CSCS
Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich

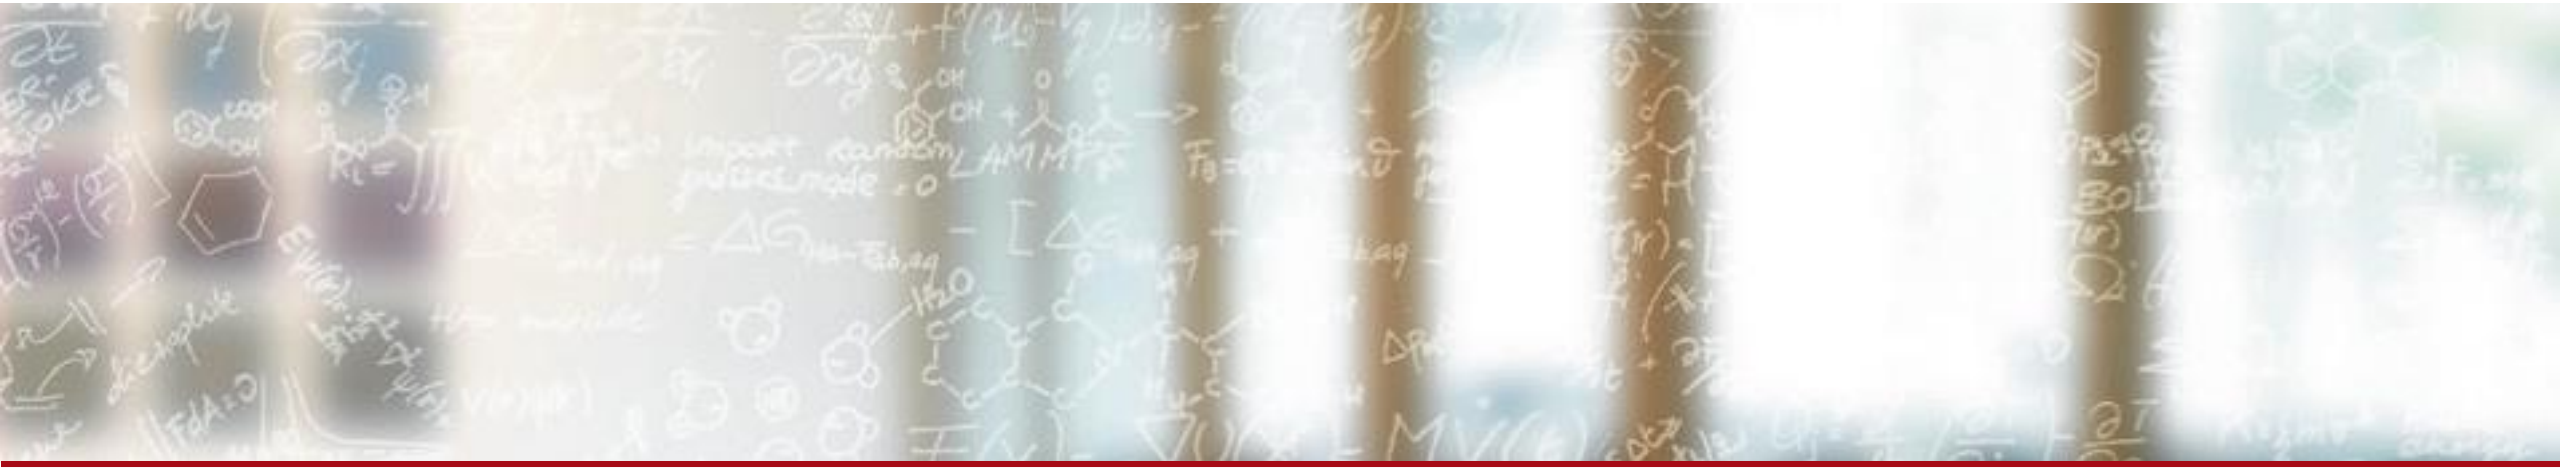

Infrastructure as a Service (IaaS) with Strong Tenant Separation on a Supercomputer

CUG25

Maxime Martinasso, CSCS

May 2025

Motivation

- Paul Scherrer Institute operates advanced experimental facilities – Third-generation Synchrotron
  - Beamlines generate **terabytes of data in real-time**.
  - This requires **immediate post-processing** and storage for experimental workflows
- Synchrotron upgrade (SLS 2.0) increases data output **by up to 40x**.
  - PSI's **on-prem data center has reached its limits**.
  - Building a **new data center is costly** and time-consuming
- PSI-CSCS history of good collaborations
- Infrastructure-as-a-Service (IaaS) on Alps
  - Cloud is too expensive – 100's of nodes + GH200

