

Table of Contents

Week of 100329	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.....	6
Thursday.....	9
Friday.....	12

Week of 100329

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 board	M/W PPSCoordinationWorkLog	WLCG Baseline Versions	WLCG Blogs Blogs	GgusInformation	Sharepoint site site - Cooldown Status Status - News News

Monday:

Attendance: local(Jan, Roberto, Patricia, Dirk, Harry, Jamie, Maria, Jaroslava Alessandro, Miguel, Andrea, Stephane, Jean-Philippe, Eva, Nilo, Miguel, Edoardo, Gavin, MariaDZ, Malik, Julia);remote(Michael/BNL, Stefano Zani (INFN CNAF), Jon/FNAL, Gonzalo/PIC, Rolf/IN2P3, Pepe/CMS, Rob/OSG, Jeremy/GridPP, Angela/KIT, Gang/ASGC, Gareth/RAL, Ronald/NL-T1, Vera/NDGF).

Experiments round table:

- ATLAS reports -
 1. ALARM ticket to CERN: CASTOR diskserver problem.
https://gus.fzk.de/ws/ticket_info.php?ticket=56793 -- FIXED
 - ◇ Inaccessible files (timeouts) on Tier-0 T0MERGE CASTOR pool due to high contention on one diskserver (1xfssl4004)
 - ◇ does CASTOR ops have an automatic way to catch those kind of issues? [Miguel - first reason given for the problem but in fact not contention - disk server got stuck. How can we automatically catch this? Saw all monitoring were working fine - no monitoring detected the issue. Don't know why server was stuck. A reboot fixed it. - Ale - please update alarm ticket. Miguel - 2nd update has this info Ale - observed same issue last night. Not so urgent so no alarm ticket (on T0merge) Miguel - maybe some timeouts. Timeout is not the same problem. This was a server problem.
 2. Central Catalog /var full -- FIXED [script to check this deployed - quattor template missing in this machine.]
 3. ALARM ticket to CERN due to DB monitoring red.

https://gus.fzk.de/ws/ticket_info.php?ticket=56790. Problem NOT due to DB but PVSS2COOL application, ATLAS DCS people fixed the issue.

◇ thanks to DB people: quick reaction and understanding of the situation (no problem on DBs)

◇ procedures to be improved, already discussed with ATLAS T0 and ATLAS DB people

4. SARA FTS got stucked (GGUS #56796) - solved at 1am today, FTS transfers from SARA failing (GGUS #56794) - still in progress, problem with disk pools. [Ale - SARA just had another small glitch on LFC - 1.5 hours - ticket submitted]
5. SRM ATLAS team ticket top priority this morning - asked for feedback from VO but nothing so far. Will downgrade priority.

• CMS reports -

◆ T0 Highlights

◇ Plans: preparing for 3.5 TeV collisions

◆ T1 Highlights:

◇ Backfill jobs were running at all sites. Now reducing the load because we've started preproduction.

◇ Some backfill outputs in Feb./Mar. went to incorrect LFN areas at some T1s. DataOps/T1s checking and deleting those backfill outputs.

◆ T2 highlights

◇ MC production ongoing

◇ T2_BR_UERJ was auto-deleting all files after transfers; fixed in the PhEDEx configuration. Some lost datasets have been subscribed again, and we plan to run a full consistency check when the re-transfers end.

◇ Transfers from PIC to Caltech are failing; PIC admins notified. In fact >20 links to Caltech are failing atm (<https://savannah.cern.ch/support/index.php?113542>).

◇ DDT-team is doing a round of T2-T2 data transfer links commissioning (<https://savannah.cern.ch/support/?113504>).

◆ 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.[Miguel - thought this problem was fixed? Pepe - will check]

◇ [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%.

◇ [CLOSED] T1_ES_PIC - Savannah #113425 - Consistency check failure for custodial data (3600 files) hosted at PIC. The PhEDEx consistency agent was reconfigured to preload libpdcap.so and all data checked is consistent. Ticket Closed.

