

Distributed Analysis in the ATLAS Experiment using Ganga

J. Elmsheuser, P. Calfayan, J. de Graat, G. Duckeck, J. Ebke, C. Galea, J. Kennedy, C. Kummer, C. Mitterer, D. Schaile, C. Serfon, and R. Walker

The distributed data analysis using Grid resources is one of the fundamental applications in high energy physics that has been addressed and realized for the start of LHC data taking. The needs to manage the resources are very high. In every experiment up to a thousand physicist will be submitting analysis jobs into the Grid. Appropriate user interfaces and helper applications have been made available to assure that all users can use the Grid without expertise in Grid technology. These tools enlarge the number of grid users from a few production administrators to potentially all participating physicists.

The Ganga job management system [1,2] has been developed as a common project between the ATLAS and LHCb experiments. Ganga provides a simple and consistent way of preparing, organizing and executing analysis tasks within the experiment analysis framework Athena, implemented through a plug-in system. It allows trivial switching between running test jobs on a local batch system and running large-scale analyzes on the Grid, hiding Grid technicalities from the users.

Support for all grids presently used by ATLAS, namely the LCG/EGEE, NDNF/Nordu-Grid, and OSG/Panda is provided. The integration and interaction with the ATLAS data management system DQ2 into Ganga is a key functionality. An intelligent job brokering is setup by using the job splitting mechanism together with data-set and file location knowledge.

The brokering uses status information from an automated system called GangaRobot [3] that regularly processes test analysis jobs at all ATLAS DQ2 supported sites. These functional tests are using different versions of the experiment software and different input file access modes like Posix I/O and a new background thread copy mode from the storage system to the worker node. In Ganga large amounts of analysis jobs can be sent to the locations of data following the ATLAS computing model, which foresees a centrally organized distribution of the data by DQ2 and the rule that analysis jobs are sent by the distributed analysis tools to the data.

Large storms of analysis jobs allow to study the detailed behavior of a test and are used for the site commissioning. Analysis jobs demand a high input rate from the local storage systems contrary to the CPU intensive Monte Carlo simulation jobs. The test results vary among the sites depending on the individual hardware and network setups and the distribution of the input files on the storage system. An example test result is shown in Figure 1.

Ganga supports amongst other things tasks of user analysis with reconstructed data and small scale production of Monte Carlo data. Jobs can be monitored with the Ganga job repository or the ARDA Dashboard for job in the LCG/EGEE grid. Submission to the Panda system is realized by the integration of the Panda client tools into the Ganga framework. This shared code is planned to ease the maintenance of the system and allows for an easy submission option for the users.




Fig. 1: Test results of a HammerCloud test: Fraction of succeeded job (left) job CPU utilization and event rate (right).

The number of daily distributed analysis jobs is at present not yet at the expected level of the LHC data taking period. Therefore all grid sites are tested with an auto-

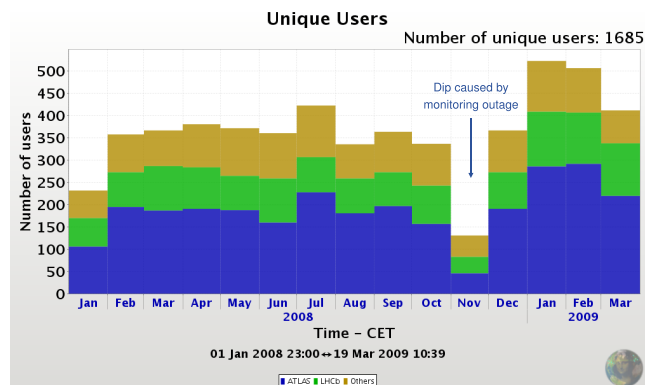


Fig. 2: Number of Ganga users per month since January 2008.

Ganga is part of the regularly scheduled CERN Analysis tutorials and has been presented and exercised in many tutorial and workshops over the last 3 years [5]. There have been over 600 people from ATLAS using Ganga in the last 6 month and since the beginning of the year 2009 there are per day 10,000-15,000 finished jobs on the LCG back-end alone. (see Figure 2).

The distributed analysis using Ganga works well at a large number of daily tested sites. Standard job configurations of user jobs are successfully automatically processed. Still the site availability and the stability of the input file access at a number of sites has to be improved to be ready for the full scale of distributed analysis expected during the LHC data taking period.

References

- [1] Ganga project web page: <http://ganga.web.cern.ch/ganga/>
- [2] F. Brochu *et al.*, Ganga: a tool for computational-task management and easy access to Grid resources, submitted to CPC, arXiv:0902.2685v1
- [3] GangaRobot overview page: <http://gangarobot.cern.ch>
- [4] HammerCloud overview page: <http://gangarobot.cern.ch/st>
- [5] Ganga Extended Tutorial: <https://twiki.cern.ch/twiki/bin/view/Atlas/GangaTutorial5>