



Draft (V1)

Amendments history:

<i>Name</i>	<i>Area</i>	<i>Date</i>

Minutes of the meeting

CERN, 12th January 2005

Agenda: <http://agenda.cern.ch/fullAgenda.php?ida=a045318>

Minutes: Jeremy Coles

Attendees: See appendix

1. Introduction and actions from October meeting (Kors Bos)

KB requested input for future meetings and feedback on suggested dates and locations. This meeting is focussed on “mass storage” systems.

LCG comprehensive review

<http://agenda.cern.ch/askArchive.php?base=agenda&categ=a045318&id=a045318s0t7/transparencies>

KB summarised the feedback from the comprehensive review. Fabric and Network were seen to be okay. Deployment and regional centres have made good progress but not yet reached “production quality”. The Application Areas need to interact more. Management and planning have a manpower issue which is being addressed.

KB went on to show D0 MC Production on LCG2 information (decoupled DO from LCG2) – see slides. The MC production sees 98-99% efficiency over 7 sites but requires a person dedicated to the task.

JG: What are the 1% of errors classed as “file not found”?

KB: Unknown at present, there is a ticket in the system to trace this problem (problems escalated). Lyon and FZK are now fine despite showing a problem at the time of the test. It was noted that often all jobs to a site fail or they all run! Wuppertal was probably down hence no jobs



LCG Grid Deployment Board Meeting



MM: How does this fit with the numbers submitted to the LHCC review (~60% efficiency) – are they misleading?

KB: These numbers were reported.

IB: The latest figures show for CMS and ATLAS better than 90% efficiency (stronger follow up on sites and stricter control of sites used). Many sites are not well managed and these either need to be improved or not used.

It was noted that these numbers are global efficiency for MC and therefore less sensitive to Grid problems.

KB: This is standard operation and NIKEF have a lot of experience running such MC so it is an example of what can be achieved.

Future meetings – see slides and discussion later in minutes.

Service Challenge meetings

KB reminded members that the Service Challenge meetings are for CERN people to meet and discuss plans with specific Tier-1s, but the meetings are open to all and it is useful for everyone who attends – these are dedicated technical meetings.

2. New GDB Working Group

<http://agenda.cern.ch/askArchive.php?base=agenda&categ=a045318&id=a045318s0t1/transparencies>

KB mentioned that there is an outstanding action to update the web-site and this will be done soon. The new site will link to the pages of the new working groups. A group is now in place for Networking and the first meeting is next week in Amsterdam.

One objective for this group will be to review the estimated aggregate data rates for each T1 (assuming all T1s same). Double these rates will be required to allow such a centre to be down for 2 days. This new working group will examine data flow estimates and usage models. In parallel with this work we need to agree who pays for the networking.

LR: CERN pays a contribution to GEANT via Switch. For the US connection CERN again pays only a contribution.

SL: Please be aware that the recent Tsunami has cut the Asia fibre in several places. For a 10GB Asia connection it is no longer possible to rely on the previous connection. Further details on this will be made available at the Amsterdam meeting next week.

3. Mass storage systems

<http://agenda.cern.ch/askArchive.php?base=agenda&categ=a045318&id=a045318s0t19/transparencies>



LCG Grid Deployment Board Meeting



The main requirement is for each centre to provide an SRM interface to their storage system.

MM: Is dCache in LCG?

IB: There is an instance at CERN but there is a list of issues still to be gone through.

5 dual Itanium machines were used for the US tests.

KB: For the first tests then we will not need all 10 machines in place?

IB: Correct – the test will be restricted by what the network can handle using multiple streams.

JG: I am concerned that the plan sees the challenge infrastructure move straight across for production use. If we (UK) get a light path for the challenge it is unlikely we will get another for production.

IB: Production should use the best connections that can be achieved.

JG: The challenges are to ensure we are on target. Light paths need to be scheduled. The UK is unlikely to have production light paths until SuperJanet 5. Use of existing networks at 95% will cause others to become unhappy and today's disk servers will be scrapped by 2007!

DP: The focus on rates is useful but not the most important figure. The real value comes in integrating many pieces to work as an ensemble. Good network figures are often just for a specific case.

Milestones for Mass Storage Deployment (Les Robertston)

<http://agenda.cern.ch/askArchive.php?base=agenda&categ=a045318&id=a045318s0t6/transparencies>

The data rates are to be confirmed after the LHCC review next week. The computing model are to be confirmed by the review. See the presentation for details.

