

Table of Contents

Week of 110613.....	1
Daily WLCG Operations Call details.....	1
WLCG Service Incidents, Interventions and Availability, Change / Risk Assessments.....	1
General Information.....	1
Monday:.....	1
Tuesday:.....	1
Wednesday.....	3
Thursday.....	6
Friday.....	7

Week of 110613

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:

- No meeting - CERN closed.

Tuesday:

Attendance: local (AndreaV, Alessandro, Guido, Maarten, Lukasz, Ken, Jan, Ignacio, Manuel, Alexandre, Eva, MariaDZ); remote (Jon, Gonzalo, Kyle, Jhen-Wei, Jeremy, Thomas, Tiju, Ronald, Foued, Rolf, Daniele; Joel).

Experiments round table:

- ATLAS reports -
 - ◆ T0/CERN
 - ◇ GGUS:71471 [srm-atlas.cern.ch](#) connectivity problems, caused by aggressive FTS settings, ATLAS relaxed settings. Related to EOS migration.
 - ◇ Same ticket was escalated to ALARM but alarm went nowhere, needs investigating
 - ◇ XROOTD daemon hanging on castoratlas. Promptly restarted by the admins. GGUS:71521 [\[Ian: this was part of a scheduled downtime.\]](#)
 - ◆ T1
 - ◇ FZK MCTAPE full, tape system fixed but endpoint excluded pending recovery of free space. GGUS:71466

- ◇ PIC srm problem, promptly solved by the site ("hanged processes in the core tape server was causing PNFS to die"). GGUS:71447 [↗](#)
- ◇ [Alessandro: Triumph is marked as not available in the dashboard because Cream CE tests fail, will contact Triumph to follow up. But the LCG CE tests are ok, so the dashboard should not mark the site as unavailable, opened a Savannah ticket about this issue.]
- CMS reports -
 - ◆ LHC / CMS detector
 - ◇ Continued running
 - ◆ CERN / central services
 - ◇ Log files of SAM tests are no longer visible through the Dashboard. Savannah:121517 [↗](#) assigned to Dashboard team.
 - ◇ Site downtime calendar interface now reads from Dashboard, and now shows CERN down for the next ten years. Dashboard team working on this too. [Alessandro: ATLAS does not see this problem, maybe this appears in CMS-specific views?]
 - ◆ Tier-0 / CAF
 - ◇ Generally keeping up with data; machine performance not as good as feared/anticipated. Switching to new CMSSW version today.
 - ◆ Tier-1
 - ◇ Many stuck transfers to RAL; being investigated.
 - ◆ Tier-2
 - ◇ MC production and analysis in progress
 - ◆ AOB
 - ◇ Ken Bloom is CRC until June 21
- ALICE reports [↗](#) -
 - ◆ General information: We are running since last week around 30K jobs, Pass-0 and Pass-1 reconstruction of a couple of runs ongoing.
 - ◆ T0 site
 - ◇ Nothing to report
 - ◆ T1 sites
 - ◇ CNAF: GGUS:71401 [↗](#). Investigating the reason why there are so many expired jobs at the site and the output of the jobs are not collected. It looks like there are some NFS problems on the client. Experts investigating
 - ◇ FZK: problems replicating data to FZK::TAPE due to a 'permission denied' error message. Experts have ben contacted already [Foued: problem is being investigated.]
 - ◆ T2 sites
 - ◇ Usual operations
- LHCb reports -
 - ◆ Experiment activities
 - ◇ More than 25 TB of data collected in the last 3 days. Stripping of 2011 data nearly achieved and re-stripping of 2010 data will start today.
 - ◆ New GGUS (or RT) tickets:
 - ◇ T0: 0
 - ◇ T1: 2
 - ◇ T2: 0
 - ◆ Issues at the sites and services
 - ◇ T0
 - A limit of jobs / DN has been relaxed now, allowing more than 2.5k jobs / DN.
 - ◇ T1
 - GRIDKA: Spikes of analysis jobs failing accessing data (GGUS:71412 [↗](#)).

Fixing the problem with dcap ports resolved the problem. Ticket closed.

· CNAF: Investigating a problem with 38 files not properly copied over to site, despite being registered in the LHCb Bookkeeping.

