

Table of Contents

Week of 100913.....	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.....	8
Friday.....	10

Week of 100913

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(Roberto, Graeme, Patricia, Jean Philippe, Ewan, Doug, Jacek, Harry, Marie Christine, Ignacio, Massimo, MariaD, Maarten, Ale, Dirk);remote(Michael/BNL, Jon/FNAL, Kyle/OSG,Gang/ASGC, Alessandro/CNAF, Gonzalo/PIC, Onno/NL-T1, Rolf/IN2P3, Joel/LHCb, Tiju/RAL, Jos/KIT).

Experiments round table:

- ATLAS reports - Doug
 - ◆ Tier-0:
 - ◇ No new data, things quiet.
 - ◆ Tier-1:
 - ◇ NDGF: Unscheduled downtime after loss of OPN connectivity and internal dCache problems.
 - ◇ INFN: Unscheduled downtime for upgrade of Storm to 1.5.4-4.
 - ◇ Taiwan disk problems on Sat., caused failures for a few hours, but then back on-line. No serious problems.
 - ◆ Tier-2:
 - ◇ Many issues, but none that are of general concerns.
- CMS reports - Marie Christine
 - ◆ Experiment activity

- ◊ Run tests on going;
- ◆ Central infrastructure
 - ◊ 3_8_3 being deployed at T0; test for AOD Commissioning Data with 3_8_3 will start today. This will result in modest increase (~10%) increase in the rate from CERN to Tier-1s when data taking resumes.
 - ◊ GGUS ticket 61706: no further news, will have more information at 4 pm today.
- ◆ Tier1 issues
 - ◊ Nothing to report
- ◆ Tier2 Issues
 - ◊ T2-TR-METU and T2-UK-Southgrid RAL back to green in SSB
 - ◊ Jobrobot errors at T2-FR-IPHC today
- ◆ MC production
 - ◊ Large abort rate for MC production at T2-TW-ASGC, <https://savannah.cern.ch/support/?116723>
 - ◊ Very large (1000M) MC production ongoing T1s and T2s; reprocessing with CMS_SW 38XX will start soon (pre production first)
- ◆ AOB
 - ◊ When we re-start with Data Taking, will run tests for Heavy Ions first, with an additional express Stream pre scaled by 100, and activate prompt reco delay (6 hours)
- ALICE reports [↗](#) - Patricia
 - ◆ T0 site
 - ◊ Good behavior of the services during the weekend
 - ◆ T1 sites
 - ◊ RAL: Over the weekend the site observed some problems with lgce01 not submitting jobs to the batch system. The problem has been solved by dropping and recreating the cream databases, which did fix the problem, however since then the site managers have been seeing log error messages reporting problems while delegating the user proxy to the system. The AliEn services have been restarted this morning solving the problem
 - ◊ CNAF: The site has been out of production during the night. Issue solved this morning and back in production
 - ◆ T2 sites
 - ◊ New bunch of T2 voboxes registrations (into myproxy server) during the weekend. Only one request (Dortmund) sent this morning is still pending
- LHCb reports - Joel
 - ◆ Suffering a lot of problems due to various space tokens getting full at CERN and at different T1. On Sunday some space has been recovered by cleaning up un-merged DST files that were input of merging production now finished. Recovered 50TB (12TB on lhcbmdst) at CERN
 - ◆ T0 site issues:
 - ◊ Due to our space-tokens lhcbuser and lhcbmdst full the operator piquet has been called twice this weekend and in turn notified the collaboration to clean up some space. Our data manager managed to free some space to make at least usable these spaces and avoid internal alarms.
 - ◊ Dirk: Would LHCb have profited from additional alerts before the weekend from castor ops? Roberto: not really - all relevant information was available to LHCb.
 - ◊ Streams replication: some blocking issues (streams processes were blocking each other) on Sunday night and LFC replication was progressing really slow. Fixed immediately.
 - ◊ Joel: not always clear when a streams issue is being picked up or fixed by T0 DBA team. Jacek: will work on improving notifications back to VOs.
 - ◆ T1 site issues:
 - ◊ CNAF: ConditionDB at CNAF times up connections form other sites WNs (and jobs

got stalled). GGUS:61989. fixed promptly by CNAF people opening to the general network the listener.

