

# **EOS 2026 Workshop**

## **Report of Contributions**

Contribution ID: 1

Type: **not specified**

## Unlocking SMR Capacity in EOS Using Zoned XFS

*Tuesday, 10 March 2026 11:00 (20 minutes)*

Recent developments in XFS have introduced native support for zoned storage devices such as Shingled Magnetic Recording (SMR) drives, offering a robust and higher-performance alternative to BTRFS for this class of hardware. With Zoned XFS EOS can now efficiently leverage the additional 20–25% storage capacity provided by SMR HDDs.

This presentation will give an overview of the architectural changes made to XFS to support zoned storage, demonstrate its performance characteristics, and provide a practical, hands-on guide for setting up, using and tuning XFS file systems on SMR HDDs.

**Author:** HOLMBERG, Hans (Western Digital Research)

**Co-author:** Mr HELLWIG, Christoph (Western Digital Research)

**Presenter:** HOLMBERG, Hans (Western Digital Research)

**Session Classification:** Development & Storage Hardware

Contribution ID: 2

Type: **not specified**

## ORCHIDE: Using EOS for space edge computing

*Wednesday, 11 March 2026 14:00 (20 minutes)*

Modern Low Earth Orbit satellites increasingly support on-board data processing using GPUs and FPGAs, enabling multi-purpose use cases beyond simple data capture. However, orchestrating computational workloads across single or multi-node satellite systems requires an efficient shared storage solutions that can handle the constraints of space environments.

This contribution presents the integration of EOS into the ORCHIDE project - an orchestrator platform for managing multi-tenant satellite workloads. ORCHIDE enables service operators to submit processing jobs that execute on sensor data stored locally, eliminating the need to download raw data to ground stations. The platform's requirement for platform-agnostic deployment across heterogeneous satellite architectures necessitated a shared filesystem with low overhead, fault tolerance, and long-term reliability.

We document the comparative analysis that led to selecting EOS over alternative solutions, present performance benchmarks demonstrating different workloads using shared filesystems and describe the integration enabling ORCHIDE to manage storage across distributed satellite compute nodes.

**Author:** WEISZ, Sergiu (National University of Science and Technology POLITEHNICA Bucharest (RO))

**Presenter:** WEISZ, Sergiu (National University of Science and Technology POLITEHNICA Bucharest (RO))

**Session Classification:** Ecosystem & Roadmap

Contribution ID: 3

Type: **not specified**

## EOS site report of the Joint Research Centre

*Wednesday, 11 March 2026 10:50 (20 minutes)*

The Joint Research Centre (JRC) of the European Commission is running the Big Data Analytics Platform (BDAP) to enable the JRC projects and scientists to store, process, and analyze a wide range and large amount of data, and to share and disseminate data products.

EOS is the main system of BDAP for storing the scientific data. The BDAP services are actively used by more than 100 JRC projects, covering a wide range of data analytics activities. The EOS instance at JRC has currently a gross capacity of 40 PB.

The talk will present the EOS service at JRC as storage back-end of the Big Data Analytics Platform. The presentation covers the EOS setup, configuration, and current status. It gives an overview about the activities over the last year, presents experiences made and issues discovered, and gives an outlook of planned activities during 2026.

**Author:** BURGER, Armin

**Co-author:** Mr EYRAUD, Franck

**Presenter:** BURGER, Armin

**Session Classification:** Site Reports

Contribution ID: 4

Type: **not specified**

## EOS Deployments through Ansible Automation

*Tuesday, 10 March 2026 14:20 (15 minutes)*

This contribution outlines the use of Ansible to automate the deployment and lifecycle management of EOS clusters. By utilizing declarative playbooks, we orchestrate the setup of MGM and FST components while ensuring consistent configuration of the QuarkDB metadata backend and XRootD framework. This automated approach eliminates manual errors, simplifies complex version upgrades, and enables the rapid scaling of storage capacity required for exa-scale environments. Attendees will see how shifting to Infrastructure-as-Code provides a repeatable and reliable foundation for managing massive-scale open storage.

