

Table of Contents

Week of 101025.....	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 101025

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(Dirk, Eddie, Edoardo, Gavin, Graeme, Harry, Ignacio, JPB, Jamie, Jan, Lola, Luca, Maarten, Maria D, Maria G, Nilo, Roberto, Simone, Stephane);remote(Gang, Gonzalo, John, Jon, Michael, NDGF, Onno, Rob, Rolf, Stefano, Xavier).

Experiments round table:

- ATLAS reports -
 - ◆ Ongoing issues
 - ◇ CNAF-BNL network problem (slow transfers) [GGUS:61440](#), [GGUS:63134](#). Failures due to transfer timeouts for the large files (>2GB, often ~4GB).
 - ◆ ATLAS
 - ◇ Data taking with stable beams since Sunday 1000.
 - ◇ Simone: testing T2-T2 inter-cloud transfers, associated failures will be handled by experts
 - ◆ T1
 - ◇ INFN-T1: Some users have problems retrieving data [GGUS:63181](#).
 - ◆ Central Services
 - ◇ Still investigating problems in accessing GOC downtime information. LHCb reported similar problems due to host certificate change on the GOC side. [GGUS:63363](#), [BUG:74295](#).

- CMS reports -
 - ◆ Experiment activity
 - ◇ Data taking, processing, analyzing, simulating
 - ◆ CERN and Tier0
 - ◇ Critical service GridMap <http://servicemap.cern.ch/ccrc08/servicemap1.html> lost its interactivity. Dashboard team redirected us to J.Casey. Item for GGUS ?
 - No clear support unit in GGUS, Maarten will follow up
 - ◆ Tier1 issues
 - ◇ Reprocessing ongoing, no issues
 - ◆ Tier2 Issues
 - ◇ large production ongoing, no issues
 - ◆ MC production
 - ◇ large production ongoing
 - ◆ AOB
- ALICE reports -
 - ◆ T0 site
 - ◇ Scheduled operations on the ALICE voalicefsXX and voalice13 nodes still ongoing
 - ◇ During the weekend some issues with CREAM have been observed, submission switched automatically to WMS. Since this morning CREAM submission is performing well
 - It deteriorated in the afternoon and 2 of the 3 CREAM CEs were observed to fail the "ops" SAM tests continuously
 - ◇ Operations on voalice11 scheduled for this afternoon: update gLite-vobox and OS updates. Those operations were already done in voalice14 and voalice12, which are expected to enter in production this afternoon
 - ◆ T1 sites
 - ◇ NIKHEF: GGUS 63400, CREAM was not working this morning. Queues were closed due to a vulnerability present in CentOS 5. Downtime was set before the queues were closed, but still did not appear in GOCDDB
 - ◆ T2 sites
 - ◇ Nothing to report
- LHCb reports -
 - ◆ Experiment activities: Normal Activities.
 - ◆ New GGUS (or RT) tickets:
 - ◇ T0: 0
 - ◇ T1: 0
 - ◇ T2: 0
 - ◆ Issues at the sites and services
 - ◇ T0
 - NTR
 - ◇ T1 site issues:
 - GRIDKA: (GGUS:63353). LHCb_DST space token running out of space. Not all 2010 pledged resources allocated.
 - Xavier, Roberto: discussion to be continued in the ticket
 - IN2P3: Shared area issue (GGUS:59880) opened July the 8th. Very urgent. (GGUS:62800) opened October the 6th. Top priority. Still to be solved. It's time for escalation.
 - Rolf: will escalate at IN2P3, please escalate in WLCG meetings as well

Sites / Services round table:

- ASGC - ntr
- BNL
 - ◆ CNAF-BNL network issue
 - ◇ on Fri the service provider of the Vienna-Amsterdam link lowered the link's priority, thereby moving the traffic to other links
 - ◇ the throughput increased and CNAF closed the ticket as solved
 - ◇ BNL reopened the ticket, since the matter is not fully understood
 - ◇ the ticket should be updated with further communications between the various service providers
- FNAL
 - ◆ Due to change in CERN technical stop - FNAL has moved its downtime from Nov 2 to Nov 8.
- IN2P3 - ntr
- KIT - ntr
- NDGF
 - ◆ transparent disk pool problem in Norway, fixed
- NLT1
 - ◆ maintenance preventing user job submission because of new kernel vulnerability, no timeline estimate
- OSG - ntr
- PIC
 - ◆ ATLAS production jobs staging errors during the weekend due to intermittent DCAP door failures, not understood; 4 doors added (total 8) to reduce potential load issues
- RAL
 - ◆ CASTOR being upgraded for ALICE + small VOs

