

CERN-RRB-2008-016

18 MARCH 2008

# PRINCIPAL LHCC DELIBERATIONS

---

13<sup>TH</sup> MEETING OF THE COMPUTING RESOURCES REVIEW BOARD

15 APRIL 2008

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 2007 and February 2008.

---

**CONCERNS FROM THE PREVIOUS COMPUTING RESOURCES REVIEW BOARD**

---

<b>SUB-SYSTEM</b>	<b>CONCERN</b>	<b>STATUS</b>
Storage Management Systems	Considerable improvement in this area has been noticed since the beginning of 2007, but more work is needed to complete fully their deployment at the LHC.	<p>Deployment of the Storage Resource Manager (SRM) Version 2.2 is approaching completion and transparent mass storage metrics are under preparation.</p> <p>Deployment of SRM Version 2.2 to all Tier-1 and many Tier-2 sites was achieved prior to the start of the February phase of the Combined Computing Readiness Challenge 2008 (CCRC08). At the beginning of February 2008 there were about 160 sites with an SRM Version 2.2 storage system in place.</p>
Service	<p>Overall stability of computing service has improved but has not yet reached the necessary level.</p> <p>Complete test of the entire chain from the DAQ to the physics analysis is still lacking.</p>	<p>The service level has constantly improved and the nominal data transfer through-put rate for 2008 has been achieved, although only for short periods. Sustained, stable operation still needs to be achieved.</p> <p>The complete test remains to be done for all experiments. This will be done during CCRC08 when all the services will be tested at full capacity <i>at the same time</i> for all the experiments.</p>

---

## LCG COMPREHENSIVE REVIEW

---

The fifth annual LHCC Comprehensive Review of the LHC Computing Grid (LCG) Project took place on 19-20 November 2007. The LHCC referees addressed the following areas: Management, Resources and Collaboration; Mass Storage and Networking; Distributed Fabric; Middleware Development and Deployment; Applications Area and Distributed Databases; and the issues of Services and Experiment Readiness. The LHCC acknowledges the considerable amount of work that has gone into the preparation of the LCG Project Comprehensive Review.

The 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 2008.

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 of 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 both the production and analysis phases and that the World-wide distributed LCG (WLCG) is becoming a reality. In particular, considerable progress has been achieved in the stability, usage and interoperability of the Grid infrastructure and in the use of the Grid by the experiments for analysis; most of the Tier centres have installed and run successfully the necessary hardware, while the Storage Resource Manager SRM Version 2.2-compliant services are being deployed, albeit after lengthy delays; the middleware services are in place and focus has shifted to ensuring stability of the installed features. *Since the Comprehensive Review, deployment of SRM Version 2.2 to all Tier-1 and many Tier-2 sites was achieved prior to the start of the February phase of CCRC08. At the beginning of February 2008 there were about 160 sites with an SRM Version 2.2 storage system in place.* The service level has constantly improved and the nominal data transfer through-put rate for 2008 has been achieved. A number of useful products have also been delivered by the Applications Areas and significant progress has been reported on the monitoring and reporting performance for both generic and experiment-specific issues.

However, the Committee did note some concerns. The Grid infrastructure has been only partially exercised and the analysis models are not yet fully defined. Although site stability and reliability have improved, they are not yet at the desired level and the required support model, especially for 24x7 operations, is not yet fully defined. *Since the Comprehensive Review, both the stability and reliability have improved and have reached the required levels.* The deployment schedule for the mass storage management and operation remains one of the most critical issues for the WLCG and its schedule is extremely tight for the upcoming Combined Computing Readiness Challenge CCRC08, which is an important milestone as all the services should be tested in their complete capacity at the same time for all the

experiments prior to the start of LHC operations. Finally, it is also important that all pledged resources for 2008 are available for the CCRC08 in order to exercise the full system. The long-term guarantee of resources, both manpower and hardware, for the long-term also remains a concern.

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.

