



LHC Computing Grid Project – LCG

LHCC Comprehensive Review

Project Status - Overview

Les Robertson
22-23 November 2004

**CERN – European Organisation for Nuclear Research
Geneva, Switzerland**

les.robertson@cern.ch

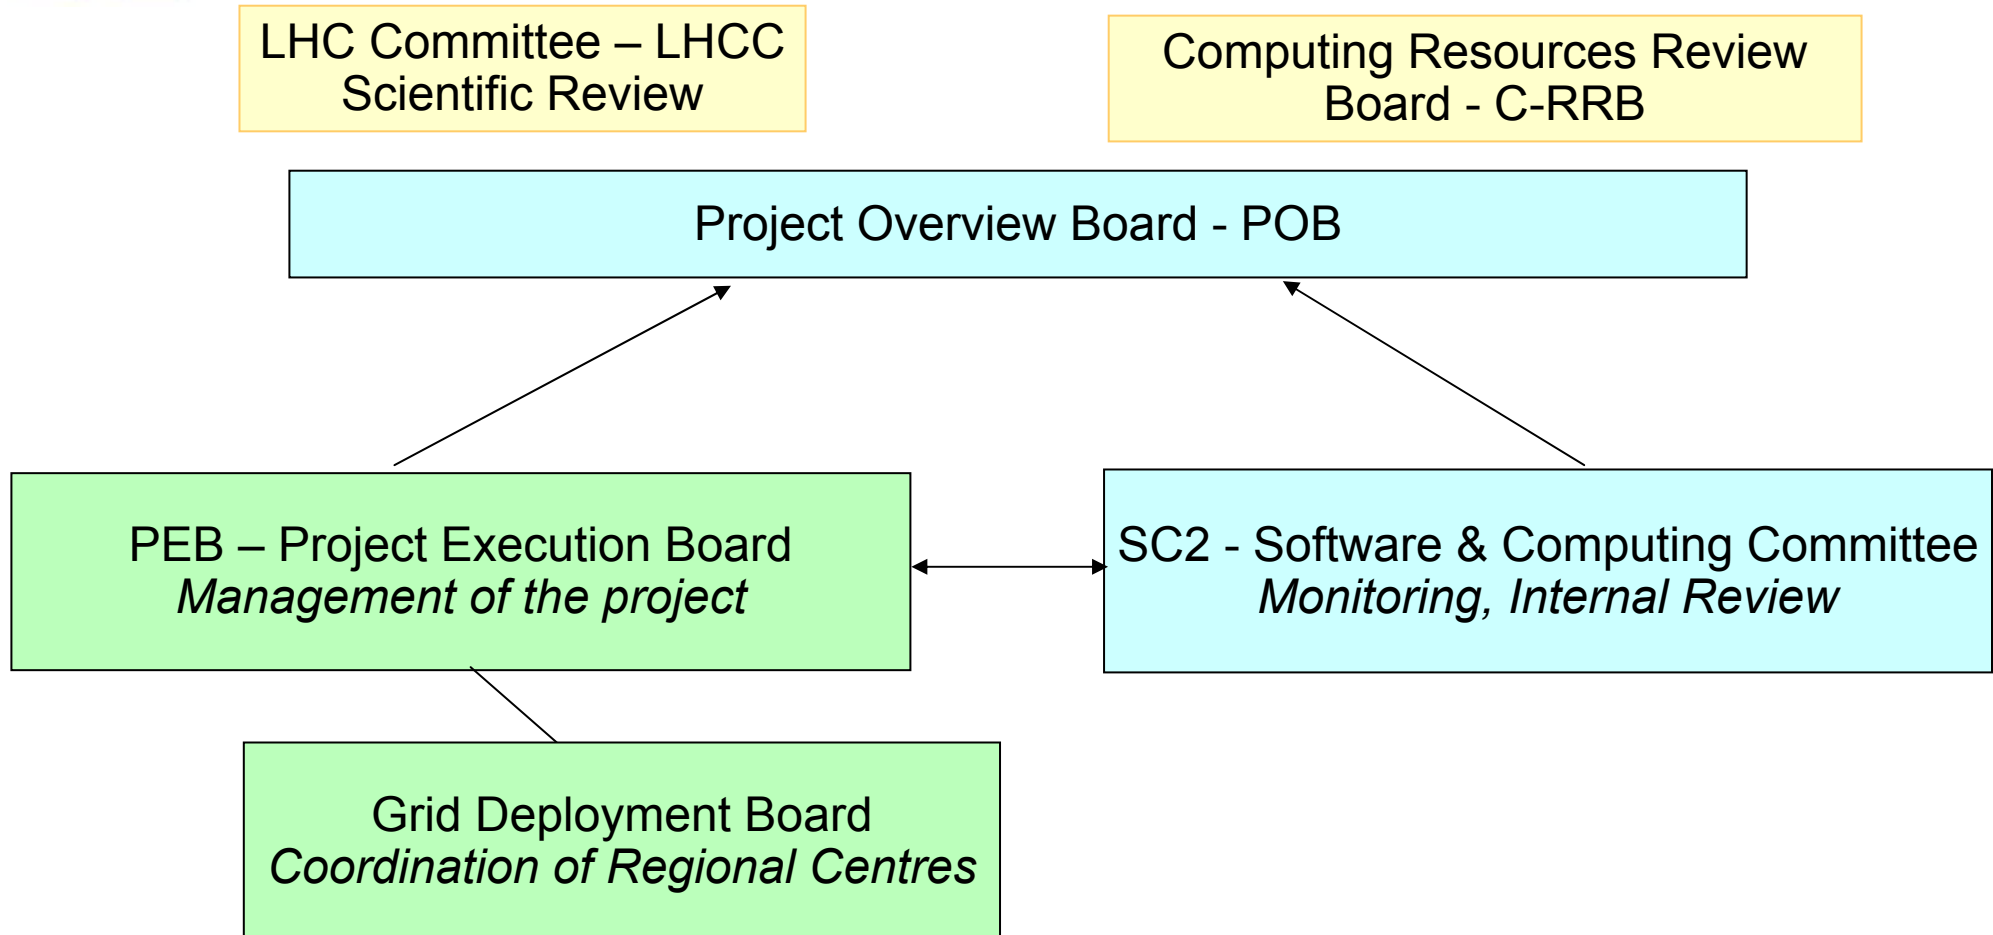


Summary

- *Structural changes*
- *Preparations for Phase 2*
- *Status & Concerns*
- *Resources*



LCG Project Organisation





PEB – Project Execution Board

- Management board of the project
 - Supervises the work of the different project activities within the broad scope agreed with the SC2 and POB
- Membership
 - Project leader / chair – Les Robertson
 - Area managers
 - Applications Torre Wenaus CERN Fabric Bernd Panzer
 - Middleware Frédéric Hemmer Grid Deployment Ian Bird
 - ARDA Project Massimo Lamanna
 - LHC computing coordinators
 - ALICE Federico Carminati ATLAS Dario Barberis
 - CMS David Stickland LHCb Nick Brook/Philippe Charpentier
 - Other members
 - Grid Deployment Board Chair Kors Bos
 - SC2 Chair Matthias Kasemann
 - EGEE Technical Director Bob Jones
 - Network planning David Foster
 - Planning Officer, PEB Secretary Jürgen Knobloch



SC2 Committee

- Standing internal review body –
 - Monitors the project on behalf of the POB.
 - In-depth analysis of the quarterly report
 - Annual review of each of the activity areas of the project
- Feedback and guidance to the PEB
- Analysis of findings to the POB

- Chairman: Matthias Kasemann /DESY
- Ex-officio members:
 - PH Deputy Division Leader, IT Division Leader
- Experiment representatives:
 - ALICE Wisla Carena ATLAS JimShank
 - CMS Albert De Roeck LHCb Gerhard Raven
- Regional Centres representatives (appointed by POB)
 - France Eric Lançon Germany Marcel Kunze
 - Italy Federico Ruggiero UK Tony Doyle
 - USA Lothar Bauerdick
- Secretary: German Cancio Melia
- Project Leader in attendance



Relation of LCG and EGEE



- LCG has been setting up the operation of a grid infrastructure for HEP experiments
 - combining national and regional grids
 - middleware certification, integration, distribution, operation management, ..
 - agreed middleware package – currently “LCG-2”
- EGEE has a goal of extending this into a general grid infrastructure for science in Europe and beyond
- EGEE brings new manpower -
 - 50% of its budget is for operation
 - 70 partners – almost all involved in HEP
- EGEE is funded for 2 years, with a *probable* extension for at least another 2 years
- LHC comes fully online just as EGEE ends



