

# Table of Contents

<b>Week of 110718.....</b>	<b>1</b>
Daily WLCG Operations Call details.....	1
WLCG Service Incidents, Interventions and Availability, Change / Risk Assessments.....	1
General Information.....	1
Monday:.....	1
Tuesday:.....	3
Wednesday.....	5
Thursday.....	6
Friday.....	8

# Week of 110718

## 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](#)

## WLCG Service Incidents, Interventions and Availability, Change / Risk Assessments

VO Summaries of Site Usability				SIRs, Open Issues & Broadcasts			Change assessments
<a href="#">ALICE</a>	<a href="#">ATLAS</a>	<a href="#">CMS</a>	<a href="#">LHCb</a>	WLCG Service Incident Reports	WLCG Service Open Issues	Broadcast archive	CASTOR Change Assessments

## General Information

General Information				GGUS Information	LHC Machine Information
CERN IT status board	M/W PPSCoordinationWorkLog	WLCG Baseline Versions	WLCG Blogs	GgusInformation	Sharepoint site - Cooldown Status - News

## Monday:

Attendance: local(Jamie, Simone, Jarka, Andrea, Ale, Stefan,Ullrich, Lukasz, Jacek, Jan, Lola );remote(Michael/BNL, Xavier/PIC, Fran/NDGF, Jon/FNAL, Tiju/RAL, Xavier/KIT, Onno/NL-T1, Giovanni/CNAF, Shu-Ting/ASGC, Rob/OSG).

Experiments round table:

- ATLAS reports - Jarka
  - ◆ Start of ES1 reprocessing campaign on Saturday morning.
  - ◆ CERN-PROD: Files deleted from CASTOR shortly after written to CASTOR, GGUS ALARM:72709. Picket operator introduced developers into issue, savannah <https://savannah.cern.ch/support/index.php?122264>.
  - ◆ CERN-PROD: cannot access some files at CERN GGUS:72697. Involved Giuseppe Lo Presti. From ATLAS side it helped to delete affected files and copy them again.
    - ◇ Jan: issue identified and bugfix release planned for Wed (fix targeted to loss of DB connections). Ale: Why were files not spread more? Deletion was done only locally as DB with global picture could not be contacted. Ale: will get back soon about intervention but do not expect a problem.
  - ◆ CERN-PROD: CERN-PROD\_DATADISK too many threads busy with Castor, TMPDATADISK locality unavailable GGUS:72714. Solved "itself" after ca 2 hrs.
  - ◆ SARA-MATRIX: On Sunday late morning SRM degraded. GGUS:72713 created, later in

the evening escalated to ALARM. Apparently, SARA has issues accepting ALARM tickets (list is too restrictive). This needs to be followed-up with GGUS.

- ◊ SARA: issues since yesterday early afternoon - restarted dCache this morning at 9:00. Service is ok since then, but this is the second time that SARA stability was affected. Now in contact with dcache developers to identify root cause of the problem.
  - Onno: Using a script monitoring for ALARM tickets - but contact address sara-matrix was not picked up. Also the list of authorised alarmers is not up to date and site would welcome an update.
  - Ale: Will be clarified issue of second contact address between site and ATLAS. ATLAS could also create a new alarm for SARA if really required.
  - Simone: list of alarmers is a list of DNs in ATLAS VOMS in alarmers group. SARA would be happy to look at VOMS list to update the alarmer list.
- ◊ FZK-LCG2: Some files not available GGUS:72702. • Xavier/KIT : unavailable files are same problem as last week - gpfs mount problem could not yet be solved.

