

bwHPC and NHR: Concepts, Infrastructures and User Support

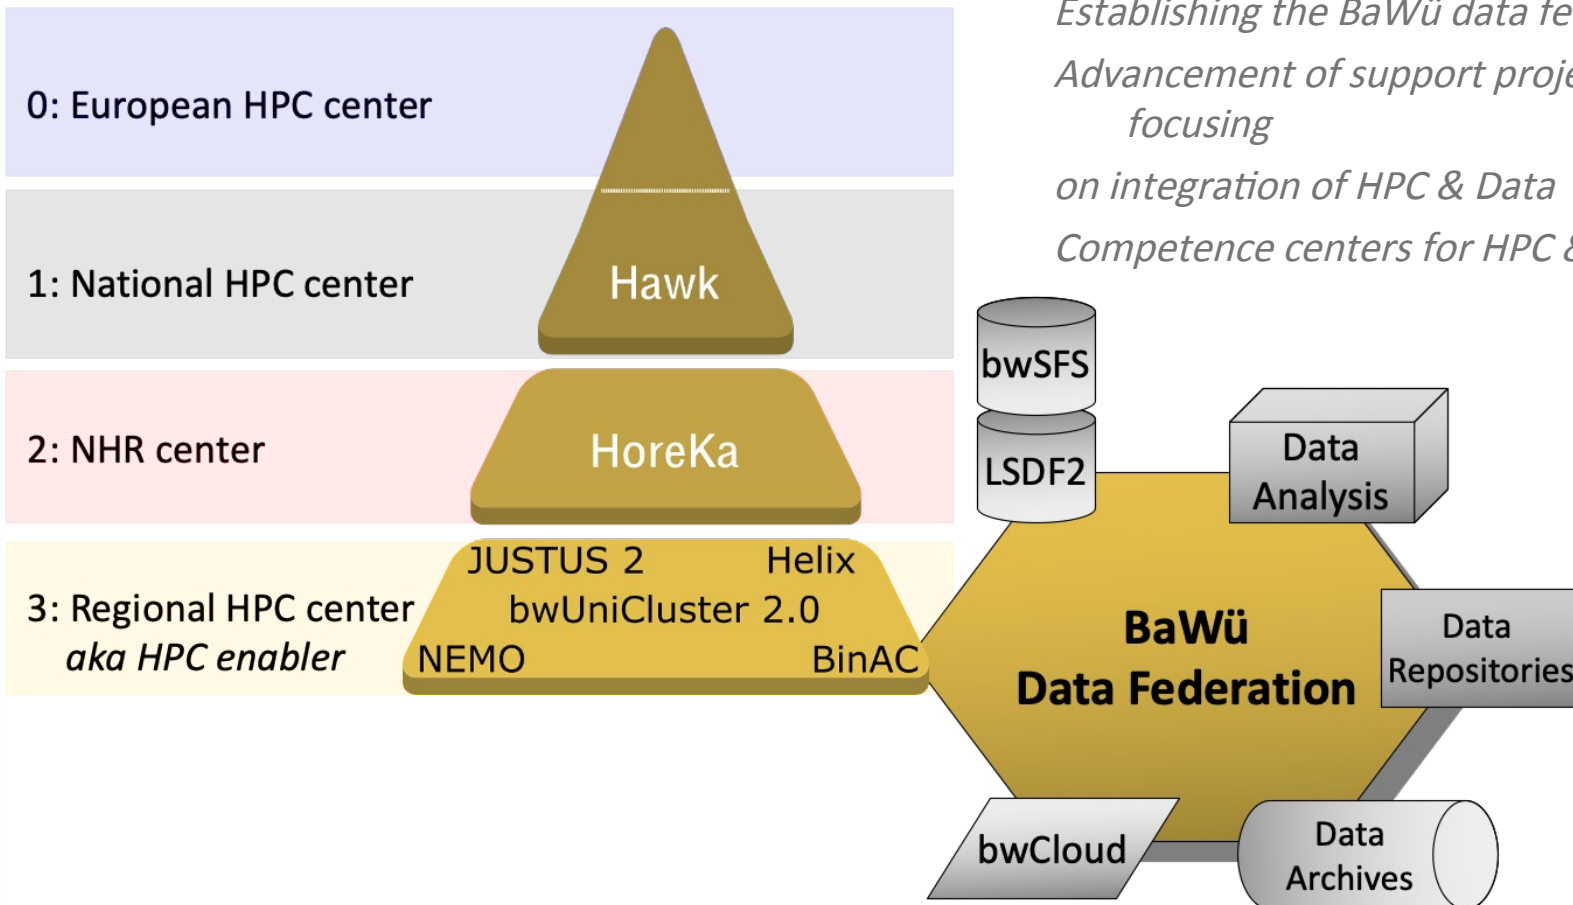
Robert Barthel, SCC, KIT



bwHPC

Baden-Württemberg's implementation strategy for HPC, Data Intensive Computing & Large Scale Scientific Data Management

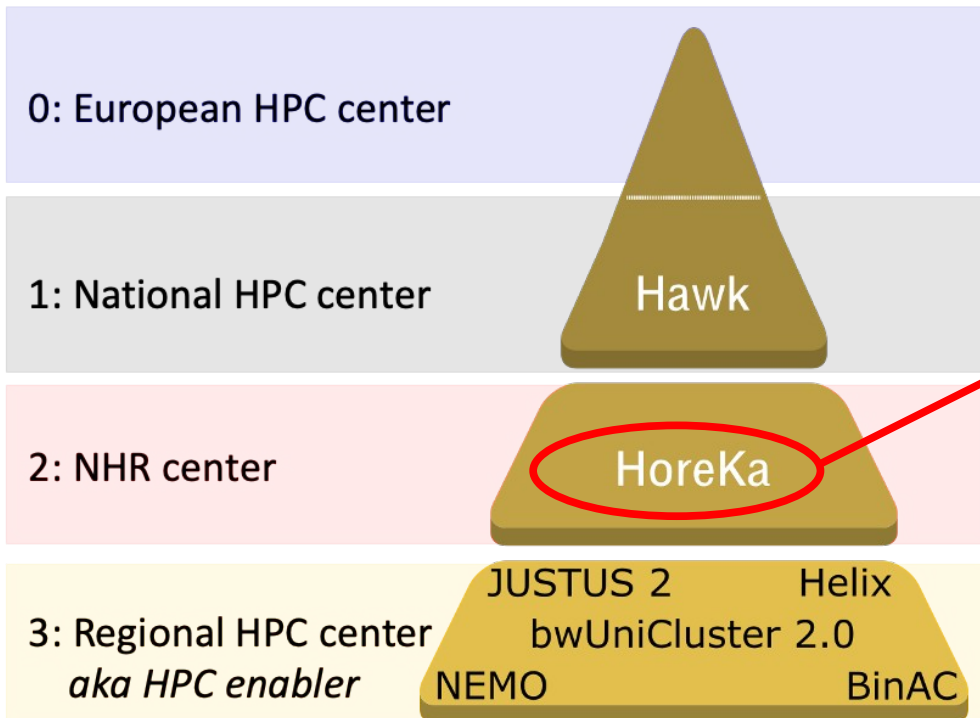
Advancement of federated HPC@tier3
Establishing the BaWü data federation
Advancement of support project focusing on integration of HPC & Data
Competence centers for HPC & Data



NHR (1)

National High Performance Computing

HPC in Baden-Württemberg



National HPC at Tier 2

Centers	Universities
NHR4CES@RWTH	RWTH Aachen
NHR4CES@TUDa	TU Darmstadt
NHR@FAU	Univ. Nürnberg-Erlangen
NHR@Göttingen	GWVG + Univ. Göttingen
NHR@KIT	KIT
NHR@TUD	TU Dresden
PC2	Univ. Paderborn
NHR@SW	Univ. Frankfurt a.M., Mainz, Kaiserslautern- Landau, Mainz, Saarland

NHR (2)

- „**Nationales Hochleistungsrechnen**“ (NHR) replaces the federal 91b DFG funding for academic Tier 2 HPC systems
- **625 million Euros over 10 Years** for 9 NHR centers
- Coordination of Applications, Method, Hardware and Operational competencies

