



Computing Resource Review Board LHC Computing Grid Project

Version 2

8 April 2004

Project Status Report

Les Robertson /CERN

General

This status report covers the six-month period from the last C-RRB in October 2003. Information on resources is available in related documents provided to the Computing RRB. Further details on progress and planning can be found in the documents linked to the LCG Planning Page on the web (<http://www.cern.ch/lcg/peb/planning>).

A revised structure of the principal management committees of the project – the Project Execution Board (PEB) and the Software and Computing Committee (SC2) – was agreed by the Project Oversight Board in November 2003 and implemented at the end of the year. The major change is that the experiment computing coordinators become members of the PEB, which now takes on the added role of proposing changes in the scope of the project. The coordinators are replaced in the SC2 by physicists who are not directly involved in the computing projects of the experiments. The SC2 retains its role as a permanent monitoring and review body and also approves work plans and changes in the project scope. The new structure is documented on the LCG web (<http://lcg.web.cern.ch/LCG/LCGProjectStructure.htm>).

A revised set of Level 1 milestones was agreed with the LHCC referees in March.

A two-day LCG workshop was held at CERN in March, covering the fabric, grid deployment, middleware and distributed analysis areas.

Estimates have been prepared of the resources required at the Tier-0 and the Tier-1 centres in 2008, together with associated networking requirements, as information for the MoU Task Force.

Grid Deployment

By the end of 2003 four grid services were delivering production capacity to LHC experiments: the LCG-1 service, opened in September 2003, was used by CMS Italy in December; US CMS was able to obtain substantial resources incorporated in the Grid 3 facility from November; Nordugrid was in use by ATLAS; the AliEn software integrated computational resources from a number of centres for ALICE. The data challenges scheduled for 2004 will extend the usage of grids beyond simulation to include significant data handling, begin to test the emerging *computing models* of the experiments, and gain experience in managing and operating a reliable grid service.

Deployment of a new release of the common LCG grid software package, LCG-2, began in February, initially to a core set of sites, and the ALICE and CMS experiments began to use this in their data challenges from the beginning of March. After resolving initial problems the grid is performing with a reasonable level of reliability. The principal problems remaining concern data handling performance and scaling. Some of these will take time to be resolved, but at present there are no critical problems outstanding.

ATLAS and LHCb are scheduled to begin their data challenges using the LCG services in May.

The sites currently participating in the LCG service are shown in the table, together providing more than 2,000 processors for LHC experiments. Mass storage services are also available at core sites including CERN, CNAF, FNAL, FZK, NIKHEF, PIC and RAL.

Regional Centres Connected to the LCG Grid - 6 April 2004			
country	centre	country	centre
Canada	TRIUMF, Vancouver	Spain	PIC, Barcelona
Czech Republic	Prague		IFIC, Valencia
France	IN2P3, Lyon		IFCA, Santander
Germany	FZK, Karlsruhe		University of Barcelona
	DESY		University of Santiago de Compostela
Holland	NIKHEF, Amsterdam		CIEMAT, Madrid
Hungary	Budapest		UAM, Madrid
Italy	CNAF, Bologna	Switzerland	CERN
	INFN, Torino		CSCS, Manno
	INFN, Milano	Taiwan	Academia Sinica, Taipei
	INFN, Roma		NCU, Taipei
	INFN, Legnaro	UK	RAL
Japan	ICEPP, Tokyo		Cavendish, Cambridge
Poland	Krakow		Imperial, London
Russia	SINP, Moscow	USA	FNAL
			BNL

The Operations Centre at RAL is now active, with a second centre at the Academia Sinica in Taipei coming online. The grid call centre at FZK has been exercised for local support and will be used more widely over the coming months. Practical experience of grid operation will take time to acquire, and full operational status will not be achieved before the end of this year.

The LCG-2 software includes the Virtual Data Toolkit (VDT) developed by a number of US projects, and supported for LCG with NSF funding. VDT is used by the Grid 3 facility in the US. There are proposals to extend Grid 3 as the Open Science Grid (OSG), which would include all US Tier-1 and Tier-2 centres. It will be important for the LHC experiments that they are able to access these US and non-US facilities in an easy and coherent manner, and discussions have been started on how to achieve this at both the technical and management levels.

It is foreseen that the resources available to ALICE through AliEn will be integrated progressively in the LCG-2 facility, transparently to ALICE users. The situation with integration of the NorduGrid resources used by ATLAS has not yet been discussed.

The *Enabling Grids for E-Science in Europe* (EGEE) project started on 1 April, including 70 partners in Europe and the United States, most of which are also involved in the LCG Project. EGEE will provide important resources for operating and managing grid services in Europe for High Energy Physics and other sciences. As explained in the October 2003 Status Report, the LCG-2 grid will be the starting point for the EGEE facility. The details of the joint operation and management of these intersecting grids will be worked out in the coming months.

Hewlett-Packard will make resources available for LHC computing through the LCG service. The first of these systems is being set up at their facility in Puerto Rico.

As part of the grid deployment activity, regional centres have agreed to provide support for the following components across LCG.

