

Table of Contents

Week of 110620.....	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.....	6
Thursday.....	8
Friday.....	10

Week of 110620

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 (AndreaV, Ken, Lukasz, Dan, Gavin, Guido, Massimo, Luca, Lola, Raja, Dirk); remote (Weijen, Michael, Jon, Gonzalo, Tiju, Xavier, Ron, Rolf, Daniele, Christian, Kyle).

Experiments round table:

- ATLAS reports -
 - ◆ T0/CERN
 - ◇ Sunday night at 2am: FTS errors when exporting from CERN SE nodes in CERN-PROD_DATADISK: "globus_xio: System error in connect" (source), "Request timeout due to internal error" (destination). Problem disappeared after a very few hours (no GGUS ticket)
 - ◇ unavailability of the dashboard SAM test page (Savannah #83420); "The server of Dashboard virtual machines had a problem over weekend". Rebooted in the morning
 - ◆ reminding 2 interventions this week:
 - ◇ today 3pm: replacement of a fibre channel between ATONR and ATLR databases (IT announcement). [Luca: the intervention already took place this morning at 11am as requested by the ATLAS online experts.]
 - ◇ wednesday 2pm: enable the Transfer Manager replacing the LSF scheduler on castorcernt3 (affecting ATLAS and CMS local pools), down for 2h (IT announcement)

- ◆ T1
 - ◇ still open issue with FZK-LCG2_ATLASMCTAPE (GGUS 71466): buffer added to have more space, migration during the week end ran well, only one pool has a still rather large queue
- CMS reports -
 - ◆ LHC / CMS detector
 - ◇ Very little data since the last meeting.
 - ◆ CERN / central services
 - ◇ CERN-PROD disappeared from the BDII on Sunday. Not clear that there was any impact on operations, but we would like to know how this happened. [Ke: see GGUS:71679. Gavin: one BDII server got stuck, Ricardo will investigate with the BDII expert when he is back.]
 - ◇ This link, which shifters use to monitor the state of critical grid services, became unavailable over Saturday night. Sent a Savannah ticket to Dashboard on Sunday morning, but was not answered until Monday morning, at least in part because relevant person wasn't listed in the Savannah dashboard team. System was recovered reasonably quickly after that.
 - ◆ Tier-0 / CAF
 - ◇ Large number of CAF pending jobs over past few days, but working our way through them.
 - ◆ Tier-1
 - ◇ No news.
 - ◆ Tier-2
 - ◇ MC production and analysis in progress.
 - ◆ AOB
 - ◇ Oliver Gutsche starts as CRC tomorrow.
- ALICE reports -
 - ◆ T0 site
 - ◇ During the weekend some users complain about the AliRoot versions in CAF. An update was needed and last night everything was back to normal
 - ◆ T1 sites
 - ◇ CNAF : GGUS:71401. There are many JAs not converting into jobs. Moreover there are many jobs with CPU/Wall < 20%. Tracing some of the jobs it was found that there were AliRoot processes hanging. ROOT experts have been involved. Under investigation. [Andrea: is this the same issue reported last week with poor CPU efficiency? Lola: yes.]
 - ◇ IN2P3: trying out torrent software installation at the site [Lola: this is to test torrent instead of AFS]
 - ◆ T2 sites
 - ◇ Usual operations
- LHCb reports -
 - ◆ Experiment activities:
 - Problem with one DIRAC central service running out of space this early morning.
 - ◇ New GGUS (or RT) tickets:
 - T0: 0
 - T1: 0
 - T2: 0
 - ◇ Issues at the sites and services:
 - T0
 - CERN: Investigation on going on the low number of LSF job slots.

GGUS:71608 [↗](#). Problem seems to have been understood and pinned down to the algorithms used by LSF for shares.

