

Table of Contents

Week of 111114.....	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 111114

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(Jamie, Torre, Ignacio, Massimo, David, Manuel, Nicolo, Maarten, Raja, Ale);remote(Michael, Gonzalo, Lisa, Jhen-Wei, Kyle, Onno, Rolf, Gareth, Maria Francesca, Douglas, Pavel).

Experiments round table:

- ATLAS reports -
- Tier-0
 - ◆ Heavy Ion running started 8am Sat. Morning, one fill of 2 bunches.
 - ◆ Issues with updates of the on-line condition replication, the databases continued to get out of sync for hours on Sat. morning. Turned out to be an upload of trigger rates for older runs, unfortunately done at the same time as new data. This threatened the processing of the first HI run.
 - ◆ More issues of SRM contact at CERN-PROD to EOS. Looks like this was reported quick enough that a ticket didn't even get made. This is a continuing problem.
 - ◆ Problems with access of files in CASTOR, using xrootd, discussion of this has been tracked by Alessandro Di Girolamo. GGUS:76283. [Massimo - being tracked between us and Ale. Looks as if this is the expected behaviour. File is on disk that should not be accessible to normal user. Ticket will probably be closed with instructions this afternoon.]
 - ◆ There was a lost of server in CASTOR (atlt3) late last week, we are waiting for a list of effected files, which should be reported today. GGUS:76237, GGUS:76250. [Massimo - we

are still draining the machine which had completely crashed over w/e so could not complete the draining. A full disk has gone - a non-tape-backed-up pool so files are gone. Files can be declared lost or you can give us 24h to try to recuperate something. Decision on ATLAS side. Torre/Ale - 24h extension perfectly fine.]

- ◆ There have been cases of spam messages in GGUS, one came in the weekend, GGUS:76282. This ticket seems to have been removed, but e-mail announcements of it when out to people. [Maarten - do GGUS people know about it? Please open a GGUS ticket(!)]
 - Tier-1
 - ◆ Remaining SRM config. issues from the BNL outage from last week. The SRM service on Friday needed to be re-config a couple times, and restarted. On Sat. early morning, new problems arised, and the SRM service on Sat. needed some changes and a restarted again. GGUS:76268. For Sat. night, and Sun. no new transfer issues to BNL were reported. [Michael - T0export which is managed by ADC people at CERN; channel was not opened after intervention, just an operational issue.] [Michael - two incidents, one Friday evening Eastern time, observed a performance degradation to almost an outage. Required a restart. Same happening Saturday morning. Noticed that JVM associated with SRM server ran out of memory so we allocated more memory to JVM. Also noticed some slowness on SRM DB. Significant wait I/O. Decided to prepare faster h/w Friday night. Switched after other incident on Saturday morning and re-indexed SRM DB. These three changes helped significantly - don't see any significant load at the moment even though hefty data replication from many sites to BNL. For last 24h observed 112TB replicated to BNL datadisk which translates to 1.3GB/s bandwidth averaged over 24h. 200K files - as we see no load in pretty good shape!]
 - ◆ There was a 10% failure rate of jobs to INFN-T1 last week, this should be helped by a new pilot version to be released today. This was also caused by a network problem between CERN and CNAF, changes were made on Friday night, that should fix this issue.
 - ◆ There will be an FZK outage for ATLAS this week, for the LFC migration of the DE cloud. This will start tomorrow morning, and continue until Thur. morning. There will be a draining of queues starting tonight. [Pavel - confirmed that this is correct. Whole DE cloud will be offline in DDM and Panda during this time. Will take until Wednesday midnight. COMPASS vo might also be affected.]
 - ◆ We would like to test transfers from CERN EOS to Tier-1s without use of SRM. We would like Tier-1 sites to verify on your FTS servers, if the endpoint eosatlassftp.cern.ch is included as an endpoint belonging to CERN-PROD. The changes made to CERN-PROD can be seen here, GGUS:76285. For more info, this issue will be managed by Alessandro Di Girolamo.
-

