CSC

ROIHU

CSC’s next supercomputer – coming 2026

DL2026 program

- **User requirements and expectations are constantly evolving**
  - Invest in an ecosystem that is flexible and meets new requirements
  - Service and functionality in focus
- **Coherent overall infrastructure for Finnish research**
  - Complementary with LUMI and other platforms

Puhti and Mahti will be replaced by a new supercomputer:

ROIHU

Cloud and **Allas services** to be invested in and improved as well

Roihu is an Eviden BullSequana XH3000 supercomputer

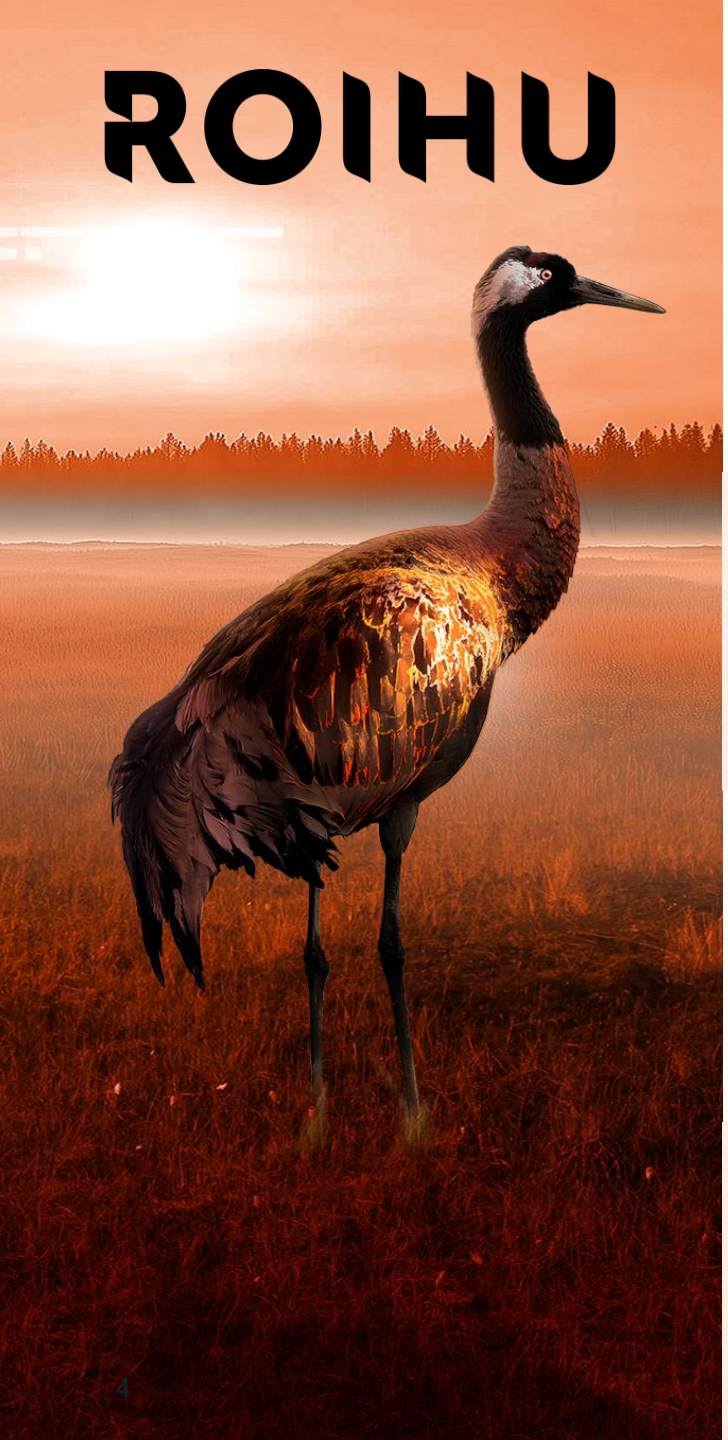

System overview
docs.csc.fi/roihu

AGGREGATE HIGH-PERFORMANCE LINPACK

34 PETAFLOPS

33.9 x 10^{15} operations per second

486 CPU nodes

2 x 192-core
AMD Turin 9965

132 GPU nodes

4 x NVIDIA
GH200 superchip

Infiniband NDR interconnect

200 Gbit/s per CPU node
200 Gbit/s per GPU

EVIDEN XH3000 SUPERCOMPUTER

9 RACKS

Fully liquid cooled – over 1 MW of power

DDN EXASCALER

6.5 PETABYTES

Fully flash-based Lustre storage

Scratch disk
6 PiB
560 GB/s read perf
280 GB/s write perf

Home & ProjAppl disk
0.5 PiB
120 GB/s read perf
100 GB/s write perf

Overview of solution – Roihu compute

- **Roihu is an Eviden XH3000 system**
  - Next generation of Mahti (XH2000)
- **CPU partition based on AMD Turin 9965**
  - 2 x 192 cores per node
  - AMD Zen 5 architecture has full AVX-512 support – twice the Flop rate vs. Mahti
- **GPU partition based on NVIDIA GH200**
  - 4 x H100 GPUs & 72-core Grace ARM CPUs per node
  - Best price/performance solution for HPC & AI
- **Special resources**
  - 4 x high memory (3 TiB) and visualization nodes (NVIDIA L40), each with 2 x 7.68 TB local disks
  - 307.2 TB disaggregate NVMe capacity

Roihu	
CPU compute nodes	486
CPU cores	186 624
Memory per node	768–1536 GiB
GPU compute nodes	132
GPUs	528
Memory per chip	120+96 GiB (CPU+GPU)
Scratch storage	6+ PiB
Home and ProjAppl	0.5 PiB
Node temp storage	960 GB
HPL aggregate perf.	33.9 PFlop/s

Overview of solution – Roihu storage

Fully flash-based Scratch storage for active data

- **10 times more bandwidth than Puhti Scratch**
  - Much better performance for difficult I/O patterns
- **6 PiB total capacity**
  - Can be expanded during lifetime of system
