EOS 2025 Workshop

Monday 24 March 2025 - Wednesday 26 March 2025

CERN

Book of Abstracts

Contents

Development / **1**

Storage Tiering in EOS

Author: Andreas Joachim Peters[1]

[1] *CERN*

Corresponding Author: andreas.joachim.peters@cern.ch

We will give an overview of new features for storage tiering in EOS version 5.3

Development / **2**

Status of the S3 Interface for EOS

Author: Andreas Joachim Peters[1]

[1] *CERN*

Corresponding Author: andreas.joachim.peters@cern.ch

We will present an overview of the current state of the S3 gateway for EOS.

Roadmap / **3**

The EOS Development Workplan & Roadmap

Author: Andreas Joachim Peters[1]

Co-author: Elvin Alin Sindrilaru [1]

[1] *CERN*

Corresponding Authors: elvin.alin.sindrilaru@cern.ch, andreas.joachim.peters@cern.ch

We will outline the EOS development roadmap, highlighting key milestones, upcoming features, and future plans. This presentation will provide insights into ongoing improvements, strategic goals, and the evolving direction of EOS.

Site Reports / **4**

EOS site report of the Joint Research Centre

Author: Armin Burger[None]

Co-author: Franck Eyraud

Corresponding Authors: franck+cern@yrnm.net, armin.burger@ec.europa.eu

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 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 been implemented in 2016 and has currently a gross capacity of 43 PB. It is composed of heterogeneous commodity hardware components which has been extended noticeably over time.

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 describes the activities over the last year, presents experiences made and issues discovered, and gives an outlook of planned activities during 2025.

Site Reports / **5**

EOS Status at IHEP

Authors: LI Haibo lihaibo[1]; Yaodong CHENG[1]; Yujiang BI[1]

[1] *Institute of High Energy Physics, Chinese Academy of Sciences*

Corresponding Authors: chyd@ihep.ac.cn, biyujiang@ihep.ac.cn, lihaibo@ihep.ac.cn

In this talk, we want to share our experiences of EOS at IHEP, including migration from CentOS 7 to Almalinux 9, construction of Alice EOS, and dual-site deployment of LHCb T1 EOS.

Site Reports / **6**

Planning an EOS Data Federation to deal with Climate Change using AI

Author: Wanderley Mendes[1]

Co-author: Paulo Nobre [1]

[1] *INPE*

Corresponding Authors: wanderley.mendes@inpe.br, paulo.nobre@inpe.br

The National Institute for Space Research - INPE (Brazil) is leading a research program: Intelligent Early Warning System for Climate Extremes - SIPEC. The project aims at predicting the likelihood of climate extremes, months in advance using a diverse source of data coming from satellites and an array of intelligent sensors spread across the country. Such data streams will feed both classical meteorological models and AI machine learning algorithms for the ultimate early warning of climate extremes.

Given the number of institutions producing large amounts of data needed to train the ML algorithms by scientists dealing with different parts of the problem, at different places, we are implementing an EOS Data Federation in Brazil. The implementation of the EOS family of tools, in addition to being capable to deal with large volumes of distributed data, also takes care of security controls for who has access to what portions of the datasets.

Site Reports / **7**

Cloud-Native EOS Deployment for ATLAS T2 on Kubernetes

Author: Ryan Taylor[1]

[1] *University of Victoria (CA)*

Corresponding Author: rptaylor@uvic.ca

I will discuss our Kubernetes-based EOS deployment as it approaches production readiness for our ATLAS T2 site, as well as evaluation of EOS for several astronomy projects.

Development / **8**

XRootD File Cloning

Author: David Smith[1]

[1] *CERN*

Corresponding Author: david.smith@cern.ch

A software development motivated by an EOS use case is explained: file cloning to facilitate updates of erasure-coded files.

9

EOS FSCK incremental improvements

Author: Elvin Alin Sindrilaru[1]

[1] *CERN*

Corresponding Author: elvin.alin.sindrilaru@cern.ch

In this presentation, we will go through the recent changes of the FSCK functionality that is now capable of addressing both RAIN and replica file layouts. Moreover, this will also include detail on the recently added best-effort functionality to address files which are in a broken state.

Development / **10**

EOS and XRootD HTTP improvements

Author: Cedric Caffy[1]

[1] *CERN*

Corresponding Author: cedric.caffy@cern.ch

With the continuous growth in the use of the HTTP protocol for file transfers within the WLCG community, several enhancements and optimisations have been introduced to the EOS HTTP and XRootD HTTP stacks.

From updates to the SciTags and packet marking specifications to addressing libcurl internal modifications, 2024 presented a number of challenges that required targeted solutions.

This presentation will provide an overview of the key changes and new features implemented to enhance the handling of HTTP file transfers in EOS and XRootD.