- Considerable progress has been achieved in the stability, usage, and interoperability of the Grid infrastructure. The LHC experiments routinely use the Grid for production activities, with approximately 20% of the production done at CERN, 40% at the Tier-1 centres, and 40% at Tier-2 centres. Great progress has also been reported in using the Grid for analysis, although the system has only been partially exercised, and the analysis models are not fully defined.
- As part of the system commissioning, the CCRC08 has been planned for 2008, with two windows, in February and in May. The CCRC08 is an important milestone, when all the services should be tested at full capacity *at the same time* for all the experiments.
- A certain number of the Tier centres have proven to be able to install and run the necessary hardware. Still, the support model, especially for 24x7 operations, is not well defined. Measures have been developed to monitor the site reliability. The improvement is constant, but the site reliability is not yet at the required level. The measure derived from both basic tests and experiment-specific tests should be published to have a better understanding of the site situation. The CERN Tier-0 centre seems to be well equipped to cope with the expected data rates. Networking is in general adequate.
- Mass storage management and operation remains one of the most critical issues for the LCG. After many delays, the SRM Version 2.2-compliant services are being deployed. The experiments have done only limited tests on the new services. The deployment schedule is extremely tight for the CCRC08. The CASTOR2 system has at the end of October 2007 released a new version fixing many issues, but it remains to be seen whether this new version has the required stability and performance. Deployment at remote sites is essential in order to reduce service difficulties. dCache v1.8 deployment has begun well, but its overall schedule remains very tight. *Since the Comprehensive Review, deployment of SRM Version 2.2 to all Tier-1 and many Tier-2 sites was achieved prior to the start of the February phase of CCRC08. At the beginning of February 2008 there were about 160 sites with an SRM Version 2.2 storage system in place.*
- Basic middleware services are in place and the focus of the past year has been more on stability than on new features. The Enabling Grids for e-Science in Europe (EGEE) Project decided to stop further development of the gLite computing element software and to focus efforts on the Web-services based CREAM computing element. The LHCC is concerned by the announced manpower cuts of European Union funds in the area of middleware development. The real test of the system will only happen with the first data, and it is essential that key experts are retained until then. A continued focus on stability and production quality software is essential for the coming year.
- The Application Area is in general in very good shape. The long-awaited migration from SEAL (Shared Environment for Applications at LHC) to the ROOT interactive tool for analysis is delayed because of unforeseen although not fundamental difficulties. The manpower level is matched to the requirements, but turnover can become a problem if the overlap is not sufficient to ensure proper knowledge transfer. Once again, in the view of the Committee, retention of

key experts until the system has been fully exercised with real data would be very prudent. The 3D distributed database project is well on track.

- The service level has constantly improved and the nominal data transfer through-put rate for 2008 has been achieved, although only for short periods. Sustained, stable operation still needs to be achieved. Planning for the CCRC08 requires a strong coordination among all sites, a robust data transfer and management service, and a clear definition of the criticality of the various services as well as a failure recovery mechanism.
- Resources for the WLCG are provided through a Memorandum of Understanding for which, however, some signatures are still missing. Not all the pledged resources for 2007 have been deployed in time and in particular the scarce disk space has caused significant problems to the production activities. In addition, a large fraction of the Tier-1 and Tier-2 sites have not confirmed the 2008 pledges. It is important that the pledged resources for 2008 are available for the CCRC08 in order to exercise the full system prior to LHC data taking. Up to 2010, the pledged resources match reasonably well the experiments' requirements, except for ALICE, where a 50% short-fall is still observed. The pledged resources for 2011-2012 are significantly below the experiments' requirements.
- Communication with remote sites has improved significantly, and a system to plan and track the progress of each site (especially the Tier-1 centres) has been set up. Many level-1 milestones are not being met, including the 24x7 support, the VO boxes support, and the site reliability. Stronger coordination is required to ensure that the sites are ready for CCRC08 and for data taking. *Since the Comprehensive Review, experience in February 2008 has demonstrated that the preparation and co-ordination has improved and is being done well.*