



## Version 1.1

*Amendments history:*

<i>Name</i>	<i>Area</i>	<i>Date</i>
<i>Jeremy Coles</i>	<i>All – draft notes</i>	<i>13/12/2007</i>
<i>“</i>	<i>Further revision</i>	<i>27/12/2007</i>

## Minutes of the meeting

CERN, 5<sup>th</sup> December 2007

Agenda: <http://indico.cern.ch/conferenceDisplay.py?confId=8508>

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)

### Detailed minutes

#### 1. Introduction (John Gordon)

John welcomed everyone to the December meeting. There was one membership change to notify – Marco La Rosa from Melbourne is moving on.

*The actions were reviewed and updated.*

For 0711-3 on proxies being stored in multiple places, Claudio Grandi noted that the middleware security group meeting was going on in Berkeley. A document is to be expected January/February. Jeff Templon commented that this will depend on the TCG meeting in January where the document will be discussed.



## LCG Grid Deployment Board Meeting



0711-3 Meeting has taken place. Ulrich has done some tests – did not find any problems so far.

0711-8 Is happening.

Philippe C: We need to confirm what are the resources available.

JG: Harry R will decide if those anticipated will meet the experiment needs

JT: Somehow in communication channels an error comes in – purchased resources end up on sheet but allocated do not. Is this just us (do other T1s see this)?

JG: In the accounting sheets there are two numbers – disk and CPU allocated to each VO and total capacity.

JT: What of the other VO specific figures?

JG: HR knows what you have installed

HR: I will take up this matter as an action

**Action 0712-1: Harry to review Jeff Templon's concern about purchased resources figures ending up in resource spreadsheet but not allocated resources.**

JT: I must have missed 0711-7, what is it about?

JG: Batch system tests

JT: There is an SA3 batch system integration meeting next week. We need to be careful of overlap.

*John continued to talk about meetings.* Since the previous GDB there has been: HEPiX; SuperComputing; SRM2.2 workshop; CCRC face-to-face meeting and a Reliability workshop. With regards to GDB's for the coming year:

January – 9<sup>th</sup>: John asked if this date should be kept as it is very close to the CERN 2008 start date. He asked if anyone had already booked flights or had a strong preference.

Luca dell'Agnello: I would prefer to keep this as the 9<sup>th</sup>.

JG: Then can we look at moving it the arrangement to Wednesday (GDB)/Thursday (pre-GDB) – this allows 1 or 2 days extra...

Jamie Shiers: On 10<sup>th</sup> this room is booked

JG: For April my gut feeling is that this meeting will move.

SL: Has Les given the dates to John?

LR: Yes.

JG: I will not decide for a few weeks in order to get more feedback.

**Action 0712-2: Gather more input before making decision on April GDB**

*On Job priorities:* Testing continues on the PPS, the implementations are not ready for production yet



Oliver Keeble: I am not sure of the current status, but when people are happy with it we will merge it into YAIM, hopefully with not much delay.

*On the persistency of storage tokens:* This was discussed at the CCRC yesterday; it will be covered in the GSSD report.

## 2. Using GGUS to connect VOs, deployers and sites (Maria Dimou)

Please refer to the slides (on agenda page) for details of the talk

JG: We no longer have single releases so communicating the status is more difficult

OK: The SA3 release team maintain a wiki page

JT: The prime focus is on the VOs. A lot of the things you mentioned are interesting to me, but the ticket content from the site point of view is not covered.

MD: This input comes via the ROCs and weekly operations meeting

JT: We are improving GGUS from a VO point of view ; what is the mechanism to do this from a site point of view? An example from last week – middleware problems get submitted by a site and end up being assigned back to the ROC

CG: They should be submitted to GGUS and the TPM (Ticket Process Manager) then needs to forward it correctly

JG: Is it not that the TPM assigns it

SB: The TPM is not perfect

MJ: We need a channel to feedback problems with GGUS

IB: The channel is through the ROCs and making sure that they speak to the GGUS people

JG: The VO situation is getting better because Maria is taking them forward

SB: TPMs are assigning the tickets

IB: It is a process issue – site managers should take problems to the ROCs

JT: This does not work – for the experiments side when it does not work it is like saying speak to the experiment coordinators!