Sites / Services round table:

- FNAL: ntr
- PIC: ntr
- OSG: ntr
- ASGC: ntr
- GridPP: ntr
- NDGF: SRM downtime today for upgrading Postgres, went ok and all should be up and running now
- RAL: ntr
- NLT1: ntr
- KIT: short downtime today to upgrade CMS dcache, should be working again now
- IN2P3: ntr
- CNAF: there were GPFS problems for about half an hour, now fixed

- Dashboard services: last week ATLAS DDM upgraded to a new hardware, there was no disruption to the service
- Database services: the third node of CMSR has just rebooted, this is being investigated
- Storage services:
 - ◆ last remaining CASTOR SLC4 head nodes are being decommissioned, some aliases are also being updated
 - ◆ tomorrow will upgrade CASTOR CERN T3 (aka Atlas T3, aka CMS T2) to upgrade to 2.1.11, this is a scheduled downtime known by the experiments
 - ◆ acron service will be upgraded on July 4, needed downtime is being evaluated, details will be posted

AOB:

- Alessandro: what is the status of the Kerberos KDC issue? Andrea: ROOT patches are being completed and will be deployed to client-side applications, but the fact that the buggy version of ROOT had been used for many months without causing KDC high load may signal that some server-side issues have appeared in the last few weeks and should be investigated. Alessandro: ATLAS was using the same ROOT 5.26 for around 6 months without problems. Ian: maybe the recent changes on the KDC side (Heimdal vs AD) could have had an impact on this issue. Manuel: some jobs were still contacting the old Heimdal server.

- MariaDZ: please send GGUS tickets that don't get sufficient attention, of concern to your experiments or sites and comments to the proposal
https://twiki.cern.ch/twiki/bin/view/LCG/WLCGCriticalServices#GGUS_Type_of_Problem_field on the implementation of a Type of Problem field for TEAM and ALARM tickets to me.

Wednesday

Attendance: local (AndreaV, Ken, Guido, Lukasz, Tim, Luca, Alessandro, Maarten, Alexandre, Massimo, Manuel, MariaDZ); remote (Jon, Gonzalo, Onno, Jhen-Wei, Daniele, Thomas, Tiju, Rolf, Foued, Kyle; Roberto).

Experiments round table:

- ATLAS reports -
 - ◆ T0/CERN

- ◇ problem with the Castor LSF scheduler; promptly reported by the Point1 computing shifter, scheduler in strange status, restarted and problem disappeared. Small effects on users/VOs
- ◇ [Andrea: KDC saw new flood of Kerberos requests from ATLAS - to be discussed in detail in the service round table]
- ◆ T1
 - ◇ still no news from PIC about GGUS 71389 [↗](#), after updating of 10th June
- CMS reports -
 - ◆ LHC / CMS detector
 - ◇ Continued running, nice long fill yesterday.
 - ◇ New HLT menu to be deployed for next fill.
 - ◆ CERN / central services
 - ◇ CASTOR intervention in progress as I write, expected to be transparent.
 - ◇ Savannah:121517 [↗](#) reported yesterday is fixed.
 - ◇ Any news about the funny CERN downtimes? At yesterday's meeting, we determined that part of the trouble was that perhaps CREAM CE's were improperly listed in a downtime that was to mark the decommissioning of LCG CE's. [Maarten/Alessandro: can the decommissioned CE's be removed from GOCDB? Manuel: will follow up.]
 - ◆ Tier-0 / CAF
 - ◇ Switch to new CMSSW took place.
 - ◆ Tier-1
 - ◇ CASTOR down at RAL due to LSF problem; they're working on it.
 - ◆ Tier-2
 - ◇ MC production and analysis in progress
 - ◆ AOB
 - ◇ CRC will be off site tomorrow and will join meeting by phone. Email to cms-crc-on-duty@cernNOSPAMPLEASE.ch always works!
- ALICE reports [↗](#) -
 - ◆ T0 site
 - ◇ Nothing to report
 - ◆ T1 sites
 - ◇ FZK: problems replicating data to FZK::TAPE due to a 'permission denied' error message. Upgrading the authorization library may solve the problem. Ongoing
 - ◆ T2 sites
 - ◇ Usual operations
- LHCb reports -
 - ◆ Experiment activities:
 - ◇ Last night more than 10TB collected with spike of luminosity at 3.7×10^{32} . Launched latest version of the stripping (Stripping14) on 2010 data.
 - ◆ New GGUS (or RT) tickets:
 - ◇ T0: 0
 - ◇ T1: 0
 - ◇ T2: 0
 - ◆ Issues at the sites and services:
 - ◇ T0
 - Running very few jobs compared the number of waiting jobs targetted to CERN. [Manuel: strange, limit was recently increased from 25000 to 8000. Roberto: sorry, indeed the number of jobs is ramping up now, so this is a non-issue.]
 - Problem with 3D replication. It appeared during the operation of restoration of historical data for one of the LHCb offline applications. Now it seems OK.

