



Version 1.11-2

Amendments history

Name	Area	Date
Jeremy Coles	All – full edit -> v1.0	11-05-07 2th March 2007
Jeremy Coles	Minor revisions -> 1.1	30-05-07
John Gordon	Minor changes section 1	13 th March 2007
Jeremy Coles	Summary and attendees	13 th March 2007

Formatted: Left
Formatted: Left

Minutes of the meeting

~~CERN~~^{CERN}, ~~2nd May~~^{7th March} 2007

Formatted: Superscript

Agenda: ~~<http://indico.cern.ch/conferenceDisplay.py?confId=8470>~~
~~<http://indico.cern.ch/conferenceDisplay.py?confId=8472>~~

Formatted: Default Paragraph Font
Field Code Changed
Formatted: Hyperlink
Formatted: Italian (Italy)

Minutes: Jeremy Coles
Attendees: Please refer to list at the end of the minutes

Meeting Summary

Formatted: Font: Bold, No underline

Prepared by John Gordon;

Formatted: Font: 12 pt
Formatted: Font: 12 pt, Not Bold
Formatted: Font: 12 pt, Not Bold

Organisation (J Gordon)

The August 1st meeting is cancelled. Proposal to move the October meeting from 3rd where it clashes with EGEE Conference to 10th. Will be confirmed or rejected soon.

Spain and Germany have changed representatives. Other countries are encouraged to keep their membership details up to date. Countries with a Tier1 should nominate a Tier2 representative too if appropriate.

Middleware Issues

SL4 (M Schulz)

WN has been through one round of testing in 32 bit mode on PPS, bugs fixed, due for release. UI and WN in 32bit are top priority – Now in PPS ready for release. Data management services are next (for SEs) as new storage hardware usually needs at least SL4. Most other components are ready for testing. 64 bit is a bigger problem. There is possible confusion over the status of gLite3.1. The components are not uniformly SL4. For example, the current experimental WMS 3.1 is SL3.



WMS&CE (I Bird)

At the March MB Ian proposed some evaluation criteria for WMS. He gave a status report. gLite WMS 3.1 has achieved 15,000 jobs/day for 7 days (criteria 10k/day for 5 days). 0.3% failures with no restart. All completed when restarted. L&B service is capable of much higher rates and is not a blocker. This is encouraging.

Current status of gLite CE. Close to 100% success of job submission – after resolving a number of timing issues with Condor. Submissions of 6000 jobs to a CE (max ~3000 at any time). Several Condor issues were found – not yet clear on a timescale for achieving the criteria

Not proven yet. Fallback proposal (a) Keep the LCG-CE “as-is” - there is no effort to port to SL4 (which implies GT4 and potentially many issues)

Deploy either on SL3 nodes (or SL4 with Xen/SL3). Contrary to previous reports, SL3 support will not stop in October 07 (only SLC3). RHEL3 security patches will continue so it is feasible to continue with LCG-CE on SL3.

We should set up a CREAM instance in parallel and subject it to the same testing procedure. However, CREAM currently lacks GRAM interface for compatibility with other projects.

Security (D Kelsey)

Dave presented

- An agreed final version of Grid Site Operations Policy – this should now be approved by the MB. It will be presented in June.
- An agreed good draft version of Grid Security Policy top-level document. Replacing the very out of date version is urgent. Also need sites and VO's to “sign”
- A Grid Policy on Handling Logged Personal Information which is relevant to user-level accounting privacy issues. This has not yet discussed by JSPG

It is foreseen that OSG and EGEE have top-level documents. Does WLCG need one too? This will be discussed by the MB.

Filesystems WG (M Jouvin)

Michel reported from the workshop held during HEPiX in April in DESY. A workplan has been agreed and evaluation started. Most Tier1s and many other sites are involved. The target is a final report for HEPiX in spring 2008.

Top 5 Issues

The afternoon was devoted to technical feedback in various areas of middleware development and deployment in response to the areas raised in the Top5 Issues presented by each experiment to the MB during April. Technical summaries of groupings of issues were presented to GDB.

Formatted: Bullets and Numbering



Castor (T Cass)

Tier0 Issues being addressed by a special task force. A new LSF Plug-in should address many of these issues. For Tier1s disk1 storage classes are the highest priority. There is no firm plan for the remaining issues yet but it is being reviewed. SRMv2.2 implementation is proceeding on track from a slow start. xrootd development is done by SLAC so negotiations on developments and support are required.

Integration & testing of data and storage management components (J Shiers)

The main outstanding issue that has not yet been addressed is multi-VO testing of Tier0-Tier1 transfers to demonstrate the nominal rates. This should be come feasible soon when ATLAS restart bulk transfers. CMS are repeatedly transferring.

