



Version 1.0

Amendments history

<i>Name</i>	<i>Area</i>	<i>Date</i>
<i>Jeremy Coles</i>	<i>All - full edit -> v1.0</i>	<i>17th November 2006</i>

Minutes of the meeting

CERN, 8th November 2006

Agenda: <http://indico.cern.ch/conferenceDisplay.py?confId=a057715>

Minutes: Jeremy Coles

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

Meeting Summary

Detailed minutes

1. Introduction (Kors Bos)

There were no comments on the previous minutes so they are approved. Kors reviewed a few of the outstanding actions. One was the on-call statistics for Brookhaven. In relation to this Jeff mentioned that at SARA/NIKHEF they wrote a script to alert operators of transitions in the SFTs – the alert was received via SMS. One test led to many alerts over night indicating that the system (the site tests) is not ready to be treated as a core production service. Security actions on Dave Kelsey are pending and will be reviewed at the next meeting.

A reminder was given that SuperComputing is November 12th-16th. A decision needs to be made if a room needs to be reserved for pre-GDB meetings in 2007. Gilbert Poulard thought it useful to have a technical meeting. Jeremy Coles asked how well attended the meetings are at present. Kors responded that at yesterday's meeting about 30 people attended. It was decided rooms should be booked – [Kors will go ahead and book the rooms. There was a proposal to use the next pre-GDB meeting for experiments to explain](#)



their numbers in the Chris Eck MegaTable. This was to be discussed after Chris's talk at this GDB.

2. ATLAS Data Transfer Functionl Tests (Alexei Klimentov)

Jeff: On slide 3, how many files? Alexi: 10

There was discussion about the status seen for Lyon. It was asked why the FTS server in Lyon should be used for T0-T1 transfers? It was understood that the problem highlighted in the summary table for T1s did not impact transfers to the Lyon cloud. In respect of the Lyon FTS server it was mentioned that the hardware changed last week. Alexei had had no confirmation of this from Lyon. Jos thought the problem was actually exaggerated after the upgrade. Alexei noted that the problem persisted for 7 days or more which is why Lyon was perceived as not working.

On slide 6, Kors asked about cases that did not work. Alexei mentioned the two cases: Taiwan is not yet configured, Lyon had FTS server problems. The infrastructure is being checked.

Some inter-SRM transfers showed problems. Jeff: Did you try the DPM at NIKHEF? We have seen some timeout problems too. Alexei: No, but ATLAS could try it. Patrick Fuhrmann: There is an issue looking via the NFS tool. From our view there is no ability to show numbers larger than 2GB. Using any other protocol gives the correct numbers. We want to get rid of the NFS tool interface. You should open a ticket. Alexei: I am not sure that we do have an NFS3 interface. Patrick: It is not possible to generate a number higher than 2 as it is a protocol limitation. There is no tool which uses the NFS view. We can only offer NFS3. Alexei: This is a question for Tier-1s. Patrick: NFS is very insecure.

Action 0611-1: Patrick to send mail to list suggesting move to NFS3 (or better move away from it) mentioning NFS security concerns.

Kors: Doing the file transfers is one thing but what about using them? Alexei: We just checksum at the moment – more tests to come. Kors: Have other experiments tried to work with larger files? Federico Carminati: No but we are planning to. Fabio Hernandez: This is a subject that could be addressed in a pre-GDB meeting.



3. Update from ALICE (Frederico Carminati)

Kors: Why are the numbers increasing? Frederico: The simulations are now more detailed and require more CPU time. The ESD has expanded as many people are adding data to it. We have balanced the new beam time requirements and the old MoU pledges, some of which have been updated since the originals. Kors: Why are you worried about SRM testing with xrootd? Frederico: There is always a long gap between first testing and functionality working properly. We need to repeat tests with DPM, dCache and CASTOR.

Les Robertson: Why do you not just rely on Tier-1s in the USA? Frederico: Service commitments and political issues are both factors. We have been surprised by how well managed Tier-2s can contribute. Les: It is not clear why things are working so well now. Frederico: We have been reporting problems for weeks and only recently have things started being fixed at higher rate. Jamie: We are now following up on a daily basis. For each problem we the next day the aim is to explain why there was a problem. The emphasis is on understanding causes.

