



Minutes from the 9th edition of the Performance and Scalability of Storage Systems workshop (Per3S), 23rd May 2025, "Maison des Mines et des Ponts", Paris

Jean-Thomas Acquaviva, Jalil Boukhobza, Philippe Deniel, Shadi Ibrahim, Philippe Raipin-Parvédy, François Trahay

► To cite this version:

Jean-Thomas Acquaviva, Jalil Boukhobza, Philippe Deniel, Shadi Ibrahim, Philippe Raipin-Parvédy, et al.. Minutes from the 9th edition of the Performance and Scalability of Storage Systems workshop (Per3S), 23rd May 2025, "Maison des Mines et des Ponts", Paris. 9th edition of the workshop Performance and Scalability of Storage Systems (Per3S), May 2025, Paris, France. pp.1-2. <hal-05116147>

HAL Id: hal-05116147

<https://hal.science/hal-05116147v1>

Submitted on 18 Jun 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License

Minutes from the 9th edition of the Performance and Scalability of Storage Systems workshop (Per3S), 23rd May 2025, "Maison des Mines et des Ponts", Paris.

Jean-Thomas Acquaviva, DDN

Jalil Boukhobza, Lab-STICC, ENSTA, Institut Polytechnique de Paris

Philippe Deniel, CEA

Shadi Ibrahim, Inria

Philippe Raipin, Orange

François Trahay, Telecom Sud Paris, Institut Polytechnique de Paris

Per3S is a workshop aiming to bring together the scientific and technological storage community to discuss and address issues and challenges associated to performance and data operations at scale. These topics cover HPC storage as well as Cloud-oriented architectures, both sharing the need for extreme scale.

The 9th edition of Per3S was built around 5 sessions of 18 presentations from young and confirmed researchers and engineers, from both academia and industry, highlighting the work of 7 labs, 4 companies, and 3 large infrastructure institutions:

- Cloud and Storage featuring 6 talks on broad topics related to KV stores, tapes integration, storage for AI, persistent memory, and caching and data placement in ICN and p2p systems
- Young Researchers Distinguished talks featuring 2 talks on I/O modeling and tracing
- Flash presentations to introduce posters featuring 4 presentations on I/O performance and diagnosis, caching and LLM optimization.
- HPC Storage and Lustre featuring 6 talks on software file systems and frameworks
- A panel discussion about hot topics in storage research.

This year's edition was rich in discussions about hot topics in storage research; here is a small summary.

New storage technologies disrupt the application layer: the integration of persistent and fast storage in the memory hierarchy disrupts the traditional way of designing applications and accessing the data. This was highlighted by the talk of Baptiste Lepers on designing a fast storage layer for persistent KV stores. New runtimes are also necessary to transparently handle data persistence as presented by Jana Toljaga in her talk on Volipmem, a system-level runtime for persistent memory.

Tapes are back (but did they ever leave?): several talks focused on how to get the tapes back in the storage hierarchy, in a more transparent way and how to manage them within existing software frameworks. This was illustrated by the talk of Valentin Honoré entitled “bridging the gap between cold and hot storage at scale” and the talk of Gauthier Evraerd on Phobos, a parallel heterogeneous object store.

AI strongly impacts the storage hierarchy: AI workloads are increasingly present in numerous domains such as HPC, cloud and edge, and it is critical to model the performance and optimize the storage stack for such applications as highlighted by the talk of Jean-Thomas Acquaviva entitled “From on prem to Cloud: an AI story” and the talk of Meriem Bouzouad entitled “Adaptive layer compression and storage with QoS-aware loading for LLM serving”.

There is still room for interesting research in data distribution and caching: to satisfy QoS requirements from the user perspectives and address the ever-changing and growing workloads, caching and data placement is still the way to go as illustrated by the presentations of Lydia Ait

Oucheggou, Alexandru Dobrila and Hocine Mahni, respectively entitled: “Bridging local efficiency and global cost: two complementary ICN caching strategies”, “Scaling a decentralization p2p storage system” and “Fabme: file-level placement policy with admission control, burst prefetching, and multi-criteria eviction for hpc multi-tier storage”.

Storage is getting more complex; tracing is still a challenge to get the most of the storage system and avoid performance glitches and interference. Two presentations focused on the tracing part, or how to get meaningful traces at a low cost such as the presentation of Louis-Marie Nicolas entitled “Modeling, advising functions and DVFS for the energy optimization OF HPC I/Os” and the one of Catherine Guelque on Pallas, a generic trace format for large HPC trace analysis. However, tracing is only the first step, analyzing and understanding the performance behavior of I/Os is also a challenge as suggested by the talk of Luan Teylo entitled “A deep look into the temporal I/O behavior of HPC applications”. Another important challenge is avoiding interferences as highlighted by Méline Trochon. In her talk on tracing-based I/O diagnosis for performance optimization of software-defined storage systems, Lina Sadi seeks to close the full loop from tracing to corrective actions in object stores, and more particularly in Ceph deployment on the edge.

Data movement is not only about device performance but also about software framework, data lifecycle management, and application-specific needs. Several talks emphasized the need for generic frameworks like Mochi, which retrospective was presented by Mathieu Dorier, or specific support, as the ephemeral storage services highlighted by Philippe Couvée. State-of-the-art systems such as Lustre are also very active in adjusting to data workflow needs, as highlighted by the talk of Sébastien Buisson. The return of experience from the Juelich Supercomputing Centre, given by Salem El Sayed, also emphasized the complex process and the need for robust software frameworks for handling petabytes of data.

Community session: Evolution of technologies and interests. As part of the workshop, an interactive session was organized to get feedback and comments from the audience. To capture and measure the different opinions expressed, we used polling software with 10 questions. The following paragraph summarizes the view of a total of 26 participants answering these 10 questions.

During the session, it appeared that, for the audience, hard drives are not considered an appealing topic of research. The interest has heavily shifted toward software-defined storage and flash storage. Tape received more votes (2 out of 26) than hard drives.

Still about the evolution of the technical landscape, two-thirds of the audience consider that Cloud storage and AI storage will remain strongly aligned, and similarly, HCP and AI storage will not diverge in the short-term future. Regarding the next technologies to look at, ZNS SSD (*Zoned NameSpace* SSD) and CXL were clear leaders. The second topic discussed during the live session was performance. A takeaway from this session could be **performance evaluation: consensus on the method but a fragmented landscape**. The audience polled about the best way to evaluate infrastructure performance. If the well-known benchmarks such as IOR and FIO won, it was by a small margin, as it seems that no solution is really satisfying except running applications against the infrastructure.

However, when asked about the best way to analyze the performance of end-user applications, traces and a Darshan-like approach made consensus. The idea of providing a simple I/O and power rating at the end of the application is also gaining widespread agreement. The question about how to implement it has not been discussed and probably will be a nice entry point for Per3S 2026.

