

User Board Meeting, EVO, 17th Sept 2008

Present: *Jeremy Coles, James Ferrando, Matt Hodges, Roger Jones, Cristina Lazzeroni, Raja Nandakumar, Andrew Sansum, Glenn Patrick (Chair), Alexandre Sousa, Fergus Wilson.*

Apologies: *Dave Colling, Alan Irving, Dave Sankey, Gianfranco Sciacca*

1. PREVIOUS MEETING

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

Concerning the actions:

1. Matt had issued the June monthly report on disk deployment/allocation to the UB list, had missed the July report due to holidays and was about to issue the August numbers (done on 17 Sept).
2. Cristina reported that the ALICE VO box security situation at Birmingham had not been resolved. Andrew advised that RAL would express similar concerns if they were approached. Roger advised that he thought that this was a GDB Class 2 problem and advised contacting Jeff Templon at NIKHEF to understand the issues (Roger to send Cristina details by email). Subsequent to the meeting, agreement was reached on the VO box deployment at Birmingham.
3. Dave Colling had provided an updated 2008 resource request for CMS although there was no further information on the new CMS computing model.
4. Most of the non-LHC experiments had started migration to Castor and this was covered later in the agenda.
5. Details of the *Tier 1 Contacts* had been updated at <http://www.gridpp.ac.uk/eb/MailList.html>. Experiments/projects had been requested to provide updates to the *Experiment Operational Contacts* (including operational mailing lists which can be used by Tier 1 staff) and *User Board Contacts* and a number of these had been received.
6. The "UB questionnaire" on experiment support would be circulated by Matt once per year in line with GridPP3 metrics.
7. Glenn had produced the Q2 report for the "other experiments".

2. TIER 1 USAGE FOR 2008/Q3

CPU

Matt reported that the CPU occupancy had only been about 50% for the last 3 months. The LHC experiments and MINOS utilisation had been low, but BaBar, Biomed and ILC had been able to profit by exceeding their fairshares.

% of CPU allocation used	Jan 2008	Feb 2008	March 2008	April 2008	May 2008	June 2008	July 2008	Aug 2008
ALICE	60.8%	3.7%	39.0%	67.3%	26.1%	0.1%	0.2%	0.3%
ATLAS	37.9%	38.1%	75.5%	30.3%	36.3%	88.7%	3.4%	5.4%
CMS	8.4%	22.6%	7.1%	36.3%	27.7%	13.8%	21.2%	3.9%
LHCb	8.0%	18.6%	21.5%	17.3%	25.6%	12.9%	21.3%	5.0%
BaBar	19.7%	12.8%	59.6%	46.7%	105.6%	67.8%	153.7%	154.5%

Within the LHC experiments there can also exist fairshare splits for production and non-production work, although the policies were not established (e.g. for throttling user jobs). Roger emphasised that it is important that random ATLAS users are not allowed to consume T1 resources and it was

important to be able to "throttle on role" (e.g. production, software installation). Raja said that LHCb needed one SGM job to be able to run, but with no user throttling.

The CPU efficiency averaged across all experiments had been on the low side falling from 80% in June to 52% in August (currently ~64% in September). This was mainly due to the LHC experiments whilst Babar had maintained good efficiency (87-95%). Roger said that ATLAS had looked at production system failures and reckoned that they had 70% efficiency with 12% of failures due to file access and 10% due to ATLAS. Andrew pointed out that the Tier 1 figures refer purely to the CPU time/wall-clock time ratio, whereas the experiments refer to overall job success. The low LHCb efficiencies were thought to be mainly due to data access and Raja reported that a ticket had been submitted about using the *rootd* protocol in preference to *rfio* on Castor.

Batch Job Efficiencies	Jan 2008	Feb 2008	March 2008	April 2008	May 2008	June 2008	July 2008	July 2008
Efficiency (%)	75.8%	72.0%	90.2%	85.1%	52.7%	79.7%	57.8%	52.0%

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

DISK

The Q3 allocations had been mostly unchanged from Q2, but there were still deployment shortfalls of ~200TB for ATLAS and ~50TB for CMS. In the ATLAS case, some of this was due to redistribution from other storage classes. There was also a shortfall for UKQCD (1.8TB deployed, 30TB allocated), but this was due to the need to migrate from NFS to Castor and discussions had just restarted on how to migrate.

On the subject of migration to Castor, ZEUS had migrated from dCache, MINOS is testing Castor and BaBar is now configured to use Castor. Alex reported that MINOS had completed migration of their ADS data and Glenn stated that Nick West had reported some Castor failures at the beginning of September (thought to be now fixed). Andrew pointed out that the on-call service provides an out-of-hours service only for the LHC experiments. It was however important that Castor problems were followed up in the daytime for the non-LHC experiments (such as MINOS).

TAPE

Overall tape usage (1360TB in August) had been fairly constant over the quarter and below allocation (1737TB in August).

3. STATUS OF TIER 1 PROCUREMENT

Andrew reported that tenders were running for CPU and disk. The disk tender would close on 22 September, followed by a three week evaluation with delivery anticipated at the end of 2008/start 2009. The CPU tender has a slightly delayed schedule with delivery expected in mid-January.