07.01.2020
Publication of
NHR Call

13.11.2020
Selection of
the 8 centres

30.07.2021
Inauguration
HoreKa

23.08.2021
Foundation of
NHR Association



Details of bwHPC

bwHPC: HPC & Data Facilities

LSDF 2 (SDS@hd/ LSDFOS)

Hot Data: Life Sciences, Medical Science,
Hydro-mechanics, Humanities,
Astrophysics, Scientific Computing / GP

bwForCluster Helix (08/2022):

Structural and Systems Biology, Medical
Science, Comp. Humanities & Soft Matter
Successor of: **MLS&WISO**

bwUniCluster 2.0 (03/2020):

General purpose,
Teaching & Education

bwForCluster JUSTUS 2 (07/2020):

Computational Chemistry,
Quantum Sciences

bwDataArchive (KA)

Cold Data: General purpose

bwSFS 2.0 (ST/HO)

bwForCluster NEMO (09/2016):


NeuroSc, Micro Systems Eng., Elementary
Particle Phys.,
Successor: **NEMO 2**, approved by DFG

bwSFS (FR/TU)


Bioinformatics, Elementary
Particle Physics, Micro
Systems Eng., NeuroSc.,
Systems Biology

bwForCluster BinAC (11/2016):

Bioinformatics, Astrophysics, Geosciences
Successor: **BinAC 2**, approved by DFG



bwUniCluster 2.0



■ Shareholders:

- Freiburg, Tübingen, KIT, Heidelberg, Ulm, Hohenheim, Konstanz, Mannheim, Stuttgart, and HAW BW e.V. (an association of university of applied sciences in Baden-Württemberg)
- Baden-Württemberg's ministry of science, research and arts (MWK)

■ Access:

- For all members of **shareholder**'s university in BW
 - For all members of the universities of applied sciences in BW
- How?** → Entitlement → Registration → Questionnaire → Login
(Details: 2. Talk today)

■ Usage:

- Free of charge
- For general purpose, teaching & education
- For technical computing (sequential & weak parallel) & parallel computing

4x bwForCluster



- Shareholders
 - German Research Society (DFG)
 - Baden-Württemberg's ministry of science, research and arts
- Access (*Details: 2. Talk today*):
 - All university members in Baden-Württemberg
 - How? → Entitlement +
Compute Project Proposal →
Registration → Login
- Usage:
 - Free of charge
 - **Approved** Compute Project Proposal to **only 1 bwForCluster** matching cluster's subject fields

bwForCluster JUSTUS 2

(07/2020):

Theoretical Chemistry, Condensed matter physics, Quantum physics

bwForCluster Helix

(08/2022):

Structural and Systems Biology, Medical Science, ...

bwForCluster BinAC

(11/2016):

Bioinformatis, Astrophysics, ...

bwForCluster NEMO

(09/2016):


Neurosciences, Micro Systems Engineering, Elementary Particle Physics, ...

Scientific Simulation and Storage Support Services (bwHPC-S5)

HPC & Data Fed. @ BaWü

Goal?

- Bridging science & HPC
- Bridging HPC tiers and Large Scale Scientific Data Facilities (LS²DM)
- Embedding services

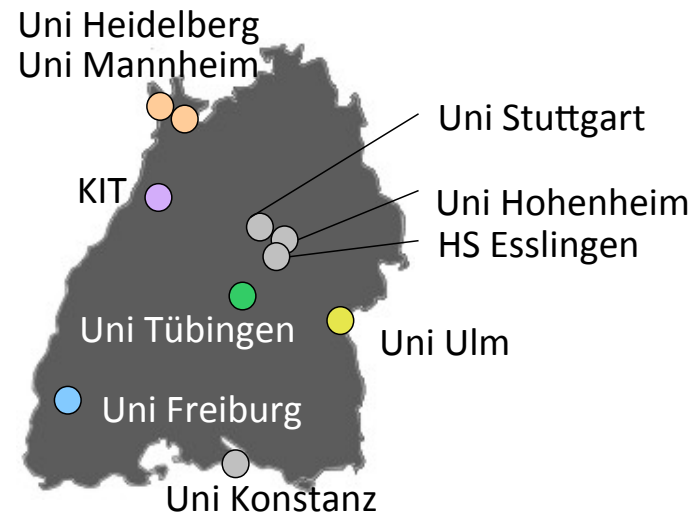


Where are these competence centers?

