

# Table of Contents

<b>Week of 110627.....</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:.....	3
Wednesday.....	5
Thursday.....	8
.....	9
Friday.....	10

# Week of 110627

## 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(Lukasz, Stefan, AndreaV, Maarten, Ignacio, Massimo, Jamie, Dan, Alessandro Steve, Daria D); remote(Todd, Ernst, Karen, Gonzalo, Oliver, Jon, Ron, Kyle, Maria Francesca, Rolf, Tiju).

Experiments round table:

- ATLAS reports -
  - ◆ T0/CERN: Transfer errors from CERN-PROD\_DATADISK (GGUS:71915) late Friday, early Saturday: related to the broken disk server with still 10% missing files. IT turned the disk server off Saturday at noon (CET). Waiting for update today. GGUS:71942 is probably a dupe.
  - ◆ T1s:
    - ◇ BNL voms service was down Saturday evening from ~17:30-20:00 UTC, was restarted (GGUS:71926)
      - The incident was caused by a known defect. Clients which cannot be served either drop off to early or stay connected indefinitely and this does not trigger proper service fail over.
    - ◇ RAL had oracle problem which brought down the SRM from Saturday ~21:00 to midday Sunday (declared unscheduled downtime Sunday morning) (GGUS:71928)
    - ◇ NIKHEF had panda jobs failing with "could not open connection to tbn18.nikhef.nl" (GGUS:71931).

- CMS reports -
  - ◆ LHC / CMS detector one long fill, otherwise not a very successful weekend
  - ◆ CERN / central services
    - ◇ Castor problems yesterday, reported 14:30 UTC with GGUS alarm ticket GGUS:71934 as suggested by Massimo on Saturday, called 75011 to verify ticket was received, was not the case, so I forwarded the ticket to computer.operations@cernSPAMNOTNOSPAMPLEASE.ch and the shifter called the Castor piquet. Jan Iven and Steve Traylen both looked into this, solution in the end came from Jan, thanks! Solution: Tape functionality has been moved to a different machine. Unclear whether the stuck "rsyslog" server or the dead tape migration/recall server were the main underlying issue. Still one incident open for Castor: INC:047937: CMST3 file acces problem with scheduling the d2d copies related to first item?
    - ◇ LSF: had trouble with low efficiency user jobs in lxbatch queue reducing the CMS share INC:047879, INC:047808 LSF bug: machines still running user jobs and the slots were taken but escaped LSF process tree, LSF is tracking these proceses (it knows they're there), but is unable to kill them even though the associated job has clearly run out of wall time - and in fact the job wrapper has gone. LSF: we see a few 12 core machines in the cmst0 queue but they are not put into production yet. When will this happen? INC:048026
      - This will ba done ASAP
  - ◆ T1 sites:
    - ◇ T1\_IT\_CNAF: FTS submit problems: GGUS:71944 -> solved: vobox proxy renewal was not working
    - ◇ T1\_UK\_RAL: SRM aborts, transfer from T0 not working GGUS:71947
  - ◆ AOB Next CRC from tomorrow: Nicolo Magini
- ALICE reports -
  - ◆ T0 site
    - ◇ Debugging low CPU efficiencies of ALICE jobs grid-wide; new jobs are being submitted with a debug flag defined to help find out how they get stuck.
  - ◆ T1 sites
    - ◇ IN2P3: Torrent fully working at last!
  - ◆ T2 sites
    - ◇ Nothing to report
- LHCb reports -
  - ◆ Experiment activities: Staging problem (DIRAC side) fixed on Friday evening releasing an emergency patch of the stager system. Things seem to be back to normality. Since Friday took 18TB of files spread on 50 runs.
  - ◆ New GGUS (or RT) tickets:
    - ◇ T0: 0
    - ◇ T1: 3
    - ◇ T2: 0
  - ◆ Issues at the sites and services
    - ◇ T0:
      - 3D Streams: Notified about one apply process aborted in the LFC streams replication. 1 hour stoppage of the service to fix the problem.
    - ◇ T1
      - RAL : Sunday received notification from RAL : Due to some problems on Oracle databases behind Castor service all castor instances have been put in downtime. Discovered this morning a lot of timeout accessing turls (SRM issue). This is also reflected in the huge amount of jobs pending the transfer of their output to RAL SEs. Open an internal ticket (84699)

