

# Planning the Planning for the LCG Service Challenges

Jamie Shiers, CERN-IT-GD

28 January 2005

## Goal of this presentation

- Many of the slides in this presentation are from previous Service Challenge presentations
  - They are included for completeness and reference
  - I would prefer to take as much time as possible to focus on a 'planning discussion' to understand what are the key issues that people would like to have clarified and in which order...
  - I think that we have heard in the last couple of days enough reasons to believe that:
    - This is a huge task!
    - That detailed plans are urgently required.
- ☹ **This is not a detailed SC3 project plan e.g. MS Project**

## Timeline

- Today: 'Long Term Planning'
- Wednesday 3<sup>rd</sup> February: review of SC3 milestones with IT-GM
- Tuesday 8<sup>th</sup> February: GDB meeting at CERN
  - **Service Challenges Milestones Document (30')**
- ...
- Timescales extremely "aggressive"
- Need to start multiple parallel threads...

## To be done over the next six months (PEB)

1. End January - Service Challenge planning paper (Service Challenge manager, Jamie Shiers)
2. February 2005 - Detailed plan for deploying reliable file transfer service (GDA)
3. March 2005 - Tier-0/Tier-1 middleware agreement - inc. workload management, catalog (PEB)
4. March 2005 - Detailed plan/agreement for deploying mass storage implementations - inc. SRM 2, light-weight disk pool manager (GDA/GDB)
5. End March - Applications area long term plan (AA manager)
6. End March 2005 - Initial plan for Tier-0/Tier-1 networking (GDB/Network WG, David Foster)
7. Early April 2005 - Final proposal for MoU (MoU Task Force)
8. Early April 2005 - Re-assessment of Tier-0, Tier-1 plans in light of new model/requirement models (P2P group)
9. Pasta IV complete (Sverre Jarpe)
10. April 2005 - Database strategy (3D Project, Dirk Duellmann)
11. Middleware rollout plan (MWA, Frédéric Hemmer)
12. End June 2005 - Production of the TDR (TDR Editor)

## Informal Overview

- In addition to planned GDB meetings, Service Challenge Meetings, Network Meetings etc:
- Visits to all Tier1 sites (initially)
  - Goal is to meet as many of the players as possible
  - Not just GDB representatives! Equivalents of Vlado etc.
- Current Schedule:
  - Aim to complete many of European sites by Easter
  - "Round world" trip to BNL / FNAL / Triumf / ASCC in April
- Need to address also Tier2s
  - Cannot be done in the same way!
  - Work through existing structures, e.g.
  - HEPiX, national and regional bodies etc.
    - e.g. GridPP (12)
- Talking of SC Update at May HEPiX (FZK) with more extensive programme at Fall HEPiX (SLAC)
  - Maybe some sort of North American T2-fest around this?

## Progress So Far...

- **High Level Overview document of Service Challenge goals**
  - The "Mandate" of the activity?
  - Useful as basis for initial discussions (and to boot-strap me...)
  - Re-cycled for LCG TDR ☺
- "Requirements" document, based on Summary of Computing Models and re-calculation of base network rates
- (Detailed) planning for SC3 (just) started - milestones document planned for next week(!)
- **Global Milestones document also required**
  - To allow partners (other IT groups, equivalents at other sites, network providers, experiments) to understand what is required of them and to produce their own plans

# LCG Service Challenges - Overview

- LHC will enter production (physics) in April 2007
  - Will generate an enormous volume of data
  - Will require huge amount of processing power
- LCG 'solution' is a world-wide Grid
  - Many components understood, deployed, tested..
- But...
  - Unprecedented scale
  - Humungous challenge of getting large numbers of institutes and individuals, all with existing, sometimes conflicting commitments, to work together
- LCG must be ready at full production capacity, functionality and reliability in less than 2 years from now
  - Issues include h/w acquisition, personnel hiring and training, vendor rollout schedules etc.
- Should not limit ability of physicist to exploit performance of detectors nor LHC's physics potential
  - Whilst being stable, reliable and easy to use

# Key Principles

- Service challenges results in a series of services that exist in parallel with baseline production service
- Rapidly and successively approach production needs of LHC
- Initial focus: core (data management) services
- Swiftly expand out to cover full spectrum of production and analysis chain
- Must be as realistic as possible, including end-end testing of key experiment use-cases over extended periods with recovery from glitches and longer-term outages
- Necessary resources and commitment pre-requisite to success!
- Should not be under-estimated!



## Initial Schedule (1/2)

- Tentatively suggest quarterly schedule with monthly reporting
  - e.g. Service Challenge Meetings / GDB respectively
  - Less than 7 complete cycles to go!
- Urgent to have detailed schedule for 2005 with at least an outline for remainder of period asap
  - e.g. end January 2005
- Must be developed together with key partners
  - Experiments, other groups in IT, T1s, ...
- Will be regularly refined, ever increasing detail...
- Detail must be such that partners can develop their own internal plans and to say what is and what is not possible
  - e.g. FIO group, T1s, ...

