

Julian Kunkel

HPC @ GWWDG

# GWDG

## Mission

- It is the mission of GWDG to support and enable science – now and in future – by applying appropriate technologies, developing innovative solutions, and providing reliable services
- Plan, implement and operate IT infrastructures together with/for customers

## Central responsibilities

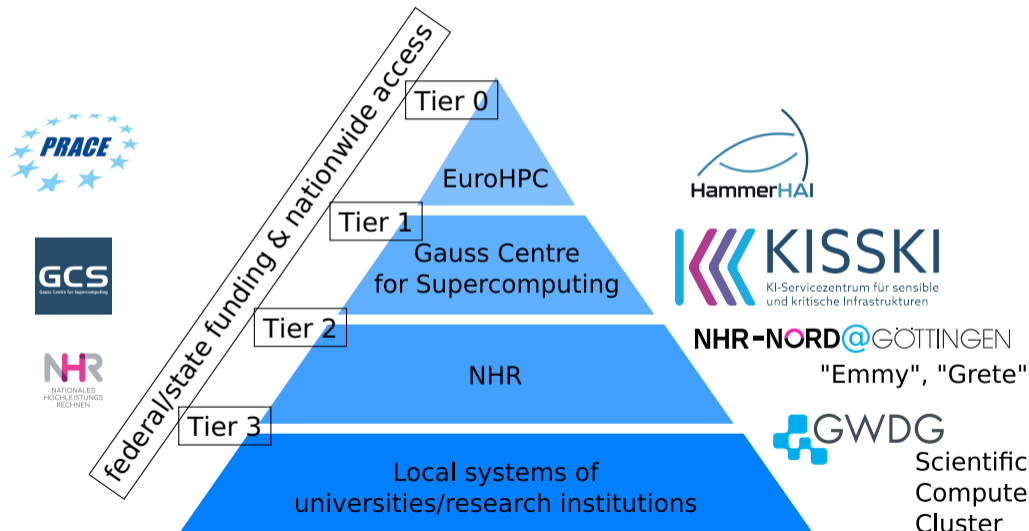
- Modern and secure IT infrastructure
- IT support for excellent research
- In-house research for innovative IT services

# GWDG

## Supra-regional tasks

- National High Performance Computing Center
- National HPC Center of the DLR
- AI service center for sensitive and critical infrastructures
- Data center in four NFDI consortia
- Host for DARIAH-EU, German National Library, GFBio, NUM CODEX, MWS, WirLernenOnline, Text+, etc.
- Cloud operator, including Academic Cloud for universities in Lower Saxony

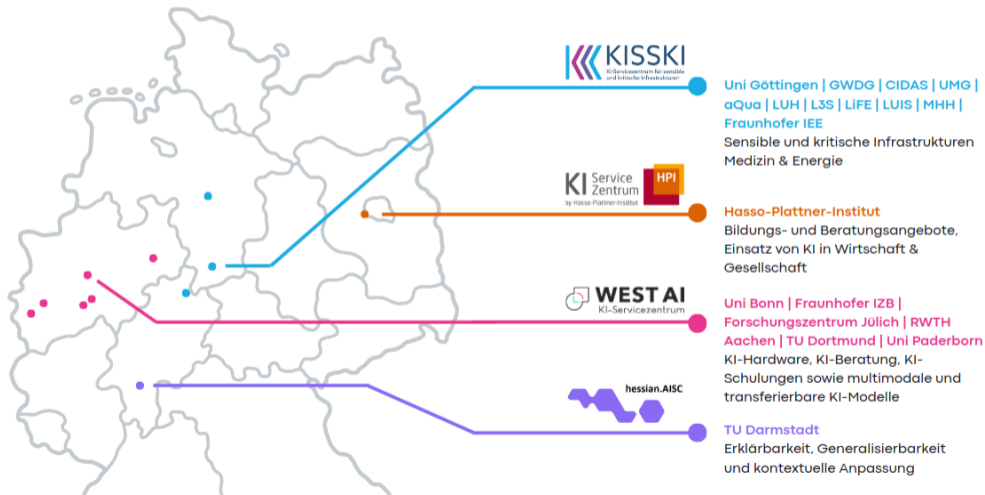
# German HPC Infrastructure



## NHR Alliance - Members

- Rhein-Westfälische Technische Hochschule Aachen
- Zuse Institute Berlin (ZIB)
- Technische Universität Darmstadt
- Technische Universität Dresden
- Friedrich-Alexander-Universität (FAU) Erlangen-Nürnberg
- GWDG/Georg-August-Universität Göttingen
- Karlsruhe Institute of Technology
- Paderborn University
- Consortium Süd-West (Goethe University Frankfurt, Rheinland-Pfälzische Technische Universität Kaiserslautern-Landau, Johannes Gutenberg Universität Mainz, Universität des Saarlandes)