- CASTOR - ntr
- dashboards - ntr
- databases - ntr
- grid services - ntr
- network - ntr

AOB:

Tuesday:

Attendance: local(Eddie, Gavin, Harry, Jamie, Jan, Lola, Luca, Maarten, Maria D, Maria G, Roberto, Stephane);remote(Dimitri, Federico, Gang, Gonzalo, Ian, John, Jon, Michael, NDGF, Paolo, Rob, Rolf, Ronald).

Experiments round table:

- ATLAS reports -
 - ◆ Ongoing issues
 - ◇ CNAF-BNL network problem (slow transfers) GGUS:61440 (updated Oct 25), GGUS:63134 (verified). Failures due to transfer timeouts for the large files (>2GB, often ~4GB): OK after timeout increased to 3600s.
 - ◆ ATLAS
 - ◇ Data taking : Collected 10 pb-1 in 48 hours (compared to 20 pb-1 since March)
 - ◆ T0
 - ◇ GGUS:63414: RAW file inaccessible on Castor. Looks like an already problematic disk server but monitoring did not spot the problem. Disk server was drained and put back in production later. In the meantime, RAW file was copied again from DAQ so that ATLAS activity can continue

Tuesday:

- Jan: machine taken out because of HW error, should not have been put back in production, procedure will be improved
- ◆ T1
 - ◇ IN2P3-CC : Suffering from same problems 10 days ago (GGUS:63431 [↗](#) and GGUS:62782 [↗](#) : 'hoppingManager process was overloaded for some time, and the hoppingManager stopped working properly'). 15 k files affected which are being migrated manually.
 - Stephane: might the problem be related to high activity (by ATLAS)?
 - ◇ TAIWAN-LCG2 serious transfer errors have been reported (GGUS:63420 [↗](#)).
 - ◇ SARA/NIKHEF queues closed.
 - ◇ Test of FTS 3.2 resumed for UK cloud : if no problem observed, will be put in production beginning of next week
- ◆ T2-T2 tests : Going on
- CMS reports -
 - ◆ Experiment activity
 - ◇ Data taking, processing, analyzing, simulating
 - ◇ Ian: high T0 farm utilization (OK)
 - ◆ CERN and Tier0
 - ◇ Critical service GridMap <http://servicemap.cern.ch/ccrc08/servicemap1.html> [↗](#) lost its interactivity. Stefano submitted GGUS:63450 [↗](#)
 - ServiceMap application is not really supported at the moment, to be followed up in the ticket
 - ◆ Tier1 issues
 - ◇ Reprocessing ongoing
 - ◇ Transfer failures to ASGC from a variety of locations, but seems to be improving
 - ◆ Tier2 Issues
 - ◇ large production ongoing
 - ◇ A large number of Savannah tickets we sent to individual sites to follow up on batch system priority issues observed at Tier-2s
 - ◆ MC production
 - ◇ large production ongoing
 - ◆ AOB
- ALICE reports [↗](#) -
 - ◆ T0 site
 - ◇ GGUS:63460 [↗](#): since yesterday afternoon, and probably Sunday evening, ce202 and ce203 are really slow in advancing the status of submitted jobs. SAM ops tests also fail for those nodes.
 - Gavin: BUG:73765 [↗](#), patch should be certified this week, workaround has been applied later in the afternoon
 - ◇ Upgrades this morning voalice13 (WMS submission, PackMan) and yesterday afternoon voalice11 (CREAM submission), both are back in production
 - ◇ MonALISA was down yesterday evening, due to the expiration of its proxy.
 - ◆ T1 sites
 - ◇ Nothing to report
 - ◆ T2 sites
 - ◇ Scheduled downtime in Bologna T2
- LHCb reports -
 - ◆ Experiment activities: Data taking. Taken in three days ~30% of the total 2010 statistics.
 - ◆ New GGUS (or RT) tickets:
 - ◇ T0: 0
 - ◇ T1: 1