◆ Weekly-scope Operations plan

◇ [Data Ops]

- Tier-0: process incoming
- Tier-1: prompt skimming, MC production and rereco, as some redigi workflows
- Tier-2: Finish current requests; wait for new production requests

◇ [Facilities Ops]

- New SAM and JobRobot Datasets at T1s and T2s. To move to use newer CMSSW versions for SAM and JobRobot.
- Weekly reports on T1 Production Performance and Resource Utilization in preparation.

· Next Thursday, T2 Support Meeting. T2s invited to join: T2_FI_HIP, T2_RU_IHEP, T2_BE_IHE, T2_RU_ITEP, T2_CN_Beijing, T2_RU_PNPI.

- ALICE reports [↗](#) - GENERAL INFORMATION: Two MC cycles have been running during the weekend with no incidents to report. Peaks over 23K jobs have been achieved.
 - ◆ T0 site
 - ◇ Responsible of the still ongoing Pass1 reconstruction jobs, good behavior observed during the weekend
 - ◇ New VOBOX entering in production (voalice07). Old VOBOX included in the xrootd partition.
 - ◆ T1 sites
 - ◇ Engaged of the MC production during the weekend mostly executed through RAL CNAF, CCIN2P3 and FZK
 - ◆ T2 sites
 - ◇ 2nd VOBOX at SPbSU entering production this morning
 - ◇ RRC-KI unblocked during the weekend
 - ◇ Problem with the local configuration at the VOBOX@PNPI. Solved on Saturday morning, the system wntered in production immediately after.
- LHCb reports - xrootd first tests at CERN. 100 user jobs and this is it today!
 - ◆ T0 sites issues:
 - ◇ CASTOR: The LHCb data written to lhcbraw service has not been migrated for several days (GGUS: 56795).
 - ◇ A vbox (volhcb26) wasn't reachable in the weekend. Remedy ticket to vbox-support. (Remedy Ticket CT0000000671634)
 - ◇ Jan (following a request on Friday) has announced that the castorlhcb head nodes are now running (the latest version of) SSL-enabled xrootd. LHCb want to stress that SRM should provide a consistent tURL for xroot for accessing data at various service classes. Preliminary tests show the problem to be still there (file open and then not read) and a close interaction with xrootd developers has been triggered.
 - ◆ T1 sites issues:
 - ◇ SARA: LFC read-only replica was unresponsive since a while (disappeared from SAMDB too) and this was at the origin of the root group-id reported last week at NIKHEF. Oracle RAC issue the reason of the problem.
 - ◇ CNAF : ConditionDB intervention: agreed for tomorrow 9-13 CET
 - ◆ T2 sites issues:
 - ◇ egee.fesb.hr shared area issue
 - ◇ ITPA-LCG2 library misisng
 - ◇ IN2P3-LPC SAM voms tests failing

Sites / Services round table:

- BNL - ntr
- INFN - ntr apart from what was explained by LHCb
- PIC - ntr. Q on situation that PIC is banned. Roberto - core developers working on it. Will ask for more info.
- IN2P3 - ntr
- KIT - in the morning network problems - now fixed. Also maintenance on the network without complications
- RAL - had a small intervention to update license keys for LSF scheduler in CASTOR - went ok. Some problems with migration to tape for CMS over w/e but now under control
- ASGC - phedex agents for CMS upgrade last Friday
- NL-T1 - "wonderful w/e" - Oracle RAC experiencing difficulties - LFCs for LHCb and ATLAS and probably also FTS. Oracle RAC restarted this morning - now aok. Issue copying files SARA to IN2P3

- configuration error of corresponding pool groups. Now fixed. All working as it should
- FNAL - ntr
- OSG - tomorrow will be doing updates of core services incl topology database. MyOSG and ticketing exchange. Should be no effect. Future GGUS change April 14 on submitting tickets to OSG
- GridPP - ntr
- NDGF - small LFC restart this morning.

