

# Table of Contents

<b>October 2009 Reports.....</b>	<b>1</b>
30th October 2009 (Friday).....	1
29th October 2009 (Thursday).....	1
28th October 2009 (Wednesday).....	1
27th October 2009 (Tuesday).....	2
26th October 2009 (Monday).....	3
23rd October 2009 (Friday).....	3
22nd October 2009 (Thursday).....	4
21st October 2009 (Wednesday).....	5
20th October 2009 (Tuesday).....	6
19th October 2009 (Monday).....	6
16th October 2009 (Friday).....	8
15th October 2009 (Thursday).....	9
14th October 2009 (Wednesday).....	9
13th October 2009 (Tuesday).....	10
12th October 2009 (Monday).....	11
9th October 2009 (Friday).....	11
8th October 2009 (Thursday).....	12
7th October 2009 (Wednesday).....	13
6th October 2009 (Tuesday).....	14
5th October (Monday).....	14
2nd October (Friday).....	15
1st October (Thursday).....	16
TEMPLATE.....	17
Date (day).....	17

# October 2009 Reports

To the main

## 30th October 2009 (Friday)

### Experiment activities:

- Production activities low.

### GGUS tickets

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

### T0 sites issues:

### T1 sites issues:

### T2 sites issues:

- Shared area problems.

## 29th October 2009 (Thursday)

### Experiment activities:

- Production activities have increased. Running a few thousand jobs.

### GGUS tickets

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

### T0 sites issues:

### T1 sites issues:

- RAL: Transfer problems. FTS transfers are extremely slow. Site has not acknowledged ticket yet.

### T2 sites issues:

## 28th October 2009 (Wednesday)

**Experiment activities:** \* Small scale production ongoing.

### GGUS tickets

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

### T0 sites issues:

### T1 sites issues:

- RAL LHCb CASTOR now configured so that DNs only mapped to LHCb accounts. This workaround has solved this particular data access problem.

### T2 sites issues:

## 27th October 2009 (Tuesday)

### Experiment activities:

- Empty system right now.
- Agreed on the LFC intervention for Thursday 29th Oct - 10AM. It will be transparent for prod-lfc-lhcb-central.cern.ch (master) and imply 30 minutes downtime for prod-lfc-lhcb-ro.cern.ch (read-only)

### GGUS tickets

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

### T0 sites issues:

- Hammering of the LSF master (from time-left utility) with too many "*bqueues*" requests: DIRAC experts will look at the code.

### T1 sites issues:

- CNAF CASTOR issue finally fixed.
- CNAF StoRM glitch yesterday affecting jobs accessing input files.
- IN2p3 issue of long jobs killed: understood to be a real underestimation of the job length for a couple of special production. No further action is required from the site but still worth to spot the coupling Physical Memory-Queue Length at Lyon (that does not make much sense)
- RAL banned yesterday for some unscheduled down. For the problem of multiple-VO users, in the interim, at RAL are looking into using specific grid map files for each instance. So, for example, the LHCb instance grid-map files will only contain entries for people in LHCB VO, plus OPS and DTEAM.

### T2 sites issues:

## 26th October 2009 (Monday)

### Experiment activities

- No scheduled activities going on at all (only SAM and user jobs).
- Part of the LHCb SAM suite (actually the one with the most relevant results, running LHCb specific application tests) is not published since a long while because using an out-of-dated version of sam clients. Upgraded the clients on DIRAC fixed the problem and new results should flow soon.

### GGUS tickets

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

### T0 sites issues:

- ce132 and ce133 (supposed to point to SL5 resources) have been found by Dirac SAM jobs to be compatible with slc4\_ia32\_gcc34,slc4\_amd64\_gcc34 only. This has to be clarified.