IB: You can specify it

JT: It is not working

IB: They are getting money for it

MJ: Why not publish the feedback

MD: There are several places for this such as the ESC. Follow the links in my talk for more information

MJ: I want to see the status of problems submitted

Markus S: This would be asking the parser to summarise savannah states but is it worthwhile?



## LCG Grid Deployment Board Meeting



MD: See the link on title page of talk: <https://twiki.cern.ch/twiki/bin/view/LCG/VoUserSupport>

JG: Take up at ROC manager's meeting. See tickets end up against sites and not chased. [There are several issues to be followed up.](#)

[Action 0712-3: Follow up concerns about site-GGUS improvement with ROC managers. What is the process for using savannah vs GGUS tickets and how can a submitter find out the state of the submission.](#)

### 3. Monitoring Working Group Update (Ian Neilson)

JT: Is this standard going to be in xml format? It would be useful to have non-trivial examples documented – this would help when we have to write local checks.

IN: You have to consume the xml

JT: Who writes the local check then? Typical examples only have one of each thing.

### 4. Visualization Ideas for Management Dashboards (John Shade)

Slide 23 : Holger Martin : Who is “we”? [This was about “user” requirements for dashboards]

John S: This was originally an EGEE project

HM: Who stated the requirement? Who is the management?

IB: It was me to start with.

JG: Slide 27 – is this a static snapshot?

John S: In much the same way that Gridmap updates. You can click on one option and keep others constant, but that is one possibility.

JG: We've looked at different ways of presenting metrics

Fabio Hernandez: The importance for sites is in being able to aggregate information for the region.

John S: Yes – static URLs for sites can be used once known

JG: There is a dashboards page

Kors Bos: For the Gridmap it strikes me that BNL and Fermilab are not there.



## LCG Grid Deployment Board Meeting



James Casey: It does not know the information, there is a mismatch of info. in the BDII and at the end you do not know the number of CPUs. Another problem for Fermi is the multitude of site names.

IB: This is a problem that needs to be resolved

KB: Regions view is nice. But in the Tiers view I would also like a VO view

James C: We need the published info not in python. How and where do you define the cloud of ATLAS? It is again a problem with naming.

IB: Due to the different configurations

James C: At least you have a structure we can use...

IN: In the longer term we do not want to patch the gridmap – we could use and change scripts for everyone but then later we would need to disentangle it!

GP: There is an ongoing ATLAS effort to use the APEL structure.

## 5. Directions for service availability calculation (Piotr Nyczyk)

FH: Yesterday our site was in downtime and we got ticketed

PN: There may still be plus or minus one hour discrepancy. We stopped doing any downtime correlation in SAM.

FH: Also sometimes in GridView Scheduled Downtime does not show

FH: Critical tests per VO are selected by the VO?

PN: yes – some people with appropriate ACLs can modify the selections

FH: I ask because at the MB yesterday ALICE said they are not interested

PN: They have their own tests and results

JT: ALICE completely ignore this

IB: So to site it looks critical but ALICE themselves are not looking at it.

JT: The first block (slide 9) is exactly the problem of mentioned by the last speaker – identifying T1s, T2s etc. It would be good if the work is not done twice. Submitting tests depending on the Tier would be useful.

PN: It is reasonable not to submit T1 tests to T2 sites. In addition in the availability calculations we should know to differentiate tests of relevance.

JT: There are other disadvantages – if the experiments have their own set of tests then sites have multiple places to look and it is not clear which is the official one.

James C: At CHEP and the workshops it was agreed that SAM is where availability would be stored. The dashboard is the experiment view, but SAM shows for all. The data should be in SAM but experiments may ask you to look at the dashboard to correlate with other information

JG: Are people publishing job success rates?

James C: This is not an availability measure at the moment.

JG: This is a good example



## LCG Grid Deployment Board Meeting



JT: There is again a disagreement between SAM and the experiment view.

JG: Do you see a load problem with the experiments getting results?

PN: Sometimes, so we started to restrict access. If there are well defined (information) channels then this should not be a problem.

JT: Do you have a plan to get the answers to your questions?

