

Extreme I/O on HPC for HEP using the Burst Buffer at NERSC

Tuesday, October 11, 2016 2:30 PM (15 minutes)

In recent years there has been increasing use of HPC facilities for HEP experiments. This has initially focussed on less I/O intensive workloads such as generator-level or detector simulation. We now demonstrate the efficient running of I/O-heavy 'analysis' workloads for the ATLAS and ALICE collaborations on HPC facilities at NERSC, as well as astronomical image analysis for DESI.

To do this we exploit a new 900 TB NVRAM-based storage system recently installed at NERSC, termed a 'Burst Buffer'. This is a novel approach to HPC storage that builds on-demand filesystems on all-SSD hardware that is placed on the high-speed network of the new Cori supercomputer. The system provides over 900 GB/s bandwidth and 12.5 million I/O operations per second.

We describe the hardware and software involved in this system, and give an overview of its capabilities and use-cases beyond the HEP community before focussing in detail on how the ATLAS, ALICE and astronomical workflows were adapted to work on this system. To achieve this, we have also made use of other novel techniques, such as use of docker-like container technology, and tuning of the I/O layer experiment software.

We describe these modifications and the resulting performance results, including comparisons to other approaches and filesystems. We provide detailed performance studies and results, demonstrating that we can meet the challenging I/O requirements of HEP experiments and scale to tens of thousands of cores accessing a single storage system.

Tertiary Keyword (Optional)

Secondary Keyword (Optional)

Storage systems

Primary Keyword (Mandatory)

High performance computing

Primary author: Dr BHIMJI, Wahid (Lawrence Berkeley National Lab. (US))

Co-authors: DALEY, Chris (Lawrence Berkeley National Lab. (US)); Dr BARD, Deborah (Lawrence Berkeley National Lab. (US)); PORTER, Jeff (Lawrence Berkeley National Lab. (US)); BURLEIGH, Kaylan (Lawrence Berkeley National Lab. (US)); GERHARDT, Lisa (LBNL); FASEL, Markus (Lawrence Berkeley National Lab. (US)); NUGENT, Peter (Lawrence Berkeley National Lab. (US)); FARRELL, Steven Andrew (Lawrence Berkeley National Lab. (US)); TSULAIA, Vakho (Lawrence Berkeley National Lab. (US))

Presenter: Dr BHIMJI, Wahid (Lawrence Berkeley National Lab. (US))

Session Classification: Track 6: Infrastructures

Track Classification: Track 6: Infrastructures