

# User Board Meeting, RAL, 20<sup>th</sup> June 2007

Present: *Jeremy Coles, David Colling, Matt Hodges, Roger Jones, Dave Newbold (Associate Chair), Raja Nandakumar, Andrew Sansum,, Glenn Patrick(Chair), Dave Sankey, Nick West, Fergus Wilson.*

Apologies: *Gavin Davies, Mike Kordosky, Paul Kyberd, James Ferrando, Simon Peeters, Dan Tovey*

## 1. PREVIOUS MINUTES

The previous minutes of the joint UB/Dteam meeting held on 22 March were accepted without alteration. Concerning actions:

- a) *Linear Collider*: Glenn had talked to Steve Worm who had agreed to act as the contact for LCFI and CALICE until someone else is identified.
- b) *Tier 2 storage accounting*: Jeremy reported that this was now stable, but waiting for developers to progress.
- c) *ADS/VTP tape service*: Andrew circulated an email to the UB on 16 April outlining the future options for experiments using the ADS/VTP service. RAL has no plans to terminate this service, but GridPP would not be funding it. This means that any experiment needing to continue with ADS beyond March 2008 would need to finance the storage costs from the experiment's budget. In the discussion, it was pointed out that SNO may wish to archive some of their data on tape outside of the robot. A cost figure was needed for this facility.
- d) *GridPP User List to be forwarded to the UB*: Stephen Burke had provided the details of this list which was set up for general user discussion/help (gridpp-users@jiscmail.ac.uk and see <http://www.gridpp.ac.uk/contact.html>). Glenn had put a link on the UB web page.

**Action 1: Provide costing for tape storage archived outside of robot - Andrew Sansum.**

## 2. TIER 1 USAGE FOR 2007/Q2

Matt said that the Tier 1 had seen 70-80% occupancy and this had been evenly split between Grid and non-Grid use. At 53%, the global CPU efficiency had been low in March (attributed to dCache and Castor problems), but had improved to over 70% in April and May. The indications were that further improvement would be seen in June (76.9% eventually reached).

	Jan 2007	Feb 2007	March 2007	April 2007	May 2007	June 2007
Efficiency (%)	63.8%	60.2%	53.4%	75.8%	71.5%	76.9%

Deployment of disk had been late for ATLAS. The NFS disk for UKQCD had been migrated to new hardware.

Monitoring tape use in Castor was difficult because of problems with repack. This meant that experiments (e.g. CMS) could appear to be way over their data allocations.

## 3. STATUS OF TIER 1 RESOURCES

A hardware procurement was underway to meet the CPU and disk commitments for 2007. The tender was launched about a month ago with the CPU expected to be available in February 2008 and the disk in March 2008.

An interim order for ~£50k of tape media had been placed to provide another 300TB of capacity.

A tender to set up a Framework purchasing agreement for future tape media over the next three years was being set up. This would start to deliver media from 2007/Q4.

A framework purchasing agreement was also being set up to deliver tape drives in November. The planning for tape drive/robot configuration was still undecided and Matt had been studying the experiment profiles.

## 4. REVIEW OF EXPERIMENT ACTIVITIES/REQUIREMENTS

### 4.1 Experiment Plans

CMS were conducting pre-tests for CSA07 until July, maintaining activity through August and with the full challenge in September for ~ one month. This would be followed by the global detector run in October. The real MC load for 2008 would become established from Feb/March 2008.

ATLAS would be conducting a full dress rehearsal test in October with another (fuller) rehearsal planned for Feb - March 2008.

LHCb were still completing the DC06 data challenge. Transfer tests would take place end July (T0-T1) and end September (T0-T1 and T1-T1).

BaBar activity was mainly focussed on analysis and the future of UK production/skimming activities was uncertain.

SNO still expected to finish activities on the Tier 1 by the end of 2007.

MINOS was continuing to ramp up activities.

H1 was making light use of the Tier 1 at the moment and had interest in using Tier 2 sites. The HERA machine would close forever on 30 June.

### 4.2 Experiment Requirements

ATLAS had no big changes to the figures from the last meeting. To understand the detector during the commissioning phase, the event size was increased to 8MB (instead of the standard 2MB/event).

CMS emphasised the importance of a ramp-up plan and not to just try and deploy everything in Feb 2008. The CMS figures were updated after the meeting (26 June) - to increase CPU and reduce storage in the September to November period. These can be found at: <http://www.gridpp.ac.uk/eb/200607/tier1expts2007june07b.xls>

LHCb had not changed their request since the last meeting. Extra temporary disk space needed for their transfer tests was thought to be within their existing disk envelope.

Glenn had agreed with DO that, due to under-use, their CPU request should be decreased to 30 KSI2K/month until activities justified an increase.

### 4.3 Experiment Allocations for 2007 Q3/Q4

Extra disk space was included in the spreadsheets to account for disk0tape1 cache in front of the storage systems. For dCache, this amounted to 25.6TB. For Castor, this currently amounted to ~106TB Castor, but at least two disk servers (18TB) could be allocated to experiments meaning this could be reduced to 88TB.

#### Cache estimates for d0t1

dCache	Castor
25.6TB	88TB

There was some discussion of what experiments routinely included in their disk requests. Both ATLAS and LHCb only include the space occupied by data. CMS put in extra components, but this may not include the overhead for their transfer buffers.

**Action 2: Experiments to list what is included in their disk requests from 2007/Q3 onwards - all experiments.**

Extra tape media space was included in the spreadsheets for disk1tape0 in ATLAS and LHCb because Castor automatically backs up such data to tape for this storage class even though it is not needed by the experiments. This was "guesstimated" from the disk requirements of each experiment.

