

Table of Contents

Week of 110905.....	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.....	6
Thursday.....	8
Friday.....	8

Week of 110905

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(Claudio, Daniele, Edward, Eva, Fernando, Ivan, Jan, Maarten, Maria D, Steve);remote(Dimitri, Giovanni, Jhen-Wei, Onno, Roger, Rolf, Tiju, Vladimir).

Experiments round table:

- ATLAS reports -
 - ◆ T0/Central Services
 - ◇ DDM Deletion machine stuck (machine not responding): Instance of DDM Deletion service stopped between 11:00 and 21:00.
 - ◆ T1 sites
 - ◇ Many transfer errors to/from Tier1s, seemed caused by high load. Specially MCTAPEs involved.
 - Transfer errors to TAIWAN-LCG2_MCTAPE: destination file failed on the SRM with error [SRM_ABORTED]. GGUS:74039
 - Transfer errors to RAL-LCG2_MCTAPE: destination file failed on the SRM with error [SRM_ABORTED]. GGUS:74041
 - SARA-MATRIX: Get error: dccp failed with output. GGUS:74043
 - PIC_DATADISK transfer failed to T1s and T2s with error AsyncWait. GGUS:74045
 - BNL-OSG2_DATADISK to FZK-LCG2_MCTAPE: GGUS:74050

- Maarten: is this mismatch between MCTAPE capacity and actual usage a cause for concern?
- Fernando: the issue will go away when the merge of MCTAPE with DATATAPE has finished (currently in progress)
- ◇ RAL had a network problem on Monday 5:00AM and no service could be contacted for some hours. AMOD blacklisted whole UK cloud. [GGUS:74052](#)
- ◇ Alarm ticket submitted for CERN by T0 expert about RAW files disappeared from CASTOR/T0ATLAS [GGUS:74053](#)
 - Jan: recent files were garbage-collected, probably based on stale information, as we also experienced about 2 months ago; currently only 1 file system on 1 server has been affected; we are looking into it
- ◆ T2 sites
 - ◇ NTR
- CMS reports -
 - ◆ LHC / CMS detector
 - ◇ Friday 2nd: end of TS (19:00), beam recovery started
 - ◇ Saturday 3rd: first ramp after TS, started working with 2 bunches/beam, 3.5 TeV, beta* 1m
 - ◇ Sunday 4th: LHC setting up. Some problem with extraction magnets in SPS, beams colliding at 16h00
 - ◇ Monday 5th: beam set-up. Collisions expected later today, but no stable beams.
 - ◆ CERN / central services
 - ◇ update on "myproxy not working from outside CERN" ([GGUS:73926](#)): No issues observed/reported over the weekend. I think Maarten can close the GGUS. We confirm we can operate with no problem now. We are interested in being kept in the loop on when the myproxy bug(s) get solved, thanks.
 - Maarten/Steve: the bug was acknowledged by the lead developer, who proposed a simple patch; some discussion then followed on the desired behavior for renewal/retrieval policies and in the end a different patch was proposed by us and accepted; the patch has been tested and is expected to be deployed on [myproxy.cern.ch](#) in a matter of days, probably the week after next week (will be announced)
 - ◇ issues with the [cmsonline.cern.ch](#) web server. IT-DB not involved in my (nightly) checks, but opened [GGUS:74033](#) to make sure somebody would have looked the following morning to confirm. Closed soon after. The issue is a CMS issue, being investigated for a stable fix still.
 - ◆ T1 sites:
 - ◇ NTR
 - ◆ T2 sites:
 - ◇ NTR
- ALICE reports -
 - ◆ T0 site
 - ◇ Nothing to report
 - ◆ T1 sites
 - ◇ CNAF: the number of jobs at the site has decreased during the past couple of weeks. The AliEn services are working smoothly. Our fair share is around 2K and we are running ~300 jobs. Site experts will look into it.
 - ◆ T2 sites
 - ◇ Nothing to report
- LHCb reports -
 - ◆ Experiment activities:

- ◊ Waiting for data.
- ◆ New GGUS (or RT) tickets:
 - ◊ T0: 0
 - ◊ T1: 1
 - ◊ T2: 0
- ◆ Issues at the sites and services
 - ◊ T1 :
 - RAL : Network connection problem (Solved)
 - ◊ Vladimir: there is a long-standing ticket GGUS:73177 [↗](#) about job failures at CERN due to timeouts in the "SetupProject" command; the correlation with big (48-core) worker nodes seems to be no longer present

