



## Version 3.0

*Amendments history:*

<i>Name</i>	<i>Area</i>	<i>Date</i>
<i>Jeremy Coles</i>	<i>All – draft notes</i>	<i>05-07-06</i>
<i>Jeremy Coles</i>	<i>All – full edit</i>	<i>07-07-06 version 1</i>
<i>Kors Bos</i>	<i>All – full edit</i>	<i>10-07-06 version 2 created</i>
<i>Dario Barberis</i>	<i>8. LHC Schedule</i>	<i>11-07-07 version 3 created</i>

## Minutes of the meeting

CERN, 5<sup>th</sup> July 2006

Agenda: <http://agenda.cern.ch/fullAgenda.php?ida=a057708>

GDB twiki page: <https://twiki.cern.ch/twiki/bin/view/LCG/GridDeploymentBoard>

Minutes: Jeremy Coles

Attendees: Please refer to list at the end of the minutes

### Meeting Summary

see GDB wiki page: <https://twiki.cern.ch/twiki/bin/view/LCG/GridDeploymentBoard>

### Detailed minutes

## 1. Introduction (Kors Bos)

Kors began the meeting by reviewing actions from the last meeting (see action list on twiki for current status).

John Gordon made a brief statement regarding the election of a new GDB chairman. At the last meeting Kors was confirmed as continuing in the role with the caveat that if anyone objected (following the last meeting and an email announcement of the result) they were to contact Les Robertson before today's meeting. Les did not receive any objections therefore Kors is indeed reappointed.

Kors continued his introduction by stating that it was now 2 years since he began the role. In that time only two votes have been taken – once for the first chairman and secondly for



## LCG Grid Deployment Board Meeting



the second chairman! He reemphasised the need for national representatives attend these meetings. On reaching slide 6, he asked is the current mode of operation how we want to continue. Ruth had asked about short reports being given which is something that used to be done. This is open for discussion. He asked does anyone want a change to the format or style. No feedback.

Kors has received little feedback so far on the dates suggested (slide 8) for future meetings. It is proposed to go to Prague in April – VRVS or a phone connection is a pre-condition. Jeff asked if March meetings can be on a Monday or a Friday. Kors replied this is an old topic.

The 5<sup>th</sup> September 2007 meeting is planned around CHEP in Victoria. The weekend before is tentatively booked for a Technical Meeting before. The next meeting this year is 6<sup>th</sup> September at BNL – given the distance it is suggested to have a 2 day meeting – Tuesday and Wednesday 5 and 6 September. Kors asked for help making arrangements since he is on leave just before the meeting. [Ruth agreed to help out.](#)

[Jeff observed that one thing missing from the agenda is interoperability](#)

[Action 0607-1 Kors to circulate link for BNL registration page](#)

[Action 0607-2 All members to register for the BNL meeting if attending.](#)

### **CERN migration to SL4:**

Tony Cass reported that there are now a couple of LXPLUS and batch nodes running SL4. It is his hope that we can switch in the autumn. Mails are arriving asking for SL4. This was discussed at the last meeting - the priority remains for worker nodes and then 64-bit.

Ian Bird gave an update: For the WNs some sites have already installed 2.7.0 and will try this for gLite. This requires some juggling of rpms. The deployment group now wants to move to SL4 64-bit but this requires the newer version of VDT. The build system now in place is at version 1.3.9. There are issues. Ian expects a distribution **on the October timescale**. If people are concerned about WNs then they can do it by hand now. BUT any move to SL4 needs effort which takes away from the testing and certifying of other pieces of middleware.

Kors: Yes but there is a pressure to move now. Ian: This is best done without adding new functionality but there is a pressure to do this. Jeff: Such as? Ian: You have seen the list.

Holger asked if a link can be added to the minutes pointing to the recipes

[Action 0607-3 Ian to request link from Michael Jouvan and provide for inclusion in the minutes](#)



## LCG Grid Deployment Board Meeting



John mentioned that this is an example of a technical issue that is not of concern to the experiments [the comment relates to some discussion at the MB on Tuesday 4<sup>th</sup>]. This is something for the EGEE Technical Coordination Group (TCG) representatives. Relative [TCG] priorities need to be set taking account of this input. Ian agreed and said moving to the newer version of gridFTP is another example. John: We don't see a complete list and nor does the TCG Ian: The TCG is dominated by users who want functionality so even if these areas are known they are of less interest to this group as a whole. **The TCG needs more site representation.**

A link to the current list of TCG areas/issues is:

the TCG PoW: <https://twiki.cern.ch/twiki/bin/view/EGEE/TCGPoW>

the priority list: <https://twiki.cern.ch/twiki/bin/view/EGEE/PriorityList>

Tony asked, should we [sites] wait or try workarounds now? Ian replied that if sites follow the workaround route there is no support – it is not being tested for hidden dependencies within RPMs.

