LCG Grid Deployment Board Meeting

Version 1.0

Amendments history:

Name	Area	Date
Jeremy Coles	*All – draft notes*	*8/10/07*
"	*Further revisiosn*	*31/10/07*

Minutes of the meeting

CERN, 10th October 2007

Agenda: http://indico.cern.ch/conferenceDisplay.py?confId=8488
GDB twiki: 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

John Gordon – not available. Please check MB webpage -
http://indico.cern.ch/conferenceDisplay.py?confId=18007

Detailed minutes

1. Introduction (John Gordon)

Sue Foffano is still waiting for replies on the composition of some Tier-2s (Finland, Norway, Sweden, Ukraine). The information is needed for the T2 accounting for the C-RRB.

The next GDB meetings will be at CERN on 7th November and 5th December.

Since the 2008 meetings were originally setup, the April meeting has had to move to the 2nd from the 9th. The February meeting will clash with either the EGEE user forum or ROC manager's meeting. A show of hands revealed that more people are likely to attend the user forum.

There have been many meetings/developments since the previous GDB – CHEP07; EGEE07; SL4 rollout, some of the more serious data transfers GSSD proposed and a weekly management meeting has started for SRM 2.2.

There are many upcoming workshops (see slides for dates): VOMS (admin); HEPiX (benchmarking); OPN; SRM 2.2 deployment; Tier-1 resilience and reliability workshop; WLCG collaboration worhsop.

Topics for the November meeting: VOM(R)S; monitoring; SL4 services; CSSD; SRM2.2; Glue. Other suggestions welcome.

2. ATLAS – data management tests (Dario Barberis)

Gonzalo Merino: The T1-T2 export was automatic or driven?
DB: We had asked which Tier-2s wanted to be involved and some asked to be subscribed to some of the data.
GM: For those that subscribed was there a high success?
DB: O(70%-80%)
GM: For PIC the T2 transfers did not make it for some reason
DB: We can follow up

DB: We were surprised about access required to raw data. It requires quasi real-time distribution of raw data plus calibration information (inc. patches). Usually a patched release is issued once a week. How to guarantee that everyone has the same status each day is a topic for internal (software) discussion.

JG: (Slide 14) What does white mean?
DB: Not done – NDGF was not transferring to many sites in this snapshot
MK: Is this run regularly?
DB: Repeated when time permits – does not involve people at sites doing anything special. It should show the background situation. This is just the functionality between sites.

Slide 17 shows a summary of the planned timescales of the data management tests.

JG: We have seen a lot of timescales for T1s to update SRMs. Does that fit with your schedule?
DB: If it is done on the forecast timescale
JG: Holger are you happy to be shifting data when doing the upgrades?
Holger Martin: We will be offline for 1-day maybe.
DB: Also for online data we have a buffer that should correspond to 5-days.
JG: That is the long-term plan.

Simone Campana: With timeouts and latency 20 files input takes many hours.
.
DB: What is the "Control stale state"?
SC: The job remained in the same state and not scheduled for longer than one day (submitted/waiting/ready). The system considers this stale, kills and resubmits.
DB: This causes trouble to shifters but does not use wall time.

JG: If you fail at stageout is that time wasted?
DB: Yes.

SC: A proxy expired is a pathological situation.
DB: What matters is that the flow continues.
JG: You can't spot pathological problems with stagein/out?
DB: This is lack of SRM response.
??: Why do you have so much at stagein – do not waste wall clock time.
DB: Each one tries – you have to try on the server that has the file. For stageout you can retry to different servers.
SC: It depends on a file by file basis.
LHCb: LCG-stagein is one approach, the other is to copy to the local job.
DB: This means copying the file to the WN.
Philippe C: So the copy is failing or the stagein on storage is failing.
PC: So the file is not on disk or you failed to copy it?
SC: I should mention that all these files are on disk. It is rare to get files from tape – either a mistake or unexpected. Lcg-cp does a stage-get and this is what fails. Not getting the file from tape that fails. It is an srm connection problem.