PAUL SCHERRER INSTITUT

PSI

Alps heterogeneous hardware

- Alps is an HPE Cray EX supercomputer being our new flagship infrastructure

  - 1024 AMD Rome-7742 nodes 256/512GB
  - 144 Nvidia A100 GPU nodes
  - 24 AMD MI250x GPU nodes (LUMI1 type)
  - 128 AMD MI300A GPU nodes
  - ***2688 Grace-Hopper nodes, 10752 GH200, ~430 PF***
  - Slingshot network (200 Gbps injection)
  - Two availability zones (HA, non-HA)
  - 100% liquid cooled
  - 100+10 PiB HDD
  - 5+1 PiB SSD (RAID10)
  - 100s of PiB tape library
  - ~10 MW (envelope for power and cooling)

Water cooled blades

Objectives

- PSI workflows and requirements
  - PSI needs to manage their own resources and users
  - Be independent but don't affect other users running on Alps
  - Security and network segregation

- CSCS IaaS should be automated and generic
  - Investigate multi-tenancy capability
    - Compute, Network, Storage
  - Explore new scenarios such as dynamically provisioned IaaS

Technology needs

- System management capable of multi-tenancy
  - Label resources using Cray System Management
  - Act only on allocated resources
  - Simplified system management actions
- Network segregation
  - Slingshot, Ethernet technology
  - Security policies
- Storage multitenancy
  - Dimensioned storage with associated performance
- Service management
  - Flexibility and independence
  - Specific services not known by CSCS
  - PSI manages node image, monitoring and production

Technology: Cray System Management

- Microservice architecture
  - System management microservices for HPE EX
  - Kubernetes
  - RESTful API support
- Microservices
  - Version Control Service (VCS)
  - Hardware State Manager (HSM)
  - Configuration Framework Service (CFS)
  - BSS, CAPMC, BOS, IMS…
- Multitenancy
  - Resources labelled into groups (HSM group)
  - Microservice acts on grouped resources
- Tenant and Partition Management System (TAPMS)