◇ PIC: Files reported UNAVAILABLE by local SRM (GGUS:62019)

◇ GRIDKA: increased user and production jobs failure rate. Jobs stalling accessing data, it seems to be a problem with dcap movers. (GGUS:61999)

· Jos: the issue is solved now.

◇ NIKHEF: CREAMCE issue, pilot jobs aborting there. (GGUS:62001)

Sites / Services round table:

- Jon/FNAL - over weekend: mc submission with very verbose logging - was difficult to keep jobs running due to high logging volume. Experiment should avoid larger submissions with high logging levels.
 - Gang/ASGC - ntr
 - Alessandro/CNAF - problem with cream CE: service reached max number of open files. The service was restarted this morning and is now working fine. Discussing appropriate limits with cream developers.
 - Gonzalo/PIC - concerning open ticket from LHCb: problems coming from single solaris pool - network / bios setting are now being reviewed by experts.
 - Onno/NL-T1 - ntr
 - Rolf/IN2P3 - unscheduled downtime sat night (21 to 1): AFS cell frozen due to overloaded.
 - Tiju/RAL - at risk tomorrow morning: intervention on firewall
 - Jos/KIT - dcap mover problem was related to wrong config on one node (lack of file descriptors).
 - Michael/BNL - ntr
 - Kyle/OSG - ntr
-
- Jacek/CERN: On Sat streaming of atlas to asgc was fixed. This morning streaming problems to all sites - still being investigated. So far the streams delay reached up to 3h.

AOB:

Tuesday:

Attendance: local(Gavin, Graeme, Patricia, Roberto, Jean Philippe, Luca, Harry, Ale, Dirk);remote(Michael/BNL, Jon/FNAL, Gonzalo/PIC, Rolf/IN2P3, Onno/NL-T1, Alessandro/CNAF, Tiju/RAL, Gang/ASGC, Andreas/KIT, Marie Christine/CMS).

Experiments round table:

- ATLAS reports - Graeme
 - ◆ ATLAS/Tier-0:
 - ◇ Faulty flow meter caused toroid fast dump yesterday - no magnets until Friday, reduced shifts.
 - ◇ No data export to T1s.
 - ◆ Central Services
 - ◇ Problem with central catalogs between 0400-0730 caused by mod_ssl RPM upgrade adding a bad configuration file. Restarted after AMOD intervention. Central services team working on permanent fix.
 - ◆ Tier-1:
 - ◇ RAL FTS drained this morning for firewall reboot.
 - ◇ Ongoing: INFN-BNL network problem GGUS:61440. Seems to be ok in the last week, but low number of transfers.
 - ◇ Ongoing: RAL-NDGF functional test GGUS:61306. Definitely still a problem.
 - ◆ Tier-2:

Tuesday:

◇ Issues, issues, ...

- CMS reports - Marie Christine
 - ◆ Experiment activity
 - ◇ Run tests on going;
 - ◆ Central infrastructure
 - ◇ Yesterday evening: high load on CMSR node 4, which caused Phedex Agents to fail; after reboot, seems Ok again. A ticket was opened about the issue: <https://savannah.cern.ch/support/?116741>.
 - Luca: user generated load caused DB nodes to swap and then to reboot - DB team is following up with users
 - ◇ Error in Central DB (Si Pixel) which triggered a series of job failures; corrected after 2hours, but CMS frontier was refreshed during the critic period, hence Frontier server was affected for about 6 hours.
 - ◇ GGUS ticket 61706: we will feed with further details as soon as the problem arises again with higher frequency.
 - ◆ Tier1 issues
 - ◇ Nothing to report
 - ◆ Tier2 Issues
 - ◇ Mostly OK, minor issues
 - ◆ MC production
 - ◇ Very large (1000M) MC production ongoing T1s and T2s; reprocessing with CMS_SW 38XX has started (pre production first)
 - ◆ AOB
 - ◇ When we re-start with Data Taking, will run tests for Heavy Ions first, with an additional express Stream pre scaled by 100, and activate prompt reco delay (6 hours)
- ALICE reports [↗](#) - Patricia
 - ◆ GENERAL INFORMATION: Pass1 reconstruction and massive MC production activities
 - ◆ T0 site
 - ◇ No issues to report
 - ◆ T1 sites
 - ◇ CNAF: expansion of the disk at CNAF to achieve 1.2PB. Establishment of the ALICE partitions under discussion with ALICE
 - ◇ IN2P3-CC: Tape system is showing bad results in ML associated to xrdcp: GGUS: 62073
 - New update: Ticket solved and verified. The origin of the problem was a lack of available space on the data server which prevented new copies
 - ◆ T2 sites
 - ◇ Operation issues reported to the site admins at: Mephi (Russia), GRIF (France) and Bologna-T2 (Italy)
- LHCb reports - Roberto
 - ◆ Experiment activities: User analysis and few MC producitions
 - ◆ T0 site issues
 - ◇ The LFC replication tests are failing consistently at all sites. The 2+2 minutes timeout too short. Verified that after 26 minutes the information was not yet propagated down to the T1 (GGUS:62078 [↗](#))
 - ◇ Luca: replication setup (and latency) was like this since the service started. Will contact LHCb offline for clarification.
 - ◆ T1 site issues:
 - ◇ IN2P3 : SRM SAM tests were failing yesterday and it has been fixed : The CRL are in a shared area on AFS here. We had an overload on our AFS cell yesterday at the timestamps you gave.

