GEORG-AUGUST-UNIVERSITÄT GÖTTINGEN IN PUBLICA COMMODA SEIT 1737

GWDG Gesellschaft für wissenschaftliche Datenverarbeitung mbH Göttingen

Julian Kunkel

Welcome

Outline

Prof. Dr. Julian Kunkel

- Deputy Head of GWDG - HPC
- Group Leader Computing Working Group
- Professor at the University of Göttingen
  - Institute for Computer Science
  - Research Group **High-Performance Storage**
- Research goals
  - Data-driven workflows
  - Storage & Parallel file systems
  - Performance analysis of parallel apps and I/O
  - Performance portability
  - Application of machine learning methods
  - Data reduction techniques
  - Management of cluster systems
  - Software engineering of scientific software

Research Group: High-Performance Storage

Research goals

- Data-centric Input/Output architectures
- Efficient execution of data-driven workflows
- Autonomous storage systems making intelligent decisions

Noteworthy involvement

- The Virtual Institute for I/O vi4io
- The Journal of High-Performance Storage J∞HPS https://jhps.vi4io.org/
- Development of IOR/MDTest/MDWorkbench
- The IO500 Benchmark IO500 https://io500.org

HPS https://hps.vi4io.org

The GWDG: service organization and data and IT service center

- works in conjunction with Universität Göttingen and Max Planck Society
- carries out independent research in the field of computer science
- provides support in preparing future professionals for a career in IT
- employs about 200 experts in 8 working groups

https://www.gwdg.de

GWDG Offered Data Services

- (Research) data Management support (DMPs, publishing)
- Offered primary storage, archive, backup
- Data sharing tools (OwnCloud, S3)
- Privacy / Security
- Consulting

The GWDG: Non-HPC storage environment (1)

- Primary storage
  - Capacity: 61.5 PB on HDDs + 2.8 PB on SSDs (gross capacity)
  - Solutions: NetApp FAS/AFF, virtualized SAN / Quantum StorNext, SDS (Ceph)
    - Different performance characteristics, requirements, features
    - Cost optimization
  - 25% growth p.a.
- Tape
  - 114 PB for backup, archival and HSM (gross capacity)
  - 4 x Quantum libraries + 2 x IBM libraries

The GWDG: Non-HPC storage environment (2)

- NetApp FAS/AFF environment
  - Critical enterprise workloads
    - VMware vSphere (3200+ VMs)
    - Kubernetes (30+ clusters)
    - Critical fileservices and apps
  - Protocols: NFS, SMB
  - 6.5 PB HDD + 0.6 PB SSD
- Virtualized SAN / Quantum StorNext environment
  - Block storage + HSM file system for various scientific use cases
  - Protocols: NFS, SMB, iSCSI, FC
  - 32.5 PB HDD + 0.5 PB SSD

The GWDG: Non-HPC storage environment (3)

- Software Defined Storage (Ceph)
  - Scientific use cases
    - Cold storage ("Data Lakes", HSM)
    - Hot storage for virtualization (OpenStack Cloud, Kubernetes)
    - S3 for global data sharing / distribution / access
  - Increased use also for non scientific use cases due to cost advantage
  - Protocols: S3, CephFS, RBD
  - 22.5 PB HDD + 1.7 PB SSD

Outline

HPC systems at GWDG

- NHR / HLRN-IV System Emmy (CPU) + Grete (GPU) → NHR@Göttingen
  - Tier-2 System
  - https://www.top500.org/system/179883/
  - https://www.top500.org/system/180092/
- Scientific Compute Cluster (SCC)
  - Tier-3 system for University of Göttingen und Max Planck Society
- DLR System CARO
  - Tier-2 System for DLR
  - https://www.top500.org/system/180038/
- Campus Institut for Dynamics of Biological Networks (CIDBN)
- Housing of various cluster systems

RZGö

- Energy-efficient data centre
- Modern concepts
  - Cooling: hot, free
  - Waste-heat usage - heating neighboring buildings
  - Cooling central with ice
  - Security - ISO27001

The GWDG: HPC storage environment

- Work storage for scratch and project storage space
  - Emmy: DDN Lustre (8.5 PiB HDD, 130 TiB NVME)
  - Grete: DDN Lustre (130 TiB NVME and LNET routing to Emmy storage)
  - CARO: DDN Lustre (8 PiB HDD, 200 TiB SSD)
  - KISSKI: VAST Data (600 TiB, 1 Lightstream dBox, 1 cBox IB, 1 cBox 100GbE)
  - SCC: BeeGFS based on DDN blockstorage (2.1 PiB HDD, 100 TiB SSD)
- Home storage
  - Tier 2: DDN GridScaler (GPFS/Storage Scale via NFS, 350TiB)
  - Tier 3: GWDG central UNIX home directory (Quantum StorNext, DLC and NFS)
- Archive
  - Tier 2: Quantum Tape Library with StorNext HSM (7PB gross)
  - Tier 3: GWDG central StorNext HSM
- Upcoming: Large central cold storage (around 20GB gross, likely Ceph based) and caching system for fast access to GWDG central S3 space

Sponsoring

The Monthly Storage Talks are powered by

Motivation for the Monthly Storage Talks

- HPC base of many research projects
- Lack of proper education on HPC topics
- Bringing together experts and users
- Improving parallel I/O in scientific workflows
- Overcoming obstacles for both hardware and the software stack
- Increase competence of HPC users
- Organization of monthly NHR Data Lakes meetings

About the Monthly Storage Talks

Key information

- Contact: Julian Kunkel, Patrick Höhn
- Location: Online (`https://meet.gwdg.de/b/pat-rru-pyl-fko`)
- Time: First Tuesday, every month, 16:00, 1 hour (exceptions are announced)
- Language: English (German, if only German participants)

Topics of interest

- Scientific workloads
- Usage characteristics (file, folders, scientific libraries)
- Performance aspects and monitoring
- ...?

About the Monthly Storage Talks

Interactivity

- Critical discussion is welcome and expected from attendees
- Discussion time slots: open topics, everyone can raise/discuss issues
  - Ultimately controlled by a moderator

Agenda

- 2024-01-10 - Informal Meeting about the monthly storage talks
  - Goal is to identify the next topics, interest in this area...
- 2024-02-06 - Best Practices in Organizing I/O – Moderator: Patrick Höhn
  - Talk: Best Practices in Organizing I/O for ML Projects – Giorgi Mamulashvili
- 2024-03-05 - Using Ceph Storage – TBD
- 2024-04-02 -

Todays Meeting

- Critical discussion is welcome and expected from attendees
- Discussion time slots: open topics, everyone can raise/discuss issues
- Presentations and moderation open to community memembers
- Feedback?