



Version 1.0

Amendments history:

<i>Name</i>	<i>Area</i>	<i>Date</i>
<i>Jeremy Coles</i>	<i>All – draft notes</i>	<i>12/03/08</i>
<i>“</i>	<i>Further revision</i>	<i>31/03/08</i>

Minutes of the meeting

CERN, 5th March 2008

Agenda: <http://indico.cern.ch/conferenceDisplay.py?confId=20227>

GDB twiki: <https://twiki.cern.ch/twiki/bin/view/LCG/GridDeploymentBoard>

Minutes: Jeremy Coles

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

Meeting Summary

(John Gordon)

Detailed minutes

1. Introduction (John Gordon)

John provided an overview of recent meetings and progress. It was noted that VRVS is no longer supported and the GDB has moved to using EVO. Upcoming issues include the SAM critical tests and whether they achieve what is needed, WN client rollout and publishing the power figures for a cluster (with appropriate use of sub-cluster).



2. CPU efficiency (Ulrich Schwickerath)

Ulrich presented results from looking at the efficiency of CPU usage seen from the CERN site and compared this with the experiment views. There was clear variation for all the experiments across the months of 2007. One reason for the discrepancy is that jobs have an overhead (for downloading components, accessing remote SEs from WNs etc.) and there is wasted time between “joblets”.

There was discussion about the LHCb case. Jeff Templon (JT) thought Condor was responsible for the overheads in this case and that it was resolved. Claugio Grandi (CG) commented that it was his understanding that DIRAC downloads more than one job if possible and this would impact the results. Philippe Charpentier (CP) answered that this was a misunderstanding. Even for “same users” LHCb have not enabled it, however it is difficult to know how much CPU time is left. On the pilot job framework he added that if the software required by the job is not installed then it may take 10 minutes to set up the needed environment.

JT: There is one thing missing from the list. What if for some reason the pilots do not want to run – i.e. the payload does not work. We used to have timeouts.

PC: Now the job stops immediately – if the TURLs for files the job needs do not succeed then the job fails.

JT: It should be looking for a “bit” first, and if it is not there then the job would not run. The 10 minute timeout is to avoid the site becoming a black hole.

L Betev (ALICE) [LB]: ALICE uses a 3 minutes timeout. It tries 6 times to contact the agent to download a payload.

JT: You probably do not count this in your efficiency while sites do.

LB: ALICE do not count jobs if the payload did not download.

US: [On Slide 8] yes this relates to LSF.

?: How do you get the time – from time of submission to the job actually running?

PC: We start when the job reaches the WN.

?: Then miss many steps such as sandbox creation.

PC: The time that the WMS wrapper takes before executing the payload is included after the job starts.

JT: To get the figures you can compare time stamps of the batch system and

?: First you need to get data from SE to local box. Slide 9 says 20GB data read – which is this.

LB: ALICE jobs are user jobs only.

It was asked why slide 11 shows a marked improvement in efficiency. Ulrich explained that this is the result of small statistics and may not be the actual trend.

Tony Cass [TC]: Users will not access data from tape, but what we are seeing is not this case. When will computing models be followed or understood to be wrong?



PC: Users are testing things where the samples are not on disk. How do you force people to use specific data sets?

TC: Looking at the WLCG accounting, this low efficiency is seen at many sites not just CERN, so is this [tape access] a reason?

JT: I have not seen non-grid job for few months.

JG: Do the experiments want to respond to Ulrich's observations?

LB: ALICE are happy with the work done with ALICE input, it has helped us to understand the batch systems. Already we have measures in the job wrappers and jobs kill themselves (but this is not always successful so there remains a need) to protect the batch systems.

JG: There is a lot of information so we need to decide how to manage the logs.

US: I would go for something simple – text based – in relative time

JG: Where do you keep the logs? Stdout or where?

Ian Bird [IB]: What's in the job wrapper now. What is the overhead due to this testing?

Markus Schulz [MS]: The overhead will reduce [compared to previous attempt where the logging led to R-GMA problems which overloaded sites] as we do not need to use RGMA anymore. We can enable the code.