→ Organised by 8 universities

Who ?

→ Experts from 10 universities



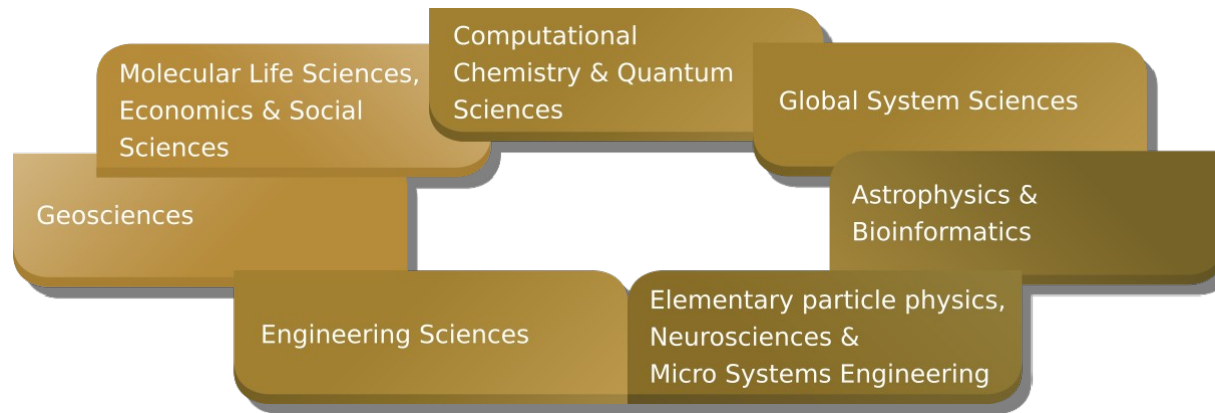
bwHPC-S5: What kind of support?

- Seminars and workshops
 - coordinated by 10 BaWü universities
- Documentation + best practices → wiki
- HPC and Data Competence Center:
 - Coordination of tiger teams
 - Help concerning: Code/workflow adaptation, porting and parallelization
 - Identify of user key topics
 - Help to access tier 2 (HoreKa) and 1 (Hawk)
 - Establish state-wide experts pool
- Providing/maintaining:
 - Community specific & HPC generic software and tools,
 - Data management tools for using e.g. [bwDataArchive](#) or [SDS@hd](#)



The screenshot shows the main page of the bwHPC Wiki. The page has a navigation bar with links for 'Main page', 'Discussion', 'Read', and 'View source', along with a search box. The main content area is titled 'Main Page' and includes a welcome message: 'Welcome to the bwHPC Wiki.' Below this, it states that bwHPC represents services and resources in the State of Baden-Württemberg, Germany, for High Performance Computing (HPC), Data Intensive Computing (DIC) and Large Scale Scientific Data Management (LS2DM). It provides the main website URL: <https://www.bwhpc.de/>. The page also features sections for 'Courses / eLearning' with links to 'eLearning and Online Courses' and 'Introduction to Linux in HPC (external resource)'. There is a 'Need Access to a Cluster?' section with a 'Registration' link. The 'HPC System Specific Documentation' section lists clusters: 'bwUniCluster 2.0' (General Purpose, Teaching), 'bwForCluster JUSTUS 2' (Theoretical Chemistry, Condensed Matter Physics, and Quantum Sciences), and 'bwForCluster' (Computational Humanities, Structural and Systems Biology, Medical Science, and Soft...

How to get support?




