21st International Conference on Computing in High Energy and Nuclear Physics **CHEP2015** Okinawa Japan: April 13 - 17, 2015

Contribution ID: **287** Type: **oral presentation**

Ceph-based storage services for Run2 and beyond

Thursday, April 16, 2015 11:45 AM (15 minutes)

In 2013, CERN IT evaluated then deployed a petabyte-scale Ceph cluster to support OpenStack use-cases in production. As of fall 2014, this cluster stores around 300 TB of data comprising more than a thousand VM images and a similar number of block device volumes. With more than a year of smooth operations, we will present our experience and tuning best-practices.

Beyond the cloud storage use-cases, we have been exploring Ceph-based services to satisfy the growing storage requirements during and after Run2. First, we have developed a Ceph back-end for CASTOR, allowing this service to deploy thin disk server nodes which act as gateways to Ceph; this feature marries the strong data archival and cataloging features of CASTOR with the resilient and performant Ceph subsystem for disk. Second, we have developed RADOSFS, a lightweight storage API which builds a POSIX-like filesystem on top of the Ceph object layer. When combined with Xrootd, RADOSFS can offer a scalable object interface compatible with our HEP data processing applications. Lastly the same object layer is being used to build a scalable and inexpensive NFS service for several user communities.

Primary author: VAN DER STER, Dan (CERN)

Co-authors: Mr PETERS, Andreas Joachim (CERN); ROUSSEAU, Herve (CERN); MASCETTI, Luca (CERN); LAMANNA, Massimo (CERN)

Presenter: Mr PETERS, Andreas Joachim (CERN)

Session Classification: Track 3 Session

Track Classification: Track3: Data store and access