

Table of Contents

Week of 080519.....	1
Daily CCRC'08 Call details.....	1
General Information.....	1
Monday:.....	1
Tuesday:.....	1
Wednesday.....	3
Thursday.....	4
Friday.....	5

Week of 080519

Open Actions from last week:

Daily CCRC'08 Call details

To join the call, at 15.00 CET Tuesday to Friday inclusive (usually in CERN bat 28-R-006), do one of the following:

1. Dial +41227676000 (Main) and enter access code 0119168, or
2. To have the system call you, click here [☞](#)

General Information

- Web archive: <https://mmm.cern.ch/public/archive-list/w/wlwg-ccrc08/> [☞](#)
- CERN IT status board: <http://it-support-servicestatus.web.cern.ch/it-support-servicestatus/> [☞](#)
- elog: <https://prod-grid-logger.cern.ch/elog/CCRC'08+Logbook/> [☞](#)

Monday:

See the weekly joint operations meeting minutes

Additional Material:

Tuesday:

Attendance: local(Olof, Roberto, James, Jamie, Simone, Flavia, Julia, Andrea, Patricia);remote(Michael, Gonzalo, JT, Derek, Vincenzo INFN/Bari).

elog review:

Experiments round table:

- LHCb: main issue since yesterday (since vulnerability issue in UK) & after AFS upgrade. Andrew, using UK-issued cert cannot run any transfers over night/morning. Restarted using Robbie's credentials. Banned IN2P3; LHCb clients based on currently in production version of gfal (1.10.8) were crashing against IN2P3. We will move then to latest version (1.10.11) and check carefully whether this fixes the problem. If not we will pass the ball to Remi&Co. Re-integrated SARA SRM endpoint and running extremely smoothly - recovering backlog quickly - until credentials problem. JT - WN there now. Not 100% yet... Jobs will start to come. SRM CASTOR instance at CNAF - endpoint down, restarted, although scheduled downtime today. Recons - PIC, RAL, CNAF, CERN smooth. Restarting at NL-T1. Lots of jobs waiting at FZK - exhausted share after many jobs over w/e. Stripping - da Vinci crashing due to mem leak. Ancestors info not available - still being investigated by book-keeping experts.
- ATLAS: castor at CERN - problem appeared about 01:00. Not noticed until this morning. Submitted ticket at 09:00. Problem cured < 11:00. Outage ~10 hours, but response time quick once reported. Meeting with castor people this morning. Discussed main two issues observed during May. 1). Error message "exhausted # threads" - patch about to be deployed. Olof - SRM down overnight, not castor instance. Same problem as late last week - cgsi problem, all threads stuck in gsoap layer. No possibility to create new threads. Still SLC3. 2nd problem - deadline problem - also slowness of

stager to putdone, srmrm, to be patched tomorrow. srm release today? Flavia - Yes. Olof will check if this will be upgraded tomorrow (or immediately?) Simone - tomorrow throughput test starts. 09 - 11 tomorrow is ok - later postpone. Simone - also affected some exports to BNL end last-week. Draining q of M7 transfers, will continue for a few hours, then setup machinery for export tests. Mail to ATLAS T1s with rates etc, also global rate (1.1-1.2GB/s out of CERN). Follows computing model as per usual. Sites can be over-subscribed to tape - also disk (except BNL...) to test if sites can take more. T2s will also participate based on shares given by clouds. T2s can also over-subscribe. MC production - finishing for FDR2. Tail - Taipei lost 170K files last week(!). Some critical for FDR2 - being recreated. Affects only 3 clouds - FR, US & CA.

- ALICE: Got srm e-p for castor area at RAL. All T1s ready for data transfer. Alien finished upgrade to 2.15 at CERN (central services) will begin installation at sites this week (not yet announced). Sites not yet fully ready: Birmingham (site manager hopes to be ok), Kosice (not expected to be ready). For remaining sites should be ok. Overlap with ATLAS? Simone - next week kept for contingency. JT: NIKHEF failing (warn) ALICE SAM test - VO box proxy renewal. Patricia - normally defined as critical. In some cases test can fail but test needs to be adapted to new VO box. Removed from critical tests for moment pending migration. Will be reinstated later. For the moment "don't worry".

Sites round table:

- NIKHEF: (part of NL-T1) we have just increased our installed capacity by 120 TB compared to yesterday. Furthermore we now have in total 88 TB of space allocated to ATLAS. Problem with cooling last week hence WNs. Migration of power from backed up to non-backup. WNs on "cheap power". Still not complete with this migration. Not clear yet how many will be able to be turned out. Replacement part "propeller like thingy" stuck on boat in harbour in Marseille and delayed by strike(!). Play-by-ear. ATLAS say storage rather than WNs please. DPM - problem with resize of space reservation. This resets lifetime to 1.7 days (1.6.7). Asked by ATLAS to upgrade - ongoing now.