- GRIDKA: Still backlog (4K waiting tasks, due to the Storage issues reported last week) but it is recovering very quickly and situation is getting much better now.
  - The meeting LHCb-KIT took place (no details yet). Nevertheless the back-tail of jobs is being reabsorbed.
- GRIDKA: Problem with the pilot jobs submitted through CREAM CEs aborting (GGUS:71952 [↗](#))
- NIKHEF: All pilot jobs submitted through CREAM CEs are aborting (GGUS:71955 [↗](#))

Sites / Services round table: Sites:

- ASGC: ntr
- CNAF: ntr
- BNL: nothing to add
- FNAL: ntr
- IN2P3: ntr
- KIT: nothing to add
- NDGF: Powercut (Saturday ~4pm. UPS could not cope with it. Total experience downtime around 90')
- NL-T1: Due to a bug in a script detecting and deleting dark data (left over uncatalogued files) a few file (10-20) have been lost (ATLAS, LHCb and ALICE effected). Experiments will be informed of the full list when the investigation is finished.
- OSG: ntr
- PIC: ntr
- RAL: The issue reported by LHCb (issue 84699) has been fixed

Central services:

- Dashboard: False "red" in the LHCb dashboard (successful tests erroneously reported as failed). Investigation going on
- CASTOR: This week upgrades: ATLAS moved to Wed afternoon. CMS suggest to move their upgrade (scheduled on Thu) to next week (allow 72h of no-beam to finish data processing)

AOB:

## Tuesday:

Attendance: local();remote().

Experiments round table:

- ATLAS reports -
  - ◆ T0/CERN:
    - ◇ ALARM to CERN-PROD due to FTS server issues: GGUS:71958 . After the fix of the problem there were many FTS jobs stuck in ready state for many hours, we canceled what we found but it would be good if this issue could be checked server side.
    - ◇ CERN-PROD issue in exporting some RAW files: GGUS:71960 . The problem seemed to be different from the FTS one, now it is not present anymore
- CMS reports -
  - ◆ LHC / CMS detector
    - ◇ fill with 60 pb<sup>-1</sup>
  - ◆ CERN / central services

Tuesday:

- ◊ CLOSED: CASTORCMS T1TRANSFER got full, blocking transfers to/from CERN. ALARM GGUS:71969 submitted at 13:48 UTC, experts started working by 14:10 UTC, ticket was updated at 17:35 UTC, solved at 07:47 UTC. This morning, a PHEDEX agent doing unnecessary staging requests to t1transfer was also stopped. Impact: ~0% transfer quality for ~20 hours.
  - The actual problem (T1TRANSFER) was solved around 17:00 UTC (reclaimed space being available on T1TRANSFER). The instabilities overnight were due to a combination of effects (slow space reclaiming, FTS instabilities GGUS:71960, one site being unstable)
- ◊ OPEN: INC:047937: CASTORCMS CMST3 file access problem with scheduling the d2d copies.
- ◊ OPEN: 12 core machines to be put in production on LSF cmst0 queue INC:048026
- ◊ OPEN: jobs on rebooted nodes need to be cleaned up manually on LSF INC:046919
- ◊ SCHEDULED INTERVENTIONS: proposing Tuesday July 5th as the date for the upgrade of CASTORCMS to 2.1.11
- ◆ T1 sites:
  - ◊ CLOSED: T1\_ES\_PIC: environment problem on PhEDEx VOBOX causing issues in transfer verification, SAV:121838
- ◆ T2 sites:
  - ◊ IN PROGRESS: T2\_IN\_TIFR: mapping issues for lcgadmin role, SAV:121822
  - ◊ OPEN: Nagios tests for SRM not submitted to T2\_ES\_CIAMAT srm.ciemat.es SAV:121790
- ALICE reports -
  - ◆ T0 site
    - ◊ Low CPU efficiencies of ALICE jobs: some progress, but the real culprit has not yet been found.
  - ◆ T1 sites
    - ◊ Nothing to report
  - ◆ T2 sites
    - ◊ Nothing to report
- LHCb reports-
  - ◆ Experiment activities: Very good fill since yesterday (took 6.6 pb-1). One issue with DIRAC-ONLINE with transfer requests to CASTOR piling up in the buffer. Issue not understood but workaround to kick these pending requests one by one to catch up the backlog.
  - ◆ New GGUS (or RT) tickets:\*
    - ◊ T0: 0
    - ◊ T1: 0
    - ◊ T2: 0
  - ◆ Issues at the sites and services
    - ◊ T0
      - NTR
    - ◊ T1
      - RAL : Many jobs in input data resolution. Suffering two different problems over there:
        - Tape recall performances issues (under discussion at RAL)
        - Bug in LRU for which files staged since some time but with fresh access requests, instead of being pinned down to the cache might be garbage collected. Shaun is thinking about a patch of this bug but would take a while to get deployed.
      - GRIDKA: The huge backlog of the previous days has been almost totally drained performing very well in the last days.

