LARGE SCALE PRODUCTION ON THE LCG SEEN BY HEPHY-UIBK

Reinhard Bischof*, Peter Oettl*, Dietrich Liko† and Gregor Mair†

Abstract. *The Large Hadron Collider at CERN will start data acquisition in 2007. The ATLAS (A Toroidal LHC ApparatuS) experiment is preparing for the data handling and analysis via a series of Data Challenges and production exercises to validate its computing model and to provide useful samples of data for detector and physics studies.*

In 2002 a series of production exercises (so-called data challenges) was planned to validate the ATLAS Computing Model and to ensure the correctness of the technical choices. The HEPHY group in Innsbruck set up a LCG Grid site in 2003 and participated in Data Challenge 2 and large scale production for the physics workshop in Rome in 2005.

In this paper we discuss the ATLAS Production System, the infrastructure of the site in Innsbruck and the experiences of the previous Data Challenges, together with the lessons learned and future plans.

1. Introduction

ATLAS is one of the four experiments at the Large Hadron Collider [2] (LHC) located at CERN. Both ATLAS and the LHC are in their construction phases and are expected to become operational in the course of 2007.

According to the ATLAS Computing Model [1] the expected volume of data recorded for off-line reconstruction and analysis will be of the order of 10 PB per year and about 1 million jobs are going to be processed per day. To process this large amount of data resources from three different Grid infrastructures, LCG [2] (LHC Computing Grid), OSG/GRID3 [3] (Open Science Grid) and NorduGrid [4], will be used.

In 2002 a series of production exercises (so-called data challenges) was planned to validate the ATLAS Computing Model and to ensure the correctness of the technical choices. While in DC1 only NorduGrid and Batch systems were used, subsequent exercises were fully Grid based.

The HEPHY group in Innsbruck set up a LCG Grid site in 2003 and participated in Data Challenge 2 and large scale production for the physics workshop in Rome in 2005. Sections 2 explains the requirements of the ATLAS experiment more detailed. Section 3 discusses the architecture of the

*HEPHY, University of Innsbruck, Austria
†CERN, European Organization for Nuclear Research, Switzerland

ATLAS Production System and section 4 relevant aspects of the Grid middleware. In section 5 the organizational structure of Grid facilities within the experiment is explained. Section 6 summarizes lessons learned from using different Grids and section 7 shows the production at HEPHY-UIBK.

2. Requirements

When the LHC will be put in operation in 2007 the ATLAS experiment will produce about 10 PB of data per year. Furthermore it will be necessary to process about 1 million jobs per day.

Using Grid sites with different Grid flavors all around the world to be able to handle such a vast amount of jobs and data requires not only a large number of Grid sites but also a strong environment to manage such a system.

The main issues will be to manage the huge amount of data and to address scalability issues with the large number of jobs on limited resources.

The requirements are in detail:

- Data Management:
  At the present time it is neither possible nor wanted to store all data in one place. To be available on demand and to avoid single points of failure the data is replicated to several Grid sites. Therefore the ATLAS data management system needs to be able to store, locate and replicate data over the Grid. In general, all the data is stored on tape and, if required, staged to disk. To avoid unnecessary network traffic ATLAS prefers to send jobs to the data. It is also important to have the data already available on disk when the job arrives, otherwise Grid resources may be blocked as jobs are waiting for their input.

- Analysis:
  The fast majority of jobs (applications) within ATLAS will be analysis jobs, i.e. jobs submitted by user-groups to analyze and evaluate data taken from the detector. These jobs are I/O intensive and have considerable short run-times. The focus is on the scalability of the underlying system and on short turn-around times. Hence, short or preemptive queues or a powerful and scalable scheduling mechanism are needed.

- Production Jobs:
  About 100.000 jobs are expected each day for Monte Carlo production. These jobs are very CPU intensive and are scheduled to run about 20 hours each. Compared to the analysis jobs short turn-around times are obviously not an issue, but due to their long run-times they tend to occupy resources quite long. Scheduling systems need to be aware that within these long run-times small jobs may be submitted and have to take care that these small jobs are moved forward in the queues.

