



## Version 0.11.01-2

### Amendments history

Name	Area	Date
<a href="#">Jeremy Coles</a>	All - full edit -> v1.0	<del>20th July 2007</del> <sup>12th March 2007</sup>
<a href="#">Jeremy Coles</a>		
<a href="#">John Gordon</a>	Minor changes section 1	13 <sup>th</sup> March 2007
<a href="#">Jeremy Coles</a>	Summary and attendees	13 <sup>th</sup> March 2007

Formatted: Left  
Formatted: Left

## Minutes of the meeting

CERN CERN, ~~46<sup>th</sup> June~~<sup>7<sup>th</sup> March</sup> 2007

Formatted: Superscript

Agenda: ~~<http://indico.cern.ch/conferenceDisplay.py?confId=8470>~~  
~~<http://indico.cern.ch/conferenceDisplay.py?confId=8485>~~

Formatted: Default Paragraph Font

Formatted: Italian (Italy)

Minutes: Jeremy Coles

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

**Meeting Summary (prepared for MB - 10<sup>th</sup> July**  
~~<http://indico.cern.ch/conferenceDisplay.py?confId=13803>~~

Formatted: Font: 12 pt, Bold

Formatted: Font: 12 pt, Bold, Superscript

Prepared by John Gordon:

Formatted: Font: 12 pt, Bold

Formatted: Font: 12 pt, Bold

Formatted: Font: 12 pt, Bold

Formatted: Font: 12 pt, Bold

Formatted: Font: 12 pt, Bold, No underline

### Accounting

Formatted: Font: 12 pt

#### Storage Accounting (Greig Cowan)

Formatted: Font: 12 pt, Not Bold

Storage accounting information is harvested reliably from the BDII but we have almost no feedback from sites on the correctness of the data. Data is collected per [GlueStorageArea](#) (disk pool). There are known to be multiple counting issues for sites which shares pools between VOs when using the default information providers for DPM and dcache. The latest DPM provider behaves correctly and there is a modified driver available for dcache.

Formatted: Font: 12 pt, English (U.S.)

Formatted: Font: 12 pt

We believe that LHC VOs at T1s do not share disk pools in dcache but cannot be sure.

Formatted: Font: 12 pt

#### Management Issues (John Gordon)

Formatted: Bullets and Numbering

- The number of T1s who publish fully to APEL has been rising. See separate presentation on July 10<sup>th</sup> to MB.



## LCG Grid Deployment Board Meeting



- The CB has been asked for details on Tier2s and their federations where appropriate. This will be used to produce consolidated reports for CRRB.
- Data is being received from OSG for their Tier1s and most Tier2s. All Tier2s will publish when they installed the latest OSG software release very soon.
- Contact has been made with relevant people in NDGF. They are implementing a central SGAS service from which they will publish to the APEL repository.
- Very few sites have switched on userDN publishing. A prototype portal to display DN information exists but requires the policy document.

### Technical Issues (Dave Kant)

There is a patch going through certification dealing with issues on Multiple CEs, User DN encryption, and VOMS roles/groups. If the experiments want to see any of these they should promote the patch #1164 at TCG.

Formatted: Font: 12 pt

### Middleware Issues

#### SL4 (John Gordon, Markus Schultz)

The SL4 WN release has been in production for a month but the uptake has been disappointing. Sites were concerned that the SL4 middleware release did not contain all the rpms required by the experiments that were previously included in SL3. Markus said this was a deliberate choice to remove operating components from the middleware packages as this had been a criticism from sites in SL3. Not all experiments had updated their VO Cards to include any extra rpms that they required. Metapackages would be even better. ACTION 1 – all experiments to have updated their VO cards by Operations Meeting on Monday 9th July. Only ATLAS was present at the weekly operations meeting. ATLAS and LHCb cards are complete. This needs to be actioned urgently as there will be no widespread deployment until the requirements are known.

Formatted: Font: 12 pt

Separate lists of rpms have a risk of circular dependencies and clashes for sites who try to all install for more than one experiment. ACTION 2 Someone (LAL and UK sites have started testing this.) to attempt installing rpms for all 4 experiments on a test box to check out the dependencies. Target date 7 days after Action 1. If there are no problems then the message to all sites is to move to SL4 as soon as possible, and by the end of August at the latest. If there are problems then ACTION 3 - SA3 should help resolve them and/or find out which combinations of experiments are problem-free. Target Action 2 + 7 days.

The experiments all said they were happy to run on SL4 with CMS expressing a strong preference for SL4 for their planned production. I think it is important to progress the actions above this week in order to give sites the all clear to deploy SL4. Otherwise we risk many sites waiting until summer is over before upgrading.



Markus was concerned about the proliferation of releases/versions he was required to support. He proposed

- Release the changes that are in the pipeline
  - In their current incarnation (mostly SL3 + VDT 1.2)
  - WMS, gLite-CE, FTS-2, gfal, lcg-utils .....
- Then concentrate for updates and new releases on:
  - SL4 and VDT-1.6
    - Services and clients
    - First 32 bits
    - Then 64bits for clients and selected services (DPM)
- What about the SL3 sites???
  - They stay with the versions that they have until they can move
  - The gLite build system will be maintained for security updates

Formatted: Bullets and Numbering

Formatted: Bullets and Numbering

Formatted: Bullets and Numbering

Formatted: Font: 12 pt

**Worker Node Disk Requirements (Kors Bos)**

Kors set out ATLAS thinking on WN disk requirements for several models. He was encouraged to try using posix i/o to read merge files from the SE which will reduce the WN disk footprint.

Formatted: Font: 12 pt

**Security**

**Operations (Romain Wartel)**

Romain gave an outline of the planned operational security work. It was good to see some disaster recovery planning in this area.

Formatted: Font: 12 pt

**Policies (Dave Kelsey)**

New top level policy document is in final round of approval by email. Should be with MB soon.

VO Operations Security Policy – draft under discussion by JSPG. Similar form to Site Operations. Should be with GDB in September or October

Grid Services Policy – too soon for generic one. Start with VO Boxes, Pilot Jobs, Portals.

VO Box Policy written last year by working group but never formally signed off. Do this now.

Pilot Jobs – draft 0.1 at JSPG

Data Policy – too generic. Start with Accounting Data



## LCG Grid Deployment Board Meeting

# GDB

### Discussion

▲ There is still a requirement seen for more coordination of how middleware and applications interpret security identity information. There is no one place to discuss this and no-one wanted to start YASG. It was decided to produce a written discussion paper and circulate it to various bodies (MB, GDB, MWSG, TCG, JSPG,.....) and see what falls out. As a starter the meeting identified 7 issues for which the solution/agreement was undecided or ambiguous.

### GLEEXEC

▲ Prepared by John Gordon

### ***Introduction (John Gordon)***

~~Future Meetings: Prague 3-4 April. Tuesday afternoon will showcase the local region. Vancouver 31 August. Transport and agenda will need coordinating with WLCG workshop. Arrive on 30<sup>th</sup> for early start on 31<sup>st</sup>. It wasn't discussed in the meeting but I give notice that I plan to hold the March 2008 GDB away from CERN to avoid the Motor Show. Suggestions welcome for either of the two GDBs to be held outside CERN in 2008.~~

~~Consultation. Pre-GDB agendas will remain flexible depending on content which will be defined well in advance. No proposal to move Face to face MB.~~

~~Countries with Tier1s should nominate a second (non-voting) representative from their Tier2 community. This is to engage these, sometimes large, sites and get their input, not to improve the information flow out to them. Will progress suggestion to delegate task of further engaging the Tier2s.~~

~~Accounting: the reaction to the MB decisions has been disappointing. Almost no feedback on how well APEL reports Tier1 use or on success of Storage accounting. JG has started comparing 2006 manual reports with APEL and will circulate a paper.~~

### ***SL4 (Markus Schultz)***

~~Markus reported that there had been some progress since the February GDB. There was now a buildable WN and UI release but it was not yet installable straight from ETICS. With tweaks it had been installed and the WN tested with an SL3 CE which will likely be a common configuration. The SAM tests were successful on 1/3.~~

~~The tarball workaround advertised in February had not been installed at many sites. Only 9/210 were advertising SL4.~~

Formatted: Font: 12 pt

Formatted: English (U.S.)



## LCG Grid Deployment Board Meeting



A solution for WN and UI looks to be in sight but the other nodes will take longer. Data Management is seen as the next priority since new disk servers which run this are also requiring SL4. In general the experiments seemed less agitated than in February. One commented that they were unhappy but realised they could not force sites to use the current solution. Many people were concerned that we could still have components running SL3 when it stops being supported in October, just before LHC data taking starts.

Another major decision will be required if the gLite WMS or CE are not considered acceptable to the experiments. Continuing with LCG versions will require porting to SL4. This is currently not planned and will take considerable effort. The GDB will continue to track but MB is also advised to monitor this situation closely. Ian's proposal on specifying performance criteria gives MB something to monitor progress against.

Markus also described the issues around providing 32bit gLite on 64bit nodes. He suggested 3 options: Provide 32bit versions of the interpreters which means managing external packages; do this only for Python using the Application Area Python version and forget Perl; ask the SL4 team to add the 32bit binaries to the distribution. Markus favoured the third.

### ***BDII (Laurence Field)***

Laurence reported on issues with the Information Service which has recently appeared to be a bottleneck and cause of many job and test failures. He highlighted load problems with sladdp and timeouts on the top level BDII. He showed correlations of timeouts with numbers of simultaneous queries and data size. Short term fixes include running the site level BDII on a standalone machine, running the CE information provider on the site level BDII and introducing regional top level BDII's. This last suggestion has been widely implemented (60 top level BDII's) but not all clients point to their regional instance. Also the CERN catchall region is too large. Longer term improvements include more caching in the client and separation of static and volatile information. Long term scalability also needs considering.

### ***VOMS***

There were three related talk on deployment of VOMS aware middleware. In a wider discussion on VOMS it was felt that there were two cases for continued coordination: firstly in user issues like registration and secondly in coordination across implementations (storage, batch, ACLs, generic attributes, etc) so that users don't see differences in behavior between sites. This would also help put an agreed WLCG view including OSG and NGDF to the TCG. I will work with people to prepare a mandate for such a group.

### ***Job Priorities (Jeff Templon)***

Mainly what was reported to MB the previous day so I won't repeat. Got agreement from most TIs to deploy this by the end of March so we should have some progress to report to April meeting.



### ***Access Control for Storage (Maarten Litmaath)***

~~Maarten had investigated how VOMS roles/groups could be used to control creation of, and access to, files in the various storage systems of interest to WLCG. In summary DPM and StoRM have full support now, dCache has significant support, Castor has minimal support, and BestMan (DRM) has none. We cannot expect grid wide consistent VOMS-ACL support this year for files or space tokens.~~

### ***Accounting (Dave Kant)***

~~Dave reported that accounting by Primary FQAN (the same as used by Job Priorities) has been deployed in APEL but to work correctly requires a patch which is currently in certification. While the UserDN information is encrypted the FQAN is currently not. While it was foreseen that VOs might eventually want to conceal their work patterns by group it was agreed that there was no reason to encrypt it just now as this would delay deployment.~~

### Detailed minutes

## **1.Introduction (John Gordon)**

~~John welcomed everyone to his first GDB as chairman. Due to this new role he announced that the new UK representative at the meeting would be Jeremy Coles. He asked to be informed of any other changes in representation.~~

~~The meeting moved on to look at future meeting dates. If anyone would like to volunteer to host a GDB John asked them to get in contact with him. For the Prague meeting the pre GDB will be based on items of interest from the Czech Republic and the neighbouring region and will not be just a technical meeting. The plan is for the GDB to finish at 16:00 on Wednesday 4<sup>th</sup> April.~~

~~Michel Jouvin asked if there will also be any pre GDB discussions ahead of the main meeting on the Wednesday itself. John said nothing is planned for this slot at the moment.~~

~~The Vancouver meeting is scheduled for 31<sup>st</sup> August at Triumph. The WLCG workshop is then 1st 2nd September. There will also be an MB the evening before the GDB since it is at the Tier 1.~~

~~Action 0603 1: John to check MB time with Les Robertson and agree intention at the MB.~~

Formatted: Bullets and Numbering



## LCG Grid Deployment Board Meeting



Since the last meeting John has consulted countries/representatives about Tier 2 representation at the GDB. The proposals put forward were:

- 1) Invite all countries with a Tier 1 to nominate a second attendee to attend on behalf of their Tier 2s.
- 2) Progress the suggestion of an individual with the task of consulting and engaging the Tier 2s. (A Tier 2 Tsar)

The consultation will continue. It is likely that different representation models will suit different countries depending on the level of engagement between Tier 1s and Tier 2s and whether a given country has a Tier 1.

