



Contribution ID: 50

Type: **not specified**

Using XRootD to Minimize Hadoop Replication

Wednesday, October 15, 2014 4:00 PM (30 minutes)

We have developed an XRootD extension to Hadoop at UCSD that allows a site to significantly free local storage space by taking advantage of the file redundancy already provided by the XRootD Federation. Rather than failing when a corrupt portion of a file is accessed, the `hdfs-xrootd-fallback` system retrieves the segment from another site using XRootD, thus serving the original file to the end user seamlessly. These XRootD-fetched blocks are then cached locally, so subsequent accesses to the same segment do not require wide area network access. A second process is responsible for comparing the fetched blocks with corrupt blocks in Hadoop, and injects the cached blocks back into the cluster. This on-demand healing allows a site admin to relax the file replication number, commonly required to ensure availability. The system has been put into production at the UCSDT2 since March of 2014, and we finished implementing the healing portion in September. The added resiliency of the `hdfs-xrootd-fallback` system has allowed us to free 236 TB in our local storage facility.

Primary author: DOST, Jeffrey (UCSD)

Presenter: DOST, Jeffrey (UCSD)

Session Classification: Storage and Filesystems

Track Classification: Storage & Filesystems