

# Table of Contents

<b>Week of 090601.....</b>	<b>1</b>
WLCG Service Incidents, Interventions and Availability.....	1
GGUS section.....	1
Daily WLCG Operations Call details.....	1
Monday:.....	1
Tuesday:.....	1
Wednesday.....	3
Thursday.....	5
Friday.....	6

# Week of 090601

## 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</a>	<a href="#">WLCG Blogs</a>
------------------------	-----------------------	-------------------------------	---------------------	----------------------------

### Monday:

*No meeting today - holiday at CERN & elsewhere*

- ATLAS (Simone) - In the morning there has been a problem with the ATLR oracle database, hosting several applications like DDM, Panda, Dashboards etc... All the affected services could not contact the oracle backend for several hours. As rough estimate, the problem last from 6:30am to 11:15am CEST. Services are fully functional since then, thanks to the intervention of the experts, with the exception of a 20 mins glitch of the ATLAS central catalogs from 13:40 to 14:00 which is possibly a side effect of the previous problem. STEP09 ramp-up activities are back to normal. The transfers now run at 50% nominal rate as foreseen. Tomorrow the incident will be discussed with the experts and more infos will be provided.

### Tuesday:

Attendance: local(Harry, Daniele, Dirk, Nick, Andrea, Patricia, MariaG, MariaDZ, Roberto, Simone, Gang,

Jason, Steve, Sophie, Miguel, Markus, Olof, Julia, Jean-Phillipe, Graeme, Jamie);remote(Brian, Michael, Jeremy, Gonzalo, Fabio, Angela, Kors, John, Jeff).

Experiments round table:

- ATLAS - Graeme: Started STEP'09 this morning. 1) There is a logbook page of the main activities and observations at <https://prod-grid-logger.cern.ch/elog/ATLAS+Computer+Operations+Logbook/>. 2) The load generators for data movement have now been ramped up to 100%. 3) User analysis has started with 100000 jobs being sent via WMS servers. 4) Pseudo-reprocessing jobs reading tapes are now running on 4 clouds and others will be defined this evening. 5) Simone reported on the DB access problems of Monday (see ATLAS report above and CERN report below). Started around 08.30, fixed at 10 but knock-on effects for some hours. Mail was sent to [phydb.support@cern.ch](mailto:phydb.support@cern.ch) but was this correct? MariaG reported that this was ok but the DB team had already been alarmed by their own monitoring (they have an informal best-efforts phone rotation). 6) There are various transfer problems - failures IN2P3 to ASGC but not vice-versa, TRIUMF to CNAF is slow, TRIUMF to 2 Canadian T2 are slow, failures to Coimbra and Lisbon and SARA to T2 are slow. 7) ATLAS had not been computing checksums but some sites have now enabled this so it has been turned on. 7) BNL DQ2 VO-box has been taking 50% cpu. Found to be a level of GFAL brought in with a DQ2 upgrade but that does not work properly. GFAL component has now been downgraded.
- CMS reports - 1) follow-up plans following broken CERN tape of 28 May are satisfactory to CMS. 2) STEP'09 starting today focussing on site readiness at T1. Plans are to trigger major activities at 16.00 CEST daily as a good time for sites world wide so today will start prestaging then tomorrow at 16.00 the reprocessing of what has been prestaged in the previous 24 hours. 3) A CRUZET (cosmics at zero tesla) run this week at CERN with data export to T1 so no major Tier 0 activities so as not to perturb this. 4) Analysis tests at T2 are being ramped up but should not interfere with other VOs.
- ALICE - Began running in STEP'09 mode yesterday with over 15000 concurrent jobs. Will begin FTS transfers next week and decide how to use FZK meantime. Expecting to upgrade Alien soon. WMS issues at CERN due to bad CE shapes at KFKI and MEPHI. Details of the STEP'09 plans for ALICE: [GDB presentation](#)
- LHCb reports - STEP'09 will start next week for them with preparations this week to clear local tape disk caches. Had CRL problems on IN2P3 worker nodes (see below) and an NFS lock on NL-T1 worker nodes (see below). In future LHCb SW will be copied locally to NL-T1 worker nodes.