Kors: What is the effect of going to Geant4? Frederico: Using our own geometrical modeller has some advantages. It allows longer steps in gas for example so there is a chance of it not being slower. With tuning it may be possible to go faster. Kors: There was mention of the CAF (CERN Analysis Facility), would that be with ALICE machines and was that in the model? Frederico: Yes it was in the model – that is part of the ALICE quota. They can be normal LXBATCH machines

Jeff: Has there been any movement on getting the ALICE VO-Box moved away from Class-2? Frederico: We are working on it. The point is there are a few things that can be done with areas such as installing AFS. The computing model plus NFS causes the class-2 problem. ALICE is experimenting with Parrot etc. We have seen nothing convincing at the moment – the tests are only on Linux and 32-bit. It may have a 20% impact on resources. Anyway, we realise it is a problem and are very actively addressing it.

Holger Martin: There is an advantage in using glxexec. Jeff: All ALICE jobs are run as one person. Glxexec would allow credentials to change to reflect the correct owner. Frederico: This leads to a safer system. Jeff: If any ALICE credentials are cracked then we would have to (effectively) stop access for the whole VO.

4. ALICE authorization for data access (Andreas Joachim Peters)

Patrick Fuhrmann: On slide 10 – is this using GSI? Andreas: The code is there and used in some places. Jeff: On slide 11 – authorization configuration for various VOs. For SARA/NIKHEF – ALICE would be the only one operating with this DPM? Andreas: We



LCG Grid Deployment Board Meeting



can specify paths. Jeff: We need to configure xrootd for just ALICE people. Frederico: Just DPM? If current plans continued then this should be in DPM. The expectation was to converge to a DPM that offers this functionality. The last thing we want is to live on a “branch” of DPM as it makes support more difficult. Markus: It is not a different version of DPM. xrootd is an additional interface. There is no plan to have this as standard authorisation. It is part of the standard distribution after testing. It is packaged and released as part of DPM.

Does this give specific file ownership within DPM – that is a single owner for files? Markus: This is a problem in practical terms. It means access via gridFTP consistent your need for an ALICE file owner. ALICE will have one data mover. Andreas: The user never puts the file directly in...Markus: Then this is a statement. There seems to be a requirement to modify gridFTP. Andreas: There is no user interface for gridFTP. Frederico: All users interface via xrootd. Markus: This means no wide area transfers for general users – there may be a site issue opening ports for xrootd. How does a user authorise FTS to transfer? Andreas: They connect to an ALICE service (via GSI). Markus: ... to transfer single file? Andreas: There is one entry point with privileges to the ALICE system. Patrick: Does this mean you switch off gsiFTP and just allow a single certificate? David: Then ALICE services run on behalf of users. These services are now subject to audit – you must ensure you provide a log to sites. The record must show things in an incontrovertible way - gsi is run by the sites. Jeff: This introduces a problem for storage that we trying to remove elsewhere with glexec. Markus: For the import stage – FTS “puts” with an ALICE server certificate something on my site – I do not know who has put this file on my site and this is a problem for the site manager. Andreas: Then we will have to do it via xrootd. Markus: Then sites will need to open ports for xrootd. Each and every server will be visible and this doubles the ports. Andreas: One port is needed for xrootd. David: And not use gridFTP?

Andreas: Each file is moved by the ALICE production manager. David: It is easy to kill CPU but not storage and the site is responsible for content it holds. Markus: There needs to be a trace that comes with authorisation. Andreas: ALICE knows the user and can tell sites immediately. Frederico: If there is a problem then we have a responsible person within ALICE. Andreas: It is not a problem if all mapped to one user. David: A site can not hand off responsibility to a VO. This is one thing that triggered Fermi to have glexec on WNs. If an auditor asks what is running then the site must know. Tony Cass: IF sites are audited they need to demonstrate that they are in control of their site. They might trust you have the information but there is a part of the process that the site control does not control. Frederico: We can give the name of the user associated with the ALICE transfer account. That is 1-person with 1 certificate. Jeff: Is this person willing to sign a legal document? What are the consequences? Ruth Pordes: At Fermilab – JSPG giving security directions – requirement is for traceability. We do allow services to transfer. There is a data management layer. We need to be able to trace back to the user. Jeff: Does this hold if the service is not hosted at Fermilab? Ruth: We are discussing this – does the service provide cached information on regular basis etc. Frederico: Many arguments here are



LCG Grid Deployment Board Meeting



hand waving. Ruth: At NERSC every DN must have a user account at the facility. Each DN needs to register.... How is ALICE dealing with that?

