Contribution ID: **45** Type: **presentation**

EOS Open Storage - evolution of an ecosystem for scientific data repositories

Thursday 12 July 2018 17:00 (20 minutes)

The EOS project started as a specialized disk-only storage software solution for physics analysis use-cases at CERN in 2010.
Over the years EOS has evolved into an open storage platform, leveraging several open source building blocks from the community. The service at CERN manages around 250 PB, distributed across two data centers and provides user- and project-spaces to all CERN experiments. Strategic collaborations with many WLCG sites and non-HEP community members such as JRC, AARnet and COMTRADE helped to broaden and generalize its usability. Another recent important milestone was the commissioning of the CERNBOX and SWAN services which enable EOS as a sync-and-share platform. In 2017 the initial EOS architecture reached its design limitations at CERN with a visible impact on service stability. In 2018 we'll put a revised architecture into production: this includes a highly scalable namespace implementation, generalized POSIX-like filesystem access and coalescence of EOS storage and the CERN tape archive. As a participating member in the extreme data cloud project and WLCG data lake R&D activity, a new important project focus is to enable cost-effective distributed storage management incorporating existing storage and dynamic cloud resources. To meet scalability requirements expected of LHC run 3, EOS is evolving to act as a hybrid unifying platform for file and object storage, which might have a significant impact on the way physics data can be handled in the future by experiment frameworks and applications.

Primary authors: PETERS, Andreas Joachim (CERN); BITZES, Georgios (CERN); SIMON, Michal Kamil (CERN); MAKAI, Jozsef (CERN); SINDRILARU, Elvin Alin (CERN)

Presenter: PETERS, Andreas Joachim (CERN)

Session Classification: Plenary

Track Classification: Track 4 - Data Handling