Sites / Services round table:

- ASGC - ntr
- CNAF
 - ◆ at-risk intervention on router this **Wednesday** (17-18 UTC)
- IN2P3
 - ◆ downtime for all services on Tue Oct 4, at least for half a day, with many services affected the whole day
- KIT - ntr
- NDGF
 - ◆ today's dCache upgrade went OK
- NLT1
 - ◆ downtime for all services on Tue Sep 13
 - ◆ downtime around the end of the month for dCache upgrade to new golden release, date to be decided
 - ◊ Maarten/Vladimir: before the end of the month would be best for LHCb, to avoid interference with the reprocessing campaign that should start around that time
- RAL
 - ◆ site router failure lasted 7 hours from 01:00 to 08:00 local time
- CASTOR/EOS
 - ◆ urgent EOS xrdcp update applied for ATLAS; also foreseen for CMS on Tue afternoon (transparent)
- dashboards - ntr
- databases
 - ◆ ATLAS replication to T1 sites was affected by RAL network problem; the problem has been fixed and the backlog is being processed
- GGUS/SNOW
 - ◆ Maria D will give a short presentation in the IT-CMS meeting on Wed about the various channels for reporting problems to CERN-IT; as the presentation is not specific to CMS, it will also be linked to this page for convenience
- grid services - ntr

AOB:

Tuesday:

Attendance: local(Alessandro, Claudio, Eva, Guido, Jan, Maarten, Maria D, Mike, Steve);remote(Dimitri, Elizabeth, Gonzalo, Jeremy, Jhen-Wei, Jon, Michael, Roger, Rolf, Ronald, Tiju, Vladimir).

Experiments round table:

Tuesday:

- ATLAS reports -

- ◆ T0/Central Services

- ◇ Castor garbage collector acting on stale information. Problem fixed temporarily, an ultimate solution on its way by developers. GGUS:74053 [↗](#) similar to old (July) GGUS:72709 [↗](#)

- Jan: the next release may not yet address the issue, the stale monitoring is not yet understood; as a workaround the file class could be changed to a persistent type or files could individually be marked protected while they have not been processed
 - Jan/Alessandro: ATLAS probably will not have to resort to such measures, since the online buffers are large enough to handle **occasional** glitches (files are not deleted there until they are seen to have been processed by the T0 and their data is stored on tape)

- ◆ T1 sites

- ◇ Many transfer errors to/from Tier1s, seemed caused by high load. Specially MCTAPEs involved.

- Transfer errors to RAL-LCG2_MCTAPE: GGUS:74041 [↗](#) in progress, site to add new hardware during week, still errors
 - Transfer errors to TAIWAN-LCG2_MCTAPE: high load, site somehow fixed, but now no activity GGUS:74039 [↗](#)
 - SARA-MATRIX: site fixed the situation (2 servers stuck, restarted) GGUS:74043 [↗](#)
 - PIC_DATADISK: fixed by site and verified GGUS:74045 [↗](#)
 - BNL-OSG2_DATADISK to FZK-LCG2_MCTAPE. Again high load, should be fixed now but no activity to test GGUS:74050 [↗](#)

- ◇ RAL had a network problem on Monday 5:00AM and no service could be contacted for some hours. AMOD blacklisted whole UK cloud. Cloud back online already on monday late morning GGUS:74052 [↗](#)

- ◆ T2 sites

- ◇ NTR

- CMS reports -

- ◆ LHC / CMS detector

- ◇ Stable beams with 200 bunch filling expected during the night or tomorrow morning
 - ◇ 250 Hz of zero-bias for Pixels test claimed to be ok for the T0

- ◆ CERN / central services

- ◇ Several Castor related issues:

- Three CMS RAW data files stuck in export from CASTOR at CERN to T1s [t1transfer pool]
 - GGUS:74072 [↗](#)
 - ◆ -> solved
 - Many rfcps hangs again [on t0streamer pool]
 - GGUS:74085 [↗](#)
 - ◆ -> last update:

The underlying problem(s) are not resolved:

- * we will need a new CASTOR release (ETA 1..2 weeks, severe)
- * we have tried to implement workarounds (i.e. auto-restart running out of file descriptors; implement timeouts for these seem not to work in all cases (i.e. we still generate). For the moment, killing and restarting "stuck" rfcps seem to be a workaround - and open a ticket for the cases where this is not possible (i.e. stuck d2d transfers)

- ◆ Jan: each pool can have a different timeout

· Furthermore there are glitches in the quality of the transfers to Tier-2s too.

