

Table of Contents

Week of 100906.....	1
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.....	7
Friday.....	7

Week of 100906

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
ALICE ↗	ATLAS ↗	CMS ↗	LHCb ↗	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(Steve, Doug, Zbyszek, Harry, Ale, Eddie, Ron, Simone, Andrea, Ignacio, Dirk, Zsolt, MariaDZ, Roberto);remote(Rolf, Alessandro, Foued Jrad, Rolf, Gang, Joel, Gonzalo, Tiju, IanF).

Experiments round table:

- ATLAS reports -
 - ◆ Ongoing issues
 - ◇ SARA oracle outage GGUS:61265 [↗](#) (plan to get NL online decided)
 - ◇ INFN-BNL network problem GGUS:61440 [↗](#)
 - ◇ RAL-NDGF functional test GGUS:61306 [↗](#) (FT transfers continue to fail)
 - ◆ Sept 4,5 and 6 (WeekEnd + Mon)
 - ◇ LHC: no physics
 - ◇ Taiwan-LCG2: downtime ended, but still issues are observed. They are for now kept out of SC RAW data distribution, and their queues are set offline.
 - ◇ NL cloud: file deletion ongoing with good pace. From SARA-MATRIX: load on LFC is low.
 - we are in the process of re-including all the Tier2s datadisk endpoint in FunctionalTest
 - NIKHEF data10_7TeV ESD (last 2 weeks) data transfers resumed: 1.3GB/s 99% eff!!

· SARA and NIKHEF are still out from SantaClaus: new data for now will not be subscribed. We will see in the next days how to proceed.

- CMS reports -
 - ◆ Experiment activity
 - ◇ Nothing to report
 - ◆ Central infrastructure
 - ◆ Tier1 issues
 - ◇ Something of a miscommunication between MC Ops and KIT. We didn't get a tape family for a new production run. Because it's simulation, we will clean out the files and retransfer with a new family. [or PIC? See below]
 - ◆ Tier2 Issues
 - ◇ Two sites today had visibility problems in BDII. T2_RU_INR and T2_UK_London_Brunel. [Maarten - normal route is GGUS ticket against site]
 - ◆ MC production
 - ◇ Very large (1000M) MC production ongoing T1s and T2s; will start the reprocessing with CMS_SW 38XX soon.
- ALICE reports [↗](#) - GENERAL INFORMATION: Pass1 reconstruction and one Pb+Pb MC cycle, principal activities during the weekend.
 - ◆ T0 site
 - ◇ There are about 2500 jobs started and running through the CREAM CE at CERN, but apparently they are not doing much. Experts have been asked to kill all jobs belonging to alisgm3.
 - ◇ from the point of view of the submission all CREAM-CE systems have been tested this morning with no incidents to report
 - ◇ the AFS issue reported last week and concerning the lack of synchronization between the ro/rw areas have been solved and confirmed by ALICE experts
 - ◆ T1 sites
 - ◇ Good behavior of the T0-T1 transfers during the last 24h when transfers to IN2P3, CNAF and FZK have been executed with no remarkable issues
 - ◇ Minor operations required this weekend at the local VOBOXES" FZK (expiration of the user proxy responsible of the production)
 - ◆ T2 sites
 - ◇ After several days of debugging, Subatech has confirmed today the good behavior of the CREAM-DB system at the site and the good synchronization between the information provided by the DB and the BDII.
- LHCb reports - Finishing reprocessing and merging of data at RAL and SARA.
 - ◆ T0 site issues:
 - ◇ CERN: none
 - ◆ T1 site issues:
 - ◇ RAL: another disk server (gdss379, lhcbuser space token) crashed (GGUS:61825 [↗](#)). No clear the reason, back again in production. requested to decrease the number of job slots having again a very poor success rate 25% (see picture) (GGUS:61798 [↗](#)).
 - ◇ GridKA: a lots of user jobs failing over the last week as the watch dog identifies them as stalled. Looking at the logs of these jobs it looks like these jobs are stalling while reading the input data via dcap servers. A ticket has been risen (GGUS:61841 [↗](#))

Sites / Services round table:

- INFN - ntr
- KIT - ntr but our FTS monitoring was updated last Friday to last version
- IN2P3 - ntr

Monday:

