


# EOS 2024 Workshop



## Report of Contributions

Contribution ID: 1

Type: **not specified**

## **CERNBox/Reva deployment at CERN**

*Thursday, 14 March 2024 14:00 (1 hour)*

**Session Classification:** Meet the CERNBox Team

Contribution ID: 3

Type: **not specified**

## EOS Operation Status at KISTI Tier-1 for ALICE experiment

*Friday, 15 March 2024 15:25 (15 minutes)*

We have been running disk and archive (replacing tape) storage for ALICE experiment with EOS for several years. In 2023, we upgraded our EOS instances to v5 from v4. More recently we deployed EOS based on Podman (instead of Docker), which is adapted as a native container runtime in EL9 distributions. In this work, we present the current status of EOS operation at KISTI Tier-1 centre based on Podman container with systemd integration.

**Primary authors:** KIM, Jeong Heon (Korea Institute of Science & Technology Information (KR)); AHN, Sang Un (Korea Institute of Science & Technology Information (KR))

**Presenter:** AHN, Sang Un (Korea Institute of Science & Technology Information (KR))

**Session Classification:** EOS Operations

Contribution ID: 4

Type: **not specified**

## Site Report Vienna Tier-2

*Friday, 15 March 2024 15:10 (15 minutes)*

We share our operational experience of running a converged EOS instance for 3 Experiments (CMS, Alice, Belle).

- In 2022 we've extended the capacity of the EOS cluster from 9 to 15 FSTs
- In 2023 we've successfully updated to EOS 5.2.2
- The system deployment and all operational tasks are fully automated with Ansible. Most recently a gap was closed with the automation of host certificate rotation.

**Primary author:** BIRNGRUBER, Erich (Austrian Academy of Sciences (AT))

**Co-author:** SEREN, Umit (Austrian Academy of Sciences (AT))

**Presenter:** BIRNGRUBER, Erich (Austrian Academy of Sciences (AT))

**Session Classification:** EOS Operations

Contribution ID: 5

Type: **not specified**

## EOS at the Fermilab LHC Physics Center

*Friday, 15 March 2024 16:25 (15 minutes)*

Fermilab has been running an EOS instance since testing began in June 2012. By May 2013, before becoming production storage, there was 600TB allocated for EOS. Today, there is approximately 13PB of storage available in the EOS instance.

The LPC cluster is a 4500-core user analysis cluster with 13 PB of EOS storage. The LPC cluster supports several hundred active CMS users at any given time.

An update of our current experiences and challenges running an EOS instance for use by the Fermilab LHC Physics Center (LPC) computing cluster. Planning and implementation of our upgrade to EOS 5 and moving to Almalinux before EOL for SL7.

**Primary author:** SZKOLA, Dan (Fermi National Accelerator Lab. (US))

**Presenter:** SZKOLA, Dan (Fermi National Accelerator Lab. (US))

**Session Classification:** EOS Operations

Contribution ID: 6

Type: **not specified**

# Comparison of High-Performance Distributed File Systems on two Platforms: Linux and Windows

Friday, 15 March 2024 14:45 (15 minutes)

Subtitle: Superiority of EOS-based Comtrade Distributed File System (CDFS) for Earth Observation Data Storage.

## Introduction

The rising quantity of data collected necessitates transitioning to the next generation of reliable, high-performance data storage solutions. Despite the clear need for high-performance storage across many industries, there has yet to be a consensus on the optimal high-performance storage technology that would suit all user needs.

The project's goal was speed-testing to provide a possibility to combine them with the currently available state-of-the-art technology results, technological capabilities and existing expertise.

## Methodological approach

We choose representative high-performance storage solutions compared on two different platforms for storage clients.

- Linux: - Ceph - EOS - IBM Spectrum Scale - Hadoop
- Windows: - Ceph - EOS-drive - EOS through Samba - EOS-wnc - Hadoop

An example of specifications of the testing environment is the following.

