

# Table of Contents

<b>Week of 100503.....</b>	<b>1</b>
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.....	9
Friday.....	11

# Week of 100503

## 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
<a href="#">ALICE</a>	<a href="#">ATLAS</a>	<a href="#">CMS</a>	<a href="#">LHCb</a>	WLCG Service Incident Reports	Broadcast <a href="#">archive</a>	CASTOR Change Assessments

## General Information

General Information				GGUS Information	LHC Machine Information
CERN IT status <a href="#">board</a>	M/W PPSCoordinationWorkLog	WLCG Baseline Versions	WLCG Blogs <a href="#">Blogs</a>	GgusInformation	Sharepoint site <a href="#">site</a> - Cooldown Status <a href="#">Status</a> - News <a href="#">News</a>

## Monday:

Attendance: local(Ueda, Harry(chair), Nicolo, Lola, Alessandro, Eva, Nilo, JanI, Timur, Akshat, Maarten, Roberto, Carles, Dirk, Manuel);remote(Jon(FNAL), Michael(BNL), Tore(NDGF), Davide(CNAF), Ron (NL-T1), Rolf(IN2P3), Gang(ASGC), Rob(OSG), Angela(KIT), Gonzalo(PIC)).

Experiments round table:

- ATLAS reports - presented in chronological order:

PIC blacklisted as source due to the disk pool lost: re-included this morning

RAL-LCG2 data transfer problem: [GGUS:57882](#), diskserver lost. Blacklisted on sunday. Waiting for more information from them.

NDGF report lost files at source when transfer fails: similar to TRIUMF <https://savannah.cern.ch/bugs/?66719> - suggest they look at this.

CERN-PROD data transfer problem: [GGUS:57872](#); ATLASDATADISK was full, error message misleading (was reporting jobs had timed out - fixed), ATLAS blacklisted the site as destination. Will reopen it when ATLAS have cleared enough space.

SARA-MATRIX storage unstable [GGUS:57879](#). Problem understood? Site still investigating, internal network problem one of the possible causes.

INFN-T1 GGUS:57892 [open](#) - not possible yet to finish the 3 repro jobs still missing. Causes to be understood.

Many o(10) FTS jobs in Active state since >2days, on several FTS servers. Yesterday and this morning we did some cleanup, sites please fix these problems! These are known issues so ATLAS did not submit tickets. JanI will discuss offline with Alessandro a CERN case.

Today Downtime: no outage downtimes foreseen for Tier1s

- CMS reports - T0 Highlights: CE SAM tests expiring April 29-30 - issues in SAM DB publication.

T1 Highlights: Running MC reprocessing and skimming

CNAF: 1) some file access/stageout failures. 2) SAM SRMv2 failures and transfer failures from all sites on May 3rd, GGUS #57913

KIT: 1) SRM SAM test errors on May 2nd. 2) waiting for a list of files to be deleted from some disk pools that are almost full.

ASGC: migrations to tape not progressing.

T2 Highlights: Running MC production and pre-production

T2\_AT\_Vienna - low number of jobs running, site admins notified.

Report on new WMS version: No job failures due to WMS. Some jobs stuck in "Submitted" state, can be detected checking logging-info. Cannot submit to ARC CE. Maarten Litmaath explained what had happened here. Last Thursday evening, though not understood till Friday evening, the CERN CMS WMS hit a hard-coded limit in the WMS Condor back-end of 100 different concurrent users with unfinished jobs, making the back-end continually exit and restart. The limit is removed in a later version and this was installed on Sunday morning, but with hacks to allow multiple library paths for other components, and appeared to work, except jobs that completed successfully were marked as still running. Another WMS was already being used with the updated Condor back-end to stress-test the fix for another issue: it does not have the problem with completed jobs remaining marked as running, and it can also submit to ARC again. It appears we will have to install the same rpms also on the other nodes.

T2\_RU\_PNPI transfers to T1\_CH\_CERN failing for timeout, will forward request to adjust timeouts on the dedicated channel on fts-t2-service.cern.ch Savannah #sr 114126

Squid issues at T2\_RU\_IHEP

SW area issues at T2\_UK\_London\_IC, CMSSW reinstalled.

Other: JobRobot submission issues during weekend (WMS problem, and agents down).

