

# Table of Contents

<b>Week of 081020.....</b>	<b>1</b>
Daily WLCG Operations Call details.....	1
General Information.....	1
Monday:.....	1
Tuesday:.....	2
Wednesday.....	3
Thursday.....	4
Friday.....	5

# Week of 081020

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/> [☞](#)
- elog: <https://prod-grid-logger.cern.ch/elog/CCRC'08+Logbook/> [☞](#)

See the weekly joint operations meeting minutes

Additional Material:

## Monday:

Attendance: local(Gavin, Jean-Philippe, Harry, Jamie, Sophie, Steve, Roberto, Simone, Andrea, Markus, Jan);remote(Gareth, Michael).

elog review:

Experiments round table:

- ATLAS (Simone) - last Friday several activities were (re-)started following decisions Thursday. Replication of MC AODs restarted; consolidation of some SRM v1 data -> SRM v2 endpoints (data not needed will be exterminated); action from last week on huge backlog of cosmic data, subscriptions cancelled, some files recalled from tape and on Friday distribution (mainly BNL & Lyon) restarted. Activity continues now... Several hiccoughs over w/e: a few instabilities due to v. high load due to above + normal data export (e.g. SRM + SE overload, FZK, IN2P3, Taiwan Fri/Sat) - relatively minor. Bigger issue Saturday with RAL: ATLAS started observing v. high failure rate. Service not completely dead but failures (Oracle error) > 90%. Fixed this morning - comments later from RAL(?). INFN cloud: problem with CNAF tape service and with one of calibration sites (Napoli). Napoli problem still there, CNAF solved this morning. JP: what was problem? A: srm preparetoput failed. Issue with transfers to Lyon: not sure if error is from FTS or SRM - will follow offline. Above activities continue.. Other news: cosmic data taking continues for one more week - probably until Monday next week. Will likely open detector after that... Might still take data with sub-detectors...
- RAL (Gareth) - 2 FE m/cs for ATLAS. Slow DB queries. Unbalanced in load. Various interventions over w/e to remove one of m/cs from load balancing went horribly wrong - now back with 2. Running ok for now... Degradation from ~04:00 Saturday morning until ~10:00 this morning - unscheduled outage in GOCDB for this time.
- ATLAS (SC) - starting from yesterday problem with unavailability of files from CASTOR at CERN - srmget - pool overload? Ticket exists..

- RAL (Brian) - transfers MCDISK - MCTAPE at RAL ATLAS maybe don't get files where they want.. Depending on order and how CASTOR puts files on tape. ATLAS explicitly making 2 copies, 1 disk 1 tape. Issue may therefore affect all CASTOR sites.
- NL-T1 (JT) - no jobs at NL-T1 since last Tuesday. Simone will check...
- CMS (Andrea) - running global run for w/e. Problems with recons jobs up to Saturday - solved with a s/w fix. Transfers working fine. Average of 600MB/s overall. Quality pretty good apart from FNAL: 50% - investigating. ASGC: now ok; were problems due to v high load on SRM v2 server. CNAF In downtime tomorrow. Lemon 2nd-ary DB down - doesn't seem to affect anything...
- Sophie - new node in production for WMS service. Informed CMS but jobs still going to wms107. New server is wms202(?) - used but still many jobs are on 107. Andrea will check (possible one is assigned to MC, one to analysis for ex;)
- LHCb (Roberto) - apart from chaotic user activity schduled activites ~zero. Activities announced last week still waiting for SAM to upgrade to latest version of GAUSS at all sites cannot start new MC activities. Problem LHCb side. How to cope with some missing sites? Being discussed... If application missing in shared area install on WN.
- NL-T1 (JT) - also no jobs from ALICE! Q: Steve - who do you have jobs from? A: nobody - been that way since Tuesday.

Sites round table:

- NIKHEF - post-mortem of post mortem NL-T1 power outage on Saturday 4th October ~16:00.

