

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



Contribution ID: 24

Type: **oral presentation**

Operational Experience Running Hadoop XRootD Fallback

Monday, 13 April 2015 17:15 (15 minutes)

In April of 2014, the UCSD T2 Center deployed `hdfs-xrootd-fallback`, a UCSD-developed software system that interfaces Hadoop with XRootD to increase reliability of the Hadoop file system. The `hdfs-xrootd-fallback` system allows a site to depend less on local file replication and more on global replication provided by the XRootD federation to ensure data redundancy. Deploying the software has allowed us to reduce Hadoop replication on a significant subset of files in our cluster, freeing hundreds of terabytes in our local storage, and to recover HDFS blocks lost due to storage degradation. An overview of the architecture of the `hdfs-xrootd-fallback` system will be presented, as well as details of our experience operating the service over the past year.

Primary author: DOST, Jeffrey Michael (Univ. of California San Diego (US))

Co-authors: MRAK TADEL, Alja (Univ. of California San Diego (US)); WUERTHWEIN, Frank (Univ. of California San Diego (US)); TADEL, Matevz (Univ. of California San Diego (US))

Presenter: DOST, Jeffrey Michael (Univ. of California San Diego (US))

Session Classification: Track 3 Session

Track Classification: Track3: Data store and access