A short-term goal is 1 month sustained service in July.

JG: The step to 5GBytes/s is more of a problem than that to 60 MBytes/s

The long term average rate for 2008 should be achieved in 2006/2007.

JG: Network equipment purchase is similar to CPUs – spending later gets more for the same money. It will be interesting to see the result of the NREN meeting; Particle Physics has not saturated previous networks. A credible claim will be supported if we start using the rates now.

LR: Are the experiment rates realistic. If so what is the ramp up model based on budget spending plans (which argues for doing things later); note that we want to test to ensure confidence as soon as possible. I urge for caution as this has not been tried and CERN tests show we can expect a lot of problems that will need to be resolved. The need to test



LCG Grid Deployment Board Meeting



and ramp up must be seen and understood by NRENs. What is the slowest ramp up we can tolerate?

FC: On the basis that we need several years to realise our stepped goals 25%; 100% 200% of data rates, the ramp up is actually quite (too?) slow.

KB: It would be good to understand from the experiments what is needed. The whole chain of reconstruction takes a long time. If we test 1GB then 10GB operation for Amsterdam it is likely we put the 1GB link into production. Our concern about usage of this test system as a production system for the data challenges is how much will it really be used?

There is a requirement for production which is negligible vs real data taking and reconstruction. The tests need a large bandwidth to test the computing models only.

DB: We have no number in mind for use in 2005 and the first half of 2006 but network requirements are low. We will be transferring data occasionally.

FC: We have not produced as much data as will be seen in a data taking year. We could send around the same data several times. Generally though the average load on the network will be low. But, we need to test the software (speed) .

KB: Setup small systems for production and a big system for testing.

IB: Even if there are not high-bandwidth requirements before data taking, we need to maintain the infrastructure to prevent it decaying

LR: Who is responsible for the overall service? Do we need a background generator to keep the load on network and servers. When experiments test we could run down this background load.

JG: Do we need that. We can produce internal loads to test robots etc.

LR: Who sets this up and who funds it?

JG: The service challenges provide knowledge and that will not decay.

We have seen systems decaying. We need background to maintain loads.

JG: The service challenges are to identify problems and then go away and solve them. The model was not to move into a production service.

DP: You can run the service challenge on production hardware. It is acceptable to have it alongside production, expand a centre and learn a lot in doing this. Each site can decide its own way forward.



LCG Grid Deployment Board Meeting



LR: The service challenges need to be long enough to be sure problems are resolved. We have assigned a long period in 2005 to fix early problems.

JG: Service challenges will guarantee traffic but the experiments will not be producing it.

SRM (Jean-Philippe Baud) – see slides on agenda page

The plan is to use SRM for the service challenges. This talk was a summary of the status. An initial comment that it has been difficult to get community replies on this topic – the GDB will be cc'ed on future requests so members can follow them up.

DP: Open Science Grid (OSG) is also doing SRM testing using dCache and ?? Implementations. Interoperability tests would be useful – we will try to open a channel to take this forward.

JG: An early series of meetings between sites on this topic seems to have decayed.

KB: Jamie will try to reinstate these meetings.

JG: How do you see SRM3 being used? Would you see LCG as defining a profile as a set of options you would expect sites to use?

JPB: The new version of gLite I/O talks to SRM implementation depending on what features are supported. You choose a site which offers you the required interface possibilities. Similarly Globus based reservation can be used, probably at the RB step.

JG: How will this information be published?

JPB: Via MDS or R-GMA which the RB can interrogate. Or as proposed in SRM3, servers can publish directly, but we see many scalability problems with MDS and R-GMA and the same is likely here.

DL: Test suites are testing setups. Lyon implemented its own SRM but would be interested in a common test suite.

JPB: The CVS repository will hold out test suite. Sites can download from next week and try to use it.

Mass storage at CERN (Bernd Panzer) – see slides on agenda page

No questions were asked.

Lunch



LCG Grid Deployment Board Meeting



Mass Storage at GridKa (Doris Ressmann) – see slides on agenda page

No questions were asked.

Mass storage at SARA (Peter Michielse (NCF)) – see slides on agenda page

No questions were asked.

Mass Storage at CCIN2P3 (Lionel Schwarz) – see slides on agenda page

No questions were asked.