Services round table:

- AFS UI - Sophie - will reconfigure. Will add new WMS node for LHCb and new one for ALICE. Brand new m/cs.
- VOMS - Steve - a pilot service is now available - what will be next version of s/w. Built against production DB - R/O access to DB. Simone: does this allow querying of e-mails? VOMS pilot twiki page. :
  - ◆ the voms core pilot is now running with the next version of the software to be deployed sometime.
  - ◆ It's running against the production database (it's a read only service (enforced by oracle as well)).
  - ◆ It's clearly important that this gets as much production like testing as possible.
  - ◆ Connection details: <https://twiki.cern.ch/twiki/bin/view/LCG/VomsPilot>
  - ◆ I'm in touch with US colleges so they can test GUMS and other things which is where it all went horribly wrong last time.

AOB:

- LHC inauguration tomorrow - Harry will run meeting(!) <https://lhc2008.web.cern.ch/LHC2008/>

## Tuesday:

Attendance: local(Harry, Gavin, Roberto,Sophie);remote(Gareth,Jeff).

elog review:

Experiments round table:

Monday:

LHCb (RS): Still running fake MC production. Alignment production tests yesterday failed with a bookeeping problem. Now the longer term plans during the LHC shutdown as decided at a meeting this morning. LHCb will run an activity FEST2009, Full Experiment System tests 2009. All components of the offline and online will be tested using realistic simulated events coming out of the HLT at 2 KHz. Data quality shifters will validate as if these were real data. Generation of the required 100M events will start in about a week from now and require several weeks at the T1 and T2. Starting next February there will be weekly functional tests leading to the full exercise in March 2009.

CMS (AS): Nothing special to report.

Sites round table: Jeff (who joined late) asked when Nikhef would get some new work and was told 1 week for the LHCb MC jobs.

Services round table: The Post-Mortem of the recent degradation in the LSF shares seen by LHCb has now been published by FIO group at <https://twiki.cern.ch/twiki/bin/view/FIOgroup/PostMortem20081013>

AOB: Roberto said he has put in a Remedy ticket to ask what made the burst of high LHCb export traffic from Castor disk servers overnight. Sophie will look into this.

## Wednesday

Attendance: local(Simone, Maria, Harry, Jamie, Jean-Philippe, Patrica, Roberto, Nick, Gav);remote(Gareth, JT, Gonzalo).

elog review:

Experiments round table:

- ATLAS (Simone) - situation relatively smooth. Yesterday and part of day before problem with ASGC - fixed this morning. Transfers now with high efficiency. (dest error from srm - no explanation or indeed confirmation problem fixed - see below!). A few things being tested or will start soon: after FTS on SLC4 modified to fix bug - Gav: now deployed at CERN on SLC4 production so expts can check - new CASTOR-SRM i/f in PPS that Jan configured for test - tests will start this week. In ATLAS DM as 'any other destination site'. Also new "CASTOR for T3" with xrootd enabled. 100TB (50 for now?) for local users at CERN. Will report on these tests next week. Last thing: deadline for closure of default pool being discussed. No deadline yet...
- LHCb (Roberto) - glexec problem testing helped understand bug in DIRAC framework - see <http://lblogbook.cern.ch/Operations/844>. Fix on DIRAC side - to be rolled out... We had yesterday in data cleanup activity some issues at CNAF. Was in scheduled downtime until 12:00 but problems still there later in afternoon. Raised ticket asking to extend downtime. Ongoing MC activity for T2 - no jobs for T1s, including NIKHEF. A few pending DC06 MC to accomplish. Particle gun and alignment starting now. VOM(R)S meeting this morning: asked to provide R/O instance at VOMS - most likely at CNAF - using 3D?
- ALICE (Patricia) - wanted to move to WMS at majority of sites and migrate from RBs. NIKHEF/SARA and IN2P3 don't have these services yet... No news from IN2P3. Being installed at NIKHEF on integration testbed. sysadmin team surprised - first time they had heard of this. Where / when was this announced? A: in this meeting a couple of weeks ago. Certain security issues with RBs. Also during weekly ops meeting ~2 weeks ago. Nick: any response from sites with CREAM CE. A: just some small T2s who wish to provide CREAM CE for ALICE. ALICE would prefer in some larger sites -> weekly OPS agenda.

