

# Table of Contents

<b>Week of 090309.....</b>	<b>1</b>
WLCG Baseline Versions.....	1
WLCG Service Incident Reports.....	1
GGUS Team / Alarm Tickets during last week.....	1
Weekly VO Summaries of Site Availability.....	1
Daily WLCG Operations Call details.....	1
General Information.....	1
Monday:.....	1
Tuesday:.....	2
Wednesday.....	4
Thursday.....	5
Friday.....	6

# Week of 090309

## WLCG Baseline Versions

## WLCG Service Incident Reports

- This section lists WLCG Service Incident Reports from the previous weeks (new or updated only).
- Status can be received, due or open
  
- INFN (Luca) - post-mortem on network incident at CNAF:  
[https://twiki.cern.ch/twiki/pub/LCG/WLCGDailyMeetingsWeek090309/network\\_incident\\_\\_INFN-T1\\_\\_a\\_pos](https://twiki.cern.ch/twiki/pub/LCG/WLCGDailyMeetingsWeek090309/network_incident__INFN-T1__a_pos)

## GGUS Team / Alarm Tickets during last week

- Tickets: see [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 - or use the following link:
  - ◆ In the GGUS Escalation reports every Monday
- Please observe the procedure to become a LHC Experiment VO TEAM member and other most recent GGUS FAQs.

## Weekly VO Summaries of Site Availability

- ALICE
- ATLAS
- CMS
- LHCb

## 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(Jean-Philippe, Julia, Harry, Alessandro, Ricardo, Roberto, MariaDZ);remote(Gareth, Daniele).

Experiments round table:

- ATLAS - Several problems during the weekend. 1) U.Toronto site is still giving intermittent dcache permission denied errors to store files since more than a week. ATLAS suggest they talk to Triumf. 2) There are srm problems to export data to Coimbra and Lisbon and both have been taken out of production. 3) FZK is down for 3 days so the german cloud is out of production. 4) Cloud validation tests were started last Friday as a precursor to starting the planned reprocessing.
- CMS reports - There is a summary of last weeks GGUS Alarm ticket tests showing response times on the reports Twiki. Response times were quite good with no showstoppers on reaching Tier 1 sites.
- ALICE -
- LHCb - 1) GGUS alarm tickets specialised for LHCb use cases will be tested this afternoon. 2) Dummy MonteCarlo production will restart using an older version of Gauss (the current one needs modifying to allow for jobs that run out of time). 3) Tier 1 reconstruction jobs are also restarting but at a low level. These have already shown several issues e.g. the issues at Nikhef trying to reconnect to the conditions database. 4) On the issue of wrong file locality being returned by the srm at IN2P3 they have followed advice from Lionel Schwarz but still have a problem.

Sites / Services round table:

CERN (by email from T.Cass): CMS has asked to postpone the already announced network intervention scheduled for the 18th of March:

<http://it-support-servicestatus.web.cern.ch/it-support-servicestatus/ServiceChangeArchive/ImportantDisruptiveNetwork>

The new proposed date is Wednesday 1st of April 2009, with the same time schedule. Please let us know before tomorrow evening (tuesday 9th/3) if anyone opposes to this new date. In this case, the original date of the 18th/3 will be kept.

TRIUMF (mail from Di Qing following up on FTS 'defective credentials' error): Sorry for the late reply. We tried to investigate it on our SE. We found the same error message in the log of gridftp door on several pool nodes. The errors already appeared one hour earlier than when we reactivated all FTS channels, so definitely it was not a problem of the FTS. However, we still have not clues why this happened.

AOB: GGUS Team Tickets (MariaDZ): Experiments are reminded there is a well documented procedure for adding members to the team ticket submitters lists and she has added the link into the GGUS section above.

RAL (Gareth): RAL and other Tier 1 sites were marked as unavailable following CE SAM test failures from about 15.00 on Saturday. Ricardo said they were aware of this and had an open GGUS ticket.

## Tuesday:

Attendance: local(Alessandro, Simone, Kors, Harry, Riccardo, Julia, Jean-Philippe, Nick, Olof, Patricia, Roberto, JT, Maria), remote(Daniele, Michael, Michel, Gareth).

