

Table of Contents

September 2009 Reports.....	1
30th September (Wednesday).....	1
29th September (Tuesday).....	1
28th September (Monday).....	2
24th September (Thursday).....	2
23rd September (Wednesday).....	3
22nd September (Tuesday).....	3
18h September (Friday).....	4
17h September (Thursday).....	4
16th September (Wednesday).....	5
15th September (Tuesday).....	5
14th September (Monday).....	8
11th September (Friday).....	8
9th September (Wednesday).....	9
8th September (Tuesday).....	9
7th September (Monday).....	9
4th September (Friday).....	10
3rd September (Thursday).....	11
2nd September (Wednesday).....	12
1st September (Tuesday).....	13


September 2009 Reports

To the main

30th September (Wednesday)

Experiment activities:

- FEST reconstruction did not start yet as expected yesterday because of a problem with options file used by DaVinci application. However transfers pit-CASTOR at CERN are taking place right now.



- In the system just few thousand jobs running from distributed user analysis activity.

GGUS (or Remedy) tickets :

- 2 new tickets since yesterday
 - ◆ T0: 0
 - ◆ T1: 0
 - ◆ T2: 2
- List of all open tickets [↗](#)

T0 sites issues :

* T1 sites issues:*

- RAL: it has been reported that opening a file takes about 1 hour causing many calibration jobs to timeout

* T2 sites issues: * shared area issue and SQLite file issue.

29th September (Tuesday)

Experiment activities:

- Today the FEST week starts. Data sent from ONLINE to the OFFLINE for being processed through grid production work flows at T0/T1 centers. As usual: first of all the XPRESS Stream will be analyzed and then- after DQ WG green light - the FULL stream that this time supposed to run at 1.8Hz HLT data acquisition rate (~60MB/s throughput to CASTOR and then to T1 accordingly a share assigned to each).
- Stripping work flow has also to be commissioned and this is the top priority task for the next weeks.

GGUS (or Remedy) tickets :

- 1 new tickets since yesterday
 - ◆ T0: 0
 - ◆ T1: 0

◆ T2: 1

- List of all open tickets [↗](#)

T0 sites issues :

T1 sites issues :

- IN2p3 is banned because of the scheduled upgrade to chimera until October 1st. It will not take part to this FEST week.

T2 sites issues :

- pilot aborting in smallish sites.

28th September (Monday)

Experiment activities: No production requested. No activity in the system apart few hundred user jobs at T1's.

GGUS (or Remedy) tickets :

- 4 new tickets since last week

◆ T0: 0

◆ T1: 0

◆ T2: 4

- List of all open tickets [↗](#)

T0 sites issues :

- A host certificate expired on one of the VOBoxes in production (volhcb04.cern.ch). No lemon-alarm triggered FIO side because in maintenance since August.

T2 sites issues :

- Shared area issues in some sites.

24th September (Thursday)

Experiment activities: EGEE'09 jobs (~4K jobs in the system) + some user jobs (~1K).

GGUS (or Remedy) tickets :

- 4 new tickets since yesterday

◆ T0: 0

◆ T1: 0

◆ T2: 4

- List of all open tickets [↗](#)

T0 sites issues :

- Agreed to proceed with SRM migration to 2.8 at 14:00 today.

T2 sites issues :

- Shared area issues in some sites.

23rd September (Wednesday)

Experiment activities: EGEE'09 fake jobs (few hundreds) + usual SAM jobs.

GGUS (or Remedy) tickets :

- 2 new tickets since yesterday
 - ◆ T0: 0
 - ◆ T1: 0
 - ◆ T2: 2
- List of all open tickets [↗](#)

T0 sites issues :

- LHCb have been requested to arrange an intervention upgrade (~1h) on the SRM service (moving to 2.8 that proper supports for xroot TURLs, and operationally has substantial improvements to the logging which should aid in debugging any problems on the service). May be this quite week could be a good time slot.
- Intervention on the LFC is over but the following errors appear in the read-only LFC server log (under investigation):

```
09/23 10:47:41 10815,2 Cns_serv: [128.142.142.150] (volhcb10.cern.ch): Could not establish an aut
```

T1 sites issues :

T2 sites issues :

- Shared area issue at some T2

22nd September (Tuesday)

Experiment activities: LHCb week in Florence: not activities going on. No scheduled activity nor user activity is running apart few MC jobs for EGEE'09 demo.

