

Table of Contents

Week of 100621.....	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.....	8
Friday.....	10

Week of 100621

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, Open Issues & Broadcasts			Change assessments
ALICE ↗	ATLAS ↗	CMS ↗	LHCb ↗	WLCG Service Incident Reports	WLCG Service Open Issues	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(Jacek, Lola, Roberto, Jamie, Alessandro, Gavin, Jean-Philippe, Harry, Jan, Elena, MariaDZ, Julia, Andrea, Simone, Oliver);remote(Michael, Gonzalo, Jon, Angela, Davide, Rolf, Rob, John, Ron, Gang).

Experiments round table:

- ATLAS reports -
 - ◆ ATLAS data taking plans changed on Saturday morning: project tag data10_7TeV, data were exported (<https://prod-grid-logger.cern.ch/elog/ATLAS+Computer+Operations+Logbook/13872>), NO ISSUES
 - ◆ GGUS problem when submitting normal/Team Tickets directly to sites: GGUS:59213 [↗](#)
 - ◆ LYON LFC: very slow access. GGUS:59219 [↗](#) [JPB - please add me in CC for LFC tickets!]
 - ◇ SAM tests not publishing LYON results: GGUS:59233 [↗](#)
 - ◇ ATLAS SAM dashboard not publishing LFC_L results: <https://savannah.cern.ch/bugs/?69016> [↗](#)
- CMS reports -
 - ◆ T0 Highlights
 - ◇ following machine schedule and taking data or not
 - ◇ latest schedule update, no collisions before Thursday

- ◊ IT will replace old wns with new ones: T0 affected with 4 cmsrepack and 108 cmst0 boxes, not a problem because no data taking is ongoing anyway
- ◊ Disk server disappearing from the CMSCAFUSER pool still under investigation: https://gus.fzk.de/ws/ticket_info.php?ticket=59190
- ◆ T1 Highlights
 - ◊ WMS issues still not completely solved, two savannah tickets for tracking:
 - <http://savannah.cern.ch/support/?115119>
 - <https://savannah.cern.ch/support/?115230>
 - ◊ CNAF still encounters transfer problems, sent GGUS team ticket: https://gus.fzk.de/ws/ticket_info.php?ticket=59224 [Wrong config on CNAF side on gridftp servers used by CMS. Root cause identified. Hope to fix by today]
- ◆ T2 Highlights
 - ◊ MC production as usual
- ALICE reports - GENERAL INFORMATION: LHC10b reconstruction activities ongoing together with the reconstruction of the recovered 900 GeV data (Pass2). In addition, there is an important analysis train ongoing and one MC cycle (pp, 7TeV). Raw data transfers have continued during the weekend with low speed rate
 - ◆ T0 site
 - ◊ Good performance of the updated CREAM service (tested and validated this morning)
 - ◆ T1 sites
 - ◊ CNAF: The site admin announced this morning that the CE log file was showing some problems (associated to the user delegated proxy into CREAM). Apparently the update of the system to CREAM1.6 was causing such a problem. After the restart of the local services, the problem has disappeared. It seems therefore associated to the local services, rather to the CREAM service itself.
 - ◆ T2 sites
 - ◊ IPNO: Local CREAM-CE not working (blparser service not alive). Alice contact person notified.
 - ◊ Trujillo: Authentication problems observed in the cream system at submission time. alice contact person at the site prevented
- LHCb reports -
 - ◆ Fairly busy week end with a lot of MC and reco-stripping (05) activities. No major issues to report.
 - ◆ During the weekend a critical SAM tests (SWDIR) has been changed and for some sites it was failing because of some wrong permissions set on some directories. This might well be a software deployment issue and has nothing to do with the shared area service itself. We rolled out to WARNING in case this happens but some sites (CNAF/GridKA) might result red for some time.
 - ◆ T0 site issues:
 - ◊ 2nd queue at CERN, supposed to be switched off last week, still appears in the site BDII and jobs been submitted through. Reopened the ticket.
 - ◊ Transparent-low-risk intervention to change the order of space token look up has it is currently done by SRM and reordering as described here. This fixes transfer problems reported last week due to a contention on the access of some files.
 - ◊ Dashboard SSB: does not update the information (input files for the feeders properly updated but information does not get propagated GGUS:59231)
 - ◊ LFC-RO disappeared from the topology in the SAMDB, no SAM test results available any longer (GGUS:59193) * T1 site issues:
 - ◊ none
 - ◆ T2 sites issues:
 - ◊ UCD and UKI-LT2-UCL-HEP too many pilots aborting there (GGUS:59216) and

