

Table of Contents

Week of 100426.....	1
Daily WLCG Operations Call details.....	1
WLCG Service Incidents, Interventions and Availability, Change / Risk Assessments.....	1
General Information.....	1
Monday:.....	1
Tuesday:.....	5
Wednesday.....	7
Thursday.....	9
Friday.....	11

Week of 100426

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 & Broadcasts		Change assessments
ALICE	ATLAS	CMS	LHCb	WLCG Service Incident Reports	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(Lola, Roberto, Patricia, Jean-Philippe, Jan, Maarten, Andrea, Jamie, Simone, Alessandro, Harry, Eva, Nilo, Nicolo, Timur, Dirk, MariaDZ, Miguel);remote(Gonzalo, Gang, Jon, Michael, Gareth, Rolf, INFN-CNAF, Angela, Jens, Rob, Joel, Ron).

Experiments round table:

- ATLAS reports -
 - ◆ Yesterday evening BNL noticed a fairly large backlog of dataset transfers from the T0 (80 production datasets + 100 functional test datasets). The full exchange and investigation in <https://prod-grid-logger.cern.ch/elog/ATLAS+Computer+Operations+Logbook/11863>
 - ◇ A large number of RAW datasets had been assigned to BNL over the last days (and the largest streams)
 - ◇ BNL has 80 FTS transfers slots. The average speed for a single transfer is approx 8MB/s (= 640MB/s total). The same performance is observed in SARA-BNL.
 - ◇ T0 resources are not congested. Same for BNL resources. Other T1s are basically idle.
 - ◇ This is one case where it would be worth to increase the number of transfers slots from 80 to i.e. 160 (or 200). We decided to wait the morning for this.
 - ◇ In the morning the backlog was gone (reported by the P1 shifter at 6:30AM CEST).
 - ◆ TRIUMF understood the problem of source files being deleted after failed transfer.
 - ◇ Explanation from Reda: "After some investigations/debugging on FTS side and dCache/SRM side, it looks like more of a dCache bug during the SRM transaction

when FTS releases the PrepareToGet request in the specific case where the transfer fails at the destination. The source file gets deleted at that very moment. We are going to upgrade to the latest dCache patch level sometime during next week's LHC technical stop. It will be broadcasted via the usual channels." [Very rare occurrence - verbosity of logging in dCache was not enough. Not easily reproducible...]

◇ TRIUMF will be upgrading tomorrow: "TRIUMF-LCG2 has a scheduled downtime planned for 2010-04-26 starting at 17:00 UTC for 2 hours. This is to perform a patch-level upgrade of the dCache to 1.9.5-17 (from 1.9.5-11).

<https://goc.gridops.org/downtime/list?id=75055437>"

◇ Not clear to me why this problem is only in TRIUMF. Is it the only site running this version of dCache? What about in the past?

◆ We agreed with CASTOR ops to re-enable the checksums in ATLASSCRATCHDISK today and test it.

• CMS reports -

◆ T0 Highlights:

◇ Weekend: data taking with squeezed beams.

◇ New data injection to PhEDEx stopped for 9 hours Sunday 01:00-10:00 for issue with PhEDEx Dataservice - no transfers; backlog quickly recovered on Sunday afternoon.

◇ FILE_EXISTS error in imports T1s-->T2_CH_CAF for 0-size files left by previously failed transfer attempts, cleaned up manually, will implement automatic cleanup.

◇ No issues to report so far following update of srm-cms.cern.ch to SRM 2.9 on Monday morning.

◆ T1 Highlights:

◇ Spring10 Monte Carlo processing running at T1s

◆ T2 highlights

◇ MC production running

1. T2_AT_Vienna had 0 jobs running - fixed.

2. Production at T2_EE_Estonia on hold while site migrates to new storage.

◇ T2_PT_NCG_Lisbon - SAM CE and JobRobot errors.

◇ T2_PT_LIP_Lisbon - SAM SRM errors.

◇ T2_RU_RRC_KI - SAM CE errors.

◇ T2_IT_Bari - SAM CE & SRM errors.

◇ T2_FI_HIP - Squid down, restarted.

• ALICE reports -

◆ T0 site

◇ CASTORALICE update scheduled for today has been completed and announced by the experts. The new version includes the deployment of the latest set of 2.1.9 patches for CASTOR (including xroot). More details available at:

<https://twiki.cern.ch/twiki/bin/view/CASTORService/ChangesCASTORALICEUpdate2195>

◇ AFS issue announced on Friday:

· Friday operation: We changed the AFS volume to exclude any possible volume corruption in the current one (this new volume was still readable/writable)

· Saturday noon: The same problem appear again with the new volume

· Weekend production: Not interruption of the Pass1 reconstruction tasks due to the tiny usage of the afs area during this weekend

· Monday morning: The problem has been understood by the experts and in order to avoid a similar situation in the future, after a meeting with the AFS experts, ALICE has agreed to go for the volume replication approach. A synchronization tool has been required to the experts to ensure a similar information in the new (and separate) readable and writable volumes that ALICE will use in the future.

