EXPERIENCE SUPPORTING THE INTEGRATION OF LHC EXPERIMENTS COMPUTING SYSTEMS WITH THE LCG MIDDLEWARE

S.Campana, A.Delgado-Peris, F.Donno, M.Lamanna, P.Mendez-Lorenzo, R.Santinelli, A.Sciabà

CERN, Geneva, Switzerland

Abstract

The LHC Computing Grid Project provides and operates the computing support and infrastructure for the LHC experiments. In the present phase, the experiments systems are being commissioned and the Experiment Integration Support team provides support for the integration of the underlying grid middleware with the experiment specific components. The support activity during the experiments' data and service challenges provides valuable information on the infrastructure performance and deeper understanding of the whole system. Results from the major activities in 2005 are reviewed and a summary of the related activities is given, In addition the support of non High-Energy Physics communities (Biomed, GEANT4 and UNOSAT) is discussed.

INTRODUCTION

The Large Hadrons Collider at CERN will start data production in 2007. The goal of the LHC Computing Grid (LCG) project is to establish a Grid infrastructure for High Energy Physics (HEP) experiments operating their detectors at CERN. Back in 2002, LCG has established the Experiment Integration and Support (EIS) team in order to provide support to the HEP experiments adapting their computing models and software environment to operate on the Grid. Since then, the team has been involved in many activities that are now considered to be fundamental by the experiments: acquiring a good and deep knowledge of the Grid middleware in order to provide solutions; participating to discussions on the deployment of middleware components critical for the experiments; creating missing and useful tools and interfaces; setting up testing Grid environments; compilation of basic documentation such as the LCG Middleware User Guide; acting as link point between experiments, middleware and physics applications developers; organization of tutorials; collecting experiments software requirements; assisting during data and service challenge operations; dissemination and public events; other collateral activities such as providing support to other scientific communities (Biomed, Geant4, Unosat) and helping during the process of making a Virtual Organization (VO) official and recognized project wide.

In what follows we give a detailed overview of the various activities carried out by EIS on a per experiment basis. We then explore what has been done for non-HEP VOs and the need for such efforts. We also outline the commonalities among the different communities that led to outline general solutions. We then conclude with some summary remark and with an outlook on future work.

THE HEP EXPERIMENTS SUPPORT

During the past years in preparation for the final data handling and analysis process the four LHC experiments went through a series of computing exercises known as Data Challenges and Distributed Productions, as well as the Service Challenges to validate their computing model and to provide useful samples of simulation data for detector and physics studies. A Data Challenge consists of the full simulation and reprocessing of data coming from the detector, carried out with the same software and computing infrastructure expected to be employed during data taking. The aim of Distributed Productions is to provide data for physics studies. However, due to the methodology and the large number of events usually produced, Distributed Productions can be considered equivalent to real Data Challenges. Together with a clear definition and implementation of the basic Grid services required by the experiments, the Service Challenges form the basis of the overall plan for establishing the full-scale global LHC Grid service for the start-up of LHC in 2007. The EIS team offered support to the four LHC experiments during the activities mentioned above, both in terms of integration between experiment software stack and Grid middleware and in terms of day-by-day operations and user support.

The **ALICE** software design relies on the AliEn lightweight Grid framework: a production environment for the simulation, reconstruction, and analysis of the physics data. The integration of the generic AliEn interface with the LCG middleware is the main activity EIS has been involved with. Every year ALICE runs a Data Challenges exercise, intended to probe and evaluate the current software and hardware solutions in increasingly demanding and realistic environments approaching the requirements of the experiment data taking and analysis. In particular, the last ALICE Data Challenge was framed in the context of the last LCG Service Challenge. During Data Challenges the EIS team has been the LCG contact between the experiment, the sites and the developers: it collected ALICE requirements, brought them to the attention of the

Middleware developers while trying to understand possible solutions to improve the system and to allow the most efficient usage of LCG resources via AliEn. The integration support consisted also in the development of new clients to facilitate the access to the information system for computational and storage resources discovery. EIS has also provided the LCG Information System plug-in for the integration with the ALICE job monitoring system, based on MonAlisa. Initially developed to accommodate ALICE requirements, such tool has been extended to accommodate various different use cases and used by several other VOs. During the Service Challenge production phase, EIS has been very active in the integration of the ALICE data transfer agents with the underlying LCG/EGEE services. For this purpose, a file transfer kit was developed. It interacts with the gLite File Transfer Service for transfer submission and status retrieval, with the Information System to gather the correct service endpoints to be used and with the MyProxy server to verify the validity of credentials. The kit is currently being integrated in the ALICE framework. In terms of operations, EIS followed up the production exercise, reporting problems to site and offering possible solutions or support. Moreover, it ensured the readiness of ALICE dedicated services at the sites, in particular troubleshooting experiment specific agent installations and configurations (VOBOX).

