

Table of Contents

Week of 110606.....	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.....	5
Thursday.....	6
Friday.....	8

Week of 110606

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 (AndreaV, Dirk, Massimo, Lola, Manuel, Jan, Jarka, Alessandro, Michal, Peter, Lukasz, MariaDZ); remote (Tiju, Ron, Rob, Brian, Rolf, Xavier, Ulf, Daniele, Jon, Michael, Gonzalo, Felix; Ian, Roberto).

Experiments round table:

- ATLAS reports -
 - ◆ RAL-LCG2: Disk server gdss135 unavailable. Site investigates. "this disk server is part of the tape buffer and therefore no files are unavailable. The worst this error could cause is a few transient errors in tape recalls."
 - ◆ FZK-LCG2: 2 tape servers broken, ~190 files lost. Files declared as lost to ATLAS DDM Ops.
 - ◆ CERN-PROD: CERN-PROD_TMPDATADISK failed to contact on remote SRM GGUS:71164. GGUS solved (2nd Jun): The SRM server for EOSATLAS locked up at 4:00 this morning (cron.daily time). Unfortunately monitoring was still working to the point that the machine didn't get immediately rebooted, this only happened around 9:20. From the plots it looks like things were OK afterwards. We don't really know why SRM decided to act like this (logs are incomplete), so this could happen again (still, we've changed some settings for updatedb and log rotation).
 - ◆ CERN-PROD: Connection refused to CERN-PROD_DATATAPE GGUS:71216, evaluated

as problem with data import into EOS (4th June). Under investigation.

- ◆ CERN-PROD: destination file failed on the SRM with error [SRM_ABORTED] at CERN-PROD_DATADISK. GGUS:71226 [↗](#) (5th June) Transfers finished after several attempts, ticket closed by ATLAS shifter.
 - ◆ CERN-PROD_TMPDATADISK: GGUS:71227 [↗](#) Unable to connect to gdss352.gridpp.rl.ac.uk (RAL). Its a problem with CERN-PROD_TMPDATADISK with a mis-leading error. [Tiju: that machine was unavailable for just one night.]
 - ◆ [Jarka: ADCR database was down today, should be ok now, any comment?]
 - ◆ [Michal: problems with dashboard tests at ASGC. Felix: following up.]
 - ◆ [Andrea: high load on kerberos service at CERN due to access from POOL tools by some ATLAS users, under investigation. This could be due to a bug in xrootd, bug #82793 [↗](#). A ROOT patch is available and will be tested to see if it solves the issue.]
- CMS reports -
 - ◆ LHC / CMS detector
 - ◇ Good Weekend for the LHC
 - ◆ CERN / central services
 - ◇ NTR
 - ◆ Tier-0 / CAF
 - ◇ Problem Saturday morning with the export disk buffer in P5. Lost a network switch which put the system in an inconsistent state. Stabilized and recovered, but working to make the impacted data useful
 - ◇ Backlog in copying files from Point5 to CASTOR. GGUS:71047 [↗](#). CMS has combined 2 Streams no backlog reported over the weekend
 - ◆ Tier-1
 - ◇ MC re-reconstruction ongoing. Restarting large scale simulation production with pile-up.
 - ◆ Tier-2
 - ◇ MC production and analysis in progress
 - ◆ AOB :
 - ◇ New CRC Ian Fisk until June 14
 - ALICE reports [↗](#) -
 - ◆ General information: Production went well during the long weekend ~30k jobs running. Pass-0 and Pass-1 reconstruction ongoing and a couple of MonteCarlo cycles. No issues to report.
 - ◆ T0 site: ntr
 - ◆ T1 sites: ntr
 - ◆ T2 sites: usual operations
 - LHCb
 - ◆ Experiment activities:
 - ◇ Taken in 24 hours 23pb-1.
 - ◇ Overlap of many activities going on concurrently in the system.
 - ◇ Reprocessing (Reconstruction+ Stripping + Merging) of all data before May 's TS (about of ~80 pb-1 data) using new application software and conditions (~80% completed)
 - ◇ First pass reconstruction of new data + Stripping + Merging
 - ◇ Tails of old reconstruction productions using previous version of the application (Reco09) almost completed.
 - ◇ MC productions at various non T1's centers.
 - ◆ New GGUS (or RT) tickets:
 - ◇ T0: 0
 - ◇ T1: 0

◇ T2: 0

◆ Issues at the sites and services:

◇ T0

· CERN the additional disk servers in the LHCb-Tape token boosted the stripping jobs now running smoothly at CERN

◇ T1

· SARA still a huge backlog of stripping jobs due to low rate of submission of pilot jobs. Decided to move to direct submission by passing the gLite WMS and the ranking expression.
· SARA reported some issue in staging files, most likely some tape driver problem. Local contact person in touch with Ron and Co.
· RAL: backlog drained by moving to direct submission. Many jobs exiting immediately because data was garbage collected. Asked RAL folks to use the policy: LRU (Last Recently Used). [Tiju: following up with local experts about the policy.]

