UB/DTEAM MEETING ~ F2F ~ 22 March 2007
Glasgow University

UB Chair: Glenn Patrick // DTeam Chair: Jeremy Coles

UB F2F

Present: Andrew Elwell, Graeme Stewart, Greig Cowan, Alessandra Forti, Pete Gronbech, Olivier van der Aa, David Colling, Derek Ross, Jeremy Coles, Andrew Sansum, Mike K, Nick West, David Kelsey, Catalin Condurache, Robert Fay, Stephen Burke, Stuart Wakefield, Roger Jones, James Ferrando, Glenn Patrick, Matt Hodges, Dave Kant, Raja Nandakumar

Apologies: David Evans, Gavin Davies, Paul Trepka

1. Minutes of last Meeting/Actions Arising

- linear collider: GP to do
- storage accounting: JC to do - it was noted that there was now an automated way of getting details for storage accounting but there was still work to be done
- tape access/VTP access: MH to do - it was noted that there was no schedule at present, information would be required over the next month regarding plans, worker node issues depend on hardware and it was noted that pool space differed on different generations of hardware - the fabric team could assist with this but it was work-in-progress. Details could be recorded on the Tier-1 website - all the nodes should eventually have 60-100 gig and this info should be published.

2. Status of Tier-1 Hardware Deployment & Usage Report

- It was noted that there was a lot of disk deployment at present, but CPU was flat at 1200 job slots. Total disk available is about 360TB, the next set of disk (+250TB) is starting to be deployed into dCache at LHCb to meet the March allocation. Some of the disk is to go into CASTOR for CMS, also to meet their allocation. ATLAS have 36TB in dCache and CASTOR.
- The second new tranche of disk will be accepted, bringing the total available to 800TB, which is in excess of requirements and relates to 2nd quarter '07.
- The fabric team at RAL are working on kick-start configuration for CASTOR, it should be simple to deploy more disk into CASTOR.
- Usage of CPU was at 80% occupancy of Farm for Jan/Feb and CPU efficiency is an issue. For CMS, employing extra disk into CASTOR is also an issue; dCache has been tested and been shown to work.

3. Review of Experiment Requirements & Allocations

It was reported that regarding the allocations for Quarter 2 at Tier-1, there had been requests for increases from BaBar, MINOS and a slight variation for LHCb. On the website, Tier-1 requests have been revised to show '07 requirements - capacity and requests have been highlighted and headroom noted. For Jan-Mar there is excess capacity, for Quarter 3 and 4 things get worse. Hardware procurements will be re-estimated this year to meet '08 requirements. It was agreed to increase MINOS for Quarter 2 but reserve judgement for Quarters 3 & 4 meantime. It was noted that an extra TB on Disk1Tape0 at CASTOR was required, however there was the possibility of NFS space being made available: the oldest hardware is due to be decommissioned at the end of March with space possibly being freed-up by mid-April. A CASTOR upgrade is scheduled which might also help space issues. It was noted that 10TB may be available in advance of allocation.

Dr S Muir Scott / s.scott@physics.gla.ac.uk

3.1 Disk

It was noted that disk will become an issue towards the end of the year; there is duplicated data in CASTOR and dCache. LHC need to provide estimates of what is required. RJ will organise this. It was agreed that GP and RN will do the same for LHCb.

3.2 Tape

It was noted that the headroom was negative for most of the year. We can only use certain tapes in CASTOR. AS will investigate this issue - tape is difficult throughout the year.

3.3 Tier-2 Planning

It was noted that BaBar had submitted a large request - there don't seem to be problems in meeting it. It was agreed that the Tier-2 Board should confirm this.

4. Castor Migration Experience

The discussion centred on migration and timescales. RJ provided an update for ATLAS. They were still in transition, no problems had been experienced at RAL but there were some at CERN. Timescale issues will be discussed over the next week.

RN noted that for LHCb the disk space was adequate and there had been no problems in transferring files into CASTOR. They had tried to run test jobs to access the data - access is possible but there are problems with the job environment. They will be running simultaneous jobs and it will take around 3 weeks to have a full system operational. Once they declare CASTOR to be a production system, all jobs will treat it as the primary storage element, then data will be moved from dCache. There is 45TB of disk and 38TB of tape to be moved.

For CMS, no-one was currently available to report, however performance problems had been noted. They have already migrated, although the system needs tuning.