3. The ATLAS Production System

The use of several Grid sites with different Grid flavors at the same time requires an environment which provides capabilities to run jobs on any site. The ATLAS production system (ProdSys) addresses this issue by layered architecture. Jobs defined in a database are picked up by a supervisor and are handed to available executors. An executor then wraps the job for a specific Grid flavor. Figure 1 illustrates this architecture.

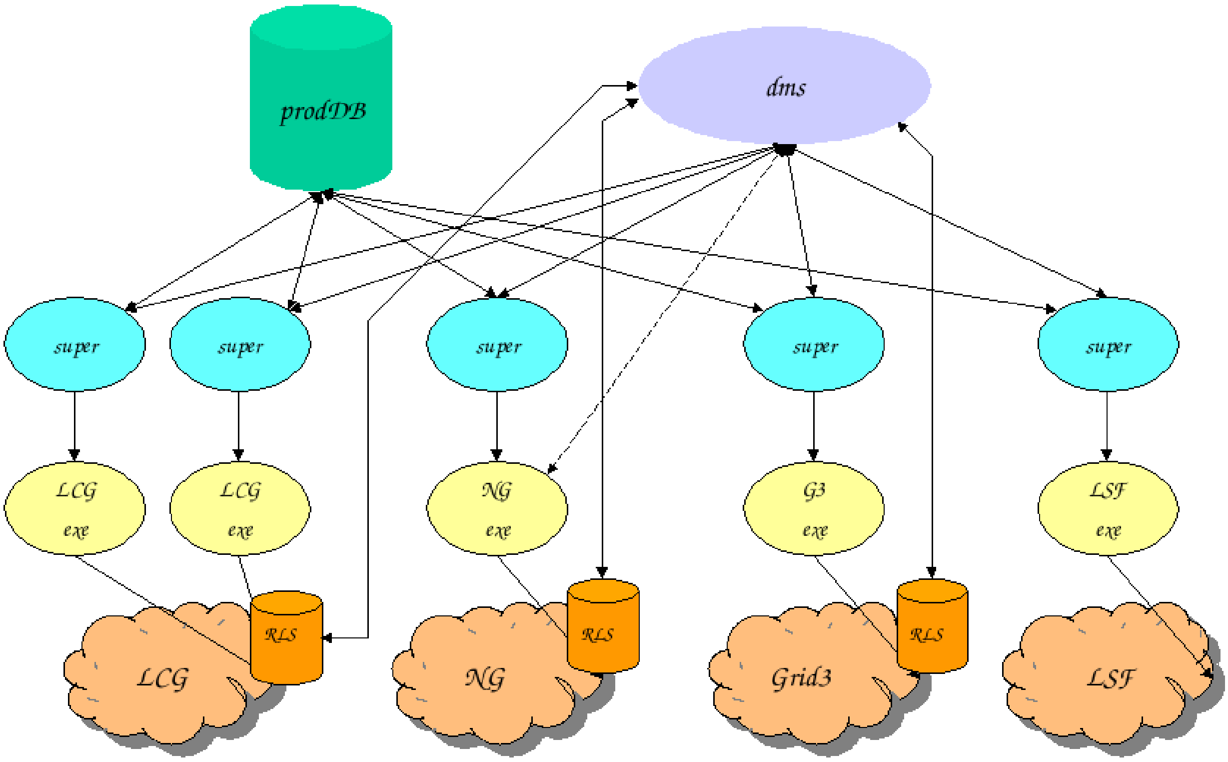

Figure 1. Architecture of ProdSys

The components of ProdSys are the following:

- Supervisor (super):
  The supervisor [5] is a second generation supervisor implementation. It picks up jobs which are defined in a database (see below) and creates Python objects which are handed to available executors.

  When a job has terminated, the supervisor validates the job. Valid output is registered in the data management system. In case that an error occurred the job may be resubmitted to the same or to another executor or it may be declared as failed.