◇ T2: ntr

Sites / Services round table:

- RAL: nta
- NL-T1: issue with tapes this weekend as reported by LHCb, fixed this morning. Things are now getting back to normal.
- OSG: ntr
- BNL: ntr
- Fermilab: ntr
- PIC: ntr
- CNAF: ntr
- NDGF: ntr
- KIT: reminder downtime tomorrow, 9-930 dcache instance at Gridka
- IN2P3: ntr
- ASGC:
 - ◆ SRM tests were failing on sunday, then fixed in the evening.
 - ◆ CE tests (LCG and Cream) are failing, will follow up.
- GridPP: has CERN or any T1 enabled the srm splitting GUC_SRMGRIDFTPSPLIT variable; and if so have they had any issues with it; especially with respect to timeouts? This variable is shown https://twiki.cern.ch/twiki/bin/view/EGEE/FtsRelease22#SRM_gridftp_split and https://svnweb.cern.ch/trac/glitefts/wiki/FTSYaimReference_2_2_4 and the variable "--preparing-ratio" in the command `gLite-transfer-channel-set`, <http://egee-jra1-data.web.cern.ch/egee-jra1-data/fts-doc/2.2/cli/glite-transfer-channel-set.1.html>.
- Dashboard services: migration of ATLAS DDM dashboard to new HW this week (3 days), should be transparent.

AOB: none

Tuesday:

Attendance: local(Michal, Lukasz, Maria, Jamie, MariaDZ, Oleg, Dirk, Massimo, Alessandro, Peter, Maarten, German, Eva, Ian);remote("can't login to alcatel").

Experiments round table:

- ATLAS reports - ATLR dashboard shows many blocking sessions for application: ATLAS_COOLOFL_SCT_W@AtlCoolCopy.exe@pc-sct-www01 [Eva - this was caused by

Tuesday:

ATLAS offline sessions which were blocking themselves. DBAs killed sessions and now ok. Ale - should ATLAS do it by themselves next time? Eva - you must contact Marcin]

- Ale - Kerberos denial of service attack - some updates in Savannah - some users not sure what do do. [Tim = we are seeing about 3 - 10 users / day launching a DOS attack on CERN Kerberos server - around 600-700 requests / second. Having to manually follow and kill the sessions. Can protect service but this requires disabling user for a certain period of time which is clearly undesirable. Root cause - not clear on. From service perspective we are just seeing effects. Dirk - some analysis between xroot security plug in but at the moment no concrete statement. A commit to the code quite some time ago which took out an infinite loop. in xroot security plug in. May not be "new" code. IT/DSS and ROOT team involved in looking for problem. Ale - ATLAS trying to understand use cases. For now appear to be all legitimate. bug:82790
- CMS reports -
 - ◆ LHC / CMS detector
 - ◇ Continued Running
 - ◆ CERN / central services
 - ◇ NTR
 - ◆ Tier-0 / CAF
 - ◇ Caught up and back to normal running
 - ◆ Tier-1
 - ◇ MC re-reconstruction ongoing.
 - ◇ Lower quality transfers CERN to FNAL. Still succeeding, but some recurring failures
 - ◇ ASGC SRM issue that was resolved
 - ◆ Tier-2
 - ◇ MC production and analysis in progress
- ALICE reports [↗](#) -
 - ◆ T0 site - Nothing to report
 - ◆ T1 sites - Nothing to report
 - ◆ T2 sites - Poznan: delegation to CREAM-CE not working. GGUS:71261 [↗](#)
- LHCb reports - Main activities: stripping of reprocessed data + reconstruction and stripping of new taken data. MC activities.
 - ◆ T0
 - ◇ CERN Tape system with a huge backlog as the reason of no migration taking place since yesterday afternoon (GGUS:71268 [↗](#)) [German - actually two problems. A misconfig in migration policies unveiled with broken tape which triggered a problem which caused some of LHCb migrations to be stuck from 09:00. Other aspect: upgrade / reconfig of tape servers which caused snowball effect; running short of tape drives. Had to reduce some internal activity such as repack. Tape drives now 90% available to experiments and problem with migration policy now understood; trying to fix.
 - ◆ T1
 - ◇ SARA: moved to direct cream submission. Part of the waiting jobs redirected to NIKHEF. Still backlog to be drained. Fixed the problem of staging files.
 - ◇ RAL: moved to LRU policy for GC