- CMS reports -
- LHC / CMS detector
 - ◆ Heavy Ions collision data taking since early Saturday morning.
- CERN / central services
 - ◆ NTR
- T0
 - ◆ Running HI express and prompt reconstruction.
 - ◆ Requesting that worker nodes migrated from LSF cmscaf to tier0 queue for pp are returned to cmscaf queue for Heavy Ions running: GGUS:76299
- T1 sites:
 - ◆ MC production and/or reprocessing running at all sites.
 - ◆ Run2011 data reprocessing.
 - ◆ Heavy Ion data subscribed to T1_FR_CCIN2P3. Subscription to T1_US_FNAL waiting for tape family deployment.
 - ◆ [T1_IT_CNAF] GGUS:76090 - files stuck in migration because they were produced by CMS with wrong LFN pattern: asked CNAF to define migration policy for these files, any update?
 - ◆ [T1_IT_CNAF] GGUS:76269 - JobRobot failures on Sat 12th for issues on one CE, fixed and

closed.

- ◆ [T1_DE_KIT]: Debugging network issues with US T2s GGUS:75985 - any update?
 - ◆ [T1_FR_CCIN2P3]: GGUS:76307 - issues with AFS software area, affecting SAM tests, JobRobot and SW deployment - in progress.
 - ◆ [T1_FR_CCIN2P3]: investigation on networking issues in progress GGUS:75983 GGUS:71864 [Rolf - still interested in establishing iperf tests between CALTECH and IN2P3. Info in 71864. Nico - ticket updated (linked Savannah) by CMS asking CalTech to run some tests - will contact them again]
 - ◆ [T1_FR_CCIN2P3]: GGUS:75829 about proxy expiration in JobRobot test: happened again on Nov 5th and Nov9th on cccreamceli02.in2p3.fr - also affecting T2_BE_IIHE SAV:124597 and SAV:124471 - CREAM developers involved GGUS:76208
 - ◆ [T1_TW_ASGC]: GGUS:75377 and GGUS:76204 about read and stageout errors on CASTOR, in progress
- T2 sites:
 - ◆ [T2_US_Vanderbilt]: enabled direct T0_CH_CERN-->T2_US_Vanderbilt link to transfer data for prompt skimming.
 - Other:
 - ◆ Josep Flix will be CRC starting from tomorrow.
-

- ALICE reports [↗](#) -
 - ◆ On Sunday it was discovered that AliEn is not permitted to write raw HI data to the tape back-end at RAL; the matter is being debugged by RAL admins and AliEn and Xrootd experts.
-

- LHCb reports -
 - Experiment activities
 - ◆ Finishing up on reprocessing last but one range of data over the week-end
 - ◆ Hopefully tomorrow the last reprocessing productions will be launched (new ConDB tag, distributed to the Tier-1s).
 - T0
 - ◆ ntr
 - T1
 - ◆ IN2P3: Significant improvement after reducing stripping load there. Hope the reprocessing backlog is finished by tomorrow morning.
 - ◆ IN2P3 : (GGUS:75158) : migration of files to the correct space token - still ongoing.
 - ◆ RAL : (GGUS:76295) : Pilots aborted at RAL across all CEs.
 - ◆ PIC : Many user jobs failing to resolve input data indicating possible SRM problem - waiting to understand in more detail before following up with site.
-

Sites / Services round table:

- BNL - nta
- PIC - ntr
- FNAL - ntr
- ASGC - ntr
- NL-T1 - 2 announcements, Nov 21 SARA MSS will be in maintenance, files from tape will have to be staged before. Not yet in GOCDB - will be done as soon as all details sorted out. SARA SRM will be on maint on Dec 12. Raja - how long will tape maint be? A: 2-3h but please check GOCDB when entry is there
- IN2P3 - nta

- RAL - main thing is that tomorrow morning outage of site while update of firewall. 08:00 - 08:30 UTC. 2 items from before; ALICE Q is ongoing, LHCb ticket - there has been an update on ticket.
- NDGF - ntr
- KIT - comment on network issues; will forward request from CMS to network team; tomorrow downtime which will affect ATLAS and prob COMPASS related to LFC migration and dCache upgrade to golden release
- OSG - ntr

- CERN - ntr

AOB:

Tuesday:

Attendance: local(Torre, Jamie, Maarten, Ignacio, David, Pepe, Eva, Alex);remote(Michael, Jeremy, Burt, Tiju, Jhen-Wei, Maria Francesca, Rolf, Ronald, Raja, Rob, Dimitri).

