

The first disk-based custodial storage for the ALICE experiment

Tuesday, 18 May 2021 11:42 (13 minutes)

We proposed a disk-based custodial storage as an alternative to tape for the ALICE experiment at CERN to preserve its raw data.

The proposed storage system relies on RAIN layout – the implementation of erasure coding in the EOS storage suite, which is developed by CERN – for data protection and takes full advantage of high-density JBOD enclosures to maximize storage capacity as well as to achieve cost-effectiveness comparable to tape.

The system we present provides 18 PB of total raw capacity from the 18 set of high-density JBOD enclosures attached to 9 EOS front-end servers.

In order to balance between usable space and data protection, the system will stripe a file into 16 chunks on the 4-parity enabled RAIN layout configured on top of 18 containerized EOS FSTs.

Although the reduction rate of available space increases up to 33.3% with this layout, the estimated annual data loss rate drops down to $8.6 \times 10^{-5}\%$.

In this paper, we discuss the system architecture of the disk-based custodial storage, 4-parity RAIN layout, deployment automation, and the integration to the ALICE experiment in detail.

Primary author: AHN, Sang Un (Korea Institute of Science & Technology Information (KR))

Co-authors: KIM, Jeongheon (Korea Institute of Science and Technology Information); HAN, Heejune (Korea Institute of Science & Technology Information (KR)); LEE, Seung Hee (KiSTi Korea Institute of Science & Technology Information (KR)); YOON, Hee Jun (KiSTi Korea Institute of Science & Technology Information (KR))

Presenter: AHN, Sang Un (Korea Institute of Science & Technology Information (KR))

Session Classification: Storage

Track Classification: Distributed Computing, Data Management and Facilities