- Executors (exe):
  Executors implement the layer between the supervisor(s) and the Grid sites. Each Grid flavor requires a specific executor (Lexor [6] and Condor-G [6] for LCG, Panda [7] for OSG and Dulcinea [8] for NorduGrid) which is able to steer the jobs, i.e. an executor implements the Grid flavor specific API to submit, delete or query jobs.

- Data Management System (dms):
  The data management system (currently Don Quijote 2) is responsible for data archiving, data allocation and data movement. Data which has been validated by the supervisor is transferred to its final destination, if necessary, and is registered in replica catalogs.

- Database (prodDB):
  A plug-in mechanism allows to connect to databases of different vendors. The used schema can be found at https://uimon.cern.ch/twiki/bin/view/Atlas/ProdDB. Currently Oracle database servers and for Monte Carlo production an additional meta-data database are in use.

4. Relevant Grid Middleware Features

Grid middleware comes along with a variety of concepts and tools to manage Grid operations. ATLAS tries to make use of as much of these as possible.

- Workload Management System
  The WMS is responsible for the management and monitoring of jobs submitted from a User Interface (UI). A set of services running on the Resource Broker (RB) machine contribute to match job requirements to the available resources (as gathered form the IS), schedule the job for execution to an appropriate Computing Element (CE), track the job status, and allow users to retrieve their job output. The Logging and Bookkeeping (LB) service keeps the state information of a job and allows the user to query its status. Each CE is the front-end to a local batch system managing a pool of Worker Nodes (WN) where the job is eventually executed. User Credentials of limited lifetime (Proxies) can be automatically renewed through a Proxy Service.

- Data Management System
  The DMS allows users to move files in and out of the Grid, to replicate them among different Storage Elements (SE), and to locate them. A number of available and supported protocols (Globus GridFTP is the most commonly used) are employed for data transfer. A central file catalogue, the Replica Location Service (EDG-RLS), keeps information about file location and about some file metadata.

- Information System
  The IS provides information about the Grid resources and their status. The information, generated by an Information Provider, are published by the Grid Resource Information Service (GRIS) running on each individual resource and propagated into a hierarchical database structure (Grid Information Index Services GIIS at every site and the Berkeley Database Information Index BDII as central collector). The GRIS is based on the Globus Monitoring and Directory Service (MDS). Information are published following a specific schema that goes under the name of Glue [9].

- Monitoring
  The GridICE [10] and R-GMA [11] services allow for application and resource status monitoring. In particular, GridICE collects information from the IS and from sensors (agents) deployed on almost all Grid machines. Such information are stored in a database which allows, via web interface, for retrieving the real time status and historical events about jobs and services. Different views focus on different aspects of job and resources monitoring. R-GMA (not used by ATLAS so far) is based on a producer-consumer model. Beside job monitoring information, R-GMA allows for application monitoring offering a variety of APIs in many programming languages. Information can be easily retrieved through SQL queries using command line tools, APIs or a web interface.

5. Facility Characterization

While the ATLAS computing model is very much Grid-based, there still remain distinct roles for different facilities. Data produced at CERN will be preprocessed at CERN Grid sites (Tier-0). The Tier-0 facility is responsible for the archival and distribution of the primary raw data received. It provides prompt reconstruction and first-pass processing of the data. The derived datasets are then distributed to Tier-1 facilities.

Remote sites are categorized either as Tier-1 or Tier-2 sites. Tier-1 sites take responsibility to host and provide long-term access and archival of a subset of the raw data. They need to provide a large number of computing elements (CE) to perform the reprocessing of the raw data.

Tier-2 facilities may take a range of significant roles in ATLAS such as providing calibration constants, running simulation and analysis. This range of roles will result in different sizes of the facilities. Tier-2 facilities also provide analysis capacity for physics working groups and subgroups. Finally, Tier-3 facilities provide access to the Grid for local user groups.

In the following the requirements and responsibilities of theses facilities are explained in detail (according to the TDR).

