



Version 1.0

Amendments history

<i>Name</i>	<i>Area</i>	<i>Date</i>
<i>Jeremy Coles</i>	<i>All - full edit -> v1.0</i>	<i>9th February 2007</i>

Minutes of the meeting

CERN, 7th February 2007

Agenda: <http://indico.cern.ch/conferenceDisplay.py?confId=8469>

Minutes: Jeremy Coles

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

Meeting Summary

Provided by Kors - <https://twiki.cern.ch/twiki/bin/view/LCG/GridDeploymentBoard>

Detailed minutes

1. Introduction (Kors Bos)

Kors reviewed actions from the last meeting. Only the action on sites to check accounting remains pending/ongoing. The next GDB meeting is on 7th March at CERN. The Prague meeting is April 3rd (Tuesday – pre-GDB will look at Tier-2/Prague issues) till April 4th. Registration is open here: <http://www.particle.cz/gdb>.

Kors reflected on the openness of the GDB as discussed at the WLCG collaboration board – it was noted that Tier-2s are under represented.

The job priority scheme was reviewed. It is installed at NIKHEF but was a concern at FZK. What is the plan? Jeff: There is some documentation and many of the Tier-1s have done the configuration by hand. Kors: Was it dealt with at the workshop?

Jeff: There was a post on rollout explaining what to do Experiments using the gLite WMS should be able to use it but VO views are there even if the WMS is not used.



LCG Grid Deployment Board Meeting



Kors: Does it need to be installed everywhere? Jeff: It would be tough to get it everywhere. If the VO view is not present then they need to match on something else (for example at an LSF site). In the MB we were asked for a plan for all Tier-1s. I looked at Holger's message but could not understand the problem.

There was a short reminder on some security issues that came up previously, such as independently generated monitoring data. Dave Kelsey thought many would or could be covered in the new services policy where the policy would depend on the sensitivity of the material. A conversation with Ian Neilson revealed that there are already requirements in the monitoring area. Use of server certificates (VOMS uses host certificate to sign) being tied to individuals was acknowledged to be an ongoing discussion in forums such as the EU PMA meetings.

2. New chairperson selection (Gonzalo Merino)

Gonzalo introduced the topic and confirmed that there were two candidates who agreed to stand: Michel Jouvin and John Gordon. Both were invited to give a brief speech to support their nomination.

John Gordon presented his background in particle physics and described a number of roles he had taken from being a contributor to the Hoffmann report to leading in e-Science work at RAL and serving within the GridPP management. Hardening of the services would be one of his aims which led to his final remark about being from Glasgow.

Michel Jouvin explained his strong participation in HEPiX and belief in the importance of such forums to make progress. He explained his role within the GRIF Tier-2 and that his nomination would strengthen the representation of Tier-2s (which need to be more involved). He observed that it is hard to say the changes which are coming.

Gonzalo explained the process and asked for clarification on some points. Fabienne Baud-Lavigne then handed out a voting envelope to each voting country representative. Votes were also received via email at and before the meeting.

After confirming the votes cast Gonzalo announced that John Gordon had won the vote and congratulated him. In return John thanked everyone for their confidence in him and said that he could not promise to do better than Kors and said that he would like to thank Kors for his contribution over the last few years – Kors received a well deserved applause.



3. SL4 Status and Plans (Markus Schulz)

Markus hoped that this would be one of the last status reports that would be needed. The last was at the 6th December meeting. Today's talk reviewed the timeline and what was failing in the ETICS build system (slide 3). Efforts have been refocused on the UI 32-bit SL4 (74% of build working) and WN 32-bit SL4 (82% of build working). There are still some missing externals and given the serial nature of the build process, the build fails until all things on which it depends succeed.

John: How do you get to 75% then?

Markus: Some packages are self contained and can succeed.

Markus posed the question: When will it end? The slide answered: not within the next 2 weeks. 4 weeks of testing are required AFTER the first successful build. On the positive side, new versions of VDT, Globus and Condor were available – there are many things to test once the build works. .