- ◊ T2: 0
- ◆ Issues at the sites and services
 - ◊ T0
 - NTR
 - ◊ T1 site issues:
 - GRIDKA: (GGUS:63353 [↗](#)). LHCb_DST space token running out of space. Waiting for a list of dark files in tokenless space. This is space that could be eventually recuperated and allocated into the DST space token
 - IN2P3: Issue with some files whose tURL is not retrievable. GGUS:63462 [↗](#)
 - RAL: GGUS:63468 [↗](#): some SRM tURLs appear truncated, transient problem

Sites / Services round table:

- ASGC
 - ◆ CMS transfer issues: OK for some T1, bad for others --> network issue, should be fixed now
 - ◆ CASTOR rmmaster daemon got stuck for a few minutes
- BNL
 - ◆ CNAF-BNL network problem: using different circuit BNL-->CNAF traffic rates improved by a factor 10, CNAF-->BNL by a factor 6; not all is understood yet, the investigations continue
- CNAF - ntr
- FNAL - ntr
- KIT - ntr
- IN2P3
 - ◆ open tickets being worked on
- NDGF
 - ◆ network problem with disk servers in Norway, but no experiment data should be unavailable
- NLT1
 - ◆ glibc security updates installed, back in production
- OSG
 - ◆ no copy received of ticket for SLAC, ticket routing being checked
- PIC
 - ◆ behavior of DCAP doors seems to have improved, but some failures remain in PanDA: different problem? Python error messages difficult to interpret; Stephane will involve Rod and Graeme
- RAL
 - ◆ yesterday's CASTOR upgrade OK (ALICE + small VOs)
 - ◆ will look into LHCb ticket
- CASTOR
 - ◆ vendor intervention for HW problem on ATLAS stager to be scheduled next week, to be arranged with ATLAS
- dashboards - ntr
- databases - ntr
- grid services - ntr
- GGUS
 - ◆ major release this Wed Oct 27
 - ◊ 20 new support units
 - ◊ better handling of VOMS groups e.g. for recognizing submitters of alarm or team tickets
 - ◊ bugfixes

AOB:

Tuesday:

Wednesday

Attendance: local(Andrea, Eddie, Harry, Jamie, Jan, Lola, Maarten, Marcin, Maria D, Maria G, Roberto, Simone);remote(Federico, Gang, Gonzalo, Ian, John, Jon, Kyle, Michael, NDGF, Onno, Paolo, Rolf).

Experiments round table:

- ATLAS reports -
 - ◆ ATLAS plans
 - ◇ most likely today (depending on SW readiness) the reprocessing campaign will start. Both detector data and MonteCarlo simulated data will be reprocessed with the same Athena release. Montecarlo reprocessing jobs will be launched one day after to facilitate task brokering to sites. All T1s will be involved + CERN (as backup)
 - See <https://twiki.cern.ch/twiki/bin/viewauth/Atlas/ADCDDataReproAutumn2010> for complete description
 - See <http://indico.cern.ch/conferenceDisplay.py?confId=72836> (A. Klimentov) for general description
 - ◆ T2-T2 full mesh test:
 - ◇ Test for small files (20MB) is basically over. Curing a few tails due to specific problems at a couple of sites.
 - ◇ Medium files (200MB) test will start tomorrow.
 - ◆ GGUS two hours downtime today created a bit of confusion among ATLAS shifters. It was surely announced and properly flagged, may be it is worth reminding people the day before during the 3PM WLCG meeting.
 - ◇ Simone: please include a reminder of important interventions in the minutes of the preceding work day - OK
 - ◇ Maria D: GGUS downtimes are also in the GOCDB, copy them into the experiment's calendar?
 - ◇ Simone: calendar is used only for sites
 - ◆ Following the discussion in ASGC, ATLAS started testing the usage of DPM in TW-FTT as the only T0D1 storage element in ASGC (removing the split between T1 and T2 storage). ATLAS asks all T1s to configure FTS so that the site TW-FTT is treated like another T1 (setting up the proper T1-T1 channels). This request will be made tomorrow at the service coordination meeting.
- CMS reports -
 - ◆ Experiment activity
 - ◇ More data expected this evening
 - ◆ CERN and Tier0
 - ◇ SLS Alarm on one of the Frontier systems. It appeared the system on critical power was monitored and alarmed before the software was enabled. Currently being repaired.
 - ◆ Tier1 issues
 - ◇ Reprocessing ongoing
 - ◇ Transfer failures to ASGC from a variety of locations. Traces to bad configuration on one of the data management agents. Improving.
 - ◆ Tier2 Issues
 - ◇ large production ongoing
 - ◇ A large number of Savannah tickets we sent to sites individual to follow up with batch system priority issues observed at Tier-2s
 - ◆ MC production
 - ◇ large production ongoing
 - ◆ AOB

