

# Kontaktpersonen (“Liaisons”) und HPC

HPC Cafe, 2022-02-08

HPC Services, RRZE / NHR@FAU



# HPC Services at FAU: NHR@FAU

- Erlangen National High Performance Computing Center (NHR@FAU)
  - Established 2021-01-01
  - Emerged out of RRZE; still close connections
  - Part of the German NHR Alliance, one of nine centers
  - Several joint activities in teaching, operations, support
  - Nationwide activity
- Focus topics at NHR@FAU
  - Atomistic simulations
  - Performance Engineering
  - Efficient building blocks for sparse solvers



# NHR@FAU structure

## Executive Board

**RRZE**  
**Technical Director**  
Marcel Ritter

**NHR@FAU**  
**Director**  
Prof. Gerhard Wellein

**NHR@FAU PIs**

**Systems**

**Dr. Thomas Zeiser**

**Training &  
Support**

**Dr. Georg Hager**

**Software &  
Tools**

**Dr. Jan Eitzinger**

**Research**

**Prof. Harald Köstler**

# Why are you here?

- High Performance Computing (**HPC**) resources at FAU
  - are currently being heavily extended (**10 M€** investment + 5 M€ legacy systems)
  - are getting increasingly **popular even with inexperienced users**
  - are often **not used effectively**
- Problems
  - **Hard** to get required **info** to **new users**
  - **Alarming lack of knowledge** with some new users
    - Handling a Linux environment from the command line
    - Handling of batch jobs
    - Data handling
    - Resource awareness
- Consequences
  - Heavy **strain** on NHR@FAU support
  - **Low utilization** of expensive **resources** (esp. GPUs)
  - **Impact** on all **users** (slowdown of file system)
- **Please help us improve the situation**

# Why is this important?

---

- HPC hardware and infrastructure are expensive resources
- Example: One high-end GPU, used for one year (we have ~500 of those)
  - 2000 € in energy consumption + another 2000 € in investment
- Many users are unaware how much money they “burn” when running their workload
- Resources blocked by user A cannot be used by user B
- We require your help in making customers aware

# How can you help?

---

- You are the **initial point of contact** for customers of RRZE and NHR@FAU
- Relevant for HPC customers
  - Help with **HPC account** applications
    - New workflow ahead!
  - **Computing time** cost and allocations
  - Software **licensing**
  - “**Getting Started**” info
  - **Specific** info on applications and workflows used by **others in your group**
  - Material on proper **file** and **resource handling**
  - **Data** management
- **In case of doubt, talk to us ([hpc-support@fau.de](mailto:hpc-support@fau.de))**

# The HPC account




# Types of HPC accounts

- With the establishment of NHR@FAU there are 2 types of HPC accounts:

- **Basic FAU HPC service (“Tier 3”)**

- Simple application form – as in the past
- Only **limited resources** (~25% of the new HPC clusters)
- **FAU has to pay** for hardware, electricity, staff, etc.



Will  
concentrate  
on this case  
today

- **NHR projects**

- Scientific **project proposal** with (external) scientific and technical **review**
- **Competition** with project applications from universities all over Germany
- **Major resources** available (~75% of the new HPC clusters)  
4k-50k-120k GPU-hours/a or 1M-10M-20M core-hours/a per project
- **No costs for FAU** – NHR pays for hardware, electricity, staff, etc.
- <https://hpc.fau.de/systems-services/systems-documentation-instructions/nhr-application-rules/>



# The HPC account for the basic FAU HPC service

---

- This is **not the same** as the user's **IdM** account
  - Form: <https://www.rrze.fau.de/files/2017/06/HPC-Antrag.pdf>
  - See <https://hpc.fau.de/systems-services/systems-documentation-instructions/getting-started/> for advice on how to fill out the fields
- You can **help your colleagues** by giving them the basic info

# HPC account form (here, “you”==“the user”)

Antrag auf Nutzung von HPC-Ressourcen am RRZE **HPC** High Performance Computing

Stand 03/2017

Antrag per FAX an 09131-85-29966  
oder als Scan an [rrze-zentrale@fau.de](mailto:rrze-zentrale@fau.de)  
oder per (Haus)post an RRZE/Serviceheke,  
Martensstr. 1, 91058 Erlangen

Neuantrag  Änderung/Verlängerung

IDM-Account:  (bestehender) HPC-Account:

**Antragsteller:**  Frau  Herr Titel

Vorname   
Nachname   
E-Mail   
Telefon   
Nationalität(en)   
HPC-Ablaufdatum bis DD.MM.JJJJ

Bei Unklarheiten vorab [support-hpc@fau.de](mailto:support-hpc@fau.de) kontaktieren

HPC-Zielsysteme   
typische Jobgröße   
inages. benötigte Rechenzeit   
benötigter Speicherplatz

**Auftraggeber:** FAU-OrgNr.

Lehrstuhl- oder  
Instituts-  
**stempel**  
und  
-anschrift

RRZE-Kontaktperson

Art der  
Anwendung /  
Name der  
Applikation

Your data

Account  
expiration date

Systems,  
requirements  
(brief!)

