

Table of Contents

Week of 120416.....	1
Daily WLCG Operations Call details.....	1
WLCG Service Incidents, Interventions and Availability, Change / Risk Assessments.....	1
General Information.....	1
Monday.....	1
Tuesday.....	3
Wednesday.....	5
Thursday.....	7
Friday.....	9

Week of 120416

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 (Andrea, Elisa, Eva, Luca, Jarka, Michael, Torre, Luc, Steve, Maarten, MariaDZ); remote (Michael/BNL, Jhen-Wei/ASGC, Alexander/NLT1, Lisa/FNAL, Tiju/RAL, Thomas/NDGF, Rolf/IN2P3, Dimitri/KIT, Giovanni/CNAF; Ian/CMS).

Experiments round table:

- ATLAS reports -
 - ◆ T0/Central Services
 - ◇ ATLAS Central Catalog: Load balancing choosing always 1 machine. [GGUS:81308](#)
 - ◆ T1s/CalibrationT2s
 - ◇ NDGF-T1 has more than 90k deletion errors in last 4 hrs (same error as last weekend) reported Fri evening. Persisted through the weekend. [GGUS:81257](#)
 - ◇ FZK-LCG2 failing transfers because of "no space left" on /var, reported Sat morning. Recurrence of last week ([GGUS:81163](#))? How to fix it for good? [GGUS:81277](#)
 - ◇ TRIUMF-US networking problems reported Friday were resolved Sat afternoon. Network problem on the dedicated circuit between TRIUMF and BNL. [GGUS:81111](#) [GGUS:81213](#)
 - ◇ BNL-ATLAS transfer errors due to expired certificate, reported Sat afternoon. Resolved in 90min. [GGUS:81281](#)
 - ◇ IN2P3-CC errors on Tier 0 export transfers. Site declared unscheduled downtime

until 8am, dCache overloaded. Escalated to ALARM. Excluded from Tier 0 exports. Problem resolved 3hr 20min after first report, Tier 0 exports restored. Closed this morning. GGUS:81286

◇ IN2P3-CC comment last Thursday that they are flooded with empty ATLAS pilots: we request more detailed information on the problem. Checking on our side, pilot rates sent to IN2P3-CC both from FR and from CERN are low. [Rolf: will follow up.]

- CMS reports -
 - ◆ LHC machine / CMS detector
 - ◇ CMS preparing for VdM (Van der Meer) scans today. Trigger rate will be very high for this period.
 - ◆ CERN / central services and T0
 - ◇ Overnight Saturday to Sunday there was a problem with LSF. About 4 hours and then recovered. [Steve: the admins did get an SMS. Will be followed up.]
 - ◆ Tier-1/2:
 - ◇ Job robot failures observed at T1_TW_ASGC
- ALICE reports -
 - ◆ NTR
- LHCb reports -
 - ◆ Prompt reconstruction and stripping going on at Tier1s and T0
 - ◆ MC simulation at Tiers2
 - ◆ T0
 - ◇ some production jobs failed due to no space left on disk. Opened a ticket, already solved: 20GB scratch disk space guaranteed for LHCb jobs (according to Vo card)
 - ◆ T1
 - ◇ IN2P3: site banned during 3 hours yesterday for SRM unscheduled down-time [Elisa: same issue as reported by ATLAS]
 - ◇ General for all Tiers1: new version of CernVMFS client available with a fix for the problem with the cache (see ticket GGUS:81181), should be deployed asap [Steve: this has not yet been released. Elisa: ok thanks, then will ask for this to be deployed as soon as it is released]
 - ◇ CNAF: ticket (GGUS:81291) for the update WMS version. Solved already.
 - ◇ PIC: ticket (GGUS:81290) for the update WMS version
 - ◇ SARA: ticket (GGUS:81289)
 - ◆ T2
 - ◇ ntr

Sites / Services round table:

- Gonzalo/PIC: Storage managers at PIC tell me that in the last weeks we are seeing a big increase in gridftp-v1 transfers, as compared to gridftp-v2. The former has the "not nice" feature that traffic is routed through the doors, and not directly to the disk pool nodes. Apparently this is a side effect of the GLOBUS_FTP_CLIENT_GRIDFTP2 env variable which existed in VDT to configure the gridftp version, but does not exist anymore in the EPEL globus. This is being tracked at GGUS:81230 . Since this might have some performance implications, I wonder if we should be trying to apply some workaround to get back to using gridftp-v2 wherever possible. [Maarten: first the IGE developers should come up with a response to the ticket, but they have not replied yet.]
- Michael/BNL: ntr
- Jhen-Wei/ASGC: looking at CMS robot hang
- Alexander/NLT1: ntr
- Lisa/FNAL: ntr
- Tiju/RAL: ntr

