

Table of Contents

Week of 100614.....	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.....	8

Week of 100614

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 & Broadcasts		Change assessments
ALICE	ATLAS	CMS	LHCb	WLCG Service Incident Reports	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(Harry(chair), Jaroslava, Laurence, Ulrich, Eva, Lola, PeterK, Andrea, JanI, Jamie, MariaG, MariaDZ, Alessandro, Jean-Philippe, Gavin, Dirk, Roberto, Simone);remote(Jon+Catalin(FNAL), Joel(LHCb), Gonzalo(PIC), Michael(BNL), Gang(ASGC), Angela(KIT), Gareth(RAL), Rob(OSG), Rolf(IN2P3), Onno(NL-T1)).

Experiments round table:

- ATLAS reports -

CERN-PROD_SCRATCHDISK: issue got solved on Saturday morning, Thanks!
https://gus.fzk.de/ws/ticket_info.php?ticket=58904

Overloaded BDIIs at several sites over the weekend resulted in failing jobs with problems retrieving inputs: Australia-Atlas, JINR, Nikhef, INFN-T1. It's becoming an issue over past 2 months. What is the underlying issue? Laurence Field asked for more details - which sites and clients. Alessandro reported CERN, RAL and INFN and that worker nodes querying their local bdii were timing out. Peter Kreuzer thought CMS are also seeing this in some SAM tests at INFN and some Tier 2s. Laurence thought the client timeout was hard-wired at 15 seconds but he will follow-up this up together with Alessandro.

Dear T1s, please take opportunity of this week to perform FTS upgrade! LHC is commissioning this week, therefore not much data to take, no data to export to T1s from ATLAS (except functional tests).

- CMS reports -

T0 Highlights 1) cosmics. 2) data taking during machine development. 3) CSM DB issue from last week : T0 Data Base Service (TOAST) ok now, after restarting DBS, moving it away from the DB-node TOAST is using. 4) Today : intervention on the storage manager DB at point 5. Afterward affected the CMS T0 (SM host down ?); T0 team changed settings and then working again.

T1 Highlights: 1) 2 GGUS Team tickets opened to INFN-T1 ((tickets :https://gus.fzk.de/ws/ticket_info.php?ticket=58920 and https://gus.fzk.de/ws/ticket_info.php?ticket=58987), may or may not be the same issue), both related to SW configuration : admins at INFN-T1 made a general fix and asked CMS to confirm things are fine now, to be followed up 2) BDII not Visible issue at INFN-T1 on Saturday Jun 12, covered in the same ticket as above : https://gus.fzk.de/ws/ticket_info.php?ticket=58987

T2 Highlights: 1) MC production as usual. 2) BDII not Visible issues at the 2 CMS Lisbon T2s : see output from the Dashboard Site Status Board :
http://dashb-ssb.cern.ch/templates/cache/bdii_log.html#T2_PT_NCG_Lisbon and
http://dashb-ssb.cern.ch/templates/cache/bdii_log.html#T2_PT_LIP_Lisbon

- ALICE reports [↗](#) -

GENERAL INFORMATION: Pass 1 reconstruction activities together with two analysis train ongoing. No MC production activities during the weekend. In terms of raw data transfers, very low activity for the moment

T1 sites: CNAF: on Sunday morning the experts detected a problem at the local ce07 (CREAM). Connections to the service were being refused at submission time. The site admin was immediately informed and he took actions in few hours (therefore there is no GGUS ticket)

FZK: also on Sunday morning the restart of all the local services at both VOBOX was necessary. ALICE experts investigating why this activity is needed from time to time at several voboxes

T2 sites: Clermont: On Saturday night experts detected a wrong information reported by the local information system of the CREAM-CE: A large amount of Alice jobs appeared in status running while the experiment had stopped submitting new agents since almost 24h. The issue was reported to the ALICE expert at the site during the weekend and this expert has confirmed this morning the problem is gone. Site is back in production

Cagliari and CyberSar-Cagliari: Both sites are out of production. The local ALIEN user proxy expired in both local VOBOXES. The responsible has been informed, waiting for his actions

- LHCb reports -

Experiment activities: Running several MC productions at low profile. Merging production. GGUS (or RT) tickets:

T1 site issues: CNAF no shared area variable defined (GGUS:58985 [↗](#)). IN2p3: SRM endpoint not available on Saturday. SAM tests confirm this outage. GGUS:58994 [↗](#)

T2 sites issues: Shared area issues at: BG05-SUGrid(GGUS:59015 [↗](#)) IL-TAU-HEP(GGUS:59007 [↗](#)).

