

CERN's Ceph infrastructure: OpenStack, NFS, CVMFS, CASTOR, and more!

Monday, October 10, 2016 3:00 PM (15 minutes)

This work will present the status of Ceph-related operations and development within the CERN IT Storage Group: we summarise significant production experience at the petabyte scale as well as strategic developments to integrate with our core storage services. As our primary back-end for OpenStack Cinder and Glance, Ceph has provided reliable storage to thousands of VMs for more than 3 years; this functionality is used by the full range of IT services and experiment applications.

Ceph at the LHC scale (above 10's of PB) has required novel contributions both in the development and operational side. For this reason, we have performed scale testing in cooperation with the core Ceph team. This work has been incorporated into the latest Ceph releases and enables Ceph to operate with at least 7,200 OSDs (totaling 30 PB in our tests). CASTOR has been evolved with the possibility to use a Ceph cluster as extensible high-performance data pool. The main advantages of this solution are the drastic reduction of the operational load and the possibility to deliver high single-stream performances to efficiently drive the CASTOR tape infrastructure. Ceph is currently our laboratory to explore S3 usage in HEP and to evolve other infrastructure services.

In this paper, we will highlight our Ceph-based services, the NFS Filer and CVMFS, both of which use virtual machines and Ceph block devices at their core. We will then discuss the experience in running Ceph at LHC scale (most notably early results with Ceph-CASTOR).

Tertiary Keyword (Optional)

Cloud technologies

Primary Keyword (Mandatory)

Storage systems

Secondary Keyword (Optional)

Object stores

Primary author: VAN DER STER, Dan (CERN)

Co-author: ESPINAL CURULL, Xavier (CERN)

Presenter: ESPINAL CURULL, Xavier (CERN)

Session Classification: Track 4: Data Handling

Track Classification: Track 4: Data Handling