GGUS (or Remedy) tickets :

- 2 new tickets since yesterday
 - ◆ T0: 0
 - ◆ T1: 0
 - ◆ T2: 2
- List of all open tickets [↗](#)

T0 sites issues :

- CASTOR LHCb upgrade to 2.1.8-12: non transparent intervention downtime 9:00-13:00. In the same slot [↗](#) also an upgrade to CASTOR Gridftp internal configuration that lcg_utils/gfal now seem to support for LHCb.
- Intervention on LFC: Upgrade to LFC version 1.7.2-4. Agreed for tomorrow morning at 10 (profiting of this inactive week)

T1 sites issues :

- CNAF: A problem with all data (intermittently) missing from StoRM fixed quickly on Saturday evening.
- IN2p3: banned due to dCache upgrade + electrical power intervention
- SARA: banned to the OUTAGE due to network intervention

18h September (Friday)

Experiment activities:

~8K MC production jobs running in the system right now and merging jobs at T0 and T1s. Very few user analysis jobs.

GGUS (or Remedy) tickets :

- 4 new tickets since yesterday
 - ◆ T0: 1
 - ◆ T1: 2
 - ◆ T2: 1
- List of all open tickets [↗](#)

T0 sites issues :

- Experiencing an unprecedented slowness in removing files through SRM and gfal. In chunks of 20 files the removal takes ~3 second per file while we remember once it was about 10Hz
- to be confirmed the intervention on SRM to move to Castor gridftp internal. Tentatively it could happen during the already agreed intervention on the 22nd if version of lcg_utils/gfal going to be used by DIRAC in production will work with this new configuration.

T1 sites issues :

- CNAF failing a fraction of pilot jobs with the "Globus error 22" message due to a problem with home directory. Pilot pool account user reached the maximum number of directories that it can create on the system. This is a typical and general problem with pilot job submission mechanism.
- GridKA is also failing a small fraction of pilot jobs with the "Globus error 10" message, that might indicate some vomns certificate not updated.

17h September (Thursday)

Experiment activities:

~7K MC production jobs running in the system, out of them 1K jobs for EGEE09 preparation. Few user analysis jobs. LHCb is using all Grid resources available and it is not longer limited by the OS.

GGUS (or Remedy) tickets :

- 4 new tickets since yesterday
 - ◆ T0: 0
 - ◆ T1: 0
 - ◆ T2: 4
- List of all open tickets [↗](#)

T0 sites issues :

T1 sites issues :

- pic was a problem with the publication of one of the queue (LHCb jobs were matched to an atlas dedicated queue)
- SARA still down and still banned. Any news on locality issue?
- RAL: has been fully reintegrated in the production mask after verification that all SEs were working happily with LHCb.

T2 sites issues :

- Still keeping the debugging activity at several large sites. We escalate a couple of tickets for Manchester and ACAD that were not progressing since three days.

16th September (Wednesday)

Experiment activities:

DIRAC has been finally certified to run in SL5 and now LHCb is changing their Configuration Service to include also SL5 resources otherwise unusable. No large activities going on in the system.

GGUS (or Remedy) tickets :

- 0 new tickets since yesterday
 - ◆ T0: 0
 - ◆ T1: 0
 - ◆ T2: 0
- List of all open tickets [↗](#)

T0 sites issues :

The proposed 22nd of September as date for the CASTOR intervention (upgrade to 2.1.8-12) is fine with LHCb

T1 sites issues :

- SARA: Status of the investigation on the locality returned by SRM? We know SARA is in downtime all day long
- RAL downtime extended. The site is out of our production.
- pic SL5 failing all jobs to ce08.pic.es.

Got a job held event, reason: Unspecified gridmanager error

People notified directly.

- GridKA SL5 is running fine.

