LCG - The Worldwide LHC Computing Grid

LHC Data Analysis

Challenges for 100 Computing Centres in 20 Countries

HEPiX Meeting
Rome
5 April 2006

Les Robertson
LCG Project Leader

The Worldwide LHC Computing Grid

- Purpose
  - Develop, build and maintain a distributed computing environment for the storage and analysis of data from the four LHC experiments
    - Ensure the computing service
    - … and common application libraries and tools
- Phase I – 2002-05 - Development & planning
- Phase II – 2006-2008 – Deployment & commissioning of the initial services

WLCG Collaboration

- The Collaboration
  - ~100 computing centres
  - 12 large centres (Tier-0, Tier-1)
  - 38 *federations* of smaller "Tier-2" centres
  - 20 countries
- Memorandum of Understanding
  - Agreed in October 2005, now being signed
- Resources
  - Commitment made each October for the coming year
  - 5-year forward look

les.robertson@cern.ch

Worldwide LCG Organisation

Collaboration Board –
chair Neil Geddes (RAL)

Overview Board –
chair Jos Engelen (CERN CSO)

Committee of the Collaboration Board
oversee the project
resolve conflicts

One person from Tier-0, Tier-1s
Experiment spokespersons

s (NIKHEF)

Architects

Experiment software architects
Applications Area Manager
Applications Area project managers

Coordinators
t representatives
nagers

Software

les.robertson@cern.ch

More information on the collaboration

Boards and Committees

All boards except the OB have open access to agendas, minutes, documents

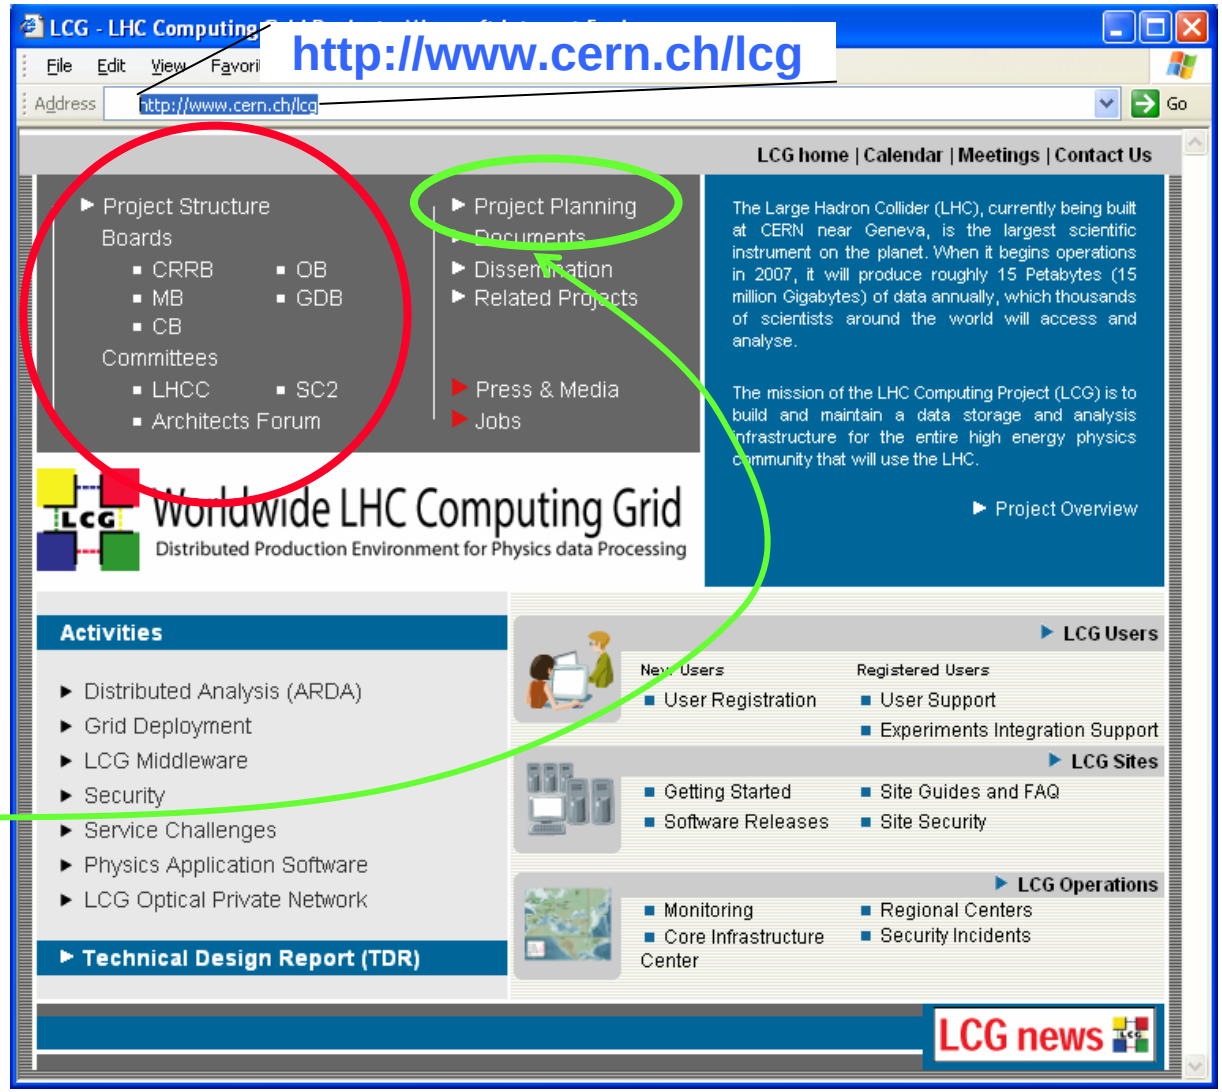

Planning data:

MoU Documents and Resource Data

Technical Design Reports

Phase 2 Plans

Status and Progress Reports

Phase 2 Resources and costs at CERN

les.robertson@cern.ch

LCG Service Hierarchy

Tier-0 – the accelerator centre

- Data acquisition & initial processing
- Long-term data curation
- Distribution of data → Tier-1 centres

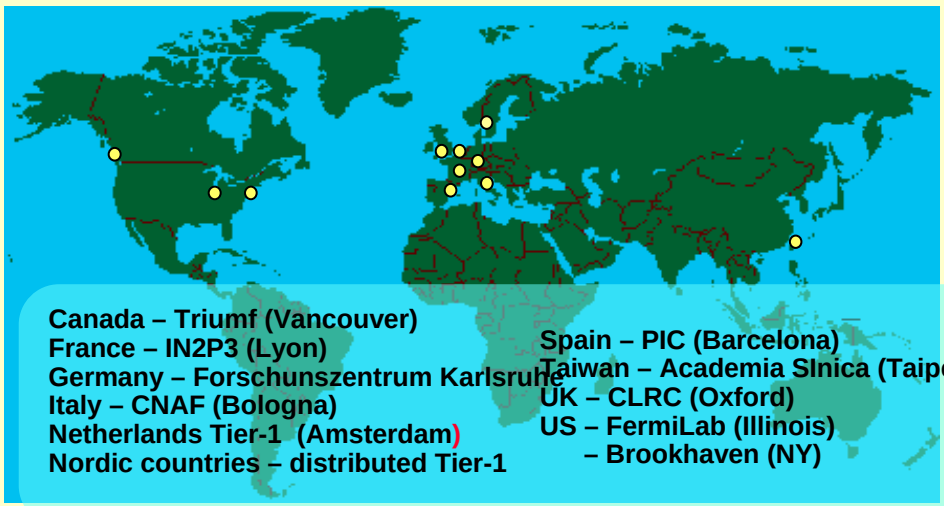

Canada – Triumf (Vancouver)
France – IN2P3 (Lyon)
Germany – Forschunszentrum Karlsruhe
Italy – CNAF (Bologna)
Netherlands Tier-1 (Amsterdam)
Nordic countries – distributed Tier-1
Spain – PIC (Barcelona)
Taiwan – Academia SInica (Taipei)
UK – CLRC (Oxford)
US – FermiLab (Illinois)
– Brookhaven (NY)