Sites / Services round table:

FNAL: Up to date with FTS.

ASGC: Will migrate to FTS 2.2.4 tomorrow.

Monday:

KIT: Migrated FTS last week. Now planning to go to FTS under slc5. Had a hardware problem with their CMS dcache headnode during the weekend - 4 hour downtime.

RAL: Had bdii issues over the weekend (had missed an upgrade) and were failing ops SAM tests. Ready to perform the FTS upgrade and proposing Wednesday. Installation problems with new disk servers last week for ATLAS and CMS - fixed by Friday. Gareth queried if the SAM to Nagios switchover is still scheduled for the 15 June. Jamie replied this date had been to mesh with an MB meeting that has now been postponed. There are differences in the test results (there are different algorithms of course) that MB members would like to fully understand so there is no new date yet. Gareth also queried if the access to the results database will change - Harry to follow-up.

OSG: Problems in the OSG-GGUS interface on the OSG side over the weekend but no updates were lost.

IN2P3: The LHCb srm crash outage was due to a known bug which is protected by an autorestart and the service was back before the ticket was created.

NL-T1: Will upgrade their FTS tomorrow.

CERN CEs: Planning to migrate four lcg-CE that submit to slc4 to submit to slc5 as soon as possible and to update the CREAM-CE to release 3.2.6. Will take some time as they need to be drained of jobs.

CERN CASTOR: Adding 128 TB to the ATLAS Tier 3 disk pool as requested by B. Panzer.

AOB: Simone reported that Friday's high OPN traffic between CERN and RAL was an ATLAS user of CERN LSF batch reading data from RAL. He was surprised that CERN worker nodes are on the OPN as he did not think other Tier 1 were configured that way. Gareth pointed out that if not on the OPN then the GPN would have been overloaded.

Tuesday:

Attendance: local(Harry(chair), Alessandro, Lola, PeterK, JanI, Ulrich, Roberto, Maarten, Gavin, Eva);remote(Angela(KIT), Michael(BNL), Joel(LHCb), Jeremy(gridpp), Catalin(FNAL), Gang(ASGC), Ronald(NL-T1), Rolf(IN2P3), Tiju(RAL), Rob(OSG)).

Experiments round table:

- ATLAS reports - CERN-PROD_SCRATCHDISK: locality of some files in CERN-PROD_SCRATCHDISK is LOST GGUS:59035. Jan Iven reported this is due to the disk being drained. The system knows the files are really NEARLINE (i.e. on tape) and cannot access them with minimal latency so flags them as LOST. Maarten thought that srm-ls should not, however, report them as LOST which it apparently does.

Today Downtime: 1) SARA-MATRIX FTS upgrade 9 - 12 CET. 2) TAIWAN FTS upgrade 10 - 11 CET. 3) BNL FTS upgrade 17-18 CET.

Last Friday and Monday the ATLAS disk space availability monitor in the CERN SLS was showing grey for dcache sites due to their CRLs not being updated. Maarten reported two problems - the master CRLs at CERN have a lot of outside accesses and are on busy AFS volumes (this is being worked on) - and the dcache client insists all lines of a CRL are parseable and recently a French site had a zero length record which killed the client.

- CMS reports - T0 Highlights: 1) cosmics. 2) data taking during machine development.

T1 Highlights: Many re-reco and skimming jobs stuck at various T1s due to central WMS issue, see <http://savannah.cern.ch/support/?115119>. identified by WMS admins as a problem with condor in wms012 at CNAF. It is stuck with 40Kjobs, the wms is now drained and not accepting jobs any longer. The admins are debugging the issue. Jobs will be aborted manually and need to be re-submitted urgently.

T2 Highlights: MC production as usual.

