CERN-LCG-PEB-2004-14
editor: Fabienne Baud-Lavigne
02 August 2004

LHC Computing Grid Project

Quarterly Status and Progress Report

Second Quarter 2004

Table of Contents

1. Applications Area –Torre Wenaus 1
2. CERN Fabric Area – Bernd Panzer 10
3. Grid Deployment Area – Ian Bird 13
4. Middleware – Frédéric Hemmer 16
5. Distributed Analysis – Massimo Lamanna 18
6. ALICE – Federico Carminati 25
7. ATLAS – Dario Barberis 27
8. CMS – David Stickland 29
9. LHCb – Philippe Charpentier 31
10. Milestone management and status – Jürgen Knobloch 33
11. Resources 36
Appendix: Milestone Tables 39

1. Applications Area –Torre Wenaus

In this quarter support for experiment deployment and production usage of AA software continued to be a top priority as all experiments conducted data challenges and production operations during the quarter. CMS, LHCb and ATLAS all successfully had AA software in DC production. GENSER was put into production for the first time, by ATLAS. Important development efforts also progressed: POOL relational DB support, the new LCG dictionary (now named Reflex), mathlib, and MCDB among others. SEAL made important progress in documentation, releasing a first version of their workbook. The collaboration with ROOT proceeded well on several fronts, and the ROOT and POOL teams worked closely to plan a smooth transition to a new major release of ROOT, version 4, in the next quarter.

Manpower was stable for the most part during the quarter but the end of LCG Phase 1 draws closer. LCG, PH/SFT and IT/DB made substantial progress during the quarter in identifying the essential resource needs to sustain the support and development of AA software beyond LCG Phase 1, and worked also at seeking out the required support. A level 1 milestone in the fall was established to have a resourced plan in place for the coming years.

Milestone performance during the quarter was fairly good. Seven level 2 milestones were completed and a further one will be met within ~2 weeks of due date. Four will be late (one for reasons beyond LCG AA control) and should be completed in the next quarter; they have no significant schedule impact. Specific attention is being given to accelerating progress with the provision of a generic Geant4/FLUGG/FLUKA infrastructure which is required for test beam physics validation studies.

WBS 1.1.1 - Software process and infrastructure (SPI)

Project Leader: Alberto Aimar, CERN IT/DI

In this quarter SPI focused on continuing to support the existing services, such as the Savannah portal, the external software service, and software testing frameworks. With some delay, SPI is preparing a collaborative work plan between LCG and EGEE. This is now scheduled for end of July. The delay is mostly due to the fact that the EGEE project in this quarter was starting up and its needs were still being defined and resources assigned. Now all information needed is becoming available for an effective collaborative work in the coming quarter. The SPI project continues to be staffed with a total of ~5 FTE, maintained constant through the period, some new effort balancing some reductions. As recommended by the LCG Internal review, SPI with the support of EP/SFT has put in place an LCG librarian who will be in charge of defining and executing common release, build and distribution procedures for all LCG software, in close coordination with the AA projects and the experiments. Milestone performance during the quarter was good: the gcc 3.2.3 compiler port was completed on time, and the new Linux certification platform port (due July 1) was completed 30 days early. The joint workplan with EGEE is late as explained below, due to delays on the EGEE side.

Savannah service: The web-based project portal based on Savannah is now in use by more than 90 projects in the LHC experiments and in other CERN divisions. New projects came from the LHC experiments and from the EGEE. After the deployment of the new version of Savannah at CERN with all the improvements, done in collaboration with the open source developers, SPI focuses on bug fixing and on providing support for the several new projects and users of the service. The registered users are now more than 900 and the service is stable and established among the users' community.

External software service: SPI continued to maintain and improve the external software service. Now there are about 60 packages and ~400 installation configurations. All the LCG software has been ported to and tested on the Scientific Linux platform derived from Red Hat Enterprise 3 that is targeted as the next standard Linux at CERN. Following the needs of the experiments as mandated in the AF the software is currently being ported to Mac OSX, and to new versions of operating systems and compilers (gcc 3.2.3). An effort started in the previous quarter to automate the compilation and deployment of external software continued to be a focus of activity, in order to support the growing number of tools, versions and platforms with minimal manpower.

Software download service: SPI continues to provide a service to allow download and installation on a local computer of the LCG software. After experience with a prototype in the previous quarter, SPI developed a system based on the pacman distribution system and is discussing with the tool authors in order to establish a service based on pacman.

Build and librarian service: A software librarian joined the project in order to improve and automate the service and develop a long term strategy for building, releasing and distributing LCG software. As requested by the internal review, SPI started to prepare for solutions that can both integrate with specific build systems (Scram and CMT) or generate build files based on standard "make and config" solutions. Also the work on a CMT configuration for LCG software is being finalized. Scram support continues at the same level as recent quarters: SPI provides the central configuration and LCG user support.

Testing frameworks: The service is focusing on user support and porting to new platforms (Mac OSX, gcc 3.2.3, etc). In the next quarter there will be the planning of common activities with EGEE; additional manpower for this activity is now due for August.

Policies and QA: In this quarter the improvement of automated QA report generation was slowed down due to the departure of the persons responsible and the delay of the arrival of the new person in charge of the service. This is now expected for August and therefore QA activities will increase in the next quarter.

CVS service: Since the IT/CVS migration last quarter the CVS service is no longer a responsibility of SPI. SPI does provide first line support to LCG developers and users.

Code documentation service: This service is stable. There were minor changes in this quarter in order to adapt it to the IT CVS service.

Coming activities: The role of software librarian is being defined and will fully start, with the task of providing build infrastructure and support for the LCG projects. Next quarter will also continue to be devoted to automating the tasks of external software installation and distribution. Work will be started in consolidating QA and test reporting and verification, as soon as new resources join the project in August.

WBS 1.1.2 - Persistency framework (POOL and conditions DB)

Project Leader: Dirk Duellmann, CERN IT/DB

POOL
The POOL project produced several production releases up to the current production version 1.6.5. The main emphasis was on supporting the experiments data challenge activities and preparations. The CMS DC04 has been passed successfully without major problems caused by the POOL software. POOL added support for new LCG RLS versions and analysed performance problems with this catalog backend. In particular the use of meta data attributes turned out to be a limiting factor and POOL reacted by temporarily removing its internal use of this feature until these performance problems have been resolved.

During the last quarter POOL prepared with the help of the ROOT team the move to the upcoming ROOT4 release. A large fraction of the code changes have been introduced, but a consistent test of the new ROOT-based streaming (in particular for STL containers) with respect to updates of ROOT3 data file and schema evolution is still outstanding and requires experiment validation.
The first POOL release based on ROOT4 (named POOL 2.0) will be produced soon after POOL 1.7. The POOL 1.x and 2.x lines will be maintained in parallel until successful validation allows dropping the 1.x builds against the older ROOT version.

Another area of significant development is the new relation abstraction layer in POOL, which will remove the tight coupling of some POOL components to MySQL/MySQL++. The imminent V1.7 release will include this new layer. This will allow to choose from a larger variety of relational database back ends (initially Oracle, MySQL via ODBC and SQLight) using a single consistent interface.
A relational file catalog (expected to later replace the MySQL catalog) has been produced based on this abstraction and successfully passed the catalog integration tests.
The work on a relational storage manager for C++ objects is proceeding well and will soon appear as a first prototype.
All new relational components are not considered production quality yet and are released to validate the approach from the experiment side and stimulate feedback.

Conditions Database
The conditions database sub-project made significant progress and produced a combined release with both conditions DB implementations supporting the LCG platform compiler. The interface review with respect to extensions in the ATLAS implementation has started. A table of extension and their interdependencies is being prepared and will be discussed with all experiments.
Even though the conditions db sub-project received significant interest from all experiments and active development contributions there still is no definite manpower commitment from the experiment side.

WBS 1.1.3 – Core libraries and services (SEAL)

Project Leader: Pere Mato, CERN EP/SFT

The main items of work in this quarter have been the release, as part of SEAL version 1.4.0, of some planned new functionality. Among the new deliverables are the first version of the SEAL workbook, the first prototype version (nearly complete) of the new C++ Reflection classes (Reflex), a new Math package for minimization and fitting and the completion of the Minuit package. In addition to the new packages, maintenance actions have included major simplifications in the Plugin Manager and a re-implementation of PyLCGDict. On the SEAL/ROOT convergence plans, the PyROOT package has been migrated successfully from SEAL to the ROOT repository and distribution. Of the three Level 2 milestones due in the quarter, one was completed (SEAL workbook); one will complete a week or two late, putting the completion in the next quarter (Reflex prototype release); and one was left undone due to relatively low priority (external software guideline document). The main highlights for each work package in this second quarter are:

- **Foundation libraries**. Major simplification in the Plugin Manager by basically reducing the number of classes the user needs to implement by a factor of two. SealBase has had a number of interface changes while the classes are still in relatively little use.
- **Math Libraries.** Minuit contains various changes with respect to the previous release and basically completes the package. FML is a new package containing a library for solving complex and customized fitting and minimization problems. FMLMinuit is new a package implementing the FML minimizer interface using the seal Minuit library. Work is progressing well in several areas: GSL testing and validation, Math libraries web pages, setting up the Mathlibs forum.
- **Framework.** A new class ComponentWrapper has been introduced which allows installing arbitrary user defined classes as a SEAL component. This was a pre-requisite for POOL adopting the SEAL component model. In addition, there has been some reorganization of the packages and some bug fixes.
- **LCG Dictionary**. A dictionary service to load dictionary files on demand has been released. Dictionary libraries are now constructed as SEAL capabilities libraries and they are handled by the Plugin Manager. Most of the current work package effort has been devoted into finishing the prototype implementation of the new reflection package (Reflex) and adaptations of the lcgdict command to generate new dictionaries.
- **Scripting Services**. The first version of PyLCGDict2 was completed. At the same time we established contacts with PyLCGDict early adopters and started to help them.
  We completed modifications to PyROOT to eliminate the dependence on Boost and moved it into the ROOT repository and distribution system starting from ROOT version 4.00/04.

- **Documentation and Education**. The initial version of the SEAL Workbook has been released. It contains topical Web pages for all SEAL work packages with very practical information of how to use the compete spectrum of the SEAL products. This completes a Level 2 milestone. The SEAL team will make an effort to keep the documentation up to date.
- **General Infrastructure.** The latest SEAL version has been released with gcc-2.3.2 and tested the new version of the CMT interfaces (LCGCMT).
- **Experiments Integration.** No major problems have been encountered during this quarter.

WBS 1.1.4 – Physicist Interface (PI)

Project Leader: Vincenzo Innocente, CERN EP/SFT

Again in this quarter the work concentrated on maintenance and user support. A bug-fix release was cut and announced to the users on May 5th (1.2.1). Apart from a small number of bug fixes, some code refactoring was done to re-use the I/O packages from SEAL, allowing the user to select the compression algorithm used from gzip to bzip2 as requested by users from CMS. This release also features the availability of (simple) ROOT tuples as an implementation of the AIDA Tuple interface.

For the fitting classes, a set of new tests has been developed and implemented by W.L.Ueng (visitor from Academica Sinica, Taiwan) re-using the design and architecture from the tests for the histogram classes. The tests are based on the CPP unit framework.

In the contrib subsystem (containing externally developed packages in an "as-is" state) the StatisticsTesting package has been upgraded to the latest release (1.1.0) of the toolkit.

Another release is planned for end of June, based on the latest release of SEAL (1.4.0).

WBS 1.1.5 – Simulation

Project Leader: T. Wenaus, BNL/CERN EP-SFT

The Generic Simulation Framework subproject (A. Dell'Acqua/ W. Pokorski) proceeded with the program agreed during the last quarter, to provide a generic Geant4/FLUGG/FLUKA infrastructure for test beam physics validation studies, and (as the last milestone in this subproject as presently constituted) to pursue GDML as an exchange format by which Geant4 based geometries can be transferred to ROOT for VMC evaluation. A first successful step had been the application of FLUKA with the FLUGG interface to the analysis of ATLAS pixel test beam data. This work has to be continued towards a generic infrastructure and manpower has been identified. The PEMF automation, which is important for a user-friendly generic infrastructure, will be added to FLUKA late this year, with the FLUKA release planned then. ALICE has already developed an in-house solution, however, and this will be examined as a possible quick solution. On the GDML front, work proceeded on supporting XML writing (as well as reading), as described in the Geant4 section below and as required for the exchange format application.

The Geant4 subproject (J. Apostolakis) released version 6.2 during the quarter. In addition to overall improvements, several specific requirements from LHC experiments for new features and revisions were included. Contributions to this minor release undertaken by SFT members (and collaborators) included addressing issues for experiment production such as improvements to navigator to cope with stuck tracks, and fixes and improvements of EM (std) for HEP showers. In addition there were fixes for high energy muon EM interactions building on the refinements in Geant4 6.1, the refinement of the PAI model, and the new ability to reflect divided volumes. Physics

improvements included the revision of binary cascade and imaginary R-matrix to aid use in LHC calorimeter simulation. Included also was the functionality to label (optionally) as a new particle an incoming particle that survives a hadronic interaction. This addresses the LHCb requirement to re-label particles in a manner that is compatible with its track reconstruction. An extension to G4Allocator to identify memory used and to release it at the user's request (and responsibility) is in. These changes were requested by CMS, to enable better accounting and use of memory. Physics lists for electromagnetic use cases are released in Geant4 6.2 for the first time.

SFT group members contributed to a number of other physics areas. One was an update of pion reaction cross-sections at low energies. An alternative data-set that reflects Barachenkov's cross-sections, provides a small improvements in shower shape. An investigation was made of the CPU performance in using Binary Cascade for calorimeter simulation. This is a preliminary step to the deployment of binary cascade for calorimeter simulation at LHC. New techniques were created to configure the complex scattering term. In particular new compile-time constructs for particle codes and particle multiplets enable us to configure it using template meta-programming. Investigated the ability of the Binary Cascade model to describe neutron production. The study concluded that experimental data can be described within error, if the need arises. A first implementation of a new variant PAI model creates also secondary photons.

Other improvements:

- The geometry modeller was extended to support the reflection of volume hierarchies containing divided volumes, following ALICE request and contribution.
- A first release was made of a module for streaming out geometry in a GDML (XML geometry description) file. The GDML schema was extended to support assembly of volumes. New requirements have been collected for enabling future Geant4 persistency via POOL or ROOT.
- Revisions in source modules and to the configuration setup were undertaken to enable support for dynamic loading on Windows. This fulfils an LHCb request.
- Revised the organization of binary releases for ‘recommended’ hadronic physics-lists in CERN AFS. This fulfils the need to deploy updates of the physics lists for a specific release of G4. It also enables someone to utilise several Geant4 releases with a particular physics list release.
- Improved Bonsai’s tag-collection capabilities, adding a "bug fix" field. Deployed Bonsai in production for Geant4 testing. Also added the capability to gather the descriptions of a group of tags, which eases measurably a part of the work for creating a release. Extended it to reflect groups of tags under testing.
- Implemented porting of extended and advanced examples to use the PI histogramming implementation (PI-lite) and updated to the AIDA 3.2.1 interface.

Release 6.1 has been built for the new Linux (cel3/sel) and this build is deployed in the SPI external software library area. Release 6.2 is currently provided via the standard Geant4 distribution mechanisms.

Positive performance results (~10% level for some cases) were obtained using O2 level optimization with gcc. This is configured as the default optimization level for the Linux/gcc setup for release 6.2. We saw first progress reported towards obtaining consistent results between optimized and debug compiles: tests confirmed that precision changes are due to the use of 80-bit registers for temporaries in place of 64-bit doubles in memory.