- CERN DB - 2 h/w failures over w/e. ALICE Online DB - a controller broke and a disk array evicted by clusterware. DB up all the time and disk array added back. LHCb online DB - router problem - affected two instances of DB. Listener was down due to failure on router. These two instances were not receiving any connection. Restarted manually.

AOB:

- GGUS & old ATLAS grid alarms list: time to review names. Yesterday we had two alarm tickets - would be useful if these were also directed to atlas-grid-alarms. MariaDZ: authorized alarmers are members of VO in VOMS. Twiki on alarms should say to put the relevant list in cc: - take offline!

Tuesday:

Attendance: local(Maarten, Maria, Jamie, Gavin, Alessandro, Stephane, Roberto, Patricia, Jean-Philippe, Malik, MariaDZ, Eva, Nilo, Manuel, Jan, Harry, Miguel);remote(Stefano Zani (INFN TIER1), Jon/FNAL, Rob/OSG, Jeremy/GridPP, Rolf/IN2P3, Ronald Starink (NL-T1)), (Tiju Idiculla (RAL)), (Pepe Flix (CMS) (Andreas Heiss), Michael/BNL).

LHC news: Beams collided at 7 TeV in the LHC at 13:06 CEST! Press Release [↗](#) - Twitter site [↗](#)

Experiments round table:

- ATLAS reports -
 1. ATLAS alarm ticket for CASTOR!
 2. Run 152166 collected to T1s and T2s. Events are reconstructed and exported to T1s and T2s. Everything running smoothly.
 3. Replication of RAW to TRIUMF-LCG2_DATADISK was missing until this morning (OK to TAPE). Replication will be done manually after Media Day.
- CMS reports -
 - ◆ T0 Highlights
 - ◇ Collecting 3.5 TeV collisions data atm! 😊
 - ◆ T1 Highlights:
 - ◇ Started preproduction.
 - ◇ 7TeV MinBias DataTaking MC run at T1 (FNAL)
 - ◆ 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. [Jan - according to Miguel this should have been closed. Pepe- no, apparently not. That was another ticket!]
 - ◇ [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%. [Gav - still open with platform]

Tuesday:

- ◊ [OPEN] T1_ES_PIC - file access problem at PIC: repeated errors trying to access one file during the 355 rereco preproduction - Savannah #113582.
- ◆ T2 highlights
 - ◊ MC production ongoing
 - ◊ T2_IN_TIFR ready to be included back into production.
 - ◊ DDT-team is doing a round of T2-T2 data transfer links commissioning (<https://savannah.cern.ch/support/?113504>).
- ◆ All
 - ◊ CMSSW_3_5_5 is broken and will be deprecated. CMSSW_3_5_6 has been promptly installed (yesterday night) on all T1s and on a lot of T2s, atm.
- ◆ [Data Ops]
 - ◊ Tier-0: process incoming
 - ◊ Tier-1: prompt skimming, MC production and rereco, as some redigi workflows
 - ◊ Tier-2: Finish current requests; wait for new production requests
- ◆ [Facilities Ops]
 - ◊ New SAM and JobRobot Datasets at T1s and T2s. To move to use newer CMSSW versions for SAM and JobRobot.
 - ◊ Weekly reports on T1 Production Performance and Resource Utilization in preparation.
 - ◊ Next Thursday, T2 Support Meeting. T2s invited to join: T2_FI_HIP, T2_RU_IHEP, T2_BE_IHE, T2_RU ITEP, T2_CN_Beijing, T2_RU_PNPI.
- ALICE reports [↗](#) - ALICE is also taking data!