SRMv2.2 (F Donno)

The main implementations tested for the functionalities requested. SRM v2.2 is available for the experiments to test. It is very important to have the experiments on the pre-production test-bed testing the environment as soon as possible in order to understand if SRM v2.2 is ready for production

Job Management (C Grandi)

As Ian Bird reported earlier, WMS 3.1 is making good progress. It addresses many of the issues. An outstanding issue (for future GDB?) is the deployment of glexec, firstly on the CE and then on the Worker Nodes. The remaining question is whether the sites will accept glexec at all. Some development for a service mode is still required.

FTS and Data Storage Management (I Bird)

FTS version 2.0 is certified – pilot service used by experiments for testing. It has interfaces to both SRM v1.1 and SRM v2.2 and includes VOMS-aware proxy renewal (ALICE) and delegation (avoids need to send passwords)

Other issues being addressed (prioritized in FTS workshops):

- Monitoring improvements – highest priority
- Deployment and configuration issues (channels etc)
- Timeouts in the client

ALICE made a statement that FTS will not be used for T1 <-> T2 transfers This has to be contrasted with WLCG:s view that FTS is the only supported transfer mechanism.

Storage Accounting has been implemented but the data needs to be checked by sites.

Information Service See Laurence Field's report to March GDB

Short term issues have been completed (mainly a deployment issue)

Formatted: Bullets and Numbering



LCG Grid Deployment Board Meeting



- Put the site-level BDII on a stand alone node
- Run the CE information provider on the site-level BDII
- Introduce regional top-level BDIIs

Medium term (work in progress)

- Improve the efficiency of the queries
- Add some form of caching in the existing tools
- Improve the query performance of the BDII service

Long Term (will focus some effort here now)

- Move away from using the LDAP client directly
- New information system client (uses caching, site BDII, multiple top level BDII)
- Re-consider the architecture
- Split information into static and dynamic
- More caching at the site-level

Top5 Summary

The LHC Experiments expressed content with the technical status reported but would await progress in coming months. There will be future reports to GDB but the Management Board should consider the management plan for this work.

Future Meetings.

Possible topics for future meetings include: CASTOR, Monitoring, T1-T2 interactions, Accounting, VOMS.

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 30th for early start on 31st. 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.~~

Formatted: Bullets and Numbering

Formatted: Bullets and Numbering

Formatted: Bullets and Numbering



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.

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 sladp 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 BDIIs. This last suggestion has been widely implemented (60 top level BDIIs) 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.



~~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~~

Formatted: Bullets and Numbering



LCG Grid Deployment Board Meeting

GDB

~~neighbouring region and will not be just a technical meeting. The plan is for the GDB to finish at 16:00 on Wednesday 4th 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 31st 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.~~

~~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-)

~~Refer to John's slides for details. The introduction covered membership (the current list was reviewed online) and meeting dates. It was agreed that the August meeting is not~~

Formatted: Bullets and Numbering

Formatted: Underline



LCG Grid Deployment Board Meeting



needed and was cancelled. Any clashes with the proposed 10th October meeting should be notified to John. Ruth volunteered that Fermilab could host the March GDB

Formatted: Superscript

Formatted: Underline

For the Vancouver meeting Jamie Shiers said that four rooms have been booked for parallel sessions on the Sunday afternoon (i.e. one per experiment) On the agenda at the moment each experiment has a reservation. The date for cancelling the booking is fast approaching – if the room is not needed the please inform Jamie as soon as possible.

Formatted: Underline

For future meetings John noted that the Quattor working group would report again on its mandate at the next meeting. There will not be a GSSD meeting on 7th May but there is a two day workshop on 3rd and 4th July.

Formatted: Underline

There was a small word change to the VOMS mandate – refer to the mail from Maria just after the April GDB. (Jeff Templon thought there were many things not covered by the current mandate and said he could elaborate if required). The job priorities working group summary is fine but there is not much overlap between this and the VOMs activity.

On Future topics Jeremy Coles thought it would be useful to examine once again major (security) incident response processes and procedures. Ian Bird thought that Romain Wartel would be the best person to invite to the meeting to tackle this topic. Dave Kelsey said this area was last presented a year ago and agreed that the GDB should look at it and that more detailed procedures need to be put in place.

Formatted: Normal

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 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 WLCG Planning (Harry Renshall Lawrence Field)

Formatted: Bullets and Numbering

Harry pointed out that the presentation was more of a “heads up” than anything else.



On slide 1 he hoped ATLAS could clarify and said that “at the last GDB I wrongly reported a request for 70% of 2008 resources to be available for Q4 of 2007 when it was in fact needed for Q1 2008.”

Nick Brook reported a resource problem for LHCb and noted that stripping and reconstruction tasks scheduled for a few months ago do not have the ability to run smoothly as accessing storage is still a problem. Harry wanted to check if this was a T0 or T1 worry. Nick said these were T1 problems.

