

The deployment of a large scale object store at the RAL Tier 1

Monday 10 October 2016 14:30 (15 minutes)

Since 2014, the RAL Tier 1 has been working on deploying a Ceph backed object store. The aim is to replace Castor for disk storage. This new service must be scalable to meet the data demands of the LHC to 2020 and beyond. As well as offering access protocols the LHC experiments currently use, it must also provide industry standard access protocols. In order to keep costs down the service must use erasure coding rather than replication to ensure data reliability. This paper will present details of the storage service setup, which has been named Echo, as well as the experience gained from running and upgrading it.

In October 2015 a pre-production service offering the S3 and Swift APIs was launched. This paper will present details of the setup as well as the testing that has been done. This includes the use of S3 as a backend for the CVMFS Stratum 1s, for writing ATLAS log files and for testing FTS transfers. Additionally throughput testing from local experiments based at RAL will be discussed.

While there is certainly interest from the LHC experiments regarding the S3 and Swift APIs, they are still currently dependant on the XrootD and GridFTP protocols. The RAL Tier 1 has therefore also been developing an XrootD and GridFTP plugin for Ceph. Both plugins are built on top of the same libraries that write striped data into Ceph and therefore data written by one protocol will be accessible by the other. In the long term we hope the LHC experiments will migrate to industry standard protocols, therefore these plugins will only provide the features needed by the LHC VOs. This paper will report on the development and testing of these plugins.

Tertiary Keyword (Optional)

Distributed data handling

Secondary Keyword (Optional)

Computing facilities

Primary Keyword (Mandatory)

Object stores

Author: DEWHURST, Alastair (STFC - Rutherford Appleton Lab. (GB))

Co-authors: CANNING, Bruno (RAL); RYALL, George (STFC); VASILAKAKOS, George (STFC); Mr JOHNSON, Ian (STFC - Rutherford Appleton Lab. (GB)); ADAMS, James (STFC RAL)

Presenter: DEWHURST, Alastair (STFC - Rutherford Appleton Lab. (GB))

Session Classification: Track 4: Data Handling

Track Classification: Track 4: Data Handling