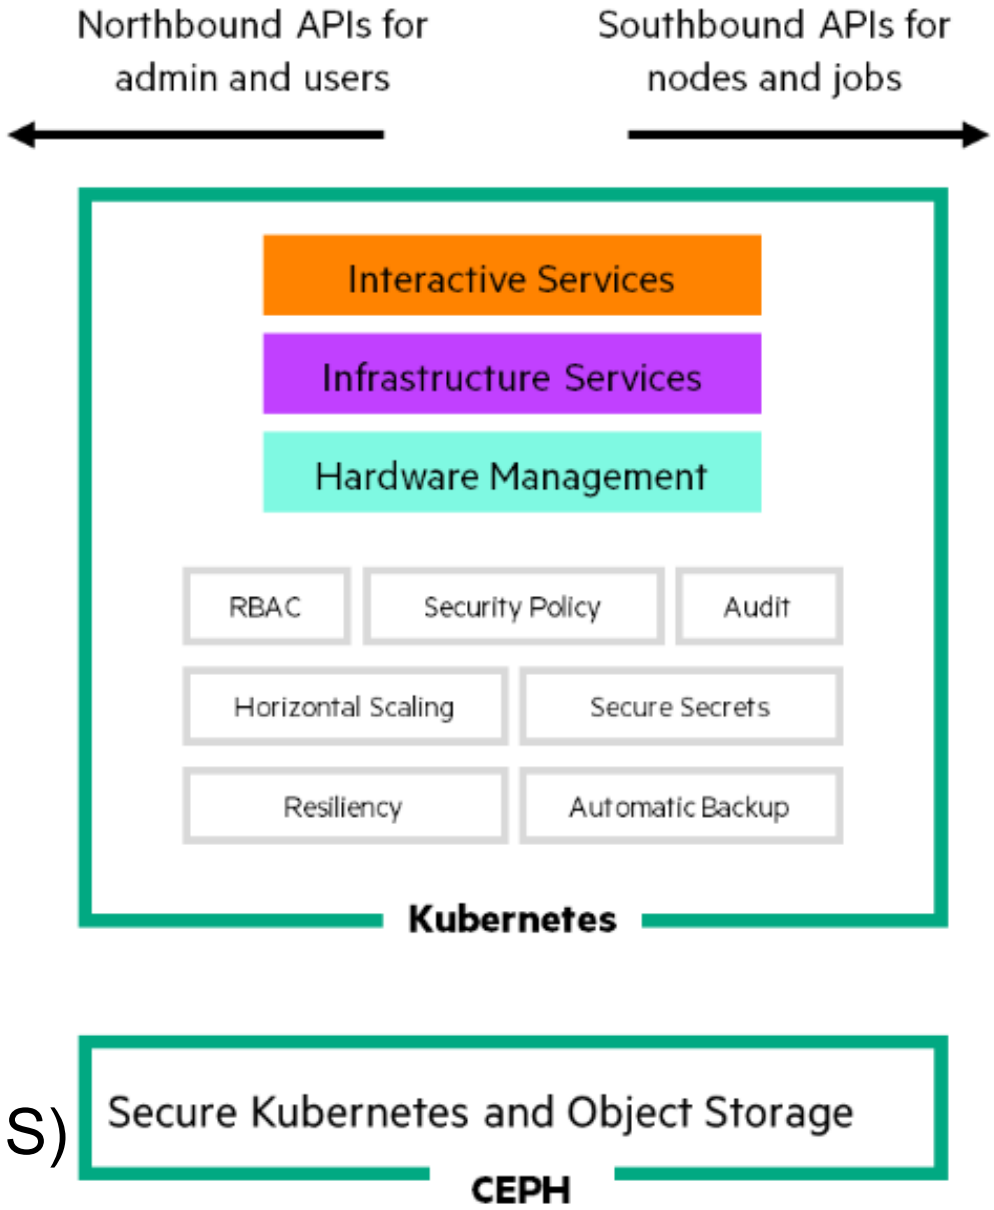

Source: HPE documentation

Technology: Manta interface and RBAC

- CSCS cannot grant CSM access to PSI
  - Security: actions can affect all Alps
  - Practicality: too complex to learn

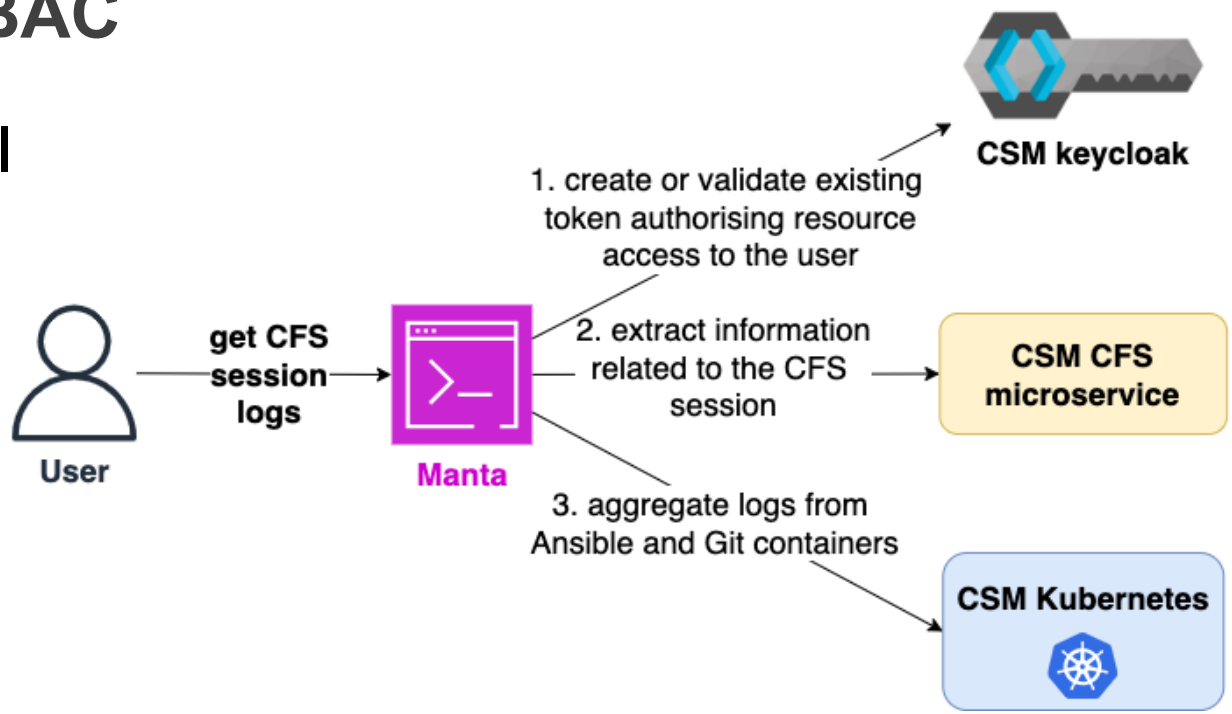

- Manta interface
  - API + CLI frontend
  - Intuitive interface, co-designed with PSI
    - Cluster configuration, power management, nodes console access, data querying, CFS logs, multi-site
- Role-based access control (RBAC)
  - Manta requires a CSM Keycloak role = HSM group
  - Each user action checks the authentication token to determine role-based permissions
  - Ensure that role matches the HSM group of resources acted upon

Technology: Network

- PSI ←→CSCS
- SWITCH network
- 2 dedicated links
  - 100Gb/s each
  - 4-5 ms latency

