

# Table of Contents

<b>Week of 080804.....</b>	<b>1</b>
Daily WLCG Operations Call details.....	1
General Information.....	1
Monday:.....	1
Tuesday:.....	2
Wednesday.....	3
Thursday.....	3
Friday.....	4

# Week of 080804

Open Actions from last week:

## Daily WLCG Operations Call details

To join the call, at 15.00 CE(S)T Monday to Friday inclusive (in CERN 513 R-068) 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-operations/> [☞](#)
- CERN IT status board: <http://it-support-servicestatus.web.cern.ch/it-support-servicestatus/> [☞](#)
- elogs: <https://prod-grid-logger.cern.ch/elog/CCRC'08+Logbook/> [☞](#)

See the weekly joint operations meeting minutes

Additional Material:

## Monday:

Attendance: local(N.Thackray, J-P.Baud, H.Renshall, R.Santinelli, S.Lemaitre, L.Canali);remote(D.Ross, M.Ernst).

elog review: One new entry from LHCb. There is a 1-2% failure rate to return the turl of a file (times-out after 30 seconds) at SARA with no obvious pattern. Since LHCb jobs typically read 20 files this translates to a high job failure rate.

Experiments round table:

ATLAS (mail from A.Klimentov): end of week 31 statistics is available at <https://prod-grid-logger.cern.ch/elog/ATLAS+Computer+Operations+Logbook/482> [☞](#) ALL Tier-1s received 100% of FT data from CERN (at the same time all cosmic data is also replicated). We do not have double registration for Tier-1s and Tier-2s during this week - thanks to DQ2 developers and all sites for keeping the system 100% efficient.

LHCb: one of the pending issues, that of Storm reporting it has deleted files when it has not, is accepted as an internal problem that will be fixed with the next release. During this week LHCb will start a low rate CCRC activity running their generic pilot jobs model. They are hitting a limit that the WMS only allows 10 delegated proxy renewals so for the moment are submitting WMS work with a 1 week proxy with no renewals. N.Thackray reported this is a limitation in vdt and he expects to get some information from them on this tomorrow. Sophie asked why there are very few jobs for LHCb at CERN the answer being they are still commissioning for the next round of Dirac2 running.

Sites round table:

Core services (CERN) report: There was a UK CA event that resulted in all UK certificate owners being sent a mail from CERN saying incorrectly that their certificates had expired. Corrective mails have now been sent.

DB services (CERN) report: LFC streaming to FZK has been stopped due to instabilities. A post-mortem of last weeks failure of Atlas offline to Tier 1s replication between Saturday and Wednesday (4 days) is being prepared, in particular why it was not spotted earlier. The issue was caused by the occurrence of a gap in the archivlogs sequence propagated to the downstream capture database.

Monitoring / dashboard report:

Release update:

AOB:

## Tuesday:

Attendance: local(H. Renshall, J. Andreeva, R. Santinelli, M. Marques Coelho, J-P. Baud, P. Mendez, S. Campana, L. Canali, S. Lemaitre, O.Barrington, D. Wojcik, );remote(D. Bonacorsi, M. Ernst, D. Ross, J. Coles).

elog review: No new entries

Experiments round table:

- LHCb: Nothing to report from the experiment. From RAL, J. Coles has asked for an specific broadcast to explain in detail the pilot role required by the experiment and that has to be implemented at service level.
- ATLAS: The experiment has announced a stress test of the new data management version of the experiment for this week (next Thursday). During this test T0-T1, T1-T1 and T1-T2 transfers will be checked with the new package for about 12h. It is not considered a site stress test, but rather a check of the new software version of the experiment. The exercise will be stopped for 12h beginning again on Friday and by the next week continuing the stable production with the new package. A broadcast will be sent to the sites by S. Campana to announced the required space which the exercise requires. In addition X. Espinal has reported that OPN was down. The question was the best procedure to report problems of this nature. The corresponding wiki page contains the mailing lists where to submit such problems, the link will be provided by H. Renshall.
- CMS: Discussion before the meeting with M. Marques on the best procedure to follow to upgrade the CMS instance of Castor at CERN and at the same time to avoid breaking the Castor@CERN rule of "no Castor upgrades on Fridays". After a further discussion also at the CMS Data Operations meeting soon after the WLCG call, an agreement was found on Tuesday, August 12th, 09h00-11h30 CERN time, as a trade-off between the CMS plans and needs (midweek global runs on Wednesdays-Thursdays, CRUZET-4 starting on August 18th for 1 week) and the Castor@CERN constraints. --- CERN IT has announced a visible job failure rates due to memory limitations in the WNs. The issue is being followed with U. Schwickerath which is checking the memory requirements of the VO. In addition certain low quality in the transfers to Lyon T1 has been observed. Following this issue directly with the site.
- ALICE: Experiment continuing the production. The experiment has not suffered by any interruption due to the Castor upgrade at CERN last week.

