

Table of Contents

Week of 100405	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.....	6
Thursday.....	8
Friday.....	11

Week of 100405

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 Availability				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:

NO MEETING! - CERN CLOSED

Experiment summaries:

- ATLAS reports -
 - ◆ 0. ATLAS collected data during the night and morning (doubled the number of collected events)
 - ◆ 1.a downtime periods for RAL (Sunday morning (few hours) and sunday evening (< 1 hour)) (GGUS #57015), Backup GOCdb when RAL unaccessible ?
 - ◆ 1.b GGUS Team GGUS:57022 to TAIWAN: Failing transfers ('ARCHIVE_LOG space had to be increased')
 - ◆ 1.c Missing?/Unaccessible files in TRIUMF (GGUS:56849). Triumf investigating but transfer was declared OK by FTS.
 - ◆ 1.d Slow transfer rate from BNL for small files (elog : 11132) : Affecting 'direct writing to the disk pools without going through the gridftp servers/doors'. To be validated with FTS support.
 - ◆ 2.a GGUS Team #57008: FTS channel IFIC-> PIC (managed by PIC) is stuck since last thursday (FTS jobs not starting). Affects MC production.
 - ◆ 2.b Temporary broken sites for transfers : INFN-MILANO-ATLASC (GGUS:57011), GOEGRID (GGUS:57006), CYFRONET-LCG2 (partially : GGUS:57025), SLAC (GGUS:57018)

- CMS reports -
 - ◆ T0 Highlights
 - ◇ Running normal 7 TeV collision processing
 - ◇ Migration to tape interrupted for 6 hours on Saturday morning and since Sunday at 13:00
 - ◆ T1 Highlights:
 - ◇ Skimming collision data
 - ◇ Storage problems at CNAF seen in transfer quality and SAM tests. The problem was probably due to the disk servers being saturated by the reprocessing on the farm (more than 1500 jobs running, with read peaks of 1.5 GB/s read and up to 600 MB/s migration to tape). The situation normalized when the number of jobs decreased. This kind of problem will probably be mitigated in the next weeks when the new storage will be deployed (the deployment of the new CPUs instead is already ongoing) (Savannah #sr113674)
 - ◇ Network problems at RAL on Sunday
 - ◇ On Saturday, large number of expired transfers (~200) at FNAL for for both FNAL -> CNAF and Buffer -> MSS migration (Savannah #sr113665)
 - ◇ At FNAL, 4 files caught on a broken storage node that have not made it yet to tape (Savannah #sr113676)
 - ◆ Detailed report on progresses on tickets:
 - ◇ [OPEN] T0_CH_CERN - Files on archived tape at T0 - Remedy #CT663457, and Savannah #112927 Update 8/3: files recovered from other sites when possible, but other files only have a copy on the damaged tape.
 - ◇ [OPEN] T0_CH_CERN - Remedy #653289 - - CERN has a high fraction of aborted JobRobot jobs with the Maradona error. Update 10/3: trying two additional options in the LSF configuration, following an advice from Platform computing. Update 17/3: additional issue identified, but changes did not solve the problem - contact with Platform continues. Update 19/3: efficiency improved. Update 23/3: efficiency back to 80% and today's values are still at 80%. Update 24/3: failures look correlated to one particular CE.
 - ◇ [OPEN] T2_CH_CAF - A burst of transfer failures at 22:30 UTC April 1st left 216 files with 0 size on cmscaf, cleaned up manually with nsrm - details in GGUS TEAM #56997 and Savannah #113652
 - ◇ [OPEN] T1_ES_PIC - file access problem at PIC: repeated errors trying to access one file during the 355 rereco preproduction - Savannah #113582.
 - ◇ [OPEN] T1_US_FNAL - Large number of expired transfers at FNAL - [Savannah #113665
 - ◇ [OPEN] T1_US_FNAL - 1 File Missing from FNAL to GRIF - Savannah #113676
 - ◇ [OPEN] T1_IT_CNAF - Transfer Errors T0, FNAL to CNAF - Savannah #113674
 - ◆ T2 highlights
 - ◇ Waiting for MC production requests
- ALICE reports [↗](#) - GENERAL INFORMATION DURING THE EASTER PERIOD: The ALICE offline shifters have not reported any remarkable issue during this period. However it seems that CAF has not been properly behaving until this evening (origin of the problem to be checked with the ALICE experts just after the vacations period). Reconstruction activities successfully ongoing following the latest shifter report. From the 13 physics runs recorded to 7TeV, 8 have been fully reconstructed and the rest are almost ready.
 - ◆ T0 site
 - ◇ Good behavior of the T0 services during the vacation period. Services have been continuously involved in reconstruction activities with no remarkable issues
 - ◆ T1 sites
 - ◇ no remarkable issues to report during the Easter period. The current main activity is concentrated in Pass1 reconstruction tasks, the T1 sites have been in production until

toda in the evening with a minimum level of production at the current moment
(remaining analysis activities)