- ◇ GridKA: removed hanging dcap movers, solved the problem. (GGUS:61999[↗](#))
- ◇ PIC: The network's server is not as stable as it should. Looking at the problem Work in progress (GGUS:62019[↗](#))
- ◇ NIKHEF: CREAMCE issue, pilot jobs aborting there. (GGUS:62001[↗](#)) . Working in close touch with developers.
- ◇ CNAF : All jobs aborted at ce07-lcg.cr.cnaf.infn.it (GGUS:62029[↗](#)). Restarted tomcat.

Sites / Services round table:

- Michael/BNL - ntr
- Jon/FNAL - ntr
- Gonzalo/PIC - ntr , question: what should T1 sites do with LFC replication ticket from LHCB ? Roberto: GGUS created individual site tickets. Sites can ignore/close those as central replication is followed up by T0.
- Rolf/IN2P3 - as pre-announced end of august, large outage next week affecting many services: starting with batch shutdown on 21 Sept (long queues closed on 19). Site will be back on 23 Sep and will aim to have dcache back rapidly (at least 1 day dcache outage though). AFS problem during weekend and yesterday again - hope outage next week will help (AFS client update will be applied which may resolve locking issues).
- Onno/NL-T1 - GGUS:62001[↗](#) - don't have more information, but since the cream CE has been restarted: Is problem still there? Roiberto: will check and update ticket. Which version of cream is running? Onno: will check and get back to LHCB.
- Alessandro/CNAF - problem with batch this morning after reconfig - batch started for unclear reasons to to clear job queue. Site is in contact with support. Also planning to improve performance of master node using new server. Will schedule downtime for the introduction of new server. Until then batch probs at CNAF may still appear.
- Tiju/RAL - two local nameserver installations for castor completed fine.
- Gang/ASGC - ntr
- Andreas/KIT - ntr

AOB:

Wednesday

Attendance: local(Massimo, Peter, Gavin, Luca, Patricia, Graeme, Harry, Ale, Dirk);remote(Gang/ASGC, Jon/FNAL, Michael/BNL, Kyle/OSG, Onno/NL-T1, Joel/LHCB, Alessandro/CNAF, Rolf/IN2P3, Tiju/RAL, Gonzalo/PIC).

Experiments round table:

- ATLAS reports - Graeme
 - ◆ ATLAS/Tier-0
 - ◇ Commissioning and testing.
 - ◇ No data export to T1s.
 - ◆ Tier-1s
 - ◇ SARA suffered a burst of stage-out failures last night, GGUS:62096[↗](#). Some confusion on the ticket, but errors have now disappeared.
 - ◇ BNL reported a namespace database load issue, which led to some transfer failures (<https://prod-grid-logger.cern.ch/eelog/ATLAS+Computer+Operations+Logbook/16972>[↗](#)). Now back to normal.
 - ◇ We would like an update on the "ongoing issues" tickets, please!
 - ◆ Tier-2s
 - ◇ TR-10-ULAKBIM back after lengthy SE instability, GGUS:61997[↗](#).