JG: Have you done any analysis to see how data is distributed?
DB: Yes, it is almost right – too many files moving around.
JG: DDM notices this?
DB: Yes – and tries to collect but not always successfully. There are files that are left around.

JG: Is there a solution – we were discussing this yesterday. It is important where there are a lot of users at T2s. Does anyone know the dCache position on ACLs?
Ian Bird: Not before 2.2
Tony Cass: It will not arrive before the data. What do you mean by quotas? Across the grid is a long way off. ACLs is one thing and quotas another.
Jeff Templon: DDM does something on our side that helps us – ATLAS production pool as long as ATLAS VOMS proxies… carved off space for local users. Users send jobs and write to local SE. Make a 70GB pool. Until this problem is solved
DB: If one can direct user output to given sites based on group/role then we can establish a policy.
JG: ACLs and quotas are worth further discussion.

3. LHCb (Philippe Charpentier)

PC: We have seen problems with shared areas at several sites – NFS? at FZK, RAL…. Issues with lhcbsgm. Later we have to go to a group based approach lhcbsgm01 etc. The changes are not made on the nfs area – we do not want all directories read/write accessible for the world to avoid users deleting software.

JG: On the shared area issue – the temporary solution was to have a shared software area.
PC: The alternative is to install on nodes or with each job.
JG: It may be time to look at this again.

JG: Was the ROOT problem found elsewhere than RAL?
PC: The compatibility library was provided elsewhere – running SLC3 on SL4.
JT: This is really looking at the password file, so we should see it fail at other sites.

IB: Is the temp files being unavailable problem at some or all sites?
PC: At all - in some cases it is temporary so we do not know if it is permanent. We need to establish a protocol so the experiment can flag file problems

SARA: We have seen this – we do send a list to the experiments, but the experiments do not update their LFCs.

PC: When we ask for 5000 files to be staged, the chance that 200 are on a single tape is better than 20. we need to balance between size of cache and request size.
JT: It is good to do but don't get attached to it – sending a list of files which have problems may work now but next year we may need FTS to transfer that data! Polling can overload servers.

TC: This is mixing up quotas and ACLs – the longer term plan is for quotas. We need to check the timeline for ACLs. Still that does not address the VOMS integration issue.
PC: Then I'm scared - all directories are world read and group read/write.

Markus Shulz: The latest middleware will be made available before certification and testing but it is used at the experiment's own risk – it is not a release

JT: Slide 13 point 3. We have seen this multiple calling problem. Given yourself a seamless system – import OS, or make sure you need as little extra stuff as possible.
PC: Recursive usage of python is an issue

4. CMS (Matthias Kasemann)

JG: Are the targets nominal rates for 2008?
MK: yes.

Michel Jouvin: One problem seen at GRIF. We were commissioned a few weeks ago but now not – it seems this was due to a change in the subscriptions and there are not enough to stay commissioned!
MK: Will feed this back.

5. ALICE (Federico Carminati)

JG: Have you done gLite WMS tests?
Patricia M: These are running continuously and going well. We have now reduced job loads so as not to impact other experiment testing.
SC: How many jobs per day
PM: 10 concurrent jobs

Luca: At CANF we will also use STORM for D1T0 – how will you cope with this?
FC: Don't know – will have to see
??: For us CASTOR is currently a priority.
Luca: Eventually you should contact the STORM developers.
FC: Yes – thanks for bringing it up

JG:On slide 13, is that a train of tasks or data?
FC: Of tasks.

JT: It was said earlier that xrootd is in production at dCache T1s. Nobody has done anything significant at our site. I'm worried – looking at other experiments it took 1-year+ to get data services up and running. I would like to see more testing
FC: WN access is done via xrootd
JT: I do not expect you will be able to compress what the others went through – you need to start soon!
FC: We are trading preparation time but with a well known product with tuning time.
JG: Doing transfers
FC: Problem is tuning data access
JG: Have the others done it?
JT: Well the preparation. If using gsidcap then others have done the work for you but not for xrootd.
??: This is the "rest of year effort"
FC: This was out gamble and it leaves less time to deploy

JG: You have two quotas at CAF (disk and CPU).
FC: We monitor with Monalisa – from this we tell usage and then communicate back priority/length of queue.