Your IdM account

Chair data & seal

Very brief  
description of  
what you want to  
do

# HPC account form

## Type of project:

- Education (master/bachelor)
- Standard (FAU employee)
- Research grant (BMBF, DFG, EU)
- Industry

Art./Finanzierung des Projekts	
Abrechnung der Rechenzeit über	<input type="checkbox"/> bestehende Kundennummer <input type="checkbox"/> neue Kundennummer für folgendes Rechenzeitprojekt
<small>für jede neue Kultur</small>	<small>RRZE intern</small> Projekt
Titel des Forschungsvorhabens	
für das Projekt insgesamt benötigte Rechenzeit	
Bewilligungszeitraum / Projektlaufzeit	
fördernde Institution und Förderkennzeichen	
Kurze Beschreibung der HPC-Aktivitäten im Forschungsvorhaben	
Wurde der Rechenzeitbedarf mit dem RRZE abgestimmt und im Antrag dargestellt? Haben die Gutachter der Förderinstitution dazu Stellung genommen?	

RRZE customer ID (ask contact person)

Personenbezogene Daten im Sinne der geltenden Datenschutzgesetze dürfen unter dieser Benutzerkennung nicht ohne Sondergenehmigung seitens des RRZE und des Datenschutzbeauftragten verarbeitet werden!

Dem Antragsteller ist bekannt, dass er sich durch eine missbräuchliche Benutzung der Informationsverarbeitungssysteme strafbar machen kann und dass beim Vorliegen eines Missbrauchs grundsätzlich Strafantrag gestellt wird. Des weiteren bemüht sich der Antragsteller, die HPC-Systeme effizient zu nutzen und gängige HPC-Praktiken zu beachten.

Benutzerrichtlinien:

<https://www.rrze.fau.de/infocenter/rahmenbedingungen/richtlinien/benutzerrichtlinien/>

Antragsteller und Auftraggeber erklärt hiermit, von den Benutzerrichtlinien sowie den ergänzenden Hinweisen auf der Rückseite dieses Antrags Kenntnis genommen zu haben

Ort, Datum	Unterschrift Antragsteller
RRZE intern	Unterschrift Auftraggeber/ Kontaktperson
Bemerkungen	idM. Freigabe Auftraggeber/Kontaktpers.

Date & place

Signatures: you and your boss or contact person

# Questions about account and resources (I)

---

## Frequent questions

- Which system(s) do I need?
  - If they don't know they should contact us first
- How to estimate resources?
  - We know this is a moving target. We just need very rough numbers (core or node hours, GB storage, parallel job size)
  - If in doubt, contact us
- What if the resources are needed for a BMBF/DFG/EU/... project?
  - Ideally, the computing center should have been contacted before the actual project proposal
  - We need to make sure we can provide the resources

# Questions about account and resources (II)

- **Does the service cost money?**
  - Baseline service (average use, standard file quotas) is free for scientific projects conducted by researchers from FAU
  - Third-party funded projects also get free baseline service (but talk to us first)
  - Better service is a matter of negotiation
    - Many solutions are possible – talk to us!
  - **Industrial** (contract) **research** is **not free**
    - Official pricing: <https://www.rrze.fau.de/infocenter/preise-kosten/#hpc>
    - Agreements are stipulated on a case-by-case basis
- **How much compute time can I draw?**
  - There are (currently) no fixed allocations
  - Users who use more CPU time get lower priority (“fairshare” mechanism)
  - Major cause of “my job does not run”

# Questions about account and resources (III)

---

- What happens after the account application was sent?
  - Notification of the user by e-mail once the account is ready for use
  - Account will appear in IdM
  - New password can be set there
  - Password changes can take several hours until they propagate to all our systems
  
- These procedures will change soon

# Upcoming changes (2022)

---

- **Electronic workflow (yay!)** for all account matters
- Access via **SSH key** instead of password
- Users must be **invited** by **liaison** or **PI**
  
- **Several project types**
  - Student projects (bachelor/master/...)
  - Hands-on accounts for lectures
  - Standard basic service (limited allotment)
  - “Normal” project with prior grant from external body (BMBF, DFG)

# Special case: HPC accounts for lectures/tutorials

- **Simplified application protocol**
  - The lecturer sends the **complete list** of **IdM accounts** of the students to [hpc-support@fau.de](mailto:hpc-support@fau.de) – no need for individual HPC account forms
  - The accounts will be created in one big bulk action
  - The **lecturer is responsible** for informing the students and **ensuring proper usage**
  