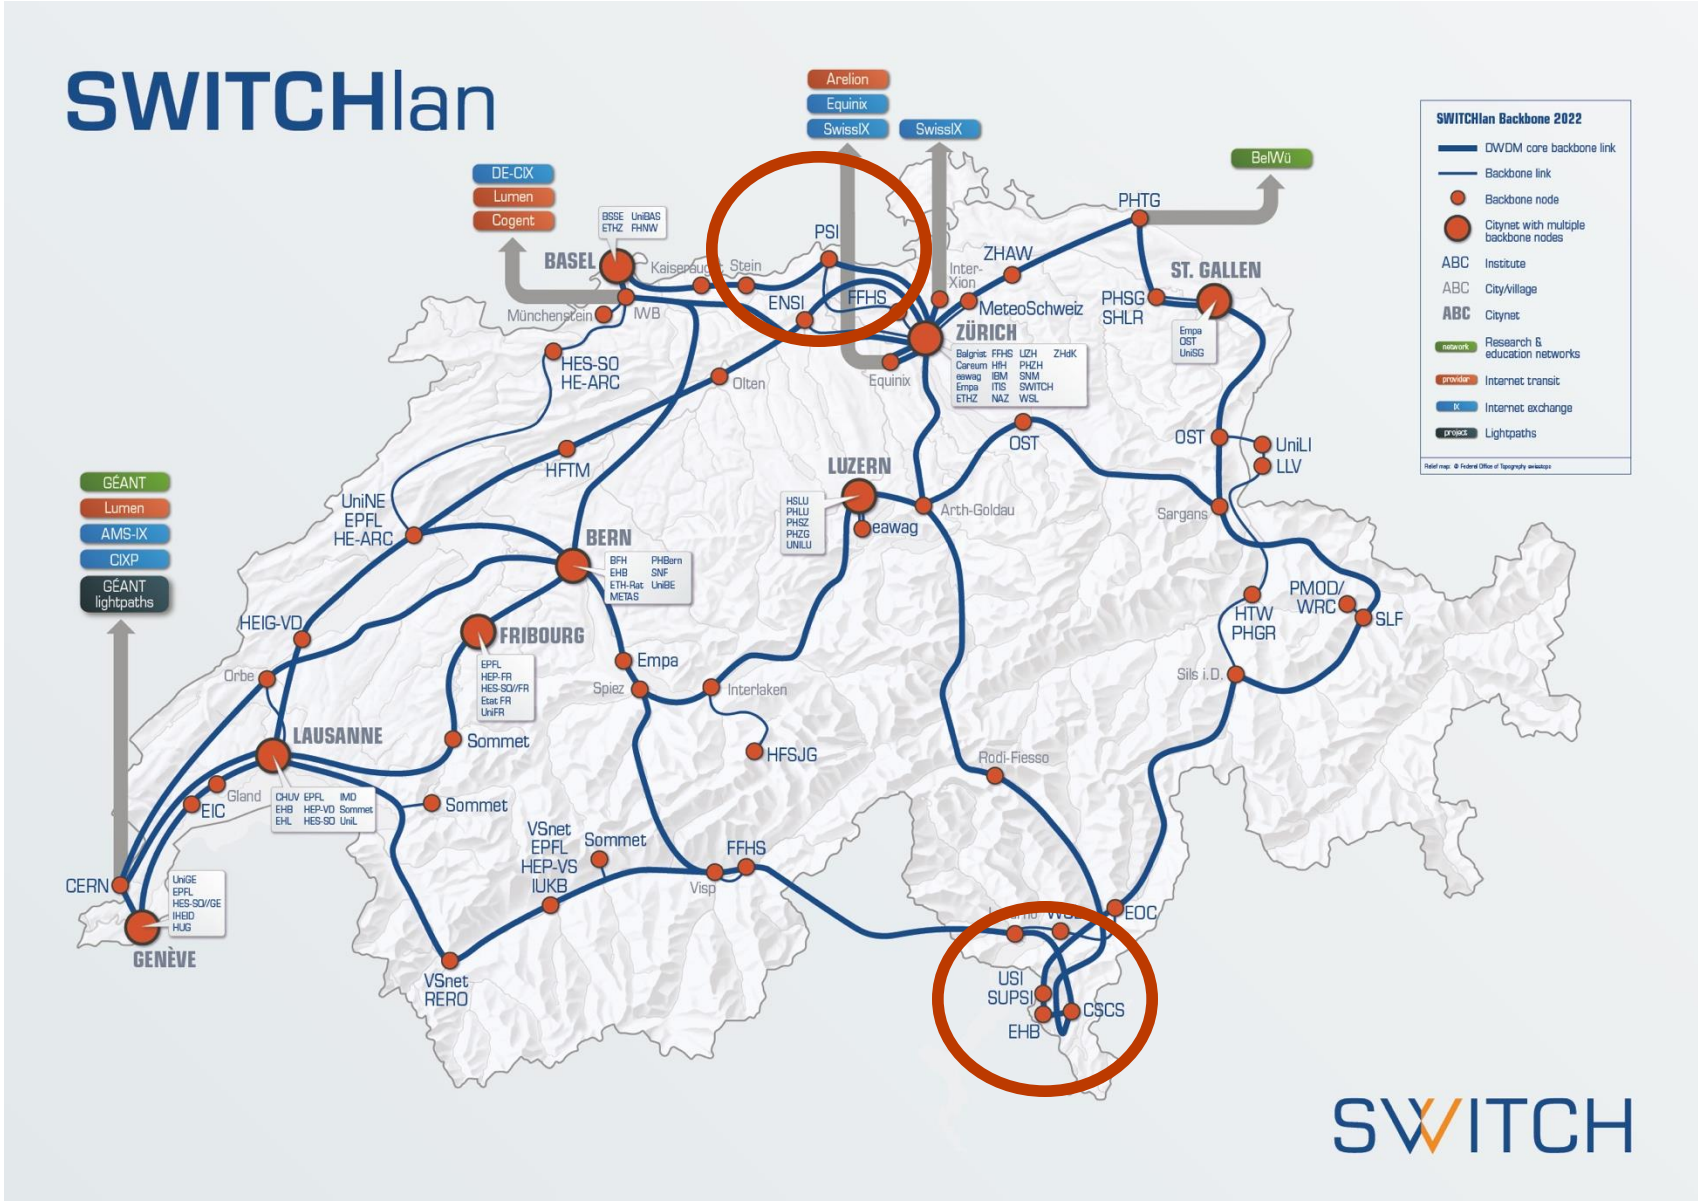

Technology: Network

- Virtual Local Area Network (VLAN)
  - enables network segregation and isolation
- Virtual Routing and Forwarding (VRF)
  - enable virtual connection and routing among VLANs
- Alps nodes automatically joining PSI VLAN
  - IP is configured at boot time to PSI network
- PSI provides its own security policies

