

LARGE HADRON COLLIDER COMMITTEE

LCG PROJECT COMPREHENSIVE REVIEW

November 2005

1. EXECUTIVE SUMMARY

The LHC Computing Grid (LCG) Project was created by the CERN Council in September 2001 with the aim of prototyping and deploying the computing environment for the LHC experiments. The formal launch of the project was at a workshop held in March 2002. Since that time, the LCG Project has demonstrated progress towards the realisation of the computing requirements of the experiments in time for LHC operation in 2007.

The LCG Project is a collaboration of the LHC experiments, the Regional Computing Centres, CERN and the physics institutes with the aim of preparing and deploying the computing environment that will be used by the LHC experiments to analyse the LHC data. The project includes support for applications and the development and operation of a computing service.

The LCG Project is divided into two phases. Phase I (2002-2005) has the objective of building a service prototype, based on existing Grid middleware, of running a production Grid service and producing the Technical Design Report for the final system. Phase II (2006-2008) will build and commission the initial LHC computing environment. The LCG is not a Grid development project and it relies on other Grid projects for the middleware development and support.

The LHCC considers that the LCG Project has shown significant progress since the last Comprehensive Review. In particular, the Service Challenges have been successful, albeit with a lower-than-expected through-put, the Grid middleware quality and stability and the interoperability between the various middleware systems on a small scale have improved, the Baseline Services have been defined and more focus has been put on practical aims matched to the requirements of the experiments. The associated LCG Technical Design Report was submitted in June 2005 and reviewed by the LHCC in October 2005 and the Computing Memorandum of Understanding was approved by the corresponding Resource Review Board in October 2005.

However, the Committee did note certain concerns. The connection between the middleware and the experiments is considered to be too weak, with the risk that the former will not satisfy the requirements of the experiments. Moreover, significant delays were reported in the deployment of the Enabling Grids for e-Science in Europe (EGEE) gLite services, in the CASTOR2 disk pool management system and in the distributed database services. The interoperability and the provision of a common interface of the various types of middleware being produced should continue to be pursued at a larger scale. The analysis models are to a large extent untested and the database service deployment is late.

The third annual LHCC Comprehensive Review of the LCG Project took place on 14-15 November 2005. The LHCC referees addressed the following areas: Overview and Global Collaboration, Applications Area, Fabric Area, Grid Deployment Area, Service Challenges, and the issues of Management and Resources. The LHCC acknowledges the considerable amount of work that has gone into the preparation of the LCG Project Comprehensive Review.

The conclusions and concerns of the LHCC are given below. They will help the Committee to follow up outstanding issues and to monitor future progress of this project in forthcoming sessions of the LHCC prior to the next LCG Project Comprehensive Review one year hence.

2. OVERVIEW

- Good progress was reported on moving towards the Worldwide LCG, with the interoperability of various Grid middleware improving markedly, but much work remains to be done. Challenges have been identified for 2006 and will be monitored by the LHCC.
- Significant progress was reported on the Applications Area, with no major concerns having been identified. Measures have been taken to stabilize the decreasing level of manpower effort in the Applications Area.
- The LHCC notes the significant progress in the Fabric Area. However, the Committee expressed concern on the reported delays related to CASTOR2, and will continue checking on the availability and performance of this system.
- Progress was reported on the Grid Deployment Area, particularly in the fields of Grid Operation and Baseline Services. The deployment of some components of the gLite middleware is very late. Continuing work on the interoperability of the various Grid systems is required.
- The latest Service Challenge 3 provided the LCG Project, computing sites and the experiments with much-needed experience prior to LHC operation in 2007. The lower-than-expected through-put observed in Service Challenge 3 and the associated problems encountered need to be addressed.
- A great deal of progress was reported since the previous LCG Comprehensive Review in November 2004. The organization structure has been adapted to the changing needs of going from the development to the service phase, and resources are being put together.

3. OVERVIEW AND GLOBAL COLLABORATION

The LCG Project does not develop and therefore does not control the middleware technology evolution. However, it does identify the starting technologies to be deployed, sets-out the evolution strategy and the long-term support technologies and works towards the future coherent, acceptable and supportable middleware solutions.