- ◆ T2 sites
 - ◇ remarkable issues at Hiroshima (still problem with the local CREAM-CE) and KFKI (still out of production. The site admin got in contact with us to follow any possible issue at the site. Checking still ongoing)
- LHCb reports -
 - ◆ 4th April 2010(Sunday)
 - ◇ Reconstruction of data received from ONLINE all weekend long. One unprocessed file out of 14 for one production. A fraction (prod. 6220) of the reconstruction jobs (originally at CERN) have been resubmitted several times because stalling (exceeding CPU time).
 - ◇ Issues at the sites and services
 - T0 sites issues: none
 - T1 sites issue:
 - NIKHEF: banned for the file access issue (GGUS ticket: 56909)
 - GRIDKA: banned for shared area issue preventing systematically all jobs to setup the environment (GGUS ticket: 57030)
 - T2 sites issue:
 - INFN-CATANIA: queue publication issue
 - INFN-MILANO-ATLASC: shared area issue
 - UKI-SOUTHGRID-OX-HEP: too many pilot aborting
 - Data reconstruction at T1's.
 - ◆ 5th April 2010(Monday)
 - ◇ Data reconstruction at T1's.
 - ◇ T1 sites issue:
 - CNAF: all jobs for one bunch of reconstruction have been declared stalled consuming 0 CPU Time. Suspicious: a shared area issue but it is not clear at the first glance.
 - RAL: Network problem affecting data upload and new jobs to be picked up by RAL. .
 - still true pb at NIKHEF: banned for the file access issue (GGUS ticket: 56909)
 - still true pb at GRIDKA: banned for shared area issue preventing systematically all jobs to setup the environment (GGUS ticket: 57030)

Tuesday:

Attendance: local(Miguel, Steve, Dawid, Ueda, Malik, Dirk, Miguel, MariaG, Jamie, Roberto, Ewan, Maarten, Flavia).

Experiments round table:

- ATLAS reports -
 - ◆ LHC - The machine stop on Thursday 8th April is confirmed and will last from 08:00 to 18:00. T1
 - ◆ BNL - FTS had a problem with source sites TECHNION-HEP (GGUS #57039) -- fixed similar error observed with TR-10-ULAKBIM (GGUS #57026, 2010-04-04)
 - ◆ FTS - some issues/worries as in the report on April 5.
 - ◆ Discussion: Steve: Do tickets to follow-up on FTS exist: Gueda: yes. Michael: BNL transfer problem to new sites - BNL uses a manual procedure for including new sites in the info system after validation. Expect to automatize this procedure soon. Brian: in UK seeing FTS transfers stuck in preparing state (preparing a ticket with more information). Are other T1

seeing this? Not so far. Gang: TEAM ticket against ASGC now closed: reason was unexpected size of archive log file - will follow-up with castor team.

