

# Table of Contents

<b>Data management issues during STEP'09.....</b>	<b>1</b>
Introduction.....	1
Daily summaries.....	1
June 2.....	1
June 3.....	1
June 4.....	2
June 5.....	2
June 6.....	3
June 7.....	3
June 8.....	3
June 9.....	4
June 10.....	5
June 11.....	5
June 12.....	6
Summary of significant issues by category.....	7
Tape systems.....	7
Issues at FZK.....	7
Issues at IN2P3.....	7
Issues at RAL.....	7
Issues at ASGC.....	8
Issues at SARA.....	8
Issues at FNAL.....	8
Issues at CNAF.....	8
dCache.....	8
CASTOR.....	9
Data losses.....	9
The "Big ID" problem.....	9
Other problems.....	9
StoRM.....	9
Middleware.....	9
GFAL.....	9
FTS.....	10
LFC.....	10
Local data access.....	10
Transfers.....	10
Other storage issues.....	10
Experiment data and workload management services.....	10
ATLAS.....	10
CMS.....	10

# Data management issues during STEP'09

## Introduction

This page collects the issues seen by experiments and sites on their data management and storage systems. They are classified as follows:

- **[TAPE]**: problems related to the tape systems
- **[MW]**: problems related to middleware (e.g. FTS, GFAL, lcg-util, root (!))
- **[CASTOR]**: problems related to CASTOR
- **[dCache]**: problems related to dCache
- **[STORM]**: problems related to StoRM
- **[STORAGE]**: unspecified storage problems
- **[TFER]**: problems related to transfers
- **[WDM]**: problems related to the experiment's data or workload management systems (e.g. PhEDEx, DQM)
- **[HW]**: problems related to hardware failures (e.g. broken disk servers)

## Daily summaries

### June 2

ALICE
Will run FTS transfers only during second week.
ATLAS
Tests are starting, including pseudo-reprocessing jobs reading tapes.
<b>[TFER]</b> Transfer failures from IN2P3 to ASGC [why?].
<b>[TFER]</b> Slow transfers from TRIUMF to CNAF [why?].
<b>[TFER]</b> Slow transfers from TRIUMF to two Canadian Tier-2 and from SARA to Tier-2.
<b>[TFER]</b> Transfer failures to Coimbra and Lisbon.
<b>[MW]</b> Bug in a GFAL version included with DQ2, had to downgrade GFAL.
CMS
AODSIM data import: tape families set at FNAL and RAL, disk-only area at ASGC, CNAF, FZK, IN2P3, PIC.
Prestaging via 1) PhEDEx agent at ASGC, RAL, PIC; 2) GFAL/SRM at CNAF; 3) manual at FNAL and IN2P3 (after HPSS downtime).
<b>[TAPE]</b> FZK not participating this week to tape tests due to SAN network problems. Still able to migrate to tape incoming AODSIM data written to tape write pool and copied to read pool.
<b>[TAPE]</b> IN2P3 not participating this week to tape tests due to HPSS upgrade.
<b>[TAPE]</b> ASGC cannot yet prestage due to a tape backlog.
<b>[WDM]</b> At RAL, prestaging happened on the wrong pool due to a mistake in the PhEDEx configuration.
LHCb
Will start during second week. This week, preparation including deletion from disk of files to prestage.
Sites
<b>[dCache]</b> NIKHEF (or SARA?): LHCb SRM tests failing due to the dCache disk pool 'cost' function not load balancing correctly across full disk pools.

