PAUL SCHERRER INSTITUT

Derek Feichtinger ::

Merlin7 - A First Tenant-Managed Alps vCluster

hpc-ch forum on HPC and Data as a Service, 5. Oct 2023

PAUL SCHERRER INSTITUT PSI

Outline

PAUL SCHERRER INSTITUT PSI

Topic

PAUL SCHERRER INSTITUT

Motivation

PSI's main data center at electrical capacity limit (300 kW). No local expansion at PSI planned

Growing Demand

- growth of demand for standard HPC/HTC resources (mostly conservative)

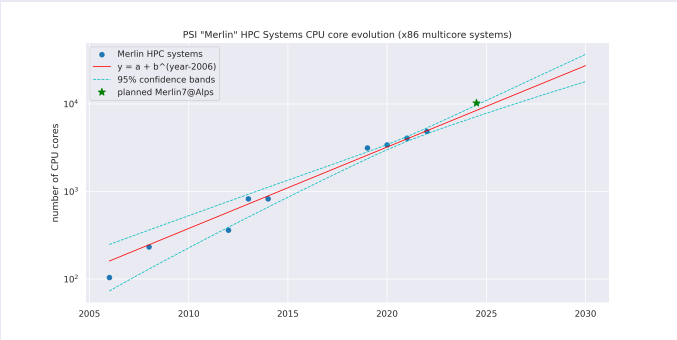

- Future experiments at the facilities predicted to require huge resources (peaky)

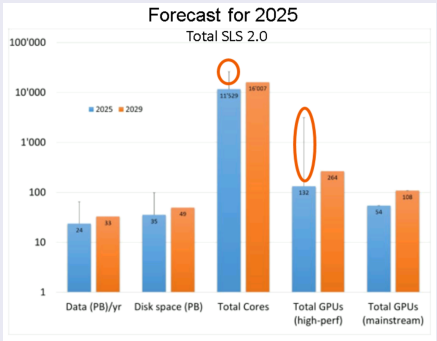

- potential impact of Machine Learning use cases

PAUL SCHERRER INSTITUT

PSI central clusters provided by SCD AWI

HPC cluster environments

System	Division	Users	CPU cores	Storage PiB	GPUs
Merlin6 (+ part Merlin5)	All	220	5212	5.5	8*A100, 70*consumer
PSI HPC@CSCS Daint	All	~40	(2579.9)	0.08	(P100 nodes avail.)
RA (SLS offline)	PSD	facility visitors	2540	11	8*A100, 4*V100
SLS (beamlines)	PSD	facility visitors	2408	3.2	2*V100S, 4*P100
SwissFEL (beaml.)	PSD	facility visitors	836	5.5	1*P100
PSI LHC/CMS Tier-3	NUM[1]	60	864	1.5	8*GTX 1080 Ti
MEG	NUM	100	540	1.1	

- Merlin6 resources are *offered centrally* to all PSI users, but also *allow buy-in* for prioritized access
  - MeG and Mu4e particle physics experiments invest in Merlin as main offline resource
  - GFA (Ansys simulations) and BIO invest in Merlin as main offline resource

[1] Tier-3 is a common infrastructure for the LHC/CMS groups of ETHZ, UniZ, PSI

PAUL SCHERRER INSTITUT

PSI central clusters provided by SCD AWI

HPC cluster environments

System	Division	Users	CPU cores	Storage PiB	GPUs
Merlin6 (+ part Merlin5)	All	220	5212	5.5	8*A100, 70*consumer
PSI HPC@CSCS Daint	All	~40	(2579.9)	0.08	(P100 nodes avail.)
RA (SLS offline)	PSD	facility visitors	2540	11	8*A100, 4*V100
SLS (beamlines)	PSD	facility visitors	2408	3.2	2*V100S, 4*P100
SwissFEL (beaml.)	PSD	facility visitors	836	5.5	1*P100
PSI LHC/CMS Tier-3	NUM[1]	60	864	1.5	8*GTX 1080 Ti
MEG	NUM	100	540	1.1	

- Merlin6 resources are *offered centrally* to all PSI users, but also *allow buy-in* for prioritized access
  - MeG and Mu4e particle physics experiments invest in Merlin as main offline resource
  - GFA (Ansys simulations) and BIO invest in Merlin as main offline resource
- the next generation of highlighted resources were selected as candidates for moving out of PSI

[1] Tier-3 is a common infrastructure for the LHC/CMS groups of ETHZ, UniZ, PSI

PAUL SCHERRER INSTITUT

PSI central clusters provided by SCD AWI

HPC cluster environments

System	Division	Users	CPU cores	Storage PiB	GPUs
Merlin6 (+ part Merlin5)	All	220	5212	5.5	8*A100, 70*consumer
PSI HPC@CSCS Daint	All	~40	(2579.9)	0.08	(P100 nodes avail.)
RA (SLS offline)	PSD	facility visitors	2540	11	8*A100, 4*V100
SLS (beamlines)	PSD	facility visitors	2408	3.2	2*V100S, 4*P100
SwissFEL (beaml.)	PSD	facility visitors	836	5.5	1*P100
PSI LHC/CMS Tier-3	NUM[1]	60	864	1.5	8*GTX 1080 Ti
MEG	NUM	100	540	1.1	

- Merlin6 resources are *offered centrally* to all PSI users, but also *allow buy-in* for prioritized access
  - MeG and Mu4e particle physics experiments invest in Merlin as main offline resource
  - GFA (Ansys simulations) and BIO invest in Merlin as main offline resource
- the next generation of highlighted resources were selected as candidates for moving out of PSI
  - PSI has a long term relationship with CSCS investing in its computational resources, and also PSI's *PetaByte Tape Archive* is run by CSCS.
  - PSI connected to CSCS through two 100Gb lines
  - PSI took strategic decision to enter a collaborative project with CSCS for establishing the next PSI general HPC cluster as a tenant cluster on the new CSCS Alps.