- CMS reports -
 - ◆ T0 Highlights
 - ◇ Processing collisions data
 - ◇ Free space on the CMSCAFUSER pool has dropped below 10%. Experts have been notified
 - ◆ T1 Highlights:
 - ◇ Skimming collisions data
 - ◇ ASGC: problems with tape migration due to problematic tape labels at some tapes, fixed (Savannah sr #113683)
 - ◆ Detailed report on progresses on tickets:
 - ◇ [OPEN] T0_CH_CERN - Files on archived tape at T0 - Remedy #CT663457, and Savannah #112927 Update 8/3: files recovered from other sites when possible, but other files only have a copy on the damaged tape.
 - ◇ [OPEN] T0_CH_CERN - Remedy #653289 - - CERN has a high fraction of aborted JobRobot jobs with the Maradona error. Update 10/3: trying two additional options in the LSF configuration, following an advice from Platform computing. Update 17/3: additional issue identified, but changes did not solve the problem - contact with Platform continues. Update 19/3: efficiency improved. Update 23/3: efficiency back to 80% and today's values are still at 80%. Update 24/3: failures look correlated to one particular CE.
 - ◇ [OPEN] T2_CH_CAF - A burst of transfer failures at 22:30 UTC April 1st left 216 files with 0 size on cmscaf, cleaned up manually with nsrm - details in GGUS TEAM #56997 and Savannah #113652
 - ◇ [OPEN] T1_ES_PIC - file access problem at PIC: repeated errors trying to access one file during the 355 rereco preproduction - Savannah #113582.
 - ◇ [OPEN] T1_US_FNAL - Large number of expired transfers at FNAL - [Savannah #113665
 - ◇ [OPEN] T1_US_FNAL - 1 File Missing from FNAL to GRIF - Savannah #113676
 - ◇ [OPEN] T1_IT_CNAF - Transfer Errors T0, FNAL to CNAF - Savannah #113674
 - ◆ Discussion - Miguel: We haven't yet received a ticket from CMS about the tape migration problems - please create one
- ALICE reports [↗](#) -
 - ◆ T0 site - Good behavior of the T0 services during the vacation period. Services have been continuously involved in reconstruction activities with no remarkable issues
 - ◆ T1 sites - no remarkable issues to report during the Easter period. The current main activity is concentrated in Pass1 reconstruction tasks, the T1 sites have been in production until today in the evening with a minimum level of production at the current moment (remaining analysis activities)
 - ◆ T2 sites - remarkable issues at Hiroshima (still problem with the local CREAM-CE) and KFKI (still out of production. The site admin got in contact with us to follow any possible issue at the site. Checking still ongoing)
- LHCb reports - (Roberto) Data reconstruction at T1's proceeding smoothly apart few errors from ONLINE (sending BEAM1 data type flag instead of COLLISION10) and few jobs stalled. Reduced the size of input raw data file to fit the available queue lengths. 8.7 million events written at the moment. 10% in process is actually physics.
 - ◆ T0 sites issues
 - ◇ CERN: LFC read-only is unreachable timing out all requests and causing many jobs failing. (GGUS:57053 [↗](#)) Ewan: being followed-up
 - ◆ T1 sites issue:

- ◇ CNAF: all jobs for one bunch of reconstruction have been declared stalled consuming 0 CPU Time. Suspicious: a shared area issue but it is not clear at the first glance.
- ◇ RAL: Network problem affecting data upload and new jobs to be picked up by RAL: CIC
- ◇ still true pb at NIKHEF: banned for the file access issue (GGUS ticket: 56909)
- ◇ still true pb at GRIDKA: banned for shared area issue preventing systematically all jobs to setup the environment (GGUS ticket: 57030) - Xavier/KIT: KIT ticket has been updated this morning, but work still in progress.
- ◆ T2 sites issue:
 - ◇ UKI-SCOTGRID-ECDF all pilot aborting there
 - ◇ UKI-SCOTGRID-DURHAM all pilots aborting there

Sites / Services round table:

- FNAL - file transfer to grif: ticked closed - faulty disk pool back. Experienced expired transfers, but suspect problem of central phedex at CERN
- PIC - GGUS ticket from ATLAS has been solved this morning by restarting the T2 SRM. In case of remaining FTS functional issues ATLAS should follow-up with FTS development
- KIT- ntr
- IN2P3 - ntr
- ASGC - SAM performance at ASGC was degraded. The firewall configuration for disk servers was corrected.
- GridPP - ntr
- NL-T1 - dcache chimera db upgrade planned for Thu 9-11 CET
- NDGF - downtime in one site, but did not effect users. High load on some T2 nodes: fixed by killing root application with large memory footprint. NDGF should follow up with experiment about expected memory size in case this problem occurs repeatedly.
- OSG - some uptake in ticket counts, but issues are still timely handled. S/w update last week went smoothly - waiting for upcoming ggus update
- RAL network interruption (2 periods of 3h and 1h on Sun) due to faulty interface card in a router - stable since the card was replaced. Will forward discussion on GOCDB failover to GOCDB team at RAL. Q: should sites retrospectively add the downtime once GOCDB is available again? Yes, please. One disk server in atlas scratch is currently unavailable.
- BNL- preparing upgrade of condor based batch system: first step today (transparent) central master will be redirected to new system (existing queue will stay unaffected). Second step will be more disruptive as all batch clients need to be upgraded. Planning for sometime in May and will be pre-announce the proposed schedule. The investigation on connection counts between LFC server and Oracle backend (JPB and BNL DBA) has concluded: lead to new LFC daemon which can use the oracle connection pool. Special thanks to Jean-Phillipe.
- **CERN FTS** - Incident. CERN-ASGC was stuck 08:00 -> 20:00 on April 1st. Incident report in progress: IncidentFTS010410

AOB:

Tuesday:

- RAL: is scheduled LHC stop for 8th confirmed? Yes. Is also the 3 day technical stop confirmed? Not yet - more news tomorrow.
- CERN/DB: CMS phedex issues: asked for more detailed feedback, but did not get any yet. FNAL: phedex was restarted which solved the issue.

