

# Table of Contents

<b>Week of 090713.....</b>	<b>1</b>
WLCG Service Incidents, Interventions and Availability.....	1
GGUS section.....	1
Daily WLCG Operations Call details.....	1
Monday:.....	1
Tuesday:.....	3
Wednesday.....	4
Thursday.....	5
Friday.....	6

# Week of 090713

## WLCG Service Incidents, Interventions and Availability

VO Summaries of Site Availability				SIRs & Broadcasts	
<a href="#">ALICE</a>	<a href="#">ATLAS</a>	<a href="#">CMS</a>	<a href="#">LHCb</a>	WLCG Service Incident Reports	Broadcast archive

### GGUS section

- GGUS Escalation reports every Monday (used for WLCG Service Report to MB)
  - ◆ Full search: [https://gus.fzk.de/ws/ticket\\_search.php](https://gus.fzk.de/ws/ticket_search.php) and select "VO" - this will give all tickets, including team & alarm
- LHCOPN Tickets in GGUS: see [https://gus.fzk.de/pages/all\\_lhcopn.php](https://gus.fzk.de/pages/all_lhcopn.php) and change your selection criteria. Future actions are also listed.
  - ◆ If network group participation is necessary, please invite them in time.
- Procedure to become a LHC Experiment VO TEAM member
- Other recent GGUS FAQs

### 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			
CERN IT status board	M/W PPSCoordinationWorkLog	WLCG Baseline Versions	Weekly joint operations meeting minutes

Additional Material:

<a href="#">STEP09</a>	<a href="#">ATLAS</a>	<a href="#">ATLAS logbook</a>	<a href="#">CMS WLCG Blogs</a>
------------------------	-----------------------	-------------------------------	--------------------------------

### Monday:

Attendance: local(Eva, Jamie, MariaG, Ueda, Andrea, JeanPhilippe, Gang, David, Julia, Harry, Roberto, Simone, Alessandro, Patricia, Jan );remote(Gareth, Brian, Angela, Alessandro).

Experiments round table:

- ATLAS (Ueda and Simone) : Problems reported at the weekend on raw datafiles not being accessible. Also a problem on Sunday morning from 10.27 to 10.33 of connection lost from the central catalogs to ATLR to be further investigated by the Physics Database Services. It appears as sites upgrade WNs to latest glite version, they start failing ATLAS production. The failure is when the ATLAS distributed python 2.5 imports the logging module Details in [https://gus.fzk.de/ws/ticket\\_info.php?ticket=50148](https://gus.fzk.de/ws/ticket_info.php?ticket=50148). There is a workaround available to the sites, as described in the ticket. This morning there has been a problem of high load caused by DQ2 on ATLR, now fixed and due to a developer's error. A post-mortem is being produced. Problems with CASTOR DB backend at CNAF occurred at the weekend. Might be ok now, but cannot be confirmed as there is currently a downtime of the local LFC (possibly due to a network problem - under investigation).

- CMS reports - No report.
- ALICE (Patricia) Production is currently on standby due the set-up of the new MC cycle. There is also a new module for job submission for WMS, which might bring in some initial inefficiencies. VOboxes will also go out of warranty in Q4 2009. Reviewing with FIO for their replacement (2 in production and 2 for standby).
- LHCb(Roberto) reports - Currently running 10,5K jobs between user private analysis (~1K) and MC09 production. The reprocessing activity announced last week is still awaiting for all sites to remove data from the cache. About 15K jobs pending the green light to be submitted to sites for testing the staging and file access at T1's. Problems at CNAF for CASTOR at the weekend (see GGUS ticket) and at FZK for shared are with problems on some of the WN. Also on T2s a wrong BDII publication altering the computation of the rank, shared are issues and too many pilots aborting.