Site Evolution / **11**

A Distributed Storage Odyssey: from CentOS7 to ALMA9

Author: Cedric Caffy[1]

[1] *CERN*

Corresponding Author: cedric.caffy@cern.ch

On the 30th of June 2024, the end of CentOS 7 support marked a new era for the operation of the multi-petabytes distributed disk storage system used by CERN physics experiments. The EOS infrastructure at CERN is composed of aproximately 1000 disk servers and 50 metadata management nodes. Their transition from CentOS 7 to Alma 9 was not as straightforward as anticipated.

This presentation will be all about explaining this transition. From the change of supported certificate and kerberos key signature lengths and algorithms, to openssl library hiccups and Linux kernel crashes, the EOS operation team had to take on different challenges to ensure a seamless operating system transition of the infrastructure while maintaining uninterrupted CERN experiments' data transfers.

Site Evolution / **12**

Refurbishing the Meyrin Data Centre: Storage Juggling and Operations

Author: Octavian-Mihai Matei[1]

[1] *CERN*

Corresponding Author: octavian-mihai.matei@cern.ch

The 50-year-old Meyrin Data Centre (MDC), still remains indispensable due to its strategic geographical location and unique electrical power resilience even if CERN IT recently commissioned the Prévessin Data Centre (PDC), doubling the organization's hosting capacity in terms of electricity and cooling. The Meyrin Data Centre (Building 513) retains an essential role for the CERN Tier-0 Run 4 commitments, notably as primary hosting location for the tape archive and the disk storage. The inevitable investments to the infrastructure (UPS and Cooling) are now triggering the refurbishment of the two main rooms where all the storage equipment is hosted. This presentation will delve into the architectural advancements and operational strategies implemented for and during the Meyrin data centre refurbishment. We will explore how these developments will impact our storage and how the storage operations team will ensure EOS's performance, scalability, and reliability in the coming years.

Development / **13**

QClient Improvements for the next fastest Metadata

Author: Abhishek Lekshmanan[1]

[1] *CERN*

Corresponding Author: abhishek.lekshmanan@cern.ch

Every operation that modifies/queries the metadata from the persistent metadata storage QuarkDB goes via QClient. We look at some current bottlenecks and improvements that v5.3 offers with various configurations.

Operational Tools & Configuration / **14**

Diagnostic tool for submitting useful information for future debugging

Author: Abhishek Lekshmanan[1]

[1] *CERN*

Corresponding Author: abhishek.lekshmanan@cern.ch

For a stuck/non responsive EOS MGM, some simple diagnostic information can go a long way. We look at a new eos-diagnostic-tool for dumping stacktraces etc. for submitting useful bug reports. We also invite discussions on how to improve the tooling for the future.

Operational Tools & Configuration / **15**

Data Federations with EOS

Author: Luca Mascetti[1]

[1] *CERN*

Corresponding Author: luca.mascetti@cern.ch

Data federations with EOS offers various approaches to seamlessly integrate and manage distributed storage across heterogeneous environments. This presentation explores multiple federation techniques and namespace aggregation with remote EOS instances. We will discuss the advantages and trade-offs of each method, considering factors such as performance, scalability, security, and ease of management. Real-world use cases and best practices will be highlighted to help organisations choose the most suitable strategy for their needs.

Development / **16**

Advancements in FSCK for EOS

Authors: Elvin Alin Sindrilaru[1]; Gianmaria Del Monte[1]

[1] *CERN*

Corresponding Authors: elvin.alin.sindrilaru@cern.ch, gianmaria.del.monte@cern.ch

One of a critical components in EOS is fsck, responsible for scanning, verifying, and repairing inconsistencies in the filesystem.

This talk will provide an in-depth exploration of fsck in EOS, covering its architecture, scanning mechanisms, and repair strategies. We will discuss recent improvements, including the introduction of a best-effort mode, and enhancements in erasure-coded file scanning, which significantly boost performance while minimizing the impact on the running instance.

Operational Tools & Configuration / **17**

A Distributed Probe for EOS: Real-Time Availability Monitoring and Alerting

Authors: Gianmaria Del Monte[1]; Octavian-Mihai Matei[1]

[1] *CERN*

Corresponding Authors: gianmaria.del.monte@cern.ch, octavian-mihai.matei@cern.ch

Ensuring the availability of EOS instances is crucial for large-scale storage operations. To enhance monitoring and incident response, we have developed a new distributed probe designed to detect and alert operators about instance malfunctions in real-time.

This talk will introduce the architecture and functionality of the probe, which runs across multiple nodes to provide redundancy and reliability. Alerts are dispatched via multiple channels, including SMS, email, Mattermost, and CERN IT's General Services Availability. Additionally, all availability events are published on a NATS-based pub-sub channel, enabling future integrations with operational tools such as EOS Diagnostic Tool.

