



Computing Resource Review Board - 18 October 2005

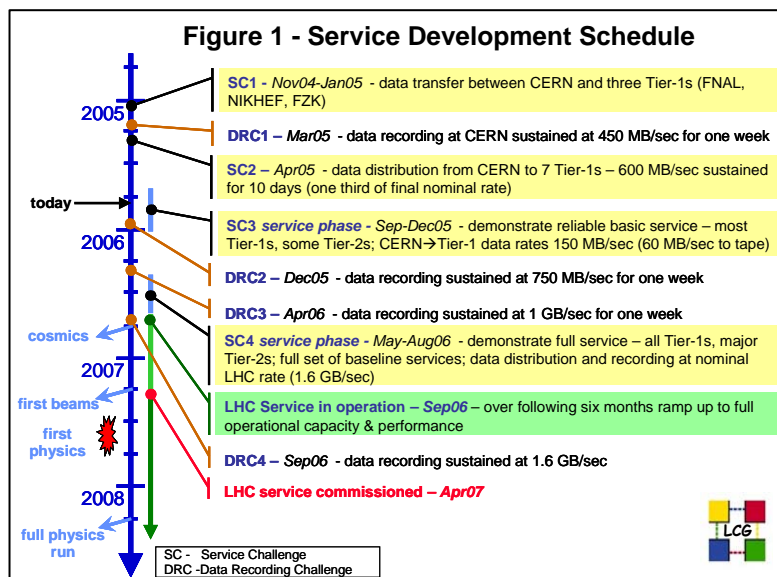
Status of the LCG Project

Introduction

This status report covers the period from April to October 2005. 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 [Technical Design Report](#) (TDR) for the initial LHC computing services was published on 20 June as scheduled. It will be reviewed along with the computing TDRs of the four experiments by the LHCC on October 7-8.

The high level plan of service and data recording challenges leading to the opening of the LCG service is shown in Figure 1. More detailed planning for Phase 2 is now being developed. The long-term plan for the Applications Area has been agreed by the Project Execution Board (PEB), as have the plans at CERN for migration to the new version of the CASTOR mass storage management system and the CERN fabric acquisition and installation programme. The Phase 2 Planning Group has completed another round of resource planning in preparation for this C-RRB. Detailed



planning for Service Challenge 4 (January-August 2006) is under way, which will for the first time include formal milestones for certain aspects of services in regional centres.

Preparations have been made to move rapidly to the new management organisation following the agreement on the MoU anticipated at this C-RRB. The first meeting of the new Management Board, which replaces the PEB and includes members from each of the Tier-1 centres, is scheduled for the first week in November, and the first meeting of the Collaboration Board is being organised for January 2006.

Applications

A detailed [plan](#) for the Applications Area in Phase 2 of LCG has been agreed.

It was decided to proceed with the option mentioned in the report to the April C-RRB of merging the ROOT analysis system with the tools and components developed by the Applications Area as the SEAL activity. While in the longer term we expect an evolution towards a single set of core software libraries, the current components used by the experiment frameworks will be maintained until they are no longer needed. However, for a number of functional components, which have been agreed by the experiments, the migration has already started. A first set of software packages from the SEAL project providing

mathematical and dictionary functionality have been moved to the code repository of the ROOT project and have been released as part of the last two development releases of version 5 in June and September.

The POOL (object persistency) project is making good progress on various fronts, including: performance optimization (collaboration with CMS), delivering the production quality conditions database (being deployed in ATLAS and LHCb), and migration to a new interface to the SEAL dictionary (Reflex).

The Simulation project is also continuing with good progress in various fronts. There has been a new release of Geant4 in June, which introduces refinements to ionization processes and includes a full implementation of gamma processes in the model and of a materials builder using NIST data. FLUKA has made a public release in July, which is the first release carried out under the INFN-CERN collaboration agreement for the maintenance and development of the FLUKA code, and is currently finishing the documentation to be published as a CERN yellow report. The physics validation is making progress with the calorimeter test-beam data and the radiation background studies. The Monte Carlo generator service has released a new version of the library in July. Some new utilities in the simulation framework have been developed such as geometry persistency.