There are a number of open GDB actions related to accounting. Issues with normalisation issues etc. are still to be tackled. John will write a paper comparing the manual accounts for 2006 with the APEL data. He will circulate this to T1s. Tier 2 accounting will be looked at from April. Everyone is encouraged to react to the existing actions!

Gilbert noted that some sites are publishing both grid and non grid work into APEL. It is useful for the experiments to know the grid vs non grid proportions.

Kors reminded the meeting that there is still a need to follow up on some policy documents in this area. Action 0703 2 John to follow up on accounting policy documents

### 2.1. SL4 status and plans Introduction (Markus Schulz/John Gordon-)

Formatted: Bullets and Numbering

There are a number of actions to be tied up and these will be followed up outside the meeting. There is one change in membership – Glenn Moloney now represents Australia. He has been on the GDB list for some time.

The next meeting is on August 31<sup>st</sup>. The March GDB will be away from CERN. More volunteers to host away meetings are welcome.

Formatted: Superscript

Vancouver arrangements: A registration page for Triumpf will open within the next few days. The meeting will start at 0900 and finish by 1600. The agenda will hopefully include American T2 presentations (to be confirmed with Ruth), OPN, and a review of the readiness of the services – a summary is needed for the workshop. Additional topic suggestions welcome.

#### **On current hot topics:**

Formatted: Font: Bold



General attributes

Formatted: Font: Italic

Jeff Templon: The problem with generic attributes is it depends on who you talk to hear whether this is playing a role or not. They have no solid definition and are being used in many ways. Erwin Laure: For the VOs it is a means for whatever. From the beginning this attribute has to stay in VO space – it must not sink into generic middleware.

Dario: Generic attributes have no impact on ATLAS scheduling. John: Long range I thought you wanted to use it for tracking by funding region. Dario: Higgs-France for example – but we do not expect that now. John: Ok, this is parked for now.

Claudio Grandi: I always said from an architectural point of view that this may have a role – it is being discussed in OGF etc now. We may have the infrastructure to use later but not for job priorities. For the time being concentrate on groups and roles. John: Is this not in contradiction to Erwin? Erwin: No we may use the attribute for something else. Markus: I am worried by the way this is being discussed. Groups and roles for jobs vs storage are different. John: There are a number of things about using VOMS and security information across the middleware, so will come out in the security discussion later.

Glexec

Formatted: Font: Italic

Dave: This was discussed at the JSPG recently (and the MWSG two weeks ago – covering status updates on the technology). It will come up later in the security discussions.

SL4

Formatted: Font: Italic

John: Are the sites happy to move. Why have not more sites moved already?  
Jeremy: In the UK we seem to have mixed messages and this stops sites moving forwards. It would be useful to have a statement from the experiments that they are ready to move to SL4 now. We also need confirmation of any rpms that are missing from the standard OS install plus gLite 3.1 that the experiments need. This information should go on the CIC portal VO ID card and ideally into meta-rpm packages that can be given to sites. Jeff: Are the experiments happy if we go completely to SL4? Latchezar Betev: Do you mean SL4 everywhere? John: If it works in one place then it works everywhere. The question is more about job streams. Dario: ATLAS need the SL3 libraries as code compiled for SL3 continues running. Olivier Keeble: This is about appropriate dependencies being there for WNs. Markus: If you want to run SL3 compiled code on SL4 then we need to find a way for both VDT versions to be available. Dario: VDT has nothing to do with experiment code. Claudio: If you are using the libraries then you are using it with GSI etc. Laura Perini: There is no link in the software. Latchezar: The experiments are building their own application software. If there is a mixed infrastructure then we build for both until all are at SL4. ALICE is happy with SLC4.  
Matthias: CMS needs SLC4 on WNs. A production release which is imminent works only on SL4. However, the old SL3 code still needs to run.  
John: Are the required rpms in the middleware? Markus: Previously we produced artificially the environment to include all dependencies. For SL4 we have stripped out non-middleware required rpms... now for some sites there will be dependency problems





## LCG Grid Deployment Board Meeting



due to obscure packages not in the middleware. So the suggestion is that the information is entered into the VO cards and then we can discuss if we then create a meta-rpm. SA3 is packaging and distributing middleware not the OS. Jeff: Dependencies in the middleware should be for the middleware. If extra packages are required then experiments need to communicate this to the sites. John: What is the experiment position on meta-pacakges?

Jeff: Do the libraries depend on the OS being run? Olivier: You can ship these together but it is better to have them as clear dependencies to make sure all are satisfied. Markus: The most straightforward solution is to create a package.

Ian: Does not the GDB need to resolve this problem?

Ulrich: At CERN we have upgraded WNs to glite 3.1. There are 2 CEs (1.1.1 and 1.1.2) available and the experiments are invited to test via these. We are planning to put additional CEs into production. Markus: This is another symptom of the problem of the PPS not functioning. It is good that you put up the software and invite testing on production – but if you put software into the PPS and there has been no testing, this just suggests that the PPS is adding latency.

>What happens with circular references?

John: Action – experiments are to get their ID cards updated with additional rpm dependencies. Do we also need an rpm package? Alessandra Forti (VRVS): A meta package would be better for sites. Markus: If a site supports multiple experiments then you may come up with a package that can not be installed! Olivier: Due to conflicting requirements... John: This can be discussed at the operations meeting. SA3 can act as a broker. Les: What is the timing? Olivier: We are not sure what we will find. John: The experiments need to check – we need someone to install the rpms. Les: So what is the timeline? John: By next Monday the experiments should update their ID cards.

Formatted: Underline

Formatted: Underline

Michel: We are running SL4 64-bit and have documented the install for ATLAS. Markus: You installed the old SL3 versions of the middleware which pulled in stuff that m/w did not need. Michel: This was installed and tested for all 4 experiments. We just removed the 7.3 libraries. John: I would like a site that tests the packages once the experiments have defined them. If the packages/details are available then a UK site will check for consistency problems. John: Someone who attends the operations meeting should raise this at the next meeting. [Jeremy will do this].

Formatted: Underline

Formatted: Underline

Formatted: Underline

Formatted: No underline

Formatted: No underline

Les: And what is the timing? I am concerned about the T2s. We seem to be on the point of being ready now. John: If no problem with all experiments then on Monday we can say okay to move. Les: So is the expectation that by the end of August all sites will have upgraded? Jeff: For some sites there are other parties/stakeholders using the facilities.

Markus: I am surprised that you have to ask now. The expectation was for all to be on SL4 by December. Matthias: CMS NEED SL4 for CSA07. John: The high-level MB decision was that T1s were to upgrade in 30 days after the WN package became available



- that time is nearly up. Dario: The ATLAS release is building now on SL3 and SL4. So we can transition for the next year.

Supported flavours (Markus)

Ian: You phrasing is wrong. We do not have the effort to support more versions. Markus: For the new WMS things can not get worse than now. Ian: The WMS is clear. I am not sure we should push out the gLite version of the CE. Ruth: Given the limited effort should any effort be put into the SL3 changes in that are in the pipeline? Markus: These are services. The WMS is able to fulfil the needs of the experiments. The WMS needs a bit of packaging and polishing. For the whole cycle it would take 2-3 months. Latcehzar: We are making an effort on the gLite 3.1 WMS. Markus: That is a check-pointed version. Michel: The only risk is with the WN. I don't see a reason not to use to SL4. Markus: Both have set of use cases. What could keep sites from moving to SL4?

Jeff: Mostly concentrating on SNs. Ian: Sites seem to be crying out for SL4. John: Markus has made a proposal ... Holger: We always wanted to have single services deployed step-by-step. As soon as the SL3 build is compatible with SL4 (interoperable) and guaranteed and tested then this is fine. Markus: You mean the frozen SL3 version. The burden is to trace all updates through all flavours. Of course they need to interoperate. John: This should come up at the operations meeting.

VOMSification (John) - this has to do with different developers usage of security information. We will come back to this in the afternoon. There are outstanding actions - mainly group actions

No August 1<sup>st</sup> meeting  
Few volunteer sites for 2008 meetings. Would like schedule to be available soon.

The update given was similar to that presented to the MB yesterday. For one week now successful UI builds have been possible. Still a lot of work required to get a fully working versions of the UI and WN middleware. There is also significant work needed to "clean" the code. Modifications of YAIM (making it more component based) are in progress.

Jeremy asked about the plans in respect of the LCG RB and gLite WMS. Tony added that support for SL3 ends in October. Markus replied that there were currently no plans to port the LCG RB and that such porting would slow down the move to software which has a longer term future. In addition the added pressure to make the WMS and gLite CE work may be useful. Ian confirmed that if there was a decision to port the LCG RB then it is not clear from where the resources/effort would come. The CE is more critical.

John: Which do we want to rely on for data taking? Are there any other components in this situation and can we set a deadline for decisions in this area? In data management for

Formatted: Font: Bold

Formatted: Normal

Formatted: Underline

Formatted: Underline

Formatted: No underline

Formatted: Font: Bold

Formatted: Normal

Formatted: Underline

Formatted: Underline, Superscript

Formatted: Underline

Formatted: Normal



## LCG Grid Deployment Board Meeting

# GDB

~~example. Markus: Work is competing with requirements for data management. The list of functional improvements currently competes with work on hardening. Jamie: I am nervous with a date in October. It is close to the accelerator start date. Markus: We will not have SRM... and everything on 64 bit in the summer. Jamie: Should we not take the accelerator schedule into account? Markus: The discussion on the CE needs to be started to allow time after any decision is made. Ian: We need to do an assessment in the MB or GDB twice a month. A recent (MB) document mentioned the performance criteria required by the middle of the year and also for the end of the year, and also for the WMS (which was pushed back to INFN) the criteria for burning this into certification. We need to follow up every other week otherwise the developers feel no pressure.~~

~~John: One month ago the experiments were unhappy with the timescales. Is this still the case? Matthias: ... what about schedule/milestone dates? Markus: For slide 9, the times are from today. Matthias: The planning all seems effort based and not milestone based. Markus mentioned that the developers are in a close loop and meeting with others (like SA3) twice per week to track progress. Ian: They are not here having to defend themselves. Although Claudio sits in the TCG, the developers are generally shielded. John: Then you have the backing of this meeting to re iterate the feeling of this meeting to Claudio and the developers. Jeff: I suggest you invite Bob Jones to participate in this discussion!~~

---

~~Markus continued with the second part of his presentation on the 64 bit challenge.~~

~~John: Who has deployed the interim solution and are the experiments happy or unhappy with it? Markus: About 9 sites are publishing SL4 [a comment was made that not all sites are publishing correctly so there may be 12 sites running on SL4]. Matthias: CMS are a little unhappy but we can not force sites to use this interim solution. Markus: There is an update on the PPS. John: I know from the UK response that sites are not happy with multiple moves to SL4. Ian: The tarball was available for some time. John: But it uses different installation methods to what many sites have now adopted. Michel: We [LAL] are running our configuration for over 1 year now. No problems from the experiment side. Main issues are with the middleware. The main problem with VO software was running on 64 bit machines. We can run the CE with some 32 bit machines. Markus: It is inevitable with users and sites having a mixture with various groups moving forward at different rates. John: The grid should be adaptable. The problems come with the data being at a site that does/does not upgrade as needed by the users.~~

~~Gonzalo: Is this the SL3 middleware on SL4? Markus: Yes, packaged in two ways—in tarballs and rpms. For a while the rpm package had a problem with updates but this is now fixed and the solution is being tested on the PPS. Gonzalo: So if a site goes to this mixed state does it need to setup special software repositories for the experiments? Markus: We need to come up with a correct convention to publish this in the information system. Jeff: What about running with other Linux variants? Markus: Use the libraries widely published by LAL~~



### 3.2. ~~BDII – the EGEE Information System Accounting T1-T2 transfers~~ (Andrew Elwell Lawrence Field)

Formatted: Bullets and Numbering

#### Storage Accounting (Greig Cowan)

Formatted: Font: Bold

Gonzalo: How do you differentiate between production and non-production areas?

