



<http://www.cern.ch/lcg>

LHC Computing Grid Project (LCG) Phase 1

Status of High-Level Planning

version 1 - 21 June 02

Prepared by the LCG Project Execution Board

This paper gives a summary of the current state of the high-level planning of Phase 1 of the LHC Computing Grid (LCG) project. It has been prepared as information for the LHCC meeting of July 2002. The paper was endorsed at the meeting of the LCG Software and Computing Committee on 19 June 2002.

The project was authorised by the CERN Council on 20 September 2001, and the key committees had first meetings before the end of the year. A workshop was held at CERN on 11-15 March 2002 at which the main directions and priorities of the project were discussed. The plan presented here reflects the conclusions of the workshop, taking account of further input and discussions that have taken place over the past three months, in particular the formal statements of requirements in several key areas that have been developed by the project's Software and Computing Committee (SC2). Most of the senior positions in the project management structure have been filled now, and the staffing for the part of the project that is CERN's responsibility is being built up, funded by CERN and special contributions from several member and observer states. The framework for planning, coordinating and managing the grid services at CERN and in the Regional Centres will be set up during the coming months. Discussions with the national and regional projects that are developing the essential "middleware" that will be used to operate the grid started at the beginning of this year and are still under way. The planning in this document is therefore very much "work in progress", and many of the key elements will evolve over the coming months.

Table of contents

1.	Scope and goals of the project		3
2.	Outline description of the LHC Computing Grid		5
3.	Background and General Requirements		9
4.	Area Descriptions		11
4.1.	Applications	11	
4.2.	Fabrics	14	
4.3.	Grid Technology	18	
4.4.	Grid Deployment	20	
5.	Level 1 Milestones		24
6.	Consolidated High Level Plan		26
7.	Risks & Problems		27
	Annexes		31
A	Data Challenge Plans		32
A.1	ALICE Data Challenge Planning	32	
A.2	ATLAS Data Challenge Plans	37	
A.3	CMS Data Challenges	41	
A.4	LHCb Data Challenges	43	
B	Resources at CERN		45
C	Resources in Experiments		46
C.1	Resources in ALICE	46	
C.2	Resources in ATLAS	47	
C.3	Resources in CMS	47	
C.4	Resources in LHCb	49	
D	Resources in Regional Centres		50
E	Management and Reporting Structure		51
F	Extracts of WBSs		54
F.1	Applications	54	
F.2	Fabrics	58	
F.3	Grid deployment – High level WBS	59	
G	Glossary		61

1. Scope and goals of the project

The job of the LHC Computing Grid Project – LCG – is to prepare the computing infrastructure for the simulation, processing and analysis of LHC data for all four of the LHC collaborations. This includes both the common infrastructure of libraries, tools and frameworks required to support the physics application software, and the development and deployment of the computing services needed to store and process the data, providing batch and interactive facilities for the worldwide community of physicists involved in LHC.

The first phase of the project, from 2002 through 2005, is concerned with the development of the application support environment and of common application elements, the development and prototyping of the computing services and the operation of a series of computing data challenges of increasing size and complexity to demonstrate the effectiveness of the software and computing models selected by the experiments. During this period there will be two series of important but different types of data challenge under way: *computing data challenges* that test out the application, system software, hardware and computing model, and *physics data challenges* aimed at generating data and analysing it to study the behaviour of the different elements of the detector and triggers. During this R&D phase the priority of the project is to support the *computing* data challenges, to identify and resolve problems that may be encountered when the first data arrives. The physics data challenges require a stable computing environment that may conflict with the needs of the computing tests, but it is an important goal of the project to arrive rapidly at the point where stability of the grid prototype service is sufficiently good to absorb also the resources that are available in Regional Centres and CERN for physics data challenges.

This first phase will conclude with the production of a Computing System Technical Design Report, providing a blueprint for the computing services that will be required when the LHC accelerator begins production. This will include capacity and performance requirements, technical guidelines, costing models, and a construction schedule taking account of the anticipated luminosity and efficiency profile of the accelerator.

A second phase of the project is envisaged, from 2006 through 2008, to oversee the construction and operation of the initial LHC computing system.

The requirements for LHC data handling are very large, in terms of computational power, data storage capacity, data access performance and the associated human resources for operation and support. It is not considered feasible to fund all of the resources at one site, and so it has been agreed that the LCG computing service will be implemented as a geographically distributed *Computational Data Grid*. This means that the service will use computing resources, both computational and storage, installed at a large number of *Regional Computing Centres* in many different countries, interconnected by fast networks. Special software, referred to generically as *grid middleware*, will hide much of the complexity of this environment from the user, giving the illusion that all of these resources are available in a coherent *virtual computer centre*. This is an emerging technology that is at present receiving substantial R&D support from agencies that fund computing developments, and is exciting considerable interest from industry.

The management structure of the project is described in some detail in *Annexe E: Management and Reporting Structure*. Two committees manage the organisation of the project: the Software and Computing Committee (SC2) and the Project Execution Board (PEB). The SC2, including representation of the experiments and some of the Regional Centres, sets the overall requirements, which are turned into a work plan by the PEB. The SC2 approves the strategy & work plan developed by the PEB, monitors progress and adherence to the requirements, and takes part in the project reviews organised by the LHCC. The PEB manages the day-to-day work of the project, brokering agreements with the various resource holders, negotiating technical solutions, preparing detailed work plans with the implementation teams, and monitoring progress of the project through agreed milestones and deliverables.

The human and material resources required for Phase I of the project come from a variety of sources, including:

1. resources at CERN funded both by the CERN base budget, and by special voluntary contributions from countries participating in the LHC programme (these are referred to in this paper as *CERN resources*);
2. industrial contributions, including resources provided by members of the *CERN [Openlab](#) for Datagrid Applications*¹;
3. resources managed by the LHC experiments, at CERN and elsewhere – this is particularly important in the area of applications software development;
4. resources provided by national funding agencies at LHC Regional Computing Centres;
5. technology research and development projects funded by the European Commission and other national and regional funding agencies².

The LCG Project has direct managerial control of the resources only in the case of the first source. In the other cases the Project has only a coordinating role supported by policies and strategies adopted in the various boards and committees of the project, unless a more direct management relationship is explicitly agreed with the resource owners.

With so many elements to be managed and coordinated, the execution of the project has been organised in four different areas:

- *applications*: This area is responsible for the management of the common applications developments and services agreed by the SC2, forming the project teams and collaborations required for implementation, and ensuring the technical and architectural coherence of the software. Staffing for this area comes largely from CERN resources (see point 1 above) and from the computing projects of the experiments (see point 3 above).
- *computing fabrics*: This is the term used for the very large computing clusters that will be installed in Regional Centres. This area is responsible for:

¹ <http://www.cern.ch/openlab>

² For a list of the major grid technology projects see <http://www.eu-datagrid.org>

- a. the provision of the tools required to manage the CERN computing fabric;
 - b. interfacing the computing fabric to the LCG grid;
 - c. providing coordination between the Regional Centres in fabric-related matters, such as: coordinating periodic computing technology studies, developing a cost model for the Phase 2 computing system, ensuring effective exchange of experience and technical plans.
- *grid technology*: This area is responsible for the provision and support of the grid middleware required by the project. It is envisaged that this technology will be supplied by projects mainly funded from other sources. Grid technology is developing very rapidly and there is therefore scope for divergence between the different projects and organisations involved. The Grid Technology Area is responsible for ensuring the coherence and interoperability of grid middleware used in different regions.
 - *grid deployment*: This area is responsible for the operation of the LCG Grid. This includes organising the integration and distribution of the grid middleware and applications support environment, providing the infrastructure and maintenance services needed to define and operate the Grid, obtaining agreement on the standards and policies to be used for authentication, authorisation, scheduling, accounting and reporting, and for coordinating the day-to-day operation of the Grid. The resources for this area will be provided largely by the Regional Centres taking part in the LCG Grid. During Phase 1 of the project LCG will be operating one of the very early data intensive production grids. Many things will be changing continuously and we shall have to be very flexible, following emerging standards and adopting new functionality while maintaining a strong focus on delivering a stable service.

Many of the activities that will contribute to the project have been going on in one way or another for some time. Many of the Regional Centres are already in operation providing services for LHC, other experiments and in some cases other sciences. A number of projects have started work over the past couple of years in developing grid middleware – the software that will be needed to knit the Regional Centres into a grid that will appear as a single facility to the applications user. The experiments have been working for many years on the frameworks and applications that are needed to analyse the LHC data. The project therefore begins its work in a very complicated environment, and it will be some time before we fully understand the major problems that will have to be solved, negotiate agreements on deliverables and services with existing projects and organisations, and construct a fully resource loaded project plan. The current target is to have a rather solid plan set up before the end of 2002.

2. Outline description of the LHC Computing Grid

Scope

It is important to consider that the primary objective of the LCG is to address the LHC computing problem via the appropriate technologies selected to create an instance of a grid. In doing so, the product is to be a system that is fully

maintained and supported, offering a comprehensive computing service to the LHC physics community. This is a big step from demonstrating that a collection of technologies can perform useful work, which is the primary product of many grid initiatives today. However, these grid initiatives are performing the proof of concept work and providing the basic technologies that will be candidates for use in the LHC computing grid, once they have reached maturity.

This section describes what the LHC computing grid will offer as a service to the users and how it will be implemented and managed as a technology. There are a number of unanswered questions, particularly concerning the operational management of such a service. These questions will be addressed more completely in Section 6 of this document, *Risks and Problems*.

The LHC Computing Grid and LHC

As the introduction to this document has outlined, the use of grid technologies matches well the thinking that has taken place to solve the computing capacity needs for LHC. The initial thoughts concluded that the complete capacity for LHC data processing could not be centralized at CERN but rather should be distributed throughout computing centres in the countries that will take part in the LHC programme.

MONARC

In 1998 the MONARC project was set up to model the distribution of jobs across a number of regional computing centres of different sizes and capacities, in particular larger Tier-1 centres, providing a full range of services, and smaller Tier-2 centres. The architecture developed by MONARC is described in the report³ of the project.

The grid is naturally seen as a technology framework with which to implement the concept of a hierarchy of Regional Centres, and this approach will be used by the LCG project to assemble these geographically disparate resources into a coherent physics service. Developing the technologies needed forms the basis of much of the work in progress in the many grid initiatives throughout the world.

One of the special characteristics of the LHC computing grid will be the quantity of data that has to be distributed throughout the Regional Centres and the way in which the data is accessed. As yet there is no clear model of which data will reside where but the raw data store will reside at least at the Tier-0 centre (CERN). The tasks of reconstructing and storing the event summary data (ESD), analysis object data (AOD) and even tag (TAG) data will be distributed among Regional Centres although the conditions under which this will occur are not yet clear. The same is true for the large amount of simulated data.

The User Perspective

From the users point of view the use of the LHC computing grid will be through one of a number of “job portals” that will hide some of the complexity of the underlying environment. Such portals will offer the possibility for the user to specify the characteristics of the job and the basic resource requirements including the datasets that are required. Suitable computer centre candidates for execution of

³ MONARC: Models of Networked Analysis at Regional Centres for LHC Experiments, Phase 2 Report CERN/LCB-2000-001 - <http://www.cern.ch/MONARC/docs/phase2report/Phase2Report.pdf>

the job will then be determined and the job passed to the appropriate batch queue. After job execution, the results will be passed back to the user.

This description of course skims over a number of major issues such as how the characteristics of the job are to be specified, how the matching to centres is performed and how the results are to be retrieved. Each of these issues is a topic that is being dealt with in the various grid initiatives.

It is very likely that the grid infrastructure will be accessed also for interactive tasks. Again here much of the details, and even of the user requirements, need to be worked out.

The Infrastructure

Underlying the user perspective described above is a grid infrastructure that includes three main architectural components, *middleware*, *network* and *fabric*.

The grid *middleware* implements the functional components that permit the interaction and cooperation between Regional Centres participating in the grid. Transparent file access across centres, authenticating and authorising jobs to use facilities and locating available resources are some of the key components that the middleware provides.

The network is the physical link between the different centres that makes their connection and inter-working possible. To realise the transparent, ubiquitous and reliable access to high-end resources that is the objective of the grid middleware, the network component will have to evolve substantially both in terms of performance and in terms of quality and quantity of services offered.

The *fabric* is the physical hardware resource that is used at each Regional Centre to host the data and execute the jobs. Typically the computing resources are some form of compute farm – a collection of computers with access to data stored on disk servers or mass storage devices. The computers and servers of the fabric are connected together through high-speed local area networks or specialized interconnects. The centres themselves are inter-connected through public networks or dedicated links.

It is not realistic within the LHC computing grid to standardize fully on the fabric components that must be present at each centre as this depends on many factors that cannot be coordinated or imposed. Such things as available hardware resources, preferred network components and architectures, staff and their skill levels as well as organizational policies and non-LHC user communities will all affect the choice of the fabric components. It will be important, however, that certain details of the architecture and the operating system, as well as certain policies and interface standards are agreed and implemented by all of the Regional Centres providing computing fabrics within the LCG grid.

Application Environment

When a job executes on a particular piece of hardware it is executing within a very specific environment. This includes the hardware environment (processor chip, memory configuration, etc.) as well as the software environment that is installed on the computer. Specific versions of each of the components of the software environment (e.g. operating system, mathematical libraries, data access package, application libraries) must be made available for a job to be able to execute correctly.

There is a clear desirability to have, as far as possible, a common software environment between the four LHC experiments, to avoid the inevitable costs associated with managing and maintaining complex and diverse software environments. The challenge will be to keep this commonality as the needs of the experiments and the technology evolve.

LHC Computing Service

The previous few sections have touched on the basic architecture and technology that form the core of the LHC computing grid. However, it will be apparent that there are a number of organizational requirements that must be addressed before the LHC computing grid will be able to provide a production resource.

The Regional Centres will provide well-managed facilities adapted to their specific environments. This includes high availability of resources, well-structured policies for maintenance and upgrades, disaster recovery policies (such as backup) and mechanisms for the determination and resolution of the complicated problems that may occur. Careful coordination between the different Regional Centres taking part in the LCG grid will be required, as foreseen in the project organisation.

Such centres manage a number of computational and data storage resources that will be made available, partially or entirely, as part of the overall LHC computing grid. Well-maintained and compatible versions of the appropriate software must be agreed and made available so that executing jobs will have the correct hardware and software environment available for them to be able to run. One issue here will be the process for upgrading the hardware and software environments, already difficult within a large computing centre, while maintaining compatibility with the other centres taking part in the grid. Adding to the complication, many Regional Centres will also be providing services to other user communities and taking part in other grids. Good configuration management at all levels will be essential, but the practical difficulties in this area should not be under-estimated. Many compromises will have to be negotiated between the managers of the Regional Centres.

The middleware software previously described provides the interoperability and the accessibility to resources at the different Regional Centres, implementing interface and protocol standards to ensure that the application has a global accessibility to and a single view of the resources, which may be realised in different forms at different centres. In a similar fashion to the job environment at the different sites, the middleware software must be tightly managed in the production environment. The coordination and support structures needed will have to be well developed and will be a major focus for the deployment teams. Given the complex and distributed nature of the overall environment the key to manageability will be the predictability of the effect of changes of state within the grid. This means that the failure of a piece of hardware or software should have a predictable effect on the overall system such that the effect of a failure can be managed. Where the system is undergoing many changes simultaneously, with unpredictable side effects, the ability to provide a managed service is significantly reduced. One of the major efforts around the creation of the computing grid will be to create a cost-effective management infrastructure that will constitute the core of the service management process.