IB: In the next few weeks (working group going for a year) there are a whole bunch of issues and possible projects to look at. In the next few weeks we need a proposal for priorities – should we come back to this in January?

FH: Are sites going to be involved in the priority setting?

IB: That should be discussed here.

[Action: 0712-4: Compile and report on list of SAM/monitoring priorities at January meeting.](#)

## 6. Architectural Models for WLCG monitoring (James Casey)

JT: Slide 8 – this shows a backdoor communications mechanism. It is interesting. A different approach is taken for ALICE - they want to be able to ticket/email sites directly and have read access to GGUS etc. At the same time we've turned on access to the testing Ian mentioned.– tests run when proxy is being renewed. The VO should say this is where/how we will publish information. How you decide to trigger sysadmins is a site responsibility.

James C: You still have ultimate control – we're developing for 100 sites and they don't always understand what turning on the text option will do, and we end up getting complaints.

JT: As a forward question, will this messaging system also be used for the dashboard? (slide 15). One would hope that if you have one then in the interests of common solutions it would be shared.

HM: I did not understand completely who needs this messaging system? Sites run their own tests as do the experiments. Who needs it?

James C: Information passing from sites to the dashboards aids reliability, and vice-versa. But presently we get patchy results. The test results should go to SAM and everyone reads out of it later. This is infrastructure to publish info between sites and the systems doing the analysis.



## LCG Grid Deployment Board Meeting



HM: Ok, I thought this was a messaging system between people!

Simon Lin: The idea that we need a messaging system seems inevitable. However, it is not the only answer for distributed systems.

James C: We'll rely on Apache – Active MQ – but don't want to roll it out without tests. One test is run via SAM and many clients are linked in, like the DDM dashboard – can we scale with a small number of agents? There is a need for synthetic testing and we also need to talk to APEL for security requirements.

### **7. CCRC'08 monitoring and reporting requirements (Julia Andreeva)**

?: I can't speak for all sites but GridIce is too aggressive – it sucks out information and publishes it all. This is over publishing with the names of all IDs for all jobs!

JA: The new version looks at logs

JT: This was also a system looked at and rejected – you can not control log files any more. This needs to be under site control.

DS: The INFN T1 use of a common file system solves this problem for several CEs situations. It is important that GridIce has access to the log files

James C: I asked Julia to write this up. There are several systems now deployed. If GridIce now has all the experience at parsing logs then we should use that. Perhaps GridIce is the wrong name.

JA: A lot of work has gone into the product.

JT: The problem is the other way around. We need to make sure we only publish info that can be published.

Stephen Burke: Part of the discussion seems to be several things are unreliable. The normal BDII is unreliable – can you fix that?

IB: Information in BDII is some statement about the state of the batch system – not how many grid jobs running.

SB: I thought Julia said the numbers were wrong. You say grid jobs but that depends on the policy at each site. VOView should divide queues by VO.

SB: How will GridIce do this?

JG: If GridIce can separate information then so should the BDII be able to do this

James C: Writing a GIP is not hard once you have got the information, but collating is ...

IB: Perhaps this is not information that should be in information system anyway – job changes

Claudio Grandi: Job history tracking is needed .... I mentioned the RB is linked to the WMS and for this reason the information it can provide is similar

IB: Can you not publish into LB from ....

CG: At the job management workshop in Rome in January we will discuss how to publish to the RB from say CREAM-CE or whatever system is in place (job wrapper built on CE).



## LCG Grid Deployment Board Meeting



JA: A job wrapper would not solve the problem of pending jobs.

CG: This is one reason we would like to integrate.

James C: Publishing is better than trawling – just take the information from the L&B

JT: Great, but this needs to be discussed with sites. There is a tension between monitoring and gathering information from sites and publishing it on websites publishing

IB: Here we are only talking about information in the L&B from the RB.

CG: Publishing means you are in the RB framework – the information is then available to the user and VO manager. You can not access RB information publicly.

IB: Anonymous stats should be published.

JT: If I ask a user what they want, they want information on their jobs. Saying the solution is GridIce....

James C: GridIce sensors – not transport, visualiation etc. Only the sensors.

