

Table of Contents

Week of 121015.....	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.....	4
Thursday.....	6
Friday.....	7

Week of 121015

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(Eva, Ivan, Jan, Maarten, Maria D, Ulrich, Xavier);remote(Alessandro, Federico, Gonzalo, Jhen-Wei, Lisa, Onno, Rob, Roger, Rolf, Salvatore, Stefano, Tiju).

Experiments round table:

- ATLAS reports -
 - ◆ T0/T1
 - ◇ CERN-PROD EOS crashed Saturday at 16:00, up and running again at 18:00. No ggus ticket, noticed immediately by EOS experts (and ATLAS users).
 - ◆ Alessandro: followup from Fri - a lower tape speed at KIT would **not** be a big problem for the upcoming reprocessing campaign; if some files cannot be recalled from tape on time, replicas at another T1 will be used instead
 - ◆ Jan: the CASTOR disk server that crashed on Fri is back, no files were lost
- CMS reports -
 - ◆ LHC / CMS
 - ◇ Physics running at full luminosity all OK so far
 - ◆ CERN / central services and T0
 - ◇ NTR
 - ◆ Tier-1:

- ◊ KIT: is hitting CERN's Frontier Launchpad bypassing squid. Not a problem, but why ? GGUS:87345 [↗](#)
- ◊ shifters opened GGUS due to HC failures. Failures gone now
 - CERN: GGUS:87331 [↗](#) was CE overload
 - TW: GGUS:87337 [↗](#) , GGUS:87347 [↗](#) (duplicate) seems some misconfiguration.
 - FNAL: SAV:132964 [↗](#) held jobs on new CE cmsosgce.fnal.gov . Appears understood
- ◆ Tier-2:
 - ◊ NTR
 - ◆ Stefano: there also is/was a problem with job submission to LSF at CERN
 - ◆ Ulrich: there was a replacement of the LSF master HW this morning, which caused the service to be degraded; it was announced on the IT Service Status Board, but not in the GOCDDB...
 - ◆ Stefano: the timing looks different, let's follow up offline
 - ◊ after the meeting: 2 CMS VOBOX nodes needed their LSF daemons to be restarted
 - ◆ Jan: the BeStMan SRM gateway to EOS-CMS was stuck for 2h both yesterday and today, apparently not noticed by CMS users
- ALICE reports [↗](#) -
 - ◆ CERN: old conditions data SE to be demoted again this afternoon, causing the load to be shifted back to EOS
 - ◊ this was done shortly before 15:00 CEST
- LHCb reports -
 - ◆ T0:
 - ◊ CERN : cleaning TMPDIR on lxbatch (GGUS:86039 [↗](#)) ongoing
 - ◆ T1:
 - ◊ GRIDKA: Staging efficiency not high enough for current reprocessing activities (GGUS:80794 [↗](#)), data access problems for jobs reading from tape cache (GGUS:87318 [↗](#))
 - ◊ IN2P3: "buffer" disk space not migrated fast enough to tape storage (GGUS:87293 [↗](#)), fixed, FTS transfers failing because of "expired proxy" (GGUS:87321 [↗](#)) also fixed and closed.
 - ◊ CNAF: "buffer" disk space increasing, transfer rate to tape storage increased
 - ◊ SARA: transfers failing due to authz issues (GGUS:87361 [↗](#)), seems to be solved already.

Sites / Services round table:

- ASGC
 - ◆ during the weekend there were CMS HammerCloud errors due to overloaded CREAM CEs; OK now, but looking into it
- CNAF - ntr
- FNAL - ntr
- IN2P3 - ntr
- NDGF - ntr
- NLT1 - ntr
- OSG - ntr
- PIC - ntr
- RAL - ntr

- dashboards - ntr
- databases - ntr

- GGUS/SNOW
 - ◆ File ggus-tickets.xls is up-to-date and attached to page WLCGOperationsMeetings. There were 3 real ALARMS since last MB, all from ATLAS, all for CERN. Drills ready and attached at the end of this page.
- grid services - nta
- storage - nta

AOB:

Tuesday

Attendance: local(Ivan, Maarten, Maria D, Stephen, Ulrich, Xavier E);remote(Doug, Federico, Jhen-Wei, Lisa, Rob, Roger, Rolf, Ronald, Saverio, Tiju, Xavier M).