- ALICE reports [open](#) - GENERAL INFORMATION: Pass1 reconstruction at the T0 site together with a MiniPass ITS-TPC tracking reconstruction. In addition, there is one analysis train also running

T0 site: Good production state since the last night when several fixes were put in place for the AFS volume replication. Both backend are currently in production The issue reported on Friday concerning the CAF nodes of ALICE should be solved by today. A discrepancy between the information provided by the nodes and the CDB information was causing some expectations. Thanks a lot to the PES group for their prompt answer and action solving the issue. Information passed to ALICE, waiting for the experiment feedback

Monday:

T1 sites: RAL T1 back in production after the drain of the system to add the pilot accounts for the rest of the LHC VOs

T2 sites Cyfronet: host certificate of the VOBOX expired. The site admins announced the renewal of this certificate

- LHCb reports - Experiment activities:

Started the 450GeV data reconstruction this morning (FULL & EXPRESS/ magup & magdown). This evening back to 3.5TeV. Due to an DIRAC agent to be restarted, a huge backlog of data transfers from NL-T1 to pic has to be drained. About 10K files to be copied to pic.

T0 site issues: LFC:RO still overloaded by some users using old version of the framework (w/o Persistency patch) and/or not using the workaround put in place in DIRAC. Maarten suggested to blacklist users who fail to use the workaround but LHCb are not yet thinking of this - such users are reminded of the simple change they need to use the workaround. Meanwhile remaining LHCb applications are being migrated over the next months to use the new framework.

T1 site issues: 1) NL-T1: still having the problem with the SE affecting also reconstruction jobs attempting to upload output (GGUS: 57812). 2) IN2p3: going to upgrade their MSS to put tape protection and only allowing lhcb production role to stage files in. Access of data staged on disk will occur through dcap protocol w/o any authorization.

T2 sites issues: 1) UK T2 continuing the investigation on data upload issue spawned time ago by LHCb. 2) UKI-LT2-IC-HEP huge amount of jobs failing there

Sites / Services round table:

- BNL: The network problems reported last Friday were resolved by that evening.
- CNAF: 1) ATLAS storm problems have now been resolved. 2) Some ATLAS jobs have been hitting our 4 GB vmem limit and being killed by LSF - seems to be at the ATLAS end of job processing. For the moment we have removed the limit and will discuss with LSF how to allow jobs to have a high vmem for a short time. Alessandro thanked for this temporary solution and reported that similar jobs are run at PIC and ATLAS will do some more investigating. 3) For CMS storm is still down - under investigation. 4) Post Mortem of last weeks storm bug incident (provoked by a full ATLAS disk) being prepared. 5) A new TSM client was found to have a serious bug and had to be rolled back.
- NL-T1: Recent dcache transfers hanging are now thought to be a network connectivity problem in their dcache nodes - trying to recover asap.
- IN2P3: The SIR of the April 24th local batch system outage has been submitted. One more SIR is in preparation.
- ASGC: CMS tape migration problem has been solved - was correlation to tape system degradation seen on 1 May. Backlog now clearing.
- KIT: Problem upgrading batch client on a CE led to having to kill 5 ATLAS jobs (3 pilot and two production).
- PIC: Last weeks problem of an inaccessible ATLAS disk pool was solved on Sunday and the pool brought back. A Solaris ZFS file system could not be mounted due to an OS problem. Some follow-on network driver performance issues being looked at.

- OSG: Over the weekend successful retransmitted some failed SAM data. Meanwhile the ticket queue is empty.

AOB:

## Tuesday:

Attendance: local(Ueda, Harry(chair), ManuelG, Lola, Alessandro, Simone, Stephane, Nicolo, Eva, Akshat, Maarten, JanI, Roberto, MariaD);remote(Jeff(NL-T1), Davide+Daniele(CNAF), Gareth+Brian(RAL), Pepe(CMS), Rob(OSG), Gonzalo(PIC), Rolf(IN2P3), Gang(ASGC), Michael(BNL), Angela(KIT), Jon(FNAL)).

Experiments round table:

- ATLAS reports -

FTS has major problems: 1) Source files deletion after failed transfers. This is a bug related to source and destination having the same file base name and means some pairs of sites (running the same mass storage software) cannot directly transfer e.g. TRIUMF to CNAF which is very painful as TRIUMF has needed data that they cannot risk being deleted. 2) Many FTS jobs in Active state for long time on several FTS servers. We did some cleanup again. 3) Can we have a patch while waiting for a fix ? ATLAS have heard these will be fixed in a next FTS release but these problems are now critical and they want an urgent fix that does not add or change functionality. This was endorsed by J.Templon at NL-T1. Harry and Maarten will contact the FTS developers and report back (initial soundings are promising).

RAL-LCG2 disk server problem GGUS:57882 [↗](#) -- enabled transfers from RAL with a less frequent retry policy.

NDGF lost files <https://savannah.cern.ch/bugs/?66719> [↗](#) -- no news

CERN-PROD data transfer to ATLASDATADISK restarted, checksum for transfers to ATLASSCRATCHDISK has been disabled.

SARA-MATRIX storage unstable GGUS:57879 [↗](#) 'solved' -- to be verified.

