

5th HPC Café

09.05.2025



Agenda HPC Café – 09.05.2025

1. Updates from the operations

- Status Procurement “HoreKa-2”
- Commissioning of bwUniCluster 3.0

2. News about NHR@KIT

- AI fast track

3. Questions and Answers

1. Updates from the operations

Samuel Braun (SCC)




Status Procurement “HoreKa-2”

■ Procurement

- Participation competition: Ongoing
- Bidding competition: Starting end of May
- Commissioning of phase 1: End of 2025
- Full commissioning: 2026
- Setup in the HPC Hall at North Campus
 - Power distribution infrastructure Upgrade in the works



HoreKa-2



Procurement Considerations: Update

- Pragmatic approach for CPU to GPU split
 - “Freezing” of the total performance of the CPU partition to HoreKa-Blue
 - Performance: ~ 2.5 PFLOP/s
 - **Bandwidth:** ~ 200 TB/s
 - „Backfill“ with GPU nodes (plus filesystem)

Commissioning of bwUniCluster 3.0 + KIT-GFA-HPC 3

- In operation since:
07.04.2025
- bwUniCluster 3.0
 - bwUniCluster 3.0 Phase 1
 - bwUniCluster 2.0 Extension

Key Numbers

- ~ 7,1 PFLOP/s
- ~ 27.000 CPU cores
- 112 GPUs




bwUniCluster 3.0: System Overview

■ bwUniCluster 3.0 Phase 1

- Hot water DLC, Campus North
- New Filesystem
- 2x Login node
- 74x CPU nodes AMD EPYC 9454
- 12x NVIDIA GPU nodes (4x H100)
- 1x AMD APU node MI300A

■ bwUniCluster 2.0 Phase 2

- Cold water DLC, Campus South
- 272x CPU nodes Intel Xeon Platinum 8358
- 15x NVIDIA GPU nodes (4x A100/H100)



bwUniCluster 3.0: Summary of Changes

Entitlement, Registration, Login and Migration

- Existing entitlements for UC2 valid
- New registration for UC3
 - Service password
- Server address: **uc3.scc.kit.edu**
- **KIT** users: **ka_** prefix in front of username
 - Also affects HOME directory path: `/home/ka/ka_OE/ka_ab1234`
 - Login: `ssh ka_ab1234@uc3.scc.kit.edu`
- New filesystem: Migration of data!

bwUniCluster 3.0: Summary of Changes

Policy

■ New quota limits

- HOME: 500GB, 5 million files (inodes)
- Workspace: 40TB, 20 million files (inodes)

■ Throttling

- The **maximum amount of cores** used at any given time from jobs running is 1920 per user (aggregated over all running jobs)

■ Access for KIT students

- Within **lecture project**: regular u-student account
- **Else**: GuP (Gäste und Partner) account (BA/MA **thesis**, HiWi, ...)

bwUniCluster 3.0: Summary of Changes

Software

- Red Hat Enterprise Linux (RHEL) 9.4
- Conda users: devel/miniconda → devel/miniforge
- EasyBuild „official beta-test“
 - Complements standard Lmod modules
 - Curated and own software inside HOME
 - Build your own HPC software stack with ease
 - Amazing!

bwUniCluster 3.0: The starting phase

■ Good

- System runs stable
- Higher performance @ lower power draw
- Data migration went exceptionally smoothly


■ Bad

- Some software modules still missing
- System got much smaller (nb. of nodes)
- Longer queueing times, fewer nodes

■ Ugly

- „I can't login ...“
- Introduction of a new support portal at the same time...

■ Announcements + Wiki

- Migration Guide: Success 
- Not all problems were anticipated
- Continuous improvement of Wiki

→ Feedback welcome* 

*) <https://bw-support.scc.kit.edu> → Feedback → bwHPC Wiki

2. News about NHR@KIT

AI fast track



Project categories at NHR@KIT and HoreKa

Project category	Duration	Min. Requested Resources	Max. Requested Resources	Review	Call
NHR Normal	1 year, extendable	2.000.000 CPUh / 30.000 GPUh	14.999.999 CPUh / 199.999 GPUh	2x Scientific	Rolling

NHR Large	1 year, extendable	15.000.000 CPUh / 200.000 GPUh	70.000.000 CPUh / 1.000.000 GPUh	2x Scientific + NHR Steering Committee	Quarterly
-----------	--------------------	--------------------------------	----------------------------------	--	-----------

Project category	Duration	Effectively Granted Resources	Review	Call
NHR Test	6 months, not extendable	500.000 CPUh / 5.000 GPUh	Technical	Rolling

NHR Starter	1 year, not extendable	360.000 CPUh / 10.000 GPUh	Technical	Rolling
-------------	------------------------	----------------------------	-----------	---------

AI Fast Track	1 year, extendable	380.000 CPUh / 20.000 GPUh	None	Rolling
---------------	--------------------	----------------------------	------	---------

AI fast track at NHR@KIT

- Aim: enabling an easy and fast access to compute resources for AI projects
- You manage the project and project members via FeLS
- Currently no new AI fast track projects are accepted
- All information also on <https://www.nhr.kit.edu/application.php>

3. Questions and Answers



Invitation to an KIT internal event

- 2. Workshop on **Research Software Engineering** at KIT
- **When?** – 7. Juli 2025, 13:00-16:15, with a get-together in relaxed atmosphere afterwards
- **Where?** – Campus Nord, Geb. 449, Room 140
- **Who?** – Everyone at KIT, who develops research software
- **What?** – Presentations of the results of the KIT-internal „Leitprojektes 9.1.4“ in the scope of the digitalization strategy
- **How?** – Program, Registration: <https://indico.scc.kit.edu/event/4819/>
- More information on RSE@KIT: <https://www.rse-community.kit.edu/>