IB: This is the same issue as for accounting

JT: Right now the experiment dashboards violate these publishing restrictions

MJ: The plan was to include SAM tests in job wrapper, what is the status of this plan?

IB: It stopped because R-GMA could not cope. The option is still there.

James C: Sites could run tests in local Nagios – this addresses how you reduce your granularity.

Notes:

JG: The room for 10<sup>th</sup> January is available – there will be a post GDB meeting.

Note also that Euro 2008 is on June 8<sup>th</sup> so there are warnings about booking well in advance. The week before that is the OGF in Barcelona. We could move the date again, or we could have the GDB in Barcelona.

[Action 0712-5 Investigate holding the June GDB in Barcelona in conjunction with OGF](#)

----- **Lunch 13:10** -----

## 8. WLCG Service Reliability Workshop (Jamie Shiers)

JT: What does the "as text" [see slides] mean?

JS: As in the title of the column

JC: But the times must imply something changes for the experiment?

HR: When the event gets escalated





## LCG Grid Deployment Board Meeting



*DNS load balancing for LFC*

FH: Is this to be used for all services?

JS: Depends if middleware can take advantage of it.

JT: We should just require this.... like the experiments do....

JS: TLA

FH: Is this going to be tested in the certification area?

JS: We should revisit the middleware components and see what can use this

JG: A grid site may not be able to use DNS

TC: we are not using any special features for this... the dynamic factor is nothing special

JS: Focus on the big red squares!

JT: This is a question for CMS – if the box goes down the experiment stops. Is there anything they requested that is less than 1 day resolution?

JS: Yes 8hrs.

JT: Should not ask of sites more than what they have signed up to.

JG: Comment on LRs ... do you think the EGEE PMB is not a stakeholder in this?

JS: [We can discuss this offline](#). We need a grid operation in 2010!

## 9. HEPiX (Michel Jouvin)

JT: You mentioned IHEPCCC

MJ: This is a formal body for organising computing in HEP. When the IHEPCCC was created they knew advice would be needed from time to time. There is no formal mandate but if they issue a request then working groups can be setup – but there is no formal relationship.

ME: It is an advisory body but without a formal relationship – IHEPCCC just seeks advice from technical groups like HEPiX.

JG: There is a slight exception – the file system group came out of HEPiX.

MJ: [The report of which I could mention at the next meeting – there is currently testing of several options on new hardware.](#)

JG: Holger's benchmarking group also came out of HEPiX

## 10. Pilot jobs (John Gordon)

*Experiment frameworks:*

*On the review team...*

JT: We know these [experiment frameworks] do things that should not be done. Given the fact that security does not register high on radar screens it doesn't make sense to have



## LCG Grid Deployment Board Meeting



the people making the mistakes making up the team! We need a “nuts and bolts” person like Steve Traylen.

JG: The EGEE2 second person is not named just yet – they would be a “nuts and bolts” type.

JG: Steve Traylen is a CERN person

IB: I agree, it would be better to have sites on the list. Sites need to contribute the people to do this review

SB: What are their terms of reference? Can they block things?

JG: They make recommendations to the MB and they decide. They also need to reassure sites that this is all good stuff.

SB: So the group does not have to reach a consensus view. Reaching a decision would be difficult with so many people!

Romain Wartel: This is design review vs implementation. I don't know why the source code review is being reviewed if the design itself has not been reviewed/settled.

ME: [Condor \(slide 6\) can be done by BNL](#)

JG: PBS pro?

MH: [We have to test it in pre-production and can do it](#)

JT: A long shot, but is anybody using PBS and not Torque. OpenPBS is still out there? It is sort of invisible because the client tools look the same. Version string computed in some non-robust way.

?: We have started tests for LSF at CERN. We have summarised results on a twiki page at CERN – [perhaps this can be used as the basis for other testing.](#)

David S: [CNAF can be added to the LSF testing.](#)

*LCAS/LCMAPS:*

CG: A C library has been produced and a prototype tested. For details we need to wait for the middleware security part. The Java part is a bit behind. There are problems for some changes as the interface (for lib) is not completely well specified. It is using glxexec directly with GUMS but not with the library that LCAS would use.

JG: glxexec calls?