Tier-1 – "online" to the data acquisition process → high availability

- Managed Mass Storage – → grid-enabled data service
- Data-heavy analysis
- National, regional support

Tier-2 – ~100 centres in ~40 countries

- **Simulation**
- **End-user analysis – batch and interactive**

les.robertson@cern.ch

Summary of Computing Resource Requirements

All experiments - 2008

From LCG TDR - June 2005

	CERN	*All Tier-1s*	*All Tier-2s*	*Total*
CPU (MSPECint2000s)	**25**	**56**	**61**	**142**
Disk (PetaBytes)	**7**	**31**	**19**	**57**
Tape (PetaBytes)	**18**	**35**		**53**

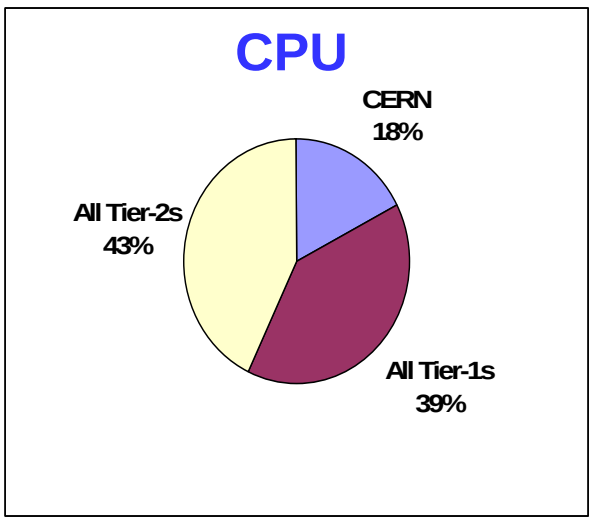

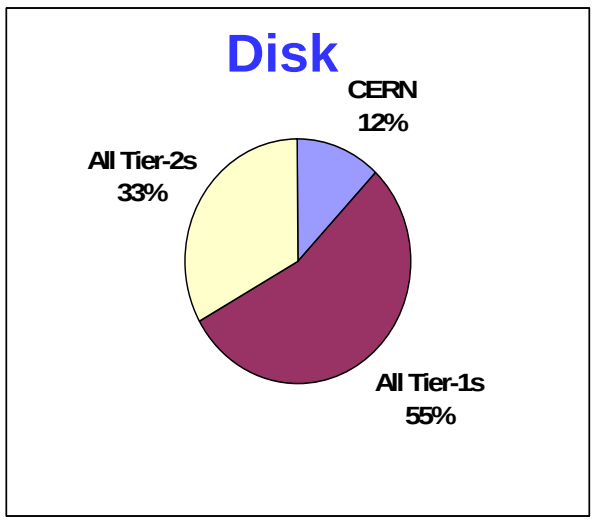

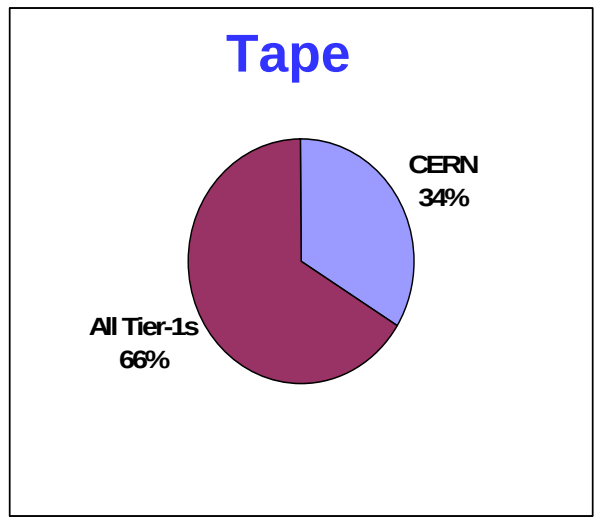

les.robertson@cern.ch

LCG depends on two major science grid infrastructures ….

