



Version 0.11.11-2

Amendments history

Name	Area	Date
<i>Jeremy Coles</i>	<i>All - full edit -> v1.0</i>	<i>14th June 2007</i> <i>12th March 2007</i>
<i>Jeremy Coles</i>	<i>Minor revisions -> v1.1</i>	<i>27th June 2007</i>
<i>John Gordon</i>	<i>Minor changes section 1</i>	<i>13th March 2007</i>
<i>Jeremy Coles</i>	<i>Summary and attendees</i>	<i>13th March 2007</i>

Formatted: Left
Formatted: Left

Minutes of the meeting

~~CERN~~CERN, ~~6th June~~^{7th March} 2007

Formatted: Superscript

Agenda: ~~<http://indico.cern.ch/conferenceDisplay.py?confId=8470>~~
~~<http://indico.cern.ch/conferenceDisplay.py?confId=8484>~~

Formatted: Default Paragraph Font

Field Code Changed

Formatted: Hyperlink, 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

Introduction (John Gordon)

The August 1st meeting has been cancelled. The October 3rd meeting has been moved to October 10th. Triumph have volunteered to organise transport from the August 31st GDB to Victoria where CHEP and the WLCG workshop will be held. A registration page will appear soon.

Tier2s

This was the first of several presentations describing how Tier2s are supported within countries. Andrew Elwell described the transfer tests being carried out between the RAL Tier1 and the various UK Tier2s.



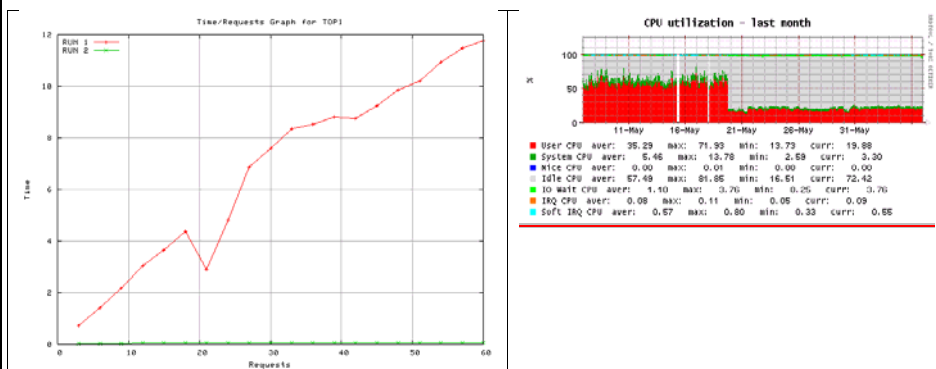
Middleware Issues

SL4 (Markus Schultz)

The gLite3.1 Worker Node built on SL4 is now in production. The corresponding UI has been certified and is currently under test on the PPS. MB milestone to rollout at T1s in 3 weeks. Sites need statements from experiments on certification.

BDII (Laurence Field)

Most of the short- and medium-term issues identified by Laurence at the March GDB have been addressed. Indexing on the database has produced the most dramatic improvements. This has removed the lcg-utils timeouts from SAM-BDII.



gLiteCE (Claudio Grandi)

The lcgCE is the only version in production. The gLiteCE built on SL3 has been performing well in tests handling 6000 jobs/day with 98% success and up to 4000 jobs in a CE at the same time. The version for production, built on SL4, should be ready for certification during the summer.

Claudio also outlined plans to retain a GRAM interface to the gLiteCE which has been requested by many countries/sites.

Job Priorities (Simone Campana)

This seems to have run into difficulties. The matching to job queues achieved by the WMS is not as originally planned. There are fixes implemented in YAIM but they were released into the wild too soon and were installed automatically at some T2 sites.

I think this work needs to be started again, perhaps with a fresh set of people. Action on MB



FTS2 (Gavin McCance)

This has been tested at CERN and RAL and should have been released in production at CERN yesterday. The plan is to roll it out to Tier1s in about 3 weeks.

VOMS

VOMRS Group (Maria Dimou)

Maria described how the group had started with a new mandate. There is a plan to replicate VOMS databases but the use-cases were not clear. At first sight this would appear to be identical to existing 3D uses but is VOMS designed for failover and/or can clients be configured to multiple servers?

Generic Attributes have been developed after a long gestation and are currently in certification. However it was not clear how they would be used and by which parts of the middleware. This led to a discussion on the use of VOMS FQAN/roles/groups/attributes could be kept consistent across the middleware. There is a need for a meeting of developers and other interested parties to thrash this out.

Access Control for Storage (Flavia Donno)

Flavia summarised from the pre-GDB. DPM has implemented Access Control and VOMS support and this is used as a model for other implementations. A Castor implementation is not likely before 2008.

Monitoring

A set of talks on various aspects of monitoring from site-tools, through SAM tests and dashboards to an overall framework for publishing. To be encouraged.