[1] Tier-3 is a common infrastructure for the LHC/CMS groups of ETHZ, UniZ, PSI

PSI Data Center Capacity Evolution

Without moving systems out

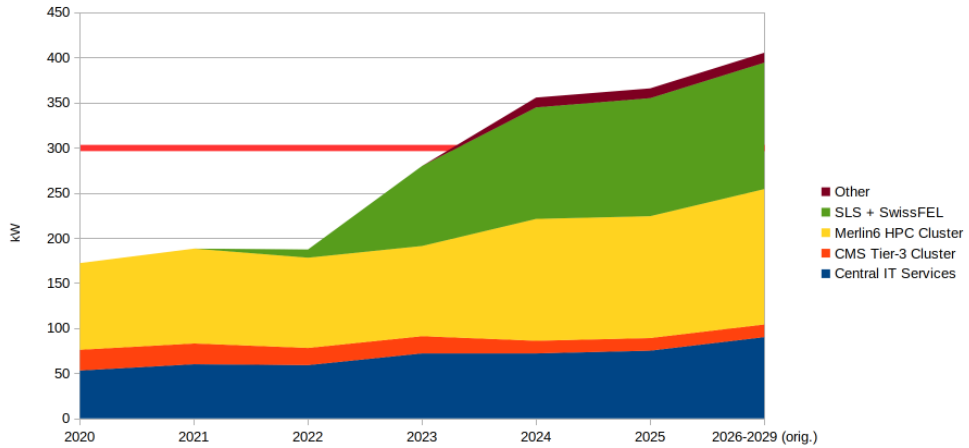

Moving systems externally

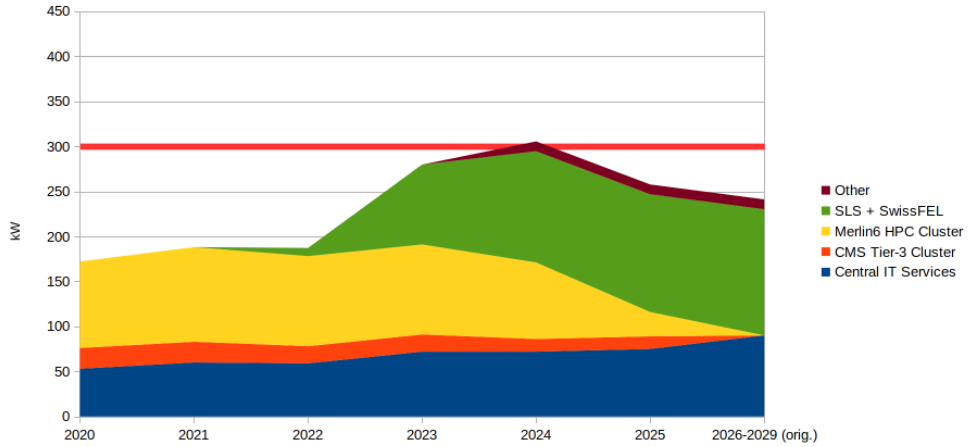

Merlin - a standard HPC Cluster with complex job loads

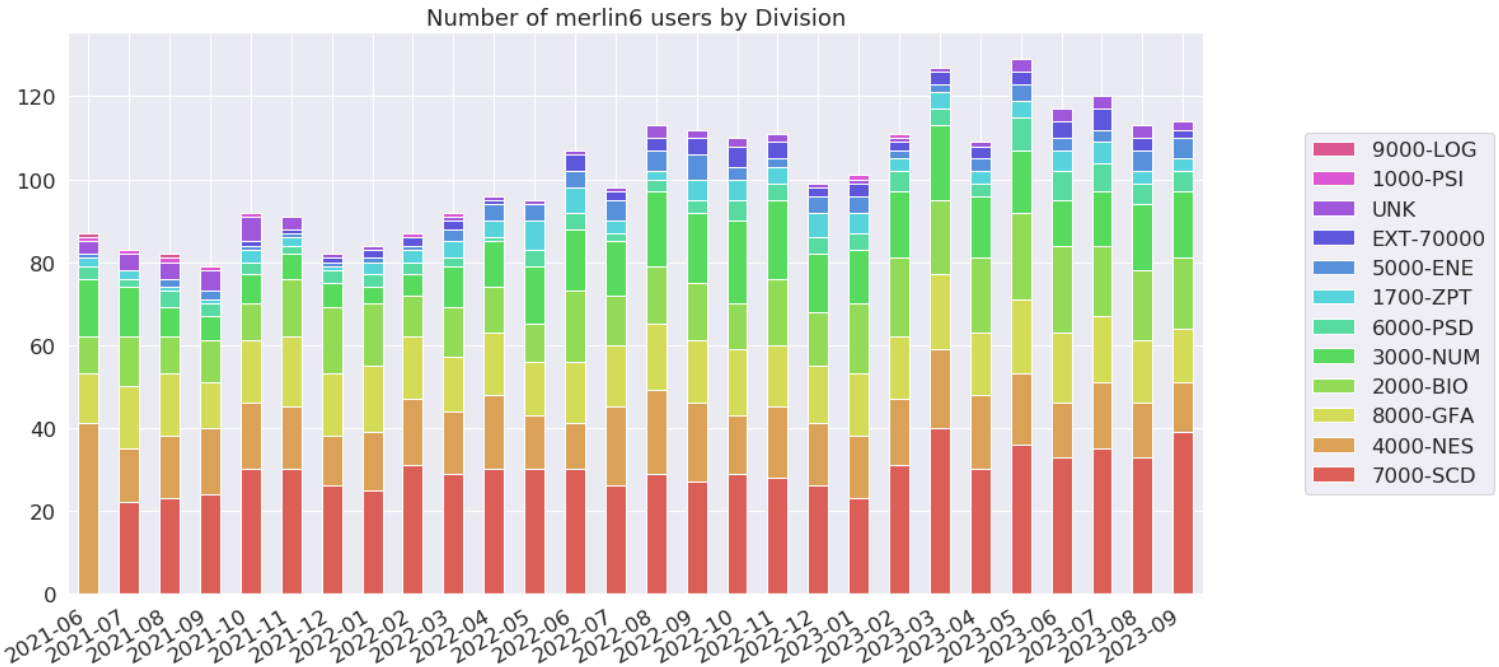

- ~320 users across 90 research groups from all 9 PSI divisions (+ external guests) over the plotted time range of 2 years.
- Complex mixture of HTC and HPC jobs. Big range of workflows and job characteristics.
- System fully integrated with PSI service and storage environment.
- multiple users/workloads may run concurrently on a node, if it has free cores. Goal is to optimally fill the available cores across the system.

