



Computing Resource Review Board – 25 April 2006

Status of the LCG Project

1. Introduction

This status report covers the period from October 2005 to March 2006. Further details on progress, planning and resources, including detailed quarterly progress reports, can be found in the documents linked to the LCG Planning Page on the web.

The new organisational structure, defined by the Memorandum of Understanding, has been put in place. The first meeting of the Management Board (MB), replacing the Project Execution Board (PEB), took place on 31 October 2005. This extends the executive of the project to include representatives of each of the Tier-1 centres, in addition to the experiments and activity leaders already present in the PEB. The Collaboration Board (CB), with representatives from each of the Tier-1 and Tier-2 sites or federations, as well as CERN and the experiments, held its first meeting on 3 February and elected Neil Geddes (RAL) as chair. The Overview Board (OB), little changed from the previous POB, had its first formal meeting on 20 March.

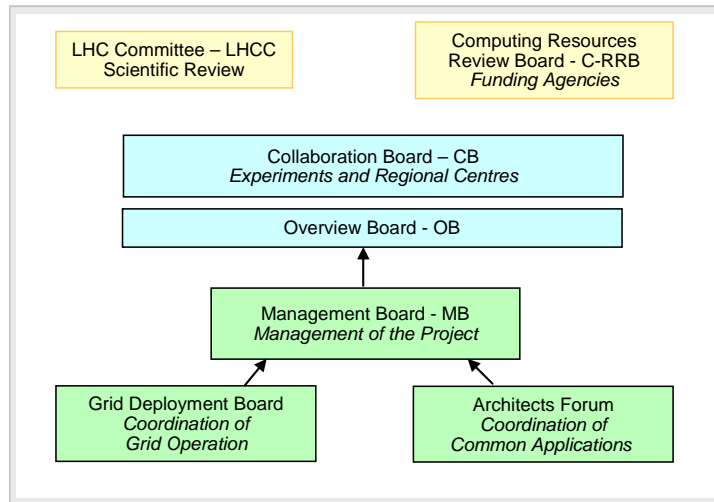


Figure 1: WLCG Organisation in Phase 2

One of the first tasks of the MB, completed before the end of the year, was to agree on the high level plan for Phase 2 of the Project, and the associated process for reporting and monitoring progress. The progress reporting for the last quarter of 2005 already used the new format and procedures. The report now includes the Tier-1 centres, and emphasises the progress in comparison with the detailed medium term planning. It is intended to extend this to include Tier-2 centres and experiment schedules at a later date.

2. Service Challenges

Service Challenge 3 (SC3) completed its formal service phase at the end of 2005. This demonstrated a marked improvement in operational reliability and availability over the situation at the end of 2004, and the formal criteria for the completion milestone were achieved.

All of the Tier-1 sites and over 20 Tier-2 sites took part, operating the agreed set of baseline services. Site reliability improved steadily during the test, and the underlying grid infrastructures showed that they are capable of running more than 20,000 jobs per day, with prolonged periods when over 10,000 simultaneous jobs are executing. During January a test of data distribution from CERN to the Tier-1s took place. The target aggregate data rate of 1 GB/sec copying data from mass storage at CERN to mass storage at the Tier-1 sites was achieved (see Figure 2 Table 1). This represents more than 50% of the nominal LHC data distribution rate for 2008.

In subsequent tests with individual sites, 8 of the 11 Tier-1 sites demonstrated that they can sustain data import rates from CERN to local disks at rates in excess of the nominal data rates for 2008.




Figure 2: Data rates from CERN to Tier-1 sites during one day of the SC3 Data Distribution test

On the other hand, all experiments experienced start-up difficulties, and there was also a general feeling by the experiments that many of the tools available were not at the level of functionality, reliability and performance that they expected. This has led to a much more formal process for planning SC4, including a clearer definition of the software to be delivered, insistence on adequate time for testing, including a 6 week period for experiment *beta* tests, and delaying the introduction of new components until after SC4. The services in SC4 will be an evolution of those available in SC3, with the emphasis on reliability rather than on functionality.

In the final plan for SC4 the service phase is scheduled to begin on 1 June 06, one month later than the schedule at the time of the October 2005 RRB. The new schedule is as follows

- SC4-1 – 15 March – Middleware and services available on the pre-production service for beta testing by the experiments.
- SC4-3.1 – 30 April – Testing completed on pre-production service
- SC4-3.2 – 30 April – SC4 disk-disk throughput tests complete
- SC4-3.3 – 1 June – Middleware and services deployed and in operation at participating sites. Start of SC4 service phase.