Re MINOS, the first TB of allocation had been received a month ago and multiple problems had been experienced. There had been low-level errors and problems with ticketing, testing had started but it was a learning experience - core rates were good.

It was noted that the dCache deadline closes on 30th June although this may need to be revised. Experiments are doing what they can to test the system.

It was asked whether there were any plans for tape purchase during the year? There was a discussion with MK re the tape UK Tier-1 rquests - tape numbers need to be analysed by GP and MK, migration will be a problem without tape space and it needs to be discussed with the Experiments. Action on GP and MK to rationalise the tape numbers and understand them better.

5. AOB

With reference to using the Grid from the small Experiments point of view, the following points were raised:

Dr S Muir Scott / s.scott@physics.gla.ac.uk

1. Having a Grid Liaison Officer would help the Experiments by providing a first point of contact for grid documentation, best practice etc. The Grid is perceived to be very large and difficult to know what procedures are. It was noted that Stephen Burke already performs the role of 'Grid Liaison' to some extent at present, and would be happy to respond to enquiries.

2. Ring-fenced NFS disk and space availability issues can give Experiments redundant NFS space? Or would it be possible to give them the space they need in CASTOR?

It was noted that authentication of documents was an issue but JC pointed out that there is a Documentation Officer in GridPP (SB) who could respond to generic questions and assist with information-flow. RJ noted that in future the Technical Co-ordinator could take on this role. There was a question about best practice and advice - it was noted there are standard routes: contact SB for Documentation, JC for deployment, etc. This feedback was noted to be useful as the smaller Experiments are not aware of the larger GridPP structure. There was a general issue here concerning unsupported small Experiments who may have difficulty in finding out things or knowing who to contact. There was discussion of a public mailing list and user support. It was agreed that the GridPP User list is to be forwarded to the UB. SB noted that there are also Instructions on the website. SB and GM to provide options to the UB on the way forward here, see if proposals work and report-back next time.

The date of the next meeting via VRVS was discussed and settled at Wednesday 20th June @ 2:00 pm.

Actions as at 22.03.07

1. RJ to organise estimates from LHC regarding disk requirements for this year.

2. GP and RN to organise estimates from LHCb regarding disk requirements for this year.

3. AS to look at tape issues for the remainder of '07.

4. GP and MK to rationalise the tape numbers.

5. The GridPP User list is to be forwarded to the UB. SB and GM to move forward the issue of small Experiment contact points.

DTEAM F2F

JC noted that Experiment issues had been discussed in the morning session and for the afternoon session it was hoped to discuss LHC plan and startup for 2007, and Experiment input would be needed for that.

1. Review of Experiment Requests & Allocations

JC asked what the Experiments needed from sites and what were their timescales?

LHCb

RN noted that all the Tier-2s will be mapped to a Tier-1 site and the data produced uploaded. Tier-2 will need to manage an upload to Tier-1, all other arrangements are via the VO box. RN reported that if one job runs for 2 days, it takes 10 minutes to transfer the data to Tier-1. JC asked whether

Dr S Muir Scott / s.scott@physics.gla.ac.uk

fail policies are understood? RN noted that if Tier-1 is unavailable they try others in rotation, if a job fails, it aborts, and exits. If the Tier-2 network is down, the job is marked as 'stalled'. GP noted that this had been done during the 'dress rehearsals' but a lot also depends on the GridPP3 action plan. DC06 has been ongoing since last May and is still not finished, but there wasn't much planning left to be done. There had been no simultaneous operation as yet, therefore the system was untested. GP noted that the Tier-2 resources in the UK can produce a lot of data, which might swamp the Tier-1 - however it doesn't matter where the data goes to within Tier-1. RN noted that Manchester and QMUL were already mapped and uploaded to CNAF at RAL. JC asked whether the Tier-1s know the mapping and expect data designated from LHCb? Yes - there had been an 'in readiness' review and data rates - it seemed relatively straightforward.

JC asked whether there were any other issues, such as worker node disk requirements? RN noted that a standard 1GB of scratch space was required. There followed a discussion about data rates. It was noted that the figures were linked to the Tier-2 pages. Not all of the documents were in the same format and the documents were incomplete. In terms of LHCb planning, there was no significant work to do. There was a question regarding memory per core? For LHCb the maximum memory needed for reconstruction was 500 MB per job.