Sites / Services round table:

ASGC (by email from Jason on 30 May):

```
- due to bandwidth limitation from prestage disk pool, we're adding
another disk servers to share the gridftp load from the only one add two
weeks ago. performance expect to drive up to 210+MB/s requesting from
pre-stage performance validation by atlas.
- scratchdisk have been enable on T1 MSS. space token confirm with Atlas
central ops.
- 60TB add to ftt dpm now, space token updated are given below:
74d80644-e027-4319-8668-91ee14ca85bc ATLASDATADISK dpm_pool
  atlas
  58.59T 58.33T Inf REPLICa ONLINE
b9d9a84d-a144-4b61-9b73-405edf14471e ATLASMCDISK dpm_pool
  atlas/Role=production
  14.65T 14.23T Inf REPLICa ONLINE
- 5 new LTO4 drives onlined this Mon., full functional in Tue., this
week. tape capacity online for the time being (waiting confirmation from
the recovery of data center, this we'll enable another tape system at
ASGC data center providing 1.4PB with old base frame) is around 0.8PB.
```

Tuesday:

performance observed from metric is around 66.02 MB/s inc. drive overhead which is dominated by CMS mainly. file per mount is around 240 according to last 4hr stats, which is good as well. avg file size around 1.7G in atlas and around 3.4G in cms enhance also the performance of restaging processes.

- installation of new tape library with two S54 high density frame set and existing 3 frame sets implemented earlier of 2009. the inventory able to provide 1.3PB with all LTO4 cartridges.
- T1 MSS metrics uploaded and confirm with Alberto (metrics generated by tapelog reporter and also xml generation toolkit provide by Tim)

CERN: 1) Report on the ATR (DB) problem of yesterday (report): The switch connecting several production databases to the GPN network failed Monday morning around 8:30 causing unavailability of the systems ATR (Atlas off-line db), COMPR (Compass), ATLDSC (Atlas downstream capture) and LHCB DSC (LHCB downstream capture). The problem was fixed and databases were available again around 10:00. Some connectivity instabilities were also possible in the case of the ATONR database (Atlas online). Issues that might have caused such instabilities were addressed around 12:30. A post-mortem from the DB services point of view will be prepared. 2) The migration of the integration databases scheduled for today has been postponed 2 weeks after an issue was found. 3) there will be (transparent) patches to castorcms and castorlhcb tomorrow and an urgent upgrade of xrootd is being planned. 3) Now freezing the linux kernel upgrades. Was to be deployed 15 June but now will be delayed till 22 June. 4) The ATLAS SRM problem of last Thursday has not re-occurred and Dirk reported that an individual user was probably the cause.

IN2P3: HPSS migration has started with metadata being moved today. They modified their mechanism to download CRLs to worker nodes last week and this did not properly finish hence the LHCb problems.

NL-T1: 1) Had an NFS problem affecting LHCb software access on worker nodes after a user somehow caused NFS locks (which should not happen). They had to reboot the lock server plus the concerned worker nodes. 2) They were failing ATLAS SAM tests over the weekend due to the dcache disk pool 'cost' function not load balancing correctly across full disk pools.

FZK: Still suffering from tape problems - looking for bottlenecks.

AOB:

- STEP'09 post-mortem workshop (<http://indico.cern.ch/conferenceDisplay.py?confId=56580>) registration opens today!

## Wednesday

Attendance: local(Harry, Daniele, Nick, Andrea, Patricia, MariaG, MariaDZ, Simone, Gang, Jason, Sophie, Olof, Julia, Jean-Phillipe, Graeme, Jamie, Oliver, Ignacio, Diano, Antonio);remote(Reda, Brian, Michael, Jeremy, Gonzalo, Kors, Joel, Andrea, NDGF).

Experiments round table:

- ATLAS - 1) Simulation at sites has now started. 2) Data distribution is going pretty well. Only 5 raw streams are being produced so not all Tier 1 get data each day - TRIUMF and NDGF were the last recipients. 3) DDM server at CERN ran out of DB connections but has since recovered. 4) CE problem at RAL where they stopped accepting pilots perhaps due to a file limit on a globus directory. Brian reported RAL have cleared one of their two ATLAS CE and are using the second for debugging this problem. 5) At ASGC they have a serious interference between many stalled analysis jobs and the production jobs - being looked into. 6) Analysis via Panda and WMS is going well though have to check why some sites are empty of pilot jobs. Larger sites can have over 1000 analysis jobs at a time. Analysis at CERN will be started tomorrow. 7) Reprocessing has started at RAL/TRIUMF/FZK and CNAF but these tasks have now been aborted to get more work into the system. 8) Some backlog

writing to tape in NDGF where they got two datasets. Writing to tape at 60 MB/sec but will probably be excluded from the next distribution. 9) They would like to be sure the data committed to tape at CERN is actually going to tape (the tapes are being recycled).