Gonzalo Merino noted that CMS CSA06 activities are scheduled for October. He asked if CMS will be happy if resources are at that time provided with SL4? Ian: Migration is more of an issue. For RHEL-7.3 etc. this was a mess and people were still using it one year afterwards. We need to agree on the migration. We'll need to support two versions in parallel. Holger suggested that sites could enable some nodes on the PPS for experiments to try - centres would like the experience.

Dario Barberis (ATLAS): We are this week finalising a production release for SL3 and then putting in place infrastructure for the next release (Sept-Oct timescale) which will be a release for SL3 and SL4. So the October timescale is approximately correct. There is a change in compiler which is more important than the OS itself. The release will be available in 32-bit version. Validation of 64-bit mode will take a little longer. In the last transition there were sites with mixed setups – service nodes and WNs at different OS. [Dario would like to encourage consistency or the setup of a second CE at sites supporting both versions. There is a need to make a migration plan.](#)

NickBrook: LHCb are in the process of migrating. Using an LXPLUS machine at SL4. LHCb software will be migrated on the timescale of October. No problems expected.

No comments from ALICE or CMS.

Kors: Following discussion with Jeff, what is the NIKHEF official supported OS (slide 12). Ian: RHEL3, SLC3 and SLF3 were thought to be clones but differences were discovered. The distribution works across CENTOS. However, Debian and SUSI are different. Almost binary compatible but local fixing has been required.

Jeff: Does anyone know about SLB compliance? It has standards for start-up scripts etc. I also disagree with the “porting” label; migrating from one set of systems to the next should increase the supported platforms. Ian: I have tried to argue for this. Test systems



## LCG Grid Deployment Board Meeting



tend to be focussed. New code is often not 64-bit clean. JRA1 and other developers need to worry about this. They are now making a build system on 32-bit and 64-bit and will reject code which is not compatible with both. More sites should get involved with the testing and certification. We need to *force* developers in this area as asking has not worked.

**Action 0607-4** Ian to provide list/contact for how to get involved with the middleware certification activity. Mail to GDB or specifically Jeff Templon.

### **Storage Class implementations**

A task force has been formed. It will (need to) have phone meetings soon. Kors would prefer someone else other than him to chair the meeting – this person will also need to call meetings and take minutes. Examples of questions the group will examine are shown on slide 14.

Nick asked about the timescale for the final report. Tony said SRM V2.2 is to be delivered in November. Kors added that September/October is the latest this can happen. Tony clarified that this output will not influence 2.2. The group are looking at how sites setup storage to support the defined classes (WAN/LAN access to disk pools, impacts for dcap etc.) John thought even this could reveal showstoppers such as a need for more cache disk. This may not be a problem now but will be in 2008. Tony thought this an implementation question not something impacting 2.2. **Kors concluded that the group needs to have a first report for the BNL meeting**

**VO Box conclusions** from the last meeting were reviewed. For xrootd it was noted that no name is against the evaluation and performance task. Ian thought xrootd is not an issue for CMS and should be removed from the list as a VO Box issue. Tony said this is a question not just about if it works but **how will it be supported long-term?** For the shared credentials issue Jeff mentioned that patched installation scripts are available at NIKHEF and these will be moved to YAIM. Markus is aware of this work.

On the topic of **Super Computing 2006** it was asked if those attending would like to have a joint demonstration. The booths are quite close together: [http://sc06.supercomputing.org/pdf/eh\\_floorplan.pdf](http://sc06.supercomputing.org/pdf/eh_floorplan.pdf). No specific action at this point but someone will need to coordinate this if it happens. Dave Kelsey asked about GIN as he thought they were aiming at SC06. Ruth understood from a TeraGrid meeting that their demo would be ready soon. NorduGrid will be present and interested in participating. They would like to show interoperability from ARC to EGEE.



## 2. LHCOPN Information (Dave)

<https://twiki.cern.ch/twiki/bin/view/LHCOPN/WebHome>

All Tier-0 to Tier-1 are in place, provided by either the European NRENs, ESNet, LHCNet, APNet or a combination of those. Currently DANTE is delivering dedicated fibers to connect CERN to the European Tier-1 centres but only some are really operational at this time. The situation of intercontinental fibres is sufficient but new contracts will be purchased soon. Cross Border Fibers are looked into to serve as a backup solution for the Tier-0 to Tier-1 connections but also for the Tier-1 to Tier-1 traffic. DANTE is currently defining an operations model, the E2ENOC, to be in place end 2006. This is more generally needed than just for the LHC OPN.

Les asked in case of problems, who initiates new routing if it is required? David answered that rerouting is automated but there may be a degradation of the service. At the user level watching this becomes a role for the ENOC (EGEE SA2). Les asked how relevant this ENOC group is to grid operations. David did not know. They will have a view at the level-3 and will know where traffic is flowing. Debugging network problems is a different issue and will be dealt with by the E2ENOC.