Experiments round table:

- ATLAS (Simone) - ATLAS observed problems in 4 T1s: RAL started unsched down yesterday, sched today - DNS problems. Lyon in sched down, ASGC too, FZK in sched down until tomorrow too! 40% of ATLAS T1! IFIC site in Spain - STORM + Lustre - need to check with STORM dev so that they can publish their info correctly and use them again!. Reprocessing campaign: pre-campaign site validation on-going. Hope validation would complete this week but let's see. Tomorrow another downtime at CNAF. The ADC review starts tomorrow. JT - timeout at NIKHEF accessing files? We have a couple of tickets re stagein of files to WNs and timeouts. We have 2Gbit network between NIKHEF and SARA - upgraded next month to 10Gb. Could limit # ATLAS jobs so that limit is not

Monday:

reached, wait, ban site or ... Simone: which files? Propose that 'hot files' are replicated to DPM in NIKHEF. Then normal input files would not congest link - will check with Hong. RAL (Gareth) - at RAL had problems since middle of yesterday related to DNS problem here. Put site into "AT RISK" which got extended through night - marked as scheduled but clearly not! CASTOR SRM for particularly ATLAS had quite severe problems - maybe should have been in unsched out. These resolved this morning and we believe we are back fully now.

- CMS reports (Daniele) - quite some details in today's CMS report (see inline wiki:[https://twiki.cern.ch/twiki/bin/view/CMS/FacOps\\_WLCGdailyreports#10\\_March\\_2009\\_Tuesday](https://twiki.cern.ch/twiki/bin/view/CMS/FacOps_WLCGdailyreports#10_March_2009_Tuesday)): Have opened about 7 CMS internal Savannah tickets, nothing yet forked to GGUS. T0: thanks for full CASTOR p-m. Tape migration - in consequence of above - now fixed. T1s: IN2P3 export errors to Fr site, fixed, closed. CNAF After some interventions on StoRM and PhEDEx at CNAF the transfers started with ~100% quality and ~10MB/s. The average rate of the last 24h is ~10MB/s so the link could be commissioned. (Savannah #107382 --> CLOSED) ; IN2P3 -> FNAL transfer issue - item fixed & closed. New tickets against FNAL, CNAF, IN2P3. Existing open tickets to CNAF & FZK. T2s: close ticket to IC re transfers to RAL; some new: KNU in Korea; Bari - problems getting files from CNAF - local storage in Bari; Pisa in downtime - heavy load on storage. Reinstall dCache to solve. MIT : problems reading from / writing to SE there. GGUS->OSG [https://gus.fzk.de/ws/ticket\\_info.php?ticket=46798](https://gus.fzk.de/ws/ticket_info.php?ticket=46798) reply last night. Reply on March 9. Assigned March 2 to affected site MIT\_CMS. No reaction for 1 week.
- ALICE (Patricia) - WMS issues: working quite well after upgrade to latest megapatch. About 3600 jobs running - detailed report to GDB tomorrow. Working hard with CREAM. Several sites working very well. Best FZK, RAL working well in the past. Putting CERN into production. 2 CEs only accessible from CERN - ce201, ce202. Small problem with info provider for ce201. First test : CREAM bug 47152 affecting CERN - in contact with developers to fix. Workaround for this problem sent to CERN; same problem seen at RAL & same workaround. Ricardo - these CEs really not production ready yet - one was reinstalled last night.
- LHCb (Roberto) - continuing T1 commissioning activity by running recons jobs using fest data. RAL DNS issue also prevented recons jobs to resolve TURL using SRM e-p at RAL. GridKA down - also IN2P3. "Wrong locality" reported by SRM at Lyon - suggest to contact SARA who had problem and fixed. NIKHEF: reconnection to conditions DB still under investigation. CERN: LHCb user space token - globus xi/o eof error. Disk server exhausted # gridftp servers running / idle. Site admin has to intervene to cleanup. Re-used GGUS ticket opened a couple of weeks ago to find longer-term solution rather than regular cleanup. glxexec tests in pps/production. Started with NIKHEF: tests failing due to LCMAPS failure - mail to mailing list. Ran GGUS tests yesterday - ok except CNAF (not recognised as valid alarmer - checking cert); problem with GGUS i/f. NL-T1: which of 2 sites to send alarm to in GGUS i/f but ticket sent to one specific site. Maria - concerning comment about NIKHEF/SARA: open ticket in Savannah - a solution foreseen for March release.

