

Table of Contents

Week of 110321.....	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 110321

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(Alessandro, Dirk, Douglas, Ewan, Fernando, Ian, Jamie, Maarten, Maria D, Maria G, Massimo, Mattia, Simone);remote(Alessandro I, Gareth, Gonzalo, Jhen-Wei, Jon, Kyle, Michael, Roger, Rolf, Ron, Xavier).

Experiments round table:

- ATLAS reports -
 - ◆ IN2P3-CC: storage problems started in the night between Fri and Sat. ATLAS shifter sent a TEAM ticket which was answered during the night. On Sat morning the storage was back (problem in the SRM database). The storage failed again at 19:00 on Sat. ALARM ticket (GGUS:68794) was sent since IN2P3-CC had been assigned a "fat" RAW dataset to be archived on tape. Ticket answered almost immediately, the problem was again fixed but IN2P3-CC experts were not confident about problem not reappearing (help of dCache devs needed). ATLAS blacklisted In2P3-CC for further assignments. The RAW dataset with transfer to In2P3-CC in progress was replicated to a second site. This morning IN2P3-CC was re-included in activities as advised by site managers on Sunday evening.
 - ◇ Rolf: the root cause was a modification introduced on Friday, recommended by the dCache developers; it caused a high CPU load which caused requests to stall; the remedy was a cron job to clean up the SRM DB regularly and update the statistics needed by the query planner (details provided by Ghita Rahal)

- ◆ T0 reconstruction suffered from problems at LSF at CERN. Job submission on saturday become very slow and then completely unresponsive. ALARM ticket was issued (GGUS:68795 [↗](#)). Problem with one (non power) user submitting 180K jobs. Problem fixed, but follow up needed.
 - ◇ Ewan: jobs were submitting other jobs!
 - ◇ Simone: no limit on number of submissions per user?
 - ◇ Ewan: no
 - ◇ Alessandro: could submission be disabled automatically for such rogue users?
 - ◇ Ewan: the matter is being followed up
- ◆ Since March 15, 0.2% of T0 jobs on lxbatch can not write output in AFS. Looks like the AFS token was lost or never acquired. (GGUS:68762 [↗](#))
 - ◇ Ewan: AFS token sometimes for the wrong user
 - ◇ Massimo: that should be fixed now
- CMS reports -
 - ◆ LHC / CMS detector
 - ◇ Good weekend for the accelerator
 - ◆ CERN / central services
 - ◇ CMS User took out LSF with 180k jobs
 - ◆ Tier-0 / CAF
 - ◇ CMS saw the LSF failure as a slow down of Express, but recovered promptly.
 - ◆ Tier-1
 - ◇ NTR
 - ◆ Tier-2
 - ◇ NTR
- ALICE reports [↗](#) -
 - ◆ T0 site
 - ◇ Many jobs get stuck saving their output using xrdcp to SEs which don't respond (e.g. overloaded). Improved timeout handling is being implemented.
 - ◇ Massimo: changes to avoid bad AFS usage?
 - ◇ Maarten: various changes are being worked on by PES and ALICE, the situation should improve in the coming days; for historical reasons the AFS issue is tied to the use of the LCG-CE: we intend to move the whole work flow to CREAM with the LCG-CE as a failover solution when the AFS issue has been fixed
 - ◆ T1 sites
 - ◇ CNAF: GGUS:68799 [↗](#). VOBox was not accessible since Sunday morning until this afternoon. Solved
 - ◆ T2 sites
 - ◇ Several operations
- LHCb reports -
 - ◆ Experiment activities:
 - ◇ MC productions running at lower pace
 - ◇ Validations of work flows for Collision 11 data taking on going.
 - ◆ New GGUS (or RT) tickets:
 - ◇ T0: 0
 - ◇ T1: 0
 - ◇ T2: 0
 - ◆ Issues at the sites and services
 - ◇ T0
 - Recovered the 4225 files on tape lost last week. Many thanks to all CASTOR people involved
 - ◇ T1