Wednesday

Attendance: local(Ueda, Jamie, MariaG, Miguel, Nicolo,David, Carlos, Przemyslaw, AndreaV, Maarten, Ignacio, Flavia, Steven, Roberto, Dirk);remote(Gang/ASGC, Jon/FNAL, Kyle/OSG, Gonzalo/PIC, John/RAL, Michel/Grif, Rolf/IN2P3, Ron/NL-T1, Xavier/KIT, Michael/BNL, Michel/Grif, Alessandro/CNAF).

Experiments round table:

- ATLAS reports -
 - ◆ T1
 - ◇ RAL -- we observed failing transfer attempts, but it was immediately understood thanks to the report yesterday (the offline diskserver). The file was successfully transferred this morning.
 - ◇ TRIUMF -- the missing files in TRIUMF (see report Apr. 5) are being recovered to resume data distribution. The issue for unexpected deletion of the files still to be followed. (ggus 56849)
 - ◇ PIC -- transfers from IFIC have been successful after the solution of ggus #57008. A ticket will be sent to FTS after collecting information concerning the FTS behavior.
 - ◇ ASGC -- some files were not available in TAIWAN-LCG2, disk server problem fixed (ggus #57098).
- CMS reports -
 - ◆ T0 Highlights - Processing collisions data.
 - ◇ Free space on the CMSCAFUSER pool has dropped below 10%. Experts have been notified
 - ◇ Gridmaps not accessible - https://gus.fzk.de/ws/ticket_info.php?ticket=57052 - now back
 - ◆ T1 Highlights:
 - ◇ Skimming collisions data
 - ◇ Running backfill
 - ◆ Detailed report on progresses on tickets:
 - ◇ [OPEN] T0_CH_CERN - Files on archived tape at T0 - Remedy #CT663457, and Savannah #112927 Update 8/3: files recovered from other sites when possible, but other files only have a copy on the damaged tape.
 - ◇ [OPEN] T0_CH_CERN - Remedy #653289 - - CERN has a high fraction of aborted JobRobot jobs with the Maradona error. Update 10/3: trying two additional options in the LSF configuration, following an advice from Platform computing. Update 17/3: additional issue identified, but changes did not solve the problem - contact with Platform continues. Update 19/3: efficiency improved. Update 23/3: efficiency back to 80% and today's values are still at 80%. Update 24/3: failures look correlated to one particular CE.
 - ◇ [CLOSED] T2_CH_CAF - A burst of transfer failures at 22:30 UTC April 1st left 216 files with 0 size on cmscaf, cleaned up manually with nsrm - details in GGUS TEAM #56997 and Savannah #113652
 - ◇ [OPEN] T1_ES_PIC - file access problem at PIC: repeated errors trying to access one file during the 355 rereco preproduction - Savannah #113582.
 - ◇ [OPEN] T1_US_FNAL - Large number of expired transfers at FNAL - [Savannah #113665