- ◇ PSNC SE problems, GGUS:62107 [↗](#). Now resolved.
- ◇ TECHNION-HEP transfer problems from a malfunctioning transceiver, GGUS:62103 [↗](#). Fixed.
- ◇ RU-PROTVINO-IHEP SRM down, GGUS:62086 [↗](#). No answer on ticket, but now recovered.
- ◇ IL-TAU-HEP low transfer efficiency to/from SARA, GGUS:62109 [↗](#). Ongoing.
- ◇ UKI-NORTHGRID-LANCS-HEP transfer failures after DPM hardware upgrade, GGUS:62105 [↗](#). Ongoing.
- ◇ CYFRONET-LCG2 poor transfer efficiency to INFN-T1, GGUS:62127 [↗](#). No response yet.
- ◇ ITEP 100% failures with SRM down, GGUS:62122 [↗](#). Latest transfers look better, but no ticket response yet.
- ◆ Michael/BNL: the planned OPN test between BNL and CNAF have been performed.
 - ◇ As was agreed CNAF and BNL engineers ran a test to compare the transfer performance between the sites using the OPN and the general IP services network (GARR-GEANT2-ESnet). Though using the OPN results in a 3 fold improvement over the general IP services network the currently achievable maximum performance of 1 Gbps (memory-to-memory) is below expectations. BNL observes a much better performance with for example KIT and CC-IN2P3.
- CMS reports - Peter
 - ◆ Central infrastructure
 - ◇ CMS requested to increase Castor stager limit (from 900 requests in 180 sec to 3000 requests in 500 sec) in order to (partially) address the large number of redundant recalls by single users (CT711598)
 - ◇ Still trying to understand the full story of the CMSR node 4 glitch from Monday night (causing Phedex Agents to fail, see <https://savannah.cern.ch/support/?116741> [↗](#)). It is important for us to understand, in order to improve the CMS Computing Shift monitoring/alarmed procedures related to CMS Oracle Rack glitches (in that particular cases, there were many elog/tickets about subsequent failures, but not about the actual source of the problem)
 - ◇ follow up on GGUS ticket 61706 (CMS Job Robot failures at CERN) : in the last 10 days, Job Robot success rates > 90%
 - ◆ Tier1 issues
 - ◇ Nothing to report
 - ◆ Tier2 Issues
 - ◇ GGUS ticket 62100 submitted for BDII not visible in T2_Estonia : now visible again, was related to a site configuration.
 - ◆ MC production
 - ◇ Very large (1000M) MC production ongoing T1s and T2s; reprocessing with CMS_SW 38XX has started (pre production first)
 - ◆ AOB
 - ◇ The CMS Computing Run Coordinator (reporting here) shift period was changed from Tuesday-to-Tuesday, starting from beginning of October
- ALICE reports [↗](#) - Patricia
 - ◆ General information: Pass1 reconstruction and massive MC production activities together with user analysis tasks
 - ◆ T0 site
 - ◇ New CREAM module has been put in production at CERN for testing purposes. The new CREAM module has deleted any reference or query to the CREAM-DB. The same module will be tested today in some T2 sites (Most probably Subatech and Torino)
 - ◆ T1 sites

- ◇ IN2P3-CC. GGUS: 62118. Same issue observed yesterday (probable lack of space on disk preventing the copy of new files? waiting for the site admin expertise)
- ◆ T2 sites
 - ◇ Usual operations
- LHCb reports - Joel
 - ◆ Experiment activities: User analysis. Fighting with shortage of disk space deciding which reprocessing has to be clean.
 - ◆ T0 site issues:
 - ◇ Issue with LFC replication (GGUS:62078 [↗](#)). It has been forgotten that for the LFC replication there is special 'real time downstream' optimization turned on in order to minimize replication latency through downstream database. It looks that during recovering of SARA it has disabled itself as there were more capture process running on the same database.
 - ◇ LHCb received a GGUS ticket about jobs not using CPU at T). Gavin: Pilot jobs in LSF not using CPU for 24h - fine but using up slots... Joel: could we get job IDs? LHCb will check
 - ◆ T1 site issues:
 - ◇ NIKHEF: Many pilot jobs aborting with exit status=999. This happens only against their CREAMCEs (GGUS:62132 [↗](#))
 - ◆ Massimo: discussing how to make it move evident when castor pool is user managed and needs cleanup (SLS will be more explicit). Joel: Still need to decide which data can be cleaned. Should to avoid piquet call as service person cannot fix the issue without LHCb.