- IN2P3 : why we are not able to recover the data at IN2P3 while it has been done for al the other sites ?
 - Rolf: a ticket would help!
 - ◆ there are various issues:
 - ◇ an extra HPSS instance would be needed (!)
 - ◇ repacks may already have happened, so it is doubtful how much of the data can still be recovered
 - ◇ the dCache DB would have to be restored from backup
 - ◆ for those reasons we are very hesitant about the operation
 - ◆ repacks have been stopped, but we cannot tolerate that for a long time
- Mattia: there were shared area SAM test timeouts for IN2P3; will check if the timeout should be increased
 - ◇ T2 site issues:
 - NTR

Sites / Services round table:

- ASGC
 - ◆ CASTOR has been slow for CMS, being investigated
- BNL
 - ◆ CNAF-BNL network ticket GGUS:61440: during the weekend a rate limitation of 75 Mbit/s (!) in layer 2 was discovered on the DANTE segment; DANTE looking into it
- CNAF
 - ◆ ALICE VOBOX is VM with 5 GB RAM: does some daemon have a memory leak?
 - ◇ Maarten: yes, the Cluster Monitor; a fix is available, we will follow up
- FNAL - ntr
- IN2P3 - nta
- KIT
 - ◆ the mistakenly deleted LHCb files have been restored
- NDGF
 - ◆ short downtime tomorrow for SRM reboot
- NLT1
 - ◆ downtime tomorrow for RAID controller firmware upgrade
- OSG
 - ◆ the corrected BNL records have been sent to the SAM team for the availabilities to be recalculated
- PIC
 - ◆ on Sat a 10 TB ATLAS disk server hung and could not be rebooted; the main board was changed today, the server is OK now
- RAL
 - ◆ transient problem with ATLAS SW area Sat evening
 - ◆ at-risk downtime tomorrow for network intervention
- CASTOR - ntr
- dashboards
 - ◆ various SAM test monitoring pages have been moved to different machines for which the ports were not reachable outside CERN; not clear if all is OK now
- GGUS/SNOW
 - ◆ VOC meeting was held this morning, minutes at https://twiki.cern.ch/twiki/pub/LCG/WLCGVOCCoordination/VObox_application_support_by_WLC
 - ◆ a team ticket does **_not_** have a higher priority than a user ticket: the perceived higher priority is due to the **_site_selection_** which triggers an e-mail sent to the admins of the site, hence an

immediate notification; that feature is also available for user tickets, but often not used there

- grid services - nta

AOB:

Tuesday:

Attendance: local(Alessandro, Douglas, Eva, Ewan, Jamie, Maarten, Maria D, Maria G, Massimo, Mattia, Pedro, Roberto);remote(Dimitri, Gonzalo, Ian, Jeff, Jhen-Wei, John, Jon, Karen, Kyle, Michael, Rob, Roger, Rolf).

Experiments round table:

- ATLAS reports -
 - ◆ T0, Central Services
 - ◇ Eowyn job scheduler for tier-0 processing failed last night. This was restarted, but job warnings continued until the morning. Experts said these warnings could be ignored, no other problems from this were seen.
 - ◇ Warnings this morning about a large backlog in files migrating to castor from Tier-0, over 10k. Experts looked at this, and said it wasn't a serious problem.
 - Massimo: a tape library was offline
 - ◆ T1
 - ◇ IN2P3 was put back into service for data exports today, as storage issue were fixed yesterday. No new issues seen so far.
 - ◇ The NL cloud has been in downtime today, there is a scheduled outage of the SRM at SARA.
 - ◇ Network issues in getting to RAL, but from somewhere outside RAL. This was diagnosed late last night as 6% packet loss to the LFC there, causing job problems in the UK cloud. There was a network at risk this morning, and this was only ticketed recently once that was over. (GGUS:68850 [↗](#))
 - John: looking into it, seems a non-local network issue, not easy; today's network maintenance is a blind alley, it was just a reboot
 - ◇ Network firewalls updated at PIC today, and site was put at risk. No actual issues reported.
 - ◆ T2/T3
 - ◇ Problems at SWT2_CPB, site taken offline (GGUS:68782 [↗](#)). Was traced to a broken nic card on a server machine, this was fixed, and site is now back online.
 - ◇ Multiple problems at GOEGRID, and site has been offline all day (GGUS:68671 [↗](#)). A number of issues solved, but currently seems to be a software issue with ATLAS software, and site is still offline.
 - ◇ OU_OCHEP_SWT2 unscheduled outage due to a movement of a server rack (Savannah:119887 [↗](#)).
- CMS reports -
 - ◆ LHC / CMS detector
 - ◇ Moving to lower energy running tomorrow
 - ◆ CERN / central services
 - ◇ NTR
 - ◆ Tier-0 / CAF
 - ◇ Prompt-Reco from the weekend has been injected. More steady use of the Tier-0 farm.
 - ◆ Tier-1
 - ◇ Poor quality transfers between CERN and FNAL. Traced back to a bad checksum on the CERN source. FNAL verifies the checksum. In the past these have not normally

Tuesday:

come as singlets. Need to keep an eye out.