- ◇ [CLOSED] T1_US_FNAL - 1 File Missing from FNAL to GRIF - Savannah #113676 - transfer completed, closed.
- ◇ [CLOSED] T1_IT_CNAF - Transfer Errors T0, FNAL to CNAF - Savannah #113674 - CNAF issues due to high load from jobs on storage, now OK, will be improved by new hw deployment.
- ◆ T2 highlights
 - ◇ MC production running
 - ◇ T2_CN_Beijing --> T1_FR_CCIN2P3 transfers timing out - added size-based timeout on FTS channel
 - ◇ T2_BE_IHEP has lost the software area (Savannah sr #113682)
 - ◇ T2_RU_IHEP - SAM CE errors (job submission failing)
 - ◇ T2_UK_London_Brunel - missing CMS site-local-config.xml on new SLC5 CE, fixed
- ◆ Discussion:
 - ◇ Nicolo: not clear if migration problem at T0 is due to CASTOR or Phedex.
 - ◇ Miguel/CERN: policies may impact on this as well and could be adapted to CMS expectation. Concerning CAF extension: starting to allocate new h/w, but please open separate ticket to track this more urgent request.
 - ◇ Jon: full functionality of gridmap has not been restored. Eg functionality to drilling down problems has not been restored yet. Nicolo: GGUS ticket is still open and recovery may still be in progress.
 - ◇ Gonzalo: see currently only backfill jobs (400 sustained) - could the amount of backfill jobs be reduced? Nicolo: will bring this up with data ops. Right now default is to fill the pledge, but this can be negotiated.
- ALICE reports [↗](#) - GENERAL INFORMATION: Reconstruction, analysis and MC production activities still going on (increase in the number of activities).
 - ◆ T0 site and T1 sites
 - ◇ Operations at the T0 site continuing smoothly. However several instabilities are observed with ce201 (CREAM-CE) is out of production (timeout error messages at submission time). no interruption in the production (the 2nd CREAM-CE is nicely performing together with the LCG-CE)
 - ◇ All T1 sites in production, no remarkable issues to report
 - ◆ T2 sites
 - ◇ No fundamental issues observed. Several issues at Wuham (wrong information provided by the local CE) and Hiroshima (wrong communication between VOBOX and CREAM-CE) are preventing these sites of entering in production. Following the configuration issues with the site admins and the CREAM developers
- LHCb reports -
 - ◆ Experiment activities:
 - ◇ Reconstruction (in total 140 jobs to reconstruct all data collected so far) and many users running their own analysis on this data. It should be worth to highlight that it is now close to a week that two important T1's are out of the production mask in LHCb : NIKHEF/SARA and GRIDKA.
 - ◆ Issues at the sites and services
 - ◇ T0 sites issues:
 - CERN: LFC read-only is unreachable timing out all requests and causing many jobs failing. This is again due to the sub-optimal Persistency LFC interface. A patch is about to come pending some test form LHCb. In the mean time the stop of LHC will allow to use frozen ConditionDB that will be deployed at the sites as local SQLiteDB and users will be able to analyze the close to 11 millions event recorded so far in the coming days.
 - ◇ T1 sites issue: b.pic.es portal has been accidentally stopped. Mistake promptly

recovered.

- PIC: the lhcbwe
- still true pb at NIKHEF: banned for the file access issue (GGUS ticket: 56909). Under investigation.
- still true pb at GRIDKA: banned for shared area issue (GGUS ticket: 57030). A first sight, it seems a degradation of performances due to the concurrent heavy activity that ATLAS software is putting on the shared area.

◇ T2 sites issue:

- none

Sites / Services round table:

- ASGC - ntr
- BNL - upgrade of condor went smoothly and had no adverse impact on running jobs.
- FNAL - reminder: open savannah ticket has nothing to do with FNAL but with some T3 and UK sites. FNAL will now close the ticket
- OSG - ntr
- PIC - ntr
- IN2P3- ntr
- NL-T1 - tomorrow short downtime (2h) for move of dcache db to new h/w
- KIT - ntr
- RAL - returned disk server to ATLAS. Tomorrow: two at risks registered for top level bddi and DB kernel upgrade
- GRIF - ATLAS dashboard shows GRIF information as not available - Michel send an email to ATLAS ops and is waiting for feedback.
- CNAF - ATLAS ticket: 57084: suspected catalog problem at CNAF turned out to be a problem on the ATLAS side. Ueda confirmed. Tomorrow: DB intervention - need support from CERN DB team for stream config. CERN DB confirmed that someone will be available to help. Email confirmation with contact name will go out soon.
- CERN: this morning between 7:19-8:56 some router instabilities occurred, which could have affected eg CNAF, IN2P3, FNAL.

AOB:

Thursday

Attendance: local(Ueda, Flavia, Maarten, Ignacio, Lola, Ewan, Nicolo, Patricia, Roberto, Sveto, Julia, Dirk);remote(Jon/FNAL, Michael/BNL, Xavier/KIT, IN2P3/Rolf, Angela/KIT, Gang/ASGC, JT/NL-T1, Jeremy/GridPP, Gonzalo/PIC, Gareth/RAL, Cristina/CNAF,Rob/OSG).