In the past years **ATLAS** produced a large sample of events on the LCG infrastructure to be used by the physicists. The EIS team provided support for the integration of the ATLAS production system with the LCG infrastructure, developing specific tools, interacting with the middleware developers and testing possible middleware functionalities. In addition, EIS followed up the production exercise, trying to understand the problems causing the large number of job failures and offering possible solutions. As a result, the last production exercise (380k jobs run on the native LCG2 system and 1.4M files registered in the EDG-RLS catalogue) has been considered a very positive experience and a big improvement with respect to previous exercises of the same kind [1]. In the second half of 2005 ATLAS was focused on the Service Challenge. EIS followed up particularly the integration of the ATLAS Data Management System with the new middleware components like the File Transfer Service and the LCG File Catalog. In addition, a prototype Data Location Interface was developed as a first step for the communication between the ATLAS Workload and Data Management Systems. EIS played a primary role in the technical coordination of the ATLAS Service Challenge activities, ensuring the readiness of the sites before and during the exercise and following up issues with the different services. For this purpose, an ATLAS specific monitoring tool was developed to test correct functioning of File Transfer servers, File Catalogues and Storage Elements, publishing the results in a web interface. In the contest of the ATLAS-LCG-EGEE task force to tackle specific issues with the middleware, EIS participated to the test for the performance of the new gLite Workload Management System. EIS is also active in the ATLAS Grid end user support, following up the requests entered via the Global Grid User Support system.

The **CMS** computing model includes, as for the other experiments, two main kinds of application: massive event simulation, and data analysis. Both have been performed on the LCG Grid infrastructure, using tools adapted to (for Monte Carlo production), or specifically developed for (CRAB) executing jobs on LCG. Especially for the case of analysis of simulated data, a rapid and substantial increase in the number of CMS users implied a correspondingly large increase in the load for the user support team that at first was almost completely absorbed by EIS. Still today, the largest active user community in LCG is represented by the CMS virtual organization. The experience of the EIS team allowed the developers of the CMS distributed production and analysis tools to substantially improve the reliability of those tools and the perception of LCG from the user's point of view.
In 2005, EIS was also involved in testing how Grid middleware can satisfy some of the CMS needs in relation to data management and cataloguing. In particular, it was investigated if the LCG File Catalogue (LFC) [2], used through its POOL catalogue interface [3], could be used as a Local File Catalogue, in the framework of the upcoming data management system, as described in the CMS Computing Model [4], and also as a PhEDEx [5] file catalogue. This involved extensive testing of the POOL Catalogue API using LFC as a backend, and performance measurements for typical read, write and query operations on files. The results were constantly fed back to the developers, which lead to drastic improvements in the usability of LFC in this context. A different, yet related activity was to study LFC as a possible implementation of the Data Location Service (DLS), a central catalogue that contains the information on the location of the replicas of the CMS data. A mapping between the proposed DLS API and the LFC API was defined and found to require minor modifications to the LFC schema. Limitations in performance were addressed by changes in the LFC protocol; currently, there is strong evidence that LFC would satisfy the needs of a DLS even at full scale. During the last LCG Service Challenge [6], EIS was involved on maintaining and running an automatic job submission system, whose purpose was to simulate access rates to data published by various CMS Tier-1 and Tier-2 sites. As a result, detailed statistics on job failures and their rates were collected, and provided a clear picture of the current level of stability of the LCG infrastructure.
It is worth mentioning that EIS is also partly involved in the management of the CMS Virtual Organisation in LCG, which includes the registration of users, the definition of VOMS [7] groups and roles and the assignment of users to them. As an interesting side effect,

this allows EIS to have a good insight of issues like policies and user privileges, which are currently seen as one of the areas of the LCG middleware requiring substantial improvements.