- Scientific specific support?
 - Choose your competence center → contact via [email](#), [trouble ticket](#)
 - Wiki: [best practice guides](#)
- Extensive support needed?
 - [Tiger team support](#)
- Cluster specific support?
 - Choose your cluster → [email list of cluster](#), [trouble ticket](#), [telephone](#)
- Complaints / policy issues?
 - [email](#) or [trouble ticket](#) @ project management
 - Contact your university member of the [LNA-BW \(User Steering Committee\)](#)

Details of NHR@KIT



NHR@KIT: HPC Facilities (1)



- **Budget:** 15 Mio. €
 - Procurement: Q3/19 – Q2/20,
Installation: Q3/20 – Q2/21
- 777 nodes:
 - **59.050 Intel CPU cores**
 - **668 NVIDIA A100 GPUs**
 - **17 PetaFLOPS Peak**
 - 16 PB Spectrum Scale, InfiniBand HDR
- 830 kW Warm Water Cooling (+Re-Use)
- **In 2021:**
 - **Green500 #13**
 - Top500 #53 (#220), EU Top 15



NHR@KIT: HPC Facilities (2)



	Standard/High-Mem	Extra-Large Mem	Accelerators
# Nodes	570+32	8	167
CPUs	2x Intel „Ice Lake“ 76 Core/152 Thread	2x Intel „Ice Lake“ 76 Core/152 Thread	2x Intel „Ice Lake“ 76 C./152 Thread
Memory	256/512 GB	4096 GB	512 GB
GPUs			4x NVIDIA A100
Local Disks	960 GB NVMe	7x 3,84 TB NVMe	960 GB NVMe
Interconnect	InfiniBand HDR200 Fat Tree		
Storage	16 PB, ~150 GB/s GPFS		

■ File systems: Quota limits, snapshots allowing fast restore of files

NHR@KIT: HPC Facilities (3)




CPU vs. GPU



Intel Ice Lake Xeon

10 nm, 38 Cores, 8x DDR4

Optimized for „General Purpose“
Average Floating Point Performance
Large memory with medium throughput
External Interconnect (InfiniBand)



NVIDIA Ampere A100

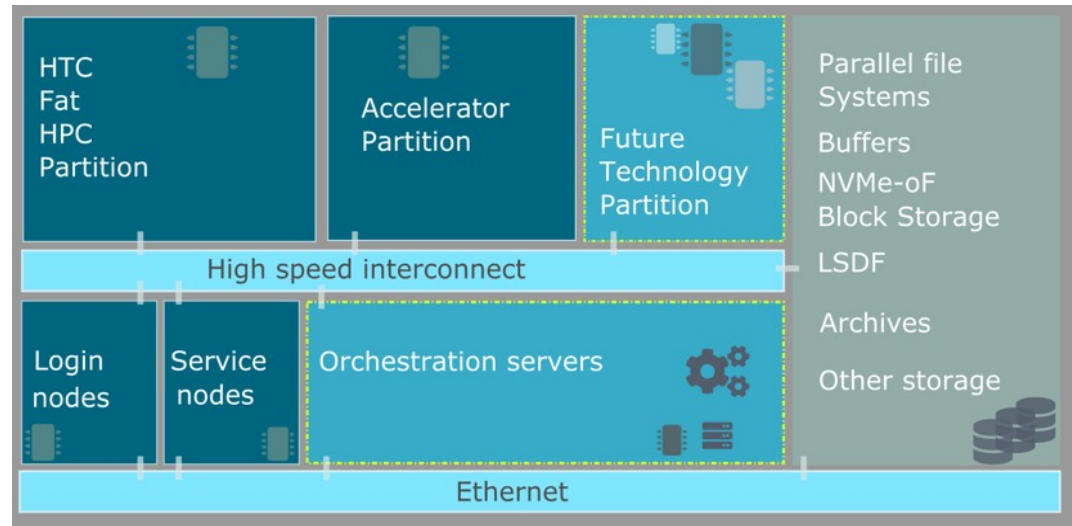
7nm, 6912 Cores, 6x HBM2

Optimized for highest performance
High floating point performance
Small, fast memory
Internal NVLink mesh for Multi-GPU

NHR@KIT: HPC Facilities (4)

Future Technology Partition

- Effective Support of scientific software development
- Porting to new hardware
- Acceleration of development cycles
- Development of efficient algorithms, libraries and applications

