

Sebastian Krey

## HPC bei der GWDG in Göttingen

# GWGD

## Structure

- Joint data and IT-comptence center for the University of Göttingen, University clinic and the Max-Planck-Society
- Established in 1970
- Shareholders Max-Planck-Society and Unversity Göttingen, each 50%

## Mission

- It is the mission of GWGD 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

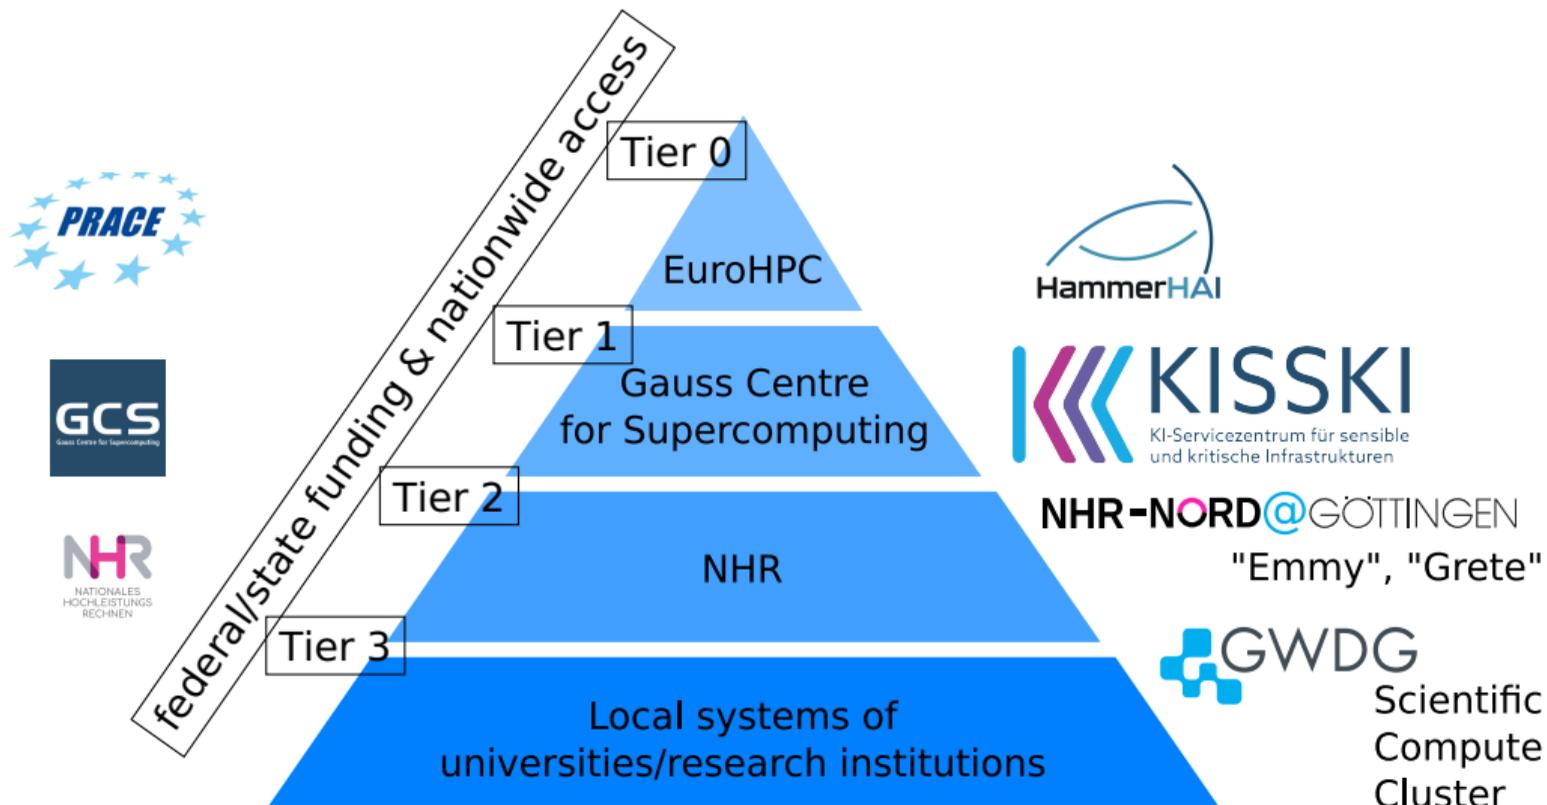
- Modern and secure IT infrastructure
- IT support for excellent research
- In-house research for innovative IT services