- **ProjData** – special disk area on Scratch filesystem for storing and sharing datasets

Home and ProjAppl storage system

- Application installations and home folders
- **Separated from Scratch**
  - Ensures responsiveness even under heavy Scratch load
- **0.5 PiB total capacity**
  - Very fast performance (100+ GB/s)

Node local storage resources

- All nodes have 960 GB local disks
  - **For temporary files**, not high-performance I/O
- Hugemem and visualization nodes have 15.36 TB high-performance NVMe storage
- 307.2 TB disaggregated NVMe capacity
  - Appears as local scratch from within a Slurm job

Features and functionality: what will **not** change

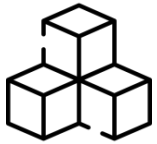

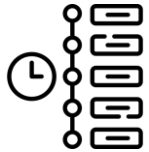

- **Roihu will have a web interface similar to Puhti & Mahti**
- **Pre-installed module environment**
  - Comprehensive stack of scientific software available
  - Programming environment will be similar to Mahti (GNU, AOCC, CUDA, OpenMPI, ...)
- **Slurm batch job scheduler**
  - Queues will have characteristics of both Puhti & Mahti
- **Apptainer containers and Tykky supported**
  - Improved documentation, base containers and container registry
  - Tykky will remain available and supported
- **Disk cleaning policy**
  - Will be enforced from the start, similar to current process on Puhti
  - CSC will provide improved tools for moving data to/from Allas

Features and functionality: **upcoming changes**

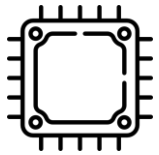

- **New hardware and OS (RHEL9)**
  - Intel vs. AMD, GPU nodes will have ARM CPUs (Grace)
- **Smaller, but much faster storage than Puhti & Mahti combined**
  - Capacity to be expanded during Roihu lifetime

- **Disaggregated fast storage service**
  - Users can request local disk mounts from centralized storage
- **Increased security**
  - Short-term certificates required for SSH access
  - Optional login nodes with per-user container isolation

- **Support for confidential data computing will be developed**
  - Integration with Sensitive Data services

- **FirecREST API for accessing Roihu compute and storage**
  - Integrate workflows and web services to HPC

Roihu timetable (tentative)

