



Contribution ID: 302

Type: Oral

Erasure Coding for production in the EOS Open Storage system

Thursday 7 November 2019 14:30 (15 minutes)

The storage group of CERN IT operates more than 20 individual EOS storage services with a raw data storage volume of more than 280 PB. Storage space is a major cost factor in HEP computing and the planned future LHC Run 3 and 4 increase storage space demands by at least an order of magnitude.

A cost effective storage model providing durability is Erasure Coding (EC). The decommissioning of CERN's remote computer center (Wigner/Budapest) allows to reconsider the currently configured dual-replica strategy where EOS provides one replica in each computer center.

EOS allows to configure EC on a per file bases and exposes four different redundancy levels with single, dual, triple and fourfold parity to select different quality of service and variable costs.

This paper will highlight tests which have been performed to migrate files on a production instance from dual-replica to various EC profiles. It will discuss performance and operational impact, and highlight various policy scenarios to select the best file layout with respect to IO patterns, file age and file size.

We will conclude with the current status and future optimizations, an evaluation of cost savings and discuss an erasure encoded EOS setup as a possible tape storage replacement.

Consider for promotion

No

Authors: SIMON, Michal Kamil (CERN); SINDRILARU, Elvin Alin (CERN); PETERS, Andreas Joachim (CERN)

Presenter: PETERS, Andreas Joachim (CERN)

Session Classification: Track 4 –Data Organisation, Management and Access

Track Classification: Track 4 –Data Organisation, Management and Access