

Table of Contents

Week of 100607.....	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.....	5
Thursday.....	7
Friday.....	10

Week of 100607

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 Blogs	GgusInformation	Sharepoint site - Cooldown Status - News

Monday:

Attendance: local(Jamie, Miguel, Nilo, Eva, Ewan, Stephen, Roberto, Maarten, Jean-Philippe, Ale, Dirk, Patricia, Simone, MariaDZ);remote(Jon, Gonzalo, Rolf, Vera, John Kelly (RAL), Alexander Verkooijen (NL-T1), Rob, Angela, Joel, Brian).

Experiments round table:

- ATLAS reports -
 - ◆ NIKHEF-ELPROD_DATADISK showing transfer errors with NO SPACE LEFT (>> 100TB free). Probably it refers to disk space in FTS.
 - ◆ Reprocessing of express stream run at CERN (for the first time) over the weekend. Ultimately successful, but details of the fair share at CERN should be understood (contention between production and analysis jobs)
 - ◆ Overall, a quiet weekend
- CMS reports -
 - ◆ T0 Highlights
 - ◇ Processing normally (some delay on Saturday with a purely CMS filesystem issue - limit on # directories, every two months we hit this limit!)
 - ◆ T1 Highlights
 - ◇ Processing normally
 - ◆ T2 Highlights
 - ◇ MC production as usual

- ◆ Can now open ALARM tickets, opened GGUS:58764 on the 3rd June, ability enabled sometime between 9am and 1pm today. MariaDZ - was a misunderstanding on procedure. There was a manual step: data is automatically extracted very night from VOMS. However, a manual step was not done due to holidays in DE. A ticket has been opened requesting that this step is automated too.
- ALICE reports - GENERAL INFORMATION: Two raw data reconstruction activities ongoing: Pass 1 reconstruction and muon data calibration. In addition there are four analysis trains. There were no further MC cycles during the last weekend. No remarkable issues reported during the weekend
 - ◆ T0 site
 - ◇ Voview information provided by the local CREAM-CE systems announces a large amount of Alice jobs in status waiting. Submitted agents are not entering in running state (not necessarily a problem)
 - ◆ T1 sites
 - ◇ FZK: Question to the site admin: The ALICE queue used at this site (cream-3-fzk.gridka.de:8443/cream-pbs-aliceXL) is announcing 0 free CPUs, with 4 jobs running and 1 waiting. In addition ALICE has not been running at this site during the weekend. Possible issue with the reported available CPUs (?) Angela - where are you looking? I see >6k job slots free. Patricia - this was situation 1 hour ago. Angela - this number should not change.
 - ◇ CNAF: Error made by ALICE during w/e - Unblocking the site in LDAP this morning
 - ◆ T2 sites
 - ◇ Trujillo (Spain): Bad performance of the local batch system. Reported to the site admin this morning
 - ◇ UCT (Cape Town): Setup of the local xrootd SE system during the weekend in collaboration with the site admins
- LHCb reports - Reprocessing (04) of old data is almost done.
 - ◆ T0 site issues:
 - ◇ CERN: jobs killed by the batch system. Remedy Ticket open under investigation.
 - ◆ T1 site issues:
 - ◇ CNAF StoRM: user reporting problems accessing data on GPFS (GGUS:58794). Problem was on static gridmap file not recognizing a particular LHCb FQAN
 - ◇ CNAF: RAL-CNAF one file systematically failing. Glitch of the network. (GGUS:58821)
 - ◇ IN2p3 : there are many jobs killed ending up in a wrong (too short) queue. Under investigation. Same it's true at other sites as well (RAL)
 - ◇ SARA - 1 dcap port (might) need to be restarted: GGUS:58838
 - ◆ T2 sites issues:
 - ◇ none

Sites / Services round table:

- FNAL - Sunday morning ~04:00 dCache admin nodes hung up - paged and fixed. ~monthly event. Latent bug that requires restart.
- PIC - issue on Sunday around midnight affected 3/10 dcap doors that died. Affected ATLAS - failing jobs trying to access cond DB root file through problematic dcap door. Fixed this morning. Believe origin of problem is pre-production version of dCache to test tape protection feature. In contact with developers who expect fix mid-week.
- IN2P3 - ntr
- NDGF - ntr
- ASGC - ntr
- RAL - ntr

