Management and Resources

+

Summary of Main Issues/Concerns/Recommendations

LCG Comprehensive Review
25th Nov 2004

Terry Wyatt. Manchester University.

- N.B. Very little claim to complete understanding or originality of conclusions

Some Significant Successes in 2004

- POOL
- LCG-2 deployed at 90 sites >9000 cpu 5PB storage
  - Approaching target number of sites
  - But initial deployment six months late
- LCG-2 used as a production batch system during Data Challenges
  - A success at some level
    - ~O(10M) events per expt: generated/simulated/reconstructed
  - But
    - Large amount of effort for experiments/LCG
    - ~60% "efficiency" due to many causes
    - Far from production quality system in many respects
    - Overhead for expts in having to deal with three separate grids

Summary of major admin. milestones

- Dec04: Experiment computing models submitted
- Jan05: Reviewed by LHCC (sub-committee)

- April05: MoUs for LCG (phase 2)
  - deployment and support of LHC computing
  - funding starts beginning 2006
  - (detailed scope is under discussion)

- July05: TDRs for Expt computing and LCG
- Our comment:
  - The order of MoU before TDR is not ideal!

- EGEE phase 2 planning (for period out to ~2010)

LCG Organizational Structures

- Role 1: provide the LCG Grid

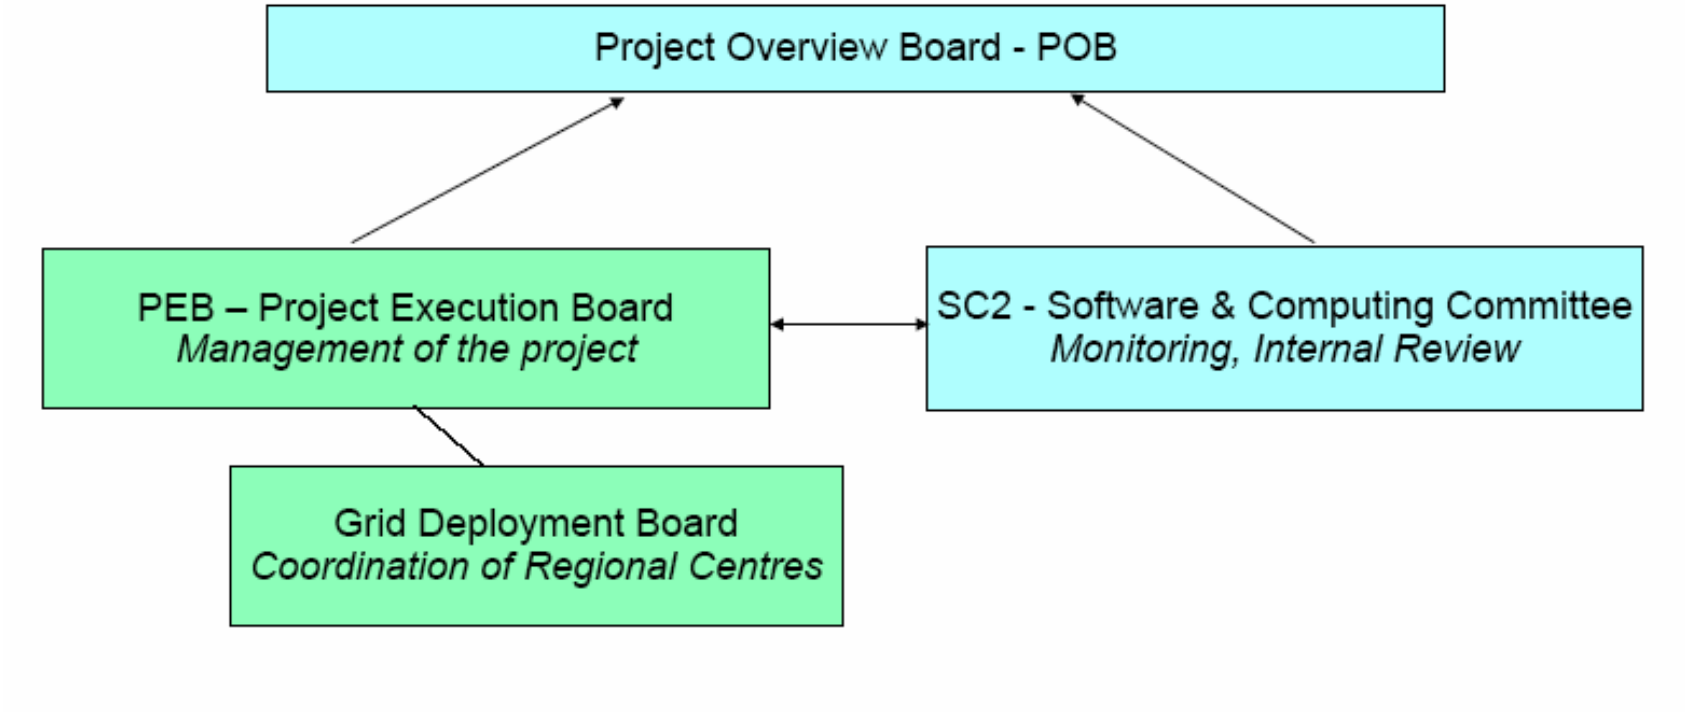

  - including coordination with EGEE for gLite middleware and other support
- Role 2: oversee/coordinate provision of Grid computing for the LHC experiments
  - including Grid3, NorduGrid as well as LCG

- Our comment:
  - Overall structure, especially lines of responsibility/authority not really clear to us
  - In particular, we did not get a good idea of how non-CERN manpower fits in