PAUL SCHERRER INSTITUT

Merlin6 Cluster Usage Statistics (multicore nodes)

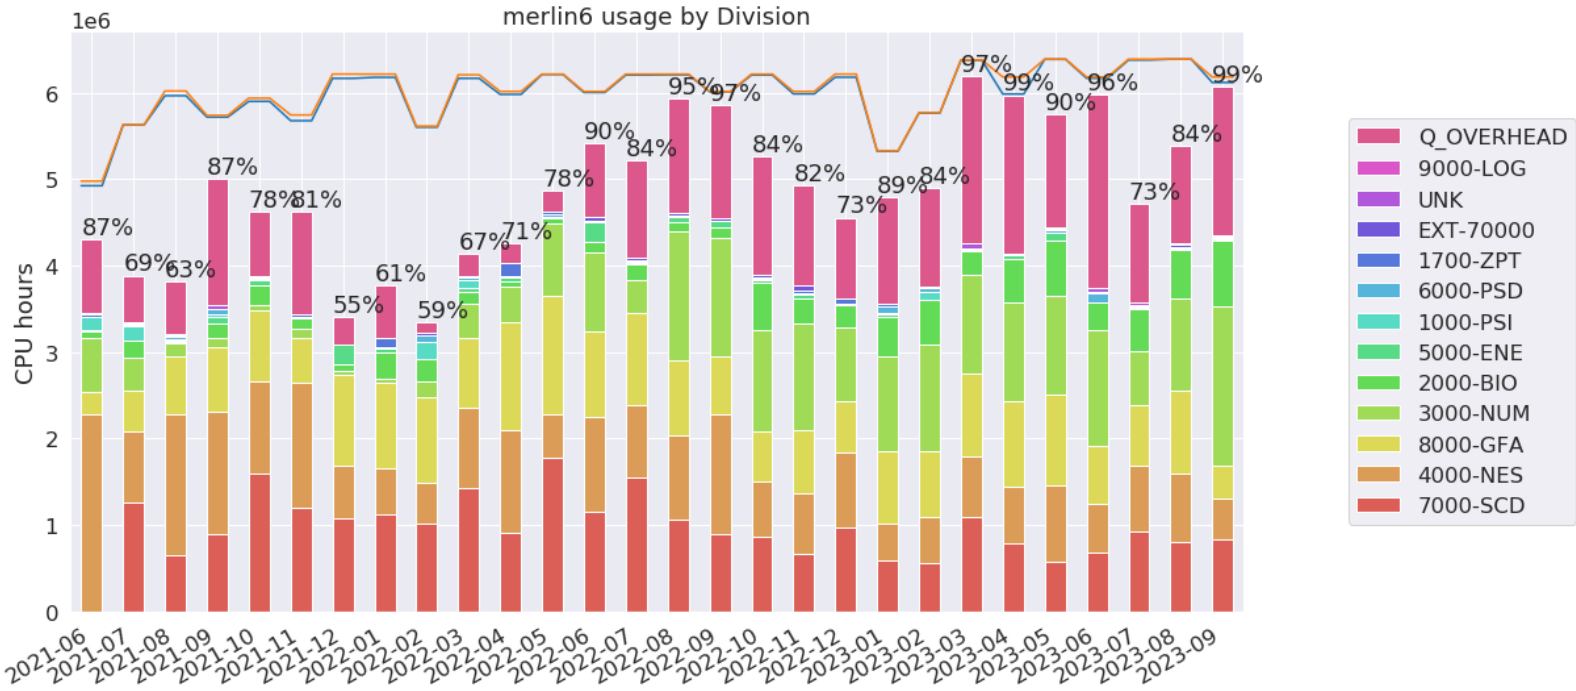

- the top lines indicate the total amount of deployed resources, and how many have been available per month.
- Usage percentage calculated in respect to all available cores of the system

PAUL SCHERRER INSTITUT PSI

(Some) Requirements for a Merlin system at CSCS

Desired Goals

- Conserve Merlin computing model as far as possible
  - support all existing workflows, access to other PSI services
  - Users should feel minimal differences from established Merlin cluster environment, to enable easy migration, same level of support effort
- allow same quality of PSI computing expert support (problem analysis) and configuration tuning/optimization

Some Important Requirements deriving from this

- PSI Admins must have sufficient privileged access
- Cluster resources must support PSI identity and access management
- Resources must be located in a PSI controlled VLAN

PAUL SCHERRER INSTITUT PSI

Topic

PAUL SCHERRER INSTITUT PSI

TransAlps Project

Collaborative project by PSI and CSCS

Near term goal

implement PSI's next HPC cluster "Merlin7" as a *vCluster* on top of the new CSCS Alps infrastructure.

vCluster: a "versatile software-defined Cluster"

Longer term goal

Prototype vClusters that will be able to elastically scale resources, to be used for addressing the large and peaky workloads predicted for the PSI facilities.