Markus: Nothing prevents the implementation of an authentication system – ship the proxy certificate with the transfer and only allow transfers with FTS. Andreas: We could offer a log file on-line and provide a secure channel to the site...David: This still does not demonstrate incontrovertibly Markus: If ALICE mirror log files to the sites in a secure manner it may be okay. Jeff: We should get an opinion from an expert? What legal negotiation has to happen to provide this service from a VO? Kors: Lawyers will only deal directly with the sysadmins. Jeff: Is the present argument convincing? Andreas: We can write the owner name in the file name. It is easy to add the DN. Jeff: How do you know the DN is correct? Kors: We have made it clear that there is a concern and suggested a solution. We should give this to the security group.

ALICE?: How do you deal with dteam transfers at the moment? Markus: This is why the next version of FTS will have a mechanism - a receiving mechanism recreates the correct ownership- it takes the user proxy but not via the correct delegation mechanism. ALICE?: Who owns the files? Markus: The person who transferred the file via FTS. You tell FTS where to find the proxy and then authenticate FTS with the individual proxy. FTS fetches a new proxy from the server – for this a password is required but this is being corrected. It is incontrovertible but less secure because the password is stored enroute – that risk is being corrected. Jeff: The dteam case is well defined and has a controlled system which is being moved away from. For the ALICE VO case ... ALICE?? We have moved away from the xrootd to an FTS discussion. If we use the xrootd framework then... Frederico: How do we make progress?

Kors: There is concern about traceability of files. ALICE suggested a not too difficult solution so we will see if the approach is acceptable. Frederico: We make logs more widely available or put the DN in the filename. Kors: The technical discussion can continue outside of the meeting. The next version can go to the operational security group for discussion.

Action 0611-2: Andreas to come up with and circulate solution to concerns surrounding site auditing of ALICE files transferred to a site.

Patrick; Why are you not using GSI? Andreas: It is unsafe to put the proxy under the same account. Now unsafe on workernode. Patrick: What is the migration path and timeline? Andreas: What is the timeline in the glxexec discussion?

5. TCG Discussion on CE Strategy & SL4 move (Markus)

Jeff: Why do you not just enable CONDOR-G? There are some services using GRAM on CE. Classic GRAM. The solution would be to go to CONDOR-G but this is not popular in project – especially JRA1



LCG Grid Deployment Board Meeting



Gonzalo: Can the applications run on SL4? Markus: We can test in compatibility mode

Jeff: The worker nodes need SL4 VDT . LCG-CE – VDT 1.2 SL3. If there are a new batch of WNs that need SL4 then... Markus: This has been tested already

Jeremy: The previous statement was that the LCG-CE would only be phased out when the new CE shown to be better. Markus: Keeping the LCG-CE is not an option beyond the end of June – it is too easy to overload it. If the project can not get the gLite-CE working by June then there is a major problem. Fabio: There is no plan to use the CREAM-CE in 2007? Markus: We need input – there is an instance on the preview testbed. This is only an indication of readiness and does not allow a clear estimate for deployment.

Ruth: For the information system and glue schema – does the TCG decision have any impact on the IS and glue work? Markus: The work is in parallel. Ruth: The main driver for the schema was to provide information. Markus: The TCG strategy – different access – works with OGF standards. The CREAM_CE needs information provider to contact it directly. The situation where one can only reach it via the WLMS is not acceptable.

Kors: So, in the second half of the month sites can start moving the WNs.

Dario: Slide 6: It suggests that WMS 3.0 is still not in a usable state for applications. We have been using it for the last 6 months. On SLC4 we can run the current ATLAS software releases.

Tony: We have timeline for software being available. CERN would like to install 64-bit hardware with SL4 in 64-bit with 32-bit compatibility libraries. We still have substantial 32-bit kit. 64-bit hardware will be available to run in 64-bit mode. Running code on 64-bit gives a 20% improvement for free. Markus: The farm will be inhomogeneous. How should we handle this with grid access routes. Blah is not able to pass information effectively.

6. glexec (Jeff Templon)

Jeff: glexec on WN: “A thin layer to change UNIX credentials based on grid identity and attribute information”.

Jeff: On Slide 5 – model 2 is like the condor mechanism and model 3 is the gLite CE model.

Jobs may run two at a time on a single WN - complementary jobs i/o and cpu bound.