### June 3

ATLAS
-------

Data distribution going well.
Soon will test xrootd analysis at CERN.
Prestaging at FZK not bad, even if with a degraded tape system.
Want to check that at CERN data committed to tape is going to tape (tapes being recycled).
[DM] Prestaging not happening at NGDF, causing backlog increase.
<b>CMS</b>
PIC sees a prestage rate of 60 MB/s, the maximum possible with only two tape drives.
PIC: stager agent was set up to make more efficient WAN outgoing transfers for data on tape.
RAL has a rate of 250 MB/s.
CNAF has a rate of 400 MB/s.
[TFER] Low quality seen in T1-T1 transfer tests.
<b>Sites</b>
[HW] RAL had some low level disk server failures.
[TAPE] ASGC has five drives this week, 7 next week, one not working.
[CASTOR] ASGC will upgrade CASTOR to 2.7-18 to fix the "big ID" issue.
[MW] BNL had a site services problem affecting transfer performance due to a GFAL library compatibility issue. After fixing this they observed up to 20000 completed SRM transfers per hour.
[MW] NDGF had to increase FTS timeouts from 30' to 90' to accommodate long running transfers.

## June 4

<b>ATLAS</b>
[MW] NDGF slow getting data from CERN due to a too low number of concurrent transfer jobs on FTS. Changed from 20 to 40.
[dCache] Problems transferring data to FZK due to a too low limit on number of file handles on a storage node.
[STORAGE] Various Tier-2 storage problems.
<b>CMS</b>
ASGC prestaged at 190 Mb/s . Some files requested by jobs before being prestaged, but this did not cause problems.
CNAF prestage at 380 MB/s. Difficult to monitor tape usage.
At IN2P3 the HPSS downtime will end before the weekend, but tape activity is not desired until next week.
At RAL CMS is competing now with ATLAS. Tape monitoring available <a href="#">?</a> .
[WDM] Not using LazyDownload cause excessive network traffic at RAL and PIC.
[WDM] T1-T1 transfers with CNAF bad due to PhEDEx not properly supporting two SEs at the same site.
[TAPE] At FZK writing to LT03 working, LT04 reading and writing not working, no estimate for fix of tape system.
<b>LHCb</b>
[CASTOR] About 7000 files were lost at CERN in the <i>lhcbdata</i> service class due to a CASTOR upgrade which turned on by accident the garbage collector (post-mortem).

## June 5

<b>ATLAS</b>
[CASTOR] At RAL the Big ID problem showed up hitting the stager and causing a CASTOR component to go down. There was a session killer to remove the offending sessions, but it was not smart enough (see <a href="#">here</a> <a href="#">?</a> ).
[TAPE] At SARA prestaging stopped at midnight, due to problems with tape backend solved with a DMF reboot.
[WDM] NDGF did not have prestage requests due to a difference in the ARC architecture where there is an extra state in the workflow.

LHCb
GGUS tickets will be sent to T1 sites with the lists of files to delete from disk buffers.
[MW] Cannot read non-root files via gsidcap at SARA and IN2P3 due to some dCache parameters controlling read-ahead which appear to be set to bad values, and not by LHCb software.
Sites
PIC is going to install four new LT04 drives on the STK robot, which will multiply by 3 the stage rate.
IN2P3 will start prestaging on June 8 for CMS and June 9 for ATLAS. A new script to optimize tape access is now in place.
[MW] RAL increased FTS timeouts which eliminated transfer errors with BNL.
[MW] TRIUMF increased timeouts, too.
[STORM] CNAF fixed some mispublishing of StoRM disk pool information which was causing slow transfers TRIUMF-CNAF <i>[how is that possibly related?]</i> .

## June 6

CMS
Prestaging at RAL ran with CMS and ATLAS together with no problem; CMS has 4 drives dedicated, ATLAS has 4 drives dedicated, LHCb has 1 drive dedicated.
[TAPE] FNAL started to recover from tape problems. Traced back to relatively small queue depth in Enstore's library manager before it overloads and the very large amount of transfers at FNAL. This increased the rate of "seeks" on tapes dramatically in combination with a 1 minute delay in reads which is 'normal' for non adjacent files on LTO4.

## June 7

CMS
FNAL is recovering tape performance.
[HW] Hardware failure on one CASTOR disk server at CNAF.
[CASTOR] At RAL, still seeing more import rate via PhEDEx than tape writing rate although more tapes have been added. It seems that Castor is only writing to a tape then stopping. James found the problem, backfill was pointed to a tape family only attached to the farm service class, T0 data importing to the Import service class was never going to migrate, fixed now.