Jeff: You talk of dependencies, does it at least compile? Markus: ETICS does try to compile. Jeff: 74% builds – this is the number of packages not the amount of code? Markus: Yes. Jos: Is there still time to iron out things which crop up in SL5? Markus: The AAM dependency journal is to be cleaned up so porting to the next OS is easier. The plan is draft and put to the PMB on Friday. It is not clear how to do structural improvement [to the process] when always on the deadline. If we delay SRM2.2 and the move to SL4 etc. then the group will have 10% of their time for this! It is expected that the PMB will recommend an approach.

Claudio Grandi: We are investing a lot in ETICS as it will allow new platform builds to be much better in the future. Also, the binary incompatibility between SL3 and SL4 is not there for SL5. Federico: It is worrying to add more external dependencies. Markus: The way the code is written you will always have a large number....the philosophy is perhaps wrong but this is how it was written.

Current status – slide 6. VDT1.6 (32-bit and 64-bit) is now available. First production tests on PPS by 3 of 4 experiments have taken place. There is a problem with APT - updates do not work leading to a situation a bit “like the hotel California”.

The problem is not ETICS. Life without ETICS would be at least as bad as it is now.

Jeff: On the last slide you said ETICS is not the problem. I want to support the idea to have effort for a plan. This is the fourth move to a new build system and the problems always remain because we have not changed the way we develop the software. Markus: I'm not so optimistic. With all the current missing functionality I don't see how developers can [free up time to] do major things. Some components have a code base over 5 years old!



LCG Grid Deployment Board Meeting



Jeff: Users need to know if we do not do this in the next 7 months then there will be less time in the future. Markus: Functionality changes should go down so perhaps it will not be so bad. Kors: Experiments are against the wall too as the functionality is needed.

Tony: Slide 9: What I am missing is some clear plan – you present a number of problems. We will not have something till April. We have FDRs (full dress rehearsals) coming up. There is an option to use SL3 software on SL4 but a problem with a single RPM upgrade. The other option is a more complex solution but these people [those who would develop the solution] could work on the SL4 build. What will we have as a stable solution for the FDRs?

Markus: The target for the FDRs is to get a generic SL4 build. We can not make a firm date. Two alternative solutions are provided that are working. Tony: But given this the earliest possibility for deployment is the end of April. Is that acceptable? These are the questions that should be being asked. There is a plan but we need to understand the implications. Sites will later come under pressure to deploy. Markus: Sites can use the tarball solution or APT. They can use SL4 WNs now. The SL3 version on SL4 has been extensively tested.

Les: The time scale is extremely tight, therefore we should implement the backup solution. If we get the proper solution then we revisit. Tony: Do all sites understand the backup? Markus: PPS is installed with the middleware and these have the same installation mechanism [as production sites]. Tony: It would be interesting to see a site-by-site list of what is expected to make sure what is needed will be available for the FDRs.

Les: Is the full solution really needed for FDRs or can it wait till ready? Is it so critical? Markus: Full solution will make situation more complex for experiments as their software needs to be built against two flavours – SL3 and SL4. Not all sites will convert at the same time. Tony: The experiments should comment. Do they have an SL4 (only) based version for their next release?

Dario: For ATLAS, we can run SL3 release on SL4 provided all compatibility libraries are installed. The March/April release will be distributed in both versions. By the end of the summer/ start of autumn, we only build on SL4 unless the transition is not completed. We would like to drop the SL3 build.

Michel: Is this really an APT problem? Markus: If you install the RPMs then you lose control over dependencies. Michel: Nothing was forced at GRIF. This is for WNs. You can reinstall WNs using APT. Markus: APT has a problem with the upgrade later. Michel: Our dependencies were satisfied. Markus: Then perhaps everyone converts to Quattor!



