
The 10th International Parallel Data Systems Workshop (PDSW'25) - updated 7/18/25

PDSW 2025 website: <https://www.pdsw.org/index.shtml>

Paper Submissions due: Aug 1st, 2025, 11:59 PM AoE

AD due: Aug 8th, 2025, 11:59 PM AoE

Paper Notification: Sep 5th, 2025, 11:59 PM AoE

Camera ready due: Sep 27th, 2025, 11:59 PM AoE

Final AD/AE due: Oct 15, 2025, 11:59 PM AoE

Submissions website: <https://submissions.supercomputing.org/>

We are excited to announce the 10th International Parallel Data Systems Workshop (PDSW'25), to be held in conjunction with SC25: The International Conference for High Performance Computing, Networking, Storage, and Analysis, in St. Louis, MO. PDSW'25 builds upon the rich legacy of its predecessor workshops, the Petascale Data Storage Workshop (PDSW, 2006–2015) and the Data Intensive Scalable Computing Systems (DISCS, 2012–2015) workshop.

The increasing importance of efficient data storage and management continues to drive scientific productivity across traditional simulation-based HPC environments and emerging Cloud, AI/ML, and Big Data analysis frameworks. Challenges are compounded by the rapidly expanding volumes of experimental and observational data, the growing disparity between computational and storage hardware performance, and the rise of novel data-driven algorithms in machine learning. This workshop aims to advance research and development by addressing the most pressing challenges in large-scale data storage and processing.

We invite the community to contribute original research manuscripts that introduce and evaluate novel algorithms or architectures, share significant scientific case studies or workloads, or assess the reproducibility of previously published work. We emphasize the importance of community collaboration for problem identification, workload capture, solution interoperability, standardization, and shared tools. Authors are encouraged to provide comprehensive experimental environment details (software versions, benchmark configurations, etc.) to promote transparency and facilitate collaborative progress.

Topics of Interest:

- Scalable Architectures: Distributed data storage, archival, and virtualization.
- New Data Processing Models and Algorithms: Application of innovative data processing models and algorithms for parallel computing and analysis.
- Performance Analysis: Benchmarking, resource management, and workload studies.

- Cloud and Container-Based Models: Enabling cloud and container-based frameworks for large-scale data analysis.
- Storage Technologies: Adaptation to emerging hardware and computing models.
- Data Integrity: Techniques to ensure data integrity, availability, reliability, and fault tolerance.
- Programming Models and Frameworks: Big data solutions for data-intensive computing.
- Hybrid Cloud Data Processing: Integration of hybrid cloud and on-premise data processing.
- Cloud-Specific Opportunities: Data storage and transit opportunities specific to cloud computing.
- Storage System Programmability: Enhancing programmability in storage systems.
- Data Reduction Techniques: Filtering, compression, and reduction techniques for large-scale data.
- File and Metadata Management: Parallel file systems, metadata management at scale.
- In-Situ and In-Transit Processing: Integrating computation into the memory and storage hierarchy for in-situ and in-transit data processing.
- Alternative Storage Models: Object stores, key-value stores, and other data storage models.
- Productivity Tools: Tools for data-intensive computing, data mining, and knowledge discovery.
- Data Movement: Managing data movement between compute and data-intensive components.
- Cross-Cloud Data Management: Efficient data management across different cloud environments.
- AI-enhanced Systems: Storage system optimization and data analytics using machine learning.
- New Memory and Storage Systems: Innovative techniques and performance evaluation for new memory and storage systems.

More details are available at: <https://www.pdsw.org/index.shtml>

Template and Submission

- A full paper up to 6 pages in length, excluding references and AD/AE appendices.
- Artifact Description (AD) Appendix is mandatory and Artifact Evaluation (AE) Appendix is optional.
- Submissions with AD and AE Appendix will be considered favorably for the PDSW Best Paper award.
- Papers must adhere to the ACM conference paper template available at: <https://www.acm.org/publications/proceedings-template>
- **UPDATE - Papers will be reviewed double-blind. Author names and affiliations should NOT be included in the submitted paper.**
- Submit your papers by Aug 1st, 2025, 11:59 PM AoE at <https://submissions.supercomputing.org/>

Reproducibility Initiative

Aligned with the SC25 Reproducibility Initiative (<https://sc25.supercomputing.org/program/papers/reproducibility-initiative/>), we encourage detailed and structured artifact descriptions (AD) using the SC25 format (<https://github.com/weidendo/sc25-repro>). The AD should include a field for one or more links to data (Zenodo, figshare, etc.) and code (Github, GitLab, Bitbucket, etc.) repositories. For the artifacts that will be placed in the code repository, we encourage authors to follow the PDSW 2025 Reproducibility Addendum (<https://www.pdsw.org/pdsw25/PDSW2025ReproducibilityInitiativeAddendum.pdf>) on how to structure the artifact, as it will make it easier for the reviewing committee and readers of the paper in the future.

Submissions website: <https://submissions.supercomputing.org/>

Organization team:

General Chair:

Suren Byna

The Ohio State University, USA

Program Co-Chairs:

Anthony Kougkas

Illinois Institute of Technology, USA

Sarah Neuwirth

Johannes Gutenberg University Mainz, Germany

Reproducibility Chair:

Ricardo Macedo

INESC TEC & University of Minho, Portugal

Publicity Co-Chairs:

Izzet Yildirim

Illinois Institute of Technology, USA

Web & Publications Chair:

Joan Digney

Carnegie Mellon University, USA