3. Level-1 Milestones

The status of Level 1 Milestones due since the last Overview Board meeting is summarised in Of the 10 milestones due, 7 have been completed, 2 have been delayed by one month, and one has been re-scheduled as part of the detailed planning of Service Challenge 4.

4. Applications Area

The work to merge the ROOT and SEAL packages, agreed last year, has made good progress and the migration from the previous set of libraries is under way.

The first public release of the new re-engineered version of the common relational database interface package (CORAL) was made available. The adaptation of POOL (object persistency) and COOL (conditions database) to this new package is ongoing and will be completed soon. This will be in time to be integrated by the experiments in their production software to be used in this year's major data challenges.

End of the year production versions have been released of ROOT and Geant4 which include a long list of new functionality required by the experiments.

. Full milestones tables are available via the LCG [Planning Page](#). Of the 10 milestones due, 7 have been completed, 2 have been delayed by one month, and one has been re-scheduled as part of the detailed planning of Service Challenge 4.

5. Applications Area

The work to merge the ROOT and SEAL packages, agreed last year, has made good progress and the migration from the previous set of libraries is under way.

The first public release of the new re-engineered version of the common relational database interface package (CORAL) was made available. The adaptation of POOL (object persistency) and COOL (conditions database) to this new package is ongoing and will be completed soon. This will be in time to be integrated by the experiments in their production software to be used in this year's major data challenges.

End of the year production versions have been released of ROOT and Geant4 which include a long list of new functionality required by the experiments.

Table 1: Level 1 Milestones due October 2005 - Marh 2006

ID	Date due	Milestone
SC3-2	31.12.05	Service Challenge 3: <i>successful completion of stable service phase</i> Complete - The criteria for success involved achieving certain reliability and availability targets at at least 5 Tier-1s and 5 Tier-2s. All 11 Tier-1s and over 20 Tier-2s took part in the service challenge.
SC3-3	28.02.06	SC-3 data distribution performance tests complete Complete - Achieved the target aggregate data rate from CERN to Tier-1s of 1 GByte/sec, with individual site targets for 5 sites: network→disk data rate of 150 MB/sec, network → tape data rate of 50 MB/sec. 8 sites have shown that they can sustain a disk data rate greater than the nominal data rate from CERN that they will have to support in 2008.
OPN-1	31.12.05	Tier-0/1 high-performance network (10 Gbps) operational at CERN and 3 Tier-1s Completed – Between CERN and FNAL, NIKHEF and IN2P3/Lyon
DRC-2	31.12.05	750 MB/s data recording demonstration at CERN Completed - 950 MB/sec was achieved
SC4-1	28.02.06	All required software for baseline services deployed and operational at all Tier-1s and at least 20 Tier-2 sites Milestone redefined - Schedule for SC4 re-deployment re-defined (see status report).
SC4-2	28.02.06	Use cases and service level support defined for SC4 Completed - outcome of the SC4 planning workshop in Mumbai in early February - document signed off by Management Board on 7 March
CAS-1	15.03.06	Castor2 Readiness Review Delayed - Re-scheduled for end of April

OPN-2	31.03.06	<p>Tier-0/1 high-performance network operational at CERN and 6 Tier-1s, at least 3 via GEANT.</p> <p>Completed - 6 sites connected - CNAF and FZK via GEANT, and FNAL, SARA, IN2P3, BNL via other links</p>
SC3-4	31.03.06	<p>All services on all Tier-1 sites monitored</p> <p>Delayed - Monitoring framework in operation with basic tests at all EGEE sites. Date for completion of full test set slipped to end April.</p>
SC3-5	31.03.06	<p>Proposal on availability levels specified in Annex 3 of the WLCG MoU (adjusted for sites that do not provide a 24 hour service)</p> <p>Completed – Agreement on initial method for measuring availability at MB meeting on 7 March.</p>

6. CASTOR 2 & Tier-0 Testing

All experiments are now using the new version of the mass storage system used at CERN, CASTOR, which has been re-developed over the past three years to be able to support the LHC data rates at CERN. High performance data recording tests at CERN at the end of the year demonstrated that the new version of CASTOR can already handle 50% of the nominal LHC magnetic tape recording requirement on a limited configuration. Further tests in January demonstrated disk storage throughput close to the requirement for 2008.

