



# *Site Update: Zuse Institute Berlin*

*DAOS User Group 2023*

*Steffen Christgau*

*Supercomputing Department, Zuse Institute Berlin*

# Where we started

Zuse Intitute Berlin operates 8 PFLOPs Tier-2 HPC system *Lise*

- most powerful system at German Tier-2 level; academic use only
- 1200+ nodes, GPU extensions (A100, PVC)

ZIB's DAOS Journey started early

- DAOS experience since July 2019 (pre v0.6)
- Research work on two-node exploration environment: Apache Pass + **Omni-Path** HFI
  - Enable DAOS usage for key HPC application *PALM* → IXPUG 2019, DUG 2019  
MPI/HDF5 middleware evaluation
- Extension of *Lise* with **DAOS system**:
  - 20 Dual-Socket CLX nodes with two **Omni-Path** HFIs, 12 AP-DIMMs, 4 NVMe each
  - about 0.5 PB total capacity → DUG'20

- Ongoing **research interest**
  - Project proposal for pre-exascale/exascale-ready code using key-value stores (rejected)
  - Interested in usage of DAOS key-value functionality → collaboration with U Potsdam + German Research Center for GeoScience (Bachelor thesis finished)
  - Recently: performance scalability tests via TCP (see below)
  
- Primary **technical hurdle: OPX + DAOS compatibility**
  - *Lise* has Omni-Path fabric
  - OPX – new native libfabric provider for Omni-Path (PSM2 replacement) by Cornelis Networks
  - OPX provider tests in summer 2023: crashes for >1 DAOS servers
  - Compatibility issue not (yet) resolved → waiting for progress

# What we learned

- Things improved: packages, documentation, (error) messages → feedback at DUG'20
- Tools with improved capabilities → easier life for admins and users
- Helpful support: Special thanks to Michael Hennecke

Thanks for your attention. Questions?