

Table of Contents

Week of 100524.....	1
Daily WLCG Operations Call details.....	1
WLCG Service Incidents, Interventions and Availability, Change / Risk Assessments.....	1
General Information.....	1
Monday:.....	1
Tuesday:.....	1
Wednesday.....	4
Thursday.....	5
Friday.....	7

Week of 100524

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 Blogs	GgusInformation	Sharepoint site site - Cooldown Status Status - News News

Monday:

- No meeting - CERN closed!

Tuesday:

Attendance: local(Harry(chair), Akshat, Ricardo, Dirk, Oliver, Jean-Philippe, JanI, Lola, Eva, Patricia, Roberto, Maarten, MariaDZ, Jamie);remote(Jon(FNAL), Gang(ASGC), Vera(NDGF), Jeff(NL-T1), John+Brian(RAL), Rolf(IN2P3), Angela(KIT), Alessandro(INFN), Jeremy(GRIDPP), Rob(OSG)).

Experiments round table:

- ATLAS reports -
- CMS reports -

T0 Highlights: Software change in the T0, bumped all datasets except RAW to new processing versions because of event content compatibility. KIT, IN2P3, PIC and ASGC have not yet approved transfer requests.

T1 Highlights: 1) power outage at PIC on Friday/Saturday with very short recovery time (thanks). No measures like moving custodial assignments of primary datasets had to be taken. 2) Request for 50 Million MinBias events will be run on T1 level as T2 level is saturated with other MC requests, almost done. 3) Expect to start pre-production for another re-reconstruction pass soon, still waiting for software and conditions.

T2 Highlights: 1) MC production as usual. 2) Starting large scale pile-up simulations at T2 sites, expect jobs with higher than normal I/O load at sites.

Weekly-scope Operations plan [Data Ops]:

Tier-0: data taking.

Tier-1: Expected re-reconstruction request for all 2010 data using new software release including skimming, possibly also re-reconstruction of corresponding MC.

Tier-2: Large scale pile-up production at most T2s

Facilities Ops]

Final webtools services migration to SL5 occurred las Friday. During this week SL4 nodes will be switched off and deprecated.

VOC working to provide CRC-on-duty access to all CMS critical machines.

Note1: Sites to provide SL5 UI/voboxes for CMS to run PhEDEx soon. By end of June it will be mandatory, as PhEDEx_3_4_0 will be released in sl5-only release. We ask the sites to provide sl5 UI/voboxes according to the deadline.

Note2: CMS Vocard (CIC-portal) to be upgraded soon.

- ALICE reports [↗](#) -

GENERAL INFORMATION: High intensive MC production running during the weekend with peaks over 18K concurrent jobs. Good behavior in general of all Grid services at all sites.

Transfers: Raw data transfers going on during the last 3 days wityh 18TB transferred and an average speed around 55MB/s.

General Issue: during the last night (from 04:00 to 6:00) no ALICE jobs were registered by MonaLisa. In fact all sites are showing a glitch around the mentioned window time. The problem comes from one of the ALICE central machines which crashed during the weekend and as result no jobs were recorded by MonaLisa.

T0 site: Peaks over 3000 concurrent jobs, good behavior of the CREAM-CE and the LCG-CE resources

T1 sites

CCIN2P3: A different environment setup was found during the weekend in both ALICE VOBOXES available at this site. The result is that one of the vobox is perfectly working while the second one is out of production due to the wrong environment setup. Reported to the ALICE expert at the site, the experts are looking for possible differences in the environment setup at both VOBOXES

NIKHEF: The local service responsible of the software installation (PackMan) is failing at this site. Same issue found this morning in Cagliari. The problem has been reported to the AliEn experts before warning the site

T2 sites

Kolkata: During the weekend, the Alien sw was updated at this site. The local ALICE site admin prevented the ALICE core team about an obsolete Alien version running at this site.

Tuesday:

IPNL: CREAM-CE out of production (submission failing). GGUS-58478

- LHCb reports -

Experiment activities: Very intense data taking activity doubling the 2010 statistics. MC production ongoing.