LCG Grid Deployment Board Meeting



Matthias: For CMS we foresee the next February build (deployable from March) should be based on SL4. If we need to build on SL3 we will do it but the intention was to build for SL4. Markus: Are sites aware of the need to move to SL4 for CMS? Les: This is not realistic now. Markus: If the software is *built* on SL4 then it will not work with the backup plan. Matthias: CMS plans will need to change.

Frederico: Our binary was built on 3, not sure yet whether it will run on 4. The problem for us may then be the need to maintain one more package.

Nick: For LHCb this is in principle not a problem. Our current software can run on both platforms. There are concerns about future builds. The GFAL libraries may cause problems in the future.

Kors: So the conclusion is that we need to ask CMS to relax their requirements. Others are okay. Les: At present there are so many open things. We need to come back to this when there is a good estimation of the end-date. We are now getting into a period where we want stable services, so rollout after March is more difficult. There needs to be an agreement between experiments and sites for a rollout strategy. The later it gets the more nervous the feeling.

Tony: There is an impact on people with machines that run SL3 and SL4. They need to support a mixed environment. Markus: In a production environment we need to expect a mix.

[Action 0702-1 Markus to send email to GDB list after 2 weeks giving a status update \(this to ensure that the information reaches the Tier-2s\).](#)

[Action 0702-2 John to reserve a slot at next GDB or perhaps pre-GDB to revisit the SL4 situation.](#)

Nick: It would be useful if this is monitored closely by the MB since their role is to manage the project. Markus was concerned at the number of places the status needs to be reported in addition to the EGEE TCG, PMB etc....

John: The report would just be a copy of reports sent elsewhere – it is not extra work.

4. Security Policies update (Dave Kelsey)

Dave explained that the policies had been exposed in a number of places (GDB, ROC managers, JSPG) but there had been zero feedback from these. Some comments were received from sites and the presentation reviewed these and the response. It was noted that the policy is common for the operational infrastructures that any given grid would have a covering letter explaining the terms and conditions in that context. The situation



LCG Grid Deployment Board Meeting



for WestGrid in Canada (similar to NDGF) was less well known than in other WLCG grid contributors. Jeff asked what it is that WestGrid thinks it is joining when they sign.

(Action already exists) [Request to GDB – comment on the Grid Site Operations Policy when it comes out. Also the new Grid Security Policy. We should aim for approval soon.](#)

In calling for feedback Kors suggested that Dave explicitly ask for the request to be forwarded to people represented but not on the GDB list.

5. Report from the Registration Task Force (Maria Dimou)

Maria outlined the history behind the registration tasks. Dave asked that the GDB thank Maria and her group for all their effort in this area over the last few years – the result is a major achievement. [Members applauded].

Kors asked how this work should go on. Claudio: There is representation where gLite integration is tackled, for example at the EGEE EMT meeting ... in addition developer and experiment exchange takes place in the TCG. There are also VO manager meetings inside EGEE. The real point is do we need very controlled effort or can we use the standard approaches? Maria then asked what we do with VOMRS. Claudio: VDT responsibilities are in the EMT meeting. We can invite Tania to be present. Not sure we need to have a task force. Jeff: The TCG was recently sent a list about VOMS development issues, many there will not even know the generic attributes. Claudio: In the VOMRS discussions during the WLCG workshop there was the same response. Jeff: I do not understand what this is all about – considering the requirements for VOMS when the situation is not known suggests the TCG alone is not enough. Claudio: A TCG task force perhaps? Markus: Jeff suggested people who do not know anything about VOMS should be included in the discussion but I do not agree with that. Ruth [VRVS]: I think there needs to be an established set of meetings. The task force has served a useful function here. The TCG is not sufficient for this. There is a “joint project” between VOMRS and VOMS that needs to be structured in some way.

Kors: I propose Dave, Maria and a few others go offline and discuss this point.

John: There is a parallel with SRM here since that also is not affiliated with just one body
Kors: Then we need to redefine the mandate.