GENERAL INFORMATION: Still some few remaining jobs from the MC production running in the system. Based on the new site publication structure available in AliEn2.18, the publication of all sites in ML has been modified yesterday night. All sites are now represented by a single spot in ML which represents all VOBOXES available at the site (independently of the number of voboxes available per site). The spots do not correspond to the single VOBOX status anymore, but they represent the global site status

The number of running jobs appearing in ML while clicking on the site spot corresponds to the total number of running jobs at the site. Internal pages of ML will keep on showing the status of each VOBOX as before

 - ◆ T0 site
 - ◊ Systems cross-checked this morning before data taking. all voboxes performing well
 - ◆ T1 sites
 - ◊ Extra checks performed this morning with no incidents to report
 - ◆ T2 sites
 - ◊ Hiroshima T2: Problems with the local CREAM-CE. Alice agents run for a very short time before dying without catching any real job from the central queue. The submission mode at this site has been modified to catch the standard output/error which will provide more information. Site admin has been asked to ensure the startup of the gridftp server at the local VOBOX
 - ◊ Cape Town: Site ready to enter in production. Small issues with the local batch system are preventing the full setup in AliEn
 - ◊ SpbSU and ITEP (Russia): These 2 sites are now running AliEnv.2.18. The corresponding configuration changes in LDAP have been applied to set the local voboxes in failover mode. no issues observed after the sites updates
- LHCb reports -
 - ◆ Started to receive COLLISION10 type data promptly migrated to tape, checksum checked and registered in LFC.
 - ◆ Actively working on commissioning real data workflows affected by some problems at application level (a patch going to be rolled out today by core developers)
 - ◆ xrootd: tests at CERN continuing: no extra variable had to be defined to manage the access of

files via xrootd that has been apparently properly configured now.

- ◆ T0 sites issues:
 - ◇ none
- ◆ T1 sites issues:
 - ◇ PIC: downtime (announced) causing some problems due to some DIRAC critical services hosted there.
 - ◇ CNAF: ConditionDB intervention (until 13:00 CET) [Stephane - discussed with ATLAS and will be done end April at next LHC shutdown]

Sites / Services round table:

- CNAF - ntr: cond DB upgrade ended ok;
- FNAL - 1) Congratulations - fantastic watching things 2) getting several questions about CPU accounting data - verified it and believe it to be correct. Verifying again... 3) OSG: started to do ticket exchange footprint at Indiana and Remedy at FNAL. Some test tickets but also some ALARM. Alarm tickets were misdirected and if real would not have been responded to - something seriously wrong! Rob - long conversation between OSG & GGUS developers. A development ticket in GGUS got transferred to production ticketing system at FNAL. Long e-mail thread - should not have got alarm ticket in prod system.
- OSG - nothing to add! In middle of service update - updating topology database and ticketing database.
- IN2P3 - ntr
- BNL - ntr
- NL-T1 - at NIKHEF since yesterday evening a failing diskserver. Lost connection via infiniband to storage. In R/O mode - no new data will be lost. Working on bringing it back but no ETA.
- RAL - ntr
- ASGC - vendor started H/A - adding one more mechanical hand to tape system. Will be tested in a day or so. Data access to tape cartridges will be affected during this period. Disk access ok - should not be a big problem - only when jobs try to access files that are not already on disk. AT RISK in GOCDB
- KIT - ntr
- GridPP - ntr

- CERN DB - high load on ATLAS offline DB. Ticket opened saying DB down - which was not true. New connections were not permitted due to high load. Being investigated. Ticket will be updated.
- GGUS/OSG - Kyle will attend for OSG and represent any GGUS-related issues in coming days. Details in <https://savannah.cern.ch/support/?109779#comment43>
- CERN LSF/web application - allows experiments to set shares. Fixed lunchtime. Affected CMS.
- CERN Storage - alarm ticket just before meeting appears to be short reboots of diskserver "not really welcome today". Ale - ticket update "glitch" after 1500s files accessible again.

