GEORG-AUGUST-UNIVERSITÄT GÖTTINGEN
IN PUBLICA COMMODA SEIT 1737

GWDG
Gesellschaft für wissenschaftliche Datenverarbeitung mbH Göttingen

Sebastian Krey

VAST at GWDG

Outline

1 About GWDG

2 Current HPC Storagesystem at GWDG

3 VAST Hardware at GWDG

4 Usage of VAST at GWDG

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

Dr. Sebastian Krey

- HPC System Architect and Admin
- Team lead of HPC admins
- Storage
- Network
- Coordination of hardware near operation
- Contact to infrastructure teams
- Background in statistics and machine learning
- Optimization of operation
  (energy efficiency and performance)

HPC systems at GWDG

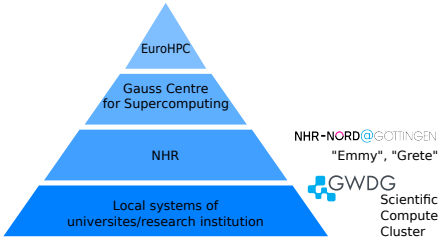

- Tier 2: **HLRN/NHR “Emmy”**
  Top500 #47 Nov. 2020, now #187
- Tier 2: **NHR/KISSKI “Grete”**
  Top500 #141 Nov. 2023, Green500 #22, now #194/#51
- Tier 2: **NHR/KISSKI “Grete Phase 3”**
  Top500 #274 Nov. 2024, Green500 #24
- Tier 3: **Scientific Compute Cluster**
- **“CARO”** for **DLR**
  Top500 #135 Nov. 2021, now #295
- Several smaller systems for Max-Planck- and Uni Göttingen Institutes

HLRN-IV “Emmy”

Uni Göttingen/GWDG

- TOP 500: #47 in 2020-11 (5.95 PFlop/s), now #187,
  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
- 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
  - CoolIT DLC

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
- CoolIT DLC with direct free cooling with outside air for residual heat

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
- CoolIT DLC

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
- CoolIT DLC

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)
- CoolIT DLC

Storage Systems

- WORK MDC: DDN ExaScaler 5 EoL Q1 2025
  - Metadata SFA7700X
  - 8 PiB HDD 2x ES14KX
  - 113 TiB NVME 2x SFA200NV
- WORK MDC new: 7 Celestica SC6100 1.3 PiB NVME (from 04/25)
- 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
- HOME SCC: 3 PiB Quantum StorNext
- HSM/Tape: Quantum StorNext HSM 60+ PiB

VAST at GWDG

After successful PoC in 2022 $\rightarrow$ candidate for the the KISSKI project.

Start of KISSKI project in November 2022, immediate tender, resulting in:

- Lightstream dBox (600TB)
- 1 cBox for IB (4x100G)) for KISSKI training platform
- 1 cBox for Eth (4x100G) for KISSKI inference platform and external services
- Unified namespace for NFS and S3

VAST Expansion

After good experiences with the VAST storage also candidate for the unified HOME storage for all HPC operations at GWDG, which should also serve the central software installations as well as storage for central services (batch system, project portal, node deployment, etc.), which was done with GPFS until then.

Open tender allowing all enterprise storage systems, winning offer:

- 2 Ceres dBoxes (2x330 TB)
- 1 cBox for IB (8x100G) for improving performance of IB

VAST for KISSKI

- AI needs an all flash platform for random read intensive workloads
- Model training with a lot of GPUs requires high performance
  $\rightarrow$ parallel filesystem
- Different hardware architecture for research hardware
  $\rightarrow$ NFS with high performance
- External users and cooperation partners require easy data ingest $\rightarrow$ S3
- Two operation sites require (synchronous) data replication (not yet implemented)
- Inference platform using cloud technologies $\rightarrow$ object storage

VAST storage platform could fulfill all criteria.

VAST as central HPC storage

- HOME storage for all user groups
- Central Spack based software installation
- Slurm state save location
- High Performance storage for read intensive workloads
- Central data pools

Strengths of VAST

- Consistent high read performance
- Easy management (Web GUI, CLI and REST API)
- Storage system updates without interruption of operation
- Completely unattended updates
- Snapshots without overhead
- Backup via off-site replication (native or S3)
- Efficient capacity usage due to compression an deduplication
- Directory based quota
- Competent and helpful support contacts
- Very good monitoring (performance, storage, etc.)

Weaknesses of VAST

- Rather low sequential write speed
- Backup via S3 has fixed infinite retention period
- Other backup options require manual scripting based on snapshots
- Only Ethernet or Infiniband as interconnect
- Very long chassis, so high rack depth needed
- Management of multiple tenants via Web UI not perfect
- Sometimes hickups in the monitoring UI

Storage efficiency

HOME:	4:1 - 5:1
Software:	Spack based software tree 10:1
WORK:	1.3: - 2:1
Total:	in our system 2:1

Performance numbers

```
IO500 version io500-sc24_v1 (extended)
[RESULT]       ior-easy-write        8.509817 GiB/s : time 485.071 seconds
[RESULT]    ior-rnd4K-write          0.199971 GiB/s : time 399.239 seconds
[RESULT]    mdtest-easy-write       48.419821 kIOPS : time 556.477 seconds
[RESULT]    ior-rnd1MB-write         5.107564 GiB/s : time 618.018 seconds
[       ]  mdworkbench-create       41.180540 kIOPS : time 71.066 seconds
[       ]            timestamp       0.000000 kIOPS : time 0.002 seconds
[RESULT]            find-easy      733.978262 kIOPS : time 36.655 seconds
[RESULT]       ior-hard-write        1.396726 GiB/s : time 351.203 seconds
[RESULT]    mdtest-hard-write        2.251237 kIOPS : time 493.604 seconds
[RESULT]                find      482.430370 kIOPS : time 64.076 seconds
[RESULT]       ior-rnd4K-read        0.687744 GiB/s : time 116.153 seconds
[RESULT]     ior-rnd1MB-read        22.772360 GiB/s : time 138.681 seconds
[RESULT]            find-hard       27.037794 kIOPS : time 41.060 seconds
[RESULT]    mdworkbench-bench       58.178553 kIOPS : time 407.360 seconds
[RESULT]        ior-easy-read       51.239343 GiB/s : time 80.636 seconds
[RESULT]     mdtest-easy-stat      206.995858 kIOPS : time 130.926 seconds
[RESULT]        ior-hard-read        4.877447 GiB/s : time 100.632 seconds
[RESULT]     mdtest-hard-stat      110.068541 kIOPS : time 11.112 seconds
[       ]  mdworkbench-delete       35.186007 kIOPS : time 83.312 seconds
[RESULT]   mdtest-easy-delete       29.257199 kIOPS : time 926.599 seconds
[RESULT]     mdtest-hard-read       68.286614 kIOPS : time 17.283 seconds
[RESULT]   mdtest-hard-delete        2.152428 kIOPS : time 517.102 seconds
[SCORE ] Bandwidth 7.382558 GiB/s : IOPS 38.815045 kiops : TOTAL 16.927916
[SCOREX] Bandwidth 3.842424 GiB/s : IOPS 50.905568 kiops : TOTAL 13.985735
```