Slide 3: Harry remarked that he will put in changes as they happen. Resources are for the first day of a quarter but the numbers are somewhat dynamic. The table presents a snapshot of what is reported at any given time. John added that this was a careful dance. Harry wanted to think again about how to incorporate changes.

It was asked how the figures related to those in Chris Eck’s table shown at the RRB. Harry said they are synchronised and should agree but Gilbert noted that some of the table was not up-to-date. Les commented that Chris’s table is updated at specific times and not as regularly as Harry’s and asked if Gilbert was saying the RRB figures were not correct as they should be the same. Frederico asked if Harry could do some external checks and inform the GDB of the outcome and he agreed.

Formatted: Underline

3. SL4 Status (Markus & Oliver)

Formatted: Bullets and Numbering

John wanted to know if the SL4 release comments related to version 3.1. Ian said that version 3.1 for the WMS is very different from 3.0 and that there is an SL4 version coming.

Flavia Donno asked how many jobs were sent concurrently for the tests mentioned. Simone thought it was of the order 200. Ricardo wanted to know how long these jobs were in the running mode. Apparently most jobs were quick “hello world” types except for those used for testing the proxy.

On the gLite CE criteria...

It was noted that RHEL3 will be patched for longer than SL3 and this may offer alternative routes/timelines for the CE changeover (LCG to gLite). Dave Kelsey thought that from the CERN SL report at HEPiX it sounded like they had taken on board that SL3 patched are needed for a longer period. Ian was not aware of any huge objections from



people so far (to proposals to support gLite in a virtual environment?). Jeff thought CentOS would be ok too but Michel Jouvin believed they would move at the same time. Security patches for RHEL3 will continue until 2010.

Markus mentioned that it had been found that rpms favour (i.e. use first) the 64-bit libraries. This is good because not all code is ported to 64-bit but only one path is available for executables. It could have been that there were install order dependencies so this finding helps – it should be possible to derive tarballs from rpm installations without problems.

Jeremy asked about DPM on SL4. Markus said the timeline depended on Acros? Who would be away shortly so it could be 1-2 weeks or more. Oliver added that gridftp problems had been identified which required globus to respond.

Markus explained that a new PPS release would be available next week whichh contains fixes to all problems identified so far. However, there is still a need to verify the complete chain – install, configure. ...Jeff asked about one of the acromyms used:SEF. Markus gave the meaning: Someone Else's Problem. In answer to another question he answered that they should not start shipping interpreters (like Perl) with middleware

Working on the SL3 build in ETICS is less of a priority. Though it was the gLite 3.0 branch with mostly the same code there was a different build system.3.1 on SL3 is not a priority and the current mixed SL3 build system is not something to keep around. John wanted to confirm that the move to SL4 was the main focus. Ian agreed adding that in the past people upgraded fairly rapidly, but now this coupled with the move to SL4 things would be slower so the experiments need to push things along.

Markus' final comment was that we need to decide if the SL4 64-bit work is more or less important than maintaining WNs on SL3. Although one can move node by node, the bad news is that with the WN and UI you touch all the client libraries. It is only the service parts that are not touched. Most components have client, server and common libraries.

4. Security Policy Update (Dave Kelsey)

This was a report on the Grid Site Operations Policy; Grid Security Policy; Grid Policy on Handling Logged Personal Information.

Grid Site Operations Policy – <https://edms.cern.ch/document/726129> (draft v1.4 19 Apr 2007). Final version.

Federico: What does it mean “create any intellectual property rights...”

Dave: If you have no rights before then you have none after. It is a standard super computing clause - Running code on your site does not give you any intellectual property rights (IPR).

Formatted: Bullets and Numbering



LCG Grid Deployment Board Meeting



Dave sent out the final version of the grid site operations policy 2 weeks ago but had not had any comments and assumed this meant everyone was happy!

John: Is it relevant for WLCG to have a policy document. Is WLCG a grid under your definition to cover sites in WLCG that are not part of EGEE, OSG etc.... ?

Dave: I am not entirely sure.

John: Apart from the covering paper [that goes with the policy document] it could something with exclusion clauses etc. indicating who should answer.

Dave: MB approval is useful; at this point we are only asking for approval of the common part.

Ruth Pordes: The signatory for OSG is agreed. OSG security team have agreed. It will be discussed at the OSG council next Tuesday. At the moment we would not ask legacy sites to sign it unless they are WLCG T1 and T2 centres. How do you see the need for re-signing happening and what are the constraints?

Dave: At some level that is an internal operational issue

Ian: The document was originally an EGEE deliverable/milestone. We will start asking sites to sign this and certainly new sites.

Dave: personally I think this is useful for all sites. May be other interoperable sites too

Ruth: But you are not blacklisting sites who do not re-sign?

Dave: So the GDB is happy this goes forward to the MB for formal approval.

