

# PReSTO in 2025!

Rapid **PR**ot**Ein** **ST**ructure **cO**mputation in an integrated **e-S**cience environment



*Cryo-EM* National facility



Ana Gonzales



Daniel Larsson



Tapu Shaikh



Tim Schulte



Karl Hörnell



Marcos Acebes



Maria Sunnerhagen



Zdenek Matej



Jie Nan



Stefan Fleischmann



Piotr Draczkowski



Sebastian Thorarensen



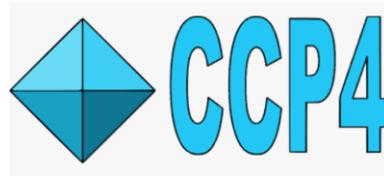
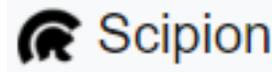
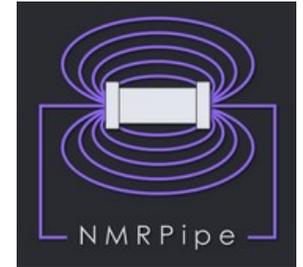
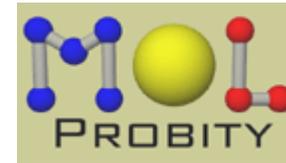
Torben Rasmussen



Björn Forsberg

More enablers 2013-2024 at <https://www.nsc.liu.se/support/presto/enablers/>

# Applications added to PReSTO



# Structural biology branches and PReSTO applications

SAXS	SSX/XFEL	micro-ED	MX	Cryo-EM	NMR
Chimera(X)	Chimera(X)	Chimera(X)	Chimera(X)	Chimera(X)	Chimera(X)
PyMOL	PyMOL	PyMOL	PyMOL	PyMOL	PyMOL
AlphaFold	AlphaFold	AlphaFold	AlphaFold	AlphaFold	AlphaFold
	LigPlot	LigPlot	LigPlot	LigPlot	LigPlot
	COOT	COOT	COOT	COOT	ARIA
	MolProbity	MolProbity	MolProbity	MolProbity	MolProbity
ATSAS	CCP4	CCP4	CCP4	CCPEM	CCPN
SasView	Phenix	Phenix	Phenix	Phenix	MDDNMR
	CrystFEL	(CoSEDA...soon)	CNS	cisTEM	CNS
		AutoLEI	XDS	RELION	CYANA
			XDSAPP3	CryoSPARC	NMRpipe
			SHARP	cryoDRGN	
			DIALS	Scipion	
	CSD (Grade2)	CSD (Grade2)	CSD (Grade2)	CSD (Grade2)	

# HPC centers for structural biology

<b>NSC Tetralith</b>	<b>NSC Berzelius</b>	<b>LUNARC Cosmos</b>	<b>MAX IV cluster</b>	<b>Euro-HPC Arrhenius</b>
MX	MX as EM dependency	MX	MX	MX as EM dependency
SSX/XFEL	NO	SSX/XFEL	SSX/XFEL	NO
<sup>1</sup> Cryo-EM too weak GPU	<sup>2</sup> Cryo-EM if KAW funded	<sup>3</sup> Cryo-EM on full nodes	NO	<sup>4</sup> Recompile Arm CPU
SAXS	NO	SAXS	SAXS	NO
NMR	NO	NMR	NMR	NO
Micro-ED	NO	Micro-ED	Micro-ED	NO

<sup>1</sup>NSC Tetralith (2018) Tesla T4 GPU nodes, 96 GB RAM, and 2 TB NVMe not sufficient in 2025

<sup>2</sup>NSC Berzelius (2021) is only accessible to KAW funded researchers => more on next slide

<sup>3</sup>LUNARC Cosmos (2022) performs well in testing by Cryo-EM researchers using an entire node (expensive solution)

<sup>4</sup>ARM CPUs in Arrhenius (2026) requires Cryo-EM software recompilation

# NSC Berzelius and the Cryo-EM use case

## + PLUS

Berzelius hardware: GPUs, NVMe:s, etc.

NSC staff

Avoid job termination by "--reservation = safe" that runs to completion but with longer queue times

✓ Node sharing by Multi-Instance GPU (MIG) lanes

- ✓ 2g.20gb, 20 GB VRAM, 3 X share GPU, OK but not selected!
- ✓ 1g.10gb, 10 GB VRAM, 7 X share GPU, too small but chosen!