Kors: how does US link into E2ENOC. David: DANTE will take on overall responsibility for monitoring European and non-European links. It should be noted that transatlantic bandwidth is 10-30 times more expensive than land based links.

Jeremy Coles asked whether all Tier-1s would have a backup link. David said that there are generally 2-links which GEANT provides via different routes. This does not take into account cross-boarder links. So far only Karlsruhe has indicated they will procure a (full) backup link. John asked about the nature of the cross-boarder links and David indicated that they are site to site dedicated links.

Kors: There was mention of (high-bandwidth) T2s to be looked at in the CMS context. Are there any for ATLAS? Dario: There will be some such sites. Kors asked Dario to ensure that David and Don are aware of these.

## 3. Policies (Dave Kelsey)

Many policy documents need updating as they get out of date very quickly. Much work involved but only a few people to do it. Feedback is still required on the accounting policy proposal from the June GDB. And Dave Kelsey will ask for approval or comments before July 14 by email for the CA policy document.



## LCG Grid Deployment Board Meeting



Ruth asked about the situation with Shibolith. Will WLCG only accept CAs that are accredited? John asked for clarification – this is to cover the case where Shibolith issues a short term certificate Dave answered that currently in the production grid we only have the approval mechanisms given on slide 9.

Kors pointed out that previous policy document review requests had failed to get feedback.

Dave mentioned a new requirement the UK from the GridPP oversight committee who represent the funding body PPARC. They want an indication of usage divided between UK and non-UK users! He asked if there was a wider requirement like this. Jeff thought the project needs to allow for more granularity than users – for example production vs analysis usage. John wondered if this would show anything meaningful as often the experiments use a single user for production work. Clearly there is a need for better granularity for debugging (who used all the quota!). VOMS roles and implementation should improve this area. Stefano Belforte added that accounting by VOMS group is definitely required by the experiments.

**Action 0607-5: Kors to mail GDB for feedback on accounting/policy**

**Action 0607-6: Dave Kelsey to mail the GDB list asking for comments/approval of the CA document**

**Action 0607-7: All to comment on/approve CA document by email before 14<sup>th</sup> July**

### **4. Accounting update (John Gordon)**

GIP = Generic Information Provider

Markus asked what is meant by used space. The SRM allows booking of space. John referred to issues with allocation versus used with the latter currently giving problems because of double counting on shared disk pools. One way to approximate would be to divide such shared disks and space used by the number of VOs using the area. There was a general comment that the Management Board (MB) are not really interested in tape storage but there could be a manual update on used tape space.

John asked if any sites currently use resilient dCache. Apparently Brookhaven and some US Tier-2s do. It is required for some UK Tier-2 sites which have put most of their storage on the WNs.

There was a comment about access control Jeff pointed out that SAs have access based control (i.e. the GlueSchema does allow it now), the problem is with the SRM implementations



## LCG Grid Deployment Board Meeting



Jeremy asked about the mentioned ‘Programme of checking’ for obvious publishing faults. Jeff thought some of these are now being checked by SAM – he recently got a ticket about glue information being inconsistent. [John said he would inform Piotr about this work and requirements.](#)

Kors: Do the experiments have their own measures for their storage usage? Dario said they did using the file catalogues. The operational figures should match with what the experiments think is there. Kors then asked about tapes: is it important to know allocated tape space when writing at full rates? John wondered if this was important as knowing if full cache copies of files would be available - WAN and LAN versions of the same file.

For the SRM GIPs:

Markus thought the dCache GIP in the new version would be in gLite 3.0.2. John reported that Jens Jensen (RAL) is writing the CASTOR2 GIP and Graeme Stewart (Glasgow) is rewriting the one for DPM.

The agenda order was changed slightly to accommodate those who arrived specifically to hear Jeff Templon’s report.

### **5. A Grid with fully functional groups and roles: VOMS deployment (Jeff Templon)**

It was asked how this area can be pushed forward. Ian disagreed with the statement that testing reveals the gLite JRA1 WLMS does not work and said this was propagating rumors. It did not work as well as people would have liked. It was agreed that the deployment process needs to be worked out. Ian thought there could be a problem if decisions are made in the GDB but not pushed down to the sites. John suggested the experiments can help push this along too.

Ruth mentioned that the job priorities working group has provided a report and asked if the GDB had seen it. It was remarked that Jeff gave a report to a previous GDB but that the document had not been distributed as it was not yet an official document. Ian added that the GDB has not seen recommendation of what needs to be implemented at sites (in particular for batch systems). John thought there was an issue for storage where ACL functionality relied on support in the SRMs. Ian said it was done for DPM and is in progress for the others. There was a question about consistency when using ACLs – Ian thought the person who creates file must set the ACL to which John replied that to start with the SRM still needs to allow writing to the storage for that user.

