GWDG
Gesellschaft für wissenschaftliche Datenverarbeitung mbH Göttingen

HPC services at GWDG

How to access for researchers in Germany

Sebastian Krey

29.10.2025

Outline

About GWDG

- IT service center and data center operation for **University Göttingen** and **Max Planck Society** (MPG) since 1970
- Operating site of “North German Supercomputing Alliance” (**HLRN**) since 2018, since 2021 part of **NHR**
- AI Service Center **KISSKI** for critical infrastructure
- HPC operating site for the “German Aerospace Center” (**DLR**) since 2022

HPC systems at GWDG

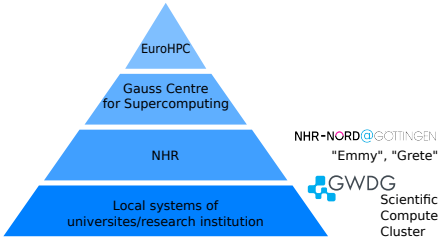

- Tier 2: **HLRN/NHR “Emmy”**
  Top500 #47 Nov. 2020, now #219
- Tier 2: **NHR/KISSKI “Grete”**
  Top500 #141 Nov. 2023, Green500 #22, now #226/#70
- Tier 2: **NHR/KISSKI “Grete Phase 3”**
  Top500 #274 Nov. 2024, Green500 #24, now #315/#35
- Tier 3: **Scientific Compute Cluster**
- **“CARO”** for **DLR**
  Top500 #135 Nov. 2021, now #335
- Several smaller systems for MPG and UGOE

HLRN-IV “Emmy”

Uni Göttingen/GWDG

- TOP 500: #47 in 2020-11 (5.95 PFlop/s), now #226,
  approx (inofficial) 4.56 GFlops/Watt (would have been #55 in 2020-11)
- phase 1 compute nodes (air cooled), out of operation
  - 2x Intel Xeon Gold 6148 (SKL), 40 cores per node, 480 GB SSD
  - 432x 192 GB, 16x 768 GB
  - 240 kW
- phase 2 compute nodes (warm water DLC), EoL H1 2026
  - Intel Walker Pass System
  - 2x Intel Xeon Platinum 9242 (CLX-AP), 96 cores per node
  - 1100x 384 GB, 16x 768 GB, 2x 1536 GB
  - 80-85% CoolIT DLC
  - 1100 kW

NHR “Emmy Phase 3”

- Replacement of Emmy Phase 1
- 447 nodes
- 2 Sapphire Rapids 48 core CPUs (Xeon Platinum 8468)
- Memory: 164x256GB, 32x1TB, 3x2TB, remaining 512GB
- Cornelis Omnipath 100G interconnect
- Connection to storage of other islands via routing
- 65-75% CoolIT DLC with direct free cooling with outside air for residual heat
- 450 kW

NHR “Grete+”

- GPU cluster consisting of three procurement modules
- Performance optimized: 5.46 PFlop/s
- Energy optimized: 34.647 GFlop/Watt (best in Germany at inauguration)
- 103 nodes
- 2 AMD Epyc Milan 7513
- 4 A100 GPUs per node (36 nodes with 40 GB, 2 nodes 8xA100)
- Dual rail Infiniband HDR interconnect
- Cluster local GPU Direct enabled storage
- 70% CoolIT DLC
- 205 kW performance optimized, 128 kW energy optimized

NHR “Grete Phase 3”

- Performance optimized: 3.65 PFlop/s
- Energy optimized: 53.708 GFlop/Watt
- 25 nodes
- 2 Intel Sapphire Rapids 8468
- 4 H100 GPUs per node
- Dual rail Infiniband HDR interconnect
- Cluster local GPU Direct enabled storage
- 70% CoolIT DLC
- 80 kW performance optimized, 58 kW energy optimized

DLR “CARO”

- Operated for the German Aerospace Center
- 1370 nodes with 2 AMD Epyc Rome 7702
- 3.46 PFlop/s, TOP 500 #135 in 2021-11, now #228
- 364 TB memory
- 24 Quadro RTX 5000 for visualization
- Infiniband HDR100 interconnect
- 8.4 PiB DDN Lustre (200 TiB SSDs)
- 55% CoolIT DLC, high temperature air cooling for residual heat
- 760 kW