Mass Storage at RAL (Jeremy Coles) – see slides on agenda page

KB: Are input and output channels separated.

JG: They are.

DP: Like Karlsruhe the read and write pools are separated. It is read once from the read pools. This functionality is obtained by administering in dCache.

JC: Is the CERN Disk Pool Manager (DPM) SRM to be available for testing soon?

IB: DPM is now available as an Alpha release

JC: Is there any indication about whether dCache will be made open source?

Various: This would be useful and we would like it to be open source. It is not clear at the moment if this will happen. **The GDB can request clarification from DESY?**

Mass storage at Fermi (Don Petravick) – see slides on agenda page

Q: For how long will the ingested data be retained?

DP: For 6 months – we do not need to retain data.

DP: [VRVS poor link] It may be useful to apply priorities as we have a hard time discriminating production traffic and service challenge traffic. It would be good to meet with **xx** people to understand roles properly, this is a good area for upfront planning.

IB: The integration work shown needs to be provided in a similar report by each Tier-1 and CERN. The whole system needs to work not just the parts. Therefore getting together with people developing is important (production users vs anonymous users). There is a lot of work here on monitoring and understanding the system; there is scope to leverage and put some of this work into the **Radian** software.



LCG Grid Deployment Board Meeting



DP: Fermilab plots were produced by a separate module. The work required 1 FTE. These are not simple systems with 32 pulls on 22 nodes. No clear acknowledgement about micro-storage management issues.

JG: At a higher level ... what is the status of the software for collections?

Don: Can use a bunch of ASCII labels to segregate. Tape primitives on tape map to dCache but grouping only bound by name space. Would be useful to have a workshop

IB: We have to arrange a workshop for this discussion.

AOB

KB: In the D0 presentation I alluded to accounting. In spite of the previous urgency a lot of sites still do not publish the information.

JG: We agreed at the last meeting to push for 2.3 and APEL (accounting) is built into it (whereas for 2.2 it was an add-on module).

IB: Deploying the latest release was raised at the last GDA. We are still asking for numbers from last year.

KB: The software is available but it seems it isn't installed?

IB: In 2.3 it is installed by default.

JG: The instructions need to be followed to install it properly – sites are not always doing this and many of the questions we receive reflect this position.

IB: We asked for log files several times

LR: Even if it is in 2.3 this probably means March before most sites publish; we may neglect numbers from last year but risk the funding agencies reaction if we do not get figures. We can not wait for the new version to be installed everywhere.

JG: We need to go to each Tier-1 and ask for usage records and remind all sites to keep logs.

KB: Of the Tier-1s present can logs be provided? Accounting was to be in place by January!

KB: Amsterdam have provided the information

KPM: Karlsruhe will provide logs

DL: Lyon – will follow up



LCG Grid Deployment Board Meeting



MM: CNAF – I will ask

RS: Triumf - will follow up

??: Fermilab – don't know the accounting schema. Will talk to Ruth

IB: Not a Fermilab specific issue as we need to know the LHC use of Grid3 for 2004 (Rob Gardner has suggested the information is available).

JG: Current problem is merger of PBS and gatekeeper logs and so there is a potential gap if PBS logs have not been kept. For these cases is it possible that we provide a spreadsheet summary?

LR: We had this discussion about all Grid vs LHC experiments usage. It is important to get a high-level (overall) view. Each Tier-1 should respond within 2 weeks. We need to know the LHC computing usage for 2004 for each Tier-1 and how this information can be provided - of course it is better if we can feed information directly into the accounting system .

KB: JG action to remind sites.

Action – all countries to provide accounting information (by setting up APEL where possible).

SL: Could JG send instructions on details of what log files should be forwarded.

JG: Yes

Future meetings:

SC meetings – at least technical people from each of the 5 Tier-1s should be present. Be aware that around 15th and 16th March there will be meetings in Lyon.

There will be no April GDB meeting.

SL (plans for workshop <http://www.twgrid.org/event/isgc2005/>): We want to schedule most international aspects towards the end of the week 25th-29th April 2005.

December 21st will be the last CERN working day for 2005 and therefore not a good day for a GDB meeting!

Action All – review second half of 2005 meeting dates and let KB know suggestions/issues.