## Initial Schedule (2/2)

- Q1 / Q2: up to 5 T1s, writing to tape at 50MB/s per T1 (no expts)
- Q3 / Q4: include two experiments and a few selected T2s
- 2006: progressively add more T2s, more experiments, ramp up to twice nominal data rate
- 2006: production usage by all experiments at reduced rates (cosmics); validation of computing models
- 2007: delivery and contingency
- N.B. there is more detail in December GDB presentations
- Need to be re-worked now!



# 2005 Q1(i)

## SC2 - Robust Data Transfer Challenge

Set up infrastructure for 6 sites

- Fermi, NIKHEF/SARA, GridKa, RAL, CNAF, CCIN2P3


Test sites individually

- at least two at 500 MByte/s with CERN

Agree on sustained data rates for each participating centre

Goal - by end March sustained 500 Mbytes/s aggregate at CERN

In parallel - serve the ATLAS "Tier0 tests" (needs more discussion)





# 2005 Q1(ii)

In parallel with SC2  
 - prepare for the next service challenge (SC3)

Build up 1 GByte/s *challenge* facility at CERN


- The current 500 MByte/s facility used for SC2 will become the *testbed* from April onwards (10 ftp servers, 10 disk servers, network equipment)

Build up infrastructure at each external centre

- Average capability ~150 MB/sec at a Tier-1 (to be agreed with each T-1)

Further develop reliable transfer framework software

- Include catalogues, include VO's



## **SC1 / SC2**

- SC1 was covered by James and site reports
- SC2 plans also discussed yesterday...
- But it is clearly an on-going activity and a necessary step towards the next one ...

 **SC3**



# 2005 Q2-3(i)

## SC3 - 50% service infrastructure

- Same T1s as in SC2 (Fermi, NIKHEF/SARA, GridKa, RAL, CNAF, CCIN2P3)
- Add at least two T2s
- "50%" means approximately 50% of the nominal rate of ATLAS+CMS


Using the 1 GByte/s challenge facility at CERN -

- Disk at T0 to tape at all T1 sites at 80 Mbyte/s
- Data recording at T0 from same disk buffers
- Moderate traffic disk-disk between T1s and T2s

Use ATLAS and CMS files, reconstruction, ESD skimming codes

Goal - 1 month sustained service in July

- 500 MBytes/s aggregate at CERN, 80 MBytes/s at each T1



## SC3 Planning

- Meetings with other IT groups at CERN to refine goals of SC3 (milestone document) and steps that are necessary to reach them
- IT-GM: definition of middleware required, schedule, acceptance tests etc
- IT-ADC: pre-production and production services required, e.g. Database backends
- IT-FIO: file transfer servers etc etc
- IT-CS: network requirements etc
- Informal discussions with experiments regarding involvement of production teams
- Many details missing: (one being identification of T2s)

## **SC3 Planning - cont**

- Base (i.e. non-experiment) software required for SC3 scheduled for delivery end-February 2005
- Targeting service infrastructure for same date
  - Database services, Gridftp servers etc.
- Acceptance testing during March
- In parallel, have to start discussions with experiments, T1s etc






## 2005 Q2-3(ii)

In parallel with SC3 prepare additional centres using the 500 MByte/s test facility

- Test Taipei, Vancouver, Brookhaven, additional Tier-2s

Further develop framework software

- Catalogues, VO's, use experiment specific solutions





# 2005 - September-December (i)


## 50% Computing Model Validation Period

The service exercised in SC3 is made available to experiments for computing model tests

Additional sites are added as they come up to speed

End-to-end data rates -

- 500 Mbytes/s at CERN (aggregate)
- 80 Mbytes/s at Tier-1s
- Modest Tier-2 traffic





## 2005 - September-December (ii)


In parallel with the SC3 model validation period,  
in preparation for the first 2006 service challenge (SC4) -

Using 500 MByte/s test facility

- test PIC and Nordic T1s
- and T2's that are ready (Prague, LAL, UK, INFN, ..)

Build up the production facility at CERN to 3.6 GBytes/s

Expand the capability at all Tier-1s to full nominal data rate





# 2006 - January - August

## SC4 - full computing model services


- Tier-0, ALL Tier-1s, all major Tier-2s operational at full target data rates (~1.8 GB/sec at Tier-0)
- acquisition - reconstruction - recording - distribution, **PLUS** ESD skimming, servicing Tier-2s

Goal - stable test service for one month - April 2006

## 100% Computing Model Validation Period (May-August 2006)

Tier-0/1/2 full model test - All experiments

- 100% nominal data rate, with processing load scaled to 2006 cpus




# 2006 - September



The SC4 service becomes the permanent LHC service - available for experiments' testing, commissioning, processing of cosmic data, etc.

All centres ramp-up to capacity needed at LHC startup

- TWICE nominal performance
- Milestone to demonstrate this 6 months before first physics data



## Timeline

- Official target date for first collisions in LHC: April 2007
- Including ski-week(s), this is only 2 years away!
- But the real target is even earlier!
- Must be ready 6 months prior to data taking
- And data taking starts earlier than colliding beams!
- Cosmics (ATLAS in a few months), calibrations, single beams, ...

## Conclusions

- To be ready to fully exploit LHC, significant resources need to be allocated to a series of service challenges by all concerned parties
- These challenges should be seen as an essential on-going and long-term commitment to achieving production LCG
- The countdown has started - we are already in (pre-)production mode
- Next stop: 2020