TC: You have a bidirectional arrow – out with xrootd but xrootd writing into CASTOR this is supposed to be via rfcp.

FC: Writing back with xrootd

6. Common Computing Readiness Challenge 2008 - Planning (Jamie)

Reminder of the "WLCG service reliability workshop". – 26th-30th November. Wiki now setup.

MJ: Does this cover Tier-2s?
JS: Logical priority – T0-T1 and so forth but T2s welcome.
JG: Anything that T2s run that T1s don't?
JS: Availability targets and response times go down. If include analysis then may get down to group clusters. Have a look at the CMS document.
JG: T2s do have people with lots of experience too….

JS: May setup a registration page [later done: http://indico.cern.ch/conferenceDisplay.py?confId=20080].

wlcg-ccrc08@cern.ch is the mailing list. T2s will be represented by experiment operations people for their sites.

JG: Are there key services other than SRM2.2 that we don't see now?
JS: Yes R/O LFC for LHCb.

JG: Where do the sites know the resources required?
JS: For February we will live with what is there. For May sites should have pledged resources in place.
JG: But without ACLs etc. they still need to be partitioned.
JT: It is also useful to know how it is going to be accessed. Discussion with LHCb revealed interesting expectations.
JS: If there are things sites need to know then good to ask now!

JT: You showed a plan at the start – please could you put these as links in the agenda? LCG area -> Workshops ….

7. Policy Update (Dave Kelsey)

- Grid Security Policy (v5.7) now approved except by OSG.

- Grid Site Operations Policy (v1.4) – now approved. Still need the covering documents for specific grids.

- Working on VO Operations Policy (v1.4) – final call soon and then GDB approval by mail.

- Pilot Jobs Policy (v0.3).

JG: On this last one we have seen encouragement from LHCb. Have the others looked at this?

JT: VOMS servers have number of properties that count as CA services….
DK: That discussion is coming up… there is not much difference between authorisation and authentication. A group is being set up to look at minimum requirements for AuthZ Attribute Authorities (AA).

Requests to GDB: Comment on VO Operations Policy (final call next week); Comment on Pilot Jobs Policy (V0.3).

Etienne: Does the VO policy include a definition on SAM critical tests?
Dave: No
Maria Dimou: This was discussed briefly by email, but we never said VO Admins responsibility should include anything other than vouching for the trustworthiness of members.
DK: Do we need to include something?
MD: It can not be part of the VO policy document.
DK: That does refer to other documentation
JT: These do not (always) define what is meant by VO admin, Vo manager, VO resource manager….
DB: Can we define how many roles need to be defined?
DK: Here we have VO manager and security contact
DB: What is the definition of the VO manager? Me or Alessandro…. Operational issues need to come to me.
JG: It is up to the VO to decide how it will meet the roles. An operations policy may have something in it to define critical tests but not define who should do it.
MD: In cases so far questions go to the VO manager and they don't know the answers.
IB: SAM critical test is completely independent of the policy.
JG: You could extend the operations policy….

HM: VO security contact and Resource manager should be defined. These are needed for the service level agreements for VO boxes for example because roles are not well defined.

DB: I want to raise this problem. The VO manager is identified as the VO administrator… the VO manager should be at the top.
IB: There should be a distinction.
DB: This is not how things started.
PC: Then the documents should be rewritten!?
IB: No
DB: Probably ok for policy but web-pages do mix up the roles.

Issue: There is a lack of clear definition of roles and responsibilities in VO hierarchies.