- NL-T1 - 1) network maintenance 2) Saturday experienced SRM problem - dCache bug which will be reported.
- KIT - ntr
- OSG - ntr

- CERN - CASTOR LHCb - small issue on internal d2d transfers - taking longer than usual to schedule. Should be "transparent"

AOB:

Tuesday:

Attendance: local(Stephen, Harry, Patricia, Jean-Philippe, Ewan, Lola, Ale, Jeremy, Miguel, Steve, Maarten, Simone, MariaDZ, Peter, Jarka);remote(Jon, Joel, Angela, Vera, JT, Gareth, Rolf, Rob, Gang, Gonzalo).

Experiments round table:

- ATLAS reports -
 - ◆ Issues between CERN-PROD_DATATAPE and several T1s (BNL, IN2P3-CC, NDGF-T1, RAL, SARA) reported in GGUS:58829 [↗](#) yesterday morning vanished at around 1am UTC tonight. Thank you for handling this. What was the solution, please? [Miguel - will have a look]
 - ◆ Slow transfers between BNL and IN2P3-CC, GGUS:58646 [↗](#). What is the status of tests and reconfiguration, please?
 - ◆ Can RAL, CNAF, NDGF please comment on their plans for FTS upgrade? Gareth - ready from next week. Other constraints: Thursday 17th? NDGF - ready when asked... (Asked!)
- CMS reports -
 - ◆ T0 Highlights
 - ◇ Processing normally
 - ◆ T1 Highlights
 - ◇ 600 jobs were stuck in WMS Savannah:114956 [↗](#), solved yesterday
 - ◆ T2 Highlights
 - ◇ MC production as usual
 - ◇ T2_RU_IHEP closed existing SE for CMS and announced a new one (Savannah:114977 [↗](#))
 - ◇ T2_TW_Taiwan 700 jobs failed with 'could not secure connection' (Savannah:115003 [↗](#)) [Gang - problem understood; all failing files on one diskserver with 125MB/s bandwidth to this server. This m/c very busy and bandwidth not enough - ordered h/w to increase bandwidth but might take some time.
- ALICE reports [↗](#) - GENERAL INFORMATION: Two raw data reconstruction activities are ongoing (increasing the number of jobs at the T0): Pass 1 reconstruction and muon data calibration. In addition there are four analysis trains also running.
 - ◆ T0 site
 - ◇ c++ headers missed in four WNs at CERN. Problem reported and described in the ticket: 58587: SOLVED. The ALICE experts have been informed. ALICE users are not reporting the problem anymore, issue seems to be solved.
 - ◇ Top priority issue: all CREAM-CE at CERN are out of production. Submission problems have been reported this morning in a GGUS ticket: GGUS:58852 [↗](#) [Ewan - parser is a SPOF for CEs - if the parser goes down all CEs will go down. Some exception monitoring has gone in but looks like a design weakness]
 - ◆ T1 sites

Tuesday:

- ◊ Minor operations applied today to FZK and CCIN2P3: Restart up of all the local services at the local VOBOXES
- ◆ T2 sites
 - ◊ Stress testing of the T2 Italian sites this morning:
 - Catania: Out of production. local batch system not performing correctly. Ticket: [GGUS:58853](#)
 - CyberSar: Out of production. local batch system reporting the right information Ticket: [GGUS:58854](#)
- LHCb reports -
 - ◆ Running several MC productions at low profile (<5K jobs concurrently).
 - ◆ Merging production.
 - ◆ GGUS (or RT) tickets:
 - ◊ T0: 0
 - ◊ T1: 1
 - ◊ T2: 0
 - ◆ Issues at the sites and services
 - ◊ T0 site issues:
 - CERN: jobs killed by the batch system.
 - ◊ T1 site issues:
 - CNAF: RAL-CNAF one file systematically failing. Under investigation RAL-CNAF people. ([GGUS:58821](#))
 - SARA: Problem fixed: was a dcap port to be restarted. SARA people will put more robust monitoring tools. [GGUS:58838](#) [JT - what kind of jobs did you send 17:00 - 22:00 last night? Big use of NIKHEF-SARA bandwidth: 300 jobs generating 10Gbps. Joel - don't know, will have to check what the activity was in monitoring.]
 - IN2P3: still problem on some pilots landing on the long queue instead of the very-long
 - ◊ T2 sites issues:
 - INFN-T2 : one CE publishing wrong information [GGUS:58850](#)

