

21st International Conference on Computing in High Energy and Nuclear Physics (CHEP2015)



Contribution ID: 323

Type: oral presentation

Disk storage at CERN

Tuesday, 14 April 2015 17:45 (15 minutes)

CERN IT DSS operates the main storage resources for data taking and physics analysis mainly via three system: AFS, CASTOR and EOS. The total usable space available for users is about 100 PB (with relative ratios 1:20:120). EOS deploys disk resources across the two CERN computer centres (Meyrin and Wigner) with a current ratio 60% to 40%. IT DSS is also providing sizable on-demand resources for general IT services most notably OpenStack and NFS clients. This is provided by our Ceph infrastructure and a few of proprietary servers (NetApp) for a total capacity of ~1 PB.

We will describe our operational experience and recent changes to these systems with special emphasis to the following items:

- Present usages for LHC data taking (new roles of CASTOR and EOS)
- Convergence to commodity hardware (nodes with 200-TB each with optional SSD) shared across all services
- Detailed study of the failure modes in the different services and approaches (RAID, RAIN, ZFS vs XFS, etc...)
- Disaster recovery strategies (across the two CERN computer centres)
- Experience in coupling commodity and home-grown solution (e.g. Ceph disk pools for AFS, CASTOR and NFS)
- Future evolution of these systems in the WLCG realm and beyond

Primary author: MASCETTI, Luca (CERN)

Co-authors: FIOROT, Alessandro (CERN); IERI, Andrea (CERN); CHAN KWOK CHEONG, Belinda (CERN); VANDER STER, Dan (CERN); LO PRESTI, Giuseppe (CERN); ROUSSEAU, Herve (CERN); GONZALEZ LABRADOR, Hugo (University of Vigo (ES)); Dr MOSCICKI, Jakub (CERN); IVEN, Jan (CERN); LAMANNA, Massimo (CERN); PONCE, Sebastien (CERN); Dr ESPINAL CURULL, Xavier (CERN)

Presenter: MASCETTI, Luca (CERN)

Session Classification: Track 3 Session

Track Classification: Track3: Data store and access