LCG & EGEE



- **Goals of LCG in working closely with EGEE**
 - Maximise the benefits of EGEE funding for LHC
 - Generalise the HEP grid to other sciences
 - Be well positioned to benefit from a future science grid infrastructure (cf GEANT research network backbone in Europe)
 - BUT - ensure that LCG can continue if there is no EGEE successor
- **LCG-2 and EGEE share the core operation**
 - LCG Grid Deployment Manager acts as EGEE Operations Manager
- **EGEE middleware development activity has a strong focus on HEP requirements**
 - Hardening, re-engineering of functionality needed by LCG (*gLite*)
 - Very close collaboration between EGEE middleware and the ARDA project of LCG (early access to development prototypes, regular technical meetings)
 - Commitment to deliver rapidly to LCG/EGEE deployment –
 - taking over, as quickly as practicable, from the LCG-2 middleware
 - EGEE middleware manager sits in the PEB
- **EGEE provides funding (4 FTEs) for ARDA**
 - Which in turn acts as the HEP pilot application for EGEE
- **Cross-representation on management boards (PEB, POB)**



Grid Operation

- The LCG project has a double role –
 - Operating the LCG-2/EGEE grid (middleware certification, integration, distribution, operation management, ..)
 - Through the [Grid Deployment Board](#) – coordinating the wider set of resources available to LHC
- Most of the resources around the world are integrated into the LCG-2/EGEE structure
- Separate infrastructures in the US (Grid-3) and the Nordic countries (NordGrid)
- Active programme aimed at compatibility/inter-working of LCG-2/EGEE and Grid3



Planning for Phase 2

- **Phase 2 Planning group**
 - end November – complete Tier-1 resource plan
 - first quarter 2005 –
 - assemble resource planning data for major Tier-2s
 - understand the probable Tier-2/Tier-1 relationships
 - initial plan for Tier-0/1/2 networking
- **Developing a plan for ramping up the services for Phase 2, linking Service Challenges –**
 - check out the infrastructure/service to iron out the problems before the experiments get fully involved
 - schedule allows time to provide permanent fixes for problems encountered
- **and Experiment Computing Challenges -**
 - checking out the computing model and the software readiness
 - **not** linked to **data challenges – which should use the regular, permanent grid service!**
- **TDR editorial board established → TDR due July 2005**



Status & Concerns



Applications Area

- All Applications Area projects have software deployed in production by the experiments
 - POOL, SEAL, Geant4, GENSER, PI/AIDA, Savannah
- 400 TB of POOL data produced
- ROOT collaboration growing
 - New I/O used by POOL; Common dictionary, maths library with SEAL

This is a key issue that we must keep our eye on

- Conditions DB containing Oracle and MySQL implementations has been released
- Geant4 successfully used in ATLAS and CMS Data Challenges with excellent reliability
- GENSER MC generator library in production (ATLAS, CMS)
- Hadronic physics validation report published
- First round of physics validation studies using Geant4 and FLUKA completed and written up



Applications Area

Concerns

- There must be clear agreement among the experiments in a number of areas before we complete plans for Phase 2
 - Participation of the experiments in POOL and the Conditions Database
 - Long-term strategy of the experiments for the use of frameworks and POOL
 - Which tools will be used for analysis – ROOT, framework-based



CERN fabric

- Fabric management automation (low-cost, reliable operation)
 - in use since beginning of the year at CERN
 - interest in using this elsewhere (GDB coordinating)
- Computer centre preparations and acquisition process at CERN are on track
- Mass storage solutions with a common interface (SRM) being deployed at Tier-1s
- Convergence on a single Linux flavour for LHC and Run II

Concerns

- At CERN - the new release of Castor is ~9 months late
- The unexpected trend towards small files has a major impact on data storage solutions - this requires more investigation as soon as the computing models appear



Grid Deployment

- The grid deployment **process** (LCG-2) is working well –
 - Integration – certification – debugging
 - Distribution - installation