There are three “grid infrastructures” used by the sites involved in the LCG Project: EGEE (Enabling Grids for E-sciencE), OSG (Open Science Grid) and NDGF (the Nordic Data Grid Facility). Each of these “infrastructures” has:

- a Grid middleware “distribution” – a self-consistent software package of libraries and tools drawn from different sources, integrated, tested, and distributed to member sites with defined maintenance and support services;
- a set of operations processes and procedures (depending on the grid infrastructure) covering the registration, authentication, authorization, security, monitoring, and problem determination, etc.;

- a system of “governance” with defined rules for membership and decision making.

OSG uses the VDT (Virtual Data Toolkit) distribution, including components from the Globus and Condor projects, and also now one component from EGEE. EGEE uses at present the LCG-2 distribution, which includes VDT and also components developed by the EGEE and LCG projects. NDGF uses the ARC (Advanced Resource Connector) middleware originally developed by the NorduGrid project. All of these distributions are founded on the Globus toolkit, although NDGF uses a different version from the others.

The middleware developed by the EGEE project is called gLite. This is a set of components, some of which are re-engineered versions of components developed by the predecessor European DataGrid project. The gLite components have been designed and implemented as a stand-alone package, and this is used by some applications communities within the EGEE project. However, it was realized over a year ago (September 2004) that an evolutionary approach would be required for the EGEE Grid that was already in operation using the LCG-2 distribution and over the past year the LCG-2 distribution has been extended to include various components of gLite as they became available, as well as components from elsewhere.

The deployment and operation of the EGEE Grid is a joint activity of LCG and EGEE. While LCG has been initially setting-up the operation of a Grid infrastructure for the LHC experiments, EGEE, through its additional manpower, has the goal of extending this into a general Grid infrastructure for science in Europe but with a strong focus on HEP requirements.

Positive steps have been taken towards the Worldwide LCG (WLCG), organizing services for LHC on all three of the Grid infrastructures. Good progress was reported on Grid interoperability as evidenced by the transfer of jobs between LCG-2 (including gLite components) and the VDT middleware distribution used by the US OSG, but the effort needs to be strengthened and exercised on a larger scale. Work is also ongoing for inter-operation between EGEE and NDGF but this is still at the conceptual level. The involvement of the LHC experiments is considered mandatory for the successful implementation of interoperability between the various Grids. These issues are acknowledged by the LCG Project management and are being addressed.

ARDA (A Realization of Distributed Analysis for the LHC) is an LCG project whose main aim is to enable LHC analysis on the Grid and to provide the interface to existing middleware in order to enable their use by the experiments. A better-defined integration with experiment activities on distributed analysis was noted. The future role of ARDA needs to be examined in view of the establishment of the LHC experiment task forces.

The challenges for 2006 include the clarification of use for the CERN Analysis Facility, end-to-end testing of the DAQ-to-Tier-0-to-Tier-1 chain, and a full Tier-1 work-load testing.

Good progress was reported on moving towards the Worldwide LCG, with the interoperability of various Grid middleware improving markedly, but much work remains to be done. Challenges have been identified for 2006 and will be monitored by the LHCC.

4. APPLICATIONS AREA

The Applications Area develops and maintains that part of the physics applications software and associated infrastructure that is shared among the LHC experiments. The scope includes common applications software infrastructure, frameworks, libraries, and tools, together with common applications such as simulation and analysis toolkits, Grid interfaces to the experiments and assisting the integration and adaptation of physics applications software in the Grid environment. The aims and scope of the Applications Area have been defined properly through the RTAGs.

The major concern at the previous Comprehensive Review, regarding the manpower effort deployed in the Applications Area, has been stabilized. The Applications Area is to be congratulated for progress in this area.

Good progress was reported on the Software Process Infrastructure (SPI). In order to optimize the use of the decreasing level of manpower, certain measures have been implemented. The measures include the introduction of the direct participation from and to the development projects, the participation from people associated with the Application Area projects, the optimization and automation of procedures and the prioritization of tasks. Moreover, a change of focus has been introduced to provide services directly to the experiments. Future areas of work include the continuing provision of the SAVANNAH service for bug tracking and task management, the installation and distribution of software services and the provision of tools for software development and testing.