· Massimo: please let us know which files gave problems

◆ Tier-2

◇ NTR

• ALICE reports [↗](#) -

◆ T0 site

◇ INC:024084 [↗](#): One CAF node with data is unreachable

· Maria D, Alessandro: better use GGUS also for those tickets, to avoid suboptimal routing by Service Desk

◇ WMS (LCG-CE) submission is disabled at CERN, will only be used as a failover mechanism until PES implement the change to prevent those jobs from getting AFS tokens.

◆ T1 sites

◇ CNAF: we applied this morning the patch to the Cluster Monitor that we are using at CERN and KIT which should solve the problem of the memory consumption. This fix will be included in the new AliEn subversion

◇ IN2P3: Tomorrow from 9AM onward we are going to try Torrent at the site.

· Rolf: note that only **tests** have been okayed for now! official approval by the IN2P3-CC management will be needed to go further

· Massimo: what about CERNVMFS?

- Maarten: not an option at this time, because it would require development on the ALICE side; the Torrent solution is supported by the AliEn framework and requires nothing to be present on the WN - it requires outbound access to alitorrent.cern.ch and the possibility for WN to serve other WN in the same site (at least the ones in the same cluster)

◇ Mattia: VOBOX proxy renewal tests are failing at KIT since the old SAM tests were replaced with the new Nagios tests at the end of the morning

· Maarten: we will look into it

◆ T2 sites

◇ Several operations

◇ GRIF_IPNO: GGUS:68849 [↗](#). The information provider is reporting wrongly (44444 jobs). Experts are looking at it

• LHCb reports -

◆ Experiment activities:

◇ MC productions running at full steam (35K jobs last 24hs) + validation of work flows for Collision11

◇ LHCbDIRAC week on going at CERN

◆ New GGUS (or RT) tickets:

◇ T0: 0

◇ T1: 0

◇ T2: 0

◆ Issues at the sites and services

◇ T0

· Validation work-flow for EXPRESS was failing at CERN apart few jobs running on pre-production nodes (CVMFS). AFS installation was corrupted, reinstalled the software now it is OK

◇ T1

· GridKA: recovered the SDST lost last week and put in the vobox at GridKA. Many thanks to GridKA people. Our T1 VOBOX responsible will re-register them on dcache.

· HPSS and d-cache masters are working to recover as much as possible of the

lost data at IN2p3. Thanks for the effort !

- ◇ T2 site issues:
- NTR

Sites / Services round table:

- ASGC - ntr
- BNL - ntr
- CNAF - ntr
- FNAL - ntr
- IN2P3 - nta
- KIT - ntr
- NDGF
 - ◆ downtime took longer than foreseen, not finished, so another 30-minute downtime has been scheduled for tomorrow
 - ◇ Douglas: in GOCDB?
 - ◇ Roger: yes
- NLT1
 - ◆ scheduled downtime going OK
- OSG
 - ◆ USATLAS T2 names for FTS yesterday?
 - ◇ Alessandro: we will follow up through e-mail
 - ◇ see "FTS channel setup for cross-cloud transfers" document on <http://indico.cern.ch/conferenceDisplay.py?confId=131670>
- PIC - ntr
- RAL
 - ◆ at-risk network intervention went OK

- CASTOR - nta
- dashboards
 - ◆ SAM tests web pages OK now
- databases - ntr
- GGUS/SNOW - ntr
- grid services - ntr
- GT group - ntr

AOB:

Wednesday

Attendance: local(Alessandro, Douglas, Edoardo, Eva, Ewan, Fernando, Jamie, Maarten, Maria D, Maria G, Massimo, Mattia, Nilo);remote(Dimitri, Gonzalo, John, Jon, Kyle, Michael, Onno, Roger, Rolf).