The FLUKA Integration subproject (A. Ferrari) A. Ferrari reported on recent FLUKA improvements and applications. The physics developments concern mainly heavy ion interactions. The interface to DPMJET-III is now fully operational. The interface to DPMJET-2.5, which is the one currently distributed, is still maintained for some time, until all possible issues with the new version will be sorted out. Electromagnetic dissociation has been implemented as a new hadronic process. This work has been performed in collaboration with G. Smirnov and under partial coverage from NASA funding. Electromagnetic dissociation is a key process concerning LHC operation with ion beams, both for experiments and machine elements.

Work is still going on to provide a new low energy neutron library and to eliminate the need for the PEMF pre-processor. PEMF is presently needed to generate material-dependent electromagnetic data files which are then read by FLUKA. Its functionalities will be embedded in the overall FLUKA initialization.

Many interesting FLUKA applications have profited from recent improvements in physics and geometry functionalities. In the Rossendorf ion therapy center work is in progress in the FLUKA validation, with the final goal of predicting the generation of beta+ emitting nuclei during the irradiation. Positron emission will allow verification in-situ of the treatment plan by means of PET. Voxel modelization has been applied to human phantoms, both for therapy and space dosimetry. For this last topic, a method to derive a voxel modelization of the human body directly from CT scan has been developed in the framework of the FLUKA-NASA collaboration.

The Physics Validation subproject (F. Gianotti) had the following principal achievements this quarter.

- Studies of hadronic interactions in the ATLAS pixel detector were completed. They provided comparisons of test beam data with G4 and with FLUKA (through FLUGG). The work is documented in LCG note CERN-LCGAPP-2004-009.
- The second simple benchmark (pion absorption) comparing data to G4 and FLUKA has been completed. A note is in preparation.
- The note documenting the first round of hadronic physics validation is completed and will shortly be published as an LCG note. This completes a level 2 milestone.
- The general infrastructure to evaluate radiation background in LHCb with G4 has been put in place and first results of fluences (neutron, charged pions, electrons) at four scoring planes have been presented. The next step is to increase the statistics and to compare with previous results obtained with GCALOR and FLUKA.
- Comparisons of EM calorimeter test beam data with G4 are ongoing in three experiments (ATLAS, CMS, LHCb). They focus on E-resolution, response linearity, and shower shapes. First results should be available in September.

Reviewing and assessing the priorities of potential further simple benchmark studies (a level 2 milestone due this quarter) had to await the completion of the pion absorption simple benchmark, which took longer than expected. With that completed, future work on simple benchmarks can be discussed and prioritized. This is expected to happen before September.

The Generator Services subproject (P. Bartalini) received the report of the internal review conducted last quarter. The document points out the critical issues of manpower (librarian, project leader) and communication with generator authors. Steps are underway to address these problems. The LCG Librarian will assist with GENSER library support but it is also important that a Monte Carlo expert work in close collaboration with the LCG librarian and generator authors; a solution is being finalized. The Russian group plans to give greater focus to MCDB, reducing their GENSER librarian role. During the quarter GENSER was put into production by an experiment (ATLAS) for the first time.

UK phenomenologists have received 1 FTE support for a project ("phenoGRID") one aspect of which is to help to get HERWIG++ incorporated into the LCG. They are eager to coordinate closely with the LCG Generator Subproject.

WP1 - GENSER 0.1.0 was released during the quarter. It contains AlpGen 1.3.2 (validated by F.Ambroglini), CompHEP 4.2p1 and 4.4.0, EvtGen alpha-00-11-07 and pdflib 804. Some of these packages were installed as external tools in LCG. Also several new versions of the already introduced sub-packages were made available in GENSER 0.1.0. There is general agreement that the current size of the HepEvt common block is not amenable for the LHC physics. A survey will be performed to check the requirements of the four LHC experiments. For the time being a special version (GENSER 0.1.1) with doubled size was made available for ATLAS DC2 production. The May generator services meeting was dedicated to the LHC version of EVTGEN. All experiments have an interest in it. LHCb agreed to make available to

LCG its private version. ATLAS agreed to help in this transition. GENSER is now on a quarterly release cycle. The LHC version of EVTGEN as well as the new Pythia 6.3, and Herwig 6.505 + Jimmy 4.1 will be in the next release.

WP2 - In May a meeting with wide attendance from Monte Carlo experts was held in Lund to decide the future development guidelines of ThePEG. The LCG Generator Subproject was represented in the meeting. Some activity concerning the doxygen documentation of ThePEG started during the quarter.

WP3 - An updated version of the MCDB plan with short- and long-term plans has been prepared. An LCG MySQL Server running MCDB was deployed during the quarter (http://mcdb.cern.ch). The corresponding milestone (1.198) is considered to be achieved. The generator level production/validation framework is still at the design stage but still looks promising. Some activity started in the Spanish groups that expressed interest in this project. The delay is due to the effort needed to coordinate with CMS. The common LCG/CMS project was presented during the recent CMS week. A 3 months delay is expected with respect to previous plans (affecting two level 3 milestones). Alternatives will be evaluated if the LCG/CMS collaboration encounters further problems.

WP4 - The work on Hijing validation was continued at JINR. A server adopting the JetWeb concept (and technology) was developed there. A presentation was given at the June LCG Generator meeting. Every effort will be made to avoid duplication of work and to merge such activity with JetWeb.

WBS 1.1.6 - ROOT participation

LCG AA and the ROOT team made further progress during the quarter on deepening the collaboration between these two projects, particularly in the dictionary and mathlib activities of SEAL and in POOL ROOT I/O preparations for ROOT4. SEAL and ROOT also completed the transfer of PyROOT from SEAL to ROOT; it will be part of the ROOT4 production release in July. ROOT and POOL worked closely on understanding and minimizing the backwards compatibility issues (or potential issues) which could arise with the planned migration of POOL to using ROOT4 and its improved support for automated schema evolution of foreign classes and STL persistency. By the end of the quarter a plan for prompt deployment of ROOT4-based POOL was in place (with concurrent support for ROOT3-based POOL to continue until the experiments are ready to drop it). This plan is accelerated with respect to last quarter's tentative plan (which foresaw ROOT4 integration extending until October), because of the success in minimizing backwards compatibility issues.

The applications area continued to support two developers on the ROOT team. Ilka Antcheva continued development of the ROOT GUI system and editing tools, and neared completion of a major documentation update for ROOT4 to be released in July. Valeriy Onouchine worked to improve the robustness of the win32gdk Windows version of ROOT, and continued development work on the ROOT GUI builder and associated GUI and editor classes. Ilka and Valeriy presented some of their recent work at an applications area meeting at the end of the quarter. Gerardo Ganis continues to work with ALICE and the ROOT team on the distributed capabilities of ROOT.

WBS 1.1.7,1.1.8 – Core/Grid Interface & Experiment Integration

Through this activity the Applications Area provides direct assistance to the experiments at the interface between core software and the grid, and supports the adaptation of physics applications to the grid environment. Activities are managed by the experiments and not the applications area. Mattias Ellert works with ATLAS on the integration and deployment of ATLAS production software on the grid, particularly the NorduGrid service. This quarter he worked on preparations for the import of the BaBar-developed xrootd and its client in ROOT, and on the extension of remote file services provided by rootd, requested by POOL and CMS.

Milestone performance during the quarter

WBS 1.1.1 - SPI:

- RH7.3/gcc 3.2.3 supported (May)
  - Completed
- SPI/EGEE collaborative workplan complete (end May)
  - Delayed to end July due to EGEE delays.
- Certification of external software for the new Linux platform (Jul)
  - Completed a month early

WBS 1.1.2 - Persistency:

- First release of POOL RDBMS abstraction layer completed (end May)
  - Completed
- RDBMS independency achieved for production versions of POOL relational components (end Jun)
  - Completed for file catalog; work still underway for collections. Expected next quarter.
- Common interface for conditions DB defined (end Jun)
  - Existing common interface is basis for present work. ATLAS (Lisbon) has been invited to present its extensions for possible inclusion and has promised to do so. Expected next quarter, depending on ATLAS.

WBS 1.1.3 - SEAL:

- External software guideline document (Jun)
  - Not started yet. A victim of higher priorities.
- SEAL workbook release (Jun)
  - Completed
- New dictionary API and reference implementation (end Jun)
  - Will be 1-2 weeks late (completing early next quarter)

WBS 1.1.4 - PI:

- No milestones this quarter.

WBS 1.1.5 - Simulation:

- Initial generic simulation prototype supporting Geant4, FLUKA via FLUGG (Feb)
  - Delayed due to communication and technical problems with FLUKA.
- First cycle of hadronic physics validation complete (Feb)
  - Completed
- Review/prioritization of simple benchmarks for physics validation (end May)
  - Delayed due to delays in completing FLUKA part of present simple benchmark study. Ready to begin at end of quarter. Expected next quarter.
- Geant4 6.2 release – performance improvements (Jun)
  - Completed
- Beta version of MCDB in production (Jul)
  - Completed slightly early

Upcoming milestones and comments on status

WBS 1.1.1 - SPI:

- LCG AA build system selection (Jul)
  - New librarian is now actively engaged in build/configuration system issues but this is likely to be late. Should be done during the quarter.

WBS 1.1.2 - Persistency:

- Conditions DB production release (end Jul)
  - Release of a version suitable for experiment evaluation is expected around that time.
- First release of the POOL relational storage manager (end Aug)
  - Should be ready.

WBS 1.1.3 - SEAL:

- Mathlib project web (Jul)
  - Coming together well; should be ready.

WBS 1.1.4 - PI:

- No upcoming level 2 milestones next quarter.

WBS 1.1.5 - Simulation:

- Agreement on parton level event generator file format (Sep)
  - No reason to believe it won't be ready
- Comparison of LHC calorimeters for EM shower development (Sep)
  - Advancing well and expected to be ready

2. CERN Fabric Area – Bernd Panzer

Major Achievements and Decisions

There has been a successful collaboration with Fermilab on a common platform for a Linux distribution, which is now called Scientific Linux. The current CERN certification is now based on this distribution, which itself is based on the source code of RH for Enterprise version 3. At the end of May the first Lxplus and Lxbatch machines have been setup with the new Linux version.

The maintaining of the Linux version 7.3 starts now to get problematic. RH stopped the support at the end of last year. We were relying on the user community (in the Fedora framework) to provide necessary updates and patches (especially security). This support has decreased over the last months. We will from the 1st of July pay a company to provide us with the RH 7.3 patches.

The HEPIX meeting in May was very useful in the area of 'coordination' between the fabric installations in the different sites (e.g. Scientific Linux, Service Data Challenges, etc.) The following link gives an overview and a summary :
http://agenda.cern.ch/askArchive.php?base=agenda&categ=a042569&id=a042569s0t1/moreinfo

The clearing of the left half of the machine room in the computer centre is progressing well. More than 500 nodes have already been moved (another 300 to go). This is a tedious and very work intensive exercise, especially as IT tries to avoid service interruptions. This was not always possible. A major move of 140 disk and stage servers from the LHS to the RHS of the machine room was successfully completed. Since many CDR machines for the running fixed target experiments were implicated, as well as stagers from practically all experiments, the move was concentrated into one day and scheduled during the last long break in SPS running before the end of the Autumn. Unfortunately there have been also about 4 small hiccups with AFS during the move.

The latest severe increase in security incidents have lead to proposals to tighten the management of Windows desktops. In this light Linux strategies for ensuring better security are under development. See minutes from the Desktop Forum :
http://agenda.cern.ch/fullAgenda.php?ida=a042526

The repair of more than 60 Elonex disk servers was finished by the end of June and the rate of failures has decreased considerably during the last 2 months.

Some experiments are now using a much larger number of smaller files than was expected (up to 25 files per event compared with one file per event assumed in earlier plans). This is and will continue to be a major problem for mass storage systems, in particular any magnetic tape layer. Much more work needs to be put into this issue to find a reasonable compromise between the experiment requirements, the restrictions of the technology and the complexity of the software system.

Milestones in Q2 2004

Mass Storage

-- Milestone 1.2.1.1.2 due 02-02-2004

Integrated prototype of the whole system

The first integrated prototype of the new CASTOR stager was successfully tested on the 26th of June. Now an intense period of testing and debugging has started, with the aim to use this setup in the upcoming ALICE DC6.

→ **milestone met 26-June-2004**

-- Milestone 1.2.1.1.3 due 01-04-2004

Production system ready for deployment

As the prototype was delayed by several months the production system is correspondingly late. The plan is still to have major production installations Ready before the end of the year.

→ **new date estimate 02.11.2004**

-- Milestone 1.2.1.3.1 due 01-03-2004

Report on LTO evaluation delivered

Report on LTO1 evaluation tests is in preparation. However, since LTO2 is now available, tests of these devices are necessary. We are in the process of establishing a collaboration with Valencia for work in this area.

→ **new estimate 17.01.2005**

-- Milestone 1.2.1.3.3 due 01-06-2004

Proposal on storage architecture and implementation of disk storage system for Phase 2 made

→ **new estimate 15.12.2005**

-- Milestone 1.2.4.1 due 01-06-2004

Baseline plan for Phase 2 agreed (architecture + technology)

→ **new estimate 15.12.2005**

Fabric Management

-- Milestone 1.2.2.4 due 01-07-2004
SMS – full production

The State Management System is now (since beginning of June) in full production. The system is a framework where the essential modules have been deployed. More modules are under development and deployed in a continuous rate.

→ **milestone met 16.06.2004**

-- Milestone 1.2.2.5 due 01-03-2004
Monitoring – Systems monitor displays operational

The LEMON monitoring system has been equipped in May with a visualization system (**http://ccs003d.cern.ch/lemon-status/**).

→ **milestone met 19.05.2004**

Computer Center Upgrade

-- Milestone 1.2.3.5 due 01-04-2004
Sub-station electrical installation starts

→ **milestone met 01-04-2004**

Networking Infrastructure

-- Milestone 1.2.5.9.5. due 01-07-2004
First phase of backbone replacement by 10 Gbit equipment

The first 10 Gbit router (N7 from Enterasys) was already integrated into the network backbone in the Computer Centre by the beginning of June.

→ **milestone met 10-06-2004**

Preview Milestones in Q3 2004

Data Recording Challenges

-- Milestone 1.2.1.2.3 due 01-10-2004
ALICE DC6 (450 MB/s)

There is a large load on the tape system during the SPS running period and the LHC experiment data challenges which will last throughout October. As the DC will need more then 50% of the available tape drives the milestone is shifted to the 15th of December. The activity will as usual start already in September and just the tape activity will be exercised at the end.

→ **milestone shifted 15-11-2004**

Computer Centre Upgrade

-- Milestone 1.2.3.8 due 02-08-2004

Left half of the machine room emptied

The milestone will shift by a few weeks, but there will be enough free space to start the refurbishment of the electrical installation in parallel to the remaining move of equipment.

Phase 2 Acquisition Process

-- Milestone 1.2.6.3 due 01-07-2004

Market survey issued for cpus and disk

The market survey papers were prepared, but will not be sent out in July. In an IT meeting with the DG he encouraged us to make a proposal for a simpler and more flexible process to deal with the acquisition of commodity equipment. A proposal for this has been prepared together with FI Department and a paper describing the new process will be sent to the Finance Committee later this year. This milestone will be re-defined to take account of the new process.

3. Grid Deployment Area – Ian Bird

Milestone Status

This first quarter of the project has seen the growth of the LCG-2 infrastructure (which also forms the initial EGEE production service), to some 60 sites, providing more than 6000 CPUs, although there are continuing issues with some site providing insufficient disk space on their Storage Elements to match the amount of CPU they provide. The list of the sites can be seen at http://goc.grid-support.ac.uk/gppmonWorld/gppmon_maps/lcg2.html, and details of their status at http://goc.grid.sinica.edu.tw/gstat/.