## June 8

ALICE
[MW] Some unspecified problem with FTD preventing the startup of transfers foreseen for this week <i>[what was it?]</i> .
ATLAS
Over the weekend, reprocessing was smooth at PIC, RAL, BNL and TRIUMF.
Last week's problem at NGDF consisting in prestage not being triggered looks like a configuration problem in Panda. Prod system doesn't realise files on tapes and hence doesn't trigger prestage.
IN2P3 restarted tapes for writing, not yet for reading.
[CASTOR] CNAF had some CASTOR problems meaning that it was slow but prestaging worked fine.
[TAPE] SARA tuned DMF <i>[what is it?]</i> and storage R/W is better balanced.
[dCache] FZK still very slow for incoming transfers, seemingly due to an overloaded SRM (GGUS #49313 <a href="#">↗</a> ).
CMS
[TAPE] FZK: tape writing is working, but there are still issues with tape reading.
[CASTOR] CNAF: stageout failures due to wrong permissions.

[TAPE] FNAL: no problems in tape writing, up to 1.4 GB/s; problems in reading, being below what was achieved last week (and previous weeks): being investigated over the weekend. Recovered the needed tape read rate on Sunday but had a large backlog since then.

#### LHCb

STEP09 starts tomorrow for LHCb.

[CASTOR] 6500 more files lost from CASTOR at CERN.

[MW] Investigating problem at NIKHEF preventing from reading non-root files via gsidcap.

[STORM] Cannot copy files out from or to StoRM at CNAF.

#### Sites

CERN: Post mortem on LHCb data loss ([link1](#) [link2](#)).

[TAPE] FZK: working on tape problem. Fixed config errors. Newer version of TSM. Consolidated stagin - saw improvement over w/e. Still looking to find bottleneck. Some of gridftp doors ran into memory problems - investigating it.

[dCache] IN2P3: reading from tape works fine since the end of the HPSS intervention, but also ATLAS was hit by the srm-ls locality problem (Fabio's email).

## June 9

#### ALICE

Began some transfer tests using a new FTD instance.

#### ATLAS

Prestaging works fine at NDGF after fix in Panda.

Prestage working fine at IN2P3 after fixing srm-ls bug.

T2s: quite a few with backlog (~12 sites >24h). Check for lack of FTS slots, slow transfers hogging slots, overload from analysis.

ASGC: good news is that late yesterday listener on old port working. Prestaging is slow, though.

[dCache] FZK healthy. Site reports that analysis activity, using excessive lcg-cps to stage data, was loading the gridftp doors. After this was stopped the SE has performed well.

[dCache] BNL has an increasing transfer backlog. They had to reboot their SRM early this morning.

[TAPE] RAL reported that they are not writing to tape.

#### CMS

FZK ready to start tape activities.

At RAL, fixed issue with CASTOR migrator. *[see Facops HN for details]*

[TFER] A file corrupted after transfer to ASGC.

[TAPE] CNAF: prestaging slow, and first files started to arrive only after 2.5 hours, due to other concurrent activities. One file never came online according to statusOfBringOnline while it was for Ls. At risk tomorrow due to late delivery of cleaning tapes.

[dCache] IN2P3 has issues at the pool level due to a bug. The workaround is to reboot the pools. A possible fix is to upgrade to a more recent version of dCache.

[dCache] FZK: many users complain and production activities report problems accessing data through dcap.

[dCache] NL-T1: problem accessing raw files still being investigated.

[MW] IN2P3 also had dcap problem. Discovered some limitation in dcache client library. New unreleased version recommended.

#### Sites

IN2P3 introduced a new component scheduling stager requests before sending to HPSS. Started yesterday for CMS - very happy with initial results, but must analyse further. Reduction of number of tape mounts is the purpose of this new component. Being activated for ATLAS.

ASGC: streams replication is working fine now.

[dCache] BNL discovered that 3 of 4 write buffer servers are idle. Under investigation. Result is BNL not keeping up.