Experiments round table:

- ATLAS reports -
 - ◆ T0, Central Services
 - ◇ No news to report.
 - ◆ T1
 - ◇ NDGF downtime today for ~1.5 hours, to finish up changes there.
 - ◇ Disk server problems at Taiwan are causing reprocessing job failures, this is getting fixed.
 - ◇ RAL network problems continue today, and is getting diagnosed from many sites in the UK cloud. Still causing LFC access issues, and no news on a fix yet.

- John: the problem is being investigated
- ◆ T2/3
 - ◇ SRM contact problems at Weizmann.
 - ◇ RHUL and QMUL in downtime today.
- CMS reports -
 - ◆ LHC / CMS detector
 - ◇ Good instantaneous luminosity. CMS trigger well behaved
 - ◆ CERN / central services
 - ◇ We appear to have a lost file in Castor. Savannah Ticket Savannah:119911 [↗](#)
 - ◆ Tier-0 / CAF
 - ◇ Prompt-Reco from the weekend has been injected. More steady use of the Tier-0 farm.
 - ◆ Tier-1
 - ◇ Poor quality transfers between CERN and FNAL. Comes from bad file reported above.
 - ◆ Tier-2
 - ◇ NTR
- ALICE reports [↗](#) -
 - ◆ T0 site
 - ◇ INC:024084 [↗](#): One CAF node with data was unreachable. It was crashing all the time due to kernel panics in the AFS module. The AFS experts advised an upgrade to the latest kernel and AFS version, which seems to have fixed the problem.
 - ◆ T1 sites
 - ◇ IN2P3: This morning we switched to Torrent as a test, but all the jobs failed since that time and nothing seems to be running at the moment; the matter is being debugged.
 - discovered after the meeting: ALICE jobs were blocked in the batch system because of excessive space usage in /tmp due to a configuration parameter being absent in the ALICE LDAP service; the configuration was switched back to AFS and jobs were running again from 17:00 onward; to be continued next week...
 - ◇ Mattia: today the Nagios test results look better; we will need to correct for yesterday's errors
 - ◆ T2 sites
 - ◇ Several operations
- LHCb reports -
 - ◆ Experiment activities:
 - ◇ MC productions and validation of work flows for Collision11
 - ◇ LHCbDIRAC week on going at CERN
 - ◆ New GGUS (or RT) tickets:
 - ◇ T0: 0
 - ◇ T1: 0
 - ◇ T2: 0
 - ◆ Issues at the sites and services
 - ◇ T0
 - NTR
 - ◇ T1
 - HPSS and d-cache masters are working to recover as much as possible of the lost data at IN2p3. Thanks for the effort (for traceability GGUS:68889 [↗](#))
 - Rolf: work ongoing; by chance a necessary older backup of the critical DB happened still to be available!

- Mattia: SAM jobs for SARA are failing with Maradona errors; will follow up with LHCb
- ◇ T2 site issues:
 - NTR

Sites / Services round table:

- BNL - ntr
- CNAF - ntr
- FNAL - ntr
- KIT
 - ◆ downtime tomorrow 9-10 AM for router reconfiguration
- NDGF - nta
- NLT1
 - ◆ will look into SARA job failures for LHCb
- OSG
 - ◆ BNL availability recalculation now tracked in GGUS:68768
- PIC - ntr
- RAL - nta

- CASTOR
 - ◆ see CMS report: no ticket yet for deleted file
- dashboards - nta
- databases
 - ◆ CMS PVSS replication has performance problems, not understood, being investigated
- GGUS/SNOW
 - ◆ see AOB
- grid services - ntr
- networks - ntr

AOB: (MariaDZ) Notes from the 2011/03/21 meeting with the WLCG VOCs on user support are in https://twiki.cern.ch/twiki/pub/LCG/WLCGVOCCoordination/VObox_application_support_by_WLCG_VOCs.txt. The SNOW-to-GGUS direction of the interface will enter production Monday 2011/03/28. Many thanks to Maite Barroso for the extensive testing.

Thursday

Attendance: local(Douglas, Eva, Ewan, Jamie, Maarten, Maria D, Maria G, Massimo, Mike, Nilo, Pedro);remote(Claudia, Foued, Gonzalo, Joel, John, Jon, Kyle, Michael, Rolf).