- GRIDKA: Problem with the pilot jobs submitted through CREAM CEs aborting (GGUS:71952). Problem still there: we only have pilot jobs submitted via LCG-CE.
- NIKHEF: All pilot jobs submitted through CREAM CEs are aborting (GGUS:71955). It was a site reconfiguration issue. Fixed.
- SARA: Spikes of failing jobs in uploading output data every night at the same time. It looks like some regular script runs at SARA. In touch via private mail with Ron

Sites / Services round table: Sites:

- ASGC: ntr
- CNAF: ntr
- BNL: ntr
- FNAL: ntr
- IN2P3: ntr
- NDGF: ntr
- NL-T1: ntr
- PIC:
- RAL: ntr

Central services:

- Dashboard: LHCb SAM test dashboard solved. Root cause unclear (the service was reconfigured a few times in the recent past)
- Databases: Short incident on CMS offline DB. A runaway application was identified and restarted.
- CASTOR: As requested by ATLAS the 2.1.11 intervention of tomorrow will be anticipated (the slot 14:00-16:00 is not suiting the experiment). The CASTOR team will try to have the intervention finished before 9:00 tomorrow.
- Grid Services: A new **AFS UI** installation will be made available on Monday July 4th. **/afs/cern.ch/project/gd/LCG-share/new\_3.2** will point at **3.2.10-0** rather than **3.2.8-0**. The re-designation of **current\_3.2** will happen at some later date, TBC.

AOB:

- GGUS "alarm" problem still under investigation. It affects ALARM tickets against CERN-PROD. In this case the submitter (having the privileges to submit alarms) is suggested to follow the usual procedure (ticket submission) and then to call the operator (75011) to countercheck the ticket arrived to them. More tests will be done in the next days to solve this issue
- New GGUS Release will take place on 6th of July. The GGUS certificate for signed ALARM email notifications expiring around that time, it will be replaced on the date of the Release. Sites that use this certificate to accept emails from GGUS should take action and replace the DN according to Savannah:121803#comment1 [↗](#).

## Wednesday

Attendance: local(Edoardo, Lukasz, Maria D, Steve, Eva, Nilo, Massimo, Maarten, Nicolo, Alessandro, Ignacio);remote(Jon, Gonzalo, Maria Francesca, Rob, Tiju, Onno, Claudia, Dmitri, Rob, Rolf).