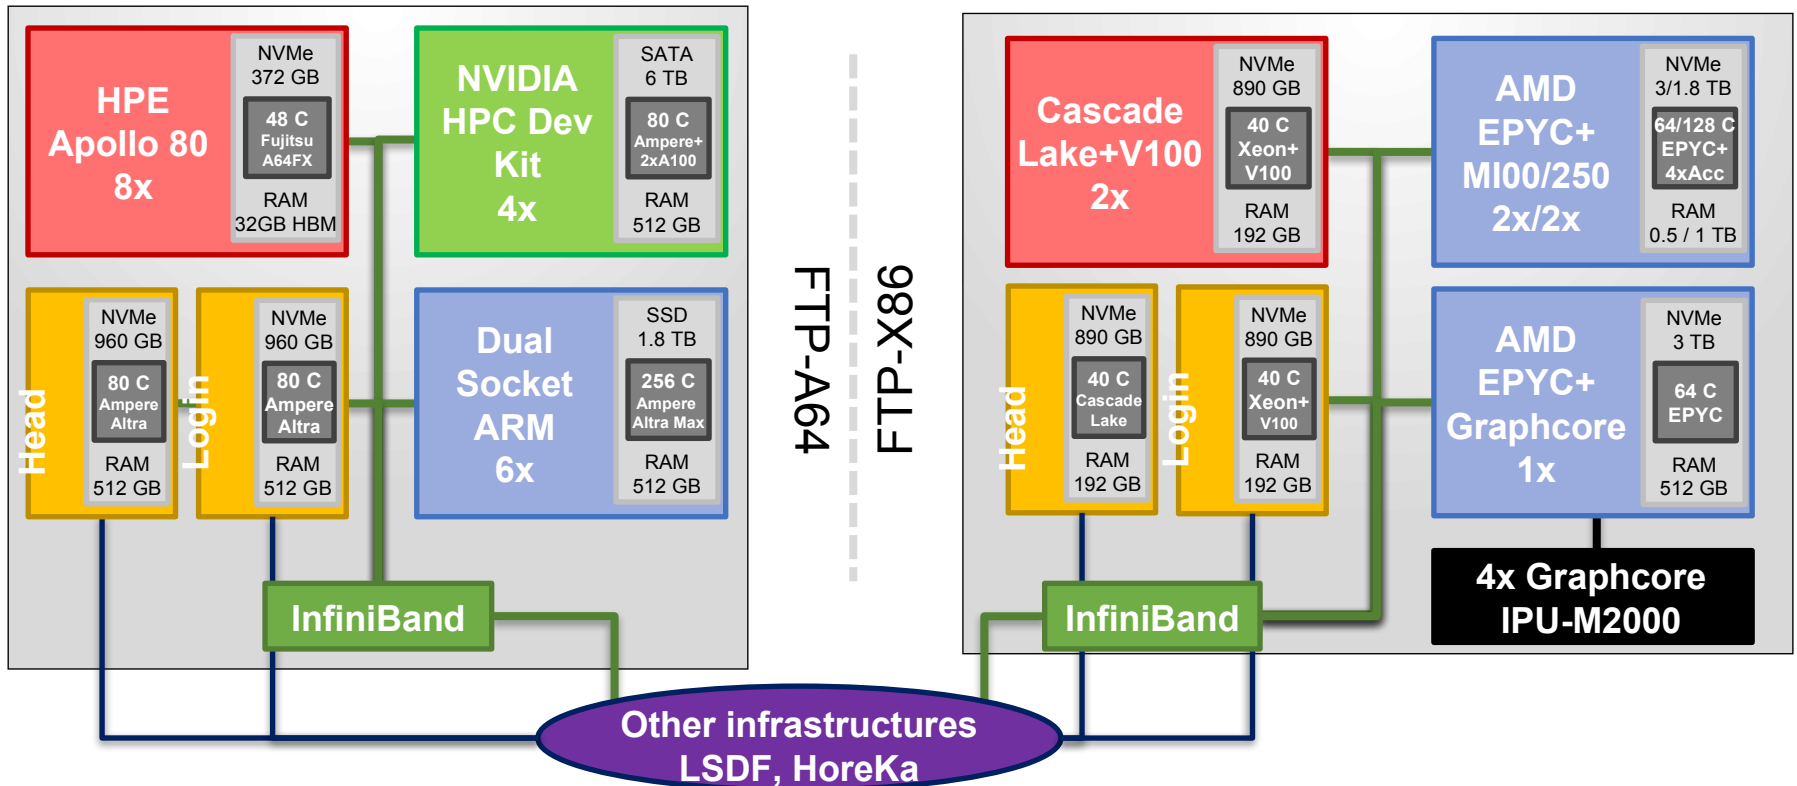


NHR@KIT: HPC Facilities (5)

Future Technology Partition

■ A64 ARM cluster

■ X86 cluster



NHR@KIT Resource Application

NHR@KIT offers four categories of projects depending on compute resources needed which implies different call and review types.