PC: Can we know from job when it starts? Internally the job wrapper knows nothing about VO work.

JG: What do you do with the information.

PC: We have information in the DIRAC database that can be put in another [more accessible] form

MS: We can measure from the first wrapper to user wrapper time.

US: This depends on the experiment framework

JT: Using a messaging system is the way to go.

[Action 0803-1 Ulrich should consult appropriate people – including Markus and those in the experiments responsible for their framework – to progress understanding of the CPU efficiencies.](#)

3. Tape efficiency (Tim Bell)

It was noted that the CMS efficiency changed due to smaller files specially used for CCRC'08.

LB: At CERN, because of local people we can have direct interaction. We have the concept of a pre-staged data set. We make a bulk submission from the UI directly to CASTOR, it is a backdoor if you like and not applicable at other sites.

TC: The answer has always been that data to be analysed by users is on disk

LB: But this suggests that the disk is the same size as tape. How do users analyse data from 3 years ago?



TC: Then you need to understand how often this will happen – what is the requirement?

?: This is for the experiments to manage

JT: In constructing the MoU I see disk and tape are identical out to 2010.

LB: We will have more data on tape than disk. It will have to be staged from tape.

JG: How you do that is down to the experiment via queuing frameworks etc.

On the LHCb figures presented on Slide 14

PC: This looks like a mistake as files in CCRC should have been on disk. I've no idea what this was.

Kors Bos [KB]: Can you discriminate users – we [ATLAS] would like to understand “the others” tape files reserved for specific tasks

TB: For “read” requests go to a queue like for everyone else. We need to understand what users are doing and work with them. We can not build the service to stop access... with tape queues is there a way to get production users to the front? This is something the CASTOR developers are investigating.

[0803-2 Action on experiments to look at their user activities reading tape. Come back in one month's time and explain findings.](#)

KB: This is primarily a T0 concern, but we also need the plots from the T1s.

JG: Yesterday [management board] there was a requirement put on T1s to produce 4 simple metrics.

IB: This is an issue that can not take long to resolve – if we need to buy more drives then we need to do it now.

KB: This is also true for T1s.

IB: The original purchase was based on targets.

4. The SRM MoU (Flavia Donno)

On use of VOMS proxies JG noted that dCache does accept VOMS proxies but does not enforce them. Claudio Grandi (CG) added that VOMS awareness is for the SRM. But the control to the actual storage should be VOMS aware. At the moment all files are saved as production manager. Directories need to be world writable and readable. The backend needs to be VOMS aware.

JG: Do all the SRMs do mapping into virtual IDs?

FD: Well DPM does it.

TC: It is in the plan for CASTOR ...

JT: The experiment missed here is ALICE. They use ROOT and will face similar problems. This worries me. CMS is also looking at xrootd.

FD: The dCache developers are thinking about re-implementing their interface with xrootd.



JT: If it is used for the access protocol and you have gone through this exercise but can't read any tape file then

IB: gsi is available

FD: If it is enforced then xrootd needs to come with VOMS roles

JG: Can we do all this in time?

FD: We are waiting for the developers to respond

LB: Can we use the security questionnaire to make the descriptions fairly uniform?

ML: We already have some issues where we can agree – evident that VOs need to document these – but the others are not so obvious.

ML: LHCb input will already cover 80% of use cases.

5. Pilot Jobs (Maarten Litmaath/John Gordon)

Maarten gave an update on the group doing the experiment framework reviews. The mission/mandate is now agreed: review security issues; define a minimum set of security requirements; advise on improvements; report to GDB & MB in coming months. The group has had two teleconferences since January. The experiments are to produce a document about their system and a security questionnaire is being discussed.

John gave an update on the situation for reviewing the experiment frameworks from a security perspective. Glexec testing is just starting and there is a new version of the policy document dealing with multiple

Fabio: I asked in the last meeting if can have a pointer to the official version of glexec

JG: I got one version recommended from the glexec people and will send it to the GDB list.

[Action: John to forward information about the official version of glexec to the GDB list](#)

LUNCH

6. Job priorities (Andrea Sciaba)

The current timescale for this to reach production is 2-3 months.