### T1 sites issues

- CNAF the configuration issue of RDST and RAW space tokens (reported the 15th of October) is still affecting transfers of data in FAILOVER system.
- Many jobs killed by the batch system because exceeding Memory. A close investigation revealed these jobs run on short queues and this may suggest that the BQS batch system does not report properly thenormalized CPU time left to the pilot which pulls longer than allocated time tasks. Also to be clarified once again why queue lengths are so strongly coupled (at IN2p3) with the memory of the machine.
- pic: lack of space on LHCb\_MC-\*-DST tokens forced to change the mapping of T2 associated sites from "es" to "ch" to avoid keeping sending data at pic.
- RAL: the non-VOMS awareness of CASTOR is at the origin of the problem observed for a user with more than one VO affiliation. CASTOR does not know how to resolve the VO a given DN belongs too because does not check VOMS attributes and only maps user DN statically.

### T2 sites issues:

mainly shared area issues.

## 23rd October 2009 (Friday)

### Experiment activities

- MC production (6K jobs)+ stripping and reconstruction work flow tests. Awaiting for FEST data coming from the pit...
- A CREAM CE campaign test against all 11 advertised endpoints is started (DIRAC submitting via gLite WMS). Many endpoints found to be not working, systematic GGUS tickets will be submitted against affected sites. This is the list of CREAM CE under tests (available ones)

LCG.Bologna-CREAM.it -BAD
LCG.GRIDKA-CREAM.de -OK
LCG.Glasgow-CREAM.uk -OK
LCG.PDC-CREAM.se -BAD
LCG.PIC-CREAM.es -OK
LCG.PSNC-CREAM.pl -OK
LCG.Padova-CREAM.it -BAD
LCG.Pisa-CREAM.it -BAD
RAL-CREAM.uk -OK
LCG.SPBU-CREAM.ru -BAD
LCG.Torino-CREAM.it-BAD

- New rank expression in place in LHCb that should awards large sites rather than smallish-free sites. This should fix a lot of problems (like under utilization of CNAF) unless sites publish wrong/misleading information (for which there is not much to do). The expression looks like :

Rank = ( other.GlueCEStateWaitingJobs == 0 ? ( other.GlueCEStateFreeCPUs \* 10 / other.GlueCEInfoT

- Yesterday the FEST reconstruction and stripping production was launched and jobs were starting to complete successfully. The data upon which these jobs were running are old FEST weeks data having not yet resolved some issues at ONLINE system (as reported yesterday)

### GGUS tickets

- 8 new tickets since yesterday
  - ◆ T0: 1
  - ◆ T1: 0
  - ◆ T2: 6
- List of all open tickets [↗](#)

### Service issue:

We have experienced a problem submitting GGUS tickets this morning.

### T0 sites issues:

- File access issue this morning (10% of jobs failing opening files some otherresolving the tURL). May be a transient problem on CASTOR (remedy ticket open)