Jeff: There are still issues of sgm vs production jobs. Everyone needs to use VOMS to make it work. 95%-98% of users at NIKHEF are using VOMS. It is the basic configuration that says sgm and users are different things. Ian: Not all users are in VOMS for each experiment. Once the registration is done, a second step is telling people only to use voms-proxy-init. Ruth wanted to clarify who defines the steps which compose the



## LCG Grid Deployment Board Meeting



deployment plan. Is it EGEE or the GDB? Ian answered that at the level of EGEE/OSG we have to provide an infrastructure that works with VOMS. As to who is responsible for getting all WLCG experiments into VOMS, this is a GDB issue.

Dario reported that the ATLAS re-registration process is still in progress. Kors thought a communication is needed which urges people to register quickly as sys-admins want to switch of the ldap servers. A deadline is required.

**Action 0607-8: Urge experiments to push users to re-register. By the next GDB implement a deadline. Before then send out document. An official presentation on job priorities is required.**

Jeff cautioned that for lcas.lcmaps we need to check it is ready. If we switch off the ldap mechanism is anyone going to get caught out? What will break!? Ian wondered about the VOMS certificates. Laurence said that the gridmap file can also be created from VOMS servers but if still using gridmap file for mapping it doesn't help.

**Action 0607-9: Jeff to make sure that the default YAIM is properly configuring the lcas lcmaps for the sgm accounts and works!**

The project needs to make sure that nobody (like CASTOR) requires the gridmapfile. Ian mentioned that DPM is the only solution that can currently support VOMS. This is not ready in dCache or CASTOR.

**Action 0607-10: Kors to arrange for a report in September on the VOMS awareness of each SRM implementation.**

There was a brief discussion on trialing the switch to VOMS before wider changes. It was thought that the PPS may present somewhere to do this work.

## LUNCH

### **6. Report from the operations workshop (Ian Bird)**

Ian went through each of the areas tackled at the OSG/EGEE/WLCG operations workshop.

On the topic of releases, Jeff asked will changes from one version of middleware to the next be documented. Ian thought the release notes should (already) cover this but Jeff described an example not covered: YAIM functions/modifications sometimes differ from





## LCG Grid Deployment Board Meeting



one place to the next. Required local changes get over written and nothing is documented to make sites aware of this situation. .

Jacko Koster asked if ARC should join this meeting. Ian replied, yes as should GGF representatives. Middleware people should feel welcome but it is not to discuss middleware. He assumed they would be there as part of EGEE anyway. All operations area should be present. Ruth pointed out that the interoperation area in this meeting was more about user support and ticketing than middleware.

Jeff was concerned by the ROCs as a link between sites and developers. Ian mentioned that they were trying to increase site representation at the TCG which would give a more direct channel for communication.

## 7. Summary of Tier-2 workshop and questionnaire (Jamie Shiers)

Kors asked if Tier-1s were at the meeting. Jamie said that most Tier-1s were represented. There were also local sessions for some countries

## 8. The LHC schedule (Jos Englen)

Jos stated that he had talked to LHC project people on schedule issues and would elaborate on the schedule but we need to work out the impacts for LCG.

One year ago it was agreed that once critical installation steps were completed, a more definite prediction for the LCG schedule would be given. It was not a new problem that led to this revision – it was a planned review.

There is a small slip compared to what was aimed at. Magnets are being installed at the rate of 25 per week and the interconnects are also being installed at that rate. This will lead to an installed complete machine in March 2007. The beam pipe will be closed at the end of August 2007. From September 1<sup>st</sup> on the machine people will decide access.

Detailed scenarios for bringing the machine online are being worked out. Beam-beam collisions are expected for the beginning of November. One other catch – the commissioning of all hardware up to full field would be prohibitive in 2007. A number of sectors will not be commissioned up to full energy and the machine will initially operate at the injection energy of 450 GeV. This has been discussed with the experiments and



## LCG Grid Deployment Board Meeting



they will find this period (up to the end of year shutdown – which will perhaps be Christmas or a little beyond) useful.

Luminosity for this period will be of order  $10^{29}$  but less than  $10^{30}$ . This is not sufficient for new physics, but is sufficient for testing the timing structure and to get tracks in the detectors. The period will be used for commissioning the machine and detectors. This still requires the full computing infrastructure but the capacity required for say storage will not need to be at the ~PB level.

There will then be a machine shutdown of say 3 months, during which the machine group will complete hardware commissioning to allow up to full luminosity. Experiments will be completing their builds at this time.

From April onwards the machine will commission at higher and higher energies. The aim is that by the end of that run (2008) statistics at a good energy will be available at the few  $\text{fb}^{-1}$  level.

For computing resources: Capacity for 2007 can be modest. For 2008 you will need to use the information above to work out requirements. The functionality should be as already scheduled (SC plans) – that is the capability of system should not be affected by these changes. A little over capacity is obviously preferred. This is the situation as currently known to us. Are there any questions/suggestions of how to map this onto LCG planning – I would be interested to hear your feedback.