INFN-T1 GGUS:57892 [↗](#) -- one of the three jobs finished.

IN2P3: 1) CC data distribution to T2s are slow GGUS:57931 [↗](#) (<https://savannah.cern.ch/bugs/?66897> [↗](#) for FR cloud) please have a look. Simone added that IN2P3 have modified their dcache pools and got a 3 times speedup but this is still not enough to clear backlogs, notably to Tokyo. One issue is the large number (800K per day) of small (50MB) analysis files. As luminosity and trigger quality improve such files should get bigger and fewer. 2) MCDISK -- the space seems to have enough free space from monitoring, but transfers fail with NO\_SPACE\_LEFT (GGUS:57632 [↗](#)). Maarten remarked this error can also happen if log space is full.

- CMS reports -

T1 Highlights: 1) Running tails of MC reprocessing - some long running jobs. 2) Running skimming preproduction. 3) CNAF: unscheduled StoRM downtime: SAM, transfer, merge job failures. 4) KIT: cleaned up temp files from disk pools. CMS Data Operations to implement automatic cleanup.

T2 Highlights: 1) Running MC production and pre-production. 2) T2\_RU\_ITEP, T2\_RU\_RRC\_KI - unscheduled downtimes, excluded. 3) T2\_RU\_JINR - job aborts due to site disappearing from BDII. 4) Upload of MC from T2\_RU\_SINP to T1\_US\_FNAL failing for permission errors. 5) T2\_RU\_PNPI transfers

Tuesday:

to T1\_CH\_CERN failing for timeout, will forward request to adjust timeouts on the dedicated channel on fts-t2-service.cern.ch Savannah #sr 114126. 6) SAM errors at T2\_BE\_IIHE - 'Best pool too high' on dCache (disk full?)

Weekly-scope Operations plan:

Data Ops: Tier-0: data taking and processing. Tier-1: finishing Spring10 re-digi/rereco campaign and starting train reconstruction model on Thursday. Tier-2: MC production of new requests .

Facilities Ops:

Improving procedures and documentation for Computing Run Coordinator this week. Exceptionally, a remote CRC is scheduled for this week. We will check the feasibility of having remote CRC's after this week.

CMS VOC started a campaign to improve the documentation on all Critical Services, together with all the service responsables.

Encouraging 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.

CMS Vocard (CIC-portal) was discussed last week at the Computing and Offline Week. Some memory issues still to be discussed with Offline this week, and it will be upgraded soon.

Note: As 13th-May is CERN holidays (Ascension day), we will have bi-weekly T2 support meeting by Wednesday 12th-May.

- ALICE reports [↗](#) - GENERAL INFORMATION: Two new MC cycles have been just started during the last night. In addition analysis train activities and pass1 reconstruction tasks at the T0 are ongoing.

T0 site: New ITCM ticket (00000000356681) has been submitted to announce that several ALICE CAF nodes are down. The list of nodes to restart are: lxbseq1415, lxfssl3311, lxfssl3313, lxfssl3315, lxfssl3317, lxfssl3318, lxfssl3401, lxfssl3403, lxfssl3404.

T1 sites: NIKHEF: Due to the fact that the local vobox provide access to exclusively 2 persons (who do not perform the automatic update of the vobox when a new version of Alien is available), a manual update of the local vobox was required this morning. (Jeff Templon reminded that this was an old problem related to the multiple mapping to the alicesgm account. ALICE to revisit this.) No issues to report for the rest of the sites.

T2 sites: 1) Clermont: Manual activities perform this morning at the vobox to bring the site back in production. 2) Birmingham: GGUS\_57847: The ticket has not been touched by the sitr admin since the 30th of April when he announced that he would look into the problem (local CREAM-CE not working). (MariaD reminded that the ggus ticket interface includes a button to escalate the issue which, in the simplest case, sends a reminder to the site).

- LHCb reports - Experiment activities: No beam no data taken since yesterday morning. Reconstruction of 450 GeV almost finished. Put in place a workaround to prevent the watchdog to kill jobs in the phase of uploading the large amount of small (streamed) reconstructed data that was causing in the last days a big fraction of jobs killed (CPU/wall clock time below the threshold).

T0 site issues: 72 files definitely lost in one of the disk server of the lhcbdata pool. List of files given to lhcb and a post-mortem already cooked up. Our data manager excludes however these data can belong to users or were effectively written in the lhcbdata pool because of the protections put in place there since before the

Tuesday:

22nd of March. It should then be closely understood how these data were perceived to sit in the lhcbdata pool and written the 22nd of March as reported as this should not have been possible.