Sites / Services round table:

- FNAL - yesterday had robotic arm failure in robot- about 11K files pending to tape - backlog cleared overnight.
- KIT - today had change in auth on ATLAS srm - short at risk - no problems seen. No longer used kpwd file but gplasma. Some groups used - asked ATLAS to send a list of all groups that should be supported.
- NDGF - ntr
- NL-T1 - ntr
- RAL - at risk this morning for rebalancing of disks behind LHCb 3D and LFC - all ok
- IN2P3 - still at risk concerning MSS. Took slightly longer than expected - should be operational at 16:00 or so.
- ASGC - nta
- PIC - ntr
- OSG - ntr
- GridPP - started to see some inefficiencies at T2s running CMS jobs. Something to do with use of lcg_cp. Remote stageout - Stephen - not anytime soon.
- CERN AFS UI - currently 3.2.1-0; new = 3.2.6-0. Installed Feb 26th. Hope to change tomorrow at 10:00. Current "current" version has bug in gLite WMS joboutput.
- CERN DSS - problem reported on CASTOR LHCb: found reason and applied corrective measure.

D2D copies. Need code change to fix. Investigation TEAM ticket from ATLAS on xroot analysis access - found some reset by peer events. More investigations on-going. Maarten - any failures reported?

AOB:

- Today announcement on new version of LCG VOMS certs 5.9.0-1 rpm. Tier1s - please ensure that you deploy new RPM before next week! Will ask sites one by one on Thursday Maarten volunteers to open a GGUS ticket against all WLCG Tier1 sites. (but next FTS version 2.2.5 will not need it 😊)

Wednesday

Attendance: local(Jaroslava, Peter, Jean-Philippe, Edoardo, Lola, Nilo, Alessandro, Patricia, Miguel, MariaD, Steve, Jacek);remote(Jon/FNAL, Joel/LHCb, Michael/BNL, Onno/NLT1, Vera/NDGF, John/RAL, Elisabeth/OSG, Angela/KIT, Rolf/IN2P3).

Experiments round table:

- ATLAS reports -
 - ◆ BNL VOMS replica has been upgraded, migrated, and is synchronized (including nicknames) twice daily against CERN's ATLAS replica. Thanks!
 - ◆ Yesterday afternoon expired VO atlas membership of one of certificates used for FTS transfers. Issue was fixed and understood yesterday evening.
 - ◆ Slow transfers between BNL and Lyon: "Network tests have passed and so neither the network nor the disk server is the cause. We are now investigating on the dCache Java process." Is there any update on the dCache Java process investigation, please? https://gus.fzk.de/ws/ticket_info.php?ticket=58646 . No update. Atlas would like to see updates at least twice a week.
- CMS reports -
 - ◆ T0 Highlights
 - ◇ Tape migration backlog last night : since it did not resolve this morning (4TB piled up on T0EXPORT pool) decided to open a GGUS team ticket https://gus.fzk.de/ws/ticket_info.php?ticket=58890.
 - problem now resolved, however CASTOR team (Miguel) checking if this was due to a potential competition between CMS Streamer files (O(100MB)) and normal T0 data migration (O(2GB)). When should a ticket be opened? Atlas had a similar problem 2 months ago. Need to inform the experiment shifters of what plots are more interesting and what levels should signal a problem.
 - ◇ critical CAF AFS area not reachable around 10:30 today : </afs/cern.ch/cms/CAF/CMSALCA>. CMS user opened ticket to CERN/IT Helpdesk (<https://remedy01.cern.ch/cgi-bin/consult.cgi?caseid=CT0000000689745&email=frank.meier>)
 - problem solved now, however need to follow up since there were several such glitches lately
 - ◆ T1 Highlights
 - ◇ CMS will keep T1s busy for next 7 days, doing re-reco of data and MC
 - ◆ T2 Highlights
 - ◇ MC production as usual
 - ◇ T2_RU_IHEP SE rename issue : will have a discussion with admins today on how to follow up, since they closed existing SE for CMS and announced a new one (<https://savannah.cern.ch/support/?115012>) however problem for CMS since need to modify all DBs where the old SE name was stored. There is a T2 support forum

meeting held weekly in CMS to discuss these issues.