CG: The timescale is based on the need to do what?

AS: ??

JT: The TCG has already said yes. The only thing that seems to be missing is documentation



AS: The documents will be in the YAIM pages.

MS: Currently, these tests been done outside the normal release structure. This also affects the YAIM core... try to release updates as infrequently as possible. In the configuration we have Maui/PBS available. We also need this for LSF and Condor.

PC: How does this mechanism allow setting of priorities? It is fairshare not priorities. For a group to get higher priority, how do you handle that?

JT: The pilot was done for ATLAS – they have the validation use-case you mention. 24hrs but use 10% of ATLAS allocation and want it now. This is a great way to submit high priority work, it will kill share for everyone.

PC: How does the VO manager know what to set?

JT: The VO can not directly configure this as it is set at the sites

CG: Batch system configuration does not provide more than you could do locally. A given group is assigned to a configuration.

MS: This is not priority but share

PC: [To ATLAS] Does it address the use case you mentioned?

KB: We have been using it continuously and it works

JG: It is mapping on to a group that has some attribute in the batch system (like a queue or)

JG: I have seen several LCG-ROLLOUT requests recently saying “I want my site to be for just production users”. Is this the solution for that?

JT/AS: It can be configured based on FQANs

MS: Do not repeat the hand waving rollout of last year (the quick fix with deny tags etc.) which ended up with large sites not being usable.

JG: We are not pushing

JT: You can do it by hand, but the second part is, I don't recommend it.

7. Monitoring (James Casey)

On slide 17 – the scale is the time during the day.

KB: Are all these tests done locally?

JC: Yes. Also, we also did some tests with SAM (not mentioned here) – still some problems getting the environment back

JT: Based on clients around the grid?

JC: Some way to test the environment when a new service starts, clients can also failover based on provided lists.

JT: Is it possible to have several networks of these? (slide 12)

JC: It has things built in like bridges. You might say connect to this broker network... and this is transparent to the client. Run one in the site, one regionally and this sends a subset up to the international level. It offers a way to scale.



JG: Have you done any modelling of rates?

JC: Yes, but in this model 1-tuple does not have to be 1 message as in the R-GMA case. Using this approach we can reduce messages

Simone Campana (SC): Imagine I want to use this for log files from jobs. Can one adapt it? Will it scale to 100k jobs per day.

JC: That depends on the size of the logs. You might want your own network.

Slide 18 – msg-sql is what rgma called an archiver.

JT: What is the MON box?

JC: It talks between the site and grid.

JT: MON box is a collector then

Helene Cordier (HC): How this can be done for ROCs? From EGEE-III there will be a moving towards this sort of ROC-site monitoring base. How will this be covered by SA1 tasks?

Maite B: This was mentioned very briefly at yesterday's ROC manager's meeting. The automation group in SA1 will translate the plan into different actors. It will affect many tools.

IB: This was explicitly in the description of work for EGEE-III.

MB: The mandate and members will be discussed at the next ROC manager's meeting.

Slide 20 – experiment specific tests –

JC: That [test] data has to go somewhere. I suggest CERN run these boxes – or those that actively run tests.

JG: The advantage for a ROC running in this way is that you may not get a slow migration. Want to have plug in for SAM tests .. run as Nagios and publish centrally. Having got that there will not be such a push for distributing the SAM database.

Slide 23

IB: This is a big step forward. If it works at NIKHEF and RAL what next?

JC: A lot of people get stuck at the level of the Nagios installation. We can support Nagios and Lemon and have a neutral xml dump

IB: So once validated do we mandate this at all T1s.

JC:Fermi already do this into their own framework.

IB: It is something to drive at the WLCG workshop.

JT: Same question different context. How big does a site have to be before we say it is worth the effort?

JG: One may expect a T2 federation to have some consumer for this information

IB: We want to get to the stage that sites react before COD tickets are issued.

JG: T1s are starting from the position of needing to have a monitoring system



IB: To be a reliable site, fabric monitoring is needed

JG: We can not mandate a fabric management system...