- Tier-0
  The Tier-0 facility at CERN is responsible for the archiving and distribution of the primary RAW data received from the Event Filter. It provides the prompt reconstruction of the calibration and express streams and the somewhat slower first-pass processing of the primary event stream. The derived datasets (ESD, primary AOD and TAG sets) are distributed from the Tier-0 to the Tier-1 facilities, and the reconstructed calibration data to the CERN Analysis Facility for data-intensive calibration. More automated calibration tasks will also be run by the Tier-0. In the event of prolonged down-time, first-pass processing and calibration must be taken over by the Tier-1 facilities. To account for failures and network outages, a disk buffer corresponding to about 5 days of data production will be required for the data flowing into the Tier-1.

- Tier-1
  Approximately 10 Tier-1 facilities are planned world-wide that will serve ATLAS. They take responsibility to host and provide long-term access and archiving of a subset of the RAW data (on average 1/10th each). They also undertake to provide the capacity to perform the reprocessing of the RAW data under their curation, and to provide ATLAS-wide access to the derived datasets, with the most up-to-date version of the data available with short latency (‘on disk’) and the previous version available but perhaps with a longer latency (‘on tape’). All of the datasets hosted are considered to be for the collaboration as a whole, and the storage and CPU pledged to be funded by the Tier-1 for that purpose.

  The Tier-1s must allow access to and provide capacity to analyse all of the hosted samples, and will provide part of the calibration processing capacity. Modest RAW data samples must be available at short latency to allow calibration and algorithmic development.

  Tier-1 facilities are expected to have a high level of service in terms of availability and response time. Given the vital role in receiving the raw data and reprocessing, down-times in excess of 12 hours become problematic in terms of catching up with processing and with the storage elsewhere of RAW data.

- Tier-2

  Tier-2 facilities may take a range of significant roles in ATLAS. This range of roles will result in different sizes of the facilities. Tier-2 facilities also provide analysis capacity for physics working groups and subgroups. This analysis activity is generally chaotic in nature. They typically will host one third of the available current primary AOD and the full TAG samples. In addition, they will provide all of the required simulation capacity for the experiment. Agreements on the primary host for the data from a given Tier-2 will be negotiated, although some flexibility will be required in the case of access problems. The relationships formed will be influenced by the ATLAS organizational plans and by the networking topology available. The primary host arrangement will help the planning of network links and may well follow the arrangements within a region for Grid operations and user support.

  The Tier-2s will also host modest samples of RAW and ESD data for code development. Some Tier-2s may take significant role in calibration following the local detector interests and involvements. In the case that larger samples need to be processed, it is planned that this would be done on the Tier-1 facilities by someone with production rights and with the agreement of the appropriate working group.

  The level of service in terms of availability and response time expected of a Tier-2 is lower than for a Tier-1 (unless it chooses to host the simulated data it generates). In principle, all members of the ATLAS VO have access to a given Tier-2. In practice (and for operational optimization), heightened access to CPU and resources may be given to specific working groups at a particular site.

- Tier-3

  Tier-3 facilities (which may be collections of desktops machines or local institute clusters) should be Grid-enabled to allow job submission and retrieval from the Grid.

  The size of Tier-3 resources will depend on the local user community size and other factors, such as any specific software development or analysis activity foreseen in a given institute, and are therefore neither centrally planned nor controlled.

6. Lessons learned from using different Grids

The ATLAS ProdSys needs reliable Grid middleware and an appropriate channel to gather information from the underlying infrastructure. Hence, improvement can be achieved as the consequence of improvements in Grid services reliability, Grid operations and in the ATLAS specific components. Still open issues are: the output files are randomly distributed around the world and not logically clustered in well defined storage locations; the production needs a substantial personnel effort, cannot run unattended and several interventions are needed. Several components of the Grid middleware and the production system still need improvements, both in terms of reliability and performance. All those issues were reported to and addressed by the middleware providers. New releases of the Grid software are now in the pre-production phase, preparing for the next massive production at the end of 2005. In general, the integration of the ProdSys framework and Grid middleware is a process under continuous evolution. Grid developers should consider the heterogeneous variety of the ProdSys, with different perspective, requirements and views of the Grid. Middleware development for such a variety of communities can be effectively carried on through the definition of generic and standard interfaces, agreed between ProdSys representatives and Grid developers, as for the SRM interface. On the other side, application developers should remember the dynamic and distributed nature of a Grid environment. Applications must be able to adapt to different runtime environments and to tolerate failures.

