

MB (Snapshot) Report on Status and Progress of SC4 activities

A weekly report is gathered from:

- Reports and issues from the Experiment Integration and Support teams given at the weekly LCG Service coordination meeting (Wednesday mornings)
- IT group reports to the weekly IT C5 internal meeting (Friday mornings)
- Reports and issues from the experiment service challenge teams at the weekly LCG Resource Scheduling meeting (Monday afternoons)

This is presented to the weekly EGEE operations meeting and appropriate site and experiment issues are discussed.

Last weeks report after this meeting is attached to the agenda.

Strategic or long-term issues are collected in a Service challenge action list twiki at

<https://twiki.cern.ch/twiki/bin/view/LCG/SCActionList>

ALICE SC Status

- From last week of July ramping up to 300 MB/sec Tier 0 to Tier 1 data export to 6 sites (CNAF, IN2P3, RAL, FZK, SARA, USA site).
- Have reached peak of 150 MB/s but only for a short period. Currently well below their target rate (less than 50 MB/s aggregate). Will continue to try and reach target.
- No end points available to them at RAL and NDGF (would be included).
- Had some srm_put problems to SARA.
- IN2P3 currently giving FTS preference to CMS.
- Main issues:
 - Site instabilities
 - Poor feedback from submitting individual problem reports to GGUS. Switched to using an ALICE internal list but will now also include GGUS.

ATLAS SC Status

- Have had post-mortem meetings on their Tier 0/Tier 1 first pass processing and data export exercise.
- Overall good with no major issues for CERN but some problematic T1 sites.
- Have just passed the milestone of transferring 1 PB of data in the last 50 days.
- Currently continuing data export at 4-500 MB/s but plan to repeat the whole exercise at the full rate of 770 MB/s in last two weeks of September.
- This overlaps with the start of CMS CSA06 and a more intensive MonteCarlo campaign by ATLAS so we are asking if this can be moved forwards.

CMS SC Status

- Planned data export exercise of disk to tape at 150 MB/s delayed while they try a higher rate Tier 0 to Tier 1 disk to disk exercise.
- Target is 500 MB/s for 1 week, minimum acceptable is 300 MB/s for 3 days
- Started badly with CERN power cut knock-on effects (Monday) then having an expired FTS proxy server Tuesday (this will be better managed when the service moves to FIO on new high reliability hardware in September)
- Reached 350 MB/s Wednesday but dropped overnight then recovered – not known why. Stopped for CERN Oracle data base upgrade on Thursday.
- Reached 300 MB/s Thursday, dropped overnight and effectively stopped Friday morning. Unexpected status replies from Castor caused Phedex to issue 80000 'prepare-to-get' requests which took hours to process.
- Reached 350 MB/s Friday afternoon and this has been maintained.
- Low rates to RAL and were not using CNAF due to their Castor2 migration. At operations meeting CNAF told them they could restart. Ran badly overnight but since 09.00 CET CNAF have been taking about 100 MB/s and CMS have reached their 500 MB/s target.
- Plan to continue this week or till they have seen sufficient stable running.
- A 30 minute Castor stop at CERN Wednesday will test recovery performance.

LHCB SC Status

- Planned to start raw data distribution and reconstruction and stripping at Tier 1 in July. In parallel continue MonteCarlo event generation at Tier 1 and Tier 2 with events stored back to CERN.
- Last Tuesday CERN changed castorsrm.cern.ch endpoint from castorgrid to srm.cern.ch to fix a gfal bug affecting LHCB transfers back to CERN.
- LHCB store batch job logs back to a central CERN classic SE (run as a VO-box service) using lcg-cp (which runs a gridftpd on the SE). This ran well during July then started having process-avalanches and crashing the box. This was traced to exceeding a 32000 directory entry limit which the lcg-cp/gridftp chain did not handle well. This was fixed and also LHCB reduced the numbers of individual files sent so the box is running well but IT remains worried about the scalability of a central gridftp file server for job logs. ATLAS run a similar service currently at lower rates. Improvements are being studied e.g. use castor2 or gridftp2 but a proper grid architecture solution would be desirable.
- LHCB observe intermittent (very) slow performance transferring some files from remote worker nodes to CERN. This happens to Castor2 and a classic SE so is not thought to be a Castor2 problem. As a workaround LHCB have increased their transfer timeout to 1000 seconds.

Issues from last week

- ALICE request a disk0/tape1 endpoint from RAL.
- ALICE request a disk0/tape1 endpoint from NDGF.
- CMS transfers to RAL run at a low rate (10 MB/s compared to target 50 MB/s).
- Medium and long-term solutions needed for the collection and access by end users of batch job logs.
- Understand limitations/useability of direct file transfer from worker nodes to CERN.
- All sites should be able to provide long term stability at their MoU performance levels.

Open Issues in SC Action list

- **General actions**
- Site monitoring of local services - to be discussed in detail at September 15 [SC Tech Day](#)
- Tools to monitor transfer activities on FTS channels; provide access to FTS logs (Gavin, Paolo) (Temporary solution being tested)
- Improve the performance and reliability of the gLite Resource Broker (critical path for CMS CSA06)
- Support multiple priority levels for grid batch jobs (critical path for LHCb distributed analysis)
- **ALICE actions**
- ALICE to report on its T1-T2 end points.
- **LHCb actions**
- Root/pool data access to SEs is needed