- ASGC - ntr [please ensure regularly participation and updates in case of problems]
- NL-T1 - 3D DBs and LHCb LFC being streamed from RAL. Started up today and will take approx 2 days. Zbyszek - copied via transportable tablespaces] Ale - FTS status? Ron - as far as we can see it is ok.
- PIC - things fine, most of weekend RED In GridView. Everything fine but regional Nagios system is currently being setup at Spanish NGI. Problem there being setup. Useful to be able to crosscheck with CERN instance of SWE region. Helped to understand that problem was outside PIC. (But for how long?) Operational problem with CMS - got message to prepare file families for new MC data. Today learned that we were receiving ~200 files before we had setup these.
- RAL - we had a problem on one of CEs that LHCb reported - quite a few jobs failing. Had to reboot machine and now things look fine.

AOB:

- GGUS tickets: often people write "fixed" when solved but no details. Please provide some details.
- Thursday is a holiday at CERN - hold meeting? If so need external volunteer to chair meeting.

Tuesday:

Attendance: local(Luca, Steve, Przemyslaw, Maarten, Edward, Doug, Harry, Jean-Philippe, Ale, Simone, Ignacio, MariaD, Dirk);remote(Massimo/CNAF, Jon/FNAL, Kyle/OSG, Michael/BNL, Joel/LHCb, Gang/ASGC, Ronald/NL-T1, Andreas/NDGF, Dimitri/KIT,Gareth/RAL Ian/CMS).

Experiments round table:

- ATLAS reports -
 - ◆ Tier-0:
 - ◇ No issues, no new data.
 - ◇ Problems with central DDM DashBoard, reporting from the central LFC not working. Expert looking at this, and his reply was that what usually takes 10 sec. is now taking 6000 sec. Not sure why at this time.
 - ◆ Tier-1:
 - ◇ NL Cloud coming back on-line. Cloud all on-line for data transfers, and now working through backlog. Cloud is still off-line for new data from tier-0, and jobs are still not getting scheduled. Hope to have them back tomorrow.
 - ◇ Taiwan Cloud now back on-line, and setup to get new data from tier-0. Some jobs failing yesterday, and a ggus about it (GGUS:61822), but we still don't get timely feedback from people there.
 - ◇ DE Cloud - LFC stopped working at FZK and all transfers stopped for the cloud. System reached a table space limit in Oracle, and this needed to get reset. Now fixed, and cloud is back on-line.
 - ◇ CCIN2P3 reports problems with AFS use there, not sure about details? No ticket on this it seems, ticket with details has been requested.
 - ◆ Tier-2:
 - ◇ In UK, Southgrid Cambridge HEP site has lost host gss credential. This is getting discussed in GridPP meeting today.
- ◆ Simone: IN2P3 LFC log shows many failure - could IN2P3 comment if this is related to AFS problems?
- ◆ Ale: found problem in the way SAM tests are run against BNL CE - test jobs seem to get stuck for unknown reason on WMS side. Investigation will continue with WMS experts, but unavailability due to test failure should not be attributed to BNL. A ticket will be created to track the issue.

Tuesday:

- CMS reports -
 - ◆ Central infrastructure
 - ◇ CMS will start producing the AOD when data taking resumes. This was always in the plan, but will be implemented now. Will result in modest increase (~10%) increase in the rate from CERN to Tier-1s.
 - ◆ Tier1 issues
 - ◇ Issue with tape families for MC at Tier-1s is related to the CMS move to doing more MC production at Tier-1s. Working on a solution.
 - ◆ Tier2 Issues
 - ◇ No issues to report.
 - ◆ MC production
 - ◇ Very large (1000M) MC production ongoing T1s and T2s; will start the reprocessing with CMS_SW 38XX soon.

