

Table of Contents

Week of 100726.....	1
Daily WLCG Operations Call details.....	1
WLCG Service Incidents, Interventions and Availability, Change / Risk Assessments.....	1
General Information.....	1
Monday:.....	1
Tuesday:.....	4
Wednesday.....	7
Thursday.....	9
Friday.....	11

Week of 100726

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(Jamie, Jean-Philippe, Stephen, Jan, Patricia, Edward, Douglas, Jacek, Stephane, Maarten, Ulrich, Alessandro, Roberto, Dirk);remote(Jon/FNAL, Michael/BNL, Gonzalo/PIC, Kyle/OSG, John/RAL, Alessander/NLT1, Tore/NDGF, Andreas/KIT, Rolf/IN2P3, Paolo/CNAF, Vladimir/LHCb, Gang/ASGC).

Experiments round table:

- ATLAS reports -
 - ◆ T1s:
 - ◇ INFN-T1 -> INFN-ROMA1 file transfer failures, FTS checksum mismatch (FTS compares the checksum of the input storage with the one from DDM,adler32). GGUS:60432. Under investigation since Friday evening.
 - ◇ INFN-T1_DATADISK FT failures: GRIDFTP_ERROR/server err.500/connection timeout.GGUS 60460 assigned at ~3am on Monday.
 - ◇ FZK<-NDGF file transfer failures due to gridftp errors on Saturday. GGUS:60437, more failures on Sunday, ticket updated. Elog 15046.
 - ◇ IN2P3<-BEIJING still some file transfer failures for large (>1GB) files. GGUS:59966, Elog.15062
 - ◇ BNL some range of WNs running analysis jobs had the name resolution problem (DNS), so a fraction of analysis jobs failed. GGUS:60540 verified. New worker nodes, attached to a new subnet, were added to the analysis queues. The new subnet

needed to be added to the DNS access control list to enable external name resolution.
Fixed at ~10:30 on Sunday.

- ◆ T2s:
 - ◇ SLACXRD transfer failures due to connection to SRM problems reported on Saturday at 1am. GGUS:60436
 - ◇ MWT2/IU transfer failures because of SE problems at MWT2/Indiana University reported on Sunday at ~2:30pm (Elog15083). Fixed by admins at ~6pm (Elog 15084).
 - ◇ PRAGUELCG2 200 file transfer failures due to expired certificate at 1pm on Sunday. GGUS:60455 verified at ~3pm: New certificate has been put in place and services rfiod and dpm-gsiftpd has been restarted.
- CMS reports -
 - ◆ Tier-0
 - ◇ Normal operation.
 - ◇ Wrong version of Transfer System was running at Point5 so processing was delayed for some data by around 12 hours
 - ◇ Data was transferred to CASTOR but not injected to Tier-0 for processing
 - ◇ Tomorrow (still I think) may run for an hour without zero suppression (HeavyIon tests) and therefore stress P5->IT link (among other things)
 - ◆ Tier-1s
 - ◇ [OPEN]Savannah #115915: T1_IT_CNAF, sometimes wrong Adler32 checksum is generated, causing failed transfers.
 - ◇ [CLOSED]Savannah #115876: JobRobot & SAM CE Errors at T1_TW_ASGC. Was CRL update problem, but some CEs reporting MaxWallClockTime incorrectly.
 - ◇ [OPEN]Savannah #115873, GGUS:60420: SAM CE prod & sft test jobs expiring for T1_DE_KIT. GlueCEPolicyMaxWallClockTime sometimes zero, similar issue to TW?
 - ◇ [CLOSED]Savannah #115745: CNAF: Transfers failing to many sites. Update process unexpectedly triggered the yaim tool using a wrong configuration file.
 - ◆ Tier-2s
 - ◇ [OPEN]Savannah #115926: T2_HU_Budapest, three files produced there also corrupted.
 - ◇ [CLOSED]Savannah #115910: T2_PT_LIP_Lisbon BDII server crashed at weekend, restarted now.
 - ◇ [OPEN]Savannah #115902: T2_RU_SINP not in BDII.
 - ◇ [OPEN]Savannah #115875: File missing at T2_US_UCSD. Subscription suspended.
 - ◇ [OPEN]Savannah #115872: Corrupted file at T2_DE_RWTH. File deleted.
 - ◇ [CLOSED]Savannah #115841: Transfers to Aachen failing. Closed, no problem at Aachen.
 - ◇ [OPEN]Savannah #115806: Transfers from T2_US_Caltech to T2_RU_JINR failing. They have PhEDEx problems.
 - ◇ [CLOSED]Savannah #115786: timeouts in transfers from T1_UK_RAL to T2_FI_HIP. FTS channel changes fixed issues.
 - ◇ No change in these tickets:
 - ◇ [OPEN]Savannah #115839: Frontier activity from Warsaw not using local cache. Was small analysis cluster, advised to use existing cache in Warsaw or setup a new one.
 - ◇ [OPEN]Savannah #115811: SAM CE errors at T2_FI_HIP. Due to only SLC4 available on certain clusters.
 - ◇ [OPEN]Savannah #115808: Analysis SAM CE error at T2_TR_METU. Reported due to a faulty disk array losing files, they are clearing up.