Future Evolution

This overview would not be complete without an indication of where the various activities are leading in the future. While it is the plan to produce the LHC computing facility through a number of well-defined deliverables, the time frame involved is necessarily long, some four years from the time of writing until the final production environment should be ready. In this area technology and costs evolve very rapidly, and while it is important to progress quickly in order to demonstrate that the basic architecture that has been selected can be implemented and satisfies the requirements, it is advantageous to delay some decisions as long as possible in order to benefit from new generations of technology or lower units costs.

The rate of evolution of the software technology used to implement a grid is particularly rapid, reflecting its immaturity. Currently the National Science Foundation and Department of Energy in the US, and the European Union and various national research programmes in Europe are making substantial investments in this area. However, there are clear indications that commercial interests from IBM and Microsoft as well as Sun Microsystems have begun, which may lead to industrial products within the timescale of the project. It is, however, unlikely that there will be a single obvious grid solution that will be acceptable to all Regional Centres within the timeframe of the next few years. There will probably be a number of options as to what the core software should be and who should provide it. The key determinant will be the availability of *production level products* with solid maintenance and support structures. The selection of the key technologies that are needed to satisfy the computing requirements of LHC and the creation of a production system will remain the responsibility of the LHC computing grid project.

As soon as good quality solutions appear that enable reliable grids to be constructed we expect to see a number of computational grids being set up for production use within the scientific community. High energy physics in many countries will certainly have access to resources in these grids, and several of the Regional Centres providing computing resources for LHC experiments will be part of more than one grid serving different communities or funded from different sources. In the shorter term this will probably look like a series of inter-connected grids, but in the longer term we would expect physical resources to be shared between grids in a rather flexible way. It is not clear today how quickly this will happen. The definition and adoption of standard protocols and service interfaces will be required, which can take place only after there is substantial practical experience of using grids, and a certain level of maturity of the technology has been achieved.

3. Background and General Requirements

The preparations for LHC data handling extend back over several years, with various development activities, requirements analyses and evaluations undertaken in the context of several projects and committees.

- During the past two years, in both America and Europe, R&D projects were proposed to extend the computational Grid technology emerging from the earlier work of the [Globus](#) project, to respond to the scale of

LHC requirements for computation and data handling. In the United States three projects have been funded, led by high energy physics: the [Particle Physics Data Grid](#) (PPDG), [Grid Physics Network](#) (GriPhyN) and the [International Virtual Data Grid Laboratory](#) (iVDGL). In Europe several significant national projects have been initiated, and the European Union has funded the [DataGrid](#), [DataTAG](#) and [CrossGrid](#) projects, all led by high energy physics. We are fortunate that these important investments are already bearing fruit, and the experience and developments of these projects are important components for building the LHC Computing Grid.

- The [Hoffmann Review](#) – The *LHC Computing Review*, chaired by S.Bethke, carried out a comprehensive review of LHC computing, covering the requirements, technology, potential computing models and costs. The Review started work in December 1999 and the final report⁴, published on 22 February 2001, contains requirements and cost estimates for constructing the global computing environment for LHC data analysis and recommendations for the development and support of the physics applications software. The conclusions and recommendations, agreed by the four experiments and presented to the LHCC, form the basic requirements input to the LCG project.
- Following the publication of the report of the Hoffmann Review, the Information Technology Division at CERN, in consultation with the experiments, prepared a technical proposal⁵ for the implementation of the CERN contribution. This was presented to the CERN management and approved as the basis for preparing a proposal for a wider project to be presented to the CERN Council.

Building on this considerable background, CERN developed a proposal⁶ for the LCG project that was presented to the Scientific Policy Committee, the Finance Committee and Council on several occasions during 2001. Phase I of the project was agreed during the CERN Council meeting of 20 September 2001.

The basic parameters used to plan the capacity and performance of the LHC computing services are given in the report of the Hoffmann review. However, several significant factors have changed since the Review, while the experiments' computing models continue to evolve. During the coming year it will be necessary to review these parameters and re-assess the computing service model. There are two questions to be addressed:

- a) The priorities for the physics that will be done at the start-up for a given luminosity, how this will evolve over the following years, and how it translates into trigger rates, event sizes, and other related parameters.
- b) The data processing and analysis models that will be used by the different experiments to achieve their physics goals.

⁴ [CERN/LHCC/2001-004](#) *Report of the Steering Group of the LHC Computing Review*

⁵ [CERN-IT-DLO-2001-003](#) *Solving the LHC Computing Challenge: A Leading Application of High Throughput Computing Fabrics combined with Computational Grids*

⁶ [CERN/2379/rev.](#) *Proposal for building the LHC Computing Environment at CERN*

Based on the answers to these questions, one has to construct the corresponding architecture for the computing service and estimate the cost, both for CERN and the Regional Centres.

Question (a) should be addressed by experiments and reviewed by the LHCC to give the computing experts a framework of basic numbers to start from. Question (b) should be addressed by the SC2 committee of the LCG project. A team of computing and applications experts should then elaborate practical computing service architectures, taking account of the likely evolution of the various technologies and the financial implications of different approaches. Several iterations may be needed to converge on the best model.

There are still large uncertainties today, both in the understanding of the experiments' needs and in estimating the evolution of the cost of computing over the next five years. Therefore, a process needs to be established to repeat this exercise at intervals as the understanding evolves.

4. Area Descriptions

4.1. Applications

The Applications Area is concerned with developing, deploying and maintaining the physics applications software and associated supporting infrastructure software that is shared among the LHC experiments. The scope includes common applications software infrastructure, frameworks, libraries, and tools; common applications such as simulation and analysis toolkits; and assisting and supporting the integration and adaptation of physics applications software in the grid environment. User support and documentation is also part of this area. For those projects defined by the SC2 which fall within this scope, the Applications Area is responsible for building a project team among participants and collaborators; developing a work plan; identifying, adopting, designing and developing software that meets experiment requirements, adheres to a coherent overall architecture, and performs within the distributed LHC computing grid; assisting in integrating the software within the experiments; and providing support, documentation and maintenance. These products will form the set of centrally supported and maintained software for LHC.

Outline of the Work Plan

Establish a project organization that effectively brings together CERN software development staff, software specialists in the experiments, the ROOT team, and participants outside the LHC community in a coordinated common effort. *First half 2002.*

By end of third quarter 2002:

Establish a software process and infrastructure supporting the development, documentation and support of high quality software within a highly distributed community of developers and users.

Establish an overall architectural approach for LHC common software, concretely expressed and realized in the initial design and development activities of the project and the persistency framework⁷ in particular.

Deliver a first prototype hybrid data store capable of providing event data persistency in a production setting, and supporting also non-event data persistency. This will be a release for deployment and testing by developers and experts.

In *mid 2003* a more functional and robust version of the persistency framework will be delivered. This will be a release properly documented and packaged for general use by the experiments. Completion of the fully functional persistency framework is expected in early 2005.

The Applications Area expects to provide a distributed production environment integrating experiment-specific software with common software components and services, grid portal services and grid middleware in late 2003. By mid 2004 this infrastructure should be extended to support the data processing and analysis environment, enabling distributed end-user interactive analysis down to Tier 3 in the LHC grid.

The Applications Area will undertake further projects according to the project mandates delivered by the SC2. Schedules and work plans will be extended to encompass new projects as they appear. The projects so far mandated and initiated are in software process and persistency framework. A new project area in mathematical libraries has recently been established and a program in that area is being developed. Further project areas expected to be established in the next few months are in detector description, simulation, analysis tools and interactive frameworks. Applications Area management suggested a program covering about 20 activity areas initiated over a two-year period, and to date the SC2's program is similar to the suggested one. Of particular importance is the establishment of a global blueprint for the development of the different areas, defining their relations and an initial implementation technology. A team to define this has recently been established and it is expected to provide results by the end of the summer.

The approach to all projects undertaken in the Applications Area will be to iteratively design and develop software on a rapid cycle, with early and frequent feedback from users ensuring a strong coupling between the experiments and the project. Experiments will also participate directly in projects. A formal mechanism to involve experiment architects in project execution through an Architects Forum was established and is operational. The Applications Area employs very open and inclusive organization and communication, with the principal mechanisms (frequent open meetings, well developed web presence and mailing lists) established and operational. Applications Area projects (such as the persistency framework) may include an R&D component investigating new and evolving technologies, which may offer, improved technical solutions in the future.

⁷ An object persistency system organises the conversion between the format in which objects are stored within an active program (the *transient* representation) and the format in which they are stored on permanent storage such as disk and tape (the *persistent* representation). This allows different applications to have consistent access to all of the information that defines the object.

Applications Area roles in phase 1 of the LCG Project include

- *Preparation of the LHC computing environment.* Common software; grid interfaces to applications software; integration of applications software into grid prototypes and validation of computing model
- *Participation in and support for experiment data challenges.* Deployment, support and evaluation of common software, infrastructure and grid interfaces in data challenge production and analysis environments
- *Documenting project outcomes in a technical report.*

The principal applications activity areas are outlined below.

Application software infrastructure

Provide the basic environment for physics software development, documentation, distribution and support. Provide general-purpose scientific libraries, C++ foundation libraries, and other standard libraries. Provide software development tools, documentation tools, quality control and other tools and services integrated into a well-defined software process. Provide central support for software, and services for convenient distribution and remote installation. Organise courses and tutorials.

Common software frameworks

Participate in the development of common software frameworks, toolkits and applications supporting simulation, reconstruction and analysis in the LHC experiments. Participation in projects of broader scope than the LHC will focus on the needs and interests of the LHC experiments, including the distribution, support, documentation and maintenance needs of the LHC. Examples are GEANT 4 and ROOT.

Support for physics applications

Support the integration and deployment of common software tools and frameworks required by the LHC experiments. Provide the 'portal' environment required to mask the complexity of the grid from researchers while providing fully distributed functionality. Provide direct assistance to the experiments at the interface between core software and the grid and support the adaptation of physics applications to the grid environment. Participate directly in the integration and deployment of common software components within the LHC experiments.

Physics data management

Provide tools for storing, managing and accessing data handled by physics applications, including calibration data, metadata describing events, event data, and analysis objects. Provide database management systems supporting both relational and object-based data management, including support for persistent storage of objects. Provide persistency support for common applications where appropriate. Provide assistance to the experiments in the integration and deployment of physics data management services and their production application. Provide data management services meeting the scalability requirements of the LHC experiments, including integration with large-scale storage management and the grid.

Teams and Staffing

The following summarises the current staffing plans for the Applications Area of the project and indicates staffing levels as of now (they are growing steadily as participation and scope of activity increase). The project manages the technical activities of personnel involved in the effort. Sub-project teams will be headed by a project leader with overall responsibility for the design, development and timely delivery of software, consistent with the overall architecture and the requirements of the experiments. Larger projects will be further subdivided into work packages.

CERN staff

- Software specialists in IT and EP: IT/API, IT/DB physics data management, ROOT team. (currently ~3.5 FTEs, 9 in near term pledges)
- Special contributions for LCG Phase 1 (currently 5 full time people)
- Software specialists on the CERN staff in the experiments. (currently ~1.5 FTEs in the persistency framework project)

Experiment collaborators

- Experiments will support the common projects through participation of their software specialists external to CERN as well as CERN staff (currently ~1.5 non-CERN FTEs active and ~5 pledged in the persistency framework project)

External collaborators

- The Applications Area is open to collaborative arrangements with participants external to the LHC collaborations. The basic criterion is that the collaborators bring at least as much to the effort as they cost in terms of added requirements and constraints.

4.2. Fabrics

Introduction

The fabric comprises the basic infrastructure required for the LHC computing facility. Each Regional Centre will have its own well-managed fabric providing a high level of availability. The larger centres will have differing customers and constraints that will largely dictate the detail of their choices of equipment and management tools and policies. The Fabric Area of the LCG project has a direct responsibility for prototyping the Tier 0/1 centre at CERN, and so much of this section is concerned with this. However, there will be a number of activities involving technical staff from all of the Regional Centres to ensure exchange of information and experience, and encourage the use of common solutions.

At CERN the fabric components are expected to be largely the same for offline and online clusters. However the architectures of these clusters will not necessarily be the same and therefore the options for inter-changeability are unclear. It is expected that the fabric architecture work that will take place during 2003 will include a study of the possibilities for closer integration of offline and online facilities.

LCG Fabric Components

The actual computing fabric is comprised of both the physical components and the organizational infrastructure surrounding it including the operations and management of the hardware, software and services. For the purposes of this document, the definition of the architecture and physical implementation has been separated from its customisation and operation. It is assumed that the CERN fabric will be operated by the CERN computer operations infrastructure.

The fabric components fall into three main groups:

- Fabric Hardware
- Fabric Management Software
- Data Management Software

The components of the LCG Fabric Area are shown below. The basic functionality is shown but expressed in a way that does not imply implementation in any particular architecture, as this will be reviewed in the coming 18 months.

- Computing Elements
- Storage Elements
 - Mass storage systems
 - Hierarchical storage management, currently the CASTOR system.
- Network Interconnects
- Infrastructure design and capacity planning
- Fabric management software
- “Gridification”, presentation of resources available to the grid information services.
- Tools for the management, maintenance and operation of the fabric.
- End user support

What is not included as part of the current LCG fabric activities but is nevertheless required for the realization of a production service is shown below. These issues will need to be treated as the project evolves.

- Wide area networking support
- Operational Management
 - Hardware preparation
 - Fabric surveillance and physical interventions
 - Backup/Restore and disaster recovery
 - Grid infrastructure monitoring
- Replacement, upgrade and purchasing policies

Dependencies

The fabric architecture will depend on a number of other technical issues. One such is the dependency on the software and approach used for data management. Issues such as replication of data objects across devices and servers to improve throughput and availability will have an effect on the design, and cost, of the fabric.

Constraints

The Fabric Area has identified a number of activities required to address the requirements of the fabric for LHC. In particular, the following activities take into

account the high-level LCG project milestones already identified. These milestones are:

1. The final LCG prototype by the end of 2004. This may be based on technologies not used today.
2. A first LCG production service in 2003. Due to funding and time constraints this will be based on the technology used for today's production clusters.
3. The technical design report for the production service will be written in 2005. This should be written with the experience gained from running the prototype in practice and therefore the prototype needs to be close to the architectural and technical solution proposed for the production system two years later. The challenge here will be to take into account the rate of evolution of technology that will potentially present alternative solutions during the intervening years.

Outline of the Work Plan

The actual work plan can be divided into 3 distinct phases as far as the Fabric Area is concerned:

1. 2002/2003: The CERN fabric work will aim to produce a production service based on the existing architecture and philosophy. Materials investment has already been made to produce this service and little additional investment is seen during 2003.
2. 2003: During 2003 the detailed studies will be made to determine the correct architectural and technological approach for the construction of the full LCG prototype. This will leverage the work currently being undertaken as part of the current technology evaluation exercise⁸ and which is due to be complete in September 2002.
3. 2004: Purchasing, construction and commissioning of the full LCG prototype will be done.