- ALICE reports [↗](#) -
 - ◆ GENERAL INFORMATION: Usual reconstruction and analysis tasks ongoing (no MC production activities)
 - ◇ Raw data transfers stopped for the moment (no activity foreseen in these terms before the weekend)
- - ◆ T0 site
 - ◇ Nothing to report
 - ◆ T1 sites
 - ◇ CNAF: The site was blocked yesterday afternoon due to some misconfigured WNs at the site (basic commands i.e., "host" were apparently missed). As prevention measurement, CNAF was blocked in the AliEn central services to avoid too many user jobs failing, while the experts were performing their own tests. Issue solved at around 18:00, the queue was reopened
 - ◆ T2 sites
 - ◇ Catania: local batch system not performing correctly. Ticket: 58853 still unsolved (site admins working on it)
 - ◇ CyberSar: Out of production. local batch system reporting the right information Ticket: 58854 (still no news from the site even if ticket is urgent)
- LHCb reports -
 - ◆ Experiment activities:
 - ◇ Running several MC productions at low profile (<5K jobs concurrently).
 - ◇ Merging production.
 - ◆ Issues at the sites and services
 - ◇ T0 site issues:
 - CERN Castor disk servers unavailable (GGUS:58886 [↗](#)) . Actually not a problem at CERN but IN2P3 slow at giving destination TURLs. Could be due to an HPSS glitch yesterday. When submitting a ticket, a filename and timestamp would be useful.
 - ◇ T1 site issues:
 - IN2P3: one CE in Lyon not publishing information correctly. Fixed now.
 - ◇ T2 sites issues:

Sites / Services round table:

- FNAL:
 - ◆ robotic arm failure yesterday. STK/SUN/Oracle contacted.
 - ◆ upgrading to FTS 2.2.4 tomorrow morning (Chicago time)
- BNL:
 - ◆ BNL-IN2P3 transfer problems: iperf tests successfully run. Having access to the FTS logs would help a lot. BNL has developed a WEB browser to look at logs; this could be used as starting point for gLite developments in this area. A repository for logs would be useful.
 - ◆ Ticket 58891 has been opened against MidwestT2 (problem seeing the LFC).
- NLT1:
 - ◆ VOMS certificates updated
 - ◆ DPM down at Nikhef for half an hour around noon after a reconfiguration (daemons hang)
- NDGF
 - ◆ FTS 2.2.4 has already been put in production last week
- RAL: ntr
- KIT: ntr
- OSG: ntr

- IN2P3: SE problem 10:30-13:00 because of HPSS: slow transfer to tape.
- CERN DBs:
 - ◆ Some problems with CMS databases last night and this morning:
 - ◇ filesystems for logs full on 3 machines. Not detected quickly because of a problem in monitoring, but the impact was minimal as the requests were redirected to other servers.
 - ◇ 3 tablespaces used by PVSS marked as needing recovery (08:00-09:00). Being investigated with Oracle support.
 - ◇ replication between online and offline databases not working
 - ◇ both DBs not responding for a few minutes around 12:00 (because of Streams?)
 - ◆ Intervention on backup databases continues. No backup. Transaction logs will be moved to somewhere else.
- CERN AFSUI:
 1. This morning at 10:00 the sl5 link or current_3.2 now points to gLite 3.2.6-0. The full list of gLite 3.2 links is now:
 - ◇ sl5 -> current_3.2
 - ◇ current_3.2 -> 3.2.6-0
 - ◇ new_3.2 -> 3.2.6-0
 - ◇ previous_3.2 -> 3.2.1-0
 2. Also the .lsc files for ATLAS's vo.racf.bnl.gov and CMS's voms.fnal.gov additional voms services have been installed.
 - ◇ This will now validate proxies from voms.cern.ch, lcg-voms.cern.ch and vo.racf.bnl.gov or voms.fnal.gov for ATLAS for CMS respectively.
 3. Note - I'll submit GGUS tickets for these in a week or so anyway if not fixed.
 - ◇ ATLAS , your vo card does have the port number for vo.racf.bnl.gov (it's not the same as CERN which is fine)
 - ◇ CMS , your vo card does have the voms.fnal.gov entry at all.
 4. Question, do atlas and cms wish to also generate proxies using these US voms servers?
 - ◇ ATLAS - NO, recent nickname situation at BNL should be understood first and understood.
 - ◇ CMS ?