EGEE - Enabling Grids for E-Science

OSG - US Open Science Grid

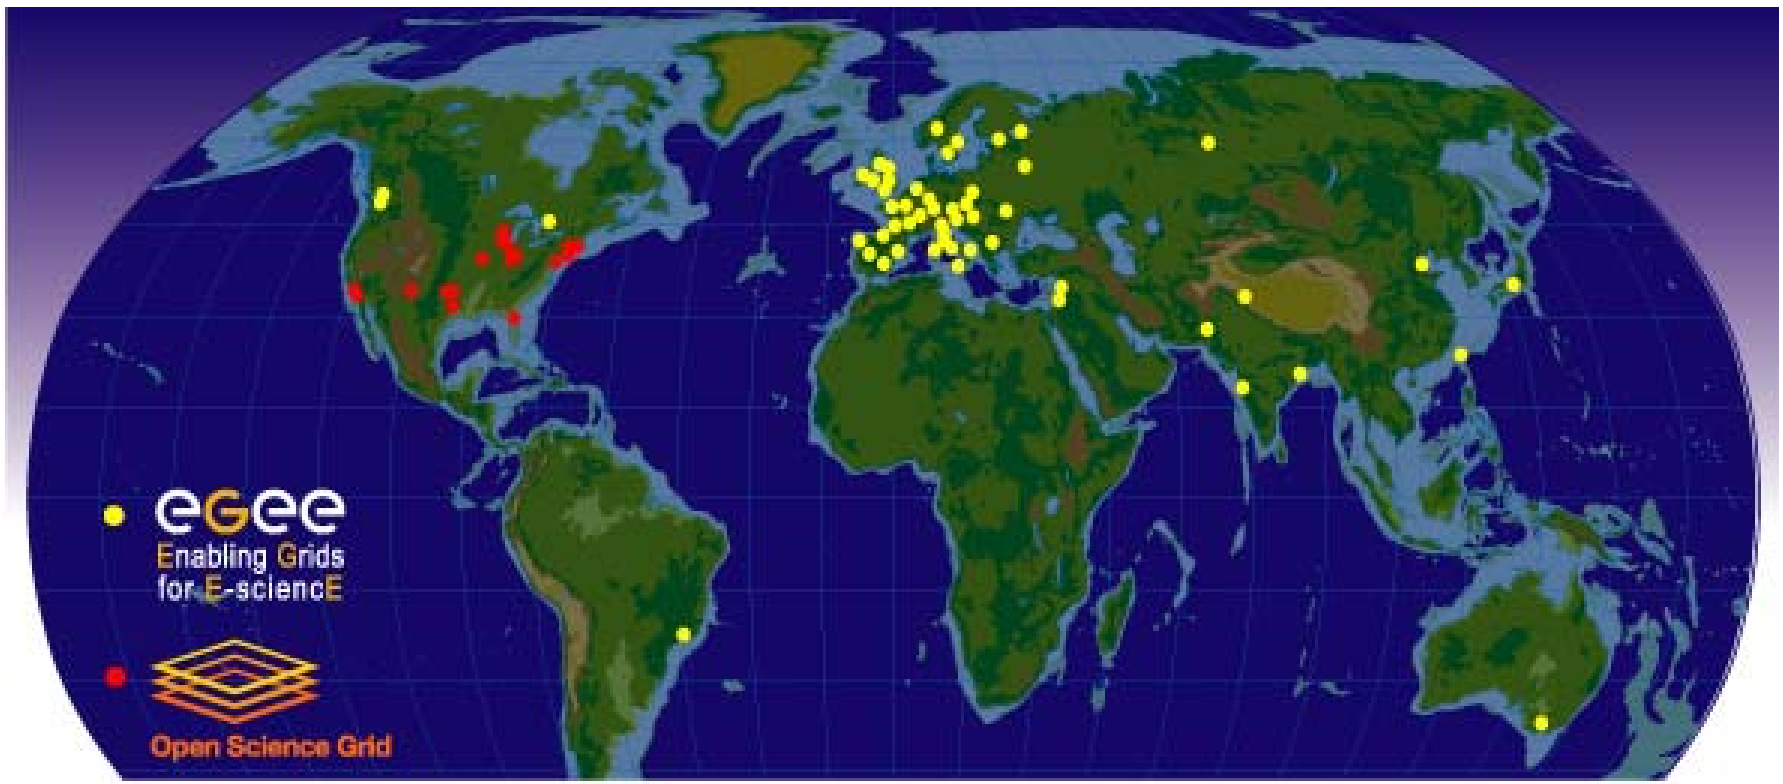

A map of the worldwide LCG infrastructure operated by EGEE and OSG.

