

# Holistic HPC I/O

Introduction

Mohammad Hossein Biniiaz, Kevin Lüdemann



# Table of contents

- 1 The GWDG
- 2 The trainers
- 3 Scalable Storage Competition (SSC)

# Learning Objectives


- Understand storage in HPC
- Working with Storage Tiering
- Learn about parallel file systems
- Investigate performance characteristic
- Perform IO Benchmarking
- Investigate Storage API

# The GWDG The Gesellschaft für wissenschaftliche Datenverarbeitung mbH Göttingen

- Collaborative Foundation: A joint venture of the University of Göttingen, Universitätsmedizin Göttingen (UMG), and the Max Planck Society.
- IT Services:
  - ▶ Hosting/Housing
  - ▶ High-Performance Computing, Cloud Solutions
  - ▶ Procurement, Consulting
- Research Excellence:
  - ▶ Computer Science
  - ▶ AI/ML, Bioscience, Digital Humanities
  - ▶ Earth System Science, Forestry
- Over 200 experts across eight innovative working groups (AG Computing).



# HPC systems at GWDG



- Tier 2: **HLRN-IV “Emmy”**  
Top500 #47 Nov. 2020, 5.95 PFlop/s
- Tier 2: **NHR “Grete”**  
Top500 #470 Nov. 2022, 1.83 PFlop/s,  
Green500 #12,
- Tier 3: **Scientific Compute Cluster (SCC)**
- **“CARO” for DLR**  
Top500 #135 Nov. 2021, now #163

# About us

## ■ Mohammad Hossein Biniiaz

- ▶ AI developer
- ▶ Trainer and AI researcher
- ▶ Working and providing LLM service
- ▶ Email: **mohammad-hossein.biniiaz@gwdg.de**


## ■ Kevin Lüdemann

- ▶ Ph.D. in physics
- ▶ Trainer and researcher
- ▶ Working on standardized training across Europe
- ▶ Email: **kevin.luedemann@gwdg.de**

# Scalable Storage Competition

**Scalable  
Storage**

**Competition**



# SSC

- Storage systems are pivotal for HPC systems
- Performance of storage significantly influences system
- It has to be scalable, adaptable and fast
- Ever growing demand for higher capacity and faster storage



# Beyond the summer school

- Form a group and take part in the competition
- Get root access to state of the art hardware
- One full month access to set up and optimize server
- Run standardised benchmarks and compete against others
- There will be an additional Summer School for onboarding (09.2024)

# What we provide

- 10 Storage servers
- 24 Compute servers
- Base installation of Rocky 8
- Admin support for you to set up own installations
- A cash prize at the end!

# Servers

## ■ Compute (24 Nodes)

- ▶ 2x Intel Xeon Gold 6148 SKL-SP @ 2.40GHz
- ▶ 192GB RAM
- ▶ 1x System S-ATA SSD disk 480GB
- ▶ 100GB/sec Omni-Path

## ■ Storage Server (10 Nodes)

- ▶ Same CPU, RAM, and system SSD
- ▶ 1x NVMe DISK 960GB
- ▶ 9x NVMe DISK 3.6TB
- ▶ 2x100GB/sec Omni-Path for storage

# SSC

- Take part! It is free!
- <https://ssc.vi4io.org/>

