

ROIHU

CSC's Next Supercomputer
Coming 2026



ROIHU



Roihu is an Eviden XH3000 supercomputer



AGGREGATE HIGH-PERFORMANCE LINPACK

34 PETAFLIPS

33.9×10^{15} operations per second

Infiniband NDR interconnect

200 Gbit/s per CPU node

200 Gbit/s per GPU

486 CPU nodes

2 x 192-core AMD Turin 9965



132 GPU nodes

4 x NVIDIA GH200 superchip



EVIDEN XH3000 SUPERCOMPUTER

9 RACKS

Fully liquid cooled – over 1 MW of power

DDN EXASCALER

6.5 PETABYTES

Fully flash-based Lustre storage

Scratch storage

6 PiB

560 GB/s Read bandwidth

280 GB/s Write bandwidth

Home + ProjAppl storage

0.5 PiB

120 GB/s Read bandwidth

100 GB/s Write bandwidth

Overview of solution – Roihu compute

CPU partition: 486 nodes

- Two 192-core AMD Turin 9965 CPUs
 - AMD Zen 5 architecture supports AVX-512 vector instruction set
- Memory per node: 768 GiB (414 nodes) or 1536 GiB (72 nodes)



• GPU partition: 132 nodes

- Four NVIDIA GH200 Grace Hopper superchips
 - Hopper H100 GPU, 96 GiB GPU memory (VRAM)
 - 72-core Grace ARM CPU, 120 GiB CPU memory
- Best price/performance solution for HPC & AI



• Additionally:

- 4 high memory (3 TiB) nodes
- 4 visualization nodes (Nvidia L40)

Overview of solution – Roihu storage

- **Home and application storage system**

- Separate from scratch to ensure responsiveness even under heavy scratch load
- For application installations and home folders
- 0.5 PiB total capacity, very fast performance (100+ GB/s)

- **Fully flash-based scratch storage**

- For data in active use
- 6 PiB in total size
- 10x more bandwidth than in Puhti storage, much better performance for difficult I/O patterns

- **All nodes have 960 GB local disks for temporary files**


- High memory and visualization nodes have 2 x 7.68 TB NVMe disks each

Features and functionality – what will or will not change?



- **Roihu web interface:** ✓

- Multi-factor authentication (MFA) will be required to login




- **Password-based SSH login:** ✗

- SSH certificate-based authentication is the target
- In case of delays: SSH keys as in Puhti and Mahti



- **Apptainer containers and Tykky:** ✓

- Improved documentation, base containers and integration of container registry
- Tykky will remain supported



- **Pre-installed module environment:** ✓

- Programming environment likely similar to Mahti (GNU, AOCC, CUDA, ...)
- Pre-installed collection of scientific software available similar to Puhti & Mahti




Features and functionality – what will or will not change?




- **Disk cleaning policy:** ✓

- Will be enforced from the start, but exact details not yet decided
- Either current Puhti process or alternative approach, **feedback is welcome!**
- Also: Improved tools for moving data to/from Allas



- **Slurm batch job scheduler:** ✓

- Partitions (queues) will have characteristics of both Puhti & Mahti
- Also: FirecREST API for accessing HPC resources using HTTP requests



- **x86 CPU architecture:** ✓ (CPU partition) ✗ (GPU partition)

- The Grace-Hopper 200 superchips have ARM CPUs(!)
- Separate Grace-Grace login node for compiling GPU programs and creating Apptainer containers with *--fakeroot* option

New functionality that is coming

- **Increased security**

- Short-term certificates for SSH access
- MFA – multi-factor authentication
- More isolation between projects/users

- **Improved support for integrations**

- APIs for accessing compute and storage resources

- **Support for confidential data computing**


- Integrated with Sensitive Data Desktop, based on encryption and isolation

- **Improved tools for Allas and container use**

- **Web interfaces will remain available and continuously improving!**


Another big change: Billing unit (BU) renewal

- **CSC's service offering and user base have grown substantially over the years**
 - More accurate resource and capacity management warranted (current model >30 years old!)
- **New model: specific units for HPC, HPC GPU, cloud and storage services**
 - To be deployed in **September 2025**
 - Affects resource applications – users must consider what resources and services are needed
 - Unused *old* BUs will be migrated = **no already granted resources will be lost**
 - **In the new model BUs will deprecate if left unused**