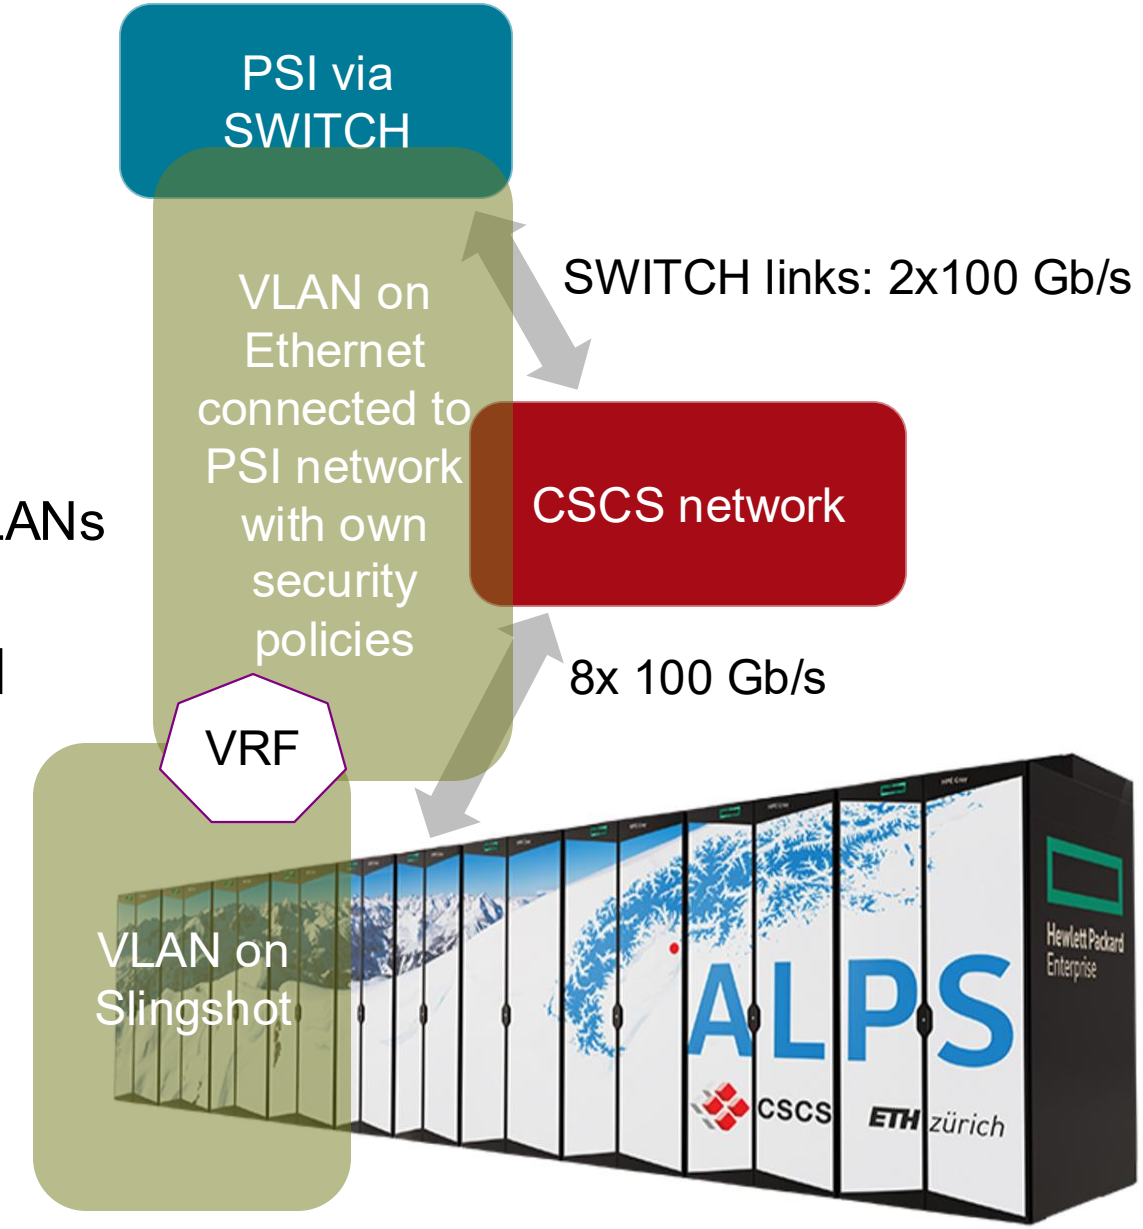

Technology: Storage

- Lack of existing HPC-compatible, multi-tenant storage solutions
- Allocate dedicated storage, HPE ClusterStor, Lustre
- PSI required node-local storage, which is not possible
- Allocated 2 pools of storage:
  - HDD for scratch (10PB)
  - SSD for high IOPS (200TB) loopback mounted on compute nodes
- PSI learning of Lustre
  - Storage fully managed by PSI
  - CSCS support for consultancy
  - Direct access to HPE support

Technology: service management

- Set of services:
  - PSI IAM LDAP, DNS, NTP, AFS, Slurm, workflow engines, Jupyter, PE with modules,…
- PSI has root access to compute nodes for full operational control
- PSI full responsibility of service management
- Services connected to PSI network and PSI Kubernetes/VM infrastructure
- Configuration automatically applied after reboot
- A key challenge was PSI's requirement for RHEL, while Alps defaulted to Suse
- Security
  - PSI security zone on the network
  - Dedicated storage
  - RBAC for actions on CSM