LCG Grid Deployment Board Meeting



Actions:

Item No.	Description	Owner	Done
0412-1	Contact Dave Kant at RAL re input of NorduGrid accounting data	A NorduGrid representative	
0412-2	Distribute details of the CERN migration to SLC3 to the GDB mail list	Tony Cass	
0412-3	To discuss and finalise dates and locations of 2005 GDB meetings	Kors Bos	
0501-1	Standing action: Update the GDB web-pages	KB/JC	
0501-2	Send mail to remind sites of logs that need to be processed for accounting records to be backdated	JG/KB	
0501-3	Ensure Tier-1 past (for 2004) accounting data is available for processing	Tier-1 country representatives	
0501-4	Provide feedback to KB on proposed meeting dates for Q3 and Q4 2005	All	



LCG Grid Deployment Board Meeting



List of Attendees

X means attended
V means attended via VRVS

Country	Member		Deputy	
Austria	Dietmar Kuhn	X		
Canada	Randy Sobie	X	Robert McPherson	
Czech Republic	Milos Lokajicek	X		
Denmark	John Renner Hansen		Anders Waananen	V
Finland	Klaus Lindberg		Jukka Klem	X
France	Denis Linglin	X	Fabio Hernandez	
Germany	Klaus-Peter Mickel	X	Holger Marten	
Hungary	Gyorgy Vesztergombi	X	Dezso Horvath	
India	P.S Dhekne			
Israel	Lorne Levinson			
Italy	Mirco Mazzucato	X	Luciano Gaido	
Japan	Hiroshi Sakamoto		Tatsuo Kawamoto	
Netherlands	Peter Michielse	X	Arjen Van Rijn	
Norway	Peter Kongshaug		Farid Ould-Saada	
Pakistan	Hafeez Hoorani			
Poland	Michal Turala		Jan Krolkowski	
			Marcin Wolter	X
Portugal	Gaspar Barreira		Jorge Gomes	
Russia	Slava Ilyin		V. Korenkov	
Spain	Manuel Delfino	V	Andreu Pacheco	V
Sweden	Niclas Andersson		Tord Ekelof	X
Switzerland	Christoph Grab		Allan Clark	
			Marie-Christine Sawley	
Taiwan	Simon Lin	X	Di Qing	
United Kingdom	John Gordon	X	Jeremy Coles	
United States	Vicky White		Bruce Gibbard	V
CERN	Tony Cass	X		
ALICE	Alberto Masoni	X	Yves Schutz	
	Federico Carminati	X		
ATLAS	Gilbert Poulard	X	Laura Perini	X
	Dario Barberis	X		
CMS	Tony Wildish			
	David Stickland			



LCG Grid Deployment Board Meeting



Country	Member		Deputy	
LHCb	Ricardo Graciani	V	Andrei Tsaregorodstev	<input type="checkbox"/>
	Nick Brook	X		<input type="checkbox"/>
Project Leader	Les Robertson	X		<input type="checkbox"/>
GDB Chair	Kors Bos	X		<input type="checkbox"/>
GDB Secretary	Jeremy Coles	X		<input type="checkbox"/>
Grid Deployment Mgr	Ian Bird	X	Markus Schulz	<input type="checkbox"/>
Fabric Manager	Bernd Panzer	X		<input type="checkbox"/>
Application Manager	Torre Wenaus		Oxana Smirnova	<input type="checkbox"/>
Communication Systems Mgr	David Foster			<input type="checkbox"/>
SC2 Chair	Matthias Kasemann			<input type="checkbox"/>
Security WG	David Kelsey	X		<input type="checkbox"/>
Quattor WG	Charles Loomis			<input type="checkbox"/>
Networking WG	David Foster			<input type="checkbox"/>

Jamie Shiers attended for Service Challenges.

The following also attended via VRVS at various points during the day:

Bonny Strong

Paul Gelissen

Jaroslva Kultan

Piot Nyczyk

Robert Macek

Alexander Kaukher

Toni Gape

Juan Jose Ortega

Willem Van Leeuon

Orhan Cakir

Doris Ressimann

Ron Trompert

Peter Elmer

Min Tsai

Jos Van Wezel

Canzio Torelli

Luis Zorzano Martinez

Tim Folkes

Sander Klous

Jules Wolfrat

Donald Petravick

Vova Moroz

Pilar Ramirez

Fred Dauger

Maurice Bounhuis

Krzysztof Piotrkowski