- Grid Operations Centre - RAL, AS Taipei
- Global Grid User Support Centre - FZK, AS Taipei
- Virtual Data Toolkit (US tools) - VDT team at Wisconsin (NSF funding)
- DataGrid resource broker - INFN/CERN
- DataGrid replica management - CERN
- DataGrid relational information system (RGMA) - RAL
- GridIce monitoring tools - INFN
- GLUE schema - INFN

- Virtual Organisation Mgt.System - INFN
- Security & VO policies and procedures - Security Group + all countries

Fabric and Networking

In October 2003 CERN and CALTECH demonstrated a record data transfer rate of 5.44 Gbps between Switzerland and California, sustained for over half an hour, similar to the effective data rates that will be needed between CERN and each Tier-1 Centre when LHC begins operation. We must now follow this up with a series of tests leading to a pilot service for high bandwidth Wide Area Networking between CERN and Tier-1 centres, as a joint activity with the national and regional research network providers. First steps in this were taken during the 1st International Grid Networking Workshop (GNEW2004) which was held at CERN in March, co-organised by CERN/DataTAG, DANTE, ESnet, Internet2 and TERENA.

The next step is to include the full chain of data distribution software, moving data between mass storage systems at high data rates. A computing data challenge operated by US CMS and CERN IT Department is starting now to demonstrate this between mass storage systems at FNAL and CERN. Further challenges will be planned as we build up to sustained operation of data transfer services between CERN and Tier-1 centres by the end of 2006.

The 2003 ALICE mass storage data challenges to test high performance data recording took place between November and January. While the individual performance targets for disk and tape were met, the tests for the full system, from the DAQ system through the High Level Trigger to mass storage, achieved only 280 MB/sec, somewhat short of the target of 300 MB/sec. The problems, principally due to the network, are understood and it is expected that this year's target of 450 MB/sec will be achieved.

The online projects of ALICE and ATLAS have both used the offline computing installation at CERN for scaling tests of their filter farms. While it has in some years been possible to fit them in to the offline schedule this is not always possible (as in the case of ALICE last year and ATLAS this year). It must be emphasised that the LCG facility at CERN is not funded or sized to support these tests.

The first version of the full suite of cluster management software – ELFms – is now in production in the CERN computer centre. Some of this technology has been adapted from the tools developed by the fabric management work package of DataGrid. The centre is now well placed to manage the configuration for LHC using the reduced staffing levels for system administration agreed last year.

The civil engineering for the power and cooling upgrade is complete and the installation of the ventilation and electrical distribution equipment is proceeding according to plan.

The first steps in the process to acquire the LCG Phase 2 equipment at CERN have begun. A series of technology choices and market surveys will be made over the next twelve months.

An agreement was reached between CERN, CNAF and PIC for the support of the full Castor mass storage management software at regional centres in Italy and Spain. DESY, in collaboration with FNAL, has agreed to support the dCache disk management package at LCG centres.

Applications

The POOL object persistency system is now integrated into the applications of three of the experiments, has been validated in production usage with millions of events, and will be used as the production event store in this year's data challenges.

The comprehensive software development infrastructure set up to satisfy the needs of the LCG Applications Area will be used to support also the middleware development of EGEE.

Important steps have been taken in the work on simulation physics validation. A first round of Geant4 electro-magnetic and hadronic physics validation has been completed, concluding that Geant 4 is "as good as or better than Geant3". The simulation physics requirements of the four experiments have been documented, and there is a good collaboration on validation work with both Geant4 and FLUKA.

The MC generator library GENSER has been developed, populated with the requested priority generators, and is being evaluated/adopted by experiments (ATLAS has fully migrated to it). There was a workshop at CERN on 25 March to discuss the future requirements and directions for support of MC generators for LHC.

The CERN Geant4 program is now squarely focused on LHC priorities and Geant4 has been successfully deployed in production in CMS and in pre-production in ATLAS.

Closer integration of the ROOT data analysis facility and the other components of the Applications Area is under discussion. ROOT provides the data store technology for POOL, and the ROOT analysis environment can in some cases be used either directly or via interfaces. Important next steps are the use of a common dictionary and mathematical library.

The adoption by the experiments of SEAL framework services is going slower than expected. This is now planned by two experiments (LHCb, ATLAS) after the data challenges have been completed.

Important data management developments planned for 2004 are conditions data storage via a common conditions database and POOL, a relational back-end to POOL and physicist-level event collections.

Most of the development programme foreseen for the Applications Area, and integration of these into the experiments, should be completed in 12-18 months. The emphasis will then move to supporting the ever increasing scale and complexity of the environments in which the software is used in the experiments as they scale up towards data taking.

Distributed Analysis

A workshop was held at CERN in January 2004 to discuss the proposals made by the RTAG on distributed analysis. Taking account of the views expressed in the workshop the PEB launched a new project in February called ARDA (see <http://www.cern.ch/lcg/peb/arda>). The project will work with teams within the experiments that are involved in prototyping distributed analysis systems, helping them to interface to grid services and coordinating between the experiments, the middleware developers and regional centres involved in distributed analysis prototyping. The project has the status of an "area" within LCG, the project leader (Massimo Lamanna) being a member of the PEB.

The ARDA project includes the HEP part of the EGEE applications activity. The project plan will be discussed in the PEB in May.

Grid Middleware

With the end of the DataGrid project a re-assignment of responsibilities for support of the DataGrid components included in LCG-2 was necessary. The first line support for these is now the responsibility of the Grid Deployment Area, which has control of the source code. It has been agreed with INFN, RAL and CERN that the original authors will retain a second level support role. The grid tools from the US projects receive first line support from the VDT team at the University of Wisconsin, with NSF support.

The EGEE project has a significant activity for middleware re-engineering and development. As explained in the status report to the October 2003 C-RRB, the EGEE middleware activity has a secondary reporting line to the LCG PEB. The middleware team that has been put together by the EGEE partners includes experts from the DataGrid, AliEn and VDT projects.

The middleware activity will work closely with the ARDA project. It is expected to provide first components within the next month.

It should be emphasized that the LCG-2 tools will remain the primary software supported for the data challenges of this year, with appropriate support for maintenance and any necessary development.