Sites / Services round table:

- BNL - ntr
- FNAL - ntr
- ASGC - our SAM tests should pass today. Guess because our CRL expired. We fix the cron table today I think should be fine now thanks.
- IN2P3 - ntr

- RAL - ntr
- NL-T1 - ntr
- CNAF - ntr
- OSG - ntr
- NDGF - ntr

- KIT (by email)
 - ◆ 1 host with 5 ATLAS disk-only pools went down yesterday around 17:45 local time. Unfortunately, this was not noticed before today 08:00, because the issues were not severe enough to trigger on-call duty during the night. After the host was rebooted, the pools came back online, too.
 - ◆ The downtime of gridka-dcache.fzk.de was successful. We changed configurations such that LHCb is not able to write outside of space tokens anymore. As a consequence, transfers that do not state a space token to be used explicitly will fail.

- CERN DB - yesterday we had problem with ATLAS ADCR offline DB. Hung around 13:30 fully recovered at 14:00. Problem started with ASM instance triggered by two disk failures in a few hours. Looking at h/w to see what is cause of disk failures as rate high compared to others. Took some extra logs to send to Oracle and continuing investigations

AOB: (MariaDZ) A presentation on the GGUS fail-safe system development will be given at tomorrow's GDB at 10:30am CEST. Agenda with EVO connection data on: <https://indico.cern.ch/events.py?tag=GDB-June11>. Development details on Savannah:113831#comment0.

- IN2P3 - yesterday someone from GridPP asked for advice on FTS config. Please can you put contact info for GridPP into minutes? Brian Davies

- FNAL - comment for KIT. FNAL networking people trying to contact KIT networking people. Please can you make sure that they are in contact on your end?

Wednesday

Attendance: local(Peter, Michal, Lukasz, Maria, Jamie, Massimo, Andrea V, Luca, Alessandro, Ian);remote(Michael, Xavier, Jon, Onno, Ulf, John, Rolf, Jhen-Wei, Foued, Rob).

Experiments round table:

- ATLAS reports -
 - ◆ Central Services:
 - ◇ No US sites in GGUS TEAM ticket form, GGUS:71348
 - ◆ T1:
 - ◇ PIC LFC/FTS scheduled downtime, cloud offline in ATLAS

- CMS reports - nothing significant to report, all running smoothly. Recovered fine from FNAL downtime. Downtime calendar red from GOCDB -> Google calendar. CERN downtime until 26 May 2026. Must be a problem somewhere... Please fix to avoid confusing shifters.

- ALICE reports -

- LHCb reports - Experiment activities:Data processing/reprocessing activities proceeding almost smoothly everywhere with failure rate less than 10% mostly for input data resolution
 - ◆ T0
 - ◇ NTR
 - ◆ T1

- ◇ NL-T1: rolling fashion intervention to update CVMFS brought a reduced capacity with number of running jobs falling to 500. Now improving. There is a backlog there for stripping not fully recovered.
- ◇ PIC: In DT today. They got over last days a bit too much jobs compared the real capacity of the site (they were well behaving in the past). For this reason a backlog of activities is formed over there too.

Sites / Services round table:

- BNL - ntr
- KIT - this morning we had some problems with shared s/w area - now problem has been fixed and shared s/w area and home directories available again
- FNAL - yesterday was hottest day on record since 2006. Had a 2h cooling outage which affected jobs but all restarted automatically at the end. Expecting even a hotter day today!
- PIC - reminder for PIC is not just LFC and FTS but a global downtime. Also remind experiment decommissioning LCG CE. From now on will be converted to CREAM.
- NL-T1 - ntr
- NDGF - ntr
- RAL - ntr
- IN2P3 - ntr
- ASGC - downtime for network construction this Friday. 06:00 - 09:00 UTC
- CNAF - ntr
- OSG - ntr

- CERN - AT RISK for site - intervention on CASTOR public and other instances foreseen next week.