- ◆ T1 sites
 - ◇ CNAF T1: On Friday night the 2nd VOBOX of the experiment was not reachable. Problem reported to the ALICE experts and the site and immediately solved. We have performed a change on the VOBOX setup at the site in order to ensure that each vobox points to a separate CREAM-CE
- ◆ T2 sites
 - ◇ Bratislava has announced during the weekend the setup of a CREAM-CE. Service tests and configuration of the site foreseen for today
- LHCb reports -
 - ◆ During the week end the increased luminosity brought up 1200 RAW files that increased a bit the load on CASTORLHCB (we took 3 times more data than integrated before). The increased production activities at the T0 range from data from the PIT, to data reconstruction (75% running at CERN) and data export to T1s. In the last week a large fraction (80%) of jobs were stalling at all sites (CERN mainly) because running out of Max CPU time in the remote queues. Asked ONLINE people to decrease from 2 to 1GB the RAW file size which is the current one. The problem is in the workflow that runs twice the same steps due to the low rejection factor with the current L0 trigger.
 - ◆ With the next release of GAUDI there will be in place a way to intercept signals from the batch system to end gracefully the jobs when a SIGTERM signal is trapped (before the fatal SIGKILL is sent by the LRMS). On going a discussion with IN2p3 whose BQS sends in advance a SIGUSR instead of SIGTERM as all other batch systems do.
 - ◆ Issues at the sites and services
 - ◇ T0 site issues:
 - Intervention on SRM to upgrade to 2.9 today 11-11:30. Preliminary tests on the PPS instance confirmed that the migration to this new version is not a problem.
 - Some slowness reported during the weekend by a user. Offline discussions seem to point the problem with the user itself that was running hundreds parallel rfio requests on the same restricted set of data in the RAW pool. The rest of users were not affected.
 - ◇ T1 sites issues:
 - CNAF: CREAM endpoint failing all pilots [Bad domain name - reserve hostname didn't match. Should be (almost) fixed now..
 - SARA: CE still banned in the production mask: to all jobs stalling due to the latency in contacting CERN ConditionDB. The problem will be fixed once the new GAUDI release with the new Persistency patch working will be in place in few days time.
 - IN2p3: SIGUSR1 signal instead of SIGTERM signal.
 - RAL: failing all jobs for a recent reconstruction production. Under investigation
 - KIT - alarm ticket against KIT
 - ◇ T2 sites issues
 - UNINA-EGEE and GRISU-UNINA: shared area issue
 - EFDA-JET: no space left.

Sites / Services round table:

- FNAL: We are rebooting our FTS server today to apply a security patch. Question: We tried FTS load-balancing about a year and it didn't work for us. Does anyone use it and does it work? [Follow-up offline]
- PIC - currently draining PIC as scheduled downtime tomorrow. ATLAS workload over w/e - all went fine but substantially higher load on shared area. Are there clear differences with this reprocessing that would explain this? Simone - will follow up

- ASGC - ntr
- BNL - to add to ATLAS' first point: network CERN-BNL - both links providing 17Gbps were utilized equally at 3 - 3.5 Gbps each. Matches well bandwidth 600-700MB/s observed. Why was BNL not getting data at a higher rate? On our action list..
- CNAF - in scheduled downtime, AT RISK for SAN; all switches redundant so just risk. Downtime due to tape library upgrade - should be over end afternoon
- KIT - since 1 week new WN installation with new kernel. Problem did not reoccur. Would like to update batch server from 10:00 on Monday. ALICE was removed from dCache instance used by LHCb - will have to restart instance later this week. Not tomorrow morning according to local expert - will check. Restart < 1 minute.
- NDGF - ntr
- IN2P3 - several things: problem with CIC portal - SIR filed on missing downtime notifications. Unscheduled outage Sunday morning 00:00 - 16:00 due to SPOF in one of basic services which crashed. Not quite sure what was underlying reason - suspect Oracle problem. For moment not complete sure - will file an SIR. Effect was that no job could be started. Normally jobs which were running should have finished OK. Effect on LHC production not too significant except for missing capacity on our side.
- RAL - on Saturday problems with SAM BDII at CERN - failing SAM tests affecting us and others. On Sunday issue with ATLAS s/w server - resolved. On Wednesday scheduled outage of CASTOR and batch - in GOCDB. Maarten - investigated the SAM BDII issue. Had been going on since 04:00. Operating "at edge of stability" - SAM BDIIs still running gLite 3.1 - more susceptible to total size of BDII. Patched a few parameters manually. SAM team have scheduled a correction of site availability numbers. Jens - seen SAM test failures where error was expired certificate on CERN site. Any comments? Maarten - open a ticket for SAM SFT.
- OSG - another top priority ticket opened falsely on Saturday afternoon. Do need to talk to users about what GGUS form priority scheme is - maybe switch some of priority mappings. This starts a chain of communication that will wake people up or get them in on w/e. This a non-critical issue. Simone - a weekly DAST meeting : put on agenda. Top priority is bottom line - maybe set by error. Maria - request also from CERN, discussed at USAG. Will put minutes of USAG in minutes of today. Changed help page. Here are the details as promised by MariaDZ: - Issue discussed at 20100121 USAG. Please see agenda <http://indico.cern.ch/conferenceDisplay.py?confId=81363> item 6, and read the point in the relevant point in the minutes. The issue was analysed according to savannah <https://savannah.cern.ch/support/?112031#comment3> which contains the text we included in the Feb 2010 GGUS Release Notes and the GGUS submit form help text https://gus.fzk.de/pages/help/help_sub_prio.php . This is the ticket submit form <https://gus.fzk.de/pages/ticket.php> . 'less urgent' is the default value, what is the point of the 'bottom line' comment above? OSG was advised at the meeting not to map tickets to 'critical' in their system, unless they are ALARM tickets.
- NL-T1: ntr
- CERN DB - intervention on ALICE online in morning (finished) to change failed controller.
- CERN SRM public will be upgraded tomorrow;

