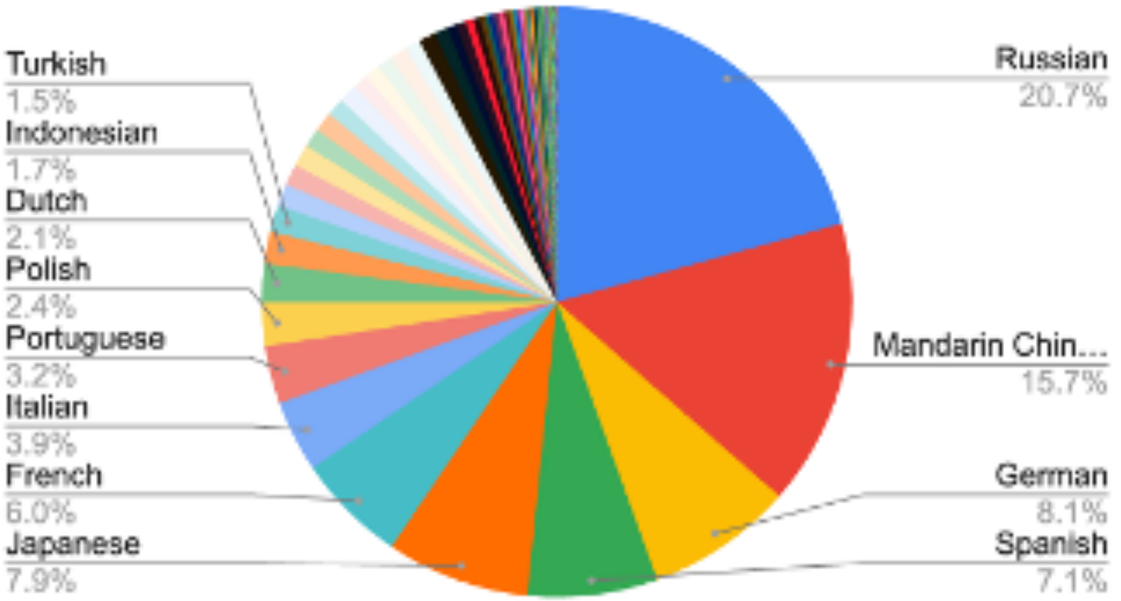

The is just the beginning of a story

- SwissAI Initiative started Aug. 2023
- Alps accepted in Jul. 2024, inaugurated in Sep. 2024
- First version of Apertus released on 2. Sep. 2025
- We will continue with Alps for at least 1-2 more years

