

# Table of Contents

<b>Week of 101115.....</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:.....	4
Wednesday.....	6
Thursday.....	9
Friday.....	11

# Week of 101115

## 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 (AndreaV, Maarten, Hurng, Oliver, Eddie, Roberto, Harry, Jan, Luca, Lola, AleDG, Massimo, Flavia, MariaDZ); remote(Michael, Jon, Xavier, Tiju, Tore, Riccardo, Onno, Rob, Rolf, Gang).

Experiments round table:

- ATLAS reports -
  - ◆ beam ramped up to 121 bunches, 113 beams collided at ATLAS. High trigger rate resulting in high disk writing rate from SFO. ~650 MB/sec for Sunday runs.
  - ◆ concerning the data production rate, RAW data export from T0 to T1s was stopped in Saturday afternoon.
  - ◆ p-p reprocessing on data taken in October was postponed to this week.
  - ◆ T0:
    - ◇ no issue during the weekend.
    - ◇ this morning we just noticed the service availability of T0MERGE Castor pool have been staying at 89% since 10th Nov (Wed.) GGUS:64249. It's due to the monitoring probe was stuck.
  - ◆ T1s:
    - ◇ BNL - disk space closes to full, "writing" activities were scheduled on a reduced number of disk servers with sufficient space for incoming data. BNL added storage and started data deletion

- ◇ RAL - transfer failures. being notified with an alarm (GGUS:64228), load being released by reducing number of concurrent transfers at FTS@RAL and number of production jobs at RAL. Site investigated today and claimed this is due to the issue mentioned here: BUG:65664. [Tiju: large number of FTS requests seen, will add 2 servers to pool.]
- ◇ SARA
  - started to drain the queue in Sunday evening for downtime on Tuesday. Reprocessing jobs at SARA went down to 0. Special request to re-enable the atlas jobs to continue until downtime.
  - power failure in one storage rack on Monday morning causing a short period of job failure and transfer failure. [Onno: power failure was observed when opened the queues for ATLAS, may be due to problem with power supply in one rack and/or increased load on storage system and network.]
- CMS reports -
  - ◆ Experiment activity
    - ◇ HI data taking: a lot of stable beam over the weekend, very good performance of machine and CMS
  - ◆ CERN and Tier0
    - ◇ Input rate from P5 (measured by looking into input rate into T0Streamer pool) is 1.4-1.5 GB/s
    - ◇ load on T0Temp peaked at 1.8 GB/s in and out
    - ◇ load on T0Export is about 5 GB/s out when farm is filled with prompt reco jobs (highest expected load). Don't expect to go significantly higher. [Jan: appreciated the email exchange from CMS about high data rates, will keep in contact.]
    - ◇ informed castor operations lists once over the weekend
    - ◇ GGUS ticket to LSF: GGUS:64229
      - saw only 2.3k of 2.8k jobs running
      - think this is due to the disk requirement for some of the jobs (repack) of 50 GB
      - maybe we can clarify that this is actually the case
  - ◆ Tier1 issues and plans
    - ◇ tails of re-processing pp data + skimming at all Tier-1s
    - ◇ tails of PileUp re-digi/re-reco
    - ◇ expect MC production and WMAgent scale testing to kick in again
    - ◇ T1\_DE\_KIT: implementation of new HI production roles: GGUS:64069 -> in progress, last update 11/12 early morning. [Xavier: increased max number of transfers, maybe it will help.]
    - ◇ T1\_FR\_CCIN2P3: transfer problems to MIT, solved? GGUS:63826 -> last reply to ticket from 8th Nov: "Concerning transfers exporting from IN2P3 to other sites ( like MIT who opened this ticket), We still see the same errors : AsyncWait, Pinning failed", contacted local site admins, they hope the changes they did for Atlas will solve the problems CMS is seeing (file access problems, staging issues, etc.) -> closed today, "The import has required some more time that the expected one. This is due to all the issues that we had facing since the downtime on 22th sept. However our dcache master is working hardly to improve the situation, thing that is illustrating from the results of our activities."
  - ◆ Tier2 Issues
    - ◇ T2\_RU\_RRC\_KI : CMS SAM-CE instabilities since several days (Savannah:117543), site admins are following up and replied yesterday, see GGUS:63820
    - ◇ T2\_FI\_HIP: Wrong gstat and gridmap information for the Finnish CMS Tier-2: GGUS:63956, last update Nov. 8th, further discussed in T1 coordination meeting, Flavia following up as well. More explanation:

