

Analyzing Parallel I/O

Jean Luca Bez (Lawrence Berkeley National Laboratory)

Julian Kunkel (Georg-August-Universität Göttingen/GWDG)

Shane Snyder (Argonne National Laboratory)


November 21st, 2024




SC24
Atlanta, GA | hpc
creates.

Motivation

- An ability to understand HPC I/O performance is critical to ensuring efficient use of storage resources.
 - However, users, system admins, and I/O researchers are often ill-equipped for understanding an increasingly complex HPC I/O ecosystem.
- Complex systems
 - Diverse application I/O workloads
 - Deep storage hierarchies utilizing emerging storage hardware
 - I/O middleware to bridge the gap
- In this BoF, we attempt to form a broad 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



Visualization of entropy in Terascale Supernova Initiative application. Image from Kwan-Liu Ma's visualization team at UC Davis.



HPE/Cray Aurora system at the ALCF

Agenda

- Talks
 - “Analyzing I/O in Deep Learning Workloads”, **Hariharan Devarajan** (LLNL)
 - “Scaling Performance Analysis Tools to Hundreds of Lustre Filesystems”, **Ellis Wilson** (Microsoft)
 - “Data analysis news from GWDG and analysis tasks in the Scalable Storage Competition”, **Christian Boehme** (GWDG)
 - “Leveraging AI: Large Language Models in HPC I/O Optimization”, **Dong Dai** (University of Delaware)
- Speaker panel

Thanks to all for attending!

BoF website will be updated with speaker materials, notes, etc.:
<https://hps.vi4io.org/events/2024/sc-analyzing-io>



Are fundamentally new tools needed to help understand/improve performance of AI/ML workloads on HPC storage systems? Or can traditional approaches be used?



How can tools do a better job of providing actionable feedback to end users, in the face of constantly changing workloads and storage subsystems?



What are some of the challenges in capturing and analyzing the energy costs of I/O in existing and emerging HPC storage systems?

Would the community benefit from a standardized way to represent the captured I/O metrics and allow seamless integration with other system metrics (e.g., scheduler, network, etc) and tools?