Sites / Services round table:

- Andreas/KIT: ntr (by email)
- Gang/ASGC - ntr * Luca: DB replication to ASGC seems down - is there a h/w problem? Gang: still checking
- Jon/FNAL - ntr
- Michael/BNL - high load on postgres DB due to activity to optimise interface between pnf and postgres (write ahead buffer and checkpoint segment being adjusted). Now good understanding of the respective settings has been achieved.
- Onno/NL-T1 - NIKHEF issue of LHCb: cream developers have been notified. Yesterday some jobs running out of memory: SARA and Nikhef now looking into implementing a memory limit for jobs.
 - ◆ Joel: Will this take into use of swap or real memory. Onno: not sure yet, but will circulate proposal for memory limit before deploying it. Joel: pls also contact VO in case of jobs with high memory footprint - VO could check if this is due to a problem with individual jobs.
- Alessandro/CNAF - ready to apply new batch config for improved cluster performance. In touch with VO's for scheduling the intervention.
- Rolf/IN2P3 - ntr
- Tiju/RAL - moved castor cms to local name server
- Gonzalo/PIC - LHCb disk only space: several pools at PIC are full since weeks which also caused sam test to fail. Roberto: will update the test code to handle lack of space properly.
- Kyle/OSG - open two ticket in GGUS
 - ◆ yesterday: BNL T1 SE not in CERN BDII
 - ◆ US CMS T1 is randomly disappearing in gridview
 - ◆ Ale: gridview problem could be because by missing entry in BDII. BNL is the only US site not present in bdii and sam tests - Maarten is already investigating
- Luca/CERN:
 - ◆ DB swap problem for CMS was caused by dataset bookkeeping application. Will review notification email list with CMS.
 - ◆ LHCb streams latency issue: indeed it turned out that during the SARA DB recovery the 'real

time downstream' optimization for LFC was switched off by mistake. The issue is now fully understood and has been resolved on Tue afternoon by the T0 DBA team.

AOB:

Thursday

Attendance: local(Peter, Harry(chair), Patricia, Roberto, Graeme, Alessandro, Jacek, Ewan, Gavin, Jan, Weizhen);remote(Tobias(KIT), Jon(FNAL), Gonzalo(PIC), Michael(BNL), Kyle(OSG), Loel(LHCb), Alessandro(CNAF), Gareth(RAL), Onno(NL-T1), Rolf(IN2P3), ASGC).

Experiments round table:

- ATLAS reports -
- Tier-0: Commissioning and testing. No data export to T1s.
- T1s:
 - ◆ Problem at NIKHEF last night with pilots using an ATLAS release version of python which was deleted while the jobs ran, GGUS:62153 [↗](#). This was our fault. Pilot factories are now updated to the latest wrappers, which use the system python. Apologies to NIKHEF.
 - ◆ ASGC: Network problem caused transfer and job failures this morning, GGUS:62161 [↗](#). Solved.
 - ◆ INFN-T1: Transfer problems (Globus XIO errors) from BNL and Liverpool, GGUS:62167 [↗](#) (initially assigned to Liverpool).
 - ◆ RAL "at risk" for nameserver switch this morning. Few errors seen, but nothing significant.
 - ◆ Ongoing BNL-CNAF networking GGUS:61440 [↗](#): Thanks to Michael for yesterday's information which we added to the ticket.
 - ◆ Ongoing RAL-NDGF networking GGUS:61306 [↗](#). No updates. Transfer efficiency is still extremely low. Gareth reported that tests yesterday indicate a problem as packets pass through CERN.
- T2s:
 - ◆ IL-TAU-HEP low transfer efficiency to/from SARA, GGUS:62109 [↗](#). Fixed.
 - ◆ ITEP 100% failures with SRM down, GGUS:62122 [↗](#). Fixed.
 - ◆ CSCS-LCG2 dead dcap door, GGUS: 62148. Fixed.
 - ◆ UKI-LT2-QMUL out of downtime, in testing.
 - ◆ UKI-NORTHGRID-LANCS-HEP out of downtime, in testing.
 - ◆ CYFRONET-LCG2 poor transfer efficiency to INFN-T1 and jobs now affected, GGUS:62127 [↗](#). In progress.
 - ◆ RU-PROTVINO-IHEP, GGUS:62086 [↗](#). Still low efficiency and no answer on ticket.
- CMS reports -
- CERN
 - ◆ follow up on GGUS ticket 61706 (CMS Job Robot failures at CERN) : since yesterday, there is again a high rate of JR failures due to Maradona errors. Ticket has been updated.
- Tier1 issues
 - ◆ FZK-LGC2 : large failure of submission on CREAM CE (see GGUS ticket 61951 from Sep 9).
 - ◇ Need confirmation from site admins that the issue was due to old version of CREAM and util-java. CMS believes that what caused the problem is the fact that the CMS Operator had a "/" in a DN field, which causes the crash with the old CREAM CE +