# GWDG

## Supra-regional tasks

- National High Performance Computing Center (NHR-NORD@Göttingen)
- National HPC Center of the DLR
- Upcoming regional HPC center for Lower Saxony (HPC.NDS)
- AI service center for sensitive and critical infrastructures (KISSKI)
- Partner in the AI Factory HammerHA1
- Data center in four NFDI consortia
- Host for DARIAH-EU, German National Library, GFBio, NUM CODEX, MWS, WirLernenOnline, Text+, etc.
- Cloud operator, including the Academic Cloud for universities in Germany

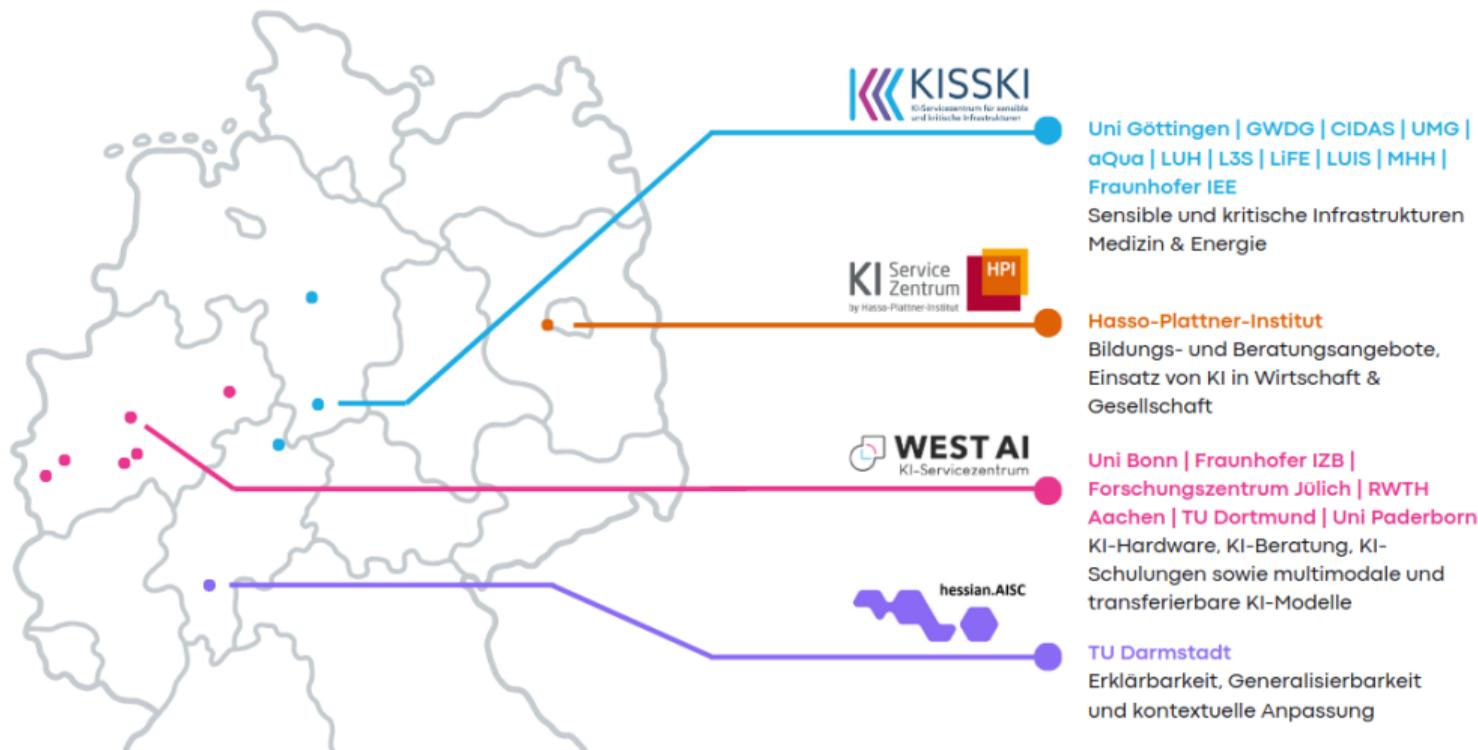
# German HPC Infrastructure



## 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 Data protection and security

5 Outlook

6 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



# CPU 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

# CPU 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: 103+3 nodes equipped with
  - ▶ 101 MEGWARE nodes (7 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, 80/40 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



# GPU System “Grete”

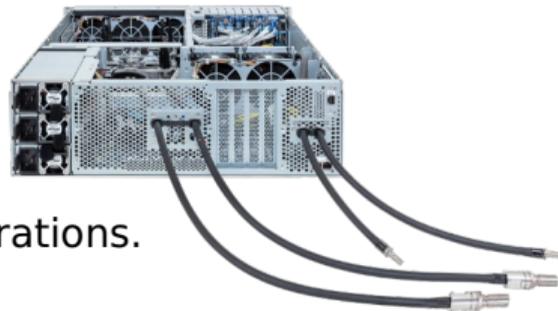
## ■ MEGWARE GPU cluster Grete

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

## ■ MEGWARE GPU Add-on 2024

- ▶ 25 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.



# “Emmy” and “Grete” NG

Next generation 2026/2027: Phase 4

- New platform for the next CPU and GPU clusters at GWDG
- Modular procurement for NHR, SCC, HPC.NDS and KISSKI
- 350+ MEGWARE Eureka CPU nodes
- 2 installation phases
- AMD Turin and Venice CPUs (2x128 cores per node)
- 768-4096GB memory
- 3,84TB NVME SSDs