AOB: Publication of FNAL tape metrics in SLS not done because of wrong URL. Could Alberto or Alessandro fix this please? (Jon has sent a mail to Alberto). Alberto has modified the URL, but apparently there is still a permission problem.

Thursday

Attendance: local(Nilo, Eva, Jaroslava, Jean-Philippe, Miguel, Ricardo, Patricia, Roberto, Harry, Andrea, Maarten, Peter, Alessandro, MariaD);remote(Jon/FNAL, Joel/LHCb, Gonzalo/PIC, Xavier/KIT, Vera/NDGF, Michael/BNL, Ronald/NLT1, Rolf/IN2P3, John/RAL, Gang/ASGC, Rob/OSG).

Experiments round table:

- ATLAS reports -
 - ◆ TRIUMF was on downtime yesterday evening in order to upgrade host certificates. When site came back from downtime we observed SECURITY_ERROR on DDM dashboards for transfers from CERN to TRIUMF. Issue was resolved once the site restart SRM after upgrade. Thanks! https://gus.fzk.de/ws/ticket_info.php?ticket=58919
 - ◆ FZK has not had up-to-date DOE CA CRLs which resulted in transfer failures between BNL and FZK. Site upgraded CRLs this morning, since then no more errors. Thanks!
https://gus.fzk.de/ws/ticket_info.php?ticket=58921

- CMS reports -
 - ◆ T0 Highlights
 - ◇ cosmics
 - ◇ data taking during machine development
 - ◇ Intervention on CMS production databases today in the afternoon will affect running
 - ◆ T1 Highlights
 - ◇ CMSSW SW tags failed to be published at INFN-T1 from June 9, 4PM. So all CMS jobs hanging
 - GGUS team ticket https://gus.fzk.de/ws/ticket_info.php?ticket=58920 opened around midnight
 - Answer by site admin this morning around 9:15AM : the file TAG is shared between all CEs and it stored on a NFS shared Filesystem. Anyway the filesystem is working fine and I didn't find any error message in the log file.
 - CMS SW Deployment team fixed the problem on June 10, around 10AM, by setting the tag manually and jobs started to run again. They will provide more feedback on what happened.
 - Lessons :
 - Answer by the INFN-T1 Site admin to did not reach the GGUS browser until noon... In parallel, the answer by the CMS site contact to the related Savannah ticket (<https://savannah.cern.ch/support/?115059>) reached us around 11:45 this morning. Question : why is there such a delay (3 hours) in getting the answer visible in the GGUS browser ? Can the CMS site contact also answer to the GGUS ticket, instead of using Savannah ? Maria says that Savannah can be used as long as the necessary option is used to update GGUS automatically.
 - In parallel, need to improve the CMS Shift procedure for opening GGUS Team tickets : currently only the CMS Comp. Run Coordinator (CRC) can open team tickets, not the Computing Shift Person (CSP). This will stay like that, however we will improve the CSP-CRC pinging procedure. In that case, we would have opened the GGUS ticket during working hours on June 9th, and potentially gained 1 night ! Atlas (Alessandro) is willing to help CMS explaining how they handle tickets and how they train shifters.
 - ◇ Standard CMS activities at T1s :
 - re-reco
 - MC production at other T1s : INFN-T1, ...
 - SW tags failed to be published at CNAF; ticket opened; fixed by hand but problem needs to be understood.
 - ◆ T2 Highlights
 - ◇ MC production as usual
 - ◇ T2_RU_IHEP SE rename issue : will have a discussion with admins today on how to follow up, since they closed existing SE for CMS and announced a new one (<https://savannah.cern.ch/support/?115012>) however problem for CMS since need to modify all DBs where the old SE name was stored
- ALICE reports [↗](#) -
 - ◆ Reconstruction activities ongoing although the principal source of running jobs are coming today from the user analysis trains
 - ◆ T0 site
 - ◇ Overload of the CAF nodes possibly coming from specific experiment software issues. Possible source of problems (memory, CU or network?) to be discussed today during the ALICE TF meeting
 - ◆ T1 sites