- ALICE reports [↗](#) -
 - ◆ General information
 - ◇ 5 RAW production cycles ongoing together with a couple of analysis trains
 - ◆ T0 site
 - ◇ GGUS:63460 [↗](#): SOLVED. Submission through ce202 and ce203 is performing well
 - ◆ T1 sites
 - ◇ Nothing to report
 - ◆ T2 sites
 - ◇ Nothing to report

- LHCb reports -
 - ◆ Experiment activities
 - ◇ No new data since yesterday, data reconstruction not problematic.
 - ◆ New GGUS (or RT) tickets:
 - ◇ T0: 1
 - ◇ T1: 2
 - ◇ T2: 0
 - ◆ Issues at the sites and services
 - ◇ T0
 - Files on CASTOR not available. Maybe a faulty diskserver. GGUS:63514 [↗](#)
 - ◇ T1 site issues:
 - CNAF: problem shown in FTS channel CNAF-IN2P3, opened first for IN2P3, now points to file not available at CNAF. GGUS:63506 [↗](#).
 - RAL: file can't be staged. GGUS:63515 [↗](#)
 - RAL: SRM reporting wrong TURL. GGUS:63468 [↗](#). Affecting ~10% of the files, not clear what caused it.
 - ◆ Round table
 - ◇ Gonzalo: check_streams SAM test failure at PIC and other T1 sites: false positive?
 - ◇ Roberto: can be a real problem e.g. at the source, will check

Sites / Services round table:

- ASGC
 - ◆ CMS transfer problem solved
- BNL - ntr
- CNAF
 - ◆ ALICE services downtime tomorrow
- FNAL - ntr
- IN2P3 - ntr
- NDGF - ntr
- NLT1 - ntr
- OSG
 - ◆ data instabilities on Indiana BDII, changes rolled back (DNS issues)
- PIC - ntr
- RAL
 - ◆ 1 disk server read-only and taken out, all its files available on other servers
 - ◆ open tickets in progress

- CASTOR
 - ◆ overload on LHCb default pool Tue evening, unexpected high activity from production user, will be investigated by LHCb
 - ◆ 1 LHCb disk server for file class "temp" essentially lost
 - ◆ reminder: does ATLAS agree with planned intervention Wed next week?
- databases

- ◆ replication of LHCb conditions data from CERN to SARA aborted yesterday, being investigated
 - ◇ Roberto: SARA DB will be avoided while the problem has not been fixed; when the cause has been found, just repopulate the SARA DB from scratch
- dashboards
 - ◆ new Site Status Board for ATLAS in production, many fixes and enhancements
 - ◆ downtime information for GGUS could be incorporated as well

AOB:

- Jamie: the T1 coordination meeting can and should be used to escalate GGUS tickets
- Maria D: Please observe the new GGUS Support Units (SUs) entering production today listed on page <https://gus.fzk.de/pages/didyouknow.php>. In the near future all SUs listed under EMI 3rd level (see list below) will be able to receive assignments **ONLY** by the Deployed Middleware Support Unit (DMSU). Savannah:117155

```
-- AMGA
-- APEL
-- ARC
-- ARGUS
-- CREAM-BLAH
-- DGAS
-- DPM
-- EMI
-- FTS
-- Gridsite
-- Information System/GIP/BDII
-- LFC
-- MPI
-- Proxyrenewal
-- STORM
-- UNICORE-Client
-- UNICORE-Server
-- VOMS
-- VOMS-Admin
-- dCache Developers
-- gLite Hydra
-- gLite Identity Security
-- gLite Java Security
-- gLite L&B
-- gLite Security
-- gLite WMS
-- gLite Yaim Core
-- lcg_util
```

Thursday

Attendance: local(Alessandro, Alexei, Andrea, Eddie, Federico, Gavin, Harry, Jamie, Jan, Lola, Maarten, Marcin, Maria D, Maria G, Simone, Suijian);remote(Gonzalo, Ian, John, Jon, Kyle, Michael, Paolo, Rolf, Ronald, Stefano, Ulf, Xavier).