- *Management node*: - 32 threads - 2x Intel® Xeon® Silver 4208 Processor - 384 GB of RAM - 2x SSDs of 2 TB.
- *Storage nodes*: - 32 threads - 2x Intel® Xeon® Silver 4208 Processor - 64 GB of RAM - 1x SSD of 2 TB for the operation system - 6x HDDs of 2 TB for data.
- *Client nodes*: - 12 threads - 1x Intel® Core™ i5-12400 Processor - 16 GB of RAM - 1x SSD with 1 TB.

## Results

The results are from the testing performed separately in an isolated environment for each high-performance solution but on the same hardware. For each platform, there were three categories of tests related to file sizes: (1) small, (2) medium, and (3) large files.

According to the results, some high-performance file systems have evident advantages, as shown in our presentation. These results should be the starting point for an even more exact comparison between these file systems. They are good starting points in choosing the right high-performance file system.

**Primary author:** MOLAN, Gregor (Comtrade 360's AI Lab)

**Presenter:** MOLAN, Gregor (Comtrade 360's AI Lab)

**Session Classification:** EOS Operations

Contribution ID: 7

Type: **not specified**

## **eosxd-csi: Mounting EOS volumes in Kubernetes**

*Friday, 15 March 2024 13:35 (10 minutes)*

Kubernetes is seeing a huge adoption across the cloud, including the one at CERN. Already used in production, the EOSxd-CSI driver exposes EOS volumes as regular PersistentVolumeClaims that can be mounted by containerized workloads. This talk will touch on storage stack in Kubernetes, how EOS falls into the mix, and finally a demo showing the CSI driver in action.

**Primary author:** VASEK, Robert (CERN)**Presenter:** VASEK, Robert (CERN)**Session Classification:** EOS Operations



Contribution ID: **8**

Type: **not specified**

## **EOS 5.2 Status**

*Friday, 15 March 2024 08:30 (20 minutes)*

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

**Session Classification:** EOS Development

Contribution ID: 9

Type: **not specified**

## EOS III

*Friday, 15 March 2024 09:25 (10 minutes)*

III = Instance Inventory Implementation

We have added recently new tools to gather statics about storage hardware, resource usage, hardware lifecycle. These tools compute the virtual cost and value of user data and the hardware of an EOS instances.

The presentation will introduce these new tools.

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

**Session Classification:** EOS Development

Contribution ID: **10**

Type: **not specified**

## **EOS ALMA9 Migration**

*Friday, 15 March 2024 14:05 (15 minutes)*

**Presenter:** VRACHNAKI, Ioanna

**Session Classification:** EOS Operations

Contribution ID: **11**

Type: **not specified**

## **EOS on ARM**

*Friday, 15 March 2024 10:55 (5 minutes)*

**Presenter:** LEKSHMANAN, Abhishek (CERN)

**Session Classification:** EOS Development

Contribution ID: 12

Type: **not specified**

## EOSXd Evolution and CFSd

*Friday, 15 March 2024 09:35 (15 minutes)*

EOSXd is the filesystem client implemented as a FUSE filesystem to provide POSIX-like access to EOS. This is a crucial component for general usage of EOS. The presentation will highlight the problems we faced and advancements since the last workshop.

EOS CFSd is a FUSE pass-through filesystem implementation allowing to add missing features to any general POSIX filesystems. The presentation will highlight some possible use-cases and performance measurements e.g. an example how CFSd adds kerberos based mapping to a CephFS filesystem without a significant performance impact.

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

**Session Classification:** EOS Development

Contribution ID: **13**

Type: **not specified**

## **EOS on SMR Status**

*Friday, 15 March 2024 11:00 (5 minutes)*

In close collaboration with the IT procurement team we have conducted R&D with EOS and SMR disks. The findings, current status and outlook on integration of SMR disks will be covered in this presentation.

