

Table of Contents

| | |
|--|----------|
| Week of 110711..... | 1 |
| Daily WLCG Operations Call details..... | 1 |
| WLCG Service Incidents, Interventions and Availability, Change / Risk Assessments..... | 1 |
| General Information..... | 1 |
| Monday:..... | 1 |
| Tuesday:..... | 2 |
| Wednesday..... | 4 |
| Thursday..... | 6 |
| Friday..... | 7 |

Week of 110711

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 (AndreaV, Lukasz, Eva, Stephen, Jarka, Lola, Miguel, MariaD); remote (Jon, John, Alexander, Paolo, Rolf, Shu-Ting, Rob).

Experiments round table:

- ATLAS reports -
 - ◆ INFN-T1: DATATAPE spacetoken ran out of space, GGUS:72473 raised to ALARM. Now seems back. Can site please comment? [Paolo: minor glitch, solved this morning. Will reply on GGUS and close the ticket.]
 - ◆ Powercut at CERN on Sunday 10th July afternoon -- distribution fault on LHC1 electric network. ATLAS magnets suffered, people working on bringing the whole detector back (matter of few days or more).
- CMS reports -
 - ◆ LHC / CMS detector
 - ◇ Ramping up detector for end of technical stop
 - ◇ Power failure yesterday caused magnet to trip. Ramping also now, will take a few days (to also cool down).
 - ◆ CERN / central services
 - ◇ Lost BDII at CERN again. Fixed again.

◇ Some files not migrated to tape after four days. GGUS:72497 [↗](#). [Miguel: looking into this.]

◆ T1 sites:

◇ Tickets still open with KIT slow transfers: Savannah:122051 [↗](#), Savannah:122052 [↗](#), Savannah:122083 [↗](#)

· Edit Xavier Mol (KIT): Even two more Tickets Savannah:122035 [↗](#), Savannah:122088 [↗](#)

◆ T2 sites:

◇ NTR

◆ AOB:

◇ NTR

• ALICE reports [↗](#) -

◆ T0 site

◇ GGUS:72484 [↗](#) : voms-proxy-init gave DB errors contacting lcg-voms.cern.ch.

Several sites failed to renew the proxy. It is a known instability of the ALICE VOMS servers which will be solved with the new upgrade. SOLVED

◆ T1 sites

◇ Nothing to report

◆ T2 sites

◇ Usual operations

• LHCb reports -

◆ No report (many people attending the WLCG Workshop in Hamburg).

Sites / Services round table:

• FNAL: ntr

• RAL: ntr

• NLT1: seeing a lot of LHCb requests for staging files from tape, the disks are getting full so the requests may become slower.

• CNAF: nta

• IN2P3: ntr

• ASGC: ntr

• OSGtr

• KIT (Xavier Mol):

◆ Correction of last week's announced downtime for one tape library - will be done tomorrow 10:00 through 14:00 local time. We will do firmware update of one tape library. As usual, reading is not possible for tapes from that library, writing will be done with other libs.

◆ Today around 5:30 local time, hick-up in the GPFS connection for both stage pools in the gridka-dcache.fzk.de dCache instance. Both pools were offline until restart at 08:40.

• Dashboard: a few machines were down this weekend dur to the power cut, now all up and running

• Databases: ALICE, CMS and LHCb online databases were down thiw weekend due to the power cut, now all ok

AOB: none

Tuesday:

Attendance: local (AndreaV, Jarka, Lukasz, Massimo, Ignacio, Lola, Peter, Eva, MariaD, Ricardo, Shu-Ting); remote (Jon, Xavier, John, Rolf, Paco, Vladimir, Rob, Paolo, Andreas).

Experiments round table:

Monday:

- ATLAS reports -
 - ◆ CERN-PROD: Some files on ATLASDATADISK not accessible. GGUS:72528 raised to ALARM ticket. Since this issue is long-lasting, and ATLAS needs to process affected data. List of affected files provided. Can CERN please comment on this issue? [Massimo: produced list of files, but also still trying to recover the machine. There is a high risk of losing files if the machine is given to the vendor for repairs, so ATLAS should decide what to do. Jarka: got the list of files, checking if the CERN copies are unique copies or there is another copy elsewhere. Will then decide whether/when to give the go-ahead for the vendor intervention.]
 - ◆ Powercut at CERN on Sunday 10th July afternoon -- ATLAS may be back tomorrow.