ATLAS

RJ reported on reconstruction and pilot jobs. 1.5 Gb per core was needed to meet higher requirements. For pilot jobs some nodes are required with higher memory. It was asked whether a job can be mapped to site by size of memory? It was noted that a large part of capacity will be used for reconstruction. Work is ongoing on reprofiling. JC asked whether it was possible to identify major ATLAS sites that might have a problem? What do sites have at present? Glasgow has 2 GB per core, upperbound on requirements. Edinburgh probably has 2 GB per core. Durham has no disk. QMUL is a problem. Imperial has 1 GB per core. LeSC is a problem. Brunel has 1 GB per core. JC noted that nearly all of the London Tier-2 have 1 GB per core. SouthGrid has 1 GB per core but has new kit. Cambridge and RAL have 2 GB per core. Liverpool has 1 and Manchester 2, Lancaster 1, Sheffield 1. The summary was that a lot of sites are under requirements and the issue needs addressed.

CMS

It was noted that CMS are happy to run at 1 GB and this should be adequate for the foreseeable future. They are inside the GB limit for 90% of their jobs, therefore it is not an issue for CMS. It was noted that for the other Experiments, there were also no memory issues.

2. 2007 Planning - now till LHC startup

Regarding the ATLAS proposal, it was noted that changes are required across sites. RJ noted that plans have been made to the end of '07 but there was work still to be done at Tier-1 and Tier-2. One issue is streaming files and access. It was noted that at Glasgow, RFIO tests to DPM had shown good rates, however, opening lots of files simultaneously had been a problem. DPM had been installed at the end of last week. They were trying to run Athena test jobs but this was ongoing - RFIO root plugin was a problem and progress was required. There may be issues at larger sites.

The CMS figure at Tier-1 is 1 MB per second; 10 MB per second will be difficult. GS noted that this was an issue for the DPM sites - they were getting decent ATLAS data but DQ2 transfers were going to the wrong place in the wrong pool. It was noted that ATLAS dress rehearsals were running in stages and would be rolled-out in phases: phase 1 by the end of June, phase 2 by end

Dr S Muir Scott / s.scott@physics.gla.ac.uk

July, phase 3 by end Sept, and phase 4 by Oct/Nov. There will be someone from ATLAS at each site.

Regarding CMS planning, SW reported that there was a 50% phased introduction with load tests and data going to Tier-1 from Tier-0, then Tier-1 to Tier-2. The data transfers were fine, they were hitting minimum rates to all Tiers. There was a problem with CASTOR running locally at Tiers, they had configuration and load problems. It was hoped that an upgrade would improve this. Tier-1 to Tier-1 transfers had not been exercised much before, they were managing 3 MB per second to each Tier 1. It was noted that there was a manpower issue with RAL PPD to be sorted out internally with CMS. PhEDEx requirements can be done from Imperial. The only issue is support at RAL PPD, but all else seems ok. It was noted that a ramp-up to more sites was likely and a clearer picture will have emerged in a month's time. The other goal was to move Monte Carlo production away from the Tier-1s, to be done over the year. Processing, re-processing and skimming would be done at the Tier-1s. MCP would be phased-over to Tier-2.

Regarding other Experiment VOs, were there any planning issues? Did they want to move jobs to Tier-2? It was noted that a production system was required and software needs to be developed for more convenient use of Tier-1. The MINOS production system should be implemented by the end of the year.

DTeam Transfer Tests: Regarding the network bandwidth tests combining and doing transfers, what was the planning/the way forward? It was noted there were bottlenecks from Tier-1 due to the disk servers allocated to DTeam testing - what were we trying to achieve? There was a difference in the way sites handle the storage elements. AS queried whether Experiments could measure their achievable rates? FTS tuning was required. There are baseline targets but the process itself identifies difficulties. JC noted that better co-ordination and communication during testing was required - the mechanism is the weekly DTeam meeting as a good forum to discuss such issues.

Fair Share Policies: Do sites have them? Manchester operates this policy but are fairshares set up to meet requests? It was noted that Glasgow and Edinburgh have fairshares reflecting the interests of each site. Glasgow also has good CPU resources. Durham do not enable fairshares, these have been disabled due to batch system failures. ScotGrid should be able to operate them ok and adjust according to need. SouthGrid - all sites except Cambridge have fairshares and the maximum CPU limit is set to enable capacity. For London (QMUL) it is first-come-first-served - fairshares are set but local interest is catered for first: LHC has priority. JC asked whether we can implement fairshares based on UB recommendations? It was noted that the Job Priorities Working Group gave recommendations, and that fairsharing between users was also an issue especially for regular users. The CMS default 50/50 policy was between analysis and production.