Experiments round table:

- ATLAS reports -
  - ◆ T0/CERN:
    - ◇ CERN-PROD failing exporting some data GGUS:72045

- ◇ follow up on FTS server outage of 27-28July night: ADC will change the way in which the switch to the latest DBRelease will be done. Under discussion technical details. ADCoS shifters will be asked to monitor the transfers of this critical file.
- ◆ T1
  - ◇ TAIWAN-LCG2 storage problems: GGUS:72048
  - ◇ IN2P3-CC FTS issue GGUS:72043
- CMS reports -
  - ◆ LHC / CMS detector
    - ◇ Machine Development testing 25 ns bunches
    - ◇ Possible short global run today
  - ◆ CERN / central services
    - ◇ CLOSED: CASTORCMS unavailable due to DB issues from ~14:00 UTC to ~18:00 UTC. GGUS:72027
    - ◇ CLOSED: Several disk servers offline on CASTORCMS T1TRANSFER from 21:00 UTC to 3:00 UTC. Correspondingly, observed drop of total space on t1transfer from 200 TB to 170 TB, and errors like "locality is UNAVAILABLE" in exports from T0 to T1s. Workaround for transfers was applied around 23:00 UTC in PhEDEx, forcing the restaging of unavailable files to t1transfer; transfer quality recovered aft
      - Note: free space still very low on T1TRANSFER, currently around 3%. Monitoring, will open new ticket if necessary.
      - Note: t1transfer dump currently not accessible. Could be useful to determine what needs to be cleaned up.
    - ◇ OPEN: INC:047937: CASTORCMS CMST3 file access problem with scheduling the d2d copies.
    - ◇ OPEN: 12 core machines to be put in production on LSF cmst0 queue INC:048026
    - ◇ OPEN: jobs on rebooted nodes need to be cleaned up manually on LSF INC:046919
  - ◆ T1 sites:
    - ◇ NTR
  - ◆ T2 sites:
    - ◇ T2\_RU\_PNPI: Need to commission uplink to T1\_TW\_AS GC for upload of custodial MC, SAV:121847
  - ◆ AOB:
    - ◇ I didn't have permission to submit ALARM tickets in GGUS, because the ALARMer role was associated to my old expired INFN certificate. Issue was fixed by GGUS support this morning: GGUS:72056
- ALICE reports [↗](#) -
  - ◆ T0 site
    - ◇ Low CPU efficiencies of ALICE jobs: investigations ongoing.
  - ◆ T1 sites
    - ◇ Nothing to report
  - ◆ T2 sites
    - ◇ Nothing to report
- LHCb reports -
  - ◆ Experiment activities: Almost all sites are well behaving for real data processing and user activities.
  - ◆ **New GGUS (or RT) tickets:**
    - ◇ T0: 0
    - ◇ T1: 0
    - ◇ T2: 0
  - ◆ Issues at the sites and services
    - ◇ T0

- NTR
- ◇ T1
  - RAL : Want to keep open these points about RAL to know how the site intends to
    - Tape recall performances issues (under discussion at RAL)
    - Bug in LRU for which files staged since some time but with fresh access requests, instead of being pinned down to the cache might be garbage collected. Shaun is thinking about a patch of this bug but would take a while to get deployed.
  - GRIDKA: Problem with the pilot jobs submitted through CREAM CEs aborting (GGUS:71952). Fixed!

Sites / Services round table: Sites:

- ASGC: Connectivity problems being investigated (?)
- CNAF:ntr
- BNL: ntr
- FNAL:ntr
- IN2P3:ntr
- KIT:ntr
- NDGF: dCache pool down to a faulty RAID controller: investigating. List of files at risk will be published tomorrow (if the problem is not solved)
- NL-T1:ntr
- OSG:ntr
- PIC:ntr
- RAL:Shaun in contact with the CASTOR developers (LHCb problem - LRU). Tape problem still under investigation

Central services:

- Dashboard:ntr
- Databases:ntr
- CASTOR:
  - ◆ Overnight problems created by a faulty switch in the computer centre infrastructure (tickets filed as network problem)
  - ◆ CMS problems: side effects of solving the CMS slow deletion problem (index creation failed blocking the entire CMS DB)
  - ◆ Upgrades/tests:
    - ◇ 2.1.11 ATLAS (done: it took about 2 extra hours due to the unexpectedly long DB scripts needed in the upgrade: elapsed time not consistent with tests and previous upgrades -notably castorpublic)
    - ◇ 2.1.11 for CMS, ALICE and LHCbnext week?
    - ◇ 2.1.11 test with ATLAS (new scheduler) and ALICE (possible rate test)
- Grid Services: ntr
- Network:ntr