- T2 FIN resources appear in BDII as part of NDGF-T1 because we need to translate the ARC information system to Glueschema and into BDII format and this is done by NDGF
- Site has an independent CMS dCache setup which is not part of the distributed NDGF T1 dCache setup so that cannot be published by NDGF, but is instead published by a CSC BDII server.
- CMS applications and WLCG accounting works fine with this setup, but the WCLG monitoring is reporting non CMS resources for the T2 Fin site which obviously is wrong.
- The ticket is about having WLCG report the right resources for us. The short term solution is to for WLCG to change configuration settings to get this corrected.
- ◇ T2\_RU\_JINR: SAM CE test fails analysis check, file not accessible: Savannah:117804 [↗](#)
- ◇ T2\_TR\_METU: SAM CE tests prod and sft-job are failing: Savannah:117824 [↗](#)
- ◇ T2\_FR\_GRIF\_IRFU: SAM test errors: CE, ... : Savannah:117836 [↗](#), admins responded, no contact to services from WAN, experts working on it
- ◇ T2\_RU\_IJEP: SRM SAM test failure since 6 hours: Savannah:117838 [↗](#)
- ALICE reports [↗](#) -
  - ◆ GENERAL INFORMATION: none.
  - ◆ T0 site
    - ◇ During the weekend there were problems with voalice09. We opened a GGUS ticket this morning (GGUS:64244 [↗](#)) because it was not possible to add a file to the SE. The problem was not on the CASTOR side, but on the ALICE, so ticket was marked as unsolved. There were several problems:
      - the file descriptors were reset to the default 1024, which lead to 'too many open files' on voalice09.Limit was increased, the redirector is working now
      - A config file missing where the firewall configuration was specified was missing
      - Still errors, under investigation
    - ◇ voalice09 has to be replaced by another vobox. This operation will take place this week as soon as the new machine is configured and ready to substitute it.
  - ◆ T1 sites
    - ◇ Nothing to report
  - ◆ T2 sites
    - ◇ Usual operations
- LHCb reports -
  - ◆ Experiment activities: Validation of the reprocessing in progress..
  - ◆ New GGUS (or RT) tickets: 1 at T0 (none at T1, T2)
  - ◆ Issues at the sites and services
    - ◇ T0:
      - Debugging the issue with accessing data on CASTOR. Some fix have been applied on Friday and seems to improve the situation. So we confirm or not during the coming days that it fix one part of the problem. [A meeting between LHCb and CASTOR was held on Monday morning to review the status, things look ok for the moment.]
      - Open another ticket (uncorrelated with xrootd issue) to track down some timeouts accessing files (GGUS:64258 [↗](#)).
    - ◇ T1 site issues: none.

Sites / Services round table:

- BNL: ntr
- FNAL: ntr
- KIT: nta
- RAL: downtime starting tomorrow for the upgrade of the CMS Castor instance to 2.1.9 (three days)
- NDGF: ntr
- INFN: ntr
- NL-T1: nta
- OSG:
  - ◆ LDAP timeout increase seems to have worked, no new failure was observed.
  - ◆ Continuing investigation of CERN-Indiana network issues, there will be a maintenance tomorrow. [Flavia: following up the network tests from Australia to Indiana].
  - ◆ Following up the upgrade of BDII to glite 3.2.
  - ◆ Observed that CERN is not using round-robin, will be followed up. [Maarten: may be due to RHEL5/SLC5 issue in glibc].
- ASGC: ntr
- IN2P3: ntr
- **CERN AFS UI:** The **current** gLite 3.2 AFS UI will be updated on Tuesday 16th November at 09:00 UTC. This is one week later than was previously advertised.

	previous 3.2	current 3.2	new 3.2
Now	3.2.1-0	3.2.6-0	3.2.8-0
After	3.2.6-0	3.2.8-0	3.2.8-0

The **3.2.8-0** has been in place now for 3 weeks for testing and is in use by CMS during this time. [Nico - some minor problem which will be reported] The very old **3.2.1-0** release will be archived and subsequently deleted in month or twos time.