util-java version (hence potentially also dangerous because the glidein-WMS has a / in a DN field too).

◇ If confirmed by site, CMS will recommend all sites to upgrade to the latest CREAM and util-java version. Tobias confirmed that their creamce1 was up to date while two others run the previous version. Experts are in Amsterdam meeting so will discuss with CMS how to upgrade next week. Peter informed that WMS glide-in jobs will have the same problem - he will take a look at the ticket.

- Tier2 Issues: Nothing to report
- MC production
 - ◆ Large 1000M events production at T1s and T2s: so far CMS Computing received requests for 451M events and ~436M events are already produced. The other half of the requests are yet to come.
- ALICE reports [↗](#) -
- GENERAL INFORMATION: Usual reconstruction, user analysis and MC activities ongoing
- T0 site: No issues have been found in the last 48h with the new CREAM module running at CERN and since yesterday at Subatech. The new module gets rid of any query to the CREAM-DB using as a single information system the resource BDII. We are planning to extend this new module to several T2 sites after Subatech to continue the testing. Planning to use this version in the next Alien release expected mid-October.
- T1 sites: IN2P3-CC: ticket 62118 submitted yesterday. SOLVED
- T2 sites: Currently working on the setup of a new site in Pakistan and for the rest of the T2 sites, the usual operations with no remarkable issues
- LHCb reports -
- Experiment activities: User analysis and few reprocessing running to the end. Smooth operations. Fixed the SRM critical tests that was failing at pic, CERN and CNAF. A first check on the availability of the space is performed before running (and then assessing) the endpoint. CNAF now going green again.
- T0 site issues: Requested CASTOR to move 40 TB from the newly created space token lhcbdst to the space token lhcbmdst that suffered lack of space in the last weeks. Despite some cleaning has been performed the lhcbmdst is still not fully pledged (309TB installed, 350TB required for 2010) while the lhcbdst is far too much (87TB).
- T1 site issues: NIKHEF: Many pilot jobs aborting with exit status=999. This happens only against their CREAMCEs (GGUS:62132 [↗](#)). They are working with the CREAM developers.

Sites / Services round table:

- BNL: Michael gave an update on BNL dropping out of the bdii as having been due to receiving an updated mixed case site spelling from a VO to which the EGEE bdii is sensitive and hence rejected the update coming from the OSG interoperability bdii. This was corrected and the site reappeared in the top level bdii.
- CNAF: Planning a scheduled configuration change on 24 September - will announce soon. Updated a creamce which then no longer appeared in the bdii - trying to understand the problem.

- RAL: Migration of CASTOR nameserver to local nameservers now done for all 4 experiments in preparation for CASTOR 2.1.9 upgrades. First will be for LHCb at the end of September. Then a CMS disk server that had been out of production was put back but started giving fsprobe errors so has been taken out again. Local CMS support has been contacted. Finally there will be an upgrade of the non-LHC firewall next Tuesday - will be put into Goddb today.
- NL-T1: Have just entered an unscheduled downtime with outside access blocked following announcement of a new kernel security vulnerability in both Redhat 5 and centos via ssh and a stack overflow. The exploit has been made public. Report is CVE-2010-3081.
- OSG: Investigating if their bdii issues are due to something in the CERN setup or inherent in the new version of the bdii software.
- ASGC: Had a network switch power outage at 07.00 UTC. Recovered in 10 minutes but hundreds of jobs failed.
- CERN lxplus/lxbatch: CERN is developing it's own kernel patch for CVE-2010-3081 though it expects one from Redhat today. This will be applied to lxplus when ready (expected shortly) then lxbatch as soon as possible as jobs drain. Given that beam is now expected on Saturday as much as possible will be done tomorrow especially for the Tier-0 workers. There will be service instabilities due to the urgency.
- CERN CASTOR: A CMS prodlogs diskserver (many small files) is showing problems so is being drained. There is sufficient space in the pool elsewhere.