• CMS reports -

- ◆ LHC / CMS detector
  - ◊ 6 fills since last Friday noon, CMS acquired ~50pb-1 luminosity.
- ◆ CERN / central services
  - ◊ Issue for CMS CRC to send GGUS Alarm tickets to INFN-T1 ("You are not allowed to trigger an SMS alarm for INFN Tier1"), see GGUS:72705 . GGUS admins acknowledged they know about this issue since the GGUS May release and they are in touch with the "Italian developers", with work in progress, see GGUS:72717
  - ◊ CMS Services Map went down on Sun Jul 17, around 00:00
    - Server restarted on Mon Jul 18 morning, see Savannah:122260 . Maybe the Dashboard team should consider deploying a back-up production server for CMS Services Map ?
      - Lukasz: waiting for new VM from cern which will solve this.
- ◆ T1 sites:
  - ◊ INFN-T1 : all transfers to the site and most transfers from the site started failing on Sat Jul 16, ~16:30 UTC
    - First problem was the failing GGUS Alarming (see above). CMS CRC-on-Duty contacted a local admin (via D.Bonacorsi, thanks !) who eventually reacted around 20:30 UTC
    - Storm End-Point was restarted
    - The local storage admin diagnosed that there actually was only 1 problematic FE - which could explain why some traffic was not affected.
- ◆ T2 sites:
  - ◊ NTR
- ◆ AOB:
  - ◊ Andrea Sciaba reporting for CRC (PK) today, thanks !

• ALICE reports - Lola

- ◆ T0 site - GGUS:72735. Mismatch between number of jobs reported by the information provider and MonaLisa.
  - ◊ Ullrich: will need more info about which jobs remain valid. Lola: will update ticket.
- ◆ T1 sites - Nothing to report
- ◆ T2 sites - Usual operations

• LHCb reports - Stefan

- ◆ Express has picked up new data.
- ◆ T0 - LFC the lookup for some users (function getusrbyuid) is failing, where the wrong uid is queried. This prevents e.g. job output to be uploaded (INC:053039)

Sites / Services round table:

- Michael/BNL - ntr
- Xavier/PIC - ntr
- Fran/NDGF - ntr
- Jon/FNAL - ntr
- Tiju/RAL - ntr
- Xavier/KIT - weeken two stage pools down for lhcb - restarted this morning
- Onno/NL-T1 - ntr
- Giovanni/CNAF - problem of sat transfers were due to storm front-end problem.
- Shu-Ting/ASGC - ntr
- Rob/OSG - ntr
- Jan/CERN: ticket from ATLAS for transfers to NDGF - problem seems on NDGF side.
- Lukasz/CERN: ATLAS dashboard was upgraded - no downtime.
- Ullrich/CERN: incident with batch farm solved after 1h.
- Reminder - **CERN - Tier2 FTS migration to SL5** On Tuesday 19th at 08:00 UTC (10:00 @ CERN) the tier2 **FTS webservice** will be migrated from SL4 to SL5 and gLite 3.1.21 to gLite 3.2.1. The operation should be utterly transparent and take around 20 minutes. Any rollback required is similarly trivial.

AOB:

## Tuesday:

Attendance: local(Peter, Lukasz, Lola, Jamie, Maria, Alessandro, JhenWei, Stefan, Jacek, Jan, Uli);remote(Ronald, Jeremy, Maria Francesca, Marc, Tiju, Michael, Jon, Dimitri, Giovanni, Gonzalo, Giovanni, Rob).

Experiments round table:

- ATLAS reports -
  - ◆ CERN-PROD 3 weeks ago diskserver with HW problems, GGUS:71866 [↗](#) , problems reappeared. [ Jan - we will declare files on this m/c to be lost. Under repair for long time so much easier to have a clean situation. Apologies. Between us and vendor. Also have a bunch of GGUS tickets open on SRM transfers which are still being investigated. A lot of transfers all come in in a bunch and all timeout. 3 tickets for ~same symptom. Will update tickets. ]
  - ◆ INFN-T1 "File Exists" errors. GGUS:72749 [↗](#) . This is not a site issue, sorry for it. It is related to a HW problem of one of the ATLAS DDM SiteServices machine, that needed to be switched to a spare one: the files already in FTS jobs launched from that box were not registered in LFC at the end of the successful transfer, thus they have been retried later from the new SS machine.
  - ◆ RAL - some files missing: D/S out of production but able to recover it
  - ◆ IN2P3 - some ongoing discussions within ATLAS/IN2P3 if issues at site. Some power problems? Maybe LFC/FTS affected? Not observing any issue right now. Any problems to be expected? [ Marc - we had incident this night. Not power but DB problem on disk array for Oracle DB. Quite serious problem - many services impacted including tables for experiments such as ATLAS, LHCb. More info in site report later. ]
- CMS reports -
- LHC / CMS detector
  - ◆ not much collisions data taking since Monday
  - ◆ expecting beam on 1380 bunches this evening
- CERN / central services