AOB: none

## Tuesday:

Attendance: local(AndreaV, Maarten, Jan, Jamie, Elena, Eddie, Simone, Harry, Roberto, Lola, MariaDZ, Flavia, AleDG);remote(Micheal, Jon, Gonzalo, Jeff, Tore, Kyle, Gang, Rolf, John, Jeremy, Riccardo, Ian, Joel).

Experiments round table:

- ATLAS reports -
  - ◆ 121b x 121b heavy ion runs with total 644.6mb-1
  - ◆ GGUS was unavailable this morning.
  - ◆ T0:
    - ◇ no problem to report
  - ◆ T1s:
    - ◇ BNL is in process of moving data from overcommitted storage pools to those which have still sufficient space to accommodate new data. We see some timeouts errors in data transfers and production jobs. Thanks for notifying ATLAS.
    - ◇ RAL: problem with MCTAPE. Disk server has been disabled temporarily and some files haven't been migrated to tape.Under investigations. No GGUS ticket (as GGUS was down) but RAL promptly reacted to email. Thanks.
  - ◆ T2s:
    - ◇ CA-ALBERTA-WESTGRID-T2: Problem with file transfers(GGUS:64293 [↗](#)).
- CMS reports -
  - ◆ Experiment activity

Tuesday:

- ◊ HI data taking: very good performance of machine and CMS. Live time of the accelerator is better than expectations.
- ◆ CERN and Tier0
  - ◊ Input rate from P5 (measured by looking into input rate into T0Streamer pool) is 1.4-1.5 GB/s
  - ◊ load on T0Temp peaked at 1.8 GB/s in and out
  - ◊ load on T0Export is about 5 GB/s out when farm is filled with prompt reco jobs (highest expected load). Don't expect to go significantly higher
  - ◊ With the high machine livetime we may develop a backlog of prompt reco. Recover during the shutdown. Rate out of castor should not exceed 5GB/s, but it could be sustained.
- ◆ Tier1 issues and plans
  - ◊ tails of re-processing pp data + skimming at all Tier-1s
  - ◊ tails of PileUp re-digi/re-reco
  - ◊ expect MC production and WMAgent scale testing to kick in again
  - ◊ Transfer problems IN2P3 to 4 Tier-2s. Savannah 117864 [Simone: is this a new problem? Similar issues are seen for ATLAS. Ian: will double check and report tomorrow.]
  - ◊ [Eddie: 3 SRM tests failing at KIT. Ian: could be an indication that we are overloading the system for skimming.]
- ◆ Tier-2 Issues
  - ◊ Analysis Jobs per day at Tier-2s is increasing. Exceeding 200k jobs per day
- ALICE reports [↗](#) -
  - ◆ GENERAL INFORMATION: none.
  - ◆ T0 site
    - ◊ Problems reported yesterday with voalice09 have been solved this morning. There was some misconfiguration in the quattor machine. Thanks to CASTOR team for helping us in following that problem
    - ◊ voalice09: as moving from CRA ("plus") to LDAP ("cern") for SLC4 quattor-managed machines, there are some changes to made in voalice09 or deprecate it now and put voalice16 to substitute it [Jamie: can we postpone the replacement of voalice09 by two weeks until the end of the HI run? Lola (after some discussion): ok, it's going to be postponed]
    - ◊ voalice16 has been from the quattor point of view configured as xrootd redirector. Not in production yet
  - ◆ T1 sites
    - ◊ Nothing to report
  - ◆ T2 sites
    - ◊ Usual operations
- LHCb reports -
  - ◆ Experiment activities: Validation of the reprocessing in progress..
  - ◆ New GGUS (or RT) tickets: 1 at T1 (none at T0, T2)
  - ◆ Issues at the sites and services
    - ◊ T0: ntr [Joel: no complaint received after the changes for CASTOR at CERN, problem seems to have been fixed.]
    - ◊ T1 site issues:
      - RAL : remove file failed for USER data (GGUS:64265 [↗](#)). [John: issue should be fixed now, the user can retry.]
      - SARA : downtime BUT the streaming for CONDDDB was still active.
      - [Joel: starting to worry about IN2P3, may need to ban it from the reprocessing this week. Rolf: still work in progress (busy with tests; local LHCb contact person is involved; AFS issue being followed up, one server

will be added; will report to T1SCM with more details). Cannot say yet whether LHCb will be able to use IN2P3 for reprocessing this week.]

