

Table of Contents

| | |
|---|----------|
| Week of 090302..... | 1 |
| WLCG Baseline Versions..... | 1 |
| WLCG Service Incident Reports..... | 1 |
| GGUS Team / Alarm Tickets during last week..... | 1 |
| Weekly VO Summaries of Site Availability..... | 1 |
| Daily WLCG Operations Call details..... | 1 |
| General Information..... | 1 |
| Monday:..... | 1 |
| Tuesday:..... | 2 |
| Wednesday..... | 4 |
| Thursday..... | 6 |
| Friday..... | 7 |

Week of 090302

WLCG Baseline Versions

WLCG Service Incident Reports

- This section lists WLCG Service Incident Reports from the previous weeks (new or updated only).
- Status can be received, due or open

| Site | Date | What | Comments |
|------|--------|----------------------------------|--|
| CNAF | 21-Feb | Network outage | Report promised |
| ASGC | 25-Feb | Fire affecting most/all services | Report expected |
| CERN | 04-Mar | CASTOR 3-hour downtime | Report in place (at above WLCG SIR link) |

GGUS Team / Alarm Tickets during last week

- Tickets: see https://gus.fzk.de/ws/ticket_search.php and select "VO" - this will give all tickets, including team & alarm - or use the following link:
 - ◆ In the GGUS Escalation reports every Monday

Weekly VO Summaries of Site Availability

- ALICE
- ATLAS
- CMS
- LHCb

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

General Information

- Web archive: <https://mmm.cern.ch/public/archive-list/w/wlwg-operations>
- CERN IT status board: <http://it-support-servicestatus.web.cern.ch/it-support-servicestatus/>
- elogs: <https://prod-grid-logger.cern.ch/elog/CCRC'08+Logbook/>

See the weekly joint operations meeting minutes

Additional Material:

Monday:

Attendance: local(Harry, Jean-Philippe, Ewan, Simone, Roberto);remote(Michel, Gonzalo,Angela, Jeff, Gareth, Jeremy, Daniele).

Experiments round table:

- ATLAS - Several reports from the weekend. 1) The strange 'defective credentials' error reported as at TRIUMF last Friday was a misleading error message, the real problem being in the FTS server. 2) There are corrupted files at IN2P3 currently under investigation. Michel reported GRIF had received 2 ATLAS tickets (one a team ticket) when in fact the failing transfers were a source error at IN2P3. Simone will follow up why (GGUS ticket 46738). 3) Central DDM services stopped for a while on Sunday and needed a restart of the dashboard Apache server to restore them. 4) Four files are missing in NDGF i.e. in the catalogue but not in local storage. Retries generated lots of individual error messages - the plan is to automatically aggregate these.
- CMS reports - Following review of the CMS reports Jeremy was given the Imperial College Savannah ticket number as 107304. Michel reported problems over the weekend at GRIF with multiple CMS jobs trying to access their MYSQL database over NFS and locking up. Some half of their worker nodes had to be rebooted. This sounds like the problem recently seen at PIC.
- ALICE -
- LHCb - 1) Dummy MC production has been put on hold pending a fix of the underlying Gauss (simulation) code. 2) The FEST'09 MC production has nearly finished and this week data transfer and reprocessing at Tier 1 will start. 3) A dcache patch addressing the srm 'file unavailable' error at IN2P3 has been deployed. 4) The CERN volhcb09 dirac server regularly runs out of swap space. A larger machine has been requested. 5) The latest version of lcg_utils is giving a 'SOAP error reading token' at the CERN srmv2 end point. Seems to be related to castor client rather than server. Jean-Philippe thought this was in fact an overloaded srm server due to the server timeout setting being commented out in its config file.

Sites / Services round table:

NL-T1 (JT): Will have a 2 hours at risk downtime tomorrow with server reboots.

FZK (AP): Had worker node queue problems at the weekend for ATLAS and CMS when 2000 batch jobs bound to the same seven worker nodes following multiple submissions. Thought to be a local batch system problem.

CERN WMS (ER): The update 41 megapatch will be applied to the CERN WMS this week one by one with all to be completed by Wednesday.

GRIF WMS (MJ): The ALICE France WMS was upgraded to level 41 last weekend.

RAL (GS): Are in the middle of a CASTOR upgrade for LHCb.

AOB: from MDZ in absentia for CNAF preparing for CMS GGUS alarm ticket tests: I've just opened a ticket regarding the DNs of the alarmers entitled to send alarm tickets from GGUS.
https://gus.fzk.de/ws/ticket_info.php?ticket=46470

It looks like that at least one DN published here is not correct:

<https://twiki.cern.ch/twiki/bin/view/LCG/OperationsAlarmsPage> For example D.Bonacorsi's DN seems to be broken. INFN T1 uses those DNs to selectively enable the translation of an alarm ticket into an SMS, so we need to be completely sure that all the DNs there are correct. This is becoming urgent given the forseen compaining test for next week. Thanks Tiziana.