Operational view

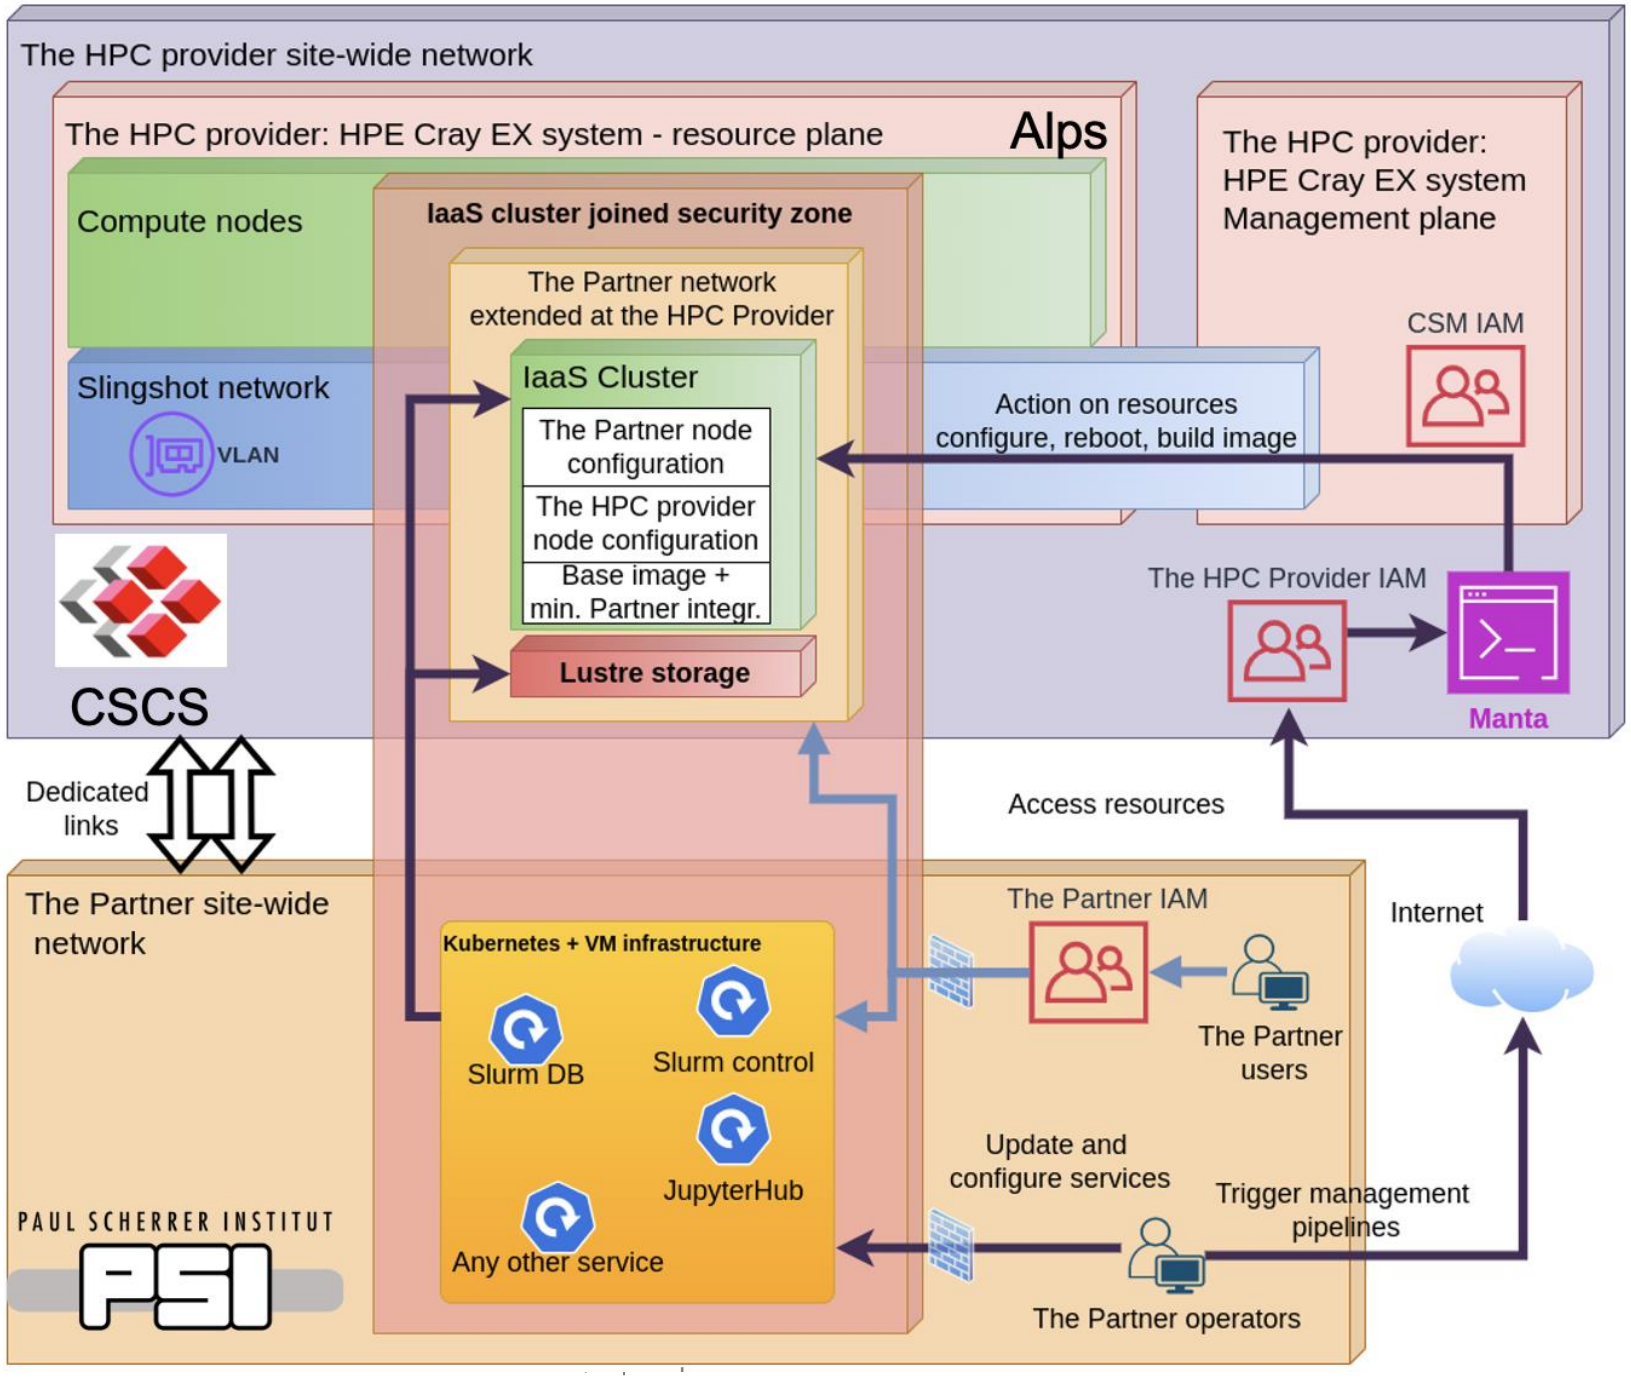

Project management and lessons learned

- Project management
  - 2 teams, regular meeting, workshop, risk, steering committee over 3 years
  - Alps was under active development, requiring frequent adjustments to account for hardware and software updates