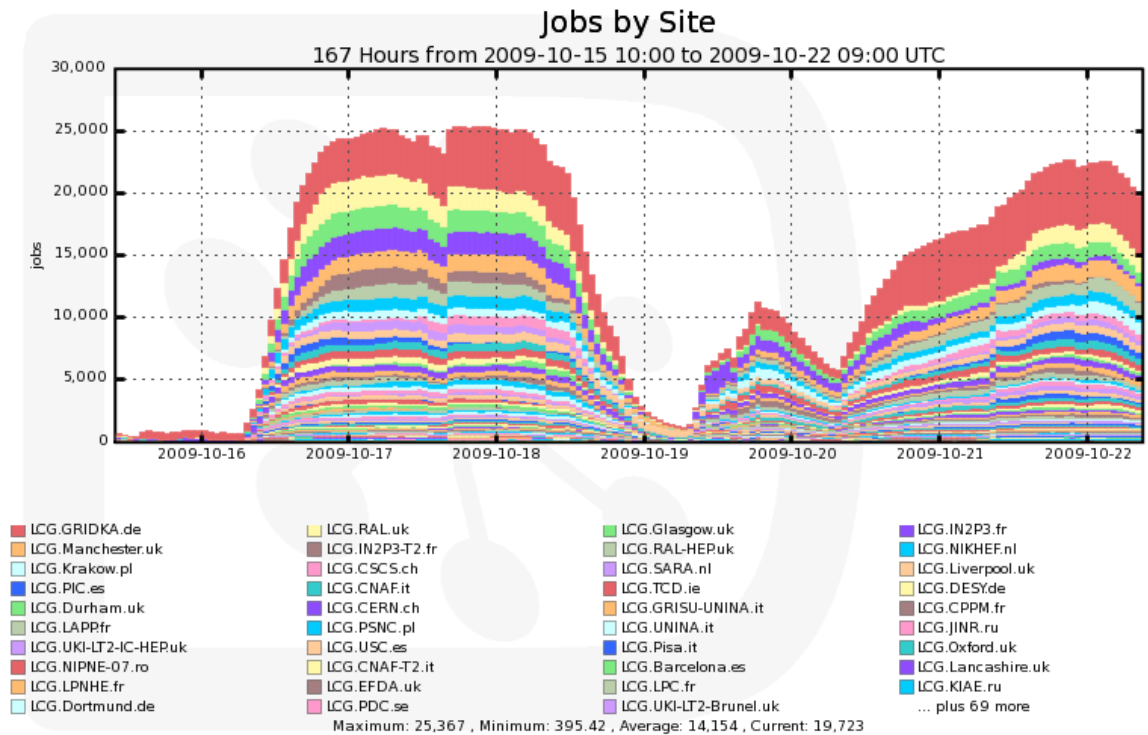
### T1 sites issues

- CNAF:issue with 60 files left in an unresolved status: looks a wrong configuration issue and mapping of pools and space tokens....still open

## 22nd October 2009 (Thursday)

### Experiment activities

- MC production and merging jobs + user analysis running in the system at the moment:18Kjobs. A profile of the last week activity in this picture
- FEST week. Data to be sent from ONLINE to CASTOR as soon as data mover problems at ONLINE are fixed



## GGUS tickets

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

## T0 sites issues:

- none

## T1 sites issues

- CNAF:issue with 60 files left in an unresolved status. Still under investigation

## 21st October 2009 (Wednesday)

### Experiment activities

- Huge amount of MC production and merging jobs + user analysis running in the system at the moment (20K)
- Run yesterday the stripping over the reconstructed data of the last FEST week. Apart 96 jobs failed at pic (because data was not there, in turn because of space problem) all jobs run at 100% efficiency at all sites (dCache sites included).

## GGUS tickets

- 1 new tickets since yesterday
  - ◆ T0: 0
  - ◆ T1: 0
  - ◆ T2: 1

22nd October 2009 (Thursday)

- List of all open tickets [↗](#)

#### T0 sites issues:

- none

#### T1 sites issues

- CNAF: issue with 60 files left in an unresolved status us under investigation by CNAF and CERN CASTOR people.