CG: local or GUMs central. Now the LCAS service is being developed using a library provided by globus that provides the interface for coding in agreement with OSG. LCAS uses the c version.

JG: Who uses the java version?

CG: There are clients – GUMS, CREAM and the new auth framework of EGEE would use java, though there is no final decision on this yet – there will be a presentation at the TCG in January. All the developers are working in java.

PC: The glxexec part uses C, so that is not holding anything up



*Timescales:*

JT: We are missing the requirements from the site side but the gut part is the same across the experiments. At the site level we need to arrange for temporary directories and change of users...this is a virtual instruction at the site level. I would not like 3 or 4 ways of doing the glxexec use [i.e. one way for each experiment]..

JG: We have these frameworks already. If they are ok then we would not ask for changes, but if there are problems asking for a common approach [to using glxexec] is reasonable

JT: This is not yet done

PC: It is being done for LHCb

JG: If nobody has done this then can we can work on an architecture together. We could setup an architects group

JG: It would seem sensible to get the experiments together

JT: We need sites involved too – they have the expertise. [I'll volunteer to be on this team.](#)

JG: I can't see the relevant people

JT: They need to tell us what they need us to do! There has to be an interaction.

JG: Good point ....

[Action 0712-6 – John Gordon to speak with experiment representatives about interfacing glxexec to the experiment frameworks and to establish an architects group \[JT volunteers to be part of team\].](#)

## 11. CCRC'08 (Jamie Shiers)

JT: This is about monitoring and support and hitting targets. It is possible that this has been discussed already. [Trying to test the actual system should also test the MoU numbers for response times.](#) (2hrs and 8hrs). [Which ones of the GGUS tickets are to be responded to within times required in the MoU and who gets the tickets etc.?](#)

JS: There was a useful discussion last week. I'm trying to go through all the areas not covered – 5 to 10 areas required follow up. What you mentioned was not explicitly mentioned but should be.

JT: This is for the wish list, you are right that this may be fiction, but some of the things we signed to do.

JS: James's point about automated reporting needs to be factored in to

JT: When we get GGUS tickets they are flat; we need priority settings

JG: With the helpdesk system we can measure some of this but the MoU has specific times that change during the run time etc.

JT: We do not need to do it perfectly...

GM: CCRC'08 involves getting the list of critical services (i.e. those not down for more than 0.5hrs) but this does not map well with MoU table 3.2 ... will this converge?

JS: There is a mismatch between 8 and 12 .... There is no expectation of 30 minute downtime. VO boxes need to go through these procedures too.



## 5. GSSD Status Report (Flavia Donno)

PC: What is a lost file for CASTOR?

FD: A file that seems to be in CASTOR but the SRM registers it as lost

*BringOnline/Get operations*

JT: It is not clear to me that in discussion anyone knew what these were supposed to do in the first place. ... there was an assumption of only one space token for each space

JG: That's the class

PC: "Paths can be used to select pools" I don't understand this statement

JT: Is pool a defined notion here?

?: no

PC: A file is in dCache if the file is in D1 space, then when you get it you will get it in a default space? There is no disk-to-disk copy?

FD: If you want you can enable disk-to-disk copy.

JG: Experiments need to be clear – if they write into storage token whereby it will go to tape, then when this is recalled from tape it will go back into a default space. Does this match use case for different streams of data?

PC: For Castor we will pass the token. For dCache the size of the space is then different. T1D0 is only used for putting. The size depends on the migration scheme and the latency. For reprocessing we need a huge default space and may need to change space dynamically.

JG: Have the others thoughts about it?

TC: A point for the PC thing, in reprocessing there is a change space for files designed to

...

PC: We do not want to change space for files for reprocessing.

TC: Change space for files was not designed to provide a schedule place for files for reprocessing... if this is not correct then we need to change CASTOR and dCache.

PC: CASTOR wishes to impose its understanding of SRM....

TC: From the use-cases – the ESD current data sets must be on D1 and may not be on tape....

JG: I thought that was shouted down when it was discussed....