[Action 0702-3 Dave, Maria et. al. to discuss the future of a VOMRS-VOMS task force and consider possible mandates for the group.](#)



6. Report from the WLCG week (Jamie Shiers)

Jamie gave a broad overview of the meeting which if all associated meetings are counted ran for almost 2 weeks (dCache workshop, tutorials, joint operations workshop etc.). All feedback was followed up. Walk-ins represented a significant overhead. Registration for sessions was not used – John thought this would help with room planning – and Jamie was not sure of the indico functionality in this area.

Jeff: The site and regional reports had a lot of information and overlap. It would be better to spread these over the week. Fewer people were present. Before arranging these it would be good to figure out what people will find interesting and useful to prepare. John: Limiting on size is perhaps the only useful guide.

Data management was a popular session. Stability of storage services at sites is a #1 issue. From the VOMRS session people were reminded of the [mid-February deadline for VOMRS requirements](#).

GGUS problems stemmed largely from a resourcing issue within the support units. Monitoring provided a popular topic as did discussion of WLCG SRM2.2: A new GSSD group was introduced to the SRM 2.2 developers – GSSD stands for Grid Storage System Deployment. Migration strategy for data was a hot topic..

From the experiment sessions it was thought that dropping the name “tutorial” would be better. All proved useful.

EGEE (All ROC managers) ARM & (CIC on Duty) COD meetings are next to be held in Stockholm 11th-12th June.

Issues raised at the joint operations workshop – provision of 24x7 services. Interventions and how they should be conducted around machine operations. Service stability. Integration of grid and experiment operations.

In conclusion the workshop proved very useful and received a lot of positive feedback. There is a proposal to hold one WLCG collaboration week per year - the next is provisionally set for April 2008.

LUNCH

7. Recent workshop updates to site and experiment plans (Harry Renshall)



LCG Grid Deployment Board Meeting



Jeff: Did anyone report about the MB decision? Harry: No Jeff: The decision was to submit figures to you [Harry] – the Harry-spreadsheets would then be the main source of resource information.

Luca: we are putting in production 400TB. Physical space is currently enough for 200TB. In 1 month we will install another 1500 KSI2K – thus an additional 700 KSI2K over the commitment.

Kors: What does RAL do with resources that are not allocated and why? John: We will not allocate them since it is hard to get the disk back – allocations are based on decisions of our User Board.

There was some discussion about ATLAS figures on slide 8. It should be 40 million for the quarter. The discussion was taken “offline”.

Jeff questioned the meaning of figures in the spreadsheet shown (Tier_1_capacity_Jan_2007_workshop_increases_and_required_2A2007.xls). Harry explained that figures in the 1Q column meant available at the start of that quarter (i.e. should be there by now). The numbers given were from baseline in accounting and then added on top were the requirements from the experiments.

[Action 0702-4 Tier-1s - please look at Harry's resource tables and try to understand them!](#)

8. Database workshop summary (Maria Girone)

On slide 11 Dario asked if there was a synchronisation requirement – Maria said there was not. Jos then asked if the upgrades required downtime. Maria: Major software upgrades (1-2 times per year) may take a few hours but that is all.

John: Are you only concerning yourself with 3D or are there experiment moves to use other technologies? Maria: Two technologies – replication streams ATLAS & LHCb are going to deploy conditions this way. CMS will use Frontier. John: No other central databases? Dirk: Some sites are using similar database application with data shipped in the other direction – in ATLAS for example. Shipping tag collections through ROOT files is possible.

Matthias: What is the status for the database hardware to be on non-interpretable PSUs?
?: All production servers are currently on this but investigating power requirements
Tony: All services are on UPS but here we are talking about.



9. Database Requirements Update (Alexandre Vaniachine) – Sacha?

Jeff: Dirk mentioned the number of nodes a while ago and they were the same then as now (slide 13). Then it was a guess. These machines are real money so we are not comfortable maxing out the resources when it is not shown that it is needed.