AOB:

Wednesday

Attendance: local(Gavin, Maarten, Jamie, Maria, Andrea, Steve, Jan, Kasia, Miguel, Ale, Jean-Philippe, Patricia, Edoardo, Malik, Roberto);remote(Stefano Zani (INFN TIER1), Michael/BNL, Jon/FNAL, Angela/KIT, Rolf/IN2P3, Gang/ASGC, Kyle/OSG, Jeremy/GridPP, Onno/NL-T1, Tiju/RAL, Greig Cowan)

Experiments round table:

- ATLAS reports -
 - ◆ 1 GGUS Alarm ticket GGUS:56848 to CERN :

- ◇ ATLAS T0 could not retrieve data from t0atlas. It was temporary. ATLAS automatically retried later.
- ◇ Castor team explained : 'there was a fast hardware intervention on a couple of servers' [Standard operations that can happen - files inaccessible for 20'. ATLAS retries worked ok. Issue understood and perfectly fine!]
- ◆ 2a GGUS Team ticket GGUS:56878 [↗](#) to CERN :
 - ◇ Temporary srm errors. Request informations. Is it related to concurrent accesses ?
- ◆ 2b GGUS Team ticket GGUS:56891 [↗](#) to CNAF :
 - ◇ Problem to access data. Solved.
 - ◇ CNAF team responded : problem on StoRM end-point (specifically the underline GPFS)
- ◆ 2c
 - ◇ GGUS Team ticket GGUS:56849 [↗](#) to TRIUMF : Files unaccessible (TRIUMF confirmed their existence of files on the storage). Still underway.
- ◆ 2d
 - ◇ GGUS Team ticket GGUS:56894 [↗](#) to SARA : Cannot write on ATLASCRATCHDISK. Solved.
 - ◇ SARA responded 'Configuration error which is fixed'
- ◆ 3a
 - ◇ ELog not available from 3 AM to 7 AM (<https://remedy01.cern.ch/cgi-bin/consult.cgi?caseid=CT0000000672570&email=suen@phys.s>).
 - ◇ Should have been restarted automatically.
- ◆ 3b
 - ◇ PhysDB issue on 30 March (<https://cern.ch/helpdesk/problem/CT672281&email=atlddmsh@mail.cern.ch>) : Problem of communication within ATLAS.
 - ◇ According to Atlas DBA, known issue close to be solved (information from yesterday late afternoon).
 - ◇ Ticket was submitted to physdb (yesterday morning). Atlas elog was filled when problem occurred and solved (few minutes later) but ticket was not updated.
 - ◇ Contrary to yesterday's report this was not "high load". Procedure updated. Monitoring has a glitch (ATLAS monitoring) that will be fixed. Kate - a new procedure was deployed today - sometimes queries (specific monitoring tasks) that give 15' high load and new connections not possible during this time. Maria - a written clarification please!
- ◆ 3c
 - ◇ SLS not always reporting availability of ATLASDATADISK at CERN. Ale Di Girolamo is investigating
- ◆ Jan - prototype SIRs for ALARM ticket and TEAM ticket - the latter seems to be stuck. Two short-lived SRM failures yesterday, 1st affected only ATLAS.
- CMS reports -
 - ◆ T0 Highlights
 - ◇ Processing collision data
 - ◇ DEFAULT service class red in SLS, seemingly due to a user who submitted 10K jobs [Maarten - should this have worked? Or a clear mistake? Miguel - investigated with Peter Kreuzer. User doing a lot of replication from data in CMS CAF to default. CMS CAF for "special users". At this moment users can do this and CMS is discussing if it should be allowed or not. User no longer allowed to copy data from CMS CAF to default.]
 - ◆ T1 Highlights:
 - ◇ Running prompt skimming
 - ◇ Running backfill jobs