- ◆ Notes
 - ◇ Tier-1s asked to configure new /cms/Role=t1production role. A few sites deployed it already (PIC, ASGC, RAL, IN2P3, FNAL, CNAF, waiting for KIT [GGUS:60055](#)). In progress (RAL completed over weekend, I was mistaken that KIT had already done it).
 - ◇ AFS at CERN slow access behaviour. GGUS team ticket still open possibly to track more issues [OPEN]https://gus.fzk.de/ws/ticket_info.php?ticket=59728
 - ◇ New Read-Only software area in AFS available but some machines don't see it properly.
 - ◇ Have imposed a rate limit in CASTOR at CERN today for normal users.

- ◆ Weekly scope plan (post-ICHEP-wave)
 - ◇ Tier-0: End of technical stop and tests ongoing
 - ◇ Tier-1: running backfill if nothing else comes up / StorageConsistency checks at T1s. Sites have been contacted.
 - ◇ Tier-2: very low MC production activity

- ALICE reports -
 - ◆ GENERAL INFORMATION: Good performance of the Grid resources during the weekend. ALICE has continued the Pass1 reconstruction activities (LHC10e) at the T0, Large number of analysis trains and two new MC cycles announced on Friday. Low T0-T1 raw data transfer rates
 - ◆ T0 site
 - ◇ Number of jobs decreasing this morning. Operations applied to the level of the AliEn services at voalice11. All CREAM services nicely performing during the whole weekend. Vobox certificate will expire soon and should be renewed: it's experiment responsibility.
 - ◆ T1 sites
 - ◇ No issues to report. All SEs checked this morning showing good results
 - ◇ number of jobs at CNAF low because of GPFS upgrade
 - ◆ T2 sites
 - ◇ Madrid and Torino: CREAM services updated to the latest version and back in production
 - ◇ SEs failing at: Trujillo. Poznan, Cagliari and Clermont (local Alice contact persons contacted)

- LHCb reports -
 - ◆ Experiment activities:
 - ◇ MC Production running at full steam (27K jobs run in the last 24 hours); reprocessing of data ongoing reasonably smoothly almost everywhere. Also many user jobs running in the system that makes for this week LHCb the most active VO in terms of CPU time consumed.
 - ◆ Issues at the sites and services
 - ◇ T0 site issues:
 - none
 - ◇ T1 site issues:
 - RAL: After having increased the memory to 3GB we did not run any longer in problems there (apart few jobs running out of wall clock time)
 - SARA: After the SE has been reintegrated they started to run reconstruction smoothly during the week end
 - IN2p3 ~250 jobs stalled. Some killed by BQS (exceeding CPU or memory) another fraction is under investigation
 - CNAF: jobs failing while pilot alive: [GGUS:60445](#). Site problem, ticket updated.

- CNAF: replication problem of LFC. Excluded that this has to do with the recent migration of the 3D database. GGUS:60458 [↗](#) (should be fixed now)
- pic : many (440 jobs) failed suddenly at 1 am on Sunday. GGUS:60451 [↗](#). This was due to a fraction of the farm that had to be halted due to a cooling problem in the center.
- ◇ T2 site issues:
 - Jobs failing at UKI-LT2-UCL-HEP and DORTMUND. Shared area issue at Grif.

