

# London Tier 2 Technical Group Meeting

## Held at:

Imperial College  
15h00 9 June 2004

## Present

Owen Maroney (Imperial HEP) - Chair  
Barry MacEvoy (Imperial HEP) - Minutes  
David Colling (Imperial HEP)  
David McBride (LeSC)  
Keith Sephton (LeSC)  
Paul Kyberd (Brunel)  
Alex Martin (QM)  
Chris Williams (QM)

## Apologies for absence

UCL  
Royal Holloway  
Steven Newhouse (LeSC)

## Agenda

### 1. Preliminaries

PK arrived at 15h20 and gave apologies for his delay.

DC suggested that a mobile contact number list should be compiled to enable OM to contact members more easily.

**Action: All members to send mobile contact details to OM ASAP.**

### 2. Site status

#### LeSC

DM reported that he had now located the source rpms for LCG2. Alas, his cluster is running RH7.2 rather than RH7.3, and the rpms cannot be re-compiled. He is therefore using User-Mode-Linux (a SourceForge project) to host VMs on his machines. He has installed a UI and a CE and SE with host certificates. GridFTP and Globus are running. The aim is to set up a host of VMs. DM stated that overhead and responsiveness appeared to be acceptable, but he had not tested against any benchmarks as yet.

AM expressed concerns as to IO capabilities on the VM platform. It was pointed out that installation of LCG2 on these VMs would be a manual process and would not involve boot via PXE or floppy.

OM asked DM to investigate the current status of Sun Grid Engine integration work for information providers which is being carried out in Durham. OM queried whether tighter co-ordination between the groups was required.

**Action: DM to speak to Durham.**

## **Imperial**

BM reported that the local LCG2 installation is fully installed on the CMS farm and working correctly. There are presently 46 WN CPUs, which should expand to 50 CPUs within a few days.

DC asked how many jobs were being sent to IC. OM stated that there were normally over 30 jobs in the PBS “infinite” queue.

## **UCL**

OM reported that the HEP front-end is in the LCG Test Zone. Ben Waugh has been attempting to configure this front-end to talk to the existing HEP PBS farm. The problems are now mostly solved and it is a question of documenting this recipe for use by other users.

**Action: BW to document recipe.**

The UCL e-Science cluster has been partially installed via LCFGng with a CE, SE and 10 CPUs of WN with PBS. It has been suggested that a total of 60 CPUs of WN is a reasonable figure for the final total. UCL is presently waiting for host certificates before testing can take place.

## **Royal Holloway**

OM reported that a CE has been installed, but there are problems with Globus. There is also an unresolved PBS issue, probably related to configurational problems with mounting a raid disk.

PK asked (tangentially) if a Tier-2 centre should provide long term-mass storage. The consensus was that there was no official policy as yet.

Royal Holloway will follow the UCL recipe as provided by Ben to patch in a 170-node PBS farm.

## **Queen Mary**

CW described the architecture at Queen Mary. This comprises a PBS farm behind NAT on Fedora 2, with the CE and SE outside NAT. The inevitable PBS routing issues have now been resolved and CW is now configuring the WNs and verifying correct installation with a test script.

A query was made as to how the SE and WNs would communicate. OM pointed out that the SE should not have to nfs mount to the WN, and that communication would be via edg-replica-manager.

CW hoped that LCG resources (320 CPUs) would be available to the Grid within 10 working days.

There was a more general discussion as to the fate of the BaBar farm(s). Given that Manchester's enthusiasm for making their farm available to LCG appears to have waned, we would have to concoct a recipe (and do the work) ourselves. A consensus emerged that we should ignore this issue, at least for the moment.

## **Brunel**

PK has now installed a SE, CE and WN and will shortly carry out the post-installation steps. He complained about the documentation available for the installation process. It was suggested that he address his (well-founded) complaints to the LCG Rollout List.

**Action: PK to contact LCG Rollout List to complain about poor or non-existent documentation.**

PK will now carry out the installation and configuration of 30 WNs, before attempting a further 30 WNs on SCSI. According to OM, Steve Traylen has reported that this problem has now been solved for LCFGng.

## **3. News from LCG**

There will be an LCG2 administrators' course (funded by GridPP) on 19-21 July 2004. Course materials for later use will be recorded by NeSC. OM requested that each site manager consider sending at least one member of personnel to the course.

**Action: Site managers to choose candidates.**

It is likely that there will be a new release within the next 2 weeks. This will contain a number of bug fixes as well as the inclusion of the long-awaited SRM dCache (a generic interface to storage systems).

OM pointed out that a migration plan would be required for this transition to the new SE.

## **4. News from GridPP**

The LCG2 developers have now been persuaded of the genuine need to release source rpms.

OM had attended HepiX, where FNAL and CERN agreed to standardise their Linux on a re-compilation of Enterprise RedHat. There is a danger that this may reduce the drive towards making middleware system-agnostic.

## **5. AOB**

There was some discussion as to the role of the 2.5 FTE support posts in London. DC was keen to provide good coverage by institute, whereas OM favoured a scheme whereby individuals took responsibility for specific areas (for example storage, or networking).

DC stated that the tender for the front-end nodes will close at the end of week 1 of July. Resources should therefore be available by late July.

The meeting was closed at 16h55.