



LHC Computing Grid Project - LCG

Management, Resources, Progress

LHCC Comprehensive Review

24 November 2003

Les Robertson – LCG Project Leader
CERN – European Organization for Nuclear Research
Geneva, Switzerland

les.robertson@cern.ch



Management



Management Team

Project Leader

Resource Manager – *Chris Eck*
Planning Officer – *Jürgen Knobloch*

Chief Technology Officer

David Foster

Overall coherence of the project
Pro-active technology watch
Long-term grid technology strategy

Applications Area

Torre Wenaus

Development environment
Joint projects, Data management
Distributed analysis

Middleware Area

Frédéric Hemmer

Provision of a base set of grid middleware
(acquisition, development, integration)
Testing, maintenance, support

CERN Fabric Area

Bernd Panzer

Large cluster management
Data recording, Cluster technology
Networking, Computing service at CERN

Grid Deployment Area

Ian Bird

*Establishing and managing the **Grid Service***
- Middleware, certification, security
operations, registration, authorisation,
accounting



Project Execution Board

*Chair – Project Leader
Scientific Secretary*

*Area Managers
CTO*

*Experiment
Computing
Coordinators*

*Chairs of the
-- GDB
-- SC2 (invited)*

Architects' Forum
Applications Area Manager
Experiment Architects


Planning, management,
architectural and technical
direction of the applications
area of the project

Grid Deployment Board
The grid “collaboration board” of the
regional centres and experiments

Members: national delegates,
experiment computing coordinators
and production managers



Management Tools

- WBS - (MS Project) - formal agreement on milestones
 - Level 1 - agreed with LHCC referees
 - Level 2 - agreed on quarterly basis with SC2
- RBS - (Excel workbook) - human resource history
 - Allocation to activities
 - Funding source
 - FTE-month granularity
- Risk analysis - using process developed by GridPP 
- Quarterly status reports → SC2 milestone analysis
- External auditing in CERN Fabric area
 - Computer centre physical infrastructure
 - Service process (SUN)
 - Total cost of ownership analysis (Openlab industrial partners)

Risk Register for the LCG Project — 8th August 2003

- Likelihood:**
- 1 - never expected to happen
 - 2 - could happen but very unlikely
 - 3 - could well happen at some point
 - 4 - will probably happen
- Impact:**
- 1 - we can deal with it, no problem
 - 2 - a bit of a hassle but not too bad
 - 3 - clearly can be dealt with, but with significant effort
 - 4 - crisis

Name	LCG		Funding		Apps		Fabric		Deploy.		GTA				
	Li	Im	Risk	Li	Im	Risk	Li	Im	Risk	Li	Im	Risk			
Operational Issues															
Recruitment/Retention difficulties	1	2	2	2	4	3	3	9	1	2	2	3	9		
Sudden loss of key staff	2	3	6	2	3	6	2	3	6	2	3	6	2	3	6
Sub-components not delivered to area	2	2	4	2	3	6	2	4	8						
Duplication of Effort	2	1	2	4	1	4	1	2	2						
Time and Budgets															
Unrealistic timeframe/schedule	3	2	6						3	3	9	3	3	9	
MC timeframe extended	3	1	3	3	1	3									
Lack of funding to meet phase 1 goals	2	3	6	2	3	6	2	4	8						
Funding agencies do not provide agreed resources	1	3	3	1	3	3									
Offered in kind contributions do not fit requirements	2	2	4	2	2	4									
Performance and Analysis															
Change in project scope	2	1	2	2	2	4	2	2	4			2	2	4	
Conflicting software requirements	2	2	4	2	3	6	2	3	6			2	3	6	
Inadequate analysis/specification	2	3	6	2	2	4	2	4	8			2	4	8	
Requirements/Functionality mismatch	2	3	6	2	2	4	2	3	6	2	4	8	1	3	3
Issues with software components															
Inadequate or late 3rd party software	3	3	9	3	2	6	1	1	1	4	4	16	2	4	8
3rd party software no longer available or not maintainable	2	2	4	3	2	6	1	1	1						
Software maintainability problems	2	3	6	1	2	2				2	4	8	2	3	6
Licensing Limitations	3	2	6	3	4	12									
Deployment/integration failure	1	2	2	1	3	3									
Resource shortfalls															
Reliability problems	3	3	9							3	3	9	3	4	12
Availability problems	2	3	6							4	3	12	3	4	12

LCG Risk Analysis - 30 Oct 2003 (v2)

