

The 1st Joint International Workshop on Parallel Data Storage and Data Intensive Scalable Computing Systems (PDSW-DISCS'16)

Held in conjunction with SC16: The International Conference for High Performance Computing, Networking, Storage and Analysis, Salt Lake City, UT, and in cooperation with SIGHPC.

Monday, November 14, 2016 9:00am - 5:30pm

General Co-Chairs:

Garth Gibson (Carnegie Mellon University)

Yong Chen (Texas Tech University)

Program Co-Chairs:

Shane Canon (Lawrence Berkeley National Laboratory)

Dean Hildebrand (IBM Research)

URL: <http://www.pdsw-discs.org>

Workshop Abstract:

We are pleased to announce that the first Joint International Workshop on Parallel Data Storage and Data Intensive Scalable Computing Systems (PDSW-DISCS'16) will be hosted at SC16: The International Conference for High Performance Computing, Networking, Storage and Analysis. The objective of this one day joint workshop is to combine two overlapping communities and to better promote and stimulate researchers' interactions to address some of the most critical challenges for scientific data storage, management, devices, and processing infrastructure for both traditional compute intensive simulations and data-intensive high performance computing solutions. Special attention will be given to issues in which community collaboration can be crucial for problem identification, workload capture, solution interoperability, standards with community buy-in, and shared tools.

Many scientific problem domains continue to be extremely data intensive. Traditional high performance computing (HPC) systems and the programming models for using them such as MPI were designed from a compute-centric perspective with an emphasis on achieving high floating point computation rates. But processing, memory, and storage technologies have not kept pace and there is a widening performance gap between computation and the data management infrastructure. Hence data management has become the performance bottleneck for a significant number of

applications targeting HPC systems. Concurrently, there are increasing challenges in meeting the growing demand for analyzing experimental and observational data. In many cases, this is leading new communities to look towards HPC platforms. In addition, the broader computing space has seen a revolution in new tools and frameworks to support Big Data analysis and machine learning.

There is a growing need for convergence between these two worlds. Consequently, the U.S. Congressional Office of Management and Budget has informed the U.S. Department of Energy that new machines beyond the first exascale machines must address both the traditional simulation workloads as well as data intensive applications. This coming convergence prompts integrating these two workshops into a single entity to address the common challenges.

The scope of the proposed joint PDSW-DISCS workshop is summarized as:

- Scalable storage architectures, archival storage, storage virtualization, emerging storage devices and techniques
- Performance benchmarking, resource management, and workload studies from production systems including both traditional HPC and data-intensive workloads.
- Programmability, APIs, and fault tolerance of storage systems
- Parallel file systems, metadata management, and complex data management, object and key-value storage, and other emerging data storage/retrieval techniques
- Programming models and frameworks for data intensive computing including extensions to traditional and nontraditional programming models, asynchronous multi-task programming models, or to data intensive programming models
- Techniques for data integrity, availability and reliability especially
- Productivity tools for data intensive computing, data mining and knowledge discovery
- Application or optimization of emerging “big data” frameworks towards scientific computing and analysis
- Techniques and architectures to enable cloud and container-based models for scientific computing and analysis
- Techniques for integrating compute into a complex memory and storage hierarchy facilitating in situ and in transit data processing
- Data filtering/compressing/reduction techniques that maintain sufficient scientific validity for large scale compute-intensive workloads

- Tools and techniques for managing data movement among compute and data intensive components both solely within the computational infrastructure as well as incorporating the memory/storage hierarchy

Paper Submissions: <http://www.pdsw-discs.org/>

Paper (in pdf format) due Wednesday, Sept. 7, 2016, 11:59PM AoE

Notification: Friday, Sept. 30, 2016

Camera ready due: Friday, Oct. 7, 2016

Slides due before workshop: Sunday, Nov. 13, 2016, 5:00 pm PDT

Paper Submission Details:

The PDSW-DISCS Workshop holds a peer reviewed competitive process for selecting short papers. Submit a not previously published short paper of up to 5 pages, not less than 10 point font and not including references, in a PDF file as instructed on the workshop web site. Submitted papers will be reviewed under the supervision of the workshop program committee. Submissions should indicate authors and affiliations. Papers must not be longer than 5 pages (excluding references). Selected papers and associated talk slides will be made available on the workshop web site; the papers will also be published in the digital libraries of the IEEE and ACM.

Work-in-progress (WIP) Submissions: <http://www.pdsw-discs.org/>

There will also be a WIP session at the workshop, where presenters give 5-minute brief talks on their on-going work, with fresh problems/solutions, but may not be mature or complete yet for paper submission. A 1-page abstract is required as instructed on the workshop web site.

WIP Submission Deadline: Tuesday, Nov. 1, 2016

WIP Notification: Monday, Nov. 7, 2016