The installation difficulties of last year have been overcome

- Reacting to problems encountered during data challenges, Incremental releases of LCG-2 have led to significant improvements in reliability, performance and scalability
 - within the limits of the current architecture
 - but scalability is much better *than expected a year ago*

Many more nodes and processors than we anticipated at this stage
→ ~80 nodes, ~7,000 processors

- Heavily used during the data challenges in 2004
 - many small sites have contributed to MC productions
 - LHCb has run up to 3,500 concurrent jobs



Grid Deployment

- The **basic issues** of middleware reliability and scalability that we were struggling with at the last review have been overcome

BUT - there are many issues of functionality, usability and performance to be resolved

→ New middleware from EGEE is awaited with impatience

- We are now moving on to tackle the issues of operations coordination and management
 - EGEE support resources now in place
 - Core operations centres → RAL, IN2P3 Lyon, CNAF, ASCC Taipei, CERN
 - Workshop at CERN 2-4 November, with representatives of Grid3 operations



- **But** – we should have been here at the start of the data challenges
 - LCG-1 last year was three months late and then little used
 - Testing really only started as the data challenges began this year
- we are about six months behind the original plan
but with a much larger grid



Middleware from EGEE



- We have a rapidly growing number of sites connecting to the LCG-2 grid -- **but** there are major holes in the functionality, especially in data management
 - EGEE has substantial funding for re-engineering, re-development - and includes key technologists from EDG, VDT, AliEn, Globus
 - The first EGEE prototype was made available in a development environment in May (6 weeks after EGEE started)
 - But there are difficulties in getting the product out
 - the first pieces are only being delivered to deployment this month
 - key components will only arrive next year
 - This leads to pressure to "deploy" the prototype - skipping the test, integration, certification process that has worked well for LCG-2
 - The evolution from LCG-2 to the new middleware can only be planned when we have seen and used the first products
- in the meantime we have to continue to maintain LCG-2, *including short term developments*



Moving Targets

- While there are difficulties in getting the middleware out and in production ..
- .. the requirements of the experiments for “grid” support is in a state of flux, particularly in the area of data management
- As soon as the initial computing models are available (at the end of the year) we have to think hard about
 - what experiments really want “from the grid”
 - What “the grid” can realistically deliver
 - and what experiments want to or will have to implement themselves



Concern – How Many Grids

- There are real pressures in some countries and regions to use similar but different middleware
- But there are also forces for convergence
 - EGEE and LCG-2 are exactly the same – by construction
 - Building compatibility between LCG-2 and Grid-3
 - Collaboration of VDT in EGEE middleware development
 - Joint security group
 - Joint operations discussions at the recent workshop
- The issues for LHC are
 - The costs to experiments and computer centres of supporting and working with multiple interfaces and management structures
 - The hassle for the end user of multiple grids offering similar functionality
 - The increasing complexity of maintaining the notion of an “LHC Grid”
- We are at a crucial stage



ARDA

- The ARDA project aims to help experiments prototype analysis systems using grid technology
 - One prototype for each of the LHC experiments
 - Two people (funded by EGEE and LCG) work closely with each experiment
- Maintain a close relationship with the EGEE middleware team
 - experience with the early versions of the new middleware (gLite)
 - feedback to developers
- ARDA is NOT a middleware or applications development project



Networking

Networking

- Latest estimates are that Tier-1s will need connectivity at ~10 Gbps with ~70 Gbps at CERN
- There is no real problem for the technology
as has been demonstrated by the *Land Speed Records*
- **But** LHC will be one of the few applications needing –
-- this level of performance as a service on a global scale
- There are many different organisations involved in providing high bandwidth networking –
between a Tier-1 centre and CERN
- LCG has to be pro-active in working with service providers
 - Pressing our requirements and our timetable
 - Exercising pilot services
 - Drawing up a “master plan”
 - that must be implemented by regional centres and research networking organisations
- Little is known at present of the networking requirements between Tier-1s,
and in support of Tier-2s



Resources



LCG-2 Support Agreements

VDT (US tools + participation in gLite design effort)

DataGrid resource broker

Replica management tools

DataGrid relational information system (RGMA)

GridIce monitoring tools

GLUE schema

Virtual Organisation Mgt System

- basic services
- management interface (VOMSrs)
- server operation

dCache storage manager

CASTOR storage manager

Operations and Support Centres

Virtual Organisation management for D0

VDT team at Wisconsin (NSF funding)