AOB:

- KIT - why did you open alarm for this space token issue? A: asked for request in 2009 and still did not get what we requested.
- Brian - saw some issues with FTS cancelling some transfers due to checksum mismatches. Can copy associated files to dest. Checksums agree with those provided by ATLAS. Issue on FTS? Any other site seeing similar issues transferring to T2s. (RAL-Cambridge and Liverpool T2s: both run DPMs. GGUS:57631).

Tuesday:

Attendance: local(Nilo, Simone, Eva, Jan, Maarten, Loca, MariaDZ, Alessandro, Ewan, Nicolo, Jamie, Harry, Flavia);remote(Michael, Gang, Jon, Gonzalo, CNAF, John, Joel, Rob, Andrea, Ronald, Rolf, Jeremy, Jens).

Experiments round table:

- ATLAS reports - Report to WLCGOperationsMeetings
 - ◆ IN2P3-CC AFS problem:
 - ◇ LFC problem related? https://gus.fzk.de/ws/ticket_info.php?ticket=57670 , atlas elog <https://prod-grid-logger.cern.ch/elog/ATLAS+Computer+Operations+Logbook/11918> [Rolf - we are investigating this issue to understand what has really happened. Many things do no work when AFS is down which might be one of reasons - cannot answer to Q about LFC relationship.]
 - ◇ IN2P3-CC disappeared from the BDII [BDII also needs AFS so not surprising that this disappeared. Failover to another site is not automatic so this explains disappearance for some time.]
 - ◆ PIC storage problem (26th @2pm) https://gus.fzk.de/ws/ticket_info.php?ticket=57662 , fixed after few minutes (thks)
 - ◆ CERN-PROD SRM source errors: https://gus.fzk.de/ws/ticket_info.php?ticket=57704
 - ◆ Interventions:
 - ◇ PIC is in sched downtime <https://goc.gridops.org/downtime/list?id=74755495> , the FTS channels are closed, no problem observed up to now
 - ◇ SRM upgrade at CERN agreed for tomorrow (the 28th) @ 11am. More details <https://prod-grid-logger.cern.ch/elog/ATLAS+Computer+Operations+Logbook/11945>
- CMS reports -
 - ◆ T0 Highlights:
 - ◇ Jobs stuck in Preparing phase on CERN-RAL FTS channel (jobs cancelled) known FTS issue, will be fixed in FTS 2.2.4
 - ◆ T1 Highlights:
 - ◇ Test jobs running at T1s
 - ◆ T2 highlights
 - ◇ MC production running
 1. T2_IT_Bari - SAM CE & SRM errors - affecting MC merge jobs.
 2. Production at T2_EE_Estonia on hold while site migrates to new storage.
 - ◇ Verifying files failing transfers at several T2s T2_IT_Rome, T2_KR_KNU, T2_US_Wisconsin, T2_UK_London_IC, T2_TR_METU
 - ◆ [Data Ops]
 - ◇ Tier-0: little activity till end of LHC stop.
 - ◇ Tier-1: little activity.
 - ◇ Tier-2: little activity.
 - ◆ [Facilities Ops]
 - ◇ This week we have an Offline and Computing Workshop. Agenda can be seen here: <http://indico.cern.ch/conferenceDisplay.py?confId=83142>.
 - NO Joint PVT/Offline/Comp/Trig Operation Meeting during Offline/Computing week (unless specifically called)
 - NO bi-weekly T2 support meeting this week. To be resumed by 13th-May.
 - ◇ Job Robot moved to CRAB 2.7.1 by 21st-April:
 - Old JobRobot datasets at the sites will be deleted soon by FacOps team.
 - ◇ T0_CH_CERN, T1_CH_CERN and T2_CH_CAF upgraded to PHEDEX_3_3_1 yesterday morning. No issues to report so far following update of srm-cms.cern.ch to SRM 2.9, which occurred as well yesterday morning.