Tuesday:

- ◆ Tier-0 ready (and hungry !) for new data
  - ◆ Working in collaboration with CERN/IT to clean-up / re-shuffle / optimize some dedicated CMS LSF queues (cmst0, cmscaf, cmsinter, ...)
  - T1 sites:
    - ◆ NTR
  - T2 sites:
    - ◆ NTR
  - AOB:
    - ◆ NTR
- 

- ALICE reports [↗](#) - General Information: Yesterday there was a problem with the Apache server in one of central services machines. Due to that the Job Broker was stuck and unable to be restated properly for few hours which provoked that the number of jobs running decreased.
  - ◆ T0 site - Change hardware of aliendb4, one of the authentication servers for Alice. (we have another 5!)
  - ◆ T1 sites - Nothing to report
  - ◆ T2 sites - Usual operations
- LHCb reports -
- Express and Full reconstruction working on latest data.

#### New GGUS (or RT) tickets:

- T0: 0
- T1: 0
- T2: 0

#### Issues at the sites and services

- T0
    - ◆ LFC the lookup for some users (function getusrbyuid) is failing, where the wrong uid is queried (INC:053039). Fixed, was due to internal LHCb LFC problem (uid reported of users that have left the VO).
  - T1
    - ◆ IN2P3: 3D streaming started failing this morning (4.30 am) and was due to a disk crash at IN2P3 (GGUS:72756) . Oracle DB at IN2P3 fixed this morning, but up to now, according to SLS, only 1 out 3 streams is working (INC:053801 [↗](#))
- 

#### Sites / Services round table:

- NL-T1 - ntr
- NDGF - ntr
- IN2P3 - 2 incidents: first one was the DB h/w problem this night. Disk array on Oracle. Happened around 03:00 had impact on several tables for LHCb, ATLAS, AMI, some grid services (FTS, LFC). Problem not totally solved - experts still working on it. But LHCb should be ok now, ATLAS DB are still degraded. AMI restarted but stream broken. Ale - we have backup server R/O at CERN. Marc - when all solved we will do an SIR. Gridengine failure last night too. GE daemon crashed also around 03:00 and restarted manually. All running and queued jobs cancelled due to some AFS token issue which happens when we restart GE. Two incidents not related!
- RAL - ntr
- BNL - ntr
- FNAL - ntr
- KIT - ntr

- ASGC - we have failed to install CMS s/w 4\_2\_7. CMS would like this asap for CMS SAM tests. We have submit install request to CMS for this. Scheduled downtime this Friday 04:00 - 09:00 for network maintenance. Will set scheduled downtime on GOCDB later.
- CNAF - ntr
- PIC - ntr
- GridPP - ntr
- OSG - ntr
  
- CERN Storage - reminder: transparent CASTOR ATLAS intervention tomorrow but haven't yet received the release and we have to test it first! Might slip.
  
- CERN Grid - FTS intervention this morning announced yesterday went ok.

AOB:

## Wednesday

Attendance: local(JhenWei, Alessandro, Jan, Maria, Jamie, Dirk, Lukasz, Jacek, Edoardo, Nicolo, Lola);remote(Marc, Onno, Tiju, Michael, Jon, Gonzalo, Maria Francesca, Dimitri, Giovanni, Vladimir, Rob).