[Luca: Streams were down one hour during a maintenance operation. This was due to a mistake by the DBAs, apologies for this.]

◇ T1

· CNAF: Investigating a problem with 38 files not properly copied over to site, despite being registered in the LHCb Bookkeeping. Has been understood: flaw in the Dirac DMS clients for certain type of operations.

Sites / Services round table:

- FNAL: ntr
- PIC: incident last night on OPN link for 14 hours due to cut fiber, rerouted through 1Gbit/s backuplink so service was degraded (mainly for CMS), eventually fixed at 4am; in a few weeks will use Internet as backup link, which will not be limited to 1Gbit/s but will not be dedicated.
- NLT1: ntr
- ASGC: ntr
- CNAF: ntr
- NDGF: ntr
- RAL: LSF performance issues as reported by CMS, fixed after a 1.5h intervention
- IN2P3: question for ALICE, why do we observe slow jobs since a few days (around 10% of all ALICE jobs)? [Maarten: observed Grid-wide for MC production, being investigated, will be discussed at ALICE task force meeting tomorrow.]
- KIT: ntr
- OSG: ntr

- Database services: problem seen with Streams apply at BNL, being followed up.
- Dashboard services: problems reported by CMS yesterday are both fixed (log files are now available; 10 year downtime is fixed in CMS-specific dashboard).

- Kerberos services: updated on the KDC flood problem
 - ◆ Tim: there was another KDC flood from ATLAS users tonight, will discuss it at the AF tomorrow.
 - ◆ Massimo: was this on the AD KDC? Tim: yes.
 - ◆ Alessandro: reported this to David Rousseau to organize the patching. Andrea: will discuss this at the AF tomorrow, too.
 - ◆ Andrea: can the bug in the SLC5 Kerberos (reported by Gerri) be relevant to this issue? Tim: yes, definitely think that the bug in the SLC5 Kerberos 1.6 may be responsible for the "Request is a replay" seen on the ATLAS xrootd redirector. The bug should be fixed in the SLC6 Kerberos 1.9.
 - ◇ Andrea: could it be useful to migrate the ATLAS xrootd redirector to SLC6 then? Massimo and Tim: will follow this up offline.
 - ◇ Jon (FNAL): note that there is another known bug introduced in the SLC6 Kerberos 1.9. This has been reported by many users of SLC6 as it prevents them from accessing the FNAL KDC. Andrea: very useful information; it would still be worthwhile to investigate whether the upgrade of the ATLAS xrootd at CERN to Kerberos 1.9 would allow users to access the CERN KDC and solve the KDC flood issue.
 - ◆ Alessandro: why does CMS not see this? Tim: this can be triggered as soon as there are more than 10 requests per second, which maybe depends on the way ATLAS groups their file accesses via xrscp. Massimo: also seeing a high load on xrootd from xrscp, could this be relevant? Tim: maybe it is related, depending on the way xrscp file requests are grouped together.
 - ◆ Alessandro: is there any hint why this started last month? There have not been major optimizations on the ATLAS side in the way files are grouped together? Tim: not clear if there is any correlation, but problem was first seen two weeks after a major change in

Kerberos encryption. No KDC flod was seen during the first two weeks after the change (though maybe users were not submitting this kind of jobs at that time).

◇ Maarten: is it conceivable to revert that change? Tim: no this would not be possible now, as many other change which depended on that one have already been done.