GGUS:59215 [↗](#))

Sites / Services round table:

- BNL - 3 points: 1) tomorrow network engineering will be upgrading core multi-level switches. 09:00 Eastern swap out H/A pair. No impact on production as switches will be swapped out gradually. 2) Need to work on publication of SE info to BDII without a CE. Concerns OSG - actively working with OSG on this point. Comes from fact that ATLAS moves files from T3 to non-OSG sites. Operational impact - working with OSG on it. Probably also an impact on LCG operations. 3) Missing site info in CERN BDII. Notice frequently (last Friday for ex) site info in OSG BDII but not at CERN BDII. Needs urgently to be resolved - has deep operations impact. Rob - had call about this Thursday evening. Not sure if site info not available at top level BDII. Seen at SAM BDII. Checking... Ale - have seen info missing in both SAM and top-level BDII. Michael - what are measures and timeline for resolution.
- PIC - ntr
- FNAL - standard dCache w/e crash where pool manager became unresponsive - typically happens ~once per month. Get paged and fix promptly. This time SAM tests check this didn't run for 2 1/2 hours - short downtime - 12% of a day charged against site.
- KIT - ntr
- CNAF - except for CMS issues ntr
- IN2P3 - since January have had WN crashes - tried to find out why. Looks like I/O error which is now corrected by last kernel of RH. Discussion also on HEPiX list - happens also at KIT and SLAC.
- RAL - intervention on disk server this morning. Part of ATLAS HOTDISK- had memory errors. Down for 2h to fix memory. Training day tomorrow so probably no phone in.
- ASGC - this morning our BDII dead, then we fixed it by restarting.
- NL-T1 - upgraded dCache today - all OK.
- OSG - saw handful of tickets over w/e. BDII issue already addressed above.

- CERN-PROD: Partial LFC SIR:
<https://twiki.cern.ch/twiki/bin/view/PESgroup/IncidentLFCRO170610>.
- CERN-PROD: Hangs on lxplus (affecting AFS UI) - AFS support are helping.

- GGUS issues:

Dear site administrators,

The DN string of the certificate used for signing GGUS alarm mails will change with the release on 23rd of June.

Current DN: /C=DE/O=GermanGrid/OU=FZK/CN=ggusmail/gus.fzk.de
New DN: /C=DE/O=GermanGrid/OU=KIT/CN=ggusmail/gus.fzk.de

The alarm test after the release will be done using the new certificate.

Please update your alarm process if necessary.

Thank you.

Kind regards,

Guenter
GGUS team

AOB: (MariaDZ) Only one Tier1 acknowledged the announcement of the GGUS certificate DN change (one attribute changes from OU=FZK into OU=KIT). Sites, if you use DN-matching to accept GGUS emails, please make the change in time for Wed 2010/06/23 any time between 6 and 9am UTC, time when GGUS is down for the Release (announced in https://gus.fzk.de/pages/news_detail.php?ID=414 [↗](#) and <https://goc.gridops.org/site/list?id=1505010> [↗](#)). Issue tracked via <https://savannah.cern.ch/support/?115222> [↗](#)

Monday:

(MariaDZ) The direct site notification unavailability in GGUS was due to a GOCDB4 problem. Some info in https://cic.gridops.org/index.php?section=cod&page=broadcast_archive&step=2&typeb=C&idbroadcast=46951 GGUS developers already took preventive action to avoid remaining with GOCDB-fed fields empty in the future.

Tuesday:

Attendance: local(Jan, Roberto, Maarten, Oliver, Jamie, Elena, Jean-Philippe, Lola, Harry, MariaDZ, Alessandro, Laurence, Ricardo, Nilo, Jacek);remote(Gonzalo, Joel, Rolf, Ronald, Gang, Michael, Ronald, Angela, Rob).