Sites / Services round table:

- BNL: consolidating disk space (with reference to the ATLAS report), should be transparent
- FNAL: power outage next Thursday from 5am to 8am Chicago time, will affect CMS
- PIC: ntr
- NL-T1:
  - ◆ Maintenance finished, now working on some software upgrades.
  - ◆ Observed failures of LCG replica manager tools for ATLAS at Nikhef (GGUS: 61349). [Elena: knew about this, so this had no effect.]
  - ◆ Announcement: on November 30 the worker nodes at Nikhef will be moved behind UPS, will need to bring them down and drain the queue.
- NDGF:
  - ◆ tomorrow downtime for dcache, officially from 1 to 2pm UTC (will be shorter in reality)
  - ◆ maintenance on Thursday affecting ATLAS and ALICE
- OSG:
  - ◆ MTU change will not be done as it was proved not to be the cause of the BDII issue
  - ◆ discovered asymmetric routing, will follow up
- ASGC: ntr
- IN2P3: nta
- RAL: downtime CASTOR upgrade for CMS (as announced), seems to be going ok
- GridPP: ntr
- CNAF: ntr
- **CERN PROD AFS gLite 3.2 UI updated.** As advertised the AFS UI was updated today at 11:00 UTC. [Maarten: new AFS version has been tested by CMS for one week, all ok, only one minor fix was needed.]

	previous 3.2	current 3.2	new 3.2
Before Today	3.2.1-0	3.2.6-0	3.2.8-0
Today	3.2.6-0	3.2.8-0	3.2.8-0

AOB: (MariaDZ)

- Serious problems, since ~08:15 CET this morning the GGUS web service plugin was down. Restored together with the mail parser around 11am. [!MariaDZ: SIR will be presented at T1SCM.]
- Answer by GGUS developer G.Grein on the CMS savannah ticket that generated multiple GGUS ones last week: "The GGUS bridge worked correctly. But as there was no site specified in the savannah ticket, the GGUS ticket was assigned to TPM. In the ticket description 4 hosts were mentioned on which problems occurred. Hence the TPM splits the original ticket into 4 tickets, one per host. This is the reason for this large number of tickets."

## Wednesday

Attendance: local (AndreaV, Elena, Simone, Maarten, Eddie, AleDG, Lola, Edoardo, Harry, Flavia, JanI, Ulrich, Dirk, Jamie, MariaDZ); remote (Michael, Jon, Rob, John, Tore, Rolf, Onno, Xavier, Dimitri, Suijian, Joel, IanF).

Experiments round table:

- ATLAS reports -
  - ◆ No Physics Runs until Friday
  - ◆ T0:

- ◊ no problem to report
- ◊ [Simone: new storage endpoint for EOS has been prepared and published in BDII. FTS for this new channel needs to be configured from all Tier1 sites and the FTS server at CERN: some sites will need to run scripts manually, elsewhere it will be done automatically. Can all T1 sites please follow this up.]
- ◆ T1s:
  - ◊ ATLAS has successfully finished Part I of Autumn re-processing campaign and is starting re-processing of October data.
  - ◊ No major problems are seen at T1's.
  - ◊ [Simone: asked 20 days ago that Taiwan T2 should be considered as a T1 for FTS. CERN changed the FTS config accordingly. Can all T1 sites follow this up and make sure that the FTS config is updated accordingly.]
- ◆ T2s:
  - ◊ No major problem
- CMS reports -
  - ◆ Experiment activity
    - ◊ Machine moving to proton-proton for 3 days. CMS does not expect to send data from P5 to T0 during this time
  - ◆ CERN and Tier0
    - ◊ After last night's run we had a lot of prompt-reco in the system and exceed 5GB/s for a short time. Mail was sent to Castor Operations as per agreement and has since dropped.
  - ◆ Tier1 issues and plans
    - ◊ tails of re-processing pp data + skimming at all Tier-1s
    - ◊ tails of PileUp re-digi/re-reco
    - ◊ expect MC production and WMAgent scale testing to kick in again
    - ◊ Transfer problems IN2P3 to 4 Tier-2s. (Savannah:117864) [IanF: this ticket is now closed.]
    - ◊ [IanF: found one more old issue with RAL, maybe will resubmit it as GGUS ticket if necessary.]
  - ◆ Tier-2 Issues
    - ◊ Analysis Jobs per day at Tier-2s is increasing. Exceeding 200k jobs per day
- ALICE reports -
  - ◆ GENERAL INFORMATION: LHC is in technical stop presently, we are recuperating (Pass1 prompt reco) the backlog of runs accumulated on the weekend/Monday. In addition, setting up a large scale MC production
  - ◆ T0 site
    - ◊ The moving from CRA ("plus") to LDAP ("cern") for SLC4 quattor-managed machines, is not going to be applied to voalice08, voalice09 and voalicefs01-05 until two weeks after the ending of the HI run [Ulrich: actually the migration took place yesterday (by accident, sorry about this), Latchezar has checked that all is ok.]
    - ◊ voalice16 is ready as a future substitute for the xrootd redirector. The migration will be scheduled for after the HI run
  - ◆ T1 sites
    - ◊ SARA/NIKHEF: Access from outside to the SE is slow, but the SE is operating fine. We are working with the site experts to identify the cause and remedy it (if possible)
  - ◆ T2 sites
    - ◊ Usual operations
- LHCb reports -
  - ◆ Experiment activities: Validation of the reprocessing in progress..
  - ◆ New GGUS (or RT) tickets: 1 at T1 (none at T0, T2)

- ◆ Issues at the sites and services
  - ◇ T0: ntr
  - ◇ T1 site issues:
    - SARA: data unavailable (GGUS:64348 [↗](#))
    - [Joel: also found out a problem with checksums at RAL. May open a ticket if necessary. The local LHCb contact at RAL (Raja) has been involved.]

Sites / Services round table:

- BNL: FTS config will be automatically updated
- FNAL: reminder, there will be a power outage tomorrow
- OSG:
  - ◆ understood and fixed the problem with the BDII, the issue was in the campus network
  - ◆ having fixed BDII, later today will ask BNL to publish full set of data again instead of a subset
  - ◆ [Flavia: sent to Rob at OSG a traceroute from Australia, which confirms that the same network is used. So the fix applied today should fix the problems from Australia too.]
- RAL:
  - ◆ CASTOR upgrade is going well
  - ◆ Will investigate what needs to be done about FTS config
- NDGF: ntr
- IN2P3:
  - ◆ Problems reported by LHCb are still under investigation, there may be several causes; some changes to the LHCb setup (the same workaround that had been applied for ATLAS) had been reported to the T1SCM last Thursday, not clear however if LHCb has accepted them or not [Joel: was not aware of the T1SCM report, will follow up]
  - ◆ Slow transfers and dcache (GGUS:64202 [↗](#), GGUS:64190 [↗](#), GGUS:63631 [↗](#), GGUS:63627 [↗](#), GGUS:62907 [↗](#)): will deploy additional servers to supply sufficient GridFTP capacity, available before the end of the week. Will also deploy import/export servers dedicated to ATLAS to separate high activity from those of the other LHC VOs, planned for beginning of next week.
  - ◆ FTS for Taiwan was set up already on November 5th
- NL-T1: LHCb problem (GGUS:64348 [↗](#)) due to a pool node that was not properly restarted, a leftover from yesterday's intervention. Will restart the node to fix the issue and will also investigate why the monitoring did not identify this immediately.
- PIC: the FTS channel to Taiwan for ATLAS is ready.
- KIT: two new roles have been created for CMS.
- ASGC: ntr
  
- Grid services (Ulrich):
  - ◆ Ongoing upgrade on WMS, should be transparent
- Frontier services (Flavia):
  - ◆ Observed high load on one of the two Frontier nodes for ATLAS, one that is presently used from FZK. The traffic is mostly coming from CERN however [AleDG: this is probably due to the HammerCloud tests ongoing from CERN]
- Storage services (JanI):
  - ◆ EOS for ATLAS is moving into extended testing
  - ◆ a CMS disk server on T0export is dead, it contains 3 HI files
  
- AOB (MariaDZ): GGUS development requests by the experiments were discussed in a dedicated meeting today. Details will be posted in Savannah and discussed at the next T1SCM. In summary:
  - ◆ TEAM and ALARM tickets now having only the default 'Type of problem' value "other" will be equipped with the full list of values (subject to review) like USER tickets.
  - ◆ The suggestion to create TEAM instead of USER tickets via the CMS-to-GGUS bridge was

rejected.