Sites / Services round table:

- FNAL: ntr
- BNL: ntr
- PIC: temperature alarm in one of the rooms. WNs automatically switched off (800 cores). Ok now but still investigating the root cause.
- RAL: at risk as lgcce01 being setup. Will also be at risk (very low) on Wednesday because of security upgrade on DB servers.
- NLT1:
 - ◆ creamce failed over the weekend and was restarted on Sunday
 - ◆ SE disk server being worked on
- NDGF: ntr
- KIT:
 - ◆ failing transfer problems fixed on Friday (DNS issue)
 - ◆ CMS job submission problem (GGUS 64020) solved (caching problem)
 - ◆ problem with one of the tape libraries: writing was ok but reading could have been problematic.
- IN2P3: transfer problems with Beijing seem to be network related. GEANT involved.
- CNAF: STORM checksum problem: developers are aware of the problem and working on it.
- OSG:
 - ◆ BNL was sending outdated data to OSG BDII from Friday 19:30 EDT to Saturday 16:30 EDT. We are still investigating the cause. (<https://ticket.grid.iu.edu/goc/viewer?id=8963> [↗](#))
 - ◆ SMS on ATLAS ALARM did not get received by Michael Ernst. We will retest this ALARM tomorrow. (https://gus.fzk.de/ws/ticket_info.php?ticket=60427 [↗](#))
- ASGC: ntr
- CERN central services:
 - ◆ DB replication problems to T1s
 - ◇ Triumph since Friday because of an upgrade at Triumph
 - ◇ CNAF: fixed this morning around 11:00
 - ◇ IN2P3?
 - ◆ SRM upgrade to 2.9.4 (at low risk)
 - ◇ castorpublic tomorrow
 - ◇ endpoints for experiments on Wednesday and Thursday (waiting for expt agreement)

AOB:

Tuesday:

Attendance: local(Stephen, Jean-Philippe, Maarten, Lola, Patricia, Jamie, Ulrich, Roberto, Jacek, Douglas, Alessandro, Edward, Jan, Dirk);remote(Michael/BNL, Jon/FNAL, Stephane/Atlas, Gonzalo/PIC, Rolf/IN2P3, Gang/ASGC, Tore/NDGF, Ronald/NLT1, John/RAL, Paolo/CNAF, Rob/OSG, Xavier/KIT, Vladimir/LHCb).

Experiments round table:

Tuesday:

- ATLAS reports -
 - ◆ Tier-0 :
 - ◇ Tests for database scaling issues performed last night. Also caused a request of 3TB of data from castor, and this caused some short term problems (small files). Cleared normally within an hour. No problems found in the database. No ticket.
 - ◆ Tier-1 :
 - ◇ Transfer failures from FZK to NDGF. These are only for a small percentage of transfers, but appear to be connection timeout issues at NDGF. This issue has existed without change for a few days now. Two GGUS tickets on this now: GGUS:60582 and GGUS:60437. May be an network issue, waiting for the expert to come back from holidays.
 - ◇ INFN-T1 issues with failures due to checksum errors. Seems to come from StoRM setup, which uses the gridftp dsi plugin to calc. checksums. Bugs were reported, new code hopefully fixed this, and it was installed today. Hopefully this is now fixed. GGUS:60432. Could this bug affect CASTOR as well? If yes, a fix will be deployed at the next technical stop.
 - ◇ SARA outage yesterday, service upgrades not correctly reported to GOCDB. This was picked up by Atlas utils, and correctly turned off transfers to the site. Site saw this quickly, but took an hour to update the GOCDB correctly, and get the site back.
 - ◇ Transfers failing INFN <-> PIC yesterday, all other transfer to these sites worked fine, only this one link. PIC reported LHCOPN link was down, and was routing through NREN/GEANT. This route was dropping JumboFrames and causing transfer failures. It has been suggested within Atlas should be a back LHCOPN link to protect from failures when main link goes out. GGUS:60472
 - ◆ Tier-2 :
 - ◇ Data storage problems at Milano causing all transfers to fail. This has been an issue for a couple weeks now, but thought to be better, and site was white-listed yesterday. All transfers failed again to for that day, and it was black-listed again this morning. GGUS:60483, GGUS:60342
 - ◇ SLACXDR failures yesterday due to an SRM outage. Lasted for about 2 hours, then was fixed, no problems since then.
 - ◇ RO-07-NIPNE set offline early this morning, lots of failures there. Site reports ldap server outage that causes SRM problems, now fixed. Site is in testing with test jobs submitted at this time. GGUS:60569
 - ◇ There are currently Atlas internal discussions about the number of parallel transfers on star channels. Hiro/BNL reported that the current number of streams between Chicago and PIC is 50 (too high). Policy being discussed.
- CMS reports -
 - ◆ Tier-0
 - ◇ Normal operation.
 - ◆ Tier-1s
 - ◇ [OPEN]Savannah #115935: T1_FR_CCIN2P3 has trouble installing new releases. AFS connection timeout.
 - ◇ [CLOSED]Savannah #115873, GGUS:60420: SAM CE prod & sft test jobs expiring for T1_DE_KIT. Due to pbs-caching was deactivated causing BDII to report nonsense.
 - ◇ No change in these tickets:
 - [OPEN]Savannah #115915: T1_IT_CNAF, sometimes wrong Adler32 checksum is generated, causing failed transfers. Should be fixed now.
 - ◆ Tier-2s
 - ◇ [OPEN]Savannah #115944: JobRobot & SAM CE Errors at T2_ES_IFCA.
 - ◇ [OPEN]Savannah #115939: Five files lost at T2_US_Nebraska.
 - ◇ [OPEN]Savannah #115938: T2_BE_IIHE is blocking transfers from T2_BE_UCL.