- ◊ FNAL running reconstruction for MinBias MC 7 TeV
- ◆ 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%.
 - ◊ [OPEN] T1_ES_PIC - file access problem at PIC: repeated errors trying to access one file during the 355 rereco preproduction - Savannah #113582.
- ◆ T2 highlights
 - ◊ MC production ongoing
 - ◊ SAM test failures in Estonia due to a power cut
 - ◊ Investigating transfer failures from several T2 to Wisconsin (Savannah #113505)
 - ◊ Transfers from IN2P3 to Bristol failed as a Lyon DN is not authorized in the RAL Myproxy (Savannah #113613)
- ALICE reports [↗](#) - GENERAL INFORMATION: Following the reports of the ALICE shifters since yesterday afternoon, there are no issues associated to the Grid or IT services at any site. 7TeV runs are being successfully reconstructed and in addition the MC production and the analysis train are also ongoing.
 - ◆ T0 site
 - ◊ Reconstruction of raw data ongoing. All T0 services behaving ok. The replication of the just recorded data to all T1 sites will (most probably) begin tomorrow
 - ◆ T1 sites
 - ◊ Responsible of MC production, still ongoing Pass6 reconstruction tasks and analysis tasks. All T1 sites are in production. It is worth to mention the case of NIKHEF which has been out of production for several weeks and it is back in production now.
 - ◆ T2 sites
 - ◊ Hiroshima-T2: Issues with the local CREAM-CE still opened. (agents die immediately after submission to WN)
 - ◊ Birmingham: Putting the site in CREAM-mode by today
 - ◊ Poznan: Issues with the local user proxy. Still under investigation
- LHCb reports - The DSTs of yesterday's data available in the book-keeping under: Lhcb -> Collision10 -> Beam3500GeV-VeloOpen-MagDown -> Real Data + RecoDST-2010-01 -> DST. These are complete DSTs without any stripping. The stripping will be deployed in the coming days, once some crashes in DaVinci are fixed .
 - ◆ T0 sites issues:
 - ◊ FTS export service: ALARM ticket issued - GGUS:56880 [↗](#) convinced that they just swapped the old (fts-t0.export.cern.ch) with the new (fts22-t0-export.cern.ch, it was pilot before) by changing the alias. (Migration to FTS22 Procedure). This "false" problem also offered the possibility to exercise the procedures for alarming problems on best effort supported services like FTS that proven to be OK.
 - ◆ T1 sites issues:
 - ◊ pic: back from the downtime yesterday, has been re-integrated in the lhcb production mask with some latency this morning due to CIC portal not sending the notification of the END downtime. CIC people report their portal overloaded due to a migration to GOC 4.
 - ◊ NIKHEF-SARA: issue with dCache accessing data. (GGUS:56909 [↗](#)) [Onno - still under investigation - involves dCache at SARA, jobs at NIKHEF.]

- ◆ T2 sites issues:
 - ◇ CBPF: Shared area issue

Sites / Services round table:

- CNAF - ntr
- BNL - ntr
- FNAL - received high priority request from CMS at 5pm yesterday to replicate 2300 pileup files to 5-6 pools each - in process of finishing that now.
- KIT - news about disk problems mainly affecting ATLAS. Firmware mismatch between disk enclosure and storage controller so failover did not work - now corrected
- IN2P3 - ntr
- ASGC - FTM / FTS online; SLC4 WNs migrated to SLC5 all T1 WNs SLC5 and onlined this week. H/A construction - delayed to tomorrow. AT RISK in GOCDB
- OSG - report that service upgrades went ok without incident yesterday
- GridPP - ntr
- NL-T1 - nothing to add.
- RAL - ntr

- CERN DB - high load on ATLAS: 08:00 or 16:00. First was user mon 2nd was closed by Florbella.

AOB:

- Recent GGUS alarm tickets: ATLAS PVSS2COOL [↗](#) timeline issues
- LHC schedule: overnight refill and ramp to 3.5 TeV. Outlook for end week and w/e. Increase intensity and test 25ns scheme. As of Friday attempt for collisions and physics until end of weekend.

Thursday

Attendance: local(Julia, Malik, Maarten, Gavin, Patricia, Andrea, Kate, Dirk, Jean-Philippe, Alessandro, Ueda, Stephane, Jan, Roberto, Jamie, Miguel);remote(Jon/FNAL, Gang/ASGC, Angela/KIT, Rolf/IN2P3, Kyle/OSG, Vera/NDGF, Gareth/RAL, Stefano/CNAF, Ronald/NL-T1, Jeremy/GridPP, Reda/TRIUMF).