- Lessons learned
  - Adaptability in project roadmap and features
  - Early user engagement drives development (Manta)

This project highlighted the importance of strong institutional collaboration and mutual understanding

Future work

- More IaaS on Alps
  - Need more automation
  - Further develop and harden Manta by integrating TAPMS
  - Harden Slingshot – identify issues and fixes
- Manta interfacing OpenCHAMI
  - Open Composable Heterogeneous Adaptable Management Infrastructure
- HPC-capable multi-tenant storage investigation
- New partners and use cases
  - 2 new partners identified to use IaaS on Alps
  - Emerging use cases include training ML models on sensitive or private datasets

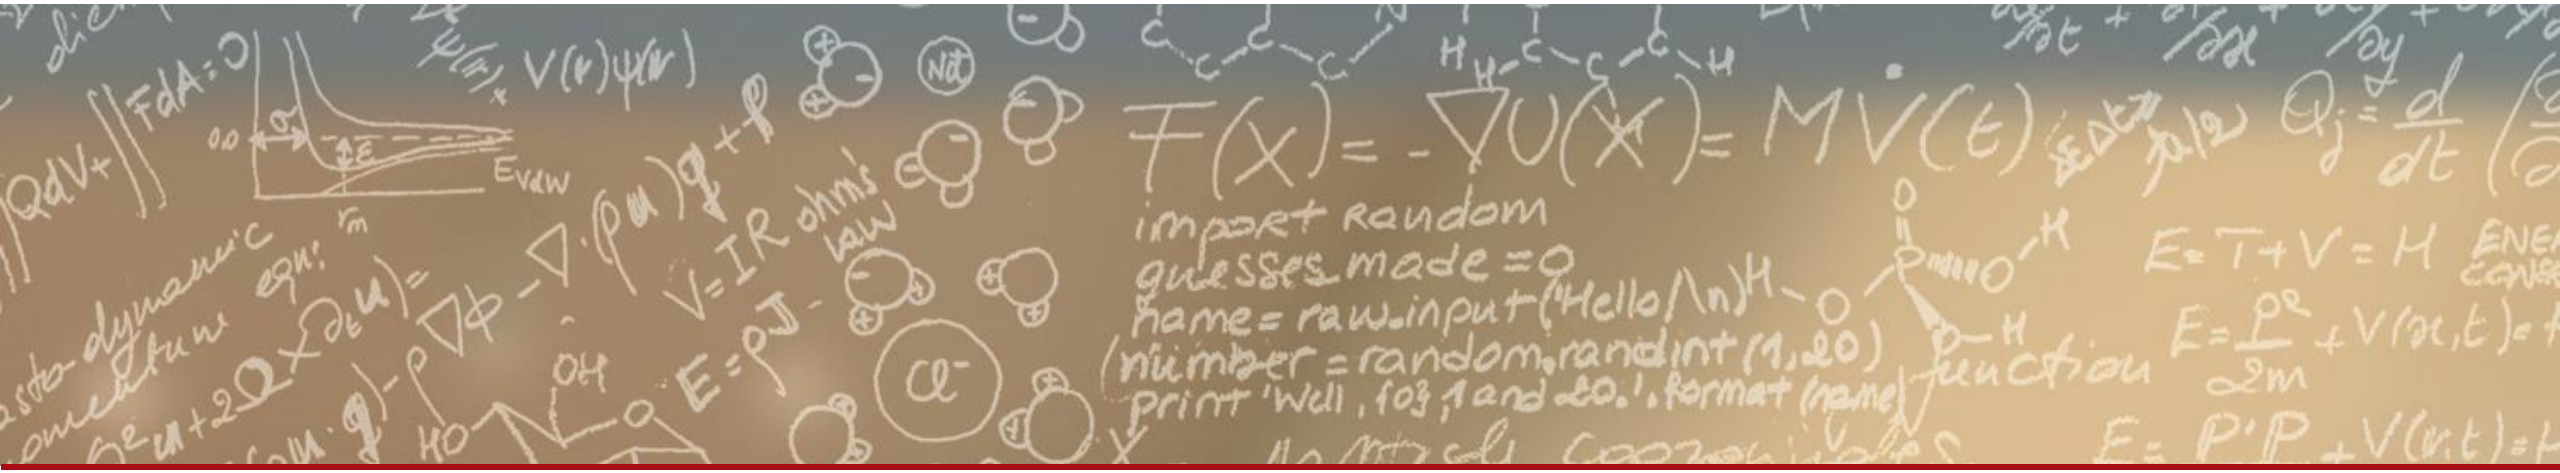

Thank you for your attention.