T0 site issues:

On Monday the M-DST space was presenting many transfers queued and users report their jobs hanging (and then killed by the watch dog). SLS effectively showed this problem yesterday today recovered. We are also seeing continued degradation of service on service class where RAW data exists. Can we have some indication what and why this happened ? Jan Iven reported that the M-DST pool is too small and LHCb should request more resources be added to it. For the raw data pool looking at the example file and job he was seeing a normal file access response of 2 seconds and the job got killed in fact on cpu time-limit.

On Saturday afternoon we had default pool overloaded triggering an alarm on SLS; it recovered by itself.

T1 site issues:

IN2p3 Open a GGUS because ~10% of the jobs were failing yesterday with a shared area issues (GGUS 58283). Reproduced the problem. Any news from IN2p3 on the AFS shared software area ? Rolf replied they were testing a workaround for the afs cache problems that vey afternoon.

RAL: 1) request to increase the current 6 parallel transfer allowed in the FTS for the SARA-RAL channel in order to clean the current backlog draining too slowly. Details of this requirement wll be discussed offline. 2) lost a disk server (same one as last time). Files have been recovered.

CNAF: some FTS transfers seem to fail with the error below. CNAF people discovered a bug in Storm to clean up failed transfers at basis of this problem SOURCE error during TRANSFER_PREPARATION phase: [INVALID_PATH] Requested file is still in SRM_SPACE_AVAILABLE state!

PIC: Power failure on the week end causing the site being unavailable.

T2 sites issues: CSCS and PSNC both failing jobs with SW area issue

Sites / Services round table:

- NL-T1: Jeff queried a report that LHCb were changing cpu processing shares at NL-T1. Roberto explained this was an internal LHCb issue - NIKHEF and SARA have different roles for them and they are preparing to move analysis storage from dcache to dpm and did not propogate a matching change in cpu shares when they should have. Jeff asked them to let NL-T1 know if cpu servers need to be moved around.
- ASGC: Seeing intermiitent failures on CMS and test jobs since a few days. Have opened ticket 114622 in the CMS Savannah.
- NDGF: Performed a successful dcache upgrade today. Queried when the new FTS would be released. Maarten reported we are still waiting on test sites (including CERN) to report but Triumpf had already put it in production. RAL will be testing the checksum case sensitivity fix. Hopefully will come out this week - more news tomorrow.
- IN2P3: Not quite understanding the ALICE problem of different VObox requirements. Patricia has sent log files to R.Vernets but not opened a ticket yet as not sure of where the problem really lies.

Tuesday:

- KIT: Over the weekend a disk partition of one of their CEs filled up and job submission was disabled till 14.00 today. Filled by CMS and ATLAS pilot job logs but all less than 5 days old so not a failure of log rotation. Investigating why a normal number of jobs managed to generate so much logging.
- CERN CASTORATLAS: Harry queried an ATLAS operator alarm that morning on degradation of the TOMERGE pool. Jan Iven reported this was in fact provoked by a log daemon being stuck. SRMATLAS may also have been affected in the period from 04.00 to 10.00.
- CERN databases: Overnight a node of the ATLAS production database rebooted - being investigated. The DB stayed available but some sessions failed over to other instances. There was also a problem with one of the nodes of the CMS production DB running out of space in the file system.

AOB:

Wednesday

Attendance: local(Miguel, Harry(chair), Lola, JanI, Eduardo, Ricardo, Jean-Philippe, Flavia, Eva, Maarten, Steve, Simone, Akshat, MariaD, Pavel);remote(Ian(CMS), Jon(FNAL), Michel(BNL), Angela(KIT), Onno(NL-T1), Tiju(RAL), Rob(OSG), Gang(ASGC), IN2P3), Andrew(LHCb)).

Experiments round table:

- ATLAS reports - Lots of MC and reprocessing activities. Several ggus tickets for data export issues. ASGC needs to allow more FTS jobs - NL-T1 has been in an intervention - RAL problem has been closed. A monitoring database server at CERN is overloaded otherwise smooth running.
- CMS reports -

T0 Highlights: 1) Preparing for 900GeV running tonight. Potentially high rates. 2) Software change in the T0, bumped all datasets except RAW to new processing versions because of event content compatibility. All sites appear to have approved the requests.