In more realistic conditions, tests carried out in collaboration with ATLAS achieved during a short period their target DAQ-Tier-0-recording-reconstruction data rates driven by 320 MB/sec from the DAQ, as shown in Figure 3. The test did not include distribution of data to Tier-1s.

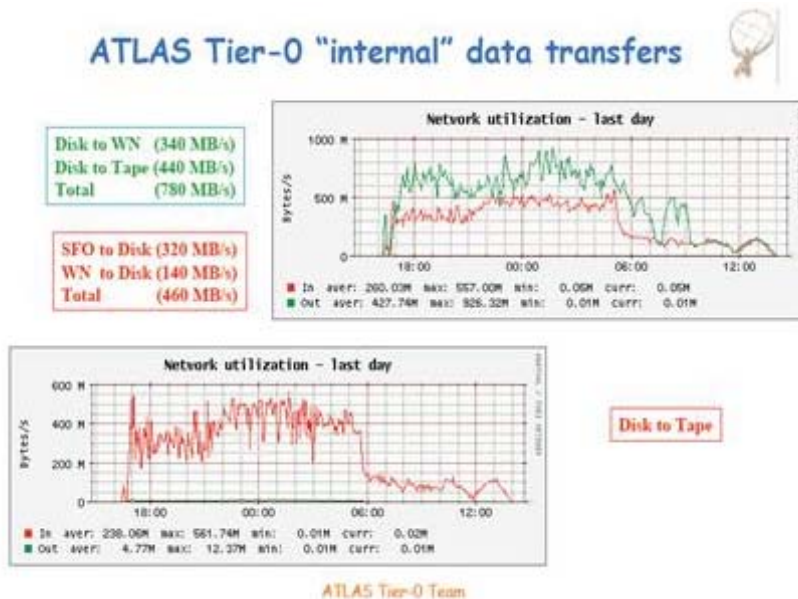


Figure 3: ATLAS Tier-0 Tests - January 2006

Deployment of the new version of CASTOR at the 4 Tier-1 sites using CASTOR is under way. CASTOR is developed at CERN, with Rutherford Lab providing the standard interface used for

grid access (SRM). The remaining 7 Tier-1 sites use the dCache mass storage system developed at DESY, with FNAL providing the SRM interface.

7. Distributed Database Deployment – 3D

An agreement on the architecture and technology for the distributed database services was reached at the 3D Project Workshop in October 2005, and the testing and deployment plan was finalised before the end of the year.

The development and deployment of the distributed database services at the Tier-1 sites will take place during 2006, in parallel with the SC4 service. The plan foresees two production phases:

- April - Sept 2006 : partial production service in parallel with the testbed, and including a subset of the Tier 1 sites: ASGC, CERN, BNL, CNAF, GridKA, IN2P3, RAL
- October '06- onwards : full production service, including the remaining tier-1 sites: PIC, NIKHEF, NDGF, TRIUMF

The timescale is tight, with the conditions database as the initial application, but so far the project is keeping to schedule..

8. SRM

An important step in completing the base storage services for LHC startup is the evolution from the current version of the Storage Resource Manager standard to a newer version (2.2) which has additional features of importance to the experiments. There are several technical issues to be resolved, and an agreement has to be made on the different classes of storage that must be supported for the experiments. As it is essential that the SC4 service phase (June through September) should be stable, it was agreed to aim for a target date for deploying SRM 2.2 at the end of that period.

9. Accounting

The accounting system has been in production since the beginning of 2005. Software is provided to filter accounting data from the compute element (CE) logs and route it to a central repository via RGMA. Data can also be submitted by other means.

The system is now seeing about 20K jobs per day from the total of 190 sites (CEs) that are reporting (see Figure 4).

There are several gaps in the accounting, even from some of the Tier-1 centres, and at present most sites only report work submitted through the grid (i.e. they do not yet include work submitted locally, direct to the batch scheduler). The LCG Overview Board has requested that all sites participating in SC4 report regularly on all usage by LHC experiments, and the results will be reported to the October RRB.





Figure 4: Accounting records and reporting sites

10. Interoperation Update – EGEE-OSG-NDGF

There has been progress in interoperability between OSG and EGEE, each grid now making its LHC resources visible to the other grid through the information service (BDII). This enables job

submission in both directions. Closer collaboration and compatibility on operational procedures is under discussion. There has also been progress in interoperability with NDGF, which now also publishes its resources via BDII.