

Information from Experiment Integration and Support teams at WLCGSCM of 2 August, from CERN group reports to an internal meeting of 4 August and from the WLCGRSM of 7 August. Updates have been made to the experiment plans at the SC Twiki (linked to Agenda) but not yet reflected to the site plans as some numbers and dates still need to be confirmed.

ALICE: Still have no end point at RAL nor NDGF. Had some srm_put problems at Sara. IN2P3 is giving preference to CMS transfers this (31 July to 6 August) week but ALICE will carry on transferring there. They reached an aggregate peak transfer rate of 150 MB/s but only for a short time. Overall their rate is well below the target of 300 MB/s but they intend to continue as long as possible until they reach it. They report poor feedback from their submitting of problems to GGUS and switched to trying to use internal site contacts. They will restart using GGUS to see if this can be improved. Corrections were made to the service challenge Twiki to reflect that only 50% (not 100%) of their network/reconstruction stress tests (running now) will be made on T1 sites, the rest being at T2 (or equivalent) sites.

ATLAS: Have had a post-mortem analysis of their July Tier0/Tier1 exercise. Overall it was good with no major issues for CERN but some problematic T1 sites. Corrections were made to the service challenge Twiki to reflect that the repeat June/July Tier 0 to Tier 1 export exercise is now scheduled for two to three weeks from the middle of September. This could be advanced if sites prefer as it would currently overlap with increased MonteCarlo production from ATLAS and the start of the CMS CSA06.

CMS: Their export exercise (the target is 500 MB/s for a week with a minimum threshold of 300 MB/s for 3 days) started badly following the CERN power cut on Saturday then an expired host certificate in FTS. They reached 350 MB/s Wednesday morning but dropped overnight for unknown reasons then recovered. From Thursday afternoon the CMS rate stabilised at over 300 MB/s only to drop again overnight and effectively stop on Friday morning. This was put down to an unexpected reply from a Phedex stager query triggering 80000 pre-stage requests. By Friday afternoon they had recovered to 350 MB/s and this was maintained over the weekend. They currently are not using CNAF and have low rates to RAL. They intend to continue these tests till the end of this week or if they see sufficient stability. Corrections were made to the service challenge Twiki to update the cpu and disk resources needed for CSA06, with bigger sites doing more and smaller sites less, though confirmation of the numbers is still needed for some sites.

LHCB: On Tuesday the castorsrm.cern.ch end point was changed to point to srm.cern.ch to fix a gfal problem seen by LHCB. Their important issues were their central job logger crashing and intermittent slow transfers to srm.cern.ch. The job logger problem was found to be a result of directories containing more than 32000 entries, not well handled by lcg-cp, which was fixed but IT remain worried about the scalability of a central gridftp server for job logs. Alternative solutions are being investigated. Meanwhile

LHCB have also reduced the number of individual files sent. The slow transfers also happen to a classic SE so are not thought to be a castor2 problem. For the moment LHCB have increased the timeout on these transfers of MonteCarlo files from worker nodes into CERN.