INFN/CERN

CERN

RAL

INFN

INFN

INFN

FNAL (to be confirmed)

NIKHEF

DESY/FNAL

CNAF/PIC/CERN

RAL, FZK, Academia Sinica Taipei, IN2P3 Lyon, CNAF

NIKHEF

These are mainly informal agreements without entries in the project milestone/WBS tables

The costs are not estimated, but they do represent substantial investments



Resources

Tier-1 Planning

- First round of collecting plans for resources completed (see Kors Bos's talk tomorrow)
 - Overall - looks like a reasonable fit requirements/planning for processors and tape storage
 - But there is a significant shortfall in disk storage
 - And sharing between experiments is not clear
 - We need to wait for the publication of initial computing models and for the outcome of the LHCC resource review in January – to understand how much of a problem this is

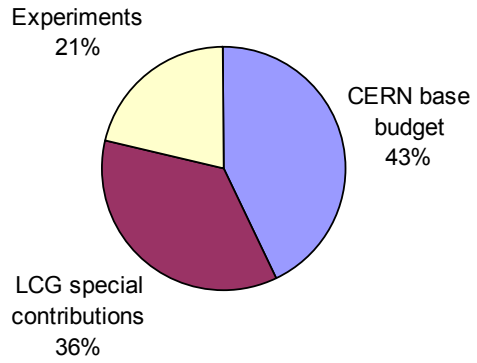
Resources at CERN for Phase 2

- Human resource plan reviewed within the original profile
 - see [CERN-C-RRB-2004-114](#)
- Shortfall in funding at CERN has been taken up by the CERN management through the C-RRB, SPC and Council
- We are proceeding on the assumption that a solution will be found during the next few months
- If this happens - the materials funding is OK within the error bars (again subject to the January review)

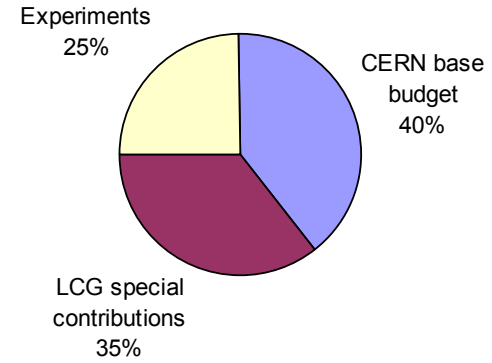
Resources from Experiments for Applications Area projects

- The resources assumed in the plan are in principle agreed with the experiments - but there is a lack of firm commitments

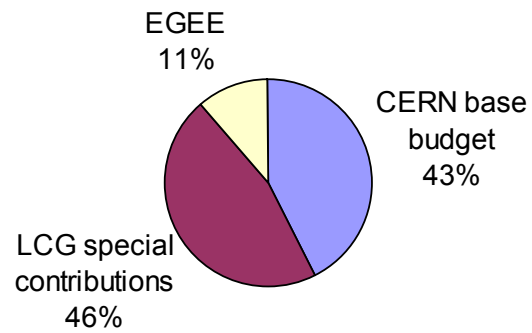
FTEs in the Applications Area - Oct-04



FTEs in the Applications Area - Oct-03



FTEs in Other Areas - Oct-04



Total FTE-years since project started

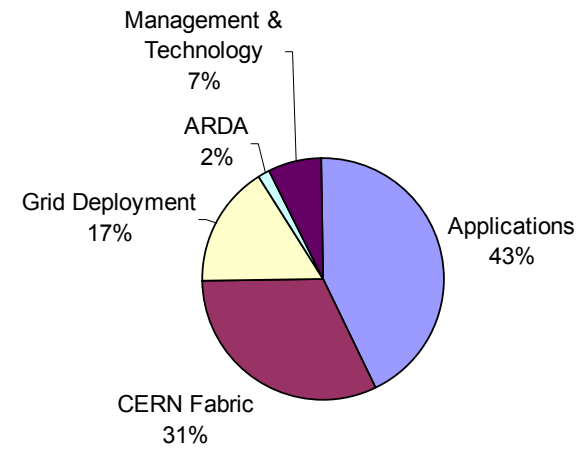


Chart 1: Number of LCG Level 1&2 Milestones (cumulative)

