LHC Computing Grid Project – Level 1 Milestones

18 November 2002

This note defines the proposed level 1 milestones for the LHC Computing Grid Project (LCG), prepared for a discussion with the LHCC referees on 25 November 2002. At the meeting draft level 2 milestones, extracted from the detailed project planning information that is in preparation, will be presented. When the detailed planning has been completed and endorsed by the SC2 committee, the level 2 milestones will be delivered to the LHCC referees so that the progress of the project towards the level 1 milestones can be monitored.

Level 1 Milestones

The high-level Project Breakdown Structure, which is reflected in the management organisation, is as follows:

1. Applications
2. Fabric management (Fabric)
3. Grid Technology (GT)
4. Grid Deployment (GD)

General information on these areas is given in CERN-LHCC-2002-022 *(LHC Computing Grid Project (LCG) Phase 1 - Status of High-Level Planning, version 1, 21 June 2002).*

The Fabric, Grid Technology and Grid Deployment areas of the project are closely coupled in providing the grid service that the project will develop, and so are synchronised around the level 1 milestones. The applications area is to a large extent independent of the other area and has separate level 1 milestones.

Ref.	*milestone description*	*target date*
M1.1	**First Global Service (LCG-1) - Initial Availability** This comprises the construction and commissioning of the first LHC Computing service suitable for physics usage. The service must offer reliably 24x7 availability to all four LHC experiments and include some ten Regional Centers from Europe, North America and Asia. The milestone includes delivery of the associated Technical Design, containing description of the architecture and functionality and quantified technical specifications of performance (capacity, throughput, reliability, availability). It must also include middleware specifications, agreed as a common toolkit by Europe and US. The service must prove functional, providing a batch service for event production and analysis of the simulated data set. For the milestone to be met, operation must be sustained reliably during a 7 day period; stress tests and user productions will be executed, with a failure rate below 1%.	July 2003

Ref.	*milestone description*	*target date*
M1.2	**General release of POOL hybrid data store** The first public, production-capable release of the persistency framework. This will be a release offering basic hybrid persistency services, documented and packaged using the SPI-defined templates and tools, for general use by the experiments in production environments. The release should support production usages with O(1M) file counts, O(50TB) data volumes, distributed operation at O(10) sites, and with O(10k) populating jobs. Previous releases were internal releases targeted at developers and experts; this is the first public release with the robustness, documentation, packaging and support requirements inherent in a public release. Specific feature set for this release defined by a milestone four months earlier**.**	June 2003
M1.3	**Distributed production environment using grid services** A distributed production environment integrating experiment specific software with common software components and services, grid portal services and grid middleware. The deliverables of this milestone will be elaborated when a Physics Interfaces work plan is developed following the conclusion of the analysis RTAG.	Nov 2003
M1.4	**Fully operational LCG-1 Service** This comprises the availability of LCG-1 as a fully operational and performant 24x7 production service. Operation must be sustained for a period of one month. This service would be used for the "5% data challenges" of the LHC experiments. LCG-1 will be operated continuously, evolving in terms of capacity, performance and functionality. It includes the addition of Regional Centres as they come on-line as defined in GDB Working Group 2. It includes the delivery of the technical service specifications and user documentation, and deployment/consolidation of an appropriate user support infrastructure. It also includes incremental releases of middleware to improve reliability, robustness and performance The service level must be as required for the 2004 data challenges. The determination and acceptance of the milestone should be done with a formal review of the service by representatives of the experiments, regional centres and LCG.	Nov 2003
M1.5	**Distributed end-user interactive analysis from a Tier 3 Regional Centre** Extension of the distributed production environment integrating experiment specific software with common software components and services, grid portal services and grid middleware. This extension will support the analysis environment, enabling distributed end-user interactive analysis down to Tier 3 in the LHC grid. A related level-2 milestone, 6 months before the due date, will provide a full specification of what this milestone includes.	May 2004

Ref.	*milestone description*	*target date*
M1.6	**Fully operational LCG-3 Service** This comprises the construction and commissioning of a fully operational full-size prototype (LCG-3) of what it will be the initial LHC Computing production service. Operation must be sustained 24x7 reliably for a period of one month for the milestone to be met. LCG-3 will be used as a proof that the LHC computing model will work, including Tier 0, 1, 2 and 3 Regional Centres, providing practical backup for the computing service TDR. LCG-3 will use the LHC Grid Toolkit, will have 50% of the components required for the 2007 production service for CMS or ATLAS, and will be used for the "20% milestones" of the experiments.	January 2005
M1.7	**Full function release of persistency framework** Completion of the fully functional POOL persistency framework. Deliverables in terms of feature set, performance and scalability for this milestone will be finalized as an outcome of the Computing TDRs of the experiments in 2004.	March 2005
M1.8	**Completion of the Computing Service TDR** The Computing Service TDR will specify the requirements for the Grid that will be used for the first production services for the four LHC experiments. It will include details of the architecture, functionality, capacity, performance, throughput and availability. It will include the Regional Centre plans that will have been developed to meet these requirements, and will provide cost estimates and an overall installation and verification schedule. It is assumed that the TDR will be approved by the LHCC within three months following its availability, and may be used to provide data for the Memorandum of Understanding for Phase 2 of the project. The full process from acquisition to service verification is expected to take 12-18 months (according to the administrative procedures of the Regional Centres). The initial service must be in full production by September 2006 (6 months before data taking). The TDR will therefore be approved after the acquisition procedures have started, but before orders are placed.	June 2005