· T1

- GRIDKA: GGUS:71572 [↗](#). Ticket escalated to ALARM as SRM became completely unresponsive and the site unusable. [Xavier: the SRM service died on Sat afternoon because the /var partition was full. After some cleanup the service was restarted this morning. The problem was due to high load and to a suboptimal setup for LHCb, will meet next week with LHCb experts to follow up.]

Sites / Services round table:

- ASGC: sheduled downtime for UPS intervention on Tue from 2 to 10am UTC
- BNL: ntr
- FNAL: ntr
- PIC: ntr
- RAL: no problems to report; tomorrow no one will connect because of a meeting
- KIT: nta
- NLT1: ntr
- IN2P3: problems with top BDII on Saturday; both servers crashed (while heavily swapping) and had to be rebooted several times; updated to the last version of BDII and will add some RAM
- CNAF: ntr
- NDGF: short downtime tomorrow to diagnose a problem with RAID controllers
- OSG: ntr
- Grid services (Gavin)
 - ◆ Please migrate the whole node scheduling tests to the new CREAM CE: creamtest001. The test CE lx7835 will be retired this week.
 - ◆ ATLAS are still using ce201 and ce202 (in drain and dure to be retired).
- Dashboard services (Lukasz)
 - ◆ Memory leak problems on dashboard12 (with virtual machines)
 - ◆ Monalisa server that checks dashboard was down
 - ◆ Problem also with ATLAS monitoring dues to one stored procedure that was slow (slower against the production DB than it was against the integration DB). [Luca: this is the LCGR production DB for all LCG dashboards].
- Storage services (Massimo/Dirk)
 - ◆ Upgrade of public Castor this morning, seeing problems with authentication that were not seen on CERN T3, could ATLAS and CMS have a look if anything changed?
 - ◆ A workaround (KRB5RCACHETYPE=none) was applied on the xrootd redirector on Friday to solve hangs in user jobs seen by ATLAS after upgrading to ROOT 5.28 with the fix for the KDC high load issue. [Summary by Andrea for the minutes of several other discussion threads:in other words the upgrade to ROOT 5.28 was necessary and reduced the load on the KDC, but it was not alone enough as these user job hangs started to appear; the workaround seems to work fine (no user job hangs with ROOT 5.28) and maybe is also useful in itself to reduce the high load (even with ROOT 5.26); this is essentially a workaround for the Kerberos 1.6 bug that still affects SLC5].

AOB: none

Tuesday:

Attendance: local(Alessandro, Dan, Gavin, Jacek, Julia, Maarten, Maria D, Massimo, Nilo, Oliver, Raja, Steve);remote(Christian, Felix, Gonzalo, Jon, Karen, Kyle, Michael, Rolf, Ronald, Xavier).

Tuesday:

Experiments round table:

- ATLAS reports -
 - ◆ T0/CERN:
 - ◇ Follow up from yesterday: the panda queues for CERN ce201 and ce202 have been turned off.
 - ◇ Outage yesterday at CERN to upgrade CASTORPUBLIC to 2.1.11 was declared to affect the SRM endpoints dedicated to Atlas. AGIS therefore blacklisted CERN DDM endpoints (which wasn't strictly necessary). The reason for declaring Atlas affected endpoints is understood (for WLCG availability), but this should be addressed so as not to trigger Atlas blacklisting of unaffected endpoints.
 - Massimo: a low-priority issue, probably followed up in autumn
 - Alessandro: meanwhile we will handle this better on the ATLAS side
 - ◇ GGUS:71715 [↗](#) CERN FTS failures blocking the T0 export starting around 14:10 UTC, ticket sent around 5pm CERN time (upgraded to alarm in the evening for accounting reasons). CERN IT responded quickly -- was found to be related to an FTS db cleanup procedure that has been ongoing for 2 weeks. Cleanup was stopped in the evening, around 9:30pm (CERN) but transfer failures happened again occasionally overnight. Tuesday morning things look fine.
 - Alessandro/Dan: SNOW updates to/from GGUS created a new alarm?
 - Maria D: no new alarm was created, just the one when the ticket was upgraded from team to alarm
 - Maria D actually had explained this in an update of the ticket, but she put the comment into the *internal* diary, which is not visible for everyone!
 - Maarten: 2nd level support only processed and routed the ticket this morning
 - Alessandro: the ticket history has some strange entries, which does not look optimal
 - all: to be sorted out offline
 - ◇ ATLAS experiences performance degradation of WLCG DB, which serves Dashboard applications (DDM dashboard, Historical views). CERN IT DBAs are investigating. DDM Dashboard had performance problems after migration of the job monitoring Dashboard to the production server. On the 8th of June the DDM Dashboard was moved to a different server from the one used by job monitoring Dashboard and since then the performance improved.
 - Julia: on May 24 the application was moved from the integration to the production DB and performance problems started; on June 8 it was moved to a less loaded node in the 4-node cluster, but the job monitoring still is not OK
 - Jacek: the cluster is used by many applications, we will need to relocate a few more
- CMS reports -
 - ◆ LHC / CMS detector
 - ◇ data taking resumed with fresh 1092 bunch fill around midnight
 - ◆ CERN / central services
 - ◇ Critical services map fixed by DashBoard team yesterday, relevant person wasn't listed in the Savannah dashboard team. System was recovered reasonably quickly after that.
 - ◇ LSF:
 - LSF not releasing jobs if node is rebooted (for example if the node was not reachable by ping/ssh), need to bkill -f manually. Unclear how to move forward to work on solution.
 - Oliver: the problem is caused by CMS jobs (huge memory usage), not by the worker nodes

- Maarten: please open a GGUS or SNOW ticket
- Gavin: we are discussing this matter with Platform
- ◇ CERN srm problems
 - according to ATLAS ticket (GGUS:71715[↗](#)), seems to be related to FTS DB overload by the background cleaning jobs while Atlas started to use gsiftp and increased the load
 - CMS saw degradation of service and opened ticket (GGUS:71718[↗](#), standard ticket at first), but we still see errors
 - Massimo: we have some timeouts on our side, not correlated with the activity level; no overload was seen; the matter is under investigation
 - ◇ public queues currently run jobs with very low CPU efficiency, we contacted the user to find out more
- ◆ Tier-0 / CAF
 - ◇ cmscaf1nw queue has > 1k pending jobs (we restricted the queue to about 10 running jobs per user because the CAF is meant for fast turnaround workflows), observing situation but not critical
- ALICE reports[↗](#) -
 - ◆ T0 site
 - ◇ Noting to report
 - ◆ T1 sites
 - ◇ IN2P3: we upgraded AliEn version at the site because the one installed had a bug that did not let us complete the torrent testing. Trying torrent this afternoon again
 - ◆ T2 sites
 - ◇ Nothing to report
- LHCb reports -
 - ◆ Experiment activities:
 - ◆ New GGUS (or RT) tickets:
 - ◇ T0: 0
 - ◇ T1: 0
 - ◇ T2: 1 (shared area problem at tbit01.nipne.ro)
 - ◆ Issues at the sites and services
 - ◇ T1
 - GRIDKA: GGUS:71572[↗](#). Continuing data access problems at GridKa. LHCb jobs are still failing due to both - problems with availability of data and access to the available data.
 - Xavier: there is a lot of stress on the system, the current setup was not foreseen for such usage, leading to a concentration of high load, several hundreds of concurrent transfers on 2 pools; we will have a meeting about this matter next week
 - Raja: what changed around June 13?
 - Xavier: nothing changed at KIT, but LHCb started submitting more work?