Sites / Services round table:

- RAL (Gareth): The problem on the ATLAS CASTOR reported last wednesday was fixed in the early afternoon and was due to a RAC node crash. Today the Oracle "BigID" patch was applied. Tomorrow a schedule CASTOR upgrade to 2.1.7-27 will take place from 5am to 4pm UTC, affecting all VOs. Question raised by Simone on the 2 disk servers which were not available last week at the time of the computer room move are now available again. Because of the intervention foreseen tomorrow, ATLAS will include RAL as of tomorrow evening for getting the fast reprocessing ESDs.
- CNAF (Alessandro): LFC problem at the moment being investigated and probably due to a network problem.
- ASGC (Gang): ATLAS reprocessing agenda to be published shortly.
- FZK (Angela): Nothing to report.
- NL-T1 (by email):

] Nikhef moving grid infrastructure to new data center

A new data center has been built at Nikhef. The existing grid infrastructure at Nikhef will be moved to this new data center between 10 August and 21 August. During the migration process, grid services will be unavailable.

This large-scale operation will take place in two phases:

1) Moving grid services and network infrastructure (10-14 August 2009)  
During this phase, all grid services at site NIKHEF-ELPROD will be unavailable.

For grid users this means that the following services cannot be used:

- Computing services (CEs gazon.nikhef.nl and trekker.nikhef.nl);
- Storage services (SE tbn18.nikhef.nl);
- Job submission services (WMS graszode.nikhef.nl, graspol.nikhef.nl, dorsvlegel.nikhef.nl);
- Requesting renewal of grid certificates via the Dutchgrid CA web site will not be possible at 10 and 11 August (requests can be submitted via mail but will not be processed);
- The web sites www.dutchgrid.nl, www.vl-e.nl and poc.vl-e.nl will be unavailable.

2) Moving compute and storage clusters (15-21 August 2009)

In this phase, the computing and storage clusters will be unavailable.

Grid users will not be able to:

- Use the computing services (CEs gazon.nikhef.nl and trekker.nikhef.nl);
- Access certain data files via SE tbn18.nikhef.nl.

Monday:

We advise all users of the grid infrastructure to:

- Request renewal of grid certificates before August 5th (only if the certificate will expire early or mid-August);
- Use the grid computing services at SARA (CE ce.gina.sara.nl)
- Submit grid jobs via the WMS at SARA (WMS wms.grid.sara.nl)
- Plan their work such, that no access is required to data files via storage element tbn18.nikhef.nl, or to copy relevant data elsewhere.

Kind regards,  
Ronald Starink

- CERN Services round table
- ♦ Databases (Eva): Migration of some production DBs (LHCb, CMS and WLCG) to RHEL5 is announced and scheduled in three weeks, after the successful migration of the validation clusters. Also three more schemas added to the ATLAS PVSS replication from the online to the offline DB.

AOB:

## Tuesday:

Attendance: local(MariaG, Eva, Jan, Jamie, Ueda, Simone, Harry, Jamie, Gang, Julia, Alessandro, Roberto, Andrea);remote(Ronald, Jeremy, Michael, Angela,Gareth, Luca).

Experiments round table:

- ATLAS -Reprocessing of cosmics is progressing at a good pace. 90% has already been done. The plan is to complete it by July 16th. Problems at CNAF reported yesterday on CASTOR and LFC are solved. Also, the problem reported yesterday on sites upgrade WNs to latest glite version causing job crash is now understood and the workaround published on GGUS (ticket number 50148) and Savannah. (see EGEE broadcast <https://cic.gridops.org/index.php?section=vo&page=broadcastretrieval&step=2&typeb=C&idbroadcast=41983>)
- CMS reports - No report.
- ALICE - No report.
- LHCb reports - LHCb will restart reprocessing tomorrow afternoon after the CERN CASTOR intervention will finish. Special care to transfer data back to NIKHEF (where it had accidentally be scratched also from tape). Asking now SARA to go again through the cache to remove data. MC09 proceeding well with 11k concurrent jobs running. Three open GGUS tickets (the one reported yesterday for a stager problem at CNAF being reopened) and DST transfers to MC-DST space at PIC failing.