PC: I think we may be not talking on the same issue. I'm talking about raw data. We'll have 500 TB to reprocess. The cache is 40TB. We do not want to migrate bunches of 40TB and then migrate it back. This is the purpose of pinning. CASTOR never wanted to implement pinning

TC: Castor implements pinning



## LCG Grid Deployment Board Meeting



PC: If we want to reprocess 400TB we want to bring on line when jobs are being run. It is a rotating pool – this is how we always worked with disk caches

JG: Do you want bigger or smaller caches on input/output. If you doing prepare to get then you'll want the same....

TC: LHC raw on space token ...we must understand how caches will be used and managed at different times.

JG: Other experiments need to think about this as much as LHCb have done.

LdA: For storm (*site status*) this is the production end-point and is used by ATLAS.

PC: Where can we get the list?

SB: The patch looked to be certified yesterday (“Not yet in production”)

JG: Is everyone happy with SRM2.2?

JG: We gave Flavia a mandate to get SRM2.2 into production. What we plan to do in January is to review the mandate. GSSD have a long list of things they want to do and the question is should we let them decide or give direction?

JT: We've gotten in the SE world new users like graduate students. They are using root to analyse data on theirdesktop... some are being told xrootd is the thing to do this... only the security model is the unsupported ALICE one. Is this the correct way?

SG: There are many plug-ins for xrootd

Jean-Philippe:The plugin for gsi is buggy and still needs to be fixed

JT: Most SEs have one headnode... these are real users who need to be able to use some protocol to access this information... root is agnostic... gsidcap etc. It would be nice if somebody could figure out how we expose an SE in a multiple VO environment

PC: You could use gsidcap

JT: The impression people [users] are getting is to use xrootd.

PC: Use gsidcap it works fine.

JT: I'm referring mainly to ATLAS users

JG: [We can ask somebody to come up with options for the next meeting](#)

FD: This is kind of on the radar – how should a user be addressed to access his data. People feel the knowledge is missing – training in this area would be really useful.

JT: What I am hearing from grad students is that they are not in a cpu crunch.

IB: What problem are we discussing?

JT: People at labs wanting to use data

JG: Tier-3?

JT: This is the grad students analyzing the data

IB: [What are the experiments telling them?](#)

JG: It is relevant to ask the experiments what they are doing in this area – area there any common solutions that need to be developed.

IB: ROOT knows how to open files – if it knows how to talk to SRM... gfal...

JT: Tell them to use gfal.

PC: gfal: SURL and that's it

FD: This issue basically comes down to training

JT: We have two students who want to do this



## LCG Grid Deployment Board Meeting



TC: I thought the experiment people are supposed to go via the experiment framework to access their data.

FD: There have been grid schools where such things were discussed, one can make available the pointers

JT: Who is doing the T3 stuff?

KB: For ATLAS, it is Dietrich.

[Action 0712-7: Schedule follow up \(GDB\) discussion on experiment user training and local access to data on SEs. \[Are there common solutions that can be shared?\]](#)

### 5. AOB

JG: This is the last meeting of 2007. It is also the end of era since when we next meet Les will not be the project leader. Under his guidance the project has made a lot of good – the experiments are for example now shipping data and a lot of this is due to Les as he has consulted a lot and been persistent - two good attributes for a project leader! We should thank Les for his contribution and wish him well. [Applause]

During this meeting we have gathered some ideas about things for the next meeting. See you in January.

**The meeting closed at 16:35.**



## 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 lcms 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



## LCG Grid Deployment Board Meeting



Item No.	Description	Owner	Status
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
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 glxec.	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 glxec 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 glxec being used in OSG already and how experiences might be shared.	John Gordon	Open
0710-5	Send statement to MB regarding pilot jobs and glxec. Request MB to consider and forward to CB for comment.	John Gordon	Open
0711-1	Query the GDB list about member feeling for holding the April 2008 GDB in Taipei	John Gordon	
0711-2	Take advice on who to ask (JSPG) about VOMS requirements	John Gordon	
0711-3	Get glxec-on-WNs field tested by some sysadmins	John White/ John Gordon	Progress
0711-4	Pass on the issue of proxies being stored all over the place to the middleware security group	John Gordon	Progress
0711-5	Put together an experiment frameworks review team	JG,IB, DK and BJ	Progress
0711-6	Ask/inform/request sites about testing glxec with the various batch systems	John Gordon	
0711-7	Talk to David Salomoni about common batch system tests	John Gordon	
0711-8	Compile summary of pilot jobs/glxec discussion for MB	John Gordon	Closed
0711-9	Confirm resources available for the CCRC as given in Harry's talk (November GDB)	Country reps	