During 2002 the work at CERN will mainly be performed as collaboration with the European DataGrid and DataTAG projects to produce a joint "production testbed". From the fabric management point of view, this will largely consist of stabilizing the existing testbed environment and increasing the number of nodes. Specifically the following activities are foreseen to take place in 2002 and 2003:

1. Extend the size of the production testbed to 200 nodes. This will then become the LCG-1 production cluster in 2003.
2. Implement increasing levels of automation and management to provide a fully automated installation and maintenance solution for the first production service.
3. Migrate existing production services (e.g. LXBATCH) to this new environment as soon as it becomes sufficiently stable.

The following activities are planned to take place during 2003. This work will be undertaken as a collaboration between the larger Regional Centres, all of whom have similar problems and challenges.

1. CPU Architecture Planning. This will address the issues and cost effectiveness of the CPU nodes that will be used. Issues such as diskless

⁸ PASTA III – for more information see <http://www.cern.ch/david/pasta/pasta2002.htm>

use, blade systems, dual/quad cpu's, manageability and upgradeability will be studied.

2. Disk Storage Planning. This will address architectural issues such as SAN vs. NAS systems, Storage over IP technologies and the potential for use of disk as a replacement for tape, leaving tape as an archive medium. Modelling of access patterns, file size and data distribution will be done.
3. Fabric Management Options. This includes understanding the technology options for creating a very large-scale cluster. Such options are the use of the European DataGrid work on fabric management and commercial products. Considerations include issues such as changing requirements based on experience, long-term maintenance and how to evolve custom developments.
4. Performance and Manageability Metrics. As part of the design of the fabric for LHC the required performance needs to be defined. The actual requirements on the Tier 0 and Tier 1 centre fabrics for analysis are not yet well defined.

It is also required to have some metrics concerning the manageability of the system in terms of the desired level of automation, interventions tolerated and personnel required. Such goals will have an effect on other decisions taken such as the required level of hardware support for remote management in the nodes and investment in automated management software.

2004 will largely be devoted to the activities required to produce the production facility including the market survey, tendering, purchasing, installation and commissioning.

Teams and Staffing at CERN

The primary responsibility of the Fabric Area is the prototyping of the local fabric at CERN. The resources for staffing come from three sources:

- CERN/IT Staff. The fabric components involve 3 main groups in the IT Division. The current estimates of manpower working on LCG activities are given:
 - Basic Services and Data Challenges - 6 FTE's
 - Data Services - 5 FTE's
 - Fabric Infrastructure and Operations - 4 FTE's
- Staff funded at CERN by special LCG contributions. These are integrated with the IT groups and sections involved in realizing the fabric milestones indicated above. In particular staff has been added to the following areas:
 - Data Services - 2 FTE's
 - Fabric Infrastructure and Operations - 3 FTE's
 - Architectural and Design Studies - 1 FTE
- European DataGrid. Three FTE's funded by the DataGrid project until the end of 2003 work on the development of fabric management tools. The people are integrated within the CERN team responsible for this work package in the DataGrid project, which is an essential part of the work of the Fabric Area.

Additional people are foreseen in 2003 to conduct the architectural studies and tests needed for the preparation of the final LCG production prototype.

4.3. Grid Technology

Introduction

The charter of the Grid Technology Area (GTA) is to deliver and support the middleware required to provide Grid services to the LCG project. In the short term the technologies emerging from the current round of Grid Research and Development projects in Europe and the US will be used.

Strategic Objectives

In the first year of the LCG, Grid services will be provided by the middleware produced by the European Data Grid project in Europe, and the Virtual Data Toolkit (VDT) of the US high energy and nuclear physics grid projects⁹ that has been adopted by the US ATLAS and CMS User Facilities projects.

A key function of the GTA will be to ensure the inter-operability of grid services deployed by the different regions. In order to achieve this the GTA will work with the middleware technology providers and the Regional Centres through the LCG deployment area to establish a support infrastructure capable of delivering reliable software versioning, release and packaging infrastructure to the LCG community. Grid middleware software integration and management for the production LCG service will be the responsibility of the Grid Deployment Area.

Middleware Components

The Grid Middleware will be used to providing the following services:

- Authentication and Authorization
- Transparent access to distributed data
 - File Replication and Caching
 - Data Movement
 - Grid - Storage Interface Specifications (allowing Grid services to access different mass storage systems in different nodes of the Grid)
- Grid Workload management
 - Job (processing and data placement) Scheduling Management
 - Grid - Fabric Interface Specifications
 - Resource Discovery, Request Planning and Management
- Monitoring and Diagnosis tools

The middleware technology is still under development and its evolution is likely to continue for several more years. The GTA will actively monitor the development of grid middleware and will evaluate alternative solutions as they emerge. Consequently, the GTA will make appropriate recommendations to the project, if the results of the evaluation indicate that the LCG project might benefit from a change in component supplier or technical direction.

⁹ PPDG – *The Particle Physics Data Grid* (<http://www.ppdg.net/>); GriPhyN – *Grid Physics Network* (<http://www.griphyn.org/>); iVDGL – the *International Virtual Data Grid Laboratory* (<http://www.ivdgl.org/>).

Relationship with other activities

Today there are a number of early producers of middleware software; open source suppliers (e.g. Condor, Globus, LCFG), Grid projects integrating, augmenting and developing middleware (e.g. GriPhyN, European DataGrid, CrossGrid, etc.), as well as strong interest in potentially commercial grid middleware suppliers such as IBM, Sun, Pallas, Microsoft and several others. The GTA will maintain close contact with all of these bodies to understand their technical and schedule roadmap to understand the potential interest of their technologies to the LCG.

Interoperability issues between the European DataGrid and VDT are expected to be addressed by the joint Intergrid¹⁰ project GLUE¹¹, endorsed by the European DataGrid, DataTAG and the other US physics grid projects. GLUE milestones will be targeted to meet the needs of the LCG production grid milestone for early 2003.

Deliverables

The GTA interacts with all the other LCG areas (Applications, Fabrics and Deployment) to ensure the supported grid middleware meets the functional and timeline needs of the project. The GTA will continuously review the set of grid middleware and services supported and propose adjustments, identify shortcomings, and survey the available market. The first middleware proposal for use by the LCG will be provided by October 2002.

The Grid Technology Area is responsible for working with the Grid Deployment Area to ensure that the grid middleware meets the production quality required for deployment and support. This will involve creating a management process including testing and validation procedures. The first version will be completed by December 2002.

The GTA will work with the Grid projects to engage resources to develop, run and support test applications that verify the functionality and throughput performance of the deployed grid services. The first version of this test suite will be completed in March 2003.

Middleware Support

Details of support mechanisms are yet to be understood. The grid projects receive funding for their own programmes of work directed towards their own objectives, and have limited lifetimes that are independent of the LCG programme. The LCG has to be able to negotiate clear support agreements for the software it uses, providing continuing maintenance during the period in which the software will be used. It is expected that initial support will be provided by a combined effort from the various grid projects and from the Grid Deployment Area of the LCG. This is currently being discussed. In the longer term we expect that the grid funding agencies will develop support programmes for the new software to ensure that it can be used by different sciences, or that industrial support options emerge. We hope that the LCG project itself does not have to undertake long-term middleware maintenance.

¹⁰ The High Energy and Nuclear Physics InterGrid Collaboration – see <http://www.hicb.org/>

¹¹ GLUE – the Grid Laboratory Uniform Environment, a project initiated by Intergrid to make recommendations on compatibility between the middleware developed in US and European grid projects.

The GTA is responsible for negotiating delivery and long-term support agreements with the Grid Middleware providers and working with the LCG management on the necessary formal agreements.

4.4. Grid Deployment

The Grid Deployment Area is concerned with managing the construction and operation of the prototype distributed computing service for LHC data analysis and the support of the experiment data challenges. This will use the innovative software provided by the Grid Technology Area to link the services operated by the many Regional Centres providing computing resources for the LHC experiments. As has been stated already, this will be a very early example of an operational grid handling large amounts of data. We expect in the longer term that the LCG grid will be part of a wider group of intersecting grids serving many user communities. However, there is much work to be done to arrive at a satisfactory conclusion of the first step. It will be hard to provide a high quality distributed virtual computing centre for LHC, and we should remain realistic about how much further we can go in Phase 1 of the project.

Grid Policy Issues

There are many items concerning the operation of the Grid that will require agreement between the Regional Centres, the collaborations, the funding agencies and other national and regional authorities. Examples include:

- acceptance of digital certificates issued by foreign authorities;
- formal rules governing the use of computing resources;
- policies for allocating resources to collaborations – authorisation, budgeting and scheduling;
- data migration and replication policies;
- software and hardware upgrade policies;
- accounting and reporting standards;
- service reliability and availability standards;
- usage of research networks;
- sharing of responsibilities and escalation procedures for supporting a distributed user community;
- sharing of the control of the local resources between the local system administrator and the grid administrator

The [LCG Grid Deployment Board](#) (GDB) will be set up, including representatives of the collaborations and the Regional Centres, to ensure that the necessary agreements and decisions can be made. While some of these can be worked out between the collaborations and the Regional Centres, in many cases it will be necessary to take account of policies and decisions made at the level of national governments, national physics institutes or funding bodies. The Regional Centre representatives on the GDB must therefore be at a senior management level.

Other technical coordination structures will be set up as needed to coordinate between Regional Centres at an operational level.

Outline of the Work Plan

The Grid Deployment Area is a joint activity of the Regional Centres, of which CERN is only one. In its role of host lab, CERN will be responsible for a number

of coordinating tasks, but it is anticipated that other Regional Centres will also provide infrastructure services.

- During 2002 the LCG staff will gain experience by working with the LHC collaborations and established Grid Technology projects in Europe and the United States, assisting them in running some parts of data challenges on their testbeds. In the fourth quarter of the year the first LCG grid service will be defined and planned, and the teams needed to deploy and support it will be established. The process for accepting and certifying software deliverables from the other areas will be established and we will begin to set up the necessary infrastructure that will be supported by a number of the larger Regional Centres.
- During the first half of 2003 an initial LCG Global Grid Service (LCG-1) will be set up, with the clear goal of providing a reliable, productive service for LHC collaborations. The service will begin with a small number of the larger Regional Centres, including sites in three continents. This is to reduce the complexity of the Grid while the service is being stabilised, but to ensure that potential issues of inter-operability between different regions are quickly identified and resolved.
- In the second half of 2003 the initial service will be stabilised and the support infrastructure strengthened. The service will be expanded to include additional sites and computing capacity. The target is to provide a round-the-clock service that can be used by the LHC collaborations for their data challenges and for their normal work, in the way that they currently use the central computing services at CERN and other centres.
- In the first quarter of 2004 a second major release of the service (LCG-2) will be made. This may include significant changes in the functionality of the Grid middleware.
- During the rest of 2004 the final service (LCG-3) will be progressively deployed, to be available by the end of the year. This service will include all the essential functionality required for the initial LHC production service, and will meet the reliability and capacity goals of the "50%" prototype¹² required for the final validation of the computing model.
- During 2005 the LCG-3 service will be operated as a reliable, available service for the LHC collaborations.

The principal areas of work are outlined below.

Supply of Grid Technology

The Grid Technology Area description has explained the problem of selecting appropriate middleware satisfying the LCG requirements in terms of functionality, support and maintenance, and inter-operability between solutions installed at different Regional Centres. The Grid Technology Area manager will propose a suitable set of middleware to the Grid Deployment Board (see below). In the short term (next few years) the suppliers are expected to be development projects funded by the European Commission or US agencies. The LCG will have to

¹² This is a prototype of the final LHC Grid that contains at least half the number of *components* that will be required for the final system in 2008 for *one* of the larger experiments. It will therefore have around 12-15% of the number of components required for the complete system.

negotiate with these projects or third parties to ensure appropriate maintenance and support.

Software Certification and Distribution

The Software Certification and Distribution team will integrate the grid middleware, standard libraries and tools, and various elements of the applications support environment into a package that can be distributed to LCG Regional Centres. The team will be responsible for managing the software release policy, defining and operating the test and certification process, packaging and distribution of the software, and providing the system for handling error reports and distributing fixes.

Grid Infrastructure

There are many infrastructure components that will be required to support an operational Grid, including:

- The Certificate Authorities (CAs) that issue digital certificates formally identifying the equipment from which the Grid is constructed and the people who may use it. In many cases the CAs are likely to be operated by national authorities, but it is anticipated that CERN and some other institutes will have to operate CAs for cases where there is no appropriate national service.
- Authorisation and registration services. The digital certificate identifies the person, but the rights of the person to use the Grid services will be governed by policies defined by physics collaborations and Regional Centre management. A service will be set up to coordinate the implementation of these policies, automating as far as possible the registration of users at the different Regional Centres.
- Grid-wide configuration management, including the management of the data defining the set of resources comprising a Grid and the users authorised to exploit it.
- The Grid requires a number of information services defining the configuration, and the computing resources and software services available. The information services will be provided by a set of high availability servers operated at several of the main Regional Centres to ensure appropriate levels of reliability and performance.

Management of the Grid infrastructure therefore involves central coordination, with the operational activity distributed across various organisations and computing centres.

Call Centre

A Grid Call Centre, or help desk, is required to support the service staff at the different Regional Centres and the deployment experts in the collaborations. The call centre is responsible for managing the problem reports, providing immediate solutions where possible and directing other problems to the appropriate experts, ensuring that they provide a timely response. The Call Centre also maintains documentation and organises training courses for its users. A global Grid will require 24-hour availability of the Call Centre, to support users in different

regions. It is anticipated that there will eventually be two or three coordinated Call Centres operating in different continents.

Integration Support for Experiments

The project expects to provide a small group of experts in the use of Grid middleware to work closely with the data challenge teams in the experiments to help them adapt their production systems to use the LCG services and give support during data challenge operation. This team may also develop tools and support facilities to mask the details of the Grid middleware and operational characteristics.

Grid Operation

Each Regional Centre is responsible for the operation of its local computing services.

Grid operation involves pro-active monitoring of the Grid – site availability, network performance, application throughput, data consistency, load sharing, etc. The operations centre detects problems and anomalies, reporting them to the local fabric management or network management as appropriate. It also takes remedial action where possible, such as dynamic reconfiguration of the Grid and changing data caching strategies. The Call Centres are kept informed of problems. As in the case of the Call Centre, the Grid Operations Centre will have to provide service on a 24-hour basis. At present the whole issue of Grid operations is little understood. The important elements will only become clear with practical experience of using a production Grid.

Teams and Staffing

The following summarises the current staffing plans for the Grid Deployment Area of the project.

CERN teams

- Software Certification and Distribution – senior team leader, with two software engineers and one or two trainees.
- Grid Infrastructure Coordination – the team at CERN would be responsible for operating the infrastructure services at CERN, including the CERN Certificate Authority, and for coordinating the global Grid infrastructure. Team leader, software engineer, senior administrative officer, one or two trainees.
- Integration Support for Experiments – senior team leader, with four or five people with experience of both Grid middleware and of running production for experiments.
- CERN Fabric Operation – Initially this would be a team dedicated to operation of the CERN prototype equipment, with a team leader, software engineer and two trainees. Later, once the Grid has proven that it can provide a stable service, this team would merge with the team operating the central physics services at CERN.

Teams in Regional Centres

- Regional Centres will provide the staff for operation of the computing fabric and implementation of the Grid deployment policies at their site. They will also ensure end-user support, and

provide an interface to the Grid Call Centre and Grid Operations Centre. The effort required will be a function of the size of the Regional Centre and the category of service that it has committed to provide.

- Certain of the Regional Centres will also operate parts of the Grid infrastructure.