Tuesday:

Attendance: local(Harry, Jean-Philippe, Ewan, Simone, Andrea, Julia, Roberto, Olof);remote(Michael, Angela, Jeremy, Brian, Gareth, Daniele, Jeff).

Monday:

Experiments round table:

- ATLAS - 1) The two PANDA servers are being migrated from BNL to CERN and also from mysql to oracle. The first one is already running at CERN but still with mysql. It will now be stopped for migration to oracle so production in the Italian and CERN clouds will be temporarily stopped. After restart tests will be made then the planning to move the second instance will be made. 2) Reprocessing tests start this week to run against tape data. Sites are asked to clear their tape disk buffers then the reprocessing framework will launch prestaging jobs. This requires a modified pilot job that will look for input data on the tape disk buffers rather than on permanent disk (via the file catalogue). 3) There has been ongoing tests of srmcp from Lxplus into CASTOR which fails when the CASTOR gridftp is running in internal mode. Reason is the url can only be accessed once and this is 'used up' in computing the checksum. LHCb use gridftp in external mode for this reason but it causes server performance problems as connections are kept open. Jean-Philippe reported that lcg_cp will also have this problem once checksumming is added. Simone will check the ATLAS use case for this functionality (but we probably cannot configure a per pool gridftp).
- CMS reports -
- ALICE -
- LHCb - An ammendment from yesterday in that dummy MC production has been halted waiting for a patch to the simulation software. Is there an understanding of the SOAP error reported yesterday at the CERN srm endpoint ? Jean-Philippe reported he had talked to Remi Mollon who would like an LHCb test job. It could be a python library error. Olof asked if this was a multi-threaded application. Roberto replied that it was and that it had been running this way since before Xmas. LHCb suspect they have lost some files on a RAL disk server. Brian asked they open a ticket that he will then follow up.

Sites / Services round table:

NIKHEF (JT): Started an emergency down about 45 mins ago following a regular emergency power cut test which caused a cooling failure. They had to shut down their worker nodes so some jobs will have been lost but expect to be back in the next hour.

BNL (ME): Yesterday had authorisation problems with their GUMS server. A redundant configuration was added last week including a load balancing switch but this did not properly close TCP connections so they built up over time leading to memory problems. The switch had to be removed for now which caused a one hour GUMS downtime.

RAL (GS): Currently upgrading the CMS CASTOR instance (LHCb was successfully done yesterday).

TRIUMF (Di Qing): From the minutes (of Monday) , TRIUMF FTS server was mentioned and said for the strange 'defective credentials' problem the real problem being in the FTS server. Basically, what we had done are only that we set all channels to inactive at the beginning of the downtime to pause the channels and set all channels to active again at the end of the downtime because of the downtime of SE. And this is the correct procedure. When the transfer started again, we also noticed this problem and the following is what we found from logs:

TRANSFER error during TRANSFER phase: [PERMISSION] globus_ftp_client: the server responded with an error 535 Authentication failed: GSSException: Defective credential detected [Caused by: [Caused by: Bad sequence size: 4]]

We believe it is a bug of FTS. Or still the bug 33449, i.e, "delegation". (Gavin Mccance thinks not the delegation bug and asked for a GGUS ticket to be raised leading to a Savannah report).

Tuesday:

ASGC status and recovery planning after UPS fire (Jason Shih): (Some changes made for readability - HR)
The severe dusting and pollutant in the data center make it difficult to restore the services rapidly. Other rack mount servers able to be up and functional normally after switch to using other network segment while the services binding to the blade system might be delayed a while until having suitable places for relocating the critical grid services (sufficient power and cooling).

If we confirm the recabling tomorrow and also if able to have sufficient power in the lab, we might have chance bringing up these services in time. Mass storage system might delay a bit till we have better location hosting 100 raid subsystem and also close to the front end disk servers direct attaching with fiber channels.

In parallel, we're cleaning up the facilities as well as data center area that we become able to bring all systems up (bypassing the UPS system as procurement might delay a while) and resume validation. We're afraid the instability of the system (especially hardware problems) might affect the availability as well. 100% full running expected to be finished in 1.5 month and maximum extend up to 2 month.

AOB:

Wednesday

Attendance: local(Jean-Philippe, Gavin, Maria, Jamie, Andrea, Harry, Julia, Diana, Olof, Simone, Roberto, Ignacio);remote(Daniele, Michael, Gonzalo, Angela, JT).

