

Table of Contents

Week of 100419.....	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 100419

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(Sveto, Maria, Gavin, Jamie, MariaDZ, Harry, Jan, Patricia, Maarten, Jean-Philippe, Nicolo, Lola);remote(Jon, Angela, Gang, Michael, Gareth, Ron, Kyle, Jeremy, Simone, Rolf, kyle).

Experiments round table:

- ATLAS reports - Report to WLCGOperationsMeetings
 - ◆ On Friday 16th at approx 9:30 CEST, BNL reported that information about BNL services were missing from the WLCG BDII. Se most serious side effect of this is that FTS channels to BNL could disappear at any time. From Michael's elog: "We observed that the CERN/WLCG BDII has dropped BNL from its information base. Though information about the site resources are complete and correct when querying the OSG BDII it is not known by the BDII at CERN". The reason of this has been explained by Michael and confirmed by Laurence/Maarten. From Michael: "Our early theory is that there were additional CEs (BNL_ATLAS_3 and BNL_ATLAS_4) recently added to the OSG BDII Interop Monitoring, this made 5 CEs and 1 SE reporting as the BNL-ATLAS Resource Group. This may have triggered a size limit in the transfer from the OSG BDII to the CERN BDII. After removing the additional CEs from the list of information to be published BNL resources are observed back in the BDII information base. OSG/Rob Quick will follow up on the technical issues with WLCG/Laurence Field et al." In fact this size limit (5MB) has been exceeded. I have open questions:
 - ◆ I still do not understand fully the issue: CERN has 20 CEs, but there is no problem exporting info in the WLCG BDII. The 5MB limit is for what exactly? And how do

we avoid this in the future?

◇ It was not clear how to escalate this problem (which was critical) to the relevant support unit in a timely manner. At the end the strategy "send email to all experts you know" worked but this might not be the case in the future.

- ◆ Gav - there is a file-based cache in the BDII.
- ◆ Maarten - could have opened a ticket against CERN BDII. This is a problem with all BDII's.
- ◆ Maria - contacting a number of experts + ticket
- ◆ Michael - on OSG side experts talked to expert at CERN who pointed out BDII not considered a critical service. Not included in set of critical services that are covered on 24x7 basis.
- ◆ MariaDZ = alarm ticket can also include cc: via mail.

• CMS reports -

◆ T1 Highlights:

◇ Monte Carlo reprocessing running at T1s

1. Some input files inaccessible at CNAF, FNAL, ASGC, site contacts notified.
2. Prestaging datasets at PIC

◇ Collision skimming running at T1s

1. KIT: failure opening 32 GB file with 'Input file too large': is there a file size limit on dCache?

◇ KIT

1. GGUS:57396 [↗](#) NFS lock blocking SW installation
2. GGUS:57417 [↗](#) Tape migration/recall slow

◆ 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_TW_ASGC - Intermittent issue with Maradona errors in CE SAM and JobRobot tests, Savannah #113850

◇ [OPEN] T1_TW_ASGC - Files with incorrect checksum on CASTOR disk pool, restaging from tape Savannah # 113897

◇ [CLOSED] T1_DE_KIT - batch system instabilities still causing intermittent failures in CE SAM tests Savannah #113824.

◇ [OPEN] T1_DE_KIT - GGUS:57396 [↗](#) NFS lock blocking SW installation

◇ [CLOSED] T1_DE_KIT - GGUS:57417 [↗](#) Tape migration/recall slow

◇ [OPEN] T1_IT_CNAF - 3 jobs failing persistently, input files corrupted Savannah #113844

◇ [OPEN] T1_IT_CNAF - 7 MC input files inaccessible at CNAF Savannah #113912

◇ [OPEN] T1_ES_PIC - 2 files continually failing exports from PIC, Savannah # 113899

◇ [OPEN] T1_ES_PIC - file with wrong checksum at PIC - probably corrupted before arriving at T1 Savannah # 113840

◆ T2 highlights

◇ MC production running

1. File access issues at T2_CN_Beijing, also in Jobrobot

2. Job failures at T2_IT_Pisa under investigation

- ◇ T2_IT_Bari SRM SAM test failures with expired certificate on gridftp server, fixed.
- ◇ T2_PL_Warsaw - site not visible in BDII