Markus: Most sites run multiple jobs on single box. Andrei Tsaregorodstev commented about complementary jobs Jeff:: Leave it to the site to optimise – they have multiple concurrent processes to optimise – if you do this you are basically taking over the site's



LCG Grid Deployment Board Meeting



role. Markus: We could pass this information through blah in the longer term. Jeff: We may be able to do it but we need to discuss with sites. Andrei: Pilot jobs could load appropriate jobs. Markus: We can ask someone running a large site if they are currently doing this... Fabio: We can specify if a job is i/o or CPU bound. Tony: We don't really do it. We look at whether the node is i/o or CPU bound but this requires the job to say it is i/o or CPU bound. Ian Neilson: There was mention of running jobs as generic VO – glide in jobs still need to be traceable. David: One user runs the pilot jobs.

Andrei: Setuid is absolutely mandatory? David: Initial policies are respected. What you lose where more than one job runs is the distinguishability between processes – and users.

Andrei: Traceability is still possible via the VO. Gonzalo: Site accounting is done at the VO level. Jeff: Wall time accounting will be correct. There are CPU time instances where people manage to escape – the time used is not attributed to the batch process.

Jeff: Status- Two guys at NIKHEF are deploying an implementation on the preview testbed. Patrick: Glexec talks to gums then LCMAPS...? David: There are two issues. There is a problem with the LCMAPS native protocol – GUMS does not require that the VO credentials are authentic (like mkgridmap file 6 months ago). We have a program to get the information ...it is an ugly interface.

Jeremy: A comment, use of glexec will prevent centralised user level accounting for VOs using it. Are the experiments aware and happy about this situation? Andrei/Frederico: LHCb and ALICE happy with this. Dario: ATLAS still wants central user level accounting (they are not currently using glexec).

Fabio: Slide 19. Is there a possibility for sites to modify this – we would like to have all glexec use associated with the site syslog. David: There is a plug-in to the underlying credential mapping framework (glexec itself is not pluggable) Fabio: Pilot jobs with no work exit quickly and do not consume resource. They do however lose a cycle of scheduling. Low rates of this are okay but high rates potentially lose hundreds of minutes of CPU time. I encourage VOs to take this into account. Jeff: Proposal – one scheduling slot should be accounted to a VO for each pilot job. Holger: what about user level accounting. We asked Dave Kelsey to create policy to ensure European laws are taken into account. It is not straightforward to see how this is done with user level accounting based on what Glexec will report back. The experiment internal accounting work should not bypass the policy work that has been done. David: If a user submits to the VO then the compliance requirement passes to the VO.

Markus: glexec does something on WNs that most batch systems do not expect to happen. With which batch systems has this (clean exit, clean up etc.) been tested? Jeff: I do not know. Without glexec it already seems to be possible to escape from job accounting anyway. Markus: If the batch system leaves something behind that is different



David: One must be careful not to break the process tree. Markus: There needs to be testing with...LSF, condor etc. David: This is ongoing now – testing is taking place on the preview testbed.

7. The use of glxec in LHCb (Andrei Tsaregorodstev)

Kors: So you do not use VOMS groups and roles? Andrei: We will move to VOMS groups and roles to describe VO polices.

Fabio: The logs shown on slide 12 are from where? Andrei: The job wrapper in this example. Kors: In summary it seems pilot jobs are run everywhere. Glxec is being tested at Fermilab. What is the timescale for wider deployment - half a year? David: We start testing on the preview testbed next week. Then there is a trade off to be made if want to wait for LCAS/LCMAPS site central interface. There is an issue of banning users. Kors: Roughly 6 months then? Markus: EGEE and GDB are making decisions without discussion. There is no automatic path that puts glxec goes into production. Jeff: Assuming something is agreed and on preview testbed, how long? Markus: We have never taken that route before. Kors: This needs to be taken up by the TCG. Markus: The technical packaging does not take much time. Testing against batch systems will take a few months. Torque/LSF perhaps quicker. It needs to be put on the list of things to role out. David: For the central site mapping service this will need to be prioritised.

Action 0611-3: Bring glxec to attention of TCG as something to be followed up (??).