Needs site to update their PhEDEx configuration.

- ◇ [CLOSED]Savannah #115926: T2_HU_Budapest, three files produced there also corrupted. XFS errors, files invalidated.
- ◇ [CLOSED]Savannah #115872: Corrupted file at T2_DE_RWTH. File deleted and retransferred back.
- ◇ [OPEN]Savannah #115839: Frontier activity from Warsaw not using local cache. New cache setup, need access from monitoring system allowed.
- ◇ No change in these tickets:
 - [OPEN]Savannah #115902: T2_RU_SINP not in BDII.
 - [OPEN]Savannah #115875: File missing at T2_US_UCSD. Subscription suspended.
 - [OPEN]Savannah #115811: SAM CE errors at T2_FI_HIP. Due to only SLC4 available on certain clusters.
 - [OPEN]Savannah #115808: Analysis SAM CE error at T2_TR_METU. Reported due to a faulty disk array losing files, they are clearing up.
 - [OPEN]Savannah #115806: Transfers from T2_US_Caltech to T2_RU_JINR failing. They have PhEDEx problems.

◆ Notes

- ◇ AFS at CERN slow access behaviour. GGUS team ticket still open possibly to track more issues [OPEN]https://gus.fzk.de/ws/ticket_info.php?ticket=59728
- ◇ New read-only software area deployed for Tier-0, Grid and CRAB@CAF usage at 10:20 today. New HEPIX script need to be deployed to take advantage of the new R/O area in all jobs.
- ◇ Green light to SRM upgrade at CERN this week.

• ALICE reports -

- ◆ GENERAL ACTIVITIES: common reconstruction and analysis train activities ongoing. Concerning the two MC cycles announced yesterday, the 40% of this production has been already completed. No raw data transfer activities in the last 24h
- ◆ T0 site
 - ◇ LCG-CE vobox (voalice13) back in production. both CREAM-CE and LCG-CE resources currently running.
 - ◇ Green light to perform the SRM-ALICE update to the latest stable version (2.9-4)
- ◆ T1 sites
 - ◇ No remarkable issues to report
- ◆ T2 sites
 - ◇ Clermont, Bologna-T2 and Mephi out of production today. Clermont is in downtime, Mephi is performing several operations today at the local VOBOX and Bologna-T2 is having some cooling problems today which have forced the sysadmins to switch off the nodes

• LHCb reports -

- ◆ Experiment activities:
 - ◇ Activities dominated by several MC production and user analysis. Data reprocessing is running to completion
- ◆ Issues at the sites and services
 - ◇ T0 site issues:
 - none
 - intervention on volhcb12 should not be done today
 - green light to SRM upgrade (in principle)
 - ◇ T1 site issues:
 - SARA: Problem with shared area (GGUS:60571)
 - SARA: Been informed they recovered the h/w hosting important data for user analysis; the list of files has been made again visible to users.