Experiments round table:

- ATLAS reports -
  - ◆ IN2P3-CC Oracle DB instabilities: ATLAS noticed some failures during the night affecting analysis and production jobs. LFC and COOL were affected. IN2P3-CC posted goddb warning [☞](#), thus we did not take any immediate action. [ Marc - still experiencing problem with Oracle. Some downtimes. FTS, LFS and CIC operations portal are affected. Broadcast of downtime might not work until tables are recovered. Ale - ATLAS is observing that LFC is working but perhaps degraded, also FTS. The service that is not working at all is FroNTier. Marc - tables are there but some data corrupted. Service working but unstable. Ale - LFC more critical than frontier for us]
  - ◆ Users jobs are creating some troubles filling up \$HOME dir in WNs where they run. RAL at least is observing the issue. User has been contacted.
  
- CMS reports - LHC / CMS detector: Single collision-run of 1.1h on 1380 bunches last night, CMS collected 5pb-1
  
- ALICE reports [☞](#) - ntr
  
- LHCb reports - Express and Full reconstruction working on latest data.
  - ◆ T1 - IN2P3: We have still 3D streaming problem at IN2P3 (GGUS:72756 [☞](#))

Sites / Services round table:

- ASGC - CMS s/w issue has been fixed - necessary version installed
- IN2P3 - still experiencing troubles with Oracle. DB experts trying to recover tables but having some difficulties - service not stable yet.
- NL-T1 - on SARA SRM one pool node crashed at beginning of afternoon. Restarted and now running fine. During period of 20' some ATLAS and LHCb files may not have been available. Followup on SRM instability issue: had contact with dCache developers who suggested better method of collecting debug info. If it happens again we will collect enough info for them to debug. Reminder: tape b/e unavailable Aug 21 - 26 and downtime is in GOCDB.
- RAL - ntr
- BNL - ntr

- FNAL - a cooling outage yesterday - had to turn off 30% of Tier1 production nodes. Production service at FNAL notified. Edoardo - 2 links from CERN to FNAL are down since two hours. Jon - received info this morning that link between FNAL and ANL was cut. Bandwidth reduced to 3.5Gbps on backup. Scheduled maintenance but not pre-notified.
- PIC - next Tuesday 26 July during 4 hours in morning we will make intervention in module hosting 1/2 of WNs. Capacity will be reduced to about 50% during 4 hours
- NDGF - this morning a security update on one of nodes in Denmark. Scheduled on GOCDB. For 1 hour scheduled but took a little bit longer - 1h 20'. Some of ATLAS data might not have been available for this time
- KIT - ntr
- CNAF - some people are correcting StoRM bug related to lcg-gt command fails and creates unreadable local file. We did gfal downgrade yesterday to attempt to correct this but this generated a lot of failures of ATLAS jobs so this operation cancelled. Ale - do you need help to debug?
- OSG - maintenance upcoming on Tue 26. Both BDIIs will be down but at separate time. Should be no affect. Each BDII will be restarted.
  
- CERN network - started replacing OPN routers at CERN. We will have to move all links to T1s to new routers. We will start tomorrow with IN2P3 at 09:00 and waiting for confirmation from RAL (Monday 10:00) CNAF (11:00) and BNL (15:00). Monday interventions not yet confirmed by T1s. No impact apart from short glitch - traffic will go to backup path. More at [https://ggus.eu/pages/ticket\\_lhcopn\\_details.php?ticket=72754](https://ggus.eu/pages/ticket_lhcopn_details.php?ticket=72754)
  
- CERN storage - update this morning on CASTOR ATLAS which took longer and less transparent than expected but overall ok.

AOB:

- No T1SCM tomorrow.

## Thursday

Attendance: local(Jamie, Maria, Matteo, Lukasz, Alessandro, Jan, Peter, Jacek, Uli, Lola);remote(Vladimir, Giovanni, Paco Bernabe, Jon, John, Marc, Maria Francesca, ShuTing, Rob).