- CMS reports - 1) Remedy ticket sent to CERN after disk only CASTOR files became inaccessible. Was due to a diskserver that crashed about 01.30. 2) CRUZET running at CERN so no STEP'09 activities at Tier 0. 3) Prestaging tests at Tier 1 started yesterday at 16.00 giving good results. The tests at FNAL have now completed successfully. 4) Workflows for processing at Tier 1 are prepared and will be triggered at 16.00.
- ALICE - 1) A new MC cycle is being prepared. 2) RAL-LCG2 CREAM-CE stopped working. Ticket was submitted and problem solved by restarting tomcat server. 3) A new Tier2 site has been added for ALICE, CESGA in Spain.
- LHCb reports - Had problems this morning ascribed to their castorlhcb upgrade but in fact this was done from 14.00 to 15.00 so they should check again.

Sites / Services round table:

CNAF: Had an LSF system error last Sunday.

BNL: 1) Had a site services problem affecting transfer performance due to a gfal library compatibility issue. After fixing this they observed up to 20000 completed srm transfers per hour, a very good rate. 2) They achieved sustained transfer rates of 700 MB/sec exporting to Tier 2 and 500 MB/sec incoming. 3) Data migration rates of 800 GB/hour from hpss to tape were observed at the same time as running merge jobs. 16 drives were used for staging and 4 were reserved for migration.

RAL: 1) Achieving 5 Gbit/sec to their Tier 2. 2) Have a couple of low level disk server failures. 3) Having transfer troubles from the Storm-SE at QMC to their CASTOR server. Will check with CNAF.

NDGF: Have had to increase FTS timeouts.

ASGC: - atlas user analysis job fill up the queue with 1.4k jobs and lower down the performance of pilot and production jobs. part of the user analysis jobs have been cleanup forcedly and apply later the reservation for different type of analysis jobs including also reasonable fair share for pilot jobs.

- rearrangement of the disk servers, pool to fit STEP requirement. changes mainly refer to Atlas specific service class and for disk cache only. this including also the new space token added and two new service class created two weeks ago, and one of it specific for pre-staging validation.

- new tape servers installed, 5 lto5 and one showing h/w error. the other one will be deployed soon next week. total online tape drives are 7 (with one offline atm, due to h/w issue):

\* tape drive hardware issue: tape can be unload normally. escalation to local IBM already

- dedication of tape drives for data migration. this is setup earlier to improve the data migration of CMS data with sufficient streams and tape drives dedication. (4 lto4 dedicated for data write)

- preparing the tape std op manual with assistant from Vlado. we will try tweaking the script provided as binding CERN local setup base on cdb. we hope to provide operation manual for on site duty staff asap.

- srm upgrade to 2.7-18 this week, from 2.7-17, concerning the big Id issue. date need to confirm later this week including also the status at the other T1s.

CERN: Database - RAL have taken over, with our thanks, DB synchronisation to ASGC.

CERN: PPS - Nothing happening in the immediate future.

AOB:

## Thursday

Attendance: local(Maarten, Gavin, Andreas U, Ignacio, Harry, Sophie, Graeme, MariaDZ, MariaG, Daniele, Gang, Edoardo M, Steve, Konstantin, Nick, Simone, Jean0-Philippe, Jamie, Diana);remote(Vera Hansper+Thomas Bellman -NDGF, P.Veronesi -CNAF, Angela, Reda, Jeremy, Dave, Gareth, Michael, Gonzalo, IN2P3).

Experiments round table:

- ATLAS - 1) Data distribution to NDGF is slow. Gavin reported they have doubled FTS concurrent transfer slots from CERN to NDGF from 20 to 40. NDGF reported they have also changed some parameters so that long running transfers do not break. 2) T1-T1 transfers were scaled down this morning which had the effect of dropping the success rate probably because it showed up the tail of problematic transfers. 3) For the T1-T2 transfers a snapshot taken at midday showed that 35 out of 62 T2 had received at least 90% of their data. Those T2 not working well had storage problems or, the more interesting case, had a large analysis load that was slowing down their WAN transfers and building up a backlog. 4) Reprocessing has started. Current inspection shows- ASGC 1 hour jobs still there after 6 hours - PIC ok with 557 jobs - SARA ok with more than 2000 jobs - NDGF only 177 jobs and don't seem to be getting prestage requests - CNAF ok with 900 jobs - FZK only 16 jobs but a very unhealthy tape system - RAL very good at 3000 jobs - TRIUMF none but doing a heavy high priority validation run - BNL also running validation - SLAC ok but slower since input and output data are via BNL. 5) Note that reprocessing logs as well as AOD are being written to tape but the success metric is number of files not number of GB. Individual site problems are now being chased up.
- CMS reports - 1) Prestaging continues at the T1 (except FZK and IN2P3) with good results. 2) Reprocessing was started yesterday at 16.00 on all 7 Tier 1. The executable was misconfigured at RAL and PIC and RAL jobs had to be killed and resubmitted. 3) Special Phedex setup at PIC to avoid WAN transfers triggering scattered tape mounts. To be tested over Spanish T1 to T2 over the next few days.
- ALICE (Patricia absent due to the ALICE weekly TF Meeting) - Alice MC new cycle creation: done. The production came back last night, currently running around 8000 jobs. MonaLisa monitoring having problems this morning, service temporarily down but experts taking care. CESGA site (Santiago de Compostela, Spain) entered production yesterday evening. Currently jobs aborting at CC-IN2P3. The WMS used for this site is that one placed at GRIF which is not able to find the resources at CC-IN2P3. The problem has been reported to the site admins and service responsible and we are looking into it (GGUS ticket #49250).
- LHCb reports - 1) Mapping problem at PIC - quickly solved. 2) Discovered a CERN CASTOR problem of inaccessible files that was thought to be due to the last minor upgrade but in fact had been there from the recent major upgrade and involves data loss after garbage collection was accidentally enabled. Attempts are being made to recover as much as possible. See post mortem at <https://twiki.cern.ch/twiki/bin/view/FIOgroup/PostMortem20090603>

Sites / Services round table:

- CERN:

- ◆ SELinux execheap check has been disabled in the SLC5 WNs (requested by ATLAS). (At CERN a ncm component is used to configure this, the same results can be achieved by executing: `setsebool -P allow_execheap=true`.)
- IN2P3: The ALICE job failures at IN2P3 are thought to be due to a problem on the GRIF bdii (the WMS used is at GRIF). The HPSS migration should be ending this afternoon so tape services will be resuming.
- TRIUMF: Seeing FTS timeouts on transfers from RAL and ASGC especially for the large (3-5 GB) merged AOD files. Have hence increased their timeouts from 30 to 90 minutes. Simone added they also see timeouts from TRIUMF to CNAF (which goes over the GPN in fact).
- RAL: The CMS executable problem was in fact picked up quickly by local CMS support. They are seeing timeouts on their RAL to BNL FTS channel.
- BNL: As regards ATLAS reprocessing BNL had 5000 jobs queued which triggered prestaging of which 4000 are already done so the reprocessing jobs will now ramp up as job slots become free.

AOB: (MariaDZ) For ALICE (also sent in email today): Please create in vomrs groups to contain your teamers and alarmers as per [https://twiki.cern.ch/twiki/bin/view/EGEE/SA1\\_USAG#How\\_to\\_register\\_your\\_TEAM\\_and\\_AL](https://twiki.cern.ch/twiki/bin/view/EGEE/SA1_USAG#How_to_register_your_TEAM_and_AL) For CMS (also in <https://savannah.cern.ch/support/index.php?104835#comment68>), please change your groups from TEAM and ALARM into "team" and "alarm". Although you were the only experiment that did what the instructions originally recommended, LHCb created Roles in lower case and ATLAS created Groups in lower case. As VOMS is case sensitive please be so kind to also go into lower case to simplify the GGUS parsing rules.

GOCDDB: is now fully moved to the backup service.

Middleware: A.Unterkircher reported a recently found lcgutils problem whereby gfal\_open gives a segmentation fault under certain conditions. Will discuss what to do offline with the experiments.

## Friday

Attendance: local(Harry, Gang, Patricia, Sophie, Lawrence, Gavin, Steve, Roberto, Simone, Graeme, Julai, Daniele, Eva, Andrea);remote(Reda, Vera, Michael, Paolo, Jeremy, Xavier, IN2P3, Dave, Gareth, Brian, Jeff).