Weekly-scope Operations plan

[Data Ops]:

Tier-0: data taking if machine provides stable beam, otherwise modest testing

Tier-1: re-reconstruction passes and MC re-digitization/re-reconstruction

Tier-2: Plan on running full-scale MC production at all T2s

[Facilities Ops]:

Fixing various bugs on the CMS WebTools front, which had to be often restarted or operated manually in last 10 days : SiteDB, PhEDEx Web

Finalizing Vidyo accounts for remote CMS Centers participating to Computing Shifts

Continue to test and integrate Critical Service recovery procedures for Computing Run Coordinator (CRC)

- ALICE reports - GENERAL INFORMATION: Four analysis trains ongoing today with no MC production activity expected for the moment.

T0 site: No issues to report

T1 sites: 1) CNAF: The site admin is still working on the local CREAM-CE issue reported yesterday. The CREAM-CE developers have suggested to migrate the current service to CREAM1.6/SL5. 2) Minimal activity today at the T1 sites

T2 sites:

Cagliari and CiberSar-Cagliari: The issue reported yesterday (Alien user proxy expired) has been solved. Both sites are back in production

RRC-KI: The ALICE responsible person has announced a cooling problem at the site. Services out of production

Kosice: CREAM-CE out of production. Issue announced to the responsible peron at the site

Hiroshima: The local SE has been taken out of production for xrootd update

- LHCb reports - Experiment activities: 1) Running several MC productions at low profile. 2) Merging production.

T0 site issues: FTS transfers out of CERN failing (GGUS:59037). Jan Iven reported there are 2 overloaded disk servers in the default service class which is the first one looked at by srm. This order can be changed so LHCb should consider this.

Tuesday:

T1 site issues: CNAF : Problems transferring to CNAF with FTS (GGUS:59038[?])

Sites / Services round table:

- BNL: Performing the FTS upgrade - a downtime of less than 1 hour is expected.
- ASGC: The FTS upgrade was completed this morning.
- NL-T1: The FTS upgrade was completed this morning.
- RAL: The FTS upgrade will be done tomorrow morning.
- CERN CE: The four CE124 to 127 will be drained tonight to be converted into submitting to slc5 worker nodes. This will still leave 8 submitting to slc4.
- CERN CREAM-CE: Planning to upgrade one CREAM-CE to release 1.6 on Thursday. Would like ALICE to confirm ahead.
- CERN CASTOR: There was a 20 minute glitch on castorlhcb this morning while moving machines.
- CERN VOMS: The host certificate of voms.cern.ch will be updated on Wednesday 16th June 10:00 CEST. If you have the lcg-vomscerts package installed on your service then you must have updated to version 5.9.0-1 of this package by this time.

AOB: Laurence Field has looked into the ATLAS bdii client timeouts reported yesterday and concluded they were mostly individual glitches. The RAL incident happened because the CERN bdii exceeded 5MB with the addition of some software tags and RAL had not yet applied the increase to 10MB that was in the last release of the bdii configuration. Laurence/Maarten have suggested that the GlueLocation object, which takes 1.5MB, may not be necessary as its information is duplicated in the SoftwareRunTimeEnvironment so would experiments please check if they use this object and let us know (wlcg-scod@cernNOSPAMPLEASE.ch).

There will be a reduced attendance for the next 3 days due to the data management jamboree in Amsterdam.

Wednesday

Attendance: local(Harry(chair), Lola, Ulrich, JanI, David, MariaD, Eduardo, Oliver, Alessandro);remote(Gonzalo(PIC), Cristina(CNAF), Gang(ASGC), Xavier(KIT), Rolf(IN2P3), Joel(LHCb), Tiju(RAL), Rob(OSG), Catalin(FNAL)).

Experiments round table:

- ATLAS reports - Follow up in <https://savannah.cern.ch/bugs/index.php?62269>[?] (currently at discussion level) for yesterday's CERN-PROD SCRATCHDISK draining files reported as LOST instead of NEARLINE - GGUS:59035[?]

Today Downtime: RAL FTS upgrade 9 - 13 CET.

- CMS reports - Tier 0 : following the LHC schedule and testing the Tier 0 processing infrastructure.

Tier 1: CNAF have patched their WMS that caused many stuck jobs yesterday and the jobs are running again but merge jobs cannot be dispatched (they need processing jobs to report back) and Tier 1 merge areas are filling up. The 175TB at FNAL is already full and processing has been stopped there. Would other sites please tell CMS if their merge areas are getting full. CMS estimate they have lost 4 days worth of processing (the ticket was raised on a Saturday but not acted on till Monday) and will be thinking again about how to handle

priority problems. Over the same weekend there was a record of 16000 concurrent jobs at Tier 2.