- Encouraging sites to provide SL5 voboxes for CMS to run PhEDEx. By end of June it will be mandatory, as PhEDEx_3_4_0 will be released, sl5-only in a mandatory upgrade at all the sites. We ask the sites to provide sl5 voboxes by beginning of June.
 - ◇ CMS VObard (CIC-portal) will be upgraded this week.
- ALICE reports [↗](#) - general information: Several issues with the ALICE central services have stopped the reconstruction and analysis trains activities for some hours tonight. The system is back in production since this morning
 - ◆ T0 site
 - ◇ Replication of the ALICE afs volumes done. Implementation of the new software in alien ongoing
 - ◇ CASTOR experts announced yesterday a non-transparent operation needed for the ALICE-SRM. Together with the experiment it has been scheduled for the 28th of April at 10:00 and it will take about 30min.
 - ◆ T1 sites
 - ◇ RAL: The site has asked for a 2-3 days time window to drain the local CREAM-CE and to enable pilot accounts for the other LHC VOs. Message passed to ALICE for a final decision.
 - ◇ Rest of T1 sites performing well
 - ◆ T2 sites
 - ◇ Catania-T2: The local CREAM-CE is showing instabilities avoiding the submission of agents to the site. ALICE contact person aware of the problem and the CREAM-CE developers have also sent a possible solution for the problem
 - ◇ GGUS:57546 [↗](#) submitted several days ago concerning the bad behavior of the CREAM-CE in Bari. The ticket has not been responded by the site although it seems the system is working this morning. In the case the problem has been solved, please update the GGUS ticket.
- LHCb reports - Reconstruction ongoing with analysis activity at low level. Commissioning new workflow Brunel-Davinci that will replace the current one Brunel-Davinci-Brunel-Davinci by removing (now) redundant passages and reducing by a factor 2 the CPU requirements.
 - ◆ T0 site issues:
 - ◇ none
 - ◆ T1 site issues:
 - ◇ IN2P3: AFS outage
 - ◇ pic: downtime
 - ◇ NIKHEF: turned off the CREAM CE. Reason (quoting Jeff) : it does not have our TMPDIR patch, so all the LHCb jobs you submitted via it, were running in \$HOME which is NFS mounted. This effectively killed our entire site. I had to kill all those jobs to get things working again.
 - ◇ GridKA: need to restart PNFS on dCache. Agreed with contact person to drain currently running jobs and then restart it. The CE will be restarted whenever this problem is addressed. The lcg-CEs are not affected.

Sites / Services round table:

- BNL - ntr
- FNAL - ntr
- PIC - ntr: scheduled intervention went fine, now back up and running
- ASGC - ntr
- RAL - ntr, intervention going ahead tomorrow
- OSG - ntr
- KIT - already mentioned that we will have short intervention of 1-2' of dCache instance for LHCb

Tuesday:

tomorrow morning - other VOs not affected.

- NL-T1 - concerning CREAM CE issue last night: problem was that CREAM CE jobs were running on home dir that was NFS Mount that blew away NFS server - this affected all other VOs. Working to restore service
- IN2P3 - ntr
- CNAF - CREAM CE problem yesterday fixed, domain name misconfig, tape lib intervention ended successfully, 2h outage tomorrow morning for STORM front-end upgrade
- NDGF - ntr
- GridPP - ntr

- CERN Grid services: transparent intervention AT RISK on CERN batch services tomorrow.
- CERN storage services: SRM public upgrade this morning discovered bug for OPS VO - special config only for this VO only on this service + not up to date gLite s/w. Failing SAM tests on replication for OPS VO - acceptable or forced roll-back? Only an issue if explicit link from OPS availability to experiment availability. (lcg-rep on SLC4 WNs) Fix for code expected soon and/or update gLite on SLC4 WNs. CERN site now appears available again. (Simone - maybe someone should review OPS tests. Maarten - replication test needs to work, internal / external gridftp difference is specific to CASTOR.)

AOB:

- T1SCM on Thursday, agenda circulated. A person from LHC operations will be present to comment on likelihood of LHC technical stops being maintained / shifted. Agenda: <http://indico.cern.ch/conferenceDisplay.py?confId=92483>

Wednesday

Attendance: local(Jamie, Maria, Brian, Patricia, Simone, Edoardo, Nilo, Eva, Jan, Maarten, Ewan, Dirk, Miguel, Julia, Lola);remote(Jon, Angela, Joel, Michael, Gonzalo, Jeremy, CNAF, Rolf, Tiju, Rob, Jens).