T1 Highlights: 1) New CMSSW release working its way to sites. Full reprocessing of the data expected to follow shortly. 2) New ticket just opened at CNAF where some MonteCarlo files do not have the expected checksums.

T2 Highlights: 1) MC production as usual. 2) Starting large scale pile-up simulations at T2 sites, expect jobs with higher than normal I/O load at sites. 3) A couple of sites have seen problems updated CRL for DOE Grids. Unclear if it's a transient or regional DNS problem. Seemed to affect Spanish and Portugese Tier 2 first.

- ALICE reports [↗](#) -

GENERAL INFORMATION: Decrease in the number of running jobs due to the end of the MC cycles started during the last weekend. In addition, the usual reconstruction and analysis activities are going on.

T0 site: Last Friday we reported about the creation of a new LanDB set including about 15 ALICE CAF nodes which required a common an specific connectivity. The name of the mentioned set had to be modified in order to follow the standard name definition set by the PES experts. As procedure to follow it was agreed to define a new set this time with the right name that as soon as it appeared properly populated with the names of the nodes, it would replace the current one. Still it appears to be empty in the network page, PES experts have been contacted.

T1 sites:

CCIN2P3: issue reported yesterday concerning different environments found in the two ALICE VOBOXES. The problem requires further and deeper investigations by the ALICE side in order to understand whether it is a problem associated to the ALICE environment.

NIKHEF: Issue reported yesterday concerning the bad behavior of the local PackMan service: Solved. The VOBOX required the update of the local AliEn version. Same procedure applied to Cagliari.

T2 site: IPNL GGUS ticket reported yesterday solved (local CREAM-CE required the restart of tomcat). Services restarted at the local VOBOX.

- LHCb reports - 26th May 2010 (Wednesday)

Experiment activities: 1) Reconstruction of recent data ongoing at T0/1s. 2) MC production ongoing at T2s.

T0 site issues: Ticket against CASTOR closed. Not a CASTOR problem. Some shared software area problems currently appearing.

T1 site issues: 1) PIC: PIC-USER space token is full. 2) NL-T1: SARA dCache is banned due to ongoing maintenance.

Sites / Services round table:

- NL-T1: The observations of ATLAS and LHCb are because we are migrating from 12 dcache pool nodes to a new 12 trying a new procedure to keep the service up. This required a dcache reconfiguration and restarts which caused some transfer failures but no more since this morning. The whole operation will take a few days and it was agreed to document the new procedure for other dcache sites.
- IN2P3: Will have a scheduled down on 8th June to move hpss and gpfs services so there will be no tape access. Also afs servers will be affected so software releases (i.e. of new afs volumes) will not be possible. A transparent Oracle intervention will also be made.
- CERN FTS: Good progress on the new version. Up at CERN in a pilot but have not exercised data transfers yet - experiments are also encouraged to exercise the pilot.
- INFN: Successfully completed Oracle and batch system upgrades.
- CERN GPN: CERN external firewall was overloaded from about 17.00 yesterday with xrootd traffic to Tier 2 sites causing packet losses especially with udp packets. This traffic was diverted to the HTAR route from 10.00 today. Maarten added this had hurt exporting of the CERN top level bdi and that 3/4 of the Tier-2 sites were failing the lcg replication test (local file to reference CERN SE). The availability statistics will be corrected and we will also look at moving bdi export to HTAR. It was understood later this traffic was from the new ALICE CAF nodes.
- CERN CASTORATLAS: An ATLAS groupdisk server has some inaccessible files and needs a file system repair.

AOB: An LHC technical stop is scheduled from 31 May to 2 June inclusive. To be confirmed.

Thursday

Attendance: local(Harry(chair), MariaG, Gavin, Steve, Maarten, Jamie, JanI, Nilo, Eva, Jean-Philippe, Roberto, Simone, Stephan, Flavia, Patricia, MariaDZ);remote(Jon(FNAL), Michael(BNL), Gang(ASGC), Rolf(IN2P3), Ian(CMS), John(RAL), Angela(KIT), Alessandro(CNAF), Rob(OSG), Tristan(NL-T1)).