- ALICE reports [↗](#) - GENERAL INFORMATION: No fundamental issues reported during the weekend. Still waiting for the raw data transfers (T0-T1), therefore the activity at the level of T1 sites has been very low. Pass 1 reconstruction at the T0 have continued during the whole weekend, together with analysis train activities.
 - ◆ T0 site - Both submission backends have been in production during the full weekend with no incidents reported by ALICE
 - ◆ T1 sites - Minimal activity, waiting for the T0-T1 raw data transfers to start the Pass2 reconstruction at these sites
 - ◆ T2 sites
 - ◇ Nagios. Site in Dublin (not an ALICE site), GGUS:57206 [↗](#): "Alice jobs filling up csTCDie DPM storage element". Strange issue since ALICE is not running any specific test associated to the SE sensor with any site. Issue discussed with Nagios experts, the problem seems to be related with Nagios. The experts just switched off nagios (service nagios stop) on this box. GGUS ticket updated. Waiting for the site admins feedback
 - ◇ New CREAM-CE implementation will be performed this afternoon in Torino. Following the developers' advices, the new approach could get rid of the gridftp server included in the voboxes, leaving the OSB at the CREAM-CE (automatic purge procedure included in CREAM, should avoid any overload of the disk)
- LHCb reports -

Sites / Services round table:

- KIT - batch system stable again after being reinstalled - booted two times on w/e due to 1 misbehaving WN. Will validate that kernel update solved underlying cause - still don't know real root cause.
- FNAL - interesting effect over w/e. 12 minute GUMS site auth outage. Detected and restored. Unfortunately SAM tests for CMS didn't run and we get charged for a 3 hour downtime because of this - crossed UTC day change and got charged 2 day downtime! *ASGC - ntr
- BNL - ntr
- RAL - 2 things: 1) end of last week problems migrating CMS files to tape - resolved. Working through backlog. 2) flag up need for interventions. Declare an outage (probably..) to CASTOR around these. Probably during LHC stop. Tuesday next week? [Experiments to give feedback tomorrow]
- NL-T1 - issue (ATLAS) with a number of jobs failing at NIKHEF. Being looked into. Reminder of network outage at SARA tomorrow - already declared.
- IN2P3 - ntr
- GridPP - ntr
- OSG - ntr
- CERN DB - issue yesterday evening on integration DB for ATLAS. Got resolved around midnight. Apply aborted - running from Michigan from Muon calibration data. Simone - was at run meeting of ATLAS this morning where this came up. Run coordinator extremely unhappy that there are critical applications in ATLAS that still use the integration database. Pushing that these move to production DB ASAP. Taken up by ATLAS management... Maria - the 3 muon centres which are streaming from them to CERN no request for production support for these activities. Need to understand support level at remote site where streaming is initiated. Run coord wants plan for this to be integrated into production. Ball in ATLAS muon court for them to come up with a plan

AOB: (MariaDZ)

Monday:

- Urgent GGUS:57131 [opened](https://gus.fzk.de/ws/ticket_info.php?ticket=57336) on 2010/04/08 mentioned last Thursday (see AOB of <https://twiki.cern.ch/twiki/bin/view/LCG/WLCGDailyMeetingsWeek100412#Thursday>) is assigned to Atlas without any update.
- To close this month's post-GGUS-release ALARM test cycle can TRIUMF please acknowledge and close their ALARM ticket https://gus.fzk.de/ws/ticket_info.php?ticket=57336 ? ALARMS' results summary in tomorrow's MB report. When should we plan a 'full-chain' ALARM test according to doc https://twiki.cern.ch/twiki/bin/view/EGEE/SA1_USAG#Periodic_ALARM_ticket_testing_ru as decided at the Tier1 coordination meeting of 2010/03/25?

Tuesday:

Attendance: local(Harry, Nicolo, Gavin, Jamie, Maria, Jean-Philippe, Maarten, Simone, Alessandro, Luca, Andrea, Steven, Flavia);remote(Federico Stagni/CNAF, Gonzalo, Jon, Angela, Michael, Rolf, Pepe, Ronald, Jeremy, Rob, Gareth, Luca).