T1 site issues: NL-T1: problem with SE (GGUS: 57812) seems to have been understood to a variety of reasons (overload SRM, network, hardware and other)

Sites / Services round table:

- CNAF: 1) Re the ATLAS jobs running out of memory queried what other sites do. Gonzalo confirmed PIC do not limit vmem. CNAF started with 2GB physical memory then ramped up to 3 GB then up to 4GB on vmem, combined physical+swap and this is now being reached, they thought in a transient stage during end of job processing. It was agreed this needs to be investigated by ATLAS. 2) Post Mortem on STORM/TSM problems is being finalised. 3) CMS problems are now solved. 4) ATLAS must clean up space on their MCDISK space token. 4) Last night a new problem was found in the WMS 3.2.14 release (which does not have the 100 user limit) affecting CMS analysis jobs. Reported to the developers but in early stage of understanding.
- RAL (Brian): ATLAS are asking our T2 to clear space in MCDISK - is this to completely empty or to return to consistency. Reply from Stephane was to completely empty RAL-PP and Birmingham but they needed to check further at QMC (London). Brian then asked ATLAS to turn off checksumming in export from T1 MCTAPE to T2 sites as there was another bug in FTS where checksumming is case sensitive and fails in these transfers. Stephane agreed but this should also be fixed in FTS.
- RAL(Gareth): Their failed diskserver has a faulty card and they cannot yet give a repair estimate. They were surprised to be blacklisted by ATLAS at the weekend due to this single diskserver failure. Ueda agreed this was not usual but there had been a huge number of errors and they had no time to discuss this internally at the weekend.
- IN2P3: Would like the explanations of the problem of slow transfers to their T2 (as given above) to be reflected in the ticket so the right people could be kept up to date.
- NL-T1: The dcache access problems at the weekend were due to a faulty port on a 4-port switch, now repaired. 5 May is a Dutch national holiday so NL-T1 will not attend the daily meeting.
- OSG: The report yesterday of no tickets was wrong - there is at least one in ggus but the search is not finding it. To be followed up.
- CERN WMS: The full upgrade of one of the CMS WMS seems to be ok and the other 3 will be done this week. Thanks to CNAF for their information on the new problem - too early to say if it will be common or a special case.

AOB: Note that 13 and 14 May are CERN holidays.

## Wednesday

Attendance: local(Ueda, Harry(chair), Eva, Nilo, Zsolt, Patricia, Akshat, JanI, DavidG, Nicolo, Alessandro, Flavia, ManuelG);remote(Jon(FNAL), Michael(BNL), Joel(LHCb), Davide(CNAF), Gang(ASGC), Rob(OSG), Gonzalo(PIC), Tore(NDGF), Rolf(IN2P3), Gareth(RAL), Andreas(KIT)).

Follow up from yesterday FTS bugs critical issue for ATLAS: email exchanges are shown at the bottom of today's report and Zsolt confirmed that the 3 critical bugs have been fixed and are in the repository which has been frozen as release 2.2.4 after removing unfinished features that will now come in 2.2.5. By the end of the week certification should be finished and new rpms ready. Maarten reminded that certification must be done carefully and he suggested a few of the most affected sites, such as CERN, try this new version first.

Experiments round table:

- ATLAS reports -

FTS -- we received the plan for FTS 2.2.4. We think it is good. Thanks for the effort.

INFN-T1 storm-fe.cr.cnaf.infn.it seems to be down (GGUS:57975 [↗](#))

RAL-LCG2: 1) disk server problem GGUS:57882 [↗](#) -- any news? There is an email report at the bottom of today's report and Gareth confirmed there is no repair estimate yet. ATLAS may look at other solutions for distributing data to UK T2s if there will be a long delay. 2) checksum is disabled for transfers to atlasproddisk in UK T2s following the request yesterday (FTS checksum issue)

NDGF lost files <https://savannah.cern.ch/bugs/?66719> [↗](#) -- any news? Tore requested details in a ggus ticket - ATLAS will create.

SARA-MATRIX storage unstable GGUS:57879 [↗](#) 'solved' -- verified (>90% success for production jobs)

IN2P3-CC data distribution to T2s slow GGUS:57931 [↗](#) (<https://savannah.cern.ch/bugs/?66897> [↗](#) for FR cloud) The number of connections for the export pool has been doubled but there is still a huge backlog. Now ATLAS-France should consider the number of concurrent files in each FTS channel to T2s. In order to see the effect the value for the channel IN2P3-TOKYO is increased from 20 to 30. Let us know if this causes any problem.