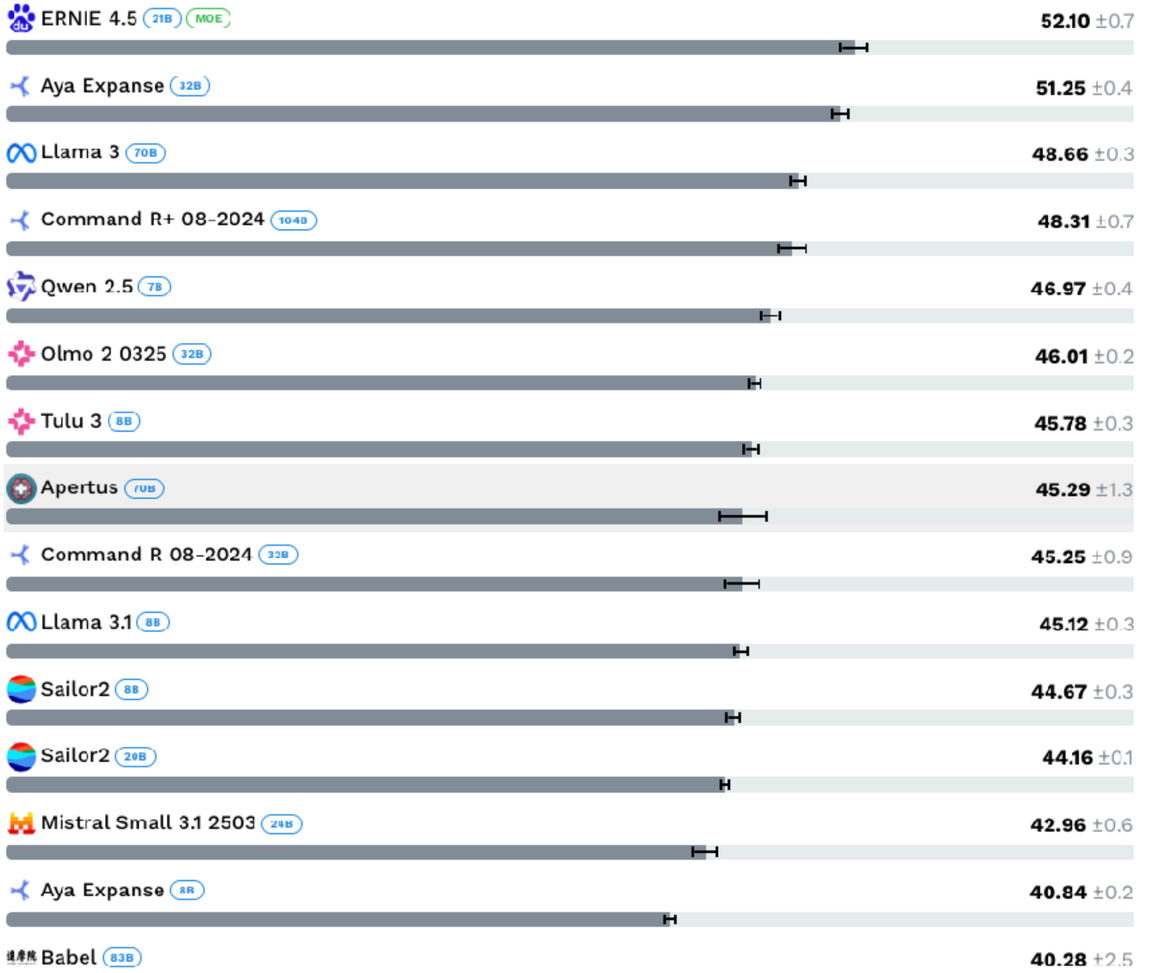

The model lifecycle

99% of resource

Foundation Models

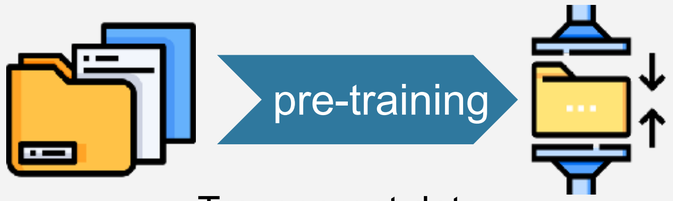

Transparent data,
energy-efficient algorithms

Infrastructure

Human-AI Alignment

Align to Swiss values for adoption in
society and economy (i.e. SME)
(in sync. with the principles of science)

