Contribution ID: **595** Type: **Poster**

Optimising XRootD access to erasure coded object stores

Tuesday 10 July 2018 16:45 (15 minutes)

Since the start of 2017, the RAL Tier-1's Echo object store has been providing disk storage to the LHC experiments. Echo provides access via both the GridFTP and XRootD protocols. GridFTP is primarily used for WAN transfers between sites while XRootD is used for data analysis.

Object stores and those using erasure coding in particular are designed to efficiently serve entire objects which are normally assumed to be small (few MBs). Some experiment workflows fit this use case extremely well, however certain work flows have been shown to be extremely inefficient when run directly against Echo. To solve this problem XRootD caches were deployed. This paper describes the performance testing of various configurations of both disk and memory caches.

Access to Echo happens via gateways. Echo has a small number of dedicated gateway machines that provide external connectivity. Jobs running on the RAL batch farm can access Echo through an XRootD gateway running inside a container on every worker node. This paper describes the setup of these worker node gateways which also provide a cache and the ability to redirect requests in the event a file is not found.

For CMS, a dedicated XCache service has been provided to allow remote jobs to access data directly from Echo. This paper describes the setup and testing of this service.

Primary authors: LAHIFF, Andrew David (STFC-Rutherford Appleton Laboratory (GB)); DEWHURST, Alastair (STFC-Rutherford Appleton Laboratory (GB)); APPLEYARD, Rob (STFC); BYRNE, Tom (STFC)

Presenters: DEWHURST, Alastair (STFC-Rutherford Appleton Laboratory (GB)); APPLEYARD, Rob (STFC)

Session Classification: Posters

Track Classification: Track 4 - Data Handling