Security

More revisions of policy documents. The top Level Policy document is hoped to be presented to GDB for approval in July. The Logged Information Policy document is proving harder than anticipated so a specific one for accounting data will be pushed through.

glxexec caused long discussion. Sites are concerned about glxexec. It was felt that it is acting like a gatekeeper without as much security analysis having been done. In the lcg-modified gatekeeper it is the same lcas/lcmaps code which is being used. In general less concern about use on CE. On WNs some sites feel strongly that the identity of the job-owner must be known for auditing, others that they do not wish sudo-like code on WN. It was originally thought that LCMAPS logging would allow accounting but it is unlikely that it can be guaranteed that all processes get the change of identity.



LCG Grid Deployment Board Meeting



The issue comes back to Pilot Jobs – are they acceptable? Jeff Templon reminded us of his summary from last year that if the pilot job downloaded subsequent payload from the same user then everyone was happy. If the pilot job downloaded payload for a different user and didn't change identity then the VO was breaking the AUP. If it were to change identity then glxexec is required. Stalemate. Referred back to MB.

Planning

Plans are maturing for the workshop preceding CHEP although some speakers/chairs are still needed.

MultiVO transfer tests are scheduled for the last week of June.

Updated schedule for dress rehearsals is urgently required.

Next Meeting

July 4th at CERN when there will be presentations on accounting and a report from GSSD following a two day workshop at CERN on 2nd and 3rd.

Formatted: English (U.S.)

~~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.~~

~~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.~~



LCG Grid Deployment Board Meeting



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 slapd 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 T1s 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 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.

Formatted: Bullets and Numbering



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)

There are outstanding GDB actions to be addressed. These are mainly group actions. Please check the actions list. There will not be an

~~No~~ August 1st meeting. We have few ~~Few~~ volunteer sites for 2008 meetings but it ~~is~~ would like be good to have a 2008 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.

Formatted: Bullets and Numbering

Formatted: Normal

Formatted: Superscript

Formatted: Normal



LCG Grid Deployment Board Meeting



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 T1-T2 transfers (Andrew Elwell/Lawrence Field)

Formatted: Bullets and Numbering

There were several questions about failure cases which have not been captured by these minutes. Andrew did point out that the failures seen were not from sequential tests so the cause was unlikely to be the test method and could be (perhaps at the 80% level) due to one of the dCache servers at the RAL T2. The future tests will roll out transfers to other T1-* and intra T2 channels.

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

3. VOMRS (Maria Dimou)

Formatted: Bullets and Numbering

Jeff Templon: Adding attributes is not a good thing to do now.
Ian Bird: Gaps should not take on what Jeff is talking about. A small focussed discussion is needed
Maria: At the operations meeting many people are asking for information

GAS:
Jeff: How does this fit with SA3?



LCG Grid Deployment Board Meeting

GDB

Maria: The ~~F~~tester is in SA3. Jeff

J: Is this ~~in~~ parallel to your certification Markus?

Markus: VOMRS is in SA1 not SA3 but it is clear ~~who does what~~.

Ian: Before getting ~~an~~ Oracle licence ~~we~~ need to understand this “need” to replicate. Who uses it?

Maria: The code is there ~~already and the-~~ ~~small~~-machine costs are small.

Ian: But effort costs ~~too~~. Can jobs switch ~~the~~ VOMS server ~~they use~~? I ~~W~~would like feedback from lcas/lcmaps ~~experts, etc.~~ ~~U~~ the use case is not clear. Answer this before more effort is put in.

Jeff: Ask Dave...

Dave: There is ~~C~~clearly a single point of failure with the server. The Oracle database replication is a way forward. Security issues of doing this have ~~not yet been~~ considered.

Jeff: Is this a ~~M~~mirror or replication? What about ~~L~~latency issues?

Michel Jouvin: I am ~~S~~surprised this is in VOMRS.

J: VOMS--

Maria: This is all ~~V~~valid input. Action 0706-1 (Maria Dimou) ~~The working group needs to~~ ~~WG~~ goes to the VOs/developers to ask about the need?. ~~The N~~next developer meeting is in 1 week in CNAF. Note that ATLAS have asked ~~three~~3 times about this functionality.

Ian: But what is the ~~u~~Use case?

Dario: We need to ~~C~~copy the database to BNL. That instance will be ~~I~~ndependent/independent. Ian

J: Dealing with ~~F~~failure is different from failover.

Claudio: It ~~M~~may be easier than it looks. Vomses file can contain >1 input - ~~C~~hecking/Checking it works needs to be done.

Jeff: Trying to get sites to manually change configuration files is difficult to progress. To put this into auto update takes time.

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

Jeff: The experiments/~~Expts~~ have asked for stability. You are ~~N~~now talking about VOMS admin v2 being implemented in mid-summer.

Maria: I'm ~~S~~saying this to give an idea when it could be ready.