HPC BU

Puhti, Mahti, Roihu
jobs *without* GPU




HPC GPU BU

Puhti, Mahti, Roihu
jobs *with* GPU



Cloud BU

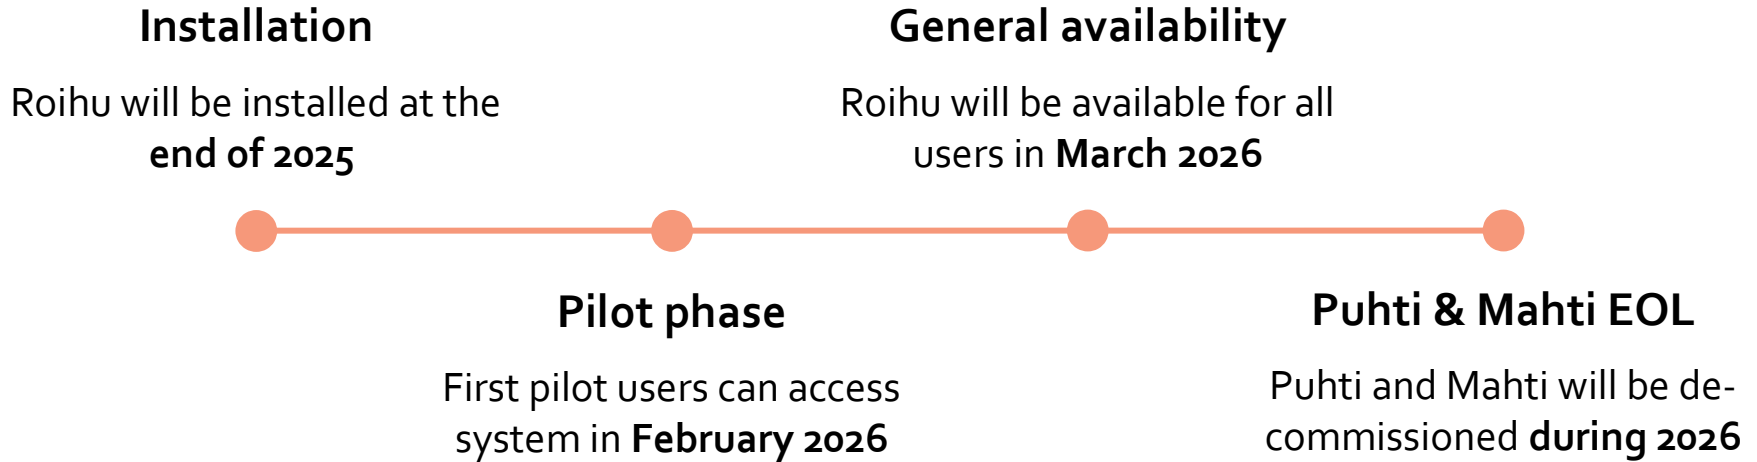
Pouta, Rahti,
Pukki, SD Desktop



Storage BU

Lustre, Allas, Shared
Fileservice, SD Connect

Roihu timetable (preliminary – subject to change!)



- Roihu installed in same datacenter as LUMI, can be brought up without disturbing Puhti and Mahti service
- Aim is to enable users to migrate to Roihu without a major break in HPC access




ALLAS

Allas and cloud services

- **New hardware for Allas and Pouta** will also be installed as part of DL2026 programme
- **Allas**
 - New version of Allas with **more storage capacity in 2026**
 - **New Allas will only support S3 interface** – swift not supported
 - **Data will not be automatically migrated** – CSC will develop tools and documentation to help with migration in 2026
- **Pouta**
 - **More GPUs to Pouta**, possibility to make available in **Rahti**
 - Also storage and CPU resources
 - **ePouta updates this year, cPouta later in 2026**

DL2026 outreach project

- The project's aim is to inform about the DL2026 infrastructure renewal and help users to transition to and use the new systems
- **How?**
 - Presentations, documentation, self-learning materials, blog posts, ...
 - <https://docs.csc.fi/computing/systems-roihu/>
- **Invite us to tell you more!**
 - We are happy to give a tailored presentation for your lab, department or university about DL2026, Roihu and what's to come
 - Contact us via servicedesk@csc.fi





**Share your needs and give us feedback
regarding the new infrastructure!**



facebook.com/CSCfi



youtube.com/CSCfi



linkedin.com/company/csc---it-center-for-science



github.com/CSCfi