Experiments round table:

- ATLAS reports - Report to WLCGOperationsMeetings
 - ◆ GGUS:57131 [opened](https://gus.fzk.de/ws/ticket_info.php?ticket=57336) has been looked into. The problem is in Panda treating files on disk and tape at CERN (and trying to access files on tape before prestage). Issue is being fixed, ticket will be closed.
 - ◆ SOURCE errors from CERN (follow up). On Friday 16th we discussed the fake SOURCE errors from CERN to T1s and we concluded that the problem is in the destination SRM being slow offering a TURL (slow means approx 3 minutes, closed to the gridftp timeouts). After some investigation, I am not sure this is the all story. In the attached plot you see a burst of errors from CERN to all T1s for T0 exports and all errors appear as below. I doubt there was a problem to all T1s at the same time. This needs more investigation especially at the FTS logs level.
 - ◇ [FTS] FTS State [Failed] FTS Retries [1] Reason [SOURCE error during TRANSFER_PREPARATION phase: [CONNECTION_ERROR] [srm2__srmPrepareToGet] failed: SOAP-ENV:Client - CGSI-gSOAP running on lxb1413.cern.ch reports Error reading token data header: Connection closed] (plot in full report)
 - ◆ SARA cloud has been stopped in DDM and Prodsys because of SARA network intervention (12h)
- CMS reports -
 - ◆ [Data Ops]
 - ◇ Tier-0: data processing.
 - ◇ Tier-1: tail of running Spring10 MC rereco.
 - ◇ Tier-2: tails of low efficiency MC production in 356 and tails of Summer09 MC production.
 - ◆ [Facilities Ops]
 - ◇ T1_ES_PIC CMS tasks backlog: T1-contact was stuck in Sweden and arrived Barcelona yesterday. Backlog to be digested in the next coming days.
 - ◇ Missing tape-family-tag request for T1s. Procedure needs be improved, so sites can setup the areas well in advance.
 - ◇ Automatic migrations of datasets to CAF DBS to be stopped by Monday April 26th. This will restrict CAF DBS to CAF-only data (express streams, user-produced datasets).
 - ◇ CMS proposal for PIC: could the site take the opportunity of the LAN switch downtime on Tuesday 27th to also schedule the intervention to disable tape protection? [will affect all VOs sharing this dCache instance]
 - ◇ CMS proposal for RAL storage intervention: Wednesday 28th.

- ◆ T1 Highlights:
 - ◇ Monte Carlo reprocessing tails running at T1s
 1. Production on hold at RAL waiting for the backlog in the migration queue to be processed.
 - ◇ ASGC
 1. Directory permission errors in import of new MC datasets from T2s, fixed. Savannah #113955
- ◆ 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
 - ◇ [CLOSED] T1_TW_ASGC - Intermittent issue with Maradona errors in CE SAM and JobRobot tests, Savannah #113850 - update 20/4: recovered automatically, closed.
 - ◇ [OPEN] T1_TW_ASGC - Files with incorrect checksum on CASTOR disk pool, restaging from tape Savannah # 113897 - update 20/4: 8 files still corrupt, deleted and invalidated globally
 - ◇ [CLOSED] T1_TW_ASGC - Directory permission errors in import of new MC datasets from T2s, fixed. Savannah #113955
 - ◇ [OPEN] T1_DE_KIT - GGUS #57396 NFS lock blocking SW installation
 - ◇ [OPEN] T1_IT_CNAF - 3 jobs failing persistently, input files corrupted Savannah #113844
 - ◇ [OPEN] T1_IT_CNAF - 7 MC input files inaccessible at CNAF Savannah #113912
 - ◇ [OPEN] T1_ES_PIC - 2 files continually failing exports from PIC, Savannah # 113899
 - ◇ [OPEN] T1_ES_PIC - file with wrong checksum at PIC - probably corrupted before arriving at T1 Savannah # 113840
- ◆ T2 highlights
 - ◇ MC production running
 1. File access issues at T2_CN_Beijing, also affecting analysis jobs and Jobrobot - update 20/4: caused by misconfiguration of dcap port in local configuration file, now fixed.
 2. VO_CMS_SW_DIR inaccessible at T2_EE_Estonia, also seen in SAM CE and JobRobot tests, Savannah #113947.
 3. Job failures at T2_IT_Pisa still under investigation.
- ALICE reports [↗](#) - GENERAL INFORMATION: No physics runs during the last night, at this moment the ALICE activity concentrates on Pass 1 reconstruction activities, together with 2 analysis trains. No MC production for the moment
 - ◆ T0 site
 - ◇ About 30 nodes belonging to CAF of ALICE have been migrated to SL5 and moved from maintenance status to production by yesterday afternoon. Machines can be now be followed by the CC operators.
 - ◇ Both LCG-CE and CREAM-CE backends nicely working for the reconstruction activities