Experiments round table:

- ATLAS reports -
  - ◆ IN2P3-CC Oracle DB instabilities: the problem persisted, site yesterday declared outage for LFC/FTS till tomorrow morning thus ADC set the whole French cloud offline in analysis/production, and blacklisted the site in DDM. FronTier and all Oracle services are down. For FronTier jobs can contact another server. That's why we'd really like to have LFC/FTS back
  
- CMS reports -
- LHC / CMS detector
  - ◆ Very long collision-run on 1380 bunches last night, CMS collected 37pb-1
- CERN / central services
  - ◆ Short degradation of availability of CMS T0EXPORT (CASTOR) instance around 11:30AM today, see [http://sls.cern.ch/sls/history.php?id=CASTORCMS\\_T0EXPORT&more=availability&period=24h](http://sls.cern.ch/sls/history.php?id=CASTORCMS_T0EXPORT&more=availability&period=24h). Also T0TEMP and CMSRPODLOGS were degraded
    - ◇ CRC decided not to open a ticket to CERN/CASTOR since CMS Tier-0 was not dramatically affected and the instances came back 100% available after 1/2h

- ◇ However, CMS Computing Shift procedures need improved alarm criteria to trigger tickets. Working on shift instructions with CMS Tier-0 experts.
  - ◇ Latest news : CMSPRODLOGS is 94% degraded now (12:50 UTC) (15:00 - now almost 100% degraded...) Jan - agreement is that you should not blindly open tickets based on SLS but only when you see an impact on your jobs. Seems to be load related.
  - T1 sites:
    - ◆ Regarding the Oracle (FTS) issue at CCIN2P3, it first looked like things are recovering and CMS started to see some working transfers on affected links around 10AM UTC today, however as of 12:00 UTC the transfer quality seems more degraded again, see for example [https://cmsweb.cern.ch/phedex/debug/Activity::QualityPlots?graph=quality\\_all&entity=dest&src\\_filter=](https://cmsweb.cern.ch/phedex/debug/Activity::QualityPlots?graph=quality_all&entity=dest&src_filter=))
    - ◇ and in fact we see the downtime has been extended by 24h : [https://goc.egi.eu/portal/index.php?Page\\_Type=View\\_Object&object\\_id=48237&grid\\_id=0](https://goc.egi.eu/portal/index.php?Page_Type=View_Object&object_id=48237&grid_id=0)
  - T2 sites:
    - ◆ The **number of slots per week used by CMS Analysis jobs at Tier-2s** has constantly increased since the beginning of the year, from ~ **5,000 per week in Jan 2011 to 17,000 now**, while the number of **active CMS users per week** went from ~ **150 to 450** at the same time
  - AOB:
    - ◆ NTR
- 

- ALICE reports [☞](#) - NTR
- LHCb reports - Express and Full reconstruction working on latest data.
  - ◆ T1 - IN2P3: Oracle problem (GGUS:72756 [☞](#)) not solved yet

Sites / Services round table:

- CNAF - ntr: Ale - yesterday you made a point about an lcg-gt issue. Discussed with cloud support of ATLAS and understood all actions already taken - tickets open against m/w. Problem on StoRM site. Due to middleware bug. Properly followed but not affecting experiment
- NL-T1 - one dCache pool node had kernel panic last night. Did someone from ATLAS/LHCb/ALICE have problems? Service restored this morning.
- FNAL - ntr
- RAL - ntr
- IN2P3 - Oracle problem still going on. Downtime extended until tomorrow 10:00 UTC. LFC/FTS/VOMS. We are observing data corruption since failure of a disk on Tuesday. As I am talking CESNET DB restored with backup from last Monday. Grid services should be available soon, maybe this afternoon. if so will end downtime. ATLSA DB should be up by now - restored with Monday night's backup. Concerning AMI DB - still down. Backup max 8 h before problem started. Peter- for FTS for CMS any expectation? Marc - DB is up, have to check services and make sure no more corruption but could be back this afternoon. Ale- we are fetching downtime from GOCDB. If you shorten downtime data will start flowing to IN2P3.
- NDGF - one of our sites in Estonia has problem for CMS: one of values was wrong - forced 14k jobs to stay in queue which created a lot of problems. Proper value now set - bunch of jobs still queuing and will be some time before cleared. Just to explain why some jobs will not get through quickly
- ASGC
- KIT: we have no incidents to report from KIT. But we have an announcement for a downtime: On 28.07.2011, starting at 10 and lasting until 11 o'clock (local time), we want to migrate the ATLAS 3D database to new hardware.
- OSG