7. Production at HEPHY-UIBK

In 2003 the High Energy Physics group at the University of Innsbruck (collaborating with ATLAS) set up a LCG-site as Austrian Tier-2 prototype. Currently HEPHY-UIBK consists of 14 Worker Nodes (WN, double Xeon machines (3.0GHz) with 2 GB RAM), 1 Computing Element (CE), 1 Storage Element (SE) with 1.7 TB disk storage and a User Interface (UI). The site supports the virtual organizations ATLAS and dteam and is managed by three administrators (in total 1 FTE). The operating system for all nodes is Scientific Linux CERN (SLC), middleware updates are done using the script based YAIM provided by LCG. The installation and validation of the ATLAS specific software is done by the ATLAS software manager team via grid jobs.

In 2004 HEPHY-UIBK participated in the Atlas Data Challenge 2 (DC2) (260000 jobs) right from the beginning in June 2004, starting with short event generation jobs (approx 30 min/job), which produced the input files (600 MB, 10000 events) for the following detector simulation jobs. In total 733 simulation jobs were done in Innsbruck consuming 11890 CPU-hours. Figures 2 and 3 show the distribution of the used cpu time and memory for all jobs (event generation, simulation, digitization, reconstruction and (only in 2005) pile up).

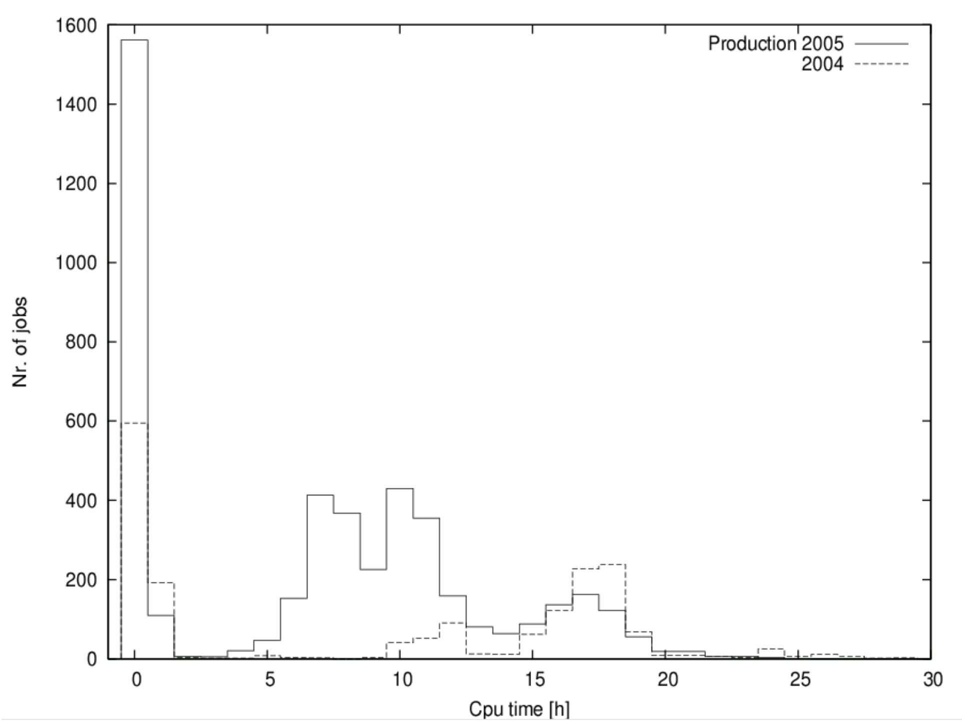

Figure 2. Distribution of cpu time

Input files for the simulation jobs were replicated to the local Storage Element (one file contains the events for several simulation jobs) and copied to worker nodes as needed. In figure 4 the first box (at) shows the file transfers to the local nodes, the replicated event files were copied several times for different simulation jobs. Most of the output files of the simulation jobs (approx. 300 MB each) were copied directly from the worker nodes to mass storage SE's in the LCG. It would be preferable to