- IN2p3 : still some evidence of shared area problem with initialization scripts timing out
- ◇ T2 site issues:
 - none

Sites / Services round table:

- BNL: ntr
- FNAL: ntr
- PIC:
 - ◆ Should PIC be involved in the discussions about OPN (ticket 60472)? No.
 - ◆ star channels issues for PIC: no ticket, internal discussion in Atlas.
- IN2P3: ntr
- ASGC: ntr
- NDGF: ntr
- NLT1:
 - ◆ LHCb disk server back in production, but still configuration error for this node: a reboot will be necessary on Friday to fix it.
 - ◆ other disk server intervention to be scheduled
- RAL: FTS server will be down on Thursday to repair a faulty disk.
- CNAF: ntr
- KIT: ntr
- OSG:
 - ◆ alarm will be tested again later today to see if the SMS is received this time.
 - ◆ BDII problem being followed up.

AOB:

Wednesday

Attendance: local(Jean-Philippe, Douglas, Jan, Edward, Patricia, Gavin, Jamie, Maarten, Andrea, Edoardo, Flavia, Alessandro, Ian, Roberto);remote(Michael/BNL, Jon/FNAL, Gonzalo/PIC, Tore/NDGF, Paolo/CNAF, Gang/ASGC, Onno/NLT1, Tiju/RAL, Xavier/KIT, Rob/OSG, Vladimir/LHCb).

Experiments round table:

- ATLAS reports -
 - ◆ Tier-0:
 - ◇ At 10:00 today there was an SRM intervention to upgrade to version 2.9-4, that lasted half an hour. This seems to have worked well, and no issues to report.
 - ◆ Tier-1:
 - ◇ Transfer failures from FZK to NDGF continue, with new failures today. Tickets GGUS:60582 [↗](#) and GGUS:60437 [↗](#) have been updated with new info. Still waiting for the expert to come back.
 - ◆ Tier-2:
 - ◇ UNI-FREIBURG with many transfer failures due to SRM access problems. Looks to be a DNS issue, people are looking at it. GGUS:60572 [↗](#)
 - ◇ Beijing transfer issues to IN2P3 continue. This issue has been around for a couple weeks now, with the same errors. Doesn't effect all transfers, only some. GGUS ticket updated with new errors today. GGUS:59966 [↗](#)
- CMS reports -
 - ◆ Tier-0
 - ◇ Normal operation.

- ◆ Tier-1s
 - ◇ [OPEN]Savannah #115935: T1_FR_CCIN2P3 has trouble installing new releases. AFS connection timeout.
 - ◇ [OPEN]Savannah #115898: T1_ES_PIC non-custodial dataset deletion <https://savannah.cern.ch/support/index.php?115898>
 - ◇ [CLOSED]Savannah #115876: Job Robot Failures at T1_TW_ASGC - CRL Issue now fixed
 - ◇ Tickets to run a data consistency check were issued to all Tier-1s on July 19th. Completed for RAL, PIC, ASGC, Open for KIT, IN2P3, CNAF. ASGC is waiting for feedback.
- ◆ Tier-2s:
 - ◇ non custodial data loss at a Russian T2 (3 disks failing in a RAID6)
- ◆ Notes
 - ◇ AFS at CERN slow access behaviour. GGUS team ticket still open possibly to track more issues [OPEN]https://gus.fzk.de/ws/ticket_info.php?ticket=59728
 - New read-only software area deployed for Tier-0, Grid and CRAB@CAF usage at 10:20 yesterday. Moving it to default for users this week.
- ALICE reports [↗](#) -
 - ◆ Decrease on the activity of the experiment today. Alice is running out of the MC production. Basically no requirements for T0 nor the T1 sites. Small user analysis activity at the T2 sites. Activity started to ramp up just before lunch.
 - ◆ T0 site:
 - ◇ update of the SRM-ALICE to the latest stable version (2.9-4) performed this morning and finished at 10:30 with no incidents to report
- LHCb reports -
 - ◆ Experiment activities:
 - ◇ Many MC production jobs running (15K in the last 24 h/s) and user analysis. Data reprocessing is running to completion mainly at NL-T1 site where the 90% of them is failing resolving data.
 - ◆ Issues at the sites and services
 - ◇ T0 site issues:
 - Volhcb12:run the intervention w/o major problems.
 - Tomorrow LHCb can give the green light to run the intervention on SRM.
 - ◇ T1 site issues:
 - SARA: Problem with shared area (GGUS:60571 [↗](#))
 - SARA:GGUS:60603 Reported timeouts retrieving tURL information of a set of files on M-DST space token.
 - IN2p3 : still some evidence of shared area problem with initialization scripts timing out
 - ◇ T2 site issues:
 - none