Jeff: Is it ~~P~~part of a proper certified release?

Maria: ~~Y~~es – testers began on it in early March.

Claudio Grandi: I've just ~~C~~hecked with the developers. ~~—f~~For match-making, ~~-intro~~ generic attributes ~~are~~is not checked in FOAN. In the ~~N~~new version of lcas/lcmaps

Formatted: Underline



everything is fine – there is no impact from generic attributes. The developer does not foresee problems here, but version in leg...

Jean-Philippe Baud: Data management depends on this – change it and queries will fail.
Jeff: No similar FQAN library is used by all products. I'm not opposed to adding it once it is checked.

4. VOMS and ACLs in Storage Services: Summary of the GSSD discussions (Flavia Donno)

Formatted: Bullets and Numbering

Jeff: The comment on “VOMS integration” in the second bullet. Although I requested this but I am uneasy about it. Suppose the site supports just LHCb and we have a person in ATLAS and LHCb. Then, when saving a file it checks and approves on the second item and so saves as LHCb.

Flavia: It uses subgroups so this is not a problem. s...

Conclusion: Action 0706-2 (Flavia) Come up with a description of the implementation and submit this to the experiments to see if this can satisfy their requirements.

Formatted: Font color: Indigo

Formatted: Underline, Font color: Indigo

Formatted: Underline

Is it okay to have ACLs at the namespace and pool level?

John: Bigger pool to have quotas

Flavia: This 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 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 BDII. 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 BDII. John: Change clients to talk to site BDII's? 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 BDII's 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 BDII. 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 priorities SL4 Status June (Jeff Templon Markus Schulz)

The WN code will be released today to production. A Tarball release is also ready. 64-bit is the same as a month ago - 70% success. The UI is ready for the PPS.

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

Dario: The Linux certification group is now starting on SL5. Will you be part of it?
Markus: No. Resources are already split. SL4 working and 64-bit to production services.
You chose what we drop but we can not add another branch with current effort.
Jeff: Having done SL4 work did you not say that SL5 would be easier?
Markus: There have been structural improvements. The Sstack should be easier to move.

Dario: The Experiments are being asked to participate in this activity
Tony: The fact that SL5 is being certified is not a criticism on Grid Deployment (GD). For some experiments we are being asked to do this so they have a platform for testing. There is a split between GD and FO (Fabric Operation). We

- Formatted: Font: 16 pt
Formatted: French (France)
Formatted: Font: Bold
Formatted: French (France)
Formatted: Bullets and Numbering
Formatted: French (France)



Will not provide services on SL5 until after the 2008 run. We would be ~~D~~reaming to think Markus should be on this now.

Markus: It is ~~B~~est to be on SL4 first.

Ian: We ~~W~~ill set up a single build machine for SL5 but as there is little effort in this area this build is— just to get an idea of problems that might impact in the future.

Stefano: The ~~T~~arball is available now. Is this a YAIM version?

Markus: It is also in release — the rpm version.

Holger: We have new machines and a ~~M~~ mixed ~~enviornement~~environment of SL3 and SL4. ~~n~~Currently~~ow~~ we have the SL3 build on SL4. Is there any functionality difference between this and the with-natively compiled versions?

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

The Information System (Lawrence Field)

Formatted: Font: Bold

This was an ~~U~~update on ~~2mths ago~~the presentation given two months ago.

Jeff: Is there an ~~I~~increase in latency?

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

John: On ~~P~~the patches — where do the others [not mentioned] stand?

Olivier Keeble: The ~~I~~indexed BDII is in the next release. Others patches are at various stages of certification.

Acceptance criteria for the CE (Claudio Grandi)

Formatted: Font: Bold

Formatted: Font: Bold

John: Does the gLite GT4 reliance have—~~Had~~ wider implications?

Markus/Ian: ~~N~~o

Claudio: The ~~I~~layer in front of the gLiteCE will not in principle be needed once we use BLAH. One scenario is to keep it.

Markus: The ~~N~~new UI and WN are all based on GT4. The ~~M~~major 3.1 release.

Claudio: The ~~N~~new WN release is going into production; the UI is on the PPS. These are Based on ~~pre-webservices~~web services GT4.

Jeff: A ~~S~~small point — the blue thing [in the diagram colouring scheme/animation] should appear before the LCG part is removed on —(slide 5!)

Formatted: Bullets and Numbering

Formatted: English (U.S.)

Formatted: English (U.S.)

6. Job priorities —(Simone Campana)

Regards~~Y~~ looking at every views published by site - the LCG RB is not concerned about this, while the—Use gLiteWMS does look at the information—~~look info in block~~. If you



LCG Grid Deployment Board Meeting



submit a job as software manager then there is no view for this -- pProduction matches 2 and 3. (see slide 3)

John: For the generic user does this not- go into the short queue but-go in with a different GID?g. Simone: The

Queues do not matter. The Scheduling is different. Different users have different priorities in the same queue for different groups.