Alexandre: We need to run 3D challenge to stress test the number of nodes needed. With 1000s of jobs accessing the same data we expect to survive with lots of memory in a small number of nodes. Jeff: So this is a small number? Alexandre: Correct, we expect to run 1 million jobs per day. Dirk: This requirement has been put forward but not absorbed. The request is 3 nodes – before tests the number of nodes is unknown – it depends on the number of queries and size of connections and these are not known. It is best not to have too long to deploy otherwise this will potentially hold up ATLAS work. We can have one more upgrade after this perhaps on the September timescale. We need experiment feedback to get the number correct. It is hoped that estimates from T0 results are half right.

Dietrich: Many of the parameters are known so we could calculate...

Alexandre: For our estimate we assume Oracle will not perform worse than MySQL.

Dirk: We need real jobs running to get the scale factor. People will not believe numbers without a number of jobs successfully running. Jeff: There are other things to remember – these things were not planned in the MoU. Resources for this come out of tape/disk/CPU allocations.

Fabio: we can proceed as with the other requirements and have a table of resources showing what and where. Dirk: Will collect similar table for each experiment. Fabio: Another site worry is the number of open connections to the sites. Dirk: At T0 we see many idle connections which is CPU expensive. Scaling depends on how successful the experiments are at getting short active connections. Kors: When do the experiments plan to start using conditions databases? Dirk: About now.

Kors: It is fair to say that the databases are in a reasonable state. Dirk promised to sit down and produce a “Harry table” with the number of connections/requirements. The experiments need to start using the infrastructure.

Alexandre: It is difficult for ATLAS to understand that the resources are shared. Dirk: Only IN2P3 has shared resources. Fabio: We will keep this unless it is shown not to work. Alexandre: Dedicated storage too?

Dirk: Any backup on security policies? Are they approved or is another round of discussion required. All database contacts have seen the proposal. Suggest we approve it unless anybody objects [relates to Maria’s talk - backup and security].



Jamie: It is important for the applications. Security patches which impact site – for example if there is a new LFC at the T0 then others will have to deploy too ... if there is a new release of one of the packages then the document requires other sites to upgrade within a given period – e.g. 2 weeks. There are two documents.

There was no disagreement. Kors considered the experts who needed to know about the document did already and so if they are happy the policies can be approved.

10. Data Access Patterns for ATLAS (Dietrich Liko)

Kors explained that this talk was originally planned for the pre-GDB discussion but Dietrich was unable to attend yesterday.

Dario: A major problem is that required software is not found on WN because AFS or similar fails. This creates a black hole effect for that WN.

Jeff: Are all required grid tools in place – if you use just these can you do reprocessing? (Slide 7). John: You are not distinguishing between jobs which expect data to be there and others that do not... Kors: The work has not developed to that level yet.

Dietrich: SAM tests should be extended to test also the posix I/O on the SE.

Fabio: What are the bandwidth needs for the WNs? Dietrich: Analysis jobs need ~2MB/s each so multiply this by the size of the site. The effective capacity needed will be a bit lower due to setup overheads etc. Analysis jobs require about half the size of production jobs. Gonzalo: Will these figures change? Dietrich: If downloading gFTP files then this adds a requirement for SE bandwidth too. This is not excluded in the future to optimise workflow. Dario: Reconstruction takes 10-20s per event. Jobs read in 10MB every 10-20s. For reconstruction the job processes all events in a file. For analysis the most likely access pattern is to access individual events in many files.

Gonzalo: Under the remit of the new GSSD this will be discussed. Spurious streams will impact on the implementation of the systems. For now we will assume short high bandwidth connections.

Greig Cowan: How do you know you have the space on the WNs. Jeff: If 4 jobs pulling in 10 files on a single node you quickly need 100GB of space per WN. Remember some sites are using WN disk as part of their SE. You need to be aware of these constraints. Dietrich: We will need to consider this. Gonzalo: Also you have to make sure the job deletes the files and frees up space for next jobs. Dietrich: The data management group are working on this. Michel: On moving to machines with 4 cores then 8 cores this becomes more of an issue. Dietrich: Assume 10 GB per job. Fabio: This should be a requirement – amount of disk – passed to the site.