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

**Session Classification:** EOS Development

Contribution ID: 14

Type: **not specified**

## EOS Namespace Locking Evolution

*Friday, 15 March 2024 08:50 (20 minutes)*

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

**Session Classification:** EOS Development

Contribution ID: 15

Type: **not specified**

## EOS FlatScheduler & Freespace Engine

*Friday, 15 March 2024 09:10 (15 minutes)*

**Presenter:** LEKSHMANAN, Abhishek (CERN)

**Session Classification:** EOS Development



Contribution ID: 16

Type: **not specified**

## **Group Balancer Free Space Engine**

**Presenter:** LEKSHMANAN, Abhishek (CERN)

**Session Classification:** EOS Development

Contribution ID: 17

Type: **not specified**

## HTTP Improvements and SciTags

*Friday, 15 March 2024 10:35 (10 minutes)*

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

**Session Classification:** EOS Development

Contribution ID: **18**

Type: **not specified**

## **EOS and fixes & low level changes for openssl, xrootd & eosxd**

*Friday, 15 March 2024 10:45 (10 minutes)*

**Presenter:** SMITH, David (CERN)

**Session Classification:** EOS Development

Contribution ID: 20

Type: **not specified**

## A new REST API Gateway for EOS

*Friday, 15 March 2024 10:10 (15 minutes)*

**Co-author:** SINDRILARU, Elvin Alin (CERN)

**Presenter:** PRIGOREANU, Andreea (IT-SD)

**Session Classification:** EOS Development

Contribution ID: 22

Type: **not specified**

## **EOS4PHysics Operations**

*Friday, 15 March 2024 13:45 (20 minutes)*

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

**Session Classification:** EOS Operations

Contribution ID: 23

Type: **not specified**

## Long Term Monitoring with Prometheus + Thanos

*Friday, 15 March 2024 14:30 (15 minutes)*

The Storage and Data Management Group at CERN manages 20 EOS instances corresponding to almost 1000 servers and 100,000 disks. Having a good monitoring and alerting system is crucial not only for day-to-day activities but also as a tool to record the evolution of our services throughout the time. In this talk an overview of the monitoring tools that are used will be presented specially in regards of long-term metric preservation.

**Presenter:** VALVERDE CAMESELLE, Roberto (CERN)

**Session Classification:** EOS Operations

Contribution ID: 24

Type: **not specified**

## EOS Benchmarks at CERN

*Friday, 15 March 2024 11:50 (10 minutes)*

Before and after Run 3 have conducted several performance benchmarks as preparation of DC'24 and Run-3 in 2024 on various EOS instances at CERN. The presentation will report the findings of these benchmarks.

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

**Session Classification:** EOS Development

Contribution ID: 25

Type: **not specified**

## EOS Hardware Procurement

*Friday, 15 March 2024 14:20 (10 minutes)*

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

**Session Classification:** EOS Operations



Contribution ID: 26

Type: **not specified**

## CERNBox Update

*Friday, 15 March 2024 15:00 (10 minutes)*

**Presenter:** BAGAKIS, Emmanouil (CERN)

**Session Classification:** EOS Operations

Contribution ID: 27

Type: **not specified**

## CIFS Update

**Presenter:** LOPRESTI, Giu

**Session Classification:** EOS Operations

Contribution ID: **28**

Type: **not specified**

## **EOS on lxplus/lxbatch**

**Session Classification:** EOS Operations

Contribution ID: 29

Type: **not specified**

## CTA Update

**Session Classification:** EOS Operations

Contribution ID: **30**

Type: **not specified**

## **Configuring EOS**

**Session Classification:** Meet the Teams

Contribution ID: 31

Type: **not specified**

## Implementing FSCK for Erasure Coded Files in EOS

*Friday, 15 March 2024 10:25 (10 minutes)*

**Presenter:** SEGRANSAN, Mano (42 Lausanne (CH))

**Session Classification:** EOS Development

Contribution ID: **32**

Type: **not specified**

## **A native S3 interface EOS/XRootD**

*Friday, 15 March 2024 11:05 (10 minutes)*

**Presenter:** SEGRANSAN, Mano (42 Lausanne (CH))

**Session Classification:** EOS Development

Contribution ID: **33**

Type: **not specified**

## Roadmap 2024

*Friday, 15 March 2024 17:10 (10 minutes)*

The presentation will summarize plans of the EOS team for 2024 and beyond.