Experiments round table:

- ATLAS reports - Last reprocessing period has finished and data is now being replicated over the grid at aggregate rates of up to 7 GB/sec. Testing a new method of determining which datasets should be where to avoid filling sites. For next weeks technical stop ATLAS will export raw data till Monday morning then ESD and AOD over Tuesday and Wednesday but sites with a scheduled downtime should continue as ATLAS will catch up later. Stephan queried if the CASTOR checksum upgrade could be rolled out at CERN. Jan Iven reported this release is not rolled out yet and also after the recent data loss any such changes are stringently reviewed.
- CMS reports -

T0 Highlights: Preparing for 900GeV running on short notice Potentially high rates

T1 Highlights: 1) Full reprocessing of the data has now been requested so loads will now go up. 2) New and Open Tier-1 tickets involve transfers to Tier-1 or Data Consistency.

T2 Highlights: 1) MC production as usual. 2) Starting large scale pile-up simulations at T2 sites, expect jobs with higher than normal I/O load at sites.

- ALICE reports [↗](#) -

T0 site: Setup of a new LanDB set for ALICE. This set will have specific HTAR exceptions (already defined by the security team) in order to avoid high loads in the outer firewall. The procedure has been the following: A LanDB firewall set was defined by the end of the last week with a "free" name definition. Associated to this LanDB set, the security team applied the agreed HTAR exceptions. Once the LanDB set was defined, the PES team warned us about the "free" name definition of the LanDB set that we had just defined. IT has specific rules concerning these names and also the responsible assignement. The strategy agreed at that moment was the following: A new and empty LanDB would be defined, this time with the proper name and responsible assignment. Once the new set was properly defined, the security team would be asked to apply the HTAR exceptions to this new set. At this moment the previous set could be deprecated. This morning still the new LanDB was not populated and PES found the reason: The CDB teampltes of the nodes had not been accordingly modified to include the new set.

Concerning the instabilities observed in the past with an important number of CAF nodes, yesterday the experiment representative had a meeting with the IT experts to establish a strategy for this problem. The conclusions were the following: All the affected nodes will be downgraded to the previous kernel version (done) and put back in production (by today) If problems still persist, these set of nodes (which have 4 disks) will be replaced by 3-disk nodes.

- LHCb reports - Experiment activities: Reconstruction of recent data ongoing at T0/1s. MC production ongoing at T2s.

T0 site issues: We are seeing many queued transfers on CASTOR LHCbMDST (GGUS 58523). See attached plots showing active and queued transfers. We are investigating if we can slightly modify how the LHCb applications handle open files, but we would also like new hardware as fast as possible. A request has been made to CASTOR.

T1 site issues: 1) IN2p3: investigating actively on the shared area shortage last week. 2) IN2p3: the CREAM CE ccreamceli01.in2p3.fr has currently 7386 Scheduled pilots while the BDII publish ~4500 free slots. This impacts the ranking mechanism attracting wrongly jobs there. GGUS ticket to be submitted.

T2 sites issues: SharedArea problem at epgr04.ph.bham.ac.uk UKI-SOUTHGRID-BHAM-HEP

Sites / Services round table:

- BNL: Planning two changes during the LHC technical stop. 1) Code changes on Force10 network switches on Tuesday. 2) Upgrade of Condor clients following the recent (transparent) server side upgrade. This requires clients to be drained from Monday evening (Monday is a US holiday) for a four hour upgrade on Tuesday morning expected to be finished by noon.
- RAL: Had disk failures on two scratchdisk servers this morning. Taken out of production while rebuilding - expected back on Saturday earliest. Monday is also a UK holiday and there will be an at-risk during UPS tests. Some planned networking changes have been postponed.
- CERN FTS endpoint on <https://fts-patch4084.cern.ch:8443/glite-data-transfer-fts/services/...> is now available with 2.2.4. CMS have already done some basic testing. Ready when ever to upgrade the production T0 and T2 service. PATCH:4084. Simone gave ATLAS plans to move their functional tests to the new endpoint tomorrow (the machine concerned is currently being upgraded to slc5) and will then participate in a decision early next week if and when to upgrade the production FTS.
- CERN CASTOR: Changing a bad network switch affecting CASTORALICE disk servers is being scheduled for some time Monday during the technical stop - expected to take 1 hour.