Planning Phase 2 for the applications area has taken quite lot of effort during the last months. Finally a [detailed plan](#) has recently been agreed by all parties.

Grid Deployment

The "Baseline Services" to be provided at CERN, Tier-1 and Tier-2 centres at LHC start-up were agreed in June. Initial versions of most of these services will already be available at the sites participating in Service Challenge 3 (SC3) this autumn, and the final versions must be ready for the start of the service phase of Service Challenge 4 (SC4) in May 2006.

For the EGEE Grid the LCG 2.6 release was made in August and has already been installed by 70% of the 160 sites. This is a pre-requisite for the service phase of SC3, and also contains a large number of bug fixes. The previous release (LCG 2.5) had been installed only at a restricted number of sites involved in the SC3 setup phase.

A mechanism is now in place to enable the client tools to be updated without the site having to do a full installation. In fact the client tools can now be updated remotely by the deployment team at CERN (in the same way as applications software).

The File Transfer Service is showing excellent reliability, reducing the hard failure rate for file transfers in SC3 to below 1%.

The operations management activity started after the November workshop at CERN is beginning to bear fruit. The System Functional Test (SFT) system is fully operational, allowing each Virtual Organisation (experiment) to select the set of tests needed to certify a site. The test framework can be used to update automatically the information service to enable the workload management system to use the latest testing status for each site. More than 120 sites now pass the full set of tests, representing over 10,000 processors. A second operations workshop was held at CNAF in Bologna from 24-26 May, including people from the US Open Science Grid (OSG) operations centre in Indiana. A third operations workshop, organised jointly by EGEE and OSG, took place at Culham (UK) from 27-29 September.

Release 1.2 of the complete EGEE middleware package (gLite) is installed on the pre-production service, and CMS and ALICE have begun to use the service. Components are upgraded to the later gLite 1.3 version as they pass the certification process.

There has been good progress on inter-operability with the OSG infrastructure in the US, and it is now possible to submit work from the EGEE Resource Broker to some OSG sites, and vice versa. There has also been an agreement on operational monitoring which should enable the operations centres in each grid to perform basic surveillance of sites in the other grid. Both the Open Science Grid and EGEE have made statements of commitment to continuing work on inter-operability.

Discussions have continued with the Nordic Data Grid Facility (NDGF) to understand how to implement inter-working with EGEE, including a workshop at CERN at the beginning of September.

An [LHC view](#) of the accounting information showing grid usage by experiment is now available on the Grid Operations Centre web site. Reporting is still not yet complete, delayed in some cases by legal difficulties in providing the detailed information.

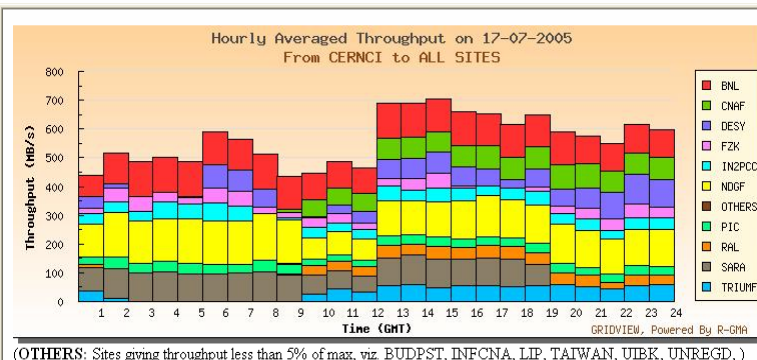
EGEE Phase 2

The proposal for Phase 2 of the EGEE project has been completed. This responds to several of the concerns expressed during preparatory meetings between the EGEE management and the experiments: the applications support for HEP will be extended to cover the full range of activities that interface to the grid – production as well as analysis; an integration, testing and deployment activity has been defined, closely linked to the deployment and operations activity, that should ensure closer cooperation between middleware developers and users in the experiments; the experiments will be represented directly in the Technical Coordination Group, which is expected to be the place where decisions on the evolution of middleware developments will be prepared. If approved by the European Commission, Phase 2 of the EGEE project will operate from April 2006 to March 2008.

