

The new systems at NHR@FAU

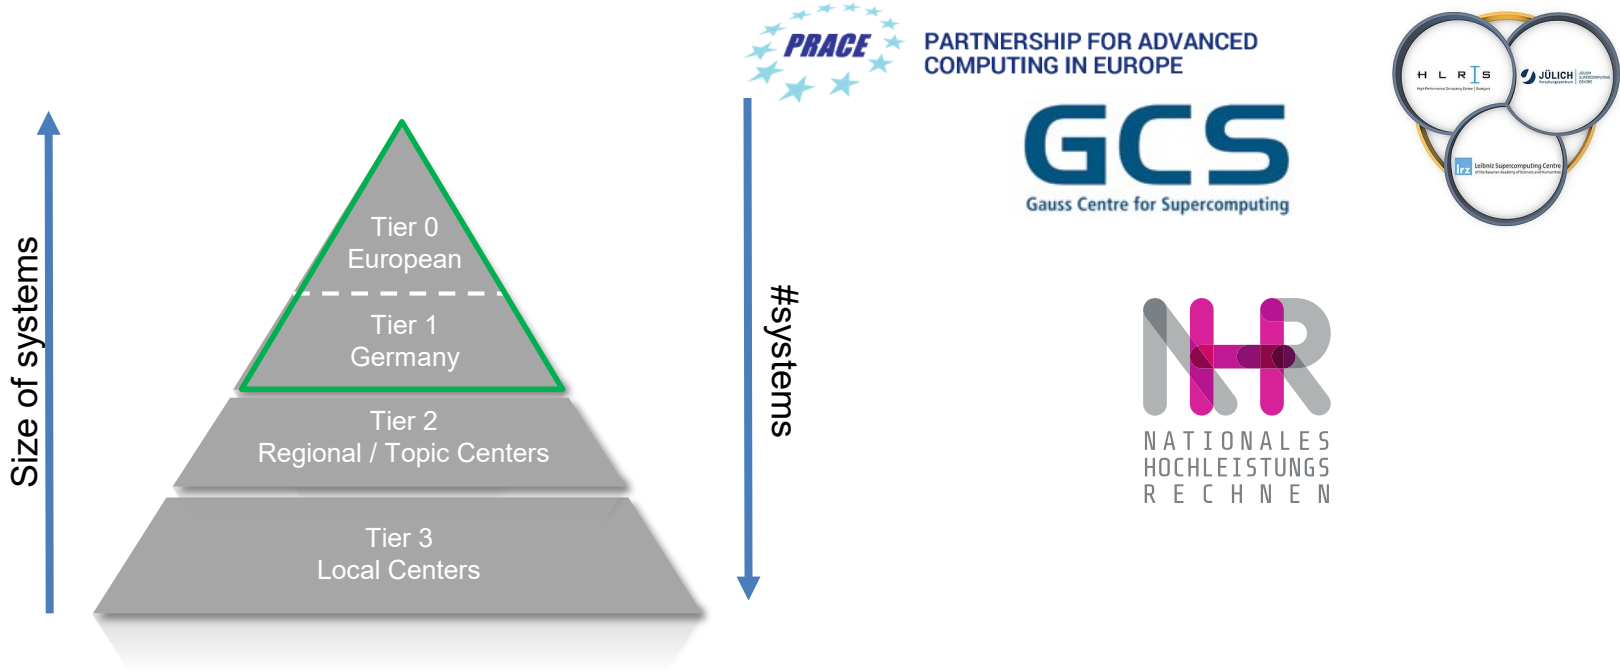
Hardware – Access – Application Process

HPC Café, 14 June 2022

HPC Services, NHR@FAU

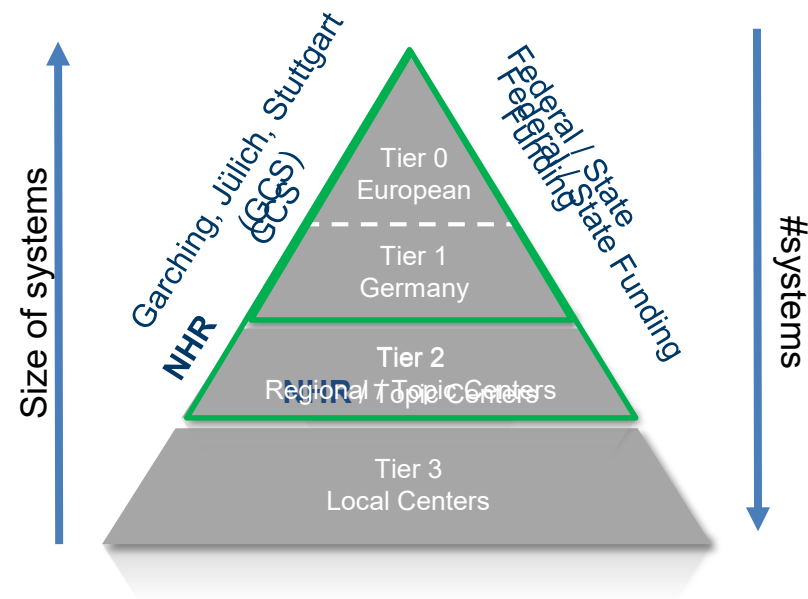
HPC supply pyramid in Germany

Recommendation by „Wissenschaftsrat“



NHR as part of HPC in Germany

HPC Supply Pyramid



- A coordinated network of **National High Performance Computing (NHR)** centers at German universities funded by federal and state governments
- The NHR network shall also strengthen **methodological competence** through coordinated **training** and continuing **education** of users and, in particular, of young scientists
- Total funding **625M Euro** (2021-2030)
- Currently **9 NHR centers**
- **Official start: Jan 1st, 2021**

The NHR centers





Im NHR-Verbund mit:

- RWTH Aachen
- ZIB Berlin
- TU Darmstadt
- TU Dresden
- U. Göttingen
- KIT
- U. Paderborn
- SW Verbund



Vorstand NHR-Verein:

- C. Schütte (ZIB),
- G. Wellein (FAU)
- C. Plessl (UP)

<https://www.nhr-gs.de>

- **Finanzvolumen (GWK): ca. 70 M€ (2021 – 2030)**
+FAU-HW (Tier-3; 2023 - 2030): ca. 7 M€ (Hardware)

- **Start: 1.1.2021**



- **Modernste HPC-Infrastruktur für numerische Simulationen, AI/ML und Datenanalyse**
- **Betriebskosten + Administration**
- **Fachexpertise **Beratung** und **Ausbildung** (~10 FTE)**
- **FAU Fachschwerpunkte im nationalen Verbund:**
 - Atomist. Simulationen (Tech. Fak./ Nat. Fak. / Medizin)
 - Performance Engineering (Informatik)
- **Fördermittel 2021/2: 14,3 M€** www.nhr.fau.de

Hardware

„Continuous investment cycles“

- **NHR funds** → nation-wide application process