Sites round table: Nothing to report

Core services (CERN) report: Nothing to report

DB services (CERN) report: Regarding the stream problem reported by the ATLAS VO, the problem has been already fixed and the post-mortem evaluation is also available here.

Monitoring / dashboard report: Nothing to report

Monday:

Release update: No news

AOB: No further news

## Wednesday

Attendance: local(Jean-Philippe, Olof, Harry, Luca, Nick, Steve, Patricia, Simone);remote(Derek, Jeremy, Michael).

elog review:

Experiments round table:

LHCb (RS): preparing an EGEE broadcast on what sites need to do to support the LHCb multi-user pilot jobs tests. LHCb are resuming part of their CCRC running and will be exporting data from CERN at their nominal rate of 70 MB/sec for about 30 minutes each 6 hours. They asked the status of the vdt fixes that would extend the number of allowed proxy delegations. Nick Thackray will check this but thought it might even make the current release 28 pack.

CMS (DB): addressing some DB-related issues, network at P5 issues, and communication improvements at the IT/CMS coordination meeting running right now (and clashing with the present phone call). Following up of the IN2P3 transfer problems reported yesterday, 1) S.Gowdy and D.Mason debugged and it seems it was related to a few stubborn files (namely, belonging to the v4 repack of CRUZET-3) which if rfc'ed just hang.. Stephen removed these files from disk and restaged them, and can rfc' them now, PhEDEx will retry automatically. 2) please note that the T1 is in unscheduled downtime now due to a AFS outage;

ATLAS (SC): asked the status of the 1.6.11 release of the LFC (this fixes the periodic crashes). Nick Thackray replied they were accelerating its passage through the PPS and he hoped for a quick release.

Sites round table:

Core services (CERN) report (ST): Monday mornings VOMRS lcg-voms.cern.ch hardware migration was not perfect after all. New registrations after 09.00 CEST on Monday 4th of August were held up at the final stage of being synchronized to voms-admins until 21:00 Tuesday 5th of August. This has now been corrected and any stuck registrations have now gone through. Why? For all VOs the synchronization was being attempted against the test VO. Preliminary tests used the test VO which is why we did not notice till a dteam user reported a problem. No long term harm done and nothing was lost.

DB services (CERN) report: streams replication to FZK of the LHCb LFC is still disabled after nearly one week. A post-mortem will be expected.

Monitoring / dashboard report:

Release update: N.Thackray asked if any experiments are testing gLexec in a non-SCAS (i.e. with the usual local LCAS mapping) configuration. Alice, LHCb and ATLAS said they are not.

AOB: H.Renshall announced he has added a link to the OPN Twiki (as requested yesterday) on the WLCG Operations web page <https://twiki.cern.ch/twiki/bin/view/LCG/WLCGOperationsWeb> under the Operations Portals and Tools section.

## Thursday

Attendance: local(Luca, Sophie, Ewan, Harry, Jean-Philippe, Roberto, Julia);remote(Daniele, Derek,

Tuesday:

Michael).

elog review:

Experiments round table:

LHCb (RS): Low rate CCRC-like activities have now started with 30-minute bursts of data from the pit through Tier-0 to Tier-1 at an aggregate of 70 MB/sec each 6 hours. An EGEE broadcast was sent requesting sites to install a new LHCb pilot role for the testing of multi-user pilot jobs.

