

# GöHPCoffe: From SCC to NHR - Emmy

Accessing and Using EMMY, The New GPU Cluster and Green500 List '22

---

Jack Ogaja  
jack.ogaja@gwdg.de

Pavan Siligam  
pavan.siligam@gwdg.de

Christian Köhler  
christian.koehler@gwdg.de

23. November 2022



Emmy (HLRN-IV) as part of NHR

Applying for NHR Resources

Miscellaneous

## Emmy (HLRN-IV) as part of NHR

---

# National High Performance Computing - NHR Alliance

- An alliance of tier 2 computing centers including, old HLRN
- Provides comp. resources, advises and trains researchers in their fields.
- Current members include:
  - RWTH Aachen (NHR4CES@RWTH)
  - TU Darmstadt - HRZ (NHR4CES@TUDa)
  - Uni Nuernberg-Erlangen -RRE (NHR@FAU)
  - Uni Goettingen - GWGDG (NHR@Goettingen)
  - KIT - SCC (NHR@KIT)
  - TU Dresden - ZIH (NHR@TUD)
  - Berlin University Alliance - ZIB (NHR@ZIB)
  - Uni Paderborn - PCP (PC2)
  - Uni Frankfurt am Main, TU Kaiserslautern, Uni Mainz & Saarlandes - South-West

# NHR's Thematic Areas of Competence and Research

- **Applications:** Physics; Engineering, Chemistry; Earth-system Sciences; Life Sciences; CFD, Social Sci. and Humanities
- **Methodological:** Programming models; Performance Analysis, Modeling and Engineering; Numerical Simulations; AI and Big-data Algorithms
- **Hardware:** Data Structures and Memory Hierarchies; GPU Programming; Vectorization Methods; MPI and Network Topologies; Memory Management and Big-data Applications; Quantum Computing; Cloud Computing; and Green IT.

# Emmy's Overview

- Over 1.2k Compute Nodes
- Large nodes with Cascade Lake CPUs
  - `standard96` (default): CLX 9242 (48 cores, Hyper-Threading), 362 GB RAM
  - `large96` (fat) CLX 9242 (48 cores, Hyper-Threading), 747 GB RAM
  - `huge96` (very fat) CLX 9242 (48 cores, Hyper-Threading), 1522 GB RAM
- Medium nodes with Skylake CPUs
  - `medium40`: SKL 6148 (20 cores cores with Hyper-Threading), 181 GB RAM
  - `large40` (fat) SKL 6148 (20 cores cores with Hyper-Threading), 764 GB RAM
- GPUs (See next slide for hardware upgrade and expansion)
  - `gpu`: 4 Tesla V100, 32 GB per GPU memory

# Current Emmy's Expansions (Grete)

## GPU expansion for the NHR system "Emmy" - Grete

- Technical specification: 36 nodes (3 racks), each equipped with
  - 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)
- Installation at RZGö
- Connection to "Emmy" storage
  - Local flash storage (Atos/DDN) for the GPU cluster at RZGö
  - Link to existing `WORK` (at MDC)
  - LNet routers for connecting to Lustre (IB↔eth↔OPA)

# Current Emmy's Expansions

## Future Technology Platform (FTP)

- Start of **F**uture **T**echnology **P**latform for testing cutting-edge HW
- Experimental applications on an alternative microarchitecture (ARM) and Security/Storage offloading (DPUs)
- Procurement of two NVIDIA ARM DevKits, each equipped with
  - 1x Ampere Altra Q80-30 (80 cores, ARM)
  - 512 GB memory (DDR4, 3200 MHz)
  - 2x NVIDIA A100 GPU (SXM4, 40 GB HBM2 memory)
  - 2x NVIDIA BlueField-2 E-Series DPU (200 GbE/HDR, 16 GB memory)
- Starting operation in late 2022



# Grete (Emmy's): Energy Efficiency and Green500 List ('22 )

Home » Lists » Green500 » November 2022 » List

## GREEN500 LIST - NOVEMBER 2022

$R_{max}$  and  $R_{peak}$  values are in PFlop/s. For more details about other fields, check the TOP500 description.

$R_{peak}$  values are calculated using the advertised clock rate of the CPU. For the efficiency of the systems you should take into account the Turbo CPU clock rate where it applies.