## 20th October 2009 (Tuesday)

**Experiment activities** About 8K jobs in the system for various samples to be simulated plus 2K jobs from users.

#### GGUS tickets

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

#### T0 sites issues:

- none

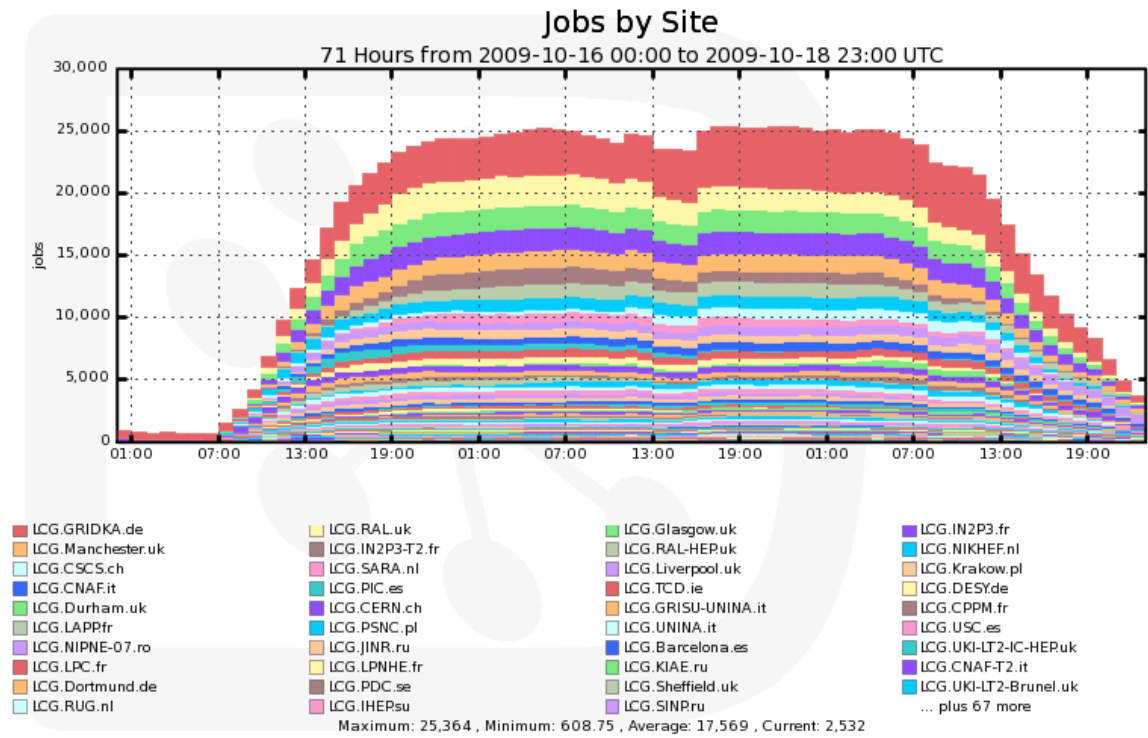
#### T1 sites issues

- GridKA: reported 6K running+ 25K jobs queued on the local batch system during the hard production period last w/e. This has been found to be due to a wrong information advertised by the GIP in the BDII at GridKA causing then (erroneously) extra pilot to be submitted there despite the protection put in place via elaborate ranking expression by DIRAC. This is most likely due to the very old version of the LCG-CE installed there (3.1.21 and/or 3.1.27) and then, under heavy load, the wrong report of the information to the IS appearing and disappearing.
- IN2p3: issue with one CE wrongly publishing and then attracting too many pilots in. The CE has been promptly removed from production by IN2p3 people
- pic:ticket wrongly open because overlooked the intervention on dCache scheduled for today.

## 19th October 2009 (Monday)


#### Experiment activities

- In the week end we had a very huge MC production activity that attained a new record in terms of concurrently running jobs (25000 between analysis and MC productions). In the pics: the increasing number of jobs on Friday night the evident plateau maintained during the week end and the slope descending since



yesterday.

Corresponding to this impressive activity (that finished on Sunday night) the gLITE WMS at CERN was suffering a lot showing the number of idle jobs increasing with the increased overall load put on the server. This is clearly shown in this pic for one of the WMS at CERN that also reports how now the status is back to normality.



### GGUS tickets

- 5 new tickets since Friday
  - ◆ T0: 0
  - ◆ T1: 0



◆ T2: 5

- List of all open tickets [↗](#)

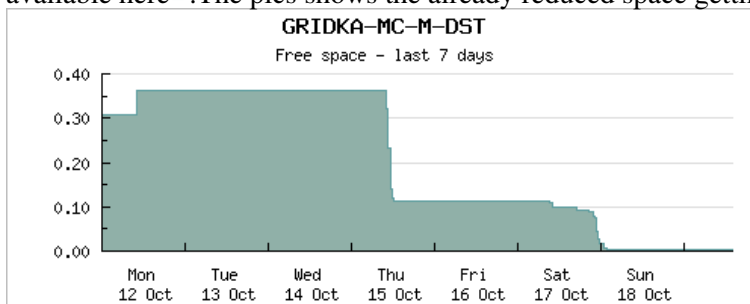
#### T0 sites issues:

- Intervention on disk-servers for LHCb to move to SL5. Agreed to run it in a scheduled slot by concentrating all disk servers in a pool to be upgraded in the same day rather than jeopardizing potential file access problems across multiple pools. We suggested Miguel to prepare a CASTORLHCB downtime for Thursday the 29th of October being this week a FEST week and some activity may overflow beginning of the next week. In that same slot also the ORACLE h/w intervention has to take place. Agreed on 29th a 8h. downtime on CASTORLHCB.

#### T1 sites issues

:

- CNAF: reported by local contact persons that despite the large number of allocated job slots on the dedicated T2 farm, LHCb did not run as expected. Most probably an issue with the ranking expression used in DIRAC that is not optimized for sites using fair share (t.b.i)
- CNAF: still not fixed the problem on 60 files reported last week
- pic and GridKA: LHCb\_MC\_M-DST space tokens full and not more transfers can take place. A monitoring page to check problematic space tokens compared with LHCb requirements is now available [here](#) [↗](#). The pic shows the already reduced space getting completely full at GridKA



#### T2 sites issues:

Shared area issues and SQLite file access issues.

## 16th October 2009 (Friday)

#### Experiment activities

- Increasing number of MC production (and merging at T1's) jobs running in the system (ramp up). Now in the system ~7500 production jobs and ~4000 user jobs concurrently running.

#### GGUS tickets

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

#### T0 sites issues:

ce130 and ce131 test jobs run happily and now the two new CEs are part of the production mask in LHCb

#### T1 sites issues:

- NL-T1: Problem for users trying to access data ("Permission denied"). Fixed (authentication problem due to the migration to ldap based system for all user mapping)
- RAL: SAM jobs failing at RAL accessing shared area (quattor issue, fixed!).
- CNAF: Issue this morning with StoRM (not responding at all). Fixed!

## 15th October 2009 (Thursday)

### Experiment activities

No much activity to report. Few MC production running. Total of 500 jobs in the system. The stripping rerun yesterday went fine at NL-T1 and IN2p3 but it has reported failure at GridKA still with jobs hanging in connection. LHCb Shifters pursuing that.

### GGUS tickets

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

### T0 sites issues:

There was a problem with a stuck server in LHCBDATA preventing some users to access data. Open (a bit tough) ALARM ticket yesterday evening.

### T1 sites issues:

- CNAF: Issue with a certificate on the CASTOR disk server at CNAF preventing few remaining transfers of RDST from last FEST week to be uploaded. Fixed very shortly. However ~ 60 files have entered a state which can not be resolved from the client. Passed the list to CNAF people.
- RAL: because of the faulty database restore we isolated as definitely lost ~200 files that have to be corrected on the various LHCb catalogs. Many more were lost during the period between the 24th of September and the 4th of October but those were just intermediate files that do not need to be corrected (already merged, users do not see them).
- GridKA still valid the problem (GGUS originally open for that must be used to track this problem down) of hanging connection
- GridKA: opened yesterday ticket about failure in listing directories (opened yesterday) is still valid.

Follows a breakdown of yesterday's stripping exercise:

site	OK	Failed
IN2p3	242	41
CNAF	329	53
NL-T1	478	58
pic	304	81
RAL	371	61
GridKA	20	543

## 14th October 2009 (Wednesday)

### Experiment activities

Re-running stripping process. After increasing the number of movers to 1500 per pool at SARA and GridKA and IN2p3 things seem a bit better and watch dog killing stuck jobs problem is not reported.

## GGUS tickets

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

### T0 sites issues:

CASTOR was not available at all. ALARM ticket open. Problem reported on the IT Service status board. Scheduled intervention on LHCBR today affecting our LFC and BK.