Formatted: Underline

Formatted: Underline

Grid Security Policy:

<https://edms.cern.ch/document/428008/4> (v5.6 - 24th April 2007). Aim to approve in June.

Ruth: Have you got OSG comments yet? Dave had not so Ruth agreed to forward them.

Formatted: Underline

Logged Information policy – <https://edms.cern.ch/document/840299> (not yet reviewed by JSPG).

Ian: This is different from the draft accounting policy?

Dave: This is the accounting policy but more general as there are different types of “monitoring” and accounting data – an appendix deals with accounting with user information

Ian: If we start to deploy accounting which has user data in it, is the appendix in hand so that we can have it approved. [Return to question]

Formatted: Underline

Dave said he would welcome comments on the general approach

Ruth: In OSG there is less talk about ownership of data and more talk about auditing policies being in place. How does this happen in JSPG and does this hold with this document.

Dave: Not sure. Not sure how it would feed in here.

Ruth: If you do not have a process of audit then the thing is not complete.



LCG Grid Deployment Board Meeting



Dave: At the top-level we say that sites audit themselves. EGEE3 talks of doing this more formally. People are expected to abide by documents.

Formatted: Underline

Ruth: We are struggling to find effort to address ... transitional services and.... Going to have to come back and inform higher level documents.

Dave: The top-level documents do not mention the VO policies that we're working on as part of an iterative process.

John: As well as policy is there something for VO managers to sign?

Dave: They sign appendix 1.... We need confirmation that these people are appointed as the VO managers. That's the next policy document being worked on.

Summary: The Grid Site Operations Policy (v1.4) is approved by the GDB. Comments are invited on the Grid Site Policy and Logged (v5.6) Information Policy (v0.3)

Questions were received at various points during the talk... John: My impression was that most regions in EGEE have a top level BDH. 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 BDH 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 toJohn: 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 BDII? Steve: No. But some countries under CERN, like Canada can have a large number of sites.

LUNCH 12:00

VOMS

3-5. Job prioritiesHEPiX DSWG Progress report (Jeff TemplonMichel Jouvin)

Ruth: Fermilab is consistent in not agreeing that this is an appropriate WG under IHEPCCC. The scope of mandate has not been agreed. What are people doing to address that?

...

John: How does IHEPCCC setup groups.

Ruth: It is not clear how WLCG relates to this WG.

John: This is just feedback – a way to improve information flow.

Michel: Many T1s are interested. The WG is for people who want to share experiences. Maybe some recommendation will come out of it but no enforcement.

Ruth: There is plenty of published information that can be used.

Are you interested in experiences with dCache at sites?

John: That was the plan.

LUNCH 12:00

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

Formatted: Font: 16 pt, Not Bold

Formatted: French (France)

Formatted: Bullets and Numbering



~~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.~~

~~4.6. Access control for storage Castor status & plans (Maarten Litmaath Tony Cass)~~

Formatted: Bullets and Numbering

~~Maarten's Kors: The situation as reported is not the full truth – after adding memory to the database server it eased the situation but there was also an issue with slots in the queue limited to 1000.~~

~~Tony: Yes – I was trying to get through the overview without going into all details. For the queue length we plan to extend this from today. It helps though if the queue does not build up in the first place. Generally we talk about the LSF plug-in but will tackle things from both directions.~~

~~John: Was 1000 a hard limit?~~

~~Tony: The limit was set because we knew we could not handle more.~~

~~Matthias: I am not so happy calling these T0 issues (slide 5) so I would call these CERN issues.~~

~~Tony: The point is the issue is understood and being dealt with. The hard limit of 1000 was not known during the design but was found in 2005/2006 when the system went live~~



and was understood at the time of the review in June. The limit was put in to make sure the system did not collapse.

Les: I talked to Richard Mount the week before last and he said he is willing to provide some support but with limited manpower – particularly for experiments for which SLAC is funded, for others more discussion is needed. Federico, it may be worth talking to him along with USALICE.

Formatted: Underline

Matthias: Does this mean that they support only USCMS and USATLAS?

Les: ... then made statement that CERN is funded to ... other funded experiments may be possible....They only have a small number of people so if it is to be used widely then we have to think about how the area can be developed. Currently only discussing this for ALICE. Discussion was specifically about CASTOR but touched on what if this becomes more widely used.

Kors: This is a fundamental issue – it is one thing to have 30000 requests in a queue. What if a poor user comes along as 30001; their job will never be dealt with. If 80% of work is for transfers and 20% users then there is an incorrect comparison – who says what is more important Higgs vs some other working group. Comparing a T0 issue (trasfers) with a CERN issue (users at CERN wanting to do analysis) and trying to decide who is going first.

Matthias: To add to that, the CMS issue in this area was not just for T0 but general priorities

Tony: Priorities can be enabled. To address Kors' point (that the number of requests is pending requests not total number of requests in queue), you can't have more pending requests than calls. The number of simultaneous requests is limited by the number of calls at the site.