The most significant activity during the quarter has been the support for the ongoing data challenges of all four experiments (although that of CMS ended in April). A number of issues have been brought to light through this activity, ranging from specific functional problems for which many fixes have either been deployed already or are being worked on, to broader architectural issues. The most significant of these is the model of job submission and the issue of matching the sophisticated abilities of modern batch systems through the simple gatekeeper (Compute Element) to the Resource Brokers. Some work-arounds to improve this situation are being deployed and tested, but it is unlikely that real solutions to this problem will be available without significant changes to the underlying model.

Two middleware releases have been certified during this quarter. The first in May consolidated many bug fixes, and provided updates to the replica manager tools in response to problems found during the CMS DC04, and needs from ATLAS and LHCb found during their preparations for their data challenges. In addition this release upgraded VDT to a version that is also used by Grid 2003 in the US, bringing interoperability closer. The second release in June, as well as more bug fixes, particularly to the resource broker, also introduces R-GMA as an additional service that will allow the experiments to use it for general monitoring needs, as well as being the platform for the accounting system being developed by RAL.

The support for operational issues has been augmented with the availability of a second Grid Operations Centre (and staff) in Taipei. Responsibility for monitoring tasks is now shared between the Taipei and RAL GOCs. The response to operational issues will gradually become the shared responsibility of these centres together with CERN. The user support portal built by FZK (http://www.ggus.org) is now available for user problem reporting, and is gradually starting to be used. A task force, with members from the experiments, support and

operational groups, has been addressing how to implement and agreed support model and workflow, with the support portal and its problem tracking system as the underlying tool.

Steps towards interoperation of several grid infrastructures have been made. Discussions are under way between Grid 2003 and LCG-2 to merge the information systems to allow job submission between both grids. Discussions with NorduGrid aimed at understanding the real technical differences are also under way. In Canada, the group at Triumf have successfully demonstrated the use of an LCG-2 node as a gateway to Grid-Canada and WestGrid, allowing LCG jobs to run on those grids.

A series of service challenges have been proposed and agreed in the PEB. These are: demonstrate a reliable data transfer service, that can sustain reliable high performance file transfers between Tier 0 and Tier 1 sites at a data rate that is a significant fraction of that anticipated at LHC startup, by the end of the year; filling the entire system with jobs for extended periods to understand system limitations and bottlenecks; security incident response challenges ("fire drills") to ensure the procedures are in place and appropriate; and finally a set of interoperability challenges to be further defined based upon the results of the discussions mentioned above. These challenges will be scheduled around the data challenges and experiment production use. A group of managers from the major centres – the GDA steering group – has been nominated to schedule and managed these service challenges.

The staffing of the CERN deployment group has been increased with the staff funded through EGEE, although the scope of work of the group has also increased correspondingly. Many of the LCG European partner sites are in a similar situation. The increased staff will help provide the needed infrastructure and operational support effort.

Detailed milestone status for Q204

Service:

1.4.5.1.3: LCG-2 operational at 30 sites (May 1 2004)

This milestone was met, with just over 30 sites within LCG-2 by the beginning of May, and 60 sites involved by the end of the quarter. This brings the amount of CPU available to over 6000. As well as the European sites, this includes 2 sites in the US, 5 in Canada, 6 in Asia and the first HP site in Puerto Rico. Up to date information on LCG sites can be seen at http://goc.grid-support.ac.uk/gppmonWorld/gppmon_maps/lcg2.html, and details of their status at http://goc.grid.sinica.edu.tw/gstat/.

Operations:

1.4.5.2.2: Second GOC in operations (June 1 2004)

This milestone was met ahead of schedule with the site at Taipei (http://goc.grid.sinica.edu.tw/goc) taking responsibility for maintaining an overview of the operational state of LCG-2 by the end of March. Responsibility is shared between the GOC in RAL and that in Taipei. Future work should extend the functionality and capabilities of both GOCs and include responsibility for resolving operational issues.

Support:

1.4.5.3.1: FZK Support portal as front line user support (June 1, 2004)

This milestone is met. The portal at FZK (http://www.ggus.org) was available by this date and ready to be used as user support tool. Some use has been made of this system, and the usage is gradually increasing, but since the site was available during the very active phase of the data challenges it was hard for the experiments to change the way they were reporting and handling problems, since during the challenges most of the problem reporting has been through specific dedicated contacts. The process to ensure all problems (even post-facto) get reported into the system is under way.

Experiment Verification of LCG-2:

1.4.5.4.1: Report on ALICE DC initial experience (June 1, 2004)
1.4.5.4.2: Report on CMS DC04 initial experience (June 1, 2004)

These milestones are partially met. ALICE and CMS have both reported on their experiences in GDB and PEB meetings. There are no written reports. In addition the GDA will produce a documented summary of the experience with the data challenges from the operational and site point of view, but this will not be written until there is more experience with the ATLAS and LHCb challenges which are continuing. The GDA report will cover experiences with all of the challenges.

1.4.5.4.3, 1.4.5.4.4 Report on ATLAS and LHCb experiences. Due to the delays in starting the Data Challenges, it was not possible to meet these milestones. These reports will be produced during the next quarter.

Middleware:

1.4.5.5.1: WN middleware ported to RHEL, report on ease of porting to other OS's (May 1, 2004)

This milestone is partially met. The Worker Node middleware is ported to and running on IA64 and IA32 under Scientific Linux, which is derived from RHEL. In fact most of the entire middleware suite of LCG-2 has been ported and is running in these environments. Testing is still continuing and at the moment only manual installation is possible, since the automated installation using LCFGng requires more work to port the tool itself. The report on the ease of porting is still outstanding, and will be delivered when the remaining installation issues have been resolved.

Upcoming milestones for Q304

Experiment Verification of LCG-2:

1.4.1.9: V_ALICE AliEn production interface to LCG-1 performance evaluation (15 July 2004)
1.4.1.10 V_ALICE LCG-1 PDC3 production environment verification (15 July 2004)
1.4.1.16 V_LHCb LHCb analysis environment (GANGA) deployed in LCG-2 (1 July 2004)

Middleware:

1.4.5.5.2: Configuration tool available (1 July 2004)

An automated tool should be provided that enables a straightforward configuration of a site.

1.4.5.5.3: Replica Manager upgrade (1 August 2004)

Upgraded data management tools to be available. This should include a distributed version of the RLS, and a replica manager service that acts on behalf of the worker nodes. This upgrade should remove the requirement of worker node outbound IP connectivity for data management.

- This milestone is no longer appropriate and will not be met. The replica manager tools have been upgraded to improve performance and to add missing functionality as requested during the data challenges. However, a distributed version of the RLS will not solve the underlying performance issues, and the strategy for providing a distributed version of the RLS will be analysed in the context of the new distributed database study.

Resources

The situation with staffing is as shown in the resource table at the end of this document. At the end of the quarter a Portuguese trainee left the deployment team as his contract ended. Another 2 INFN funded fellows will leave the Certification team in the next quarter (one in August and one in September). Between April and June the EGEE funded staff joined the group. EGEE SA1 (operations activity) funded 8 FTE. Two of these (not counted in the table) are associated with EGEE project overhead, 1 is located in the DB group to provide dedicated database operational support. Five of the EGEE funded FTE provide additional support for grid operations support in the areas of supporting the deployment and service operations, as well as a new activity of operational monitoring which CERN is obliged to provide as part of its commitment to EGEE.

4. Middleware – Frédéric Hemmer

The main middleware activities during the reporting period were:

- Set up of the development, integration, and testing infrastructure
- Continue the design and development of the next generation Grid middleware (named "*gLite*")
- Provide a prototype installation of the gLite middleware to the ARDA project
- Training and dissemination activities
- Coordination with other projects and contributions to standardization efforts

Set up of the development, integration, and testing infrastructure

The three development clusters (UK, IT/CZ, and CERN) have put in place a distributed development and testing infrastructure fulfilling Milestone 1.5.2.7. These testbeds are expected to grow in size and geographic distribution in the following months to allow software development in a realistic Grid environment.

The integration infrastructure consisting of a common CVS repository and continuous integration servers for the middleware has been put in place at CERN together with the relevant guidelines and procedures This together with the provisioning of testing tools fulfils Milestone 1.5.2.6.

A distributed testing testbed has been put in place at CERN, RAL, and NIKHEF. The support for the RAL and NIKHEF sites is provided by EGEE SA1. The initial testbed consists of 18 machines and is expected to grow in size and maybe also geographic distribution in the following months.

Continue the design and development of the next generation Grid middleware

The architecture of the next generation Grid middleware (named "*gLite*") has been defined (https://edms.cern.ch/document/476451/). The architecture comprises 5 logical service groups:

- Access Services
- Security Services
- Information & Monitoring Services
- Data Services
- Job Management Services

Initial services for each of these groups have been selected from existing components originating from AliEn, EDG, and VDT. These services are being re-engineered to follow a

Service Oriented Architecture based on Web Services and services not yet existing are being written.

The gLite services will follow established standards as much as possible; however not necessarily conform to proposed, but yet to be finalized standards, like WSRF. We consider WSRF an important development in the field of Grid computing; but try to avoid major modifications imposed by changing standards. We will follow the development of WSRF in order not diverge from it and will conform to it once it becomes an established standard. Conformance with WS-I is an important goal for gLite to ensure interoperability.

A detailed release plan has been worked out (https://edms.cern.ch/document/468699).

Provide a prototype installation of the gLite middleware to the ARDA project

Prototype versions of the gLite Grid services are installed on a distributed testbed spanning CERN and Wisconsin-Madison. On this testbed, members of the ARDA project can assess the service semantics and interfaces. Since the opening of the testbed to the ARDA project on May 18th, many bug fixes and upgrades have been applied. A detailed status overview of the testbed can be found at http://egee-jra1.web.cern.ch/egee-jra1/Prototype/testbed.htm.

Training and dissemination

The activity manager and his deputy have been active in the project training activities and contributed to the dissemination of the project results through presentations at international conferences and meetings.

Coordination with other projects and contributions to standardization efforts

Links to related projects like OGSA-DAI, Unicore, and GEMSS have been established. Various activity members contribute to ongoing standardization efforts, in particular through GGF research and working groups (OGSA-WG, OREP-WG, GSM-WG, INFOD-WG).

The major LCG events in the reporting period were the 2nd ARDA workshop, 21st-23rd June, CERN (http://lcg.web.cern.ch/LCG/peb/arda/LCG_ARDA_Workshops.htm) and the LCG PEB meetings (http://agenda.cern.ch/displayLevel.php?fid=3l132).

The work is coordinated with the EGEE project through participation to the 1st EGEE Conference, 19th-22nd April 2004, University College, Cork, Ireland (http://public.eu-egee.org/conferences/kickoff/), the EGEE All Activity meetings, 18th June 2004, CERN, Geneva, Switzerland (http://agenda.cern.ch/age?a042415), and the EGEE PEB meetings (http://agenda.cern.ch/displayLevel.php?fid=168).

Internal coordination takes place through weekly meetings of the Engineering Management Team (EMT) (http://agenda.cern.ch/displayLevel.php?fid=306), monthly Design Team meetings (P. Buncic, S. Fisher, D. Groep, F. Hemmer, P. Kunszt, E. Laure, M. Livny, and F. Prelz), and an All-Hands meeting, 28th-30th June 2004, Cosener's House, Abingdon, UK (http://agenda.cern.ch/fullAgenda.php?ida=a041852). Focused meetings on testing, integration, and development issues are being held as well.

Other activities included presentations at the Condor week, 14th-16th April 2004, Madison, Wisconsin, USA, the GridPP Collaboration Meeting, 2nd-4th June, CERN, and attendance at GGF11, 6th-10th June, Honolulu, Hawaii USA.

Seven papers have been accepted at Computing in High Energy Physics (CHEP) conference, September 2004, Interlaken, Switzerland:

- "A Pattern-based Continuous Integration Framework for Distributed EGEE Grid Middleware Development", accepter for poster presentation
- "Middleware for the next generation Grid infrastructure", accepted for oral presentation
- "Distributed Testing Infrastructure and Processes for the EGEE Grid Middleware", accepted for poster presentation
- "Distributed Tracking, Storage, and Re-use of Job State Information on the Grid", accepted for poster presentation
- "Data Management in EGEE", accepted for oral presentation
- "Practical approaches to Grid workload and resource management in the EGEE project", accepted for oral presentation
- "Information and Monitoring Services within a Grid Environment" accepted for oral presentation

Milestones

The Q2 milestones foreseen were:

WBS	ID_Unique	Name	Date_MS	MS_Done
1.5		**Middleware**		
1.5.2.6	5.1912	First version of prototype available for experiments	16-06-04	18-05-04
1.5.2.7	5.1889	Development and integration tools deployed	30-06-04	30-06-04
1.5.2.8	5.1888	Software cluster development & testing infrastructure in place	30-06-04	30-06-04
1.5.2.9	5.1887	Architecture & Planning Document for release candidate 1	30-06-04	30-06-04

The milestones due in the second half od this year are listed below. It is expected that these will be met on time.

WBS	ID_Unique	Name	Date_MS	MS_Done
1.5		**Middleware**		
1.5.2.10	5.1886	Integration & Testing infrastructure in place; test plan	31-08-04	
1.5.2.11	5.1885	Grid Services design document for release candidate 1	31-08-04	
1.5.2.12	5.1913	Second version of prototype available for experiments	31-08-04	
1.5.2.13	5.1884	Release Candidate 1	31-12-04	

5. Distributed Analysis – Massimo Lamanna

Introduction

This is the first full report of the ARDA project.
In the initial phase, most of the activity was devoted setting up the group (organized as a CERN section for the CERN based part), to the logistics and discussing the project with the experiments and with the interested part of the HEP community.

Currently the web site (http://lcg.web.cern.ch/LCG/peb/arda) contain information on the project itself, practically all the presentations (CERN bodies such as the PEB and GAG, EGEE events such as the Cork conference and external events such seminars at DESY and at RAL). Mailing lists are also provided. In particular the arda@cern.ch mailing list hosts technical discussions about the gLite software across the LHC experiments community (this list re-factorises the several pre-existent lists).

The group is reasonably installed in bd.510 since end of April. The CERN section comprises 4 persons hired under the posts EGEE-NA4/HEP (2 year staff contracts) and 6 CERN staffs and fellows. The section includes 4 collaborators visiting CERN on a temporary basis (Russia and Taiwan).

We have received several invitations to present the project (DESY, RAL, gridPP meeting at CERN, GGF). Only the GGF offer was declined (due to logistics problems – we look forward to present ARDA at the next GGF).

The definition of the detailed work plan is under control (Milestones 1.6.2, 1.6.3, 1.6.5), although there is no final proposal from CMS for their prototype (Milestone 1.6.4). However, the activity is organized in a way that the persons attached to the CMS part of the project are contributing useful preparatory work (see CMS section of this report).

The activity with external collaborators (a total of 4 people from Russian institutes and Taipei) is progressing (note that these colleagues are only part time at CERN).

All the public documentation (documents, presentations) is available from the ARDA web site (http://lcg.web.cern.ch/LCG/peb/arda/LCG_ARDA_Docs_Public.htm).

Activity status

The focus of the activity is to deliver 4 end-to-end prototypes, one per LHC experiment, based on the EGEE middleware (gLite). This section reviews the status of the 4 parallel projects and then a summary is given from the perspective of the middleware (tested functionality, main problems).

ALICE

ALICE consider that the most crucial parts for running analysis applications on GRID platforms is the API for interaction of the applications with the GRID services. The requirement is to access these services, for example a GRID file catalogue, from within an application as from the conventional shell environment.
We have been working on a solution to provide a pseudo-stateless implementation of a grid access service (GAS) to be used first in the framework of the AliEn grid and later on within the EGEE gLite prototype framework.
In general there are two main implementation possibilities for a grid access service:

- a factory-like implementation with one or several user-exclusive instances of a stateful GRID daemon per API or UI session (a pattern present in the OGSI specification)
- a pool of pseudo-stateless grid access daemons working on a call-by-call basis shared by the user community

From experience with middleware development in AliEn the first possibility has shown many traps and problems if many users or many jobs of a single user instantiate a user session in parallel. In a production setup with more than 1000 parallel executing jobs the resource and memory limit of the access server could be easily reached. For the implementation with a server pool one can limit the number of shared daemons to the resources of the server machine.
The call-by-call approach implies a special layout of the access service with respect to authentication and performance.

