


Mid-term hardware strategy at NHR@FAU

HPC-Café 25.04.2023

Prof. Dr. Gerhard Wellein, Dr. Thomas Zeiser
Zentrum für Nationales Hochleistungsrechnen Erlangen (NHR@FAU)



Agenda

- Hardware: Status/Update The logo for NHR FAU features the text 'NHR' in blue, a green hexagonal shape with a circuit-like pattern, and the text 'FAU' in blue.
- Investment cycle 2023 – 2025
 - NHR investments
 - Procurement process
 - FAU(-only) needs
- Terms:
 - **NHR** → NHR-funds → accessible through NHR-application process („**Tier-2**“)
 - **FAU** → FAU-funds (91b) → FAU/RRZE access only („**Tier-3**“)

NHR@FAU – Hardware: Status / Update



NHR@FAU: HPC Infrastructure

- **HPC systems** accessible through **NHR+FAU**
 - Fritz: CPU-Cluster (11/2022: TOP500, rank 151) with approx. 70,000 cores
 - Alex: GPU-Cluster (11/2022: TOP500, rank 174) with 304 (+352) NVIDIA GPUs
- **Central** (online/offline) **storage**: 8 PB / >3300 tapes possible (**NHR+FAU**)
 - HOME, WORK, VAULT (ARCHIVE)
- HPC systems for **FAU** (local access only)
 - Approx. 17,000 cores (meggie, woody, tinyFAT)
 - Approx. 150 NVIDIA GPUs (tinyGPU)

Fritz & Alex: Fact Sheet

	#nodes	Node conf.	Storage	Typical job sizes	Peak (FP64)
Fritz	992 Intel ICL (71,424 cores) (NHR:76%; FAU: 24 %)	2 * 36 c (8360Y) 256 GB 1 x HDR100	Shared PFS • 3 PB • >20 GB/s	1 – 64 nodes	5.9 PF/s (4.1 PF/s)
	64 Intel SPR (6,656 cores) (NHR:87%; FAU: 0%)	2 * 52 c (8470) 1 TB / 2 TB 1 x HDR100		1 – 4 nodes	n.y.a.



151

Fritz - Megware D50TNP, Xeon Platinum 8360Y 36C 2.4GHz,
InfiniBand HDR100, MEGWARE
Universitaet Erlangen - Regionales Rechenzentrum Erlangen
Germany

71,424

3.58

5.45

672

Power consumption (kW)
for LINPACK

Fritz & Alex: Fact Sheet

	#nodes	Node conf.	Storage	Typical job sizes	Peak (FP64)
Alex	38 AI/ML (304 NVIDIA A100) (NHR:61%; FAU:18%)	8 * NVIDIA A100 2 * 64 c (AMD) 1 TB 2 x HDR200	Node local 14 TB NVMe	1 – 8 GPUs	5.6 PF/s
	44 MD (352 NVIDIA A40) (NHR:78%; FAU:22%)	8 * NVIDIA A40 2 * 64 c (AMD) 0.5 TB	Node local 7 TB NVMe	1 – 8 GPUs	---



174

Alex - MEGWARE NF5488A5, AMD EPYC 7713 64C 2GHz, NVIDIA A100 SXM4 80 GB, Infiniband HDR, MEGWARE
Universitaet Erlangen - Regionales Rechenzentrum Erlangen
Germany

34,720

3.24

5.60

149

Power consumption (KW)
for LINPACK

NHR@FAU – Investment cycle 2023 – 2025

NHR investments

Procurement process

FAU(-only) needs



NHR-investment funds

- Annual budgets
 - complex (forward looking) procurements
 - complex payment plans (annual budget must be spent by Dec. 31)
- Available for HPC investments in 2023/24: approx. 6.5 M€
- **Contract must be signed in 2023**
 - Payments: 2023 and 2024
- Hardware installation ~ Q4/2024 – Q1/2025
- Infrastructure extension 2023/4 – Home of Fritz / warm water cooling

NHR investment

- Basic idea: Fritz.2 & Alex.2
 - CPU-only cluster – homogeneous
 - GPU-throughput cluster – homogeneous

→ But open procurement!

- Open: Investment ratio CPU vs. GPU
 - Fritz / Alex: 2:1 (initially)
 - Now: 2:1 – 3:1 – ???
- NHR project volume – as of today:
 - Fritz: 1,3 years
 - GPU: 1 year

NHR investment

- Fritz.2
 - Compute nodes: 2 x n-core Chips + 512 GB memory (min.)
 - Network: Min. 100 Gbit/s – 200 Gbit/s (if n is large)
- Alex.2
 - Compute nodes: 4 – 8 GPUs + Host CPUs
 - Network: Min. 100 Gbit/s
- Shared parallel file system
 - All flash – min. 1 PB **or**
 - HDD – min. 3 PB
- High speed network architecture: blocking allowed

NHR@FAU – Investment cycle 2023 – 2025

NHR investments

Procurement process

FAU(-only) needs



Procurement process

- Wettbewerblicher Dialog mit Teilnahmewettbewerb
- Zeitplan des Verfahrens bis Zuschlag
 - 16.4. Veröffentlichung Teilnahmewettbewerb
 - 15.5. Ende Teilnahmewettbewerb
 - Anfang Juni: Dialogrunde
 - Mitte Juni: Finale Ausschreibungsunterlagen → Angebote
 - Mitte Juli: Dialogrunde
 - Anfang September: Finale Angebote
 - Anfang November: Bekanntgabe Zuschlag

Procurement process – user involvement

- Mitglieder im Ausschreibungsteam:
 - Max. 2 externe Mitglieder
 - NDA!
 - Zeitaufwand: ca. 4-5 Wochen → siehe vorherige Folie
 - Technischer Hintergrund wichtig!
- Auswahlausschuss
 - Ausschreibungsteam berichtet dem Ausschuss über Fortschritt des Verfahrens
 - Diskussion offener Punkte im Verfahren
 - Finale Entscheidung
 - ca. 4 – 5 Sitzungen
- Email an gerhard.wellein@fau.de oder thomas.zeiser@fau.de (bis 3.5.)

NHR@FAU – Investment cycle 2023 – 2025

NHR investments

Procurement process

FAU(-only) needs



FAU needs – „Tier-3“

- Access without NHR proposal for FAU members
- Current FAU/Tier-3 resources
 - Approx. 20% share of Fritz / Alex
 - Woody, tinyFAT, tinyGPU (partially funded by research groups)
 - Meggie (in operation since Q4/2016)
- Meggie / tinyGPU may be replaced by share of new system
 - Contract will allow extensions in 2024/5
 - DFG Art 91b application (1 – 3 Million €) needs to be written – together with potential FAU-users
 - Demand?

Background Meggie

- Meggie
 - Approx. 700 nodes with 14,000 cores – Intel Broadwell generation
 - OmniPath Interconnect – Intel split off company...
 - In operation since Q4/2016!!!
 - No guarantee for stability – failure of a central component → immediate shutdown
 - Power consumption: 100 KW – 150 KW
 - May be shut down in Q3/Q4 because of energy costs
- Please port codes and move project from Meggie to Fritz
 - Simple form via webpage: <https://hpc.fau.de/tier3-access-to-fritz/>
 - In-person assistance / support available for code porting / optimization
Talk to us or contact us via hpc-support@fau.de

-
- Comments / recommendations / complaints ????
 - Next HPC-Cafe: 9.5.2023