Greig: This will be done with space tokens in SRM2.2. John: Common storage tokens....

Kors: What is the query resolution? Greig: 8 hours. John: For free space the problem is that you get the aggregate if you sum over all - ATLAS only want the sum for ATLAS not the shared. Jeff: The space is available to several VOs. John: Even within a VO. Greig: The information is in the MySQL database. John: Can you give a query for ATLAS to run directly? Greig: They can use an lcg-infosites query. Kors: We would like to use this in the job submission to see if space is available in the cloud allowing for some contingency. So for example if there is less than a couple of TB then we may not submit to that cloud or site... we would like to automate this part. John: Can't you do this with an ldap query now to T1s?

Michel: Does the query give the real-time situation at the site? Greig: I agree we can't use the accounting database for real-time (job matching) decisions. Kors: Another question relates to user based accounting for storage. Most of production is one user so that's fine, but I'm scared of real users if we do not have separate pools. When the pools are not separate then we are interested.

Jean-Philippe: On the "SRM2.2 and Glue 1.3" slide, it mentions accounting at the space token level. Is that a new requirement? In the pre-GDB in June we agreed we will provide accounting at the disk pool level not at the Space Token level. Is it needed?

John: I thought it was a use case - to account for ESD and production data across the grid. The only way to do this via the space token. Jean-Philippe: You can provide space by groups but can't say if it is for AOD or something else. Kors do you need that?

Claudio: Connects with the authorisation to write on a given partition and then write a given type of data. You can reserve space for a given task. Can reserve the space but what about name? You can use that given tag for storing the data. This should not be mixed with the authorisation level information. Jean-Philippe: Do we need accounting at this level AOD vs ESD? John: We are reserving space at this level so must be accounting it.... Can't this then be published? Space is in the glue schema.

Formatted: Underline

Questions about failure cases.

Failures were not from sequential tests. Perhaps issue (80%) with one of the dCache servers at RAL T2.

Future is to roll out tests to other T1\* and intra T2 channels.

Kors: Many of the problems being seen are also seen by ATLAS



John: Most relate to CASTOR

**Accounting Update (John Gordon)**

Formatted: Font: Bold

Kors: What is your definition of correct? John: That there is agreement between the APEL figures and the institutes own numbers. Kors: What about NorduGrid accounting? John: We will come back to this.

Ruth: OSG report to WLCG not EGEE. Ian: It is a tool. John: OSG are publishing into a repository and WLCG has a view on the data. Claudio: Some sites are comparing DGAS and APEL figures.

Fabio: For sites not using APEL is the encryption algorithm published? John: Yes – on the wiki

Jeff: We reported a discrepancy at the last WLCG MB but never heard back. What is the correct route for dealing with this problem? John: Via a GGUS ticket. Les: If the figures are not correct you need to report this to the system. You've got to report – did you report and it was not correct or is this a problem with the [publishing] mechanism?

Fabio: A comment on Jeff's comment. We have the same problem. We have no way to control that what is sent is visible to the portal. Some recipes are given but we see time-outs. Please describe the protocol we have to use to check the data

- Formatted: French (France)
- Formatted: French (France)
- Formatted: French (France)
- Formatted: English (U.S.)
- Formatted: English (U.S.)
- Formatted: English (U.S.)
- Formatted: English (U.S.)
- Formatted: English (U.S.)
- Formatted: English (U.S.)
- Formatted: Underline
- Formatted: Underline, English (U.S.)
- Formatted: Underline

Les: For T2s we will start reporting CPU to the RRB from September (there next meeting is in October). John: We could go back further.

Les: Currently we don't know who sites are as we need to guess from the name in the GOCDB. We need to know the GOCDB names for sites and how these accumulate into federations. Then we need a simple name for the report – for example a country code with name (it needs to be self-describing and start with 2 character code). Sue is now going through a list of people who may be associated with sites. Make sure CB members reply to this message. We will circulate figures from July, and then from September we will use this information for the RRB report. We will not give a chance to change the figures (as has been done for T1s). Fabio: Then you have to make sure the portal works.

Les: People should complain if numbers are not correct – it is up to sites to say this does not work. John: Some T2s may not have their own accounting to check against. Fabio: Today (and yesterday) there is a problem with the aggregation. If you go to the portal now under France there are Russian sites listed!

Michel: Is the country code just for what is presented on the portal? John: No – it is for the Tier-2 federation. With a site like GRIF that \*is\* a Tier-2 should they give themselves a name like F-GRIF? Les: The names were put into the MoU under countries – no thought was given to the mapping. We can prefix the federation name with a country code.... *Something about new name matching DB name*

Formatted: Font: Italic



Ruth: Will you speak to the requirements on other grids? John: I did ask about the position of user DNs. Ruth: In terms of the GDB. This is WLCG talking about requirements and a question of details. What you talk about here is the grids' need to provide information for whom? (Kant slide 3)

Gonzalo: Will the test on WNs become a critical one for ops? John: The site will not be blacklisted but the COD may raise a ticket. It is not critical in the sense that it may stop the site running.

Jeff: This compares the site APEL DB with the GOC DB. Why are we injecting data into a local DB if we push the data? We use R-GMA to publish to the GOC. What is the site APEL database? Dave Kant: Every site is running APEL ...it is the archiver on the monbox that contains a persistent archive.

Fabio: Will the SAM test work for sites not using the publisher? Dave: OSG currently publish anonymous summaries.

Formatted: Underline

Formatted: English (U.S.)

### 3. Disk space on WNs-VOMRS (MariaKors Bos)

Formatted: Bullets and Numbering

Kors had already mentioned many of the details on the GDB mailing list.

Jeff: I checked our most recent WNs and they have 130 GB available disk space. What will help most with modern batch systems would be to check batch system can pass JDL requirements. The mechanism is in the gLite-CE. John: I understood the gLite-CE is set for memory and CPU requirement passing... Claudio: You can pass whatever the WMS is aware of. If it is in the glue schema then the gLite-CE is able to pass it through! Then there is a need to modify blah for your batch system.

Claudio: Put in the requirement of the job – for example the minimum amount of memory required – and this is used for matchmaking but also passed through as ENV for blah. The name of the parameter should be known to the WMS

John: So only things used for matchmaking can be passed? Michel: My WN has this amount of memory and disk etc. To use blah in this way need you need to have a maximum and that blah with local batch system matches only to those that meet the maximum. John: If WNs can support different levels then the batch system can understand that.

Kors: The disk space price is 3% of the price of the node. 160 GB on a 4-node machine seems small. ??: There is also the architectural thing to consider. What is the problem



## LCG Grid Deployment Board Meeting



with the infrastructure. Where is the problem? Kors: The reported error is “no disk space left” Jos: From the slide is access through a network? You read from the SE and make a local copy on the WN. Why can’t you use the local SE? Kors: Sometimes it works ... but there have been problems with small files. Jos: So you moved to copying to the WN as the network is too slow.

CERN: There is a different architecture for the two schemes. Disk space management on WNs. Make sure there are no clashes on the nodes. Both models are feasible but we need to test them. In principle this has nothing to do with file merging. What works best here will be decided by the access patterns. Directly read/write to mass storage vs copying. The architecture in these cases is different. John: This opens a debate – the sites are aware.

Formatted: Underline

CERN: For new WNs there should not be a problem but these are all going to SL4. For the blah schema we have the possibility in the new release to pass memory requirements (not for LSF) but not disk requirements, however that should be easy to do.

Jeff: Adding attributes is not a good thing to do now.

Formatted: English (U.S.)

Ian: Gaps should not take on what Jeff is talking about. A small focussed discussion is needed

Maria: At ops meeting many people asking for information

GAS:

J: How does this fit with SA3?

M: Tester is in SA3

J: In parallel to your certification Markus?

M: VOMRS is in SA1 not SA3 but it is clear

Ian: Before getting Oracle licence need to understand need to replicate. Who uses it?

M: The code is there – small machine costs

Ian: But effort cost. Can jobs switch VOMS server? Would like feedback from leas/lemaps etc. Use case is not clear. Answer this before more effort

J: Ask Dave...

D: Clearly single point of failure. Oracle db way forward. Security not yet considered.

J: Mirror or replication? Latency issues?

Michel: Surprised this is in VOMRS

J: VOMS..

M: Valid input. WG goes to VOs to ask? Next developer meeting in 1 week in CNAF. ATLAS asked 3 times

Formatted: Underline, English (U.S.)

Formatted: English (U.S.)

I: Use case?

D: Copy database to BNL. Independent

I: Failure is different from failover

Claudio: May be easier than looks. Vomses file can contain >1  
Checking works



~~J: Trying to get sites to manually change config files difficult to progress. To put into auto update takes time.~~

~~Markus: In environment for production quality service to ensure chain not security patch in less than 4-6 weeks is strange. Oracle does not have a turn-around time quicker than this.~~

~~J: Expts have asked for stability. Now talking about VOMS admin v2 mid summer.~~

~~M: Saying to give idea~~

~~J: Part of proper certified release?~~

~~M: yes — testers began on it early March.~~

~~Claudio: Checked with developers — for match making intro generic attributes is not checked in FQAN. New version leas/lemaps fine — no impact from generic attributes. Does not foresee problems but version in leg...~~

~~Jean-Philippe: Data management depends on this — change and queries will fail.~~

~~Jeff: No similar FQAN lib used by all products. Not opposed to adding once checked.~~

**4. VOMS and ACLs in Storage Services: Summary of the GSSD — discussions WLCG Service Collaboration (Flavia Donno Jamie Shiers)**

Formatted: Bullets and Numbering

~~SL5 did not get discussed.~~

~~Jeff: Comment “VOMS integration”. Second bullet. Requested this but uneasy. Suppose site supports just LHCb and have person in ATLAS and LHCb. Checks second item and saves as LHCb~~

~~Flavia: It uses subgroups....~~

~~Conclusion: Come up with a description of implementation and submit to experiments to see if this can satisfy their requirements.~~

Formatted: Underline, Font color: Indigo

Formatted: Underline

~~Okay to have ACLs at the namespace and pool level?~~

~~John: Bigger pool to have quotas~~

~~F: Requires at least a VOMS aware implementation~~

~~Questions were received at various points during the talk... John: My impression was that most regions in EGEE have a top level BDII. The question is how to get resources pointing (leg-utils and RBs etc) at them. Is this for regional coordination? Steve Traylen: We asked the sites to do this recently. User’s select their own top level BDII so they are more difficult to change... ATLAS mentioned that they changed their approach last week — i.e. away from the default configuration on UIs and batch workers. Users can override~~





default settings. One reason users sometimes select alternatives is that some top level BDHs contain extra sites.