- CMS reports -
 - ◆ LHC / CMS detector
 - ◇ Once LHC will be back in collisions mode, CMS will run with magnet off for at least 2 days. CMS getting prepared for running in that mode.
 - ◆ CERN / central services
 - ◇ Some files not migrated to tape after four days. CASTOR team stated they found the problem and will deploy a permanent solution, see GGUS:72497. [Ignacio: problem is due to missing database constraints. This probably happened during a recent export/import intervention that was done to improve performance. Apologies about the problem. Now fixed the issue for the specific files that were affected, but the more general problem of missing constraint is there and could affect other files. Need a 2h intervention with the database offline to recover completely (joint intervention with the DBAs from IT-DB). Propose 2pm tomorrow. Peter: data taking will soon resume, but tomorrow should still be ok, and this intervention is needed. There is also an SRM intervention scheduled for tomorrow, could reuse the same downtime. Will circulate the proposal to T0 experts and confirm this.]
 - ◆ T1 sites:
 - ◇ Decided in yesterday's weekly CMS Facilities Operations meeting to make a collaborative effort between CMS and KIT in trying to understand the transfer rate limitations observed lately. It is not completely clear if the observed low transfer rates were due to a large number of pool-to-pool copies from sites with different performance profiles using the STAR channel, or due to insufficient capacity to drive production and transfers at the same time, or a missing patch at 2 Gridftp doors, or a combination of all that.... To be investigated.
 - ◆ T2 sites:
 - ◇ NTR
 - ◆ AOB:
 - ◇ New CMS CRC-on-Duty for 2 weeks : Peter Kreuzer

- ALICE reports -
 - ◆ T0 site
 - ◇ Nothing to report
 - ◆ T1 sites
 - ◇ Nothing to report
 - ◆ T2 sites
 - ◇ Usual operations

- LHCb reports -
 - ◆ Experiment activities: Restripping and MC jobs. [Vladimir: this is a massive operation and is causing the overload of stagers everywhere.]

Sites / Services round table:

- FNAL: ntr
- KIT:
 - ◆ this morning at 3am one power supply was malfunctioning; this affected ATLAS pools, was fixed at 10am
 - ◆ this morning at 5am there were also problems with several CEs; still working on one of them
- RAL: ntr
- IN2P3: ntr
- NLT1: ntr
- OSG: there will be a short 5 minute intervention in approximately 45 minutes
- NDGF:
 - ◆ dcache was upgraded 20 minutes ago, was in GOCDB, now completed
 - ◆ planned outage tomorrow from 16 to 17 for hardware replacement, network to Finland will be down, this will affect ALICE
 - ◆ one storage server failed this morning, some ATLAS data was unavailable
- ASGC: ntr

- Databases: one new node was added to LCGR, users were contacted to restart their applications
- Storage: nta
- Dashboard: ntr
- Grid: ntr

AOB:

- MariaD: waiting for suggestions from Jon/Rob about the GGUS ALARM tests. Rob: 2pm Eastern time, 1pm central time will be ok.

Wednesday

Attendance: local (Andrea, Massimo, Alessandro, Eva, Peter, Lukasz, Jarka, MariaD, Lola); remote (Alexander, Rob, Jon, Rolf, Tiju, Andreas, Paolo, Shu-Ting, Vladimir).

Experiments round table:

- ATLAS reports -
 - ◆ CERN-PROD: Issue with symlink CERN-PROD_TZERO, GGUS:72585 [↗](#). Now fixed, thank you!
 - ◆ VOMS: Hit bug described in GGUS:72595 [↗](#) -- I have 2 certificates with same member DN, but different CA DN. Was not able to submit TEAM nor ALARM tickets to GGUS this morning. GGUS replaced DN list from this morning by yesterday's list, it fixed issue with GGUS ticket submissions. I've removed old (about to expire) certificate from VO, hopefully this will "fix" my issues. This bug is supposed to be fixed in a new VOMS release. Is there any plan to upgrade ATLAS VOMS at CERN? Thank you very much! [!MariaD: the daily VOMS-GGUS synchronisation gave an error today. GGUS developer Helmut Dres restored yesterday's TEAMers' and ALARMers' groups to continue operation. Steve identified a known and fixed VOMS bug (tolerates 2 certificates with identical DNs from 2 different CAs). The fix will be put in production in the near future. Should be discussed with Steve when he is back.]
- CMS reports -
 - ◆ LHC / CMS detector
 - ◇ CMS started collision data taking with magnet off. Besides a few application failures (DQMHarvest), so far so good. CMS Magnet will stay off until Thu evening.
 - ◆ CERN / central services