AOB: (MariaDZ)

- Rephrased and repeats here yesterday's AOB due to its importance: New GGUS Release will take place on 6th of July. The GGUS certificate for signed ALARM email notifications expiring around that time, it will be replaced on the date of the Release. Sites that use this certificate to accept emails from GGUS should take action and replace the DN according to Savannah:121803#comment1 [↗](#). Please declare in the savannah ticket if you do use certificate-recognition.
- Non-delivery of ALARM emails to the T0 is traced down to a change of the sender's email address



from helpdesk@ggusSPAMNOTNOSPAMPLEASE.org (which is allowed to post) to apache@ggusSPAMNOTNOSPAMPLEASE.org (which was unknown to the e-groups so far). This change was introduced surreptitiously with the introduction of the new KIT exim mailer. The change of behaviour didn't show during the ALARM tests that follow every release because they are submitted via a script (which has control over the sender address) and not via the web ticket submission form. apache@ggusNOSPAMPLEASE.org is now one of the administrators of the 4 e-groups atlas-operator-alarm@cernNOSPAMPLEASE.ch, alice-operator-alarm@cernNOSPAMPLEASE.ch, cms-operator-alarm@cernNOSPAMPLEASE.ch, lhcb-operator-alarm@cernNOSPAMPLEASE.ch hence, allowed to post. Test ALARMS will be sent tomorrow 2011/06/30 during the day to the 4 LHC VOs. T1s please **IF you check the sender's address of GGUS emails please add apache@ggusNOSPAMPLEASE.org IN Addition to the helpdesk@ggusNOSPAMPLEASE.org and support@ggus.org** that you already know.

- The SNOW bug attributing comments to the wrong supporter when synchronising with GGUS is now submitted as <https://cern.service-now.com/service-portal/view-request.do?n=RQF0017147>

## Thursday

Attendance: local(Steve, Massimo, Nicolo, Shu-Ting, Maarten, Nilo, Fernando, Alessandro, Lukasz);remote(Michael, Karen, Jon, Gonzalo, Dmitri, John, Roland, Maria Francesca, Claudia, Rolf).

Experiments round table:

- ATLAS reports -
  - ◆ T0/CERN:
    - ◇ CERN-PROD failing export GGUS:72085 . One every day?
  - ◆ T1
    - ◇ NIKHEF-ELPROD issue exporting data GGUS:72083
- CMS reports -
  - ◆ LHC / CMS detector
    - ◇ Machine Development
    - ◇ Most subdetectors out
  - ◆ CERN / central services
    - ◇ free space recovering on T1TRANSFER. Also t1transfer dump available again, checking content. Thanks!
    - ◇ IN PROGRESS: myproxy.cern.ch unreachable from outside CERN GGUS:72090
    - ◇ OPEN: INC:047937: CASTORCMS CMST3 file acces problem with scheduling the d2d copies.
    - ◇ IN PROGRESS: 12 core machines to be put in production on LSF cmst0 queue INC:048026
    - ◇ OPEN: jobs on rebooted nodes need to be cleaned up manually on LSF INC:046919
  - ◆ T1 sites:
    - ◇ FNAL: ~2 hour degradation in transfer quality with gridftp errors (host certificate/DNS name mismatch), recovered with no ticket submission.
  - ◆ T2 sites:
    - ◇ T2\_ES\_IFCA scheduled maintenance, not reflected in Dashboard downtime calendar: SAV:121854
    - ◇ T2\_BE\_IIHE: stageout issues for MC production for permission problems, CMS MC operators in contact with local admins to fix.
  - ◆ AOB:
    - ◇ CMS critical service map unavailable, submitted SAV:121896
    - ◇ GGUS:72096 ALARM test submitted to CERN to test if GGUS mails are received by CERN operators.