Kors: ATLAS checks the top level BDH in region and then goes to the CERN BDH. Users try the default setting first. Fabio: Are all regional top level BDHs supposed to refresh from same source? Lawrence: Yes from the FCR. This is just a web page so should scale it only needs to support the number of top level BDHs (about 60).

John: Ian put forward a document suggesting 200,000 jobs per day per large experiment by the end of the year. Can it cope? Lawrence: I looked at the accounting yesterday. The problem is the clients all querying the same BDH. With deployment changes we can meet these requirements. Ian: Is there something we can do in the next few months to split the load between the static and dynamic information? Lawrence: It depends on priority and effort but could be done. For queries the work needs to be done on the client side so we need to rethink the site level BDH. John: Change clients to talk to site BDHs? Query more locally? Ian: Like a squid cache. John: And this helps because in the LHCb example many of the queries are for static information like the port for gridFTP. Lawrence: And the priority of slapd is so low that when the CE gets loaded it [slapd] gets killed. John: Regional BDHs also get overloaded. Ian: So we should cache information at sites so queries are not going to .... John: So you have a top level and bottom level querying mechanism, will there be a timing issue?

Olivier van der Aa: Is the gLite CE still running the MDS? Lawrence: Yes, we would like to use the BDH. John: What is the action plan? Lawrence: On slide 18 we have started already on the short term issue. Medium term will start soon. Ian: Some items are done - caching for example... leg utils and gfal changes will be done after SRM 2.2 changes. Kors: Is there any region without a top level BDH? Steve: No. But some countries under CERN, like Canada can have a large number of sites.

LUNCH 12:00

### VOMS

#### 3.5. Job priorities SL4 Status June French Tier-2 - (Jeff Templon Markus Schulz Fabio Hernandez)

Jeremy: When you refer to failover for the top-level BDH is it automatic? Fabio: No, it is a synchronised change to use a different BDH.

John: Does the aggressive ramp up include T2s? Fabio: Yes. WN will be released today to production. Tarball release is also ready. 64 bit is same as a month ago - 70% success. UI ready for PPS.

- Formatted: Font: 16 pt
- Formatted: German (Germany)
- Formatted: Font: Bold, German (Germany)
- Formatted: German (Germany)
- Formatted: Bullets and Numbering
- Formatted: German (Germany)



Jeff: Plan for how to handle slac on 64 bit problems. Underscores Lawrence: What problem? In hand

Dario: Linux certification group now starting SL5. Will you be part of it?

Markus: No. Resources split. SL4 working and 64 bit to production services. You chose what we drop but can not add another branch.

J: Having done SL4 work did say SL5 would be easier

M: Structural improvements. Stack should be easier to move.

D: Expts are being asked to participate in this activity

Tony: Fact SL5 certified not criticism GD. In some expts being asked to do this so they have platform for testing. Split between GD and FO. Will not provide services on SL5 until after 2008 run. Dreaming to think Markus should be on this now.

M: Best to be on SL4 first

I: Will set up single build machine for SL5 but little effort in this area — just get idea of problems

Stefano: Tarball available now. YAIM version?

Markus: It is also in release — rpm version.

Holger: new machines. Mixed environment SL3 and SL4. now have SL3 build on SL4. is there any functionality difference with natively compiles.

Markus: We will not update/maintain the SL3 on SL4 distribution.

## 6. Report from GSSD Storage Workshop (Flavia)

Formatted: English (U.S.)

Formatted: Bullets and Numbering

Jeff: On slide 6, how did the other protocols perform? Only GPFS is mentioned. Flavia: All were similar.

Jeff: Slide 8 – this disturbs me (testing v2.2 in production). We are going between two routes – we were told to do things in the Pre-Production System (PPS) and now in production. John: Is this from different sources – typically the experiments? Dario: The PPS sites have few resources behind them. Jeff: Is the plan to test priorities on the PPS ....? Ian: This is access to data – catalogues etc. The manpower involved is setting up catalogues etc. Jeff: There is (experiment) manpower to do tests in the PPS – yes or no? Ian: There is only one copy of the production catalogues/data. Testing jobs and with division between priorities this can be done. John: Scale vs functionality. Kors: In the slides yesterday was reported the stress test of BNL. We make it part of our T0 test. We see with the new SRM that we get better performance and setup end point for this. Jeff: So there is manpower for testing on the PPS. Nick: We would like [the tests] to coincide – to run things in parallel is not a good use of limited resources. Michel: The experiments have said the resources to do the tests is more than that available in the PPS. To change this needs an increase in manpower and resources in the PPS.



Jeff: The slide seemed to imply there is no manpower to do PPS testing so we should not bother deploying things there.

Ruth: I thought the PPS was for system tests of new releases of the gLite software. SRM v2.2 tests are being done now with P2P test systems. People bring up separate SE instances and declare them as 2.2 storage end-point. John: We could go to the MB and say everyone should do this testing through the PPS. Les: We are stuck in the terminology about the PPS. Who will setup the environments for experiments to test. We do not need to talk about the PPS. The practical thing is a beta test to be run by the experiments.

Fabio: You can use them and not publish in the production BDIL. John: If EGEE has a model for testing before components reach production then it should be used. Les: We are really talking about putting dCache through the gLite certification process? This is more a hardware issue. Jeff: If you say run this test in the production environment then you have to accept that this may destabilise the production environment.

Michel: This problem is mainly at T1. For DPM there is not such an issue. Why not sacrifice one T1 to make the move? Patrick: Only in the real production system will we find some of the problems but this not a good argument for not testing before. Flavia's tests are very specific. We don't know if the access profiles of the experiments will work with the implementation. At least some basic experiment tests need to be done.

Les: What is an acceptable confidence level. Take one site first. How much resource do you need to make progress? Flavia: A plan was agreed yesterday [therefore we do not need to deal with this here]. Check the twiki page by the end of week for the rollout plan. – people should look on Monday (Flavia to send a reference – pointer)

Raise at MB next Tuesday.

## 7. Security

### Security Incidents Management in EGEE (Romain Wartel)

Kors: What does “forensics” mean?

Romain: Looking closely for evidence when trying to understand how the incident happened – look at logs etc.

Romain mentioned the intention to create a named list for a “Management crisis room”

Jeremy: How do users stay informed during an incident? How will they know what they can expect? Romain: There are many VOs so not all can be directly involved....Dario: They may want access to data... Dave: Communication is important so we need an 2Out-

Formatted: Underline

Formatted: Bullets and Numbering

Formatted: Font: Bold

Formatted: Font: Not Bold

Formatted: No underline

Formatted: Underline

Formatted: No underline



of-bounds” contact point and perhaps a paper document stating who is in control, who makes decisions etc. Romain: Information flow is massive surrounding an incident.

Ian: On Dario’s questions. Contacting VOs is not just a security group decision. How do you want us to contact you from an operational standpoint? Contact people do not always know what to do. There is responsibility on VOs too. John: In a recent HEPIX talk a major failure was mentioned and that having phone numbers was essential as there was no email. Do you know who to contact? Romain: There are several channels that can be used but we also need to understand responsibilities.

J: Who is going to decide the management crisis room membership? Romain: We [OSCT?] had just started. John: Will you come back with a proposal then.

Jeff: It is useful to have categories for incidents – this needs to be clear – having just one site down is very different from a major grid wide incident. John: Even for a non-grid incident, for example passwords compromised, then certain VOs will want/need to know. Romain: With a credential issue we would contact the VO very quickly.

**Security Policies (Dave Kelsey)**

Dave indicated that he would like approval via email. John: You should differentiate for those people with comments.

*VO Operations Policy*  
Please comment

Formatted: Font: Italic

Formatted: Font: Italic, Underline

Formatted: Underline

Formatted: Font: Italic

*Grid Services Policy*

Dave: For the VO Box Policy – can we just approve it?

Jeff: It probably has nothing (no text) about the service classes. These are the conclusions that Cal wrote for the task force. Kors: No, that exists only in the slides. It needs to be put in a document. John: After those conclusions came out sites were happier running VO Boxes. Dave: It looks like we need to revisit this area. John: There was discussion on SLAs yesterday. The SLA between sites and VOs. There were points about security in that. Jeff: Such as what we do to bring the service back up. John: For operations, Holger suggested security responsibilities etc. Jeff: We are running one VO-box in the “bad class”. We have an unofficial agreement with the VO that if there is any problem we pull the box. Dave: So if sites have SLAs do they need this policy? John: SLAs deal with responsibilities not policies.

Formatted: Underline

Formatted: No underline

*Pilot Jobs Policy*

SLA could cover some of this area too. John: The policy is an overall acceptable use document, a VO policy would be a sub-part of that not the other way around.

Formatted: Font: Italic



On glexec/pilot jobs:

Jeff: Some sites don't want users to do a user switch. Can run in one of two modes.

John: On point 4 – is there another way to do it? Jeff: If this policy gets published before glexec works then there will be a hole. It says we have to use the system provided utility. ALICE does this already.

User-level job accounting

Dave: This is an early draft but comments are welcome.

Formatted: Font: Italic

Security discussion (John Gordon)

Formatted: Font: Bold

John: I mention a couple of things from previous discussions. Do we need a group to look at the different ways security is being applied – e.g. ACLs in GSSD within storage? What about job access. We have the portals group and other people taking VOMS without clear instruction on how to apply roles and groups. Implementations are using the same thing with different access mechanisms. Where should this discussion take place. Already there are many security bodies. The MB or whoever needs to rubber stamp this. Do we extend the MWSG ?

Jeff: I don't think it is a security issue. Are the keys secure enough is a security concern. How do you use them is not security. This is about using stuff in VOMS proxy in different ways Ruth: It is about usability. Jeff: Not just usability. John: In this case it is not just security. Romain: For the security issues it is difficult to put them in one group – for example vulnerabilities. However, there is an overlap of people in the groups. For example the security coordination group. John: This is wider than just relevant to security – it involve users, middleware developers and so on and needs use cases... Ian: Which problem are we trying to solve? John: How we agree to use VOMS credentials. Does the middleware do what users expect?

Ian: I agree with Jeff. This is not a security discussion and we are tying up many things. A technical forum is needed. The job priorities work was not thought out end-to-end because it missed such a group. Jeff: That was done by design – we wanted to try it to progress with it. Ian: There was no overall control; it was not fully thought out how it would be implemented by the many areas. We have so many groups. Should we re-use the middleware security group? Perhaps, but it does not have not all the right people.

Jeff: This gets abstracted into entities. Ian: It is a design team. Ruth: It is an EGEE deisgn team. Ian: Miron is on it. Michel: I thought two issues concerned the FOAN and general attributes. For the discussion of ACLs, should we allow similar roles at the same time?

John: If we have the right people together they may resolve this question. Jeff: Also, the different middleware bits make assumptions about the semantics for matching. The WLMS gathers all and evaluates. lcms is different. There are no VOMS based access rules. There is no unique way to say FOAN A is the same as FOAN B from different implementations. Ruth: In the GDB we can ask someone to go and write up the issues



and test the implementations. Jeff: There is a wiki page that has started on this. Ruth: I am interested in the interoperability. John: We should try to get this list together – Jeff to distribute wiki link or setup a new one.

