

# Table of Contents

<b>Week of 091109.....</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.....	7

# Week of 091109

## 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.
- OSG items selected from the GGUS escalation reports.
- 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](#)
3. The scod rota for the next few weeks is at [ScodRota](#)

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(Julia, Massimo, Jamie, Patricia, Gang, Graeme, Andrea, Wei, David, Olof, Jean-Philippe, Gavin, Roberto, Giuseppe);remote(Alessandro, Michael, Angela, Gareth, Brian, Gonzalo, Jason, Daniele, Ron).

Experiments round table:

- ATLAS - a big list! ASGC: failing all transfers as reported Friday - failed all w/e, started recovering Sunday night but not properly fixed until this morning. Added back into T1-T1 transfer matrix and continued func. tests. SARA started to get out of space errors Sunday morning on tape endpoint. Tape migration problem? Alarm ticket - problem resolved within a few hours. Bravo! Still worried that SARA has too small a buffer - 4.5 GB - thus any problem quickly becomes serious. Production qs at RAL offline all w/e as diskserver out of production - holding up merging of MC production. Tokyo: problem with Japanese CA - resolved Monday morning A-P time. A "bad moon" arose over CERN Sunday night - DB overload on ATLR? Instance 3 crashed - Panda server in bad state unable to accept

client callbacks. 10-12h affected production and analysis severely. This morning suffered degraded T0 exports - tracked on GGUS 53113. Looks like corrupted delegated credentials?! Destroyed and recreated and looks ok for the mo. Like to gain a better insight into problem. At same time seeing problems on CASTOR disk servers - gridftp failing to connect. Some downstream blockages in CASTOR? IN2P3: down for upgrade to golden dCache - hope this will enable tape system protection (from user recalls)! SARA reported 1.9.4 that this didn't work - critical issue for ATLAS. Andrea - tape protection was a doc bug or s/w bug? A: don't know! Ron - got reply from dCache devs - issue with SRM with tape protection. OK for gsidcap and gridftp but not srm. Will be fixed. Doesn't work in 1.9.4. 1.9.5??

- CMS reports - Some MC prod issues with gLite WMS. Global error 10 problem - still pending with condor support team. Trying bypasses... 3 WMS moved into prod to attempt to relieve pain. ASGC still not fully in game. Some updates at fac meeting later. Closed problem in transferring t0-chicago - FTS config at FNAL. Some progress with backlog of migration at RAL. Transfer links: export from CNAF-CALTECH probs and import to FNAL from TIFR. Quite some T2 issues but some closed.. Staging out problems at several sites. Some SAM job robot failures - some sites have not put site in sched down whilst fixing kernel patch. In summary quite some T2 issues, not so bad at T1. Gang-most tickets to ASGC closed. Transfers to ASGC started yesterday evening around 20:00.
- ALICE - quite relaxed w/e. T0: finished list of official DNs in myproxy server. Put VO alice 13 back in prod after kernel upgrade.
- LHCb reports - During w/e all remaining mc prod ran. Last week stripping exercise ran very smoothly. Problem with hanging connections thus fixed. For coming week no envisaged activities, some low level FEST jobs - a quiet week. CASTORLHCb issue Friday evening fixed - LSF mapping problem. (Unresponsive CASTOR). NIKHEF: some WNs misconfigured after kernel upgrade - promptly fixed. IN2p3: down for upgrade of dCache

CNAF: issue listing directories (very slow) under investigation. Olof - working on post-mortem for Friday's incident.

