



Derek Feichtinger ::


HPCE group update

22. July 2024



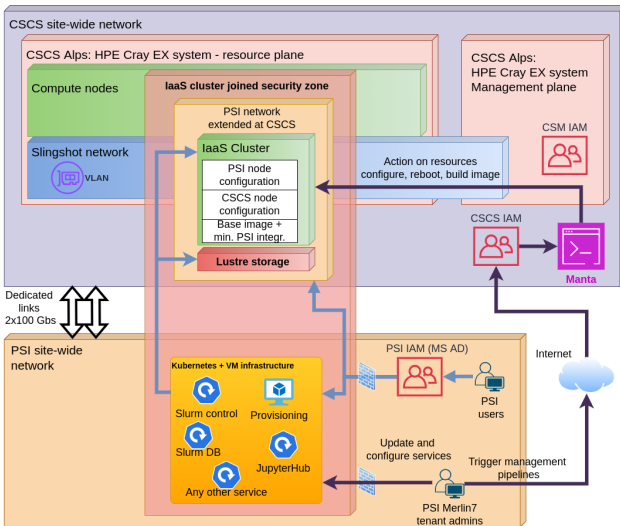
## 1 TransAlps

## Status – Timeline as of July 7th



- **1st of August: Preproduction:** for x86 multicore nodes. Not yet clear for Grace Hopper GPU nodes
  - Selected Test Users and Stakeholders
- **1st of October: Production**
  - Migration of projects and project storage areas
- **End of December:** Decommissioning of Merlin6 at PSI

# Architecture





- System not available for most of June
  - Major Downtime for recabling that CSCS had to extend
  - Power cut affecting CSCS and Lugano area with some fallout
- In first week of July CSCS made the newly defined Alps vClusters available
  - CSCS added functionality to the *manta* management tool that will allow us to move systems between vClusters
- **Grace Hopper nodes** will become available some time in August

## "Burst" vCluster: merlin7 + gmerlin7

79	2 x AMD EPYC 7742 64 cores
3	AMD EPYC 7713 64 cores, NVidia A100

## "Sichle" vCluster: Test system

3	2 x AMD EPYC 7742 64 cores
1	AMD EPYC 7713 64 cores, NVidia A100



## PSI Merlin7 platform deployment

- finishing initial version of PSI provisioning layers
- Slurm
  - V23.11.8 now compiled with HPE Cray *Slingshot* integration
  - *configless* mode allows dynamic joining/leaving of nodes
- Scientific SW deployment on AFS
  - efficient local caching of AFS allows us to offer end users to install their own python environments (metadata-heavy, problematic on lustre)
  - initial pmodule + Spack based deployments
- dedicated Lustre storage appliance "PSISTOR"
  - ~9PB **User and Project areas**
  - 1TB **"local" scratch** per node (using SSD pool)
  - ~20TB fast **cluster shared scratch** (using SSD pool)
- SSH and nomachine access

## Stakeholders

Stakeholder	Multicore nodes	GPU nodes	Lustre Storage/PB
CNM/MeG (Ritt)	5		2.0
SCD/LTC (Läuchli)	12		
CLS (Steinmetz)			3.5
CAS [GFA] (Kiselev)	10		
CNM/mu3e (Ritt)	4		1.2
SCD/LMS (Marzari)	3		
CENTRAL	48	9	1.3
TOTAL	82	9	8.0

- Initial Test Users x86: SCD/LMS: M. Krack (CP2K), GFA (Ansys), CNM/MeG S. Ritt
- Initial Test Users GPU: CLS (=BIO) may want to test current x86/A100 nodes



- Contract with CSCS for IaaS
  - to be iterated and finalized with CSCS over August
  - no standard SLA, but agreed to define a number of expected SLOs
- agreement in PSI/CSCS Steering board
  - CSCS will continue to offer same level of development support until end of year (compensation of delays)
  - PSI will receive better access to HPE Cray information through CSCS