LCG Grid Deployment Board Meeting



John: In terms of how much is available or configured, nothing is passed on to the batch system at the moment. Fabio: Disk space is also a requirement so it should also be in the JDL and passed to the batch system. Dietrich: One should agree on some reasonable minimum. Michel: If the space is wrong then the job crashes so this is worse than the memory issue.

Greig: I agree that the SAM tests should test posix access.

11. Report from the storage classes task force (Flavia Donno)

The task force has been using pre-GDB slots and it was agreed that this can continue.

Jeff: I thought we were moving to SRMv2 because there were things we could not do without it. Your bullet (on the work in progress slide) suggests we can use v1 and v2. John: One use case – anything in v1 is readable with v2. Maarten: For all ongoing production activities the v2 end-point is used (for CASTOR). You could move stuff in the v1 instance to v2.2. Not all the data need be accessible by a single end-point. New data goes into the new instance. [The version field is in the information system]

Tony: This is theoretical vs practical possibilities. We will not support the v1 instance beyond the minimum time to test v2. We will close v1 as soon as data is recataloged. Maarten: It is not a technical requirement. You can copy stuff you want and leave the rest. A mixed environment is tolerable. Flavia: Each has an entry in the catalogue – you could also have a replica in the catalogue.... Jeff: That is the point - different endpoints even if it is the same file.

Michel: The point about backdoors was discussed during the MB some weeks ago; the SRM is not an access protocol – the only generally available protocol is rfiio etc. Do we need any other access protocol? Maarten: Is SRM really managing the data. These user data or experiment data are output of user jobs (not random files) – so the file in principle appear in catalogues. Alexandre: SRM does not support local access protocols? Flavia/Maarten – those examples were chosen as the most popular – the SRM is not limited like that. Dietrich: As a user I can also open a file locally, do I gain something by using SRM to access data on disk? Maarten: If you can do that is you can predict the TURL! One reason is to avoid a particular disk server getting overloaded. Tony: It has been explicitly stated that all access to data is via SRM!

Tony: The conclusion and important thing for sites and the discussion group is whether disk buffers are sized correctly. If this is so we will not get into all these arguments. Kors: How? Tony: The experiments say how often they intend to reprocess and provide information about the jobs. Kors: If you underestimate T1D0 buffers the risk is the job



LCG Grid Deployment Board Meeting



may be waiting for a while? Maarten: If the data is on Tape it can go to T1D0 or T1D1. If T1D0 then the system can remove files and maintain space. John: If you try to use T1D1 for pinning then you have to manually maintain the space – if the disk gets full you will fail.

Kors: I propose we use the next pre-GDB to discuss more. John: I need to take consultation on times for the pre-GDB in order to find parallel rooms etc. Flavia: Most of the work can be done in the working groups and then the meeting used to exchange progress. John: Yes, the pre-GDB should be to exchange technical progress.

12. AOB

Milos invited everyone to the Prague GDB meeting. Registration is open

Dario: Milos you should state how to enter the date on the form! [ed. Looks like two text fields].

MEETING CLOSED AT 16:40

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
0605-5	Tier-1s to gather and publish (to the GDB) storage data. 8 numbers required for disk (allocated and used per experiment). 4 numbers for tape (allocated per experiment). This data is to be gathered at the end of each month.	Tier-1 managers	Ongoing
0606-6	Raise package management as a high priority for WLCG at the next TCG meeting	Erwin Laure	Open
0606-7	Take up and discuss technical solutions for removing shared credentials from the VO boxes	Markus Schulz	Open
0607-8	Urge experiments to push users to re-register and inform tha a deadline will be imposed at the next GDB. Circulate job priorities document.	Kors Bos	Open
0607-9	Ensure the default YAIM is properly configuring lcas lmaps for the sgm accounts (and that it works!)	Jeff Templon	Open
0607-10	Arrange a report for the September meeting on VOMS awareness in each of the SRM implementations to date	Kors Bos	Open
0609-1	Follow up on NDGF security policy position	Les Robertson	Open