Core services (CERN) report:

- Some details about upgrades tomorrow and Thursday - CMS asked for more details of what is involved. Fixes slow stager rm - hit ATLAS hard and could hit also CMS. Problem - index not properly handled and so "hint" had to be added to avoid full table scan. Simone - problems also caused by internal FTS logic (cleanup actions triggered if file transfer fails.)

DB services (CERN) report:

Monitoring / dashboard report:

- Two instances of collector running in parallel by mistake - too much load on DB. Affected CMS dashboard. Killed and now ok. CMS running analysis exercise. Want to follow i/o rate between storage and WNs. Request to setup Mona Lisa for this. Hope to setup this week.
- ATLAS dashboard upgrade went smoothly yesterday as scheduled.

Release update:

AOB:

- elog now available? Some problems due to config change. Down 14:00 - 14:30 - should be ok now.

Wednesday

Attendance: local(Nick, Jamie, Roberto, Patricia, James, Julia);remote(Michael, Derek, Simone).

elog review:

Experiments round table:

- ATLAS (Simone): throughput test started this morning - generation of data at 09:00. Problem with registration in ATRLAS catalog. First data after lunch. RAL/CNAF problems - solved in 'no time'. In last 1 hour all running smoothly. Throughput as expected. SARA endpoint missing -in scheduled downtime. Channel inactive. Roberto - intervention over. Simone - transfers queued in FTS - will contact SARA people. Michael - are we currently at nominal? Only see 120MB/s to BNL. CERN-BNL channel rather full. 4 active transfers, 26 ready. See from FTS logs if t/p per transfer ok. FTS logs published but only from day before. RAL - problem with CA. Did not upgrade CA. Graeme Stewart sent top priority e-mail now fixed. CNAF - f/s behind STORM f/e not mounted. Re-mounted and now works.
- LHCb (Roberto): observations to report - LHCb wants to confirm same issue at RAL as ATLAS. Echo of UK CA upgrade. Local contact confirmed upgrade of certs into CASTOR diskservers later today. GV plots - not transferring out of CERN. Problem with LHCb online servers at pit and thence out to T1s. Not known when will be fixed. Huge backlog of jobs to SARA & IN2P3. Smooth transfer rate over last few days. IN2P3 - after 4 days kept off production mask due to gfal issues. Still facing issues with gsidcap server. Filled elog entry with plot showing CPU efficiency over last week for IN2P3. Essentially 0 which is not considered acceptable for a T1(!). Upgraded gfal to recommended level but issues still there. Stripping - run happily test jobs with patched version of da Vinci. Hope to restart stripping this week. recons fine at all T1s + T0 except IN2P3 and SARA. Few space token issues at several sites - ggus tickets opened. Site admins very responsive. Policy for deploying new CAs - policy at JSWG. Should be fully transparent to VO - clearly not the case. Keep old and new for transition period. Not case.
- ALICE (Patricia): not making any production. Transfers. Still making transfers with FTD from Alien 2.14. New version - 2.15 - not being used. 2 groups - 1) dummy files to check channels and keep FTS active. 2) Some real files to FZK and NDGF. working ok. Dummy files to CNAF - several issues, mostly to SE name in FTD config (not site) IN2P3 and SARA - Lyon upgrading VObox only certified today. Hope will start once 2.15 in case. + for SARA will not make transfers until 2.15 in place. RAL (new since Feb) - SRM endpoints provided yesterday. elog issue can be closed. Info in config file - test channel for first time.

Sites round table:

Core services (CERN) report:

- LCG OPN: We need to upgrade the software of the two CERN routers that connect to the Tier1s (there is a bug that affect the correct routing for the backup paths in some cases). It is necessary to reboot them, that means a downtime of 5 minutes each. We will first upgrade one, than few days later we'll proceed with the second.

Would it be OK if we schedule one reboot for Monday 26/5 and one for Wednesday 28/5? Or do you prefer we postpone them to later dates?

DB services (CERN) report:

Monitoring / dashboard report:

Wednesday

- Moved gridmap onto new h/w - modern mid-range server. Gets about 25K hits a day. Should be faster!
- A draft of the report to next week's MB is attached to this page. See the break-down of various service issues seen last week (and reported on this). Did the existing monitoring / logging / reporting pick up the issue? Are there specific actions?