- ◆ T1 sites
 - ◇ FZK: Local operatios required this morning in pps-vobox to restart the services (site independent issue)
 - ◇ CNAF: Ticket submitted to the site: GGUS:57471 [↗](#). The local CREAM-CE is not performing well (timeout error messages observed at submission time from one of the ALICE VOBOXES placed at the site). The ALICE contact person at the site has been also warned. CNAF is already configured in full CREAM-mode. The production is not stopped thanks to the 2nd CREAM-CE available at the site and perfectly performing
- ◆ T2 sites
 - ◇ Bratislava, KFKI and MEPHI contacted to follow the status of the CREAM-CE setup at this sites. These are the last T2 sites still without the service available for ALICE.
 - ◇ Madrid-CIEMAT has entered in CREAM-mode this morning
 - ◇ Wuham and Cape Town: both sites have already entered production via CREAM-CE
- LHCb reports -
 - ◆ Experiment activities:
 - ◇ Data reconstruction, stripping and user analysis continuing.
 - ◇ T1 sites issue:
 - SARA: currently in downtime and banned for LHCb. GOC-DB announcement arrived hours after downtime had commenced.
 - ◇ T2 sites issue:
 - SAM jobs stalled at INFN-TRIESTE

Sites / Services round table:

- CNAF - ntr except next Monday "at risk" intervention on SAN. Should be completely transparent to experiments. On Tuesday & Wednesday downtime on ATLAS DB (LFC & conditions) to migrate to new h/w for RAC. Simone- 2 days downtime? A: unfortunately yes. Hope to end in 1 day but scheduled for 2.
- PIC - noticed yesterday a bunch ~800 jobs from ATLAS (pilot job role) stayed in batch for 12h with 0 CPU consumption. Staged correctly input file and then stopped. Killed around noon and reappeared again but ran successfully the 2nd time. What happened? Which channel can we ask about this? Simone - first expert of Panda is Xavier Espinal. He can start drilling down on this, navigating logs etc. Otherwise escalate to Panda devs if needed.
- FNAL - ntr
- KIT - batch system still stable.
- BNL - ntr
- IN2P3 - ntr
- NL-T1 - SARA currently down for network maintenance. Proceeding well and expect to finish before end of slot.
- RAL - AT RISK for some DB maintenance work behind LFC, 3D and FTS. Backlog of migrations to tape for CMS: 3/4 way through. Q about outage during next week's tech stop? Response in CMS report suggesting Wednesday would be a good day - only response seen so far.. Will take on board.
- ASGC - ntr
- GridPP - ntr
- OSG - BDII discussions yesterday: working with BNL to reduce size of info sent so can register all ATLAS CEs and stay withing 5MB limit of acceptable filesize. Escalation procedures yet to be defined as to what to do for such emergencies. Able to contact BDII experts via informal communication lines. Security ticket opened this morning but seems not to be security related. GGUS:57472 [↗](#) Submitter education issue?
- CERN LSF: slowness still an issue for ATLAS - following with platform 2 distinct problems: Thursday which triggered alarm - this was something else than slightly more chronic problem to

which I was referring in my previous statement.

- CERN SRM 2.9 - prelim tests for CMS went fine - green light for upgrade on Monday 26th (not yet officially scheduled)
- CERN DB - running reboot of ALICE Online .

AOB: (MariaDZ)

- The last USAG meeting will take place tomorrow, Wednesday April 21st @ 16hrs CEST (room 28-R-014 or by telephone). Main discussion item is its re-incarnation. Details on agenda: <http://indico.cern.ch/conferenceDisplay.py?confId=88353>
- The ALARMS' test round this month ended with conclusion <https://savannah.cern.ch/support/?113228#comment10> Will repeat TRIUMF alarm as pager did not go off. Not understood.