- Omni-Path CN5000 400G interconnect
- 15+ GPU nodes with 8xB200 Nvidia GPUs

# 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!

## 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 and Starter projects with 300k CPU core hours or 2k GPU (Nvidia A100) hours per quarter
- Proposal preparation, Software setup
- Scientific review of compute time projects
- Own research projects

# 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



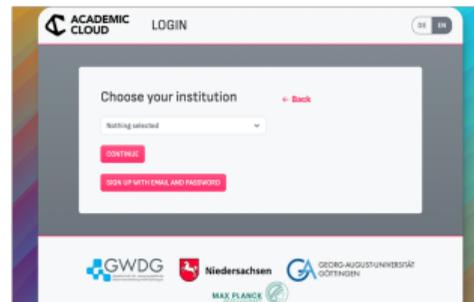
Invitation email



Login to hpcproject.gwdg.de



Federated login via Academic Cloud



Self registration is possible, too

# Project management project detail

The screenshot displays the GWDG Project Management web application. On the left is a dark sidebar with navigation options: PROJECTS (Active, Open, Archived), Join Requests, Actions logs, Users, gwdo (Privacy notice, Imprint), and a footer with the ID 127.0.0.1-8000/projects/96dd99b-d53c-4176-a90-6733a6534aa5#. The main content area shows the breadcrumb 'HPC Project Management' and search filters. The project title is 'AG Example Professor'. Below it, a 'Project tree' section states 'This project has no sub-projects.' The main project details include: Name: AG Example Professor, ID of HPC project: ag\_example\_professor (with a link to view actions file), and dates 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. A 'Share join link for this project' button is present. Below this is an 'Add users to project' section with a 'Show form' button. The 'Active project members' section contains a table with two members.

