

**** CALL FOR EXTENDED ABSTRACTS ****

2nd International Workshop on the Lustre Ecosystem:
Enhancing Lustre Support for Diverse Workloads

Baltimore, Maryland, USA
March 8-9, 2016

<http://lustre.ornl.gov/ecosystem>

**** IMPORTANT DATES ****

Note: all submission deadlines are midnight anywhere-on-earth

Extended abstracts due: January 15, 2016
Acceptance notification: February 5, 2016
Technical presentations: March 9, 2016
Full papers due: March 14, 2016

**** SCOPE and TOPICS ****

The Lustre parallel file system has been widely adopted by scientific high-performance computing (HPC) centers as an effective system for managing large-scale storage resources. Lustre achieves unprecedented aggregate performance by parallelizing I/O over file system clients and storage targets at extreme scales. Today, 7 out of 10 fastest supercomputers in the world use Lustre for high-performance storage.

Traditionally, Lustre development has focused on improving the performance and scalability of large-scale scientific workloads. In particular, large-scale checkpoint storage and retrieval, which is characterized by bursty I/O from coordinated parallel clients, has been the primary driver of Lustre development over the last decade. With the advent of extreme scale computing and Big Data computing, many HPC centers are seeing increased user interest in running diverse workloads that place new demands on Lustre.

This workshop series is intended to help explore improvements in the performance and flexibility of Lustre for supporting non-scientific application workloads. The 2015 workshop was the inaugural edition, and the goal was to initiate a discussion on the open challenges associated with enhancing Lustre for diverse applications, the technological advances necessary, and the associated impacts to the Lustre ecosystem. The workshop program featured a day of tutorials and a day of technical paper presentations.

In this workshop, we seek contributions that explore improvements in the performance and flexibility of the Lustre file system for supporting diverse workloads. This will be a great opportunity for the Lustre community to discuss the challenges associated with enhancing Lustre for diverse applications, the technological advances necessary, and the associated ecosystem.

Specific topics of interest include, but are not limited to:

- Using Lustre as a Shared Resource
- Adaptability and Scalability of Lustre for Diverse Workloads
- Resilience and Serviceability of Lustre
- Knowledge Provenance in Lustre
- Application-driven Lustre Benchmarking
- Integrating Big Data Technologies with Lustre
- Performance Monitoring Tools for Lustre

**** SUBMISSION GUIDELINES ****

Authors should electronically submit extended abstracts of previously unpublished work in PDF format. Abstracts may consist of up to 5 US letter-size (8.5 by 11in.) pages, including figures, tables, and references. Submissions should be formatted in IEEE Computer Society conference proceedings style.

Accepted abstracts are expected to be expanded into technical papers of at most 10 pages in length, all-inclusive. Full papers will be due the week following the workshop, and will undergo shepherding in preparation for publication in the workshop proceedings.

**** ORGANIZING COMMITTEE ****

==PROGRAM CO-CHAIRS ==

Neena Imam, Oak Ridge National Laboratory, USA
Michael Brim, Oak Ridge National Laboratory, USA
Sarp Oral, Oak Ridge National Laboratory, USA

== TUTORIALS CHAIR ==

Richard Mohr, University of Tennessee, USA

**** TECHNICAL PROGRAM COMMITTEE ****

<TBD>