Contribution ID: **144** Type: **Poster**

C3PO - A Dynamic Data Placement Agent for ATLAS Distributed Data Management

Tuesday 11 October 2016 16:30 (15 minutes)

This contribution introduces a new dynamic data placement agent for the ATLAS distributed data management system. This agent is
designed to pre-place potentially popular data to make it more widely available. It uses data from a variety of sources. Those
include input datasets and sites workload information from the ATLAS workload management system, network metrics from different
sources like FTS and PerfSonar, historical popularity data collected through a tracer mechanism and more. With this data it decides
if, when and where to place new replicas that then can be used by the WMS to distribute the workload more evenly over available
computing resources and then ultimately reduce job waiting times. The new replicas are created with a short lifetime that gets
extended, when the data is accessed and therefore the system behaves like a big cache.
This paper gives an overview of the architecture and the final implementation of this new agent. The paper also includes an
evaluation of different placement algorithms by comparing the transfer times and the new replica usage.

Primary Keyword (Mandatory)

Distributed data handling

Secondary Keyword (Optional)

Tertiary Keyword (Optional)

Author: GARONNE, Vincent (University of Oslo (NO))

Co-author: BEERMANN, Thomas (CERN)

Presenter: BEERMANN, Thomas (CERN)

Session Classification: Posters A / Break

Track Classification: Track 4: Data Handling