John:SRM does give the ability to do prepare-to-get. One job can process a whole dataset.

?: Not all prepare to get commands/jobs Les: This may not be workable.

John: What is the typical number of files in a dataset? Kors: Thousands. Can this be done with separate pools or servers?

Tony: The queue is per stager. There is no belief that it will approach anywhere near the level possible at the moment. CMS open many files.... Queues here are the pending requests in the cycle.

John: What about quotas? This is not a DPM issue but certainly a tape issue. Does the glueschema handle this – i.e. disk and tape components?

Flavia: As far as I know glue 1.3 does not but 2.0 will

John: So this is limiting the amount of cache disk, tapes disk offline etc.

Formatted: Bullets and Numbering

7. Experiment top 5 issues (Jamie Shiers)

Formatted: Normal, Left, No bullets or numbering

Integration & testing of data and storage management components

Formatted: Font: Bold



LCG Grid Deployment Board Meeting



John: What is the difference between heartbeat transfers?

Jamie: Network problems have been intermittent and hard to debug. For example you start to get packet loss etc. ... should checking this be what the OPN group does. In Matthias' slide he want to see if queues are full, stalled etc. Should agree what it is we want.

John: Several monitoring frameworks are in place. In the UK we use one from EDG and this is intended to monitor such things as packet loss... is it that we want this for network T0-T1 tests?

Kors: The networks are reliable. It is end-to-end that often does not work and at the ends are the SRMs.

John: But you run background transfer tests?

?: There are 4 to 5 layers to be tested from the top layer. Network testing is just testing the bottom layer.

Kors: This is how we learnt a lot – iperf.

John: I'm talking about site to site.

Jamie: There have been network failures. There was one case this year, one last year.... We had three weeks of poor transfers and high failure rates and it took a lot of effort to debug the problem.

Tony: There is an ongoing problem with packet loss on the CERN network!

Jamie: Was there agreement on the basic goals?

John: We were asking for such things in the T1 milestones.... Is this because it has not come out of the experiments group or that the T1s have not signed up?

Jamie: They were made into experiment milestones, not T1 milestones. ... this needs to be centrally coordinated somehow.

Kors: During March we had to admit that we could not do tests for the multi-VO rates.

John: There were some results but none at the target rates

Jamie: This is tight for the FDRs. In July the experiments will have target rates to meet. The tests are more than just T0 export.

John: CMS demonstrated. Don't see ATLAS stressing any part.

Kors: No

Jamie: (slide 10)

John: Where do we go with this multi-VO test – whose milestone is it?

Les: CERN and T1s.

Matthias: It is to be driven by the experiments and they must agree if they will do it.

John: The experiments have agreed then. Have the T1s agreed?

Les: The schedule of tests is to become more and more continuous – to get each working up to the individual required rates and then run simultaneously. The secondary thing is which T1s are involved. Those that need to be supporting the experiments need to be ready when the experiments are able to run.

John: It is the ramp up part. When do you need a stable service vs installing new kit. The real issue is when ATLAS can drive things at a higher rate.

Les: Jamie should also identify times when things will overlap

Formatted: Underline



Kors: As a rough guess – for all the month of May testing of the stager so the March tests for T0-T1 transfers will be done in June. This is not shutting down after June/July. We do not need to specifically schedule the overlap. No FDR tests will stop this.

SRM v2.2 Status

Formatted: Font: Bold

Slide 3

Jeff: This thing about namespace is it the same for DPM? (the connection between path and backend)

John: You can have multiple instances. For files with similar names but different tokens they should be in different paths.

Jeff: As an example – if ATLAS had corrupt files and the path suggested they are all on the same disk server then this is not the case. The experiment may think the paths mean something – characteristics of backend.

John: It should be the token that does that

Jeff: About the VOMS issue – don't spend too much time developing something – get it out to testing. We found problems in job priorities group with different VOMS views.

John: How do the experiments feel about this feedback. Are they ready to start testing?

Michel: At GRIF we are part of the testbed (DPM) but it is not clear how to be part of PPS. Can system be part of both infrastructures?

Flavia: We are waiting for input from experiments on how they intend to do the tests and how we should configure the PPS to support this.

Formatted: Underline

John: The experiments probably want to test on real data but sites are reserved about installing "non-certified" software.

Flavia: You can copy data into the PPS and then do some SRMv2 tests. The experiment needs to declare they will use the production catalogue so files generated on the PPS will be registered in the production catalogue.

Michel: It is not easier to say they are running these tests on the production service.

Flavia: In PPS only 2 sites have storage in production.

Flavia: DPM at the moment does not have a space garbage collector. We were filling up DPM production instances and had to stop due to this limitation.