T2 sites issues :

15th September (Tuesday)

Experiment activities:

Running smoothly in the Grid various MC production and relative merging activities. Many sites completely moved or moving to SL5 while Dirac is still pending the certification to support also in production the new

17h September (Thursday)

OS (only on 64bit)

GGUS (or Remedy) tickets :

- 5 new tickets since yesterday
 - ◆ T0: 1
 - ◆ T1: 0
 - ◆ T2: 4
- List of all open tickets [?](#)

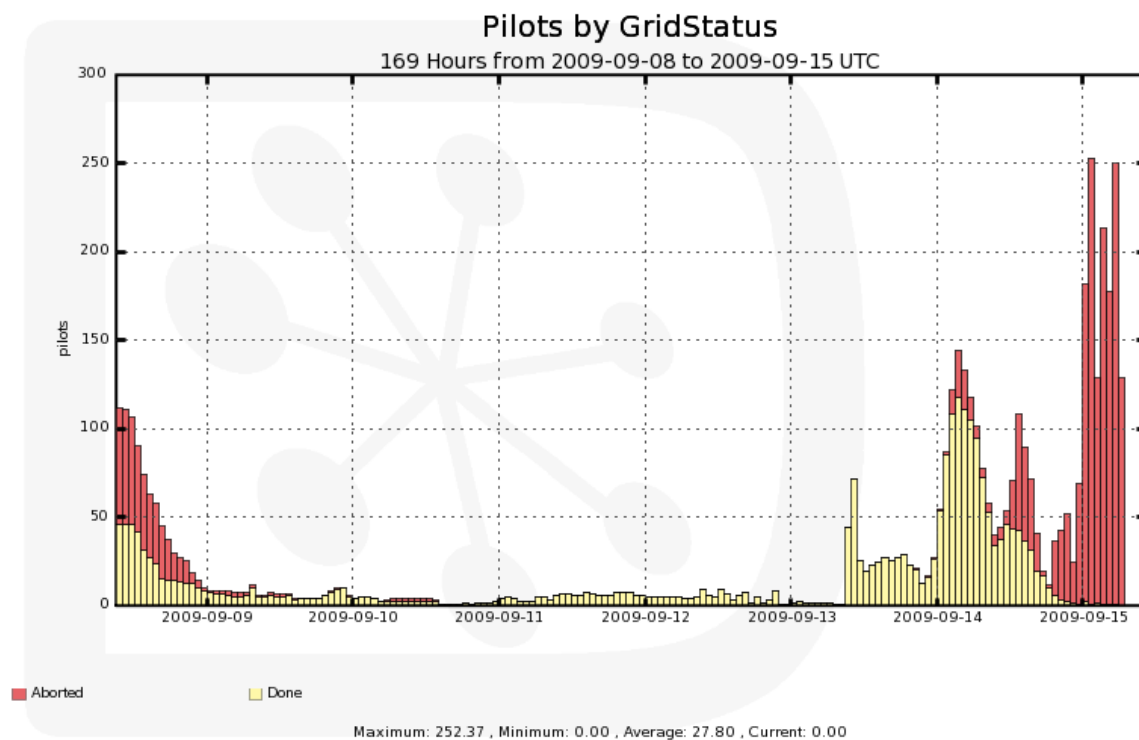
T0 sites issues :