Sites / Services round table:

- ASGC - ntr
- BNL - ntr
- CNAF - ntr
- FNAL - ntr
- IN2P3 - ntr
- KIT - ntr
- NDGF

- ◆ GPFS failure at 1 site this morning, OK now
- ◆ short downtime tomorrow to reboot 1 server behind the SRM
 - ◇ Alessandro: what space tokens are affected?
 - ◇ Christian: not clear, but the server designation is HPC2N_UMU_SE_026
- NLT1
 - ◆ LFC server home directories unavailable, being worked on
 - ◇ Raja: which VOs are affected?
 - ◇ Ronald: ATLAS and possibly LHCb
- OSG - ntr
- PIC - ntr

- CASTOR - nta
- dashboards - nta
- databases - nta
- grid services
 - ◆ the FTS DB cleanup procedure has been made less aggressive
 - ◆ LSF share issue reported by LHCb yesterday: the scheduling algorithm has been changed to take into account only the consumed **wall-clock** time; the algorithm used to look at the consumed CPU time instead, which caused free job slots to be given preferentially to ALICE, due to the low CPU/wall ratio ALICE jobs have been suffering lately

AOB:

Wednesday

Attendance: local(Alessandro, Dan, Julia, Luca, Lukasz, Maarten, Maria D, Massimo, Oliver);remote(Christian, Dimitri, Elizabeth, Felix, Gonzalo, Joel, John, Jon, Karen, Onno, Raja, Rolf).

Experiments round table:

- ATLAS reports -
 - ◆ T0/CERN:
 - ◇ Ongoing DB issues:
 - dashb issue (reported yesterday) -- plan is to investigate moving SWAT to another RAC.
 - A particular DDM database query required for T0 export has been occasionally very slow in the past days. Currently stabilized thanks to a new index hint, but the stability is being monitored.
 - ◆ T1s:
 - ◇ SARA: GGUS:71759 [↗](#) Reported during yesterday's daily meeting: LFC outage caused by home directory outage. Errors lasted from around 12h-14h UTC.
 - Onno: details in the ticket, apologies - an at-risk downtime should have been posted
 - Alessandro: LFC test results are not included in T1 availability calculations for ATLAS
- CMS reports -
 - ◆ LHC / CMS detector
 - ◇ record 18 hour fill ended, waiting for new physics fill (not before tomorrow morning due to collimation qualification)
 - ◆ CERN / central services
 - ◇ LSF:
 - LSF not releasing jobs if node is rebooted (for example if the node was not reachable by ping/ssh), need to bkill -f manually. Under investigation, ticket

opened as requested INC:046919

◇ GGUS:71718 closed, problems solved, thanks.

· Still seeing problems on t1transfer pool but not related to Castor itself, we are investigating why this pool holds many files to be migrated currently. Don't see this moving ahead, number of files to be migrated increasing rather than decreasing. Turns out to be lots of user files going to /castor/cern.ch/user. Opened INC:047103 just to be sure that there is no problem with the migration.

- Massimo: we have no comments yet, the matter is under investigation

- ALICE reports -

- ◆ T0 site

- ◇ Nothing to report

- ◆ T1 sites

- ◇ IN2P3: Torrent was working at the site, but jobs were doing installation at the wrong place (/tmp) due to a variable misconfiguration. For the moment ALICE jobs have been blocked at the site. Should be fixed later today.

- ◆ T2 sites

- ◇ Nothing to report

- LHCb reports -

- ◆ Experiment activities:

- ◆ New GGUS (or RT) tickets:

- ◇ T0: 0

- ◇ T1: 0

- ◇ T2: 0

- ◆ Issues at the sites and services

- ◇ T1

· GRIDKA: GGUS:71572. Continuing data access problems at GridKa. LHCb jobs are still failing due to both - problems with availability of data and access to the available data. Site remains flaky even with a much lower load than before. We are concerned if this is related to the LHCb space token migration earlier this month, with the disk servers not being re-assigned properly, as needed.

- Joel: the number of jobs is tending toward zero, but a big activity is still observed at KIT - please indicate what kinds of traffic, from where etc.
- Joel: the space token change was not implemented the way we requested - might this be related to the problem?
- Dimitri: will try to push this forward, have the right experts in the loop, the ticket updated etc. but note that tomorrow is a holiday and on Friday not all the right people may be present either
- Joel: a T1 should have 24x7 coverage; this matter is very urgent for LHCb because of the upcoming summer conferences!