Sites / Services round table:

- CNAF - currently in downtime to update kernels. Qs closed yesterday - down until Thursday 12:00(!). During this period will also do CASTOR upgrade previously scheduled but postponed.
- BNL - ntr. Starting transparent o/s upgrade of ~60 storage servers.
- FZK - one of tape libs down at moment. Waiting for spare parts -expect to be fixed Wednesday. Problem with ATLAS tape read pool. Config error so SAM tests failed - fixed now.
- RAL - Still have ongoing issue with batch work ending up on wrong node. A couple of "at risks" for Oracle quarterly patches. Brian - castor diskserver with bad memory - corrupted two files - disk server out. In process of calculating checksums and comparing with ATLAS. So far haven't found any corrupted files. Will put diskserver back in production. Continue to checksum remaining files.
- PIC - ntr
- ASGC - during w/e had some tech problems with srm settings. Reversed settings. Phedex and DDM had major errors. Continued some probing. One of errors due to vo to stager mapping, Disabled vo mapping and try to use srm stager. Other errors from wrong notification in srm settings.
- NL-T1 - busy rebooting WNs and installig new kernels. NIKHEF will reboot various Grid services Wed. Hanging daemon on WNs which caused some issues. Early afternoon today problem with central router - hence NIKHEF unavailable for a whie. Being investigated. This w/e had issue with ATLAS dCache pools & tape - fixed yesterday.
- CERN: scheduled Linux upgrade ongoing - please report any problems. Delegation issue with FTS - also seen similar issue from CMS - not understood.
- DB: transparent security upgrades at RAL on Nov 11 and Nov 17 at IN2P3. BNL < 4h intervention on

Nov 12 for ATLAS conditions -> new h/w. Migrating ATLAS dashboards (2 applications) tomorrow. Next week expect to migrate CMS dashboard.

- Network: fibre cut affecting connection for PIC - both main and backup link affected. 11:30 - 17:00 last Friday. This morning outage affecting both again - power cut? Should not be stable.

AOB:

## Tuesday:

Attendance: local(Dirk, Eva, Maria, Jean-Philippe, Gang, Nicolo, Wei, Roberto, Lola, Maria, Maria, Nick, Graeme, Miguel, Olof);remote(Gareth, Elisabetta, Angela, Brian, Jason, Ronald, Jeremy, Daniele, Gonzalo).

Experiments round table:

- ATLAS (Graeme) - RAL diskserver back, queues back, few files lost. CCIN2P3 failing transfers to other T1. pools overloaded. Configuration changes applied this morning did not help. FTS issues mentioned yesterday now fixed. For CNAF: access to storage not available at the moment. Request to get a clearer message in the GOCDB downtime description and on when the storage will be accessible again.
- CMS reports (Daniele) : All tickets to ASGC closed but two on data consistency; data operations can resume. CNAF and IN2P3 data access errors reported to sites. CNAF 4 custodial files not accessible, IN2P3 1 custodial file missing to be retransferred. RAL still problems in data movement. To be checked by the site if fixed.

A number of closed tickets on T2s. For more details see the attached report.

- ALICE (Lola): A complete list of VOboxes sent to px.support to be authorized on the ALICE myproxy server.
- LHCb reports (Roberto) Two voboxes running out of warranty on 15th Nov. A request for new machines has been sent by LHCb. CNAF: problems deleting some remote directories on CASTOR (an intervention is ongoing). GridKA: problems deleting three directories over dCache instance.

Sites / Services round table:

- CNAF: Currently upgrading kernel versions on disk servers and upgrading to CASTOR version 2.1.7-19. Storage unavailable at the moment (see report by ATLAS). Needs to be clarified at the level of the GOCDB announcement.
- ASGC (Jason): Part of the new installed drives showing mechanical problem and can't access the tape system normally. except for the new 14 drives; still have 6 drives online serving recall and migration requests. Still working on the SIR. Trying to migrate the TSM backup server into the SAN zone together with the raid subsystem serving for CASTOR backup, to perform tape backup from the allocated subsystem.
- FZK (Angela): one of tape libs down at moment. Waiting for spare parts to be received on Wednesday and installed on Thursday.
- NL-T1 (Ronald): Since the upgrade of the worker nodes to CentOS 5, they suffer from frequently hanging nsd daemons. That kills the performance of the nodes and causes job failures. We are investigating the cause of these hangs and keep restarting the daemons if we detect the problem. Yesterday we added the trust relation to VOMS server lcg-voms.cern.ch. Due to an error in the configuration, many SAM tests failed last night. This issue was resolved this morning. The software

Tuesday:

area could not be mounted on the VO box for Alice as result of a bug in a recently updates package. We've applied a workaround for this.