8. Job priorities WG (Dietrich Liko)

PC: Are you using groups or roles?
DL: Groups – but as an abstract string
JG: Are you saying what is required at each step? You are now installing on the certification testbed - what is your estimate to finish testing?
JT: Simone should answer.
SC: Recently tried on certification testbed all elements to be tested and the functionality worked. Implementing was via YAIM so now it is possible to configure it [the installation via YAIM]. However, there is still a problem with obsolete FQANs.
JG: Is this YAIM configuration independent of the rollout of YAIM4?

JC: Are these results for all batch systems that are in use across WLCG?
SC: No. Torque/maui only have been tested. There are batch system dependent things which would require further specific testing.
JT: There are two issues – configuration of the batch system is dependent while other parts are independent. SGE need changes to dynamic scheduler… for major batch systems the situation is okay. Has France got it working on pqs?

DL: The hope is proper testing on PPS.
MD: Job priorities are based on roles. ATLAS insisted on 20+ groups.
DL: This solution can not scale so not using many roles in ATLAS.
DB: There is no way to set up 20+ groups. Will use information internally but there is too much manual intervention required to automate it.

PC: What you are addressing is not priority but fairshares. If a user submits 20,000 jobs and another 1 job then they wait. You don't address the issue of the guy submitting a few jobs having any priority over the one with the larger number.

Steve Traylen: This is two users in the same group. There is always a user fair share configured.

PC: This is done at the level of the site not WMS.

SL: We can implement priorities across groups.

PC: Yes, up to three.

JG: Do we want to stop this now or get it into production?

SC: Has to be in otherwise if one user submits 1 million jobs then how do you protect others?

IB: The WMS only submits a certain fraction of jobs.

JT: If you have this implemented … single user caps come with the CREAM CE

JT: I encourage people to remember the short term nature of this solution.

9. VOMS attributes usage – open issues (Christoph Witzig)

JT: There was a suggestion that there be a test suite associated with this document (slide 4) and anyone implementing would test against this.

Slide 6: JT: I understood this the other way around. There are two implementations already. I prefer VOs can use this [free text attribute] in their own frameworks but that the generic middleware is blind to it.

Slide 7 – point 5. The primary FQAN can only have one role. – if it is important to have more then speak to Christoph.

ML: There is still a problem that the FQAN used to submit the job is wrong for data management.

Akos Frohner: This slide is about authorisation whereas what you refer to is about ownership. All FQANs should be considered equal for authorisation.

ML: LHCb asked for a group as primary FQAN and got it.

JG: Are you now saying you can have multiple roles? By this statement the m/w needs to decide on what access is granted.

IB: Agreed we have to go back and look at the auth process to define rules. There is a specturm of possibilities.

JT: So, it is clarified that it is recognised that there are problems!

Study on authorization

PC: A remark – currently authorisation is done with mapping on CE and SE. Has there been any discussion with these people on how these things interact with this area.

CW: I understand dCache etc. will move to this model.

PC: What is the consequence for files? That there is no direct back door access?

JG: You mentioned an authorisaion service earlier – is this how you will implement it?

CW: It is not the right approach

JG: Any hints on how you will implement it

CW: We will deliver a set of recommendations but would rather listen for now.

JG: Then the recommendations have to be implemented?

CW: That has to be agreed by the TCG

JG: Will this be implement by EGEE-III?

IB: You mean the recommendations that we can make use of?

MS: There is no automatic process that everything out of this study has to be implemented.

JG: The last two talks have referred to autnZ.

JT: This is an area we have been blind to in the past. Clearly a deployment focus. Another issue – we have been running out of production manager pool accounts. 12 accounts are being used by 5 other people. Depending on how the permutation of groups/user accounts, they get mapped to different production manager accounts. So we could also look at how to rationalize this situation. Perhaps turn off the gridmapfile for anything other than vanialla users.

MJ: We've had no complaints in the last year and rely on use of VOMS proxy.

MD: Would like to map VO members to a flat structure. Is this ok for VOs? How ready are other services for VOMs proxies? Groups with roles operator/site. There would have to do a lot of changes.