· RAL : Slow staging rates.

- Raja: I am in contact with the CASTOR team at RAL, will open a ticket later

Sites / Services round table:

- ASGC

- ◆ the transfer rate to CERN is limited, being investigated

- CNAF - ntr

- FNAL
 - ◆ probably unable to join Thu and Fri because of concurrent CMS T1 meetings
- IN2P3 - ntr
- KIT
 - ◆ tomorrow is a holiday and on Friday it is not clear if anyone can connect to the meeting
- NDGF
 - ◆ yesterday afternoon various routers in Sweden failed, all were back in 5 minutes except the one connecting KTH, still not back at the time of this meeting; 10% of ATLAS jobs are failing due to missing data served by KTH
- NLT1
 - ◆ LFC for LHCb was also affected by yesterday's home directory problem, but LHCb have an automatic failover mechanism
 - ◆ also the FTA was affected, which may have caused some FTS transfers to fail
- OSG - ntr
- PIC - ntr
- RAL - ntr

- CASTOR
 - ◆ T3 intervention ongoing OK
- dashboards
 - ◆ ATLAS job monitoring performance: SWAT application will be migrated
 - ◆ also the cache was cleaned, which boosted the current performance, but we will have to see if it stays like that
 - ◆ this morning the dashboard home page was down due to a disk failure, OK now
 - ◆ tomorrow there will be a 30 minute downtime for the ATLAS DDM v2 dashboard
- databases - nta
- GGUS/SNOW
 - ◆ the issues with yesterday's ATLAS ticket (GGUS:71715) are being investigated with the developers

AOB:

Thursday

Attendance: local(Alessandro, Dan, Edoardo, Eva, Gavin, Lukasz, Maarten, Maria D, Massimo, Raja);remote(Christian, Jeremy, John, Karen, Kyle, Michael, Rolf, Ronald, Todd).

Experiments round table:

- ATLAS reports -
 - ◆ LHC:
 - ◇ 16:24 Power cut (Lightning strike). No data since the record run.
 - ◆ T0/CERN:
 - ◇ dashb issue -- moving SWAT didn't help performance. Today we also noticed a possibly related slowdown in DDM Vobox callbacks to the dashboard.
 - Lukasz/Eva: collectors were set up on other machines to compare the performance, but they use the integration DB instead of the production DB, which cannot really be compared
 - Eva: the matter is still being investigated
 - ◆ T1s:
 - ◇ NDGF powercut (3h45-8h UTC). Largely transparent due to no t0 exports.
- CMS reports -

- ◆ Oli has to be at a workshop and has to give a talk right at the time of the meeting, apologies for not being at the meeting in person.
- ◆ LHC / CMS detector
 - ◇ cryo problem yesterday after power cut should be recovered by this evening, not clear if machine will run the collimation qualification or go to collisions with 1236 bunches directly
- ◆ CERN / central services
 - ◇ T1TRANSFER castor pool: migration issue was solved, migration restarted, ticket INC:047103 [↗](#) closed, thanks
 - ◇ Question to Castor: we frequently see drops in free space in our pools, for example T0EXPRESS:
https://sls.cern.ch/sls/history.php?id=CASTORCMS_T0EXPRESS&more=nv:Percentage+Free
and T0EXPORT:
https://sls.cern.ch/sls/history.php?id=CASTORCMS_T0EXPORT&more=nv:Percentage+Free
. I think these are disk servers dropping out of the pool and then being restored.
Correct? Nothing to worry about?
 - Massimo: yes, short interventions by sysadmins
 - Raja: examples? I see the same at RAL
 - Massimo: firmware upgrade followed by a reboot, disk replacement sometimes followed by a reboot
- ALICE reports [↗](#) -
 - ◆ T0 site
 - ◇ Nothing to report
 - ◆ T1 sites
 - ◇ IN2P3: AliEn configuration for Torrent not fully done yet.
 - ◆ T2 sites
 - ◇ Nothing to report
- LHCb reports -
 - ◆ Experiment activities: Many fixes and updates to DIRAC over last 24 hours. Bug which affected users hotfixed. Overnight problem with DIRAC component doing staging of files at Tier-1s fixed.
 - ◆ New GGUS (or RT) tickets:
 - ◇ T0: 0
 - ◇ T1: 0
 - ◇ T2: 0
 - ◆ Issues at the sites and services
 - ◇ T1
 - GRIDKA: GGUS:71572 [↗](#). Situation slightly better in the last 24 hours. Very low (~200) numbers of running jobs. Effort on follow up with GridKa about the space token migration that was done a few weeks ago - worry that it was not completed.
 - RAL : Slow staging rates - need to wait and see the effects of the changes made yesterday. Significant fractions of jobs still failing with "input data resolution" indicating continuing problems with garbage collection.