- RAL: (Gareth): CASTOR info provider problems over night (SAM tests failing from midnight to 3am). Solved now. Site at risk for database patch upgrade for ATLAS (OGMA) LHCb (LUGH). Some instability caused by it at the server level. Data movement problems from CMS to be checked if already fixed.
- PiC (Gonzalo): Scheduled interventions for re-cabling. Storage working on degraded mode (20%). Experiments seem to cope with this. Should be over around 6pm.
- FIO: SIR for LHCb outage of last Friday submitted. At <https://twiki.cern.ch/twiki/bin/view/LCG/WLCGServiceIncidents>

AOB:

## Wednesday

Attendance: local(Jamie, Maria, Graeme, Eva, Nicolo, Gang, Simone, Jean-Philippe, Roberto, Olof);remote(Angela, Michael, Onno, Tiju, Jason, Elisabetta, Daniele).

Experiments round table:

- ATLAS (Graeme) - 4 issues: Lyon still see some severe degradation on transfers, in particular into Lyon. GGUS 53135. No news from Lyon - holiday in Fr! CNAF: storage taken out of ATLAS site services. Clarification on storage downtime has not appeared 😞 RAL: fallout from 1month ago? Substantial # of missing files in 1 d/s produced during period of significant CASTOR instability. Why did ATLAS procedures not pick these up? Working with ATLAS contacts at site for clean-up. New for today: BNL seems to have SRM down. P1 shifter put in GGUAS ticket 53106. Michael - resolved now. transfer picking up. Elisabetta: asked colleagues to update message on GOC DB - message has been updated. Reply to ticket came from ATLAS providing same explanation. Graeme - were expecting more expansive explanation on wlcg-operations list. e.g. a clear statement as to when downtime would end. Downtime is scheduled for 4 days(!) for kernel upgrade. Not acceptable during data taking. An intervention like this should be discussed with VOs.
- CMS reports (Daniele) - MC production issues & globus error 10. Some suggestions about retry count and shallow retry count - being applied now - hope to see some improvements. No reply to condor support. T1: 2 tickets to CNAF & IN2P3 custodial & not accessible were due to CMS mistakes - not properly injected in PhEDEx - not a site issue. Sites not guilty! Ticket closed. T2s: no tickets closed in last 24h. Some progress on a few tickets - details in twiki. Nicolo - new Savannah ticket - frequent quality degradation CERN to FNAL. Some stats: main SRM copy error message - gather stats and will submit ticket for CERN & FNAL to investigate.
- ALICE (Patricia before the meeting) -Helping sites to deprecate the current gLite3.1 VOBOX. The upgrade of the VOBOX has been confirmed by SARA for next Wednesday. New bunches of DN registrations into myproxy server may arrive due to this upgrade operations. CERN: voalice11 has been installed with the gLite3.2 version using Quattor. At this moment several Lemon sensor packages are being installed to complete the alarm system foreseen at the T0 for all ALICE VOBOXES. In terms of the production, the experiment is currently running about 5000 concurrent jobs with no special incidents to report. An email has been sent to the SAM experts at CERN. It seems one of the SAM UIs at CERN is not properly publishing the VOBOX sensor results. Issues triggered by the regional experts of Alice in Russia. Waiting for expert's answer.