Release update:

AOB:

Thursday

Attendance: local(Roberto, Gavin, Simone, Daniele, James, Jamie, Julia);remote(Michael, Derek, Gonzalo).

elog review:

Experiments round table:

- CMS: (Daniele) transfer out of CERN to T1s - some authorization issues at RAL & IN2P3 - different issues - CA (RAL) proxy expired (IN2P3); Still relying on test traffic, not only prod from CSA (not high enough...) Harder than Feb to keep rate high! Like to check in more details FTS server logs. Gav - to look at rates? Rates, status of files - to be sure filling qs -> more offline. Long tail of T1-T1 testing. Collecting all plots for post-mortem. T1-T2s to test also non-regional traffic - ok. Analysis = moved to 2nd phase, random job submission at some T2s to other T2s. Mostly rely on dashboard to see successes / failures. Some time spent on CASTOR wrt garbage collection. Issue with myproxy at CERN - not renewing proxy. Julia - monitoring for T0 (A+A) can it be used for other transfers as well? Gav - coupled to SL(C)4 version. Can this be used to help CMS? Maybe add more info? Daniele - good if done centrally and in common. Access also to logs of pilot server.
- CMS readiness - T0 workflows tested in Feb. Prompt recons etc also done in Feb. Full chain still due. Like to complete this within days. T1 workflows - reprocessing and skimming both being exercised. Get all measurements from this exercise? Full exercise with some overlap with others. Transfers - T0-T1 stable, try to understand and measure behaviour with others. T1-T1 good shape, T1-T2 being done now. Full mesh hard - T2-T1 easier (ok). T2 workflows (bulk MC and analysis) - analysis already - extend with chaotic analysis. For MC prod - partially by T2s but also T1s. >8K slots used for MC prod.
- ATLAS: (Simone) - sent a few slides (to be attached) - current status of throughput - quite good except: 1) 2 castor issues yesterday. 16:40 - 17:20 - connection to DB died. Exhaustive explanation from castor team, attached to elog. (srm request db - maybe attached wrong way round!). 2) Started during football game(!) - ticket sent. Spotted by Michael. Some problem at gridftp level - 'size error'. Full error message in elog. Disappeared by time ticket submitted! (22:30 - 00:00). Thought to be backup, whereas previous issue not yet understood. 3) Fix for 'Monday issue' - watchdog - + patch for castor yesterday morning (slow stagerrm thing). elog entries will be updated... Other issues (RAL, CNAF) were fixed during meeting. Performance problem at BNL - backup link limited to 1Gbps. BNL now running at a bit more than nominal. Just after midnight one write pool in TRIUMF offline - generating errors. Hopefully fixed(!) More details in slides. Delivering data also to T2s. Report from Alexei also in the attached slides. Monitoring in clouds - contacts to follow up. Planning for next week - full exercise all in parallel. Explained in slides. Start from Monday. Throughput exercise can finish this Saturday. Part of next week's exercise would be reprocessing - 'proof of concept' of calib db access, but not full stress load. Daniele - all T1s and all T2s next week? A: T1s, yes. Most T2s - some in Spanish cloud have not finished deploying SRM v2.2. 95% of T2s.. Tapes at T1s: no - all from disk. Tape to write raw from T0. Reprocessing is from tape. T0 disk -> T1 tape etc.

- ATLAS readiness - transfers in v good shape, MC quite robust, reprocessing still to be demonstrated.
- LHCb: (roberto) - internal problem with online, fixed, backlog transfer at good rate. 120MB/s out of CERN. Another problem with mounting of daq area. Ramp-up again. Reconstruction - smooth. IN2P3 better - problem fixed by restarting gsidcap door. Unban site & restart recons & stripping. NIKHEF - cleaning up large backlog of recons jobs. working fine. Banned RAL as recons jobs crashing with rfio connection failure. Nothing strange reported- follow offline. Stripping - 9K jobs pending. Many for RAL & IN2P3. 2.5k jobs failed with 'ancestors' problem. 1K jobs CERN, PIC, NIKHEF, FZK. CNAF shared area - input of python modules timing out. On-going... Analysis: ganga i/f to dirac3. End this week or next will include analysis.
- LHCb readiness: infrastructurewise we are ready - still polishing on lhcb side. all steps tested. Not yet analysis...

Sites round table:

- RAL: Please note that RAL is closed for a public holiday on Monday 26th. This means that any problems caused by this (LCG OGN intervention above) will be dealt with by the on-call system and responses will not be as swift as on a normal working day.

