



SC22

Dallas, TX | hpc
accelerates.

Analyzing Parallel I/O BoF

Shane Snyder (ANL)

Julian Kunkel (GWDG)

November 17, 2022

Motivating this BoF

- An ability to understand parallel I/O performance is critical to ensuring efficient use of HPC resources
 - However, users, system admins, and I/O researchers are often ill-equipped for understanding an increasingly complex HPC I/O ecosystem
- Complex systems
 - Increasingly diverse application I/O workloads
 - Systems employing deep storage hierarchies utilizing emerging storage hardware
 - I/O middleware to bridge the gap
- In this BoF, we attempt to form a diverse community of parallel I/O stakeholders to discuss state-of-the-art in analyzing parallel I/O, with the following goals:
 - Inform community of recent advances in tools/techniques for I/O monitoring
 - Discuss experience/limitations of current approaches
 - Derive a roadmap for future I/O tools to capture, analyze, predict, and tune I/O



Agenda for today's BoF

- Talks
 - “Detecting Data Races on Storage Systems Using Recorder”, Chen Wang (LLNL)
 - “Visualizing I/O Bottlenecks with DXT Explorer 2.0”, Jean Luca Bez (LBL)
 - “Darshan I/O Runtime Monitoring”, Ann Gentile (SNL)
 - “Non-Intrusive Monitoring and I/O Classification with IOFS”, Christian Boehme (GWDG)
 - “I/O Monitoring with VAST”, Rob Mallory (VAST)
- (Audience-driven discussion)



Other events of interest @ SC'22

- Technical program:
 - “A Taxonomy of Error Sources in HPC I/O Machine Learning Models”, Isakov et al. (Tuesday, 4:00 PM)
- Workshops:
 - PDSW (Monday)
 - “Drishti: Guiding End-Users in the I/O Optimization Journey”, Bez et al.
- BoFs:
 - “IO500: The High-Performance Storage Community”, Hildebrand et al. (Tuesday, 5:15 PM)
- Tutorials:
 - “Parallel I/O in Practice”, Latham et al. (Sunday)



Thanks to all for attending!

BoF website will be updated with speaker materials, notes, etc.:

<https://hps.vi4io.org/events/2022/sc-analyzing-io>