- ALICE reports [↗](#) -
 - ◆ Pass1 reconstruction cycle LHC10e finishing. Starting a new LHC10f cycle in addition with a new bunch of TPC gain calibration jobs. Also a new MC cycle started during the night.
 - ◆ T0 site
 - ◇ Issue reported yesterday (about 2500 jobs running in the system, consuming 0 CPU) seems to be solved. Site is performing well today
 - ◆ T1 sites
 - ◇ SARA: Yesterday the site admin announced an update of the ALICE VOBOX foressen for 20/09. A backup of the full \$HOME area of the experiment needed before any operation. Green light of the experiment for the update
 - ◇ CCIN2P3: AFS problem reported by the site admin this morning. Currently all jobs are staying on the local queue and will be spawned on WN's as soon as AFS is back.
 - ◆ T2 sites
 - ◇ no remarkable issues to report

- LHCb reports -
 - ◆ Lot sof jobs failing because our Bookkeeping service was overloaded and was not able to serv requests.
 - ◆ T0 site issues:
 - ◇ CERN: Some VOBOXes have been blocked by SAM people without any warning nor discussion with LHCb.
 - ◆ T1 site issues:
 - ◇ CNAF: GGUS:61854 [↗](#) SRM failing (human error)
 - ◇ IN2P3 : GGUS:61874 [↗](#) gLexec failing
 - ◇ CNAF : GGUS:61876 [↗](#) gLexec failing
 - ◇ CERN : GGUS:61878 [↗](#) gLexec failing

Sites / Services round table:

- Massimo/CNAF: ntr
- Luca/CNAF: Still open: data transfer issues between BNL and CNAF - should do test with two dedicated servers over OPN. Luca will contact BNL to plan this joint test. SRM endpoint upgrade for all VOs planned for early next week.
- Jon/FNAL: ntr
- Michael/BNL : ntr
- Gang/ASGC : working on ATLAS ticket - will get back to VO soon
- Ronald/NL-T1: data for 3D databases transferred from RAL until tomorrow. Afterwards a short restart of the DBs will be required resulting in a few minutes of LFC and FTS outage for ATLAS.
- Andreas/NDGF: performing kernel upgrades (SRM outage)
- Dimitri/KIT: ntr

- Garethj/RAL: working on understanding throughput issues reported by LHCb. BDII updates were taking long, resulting sometimes in missing information. - now information is available but may get occasionally stale - issue is still being worked on.
- Kyle/OSG: ntr

AOB: (MariaDZ) The tool mentioned at yesterday's meeting that calculates progress of tickets when a specific site is notified is the **GGUS Ticket Timeline Tool** <https://gus.fzk.de/stat/ttt.php> linked from the GGUS homepage.

Wednesday

Attendance: local(Doug, Jean-Philippe, Zbyszek, Maria, Alessandro, Steve, Edward, Edoardo, Gavin, Marie-Christine, MariaD, Andrea);remote(Massimo/CNAF, Jon/FNAL, Michael/BNL, Onno/NLT1, Joel/LHCb, Foued Jrad, Rolf/IN2P3, Tiju/RAL, Andreas/NDGF, Kyle/OSG).

Experiments round table:

- ATLAS reports -
 - ◆ Ongoing issues
 - ◇ INFN-BNL network problem GGUS:61440 (Some discussion yesterday about this).
 - ◇ RAL-NDGF functional test GGUS:61306 (More failures happened again this morning).
 - ◆ Tier-0
 - ◇ No new data. (One stuck file in Castor, not a major issue.)
 - ◆ Tier-1
 - ◇ BNL had issues with their SRM today, and caused transfer failures in the US cloud. This was fixed within an hour or so, and things seem to be working there. Michael: it was not an SRM issue but the Name Server was stuck and had to be restarted (problem detected automatically).
 - ◇ Alessandro: SAM tests failing at BNL because they are submitted with Alessandro DN and lcgadmin role which is not recognized as high priority role by OSG. To be added to the high priority job queue. As it is not a site issue and as the CE SAM results are taken into account for site availability, Michael would like to see the statistics corrected on the MB slides (Maria agreed).
 - ◇ Jobs now being brokered to the NL cloud, although analysis is still offline there. Should be fully up tomorrow.
 - ◆ Tier-2
 - ◇ A number of site issues, all seem to be site specific, and not worth reporting (still they take up time...).
- CMS reports -
 - ◆ Experiment activity
 - ◇ Run tests on going
 - ◆ Central infrastructure
 - ◇ CMS will start producing the AOD when data taking resumes. This was always in the plan, but will be implemented now. Will result in modest increase (~10%) increase in the rate from CERN to Tier-1s.
 - ◆ Tier1 issues
 - ◇ Nothing to report
 - ◆ Tier2 Issues
 - ◇ No issues to report.
 - ◆ MC production