IB: No, but we can mandate on availability. What is the difference to running certain middleware?

MS: Functionality of middleware is for users. Leverage where availability is low.

JG: In your keynote speech then you can say that as WLCG we expect in order to deliver reliability that fabric monitoring plus a Grid service plugin/consumer is needed.

JG: What is the timescale?

IB: It would be good to report for the April GDB and then for the workshop

JT: I don't think there should be YAIM for Nagios.

MS: To make something mandatory then some deployment help is needed.

Slide 26:

MS: So, you recommend SAM CE job wrapper tests would then go through the L&B?

JC: Yes.

MJ: Is this a problem for non-RB submitted jobs?

ML: It is reported by the submission tools of the experiments.

IB: Can you publish into L&B from anywhere other than RB.

MS: L&B never tested for messaging on that scale.

SC: WNs already send messages. For those not using Condor-G or WMS there will still be some way using the normal L&B API.

Slide-31

JT: Nagios tells him there is a problem. Site admin clicks on acknowledge box in Nagios. The site admin does not want to have to look at multiple systems.

JC: Measuring metrics from VOs and presenting the data in a way that sites can get a feel for it is needed. Use of a messaging system is a requirement. Service maps will become the base for presentation.

JT: The critical thing is to also see the VO stuff in Nagios.

8. CCRC – How did it go? (Jamie Shiers)

Luca dell'Agnello: The schedule changed often

JS: Perhaps in the Monday meeting we should mention updates to the schedule.

JT: Were there any VO box problems in CCRC?

JS: ALICE – 5 sites – export 250MB/s well above target but 2 weeks for first site and then each site another week.

JT: SLAs are to make sure the box is up and running.

JS: Patricia is writing a document which better defines what is required.

IB: For May one issue may be whether resources are in place

JG: The maximum rates/processing will be bigger but already we have reached nominal rates.



LCG Grid Deployment Board Meeting



IB: a) will 2008 resources be in place and b) can experiments make use of them?

Harry Renshall (HR): Every site will have enough for 1-month running

JG: On critical service follow-up times and targets for the T0, the question is what do we do at the T1s? If I can see no data flowing from GridKa to CERN. Is there a method to follow? It is quite easy to have Nagios active monitoring and callout. For T1s – what is the process to get this implemented. Does it depend on part of the service, who asks for improved cover etc. We need feedback from T1s about calling people over night. Are T1s happy to have external triggers initiating call outs?

JT: For NIKHEF no. This will soften by 2009. Experience has shown that half of callouts are not a problem with the site.

JG: Is this a trust issue?

KB: I would be uncomfortable calling people out

JS: One issue here is that email/SMS is not synchronous

JT: If tests are accurate then we would link them to a call out otherwise email is preferred.

Other site manager's were asked:

Karlsruhe – we will consider

ASGC - Simon we will also consider it

Fabio – ??

RAL – JG: the site has started to discuss it so maybe

INFN – LdA: – we are also discussing this, regular callouts will be a problem.

JG: So should I summarise the position for the MB?

IB: Yes.

MD: Will all the T1s have the same naming conventions?

JG: I was not thinking about how this works but just want the principal in order

KB: We do need this with some urgency. Can we discuss at the next MB?

JG: I would be surprised if many would implement this now.

IB: Can we have the alarm email list covered over the weekend – we have been asking for this for over a year?

JT: Can email trigger Nagios

JC: Yes

IB: Write down what that clear route is from alarm through email and how it is implemented. Someone needs to respond within 1hr – 12hrs or whatever is agreed. Friday afternoon problems being picked up on Monday and fixed by Tuesday is not good.

9. VOM(R)S WG activity planning (Maria Dimou)



LCG Grid Deployment Board Meeting



GG: Is the membership of VOMRS regularly refreshed

MD: Yes – every time there is a change in the CERN HR database a resynchronisation is made.

JG: So this [the presentation] is the position of the VOMS admin team. Any discussion?