- **Survey** on the general demand of **HPC** resources for **lectures/tutorials**
  - <https://hpc.fau.de/umfrage-hpc-ressourcen-fur-lehrveranstaltungen/>



# Software licensing



# Software licensing

---

- **If commercial software is required, users need to bring their own licenses**
  - Exception: standard HPC tools (compilers, libraries, performance tools)
- Very software specific
  - Local license files in \$HOME
  - License server @ computing center
  - License server @ your chair
  - License sever somewhere on the internet
  - Parallel or sequential code? Are parallel licenses available?
- In case of doubt, talk to us

# New users info



# New-user intro and knowledge transfer (I)

- **Required skills** for HPC users
  - Linux on the command line
  - Handling logins via Secure Shell (only way to access our systems)
  - (Efficient) file handling
  - Batch processing
- **Official docs & training**
  - “HPC Intro for beginners,” live talk on Wednesday after the monthly HPC Cafe
  - <https://hpc.fau.de/systems-services/systems-documentation-instructions/getting-started/>
  - “Logging in to RRZE HPC systems” video: <https://youtu.be/J8PqWUfkCrl>
  - <https://hpc.fau.de/systems-services/systems-documentation-instructions/hpc-storage/>
  - "Handling file systems properly" video: <https://www.fau.tv/clip/id/40199>
  - <https://hpc.fau.de/systems-services/systems-documentation-instructions/batch-processing/>
- **More extensive material is in the making**

# New-user intro and knowledge transfer (II)

---

- **Information specific to your group**
  - Special software
  - Common input data sets
  - Job scripts and workflows that are “known to work”
  - Data handling, pre- and postprocessing best practices
- **Do not leave your colleagues alone with this**
  - Set up a process for knowledge transfer within the group
  - Experienced HPC users should educate newbies

# File storage and data management



# Data/file storage

## NHR@FAU provides several data storage options

Access	Purpose	Technology	Backup	Snapshots	Data lifetime	Quota
\$HOME	Source, input, important results	NFS on central servers, small	YES	YES	Account lifetime	50 GB (fixed)
\$HPCVAULT	Mid-/long-term storage	Central servers	YES	YES	Account lifetime	500 GB (flexible)
\$WORK	Short-/mid-term storage, General-purpose	Central NFS server	(NO)	NO	Account lifetime	500 GB (flexible)
\$FASTTMP <small>(cluster-local, only some clusters)</small>	High performance parallel I/O	Lustre parallel FS via InfiniBand	NO	NO	High watermark	Only inodes
\$TMPDIR	Node-local job-specific dir	HDD/SDD/ramdisk	NO	NO	Job runtime	NO

Shared file systems

More is possible, up to negotiation (talk to us!)

# Shared file systems

- What does “shared” mean?
  - The actual storage devices (HDDs, SSDs) are part of a server infrastructure
  - The servers can be accessed by many users (concurrently)
  - The servers can be accessed from many client systems (concurrently)
- Shared file systems pros & cons
  - Available from “everywhere”
  - Data exchange made easy (among users & systems)
  - Shared resource == shared performance
  - A few users can slow down operations for everyone
- <https://hpc.fau.de/systems-services/systems-documentation-instructions/hpc-storage/>
- “Handling file systems properly” video: <https://www.fau.tv/clip/id/40199>



# Long-term data management

---

- **It is the responsibility of the group to manage data handover**
  - Contract termination, finished thesis
  - Do the handover while the leaving colleague is still in reach
  - We can help with data copying but we need explicit permission from the owner
- **Do not wait until your colleague has left the university**
  - People are hard to reach sometimes

# Upcoming hardware & software changes 2022



# Upcoming hardware & software changes 2022

- **EOL for parallel computer Emmy** (2013-2022) late spring?
  - **Partial upgrade** of throughput cluster **Woody** June-August
    - Nodes with only 4 GB/core will be retired
    - OS upgrade from Ubuntu 18.04 to 20.04
    - Batch system change from Torque to SLURM
    - Granularity will change from nodes to cores
  - **EOL for GTX1080/GTX1080Ti** nodes in **TinyGPU** summer?
  - **OS upgrade** on parallel computer **Meggie** early summer?
- no **CentOS7** any more; AlmaLinux8 on all parallel computers + Alex
- no **Ubuntu 18.04** any more; Ubuntu 20.04 on Woody and all Tiny\*
- all clusters will run **SLURM** as batch system

# Conclusion



# Thank you for your time

---

Rule #1: Talk to us!

[hpc-support@fau.de](mailto:hpc-support@fau.de)



<https://hpc.fau.de>