Job Management (Claudio Grandi)

Formatted: Font: Bold

John: Could you say more about the LCAS/LCMAPS package – is it available as a node or part of distribution?

Claudio: It is available as a library and can be used by processes on that node. To serve more than one node you need to build an interface on the network. glexec needs to be deployed on all WNs and the libraries are linked in glexec.

John: Is there any issue in populating that database?



Claudio: The service may use the same file for configuration as used now for the library. It is a service which needs to be developed.

John: The distributed bit of it?

Jeff: Can I try to answer. LCMAPS is simple – present it with the DN and FQAN and get back a UID and GID. There are rules on how to do the transform but this is simple. There is only one thing you need to do and if it is on one machine you call LCMAPS on it, so you have to expose an interface – contact on port whatever to get the UID and GID. You just need to be able to expose a port and interface

John: That is not my point. Where is it deployed?

Claudio: It is deployed on all EGEE sites....

John: What rules do you put in it then?

Claudio: These are put in at configuration time. You are concerned about dynamic updates of that information? Yes, then for that you could use the GP-Box plugin but this is far from production.

Kors: Third point (slide 5). If a site does not support glxexec then does this mean no glide in jobs?

Claudio: Nobody can prevent a VO from running what they want within a job – the payload could have multiple users in the same job, but this is not wise for security reasons

Kors: Have all the TIs agreed to run glxexec?

John: Have they been asked?

Claudio: From a security point of view it is okay if the payload runs jobs by the same user.

Jeff: Sites that looked into this realised it impacts their ability to answer on legal responsibilities.

DCache (Flavia)

Formatted: Font: Bold

Matthias: With regard to the CMS request for ACLs and consistent behaviour. If this can only come in 2008 then I will take the issue back to CMS.

FTS, DM, IS (Ian Bird)

Formatted: Font: Bold

Kors: The error reporting is now much better in FTS and this is much appreciated

Federico: We would prefer xroot in DPM in the standard distribution

Ian: The issue is if you have to restructure the existing plug-in

Information system:

Maarten: Sites have reported speed improvements when more indices have been added – there is conflicting evidence but this would enable us to run for longer

Stefano: Is this work by SA1, LCG or

Ian: It is work done in the CERN group



LCG Grid Deployment Board Meeting



Stefano: Are there enough people for the work to be done?

Ian: We need to come up with a plan for this.... A lot of things are blamed on the Information system even if they are not.

Flavia: One further point on the LHCb issue. Root and SRM are not well integrated in the sense that a TURL returned from the SRM may not be useful with dcap.....solution available soon.

Discussion

Formatted: Font: Bold

John: Of the issues raised by the experiments are any at odds with the Tier-1s? For RAL DM and CASTOR are issues. And then what is the immediate feedback of the experiments. Are they happy with everything they heard today and that the issues being addressed?

Federico: You mean happy that things raised are answered?

Matthias: Not all of the Top-5 are of the same urgency. Things will be clearer in 4-6 weeks. Not everything is expected to be resolved quickly.

Ricardo: Most of the issues are being addressed. One point though is the stability of services in general is poor and they can be brought down by a single user. The current system depends on users behaving. The more chaotic future use is likely to lead to less stability, not under high-load but with non-standard requests.

Dario: I'm happy some of the issues are being addressed but still worried. I'm unhappy that quotas are postponed until 2008 (quotas per site per VOMS group).

John: Is this not also a management issue? Even if we implement it, how do you manage it?

John: So we have four happy experiments where next?

Les: The important thing to point out is where the experiments are not happy with the response. Today was largely about status. Quotas and access rights have been talked about for a long time and are not high priority for the developers, so if this is a higher priority then we need the experiments to raise it at the expense of something else. If there are concerns about things missing then experiments should make this known. The MB needs to make sure that things are monitored and progressing.

John: We need a monitoring plan.

Formatted: Underline

Matthias: The CMS issues were understood. I understand how they are prioritised and that is probably appropriate.

John: Perhaps it is not obvious when things are implemented and really solved. The management plan – is that for the MB to extract from these technical responses?

Formatted: Underline

Les: Yes

Formatted: Underline



LCG Grid Deployment Board Meeting

GDB

Jeff: I am a little surprised that accounting, groups and roles has gone so low on the priority scale. ATLAS problems are about storage usage. The thing that Ricardo talks about is going to hurt bad with lots more people using the system

John: Transferring the data from T1s is rightly a higher priority. These other things come afterwards. The basics needed are to get data out to the T1s – rates and reliability are required.

Jeff: Yes

John: Let's worry first about the FDRs now. We then need a plan for 2008

John: The GDB actions need to be reviewed checked between the wiki and minutes and those outstanding discussed with the owners. An update will be given next time. Many things from today need to be reviewed again.

Formatted: Underline

John: Are there any other ideas for presentation topics?