Experiments round table:

- ATLAS reports -
 - ◆ ATLAS data taking plans : project tag data10_7TeV, data are exported
 - ◆ No major problem
 - ◆ LYON LFC: very slow access. GGUS:59219, reported yesterday, the problem is not reproducible. The ticket is closed. It is not understood what caused the problem. [JPB - need to look at logs. Ale - can we put JPB in contact with local experts? Rolf - can follow-up - don't know if it has already been investigated.]
 - ◆ SLS ATLAS storage monitor turning grey GGUS:59237. The problem is related to CRL not updated on CERN AFS. Reply from CERN in GGUS that the problem "fixed by itself", but investigations are ongoing. Re-appeared yesterday afternoon and today again.
 - ◆ Two spikes of errors in transfers from T0 to all Tiers-1's were observed this afternoon. GGUS 59262. All transfers completed from second attempt.
- CMS reports -
 - ◆ [Data Ops]
 - ◇ Tier-0: expect new data end of this week (today?) and continuing high intensity data taking for the coming days after collisions have been established
 - ◇ Tier-1: in the tails of data re-processing, MC reprocessing still ongoing
 - ◇ Tier-2: busy with new MC requests
 - ◆ [Facilities Ops]
 - ◇ Sites are encouraged to migrate their VOBOXes to SLC5 and install latest PhEDEx release.
 - At some point we won't be able to build PhEDEx SLC4 releases, and additionally we could have a change on the PhEDEx DB schema, with a mandatory PhEDEx release available for SLC5 occurring.
 - CMS would like sites to deploy these new VOBOXes during July.
 - ◆ T0 Highlights
 - ◇ following machine schedule and taking data or not
 - ◇ latest schedule update, possible collision tonight, 3X3 bunches (2 colliding bunches in the experiment), inst. lumi up to 5E29 (estimate)
 - ◇ one of the two CMS Frontend WebTools servers is not available since last Friday June 18. We need this machine because it is one of our two load balanced web frontends which serve all kinds of critical web services. GGUS ticket: https://gus.fzk.de/ws/ticket_info.php?ticket=59254, Remedy ticket <https://remedy01.cern.ch/cgi-bin/consult.cgi?caseid=CT0000000692341&email=gowdy@cern> [This needs follow-up: acceptable threshold for data taking in terms of redundancy?]
 - ◆ T1 Highlights
 - ◇ production rolling
 - ◇ T1 sites are asked to report about space in their unmerged diskpools, how much we are currently using of the available space. Please report problems.