#### Sites / Services round table:

- RAL - see above under ATLAS. [ Had significant problems for 24h - will do our own brief pm - first steps on disaster recovery plan as situation went on too long! Severely degraded all night and good chunk of yesterday and this morning. One thing to add - sorry for confusion over scheduled at-risk! ]
- GRIF (Michel) - very severe cooling problem at end of last week - lost 2K jobs or so! Power cut unexpected during 2 consecutive nights.
- NL-T1 (JT) - where to raise where one address or 3? -> MB

AOB:

Tuesday:

## Wednesday

Attendance: local(Alessandro, MariaG, MariaDZ, Jean-Philippe, Jamie, Ricardo);remote(Daniele, John).

Experiments round table:

- ATLAS (Alessandro) - Observed unsched downtime in Lyon - power? Ended a couple of hours ago. Unsched->sched downtime in RAL finished this morning too. -> ATLAS operations going better but still FZK & ASGC down. Disk-server broken in INFN-Roma1 - site admin advises have to wait for repair. 200 "lost files" - hopefully available again.
- CMS reports - Top 3 issues: 1) 1 file not migrating to tape at CERN since ~5 days. 2) quite a few sites recovering from storage issues, e.g. Pisa & Bari. Pisa: reporting highload on dCache headnode -> reinstallation & restart of Phedex. Transfers now going ok but not yet to IN2P3. Bari: not major storage but h/w on specific diskserver -> data moving back to server now h/w intervention over. 3) requests for invalidating files, e.g. CNAF. Maybe inaccessible or corrupted. Possible corruption or loss of full load test sample in Nebraska. All additional details in [https://twiki.cern.ch/twiki/bin/view/CMS/FacOps\\_WLCGdailyreports#11\\_March\\_2009\\_Wednesday](https://twiki.cern.ch/twiki/bin/view/CMS/FacOps_WLCGdailyreports#11_March_2009_Wednesday).
- ALICE -
- LHCb (Roberto)
  1. Issue with the new French CA whose certificate does not appear under the usual grid-security path on the AFS UI preventing some French users that have renewed their certificates (or registered in the LHCb VO) before the change was done by Steve and Pierre. UI deployment responsables are working on that and Steve is checking whether their procedures are updated for next time this happens. [MariaDZ - known problem - they should open a GGUS ticket. AFS UI has to get their CAs reinstalled. ] Ricardo - we just got ticket from Steve T to do this proposed action - will follow up.
  2. the issue of hung processes in gridftp castor servers (globus\_xio problems) is still under a deep investigation from Shaun and Giuseppe Lopresti. This seems to be correlated with clients abruptly killed when timing out and leaving then processes on the server in a CLOSE\_WAIT status and sitting there forever (then exhausting the number of threads).

Sites / Services round table:

- INFN (Luca) - post-mortem on network incident at CNAF:  
[https://twiki.cern.ch/twiki/pub/LCG/WLCGDailyMeetingsWeek090309/network\\_incident\\_\\_INFN-T1\\_\\_a\\_pos](https://twiki.cern.ch/twiki/pub/LCG/WLCGDailyMeetingsWeek090309/network_incident__INFN-T1__a_pos)
- ASGC (Suijian Zhou) - FTS is back at ASGC, and also voms, ui, BDII etc. The recovery of castor still need some time while we are trying to speed it up. See GDB agenda [☞](#) for a report on the incident.
- DB (Maria) - streams propagation stopped for Lyon site this morning for 2 hours to allow for planned intervention for Oracle patch. Unplanned one yesterday until 16:00 caused by intervention on electrical power also Lyon. Streams was restarted ~16:00.
- CNAF (Daniele) - FTS on Oracle at CNAF & LFC for ATLAS has finished and all will be back in production.

