



Contribution ID: 23

Type: **Presentation**

## Converging Storage Layers with Virtual CephFS Drives for EOS/CERNBox

*Wednesday, 26 January 2022 10:00 (20 minutes)*

The CERNBox service is currently backed by 13PB of EOS storage distributed across more than 3,000 drives. EOS has proven to be a reliable and highly performing backend throughout. On the other hand, the CERN Storage Group also operates CephFS, which has been previously evaluated in combination with EOS as a potential solution for large scale physics data taking [1]. This work seeks to further explore the operational benefits of a combined EOS/CephFS solution as a CERNbox backend. First, we present the functional validation work done using a canary instance and existing micro benchmarks. Next, we show how the solution was gradually introduced to production, observing the relative impacts of metadata and backend storage on user perceived small op performance. Finally, the qualitative impact of the solution is discussed: potential for enhanced QoS (e.g. policy driven low latency vs low-cost areas), simplification of hardware operations across the entire lifecycle, and how the work may enable future cloud-based deployments.

[1] <https://doi.org/10.1007/s41781-021-00071-1>

**Primary authors:** PETERS, Andreas Joachim (CERN); VAN DER STER, Dan (CERN); VALVERDE CAMESELLE, Roberto (CERN)

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

**Session Classification:** Scalable Storage Backends

**Track Classification:** Main session: Scalable Storage Backends for Cloud, HPC and Global Science