- ◊ Several sites observed low CPU efficiencies of CMS jobs. This is most probably due to event sorting in the executable. Experts are looking into this.
- ◆ T2 Highlights
 - ◊ MC production as usual
- ALICE reports [↗](#) - GENERAL INFORMATION: Activities described yesterday in terms of MC production, analysis trains and Pass1 and 2 reconstruction are ongoing today.
 - ◆ T0 site
 - ◊ in terms of the 900 GeV recovered data, Alice has announced today that presently the offline team is reconstructing the data in Pass 2 with the latest AliRoot revision and updated (with respect to the May Pass 2). The full 900 GeV set has been successfully reprocessed (and replicated) and is now being analyzed.
 - ◆ T1 sites
 - ◊ CNAF: Half of the site in production (ce07, 2nd CREAM timing out at submission time). Alice contact person at the site already contacted.
 - ◆ T2 sites
 - ◊ Subatech: Internal network failure on one module of a core router. Discovered around 18:00, the module has been replaced around 20:00. 2 hours unavailability but the services have recovered without manual intervention
 - ◊ IPNO: the problem reported yesterday concerning the blparser service of the CREAM service has been solved. However there is still some pending issue with the system, although the submission procedure works fine, the jobs fail (FailureReason = [pbs_reason=-1]). Alice contact person at the site contacted
 - ◊ Trujillo: Site admin contacted again, the site is out of production since yesterday due to authentication problems shown by the local CREAM system
- LHCb reports -
 - ◆ A lot of MC and reco-stripping (05) activities. Reco-stripping is having problem in running jobs at CNAF and this slows down a bit the progress
 - ◆ We need to coordinate an upgrade on all gLite WMS used by LHCb at T1's to fix the annoying problem of jobs stuck in running status as reported by glite-wms-job-status. Past weeks test on the pilot gswms01 with the patch (<https://savannah.cern.ch/patch/index.php?3621> [↗](#)) installed showed a net improvement of this problem. [Roberto / Maarten : stressed tested recent WMS recently. Can schedule upgrade at CERN soon. Patch not yet in staged rollout so external sites cannot so easily follow. Upgrade at Tier1s once patch officially release - expected in a few days. Maarten - depends on overhead in release process. Patch declared certified. CMS have already deployed this on CMS WMS at CERN.] * T0 site issues:
 - ◊ LSF seems to report wrong information on the used CPU Time (GGUS:59247 [↗](#)) [Once transferred to Remedy no info in GGUS] Ricardo - assigned to batch experts this morning, will have a look asap.
 - ◊ LFC-RO disappeared from the topology in the SAMDB, no SAM test results available any longer (GGUS:59193 [↗](#)). Any news? [Assigned to SAM and not BDII - will pick up - updates this afternoon. Roberto - since beginning of June don't have results from LFC local.]
 - ◆ T1 site issues:
 - ◊ CNAF LSF does not report properly the time left and all pilot abort immediately believing there is not time left
 - ◊ CNAF: one CREAM CE had problem submitting pilots. Jobs failed status was then wrongly reported by our gLite WMS (GGUS:59253 [↗](#))
 - ◊ SARA: problem accessing data via dcap port (GGUS:59252 [↗](#))

Sites / Services round table:

- BDII - Alessandro: GGUS ticket opened to some problem in SAM BDII - no entry found for one host.

More in general: in last few weeks ATLAS DCO observed an increase of # of failures related to (possible) BDII issue. Not sure if related to OSG - think could affect all sites. Started submitting GGUS tickets for all such problems to track better. (Don't know if there was a service change). Decouple issues: most errors "information entry" missing in BDII. Service provider? Load? Very difficult for experiments to tackle these issues. For ADC we think that this is a wider spread problem than OSG. Michael - let's ask Rob to summarize what we know OSG-side. Laurence :SAM problem - tracked in a ticket - James managing this. Other problems with missing info: last week given 4 example tickets - all unrelated to each others. Different sites that use different top-level BDII's. More detailed follow-up on 1 specific issue: was a service-related problem. gstat-prod can monitor top-level and site-level BDII's and hence infer status of resource level BDII's. In this case problem on storage system. Can identify time when it happened but don't know why. local sysadmin has to investigate further - local monitoring and logs. Will produce trouble-shooting guide so that these issues can be followed-up. Michael - will this also cover SAM test issues? Lawrence - yes.

- IN2P3 -
- FNAL - very busy right now. SRM overload and had to restart again yesterday.
- NL-T1 - dCache upgrade not executed on some pool nodes. Done this morning - some pool nodes not available during this time. Ticket is waiting for reply - waiting for LHCb to confirm AOK.
- ASGC - ntr
- BNL - in middle of swapping core routers
- KIT - ntr
- PIC - since Saturday see jobs from ATLAS production - mainly seg faulting. Known issue? ATLAS looking into it.
- OSG - nothing other than BDII's
- CNAF - LHCb: have updated config for LSF - should fix their problem - has to be tested (LHCb contact informed). CREAM CEs - seen a number of problems related to heap space exhausted in Tomcat. Will increase and restart. Roberto - do you think LSF now back to "normality"? A: yes, should be.
- CERN DB - tomorrow at 14:00 int11R db will be migrated to new h/w. On Thursday int2r of CMS also to new h/w. Both will take about 1 H.[Ricardo - was this the same problems as at CERN? And what was change at CNAF? (CPU factor)]

AOB: (MariaDZ) GGUS:52383 [↗](#) is pending since October 2009. Correctly assigned at the beginning to VOSupport(CMS), it ended up in the generic VO Management Support Unit. Now put back on the CMS VO Admin lap.