Experiments round table:

- ATLAS reports -
 - Tier-0
 - ◆ Heavy Ion running proceeding
 - ◆ CERN-PROD problems accessing files via SRM and xrscp: solved and verified. xroot redirectors did not have the proper "localgridmap" entries for ATLAS power users. Same issue was discovered and solved 2 months ago, but patch lost; configuration now saved persistently. GGUS:76283
 - ◆ Anticipating today a list of affected files from loss of server in CASTOR (atlt3) late last week. GGUS:76237, GGUS:76250. [Ignacio - Massimo provided list for ATLT3 files to Alessandro and Guido but would like to know if these files could be found somewhere else.]
 - Tier-1
 - ◆ Planned FZK outage for ATLAS underway, for the LFC migration of the DE cloud. DE cloud taken offline last night to drain, blacklisted this morning in DDM.
 - ◆ Brief interruption in CERN-BNL transfers last night due to wrongly replaced X509 host certificate at BNL causing authorization failures. Quickly resolved. GGUS:76361
 - ◆ Planned RAL firewall intervention this morning proceeded without incident, connectivity break less than 15min
 - ◆ Planned PIC firewall intervention 7:30-8:30 this morning. Followed by (unrelated?) problems with PT site connectivity. Report from PIC midday that "Spanish network provider (RedIRIS) has a problem in Seville. That's why there is no connection between PIC and PT sites. The issue is reported and actions are anticipated."
-

- CMS reports -
- LHC / CMS detector
 - ◆ Heavy Ions collision data taking since early Saturday morning.
- CERN / central services
 - ◆ NTR
- T0
 - ◆ Running HI express and prompt reconstruction.
 - ◆ Requesting that worker nodes migrated from LSF cmscaf to tier0 queue for pp are returned to cmscaf queue for Heavy Ions running: GGUS:76299 (This ticket should be for LSF, not CASTOR --> Ignacio assigned it to LXBATCH 2nd Line Support)
- T1 sites:
 - ◆ MC production and/or reprocessing running at all sites.

Tuesday:

- ◆ Run2011 data reprocessing.
 - ◆ Heavy Ion data subscribed to T1_FR_CCIN2P3. Subscription to T1_US_FNAL waiting for tape family deployment. [solved]
 - ◆ [T1_DE_KIT]: Debugging network issues with US T2s GGUS:75985 - any update?
 - ◆ [T1_IT_CNAF] GGUS:76090 - files stuck in migration because they were produced by CMS with wrong LFN pattern: asked CNAF to define migration policy for these files. **Ticket closed.**
 - ◆ [T1_IT_CNAF] GGUS:76426 - files stuck in migration since 9 days. New ticket, some blocks are not closing... checked some files via lcg-ls and we get all the time: "[SE][Ls][SRM_INVALID_PATH] No such file or directory". Asked the site to check the files...
 - ◆ [T1_FR_CCIN2P3]: GGUS:76307 - issues with AFS software area, affecting SAM tests, JobRobot and SW deployment. AFS went back, but the problem was not properly understood from the site. **Ticket closed.** We can re-open if issues re-appear.
 - ◆ [T1_FR_CCIN2P3]: investigation on networking issues in progress GGUS:75983 GGUS:71864
 - ◆ [T1_FR_CCIN2P3]: GGUS:75829 about proxy expiration in JobRobot test: happened again on Nov 5th and Nov9th on cccreamceli02.in2p3.fr - also affecting T2_BE_IHHE SAV:124597 and SAV:124471 - CREAM developers involved GGUS:76208
 - ◆ [T1_TW_ASGC]: GGUS:75377 and GGUS:76204 about read and stageout errors on CASTOR, in progress
 - T2 sites:
 - ◆ [T2_US_Vanderbilt]: enabled direct T0_CH_CERN-->T2_US_Vanderbilt link to transfer data for prompt skimming.
 - Other:
 - ◆ Hanging SAM Nagios test in T2_ES_CIEMAT (probably other sites affected): GGUS:76432 [↗](#). Temporary problem in a message broker causing CMS Nagios tests to hang. SAM tests were trying to use the alias egi-2.msg.cern.ch, which pointed to gridmsg106.cern.ch, which was in maintenance this morning. Maybe the machine suffered from an unexpected problem and was not removed from the alias quickly enough? We don't see any information in ITSSB (<http://itssb.web.cern.ch/service-incidents> [↗](#)).
-