Timeline

TransAlps Project Timeline

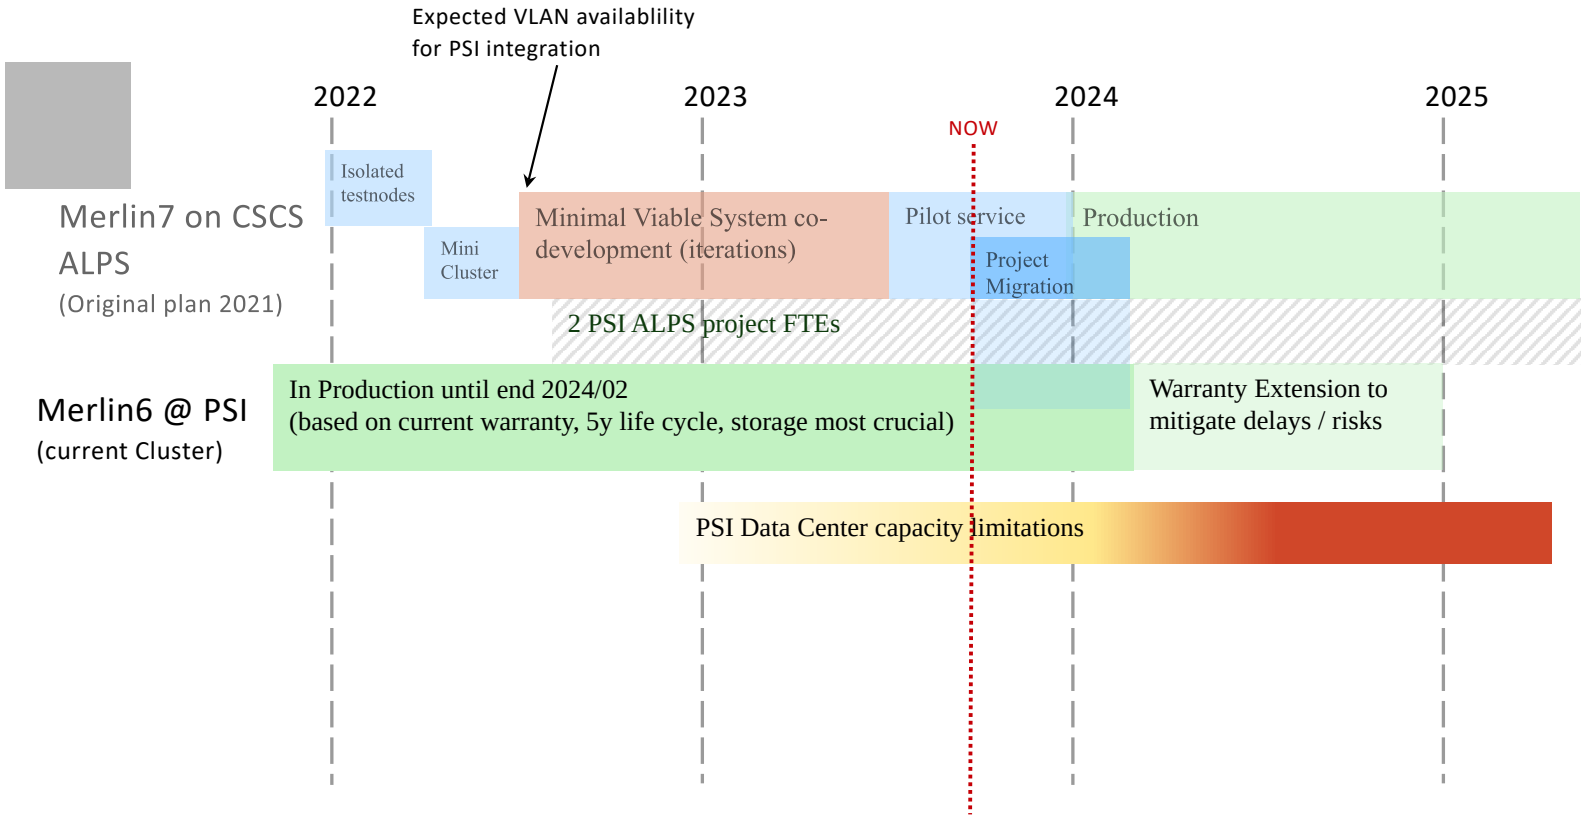

Timeline

TransAlps Project Timeline

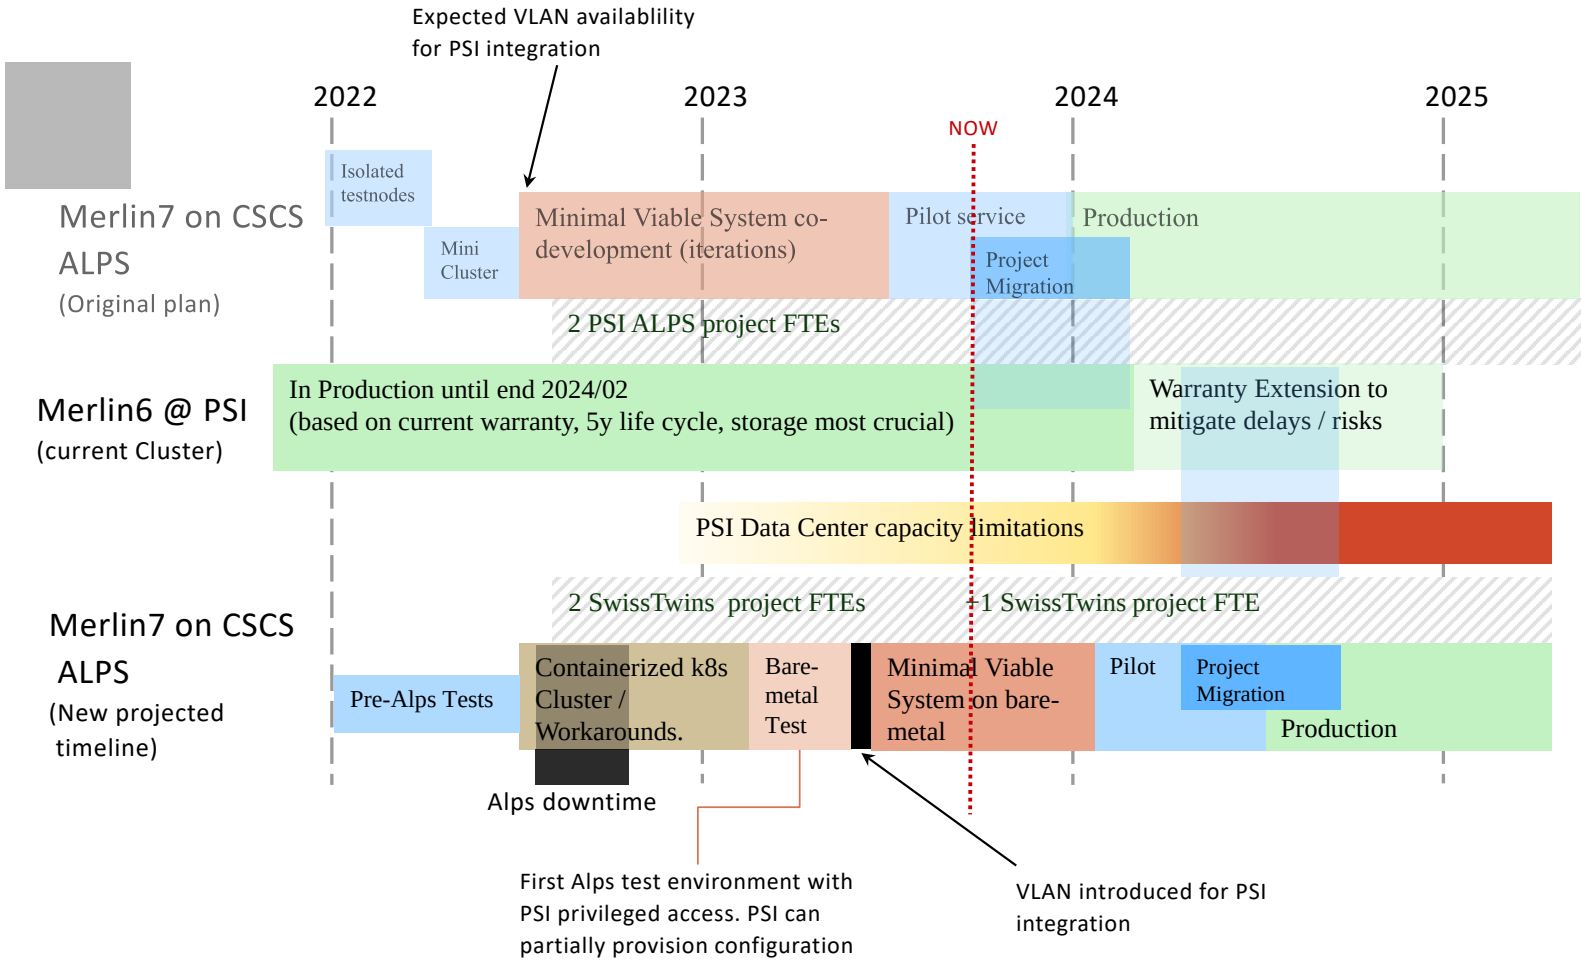

PAUL SCHERRER INSTITUT

Evaluation of two vCluster solution proposals

Initial solution proposals for vCluster installation

Bare-Metal Cluster

- Tenant running directly on the basic OS layer of a HW compute node.
- requires Tenant to have some access to HPE Cray Shasta management plane.

Containerized Cluster

- k8s pods run as "logical compute nodes" (simplest case: 1 pod per physical node)
- basic OS layer of the HW node controlled by CSCS

PAUL SCHERRER INSTITUT

Evaluation of two vCluster solution proposals

Initial solution proposals for vCluster installation

Bare-Metal Cluster

- Tenant running directly on the basic OS layer of a HW compute node.
- requires Tenant to have some access to HPE Cray Shasta management plane.

Containerized Cluster

- k8s pods run as "logical compute nodes" (simplest case: 1 pod per physical node)
- basic OS layer of the HW node controlled by CSCS
- This solution was discarded for our TransAlps project after some weeks of prototype development between CSCS and PSI. Could be taken up again for later vClusters.

PAUL SCHERRER INSTITUT PSI

Merlin7 sizing

Initial Merlin7 production system in 2024

Component	Amount	Details
x86 multicore nodes	82	AMD EPYC 7742 128 core, 512GB RAM
cores	10496	
GPU nodes	9	4 GraceHopper CPU/GPU node 256 cores, 512 GB RAM
GPUs	36	
TB HDD (Lustre)	8000	
TB SSD (Lustre)	100	