Ian: This glxec business. Did we agree that we would ask sites if they will deploy glxec as is now?

John: We will revisit this before the next meeting

Jeff: The conclusion will be based on a mixed population – some will not accept glxec in gLite and some will only deploy pilot jobs with glxec. So, if this is an issue for ATLAS we should know.

John: The [glxec] surveys need to provide an unbiased response. When the consensus is known then others may change their response. Did this originate from the TCG? Who is the action on?

Jeff: This is perhaps more of a hot topic than SL4!

Action 0705-1 – John to get feedback from Markus and Alessandra on previous feedback from sites on glxec.

Formatted: Underline

Formatted: Underline

Kors: This is just as much an issue for LHCb as ATLAS. Ask the GDB list – at least the T1 sites should respond

Maarten: There is an issue with glxec – batch systems may not be able to handle job trees. I have not seen anyone discuss this to completion. There is still technical work to be done to complete this.

Maarten: It is fine if jobs behave – what if the batch system needs to clean them up. Once this is UID of process then anything below this will be impacted

John: Also use process trees...

Maarten: This is where we started thinking but there was no conclusion on the discussion. There may be still a bit of work before we can go ahead with the idea.

Ian: This is not changing the user ID but deployment and impact on user auditing. Use-cases are all solved with correct logging – there is no need to do set UID on WN...

John: We can build in auditing

Ian: This has to be done by the system and not the user.



LCG Grid Deployment Board Meeting



~~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 T1s 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~~



LCG Grid Deployment Board Meeting



~~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. Tacking 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)~~

Formatted: Bullets and Numbering

~~There was a brief discussion about using the Tape1Disk0 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?~~

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

Formatted: Bullets and Numbering

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

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



LCG Grid Deployment Board Meeting



- ~~-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~~

Formatted: Bullets and Numbering

9.8. AOB

Formatted: Bullets and Numbering

There was no other business.

MEETING CLOSED AT ~~16:35~~16:50

Actions:

Item No.	Description	Owner	Status
0602-4	Phrase the requirement on how to use policies in the WLMS	Cal Loomis	Open
0603-3	Follow up to ensure all sites in country are publishing accounting data or contact John Gordon with issues preventing this happening	Country representatives	Open
0604-6	Drive forward discussions on the VOMS and protocol issues	Ian Bird	Open
0605-3	Provide feedback (with reasons) to Dave Kelsey or Kors Bos on whether the security policy presented by Dave is acceptable.	All	Open
0605-4	Tier-1s to report back to GDB on what proportion of their current WLCG work is not reported/accounted within WLCG	Tier-1 managers	Open
0606-7	Take up and discuss technical solutions for removing shared credentials from the VO boxes	Markus Schulz	Open

Formatted Table



LCG Grid Deployment Board Meeting



Item No.	Description	Owner	Status
0607-9	Ensure the default YAIM is properly configuring lcas lcms for the sgm accounts (and that it works!)	Jeff Templon	Open
0609-1	Follow up on NDGF security policy position	Les Robertson	Open
0609-2	Look up statistics for automated on-call system and send information to GDB	Bruce Gibbard	Open
0609-6	Send storage type sampling script to John Gordon.	Jeff Templon	Open
0609-7	Move accounting to work in decimal units	Tier-1s/sites	Open
0610-5	Provide more detail on who is supposed to sign the site policy for each "organisation" mentioned in the security policy document	Dave Kelsey	Open
0610-6	Send the site operational procedures policy to the list again for comment ahead of approval and ensure lawyers at sites have a chance to review the document	Dave Kelsey	Open
0701-3	Check the CPU and storage accounting figures being published for the site	Sites	Open
0702-3	Discuss the future of a VOMRS-VOMS task force and consider possible mandates for the group	Dave Kelsey, Maria Dimou et. al.	Open
0702-4	Check Harry' resource tables and understand what they mean	Tier-1 sites	Open
0703-1	Check the Victoria MB time with Les Robertson and agree intention at the MB	John Gordon	Open
0703-2	Follow up on accounting policy documents	John Gordon	Open
0703-3	Send out a link to the latest patch	Jeff Templon	Open
0703-4	Follow up on the VOMS coordination group mandate	John Gordon	Open
0703-5	Refer Cal Loomis to Marian Dimou concerning the representation of smaller VO requirements in TCG discussions	John Gordon	Open
0704-1	Update slide 17 of presentation and formulate a request for documentation to be provided by the middleware developers to explain options with components (needed by Quattor maintainers)	Michel Jouvin	Open
0704-2	Follow up on VOMS coordination group mandate wording with Maria Dimou	Ian Bird	Done
0705-1	Get feedback from Markus and Alessandra on previous feedback from sites on glexec.	John Gordon	Open
0705-2			

