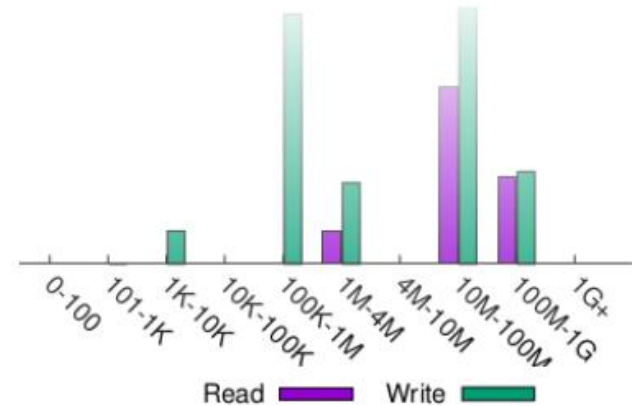


NOVEMBER 13, 2018

ANALYZING PARALLEL I/O BOF

PHIL CARNS
Argonne National Laboratory

JULIAN KUNKEL
University of Reading



WHY IS ANALYZING PARALLEL I/O SO IMPORTANT?

(and what are we doing here?)

- Application performance improves dramatically when we understand and tune it's I/O for a given system.
- But HPC storage systems aren't getting any simpler!
 - **Summit** (ORNL): 250 PiB of GPFS storage, node-local SSDs, IB fabric
 - **Perlmutter** (LBNL): all-flash Lustre storage, Slingshot fabric
- Upcoming exascale machines will push the frontiers of storage architecture even further.



We must continue to improve our tools, techniques, and best practices for new users to get the most out of new storage systems.

TODAY'S TOPICS

To that end, we have assembled expert guests from industry, facilities, and academia to share their experiences and address the following questions:

- 1. *What lessons have you learned from your work?***
(for tool developers, facility operators, or users)
- 2. *What is needed next from the analyzing parallel I/O community?***
(what would make your work easier in the future)
- 3. *<INSERT YOUR BIG QUESTION HERE>***
(audience participation in the panel discussion!)

We hope that by sharing new innovations and common goals, we can continue to keep pace with emerging systems and applications.

TODAY'S SPEAKERS AND PANELISTS

- **Using Benchmarks to Understand Performance Behavior**
 - Julian Kunkel (University of Reading)
- **A Peek into Workflow I/O**
 - Jakob Lüttgau (DKRZ)
- **I/O profiling in the field with Ellexus**
 - Rosemary Francis (Ellexus)
- **Monitoring of I/O with Lustre**
 - Andreas Dilger (WhamCloud)
- **IO Workload Throttling on Supercomputers**
 - Si Liu (TACC)



Whamcloud



RELATED EVENTS AT SC18

Analyzing parallel I/O is a hot topic!

- HUST workshop (**Sunday**):
 - Huang et al. “OOOPS: An Innovative Tool for IO Workload Management on Supercomputers”
- PDSW workshop (**Monday**):
 - Chien et al. “Characterizing Deep-Learning I/O Workloads in TensorFlow”
 - Luetzgau et al. “Toward Understanding I/O Behavior in HPC Workflows”
- Related topic BoFs (**Wednesday, 5:15pm**):
 - Gentile et al. “Monitoring Large-Scale HPC Systems”
 - Kunkel et al. “The IO-500 and the Virtual Institute of I/O”
- SC technical program (**Thursday, 4:30pm**):
 - Lockwood et al. “A Year in the Life of a Parallel File System”

THANK YOU!

PRESENTATIONS AND NOTES WILL BE ONLINE:

<https://hps.vi4io.org/events/2018/bof-analyzing>



Argonne National Laboratory is a
U.S. Department of Energy laboratory
managed by UChicago Argonne, LLC.