Wednesday

Attendance: local(Jamie, Maria, Lola, Nicolo, Simone, MariaDZ, Harry, Jean-Philippe, Steve, Carlos, Jan, Eva, Maarten, Alessandro, Gavin, Nilo);remote(Gonzalo, Gang, Jon, Onno, John, Rolf, Andreas, Michael, Luca, Rob, Greig).

Experiments round table:

- ATLAS reports - Report to WLCGOperationsMeetings
 - ◆ Still SOURCE errors during TRANSFER phase, today from CERN to TRIUMF. This is becoming a bit difficult to handle since we are not sure anymore if the problem is source or destination. Last ticket TRIUMF GGUS:57508 has been assigned to TRIUMF - needs to be understood asap to whom one should assign such tickets. These problems occur now every day - previously was ~once per month! [On DDM retry eventually goes..]
 - ◆ RAL-QMUL transfers piling up because of stuck FTS job. GGUS:57497. (A single transfer with 200 files hence kill the lot or not...)
- CMS reports -
 - ◆ T0 highlights
 - ◇ some files waiting for export for a long time - details will be added
 - ◇ even if export activity from CERN for CMS <300MB/s quite a lot of "source error during transfer preparation phase" - a few %. Collect SURLS and open ticket.
 - ◆ T1 Highlights:
 - ◇ Data and MC reprocessing with CMSSW_3_5_7 running at FNAL
 - ◇ Tails of Spring10 Monte Carlo processing running at T1s
 1. Green light to restart production at RAL after the backlog in the tape migration queue was processed.
 - ◇ CNAF
 1. Backfill workflow injected at full scale to test temporary solution for storage load issues seen during the last few weeks. [Luca - have already installed 2010 CPU pledges whilst storage not here yet. Load issues due to this - many more jobs running than previously and bandwidth to storage not enough. Separated meta-data/data access to optimize access. Real solution is to put in production storage that is being delivered these days. Q: many jobs running on farm overnight - up to 1080 jobs. Did not observe and serious errors. Feedback? A: not yet a report on test results so can't comment at moment. More info tomorrow...
 - ◆ T2 highlights

- ◇ MC production running
- ◆ 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] T0_CH_CERN - Gridmaps not accessible - GGUS #57052 - now back, but not yet with full functionality. Update 20/4: functionality restored, closed.
 - ◇ [OPEN] T1_TW_ASGC - Files with incorrect checksum on CASTOR disk pool, restaging from tape Savannah # 113897 - update 20/4: 8 files still corrupt, deleted and invalidated globally
 - ◇ [CLOSED] T1_DE_KIT - GGUS #57396 NFS lock blocking SW installation
 - ◇ [OPEN] T1_IT_CNAF - 3 jobs failing persistently, input files corrupted Savannah #113844
 - ◇ [OPEN] T1_IT_CNAF - 7 MC input files inaccessible at CNAF Savannah #113912
 - ◇ [OPEN] T1_ES_PIC - 2 files continually failing exports from PIC, Savannah # 113899. Update 20/4: files are available at PIC, but still failing transfers.
 - ◇ [OPEN] T1_ES_PIC - file with wrong checksum at PIC - probably corrupted before arriving at T1 Savannah # 113840
- ALICE reports [↗](#) - GENERAL INFORMATION: ALICE activity keeps on performing on Pass 1 reconstruction activities, together with 2 analysis trains. Still no MC production
 - ◆ T0 site
 - ◇ Both LCG-CE and CREAM-CE backends niely wortking for the reconstruction activities
 - ◆ T1 sites
 - ◇ CNAF: GGUS:57491 [↗](#) submitted this morning. ce01 (CREAM-CE) is failing at submission time. The ALICE expert at the site has been also informed via email. Production/analysis activities continuing through ce07 [Luca - going to reinstall failing CEs. Not yet sure if this can be done tomorrow but a matter of days. In the meantime the other CREAM CE is available.]
 - ◆ T2 sites
 - ◇ Madrid-T2: CREAM-CE is failing at submission time at the site. The analysis activities at the site are therefore stopped. The local ALICE expert has been informed and waiting for his action
- LHCb reports -
 - ◆ Data reconstruction, stripping and user analysis continuing.
 - ◆ Issues at the sites and services
 - ◇ T1 sites issue:
 - SARA: currently in downtime and banned for LHCb. Problem with application configuration. Waiting on new release of software.
 - ◇ T2 sites issue:
 - SAM jobs stalled at INFN-TRIESTE

