EXPERIENCE ON LARGE SCALE PRODUCTION ON THE GRID

G. Poulard, CERN, Geneva, Switzerland
For the ATLAS Data Challenges, Grid and Operations teams.

Abstract

The ATLAS Collaboration is preparing for data taking and analysis at the CERN LHC, scheduled to start operating in 2007. As part of this preparation, a series of computing Data Challenges (DCs) was started in 2002, with the goal of validating the Computing Model, the complete software suite, and the data model, and to ensure the correctness of the technical choices to be made for the final offline computing environment. DC1 was conducted during 2002-03; the main goals achieved were to set up the production infrastructure in a real worldwide collaborative effort, and to gain experience in exercising an ATLAS wide production model. DC2, run from May until December 2004, was a similar exercise but deployed a new production system able to submit jobs on three different Grids: LCG, OSG/GRID3 and ARC/NorduGrid. Right after DC2, a new Monte Carlo production was necessary to produce the needed event samples for an ATLAS Physics workshop (called Rome workshop) taking place in June 2005. In this paper we present our experience with these large-scale productions conducted on the three aforementioned Grids.

INTRODUCTION

In 2007 the Large Hadron Collider (LHC) [1] is due to come into service at the European Particle Physics Laboratory (CERN) in Geneva. In the LHC two proton beams, each of energy 7 TeV, are steered to collide head-on at the centre of large complex detectors. These collisions are expected to reveal fundamental new processes in particle physics.

The **ATLAS** (A Toroidal LHC ApparatuS) detector [2] is one of several that are being constructed to exploit the LHC. According to its computing model, after reduction of the data by the online trigger processor farms, the expected volume of data recorded for offline reconstruction and analysis will be of the order of few Petabytes (10^{15} bytes) per year to be analysed by institutes around the world. To ensure that the required computing resources are available, the experiment will need a worldwide distributed data management and computing system, on a scale that is one to two orders of magnitude larger than previous experiments in particle physics.

To prepare for this task, the LHC Computing Review [3] in 2001 recommended that the LHC experiments should carry out **Data Challenges** (DCs) of increasing size and complexity. The goals of the ATLAS Data Challenges are the validation of the ATLAS computing model, of the complete software suite, of the data model, and to ensure the correctness of the technical computing choices to be made.

The ATLAS collaboration decided to perform these DCs in the context of the LHC Computing Grid project LCG [4], to which ATLAS is committed, but also to use both the middleware and the resources of two other Grid projects, OSG/GRID3 [5] and ARC/NorduGrid [6]. The task of the LCG project is to prepare the computing infrastructure for the simulation, processing and analysis of the LHC data for all four LHC collaborations. The main emphasis of the LCG project is the deployment of Grid technologies for LHC computing. Both OSG/GRID3 and ARC/NorduGrid have similar approaches using the same foundations (GLOBUS) as LCG but with somewhat different middleware suites.

ATLAS Data Challenge 1, in 2002-3, was run on conventional non-Grid infrastructure in Europe, Asia and North America and fully on Grid in the Nordic and associated countries. Some early testing of the Grid infrastructures was done on LCG and Grid3.

Data Challenge 2, in the second half of 2004, and the Rome production, in 2005, were completely run on the Grid infrastructure.

ATLAS DC1 is fully described, with emphasis on the computing aspects, in an internal CERN report [7]. In the following sections are described some general considerations on the scope and the scale of the DC2 and Rome productions; the ATLAS Production System; some statistical information and general observations on DC2 and Rome productions; the main lessons and suggestions for improvements or future developments and, finally some conclusions and outlook.

DC SCOPE AND SCALE

The design and construction of an experiment like ATLAS requires a large amount of simulated data in order to optimise the design of the detectors, to estimate the physics performance that will be achieved, and to test the software and computing infrastructure. Samples of simulated data consist of a large number ($\sim 10^7$) of simulated "events". Each "event" is the result of a collision between two protons, and the full simulation requires the following steps:

- "Event Generation": the final state configurations of proton-proton collisions are generated using programs relying on theoretical calculations, phenomenological models and experimental inputs.

- “Detector Simulation”: the interaction of the generated particles inside the detector is simulated, taking into account the real geometry and distribution of the detector material.
- “Digitization”: the detector response is derived from the particle interactions and it is written in a format compatible with the real output of the detector. In addition, because of the high rate of collisions at LHC, digitized signals from several simulated events can be piled-up to create samples with a realistic experimental background. The digitized events (with or without pile-up) can now be used to test the software suite that will be used on real LHC data.
- “Reconstruction”: particle trajectories and energies from the detector data are reconstructed.