LCG Grid Deployment Board Meeting



Item No.	Description	Owner	Status
0609-2	Look up statistics for automated on-call system and send information to GDB	Bruce Gibbard	Open
0609-3	Follow up VO-Box concerns with ALICE (a class-2 solution was put forward but not taken up by ALICE)	Kors Bos	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
0702-1	Send mail to GDB list after 2 weeks giving a status update on the SL4 work	Markus Schulz	Open
0702-2	Reserve a slot at the next GDB or pre-GDB to revisit the SL4 situation	John Gordon	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

List of Attendees

X means attended

V means attended via VRVS

Country	Member		Deputy	
Austria	Dietmar Kuhn	<input type="checkbox"/>		<input type="checkbox"/>
Canada	M Vetterli	<input type="checkbox"/>	R Tafirout	<input type="checkbox"/>
Czech Republic	Milos Lokajicek	<input type="checkbox"/>	Jiri Kosina	<input type="checkbox"/>
Denmark	John Renner Hansen	<input type="checkbox"/>	Anders Waananen	<input type="checkbox"/>
Finland	Klaus Lindberg	<input type="checkbox"/>	Jukka Klem	<input checked="" type="checkbox"/>
France	Fabio Hernandez	<input type="checkbox"/>	Dominique Boutigny	<input type="checkbox"/>
Germany	Klaus-Peter Mickel	<input type="checkbox"/>	Holger Marten	<input type="checkbox"/>
		<input type="checkbox"/>	Jos van Wezel	<input checked="" type="checkbox"/>
Hungary	Gyorgy Vesztergombi	<input checked="" type="checkbox"/>	Dezso Horvath	<input type="checkbox"/>
India	P.S Dhekne	<input type="checkbox"/>	B. Vinod Kumar	<input type="checkbox"/>
Israel	Lorne Levinson	<input type="checkbox"/>		<input type="checkbox"/>
Italy	Mirco Mazzucato	<input type="checkbox"/>	Luciano Gaido	<input type="checkbox"/>
Japan	Hiroshi Sakamoto	<input type="checkbox"/>	Tatsuo Kawamoto	<input type="checkbox"/>
Netherlands	Jeff Templon	<input checked="" type="checkbox"/>	Ron Trompert	<input type="checkbox"/>