The merger of the ROOT interactive tool for analysis and SEAL (Shared Environment for Applications at LHC) is well-underway. The merger has been initiated in order to optimize resources and ease the long-term maintenance by avoiding duplicate development, to provide a coherent set of products developed and maintained by the Applications Area for the experiments, and to have the development of ROOT fully integrated in the LCG organization. The LHCC noted that the transition to the implemented changes will be gradual and the experiments request to continue maintenance of the SEAL libraries in the meantime, which may lead to duplication of effort. In order to complete the merger, it is important that agreement is reached on the contentious issues and that a strict set of milestones are set for the final conversion in the near future.

The development of PROOF, a prototype extending ROOT to run on a distributed, heterogeneous system, is progressing. However, the intended use by the experiments is not yet clear. The LHCC recommends that the experiments should make a decision on their intended use of PROOF so that the project can be adjusted accordingly and the relatively large manpower effort can be deployed effectively.

Good progress was reported on POOL (Pool of persistent Objects for the LHC), which handles the data persistency, data management, deployment on the Grid and relational databases in general, with no major concerns having been identified.

Progress was reported on the simulation packages. New releases of the GARFIELD simulation package for gaseous detectors, Geant4 and FLUKA have been or will be made available imminently. Good progress was also reported on the physics validation of simulation results using data from experimental measurements.

Significant progress was reported on the Applications Area, with no major concerns having been identified. Measures have been taken to stabilize the decreasing level of manpower effort in the Applications Area.

5. FABRIC AREA

The goal of the computing Fabric Area is to prepare the Tier-0 and analysis facilities at CERN. The former focuses on the mass storage of the raw data, the first processing phase and the data export, while the latter concentrates on the data analysis task.

CASTOR2 is the last major step in the process of enabling the CASTOR disk management system for use with LHC data. The Committee expressed concern at the delayed deployment of CASTOR2 and its lack of full functionality. Part of the delay in putting CASTOR2 in production is the weak interaction with the experiments. The development and deployment of CASTOR2 needs to be done in collaboration with the experiments and must be followed very closely. Clear migration strategies must be defined, taking into account the required level of support and the interactions with other Tier-1 centres using the product.

The alternatives to tape storage are becoming attractive but are still immature. Alternatives and evolutionary products, such as disk storage, need to be evaluated and tested in realistic scenarios and in the meantime tape storage will continue to be used for the next few years. The LHCC considers that this conservative approach is reasonable, given the numerous uncertainties in the alternative options.

The collaboration between the experiments and the Tier-1 centres needs to continue and improve, in particular in order to clarify the implications of the Computing Technical Design Reports for the Tier-1 design and operation requirements.

Procurement of equipment remains a problem area in several centres, including at CERN. The financial, management and technical issues that can result from lengthy purchasing procedures should not be underestimated.

Good progress was reported on the Extremely Large Fabric management system (ELFms), which manages and controls the heterogeneous CERN Computing Centre environment.

The Committee requests that a presentation of the strategies for the implementation of the computing fabric, in preparation for the LHC start-up, be discussed in future reviews. These strategies will form part of the LHCC reviewing process of the LCG Project.

The LHCC notes the significant progress in the Fabric Area. However, the Committee expressed concern on the reported delays related to CASTOR2, and will continue checking on the availability and performance of this system.

6. GRID DEPLOYMENT AREA

The aim of the LCG Grid Deployment Area is to deploy and operate a pre-production LHC computing environment. Its scope is to integrate a set of middleware, to coordinate and support its deployment in the Regional Centres, to provide operational services to enable running as a production quality service, to provide assistance to the experiments in integrating their software and to provide direct user support. The LCG Grid Deployment Area relies upon the close interaction with the EGEE project.

Significant progress was reported on the Grid operation. A framework for Site Functional Tests (SFTs) is in place, although not all tests have been implemented, the verification of integrity and consistency of published information is available, the Regional Operation Centre and Core Infrastructure Centres are on duty and the user support is in place. Much of the manpower in this area is funded through the EGEE, and the implications of the end of the EGEE-2 and the transition to the next phase should be considered very seriously. The LHCC encourages the European Grid Organisation (EGO) initiative of CERN and other HEP institutes, which proposes a longer term structure for supporting international science Grid infrastructure.