- ALICE reports [↗](#) -
 - ◆ The problem at RAL reported yesterday has not yet been resolved; experts will continue looking into it.
-

- LHCb reports -
 - Experiment activities
 - ◆ Next round of reprocessing productions have just been launched.
 - T1
 - ◆ IN2P3: Slow decrease of the reprocessing backlog (not yet finished).
 - ◆ PIC : Many user jobs failing to resolve input data indicating possible srm problem. Failure rate ~25% at present. LHCb PIC contact looking at it
-

Sites / Services round table:

- BNL - nta
- FNAL - ntr
- RAL - ntr
- ASGC - ntr
- NDGF - ntr
- IN2P3 - ntr

- NL-T1 - ntr
- KIT - asked again for update from network group for CMS ticket - hopefully they will update it today
- GridPP - ntr
- OSG - ticketing changes that allow attachments to be exchanged will be tested for 1 week in integration systems; will start with BNL and eventually FNAL; will go in production 1 week if AOK
- CERN - ntr

AOB:

- LHC talk this Thursday at T1SCM from Mike Lamont. "LHC - Future Running Scenarios" : <http://indico.cern.ch/conferenceDisplay.py?confId=162564>

Wednesday

Attendance: local(Jamie, MariaDZ, Maarten, Edoardo, Alessandro, Jacek, David, Pepe, Raja, Torre, Ignacio, Manuel);remote(Michael, Ron, Lisa, Gonzalo, Jhen-Wei, Tiju, Maria Francesca, Rolf, Rob, Xavier).

Experiments round table:

- ATLAS reports -
- Tier-0
 - ◆ Heavy Ion running proceeding
 - ◆ Received from Massimo a list of affected files from loss of server in CASTOR (atlt3) late last week. At Massimo's request we are investigating availability of replicas, will report back on any we would like the Castor team to try to recover. GGUS:76237, GGUS:76250.
- Tier-1
 - ◆ Planned FZK outage ongoing as scheduled (at time of writing; due to complete at noon today) for the LFC migration of the DE cloud.

-
- CMS reports -
 - LHC / CMS detector
 - ◆ Heavy Ions collision data taking since early Saturday morning. Later today, p-Pb Collisions.
 - CERN / central services
 - ◆ (<http://itssb.web.cern.ch>): Tonight "External DNS problem" affected SAM tests, JobRobot and some transfers at the sites, for about 1h.
 - T0
 - ◆ Running HI express and prompt reconstruction.
 - ◆ Nodes which migrated from LSF cmscaf to tier0 queue for pp are returned back to cmscaf queue for Heavy Ions running: GGUS:76299 (closed). 48 machines (384 cores) added. Total slots atm 624.
 - T1 sites:
 - ◆ MC production and/or reprocessing running at all sites.
 - ◆ Run2011 data reprocessing.
 - ◆ [T1_DE_KIT]: Debugging network issues with US T2s GGUS:75985. IPerf tests ongoing with some US-sites.
 - ◆ [T1_IT_CNAF] GGUS:76426 - files stuck in migration since 9 days as some files are not existing. In fact, the problem was related to a crash of the PhEDEx agents running on a VM and the movement of the agents to another VM. This created an interference, probably not all agents were full dead in the original machine, and some files were published to be at the site, while they were indeed deleted by the new instance... CMS already deleted all files from CNAF for retransfer.