Experiments round table:

- ATLAS reports -
 - ◆ CERN-PROD srm-atlas problems exporting data: source errors, https://gus.fzk.de/ws/ticket_info.php?ticket=57737 - this error was rather rate but now occurs ~once per week. Why do we have it now and not before? [Jan - linked to SRM ATLAS becoming unavailable due to running out of front-end threads in SRM server. Slow scheduling due to SRM not being available. If FTS retries at a bad time get this error. Should(?) be fixed as from intervention (11:20 or so) this morning.]
 - ◆ INFN-T1 problems in exporting/importing data: https://gus.fzk.de/ws/ticket_info.php?ticket=57733 (was before the downtime) [We noticed problem in STORM backend. Under investigation - as we know something more will be added to GGUS ticket.]
 - ◆ TRIUMF-LCG2 problem in exporting data due to missing files: https://gus.fzk.de/ws/ticket_info.php?ticket=57734. The problem is similar to the one reported in https://gus.fzk.de/ws/ticket_info.php?ticket=57028, the pnfs patch should fix this problem. Since 57028 is onhold, we pointed 57734 to the onhold one and we closed it. ATLAS will just need the list of files missing.
 - ◆ Weizmann transfer errors (https://gus.fzk.de/ws/ticket_info.php?ticket=57710): errors seem to be due to FTS server /var full (similar to the one of few days ago https://gus.fzk.de/ws/ticket_info.php?ticket=57619). Ticket redirected.
 - ◆ CERN-PROD yesterday morning upgrade of xrootd caused 100% failures of analysis and production jobs. https://savannah.cern.ch/bugs/?func=detailitem&item_id=66485 Problem fixed in the afternoon. (Patch cured problem)

- ◆ Today Downtime:
 - ◇ RAL-LCG2 storage and batch 11:15-15:30
<https://goc.gridops.org/downtime/list?id=74905445>
 - ◇ INFN-T1 storm upgrade 13:30-15:30 CET + DB migration (LFC down) 9:00-17:00
 CET <https://goc.gridops.org/downtime/list?id=75205438> and
<https://goc.gridops.org/downtime/list?id=74905444>
 - ◇ CERN-PROD SRM upgrade 11:00-11:05
- CMS reports -
 - ◆ T1 Highlights:
 - ◇ Spring10 MC production running.
 - ◇ Test workflows running
 - ◇ CNAF
 1. 31 Files stuck in migration to tape for a long time (up to 3 weeks).
 - ◇ PIC
 1. 2 Files with bad checksum repeatedly failing transfers to FNAL (known issue, invalidating files).
 - ◇ KIT
 1. Site contact reported that disk-only pools are filling up, possibly issue with automatic cleanup of temp files in prompt skimming workflows - CMS DataOperations following up.
- ALICE reports - GENERAL INFORMATION: Pass1 and Pass2 (cosmics) reconstruction activities at the T0 and T1 sites are going on. In addition, there are two analysis trains also in production
 - ◆ T0 site
 - ◇ SRM-ALICE update finished this morning. No issues or problems observed by ALICE after the update
 - ◇ AFS issue: Replicated volumes already in place. Implementation in aliEn finished yesterday evening and put in production for testing purposes. Still the R/W area is being used a lot (more than the replica) but this is the expected situation until the current batch jobs finish and new ones recognize the new readable volume. Hence we'll have to wait until the R/W is really "calm" before we can draw conclusions.
 - ◆ T1 sites
 - ◇ RAL T1 out of production draining the CREAM-CE to include the pilot account for the rest of LHC VOs. Site is out of production, waiting for their confirmation to put it back
 - ◆ T2 sites
 - ◇ Catania-T2: Irregular behavior of the CREAM-CE system reported yesterday to the site admin.
 - ◇ Bratislava: The sw area issue reported few days ago and observed at the local VOBOX (the sw area of ALICE was not defined inside the local VOBOX) has been solved. Making the last checks before putting the new system in production by today
 - ◇ Bari CREAM CE issue of yesterday solved.
- LHCb reports - Reconstruction ongoing with analysis activity at low level. New LHCbDirac release put in production
 - ◆ T1 site issues:
 - ◇ pic: downtime finished (still some stalled jobs to be checked)
 - ◇ NIKHEF: CREAM CE problem has been fixed. Back in production.
 - ◇ GridKA: need to restart PNFS on dCache. (on going) - no info from GridKA when things are done - can info be improved? Angela - AFAIK we told LHCb that we would start at 09:00 and have updated status page. Took a bit longer than expect but now AOK. Will try to inform next time. LHCb-grid@cernNOSPAMPLEASE.ch - also special LHCb GridKA contact.