ID	Description	Risk Analysis				
		Likelihood		Impact		Compu
		Avg	Stdev	Avg	Stdev	Risk
Security Issues - Intentional or malicious attacks						
Misuse of LCG resources - CPU, storage, network etc						
M1	Resources used to launch online attacks on other sites via DOS, Virus, Worms, SPAM etc	3.0	0.0	3.0	0.0	9.0
M2	Resources used for offline attacks on other sites, e.g. to crack passwords or pass phrases	2.0	0.0	2.0	0.0	4.0
M3	Resources used to distribute or share non-LCG data, e.g. copyrighted, illegal, or inappropriate material	3.0	0.0	3.0	0.0	9.0
M4	Resources misused by inappropriate setting of access control or priority	3.0	0.0	1.0	0.0	3.0
M5	Use of LCG resources by unauthorized parties	3.0	0.0	1.0	0.0	3.0
M6	Use of LCG resources for unauthorized purposes, e.g. financial gain	2.0	0.0	2.8	0.5	5.5
Confidentiality and Data integrity issues						
C1	Theft of credentials, e.g. private keys	3.0	0.0	2.0	0.0	6.0
C2	Data or passwords/pass phrases exposed, e.g. in unprotected files or on the network	3.0	0.0	2.0	0.0	6.0
C3	Falsification of scientific data, analysis and/or results	1.8	0.5	3.0	0.0	5.3
C4	Unauthorized monitoring of network communications	2.0	0.0	2.0	0.0	4.0
C5	Unauthorized access to data	3.0	0.0	1.0	0.0	3.0
C6	Unauthorized distribution or exposure of data	2.8	0.5	2.0	0.0	5.5
C8	Identity or usage information is harvested by unauthorized persons	2.0	0.0	1.0	0.0	2.0
Disruption of LCG infrastructure for political or other reasons						
D1	Disruption via exploitation of security holes	3.0	0.0	3.0	0.0	9.0
D2	Corruption of or damage to data	1.0	0.0	2.8	0.5	2.8
D3	DOS attacks towards LCG to prevent normal working of network or services	1.8	0.5	3.0	0.0	5.3
D5	"Poisoned" resources are deployed on LCG to confuse operations, debugging or results	1.0	0.0	2.8	0.5	2.8



Management Tools

- WBS - (MS Project) - formal agreement on milestones
 - Level 1 - agreed with LHCC referees
 - Level 2 - agreed on quarterly basis with SC2
- RBS - (Excel workbook) - human resource history
 - Allocation to activities
 - Funding source
 - FTE-month granularity
- Risk analysis - using process developed by GridPP
- Quarterly status reports - SC2 milestone analysis
- External auditing in CERN Fabric area
 - Computer centre physical infrastructure
 - Service process (SUN)
 - Total cost of ownership analysis (Openlab industrial partners)





Resources



LCG Project - Human Resources

unweighted FTEs in October-03

Applications

Software Process Infrastructure	6.0
Object Persistency (POOL)	12.3
Core Libraries and Services (SEAL)	5.4
Physics Interfaces (PI)	5.1
Simulation	16.0
GRID interfacing	4.2
Architecture	0.2
Management	1.7
Total	50.9
ROOT	5.7

Includes resources at CERN and at Other Institutes



CG Project - Human Resources

unweighted FTEs in October-03

CERN Fabric

System Management & Operations	10.4
Development (e.g. Monitoring)	9.6
Data Storage Management	7.5
Grid Security	4.0
Grid-Fabric Interface	8.0
Internal Networking for Physics	1.0
External Networking for Physics	0.6
Management	1.3
Total	42.4