- ◆ [T1_IT_CNAF]: GGUS:76491. Trying to understand why the BlockDownloadVerify PhEDEx agents get stuck from time to time...
 - ◆ [T1_FR_CCIN2P3]: investigation on networking issues in progress GGUS:75983 GGUS:71864 [Rolf - in contact with Renater about the issue with CalTech]
 - ◆ [T1_FR_CCIN2P3]: GGUS:75829 about proxy expiration in JobRobot test: happened again on Nov 5th and Nov9th on ccreamceli02.in2p3.fr - also affecting T2_BE_IIHE SAV:124597 and SAV:124471 - CREAM developers involved GGUS:76208
 - ◆ [T1_TW_ASGC]: GGUS:76498. 'Could not submit to FTS' errors for data transfers to the site. Is the proxy expired in the VOBOX?
 - ◆ [T1_TW_ASGC]: GGUS:75377 and GGUS:76204 about read and stageout errors on CASTOR, in progress... Any news on these old-long-standing-tickets? [Jhen-Wei : these problems should be fixed; please verify; also replied about proxy; did expire - trying to replace. Will send mail offline]
 - T2 sites:
 - ◆ NTR (at least, relevant for this meeting)
 - Other:
 - ◆ Hanging SAM Nagios test in T2_ES_CIEMAT (probably other sites affected): GGUS:76432. Temporary problem in a message broker causing CMS Nagios tests to hang. SAM tests were trying to use the alias egi-2.msg.cern.ch, which pointed to gridmsg106.cern.ch, which was in maintenance this morning. Maybe the machine suffered from an unexpected problem and was not removed from the alias quickly enough? We don't see any information in ITSSB (<http://itssb.web.cern.ch/service-incidents>). Ticket assigned, but no further comments.
-

- ALICE reports [↗](#) -
 - ◆ The problem at RAL has not been solved yet. There is a problem with alienmaster's permissions. Ongoing (looks like something wrong on alien side)
-

- LHCb reports -
 - Experiment activities
 - ◆ Next round of reprocessing productions going on.
 - ◆ Expect to launch next round of MC simulations by end of November.
 - T1
 - ◆ IN2P3: Old reprocessing backlog down to ~200 jobs now. Problems with access to conditions DB at IN2P3 - GGUS ticket being submitted. [Rolf - did you already open or will do? Raja - Joel and new IN2P3 contact are in the process of doing this now but not yet opened]
 - ◆ PIC : Many user jobs failing to resolve input data indicating possible srm problem. Failure rate ~25% at present. LHCb PIC contact looking at it. Also ongoing investigations of very high memory consumption by user jobs.
-

Sites / Services round table:

- BNL - ntr
- NL-T1 - ntr
- FNAL - ntr
- PIC - ntr; acknowledge what was said by Raja; seeing pretty high rate of errors and investigating
- ASGC - nta
- RAL - ntr
- NDGF - ntr
- IN2P3 - nta
- KIT - ntr
- OSG - ntr

- CERN network - problem with DNS last night; Due to vulnerability of s/w used for DNS. No fix. Producer of s/w working on a fix. Problem may be an attack or by chance - several reports of problems at same time on Internet so may happen again. Error message in logs but DNS fails; From outside is ok as 2nd-ary DNS at Switch. External sites could resolve names except in sub-domains. Dynamic load balancers lost. Ale - didn't report; production noticed a small drop in # running jobs. Manual operation to restart. Asking why Switch server did not get affected. Was not a general attack.