Sites / Services round table:

- FNAL: About ~month ago, we retired 1000 job slots, reducing the available slots from 8000 to 7000 for Tier-1 production. We forgot to reduce the number of allowed running Tier-1 jobs in our condor configuration. Yesterday we had 2 dataops groups submit 12000 jobs to FNAL, and condor gave up more jobs than we had job slots. This caused lots of confusion and trouble, including not having enough job slots to run the SAM ops tests and the OSG rsv test. We corrected the condor config once we figured this out and much work did get finished yesterday.
- KIT - restart pnfs on LHCb instance as mentioned. Tomorrow one tape library will be realigned - from 13:00 - 16:00 recall of files will be stalled.
- BNL - yesterday in early pm we observed processing of data transfers slowed down. Preparetoget/put took longer than usual. Slowness in SRM server - reason lots of entries in SRM postgres db that were obsolete. Vacuum process had not removed these - had to clean DB. Minor but necessary intervention to restore performance.
- PIC - ntr
- GridPP - ntr
- CNAF - Storm upgrade in progress, ATLAS LFC in progress - new h/w for DB cluster, finished successfully. CMS: please send us list of stuck files so we can check
- IN2P3 - ntr
- RAL - completed intervention and bringing up all services
- OSG - ntr
- NDGF - 2 small issues: 1) power outage on 1 of storage sites - some of ATLAS data temporarily unavailable, power back in a few hours 2) detected bug of gridftp door of dCache - just some errors detected (authentication errors) - will do a quick fix and restart of dCache services later today
- CERN Storage Services: before update one of SRM ATLAS nodes got stuck again - 04:00 - 11:00 intermittent. In addition reenabling gridftp checksums. Got disabled after bug found which has now been fixed.

AOB:

- From tomorrow ATLAS will start black-listing sites on DM based on downtimes - currently a manual intervention that sites are not used for transfers. Now an automatic system that takes info from GOCDB. As this is now automatic scripts has no mercy - please make sure you register site downtimes and type (LFC, SRM etc.) correctly.

Thursday

Attendance: local(Roberto, Patricia, Harry, Andrea, Maarten, Ewan, Jamie, Maria, Gavin, Jean-Philippe, Simone, Zbyszek, Alessandro, Jan, MariaDZ);remote(Jon, Kyle, Rolf, Brian, Angela, Jeremy, Jens, Michael, JT, CNAF, Gonzalo, Joel).

Experiments round table:

- ATLAS reports -
 - ◆ Yesterday CERN SRM upgrade to 2-9.3 went fine. We did not observe source errors after the upgrade.
 - ◆ CERN-PROD srm-atlas not able to export data >4GB: ALARM ticket GGUS:57771 [\[Jan - 2nd time we had a gridftp checksum error. Thought we had fixed it but not so. Had intended to roll out for all other CASTOR instances but will now not. Simone - 3 teams of people tested this, CASTOR, experiment and experiment support. Largest file was 3.8GB. "Bad luck" \]](#)

- ◆ CERN-PROD srm-atlas ineff to FZK and PIC, team ticket
https://gus.fzk.de/ws/ticket_info.php?ticket=57787
 - ◆ PIC problem storage ATLAS pool down 00:57
https://gus.fzk.de/ws/ticket_info.php?ticket=57789 [Gonzalo - not much news but had quite a big incident - lost an ATLAS pool - almost 100TB pool. Problem with filesystem - in hands of SAN support (zfs filesystem). Almost 600K files.]
 - ◆ INFN-T1 MCDISK out of space
 - ◆ SARA-MATRIX SRM connection error at 3:50 CET
https://gus.fzk.de/ws/ticket_info.php?ticket=57793 [JT- 1) filesystem which is supposed to be highly redundant giving problems. Hence problems at dCache level. 2) All 24 dcap servers had 500 active logins(!) - also froze dCache. Restarted all doors - corresponding jobs probably lost. Why so many active connections? Maarten - logs? Simone - WN running ATLAS jobs? Intensive file merging activity. Many inputs (50+) and 1 output. One job has 100 input jobs. This can give many concurrent connections. JT - could also explain load on experiment s/w area.]
 - ◆ BNL - burst of errors. GGUS:57801
 - ◆ INFN-T1 problems of jobs failing due to storage unavailability: GGUS:57813
 - ◆ Today Downtime:
 - ◇ no outage downtimes foreseen for Tier1s
- CMS reports -
 - ◆ T1 Highlights:
 - ◇ Running skimming and rereconstruction jobs
 - ◇ WMAgent is being commissioned at FNAL
 - ◇ CNAF: 31 files stuck in migration to tape for a long time (up to 3 weeks). [Had list of files "stuck" noticed that 12 do not exist any more, 18 didn't have putdone command. Only 1 really stuck. Under investigation why]
 - ◇ KIT: waiting for a list of files to be deleted from some disk pools that are almost full
 - ◆ T2 Highlights:
 - ◇ Running 7 TeV simulation
 - ◇ T2_US_Nebraska had an unexpected downtime due to LDAP server
 - ◇ T2_RU_PNPI transfers to T1_CH_CERN failing for timeout, will forward request to adjust timeouts on the dedicated channel on fts-t2-service.cern.ch Savannah #sr 114126
 - ◇ Still intermittent SAM errors at T2_EE_Estonia Savannah #sr 113947
 - ALICE reports - GENERAL INFORMATION: Low production at this moment. Several Pass 1 reconstruction jobs together with one train analysis are the current activities of the experiment.
 - ◆ T0 site
 - ◇ Implementation of the AFS volume replicas inside AliEn in production since yesterday evening
 - ◇ CAF: ALICE reported yesterday that around 50% of the specific CAF nodes for ALICE were in status maintenance or really down. The issue is being followed with the PES experts. ITCM tickets have been submitted for all those nodes not performing well
 - ◆ T1 sites
 - ◇ CNAF-T1: GGUS ticket concerning the bad behavior of one of the local cREAM-CE has been verified and closed.
 - ◇ RAL: Still out of production (CREAM-CE drain operation still ongoing)
 - ◆ T2 sites
 - ◇ BARI-T2: GGUS ticket concerning the bad behavior of one of the local cREAM-CE has been verified and closed.
 - ◇ GRIF_IPNO: ticket 57803: Local CREAM-CE not working (problems observed at submission time)