Sites / Services round table:

- ASGC - ntr
- BNL - ntr
- CNAF - ntr
- FNAL - ntr
- GridPP - ntr

- IN2P3 - ntr
- NDGF
 - ◆ KTH site OK since ~midnight, then power cut (see ATLAS report), all OK now
- NLT1 - ntr
- OSG - ntr
- RAL - ntr

- CASTOR
 - ◆ CMS backlog processed yesterday evening
 - ◆ Wed June 29 upgrade to 2.1.11 for ATLAS
 - ◆ Thu June 30 upgrade to 2.1.11 for CMS
 - ◆ main features of 2.1.11:
 - ◇ allows new, in-house scheduler to be used instead of LSF
 - ◇ tape gateway
 - ◇ fixes for various bugs affecting tape handling
- dashboards
 - ◆ yesterday's ATLAS DDM dashboard upgrade went OK
- databases - nta
- GGUS/SNOW - ntr
- grid services
 - ◆ Whole node test CE -> lxb7835 (out of warranty) will be retired on 30 June, in order to make space for the new batch nodes. Its replacement creamtest001 is now available.
- networks - ntr

AOB:

Friday

Attendance: local (AndreaV, Dan, Alessandr, Maarten, Lukasz, Raja, Eva, Nilo, Massimo); remote (Michael, Xavier, Alexander, Jhenwei, Maria Francesca, Rolf, Gareth, Kyle).

Experiments round table:

- ATLAS reports -
 - ◆ T0/CERN:
 - ◇ Jobs failing and FTS transfers failing due to a failed disk server (GGUS:71866 [↗](#) GGUS:71869 [↗](#)). 90% of files are already restored. [Massimo: vendor call was opened to recover the other files.]
 - ◇ SRM_FAILURE Unable to issue PrepareToGet request to Castor. (GGUS:71895 [↗](#)) [Massimo: following up, with less priority than the previous issue.]
 - ◇ Alarm (GGUS:71904 [↗](#)) since about 12:30h all attempts of reading from and writing to CASTOR pools T0MERGE and T0ATLAS fail, either because of client-side timeouts (after 15min of no response) or "Device or resource busy". [Massimo: this should be solved now.]
 - ◇ [Dan: also following up on one user who has been doing very heavy writes - the scratch disk was blacklisted because it became full, and this also triggered some file deletions. Massimo: noticed this, but the alarm chain from GGUS did not work completely as expected. Maarten: please contact MariaD if there is a potential GGUS issue.]
 - ◆ T1s:
 - ◇ BNL had a short dcache outage ~8pm CERN time and was auto-excluded for 1 hour for analysis. [Michael: 15 minute intervention was triggered by an issue discovered in the dcache nameserver. There are two analysis queues, only the short queue was excluded from job brokerage, while the long queue remained available. With 1000's

of jobs in the short queue all job slots remained occupied during the time the queue was excluded from brokerage]