Project Technology

Grid Technology, modelling, evaluations

Total 5.0

Grid Deployment

Integration and Certification 14.0

Grid Infrastructure, Operations and User Support 10.0

Total 24.0

LCG Management 5.0

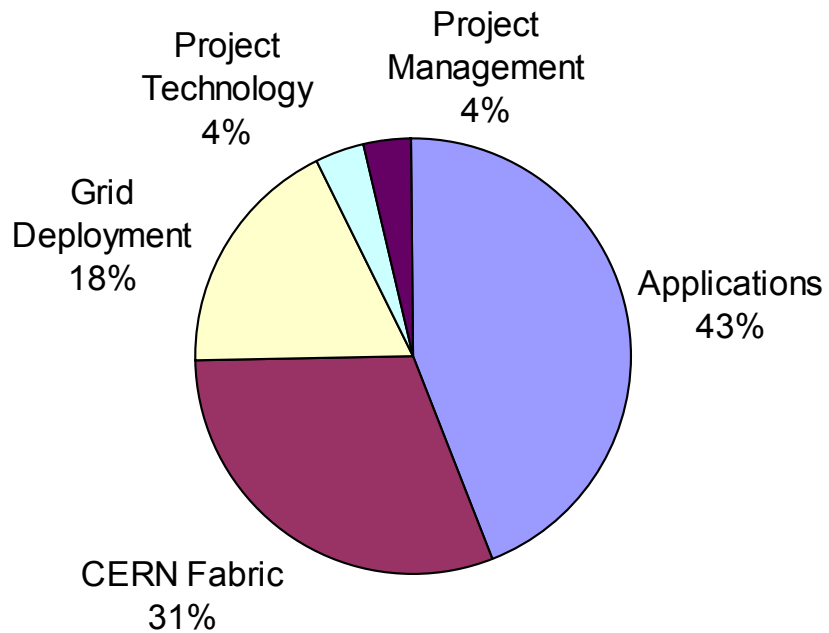
*includes all staff for LHC services at CERN
Does NOT include EDG middleware*



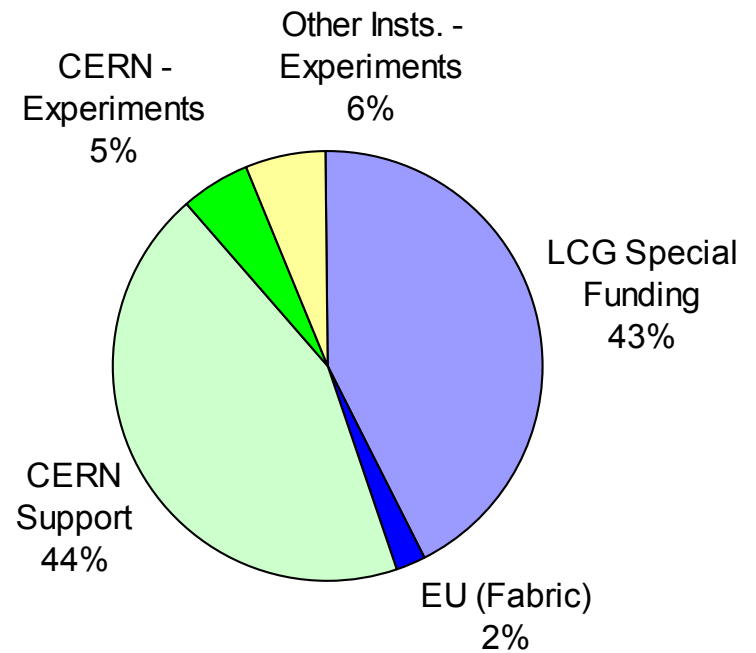
Personnel resources

Includes –

- all staff for LHC services at CERN (inc. networking)
- Staff at other institutes in applications area



Current Staff (FTEs) by Project Area



Current Staff (FTEs) by Funding Source

without Regional Centres,
EDG middleware



Human Resources used to October 2003

CERN + apps. at other institutes

LCG Project - Human Resources

Experience-weighted FTE-years

Resources used (FTE-years)

	2002	1Q03	2Q03	3Q03	Total
Applications (with ROOT)	37.7	13.3	15.6	15.7	82.3
CERN Fabric	32.2	10.4	10.3	10.7	63.5
Grid Deployment	7.5	4.1	4.6	4.9	21.1
Project Technology	1.2	0.3	0.8	1.4	3.6
LCG Management	5.9	1.8	1.7	1.6	11.0
Total	84.5	29.9	33.0	34.2	181.6



Human Resources at Other Centres

Resources reported to C-RRB

- Formal projects with India, Russia (protocols signed with CERN)
 - Review of projects with India foreseen for early next year

Essential part of the project, but only partially reported to C-RRB

- Regional Centres
 - Grid operations centre - RAL
 - Grid call centres - FZK
 - Significant human resources for operating services in Regional Centres
- These resources are provided as part of a collaboration
 - No formal MoU - non-binding agreements made in the Grid Deployment Board
 - Optional reporting of human resources
 - Moving towards common planning, reporting of computing resources

Contributing to the project, but not managed by LCG

- Grid technology developed by other projects
- EGEE will provide funding for operations, middleware at many centres in Europe