AOB:

- T1SCM tomorrow; special talk on future LHC running
- (MariaDZ) Service managers who receive ALARM notifications via sms asked to be spared from the formatted frames with asterisks in order to see immediately which service is affected. This will be done with the next monthly GGUS Release. Details in Savannah:124169

Thursday

Attendance: local(Kasia, Manuel, Pepe, Jamie, Maria, Torre, Andrea, Stephane, Massimo, Maarten, Alessandro, MariaDZ, Ian);remote(Gonzalo, Lisa, John, Ronald, Jhen-Wei, Joel, Aresh, Andreas, Rolf, Maria Francesca, Rob, Paolo Franchini).

Experiments round table:

- ATLAS reports -
 - Tier-0
 - ◆ Heavy Ion running proceeding
 - ◆ Alarm ticket placed yesterday afternoon for slow AFS response time on volumes critical for Tier-0 operation -- addressed promptly, closed and verified. GGUS:76519 [Massimo - this was an overload which was solved almost spontaneously. Trying to have a setup which would avoid this in the future. Volume in question was extremely small - co-location with another volume(s) caused problems. SIR-lite will be produced.]
 - ◆ SRM-EOSATLAS problems: many failed transfers between srm-eosatlas.cern.ch and different sites. Resolved and verified. GGUS:76507
 - ◆ Received the incident report we requested on Monday describing recent SRM-EOSATLAS instabilities, thank you
 - ◆ ~30min interruption of landline phone service at CERN on very short notice! Exemptions received for a few crucial shift phones. No noticeable impact on computing ops.
 - Tier-1
 - ◆ FZK outage for LFC migration completed after an extension (LFC migration operation took longer than anticipated), DE cloud fully online for production since 10am today
-

- CMS reports -
- LHC / CMS detector
 - ◆ Heavy Ions collision data taking since early Saturday morning. Maybe tomorrow p-Pb Collisions, but not certain atm, due to septum problems in PS.
- CERN / central services
 - ◆ NTR
- T0
 - ◆ Running HI express and prompt reconstruction.
- T1 sites:
 - ◆ MC production and/or reprocessing running at all sites.
 - ◆ Run2011 data reprocessing.

- ◆ [T1_DE_KIT]: Debugging network issues with US T2s GGUS:75985. IPerf tests ongoing with some US-sites. **re-opened the ticket**
 - ◆ [T1_IT_CNAF] GGUS:76426 - files stuck in migration since 9 days as some files are not existing. **after actions taken, it's solved**
 - ◆ [T1_IT_CNAF]: GGUS:76491. Trying to understand why the BlockDownloadVerify PhEDEx agents get stuck from time to time... In particular, it's not stuck, it's working hard doing checksums and the agent is not able to communicate to central that is alive... This agent wasn't designed to do checksums vastly. Asked clarification...
 - ◆ [T1_FR_CCIN2P3]: investigation on networking issues in progress GGUS:75983 GGUS:71864 [Rolf - net work issue still worked on, transfer speed going up.]
 - ◆ [T1_FR_CCIN2P3]: GGUS:75829 about proxy expiration in JobRobot test: happened again on Nov 5th and Nov9th on cccreamceli02.in2p3.fr - also affecting T2_BE_IIHE SAV:124597 and SAV:124471 - CREAM developers involved GGUS:76208 [Rolf - Proxy: problem arrived only with 1 of the CEs which is in front of old batch system. We will accelerate migration and switch WNs still on this CE to the grid engine CEs so problem should have gone by next week.]
 - ◆ [T1_TW_ASGC]: GGUS:76498. 'Could not submit to FTS' errors for data transfers to the site. Proxy was expired in VOBOX. CMS VO-manager granted 'production' role to a new certificate for being used in ASGC and proxy was re-created. **closed**
 - ◆ [T1_TW_ASGC]: GGUS:75377 and GGUS:76204 **closed** about read and stageout errors on CASTOR, in progress... Concerning GGUS:75377, needs confirmation from CMS side. But the check will take some time since CMS is not running at full capacity at ASGC atm because of the CASTOR migration, and the ReReco is taking all available slots atm. The Redigi jobs which are the ones which use those "Corrupted Files" are behind the queue atm...
 - T2 sites:
 - ◆ NTR (at least, relevant for this meeting)
 - Other:
 - ◆ Hanging SAM Nagios test in T2_ES_CIAMAT (probably other sites affected): GGUS:76432. Confirmed there was a problem with the Message Broker (gridmsg106) on Monday between 08:00 and 11:00. The broker-list file is refreshed every hour, so we could hit the moment when the broker was down but still in use by Nagios. Clarified and **closed**.
-