## - MINUS

a) Berzelius compute allocations has too little disk-space for Cryo-EM raw data => Separate storage application required.

b) GPU Usage Efficiency Policy => less GPU intense jobs terminates

c) Berzelius is in high demand

1g.10gb MIG lane too small for Cryo-EM

Only KAW funded researchers allowed to use NSC Berzelius

# In 2026 add PReSTO to EuroHPC Arrhenius

Feature	Berzelius	Arrhenius
<b>GPU Architecture</b>	A100 (Ampere) + H200 (Hopper)	GH100 Hopper in Grace Hopper Superchip
<b>GPU Memory per Unit</b>	40–80 GB (A100), 141 GB (H200)	192 GB HBM3(e)
<b>CPU<sup>1</sup></b>	AMD EPYC or Intel Xeon	NVIDIA Grace CPU (72 Arm cores)
<b>Memory Model<sup>2</sup></b>	Separate CPU/GPU memory spaces	Unified, coherent CPU–GPU memory
<b>Interconnect</b>	HDR/NDR InfiniBand	NVLink-C2C (900 GB/s) + NVLink Switch
<b>Node Composition</b>	8 GPUs per DGX node	1 GPU + 1 CPU per superchip
<b>Total GPUs</b>	~880 (A100 + H200 combined)	1,536 Hopper GPUs

<sup>1</sup> PLUS: Arm CPUs have lower power consumption vs x86-64

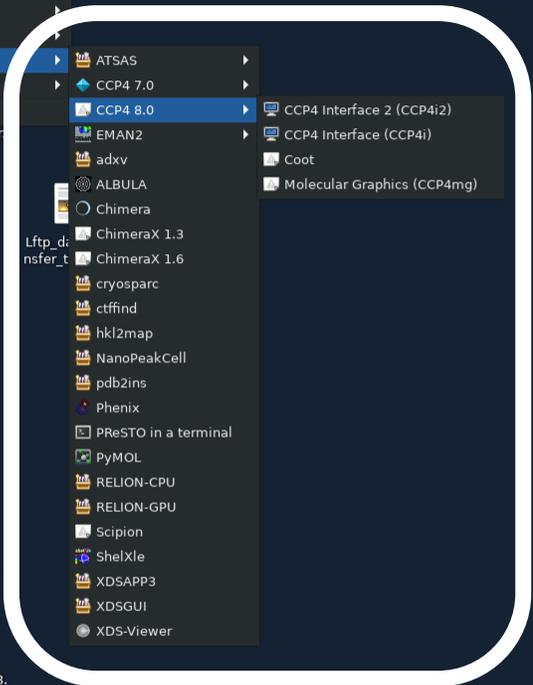
MINUS: Some CryoEM software to be recompiled from source with Arm compatible toolchain, compilers, dependencies...

<sup>2</sup>Arrhenius CPU and GPU shares memory (fast)

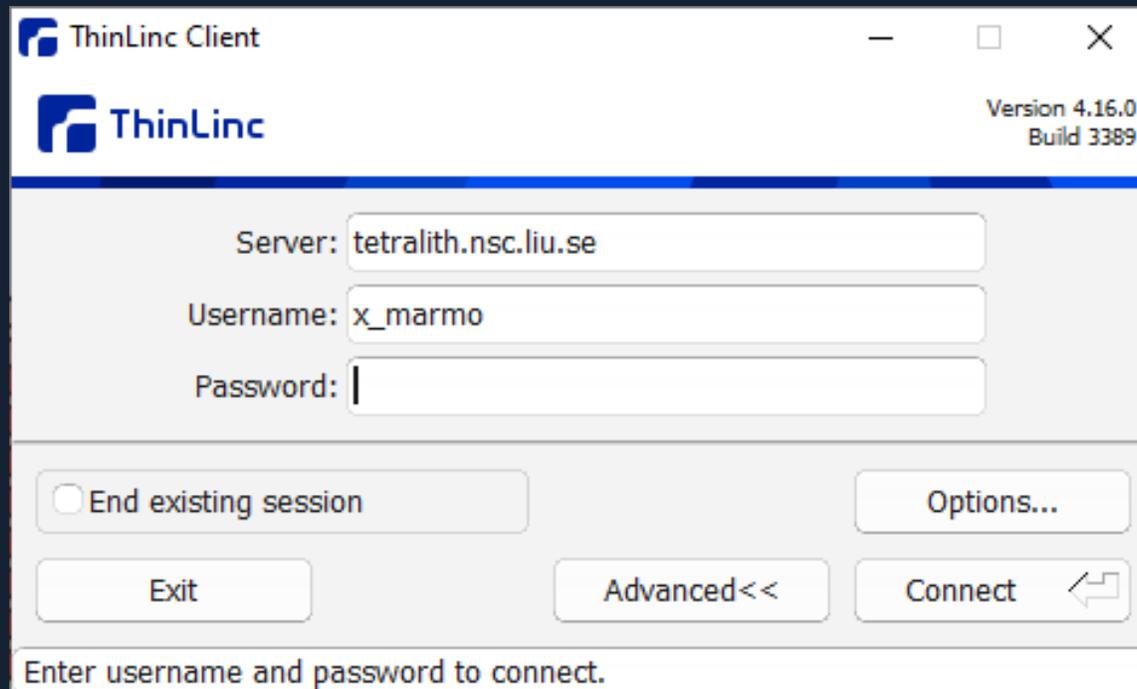
Berzelius CPU RAM separate from GPU VRAM (slower)

# Thinlinc Remote Desktop

- Available from any PC, Linux, MAC computer



PReSTO desktop menu  
-software with GUIs



# Use of old software versions

Terminal - x\_marmo@b...