AOB:

Friday

Attendance: local(Peter K, Harry(chair), Nicolo, Graeme, Jan, Ewan, Simone, Roberto, Jacek);remote(Tobias(KIT), Gang(ASGC), Joel(LHCb), Michael(BNL), Gonzalo(PIC), Jon(FNAL), Rolf(IN2P3), Onno(NL-T1), Andrea(CNAF), Gareth(RAL)).

Experiments round table:

- ATLAS reports - Reduced shifts until Saturday morning. No data export to T1s.
- T0/CAF: CAF patched and rebooted. T0 50% done (13:00 CERN time). Rest should be done now.
- Central Services: Patch and reboot went smoothly - DDM Central Catalog, DDM Site Services, Panda Servers
- T1
 - ◆ Downtimes at NIKHEF, SARA, KIT, NDGF, CC-IN2P3.
 - ◇ News of any facilities which might be down at the weekend would be welcome.
 - ◆ BNL namespace overload in MCDISK, GGUS:62206 [↗](#). Solved.
 - ◆ INFN-T1 transfers timing out, GGUS:62167 [↗](#). Slow network link from INFN-T1 - Liverpool? Being investigated.
 - ◆ Ongoing BNL-CNAF networking GGUS:61440 [↗](#): No change since yesterday.
 - ◆ Ongoing RAL-NDGF networking GGUS:61306 [↗](#): No change since yesterday.
- T2:
 - ◆ Many clusters in unscheduled downtimes for reboots
 - ◇ Friendly reminder to use downtimes and communicate with your cloud support.

- ◆ UKI-NORTHGRID-LANCS-HEP back online.
 - ◆ TOKYO-LCG2, network link unstable, GGUS:62221 [↗](#). Solved.
 - ◆ UKI-SOUTHGRID-RALPP offline, trying to recover a dCache pool. GGUS:62208 [↗](#).
 - ◆ CYFRONET-LCG2, GGUS:62127 [↗](#). Ticked marked "In progress", but no update for 2 days.
 - ◆ RU-PROTVINO-IHEP, GGUS:62086 [↗](#). Site reported problem with WAN link, seem ok now.
- CMS reports -
- T0:
 - ◆ Known Kernel issue keeping everyone busy...
 - ◇ CMS Core Operators were informed to expect service and site instabilities in their workflows during the coming hours/days.
 - ◇ draining of dedicated Tier-0/CAF queues + upgrade/reboot started Thursday night and finished as expected on Friday at 12:00. No major incidence on CMS operations.
 - ◇ upgrade/reboot of all CMS VoBox's on-going, managed by CMS VOC. CMS would like to get as much as possible done before the start of LHC collisions running foreseen for Sunday.
 - ◆ Clean-up of Castor Pool CMSPRODLOGS:
 - ◇ The CMS Tier-0 team has been requested to take action.
 - ◇ Usually log files are flushed on a 6 months basis, but given the need it will be done now
 - ◆ Follow up on GGUS ticket 61706 (CMS Job Robot failures at CERN):
 - ◇ tests are back to 100% since Sep 17, 01:00 AM.
 - ◇ Concerning the failures the day before, trouble shooting on-going, checks done on CERN/IT side + CMS sent the list of jobs that failed with Maradona errors in the last 2 days.
 - ◆ Follow up of CMSR failure from Monday when one overloaded service went down and affected others:
 - ◇ CMS and IT experts are currently discussing on how to improve dependencies of various CMS Services on CMSR; in particular for DBS, on how to reduce "JDBC" connections to the DB, while for Frontier, on the tuning of the "tcp_keepalive_time" parameter.
 - ◇ More news next week.
 - Tier1 issues:
 - ◆ Several sites expressed the need for a central recommendation regarding the general kernel issue. Harry reported that the CERN security service had mailed information and invited questions to all NGI and WLCG site security contacts.
 - ◇ As of Friday Sep 17, noon, 14 CMS sites (T1+T2) went in Downtime for that matter
 - ◆ Large re-RECO and Skims of 2010 CMS collision and cosmic data starting now. However, a few Tier-1s are/will be out :
 - ◇ FZK-LCG2 : for the known kernel issue
 - ◇ INFN_T1 : for the known kernel issue + queues to be drained during week-end for scheduled downtime (Storm upgrade) on upcoming Monday
 - ◇ IN2P3_CC : for a scheduled a downtime on Tue Sep 21; not clear when queues draining will start (???). Rolf replied that draining of very long jobs will start on 19 September preparing for the complete outage on 21st September. A mail with details has been sent.
 - ◇ so for the moment, re-processing planed at FNAL, PIC, ASGC and RAL (while more downtimes might follow due to the kernel issue)
 - Tier2 Issues:
 - ◆ Same remark as above by several Tier-2s about the need for central guidance regarding the general kernel issue