Formatted: Underline

*List of uncertainties and issues – VOMS.*

Formatted: Font: Italic

Ian: Is this just VOMS? John: Security. Ian: Architecture problems too? John: Perhaps you can come up with a list of issues too. Ruth: Where do you bring up issues now? Ian: Where I can. Some get addressed. On the VOMS topic, where have you discussed how job priorities get implemented for ATLAS in OSG? Ruth: I was told it had no implications Ian: There is no ATLAS request to change priorities in OSG? Ruth: They use the VO task queue in Panda. OSG does not restrict the applications that run. (Jim Shank) said that we do not have to worry about this. John: A test of this would be that production roles work in OSG with the same semantics – same effect. Jeff: Is that not a VO internal issue? John: We need to define the problem better – we will do this by creating a list of issues.

Formatted: Underline

Formatted: Underline

Formatted: Underline

Formatted: Underline

Formatted: Underline

Formatted: Underline

▲ Korsi: Also, NDGF do not have same scheme. Ian: We need to understand what ATLAS is doing. Cronos also has an internal job submission queue. We need to decide what your priorities are here. Dario: We want to try out things... Ian: I need a statement – what is your strategy for job submission to EGEE? “We will test everything” means lots of work and if it is not required in the end then... Roger Jones: Most user job submission is via the RB. This is more than just production. Laura: ATLAS has no intention to stop using the WMS.

John: Job priorities have just highlighted the inconsistencies that come up elsewhere. Dario: People need to be able to submit jobs independently without a single task queue for the VO, unless there is no interference. John: With several instances of Panda then you are subverting a single priority system. Dario: We have two – one for production work and one for users. Both send jobs to sites and somewhere on the sites we need to decide the priorities between the two. Ian: So OSG will need to implement job priorities. Ruth: This is VOMS roles based mapping to accounts, then accounts to batch queues and these have different priorities. Jeff: If you go passed this and have more than 2 or 3 FQAN shares it becomes a lot of work (more than production vs user). With a scheme of 40 different groups and roles, if ATLAS needs this at the batch level then we have a problem. Dario : For this year we have two. Jeff: So handling of the other 38 is in your task queues? It won't work as we are doing it now. If we need to do this in generic middleware then we have a big problem.

John: We will produce a list and air it in a few areas such as the GDB,MB and OSCT.

**The Information System (Lawrence Field)**

Formatted: Font: Bold

[Update on 2mths ago presentation.](#)



Jeff: Increase in latency?

Lawrence: Only if CE drops out. GIP caches dynamic information. If these time out can use cache.

John: Patches—where do the others stand?

Olivier: Indexed BDII in next release. Others at various stages of certification.

Acceptance criteria for the CE (Claudio Grandi)

John: gLite GT4 reliance. Had wider implications?

Markus/Ian: no

Claudio: layer in front of gLiteCE will not in principle needed once use BLAH. One scenario is to keep it

Markus: New UI and WN are all based on GT4. Major 3.1 release.

Claudio: New WN release into production; UI PPS. Based on pre- webservices GT4.

Jeff: Small point—blue thing should appear before LCG part is removed. (slide 5)

—Job priorities (Simone Campana)

Y looking at every view published by site. LCG RB not concerned about this. Use gLiteWMS—look info in block. If submit job as software manager then no view for this. Production matches 2 and 3. (slide 3)

John: generic user—go into short queue but go in with different gi.

Queues do not matter. Scheduling different. Different users have diff priorities in same queue for diff grps.

John: Why

Jeff: Not even clear why mapping was wrong. NIKHEF had problem for 2 weeks on production site. No easy way to notice problem. Capital letters in config file while YAIM looked for all capitals. Broken without knowing.

Sim: 1 not matching. 2. attract jobs pile up and remain until time out. Some work to keep looking for these gatherings.

Jeff: SAM test for this—submit job with/out VOMS role and

John: Implementation—can see shares between VOs implemented. No reason to think it won't work. SGE plug—allows to split shares.

John: Easy to say keep for T1s. But use-cases for T2s who want to split

Ian: As Simone says there are short term issues to fix so ATLAS production can go on. Apart from this it seems mechanism can't work—150+ sites will not keep up with

Formatted: Font: Bold  
Formatted: Font: Bold

Formatted: English (U.S.)  
Formatted: English (U.S.)  
Formatted: Bullets and Numbering



requirements. Not obvious we are doing the right thing. Exclusivity is not what batch system sees — ATLAS should be the sum of everything.

John: Who should revisit the design?

Ian: Someone fresh with an understanding of the batchsystem. Need to consider how do this — asking YAIM to do config is not going to work.

Jeff: Got it out fast because we knew there would be problems. Some people then took it...:

Ian: reason not to do this on the production system. PPS first.

John: Stress testing to check shares etc.

Ian: Need just one test

Erwin: Can do a mock up. Needs to be part of .... Hit production can be damaging.

John: Use cases?

Kors: Brought on to production system What went wrong..

Claudio: Was a short term solution.

Jeff: WLMS did not work out of box. Exposed to — fix bugs... and by this time many T1s deployed.

Markus: Point is (inc coop with expts) should do on PPS. Under utilised. Expensive

John: LHCb case yesterday — not enough data.

Claudio: For some tests but not job priorities.

Ian:

John: Summary from Ian that someone needs to look at this again.

Claudio: JPWG can address membership

Markus: Toruqe/LSF next solution is Condor.

Action 0706 1: Erwin — look at membership and approach of JPWG.

Simone: Short term? Ask T1s to fix and T2s not to deploy.

Fabio: Wait for receipe before deploying?

**FTS 2.0 deployment and testing status (Gavin mcCance)**

John: Do dates clash with FDRs?

Stefano: We have to live with it.

John: Is that the position of other expts?

Yes

Fabio: Sites — the message is that sites should wait till 3 weeks after 18<sup>th</sup> June or then?

Gavin: New client is backward compatible — upgrade as soon as possible.

Fabio: summer period

**CASTOR task force update (Bernd Panzer-Steindel)**

- Formatted: Underline, Font color: Blue
- Formatted: Font color: Blue
- Formatted: Underline, Font color: Blue
- Formatted: Font color: Blue
- Formatted: Underline, Font color: Blue
- Formatted: Font: Bold

Formatted: Superscript

- Formatted: Font: Bold
- Formatted: Bullets and Numbering
- Formatted: English (U.S.)
- Formatted: English (U.S.)





~~John: Flavia said you would mention ACLs and VOMS. VOMS support is low down the list and ACLs require this?~~

~~B: Our focus is as given in slide 10.~~

~~Flavia: Sebastian also made clear ACLs not likely till later in 2008.~~

### ~~Monitoring (Markus Schulz)~~

#### ~~Grid Services Monitoring (James Casey)~~

~~Jeff: Site runs this and the information remains at the site?~~

~~James: Yes. Haven't talked much about the publishing.~~

~~John: And a high level probe passes tes~~

~~James: Up according to SAM and up according to~~

~~Jeff: (Slide 24) Thought only sending high level info off site.~~

~~James: Yes and no. SAM deals with all the gFTP logs too. Every FTS will soon be shipping....~~

~~Jeff: When did someone say this would be done?~~

~~Ian: Two years ago. All the gridview plots relied on this. No DNs shipped.~~

~~Jeff: So long as site can switch on....~~

~~James: No private and only personal is the DN.~~

~~Ian N: Which is not private because everyone agreed it can be used when taking the certificate.~~

~~<https://twiki>~~

~~Jeff: For site Admins. If site has something that works then~~

~~James: Believe probe set adds value to any site. Interface system provided but only ... we do not tell sites to run it.~~

~~Jeff: Have our things that do some of this... don't want to install something that is run remotely. So can turn it off.?~~

~~Ian: Will be asked if service running~~

~~Jeff: So many things monitoring us~~

~~Ian: Start collecting sensors... available to everyone. Sites can run. Have to ship some information about site to repository for analysis to determine if site is available....~~

~~Ruth: James says OSG is collaborating with monitoring group. Model follows what we have been discussing for a while. Can run own probes and have a well defined interface to repository for WLCG. James change specification should involve OSG and EGEE not just as stakeholders but as part of working group. Thought would work on specifications. Prototype and implementation should not be driving specification.~~

~~James: Agree wikis not best place to keep documents. Worked through those problems so more aligned with Erwin....etc~~

Formatted: English (U.S.)

Formatted: Bullets and Numbering

Formatted: English (U.S.)

Formatted: Font: Bold



Ian: Prototype is an intimate part of your mandate. This is not just a discussion group on standards.

Ruth: Would expect James to report also to OSG project.

James: Fair — not at level of information back. Probes written but no infrastructure to run these at site at thr moment. Will pick up at the operations workshop.

Ruth: So there is a plan in the OSG document database to have v1 released soon. Should have as much attention in group.

James: Useful to have another outside view. Different middleware implementations need to appear...

Ruth: You said from outside. My point is that this is from the inside.

Alessandra: Jeff's point. Should have

James: Tests will be contributed from sites. Tests -> availability number. Many tests will not affect availability.

Oxana: NDGF not in the loop so far. Implementation is not the same. Tests may not be applicable. Do not run RB for example. Is it always failing or ??.

James: If you could nominate someone we'll integrate them soon

Formatted: Font color: Indigo

[Chat comments about quality of video conferencing]

**System ... Dashboards and Monitoring (Julia Andreeva)**

Formatted: Font: Bold

Jeremy: Site efficiency reports — which jobs?

Formatted: Font: Bold

Julia: These are only jobs going via RB

Formatted: Font: Bold

Jeremy: Many of the reasons are “unknown”

Formatted: Font: Not Bold

Julia: This is because there is no information back from the RB.

?: Monthly site summary uses just specific applications. ATLAS and CMS this is production. LHCb and ALICE uses pilot jobs...

**Systems Management Working Group (Alessandra Forti)**

Formatted: Font: Bold

John: No questions

Formatted: Font: Bold

**Latest on SAM (Piotr Nyczyk)**

Formatted: Font: Bold

HONE?

Formatted: Font: Bold

Holger: Lots of MC and analysis

Jeff: VO specific tests — how are they run?

Piotr: ALICE run own test suite — run as a cron job. Gssi cron job. LHCb is also cron job but using SAM submission framework.



Jeff: Seen jobs forking on CE and kill them. Do you have an account call SAM ALICE?

Piotr: Yes. Falls back.

Jeff: Irritating log messages.

SAM runs every hour...

Access to database — write only?

Change of tests — should happen after broadcast of changes!

ALICE tests.

John: When do we get back the interface? 1 month or so to finish stanardisation. Many expensive queries

Fabio: Should as a site we be concerned with the current result of those tests? For ALICE site is very red.

Only ops VO published at the moment. Definition of which tests are critical do not make much sense.

Fabio: So sites will be notified when to start taking notice of the expetiment tests?

P: VO tests — more VO and sites.

F: Who negotiate with?

Nick: Someone centrally should decide this — not 1-1 communication with sites. VO says take seriously and then ...

Jeff: SAM review recently. Couple things questioned had to do with interface. Critical tests regardless of test and all ...

P: Visualisation part will be moved to GridView.

John Can we also have upgrades in vis. For GridView. How control this — told sites not to talk to GridView dev.

Ian: Should set up group to gather requirements and priortiesd

**Security (Dave Kelsey)**

Exposure of DN in monitoring — are sites happy?

Jeff: Do users agree to this?

Dave: AUP says “generally yes”

Formatted: Underline, Font color: Indigo

Formatted: Underline

Formatted: Underline, Font color: Indigo

Formatted: Font color: Indigo

Formatted: Font: Bold

Formatted: Bullets and Numbering



## LCG Grid Deployment Board Meeting



Jeff: Two levels — one level you can't turn off.

Julia: Job submission tools also give DN as well as SAM. Expts want this information shown.