INFN-T1 the memory info in the log of the 'finished' problematic jobPy:

```
PerfMonSvc      INFO <vmem>:  (    999.996 +/-      0.000 ) MB
Py:PerfMonSvc  INFO <vmem>:  (   1572.972 +/-     1.089 ) MB
Py:PerfMonSvc  INFO <vmem>:  (   1679.324 +/-     0.000 ) MB
Py:PerfMonSvc  INFO <vmem>:  (    732.598 +/-     0.000 ) MB
Py:PerfMonSvc  INFO <vmem>:  (   1118.802 +/-     0.324 ) MB
Py:PerfMonSvc  INFO <vmem>:  (   1158.582 +/-     0.000 ) MB
Py:PerfMonSvc  INFO <vmem>:  (    698.312 +/-     0.000 ) MB
Py:PerfMonSvc  INFO <vmem>:  (    767.703 +/-     0.182 ) MB
Py:PerfMonSvc  INFO <vmem>:  (    802.273 +/-     0.000 ) MB
```

Note, this does not show possible peaks.

- CMS reports - T0 Highlights: High load on CMSCAFUSER node, activity went down by itself.

T1 Highlights: 1) Running tails of MC reprocessing - some long running jobs. 2) Skimming running - 1 job failed at FNAL hitting 30 GB limit on WN disk Elog #2208. 3) FNAL and other OSG sites disappeared from WLCG BDII for issue in propagation from OSG BDII, fixed. 4) CNAF back from unscheduled StoRM downtime - please close GGUS TEAM #57913 if issue solved. 5) 3 jobs stuck on CERN-RAL channel on FTS-T0-EXPORT, canceled by site contact.

T2 Highlights: 1) Running MC production. 2) T2\_CH\_CSCS, T2\_BE\_UCL - scheduled downtimes, excluded. 3) T2\_AT\_Vienna - still low number of batch slots for CMS jobs, site admins investigating. 4) Upload of MC from T2\_RU\_SINP to T1\_US\_FNAL failing for permission errors. 5) T2\_RU\_PNPI transfers to T1\_CH\_CERN failing for timeout, forwarded request to adjust parameters on the dedicated channel on [fts-t2-service.cern.ch](https://fts-t2-service.cern.ch) Savannah #sr 114126, Remedy #681529. 6) SAM CE errors at T2\_RU\_JINR - SAM jobs aborting. 7) CMSSW installation issues at T2\_IT\_Pisa, site contact following up.

- ALICE reports [↗](#) - GENERAL INFORMATION: Pass1 rconstruction activities ongoing at the T0 (the minipass ITS-TPC tracking reconstruction mentioned yesterday has been completed). In addition, an analysis train and the already mentioned MC cycles are also running in the Grid



T0 site: ITCM ticket 356681 mentioned yesterday has been completed and the machines are back up and running. Thanks to the experts for the prompt action. Putting the machines back in production (sms) status is the pending action to do today. In fact 3 of the nodes went down early afternoon - a new ticket has been created.

T1 sites: No issues to report

T2 sites: 1) GGUS\_57847 ticket for Birmingham. Escalation procedure starting today. 2) GGUS\_57851 ticket for Cyfronet. The VOBOX required a new host certificate. No news since the 30th of April.

- LHCb reports - Experiment activities: No activity apart of user analysis

T0 site issues: 1) Yesterday's file deletion problem wasn't on lhcbdata but on default pool that makes much more sense. 2) During one hour time windows yesterday (from 13:30 to 14:30 UTC) all reconstruction jobs at all sites were failing to access a table on the condition DB. Open a GGUS ticket against DB people at CERN and it results that the problem wasn't service related but a user altered a table accidentally and then (the wrong) information correctly propagated down to T1s affecting jobs. Eva provided all details to chase this up LHCb side.

T1 site issues: 1) NL-T1: dramatic improvement since yesterday transferring the backlog of RAW data to PIC. This confirms that the variety of problems reported by Ron have been properly addressed. 2) CNAF: LHCb\_RDST space token issue. The SLS sensor indeed got stuck trying to parse the output of the metadata query over there that contains a SRM error message. (GGUS: 57948). Space Reservation with token=499eb057-0000-1000-8366-b54ffea05e44 StatusCode=SRM\_INTERNAL\_ERROR explanation=stage\_diskpoolsquery

T2 sites issues: USC-LCG2: failing voms SAM tests

Sites / Services round table:

- FNAL: Good work by OSG to resolve OSG-WLCG BDII non-transmittal issue last night. Problem affected more than FNAL, approximately 3/4 of OSG sites were affected. Rob reported that this seems to be an CEmon client problem so have turned off an CEmon collector at the OSG bdii but investigation continues.
- BNL: 1) One of their 3 1MW UPS systems is down with a bad battery block and has been bypassed. They are looking for a temporary solution over the next few days but till then some services will be at risk. 2) There is an FTS parameter option urlcopycopy-tx-to-per-mb which could be useful to detect and kill slow transfers hence free up FTS slots and help us optimise reconstruction work so advice on this would be appreciated. Zsolt will follow-up.
- CNAF: 1) For today's ATLAS ggus ticket they suspect network issues as they had a period of receiving no external input for 5 minutes. 2) LHCb have an open ticket for an RDST space token but they seem to be querying CASTOR instead of STORM - there is no CASTOR service in the CNAF bdii and they will update the ticket. Joel will cross-check what LHCb is querying. Flavia queried if CASTOR is still alive at CNAF as what LHCb get is an srm internal parsing error. Davide will check on his side.
- ASGC: Yesterday at about 16.30 ASGC FTS was found to be very slow, appearing down. Was due to database backup activity causing high load. Service recovered when database processes were cleaned up.
- CERN DBs (Eva): There was a high load during the night on the ATLAS archive database caused by ATLAS activities. The latest Oracle security patches have been received and we will be scheduling

their deployment, as usual firstly on the integration databases.

- CERN WMS (ManuelG/Maarten) : Upgrading of the 3 remaining CMS WMS is planned. Maarten reported that the one already done is performing very well having handled 43K jobs in one day with some capacity in reserve. The problem seen at CNAF with this version may be special as they had 20K unfinished jobs on their WMS at the time.
- CERN Frontier/Squid (Flavia): On 20 April a new release of the Frontier/Squid cacheing has been announced as the recommended release. ATLAS would like to have this installed at their T1 and T2 sites. She will add a url containing instructions into these minutes.
- CERN FTS (Maarten by email): Proposed plan - we have spoken with FTS developer Zsolt Molnar and other members of the data management middleware team: 2.2.4 will be frozen now and sent to internal certification today or tomorrow, and there is manpower to handle that process with high priority. If all goes well, certified rpms could be available for sites to upgrade next week. FTS 2.2.4 also contains a few other fixes and enhancements, but according to Zsolt it might take a significant amount of work and time to recode the wanted bug fixes against the pure 2.2.3 code. Therefore I think we should certify what we have now and do a staged deployment, e.g. first at 1 or 2 sites that are the most affected.

Atlas agree with this plan. CMS reply (from Nicolo) : before the fix for this issue is deployed, I would just like to get confirmation that FTS is not going to trigger deletion at the destination when a transfer fails for the following reason:

DESTINATION error during TRANSFER\_PREPARATION phase: [FILE\_EXISTS] File exists and overwrite option is OVERWRITE\_NEVER

In that case obviously we don't want to cleanup the DESTINATION file (unless we verify that it's corrupt...)

Reply from Zsolt: According to FTS logic, FTS deletes destination files only if it did not exist before the transfer. In preparation phase, FTS checks if the file exist, in such a case it does not delete files even in case of corrupt transfer. File removal is sent only in the finalization phase, never in preparation or transfer phases.

- RAL ATLAS diskserver (A. Dewhurst by email): As I mentioned earlier in the phone meeting, there is currently a disk server in downtime at RAL. One of the disk failed and there also appears to be a problem with the RAID card. There is no evidence to suggest that there has been any data loss but it will take quite a while to get it back in production, probably by early next week. A list of files that were on the disk server can be found at: /afs/cern.ch/user/d/dewhurst/public/gdss397.txt.gz . Is there anything that can be done in the meantime?

AOB:

- Release 2.7.STABLE9-3.9 of the frontier-squid distribution has been announced on April 26, 2010. Details can be found at the following URL: <https://twiki.cern.ch/twiki/bin/view/PDBService/SquidRPMsTier1andTier2>. This is the version ATLAS recommends for its Tier-1 and Tier-2 sites.

## Thursday

Attendance: local(Harry(chair), Ueda, Eva, Nicolo, JanI, ManuelG, Maarten, Akshat, Zsolt, Patricia, Lola, Alessandro, Douglas, MariaD);remote(Jon(FNAL), Tore(NDGF), Roberto(LHCb), Gang(ASGC), Gonzalo(PIC), Jos(KIT), Rolf(IN2P3), Davide(CNAF), John(RAL), Michael(BNL), Jeff(NL-T1), Rob(OSG), Jeremy(GridPP)).

Experiments round table:

- ATLAS reports -

RAL-LCG2 -- disk server problem -- We have started working on the list of the lost files, if we have other replicas or not, which ones we need to recover immediately and which ones we can wait, etc.

NDGF -- lost files (GGUS:57992) -- the symptom looks very much like the one we observed with TRIUMF (FTS deleting the source files).

SARA-MATRIX -- some files were unavailable since yesterday evening until this morning. Now they seem back online so no ticket was sent.

TAIWAN-LCG2 -- transfer failures -- fixed by now.