- ALICE reports [↗](#) -

T0 site: Concerning the migration to CREAM1.6 of the current CREAM services at CERN, ALICE gives green light for this migration. Please inform the experment about date-time when to stop sending jobs to these systems.

T1 sites: Raw data transfer yesterday to FZK and CNAF with no incidents to report . All T1 sites in production today with a minimal activity.

T2 sites: KFKI: last site without CREAM system. The system has been fully configured and announced to Alice. currently working in the testing phase to put it in production

- LHCb reports - 1) So far no replies on the CERN and CNAF tickets from yesterday. For CERN Jan Iven clarified that LHCb is best placed to decide on the order in which srm searches CASTOR service classes. Joel would, however, like to know if there are any side effects so Jan will reply to that. For the CNAF FTS problems Cristina reported there have a disk problem that is being worked on then the ticket will be updated. 2) Ok for CERN to go ahead with the CREAM-CE upgrade on ce201. 3) At CERN bjobs -N does not seem to return the right normalised cpu times used by a job - Ulrich will take this offline. 4) Will reply to PIC about failing LHCb jobs on Monday. 5) A CERN ticket on failure to receive an alarm SMS had been closed without comment but now seems to be working as it should. Harry will follow up.

Sites / Services round table:

- PIC: FTS 2.2.4 was installed last week.
- CNAF: Finishing FTS 2.2.4 testing - hoping to make the upgrade tomorrow.
- BNL (by email): FTS was successfully upgraded at BNL to v 2.2.4.
- NT-T1(by email): NTR.
- RAL: FTS upgrade completed this morning.
- FNAL: CMS merged area full and failed a few SAM tests.
- CERN CEs: The 4 lcg-ce submitting to slc4 are now being drained. Will try and migrate them tomorrow in the same time slot as the CREAM-CE.
- CERN Databases: Addressing what seems to be an Oracle bug affecting the LHCb dashboard.
- CERN VOMS: Host certificate for voms.cern.ch was updated at 10:45 this morning.
- CERN CASTOR: ATLAS SCRATCHDISK draining has finished and all files are accessible again. Are replacing some hardware in the LHCb MDST space token which will provoke some disk to disk copies.

AOB: Alessandro reported that ATLAS had closed the ggus ticket as solved but left open the Savannah report. MariaD reminded that normal policy is to only close a ggus ticket that led to a Savannah bug report when the bug is closed in Savannah but agreed with Alessandro that this ticket was different in that the issue was inaccessible files which are now accessible so it makes sense to close this ggus ticket. Should a similar incident happen a fresh ggus ticket will be raised.

Thursday

Attendance: local(Harry(chair), Oliver, JanI, Akshat, Lola, Miguel, Eva, Ulrich, MariaD);remote(Gonzalo(PIC), Xavier(KIT), Joel(LHCb), Ricardo(INFN), Catalin(FNAL), Gang(ASGC), Rob(OSG), John(RAL), Ronald(NL-T1), Rolf(IN2P3), Jeremy(GridPP), Tore(NDGF), Alessandro(ATLAS)).

Experiments round table:

- ATLAS reports -

No major site issue to report. Security Service Challenge 4 ongoing - will be reported on when complete. INFN-T1 FTS upgrade completed.

Tier2s: US Functional Test Data Distribution is slow: under investigation but errors seem to be disappearing - transfers not starting to ILLINOISHEP_DATADISK, NET2, SLACXRD and WISC.

- CMS reports - 17 June 2010 (Tuesday)

T0 Highlights: Following machine schedule and taking data or not.

In between some test of new developments for the T0 processing infrastructure. Yesterday, tested switch to cmsprod account (running the T0) shadow afs volume. If there is a problem with the cmsprod account's afs volume, we can switch over to a shadow volume to continue processing. The switch over was tested with cmsprod2 account but caused degradation of CASTORCMS up to unavailability, GGUS ticket: https://gus.fzk.de/ws/ticket_info.php?ticket=59106.

Explanation: CASTOR processes got stuck (each for several minutes) on this home directory, which ended up with new jobs no longer being submitted into the scheduling process (hence the effect on the whole instance). Jan Iven reported this was an effect of the 30 minute 'human time scale' to invoke the switchover during which the cmsprod2 account was disabled. This afs dependency is being investigated and the problem has been taken offline.