Sites / Services round table:

- PIC - ntr

- ASGC - ntr
- FNAL - ntr
- NL-T1 - announcement: Monday 26th April we will have at SARA downtime for SRM and BDII. Already in GOCDB.
- RAL - ntr
- KIT - ntr: batch system still stable
- BNL - ntr
- CNAF - correct announcement yesterday on intervention for next week. 4h downtime on tape library next Monday was well as AT RISK on SAN. (All interconnections redundant). Tue - no intervention on DB, only on Wed. Took advice on how to minimize intervention. CREAM CE reinstall will be done this Friday.
- OSG - ntr
- IN2P3 - ntr

- CERN - meeting with ATLAS T0 next week to discuss how submission slowness issue can be addressed. Specific issue on Thursday night "dispatch issue" - only 1 job per host and not 8 - being followed up with platform.

- CERN DB - problem with 2 LCG applications: SAM and GRIDVIEW. Caused by user trying to run some tests on production database. User has been contacted.

AOB:

- GOCDB - is there a problem with the GOCDB sending downtime announcements too late? SARA announced downtime and we got email 6 hours after starting time. GGUS:57515 [ticket](#) submitted at the end of the meeting. Response from Rolf Rumler:

Downtime notifications basically depend on two components, the GOCDB (hosted at RAL) and the EGEE Operations portal (hosted at IN2P3-CC). In the case of the delayed downtime notification for SARA (the exact message is probably this one: "START of SARA-MATRIX/ SCHEDULED downtime (OUTAGE) [20-04-2010 06:00 to 20-04-2010 18:00 UTC]" it seems that an intervention on the underlying database of the Operations portal induced the late delivery. The portal team will try to inform sites and VOs more promptly in case of problems on their side.

Thursday

Attendance: local(Harry(chair), Patricia, Dirk, Simone, Nicolo, Gavin, Lola, Andrea, Maarten, Flavia, MariaDZ, Jan, Steve, Eva);remote(Jon(FNAL), Gang(ASGC), Michael(BNL), Gareth(RAL), Ronald(NL-T1), Gonzalo(PIC), Luca (CNAF), Jeremy(GridPP), Rob(OSG), Rolf(IN2P3)).

- AOB from yesterday - see above for information on downtime notification issue raised by LHCb.

Experiments round table:

- ATLAS reports - INFN-T1 SCRATCHDISK filled up. DDM operations is taking care of it. Luca queried which action ATLAS would prefer at CNAF. The reply was ATLAS will clean up their scratch space then decide if they want the space token increased. Local ATLAS contacts will hence get in touch later.

- CMS reports -

T0 Highlights: CASTOR tape degraded in SLS + large backlog of migration in T0EXPORT GGUS #57556 later escalated to an alarm ticket #57560. Jan explained that the backlog was due to a tape robot intervention

that took longer than expected but that the backlog was now clearing. Nicolo confirmed later that the ticket had been closed as solved.

T1 Highlights: 1) Data and MC reprocessing with CMSSW_3_5_7 running at FNAL. 2) FastSim running at FNAL. 3) Tails of Spring10 Monte Carlo processing running at T1s. 4) CNAF - Backfill workflow injected at full scale to test temporary solution for storage load issues seen during the last few weeks - no job failures so far. Large file timeouts in import from FNAL, GGUS #57570. Luca reported that their gridftp servers were not overloaded and requested more information which Nicolo will provide. 5) IN2P3 - Transient failure in SRM SAM test with CRL error, disappeared on next submission.

T2 highlights: 1) MC production running. 2) Test workflow running in CNAF T2 region using new WMS version.

- ALICE reports [☞](#) - GENERAL INFORMATION: There are two new MC cycles started up yesterday evening. In addition the analysis train activities are ongoing.