AOB: (MariaDZ) Conclusions from the ALARMS' test exercise assembled in <https://savannah.cern.ch/support/index.php?105104#comment36> [☞](#). Pending action to sort out for the next round **in one month** the Dutch case in <https://savannah.cern.ch/support/?107440> [☞](#). Progress will be monitored via <https://savannah.cern.ch/support/?107452> [☞](#) Testing rules updated with dates in [https://twiki.cern.ch/twiki/bin/view/EGEE/SA1\\_USAG#Periodic\\_ALARM\\_ticket\\_testing\\_ru](https://twiki.cern.ch/twiki/bin/view/EGEE/SA1_USAG#Periodic_ALARM_ticket_testing_ru) Other

interesting story: the routing to USA T1s to be discussed tomorrow 12 March (ATLAS, CMS please attend!), details in <http://indico.cern.ch/conferenceDisplay.py?confId=54492> .

## Thursday

Attendance: local(Alessandro, Nick, Roberto, Maria, Diana, Jamie, Gavin, Harry, Ricardo, Jean-Philippe, James, Stephen, Ignacio, MariaDZ);remote(Gareth, Daniele).

Experiments round table:

- ATLAS (Alessandro) - ATLAS asks T1s to run cron to update config file of FTS server (should be done daily anyway) - FZK should send a broadcast when things are fully ready and working. Hopefully this afternoon / evening. (GridKa 's new srm for ATLAS : [atlassrm-fzk.gridka.de](mailto:atlassrm-fzk.gridka.de)). Many failures in transfers RAL->FZK yesterday. Site services stopped. Yesterday evening many Twiki pages at CERN were inaccessible around 22:00. Birger Koblitz following up. Brian@RAL: we have now updated FTS channel for FZK. We don't run cron job daily as we have cloud channels but do try to regularly check in case of new e-ps.
- CMS reports - (Daniele) - top 3 site related comments: 1) As from yesterday files not going to tapes at CERN > 6 days now. Reply from Miguel - not completely clear if understood yet, keep open. Was the path we followed for support ok? 2) On several sites trouble shooting status of specific files of a d/s. Often < 1000 files or blocks of a few hundred. Trying to understand if we are 'wasting time' or if this hints to some underlying problems. 3) Tickets tracking internally getting into discussions. 4 new Savannah tickets. Details in twiki. Only ticket to comment on was FZK (GGUS) T2 in Italy reported that most transfers fail with prepare to . Currently >10 attempts for 1 successful transfer! Full details in Twiki
- ALICE -
- LHCb (Roberto) - CERN related issues: one of WMSs seems to be stuck. Any submission attempt hangs WMS216. Some files 'temporarily lost' as problem with mother board. More info on Steve stp on new French CA not probably propagated to AFS UI. Ricardo - got ticket from Steve -> Sophie to fix and change procedure. "In the right place". Long-standing globus xi/o issue. GGUS ticket [GGUS ticket](#) Due to hanging gridftp processes on server. #streams from some sites > 1 - clients hanging and killed by wrapper. Caused gridftp processes to hang e.g. at CERN, RAL. lcg\_cp command with more than one stream should open corresponding port range. See GGUS for analysis. Ignacio - problem is that this leaves hanging processes. We reduced time for cleanup from 4 to 2 hours. Real workaround is to use internal CASTOR gridftp - only external is using the external. Olof arrives "just in time" - from site perspective problem is not with installation. gridftp transfer stuck as funny abort. After 100 or so stuck new connections are rejected - before it reaches gridftp application. Stale gridftp sessions - last thing they log is "aborting" - process remains. Something in ACK of abort doesn't work. However these servers are "attractive" to new connections... LHCb recommended to use standard internal gridftp version but are waiting for new gfal release(?) Only seen with certain sites connecting to CERN SE. Roberto - using lcg\_cp command line using multiple streams which works mostly but not always - firewalls not always configured correctly for WN callback. Brian - what expected port range normally is? Jean-Philippe - 20K - 25K (globus default). Lyon - wrong locality returned by srm - dcache provide a patch. Could be site config? (Still as before). Steve - French members of LHCb were added with new identities. Could do same for CMS... MariaDZ - do you have answer from CNAF about failed alarm? No - a suivre.

