

PRINCIPAL LHCC DELIBERATIONS

10TH MEETING OF THE COMPUTING RESOURCES REVIEW BOARD

24 OCTOBER 2006

EMMANUEL TSESMELIS

SCIENTIFIC SECRETARY, LHCC

GENERAL

This document summarises the principal LHCC deliberations concerning the LHC Computing Grid (LCG) Project at the Committee's sessions in May, June and September 2006.

CONCERNS FROM THE PREVIOUS COMPUTING RESOURCES REVIEW BOARD

SUB-AREA	CONCERN	STATUS
Fabric Area	Delays in the deployment of the CASTOR2 disk pool management system. The database service deployment is late.	CASTOR2 has been successfully deployed at the Tier-0. Rolling it out at several Tier-1 sites went less smoothly. The distributed database infrastructure is now ready for production at CERN and at seven Phase 1 Tier-1 sites. The remaining Phase 2 Tier-1 sites will not meet the October 2006 milestone and this is a concern.
Middleware Projects	Grid interoperability needs to be strengthened and exercised on a larger scale. Delays in the deployment of the EGEE gLite services. Connection to the requirements of the experiments is considered to be weak.	Much work has been invested to support authentication, job submission, and mass storage access across the Grid borders. The gLite 3.0 services have been successfully deployed. The experiments have been included in the middleware planning group of the EGEE.

LHCC COMPREHENSIVE REVIEW

EXECUTIVE SUMMARY

The fourth annual LHCC Comprehensive Review of the LCG Project took place on 25-26 September 2006. The LHCC referees addressed the following areas: Management and Global Collaboration, Middleware, Applications Area, CERN Fabric, Distributed Fabric, Grid Deployment & Operations, Service Challenges, and Project Planning & Communications. The LHCC acknowledges the considerable amount of work that has gone into the preparation of the LCG Project Comprehensive Review.

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 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) had 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) is building and commissioning 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 interoperability between the various Grids, particularly the Enabling Grids for e-Science in Europe (EGEE) and the Open Science Grid (OSG), has improved considerably, the Worldwide LCG (WLCG) has successfully deployed the gLite 3.0 middleware, the CERN computer centre is on schedule to be ready to handle the requirements for the Tier-0 and Calibration Analysis Facility (CAF) for LHC start-up in 2007, the data transfer rates achieved in the Service Challenges, although not meeting fully the original targets, are compatible with the latest estimated requirements from the experiments, and the planning, reporting and reviewing system has improved significantly.

Moreover, the Collaborations are updating their resource requirements in view of the revised LHC schedule for the years 2007-2010. Initial indications show that the modification in the schedule reduces the gap between the required and available resources for these years, but the total level of funding requested will be required to complete the LCG Project.

However, the Committee did note some concerns. The overall stability of the service needs to be improved, and a reliable, unattended operation has not yet been achieved, particularly for the Tier-1 and Tier-2 sites. A complete test of the entire chain from the DAQ to the physics analysis is still lacking, although many pieces have been tested. Attention should be given to the storage management systems – the CASTOR2 performance and stability needs to be improved both for the Tier-0 and Tier-1 sites, the new version of the Storage Resources Manager interface needs to be implemented in all of the mass storage systems.

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.

OVERVIEW

- Good progress was reported on the interoperability between the various Grids, with much effort being put into authentication, job submission and mass storage access across the Grid borders. The future role of NorduGrid within the WLCG needs to be clarified. The Collaborations are updating their resource requirements in view of the revised LHC schedule for the years 2007-2010. Initial indications show that the shortfall in resources is reduced for these years, but the total level of funding requested will be required to completed the LCG Project.
- WLCG has successfully deployed gLite 3.0 and is improving the usage of the EGEE and OSG Grids. The experiments presented a coordinated plan of the necessary functions and services that have still to be deployed. Adequate manpower should be devoted to fixing of bugs and to consolidating the existing code.
- Good progress was reported on the Applications Area, with no major concerns having been identified. Action must be taken soon to ensure the required level of manpower remains available beyond 2007. The experiments are encouraged to make decisions on the use of PROOF.
- The CERN computer centre infrastructure is on track to be ready to handle the known Tier-0 and CAF requirements for LHC start-up in 2007. The funding and manpower situation of the CERN Fabric has improved considerably.
- Attention should be given to implement the Storage Resource Manager interface to the storage management systems. The performance and stability of CASTOR2 needs to be improved for both the Tier-0 and Tier-1 sites.
- The Distributed Fabric is well suited to the challenges it faces for LHC production and analysis. Tier-1 sites should improve communication with the experiments and local Tier-2 sites, and focus on stability. Middleware software developers must concentrate on stability over functionality. The 3D Phase 2 Tier-1 sites should move rapidly to deployment. Experiments should provide realistic estimates of their database requirements.
- Considerable progress was reported on the deployment of the Grid services and operations, highlighted by the rapidly increasing usage of the EGEE and OSG Grids. The Committee recommends that emphasis is put on improving further the stability and reliability of EGEE and OSG services and that the experiments get further involved in the first line support for users in their organisation.
- There has been a significant amount of work and progress in the services during the experience gained through the Service Challenges. The rates achieved, although not meeting fully the original targets, were compatible with the latest estimated nominal requirements from the experiments, and are considered to be an important result. The experiments have successfully used the Grid for many tasks, albeit with some problems that need to be corrected.
- The planning, reporting and reviewing system of the LCG has improved significantly in Phase 2 of the WLCG Project.