Action 0611-4: Bring up security issues of VO accounting with Dave Kelsey (Kors)

TCG is discussing what? Kors: The TCG sets priorities for building and distribution of components into the EGEE release. Andrei: What is at stake is analysis on the grid. Starting in 2007 is too late. 6-8 months too late. Markus: What should we stop? The EGEE WLMS? Andrei: We keep saying what we need. Andrei: Are there sites able to test this now? Andrei: Only Tier-1s are relevant to LHCb. Currently 30 minute jobs are competing with 24hr jobs.

Suggestion, Andrei mail the Tier-1 sites whether they will accept this now or raise at LHCb-Tier-1 meetings.

8. Security vulnerability and response procedures (Ian Neilson)

Kors: I note you called the vulnerability “extremely critical” but leave it sites to decide what to do – some may understand and they shut down while others do not understand and wait for a patch! Ian: It is a site responsibility. Jeff: It is fine if it is up to sites but the extent of the issue needs to be clear. Ian: More information in advisory perhaps but could



LCG Grid Deployment Board Meeting



be dangerous - who is responsible if the OSCT recommend something? Kors: Site admins are not just responsible for their own site but also for keeping the grid secure. Markus: It needs to be made clear there is no overall management control.

Jeff: Doing an ldap search on list now suggests the most frequent version of torque on the grid is the unpatched one! Maarten: Is the information dynamic? Markus: No – the situation may be better than it looks.

Markus: To identify a correction to a problem, build, implement and then verify software and to get this done in 5 hrs is good.

Jamie Shiers: I keep a hardcopy of contacts and a contacts list in my phone.

Jeremy: How and on what timescale should sites announce shutting down/closing

Jamie: Nick Thackray has a proposal which included backup procedures. Jeremy: But if it is a security incident you may not wish to publish to broadcast lists? Jamie: The announcement does not have to mention the reason for closing the queues.

9. Summary Storage Classes Meeting (Flavia Donno)

Actions 0611-5: Experiments should agree on the VO part of the structure of their namespace (to allow space descriptions to be the same at all sites) and publish it.

Jeff: There is an argument not to have this requirement – it is up to sites to define underneath how they setup their infrastructure. Maarten: SURLS are precalculated by VOs – they use generic services. We can find out what will be that part. SRM itself does not require this but it will probably mean that handling data in optimum way is lost. Random user data should not end up scattered everywhere. *This* is how we think we will lay out namespace.

Andrei: At what level do you want this specified? The first part already varies. We do not assign directories randomly. What will sites do with this? Flavia: If you want data to be associated with a particular dataset or type of hardware this is required. Maarten: If data is not nicely organised and relies on the space token then sites will not be able to offer optimal staging. To avoid a surprise, I would recommend careful organisation of the experiment namespace. For the production manager this is not likely to be a problem.

Tony: Take RAW and ESD – organised by run number or physics channel etc. If you are reprocessing later by channel and the data is organised in some orthogonal way there will be a problem. Andrei: When we do reprocessing of raw data we do not want to have to care about the order in which files have been put on the tapes. Maarten: If well organised, a reprocessing of data for the last month will be able to read tapes for that month. Each tape to be loaded is full of relevant data.



LCG Grid Deployment Board Meeting



Andrei: The cache has to be large as the data set. We are not able to guarantee the order of reprocessing is the same as that on the tape. Maarten: Some types of operation will be more expensive than others.

There needs to be more experiment input at subsequent meetings to address such questions as the disk buffer size required by each experiment. [Kors will ensure a request is circulated again ahead of the next meeting.](#)

Tony: The phrasing of “rename a space token” is not implementable without contact between sites – it is not part of the SRM. It is not clear that there is a use case for this. Maarten: A sub-classification in raw – not everything is classified correctly at the start. Tony: That is creating a new token and migrating. I just want to be sure someone is not trying to sneak in another SRM2 feature. Patrick: It is best not to phrase it as it is in the slides as our US colleagues will be shocked!

Fabio: What fraction of disk should be used? Kors: What do experiments assume about data movement. Andrei: Dario’s did a back of the envelope calculation of how large disk buffers need to be and whether this disk is accounted for in allocations. It was a small number in the background was the conclusion.

Tony: Note that dCache and CASTOR have different approach to pins over requests. Kors: The disk is what is in the MoU.

10. Site Availability Monitoring (Piotr Nyczyk)

Talk and discussion postponed until a future meeting.