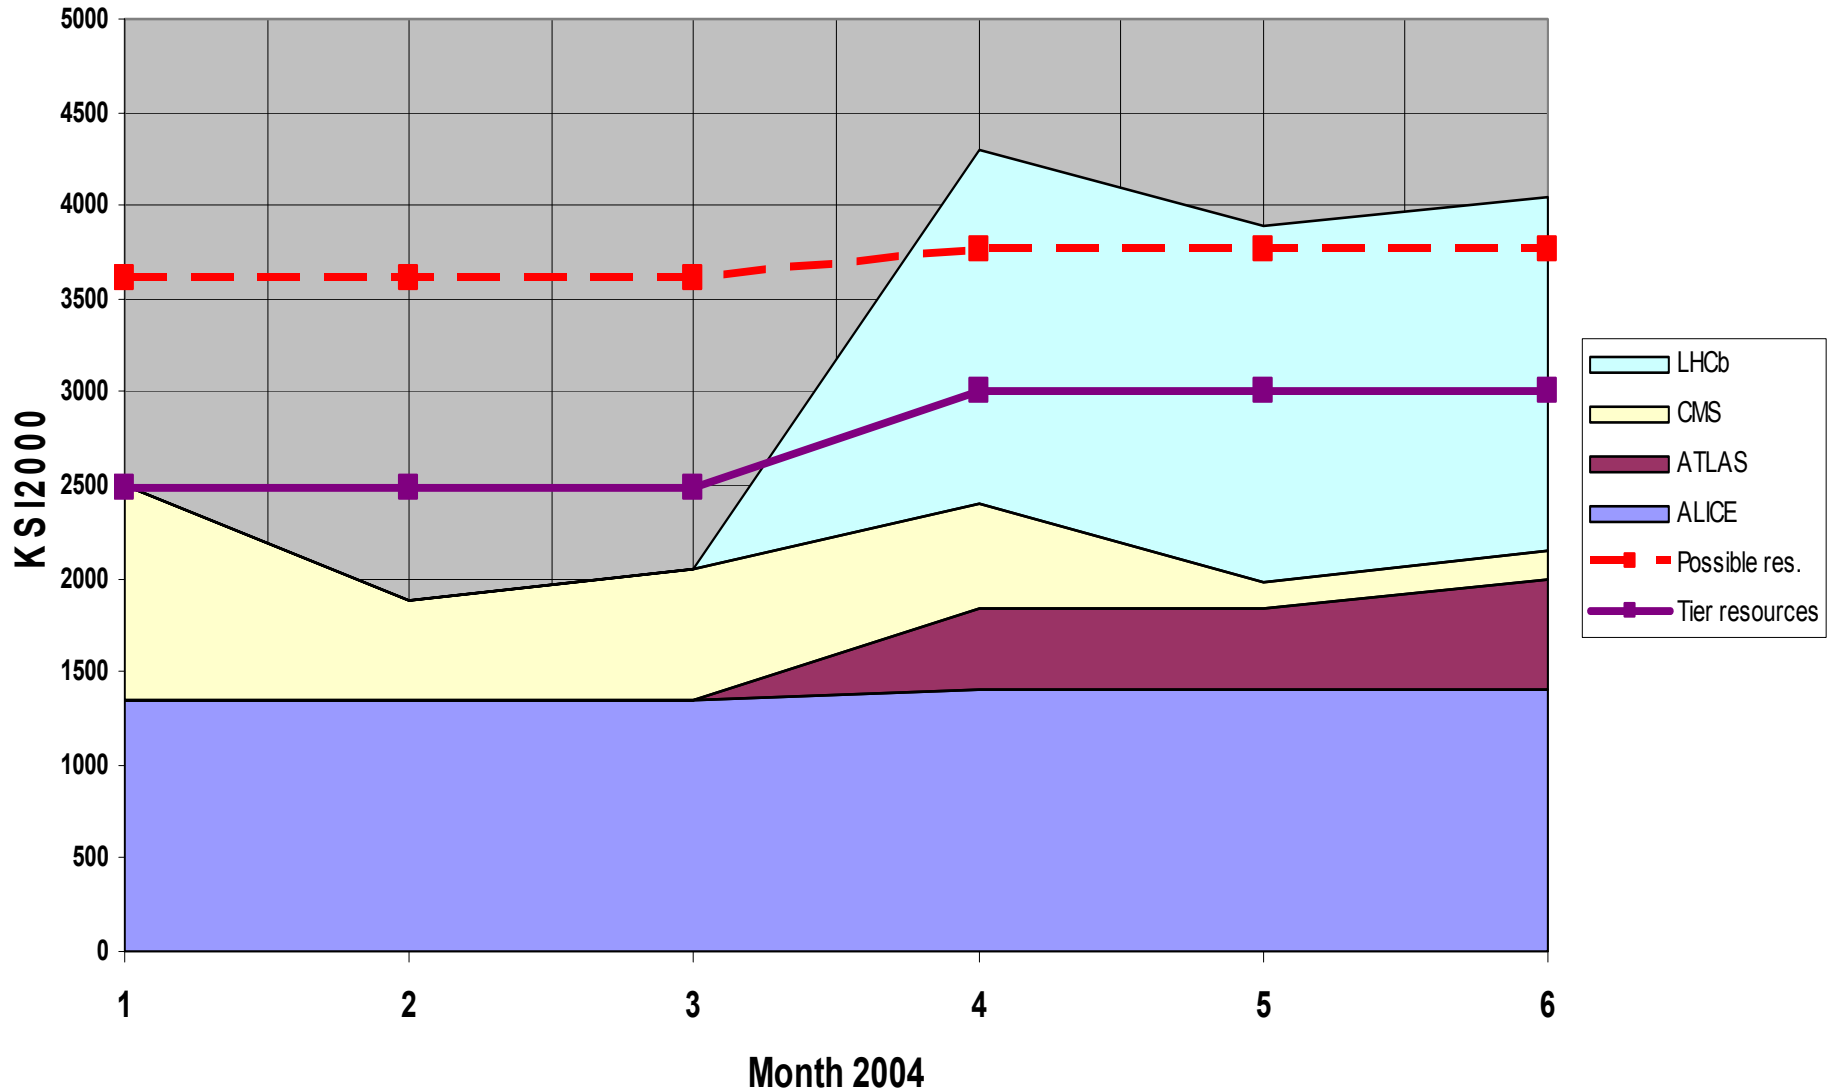


LCG Regional Centre Capacity 2Q2004

Country	CPU Capacity kSI2k	Disk Capacity TB	LCG Support FTE	Online Tapes TB	Diff. to CPU Capacity announced May 200303
CERN	700	100	10.0	1000	0
Czech Rep.	30	6	1.0	2	-30
France	150	24	1.0	160	-270
Germany	305	44	9.3	74	98
Holland	30	1	3.0	20	-94
Hungary					-70
Italy	895	110	13.3	100	388
Japan	125	40		50	-95
Poland	48	5	1.8	5	-38
Russia	102	12.9	4.0	26	-18
Taipei	40	5.8	4.6	30	-180
UK	486	55.3	3.4	100	-1170
USA ¹	80	25		100	-721
Switzerland	18	4	1.4	20	-8
Spain	150	30	4.0	100	0
Sweden					-179
Sum	3159	463	57	1787	-2388



Resources and Requests






Progress



Milestone Performance



- New milestone process introduced in July
- SC2 monitors Level 1 and Level 2 milestones
- Baseline for level 2 milestones agreed in July
- Additional verification milestones defined - implemented by experiments - to measure experiment acceptance of LCG deliverables
- Currently 188 milestones defined
- 90% of the milestones due by October have been met