- CMS reports - T0 Highlights: 3 jobs stuck in Preparing phase on CERN-ASGC channel on fts22-t0-export.cern.ch - cancelled.

T1 Highlights: 1) Running tails of MC reprocessing - some long running jobs. 2) Running preproduction for next ReReco pass. 3) Testing train workflows - instead of waiting for a request, DataOps will update the workflow with new events and include new releases and global tags.

T2 Highlights: 1) Running MC production. 2) Upload of MC from T2\_RU\_SINP to T1\_US\_FNAL failing for permission errors. 3) T2\_RU\_PNPI transfers to T1\_CH\_CERN failing for timeout, forwarded request to adjust parameters on the dedicated channel on fts-t2-service.cern.ch Savannah #sr 114126, Remedy #681529. 4) SAM CE errors at T2\_US\_Florida - test jobs aborting. 5) CMSSW installation issues at T2\_UK\_London\_Brunel, installation jobs aborting.

- ALICE reports - GENERAL INFORMATION: Pass1 reconstruction, one analysis train and 2 MC cycles running smoothly with no major incidents to report.

T0 site: No news on the ITCM ticket 357112 submitted yesterday concerning 6 ALICE CAF nodes which seem to be down. Email to sysadmins to change the priority of this ticket to urgent. Feedback has been that these nodes are unreachable but not down so perhaps have an internal high load or memory use. So far the admins restarted xrootd but they are very busy and would like more information from ALICE - could there be an ALICE configuration change ?

T1 sites: GGUS\_58011 to CNAF: ce01 (one of the 2 CE-CREAM) failing (connection refuse) at the site. Message observed at submission time. Site admin also warned.

T2 sites 1) Cagliari-VOBOX: Proxy renewal daemon of the VOBOX dead. site admin contacted. 2) 2nd Kolkata VOBOX not accesible from CERN, site admin warned

- LHCb reports -

T0 site issues: The default pool was yesterday under heavy load and also SLS showed this unavailability (due to high number of concurrent transfers see the snapshot available here - limit is 200). This was affecting users running through the LXBATCH and LXPLUS. (GGUS: 57982). After investigation may have been an LHCb user triggering pool to pool copies.

T1 site issues: 1) CNAF: none of the transfers at around 8pm yesterday were going on. Open an ALARM ticket (57996) perceived as a show stopper. Problem fixed very quickly after one hour. Much appreciated. 2) LHCb\_RDST space token issue. Found a bug in the LHCb clients used. Corrected immediately.

T2 sites issues: UK T2 sites upload problem: seems to be due to a disruptive interference of their firewall.

Sites / Services round table:

- IN2P3: The last of 3 SIR reports, on the AFS outage, has now been filed.
- CNAF: 1) The failing ALICE CE is fixed by increasing the open files limit on the box. Final fix is in CREAM-CE 2.6 (Maarten reminded this is already released for SLC5 and is expected next week for SLC4. CNAF is still running CREAM under SLC4.) 2) gpfs had a problem - information sent to IBM.
- RAL: The failed ATLAS disk server is now about 40% rebuilt - had a disk failure but then crashed so not confident to put it back in production while rebuilding. Probably not back before Monday. Meanwhile Brian has been checking checksums and finding them to be correct.
- NL-T1: Today's ATLAS problem was due to a crashed dcache pool from a backup process that exceeded memory - now fixed. ALICE have jobs at Nikhef using zero cpu - ALICE will check.
- OSG: trying to recreate the bdi retransmission failure of CERN to OSG on a testbed.
- CERN WMS: all CMS WMS have now been upgraded and are working well - thanks to PES.
- CERN FTS: New version 2.2.4 had a small glite problem but should enter certification tomorrow.
- CERN CE: A new CREAM-CE (ce203) has been added into production.

AOB: (MariaDZ) Reminder: NB! Meeting next Monday 2010/05/10 @ 15:30 CEST in this room (same tel. number) on services involved in successfully delivered and handled ALARMS' tickets. Agenda <http://indico.cern.ch/conferenceDisplay.py?confId=93347>

## Friday

Attendance: local(Ueda, Harry(chair), Dirk, Eva, Nilo, Roberto, Maarten, Lola, JanI, Zsolt, Alessandro, Flavia, Douglas, Nicolo, Akshat);remote(Jon(FNAL), Vera(ndgf), Michael(BNL), Gang(ASGC), Rolf(IN2P3), Daniele(CNAF), Rob(OSG), Tiju(RAL), Brian(RAL), NL-T1).

Experiments round table:

- ATLAS reports -

FTS -- "DESTINATION error" for the files UNAVAILABLE at the source (GGUS:58042) (seen at RAL) Not as urgent as the previous FTS issues.