## June 10

ALICE
[MW] At NDGF the problem was with the destination path in FTD. After changing was ok.
[CASTOR] RAL: site admins needed to fix write permissions.
ATLAS
[dCache] dCache at IN2P3 went down at 0100 UTC. Resolved at 0800 UTC after a team <a href="#">☞</a> and an alarm <a href="#">☞</a> ticket where sent. The reason was that the dCache internal routing to gridftp server was down.
[dCache] SARA_MCDISK almost ran out of space.
[TAPE] In the morning RAL reported severe tape robot problems.
[TAPE] CNAF has reported drives disabled awaiting fresh cleaning cartridges, so degraded tape performance.
[CASTOR] CASTOR at CERN was discovered to be limited to about 3 GB/s for exporting data. This was expected given the pool configuration and the hardware.
CMS
[TAPE] Only two tape drives available at RAL because of a stuck cartridge. Rebooting the robot recovered all drives apart from the stuck one. Performance was enough to absorb the backlog.
[WDM] Could not transfer AODSIM to FNAL because their FileDownloadVerify agent failed after the PFNs where changed to avoid putting too many files in a single directory. Solved by reverting to the standard PFNs.
LHCb
[TAPE] At CNAF, stage attempts failed due to the problem with the cleaning cartridges.
[MW] Investigation with Jeff has discovered a middleware issue with dcap plugin for root. Enabled in DIRAC the configuration to read data after a download into the WN (ATLAS approach). Running smoothly now.
[MW] LFC issue: due again to Persistency interface.
[STORM] MC_DST and DST space tokens at CNAF. It seems that the space token depends on the path while StoRM should guarantee the independence between space token and namespace path.
Sites
[TAPE] FZK made available a post-mortem for their tape system problems. The main problem was an unstable SAN connectivity, but also memory problems in the disk servers. Transfers seem better now.
[dCache] IN2P3: Since 3:00 AM this morning all gridFTP transfers were aborting. This was due to a failure with the component in charge of routing requests to gridFTP servers. The corresponding service has been restarted and Dcache is full available since 11:00 AM. Probably due to some load - many pending requests

## June 11

ALICE
[MW] Having problems with the <code>-o</code> option of <code>glite-transfer-submit</code> , as the "File exists" error message is received.
ATLAS
[TFER] Geneva-Frankfurt fibre cut. Due to rerouting, ASGC have a serious problem with <code>gridftp_copy_wait</code> : Connection timed out from other T1s and from CERN.
CMS
[TFER] OPN cut impact on CMS, mainly links to/from ASGC.
[TAPE] At IN2P3 the Tier-2 was closed because the analysis jobs caused files to be staged from HPSS and thus interfering with Tier-1 operations. This was a known issue and some actions were already identified for strictly separating the storage spaces of the Tier-2 from those of the Tier-1 and to modify the configuration of the Tier-2 to make it a disk-only site. Unfortunately, this could not be finished before the beginning of STEP'09.
LHCb
[MW] STEP09 reprocessing and staging confirmed that Persistency interface to LFC needs urgently to be

updated. A lot of jobs failing to contact LFC server because brought down by the inefficient way Persistency queries it. A Savannah bug open (ref. <https://savannah.cern.ch/bugs/?51663>). This prevents to continue STEP09. Killing currently running jobs. LHCb will ask again sites to wipe data from the disk cache and rerun the reprocessing using SQLite slices of the Condition DB (instead of using it directly) in order to compare the two approaches.

**[TAPE]** Staging at SARA. All files to be reprocessed failed after 4 attempts.

**[CASTOR]** File access problem on CASTOR at CNAF.

**[dCache]** Issue transferring to MC\_DST at IN2P3. Disk really full.

#### Sites

ASGC: a problematic LTO4 drive has been fixed and all new tapes are now functional.

**[TAPE]** RAL: Still have 78k files refusing to budge. After running last night, the migrator process dealing with these started impacting the server, taking >90% of memory and still using 20%+ CPU time with no sign of a migration. Restarted the mighunters again just to free up resources on the machine, and will continue to monitor (see [here](#)).

