

Data Management Software Development for ATLAS

C. Serfon and M. Klinger

To deal with the huge amount of data which will be generated by the experiments at the LHC, computing grids have been set up. These grids allow us to split the computing resources as well as data over various computing centers. To manage the data, a software package called dq2 [1] has been developed within the ATLAS collaboration. This software allows us to transfer files between the different sites. However some important functionalities are currently missing in dq2. This concerns mainly the Data Management at the site level : data deletion, archiving, staging and ensuring consistency between files physically on sites and the information provided by the local catalogues (LFC or LCG File Catalog).

Therefore a set of tools has been developed at LMU to manage these tasks. These tools are grouped into a framework [2]. Although it has been primarily developed to manage operations in the GridKa cloud [3], it can be used at every site, that means it can be used at all sites independant of the Storage Element flavour (dCache, dpm, Castor).

The main features of this framework are :

- Deletion tool : This allows the clean-up of obsolete data from the Storage Elements and from the local catalogue. Various modes are available to clean data remotely/locally.
- LFC/Storage Element consistency check : One of the biggest problems is the consistency between the data stored on a given site and what is stored on the local catalogue. A consistency error can be a lost file, a corrupted file, etc. About 15 different types of error have been identified. For the most common ones (lost files, duplicated files) a correction procedure has been implemented.
- Monitoring tool : To monitor the data replication over the different sites (Figure 1).




Fig. 1: Number of files from a specific dataset which arrived at FZK vs time.

- Archiving tool : This tool is used to archive data on tape when it is not needed on disk pools anymore.
- Staging tool : This tool is used to send staging request for files on tape and was used to test the tape system at GridKa.

All these tools take advantage of local access which allows better efficiency (less failures because it uses low level commands) and increased speed compared to using grid-middleware tools. Figure 2 for instance shows the deletion rate for files at GridKa using the deletion tool. A rate of about 7-8 files deleted/second can be sustained over many hours, which is much faster than the deletion rate obtained using remote access (between 1 and 2 files/second).




Fig. 2: Number of files deleted per minute using the deletion tool. About 46000 files (representing more than 4 TB) have been deleted in less than 3 hours. The holes represent the interactions with the local catalogue (LFC).

These tools have been heavily used for months within the GridKa cloud and in other sites (CERN, France, Italy). The consistency tool in particular leads to the detection of a huge number of discrepancies (260 000 at this date) which have subsequently been corrected. Work is currently ongoing to include these tools in dq2.

References

- [1] ATLAS Distributed Data management Operations, D. Barberis *et al.*, ATL-SOFT-PUB-2006-006.
- [2] <https://twiki.cern.ch/twiki/bin/view/Atlas/DDMScriptsFrameWork>
- [3] C. Serfon *et al.* GridKa operations, Annual report 2006