11. The MegaTable (Chris Eck)

[Action 0611-6: All sites \(+ funding agencies\) to give feedback \(see actions slide in Chris Eck’s talk for areas to be considered\) to megatable by next meeting.](#)

Holger: Isn’t this late? Kors: The proposal is to use the next pre-GDB meeting for further discussion. Chris: It is best to get it sorted out as soon as possible to prevent difficulties later. Dario: I am reluctant to tell sites how much money they have to spend – it depends on the sites. To do this would require us to have a knowledge of prices in each country. It is better for sites to use the experiment ratios when purchasing. Chris: Roger [Jones] has complained that the numbers are not yet accurate. Dario: Perhaps he is trying to do this to the 1% level. The pledges for 2009 are very different to 2008. Do we need to rebuild the megatable every year since the ratios change!? Chris: It would be useful to have requirements clear for sites over a number of years. Dario: Information is provided to



LCG Grid Deployment Board Meeting



2012 but what is needed depends on how the machine works. Kors: Agree these numbers for 2008. We will not buy equipment for 2009 now anyway. Les: Funding agencies may change their profile each year which means the experiment models may need to change again. There are many assumptions including about the machine. One has to go through a process every year – we need to do this regularly and some numbers will definitely change. Jeff: What exactly is the problem this [table] is trying to solve? Chris: It is to give sites information about what services are required by the experiments at each site. Jeff: Then iterative process is correct – profiles will skew and need to be corrected. The TDRs are like holy books –high priests [from the experiments] are needed to interpret the requirements for sites.

Kors: I propose Chris gets the full table out later this week. We'll send it to the list and ask the questions that have been asked here. Then at next the next GDB, we will use the pre-GDB meeting (5th December) to discuss what the numbers mean for sites.

Frederico: Also because of the process flexibility will be needed on both the site and experiment sides. Fabio: Do you expect that by 5th December all questions will have answers – including for example on the required bandwidths? Chris: Some problems will not be solved within that timescale. Yesterday – South Africa moved from CERN to IN2P3 and I assume this is why you are asking about that! Les: Sites should do as much as they can before the meeting but it is not possible to do everything. Fabio: I do not see how the experiments can calculate the cache for each site – even for a site to do this it is difficult. Chris: Sites will need to know the pattern of access. Kors: It is not clear what is meant here.... it will be discussed at the next meeting.

Holger: What is the current completeness of the table? Chris: By Sunday/Monday a version will have the full set of numbers for all four experiments for all sites which are now in the MoU and some figures for others.

12. AOB

No other business was discussed.

MEETING CLOSED AT 17:00



LCG Grid Deployment Board Meeting



Actions:

Item No.	Description	Owner	Status
0602-4	Phrase the requirement on how to use policies in the WLMS	Cal Loomis	Open
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
0604-6	Drive forward discussions on the VOMS and protocol issues	Ian Bird	Open
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
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
0607-3	Request link from Michael Jouvan on running middleware on SL4 and forward for inclusion in the minutes	Ian Bird	Done
0607-5	Mail GDB on feedback on the accounting/policy document	Kors Bos	Done
0607-6	Mail the GDB list asking for comments/approval on the CA document	Dave Kelsey	Done
0607-7	Comment on/approve the CA document before 14 th July	All	Done
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 lcms 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	Done
0607-12	Develop a schedule for agreeing the impacts of the new machine dates	Les Roberston	Done
0609-1	Follow up on NDGF security policy position	Les Robertson	Open
0609-2	Look up statistics for automated on-call system and send information to GDB	Bruce Gibbard	Open
0609-3	Follow up VO-Box concerns with ALICE (a class-2 solution was put forward but not taken up by ALICE)	Kors Bos	Open
0609-4	Those with comments on, and suggestions for, the storage class implementations work (many raised during the meeting discussion) should send them to the GDB list.	All	Open
0609-5	Reserve time at the October GDB for a continuation of the storage class implementation work.	Kors Bos	Open
0609-6	Send storage type sampling script to John Gordon.	Jeff Templon	Open
0609-7	Move accounting to work in decimal units	Tier-1s/sites	Open
0610-1	Send mail reminder to GDB list on maximum file sizes	Kors Bos	Open
0610-2	Provide Kors with a URL for the e2emonit work in EGEE	Jeremy Coles	Closed
0610-3	Arrange November pre-GDB meeting to further discuss storage classes	Kors Bos	Open
0610-4	Review storage classes summary slides and send comments/suggestions to the list	All	Open
0610-5	Provide more detail on who is supposed to sign the site policy for each "organisation" mentioned in the security policy document	Dave Kelsey	Open
0610-6	Send the site operational procedures policy to the list again for comment ahead of approval and ensure lawyers at sites have a chance to review the document	Dave Kelsey	Open