Andrew said that he expected disk1tape0 to be working properly in Castor at RAL by September and hopefully this overhead can soon be discarded.

There was some discussion of repacking overheads for tape storage. In the longer term this should be less of an issue since data should be written at a steady state with few deletions.

There was also some discussion on whether data compression was used in the experiments.

**Action 3: Experiments to establish what compression (if any) is applied to data - all experiments.**

Glenn summarised the allocations for Q2 as follows:

1. CPU - For July and August CPU is OK. There is an over-allocation in September where ATLAS and CMS peak, but this may get compensated by the vacation/post conference period in actual utilisation.
2. Disk still has headroom through Q3.
3. With tape (and the vagaries of Castor), it was difficult to be certain of the physical headroom, so in this sense any allocations are nominal! The latest media purchase (assumed August) is the last until the end of the year, so experiments should take care to delete unwanted data and be vigilant when they write large amounts of data.

The 2007/Q2 allocations can be found at <http://www.gridpp.ac.uk/eb/200607/tier1allocs2007Q3.xls>

No indicative allocations were given for 2007/Q4 because of the uncertainties with the procurement and planning at the end of the year.

#### **4.4 Experiment Allocations for 2008**

Glenn stated the importance of now looking to the planning for 2008 and matching this onto the 2007/Q4 requirements in the light of the new LHC schedule. He would shortly be asking experiments to give an estimated profile of their 2008 requirements (as was done for 2007) with a view to getting a complete picture by end August.

**Action 4: Experiments to provide 2008 profile for CPU/disk/tape - all experiments.**

### **5. SL4/SL3 SERVICE**

Matt had already discussed with the LHC experiments their requirements for a SL4 service at the Tier 1. The other experiments made comments as follows:

- BaBar - Code runs on SL4. Not built. Would adapt to any Tier 1 changes.
- MINOS - Code runs fine on SL4 (at FNAL). Not an issue.
- SNO - Root based. Would need SL4 queue to build.
- H1 - Moving main production to 64 bit SLD4 (Scientific Linux DESY).

It was also noted that LHCb had requested a dedicated CE for SL4 so that specific platforms can be efficiently matched with jobs.

SL4 worker nodes would be set up on the Tier 1 so that experiments could test their code.

Tier 2 sites appeared to be waiting for the native SL4 WN and UI.

### **6. CASTOR**

Glenn did not want to revisit the full history of Castor, but had instead asked David Corney to prepare a revised timetable of Castor plans which can be found at:

<http://www.gridpp.ac.uk/eb/200607/castorplansjun-oct07.doc>

There is also a web page devoted to Castor Experiment Technical Issues at:

[http://www.gridpp.ac.uk/wiki/RAL\\_Tier1\\_CASTOR\\_Experiments\\_Technical\\_Issues](http://www.gridpp.ac.uk/wiki/RAL_Tier1_CASTOR_Experiments_Technical_Issues)

The major item to note is that there are going to be separate Castor instances for ATLAS, CMS, LHCb and the production platform. As part of this plan, the LHC experiments will be migrating from version 2.1.2 to 2.1.3 at the earliest opportunity.

MINOS said that the Castor service had been lost for one week and asked if version 2.1.3 was going to be any better. Andrew said that it looked to be a significant step forward as the CMS tests looked promising and advances had also been made in CERN discussions and securing new contractor effort at RAL.

## 7. FUTURE OF dCache

It was agreed that, due to the problems encountered with Castor, the original timescale of terminating the dCache service on 30 June was now unrealistic.

It was proposed and agreed that dCache would not be terminated until at least the end of 2007 and that, in addition, six months warning would be given of the eventual closure date.

## 8. GRID ONLY ACCESS TO TIER 1

Andrew reminded the meeting that there was a milestone for the end of August to offer only Grid access to the Tier 1. This had been discussed at the January UB meeting (<http://www.gridpp.ac.uk/eb/040107/nongridaccess.pdf>) and at that time it was felt that the proposal was premature and the milestone needed to be adjusted in the light of further experience.

It was agreed that since the Tier 1 could not yet offer a full Grid service (e.g. SRM enabled Castor) that this milestone should be delayed until the first quarter of 2008.

MINOS emphasised the need for a UI to continue on the Tier 1.

It was agreed that Andrew would produce a technical document on the full implications of a Grid only service and circulate it to the UB so that an informed decision could be made.

**Action 5: Provide technical document outlining Grid only service at Tier 1 - Andrew Sansum.**

## 9. AOB

Nick West raised two points concerning MINOS:

- The RAL Tier 1 (190GB, 84% full) and Tier 2 (162GB, 90% full) software disks are quite small and need to be expanded. This would be looked at, but there was not thought to be a problem with the Tier 1 disk and CMS had moved from the shared disk.
- There is a need to replace the *minos VO* with a new name called *minos.vo.gridpp.ac.uk* to avoid a clash with the *minos VO* hosted at Fermilab. Nick was encouraged to discuss this with Sergey Dolgobrodov (VOMS management, Manchester) and to keep Jeremy, Andrew and Glenn informed to ensure all the implications were understood.

## 10. NEXT MEETING

To be arranged before the end of September. A non-RAL venue (e.g. Tier 2 site) would be considered.

## **ACTIONS**

*Action 1: Provide costing for tape storage archived outside of robot - Andrew Sansum.*

*Action 2: Experiments to list what is included in their disk requests from 2007/Q3 onwards - All Experiments.*

*Action 3: Experiments to establish what compression (if any) is applied to data - All Experiments.*

*Action 4: Experiments to provide 2008 profile for CPU/disk/tape - All Experiments.*

*Action 5: Provide technical document outlining Grid only service at Tier 1 - Andrew Sansum.*