LCG Grid Deployment Board Meeting



Country	Member		Deputy	
Norway	Jacko Koster	<input type="checkbox"/>	Farid Ould-Saada	<input type="checkbox"/>
Pakistan	Hafeez Hoorani	<input type="checkbox"/>		<input type="checkbox"/>
Poland	Ryszard Gokieli	<input type="checkbox"/>	Jan Krolkowski	<input type="checkbox"/>
Portugal	Gaspar Barreira	<input type="checkbox"/>	Jorge Gomes	<input type="checkbox"/>
Russia	Alexander Kryukov	<input type="checkbox"/>	Vladimir Korenkov	<input type="checkbox"/>
Spain	Manuel Delfino	<input type="checkbox"/>	Xavier Espinal	<input checked="" type="checkbox"/>
Sweden	Niclas Andersson	<input type="checkbox"/>	Tord Ekelof	<input type="checkbox"/>
Switzerland	Christoph Grab	<input type="checkbox"/>	Marie-Christine Sawley	<input type="checkbox"/>
Taiwan	Simon Lin	<input checked="" type="checkbox"/>	Di Qing	<input checked="" type="checkbox"/>
United Kingdom	John Gordon	<input checked="" type="checkbox"/>	Jeremy Coles	<input checked="" type="checkbox"/>
United States	Ruth Pordes	<input checked="" type="checkbox"/>	Bruce Gibbard	<input type="checkbox"/>
CERN	Tony Cass	<input checked="" type="checkbox"/>		<input type="checkbox"/>
ALICE	Alberto Masoni	<input type="checkbox"/>	Yves Schutz	<input type="checkbox"/>
	Federico Carminati	<input checked="" type="checkbox"/>		<input type="checkbox"/>
ATLAS	Gilbert Poulard	<input checked="" type="checkbox"/>	Laura Perini	<input checked="" type="checkbox"/>
	Dario Barberis	<input checked="" type="checkbox"/>		<input type="checkbox"/>
CMS	Lothar Bauerdick	<input type="checkbox"/>	Tony Wildish	<input type="checkbox"/>
	Stefano Belforte	<input checked="" type="checkbox"/>		<input type="checkbox"/>
LHCb	Ricardo Graciani	<input checked="" type="checkbox"/>	Andrei Tsaregorodstev	<input type="checkbox"/>
	Nick Brook	<input checked="" type="checkbox"/>		<input type="checkbox"/>
Project Leader	Les Robertson	<input checked="" type="checkbox"/>		<input type="checkbox"/>
GDB Chair	Kors Bos	<input checked="" type="checkbox"/>		<input type="checkbox"/>
GDB Secretary	Jeremy Coles	<input checked="" type="checkbox"/>		<input type="checkbox"/>
Grid Deployment Mgr	Ian Bird	<input type="checkbox"/>	Markus Schulz	<input checked="" type="checkbox"/>
Fabric Manager	Bernd Panzer	<input type="checkbox"/>		<input type="checkbox"/>
Application Manager	Pete Mato Vila	<input type="checkbox"/>		<input type="checkbox"/>
Security WG	David Kelsey	<input checked="" type="checkbox"/>		<input type="checkbox"/>
Quattor WG	Charles Loomis	<input type="checkbox"/>		<input type="checkbox"/>
Networking WG	David Foster	<input type="checkbox"/>		<input type="checkbox"/>
Planning Officer	Alberto Aimar	<input checked="" type="checkbox"/>		<input type="checkbox"/>

The following also attended:

Name	Area	Name	Area
Matthias Kasemann	CERN	Leif Nixon	Sweden
Oxana Smirnova	NDGF	Claudio Grandi	INFN-JRA1
David Colling	London Tier-2	Dietmar Kuhn	CERN



LCG Grid Deployment Board Meeting



Stephen Gowdy	ATLAS	Luca dell' Aguello	INFN-CNAF
Michel Jouvin	GRIF	Gonzalo Merino	PIC
Maria Dimou	CERN		

Attending via VRVS:

Nick Brook (LHCb)
Ruth Pordes (OSG – Chicago)
Jose Hernandez (Madrid)
Christoph Grab
Reda Tafirout
Frederique Chollet (Annecy)
Stefano Belforte
Peter Gronbech (Oxford)
Mirco Ciriello

Ricardo Graciani (Barcelona)
Davide Salomoni (Bologna)
Alexandre Vaniachino

At meeting during chairman vote:

Jeremy Coles (RAL – GridPP)
Matthias Kasemann
Oxana Smirnova (NDG)
Leif Nixon (Sweden)
Claudio Grandi (INFN)
David Colling (London Tier-2)
Dario Barberis (ATLAS)
Alberto Aimar (CERN)
Dietmar Kuhn (CERN)
G Veszterponbi (Budapest)
Alberto Masoni (ALICE)
Stephen Gowdy (ATLAS)
Jos van Wezel (FZK/GridKa)
Jukka Klem (HIP Helsinki)
Tony Cass (CERN)
Jeff Templon (NIKHEF)
Les Robertson (CERN)
Dave Kelsey (RAL)
Simon Lin (ASGC)
Di Qing (ASGC)
Luca dell' Aguello (INFN-CNAF)
Laura Perini (ATLAS)
Gilbert Poulard (ATLAS)
John Gordon (UK)
Michel Jouvin (GRIF)



LCG Grid Deployment Board Meeting



Xavier Espinal (PIC/IFAE)
Gonzalo Merino (PIC)
Jamie Shiers (CERN)
Federico Carminati (CERN)