

CERN - C - RRB - 2005 - 003

6 APRIL 2005

PRINCIPAL LHCC DELIBERATIONS

7TH MEETING OF THE COMPUTING RESOURCES REVIEW BOARD

19 APRIL 2005

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 November 2004 and February and March 2005.

LHCC COMPREHENSIVE REVIEW

The second annual LHCC Comprehensive Review of the LCG Project took place on 22-23 November 2004. The LHCC referees addressed the following areas: Status Update, Management and Planning; Middleware and Interaction with EDG/EGEE and other Developments in Middleware; Fabric Area and Wide Area Networking; Grid Deployment and Regional Centres; Applications Area. 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.

- Progress was reported in the development and use of the middleware. In particular, the LCG-2 middleware has been deployed at around 90 remote sites and used in the Data Challenges 2004. However, the LHCC noted outstanding issues concerning the LCG-2 low job success rate, inadequacies of the workload management and data management systems, as well as delays in the release of the gLite (Lightweight Middleware for Grid Computing) services of the EGEE (Enabling Grids for e-Science in Europe). Continued delays in gLite may hinder future progress in ARDA (A Realization of Distributed Analysis for the LHC). Since the Comprehensive Review, all requested gLite modules have been made available as prototypes and are on hand for testing by users. The connection between LCG, gLite and the experiments is considered to be too weak, with the risk that the middleware will not satisfy the requirements of the experiments. LCG-2 has been used as a production batch system, but Grid-based analysis of the simulated data is only just starting. The interoperability and availability of a common interface of the various types of middleware being produced should be pursued and developers of the gLite middleware should remain available for the support phase. Since the Comprehensive Review, the LCG Grid team and the US Grid team have started to explore the technical aspects for interoperability and several points of collaboration have been identified showing that real interoperability is considered to be feasible.
- The LHCC has no major concerns regarding the Fabric Area and Wide Area Networking. A well-defined set of tools for the Fabric Area and Wide Area Networking has been chosen. However, in view of the reported delays, the Committee will continue checking on the availability and performance of the CASTOR disk pool management system. Since the Comprehensive Review, significant effort in the development and testing has been performed and the CASTOR team is close to deploying the new version.
- Good progress was reported on the installation of Grid software in remote sites. A large amount of data has been processed on the LCG-2 Grid as part of the Data Challenges and the LCG-2 Grid has been operated successfully for several months. However, the LHCC noted that the service provided by LCG-2 was much less than production quality and the experiments and LCG Project expended a large amount of effort to be in a position to use the service.

- The LHCC noted the good progress in the Applications Area with all projects demonstrating significant steps in the development and production of their respective products and services. The major outstanding issues lie with the insufficient coordination between the Applications Areas and ROOT and with the imminent reduction of manpower due to the transition from the development to the deployment, maintenance and support phases.
- The LHCC took note of the upcoming milestones for the LCG and noted that discussions are currently underway to secure the missing manpower to develop, deploy and support the Grid services. The lines of responsibility and authority in the overall organization structure need further clarification.

LHCC REVIEW OF COMPUTING RESOURCES FOR THE LHC EXPERIMENTS

There has been a significant increase in the required computing resources compared to the LHC Computing Review of 2001 with the most significant increase for all experiments being in the amount of the required disk storage. The Committee believes that the increased requirements almost certainly make the present computing models more credible, particularly when considering the support of a large number of simultaneous users and processors. In addition, the experiences from the recent data challenges have provided a foundation for testing the validity of the computing models but tests are at this moment incomplete. There is a concern about the validity of the data analysis components of the computing models and there is a risk that resource requirements for disks and CPUs will increase if the Grid functionality is not properly used.

The principal recommendations from the Review Committee are: a) the average and peak computing requirements of the four experiments require further detailed study, b) plans for the distributed analysis during the initial LHC period should be developed, c) the dependence of the computing model on factors such as the raw event size and reconstruction time should be addressed for each experiment, d) details of the ramping-up of the system in the years 2006-2008 should be determined and a plan for the evolution of the required resources should be worked out, e) a complete accounting of the offline computing resources required at CERN is needed for the period 2006-2010, and f) requests for Tier-1 / Tier-2 functionality at CERN should be supported and the associated planning should be refined for the four experiments.