PAUL SCHERRER INSTITUT

Basic Architecture

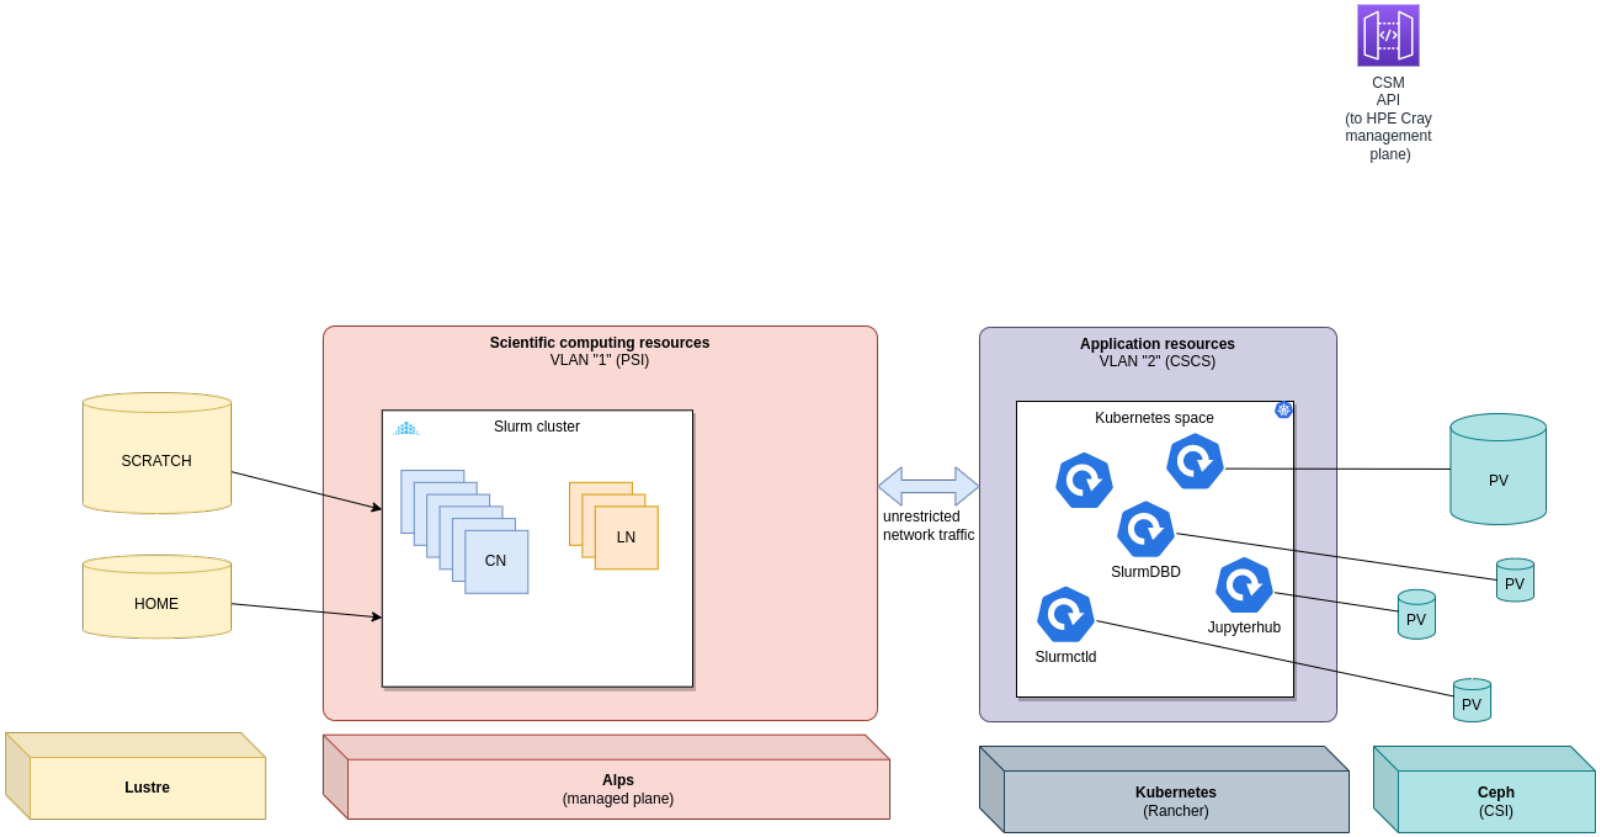

Architecture graph courtesy of Miguel Gila, CSCS

PAUL SCHERRER INSTITUT

Adding k8s resources

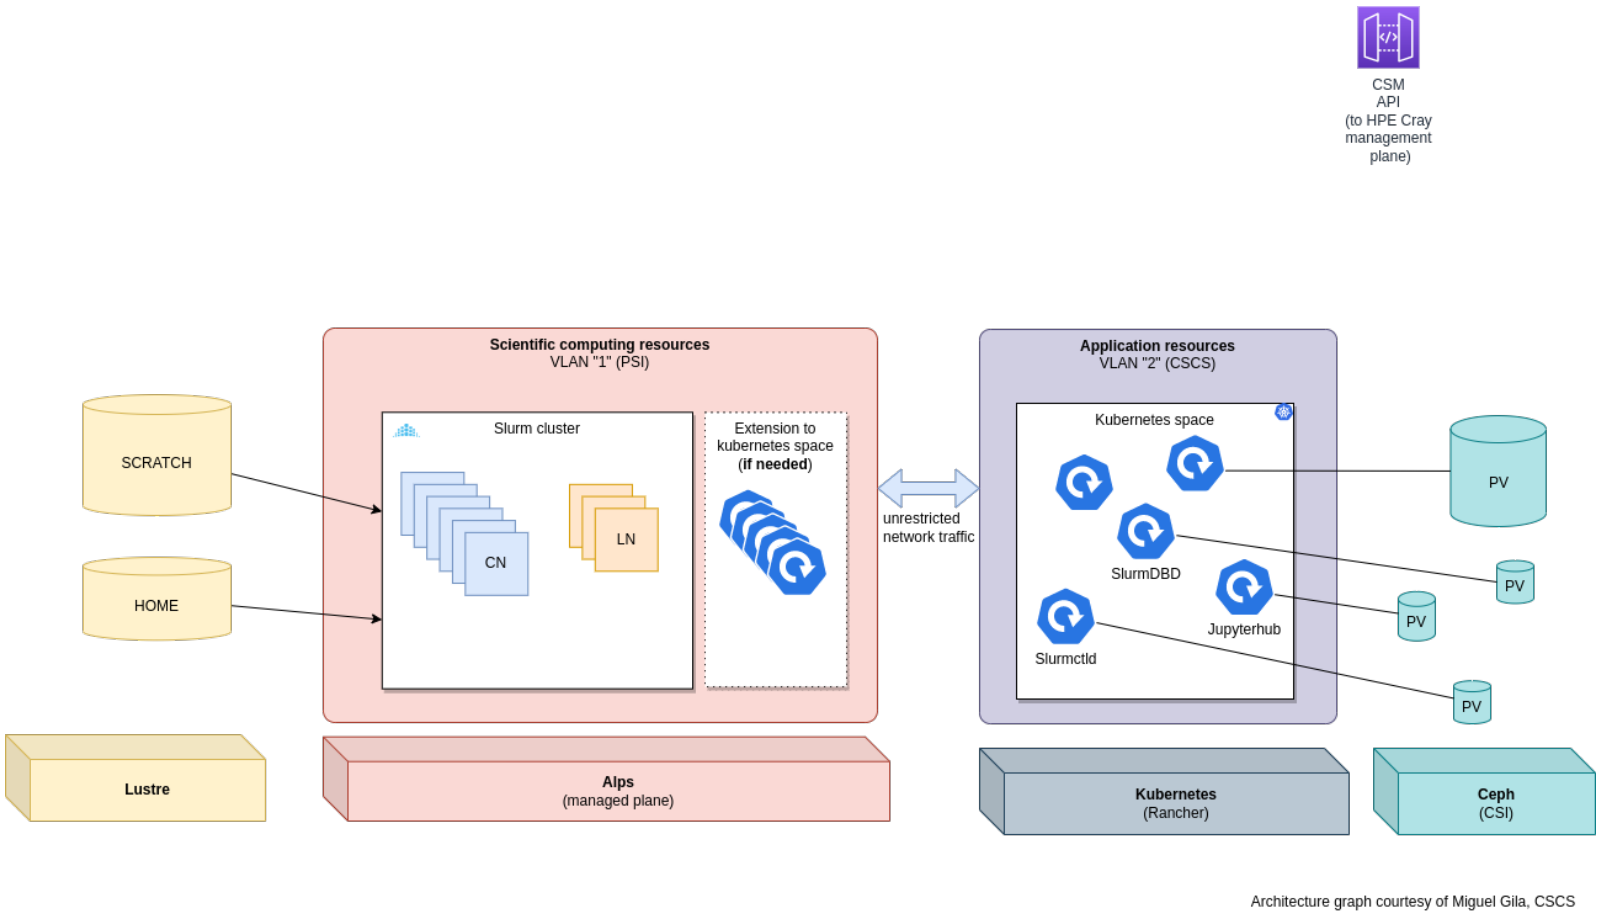

Architecture graph courtesy of Miguel Gila, CSCS

PAUL SCHERRER INSTITUT

PSI user access

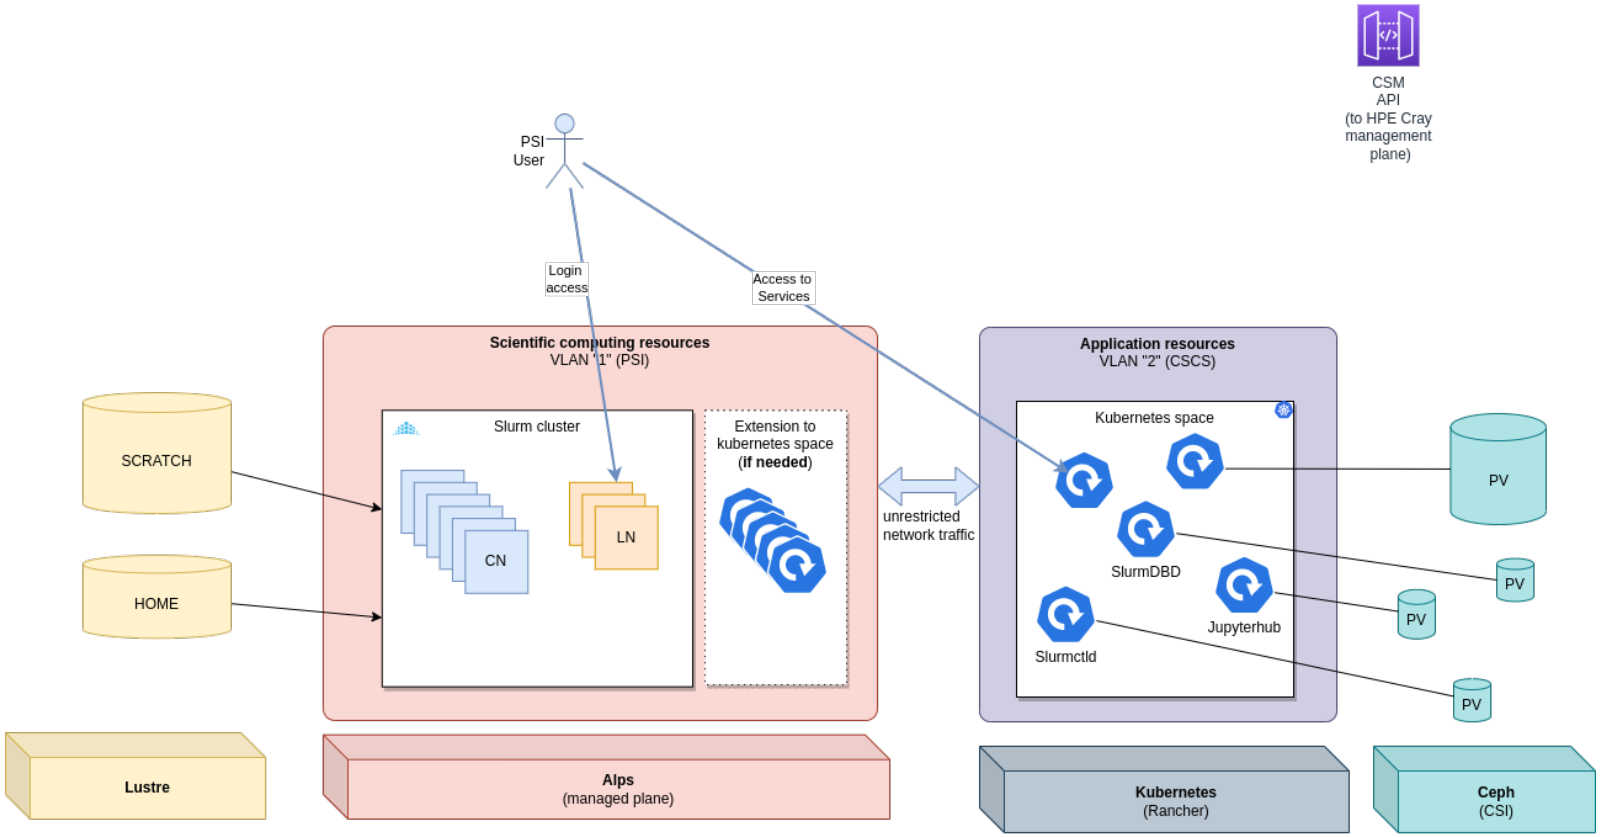

Architecture graph courtesy of Miguel Gila, CSCS

PAUL SCHERRER INSTITUT

PSI tenant admin access

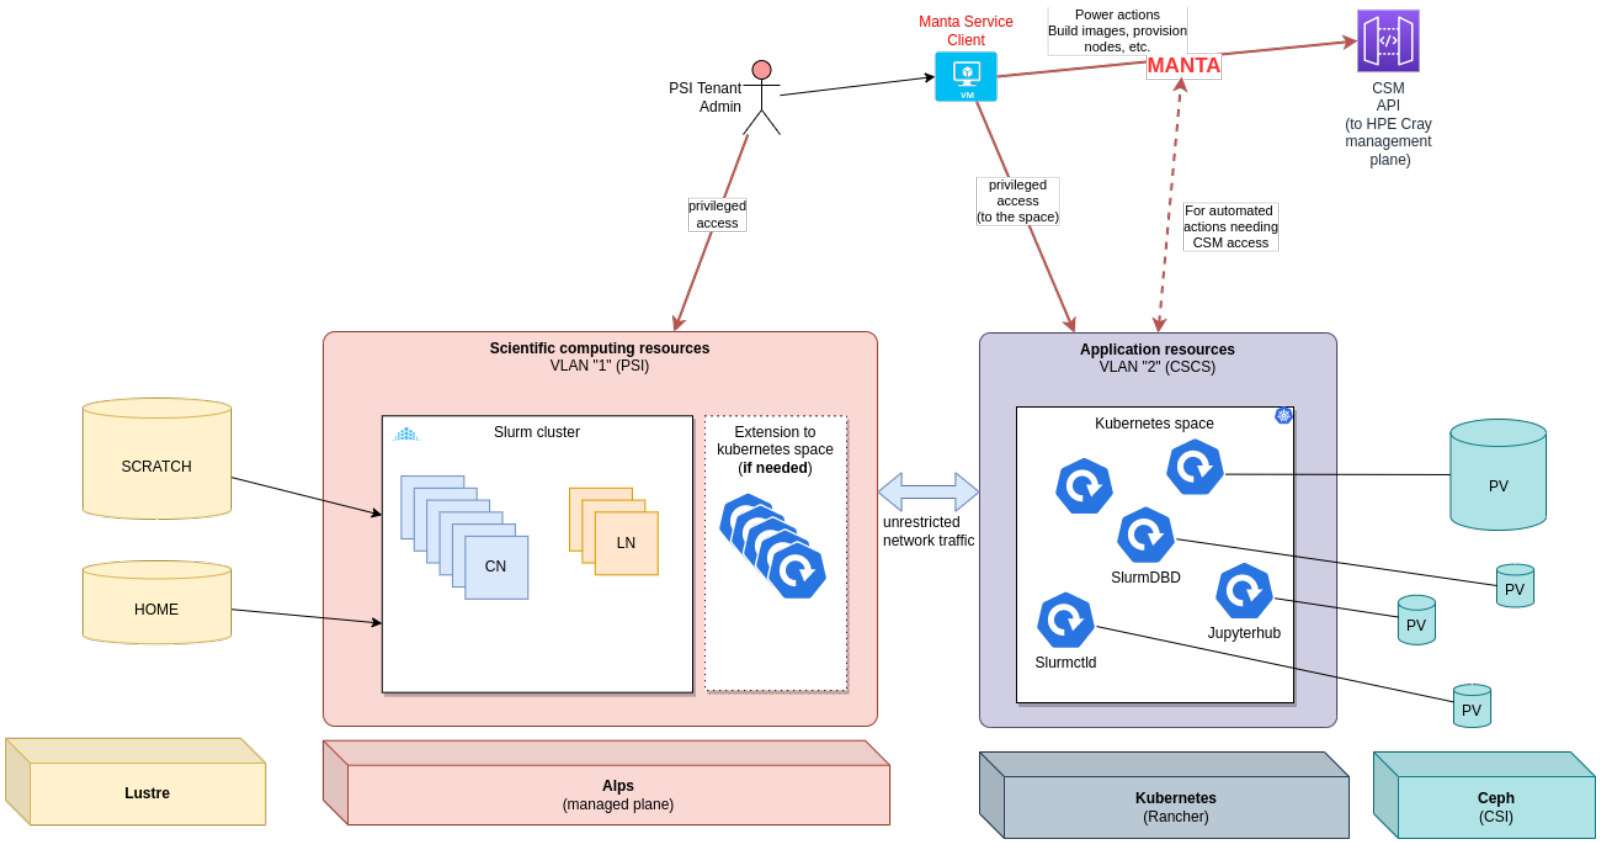

Architecture graph courtesy of Miguel Gila, CSCS

PAUL SCHERRER INSTITUT

Node OS image configuration

PSI tenant can add configuration layers that make up the OS image of an Alps PSI node

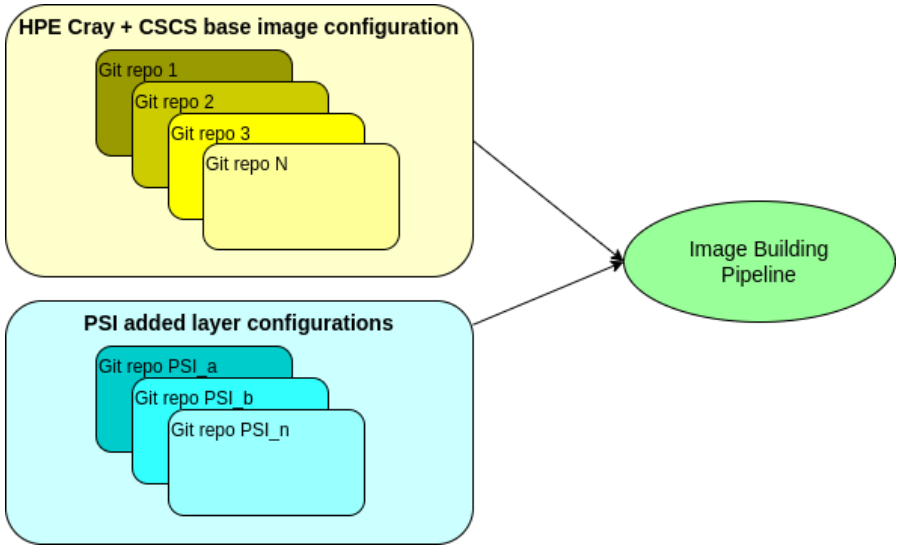

- Layers are added in the form of git repositories containing Ansible playbooks/roles
- part of the configuration is deployed onto the live node immediately after a boot.

PAUL SCHERRER INSTITUT

Current Development Cluster and Testing Environment

- Two vCluster environments. One targeted for enabling CI/CD based operations (testing node configurations before deployment)
- current cluster nodes based on node types
