Contribution ID: **43** Type: **Presentation**

EOS the CERN disk/cloud storage for science

Tuesday, January 30, 2018 11:50 AM (20 minutes)

EOS, the high-performance CERN IT distributed storage for High-Energy Physics provides now more than 250PB of raw disks and supports several work-flows from LHC data-taking and reconstruction to physics analysis.
The software is developed at CERN since 2010, is available under GPL license and it is also used in several external institutes and organisations.

EOS is the key component behind the success of CERNBox, the CERN cloud synchronisation service for end-users which allows syncing and sharing files on all major mobile and desktop platforms aiming to provide offline availability to any data stored in the EOS infrastructure.

Today EOS provides multiple views/protocols to the same namespace and storage backend - via the OwnCloud synchronisation client, via xrootd protocol for physics data analysis applications access, as a mounted filesystem, with latency optimised wide-area file access protocols or via SAMBA endpoint for Windows clients.
In addition is possible to interact with the system using Jupyter Notebooks provided at CERN by the SWAN (Service for Web based ANalysis) platform which offers to scientists a web based service for interactive data analysis.

We report on our experience with this technology and applicable use-cases, also in a broader scientific and research context and its future evolution with highlights on the current development status and future roadmap.

Primary author: MASCETTI, Luca (CERN)

Presenter: MASCETTI, Luca (CERN)

Session Classification: Scalable Storage Backends for Cloud and HPC: Foundations