



Meeting Summary

July 2008 GDB meeting @CERN

Introduction (John Gordon)

August's GDB has been cancelled. September's will concentrate on Tier2 issues. Ideas gratefully received.

I asked for feedback on dates for 2009. Should we stay with the second Wednesday of the month or revert to the first?

In addition to the issues highlighted below we heard presentations from NorduGrid and OSG on their operations security and from James Casey on the distributed monitoring plans of EGEE which are starting to be implemented.

Meeting Agenda <http://indico.cern.ch/conferenceDisplay.py?confId=20231>

Benchmarking (Helge Meinhard)

Helge repeated the proposal for benchmarking that he gave to the MB but with more detail. This has now been agreed by MB (any contradictions?) so we need to work on details. I am setting up a small group to work on this. I got no volunteers but Jeff suggested that all experiments should be involved with a matching number of sites and some of the original HEPiX group. This seems a good suggestion and I will progress it.

The group should at least look at:-

- agreeing the exact conditions under which the chosen benchmark will be run.
- producing an understandable paper for wider circulation;
- producing a proposal for evaluating existing capacity
- producing a proposal for translation of requirements

Agreement is required in time for input to the October C-RRB.

Reporting on Installed Capacity (Flavia Donno)

Flavia expanded on the technical details of her proposal to gather information on cpu and storage installed at sites for comparison with MoU pledges. She has a student working on collecting data and various people are developing new version of the storage information providers. The storage accounting portal is undergoing an overhaul and needs to work with OSG on publishing their data.

Issues outstanding include:-



- a) There seems to be no standard advice on how sites with clusters shared with non-LHC work should publish their capacity. Some publish everything on the grounds that it is all theoretically available so their installed capacity can greatly exceed their pledge. Others publish a VO's share so usage can exceed installed capacity.
- b) The MoU definition of 'installed capacity' includes disk that is in the machine room but has not been deployed/configured for any experiment. It may be allocated to one but may not be powered on. Such disk is not currently known to the Information Service. One proposal was to publish this capacity in a virtual storage space so that it could be included in any capacity measurement. I saw this as being difficult to maintain manually. Another proposal was to publish the 'total installed capacity' in a virtual storage space. This would only be used for the collection of management information. OSG had problems with this concept.

Flavia says she has all technical expertise required in her working group. I propose that they work up a proposal and bring it back to MB for approval.

Tier1 Readiness (John Gordon)

My summary of Tier1 readiness. Generally T1s feel fairly confident due to experiences in CCRC08 in May. Things will go wrong but in general service was recovered quickly. Still concerns over communications and middleware limitations.

Kors Bos took issue with this complacency. He stressed:-

- a) ATLAS recognized that they had only addressed half of the problem. When users run jobs in large numbers a different set of problems will be seen.
- b) Sites did not recognize the effects of even very short breaks in service on the experiments. It took a lot of manual work by the experiment to recover. In my opinion this is a limitation of the experiment frameworks. Grids are inherently unreliable and frameworks must take this into account. That said, little is going to change before data taking so sites should try to minimize the number of breaks as well as maximizing reliability. One part of the rise in reliability has been through better monitoring and call-out. The next step needs to be in resilience to reduce the number of breaks in service.

Tier2 Readiness (Mike Vetterli)

Mike identified a problem with communications, specifically between Tier2s in a different country from their associated Tier1. A concern was how user support would be done for end-user analysis.

Machine Readiness (Jamie Shiers)

Jamie reported that the date for first circulating beam was August 8th. The experiments expect to start taking data for calibration with single beams so everything should be in place before then. Tier1s would like a period of stability before beam but we are almost too late for this already.



LCG Grid Deployment Board Meeting



I would like the MB to track outstanding issues and apply pressure where necessary. In particular:-

- a) Storage Tokens to be defined
- b) Storage required for each token
- c) Complete list of minimum middleware releases required.
- d) Anything else?

Middleware (Markus Schulz)

Markus reported on forthcoming developments:-

- a) SCAS is still with the NIKHEF developers for 'deep certification'. This is the holding item on Multiuser Pilot Jobs for LHCb.
- b) The WMS/LB for gLite3.1/SL4 was released. This is recommended as being much better.
- c) Job Priorities were released but no-one seems to have used them yet. ATLAS are organizing further tests with some of their Tier2s.
- d) Within six months we should see a new version of FTS addressing: split of SRM; negotiations and gridFTP: improved throughput and logging: full VOMS support
- e) Ulrich Swickerath reported on his tests of the Cream CE. Its basic functionality was OK but it had a number of limitations and it failed to meet the target reliability.