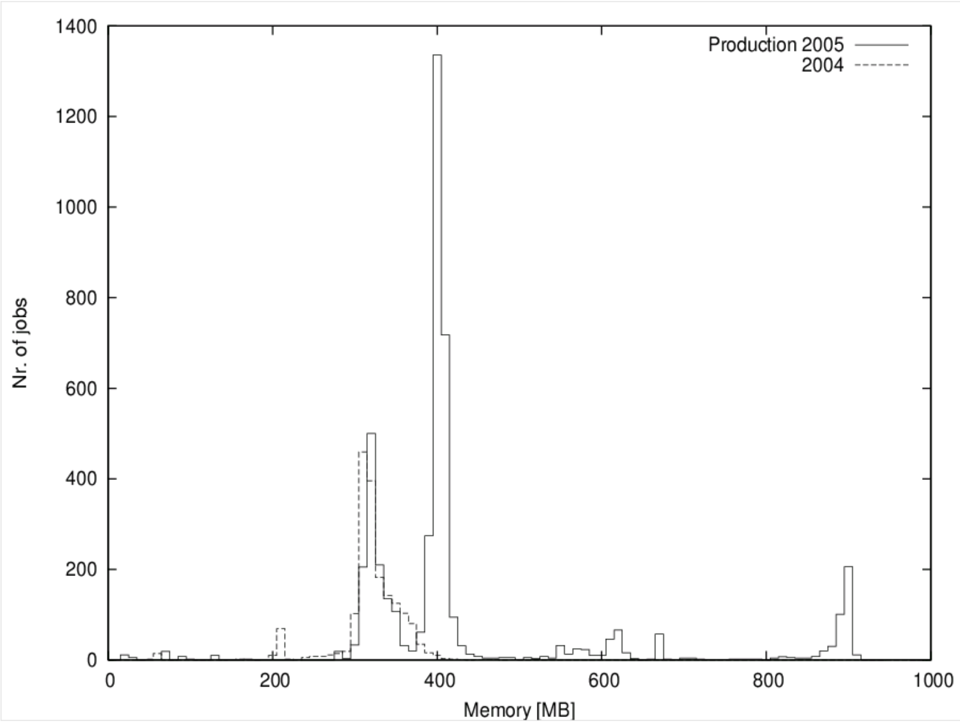

Figure 3. Distribution of used RAM

copy output files to the local SE in order not to block worker nodes with file transfers. Several times the transfer time either for output or input files exceeded the wallclock time limit of the atlas queue (72 h), especially for the digitization jobs, because a lot of jobs in the grid were requesting input files within a short time interval, thus overloading the SE in Barcelona at the beginning of the digitization phase.

In 2005 HEPHY-UIBK participated in the production for the physics workshop in Rome (573315 jobs in LCG, Nordugrid, grid3) with a slightly higher part in the production as in 2004, details are shown in figure 2 and 3.

Several short downtimes of HEPHY-UIBK during the productions were necessary for the middleware updates, two downtimes were caused by problems with the raid array of the SE. Overall the cluster was running stable with rare configuration or software problems that were fixed in a reasonable time.