The typical use case for a grid service access is that a client specifies a certain input command with input arguments. These are transferred to the grid access service and executed directly there or by other remote services. The results are fetched back to the client. The amount of data returned by these services is small in terms of data volume but large in terms of data items. This has direct implications to the useable protocols for this information exchange. Several SOAP encoding schemes with gSOAP (a C++ implementation for the SOAP protocol) have been tested and compared, but evidently the encoding mechanisms for arrays and key/value pairs as strings, floating point and integer numbers didn't show a good performance.

Based on this experience we have implemented a C++ class to en/decode scalar, array and hash variables into/from a single string object which is transferred via SOAP as a standard SOAP string. To fulfill security requirements the encoded string object is encrypted using symmetric CIPHER mechanisms provided by OpenSSL. Every client needs to authenticate once at the beginning of an API or UI session to a dedicated authentication service, which runs on the same machine as the pool servers. During this step the symmetric CIPHER encryption mechanism is initiated and synchronized between the client and the server and stored in a shared memory area (accessible to each pool server process). Every client call uses the assigned client session identifier to indicate to the server which decryption CIPHER has to be used for the received SOAP message. Since the CIPHER is changing with every client call, it is impossible to spy a SOAP message and reuse again to execute a certain task on the grid access service. For the first prototype we implemented Kerberos authentication. A job token and GSI based authentication will be implemented in the future.

The client/server/authentication architecture and protocol has a common layout and can be used for all kind of GRID services or database access services. In our implementation we execute directly on the server side via a C to PERL interface class the PERL API of AliEn (in the future gLite). With a very small modification in the AliEn PERL code it was possible to let the server execute calls with the identity of different users as indicated by the session identifier.

The client code is provided as a C++ shared library which has been used already to write a TGrid plug-in for ROOT. Moreover a set of standard commands like “find”,”ls”,”ps” for the GRID framework are implemented as shell commands i.e. “glite_find”,”glite_ls”,”glite_ps”.

First tests have shown that the call-by-call and session model for a grid access service overcomes the scalability problems of other implementations and provides a very lightweight framework for client APIs and User Interface services. A first demo of the gLite C++ access library has been shown at the ARDA workshop (21-23 June 2004).

ATLAS

The ATLAS prototype will follow the DIAL analysis model, which is part of the distributed analysis strategy of ATLAS (http://www.usatlas.bnl.gov/ADA/). The present DIAL system (version 0.9) is being used as a starting point.

Test jobs have been run and feedback provided to the EGEE developers. In general, the EGEE test bed is not yet ready for stress or performance testing and the ATLAS requirements to run ATHENA applications are not yet met from the present prototype:

- Lack of resources. ATHENA requires considerable memory (when ran in analysis mode). This can not be done efficiently on 256 MB machines. We expect more modern machines to be installed beginning on next quarter.
- Software distribution. Presently ATHENA is not installed on the worker nodes but seen via AFS: this is clearly a workaround and a general solution must be found (either from ATLAS or from LCG-2).

The present status is that ATHENA does run on the gLite prototype and the submission exploits part of the DIAL system. The above mentioned limitations are being addressed.

The collaboration with the DIAL team will include (on the ARDA side) setting up a Savannah site for DIAL (CVS, documentation and bug tracking infrastructure (using the service is ran by LCG).

The second line of activity is the testing of the AMI metadata catalogue. This is very encouraging and a report has been prepared (URL). This report serves as a basis of discussion and further iteration with the ATLAS AMI development team (ATLAS Grenoble). Several limitations observed have been fixed and more general issues (technology choices) are under discussion. In the ARDA workshop, the session on metadata services on the grid was animated by S. Albrand, the AMI leader. Since the AMI team is interested to evaluate Oracle as a possible back end, this activity is under the responsibility of our Taiwan colleagues (in ARDA and local know how).

In the first phase of the project (in particular before the agreement on the ATLAS prototype) the group contributed effort on the ATLAS production side, developing the job submission interface AtCom. This activity had interesting side effects in order to understand the current organization and management of data in ATLAS, understand the current metadata catalogue (AMI) used by ATLAS.
Some effort is still needed to provide the required extensions necessary to work with the database schema changes for Combined Test Beam.

CMS

The ARDA project and the CMS experiment have agreed on the main technology areas and software tools in the ARDA-CMS collaboration towards the first analysis prototype. In the initial phase, the ARDA project will concentrate on data management aspects and executing CMS analysis software in the EGEE-gLite middleware environment. This is an interim agreement waiting for the final decision.

The RefDB (“Reference Database”) software has a central position in the Monte Carlo production within the CMS experiment. Even though originally designed for production book-keeping in parallel to the McRunjob submission tool, the data contained in this system are vital for any CMS analysis framework. Within the reporting period, a comprehensive stress test was carried out in order to ensure that the RefDB system can meet the performance requirements of the first analysis prototype. The draft report is available on the ARDA web site.

In addition, a system for publishing data collections, produced during data challenge, was implemented in the McRunjob and the RefDB software. These software tools were also extended to support the new structure of COBRA META files in order to improve handling of COBRA META information in CMS analysis and production. Moreover, a set of tools for analysis support has been developed and put available under the CVS repository (OCTOPUS/AnalysisTools). These tools mainly provide the interface to RefDB for extracting of the XML POOL catalogue, or list of logical file names or COBRA META files from RefDB for any available data collection. Finally, the ARDA project has started together with the CMS experiment to contribute to the redesign of the RefDB system in order to improve the data management functionality and the user interface. The final role of ARDA in this development has not been decided yet.

The pre-production EGEE-gLite middleware test bed has been available for the ARDA project since mid May. Even though the system should not be used for massive production, the very first CMS analysis jobs have been successfully run on the test bed. For the moment the main problem is the unavailability of a storage element attached to the CASTOR tapes (worked around by manually staging in input files).

We have also started to register files, which were originally produced in the CMS DC04 data challenge, into the gLite file catalog. The next step is to understand the functionality and the performance of the gLite metadata catalog in relation to the CMS analysis framework.
Now the gLite file catalogue is being exercised with the entries of the DC04 RLS.

LHCb

The ARDA activities in collaboration with the LHCb experiment focus around three main topics. These include the GANGA (http://ganga.web.cern.ch) project a front-end for the configuration and management of physics jobs running in the Gaudi/Athena framework. Gaudi/Athena is the base of all physics analysis relevant activities, e.g., Monte-Carlo data production, detector simulation, data analysis in LHCb.

The second topic is oriented towards the DaVinci project, which is the physics analysis package of LHCb. The ARDA activities include understanding how to integrate and deploy DaVinci in the EGEE middleware prototype as well as to interface it to GANGA. Finally, the LHCb metadata catalog was exported and tested on a separated test-bench.

In the second quarter, ARDA participated in a 2-day GANGA workshop, which defined the goals of the GANGA project for the near future. One of the goals defined was that project had to embrace the end-user in order to receive more immediate feedback. ARDA since has had an active role in the deployment of the GANGA project within LHCb by taking over the role of release manager of the project, thus allowing more scheduled and predictable releases. In addition, ARDA contributes by testing the GANGA software assuming the role of a physics user and filing bug reports, both for real functionality misbehaviors as well as inconsistencies within the software.

DaVinci is one of the LHCb software projects, based on the Gaudi/Athena software framework. As such, it is suitable as an example to explore the use of LHCb physics software within the activities of ARDA. After the initial exposure to the software, a small simple module “called an algorithm” was written. Algorithms are user-supplied packages, which run within the framework to do the actual data analysis, to be performed by the physicist. This user-supplied algorithm was used to mainly test the functionality of GANGA, since GANGA is has direct support for handling the submission DaVinci jobs.

The metadata catalog of LHCb, called a “bookkeeping service”, has the purpose to allow storage and retrieval of data not directly related to experimental data taking. The focus of the bookkeeping database is to store/select/retrieve the location of data files. The CERN Oracle server hosts the bookkeeping database. A central server that deals with both web pages and XML-RPC services hosts the bookkeeping services. The Bookkeeping server is based on code providing both an XML-RPC server and a web server, written in python. In order to make detailed tests the database was exported and the service was set-up on a different machine (in Taiwan). The system was stress tested and revealed certain flaws, which are not acceptable for an experimental wide service to be used by anybody, but which was adequate for the original goal as a production tool. The Taiwan group is using local Oracle expertise to study and possibly improve the data base backend organization and performance.

Collaboration with outside institutes

The activity with external collaborators (a total of 4 people from Russian institutes and Taipei) is progressing (note that these colleagues are only part time at CERN).
During the present quarter, the activity was focused on data-management related issues: detailed survey of protocols (used directly by the experiments for efficient data access; one document publicly available, one in preparation) and data base technology (see ATLAS and LHCb sections).

gLite functionality tests

This is a list of tests done within the ARDA project. This is not a formal list but it is intended to give a preliminary idea of the coverage of the investigation effort.

Component	Description	State	Comments
File catalogue	Registering files	OK	Unstable; Limitation on filename length affects CMS
File catalogue	Registering external links	OK	
File catalogue	Stress test: registering 100,000 files in the same catalogue directory	OK	Slowdown of registering time per file was noted (no quantitative results yet)
Metadata catalogue	Applying metadata tags to directories. Setting and retrieving values and searching for files.	OK	
Metadata catalogue	Stress test: Applying a metadata tag to a directory containing 100,000 files	Catalogue corruption	
SE	CERN Castor via SRM	OK	Only disk back-end available
SE	CERN dCache via SRM	Not yet activated	
SE	External Site via SRM (Wisconsin)	Not yet activated	Recently become available
CE	Simple jobs	OK	
CE	Jobs with input file staging from catalogue SEs	Failed	Cannot test currently because of catalogue corruption
CE	Jobs with input file staging from local machine	OK	
CE	Jobs using files from AFS via catalogue references	OK	
CE	Jobs using files from CERN CASTOR via catalogue references	Failed	Direct CASTOR access still missing (critical)
CE	External site (Wisconsin)	OK	Recently become available

The following list summarises the status of the interaction of the experiments' application software within the gLite prototype. This is not a formal list but it is intended to give a preliminary idea of the status of the investigation effort.

Experiment	Description	State
ALICE	Developing new C++ API and ROOT plug-in for gLite (in form of a grid access service)	Ok (In progress ca 90%)
ATLAS	Stripped-down Athena submitted to gLite (via DIAL)	Ok
CMS	Simple ORCA job via executable and data from AFS	Ok
CMS	ORCA job with file staging from CASTOR	The two problems due to long CMS filenames and the missing CASTOR access still missing circumvented with workarounds
LHCb	A first GANGA plug-in able to submit jobs and poll their status	Ok
LHCb	GAUDI jobs via executable from AFS	Ok (User-defined algorithm executed)

ARDA Workshop activity

A workshop was organized and ran on June 21-23. The workshop was by invitation only to allow technical discussions (very difficult if the size would have been as the first ARDA workshop (organized by LCG at the beginning of 2004). The description of the workshop, participants, agenda and corresponding presentations can be found on the web under http://lcg.web.cern.ch/LCG/peb/arda/LCG_ARDA_Workshops.htm. A summary (whose important part is a prioritized list of functionality) is in preparation (edited by the ARDA project leader together with the experiments interfaces).

A new open round of discussion (3rd ARDA workshop) has been agreed and is being organized by October. ARDA is organizing the join NA4/JRA1 and NA4/SA1 (middleware and operations) session in the EGEE event mid of July, which will serve as an interim workshop (agenda: http://agenda.cern.ch/fullAgenda.php?ida=a041952).

Status of Milestones and Outlook

The milestones belonging to this quarter are listed in the following table. As explained, one out of 3 (1.6.4) is late (late definition of the ARDA prototype activity within CMS). The impact of this delay is partially mitigated by the activity on CMS-related components which would have been needed anyway.

WBS.	Task	Date due	Date done
1.6.2	E2E ALICE prototype definition agreed with the experiment	31-05-2004	31-05-04
1.6.3	E2E ATLAS prototype definition agreed with the experiment	31-05-2004	31-05-04
1.6.4	E2E CMS prototype definition agreed with the experiment	31-05-2004	Late
1.6.5	E2E LHCb prototype definition agreed with the experiment	31-05-2004	31-05-04

Next quarter milestones are listed in the next table. All this activities have as principal components the usage of the gLite prototype together with experiment-specific software. All these activities have started and look promising. The main difficulties we envisage are:

1. Stability of the gLite prototype
2. Availability of significant resources (o(100) CPU sites, at least 3 sites connected) to attract serious users
3. Effective data access to significant fraction of the experiments data store (CASTOR storage element not yet available)
4. Coherent software distribution mechanism for the experiment software

WBS.	Task	Date due
1.6.6	E2E ALICE prototype using basic EGEE middleware	30-09-04
1.6.7	E2E ATLAS prototype using basic EGEE middleware	30-09-04
1.6.8	E2E CMS prototype using basic EGEE middleware	30-09-04
1.6.9	E2E LHCb prototype using basic EGEE middleware	30-09-04

6. ALICE – Federico Carminati

The AliRoot framework deals with the event generation, particle transport, and simulation of detector response, creation of summable digits, event merging and digitization, clusterization/creation of reconstructed points, track reconstruction, particle identification and analysis.

The AliEn GRID system developed in ALICE implements a lightweight but complete grid solution that provides the functionality needed for production and analysis of HEP data.

During the second quarter of 2004, the AliRoot framework has continued evolving. The version used in Physics Data Challenge (PDC04) has been frozen, however several bug fixes and "back porting" of improvements has taken place. These are necessary for the analysis of data produced by PDC04. New developments are going on in the event generators, particle transport and reconstruction to be finalized for a next short Physics Data Challenge (PDC05) in mid-2005.

The AliEn system is actively maintained and improved to meet the needs of the '04 physics data challenge. All modifications are discussed with the gLite developers to assure compatibility and portability into the gLite framework.

Code development

The following development has been successfully carried on.

Event generation

- A model for jet quenching has been implemented in Pythia.
- Collaborating classes have been designed for the description of nuclear geometry (Glauber model) and quenching weights (Wiedemann-Salgado model).
- The forced decay of B mesons (B into D+e, and D into p+K) has been implemented in Pythia.
- The DPMjet generator has been added to the generator list and comparison with Pythia are going on.

Particle transport

- The new tagged version v0-6 of the Geant3 package is used in PDC04.
- The FLUKA interface has been completely implemented and detailed testing is in progress.
- The Geometry Modeler is ready to be used as unique description tool of the detector geometry in the AliRoot framework. Validation is essentially finished for GEANT3 and is starting with FLUKA.
- A more realistic geometry and improved detector response has been implemented for TRD.
- A new geometry has been implemented for V0 and a pp test production has been performed.

Event merging/digitization

- The new steering class for all the simulation stages (event generation, particle transport, creation of summable digits, event merging, and digitization) is used in PDC04.

Reconstruction

- The new steering class for the reconstruction is used in PDC04.
- The new improved and more efficient TOF reconstruction has been implemented.
- A version tuned for efficiency has been implemented for the ITS and TPC reconstruction, including the reconstruction of secondary tracks.
- Work is progressing in improving the ITS reconstruction following the experience gained with the TPC reconstruction.
- The HLT reconstruction has been included in the standard reconstruction chain and the resulting tracks are stored to the ESD.
- Root Trees have been adopted and implemented as data container for the ESD.
- The global PID is being improved to handle mis-matching information extracted from various subdetectors.
- The framework to take into account the RAW data format has been developed. Algorithm to transform simulated digits into raw data and to start reconstruction from raw data is being developed subdetector by subdetector.