LCG Level 1 Milestones


applications

General release of POOL hybrid data store

Distributed production environment using grid service

Distributed end-user interactive analysis

Full Persistency Framework



launch workshop

LHC Global Grid TDR

“50% prototype” (LCG-3) available

grid service

LCG-1 reliability and performance targets

First Global Grid Service (LCG-1) available



LCG Level 1 Milestones


applications

General release of POOL hybrid data store

Distributed production environment using grid service

Distributed end-user interactive analysis

Full Persistency Framework



launch workshop

LHC Global Grid TDR

"50% prototype" (LCG-3) available

grid service

LCG-1 reliability and performance targets

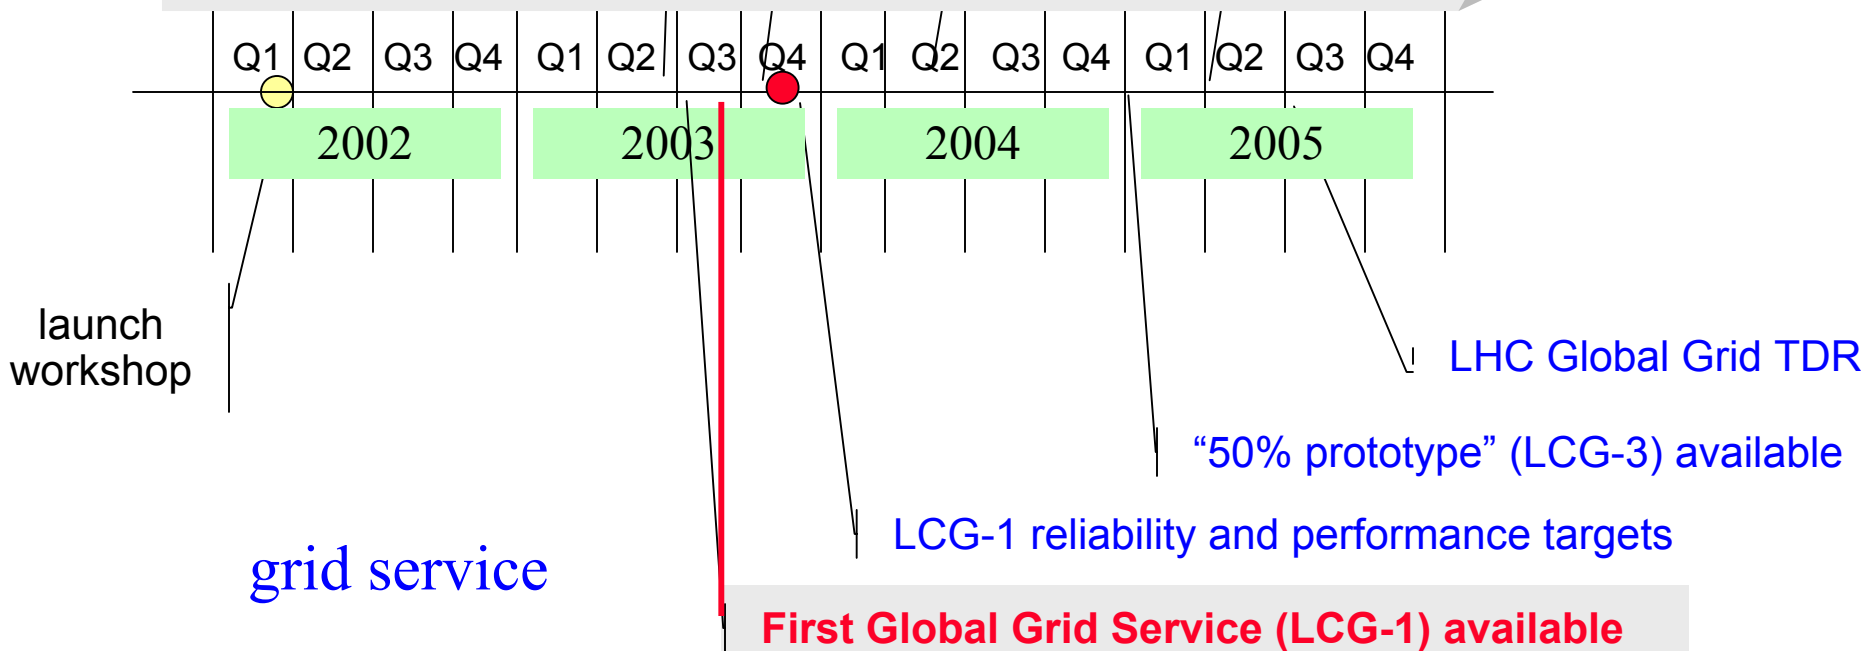
First Global Grid Service (LCG-1) available



- Due on 1 July
 - Delays in delivery of middleware
 - Milestone achieved 15 September
 - Lower functionality than planned
- Testing/usage by experiments only starting now
(planned for end August)

g grid service
analysis
framework

2.5 month delay






Strategy for getting back on track –

the LCG-2 availability date will be maintained – December 2003
-- this is the service for the data challenges in 2004

However –

- decision on middleware for LCG-2 taken without experience of production running
- reduced time for integrating and testing the service with experiments' systems before data challenges start next spring
- additional functionality will have to be integrated later






LCG Level 1 Milestones

The level-1 milestone for verification of the LCG-1 service

- 30-day service stability
- Requires production use by experiments
- Due end November
 - LCG-1 slip pushes this back to mid-February

BUT - will now be verification of the **LCG-2** service

service
sis
work





Suspended Level 1 Milestones

- Two Applications Area Level-1 milestones -
 - Distributed production environment (November 2003)
 - Distributed analysis environment (May 2004)

Have been suspended waiting for specification of requirements by SC2

Requirements defined on 7 November 2003

- grid & distributed analysis requirements
- *HEPCAL 2, ARDA*



LCG Level 1 Milestones


applications

General release of POOL hybrid data store

Distributed production environment using grid service

Distributed end-user interactive analysis

Full Persistency Framework



launch workshop

LHC Global Grid TDR

“50% prototype” (LCG-3) available

grid service

LCG-2 reliability and performance targets

First Global Grid Service (LCG-1) available



Level 1 Milestones Summary

- POOL released on time
- First LCG service 2-3 months late
- Recovery strategy
 - maintain the timetable for the second release
 - → essential to serve the data challenges of 2004
- Two distributed environment milestones invalidated by lack of requirements