- ◆ The request to be able to convert TEAM tickets into ALARMS was accepted.

## Thursday

Attendance: local (Andrea, Elena, Eddie, IanF, Simone, Stephane, Jacek, MariaG, Harry, Ulrich, Miguel, AleDG, Steve); remote (Michael, Jon, Tore, Xavier, Gareth, Jeff, Foued, Rolf, Luca, Kyle, Jeremy, Suijian, Joel).

Experiments round table:

- ATLAS reports -
  - ◆ No Physics Runs until Friday
  - ◆ T0:
    - ◇ FTS channel CERN-SARA stucked (GGUS:64352 [↗](#)). According to channel definition it should have 45 jobs. Channel queue was full (10/10) and CERN increased the number of jobs from 10 to 11 yesterday. [Steve/Simone: the problem was that the maximum number of files per VO was higher than the overall maximum number of files in that channel; the latter will be increased.]
  - ◆ T1s:
    - ◇ INFN: storm was down yesterday, the problem was noticed and fixed by the site in 1.5h. Thanks. Today: srm is down (GGUS:64381 [↗](#)). Restarted.
    - ◇ BNL: problem with dCache's name space service was noticed and fixed by the site. Thanks.
    - ◇ LYON: Since there is data reprocessing which has high priority and Lyon is not able to explain the reported 'High load on dcache from ATLAS', no activity except 'data reprocessing' will be restarted for the moment even if Lyon deploys more Gridftp hardware.
  - ◆ T2s:
    - ◇ No major problem
- CMS reports -
  - ◆ Experiment activity
    - ◇ Continued proton-proton work. CMS does not expect to send data from P5 to T0 during this time
  - ◆ CERN and Tier0
    - ◇ Quiet at CERN.
  - ◆ Tier1 issues and plans
    - ◇ [IanF: two of the seven CMS T1 sites (RAL and FNAL) were down due to scheduled maintenance]
    - ◇ re-processing pp data is wrapped up. Postmortem in progress.
    - ◇ skimming is completing
    - ◇ tails of PileUp re-digi/re-reco
    - ◇ expect MC production and WMAgent scale testing to kick in again
    - ◇ 3 RAW Data files at CNAF are failing to be read (sr #117902: 3 corrupt MuMonitor files? [↗](#)). Checksum of these 3 files do not match the bookkeeping system, looks like real corruption. Retransferred from CERN. Files were corrupted before the checksum validation after transfer was deployed at CNAF. [Luca: transferred before or after June? IanF: before (in May), hence corruption was not found by the checksum validation. It might help if CNAF could revalidate all data transferred before June. Luca: will be difficult because this is over 1PB. IanF: ok, then any corrupted files will be found as we go along.]
  - ◆ Tier-2 Issues
    - ◇ No big Tier-2 Issues

- ALICE reports [↗](#) -
  - ◆ GENERAL INFORMATION: none.
  - ◆ T0 site
    - ◇ Nothing to report
  - ◆ T1 sites
    - ◇ Nothing to report
  - ◆ T2 sites
    - ◇ Usual operations
  
- LHCb reports -
  - ◆ Experiment activities: Validation of the reprocessing in progress..
  - ◆ New GGUS (or RT) tickets: 1 at T1 (none at T0, T2)
  - ◆ Issues at the sites and services
    - ◇ T0: ntr
      - [Joel: found a new xrootd problem, will open a new ticket.]
    - ◇ T1 site issues:
      - RAL : checksum is available only for new file. So If we check the "old" files, the checksum is 0. [Luca: is LHCb using checksums at all sites? It seems that at CNAF the checksum is not used. Joel: will check for non-CASTOR sites, at RAL checksums were disabled because of an issue with CASTOR. At CERN Massimo runs a tool to identify files that may be problematic; if you suspect problems with CNAF files, please send a few examples so that they can be checked. Miguel: at CERN, the problem is related to aborted transfers. Joel: this is not specific to CASTOR.]
      - CERN : webservice serving the LHCb distribution application was not available.
      - [Joel: at SARA, no new problems have been observed. It was actually all pools that were restarted after the maintenance.]