T0 site: No issues to report in terms of workload services. Deployment of the latest set of 2.1.9 patches for CASTOR (including xroot) has been announced yesterday evening. It will take place on the 26th of April at 9:00 (it should finish at 10:30 and should be a transparent action). More details available at: <https://twiki.cern.ch/twiki/bin/view/CASTORService/ChangesCASTORALICEUpdate2195>

T1 sites: 1) GGUS ticket opened yesterday concerning the bad behavior of one of the CREAM-CE at CNAF has been closed after checking the good behavior of the system (announced by the ALICE expert at the site). 2) All T1 sites performing MC production activities for the moment.

T2 sites: 1) Bari-T2: GGUS ticket: 57546. The local cREAM-CE seems to be out of production. Submission is not working and in addition, the query to the resource BDII does not respond. ALICE contact person at the site also contacted via email. 2) Manual intervention required at Bologna-T2. After the update of the site to AliEnv.218, the services were not properly restarted. Site back in production now. 3) Madrid-T2 site solved the problems reported yesterday and concerning the local CREAM-CE. Site is back in production.

- LHCb reports - Data reconstruction, stripping and user analysis continuing at low level. No new data taken for couple of days due to ongoing LHC tests.

Sites / Services round table:

- FNAL - 1. The 0% quality from PIC to FNAL is because we reject all the files after transfer because our checks find that the files have a bad checksum, compared to cms tmdb entry - Therefore all transfers fail. The files evidently for this block are bad at PIC. We suspended the block. (2). The 20 minute phedex file-stager downtime should be ignored - the phedex agents do not report reliably. The minimum amount of time before a shifter should report any phedex problem is 2 hours, as stated in the shift instructions.
- BNL: About 2 weeks ago BNL reported some transfer failures with small files traced to a gridftp problem with the dcache doors. A patch has now been received from dcache and deployed on one door overnight and shows an improvement so this will be deployed (transparently) on the remaining doors, expected to be completed by local midday.
- CERN SRM: Will be scheduling srm upgrades next week. srmpublic will be done on Tuesday morning so there will be a 30 minute gap in the CERN srm SAM monitoring tests.
- CERN myproxy: There is a pending hardware intervention on myproxy.cern.ch server to be made. This service does have a failover but is a critical service so should be at its best. Possibly next week, to be scheduled with the vendor, but no date has been fixed yet.

- CERN Databases: We lost connectivity at 2.30 today with the CMS online cluster in their pit. It seems they were running a GPN disconnection test of which we were not notified.

AOB:

Friday

Attendance: local(Harry(chair), Steve, Nicolo, Patricia, Lola, Maarten, Simone, Alessandro, Ueda, Eva, JanI, Andrea, Gavin);remote(Jon(FNAL), Michael(BNL), Gang(ASGC), Gareth(RAL), Onno(NL-T1), Andreas(KIT), Rob(OSG), Rolf(IN2P3), Jens(NDGF), Alessandro(CNAF)).

Experiments round table:

- ATLAS reports -

TAIWAN-LCG2_DATADISK, DESTINATION error during TRANSFER_PREPARATION phase: [NO_SPACE_LEFT] Transfer errors observed with source at TAIWAN-LCG2 for FT and T0 exports (T1-T1). Problem fixed after DBA at TW cleared the archive log. Gang reported that the archive log lifetime had already been reduced from 31 to 15 days and they would look again at this issue.

Reprocessing campaign started (supposed to last 2 weeks). 9 clouds are concerned - CERN will only be used if there is a problem elsewhere. Datavolume to reprocess is ~166 TB (164 datasets and 91k files). Estimated output volume of ~90TB. Merged ESD and merge AOD will be written on DATADISK and DATATAPE at Tier-1 where dataset was reprocessed. All other data types, including log will be written on DATADISK. Details in <https://twiki.cern.ch/twiki/bin/view/Atlas/ADCDataReproSpring2010>. AOD,DESD will be distributed within clouds according to the current merge AOD, DESD shares: http://atladcops.cern.ch:8000/drmon/repromon_TiersInfo.html A problem in job definition caused several reprocessing jobs to remain in WAITING state. The problem is being cured.

- CMS reports -

T0 Highlights: 1) Notification from IT-DB: GPN interruption at P5 --> Conditions streams were stopped Apr 22 15:00-16:00. 2) CASTORCMS 50% degraded on Apr 22 16:00 CEST - impacting transfers P5-->T0 and T0-->T1. At the same time large amount of user requests Remedy #CT0678533, correlated? Went away in ~1 hour, happened again on Apr 23 13:00 CEST., GGUS #57600. Jan Iven reported that there were in fact 20000 requests from a single user who has since been banned and they are looking into ways to protect against such abuse.