Green500 Data

← 1-100 101-200 201-300 301-400 401-500 →

TOP500				Energy Efficiency		
Rank	Rank	System	Cores	Rmax (PFlop/s)	Power (kW)	(GFlops/watts)
1	405	Henri - Lenovo ThinkSystem SR670 V2, Intel Xeon Platinum 8362 2800Mhz (32C), NVIDIA H100 80GB PCIe, Infiniband HDR, Lenovo Flatiron Institute United States	5,920	2.04	31	65.091
2	32	Frontier TD5 - HPE Cray EX235a, AMD Optimized 3rd Generation EPYC 64C 2GHz, AMD	120,832	19.20	309	62.684
11	349	SSC-21 Scalable Module - Apollo 6500 Gen10 plus, AMD EPYC 7543 32C 2.8GHz, NVIDIA A100 80GB, Infiniband HDR200, HPE Samsung Electronics South Korea	16,704	2.27	103	33.983
12	470	Grete - MEGWARE GPU Server System, AMD EPYC 7513 32C 2.6GHz, NVIDIA A100 SXM4 40 GB, Infiniband HDR, MEGWARE NHR@Göttingen Germany	17,856	1.83	69	32.149
13	353	Tethys - NVIDIA DGX A100 Liquid Cooled Prototype, AMD EPYC 7742 64C 2.25GHz, NVIDIA A100 80GB, Infiniband HDR, Nvidia NVIDIA Corporation United States	19,840	2.25	72	31.538



- Ranked 12th on the current Green500 list (Nov. 2022)
- Currently the most energy efficient system in Germany.
- Compute performance of 1.83 PFLOPs (HPL)
- 32.149 GFLOPs/W with Direct Liquid Cooling (DLC).
- 2nd most energy efficient among the listed NVIDIA A100 40 GB GPU systems.

## Some of the Major Applications ported to Emmy

- GROMACS, NAMD, VASP: Molecular Dynamics
- ABAQUS, ANSYS, Star-CCM+, OpenFOAM: Engineering and CFD
- ParaView: Visualization
- TURBOMOLE, RELION: Quantum Chemistry, and Electron-microscopy
- Singularity: Containerization
- Development and Performance Tools:
  - Arm DDT
  - LIKWID
  - VTune
  - Julia, Valgrind, Charm++
- More at... <https://www.hlrn.de/doc/display/PUB/Software>

# Applying for NHR Resources

---

## Emmy: Reasons to consider applying for Emmy Resources

- Part of the 9-member alliance of tier-2 supercomputing centers in Germany.
- Larger capacity - more compute nodes and storage, newer technologies and hard..
- New pool of GPUs - over 30 nodes with one of the latest architecture.
- Wider reach for consultants, incl. Berlin and tier 3 centers in the northern states.
- Wider range of applications ported: Molecular Dynamics, Chemistry, Physics, CFD, Life Sciences, Earth-system Sciences, etc

# Emmy: Applying for Personal and Project Accounts

- Requirements: Research at German universities
- Apply for an account at <https://zulassung.hlrn.de>
- Starting from 300.000 Coreh per quarter
  - Project application for further compute time
  - review by scientific board (+techn. review by HPC sites)
  - sci. review can be omitted , if project has been granted by BMBF/DFG/EU/... (Whitelisting)
- Compute time usable at NHR@ZIB (“Lise”) and NHR@Göttingen (“Emmy”)

# Miscellaneous

---

## Miscellaneous

- Data transfer: SCC to Emmy

As a first step load the ssh private key(s) with the ssh agent

(*ssh - add < privkey >*)

- For small files:

```
scp -A scc :< path_to_file > emmy :< location >
```

- For large files: on the local machine

login into transfer node which is dedicate to do the remote I/O

```
ssh -A < gwdguser > @transfermdc.hpc.gwdg.de or
```

```
ssh -A < gwdguser > @gwdud108.gwdg.de
```

transfer the data using *rsync*

```
rsync --rsh='ssh -A' -aH < local_gwdg_path > < hlrn-  
user > @glogin1.hlrn.de :< hlrn_path >
```

credits: Tim Ehlers