## LCG Grid Deployment Board Meeting



Item No.	Description	Owner	Status
0712-1	Review Jeff Templon's concern about purchased resource figures ending up in resource spreadsheet but not allocated resources.	Harry Renshall	
0712-2	Gather more input before making a decision on April GDB	John Gordon	
0712-3	Follow up concerns about site-GGUS improvement with ROC managers. What is the process for using savannah vs GGUS tickets and finding out their status	John Gordon	
0712-4	Compile and report on list of SAM/monitoring priorities at January meeting	Piotr	
0712-5	Investigate holding the June GDB in Barcelona in conjunction with OGF	John Gordon	
0712-6	Speak with experiment reps about interfacing glxec to the experiment frameworks and establishing an architects group [nb. JT volunteers to be part of team]	John Gordon	
0712-7	Schedule follow-up (GDB) discussion on experiment user training on the topic of local access to data on SEs [Are there common solutions that can be shared?]	John Gordon	

### List of Attendees

X means attended  
V means attended via VRVS

Country	Member	Present?	Deputy or Technical Assistant	Present
Austria	Dietmar Kuhn	X		
Canada	Reda Tafirout		Mike Vetterli	
Czech Republic	Milos Lokajicek			
Denmark	John Renner Hansen		Anders Waananen	
Finland	Klaus Lindberg		Jukka Klem	X
France	Fabio Hernandez	X	Dominique Boutigny	
Germany	Klaus-Peter Mickel		Holger Marten, Jos van Wezel	X
Hungary	Gyorgy Vesztergombi	X	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	Jacko Koster		Farid Ould-Saada	
Pakistan	Hafeez Hoorani			
Poland	Ryszard Gokieli	V	Jan Krolkowski	



## LCG Grid Deployment Board Meeting



Portugal	Gaspar Barreira		Jorge Gomes	
Romania	Mihnea Dulea			
Russia	Alexander Kryukov		Vladimir Korenkov	
Spain	Jose Hernandez	V	Xavi Espinal	
Sweden	Leif Nixon		Tord Ekelof	
Switzerland	Christoph Grab	V	Allan Clark, Marie-Christine Sawley	
Taiwan	Simon Lin	X	Di Qing	
United Kingdom	Jeremy Coles	X	John Gordon	
United States	Ruth Pordes		Michael Ernst	X
CERN	Tony Cass	X		
ALICE	Alberto Masoni	X	Yves Schutz	
	Federico Carminati			
ATLAS	Kors Bos	X	Stephen Gowdy	X
	Dario Barberis			
CMS	Matthias Kasemann	X	Patricia McBride	
LHCb	Ricardo Graciani		Andrei Tsaregorodstev	
	Nick Brook			
Project Leader	Les Robertson	X		
GDB Chair	John Gordon	X		
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			
Quattor WG	Michel Jouvin	X		
Networking WG	David Foster			
Planning Officer	Alberto Aimar	X/V		

### Others present at CERN

F. Chollet – IN2P3  
 I. Ueda – Tokyo  
 Jamie Shiers – CERN  
 Nick Thackray – CERN  
 John Shade – CERN  
 James Casey – CERN  
 Gonzalo Merino – PIC  
 Patricia McBride – CMS/FNAL  
 Dietmar Kuhn – Innsbruck  
 Harry Renshall – CERN  
 P. Charpentier – CERN/CMS  
 Sue Foffano – CERN  
 Gilbert Poulard – ATLAS  
 Claudio Grandi – INFN



## LCG Grid Deployment Board Meeting



David Salomoni – INFN  
Luca dell'Agnello – INFN  
Stephen Burke – RAL  
Oliver Keeble – CERN  
Ian Neilson – CERN  
Kors Bos - ATLAS  
Jason Shih - ASGC

### **Others on VRVS:**

Duncan Rand - UK