Analysis

- Analysis from ESD object is being generalized.
- A new module for jet analysis has been added.
- A general purpose AOD class has been developed and is under test.

Collaboration with the HLT Project

- The Hough reconstruction algorithm has been significantly improved for speed and efficiency.

General improvements

- Makefile improving the handling of dependences and allowing for a safe parallel compilation has been developed.

Data Challenge

In preparation of PDC04, the AliEn procedures for bookkeeping, job submission and job information services have been improved.

The monitoring has been extended using both AliEn native tools and through an implementation in the MonALISA framework.

The PDC04 main physics objective is to study in details the hard physics for the PPR. The required statistics is set by the needs for jets simulation and by the goal of verifying the capability to simulate and process events, at a 10% scale, of a standard data-taking year. The Data Challenge will help in improving code performances and will provide a validation for the computing model.

The DC consists of three phases. In the first phase, 10^5 simulation events are produced in a distributed way and stored at CERN. This phase is has been successfully completed in June. In the second phase, the events are distributed from CERN to remote centers for reconstruction. Produced data are stored in the remote local SE and a zipped backup copy is stored in CASTOR at CERN. This phase has started beginning of July and will last at least until mid September, provided the requested resources are continuously available. The third phase consists of distributed analysis (both chaotic and ordered) including the parallelization of analysis tasks with PROOF. This phase will start mid-September.

In phase 1 a maximum of 1500 jobs ran concurrently, distributed worldwide. The average number of jobs over the three months duration of this phase was 650. This corresponds to a computing capacity of 285M SI-2K. 26 TB of data have been produced and stored in CASTOR at CERN as 1.3M files. 3.8M files have been stored in the AliEn catalog. There were 12 major (more than 2% 0f total amount of jobs) production sites and 24 sites all together. The LCG2 share was 24%.

Phase 2 started 2 weeks ago and has processed 10K events out of 1.4M events to be run with various physics configuration. The output of the jobs is in total 8 TB and will be stored on local SE.

Contribution to the Computing MoU

ALICE continues to provide input to the draft MoU. ALICE has also produced a very preliminary draft of the ALICE dependent MoU which will take the form of an addendum to the existing M&O MoU. Version 3 of the C-MoU and of the addendum will be discussed by the ALICE Management Board in September.

7. ATLAS – Dario Barberis

Application Area

During the second quarter of 2004, the whole ATLAS software developers' community has focused on getting the software ready for the support of the Combined Test Beam (May-October 2004) and the operation of Data Challenge 2 (June-October 2004). Very little new functionality coming from LCG Application software developments has been added to the ATLAS software base, all attention being given to integration and performance problems of existing ATLAS developments. ATLAS collaborators during this quarter contributed work to the LCG Applications Area on:

- POOL event collections;
- Conditions Database (on API extension refactoring, extraction tools, hierarchical versioning system);
- setting up CMT interfaces for LCG External packages;
- porting LCG and external code to the MAC OSX platform
- physics validation of Geant4 and Fluka simulation.

LCG-2 Deployment

We started the operation of Data Challenge 2 in June 2004. It makes use of the LCG-2 Grid for ~60% of the production (the rest being distributed between Grid3 in the USA and NorduGrid in the Scandinavian countries and elsewhere). We had considerable problems of robustness of the system (including our own parts), which delayed considerably the start and also the ramp-up of the production rate. Exchange of information with the LCG Grid Deployment Area is generally good but the interactions with some other services provided by CERN IT Department have been less than ideal. Reported faults have been addressed but not always solved on the necessary timescale; work-arounds for some of the problems are now implemented as part of the ATLAS production system. Much more human work turns out to be needed to manage the production on the LCG-2 system than we anticipated.

Here are some major problems that occurred at the start of DC2 operation on the LCG-2 Grid:

- Tuesday June 29: Service machine, hosting ATLAS BDII, rebooted: configuration files lost, no more BDII, 2 days lost.
- Wednesday July 7: 2 new machines were allocated to ATLAS to improve the job submission rate. It took 12 days to get them visible from outside CERN, as the necessary ports were blocked by the CERN firewall (not a very useful configuration for a Grid machine).
- WP1: a problem was found in the C++ API in WP1. After discussion help was provided by people who were working previously in EDG. It took some time to identify the problem. The fix was in the end provided in an efficient way.
- Non-scalability of GRIS: well-known problem, which was fixed by a patch in LCG m/w, included in the latest release, but not yet installed everywhere.

ARDA

The ARDA project started up in earnest this quarter. ATLAS and ARDA reached agreement on a plan for delivery of an ATLAS-ARDA prototype. See details in http://www.usatlas.bnl.gov/ADA/docs/ATLAS_ARDA_4.doc. The ARDA team started deploying a DIAL service on a machine at CERN and made preliminary connection to the gLite prototype. The ARDA team has put up a Savannah site for DIAL. They have also been studying the connection to AMI, the ATLAS book-keeping service.

Verification milestones

1.1.2.1 10196 V_ATLAS ATLAS POOL validation with DC1 data 19-01-2004
Superseded by events and lack of manpower: validation was done at a lower scale with DC2 pre-production tests. Now obsolete, should be removed from the list.

1.1.2.1 10199 V_ATLAS ATLAS validation of POOL Metadata/event collections 01-02-2004
Integration into ATLAS software scheduled for ATLAS Release 9 (11 August 2004). Completion date for this milestone should be end of DC2 (i.e. successfully exercising the functionality in the course of DC2 will constitute the validation).

1.1.2.1 10197 V_ATLAS ATLAS POOL validation with complete Event Data Model 01-04-2004
Integration into ATLAS software scheduled for ATLAS Release 9 (11 August 2004). Completion date for this milestone should be end of DC2 (i.e. successfully exercising the functionality in the course of DC2 will constitute the validation).

1.1.3 10200 V_ATLAS ATLAS integration: SEAL Integration into Athena 31-12-2003
Dictionary, Python and all that is necessary for POOL already there since POOL integration (Release 7, Autumn 2003)
Plug-in Manager and other components now scheduled for ATLAS Release Cycle 10 (October 2004).

1.1.5.1 10203 V_ATLAS ATLAS integration: Generic simulation framework validated by ATLAS 01-05-2004
The Apps Area has essentially descoped this to the point that it's not sure that we'll use it for Fluka simulation/validation. It should not be a milestone any longer.

1.4.1.13 4.1169 V_ATLAS ATLAS production on LCG-1 29-02-2004
ATLAS production started on LCG-2 on 24-06-04. Delays have been introduced both in the preparation of the ATLAS software and of the ATLAS production system.

8. CMS – David Stickland

In March and April 2004 CMS undertook a Data Challenge (DC04) whose purpose was to demonstrate the ability of the CMS computing system to cope with a sustained data-taking rate equivalent to 25Hz at a luminosity of 2x10**33 for a period of 1 month. This corresponds to 25% of the LHC startup scenario. The Challenge consisted of the three phases: (a) Reconstruction at Tier-0 (CERN) at 25 Hz; (b) Distribution of reconstructed data to Tier-1's; (c) Physics-groups analysis at the Tier-1's as data arrived. Six of the CMS Tier-1 centers and a few of the Tier-2 centers participated in DC04. The Tier-1 operations were at CNAF(Bologna), FZK (Karlsrhue), FNAL(Chicago), IN2P3(Lyon), RAL (Rutherford), PIC (Barcelona); INFN (Legnaro) played a major role as a Tier-2 center, Florida, CIEMAT and others also actively participated to the Tier-2 Analysis studies.

The data challenge was preceded by a pre-challenge production phase, started in July 2003. About 50 millions events were required to match the 25 Hz rate for a month. The CMS physicists requested actually more than 70 millions events. These were simulated and about 35 million of them were digitized and transferred to CERN in time for the start of DC04. The total amount of data produced up to the start of DC04 was of about 100 TB in more than 700,000 files, corresponding to about 400 KSI2000 years (more than 500 years of CPU on a Pentium 4 2.4 GHz). About 35 computer centers in Asia, Europe and US participated in data production. A set of tools, OCTOPUS, provided the needed functionality to run in the heterogeneous environment, where some of the computing centers did not make use of grid tools and the others used two different grid systems: LCG in Europe and Grid3 in the USA.

The components deployed in DC04 varied widely in terms of engineering, development, and stability. Several challenge components, such as the CMS software framework and the production tools, have been developed over several years and demonstrated at scale. A number of grid components, including those used for wide area transfer and file cataloging, were engineered to meet specific use-cases, but had not been integrated and demonstrated in a system of this scale prior to the challenge. Components such as the database used to manage transfers and several of techniques used for real time analysis were developed specifically and rapidly for DC04 to solve specific cases of missing functionality.

During DC04 digitized events (raw data) were staged from CERN Castor Mass Storage System (MSS) and were processed by jobs submitted to a computer farm consisting of about 500 CPU's (about 300 KSI2000 CPU power). The overall rate was determined by the job submission speed: the required 25 Hz is given by more than 2000 jobs (1000 events each) per day, corresponding to 40 MB/s staged from the Castor MSS, 380 CPU's running at 100%, and 4 MB/s of output.

Reconstructed data (DST) were automatically archived to tape and distributed to each of the Tier-1 centers. An ad-hoc developed Transfer Management DataBase (TMDB) has been used to steer data distribution. The RLS catalogue at CERN was used to locate the physical instances of the data files CMS-wide and to classify them as a function of their content.

Several data transfer tools have been supported: the LCG Replica Manager tools, native SRM (Storage Resource Manager) and SRB (Storage Resource Broker). Each tool had its own dedicated export buffer. All transfer tools used the RLS catalogue to store information about physical instances of the files.
Each data transfer tool had a dedicated agent running at the Tier-0 responsible for making data available on the appropriate Export Buffer. At each Tier-1 there was an agent responsible for discovering files to be transferred to that center from the TMDB and performing the replication using the preferred transfer tool to a safe location (MSS). In some cases selected files were simultaneously replicated from the Tier-1 center to Tier-2 centers. The TMDB always had a precise picture of the file status starting from its appearance on the computer farm up to its replication to Tier-1's MSS systems.

Selected Tier-1 sites ran Physics Group Analysis in real-time. The processing jobs were submitted as soon as all the files of a given run were available at the center. Jobs were either submitted to the local batch queues on the Tier-1 centers or submitted through grid interfaces. Grid enabled jobs were submitted to the LCG Resource Broker (RB) and ran on Computing Elements close to the Storage Element (SE) hosting the data for analysis. In particular in some regions data were made available for analysis on dedicated disk Storage Elements at the Tier-1 and at some Tier-2 centers, depending on the data content. The RB has been able to select the appropriate site where to submit the jobs. The job output was stored on a SE and registered to the RLS and thus made available to the whole collaboration. During the last days of DC04 running an average latency of 20 minutes was measured between the appearance of the file on the CERN computer farm and the start of the analysis job at the remote sites. On the last day of the challenge the network usage from CERN achieved a new high for production use of almost a Gb/s, with major data flows to Italy, Spain and the United States.

The main activity after the end of DC04 will be the continuous analysis of the DST's as well as of the simulated events. The main task for the Tier-2 centers is to provide to the end-users the ability to access the computing resources and the data (either locally or at some remote site, e.g. a Tier-1). CMS is testing several prototypes of tools that act as an interface to the physicist. They allow creation and submission of jobs to local and remote resources, possibly via the LCG Resource Broker, starting from a predefined syntax that defines the data to be analyzed.

The analysis of the Data Challenge is now underway with plans to present many of the results at this September CHEP conference in Interlaken.

The submission of the Computing TDR of CMS has been moved back to July 2005 to match that of the LCG-TDR with which it must be closely aligned. In the next phase of CMS Computing we will need to meet the requirements of producing, distributing and making useable to the physicists in the PRS groups about 10 million events per month. The next Data Challenge will take place after the completion of the Physics TDR at the end of 2005. For the next 18 months CMS will be operating its productions and analysis in essentially continuous and concurrent mode.

9. LHCb – Philippe Charpentier

Applications area

This quarter has been devoted to the LHCb Data Challenge, DC04, and because of this no major integration of LCG-AA software has taken place until the start of DC04, when the software was frozen, LHCb has been following the releases of SEAL and POOL. The versions currently being used for DC04 are SEAL 1.3.4 and POOL 1.6.3. Tests have started with the latest versions of POOL, in particular those based on Root 4.

The Gauss simulation application, based on Geant 4 version 6.1, is being used for the first time as part of DC04. A few bugs have been reported to the Geant 4 team that were causing crashes during production and patches have been applied to the running software.

DC04 summary

DC04 was started on May 6th. It consists for this first phase of massive data production, running within a single job: simulation, digitisation (including spill-over) and reconstruction. Only the DST produced by the reconstruction stage is then stored. A typical job produces 500 signal events or 500 background events together with 900 minimum bias events. Their duration is slightly more than 2 days NCU and each job produces 350 MBytes of data.
The production is based on the LHCb-developed Dirac workload management system. It allows submission to Dirac-CEs as well as to the LCG2 Resource Broker. The file metadata is recorded in the LHCb Bookkeeping Database (BKDB) that was proven very effective in DC03. Two implementations of the file and replica catalogs are populated: a replica table within BKDB and an AliEn file catalog. Data are transferred using bbftp, sftp or gridFTP depending on the site. All data are replicated at a Tier1 (assigned for each site) and at the CERN-Tier1.
During the first few weeks of DC04 LHCb concentrated on using Dirac-CEs while debugging Dirac and testing submission to LCG2. Actual production was started on LCG2 at the end of May. Typically LHCb was running concurrently 800 jobs on LCG with peaks above 1000 jobs. Several difficulties however were encountered that needed fixes in the LCG middleware. Thus the production on LCG2 was suspended on June 27th to allow fixes to be implemented and new middleware to be deployed. It was resumed on July 17th; experiences with the upgraded system will be reported in the next quarter.
An interim report on LHCb experience with LCG2 has been submitted to the LCG-GDA when pausing the DC04. It is the basis for all modifications on LCG2 requested by LHCb. The support LHCb received from the GDA team during DC04 has been exemplary.
From the start of DC04 to the end of Q2, 100,000 jobs have been successfully run amongst which 16,000 on LCG resources. They correspond to 63 million events. At the time DC04 was paused on LCG2, about 50% of the jobs were running on each system. The efficiency however of LCG2 was around 60% while that of Dirac was 95%. Most losses on Dirac-CEs were due to system problems (network failures, site mis-configuration e.g. firewalls; less than 1% of inefficiency is caused by crashes in the LHCb applications.
In order to thoroughly test the improvements introduced in LCG2, it was requested to extend the DC04 production phase until end of August.
The next phases of DC04 are:

- September 04: analysis production. Tagging and selection algorithms will be run as production jobs on LHCb Tier1s in order to prepare data for analysis by physicists. The stripped DSTs will be deployed on all Tier1s
- October 04 onwards: physics analysis. Physicists will be able to fully analyse these data, progressively using the GANGA interface to the Grid.

ARDA and gLite

The two members of the ARDA group associated with LHCb were extremely effective during Q2. They started work on testing and deploying GANGA (the LHCb-ATLAS analysis user interface to the Grid). An agreement on plans for delivering an LHCb ARDA prototype was agreed upon at an early stage of the project.
As gLite was made available, LHCb-ARDA team started the integration with Ganga in order to fully test its functionality and performances.

The LHCb BKDB was made available to the ARDA team for evaluation as a file-metadata catalog; this evaluation was undertaken in Taiwan. The LHCb file and replica catalog interface was also made available to ARDA as a possible interface for gLite. It is already used within LHCb for interfacing both to our BKDB and AliEn FCs.

LHCb made a substantial contribution to the ARDA workshop in June.

The LHCb Dirac team is following very closely the developments in the gLite area and are willing to contribute, particularly in the area of the Workload Management System and in testing gLite components. We look forward to being able to contribute to the gLite project and hope for a closer working relation with the EGEE-gLite team.