Project category	Duration	Min. Resources	Max. Resources	Review	Call
NHR Test	6 months, not extendable	- / -	500.000 CPUh / 5.000 GPUh	Technical	Rolling
NHR Starter	1 year	- / -	360.000 CPUh / 10.000 GPUh	Technical	Rolling
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

NHR@KIT: Support for users (1)

By

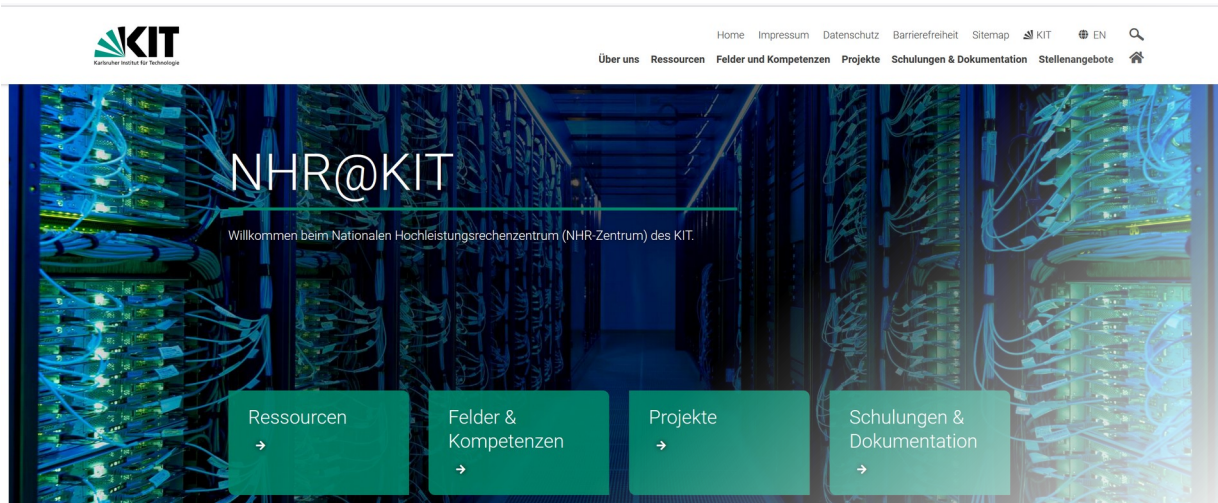
- Software Sustainability and Performance Engineering Team (SSPE)
 - Porting codes to new Programming Languages, Environments, Libraries
 - Porting codes to new hardware (Accelerators, ARM CPUs etc.)
 - Support with Continuous Integration/Testing/Benchmarking/Deployment (Cx)
- Simulation and Data Life Cycle Labs (SDLs)
 - Support with scientific data mgnt., data intensive computing & analysis

What:

- Training
 - More trainings, workshops, Hackathons etc. at KIT
 - NHR@KIT Training Overview: <https://www.nhr.kit.edu/english/66.php>
 - NHR Alliance Training Overview: <https://www.nhr-verein.de/en/courses-and-workshops>
- Call for Collaboration Projects
 - Rounds of calls for proposals for collaborative research projects between researchers in NHR@KIT and from the user communities earth system science, materials science, engineering in energy and mobility, as well as particle and astroparticle physics ...

NHR@KIT: Support for users (2)

- Website, <https://www.nhr.kit.edu>
 - Resources, Documentation, Consulting, Training, Support ...



- Voucher Projects
 - To apply with a specific project via the application form to SDLs and SSPE team, to work on questions in the area of High Performance Computing as well as on the topic of Data Science.
 - Exploration Voucher (up to 6 weeks)
 - Realisation Voucher (up to 6 months)

Thank you for your attention!

Questions?