John Gordon: The long term requirements are changing for the experiments since the MoU. We [sites] would like less required. How can we set the level for 2007? As a site – can we say what we want to install? If we provide less in 2007 then the carrot is the experiments get more in 2008. Jos: It is a picture of capacity? John: Or will the experiments open up their triggers anyway? Jos: It will be difficult at these luminosities....Kors: We have not had the discussion in this forum – seeing this new schedule (2 months data taking to end 2007). What does it mean to “shake out” the system? Does the data need to be kept?

Dario: ATLAS has started discussions. We need to rework input numbers in our computing model. There will be no reprocessing of 2007 data after 2008 data is available. The computing model was targeted at steady state but now that we have a machine start up scenario we can compute the 2007 requirements for resources. We thought we could do this during July - what is the timescale for the GDB presentation? John: The MoU says resources for April 2007. Dario: That is obviously now unnecessary. But we need to evaluate what is required. There are two key components 1. Storage – capacity is proportional to days of data taking. 2. Throughput – the rate will be similar so little change for the Tier-0 requirements. Kors: Will you distribute data? Dario: No doubt. In addition, as the CERN analysis facility may be reduced in size this would still be useful.



## LCG Grid Deployment Board Meeting



Jos: To be explicit, collisions will be provided concurrently for all 4 experiments. Then in 2008, once at full rates, their will be reduced luminosity for LHCb as planned. The aim is to get to  $10^{33}$  as soon as possible. There is no official decision on heavy ion running - this will be driven by physics arguments. The goal for a few  $\text{fb}^{-1}$  will bring LHC into the region it needs to be. By that time Fermilab will also have upgraded. Heavy ion should follow but this is a personal thought. At a guess this will probably be in early 2009.

Kors: Will there be a December shutdown at Christmas?

Jos: Perhaps a brief stop as the CERN management changes during that period! There are no technical/objective reasons for a shutdown.

Tony: If heavy ion running starts after a few  $\text{fb}^{-1}$ , how will this impact ALICE? Most of their reprocessing takes place during the shutdown.

Jos: Perhaps that is a factor – a few days may be sufficient to get them started. We need to decide if a few days of heavy ion running are efficient for the machine. We also need to ask ALICE.

Les: We should talk about the timing. There are reasons to act quickly – at least at the Tier-0 we need decisions to be made on the September timescale. What we have heard is a 2 month slip. This is not much. It is important to get through the first round thinking about what this means and to get sites involved. The next GDB is a good target for the first throughput of responses. We have an overview board the week before that GDB. A third reason is that there is an LHCC meeting at the end of September – they will be concerned about this. Papers for RRB (October) need to be in for the end of September. We do need to press ahead even though it is holiday season. We have to ask the others if it is realistic for the end of this month – then arrange a discussion with the sites via VRVS during August.

Stefano: A two month delay will not have significant impact on CMS needs. Les: The important thing to understand is the number of days of running in 2008. Jeff: It is also an issue for 2007. Two things – the schedule slipped once before since the MoU was agreed and a 2 month slippage is significant in 2007 run time. Stefano: We never required 2007 resources based on the amount of data. The resource requirements were based on 100 days of nominal running in 2008. Sites should be at a significant scale by the end of 2007 given readiness needs for April 2008. John: So CMS is not expecting to do much in 2007? Les: In addition the machine efficiency is likely to be less than 100%. John: The project wants everything in April to test 6 months before beams? Les: You can not assume you can wait. Stefano: Cosmics will mean we will be taking data earlier than the machine ready date. Nick: I agree with Stefano's comments. The running schedule for 2008 needs to be defined better. April from when to when, what is the data taking time? It will help with the ramp up profile. Jos: The aim is for a few inverse fb by the end of the year. This will depend on the ease with which the machine is brought online. I can not be precise.



## LCG Grid Deployment Board Meeting



Les: Can LHCb respond in the next few weeks? Nick: We don't have the information. We can guess the ramp up scenario. What can be done with 450 GeV? I can't give an answer by the end of this month.

Do we agree to delay the target date from April to August? The target was supposed to be 3 months before beams. Les: It depends on whose plans you take – assume beams 1<sup>st</sup> November. Kors: The original planning had beams colliding in July. Tony: The two month delay only makes our plans more realistic. We see things slip and should not be planning for a later date. John: It is a question of resources in 2007 vs 2008 which impacts on value for money. Les: It depends what we mean by being ready. We are not going to get more information unless we go and get it. Do we assume all is going well in 2008 with 150 days running in 2008 or take another position based on the presentations given by machine people? Nick: This is the first I have heard of a run April to the end of the year and this opens up 3-4 months of running... Tony: Why does this require more processing power? We just need to handle it [data] for longer. Nick: Storage increases linearly with data taken. Tony: So Disk/Tape storage increases?