- Thomas/NDGF: scheduled downtime tomorrow afternoon, will affect data availability for ATLAS and ALICE (but will still be able to write)
- Rolf/IN2P3: nta
- Dimitri/KIT: ntr
- Giovanni/CNAF: is there any update on EMI worker node tests? [Maarten: some news but still need more progress. CMS did manage to fix the main issue seen one month ago, but it is not yet clear if there are other issues. Not much progress in ATLAS. Will try to follow up at next T1SCM.]
- Jarka/Dashboard: **WLCG transfer monitor**: Since Friday 13th April 2012 we experience problem with reporting from the following FTS instances: `fts.triumf.ca`, `lcgfts.gridpp.rl.ac.uk`, `cclcgftsprod.in2p3.fr`, `fts.grid.sara.nl`
 - ◆ PIC was also in the list, but it was fixed before weekend.
 - ◆ In order to fix the problem `msg-bulk` daemon should be restarted: `service glite-msg-bulk restart`
 - ◆ Michail Salichos contacted FTS responsables asking to do it on Friday morning, but the problem still persists.
 - ◆ [Michael: messaging daemon got hit by an ActiveMQ bug, a workaround will be attempted]
 - ◆ [Rolf: on FTS problem, IN2P3 restarted the service but did not notice any attempt to contact IN2P3 on Friday. Jarka: will follow up.]
- Steve/Grid: ntr
- Luca/Storage: ntr
- Eva/Databases: one node of the LHCb online cluster rebooted spontaneously today, being investigated

AOB:

- MariaDZ: Up-to-date file `ggus-tickets.xls` is attached to page `WLCGOperationsMeetings`. There was 1 ALARM last week (ATLAS to IN2P3-CC yesterday, Sunday). There are 7 real ALARMS subject to drills for the MB next week.

Tuesday

Attendance: local (Andrea, Luc, Oliver, Alessandro, Jarka, Maarten, Zbiszek, Elisa, Luca); remote (Lisa/FNAL, Kyle/OSG, Pavel/KIT, Giovanni/CNAF, Ronald/NLT1, Thomas/NDGF, Jhen-Wei/ASGC, Rolf/IN2P3, Gonzalo/PIC, Tiju/RAL).

Experiments round table:

- ATLAS reports -
 - ◆ T0/Central Services
 - ◇ ATLAS Central Catalog: Load balancing choosing always 1 machine. GGUS:81308 [↗](#). Machine not selected did get requests. Ongoing
 - ◇ [Luc: just got an alarm for ATLAS at Castor T0]
 - ◆ T1s/CalibrationT2s
 - ◇ NDGF-T1 in DT 11:00-17:00 OPN link to one of subsites will be down.
 - ◇ LRZ-LMU (Calib T2) had transfer failure. Downtime has been extended. GGUS:81324 [↗](#)
- CMS reports -
 - ◆ LHC machine / CMS detector
 - ◇ Waiting for 1300 bunch beam injection
 - ◆ CERN / central services and T0
 - ◇ GGUS:81312 [↗](#) => SRM tests for CMS not running due to wrong PYTHONPATH on `sam-cms-prod.cern.ch`

- Andrea patched and restarted NAGIOS, still need to figure out what happened [Oliver: no reply from Nagios experts on this ticket yet]
- ◇ GGUS:81199 => GEANT Networking Problems to Baltics and Norway
 - on hold, still waiting for update from GEANT? [Maarten: this is assigned to network people and officially resolved last Friday. You may post a comment suggesting to close the ticket.]
- ◆ Tier-1/2:
 - ◇ T1_TW_AS GC: JobRobot problems continue, error is: "Globus error 25: the job manager detected an invalid script status", response yesterday indicated problematic WNs [Maarten: to fix this error it is probably enough to restart the gatekeeper. Jhen-Wei: thanks, will follow up.]
- ALICE reports -
 - ◆ CNAF: job submission stopped yesterday evening due to issue with resource BDII on the CEs: each reported 444444 jobs waiting. Looks fixed since ~noon.
- LHCb reports -
 - ◆ Prompt data reconstruction, data stripping and users analysis going on at Tier1s and T0
 - ◆ Almost no activity at Tiers2 (very few MC simulation jobs)
 - ◆ T0
 - ◇ ntr
 - ◆ T1
 - ◇ GRIDKA: 70 files of the last 2011 re-stripping are missing from GRIDKA SE, investigations ongoing. Opened a ticket (GGUS:81322).
 - ◆ T2
 - ◇ ntr