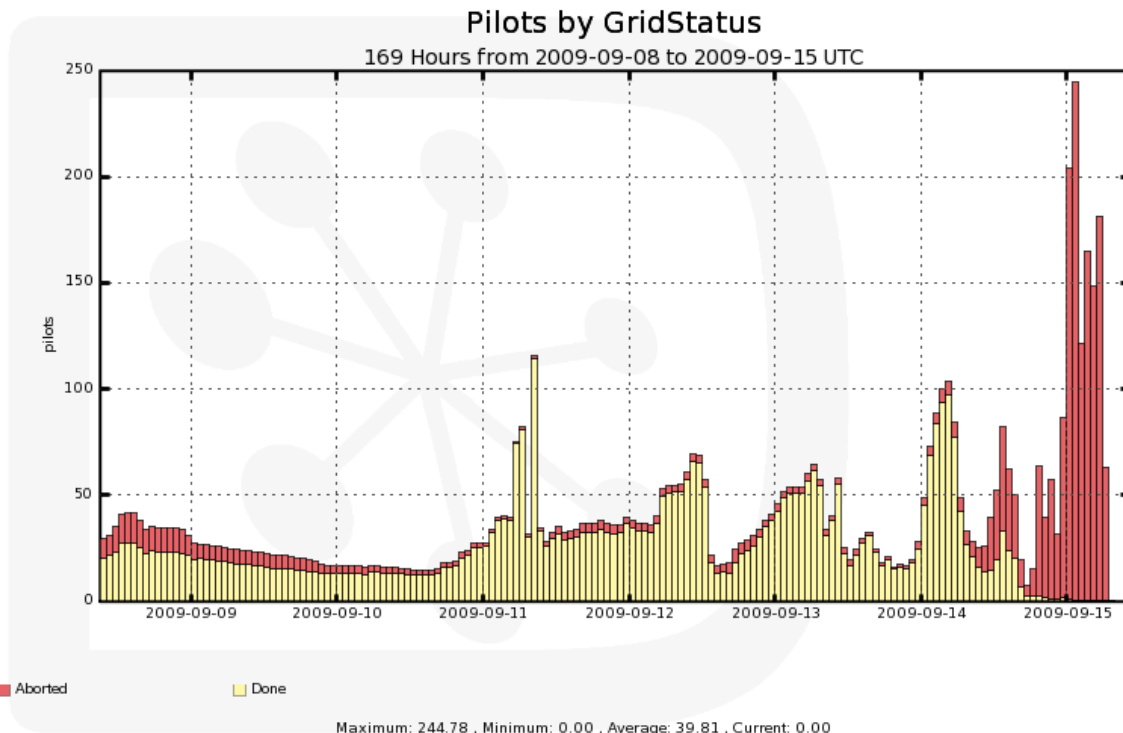
One faulty WN failing LHCb SAM jobs

T1 sites issues :

T2 sites issues :

Manchester QMUL GRIF and other centers are failing a huge amount of pilot. GGUS tickets open against the top 3 sites. In these pictures the aborting pilots at Manchester and QMUL





14th September (Monday)

Experiment activities:

Several active MC production going on on the system. During the week end an issue with the DIRAC WMS (overloaded?) and jobs stuck in Received status despite the large amount of pilot jobs on remote sites. Now running ~7K jobs concurrently in the grid. Debugging many T2 sites with large fraction of pilots aborting. MC09 data validation also going on.

GGUS (or Remedy) tickets:

- **14** new tickets since Friday
 - ◆ T0: 0
 - ◆ T1: 2
 - ◆ T2: 12
- List of all open tickets [↗](#)

T0 sites issues :

Noticed the network issue this morning affecting some services at CERN like castor instances and BDII. This has introduced some perturbation in various activities.

T1 sites issues :

pic: A network issue (apparently different than the one at CERN) to pic this morning making all DIRAC web portal services unavailable.

pic: Stalled lhcb user jobs at pic

CNAF: open ticket to Storm people for a permission on a directory. SARA issue of UNAVAILABLE status is blocking the MC09 validation trusting DIRAC the status of the file as reported by SRM

T2 sites issues :

at some sites there are failing a lot of pilot jobs

11th September (Friday)

Experiment activities:

- Overnight a lot of production managed to clear and then activity dropped down. This morning they have been extended and number of jobs are rumping up
- Data checking of all done productions at Tier1 spawning the problem at SARA (see below)

GGUS (or Remedy) tickets since yesterday:

- 5 new tickets since yesterday
 - ◆ T0: 0
 - ◆ T1: 2
 - ◆ T2: 3
- List of all open tickets [↗](#)

T0 sites issues :

Compiled a detailed report on the WMS issue reported last week available here: [↗](#)

T1 sites issues :

SARA all MC09 data file not available (GGUS open).
CNAF moving also M-DST space to M-MC-DST space token and preparing to migrate to SL5.
CNAF many pilot jobs aborting.

T2 sites issues :

at some sites there are failing a lot of pilot jobs

9th September (Wednesday)

Experiment activities:

- Running steadily about 15K jobs in the system serving various MC physics production. The WMSes at CERN seem to have recovered since yesterday despite nothing has changed LHCb side. (see CERN issue).