**[dCache]** Stage throughput cut in half for CMS due to a stager down.

**[dCache]** FZK: Failing stageout for ATLAS: one GridFTP door had to be restarted because out of memory. One library down for 2 hours. Fixed. Currently some drives used for writing are not accessible because of SAN zoning problem (a FW bug).

**[dCache]** BNL: SRM server issue: overloaded due to analysis jobs staging out output files via lcg-cp so to write to the correct space. Need a version of dcpp capable of that. Developers say dCache 1.9.2 but becoming an immediate and urgent issue.

**[TFER]** BNL: Networking: big fibre cut in Frankfurt area affecting also US. For BNL bandwidth reduced from 10 to 4.2 Gbps.

**[TFER]** Network: Two independent fibre cuts .

## June 12

#### ALICE

**[MW]** The problem with FTS reported yesterday was due to not using fully qualified URLs to force SRM 2.2 (SRM 1 does not have an overwrite option). After using proper URLs the problem was solved.

#### ATLAS

**[dCache]** Many jobs died at FZK due to a gridftp door problem.

**[STORM]** Problem with StoRM at CNAF overnight - fixed this morning.

#### CMS

**[TAPE]** FZK tapes fine now, but the number of recalled files marked as being on disk was still quite low. This was due to two reasons: 1. One of the 2 CMS stager nodes had limited throughput due to copying data to itself (stager runs on same as for some other pools). 2. Config error on the stager pools. Majority of the data staged back was copied to the tape write pools instead of the tape read pools. (Some behaviour which seems to show up under heavy load of the system).

#### LHCb

During the weekend, staging and reprocessing data will be done w/o the remote ConditionDB access (that currently would require a major re-engineering of the LHCb application). but via available SQLite CondDB slices.

**[dCache]** The issue with staging at SARA seems related to the fact that all directories to be staged were also removed from tape.

#### Sites

**[STORAGE]** GRIF: ATLAS observed lcg-cp timeouts at all French sites (Hammercloud tests). Suspect not site nor load - maybe some race condition? Similar problem reported also with STORM during pre-GDB. Investigating this.

FZK: SAN switch reset after contact was lost to 8 LTO3 drives. Datasets from LHCb were unavailable. Some data from CMS. Restarted movers that stopped because of this outage. Restarted a mover for ATLAS which was stuck for more than 5 hours. Trigger went unnoticed because of intense other activities. Rates picked up

to >100MB/s after restart. Also a CMS pool was recalling data but then never showed them when they arrived. Restart fixed this. 3 of 4 gftp proxies (doors) locked up. Cause unknown. Restarted. We ave now 2 more stager nodes ready that contain no other pools.

## Summary of significant issues by category

### Tape systems

Given the specificity of tape systems, the issues related to these are broken down by site. As a general observation, tape systems seem to be among the most delicate and complex components and the related issues can easily cause serious trouble to a site.

#### Issues at FZK

FZK could not participate to tape-related activities in the first week of June due to extremely bad performances of their tape system.

An update of the tape library manager resulted in random failures in accessing the drives. At first these failures were diagnosed as SAN problems, and a second emergency SAN was set up, which although would provide only limited rates. At the same time a configuration parameter in the tape library manager was changed, which restored the normal behaviour. Just before the start of STEP'09 it was decided to leave the tape system to ATLAS alone for the first week.

In the first week also other problems were observed, including hardware problems in disks, library and tape drives, frequent memory exhaustions on pool nodes, the TSM software crashing.

Things improved in the second week, when two dedicated stager hosts were added, to reduce the pool-to-pool traffic on the existing four nodes. This greatly improved the stability and allowed to achieve stage rates of 100-150 MB/s. However a configuration error caused the data on the CMS stager to be trashed, with files being deleted before being copied to the proper read pools. This is still being investigated. Finally, tests done with ATLAS showed rates compatible with the nominal ones.