Experiments round table:

- ATLAS reports -
 - ◆ T0 : "checksum" error for a RAW (calibration) data file from the last year. ATLAS Tier-0 and DAQ people informed (ggus #57101)
 - ◆ INFN-T1 -- LFC downtime, the whole cloud is set offline in the production system, and subscriptions for transfers are stopped.
 - ◇ announcement came only yesterday, a little short for us to treat properly. was it due to the late confirmation of the LHC stop?
 - ◆ SARA-MATRIX -- SE downtime (gocdb:74655452,74605455), we did not set the site offline because the downtime was short.
 - ◆ BNL -- ggus #57121 has been sent by a shifter for files not transferred. The response to the ticket suggests a possible issue in ATLAS DDM. We will look into this.

- ◆ Discussion: Michael/BNL - looking at transfer delays. No explanation so far but working closely with ATLAS.
- CMS reports -
 - ◆ T1 Highlights:
 - ◇ FNAL
 - 0% transfer quality in all transfers to FNAL from 15:00 to 19:00 UTC on 2010-04-07 - FileDownloadVerify script issue?
 - ◇ Skimming collisions data
 - One failure at KIT for input 'File too large' (36 GB)
 - ◇ Running backfill
 - ◆ Detailed report on progresses on tickets:
 - ◇ [OPEN] T0_CH_CERN - Files on archived tape at T0 - Remedy #CT663457, and Savannah #112927 Update 8/3: files recovered from other sites when possible, but other files only have a copy on the damaged tape.
 - ◇ [OPEN] T0_CH_CERN - Remedy #653289 - - CERN has a high fraction of aborted JobRobot jobs with the Maradona error. Update 10/3: trying two additional options in the LSF configuration, following an advice from Platform computing. Update 17/3: additional issue identified, but changes did not solve the problem - contact with Platform continues. Update 19/3: efficiency improved. Update 23/3: efficiency back to 80% and today's values are still at 80%. Update 24/3: failures look correlated to one particular CE.
 - ◇ [OPEN] T0_CH_CERN - Gridmaps not accessible - GGUS #57052 - now back, but not yet with full functionality
 - ◇ [OPEN] T1_ES_PIC - file access problem at PIC: repeated errors trying to access one file during the 355 rereco preproduction - Savannah #113582.
 - ◇ [CLOSED] T1_US_FNAL - Large number of expired transfers at FNAL - [Savannah #113665
 - ◆ T2 highlights
 - ◇ MC production running
 - ◇ T2_UA_KIPT (Kharkov-KIPT-LCG2) - BDII publication error
 - ◇ Squid issues at T2_US_Florida, T2_RU_PNPI
 - ◆ Discussion:
 - ◇ FNAL transfer failure might be do to verify scripts. Jon/FNAL: confirmed. A problem was introduced in the verification scrips as part of an effort to optimize tape efficiency. Fixed as soon as the script responsible realized. This affected only the transfer verification status, but not the actual transfers.
 - ◇ CMS is preparing a ticket for castor tape migration issues.
- ALICE reports [↗](#) -
 - ◆ GENERAL INFORMATION: MC cycle completed. Analysis train and reconstruction activities going on (in particular at the T0, pass1 reconstruction tasks and PHOS calibration tasks)
 - ◆ T0 site
 - ◇ ce201.cern.ch fails providing timeout errors. GGUS ticket submitted (57132). The issue is not stopping the ALICE tasks at all since the 2nd CREAM is perfectly performing and in addition, all reconstruction activities are beeing executed also through the LCG-CE.
 - ◆ T1 sites
 - ◇ All T1 sites in production. Checking in particular SARA (site has not been taking jobs for several days). Agents die immediatelly after the submission to the WN. Two error messages appearing: The agent cannot contact the Clustermonitor service running in the VOBOX, and in addition, the specific alice sw installed at the vobox seems not to be found (sw area issue). Getting more information from Alice before