John: Why?

Jeff: It is Nnot even clear why the mapping was wrong. NIKHEF had the problem for 2 weeks on the production site. There is Nno easy way to notice the problem. In the NIKEF case, cCapital letters in the config file did not match what -while-YAIM looked for (which was all capitals). The site was -Bbroken without us knowing.

Simone: Consequences. One there is -not matching. Two, the site2- attracts jobs which pile up and remain until they time out. Some work is required to keep looking out for these gatherings.

Jeff: There is a SAM test for this - submit job with and with/out a VOMS role and look at outcomes.

John: Implementation--eCan see shares between VOs implemented. There is Nno reason to think it won't work. SGE for example plug--allows us to split shares.

John: It is Eeasy to say keep this for T1s. But what use-cases are there 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 the mechanism can't work. --150+ sites will not keep up with the requirements. It is Nnot obvious we are doing the right thing. ExlusivityExclusivity is not what the batch system sees - ATLAS should be the sum of everything.

John: Who should revisit the design?

Ian: Someone fresh with an understanding of the batch systems. We Nneed to consider how to do this - asking YAIM to do the configuration is not going to work.

Jeff: We Ggot it [the current implementation] out fast because we knew there would be problems. Some people then took it when it was not ready! t--

Ian: This is a reason not to do this on the production system. The PPS should do it first.

John: What about Sstress testing to check shares etc.

Ian: We Nneed just one test.

Erwin: We Ccan do a mock up. This nNeeds to be part of wider to hit the Hit production sites with this can be damaging.

John: What about the Uuse cases?

Kors: This was -Bbrought on to the production system, then -Wwhat went wrong?:-

Claudio: It Wwas a short term solution.



Jeff: The WLMS did not work out of the box. The solution was exposed to users Exposed to—and fix-bugs needed to be fixed... and by this time many T1s had deployed it.

Markus: The Ppoint is (including cooperation with the experimets) we should do this on the on-PPS which is current under utilised anyway. S. Under-utilised. Expensive

John: For the LHCb case yesterday the problem was– not enough data being available on PPS.

Claudio: For some tests that matters but not job priorities.

Ian:

John: The Ssummary from Ian is that someone needs to look at this again.

Claudio: The JPWG (Job Priorities Working Group) can address its membership

Markus: Torque/LSF are done. The next solution is for Condor.

Action 0706-34: (Erwin)— Llook at membership and approach of JPWG

Simone: What is the Sshort term fix here?? Ask T1s to fix and T2s not to deploy?;

Fabio: Do we Wwait for a reeeiperceipe before deploying?

FTS 2.0 deployment and testing status (Gavin mcCance)

John: Do the deployment dates clash with FDRs?

Stefano: Even if they do, Wwe have to live with it.

John: Is that the position of the other experiments?

Yes

Fabio: SitesIs— the message-is that sites should wait till 3 weeks after 18th June-or then?

Gavin: The Nnew client is backward compatible – upgrade as soon as possible.

Fabio: This may be a problem given the coming summer period

7. CASTOR task force update -(Bernd Panzer-Steindel)

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

Bernd: Our focus is as given in slide 10.

Flavia: SebasatianSebastian also made it clear that ACLs [for CASTOR] are not likely till later in 2008.

--- Lunch ---

Formatted: Underline, Font color: Blue

Formatted: Font color: Blue

Formatted: Font color: Blue

Formatted: Underline, Font color: Blue

Formatted: Underline, Font color: Blue

Formatted: Underline, 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.)

Formatted: Font: Bold

Formatted: English (U.S.)



8. Monitoring (Markus Schulz)

Grid Services Monitoring (James Casey)

On the prototype Nagios monitoring framework proposal:

Jeff: A site runs this and the information remains at the site?

James: Yes. I 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) I thought you were only sending high-level information off site? -

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

Jeff: When did someone say this would be done?

Ian: Two years ago. All the GridView plots relied on this. No DNs are shipped.

Jeff: So long as site can switch on

James: No private data is included and the only personal data is the DN.

Ian Neilson: 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: We believe the probe set adds value to any site. An interface system is provided but only we do not tell sites to run it.

Jeff: We have our things that do some of this and s--- don't want to install something that is run remotely. So can we turn it off:?

Ian: If you are asked Will be asked if a given service is running can you answer?

Jeff: There are So many things monitoring us.

Ian: We started collecting sensors...these are available to everyone. Sites can run there own checks. We have to ship some information about the site to a repository for an analysis to determine if the site is available....

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

James: I Agree - wikis are not the best place to keep documents. Worked through those problems so more aligned with Erwin...ete.

Ian N: The 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.

Formatted: Bullets and Numbering

Formatted: English (U.S.)

Formatted: Font: Bold



James: That is fair – we are not yet at a level to give information back. The probes are written but we have no infrastructure to run these at sites at the moment – at the moment. Will pick up on this at the operations workshop.