- CERN EOS - minor update to EOS ATLAS this morning. Not on status board as status board down at the time!
- CERN DB - helping IN2P3 DBAs to re-enable streaming of ATLAS data. Not excluded that full copy of 600GB of data will be need
- CERN dashboard - got new VMs from CERN and will be migrating services. GridMap and CMS critical services tomorrow

AOB:

## Friday

Attendance: local(Peter, Alexei, Jamie, Lukasz, Jan, Alessandro, Jacek, Uli);remote(Michael, Giovanni, Jon, John, Vladimir, Onno, ShuTing, Marc, Rob, Xavier, Onno).

Experiments round table:

- ATLAS reports -
    - ◆ IN2P3-CC Oracle DB: Information from Pierre Girard to ATLAS Computing this morning: LFC data recorded after Monday at 12:00 are not available. Actions taken are summarized in the ATLAS elog 27536
    - ◆ All UK cloud blacklisted in analysis by HC due to the fact that one dir in LFC (the one used by Johannes, the HC jobs owner) reached the 1M entry. Cloud support and DDM team are cleaning it, automatic cleaning of the empty dir should be applied within ATLAS tool
    - ◆ Express Stream reprocessing has started.
    - ◆ ATLAS is in data taking: last fill (#1962) lumi peak 1.28e33, integrated luminosity 46.3pb-1, lasted for 15h26' !
  - CMS reports -
  - LHC / CMS detector
    - ◆ NTR
  - CERN / central services
    - ◆ NTR
  - T1 sites:
    - ◆ Transfers from/to IN2P3 ok again
  - T2 sites:
    - ◆ NTR
  - AOB:
    - ◆ NTR
- 

- ALICE reports [↗](#) - NTR
- LHCb reports - Express and Full reconstruction working on latest data.
  - ◆ T1 - IN2P3: Oracle problem (GGUS:72756 [↗](#)) not solved yet

Sites / Services round table:

- BNL - ntr
- CNAF - ntr
- FNAL -
- RAL - ntr
- ASGC - ntr
- NL-T1 - this morning we had a problem with Oracle DB that became unresponsive affecting LFCs and FTS. OK again after complete restart. No clues in logs.

- IN2P3 - update on Oracle problem. Grid services now ok. Ended downtime yesterday. Some details on DB restore. Confirm that cesnet base for LFC Monday evening. ATLAS DBs from Monday 07:00 and AMI from Monday 20:55. ATLAS stream now ok - opened yesterday and synchronised this morning. AMI stream still down. Vladimir - LHCb services? A: re-imported DB directly from CERN so all data should be synchronised. Conditions and LFC? Alexei - received email from IN2P3 director Dominique Boutigny. Appreciate very much this email.  
Details from Marc:

CHESMET (LFC, FTS): restored from Mon Jul 18 12:00 CEST (noon)  
Correction is needed there. The backup version is not from Monday evening !

AMI: restored from Mon Jul 18 20:55 CEST  
DB\_ATL: restored from Mon Jul 18 07:00 CEST

- KIT - ntr
- OSG - ntr
  
- CERN - ntr

AOB:

-- JamieShiers - 27-Jun-2011

---

This topic: LCG > WLCGDailyMeetingsWeek110718

Topic revision: r21 - 2011-07-22 - JamieShiers



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