RAL-LCG2 -- unavailable files -- shifters did not well understand the situation and sent a new ggus ticket and updated the existing tickets. sorry for the bother. maybe putting a downtime could help (shifters are trained not to send a ticket when the site is in downtime).

NDGF -- files not exist on SE (GGUS:57992) -- If possible, please provide a list of the files not existing on SE but registered on LFC.

ATLAS internal: RAL-LCG2: replicated some datasets, which were available in other T1s, into scratchdisk, so that the number of errors decrease. Brian reported that files were still trying to be accessed from datadisk even when there were copies on scratch - apparently this is a feature of the ATLAS DDM.

- CMS reports -

T1 Highlights: 1) Running tails of MC reprocessing. 2) 2 jobs running for 3 days at CNAF with 0 CPU time used, site contact notified - being looked at. 3) List of files leftover from failed production submitted to CNAF for cleanup. 4) ReReco train workflows running. 5) Skimming running. 6) WMAgent commissioning. 7) Backfill tests running. 8) RAL - some files waiting for tape migration for a long time (related to low volume of files to migrate).

T2 Highlights: 1) Running MC production. 2) Thanks to Maarten for providing method to detect jobs stuck in "Submitted" state in WMS for loss of input sandbox. Maarten reported that only a small fraction of jobs were left in this state. In fact a whole job collection was aborted with an error leaving some jobs in this state so this may not in fact be a WMS problem. He is waiting for some feedback from CMS. 3) Upload of MC from T2\_RU\_SINP to T1\_US\_FNAL failing for permission errors. 4) T2\_RU\_PNPI transfers to T1\_CH\_CERN failing for timeout, forwarded request to adjust parameters on the dedicated channel on fts-t2-service.cern.ch Savannah #sr 114126, Remedy #681529. Some tuning has now been done. 5) CMSSW installation issues at T2\_UK\_London\_Brunel, installation jobs aborting. 6) T2\_EE\_Estonia - GlueServiceUniqueID tag for new SRM endpoint not propagating from BalticGrid BDII to top-level WLCG BDII GGUS #58016, preventing use of new SRM for production transfers with FTS.

- ALICE reports [ALICE](#) - GENERAL information: Short run (118903) recorded with stable beam tonight. Pass1 reconstruction tasks together with one train analysis and two MC cycles are smoothly running on the Grid.

T0 site: Last night voalicefs01 to voalicefs03 have had a linear growth in load factor - now up to 120 so they will soon stop responding. Investigating the issue with Alice, also to find a possible parallelism with the ALICE CAF nodes which shown connection problems during the last days (thank so much to Fabio for his help these days solving these issues).

T1 sites: CNAF: Problem with the CREAM-CE reported yesterday solved. system back in production. Site admin can close the ticket.

NL-T1: GGUS\_58038 ticket submitted this morning: The local CE was providing wrong information and this issue was affecting the Alice production at the site. The issue has been already solved and confirmed by Alice. New agents submission can be now followed up to investigate the issue reported yesterday by Jeff

T2 sites: Configuration changes performed this morning at JINR to put the site in pure CREAM-CE mode

Authentication issues observed at the vobox in Catania yesterday. The ALICE expert proxy was corrupted at the vobox. The issue has been solved and the site is back in production

- LHCb reports - Experiment activities: Launched yesterday few production to validate new workflows. No big deal. One production at 91.6% with few stalled jobs under investigation and the remaining 1546 RAW files expected to finish later in the afternoon. Few small MC production requested and about to come in the weekend together with high intensity data from the LHC.

T1 site issues: pic: DIRAC services at pic are down (GGUS 58036)

T2 sites issues: RO-15-NIPNE failing all jobs.

CIC Portal: still some instabilities with bunch of notifications (as backlogs formed on the server than sent all at once) and/or missed announcements.

Sites / Services round table:

RAL: Failed ATLAS diskserver should be repaired by tonight then put back in production tomorrow.

OSG: bdii retransmission from CERN still under investigation. emon collector on one server (of two) died and this machine will probably be taken offline on Monday but will be checked with CMS and ATLAS first. Maarten queried if a single server could cope with the load pointing out that the OSG queries are somewhat inefficient with the same query being done multiple times. The SAM team have an open Savannah bug on this, number 66120. Rob was confident the machine could handle more than the total load they see now.

CERN FTS: Going to review the modifications in the certification testbed to be sure the recent critical errors were covered. Maarten reported that there had been a delay in the certification process due to a change in procedure in moving from EGEE to EGI but he thought a release could be made early next week.

AOB:

-- JamieShiers - 29-Apr-2010

---

This topic: LCG > WLCGDailyMeetingsWeek100503

Topic revision: r19 - 2010-06-11 - PeterJones



Copyright &© 2008-2021 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](#)