Wednesday

Attendance: local(Elena, Jean-Philippe, Zbigniew, Markus, Lola, MariaG, Maarten, Alessandro, Gavin, MariaD, Ian, Roberto);remote(Alessandro/CNAF, Michael/BNL, Gonzalo/PIC, Joel/LHCb, Gang/ASGC, John/RAL, Angela/KIT, Onno/NLT1, Christian/NDGF, Rob/OSG, Catalin/FNAL).

Experiments round table:

- ATLAS reports -
 - ◆ Two problems were seen: transfers from T0 to all Tiers-1's were failing at first attempt because of srm failure but succeeded at second attempt GGUS:59262 [↗](#) and timeouts in retrieving files from T0ATLAS GGUS:59269 [↗](#). The source of both problem is overloaded CERN LDAP service. This has been fixed
 - ◆ Ongoing issue with SLS ATLAS storage monitor turning grey related to CRL not updated on CERN AFS GGUS:59237 [↗](#)
 - ◆ No major problem related to Tier-1's or Tier-2s

- CMS reports -
 - ◆ T0 Highlights
 - ◇ LDAP related issues as well
 - ◇ following machine schedule and taking data or not
 - ◇ latest schedule update, possible collision tonight, 3X3 bunches (2 colliding bunches in the experiment), inst. lumi up to 5E29 (estimate)
 - ◆ T1 Highlights
 - ◇ production rolling
 - ◇ T1 sites are asked to report about space in their unmerged diskpools, how much we are currently using of the available space. Please report problems.
 - ◆ T2 Highlights
 - ◇ MC production as usual
- ALICE reports [↗](#) - GENERAL INFORMATION: Raw data transfer activities restarted. Presently transferring to NDGF, RAL, CNAF, FZK
 - ◆ T0 site
 - ◇ No issues have been found by ALICE after the upgrade of ce201 to the latest 3.2.6 version. Green light for the upgrade of ce202 and ce203 tomorrow. As soon as the services are migrated and put back in production, feedback will be provided to the service responsible
 - ◆ T1 sites
 - ◇ CNAF: ce07 back in production after the timeouts messages found yesterday and reported to the site admin.
 - ◆ T2 sites
 - ◇ MEPHI: Last russian site still blacklisted (still running SL4) has been put back in production today. The site is providing also a CREAM service. Submission services tested and no issues found. Pending service: xrootd SE still to be configured
 - ◇ Bratislava: Manual restart of all local services at the site. Most probably the proxy renewal mechanism of the vobox has failed at a certain point killing all AliEn services. The vobox will be observed in the next hours to see if the problem persists
- LHCb reports -
 - ◆ A lot of MC and reco-stripping (05) activities. Reco-stripping is having problem in running jobs at CNAF and this slows down a bit the progress
 - ◆ T0 site issues:
 - ◇ still valid the issue of LFC-RO not appearing in SAMDB (GGUS:59193[↗](#))
 - ◆ T1 site issues:
 - ◇ CNAF LSF does not report properly the normalization factor used by CPU Time left utility to run in filling mode. The value is ~60 times larger than expected if the reference CPU is the machine advertised in the BDII (2373 2KSI) (GGUS:59274[↗](#))
 - ◇ SARA the file was stored on a server that was unavailable for a few minutes (GGUS:59252[↗](#))
 - ◆ T2 sites issues:
 - ◇ Legnaro: similar (to CNAF) problem in the value of the cpuf reported by LSF (GGUS:59272[↗](#))
 - ◇ also shared areas problems

Sites / Services round table:

- CNAF
 - ◆ ALICE: CREAM-CE problems related to memory exhaustion. In contact with support
 - ◆ LHCb: CNAF does not publish normalized values. In contact with LHCb about this issue.
- BNL: ntr
- PIC: ntr