**Author:** VALA, Martin (Pavol Jozef Safarik University (SK))

**Presenter:** VALA, Martin (Pavol Jozef Safarik University (SK))

**Session Classification:** Operational Tools & Configuration

Contribution ID: 5

Type: **not specified**

## Optimizing High-Dimensional Data Storage: NDMSPC with THnSparse and TTree on EOS

*Wednesday, 11 March 2026 14:20 (20 minutes)*

This contribution presents the NDMSPC (N-Dimensional Space) framework, designed for efficient management and analysis of high-dimensional datasets within the CERN EOS environment. We explore the integration of ROOT's THnSparse for memory-efficient multi-dimensional histogramming alongside TTree for robust data storage. By leveraging EOS as the underlying storage layer, the framework achieves the high-throughput I/O necessary for processing exa-scale sparse data. The discussion focuses on optimization techniques for data layout and access patterns that maximize performance when handling complex N-dimensional physics structures in a distributed storage architecture.

**Author:** VALA, Martin (Pavol Jozef Safarik University (SK))

**Presenter:** VALA, Martin (Pavol Jozef Safarik University (SK))

**Session Classification:** Ecosystem & Roadmap

Contribution ID: 6

Type: **not specified**

## Apollon & Hermes: Next step in EOS disk operations

*Tuesday, 10 March 2026 14:00 (20 minutes)*

This presentation introduces Apollon and Hermes, two complementary tools that automate the complete disk lifecycle. Apollon is a centralized gRPC service that manages automated draining, intelligent repair of failed drains through dynamically-loaded plugin strategies, and systematic leftover file handling. Hermes runs on each storage node, interfacing with repair catalogue systems to detect disk failures, coordinate draining with Apollon, and automatically register replacement disks.

Together, these tools transform disk operations from a multi-step manual process into an automated workflow requiring intervention only for physical disk replacement, significantly reducing operational fragmentation while enhancing reliability.

**Authors:** DEL MONTE, Gianmaria (CERN); MATEI, Octavian-Mihai (CERN)

**Presenter:** MATEI, Octavian-Mihai (CERN)

**Session Classification:** Operational Tools & Configuration

Contribution ID: 7

Type: **not specified**

## EOS-Alarms - Status Automatic Monitoring

*Tuesday, 10 March 2026 14:50 (20 minutes)*

EOS-Alarms is managed by a multitude of subsystems. This talk will focus on Winston, which is an alarm management system that automates EOS health monitoring by continuously analyzing core subsystems through parsing JSON output from EOS commands.

Winston understands EOS-specific operational patterns, distinguishing spare scheduling groups from active ones, correlating drain status with filesystem health, tracking quota-space view consistency, and monitoring node uptime patterns. Instance-specific configurable thresholds adapt to different deployment patterns across production, development, and archive storage. Operators manage alarms through Mattermost slash commands for diagnostics, suppression during maintenance windows with reason tracking, and status queries. Dynamic alarm suppression prevents alert fatigue. ServiceNow integration automatically creates tickets with EOS-specific context (FSID, hostport, scheduling group), updates them as conditions evolve, and closes them when metrics normalize.

**Author:** MATEI, Octavian-Mihai (CERN)

**Presenter:** MATEI, Octavian-Mihai (CERN)

**Session Classification:** Operational Tools & Configuration

Contribution ID: 8

Type: **not specified**

## CERN Disk Benchmark Tool

*Tuesday, 10 March 2026 14:35 (15 minutes)*

The CERN Disk Benchmark Tool benchmarks Linux block devices and filesystems by measuring sustained large-block write speed and in-place file rewrite performance at a chosen mount point. It supports parallel I/O and captures device statistics with iostat. The tool generates throughput-versus-usage plots and produces a self-contained PDF report, including configuration metadata for reproducible analysis.

**Author:** PETERS, Andreas Joachim (CERN)

**Presenter:** PETERS, Andreas Joachim (CERN)

**Session Classification:** Operational Tools & Configuration

Contribution ID: 9

Type: **not specified**

## NFS4 R&D for EOS

*Monday, 9 March 2026 14:50 (20 minutes)*

CERN-NFS v4.0 is a prototype NFS server implemented in modern C++ that currently runs as a standalone service on a local filesystem backend. The server is built around a handle-first VFS abstraction, enabling filesystem backends to be cleanly plugged in without impacting core NFS protocol logic. While the current implementation targets POSIX and in-memory filesystems, the architecture is designed to integrate EOS as a backend. To fully support the EOS access model, the protocol must be extended from NFSv4.0 to NFSv4.1, enabling file layouts with flex files to map MGM-to-FST redirection and avoid routing data traffic through gateway nodes. In addition, the design anticipates support for NFSv4 security mechanisms such as Kerberos and certificate-based authentication, both natively supported by Linux kernel clients, as a further step toward full NFSv4 integration with EOS.