AOB:

Thursday

Attendance: local(Ian, Jamie, Maria, Peter, Gavin, Maarten, Michal, Lukasz, Alessandro, Andrea V, Jan);remote(Xavier Espinal/PIC, Michael, Roberto, Jon, Xavier Mol/KIT, Ulf, Jhen-Wei, Rolf, Lorenzo, Gareth, Jeremy, Eva, Rob).

Experiments round table:

- ATLAS reports -
 - ◆ T1:
 - ◇ PIC GGUS:71365 cannot contact srmatlas.pic.es, OPN ACLs on pic router fixed. Transfers managed by IN2P3, SARA and NDGF were OK which indicates they are not going via OPN. Tickets to check OPN GGUS:71389 GGUS:71390 GGUS:71391
 - ◇ Following ACL fix, PIC FTS transfers stuck, quickly fixed by agent restart GGUS:71384

- CMS reports -
 - ◆ LHC / CMS detector
 - ◇ Continued Running
 - ◆ CERN / central services
 - ◇ NTR
 - ◆ Tier-0 / CAF
 - ◇ Tier-0 running with Prompt Reco from a few days ago. High utilization. New version of CMSSW will be deployed after the holiday weekend and may improve the memory situation and allow us to increase the number of cores.
 - ◆ Tier-1

- ◊ MC re-reconstruction ongoing.
- ◆ Tier-2
 - ◊ MC production and analysis in progress
- ◆ gLite UI: recommending sites not to upgrade to latest UI until there can be a fix to CRAB to address a backward incompatibility
- ALICE reports [↗](#) -
 - ◆ T1 sites - CNAF: investigating the reason why there are so many EXPIRED jobs at the site
- LHCb reports - Experiment activities: Took some data last night (not so much). flooded in the last week GridKA whose share has been increased too much and almost all RAW data ended up there leaving other centers (RAL and CNAF) almost idle.
 - ◆ T0
 - ◊ Changed one field in the back-end of LFC (hostSE) [new NL-T2 SE]
 - ◆ T1
 - ◊ PIC: Back from the DT yesterday a problem with SRM has been spotted. Now fixed. Open a GGUS ticket for a CE that has been de-commissioned after the DT (GGUS:71383) [Xavier - we were very surprised about this ticket as it had been announced for 1 week and was reported in the minutes. Roberto - can just apologies for this : our shifter made a mistake.]

Sites / Services round table:

- pic: OPN transfers failure 20:00_08062011 to 8:30_09062011:
 - ◆ During the Scheduled downtime the SRM headnode was migrated to a new hardware with a new IP. The firewall rules were updated but there was an ACL that was not set for this new IP at router level, preventing OPN traffic to flow. The Manager on Duty saw traffic during the night : T2s, some T1s (IN2P3, SARA and NDGF) and functionality tests (lcg-cp's, etc.) and thought the problem was related to a DNS freshness issue. The real problem was not correlated until this morning when the network expert "discover" the missing ACL at router level for the OPN VLAN.
 - ◆ Related actions: GGUS ticket opened to investigate non-OPN route being used from NDGF, IN2P3 and SARA to PIC.
 - ◊ Ticket to IN2P3: https://ggus.eu/ws/ticket_info.php?ticket=71389 [↗](#)
 - ◊ Ticket to NDGF: https://ggus.eu/ws/ticket_info.php?ticket=71390 [↗](#)
 - ◊ Ticket to SARA: https://ggus.eu/ws/ticket_info.php?ticket=71391 [↗](#)
 - ◆ Alessandro - also included Edoardo Martelli into the ticket.
- BNL - ntr
- FNAL - ntr; however still waiting for KIT networking experts to contact FNAL networking experts. Still trying to understand some decisions that don't make any sense to us.
- KIT - Xavier, already triggered network experts. No issues to report. Downtime for dCache CMS-SRM next Tuesday 09:00 - 09:30 to update specific dCache services to fix some bugs
- NDGF - we will look into above ticket, normal that not all traffic goes over OPN.
- ASGC - ntr
- IN2P3 - ntr; note the problem with network and will contact network people
- CNAF - ntr
- RAL - ntr
- NL-T1 - ntr
- GridPP - ntr
- OSG - ntr
- CERN KDC: During the last couple of weeks high load on the Kerberos (KDC) service at CERN has been observed originating from batch jobs run by some ATLAS users. The problem was traced back

to a large number of 'kinit' calls issued via the Xrootd Krb5 authentication plug-in. The plug-in was issuing these calls in the attempt of reinitializing the credentials, after a failure reported by the server. The problem was already fixed in Xrootd, but the the fix was not included in the version used by the affected ATLAS users.