Kors: The biggest impact is on ALICE. Nick: Others were planning to do reprocessing during shutdown. Les: We still want an initial response by the end of the month and we arrange to meet with the machine people to get their view. CERN can not be relaxed about this change. The experiments need to be part of this and it is useful to get Tier-1s involved for August.

Jeff: I encourage the experiments to be realistic. The statement suggests 2007 is being used to ramp up. It seems like it is not required for physics; the situation is different from when the MoU was written. Sites do have to carefully consider this ramp up – it is a serious deal – so they need the best information from all sides. Tony: Installing machine batches of 1200 compared to 400 was not easy. If we do not do anything in 2007 we will struggle in 2008. Jeff: Better to have realistic planning – plan for reality not the dream. Les: Each experiment needs to reach a conclusion on what is realistic. I propose a round with the experiments in 4 weeks and a Tier-1 meeting in August.

Action: 0607-11: Kors to send link around for machine presentation (Evans talk)

Action: 0607-12 Les to come up with schedule for agreeing impacts of new machine dates – will depend on availability of machine people.

John: Are we expecting to get efficiency figures for 2008? Les: They [the machine group] are usually open. Kors: Does the MB continue over July/August? Les: yes.

## 9. T1-T2 data/relationships (Les Roberston)

Kors: I am confused about operational relationships in WLCG vs EGEE. Les: We moved away from the hierarchical model. It is a complicated picture. We need to separate data



## LCG Grid Deployment Board Meeting



flow from support models. Kors asked about dates. Les: Aim to get first understanding on September timeframe. John: the original plan was to review MoU pledges in September  
Les: A review will be done by the RRB in October. Will sites be ready before then?  
Perhaps the default is for them to remain at the current pledge levels unless updated.

### 10. Experiment activities

#### LHCb (Nick Brook)

Seeing issues where sites are down but not scheduled as down. Some SEs are getting overloaded – this was a gridFTP issue and now resolved.

Jeff asked if there was a problem getting jobs to sites? Nick replied that in the VO view they see for some sites several hundred jobs queued and they could put more jobs in but with no confidence they will run. Jeff: So you don't get the jobs to the sites? John: A lot of sites have a queue for the VO. Jeff: This is a problem in the middleware. The RB does not match the VO view. The VO specific queue is a hack. The new one can do it (there is a patch in CVS). The thing at fault is the WLMS. Tony: If CERN are using Castor....Jeff: You are only supposed to be talking to the SRM. Nick: TURL access has problems.... CASTOR2 was tried, but it gave a pseudo TURL pointing at CASTOR1. Tony: This due to the application set up to use CASTOR1. The return is the default. Jeff: It sounds like a client libraries issue. Tony: The default configuration on the grid at CERN is CASTOR-1. Nick: This requires inbound access for dcap call back Jeff: An incoming connection to the WNs is not acceptable. Les: NIKHEF is operating as one site but it is actually two. dCache did not know about this requirement when first being developed.

Jeremy: What are the examples of LHCb experiment requirements not being propagated?  
Nick: Local protocols – would have been brought out earlier. Setting up the Tier-1 transfer matrix. There have been other tickets in GGUS. Les: were the tickets were ignored? Nick: the dCache implementation rolled out with gLite had a number of CAs limited to 56. Ian: That is a site configuration issue. Nick: There are two issues: accessing data at dCache sites using dcap and only able to deal with 56 CAs. Ian: The latter is a general Globus issue. Nick: The other is the problem at NIKHEF. Jeff: Originally dcap was switched off due to security issue (only off at NIKHEF). Ian: Is this a site configuration or middleware problem? The 56 limit is now fixed. Gsidcap is known to dCache but they have a priority of developing SRM v2.2. Given this NIKHEF have no option but to enable the access? Jeff: We should reorganize our network infrastructure to fix a problem in the software! Another possibility is that we view the source and attempt to fix it. Tony: The source for dCache is only available to collaborators. John: This is a distributed site problem; NorduGrid will run into a similar problem. Is it possible to VPN?

Jeff: NIKHEF is working with LHCb to find a solution.



## LCG Grid Deployment Board Meeting



On the issue of LHCb-Tier-1 meetings:

Jamie: It has taken time to get requirements in place. Harry presented combined action list last week but found it was incomplete. Problems like these need to be resolved and the service coordination meeting is the best place for this. With some pushing of GGUS and site negotiation we should make progress. Experiment specific meetings may be useful to reinforce ideas but not at the expense of other procedures. Nick: the requirement to read directly from local the SE was made in January. Tony: Site specific configuration is available – GID site scripts are available.

