

NAISS compute resources

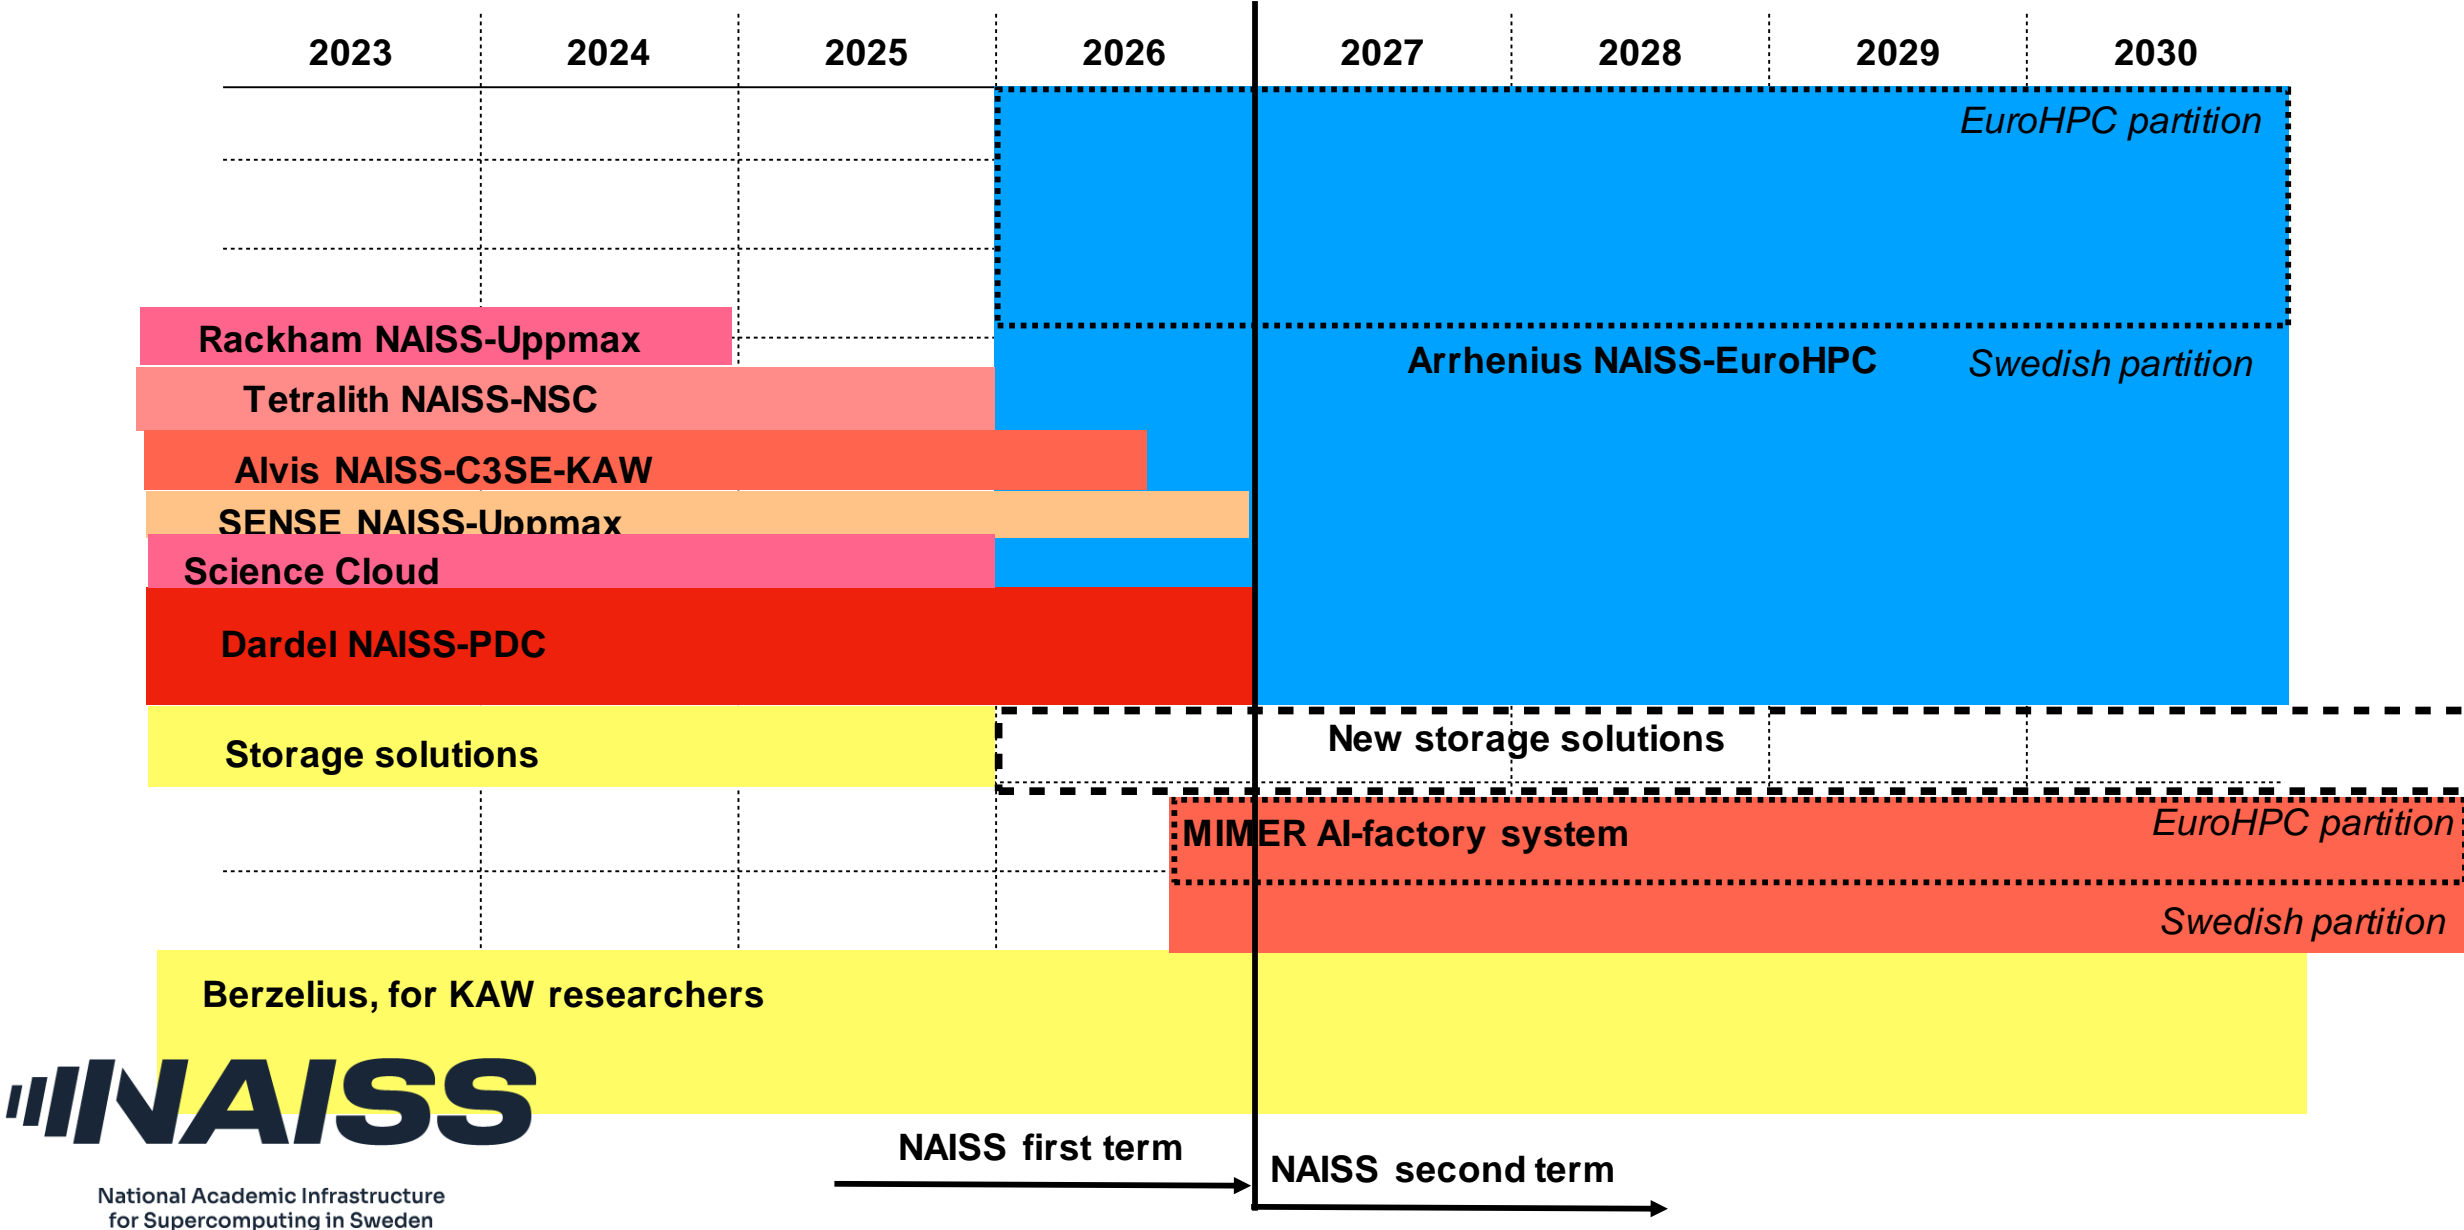
Björn Alling, Deputy NAISS Director
Professor of Theoretical Physics