The purchase of a new tape robot and double density drives had just been approved by DIUS as a single tender. Delivery expected by end of the year.

4. REVIEW OF EXPERIMENT REQUIREMENTS FOR 2008/Q4

Glenn said that (unusually) the experiment requirements/allocations did not warrant extensive discussion for Q4. The only update that he had received was a renewed request from CMS (see http://www.gridpp.ac.uk/eb/0908/request2008q4_CMS.xls). This was discussed and it was agreed that after meeting all other allocations for Q4, Glenn would distribute any remaining headroom to the LHC experiments on a proportionate basis. This would not prejudice the 2009 allocations (which depend on final deployment figures from current procurement) and any extra resources granted in 2008 would be seen as an "advance" on 2009 in the first instance.

Roger informed the meeting that a request from T2K was expected that afternoon (received 21 September).

Glenn reminded the meeting that he had already agreed resources for NA48 and it was now up to the experiment to make the necessary arrangements. Cristina said that they had been rehearsing their setup at Birmingham before making use of the Tier 1.

Allocations would be circulated by email in the usual way.

5. 2009-2013 WLCG PLEDGES

Glenn reported that a forward look for resource pledges to the WLCG for 2009-2013 had to be submitted by 22 October. ATLAS/CMS/LHCb would provide their own information, but it would be useful to have some feedback from the other experiments on their future profiles - even if it was only information on anticipated future funding/experiment lifetime. He would like to receive any input in the next week or so.

6. MASS STORAGE STATUS

Andrew outlined the current Castor situation. The upgrade to 2.1.7-15 had been made in August. The ATLAS/LHCb/CMS experiment instances had since had problems accessing data that had been written and this resulted in Castor being unavailable for extended periods between 14 August and 27 August. This was eventually solved by a parameter change in the Oracle configuration. Full details can be found at http://www.gridpp.ac.uk/wiki/RAL_Tier1_Incident_20080814.

There had also been further Castor problems with the LSF scheduler dispatching tasks extremely slowly, leading to delays in the receipt of ATLAS raw data. This happened on 10/11 September and impacted the LHC "red-button" day. Restarts and some throttling of the CERN load had fixed the problem. LHCb also reported that ~200 jobs had recently got stuck which resulted in a restart of LSF.

As a consequence of the (mainly) ATLAS problems, more auto-restarts were being deployed. It was also hoped that information flow to experiments would improve with the start of Gareth Smith as Tier 1 Production Manager from 1 October.

It was still planned to close dCache at the end of 2008 with unpaid access to the ADS terminating by the end of the financial year. ATLAS, CMS and LHCb had all migrated to Castor and MINOS is currently "in flight". The data from other experiments was considered small, but a list should be generated and circulated to remind all experiments of their data situation. Glenn suggested that there should be a graceful closure of dCache and ADS with data being marked inaccessible for 3 months before any actual catalogue deletions.

Status of Castor planning can be found at:

http://www.gridpp.ac.uk/wiki/RAL_Tier1_CASTOR_planning

Current Castor-experiment actions can be found at:

https://www.gridpp.ac.uk/wiki/RAL_Tier1_CASTOR_Experiments_Actions

7. GRIDPP3 QUARTERLY REPORTING

Glenn thanked experiments for the input to the Q2 report. As Q3 would soon end, he would be asking for further input on experiment success rates and examples of Grid failures for the next report. All of the "other experiments" are encouraged to keep statistics along the lines of MINOS:

http://www-pnp.physics.ox.ac.uk/~west/minos/grid_metrics/jun_2008.html

8. FUTURE FORMAT OF WEEKLY CASTOR/T1 MEETINGS

Following on from the discussion at the last UB meeting, it was planned that on the first Wednesday of every month a general Tier 1 meeting would be held in addition to the normal Castor meetings. The first such meeting will be on 1st October.

It was also envisaged that improvements to the current telephone conferencing would need to be made to make this more useful to external participants.

9. NEXT MEETING

The next meeting is scheduled for Monday 1st December at 10:00am using EVO (note unusual date/time!).

ACTIONS

Action 1: ALICE VO Box at Birmingham. Cristina to keep Glenn informed of issues and Roger to email GDB details to Cristina - **Cristina/Roger/Glenn**.

Action 2: Other experiments (i.e. not ATLAS/CMS/LHCb) to provide Glenn with any background on UK computing requirements for 2009-2013 by 3 October - **All "other" experiments**.

Action 3: Distribute remaining resource headroom to LHC experiments for Q4 - **Glenn**.

Action 4: Circulate list of data remaining in dCache and ADS and remind experiments of migration schedule to Castor - **Andrew/Matt/T1**.

Action 5: Advertise the new monthly Tier 1 meetings (part of the Castor series) and arrangements to the UB list (and beyond?) - **Andrew**.

Action 6: Update List of Experiment Operational Contacts on UB web pages - **Glenn**

Action 7: Q3 report on "other experiment" metrics - **"Other" experiments/Glenn**

Action 8: Prepare and circulate UB questionnaire on experiment support by end March 2009 - **Matt**