### CMS (Stefano Belforte)

John: Why did you use 15 RBs? Stefano: The summary plot was for all experiments not just CMS. We never find an RB can exceed 5000-6000 job submissions per day. Ian: This is a false conclusion- on EGEE we see 30,000-40,000 jobs total so there is no need for higher submission rates per RB. Stefano: We tried to run jobs using just 4 RBs and were having problems submitting the jobs. Ian: Was that with the LCG or gLite RB? Stefano: LCG.

Ian: You mention production infrastructure vs public infrastructure – the latter is not correct.

### ATLAS (Dario Barberis)

An issue with rates at RAL triggered a brief discussion on how the data is being handled in the experiment tests. Who should be deleting the data? Tony: Having sites delete the data is not the production environment you are aiming for. John: At RAL the load on dCache is rather high. Simone : 2 machines were added to the CERN CASTOR pool. Agents required for the Tier0 and Tier-1 to Tier-2 transfers so the workload scales with the number of agents.

Tony: You can ask to throw away data on disk but not on tape. Disk0Tape1 just gets deleted from disk as soon as data is written to tape. In either case it deletes. Tape is .....

John: If the cache is getting full then it could be a site issue.

### **AOB**

Jeremy mentioned that the minutes name list circulated at the meeting was the old version hence the wrong names in some areas.

The meeting closed at 17:30



## Actions

Item No.	Description	Owner	Status
0602-1	Bring the VO enablement issues/concerns again to the attention of the ROC managers	Maria Dimou	Done
0602-2	Provide written requirements of what ATLAS needs from user level accounting.	Laura Perini	Open
0602-3	Ensure the topic of user level accounting is an item at the next GDB	Kors Bos	Done
0602-4	Phrase the requirement on how to use policies in the WLMS	Cal Loomis	Open
0602-5	Distribute request to GDB list for input on the SL upgrade window for 2006/2007.	Kors Bos	Done
0602-6	Announce the work on Storage Information for the next meeting	Kors Bos	Done
0602-7	Ask for not further increase the usage of VOBox until next workshop	Kors Bos	Done
0603-1	Remind experiments of timescale for transition to SL4	Kors Bos	Done
0603-2	Start the process for electing a new GDB chairman	Les Robertson	Done
0603-3	Follow up to ensure all sites in country are publishing accounting data or contact John Gordon with issues preventing this happening	Country representatives	Open
0603-4	Schedule a GFAL talk for the next GDB – possibly to be given by someone who has attended the VO Box workshop	Kors Bos	Done
0604-5	Gather together (document) VOMS groups and roles status – what is possible, what are the limitations in current implementations and the main issues	Jeff Templon	Done
0604-6	Drive forward discussions on the VOMS and protocol issues	Ian Bird	Open
0605-1	VOBox task force to produce a work priority list before Easter	Cal Loomis	Done
0605-2	Based on output of 0605-1 coordinate a response on what is feasible – to be revisited at June GDB	Kors Bos	Done
0605-3	Provide feedback (with reasons) to Dave Kelsey or Kors Bos on whether the security policy presented by Dave is acceptable.	All	Open
0605-4	Tier-1s to report back to GDB on what proportion of their current WLCG work is not reported/accounted within WLCG	Tier-1 managers	Open
0605-5	Tier-1s to gather and publish (to the GDB) storage data. 8 numbers required for disk (allocated and used per experiment). 4 numbers for tape (allocated per experiment). This data is to be gathered at the end of each month.	Tier-1 managers	Ongoing
0605-6	Remind Tier-1 centres (via email) about the need to gather and publish storage accounting data.	John Gordon	Done
0605-7	Formulate the accounting discussion (from this meeting and other discussions) into a position statement and circulate to the GDB	Kors Bos	Open
0606-1	Provide for the next meeting an estimate on the middleware porting timescales for SL4	Ian Bird	Done
0606-2	Check the list of proposed meeting dates for 2006 & 2007 – see Kors' slides	All	Open
0606-3	Review the list of topics on slide 10 of Kors's talk and provide feedback	All	Done
0606-4	Mail the GDB list with a summary of the NDGF announcement	Oxana	Done
0605-5	Form group to review site implications of the storage interfaces defined in the Fermilab SRM meeting	Kors Bos	Done
0606-6	Raise package management as a high priority for WLCG at the next TCG meeting	Erwin Laure	Open
0606-7	Take up and discuss technical solutions for removing shared credentials from the VO boxes	Markus Schulz	Open
0606-8	Circulate summary email of the VO Box discussions	Kors Bos	Done
0606-9	Write a summary of the main issues and conclusions in the area of VO boxes	Jeff Templon	Done