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

**Session Classification:** Final Session



Contribution ID: 34

Type: **not specified**

## **XRootD Status and Plans**

*Friday, 15 March 2024 11:15 (15 minutes)*

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

**Session Classification:** EOS Development

Contribution ID: 35

Type: **not specified**

## **CERN Storage Day: Welcome to the IT Department**

*Thursday, 14 March 2024 09:00 (10 minutes)*

**Presenters:** DUELLMANN, Dirk (CERN); MOSCICKI, Jakub (CERN)

**Session Classification:** CERN Storage Technologies

Contribution ID: 36

Type: **not specified**

## **CERN IT Storage and Data Management Group Overview**

*Thursday, 14 March 2024 09:10 (20 minutes)*

**Presenter:** KEEBLE, Oliver (CERN)

**Session Classification:** CERN Storage Technologies

Contribution ID: 37

Type: **not specified**

## Large Scale Disk Storage: EOS

*Thursday, 14 March 2024 09:30 (30 minutes)*

**Primary author:** PETERS, Andreas Joachim (CERN)

**Co-author:** SINDRILARU, Elvin Alin (CERN)

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

**Session Classification:** CERN Storage Technologies

Contribution ID: **38**

Type: **not specified**

## **Large Scale Tape Storage: CTA**

*Thursday, 14 March 2024 10:00 (30 minutes)*

**Presenter:** AFONSO, Joao (CERN)

**Session Classification:** CERN Storage Technologies

Contribution ID: 39

Type: **not specified**

## User and Collaboration Storage: CERNBox

*Thursday, 14 March 2024 10:50 (20 minutes)*

**Presenter:** LO PRESTI, Giuseppe (CERN)

**Session Classification:** CERN Storage Technologies

Contribution ID: 40

Type: **not specified**

## Infrastructure Storage: Ceph

*Thursday, 14 March 2024 11:10 (20 minutes)*

**Presenter:** BOCCHI, Enrico (CERN)

**Session Classification:** CERN Storage Technologies

Contribution ID: 41

Type: **not specified**

## **More Filesystems: AFS,FILER,DFS,SAMBA,,...**

*Thursday, 14 March 2024 11:30 (15 minutes)*

**Presenters:** BOCCHI, Enrico (CERN); BUKOWIEC, Sebastian (CERN)

**Session Classification:** CERN Storage Technologies



Contribution ID: 42

Type: **not specified**

## Content Delivery Storage: CVMFS

*Thursday, 14 March 2024 11:45 (15 minutes)*

**Presenter:** VOLKL, Valentin (CERN)

**Session Classification:** CERN Storage Technologies

Contribution ID: 43

Type: **not specified**

## File Transfer Service (FTS)

*Thursday, 14 March 2024 12:00 (15 minutes)*

**Presenter:** LOPES, Joao

**Session Classification:** CERN Storage Technologies

Contribution ID: 44

Type: **not specified**

## Scientific Data Management: Rucio

*Thursday, 14 March 2024 12:15 (15 minutes)*

**Presenter:** BARISITS, Martin (CERN)

**Session Classification:** CERN Storage Technologies

Contribution ID: 45

Type: **not specified**

## EOS site report of the Joint Research Centre

*Friday, 15 March 2024 15:55 (15 minutes)*

The Joint Research Centre (JRC) of the European Commission is running the Big Data Analytics Platform (BDAP) to enable the JRC projects to store, process, and analyze a wide range of data and disseminate data products. The platform evolved as a core service for JRC scientists to produce knowledge and insights in support of EU policy making.

EOS is the main storage system of the BDAP for scientific data. It is in increasing use at JRC since 2016. The Big Data Analytics Platform is actively used by more than 90 JRC projects, covering a wide range of data analytics activities. The EOS instance at JRC has currently a gross capacity of 34 PB with an additional increase planned throughout 2024.