Experiments round table:

- ATLAS reports -
 - ◆ T0/T1
 - ◇ FZK T0 export errors due to filling disk - in progress - ATLAS DE cloud support working on the problem - GGUS:87385 [↗](#)
 - Xavier M: the trouble is with full transfer queues, not disks; 1 user may be blocking a lot of transfer slots; we have banned 1 DN temporarily, but that does not deal with previously submitted transfers; we rely on the ATLAS contact for the German cloud (Rod Walker) in this matter
 - ◇ INFN Worker node errors - nodes reconfigured (solved) GGUS:87346 [↗](#)
- CMS reports -
 - ◆ LHC / CMS
 - ◇ Physics running at full luminosity all OK so far
 - ◆ CERN / central services and T0
 - ◇ CASTOR issues overnight, data was stacking up at Point5. Perhaps better now. GGUS:87382 [↗](#).
 - Xavier E: the DB was overloaded due to high activity on top of a changed execution plan; to cure the problem a lot of pending transfers were killed, disk server draining was stopped and the execution plan was corrected; there appears to be a lot of srmRm activity as well
 - Stephen: srmRm?! unexpected, we will follow up
 - ◆ Tier-1:
 - ◇ KIT: is hitting CERN's Frontier Launchpad bypassing squid. HEPiX tests. Solved. GGUS:87345 [↗](#)
 - ◆ Tier-2:
 - ◇ NTR
- ALICE reports [↗](#) -
 - ◆ CERN: EOS-ALICE so far looking OK with a significant conditions data read load
 - ◆ KIT: ALICE disk SE not working since Fri evening; experts notified
 - ◇ looks OK again since ~18:00 CEST
- LHCb reports -
 - ◆ T0:
 - ◇ CERN : cleaning TMPDIR on lxbatch (GGUS:86039 [↗](#)) ongoing
 - ◆ T1:
 - ◇ GRIDKA: Staging efficiency not high enough for current reprocessing activities (GGUS:80794 [↗](#) not urgent but opened since April), data access problems for jobs

reading from tape cache (GGUS:87318 [↗](#)), Staging failures, can be closed:
GGUS:87061 [↗](#)

Sites / Services round table:

- ASGC - ntr
- CNAF - ntr
- FNAL - ntr
- IN2P3 - ntr
- KIT
 - ◆ between 07:30 and 08:30 CEST authorization failed on the LHCb SE
 - ◆ the FTS expired proxy patch has been applied
- NDGF - ntr
- NLT1 - ntr
- OSG - ntr
- RAL
 - ◆ the CASTOR-CMS upgrade went OK this morning

- dashboards - ntr
- GGUS/SNOW
 - ◆ Experiments/sites with GGUS tickets that do not receive adequate support, please submit to MariaDZ for presentation at the WLCG Operations' meeting this Thursday.
- grid services
 - ◆ the site BDII nodes have been upgraded to EMI-2
- storage - nta

AOB:

- T0 and all T1 should apply the "expired proxy" FTS patch, as detailed in GGUS:81844 [↗](#):
 - ◆ the repository is http://etics-repository.cern.ch/repository/pm/volatile/repomd/name/FTS228_proxy_expired_patch/index
 - ◆ ONLY transfer-fts package should be picked for update and tomcat must be restarted
- whenever the FTS is found suffering from the remainder of the bug, a simple restart of tomcat then is sufficient to cure the service

Wednesday

Attendance: local(Alexey, Ivan, Maarten, Maria D, Stephen, Ulrich, Xavier);remote(Doug, Federico, Jhen-Wei, John, Lisa, Pavel, Rob, Roger, Rolf, Ronald, Salvatore, Vladimir).

Experiments round table:

- ATLAS reports -
 - ◆ T0/T1
 - ◇ PIC T0 - Failed transfers to DATATAPE due to stuck tape in drive GGUS:87485 [↗](#)
 - ◇ Triumph: stage in failures GGUS:87459 [↗](#)
 - ◇ ND-ARC - ND.ARC: jobs failing with "Transformation not installed in CE" GGUS:87378 [↗](#)
 - ◆ **ATLAS express stream reprocessing has started this week**
 - ◆ Xavier: yesterday a GGUS ticket was opened about slow recalls from tape affecting T0 operations; that was due to a misconfiguration, fixed now