AOB: none

Thursday

Attendance: local(Massimo, Maria D, Maria G, Guido, Maarten, Manuel, Alexandre, Jamie, Mike);remote(Jon, Roberto, Jeremy, Ken, Jhen Wei, Thomas, Gareth, Roland, Andreas, Gonzalo, Daniele, Alessandro).

Experiments round table:

- ATLAS reports -
 - ◆ T0/CERN
 - ◇ problem with the privileges of user ddmusr03 on the castorcernt3 stager; problem promptly fixed ("the problem was due to the move to the SLC5 nodes, unfortunately the regexp in the CUPV configuration was not covering them")
 - ◆ T1
 - ◇ still an open issue with FZK-LCG2_ATLASMCTAPE (GGUS 71466): ATLAS writing data faster than tape migration speed
- CMS reports -
 - ◆ LHC / CMS detector
 - ◇ No real fill until just this morning.
 - ◇ New HLT menu is deployed, looks fine so far.
 - ◆ CERN / central services
 - ◇ CASTOR intervention was no problem.
 - ◆ Tier-0 / CAF
 - ◇ It's running.
 - ◆ Tier-1
 - ◇ No problems.
 - ◆ Tier-2
 - ◇ MC production and analysis in progress
- ALICE reports [↗](#) -
 - ◆ T0 site
 - ◇ AFS disabled from ALICE xrootd servers
 - ◆ T1 sites
 - ◇ FZK: problems replicating data to FZK::TAPE due to a 'permission denied' error message. Upgraded the authorization library may solve the problem. Solved
 - ◇ CNAF : GGUS:71401 [↗](#) many expired jobs at the site. Under investigation
 - ◆ T2 sites
 - ◇ Usual operations
- LHCb reports -
 - ◆ T0
 - ◇ CERN: Ticket opened against LSF support: still too few jobs running in the system (at most 800) after the increase of the limit to 8000. Yesterday at the time of the same issue was reported it was noticed a ramp up (from 300 to 800) then it has been thought the problem was not longer there. GGUS:71608 [↗](#)
 - ◆ T1

- ◇ GRIDKA: more than half of user and production jobs accessing data via dcap protocols failing because not consuming CPU. Preliminary investigations from Doris seem to pin down a load problem with the disk pools (not tokenized) in front of the tape. (GGUS:71752). We are not entirely convinced with this argument as user jobs do not use tape system.

Sites / Services round table:

- FNAL: ntr
- PIC: ntr
- NLT1: ntr
- ASGC: ntr
- CNAF: investigating on the ALICE issue
- NDGF: This morning the SRM suffered a 5-minute interruption due to the DB being inadvertently restarted
- RAL: ntr
- KIT: Investigating on the reported problems. The slow writing seems related to cartridge handling (being investigated)
- OSG: ntr

Central services:

- Dashboard: ntr
- Storage: ntr
- Batch: CEs retired from GOCDB (the one generating issues discussed in the last two weeks on the GOCDB side)

AOB: Maria Dimou mentioned ticket GGUS:71047. From a routing problem it suffer of a known issue: Once a GGUS ticket is reopened after 3 days since SNOW goes in resolved status it cannot reopen the SNOW ticket (because the SNOW ticket reaches its final state -closed- after 3 days being -resolved-). The suggestion (guideline/workaround) is to start a fresh ticket in these cases.

Friday

Attendance: local (Lukasz, Ken, AndreaV, Jamie, Maria, Eva, Dirk, Guido, Massimo, Maarten, Alessandro); remote (Onno, Dmitri, Michael, Roberto, Jon, Gonzalo, Felix, Kyle, Daniele, Gareth).

Experiments round table:

- ATLAS reports -
 - ◆ T0/CERN
 - ◇ data transfer from CERN to BNL had one file failing and reporting problems with server lxfsrc3707.cern.ch (t0atlas pool); the server seems ok, waiting for the next attempt to see if it was just a temporary glitch (otherwise escalate the problem to SRM/Castor)
 - ◇ transparent intervention on the castoratlas xrootd redirector in the morning to disable the Kerberos replay cache: fast and harmless!
 - ◆ T1
 - ◇ still open issue with FZK-LCG2_ATLASMCTAPE (GGUS 71466): problem with the connection to the tape backend, more tape drives added as a temporary solution
 - ◇ errors in contacting CNAF Storm storage element (GGUS 71618), problem reported at 8:30pm and fixed at 9pm (just a glitch of storm); no problems ever since
- CMS reports -