The presentation will give an overview about EOS as storage back-end of the Big Data Analytics Platform. It covers the general setup and current status, experiences made, issues discovered, and an outlook of planned activities and changes in 2024.

**Primary author:** BURGER, Armin (EC Joint Research Centre)

**Co-author:** Mr EYRAUD, Franck (EC Joint Research Centre)

**Presenter:** BURGER, Armin (EC Joint Research Centre)

**Session Classification:** EOS Operations

Contribution ID: 46

Type: **not specified**

## NDMSPC - EOS and N-Dimensional Analysis with ROOT, Enhanced by Web Interface and VR Visualization

*Friday, 15 March 2024 11:30 (10 minutes)*

This talk delves into the synergy of NDMSPC (NDimensional SPaCe), EOS at CERN, and N-dimensional histograms via the ROOT framework. Discover how a web interface serves as a powerful analysis tool, enabling dynamic queries on projections in N-dimensional space. Interact and visualize the data seamlessly with JSROOT and VR (aframe), ushering in a new era of immersive exploration in high-energy physics.

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

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

**Session Classification:** EOS Development

Contribution ID: 47

Type: **not specified**

## Purdue EOS status report

*Friday, 15 March 2024 16:40 (15 minutes)*

Purdue University switched to EOS storage a couple of years back, and we are still exploring/appreciating the benefits of this distributed file-system. In this talk we will give a brief status report and will outline our plans for using EOS at the CMS Tier-2 center at Purdue.

**Primary author:** PIPEROV, Stefan (Purdue University (US))

**Presenter:** PIPEROV, Stefan (Purdue University (US))

**Session Classification:** EOS Operations

Contribution ID: 48

Type: **not specified**

## Meet the team: EOS & XRootD

*Thursday, 14 March 2024 14:00 (3 hours)*

**Session Classification:** Meet the EOS/XRootD Team

Contribution ID: 49

Type: **not specified**

## Meet the team: Ceph FS

*Thursday, 14 March 2024 14:00 (30 minutes)*

Discussion on the Ceph Filesystem.

**Session Classification:** Meet the CEPH Team



Contribution ID: 50

Type: **not specified**

## **Meet the team: CERN Tape Archive (CTA) Operations**

*Thursday, 14 March 2024 14:00 (1h 30m)*

This session is open to all, but is geared towards sites with specific questions about CTA and tape operations.

**Session Classification:** Meet the CTA Team

Contribution ID: 51

Type: **not specified**

## EOS Status at IHEP

*Friday, 15 March 2024 16:10 (15 minutes)*

The Institute of High Energy Physics (IHEP) has been utilizing the EOS storage system since 2016, with the current capacity nearing 60PB. We use EOS to provide both disk and tape storage solutions for HEP experiments. This report will introduce some of work we carried out in 2023, which include:

1. The adoption of EOS as SE for the WLCG Tier2, replacing the previous DPM system, and the provision of SE and CTA services for the LHCb Tier1 via EOS.
2. The upgrade of our system from EOS V4 to V5, which addressed the issue of high concurrency demands posed by the LHAASO experiment instance.
3. The establishment of a new ARM EOS cluster, which has reached a capacity of 2PB.

Additionally, we have conducted extensive testing and validation of EOS on AlmaLinux. Our plan is to migrate our entire computing cluster to AlmaLinux 9.2 by the end of June this year.

**Primary author:** LIHAIBO, LI Haibo

**Presenter:** LIHAIBO, LI Haibo

**Session Classification:** EOS Operations

Contribution ID: 52

Type: **not specified**

## Roadmap Discussion

*Friday, 15 March 2024 17:20 (10 minutes)*

**Session Classification:** Final Session

Contribution ID: 53

Type: **not specified**

## Shared EOS instance at JINR

*Friday, 15 March 2024 16:55 (15 minutes)*

The Joint Institute for Nuclear Research (JINR) utilizes a diverse storage ecosystem, encompassing a number of dCache, EOS and Ceph storage instances to address varied scientific needs. In this report we focus on one of the largest EOS instances at JINR, operational since 2019 as a shared storage system for experimental data. Presently, it has the capacity of 22 PB and hosts over 7 PB of data from numerous experiments and projects. We give an overview of the current setup, development plans and share our experience of operating it.