- Grid Operations Centre – As noted above, this activity is at present not well understood, and could be the subject of future R&D investment by funding agencies. In the short term the LCG project will appoint a senior manager to set up the initial structure needed to provide operational coordination between the Fabric Operation activities in the different Regional Centres.

- Call Centre – We are at present discussing with organisations that may be willing to extend the services they provide for existing Grid projects to cover the initial needs of the LCG. Due to the global nature of the LCG service it will be necessary to have extended opening hours or have more than one call centre in different time zones.

- Certificate Authorities – In many cases these will be operated as national services for high-energy physics or a wider community. In other cases Certificate Authorities may have to be established specifically for the LCG.

5. Level 1 Milestones

The Level 1 milestones of Phase 1 of the project are summarised in this section. While there are specific high level milestones for the Applications Area, the progress of the work of the other three areas is marked by a single set of milestones concerned with the deployment of the LCG Grid service which, as explained in the next section, integrate deliverables from all four areas of the project.

M1.1 – Availability of the first LCG Global Grid Service (LCG-1)

June 2003

LCG-1 is a reliable Global Grid Service offering 24 hour, 7 day a week availability, including about ten Regional Centres in Europe, Asia and North America that are ready to take part. The grid provides a batch service for all four experiments for event production and analysis of the simulated data set. The middleware deployed should be the "converged" European-US toolkit recommended by the GLUE collaboration. This milestone is a functionality and existence test, with a related milestone (M1.3) 5 months later for which throughput and reliability will be the key measures. A level 2 milestone will define the target performance levels for this and milestone M1.3 six months before the due date of this milestone. All of these targets must be sustained during a 7-day period for the milestone to be considered met.

M1.2 – General release of hybrid event store **June 2003**

The first release of the hybrid data store capable of providing event data persistency in a production setting, and supporting also non-event data persistency. This will be a release offering robust basic hybrid persistency services, properly documented and packaged for general use by the experiments, but it will not yet have full functionality.

M1.3 – Fully operational LCG-1 Service & distributed production environment **November 2003**

This milestone corresponds to the LCG-1 service achieving defined target levels of capacity, throughput, reliability and availability. The targets, which will have been defined in the specification of the LCG-1 service, must be sustained during a 30-day period for the milestone to be considered met.

The milestone includes the availability of the technical service specifications and user documentation, and an appropriate user support infrastructure. The service will by this point be ready to expand to include additional Regional Centres.

It also includes the availability of a distributed production environment integrating experiment-specific software with common software components and services, grid portal services and grid middleware.

M1.4 – Distributed end-user interactive analysis from a Tier 3 Regional Centre **May 2004**

Extension of the distributed production environment integrating experiment specific software with common software components and services, grid portal services and grid middleware. This extension will support the analysis environment, enabling distributed end-user interactive analysis down to Tier 3 in the LHC grid. A related level-2 milestone, 6 months before the due date, will provide a full specification of what this milestone includes.

M1.5 – Fully operational LCG-3 Service **December 2004**

LCG-3 will include all essential functionality required for the initial LHC production service. LCG-3 will be used as a practical demonstration of the LHC computing model, including Tier 0, 1, 2 and 3 Regional Centres, providing input for the computing service TDR. LCG-3 will use the LHC Grid Toolkit, will have 50% of the components required for the 2007 production service for CMS or ATLAS, and will be used for the "20% milestones" of the experiments.

This milestone will be met when specified levels of performance and reliability have been met for a period of 30-days. These target levels will be defined 6 months before the due date of this milestone by a related level-2 milestone.

M1.6 – Full function release of persistency framework **March 2005**

Completion of the fully functional persistency framework

The Computing Service TDR will specify the requirements for the Grid that will be used for the first production services for the four LHC experiments. It will include details of the architecture, functionality, capacity, performance, throughput and availability. It will include the Regional Centre plans that will have been developed to meet these requirements, and will provide cost estimates and an overall installation and verification schedule. It is assumed that the TDR will be approved by the LHCC within three months following its availability, and may be used to provide data for the Memorandum of Understanding for Phase 2 of the project. The full process from acquisition to service verification is expected to take 12-18 months (according to the administrative procedures of the Regional Centres). The initial service must be in full production by September 2006 (6 months before data taking). The TDR will therefore be approved after the acquisition procedures have started, but before orders are placed.

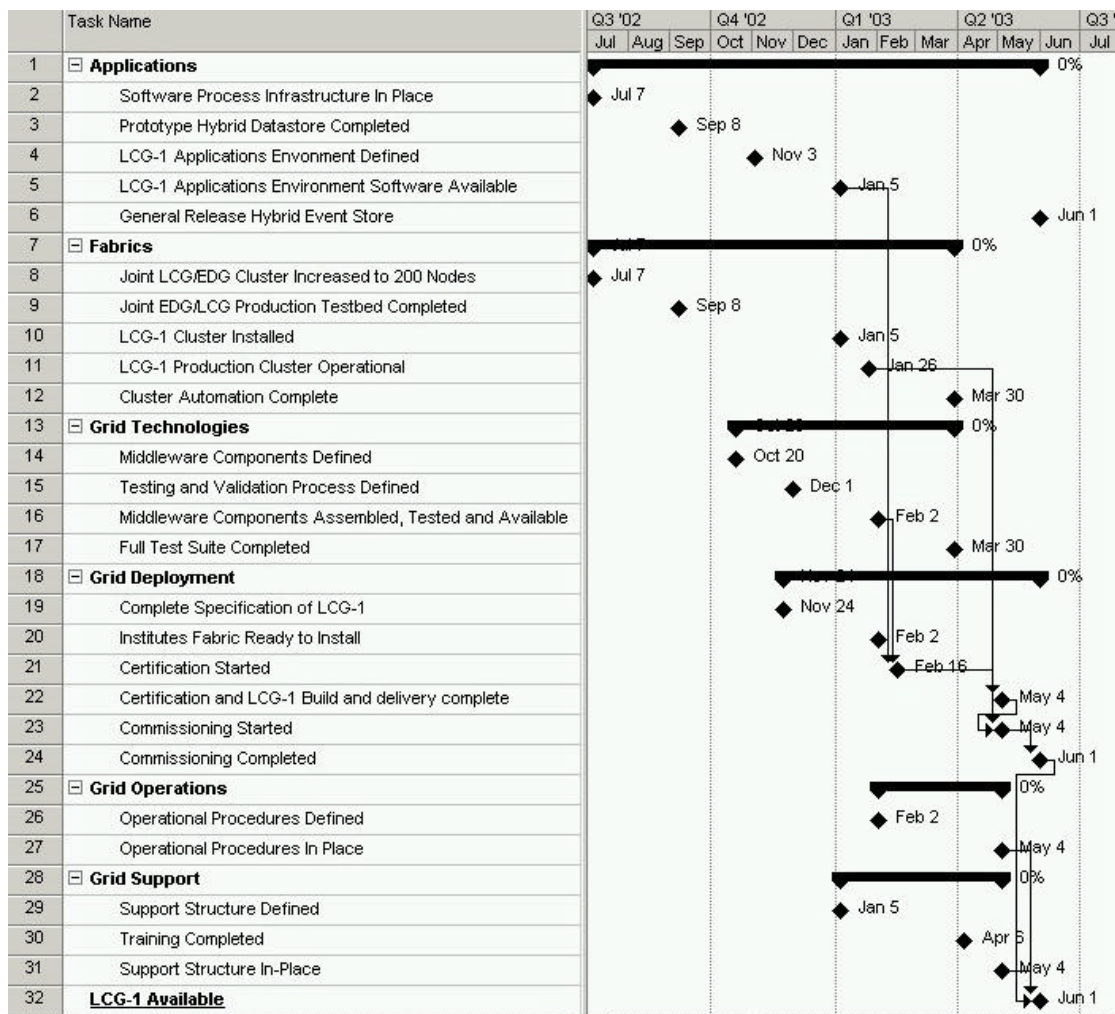
6. Consolidated High Level Plan

Despite all the current research activities associated with the many grid activities there are few, if any, examples of real production grids. The LCG recognises that there are many issues in taking a research activity through to a fully supported production service not all of which have yet been dealt with, such as overall grid management and user support. In order to expose and deal with the potential problems and issues it is desirable to move towards a production service as soon as it is practical to do so.

The following plan gives an initial schedule for the first production service, showing only the key milestones and deliverables. It deals with all the areas of the project foreseen by the LCG (Applications, Deployment, Technology and Fabrics) and the additional objectives and milestones in order to create a production service for Operations and User Support.

In the Applications Area, convergence on an agreed software environment and a stable version of the required software must be achieved. Fabrics will create an initial production test-bed with the European DataGrid and evolve that to become the first production facility for the LCG. Operations will define the operational procedures and implement them. Grid Technology will provide a recommendation and a set of software components for the middleware software to be used. Grid deployment will organize the participating institutes to ensure that systems and infrastructure are available to create the grid and will provide the agreed software to them. A user support structure will be needed that will require both training and procedures to be put in place for problem determination and resolution.

While this plan deals with deliverables there are a number of other policy issues that will need to be clarified during this time period. For example, security policies, management of maintenance, software environment update policies, hardware upgrade policies etc. These will be treated as part of the work of Grid Deployment.



7. Risks & Problems

The major sources of risk that have been identified so far for Phase 1 of the project are in the following areas.

- Management and coordination
- Grid technology
- Materials funding at CERN

Each of these areas is looked at below.

Management, Coordination and Priorities

As was noted in Section 1 *Scope and Goals of the Project*, the resources that will be used to perform the work of the project come from a variety of sources. The different sources have different levels of commitment to the project, and the management of the project will have different degrees of control and influence over their use. In many cases the project has only a coordinating role, and in others the project is only one of the users of the resources. It will be important for the project to build effective collaborations with these external resource holders. This is a common situation in high energy physics experiments, but the project

will obtain major resources from Regional Computer Centres and from Grid R&D projects, which may not have experience of such collaborations. An additional factor in this area is the fact that much of the special funding that has been provided for the project has been given because of a general interest in advancing Grid computing, rather than as investment in high energy physics. This may lead to conflicts in priorities.

Grid Technology Development

As has been mentioned, there are many producers of grid technology at the present time, most of them research organisations or projects supported by public R&D funding. The projects that are expected to supply solutions for the LCG are all based on the Globus toolkit, developed as a collaboration between the Argonne National Laboratory and the Information Sciences Institute at the University of Southern California in Los Angeles. Several industrial companies are showing interest in this technology, and some of them are also basing their developments on the ideas developed in the Globus project.

While this adds to the richness of choice, and gives some assurance that there will be grid technology available, most of these projects are very much research in nature, studying the basic techniques and functionality that will be required to create a grid. At this stage of maturity even the basic architecture used by the different developers is not stable. At the beginning of this year the Globus Project, along with IBM and some other companies, announced that they would develop a new architecture called OGSA (the Open Grid Services Architecture), based on the Web Services architecture developed by Microsoft. This may lead to solid industrial solutions, but in the short term it creates uncertainty for the projects that are based on Globus. The future of the current Globus toolkit is now unclear, while the capabilities and availability of the future OGSA toolkit are not well defined. It will also be a considerable time before we know whether any of the industrial solutions will provide the functionality needed for LHC.

A second problem with the immaturity of the technology is that there is, as yet, little emphasis on software quality and maintainability, or on the manageability of the grid as a system. This situation will certainly improve with time, as experience with real production grids is obtained, but manageability, like quality, are best factored into the software development from the design stage. We can expect the operation of a system built using first generation grid technology to encounter serious difficulties in this area.

A third difficulty is that the wealth of Grid projects will lead to different solutions, each with its strengths and weaknesses, complicated by national and regional considerations. It will be difficult to obtain consensus on a single set of middleware, or even on the standards needed to obtain compatibility between different systems, while it will not be possible to deploy a global grid service if a suitable agreement is not reached. This has been discussed in section 4.3 *Grid Technology*. We hope that the GLUE initiative will succeed in finding a solution to the problem acceptable to the main middleware suppliers, but there is a risk that the LCG project will have to make the very hard decision of selecting and trying to impose one solution.

Materials Funding at CERN

Annex B *Resources at CERN* refers to the resource overview presented to the project's Computing Resource Review Board on 23 April. The situation with personnel for Phase 1 of the project is rather good, with agreed or probable contributions covering almost all of the needs of the project. The situation with materials funding on the other hand is very poor, and only about 50% of the funding required has so far been identified. If this cannot be resolved over the coming months, we will not be able to cover the materials costs of the project in 2003.



<http://www.cern.ch/lcg>

LHC Computing Grid Project (LCG)

Status of High-Level Planning

Annexes

version 1 - 21 June 02

Prepared by the LCG Project Execution Board

A Data Challenge Plans

A.1 ALICE Data Challenge Planning

Computing Data Challenges

In 1997, the ALICE experiment and the IT division decided to initiate the Mass Storage Project. The goal of this project was to evaluate the LHC requirements in term of Mass Storage and the available products. After several tests, the goals were extended to cover more realistic and complete test-cases. Eventually, the scope has been extended to cover a significant part of the functions of the ALICE DAQ and offline projects and of the CERN Tier 0 and became the ALICE Data Challenge. Since 1998, ALICE experiments and the IT division have jointly executed several large-scale high throughput distributed computing Data Challenges. The goals of these are to prototype the data acquisition and the offline computing systems, to verify their integration, to assess technologies and computing models, to test hardware and software components of these systems in realistic conditions and to realize an early integration of the overall ALICE computing infrastructure. The ALICE Data Challenge 3 (ADC III) in 2001 included prototypes of the following systems: the ALICE data acquisition system, the ALICE offline computing, the CERN Tier 0 system and the Mass Storage System. The ALICE data acquisition software (DATE) was used for injecting simulated physics raw data of the ALICE TPC, produced with the AliRoot simulation program, in the data acquisition fabric, for event building and for controlling the data challenge. The data were then formatted with the ROOT I/O package and a data catalogue based on MySQL was established. The Mass Storage System CASTOR was used to store all the data.

The components tested during the ADCIII are shown in Figure 1.