 - FAU/DFG funds → FAU + regional universities only
 - Others → e.g. investment of specific projects
-
- **Compute systems:**
 - Investment cycle: every 3-4 yrs.
 - **Fritz & Alex: approx. 70-80% of cycles – NHR share**

 - **Central storage systems**
 - Investment cycle: every 5-6 yrs.


NHR@FAU: Hardware FAU-only

To be switched off in the coming months

- **Legacy FAU Systems** (Installation before 2021)
 - Emmy (2013) & Meggie (2016): IB-Cluster with ~ 10k + 15k Intel cores (Intel)
 - TinyGPU-pbs, TinyGPU-slurm: 68+142 GPUs (GTX1080*, 2080Ti, 3080, A100)
 - Woody-hw, Woody-sl/kl, TinyFAT 72 + 176 + 47 nodes

- **New system 2022:** - Partial renewal of Woody (not yet available)

- **Fundus** (Central **Online/Offline Storage**): 5 PB & space for 3300+ Tapes + diverse price-optimized NFS server (“\$WORK”)
- **New Tier2 systems 2021: „Fritz“ (CPU) & „Alex“ (GPU)**
 - Investment: 11,2 M€
 - Vendor: MEGWARE
 - NHR vs. FAU (not including project nodes): 3:1
 - Total power consumption: approx. 1 MW
- **New system 2022: 64 additional fat CPU-nodes (1 TB)**



NHR@FAU Fritz & Alex – Fact Sheet (NHR+FAU)

	# nodes	Node configuration	Storage	Typical job sizes	Peak (FP64)
Fritz	944 Intel ICL (67,968 cores)	2 * 36 c (8360Y) 256 GB 1 x HDR100	Shared PFS <ul style="list-style-type: none">• 3 PB• 80 GB/s	1 – 64 nodes	5.2 PF/s
Alex	32 NVIDIA A100 (256 A100) <i>(including 40 A100 dedicated to HSC)</i>	8 * A100 2 * 64 c (AMD) 1 TB 2 x HDR200	Node local 14 TB NVMe	1 – 8 GPUs	5.0 PF/s
	38 NVIDIA A40 (304 A40)	8 * A40 2 * 64 c (AMD) 0.5 TB	Node local 7 TB NVMe	1 – 8 GPUs	---