Sites / Services round table:

- BNL: ntr
- FNAL:
  - ◆ switch failure yesterday afternoon, was fixed during the night
  - ◆ update on power outage: brining services back up
- NDGF: ntr
- PIC: ntr
- RAL:
  - ◆ CASTOR upgrade for CMS was completed this morning
  - ◆ one diskserver was out for LHCb this morning, will be back in the afternoon
  - ◆ question about EOS and FTS: should we not have heard this through other channels? [Simone: it was just a reminder, the BDII was updated and everything should have happened transparently. In any case, realized that yesterday's request was wrong: since the end point was at CERN, there was no need to do anything at T1's.]
  - ◆ same question about the Taiwan issue: was the T1SCM the right place to discuss this? [Simone: the request stands, and it was correct to raise it at the T1SCM. Will review the T1SCM minutes to make sure that this was correctly documented there. Maria: yes, discussing this at the T1SCM was the correct procedure.]
  - ◆ the changes requested for Taiwan have not been made yet
- NL-T1: ntr
- KIT: ntr
- IN2P3: ntr
- CNAF: problems with storm for ATLAS this morning, no requests were processed (did not trigger any alarms unfortunately, service was restarted when realized); next week storm will be upgraded,

this should fix the bug.

- OSG: network seems fixed, but waiting for a full report by the network engineer before increasing BNL load; when this is fully fixed, will also check again the CERN timeout problem again
- GridPP: ntr
- ASGC: ntr
  
- Storage services (Miguel): will send an announcement today about a CASTOR upgrade next week on Wednesday, affecting ATLAS and CMS

AOB (MariaG): the next T1SCM will be in two weeks, to realign it with the Architects Forum. In the next T1SCM, Roger Bailey will present the latest plan about the LHC shutdown and 2011 startup.

## Friday

Attendance: local (AndreaV, Elena, Jacek, Nilo, Ignacio, Harry, MariaG, Eddie, Roberto, Ulrich, AndreaS, Lola, IanF, Simone); remote(Ulf, Michael, Rolf, Xavier, Xavier, Suijian, Onno, Tiju, Kyle).

Experiments round table:

- ATLAS reports -
  - ◆ HI runs starts tomorrow
  - ◆ T0:
    - ◇ Oracle problem with PANDA DB. Solved, thanks to PhyDB team. [Jacek: all is working well now. Problem was due to contention caused by massive data deletion launched yesterday afternoon. As a workaround, PANDA was moved to a different node of the RAC cluster yesterday evening. The problem was back this morning, but then all went back to ok. An additional problem was caused yesterday by the rollback resulting from the killing of the deletion process. MariaG: are you in contact with the PANDA team in ATLAS? Jacek: yes, in contact with several people from ATLAS. It was actually the ATLAS DBAs that started the massive deletion to fix a problem with monitoring. Nilo: if you need to delete many more rows than those you want to keep, it is faster to drop the table and recreate it.]
  - ◆ T1s:
    - ◇ SARA: srm was down this morning due to s full file system, GGUS:64407. Fixed in 1.5. Thanks. [Onno: the filesystem is monitored, so Nagios sent an alert this morning. The problem happened during log rotation. Ususally dcache logs are written to a large partition, but after the dcache upgrade the logs were written to the small root partition again. Now the issue has been fixed and the logs are written to the correct partition.]
    - ◇ FZK: failed srm-bring-online requests, GGUS:64365, on hold, linked to GGUS:61083, on hold from August. [Simone: does 'on hold' mean that it is waiting for ATLAS or for the site? Elena: problem comes from a timeout that needs to be increased. Xavier: there were two problems; one limit had to be increased and the system rebooted; also, some files were already back from tape but could not be copied in time to the required destination and the timeout was exceeded.]
    - ◇ [Eddie: SAM SRM tests were failing yesterday for INFN (for ATLAS and LHCb) and are still failing today.]
  - ◆ T2s:
    - ◇ No major problem
- CMS reports -
  - ◆ Experiment activity
    - ◇ Continued proton-proton work. CMS does not expect to send data from P5 to T0 during this time. HI not expected back until Saturday morning.
  - ◆ CERN and Tier0