### T1 sites issues:

- Requested NL-T1 to clean up some old data being apparently not possible from remote.
- Failure to list directories on GridKA dCache SRM

### T2 sites issues:

shared area issues

## 13th October 2009 (Tuesday)

### Experiment activities

- 2.5K jobs running in the system right now between production/merging and distributed user analysis activity.
- Decided to relaunch some stripping production activities to demonstrate and to debug (or to verify the solution adopted about) the problem of hanging connections at the 3 dCache faulty sites (NL-T1, GridKA, IN2p3). These jobs will have 10 input files to be open/processed and we invite all site admins to follow closely their servers on the storage systems to check the LHCb activities (dcap movers, firewall, load etc.etc). As soon as they notice something please get in touch both directly with LHCb and trace this down in already open tickets. This is by far the largest issue we are currently running at dcache sites.
- Discussed the possibility to reshuffle the allocation of space tokens at sites. Sites provided (or are close to provide) 2009 pledged resources. Nonetheless this is not enough if we want to keep at least 3 replicas for DSTs. Real data size is not clear but some disk server now in (M)DST could be moved to MC-(M)DST tokens. Decision not taken yet.

## GGUS tickets

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

### T0 sites issues:

A lot of pilot jobs are failing to be submitted through ce128 and ce129 that still represent the only way to get SL5 batch nodes at CERN.

### T1 sites issues:

- Problem deleting a directory on StoRM at CNAF
- Issues at NL-T1, GridKA and IN2p3 were the efficiency in accessing data is below 50% (jobs mainly killed by the watchdog after being correctly found to be stalled).

Fair to report that at NL-T1 ,after they have increased the number of movers per pool, the problem is not there since yesterday.

## 12th October 2009 (Monday)

### Experiment activities

- Few jobs of various still remaining MC production running in the system now and few (K jobs) of user analysis. Validation of these production on going. During the week end the disk space was filling quickly at various T1s MC-M-DST space tokens (at the pace of ~5TB/day) as reported here [☞](#). Only at few sites 2009 pledge resources were allocated (pic) while notifications have been sent to all sites that did not (GridKA, RAL, CNAF, IN2p3). Either reducing the activity or reducing the number of copies to be replicated (currently 3) as a temporary solution in case requested resources do not

### GGUS tickets

- 4 new tickets since Friday
  - ◆ T0: 0
  - ◆ T1: 3
  - ◆ T2: 1
- List of all open tickets [☞](#)

### T0 sites issues:

### T1 sites issues:

- Disk space issues at RAL IN2p3 GridKA and CNAF
- Problem of hanging connection understood at all dCache sites reported on Friday and due various reasons described in the tickets open after the meeting

And this is the summary of this issue at dCache sites:

SARA:

It was an oops on their side. Although they allow for 1500 connections (dcap) to their pool for the LHCb t1d1 pools it was set to a ridiculously low number of 100!

This is by accident. This caused a huge queue to form causing our watchdog to kill the job investigation time was still happening. Fixed! This is probably the same cause of previous closed).

GridKA:

Killed almost 1500 movers from one pool. To be confirmed.

IN2p3:

it was a firewall misconfiguration issue on a file server. The jobs could not receive its

- Issues at T1's still present (IN2p3 and GridKA) while it shifted to tURL resolution issue at NIKHEF. Ticket kept open and upgraded.

## 9th October 2009 (Friday)

### Experiment activities

- Stripping + user analysis + MC & merging productions. Still impressive number of 18K jobs running concurrently in the system. Yesterday running happily EXPRESS stream reconstruction jobs for FEST data transferred from the pit last week. We had (perhaps for the first time) all type of jobs running happily in the system at the same time: stripping, reconstruction, user analysis and MC &

merging productions.