Formatted Table

List of Attendees

X means attended
V means attended via VRVS

Formatted: Centered

Country	Member	Deputy
---------	--------	--------



LCG Grid Deployment Board Meeting



Country	Member		Deputy	
Austria	Dietmar Kuhn	X	-	
Canada	M-Vetterli		R-Tafirout	X
Czech Republic	Milos Lokajicek		Jiri Kosina	
Denmark	John Renner-Hansen		Anders Waananen	
Finland	Klaus Lindberg		Jukka Klem	X
France	Fabio Hernandez		Dominique Boutigny	
Germany	Klaus Peter Mickel		Holger Marten	
			Jos van Wozel	
Hungary	Gyorgy Vesztegombi	X	Dezso Horvath	
India	P.S Dhekne		B. Vinod Kumar	
Israel	Lorne Levinson	V	-	
Italy	Mirco Mazzucato		Luciano Gaido	
Japan	Hiroshi Sakamoto		Tatsuo Kawamoto	
Netherlands	Jeff Templon	X	Ron Trompert	
Norway	Jacko Koster		Farid Ould-Saada	
Pakistan	Hafeez Hoorani		-	
Poland	Ryszard Gokioli	V	Jan Krolkowski	
Portugal	Gaspar Barreira		Jorge Gomes	
Russia	Alexander Kryukov		Vladimir Korenkov	
Spain	Manuel Delfino		Xavier Espinal	
Sweden	Niclas Andersson		Ford Ekelef	
Switzerland	Christoph Grab	X	Marie-Christine Sawley	
Taiwan	Simon Lin	A	Di Qing	X
United Kingdom	John Gordon		Jeremy Coles	
United States	Ruth Pordes	A	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	V	-	
Project Leader	Les Robertson		-	
GDB Chair	Kors Bos	X	-	
GDB Secretary	Jeremy Coles	X	-	
Grid Deployment Mgr	Ian Bird	X	Markus Schulz	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	
Fabric Manager	Bernd Panzer	-	
Application Manager	Pete Mato-Vila		
Security WG	David Kelsey	-	
Quattor WG	Charles Loomis		
Networking WG	David Foster	X	
Planning Officer	Alberto Aimar	X	

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



LCG Grid Deployment Board Meeting



	<u>Nick Brook</u>	<u>V</u>		
<u>Project Leader</u>	<u>Les Robertson</u>	<u>X</u>		
<u>GDB Chair</u>	<u>John Gordon</u>	<u>X</u>		
<u>GDB Secretary</u>	<u>Jeremy Coles</u>	<u>X</u>		
<u>Grid Deployment Mgr</u>	<u>Ian Bird</u>	<u>X</u>	<u>Markus Schulz</u>	<u>X</u>
<u>Fabric Manager</u>	<u>Bernd Panzer</u>			
<u>Application Manager</u>	<u>Pere Mato Vila</u>			
<u>Security WG</u>	<u>David Kelsey</u>	<u>X</u>		
<u>Quattor WG</u>	<u>Michel Jouvin</u>	<u>X</u>		
<u>Networking WG</u>	<u>David Foster</u>			
<u>Planning Officer</u>	<u>Alberto Aimar</u>	<u>X</u>		

Others present at CERN

Jamie Shiers

Harry Renshall

M. Lameme

Flavia Donno

Sue Foffano (CERN)

Simone Campana (CERN)

Nechaerskry Andrey (CERN)

Steve Traylen

Also present in the meeting room:

Steve Traylen (CERN)

Matthias Kasemann (CMS/CERN)

Michel Jouvin (France)

Oliver Keeble (CERN)

Jamie Shiers (CERN)

Stephen Gowdy (ATLAS/SLAC)

J Knobloch (CERN)

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

Formatted: Font: Bold, Not Italic

Formatted: Font: Bold, Not Italic

Formatted: Font: Not Bold

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

Formatted: Font: Not Bold

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

Formatted: Italian (Italy)

Formatted: Italian (Italy)

Formatted: Italian (Italy)



LCG Grid Deployment Board Meeting



~~Pete Gronbech—Oxford~~
~~Gabriel Stoeica~~
~~Lief Nixon—Linköping~~

~~Frederique Chollet~~
~~Greig Cowan~~
~~Stefano Belforte~~

~~Gonzalo Merino~~

~~Richard Gokieli~~

~~Alvaro Fernandez (IFIC)~~

~~Juergen Knobloch~~

~~Additionally on VRVS PM:~~

~~Pierre Girard—Lyon~~

~~Paul Gelissen—Bern~~

~~Jos Van Wezel—Karlsruhe~~

~~Les Robertson—CERN~~

~~Elizabeth Sexton Kennedy—Switzerland~~

~~Helene Cordier (Lyon)~~

~~Owen Synge (DESY)~~

Formatted: English (U.S.)

Formatted: Italian (Italy)

Formatted: English (U.S.)

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

Formatted: English (U.S.)