JG: This is more an issue for sites, do they want to enable groups

ML: The gridmapfile is a hack – one choice per DN. VOMS would anyway replace it. So Jeffs question – can we have an intermediate step and just have gridmapfile show plain VO membership.

JG: How do we do it?

MJ: What services are not VOMS aware? Gridftp – but that is not a problem. Gsissh – again not a problem.

PC: We can not accept it until it runs on all sites with VOMs proxies. User, lhcbprod and lhcbsgm.

MJ: You are wrong – you are running at GRIF. It works for software installation.

PC: There are many sites where it does not work.

PC: How can we use VOMS proxies with more than just 3 roles – there are many site configuration problems. LHCb are fed up.

MD: LHCb tried to move the production group to a role. This has nothing to do with it.

PC: VOMS should be able to use groups and roles. LHCb have given up trying to use different groups as sites are not configuring themselves correctly. LHCb and production

and person mapped as sgm -> we are not discussing theory. Group needs to discuss practice.

Action – Either circulate around this group or do something else off line

10.   glexec & pilot jobs (John Gordon)

*The aim is not to get glexec but to enable pilot jobs.

IB: A lot of the reluctance with pilot jobs is coming from the UK.
JT: I have a different reservation, I'm worried if about giving VOs the freedom to glexec anything. Some thought needs to go into how it is used. This is a deployment issue. Job has a tmp working direcroty. But with glexec there would not have working directory. Glexec would need to handle the permission on the file.

DK: Policy doc says VO should make public how they will use pilot jobs
JT: And if a site does not like it?
DK: They do not need to run that VO. It would be useful if experiments running pilot jobs do document their framework.
JG: What is the common framework?
JT: How to get from the job coming in to running the job
MS: One action to change id
JG: Would it be enough to put experiment pilot job experts together.
MS: We have not done the ID change yet. So this is a unique time to do this. Perhaps use what is used in the US.

[JG note: MS comment on fact this is implemented in US and we might learn from their implementation].

??: Having both LCAS/LCMAPS on each worker node does not scale. The library version has been in the release for some time. glexec is now part of some VDT release and active in OSG.

MS: The TCG still has to discuss where this will go. Already it is on the CE. It is a simple library with no functionality – can add and test within 2 weeks (with top priority).
JT: I would worry about rolling this out quickly. People start to adapt their frameworks.
IB: Perhaps we can stop discussing pilot jobs and agree that we'll support them?
JG: So the statement is that we recommend that pilot jobs run everywhere. Is there anything we can do to stop sites saying they will not use glexec.
MS: Most sites now run pilot jobs – but it is up to them to choose.
??: Somehow we need to express the glexec mode in the information system.

PC: The whole point is to know if we can run generic pilot jobs.

TC: glexec needs to have a 4th mode which is to abort the job.

PC: Pilot jobs are a fact – the only way to prevent them is to cut outbound access. We use pilot jobs now.

IB: And this issue is about identity changes.

MJ: I'm not sure it is a good idea to allow sites to accept a VO but not pilot jobs. At CHEP …

MS: You realise and I realise this is the best way forward. But you can not ignore that sites autonomy.

JT: Could it be said that pilot jobs would be earmarked with a special role and then a site just does not accept that role?

MS: That can be bypassed very quickly.

JT: But that would be a policy violation. We could go back to another solution now – one user runs jobs for themselves.

Les Robertson: The current situation is that pilot jobs are running. Do you have to run pilot jobs if you want to take part in this work? It appears we have a solution with glexec. If you do not run it then you only run a subset of the VO work. There is a solution that you can use – if don't want it then sites have to propose something else …we should conclude something now.

PC: What is a pilot job?

LR: You [LHCb] run jobs run as different users

PC: We are not breaking the policy.

IB: Only ALICE is breaking the policy. With them a general user runs jobs for anybody.

JT: One set of people submitting to task queue

John: Action: Then I am to send a proposal to the MB and the MB is to forward one to the CB.

JT: Before rolling this out there has to be some agreement on how it will be used.