The development of the gLite middleware is late and is at risk of interfering with the Service Challenge operations. The gLite middleware was originally scheduled to be released a few weeks following the November 2004 LCG Comprehensive Review, but is only now being deployed. The Committee regrets that no quantitative results were shown during the Comprehensive Review in order to gauge the reliability and performance of the current middleware.

A set of Baseline Services provided by the LCG Project to the experiments has been developed since the previous Comprehensive Review. A measure of evaluating the performance and reliability of these services is being implemented.

Progress was reported on the Grid Deployment Area, particularly in the fields of Grid Operation and Baseline Services. The deployment of some components of the gLite middleware is very late. Continuing work on the interoperability of the various Grid systems is required.

7. SERVICE CHALLENGES

Service Challenges have been instigated to test the Grid service prior to the start of LHC operation in 2007.

Much experience was gathered from the latest exercise Service Challenge 3. The collaboration between the LCG Project, the computing sites and the experiments has been strengthened, all of the Tier-1 centres and more than twenty Tier-2 centres participated, many of the Baseline Services including the LCG File Catalogue were deployed and many problems were resolved and improvements made. However, data through-put for the Service Challenge 3 has been much lower than planned. Problems related to configuration issues, instabilities of CASTOR and the dCache hierarchical storage manager and the lack of support for the SRMcp protocol by the gLite File Transfer Service (dCache performs better with this protocol) contributed to the lower-than-expected performance. Corrective action has been taken in a number of areas and the through-put tests have been re-scheduled for early 2006. The above concerns underlined the complexity of the exercise and the necessity of carrying-out successfully the programme of Service Challenges before 2007.

Planning for Service Challenge 4 is underway with the aim of demonstrating the full service. The Challenge will include all Tier-1 centres and all major Tier-2 centres as well as the full set of Baseline Services and the data distribution and recording at the nominal LHC rate.

The latest Service Challenge 3 provided the LCG Project, computing sites and the experiments with much-needed experience prior to LHC operation in 2007. The lower-than-expected through-put observed in Service Challenge 3 and the associated problems encountered need to be addressed.

8. MANAGEMENT AND RESOURCES

The LCG Project has started the transition from the development phase to the service phase. The associated changes in the organization structure reflect this shift of emphasis in the LCG Project and include the introduction of a Management Board to replace the Project Execution Board, the creation of a Collaboration Board and focused Task Forces organized by each experiment. As a result, the planning, monitoring and reviewing process has been greatly improved as the project milestones are now adequately monitored, the Tier sites and experiments are involved in the management structure and a satisfactory resource sharing has evolved. The set of high-level milestones is being defined and need to be revisited once the planning process is complete and once each milestone has a clear deliverable. The associated LCG Technical Design Report was submitted in June 2005 and reviewed by the LHCC in October 2005. The Computing Memorandum of Understanding was approved by the corresponding Resource Review Board in October 2005.

Good progress was reported on securing resources for the LCG Project. The level of computing hardware is globally satisfactory. But the balance amongst the experiments for resources outside of CERN will be difficult to achieve. The short-term deficit in manpower at CERN is being addressed through a re-direction of budgets and funding from external sources. The issue of long-term support remains a concern, particularly for manpower regarding the middleware and operational support.

Distributed databases are needed for condition, calibration and alignment data and for event tag information. At present, the experiments do not have clear access patterns and estimates for the quantity of data to be distributed in databases. In view of this, database service deployment is late and it is only now that services are being set up at the Tier centres. The service requirements for the Tier-1 and Tier-2 centres need to be agreed. The LCG 3D database common project is scheduled to provide a full service for the LHC as of September 2006.

An extensive and redundant network is available world-wide for the LHC computing needs. This network provides the central architecture for linking the various Tier centres. There are no major Tier-0 to Tier-1 networking technical issues, and the models from Tier-1 to Tier-2 centres are evolving.

A great deal of progress was reported since the previous LCG Comprehensive Review in November 2004. The organization structure has been adapted to the changing needs of going from the development to the service phase, and resources are being put together.

E. Tsesmelis