Name	Username	Email	Added	Project Username	Role
HLRN Testaccount05	hrlntest05	hrlntest05@gwdg.de	Active since: 2024-06-20, 00:00:00	3023412	PI
HLRN Testaccount04	hrlntest04	hrlntest04@gwdg.de	Active since: 2024-06-20, 00:00:00	6564379	Member

# JupyterHPC

- Quick and easy way to access HPC resources via
  - ▶ No ssh setup necessary
  - ▶ Frictionless and comfortable even for beginners
  - ▶ Very useful for courses, trainings and fallback for
- 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

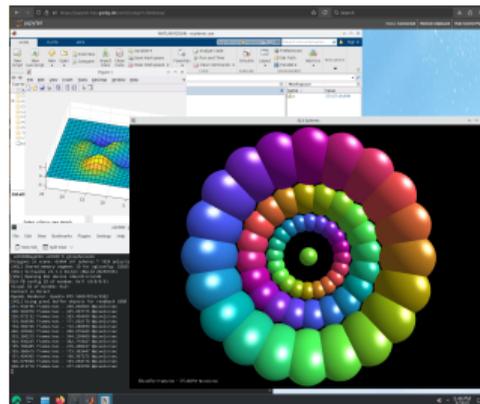
HPC Device

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 llvmpipe 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



# Data protection and security

- Secure HPC environment for batch processing of medical data
- Workflow adapted for other types of critical data
- Further improvements to simplify deployment and scale out in development
- Certification of processing data of highest non-classified data protection levels in preparation
- Critical backend system administration only by dedicated privileged admin workstations (PAW)
- Central logfiles protected against manipulation
- Next step: Securing HPC deployment with Secure Boot

# Trusted Research Environments (TREs)

- TREs provide **highly secure** environments for analyzing sensitive data
- All computations and analysis are conducted entirely within the TRE
  - ▶ Results are manually reviewed pre-publication
- Based on the "5 SAFE" framework from the UK Office of National Statistics
- Also known as:
  - ▶ **Secure Processing Environments (SPEs)** (Common in DE/EU)
  - ▶ Secure Data Environment (SDE) (Common in England; Especially NHS)
  - ▶ Data Safe Havens (Common in UK; Especially Scotland)

## Other Reasons for a TRE

### ■ Reproducibility

- ▶ Controlled Environments
- ▶ Container-Images are publishable

### ■ HPC-Resources (incl. GPU-Access)

- ▶ Big Data, for example Genome Analysis
- ▶ Machine-Learning Models

### ■ Outsourcing and Centralizing Security/Administration Burden

### ■ Legislative Pressure

- ▶ Europe: European Health Data Space
- ▶ Germany: Gesundheitsdatennutzungsgesetz

## 5 Safes

### 1 Safe People

- ▶ Trustworthy institutions and “Know your Customer”
- ▶ Completed training, Signed user agreement

### 2 Safe Projects

- ▶ No access without data owner's approval

### 3 Safe Settings

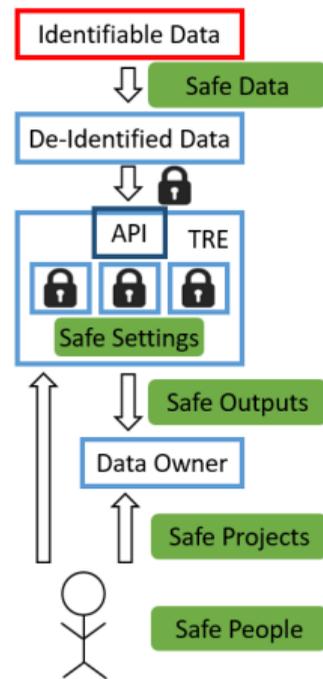
- ▶ Data is only processed within the TRE
- ▶ No internet access, no download possible

### 4 Safe Outputs

- ▶ Manual approval of exports before publication
- ▶ Focus: No re-identification; Only aggregated data