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

James: It is 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 will then be used to give an availability number. There are many tests – this will not affect availability.

Oxana: NDGF is not in the loop so far. The implementation is not the same. Tests may not be applicable. We do not run an RB for example. Would it always be "failing" or given some other status???

James: Action 0706-4 (Oxana) If you could nominate someone we'll integrate them soon.

[Chat comments about quality of video conferencing]

System ... Dashboards and Monitoring (Julia Andreeva)

Jeremy: In the site efficiency reports – which jobs are used as a basis for the calculations?

Julia: These are only jobs going via an RB

Jeremy: Many of the failure reasons are "unknown". Can this be more informative?

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

The monthly site summary uses just specific applications. For ATLAS and CMS this is production. LHCb and ALICE uses pilot jobs.

Systems Management Working Group (Alessandra Forti)

John: No questions

Latest on SAM (Piotr Nyczuk)

HONE?

Holger: Lots of MC and analysis

Jeff: The VO specific tests – how are they run?

Piotr: ALICE run their own test suite – these run as a cron job. Gssi cron job. LHCb is also using a cron job but via the SAM submission framework.

Formatted: Font color: Indigo

Formatted: Font color: Indigo

Formatted: Font: Bold

Formatted: Font: Bold

Formatted: Font: Bold

Formatted: Font: Not Bold

Formatted: Font: Bold

Formatted: Font: Bold

Formatted: Font: Bold

Formatted: Font: Bold



LCG Grid Deployment Board Meeting



Jeff: We have seen ~~Seen~~ jobs forking on the CE and then kill them. Do you have an account call-~~ed~~SAM ALICE?

Piotr: Yes. ~~Falls back-~~

Jeff: They produce ~~h~~irritating log messages.

SAM tests currently ~~runs~~ every hour. ~~r---~~

~~Access to database — write only?~~

Jeremy: Even today we have seen a new SAM test introduced which produces error messages for sites to follow up before the fix (in this case latest CA update) is available to download. Can the process make the update possible before the test changes and sites start investigating? Piotr:

Change of tests— should happen after broadcast of changes but there are many discussions ongoing about this. †

ALICE tests:

John: When do we get back the interface? Piotr: ~~One+~~ month or so to finish ~~standardisation~~standardisation. It was removed due to ~~M~~many expensive queries slowing things down.

Fabio: Should we as a site ~~we~~ be concerned with all the current results ~~of those tests~~? For ALICE ~~our~~ site is very red. Piotr:

Only the ops VO information is published at the moment. ~~Defin~~Current ~~it~~iondefinitions of which tests are critical do not make much sense.

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

~~P-~~VO tests is— something to be followed up between the ~~more~~VO and sites. Fabio:

~~F-~~Who do we negotiate with?

Nick Brook: ~~Action_0706-5~~ Someone centrally should decide this —~~not~~ 1-1 communication with sites is not possible. If a VO says something should be taken seriously then ...~~and then~~

Jeff: I carried out a SAM review recently. A ~~C~~ouple of things questioned had to do with the interface. ~~Critical tests regardless of test and all~~ Piotr:

Formatted: Font color: Indigo

Formatted: Underline, Font color: Indigo

Formatted: Underline



P: ~~The~~ visualisation part will be moved to GridView.

John Can we also have upgrades in visualisation? For GridView: ~~How control this~~ told sites ~~have been told~~ not to talk to the GridView developers.

Ian: ~~Action 0706-6~~ We ~~S~~should set up a group to gather requirements and decide ~~priorities~~ priorities.

9. Security - (Dave Kelsey)

Dave: On the issue of the ~~E~~xposure of the DN in monitoring, — are sites happy?

Jeff: Do users agree to this?

Dave: The AUP says “generally yes”.

Jeff: There are ~~T~~two levels and of these levels — ~~one level~~ you can’t turn off.

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

Dave: Cases are well made ~~for~~ debugging and ~~for~~ to provide live ~~in~~ information.

Jeff: If ~~U~~Using the EGEE RB then the user publishes DNs across world! ;

Julia: It is ~~A~~also the same with the experiment submission tools.

Dave: What about ~~J~~jobs running at sites and the DN being published.

Jeff: Wow, I am ~~then humm...~~ not sure about the DN being published.

Don: Users agree to the use of information for operational/managerial and security purposes in the AUP. That phrase must include the ~~disclosure~~ disclosure of ops data with ~~traditonal~~ traditional data attached.

Dave: Did users understand this when they signed, and that this meant publishing to the world? ;

Don: The ~~L~~language could be constructed in either way. We have done it one way and must be committed to it.

Dave: It would be ~~U~~useful to have site views

VO/site trust

VO Ops Policy

Jeff: It is ~~N~~not true that we did not discuss this. I have the ~~N~~notes from 2006. The GDB did this discuss this.