**Author:** PETERS, Andreas Joachim (CERN)

**Presenter:** PETERS, Andreas Joachim (CERN)

**Session Classification:** Development

Contribution ID: **10**Type: **not specified**

## EOS File Notifications

*Monday, 9 March 2026 16:10 (15 minutes)*

EOS File Notification adds a customizable way to the EOS Workflow Engine to send file change notifications to other systems. It detects changes in watched folders and sends out alerts using different methods, like HTTP(S), gRPC, ActiveMQ, or Redis. Each alert uses a consistent JSON message that contains key file info, such as metadata, folder details, checksums, times, and user access info. This setup helps other services respond to file changes, track activity, or automate tasks, making it easier to share EOS file event updates outside the storage system.

**Author:** PETERS, Andreas Joachim (CERN)

**Presenter:** PETERS, Andreas Joachim (CERN)

**Session Classification:** Development

Contribution ID: 11

Type: **not specified**

## **EOS 'rclone'**

*Tuesday, 10 March 2026 09:40 (20 minutes)*

The EOS rclone tool is an EOS-integrated file replication and synchronization utility that provides one-way copy and bi-directional sync operations between directories within EOS or across EOS-accessible endpoints. It supports dry-run execution, selective updates, verbosity control, and optional inclusion of EOS-specific semantics such as atomic files, versioned files, and hidden entries. Synchronization decisions are based on file modification times, enabling efficient incremental transfers. The tool is designed for controlled data propagation, validation, and migration workflows within EOS environments while offering predictable, script-friendly behavior.

**Author:** PETERS, Andreas Joachim (CERN)

**Presenter:** PETERS, Andreas Joachim (CERN)

**Session Classification:** Development

Contribution ID: 12

Type: **not specified**

## The EOS Development Workplan & Roadmap

*Wednesday, 11 March 2026 14:40 (20 minutes)*

This presentation outlines the EOS development plan and roadmap toward EOS 6, detailing key architectural milestones, planned features, and strategic objectives that will shape the next major evolution of the system. It provides an overview of ongoing engineering efforts, highlights upcoming enhancements across scalability, performance, and interoperability, and discusses longer-term directions guiding EOS development. The roadmap is defined to align EOS with the HL-LHC's required use cases and the evolving hardware landscape.

**Authors:** PETERS, Andreas Joachim (CERN); SINDRILARU, Elvin Alin (CERN)

**Presenters:** PETERS, Andreas Joachim (CERN); SINDRILARU, Elvin Alin (CERN)

**Session Classification:** Ecosystem & Roadmap

Contribution ID: 13

Type: **not specified**

## News in HTTP's server-side support

*Monday, 9 March 2026 15:50 (20 minutes)*

- WLCG data integrity check
- Enable web-browser access to an EOS instance

**Author:** CAFFY, Cedric (CERN)

**Presenter:** CAFFY, Cedric (CERN)

**Session Classification:** Development

Contribution ID: 14

Type: **not specified**

## One File, Many Hashes: The New Alternative Checksums Feature

*Tuesday, 10 March 2026 09:00 (20 minutes)*

This presentation introduces Alternative Checksums in EOS 5.4.0, a feature that enables the storage of multiple hash types for a single file alongside the default system checksum. By allowing for various algorithms, such as MD5 or SHA-256, on a per-directory basis, EOS provides the necessary flexibility to meet diverse experiment requirements.

**Author:** DEL MONTE, Gianmaria (CERN)

**Presenter:** DEL MONTE, Gianmaria (CERN)

**Session Classification:** Development

Contribution ID: 15

Type: **not specified**

## A High-Performance S3 Gateway for EOS

*Monday, 9 March 2026 15:10 (20 minutes)*

This presentation introduces a new S3 interface for EOS developed as a lightweight plugin for VersityGW. This “thin layer” provides an efficient bridge between the S3 protocol and the EOS backend by natively translating S3 requests into gRPC and HTTP calls.

