

User Board Meeting, RAL, 29th March 2006

Present: Catalin Condurache, Jeremy Coles, Matt Hodges, Roger Jones, Raja Nandakumar, Dave Newbold (Chair), Glenn Patrick (Minutes), Dave Sankey, Andrew Sansum, Dan Tovey, Peter Richardson, Nick West, Fergus Wilson

1. PREVIOUS MINUTES

These were accepted without alteration.

Action 1: Update website with resource schedules, minutes, etc. - Dave Newbold.

2. EXPERIMENT REQUESTS FOR 2006

2.1 2006 Quarter 1 Usage

The latest numbers on quarter 1 usage were presented by Andrew Sansum. BaBar had been the largest CPU consumer of the Tier 1, with the LHC experiments starting to ramp up production. The farm had been running full for the first two months of 2006 with an averaged utilised capacity of 540 kSI2k out of an available 830 kSI2k - the difference attributed to the efficiencies of jobs, wait times, etc.

On the disk front, as of March 2006 BaBar were using 86.8TB (98.4TB available, 95TB allocated), ATLAS 12.1TB (18.1TB available, 14TB allocated), CMS 24.2TB (31.9TB available, 14TB allocated) and LHCb 7.3TB (9.65TB available, 10TB allocated). BaBar still needed to release one disk server and CMS were still using more than their allocation. Experiments need to maintain a dialogue with Matt and Catalin on how to effectively release servers and not just delete disk files. It was noted that, as a result, disk had not been made available to the smaller experiments as planned.

Action 2: CMS to review disk space use and come within allocation - Tim Barrass/Dave Newbold.

Action 3: Liaise with Matt and Catalin on disk servers - all experiments.

2.2 Tier-1 CPU Resources

The total requests for Q2 came to 1,256 kSI2k out of a potential capacity of 1,012 kSI2k (including cpu being deployed and due to be available in May/June). It was therefore decided to overallocate for Q2 and scale the experiment requests to fit with the available capacity. There was still a lot of Tier 1 usage that is not through LCG.

Action 4: Ratio of Grid/non-Grid experiment usage to be reported for next meeting - Andrew Sansum et al.

2.3 Tier 2 CPU Resources

The Tier 2 CPU requests for Q2 totalled 3,800 kSI2k out of a total of 4,340 kSI2k. Jeremy said that there were some anomalies in the accounting (e.g. Manchester wrongly accounted, Cambridge Condor use not included), but a capacity of 4,000 kSI2k was a good estimate. All the main experiments planned large-scale use of the Tier 2 sites consistent with the published capacity from Q2 (ATLAS 950 kSI2k, CMS 400 kSI2k, LHCb 1500 kSI2k, BaBar 520 kSI2k) and all the way through to Q4.

2.4 Tier 1 Disk Resources

The main requests were BaBar (120TB), ATLAS (30TB), CMS (50TB) and LHCb (40TB). A total of 296.5TB was requested for Q2 out of a current availability of 177TB. However, the extra disk already at RAL would bring the total up to 308TB and the aim should be to have this available at the end of May so that it can be fully utilised for the whole of June. The capacity will then fall to

280TB in August with the retirement of 36TB of old equipment. It was noted that the BaBar disk numbers need to be confirmed before June for the 2007 budget discussions. The LHCb profile is calculated on the requirements to store and analyse data from DC06/SC4.

It was agreed to leave the current Q2 disk allocations unchanged from Q1 until the new disk is deployed when Dave will release new allocations.

Meeting the total disk requests for Q3 (413.5TB) and Q4(448.5TB) was considered to be vital and the meeting was in agreement that spend should be brought forward from 2007 to provide up to 600TB of disk ("early spend" scenario).

Action 5: Timeline for early disk/cpu purchase needs to be agreed - Andrew Sansum/Dave Newbold.

