

# Distributed Analysis Tools within the ATLAS Experiment

J. Elmsheuser, O. Biebel, G. Duckeck, J. Kennedy, G. Kroboth, B. Ruckert, T. Mahmoud, D. Schaile, M. Schott, and C. Serfon

The distributed data analysis using Grid resources is one of the fundamental applications in high energy physics to be addressed and realized in the near future. In every experiment up to a thousand physicist will be submitting analysis jobs into the Grid, namely LCG [1], the grid flavor developed especially for the large hadron collider LHC. Appropriate user interfaces and helper applications have to be made available to assure that all users can use the Grid without too much expertise in Grid technology. These tools enlarge the number of grid users from a few production administrators to potentially all participating physicists. There have been previous reports on these topics and here the focus will be put on the progress made during the recent months [2,3].

Within the D-Grid High Energy Physics Computing Grid Working Package 3 [4] we are working on distributed and interactive data analysis on the Grid. In this context we have been extensively contributing to the development and support of the the job and scheduling manager GANGA [5]. GANGA (Gaudi / Athena and Grid Alliance) is an interface to the Grid that is being developed jointly by the ATLAS and LHCb experiments. GANGA is a front end for job definition and management of analysis jobs to run in a distributed environment.

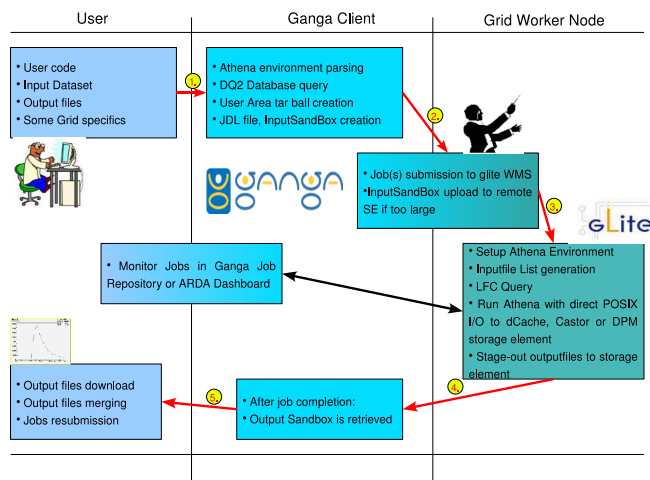


Fig. 1: Work-flow of a Grid job managed by the job scheduler GANGA.

Figure 1 shows the work-flow of a typical user job using the Athena and LCG plug-in in GANGA. GANGA acts as an abstraction layer between the user and the Grid. In a first step it bundles the necessary user code, configuration and input data-set information for the Grid submission. The job is sent via the resource broker to the Grid worker nodes where the input data is located. It can be split into several parallel sub-jobs to process a subset of the input data-set files. On the Grid worker node the Athena environment is setup, the user code is being compiled if necessary and the input data is directly accessed on the close storage element via POSIX I/O. After Athena completion the output files produced are stored on a close storage element. During Grid job execution the jobs are

monitored for their status in the GANGA job repository and the ARDA dashboard [6]. After the Grid job completion the log-files are automatically transferred back to the users local desktop. The output files can be downloaded and eventually merged if a job was split to several parallel sub-jobs.

We have successfully extended the functionality of GANGA in numerous areas. The most important is the integration and on-going improvement of the ATLAS data management system DDM/DQ2 with direct access to the input data files via POSIX I/O. A new service offered by the DDM/DQ2 team to manage incomplete replicated data-sets has been integrated. This service called site-index/tracker helps to fully process all files of a data-set across many sites. Incompletely replicated data-sets at various grid sites is a common problem. This new tool helps to automatically generate jobs based on the file locations and not only based on the data-set locations.




Fig. 2: Number of GANGA users per month since January 2007.

GANGA has been used at least once by over 1300 users in total and over 800 ATLAS user since the beginning of 2007. There are over 60 users in the ATLAS experiment per week (see Figure 2). Over half of the users are using GANGA at CERN and all other users are almost evenly spread among other sites. Several tutorials have been held during 2007 with an attendance of 30-40 people at each occasion to teach physicists how to use the distributed analysis tools. During July until September 2007 Grid jobs sent out by GANGA to the LCG Grid have been running with 8% at CERN and 37% at Tier-1 sites where the majority of the Tier-1 jobs ran at Lyon and GridKa. After the successful data-set replication among most of the Tier-2 sites now 40% of the Grid jobs have been running at Tier-2 and 15% at Tier-3 sites, respectively, following the ATLAS user analysis computing model.

## References

- [1] LCG project web page: <http://lcg.web.cern.ch/LCG/>
- [2] J. Elmsheuser *et al.*, Annual report 2005, p. 99
- [3] J. Elmsheuser *et al.*, Annual report 2006, p. 95
- [4] D-Grid HEP/CG project web page: <https://www.d-grid.de/>
- [5] GANGA project web page: <http://ganga.web.cern.ch/ganga/>
- [6] ARDA dashboard web page: <http://dashboard.cern.ch/atlas/>