Sites / Services round table:

- BNL: ntr
- FNAL: power outage last night, but UPS worked, so no service interruption.
- PIC: ntr
- NDGF: ntr
- CNAF: ntr
- ASGC: ntr
- NLT1:
 - ◆ SARA SRM: one pool node does not copy data from disk to tape. Will be at risk tomorrow to

fix it.

- ◆ tomorrow also, a disk server will be rebooted to solve the performance problem with the SW area. This should be transparent.
- ◆ Atlas stage disk problem being investigated
- ◆ Unable to reproduce LHCb problem (GGUS 60603): may be due to the SRM restart yesterday afternoon
- RAL:
 - ◆ top level BDII upgrade and DB upgrade were done successfully
- KIT:
 - ◆ creamce2 failing: will be taken out of production for investigation
- OSG:
 - ◆ alarm test will be redone this afternoon as it failed yesterday because BNL was mis-spelled in GGUS
- Flavia: a new Frontier instance has been put in production for Atlas. It will be added to the DNS load balancing.

AOB:

- Open GGUS tickets for tomorrow's WLCG T1SCM? The list should be provided now. Be also prepared to comment on the SIRs.

Thursday

Attendance: local(Jean-Philippe, Douglas, Hurng-Chun, Patricia, Alessandro, Gavin, Edward, Lola, Maarten, Jamie, Jacek, Ulrich, Jan, Andrea, Ian);remote(Gonzalo/PIC, Jon/FNAL, Michael/BNL, Xavier/KIT, Paolo/CNAF, Gang/ASGC, John/RAL, Rob/OSG, Onno/NLT1, Vladimir/LHCb, Christian/NDGF).

Experiments round table:

- ATLAS reports -
 - ◆ Tier-0:
 - ◇ Beam came on last night, and Atlas recorded a special run of data, and created a very large back-log of files to migrate to castor. Experts were informed, and feedback from castor people is that things are working, and the backlog is because of the trigger settings from Atlas creating a large number of files. Hopefully this will be different when they change the settings for physics beam today. Tier-0 support says the file backlog can continue for a few days and be fine, so we will watch and see.
 - ◆ Tier-1:
 - ◇ SARA had an SRM outage starting at 5:30am today CERN time, and caused transfer failures for anything involved with site. A ticket was created, but no response. An alarm ticket was created at 8:30am, and got a response within the hour. A disk partition was full and blocking the SRM service, this was cleaned out and service came back immediately. Two hours later there were more SRM outage failures, but this was short term and not reported. GGUS:60637 [↗](#) GGUS:60642 [↗](#)
 - ◇ NDGF problems with UNICPH-NBI in downtime. This was not reported to GOCDB, because I was told site is not in GOCDB. Can this be checked by NDGF people and fixed if true, so we can correctly get site downtime reports? GGUS:60683 [↗](#) Savannah #69843, Site is currently blacklisted.
 - ◇ Alarm ticket testing to BNL finally worked, but alarm did not ring Michael's phone. Brings up issue with GGUS alarms, it took me TWO DAYS to issue this alarm ticket. How often are GGUS alarms tested? (once a month after each GGUS release). Can this be more often to insure that alarm ticket will work when needed? GGUS:60625 [↗](#)
 - ◆ Tier-2:

- ◇ More transfer failures to IN2P3 from Beijing. Ticket now open for over two weeks, and updated almost daily. Lots of info to fix the problem, does the site need more? GGUS:59966 [↗](#)
- ◇ UKI-LT2-RHUL problems with SRM failures. Site blacklisted for now. GGUS:60631 [↗](#)
- ◇ Milano-INFN tested site, and asked to be whitelisted. All transfer to site started to fail right away. Reported new failures, and site is back in blacklist. GGUS:60483 [↗](#)
- CMS reports -
 - ◆ Tier-0
 - ◇ Normal operation.
 - ◆ Tier-1s
 - ◇ [OPEN]Savannah #115935: T1_FR_CCIN2P3 has trouble installing new releases. AFS connection timeout. Progress today and waiting on confirmation to close the ticket.
 - ◇ [OPEN]Savannah #115898: T1_ES_PIC non-custodial dataset deletion <https://savannah.cern.ch/support/index.php?115898> [↗](#)
 - ◇ [CLOSED]Savannah #115876: Job Robot Failures at T1_TW_ASGC - CRL Issue now fixed
 - ◇ Tickets to run a data consistency check were issued to all Tier-1s on July 19th. Completed for RAL, PIC, Open for KIT, IN2P3, ASGC, CNAF. Data Ops sent a reminder to indicate what to do with orphaned files.
 - ◆ Notes
 - ◇ AFS at CERN slow access behaviour. GGUS team ticket still open possibly to track more issues [OPEN]https://gus.fzk.de/ws/ticket_info.php?ticket=59728
 - Moving it to default for users this week.
- ALICE reports [↗](#) -
 - ◆ MC production stopped, waiting for new cycles. Pass1 reconstruction activities running at CERN and basically no job requirements for the T1 sites. Low user analysis activity at the T2 sites
 - ◆ T0 site
 - ◇ GGUS:60650 [↗](#). One of the CREAM-CE: ce202 not working at submission time. Current status: SOLVED (tomcat needed to be restarted).
 - ◆ T1 sites
 - ◇ cream-2-fzk.gridka.de at FZK taken out of production
 - ◇ Alice still running a low number of jobs at CNAF. We assume it is due to the GPFS migration at CNAF. Status of the migration (?). 70 % done, but the low number of jobs is due to the fact that ALICE has used 168% of its share, so the job priority is currently low.
- LHCb reports -
 - ◆ Experiment activities:
 - ◇ Received some data last night (3 out of 7 runs sent to OFFLINE). Expected much more tonight. False alarm when data was not yet migrated to tape on CASTOR. No disturbance following the intervention on SRM .
 - ◆ Issues at the sites and services
 - ◇ T0 site issues:
 - Intervention on SRM to 2.9.4.
 - ◇ T1 site issues:
 - GridKA: GGUS:60647 [↗](#) Pilot aborting at GridKA. Problems seems related with 2 (out of three) CREAMCEs there. Ticket updated.
 - SARA: GGUS:60603 [↗](#) Reported timeouts retrieving tURL information of a set of files on M-DST space token. It seems to be a general problem on their

SE from 2am till 7am UTC and they claim the timeout on DIRAC is a bit too aggressive. We never saw this problem before. Seems to coincide with peak of staging for ATLAS. Should SARA use dedicated instances for experiments as at KIT?

- ◇ T2 site issues:
 - none

Sites / Services round table:

- PIC: next Wednesday there will be a downtime for the compute service as the network will be restructured. WAN transfers should be ok.
- FNAL: ntr
- BNL: ntr
- KIT: ntr
- CNAF: ntr
- ASGC: ntr
- RAL: lost 2 disk servers (ATLASMCDISK): they are currently with fabrics for investigation. As soon as the duration will be known a list of files will be provided.
- NLT1:
 - ◆ full partition last night. dCache was logging too much. Partition size has been increased and logging is back to normal. But needs to understand why this bad logging started.
 - ◆ normal GGUS tickets are only processed during normal working hours
 - ◆ Atlas stage problem (GGUS:60587 [↗](#)) being investigated.
- NDGF:
 - ◆ ntr
 - ◆ comment on the Tier2 status: could not change the site status in GOCDB because of a missing role. Will use an alternative way to update GOCDB. Site should be back tomorrow.
- OSG: ntr

- CERN central services: ntr

AOB:

Friday

Attendance: local(Jan, Douglas, Jamie, Edward, Patricia, Jacek, Huang, Nilo, Ulrich, Simone, IanF);remote(Michael, Stefano, Jon, Gonzalo,Xavier, Gang, Ronald, Gareth, Kyle, Vladimir, Christian/NDGF).