Sites round table:

Tuesday:

- NL-T1: might fail SAM test or two - rogue biomed user. Load on b/e fileserver ~150. Jobs going through - working on it right now. Simone - looked at status of production: MC production v low except in a couple of clouds with backlog. Hence no jobs.
- NDGF announced 07:00 - 14:00 UTC on 27th October.
- ASGC - Sebastien - I can confirm (from ASGC) that the CASTOR instance was degrading this morning (your morning). This was the consequence of an accumulation of requests in the stager due to extremely high SRM activity. It was solved by Hungche and Giuseppe around 11am Geneva time

Services round table:

- DB (Maria) - streams problem on Sunday. DB from NDGF was in strange state propagation for ATLAS to T1 started to fill correction - pileup) up on Saturday LCRs not propagated and paused on Sunday morning. Resumed in operation on Monday morning when the DB at NDGF was restated. One T1 can effectively block streams to all (when LCRs pileup to extent that queue memory for them is exhausted). Not source due to down stream capture box. Raised to Oracle as an enhancement - site with problem would be detected and split & re-merged when back. Not even in Oracle 11g - down by hand e.g. when site in maintenance mode. Can live for 5 days with buffer(s). Adding automatic detection of pileup so alarm can be sent. Simone: channel to alarm T1s? A: no - will implement alarm in monitoring but cannot expect T1s to intervene at w/es. Streams currently on working hours. Splitting operation - and merging - requires a lot of manual work. Done only when window known to be >> 8 hours - e.g. 1 day or so. Power cut in CERN computer centre on 18 Nov - integration RACs will be affected (vault - hence tape robots). Simone - another intervention for robots in November - is this the same? Check. Maria - started with Oracle critical patch updates for October. integration being done, others being scheduled and then production. Have also to schedule interventions on integration clusters as some run production applications, e.g. LHCb booking on validation ATLAS integration used for online applications(!)

AOB:

## Thursday

Attendance: local(Maria, Harry, Jamie, Gavin, Patrica, Nick, Roberto, Patricia, Simone, Steve);remote(Michael, Gareth, Jeremy, Julien, JT).

elog review:

Experiments round table:

- ALICE (Patricia) - WMS request: daily meeting / weekly forum. Fabio reported that Lyon will setup a CREAM CE. Patricia able to test if required. ALICE smooth running. Q: only for ALICE? Julien - no info - sorry! Jeremy - open Q about VO box at B'ham. No pool accts on VO box at B'ham? VO box in CC at B'ham need direct mapping between users / accts. JT - could use NL-T1 solution - only allow 1 VO box admin. Patricia - 2 people are needed for holiday / conference cover. JT - true it is an open issue. Not addressed by "VO box s/w". Pool accts on VO boxes still open - take issue to weekly joint ops. Roberto - what is status of CREAM deployment? A: just FZK for now. Preparing documentation regarding ALICE requirements - once green light will circulate. RAL also planning (Nick).
- ATLAS (Simone) - Standing issues: (solved or partially) - Lyon LFC went into trouble (generic impossible to contact) midnight. Solved ~9:30 this morning. Not clear. Julien - problem with Oracle. Patch on Oracle cluster - all fine until 22:00 when cluster machine crashed. Maria - which patch? A: security. Impacted LFC, FTS & CIC portal. Maria - was this intervention announced? A: no - not announced as normally no service interruption. Maria - agreement is that "transparent interventions"