contacting the site admin

- ◆ T2 sites
 - ◇ Issues reported during several days concerning the bad behavior of the CREAM-CE in Hiroshima has been solved. The site is back in production
 - ◇ GRIF has provided a CREAM-CE for ALICE at the site just before the Easter break. The system has been fully tested this morning and the site has been put in CREAM submission mode
 - ◇ CapeTown T2 has solved the firewall issues reported to the site admin several days ago. The site has been put in production
 - ◇ GSI voboxes in bad shape. The VOBOX currently in production cannot be reached via gsissh. The 2nd VOBOX (out of production for a long time) is intended to enter again production as the 2nd VOBOX (failover mechanism established for ALICE with the latest AliEnv2.18). However the VOBOX service available at this node is obsolete.
- LHCb reports -
 - ◆ Experiment activities:
 - ◇ No data to be processed but user analysis activities continuing.
 - ◇ Workaround to the currently substandard Persistency interface has been put in place by LHCb/Dirac. Appears to be working.
 - ◇ Many problems at T1s being investigated.
 - ◆ Issues at the sites and services
 - ◇ T1 sites issue:
 - NIKHEF/SARA: banned for the gsidcap file access issue (GGUS ticket: 56909). Under investigation.
 - GRIDKA: banned for shared area issue (GGUS ticket: 57030). SAM jobs still failing.
 - CNAF: banned due to user jobs failing after rename of COnditionDB instance the 30th of March. 42% failure rate over past two days, under investigation with local contact.
 - IN2P3: LHCb VO box had to be restarted.
 - ◆ Discussion:
 - ◇ shared area @ GRIDKA: high load seems to be due to high load shared area from other VOs. site is looking at extending the h/w resources.
 - ◇ Conditions DB access at RAL needs DIRAC configuration change to update the DB connect string.

Sites / Services round table:

- FNAL - ntr
- BNL - ntr
- KIT - ntr
- IN2P3 - ntr
- ASGC - ntr
- NL-T1 - h/w upgraded for chimera, Planned: network maintenance for April 20 (1day downtime)
- GridPP - ntr
- PIC - ntr
- RAL - h/w behind top-level bdii was updated successfully. The kernel upgrade for atlas conditions DB required to extended the intervention slot. RAL observed yesterday for the first time the saturation of the fibre link to CERN (9.6 Gb/s for some hour).
- CNAF - transfer checksums still disabled: working on resolving the remaining problems
- OSG - ntr
- Ignacio/CERN 36 TB have added to CMS CAF

- Sveto/CERN : Intervention ongoing with LHCb ONLINE router - expect network outages on the LHCb ONLINE DB

AOB:

Friday

Attendance: local(Ueda, Lola, Miguel, DB, Ewan, Flavia, Malik, Dirk);remote(Jon/FNAL, Xavier/KIT, Michael/BNL, Joel/LHCb, Gang/ASGC, Rob/OSG, Onno/ NL-t1, Gareth/RAL, Gonzalo/PIC, Tore/NDGF).

Experiments round table:

- ATLAS reports -
 - ◆ Report to WLCGOperationsMeetings
 - ◇ FTS
 - files removed after failed transfer attempts (ggus#56849) -- we could not find any explanation except that FTS might have deleted the sources file after the failures. We need a help from the FTS developers to understand the issue.
 - transfer jobs stuck in the channel IFIC-PIC between 2010-04-01 midnight to 2010-04-06 morning until the FTS and the source SE were restarted (57158).
 - ◇ CERN
 - checksum error reported yesterday : ATLAS DAQ calculated the checksum and did not verify it after copying the file to castor pool. therefore, it is possible the file was corrupted before being transferred to castor. the ggus ticket will be closed. (ggus #57101)
 - alarm ticket 57083 (2010-04-07) should have been reported on Wednesday
 - ◇ BNL -- investigation for the case ggus #57121 on going.
- CMS reports -
 - ◆ T0 Highlights
 - ◇ CASTORCMS DEFAULT degraded in SLS, caused by 15k user requests
 - ◆ T1 Highlights:
 - ◇ Skimming collisions data
 - New version preproduction running at FNAL
 - ◇ Running backfill
 - ◆ Detailed report on progresses on tickets:
 - ◇ [OPEN] T0_CH_CERN - Files on archived tape at T0 - Remedy #CT663457, and Savannah #112927 Update 8/3: files recovered from other sites when possible, but other files only have a copy on the damaged tape.
 - ◇ [OPEN] T0_CH_CERN - Remedy #653289 - - CERN has a high fraction of aborted JobRobot jobs with the Maradona error. Update 10/3: trying two additional options in the LSF configuration, following an advice from Platform computing. Update 17/3: additional issue identified, but changes did not solve the problem - contact with Platform continues. Update 19/3: efficiency improved. Update 23/3: efficiency back to 80% and today's values are still at 80%. Update 24/3: failures look correlated to one particular CE.
 - ◇ [OPEN] T0_CH_CERN - Gridmaps not accessible - GGUS #57052 - now back, but not yet with full functionality
 - ◇ [OPEN] T0_CH_CERN - Remedy #675138 - Investigation on tape migration delay on April 3-4
 - ◇ [OPEN] T1_ES_PIC - file access problem at PIC: repeated errors trying to access one file during the 355 rereco preproduction - Savannah #113582.
 - ◇ [OPEN] T1_TW_ASCG - Transfer errors T2_RU_SINP-->T1_TW_ASCG Savannah #113615 - authentication issues on source site solved, now 'FILE_EXISTS'

errors on destination.