- MC production:
 - ◆ Large 1000M events production at T1s and T2s on-going
- AOB
 - ◆ CMS will continue Offline Computing Operations during the Christmas Shutdown at CERN, in reduced mode
 - ◇ only 1-8-hour Computing shifter / day
 - ◇ Computing Run Coordinator (24/7 on-call) maintained, but not required to be at CERN
 - ◇ still need to organize the operation mode in terms of expected attendance with all Sites admins and Service Application managers
- ALICE reports [↗](#) -
- T0 site: Good performance of the services, nothing to report
- T1 sites: INFN-T1: GGUS: 62211. Local ce07 (CREAM-CE) reporting wrong information in the info system. Also reported to the Alice contact person
- T2 sites:
 - ◆ Poznan: Site out of production (operations with the local VOBOX)
 - ◆ Usual operations with the rest of the sites
- LHCb reports -
- Experiment activities: User analysis mainly. PPG is currently defining the plan w.r.t. trigger; therefore even if LHCb take data, they should first commission the new workflow with stripping10 (10th reprocessing). Towards a global shutdown of the Grid with the security issue ?
- T0 issues: none. J.Closier is consulting the service providers on the LHCb VOboxes with a view to security patching/rebooting before the weekend.
- T1 site issues:
 - ◆ RAL: half of the jobs are stalling. Shifter will look at the problem
 - ◆ SARA: 92% of the (production) jobs aborted resolving tURL. Also SAM reflected the issue with the remote SRM discovered to be in downtime.
 - ◆ PIC: 20% of failure rate due to a variety of problem related with the storage. Shifter will look at the problem.

Sites / Services round table:

- BNL: Working on optimisation of the performance of their postgres database while there is no fresh raw data - a complex issue that requires full exposure of the SE to the typical ATLAS workload and data replication activities. Most likely explains some transfer failures observed by ATLAS last night..
- IN2P3: Down due to the security issue and currently meeting to decide if they come back before Monday.
- NL-T1: Nikhef is up and running apart from the VOboxes which will not get the security upgrade before Monday. At SARA they are still upgrading (which explains the srm errors seen by LHCb) but should be back this afternoon. They are using the CERN kernel mods (no official release from Redhat yet) so cannot guarantee stability. Ewan reported CERN had had a problem that was fixed this morning. (It was in the CCISS disk driver used on some newer HP boxes.)

- CNAF: Are in scheduled downtime till Monday. Storage should stay available but other services not yet known. Do not want to use the CERN kernel as they need gpfs support but may try it if nothing comes from Redhat over the weekend. They will also be upgrading LSF and Oracle next week.
- RAL: After yesterdays CMS disk server failure they have found 30 non-migrated files with checksum errors. Local CMS support is checking files that were migrated. They are working with E.Martelli at CERN on the problematic link to NDGF.
- CERN lxplus/lxbatch: All lxplus and T0 batch have been patched/rebooted. Some 50% of lxbatch is done and the rest are draining to be done on Monday.
- CERN CASTOR: The ATLAS T3 cluster has lost a disk server containing non-migrated files - still being worked on. The CASTOR disk servers have no users so should not be exposed to the kernel bug but they will be working with the experiments to coordinate their upgrades.

AOB: Harry pointed out that Tier-1 operations are likely to be difficult if collisions are resumed on Sunday.

-- JamieShiers - 06-Sep-2010

This topic: LCG > WLCGDailyMeetingsWeek100913

Topic revision: r12 - 2010-09-17 - MichaelErnst



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

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