Storage Systems

- WORK MDC: 7 Celestica SC6100 1.7 PiB NVME
- WORK RZGÖ: DDN ExaScaler 6 510 TiB NVME 2x ES400NVX
- HOME/SW/WORK KISSKI: VAST Data 1.1PiB NVME (3x dBox, 3x cBox)
- WORK SCC: 2.2 PiB BeeGFS based on DDN SFA7990 block storage
- WORK Ceph: 600 TiB NVME CephFS and S3
- COLD: 20 PiB HDD CephFS and S3
- HSM/Tape: Quantum StorNext HSM 60+ PiB

Network for National High-Performance Computing

9 national centers für High performance computing on level Tier 2.

NHR-NORD@GÖTTINGEN

UGOE + GWDG

- Since 2021: Funding for national Tier-2 supercomputing (62.5M=C p.a.)
  - Nine centres
  - Annual funding 7,3M€ p.a.
- Usable for researchers at all German universities

Getting access

- Affiliation with German university needed
- PI of cmopute project can invite external researchers (MPG, Fraunhofer, Leibnis, other countries)
- Up to 1,200k CPU core/1500 GPU hours p.a. usable without application (test project)
- Secure Workflow for processing sensitive data (medical, financial, etc.)
- Larger projects require application
  → https://docs.hpc.gwdg.de/start_here/nhr_application_process/index.html

Getting access

Project types:

Test account: Small single user access, no formal application, just support mail, easy and fast access but no KPI benefit for the NHR site, not all storage systems usable

Starter: Small compute project, review mainly technical, full access to all storage resources, immediate start, duration limited to one year

Standard: Normal compute project, full access to all resources, scientific and technical review, quarterly start, standard runtime 4 quarters, can be extended, usually granted at preferred NHR site

Large: High compute requirements, will be located to best fitting NHR site by central authority

Application

- Scope of work for the project
- Necessity for compute ressources
- Working plan including distribution of compute usage over project lifetime
- Methods and degree of parallelization
- Required technical infrastructure (compute, storage, network, etc.)
- Whitelisting of third party funded projects, that have already shown the compute requirement in reviewm, possible, then only technical requirements

NHR at GWDG

- Possible jobsizes:
  - CPU: 1 to 256 nodes with 96 cores each, 256-2048 GB memory
  - GPU: 1 to 64 GPUs (4 per node), A100 and H100 GPUs, 40-96 GB GPU and 512-1024 GB system memory
- Storage:
  - Home: Central storage, NVME, 60 GB for every user
  - Performance: NVME, compute island specific, workspaces up to 90 days, 3 TB per user, 40+ TB per project
  - Capacity: Central HDD storage, 12+ TB per project, workspaces up to 1 year for users and projects
  - Tape: Central tape based storage for inactive data that is required again at a later stage of the project, 12+ TB per project

NHR at GWDG

Usage model

- Batch operation with Slurm (majority)
- Interactive CLI via Slurm (for testing)
- Jupyter notebooks running on CPU and GPU nodes (development)
- Virtual Desktop environment (webbrowser based), GPU hardware acceleration possible (visualization, development)

Support

- E-Mail via ticket system
- Matrix chat channel (user community with staff assistance)
- Regular webinars and VCs for different topics
- GWDG Academy courses for different skill levels

KISSKI at GWDG

AI service center for critical infrastructure

- Access to AI services with low entry level
- ChatAI via academic cloud and DFN very popular
- Chatbot, image generation, trancription, translation, coding assistance, etc.
- OpenAI compatible API
- Two hardware platforms (batch with Slurm or browser based AI services)
- Separated work environment for critical data (e.g. medical or financial data)
- Secure environment will be upgraded for data up to risk class D
- Not completely free of charge

Summary

- Wide collection of compute and storage ressources available
- Application not difficult
- Test accounts and starter projects usually within 2-3 working days
- Different support models help with usage
- No need to restrict scope of work in research due to compute ressource constraints