    - FTEs, management stucture (LCG, EGEE, etc)

- Our recommendations:
  - A specific forum to review/coordinate all grids available to LHC expts is required
    - How this forum would fit into current LCG structure is not clear to us
    - Perhaps better that two roles explicitly separated?
    - Perhaps another body is required?
      - with approx equal representation from the four experiments and the three Grids?
  - All three Grids should be accessible via a common interface!

Grid Commissioning Strategy

- Introduce “service challenges”
  - Test performance of a specific aspect of Grid
  - (complementary to expt Data Challenges)
- Service Challenge for robust file transfer
  - Very challenging to get such a large system in ultra reliable production
- From Dec 2004 to Nov 2006:
  - In many steps ramp up to full functionality of entire system
- Looks like an excellent way of planning commissioning
- Other service challenges planned for security, operations, and user support

Tier-1 Resources Planning

Resource type	requested	available	(for 2008)
CPU (MSI2K)	47.8	OK	
Disk (PBytes)	22.2	factor 2 low	(very preliminary)
Tape (PBytes)	16.6	OK	

- Tier-2 plan by March 2005

- Our comment:
  - Actual resource needs will depend on actual analysis patterns in 2008
    - which will probably look nothing like what gets put into the expt computing model documents in Dec04!
      - (despite the best efforts of all involved)
      - e.g., from experience in Run 2

Manpower at CERN

year	2004	2005	2006	2007	2008
REQUIREMENTS					
Applications	51.1	43.8	41.1	32.6	28.7
Physics Services	46.5	48.1	46.6	45.9	42.6
Grid	30.8	30.6	32.5	30.9	23.4
LCG Project Management	5.8	5.7	4.9	3.6	2.9
TOTAL REQUIREMENTS	**134**	**128**	**125**	**113**	**98**