Dave: Cases are well made — debugging and for live info

Jeff: Using EGEE RB then user publishes DNs across world.

Julia: Also expt submission tools

Dave: Jobs running at sites and DN being published.

Jeff: Wow then humm... not sure about DN being publishes.

Don: Users agree to use of information for operational/managerial and security purposes in AUP. That phrase must diselcourse of ops data with tradivional data attached.

Dave: Did users understand when signed that this meant publishing to the world.

Don: Language could be constructed in either way. We have done it one way and must be committed to it.

Dave: Useful to have site views

Formatted: Underline

*VO/site trust*

Formatted: Font: Italic

*VO Ops Policy*

Formatted: Font: Italic

Jeff: Not true did not discuss this. Notes from 2006. GDB did this discuss this.

John: What you say about trusting VO software is like what happened with the VO boxes. Have we been through the same liop with glxexec?

Erwin: VO run submission systems — some use RB provided by ... RB keeps binding between user and job. Don's question — how are we sure that these other submission systems keep the binding.

DON: If someone on street picks up one of these and runs with it....

Ian: The way Don presents at Fermi has trust relationship with VOs. Trust CMS that nobody breaks in to the infrastructure. Fine for a ingle site. Issue is how does a given VO make such a relationship with 200 sites.

John: Dave has talked about policy that may cover this and then there is a technical implementation

Ian: If all jobs accept single job to run as multiple users then this will allow logging — glxexec can do logging — but need to trust use

Claudio: If logging done in a central service not a local machine

Erwin: Still trusting information that is being published.

Dario: What is different from now? Now production jobs been put in by many VOs. Trust for Vos already.

John: User is submitting payload to system

Oxana: Differene now is that I ban single user. But if has ...

Dario: ban prod manager = ban VO.

Agreed exec by VO.



## LCG Grid Deployment Board Meeting



Jeff: Strictly using policies defined what they are doing is not allowed. Need to be consistent. Need to formalise trust. Supposed to be doing own work.

Dave: No sharing of credentials.

Ian: Production manager is responsible — not code from 500 diff people.

John: Not well defined — how ATLAS agrees to code

Jeff: Due diligence —

Dave: Need to cope with scaling issue — policy document to address VO services — pilot jobs is just one example.

Jeff: Let them take responsibility

Don: Accountability is the first step. Site also needs to be able to authenticate submitting party — that is a technical issue. Think spooler not pilot jobs.

John: Pilot jobs come in with user credentials.

Formatted: Underline

### **Glexec (John)**

CE acting like a gatekeeper.

Maarten: You say trust gatekeeper but it is the same code — leas/lemaps.

John: Does not use gridmapfile?

M: Leg compute element has been configured to use VOMS or fall back to gridmap...

Jeff: Run gatekeeper as root — more code then glexec.

Jeff: Looking at the EDG gatekeeper!

John: In due diligence should be writing this down.

Erwin: Point on switching UID. What is establish for trust is

John: Moving on to WNs

Formatted: Font: Bold

Formatted: Font: Bold

Erwin: Change of identity on WNs has nothing to do with trust VO and site. Logging point. It's a means to have some sandboxing.

Claudio: Purpose to change is to isolate the environment

User can kill the pilot

John: Stes argue on both — jobs running as other user and

Jeff: If no change then every can not trace back. We've tried this. Not possible to trace all the way back to the user.

Claudio: Can start new thread with new identity which may be used to do something nasty...



Slide 6:

Questions:

a)

Nick: Don't think we care. In both cases the site should have a tool that is the same regardless.

John: If site will not run pilot jobs then will not use?

LHCb: Yes. The if sites fo not then question to funding agencies.

b) Jeff: WN consistency. If looked at ATLAS0017 on WN4 as user on WN8. Datamanagement does not care about unix ID.

Michel: At GRIF

Fabio: leas/lemaps also used to ban

Jeff: Need to have on each

Claudio: Consistency if have shared file system

Fabio: If no web service then need to do something

Jeff: Banned list is nearly static.

F: Distribute to all WNs in this way is not a long term solution

Unique service will come.

Jeff: Not pretty but short term. Consistency of mapping across WNs is dynamic.

Gridmapdir on all WNS — share gridmap dir across all WNs or need a web service.

e)

John: Is there a compromise? ATLAS?

Dario: What matter to us is accounting done correctly

Claudio: Accounting on batch system level. Accounted to pilot not user.

John: If lemmaps logging — could also benefit accounting

Claudio: Keep alive a thread tthen get a usage for batch slot. Can can you know how much thread uses.

John: ID change logged as one user. Another after.

Claudio: Even then can only assign wall clock time.

Ian: If do this switech then previous discussion expts must take account of user level accounting.

Jeff: Only wall clock time of pilot jobs. Yet to see cpu time is properly accounted.

Condor—

Nick: As a site do you care . Audit and traceability. You do not care how we account.

John: Will discuss with Dave. Will think of a compromise and come back.

Ian: Need to get out of way. All sites running pilot jobs — need to decide how going to stop.

John: Most sites see reasonable consistency.

Nick: What are the timescales now?

- Formatted: Underline
- Formatted: Underline, Font color: Indigo
- Formatted: Font color: Indigo



John: Raise at ops workshop?

Formatted: Underline, Font color: Indigo

Ian: yes — what is the timescale for deploying glxec? Question for JRAI.

Formatted: Font color: Indigo

Claudio: In workplan but not priority.

Erwin: From this discussion this priority goes up

Claudio: needs to go back to those doing leas/lemaps dev. OSG timescale — weeks. Of order couple of months.... Till then use shared file system.

Michel: Draw baeks of shared file system workaround

Jeff: Don't know how far it will scale — 10,000 nodes...

Michel: No security draw backs.

Jeff: Need to make a decision — December. Write down — everyone currently running pilot jobs.

John: Pilot jobs pulling in more jobs for same identity.

Ian: Sites seem to say no pilot jobs.

Formatted: Underline

Jeff: Will send around decision from previous meeting.

### Workshop (Jamie Shiers)

Formatted: Font: 16 pt, Bold

One presentation on all expts status — no feedback.

Need names for sessions

John to report on residual services.

Two panel discussions. Site reviews — INFN/GridPP/ATLAS.

Need points for short summaries for CHEP summary.

John: Full scale dress rehearsals

Dress rehearsals — need to pin down dates and plans.

±

### Postscript (John)

Formatted: Font: 16 pt, Bold

Glxec — something came out.

Follow on on Dave on policy issues

Erwin — job priorities group See slides. A registration page for Vancouver will be available soon.

Auditorium:

Jeremy Coles — RAL/GridPP

Tony Cass — CERN

Ian Bird — CERN



# LCG Grid Deployment Board Meeting



- [Harry Renshall—CERN](#)
- [Oxana Smirnova—NDGF](#)
- [Xavier Espinal—PIC](#)
- [Milos Lokajicek—FZU](#)
- [Jim Shank—Boston uni.](#)
- [Hiroshi Sakamoto—Tokyo](#)
- [Dave Kelsey—RAL/JSPG](#)
- [Fabio Hernandez—CC-IN2P3](#)
- [Michel Jouvin—LAL](#)
- [Etienne Urbah—LAL](#)
- [George Vesztegombi—Budapest](#)
- [Luca Dell’Agnello—INFN](#)
- [Holger Martin—FZK/Germany](#)
- [Laura Perini—ATLAS/Milano](#)
- [Gilbert Poulard—ATLAS/CERN](#)
- [Stephen Gowdy—ATLAS/SLAC](#)
- [Dario Barberis—ATLAS](#)
- [Claudio Grandi—CERN/INFN](#)
- [Latchezar Betev—ALICE](#)
- [Sue Foffano—CERN](#)
- [Andrew Elwell—Glasgow/GridPP](#)
- [Greig Cowan—Edinburgh/GridPP](#)
- [Dietmar Kuhn—Innsbruck](#)
- [Erwin Laure—EGEE](#)
- [Frederique Chollet—LAPP/IN2P3](#)
- [Jean-Philippe Baud—CERN](#)
- [Sophie Lemaitre—CERN](#)
- [Alberto Masoni—ALICE](#)
- [Jos van Wezel—FZK/GridKa](#)
- [Alberto Aimar—CERN](#)
- [Jeff Templon—NIKHEF](#)
- [Kors Bos—NIKHEF](#)
- [Simone Campana—CERN](#)
- [Markus Schulz—CERN](#)
- [John Gordon—RAL/Chair](#)
- [Flavia Donno—CERN](#)
- [Jamie Shiers—CERN](#)

## VRVS

- [Ricardo Graciani](#)
- [Alessandra Forti](#)

Formatted: Italian (Italy)

Formatted: English (U.S.)

Formatted: Italian (Italy)

Formatted: English (U.S.)

Formatted: French (France)

Formatted: Italian (Italy)

Formatted: English (U.S.)

Formatted: English (U.S.)

Formatted: English (U.S.)

Formatted: French (France)

Formatted: Italian (Italy)





## LCG Grid Deployment Board Meeting

# GDB

[Jose Hernandez](#)  
[Stefano Belforte](#)  
[Richard Gokieli](#)  
[Donald Petravick](#)

[Pm:](#)  
[Ruth Pordes](#)  
[Christoph Grab](#)

~~Jeff gave an update on progress in the job priorities area. He noted that some of the answers to questions about site setups had strange groups showing up in shares which indicate a country priority (e.g. /ATLAS/country). Fabio questioned what the final stage would be for this "temporary solution". John: This is a short term evaluation of a longer term solution. Jeff: I support what John said. We do not have guarantee this is a final and permanent solution. We are pushing this deployment to avoid mistakes made in the past, which is to design a complete solution before having wider experience. Does it do what is required? ATLAS was clear about the requirements. CMS were similar in their requests. LHCb and ALICE do not care so much (with their generic user ID approach). Frederico: It is not 100% irrelevant for ALICE. A small number of roles are needed but it is not on the critical path. Fabio: Is the CMS information available somewhere? Jamie: It is not known to me (ECM). Fabio: Then we need other roles enabled? Maarten: For the longer term we will probably need something different. There are many worries that this implementation will not scale at all. Do batch systems honour these shares....we needed a workaround for the most urgent issues.... Fabio: I just wanted to make sure this is understood. Jeff: I'm not convinced this will scale—but this is a prototype.~~

~~Kors: This came out of the requirements we posed to solve a few problems like how to set user Monte Carlo with a lower priority than reconstruction. It solves incidents like that where a general user used many hours of the ATLAS T1 share.~~

~~Luca: CNAF deployed a few days ago. I spoke to the LSF plug in developer who confirmed it was working. [Jeff checked but could not see it]. Gonzalo: PIC are deploying the new information provider in the PPS? John: ASGC information system is setup but not publishing correctly. John: RAL has it implemented but not publishing—say 2 weeks. Ulrich: CERN were late in deploying because we were hit by scalability issues. We have shares in production already. We are not yet publishing but can do this quickly after some more checks on things that may not work. Fabio: What is the scale issue? Ulrich: It was with the plugin when there are 15,000-20,000 jobs in the queue. We needed to filter out local jobs. The new plug in provided by Jeff is 2-3 times faster. Fabio: What is this version and where is it!? Jeff: It is listed as an official patch (g-Lite middleware contributed patch) and is now in certification and testing. There is no functionality difference just the way queries are done.~~



Jeff: Having just checked, INFN are not publishing but ASGC seem to have fixed some of their problems.

Action 0703-3: Jeff to send out link to latest patch.

#### 4. Access control for storage (Maarten Litmaath)

Formatted: Bullets and Numbering

Maarten's talk contained the background status and some revealing questions. On slide 4 he mentioned that Grid wide consistent VOMS ACL support is not [expected] for this year. How much [of the functionality] will be required for next year? Can we survive with what we have?