Sites / Services round table:

- NIKHEF (Ronald): WN capacity will be ramped down by 30% because of some cooling problems that will be solved only with the move to the new Computer Center (see Monday's report). The full pledge will be reached by Oct/Nov 2009.
- BNL (Michael): The HPSS is currently being upgraded (just started) and will be off line for the entire business day (till 5pm local time). During this window, no data transfers into or out of the system will be possible.

Tuesday:

- FZK (Angela): Since Sunday evening one of the pair of NFS servers serving the ATLAS software repository is suffering from a high load. This causes ATLAS repro jobs to fail. Trying to identify the reason for this but possibly some worker node(s) are in a loop causing excessive access. At the moment the server with the high load is no longer selected by the WNs. ATLAS will try to get the fast reprocessing going using the remaining server. One of the tape libraries shows errors contacting drives. We may have to restart the library control to try to fix this. The assembly of a dedicated 'tape stager' cluster is progressing. Early next week we can continue recall tests using the new staggers. We have observed a sustained 100MB/s for CMS recalls (using 2 LTO drives) and 300-350MB/s for ATLAS recalls (using 6 LTO drives). The new staggers will provide redundancy and an increased bandwidth. Planned intervention on tape as announced in CIC site report for this Thursday (16th July) has to be postponed to next week Tuesday (21st July). We will change some tapes and add two tape drives.
- RAL (Gareth): Currently in scheduled intervention for the CASTOR upgrade to 2.1.7.27. Should finish within announced time window. Next Monday there will be an intervention affecting the ATLAS LFC to separate its database backend from other non-ATLAS ones.
- CNAF (Luca): Still investigating the root cause of yesterday's ATLAS LFC outage (possibly due to network or exhausting DB number of concurrent connections).
- ASGC (Gang): Just out of a scheduled CE downtime of a few hours for a upgrade.

\* FIO Operations (Jan): Intervention to upgrade CASTORPUBLIC to version 2.1.8-9. This will affect all CASTOR users from 9:00 to 13:00. Upgrades on CASTOR tape service. All CASTOR users affected from 8:00 to 17:00.

- Databases (Eva): ATLR services have suffered DB service stop with session kill for several connected session on Sunday 12th at 10:26. Full service connectivity was restored at 10:28. A second similar issue happened on 10:32 and full service connectivity was restored on 10:36. Applications with 'automatic reconnect' should have experienced 2 glitches in the time window mentioned. The event was triggered by what should have been a transparent change to fix an error message from the monitoring system. Namely the flash recovery area was filling up due to DB growth and the parameter `db_recovery_file_dest_size` needed to be increased accordingly. An erroneous setting of the parameter by the DBA triggered Oracle to perform service relocation as described above. The parameter has been set correctly at 10:36 by the DBA and services restarted and made fully available since then. Yesterday afternoon the streams replication from the ATLAS production to Tier1 sites affected 8 sites to a deadlock triggered by a code executed by ATLAS people at CERN in order to migrate some COOL schemas. Problem understood and solved. Also, this morning the apply processes at BNL, IN2P3 and ASGC were blocked for the apply status. The problem seemed to be caused by the apply servers blocking themselves. In order to fix it, the apply processes were stopped, changed the apply parallelism to 1 and restarted. The apply was able to resume.

AOB:

## Wednesday

Attendance: local(Ueda, MariaG, Jamie, Harry, Simone, Alessandro, Julia, David, Miguel, Roberto, Gang, Andrea, Nick) ;remote(Michael, Angela, John, Luca).

Experiments round table:

- ATLAS - File access to CASTOR problems reported on Monday solved. Some problems with RAL GridMap files reported (see RAL report below). Also reported and under investigation with CNAF some instabilities in accessing the LFC (see CNAF report).