### GGUS tickets

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

### T0 sites issues:

- One of the AFS servers got nut @ ~10:45am affecting many other services (like SVN, lxbatch, DIRAC servers, TWIKI,web servers). 1 hour impact OUTAGE.
- ce128 and ce129 are failing since yesterday a huge amount of pilot jobs with error (configuration issue on the gatekeeper):

Reason = 47 the gatekeeper failed to run the job manager

### T1 sites issues:

- RAL: awaiting for the list of files that may be lost following the ORACLE outage last week in order to take a decision on whether recovering them or not.
- At all dCache sites ( **except pic** ) we are suffering a huge amount of user and stripping jobs failing because killed by the watch dog (being stuck in an open connection or failing retrieving tURL). The issue is now being investigated before contacting sites. Here the snapshot site by site about the stripping although the user analysis activity is not much different:

NIKHEF: 445 failed / 129 done + 4 completed - all failed due to the watchdog killing the job (at

GRIDKA: 192 failed / 508 done + 21 completed - failed are a mixture of TURL resolution failures a

IN2P3: 431 failed / 1322 done + 65 completed - similar to NIKHEF the watchdog kills the jobs when

this is compatible with user analysis activity whose summary for one guinea pig user is given in the following table

Site	Done	Failed
CERN	210	10
CNAF	333	0
pic	360	3
GridKA	6	404
IN2p3	157	347
NL-T1	7	410
RAL	-	-
Total	1073	1176

### T2 sites issues:

## 8th October 2009 (Thursday)

### Experiment activities

9th October 2009 (Friday)

- Second round of GGUS Alarm test results available here
- Impressive number of concurrently running jobs: 21K corresponding to various MC production requested by the physics group and Stripping of  $10^9$  MB sample. In this plot a snapshot of the last 12 hours activity in terms of running jobs.

### GGUS tickets

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

### T0 sites issues:

### T1 sites issues:

- RAL SE still not available at the time of this report compilation.
- IN2p3: PinManager outage this morning. Now SE reintegrated.

### T2 sites issues:

## 7th October 2009 (Wednesday)

### Experiment activities

- MC production and stripping + user activity going on at the pace of 11K concurrent jobs
- Rerun the GGUS Alarm test (different submitter this time)

### GGUS tickets

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

### T0 sites issues:

- UI issue yesterday fixed by hacking the environment set by the Grid UI and removing all PYTHONPATH references.
- CASTOR slowness issue: CASTOR people will generate a 2.1.8-13 release containing the fix of the multiple stagers logic and soon the srm-lhcb endpoint will be upgraded.

### T1 sites issues

:

- RAL SE still not available at the time of this report compilation.
- CNAF Storm another issue with ACLs set that were preventing to access data from pilot pool account on the WN. Fixed.!
- IN2p3: Stripping jobs (2 out of 60) failed accessing data and this seems to be due to a load problem.

### T2 sites issues:

8th October 2009 (Thursday)

- Shared area issue at smallish sites in Bulgaria Russia and Spain

## 6th October 2009 (Tuesday)

### Experiment activities

- A new Dirac release has been put in production v4r19p4 which solved the problem of bookkeeping registration
- Ramp up of MC production and Stripping activities.
- FEST: data taken until Thursday then a major issue with the ONLINE. The EXPRESS streams have not been processed so far because of a bug with DaVinci. Awaiting for the next release of DaVinci.

### GGUS tickets

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

### T0 sites issues:

- The new UI put in production yesterday (3.1.38-0) resulted to have broken some link/path preventing DIRAC services to load lfc/gfal/lcg\_utils libraries. Sophie has been informed and is looking at this problem. In the mean time being a problem that affect all VO indifferently we propose to roll back to previous version until this new is fixed.
- Issue about slowness reported a couple of weeks ago sorted into a bug [↗](#) whose the quick solution would be to roll a new version of the stager with a lower value for the idle connection timeout (from 60 to 10 seconds) hard-coded. That would already improve the worse of the scenarios bringing the deletion rate from a unworkable 0.3Hz to a more acceptable 2 Hz.

### T1 sites issues

:

- RAL SE still not available at the time of this report compilation
- CASTOR at CNAF intervention. Temporarily put out of production RAW and RDST space tokens.