- ◊ Quiet at CERN.
- ◆ Tier1 issues and plans
  - ◊ RAL Back
  - ◊ FNAL powered down again with power issues. This is an issue with reprocessing, but not an issue for HI data
  - ◊ [IanF: there was a Savannah ticket about two files at ASGC, not yet clear if this was a site issue or a CMS issue. Suijian: will check.-->Solved: "When the user created this ticket. She didn't assign it to ASGC, thus we couldn't see it".]
- ◆ Tier-2 Issues
  - ◊ No big Tier-2 Issues
- ALICE reports [↗](#) -
  - ◆ GENERAL INFORMATION:
    - ◊ Average rate from P2 to CASTOR 2GB/sec with a max rate of 2.5 GB/sec
    - ◊ Average rate copy rate from aliceT0 to alicedisk 2.5 GB/sec
    - ◊ Waiting for Saturday for the HI
  - ◆ T0 site
    - ◊ 8308 jobs running currently with picks of 10k! [IanF: are you going for opportunistic T0 resources (getting them from other people)? Lola: no, just using ALICE resources (Grid resources).]
  - ◆ T1 sites
    - ◊ SARA/NIKHEF: Issue reported on Wednesday has disappeared "mysteriously". We are going to do some tests if we can reproduce it and find the cause. [Onno: was problem at SARA or at Nikhef? Lola: it was at SARA. Onno: then it is probably due to the dcache upgrade on Tuesday, some pools were not started properly.]
    - ◊ FZK: CREAM submission has stopped. We are in contact with the site admin
  - ◆ T2 sites
    - ◊ Usual operations
- LHCb reports -
  - ◆ Experiment activities: Validation of the reprocessing in progress. Migrated SAM critical tests to a new OS VOBOX and some temporary errors (not related to the site services) may have been experienced during this intervention this morning. Situation restored.
  - ◆ New GGUS (or RT) tickets: 1 at T1 (none at T0, T2)
  - ◆ Issues at the sites and services
    - ◊ T0: ntr
    - ◊ T1 site issues:
      - CNAF : problem of lcas authorization on the gridftp servers. There was an update on Nov 17 and something went wrong (fixed)
      - RAL: Observed issues with the SRM at RAL yesterday (read-only file system cannot write)
      - IN2p3: shared area issues reported by SAM
      - [Eddie: Investigating why a SAM CE test, the lhcb-availability, is incorrectly included in the results of the critical tests although that it is defined as a non-critical one.]

Sites / Services round table:

- FNAL: The scheduled power outage went well & all services were restored on schedule. Around 3 PM, the main computing building at FNAL had a 1400 Amp breaker trip that took out the largest UPS circuit. This was not the breaker that was replaced during the morning shutdown. An investigation started and was expected to last until the late Thursday evening hours. Since most computing people had been in since around 3 AM for the scheduled outage, it was decided to keep computing equipment down until at least 6 AM. It is expected the circuits will be re-energized at 6 AM if the fault is

## WLCGDailyMeetingsWeek101115 < LCG < TWiki

understood. I have decided that CMS will start restoring services at 8 AM Friday if possible. I won't be attending the Friday WLCG OPS meeting. This info is accurate as I know it, but of course subject to corrections for any misunderstanding I may have...Jon

- NDGF: ntr
- BNL: ntr
- IN2P3: ntr
- KIT: nta
- PIC: ntr
- ASGC: nta
- NL-T1: nta
- RAL: ntr
- OSG: FNAL power outage has affected some of our systems, we are recovering

AOB: none

-- JamieShiers - 10-Nov-2010

---

This topic: LCG > WLCGDailyMeetingsWeek101115

Topic revision: r23 - 2010-11-22 - SuijianZhou



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