- ◊ Large abort rate for bulk submission in CMS production at T1/T2 signaled and ticket open: <https://savannah.cern.ch/bugs/?72423>
- ◊ Very large (1000M) MC production ongoing T1s and T2s; will start the reprocessing with CMS_SW 38XX soon.
- ◆ AOB
 - ◊ New CRC on duty: Marie-Christine Sawley
- ALICE reports [↗](#) -
 - ◆ GENERAL INFORMATION: Same situation since yesterday in terms of reconstruction and MC cycles.
 - ◆ T0 site
 - ◊ Pass1 reconstruction activities continuing with no remarkable issues to report
 - ◆ T1 sites
 - ◊ IN2P3: Problem reported yesterday and concerning the access to the AFS area (problems to access the area): SOLVED. site back in production
 - ◆ T2 sites
 - ◊ Usual operation activities with no remarkable issues
- LHCb reports -
 - ◆ Experiment activities:
 - ◊ Production jobs and user jobs are failing due to the corruption of our CONDDDB database. Patch being applied and will be propagated to the sites in the next hour.
 - ◆ Issues at the sites and services
 - ◊ T0 site issues:
 - CERN:
 - ◊ T1 site issues:
 - IN2P3 : GGUS:61904 [↗](#) transfer to lhcb_dst are failing (disk full although free space is reported). Rolf: looks similar to an Atlas problem seen a few weeks ago.
 - RAL : we confirm that we accept the upgrade of CASTOR proposed for the 27th of september.
 - NLT1: DBs ok, STREAMS set up but need to recover the missing data.

Sites / Services round table:

- CNAF: ntr
- FNAL: ntr
- BNL: ntr
- NLT1:
 - ◆ Oracle DB running happily on a single machine
 - ◆ STREAMS needs to catch up
 - ◆ Vendor investigating the cause of the problem
 - ◆ There was this morning a scheduled maintenance on one dCache disk pool node at SARA (successful).
- IN2P3: ntr
- RAL: ntr
- NDGF: ntr
- OSG: ntr
- KIT:
 - ◆ LHCb Cond DB: there was a short network interruption and Streams replication had to be restarted (no data loss).
 - ◆ PBS server rebooted. WNs restarted by engineer (no job lost).
- ASGC: CMS transfer problem reported yesterday was understood and solved. It's because the NFS system mounted on cmsvobox became invalid in the morning. We fixed it by remounting the NFS

system.

- CERN:
 - ◆ Edoardo: tests between BNL and CNAF are ongoing (they started at 15:00).

AOB:

Thursday

Attendance: local();remote().

Experiments round table:

- ATLAS reports -
- CMS reports -
- ALICE reports [↗](#) -
- LHCb reports -

Sites / Services round table:

AOB:

Friday

Attendance: local(Harry(chair), Doug, Steve, Roberto, Marie-Christine, Przemek, Maarten, Alessandro);remote(Alain(OSG), Dimitri(KIT), Michael(BNL), Jon(FNAL), Onno(NL-T1), Massimo(CNAF), Gareth(RAL), Rolf(IN2P3), Gang(ASGC), Gonzalo(PIC)).

Experiments round table:

- ATLAS reports -
- Tier-0:
 - ◆ No new data, things quiet.
 - ◆ Problems with DDM Dashboard, and reporting of transfer errors went away for 3-4 hours. An update agent that normally takes 10sec. was taking 2-3 hours. Cause is uncertain, and we had to just wait it out.
 - ◆ The Express Stream is now being transferred to Tier-1 sites, this will increase data transfer to tier-1 a few percent for new data. Started today, existing express stream data is being transferred, so far without problems.
- Tier-1
 - ◆ All issues with NL cloud outage resolved, all service back online at this time.
 - ◆ Transfer timeout errors continue between RAL and NDGF, there is an open ticket on this starting its fourth week now. Gareth reported that he knows this ticket is being worked on and he will try and update it with more information.
- Tier-2
 - ◆ Many tier-2 issues, but all rather specific to each site. Little commonality.

- CMS reports -

Experiment activity: Run tests on going

Central infrastructure: CMS will start producing the AOD when data taking resumes. This was always in the plan, but will be implemented now. Will result in modest increase (~10%) increase in the rate from CERN to Tier-1s.