Verification milestones

1.4.1.14 V_LHCb Dirac deployed in LCG1 (now LCG-2) 28-11-03
Was achieved on 24-05-04 when DC04 started to use LCG2. Mainly due to the late delivery of the middleware and Dirac

1.4.1.15 V_LHCb 50% of LHCb DC04 production executing in LCG-2 01-06-04
Was approached towards after 22-06-04 with constant 40% running on LCG and peak at 57%. However with too high an inefficiency on LCG2. The milestone full text specifies “with LCG2 efficiency higher than 95% of that of Dirac sites”.

1.4.1.16 V_LHCb LHCb analysis environment (GANGA) deployed in LCG2 01-07-04
Although formally in Q3, this milestone should be rephrased as “GANGA deployed on LCG2 through a Dirac interface” as this is the baseline of LHCb. It is foreseen to be met for August 31st.

1.4.5.4.4 Report on LHCb initial experience 01-05-04
Achieved on 30-06-04 with a report submitted to the LCG-GDA team. Was late due to the late start of LCG2 usage.

10. Milestone management and status – Jürgen Knobloch

The time distribution of milestones and the current achievement is shown in Chart 1.

Chart 1: Number of LCG Level 1&2 Milestones (cumulative)

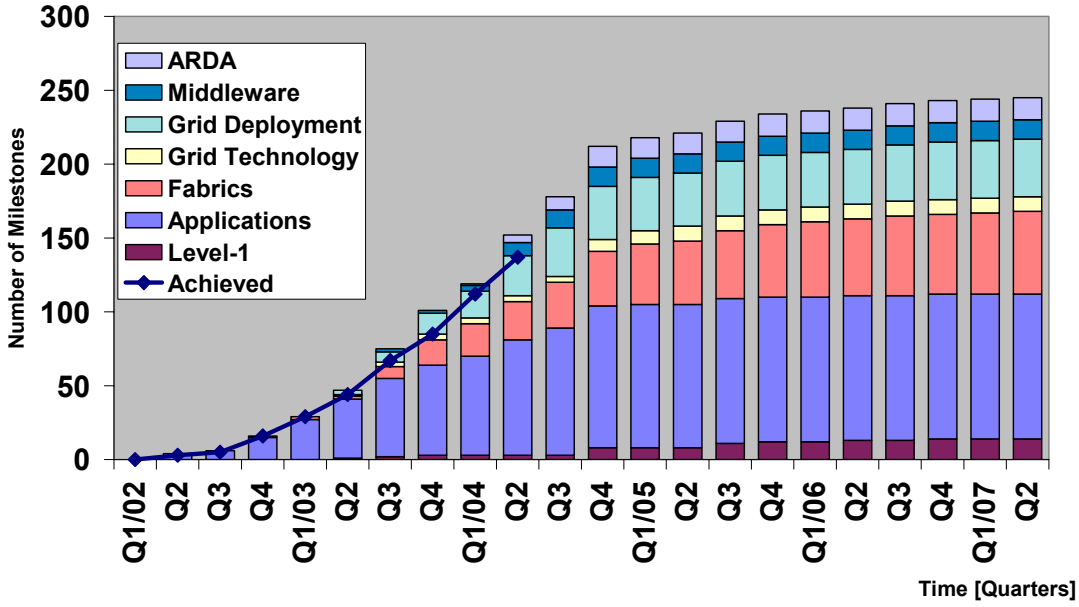

The delay of milestone achievement is shown in Chart 2. The fraction of milestones with a delay larger than 4 weeks is 17%.

Chart 2: Delay of achieved milestones [weeks]

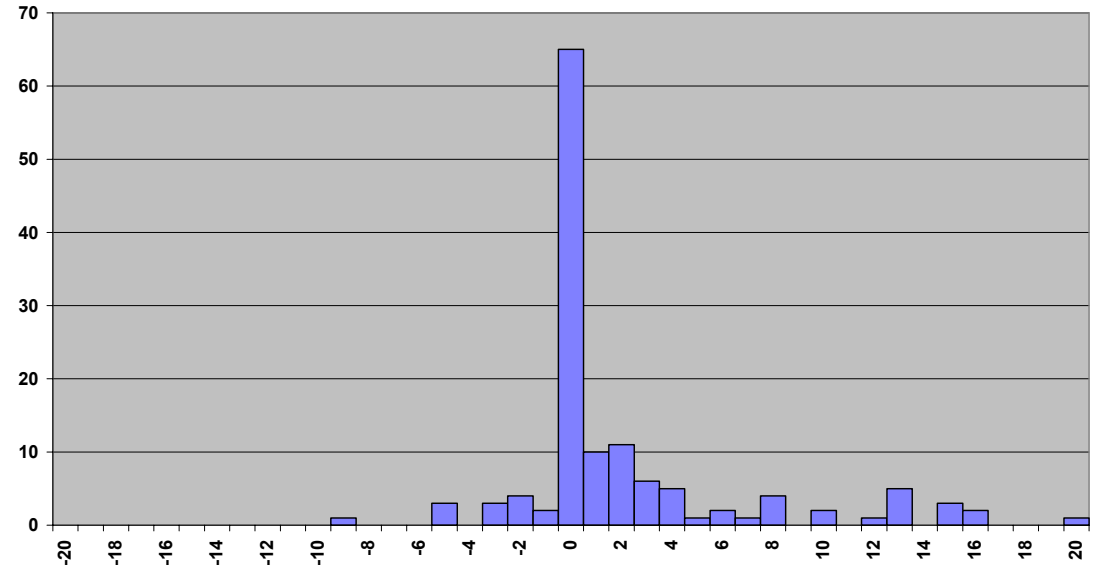

The number of milestones in the different areas and milestone levels is given in Table1:

Milestones		Number
Level-2	Applications	98
	Fabrics	57
	Grid Technology	10
	Grid Deployment	39
	Middleware	13
	ARDA	15
Level-1		15

Table 1: Number of milestones

A summary of all currently defined level-1 and level-2 milestones is given in Appendix 1. The tables are also available from the Web providing links to a more detailed description of milestones. The web address of the milestones table is:

http://lcg.web.cern.ch/LCG/PEB/Planning/MilestonesJuly2004/

Milestones that have not yet been achieved by now are given in the following Table 2. Details of the missed milestones are given in the individual area reports:

WBS	ID_Unique	Comment	Name	Date_MS
1.1.1	1.212	i_apr04	SPI/EGEE collaborative workplan complete	31-05-2004
1.1.2.1	1.22		RDBMS independency achieved for POOL relational components	30-06-2004
1.1.2.2	1.221	i_apr04	Common interface for Conditions DB defined	30-06-2004
1.1.3	1.203	i_apr04	External software guideline document	31-05-2004
1.1.3	1.205	i_apr04	New Dictionary API and reference implementation	30-06-2004
1.1.5.4	1.213	i_apr04	Review/prioritization of simple benchmarks for simu physics validation	31-05-2004
1.1.2.1	10199	V_ATLAS	ATLAS validation of POOL Metadata/event collections	01-02-2004
1.1.2.1	10197	V_ATLAS	ATLAS POOL validation with complete Event Data Model	01-04-2004
1.1.3	10200	V_ATLAS	ATLAS integration: SEAL Integration into Athena	31-12-2003
1.2.1.1.3	2.519		Production system ready for deployment	01-04-2004
1.2.1.3.3	2.524		Proposal on storage architecture and implementation of disk storage system for Phase 2 made	01-06-2004
1.2.4.1	2.564		Baseline plan for Phase 2 agreed (architecture + technology)	01-06-2004
1.4.1.15	4.1185	V_LHCb	50% of LHCb DC04 production executing in LCG2	01-06-2004
1.4.5.4.3	4.1628	i_jan04	Report on ATLAS initial experience	01-05-2004
1.4.5.4.4	4.1627	i_jan04	Report on LHCb initial experience	01-05-2004
1.4.5.5.1	4.1644	i_jan04	WN middleware ported to RHEL, report on ease of porting to other OSs	01-05-2004
1.6.4	6.1897	i_apr04	E2E CMS prototype definition agreed with the experiment	31-05-2004

Table 2: Milestones not yet achieved

In the course of the project, certain tasks have become obsolete and some milestones have been superseded by more refined ones. We have therefore removed such obsolete milestones from our system. They are given in the following Table 3:

No.	WBS	Comment	Name	Date_MS
1	1.1.5.1		Generic simulation framework prototype available (G4 and FLUKA)	31-12-2003
2	1.1.2.1	V_ATLAS	ATLAS POOL validation with DC1 data	19-01-2004
3	1.1.3	V_LHCb	LHCb integration: SEAL plugin manager integrated in Gaudi	19-12-2003
4	1.1.5.1	V_ATLAS	ATLAS int: Generic simulation framework validated by ATLAS	01-05-2004
5	1.2.1.2.1	V_ALICE	ALICE DC5 (300 MB/sec)	01-10-2003
6	1.2.1.2.2	V_ALICE	ALICE Verification of target performance for ADC V	15-11-2003

Table 3: Milestones deleted

11. Resources

Materials Resources at CERN

The negative balance of 1.5 MCHF announced at the April 2004 C-RRB for the year 2004 could be reduced to 1 MCHF by shifting payments for certain contracts into 2005.
Recent discussions with the CERN directorate confirmed that a solution for the remaining 1 MCHF will be found. This has however delayed the installation of equipment at CERN.

Human Resources

The following tables summarise the human resources working on the project in July 2004 at CERN (Fabric, Grid Technology, Grid Deployment, Experiment Support Areas) and at CERN and in outside institutes (Applications Area). These tables contain now FTE numbers for human resources provided by EGEE to the LCG project. Overheads and effort allocated to non-LHC activities in EGEE are not counted as well as people working in areas outside the scope of LCG. In particular the EGEE funded effort on middleware development, which takes place at CERN, in Italy and in the UK, is not included. There are of course also significant human resources working in Regional Centres providing computing services and involved in grid deployment. The Rutherford Appleton Laboratory (RAL) also operates a Grid Operations Centre, and the Forschungszentrum Karlsruhe (FZK) operates a Grid Call Centre. The Academia Sinica in Taipei has opened a second combined operations and call centre on May 24, 2004. There are also resources in India and Russia working on projects for which formal agreements have been put in place and which are included in the resources reported to the Computing Resource Review Board (see Report on Resources to the April 2004 C-RRB).

LCG Project - Human Resources

Experience-weighted FTEs

Funding source: ***PH, IT*** - funded from CERN base budget as a CERN support activity; ***LCG*** - funded at CERN by special contributions; ***EGEE*** - funded from the EU via the EGEE Project (net value); ***Experiments*** : funding allocated through experiments - ***CERN*** - funded in EP from CERN base budget, ***Insts.*** -funded by an external institute

Experience-weighting factors : trainee, student 0.5; fellow, junior staff to 2 yrs experience 0.85; staff, associates 1.0; experienced staff, associates (>5 years) 1.2

	Experience-weighted FTEs in July-04								Resources used (FTE-years)			
	funding source											
			CERN			*Experiments*						
Applications	*PH*	*IT*	***Support***	***EGEE***	***LCG***	***CERN***	***Insts.***	***Total***	***2002***	***2003***	***1H04***	***Total***
Software Process Infrastructure	*1.4*	*1.2*	2.6		2.0		0.1	4.7	3.9	7.1	2.8	13.7
Object Persistency (POOL)		*2.5*	2.5		2.9	1.0	3.4	9.9	9.4	12.8	5.2	27.3
Core Libraries and Services (SEAL)	*1.4*		1.4		1.1	0.5	2.2	5.1	5.7	5.7	2.8	14.1
Physics Interfaces (PI)	*1.0*		1.0			0.7		1.6	0.6	2.3	1.3	4.2
Simulation	*11.0*		11.0		4.1	1.3	0.1	16.5	11.2	17.1	8.3	36.6
GRID interfacing					1.0			1.0	1.0	2.6	0.8	4.4
Architecture						0.2		0.2		0.2	0.1	0.3
Management	*1.0*	*0.2*	1.2		0.4	0.2	0.1	2.0	0.9	2.0	1.1	4.0
Total	*15.7*	*4.0*	**19.7**		**11.4**	**4.0**	**6.0**	**41.0**	**32.6**	**49.7**	**22.4**	**104.7**
ROOT	*2.4*		2.4		3.2	0.8	0.7	7.1	4.5	7.8	3.9	16.1
CERN Fabric												
System Management & Operations		*8.5*	8.5	2.0	2.4			12.9	5.8	11.3	5.6	22.7
Development (e.g. Monitoring)		*3.4*	3.4		2.9			6.2	7.2	8.8	3.4	19.3
Data Storage Management		*8.2*	8.2		2.0			10.2	5.2	8.4	4.9	18.5
Grid Security					0.9			0.9	2.1	3.0	0.4	5.6
Grid-Fabric Interface		*1.0*	1.0		1.2			2.2	5.7	5.5	1.7	12.9
Internal Networking for Physics		*1.2*	1.2					1.2	1.1	1.1	0.6	2.9
External Networking for Physics		*0.9*	0.9					0.9	0.7	0.7	0.4	1.7
Management		*2.1*	2.1					2.1	2.1	2.1	1.1	5.3
Total		*25.4*	**25.4**	**2.0**	**9.3**			**36.6**	**29.9**	**40.9**	**18.1**	**88.9**
Project Technology												
Grid Technology, middleware, modelling, evaluations		*4.6*	4.6					4.6	1.2	3.7	1.7	6.6
Total		*4.6*	**4.6**					**4.6**	**1.2**	**3.7**	**1.7**	**6.6**
Grid Deployment												
Integration and Certification					9.5			9.5	2.1	9.6	5.1	16.8
Grid Infrastructure, Operations and User Support		*3.6*	3.6	4.5	14.2			22.3	5.6	11.0	9.3	25.9
Management					1.2			1.2	0.4	1.2	0.6	2.2
Total			**3.6**	**4.5**	**24.9**			**33.0**	**8.1**	**21.8**	**15.0**	**44.9**
Experiment Support												
Distributed Analysis (ARDA)		*1.2*	1.2	4.0	5.5			10.7			1.7	1.7
Support in Experiments					2.4			2.4		0.3	0.6	0.9
Total		1.2	1.2	4.0	7.9			13.1		**0.3**	**2.3**	**2.6**
LCG Management		*4.1*	**4.1**		**1.6**			**5.7**	**5.5**	**5.3**	**2.6**	**13.4**
LCG Project Total (inc.Root)	*18.1*	*42.8*	**60.9**	**10.5**	**58.2**	**4.8**	**6.7**	**141.0**	**81.7**	**129.5**	**66.1**	**277.3**

LCG Project - Human Resources

Unweighted FTEs

Funding source: PH, IT- funded from CERN base budget as a CERN support activity; ***LCG-*** funded at CERN by special contributions; ***EGEE*** - funded from the EU via the EGEE Project (net value); ***Experiments:*** funding allocated through experiments - ***CERN -*** funded in EP from CERN base budget, **Insts.-**funded by an external institute

Unweighted FTEs in July-04

			funding source					
			CERN			*Experiments*		
Applications	*PH*	*IT*	***Support***	***EGEE***	***LCG***	***CERN***	***Institutes***	***Total***
Software Process Infrastructure	*1.2*	*1.0*	2.2		2.0		0.1	4.3
Object Persistency (POOL)		*2.4*	2.4		2.7	1.0	3.2	9.3
Core Libraries and Services (SEAL)	*1.5*		1.5		1.2	0.4	1.9	5.0
Physics Interfaces (PI)	*0.9*		0.9			0.7		1.6
Simulation	*10.2*		10.2		4.1	1.1	0.1	15.5
GRID interfacing					0.8			0.8
Architecture						0.2		0.2
Management	*0.8*	*0.2*	1.0		0.4	0.2	0.1	1.7
Total	*14.6*	*3.6*	**18.2**		**11.1**	**3.6**	**5.4**	**38.3**
ROOT	*2.0*		2.0		3.0	0.7	0.7	6.4
CERN Fabric								
System Management & Operations		*8.0*	8.0	2.0	3.0			13.0
Development (e.g. Monitoring)		*3.2*	3.2		3.0			6.2
Data Storage Management		*7.4*	7.4		2.0			9.4
Grid Security					1.0			1.0
Grid-Fabric Interface		*1.0*	1.0		1.0			2.0
Internal Networking for Physics		*1.0*	1.0					1.0
External Networking for Physics		*0.8*	0.8					0.8
Management		*1.8*	1.8					1.8
Total		*23.2*	**23.2**	**2.0**	**10.0**			**35.2**
Project Technology								
Grid Technology, middleware, modelling, evaluations		*4.0*	4.0					4.0
Total		*4.0*	**4.0**					**4.0**
Grid Deployment								
Integration and Certification					9.0			9.0
Grid Infrastructure, Operations and User Support		*3.0*	3.0	4.5	17.0			24.5
Management					1.0			1.0
Total		*3.0*	**3.0**	**4.5**	**27.0**			**34.5**
Experiment Support								
Distributed Analysis (ARDA)		*1.0*	1.0	4.0	5.6			10.6
Support in Experiments					2.0			2.0
Total		*1.0*	**1.0**	**4.0**	**7.6**	**0.0**	**0.0**	**12.6**
LCG Management		*3.6*	**3.6**		**1.5**			**5.1**
LCG Project Total (inc.Root)	*16.6*	*38.4*	**55.0**	**10.5**	**60.2**	**4.3**	**6.1**	**136.0**

Appendix: Milestone Tables

In the “Comment”-field, the date of introduction (i_apr04 = introduced April 2004) as well as experiment verification milestones (e.g. V_ALICE = ALICE verification milestone) are indicated.