Just noticed that the CERN CMS dashboard has stopped showing data since 7 am. A ticket will be raised.

T1 Highlights: Still recovering from WMS issue over the weekend and filled up unmerged data spaces especially at FNAL which is still 92% full. Still many thousands of jobs in an unclear state - being followed-up with CNAF.

T2 Highlights: MC production as usual

Other issues: Authentication to CRAB servers on vo boxes here at CERN failed under certain circumstances. Traced back to changes related to voms: works with lcg-voms.cern.ch, does not with voms.cern.ch. Needed lcg-vomscerts update which was announced long time ago on voboxes and done now through CMS VOC on all the voboxes so considered solved.

- ALICE reports [↗](#) - GENERAL INFORMATION: Pass 1 reconstruction activities ongoing at CERN. Together with four wide distributed analysis trains. No MC or raw data transfers activities for the moment.

T1 sites: FZK: proxy renewal mechanism of pps-vobox stopped at a certain moment in the last hours. As result the alien user proxy has expired stopping all local services. Intervention needed by local ALICE VObox support.

T2 sites: 1) Hiroshima-T2: Site has successfully migrated to the latest xrootd version. 2) Kolkata-T2: yesterday the site admin announced a disk full issue observed in the local VOBOX. Problem solved this morning.

- LHCb reports - Have just started a large (200 million events) MonteCarlo production.

T0 site issues: 1) grid_2nd_lhcb remove from the BDII. 2) CREAM CE upgrade to the latest release. 3) CE attached to SLC5 subcluster 4) LFC intermittent connection problem for analysis users - message is 'could not secure the connection'.

T1 site issues: CNAF FTS transfer (GGUS:59143) - what is the status please ? Ricardo reported the FTS servers are correct and they will update the ticket. He suspected there may be an out of date domain name server (secondary ?) record in use at CERN - they changed name 2 weeks ago and informed CERN. To be checked (Harry).

Sites / Services round table:

- CNAF: FTS upgrade to 2.2.4 done at INFN-T1 - their 3 FTS servers are now on the same subnet. They currently have a bdii issue - their server is not accepting connections - under investigation.
- RAL: WMS server lcgwms02 had a corrupted database and was down overnight - now back in production.
- NL-T1: Had a small problem with a storage node at SARA yesterday - it was down for 30 minutes for a hardware change.
- CERN Databases: CMS pvss replication is recovering from latency induced by CMS users running very large transactions - they have been contacted.
- CERN CEs: The 4 lcg-ce submitting to slc4 were drained overnight but found to have 4 new CMS jobs this morning (sent by direct submission) that had to be cancelled. For the upgraded CREAM-CE, ce201, LHCb jobs are going fine and they would like some ALICE jobs also. When the experiments are happy they will schedule upgrades of the other 2 CREAM-CE, hopefully next week.
- CERN CASTOR: Requesting LHCb when they would like the srm search order on space tokens (first is default) to be changed - a low risk operation. Joel requested next Monday when experts are back from the Amsterdam data management meeting.

AOB: (MariaDZ) A CERN Remedy PRMS migration is announced for today 17:30 (detailed info here). If any problem is seen tomorrow with GGUS tickets to the Tier0, please inform ggus-info@cernNOSPAMPLEASE.ch for workflow investigation.

Friday

Attendance: local(Harry(chair), Oliver, Roberto, Ulrich, JanI, Lola, Alessandro, Elena, Eva, Nilo);remote(Catalin(FNAL), Rolf(IN2P3), Onno(NL-T1), Gang(ASGC), Rob(OSG), Joel(LHCb), Gareth(RAL), Davide(CNAF), Cristian(NDGF)).

Experiments round table:

- ATLAS reports -

CERN-PROD locality LOST for some file: known issue, yesterday another GGUS:59188. Now closed as understood.

TRIUMF will be re-included in the INFN-T1 SiteServices instance since FTS @INFN-T1 has been upgraded.

ATLAS DistributedComputing issue: Automatic exclusion for site in downtime to be checked: INFN-T1 was blacklisted also if the downtime was 'not for ATLAS'. savannah bug to AGIS (ATLAS internal). CNAF are no longer blacklisted.