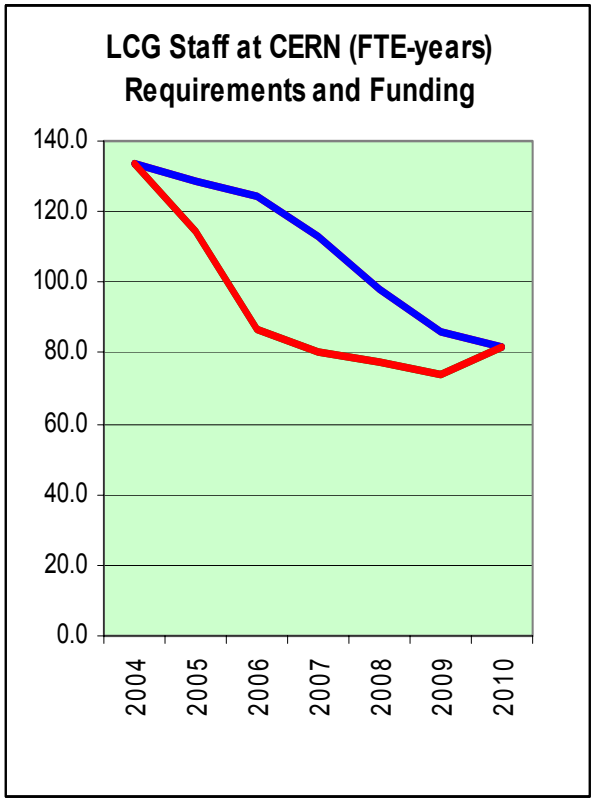

- Missing 107 FTE*years
  - under assumption that EGEE funding for “HEP” continues in phase 2 at present level
- Serious discussions to fill shortfall are ongoing
- Planned ramp down as development → support
- Our comment:
  - Not convinced that all projects will be out of development mode by 2008 (let alone 2006)

EGEE Middleware - gLite

- Relied upon to provide much improved performance/functionality

  (wrt EDG middleware currently used in LCG-2)

  - error messages/traceback/monitoring
  - security
  - scalability
  - workload management
  - support for bulk operations
  - data management
- Very little of this is currently available in anything like pre-production quality on the test-bed

Our comments/recommendations:

- Need to ensure gLite is deployed asap on test-bed and then moves to certification and release preparation on "pre-production" system
- Very important to continue to require rigorous testing
- Need to ensure that gLite developers take this debugging/commissioning as their main priority
  - rather than developing next version
  - this is agreed by the gLite manager
  - need to ensure this policy is respected by remote developers
- Individual pieces should be made available for deployment onto production LCG-2 sites
  - (Some individual pieces will provide useful functionality)
  - (N.B. nobody is forced to use gLite at this stage, because …..
- Current LCG-2 system has to be properly maintained whilst all of this is going on!
  - about 2 FTEs for short term support have been provided

Our concerns:

- When is gLite actually going to be deployed?
  - Many things are claimed to be very close to being ready, but …..
  - overall timescale is very uncertain
- Not at all clear what is the correct solution for data management
  - (even assuming current gLite work succeeds)
- Severe concern:
  - Grid-based analysis of Data Challenge data and ARDA developments are waiting for gLite
  - This will take a significant time to get going after gLite available
  - Getting the expt computing TDRs written in time for July05 will be very difficult

Major Items to Watch Early in 2005

- gLite middleware deployment:
  - We shall expect to see major progress (at least at level of test bed) by Jan Meeting!

$\Rightarrow$ real progress in Grid-based analysis of DC04 data and in ARDA

- Service Challenge for robust file transfer
- CASTOR improvements
- Applications Area$\Leftrightarrow$ROOT coordination/cooperation/communication
  - e.g., dictionary incompatibilities that prevent POOL files being read by ROOT
- General reliability of Grid operation
  - quantified in terms of metrics
- But we have not reviewed in detail formal project Milestones

- After computing models are agreed, but before TDRs are written, take stock:
  - Does middleware do what the expts need?
  - Are common software projects working on the right tools?

Final (rather vague) worries

- Are:
  - the experiments
  - the HEP world outside CERN
  - other fields of science

  as tightly coupled into LCG as they need to be?
- Disappointing for us that EGEE conference was scheduled at same time as this review:

  ⇒ e.g., many relevant people missing from this session
- How should computing reviews be structured in future?
  - How to get all referees involved?
    - (Disappointing level of attendance/involvement by LHCC as a whole in this review --- much less than in typical expt review)
  - Computing TDRs will definitely need review from entire LHCC!

Personal optimism that it will all work:

- c.f. currently running experiments
- e.g., DØ SAMGrid
  - currently generates/simulates/reconstructs

    5-10M full MC events worldwide per month
  - will reprocess entire run 2 data sample offsite in 2005:
    - 10^9 events, PB data, several THz CPU for six months