ATLAS (SC): Throughput tests (full nominal data rates) started at 10.00 today and will continue for 12 hours. Functional tests (10% rates) will resume at 10.00 on Friday. This weekly pattern will probably be repeated from now on. There are 4 outstanding problems 1) ASGC has a 40% timeout failure rate in srmpu 2) srmpu is also timing out at FZK apparently due to excess load within nfs 3) SARA has permanent srmpu failures thought to be an old dcache bug where a new directory is created with ownership by root. This has to be reset manually 4) RAL LFC has been down since 1 hour ago. Transfers have only just resumed. D.Ross said there was no scheduled downtime today so he would check on this.

CMS (DB): are running the third of their planned four midweek global cosmics runs and have detected no major problems. FZK is the custodial site for this exercise. Preparations are going ahead to start coordinated shift running at CERN and FNAL.

Sites round table: a local network provider failure at BNL yesterday did not lead to any outage as there is a redundant path.

Core services (CERN) report: ATLAS request to know the status of their request for 2 new Castor pools. Sophie will check with Miguel.

DB services (CERN) report: Streams replication from CERN to Gridka of LFC for LHCB was down from Thursday 31-7 till Wednesday 6-8. The issue was caused by a problem on the Gridka RAC cluster. Currently the replication has been restored using the work-around of using Gridka cluster node 2 only. The problem with RAC node 1 is still under investigation by Gridka DBAs.

Streams replication from atlas offline to online has stopped in the morning of 7-8 because of an unsupported update of conditions data. The replication has been restarted in the early afternoon and errors have been manually cleaned.

Monitoring / dashboard report:

Release update:

AOB:

## Friday

Attendance: local (Harry - competing with Olympics opening!);remote (Derek, Simone, Gonzalo, Jeremy, Daniele).

elog review:

Experiments round table:

ATLAS (SC): The throughput test ran as announced yesterday starting at 10 am. It took 3-4 hours to have the whole machinery running (LSF jobs have to finish) so data export started about 14.00. Fresh injections were

Thursday

stopped at 10 pm and the data preparation had drained by 3 am today. There were two problematic sites namely FZK which is suffering from an overloaded nfs, as yesterday, and RAL where transfers only started at 14.00 today. An operator alarm ticket was raised but they had to wait for expert help. D.Ross explained that the ATLAS SRM is not coping with high transfer load as well as the CMS one does so they cut the number of concurrent transfers from 40 to 20. They are currently receiving data at about 60 MB/sec but this is insufficient for fast catchup. Their SRM expert is on holiday for another week. ATLAS also saw an unexplained gap in transfers to BNL from 09.00 to 11.00 CEST today. They also had problems with 2 of their 4 calibration sites - Naples ran out of disk space and Munich was problematic. ATLAS have now resumed functional tests at 10% rate but RAL and FZK will give errors on retransmitting data until they have caught up.

CMS (DB): Work is running smoothly. They have some internal operations tickets but nothing unusual.

Sites round table:

RAL (DR): As regards the LFC problems reported by ATLAS yesterday we attribute this to the fact that in the reported time window we were doing dcache to castor data migration and this overloaded our LFC.

PIC (GM): are having overload problems of their PBS master node with timing out errors affecting the stability of their batch system. The node had been replaced on Tuesday so they decided to revert to the previous master node (done yesterday evening) but this also shows the same behaviour. They suspect issues in the configuration of the CE that is forwarding jobs to PBS and have scheduled a 2 hour downtime on Monday morning to investigate. They will drain the batch queues as normal.

Core services (CERN) report:

DB services (CERN) report:

Monitoring / dashboard report:

Release update:

AOB:

-- JamieShiers - 30 Jul 2008

---

This topic: LCG > WLCGDailyMeetingsWeek080804

Topic revision: r14 - 2008-08-08 - HarryRenshall



Copyright &© 2008-2021 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

or Ideas, requests, problems regarding TWiki? use Discourse or Send feedback