Core services (CERN) report:

- C2CMS intervention
- Jan has updated castor problems page. 5 cases with service outtages - workarounds in place for 4. Vulnerable to problems like yesterday pm (too many connections) - in elog link [☞](#)

DB services (CERN) report:

Monitoring / dashboard report:

Release update:

AOB:

Friday

Attendance: local(Jamie, Roberto, Julia, Nick, Simone, Patricia, Andrea); remote(Derek, anonymous).

elog review:

Experiments round table:

- CMS (Andrea) - transferring to-t1 reco data (production transfers). CASTOR G/C problem - solved. Thanks to castor team. All other exercises going well. Next week all functional blocks simultaneously
- LHCb (Roberto) - pit-t0-t1: online problem seems fixed. accumulated huge backlog - spike of 160MB/s. 6 hour duty cycle transfers - hope no more online probs! t0-t1: some probs with gridftp timeouts cern-sara cern-pic (degrading performance of channels). Same as reported by simone. Looks to be fixed at PIC. (Reconfigured dcache implementation - all pools full except 2 - redirector didn't take into account). recons: running 'moderately fine' - 5 sites. in2p3 & ral banned as recons app crashing. LHCb or site config issue? Under investigation. Up to 500 jobs successfully, 300 @ CNAF, >200@FZK, >600@NIKHEF - 200 failing during tmpwatch deleting files or running jobs. PIC: ~200 jobs. Reflects T1 share. Stripping: 6K pending jobs - still ancestors problem - under investigation.

Upto 2K jjobs at CERN, 600@CNAF, 1/2 failing, 700@FZK, 0@NIKHEF - mismatch failing. 450@PIC. Plans for next week: stop stripping (degradation of recons performance - never run up to 1800 jobs as before). Need to concentrate on recons, re-integrate in2p3 & ral. Once ok, 'great success of last week' - go back to stripping. High throughput of ATLAS+CMS affects T1 data access? (Guess...) Local protocols - not obvious that concurrent access would affect this.

- ALICE (Patricia) - moved start of 3rd commissioning exercise to start of June (2-22) - problems with detectors. Still making some transfers. Increase size from 2GB to 10GB. Several issues with CNAF, 99% efficiency transferring to CNAF. Still to deploy Alien 2.15 (SARA, IN2P3). Still to install Alien in VO box.
- ATLAS (Simone) - throughput as expected (i.e. what it should be!) - zero backlog! tiny backlog for cnaf. Everything flowing smoothly. Next week: finally 'all inclusive' exercise. Full steam also production system. Data transfers as described yesterday. Only aspect still missing is reprocessing, starting from Monday. staging agent not well tested. Big step forward yesterday was understanding SARA 'slow-ness' - problem was due to dq2 bug - same file transferred several times. Plus parallel import from reprocessing into SARA tape - 110MB/s - quite impressive! Now 2nd best performing site (after BNL).

Sites round table:

Core services (CERN) report:

- Several problems and actions on Data services in relation with the ongoing CCRC tests
 - ◆ Castor and SRM services suffered from instabilities, which are actively being followed with the developers. Several of these issues are understood and fixed in software, and we have deployed a number of bug fix releases during this week Castorcms and Castoratlas now run the latest 2.1.7-7 release, addressing the slow stager_rm and putDone commands.
 - ◆ the Castor nameserver daemon has been upgraded to 2.1.7-7, to work around the overloaded service during backup runs of the databases
 - ◆ Castor SRM has been upgraded to 1.3-22, addressing some deadlock issuesAll these upgrades were transparent.
- Long standing (since beginning of CCRC) CASTORCMS/t1transfer GC problem was finally understood by the castor development team yesterday: the SRM v22 prepareToGet request was giving a too high weight for files already resident in the pool while files being recalled from tape (like in early May) or copied from other diskpools (like the last 2 weeks) were given a too low weight. This caused a very high rate of internal traffic and it is quite impressive that the system managed to keep up given the load it generated on itself...

DB services (CERN) report:

- Still experiencing highload from CMS dashboard - plan review after CCRC'08.

Monitoring / dashboard report:

Release update:

- TRIUMF may go offline next week. Escalate if necessary to TRIUMF. Alot of pre-staging starts from TRIUMF so this would be painful!

AOB:

-- JamieShiers - 16 May 2008

This topic: LCG > CCRC08DailyMeetingsWeek080519

Topic revision: r10 - 2008-05-23 - JamieShiers



Copyright &© 2008-2020 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

Ideas, requests, problems regarding TWiki? Send feedback