.. and an excellent Wide Area Network

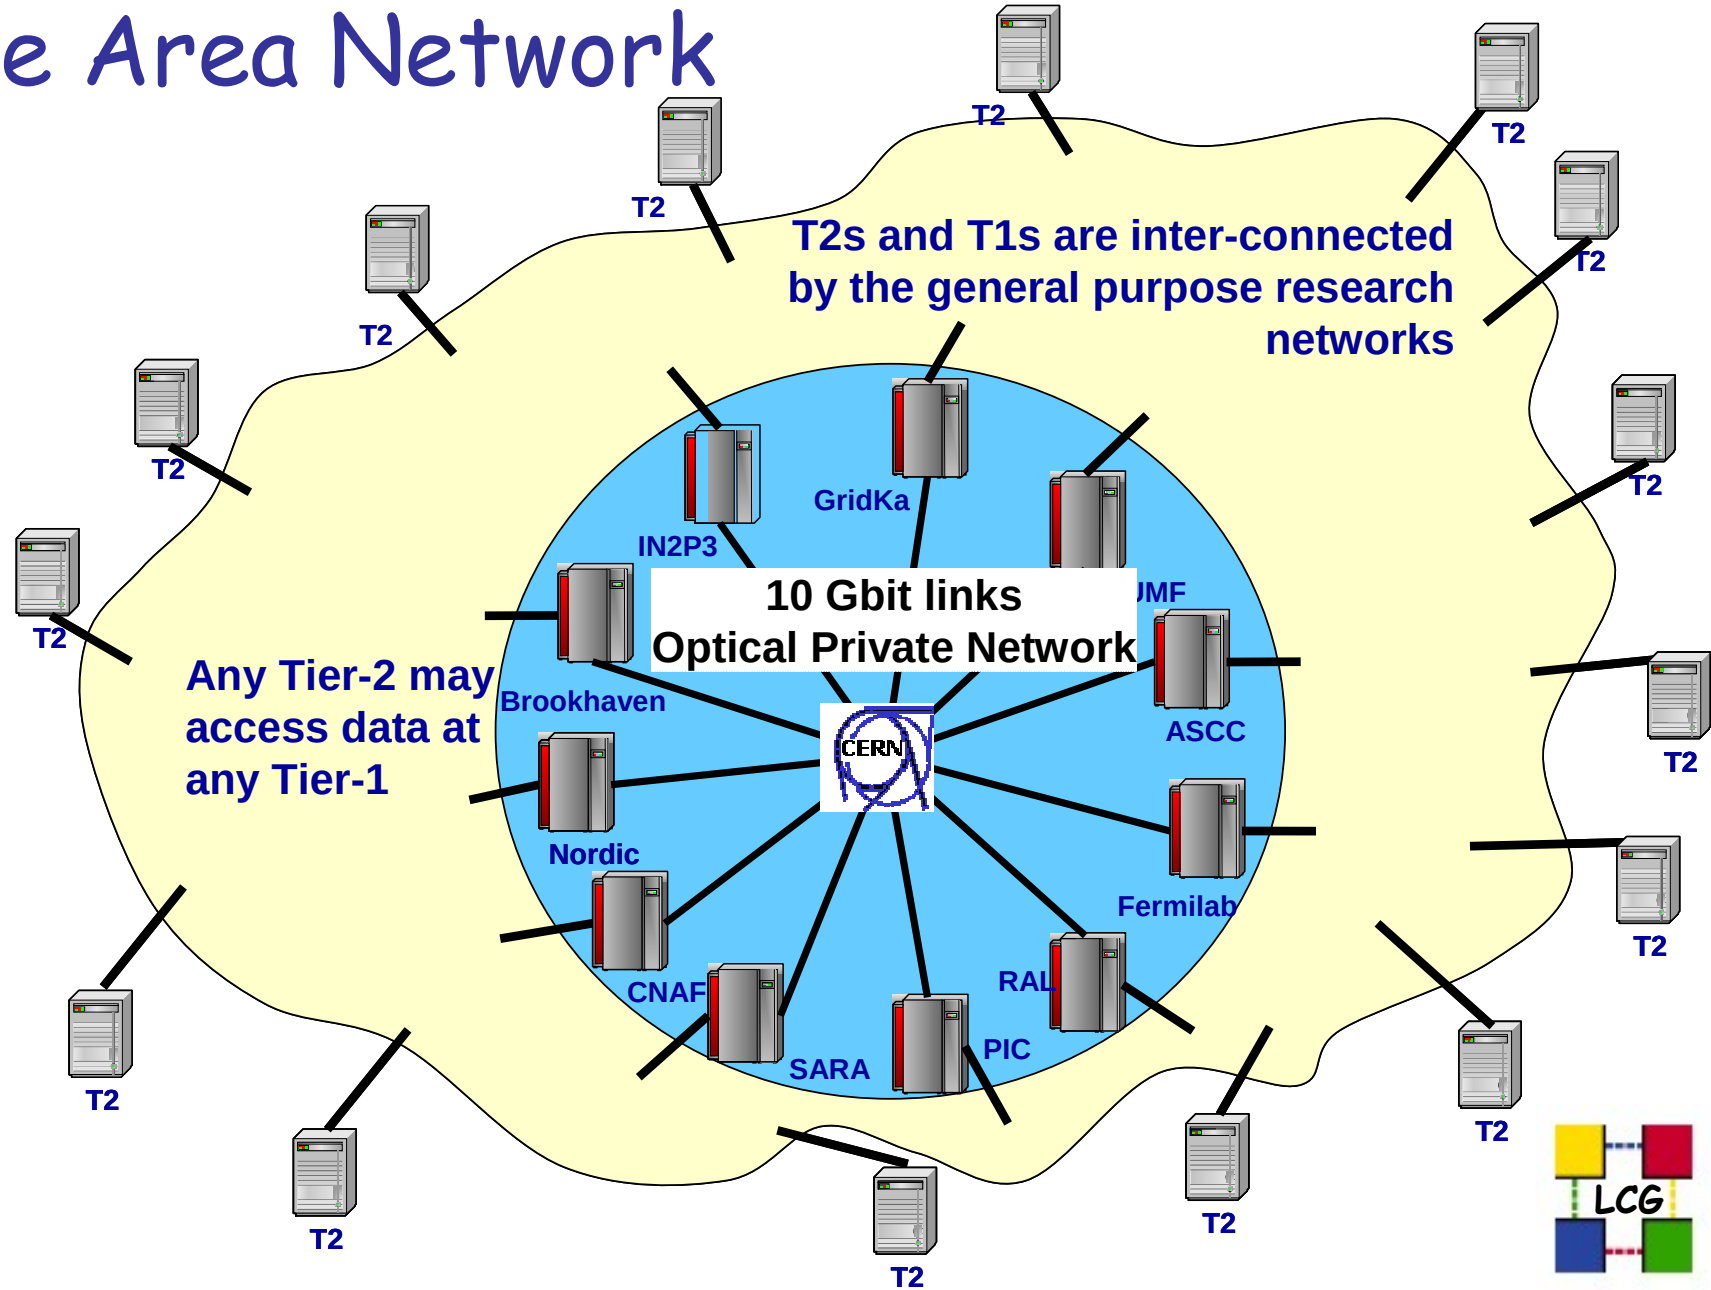

Sustained Data Distribution Rates CERN → Tier-1s

Centre	ALICE	ATLAS	CMS	LHCb	Rate into T1 MB/sec (pp run)
ASGC, Tai			X		100
CNAF, Italy	X	X	X	X	200
PIC, Spain		X	X	X	100
IN2P3, Lyon	X	X	X	X	200
GridKA, Germany	X	X	X	X	200
RAL, UK		X	X	X	150
BNL, USA		X			200
FNAL, USA			X		200
TRIUMF, Canada		X			50
NIKHEF/SARA, NL	X	X		X	150
Nordic Data Grid Facility	X	X			50
Totals					1,600

