



CLAIX-2025 Overview

Hardware and Software Setup

Dr. Marc-André Hermanns, Philipp Martin

Hardware of CLAIX 2025

CLAIX-2025 Nodes Types

	HPC	ML
CPU	2x AMD EPYC 9654 (Genoa), 96 cores each	2x AMD EPYC 9535 (Turin), 64 cores each
Memory	384 GB (317 nodes) 786 GB (104 nodes)	786 GB (20 nodes)
GPU	– none –	4x NVIDIA H100 (80 GB HBM3, NVLink)
Interconnect	Omni-Path Cornelis CN5000	
Storage	1.6 TB SSD / node	

- Different CPU types on HPC and ML nodes
 - Be aware when reusing Slurm scripts!
- Two different memory configurations on HPC nodes
 - Small configuration with 2 GB per core
 - Medium configuration with 4 GB per core
- GPUs on ML Nodes have slightly smaller memory than on CLAIX-2023
 - HBM3 has higher memory bandwidth
- **Focus of this workshop is the HPC segment**

CLAIX-2025 Nodes Types

Hostname	Node Type	Recommended Use
login25-1.hpc.itc.rwth-aachen.de	HPC	General Use & CPU-centric Development
login25-2.hpc.itc.rwth-aachen.de	HPC	
login25-3.hpc.itc.rwth-aachen.de	HPC	
login25-x-1.hpc.itc.rwth-aachen.de	HPC	FastX-based Use & Development
login25-g-1.hpc.itc.rwth-aachen.de	ML	GPU-enabled Development

- Use the default login nodes for standard development work
- Use the X login node for FastX connections
- Use the G login node for GPU development

Partition Setup of CLAIX-2025

CLAIX-2025 Nodes Types

Partition (-p / --partition)	Config	Recommended Use
c25ms	Small Memory	General Use / Default Partition
c25ml	Large Memory	Larger per core demands
c25g	GPU	GPU-based computing

- As with CLAIX-2023 (-low) handling is automatic depending on available quota
- Only select a partition if default (small memory) is insufficient
 - More small-memory nodes available
 - Large-memory nodes are chosen automatically if all small-memory nodes are occupied

Workshop Reservation

- **Reservation aixcelerate**
 - Short turn-around time during the workshop
 - 10 HPC nodes total for all teams
- **Compute Budget lect0200**
 - Check access to compute budget by UNIX group affiliation

```
$ groups
.... lect0200 ...
$ sbatch --reservation=aixcelerate -A lect0200 <jobscript.sh>
...
$ salloc --reservation=aixcelerate -A lect0200 -N <#nodes> -n
<#procs> ...
```

Software and Module System on CLAIX 2025

One Cluster – Multiple Software Stacks

- **Managed with EasyBuild**
 - Mostly same toolchains available as on CLAIX-2023
 - Focus on newer toolchains (intel/2024a and foss/2024a and above)
- **Software Modules based on Lmod**
- **Module Tree is separate from CLAIX-2023**
 - Different Software Versions may be installed for either system
 - Contact servicedesk@itc.rwth-aachen.de if specific software is missing/needed

Workshop Environment

- **Default environment (after login) is intel/2024a (same as CLAIX-2023)**
 - Consider default toolchains like LTS toolchains
- **Use 2025b toolchains for workshop**
 - Newer toolchains may improve software stack performance
 - Test and move to newer toolchains on regular intervals
 - Newest AMD Software installed for GCCcore/14.3.0

```
$ module purge  
$ module load intel/2025b  
...  
  
$ module purge  
$ module load foss/2025b
```

AMD-specific Software Modules

- **AMD Optimized CPU Compiler**
 - AOCC/5.1.0 installed for both toolchain families
 - Optimizers C/C++/Fortran Compilers
- **AMD Optimized CPU Libraries**
 - AOCL-BLAS/5.1 installed only for GCC toolchain » foss/2025b
 - Optimized BLAS Libraries for CPU

```
$ module purge  
$ module load foss/2025b  
$ module load AOCC  
[ ... ignore dependency warning for now ... ]  
$ module load AOCL-BLAS
```

AMD-specific Performance Tool Installations

- **AMD uProf**
 - AMD - uProf/5.2 installed independent of toolchains
 - Sampling-based profiling tool
- **Score-P 10.1**
 - Installation for AOCC available in foss/2025b toolchain

```
$ module load AMD-uProf  
$ module load Score-P/10.1-AOCC-5.1.0-CUDA-12.9.1
```

Lmod – Short Recap


- **module purge**
 - Unload all loaded modules
- **module load**
 - Load specified modules and dependencies
- **module unload**
 - Unload specified modules
- **module spider**
 - Search for modules (case-insensitive, partial name allowed)

Storage and I/O Technologies for CLAIX 2025

Storage Overview CLAIX-25

Access	Filesystem	Capacity	# Files	Backup/Snaps	Pro	Con
\$HOME	GPFS	250 GiB	1 mio.	Tape (off-site)	backup	quota
\$WORK	GPFS	250 GiB	3 mio	Snapshots	snapshots	quota
\$HPCWORK	Lustre	1 TiB	1 mio.	None	Bandwidth	Metadata performance
\$BEEOND	BeeGFS	1.4 TiB / Node	unlimited	None	performance	temporary
\$TMP	XFS	1.4 TiB / Node	unlimited	None	performance	not shared & temporary

Filesystem Choice on CLAIX-25



- **Additional Information**

- If you need more quota on a specific filesystem, open a support ticket and we will help you
- BeeOND size scales with number of nodes, pay attention to this in your jobs
- If your application does not do a significant amount of I/O, default to \$WORK or \$HPCWORK

- **Differences from CLAIX-23**

- \$BEEOND now has the same SSD size for CPU and GPU nodes (1.4 TiB/Node)
- NVIDIA GPU Direct Storage is now available and may be used by e.g. pyTorch applications
- Lustre access needs to pass over gateway nodes to switch from Omnipath to Infiniband

Striping settings on CLAIX-25

- **Striping**
 - Across how many targets files are split up
 - More targets can improve throughput but add some latency
 - For advanced / very specific usage scenarios, **contact us** to figure out striping settings
 - If you know what you are doing, you can set the striping yourself for \$HPCWORK and \$BEEOND
- **Lustre**
 - 0 - 1 GiB unstriped
 - 1 - 64 GiB striped across 4 nodes (16 MiB stripe size)
 - 64+ GiB striped across 16 nodes (16 MiB stripe size)
 - `lfs setstripe` (https://wiki.lustre.org/Configuring_Lustre_File_Striping)
- **GPFS**
 - Striped across 2 disks with 8 MiB stripe size
 - cannot be adjusted by users
- **BeeOND**
 - Striped across 4 nodes (if available), with 512 KiB stripe size
 - `beegfs --tls-disable` entry set (https://doc.beegfs.io/latest/advanced_topics/striping.html)

Thank you for your attention.