

Darshan Activities at Supercomputing 2016

Darshan demonstrated a strong presence at SC this year, referenced in a number of papers and presentations on parallel I/O analysis. I/O researchers, application developers, and facilities staff all recognized the utility of Darshan and expressed interest in using it to help understand and tune application I/O workloads, as well as to characterize I/O trends at HPC facilities.

Papers:

- 1- Shane Snyder, a software developer at ANL and one of the Darshan developers, [presented a paper on Darshan's newly modularized architecture](#) at the ESPT workshop. This work explains how Darshan instrumentation can be extended using a new plug-in style architecture and demonstrates new insights that are possible from a more comprehensive instrumentation of HPC I/O workloads.
- 2- James Dickson, PhD student from University of Warwick, [presented a paper on a flexible I/O proxy application \(MACSio\)](#) that can be used to automatically replicate the I/O workloads of HPC applications at the PDSW/DISCS workshop. Darshan is an integral part of this framework, with its detailed I/O characterization data being used to generate corresponding MACSio I/O workload parameters.

Posters:

- 3- Matt Bryson, a PhD student at UCSC, [presented a poster](#) at the *ACM Student Research Competition* analyzing the causes for reduced I/O performance of the popular VPIC code on the new Cori system at NERSC. This research used Darshan for characterizing VPIC I/O workload performance under different configurations to determine the root cause of reduced I/O rates.

Analyzing Parallel I/O Birds-of-a-Feather Session:

- 4- Philip Carns, a software developer at ANL and one of the Darshan developers, co-organized this session along with Julian Kunkel of DKRZ. [The purpose of the session was to share information with the broader I/O characterization community and get feedback on our ongoing research.](#)
- 5- Shane Snyder [presented on new Darshan design features that could be of interest to the community at large.](#) One upcoming feature of particular interest was an extended I/O tracing capability contributed to Darshan by Intel's High Performance Data Division (HPDD).
- 6- Glenn Lockwood, a systems staff member at NERSC, [presented ongoing research on characterizing burst buffer I/O behavior using a holistic I/O characterization framework called TOKIO.](#) The [TOKIO \(TTotal Knowledge of I/O\) project](#) is a joint collaboration between ANL and LBNL aiming to provide complete characterization of HPC I/O workloads across all components of the I/O stack, with Darshan being the component responsible for application-level characterization. The community was greatly interested in this research and applying/leveraging it within their own facilities, potentially leading to a much wider adoption of Darshan.

Parallel I/O in Practice Tutorial

- 7- Rob Ross of ANL presented Darshan as part of a parallel I/O tutorial, explaining how Darshan can be utilized to understand the behavior of HPC I/O workloads and to drive potential tuning efforts.

Evidence of Community Impact

- 8- Rajeev Thakur of ANL presented an Exascale Computing Project (ECP) overview in his keynote talk at the ExaMPI workshop, where he noted that Darshan is one of the most widely used software packages at DOE facilities.
- 9- Bilel Hadri of KAUST [presented on early burst buffer experiences of KAUST's Shaheen II system](#) at the "Burst Buffers: Early Experiences and Outlook" Birds-of-a-Feather session. He mentioned using Darshan to help track down file system performance problems on this system.