Experiments round table:

- ATLAS reports -
 - ◆ T0, Central Services
 - ◇ Eowyn job scheduler crashed again last night. The shifter handled this in discussion with experts, and left the note that it was because of LSF. Not sure what the problem was with LSF last night.
 - Ewan: there was an issue with various overloaded WN ~1-2 AM
 - Douglas: that does not look related
 - ◆ T1
 - ◇ Network issues with RAL continue. The ggus ticket (GGUS:68850) was updated with more feedback on failures to contact the LFC, but no more feedback from RAL networking about what the problem might be, or further actions that have been taken.
 - John: the issue is being pursued actively

- ◆ T2/3
 - ◇ UTD outage, and request that shifters need to put the queues offline.
 - ◇ TW-FTT most jobs are failing here, about 250 in the past 12 hours, all with put errors. There has been a ticket on this since 9am (GGUS:68914[↗](#)), but no response yet.
 - ◇ SLACXRD sched. downtime, this was announced and no problems here.
- CMS reports -
 - ◆ LHC / CMS detector
 - ◇ Good instantaneous luminosity over a long unexpected fill at 7TeV overnight. CMS trigger well behaved.
 - ◆ CERN / central services
 - ◇ We appear to have a lost file in Castor. Savannah Ticket Savannah:119911[↗](#)
 - Massimo: this case does not look encouraging so far; any extra information would be useful
 - Ian: we will dig up the logs
 - ◇ CMS Issue with a transfer component sending release validation from FNAL to CERN. No indication it's a CERN central component. Looks like a data management component, if anything.
 - ◇ Massimo: yesterday around 6 PM the default pool had a high load due to user "cmsprod": is that OK?
 - ◇ Ian: such usage of the default pool is unexpected, we will look into it
 - ◆ Tier-0 / CAF
 - ◇ Since March 20 reasonable utilization of the CERN CPU.
 - ◆ Tier-1
 - ◇ KIT in scheduled downtime
 - ◆ Tier-2
 - ◇ SAM CE Failure at T2_FR_IN2P3. Savannah ticket Savannah:119956[↗](#). Just the production test, but logs just indicate the jobs abort.
- ALICE reports [↗](#) -
 - ◆ T0 site
 - ◇ The upgrade of the kernel and the openafs version has been done on 2/3 of the machines. The rest of them were in production and have sensitive data, so we will do the upgrade later. This seemed to solve the problems reported during this week (nodes crashing several times)
 - ◆ T1 sites
 - ◇ Noting to report
 - ◇ Mike: the tests that failed yesterday are OK now
 - ◇ Maarten: we switched off some tests that look irrelevant to some extent; we will improve the tests in the course of this year
 - ◆ T2 sites
 - ◇ Usual operations
- LHCb reports -
 - ◆ Experiment activities:
 - ◇ MC productions and validation of work flows for Collision11
 - ◆ New GGUS (or RT) tickets:
 - ◇ T0: 0
 - ◇ T1: 0
 - ◇ T2: 0
 - ◆ Issues at the sites and services
 - ◇ T0
 - NTR

◇ T1

- HPSS and d-cache masters are working to recover as much as possible of the lost data at IN2p3. Thanks for the effort (for traceability GGUS:68889 [↗](#))
 - Rolf: Pierre Girard needs to know where to restore the files
 - Joel: we will follow up in the ticket

◇ T2 site issues:

- NTR

Sites / Services round table:

- BNL
 - ◆ GGUS:61440 [↗](#): bandwidth cap has been removed, the link looks OK now
- FNAL - ntr
- CNAF
 - ◆ 2 downtimes next week during the LHC technical stop to upgrade StoRM to version 1.6.2: on Tue Mar 29 14-18 UTC for ATLAS, on Wed Mar 30 14-18 UTC for LHCb
- IN2P3 - ntr
- KIT
 - ◆ downtime today 15-16 UTC for dCache security update
- OSG
 - ◆ the OSG and EGI ticketing systems got into mail loop causing tickets to be opened on both sides until admin interventions on both sides; the cause was the accidental inclusion of the OSG ticketing address in the CC of an EGI ticket; the OSG system will be configured to recognize the address of the EGI system and avoid such loops
- PIC - ntr
- RAL - ntr

- CASTOR - nta
- dashboards - nta
- databases - ntr
- GGUS/SNOW - ntr
- grid services
 - ◆ a SIR has been created for the LSF incident during the weekend; Platform informed us of an option to limit the number of pending jobs per user, which would have avoided the problem; a limit of 20-30k seems reasonable
- GT group - ntr

AOB:

Friday

Attendance: local(Dirk, Douglas, Eva, Ewan, Jamie, Jan, Maarten, Maria G, Massimo, Mattia, Nilo, Simone);remote(Alexander, Christian, Felix, Gonzalo, Ian, Joel, John, Jon, Karen, Michael, Rob, Xavier).