- ◊ Nothing to report
- ◆ T2 sites
 - ◊ Subatech: Strange performance of the agent submissions observed yesterday at Subatech. One of the voboxes was not respecting the upper limit defined by the ALICE DB in terms of max number of waiting jobs and therefore submitting new agents all the time. In addition the proxy renewal mechanism was not behaving as expected (user proxy renewed at each agent submission while the ALICE code triggers such a mechanism once per hour only). The problematic VOBOX has been stopped until this morning to drain the site and restrt the proxy renewal mechanism.
 - ◊ Catania: local batch system not performing correctly. Ticket: 58853 SOLVED
 - ◊ CyberSar: Out of production. local batch system reporting the right information Ticket: 58854. Problem solved although not announced by the site nor through the GGUS ticket
- LHCb reports -
 - ◆ Experiment activities:
 - ◊ Running several MC productions at low profile (<5K jobs concurrently).
 - ◊ Merging production.
 - ◊ RecoStripping-04 for 450GeV started today
 - ◆ Issues at the sites and services
 - ◊ T0 site issues:
 - none
 - ◊ T1 site issues:
 - none
 - ◊ T2 sites issues:
 - Shared area problems at CBPF (GGUS:58922 [↗](#)), UFRJ (GGUS:58923 [↗](#)), BHAM (GGUS:58925 [↗](#)), USC-LCG2 (GGUS:58940 [↗](#)).
 - Jobs failing at EFDA-JET (GGUS:58932 [↗](#))

Sites / Services round table:

- FNAL:
 - ◆ FTS upgrade to 2.2.4 being done
 - ◆ tape metrics problem fixed
 - ◆ robot problem again
- PIC: ntr
- KIT: ntr
- NDGF: ntr
- BNL: ntr
- NLT1: ntr
- IN2P3: ntr
- RAL: ntr
- ASGC:
 - ◆ 50 WNs went down last evening, 80 jobs killed. Nodes restarted this morning, everything ok now.
- OSG: ntr
- As a follow-up to the problem with last security patch on Oracle DBs, here is the status:
 - ◆ April PSU applied and rolled back:
 - ◊ LHCB online and offline, ATLAS online and offline, CMS online and offline
 - ◆ April PSU applied:
 - ◊ All test and integration/validation databases ALICE online, ATLAS archive, CMS archive, WLCG and downstream databases (for ATLAS and LHCb)
 - ◆ April PSU never applied:

◇ PDBR and COMPASS databases

- CMS DBs: the roll back of the security patch is being done now. A new patch has been received from Oracle, but we need to reproduce the problem to validate the patch.
- Follow-up on LHCb LSF problem at CERN (Ricardo):
 - ◆ At least some of the jobs were killed because they exceeded the CPU time limit. A more comprehensive reply will be given in the ticket which was opened by LHCb.
 - ◆ Regarding the request to receive an e-mail when a job is killed for excessive CPU time we will investigate this possibility, but since this is an internal LSF mechanism it might not be trivial.

AOB:

Friday

Attendance: local(Eva, Nilo, Miguel, Andrea, Patrica, Jamie, Maria, Jerka, Simone, Ale, Harry, Ewan, Peter, Maarten);remote(Jon, Gonzalo, Onno, Michael, Xavier, Joel, Rolf, Gang, Rob).