- LHCb reports (Roberto) - Only user jobs running in the system at T1 mainly (~3K concurrently). No relevant production activities running right now. T1 issues: RAL: Observed some perturbation at RAL with timeouts accessing software on shared area. Awareness about at RISK intervention at RAL but the problem occurred outside the scheduled intervention (last night); NIKHEF: yesterday still issues with ROOT application resolving HOME apparently still related with the nscd server stuck on some WN following the OS kernel upgrade.; SARA: unbanned the CE and again jobs started to fail with some specific application error (134 Gauss). Problem affecting jobs there and few other (mainly Russian) sites. CERN: volhcb02 and volhcb06 issue: since FIO need to install new hardware that is arriving soon in the racks that these machines are currently placed, they need to quickly move these to a temporary place until their services are migrated. LHCb looking for a possible slot...

Sites / Services round table:

- FZK (Angela) : ntr
- NL-T1: 1) NIKHEF rebooted all service m/cs to upgrade kernel. WNs and UI were already upgraded before this. At SARA UI and WNs upgraded earlier too. Service nodes to be rebooted. 2) NIKHEF nscd daemon hanging - causes WN problems - under investigation, reported already yesterday 3) ATLAS has reported some jobs fail due to incomplete file transfers from SARA SRM. Caused by dCache bug - occurs when 2 jobs try to read same file from same pool node. If they end around same time dCache (sometimes) closes wrong connection. Known about by developers- awaiting a fix. Simone - problem occurs on reading? Get an incomplete file. Yes, when several jobs run at same time accessing same time. Roberto - nscd issue - should we ban farm? How many WNs affected? Plan? A: NIKHEF monitoring WNs and restarting daemon when it appears to have crashed. Don't know how many WNs affected but started after kernel upgrade. If related then all WNs affected...
- RAL: ntr
- BNL: as discussed earlier SRM problem. Probably have to upgrade to golden release - server gets into loop throwing exception. Network problems overnight caused by switch issuing broadcast storm. Probably caused by last week's network upgrade. Switch taken out and network stable again. Services stayed up but were impacted. DB migration to new hardware planned for tomorrow 12.11 for less than four hours.
- ASGC: mechanical problem of tape system due to redundant gripper. Have to shut whole system down 1-2 h and recalibrate. Maintenance should be transparent - shutdown q to avoid staging recall. 1 diskserver overloaded. Affecting for 2 hours ATLAS data transfers in early morning.
- CNAF: ntr
- CERN: 1 intervention today on myproxy-fts service - moved to new h/w and alias flipped. Should now be updated at all sites - will check and switch old off.
- DB: 3D OEM migrated to 10.2.0.5. Part of ATLAS dashboard applications migrated yesterday to LCGR. Preparing the rest. Working on apply performance problems for ATLAS PVSS data.

Release report: deployment status wiki page

AOB:

## Thursday

Attendance: local(DIRk, Jamie, Maria, Gang, Harry, Eva, Alessandro, Lola, Andrew, Julia, Nick, Jan, Simone, MariaDZ);remote(Michael, Daniele, Jason, Angela, Fabio, Ronald, Gareth, Elisabetta,Roberto).

Experiments round table:

- ATLAS - Good afternoon! ATLAS has 2 T1 problems: IN2P3 still having problems as in GGUS ticket. This ticket answered on Monday when submitted but left untouched for ~48h. Still experiencing data transfer errors T1-T1 including Lyon. Contacted T1 expert directly today... INFN: understood downtime was 4 days and involved also storage. Discussion with ATLAS computing

coordinator: critical patches should be applied but other upgrades should be discussed with experiments as we are in data taking - mail sent from Kors to T1 coordinators. Good news - dashboard DB migration went perfectly smooth - shorter than expected. During past days ATLAS has seen some errors from filename length - different instances have different limitations. Trying to find a solution - i.e. max # chars allowed for user datasets. Fabio: ticket has not been ignored - people are working on it but not an issue problem. dCache people have made some changes but this has not been enough to solve problem. All that can be said so far.. Don't see relevance of MoU here.. Ale - only understood after discussing with Ghita that people were working on it. Maria - would be use in case of problems to connect regularly to meeting and report on status of problem. Fabio - yesterday was a public holiday in France. This morning people connected to daily ATLAS meeting and are collecting more for this meeting.