Design target is twice these rates to enable catch-up after problems

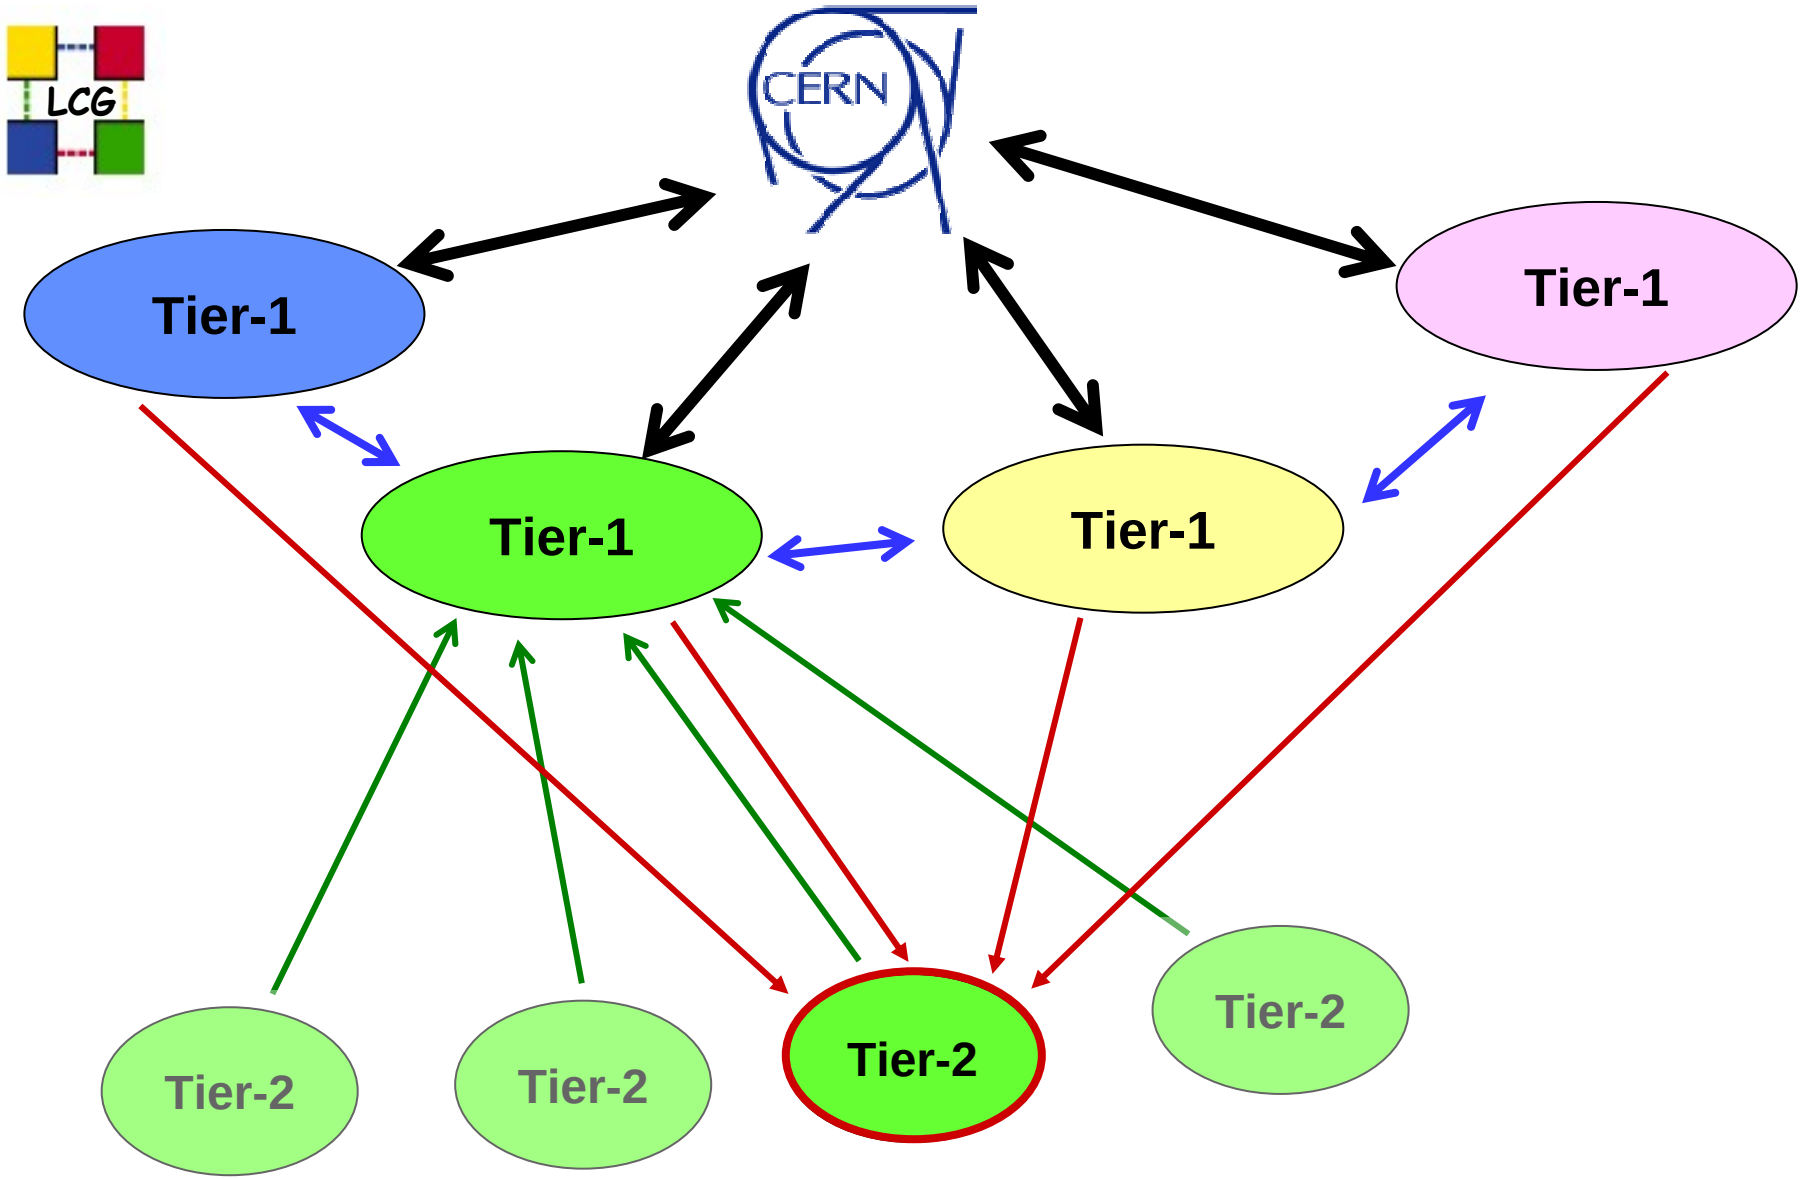

Experiment computing models define specific data flows between Tier-1s and Tier-2s

ATLAS "average" Tier-1 Data Flow (2008)