**Author:** DEL MONTE, Gianmaria (CERN)

**Presenter:** DEL MONTE, Gianmaria (CERN)

**Session Classification:** Development

Contribution ID: 16

Type: **not specified**

## EOS Implementation Phase at INPE

*Wednesday, 11 March 2026 09:00 (20 minutes)*

The Intelligent Early Warning System for Climate Extremes (SIPEC/SisMOM), led by the National Institute for Space Research (INPE, Brazil), integrates satellite and in-situ sensor data to support early prediction of climate extremes using numerical models and machine learning techniques. To enable distributed access to large and heterogeneous datasets produced by multiple institutions, INPE is implementing a data federation based on the EOS ecosystem.

Over the last six months, the project has entered an implementation and validation phase, supported by a joint team of developers and operators from CERN Storage Group and two professionals from INPE. This work included the deployment of the software and network infrastructure needed by EOS core services, and configuration of authentication and authorization mechanisms enabling fine-grained access control across federated sites.

These results will demonstrate the feasibility of a secure and scalable EOS-based data federation for climate research and highlight the value of close collaboration with CERN Storage Group in data-intensive scientific environments to become operational in the current year.

**Author:** Mr OLIVEIRA MENDES, Wanderley (Instituto Nacional De Pesquisas Espaciais - INPE (BR))

**Co-author:** Dr LUIS GOMES, Jorge (Instituto Nacional De Pesquisas Espaciais - INPE (BR))

**Presenters:** Dr LUIS GOMES, Jorge (Instituto Nacional De Pesquisas Espaciais - INPE (BR)); Mr OLIVEIRA MENDES, Wanderley (Instituto Nacional De Pesquisas Espaciais - INPE (BR))

**Session Classification:** EOS based Site & Software Evolution

Contribution ID: 17

Type: **not specified**

## EOS recycle bin updates

*Monday, 9 March 2026 16:25 (15 minutes)*

This presentation will focus on the recent improvements done to the recycle bin implementation. This will cover the internal changes with respect to recycle bin clean-up, the addition of recycle project concept and the use of the GRPC interface.

**Author:** SINDRILARU, Elvin Alin (CERN)

**Presenter:** SINDRILARU, Elvin Alin (CERN)

**Session Classification:** Development

Contribution ID: **18**

Type: **not specified**

## **EOS development updates 5.3/5.4**

*Monday, 9 March 2026 14:10 (20 minutes)*

This presentation will go through the main developments included in the EOS code base in the past year.

**Author:** SINDRILARU, Elvin Alin (CERN)

**Presenter:** SINDRILARU, Elvin Alin (CERN)

**Session Classification:** Development

Contribution ID: **19**

Type: **not specified**

## **XRootD 6.0**

*Monday, 9 March 2026 14:30 (20 minutes)*

In this talk we will discuss the new features in XRootD 6.0.

**Author:** AMADIO, Guilherme (CERN)

**Presenter:** AMADIO, Guilherme (CERN)

**Session Classification:** Development

Contribution ID: 20

Type: **not specified**

## EOS Quotas by Logical Space

*Monday, 9 March 2026 16:40 (10 minutes)*

In this contribution, we will discuss recent developments in how EOS quotas are tracked, some features which are still missing, and our future plans for further improvements.

**Author:** AMADIO, Guilherme (CERN)

**Presenter:** AMADIO, Guilherme (CERN)

**Session Classification:** Development

Contribution ID: 21

Type: **not specified**

## EOS operations at CERN - 2025 status report

*Wednesday, 11 March 2026 11:30 (20 minutes)*

This work provides an overview of EOS operations at CERN, highlighting its critical role in supporting large-scale physics data processing and storage. EOS is a high-performance distributed storage system engineered to manage the massive volumes of scientific data produced by CERN experiments. This presentation highlight recent operational achievements, and strategic objectives for the last and current year, with a particular focus on advances in efficiency, reliability, and scalability. It also examines the impact of EOS on physics workflows, emphasising its contribution to seamless data access and analysis.