There is a detailed post-mortem.

#### Issues at IN2P3

Lyon could not participate to the first week of tape-related activities due to a scheduled HPSS upgrade, started on June 1 and completed on June 4. After its completion, tape access was fully restored, but ATLAS thought it was only for writing due to the "srm-ls problem" (consisting in the wrong locality being reported), which made them think that the files could not be staged.

After the HPSS, the experiments were asked to avoid the prestage tests before June 9, in order to be able to introduce a new component responsible for scheduling the tape staging requests sent to HPSS by dCache.

On June 11, the CMS Lyon Tier-2 was closed to analysis jobs because they were causing files to be staged from HPSS and interfering with Tier-1 operations.

#### Issues at RAL

A first migration problem was observed for CMS, and fixed on June 7. It was due to only having previously set up the tape family used for the CMS "backfill" data on the Farm service class. This meant that Tier-0 data importing to the Import service class was never going to migrate, no matter how many drives got thrown at it (HN) [↗](#).

A bug in the CASTOR migrator was then discovered. As RAL had CASTOR set up, they would have no control over the number of migration streams per tapefamily: the policies for managing migration stream numbers are not being respected by CASTOR (HN) [?](#).

On June 9 RAL reported that they were not writing to tape due to migration issues.

On June 10, severe tape problems due to a stuck cartridge which caused all tape drives but two to become unavailable. A reboot of the robot recovered all drives but the stuck one. A fix for the migrator bug was applied.

On June 11 the migrator process was consuming more than 90% of the memory with little CPU usage and no migration happening. Fortunately there was enough free space on the disk buffers for the service not to be affected.

### **Issues at ASGC**

At the start of STEP'09 ASGC could not start prestaging data for CMS due to a previous tape backlog.

A tape driver had a problem, but it was fixed.

### **Issues at SARA**

A couple of issues were reported, related to the Data Migration Facility. One was solved by a server reboot, one with a configuration tuning. On June 11 LHCb reported failures in staging data.

### **Issues at FNAL**

On June 6 FNAL solved a performance problem seen on June 5 due to a too small queue depth in Enstore's library manager combined with a very large of transfers. The rate of "seeks" on tapes increased dramatically and with a 1 minute delay in reads typical for LTO4 drives led to very low stage rates.

### **Issues at CNAF**

The most significant problem was related to the unavailability of cleaning cartridges combined with a bug in the firmware of the Sun T10000 drives, causing an excessive usage of those cartridges. As a consequence, several tape drives became unavailable.

## **dCache**

Some of the problems seem related to configuration or operational issues:

- A wrong configuration of the cost function caused an incorrect balancing of full disk pools causing LHCb tests to fail at SARA.
- A too low limit on the number of file handles on a storage node (FZK).
- srm-ls reporting the wrong locality (NEARLINE) for files successfully staged (IN2P3).
- A space token becoming full (SARA).

Many stability and overload problems were observed:

- An overload of the SRM server (FZK).
- SRM required rebooting (BNL).
- One caused by too many concurrent lcg-cp (FZK).
- Problems reported with dcap access (FZK).
- 3 over 4 write buffer servers were idle (BNL).

- The dCache internal routing to gridftp server was down, solved by a service restart and probably load related (IN2P3).
- Frequent memory exhaustions on pool nodes or GridFTP doors (FZK).

Finally, these problems should be related to dCache bugs:

- Pools needed to be rebooted regularly, the proper fix being a dCache upgrade (IN2P3).

## CASTOR

In the case of CASTOR, not many problems were seen; in particular stability did not seem to be an issue, apart from the "Big ID" bug, which nevertheless did not have a significant impact.

### Data losses

A rather serious problem was the loss of thousands of LHCb files at CERN on June 3. During the CASTOR 2.1.8-7 upgrade on the 27th of May, the garbage collection was accidentally re-enabled in the *lhcbdata* service class. This service class had tape backup enabled since several months, but ~65000 files were created before and had no tape copy. Of those, about 7600 had no other disk copies and were utterly lost. The cause was a bug in the upgrade script: it caused D1T1 service classes to become D0T1 after the upgrade.