AOB (Jamie): The WLCG does not want to have more than one-third of an experiments resources in overlapping scheduled downtimes so please report your downtime plans for next weeks LHC technical stop here for any coordination.

AOB (MariaDZ): A round of the Periodic ALARM tests full chain will take place next week Mon-Wed. The **Tier0 ONLY** is concerned this time. Steps to follow are in <https://savannah.cern.ch/support/?114705> and services to test are in <https://twiki.cern.ch/twiki/bin/view/LCG/WLCGCriticalServices>

AOB (Harry): The technical stop from 31 May to 2 June is confirmed.

Friday

Attendance: local(Harry(chair), Simone, Jean-Philippe, Stephan, Maarten, Lola, JanI, Eva, Nilo, Patricia, Ian, Akshat, JohnG, Pavel, Ricardo);remote(Jon(FNAL), Rolf(In2P3), Xavier(KIT), Gonzalo(PIC), Michael(BNL), Alexander(NL-T1), Roberto(LHCb), Rob(OSG), John+Brian(RAL), Jeremy(GridPP), Tore(NDGF), Gang(ASGC)).

Experiments round table:

- ATLAS reports -

Data replication of reprocessed datasets going on - should finish over the weekend. Many sites are almost full (INFN-T1, PIC is critical if data taking is efficient over the week-end). ATLAS ADC Ops will trigger deletions when possible. 1 problem discussed by mail : Transfers from FZK were failing. Seems to be solved after FZK intervention. Gonzalo queried which space tokens at PIC are full - answer was MCDISK and DATADISK. ATLAS will stop transfers till they have cleaned enough disk.

FTS in PPS is used for Functional Tests CERN-> T1 since this morning. No obvious errors so far.

FTS channel ATLAS T1 -> Tokyo setup for tests by IN2P3 (ATLAS T1 is the sum of all ATLAS T1s except IN2P3). Brian queried how many concurrent transfers - thought to be 10 but ckeck with L.Schwarz at IN2P3.

- CMS reports -

T0 Highlights: Preparing for stable running over the weekend. Ian queried status of failed Oracle RAC nodes - Eva reported one still down in CMS offline. Currently not a problem but would be a worry if load increases.

T1 Highlights: Data Reprocessing has started. Also a few validation samples. Backfill activities are being ramped down to avoid interfering with production work. New and Open Tier-1 tickets involve transfers to Tier-1 or Data Consistency.

T2 Highlights: 1) MC production as usual. 2) Starting large scale pile-up simulations at T2 sites, 21 sites and counting. expect jobs with higher than normal I/O load at sites.

- ALICE reports [↗](#) -

T0 site

One of the ALICE users reports today that his simple analysis jobs fail to compile because of some missing header files: (stdarg.h, limits.h, typeinfo, vector, string, iostream, float.h, etc). It seems the c++ compiler is not correctly installed in the WNs: lxbsu1243, lxbsq0621, lxbsu0818, lxbsu1418. The same analysis tasks compile successfully elsewhere. GGUS-58587

lanDB set defined for ALICE and reported yesterday still appears to be empty. Experts confirm the nodes appear on the server side, but still nothing is found nor updated to the lanDB set. PES experts following it. Ricardo reported this is not in fact causing any operational problems.