The **LHCb** computing model relies on the Dirac framework. Originally based on remote agents installed at the sites, Dirac automates most of the production tasks, including job definition, submission scheduling and monitoring, as well as software installation and data management. At the very beginning of the Data Challenge in April 2004, LHCb was running via Dirac more than 85% of the production on non-LCG sites. EIS has been fully involved in the evolution of Dirac toward a more effective use of the LCG/EGEE infrastructure offering expertise and assistance in the integration of new middleware like SRM storage elements, the LCG File Catalog and the Data Location Interface. LHCb has been the first community using customized experiment specific tests on the Grid Site Functional Tests (SFT) for monitoring and automatic management of sites in production; it has provided a list of issues/requirements to the SRM [8] developers for further evolution of the service; it is currently looking forward for new solutions coming from the middleware providers like the new generation Computing Element (CREAM), a new Information System component (CEMon) and the new system for enforcing grid policies (G-PBOX).
The first phase of the LHCb Data Challenge was concluded by running the majority of the jobs on the LCG Grid infrastructure with a reasonable efficiency reaching a peak of 6K jobs running simultaneously per day.
EIS gave his contribution by providing customized distribution of standard LCG client tools, like gridftp and lcg_utils, servers like gridftpd on special non-LCG storages server and experiment specific tools for querying directly the Information and the Storages Systems in order to integrate the Dirac production framework with LCG.
Furthermore, EIS has always been involved in Operational and User Support activities devoted to chasing/tackling non functioning sites and middleware related problems and to answering questions and giving hints, suggestions or recipes for an optimized middleware usage, making them available for consultation by everybody. EIS has also given important contributions in proposing, testing, evaluating and integrating new architectural solutions into the LHCb computing framework A good example is the integration of the new automatic DIRAC submission system for the job analysis management with LCG middleware by providing mechanisms for authorization/authentication permissions renewal on non-LCG service machines, by writing PYTHON APIs for job management on LCG and by writing tools for estimating the normalized CPU and Wall Clock time within a job independently of the underlying batch system.

SUPPORT FOR NON-HEP VIRTUAL ORGANIZATIONS

The computational and storage capability of the Grid is attracting other research communities, beside HEP.

The **Biomed** community has certain peculiarities that make it different from the HEP experiments. Firstly, this VO is not yet as mature as the physics VOs. During the last months, the EIS team has collaborated with them in the developing of a Grid culture. Additionally, it is a much more scattered community, in that each of the different applications that form part of it are independent, lacking the centralized job submission and data management framework that the physic VOs present. This diversity of applications brings in some new requirements to the Grid. One of the main taks of the EIS team was to identify and evaluate these necessities and communicate them to the appropriate people in the deployment group.
Finally, the great achievement of the Biomed VO was the realization of the first Biomedical Data Challenge in summer 2005, in which 72.751 jobs were run, 1 TB of data was generated, and the equivalent of ~70 CPU years were computed. This was the first occasion in which the community really worked organised as a whole, and the EIS team has collaborated to improve their overall job submission strategy as well debugging some site problems.

The **Geant4** collaboration asked the LCG support to perform several simulations per year on the LCG resources. Geant4 is a toolkit for the Monte Carlo simulation of the interaction of particles with matter. It is applied to a wide field of research including HEP, nuclear experiments, medical, accelerator and space-physics studies. ATLAS, CMS, LHCb, Babar and HARP are actively using Geant4 in production.
In December 2004, the Geant4 team performed a first complete production using the LCG resources to test the new versions of their software. In order to perform this production, the EIS team has provided an initial toolkit able to control job execution on the Grid. At the end of the first Geant4 production, the toolkit has been extended and generalized. It can be used now by any collaboration or experiment to perform their initial tests.
The second production of Geant4 (May 2005) was performed using already the generalization of the toolkit.
Users who already have developed their own software can also get advantage of the use of this toolkit.
Geant4 became then a fully recognized and supported VO in the EGEE project.

In summer 2005, another community, **UNOSAT,** asked for support to use the LCG resources. UNOSAT is a United Nations initiative to provide the humanitarian community with access to satellite imagines and Geographic System services. UNOSAT is promoted by the UN Institute of Training and Research (UNITAR) and managed by the UN Office for Project Services (UNOPS).
EIS has covered an important role helping UNOSAT to integrate their software environment with the LCG Grid

infrastructure: the satellite images provides by UNOSAT have been stored in Storage Systems at CERN and registered in the LCG Catalog (LFC). The files so registered have been identified with an easy-to-remember Logical File Name (LFN). The UNOSAT users provide as input information the coordinates of each image. AMGA, a database developed by the LCG/ARDA team, maps these coordinates to a corresponding LFN and therefore to the physical location of the images.