- ASGC: ntr
- RAL: GGUS alarm ticket exercise. 2 hours delay between the ticket and the mail (normally a few minutes). Same problem seen at SARA. Problem at GGUS?
- KIT: ntr
- NLT1:
 - ◆ SRM at SARA: tried new configuration parameter yesterday to shift the load out of overloaded servers but noticed problems when reading files; went back to the previous configuration; should be ok. GGUS:59283 [↗](#) opened.
 - ◆ GGUS alarm ticket exercise: some confusion at NIKHEF where they try to do automatic processing of the tickets. The problem was due to the fact that the ticket was not signed by GGUS but by Guenter himself. This is wrong and should not happen. GGUS:59327 [↗](#).
- NDGF: ntr
- FNAL: ntr
- OSG: To investigate the issue of sites disappearing from SAM BDII, this BDII will be monitored carefully by OSG. The feeling is that this BDII is less stable than the central BDII. Alessandro: Atlas does not blacklist sites just according to the results of the SAM tests; they check that there are actual failures affecting Atlas jobs. Michael: as the SAM BDII is part of the verification of the sites, this service must be reliable.
- CERN:
 - ◆ Ian: CASTOR failures were noticed starting around 10:00 and jobs started failing around 13:00. This was investigated and fixed quickly: it was due to the overload of the CERN LDAP service used by the CASTOR disk servers. The root cause was a stress test done from virtual machines.
 - ◆ Gavin: this problem affected the batch service as well. Problem understood. No large scale stress test should target a production service. ce202 and ce203 upgrade tomorrow (announced in GOCDB).
 - ◆ DBs: migration of int11r instance to new hardware as announced.

AOB:

- - ◆ (MariaDZ) GGUS:52383 [↗](#) is pending since October 2009. Correctly assigned at the beginning to VOSupport(CMS), it ended up in the generic VO Management Support Unit. Now put back on the CMS VO Admin lap.
 - ◆ Please see <https://savannah.cern.ch/support/?115137> [↗](#) about the ALARM tests
 - ◆ Joel reports a discrepancy between VOMRS and VOMS. A GGUS ticket must be submitted and assigned to VOMRS.

Thursday

Attendance: local(Elena, Jean-Philippe, Harry, Nicolo, Alessandro, Gavin, Ian, Maarten, MariaD, Jacek, Edoardo);remote(Jon/FNAL, Michael/BNL, Gang/ASGC, John/RAL, Jeff/NLT1, Angela/KIT, Davide/CNAF, Rob/OSG).

Experiments round table:

- ATLAS reports -
 - ◆ LYON LFC: some atlas users cannot access LFC. The problem is not seen permanently GGUS:59344 [↗](#), investigation was done yesterday evening by IN2P3: problem to establish authenticated connection. The ticket is re-assigned to LYON LFC team. Similar problem was observed Saturday GGUS:59219 [↗](#). JPBaud, in cc, would like to see the logs of the LFC servers if possible.
 - ◆ GGUS: problem when submitting tickets directly to LCG sites for normal/Team tickets (the problem was also seen on Saturday) GGUS:59213 [↗](#) re-opened, email to

ggus-info@cernSPAMNOTNOSPAMPLEASE.ch. The problem was fixed.

- ◆ Ongoing issue with CRL not updated on CERN AFS GGUS:5923 [↗](#)
- ◆ ATLAS data taking plans: Start physics run at 18:00 (data10_7TeV). There will be the possibility to have up to 550MB/s RAW data from SFO to T0. More details in [elog](#) [↗](#)
- CMS reports -
 - ◆ Issue at CERN
 - ◇ AFS slowness yesterday afternoon especially on AFS122 (CMS software). Back to normal performance at 8pm. No ticket was sent because nothing seemed to be broken.
 - ◇ Empty CRL file in AFS UI area: GGUS:59349 [↗](#); fixed this morning
 - ◆ T0 Highlights
 - ◇ following machine schedule and taking data or not
 - ◇ latest schedule update, possible collision tonight, 3X3 bunches (2 colliding bunches in the experiment), inst. lumi up to 5E29 (estimate)
 - ◆ T1 Highlights
 - ◇ production rolling
 - ◇ T1 sites are asked to report about space in their unmerged diskpools, how much we are currently using of the available space. Please report problems.
 - ◆ T2 Highlights
 - ◇ MC production as usual
- ALICE reports [↗](#) - GENERAL INFORMATION: Raw production ongoing: LHC period LHC10c (calibration events). MC production: 7TeV pp events. 4 analysis trains
 - ◆ T0 site
 - ◇ Upgrade of ce202 and ce203 to the 3.2.6 version foreseen for this afternoon. Feedback will be provided to the system experts as soon as the nodes are put back in production
 - ◆ T1 sites
 - ◇ No issues to report
 - ◆ T2 sites * To reinforce WN nodes under CREAM-CE, Hiroshima site will be in a drain mode at around 0300 (CEST) on June 27, and shutdown after 24H draining. * Manual interventions needed this morning in Catania and TriGrid-Catania, to put the sites back in production
- LHCb reports -
 - ◆ A lot of MC and reco-stripping (05) activities. Reco-stripping is running close to its completion