- ALICE reports [↗](#) -
  - ◆ General information: Proxy renewal connection time out to myproxy.cern.ch occurred in several sites. Under investigation.
  - ◆ T0 site
    - ◇ Nothing to report
  - ◆ T1 sites
    - ◇ Nothing to report
  - ◆ T2 sites
    - ◇ Nothing to report
  
- LHCb reports -
  - ◆ Experiment activities: Quite day. Nothing to report. LHCb week ongoing.
  - ◆ **New GGUS (or RT) tickets:**
    - ◇ T0: 0
    - ◇ T1: 0
    - ◇ T2: 0
  - ◆ **Issues at the sites and services**
    - ◇ T0
      - NTR
    - ◇ T1
      - NTR

Sites / Services round table: Sites:

- ASGC: ntr
- CNAF: ntr
- BNL: ntr
- FNAL: gFTP problem (CMS) due to a side-effect of a certificate replacement
- IN2P3: ntr
- KIT: ntr
- NDGF: dCache problem solved (no file loss). Volume still in R/O for observing its behaviour.
- NL-T1: Disk server lost connection to the metadata server: restarted. Put in place a procedure to do it automatically.
- OSG: ntr
- PIC: ntr
- RAL: Several intervention scheduled on the 5th of July (Tuesday). 2h downtime in the morning

Central services:

- Dashboard: CMS problem (service map) solved.
- Databases: ntr
- CASTOR: Satisfactory solution to the sapce reclaim (garbage collection) problem. Added as an hotfix to 2.1.11
- Grid Services: myproxy problem solved (firewall rules expired and the warning was lost)

AOB: (MariaDZ) GGUS:72094 [↗](#) Test ALARM for Atlas was successful (the whole chain - operator and SNOW interface). Other experiments are done by their supporters. LHCb is GGUS:72098 [↗](#) and went well. CMS is GGUS:72096 [↗](#) (SMS is late but the GGUS email was delivered because the operators replied. ALICE is GGUS:72099 [↗](#). Operators' replied. Please stop telling submitters to call 75011.

## Friday

Attendance: local(Massimo, Nicolo, Lola, Alessandro, John, Lukasz, ShutLin, Eva, Steve);remote(Roberto, Michael, Ulf, Kyle, Onno, Xavier, Jon, John, Rolf, Gonzalo, Giovanni).

Experiments round table:

- ATLAS reports -
  - ◆ T0/CERN:
    - ◇ CERN-PROD ALARM GGUS:72132 [↗](#) writing into T0Merge pool
      - At first order it is due to high rate on a busy D1 pool: concurrent writes can exhaust the space. Retry solves the problem. More investigation going on.
    - ◇ CERN-PROD failing export GGUS:72085 . One every day? -- Please note that it's not an error of copy from the previous day, the problem reappeared yesterday night again.
      - CASTOR diskserver (in case of failures) are not even contacted by FTS to initiate the transfer. The network logs are being studied but it is unlikely this comes from faulty network equipment or similar reasons. Since FTS takes long (minutes instead of tens of seconds) to initiate the transfer (creating the TURL) it is possible there is some problem in FTS itself (investigating).
    - ◇ CERN-PROD, since 5 days GGUS:71915 , it seems that there are still HW issues not solved.
      - In order to avoid to reinstall all 22 data disk on a new machine (error prone) sysadmins are trying to change component by component (data cabling, back plane, controller, power distribution... to nail down the problem).
    - ◇ CERN-PROD GGUS:72145 : submission not working to 2 cream ce. Moreover it seems that ATLAS cannot run more than 500jobs in parallel.
    - ◇ ATLAS autoblacklisting system of DDM endpoints blacklisted srm-atlas endpoints due to a scheduled downtime of srm-eosatlas. This should be fixed within the ATLAS auto-blacklisting system.
- CMS reports -
  - ◆ LHC / CMS detector
    - ◇ Machine Development
    - ◇ ECAL and CSC in global run
  - ◆ CERN / central services
    - ◇ Migration of production data from CASTORCMS CMSCAF to EOSCMS in progress. Peak rates > 1 GB/s with xrscp from CASTOR, and > 700 MB/s through FTS+SRM from T1s. Currently 260 TiB (~50%) migrated, should finish by next week. Then need to migrate other disk pools and user data.
    - ◇ vocms110 unreachable: urgent to get it back. Any estimation?
      - Intervention priority risen. Machine is back now
  - ◆ T1 sites:
    - ◇ RAL: CMS data operations to replace decommissioned LCG-CE with CREAMCE in the list of queues used at RAL
  - ◆ T2 sites:
    - ◇ NTR
- ALICE reports [↗](#) -
  - ◆ T0 site
    - ◇ Nothing to report
  - ◆ T1 sites
    - ◇ IN2P3: Problem with VOBOX AliEn services. Under investigation.
  - ◆ T2 sites