Experiments round table:

- ATLAS reports -
 - ◆ 0. ATLAS collected data during 6 hours last night (Statistic increased by factor 4)
 - ◆ 1.a GGUS Team ticket GGUS:56936 [↗](#) to TAIWAN:
 - ◇ SE not accessible during few hours
 - ◇ Reason : 'Our srm interface is down due to database space is full'
 - ◇ FTS servers was also stuck
 - ◇ Now everything OK
 - ◆ 1.b GGUS Team ticket GGUS:56946 [↗](#) to CERN :
 - ◇ srm failures (temporary issue). Response: 'we were again running out of frontend threads'. Problematic periods provided within ticket
 - ◆ 2.a INFN-NAPOLI-ATLAS SE not working during few hours (GGUS Team # 56944). No entry in Gocdb.
 - ◆ 2.b LIP-COIMBRA : 1389 lost files (Savannah : 65162) . Reason : 'Side effect of raid-controller problem'. Files being replicated from Grid when possible
- CMS reports -
 - ◆ T0 Highlights

- ◇ CASTORCMS fully recovered after used was banned. He was triggering disk-to-disk copies from the CMSCAF to the DEFAULT service class
- ◇ At around 22 UTC all Tier-1 sites experienced a few "SOURCE error during TRANSFER_PREPARATION phase: [GENERAL_FAILURE] Timed out" with transfers from the Tier-0. It was correlated to an increase in the number of active transfers in the T0EXPORT service class to ~ 400, which should not be a problem in principle. [GGUS ticket request]
- ◆ T1 Highlights:
 - ◇ Backfill jobs affected by the WMS-ICE bug; CREAM CEs will be excluded from production activities until the new WMS release is available
 - ◇ Some failures in backfill jobs at ASGC, probably due to too large output files
 - ◇ PIC is fully back in production after the scheduled downtime
 - ◇ ASGC shows several failures in SAM tests which use the storage and transfers have a very bad quality. The cause is the SRM interface being down due to a full database space. Progress in tracked in GGUS:56936 (submitted by ATLAS) an in Savannah #sr 113643. The problem is supposedly fixed, to be verified
- ◆ 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%.
 - ◇ [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_ASGC - Transfer Errors from T0 -> ASGC - Savannah #sr 113643
- ◆ T2 highlights
 - ◇ MC production ongoing
 - ◇ SAM test failures in Estonia still persisting: the CMS software area is unavailable (Savannah #sr 113577)
 - ◇ Unscheduled downtime in Florida related to dCache
 - ◇ A few hours of unscheduled downtime at Caltech due to network problems
- ALICE reports - GENERAL INFORMATION: Successful data taking last night. All 7TeV runs have been successfully reconstructed but one (experts looking at it now, no indications this issue can be associated to any issue with Grid/IT services). In terms of the current MC production, all 7TeV simulation is done (6Mio events with two different generator settings). What is running today in the system are the remaining jobs of the last night raw reconstruction, analysis trains and user analysis. Finally and in terms of T0-T1 raw data transfers, they are still pending [Q on end-points in use? To be confirmed at task force meeting]
 - ◆ T0 site
 - ◇ Performing the last reconstruction jobs (raw data) successfully
 - ◆ T1 sites
 - ◇ Finishing the remaining analysis jobs
 - ◆ T2 sites
 - ◇ GRIF_IRFU (France): New CREAM-CE setup at the site. Problems with the information provider of the system. It has been solved this morning. Checking the system before putting it in production
 - ◇ Hiroshima-T2: Still observing problem with the local CREAM-CE: The gridftp service at the VOBOX is not able to catch the output sandbox of the agents. In addition the CREAM-DB showed several problems yesterday night losing all the

submitted jobs

◇ RRC-KI, ITEP and SpbSU have been blocked this morning at the ALICE central services in order to perform several Alienv2.18 tests.