On June 5, another incident caused the loss of another ~6600 LHCb files. This time it happened when trying to make T0D1 files into T1D1 ones with a script that unveiled a bug in CASTOR such that files with more than one disk copy had all their copies deleted.

### The "Big ID" problem

This infamous problem hit the ATLAS RAL stager for a few hours, causing a CASTOR component to go down. There was a session killer to remove the offending sessions, but it was not smart enough. At ASGC, they applied a CASTOR patch for it on June 3.

### Other problems

Disk server failures were observed at CNAF.

Wrong permissions of directories at CNAF and RAL.

A bug in the CASTOR migrator was discovered at RAL (see Tapes section).

## StoRM

LHCb had some failures moving data to and from StoRM at CNAF, as seen by their SAM tests. No more details available.

LHCb reported that apparently the space token for a file depends on the file path instead of being orthogonal to it.

## Middleware

### GFAL

Right at the start of STEP'09, the GFAL version distributed with DQM was found to have library files in a different location, compared to the PYTHONPATH set by the Site Services installation kit. ATLAS decided to revert to a previous GFAL version.

## FTS

Problems with FTS were largely due to configuration problems or user mistakes:

- Due to the large size of some ATLAS files, FTS timeouts on transfers had to be increased at NDGF, BNL, TRIUMF and RAL.
- Slow data transfers to NDGF due to a too low number of concurrent transfer jobs.
- ALICE was using the `-o` option of `glite-transfer-submit` but not using fully qualified SRM v2 SURLS, which cause FTS to use SRM v1 SURLS, for which the overwrite option does not work.

## LFC

A show-stopper for LHCb was the fact the the LFC interface in Persistency was rather obsolete and inefficient, and using it killed the LFC server. This prevented from directly accessing the ConditionDB and SQLite slices of it had to be used instead.

## Local data access

LHCb could not read non-root files via `gsidcap` at SARA, NIKHEF and IN2P3. After an investigation, Jeff Templon discovered a middleware issue with the `dcap` plugin for root. The workaround was to enable in DIRAC the configuration to read data after a download into the WN.

## Transfers

The biggest problem in this area was caused by a fibre cut between Geneva and Frankfurt affecting several LHCOPN links (from CERN to ASGC, CNAF, FZK, NDGF and TRIUMF). The consequent rerouting caused GridFTP timeouts in transfers to ASGC. In addition, fibres used for USLHCNet were lost and hence FNAL and BNL were also affected, the effect being a big reduction in the bandwidth. The traffic was automatically rerouted via backup links.

## Other storage issues

ATLAS observed `lcg-cp` timeouts at all French sites (Hammercloud tests). Suspect not site nor load - maybe some race condition caused by a kernel bug? Similar problem reported also with STORM during pre-GDB. Under investigation.

## Experiment data and workload management services

### ATLAS

Prestaging at NDGF did not happen initially because of a problem with Panda related to the special architecture of ARC sites.

### CMS

At some point at RAL prestaging happened to the wrong pool due to a mistake in the PhEDEx configuration.

Again at RAL, an excessive internal network usage was traced back to the "lazy download" option being disabled in production jobs.

At CNAF, Tier-1-to-Tier-1 transfers with CNAF were bad due to the PhEDEx inability to correctly deal with the case of two different SE endpoints at the same site: exports did not work because the stager agent did not work with StoRM, imports did not work because there was no way for PhEDEx to correctly assign a transfer to the FTS channel with StoRM rather than to the FTS channel with CASTOR.

## Step09ExplIssues < LCG < TWiki

At FNAL, incoming AODSIM transfers failed because of a change in the PFNs to avoid having too many files in a single directory, which caused the FileDownloadVerify agent to fail on the new PFNs.

-- AndreaSciaba - 30 Jun 2009

---

This topic: LCG > Step09ExpIssues

Topic revision: r7 - 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](#)