- ALICE reports [↗](#) -
 - ◆ The problem with raw data storage on tape at RAL was resolved yesterday: "it was a detail in the envelope that was not correctly set any more, we've reverted the code and now the transfers look fine."
-

- LHCb reports -
- Experiment activities
 - ◆ Next round of reprocessing productions going on.
 - ◆ Expect to launch next round of MC simulations by end of November.
- T0
 - ◆ 10 disk server added to lhcb-tape
- T1
 - ◆ IN2P3: (GGUS:76515) not assigned after 18 hours . Is it normal ? [Rolf - will investigate]
 - ◆ IN2P3: (GGUS:76528) can not do a AFS release on ccage.in2p3.fr. All WNs now moved to cvmfs so no longer an issue.
 - ◆ RAL - issue this morning with network, seems again right now that there is an issue too. [John - this morning at 7:50 local time one of routers died. Was down 07:50 - 09:00. Power supply replaced and came back up. Not aware of any on-going problems. Lost about 300 jobs as a result. Andrea S - got report from GEANT-4 about not being able to submit jobs to CREAM CEs at RAL. (About 1h ago).

Sites / Services round table:

- PIC - ntr
- FNAL - ntr
- RAL - nta
- NL-T1 - ntr
- ASGC - ntr
- KIT - ntr
- IN2P3 - nta
- NDGF - ntr
- CNAF - ntr
- OSG - ntr

- CERN DB - migration of ATLAS agis application from INT to PRO. Ale - results also verified.

AOB:

Friday

Attendance: local(Torre, Maria, Jamie, Pepe, Raja, Massimo, David, Lola, Manuel, Eva);remote(Michael, Gonzalo, Ulf, Rolf, Onno, Lisa, John, Jhen-Wei, Xavier, Rob).

Experiments round table:

- ATLAS reports -
 - Tier-0
 - ◆ Heavy Ion running proceeding
 - ◆ DDM catalog cleanup throughout the week has led to high loads on ADCR, cleanup not yet completed but will be suspended for the weekend
 - Tier-1
 - ◆ Nothing to report
-

- CMS reports -
- LHC / CMS detector
 - ◆ Heavy Ions collision data taking.
- CERN / central services
 - ◆ Yesterday night around 23h, the CMSR instance 1 hanged and was manually rebooted. Load on instance was not significantly bigger yesterday. DB-experts suspect an Oracle bug related to the activity by the application running on node and statistics gathering. They will change the time when statistics are gathered to let them run during the day, to better check the hypothesis and make debugging easier.
- T0
 - ◆ Running HI express and prompt reconstruction.
- T1 sites:
 - ◆ MC production and/or reprocessing running at all sites.
 - ◆ Run2011 data reprocessing.
 - ◆ [T1_DE_KIT]: Debugging network issues with US T2s GGUS:75985. IPerf tests ongoing with some US-sites. **re-opened the ticket**
 - ◆ [T1_IT_CNAF] GGUS:76426 - files stuck in migration since 9 days as some files are not existing. **Closed**
 - ◆ [T1_IT_CNAF]: GGUS:76491. Trying to understand why the BlockDownloadVerify PhEDEx agents get stuck from time to time... **Agent was habilitated to checksum by**