- The Level-1 milestones no longer really pace the project
 - Better understanding of the reality of the LCG service
 - Plans being formulated to prepare for computer model testing (2004) and validation (2005)
 - 2nd generation middleware planning
- Aim for revised Level 1 milestones for discussion with the referees in March



LCG Service Time-line

computing service

physics

open LCG-1 (achieved) – 15 Sept

LCG-2 - upgraded middleware, mgt. and ops tools


second generation middleware prototyping, development

LCG-3 – second generation middleware

TDR for the Phase 2 grid

Phase 2 service acquisition, installation, commissioning

Phase 2 service in production



Testing, with simulated event productions

principal service for LHC data challenges

Computing models

validation of computing models

experiment setup & preparation





Project Evolution

- Scope (formal requirements) of project now largely defined
- First applications products delivered
- Grid service opened
- Focus for next year:
 - Take-up by experiments of applications products, grid service
 - Development of computing models
 - Second round of grid technology - for distributed analysis
- Long-term planning for applications area -
 - Development
 - Long term support
 - Staffing - balance of experiments, institutes, CERN
- Re-organisation of the management committees



Issues



Collaborating Regional Centres

- Concern about staffing and priorities in Regional Centres during the initial service roll-out
- The close collaboration needed to establish and operate a grid is new to computing centres
- Fundamental work still to be done on operations-related tasks - that will require active and close collaboration between centres
- This is a major challenge for the Grid Deployment Board
 - Effective collaboration : agreements → commitments
 - Participation of technical service management



Grid Technology

- Reliability and Scalability are still major concerns
- Two major flavours - AliEn (ALICE) and Globus
- At least three versions of the Globus stack
 - Globus+Nordugrid
 - Globus+Condor+.. → VDT
 - VDT+EDG (resource broker, data management) → LCG
- Longer term - essential to live with multiple grids, different middleware
 - Standards? Gateways? ..
- May be premature to tackle this problem before we have one reliable, scalable grid



Grids - Funding - Goals

- **Grid technology & services**
 - A potential solution for LHC
 - An opportunity for working closely with computer scientists, other sciences
- **Funding**
 - Funding agencies like grids
 - HEP is seen as a ground breaker, risk taker
 - LHC is seen as an ideal demonstrator application
→ new (non particle physics) funding
- **Leading edge** → **Bleeding edge**
- **Must keep our eye on the goal of LHC data handling**
 - **Simplify** rather than complicate
 - Grid as a continuous **service** – not just data challenges



Computing Models

- Computing models of the experiments must be known by end of 2004
 - so that the LCG TDR can be prepared by mid-2005
 - so that the acquisition processes of the Tier 0, Tier 1 and large Tier 2 centres can complete by mid 2006
- These models will depend on the **reality** of the technology
 - Target functionality that looks feasible for 2007-08
 - A baseline model that we can be sure of
- We are late in getting analysis experience with grids
- Essential to organise **joint LCG-experiment tests early in 2004**
 - Tier 0+1+2 - batch - ESD analysis - production
 - Tier 1+2+3 - end-user analysis



Experiments and the Applications Area

- Increased emphasis on integration of common developments in experiment applications
- Participation of experiments in LCG common activities
 - Relative priorities of common projects and experiment-private activities
 - Establishing the "value" to an experiment of common tools
 - Long-term commitment - support and maintenance
- The full scope of the applications area can now be established (distributed analysis, conditions database, event level metadata, simulation)
 - Long term responsibilities and resource plan in 1Q04



Phase 2



Phase 2 Preparations

- LCG authorised only for Phase 1
 - To end 2005
 - Applications environment development
 - Support and maintenance in Phase 2?
 - Blueprint for the Phase 2 Tier 0/1/2 service
 - Proof of concept for this blueprint
 - Preparation of CERN computing centre for Phase 2
 - Forward resource planning at CERN
- Last two points assumed also to be under way at all regional centres




Transition to Phase II

- Have not agreed yet with experiments on applications support model, staffing profile
- Tier 0/1 facility at CERN
 - Infrastructure being upgraded
 - Equipment acquisition for Phase II starting now
- Phase II CERN Fabric and Grid staffing requirements estimated
 - ~20M CHF more than funding in the Medium Term Plan (MTP) Of which ~16M is missing manpower
No change from situation reported to Council in April 2002
- Special funding for staff falls off in 2005 - too soon
- MoU for Phase II to be developed over next 12 months