Maarten: At present, only primary FQANs are looked at. John: Compare this with unix where the file is owned by a unix group attached to your shell. But when reading access is based on any of which you are a member....Maarten: It uses the primary group ID unless the directory has a secondary group ID, in which case that is inherited. The ACL says who is allowed to do it but ... Jean Philippe: For the permission to create a file the primary group and secondary groups are used. For directories, then it is either yours or the parent. For space tokens or namespace, DPM checks all primary and secondary groups. We do not have space tokens—files are in the space where placed at put time. For reading only permissions in the namespace are checked.

Kors: Is there a hierarchy? Can an admin remove files from say the Higgs group. JP: The permission to remove is from the namespace. So for "Group Higgs" only people in that group can remove the file. Maarten: Is it a problem to have ATLAS admins to be a member of all groups? Kors: So, it is impossible for a general Higgs user to write in the production area? JP: Yes by default

The talk continued onto service priorities—privileged groups/roles for QoS, higher bandwidth—and matters such as quotas not being an SRM feature. Maarten was asked if he could circulate the report mentioned on slide 7, he said that Flavia would be forwarding it to the list.

John: I would like to know the experiment requirements—can you work with what is available now? Maarten: There is a monitoring subgroup looking at what is missing too. They should have some interaction. Jeff: Do we define the semantics of glue such that it publishes information or move to an accounting sensor on the SE? This needs a decision. Maarten: There has been a lot of discussion. We thought we had allowed for these things to be published by the schema. We can do an LCG schema addition but this may create more trouble than it solves and then it is better to have dedicated sensors.

Maria Dimou: A generic attribute was requested to give priority on transfers for VOMS aware services. It is to be used in one case to identify the path to the storage. We have



~~struggled with getting the requirements in this area. The implementation is promised for March. Maarten: We may use generic attributes to implement some of the things discussed. John: We will have a discussion after the third talk.~~

~~Kors: Slides 2 and 3 show things we can use. No timescale is given for the others. Maarten: This year we can forget about consistent ACL VOMS management. It is not unthinkable that it could even take another year. To get an impression, how nasty would it be if had to wait for availability everywhere? DPM is fine, but the TIs will not have certain features for a while. JP: Different SEs will not support ACLs for example. For this year we provide a service to replicate ACLs from one SE to another.~~

## ~~5.Accounting Using VOMS roles and groups (Dave Kant)~~

Formatted: Bullets and Numbering

~~Encrypted data is now implemented at FZK and RAL~~

~~John: In Maarten's talk, there were things in the glue schema about who was allowed access to storage space. Can that be used? Dave: Yes we would extract it. John: Into something like a Tier 2 tree view? Dave: I need to look into it.~~

~~Ian: This issue of encrypting the FQAN. Has anyone posited this as a problem that needs to be solved? Is there a requirement for doing this? John: It is up to the experiments. The information can be used to identify individuals. Also this would show how much CPU individual physics groups are doing. When will you [the experiments] be worried? Is it a requirement at all? Luca: It could be an FQAN for only one user! Maria: Everyone understands, it was never said the user DN itself should not be public. Ian: The role group part. Maria: In VOMS today the information is viewable! Ian: From the FQAN can you determine the user? Is it a real risk? Do we want to get stuck on this? If worried then we will need to encrypt. John: nobody wants it short term. Maarten: Probably in the longer term we will want to encrypt.~~

~~Jeff: If it is implemented it is important to have the full chain whether encrypted or not. There is proliferation of groups and it will be unpredictable what a first FQAN will be. The APEL system, LCMAPS, gPlazma and DPM will all interpret the outcome slightly differently. LCMAPS uses the first group but wild carding is also possible. DPM starts with the primary and steps through the FQANs until it matches. There are different frameworks for matching so the outcome is arbitrary. Maarten: How can APEL then tell anything at all? It has to be the primary! Jeff: It is obvious to me need to use the same mapping route. John: We need VOMS use cases that have to specify the role they want to take. I want to run this job in role of production manager. It is what the user specifies. The middleware should not be taking account of all possibilities~~

~~Jeff: If you provide a list of different FQANs the request may come into a site where there is an exact match on the primary one. At a second site there is not an exact match. Some things have wild cards that will match /ATLAS/?. Storage may look at the second or third FQANs and come up with different results.~~



**Discussion:**

~~Maarten: It is important sites do this mapping. Users should not be able to shop around if there is not an exact match. In LCMAPS, if it can not be mapped then a fatal error results. We should require other such matching mechanisms to have the same result. JP: For permissions to have correct ownership it must consider all primary and secondary FQANs. Accounting must only be done on the primary group. Permissions must be done using all FQANs.~~

~~John: How is the situation viewed by the experiments? Lat: We have a problem with proxy renewal but this is not really VOMS. Maarten: It is a bug being fixed. Stephen (ATLAS): We have a secondary groups issue much of the time. John: Secondary groups here means those you are a member of but not using. Ian: You are perhaps referring to DPM which supports VOMS but not secondary groups which is in a new version. John: This is implementing ACLs across the site. Jeff: It is also a user education issue. Writing alone is not enough, the user also needs to turn it on when using a proxy. JP: Secondary groups are all FQANs except the first. Nick (LHCb): We want glexec so we can select priorities.~~

~~John: Back to the ATLAS issue. There is no public explanation for a third dimension/view covering the funding agency. What is the use case? Stephen: I think this came up in conversation with French members where they request resources to be set aside. Site resources are not all pledged in MoU and they want to set some aside for specific users. Maria: Attributes were introduced to represent this dimension. It was a surprise but implemented. The problem is how it will work given such a vague requirement. John: The attribute is a random string that can be attached to an individual and this is persistent when a VOMS proxy is obtained. Gilbert: This dimension may also be a physics group for example for a physics conference. John: We can not do "French and Higgs Group" scheduling but can deal with "French Higgs group". Maria: LHCb wanted it [general VOMS attribute] to associate the user DN with their AFS login ID ... after this other VOs were asked if they would use it. Then came nine months of silence. Now everybody wants it but for different reasons. CMS want to use it to give access to specific web pages, perhaps ALICE do not want anything. John: How do you use VOMS proxy on the web? Maria: ... Stephen: A Tier 2 site also asked for priority for their users. John: Are multiple attributes allowed? Maria: Just one that can have different parameters for each VO. Jeff: This underlies the importance of what I was saying. We need one implementation. Tackling on attributes may not be implemented outside a given region.... Kors: It is important to get something out with basic functionality to test that is prototype early. Maarten: Most users will use one VOMS proxy, it is a sparse matrix. Most users do not have Admin needs. There may be 20 groups but any individual may be in 2 perhaps.~~



~~John: Is there still space for a coordination group. There was an action for a group to come up with a new mandate.~~

~~Action 0703-4 John Gordon to follow up on a VOMS coordination group mandate.~~

~~John: Are we happy? The TCG is well defined but missing Nordugrid and OSG etc. Are the experiments happy that all things are being fed through? The TCG is more about setting priorities but does not commission work too.... Ian: It does!~~

~~Nick: The requirements from the GDB could be useful expressed directly to the TCG. John: How do we take this forward. Set up a sub group? Ian: The issue here is that there are different people in the TCG and GDB representing the same group, so the two see different priorities based on the personal input. It is good to see the GDB requirements but then we need to avoid the TCG experiment representatives coming up with different priorities. John: How do we formally take this forward? There are no volunteers to setup sub group. Maria: At a workshop last week the smaller VOs did not know about the TCG as being the place to submit requirements. John: The meeting here is essentially for WLCG stakeholders, it is not a GDB for everyone. Ian: NA4 is setup for smaller VOs — Cal is vocal in the TCG about opinions expressed to him. John: Maria, perhaps this is feedback for Cal.~~

~~Action 0703-5: John to refer Cal to Maria concerning the representation of some smaller VOs.~~

## ~~6.GDB March 07 News of reporting and resource tables (Harry Renshall)~~

~~There was a brief discussion about using the TapeIDisk0 terminology in respect of ALICE. Harry agreed to change slide 3 wording. [His point was that ALICE manage tape space and that impacts disk but they do not manage the disk — point 3].~~

~~For the ATLAS tests: RAL — has not said when it will be ready. It is currently testing CASTOR with ATLAS. ASGC will be in but taken out for a power upgrade.~~

~~Gonzalo: PIC disk put in place gets filled quickly. It is now at 99% used.~~

~~On the CMS part:~~

~~Fabio: Is it the responsibilities of sites to clean tapes?~~

~~Harry: The experiments will not recycle tapes so this is up to the sites. They will clean the catalogues but I am not sure about disk.~~

~~Gilbert: Not all T2s have signed the MoU. Can we get a clear view on those that have yet to sign?~~

Formatted: Bullets and Numbering



### 7. ~~Grid Storage System Deployment (GSSD) (Maarten)~~

~~There will be a continuation of the storage classes working group with an enlarged scope.~~

#### ~~John's postscript on topics for future meetings:~~

- ~~-We hope SL4 is not an issue next time but would like to hear that—status report.~~
- ~~-Taken an action to review working groups. An update on the status of the Quattor working group is overdue.~~
- ~~-Progress towards SRM 2.2~~
- ~~-Progress on job priorities~~
- ~~-Mechanism for GDB input to reach the TCG/developers~~

### 9.8. AOB

There was no other business.

MEETING CLOSED AT ~~17:00~~16:50

#### Actions:

Item No.	Description	Owner	Status
0602-4	Phrase the requirement on how to use policies in the WLMS	Cal Loomis	Open

Formatted: Bullets and Numbering

Formatted: Bullets and Numbering

Formatted: Bullets and Numbering

Formatted Table



## LCG Grid Deployment Board Meeting



Item No.	Description	Owner	Status
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 lmaps 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
<a href="#">0704-1</a>	<a href="#">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)</a>	<a href="#">Michel Jouvin</a>	<a href="#">Open</a>
<a href="#">0704-2</a>	<a href="#">Follow up on VOMS coordination group mandate wording with Maria Dimou</a>	<a href="#">Ian Bird</a>	<a href="#">Done</a>
<a href="#">0705-1</a>	<a href="#">Get feedback from Markus and Alessandra on previous feedback from sites on glexec.</a>	<a href="#">John Gordon</a>	<a href="#">Open</a>
<a href="#">0706-15-2</a>	<a href="#">Check use cases and VOMS need for failover with the developers and VOs</a>	<a href="#">Maria Dimou</a>	<a href="#">Open</a>
<a href="#">0706-2</a>	<a href="#">Provide description of implementation(s) of VOMS based ACLs and submit this to the experiments to confirm it satisfies their requirements.</a>	<a href="#">Flavia Donno</a>	<a href="#">Open</a>
<a href="#">0706-3</a>	<a href="#">Review the membership and approach of the Job Priorities Working Group</a>	<a href="#">Erwin Laure</a>	<a href="#">Open</a>

Formatted Table



## LCG Grid Deployment Board Meeting



Item No.	Description	Owner	Status
<a href="#">0706-4</a>	<a href="#">Nominate someone to join the grid services monitoring work</a>	<a href="#">Oxana Smirnova</a>	<a href="#">Open</a>
<a href="#">0706-5</a>	<a href="#">Follow up on how best to proceed with site-experiment negotiation on what VO SAM tests are to be monitored</a>	<a href="#">John Gordon</a>	<a href="#">Open</a>
<a href="#">0706-6</a>	<a href="#">Setup group to gather and prioritise GridView requirements</a>	<a href="#">Ian Bird/ John Gordon</a>	<a href="#">Open</a>
<a href="#">0706-7</a>	<a href="#">Follow up c) with Dave Kelsey</a>	<a href="#">John Gordon</a>	<a href="#">Open</a>
<a href="#">0706-8</a>	<a href="#">Raise glexec questions at the Stockholm operations workshop</a>	<a href="#">Ian Bird</a>	<a href="#">Open</a>
<a href="#">0707-1</a>	<a href="#">July actions TBC</a>		
<a href="#">0707-2</a>			
<a href="#">0707-3</a>			
<a href="#">0707-4</a>			

Formatted Table

### List of Attendees

X means attended  
V means attended via VRVS

Country	Member		Deputy	
Austria	<a href="#">Dietmar Kuhn</a>	X	-	
Canada	<a href="#">M Vetterli</a>		<a href="#">R Tafirout</a>	X
Czech Republic	<a href="#">Milos Lokajicek</a>		<a href="#">Jiri Kosina</a>	
Denmark	<a href="#">John Renner Hansen</a>		<a href="#">Anders Waananen</a>	
Finland	<a href="#">Klaus Lindberg</a>		<a href="#">Jukka Klem</a>	X

Formatted: Centered

Formatted: English (U.S.)