Sites / Services round table:

- Lisa/FNAL: ntr
- Kyle/OSG: ntr
- Pavel/KIT: announce downtime April 25 for tape service from 6 to 10 UTC, writing and reading will not be available
- Giovanni/CNAF: upgrade of tapes tomorrow from 5.30 to 13 UTC
- Ronald/NLT1: ntr
- Thomas/NDGF: observed some short (1 minute) network interruptions earlier today; network provider is escalating this with the vendor
- Jhen-Wei/ASGC: LCG CEs will be removed
- Rolf/IN2P3: prepared some statistics for short jobs as suggested by ATLAS
 - ◆ Since Friday 4pm: ATLAS submitted 46k jobs, with 43% short (less than 200 seconds); LHCb is even worse, 59k jobs with 90% of them short; ALICE also has a high fraction of short jobs, but has much fewer jobs altogether. It seems that many ATLAS and LHCb jobs are doing nothing.
 - ◆ [Alessandro: number of jobs seems ok, taking into account that IN2P3 has around 5k slots for ATLAS. Understand the concern about short jobs, but expect that at CERN the fraction may also be around 40%, this is because ATLAS sends more jobs than slots available. Would therefore separate two issues: first, is there anything that changed recently; second, we should in any case do some optimizations, but this is not urgent if no significant change has been seen, and we should quote wall clock time efficiency rather than number of job efficiency. Please report to the T1 contact for ATLAS so that we can follow up offline.]
 - ◆ [Rolf: reminder, raised this last week because a peak was observed of short job fraction, around 60%. This was higher than in the past.]
- Gonzalo/PIC: uploaded a couple of SIRs for file loss, already discussed some time ago
- Tiju/RAL: ntr

- Luca/Storage: will follow up the ATLAS alarm for Castor T0
- Zbiszek/Databases: observed a problem last week with Streaming with last Oracle version, whenever there is a network glitch. Must restart the apply process on the target side, which is inconvenient. Working on the monitoring to make this more automatic.
 - ◆ This was observed again this morning around 7am. [Maarten: was there a network glitch? Elisa: SSB says there was a problem with the firewall, not yet resolved.]
 - ◆ [Andrea: is this a bug in Oracle? Zbiszek: they probably consider this a feature (judging from the reports from other users), but will follow it up in a Service Request.]
- Jarka/Dashboard: about the issue discussed yesterday, we had not contacted Lyon last week; apologies for this, we'll make sure the communication is improved.

AOB: none

Wednesday

Attendance: local(Elisa, Luc, Luca C, Luca M, Maarten, Mike, Steve);remote(Burt, Giovanni, Gonzalo, Jhen-Wei, Oliver, Pavel, Rob, Rolf, Ronald, Thomas, Tiju).

Experiments round table:

- ATLAS reports -
 - ◆ T0/Central Services
 - ◇ ALARM ticket GGUS:81352 [↗](#): Problem in retrieving RAW data from CASTOR/T0ATLAS. Fixed very rapidly.
 - Luca M: was an alarm ticket justified for that disk server?
 - Maarten/Luc: such disk servers are critical for the ATLAS data taking workflow, viz. for merging files before they go to tape
 - Luca M: what if all its data were lost?
 - Maarten/Luc: copies are kept online until the data is known to be on tape; however, an unavailable disk server with unprocessed data leads to a backlog that may become quite annoying, so the fate of that machine should be decided ASAP
 - ◇ ATLAS Central Catalog: Load balancing choosing always 1 machine. GGUS:81308 [↗](#). Idle for the moment.
 - ◆ T1s/CalibrationT2s
 - ◇ FR-Tier1 pilots issue. Experts working. Investigating communication with panda server & Interplay between long & very long queues.
- CMS reports -
 - ◆ LHC machine / CMS detector
 - ◇ still waiting for 1300 bunch beam injection
 - ◆ CERN / central services and T0
 - ◇ new:
 - ◇ old:
 - GGUS:81312 [↗](#) => SRM tests for CMS not running due to wrong PYTHONPATH on sam-cms-prod.cern.ch
 - Andrea patched and restarted NAGIOS, still need to figure out what happened
 - in progress since yesterday, 4/17, but no comment yet
 - GGUS:81199 [↗](#) => GEANT Networking Problems to Baltics and Norway
 - was set to "in progress" on 4/17, waiting for final confirmation that problem is fixed
 - ◆ Tier-1/2:
 - ◇ new:

- KIT: GGUS:81355 => pilots fail immediately with CREAM delegation errors
 - Pavel: CREAM services were restarted, but no problem was seen in the logs, please check

◇ old:

- ALICE reports -
 - ◆ NTR
- LHCb reports -
 - ◆ Prompt data reconstruction, data stripping and users analysis going on at Tier1s and T0
 - ◆ MC simulation at Tiers2
 - ◆ **New GGUS (or RT) tickets**
 - ◆ Central services:
 - ◇ waiting for the release which fixes the problem of CernVm-fs stale cache (GGUS:81181)
 - Steve: fix released today, see central services report below
 - ◆ T0
 - ◇ waiting for an update about the lost files due to a broken Castor disk server (GGUS:80973)
 - Luca M: we will update the ticket
 - ◆ T1
 - ◇ GRIDKA: 70 files of the last 2011 re-stripping missing from GRIDKA SE (GGUS:81322)
 - ◆ T2

Sites / Services round table:

- ASGC
 - ◆ CMS job robot tests should be OK now
- CNAF
 - ◆ tape system upgrade finished OK
 - ◆ StoRM back-end for ATLAS crashed and was rebooted; being investigated
- FNAL - ntr
- IN2P3
 - ◆ pilot job situation got much better this morning, but not yet OK
- KIT - nta
- NDGF - ntr
- NLT1 - ntr
- OSG
 - ◆ investigating MyOSG client query spikes at the top of each hour; looks like a cron job at some European sites, e.g. CERN and IN2P3; will open tickets to get the load to be spread better
 - ◇ Maarten: each top-level BDII needs to query MyOSG (and the GOCDB) regularly; that used to be much more frequently than once per hour; it probably depends on the exact version
 - ◇ after the meeting: the EMI top-level BDII indeed has an hourly cron job for that; the developers will need to insert a random delay
- PIC - ntr
- RAL - ntr
- dashboards - ntr
- databases - ntr
- grid services
 - ◆ **CVMFS**

- ◇ A bugfix release 2.0.13 is published. Among some minor fixes, it fixes a bug that can lead to stale negative cache entries in the CernVM-FS memory cache as LHCb describes above.
- ◇ The release has been running on the CERN pre-production nodes since last Friday. Please update with caution anyway.
- ◇ CERN will deploy to all production nodes tomorrow, Thursday morning.
- ◇ See updated release notes at <http://cernvm.cern.ch/portal/cvmfs/release-2.0>

- storage - ntr

AOB:

Thursday

Attendance: local (Andrea, Luc, Mike, Eva, Maarten, Steve); remote (Gonzalo/PIC, Jhen-Wei/ASGC, Thomas/NDGF, Lisa/FNAL, Gareth/RAL, Ronald/NLT1, Giovanni/CNAF, Marc/IN2P3, Rob/OSG; Elisa/LHCb, Oliver/CMS).

Experiments round table:

- ATLAS reports -
 - ◆ T0/Central Services
 - ◇ ATLAS new T0 expert submitted (perhaps w/o enough info to debug) an ALARM, we then investigated the situation that indeed was alarming, one out of 4 the LFC frontend was down. LFC ops are working on it.
 - ◆ T1s/CalibrationT2s
 - ◇ FR-Tier1 pilots issue. Various queues pointing to the same panda site had non-consistent parameters.
- CMS reports -
 - ◆ LHC machine / CMS detector
 - ◇ still waiting for 1300 bunch beam injection
 - ◇ Machine Development postponed to start Saturday
 - ◆ CERN / central services and T0
 - ◇ new:
 - VOMScore seems to show more glitches recently <https://sls.cern.ch/sls/history.php?id=VomsCore&more=availability&period=24h>, anything wrong? [Steve: glitches and restarts are a known issue, it is a problem only if the frequency is increasing. Will have a look anyway and will soon deploy a fix from EMI. Oliver: Ok will tell the shifters to only submit tickets if VOMS is down for a long period. Maarten: note that VOMS will normally work even if the system is degraded.]
 - ◇ old:
 - GGUS:81312 => SRM tests for CMS not running due to wrong PYTHONPATH on sam-cms-prod.cern.ch
 - Andrea patched and restarted NAGIOS, still need to figure out what happened
 - in progress since yesterday, 4/17, but no comment yet! [Maarten: will ping the Nagios team to update the ticket.]
 - GGUS:81199 => GEANT Networking Problems to Baltics and Norway
 - was set to "in progress" on 4/17, waiting for final confirmation that problem is fixed
 - ◆ Tier-1/2:
 - ◇ new:

- GGUS:81424 => JobRobot problems at T1_TW_ASGC, seems to be related to old lcg-CEs, which will be gone soon [Jhen-Wei: LCG CEs were removed this morning, this should fix the issue.]

◇ old:

- ALICE reports -
 - ◆ Low activity led to pilots running without doing any work, as reported by IN2P3; ramping up again now.
 - ◆ [Andrea: was this a new problem only for ATLAS then? Indeed Rolf had reported that also LHCb and ALICE had abnormally high numbers of short pilots. Marc: not really new, for ALICE this is observed since some time. Maarten: not really clear what changed; what Luc reported about panda in ATLAS is something that has just been understood, but that configuration has not been changed. Elisa: also talked about this with the LHCb experts today, last week there was a change in the job submission rate but it is not clear if this is related; will follow up with the LHCb contact person at Lyon.]
- LHCb reports -
 - ◆ ◇ Prompt data reconstruction, data stripping and users analysis going on at Tier1s and T0. Yesterday stopped productions for reconstruction and stripping with the previous application version, which were taking very long time, causing jobs to hit the end of queues at sites and being rescheduled.
 - ◇ MC simulation at Tiers2
 - ◇ T0
 - waiting for an update about the lost files due to a broken Castor disk server (GGUS:80973) (last update was on Monday)
 - ◇ T1
 - GRIDKA
 - 70 files of the last 2011 re-stripping missing from GRIDKA SE (GGUS:81322).
 - job submission to Gridka WMS are failing (GGUS:81405). Ongoing.
 - some FTS transfers failing, fixed now (GGUS:81398). Promptly fixed yesterday.
 - CNAF upgrade WMS version ongoing (GGUS:81291)

Sites / Services round table:

- Gonzalo/PIC: ntr
- Jhen-Wei/ASGC: another CMS server had a partition error due to a hardware failure (vendor call ongoing) and some files need to be recovered.
 - ◆ [Oliver: a bit worried, files had already been lost last week and now more files are being lost, is the system stable enough? Jhen-Wei: will follow up offline and report more tomorrow.]
- Thomas/NDGF: ntr
- Lisa/FNAL: looking for advice about FTS, had three overload incidents since the upgrade to the latest version, due to the large number of log files (many files per transaction, all in one directory). Can we fix this for instance by reducing verbosity?
 - ◆ [Steve: probably can fix this by reducing from debug to info, but you should ask the experts. Maarten: please open a ticket or contact the support list (fts-support at cern.ch).]
- Gareth/RAL: added back two nodes to the FTS frontend (had already added them weeks ago but then had temporarily removed them)
- Ronald/NLT1: ntr
- Giovanni/CNAF: ntr
- Marc/IN2P3: just checked with the operations team, the ALICE pilot jobs are ok again
- Rob/OSG: problem reported yesterday is being fixed and all seems to be going ok, thanks to Maarten

again; will keep an eye on the evolution

- Mike/Dashboard: ntr
- Eva/Databases: the April security patch for Oracle has been released. The validation will start in upcoming weeks. We will not be able to apply these patches during the technical stop unfortunately, so we'll do it as soon as possible.]
- Steve/Grid: ntr

AOB: (MariaDZ) The interface developers of local ticketing systems were informed about some GGUS fields being withdrawn (Savannah:127148) and others becoming mandatory (Savannah:127146).

