



Contribution ID: 96

Type: **Poster presentation**

di-EOS - "distributed EOS": Initial experience with split-site persistency in a production service

Monday, 14 October 2013 15:00 (45 minutes)

After the strategic decision in 2011 to separate tier-0 activity from analysis, CERN-IT developed EOS as a new petascale disk-only solution to address the fast-growing needs for high-performance low-latency data access. EOS currently holds around 22PB usable space for the four big experiment (ALICE, ATLAS, CMS, LHCb), and we expect to grow to >30PB this year. EOS is one of the first production services to be running in CERN's new facility located in Budapest: we foresee to have about a third of total EOS storage capacity in the new facility in 2013, making it the largest storage service in the new CERN computer centre.

We report on the initial experience running EOS as a distributed service (via the new CERN IT Agile Infrastructure tools and a new "remote-hands" contract) in a production environment, as well as on the particular challenges and solutions. Among these solutions we are investigating the stochastic geo-location of data replicas and countermeasures against "split-brain" scenarios for the new high-availability namespace. In addition we are considering optimized clients and draining procedures to avoid link overloads across the two CERN sites as well as maximising data-access and operations efficiency.

Summary

Primary authors: Mr PETERS, Andreas Joachim (CERN); IVEN, Jan (CERN); MASCETTI, Luca (CERN); ESPINAL CURULL, Xavier (CERN)

Presenter: ESPINAL CURULL, Xavier (CERN)

Session Classification: Poster presentations

Track Classification: Data Stores, Data Bases, and Storage Systems