## LCG Grid Deployment Board Meeting



Country	Member		Deputy	
France	Fabio Hernandez		Dominique Boutigny	
Germany	Klaus Peter Mickel		Holger Marten	
			Jos van Wezel	
Hungary	Gyorgy Vesztergombi	X	Dezso Horvath	
India	P.S Dhekne		B. Vinod Kumar	
Israel	Lorne Levinson	√	-	
Italy	Mirco Mazzucato		Luciano Gaide	
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	
Russia	Alexander Kryukov		Vladimir Korenkov	
Spain	Manuel Delfino		Xavier Espinal	
Sweden	Niclas Andersson		Tord Ekelef	
Switzerland	Christoph Grab	X	Mario-Christine Sawley	
Taiwan	Simon Lin	▲	Di Qing	X
United Kingdom	John Gordon		Jeremy Coles	
United States	Ruth Pordes	▲	Bruce Gibbard	
CERN	Tony Cass	X	-	
ALICE	Alberto Masoni	X	Yves Schutz	
-	Federico Carminati	X	-	
ATLAS	Gilbert Poulard	X	Laura Perini	
-	Dario Barberis		-	
CMS	Lothar Bauerdick		Tony Wildish	
-	Stefano Belforte	X		
LHCb	Ricardo Graciani		Andrei Tsaregorodstev	
-	Nick Brook	√	-	
Project Leader	Les Robertson		-	
GDB Chair	Kors Bos	X	-	
GDB Secretary	Jeremy Coles	X	-	
Grid Deployment Mgr	Ian Bird	X	Markus Schulz	X
Fabric Manager	Bernd Panzer		-	
Application Manager	Pete Mato Vila			
Security WG	David Kelsey		-	
Quattor WG	Charles Loomis			
Networking WG	David Foster	X		

Formatted: English (U.S.)

Formatted: English (U.S.)

Formatted: English (U.S.)

Formatted: English (U.S.)

Formatted: English (U.S.)

Formatted: English (U.S.)

Formatted: English (U.S.)

Formatted: English (U.S.)

Formatted: English (U.S.)

Formatted: English (U.S.)

Formatted: English (U.S.)

Formatted: English (U.S.)

Formatted: English (U.S.)



# LCG Grid Deployment Board Meeting



Country	Member		Deputy	
Planning Officer	Alberto Aimar	X		

Country	Member	Present?	Deputy or Technical Assistant	
-	-	-	-	-
Austria	<a href="#">Dietmar Kuhn</a>	XX	-	-
Canada	<a href="#">Reda Tafirout</a>	V	<a href="#">Mike Vetterli</a>	-
Czech Republic	<a href="#">Milos Lokajicek</a>	-	-	-
Denmark	<a href="#">John Renner Hansen</a>	-	<a href="#">Anders Waananen</a>	-
Finland	<a href="#">Klaus Lindberg</a>	-	<a href="#">Jukka Klem</a>	XX
France	<a href="#">Fabio Hernandez</a>	XX	<a href="#">Dominique Boutigny</a>	-
Germany	<a href="#">Klaus-Peter Mickel</a>	-	<a href="#">Holger Marten, Jos van Wezel</a>	Formatted: German (Germany)
Hungary	<a href="#">Gyorgy Vesztegombi</a>	-	<a href="#">Dezso Horvath</a>	Formatted: German (Germany)
India	<a href="#">P.S Dhekne</a>	-	-	-
Israel	<a href="#">Lorne Levinson</a>	V	-	-
Italy	<a href="#">Mirco Mazzucato</a>	-	<a href="#">Luciano Gaido</a>	-
Japan	<a href="#">Hiroshi Sakamoto</a>	X	<a href="#">Tatsuo Kawamoto</a>	-
Netherlands	<a href="#">Jeff Templon</a>	XX	<a href="#">Ron Trompert</a>	-
Norway	<a href="#">Jacko Koster</a>	-	<a href="#">Farid Ould-Saada</a>	-
Pakistan	<a href="#">Hafeez Hoorani</a>	-	-	-
Poland	<a href="#">Ryszard Gokieli</a>	X	<a href="#">Jan Krolkowski</a>	-
Portugal	<a href="#">Gaspar Barreira</a>	-	<a href="#">Jorge Gomes</a>	-
Romania	<a href="#">Mihnea Dulea</a>	-	-	-
Russia	<a href="#">Alexander Kryukov</a>	-	<a href="#">Vladimir Korenkov</a>	-
Spain	<a href="#">Jose Hernandez</a>	-	<a href="#">Xavi Espinal</a>	-
Sweden	<a href="#">Leif Nixon</a>	X	<a href="#">Tord Ekelof</a>	-
Switzerland	<a href="#">Christoph Grab</a>	V	<a href="#">Allan Clark, Marie-Christine Sawley</a>	Formatted: French (France)
Taiwan	<a href="#">Simon Lin</a>	X	<a href="#">Di Qing</a>	-
United Kingdom	<a href="#">Jeremy Coles</a>	X	<a href="#">John Gordon</a>	-
United States	<a href="#">Ruth Pordes</a>	XX	<a href="#">Michael Ernst</a>	V(om)
-	-	-	-	-
CERN	<a href="#">Tony Cass</a>	XX	-	-
ALICE	<a href="#">Alberto Masoni</a>	XX	<a href="#">Yves Schutz</a>	-
-	<a href="#">Federico Carminati</a>	X	-	-
ATLAS	<a href="#">Kors Bos</a>	XX	<a href="#">Stephen Gowdy</a>	VX
-	<a href="#">Dario Barberis</a>	X	-	-
CMS	<a href="#">Matthias Kasemann</a>	XX	<a href="#">Patricia McBride</a>	-
LHCb	<a href="#">Ricardo Graciani</a>	V	<a href="#">Andrei Tsaregorodstev</a>	-
-	<a href="#">Nick Brook</a>	V	-	-
Project Leader	<a href="#">Les Robertson</a>	XX	-	-
GDB Chair	<a href="#">John Gordon</a>	XX	-	-
GDB Secretary	<a href="#">Jeremy Coles</a>	X	-	-
Grid Deployment Mgr	<a href="#">Ian Bird</a>	XX	<a href="#">Markus Schulz</a>	XX
Fabric Manager	<a href="#">Bernd Panzer</a>	-	-	-



# LCG Grid Deployment Board Meeting



<a href="#">Application Manager</a>	<a href="#">Pere Mato Vila</a>	-	-	-
<a href="#">Security WG</a>	<a href="#">David Kelsey</a>	<del>XX</del>	-	-
<a href="#">Quattor WG</a>	<a href="#">Michel Jouvin</a>	<del>XX</del>	-	-
<a href="#">Networking WG</a>	<a href="#">David Foster</a>	-	-	-
<a href="#">Planning Officer</a>	<a href="#">Alberto Aimar</a>	<del>XX</del>	-	-

## Others present at CERN

- [Greig Cowan - UK](#)
- [Gonzalo Merino – Spain](#)
- [Schwickerath Ulrich – CERN](#)
- [Miguel Santos – CERN](#)
- [Latchezar Betev – ALICE](#)
- [Claudio Grandi – INFN](#)
- [Luca dell’Agnello – INFN](#)
- [Alberto Masoni – ALICE](#)
- [Laura Perini – ATLAS](#)
- [Gilbert Poulard – ATLAS](#)
- [Oliver Keeble – CERN](#)
- [Sue Foffano – CERN](#)
- [Patricia McBride – CMS](#)
- [Frederique Chollet – IN2P3](#)
- [Jean-Philippe Baud – CERN](#)
- [Erwin Laure – EGEE](#)
- [Jamie Shiers -EGEE](#)

**Formatted:** Font: (Default) Arial, 10 pt, Bold, French (France)

**Formatted:** Left

**Formatted:** English (U.S.)

**Formatted:** Italian (Italy)

**Formatted:** German (Germany)

**Formatted:** English (U.S.)

**Formatted:** French (France)

**Formatted:** French (France)

## On VRVS :

- [Dave Kant](#)
- [Alessandra Forti](#)
- [Etienne Urbah](#)
- [Alvaro Fernandez](#)
- [David Colling](#)
- [Martin Gasthumber](#)
- [Oscar Koseruo](#)
- [Don Petravick](#)

**Formatted:** Italian (Italy)

## Others present at CERN

**Formatted:** Font: Bold, Not Italic

**Formatted:** Font: Bold, Not Italic



## LCG Grid Deployment Board Meeting



Jamie Shiers

Formatted: Font: Not Bold

Harry Renshall

M. Lameme

Formatted: Font: Not Bold, English (U.S.)

Flavia Donno

Formatted: Font: Not Bold

Sue Foffano (CERN)

Simone Campane (CERN)

Formatted: Font: Not Bold, Italian (Italy)

Neechaersky Andrey (CERN)

Formatted: Font: Not Bold

Steve Traylen

Formatted: Font: (Default) Times New Roman, 12 pt, Not Italic, English (U.S.)

### **Also present in the meeting room:**

Steve Traylen (CERN)

Matthias Kasemann (CMS/CERN)

Michel Jouvin (France)

Formatted: English (U.S.)

Oliver Keeble (CERN)

Jamie Shiers (CERN)

Stephen Gowdy (ATLAS/SLAC)

J Knobloch (CERN)

Formatted: English (U.S.)

Luca del'Agello (INFN CNAF)

Gonzalo Merino (PIC)

Harry Renshall (CERN)

Ulrich Schwickerath (CERN)

T Kleinwort (CERN)

Fabio Hernandez (CC-IN2P3)

### **Other on VRVS**

Jose Hernandez — Madrid

Frederique Chollet — Annecy

Marek Domaracky — Bern

Olivier van der Aa — London

David Colling — London

Dave Kant — RAL

Pete Gronbech — Oxford

Gabriel Stocica

Lief Nixon — Linkoping

Frederique Chollet

Greig Cowan

Formatted: Italian (Italy)

Stefano Belforte

Gonzalo Merino

Richard Gokieli

Formatted: Italian (Italy)

Alvaro Fernandez (IFIC)

Juergen Knobloch

Formatted: English (U.S.)

### **Additionally on VRVS PM:**

Pierre Girard — Lyon

Formatted: Font: Not Bold, English (U.S.)

Paul Gelissen — Bern

Formatted: English (U.S.)

Jos Van Wezel — Karlsruhe



**LCG Grid Deployment Board Meeting**



~~Les Robertson — CERN~~  
~~Elizabeth Sexton Kennedy — Switzerland~~  
~~Helene Cordier (Lyon)~~  
~~Owen Syge (DESY)~~