- ◆ T2 highlights
 - ◇ MC production running
 - Transfer errors T2_RU_SINP-->T1_TW_ASGC Savannah #113615 - authentication issues on source site solved, now 'FILE_EXISTS' errors on destination.
 - ◇ T2_TR_METU (TR-03-METU) - BDII publication error
 - ◇ Squid issues at T2_RU_PNPI (ru-PNPI)
 - ◇ Reminder for sites running DPM : workaround for globus rfio library issue on SLC5
- ALICE reports [↗](#) -
 - ◆ GENERAL INFORMATION: No physics yesterday, analysis train activities ongoing. The current ALICE situation is the following:
 - ◇ 99.1% if the 7TeV events successfully reconstructed
 - ◇ 5.72TB of raw data
 - Still recorded only at the T0. Distribution of these raw data to T1 needs still discussion to the level of the experiment
 - ◇ 3.3 TB of ESD data
 - ◇ 10Mio event
 - ◆ ANNOUNCEMENT FOR THE SITES: The distribution of the latest AliEnv2.18 version is foreseen for Monday the 12th of April. This software installation will be performed centrally from CERN. Regional experts are encouraged to check the status of the local services at their sites once the distribution of the new AliEn version will be done (announcement through the ALICE TF mailing list)
 - ◆ T0 site
 - ◇ Issue reported yesterday with one of the local CREAM-CE (ce201.cern.ch) SOLVED. The system is back in production
 - ◆ T1 sites
 - ◇ minimal activity at these sites currently
 - ◆ T2 sites
 - ◇ Still problems with CapeTown and the local CREAM-CE. While specifying an InputSandbox field into the jdl, communication problems between the vobox and the local cream appears which prevents the submission of the inputsandbox to the CREAM-CE. following the issue with the site admin
 - ◇ VOBOX configuration problems at Wuham. Authentication problems while registering the user proxy into the VOBOX. Working together with thre site admin to find the problem
- LHCb reports -
 - ◆ Experiment activities:
 - ◇ No data to be processed but user analysis activities continuing.
 - ◇ On Monday 2009 real data reprocessing will be launched and few MC production correlated.
 - ◆ Issues at the sites and services
 - ◇ T1 sites issue:
 - NIKHEF/SARA: gsidcap file access issue (GGUS ticket: 56909). Reproducible: seems a clash deep in the libraries used by gsidcap and condition DB libraries.
 - GRIDKA: shared area issue (GGUS ticket: 57030). It seems OK
 - CNAF: banned until the new configuration (with right endpoints for CondDB) is deployed.
 - RAL: banned until the new configuration (with right endpoints for CondDB) is deployed.

Sites / Services round table:

- FNAL - ntr
- BNL - investigated recent transfer issues and found handshake problem with gridftp adapter which shows up with small files (few MB). Dcache developers have been informed and a temporary workaround using gridftp V1 has been put in place. Since then transfers (T2) are work fine also for small files. BNL points out that similar problems may also happen also for T1-T1 transfers. Ueda: reminded at proposal to adjust timeout in FTS. Dirk will follow up with FTS team.
- KIT - ntr
- ASGC - ntr
- OSG - will test disaster recovery on ticketing system - should be transparent (primary server will go down and service will fail over to secondary). More details at <http://osggoc.blogspot.com/2010/04/saturday-april-10-2010-ticket-system.html>
- NL-T1: data issues at sara/nikhef - is ticket up to date? Joel: will update with recent info.
- RAL : ntr
- PIC: ntr
- NDGF: ntr
- IN2P3: reminder: outage of tape system on Apr 13 (tue) for installing a new library. Correction: yesterday LHCb vobox restart - was rather a restart of the LHCb services running on that box due to wrong version of dirac (Joel confirms).

AOB:

-- JamieShiers - 30-Mar-2010

This topic: LCG > WLCGDailyMeetingsWeek100405

Topic revision: r16 - 2010-06-11 - PeterJones



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