Applications Area Milestones

WBS	ID_Unique	Comment	Name	Date_MS	MS_Done
1		___	**LHC Computing Grid Project**		
1.1			**Applications**		
1.1	1.3		Development, support and resource plan through 2008	31-12-04	
1.1	1.4		Phase 1 AA software complete and deployed	30-09-05	
1.1.1			**Software Process and Infrastructure (SPI)**		
1.1.1	1.156		Initial process and infrastructure team in place	15-04-02	14-07-02
1.1.1	1.158		Basic LCG code development and repository services in place	01-05-02	02-05-02
1.1.1	1.159		Software process workplan to SC2	11-10-02	11-10-02
1.1.1	1.16		All infrastructure tools and services in beta	31-12-02	09-02-03
1.1.1	1.161		All infrastructure tools and services released (SPI V1.0)	28-02-03	28-02-03
1.1.1	1.162		2003 SPI workplan complete	30-05-03	29-06-03
1.1.1	1.163		Software library complete	31-05-03	01-05-03
1.1.1	1.164		Savannah based portal complete	31-05-03	22-12-02
1.1.1	1.17		SPI support for Windows binary version of LCG software	15-08-03	15-08-03
1.1.1	1.188		SPI tools operational on IT CVS service	30-11-03	29-01-04
1.1.1	1.199	i_jan04	savannah.cern.ch migrated to GNU savannah	20-02-04	04-02-04
1.1.1	1.22	i_apr04	RH 7.3 gcc 3.2.3 supported	15-05-04	15-05-04
1.1.1	1.212	i_apr04	SPI/EGEE collaborative workplan complete	31-05-04	
1.1.1	1.216	i_apr04	Certification of external software for the new Linux platform	01-07-04	01-06-04
1.1.1	1.201	i_jan04	LCG AA build system selection	15-07-04	
1.1.2			**Physics data management**		
1.1.2	1.103		Persistency framework workplan to SC2	12-07-02	02-08-02
1.1.2	1.109		2003 persistency framework workplan complete	24-02-03	05-03-03
1.1.2	1.2		General release of POOL hybrid data store	15-06-03	30-06-03
1.1.2	1.117		2004 persistency framework workplan complete	15-12-03	03-04-04
1.1.2.1			**POOL Persistency Framework**		
1.1.2.1	1.101		Hybrid event store workshop	05-06-02	05-06-02
1.1.2.1	1.102		Persistency framework status report	18-06-02	18-06-02
1.1.2.1	1.104		V0.1 POOL internal release	30-09-02	02-10-02
1.1.2.1	1.105		V0.2 POOL internal release	31-10-02	15-11-02
1.1.2.1	1.106		First prototype (V0.3) release of POOL hybrid data store	15-12-02	18-12-02
1.1.2.1	1.108		SPI compliance of POOL	28-02-03	07-03-03
1.1.2.1	1.11		POOL V0.4 'interface-complete' release	28-02-03	07-03-03
1.1.2.1	1.111		Define feature set of first general release of POOL	01-03-03	05-03-03
1.1.2.1	1.112		POOL V1.0 pre-production release	30-04-03	13-05-03
1.1.2.1	1.113		Nightly builds deployed in POOL	30-05-03	29-06-03
1.1.2.1	1.176		CMS POOL integration: POOL persistency of CMS event	31-07-03	31-07-03
1.1.2.1	1.177		ATLAS POOL integration: POOL persistency in Release 7	10-09-03	11-09-03
1.1.2.1	1.116		POOL RDBMS independence layer in beta	30-09-03	02-01-04
1.1.2.1	1.181		POOL support for Windows binaries	30-09-03	23-01-04

WBS	ID_Unique	Comment	Name	Date_MS	MS_Done
1.1.2.1	1.114		Initial POOL deployment on LCG-1	15-11-03	24-11-03
1.1.2.1	1.121		POOL hierarchical cataloging production release	15-03-04	09-03-04
1.1.2.1	1.219		POOL RDBMS abstraction layer completed	31-05-04	31-05-04
1.1.2.1	1.22		RDBMS independency achieved for POOL relational components	30-06-04	
1.1.2.1	1.223		First release of the POOL Relational Storage Manager	31-08-04	
1.1.2.1	1.122		POOL meets scalability requirements	31-10-04	
1.1.2.1	1.224		POOL integrates ROOT4	31-10-04	
1.1.2.1	1.218		Full function release of POOL persistency framework	01-03-05	
1.1.2.2			**Conditions database**		
1.1.2.2	1.221	i_apr04	Common interface for Conditions DB defined	30-06-04	
1.1.2.2	1.222	i_apr04	Conditions DB production release	31-07-04	
1.1.3			**Core libraries and services (SEAL)**		
1.1.3	1.123		Establish core libraries and services (SEAL) project	30-10-02	30-10-02
1.1.3	1.125		Define the V1 SEAL software suite	30-11-02	30-11-02
1.1.3	1.107		Prototype object dictionary service released	01-12-02	18-12-02
1.1.3	1.126		Present the initial SEAL workplan to SC2	10-01-03	10-01-03
1.1.3	1.127		SEAL V1 essentials in alpha	31-03-03	07-04-03
1.1.3	1.129		SEAL V0.3 internal release	16-05-03	24-05-03
1.1.3	1.128		Nightly builds deployed in SEAL	30-06-03	10-07-03
1.1.3	1.13		SEAL V1 release	30-06-03	18-07-03
1.1.3	1.187		SEAL support for Windows binaries	15-09-03	09-10-03
1.1.3	1.203	i_apr04	External software guideline document	31-05-04	
1.1.3	1.204	i_apr04	Workbook for SEAL	15-06-04	30-06-04
1.1.3	1.205	i_apr04	New Dictionary API and reference implementation	30-06-04	
1.1.3.8			**Mathematical libraries**		
1.1.3.8	1.167		Delivery of first round of GSL enhancements	30-05-03	30-05-03
1.1.3.8	1.184		Math library workplan in place	31-07-03	27-03-04
1.1.3.8	1.124		Statement on GSL and NAG usage for math library	30-09-03	24-10-03
1.1.3.8	1.206		mathlib project web	15-07-04	
1.1.3.8	1.207		First version of the C++ mathlib package	01-10-04	
1.1.4			**Physicist Interface (PI)**		
1.1.4	1.149		Establish physicist interface project	19-11-02	25-11-02
1.1.4	1.15		Report on PI experiment consultations	20-01-03	20-01-03
1.1.4	1.151		Initial PI plan presented to SC2	29-01-03	07-02-03
1.1.4	1.153		Physicist interface (PI) workplan completed	01-07-03	04-07-03
1.1.4.1			**Analysis services**		
1.1.4.1	1.152		AIDA interface review (developers) completed	30-05-03	30-05-03
1.1.4.1	1.171		AIDA interface review (users) completed	15-09-03	09-10-03
1.1.5			**Simulation**		
1.1.5	1.131		SC2 mandate for simulation project	07-12-02	07-12-02
1.1.5	1.132		Simulation project structure defined	28-01-03	28-01-03
1.1.5	1.133		Simulation project draft workplan	28-02-03	28-02-03
1.1.5	1.134		Simulation project initial workplan approved	14-03-03	19-03-03
1.1.5.2			**Geant4 development and integration**		
1.1.5.2	1.189	i_apr04	SPI-G4 collaborative infrastructure pilot	16-02-04	16-02-04
1.1.5.2	1.209	i_apr04	Geant4 6.2 release - resource usage refinements	25-06-04	25-06-04
1.1.5.2	1.21		Consolidated Geant4 acceptance suite for LHC	15-10-04	
1.1.5.2	1.211	i_apr04	Geant4 7.0 release - physics models and geometry	17-12-04	
1.1.5.4			**Physics validation**		

WBS	ID_Unique	Comment	Name	Date_MS	MS_Done
1.1.5.4	1.136		Test beam setup of the four experiments available in G4	31-05-03	31-05-03
1.1.5.4	1.143		First cycle of EM physics validation complete	30-09-03	30-09-03
1.1.5.4	1.14		Simulation physics requirements revisited	31-12-03	25-02-04
1.1.5.4	1.145		First cycle of hadronic physics validation complete	01-02-04	30-06-04
1.1.5.4	1.213	i_apr04	Review/prioritization of simple benchmarks for simu physics validation	31-05-04	
1.1.5.4	1.214	i_apr04	Comparison of LHC calorimeters for EM shower development	15-09-04	
1.1.5.4	1.215	i_apr04	Second iteration of hadronic physics validation complete	15-10-04	
1.1.5.4	1.146		Simulation test and benchmark suite available	31-12-04	
1.1.5.4	1.148		Final physics validation document complete	31-12-04	
1.1.5.6			**Generator services**		
1.1.5.6	1.138		Generator librarian and alpha version of support infrastructure in place	30-06-03	30-06-03
1.1.5.6	1.197		Agreement on formats for event generator common samples	01-03-04	01-03-04
1.1.5.6	1.198		Beta version of MCDB in production in the LCG environment	01-07-04	28-06-04
1.1.5.6	1.202	i_apr04	Agreement on parton-level event generator file format	01-09-04	
1.1.5.6	1.208	i_apr04	Generator production framework beta	01-12-04	
			Experiment verification milestones - AA		
1.1.2.1	1.18	V_CMS	CMS POOL integration: POOL persistency of CMS event	31-07-03	31-07-03
1.1.2.1	1.18	V_ATLAS	ATLAS POOL integration: POOL persistency in Release 7	10-09-03	11-09-03
1.1.2.1	1.18	V_CMS	CMS POOL validation with PCP data	31-10-03	21-10-03
1.1.2.1	10198	V_LHCb	LHCb POOL integration: Gaudi persistency replaced by POOL	19-12-03	04-02-04
1.1.2.1	10199	V_ATLAS	ATLAS validation of POOL Metadata/event collections	01-02-04	
1.1.2.1	10197	V_ATLAS	ATLAS POOL validation with complete Event Data Model	01-04-04	
1.1.3	10200	V_ATLAS	ATLAS integration: SEAL Integration into Athena	31-12-03	
1.1.4	10202	V_ATLAS	ATLAS int: ROOT implementation of AIDA histograms in Athena	30-11-03	30-11-03

Fabrics Area Milestones

WBS	ID_Unique	Comment	Name	Date_MS	MS_Done
1.2		___	**Fabrics**		
1.2.1			**Repositories and Managed Storage**		
1.2.1.1			**Castor**		
1.2.1.1.1	2.521		Concept of pluggable scheduler and high rate request handling demonstrated	01-10-03	22-10-03
1.2.1.1.2	2.52		Integrated prototype of the whole system	02-02-04	24-06-04
1.2.1.1.3	2.519		Production system ready for deployment	01-04-04	
1.2.1.1.4	2.518		Full functional CASTOR release available	01-03-05	
1.2.1.1.5	2.1167	V_ATLAS	Full CASTOR support for ATLAS DC3 (2nd half 2005)	30-06-05	
1.2.1.2			**Data Recording Challenges**		
1.2.1.2.3	2.515	V_ALICE	ALICE DC6 (450 MB/sec)	01-10-04	
1.2.1.2.4	2.1172	V_ALICE	ALICE Verification of target performance for ADC VI	15-11-04	
1.2.1.2.5	2.514	V_ALICE	ALICE DC7 (750 MB/sec)	03-10-05	
1.2.1.2.6	2.1173	V_ALICE	ALICE Verification of target performance for ADC VII	15-11-05	
1.2.1.3			**Evaluation of new tape hardware**		
1.2.1.3.1	2.526		Report on LTO evaluation delivered	17-01-05	
1.2.1.3.2	2.525		Report on high-end tape drive evaluation delivered	01-12-04	
1.2.1.3.3	2.524		Proposal on storage architecture and implementation of disk storage system for Phase 2 made	01-06-04	
1.2.2			**Fabric Management**		
1.2.2.1	2.538		SPMA – full production	01-08-03	18-08-03
1.2.2.2	2.537		NCM – full production	03-11-03	03-11-03
1.2.2.3	2.536		SMS – initial system	01-10-03	01-10-03
1.2.2.4	2.535		SMS – full production	01-07-04	01-07-04
1.2.2.5	2.534		Monitoring – Systems monitor displays operational	01-10-03	01-06-04
1.2.2.6	2.533		Monitoring – No further need for private (experiment) monitoring	01-03-04	30-03-04
1.2.2.7	2.532		FT – development started	01-09-03	01-09-03
1.2.2.8	2.531		FT – Initial system available	01-03-04	01-03-04
1.2.2.9	2.53		FT – Deployment complete	01-09-04	
1.2.2.10	2.529		HSM ops display – development started	01-10-03	01-10-03
1.2.2.11	2.528		HSM ops display – system complete	01-12-04	
1.2.3			**Computer Centre Upgrade**		
1.2.3.1	2.553		Vault converted to machine room	01-11-02	01-11-02
1.2.3.2	2.552		Right half of m/c room migrated to vault	01-08-03	15-08-03
1.2.3.3	2.551		Sub-station civil engineering starts	01-09-03	18-08-03
1.2.3.4	2.55		Sub-station civil engineering finishes	01-03-04	01-03-04
1.2.3.5	2.549		Sub-station electrical installation starts	01-04-04	01-04-04
1.2.3.6	2.548		Sub-station commissioned	07-01-05	
1.2.3.7	2.547		Electrical distribution on right side of m/c room upgraded	02-02-04	02-02-04
1.2.3.8	2.546		Left half of m/c room emptied	02-08-04	
1.2.3.9	2.545		Electrical distribution on left side of m/c room upgraded	01-06-05	
1.2.3.10	2.544		Machine room HVAC upgraded	01-03-05	
1.2.3.11	2.543		New 800KW UPS for physics installed	03-04-06	
1.2.3.12	2.542		Current UPS area re-modelled	01-01-07	
1.2.3.13	2.541		2nd 800 KW UPS added	02-04-07	
1.2.3.14	2.54		3rd 800 KW UPS added	01-04-08	
1.2.4			**Batch Computing**		