**Primary author:** BALASHOV, Nikita (Joint Institute for Nuclear Research (RU))

**Presenter:** BALASHOV, Nikita (Joint Institute for Nuclear Research (RU))

**Session Classification:** EOS Operations

Contribution ID: 54

Type: **not specified**

## ALICE File Consistency Check System: A new solution based on EOS FSCK

*Friday, 15 March 2024 11:40 (10 minutes)*

Undoubtedly, the processing of ALICE experiment data relies on the quality and integrity of data. Currently, ALICE uses a distributed file crawler that periodically evaluates samples of files from each storage element in order to gather statistics about the number of corrupted or inaccessible files. The main issue with this solution lies in its inability to provide a comprehensive overview of a storage element status, as the analysis results are based on examining a random selection of files. This presentation will describe a new solution for the ALICE File Consistency Check System. The new approach will overcome the limitations of the file crawler by using the powerful consistency checking tools provided by EOS. The idea behind this project is to collect all the existing errors on an EOS instance from the reports generated by the FSCK command with the goal of reconciling the contents of the local storage with the central catalogue and, where possible, recover the lost content from other replicas. The output of the FSCK report command will be accessed through the new HTTP interface available in the latest versions of EOS.

Hence, this solution not only produces a more accurate integrity analysis but automates the recovery of data loss as well.

**Presenter:** PRIGOREANU, Andreea (IT-SD)

**Session Classification:** EOS Development

Contribution ID: 55

Type: **not specified**

## **Meet the team: FTS/CVMFS/DFS**

*Thursday, 14 March 2024 14:00 (3 hours)*

CVMFS: 513/R-070 - Openlab Space

FTS: 28/S-029

DFS: 31/2-28

**Session Classification:** Meet the FTS/CVMFS/DFS Team

Contribution ID: 56

Type: **not specified**

## **Hands-on CERNBox/Reva and Cephfs**

*Thursday, 14 March 2024 15:00 (1 hour)*

**Session Classification:** Meet the CERNBox Team

Contribution ID: 57

Type: **not specified**

## **Hands-on CERNBox/Reva and EOS**

*Thursday, 14 March 2024 16:00 (1 hour)*

**Session Classification:** Meet the CERNBox Team



Contribution ID: 58

Type: **not specified**

## Meet the team: Ceph General

*Thursday, 14 March 2024 15:30 (1h 30m)*

Ask us anything about Ceph.

**Session Classification:** Meet the CEPH Team

Contribution ID: 59

Type: **not specified**

## Meet the team: Ceph S3

*Thursday, 14 March 2024 14:30 (30 minutes)*

Discussion on ceph-provisioned S3 object storage.

**Session Classification:** Meet the CEPH Team

Contribution ID: **60**

Type: **not specified**

## **Meet the team: Ceph RBD**

*Thursday, 14 March 2024 15:00 (30 minutes)*

Discussion on Ceph RADOS Block Devices.

**Session Classification:** Meet the CEPH Team

Contribution ID: **61**

Type: **not specified**

## **Meet the team: CERN Tape Archive (CTA) Newcomers**

*Thursday, 14 March 2024 15:30 (1h 30m)*

This session is open to all, but is geared towards sites who are newcomers to CTA or are interested in evaluating CTA for their archival storage needs.

**Session Classification:** Meet the CTA Team

Contribution ID: **62**

Type: **not specified**

## **CVMFS: Openlab room**

**Session Classification:** Meet the FTS/CVMFS/DFS Team

Contribution ID: **63**

Type: **not specified**

## **Group Photo and Logistics**

*Thursday, 14 March 2024 12:30 (10 minutes)*

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