GGUS (or Remedy) tickets since yesterday:

- 4 new tickets since yesterday
 - ◆ T0: 1
 - ◆ T1: 1
 - ◆ T2: 2
- List of all open tickets [↗](#)

T0 sites issues :

WMS did not processed many thousand of jobs whose status had to be changed since a week or so. To be investigated (there is room for another bug of the gLite WMS software) but it looks that now the WMS at CERN are happily processing the huge amount of LHCb jobs. What's happened ?

T1 sites issues :

- CNAF: migrated DST space token to MC-DST space token
- IN2p3: some user data not available: dodgy disk server?
- IN2p3: some files hadnot tURL returned because of the following:

```
2009-09-08 18:55:47 UTC Wrapper_5166473 ERROR: SRM2Storage.__gfal_exec: Failed to perform gfal_tu
[SE][StatusOfGetRequest] httpg://ccsr.in2p3.fr:8443/srm/managerv2: org.xml.sax.SAXParseException
XML character (Unicode: 0x2) was found in the element content of the document.
```

suspicious is some libraries on some WN being not the file it self a problem.

8th September (Tuesday)

Experiment activities:

- After the problems reported yesterday of increasing number of sites flooded , the system has been drained and then the activities have been resumed. The investigation on the problem boiled down a new bug for the gLite WMS and highlighted an old bug that hasn't been fixed yet. Summary of the investigation carried yesterday is available on the LHCb e-logbook [?](#)
- Running about 15K jobs in the system serving various MC physics production and the system seems to run at the edge as during the week end..

GGUS (or Remedy) tickets since yesterday:

- 2 new tickets since yesterday
 - ◆ T0: 0
 - ◆ T1: 1
 - ◆ T2: 1
- List of all open tickets [?](#)

T0 sites issues :

Around noon one of the cluster head nodes of the Database serving CASTOR suffered an h/w problem and had to be rebooted. Some users were affected by this shortage of service

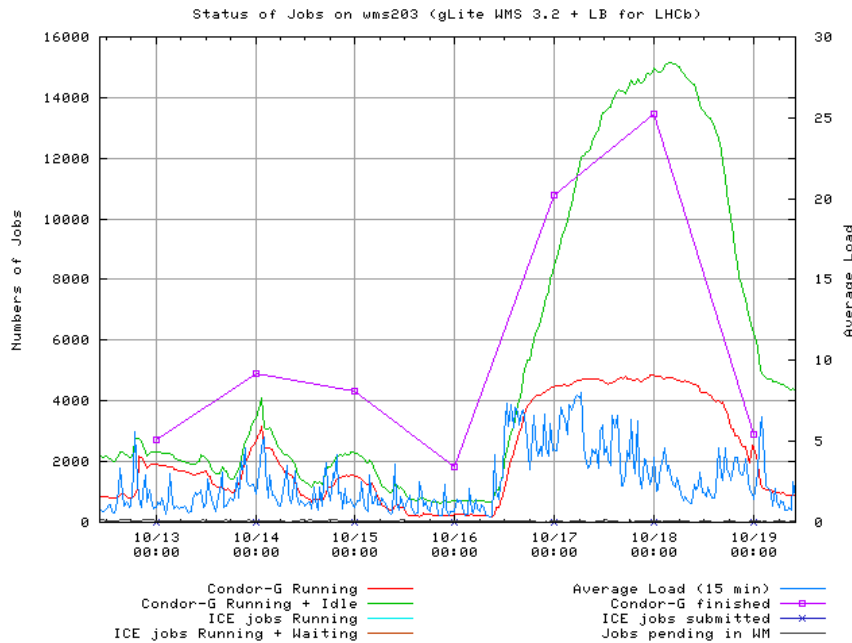
T1 sites issues :

- CNAF: shared area instability causing production jobs to fail.