This requires running a chain of different programs with different characteristics in terms of memory usage and CPU time consumption. Typically a simulation (long) job runs for ~24 hours while a digitization or reconstruction (short) job runs for 3 to 4 hours.

ATLAS DC2 PRODUCTION SYSTEM

In order to handle the task of ATLAS DC2 an automated production system (ProdSys) [8] was designed. All jobs were defined and stored in a central database (prodDB). A supervisor agent (Windmill) [9] picked them up, and sent their definition as XML messages to various executors, via a Jabber server. Executors are specialised agents, able to convert the XML job description into a Grid-specific language (e.g. JDL, job description language, for LCG). Four executors were developed, for LCG (Lexor) [10], NorduGrid (Dulcinea) [11], Grid3 (Capone) [12] and legacy systems [13], allowing the Data Challenge to be run on different Grids.

When a LCG job was retrieved by Lexor, it built the corresponding JDL description, created some scripts for data staging, and sent everything to a dedicated, but otherwise standard, Resource Broker (RB) through a Python module built over the workload management system (WMS) API. The requirements specified in the JDL let the RB choose a site where ATLAS software was present and the requested amount of computation (expressed in SpecInt2000*Time) was available. An extra requirement was a good outbound connectivity, necessary for data staging.

The actual executable was wrapped in a script that performed various tasks: check the ATLAS software installation on the worker nodes; download and install the pacman package [14] for the required application; set up the ATLAS environment; stage in the input files, perform the transformation and stage out the results.

For data management, a central server, Don Quijote (DQ) [15] offered a uniform layer over the different replica catalogues of the 3 Grid flavours. Thus all the copy and registration operations were performed through calls to DQ. The software was installed using LCG tools.

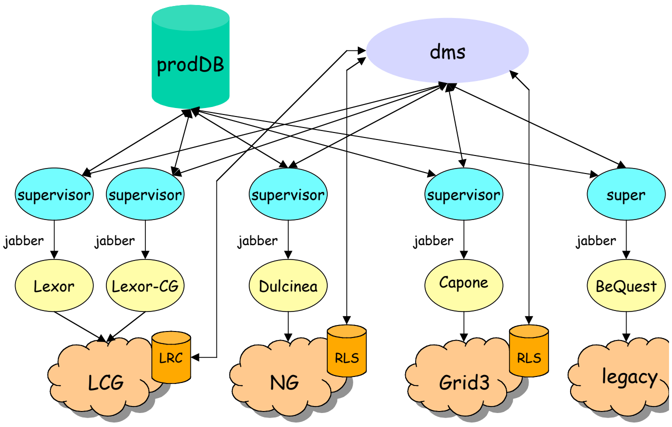

Figure 1: The ATLAS production system architecture.

ACHIEVED PRODUCTIONS

DC2 started in July 2004 and lasted until December, immediately followed by the Rome production that lasted until May 2005.

For DC2 in total 10 million events were processed in ~260 thousand jobs, consuming ~200 kSI2k years of CPU and producing ~60 TB of data; 69 sites spread over 20 countries participated. For the Rome production 8.5 million events were simulated and over 6 million reconstructed, in ~570 thousand jobs. 84 sites in 22 countries participated.

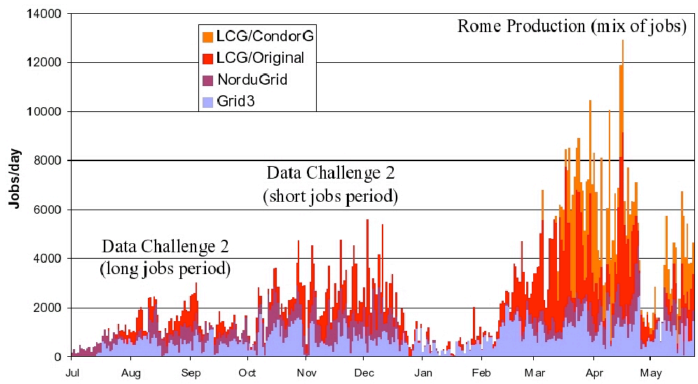

Figure 2: Number of jobs run per day.

Figure 2 shows that the total number of jobs run per day depended primarily on the type of jobs that were run (“long jobs period” corresponded to simulation jobs running for about 24 hours; “short jobs period” to “digitization” and “reconstruction” jobs running typically for 4 to 5 hours and finally “mix of jobs” corresponded to all kind of jobs running in parallel); the introduction in March 2005 of a new executor in LCG allowed a more efficient job submission as described in the following section.