Wednesday

also announced. CERN announces this - good if others do too. Nick - IN2P3 site is obliged(!) - part of agreement - to give announcement. Simone - special category "at risk". Nick - should get into habit of announcing - move to RSS feed. 2) Taipei - phone call this morning with Jason. Error with Oracle b/e - looks like its working now. CASTOR experts at HEPiX might be able to help! 3) FTS on SLC4 now being used for CERN-Lyon channel. Should be watched - errors now reported with phase. Error reporting now as on SLC3. Green light for production? Gav - patch in cert or PPS. N) New SRM for CASTOR - also PPS - now tested CERN-CERN using new SRM e-p. Not yet in dashboard - should be included as new site. Also tested new castor instance with xrootd i/f - bring more ATLAS people into loop. Team from Markus Elsing (test application-level access). (ATHENA which uses root which accesses data via xrootd.) Migration of user data into group pool on-going - ready for shutdown of default pool. Last point is for CASTOR - hot disk server(s) still there. Following request for help from user files in default pool. Could not access via rfiio - not even 1st byte. Lemon plot for m/c 100% all in IOWAIT. High mem consumption >100 load. (How much normal?) Open tickets for this...

- LHCb (Roberto) - rampup of fest09 (acronym see yesterday). 100M events, 100K jobs, 1-2TB of output data. Going to happen now. Integrity checks: up to 300K files at CERN without any problems. From system point of view all went fine - don't know yet about results themselves. Pilot group defined at CERN doesn't work - looks like local Q not properly configured. Uli working on it. PPS instance of new castor - lhcb will run some stress tests.

Sites round table:

- ASGC - see addendum above regarding degradation seen by ATLAS at ASGC.
- ASGC - annual campus shutdown Oct 26 00:00 - 08:00 UTC
- NL-T1 (JT) - high load on dCache at SARA. gplazma has changed behaviour, checks all possible fqans. Users in many groups - can spend alot of time!
- RAL (Gareth) - heavy load on 2 machines that provide ATLAS SRM. Maybe "at risk" coming up to add more resources.

Services round table:

- DB (Maria) - Post mortem on ATLAS Streams problems over last weekend. Details on error and suggest some improvements. All problems caused by high load at NDGF - single instance small server. Asked many times ATLAS not to run stress tests on Friday evening without any warning. Noone available to look at problem, no resolution can be expected before Monday. Streams henced paused at all T1 sites with associated problems. Q: affecting other VOs? A: no - only ATLAS.

AOB:

- Clarification of what goes to weekly / daily meetings? Time scale dictates..

## Friday

Attendance: local(Eva, Sophie, Jamie, Harry, Simone, Gav, Jean-Philippe, Patricia);remote(Jeremy, Michael, Gareth, JT).

elog review:

Experiments round table:

- ATLAS (Simone) - announcements: functional tests Tx-Ty resumed. Paused last week to drain cosmics backoo. Throughput test next week. ATLAS T0 people to be contacted - has to be agreed with them. Disk space? Will check this afternoon and follow up Monday. Cosmic data taking stops Monday and schedule for next months being discussed. Status of services: ATLAS T0 dashboard showed no activity 04:30 - 09:00 - config problem with one site - solved as soon as dashboard experts. 1 h SRM problem with FZK - fixed after notification from ATLAS shifter. Problem delivering data (srmpub problem) to Great Lakes T2 - Michael problem solved since 30'. Simone - still high inefficiency from everywhere into ASGC. Degraded since quite some while now - looks like gridftp error (not srm, not Oracle) - follow up. Lyon having problem getting a few files from CERN - hot filesystem / diskserver problem? Any news? Persisting since quite a while... More aggressive followup. New SRM being tested. SRM ATLAS PPS e-p might not be in BDII - FTS doesn't know? Please publish - Sophie will follow up. Will stop transfers for now and reactivate when done. Notification problem - Graeme on call observed problem with CNAF. Notified with both team and alarm ticket. Team ticket never answered, alarm degraded to team ticket. Why not treated as alarm? Why VO not acknowledged? Site in scheduled down time - 6H after end still not fully running hence ticket. Why? Registration of files in LFCs. For many sites registering files is v. slow. DDM site services for some LFCs can register 1K files / hr - 0.3 - 0.5Hz vs expectation of 20Hz. This is true for ATLAS service. VO box overload? Try one VO box to serve 1 site - no good, try stopping central deletion etc. Still no good. Logs from TRIUMF + PIC -> JPB to look for problems. What does DDM do to register files in LFC? Bulk registration should be even faster than 20Hz. (20Hz demonstrated 1 year & half ago -> bulk methods.) Harry - some site dependancy? A: some LFCs show problem but not all. Q: are they configured in a different way? A: all deployed differently... LFC at T1 serves whole cloud. Backlog for 1 site in cloud but not others - hence maybe not LFC. Backlog now for almost all sites (some increase, some decrease...) JPB - timestamp shows LFC replies very quickly on a given request. Need log of client to see if LFC replies to client request. Simone - DDM logs. Need news from Miguel to see if all info is logged or just sparse info. Can also ask to start VO box in debug mode. JT: does it look at first FQAN or all of them? A: doesn't increase processing time much... 2 or 4 int comparisons in CPU are ~nothing. LFC doesn't use string comparison but VUID - hence extremely fast. Sophie - is CERN LFC one of the affected ones? Simone - not so far, but backlog of 1.5hours since today. If you have to restart VO box and scratch local DB lose info on all files transferred but not registered - could loose e.g. 20K transfers. JPB - can check logs at CERN from a couple of hours ago. Sophie will check. JT - request for more disk space at NL-T1? Didn't receive request yet so might not come until Monday. Disk space at SARA. Simone to follow up.
- ALICE (Patricia) - Migration of WMS. All Russian sites, CERN and Italy. All working fine. Today with Germany, NIKHEF & SARA. Q: use SARA WMS at NIKHEF? A: ok. Main sites will be done today, other sites next week. CREAM doc sent this morning to Nick and OPS teams -> sites and other teams. One issue observed last night at CERN: sgm accts not mapped correctly to accts on CE, hence jobs aborting. Found by Maarten overnight, fixed by Uli
- LHCb (Roberto) - just some MC simulation going on for some small physics groups like LSS-Beam7TeV. Joel has submitted the first test jobs for FEST09 production and depending on the results LHCb will ramp-up next week this not-negligible production. I've also to report that Andrew did open a remedy ticket two weeks ago (or so) declaring that some files in a disk server (faulty) were not accessible (space T0D1) Andrew Smith right now tells me it's the disk server lxf4302 (or something, CASTOR people know) and some data are only available there and its unavailability is delaying by two weeks the processing of these data. Could we escalate and follow closely this issue? Sophie - vendor not respecting delays for fixing machines, several disks not available. Escalated by sysadmins. Harry - did this info get reflected back into GGUS ticket and hence back to Andrew? Will check

Sites round table:

- NL-T1 (JT) - problem around time of ops meeting with fileserver - look like denial of service

attack(!) but turned out to be VDT update. Destroyed local patches to job manager -> standard PBS. Degraded for ~4 hours, now fixed.

Services round table:

- DB (Eva) - streams replication for ATLAS online-offline (ATLAS\_CONF\_TGC). Delete operation on all rows in table. Table had no indices nor primary keys. Apply server trying to apply this big transaction but bound by I/O activity in system. ATLAS DBA unable to contact person for this schema. Stopped apply and create indices, then apply went through backlog. Normally all tables involved in streams replication should have primary key and indices. Found several times that developers don't create primary keys! Simone - reported yesterday to ATLAS point about no unannounced stress tests Friday night. Agreed by those present that this is not a good idea... Should go into ATLAS OPS minutes.

AOB:

---

This topic: LCG > WLCGDailyMeetingsWeek081020

Topic revision: r11 - 2008-10-31 - SteveTraylen



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