DK: We talk about 3 modes. Actually we want 2 modes. We –centrally- want the audit logs.

IB:The conclusion must be if you want to run the LCG VOs then you have to run pilot jobs.

11. CERN, T0 and CAF resources (Bernd-Panzer Steindel)

JT: If you use smaller disks and buy *more*, then you'll need more racks and cooling.

12. GSSD (Flavia Donno)

??: What does 75% efficiency mean in this context?
FD: Out of 4 files 3 copied
SC: Some files were on tape so stagein may have failed, so in some sense the numbers are arbitrary!

JT: There is an issue with regular expressions for FQANs, have you thought about changing to using the standard?
FD: no.

JG: Does "site catalogues" mean LFCs?
SC: Not only, given the OSG and NorduGrid differences – people may want to do it in some other way.

IB: It is best not to ask for a "wish list". The strategy to check consistency is a valid request.
TC: Also changes for the experiments and a number of other things need to be done to improve the tape service at CERN … new requirements will distract from that work.

13. Monitoring (John Gordon)

Introduction (John Gordon)

There is a summary of VO specific SAM tests – https://twiki.cer.ch/....

JT: PC's comment was more extreme. 25 raid arrays might support the VO but if the user can not get the file he wants then the site is down!

What next? (James Casey)

MJ: People want to do things for integration.
JT: Reliability of message transfers is a problem with R-GMA
JT: The dashboard monitoring uses L&B data, the other uses code?
??: It would be useful to have a standard report on ACLs/mapping.

5. AOB

Announcement – Les Robertson will be returning next year. Ian Bird will be appointed by OC board from January.

The meeting closed at 17:30

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
0606-7	Take up and discuss technical solutions for removing shared credentials from the VO boxes	Markus Schulz	Open
0607-9	Ensure the default YAIM is properly configuring lcas lcmaps for the sgm accounts (and that it works!)	Jeff Templon	Open
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
0609-6	Send storage type sampling script to John Gordon.	Jeff Templon	Open
0609-7	Move accounting to work in decimal units	Tier-1s/sites	Open
0610-5	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
0701-3	Check the CPU and storage accounting figures being published for the site	Sites	Open
0702-3	Discuss the future of a VOMRS-VOMS task force and consider possible mandates for the group	Dave Kelsey, Maria Dimou et. al.	Open
0702-4	Check Harry' resource tables and understand what they mean	Tier-1 sites	Open
0703-1	Check the Victoria MB time with Les Robertson and agree intention at the MB	John Gordon	Open
0703-2	Follow up on accounting policy documents	John Gordon	Open
0703-3	Send out a link to the latest patch	Jeff Templon	Open
0703-4	Follow up on the VOMS coordination group mandate	John Gordon	Open
0703-5	Refer Cal Loomis to Marian Dimou concerning the representation of smaller VO requirements in TCG discussions	John Gordon	Open
0704-1	Update slide 17 of presentation and formulate a request for documentation to be provided by the middleware developers to explain options with components (needed by Quattor maintainers)	Michel Jouvin	Open

LCG Grid Deployment Board Meeting

Item No.	Description	Owner	Status
0704-2	Follow up on VOMS coordination group mandate wording with Maria Dimou	Ian Bird	Done
0705-1	Get feedback from Markus and Alessandra on previous feedback from sites on glexec.	John Gordon	Open
0706-1	Check use cases and VOMS need for failover with the developers and VOs	Maria Dimou	Open
0706-2	Provide description of implementation(s) of VOMS based ACLs and submit this to the experiments to confirm it satisfies their requirements.	Flavia Donno	Open
0706-3	Review the membership and approach of the Job Priorities Working Group	Erwin Laure	Open
0706-4	Nominate someone to join the grid services monitoring work	Oxana Smirnova	Open
0706-5	Follow up on how best to proceed with site-experiment negotiation on what VO SAM tests are to be monitored	John Gordon	Open
0706-6	Setup group to gather and prioritise GridView requirements	Ian Bird/ John Gordon	Open
0706-7	Follow up c) with Dave Kelsey	John Gordon	Open
0706-8	Raise glexec questions at the Stockholm operations workshop	Ian Bird	Open
0708-1	Provide feedback on the VO Operations policy	Reps/All	Open
0710-1	Comment on VO Operations Policy (final call next week); comment on Pilot Jobs Policy (v0.3)	All	Open
0710-2	Seek better definitions of VO roles – such as VO manager, VO operator etc – as they relate to policies.	Dave Kelsey	Open
0710-3	Circulate more requirements/issues information to the VOMS attributes group	??	
0710-4	Follow up on Markus's comment about glexec being used in OSG already and how experiences might be shared.	John Gordon	Open
0710-5	Send statement to MB regarding pilot jobs and glexec. Request MB to consider and forward to CB for comment.	John Gordon	Open