COMMON SOLUTIONS

Many EIS activities are not focused to the support of a particular VO. Even though developed in the context of providing support to the needs of a specific scientific community, EIS provided several new tools to simplify the middleware usage or facilitate the integration of the experiment framework with the middleware itself. As examples we can mention: customized versions of the LCG Data Management clients; tools to facilitate the collection of information from the Information System; utilities to inspect the standard error and standard output of running jobs, scripts to evaluate the remaining WallClock and CPU time of jobs on a batch system, a generic C++ interface for the information system and many others. It is relevant to notice that several functionalities provided by those tools have been then integrated in the standard LCG middleware proving the importance of the developed utilities. In addition, EIS provided complete middleware components and software toolkits: the first (and only) prototype of the VO-BOX node to run experiment specific agents and currently deployed in many LCG. This includes services such the Globus GSI-enabled *ssh* server and relative configuration, an ad-hoc proxy renewal service with relative user level tool, a VO service publication mechanism and the overall configuration of the node type. Moreover, EIS has developed a full solution for automatic experiment software installation and configuration at a site from a central control station. This consists of a high level interface that drives the installation from the LCG User Interface and a set of services deployed at every site to manage the installation, validation, removal and general management of the software and to publish the relative software tags in the Information System. In addition, a fabric level client-server application (*Tank-and-Spark*) has been designed and developed in order to overcome limitations of the current mechanism and accommodate site and experiment requirements. This relies on a *MySQL* database to hold the state of the installation for different nodes and a *rsync* server to distribute the software [9].

CONCLUSIONS

For any new community exploiting the Grid infrastructure it is fundamental to individuate a support team to assist the new collaboration during the implementation, integration and execution of specific applications on the Grid. This guarantees a successful, smooth and efficient entrance in the Grid environment. The Grid support team will be very often also the Virtual Organization contact person with the EGEE/LCG infrastructure. This person will work together with the EGEE/LCG developers and the deployment team and with the responsibility of helping the new community with all aspects of support: providing solutions, developing missing integration tools, ensuring that adequate security policies are enforced, assisting during operations, providing user support with the Grid middleware usage, compiling specific documentation, organizing tutorials specific for the community of users and applications. Once these new communities have reached a good level of maturity and confidence with the Grid infrastructure, the support can be reduced to normal operations.
The EIS group has worked actively to support these new Virtual Organization providing the support and the infrastructure necessary to make them effective on the Grid in a very short time.

ACKNOWLEDGEMENTS

We are thankful to the LCG Project at CERN which has made this work possible. This work has been partially funded by Istituto Nazionale di Fisica Nucleare, Rome – ITALY; Ministero de Educaciòn y Ciencia, Madrid – SPAIN and CERN, Geneva - SWITZERLAND.

REFERENCES

[1] Analysis of the ATLAS Rome Production Experience on the LHC Computing Grid. S. Campana, D. Barberis, F. Brochu, F. Donno, A. De Salvo, L. Gossens, S. G. de la Hoz, T. Lari, D. Liko, J. Lozano, G. Negri, L. Perini, G. Poulard, S. Resconi, D. Rebatto, L. Vaccarossa, e-Science 2005, Dec. 5 - 8, 2005, Melbourne, Australia
[2] http://wiki.gridpp.ac.uk/wiki/LCG_File_Catalog
[3] http://lcgapp.cern.ch/project/persist/
[4] http://lcg.web.cern.ch/LCG/peb/LHCC /expt_reqts/CMScomputingmodel.pdf
[5] http://cms-project-phedex.web.cern.ch/cms-project-phedex/
[6] https://twiki.cern.ch/twiki/bin/view/LCG /LCGServicChallenges
[7] VOMS, an Authorization System for Virtual Organizations, R. Alfieri ,R. Cecchini ,V. Ciaschini, L. dell'Agnello,A. Frohner,A. Gianoli ,K.L Lorentey, F. Spataro, 1st European Across Grids Conference, Santiago de Compostela, February 13-14, 2003
[8] http://sdm.lbl.gov/srm-wg/
[9] Installing and Configuring Application Software on the LHC Computing Grid. R. Santinelli, F. Donno, e-Science 2005, Dec. 5 - 8, 2005, Melbourne, Australia