CG: I checked with them after seeing the slides. They can do this and estimate of order 3 weeks to implement the code changes. Of course on top of this are the requirements. Also in EGEEIII the developers have other duties. VOMS admin uses VOMS admin directly. So all the small security changes need to be made.

IB: 3 weeks or 3 months, I can't imagine that we would change this this year.

MD: Due to manpower limitations it would be good. It would not disturb the working system

CG: It can not happen before the 2.5 release (in August?).

IB: There are lots of conditions so to get access for developers (to HR database) will take time. This brings me to the point that FERMILAB have a contribution to make.

IB: The idea of having a simpler system and addressing operational availability is a good one.

MD: The way things go then will affect how difficult it is to get VOMRS support at FERMI.

IB: They are supposed to be supporting VOMRS.

JG: We used to have OSG-WLCG-EGEE meetings for this sort of discussion.

10. AOB

Will have CCRC on agenda again next time. Also we expect to hear back from Tier-1s

Next meeting 2nd April. Pre-GDB on 1st April

The meeting closed at 16:35.



LCG Grid Deployment Board Meeting



Actions:

Item No.	Description	Owner	Status
0602-4	Phrase the requirement on how to use policies in the WLMS	Cal Loomis	Open
0603-3	Follow up to ensure all sites in country are publishing accounting data or contact John Gordon with issues preventing this happening	Country representatives	Open
0604-6	Drive forward discussions on the VOMS and protocol issues	Ian Bird	Open
0605-3	Provide feedback (with reasons) to Dave Kelsey or Kors Bos on whether the security policy presented by Dave is acceptable.	All	Open
0605-4	Tier-1s to report back to GDB on what proportion of their current WLCG work is not reported/accounted within WLCG	Tier-1 managers	Open
0606-7	Take up and discuss technical solutions for removing shared credentials from the VO boxes	Markus Schulz	Open
0607-9	Ensure the default YAIM is properly configuring lcas lmaps for the sgm accounts (and that it works!)	Jeff Templon	Open
0609-1	Follow up on NDGF security policy position	Les Robertson	Open
0609-2	Look up statistics for automated on-call system and send information to GDB	Bruce Gibbard	Open
0609-6	Send storage type sampling script to John Gordon.	Jeff Templon	Open
0609-7	Move accounting to work in decimal units	Tier-1s/sites	Open
0610-5	Provide more detail on who is supposed to sign the site policy for each "organisation" mentioned in the security policy document	Dave Kelsey	Open
0610-6	Send the site operational procedures policy to the list again for comment ahead of approval and ensure lawyers at sites have a chance to review the document	Dave Kelsey	Open
0701-3	Check the CPU and storage accounting figures being published for the site	Sites	Open
0702-3	Discuss the future of a VOMRS-VOMS task force and consider possible mandates for the group	Dave Kelsey, Maria Dimou et. al.	Open
0702-4	Check Harry' resource tables and understand what they mean	Tier-1 sites	Open
0703-1	Check the Victoria MB time with Les Robertson and agree intention at the MB	John Gordon	Open
0703-2	Follow up on accounting policy documents	John Gordon	Open
0703-3	Send out a link to the latest patch	Jeff Templon	Open
0703-4	Follow up on the VOMS coordination group mandate	John Gordon	Open
0703-5	Refer Cal Loomis to Marian Dimou concerning the representation of smaller VO requirements in TCG discussions	John Gordon	Open
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