All the ROOT versions prior to 5.28/00a are potentially affected by the problem. This includes the version 5.26/00d used by ATLAS and the version 5.27/06d used by CMS. We (ROOT) are back-porting to the relevant branches (5-26-00-patches - done; 5-27-06-patches - to be done later today). New tags on these branches are foreseen to ease the deployment. [Ale - other sites such as IN2P3 might also see this problem. Maarten - and the other VOs should take note. Andrea - this will be discussed at the AF next week.]

- CERN storage: at risk next Tuesday 09:00 - 12:00 whilst moving off SLC4 h/w; Wednesday downtime on CERNT3 instance, 09:30 - 11:30.

AOB:

Friday

Attendance: local(Peter, Jan, Alessandro, Jamie, Michal, Lukasz, Maarten);remote(Michael, Xavier, Jon Felix, John, Rolf, Rob, Onno, DImitri, Lorenzo).

Experiments round table:

- ATLAS reports -
 - ◆ T1
 - ◇ TW scheduled downtime ended, cloud now online and things look OK
 - ◇ Checking OPN at IN2P3 GGUS:71389 and NDGF GGUS:71390, SARA confirmed OK. Comments from OPN guys [Xavier - pool should be inside OPN but storage maybe outside OPN. Peter - latest in ticket was traceroute between gridftp doors which goes over OPN. Rolf - heard same thing from expert that control data flow might well be outside OPN but data transfer goes through OPN]
- CMS reports -
 - ◆ LHC / CMS detector
 - ◇ Continued Running
 - ◆ CERN / central services
 - ◇ NTR
 - ◆ Tier-0 / CAF
 - ◇ Depending on the machine performance a high load is expected over the holiday weekend
 - ◆ Tier-1
 - ◇ Ticket open for transfers KIT to FNAL and CCIN2P3 to FNAL. Both of these appear to be related to files with checksum issues. Site contacts are following up.
 - ◆ Tier-2
 - ◇ MC production and analysis in progress
- ALICE reports [↗](#) - General information: CPU ratio is low grid-wide for the ongoing MC production, under investigation
 - ◆ T0 - Nothing to report
 - ◆ T1 - 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
 - ◆ T2 sites - Usual operations

- LHCb reports - Experiment activities: Taken last night 11pb-1 data everything well transmitted to CASTOR and final destination with no problems. Expected a quite hot w/e in terms of new physics.
 - ◆ T0
 - ◇ Observed problems in some lxplus node with CVMFS. Stefan reports a too aggressive caching mechanism with the current CVMFS (0.68) preventing to pick up new software. [Maarten - is there a work-around? Roberto - a new release that fixes this problem. Steve will deploy but don't know when. Maarten - don't see at all sites as they are not running the same version? Roberto - yes, "guess" that it is due to this version as its the same one at CERN and PIC. Xavier - during downtime upgraded to 0.68. Noticed on Wednesday that there is a 0.70 release. We did not install as didn't have time to test. Good to have some news on this.]
 - ◆ T1
 - ◇ GRIDKA: Spikes of analysis jobs failing accessing data (GGUS:71412). Problem with dcap ports restarted over night. [Dimitri - see ticket was updated by Doris that problem was fixed. Please can you also update ticket so that it can be closed?]
 - ◇ PIC: Received a GGUS ticket (GGUS:71399) about "abusive use of the scratch area" in the WN. It boiled down to a problem with CVMFS there (instead LHCb fault): the S/W is not picked up by the CVMFS clients but installed locally on the WN (for each job ~10GB).

Sites / Services round table:

- BNL - ntr
- PIC - nta
- FNAL - ntr
- ASGC - ntr
- RAL - ntr
- IN2P3 - nta
- NL-T1 - ntr
- CNAF - ntr
- KIT - ntr
- OSG - ntr

AOB:

- No call on Monday - CERN will be closed.

-- JamieShiers - 31-May-2011

This topic: LCG > WLCGDailyMeetingsWeek110606

Topic revision: r13 - 2011-06-10 - JamieShiers



Copyright &© 2008-2021 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

or Ideas, requests, problems regarding TWiki? use Discourse or Send feedback