- CMS reports - No report.
- ALICE (reported by Patricia before the meeting)- Small issues regarding the latest (pilot) version of the specific experiment submission module is causing some inefficiencies in the submission procedure through one of the ALICE VOBOXES at CERN. The situation was expected since this new module is still on testing phase. The current ALICE VOBOXES will be out of warranty at the end of the year, the experiment has asked for 4 new VOBOXES and the support will by this week cover the configuration requirements with the FIO responsible.
- LHCb reports - Reprocessing will be re-established this afternoon, with tests on all Tier1s with small job samples. MC09 proceeding smoothly (with 8k simultaneous jobs) + the usual 1K for analysis. Open still a GGUS ticket with CNAF for wrong BDII publication while the problem at PIC for DST transfers is solved.

Sites / Services round table:

- BNL (Michael): The HPSS upgrade went well (although it took 4 hours more than announced: format conversion took longer than expected and small hardware fault). Currently experiencing some timeouts when staging out of a SE (a massive deletion is currently in progress, with some db locking issues with the scheduling).
- RAL (John): two disk server failures overnight which caused some file transfer problems. Now back in production. Also an issue with GridMap files still under investigation as they seem to be still seen by ATLAS. To be checked.
- FZK (Angela): Still investigating the NFS problem reported yesterday.
- CNAF (Luca): ongoing investigations of the LFC problem that occurred over the weekend. Two network glitches of one minute today between midday and 2pm might explain the LFC instability reported by ATLAS.
- ASGC (Gang): CMS have cleaned out their jobs and CE upgrade is being scheduled.
- Services:
  - FIO Ops: Successful upgrade of CASTORPUBLIC (intervention finished on time). Some clarification is required for the availability calculations from SAM tests for the LHC experiments (which are not affected by CASTORPUBLIC, but the OPS VO is). Ongoing tape service upgrade.
  - Databases: nothing to report.
  - gLite release issues: see EGEE broadcast concerning the workaround for the recent glite upgrade <https://cic.gridops.org/index.php?section=vo&page=broadcastretrieval&step=2&typeb=C&idbroadcast=41983>

AOB:

## Thursday

Attendance: local(Jan, Jamie, Roberto, Eva, MariaG, Simone, Andrea, Alessandro );remote(Luca, John, Angela, Michael, Jeremy, Daniele).

Experiments round table:

- ATLAS (Simone) Fast reprocessing is almost complete - 99.7% done. Still ~400 batch jobs to be done at NDGF. In touch with them. Confirmed that LFC problems at CNAF coincided with network

glitches reported. Central catalog upgrade postponed to this afternoon for 1 server this afternoon and 2nd on Monday (to be installed with latest version). This will decrease load on ATLR due to optimized DQ2 version.

- CMS reports - Apologies for absences due to exams. For T1s: Transfers to CNAF failing for proxy expiration. T2s: Phedex agents down at a T2 in Pakistan and Bristol. Jobs running for a long time of more than 20 hours under investigation at Pisa. Investigating a problem at Purdue with job aborting and a debug link from the UK SouthGrid RALPP to US UMD T3. Full details in report. Question from Jan regarding CASTORCMS interruption 05:50 for 10' - high load - (no impact seen by experiment)
- ALICE - no report
- LHCb reports - reprocessing restarted yesterday with total of 2K concurrent jobs at T1s, 8K for MC, 3K for analysis: total 13K. Quite impressive number! [ Full details in report ]. Problem observed with CNAF and RAL with high # of jobs stalling. Discussion on configuration: CNAF: limit on # of LSF slots. LHCb should try to reduce load at least for time being. RAL confirms that they had to increase # of root connections from 200 -> 400 thus solving this problem. CNAF will investigate. Problems with transfers to GridKA to be fully explained in new GGUS ticket.

Sites / Services round table:

- BNL (Michael) - Mass deletion of files from obsolete d/s. File system ran out of inodes. Configured for 1M inodes – caused pnfs latencies and timeouts. Created new f/s – move probably today. 100M inodes, hopefully will suffice! Simone – was this deletion via SRM or directly against pnfs? Yes, the latter. Cleanup operations performed by Hiro.
- CNAF (Luca) - see above for comments inline.
- RAL (John) - ditto; confirmation that ATLAS file transfer problem was fixed yesterday.
- FZK (Angela) - data access problems reported for a server with with disk only files. Now fixed. LHCb problem - required GGUS ticket as above.
- ASGC (Gang) - progress on CE: will be ready for ATLAS later today (TBC)
- Services: tape upgrade yesterday went fine, finished around 17:30. Glitch reported re CASTORCMS turned out to be 'transparent'.