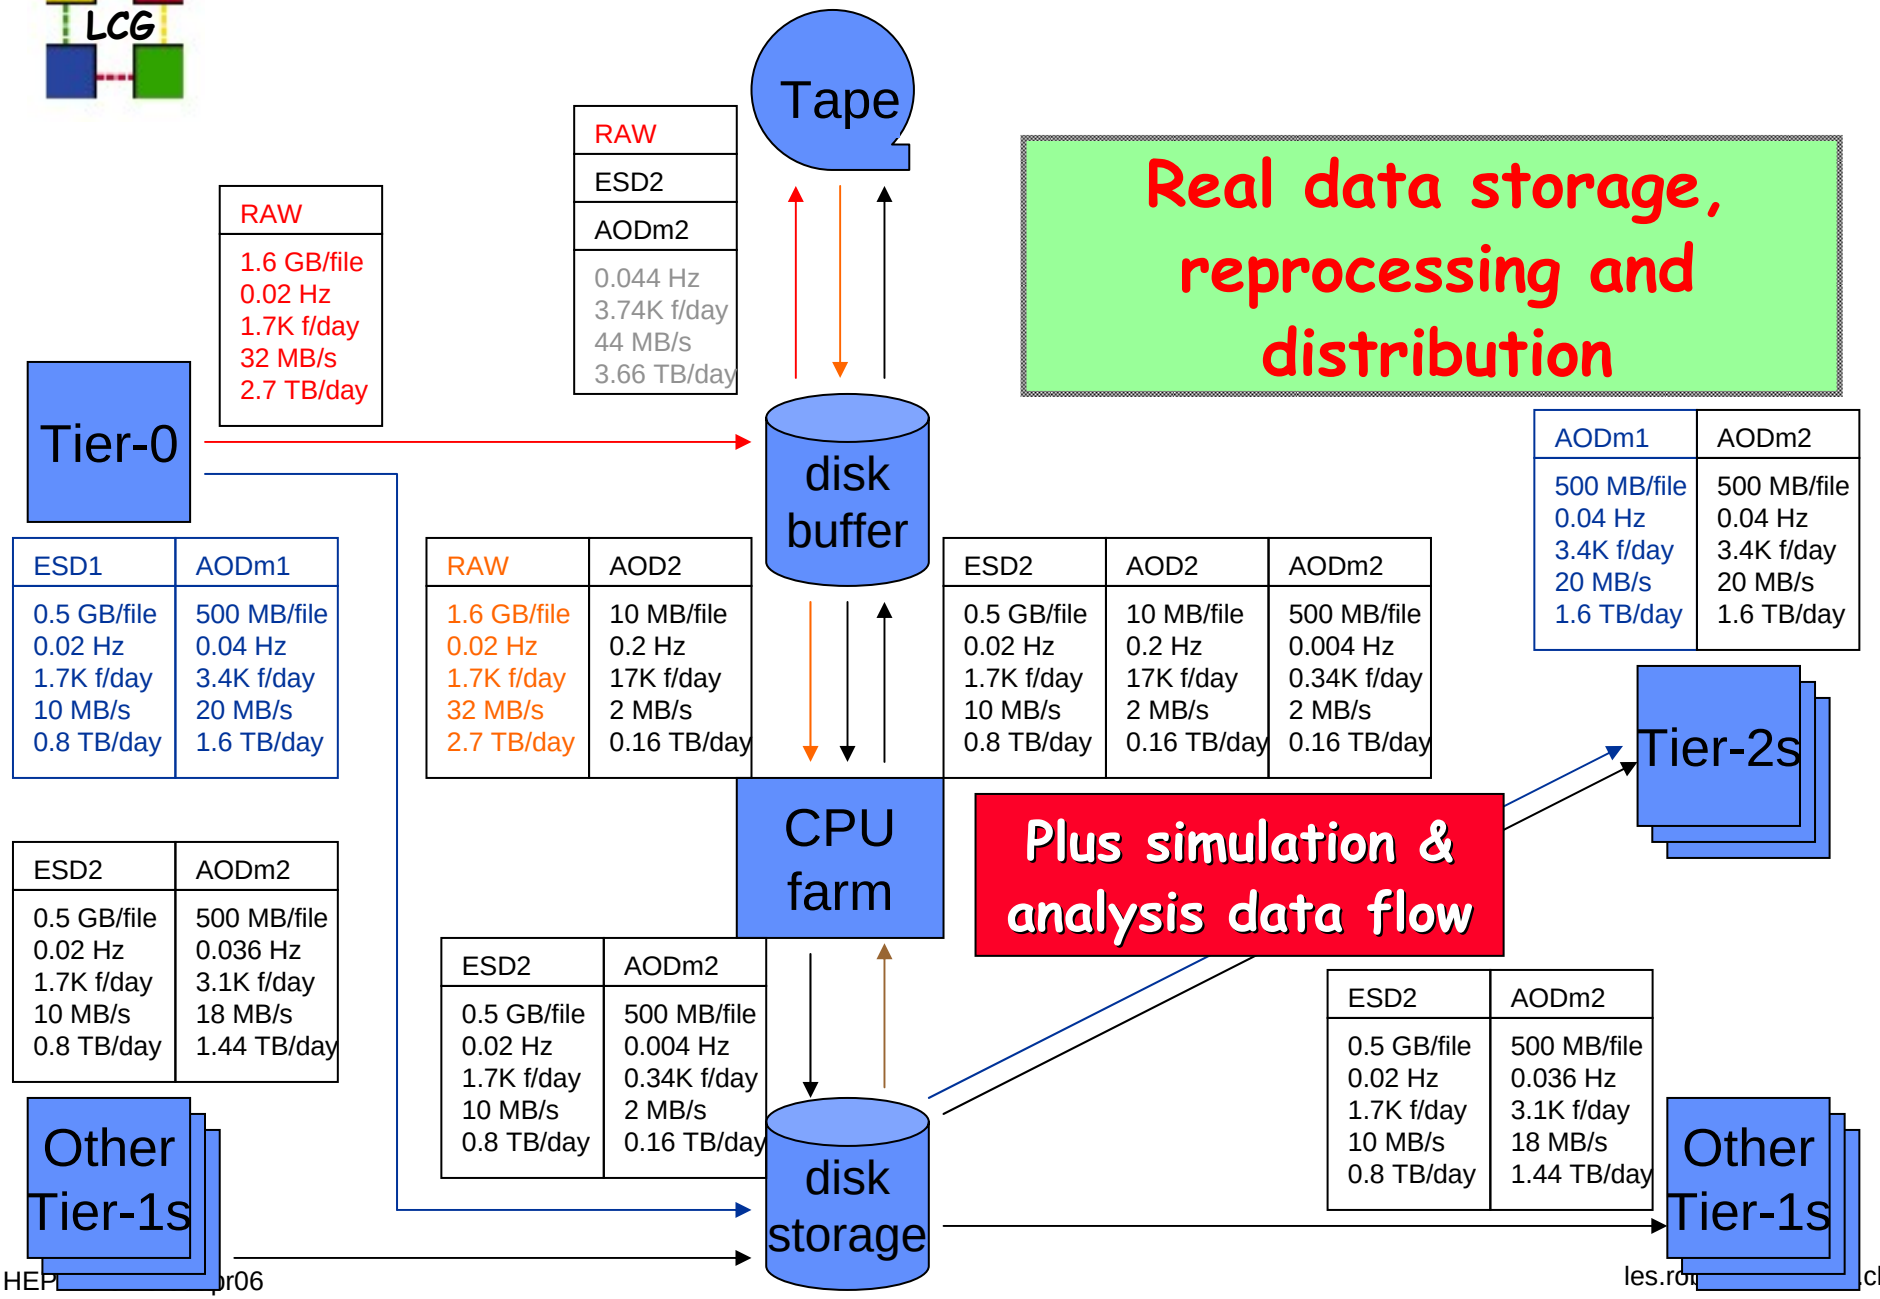

More information on the Experiments' Computing Models

- **LCG Planning Page**

Technical Design Reports

- LCG TDR - Review by the LHCC
- ALICE TDR *supplement*: Tier-1 dataflow diagrams
- ATLAS TDR *supplement*: Tier-1 dataflow
- CMS TDR *supplement* Tier 1 Computing Model
- LHCb TDR *supplement*: Additional site dataflow diagrams

- **GDB Workshops**

→ Mumbai Workshop - see GDB Meetings page
Experiment presentations, documents

→ **Tier-2 workshop and tutorials**
CERN - 12-16 June

les.robertson@cern.ch

Problem Response Time and Availability targets Tier-1 Centres

Service	Maximum delay in responding to operational problems (hours)			Availability
	Service interruption	Degradation of the service		
		> 50%	> 20%	
Acceptance of data from the Tier-0 Centre during accelerator operation	12	12	24	99%
Other essential services – prime service hours	2	2	4	98%
Other essential services – outside prime service hours	24	48	48	97%

les.robertson@cern.ch

Problem Response Time and Availability targets Tier-2 Centres			
Service	***Maximum delay in responding to operational problems***		***availability***
	Prime time	***Other periods***	
End-user analysis facility	**2 hours**	**72 hours**	**95%**
Other services	**12 hours**	**72 hours**	**95%**

les.robertson@cern.ch

Measuring Response times and Availability

Site Functional Test Framework:

- monitoring services by running regular tests
- basic services – SRM, LFC, FTS, CE, RB, Top-level BDII, Site BDII, MyProxy, VOMS, R-GMA, ….
- VO environment – tests supplied by experiments
- results stored in database
- displays & alarms for sites, grid operations, experiments
- high level metrics for management
- integrated with EGEE operations-portal - main tool for daily operations