Countries/organisations concerned are DE, ES, FR, IT, CERN and OSG. This was announced already in the AOB of WLCGDailyMeetingsWeek120402#Wednesday The changes will take effect at the next GGUS Release on 2012/04/25. Announcement in GOCDB done. See news on https://ggus.eu/pages/news_detail.php?ID=458

Friday

Attendance: local (Andrea, Luc, Steve, Mike, Maarten, Massimo, Oliver, Jarka, Jamie); remote (Mette/NDGF, Lisa/FNAL, Gonzalo/PIC, John/RAL, Jhne-Wei/ASGC, Xavier/KIT, Onno/NLT1, Rob/OSG; Elisa/LHCb).

Experiments round table:

- ATLAS reports -
 - ◆ T0/Central Services
 - ◇ LFC back in business
 - ◇ LSF down at 8pm. Alarm ticket GGUS:81445. Still ongoing.
 - ◆ T1s/CalibrationT2s
 - ◇ ntr
- CMS reports -
 - ◆ LHC machine / CMS detector
 - ◇ still waiting for 1300 bunch beam injection
 - ◇ machine development starts Saturday, April 21st, 6 AM CERN time
 - ◆ CERN / central services and T0
 - ◇ new:
 - LSF problems over night, didn't open ticket because Atlas opened ALARM ticket GGUS:81445, any update?
 - ◇ old:
 - GGUS:81312 => SRM tests for CMS not running due to wrong PYTHONPATH on sam-cms-prod.cern.ch CLOSED, solution was proposed
 - GGUS:81199 => GEANT Networking Problems to Baltics and Norway
 - last update 4/17, waiting for final confirmation that problem is fixed
 - ◆ Tier-1/2:
 - ◇ new:
 - GGUS:81453 => CNAF WMS wms002.cnaf.infn.it does not match anything, seems to have recovered mostly but still seeing problems from time to time, would appreciate feedback
 - ◇ old:
 - T1_TW_AS GC disk server loss, facility is in contact with vendor and also works with CMS experts to identify lost files, update from Jhen-Wei, all mc and data files can be re-staged from tapes without problem. Temporary files will be deleted if too old, otherwise reported.
 - issue not related to Castor DB loss last week

- ALICE reports [↗](#) -
 - ◆ NTR

- LHCb reports -
 - ◆ Prompt data reconstruction, data stripping and users analysis going on at Tier1s and T0.
 - ◆ MC simulation at Tiers2
 - ◆ T0
 - ◇ 2500 LHCb jobs SIGSTOPed during last night. Not sure, but they were suspected of being responsible for killing the batch system with too many queries. Though on DIRAC side nothing has been changed with respect to queries to the batch system in the last year. Jobs re-started this morning.
 - ◇ waiting for an update about the lost files due to a broken Castor disk server (GGUS:80973 [↗](#)) (last update was on Monday). [Massimo: will update the ticket and provide details about the lost files.]
 - ◆ T1
 - ◇ GRIDKA: job submission to Gridka WMS are failing (GGUS:81405 [↗](#)). Ongoing.
 - [Xavier/KIT: missing files seem to have all been removed by the LHCb representative. Elisa: thanks for the info; this was not intentional, so there must be a bug to fix in the LHCb infrastructure.]
 - ◇ SARA:
 - ticket (GGUS:81457 [↗](#)) for pilots aborted with Reason=999. [Onno: pilot abort was caused by Torque upgrade to 2.5.11 that affected Cream CE polling.]
 - asked to upgrade to last CernVM-FS version (GGUS:81462 [↗](#)). [Onno:new cvmfs software has been installed but not yet restarted; next week will do a rolling reboot of inactive nodes.]

Sites / Services round table:

- Mette/NDGF: problem with dcache tonight (GGUS:81447 [↗](#)), all seems ok now
- Lisa/FNAL: ntr
- Gonzalo/PIC: ntr
- John/RAL: ntr
- Jhen-Wei/ASGC: nta
- Xavier/KIT: nta
- Onno/NLT1: nta
- Rob/OSG: issues reported last week is now fixed and can be closed. [Maarten: a patch will also soon be provided by BDI developers.]

- Massimo/Storage: nta
- Mike/Dashboard: ntr
- Steve/Grid: **CERN Batch** summary is, problems yesterday evening and again this morning seemed to be correlated with grid_lhcb jobs querying the system. lhcb jobs have been resumed after being paused and the problem didn't occur again. We are still investigating, along with Platform support.

AOB: none

-- JamieShiers - 22-Mar-2012

This topic: LCG > WLCGDailyMeetingsWeek120416

Topic revision: r24 - 2012-04-20 - AndreaValassi



Copyright &© 2008-2020 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