mistake. Corrected. Closed

- ◆ [T1_FR_CCIN2P3]: investigation on networking issues in progress GGUS:75983 GGUS:71864
- ◆ [T1_FR_CCIN2P3]: GGUS:75829 about proxy expiration in JobRobot test: happened again on Nov 5th and Nov9th on cccreamceli02.in2p3.fr - also affecting T2_BE_IIHE SAV:124597 and SAV:124471 - CREAM developers involved GGUS:76208 **CMS will drain cccreamceli02.in2p3.fr over the next week from production tools to sync. with the decommissioning of the CE**
- ◆ [T1_TW_ASGC]: GGUS:76558. Some transfers to ASGC failing, caused by multiple VOs in Felix's DN. Some transfers are failing because DN mapping is incorrect (dteam and CMS). **Fixed and Closed**
- ◆ [T1_TW_ASGC]: GGUS:75377, needs confirmation from CMS side. But the check will take some time since CMS is not running at full capacity at ASGC atm because of the CASTOR migration, and the ReReco is taking all available slots atm. The Redigi jobs which are the ones which use those "Corrupted Files" are behind the queue atm... [Jhen-Wei yesterday increased CPU slots to 90% of capacity. Please check]
- T2 sites:
 - ◆ NTR (at least, relevant for this meeting)
- Other:
 - ◆ NTR

- ALICE reports [↗](#) -
 - ◆ Problem with two CAF nodes - have to be reinstalled; ready later today or Monday

- LHCb reports -
- Experiment activities
 - ◆ Next round of reprocessing starting to tail off. Backlog primarily now at IN2P3 which has most of the last round of data.
 - ◆ Still expect to launch next round of MC simulations by end of November.
 - ◆ Waiting for new templates from WLCG to give updated definitions of service criticality for LHCb.
- T1
 - ◆ CNAF : Job failures probably related to power cut there.
 - ◆ SARA : Early signs of job failures to resolve turls on WNs. Waiting to see how the situation develops.
 - ◆ dCache : GGUS ticket (76561) opened for assistance migrating data from one space token to another.
- T2
 - ◆ Shared software area problems at WCSS64(Poland:76569) and Auver(France:76586)

Sites / Services round table:

- BNL - ntr
- PIC - it seems that the issue that affected us for LHCb jobs having difficulties in accessing input data have been solved; configuration of pool costs. Solved this morning and we are watching it; success rate now much higher; Raja - should have put it in report
- NDGF - ntr
- IN2P3 - network ticket initially opened for CMS - has some news; please can CMS people look at this. Appears packet losses which created loss of transfer speed have disappeared; this is good news also for other T2 sites; also non-CMS. Pepe- will check
- NL-T1 - ntr; problem reported by LHCb please open a ticket!
- FNAL - ntr

- RAL - ntr
- ASGC - nta
- KIT - one tape lib down since yesterday; waiting for vendor; done just before this call and so should soon be able to work through backlog; 5 or 6 CMS disk-only pools crashed just before meeting; working on bringing them back online
- OSG - ntr
- CASTOR: **SIR-lite** for GGUS:76519:
 - ◆ **16-NOV-2011 17h30** An ATLAS alarm ticket on slow AFS response times impacting their T0 monitoring was created. The identified cause was a co-located LHCb volume, for which too many concurrent connections were opened between 17h00 and 17h26. The overloaded volume was moved off to another server. In general, server-side request throttling is in place to protect a share of the available threads in case of heavy access to a specific volume. However, all remaining clients continue to compete for the protected share of threads which can result in reduced response times in case multiple volumes see an very high number of concurrent connections. We have already in place mechanisms to confine volumes and we are studying how to implement this for critical volumes (in general confinement is used to quarantine volumes with problems, here it would be to protect given volumes from legitimate user activity on other volumes). The difficulty is that the volumes to be protected are very small (in this case ~ 2 GB) so we are studying ways to do it efficiently.
- CERN DB: ATLAS integration DB and LCG db will be migrated to new h/w next Tuesday. Will require 2 h downtime

AOB:

-- JamieShiers - 31-Oct-2011

This topic: LCG > WLCGDailyMeetingsWeek111114

Topic revision: r20 - 2011-11-18 - MaartenLitmaath



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