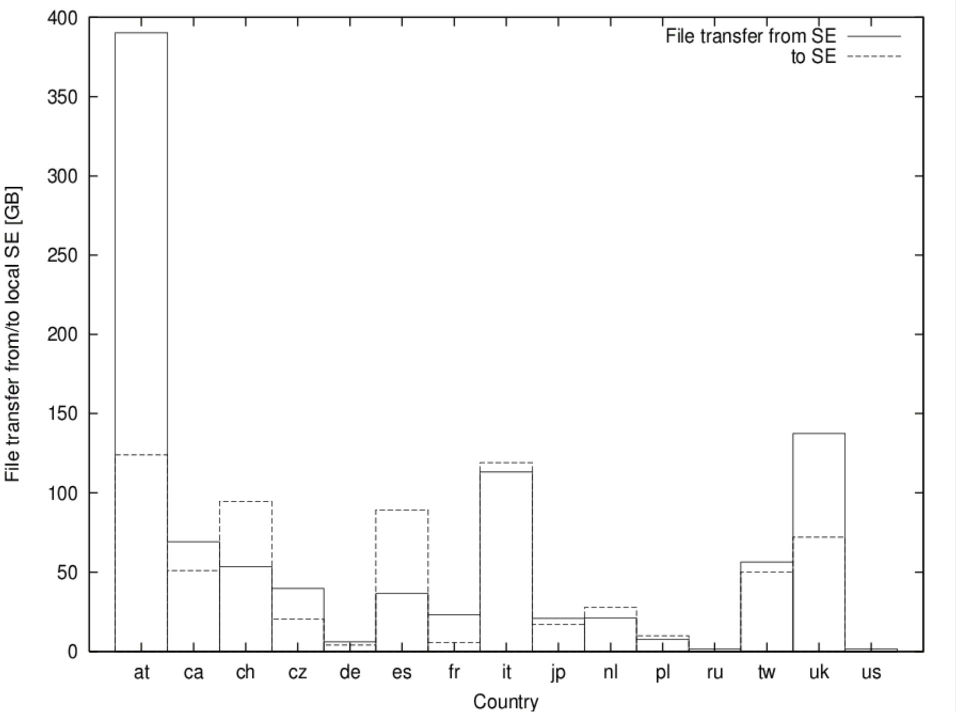

Figure 4. File transfer to local SE sorted by countries (box at shows transfer to local worker nodes)

Conclusion

The Austrian Tier-2 prototype, HEPHY-UIBK, has successfully participated in the large scale productions of ATLAS on the grid in the last two years. The part in the production was smaller than expected from a Tier-2 due to the small amount of hardware resources. Assuming that there will be 40 Tier-2s in ATLAS, the mean value for Tier-2 requirements in 2008 should be 500 kSI2k and 220 TB. For the Austrian High Energy Physics groups in Innsbruck (ATLAS) and Vienna (CMS) it is of great interest to establish an adequate Austrian Tier-2 supporting ATLAS and CMS.

Acknowledgments

We would like to thank the members of the HEPHY group at the University of Innsbruck for their important work in Innsbruck; especially Dietmar Kuhn who pushed us to have this paper realized. Furthermore many thanks to the ATLAS collaboration and project associates, especially Luc Goossens, Armin Nairz and Santiago Gonzalez de la Hoz, for their help to gather all the information and their useful feedback.

References

[1] The ATLAS Computing Group, "ATLAS Computing Technical Design Report", *ATLAS–TDR–017 (CERN-LHCC-2005-022) (2005)*

[2] The LCG Editorial Board, "LHC Computing Grid Design Report", *LCG–TDR–001 (CERN-LHCC-2005-024) (2005)*

[3] Open Science Grid (OSG) Web Page, http://www.opensciencegrid.org/

[4] Nordic Data Grid Facility (NDGF) Web Page, http://www.ndgf.org/

[5] EOWYN Web Page, https://uimon.cern.ch/twiki/bin/view/Atlas/Eowyn

[6] LCG Executors Web Page, https://uimon.cern.ch/twiki/bin/view/Atlas/ExecutorsCommon

[7] Panda Web Page, https://uimon.cern.ch/twiki/bin/view/Atlas/Panda

[8] Dulcinea Web Page, http://www.nordugrid.org/applications/dulcinea.html

[9] S. Andreozzi and M. Sgaravatto and C. Vistoli, "Sharing a conceptual model of Grid resources and services", *Conference on Computing in High Energy and Nuclear Physics (CHEP 2003)*

[10] GridICE Homepage, http://grid.infn.it/index.php?gridice0

[11] R-GMA Homepage, http://www.r-gma.org/