Contribution ID: **42** Type: **Oral**

Making the most of cloud storage - a toolkit for exploitation by WLCG experiments.

Monday, October 10, 2016 2:15 PM (15 minutes)

Understanding how cloud storage can be effectively used, either standalone or in support of its associated compute, is now an important consideration for WLCG.

We report on a suite of extensions to familiar tools targeted at enabling the integration of cloud object stores into traditional grid infrastructures and workflows. Notable updates include support for a number of object store flavours in FTS3, Davix and gfal2, including mitigations for lack of vector reads; the extension of Dynafed to operate as a bridge between grid and cloud domains; protocol translation in FTS3; the implementation of extensions to DPM (also implemented by the dCache project) to allow 3rd party transfers over HTTP.

The result is a toolkit which facilitates data movement and access between grid and cloud infrastructures, broadening the range of workflows suitable for cloud. We report on deployment scenarios and prototype experience, explaining how, for example, an Amazon S3 or Azure allocation can be exploited by grid workflows.

Tertiary Keyword (Optional)

Storage systems

Primary Keyword (Mandatory)

Cloud technologies

Secondary Keyword (Optional)

Object stores

Primary author: KEEBLE, Oliver (CERN)

Co-authors: ALVAREZ AYLLON, Alejandro (CERN); MANZI, Andrea (CERN); FURANO, Fabrizio (CERN); BITZES, Georgios (CERN); ARSUAGA RIOS, Maria (CERN)

Presenter: KEEBLE, Oliver (CERN)

Session Classification: Track 4: Data Handling

Track Classification: Track 4: Data Handling