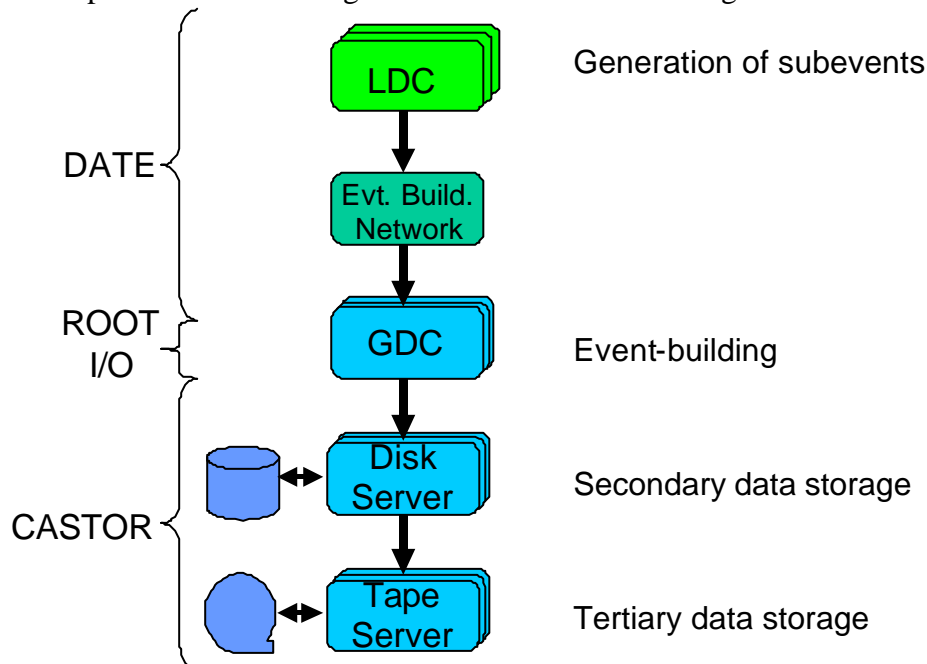


Figure 1 – Components tested during ADCIII

During the ADC III, extensive measurements have been performed and major performance milestones have been achieved: up to 550 Mbytes/s in the computing fabric, up to 85 Mbytes/s through the whole chain during a week and more than 110 TBytes of data stored into mass storage. The fabric has proven to be stable and to deliver stable performance for periods of up to a week. The LHC computing test-bed has been instrumental for the ADC III and ALICE plans to make further use of it for the larger tests scheduled in the future. The Data Challenges are the only opportunity to diagnose problems of implementation and scalability of all the components, in particular of the Mass Storage System. During the ADC III, in 2001, the planned milestone of 100 MB/s sustained to tape has not been achieved. The two main causes were the limited performance of the then available generation of IDE-based disk servers and the inefficient load balancing performed by CASTOR amongst the disk servers. The overall performance problems are not yet well understood. The ADC III key results are shown in Table 1.

Max throughput in DATE	550 Mbytes/s
Max throughput in DATE (ALICE-like traffic)	350 Mbytes/s
Max throughput in DATE+ROOT	240 Mbytes/s
Max throughput in DATE+ROOT+CASTOR	120 Mbytes/s
Average throughput during a week	85 Mbytes/s (> 50 TB/week)
Total amount of data through DATE	500 TBytes in DAQ (2*10 ⁷ events)
Total amount of data through DAQ+ROOT	200 TBytes
Total amount of data through DATE+ROOT+CASTOR and recorded onto tape	110 TBytes (100.000 files of 1 GBytes)
Number of DATE runs	2200 runs
Longest DATE run	86 hours, 54 TB
Metadata database	10 ⁵ entries

Table 1 – ADC III results

The ALICE Data Challenge IV configuration in 2002 is shown in Figure 2.

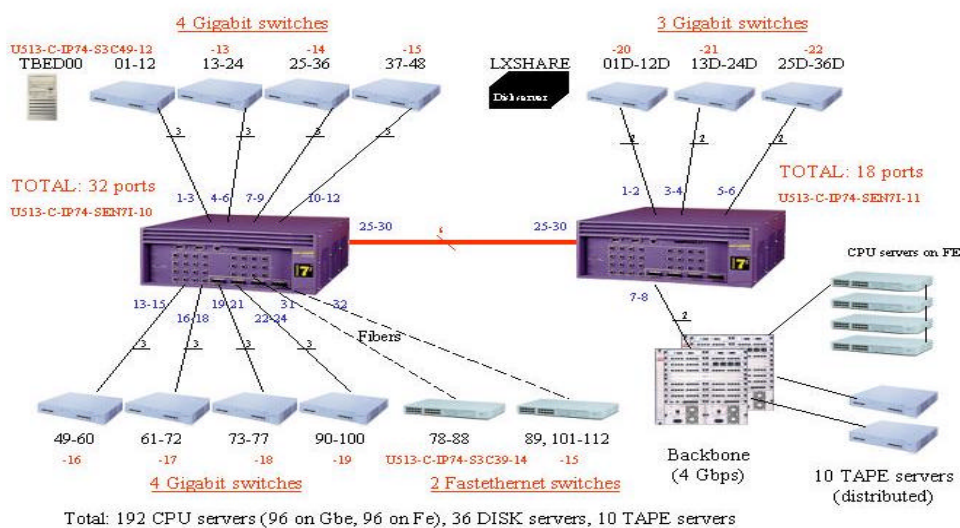


Figure 2 – ADC IV configuration

The ALICE Data Challenges milestones are listed in Table 2.

Date	MBytes/s	TBytes to MSS	Offline milestone
10/2002	200	200	Rootification of raw data -Raw data for TPC and ITS
5/2003	300	300	Integration of single detector HLT, at least for TPC and ITS - Quasi on-line reconstruction at CERN - Partial data replication to remote centres
5/2004	450	450	HLT prototype for all detectors that plan to use it - Remote reconstruction of partial data streams - <i>Raw digits</i> for barrel and MUON
5/2005	750	750	Prototype of the final HLT software Prototype of the final remote data replication (<i>Raw digits</i> for all detectors)
5/2006	750 (1250 if possible)	750 (1250 if possible)	Final test (Final system)

Table 2 – ALICE Data Challenges milestones

The requirements in terms of bandwidth to the Mass Storage System are shown in Figure 3.

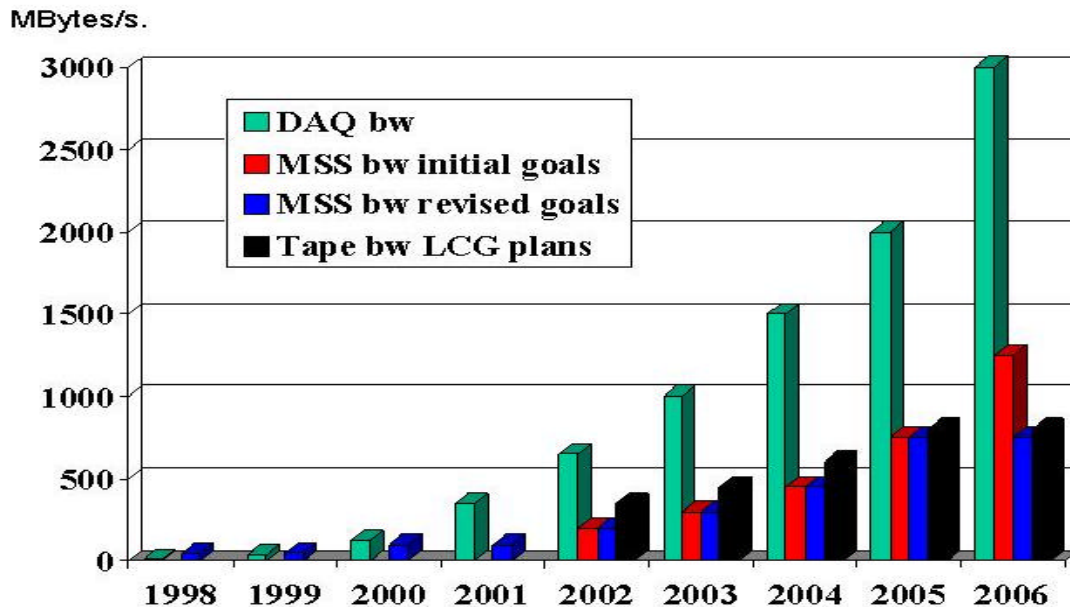


Figure 3 – ALICE DC bandwidth goals

The “DAQ bw” plot only concerns data acquisition software and network objectives. The “MSS bw initial goals” plot represents the ALICE initial goals for the bandwidth to mass storage system before the delay of LHC machine, as they have been requested to COCOTIME in June 2001.

The "MSS bw revised goals" numbers represents the ALICE goals for the bandwidth to mass storage system, taking into account the LHC startup delay. The assumption is that the first heavy ion run will be at the end of the first proton physics run, in March/April 2008 if there is no winter shutdown in 2007 like in the previous plan, or in November 2007 if there is a long winter shutdown. The revised goals also take into account the document "LCG Prototype – Capacity and Guidelines for Scheduling Usage in 2002/3" by the PEB (27th of February 2002), documenting the agreement reached in the LCG Project Execution Board (PEB) concerning the capacity profile and scheduling guidelines for the LCG prototype equipment at CERN.

The model described in this document, which has been endorsed by all the experiments, is plotted as "Tape bw LCG plans". It foresees to have 350MB/s in 2002 and 450MB/s in 2003, and this hopefully gives to ALICE the margin to obtain the sustained target performances (200MB/s in 2002 and 300MB/s in 2003). The document describes the CERN prototype as sized to demonstrate at least 2/3 of the data recording rate required by ALICE during the first ion run.

The capacity requirements are shown in Figure 4.

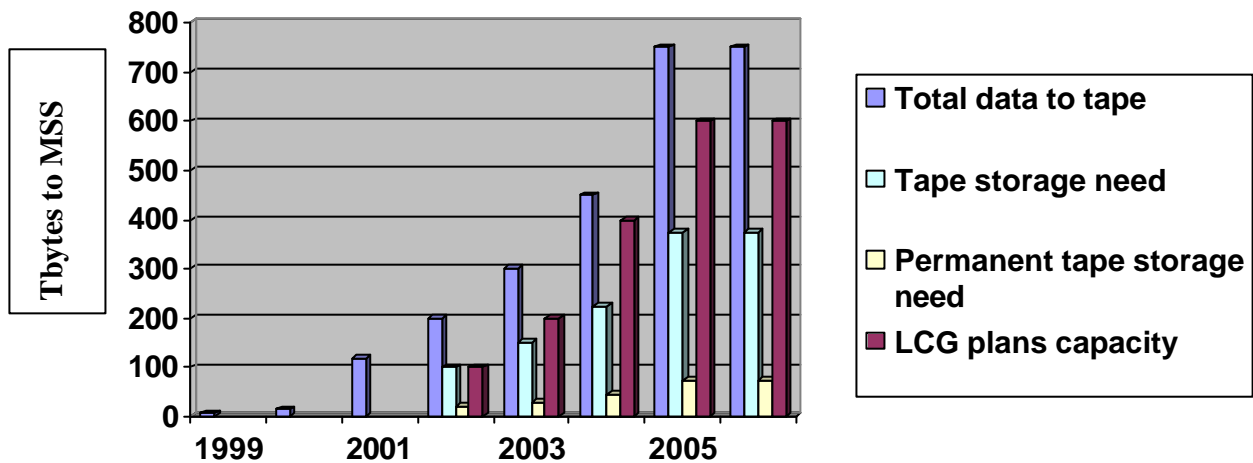


Figure 4 – Capacity requirements

Tapes will be recycled during the Computing Data Challenges, so we actually need ~50% of the total data to tape, as plotted in the "Tape storage need". Data will be used for testing purposes during one to two months after the Data Challenge, then they can be thrown away; only around 10% of data will be kept for longer, as indicated in the "Permanent tape storage need" plot.

Physics Data Challenges

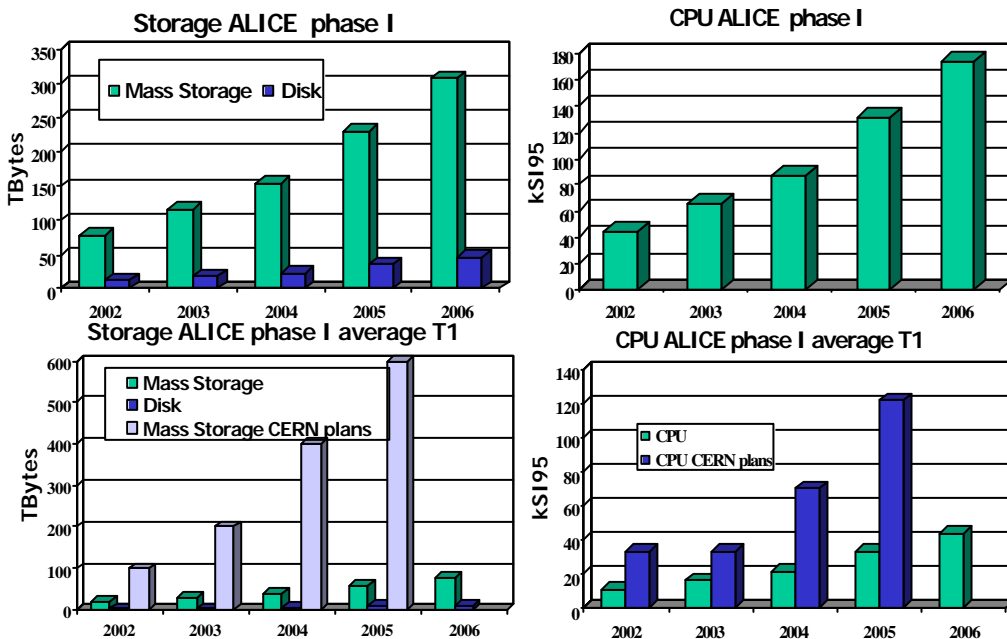
The goal of the Physics Challenges are:

- Determine readiness of the offline framework for data processing
- Validate the distributed computing model
- Percentages refer to the amount of simulation that we will perform during data taking (= 1/10 of real data)

The Physics Challenges milestones are reported in the table below.

PDC Period (milestone)	Size	Physics Objective
<u>06/01-12/01</u>	1%	pp studies, reconstruction of TPC and ITS
<u>06/02-12/02</u>	5%	First test of the complete chain from simulation to reconstruction for the PPR Simple analysis tools. Digits in ROOT format.
<u>01/04-06/04</u>	10%	Complete chain used for trigger studies. Prototype of the analysis tools. Comparison with parameterised MonteCarlo. Simulated raw data.
<u>01/06-06/06</u>	20%	Test of the final system for reconstruction and analysis.

The next series of plots show the evaluation of ALICE needs for the Physics Data Challenges. At present four participating centres (including CERN) are considered. The estimates are done assuming the target of simulating 5%, 10% and 20% of the events simulated in 1 year of standard data taking (which corresponds to 10% of the real events collected). The ramp up between challenges can change. These estimates are based on the resources strictly needed for the Physics Challenges and represent a minimum. Code development special studies and Grid experimentation are not taken into account. This evaluation is based on several large productions already done. Nevertheless the uncertainties are still sizeable. A 50% uncertainty factor should be considered.



A.2 ATLAS Data Challenge Plans

Introduction

The goals of the ATLAS Data Challenges are the validation of the Computing Model, of the complete software suite, of the data model, and to ensure the correctness of the technical choices to be made. It is understood that these Data Challenges should be of increasing complexity and will use the software which will be developed in the LCG project, to which ATLAS is committed, as well as the Grid middleware being developed in the context of several Grid projects like EU Data Grid or GridPP. The results of these data challenges will be used as input for a Computing Technical Design Report and for preparing a Memorandum of Understanding in due time.

The Data Challenges in Brief

We are considering here the first data challenges, DC0, DC1 and DC2, which will be run from now to 2004. For all of them it is essential to have physics content in order to bring the physicists community into the exercise and for a more sensible validation of the software. It is important to mention that for DC1, in 2002, a major goal is to provide simulated data to the High Level Trigger (HLT) community that has to prepare its own Technical Design Report (TDR) by mid 2003.

Data Challenges will of course continue beyond Data Challenge 2, probably annually, but detailed discussion of these has not yet begun.