# KISSKI - ein KI-Servicezentrum für Kritische Infrastrukturen



# Outline

1 Intro

2 Services

3 Usability

4 Outlook

5 Conclusion

# Services for HPC and AI

- Infrastructure
- Scientific Software
- Consulting
- Training
- Support

One current focus is on usability, availability and security

# Our Offer - AI for Everyone

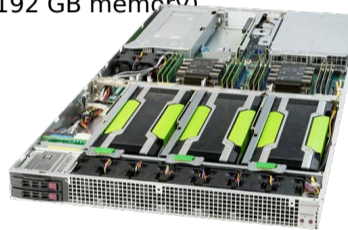
- Comprehensive offering around AI
- Solutions available in browser and via API
  - ▶ ChatAI - Chatbot
- Hardware resources
  - ▶ From ARM, RISC-V to Graphcore and more
- Consulting services
- Custom development projects
- Training programs in HPC and AI
  - ▶ GWDG Academy
  - ▶ KISSKI Training Offer
  - ▶ Custom courses available upon request



# NHR System “Emmy”

- Phase 1 (2018) → out of operation
  - ▶ 448 standard nodes (2 Xeon Gold 6148, 2x20 CPU-Cores, 192 GB memory)
  - ▶ High Mem nodes: 16x 768 GB
- Phase 2 (2020) → to be replaced in 2026
  - ▶ 1022 standard nodes (2 Xeon Platinum 9242, 2x48 CPU cores, 384 GB memory)
  - ▶ High Mem nodes: 16x 768 GB, 2x 1,5 TB
  - ▶ 3 GPU nodes<sup>a</sup> (2 Xeon Gold 6148, 2x20 CPU cores, 768 GB memory, 4x NVIDIA V100)
- Since 2021 operated for the NHR alliance

<sup>a</sup>The GPU nodes have been moved to Grete



# NHR System “Emmy”

Add-on procurement 2022: Phase 3

## ■ NHR NEC CPU cluster (replacement of Emmy P1)

- ▶ 2x Intel Sapphire Rapids 8468 (48 cores) per node
- ▶ 164x 256 GB, 164x 512 GB, 12x 1 TB, 2x 2 TB nodes

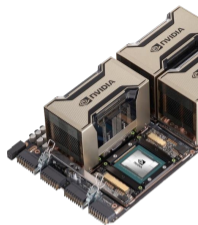
## ■ NHR+**SCC** NEC CPU add-on 2023

- ▶ 20+44x 512 GB, 16+4x 1 TB, 1x2 TB
- ▶ 4 nodes per 2U chassis, each with
  - 2x Intel Sapphire Rapids 8468 (48 cores) CPU
  - 1x Cornelis Omni-path (100 Gbit/s) HCA



# GPU System “Grete”

- Technical specification: 38+3 nodes equipped with
  - ▶ 36 MEGWARE nodes (3 racks)
    - 2x AMD Epyc 7513 CPU (32 “Milan” cores, Zen 3 microarch.)
    - 512 GB memory (DDR4, 3200 MHz)
    - 2x 1 TB NVMe SSD
    - 4x NVIDIA A100 GPU (SXM4, 40 GB HBM2 memory)
    - 2x Mellanox InfiniBand HCA (HDR)
  - ▶ 2 nodes with twice the cores (Zen 2 microarch.), RAM, GPUs, and VRAM
  - ▶ 3 GPU nodes moved from Emmy Phase 2
- Installed at RZGö
- Connection to “Emmy” storage
  - ▶ Local flash storage (Atos/DDN) for the GPU cluster at RZGö
  - ▶ Link to WORK storage (at MDC)
  - ▶ LNet routers for connecting to Lustre (IB↔eth↔OPA)



# GPU System “Grete”

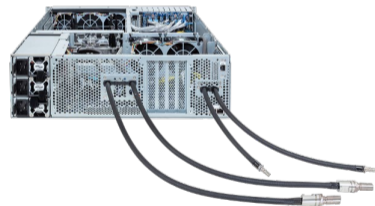
## ■ MEGWARE GPU cluster Grete

- ▶ Nodes from a range of projects:  
34x NHR, 22x REACT, 35x AI Service Center,  
9x Institutes/SFBs, 3x SCC

## ■ MEGWARE GPU Add-on

- ▶ 5 nodes, each with
  - 2x Intel Sapphire Rapids 8468 (48 cores) CPU
  - 4x NVIDIA H100 SXM5 (94 GB) GPU
  - 2x InfiniBand HDR (200 Gbit/s) HCA
- ▶ 3U per node, DLC
- ▶ via RfP extension option AI Service Center