Fritz: <https://hpc.fau.de/systems-services/systems-documentation-instructions/clusters/fritz-cluster/>

Alex: <https://hpc.fau.de/systems-services/systems-documentation-instructions/clusters/alex-cluster/>


NHR@FAU next Gen Woody– Fact Sheet (FAU-only)

	# nodes	Node configuration	Storage	Typical job sizes	Peak (FP64)
Woody-NG	70 Intel ICL (2,240 cores)	2 * 16c (Xeon 6326 @2.9 GHz) 256 GB 1 x 25 GbE	Local SSD	1core – 1 node	
Woody	112 Intel Skylake (448 cores)	1 * 4c (E3-1240 v6 @ 3.70GHz) 32 GB 1x GbE	Local SSD	1core – 1 node	
	64 Intel Kabylake (256 cores)	1 * 4c (E3-1240 v5 @ 3.50GHz) 32 GB 1x GbE	Local SSD	1core – 1 node	

Which systems to use?

Fritz (CPU) / Alex (GPU) through **NHR@FAU** review process:

- **External researchers** (Non-FAU)
- **FAU** researchers with **large demand**



Which systems to use? FAU-only

- **FAU users with low demands**
 - Woody => single core/single node throughput – about 50% of the nodes have been financed by special projects
 - TinyFAT => huge memory demands – but keep in mind that all of these nodes have been financed by special projects
 - TinyGPU => general GPU workload – but keep in mind that many of these nodes have been financed by special projects
- **Experienced FAU users with extended demands (but below the NHR threshold)**
 - Woody => single core/single node throughput
 - Alex + Fritz within the FAU contingent => requires extra application and justification; **not for new comers or casual users!**

(Coming changes to the) Software environment

- **Alex + Fritz:** AlmaLinux8 (RHEL8 clone) + Slurm
- **Woody-NG:** AlmaLinux8 (RHEL8 clone) + Slurm (ETA: Q3)
- **TinyGPU-slurm / TinyFAT:** Ubuntu 20.04 + Slurm
=> to be migrated to AlmaLinux8 (RHEL8 clone) + Slurm (ETA: Q3/Q4?)
- **Meggie:** CentOS7 + Slurm
=> to be reinstalled with AlmaLinux8 (RHEL8 clone) + Slurm (ETA: Q3?)
- **Emmy:** CentOS7 + Torque => to be turned off in Q3/2022
- **Legacy Woody:** Ubuntu 18.04 + Torque => to be turned off in Q3/2022
- **TinyGPU-PBS:** Ubuntu 18.04 + Torque => to be turned off in Q3/2022

The software environment will be more uniform – making the movement between clusters easier! Also moving to spack-0.18.0.

Some special RRZE things will be abandoned (e.g. own wrappers for Intel MPI)

Application Process

Application for NHR resources: Basics

- NHR resources are **open to researchers at German universities**
- Researchers may chose center which meets their requirements best
- **Application** for resources through **central portal** (available in Q3/2022)
- **Peer review process** – scientific quality / need for Tier-2 resources
- **Simplified review process** for projects reviewed by DFG (and others)
- **User/project support** available by scientific and technical experts

Admission to NHR and NHR@FAU systems

- There will be a **central application platform** for all NHR funded systems
 - Upon submission, you will be asked for your “target center”
 - Available Q3/2022
- Meanwhile apply at each center locally

- Our **NHR@FAU application form** is available at <https://hpc.fau.de/systems-services/systems-documentation-instructions/nhr-application-rules/>
- Fill application form and submit it by email to: hpc-support@fau.de

- Any questions? Feel free to ask us: hpc-support@fau.de

Admission to NHR and NHR@FAU systems

- Application evaluation process
 - **Technical reviewing** by NHR@FAU staff
 - **Local Steering committee assigns reviewers** (local and external)
 - **Scientific reviewing** by local/external reviewers (max. 3 weeks)
 - Allocation decision by **local steering committee**
 - For large projects: approval by **central NHR board** necessary
(*“Zentraler NHR-Nutzungsausschuss”*)
 - Accepted projects get Liaison Scientist / support expert assigned
- **From application to decision: max. 3 months!**

NHR@FAU project types

Project type	GPU limits (GPU-hrs)	CPU limits (core-hrs)	Review	Deadline
Test projects (not FAU!)	< 3.000 (A40 or A100)	< 500 k core-hrs	technical review	rolling call
Normal	A40: 6.000 – 60.000	1 – 10 million	technical review; two external scientific reviews	rolling call
Pre-reviewed Normal	A100: 4.000 – 40.000		technical review; one simplified scientific review	rolling call
Large Scale*	A40: 60.000 – 180.000 A100: 40.000 – 120.000	10 – 30 million	technical review; two external scientific reviews	quarterly cutoff