EXPERIENCE

By design the ATLAS production system had a very high dependency on the Grid middleware. During DC2 many problems were identified and reported to the Grid developers, forming the basis for improvements. For some problems a workaround needed to be included in the production system. One example of the latter was the slowness of the job submission in the LCG. Even when using several submission points, it was not possible to saturate the LCG resources. This led to the development of a second executor, called Lexor-CG, using the Condor-G service [16], that submits jobs directly to the LCG computing elements, bypassing the resource broker. Introduced in March 2005, it promptly doubled the number of jobs we were able to run on LCG resources (see figure 2).

Also, on the ATLAS side, DC2 started before development was finished and many problems were found and corrected while running the DC. In general, our software was not prepared for the unreliability of the middleware: error handling, logging, etc. were too minimal. In addition the data management component DQ was lacking important functionality to manage the huge amount of files involved. Due to manpower shortage it was not possible to develop and implement these functionalities in due time.

CONCLUSION AND OUTLOOK

The three Grids in use in ATLAS have proven to be usable for large-scale productions. It was possible, however not easy, and requiring a non-scalable amount of manpower. All involved components both on the Grid and the ATLAS side showed to be clearly in need of substantial improvements. Problems were coped with, but often with work-arounds or *ad hoc* fixes

ProdSys II is in deployment phase. It features the same architecture, but has a new implementation of the crucial supervisor component (Eowyn) [17]. Effort has also been spent on factorising out the common pieces of executor implementations.

The data management project has been split off and now has more and separate manpower. Its new version, DQ2, features a new design that relies less on the middleware. It is also in a deployment phase.

Both ProdSys II and DQ2 systems are already used in the on-going productions but the integration is not completed. It is well advanced in the new executor (Panda) [19] developed in OSG/Grid3 but not yet in LCG and NorduGrid.

These new components form the basis of the ATLAS Tier-0 Management System that orchestrates the operation of the ATLAS Tier-0 at CERN. The scaling tests performed recently in the context of Service Challenge 3 were satisfactory. The results are presented at this conference [20].

The main goal currently is to complete their integration for the beginning of Service Challenge 4 in June 2006.

ACKNOWLEDGEMENTS

We thank all members of the ATLAS collaboration, more specially those of the ATLAS Data Challenges, Grid and Operations teams, who participated to this effort, the supporting people of the three Grids and the people involved in the maintenance of each site.

REFERENCES

[1] CERN-LHC: http://www.cern.ch/lhc
[2] ATLAS: http://atlas.web.cern.ch/Atlas/index.html
[3] LHC Computing Review: http://cern.ch/lhc-computing-review-public/Public/Report/Report_final.pdf
[4] LCG: http://lcg.web.cern.ch/LCG
[5] The Grid 2003 project: http://ww.iddgl.org/grid2003/index.php
[6] NorduGrid: http://www.nordugrid.org/
[7] CERN-PH-EP/2004-028
[8] L. Goossens et al., ATLAS Production System in ATLAS Data Challenge 2, proceedings of CHEP 2004, Interlaken, contribution no. 501
[9] The Windmill project: http://hepcc12.uta.edu/windmill/
[10] Lexor: A. De Salvo, G. Negri, D. Rebatto, L. Vaccarossa: "LEXOR, the LCG-2 Executor for the ATLAS DC2 Production System", Conference on Computing in High Energy and Nuclear Physics (CHEP04), September 2004, Interlaken (Switzerland).
[11] Oxana Smirnova et al., Performance of the NorduGrid ARC and the Dulcinea Executor in ATLAS Data Challenge 2. CHEP 2004, Interlaken, contribution no. 499
[12] Robert Gardner et al., "ATLAS Data Challenge Production on Grid3", CHEP 2004, Interlaken, contribution no. 503
[13] J. Kennedy, "The role of legacy services within ATLAS DC2", CHEP 2004, Interlaken, contribution no. 234
[14] Pacman: http://physics.bu.edu/pacman/
[15] M. Branco: "Don Quijote - Data Management for the ATLAS Automatic Production System", Conference on Computing in High Energy and Nuclear Physics (CHEP04), September 2004, Interlaken (Switzerland).
[16] Condor, High Throughput Computing, http://www.cs.wisc.edu/condor/condorg/
[17] https://uimon.cern.ch/twiki/bin/view/Atlas/Eowyn
[18] D. Cameron et al: A scalable Distributed Data Management System for ATLAS, CHEP 2006, Mumbai, contribution 75.
[19] K. De: "Panda" – Production and Distributed Analysis System for ATLAS, CHEP 2006, Mumbai, contribution 347
[20] L. Goossens et al. ATLAS Tier-0 Scaling Test, CHEP 2006, Mumbai, contribution 341.