

Oxford

before I left there was still the problem of failures due to files being deleted before some action in the pilot happened. Site is now in downtime to upgrade condor. Hopefully the upgrade will fix the problem.

Glasgow OK

nodes seem to have worked fine for a week and the software has been installed. First prod jobs started to run last night. They are running both single core and multicore jobs at the same time since the UKI-SCOTGRID-GLASGOW_TEST is active and sending a number of tests. Plan is to add 96 extra cores.

Liverpool OK

is working since last week is still validating releases and got activated jobs last night, doesn't run a mixture of single and multi core jobs. Will add more core this week and add to single core production.

Lancaster OK

* We hope to maintain separate multi-core and single-core queues - in order to keep things sane in our shared cluster of multi-generation hardware this is a needed and hopefully atlas (and other users) won't mind splitting the single and multi job submissions between them.

* Stable number of multicore jobs due to the multicore SGE ability to weight by number of cores (i.e. more cores have preferential treatment) maybe worth documenting this feature. Matt will document this configuration in the wiki. The dedicated nodes in torque help with a 3rd of the throughput and with picking up after a drop in submitted jobs.

* "Plans to try a QMUL-like strategy of "dynamic" job reservation (QM run a cron that uses sge tools to reserve slots when it notices waiting multicore jobs)."

This has been put on the backburner as the shared cluster does not have advanced_reservation enabled, and we're not sure we want to change that.

Also Lancaster is jumping on the "more cores for multi-core" bandwagon, opening over 600 new job slots to multicore work today (with any luck). These are 16-core nodes so I'm hoping that helps with job tessellation (some how). My usual procedure to help multicore along is enable the multicore queue first, have a cup of tea, then enable the serial queue on the nodes. This gives a chance for multicore jobs to grab the slots and the factory to line up some more, and is another example of the usefulness of a separate queue.

Manchester OK

Looking at what happens when ATLAS submission drops <http://tinyurl.com/o4ojk96> Lancaster restarted immediately and Manchester took a bit longer to pick up but it probably had also less empty slots since lots of smaller VOs are filling the ATLAS gaps.

QMUL OK

was running fine but since last week it has stopped running due to a local user requiring resources to

finish a paper. QMUL added their configuration to the docs in the wiki
<https://www.gridpp.ac.uk/wiki/Category:Multicore>

Next?

RALPP, ECDF and Sheffield are all good candidates. They all have inactive multicore queues that were setup at some point in history.