Experiments round table:

- ATLAS reports -
 - ◆ T0, Central Services
 - ◇ Serious problems with castor today, starting a little after noon. First noticed in the tier-0 system, and an alarm ticket was created (GGUS:68949 [↗](#)). Fairly quickly the problem was found, and services were fixed. Tier-0 systems showed things working right away. Not sure what the error actually was at this time.
 - Jan: a wrong configuration file was copied by accident when some new disk servers were moved in for ATLAS and CMS

- ◆ T1
 - ◇ LFC access problems caused many jobs failures overnight in the FR cloud.
(GGUS:68969 [↗](#)), most around 2:00-3:00, but a few seen up until 9:00 this morning.
 - Rolf: ticket in progress, the cause is not clear yet; also some other problems were seen for ATLAS jobs
 - ◇ RAL network issues persist today, but no feedback on the issue in the last 24 hours.
(GGUS:68850 [↗](#))
 - John: some packet losses were seen on some links, but that seems to have gone away and may be unrelated; our network experts are looking further into the matter
 - ◇ Mattia: to clear up some confusion - since Tue last week ATLAS are using the (new) SAM-Nagios tests instead of the old SAM tests
- ◆ T2/3
 - ◇ Weizmann back down again with SRM contact issues.
 - ◇ RHUL showing SRM contact issues.
- CMS reports -
 - ◆ LHC / CMS detector
 - ◇ Running at 2.75TeV
 - ◆ CERN / central services
 - ◇ Savannah Ticket Two blocks with files containing bad checksums
Savannah:119969 [↗](#)
 - Massimo: please open a SNOW ticket
 - ◇ Alarm ticket issues on Castor GGUS:68952 [↗](#)
 - see ATLAS report
 - ◆ Tier-0 / CAF
 - ◇ Record lumi runs from earlier in the week working through the system now. CPU efficiency looks good
 - ◆ Tier-1
 - ◇ NTR
 - ◆ Tier-2
 - ◇ NTR
- ALICE reports [↗](#) -
 - ◆ T0 site
 - ◇ Nothing to report
 - ◆ T1 sites
 - ◇ Nothing to report
 - ◆ T2 sites
 - ◇ Usual operations
- LHCb reports -
 - ◆ Experiment activities:
 - ◇ MC productions and validation of work flows for Collision11
 - ◆ New GGUS (or RT) tickets:
 - ◇ T0: 0
 - ◇ T1: 0
 - ◇ T2: 0
 - ◆ Issues at the sites and services
 - ◇ T0
 - CVMFS : put back in prod with 5 GB cache
 - ◇ T1
 - NTR
 - ◇ T2 site issues:

· NTR

Sites / Services round table:

- ASGC - ntr
- BNL - ntr
- CNAF - ntr
- FNAL - ntr
- IN2P3 - nta
- KIT - ntr
- NDGF
 - ◆ tape system outage on Monday, some ATLAS and ALICE data temporarily unavailable
- NLT1 - ntr
- OSG - ntr
- PIC - ntr
- RAL
 - ◆ 2 outages next week for CASTOR 2.1.10 upgrades: Mon for CMS, Wed for all others

- CASTOR - nta
- dashboards - ntr
- databases
 - ◆ CMS PVSS replication performance problems: cause unknown, being followed up with Oracle support; in the meantime the setup has been split and the performance is OK now
 - ◆ Oracle bug affecting CMS DB: a patch is available but will not be applied at this time, because the query has been changed such that the bug is avoided
- grid services - ntr

AOB:

-- JamieShiers - 18-Mar-2011

This topic: LCG > WLCGDailyMeetingsWeek110321

Topic revision: r15 - 2011-03-25 - 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