

Distributed Data Management on the petascale using heterogeneous grid infrastructures with DQ2

Wednesday, 13 February 2008 16:25 (25 minutes)

DQ2 is specifically designed to support the access and management of large scientific datasets produced by the ATLAS experiment using heterogeneous grid infrastructures. The DQ2 middleware manages those datasets with global services, local site services and enduser interfaces. The global services, or central catalogues, are responsible for the mapping of individual files onto DQ2 datasets. The local site services are responsible for tracking files available on-site, managing data movement and guaranteeing consistency of available data. The enduser interfaces provide users with the ability to query, manipulate and monitor datasets and its transfers. The distinction between global and local services is a core design decision as it clearly separates site-specific information, e.g. local site storage management, from global information. With this separation, any change within site infrastructures does not affect global reliability of the system and QoS requirements can be guaranteed.

3. Impact

Data movement is driven from the destination site using a unique pull-based subscription methodology. A user subscribes a dataset to a site and the system keeps track of all changes. The site services then fulfill the subscription by enacting the data movement in an intelligent and optimised way. The enacting layer relies on the EGEE gLite-FTS, glite-LFC, gLite-BDII, NorduGrid-RLS and OSG-LRC to interconnect the EGEE, NorduGrid and OSG infrastructures transparently. This allows scientists to work with all three grid infrastructures without specialised knowledge and eases the way they can store and access their data. The integration of all three grid infrastructures and the support for multiple grid storage systems (CASTOR, dCache, StoRM, DPM) is therefore one of the key points of the systems. The other key points are the systems proven scalability to the petascale, its non-invasiveness to existing services and its fault-tolerance to support heavily data-dependent sciences on the grid.

URL for further information:

<https://twiki.cern.ch/twiki/bin/view/Atlas/DistributedDataManagement>

4. Conclusions / Future plans

DQ2 is used within ATLAS, handling bookkeeping and data placement requests across large, medium and small computing centres worldwide. Large-scale dedicated tests are routinely run in preparation of live data-taking and DQ2 already manages millions of files with storage requirements in the petascale. Data movement peaked at stable 1.2 GB/sec for multiple days already and thus proved the systems scalability. Future plans involve optimising data placement, performance and enduser experience.

Provide a set of generic keywords that define your contribution (e.g. Data Management, Workflows, High Energy Physics)

Data Management, Petascale, Distributed Computing

1. Short overview

We describe Don Quijote 2 (DQ2), a new approach to the management of large scientific datasets by a dedicated middleware. This middleware is designed to handle the data organisation and data movement on the petascale for the High-Energy Physics Experiment ATLAS at CERN. DQ2 is able to maintain a well-defined quality of service in a scalable way, guarantees data consistency for the collaboration and bridges the gap between EGEE, OSG and NorduGrid infrastructures to enable true interoperability.

Primary authors: Mr LASSNIG, Mario (CERN & University of Innsbruck, Austria); Mr BRANCO, Miguel (CERN); Mr SALGADO, Pedro (CERN); Dr GARONNE, Vincent (CERN)

Presenter: Mr LASSNIG, Mario (CERN & University of Innsbruck, Austria)

Session Classification: Data Management

Track Classification: Existing or Prospective Grid Services