This should be a calm weekend - any data taken will not be exported.

- CMS reports -

T0 Highlights: 1) following machine schedule and taking data or not. 2) latest LHC schedule update, no collisions before Wednesday. 3) IT will replace old wns with new ones: T0 affected with 4 cmsrepack and 108 cmst0 boxes, not a problem because no data taking is ongoing anyway. 4) Disk server might have disappeared from the CMSCAFUSER pool, ticket: https://gus.fzk.de/ws/ticket_info.php?ticket=59190. Under investigation, answer: "Example file is on a machine with network trouble (which also explain the small spikes in the SLS capacity graph - the machine is jumping between "available" and "gone")"

T1 Highlights: 1) FNAL unmerged situation under control again. 2) WMS issues still not completely solved, two savannah tickets for tracking: <http://savannah.cern.ch/support/?115119> and <https://savannah.cern.ch/support/?115230>. 3) CNAF declared unexpected downtime due to storage issues <https://goc.gridops.org/downtime?id=78805471>. Downtime now over, CNAF is reactivating transfers after tests checked out.

T2 Highlights: MC production as usual, very nice performance

Other issues: DashBoard issue from yesterday traced back to update process locking the tables so no updates could be processed, after update process was done, entries were visible again.

- ALICE reports - GENERAL INFORMATION: Some distributed analysis trains. No MC activities for the moment.

T1 sites: FZK: issue reported yesterday about Alien user proxy has been solved.

- LHCb reports -

Experiment activities: 1) 200M MC production is launched. 2) RecoStripping-05 launched

T0 site issues: CERN : LFC-RO (could not secure connection (GGUS:59155) converted in an ALARM ticket (GGUS:59174). LFC-RO disappeared from the topology in the SAMDB, no SAM test results available any longer (GGUS:59193). This was due to an expired certificate on an LHCb machine which was rapidly fixed when the right people were alerted. LHCb would like a post-mortem on this incident.

T1 site issues: 1) CNAF GGUS:59038 Still problem transferring into CNAF from all T1s despite the site claims to be fixed. LHCb would also like a postmortem on this incident. Davide agreed to answer the ticket as soon as possible and open a post-mortem. 2) GOCDB problem, can not retrieve information : (GGUS:59172). Apparently gocdb switched to another instance. 3) Tested a ggus operator alarm ticket at RAL which correctly sent an SMS to the LHCb expert on call.

T2 sites issues: 1) USC-LCG2 GGUS:59189 (pilots aborting). 2) UKI-SCOTGRID-DURHAM GGUS:59163 shared area

Sites / Services round table:

WLCGDailyMeetingsWeek100614 < LCG < TWiki

- IN2P3: Last months site availability figures (now Nagios based) came in at 100% rather than the expected 94% or so. Under investigation.
- NL-T1: The dcache serving the SARA srm will be upgraded to version 1.9.5-19 next Monday. This includes a bug fix for the CRL issue they have had for the last few weeks.
- CNAF: Had a problem with their site bdii - have upgraded it to the latest version yesterday - no problems so far. The unexpected WMS downtime already reported by CMS was due to an overloaded WMS machine leading to jobs stuck for some days. Now adding more WMS to the CMS production pool. Oliver queried if CMS needed to change any configurations for this - will be checked with Daniele Bonacorsi.
- OSG: Two issues 1) ATLAS SAM tests are occassionally failing - not yet clear if this is a top level bdii problem or due to what the OSG sites are publishing. 2) Some ggus tickets have weak descriptions that we enhance in our ticketing but our changes do not get reflected back into ggus and updates come back with the weak description. Being followed with ggus.
- CERN Databases: Migrations to new hardware for the CMS integration database have been scheduled for next week.
- CERN dashboard: Update on the CMS Dashboard reporting failure reported yesterday: Oliver was referring to the distribution of the number of running jobs for the last 24 hours in the historical view of the job monitoring. There was a heavy process recalculating job monitoring statistics which put a lock on the table used for this distribution. It is back to normal now and apart from that all CMS Dashboard applications were working properly.

AOB:

-- JamieShiers - 08-Jun-2010

This topic: LCG > WLCGDailyMeetingsWeek100614

Topic revision: r13 - 2010-06-18 - HarryRenshall



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