



Derek Feichtinger ::

## HPCE AWI Group Update 2023-12-11



**1** HPCE news

**2** TransAlps



1 HPCE news

2 TransAlps



## general news

- Pmodules-1.1.16 ready for Beta testing (Achim)
  - improvements on overlays, offers YAML based configuration files
- AFS status definition and future: intensive collection of current use cases (Achim)

## BIO Computing

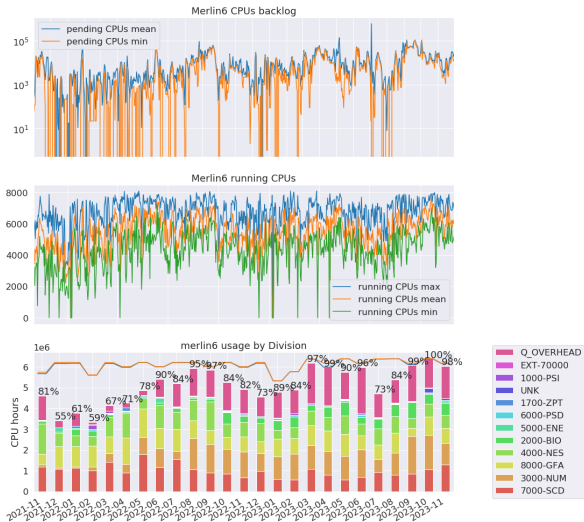
- Greta taking over former responsibilities of Spencer. She already is acting as principal point of contact for BIO HPC and is getting into the operational and project work
  - onboarding process: learning about PSI HPC environments and BIO related operations
- Greta contributing to collaboration of PSI within TransAlps project with CSCS/NVidia/Lindahl-Group on optimization of Relion CryoEM Analysis SW for NVidia Grace Hopper architectures
  - benchmarking of PSI reference data sets

## LHC Computing:

- Tier-2: massive downtimes due to problems of CSCS Ceph Storage
- Tier-3: Move to CSCS could not start yet due to CSCS manpower required to fix current infrastructure problems



# Merlin6 usage high - resource contention



- Q\_OVERHEAD
- EXT-70000
- 1000-PSI
- UNK
- 1700-ZPT
- 6000-PSD
- 5000-ENE
- 2000-BIO
- 4000-NES
- 8000-GFA
- 3000-NUM
- 7000-SCD

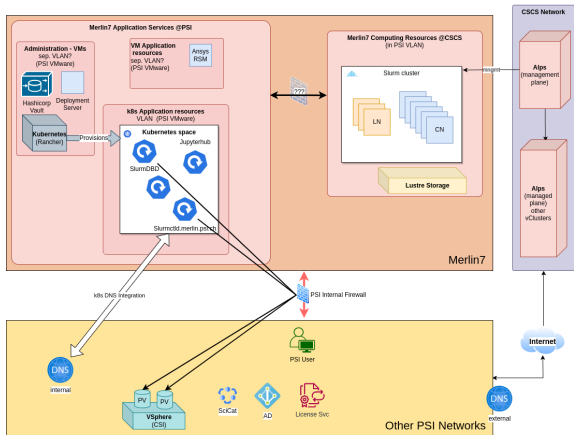


1 HPCE news

2 TransAlps

# Merlin7 core services will run at PSI

More Detailed Scenario 2 with Application Resources at PSI, k8s service and VMs



**Desired solution for application services on k8s and VMs:**

- try to have minimal FW barriers between application services and compute nodes
- more restrictive FW barriers between PSI internal net and application services, but must be able to resolve internal PSI addresses and interact with some well defined internal services

21. 11. 2023 D. Feichtinger / M. Caubert

## Building up of PSI side core service infrastructure

- a service subnet for Merlin7 core services has been created
  - VMWare cluster supplied by AIT (Peter Hüsser's group)
  - a kubernetes production level system will be deployed on VMs using *Rancher* (Hans)
    - Slurm scheduling k8s deployment has been implemented and slurm will run from there
  - secret handling will be moved to a *Hashicorp Vault* installation (Krisz)
  - we already saw that some services are not suitable for containerization (Krisz work on Anysis RSM service). Therefore we foresee running such services from traditional VMs.

## PSI Environment deployment (Marc, Elsa, Hans)

- currently Cray/HPCE deployment chain is complex and difficult outside tenants to deploy SW into the Alps vCluster boot images. Many problems detected.
  - Also: very restricted HPE/Cray deployment environment, ancient ansible
  - we made the decision to only make minimal use of the HPE/Cray image provisioning mechanism, only using it for time intensive tasks like RPM installations
  - rest handled by pushing from a PSI-side ansible deployment host
  - gives us more independence and isolates us from changes in the Cray/HPE provisioning chain
    - price to pay: some longer node boot times, need to implement an post-boot trigger mechanism
- CI/CD (Elsa)
  - Implementation of workflows using development/pre-production/production branches with our test (psi-tds) and main alps vClusters (psi-dev)





# Timeline