- CMS reports - No tickets for ASGC ☺ All now taken care of now site back to operations. One pending was reassigned so no action pending - other from yesterday fixed and closed. RAL: backlog understood possibly related to 10TB transfer limit window in PhEDEx - once routed files to RAL go below this should be improving. Nothing more at T1s.. T2s: progress at CSCS - some errors on CEs, some CMS-specific tests since Sunday - now back to green. Estonia - massive set of errors both SAM tests & job robot errors, also presence in BDII unstable. Seems ok since yesterday 6GMT. PhEDEx components down at 4 RU T2s - all recovered. Minor progress on other tickets - will report when finished. Fabio: hit by some SAM tests failing due to CMS FroNtier server at CERN. Daniele - yes, thread on this but no ticket for now.
- ALICE - Not much to report, preparing test RPMs for voalice11 when finished will be put in production.
- LHCb reports - Just 500 jobs in the system .between user/sam and few remaining stripping. Couple of large production coming soon (50M events) T0: volhcb02 and volhcb06 issue: agreed on moving them today (14:00); T1: RAL: shared area issue under investigation (requested this morning more info); SARA: Exit code 134 (jobs killed abruptly by the system): under investigation the reason LHCb application side

#### Sites / Services round table:

- BNL: ntr (db migration t onew h/w as foreseen)
- ASGC: minor issue of FTS ASGC-BNL. Ticket opened by ATLAS. Channel agents ?? Transfer now ok. Maria: any progress on SIR for CASTOR downtime. Jason - reviewing with CASTOR support team then DB.
- KIT: tape work on-going but not yet finished - looking good. LATE NEWS - repair has just finished!
- IN2P3: deployed golden release of dCache last Monday. CMS, ALICE and LHCb only using SL5 at site.
- NL-T1: SARA: next week Wed network maintenance to upgrade storage-computing to 40Gb/s. Scheduled downtime for upgrade to LFC/FTS (migration to new h/w scheduled next week) postponed to Jan 2010. dCache problem reported yesterday - developers have found explanation and working on fix end week for 1.9.4/5. Not quite as reported yesterday but two transfers scheduled for exactly the same ms. NIKHEF: VO box migration to SL5 and gLite 3.2.
- RAL: will need to apply updated kernel to one of FTS nodes - will look for downtime early next week.
- CNAF: completed activity on kernel upgrading and CASTOR one. Have new problem with new version of CASTOR - recalls from tape run but some problems on disk server. Waiting feedback from CASTOR developers. Simone - then site should get no data. Maria - is there an entry in GOCDB? No. will put one. Roberto - can CE be put back in production? Y. [ should be a new downtime ]
- Dashboards: confirm migration of ATLAS went smoothly and fast. Last application scheduled for Monday. Scheduling CMS migration - waiting for confirmation. Preparations for migration will start tomorrow.

- DB: performance of apply for ATLAS PVSS has improved and now recovering backlog. 15h behind. Tomorrow morning will be up to date.
- FIO: SRM - last weeks performance problem for ATLAS - snapshot release that might not fix but should improve things. Can we deploy - transparent - on one of servers? Simone - waiting for full green light from inside ATLAS. SRM - ALICE & CMS to 2.8.2 - waiting for CMS (OK from ALICE) - proposed date Monday morning. Last: myproxy migration issue from last week understood - DNS time to live internal / external. Was ok. Daniele - for upgrade someone will reply but prefer Tuesday 17th. Confirmed by Daniele after the meeting on Monday 16.11 as agreed date for upgrade.

AOB:

## Friday

Attendance: local(Jamie, Alessandro, Julia, Andrew, Jean-Philippe, Patricia, Edoardo, Roberto, Eva);remote(Daniele, Angela, Gonzalo, Brian, Gareth, Michael, Fabio, Onno, Jason).