- CMS reports -
 - ◆ LHC / CMS

- ◊ Some physics running
- ◆ CERN / central services and T0
 - ◊ Calibration issues due to way we recovered from CASTOR issues yesterday (CMS internal problem).
- ◆ Tier-1:
 - ◊ NTR
- ◆ Tier-2:
 - ◊ NTR
- ALICE reports [↗](#) -
 - ◆ CERN: EOS-ALICE head node crashed yesterday ~16:00, back in ~1h
 - ◆ KIT: ALICE disk SE OK again since yesterday 18:00; looking into monitoring improvements (working with ALICE MonALISA/Xrootd manager and with advice from EOS team) that may also benefit other Xrootd installations for ALICE
- LHCb reports -
 - ◆ T0:
 - ◊ CERN: pilots aborted (GGUS:87447 [↗](#)) and redundant pilots, asking to be deleted (GGUS:87848 [↗](#))
 - ◆ T1:
 - ◊ GRIDKA: Staging efficiency not high enough for current reprocessing activities (GGUS:80794 [↗](#) not urgent but opened since April), data access problems for jobs reading from tape cache (GGUS:87318 [↗](#)). Reprocessing going slow due to that.
 - ◊ SARA: running out of Tape. Banned for writing GGUS:87486 [↗](#)
 - Ronald: more tapes have been added, but LHCb are already using more than what was pledged...
 - Federico: thanks! we are discussing what to do about the usage

Sites / Services round table:

- ASGC - ntr
- CNAF - ntr
- FNAL - ntr
- IN2P3 - ntr
- KIT - ntr
- NDGF - ntr
- NLT1 - nta
- OSG - ntr
- RAL
 - ◆ 1 ATLAS disk server has been unavailable, but is ready to go back in
- dashboards - ntr
- GGUS/SNOW
 - ◆ We need a new CMS technical contact for the Savannah-GGUS bridge as the interface experts from the experiment side moved on to other things. Details in Savannah:131565 [↗](#) (Answer is: Oliver Gutsche will be our contact for this from now on).
 - ◆ GGUS:85197 [↗](#) is waiting for OSG input on site name. (Rob updated the ticket after the meeting).
 - ◆ **NB!!** GGUS Release next Wednesday 2012/10/24 with the usual test tickets (ALARMS and attachments). This is published on the GGUS homepage [↗](#). Details in https://ggus.eu/pages/news_detail.php?ID=471 [↗](#)
- grid services
 - ◆ FTS: the "expired proxy" patch has been applied
 - ◆ BDII: transparent upgrades to EMI-2 ongoing

- storage - nta

AOB:

Thursday

Attendance: local(Alexey, Felix, Ivan, Kate, Maarten, Stephen, Ulrich, Xavier);remote(Doug, Federico, Gonzalo, John, Kyle, Lisa, Roger, Rolf, Ronald, Saverio, WooJin).

Experiments round table:

- ATLAS reports -
 - ◆ T0/T1
 - ◇ FZK - T0 transfer to DATATAPE failures - GGUS:87510 [↗](#)
 - ◇ FZK - Failures of staging of files from DATATAPE and MCTAPE - GGUS:87526 [↗](#)
 - ◇ IN2P3-CC - Failure of staging of files from MCTAPE - Site reports it fixed - GGUS:87529 [↗](#)
 - ◇ SARA-MATRIX - Failure of staging of files from DATATAPE - GGUS:87531 [↗](#)
 - ◇ CERN PROD EOS - missing files GGUS:87530 [↗](#)
 - Xavier: those files got written a first time, then truncated on a second write which failed; they will have to be uploaded yet again
 - ◇ Xavier: there was a CASTOR ticket about stuck transfers from ATLCAL to T0ATLAS; we killed them to unblock the situation and are investigating the cause
 - ◆ **ATLAS express stream reprocessing still continuing**
- CMS reports -
 - ◆ LHC / CMS
 - ◇ Some physics running
 - ◆ CERN / central services and T0
 - ◇ EOS namespace crashed yesterday evening. Recovered quickly. INC:179247 [↗](#)
 - Xavier: also the BeStMan SRM was stuck a few times today, looking into it
 - Stephen: we do not understand the srmRm activity you reported on Tue, please provide examples
 - Xavier: OK
 - ◆ Tier-1:
 - ◇ NTR
 - ◆ Tier-2:
 - ◇ NTR
- ALICE reports [↗](#) -
 - ◆ NTR
- LHCb reports -
 - ◆ T0:
 - ◇ CERN: pilots aborted (GGUS:87447 [↗](#)) and redundant pilots, asking to be deleted (GGUS:87448 [↗](#)). Still discussions going on this, not yet understood.
 - ◆ T1:
 - ◇ GRIDKA: Staging efficiency not high enough for current reprocessing activities (GGUS:80794 [↗](#) not urgent but opened since April), data access problems for jobs reading from tape cache (GGUS:87318 [↗](#)). Reprocessing going slow due to that.