Experiments round table:

- ATLAS reports - ATLAS detailed datataking plans for the next 10 days is reported in slide 2 of <http://indico.cern.ch/getFile.py/access?contribId=0&resId=1&materialId=slides&confId=92925>. During these 10 days ATLAS project tag will be data10_comm, thus NO data export is foreseen for those data.
- CMS reports -
 - ◆ T0 Highlights
 - ◇ cosmics
 - ◇ data taking during machine development
 - ◇ Intervention on CMS production databases yesterday in the afternoon : was not as smooth as foreseen:
 - patches did not work as expected with online DB, however thanks to the fast intervention of the CERN DB group fixed the issue
 - other issue with the T0 Data Base Service (TOAST): during the roll-back, other CMS services (DBS, PhEDEX) have moved to the DB-node TOAST is using. As a consequence, the load on this particular node is too high now affecting TOAST. Solution for CMS is to restart these other services. CMS Will schedule that today. [Eva - emergency was due to CMS not doing anything special with DB and wanted to roll-back before end of week. Preferred Thu than Friday! On Wed both CMS DBs hitting a bug in 10.2.0.4. Not related with security patches but CMS felt more comfortable with roll-back.]
 - ◆ T1 Highlights
 - ◇ re-reco, MC production
 - ◆ T2 Highlights
 - ◇ MC production as usual
 - ◇ T2_RU_IHEP SE deprecation issue addressed : will finalize data consistency check on the old SE to make sure no data is lost and then switch to the new SE
- ALICE reports [☞](#) -
 - ◆ T0 site
 - ◇ Yesterday we reported an overload problem observed in the CAF nodes, possibly coming from specific experiment software issues. Possible source of problems (memory, CPU or network?). Harry informed us that the master node (1409) had shown swap errors before becoming unreachable. Discussed with the CAF experts

during the ALICE TF meeting, the issue was coming from one of the analysis jobs that due to some misconfigurations was using a huge amount of memory.

◇ Bug found in the CREAM-Ce setup of the LDAP DB of Alice (service independent issue). This was causing a lack of production through CREAM at the T0. Solved this morning

- ◆ T1 sites
 - ◇ All T1 sites in production
- ◆ T2 sites
 - ◇ Grenoble: alien user proxy (local expert at the site) expired. Site out of production, expert has been warned.
 - ◇ Bologna T2: Restart up of all services this morning to bring the site back in production

• LHCb reports -

- ◆ Running several MC productions at low profile (<5K jobs concurrently).
- ◆ Merging production.
- ◆ T0 site issues:
 - ◇ ask LSF support to keep the short grid queue for SGM and the LONG queue for the rest.
- ◆ T1 site issues:
 - ◇ IN2P3 : finish some intervention on afs to cure the sw area problem
- ◆ T2 sites issues:

Sites / Services round table:

- FNAL - ntr
- PIC - ntr
- NL-T1 - 1 issue: SARA SRM yesterday had transfers that were stuck. The hanging transfers were all from ATLAS jobs and all going to same pool node. Restarted pool node since when new transfers seem fine, ATLAS jobs stuck - killed around 1500 hanging jobs. Think it was a network problem - many dropped packets on this node. Don't see # dropped packets increasing so problem maybe disappeared. Asked network experts to investigate.
- KIT - ntr
- ASGC - ntr
- IN2P3 - ntr
- BNL - 1 GGUS ticket submitted against mid-West T2, experts looking at it.
- OSG - ntr

- CERN - DB more details on bug affecting CMS DBs Wed - followed up with Oracle support. Affects 10.2.04. Fixed in 10.2.0.5. Workaround for 10.2.0.4. Analyzing impact - once understood will send info to T1s. Should not be related to patch - it is something done at DB level.

- CERN Dashboards - Experiment site availability reports sometimes blocked (links at top of this page) - We are currently running 4 applications for 4 different experiments at the same host (we do not have 4 spare hosts one per experiment). We are using virtual hosts in httpd configuration. Apparently this make httpds get screwed up rather often. Restart of httpd helps but only for some time, then it gets crazy again. We have a savannah bug <https://savannah.cern.ch/bugs/?68224> and hope for a fix on a relatively short timescale.

AOB:

-- JamieShiers - 03-Jun-2010

This topic: LCG > WLCGDailyMeetingsWeek100607

Topic revision: r15 - 2010-06-11 - OnnoZweersExCern



Copyright &© 2008-2022 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