- ◇ All CMS data transfers from/to CERN failing this morning, with FTS error. Opened GGUS:72588 (Team), but no answer. Probably same issue as observed by ATLAS (GGUS:72585)
- ◇ CMS ready for CASTORCMS intervention today. Stopped T0 processing for 2 hours (<http://itsb.web.cern.ch/planned-intervention/castor-cms-maintenance>).
[Massimo: DB part of upgrade is completed, services are restarting.]
- ◇ CMS migrated 550TB data from CASTOR (CMSCAF) to EOS. It would be interesting to know how often data "reader" falls back to castor (via re-director).
[Massimo: took note, indeed it will be interesting to monitor this.]
- ◆ T1 sites:
 - ◇ NTR
- ◆ T2 sites:
 - ◇ NTR
- ◆ AOB:
 - ◇ NTR
- ALICE reports -
 - ◆ T0 site
 - ◇ CAF:
 - Problem with master CAF node lxbsq1409. During user session some files are created in user's sandbox on master node and workers in /tmp. From time to time (this happens randomly) these files have as UID some random number which can not be resolved which can not be access by user due to permission denied. After a bit of investigation, came out that this could be due to a corrupted file system. Under investigation
 - spma_ncm error reports a Hardware DBus error while running smartd component in a couple of CAF nodes (lxbsq1409 and lxbsq1410). This started happening since last week after and update of 12 packages. SNOW ticket INC052246
 - [Andrea: could these two issues be related, they both concern lxbsq1409. Lola: could be, the two timestamps are the same.]
 - ◆ T1 sites
 - ◇ Nothing to report
 - ◆ T2 sites
 - ◇ Usual operations
- LHCb reports -
 - ◆ Experiment activities: Restripping and MC jobs.

Sites / Services round table:

- NLT1: ntr
- OSG: ntr
- FNAL: ntr
- IN2P3: Rolf on holiday next week, will be replaced by Mark.
- RAL: ntr
- NDGF: ntr
- CNAF: ntr
- ASGC: ntr

- Dashboard: ntr
- Databases: ntr
- Storage: nta

AOB:

- MariaD: on holiday for a month from tomorrow. Please use the normal contacts to report GGUS issues, i.e. open a GGUS ticket asking for assignment to the 'GGUS' Support Unit or (in case of GGUS total unavailability) write to ggus-info@cernNOSPAMPLEASE.ch

Thursday

Attendance: local (AndreaV, Steve, Jarka, Alessandro, Massimo, Lukasz, Stefan, Peter, Maarten); remote (Michael, Xavier, Jon, Ronald, Jos, Tiju, Shu-Ting, Rob, Paolo).

Experiments round table:

- ATLAS reports -
 - ◆ RAL-LCG2: [SRM_ABORTED] errors due to high load on castor at RAL. Throttled back FTS and Batch farm limits and SRM appears to have recovered. Thank you! GGUS:72618 [↗](#). [Tiju: removed limits, throttling still in place.]
 - ◆ [Jarka: ATLAS is up and running today since 1pm after the power cut.]
- CMS reports -
 - ◆ LHC / CMS detector
 - ◇ CMS magnet should be ramped up at around 12:30. In the afternoon will switch back HLT menu to normal data taking.
 - ◆ CERN / central services
 - ◇ Problem setting the "lense symbol" in the CMS SSB, see Savannah:122212 [↗](#)
 - ◆ T1 sites:
 - ◇ NTR
 - ◆ T2 sites:
 - ◇ NTR
 - ◆ AOB:
 - ◇ NTR
- ALICE reports [↗](#) -
 - ◆ T0 site
 - ◇ Instabilities in ALICE CAF are still under investigation
 - ◆ T1 sites
 - ◇ Nothing to report
 - ◆ T2 sites
 - ◇ Usual operations
- LHCb reports -
 - ◆ Finishing up on old stripping productions.
 - ◆ New GGUS (or RT) tickets:
 - ◇ T0: 0
 - ◇ T1: 0
 - ◇ T2: 0
 - ◆ Issues at the sites and services
 - ◇ T0
 - [Stefan: read-write instance of LFC at CERN is being moved to SLC5 as we speak.]
 - ◇ T1
 - Gridka: Problems with file staging because of problems with disk pools. The copying from the staging pools into the final pools does not work properly. Currently being investigated.

- CNAF: Short glitch with the shared software area because of broken network card. Discovered and solved this morning
- RAL: New increased disk pools for Tape have been put into production. Old pools are in draining.

Sites / Services round table:

- BNL: need an intervention today on the SRM database. Will require outage of 1h, will apply measures to limit job/transfer failures, will inform ATLAS via elog.
- PIC: ntr
- FNAL: ntr
- NL-T1: ntr
- KIT:
 - ◆ staging possible again as reported by LHCb
 - ◆ still recovering from a disk outage on Tuesday, causing issues for ATLAS dcache, being worked on
- RAL: nta
- ASGC: ntr
- OSG: early morning call from an ATLAS user that resulted in ticket GGUS:72638 [↗](#). [Alessandro: were not aware of it, now having a look. This is not a standard ATLAS procedure that a single user may contact you, will follow up offline.]
- CNAF: ntr
- Grid services [Steve]:
 - ◆ **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. Notification will be posted to GOCDB and CERN SSB after tomorrow's meeting to allow any feedback. The tier2 FTA channel agent nodes remain mostly on SL4 at this stage. Upgrading to SL5 is more pressing now as the CERN SL4 infrastructure begins to degrade, e.g. yesterday's GGUS:72585 [↗](#) as reported by atlas.
 - ◆ **CERN - CERN VOMS** concerning yesterday's comments on VOMS and the presence of GGUS:72595 [↗](#) which causes problems resolving group members when a user exists with same DN but different CA. The user in question has removed their redundant identity so the situation is alleviated. Upgrading to VOMS from the last gLite version to the first EMI version is wanted anyway shortly to resolve VOMS's leaking oracle cursors. Estimate 2nd half of August and will be transparent with easy rollback. [Andrea: is there any change required on the client side? Steve/Maarten: no, client is compatible (extensively tested), and there is also no schema change.]
- Dashboard services: ntr
- Storage services [Massimo]:
 - ◆ Upgraded ATLAS and CMS instances of EOS. Main motivation is better fail over from one head node to another. Also better self-healing and handling of file copies.
 - ◆ ATLAS instance of CASTOR is presently overloaded on the database. Under investigation.

AOB: MariaD: on holiday for a month from today. Please use the normal contacts to report GGUS issues, i.e. open a GGUS ticket asking for assignment to the 'GGUS' Support Unit or (in case of GGUS total unavailability) write to ggus-info@cern.NOSPAMPLEASE.ch

Friday

Attendance: local(Alessandro, Jarka, Maria, Jamie, Peter, Michal, Romain, Stefan, Jacek, Ricardo, Miguel, Stephane, Ueda);remote(Michael, Jon, Xavier, Maria Francesca, Tiju, Onno, ShuTing, Rob).

Experiments round table:

- ATLAS reports - Security issue at a Tier2 site. ATLAS is actively investigating together with CERN and WLCG security contacts. [Romain - incident affecting 2 T2s in CA. Sites given pattern to check for. Incident not contained nor understood. All security teams still actively investigating. No evidence that hacker used or abused ATLAS grid. Ale - identified the users that accessed these sites. Pilot jobs with different roles. Should we always revoke cert? A: No, not in general. Will be stated explicitly. Romain - priority was that ATLAS stopped sending jobs there asap which was done. Jarka - this sites was in scheduled downtime since a while.]
 - CMS reports -
 - LHC / CMS detector
 - ◆ CMS resuming normal collisions+Magnet-On data taking
 - CERN / central services
 - ◆ 1 Dashboard monitoring issue :
 - ◇ Second occurrence of RRC-KI (Kurchatov Institute) that is in official downtime, yet not appearing as such on the dashboard downtime monitoring, hence neither on the CMS Downtime GCal that now inherits from the former, see Savannah:121625
 - T1 sites:
 - ◆ NTR
 - T2 sites:
 - ◆ Positive note : CMS T2 Readiness very stable since 1 month : 84% sites with readiness fraction > 80% !
 - AOB:
 - ◆ NTR
-

- ALICE reports [↗](#) -
- LHCb reports -
 - ◆ Finishing up on stripping productions
 - ◆ Validation of stripping productions for new data went fine. Express has picked up new data.
 - ◆ T0
 - ◇ Migration of LFC R/W instances yesterday to slc5 virtual machines went smooth
 - ◆ T1
 - ◇ Gridka: Problems with staging of files for stripping production (GGUS:72682 [↗](#))
 - ◇ NL-T1: Problems with staging of files as of tonight (GGUS:72666 [↗](#)). Jobs started again slowly to access files and running late morning.

Sites / Services round table:

- BNL - ntr
- FNAL - ntr
- PIC - ntr
- NDGF - ntr
- RAL - yesterday 15:00 - 17:00 ATLAS castor instance down due to Oracle problem.
- ASGC -
- NL-T1 - SARA SRM issue: not only with staging of files but with all SRM transactions. At first we thought just bad performance due to load. Lots of timeouts but also some successful transactions. Didn't find cause but then around noon today restarted dCache on SRM node and then all was fine. Probably dCache just got stuck and a restart fixed that. Apart from GGUS ticket mentioned by LHCb another ticket on this issue: GGUS:72663 [↗](#). Unless we hear otherwise we will close. Announcement: between Aug 21 - 26 we have maintenance on MSS. In this week no tape access. Will submit to GOCDB soon.

WLCGDailyMeetingsWeek110711 < LCG < TWiki

- KIT: Still working on a broken disk-only pool of ATLAS (f01-032-105-e_1D_atlas). The filesystem is corrupt and a possible fix will take at least until Monday.
- OSG - ntr

- CERN - ntr

AOB:

-- JamieShiers - 27-Jun-2011

This topic: LCG > WLCGDailyMeetingsWeek110711

Topic revision: r19 - 2011-07-19 - AndreaValassi



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](#)