See e.g. today at 9.00 UTC:

https://cmsweb.cern.ch/phedex/debug/Activity::QualityPlots?graph=quality_all&entity

- ◆ T1 sites:
 - ◇ 4 hours intervention on CNAF MSS tomorrow morning. No recalls/migration
- ◆ T2 sites:
 - ◇ NTR
- ALICE reports [↗](#) -
 - ◆ T0 site
 - ◇ Nothing to report
 - ◆ T1 sites
 - ◇ IN2P3: The problem seems to be with the version of openssl used by AliEn.
 - added after the meeting: reverting to the previous openssl (from 0.9.8o back to 0.9.8l) **solved** the problem!
 - ◇ CNAF: the issue of the decrease of the jobs have been solved. The old configuration of the WNs have been restored and the problem seems to be solved
 - ◆ T2 sites
 - ◇ Usual operations
- LHCb reports -
 - ◆ Experiment activities:
 - ◇ Waiting for data.
 - ◆ New GGUS (or RT) tickets:
 - ◇ T0: 0
 - ◇ T1: 0
 - ◇ T2: 0
 - ◆ Issues at the sites and services
 - ◇ Nothing to report

Sites / Services round table:

- ASGC
 - ◆ various CASTOR operations timed out in the Transfer Manager (internal scheduler), under investigation
- BNL - ntr
- FNAL - ntr
- GridPP - ntr
- KIT - ntr
- NDGF - ntr
- NLT1
 - ◆ network maintenance Tue Sep 13 for SARA **and** NIKHEF; NLT1 computing capacity reduced to ~10% during the intervention
- OSG - ntr
- PIC
 - ◆ downtime next Tue Sep 13 for dCache upgrade + network intervention
- RAL
 - ◆ at-risk downtime tomorrow (Wed Sep 7) to apply Oracle patch
- CASTOR/EOS
 - ◆ ATLAS EOS unstable, waiting for the next release to fix that
 - ◇ Alessandro: from Tue Sep 13 onward ATLAS will fully rely on EOS for disk-only storage; in the last days about 25% of the time EOS was found **unavailable** when a connection was attempted from the analysis queue, while for CASTOR it was 0%

- ◆ CMS EOS upgrade on hold, because the foreseen improvements are not urgent
- dashboards
 - ◆ looking into ATLAS SRM test failures for KIT
 - ◇ Alessandro: sometimes the results are "none" only for KIT; there is an issue on the SAM-Nagios side as well as on the Dashboard side
- databases
 - ◆ during the night ATLAS replication to IN2P3 was affected by a memory issue at IN2P3 which has been fixed in the meantime
- grid services
 - ◆ px306.cern.ch is available for testing the fix for the MyProxy policy regular expression issue encountered by CMS; tested OK by IT-ES and IT-PES
 - ◆ during the last 5 or 6 days some 100 to 200 WN crashed per night due to excessive swapping and/or disk usage; at least some class of ALICE (user) jobs are involved; under investigation

AOB: (MariaDZ) The presentation for the IT/CMS meeting to take place tomorrow (see last bullet in yesterday's notes) is available in
<https://twiki.cern.ch/twiki/pub/LCG/VoUserSupport/IT-CMS-Tickets-20110907.ppt> or
<https://twiki.cern.ch/twiki/pub/LCG/VoUserSupport/IT-CMS-Tickets-20110907.ppt>

Wednesday

Attendance: local(Claudio, Guido, Ivan, Jan, Maarten, Maria D, Steve, Zbyszek);remote(Daniele, Jeremy, Jhen-Wei, Jon, Kyle, Michael, Pavel, Roger, Rolf, Ron, Tiju).

Experiments round table:

- ATLAS reports -
 - ◆ T0/Central Services
 - ◇ Synchronization of Castor and EOS was proceeding slowly, had a chat with the experts and better defined procedures/strategies
 - ◆ T1 sites
 - ◇ NTR
 - ◆ T2 sites
 - ◇ minor problems promptly spotted by shifters and fixed by the sites
- CMS reports -
 - ◆ LHC / CMS detector
 - ◇ Taking data. Pixel test mentioned yesterday will happen later today
 - ◆ CERN / central services
 - ◇ cmsdoc DNS alias problem (GGUS:74099): cmsdoc was not accessible from outside CERN. This created problems to CRAB outside CERN.
 - Fixed in 2 hours. Actually this was raised to alarm by mistake but it turned out to help!
 - ◇ castor problems (GGUS:74085): rfcv's hanging degrade castor pools. Strong impact on T0 operations and P5-T0 transfers
 - Raised to alarm. We had a meeting with CASTOR team this morning. More details from Jan
 - Need a contact or an agreed procedure if a blocking problem happens during the night or holiday/week-ends
 - Jan: we are waiting for a hotfix release that is pending; in principle we could even switch CMS back to LSF if desired
 - Claudio: we need to avoid going back to LSF, better continue debugging and fixing the Transfer Manager progressively
 - Claudio: an additional ticket was just opened (GGUS:74124)

