



Contribution ID: 324

Type: oral presentation

Federating LHCb datasets using the Dirac File Catalog

Monday, April 13, 2015 4:45 PM (15 minutes)

In the distributed computing model of LHCb the File Catalog (FC) is a central component that keeps track of each file and replica stored on the Grid. It is federating the LHCb data files in a logical namespace used by all LHCb applications. As a replica catalog, it is used for brokering jobs to sites where their input data is meant to be present, but also by jobs for finding alternative replicas if necessary.

The LCG File Catalog (LFC) used originally by LHCb and other experiments is now being retired and needs to be replaced. The DIRAC File Catalog (DFC) was developed within the framework of the DIRAC Project and presented during CHEP 2012. From the technical point of view, the code powering the DFC follows an Aspect oriented programming (AOP): each type of entity that is manipulated by the DFC (Users, Files, Replicas, etc) is treated as a separate 'concern' in the AOP terminology. Hence, the database schema can also be adapted to the needs of a Virtual Organization. LHCb opted for a highly tuned MySQL database, with optimized requests and stored procedures.

This paper will present the improvements brought to the DFC presented at CHEP 2012, its performance with respect to the LFC, as well as the migration procedure used to migrate the LHCb data from the LFC to the DFC. Finally it will show how a combination of the DFC and the LHCb framework Gaudi allow LHCb to build a data federation at low cost.

Primary author: HAEN, Christophe (CERN)

Co-authors: Dr TSAREGORODTSEV, Andrei (CPPM, Aix-Marseille Université, CNRS/IN2P3, Marseille, France); FRANK, Markus (CERN); CHARPENTIER, Philippe (CERN)

Presenter: HAEN, Christophe (CERN)

Session Classification: Track 3 Session

Track Classification: Track3: Data store and access