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

Erwin: VOs run submission systems, and — some use the RB provided by sites. ~~---~~The RB keeps the binding between user and job. In relation to Don’s question, — how are we sure that these other submission systems keep the binding.

- Formatted: Font color: Indigo
- Formatted: Underline, Font color: Indigo
- Formatted: Font color: Indigo
- Formatted: Font: Bold
- Formatted: Bullets and Numbering

- Formatted: Underline
- Formatted: Underline
- Formatted: Underline
- Formatted: Font: Italic
- Formatted: Font: Italic



LCG Grid Deployment Board Meeting



DonON: If someone on the street picks up one of these and runs with it what happens?...

Ian: The way Don presents this, at-Fermi has a trust relationship with VOs. They Ttrust CMS that nobody breaks in to the infrastructure. This is Ffine for a single site. The fissue 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 a single job to run as multiple users then this will allow logging – glxec can do logging – but we need to trust the use.

Claudio: What iif logging is done in a central service not a local machine?

Erwin: You are Sstill trusting the information that is being published.

Dario: What is different from now? Now production jobs beenare put in by many VOs. This indicates Ttrust for the VOes already.

John: The Uuser is submitting a payload to the system.

Oxana: The Ddifferencee now is that I ban a single user if there is a problem.–But if has ...

Dario: To ban the production manager is to= ban the VO.

Agreed exec by VO:

Jeff: Strictly, using the policies as currently defined, what they are doing is not allowed. We Nneed to be consistent. We Nneed to formalise trust. A user is Ssupposed to be doing their own work.

Dave: No sharing of credentials.

Ian: The Pproduction manager is responsible – it is not code from 500 different people.

John: This is noNet well defined – how ATLAS agrees to the code.

Jeff: Due diligence is required.–

Dave: We Nneed to cope with the scaling issue – this calls for a policy document to address VO services – pilot jobs is just one example.

Jeff: Let them take responsibility

Don: AaccountabilityAccountability is the first step. A Ssite also needs to be able to authenticate the submitting party – that is a technical issue. Think spooler not pilot jobs.

John: Pilot jobs come in with user credentials.

Glxec (John Gordon)

Formatted: Underline

Formatted: Font: Bold

Formatted: Font: Bold



LCG Grid Deployment Board Meeting

GDB

The CE is acting like a gatekeeper.

Maarten: You say trust the gatekeeper but it is the same code – lcas/lcmapi.

John: Does it not use the gridmap file? Maarten:

M: The Leg LCG compute element has been configured to use VOMS or fall back to the gridmap file.---

. Jeff: Running the gatekeeper as root – this has more code than glxexec.

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.

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

Claudio: The Purpose to change is to isolate the environment. The User can kill the pilot.

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

Jeff: If there is no change then every job can-not be traced back. We've tried this. It is not Not possible to trace all the way back to the user.

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

John's Slide 6:

Questions:

a)

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

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

LHCb: Yes, and i- The if sites do not then we will ask questions to the funding agencies/agencies.

b) Jeff: This is WN consistency. If I looked at ATLAS0017 on WN4 as a user on WN8 would it be the same?-. Data managementmanagement does not care about the unix Unix ID.

Michel: At GRIF

Fabio: lcas/lcmapi can also be used to ban the user.

Jeff: You Need to have mapping on each.



LCG Grid Deployment Board Meeting



Claudio: There is consistency if you have a shared file system.

Fabio: If there no web service then we need to do something.

Jeff: The Banned list is nearly static.

Fabio: To distribute to all WNs in this way is not a long term solution. ??: A

Unique service will come.

Jeff: It is not pretty but it is a short-term solution. Consistency of mapping across WNs is dynamic.

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

c)

John: Is there a compromise? ATLAS?

Dario: What matters to us is that accounting is done correctly.

Claudio: Accounting on the batch system level - usage is -Accounted to the pilot not the user.

John: So if lcms is logging this- could also benefit accounting?

Claudio: Keep alive a thread then get-a usage for batch slot. Can you know how much a thread uses?:-

John: The ID change is logged as one user. -Another after.

Claudio: Even then you can only assign wall clock time to it.

Ian: If you do this switch then from the previous discussion, the experiments must take account of user level accounting.

Jeff: Only wall clock time of pilot jobs. We have yet to see if cpu time is properly accounted. Condor.... -

Nick: As a site do you care? You are more concerned with -Audit and traceability. You do not care how we account.

John: Action 0706-7 (John) Will discuss with Dave. Will think of a compromise and come back.

Ian: We need to get this out of the way. All sites are running pilot jobs and we- need to decide how we going to stop this.

John: Most sites see reasonable consistency.

Nick: What are the timescales now?

John: Action 0706-7 (Ian) Can this be Raised -at the operations workshop?

Ian: yes - what is the timescale for deploying glxec? This is a Question for JRA1.

Claudio: It is in the In-workplan but not a priority.

Erwin: From this discussion this priority goes up.