Regarding scaling the Tier-2 hardware resources, it was asked whether Tier-2 can meet the requirements?

JC asked if there were any other issues to raise with respect to planning LHC startup? There were no other issues.

3. OS Migration

Regarding SL4 worker nodes, what was the timescale etc? Experiment software was 32GB vs 64GB.

For ATLAS RJ reported that the next software release would be available next week and was compatible with SL3 & 4, and can be chosen according to worker nodes. There were however

Dr S Muir Scott / s.scott@physics.gla.ac.uk

compatibility issues with SL4-3, 32GB guarantees to be validated, but 64GB not yet - testing was still to be done. It was noted that SL3 would be ok until October, but by then SL4 needs to be available, and there were compatibility issues. It was asked when Tier-2 sites would move to SL4 - Sept? September was considered too late, it needs to be done sooner.

For LHCb, it was noted that SL3 & 4 works on all machines - they will use SL4 if possible, 64GB machines can run SLC32GB. There were random number generation problems and it was hoped that once these were fixed, SL4 will be ok. The plan was that over the next few months SLC3 will be supported (and for the remainder of the year), but next year SLC4 will be used. Was it a specific requirement to run SL4? No - if the site has both machines then separate OS configurations are required - gLite middleware is not SLC4-compatible yet.

It was noted that CMS are building on SL3 at the moment, SL4 is wanted but is not available yet.

The meeting broke for lunch.

4. Site Management and Performance Issues

JC asked if there were variations in site efficiency - experiments had work over the past few months at Tier-1 and Tier-2. Comparing wall clock and CPU time gave ratios.

Observed Efficiency Variations:

CMS were at 76.3% for March, efficiency here has been falling and this was dominated by CASTOR issues. Was the fall-off in November and December due to low status tics? CMS had reconfigured CASTOR in the last month or two and had fewer servers, therefore issues relating to this were currently being worked-through - these were mainly hardware and CASTOR configuration.

For LHCb, efficiency has moved from 38, to 42, to 54% - things are improving. It was noted that for most of last year efficiency was abolve 84%. In December, dCache became unstable and softlinks are a problem. Across London, it was noted that Tier-2 was also quite low and was related to the above issues.

For ATLAS it was noted that efficiencies also dropped, and this was directly related to dCache, efficiencies were down across the board, possibly also related to code changes.

MINOS was steady at around 70-80%, through '06 it was noted that it was low at periods but not much work went through at those times. ZEUS went from good efficiencies early on to low efficiencies in September (there were low levels of work). The current status is low and ZEUS efficiency has gone down. It was noted there was no dCache allocation at RAL. BaBar efficiencies were deteriorating due to doing tape access. It was noted that NorthGrid had stable efficiency.

It was asked what DTeam should be checking? A list should be put in one area for SysAdmin to check. A list of monitoring links was required. At present ATLAS was doing production monitoring, but there was none for CMS - CMS will be added to Wiki. It was noted there was also nothing for ALICE. The smaller Experiments do not have formal monitoring.

It was asked whether the SAM test were being implemented? Yes, SAM tests were running for LHCb now and it was noted that it was now possible for each Experiment to implement its own SAM tests. Did CMS have exclusion parameters?

Dr S Muir Scott / s.scott@physics.gla.ac.uk

Use of Disk:

It was noted that generic pools were moving to dedicated setup. As knowledge increases regarding the SAM set up, the disk servers will be assigned to VOs. It was noted that in DPM there is limitation - a pool can have all VOs or a single VO can sign-in - a future version will enable specifying. We can't currently do quota-ing either - some of the disks are too full. Clearing data from the storage elements is required. A mechanism is required by the VOs to clear data from their own areas. Experiments' VOs need to monitor disk usage.

Regarding SRM2.2, one ATLAS pool could be possible with permission tags - it is not clear how to operate this. At Tier-1, ATLAS is full in dCache.

Local vs Grid Usage of Resources:

At LHCb, all usage is Grid usage as far as possible. For CMS, the recommended model is Grid access, this is so also for the people working at CERN and RAL. An audit has been done - even when non-trivial, there is an amount of non-Grid usage. The ATLAS problem was least problematic.