Communal Red-teaming

Measure & mitigate bias, truthfulness,
privacy, misinformation, societal risks

Deployment & Feedback

Release open models, data, source

Soon Alps will be too small for pre-training LLMs


ALPS


ALPS


ALPS

Alps data hyper-cube at ECMW in Bologna

IFS running @ ECMF in Bologna

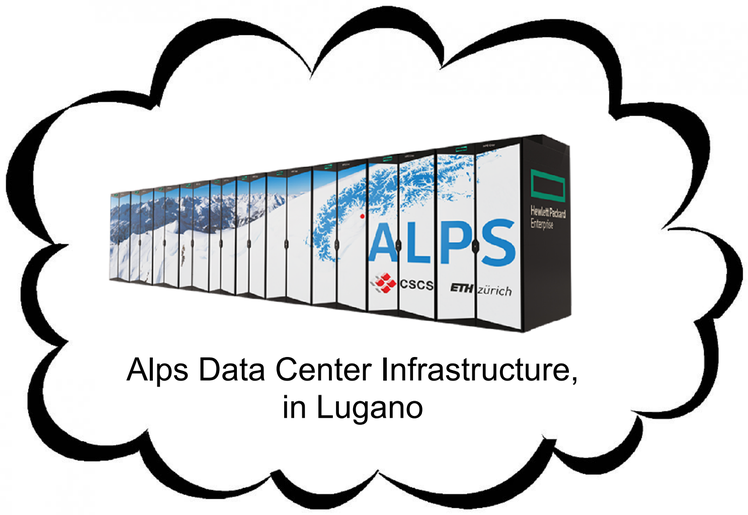

Alps Data Center Infrastructure,