- CMS reports -
 - ◆ Oli has to be again at a workshop and has to give a talk right at the time of the meeting, apologies for not being at the meeting in person.
 - ◆ LHC / CMS detector
 - ◇ machine injected 1236 bunches in the morning, hope that collisions start soon
 - ◆ CERN / central services
 - ◇ CASTORCMS went down 9:40 AM CERN time:
<http://it-support-servicestatus.web.cern.ch/it-support-servicestatus/IncidentArchive/110624-Ca>
 , fixed at 10 AM, thanks [Massimo: one head node went 'funny' and was fixed by reboot - exact causes are being investigated]
 - stopped job submission on T0 for 20 minutes, not a big problem
 - confusing: main IT SSB list problem fixed at 10 AM but detailed page does not:
<http://it-support-servicestatus.web.cern.ch/it-support-servicestatus/IncidentArchive/11>
 [Massimo: following up with SSB team]
- ALICE reports [↗](#) -
 - ◆ General Information: Low Job efficiency is being investigated, Maarten got root acces to some worker nodes to do the debugging of the jobs. Suspicion that it is not just ROOT processes hanging, but multiple issues.
 - ◆ T0 site
 - ◇ Nothing to report
 - ◆ T1 sites
 - ◇ IN2P3: Torrent was working, but there still is an unsolved configuration problem causing the software tar balls to be put in /tmp and filling up that file system, while they get unpacked in the right place. We backed out of using Torrent for now and intend to continue debugging this next week.
 - ◆ T2 sites
 - ◇ Nothing to report
- LHCb reports -
 - ◆ Experiment activities:
 - ◇ Further problems with pre-staging of files within DIRAC. Being handled by hand at present while debugging of the problem by the experts continues.
 - ◆ New GGUS (or RT) tickets:
 - ◇ T0: 0
 - ◇ T1: 0
 - ◇ T2: 0
 - ◆ Issues at the sites and services:
 - ◇ T1
 - GRIDKA: GGUS:71572 [↗](#). Job success rate fine now. However we need to know whether we can ramp up the level of jobs running there. Currently we have ~200 jobs running with >6100 waiting jobs for GridKa. This is now critical for us and needs to be addressed soon. About 100TB of free space has also been transfered into LHCb-DISK from LHCb-DST. However the space token migration that was done a few weeks ago has not really been completed. [Xavier: following up, but not ready yet to resume full load during the weekend, still running at reduced maximum load]
 - RAL : Believe problems with job failures due to "input data resolution" have been understood now - large latency [Raja: a few days] between pre-staging and actual running of the jobs.

Sites / Services round table:

- FNAL [by email]: we moved back to our read-only NFS code servers for CMSSW distributions for worker and interactive nodes.
 - BNL: nta
 - KIT: nta
 - NLT1: short downtime this morning to replace one server CPU
 - ASGC: yesterday lost network connection of DPM disk, now back to normal
 - NDGF: next Monday 8 to 9 UTC reboot of some dcache nodes, some data may be unavailable
 - IN2P3: ntr
 - RAL: nta
 - OSG: ntr
-
- Storage services: next week two upgrades to 2.1.11, on Wed for ATLAS and Thu for CMS. Tried to register this in GOCDDB but failed (opened [GGUS:71902](#) to report the GOCDDB problem).
 - Dashboard services. Raja: LHCb dashboard is pretty unstable. Maarten: similar problem seen for ALICE, may be due to use of an older version of Dashboard (newer version currently available only for ATLAS and CMS). Lukasz: will follow up.
 - DB services: multiple disk failures for PDBR last night, moved to standby, due to hardware problems. Next week will move PDBR physically to another CC rack, three standby's will be available.

AOB: none

-- JamieShiers - 17-Jun-2011

This topic: LCG > WLCGDailyMeetingsWeek110620

Topic revision: r18 - 2011-06-25 - MaartenLitmaath



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](#)