LCG Grid Deployment Board Meeting



Item No.	Description	Owner	Status
0705-1	Get feedback from Markus and Alessandra on previous feedback from sites on glxec.	John Gordon	Open
0706-1	Check use cases and VOMS need for failover with the developers and VOs	Maria Dimou	Open
0706-2	Provide description of implementation(s) of VOMS based ACLs and submit this to the experiments to confirm it satisfies their requirements.	Flavia Donno	Open
0706-3	Review the membership and approach of the Job Priorities Working Group	Erwin Laure	Open
0706-4	Nominate someone to join the grid services monitoring work	Oxana Smirnova	Open
0706-5	Follow up on how best to proceed with site-experiment negotiation on what VO SAM tests are to be monitored	John Gordon	Open
0706-6	Setup group to gather and prioritise GridView requirements	Ian Bird/ John Gordon	Open
0706-7	Follow up c) with Dave Kelsey	John Gordon	Open
0706-8	Raise glxec questions at the Stockholm operations workshop	Ian Bird	Open
0708-1	Provide feedback on the VO Operations policy	Reps/All	Open
0710-1	Comment on VO Operations Policy (final call next week); comment on Pilot Jobs Policy (v0.3)	All	Open
0710-2	Seek better definitions of VO roles – such as VO manager, VO operator etc – as they relate to policies.	Dave Kelsey	Open
0710-3	Circulate more requirements/issues information to the VOMS attributes group	??	
0710-4	Follow up on Markus's comment about glxec being used in OSG already and how experiences might be shared.	John Gordon	Open
0710-5	Send statement to MB regarding pilot jobs and glxec. Request MB to consider and forward to CB for comment.	John Gordon	Open
0711-1	Query the GDB list about member feeling for holding the April 2008 GDB in Taipei	John Gordon	
0711-2	Take advice on who to ask (JSPG) about VOMS requirements	John Gordon	
0711-3	Get glxec-on-WNs field tested by some sysadmins	John White/ John Gordon	Progress
0711-4	Pass on the issue of proxies being stored all over the place to the middleware security group	John Gordon	Progress
0711-5	Put together an experiment frameworks review team	JG,IB, DK and BJ	Progress
0711-6	Ask/inform/request sites about testing glxec with the various batch systems	John Gordon	
0711-7	Talk to David Salomoni about common batch system tests	John Gordon	
0711-8	Compile summary of pilot jobs/glxec discussion for MB	John Gordon	Closed
0711-9	Confirm resources available for the CCRC as given in Harry's talk (November GDB)	Country reps	
0712-1	Review Jeff Templon's concern about purchased resource figures ending up in resource spreadsheet but not allocated resources.	Harry Renshall	
0712-2	Gather more input before making a decision on April GDB	John Gordon	
0712-3	Follow up concerns about site-GGUS improvement with ROC managers. What is the process for using savannah vs GGUS tickets and finding out their status	John Gordon	
0712-4	Compile and report on list of SAM/monitoring priorities at January meeting	Piotr	
0712-5	Investigate holding the June GDB in Barcelona in conjunction with OGF	John Gordon	



LCG Grid Deployment Board Meeting



Item No.	Description	Owner	Status
0712-6	Speak with experiment reps about interfacing glexec to the experiment frameworks and establishing an architects group [nb. JT volunteers to be part of team]	John Gordon	
0712-7	Schedule follow-up (GDB) discussion on experiment user training on the topic of local access to data on SEs [Are there common solutions that can be shared?]	John Gordon	
0803-1	Consult appropriate people – include Markus and those in experiments responsible for the frameworks – to progress understanding of the CPU efficiencies	Ulrich	
0803-2	Look at user activities reading tape. Come back to GDB (in 1 month) to explain findings	Experiments	
0803-3	Forward information to the GDB list about what is to be considered the official version of glexec	John Gordon	

List of Attendees

X means attended

V means attended via VRVS

Country	Member	Present?	Deputy or Technical Assistant	Present
Austria	Dietmar Kuhn	X		
Canada	Reda Tafirout		Mike Vetterli	
Czech Republic	Milos Lokajicek			
Denmark	John Renner Hansen		Anders Waananen	
Finland	Klaus Lindberg		Jukka Klem	X
France	Fabio Hernandez	X	Dominique Boutigny	
Germany	Klaus-Peter Mickel		Holger Marten, Jos van Wezel	X
Hungary	Gyorgy Vesztergombi	X	Dezso Horvath	
India	P.S Dhekne			
Israel	Lorne Levinson			
Italy	Mirco Mazzucato		Luciano Gaido	
Japan	Hiroshi Sakamoto	X	Tatsuo Kawamoto	
Netherlands	Jeff Templon	X	Ron Trompert	
Norway	Jacko Koster		Farid Ould-Saada	
Pakistan	Hafeez Hoorani			
Poland	Ryszard Gokieli	V	Jan Krolkowski	
Portugal	Gaspar Barreira		Jorge Gomes	
Romania	Mihnea Dulea			
Russia	Alexander Kryukov		Vladimir Korenkov	
Spain	Jose Hernandez	V	Xavi Espinal	
Sweden	Leif Nixon		Tord Ekelof	



