

User Board Meeting, EVO, 1st Dec 2008

Present: *Stephen Burke, Matt Hodges, Raja Nandakumar, Andrew Sansum, Glenn Patrick (Chair), Philip Rodrigues, Gianfranco Sciacca, Nick West Fergus Wilson.*

Apologies: *Jeremy Coles, Dave Colling, Cristina Lazzeroni, Roger Jones*

1. PREVIOUS MEETING

There were no changes to the draft minutes of the previous meeting held on 17 September.

Concerning the actions:

1. It was understood that the situation with the ALICE VO Box at Birmingham had been resolved after consultations between Jeremy, Glenn and Cristina.
2. The WLCG pledges for 2009 and resource forward look up until 2013 had been submitted by GridPP.
3. The remaining resource headroom at the Tier 1 Centre had been distributed by Glenn. For disk, this had largely been swallowed up by the reduced capacity after a server audit and the decision to compensate experiments for the 5% Castor overhead in storing data.
4. The list of data remaining in dCache and ADS still needed to be circulated. A reminder should also be sent to experiments about the final migration schedule to Castor. **Action - Matt/Andrew.**
5. The new monthly Tier 1 meetings had been advertised to the UB list. These are held on the first Wednesday of each month at 13:30.
https://www.gridpp.ac.uk/wiki/RAL_Tier1_Experiments_Liaison_Meeting
6. The list of Experimental Contacts on the UB web pages had been updated by Glenn. These are used by Tier 1 Staff to contact experiments in case of problems and should be kept upto date.
<http://www.gridpp.ac.uk/eb/MailList.html>.
7. Glenn had produced the Q3 GridPP report for the "other experiments" and thanked those who provided input.
8. The UB questionnaire on experiment support was due to be circulated in 2009/Q1 by Matt, but this may overlap with the lightweight Tier 1 review proposed by Andrew (see also Section 8).

2. TIER 1 USAGE FOR 2008/Q4

CPU

Matt reported that the CPU occupancy had generally been low at 35-40% with very few LHC jobs, although there had occasionally been peaks when the farm was full.

% CPU Alloc.	Jan 2008	Feb 2008	March 2008	April 2008	May 2008	June 2008	July 2008	Aug 2008	Sep 2008	Oct 2008	Nov 2008
ALICE	60.8%	3.7%	39.0%	67.3%	26.1%	0.1%	0.2%	0.3%	0.2%	56.1%	
ATLAS	37.9%	38.1%	75.5%	30.3%	36.3%	88.7%	3.4%	5.4%	16.8%	21.4%	
CMS	8.4%	22.6%	7.1%	36.3%	27.7%	13.8%	21.2%	3.9%	7.0%	8.4%	
LHCb	8.0%	18.6%	21.5%	17.3%	25.6%	12.9%	21.3%	5.0%	1.9%	17.5%	
BaBar	19.7%	12.8%	59.6%	46.7%	105.6%	67.8%	153.7%	154.5%	149.8%	71.2%	

The overall CPU efficiency was 60-70% and had been good for the LHC experiments, but lower for the smaller experiments. Although the LHC experiments are usually tracked and contacted when efficiency problems are observed, this is not usually the case for the other experiments.

Batch Efficiencies	Jan 2008	Feb 2008	March 2008	April 2008	May 2008	June 2008	July 2008	July 2008	Aug 2008	Sep 2008	Oct 2008	Nov 2008
Efficiency (%)	75.8%	72.0%	90.2%	85.1%	52.7%	79.7%	57.8%	52.0%	52.0%	72.4%	72.8%	58.3%

Current efficiencies can be found at <http://www.gridpp.rl.ac.uk/stats/eff/RAL/All/archive/summary.html>

The throttling of BaBar jobs (500 cpu limit) was raised and Glenn said it was important that idle capacity in the farm could be used because of the lack of LHC data taking. Andrew said he would investigate what could be done to raise or remove this limit. It was also subsequently agreed by the PMB (see item 5 in <http://www.gridpp.ac.uk/pmb/minutes/081201.txt>) that BaBar should be encouraged to use free resources.

DISK

The Q4 allocations were changed by Glenn to take account of slightly reduced capacity after a disk audit and the 5% overhead from Castor. A new deployment process was now in place that should be faster.

At the moment, ATLAS are just ~30TB short of allocation (1,056TB) and CMS just ~20TB short of allocation (650TB). LHCb have their complete allocation of 169 TB (after 5% reduction from Programmatic Review).

Glenn raised the question of how the smaller experiments would keep within their individual allocations if they shared disk pools on the GEN instance of Castor. At the moment, there are no quotas enforced and this would need monitoring to ensure that each experiment kept within their limit.