Experiments round table:

- ATLAS reports -
 - ◆ ATLAS plans: starting of reprocessing campaign postponed, waiting for SW validation. Most likely will start on Friday. Sorry for that.
 - ◆ CASTOR intervention: the sooner the better, well before HI data taking
 - ◆ IN2P3-CC showc 40% failures writing in DATADISK. Message implies the space token is full ("an end-of-file was reached globus_xio: An end of file occurred (possibly the destination

disk is full)") but in fact it is not according to SLS

http://sls.cern.ch/sls/service.php?id=IN2P3-CC_ATLASDATADISK. GGUS:63555 has been created.

◇ Rolf: ticket has some unclear/irrelevant contents (acknowledged), but the real problem is that the dCache request rate is 11 times higher than normal !

◇ Simone: please let us know the user DN and the types of SRM calls

• CMS reports -

◆ Experiment activity

◇ More Data taking. Emergency access yesterday evening for CMS water alarm.

◆ CERN and Tier0

◇ Smooth operations

◇ GridMap issue persists. Shift crews given a work around, but would like to get it fixed.

· ticket is in progress

◆ Tier1 issues

◇ Reprocessing ongoing

◇ CMS is about to start a 500M event simulation campaign for events with pile-up. Relatively short from the CPU standpoint, but IO intensive.

◆ Tier2 Issues

◇ large production ongoing

◆ MC production

◇ large production ongoing

◆ AOB

• ALICE reports -

◆ General information

◇ 5 RAW production cycles ongoing together with a couple of analysis trains

◆ T0 site

◇ Nothing to report

◆ T1 sites

◇ We observed in MonALISA this morning the failure of RAL - CASTOR2 and RAL - TAPE SE

· RAL: old xrootd servers being checked? should no longer matter after CASTOR upgrade; please open a GGUS ticket if there is a real problem

◆ T2 sites

◇ Nothing to report

• LHCb reports -

◆ Experiment activities: Data taking. Reconstruction proceeding. No MC.

◆ New GGUS (or RT) tickets:

◇ T0: 0

◇ T1: 1

◇ T2: 0

◆ Issues at the sites and services

◇ T0

· GGUS:63514 (opened yesterday). Diskserver down. put "on hold". Files re-replicated manually.

◇ T1 site issues:

· IN2P3: Very low pilots efficiency for CREAM CE. Asked to investigate. GGUS:63559

· CNAF: GGUS:63506 (Opened yesterday) A number of files are corrupted. Space token full. Corrupted in the replication process. They will be re-replicated manually.

- RAL: GGUS:63515 [↗](#) (Opened yesterday) File could not be staged. Quickly fixed and closed.
- RAL: GGUS:63468 [↗](#) (Opened 2 days ago) SRM reporting wrong TURL. Not clear what caused it, for the moment we asked our users not to use RAL.

Sites / Services round table:

- ASGC - ntr
- BNL - ntr
- CNAF - ntr
- FNAL - ntr
- IN2P3 - ntr
- KIT - ntr
- NDGF - ntr
- NLT1
 - ◆ conditions DB for ATLAS and LHCb were missing information needed for replication to work, fixed by 3D experts at CERN, now all OK
 - ◇ see databases report below
- OSG
 - ◆ GGUS alarms for yesterday's upgrade were received and routed OK
- PIC - ntr
- RAL
 - ◆ incomplete tURLs for LHCb: hopefully understood and fixed next week when Shaun is back
 - ◆ scheduled downtime next Tue: tape robot will be down + other interventions
- CASTOR - ntr
- GGUS
 - ◆ all OK after yesterday's upgrade
- dashboards - ntr
- databases
 - ◆ yesterday's incident with ATLAS and LHCb conditions DB at SARA was due to a missing step in the documentation, now fixed
- grid services - ntr

AOB:

Friday

Attendance: local(Alessandro, Eddie, Gavin, Harry, Jacek, Jamie, Jan, Lola, Maarten, Maria G, Simone);remote(Andreas, Foued, Gang, Gonzalo, Ian, John, Jon, Michael, Onno, Renato, Rob).

Experiments round table:

- ATLAS reports -
 - ◆ IN2P3-CC overload (follow up): the problem went away but is not fully understood.
 - ◇ The load seems caused by SRM+gridftp writes. Some Panda I/O intensive jobs have been spotted, but not so many (few hundreds). At the same time, Lyon provided a list of DNs as top users, but there is no track in the analysis monitoring about activity from those users (users have now been contacted privately).
 - ◇ Yesterday Lyon has been put offline in Panda and the problem disappeared. It has then been re-put online with a cap on the number of running jobs (1400)
 - ◇ Currently there are several dccp timeouts and staging is being problematic: dCache seems still not healthy (despite the cap).