Claudio: This needs to go back to those doing the lcas/lcms development. OSG timescale - weeks. Of order couple of months.... Till then use the shared file system method.

Michel: Are there Draw-backs of the shared file system workaround?

Jeff: We Don't know how far it will scale - 10,000 nodes or more.s.-

Michel: No security draw backs? :-

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

Formatted: Underline

Formatted: Font color: Indigo

Formatted: Underline, Font color: Indigo

Formatted: Font color: Indigo

Formatted: Font color: Indigo

Formatted: Underline, Font color: Indigo

Formatted: Font color: Indigo



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

Ian: Sites seem to say no pilot jobs.

Jeff: Will send around decision from previous meeting.

Formatted: Underline

Workshop (Jamie Shiers)

Formatted: Font: 16 pt, Bold

There is One presentation on the status of all experiments - status - no feedback. We need Need names for the sessions.

John to report on residual services. There are two

Two-panel discussions and - sSite reviews - INFN/GridPP/ATLAS.

We then Need to collect the main points (for short summaries) for the CHEP summary.

John: What is happening with the Full scale dress rehearsals?

Jamie:

Dress rehearsals - We need to pin down dates and plans.

.

Postscript (John)

Formatted: Font: 16 pt, Bold

Glexec - something came out.

Follow on on Dave on policy issues

Erwin - job priorities group See slides.

Auditorium:

Jeremy Coles - RAL/GridPP

Tony Cass - CERN

Ian Bird - CERN

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 Vesztergombi - Budapest

Luca Dell'Agnello - INFN

Formatted: Italian (Italy)



LCG Grid Deployment Board Meeting



[Holger Martin](#) — FZK/Germany

[Laura Perini](#) — ATLAS/Milano

[Gilbert Poulard](#) — ATLAS/CERN

[Stephen Gowdy](#) — ATLAS/SLAC

[Dario Barberis](#) — ATLAS

[Claudio Grandi](#) — CERN/INFN

[Latehezar 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](#)

[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

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)



~~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)

~~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?~~

Formatted: Bullets and Numbering



~~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



LCG Grid Deployment Board Meeting



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:~~

- ~~-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~~

Formatted: Bullets and Numbering



LCG Grid Deployment Board Meeting



-Mechanism for GDB input to reach the TCG/developers

~~8-10~~-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
0603-3	Follow up to ensure all sites in country are publishing accounting data or contact John Gordon with issues preventing this happening	Country representatives	Open
0604-6	Drive forward discussions on the VOMS and protocol issues	Ian Bird	Open
0605-3	Provide feedback (with reasons) to Dave Kelsey or Kors Bos on whether the security policy presented by Dave is acceptable.	All	Open
0605-4	Tier-1s to report back to GDB on what proportion of their current WLCG work is not reported/accounted within WLCG	Tier-1 managers	Open
0606-7	Take up and discuss technical solutions for removing shared credentials from the VO boxes	Markus Schulz	Open
0607-9	Ensure the default YAIM is properly configuring lcas lcms for the sgm accounts (and that it works!)	Jeff Templon	Open
0609-1	Follow up on NDGF security policy position	Les Robertson	Open
0609-2	Look up statistics for automated on-call system and send information to GDB	Bruce Gibbard	Open
0609-6	Send storage type sampling script to John Gordon.	Jeff Templon	Open
0609-7	Move accounting to work in decimal units	Tier-1s/sites	Open



LCG Grid Deployment Board Meeting



Item No.	Description	Owner	Status
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
0706-15-2	Check use cases and VOMS need for failover with the developers and VOs	Maria Dimou	Open
0706-2	Provide description of implementation(s) of VOMS based ACLs and submit this to the experiments to confirm it satisfies their requirements.	Flavia Donno	Open
0706-3	Review the membership and approach of the Job Priorities Working Group	Erwin Laure	Open
0706-4	Nominate someone to join the grid services monitoring work	Oxana Smirnova	Open
0706-5	Follow up on how best to proceed with site-experiment negotiation on what VO SAM tests are to be monitored	John Gordon	Open
0706-6	Setup group to gather and prioritise GridView requirements	Ian Bird/ John Gordon	Open
0706-7	Follow up c) with Dave Kelsey	John Gordon	Open
0706-8	Raise glexec questions at the Stockholm operations workshop	Ian Bird	Open

Formatted Table



LCG Grid Deployment Board Meeting



List of Attendees

X means attended
V means attended via VRVS

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 Wezel	
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 Kester		Farid Ould-Saada	
Pakistan	Hafeez Hoorani		-	
Poland	Ryszard Gokieli	V	Jan Krolikowski	
Portugal	Gaspar Barreira		Jorge Gomes	
Russia	Alexander Kryukov		Vladimir Korenkov	
Spain	Manuel Delfino		Xavier Espinal	
Sweden	Niclas Andersson		Tord Ekelof	
Switzerland	Christoph Grab	X	Marie-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	V	-	