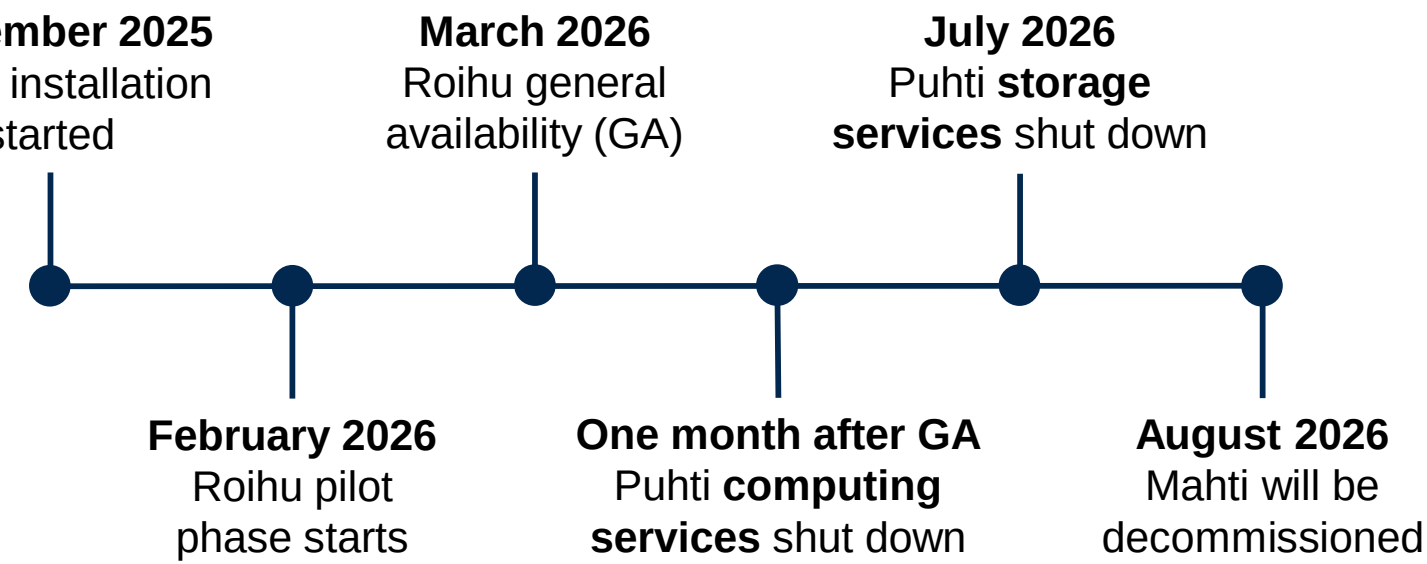

- Roihu installed in same datacenter as LUMI – will be brought up without disturbing Puhti & Mahti
- Aim is to enable users to migrate to Roihu without any break in HPC access

Roihu pilot phase

- **The pilot phase is a key part of Roihu's acceptance testing**
  - A heterogeneous and realistic workload will help CSC to improve and ensure the system's readiness
- **Selected pilot projects will get:**
  - Early access to Roihu and large-scale resources which will only be available during the pilot phase
  - Opportunity to do impactful science while helping CSC shape the permanent user environment of Roihu
  - Each project will be assigned a CSC specialist as a direct line for questions and guidance
- **Proposals will be reviewed for technical suitability, readiness and scientific potential**
  - Projects must be ready to setup their environments quickly, run jobs independently, and provide feedback to CSC

Apply for pilot access by 28th of November!

research.csc.fi/roihu-pilot

The pilot phase will run for about four weeks starting at the end of February 2026

How to prepare for Roihu migration?

Data migration

- Clean up your data before migration:
  only move what you really need!
  - Roihu has less disk space than Puhti & Mahti combined, and stricter criteria for granting extended quotas
  - Existing quota extensions will not be automatically moved to Roihu
- **As a general rule, data should be moved directly to Roihu (not via Allas)**
  - Familiarize with data transfer utilities in advance: docs.csc.fi/data/moving
  - Detailed Roihu migration guide coming later

Application migration

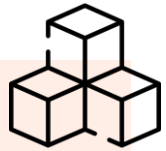

- **New hardware, OS, module stack**
  - Prepare to recompile your own software
  - Check if/how your GPU code can run on the GH200 nodes (ARM)
  - Containerized Python environments (Tykky) and other Apptainer containers might work as is (except on GPU nodes)
- New Slurm partitions – **old batch scripts should not be expected to work**
  - CSC will develop documentation and provide templates and support for setting up new scripts and workflows

Changes to Allas and cloud services

- **New hardware for Allas and Pouta** will also be installed as part of DL2026 program
- New version of Allas with **more storage capacity** coming in 2026
  - **”Allas 2” 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**, possibility to make available in **Rahti**
  - Also storage and CPU resource
  - ePouta updated this year, cPouta later in 2026

DL2026 outreach project

- Project’s aim is to inform about the DL2026 infrastructure renewal and help users to migrate to and use the new systems
- **How?**
  - Presentation, documentation, self-learning materials, blog posts, ...
  - docs.csc.fi/roihu
- **Invite us to tell you more!**
  - We are happy to give a tailored presentation for your lab, department or university about DL2026, Roihu, or CSC services in general

Contact us via:
servicedesk@csc.fi

Share your needs and give us feedback!

Follow us

LinkedIn

Instagram

Facebook

YouTube

csc.fi