Experiments: Problems at Sites:

Most of the observed efficiency was down to CASTOR, there was a problem with LHCb at the London Tier-2. It was asked whether something had changed in the ATLAS software to make it less efficient? Efficiency of 90% in June/July fell away to 50% in Oct/Nov. RJ reported that the data movement tools were a problem, if the job were hanging and retrying it would affect efficiency, if it doesn't shift, it times-out and retries. Input from the Experiments was required - are they seeing specific problems at sites, eg: ATLAS at UCL? It was noted that at Brunel, during scheduled downtime, jobs abort. At UCL Central, release date is unknown, but should be finished by next week. At Liverpool an old release of ATLAS existed. The same problems are not experienced at CMS and LHCb. JC asked if tests were being run to see if sites are targetting jobs?

SE availability:

Was there a problem at Liverpool and UCL from the Experiments regarding SE availability? This was not an issue for LHCb. CMS experienced occasional difficulty. Do the Experiments raise storage issues with sites? Yes - directly through GGUS. For ATLAS it can be both: direct contact as well as GGUS.

Unavailability of Required Software:

Whose responsibility was this? It was the responsibility of ATLAS to install software and keep it up-to-date. Regarding procedure, should a GGUS ticket be raised regarding software upgrade issues? For ATLAS ACL change it was noted that sites can't test themselves and then discover afterwards that there is a bug. There was no co-ordinated rollout and it was not properly tested. How do we deal with configuration change requests? RJ asked whether this went to the EGEE team? SRM2.2 should help some of these issues to disappear. If ATLAS were to request another change to ACL, how should they deal with this? Someone should take charge and run tests - if we had said no to doing an ACL change, ATLAS would not have been able to do production for the UK. It was agreed that we can't delegate a VOMS role for ACL. Sites could be targetted for production work. It was noted that GGUS tickets had been raised. It was reported that Glasgow would be happy to test on everyone's behalf next time, but this is an informal/unofficial structure. It

Dr S Muir Scott / s.scott@physics.gla.ac.uk

was noted that we also need to decide how to deal with EGEE broadcasts. It was noted that during the database backup/ACL change, no advice had been provided.

It was noted that Python 2.4 could be an issue (was this standard in SLC4?). It could be 2.3.4. Had there been any other packages/requests? UI at RAL was not up-to-date, SB had sent in a ticket over a year ago. It was asked how many people relied on RAL UI? MINOS and UCL Oxford. Were MINOS not using the VOMS server at Manchester? They were trying to move to VOMS at Fermilab.

Support for Smaller VOs:

Was this lacking anywhere? Issues of documentation and contact had been discussed earlier. Were there any other issues needing support? No specific requests were forthcoming. The UB is not usually involved in MINOS, or snow?, or CDF. There were communication issues. There were problems at Cambridge, Sheffield, Liverpool, Bristol - QMUL has 64MB migration - different versions can cause jobs failure. We can't discover which version LHCPI? talk to.

5. Experiment T1/T2 Engagement

Representation at Sites:

On the DTeam there should be one person from each Tier-2, the Tier-2 Co-ordinator plus a specialist in storage. Each Experiment should send someone to the DTeam each Tuesday as this is a useful forum to deal with problems. There is a list of contacts on Wiki to help with installations or problems. It was noted that the smaller Experiments need a contact point for each Tier-2 - disconnect has hindered work and there has been little exchange of information. Perhaps a top-level webpage is required?

Areas that Need Improvement:

It is possible to email the UB with appropriate links - how else could we improve response to Operations?

GGUS and Ticketing:

It was noted that Tier-1 has its own ticketing system. Ticket numbers are now preserved for identification and closed at different levels. It was noted however that tickets don't reach the places where they should, and are not responded to. Problems should be raised with the site but also with GGUS. RN noted that at RAL Tier-1 - if he opens a ticket to GGUS, should he open a local ticket as well? This doesn't automatically translate - tickets are looked through weekly and assigned. If a ticket is raised twice (at two different places), this causes confusion and extra work. It was noted that local and GGUS systems have the same priority but one ticket is tracked at various levels. It was recommended to use GGUS. SB noted that it was not the system which was at fault, rather it was an issue of the manpower required to run it. It was noted that GGUS has improved, but there are difficulties with the support route behind it. The function is not fully defined and it is not a 24-hr service. The practicalities are that nearly everyone is based in Europe, and the initial assignment of tickets is different from the support given.

The meeting closed at 3.00 pm.

Dr S Muir Scott / s.scott@physics.gla.ac.uk