• LHCb reports -

- ◆ "We are very proud to announce that we collected 1.3 million collision events last night in nominal conditions, i.e. Velo fully closed and all detectors IN. The runs 69353,54,55 have been sent out for production and should become available for analysis later today"(O.Callot)
- ◆ Reconstruction jobs will be running this afternoon and over the weekend.
 - ◇ T0 sites issues:
 - For the last hour the LHCb Castor stager has been blocked by a inadvertent DOS attack. Any jobs that have attempted to access files in the last hour may have seen problems. [Situation back to normal now. Jan - 40K migration requests at same time - overwhelmed staging. Jobs largely killed. Activity legitimate but a bit much in one go! Small chunks would help. Roberto - size of files staged also an issue - histogram files]
 - ◇ T1 sites issues:
 - CNAF: many users complain about the ConditionDB unreachable because the connection string has changed. This causes their jobs to time out. Fix has been applied LHCbApp side.
 - NL-T1: Ongoing data access issues with gsidcap for user analysis jobs.
 - ◇ T2 sites issues:
 - GRISU-UNINA: shared area issue

Sites / Services round table:

- TRIUMF - ntr
- BNL - ntr
- FNAL - ntr
- ASGC - SRM i/f down 04:30 as disk space for SRM DB full. Added one more diskserver. Also crashed job manager. Everything ok 09:30. Transfers to ASGC ran fine up to now. Closed CMS Savannah ticket.
- KIT - ntr
- IN2P3 - pre announcement of sched AT RISK 13 April for installation of 3rd robot library. Access to tapes will not work all day.
- OSG - ntr
- NDGF - ntr
- CNAF - ntr
- NL-T1 - NIKHEF lost 2 disk servers this morning. In batch put last week in R/O mode. Got engineer on site. Managed to restore 1. Looking at 2 others. Escalated to h/w vendor and actions to improve stability.
- GridPP - ntr
- RAL - seeing some load issues on top BDII

- CERN DB
 - ◆ ATLAS issue, Tuesday March 30th, 8:00
 - ◇ Atlas experienced a transient spike of high load which has affected also the time needed to establish new sessions to the DB, with a few timeouts reported, for a duration of about 10 minutes. At the DB level the high load in question manifested itself as concurrency contention across the atlas offline cluster. Upon analysis, the root cause of the issue was attributed to the combination of the high activity of atlas_t0 service and the running of a custom job used to clean-up the DB audit table (auditmon). Sporadic occurrences of this issue had already been observed in previous occasion and mitigated with a change of the preferred node where the auditmon was

running. To provide a more stable fix for this issue, given that a change of node where the job runs has proven to be not sufficient to cover all cases, a new version of the auditmon job in question has been developed and tested. It has been deployed in production and scheduled to run for the first time on Thursday morning.

- ◆ ATLAS issue, Tuesday March 30th, 16:40 (Q" who was contact? A: Florbella -> T0 services)
 - ◇ A transient high load of ATLR node N.2 was observed, due to high utilization of the node, which could potentially affect service time of DB services on that node. Upon investigation the issue was correlated to the fact that a certain fraction of the Atlas cool jobs were using the service atlas_t0 (which runs on node N.2 only and is intended for atlas_t0 accounts only) instead of atlas_coolprod (which is load balanced and intended for cool reader activity). This fact has been discussed with Atlas who has since changed the connect string where appropriate.
- ◆ ATLAS PVSS alarm ticket - timeline understood: actions should (preferably) be logged in GGUS ticket (there were actions between the operator alarm and first response several hours later).

- CERN CASTOR+SRM: tried a new combo of f/s and o/s:

AOB:

Friday

NO MEETING! - CERN CLOSED

-- HarryRenshall - 26-Mar-2010

This topic: LCG > WLCGDailyMeetingsWeek100329

Topic revision: r13 - 2010-04-01 - JamieShiers



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