Formatted: Centered

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	
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		
Planning Officer	Alberto Aimar	X		

Country	Member	Present?	Deputy or Technical Assistant	
Austria	Dietmar Kuhn	XX	-	-
Canada	Reda Tafirout	-	Mike Vetterli	-
Czech Republic	Milos Lokajicek	X	-	-
Denmark	John Renner Hansen	-	Anders Waananen	-
Finland	Klaus Lindberg	-	Jukka Klem	X
France	Fabio Hernandez	XX	Dominique Boutigny	-
Germany	Klaus-Peter Mickel	-	Holger Marten, Jos van Wezel	X, XV
Hungary	Gyorgy Vesztergombi	X	Dezso Horvath	-
India	P.S Dhekne	-	-	-
Israel	Lorne Levinson	-	-	-
Italy	Mirco Mazzucato	-	Luciano Gaido	-
Japan	Hiroshi Sakamoto	X	Tatsuo Kawamoto	-
Netherlands	Jeff Templon	XV	Ron Trompert	-
Norway	Jacko Koster	-	Farid Ould-Saada	-
Pakistan	Hafeez Hoorani	-	-	-
Poland	Ryszard Gokieli	X	Jan Krolikowski	-
Portugal	Gaspar Barreira	-	Jorge Gomes	-
Romania	Mihnea Dulea	-	-	-
Russia	Alexander Kryukov	-	Vladimir Korenkov	-
Spain	Jose Hernandez	V	Xavi Espinal	X
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	X
United States	Ruth Pordes	VV	Michael Ernst	V(em)
CERN	Tony Cass	XX	-	-
ALICE	Alberto Masoni	X	Yves Schutz	-



LCG Grid Deployment Board Meeting



	Federico Carminati	X		
ATLAS	Kors Bos	XX	Stephen Gowdy	XX
	Dario Barberis	X		
CMS	Matthias Kasemann	V	Patricia McBride	
LHCb	Ricardo Graciani	VV	Andrei Tsaregorodstev	
	Nick Brook	XV		
Project Leader	Les Robertson	X		
GDB Chair	John Gordon	X		
GDB Secretary	Jeremy Coles	X		
Grid Deployment Mgr	Ian Bird	XX	Markus Schulz	XX
Fabric Manager	Bernd Panzer			
Application Manager	Pere Mato Vila			
Security WG	David Kelsey	XX		
Quattor WG	Michel Jouvin	X-X		
Networking WG	David Foster			
Planning Officer	Alberto Aimar	XX		

Others present at CERN

[Harry Renshall – CERN](#)

[Jamie Shiers – CERN](#)

[Oxana Smirnova – NDGF](#)

[Jim Shank – Boston uni.](#)

[Etienne Urbah – LAL](#)

[Luca Dell’Agnello – INFN](#)

[Laura Perini – ATLAS/Milano](#)

[Gilbert Poulard – ATLAS/CERN](#)

[Claudio Grandi – CERN/INFN](#)

[Latchezar Betev – ALICE](#)

[Sue Foffano – CERN](#)

[Andrew Elwell – Glasgow/GridPP](#)

[Greig Cowan – Edinburgh/GridPP](#)

[Erwin Laure – EGEE](#)

[Frederique Chollet – LAPP/IN2P3](#)

[Jean Philippe Baud – CERN](#)

[Sophie Lemaitre – CERN](#)

[Alberto Masoni – ALICE](#)

[Simone Campana – CERN](#)

[Flavia Donno – CERN](#)

Others on VRVS

[Ricardo Graciani](#)

[Alessandra Forti](#)

[Stefano Belforte](#)

[Donald Petravick](#)

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

Formatted: Left

Formatted: English (U.S.)

Formatted: French (France)

Formatted: Italian (Italy)

Formatted: English (U.S.)

Formatted: Italian (Italy)

Formatted: English (U.S.)

Formatted: Font: Bold, English (U.K.)

Formatted: English (U.S.)



Others present at CERN

Jamie Shiers

Harry Renshall

M. Lameme

Flavia Donno

Sue Foffano (CERN)

Simone Campane (CERN)

Nechaersky 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'Agnello (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 — Ancecy

Marek Domaraeky — Bern

Olivier van der Aa — London

David Colling — London

Dave Kant — RAL

Pete Gronbech — Oxford

Gabriel Stocica

Lief Nixon — Linkoping

Frederique Chollet

Greig Cowan

Stefano Belforte

Gonzalo Merino

Richard Gokieli

Alvaro Fernandez (PIC)

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: Not Bold, Italian (Italy)

Formatted: Font: Not Bold

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

Formatted: English (U.S.)

Formatted: English (U.S.)

Formatted: Italian (Italy)

Formatted: Italian (Italy)



LCG Grid Deployment Board Meeting

GDB

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

Formatted: English (U.S.)