Experiments round table:

- ATLAS - 1) This morning many dataset transfers failed to complete due to a set of scripts that got stuck overnight. This arises from the 'open datasets' way of working where ATLAS no longer wait for a complete dataset to be ready before starting to transfer it. The problem is version numbers are changing frequently and, for historical reasons, there is a call to the central catalogue containing all recent version numbers and this has exceeded the allowed string size. Immediate fix was to run fewer cron jobs changing the version number (currently 90% of transfers have recovered) then shortly after this meeting the definitive fix will be to only send the currently valid version number. 2) T0-T1 transfers are always running and new T1-T1 and T1-T2 transfer subscriptions have been made so there will be a burst of data movement. 3) Some of the residual failed T1-T1 transfers have been understood as due to FTS timeouts so these are being systematically increased. 4) SARA prestaging stopped at midnight but has since recovered. 5) NDGF are not getting prestage requests and this is thought to be a difference in the ARC architecture where there is an extra state in the workflow. They are working with the local experts. 6) The problem at ASGC of worker nodes failing to connect to the conditions database is now understood to be due to a port number change that happened some time ago. This may need recompilation of some ATLAS modules which will take time though there may

be a workaround using environment variables. 7) Reprocessing at TRIUMF and BNL is picking up as expected after their validation jobs. 8) Simulation is also running well except that NDGF have found a workflow issue that reprocessing jobs are blocking slots that should be available for simulation. ASGC is catching up on their simulation load since their reprocessing is not running. 9) User analysis is running well with 275000 jobs run since last Tuesday. Daniele queried the nature of these jobs - there are two types - a short 15 min job reading 2GB and a longer 45 min job reading 15 GB. Equal times go into each flavour. 10) Merging of 5 million events is about to start - typically 20 files of 100 MB are merged into a single 2 GB file. 11) ATLAS will not hesitate to send site alarm tickets during STEP09.

- CMS reports - 1) Last day of CRUZET run. Tier 0 tape writing tests will start tomorrow (Simone - ATLAS already writing to tape at raw data rates - physical commit to tape needs checking) 2) T1 prestaging tests going well and some sites would like to try other methods. CMS trying to discover what ATLAS activities are overlapping at their multi-VO sites. 3) Reprocessing is running well with tests without prestaging in the planning. All 7 T1 are getting reprocessing jobs and they have been submitted based on expected site resources so some sets are not completing and are being killed. CNAF are having problems using both CASTOR and STORM at the same time so a new approach will be tried. 4) IN2P3 prefer to start major activities next Monday to give time for consolidation of the recent changes. 5) ASGC reported they are not getting enough CMS jobs.
- ALICE - The IN2P3 ticket from yesterday (failing batch jobs) has been closed as solved but details are not yet known. RAL are asking ALICE to test batch jobs under SLC5 after which they will open some more job slots.
- LHCb reports - 1) GGUS tickets will be sent to sites next week to clear tape disk caches. 2) Several thousand files have been lost from the CERN lhcbdata pool when garbage collection was accidentally turned on. 3) the gsidcap plugin (as used at NL-T1 and IN2P3) cannot read non-root native files.

#### Sites / Services round table:

- RAL - After increasing the FTS timeout there have been no more transfer failures with BNL. At the RAL Tier1 we have created a blog of our experiences so far with STEP09. This has been linked to the WLCG aggregate blog of Steve Traylen which is to be found as 'WLCG Blogs' under 'STEP09 Additional Material' above. Reminded that at the end of this month they will be moving to their new building (photos on the Blog) and as a start the CA service will be moved next Monday-Wednesday.
- TRIUMF: Systematically increasing FTS timeout settings.
- BNL: Reprocessing is ramping up and at the same time there is significant analysis activity.
- CNAF: Problem of slow transfers TRIUMF-CNAF fixed - was due to a mispublishing of Storm disk pool information.
- NDGF: Have been tuning their LHCOPN connection. Simone reported transfers from T0 running at 200 MB/sec and their backlog has been cleared.
- ASGC: Read-only access to the ATLAS conditions database from the worker nodes has been failing inhibiting STEP09 activities. Problem now understood - see ATLAS report.
- IN2P3: End of HPSS migration at CC-IN2P3 - Service is restarted gradually. Staging is allowed but load is kept low until Monday (writing on disk only for the moment).
- CERN - Oracle has released a first version of the patch for the "BigId" issue (cursor confusion in some condition when parsing a statement which had already been parsed and for which the execution

plan has been aged out of the shared pool). It was immediately identified that the produced patch is having a conflict with another important patch that we deploy. Oracle is working on another “merge patch” which should be made available in the coming days. In the meantime a work-around has been deployed in most of the instances at CERN as part of CASTOR 2.1.8-8.

---

This topic: LCG > WLCGDailyMeetingsWeek090601

Topic revision: r23 - 2009-06-08 - HarryRenshall



Copyright &© 2008-2022 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](#)