  - 2x AMD EPYC 7742, 512 GB RAM
  - 1x AMD EPYC 7713 + 4 NVidia A100 GPUs, 512 GB RAM
- Future production GPU nodes will be based on NVidia GraceHopper

PSI-dev - main development environment

- 14 Compute nodes
- 1 Login node

PSItds - testing environment

- 2 compute nodes
- 1 Login Node

Own k8s test cluster ("Morgana") at PSI for developing services and acquiring knowhow

PAUL SCHERRER INSTITUT

Current Status: Multi-Tenancy

- PSI vCluster in own VLAN since August of this year, following major updates to HPE Cray Shasta system and network (Slingshot)
- Tenant admin operations PSI admins can do by themselves
  - Boot/Reboot nodes
  - provision nodes with prepared OS images
  - Access to boot and configuration logs for debugging.
- CSCS assistance required for
  - building of new OS images (automatic syncing of PSI layer into pipeline tbd)
  - all storage related configuration (our prod. storage not yet available)
- new k8s system at CSCS for services still needs to be deployed: separate service needed after moving to VLAN

PAUL SCHERRER INSTITUT

Current Status

PSI Environment

- PSI IdM integration: OK
- AFS integration: OK
- Merlin base environment through Ansible playbooks/roles: ongoing
- PSI HPC SW Stack: test deployments using both env-modules (PSI pmodules) and SPACK

Major topics that that need to be addressed

- security policies (VLAN and FWs between institutions)
- dynamic node inventory (currently clusters are defined statically)
- Optimized usage of Slingshot network for specific SW (e.g. pre-compiled/commercial SW by runtime substituting fabric libraries)
- monitoring
  - making useful low level monitoring information from CSCS Alps management plane available and integration with PSI monitoring

PAUL SCHERRER INSTITUT PSI

Topic

PAUL SCHERRER INSTITUT

Credits - TransAlps Project Team

PSI (SCD/HPCE)

- Achim Gsell
- Derek Feichtinger (project lead)
- Elsa Germann
- Hans Viessmann
- Marc Caubet
- Spencer Bliven

CSCS

- Cerlane Leong (former project lead)
- Riccardo Di Maria (project lead)
- Chris Gamboni
- Hussein Harake
- Manuel Sopena
- Marco Passerini
- Theofilos Manitaras

- CSCS Miguel Gila (new Architect) and Pablo Fernandez (Finances)
- All colleagues in supporting teams of both CSCS and PSI.
- Part of this work is funded through SwissTwins by SERI