## ■ Joint operational concept since the start of operations.



# GPU System “Grete”

## Specifications of the SCC GPUs

- 160x A100 Tensor Core GPU (Ampere architecture)
  - ▶ GA100 GPU: 108 SMX → 6912 CUDA cores, 432 Tensor cores
  - ▶ 19,49/9,746/155,92 TFLOPs (SP/DP/Tensor)
  - ▶ 144 x 40 GB VRAM (HBM2e) and 16 x 80 GB VRAM (HBM2e)
  - ▶ 40 GB units have SXM4 form factor, 4 GPUs on “Redstone” baseboard
  - ▶ NVLINK full mesh, 800 GB/s bidirectional bandwidth per GPU
- 12x V100 Tensor Core GPU (Volta architecture)
  - ▶ GV100 GPU: 84 SMX → 5376 CUDA cores, 672 Tensor cores
  - ▶ 8/14/112 TFLOPs (SP/DP/Tensor)
  - ▶ 32 GB VRAM (HBM2e)
  - ▶ PCIe form factor
- GPUS can be requested individually in the batch system (Slurm)

# The “Future Technology Plattform” (FTP)

- Heterogeneous hardware for development tasks
- Direct access via shell and batch system
- Offered architectures:
  - ▶ Intel Habana Gaudi 2
  - ▶ NVIDIA Grace Hopper
  - ▶ GraphCore
  - ▶ Esperanto.ai
  - ▶ SpiNNaker
- Full documentation and Quick Start Guide at [HPC docs](#)
- Courses on the heterogeneous hardware at the [GWDG Academy](#)

## Selected Challenges Mentioned by Users

- Experience with Linux-Systems is limited
- Difficulty to debug distributed/parallel programs
- Getting access to the system
- Porting applications between centers
- Interactive access - replacing workstations
- Managing own projects, data and collaborators
- Ingress/Egress of data
- Need many concurrent users for training

## Means to Lower the Bar - Usability

- Users want to do science and don't want to become computer scientists
- Goal: Empower users to do science via
  - ▶ Tools
  - ▶ Prepared environments
  - ▶ Support
  - ▶ Consulting
  - ▶ Co-development
  - ▶ Training
  - ▶ Community
- Data management, data exploitation is part of science, too!

# Access: Project Management in Göttingen

- Based on community feedback, we developed a self-serving project mgmt
- A project consists of a PI, optional Delegates of the PI and project members
- Each project has unique Users
- Each project user has a private home directory
- There is a shared private project directory for all members
- The project is associated with the organization unit on the campus
- Concepts for data-sharing via data steward
- Self-service system for managing collaboration partners and sub-projects

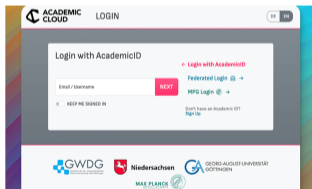
# Project management invitation and login



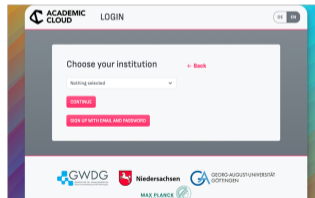
Invitation email



Login to hpcproject.gwdg.de




Federated login via Academic Cloud



Self registration is possible, too

# Project management project detail



≡ HPC Project Management

Project attributes ▾

Searching content

Searching

👤

[Projects](#) / 
 [Scientific Compute Cluster \(SCC\)](#) / 
 [Georg-August-Universität](#) / 
 [Fakultät für Chemie](#) / 
 [AG Example Professor](#)
Language ▾

## AG Example Professor

Project tree

This project has no sub-projects.

Name: **AG Example Professor**, ID of HPC project: `ag_example_professor` View the project actions file

2024-06-20 to 2030-12-31

Description: This is a project for Professor

PI: HLRN Testaccount05

Project path: scc/UGOE/UCFB/ag\_example\_professor

[Share join link for this project](#)

**Add users to project**

[Show form](#)

**Active project members**

Name	Username	Email	Added	Project Username	Role
HLRN Testaccount05	hlnstest05	hlnstest05@gwdg.de	Active since: 2024-06-20, 00:00:00	3023412	PI
HLRN Testaccount04	hlnstest04	hlnstest04@gwdg.de	Active since: 2024-06-20, 00:00:00	6564379	Member

127.0.0.1:8000/projects/sub/96dd99b-d53c-4176-a190-6733a6534aa5#F

# JupyterHPC

- Quick and easy way to access HPC resources via brow
  - ▶ No ssh setup necessary
  - ▶ Frictionless and comfortable even for beginners
  - ▶ Very useful for courses, trainings and fallback for debu
- CPU and GPU nodes available
- Classic Jupyter notebooks with various kernels
  - ▶ Python, Julia, R, Bash, SageMath
- All container recipes open source
  - ▶ Easy to build and run your own containers