AOB:

## Friday

Attendance: local(Jamie, Maria, Luca, Simone, Jan, Miguel, Gang, Alessandro, Jean-Philippe, Graeme, Patricia);remote(Joel, Daniele, Michael, Gareth, Angela, Brian, Tiji, Jeremy)

Experiments round table:

- ATLAS (Simone and Graeme) - Announcements about glitches with dCache; tests for prestaging and reprocessing at ASGC resumed but immediately stopped – error “ No default service classes defined” and the GGUS ticket is 50363”. Being investigated. Also investigating some corrupted files produced in February at NDGF.

- CMS reports - (Daniele) Most of problems mentioned yesterday were understood and solved (CNAF transfer problems, PhEDEx agents down at Pakistan NCP and Bristol, problems at Pisa with jobs running very slowly and Budapest for SAM CE analysis test). Still investigating on two problems at London QMUL and Purdue.
- ALICE (Patricia) Production running smoothly. Expecting instabilities at the weekend for VOBoxes upgrades, including the WMS 3.2 latest release at CERN to be tested in production also this weekend. A new Alien version might be coming soon. More news next week. From Latchezar: The AliEn central services servers will be shut down and relocated to new rack and power infrastructure on Tuesday 21 July from 09:00 to 17:00 CEST. No production will be running, the user access and analysis will also be closed.
- LHCb reports (Joel): 20K jobs concurrently running yesterday (a new record)! IN2P3 transfers completed and migrated to MSS this morning. Problems with the limited number of disk server connections reported yesterday by CNAF and RAL.

Sites / Services round table:

\* BNL (Michael): 1M files deleted in one chunk & resulted in numerous problems, 1 reported yesterday re trash info in dCache – seen subsequent problems with pnfs manager. Still sluggish yesterday morning. Needed more rigorous cleanup -> unsched. downtime. Work has taken longer than foreseen (4h rather than 1h). Some disk servers had problems coming back up. At about 2:00 Eastern “vacuum” process kicks in and removes obsolete entries. A lot of vacuuming hence v high load -> stuck. Lesson learned: deletion should not be done in such large quantities! (<<1M) Will look at improving i/o performance of DB. Use SSD instead of mag disk ?? JPB – do you still see LFC probs? A: no – a 1 off. Michael – attaching 2PB of disk to system; some preparations done during yesterday’s downtime so no more downtime expected due to this addition. Last but not least: data integrity checking in US. Also at T2s.

\* FZK (Angela): probs with one of tape libs. Enabling drives makes others become unavailable! 8 drives available, trying to get another 6 before w/e.

\* RAL (Gareth) ATLAS LFC intervention announced for Monday postponed to 27th July.

\* ASCG (Gang): from point of view of expts due to power cut earlier this morning – 3” – all services down as no UPS. Don’t want to use big UPS as big problems! 2nd power route will be established. Backup power in 50ms.

\* Services:

\* FIO OPS (Miguel and Jan): transparent intervention on CERN Castor DB network switches announced on IT Status Board (and entered on GoCDB) at 10am Monday and Tuesday 20th and 21st July.

\* Databases for Physics (Luca): Propagation to ASGC stopped for one and half hour this morning following the site power cut. Announced on IT Service Status Board and to the experiments the CMSR and LCGR RAC clusters migrations to RHEL5 for 20th and 22nd July respectively. AOB:

-- JamieShiers - 10 Jul 2009

---

This topic: LCG > WLCGDailyMeetingsWeek090713

Topic revision: r17 - 2009-07-17 - MariaGirone



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