LCG Grid Deployment Board Meeting



Item No.	Description	Owner	Status
0611-1	Patrick to send mail to the list suggesting a move to NFS3 (or better to move away from NFS) mentioning the NFS security concerns	Patrick Furhmann	Open
0611-2	Come up with and circulate a solution to the concerns surrounding the site auditing of ALICE file transfers to sites	Andreas Joachim Peters	Open
0611-3	Bring GLEEXEC to the attention of the EGEE TCG as something to be followed up.	??	Open
0611-4	Bring up the security issues of VO (internal) accounting with Dave Kelsey	Kors Bos	Open
0611-5	Agree the VO part of the experiment namespace (to allow space descriptions to be the same at all sites) and publish it to the list	Experiments	Open
0611-6	All sites (plus funding agencies) to give feedback (see actions slide in Chris Eck's talk) for areas to be considered) to megatable by the next meeting	Country representatives /sites	Open

List of Attendees

X means attended

V means attended via VRVS

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	X
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		Tatsuo Kawamoto	
Netherlands	Jeff Templon	X	Ron Trompert	
Norway	Jacko Koster		Farid Ould-Saada	
Pakistan	Hafeez Hoorani			
Poland	Ryszard Gokieli	X	Jan Krolkowski	
Portugal	Gaspar Barreira		Jorge Gomes	
Russia	Alexander Kryukov		Vladimir Korenkov	
Spain	Manuel Delfino		Andres Pacheco	
Sweden	Niclas Andersson		Tord Ekelof	
Switzerland	Christoph Grab		Marie-Christine Sawley	



LCG Grid Deployment Board Meeting

GDB

Country	Member		Deputy	
Taiwan	Simon Lin	<input type="checkbox"/>	Di Qing	<input checked="" type="checkbox"/>
United Kingdom	John Gordon	<input type="checkbox"/>	Jeremy Coles	<input checked="" type="checkbox"/>
United States	Ruth Pordes	<input type="checkbox"/>	Bruce Gibbard	<input type="checkbox"/>
CERN	Tony Cass	<input checked="" type="checkbox"/>		<input type="checkbox"/>
ALICE	Alberto Masoni	<input checked="" type="checkbox"/>	Yves Schutz	<input type="checkbox"/>
	Federico Carminati	<input checked="" type="checkbox"/>		<input type="checkbox"/>
ATLAS	Gilbert Poulard	<input checked="" type="checkbox"/>	Laura Perini	<input type="checkbox"/>
	Dario Barberis	<input checked="" type="checkbox"/>		<input type="checkbox"/>
CMS	Lothar Bauerdick	<input type="checkbox"/>	Tony Wildish	<input type="checkbox"/>
	Stefano Belforte	<input type="checkbox"/>		<input type="checkbox"/>
LHCb	Ricardo Graciani	<input type="checkbox"/>	Andrei Tsaregorodstev	<input checked="" type="checkbox"/>
	Nick Brook	<input type="checkbox"/>		<input type="checkbox"/>
Project Leader	Les Robertson	<input checked="" type="checkbox"/>		<input type="checkbox"/>
GDB Chair	Kors Bos	<input checked="" type="checkbox"/>		<input type="checkbox"/>
GDB Secretary	Jeremy Coles	<input checked="" type="checkbox"/>		<input type="checkbox"/>
Grid Deployment Mgr	Ian Bird	<input type="checkbox"/>	Markus Schulz	<input checked="" type="checkbox"/>
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	<input type="checkbox"/>		<input type="checkbox"/>
Quattor WG	Charles Loomis	<input type="checkbox"/>		<input type="checkbox"/>
Networking WG	David Foster	<input type="checkbox"/>		<input type="checkbox"/>
Planning Officer	Alberto Aimar	<input checked="" type="checkbox"/>		<input type="checkbox"/>

The following also attended:

Name	Area	Name	Area
Jamie Shiers	CERN	Luca Dell'Acnpello	INFN
Maarten Litmaath	GD CERN	Olaf Barring	CERN
David Groep	SARA/NIKHEF	Patrick Fuhrmann	DESY
Andrew Peters	CERN	Flavia Donno	CERN/INFN
Gonzalo Merino	PIC		