- Data Challenge 0 is ending just now. It is mentioned here because it was both a continuity test of the code chain and a readiness test for the subsequent data challenges. The data was generated using “traditional” sequential file format that was converted to Objectivity, the persistency mechanism used up to now, before the reconstruction step. Note that we used both “old” data, simulated for the Physics TDR in 1998-99, and “new” data, simulated with a more up-to-date version of the ATLAS detector geometry.
- Data Challenge 1 runs from April 2002 to early part of 2003. It is divided into two phases:
 - In the first phase, May to July 2002, we will put in place the infrastructure and the production tools to be able to run the production worldwide, we also intend to generate and simulate several data sets, some as large as 10^7 events, for high-level trigger (HLT) studies and for physics analysis.
 - In the second phase, starting in September, we will also produce some physics samples, but will be more oriented to the testing of the new software, including Geant4, the new event data model and the evaluation of the new persistency mechanism (developed in the context of the LCG project). One of the aims is to be able to reconstruct and analyse worldwide a large sample of events over a short period of time.

As stated before, DC1 will run on sites worldwide, and will make as much use of Grid tools as is feasible.

- Data Challenge 2 runs from the last quarter of 2003 and will continue in the first quarter of 2004. The scope will depend on what has been or has not been achieved in DC1. The main goals will be to:
 - use Geant4 in a more intensive way for simulation
 - fully deploy the complete Event Data Model and the Detector Description
 - test the calibration and alignment procedures
 - perform large-scale physics analysis
 - use as much as possible the LCG common software
 - use widely the Grid middleware

To summarize:

Data Challenge	Size	Date
DC0	10^5 events	December 2001 – May 2002
DC1	A few samples of $< 2 \times 10^{**6}$ selected events	May 2002 – Early 2003
DC2	To be defined (10^8 events?)	September 2003 – June 2004

Table 1: DC plans

Estimation of the resources

An estimation of the resources needed for DC1, especially Phase I, has been elaborated on the basis of the previous large-scale jet production performed in 1997. Since it corresponds to specific physics channels for the purpose of the High Level Trigger TDR we name it DC1-HLT. Details on the production, such as type of generated events, selection criteria can be found in the ATLAS Physics note ATL-PHYS-97-102 (30 July 1997). The new production will be performed in a similar way, with a filtering between the event generation and the detector simulation. We foresee to apply an acceptance cut in the detector simulation at some η value; several sets of data will be produced with different parameters. The event generation and the detector simulation will be performed over a short period of time, typically in a 2-3 weeks period, using few 1000 CPUs worldwide. The numbers given in the tables below are for one dataset that will be produced with ‘reasonable’ assumptions.

In addition several samples of physics data with smaller statistics will be produced.

	Number of events	h range	Time per event SI95 sec	Total time SI95-sec	Total time SI95-hours
Generation	10^7				
Simulation	1.5×10^6	$ \eta < 6.$	16000 (*)	2.4×10^{10}	7×10^6
Reconstruction	1.5×10^6		280	0.4×10^9	2×10^5

Table 2: DC1-HLT CPU resources (for one sample)

(*) The CPU time per event is strongly correlated to the h coverage. The previous production was performed with $|h| < 2.7$.

	Number of events	Event size MB	Total size GB	Total size TB
Simulation	1.5×10^6	2	3000	3
Reconstruction	1.5×10^6	0.5	750	0.75

Table 3: DC1-HLT Volume of data (for one sample)

As stated before, pile-up will be added for two luminosities. It is not yet clear if it will be produced for the full sample and for the full h range. The values given below assume that the pile-up is performed for 1.5×10^6 events with $|h| < 2.7$.

L	Number of events	h range	Event size MB	Total size GB	Total size TB
2×10^{33}	1.5×10^6	$ h < 2.7$	2.6	4000	4
10^{34}	1.5×10^6	$ h < 2.7$	6.5	10000	10

Table 4: DC1-HLT volume of data with pile-up for 2 luminosities (for one sample)

The numbers are given for the events to be simulated for the HLT trigger. It assumes that only Monte Carlo ‘digits’ will be written to the output stream.

The Participating institutes

As of today several computing centres or institutes have expressed their interest to participate in the Data Challenges. Some will participate to the ‘massive’ production of data; some are more interested in concentrating to specific physics channels with smaller statistics.

The following institutes are already engaged in the phase I of DC1 or (*) will join soon:

- Czech Republic: Prague
- France: CCIN2P3 Lyon
- Germany: FZK Karlsruhe
- Greece: Thessaloniki (*)
- Italy: CNAF Bologna, Milan, Naples, Roma
- Nordic Cluster: Copenhagen, Oslo, Stockholm
- Spain: IFIC Valencia

- Switzerland: Manno (*), CERN
- UK: RAL, Birmingham, Cambridge, Glasgow, Lancaster, Liverpool
- Australia: Melbourne
- Canada: Alberta, TRIUMF
- Israel: Weizmann (*)
- Japan: Tokyo
- Russia: JINR Dubna, ITEP Moscow, SINP MSU Moscow, IHEP Protvino
- Taiwan: Taipei
- US: BNL

A.3 CMS Data Challenges

It is widely agreed that a relevant measure of a Computing Testbed is its complexity as opposed to its capacity; the LCG thus defines its Testbeds in terms of a fraction of the complexity of the final system. However, data challenges are more simply defined as a fraction of the computing capacity in the final distributed computing system (designed for $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ luminosity). Our 20% challenge is then sized to be equivalent to the planned 100Hz data-taking rate at the 20% luminosity expected for initial LHC running. The 10% and 5% challenges are thus planned for 50Hz and 25Hz respectively.

Using CMS tables as submitted to the LHC Computing Review and making PASTA-II based extrapolations, the 50% complexity CERN T0/T1 testbed proposed in LCG-3 is about 18% of a single experiments computing capacity requirements in 2009, and thus consistent within current estimates, with the 20% goal.

These CMS data challenges are distributed computing challenges, and thus place requirements on the CERN based computing, the offsite computing and the network and grid software. We assume that the CERN based computing prototype will reach the eventual 50% complexity/20% capacity target, but we do not assume that all external sites contributing to the challenges will meet these exact targets; which would not be sound economics. We instead assume an increasing involvement of centres as they come online and specify the integrated T1 and T2 complexity/capacity goals.

CMS productions and other challenges: In addition to the Level-1 and Level 2 challenges listed below, there are ongoing productions for various Physics uses that can typically run in the background on the resources available worldwide. There will also be Level-3 milestones/challenges addressing specific components of the system (such as single-site challenges, multi-site challenges, analysis challenges) but not requiring long-term majority use of the Testbeds.

5% Data Challenge. To be completed April 1 2004. A level-2 CPT Milestone
Purpose: Demonstrate validity of software baseline to be used for Physics TDR and for definition of Computing TDR.

This comprises the completion of a 5% data challenge, which successfully copes with a sustained data-taking rate equivalent to 25Hz at a luminosity of $0.2 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ for a period of 1 month. (Approximately 5×10^7 events)

The emphasis of the challenge is on the validation of the deployed grid model on a sufficient number of T0-2 sites.

We assume that 2-3 of the Tier-1 centres and 5-10 of the Tier-2 centres intending to supply computing to CMS in the 2007 first LHC run would participate to this challenge.

10% Data Challenge. To be completed April 1 2005. A level-2 CCS Milestone
Purpose: Validate LCG-3 hardware and middleware infrastructure in preparation for the LCG TDR.

This comprises the completion of a 10% data challenge, which successfully copes with a sustained data-taking rate equivalent to 50Hz at a luminosity of $0.2 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ for a period of 1 month. . (Approximately 10^8 events)

The emphasis of the challenge is on the validation of the deployed grid model on a sufficient number of T0-2 sites.

We assume that most Tier centres intending to supply computing to CMS in the 2007 first LHC run would participate to this challenge.

The output of this challenge will be used to test the full analysis chain and comes at the time when the Physics TDR is being finalized.

20% Data Challenge. To be completed April 1 2006. A level-1 CPT Milestone
Purpose: Demonstrate CMS readiness for Data-taking

This comprises the completion of a 20% data challenge, which successfully copes with a sustained data-taking rate equivalent to 100Hz at a luminosity of $0.2 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ for a period of 1 month. . (Approximately 2×10^8 events)

The challenge should include the full range of operations required for the analysis of CMS data under realistic conditions and must demonstrate the effective use of the distributed computing systems.

The challenge does not necessarily require the commissioning of 20% of the computing capacity dedicated to CMS, since common time-shared facilities, such as the T0 prototype, will be exploited.

We assume that all Tier centres intending to supply computing to CMS in the 2007 first LHC run would participate to this challenge. Of these we assume a minimum integrated offsite capacity of two 50% complexity T1's and 5 50% T2's.

A.4 LHCb Data Challenges

Computing Data Challenges

LHCb Computing Data Challenges (CDC) will concentrate on testing the various aspects of making collaborate the Grid Middleware (EUDG, since LHCb has no US collaborating institutes) with the LHCb software framework Gaudi.

The Gaudi framework will be used for all software applications in LHCb. Hence it will have to replace all its current interfaces to file storage systems with those to grid services. Users will submit their jobs to the Grid using a Gaudi-based user interface which will take care of preparing the job's JDL, submit the jobs, follow their execution and retrieve the output. This interface (called GANGA) will extensively use Grid services. GANGA is currently being developed in collaboration with ATLAS who also uses Gaudi as a framework.

CDC's will be used to test the functionality in production context of GANGA and the associated Grid services. It is foreseen that useful simulation data be produced during these Data Challenges, most probably "inclusive B" samples for which a very large number of events are required for trigger and high level trigger studies.

In the medium term, it is also foreseen to use Grid services for data analysis, possibly through GANGA, but in any case in an interactive environment. This analysis will consist in AOD production and AOD analysis by individual physicists.

Physics Data Challenges

LHCb will have to produce a large number of simulated events on one hand for studies needed for specific TDR's (LHCb-light, trigger...) and on the other hand in order to test their reconstruction and analysis software, as well as the analysis model. Although the demand is small compared with the two large experiments, a fair amount of storage and CPU power will be required in order to achieve the production goals in a reasonably short time.

The PDC's will be based on the existing LHCb production scheme, which partly uses already Grid services. Additional Grid usage will be incorporated in the production environment as soon as it will be validated in CDC's.

Participating sites

Many of the LHCb computing sites have expressed a firm interest and commitment in participating in both the CDC's and PDC's as from 2002. It is of the utmost importance that as many sites as possible get rapidly started with the installation and running of both Grid software and LHCb software. This will ensure their participation in the development of the necessary tools since the requirements can be quite different from one site to the other.

Plans for Data Challenge

The table below summarizes the currently planned DC's in LHCb as well as the participating sites.

Year	Participating sites	Planned DC's	
2002	Tier 0/1: CERN Tier 1 : Bologna, IN2P3, RAL, Nikhef, Moscow (September) Tier 2 : Liverpool, Oxford, Bristol, Edinburgh/Glasgow	CDC	<ul style="list-style-type: none"> • Evaluate installation procedures • Test job submission to fixed site with EDG • Control with PVSS • Prototype of new data management database (bookkeeping) • Prototype of configuration database • Measure production rates and network throughput at RCs • Implement data quality checking • First tests of job submission with resource broker • Test with analysis
		PDC	June/September June/July • LHCb-light, trigger TDRs
2003	+ Barcelona, Switzerland, Germany, Poland	CDC	<ul style="list-style-type: none"> • Production with job submission via the resource broker • Run new data management DB in production • First prototype of Ganga
		PDC	<i>June/July</i> HLT studies
2004	+ ????	CDC	<ul style="list-style-type: none"> • Production version of Ganga • Test the analysis model • Full use of Grid • Use of Gauss (G4 based simulation) in production • Tests of data quality checking in production

B Resources at CERN

See [CERN-C-RRB-2002-01-Rev](#) *LCG Phase 1 Project Resources at CERN, Report to the Computing RRB of 23 April 2002.*

C Resources in Experiments

C.1 Resources in ALICE

This is a first estimate of both hardware and human resources available within ALICE for the LCG Project Phase I.

Hardware Resources

As far as the hardware resources are concerned the information is contained in the attached table. In the table the current situation is shown. It reports the list of the planned Tiers-1 and 2 with resources declared available for 2001 and 2002. Numbers are continuously updating, plans for the ramp-up in the next years will be given in the context of the evaluation performed for Regional Centres within the Grid Deployment Project Area.

<i>TYPE</i>	<i>COUNTRY</i>	<i>LOCATION</i>	<i>CPU IN 2001/2 (SI2K)</i>	<i>DISK IN 2001/2 (TB)</i>
TIER-1	CERN	CERN	64	4.5
TIER-1	France	Lyon	100	10.5
TIER-1	Germany	Karlsruhe	10	7
TIER-1	Italy	Bologna (CNAF)	35	
TIER-1	UK	RAL		1
TOTAL TIER-1			209	23
TIER-2	Armenia	Yerevan	1	0.1
TIER-2	Croatia	Zagreb	3.4	0.2
TIER-2	Germany	GSI Darmstadt	48	
TIER-2	India	Kolkata		
TIER-2	Italy	Bari		
TIER-2	Italy	Bologna		
TIER-2	Italy	Catania	10	5
TIER-2	Italy	Torino	14	6
TIER-2	The Netherlands	NIKHEF		
TIER-2	Poland	Warsaw		
TIER-2	Rumania	Bucharest		
TIER-2	Russia	Moscow		
TIER-2	USA	OSC/OSU	20	3.4
TOTAL TIER-2			96.4	14.7
GRAND-TOTAL TIERS -1,2			305.4	37.7

Human Resources

ALICE has decided to participate, with an active role, to the LCG project. This has resulted in a full commitment to the SC2 and PEB work and in the leading of three Project RTAGs (Grid Applications, Mass Storage and Software Process).

As far as the different Project Areas are concerned, here follows an estimate of the ALICE contributions.

Area	FTE's	Note
Data Challenges	1(DAQ)+1(Offline)	This is an ALICE specific activity, but all experiments count on ALICE to lead the way for I/O rate and MSS testing.
Grid activities	10 (~30 people)	A large part of this is part of the common activities on the testbed together with the WP8 of DataGrid
Applications	0.5	This is ALICE's contribution to ROOT support.

C.2 Resources in ATLAS

ATLAS is committed to participate actively in the common projects that will be launched by the LCG, especially in the Applications Area. It should also be noticed that the collaboration is heavily involved in several of the Grid projects, worldwide.

Up to now ATLAS has been active in the first RTAGs launched by the SC2 and, once the projects are in place, expects to divert some of its resources from ATLAS-specific development to the corresponding common LCG effort. As an example, we expect contribute the equivalent of 2-3 FTEs to the project on Persistency. Obviously it is not possible to establish a general rule at this stage: the contribution will depend on the project itself, and the level and duration of participation will be defined case by case.

In addition, in common with the other LHC experiments, the ATLAS Chief Architect expects to devote up to ~20% of his time to the LCG project through participation in the Architects' Forum.

Concerning the participation in the Data Challenges the contribution of the institutes and Regional Centres is dealt with by the ATLAS National Computing Board (NCB). Based on information received so far, we believe that the effort requirements needed to run the first data challenges are well covered. Typically it is of the order of 2-4 FTEs for the major Regional Centres and institutes and of 0.5-1 FTE for the smaller ones.

C.3 Resources in CMS

Human Resources

CMS plans to contribute with 10 FTE to the application area, corresponding to about 50% of our person-power involved in similar activities in CMS.

Contribution to GRID deployment and development is more involved. CMS is currently contributing to:

- EU Data Grid project
- Other Grid projects, such as PPDG and GriPhyN

- Operation of CMS data production in Regional Centres
- Development of data production tools not presently covered by grid middleware
- Integration and deployment of grid middleware as it comes available

The current total effort in these areas is estimated to be 29 FTEs. We expect person-power to be available at the same level for the coming years

Hardware Resources

The following table summarizes the hardware resources used in the current data production challenge focused on producing simulated events for the DAQ TDR (mainly high-level-trigger studies).