- ◆ T1 sites:
 - ◇ 4 hours intervention on CNAF MSS today. No recalls/migration
- ◆ T2 sites:
 - ◇ NTR
- ALICE reports [↗](#) -
 - ◆ T0 site
 - ◇ Nothing to report
 - ◆ T1 sites
 - ◇ IN2P3: Problem has been understood and patch has been provided. Site is back in production
 - ◇ RAL: GGUS:74098 [↗](#). There is an issue with CREAM for ALICE. The NAGIOS CREAM tests are failing and the cause is under investigation. Experts have been contacted
 - ◆ T2 sites
 - ◇ Usual operations
- LHCb reports -
 - ◆ Experiment activities:
 - ◇ Waiting for data.
 - ◆ New GGUS (or RT) tickets:
 - ◇ T0: 0
 - ◇ T1: 0
 - ◇ T2: 0
 - ◆ Issues at the sites and services
 - ◇ SARA: We have ticket (GGUS:73244 [↗](#)) opened 1 month ago, it seems we need help from experts.
 - Maarten: will have a look and update the ticket

Sites / Services round table:

- ASGC
 - ◆ site BDII unstable since this morning, being investigated
 - ◇ Maarten: check that it uses openldap 2.4
- BNL - ntr
- CNAF - ntr
- FNAL - ntr
- GridPP - ntr
- IN2P3 - ntr
- KIT - ntr
- NDGF
 - ◆ a fiber break in Norway affects the availability of data stored there, not clear yet which data for which experiments; no time estimate yet
- NLT1 - ntr
- OSG - ntr
- RAL - ntr
- CASTOR/EOS
 - ◆ SLS monitoring issue: all SRMs marked red, but no tickets other than for real problems; under investigation
- dashboards - ntr
- databases - ntr
- GGUS - ntr
- grid services

- ◆ WN stability improving: only 19 nodes crashed overnight, which is at the usual level; exact cause of yesterday's problem still has not been found, but ALICE jobs now look **less** implicated; no changes were made on the ALICE side

AOB:

- Jan: might e-mail replies to GGUS tickets cause new tickets to be created?
- Maria D: should not happen, the subject line is parsed and a reply should get turned into a comment in the existing ticket
- NOTE: because of a CERN holiday there will be **no** meeting on Thu Sep 8, while there will be a meeting on Fri Sep 9.

Thursday

Canceled because of CERN holiday.

Friday

Attendance: local(Claudio, Guido, Ivan, Jan, Maarten, Manuel, Zbyshek);remote(Andreas, Christian, Daniele, Gareth, Gonzalo, Jhen-Wei, Jon, Lisa, Mette, Michael, Onno, Rolf, Vladimir).

Experiments round table:

- ATLAS reports -
 - ◆ T0/Central Services
 - ◇ Tier0 Oracle DB stuck for ~40 minutes on Thu 8th around 2pm; too heavy a load by jobs accessing directly COOL; required reboot
 - ◇ EOS: found a bug (command "find" on directories not readable created a deadlock in the namespace); problem fixed in the eosatlas software (to be propagated to EOS release)
 - Jan: EOS for ATLAS was updated ~10 min ago
 - ◆ T1 sites
 - ◇ INFN-CNAF LFC not accepting new registrations (Thu night): the DB reached a file allocation limit. Limit raised, problem fixed GGUS:74159 [↗](#)
 - ◇ SARA-MATRIX SRM problems (8 pool nodes crashed, GGUS:74155 [↗](#)); problem should be fixed (or at least workaround in place)
 - Onno: the 8 nodes all have the same HW, OS and configuration and are different from the rest of the storage cluster; due to the crashes the dCache head node ran into memory problems and in the end the whole cluster needed to be restarted; the vendor has suggested changing a particular timeout value; the nodes had already crashed 2 weeks ago (Thu Aug 25)
 - Gonzalo: what kind of HW?
 - Onno: Dell PowerEdge with storage from DDN
 - ◆ T2 sites
 - ◇ NTR
- CMS reports -
 - ◆ LHC / CMS detector
 - ◇ Physics production in 1380 bunch fill
 - ◆ CERN / central services
 - ◇ Castor patch installed yesterday. Since then we didn't see new stuck files. Actually many old stuck files became visible and are being cleaned by the castor people.