Experiments round table:

- ATLAS - CERN: problem in transfer from 05:00 - 08:00/09:00. Solved itself! Should be solved by patch to SRM scheduled for next Monday. (Avoid lock-contention). ATLAS happy to have this patch Monday as foreseen! Yesterday problem with application delivering data (which data to which T1) crashed after kernel upgrade. 12:00 - 16:00 yesterday no d/s subscription. Restarted application - ok & understood. Good news from Lyon and CNAF - working much better. CNAF back in production in DDM.
- CMS reports - Not much at T1 level - following up with RAL ticket on 10GB q limit on transfers. Explains why some files not routed to RAL. T2: closed 2 pending tickets - US T2 in Wisconsin (files lost - understood to be related to problem on site: now sorted out); San Diego - understood and fixed - GSI errors on single WN with problem directory (empty). Issue that is getting worse: MC activities via gLite WMS. At least 2 T2s flooded by jobs possibly being resubmitted by WMS itself. Still globus error 10 issue open. MC production operations trying to implement retry count and shallow retry count to 3 - using this for MC production. Now see massive # of jobs causing problems to these two T2s- related to by-pass? Investigating. T2 at IN2P3 closed qs until sorted out.
- ALICE - In last days have sent 4 new requests to mx.support to register 4 new DNs in myproxy - request pending. dns please register asap. Budapest, Hiroshima, RAL and FZK. Sites cannot enter production until this is done. Issues last night: Legnaro T2 - error in publication of ALICE q in IS. Site now ok. Small issue with latest VO box & SL5: generic service associated to user proxy registration giving good results in SAM but bad in MonaLisa. Issue coming from environment in MonaLisa which conflicts with that of VO box.
- LHCb reports - the two MC production announced yesterday have been submitted. We expect a week end fairly busy; Stripping jobs at SARA (still failing with the application error code 134) ; DC06 DST redistribution across T1. It is about 75TB of data distributed in equal weight across all T1 out of CERN. No major problems observed in the first round of transfers (finished all at the same time) apart some issue at CERN (logged in the T0 section of this report). T0 issues: T0 sites issues: Some AFS volumes were not accessible this morning and since yesterday 17:30 affecting LHCb users of SL5 cluster. The problem has disappeared at around 10. Discussions about the criticality of AFS and its coverage; Reported some contention activity on CASTOR, activity coming from the LHCb PIT machines (outside CERN LAN); Problems transferring to SARA at some point on time last night due (most likely) to some disk server at the source. This problem disappeared this morning (no GGUS open). T1 issues: GridKA: issue with shared area yesterday.



Sites / Services round table:

- KIT: ntr
- PIC: ntr
- RAL: ntr - except declaring at risk Tue/Wed due to preventative maintenance on UPS.
- IN2P3: ntr
- ASGC: ntr - working on SIR.
- BNL: 1) announced DB migration to new h/w postponed - backup took longer than expected: file sizes created too large; creating another backup would have gone beyond announced time window. Proposal is to resume migration Monday. q to ATLAS regarding coordinators message of yesterday. Preference is to postpone until Xmas break. 2) As announced last week upgrade 60 disk servers. Running new Solaris version. Announced as transparent and worked perfectly. No transfer failures or jobs dying. Now completed. Jean-Philippe - please can you send me a mail regarding Solaris version?
- NL-T1: New dCache version with 2 fixes: 1 fixes incorrect filesizes and other fix is for tape protection. Tested and it works! Install on production instance on Wednesday.
  
- Network: fibre cut in France, link to PIC down (both main and backup) traffic via GPN. Down since 10:30. No ETR.
  
- Dashboards: migration for ATLAS DDM Monday; CMS dashboard Tuesday from 09:00. Estimated time ~3 hours.

AOB:

---

This topic: LCG > WLCGDailyMeetingsWeek091109

Topic revision: r3 - 2009-11-13 - unknown



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