*T0/CAF * problem with one diskerver node yesterday evening; node lxsrl5106 had to rebooted around 22.00. see Team ticket https://gus.fzk.de/ws/ticket_info.php?ticket=61966all

Ggus Ticket https://gus.fzk.de/ws/ticket_info.php?ticket=61706 has been opened for large number of jobs failing with Maradona errors at CERN. Will report further on Monday.

Tier1 issues: Nothing to report

Tier2 Issues:

T2-TR_METU (Turkey) has been unavailable since more than 48 hours (all SAM tests failing)

T2-EE-Estonia has BDII and Job robot failure (was previously in scheduled downtime).

MC production:

Solved: Large abort rate for bulk submission in CMS production at T1/T2 signaled and ticket open: <https://savannah.cern.ch/bugs/?72423>, bug of CREAM CE ;

Very large (1000M) MC production ongoing T1s and T2s; reprocessing with CMS_SW 38XX has started.

- ALICE reports -

GENERAL INFORMATION: Massive MC production with 7 active MC cycles.

T0 site: Specific CREAM-CE tests executed this morning with no incidents to report

T1 sites: All T1 sites in production

T2 sites: Usual procedures needed at several T2 sites and followed directly with the site admins

- LHCb reports -

[Experiment activities]:

- Production jobs and user jobs are failing due to the corruption of our CONDDDB database. Problem fixed in LHCb. No huge production activity on going, mainly users.

Issues at the sites and services

- T0 site issues:
 - ◆ CERN: CREAMCE failing all pilots. GGUS:61957 due to CE log files becoming bigger than 2 GB. Maarten reported there is already a bug open for this and this was due to CERN having set the logging level as high to debug other issues.
 - ◆ We had the lhcbraw class overloaded on Thursday due to a super user running a burst of jobs accessing EXPRESS data. For a couple of hours SLS was reporting the service unavailable.

- T1 site issues:
 - ◆ CNAF: ConditionDB at CNAF times out connections from other sites WNs (and jobs got stalled). GGUS:61989. Apparently listener only set to listens to the OPN.
 - ◆ IN2P3 : GGUS:61904 transfer to lhcb_dst are failing. Moved disks from another space token to allow jobs to write in.
 - ◆ RAL (SARA) Observed failures in the Oracle Streams Apply process for LFC (since resolved). The problem seems to be due to some entries in the T1 instance at RAL (and SARA as consequence) that do not have the corresponding ones in the central catalog at CERN. Digging into details it looks like some DNs and FQANs have been added manually in the users and groups tables locally on the site to allow UKI NGI to test the LFC (a ticket dealing with this problem is GGUS:60618). It is worth to remind that any update done on the read-only instance of LFC at T1' s becomes an inconsistency in the replication and would compromise the whole replication of information and then **has to be avoided**. Sites should ask central LHCb operation managers to add new users or groups as done for IN2p3.

Sites / Services round table:

- KIT: Responding to ggus ticket 61636 from LHCb about pilot jobs failing when the CREAM-CE gets a timeout connecting to the PBS server a workaround has been put in place in the submission script. They have been in touch with the developers (as recommended by Maarten for all such issues).
- CNAF: They are working to open the worker node GPN ports which are blocking access to the conditions DB reported by LHCb.
- ASGC: SIR on ASGC DB issue last week has been updated to <https://twiki.cern.ch/twiki/bin/view/LCG/WLCGServiceIncidents>.
- RAL: Concerning the high LHCb load on their disk servers they are now running stably at 800 batch jobs and will stay this way over the weekend. Roberto reminded this was due to an exceptional amount of merging jobs that should finish next week after which RAL could return to a normal number of job slots.
- IN2P3: One of their CE (for CMS) stopped at 20.00 yesterday and was restarted at 01.00. Existing jobs continued but no new ones started. Still under investigation.
- Streams Replication: An apply process has just crashed in Taiwan with an apparently corrupted database file. CERN could not yet contact anyone there. Gang reported that he was aware and the problem has been forwarded to their site operations. He is expecting a reply soon. Streams propagation to SARA is now up to date but now has to be merged back into the main streams replication.

AOB:

-- JamieShiers - 03-Sep-2010

This topic: LCG > WLCGDailyMeetingsWeek100906

Topic revision: r11 - 2010-09-10 - HarryRenshall



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