Service Challenges

The setup and throughput phase of Service Challenge 3 was completed during August. The throughput test – CERN to the Tier-1s - used the new gLite File Transfer System (FTS). FTS itself was quite stable but the service throughput and reliability objectives were not fully achieved. The sustained aggregate throughput achieved was 500 MBytes/sec compared with the target 1 GByte/sec, and many Tier-1s did not reach the sustained 150 MByte/sec target. Subsequent work has been undertaken to investigate the problems at specific sites with promising results. 15 sites took part in the throughput test, which will be repeated later in the year when the investigative work has been completed. This does not affect directly the service phase of SC3, where the aggregate data rate from CERN is limited to 200 MBytes/sec.

Analysis of GridFTP Log for Service Challenge 3



The service phase began on 1 September with ALICE, with CMS beginning in mid-September. The other experiments will begin according to an agreed schedule.

ARDA

The activity of the LCG ARDA group focuses on the development of four analysis prototypes, one agreed with each experiment. Although these developments are independent, the evolution of the four prototypes benefits from the interchange of information inside the ARDA team. The work with ALICE is concerned with delivering services to allow fast access to grid resources from the ALICE framework. The activity with ATLAS has provided different prototypes integrating components developed inside the experiment, and in particular using the *GANGA* user-grid interface to execute analysis jobs. With CMS the activity has produced a full working prototype that is now being evolved and integrated with CMS analysis components. Within LHCb the work has involved participation in the development of the *GANGA* system. The ARDA team also provides tutorials, support for pilot users and participates in testing of the EGEE

gLite middleware package. ARDA has developed a common metadata system which will be soon available with gLite package.

CERN Fabric

The strategy for building up the magnetic tape service to handle the LHC data recording rates has been agreed and the acquisition process was approved at the September meeting of the Finance Committee (CERN/FC/4965). In order to maximise performance and media capacity, while minimising cost, the selection of hardware will be delayed for as long as possible. However, past experience shows that, since the HEP environment differs significantly from the commercial environment for which tape drives are optimised, extensive testing of equipment in the HEP environment is required. In order to balance the need to delay the choice of tape drives and the need to provide a tape subsystem with the demonstrated ability to meet the requirements for LHC data taking, equipment from two competitive suppliers will be installed under a rental contract at the beginning of 2006 and evaluated in a production environment during Service Challenge 4. A final decision on the hardware to be purchased will be taken in September 2006.

The new electrical installation in the Computer Centre is being brought into operation and the first stage of the uninterruptible power supply (UPS) commissioned. There are continuing delays in the installation of the cooling and ventilation upgrade, and we are now approaching the point at which this may impact other aspects of the service schedule.

A detailed schedule for the acquisition of processors and disk equipment has been prepared based on a repetitive tender cycle selecting always two independent suppliers. The strategy, constrained by CERN's purchasing rules, aims to minimise the risks inherent in dealing with relatively small suppliers while delaying purchasing decisions as long as possible to benefit from improved technology and price evolution.

Networking

The tendering process for the internal network equipment for the CERN facility was completed on schedule and the first components are already installed. The network working group of the LCG Grid Deployment Board, including the major regional centres and the national and regional research and education network providers (NRENs), has reached a satisfactory agreement on an overall architecture for the high performance network interconnecting the Tier-1 centres and CERN. Implementation of the network is in most cases the responsibility of the research network organisations, independent of the LCG structure. Collaboration has been excellent and at present we do not anticipate problems in meeting the LCG schedule. In particular the work has benefited from close cooperation with the people implementing a major upgrade to the European research network backbone, GÉANT.

The networking between Tier-2 and Tier-1 centres is assumed to be satisfied by the "normal" network connectivity of the centres concerned, provided by the local NREN. While this is a reasonable assumption for start-up, based on the requirements documented in the experiments' Computing TDRs, in the longer term network connectivity of Tier-2s is likely to be a significant factor in the way in which end-user analysis is performed.