**Author:** MASCETTI, Luca (CERN)

**Presenter:** MASCETTI, Luca (CERN)

**Session Classification:** Site Reports

Contribution ID: 22

Type: **not specified**

## Some recent XRootD changes for EOS

*Tuesday, 10 March 2026 09:20 (20 minutes)*

A selection of changes in XRootD with impact on EOS is presented. This covers bug fixes, associated refactoring and some new features. In particular the file-clone feature in XRootD release 6 and ongoing development in EOS that makes use of the feature is discussed.

**Author:** SMITH, David (CERN)**Presenter:** SMITH, David (CERN)**Session Classification:** Development

Contribution ID: 23

Type: **not specified**

## EOS at the Fermilab LHC Physics Center

*Wednesday, 11 March 2026 11:10 (20 minutes)*

The Fermilab CMS LPC center has operated an EOS instance since initial testing began in June 2012, transitioning to production storage in 2013 with an initial capacity of 600 TB. Today, the system provides approximately 16 PB of storage to support the 4,500-core LPC user analysis cluster, which serves several hundred active CMS users at any given time.

In this talk, we will provide an update on our system setup and outline near-term plans for the CMS LPC production EOS system. We will share operational experiences from running our EOS system over the past year. Additionally, we will briefly discuss our recent efforts within the Fermilab Disk Evolution Project to consider the future architecture of HEP data storage at Fermilab for the HL-LHC era.

**Author:** WU, Yujun (Fermi National Accelerator Lab. (US))

**Presenter:** WU, Yujun (Fermi National Accelerator Lab. (US))

**Session Classification:** Site Reports

Contribution ID: 24

Type: **not specified**

## EOS ChatBot - Pilot Evaluation & Campaign Proposal

*Tuesday, 10 March 2026 15:10 (20 minutes)*

A brief, transparent update on EOS ChatBot, an AI assistant prototype integrated into Mattermost and the EOS Community forum. This session covers what EOS ChatBot is designed for, why a full evaluation is premature, what we are measuring, and a concrete plan leading to a data-driven decision in May. Includes a practical walkthrough of how to use it today.

**Author:** Dr ARSUAGA RIOS, Maria (CERN)

**Presenter:** Dr ARSUAGA RIOS, Maria (CERN)

**Session Classification:** Operational Tools & Configuration

Contribution ID: 25

Type: **not specified**

## Workshop Introduction

*Monday, 9 March 2026 14:00 (10 minutes)*

### **Abstract —EOS Workshop Introduction**

We open the workshop with a short practical overview: how the days are structured, the timetable, session flow, technical details, and where to find coffee breaks and the workshop dinner. We also briefly introduce the people in the room before starting, to make it easy for participants to meet each other and have useful discussions.

**Author:** PETERS, Andreas Joachim (CERN)

**Presenter:** PETERS, Andreas Joachim (CERN)

**Session Classification:** Development

Contribution ID: 26

Type: **not specified**

## I/O Shaping in EOS

*Tuesday, 10 March 2026 10:00 (20 minutes)*

**I/O shaping** in EOS is a cluster-wide mechanism that helps share storage bandwidth fairly across a whole instance. It continuously collects near-real-time I/O activity from all nodes and builds a global picture of who is reading and writing - grouped by user, group, or application. Using this global view, EOS can automatically apply scheduling and priority rules that decide how the available I/O capacity is shared.

Unlike local throttling on a single machine, global I/O shaping coordinates decisions across the entire system. This prevents one heavy consumer from slowing everyone else down and keeps latency stable when different workloads run at the same time. Administrators can prioritize critical workloads, guarantee minimum throughput for selected users, or limit overly aggressive activity, all while keeping overall throughput high. The result is more predictable performance and better sharing of storage resources without requiring changes in applications.

We will present a report on its current status and the technical aspects of its implementation.

**Presenter:** OBIS APARICIO, Luis Antonio (CERN)

**Session Classification:** Development

Contribution ID: 27

Type: **not specified**

## EOSHPM Evolution

*Wednesday, 11 March 2026 09:20 (20 minutes)*

CERNBox, CERN's cloud collaboration platform, currently serves more than 27,000 users worldwide and manages over 4.1 billion files across multiple petabytes of data. Behind this service, EOS HPM (EOS Home-Project-Media) provides a large-scale, multipetabyte storage infrastructure that enables reliable access to both personal and project spaces.