LCG Grid Deployment Board Meeting



Switzerland	Christoph Grab	V	Allan Clark, Marie-Christine Sawley	
Taiwan	Simon Lin	X	Di Qing	
United Kingdom	Jeremy Coles	X	John Gordon	
United States	Ruth Pordes		Michael Ernst	X
CERN	Tony Cass	X		
ALICE	Alberto Masoni	X	Yves Schutz	
	Federico Carminati			
ATLAS	Kors Bos	X	Stephen Gowdy	X
	Dario Barberis			
CMS	Matthias Kasemann	X	Patricia McBride	
LHCb	Ricardo Graciani		Andrei Tsaregorodstev	
	Nick Brook			
Project Leader	Les Robertson	X		
GDB Chair	John Gordon	X		
GDB Secretary	Jeremy Coles	X		
Grid Deployment Mgr	Ian Bird	X	Markus Schulz	X
Fabric Manager	Bernd Panzer			
Application Manager	Pere Mato Vila			
Security WG	David Kelsey			
Quattor WG	Michel Jouvin	X		
Networking WG	David Foster			
Planning Officer	Alberto Aimar	X/V		

Present:

J Coles (UK)
 J. Klem (Helsinki HIP)
 K-P Mickel (GridKa, Germany)
 J Templon (NIKHEF)
 I Bird (CERN)
 P Charpentier (CERN/LHCb)
 M. Girone (CERN)
 J. Gordon (RAL)
 T. Bell (CERN)
 H-C Lee (ASGC)
 L. Betev (CERN/ALICE)
 A. Sciaba (CERN)
 U. Schwickerath (CERN)
 M. Barroso (CERN)
 S. Foffano (CERN)
 E. Laure (CERN/EGEE)
 A. Aimar (CERN)
 I. Ueda (Tokyo)
 M. Litmaath (CERN)
 R. Gokieli (Poland)



LCG Grid Deployment Board Meeting



Kors Bos (CERN/ATLAS)
D. Kuhn (Austra)
M. Lokajiceh (Prague)
C. Grandi (INFN)
G. Veszterpondi (Budapest)
G Qin (ASGC)
M. Jouvin (LAL/GRIF)
T. Cass (CERN)
G. Carlino (INFN/Napoli)
L dell'Ágnello (INFN/CNAF)
H. Renshall (CERN)
M. Schulz (CERN)
S. Lin (ASGC)
M. Kasemann (CMS)

*Hiroshi Sakamoto
Frederique Chollet
Piotr Nyczyk
Kai Neuffer
Leif Nixon
Xavier Espinal
Matthias Kasemann
Andreas Heiss
Jose Hernandez
Steve Traylen
Fabio Hernandez
Ricardo Silva
Marek – EVO support
Davide Salomoni*

*PM-
Ioannis Liabotis
Jules Wolfrat
Paolo Veronesi
Marc Rodriguez Esp*

Others present at CERN

F. Chollet – IN2P3
I. Ueda – Tokyo
Jamie Shiers – CERN
Nick Thackray – CERN
John Shade – CERN
James Casey – CERN
Gonzalo Merino – PIC
Patricia McBride – CMS/FNAL
Dietmar Kuhn – Innsbruck



LCG Grid Deployment Board Meeting



Harry Renshall – CERN
P. Charpentier – CERN/CMS
Sue Foffano – CERN
Gilbert Poulard – ATLAS
Claudio Grandi – INFN
David Salomoni – INFN
Luca dell’Agnello – INFN
Stephen Burke – RAL
Oliver Keeble – CERN
Ian Neilson – CERN
Kors Bos - ATLAS
Jason Shih - ASGC

Others on VRVS:

Duncan Rand - UK