Sites / Services round table:

- RAL (Gareth) - reprocessing - is that waiting on FZK problem to be resolved? A: still doing site validation. Not clear what kind of conditions DB release will be used - one single big file or many

files... Doing tests of both..

- (Daniele) - ask to enable FQAN Tier1 access role on CERN WMS. This is the role we will use to access data on T1 sites. Asked to CNAF WMS support - same request now to CERN WMS support. Ricardo - would be nice to have ticket to WMS support. (& VO ID card)
- DB (Maria) - downtime yesterday of LHCb online cluster - accidental power cut at point 8. Power back quite soon but DB couldn't access disk arrays. Fixed by LHCb sysadmin ~16:30. Worked on bringing up DB. Took one hour due to problem with ASM disk group which had to be recreated. Known Oracle bug for which patch is deployed on all production DBs including this one! Talking to Oracle Support for a working patch...
- GridMAP (James) - new gridmap yesterday with all WLCG sites including OSG.

AOB:

Question(MariaDZ) - new MoU values - required? Can we close ticket? Network intervention moved to March 19th - has all announcement been done? Action: SMOD. Downtime in GOCDB for old schedule - must be moved.

## Friday

Attendance: local(Roberto, Patricia, Simone, Jamie, Harry, Nick, Ricardo);remote(JT, Gareth, Michael).

Experiments round table:

- ATLAS (Simone) - end of intervention in FZK yesterday. Still an issue with version of O/S (SLC4/5) - have to downgrade part of services from 5->4. Being done yesterday / today. Still problems accessing FZK - was not online last time checked. This morning realised that functional tests not running. They generate fake files using batch slots on lxbatch Tier0 ATLAS Q. Alessandro followed up with Ulrich - Qs being drained for next week's major network intervention. Agreed to have workaround to have these jobs running - for these jobs not an issue if jobs fail e.g. when network intervention starts. Now jobs running again... Noticed several times - but hard to reproduce - deleting files in CASTOR@CERN v. slow (~1 file / minute) (srmRm) under certain circumstances. Reproduced, ticket to CASTOR support - most likely problem at level of CASTOR SRM s/w. Opened ticket against developers. Hope for some follow up in coming days... Reprocessing campaign will now most likely wait until next week to start.
- CMS reports - (no report - apologies from Daniele)
- ALICE (Patricia) - again had to open ticket about WMS in France. Site: GRIF; service again very unstable - 2nd time this week! (Ticket number: ) Important that site realizes how important WMS is for ALICE production - maybe ROC should take some action on this? (Megapatch has been applied). CNAF: CREAM-CE issue as mentioned during GDB now fixed. VO BOX in IPNO - T2 in Fr - access denied. CERN: ce201 problem fixed - ce202 was in any case handling load. Submission across multiple CREAM CEs has to be done at application level. Not seen at service level(!)
- LHCb (Roberto) - [Q for Ricardo - 2 CEs at CERN for SLC5 - 128/129? Have to check.] Main point for today: LHCb activity seriously interrupted as all instances for WMS at CERN are hanging for job submission or querying. Opened ticket yesterday - all services restarted and ticket closed. Now service is worse... Has something been introduced by megapatch? Now using WMSs at Tier1s. Nick - known issues with these services, there is a workaround - has it been implemented on these nodes? Will followup with Sophie. SAM unit test failing systematically against StoRM endpoint. GGUS ticket [open](#) w/o any answer

Sites / Services round table:

AOB:

- Network outage - March 19th! There has been an EGEE broadcast. CASTOR nameserver is scheduled down on this day too. March 19 - 13h Nameserver (affects all instances, downtime, after network intervention, 3h) (from IT status board). To be clarified...

-- JamieShiers - 06 Mar 2009

---

This topic: LCG > WLCGDailyMeetingsWeek090309

Topic revision: r13 - 2009-03-13 - RobertoSantinel



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](#)