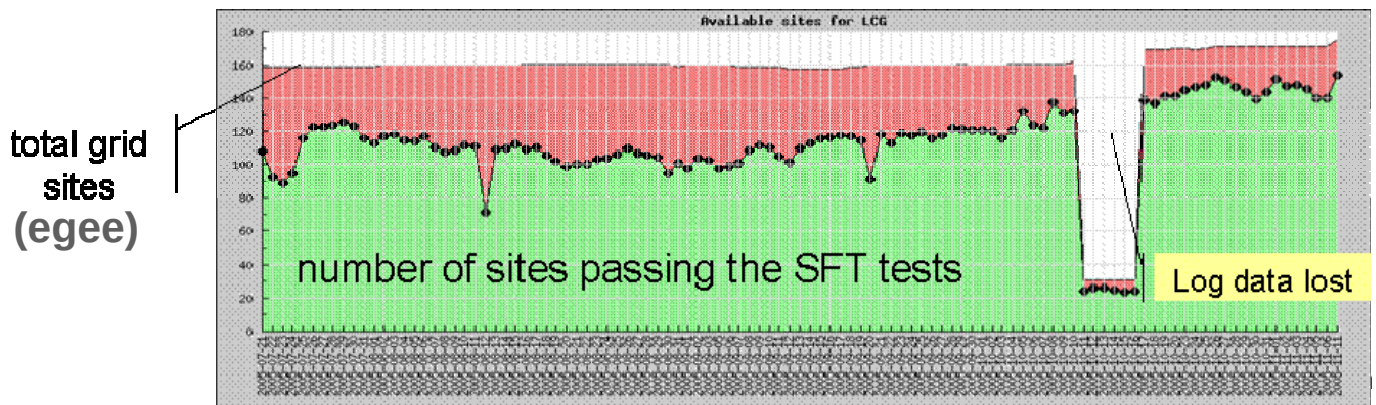

Site Functional Tests

- Tier-1 sites without BNL
- Basic tests only

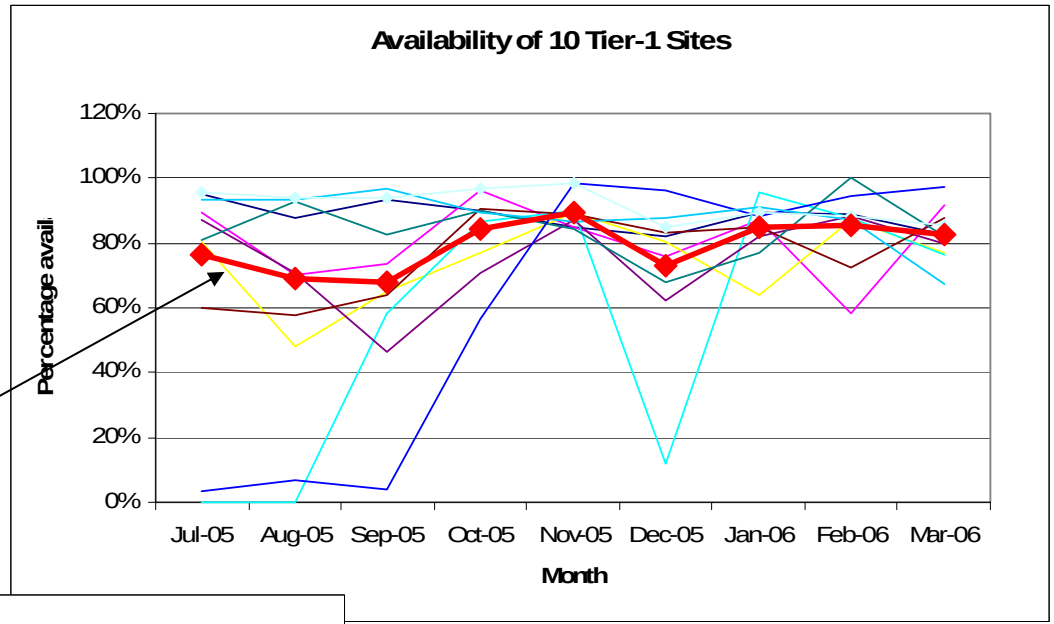

average value of sites shown

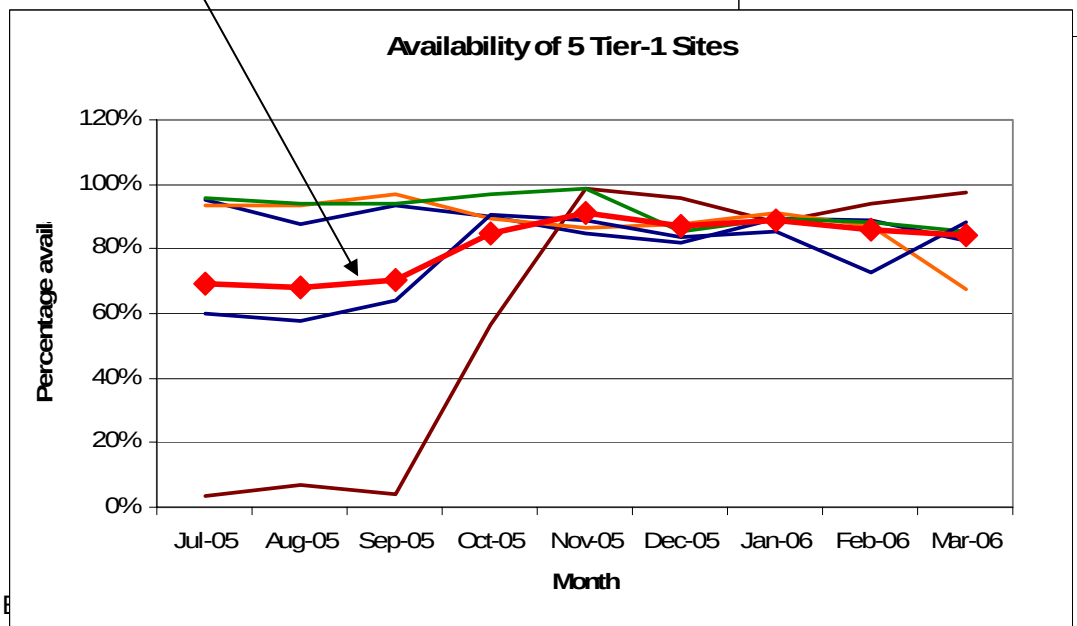

- Only partially corrected for scheduled down time
- Not corrected for sites with less than 24 hour coverage

les.robertson@cern.ch

Availability Targets

- End September 2006 - end of Service Challenge 4
  - 8 Tier-1s and 20 Tier-2s
    > 90% of MoU targets
- April 2007 – Service fully commissioned
  - All Tier-1s and 30 Tier-2s
    > 100% of MoU Targets

les.robertson@cern.ch

Service Challenges

- Purpose
  - Understand what it takes to operate a real grid service – run for weeks/months at a time (not just limited to experiment Data Challenges)
  - Trigger and verify Tier1 & large Tier-2 planning and deployment –
    - tested with realistic usage patterns
  - Get the essential grid services ramped up to target levels of reliability, availability, scalability, end-to-end performance

- Four progressive steps from October 2004 thru September 2006
  - End 2004 - SC1 – data transfer to subset of Tier-1s
  - Spring 2005 – SC2 – include mass storage, all Tier-1s, some Tier-2s
  - 2nd half 2005 – SC3 – Tier-1s, >20 Tier-2s –first set of baseline services
  - **Jun-Sep 2006 – SC4 – pilot service**

→ Autumn 2006 – LHC service in continuous operation – ready for data taking in 2007

les.robertson@cern.ch

SC4 – the Pilot LHC Service from June 2006

A stable service on which experiments can make a full demonstration of experiment offline chain

- DAQ → Tier-0 → Tier-1
  data recording, calibration, reconstruction
- Offline analysis - Tier-1 ←→ Tier-2 data exchange
  simulation, batch and end-user analysis

And sites can test their operational readiness

- Service metrics → MoU service levels
- Grid services
- Mass storage services, including magnetic tape