* If more resources are required, consider to apply at GCS (www.gauss-centre.eu)

Application form

- Word document (will be replaced by central application portal)
- Structure
 - Part A: **Administrative Data**
 - Part B: **Scientific Project Information**
- The application form addresses:
 - scientific quality
 - justification of requested resources
 - performance / scalability of the code
 - administrative data

Project Application Form: Part A

- **Part A – Administrative Data**

- **1 General Information**

- Principal Investigator, Project Manager, institution
 - Scientific field, HPC experience

- **2 Project Information**

- Title, acronym, project type (test/normal/large) and duration
 - Pre-existing reviews, sensitive data

- **3 Technical Description**

- Requested resources (CPU, GPU, storage)
 - Application software and tools
 - Proof of efficient HPC usage

Project Application Form: Part B

- **Part B – Scientific Project Information**

- **Test/porting projects**

- Please provide a short summary of your test or porting project which justifies the use of the HPC resources. The total text should not exceed one page.
- A half page “project description” is sufficient. Focus on HPC aspects.

- **Normal projects with granted/reviewed DFG/BMBF/EU application**

- As the scientific part of your project has already been reviewed, it is sufficient to provide a project summary. Anyway you should make clear the amount of HPC resources you apply for. The total text of this chapter **should not exceed two pages**.

- **Large projects and normal projects without existing review**

- Please provide a scientific description of your project. The total should not exceed 3 to 6 pages for a normal project, and 5 to 10 pages for a large project, respectively.

Project Application Form: Part B

- **Part B – Scientific Project Information**

- **4 Scientific Project Description**

- Abstract, description (all types)
- State of the art and preliminary work, scientific goals, detailed scientific project description, detailed project plan (only large projects and normal projects without existing review)
- Summary/statement of existing scientific reviews and justification for applied HPC resources (normal project with grant/review by DFG, BMBF, EU, etc)

Please note: **Attaching an existing review (or essential parts of it) may considerably speed-up the scientific reviewing process!**

- References

Project Application Form: Part B

- **Part B – Scientific Project Information**

- **5 Follow-up Projects**

- Resources used and results obtained so far
 - Short intermediate report and outreach

- **6 Optional Information**

- Reviewer suggestion
 - Awareness in the community

- **7 Final Notes**

- Acknowledgement of granted resources in publications/applications
 - Providing of publications to NHR@FAU
 - Possibility to future serve as a reviewer

Need Help?

We will **help you answering your questions** about the compute time application, the submission and the reviewing process

We will also **support you porting your applications** to run effectively on our systems

In any case, contact us via hpc-support@fau.de

After submitting the application, there is the possibility to immediately start doing calculations (via a test/porting project account)

Standard project runtime: 12 months – **project extension with reduced effort** is easily possible!

Access

Access

- New HPC portal: <https://portal.hpc.fau.de>
<https://hpc.fau.de/systems-services/systems-documentation-instructions/getting-started/nhrfau-hpc-portal-usage/>
- Login with SSO (DFN-AAI/eduGAIN) to portal
- **Access only with SSH-Keys to the HPC systems**
- Already in use for NHR projects
 - If your university does not yet release the required attributes to our service provider, you'll see an error page asking you to contact your local IT (i.e. your identity provider)
- In use also for first regional universities
- Will hopefully replace paper forms also for FAU projects by end of the year!

Export Control and Political Sanctions

The principal investigator is responsible for **compliance with all regulations concerning export control and dual-use goods**. In this context, check for any **existing restrictions on the use of or access to HPC systems for foreign nationals**.

Currently, computing time or consulting provided through NHR **may not be used for projects involving Russian partners**, even if you as the applicant are a German citizen.

There is **no general ban** on the use of resources by Russian nationals, so that employees with Russian nationality at German universities may continue to use the HPC systems, provided that no specific sanctions apply to them.

Final steps

- For accepted projects please do not forget the contribution of NHR@FAU
 - Acknowledge support by NHR@FAU using our standard text:
<https://hpc.fau.de/systems-services/hpc-usage-reports/>
 - Send us a copy of your paper to nhr-redaktion@lists.fau.de
 - Results should also be briefly presented on NHR@FAU and/or NHR-Verein websites!!!
- Accepting NHR@FAU large scale project proposals until July 1st, 2022 <http://tiny.cc/NHRFAU-Application>



<https://hpc.fau.de/systems-services/systems-documentation-instructions/nhr-application-rules/>

THANK YOU! – Questions?

NHR@FAU

<https://hpc.fau.de>