Run Program...  
Terminal Emulator  
File Manager  
Mail Reader  
Web Browser  
Settings  
Administration  
Accessories  
Development  
Graphics  
Internet  
Office  
Other

PreSTO menu  
load default modules

System  
Log Out  
graphics\_berz elius.txt  
TM.txt  
Lftp\_d... nsfer\_t...  
hkl2map  
README  
Old Firefox Data  
NC1-2024.03. workshop-0...  
Lftp-Umeå

- ATSAS
- CCP4 7.0
- CCP4 8.0
- EMAN2
- adsv
- ALBULA
- Chimera
- ChimeraX 1.3
- ChimeraX 1.6
- cryosparc
- ctffind
- hkl2map
- NanoPeakCell
- pdb2ins
- Phenix
- PreSTO in a terminal
- PyMOL
- RELION-CPU
- RELION-GPU
- Scipion
- ShelXle
- XDSAPP3
- XDSGUI
- XDS-Viewer

CCP4 Interface 2 (CCP4i2)  
CCP4 Interface (CCP4i)  
Coot  
Molecular Graphics (CCP4i)

mochma@cosmos2:~

```
[mochma@cosmos2 ~]$ module avail Phenix
```

----- /sw/pkg/presto/modules/all -----	
Phenix/1.19.2-4158-Rosetta-3.10-PreSTO-11.4	Phenix/1.21.1-5286-Rosetta-3.10-PreSTO-11.0
Phenix/1.20.1-4487-Rosetta-3.10-PreSTO-9.0	Phenix/1.21.1-5286-Rosetta-3.10-PreSTO-11.1
Phenix/1.20.1-4487-Rosetta-3.10-PreSTO-9.2	Phenix/1.21.1-5286-Rosetta-3.10-PreSTO-11.2
Phenix/1.20.1-4487-Rosetta-3.10-PreSTO-9.3	Phenix/1.21.1-5286-Rosetta-3.10-PreSTO-11.3
Phenix/1.20.1-4487-Rosetta-3.10-PreSTO-9.4	Phenix/1.21.1-5286-Rosetta-3.10-PreSTO-11.4 (D)
Phenix/1.21-5207-Rosetta-3.10-PreSTO-10.0	

Where:  
D: Default Module

Phenix 1.19.2 for joint X-ray/Neutron refinement

Terminal - x\_marmo@tetralith2:~

```
[x_marmo@tetralith2 ~]$ module avail PyMOL
```

----- /software/presto /modules/all -----		
PyMOL/2.1.0-2-PreSTO	PyMOL/2.5.0-PreSTO-9.6	PyMOL/2.5.0-2-PreSTO
PyMOL/2.3.0-1-PreSTO	PyMOL/2.5.0-PreSTO-10.0	PyMOL/2.5.0-3-PreSTO
PyMOL/2.3.0-2-PreSTO	PyMOL/2.5.0-PreSTO-11.0	PyMOL/2.5.0-4-PreSTO
PyMOL/2.5.0-PreSTO-toolchain-foss-2021a	PyMOL/2.5.0-PreSTO-11.1	PyMOL/2.5.0-5-PreSTO
PyMOL/2.5.0-PreSTO-8.0	PyMOL/2.5.0-PreSTO-11.2	PyMOL/2.5.0-7-PreSTO
PyMOL/2.5.0-PreSTO-9.0	PyMOL/2.5.0-PreSTO-11.3	
PyMOL/2.5.0-PreSTO-9.4	PyMOL/2.5.0-1-PreSTO	

----- /software/sse2/tetralith_el9/modules -----	
PyMOL/recommendation (D)	PyMOL/2.5.0-hpcl-gcc-2022a-eb

Where:  
D: Default Module  
[x\_marmo@tetralith2 ~]\$

Use of ATSAS-PyMol plugin called sas.py

# In 2025 LUNARC Cosmos for national Cryo-EM use?

Feature	Berzelius thin/fat DGX-A100 Node	Cosmos Intel/AMD NVIDIA A100 Node*
GPUs per node	8× NVIDIA A100 40/80 GB (thin/fat) HBM2	1× NVIDIA A100
GPU Memory	40/80 GB per GPU for thin/fat Node	40 GB per GPU
CPU	Dual AMD EPYC, 128 cores	32/48 cores for Intel/AMD Node
RAM	~1/2 TB for thin/fat Node	384/512 GB for Intel/AMD Node
Local Storage	15/30 TB NVMe for thin/fat Node	1/2 TB SSD for Intel/AMD Node
Number of Nodes	60 thin and 34 fat nodes	6 AMD and 4 Intel NVIDIA A100

\*Cosmos also have 6 AMD NVIDIA A40 and 5 Intel NVIDIA A40 nodes without MIG support

LUNARC Cosmos testing of Cryo-EM software

+ a single CryoSPARC instance per compute project.

=> PI master instance overview like on standard workstation

+ RELION and CryoSPARC runs ok, CryoDRGN not tested

- Use of a **full Cosmos node** (no MIG lane sharing) considered an expensive solution

**THE END**