◇ GRIF_IRFU: ticket 57805: Local CREAM-CE not working (problems observed at submission time)

- LHCb reports - Reprocessing activity launched today with analysis activity at low level. MC production (10M events) launched.
 - ◆ T1 site issues:
 - ◇ RAL : downtime at RISK for intervention on WN for installing 32 bit CASTOR clients
 - ◇ NL-T1; Storage at SARA failed all SAM tests since yesterday night and also users reported a lot of problem up/downloading data out of there. It looks like SRM was switched off. (GGUS:57812)

Sites / Services round table:

- ASGC - ntr
- FNAL - ntr
- IN2P3 - ntr
- RAL - AT RISK for WN installation today went well. Do ATLAS have an idea of what max filesize will be??? Simone: guideline of 5GB. Enforced in Athena but not in data management system. Logfile can be e.g. 6-7GB.
- KIT - ntr
- NDGF - still have some problems this morning with pools being down after power outage yesterday but fixed now. Problem with door on tape library which opened itself over night! Caused tape traffic to stop for a few hours.
- BNL - 1) as mentioned by Alessandro related to a certain slowness of our SRM seen with write operations. About 1GB/s out of dCache using SRM methods see slowness of write - experts investigating. No reason to send tickets - we are working on it. 2) Slow file transfer rates IN2P3 -> BNL: asked network experts in Lyon for assistance. Setup iperf server, people at Lyon ran iperf tests against BNL last night and rates << expectations. Working on this with network providers.
- NL-T1 - other than ATLAS issues above: national holidays tomorrow and May 5 - no representation on those days
- CNAF - ntr
- PIC - ntr
- OSG - ntr
- GridPP - had meeting with LHCb today to understand transfer problems from T2s. Jobs try to transfer jobs back from T2s to T1s at end of job - efficiency very low: 50% down to 2%. Only reproduce with not-DPM sites. Struggling to understand. Port or firewall issue?
- CERN Storage - change assessment? No - had announced as outstanding action on s/w rollout. Maybe something to add to template (test below, at and above magic limits)
- CERN DB - LHCb replication issue. From noon until now replication offline-T1s frozen due to unknown problem with downstream capture process. Only conditions affected. Maybe related to stuck transaction? Since 11:00 no data propagated to T1s for conditions.

AOB:

Friday

Attendance: local(Jamie, Maria, Harry, Maarten, Roberto, Jean-Philippe, Lola, Eva, Nilo, Dirk, Alessandro, Simone);remote(Jon, Gang, Gonzalo, Rolf, Michael, John, Andrea, Rob, Jan).

Experiments round table:

- ATLAS reports -

- ◆ INFN-T1 unscheduled downtime goddb downtime GGUS:57813 [↗](#) [Davide - comment later on about problems yesterday]
- ◆ PIC problems in data transfer GGUS:57848 [↗](#)
- ◆ ATLAS Reprocessing Status
 - ◇ ~ 97% done. Details ATLAS Reprocessing Status
- ◆ Today Downtime:
 - ◇ no outage downtimes foreseen for Tier1s
- CMS reports -
 - ◆ T1 Highlights:
 - ◇ Ready for the upcoming high intensity 900GeV fill.
 - ◇ Running skimming and re-reconstruction jobs.
 - ◇ WMAgent is being commissioned at FNAL. Processing is done, merges are being forced out; no failures. Everything looks good.
 - ◇ CNAF: due to an unexpected problem the storm service was unavailable, since yesterday 5 pm. Fixed today around 12 am.
 - ◇ KIT: waiting for a list of files to be deleted from some disk pools that are almost full.
 - ◆ T2 Highlights:
 - ◇ Running 7 TeV simulation
 - ◇ We will be doing central space cleanup next week (deletion requests for >500TB worth of data across central space).
 - ◇ T2_US_MIT had a power outage. All Tier-2 services were offline. They restored operations shortly.
 - ◇ T2_RU_PNPI transfers to T1_CH_CERN failing for timeout, will forward request to adjust timeouts on the dedicated channel on fts-t2-service.cern.ch Savannah #sr 114126
 - ◇ Still intermittent SAM errors at T2_EE_Estonia Savannah #sr 113947
- ALICE reports [↗](#) - GENERAL INFORMATION: Pass1 reconstruction tasks together with one analysis train
 - ◆ T0 site
 - ◇ A detailed list of all ALICE CAF nodes and their status have been sent to PES experts. About 50% of these nodes are out of production for different reasons.- Update Monday with new information
 - ◆ T1 sites
 - ◇ RAL announced yesterday a delay for final phase of the CREAM-CE intervention. It should be finished by for today mid afternoon.
 - ◆ T2 sites
 - ◇ Birmingham: The local CREAM-CE is failing at submission time in the case that the submitted jdl contains an ISB (globus_ftp_client errors). GGUS_57847
 - ◇ GRIF_IRFU: Ticket GGUS_57805 reopened. The resource BDII of the local CREAM-CE is still reporting the wrong information
 - ◇ Cyfronet: local ALICE VOBOX is not accessible. GGUS_57851
 - ◇ Grenoble: Strange situation observed at the local VOBOX. The \$HOME directory of the mentioned service shows hundreds of empty directories called home_cream_XXXXXXX. Reported to the ALICE contact person at the site
- LHCb reports - 2 MC Simulation running smoothly (3K jobs) and 2 Data Reconstruction activities (MagUp and MagDown). Discovered yesterday a problem affecting many jobs (50% failures rate) internal in LHCb. Upload issues at SARA and (temporarily in the afternoon yesterday) at CNAF
 - ◆ T0 site issues:
 - ◇ High memory consumption from one of the Streams queue monitor processes observed on lhcb downstream capture. Does not affect replication. Problem being investigated by Oracle.

◇ Lhcb downstream capture for conditions was stuck since midday yesterday, DB people looking at the problem. Any news? [Problem with downstream capture fixed after meeting - big transaction on source DB - didn't fit in memory. Capture should have aborted with error but got stuck. Increased amount of memory for capture. Affected only conditions. Roberto to be added to mailing list to get info about such interventions.]

◆ T1 site issues:

◇ CNAF: a temporary glitch on StoRM yesterday afternoon affecting one hour SAM tests jobs and few reconstruction jobs to upload data output. Now it is OK since yesterday 16:00 UTC.

◇ NL-T1: still actual the problem with the SE affecting also reconstruction jobs attempting to upload output (GGUS: 57812)

◆ T2 sites issues:

◇ UK T2 continuing the investigation on data upload issue spawned time ago by LHCb.

◇ Shared area problems at 3 different sites: INFN-NAPOLI, UKI-LT2-UCL-CENTRAL ESA-ESRIN

Sites / Services round table:

- NL-T1: NL-T1 will not dial in today because today it is a public holiday in the Netherlands. Yesterday, we have had some performance issues with dCache and we have spent a lot of time in fine tuning. This has helped to some extent but there are still issues.
- FNAL - ntr
- ASGC - ntr
- PIC - situation hasn't changed: pool we lost from ATLAS still cannot be mounted. In contact with Sun support. Localized with zfs filesystem. Can see filesystem and data but at time of mounting "or importing" in zfs jargon hangs. Investigating with Sun experts
- IN2P3 - filed SIR for problem with downtime notification
- BNL - investigating transfer rate problems. Not restricted to international connectivity - investigated transfer rates achievable from CERN-BNL: limited to 300MB/s. Also for BNL-T2 centres. In contact with network providers. Debugging.[Alessandro - did you include Edoardo in the loop? Or else should we ask that he joins meeting in Monday. A: escalated to high level in US - looks like something wrong close to BNL.]
- RAL - ALICE mentioned CREAM CE outage - hope to have it fixed by end of day. Monday is holiday in UK - noone from RAL.
- CNAF - problems yesterday with CMS and ATLAS. CMS problems seem to be solved; following up on ATLAS. part of problem is full MCDISK. Also some gpfs inconsistencies. Hope to be solved later this afternoon. Post-mortem on Monday. [Roberto - also something for LHCb? A: no GGUS ticket. Happened around 16 UTC - matches same timeframe for gpfs problems]
- KIT - had problem with ATLAS disk pool server this morning and had to reboot. Seen by ATLAS? No 😊
- CERN Grid Services - yesterday at end of afternoon all LCG CEs started failing SAM tests. Missing attribute required by SAM tests to target a particular CE. Around about midnight fixed by Ulrich - he did something manually to prevent problem from happening again but still not understood. (glue_ce_hostname)

AOB:

-- JamieShiers - 23-Apr-2010

This topic: LCG > WLCGDailyMeetingsWeek100426

Topic revision: r17 - 2010-06-11 - PeterJones



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