T1 Highlights: 1) Tails of Spring10 Monte Carlo processing running at T1s. 2) Some jobs failing for input file loss at CNAF, FNAL - will be invalidated. 3) CNAF - No site-related failures from Backfill workflow injected to test temporary solution for storage load issues seen during the last few weeks - now ramping down submission to free batch slots for new production requests. 4) Transfer quality on FNAL-CNAF improved after setting size-based timeout on channel. Luca reported that he was not happy that individual file transfer rates are not exceeding 3 MB/sec on this link (aggregate is 60 MB/sec) as they had seen better in the past. He will follow up. 5) CMS Squid cache clearing has been recommended by the experts at CNAF to increase efficiency.

T2 highlights: 1) MC production running. 2) Production workflow running in CNAF T2 region using new WMS version. 3) Production at T2_EE_Estonia on hold while site migrates to new storage. 4) Merge failures at T2_FR_IPHC, now stable. 5) T2_CH_CSCS - SAM CE tests: Maradona errors.

- ALICE reports -

GENERAL INFORMATION: Pass2 reconstruction activities started at T1 sites (cosmics). In addition there are four analysis trains together with two MC cycles in production

T0 site: 1) Pass1 reconstruction activities ongoing at the T0 with no incidents to report. Both CREAM-CE performing well. 2) CAF ALICE nodes have been put in maintenance status yesterday afternoon to apply some kernel upgrades. 3) The access to the software area of ALICE at the T0 is very slow. A remedy ticket has been opened. The AFS team have reported that one or more hosts are blocking callbacks from the ALICE afs file server which hence accumulate before they timeout. This could be a return path firewall issue for example. Under investigation. Meanwhile ALIICE have defined a new software area and made symlinks to it.

T1 sites: CNAF-T1: ce01-lcg.cr.cnaf.infn.it (CREAM-CE) still failing. When the issue was reported on Wednesday (via GGUS ticket 57491), the ALICE contact person confirmed several hours later, the system was back in production. This was confirmed by ALICE who tested the system. However, the CREAM-CE is failing again this morning (new ticket 57585). The 2nd CREAM-CE ce07-lcg.cr.cnaf.infn.it has been put in production. CNAF then reported that the CREAM-CE has meanwhile been reinstalled on new hardware so ALICE will test this.

T2 sites: No news concerning the GGUS ticket (57546) submitted yesterday concerning the CREAM-CE system in Bari.

- LHCb reports - Data reconstruction, stripping and user analysis continuing at low level. No new data taken for couple of days due to ongoing LHC tests.

Sites / Services round table:

- BNL: As announced yesterday they have now patched all their gridftp dcache doors and run a successful large scale test with small (~ 1MB) files. Transfers to T1 and T2 sites have hence now been reverted to use gridftp2.
- IN2P3: The GOCDB downtime notifier is itself down following an Oracle intervention which crashed the DB server. Under investigation but not expected to be back before Monday.
- CNAF: There will be an at-risk hardware intervention on a core switch on Monday. Also a tape library will be down.
- OSG: A top priority problem ticket was opened last night from a mid-west T2 site which on investigation looked like a standard type of job failure. They plan to review their policy of how ticket priorities are allowed to be defined to avoid such low priority items receiving higher priority. Michael added that in the ATLAS elogger this problem had correctly appeared as 'less urgent'.
- CERN-PROD: Now that the last WLCG FTS server has been upgraded to the 2.* series and all FTS clients should be using delegation, we would like to decommission the `myproxy-fts.cern.ch` service as soon as possible. Can we agree this with the experiments and plan a date? See: [MyProxyFTSRetire](#)

AOB: Stable LHC beam is planned for most of this coming weekend before a 3-day technical stop scheduled from 08.00 Monday to 20.00 Wednesday next week.

-- JamieShiers - 19-Apr-2010

This topic: [LCG > WLCGDailyMeetingsWeek100419](#)

Topic revision: r21 - 2010-04-23 - HarryRenshall



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

WLCGDailyMeetingsWeek100419 < LCG < TWiki

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