



Analyzing Parallel I/O BoF

Julian Kunkel (GWDG)
Shane Snyder (ANL)



November 16, 2021

SC'21, St. Louis, MO

Motivation for being here

- The ability to understand parallel I/O performance is critical to ensuring efficient use of HPC resources, yet users, system admins, and other I/O researchers are often ill-equipped for understanding an increasingly complex HPC I/O ecosystem
- Complex systems
 - Systems employing deep storage hierarchies utilizing emerging storage hardware
 - I/O workloads becoming more diverse to support application use cases
 - checkpoint-restart, AI/DL, in-situ, etc.
 - 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 in mind:
 - 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
 - "Darshan traces visualization", Jean Luca Bez (LBNL)
 - "Experience with the NERSC all-flash system", Glenn Lockwood, Alberto Chiusole (NERSC)
 - "Exploring ways to use the IO500 benchmark and IO500 list", Radita Liem (RWTH Aachen)
 - "Determining I/O Weather", Julian Kunkel (GWDG/Uni-Göttingen)
 - "I/O Monitoring with PIKA", Sebastian Oeste (TU Dresden)
- Speaker panel, led by Julian
 - Please let us know what questions or comments you have for us in the chat!

Other events of interest at SC

- Technical program:
 - “Systematically Inferring I/O Performance Variability by Examining Repetitive Job Behavior”, Costa et al. (**Tuesday, 4:30 PM**)
- Workshops:
 - PDSW (**Monday**)
 - Numerous relevant talks, including from a number of our speakers (Glenn, Alberto, Radita, and Jean Luca)
 - DRBSD (**Sunday**)
 - “Unbalanced Parallel I/O: An Often-Neglected Side Effect of Lossy Scientific Data Compression”, Wang et al.
- BoFs:
 - “The IO-500 and the Virtual Institute of I/O”, Kunkel et al. (**Thursday, 12:15 PM**)
- Tutorials:
 - “Parallel I/O in Practice”, Latham et al. (**Monday**)
- Doctoral showcases:
 - “Accelerating I/O for Traditional HPC and Modern ML Workloads on Emerging HPC Systems”, Wei Der Chien (**Thursday, 4:30 PM**)

Thanks to all for attending!

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

<https://hps.vi4io.org/events/2021/sc-analyzing-io>