### 5 Safe Data

- ▶ Data is pseudonymized when ingested into the TRE



# On-Premise HPC-based TRE: GÖTRE

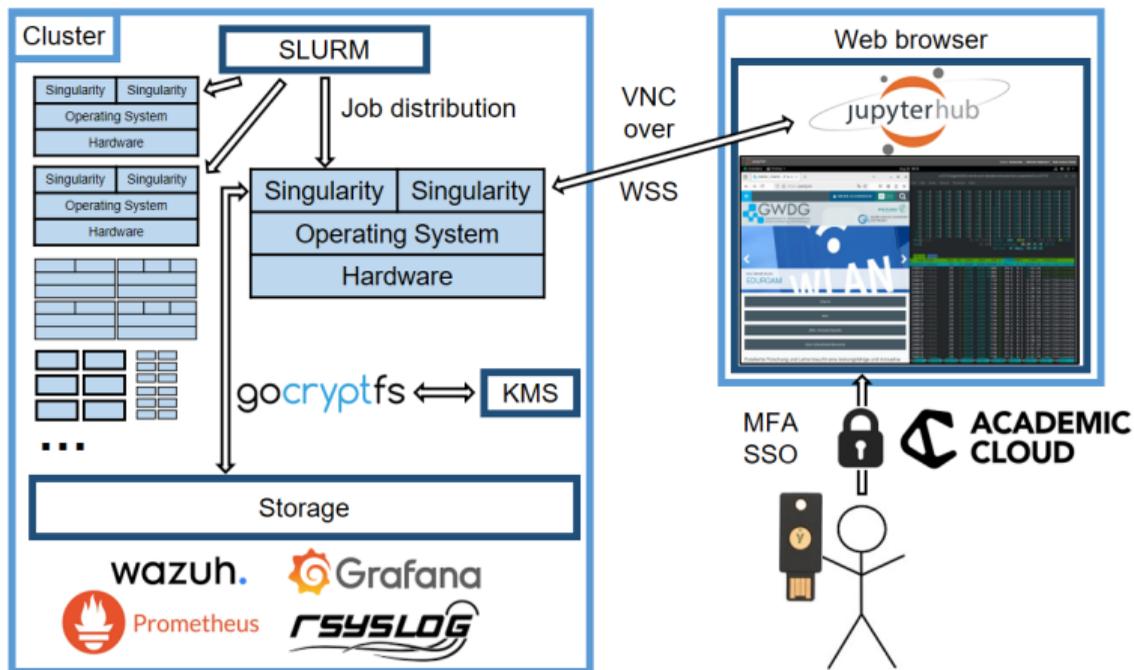
## Technical Features

- Zero-Trust Architecture
- End-to-End Encryption on file system level
  - ▶ Transparent to the user
- Multi-Factor Authentication
- Collaborative Work possible
  - ▶ Analogous to Windows Fileserver
- Complete Audit-Log
- One Key per File
  - ▶ Secure Sharing via Key Management System

## Organizational Features

- *High Usability despite Strong Security*
  - ▶ Access via web interface
  - ▶ No client software required
  - ▶ Graphical desktop environment
  - ▶ Supports non-web-based applications
- Embedded in HPC Infrastructure
  - ▶ High Performance
  - ▶ Machine-Learning support via GPUs

# GÖTRE: Architecture



# GÖTRE: Interface

The screenshot shows the GÖTRE interface within a JupyterLab environment. The interface is divided into several panes:

- Left Pane (Navigation):** A sidebar titled "HOME" with a search bar and a list of items:
  1. Start Here
  2. How to use...
  3. Software Stacks
  4. Services
  5. Project Management
  6. Training
- Main Content Area:** A "HOME" page with introductory text and an "Info" section. The text states: "This is the GWGD HPC Documentation, organized by topic both in the navigation bar to the left and below. In case you are new, welcome and you will most likely find what you are looking for under [start here](#). In case you returning or from another compute cluster or center, also welcome and you will most likely find what you are looking for under [how to use](#). If you are searching for something specific, feel free to use the search function or browse our large catalogue of content directly." The "Info" section mentions transitioning to a unified software stack and SLURM controller.
- Right Pane (RELION Job Management):** A window titled "RELION-4.0.1: /mnt/vast-standard/home/fars.quentin01/u15733" showing a 3D molecular model and a list of job actions:
  - Initial
  - Maton correction
  - CTF estimation
  - Manual picking
  - Auto-picking
  - Particle extraction
  - 2D classification
  - 3D initial model
  - 3D classification
  - 3D auto-refine
  - 3D multi-body
  - Subset selection
  - CTF refinement
  - Bayesian polishing
  - Mask creation
  - Join star files
  - Particle subtraction
  - Post-processing
  - Local resolution
  - External
 Below the list are controls for "Current" (Give\_alias\_here), "Display", "Finished jobs", and "Running jobs".
- Terminal Window:** A terminal window at the bottom showing system statistics and a table of running jobs. The statistics include:
 

```

      Mem: 13.96/7556 Tasks: 175 1222 thr: 805 kthr: 9 running
      0K/0K Load average: 16.09 15.01 14.32
      uptime: 38 days, 13:19:17
      
```

 The table lists jobs with columns: PID, USER, PRI, NI, VIRT, RES, SHR, S, CPU%, MEM%, TTY, and COMMAND.

# GöTRE: Interface

The screenshot displays a Jupyter Remote Desktop environment. The top navigation bar includes 'jupyter' and 'Status: Connected'. The main window shows a code editor with the following Python code:

```
[3]: from matplotlib import pyplot as plt
plt.suptitle('First 9 images')
for i in range(9):
    plt.subplot(330 + 1 + i)
    plt.imshow(train_X[i], cmap=plt.get_cmap('gray'))
plt.show()
```

Below the code, a plot titled 'First 9 images' displays a 2x3 grid of handwritten digits: 5, 0, 4, 1, 9, 2.

The right-hand terminal window shows the following output:

```
Epoch 4/15
422/422 ----- 6s 13ms/step - accuracy: 0.9791 - loss: 0.0695 -
val_accuracy: 0.9893 - val_loss: 0.0408
Epoch 5/15
422/422 ----- 6s 13ms/step - accuracy: 0.9825 - loss: 0.0588 -
val_accuracy: 0.9882 - val_loss: 0.0406
Epoch 6/15
422/422 ----- 6s 13ms/step - accuracy: 0.9836 - loss: 0.0537 -
val_accuracy: 0.9902 - val_loss: 0.0344
Epoch 7/15
422/422 ----- 6s 13ms/step - accuracy: 0.9842 - loss: 0.0502 -
val_accuracy: 0.9917 - val_loss: 0.0347
Epoch 8/15
422/422 ----- 6s 13ms/step - accuracy: 0.9856 - loss: 0.0462 -
val_accuracy: 0.9913 - val_loss: 0.0326
Epoch 9/15
422/422 ----- 6s 13ms/step - accuracy: 0.9858 - loss: 0.0446 -
val_accuracy: 0.9927 - val_loss: 0.0309
Epoch 10/15
422/422 ----- 6s 14ms/step - accuracy: 0.9876 - loss: 0.0399 -
val_accuracy: 0.9898 - val_loss: 0.0329
Epoch 11/15
422/422 ----- 6s 15ms/step - accuracy: 0.9872 - loss: 0.0389 -
val_accuracy: 0.9927 - val_loss: 0.0319
Epoch 12/15
422/422 ----- 6s 15ms/step - accuracy: 0.9884 - loss: 0.0354 -
val_accuracy: 0.9912 - val_loss: 0.0328
```

# Outline

1 Intro

2 Services

3 Usability

4 Data protection and security

**5 Outlook**

6 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 Data protection and security

5 Outlook

**6 Conclusion**

# Summary

- GWDG offers various HPC resources
- There are many ways to get access
- Besides performance, usability and availability can matter
- Providing highest data protection standards important esp. for medical research
- Combination of data protection and good usability possible
- If in doubt, ask our [Support](#)