List of Attendees

X means attended
V means attended via VRVS

Country	Member	Present?	Deputy or Technical Assistant	Presen
Austria	Dietmar Kuhn			
Canada	Reda Tafirout		Mike Vetterli	
Czech Republic	Milos Lokajicek	X		
Denmark	John Renner Hansen		Anders Waananen	
Finland	Klaus Lindberg		Jukka Klem	X

LCG Grid Deployment Board Meeting

France	Fabio Hernandez		Dominique Boutigny	
Germany	Klaus-Peter Mickel		Holger Marten, Jos van Wezel	X
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		Jan Krolikowski	
Portugal	Gaspar Barreira		Jorge Gomes	
Romania	Mihnea Dulea			
Russia	Alexander Kryukov		Vladimir Korenkov	
Spain	Jose Hernandez		Xavi Espinal	
Sweden	Leif Nixon		Tord Ekelof	
Switzerland	Christoph Grab		Allan Clark, Marie-Christine Sawley	
Taiwan	Simon Lin		Di Qing	
United Kingdom	Jeremy Coles	X	John Gordon	
United States	Ruth Pordes		Michael Ernst	
CERN	Tony Cass	X		
ALICE	Alberto Masoni	X	Yves Schutz	
	Federico Carminati	X		
ATLAS	Kors Bos		Stephen Gowdy	X
	Dario Barberis	X		
CMS	Matthias Kasemann	X	Patricia McBride	
LHCb	Ricardo Graciani		Andrei Tsaregorodstev	
	Nick Brook			
Project Leader	Les Robertson	X		
GDB Chair	John Gordon			
GDB Secretary	Jeremy Coles	X		
Grid Deployment Mgr	Ian Bird	X	Markus Schulz	X
Fabric Manager	Bernd Panzer			
Application Manager	Pere Mato Vila			
Security WG	David Kelsey	X		
Quattor WG	Michel Jouvin	X		
Networking WG	David Foster			
Planning Officer	Alberto Aimar	X		

Others present at CERN

At CERN:

Maarten Litmaath – CERN
Tadaaki Isobe – Tokyo
Jim Shank – ATLAS/US
Marck ud Sanden – NL
Jamie Shiers – CERN
F Chollet – IN2P3
G Poulard – CERN/ATLAS
Luca Dell'Agnello – INFN
Sue Foffano – CERN
Harry Renshall – CERN
Gonzalo Merino – PIC
Francisco Martinez – PIC
Gerard Bernabeu – PIC
Oxana Smirnova – NDGF
Philippe Charpentier – CERN/LHCb
Nick Thackray – CERN/GD
Gavin McCance – CERN/GD
Paolo Pedebco – CERN/GD
Simone Campana – ATLAS
Dirk Duellmann – CERN
Maria Girone – CERN
Patricia Mendez – CERN
Maria Dimou – CERN

Pm:
Steve Traylen
Romain Wartel
Flavia Donno
Lawrence Field

On VRVS:
Derek Feichtinger
Etienne Urbah
Stefano Belforte
Jos Van Wezel
Richard Gokieli
Andrei Tsarago.....