2.5 Tier 1 Tape Resources

A total of 304TB out of a capacity of 318TB was being used at the end of February. Requests for Q2 (372TB) and Q3 (623TB) exceeded the anticipated capacity of 336TB.

As BaBar need to use tape because of the disk shortfall, it was decided to overallocate for Q2 and to ask the LHC experiments to confirm their tape profiles on a month-by-month basis for the rest of 2006.

Andrew reminded the meeting that the new T10k drives work with Castor, but are not integrated with ADS at the moment. The Tier 1 Board had agreed that an extra 100TB of media for the T10k drives could be bought which would bring the total capacity to 446TB. The "early-spend" scenario may be needed to meet the requests towards the end of the year.

Tape bandwidth requirements had been provided by CMS (100MB/s) and BaBar (17MB/s) with CASTOR able to provide 60 MB/s per drive. It was felt that this discussion should be deferred so that the numbers can be understood in terms of the experiment computing models at the dataflow meeting on 21st April.

Action 6: Confirm tape profiles on monthly basis for rest of 2006 - all LHC experiments.

2.6 Tier 2 Disk Resources

It was estimated (by a nightly cron job) that 55TB out of 144TB of disk was currently being used at the Tier 2 centres. The planned use of Tier 2 disk by experiments in Q2 and Q3 is consistent with the capacity although it was not clear how this would get accounted.

Action 7: Publish resource Tier 1 and Tier 2 allocations for Q2/Q3 asap - Dave Newbold.

3 UB Structure

There was a discussion on the role of the UB. It was agreed that:

- We should go back to having formal representatives from each experiment and the list of contacts would be resurrected.
- The meetings should be kept open to anyone from the particle physics community.
- The resource planning of the UB did not really work. This needs to be a separate function with someone empowered to take this on board as a continuous responsibility not just at quarterly reviews. There was also a need for more Tier 2 input and the involvement of other people involved with the financial planning.
- The role of the UB chair in GRIDPP3 was not clear and would almost certainly need to be changed.

Action 8: Resurrect list of experiment representatives - Dave Newbold.

4 GRIDPP3 Timeline

A single bid had been proposed for particle physics computing until 2010. The PPARC call had not appeared when expected, but was still supposed to appear before the end of March. There were many unknowns, but the current plan was:

Early April - Capture user requirements

Mid April - Scoped initial plan

28 April - Proforma to define work plans.

8 May - Metrics and milestones review.

15 May - CB meeting.

Mid June - Scoped final plan.

30 June - OC meeting.

End July - Proposal.

This would get reviewed by the PPRP and be discussed at the first meeting of the Science Committee in 2007 with the result announced in March 2007. As part of this timetable, GridPP2 would be extended until April 2008.

Input would be sought from the LHC experiments and other main users (up to 2010). This would include resource requirements for the period 2007-2012 divided into classes such as "simulation", "reconstruction" and "analysis". Service level requirements and strategy would also need to be connected with the simplified computing models.

In order to have a standard format, Dave would send round a questionnaire in the next few days. An opportunity for discussion might also be possible at the Warwick IOP meeting.

Action 9: Circulate questionnaire to all experiments/projects - Dave Newbold.

Action 10: Start thinking about input - All experiments/projects.

5 TRAINING and PUBLICITY

At the last GRIDPP Oversight Meeting it was felt that not enough training opportunities were being made available and experiments were encouraged to make proposals. Roger Jones pointed out that it was not clear what training facilities and charges are available (e.g. NeSc).

It had been noticed that very few of the GRIDPP press releases were connected with physics. It had been suggested that experiments try to think of something every quarter and send something to Sarah Pearce. Mention was made of the new Bs mixing results from CDF/D0 as being a suitable topic.

Action 11: Identify training Grid opportunities to PMB - all.

Action 12: Identify/produce physics related press releases and inform Sarah Pearce - all.

6 NEXT MEETING

It was agreed to hold this on the afternoon of 28 June as part of GRIDPP16 at QMUL.