TAPE

Overall tape usage (1,294TB in October) was well under the total allocation of (2,092TB in October).

3. STATUS OF TIER 1 PROCUREMENT AND NEW MACHINE ROOM

Andrew reported that the disk order went out a few weeks ago with delivery expected in early January. Any extra capacity (over the original plans) would probably just offset the 5% overhead due to Castor. It was expected that the disk would be available for deployment to experiments from April after testing.

The CPU order was just about to go into the system. Delivery anticipated for second week of February, but faster deployment/testing should mean that it is also available from April alongside the disk.

For tape, the planning figures would be achieved. A new robot would be delivered the first week of January and the particle physics drives would be installed end January/early February.

The current estimate for the new machine room to be handed over is 22 December. This date is not certain and the building would also have to be accepted and any snags resolved. The external planning for moving the equipment had yet to go to the PMB, but it was expected that this would be done towards the end of January. This would be done in stages with minimum downtime for critical services (e.g. FTS, LFC).

4. TIER 1 STORAGE REVIEW

Glenn reported that a review of the mass storage system at the Tier 1 centre had taken place on 21 November and thanked those experiments/projects which had provided input. ATLAS/CMS/LHCb had been part of the review, but it was accepted that the main problem for the other experiments was the available support and documentation. Prioritisation against the main LHC experiments was also difficult. A summary of the meeting can be found at <http://www.gridpp.ac.uk/pmb/minutes/081121.txt>

5. STATUS OF CASTOR MIGRATION

Glenn reported that ATLAS, CMS and LHCb had all now migrated to Castor and everything was still on track to close dCache at the end of 2008, with Castor as the only mass storage system for particle physics.

Nick reported on the MINOS migration. The experiment had 3TB of legacy data in ADS, 2TB in dCache and 6TB in NFS. The ADS data was now not needed and could be signed off. In the case of dCache, confirmation was awaited from one user, but this should be sorted out on the timescale of ~2 weeks. The NFS data was harder to migrate as it represented a legacy of 7-8 years work and then there was the question of the ~20 users still needing interactive login to the Tier 1.

It was decided that the priority for MINOS was to complete the dCache migration to Castor as soon as possible. This would enable dCache to be shutdown early in the New Year as planned. The NFS data could take a lower priority and the interactive users could continue and the situation reviewed at the next UB meeting.

Fergus confirmed that BaBar did not need to migrate any data from the ADS. There was however one user (Tim Adye) who would like some notice of the eventual closure as he still used the facility for backup.

It was agreed that final dataset lists and notification of closure for ADS and dCache needed to be circulated to UB and experiments.

6. 2009 PLANNING AND ALLOCATIONS

Glenn reported that he planned to carry the 2008/Q4 allocations forwards into 2009/Q1, as there was no new hardware to deploy. The CPU allocations may need some adjustment depending on which experiments were active.

Attention would now switch to updating the 2009 resource profile for all experiments and matching this to the new capacities. He thanked those experiments which had provided numbers and reminded the others that the lack of a response is interpreted as a request for zero resources.

7. USER SUPPORT

Stephen reported on which experiments he had supported this quarter. The myProxy discussion with MINOS had been appreciated. Help and advice had been provided to ILC to help them with teething problems on Castor. Janusz was not available to report on his work.

8. TIER 1 MINI REVIEW

Andrew reported that he had been actioned to conduct a lightweight review of the Tier 1 Centre and to get experiment feedback on the operation and procedures. The intention was to improve the interaction with experiments/projects and to identify those things that needed improvement. There was obviously some overlap with the survey that Matt was intending to conduct.

It was agreed that Andrew should draft a plan for this review and circulate it around the UB contacts for comment. A dedicated meeting early in 2009 (either EVO, phone or F2F) could be organised.

9. NEXT MEETING

The next (standard) UB meeting would be in March 2009 and Glenn would canvass for a suitable date in the New Year.

ACTIONS

Action 1: Experiments to provide Glenn with updated resource profiles for 2009 if they have not already done so - **Experiments**.

Action 2: Plan preliminary allocations throughout 2009 using request versus capacity profile - **Glenn**

Action 3: Circulate lists of datasets remaining in dCache and ADS and remind experiments of migration schedule to Castor - **Matt/Andrew**.

Action 4: Complete migration from dCache to Castor - **MINOS**

Action 5: Investigate removal of BaBar job throttling - **Andrew**.

Action 6: Prepare and circulate proposal for Tier 1 Mini-Review and how this fits in with planned UB questionnaire - **Andrew/Matt**