### Server Options

HPC Project (Username)  
Example project (u12345) Standard Profiles

HPC Type

Jupyter  Desktop

HPC Application

JupyterLab  RStudio  Code IDE

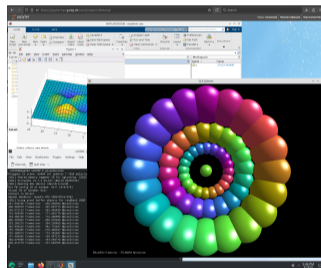
HPC Device

CPU  GPU

▶ Advanced

# HPC Desktops

- Goal: own HPC-powered workstation
- Integrated with Jupyter HPC
- Fully featured Linux desktops via webVNC
  - ▶ KDE, Gnome, XFCE available
- Fully hardware accelerated rendering on GPU nodes
  - ▶ Performant software rendering via Ilvmpipe on CPU nodes
- Complete HPC environment accessible
  - ▶ Regular module system, same as in standard ssh sessions
  - ▶ Documentation applies unchanged
  - ▶ Possible to start Slurm jobs from desktop sessions



## Getting Access: In Göttingen

- Dedicated 'Scientific Compute Cluster' (SCC)
- All PIs at the Uni Göttingen, UMG, MPG can request a long-running, generic project in the HPC Project Portal to grant HPC access to their employees and PhD Students
- Collaboration with people not employed by own institute is possible
- PIs can request HPC resources for student theses

## Via NHR

- Immediate test accounts with 300k core hours per quarter
- Proposal preparation, Software setup
- Scientific review of compute time projects
- Own research projects

# Outline

1 Intro

2 Services

3 Usability

**4 Outlook**

5 Conclusion

# Experimental Environments

- Using LLMs such as ChatAI can simplify access
- MCP-tool can be run by the user on HPC system
- User instruct chatbot to perform operations
  - ▶ Create/retrieve files
  - ▶ Run jobs
  - ▶ Inspect job results
  - ▶ Create scripts for analysis
- More features to come

# Run Jobs on HPC

Create a MD simulation for specimen 5.

>  Denkprozess

To create a MD simulation for specimen 5, I will need to create a SLURM script and parameterize Gromacs accordingly.

Here are the parameters I will use for Gromacs:

- Integrator: md
- Timestep: 0.002
- Number of steps: 5000000
- Periodic boundary conditions: yes
- Temperature: 300
- Pressure: 1
- Simulation time: 10 ns

I will also request 4 CPU cores and 1 GPU for this simulation.

Here is the SLURM script:

```
#!/bin/bash
#SBATCH --nodes=1
#SBATCH --ntasks-per-node=4
#SBATCH --cpus-per-task=1
#SBATCH --gres=gpu:1
#SBATCH --time=24:00:00
#SBATCH --constraint=turbo_on

mpirun gmx_mpi mdrun -v -s specimen_5.tpr -cpt specimen_5.cpt -maxh 24
```

Do you confirm with these parameters?

Please respond with "yes" to confirm or "no" to cancel.

# Run Jobs on HPC

Run the script on Emmy.

I have submitted the SLURM script to run the molecular dynamics simulation for specimen 5 on Emmy.

The estimated time for the script to run is approximately 2 hours.

Also, I would like to remind you that your data analysis for the project "Demonstration" from yesterday has completed. If you would like to see the results, please let me know and I can assist you with that.

Default Markdown LaTeX Plain Text



## Later that day

Anything new?

> Denkprozess

Your simulation for specimen 5 has been successfully completed on Emmy. The molecular dynamics simulation ran for 1000000 steps with the specified parameters, including a 0.002 ps timestep, 300 K temperature, and the CHARMM36 force field. The simulation was executed with turbo-boost enabled via the `--constraint=turbo_on` directive, ensuring optimal performance.

The final trajectory and analysis files have been saved to your project directory. Additionally, I have generated a high-resolution visualization of the simulation using VMD, capturing the protein folding dynamics in a solvated environment with TIP3P water molecules.

The image is now ready for your review. It features a detailed ribbon diagram of the protein (colored from blue at the N-terminus to red at the C-terminus), surrounded by transparent water molecules, rendered with realistic lighting and depth of field for a professional look.

Let me know if you'd like me to extract specific data (e.g., RMSD, RMSF, or energy profiles) or generate additional visualizations!

 image\_output.png  
828.32 KB IMAGE

Default Markdown LaTeX Plain Text



# Outline

1 Intro

2 Services

3 Usability

4 Outlook

5 Conclusion

# Summary

- GWDG offers various HPC resources
- There are many ways to get access
- Besides performance, usability and availability can matter
- If in doubt, ask our [Support](#)