### T2 sites issues:

## 5th October (Monday)

### Experiment activities:

- All MC physics productions and stripping of  $10^9$  Minimum bias production has been stopped because of a couple of LHCb internal issues: an issue with the BK service still not fixed and an incompatibility issue between Brunel application and AppConfig used in some production. As soon as these issues are fixed and the failover accumulated requests are drained these production can go ahead. Distributed analysis activity proceeding at the pace of 2K concurrent jobs.

### 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 :**

- Issue with CASTOR instance at RAL since Sunday. User jobs and also SAM jobs are reporting correctly this outage. Under investigation it seems related to a problem on the h/w beneath Oracle databases.
- Held the scheduled intervention on 3D RAC at RAL that went fine. Check the LFC over there and it is OK.
- Issue with Stripping at CNAF reported on Friday understood to be due to wrong ACL on files on StoRM.

**T2 sites issues :**

- too many pilots aborting in several sites + shared area in another couple of sites.

## 2nd October (Friday)

**Experiment activities:**

- MC production drained for the issue with the BK and now restarted (~1-1,5K jobs, ramp up). Fest not started yet but pit-CASTOR transfers are ongoing.
- Stripping has been also commissioned and it has been launched against previous large MC production (the  $10^9$  Minimum Bias production). Large number of jobs expected to come soon at T0 and T1 for data reprocessing (25K merged files to be processed).
- GGUS ticket alarm test results have to be found here

**GGUS (or Remedy) tickets since yesterday:**

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

**T0 sites issues:**

- Issue at CERN (file access slowness) found to be due to high I/O activity following a concurrent number of jobs (in turn due to the ce128 status)

**T1 sites issues :**

- Stripping jobs failing to access data at CNAF
- RAL has increased to 300 the number of slot per disk server. pic and GridKA clean up the dcap movers and no more jobs failed because stuck in opening connection. pic had open a ticket against dCache devs for that problem

**T2 sites issues :**

5th October (Monday)



- too many pilots aborting in several sites + shared area in another couple of sites.

## 1st October (Thursday)

### Experiment activities:

- As agreed at WLCG operations last week LHCb have run this morning a GGUS ALARM ticket test and collecting results
- FEST production still awaiting for a fix about options file of Davinci. Only transfers pit-CASTOR

going on the system

- Since yesterday various MC productions in the system launched that paired user analysis activities: up to 19K jobs concurrently running in the system (this morning). Because of an issue with the BK service, the output data registration was failing and 80k requests were accumulated in the failover mechanism. Decided to suspend these productions until the problem is understood.
- A peak of user activities between yesterday and today resulting in a general scalability problem at

T1's and CERN

- Stripping has been also commissioned and ready to be launched against previous large MC production (like the 1billion MB production)

### GGUS (or Remedy) tickets since Friday:

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

### T0 sites issues:

- Issue with one of the CEs (ce128) that was in a fanny state. Once restarted the daemons the situation went back to normality
- Problem with user jobs accessing data: no tUrl returned and jobs left hanging until the watchdog kills it. This is due (as seen at pic and RAL) to connection never closed from previous attempts.

### T1 sites issues :

2nd October (Friday)

- RAL: slowness opening some files in CASTOR from some user jobs: exhausted all connections due to left over pending requests
- pic: also problems reported by users in opening file through gsidcap. Not a systematic problem it was due to number of connection exhausted.
- GridKA: apparently the same problem as pic
- NIKHEF: apparently the same problem as pic

T2 sites issues :

## TEMPLATE

### Date (day)

Experiment activities

•

Top 3 site-related comments of the day:	1.
	2.
	3.

GGUS tickets

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

T0 sites issues:

T1 sites issues:

- [OPEN - since February 2009 ]

T2 sites issues:

- [CLOSED - since April 08 ]

-- GreigCowan - 02-Nov-2009

---

This topic: LHCb > ProductionOperationsWLCGOct09Reports

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



Copyright &© 2008-2019 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

Ideas, requests, problems regarding TWiki? Send feedback