NAISS compute resources

- From 2027, NAISS second term, NAISS will have retired all old SNIC systems and entered the NAISS era of fewer and larger compute systems.
- Arrhenius will serve as CPU and GPU HPC, Sensitive data, cloud resource, - a distinct increase in resources in 2026
- Mimer AI-factory resource will add AI-capacity, fall 2026.
- KAW-supported AI researchers will have continued access to the KAW funded Berzelius AI system.
- Project storage will continue but its usage as long-term storage will not be possible. New object storage will arrive Q1 2026 to replace SweStore. A storage-for-payment services.

Time line NAISS systems



Arrhenius

- Delivery in December 2025. Start of operations Q1 2026.
- CPU-partition: 424 AMD Turin 128-core
- GPU-partition: 382 nodes x 4 = 1528 Grace Hopper Superchips
- 29 PB fast parallel file storage.
- Sensitive data partition
- Cloud partition
- > 60 PFLOPS (~ 7xDardel)

Arrhenius

- Delivery in December 2025. Start of operations Q1 2026.
- CPU-partition: 424 AMD Turin 128-core
- GPU-partition: 382 nodes x 4 = 1528 Grace Hopper Superchips
- 29 PB fast parallel file storage.
- Sensitive data partition
- Cloud partition
- > 60 PFLOPS (~ 7xDardel)

Alvis for comparison

#GPUs	GPUs	Capability	CPU	Note
44	V100	7.0	Skylake	
160	T4	7.5	Skylake	
332	A40	8.6	Icelake	No IB
296	A100	8.0	Icelake	Fast Mimer
32	A100fat	8.0	Icelake	Fast Mimer

Arrhenius for simulations - the example VASP

- Arrhenius is a large increase in FLOPS thanks to the GPU-partition. Obviously good for e.g. AI applications.
- But for CPU:s the situation is different.
- What about traditional simulations codes?
- Important to port current CPU-load to Arrhenius GPUs
- The example of VASP - materials science DFT code
The most used software on Tetralith.

Arrhenius: GPU-part

Benchmark by Weine Olovsson


- **GPU node:**

382 x • 4 x Nvidia GH200 Grace Hopper Superchip

- H100 GPU
- 72c Arm Neoverse v2 cores

- **Quick test: GaAsBi 512 atoms 4 k-pts, VASP PBE**

- 1 x Tetralith CPU node: 3 084 s [8 nodes -> 471 s]
- 1 x A100 @Berzelius: 704 s (x4.4)
- 1 x GH200 @Dardel(testing): 404 s (x7.6)
(Same as Arrhenius)



Arrhenius: GPU-part

Benchmark by Weine Olovsson

- **GPU node:**

382 x • 4 x Nvidia GH200 Grace Hopper Superchip

- H100 GPU
- 72c Arm Neoverse v2 cores


- **Quick test: GaAsBi 512 atoms 4 k-pts, VASP PBE**

Tetralith has 1844 Nodes • 1 x Tetralith CPU node: 3 084 s [8 nodes -> 471 s]

- 1 x A100 @Berzelius: 704 s (x4.4)

Arrhenius will have 1528
GH200

- 1 x GH200 @Dardel(testing): 404 s (x7.6)
(Same as Arrhenius)



Arrhenius: GPU-part

- **Quick test: Fe 128 atoms 8 k-pts, VASP MD non-collinear**

Tetralith has 1844 Nodes

- 1 x Tetralith CPU node: 101 303 s

- 1 x A100 @Berzelius: 10 390 s (x9.8)

Arrhenius will have 1528
GH200

- 1 x GH200 @Dardel(testing): 7 482 s (x13.5)

- **Quick test: SiC 576 atoms 8 k-pts, VASP PBE**

- 1 x GH200 @Dardel(testing): runs (others OOM)
(Same as Arrhenius)

MIMER - AI-factory



Hardware parameters

Procurement discussions with EuroHPC underway.

Cloud first.

Sensitive data ready.

GPU for AI training and inference.

Massive object storage, fast NVMe cache.

MIMER - AI-factory



-
- Procurement and recruitment ongoing
 - Dir & executive board in place
 - Jul 2026 Expected system delivery.
 - Oct 2026 Expected system operational.

Timelines

Director and executive board in place.

Recruiting 52 FTE AI experts, kick-started by ENCCS, NAISS, and affiliated parties.

Procurement discussions with JU underway.

Expected hardware delivery in Jul 2026.

Expected operational in Oct 2026.



Questions?

Erik Lindahl, Director <erik.lindahl@naiss.se>

Björn Alling, co-Director <bjorn.alling@naiss.se>

NAISS