- ◆ LHC / CMS detector
 - ◇ Took data until a beam dump a little while ago, including a VdM scan this morning.
- ◆ CERN / central services
 - ◇ I think OK.
- ◆ Tier-0 / CAF
 - ◇ Have been encountering a greater rate of "stuck" T0 jobs than usual over the past few days; operators are keeping an eye on it.
- ◆ Tier-1
 - ◇ PIC seeing inefficient production due to issues staging from tape, now appears to be resolved.
 - ◇ Transfers from CERN to IN2P3 had been failing, but seems to have recovered in last few hours.
- ◆ Tier-2
 - ◇ MC production and analysis in progress; people want production to happen faster, of course.
- ALICE reports [↗](#) -
 - ◆ General Information: Low efficiency of the jobs is being investigated.
 - ◇ [Maarten: this happens since a few weeks. Detected at IN2P3 but grid-wide. High priority item for investigation but not understood. Trouble is accessing remote files. Patterns involved have not changed. Recently for some unknown reason these accesses get stuck and jobs timeout, considered lost. However quite a lot of jobs finish ok. Trouble seems to be on xrootd layer but no recent changes tied to onset of this problem.]
 - ◇ [Massimo - files are opened or xrdcp-ed to WN? Maarten: opened directly from ROOT but in a multi-threaded program. Some sort of race condition that under certain circumstances thread enter deadlock like situation from which they sometimes emerge but meanwhile waste a lot of wall clock time.]
 - ◆ T0 site - Nothing to report
 - ◆ T1 sites
 - ◇ CNAF : GGUS:71401 [↗](#) many expired jobs at the site. Under investigation
 - ◆ T2 sites - Usual operations
- LHCb reports -
 - ◆ Experiment activities: Smooth data taking (order of 10pb-1 recorded last night and this morning).
 - ◆ New GGUS (or RT) tickets:
 - ◇ T0: 0
 - ◇ T1: 0
 - ◇ T2: 0
 - ◆ Issues at the sites and services
 - ◇ T0
 - CERN: Investigation on going on the low number of LSF job slots. Lowered the slots for Alice, may be this is going to help. GGUS:71608 [Maarten: maximum raised to much more than nominal (x 2) so that ALICE could take advantage of spare capacity. If LHCb is submitting jobs LSF should recognise that LHCb has not been receiving fair share and should eat up access being consumed by ALICE or indeed others. Still bug or other issue in LSF not understood.]
 - ◇ T1
 - GRIDKA: GGUS:71572 The situation is not good yet. We temporary banned Gridka SE to let the system draining the backlog of SRM requests. They put in place new h/w in the read-only pools in order to alleviate this issue.

Sites / Services round table:

- NL-T1 -
 - ◆ yesterday SARA SRM in unscheduled downtime as disk controller had to be replaced. Some files may have been unavailable. Fixed yesterday afternoon.
 - ◆ GGUS:71599 [↗](#) ticket from ATLAS assigned to NIKHEF. NIKHEF waiting for more information. Any news from ATLAS? [Ale - will update ticket.]
 - KIT - problems with LHCb pools during night and maybe problem memory leak. Will ask Doris if any news on LHCb problems (GGUS:71572 [↗](#)). [Roberto: jobs are not consuming CPU and get killed by Dirac. Could be related to high activity on RO disk pools, but this is not clear yet.]
 - BNL - we will have a transparent network intervention on Monday morning. gridftp doors will be moved from one switch to another "real fast" and don't expect any impact.
 - FNAL - ntr
 - PIC - ntr
 - ASGC - announcement: going to have scheduled UPS maintenance on Tuesday 21 June. UPS to protect network switch and data storage. 12h downtime from 02:00 - 10:00 UTC.
 - CNAF - ntr
 - RAL - ntr
 - OSG - ntr
-
- CERN DB - there will be an intervention on Monday on COMPASS DB which will be switched to different h/w - 1h of downtime.

AOB: none

-- JamieShiers - 08-Jun-2011

This topic: LCG > WLCGDailyMeetingsWeek110613

Topic revision: r12 - 2011-06-17 - OnnoZweersExCern



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