in Lugano

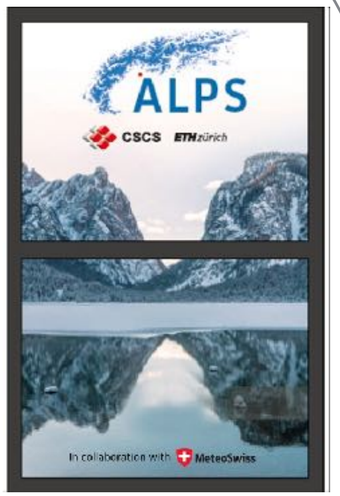

Alps Data Hypercube
in Bologna

ECMWF MARS archive in Bologna

CSCS Electricity Consumption 2008-2024

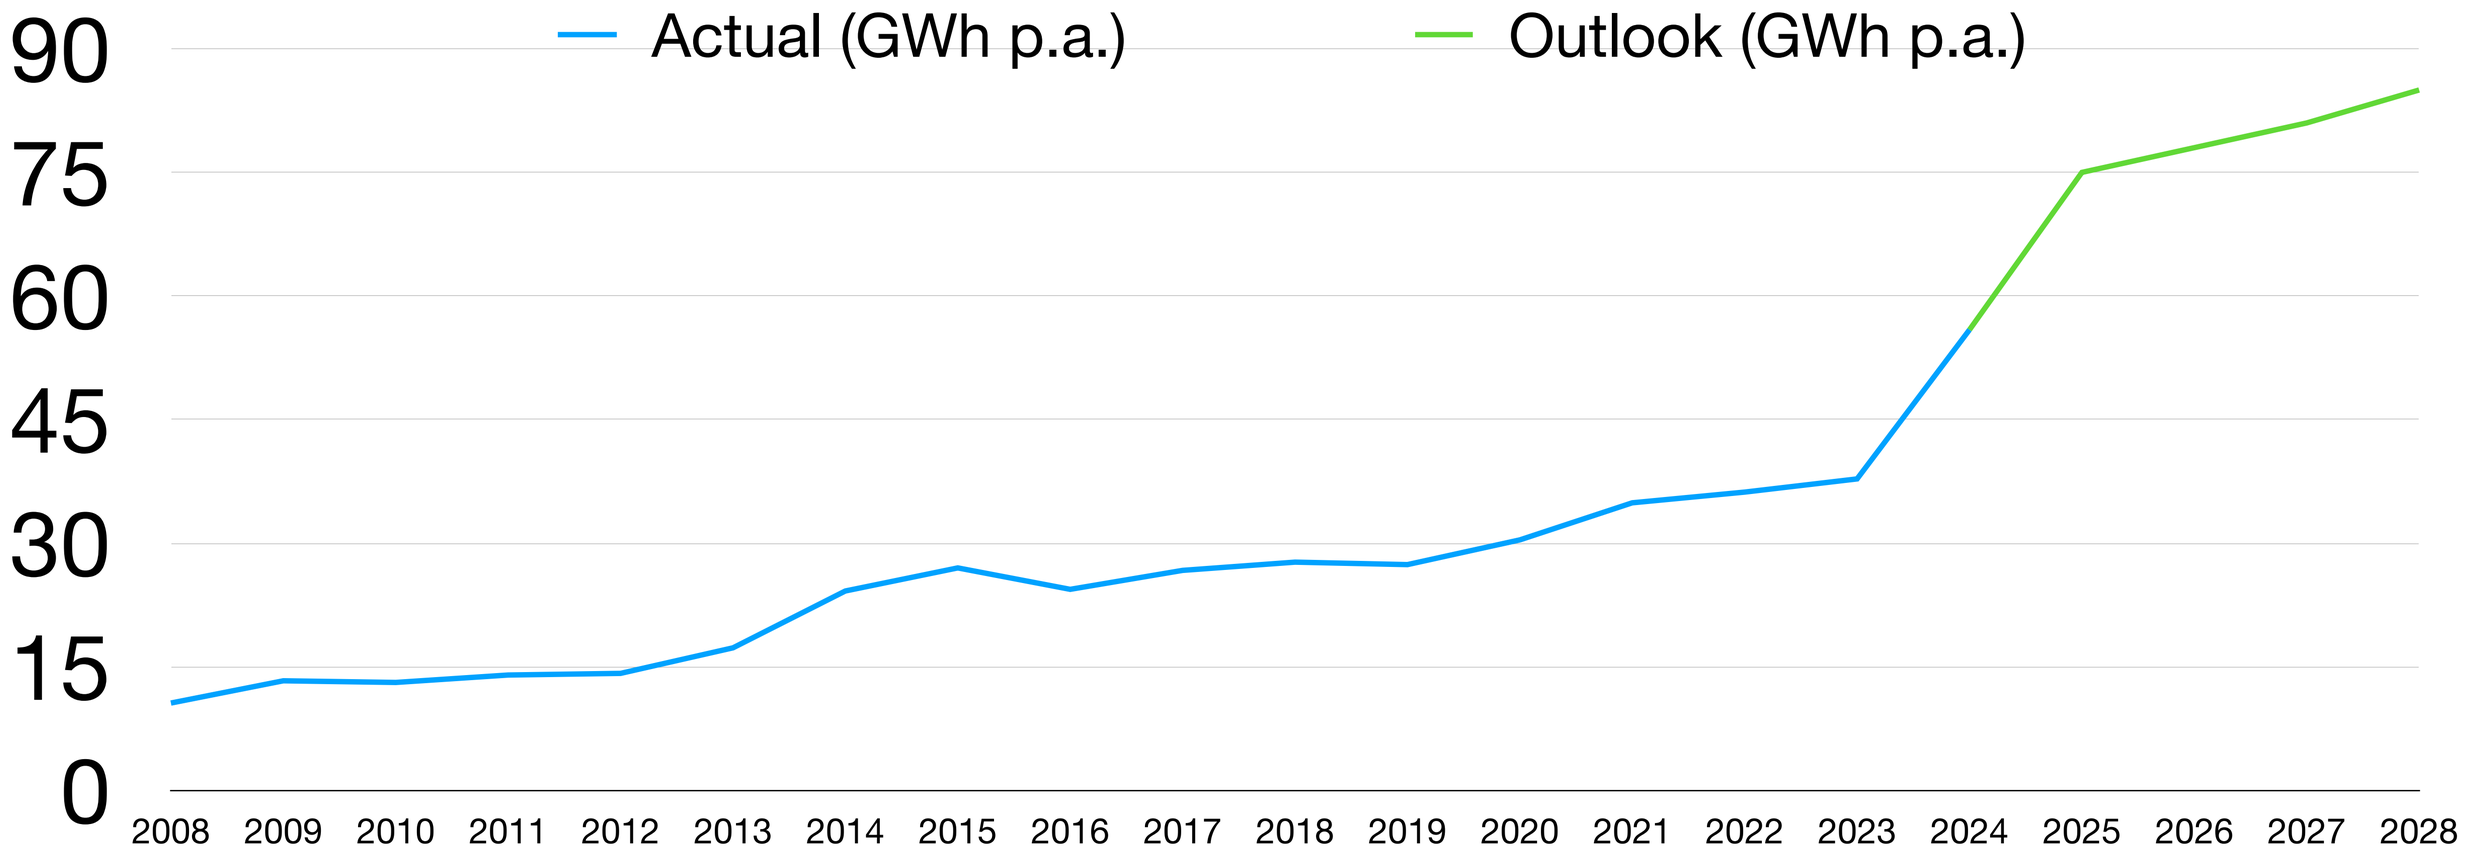

LUMI is a consortium that has deployed the first EuroHPC pre-exascale supercomputer

Scaling from 10s to 100s of MW

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