Extension to most Tier-2 sites

Evolution of SC3 rather than lots of new functionality

In parallel –

- Development and deployment of distributed database services (3D project)
- Testing and deployment of new mass storage services (SRM 2.1)

les.robertson@cern.ch

Medium Term Schedule

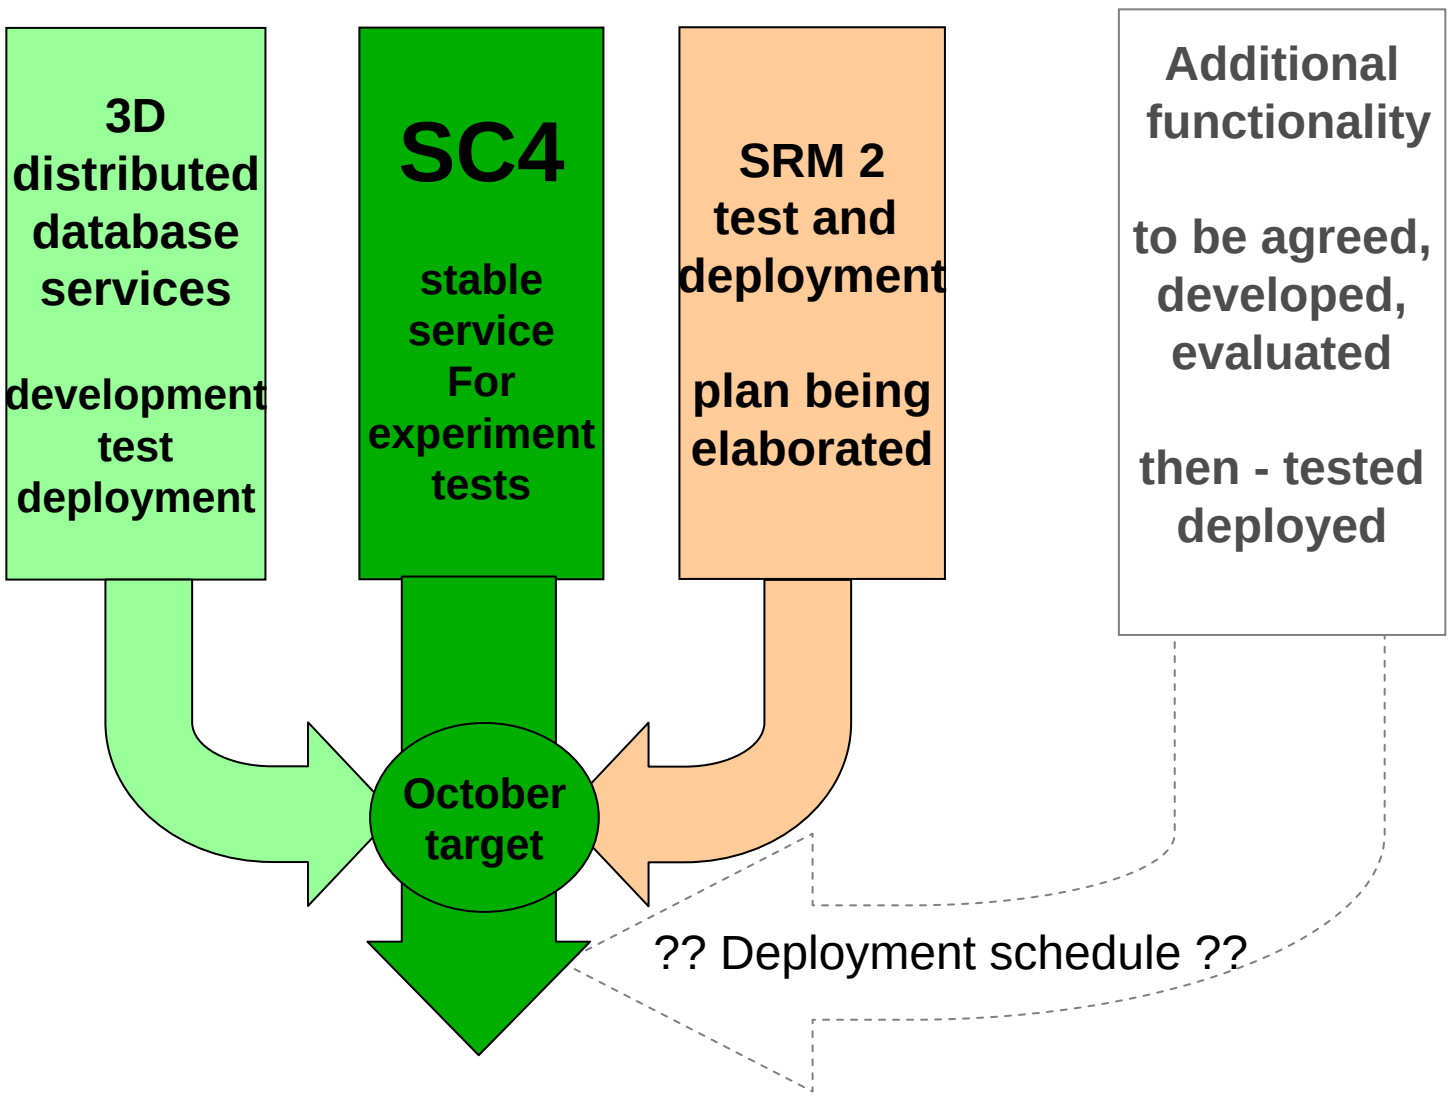

les.robertson@cern.ch

LCG Service Deadlines

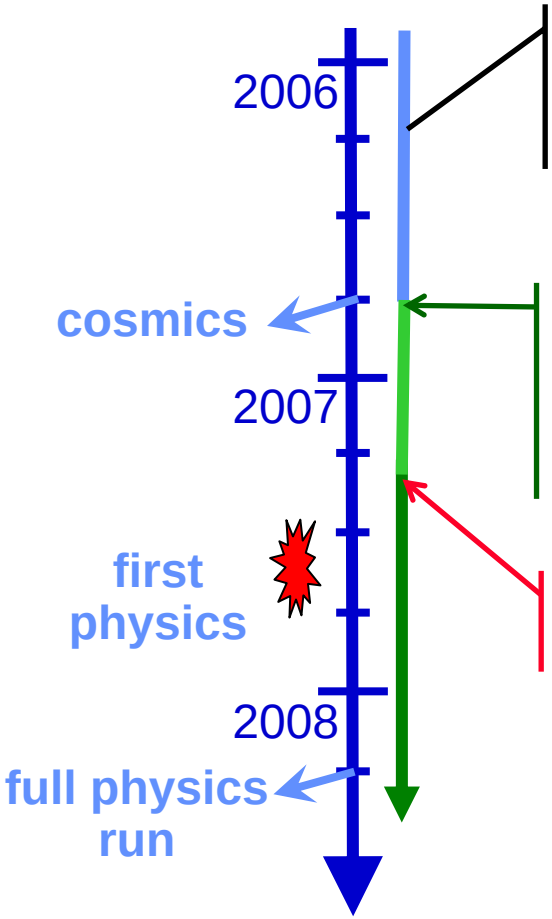

Pilot Services –
stable service from 1 June 06

LHC Service in operation – *1 Oct 06*
over following six months ramp up to full operational capacity & performance

LHC service commissioned – *1 Apr 07*

les.robertson@cern.ch

Conclusions

- LCG will depend on
  - ~100 computer centres – run by you
  - two major science grid infrastructures – EGEE and OSG
  - excellent global research networking
- We have
  - understanding of the experiment computing models
  - agreement on the baseline services
  - good experience from SC3 on what the problems and difficulties are
- Grids are now operational
  - ~200 sites between EGEE and OSG
  - Grid operations centres running for well over a year
  - > 20K jobs per day accounted
  - ~15K simultaneous jobs with the right load and job mix

BUT – a long way to go on reliability

les.robertson@cern.ch

- The Service Challenge programme this year must show that we can run reliable services

- Grid reliability is the product of many components
  – middleware, grid operations, computer centres, ….

- Target for September
  - 90% site availability
  - 90% user job success

Too modest?
Too ambitious?

- Requires a major effort by everyone to monitor, measure, debug

First data will arrive next year

NOT an option to get things going later