- Simone: reprocessing campaign starting tonight will have lots of storage activity!
- ◆ Staging problems in NDGF-T1: the status of SRMBringOnline request (used by DDM to determine if the file is online) reports the request as PENDING, even if the file is actually online according to SRMLs. GGUS:63583 [↗](#) has been submitted. It will be escalated to ALARM shortly, since the functionality is needed for reprocessing and the reprocessing is about to start (most likely this evening).
- CMS reports -
 - ◆ Experiment activity
 - ◇ More Data taking.
 - ◆ CERN and Tier0
 - ◇ Smooth operations
 - ◆ Tier1 issues
 - ◇ Reprocessing ongoing
 - ◇ CMS is about to start a 500M event simulation campaign for events with pile-up. Tape family requests sent yesterday. Tickets open at RAL, KIT, ASGC
 - ◆ Tier2 Issues
 - ◇ large production ongoing
 - ◇ Still many batch system tickets open
 - ◆ AOB
- ALICE reports [↗](#) -
 - ◆ General information:
 - ◆ T0 site
 - ◇ Nothing to report
 - ◆ T1 sites
 - ◇ RAL: GGUS:63612 [↗](#) concerning two SE's (RAL-CASTOR2 and RAL-TAPE SE). Both are reporting the same error that files can not be added, moreover RAL-TAPE is pointing to be restarted.
 - possibly AliEn at RAL or the AliEn LDAP service needs to be reconfigured to pick up the new CASTOR situation
 - ◆ T2 sites
 - ◇ Nothing to report
- LHCb reports -
 - ◆ Experiment activities: Data taking. Reconstruction proceeding. Few MC productions.
 - ◆ New GGUS (or RT) tickets:
 - ◇ T0: 0
 - ◇ T1: 1
 - ◇ T2: 0
 - ◆ Issues at the sites and services
 - ◇ T0
 - GGUS:63514 [↗](#) (opened Wednesday). Diskserver down. put "on hold". Files "invisible" for user, who can use replicas. Waiting update.
 - Jan: disk server almost dead, we do what we can
 - ◇ T1 site issues:
 - IN2P3: A number of jobs seg-faulting. GGUS:63573 [↗](#) probably related with GGUS:62732 [↗](#)
 - RAL: GGUS:63468 [↗](#) (Opened 3 days ago) SRM reporting wrong TURL. SRM instance went in half an hour downtime to apply a patch. Hopefully problem solved.

Sites / Services round table:

- ASGC - ntr
- BNL - ntr
- FNAL - ntr
- KIT
 - ◆ computing cluster was split yesterday, 1 LCG-CE and 1 CREAM CE no longer support ALICE (the others do), but ALICE kept sending jobs that immediately failed and caused unnecessary load
 - ◇ AliEn configuration fixed Fri morning, all OK now
- NDGF
 - ◆ disk pools in Norway have been set read-only because of a faulty switch, issue is expected to be fixed on Mon
 - ◆ intervention Mon 09:00 UTC: all pools offline for ~10 min
- NLT1
 - ◆ NIKHEF top-level BDII slapd went 100% CPU yesterday evening, fixed by reboot; upgrade to openldap 2.4 foreseen, first being checked on the site BDII
- OSG - ntr
- PIC - ntr
- RAL
 - ◆ CASTOR for LHCb restarted to cure GGUS:63468 [↗](#), did not work due to another problem, fixed; now probably OK, waiting for SAM test results
 - ◆ ALICE ticket being investigated
 - ◆ CMS tape classes ticket?
 - ◇ Ian: may have been only in Savannah, will follow up
- CASTOR
 - ◆ vendor intervention on ATLAS stager now planned for Tue Nov 2 at 14:00
- dashboards - ntr
- databases
 - ◆ replication to SARA failed again yesterday late afternoon due to incomplete fix, then OK since ~19:00
- grid services - ntr

AOB:

-- JamieShiers - 22-Oct-2010

This topic: LCG > WLCGDailyMeetingsWeek101025

Topic revision: r13 - 2010-10-29 - MaartenLitmaath



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