All Regional Centres have committed to ramp up their resources to match the data challenge plan described in section B.3

Number of Regional Centres	11
Number of Computing Centres	21
Number of CPU's	~1000
Largest Local Fabric	176 CPUs
Number of Production Passes for each Dataset (including analysis group processing done by production)	6-8
Number of Files	~11,000
Data Size (Not including fz files from Simulation)	17TB
File Transfer by GDMP and by perl Scripts over scp/bbcp	7 TB toward T1s 4TB toward T2s
Intra-Fabric throughput	250 GB/hour 1TB/day

C.4 Resources in LHCb

Contribution to Data Challenges

Several sites are contributing significantly to the coming LHCb data challenges, as stated in the corresponding section of this document. Currently, besides the coordination of LHCb grid activities, with one FTE from Oxford, the estimation of the number of FTEs is as follows.

LHCb DC site	Number of FTEs
Barcelona	0.5
Bologna	2
Bristol	0.5
Cambridge	0.5
CERN	2
Edinburgh	1
Liverpool	1
Lyon	1
Moscow	0.5
NIKHEF	2
Oxford	2.5
RAL	1.5

The same people are preparing for the PDCs and the CDCs. It is to be noted that a joint effort has been launched with ATLAS on common software for interfacing the Grid to the experiment framework Gaudi, called GANGA. Manpower is available essentially from GridPP for this work and the LHCb contribution in this effort is included in the above table.

Contribution to common software

Depending on the outcome of coming RTAGs, LHCb is fully committed to contribute to common software activities that could emerge. Already manpower has been allocated on the persistency project. It is to be noted however that the LHCb core software team is relatively small and is already contributing very significantly to the development and maintenance of the Gaudi software framework that is used amongst others by ATLAS.

Whenever new common projects emerge, LHCb will make all efforts to contribute to those projects. LHCb is also prepared to possibly house some projects and to house LCG new hires for the development or integration of common project software.

D Resources in Regional Centres

The current understanding of the planning of resources in Tier 1 and Tier 2 Regional Centres during the first phase of the project is maintained on the Web. See http://www.cern.ch/lcg/peb/rc_resources. This information is updated regularly, and people responsible for Regional Centres are invited to provide missing information.

The state of planning and authorisation varies from centre to centre, and the resources available are shared between different activities, such as grid prototyping, LHC physics data challenges, and other experiments. This data is provided for information to the LHCC, to show the scale of the resources being built up in Regional Centres. The resources in Regional Centres that will be made available for the LCG services will have to be determined as part of the detailed planning process of the project, but it seems clear that there will be sufficient resources available, both in terms of capacity and manpower, to operate the LCG services proposed for Phase 1.

E Management and Reporting Structure

The project is managed by three boards, each with a clearly separate function.

The Software and Computing Committee (SC2)

The Software and Computing Committee defines the requirements and approves the major strategic directions of the different areas of the project. It approves the work plan devised by the Project Execution Board (see below), and monitors the progress of the work and the adherence to the requirements. The CERN Director General appoints the SC2 chair (the present chair is Matthias Kasemann from Fermilab), and the membership includes two coordinators from each LHC experiment, technical managers from each region represented in the Project Overview Board (see below), a representative from CERN/EP Division, the leader of CERN/IT Division, and the Project Leader. The chair of the POB is invited. The SC2 normally meets once per month.

The SC2 gets technical advice from Requirements and Technology Assessment Groups (RTAGs). These are short-lived committees, set up with a restricted mandate to give advice and proposals on specific topics. Each RTAG includes one member from each experiment, and experts selected according to the topic. The RTAG would normally complete its work and report within two months.

The Project Execution Board (PEB)

The Project Execution Board is responsible for carrying out the work of the project which, as explained in section 1 (*Scope and goals of the project*), is organised as four areas: Applications, Computing Fabric, Grid Technology and Grid Deployment. The chair of the PEB is the Project Leader, appointed by the CERN Director General (the present Project Leader is Les Robertson from CERN). The Project Leader appoints the Area Managers. The members of the PEB include one delegate from each of the LHC experiments, the Area Managers, the Resource Manager, the Planning Officer, and the Chief Technology Officer. In addition external members, representing related projects and organisations may be appointed by the Project Leader.

The PEB meets generally once per week. It organises the planning of the implementation of the project, including obtaining agreement on the resources, schedule and milestones for each activity among the various teams and institutes concerned. It manages the progress and direction of the project, ensuring that the results conform to the requirements expressed by the SC2. It may also identify areas that require further definition, or issues that require resolution by the SC2. The work plan and progress are regularly presented to the SC2.

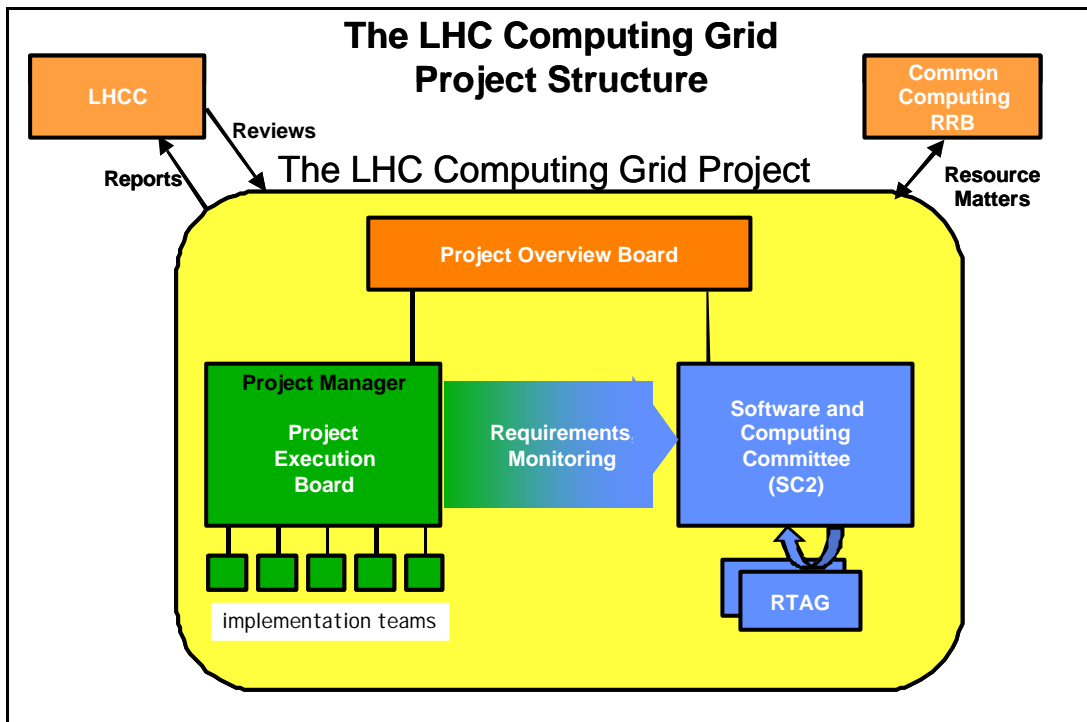
In the Applications Area an Architects *Forum*¹³ has been established to discuss and reach agreement on strategic issues affecting the work of the project. The membership includes the experiments' computing architects, with the computing coordinators invited.

¹³ See http://www.cern.ch/lcg/peb/documents/architects_committee.htm

In the Grid Deployment Area a *Grid Deployment Board*¹⁴ will be established in which the computing management of the Experiments and the Regional Centres can discuss and take, or prepare, the decisions necessary for planning, deploying and operating the LCG Grid.

The Project Overview Board (POB)

Project Overview Board is responsible for monitoring the overall progress and performance of the project, taking decisions on issues that are beyond the competence of the PEB and SC2, and dealing with differences or disputes that cannot be resolved at a lower level. Chaired by the CERN Director for Scientific Computing, the membership of the POB will include the CERN Director for Collider Programmes, the CERN/IT Division Leader (secretary), the LHC experiment spokespersons, and representatives of participating Regional Centres (this representation has not yet been decided). The chairperson of the SC2 and the Project Leader will be in attendance.



Reporting and Relationships

The Project Leader reports on the progress of the project at each SC2, which also hears regular status reports from the Area Managers of the PEB. The aim is to maintain a good bi-directional communications channel between the SC2 and the PEB, with a view to identifying problems at an early stage and finding mutually satisfactory solutions.

The SC2 chair and the Project Leader report to the POB.

The project provides reports on the human and materials resource situation to the Computing Resource Review Board, which meets twice per year and is chaired by

¹⁴ See <http://www.cern.ch/lcg/peb/ Documents/Grid-deployment-board.doc>

the CERN Director for Scientific Computing. At the autumn meeting the budget for the following year will be presented.

The project reports on planning and progress to the LHCC, as requested by its chairperson. The LHCC will carry out a review of the project about once each year.

A project management plan will be produced when the full WBS is completed and resource loaded by the first quarter of 2003.

F Extracts of WBSs

The planning of the different areas of the project is in progress. These extracts represent a high level view of the current state.

F.1 Applications

LCG
PBS

Description

Project PBS

1.1

Applications

ID: lcg:/lcg/app

1.1.1

Application software infrastructure

ID: lcg:/lcg/app/infra

Provide the basic environment and infrastructure for physics software development, including standard tools and libraries, software process with associated tools and infrastructure, and infrastructure for software distribution, documentation, and communication.

2002/4/15: #10204: Initial process and infrastructure team in place

2002/4/20: #10205: Process and infrastructure planning meeting Done!

2002/5/1: #10206: Basic LCG code development and repository services in place Done!

2002/6/15: #10209: Software process workplan to SC2

2002/7/1: #10210: Functional software process infrastructure in place

1.1.1.1

Software process

ID: lcg:/lcg/app/infra/process

1.1.1.1.1

Repository management

ID: lcg:/lcg/app/infra/process/cvs

1.1.1.1.2

Configuration management

ID: lcg:/lcg/app/infra/process/config

1.1.1.1.3

Release management

ID: lcg:/lcg/app/infra/process/release

1.1.1.1.4

Librarian

ID: lcg:/lcg/app/infra/process/librarian

1.1.1.1.5

Software development tools

ID: lcg:/lcg/app/infra/process/swdev

1.1.1.1.6

Testing and quality assurance

ID: lcg:/lcg/app/infra/process/qa

1.1.1.1.7

Performance, scalability and benchmarking

ID: lcg:/lcg/app/infra/process/test

1.1.1.2

Standard libraries and packages

ID: lcg:/lcg/app/infra/lib

1.1.1.2.1

Mathematical libraries

ID: lcg:/lcg/app/infra/lib/math

1.1.1.3

Software distribution

ID: lcg:/lcg/app/infra/dist

1.1.1.4 **Documentation infrastructure**

ID: lcg:/lcg/app/infra/write

1.1.1.5 **Communication and reporting tools**

ID: lcg:/lcg/app/infra/comm

1.1.2 **Physics data management**

ID: lcg:/lcg/app/data

Provide tools for storing, managing and accessing data handled by physics applications.

1.1.2.1 **Persistency framework**

ID: lcg:/lcg/app/data/frame

2002/6/5: [#10207: Hybrid event store workshop](#)

2002/6/9: [#10208: Persistency framework status report](#)

2002/7/12: [#10211: Persistency framework workplan to SC2](#)

2002/9/15: [#10212: First prototype release of hybrid data store](#)

2003/1/31: [#10214: Persistency framework 'interface-complete' release](#)

2003/6/1: [#10215: General release of hybrid event store](#)

2005/3/1: [#10218: Full function release of persistency framework](#)

1.1.2.1.1 Architecture and infrastructure

ID: lcg:/lcg/app/data/frame/inf

1.1.2.1.2 Object streaming core software

ID: lcg:/lcg/app/data/frame/core

1.1.2.1.3 Object streaming components

ID: lcg:/lcg/app/data/frame/stream

1.1.2.1.4 Relational database infrastructure

ID: lcg:/lcg/app/data/frame/rdbinf

1.1.2.1.5 Relational components

ID: lcg:/lcg/app/data/frame/rdb

1.1.2.2 **Conditions database**

ID: lcg:/lcg/app/data/cond

1.1.3 **Common frameworks**

ID: lcg:/lcg/app/frame

Participation in the development of common software frameworks, toolkits and applications supporting simulation, reconstruction and analysis in the LHC experiments.

1.1.3.1 **Framework services**

ID: lcg:/lcg/app/frame/service

2002/12/1: [#10213: Prototype object dictionary service released](#)

1.1.3.1.1 Object dictionary

ID: lcg:/lcg/app/frame/service/dict

1.1.3.2 **Event processing**

ID: lcg/lcg/app/frame/event

1.1.3.3 **Simulation**

ID: lcg/lcg/app/frame/simu

1.1.3.4 **Detector description and geometry model**

ID: lcg/lcg/app/frame/desc

1.1.3.5 **Interactive environment**

ID: lcg/lcg/app/frame/inter

1.1.3.6 **Physics analysis**

ID: lcg/lcg/app/frame/analysis

1.1.3.7 **Visualization**

ID: lcg/lcg/app/frame/visual

1.1.4 **Physics applications support**

ID: lcg/lcg/app/physics

Support the integration and deployment of common software tools and frameworks required by the LHC experiments.

1.1.4.1 **Grid portals**

ID: lcg/lcg/app/physics/portal

Provide the 'portal' environment required to mask the complexity of the grid from researchers while providing fully distributed functionality.

1.1.4.1.1 Distributed production

ID: lcg/lcg/app/physics/portal/prod

2003/11/1: [#10216: Distributed production environment using grid services](#)

1.1.4.1.2 Distributed analysis

ID: lcg/lcg/app/physics/portal/analysis

2004/5/1: [#10217: Distributed end-user interactive analysis from a Tier 3](#)

1.1.4.2 **Core/grid interface**

ID: lcg/lcg/app/physics/gridint

Direct assistance to the experiments at the interface between core software and the grid and support the adaptation of physics applications to the grid environment.

1.1.4.3 **Experiment integration**

ID: lcg/lcg/app/physics/expint

Direct participation in the integration and deployment of common software components within the LHC experiments.