Experiments round table:

- ATLAS reports -
 - ◆ Tier-0:
 - ◇ Physics run from last night (finally...), and good data were processed. This caused an over-load of AFS use, but seems to be fault of internal systems System experts made some tweaks, and seems to be working fine now.
 - ◆ Tier-1:
 - ◇ SARA having problem with access to tape for certain files, causing large spikes of failures. This has been happening every 6 hours for a day and a half now. Some confusion at first if these should be ignored or not. dCache dev. have now been contacted to see if they can help. GGUS:60587 [↗](#)
 - ◇ TRIUMF was failing all jobs of a certain type. Was related to conditions pool file access, and corrupted PFC. The PFC was rebuilt, and no more reports for the day. GGUS:60675 [↗](#)

- ◊ Many FTS errors at FZK today, because of a short interruption due to a hanging GPFS server. This has now been fixed. GGUS:59968 [↗](#)
- ◆ Tier-2:
 - ◊ NTR
- ◆ Tier-3:
 - ◊ Transfers to NERSC failing in US cloud, but could not submit a ticket to them because they are not in the GGUS system. They are listed in the OIM, and it seems that GGUS hasn't updated their info. Savannah #116013 [Michael - NERSC is not even a T2 - its a T3! Douglas - told even as T3 should be able to get GGUS tickets. Michael - not a critical site. Ale - need some periodic update in GGUS of sites. Jan - several SLS reds for ATLAS T3. Looks like high read load on a small # of files - could spread out to whole pool?]
- CMS reports -
 - ◆ Tier-0
 - ◊ Normal operation. Smooth operations with 25x25 bunches
 - ◆ Tier-1s
 - ◊ [CLOSED]Savannah #115935: T1_FR_CCIN2P3 has trouble installing new releases. AFS connection timeout. Progress today and waiting on confirmation to close the ticket.
 - ◊ [OPEN]Savannah #115898: T1_ES_PIC non-custodial dataset deletion <https://savannah.cern.ch/support/index.php?115898> [↗](#)
 - ◊ [CLOSED]Savannah #115876: Job Robot Failures at T1_TW_ASGC - CRL Issue now fixed
 - ◊ Tickets to run a data consistency check were issued to all Tier-1s on July 19th. Completed for RAL, PIC, Open for KIT, IN2P3, ASGC, CNAF. Data Ops sent a reminder to indicate what to do with orphaned files.
 - ◆ Notes
 - ◊ AFS at CERN slow access behaviour. GGUS team ticket still open possibly to track more issues [OPEN]https://gus.fzk.de/ws/ticket_info.php?ticket=59728
 - Moved users to default read-only AFS yesterday
- ALICE reports [↗](#) - GENERAL INFORMATION: Central services of MonaLisa are down this morning (one of the central machines died today). ALICE experts are working on it. For the time being, we are checking manually the status of the production at the most crucial sites
 - ◆ T0 and all T1 sites manually checked: tiny production at these sites for the moment
 - ◆ Raw data transfers: small amounts of transfers currently ongoing with CNAF, NDGF and FZK
- LHCb reports - Received some data. Reconstruction, MonteCarlo and user analysis jobs. Otherwise NTR. Except SARA - problem with file access.

Sites / Services round table:

- BNL - ntr
- CNAF - ntr
- FNAL - ntr. JB will be on vacation but will be replaced.
- PIC - ntr
- KIT - ntr
- ASGC - ntr
- NL-T1 - sched restart of dCache on a pool node to correct config error.
- RAL - ntr
- NDGF - Copenhagen pools coming back online. Doug - in GOCDB now? No, not fixed now.
- OSG - ntr

- CERN - ntr

AOB:

-- JamieShiers - 22-Jul-2010

This topic: LCG > WLCGDailyMeetingsWeek100726

Topic revision: r13 - 2010-07-30 - 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](#)