Site Reports / **18**

EOS for Physics at CERN: Operational Insights, Achievements, and Future Directions

Author: Maria Arsuaga Rios[1]

[1] *CERN*

Corresponding Author: maria.arsuaga.rios@cern.ch

This work presents an overview of the EOS operations at CERN, focusing on its role in supporting physics data processing and storage. EOS is a high-performance distributed storage system designed to handle the vast volumes of scientific data generated by CERN experiments. This study examines key performance metrics, recent achievements, and strategic objectives for the current year, emphasizing improvements in efficiency, reliability, and scalability. Special attention is given to the impact of EOS on physics workflows, ensuring seamless data access and analysis. By evaluating past accomplishments and future goals, this work highlights the continuous evolution of EOS to meet the growing demands of physics research at CERN.

Site Evolution / **19**

Evaluating Jumbo frames performance across LHC experiments

Author: Maria Arsuaga Rios[1]

[1] *CERN*

Corresponding Author: maria.arsuaga.rios@cern.ch

This work presents an evaluation of JUMBO frame tests conducted at CERN to assess their impact on data transfer performance across different physics workflows. Preliminary internal tests were carried out to analyze potential benefits and challenges, followed by collaborative testing involving the ATLAS, CMS, and LHCb experiments. The goal was to measure the advantages of JUMBO frames in terms of efficiency and throughput while identifying and resolving any issues arising from their deployment. The study provides insights into the feasibility of JUMBO frames for large-scale scientific data transfers, aiming to optimize network performance for high-energy physics experiments.

Operational Tools & Configuration / **20**

Deploying an EOS Instance from Scratch: A Practical Guide

Author: Gianmaria Del Monte[1]

[1] *CERN*

Corresponding Author: gianmaria.del.monte@cern.ch

EOS is a powerful and flexible storage system, but setting up a new instance from scratch requires a solid understanding of its configuration and operational best practices. This talk will provide a step-by-step guide to deploying EOS, covering key components and essential configurations.

We will walk through the setup process, including storage provisioning, replication, erasure coding, and balancing strategies. The session will also touch on best practices for performance tuning and ensuring reliability in production environments.

This talk is ideal for system administrators and operators looking to gain practical insights into EOS deployment, whether for testing, small-scale clusters, or large production environments.

Site Reports / **21**

CERNBox and EOSHPM status update

Authors: Jan Iven[1]; Diogo Castro[1]

Co-author: Pablo Medina Ramos

[1] *CERN*

Corresponding Authors: diogo.castro@cern.ch, jan.iven@cern.ch, pablo.medina.ramos@cern.ch

CERNBox and EOS HOME/PROJECT(/MEDIA) operational issues seen in 2024 and expected in 2025.

Development / **22**

EOS 5.2 and 5.3 Status / Overview

Corresponding Author: elvin.alin.sindrilaru@cern.ch

This presentation will give a short overview of the past releases and significant changes, new features and bug fixes.

Benchmarking & Hardware Evolution / **23**

Storage Hardware at CERN

Corresponding Author: luca.mascetti@cern.ch

Current & future storage hardware at CERN.

Benchmarking & Hardware Evolution / **24**

Analysis Benchmarking with EOS/RNTuple

Corresponding Author: andreas.joachim.peters@cern.ch

This presentation will report about the benchmarking results of various EOS setups at CERN using the new RNTuple framework.

Benchmarking & Hardware Evolution / **25**

XRootD Update & Parallel Socket Benchmarking

Corresponding Author: amadio@cern.ch

Roadmap / **26**

Discussion, Proposals, Feature Requests

Survey Topics
- Ansible Configuration for EOS
- SquashFS as small File Repository

Hands On / **27**

How to benchmark EOS with bash?

Corresponding Author: andreas.joachim.peters@cern.ch

Hands On / **28**

How to setup authentication Front-ends?

Corresponding Author: elvin.alin.sindrilaru@cern.ch

Hands On / **29**

How to configure TLS & ZTN in XRootD?

Corresponding Author: amadio@cern.ch

Hands On / **30**

How to transtion from MQ to no MQ?

Corresponding Author: elvin.alin.sindrilaru@cern.ch

Site Reports / **31**

You still have those QDB backups, right? (Practical example of disaster recovery of EOS deployment)

Author: Stefan Piperov[1]

[1] *Purdue University (US)*

Corresponding Author: stefan.piperov@cern.ch

In December of 2024 the EOS cluster at Purdue University suffered a security incident which wiped out all metadata of our production deployment. In this brief talk we will give a step-by-step example of what it takes to recover from such setback, and discuss the best backup practices.