SC4 Weekly Report and upcoming activities   24 July 2006.

Information from Experiment Integration and Support teams at WLCGSCM of 19 July,
from internal computing meetings of LHCB, CMS and ATLAS,
from CERN group reports to an internal meeting of 21 July and from the WLCGRSM of
24 July.

ALICE: As of 19 July no T0 to T1 transfers were working. CNAF had been running well
but was turned off for their migration to CASTOR2. ALICE generated 14 TB of MC
events last week with jobs opening calibration files at CERN using xrootd over Wan.
Current main concern is lack of SRM end point at RAL for T0-T1 transfers. Reportedly
due to lack of disk space though Alice may be offered 1.6TB of the 2 TB they need (this
will allow 9 hours of transfers). Another problem is that only Nikhef fully implement a
tape SE with garbage collection. Alice will concentrate on using CNAF to build up their
transfer service and will organise a meeting with the CERN FTS team.

ATLAS: Many of their T1 sites were down with cooling problems which gave poor
performance when they tried a short T0 data expert test on Friday 21st. ATLAS have now
switched to using their simpler batch job load generator for data export and in the near
future will be testing the effects of reducing the size of the CERN export disk pool and of
exporting smaller or larger file sizes (simulating different mixes of raw, aod and esd). A
major change of plan is they will eventually return to full nominal export rate and
continue this up to their next full T0 DAQ simulation in September. Will this pose a
problem for Tier 1 sites (e.g. in tape recycling)?

CMS: About to start generating the 1 M events/day needed for CSA06. The CSA06
metrics will be measured by CMS and are summarised in talks given at
http://indico.cern.ch/conferenceDisplay.py?confId=4669. Current main concern is very
poor job submission reliability of the gLite RB (below 10% for single submission and at
64.2 % for bulk submission while that of the LCG RB was 99.4%). The CMS job robot
was used for submission and the jobs were short analysis jobs which read and parsed data
from a local SE. A dedicated RB to receive all the latest gLite patches has now been
setup. There will be a close collaboration with ATLAS in this work. A second major
concern for CMS was that they planned to need 3D services to be ready by mid-
September with Oracle streams running at CERN and squid at T1 sites and streams at a
few T1 sites. They have now relaxed this requirement and if any merging is needed they
will do it by hand.

LHCB: Most of the outstanding problems have been understood and fixed or have a
workaround. There is a long LHCB to do list as a result. They had switched to using
srmcp instead of lcg-cp which was pointing them to the wrong castor srm at CERN due to
a bug in gfal but are now using lcg-cr plus lcg-uf. This has also shown up that CERN
needs to redefine the castorsrm alias to point to srm.cern.ch rather than castorgrid.cern.ch
(which circumvents the bug) and this will be done tomorrow. The current main concern
of LHCB is the slow data rates for transfers of montecarlo data back to CERN. This
cannot be reproduced except intermittently by James and Roberto at CERN and the

LHCB workaround has been to increase a timeout to 1000 seconds. The problem of file access to SARA from Nikhef worker nodes has been worked around by only using Nikhef in Tier 2 mode i.e. as a source of montecarlo events.