WBS	ID_Unique	Comment	Name	Date_MS	MS_Done
1.2.4.1	2.564		Baseline plan for Phase 2 agreed (architecture + technology)	01-06-04	
1.2.4.2			**Grid Integration**		
1.2.4.2.1	2.562		10 LXBATCH nodes integrated in LCG-1 (isolated from main LXBATCH)	01-08-03	20-01-04
1.2.4.2.2	2.561		100 LXBATCH nodes integrated in LCG-1 (isolated from main LXBATCH)	01-10-03	20-01-04
1.2.4.2.3	2.56		Full integration of LXBATCH worker nodes in LCG-1 completed	01-12-03	15-03-04
1.2.4.3			**Batch Scheduling**		
1.2.4.3.1	2.557		Decision on Phase 2 Batch Scheduling System made	01-12-04	
1.2.5			**Installation Schedule (short term)**		
1.2.5.9			**Networking infrastructure**		
1.2.5.9.3	2.58		Integration of 10 Gbit equipment into the Prototype done	01-07-03	01-07-03
1.2.5.9.5	2.582		First phase of backbone replacement by 10 Gbit equipment	01-07-04	
1.2.5.9.6	2.583		Backbone fully upgraded to 10 Gbit	01-07-05	
1.2.6			**Phase 2 Acquisition Process**		
1.2.6.1	2.591		Decisions on the equipment selection criteria made	01-12-03	01-12-03
1.2.6.2	2.59		Agreement with SPL on acquisition strategy ready	01-12-03	01-12-03
1.2.6.3	2.589		Market survey issued for cpu and disk	01-07-04	
1.2.6.4	2.588		Market survey issued for tape system	01-12-04	
1.2.6.5	2.587		Phase 2 TDR ready	01-07-05	
1.2.6.6	2.586		Finance Committee for cpu & disk vendor selection passed	01-09-05	
1.2.6.7	2.585		Finance Committee for tape system selection passed	01-12-05	
1.2.6.8	2.584		All orders placed for 2006 equipment	01-03-06	
1.2.6.9	2.575		Order placed for tape system	09-01-06	
1.2.6.10	2.574		2006 system installation and commissioning complete	01-08-06	
1.2.6.11	2.573		Phase 2 service in production	01-09-06	
1.2.7			**Milestones prior to June 2003**		
1.2.7.1	2.599		L2M Production Pilot I starts	15-01-03	06-01-03

Grid Technology Area Milestones

WBS	ID_Unique	Comment	Name	Date_MS	MS_Done
1.3		___	**Grid Technology Area**		
1.3.1			**Technology Review**		
1.3.1.3	3.496		Report on the Review of Grid Technologies 2003 ready	06-10-03	06-10-03
1.3.1.5	3.498		Report on the Review of Grid Technologies 2004 ready	02-11-04	
1.3.2			**System Modeling**		
1.3.2.2	3.337		Modeling plan for work during 2003 created	15-07-03	15-07-03
1.3.3			**LCG Technology Evolution**		
1.3.3.2	3.344		Report on GT3 capabilities ready	01-09-03	01-09-03
1.3.5			**TDR Planning**		
1.3.5.2	3.347		Plan and outline for the TDR available	01-10-04	
1.3.5.3	3.348		Proposal for review created	24-02-05	
1.3.6.4	3.1188	V_LHCb	Computing model chapter of LHCb computing TDR available	01-12-04	
1.3.5.4	3.349		LCG TDR completed	01-06-05	
1.3.7			**Milestones prior to June 03**		
1.3.7.3			**Deliver initial software packages**		
1.3.7.3.2	3.494		Initial software packages delivered	30-05-03	30-05-03
1.3.7.3.3	3.1174	V_ALICE	ALICE AliEn analysis interface to LCG-2 performance evaluation	15-10-04	

Grid Deployment Area Milestones

WBS	ID_Unique	Comment	Name	Date_MS	MS_Done
1.4		___	**Grid Deployment**		
1.4.1			**First Global Service Initial Availability**		
1.4.1.1	4.353		Initial middleware for LCG-1 delivered	30-04-03	31-07-03
1.4.1.2	4.354		LCG-1 security model implemented	30-06-03	30-06-03
1.4.1.3	4.355		First prototype operations & user support service in place	30-06-03	30-06-03
1.4.1.4	4.356		LCG-1 m/w to agreed 10 Tier 1 sites deployed	15-07-03	15-09-03
1.4.1.6	4.956		LCG-1 certified and commissioned	01-08-03	29-09-03
1.4.1.7	4.1175	V_ALICE	ALICE AliEn production interface to LCG-1 readiness review (prototype design and functionality)	14-11-03	05-01-04
1.4.1.8	4.1176	V_ALICE	ALICE AliEn analysis interface to LCG-2 readiness review (prototype design and functionality)	16-08-04	
1.4.1.9	4.1177	V_ALICE	ALICE AliEn production interface to LCG-2 performance evaluation	17-08-04	
1.4.1.10	4.1178	V_ALICE	ALICE LCG-2 PDC3 production environment verification	15-07-04	
1.4.1.11	4.1179	V_ALICE	ALICE LCG-2 PDC3 analysis environment verification	15-10-04	
1.4.1.12	4.358	V_ATLAS	ATLAS validation of LCG-1 (now LCG-2)	15-10-04	
1.4.1.13	4.1169	V_ATLAS	ATLAS production on LCG-1 (now LCG-2)	29-02-04	24-06-04
1.4.1.14	4.1186	V_LHCb	DIRAC deployed in LCG1 (now LCG-2)	28-11-03	24-05-04
1.4.1.15	4.1185	V_LHCb	50% of LHCb DC04 production executing in LCG-2	01-06-04	
1.4.1.16	4.1187	V_LHCb	LHCb analysis environment (GANGA) deployed in LCG-2	01-07-04	
1.4.2			**Full Operational LCG-1 Service**		
1.4.2.2	4.957		LCG-1 performance goals defined	01-09-03	09-09-03
1.4.2.3			**Expand Service to include Tier2 centres**		

WBS	ID_Unique	Comment	Name	Date_MS	MS_Done
1.4.2.3.1	4.958		Tier2 centres included in service	18-10-03	18-10-03
1.4.2.4			**Middleware functionality complete**		
1.4.2.4.6	4.965		Middleware functionality complete	01-10-03	15-12-03
1.4.2.5	4.369		Job Execution model defined	30-09-03	15-09-03
1.4.2.8	4.963		Deployed to agreed sites	15-10-03	15-01-04
1.4.2.9			**Security Procedures for 2004 in place**		
1.4.2.9.3	4.969		Security Procedures for 2004 in place	30-11-03	10-11-03
1.4.2.11	4.966		Service Review completed (LHCC and internal review)	24-11-03	24-11-03
1.4.3			**Fully Operational PDC4 environment**		
1.4.3.7	4.118	V_ALICE	ALICE LCG-2 PDC4 production environment verification	15-07-06	
1.4.3.8	4.1181	V_ALICE	ALICE LCG-2 PDC4 analysis environment verification	15-10-06	
1.4.4			**Completion of Computing Service TDR**		
1.4.4.2	4.382		Experience with LCG-2 reviewed	01-07-05	
1.4.5			**LCG-2 in 2004**		
1.4.5.1			**Service**		
1.4.5.1.1	4.164	i_jan04	LCG-1 Fully operational: LCG-2 operational at Core sites	01-02-04	15-02-04
1.4.5.1.2	4.1639	i_jan04	30 day verification run successfully terminated	15-03-04	30-03-04
1.4.5.1.3	4.1638	i_jan04	LCG-2 operational at 30 sites	01-05-04	23-04-04
1.4.5.1.4	4.1637	i_jan04	50% prototype available	20-12-04	
1.4.5.2			**Operations**		
1.4.5.2.1	4.1635	i_jan04	Basic accounting service in place	01-03-04	30-03-04
1.4.5.2.2	4.1634	i_jan04	Second GOC in operation	01-06-04	30-03-04
1.4.5.3			**Support**		
1.4.5.3.1	4.1632	i_jan04	FZK Support portal as front line user support operational	01-06-04	01-06-04
1.4.5.4			**Experiment verification of LCG-2**		
1.4.5.4.1	4.163	i_jan04	Report on ALICE DC initial experience	01-04-04	01-04-04
1.4.5.4.2	4.1629	i_jan04	Report on CMS DC04 initial experience	01-05-04	15-06-04
1.4.5.4.3	4.1628	i_jan04	Report on ATLAS initial experience	01-05-04	
1.4.5.4.4	4.1627	i_jan04	Report on LHCb initial experience	01-05-04	
1.4.5.5			**Middleware**		
1.4.5.5.1	4.1644	i_jan04	WN middleware ported to RHEL, report on ease of porting to other OSs	01-05-04	
1.4.5.5.2	4.1643	i_jan04	Configuration tool available	01-07-04	
1.4.5.5.3	4.1642	i_jan04	Replica Manager upgrade	01-08-04	

Middleware Milestones

WBS	ID_Unique	Comment	Name	Date_MS	MS_Done
1.5		___	**Middleware**		
1.3.4			**LCG - EGEE Coordination**		
1.3.4.1	3.1657		EGEE senior management appointed	15-07-03	15-07-03
1.3.4.2	3.1658		Technical design team established	01-09-03	04-12-03
1.3.4.8	3.1659	i_jan04	EGEE Middleware people hired	29-02-04	23-03-04
1.3.4.9	3.166	i_jan04	EGEE Middleware execution plan available	29-02-04	17-03-04
1.3.4.10	3.1661	i_jan04	EGEE Contract signed	01-04-04	01-04-04
1.5.2.6	5.1912	i_apr04	First version of prototype available for experiments	16-06-04	18-05-04
1.5.2.7	5.1889	i_apr04	Development and integration tools deployed	30-06-04	30-06-04
1.5.2.8	5.1888	i_apr04	Software cluster development & testing infrastructure in place	30-06-04	30-06-04
1.5.2.9	5.1887	i_apr04	Architecture & Planning Document for release candidate 1	30-06-04	30-06-04
1.5.2.10	5.1886	i_apr04	Integration & Testing infrastructure in place; test plan	31-08-04	
1.5.2.11	5.1885	i_apr04	Grid Services design document for release candidate 1	31-08-04	
1.5.2.12	5.1913	i_apr04	Second version of prototype available for experiments	31-08-04	
1.5.2.13	5.1884	i_apr04	Release Candidate 1	31-12-04	

ARDA Milestones

WBS	ID_Unique	Comment	Name	Date_MS	MS_Done
1.6		___	**ARDA**		
1.3.7.2	3.1663	i_jan04	ARDA project organized	31-01-04	19-02-04
			E2E prototype definition		
1.6.2	6.1899	i_apr04	E2E ALICE prototype definition agreed with the experiment	31-05-04	31-05-04
1.6.3	6.1898	i_apr04	E2E ATLAS prototype definition agreed with the experiment	31-05-04	31-05-04
1.6.4	6.1897	i_apr04	E2E CMS prototype definition agreed with the experiment	31-05-04	
1.6.5	6.1896	i_apr04	E2E LHCb prototype definition agreed with the experiment	31-05-04	31-05-04
			Prototypes using basic middleware		
1.6.6	6.1895	i_apr04	E2E ALICE prototype using basic EGEE middleware	30-09-04	
1.6.7	6.1894	i_apr04	E2E ATLAS prototype using basic EGEE middleware	30-09-04	
1.6.8	6.1893	i_apr04	E2E CMS prototype using basic EGEE middleware	30-09-04	
1.6.9	6.1892	i_apr04	E2E LHCb prototype using basic EGEE middleware	30-09-04	
			Prototypes capable of analysis		
1.6.14	6.1904	i_apr04	E2E prototype for ALICE, capable of analysis	31-12-04	
1.6.15	6.1903	i_apr04	E2E prototype for ATLAS, capable of analysis	31-12-04	
1.6.16	6.1902	i_apr04	E2E prototype for CMS, capable of analysis	31-12-04	
1.6.17	6.1901	i_apr04	E2E prototype for LHCb, capable of analysis	31-12-04	
1.6.18	6.19	i_apr04	E2E prototype for each experiment (4 prototypes), capable of analysis (or advanced production)	31-12-04	
1.6.19	6.189	i_apr04	E2E prototype for each experiments (4 prototypes), capable of analysis and production	31-12-05	

Global Level-1 Milestones

The level 1 milestones have been revised and most of the previous milestones have been replaced.

Note:
The workplan for the applications area being established now, we plan to provide additional milestones in this area for the next quarterly report.

#	Milestone	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 reliable 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%.* **The milestone was met with 2.5 months delay on September 15, 2003.**	**July 2003**
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.* **The milestone was met on June 30, 2003**.	**October 2004**

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. **The milestone was met on March 1, 2004**.	**November 2003**
M1.5	**Baseline plan for Phase 2 agreed (Tier-0 & Tier-1s)** *A document will describe the required resources and level of service to be provided by the Tier-1 and Tier-0 centres for Phase 2, and include estimates for continuing support beyond Phase 2. This document will provide the initial resource and service requirements for the LCG MoU, and provide a model of the base services to be expanded in the LCG TDR.*	**October 2004**
M1.6	**AA development & support plan & resources through 2008** *With the LCG-funding of manpower in the applications area running down during 2004 and 2005, a plan is required to establish the level of long-term support that is required for the products that are essential for the experiments. This plan will lead to the definition of more detailed milestones.*	**December 2004**
M1.7	**Baseline distributed batch environment using grid service - four experiments using LCG-2 level of functionality (general batch analysis service going beyond MC production)** *Reliable batch service incorporating Tier 0, all Tier 1s and several large Tier 2s (about 20 sites, 2000 systems) in regular use by all four experiments. Basic scheduling and data distribution functionality, including synchronisation of metadata catalogues. The system must be demonstrated with a scale and performance of 25% of that required at LHC startup. This validates a fallback solution in the event that there are delays with the development of new middleware.*	**December 2004**

M1.8	**Distributed end-user batch and interactive analysis prototype (ARDA) available and successfully tested in four experiments** *Four instances of end-to-end distributed analysis solutions implementing the ARDA middleware into the experiments' frameworks are operational in the form of prototypes.*	**December 2004**
M1.9	**ALICE target performance achieved and verified – DC-6 (450 MB/s)** *The full chain of central data recording to tape achieves a sustained operation at 450 MB/s for a period of at least 10 days.*	**December 2004**
M1.10	**Coherent operation of the LHC computing Tier-1 centres** *The Tier-1 centres for LHC computing have operated as a reliable grid service for three months, with a full set of management, monitoring, accounting and data management tools. This is a milestone of the Tier 1 centres and ancillary services (operations centres, call centres, virtual organization management, etc.).*	**July 2005**
M1.11	**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. The full process from acquisition to service verification is expected to take 12-18 months (depending on 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.*	**July 2005**
M1.12	**Phase 1 applications area software – full required functionality – available and successfully used in ATLAS, CMS & LHCb** *Products include POOL, SEAL, Geant4, ROOT, conditions database, event collections, analysis services, generator services, etc. A detailed description of the milestone as a union of specific Level 2 technical and adoption milestones will be defined with Milestone #2 in December 2004.*	**September 2005**

M1.13	**ALICE target performance achieved and verified – DC-7 (750 MB/s)** *The full chain of central data recording to tape achieves a sustained operation at 750 MB/s for a period of at least 10 days.*	**December 2005**
M1.14	**Initial production LCG middleware providing the functionality for the first LHC data taking in operation (Tier-0, Tier-1 Tier-2)** *The middleware required for the first data taking at LHC must be available at Tier-0, Tier-1 and large Tier-2 centres one year before the start of the run.*	**April 2006**
M1.15	**Initial LCG system functionality for first LHC run complete (Tier-0 & Tier-1s)** *Tier-1 and Tier-0 centres are fully operational and prepared to increase the capacity as required for data analysis.*	**October 2006**