

# EOFS 2026 Workshop on Open Source Parallel Filesystems

Thursday, 12 March 2026 - Friday, 13 March 2026

Maison des Mines et des Ponts et Chaussées



## Book of Abstracts



# Contents

DAOS Update . . . . .	1
Lustre status and current developments . . . . .	1
RobinHood: storing and querying a filesystem's metadata for fast access . . . . .	1
Phobos: A Flexible, Open-Source Tape Storage System for HPC and Beyond . . . . .	1
BeeGFS Auto-Tiering for Data-Heavy Environments . . . . .	2



**Session II: Status and Development Trends / 1****DAOS Update****Author:** Michael Hennecke<sup>1</sup><sup>1</sup> *HPE*

This presentation will provide an update on activities in the DAOS Foundation including the upcoming DAOS 2.8 release, first experiences of deploying DAOS on 400Gbps fabrics (IB-NDR, Slingshot 400, OPA 400, RoCE), a preview of the HPE Cray Supercomputing Storage K3000 product based on DAOS, and other notable research and development topics from the DAOS community.

**Session II: Status and Development Trends / 2****Lustre status and current developments****Author:** Marc Vef<sup>1</sup><sup>1</sup> *DDN/Whamcloud*

This talk will present the current status of Lustre development, upcoming features, and roadmap. This will include topics, such as:

- the Lustre nodemap feature that was significantly extended recently;
- the Erasure Coding effort status and next steps with Immediate Write Mirroring;
- the Lustre Trashcan/undelete feature; and
- the Lustre quota aggregation feature.

**Session II: Status and Development Trends / 4****RobinHood: storing and querying a filesystem's metadata for fast access****Author:** Yoann Valeri<sup>1</sup><sup>1</sup> *CEA*

As supercomputers are becoming faster and faster, so does their data output. Since the regularly accessed data must be stored and available quickly to users, it is important to put it on fast storage systems. However, these tend to have a low capacity, meaning we must be able to choose the data which should remain on those types of storage systems, and which can be placed on slower but more capacitive systems. As such, it is important to be able to accurately know the state of a filesystem at any point, but using the conventional means provided by the operating system for this, for instance to do filesystem traversals, can be time consuming if done regularly. Moreover, these operations impose a heavy load on the filesystem, making it slower. To counter these problems, we created a suite of tools called RobinHood that aims to mirror a filesystem in a database, and use the latter to define policies that will manage data placement according to their usage.

**Session II: Status and Development Trends / 5**

## **Phobos: A Flexible, Open-Source Tape Storage System for HPC and Beyond**

**Author:** Thomas LEIBOVICI<sup>1</sup>

<sup>1</sup> CEA

In the era of exascale computing and data-intensive workflows, efficient tiered storage architectures are essential to balance performance, capacity, and cost. While parallel file systems like Lustre, BeeGFS, and DAOS excel at handling high-throughput I/O, the seamless integration of high-capacity, long-term storage solutions such as tape libraries remains a major challenge for long-term data retention and cost-effective archiving.

This talk presents Phobos, an open-source storage system developed by CEA, specifically designed to address these challenges by providing a highly efficient, scalable, and vendor-neutral solution for managing tape-based archives and large robotic libraries. Built on the Linear Tape File System (LTFS) —an open, standardized format—Phobos ensures interoperability, long-term data preservation, and independence from proprietary formats and software, making it a cornerstone for data sovereignty. Phobos offers multiple front-ends, including Lustre/HSM, S3, and iRODS, enabling seamless integration in diverse HPC, cloud, and data management environments.

**Session II: Status and Development Trends / 6**

## **BeeGFS Auto-Tiering for Data-Heavy Environments**

**Author:** Philipp Falk<sup>None</sup>

Data-heavy workloads demand flexible and efficient storage strategies. In this session, the BeeGFS VP of Engineering will present BeeGFS auto-tiering capabilities, demonstrating how data can be dynamically placed across different storage tiers, including internal BeeGFS systems and external S3-based storage to balance performance, scalability, and cost.