Unfortunately the instabilities of the past days created an unmanaged situation between the T0 machinery and the DBS and no new runs from the express stream are made available for analysis. Need a patch from the CMS side (expected later today).

- Jan: CASTOR update deployed yesterday still has issues, so you may need to open further tickets; the cleanup still continues; the servers have debugging ongoing, which may lead to occasional restarts of daemons; the next CASTOR update is expected to become available next week
- Claudio: so far no new tickets had to be opened
- Jan: there were problems observed for CMSCAF, which would not give rise to "automatic" tickets
- ◊ Exports from CERN have been failing yesterday. Probably also a consequence of the castor problems. They are ok since yesterday ~22:00
- ◆ T1 sites:
 - ◊ NTR
- ◆ T2 sites:
 - ◊ UCSD (hosting also one of the CRAB servers) unavailable because of a major blackout in California, Arizona and Mexico. No answers can be expected from them in the next hours (or days)
- ALICE reports [↗](#) -
 - ◆ T0 site
 - ◊ Nothing to report
 - Jan: an alarm ticket was opened yesterday because of a severe degradation in the performance of data replication from the experiment into CASTOR; the cause lay in the shutdown of an Xrootd redirector after it received a connection from a particular machine; there was also an issue with the alarm procedure itself: the online expert in question was not allowed to post to the relevant mailing lists; the operator alarm list membership for ALICE will be looked into
 - ◆ T1 sites
 - ◊ RAL: [GGUS:74098](#) [↗](#). Authorization problem disappeared after various Quattor components were run, but the SAM tests still fail; under investigation
 - ◆ T2 sites
 - ◊ Usual operations
- LHCb reports -
 - ◆ Experiment activities:
 - ◊ Data processing. Due to error during software distribution procedure, last tag for Conditional DB missed in Oracle DB. As result all reconstructions jobs failed yesterday. Fixed in few hours.
 - ◆ New GGUS (or RT) tickets:
 - ◊ T0: 2
 - ◊ T1: 0
 - ◊ T2: 5
 - ◆ Issues at the sites and services
 - ◊ CERN:
 - ([GGUS:74169](#) [↗](#)) few jobs failed with "Payload process could not start after 10 seconds"
 - ([GGUS:74175](#) [↗](#)) Several files that are supposed to be on Castor are not accessible
 - Jan: the matter is being investigated, but we already found that the disk server in question was not aware of the change from LSF to the Transfer Manager (internal scheduler)

Sites / Services round table:

- ASGC
 - ◆ upgrade to openldap 2.4 appears to have fixed the site BDII instability
- BNL - ntr
- CNAF - ntr
- FNAL
 - ◆ Lisa Giacchetti will replace Jon in these meetings
- IN2P3 - ntr
- KIT - ntr
- NDGF
 - ◆ Mette Lund will replace Christian in these meetings
 - ◆ Norway fiber break: the vendor said all is OK again, but we still have difficulties and therefore extended the downtime to Monday
- NLT1
 - ◆ yesterday afternoon a FroNTier Squid server stopped working due to a full partition
 - ◆ reminder: downtime on Tue
- PIC - ntr
- RAL - ntr

- CASTOR/EOS
 - ◆ ATLAS are asked to open a test alarm for EOS this afternoon to test the workflow
 - ◇ done: all OK
 - ◆ ATLAS Xrootd redirector is causing "atcal" files to be staged into the default service class; being investigated
- dashboards - ntr
- databases
 - ◆ as a side effect of rebooting the ATLAS offline DB the replication to all T1 was frozen for 2h
- grid services
 - ◆ issue with ALICE jobs running "ifconfig" still being investigated (the commands are denied and cause tons of errors to go to the central audit logs and the serial console, which slows down the worker node)

AOB:

-- JamieShiers - 25-Aug-2011

This topic: LCG > WLCGDailyMeetingsWeek110905

Topic revision: r11 - 2011-09-09 - MaartenLitmaath



Copyright &© 2008-2020 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