CASTOR news: 980 TB in 46 disk servers have been received for CASTORALICE. The allocation per pool proposed by ALICE has been the following: 750TB for ALICEDISK and 190 TB in T0ALICE

Intervention: there is a flakey router affecting 11 disk servers in castoralice/alicedisk and 2 disk servers in t0alice with network errors and disconnects. The network people asked when could an intervention be scheduled to sort out this problem. Online and Offline representative have agreed with the network team to have this operation the next Monday 31st of May at 9:30 (it should take around 1h)

T1 sites: T1 sites in production. During yesterday's ALICE TF meeting it was agreed that sites should migrate the local CREAM-CE services to CREAM1.6/gLite3.2 as soon as possible. We will contact individually with each T1 site to have their plans for this update

T2 sites

Kolkata: user proxy delegation procedure was not working at the site after updating to CREAM1.6. Issue has been found: the repository chosen by the site to download the corresponding rpms was not the glite standard repository. Issue solved by the site admin. Maarten explained further - there is a Redhat 'development' repository called epel which contains some uncertified gLite software in particular a trust manager and this was in the installation path. In discussion with GT to clarify their installation documentation on this point.

Clermont: CREAM local system is showing authentication issues. GGUS-58590

PNPI: Restart of all services this morning to bring the site back in production

Cagliari: Local CREAM-CE not performing well. Connections refused at submission time. GGU-58591

- LHCb reports -

Experiment activities: Reconstruction of recent data ongoing at T0/1s.

T0 site issues: In the last days have been observed many queued transfers on CASTOR LHCbMDST causing

a large fraction of user jobs to fail at CERN as well as many transfers of DST outgoing to T1's (GGUS 58523). As a temporary measure CERN has been banned yesterday to accept further user jobs and as soon as the backlog has been drained jobs re-started to succeed and transfers of DST started to succeed again (see picture). There are available 490 TB of new h/w delivered. Ignacio started to install as an emergency measure - 4 new diskserver on lhcbmdst (the loaded service class root of these problems) and today we do expect them in production (once few issue with the configuration will be sorted out). Also asked to disable the GC on this SC (it is T1D1 and not T1D0). Jan Iven reported these boxes require a new cdb hardware profile but he hoped to have them deployed by the end of today.

T1 site issues: 1) IN2p3: huge amount of queued jobs and very few running. This is throttling one of the reprocessing activity backlog to be drained. Two-fold problem: 1 the very-long queue recently open to lhcb accepts too few concurrently running jobs. 2. the creamce is publishing wrong information and attracts jobs despite the lhcb rank would prevent so. (GGUS 58572). Maarten reported this looks like a known bug between WMS and CREAM where WMS reports the wrong job status. A new version of the WMS is in preparation hopefully ready by the end of next week. Roberto confirmed this explains the apparent long queues the Dirac system thinks are at IN2P3 - as a work-around they have temporarily stopped using the CREAM-CE at IN2P3. 2) KIT: many jobs consuming no cpu - probably a shared software area issue.

Sites / Services round table:

- FNAL: The problem of a CERN bdii server publishing out of date information (and DNS balancing is round robin) that happened a few months ago has recurred. Ricardo reported that preparing a sensor for this is in their work list. Rob wondered if it was picked up in the OSG monitoring review.
- PIC: SIR on last weeks power problems is in preparation.
- BNL: The interventions scheduled for Tuesday will now include reorganising network connections to worker nodes but within the 8 am to noon time window.
- RAL: One of the two failed scratch diskserver has been returned to production, the second one will not be before midday tomorrow. ATLAS software server will be upgraded Monday evening - should be transparent. At-risk Tuesday with UPS tests. NB - Monday is UK holiday. Brian reported they had successfully tested the case sensitive checksum fix in FTS 2.2.4. Investigating backlog of transfers TRIUMF to RAL and suggest RAL increase the number of concurrent transfers.
- OSG: No attendance Monday due to US holiday.
- ASGC: Had Maradona errors traced to a bad worker node (a restart had failed to mount a file system). This was fixed but they are now getting more errors over different worker nodes - under investigation. Maarten will also investigate.
- CERN CASTORATLAS: The ATLAS groupdisk that went down on Tuesday is back with its files available.
- CERN Databases: Between Monday and Tuesday they will patch all the production databases. For CMS offline their failed node is still down though the MoU says 12 hours - it should be done on Monday. They are trying to restart the node without its failed memory but this is not sure to work. There is enough remaining capacity to run the database.

AOB:

-- JamieShiers - 20-May-2010

This topic: LCG > WLCGDailyMeetingsWeek100524

Topic revision: r13 - 2010-06-02 - BrianDavies



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](#)