◇ Nothing to report

• LHCb reports -

- ◆ Experiment activities: Processing of last taken data going smoothly. Closed previous reprocessing activities. Validating another big reprocessing (FemtoDST production) that ideally will take place during this machine TS. \* **New GGUS (or RT) tickets:**

◇ T0: 1

◇ T1: 1

◇ T2: 0

- ◆ **Issues at the sites and services**

◇ T0

- CASTOR: There are 9 files from one production (SDSTs) that are unavailable in Castor (status INVALID). Ticket open: INC048765 Any news?
- VOBOX: received this morning an ALARM for one of our voboxes running out of /opt space (at 99% occupancy). Triggered DIRAC expert to clean up a bit this space.
- LFC: upgrade to SL5: some time next week.

◇ T1

- RAL: They had to revert back the xrootd interface (put in production yesterday) as it was presenting still some instabilities. Planned major intervention on their network on Tuesday morning planned to use this DT slot to also upgrade to CASTOR 2.1.10-1 .
- GRIDKA: Open a ticket to ask them to clean up old test files outside any token but wasting 4-5 TB of space. (GGUS:72114)

Sites / Services round table: Sites:

- ASGC: ntr
- CNAF: ntr
- BNL: ntr
- FNAL:
  - ◆ Due to high (external) temperature, part of the farm will be switched off to keep up with the available cooling power
  - ◆ Monday FNAL won't connect (national holiday)
- IN2P3: Excessive load (making SAM tests fail) due to an ATLAS Tier3 downloading too much data to restored its data store after a crash. ATLAS pointed out that sites should follow the procedures prepared by the experiment.
- KIT: Communication problem (dCache gdfs). 9 pools went offline overnight. They were back in production at 8:00am
- NDGF: RAID problem. No data loss (confirmed) but the volume will stay in RO for a longer while to understand the problem
- NL-T1: Observed higher failure rate of LHCb jobs for overnight jobs than for those ran during day time. Investigating. Need for more information from the experiment. Could it be a side effect of more production jobs running at night and the day having more analysis jobs?
- OSG: Monday US colleagues won't connect (national holiday)
- PIC: ntr
- RAL: ntr

Central services:

- Dashboard: ntr
- Databases: Yesterday intervention (RAC6) OK. More intervention next week. CMS noticed a problem on streamer replication (overnight). It is confirmed and it was due to an unusual

(physiological) peak activity.

- CASTOR:
  - ◆ 2.1.11: CMS on Tuesday, LHCb and Alice on Wed/Thu (tbc)
  - ◆ ATLAS stress test: Monday 9:00 - 13:00
  - ◆ ALICE transfer test (upgrade of CDR at the pit): Friday (tbc)
- Grid Services:
  - ◆ ACRON upgrade: 4th of July 10:00 (Authenticated CRON)
- Network:

AOB:

-- JamieShiers - 27-Jun-2011

---

This topic: LCG > WLCGDailyMeetingsWeek110627

Topic revision: r30 - 2011-07-01 - MassimoLamanna



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