Sites / Services round table:

- ASGC - ntr

- CNAF
 - ◆ yesterday between 16:00 and 17:00 CEST all transfers and data accesses by jobs failed due to a network issue
- FNAL
 - ◆ secondary FTS instance upgraded with "expired proxy" patch, primary instance will be done next week
- IN2P3 - ntr
- KIT
 - ◆ looking into the tape performance issues
 - ◆ next Mon Oct 22 starting at 12:00 UTC the CEs cream-{1,2,3}-fzk.gridka.de will enter downtime for retirement; new CEs cream-{6,7,8}-kit.gridka.de (sic) are already available
- NDGF - ntr
- NLT1 - ntr
- OSG - ntr
- PIC
 - ◆ an issue affecting ATLAS tape reads was fixed by a firmware upgrade yesterday evening
 - ◆ Stephen: you may soon get a CMS ticket about tape backlogs, possibly related?
 - ◆ Gonzalo: will look into it
- RAL - ntr

- dashboards - ntr
- databases
 - ◆ 1 ATLARC DB node needed to be rebooted to cure a shared memory problem, looking into it
 - ◆ the CASTOR-LHCb stager DB needs to be migrated to new HW, requiring 3h downtime; we propose Nov 22, i.e. during the LHC machine development foreseen for that week
 - ◇ Federico: OK (the reprocessing campaign will be made to deal with it)
- grid services - ntr
- storage - nta

AOB:

- SIR report for CASTORCMS incident on Monday 15 SIR

Friday

Attendance: local(Alexey, Eva, Felix, Ivan, Maarten, Stephen, Ulrich, Xavier E);remote(Dimitrios, Doug, Federico, Gareth, Kyle, Onno, Rolf, Salvatore, Xavier M).

Experiments round table:

- ATLAS reports -
 - ◆ T0/T1
 - ◇ Nothing new to report
 - ◆ **ATLAS express stream reprocessing mostly complete**
 - ◇ Handful of jobs to complete in Most clouds
 - ◇ FZK has 20% of jobs still to complete

- CMS reports -
 - ◆ LHC / CMS
 - ◇ Some physics running since 11am
 - ◆ CERN / central services and T0
 - ◇ Transfers to EOS from CASTOR problems since ~noon yesterday. INC:179785 [↗](#).
 - Xavier E: looking into it
 - Maarten: was the srmRm traffic understood in the meantime?

· Stephen: not yet

- ◆ Tier-1:
 - ◇ A transfer to PIC was suspended, not sure why. Now enabled. SAV:133094
 - ◇ KIT has stopped production and custodial imports due to tape system issues.
- ◆ Tier-2:
 - ◇ NTR
- ALICE reports -
 - ◆ One lcg-voms.cern.ch machine got its host cert renewed with the wrong subject, thereby causing VOMS client authentication errors: GGUS:87589
 - ◇ Ulrich: this was fixed shortly before 15:00 CEST, a new lcg-voms certificate has been deployed on all lcg-voms nodes
 - ◇ Maarten: this change *should* be transparent, but there may still be a few legacy services depending on having a copy of the host cert, as used to be distributed via the lcg-vomscerts rpm; I will make a new version available in the ETICS lcg-vomscerts repository
- LHCb reports -
 - ◆ Nothing new to report

Sites / Services round table:

- ASGC - ntr
- CNAF - ntr
- IN2P3 - ntr
- KIT
 - ◆ today between 02:30 and 04:30 CEST authorization failed on the LHCb SE
 - ◆ Mon Oct 22 between 05:00 and 07:30 CEST there will be frequent network interruptions
- NLT1 - ntr
- OSG - ntr
- RAL
 - ◆ interventions on Tue Oct 23:
 - ◇ CASTOR-LHCb upgrade
 - ◇ replacement of old gLite CREAM CEs with new EMI CREAM CEs; experiment contacts will be notified
- dashboards - ntr
- databases - ntr
- grid services - nta
- storage
 - ◆ EOS-ALICE head node has been rebooted as agreed, took ~15 min; will be upgraded on Mon Oct 22 at 14:00 CEST
 - ◆ Stephen: EOS-CMS upgrade plan for next week?
 - ◆ Xavier E: let's discuss that offline

AOB:

-- JamieShiers - 18-Sep-2012

This topic: LCG > WLCGDailyMeetingsWeek121015

Topic revision: r16 - 2012-10-19 - MaartenLitmaath



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