

Site Update

HLRN-IV "Lise" @ ZIB

North German Supercomputing Alliance


Thomas Steinke, Zuse Institute Berlin

2020 IXPUG US Annual Meeting, October 13-16, 2020

HLRN: 7 Federal States – 1 HPC System



founded in 2001



HLRN-IV "Lise"

final configuration April 2020



- 1,270 nodes, Intel CLX AP
- 121,920 cores,
- 1.3+ MWs
- 8 PFlop/s peak perf.
- 8 PB Lustre
- 1,5 t / rack
- warm water cooling
- #48 in TOP500 (06/2020)

Key Properties of the HLRN-IV System



8 (16) PFlop/s peak performance (5 ×)

3.5 × sustained application performance



121,920 (>203,000) compute cores (2.5 ×)




8 (16) PByte on-line permanent storage capacity (2 ×)

Total numbers for both HLRN compute complexes at ZIB and Univ. Göttingen



HLRN User Community

- 3,000+ users
- 200+ large projects
- **Physics** (astro, high-energy)
- **Chemistry** & material science
- **Engineering**
- **Earth Science** (incl. climate)
- **Life Science** (biology, medicine)



COVID-19



14.10.2020



steinke@zib.de



... and then came



DAOS IO Testbed @ZIB

Total capacity:

0.5 PB SSD

30 TB DCPM

Expected Performance:

R/W: 200 / 174 GB/s ($\approx 2\times$ over 8 PB Lustre)

9 Mio IOPS ($10^x \times$ over Lustre w/o BB)

- 20x Intel Server R1000WF, attached to OPA fabric of HLRN-IV "Lise"
 - ❖ 2x Cascade Lake Gold 6240R, 24c@2,40 GHz • DRAM: 192 GB DDR4
 - ❖ 1,5 TB DCPM
 - ❖ 25,6 TB NVMe NAND SSD
 - ❖ 2x 100 Gb/s OPA HFI



Why DAOS?

- Our current Lustre installation w/o Burst Buffer
 - Poor IOPS performance
- Complementing or alternative solutions to Lustre:
 - ❖ add BB, BeeGFS/BeeOND on SSDs, DAOS, ... (within our horizon)
- Workloads that can benefit from DAOS
 - ❖ Large eddy simulations with PALM (MPI-IO / netCDF DAOS container)
 - ❖ Radiative transfer simulations for stellar atmospheres (PHOENIX)
 - ❖ CFD/FEM workloads with frequent I/O on many small files
 - ❖ AI/DL workflows



THANKS!

???