			Processed (%)	running/wai
3.5 Magnet Down				
6979	no prescaling, runs up to 71816	100%	1/0	337
6981	prescaling, runs from 71817	94.5%	119/2	227
3.5 Magnet Up				
6980	no prescaling, runs up to 71530	99.9%	0/0	2608
6982	prescaling, runs from 71531	99.8%	2/1	34

- - ◆ All SAM jobs failed in test the LHCb Applications due to a bad options on the LHCb side.
 - ◆ VOMS: GGUS:59337 [↗](#) Inconsistencies observed between VOMRS and VOMS
 - ◆ T0 site issues:
 - ◇ The USER space token got filled up (as reported by SLS and service managers). LHCb Data Manager has to clean it up.
 - ◆ T1 site issues:
 - ◇ SARA: dcap access issue reported at GGUS:59252 [↗](#) has been confirmed to be solved.

◇ CNAF (and Legnaro): LSF returning inconsistent (with the BDII) information (GGUS:59274). The discussion is proceeding and it looks like site managers believe that this is not matter to be exposed to experiments. However the problem is a real problem, the information returned by LSF clients on the WN are used to run pilots in filling mode.

- ◆ T2 sites issues:
 - ◇ RO-07-NIPNE, problem with shared objects (GGUS:59339)
 - ◇ GRISU-CYBERSAR-CAGLIARI: shared area issue (GGUS:59334)

Sites / Services round table:

- FNAL: ntr. Question from Jon: how many alarms should have been received by a site for the alarm test? Answer: one.
- BNL: ntr
- ASGC: ntr
- RAL: disk problem GDSS547
- NLT1:
 - ◆ SARA SE: dCache 1.9.5-19 does not accept old style proxies
 - ◆ alarm test: SARA did not receive any mail notification. Jeff would like to get from GGUS support the mail timestamp and mail id to investigate.
- KIT: ntr
- CNAF: ntr
- OSG: alarm tickets exercise went well in OSG: tickets received and processed quickly.

- CERN-PROD: SIRs for LDAP overload incident: IncidentLDAP220610 and IncidentsLdapOverloaded22Jun2010

AOB: (MariaDZ)

- ALARM test script hickups are the reason for yesterday's personal certificate signature. Tier 1 sites should only consider the alarm tickets submitted 2010-06-23 12:27:51 UTC and check they received the alarm mail correctly signed with the GGUS certificate or not. Please enter comments in <https://savannah.cern.ch/support/?115137> if possible.
- GGUS:59337 corresponds to LHCb report of vomrs-voms non-sync. Should be repaired now (?)
- SIR by the GGUS developer on the GOCDB-GGUS data copy problem again in <https://savannah.cern.ch/support/?115297#comment1>. Thanks to Tom Fifield for reporting this.

Friday

Attendance: local(Elena, Harry, Jean-Philippe, Alessandro, Jamie, MariaG, Rafaelo, Maarten, Ian, Jacek, Akshat);remote(Michael/BNL, Jon/FNAL, Rolf/IN2P3, Robert/LHCb, Gang/ASGC, Vladimir/CNAF, John/RAL, Jos/KIT, Onno/NLT1, Rob/OSG).

Experiments round table:

- ATLAS reports -
 - ◆ ATLAS data taking: physics run during the night (data10_7TeV) with 300MB/s RAW data from SFO to T0. New physics run is expected tonight.
 - ◆ LYON LFC: ongoing problem with accessing IN2P3 LFC GGUS:59344, GGUS is updated with client side error messages with more infos from logs of ATLAS Site Services. Rolf: expert working on it, should get a reply tonight. Alessandro: is alarm ticket ok if problem occurs again? Yes.