7th September (Monday)

Experiment activities:

- Pilot job submission has been stopped temporarily (and then MC production) to allow the system to drain. In this picture the number of jobs treated by one instance of WMS at CERN



Many affected sites (T1 and T2) seem indeed (also accordingly previous GGUS tickets open) to report wrong information about the number of effective jobs waiting/running in the local queue and causing so the WMS ranking expression to be always OK and the site attractive. This turned into sites literally flooded by LHCb. A symptomatic case is for example Manchester (GGUS open). At other sites like pic or CNAF the problem seems to be introduced when they moved to VOVIEWs per FQAN for which they also introduced wrong configuration coming with YAIM. Before chasing this issue at all sites we would like to understand whether this has something to do with the recent upgrade of the WMS at CERN or rather with the upgrade that in many sites took place end of August.

GGUS (or Remedy) tickets since yesterday:

- 1 new tickets since yesterday
 - ◆ T0: 0
 - ◆ T1: 0
 - ◆ T2: 1
- List of all open tickets [↗](#)

T0 sites issues :

- Since about 6:10 this morning we have a problem affecting all LHCb CASTOR instances and fixed at 7:45 as traced back to a database problem.

T1 sites issues :

T2 sites issues:

Manchester: open a GGUS for debugging this issue inclose touch with Alessandra Forti presenting the worse case (8K jobs reported by Dirac locally queued and just 11 reported by the BDII)

4th September (Friday)

Experiment activities:

- Still smoothly running in the system various MC productions (Simulation + reconstruction + merging @T1s [↗](#)) and user distributed analysis (12-13K) jobs steadily running in the syste
- Testing of SL5 is also proceeding in parallel

GGUS (or Remedy) tickets since yesterday:

- 1 new tickets since yesterday
 - ◆ T0: 0
 - ◆ T1: 0
 - ◆ T2: 1
- List of all open tickets [↗](#)

T0 sites issues :

- Issue with CASTOR at CERN (moving data in)--> **understood** to be related to sick diskservers
- vomscert problem at CERN --> **fixed** by installing it also on SL4 WN (where LHCb does still run)

T1 sites issues :

- CNAF: in touch with CNAF people to change locally the space tokens definition

3rd September (Thursday)

Experiment activities:

MC production on going very smoothly under the grid point of view. Right now 14K jobs concurrently running in the system

GGUS (or Remedy) tickets since yesterday:

- 3 new tickets since yesterday
 - ◆ T0: 1
 - ◆ T1:1
 - ◆ T2:1
- List of all open tickets [↗](#)

T0 sites issues :

- Issue with CASTOR at CERN (moving data in) open/create error: Device or resource busy. Investigation on going through GGUS ticket
- SAM critical test jobs failing at CERN because of a voms server certificate not properly updated

T1 sites issues :

- CNAF transfers failing: was because we exhausted the disk space on the destination space token. The problem is that this space token was not the right one and this triggered a new issue: migrating all MonteCarlo data stored so far in these space tokens (as matter of fact supposed to host **real** data) to MC space tokens as trusted by LHCb (that provide up to 130 TB).In details the idea is to map the space token definition on StoRM at CNAF as follows:

LHCb_M-DST --> LHCb_MC-M-DST and LHCb_DST --> LHCb_MC-DST and the real data sace tokens to point to another path accordingly the LHCb conventions.

- pic: BDII wrongly publishing the number of jobs (this was the cause of the huge amont of jobs going to pic). This seems to be due to the VOView that do now rely on FQAN and are not unique per VO
- CNAF: similar problem as pic. VOViews per FQAN altering the view of LHCb in terms of jobs running and then the ranking expression that is always very attractive (flooding sites then)

2ndSeptember (Wednesday)

Experiment activities:

About 15K jobs concurrently running in the system for various MC productions.

GGUS (or Remedy) tickets since yesterday:

- 6 new tickets since Friday
 - ◆ T0: 1
 - ◆ T1:4
 - ◆ T2:1
- List of all open tickets [↗](#)

T0 sites issues:

We are seeing a degraded performance of transfers to CERN with transfers failing with errors like (no GGUS yet open):

```
TRANSFER error during TRANSFER phase: [GRIDFTP_ERROR]

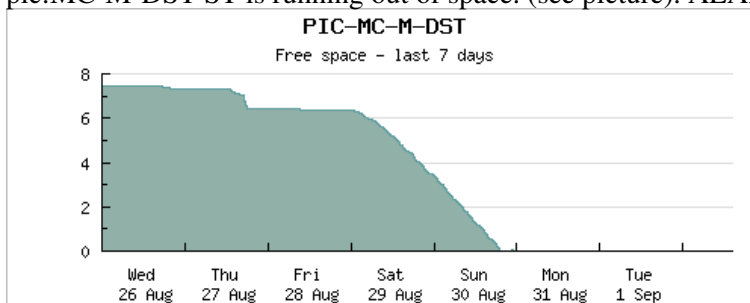
globus_ftp_client: the server responded with an error 500 Command
failed. : open/create error: Device or resource busy
```

T1 sites issues :

- The data integrity check that LHCb runs after each production is over to commission it, has pointed some data at RAL being UNAVAILABLE according SRM. This was due to some dodgy disk server. Fixed!

```
<font><font size="2">RAL_MC-DST      srm://srm-lhcb.gridpp.rl.ac.uk/castor/ads.rl.ac.uk/prod/lhcb
```

- LHCb would kindly ask also IN2p3 and GridKA (GGUS ticket sent) to evaluate the possibility to install on their WN's d-Cache clients > 1.9.3 as requested at pic and NL-T1. The gLite m/w stack comes with a very old version of them and sites should rather to get in touch with dCache developers to pick up most recent ones.
- pic:MC-M-DST ST is running out of space. (see picture). ALARMING system temporarily disabled.

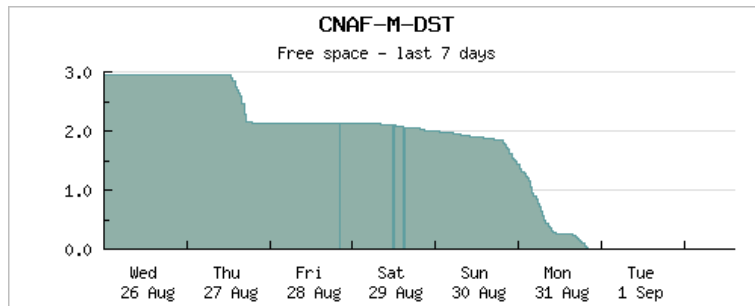


- CNAF: M-DST space token is full and for this reason further FTS transfers are failing with

```
</img><im
globus_ftp_client: the server responded with an error
500 500-Command failed. : callback failed.
500-globus_xio: System error in write: Disk quota exceeded
```

500-globus_xio: A system call failed: Disk quota exceeded

500 End.



1st September (Tuesday)

Experiment activities:

MC production on going very smoothly under the grid point of view. Over the week end roughly produced (and registered correctly) about 50 million events (118K jobs since Friday). Going to ramp up further.

GGUS (or Remedy) tickets since Friday:

- 2 new tickets since Friday
 - ◆ T0: 0
 - ◆ T1:0
 - ◆ T2:2
- List of all open tickets [↗](#)

T0 sites issues:

T1 sites issues :

iHidden="false" UnhideWhenUsed="false" Name="Medium List 2 Accent 4"/> <![endif]-->

Issue at IN2p3 reported in August about SAM jobs failing because of memory exceeding in the Q-queue. This issue pointed out other relevant sub-points rather regarding VO ID-Card and memory requirement more in general (to be escalated to GDB most probably)

1. It is not relevant to set requirements on physical memory that but rather to virtual memory that is effectively used. The physical memory is a fabric specific matter and nothing to do with end-users.
2. It is not conceivable that short queues (it seems to be the case at Lyon) had correlated with low memory hardware. We understand that sysadmins prefer to put old h/w on dedicated queues but it is not right the assumption that short jobs are correlated with low memory demanding jobs.
3. We understand that we can drive our memory requirements via specific JDL but we also found sites that publish 0 MB. This might steer to a waste of resources not matching requirements and the reason why LHCb is reluctant in doing so.

-- Main.RobertoSantinel - 2009-09-30

This topic: LHCb > ProductionOperationsWLCGSep09Reports

Topic revision: r3 - 2010-01-29 - unknown



Copyright &© 2008-2020 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

Ideas, requests, problems regarding TWiki? Send feedback