This presentation reviews the current infrastructure, available resources, and the recent evolution of EOS HPM in production. We discuss the progressive upgrade path from EOS 5.2.x to 5.4.0, the transition to an MQ-less (Pub/Sub) architecture, and a major refurbishment campaign involving disk server decommissioning in MDC and new node deployments in PDC.

Operational improvements are also presented, including the redesign of the redirector architecture, the deployment of EOSBACKUP, the evolution of the quota model, monitoring enhancements, and improved observability.

Finally, we review operational incidents such as OOM events, SSD failures, and redirector outages, as well as the lessons learned from them.

**Authors:** DEL MONTE, Gianmaria (CERN); IVEN, Jan (CERN); MEDINA RAMOS, Pablo

**Presenter:** MEDINA RAMOS, Pablo

**Session Classification:** EOS based Site & Software Evolution

Contribution ID: 28

Type: **not specified**

## CERN Tape Archive Status and Roadmap

*Wednesday, 11 March 2026 09:40 (20 minutes)*

The CERN Tape Archive (CTA) manages the archival and retrieval of more than one Exabyte of physics data produced by the many experiments at CERN. CTA has proved adequate to the operational demands of LHC Run-3. Preparations have already started for Run-4 and the High-Luminosity LHC era, with significant performance, scalability, and operational improvements foreseen. This contribution reports on the current production status of CTA and outlines the roadmap for LS3.

**Author:** DAVIS, Michael (CERN)

**Presenter:** DAVIS, Michael (CERN)

**Session Classification:** EOS based Site & Software Evolution

Contribution ID: 29

Type: **not specified**

## **EOS in ALICE T2 deployments at ORNL and LBNL**

*Wednesday, 11 March 2026 11:50 (20 minutes)*

This talk will present the current EOS storage architecture and current EOS issues for the ALICE-USA T2 sites at Oak Ridge National Laboratory and Lawrence Berkeley National Laboratory.

**Author:** MOULTON, Steve (Oak Ridge National Laboratory - (US))

**Presenter:** MOULTON, Steve (Oak Ridge National Laboratory - (US))

**Session Classification:** Site Reports

Contribution ID: **30**

Type: **not specified**

## **EOS Site report of IHEP**

*Wednesday, 11 March 2026 12:10 (20 minutes)*

I'll give a brief overview of EOS status at IHEP.

**Authors:** HOU, Siqu; CHENG, Yaosong (Institute of High Energy Physics Chinese Academy of Sciences, IHEP); Dr BI, Yujiang (Institute of High Energy Physics, Chinese Academy of Sciences)

**Presenter:** Dr BI, Yujiang (Institute of High Energy Physics, Chinese Academy of Sciences)

**Session Classification:** Site Reports

Contribution ID: 31

Type: **not specified**

## New Technologies for Environmentally Responsible Data Growth

*Tuesday, 10 March 2026 11:20 (20 minutes)*

As the global data sphere expands, the storage industry faces a dual challenge: satisfying the insatiable demand for capacity while drastically reducing the power and environmental footprint of data centers. This presentation explores the intersection of areal density leadership and next-generation “regenerative” drive technologies as the primary levers for solving these challenges and achieving optimal total cost of ownership.

**Author:** STRONG, Ed

**Co-author:** CRATON, Dave

**Presenter:** STRONG, Ed

**Session Classification:** Development & Storage Hardware

Contribution ID: 32

Type: **not specified**

## Next-Generation Exascale Flash Storage

*Tuesday, 10 March 2026 11:40 (20 minutes)*

This project aims to evaluate next-generation, high-density flash-based storage technologies through a strategic CERN openlab –Pure Storage collaboration. By combining CERN’s exascale operational expertise with Pure Storage’s high-efficiency DirectFlash platform, the initiative will assess performance, scalability, energy efficiency, cost, and reliability. The overall goal is to determine whether such technologies can sustainably and cost-effectively support future scientific data volumes at exabyte scale.

**Presenter:** CHOUDHURY, Ruhi**Session Classification:** Development & Storage Hardware