1.1.6 **QA/QC and testing**

ID: lcg/lcg/app/qa

1.1.7 **Documentation and training**

ID: lcg/lcg/app/doc

1.1.8 **Support and maintenance**

ID: lcg/lcg/app/support

[1.1.9](#) **Project planning**
ID: lcg/lcg/app/plan

Items: 42

F.2 Fabrics

The following are the level 1 and level 2 milestones for fabrics:

ID	Task_Name	Duration	Start_Date	Finish_Date	Predecessors
1	Technology Working Groups	250 days	Jan 1 '03	Dec 16 '03	
2	CPU Architecture proposal for full size prototype	220 days	Jan 1 '03	Nov 4 '03	
3	Disk Storage Architecture proposal for full size prototype	220 days	Jan 1 '03	Nov 4 '03	
4	Fabric Management proposal for full size prototype	30 days	Nov 5 '03	Dec 16 '03	2,3,5
5	Management Working Group	90 days	Jan 1 '03	May 6 '03	
6	Performance and Maintainability metrics defined for Full size prototype	90 days	Jan 1 '03	May 6 '03	
7	Joint EDG/LCG Production Testbed	100 days	May 14 '02	Sep 30 '02	
8	Low level cluster automation	70 days	May 14 '02	Aug 19 '02	
9	Increase to 200 Nodes	70 days	May 14 '02	Aug 19 '02	
10	Testbed technical description	70 days	May 14 '02	Aug 19 '02	
11	Testbed construction complete	30 days	Aug 20 '02	Sep 30 '02	8
12	First LCG Production Service	156 days	Oct 1 '02	May 6 '03	
13	Production service technical description	90 days	Jan 1 '03	May 6 '03	
14	Full Cluster Automation	120 days	Oct 1 '02	Mar 17 '03	7
16	Full Size Prototype	240 days	Jan 1 '04	Dec 1 '04	
17	Market Survey/Tendering Process	90 days	Jan 1 '04	May 5 '04	1
18	Additional Equipment purchased	90 days	May 6 '04	Sep 8 '04	17
19	Full size prototype construction	60 days	Sep 9 '04	Dec 1 '04	18

F.3 Grid deployment – High level WBS

WBS id	Title	Description	start	end	weeks
1.4	Grid Deployment				
1.4.1	Establish Grid Deployment Teams		1-Jun-02	15-Sep-02	15
		<i>Define the teams needed to deploy and operate the LCG service, negotiating sharing of responsibilities between LCG, grid technology projects, and Regional Centres</i>			
1.4.2	Certification and Build		1-Aug-02	31-Dec-05	178
		<i>LCG certification, build, distribution & release management</i>			
1.4.2.1	Initial training of CERN team, working with DataGrid WP6		1-Jul-02	30-Nov-02	22
1.4.2.2	Define certification & build process		1-Aug-02	30-Sep-02	9
1.4.2.3	Assemble test suites		1-Sep-02	31-Dec-02	17
1.4.2.4	Negotiate/define LCG-1 content, config, participants, delivery process & schedule		1-Sep-02	30-Nov-02	13
1.4.2.5	Define release process & schedule for LCG-1		1-Dec-02	31-Jan-03	9
1.4.2.6	Certification & build - LCG-1		1-Jan-03	30-Apr-03	17
1.4.2.7	Prepare release documentation - LCG-1		1-Mar-03	30-Apr-03	9
1.4.2.8	Release cycle LCG-1		1-May-03	30-Apr-04	52
1.4.2.9	Define release process & schedule for LCG-2		1-Sep-03	31-Oct-04	61
1.4.2.10	Certification & build - LCG-2		1-Jan-04	31-Mar-04	13
1.4.2.11	Prepare release documentation - LCG-2		1-Feb-04	31-Mar-04	8
1.4.2.12	Release cycle LCG-2		1-Apr-04	31-Dec-04	39
1.4.2.13	Define release process & schedule for LCG-3		1-Jul-04	31-Aug-04	9
1.4.2.14	Certification & build - LCG-3		1-Sep-04	30-Nov-04	13
1.4.2.15	Prepare release documentation - LCG-3		1-Oct-04	30-Nov-04	9
1.4.2.16	Release cycle LCG-3		1-Dec-04	31-Dec-05	56
1.4.3	Grid infrastructure		1-Jul-02	31-Dec-05	178
		<i>Coordination and operation of the LCG grid infrastructure at CERN and at several Regional Centres</i>			
1.4.3.1	Initial training of CERN team, working with DataGrid WP6		1-Jul-02	30-Nov-02	22
1.4.3.2	Develop infrastructure plan for LCG-1		1-Oct-02	31-Dec-02	13
1.4.3.3	Set up coordination infrastructure for CAs		1-Dec-02	31-Jan-03	9
1.4.3.4	CA coordination and operation		1-Feb-03	31-Dec-05	152
1.4.3.5	Set up registration & authorisation infrastructure		1-Dec-02	28-Feb-03	13
1.4.3.6	Coordination & operation of registration & authentication infrastructure		1-Mar-03	31-Dec-05	148
1.4.3.7	Set up infrastructure for LCG-1		1-Dec-02	28-Feb-03	13
1.4.3.8	Operate infrastructure for LCG-1		1-Mar-03	30-Apr-04	61
1.4.3.9	Develop infrastructure plan for LCG-2		1-Oct-03	30-Nov-03	9
1.4.3.10	Set up infrastructure for LCG-2		1-Dec-03	31-Jan-04	9
1.4.3.11	Operate infrastructure for LCG-2		1-Feb-04	31-Dec-04	48
1.4.3.12	Develop infrastructure plan for LCG-3		1-Jun-04	31-Aug-04	13
1.4.3.13	Set up infrastructure for LCG-3		1-Sep-04	31-Oct-04	9
1.4.3.14	Operate infrastructure for LCG-3		1-Nov-04	31-Dec-05	61
1.4.4	LCG Grid operation		1-Oct-02	31-Dec-05	170
		<i>This is the distributed Grid-level operation of the LCG services. This is not yet understood.</i>			0
1.4.4.1	Grid operation requirements		1-Oct-02	31-Dec-02	13

1.4.4.2	Grid operation plan for LCG-1	1-Jan-03	28-Feb-03	8
1.4.4.3	Commissioning of LCG-1	1-May-03	30-Jun-03	4
1.4.4.4	Operation of LCG-1	1-May-03	30-Apr-04	52
1.4.4.5	Long term grid operation plan	1-May-03	31-Aug-03	17
1.4.5	Regional Centre Technical Liaison	1-Oct-02	31-Dec-05	170
1.4.6	CERN Fabric operation	1-Jun-02	31-Dec-05	187
	<i>This is the operation of the fabric providing Grid services at CERN. Initially this will include only the LCG prototype, but progressively most of the "LXBATCH" service should be integrated in the CERN Grid fabric</i>			
1.4.7	Call Centre	1-Nov-02	31-Dec-05	165
	<i>Help desk, error report management, support for systems staff at Regional Centres.</i>			
1.4.8	User Support	1-Sep-02	31-Dec-05	174
	<i>Support for data challenge, production office, end users</i>			
1.4.9	Liaison with Grid technology suppliers	1-Jun-02	31-Dec-05	187
1.4.10	Grid Deployment Board	1-Jun-02	31-Dec-05	187

G Glossary

A

ALICE

A Large Ion Collider Experiment at CERN

AliEn

ALice production ENvironment

AliROOT

AliRoot is the name for the ALICE Off-line framework for simulation, reconstruction and analysis.

AOD

Analysis Object Data

API

Application Programming Interface

ATLAS

A Toroidal LHC Apparatus (experiment at the LHC accelerator at CERN)

B

BNL

Brookhaven National Laboratory (at Upton, NY, USA)

C

C++

C++ Programming Language

CA

Certificate Authority

CACR

Centre for Advanced Computing Research (at Caltech)

CARF

CMS Analysis and Reconstruction Framework

CASTOR

CASTOR is an implementation of a Managed Storage system, developed at Cern

CDC

Computing Data Challenge

CDR

Central Data Recording (system at CERN)

CERN

European Laboratory for Particle Physics (in Geneva, Switzerland)

CESR

Cornell Electron Storage Ring (in Ithaca, NY, USA)

CMS

Compact Muon Solenoid (experiment at the LHC accelerator at CERN) + Code Management System + Compiler Monitor System + Conversation Monitor System

CNAF

CNAF is the INFN National Centre for R&D into Informatics and Telematics Technology

COCOTIME

Is the CERN meeting for physics data processing capacity planning

CPU

Central Processing Unit

CSNET

Computer Science Network

CTP

Computing Technical Proposal

CVS

Concurrent Versions System

D

DAQ

Data Acquisition

DataGrid

European Union Project on Grid Technologies and Applications

DataTAG

European Union Project to create a large-scale intercontinental Grid

DB

Data Base + Data Buffer

DBM

Data Base Manager

DBMS

Data Base Management System

DBS

Data Base Server

DC

Data Collection + Data Communication +
Data Control + Device Control + Direct
Current + Data Challenge

DESY

Deutsches Elektronen-Synchrotron (in
Hamburg, Germany)

DFS

Distributed File System

DM

Distributed Memory + Data Management

DMA

Direct Memory Access/Addressing

DoE

Department of Energy (US Funding
Agency)

DS

Data Services (CERN-IT group)

DST

Data Summary Tape

E**EDG**

European DataGrid Project

EDM

Event Data Model

ESA

European Space Agency

ESD

Event Summary Data

ESNET

Energy Science Network (in the USA)

EUDG

European Union DataGrid - Grid
middleware project
<http://www.eu-datagrid.org>

F**FIO**

Fabric Infrastructure and Operations
(CERN-IT group)

FNAL

Fermi National Accelerator Laboratory
(Batavia, Illinois, USA)

FOCUS

Forum on Computing Users and Services
(at CERN) + Forum of Control Data Users

FTE

Full-Time -Equivalent

FTP

File Transfer Protocol (Internet)

FZK

Forschungszentrum Karlsruhe

G**GANGA**

GAUDI/ATHENA and Grid Alliance

GAUDI

The Gaudi project is an open project for
providing the necessary interfaces and
services for building HEP experiment
frameworks in the domain of event data
processing applications.

Gb

Gigabit (1,024 Megabits)

GB

Gigabyte (1,024 Megabytes)

GD

Grid Deployment

GDB

LCG Grid Deployment Board

GDM

Grid Data Management

GGF

Global Grid Forum

GIOD

Globally Interconnected Object Database
(Caltech/CERN/HP joint project)

Globus

Grid Middleware Provider (USA)
<http://www.globus.org>

GLUE

Grid Laboratory Universal Environment

GML

Generalized Markup Language

GridPP

Grid for Particle Physics – UK national
project for grids in high energy physics

GriPhyN

Grid for Physics Network – Grid
middleware project (USA)
<http://www.griphyn.org>

GUI

Graphical User Interface

H**HEP**

High Energy Physics

HEPIX

HEP Unix environment

HICB

High Energy Physics Intergrid
Collaboration Board

HLT

High Level Trigger

HSM

Hierarchical Storage Management

HPSS

High Performance Storage System

HTML

HyperText Markup Language

HTTP

HyperText Transfer Protocol

I**IBM**

International Business Machines
(Corporation)

IDE

Imbedded Drive Electronics + Integrated
Development Environment (Borland) +
Integrated Drive Electronics + Intelligent
Drive Electronics + Interactive Design and
Engineering + Interface Design
Enhancement

IEEE

Institute of Electrical and Electronics
Engineers

IGUANA

Interactive Graphical User ANALysis
(CMS software project)

IHEP

Institute for High Energy Physics, Russia

INFN

National Institute for Nuclear Physics,
Italy

IN2P3

National Institute for Nuclear and Particle
Physics, France

IP

Internet Protocol

IT

Information Technology

ITEP

Institute for Theoretical and Experimental
Physics, Moscow

iVDGL

International Virtual Data Grid Laboratory
– Grid project (USA)
<http://www.ivdgl.org>

J**JANET**

Joint Academic Network (UK)

JAR

Java Archive File (format)

JCL

Job Control Language

JDL

Job Description Language

JINR

Joint Institute of Nuclear Research,
Moscow

JPL

Jet Propulsion Laboratory

JTB

Joint Technical Board

JVM

Java Virtual Machine

K**kb**

kilobit (1,024 bits)

kB

kilobytes (1,024 bytes)

Kbps

Kilobits Per Second

KBps

Kilobytes Per Second

KEK

High Energy Accelerator Research
Organisation (in Tsukuba, Japan)

KHz

Kilohertz

L

LAN

Local Area Network

LCB

LHC Computing Board (CERN committee ceases activity in 2000)

LCFG

Automated Installation and Configuration System developed at the University of Edinburgh

LCG

LHC Computing Grid Project

LEP

Large Electron Positron Collider (at CERN)

LHC

Large Hadron Collider (at CERN)

LHCb

Large Hadron Collider Beauty experiment

LHCC

LHC experiments Committee

LINUX

Operating system named after Linus Torvalds

LXBATCH

The main Linux batch service at CERN

M

Mb

Megabit (1,024 kilobits)

MB

Megabyte (1,024 kilobytes)

MBps

Megabytes Per Second

Mbps

Megabits Per Second

MONARC

Models of Networked Analysis at Regional Centres (HEP common project)

MoU

Memorandum of Understanding

MSS

Mass Storage System + Multiprotocol Switched Services (IBM)

MySQL

Open Source version of an SQL database

N

NAS

Network Attached Storage

NASA

National Aeronautics and Space Administration

NCSA

National Center for Supercomputing Applications

NFS

Network File System (Sun)

NIKHEF

The National Institute for Nuclear Physics and High Energy Physics, Holland

NSF

National Science Foundation

NSFNET

National Science Foundation Network

O

Objectivity

Object Database Commercial Product

ODL

Object Definition Language

ODBMS

Object Database Management System

ODMG

Object Database Management Group

OGSA

Open Grid Services Architecture

OO

Object Oriented

P

PASTA

Technology Review Organised by CERN

PAW

Physics Analysis Workstation

Pb

Petabit (1,024 Terabits)

PB

Petabyte (1,024 Terabytes)

PBS

Product Breakdown Structure

PC

Personal Computer

PDC

Physics Data Challenge

PDM

Physics Data Management

PEB

LCG Project Execution Board

Persistency

The ability to store data objects

POB

LCG Project Overview Board

PPDG

Particle Physics Data Grid (USA)
<http://www.ppdg.org>

Q**QA**

Quality Assurance

QC

Quality Control

R**RAID**

Redundant Arrays of Independent Disks +
 Redundant Arrays of Independent Drives
 + Redundant Arrays of Inexpensive Disks

RAL

Rutherford Appleton Laboratory

RAM

Random Access Memory

RC

Regional Centre + Readout Crate

RCC

Regional Computing Centre

R&D

Research and Development

RD45

Common project addressing the issue of
 object persistency

RDMS

Russia Dubna Member States

Remedy

Commercial product for building trouble
 ticketing systems

ROOT

An Object Oriented Data Analysis
 Framework, including a data persistency
 service

RTAG

Requirements and Technology
 Assessment Group

S**SA/SD**

Structured Analysis/Structured Design

SAN

Storage Area Network

SCADA

Supervisory Control And Data Acquisition

SCB

Software and Computing Board (of CMS)

SC2

LCG Software and Computing Committee

SCRAM

Software Configuration, Release And
 Management (CMS software product)

SCSI

Small Computer Systems Interface

SCTB

Software and Computing Technical Board
 (of CMS)

SI95

SpecInt 95 – A measure of processor
 performance

SI2K

SpecInt 2000 – A measure of processor
 performance

SLAC

Stanford Linear Accelerator Center
 (California, USA)

SPC

Scientific Policy Committee

SQL

Structured Query Language

SWIB

Software Institution Board (of US-CMS)

T		V	
Tb	Terabit (1,024 Gigabits)	VDT	Virtual Data Toolkit – the grid middleware package being used by the US high energy physics community
TB	Terabyte (1,024 Gigabytes)	VO	Virtual Organisation
TCP	Transmission Control Protocol	W	
TCP/IP	Transmission Control Protocol/Internet Protocol	WAN	Wide Area Network
TDR	Technical Design Report	WBS	Work Breakdown Structure
Tier-0, Tier-1 etc.	Indicator of the hierarchical level of a regional centre	WPn	Work Package number
U		WWW	World-Wide Web (Internet)
UNIX	(AT&T Bell Laboratories Operating System)	X	
URL	Uniform Resource Locator	XML	Extensible Markup Language
		Y	
		Z	