## LCG Grid Deployment Board Meeting



Item No.	Description	Owner	Status
0607-1	Circulate link for BNL registration page	Kors Bos	Open
0607-2	Register for the BNL meeting via the registration page (if attending)	All	Open
0607-3	Request link from Michael Jouvan on running middleware on SL4 and forward for inclusion in the minutes	Ian Bird	Open
0607-4			
0607-5	Mail GDB on feedback on the accounting/policy document	Kors Bos	Open
0607-6	Mail the GDB list asking for comments/approval on the CA document	Dave Kelsey	Open
0607-7	Comment on/approve the CA document before 14 <sup>th</sup> July	All	Open
0607-8	Urge experiments to push users to re-register and inform tha a deadline will be imposed at the next GDB. Circulate job priorities document.	Kors Bos	Open
0607-9	Ensure the default YAIM is properly configuring lcas lmaps for the sgm accounts (and that it works!)	Jeff Templon	Open
0607-10	Arrange a report for the September meeting on VOMS awareness in each of the SRM implementations to date	Kors Bos	Open
0607-11	Circulate the presentation on LHC machine readiness/expectations written by Evans.	Kors Bos	Open
0607-12	Develop a schedule for agreeing the impacts of the new machine dates	Les Roberston	Open

## List of Attendees

X means attended

V means attended via VRVS; P means attended via Phone; A number means only for that item

Country	Member		Deputy	
Austria	Dietmar Kuhn	X		
Canada	M Vetterli		R Tafirout	
Czech Republic	Milos Lokajicek	X	Jiri Kosina	
Denmark	John Renner Hansen		Anders Waananen	
Finland	Klaus Lindberg		Jukka Klem	
France	Fabio Hernandez	X	Dominique Boutigny	
Germany	Klaus-Peter Mickel		Holger Marten	X
			Jos van Wezel	
Hungary	Gyorgy Vesztergombi		Dezso Horvath	
India	P.S Dhekne			
Israel	Lorne Levinson			
Italy	Mirco Mazzucato		Luciano Gaido	
Japan	Hiroshi Sakamoto	X	Tatsuo Kawamoto	
Netherlands	Jeff Templon	X	Ron Trompert	
Norway	<a href="#">Jacko Koster</a>	X	Farid Ould-Saada	X
Pakistan	Hafeez Hoorani			
Poland	Ryszard Gokieli		Jan Krolkowski	





## LCG Grid Deployment Board Meeting



Country	Member		Deputy	
Portugal	Gaspar Barreira	<input type="checkbox"/>	Jorge Gomes	<input type="checkbox"/>
Russia	Alexander Kryukov	<input type="checkbox"/>	Vladimir Korenkov	<input type="checkbox"/>
Spain	Manuel Delfino	<input type="checkbox"/>	Andres Pacheco	<input type="checkbox"/>
Sweden	Niclas Andersson	<input type="checkbox"/>	Tord Ekelof	<input type="checkbox"/>
Switzerland	Christoph Grab	<input type="checkbox"/>	Marie-Christine Sawley	<input type="checkbox"/>
Taiwan	Simon Lin	<input type="checkbox"/>	Di Qing	X
United Kingdom	John Gordon	X	Jeremy Coles	X
United States	Ruth Pordes	X	Bruce Gibbard	<input type="checkbox"/>
CERN	Tony Cass	X		<input type="checkbox"/>
ALICE	Alberto Masoni	<input type="checkbox"/>	Yves Schutz	<input type="checkbox"/>
	Federico Carminati	<input type="checkbox"/>		<input type="checkbox"/>
ATLAS	Gilbert Poulard	X	Laura Perini	X
	Dario Barberis	X		<input type="checkbox"/>
CMS	Lothar Baurdick	<input type="checkbox"/>	Tony Wildish	<input type="checkbox"/>
	Stefano Belforte	P		<input type="checkbox"/>
LHCb	Ricardo Graciani	<input type="checkbox"/>	Andrei Tsaregorodstev	<input type="checkbox"/>
	Nick Brook	X		<input type="checkbox"/>
Project Leader	Les Robertson	X		<input type="checkbox"/>
GDB Chair	Kors Bos	X		<input type="checkbox"/>
GDB Secretary	Jeremy Coles	X		<input type="checkbox"/>
Grid Deployment Mgr	Ian Bird	X	Markus Schulz	5
Fabric Manager	Bernd Panzer	<input type="checkbox"/>		<input type="checkbox"/>
Application Manager	Pete Mato Vila	<input type="checkbox"/>		<input type="checkbox"/>
Security WG	David Kelsey	X		<input type="checkbox"/>
Quattor WG	Charles Loomis	<input type="checkbox"/>		<input type="checkbox"/>
Networking WG	David Foster	X		<input type="checkbox"/>
Planning Officer	Alberto Aimar	X		<input type="checkbox"/>

The following also attended:

Name	Area
Jamie Shiers	CERN
Gonzalo Merino	Spain/PIC
Luca dell' Aguello	INFN-CNAF
Harry Renshall	CERN
I Ueda	Japan/ICEPP
Jean-Philippe Baud	CERN – for 5
Laurence Field	CERN – for 5
Oliver Keeble	CERN – for 5



**LCG Grid Deployment Board Meeting**