- CMS reports -

- ◆ T0 Highlights
 - ◇ recorded some data in last nights fill
 - ◇ expect more data this weekend
 - ◇ latest schedule update, possible collision tonight, 3X3 bunches (2 colliding bunches in the experiment), inst. lumi up to 5E29 (estimate)
- ◆ T1 Highlights
 - ◇ production rolling
 - ◇ T1 sites are asked to report about space in their unmerged diskpools, how much we are currently using of the available space. Please report problems.
- ◆ T2 Highlights
 - ◇ MC production as usual
- ALICE reports [↗](#) - GENERAL INFORMATION: No raw data reconstruction activity at this moment. Production concentrated on user analysis activities and MC.
 - ◆ T0 site
 - ◇ ce202 and ce203 migrated to 3.2.6 version yesterday afternoon. The systems were tested and validated by Alice. No issues found. Feedback sent to the service responsible
 - ◆ T1 sites
 - ◇ No issues to report
 - ◆ T2 sites
 - ◇ Legnaro announced some problems during the migration of the local CREAM to 1.6 (site was yesterday in downtime to proceed with the upgrade of the system). The site has downgraded to CRE!M1.5 until the problems are solved. Site back in production with CREAM1.5
- LHCb reports -
 - ◆ Received a run (10 files in total) last night.
 - ◆ Reco-stripping-05 is running close to its completion while reco-stripping-04 a bit late.
 - ◆ MC productions still ongoing with some problem application side.
 - ◆ T0 site issues:
 - ◇ 3D stream replication failures to pic since yesterday (processes got stuck on one of big transaction last night, unknown the reason). This means no new ConditionDB data replicated to the Iberian T1; today at 10.22 it affected also the capture process; for one hour no new data were replicated to all T1s. So far the problem has been fixed. So replication to T1s is on track and all data gaps had been recovered. Open a GGUS ticket for traceability (GGUS:59404 [↗](#)) considering that this problem - in data taking period - would have compromised the data reprocessing at T1's letting all jobs to crash and would have been important enough to raise an ALARM. Reason of the process stuck has to be further investigated.
 - ◇ wms203 in drain to apply the patch that cures issues against CREAMCEs
 - ◇ upgraded the two remaining CREAMCEs ce202 and ce203 to 3.2.6
 - ◇ The FAILOVER space token got filled up yesterday has been cleaned by the data manager. Crash files uploaded in the debug area (served by the same service class)
 - ◆ T1 site issues:
 - ◇ PIC: MC_M-DST space token completely full bringing all transfers to there failing. PIC provided their quota to LHCb so this becomes an issue for LHCb to steer upload of MC output data somewhere.
 - ◇ RAL: Observed some degradation in the SARA-RAL transfer channel. The replication takes very long and the channel is particularly slow if compared with any other channel setup at RAL (or at CERN) to RAL (GGUS:59397 [↗](#)).
 - ◆ T2 sites issues:
 - ◇ none

Sites / Services round table:

- BNL: ntr
- FNAL: ntr
- IN2P3: ntr
- ASGC: ntr
- CNAF: ntr
- RAL: site at risk Monday-Wednesday next week because of work on a transformer.
- KIT: pre-announcement of 1 hour downtime on 6th July for networking maintenance (not yet in GOCDB)
- NLT1:
 - ◆ second fileserver with problems (Infiniband), investigating with vendor the issue. This is only affecting LHCb.
 - ◆ cable swap announced in GOCDB, but a 5 hours delay was observed before GOCDB actually announced it. ASGC and other sites have seen similar problems. Could NLT1 submit a GGUS ticket?
- OSG: Improved publication of Storage Elements in OSG. BNL happy.

- CERN DBs: investigating the STREAMS replication problem reported by LHCb, but no trace file, so not very optimistic about the possibility to find the reason.

AOB:

-- JamieShiers - 18-Jun-2010

This topic: LCG > WLCGDailyMeetingsWeek100621

Topic revision: r22 - 2010-06-25 - JosVanWezel



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