Phase 1+2 - Human Resources at CERN


LCG Human Resources Budget & Funding Estimates





Phase II Human Resource Strategy

- Establish Applications Area scope, estimate human resource requirements, develop staffing plan
- Phase 2 MoU -
 - Task force being convened (mandate to POB on Thursday)
- Strategy for finding missing resources at CERN
 - Subject to policies of new CERN management
 - Explore further EU sources (not very optimistic)
 - Explore further special contributions (March C-RRB, after MoU drafted)
- Re-examine scope of Phase 2, reduce computing capacity at CERN





Back pocket foils



Key Progress - Applications

- First POOL public release - and integration in mainline applications of CMS, ATLAS
- Integrated simulation project
 - Physics validation
 - GEANT 4, FLUKA, MC generators
 - Generic simulation framework
- Increasing participation of external institutes



Key Progress - Fabric

- Deployment of main components of Quattor automated management system (from EDG)
- CERN Computer Centre upgrade on track
- Tape service upgrade, and 1 GByte/sec demonstration
- CASTOR development plan agreed

- Phase 2 acquisition process starting



Key Progress - Grid Technology

- Starter set of grid technology agreed (February 2003) - components from -
 - European DataGrid (EDG)
 - US (Globus, Condor, PPDG, GriPhyN) → the Virtual Data Toolkit
- *Grid technology review*
 - Gap analysis - from UK e-Science
 - OGSA engineering activity to gain experience in grid web services - preparation for 2nd generation toolkits
- Distributed analysis requirements specified - HEPCAL 2 and ARDA
- New EU project established -
 - EGEE - Enabling Grids for e-Science in Europe
 - Funding for re-engineered middleware



Key Progress - LCG Service

Certification and distribution process established

Expert debugging team set up - established effective relations with middleware suppliers on both sides of the Atlantic

Agreement reached on security policies, principles for registration and security

Rutherford Lab (UK) to provide the initial Grid Operations Centre

FZK (Karlsruhe) to operate the Call Centre

The initial service was opened on
4 September - 12 Centres

- now 24 sites active

Experiments have started testing

- but not yet in production

LCG-2 service defined





Requirements for Distributed Analysis

- Formal requirements - delivered last week!
 - HEP Common Applications Layer - HEPCAL II
 - ARDA - blueprint for distributed analysis
 - Analysis of basic services
 - Web Service architecture
 - Emphasis on initial ~6 month prototyping phase
- Project management preparing a skeleton work-plan
- Have to understand the responsibilities of
 - Generic middleware
 - Common HEP applications
 - Experiment-specific
- Must not miss the opportunity to bring together AliEn, EDG, US experience



LCG-2 and ARDA

- Important that the ARDA prototyping involves real users, real applications, all experiments
- There has been a lot of investment in stabilising Globus/VDT/EDG middleware - the components of LCG-2
- LCG-2 will be the main service for the 2004 data challenges for the large experiments
- This will provide essential experience on operating and managing a global grid service
- The ARDA post-prototype implementation will have to catch up with this experience before it can replace LCG-2



EGEE-LCG Relationship

Enabling Grids for e-Science in Europe - EGEE

- EU project approved to provide partial funding for operation of a general e-Science grid in Europe, including the supply of suitable middleware
- EGEE provides funding for 70 partners, large majority of which have strong HEP ties

Agreement between LCG and EGEE management on very close integration

OPERATIONS

- LCG operates the EGEE infrastructure as a service to EGEE
 - ensures compatibility between the LCG and EGEE grids
- In practice - the EGEE grid will *grow out of* LCG
- The LCG Grid Deployment Manager (Ian Bird) serves also as the EGEE Operations Manager



EGEE-LCG Relationship (ii)

MIDDLEWARE

- The EGEE middleware activity provides a middleware package
 - satisfying requirements agreed with LCG (..HEPCAL, ARDA, ..)
 - and equivalent requirements from other sciences
- **Middleware - the tools that provide functions -**
 - that are of general application ..
 - not HEP-special or experiment-special
 - and that we can reasonably expect to come in the long term from public or commercial sources (cf internet protocols, unix, html)
- **Very tight delivery timescale dictated by LCG requirements**
 - Start with LCG-2 middleware
 - Rapid prototyping of a new round of middleware. First "production" version in service by end 2004
- **The EGEE Middleware Manager (Frédéric Hemmer) serves also as the LCG Middleware Manager**