Experiments round table:

- ATLAS (Simone) - yesterday expert on call (Stephane) sent test alarm to all T1s except Taipei (missed morning slot - sent today). Workflow for all ok - tickets closed. 1 thing still to understand for CERN - something? didn't get to right person. Olof - same problem as in past - SMS message was not original one from GGUS. We get a reply from operators. For some reason feed through e-groups to phone gateway still doesn't work. Get operator replies so see ticket anyway. Signed messages not forwarded to mobile phones. Upgrade announced for gateway - did it ever happen? Harry - still pending. (Test will take place every 3 months). Simone - testing of FTS in PPS: no problems with this (cures delegation problem). Fix timescale of end week - if no problems = working. 3rd point: Birger reported CASTOR ATLAS unavailable. Olof - under service report later...
- CMS reports (Daniele) - last 24H 4 new CMS Sav. & no new GGUS. Focus now on followup on CASTOR. Main impact on CMS: some files at P5 waiting for transfers. Ignacio commented on all Qs - N/S restart allowed data to flow through. n/s information not up to date - T0export pool - might see some wrong filesize (0) for entries when n/s stuck. Cleanup needed - q to CMS to provide list of files affected - in progress. T1s: 3 new tickets + 1 for T2. CNAF: some files either corrupted or missing - need to check. FZK: set of transfer errors from T1 to CIEMAT (in preparation phase hence in FZK). Some transfer errors FNAL - IN2P3: may point to lack of disk at dest? Ticket for IN2P3. 3 tickets opened since several days: transfers to FNAL from IN2P3 showing some errors due to source issues. Some d/s not moving to tape at IN2P3. T2s: closing ticket to Brazil T2 (network issue). New ticket to US T2 in Florida. Deletion request approved last week but still pending. Why? Ticket to Pisa: detailed report - h. load in dCache. Reinstallation of dCache head node foreseen. MIT:
- ALICE -
- LHCb (Roberto) - confirm what Daniele and Simone reported - LHCb CASTOR instance not available all morning. Another issue reported yesterday evening. File waiting >12h. Olof - mig streams stuck as request to vol mgr for 0 size volume. This was the original problem. Under investigation...

Sites / Services round table:

Wednesday

NL-T1 (JT) - problem with cooling yesterday. Short summary: valve should switch from primary to backup cooling water. Got stuck - short pm - longer pm will be sent. Found some holes in procedures as regular people not here. Running at about 80% capacity - will switch other WNs on after meeting. More:

Around 13:30 on 3 March, the facilities people at Nikhef did a routine test of the emergency power. The test is realistic : it is done by cutting the primary power to the chain, just as would happen during a real outage. During the test yesterday, a mixing valve got stuck in the cooling system, resulting in no new cooling water. Today we are slowly turning machines back on, and monitoring the temperature as we do so. Positive consequences of the incident are that we have improved documentation for the remote management of the cooling system.

CERN (Olof) - outage of complete castor system this morning due to switch intervention on private switches behind Oracle head-nodes -> NAS storage. Same intervention as scheduled for two days ago? No announcement of anything for today... All CASTOR DBs lost connections to storage - caused head nodes to reboot. Ignacio - RAC s/w told head nodes to reboot. Olof - once stager DBs back application automatically reconnected. N/S does not have this auto-reconnect. Had to restart N/S & vol mgr. Took most of morning up to ~12:30 - 13:00 (according to SLS). Should have been functional for expts about this time. Will write pm about this incident. Daniele - please link post-mortem to this twiki.

TRIUMF :follow up report on 'defective credentials ' problem from Akos Frohner (CERN/IT/DM): The problem has been reported to the fts-support@cern.NOSPAMPLEASE.ch list as well, where after a few mail exchanges Stephane Jezequel <JEZEQUEL@LAPP.IN2P3.FR_moz-userdefined=""> reported that the problem seems to be related to the TRIUMF SE and transfers to other storage elements were successful, so the problem is not likely to be in FTS. Since the error message above is from the transfer phase that means that the credential has been already used to contact the SRM endpoints, which supports the hypothesis that the credentials were also correct. So I do not believe it was a problem of FTS. Unfortunately we could not catch this transient error, so we could not debug the problem properly to discover the real cause.

LHC VOMS

Following the French CA change from /C=FR/O=CNRS/CN=GRID-FR to /C=FR/O=CNRS/CN=GRID2-FR secondary identities for French VO members have been added. This has been completed so far **dteam** and **atlas** on the authority of Pierre Girard. Other VOs will probably follow.

- Release (Antonio) - One important news is that a tentative date for the release of the WN on SLC5 was set for the 16th of March.
- DB (Maria) - downtime of CMS online Monday pm frm 15:00 - 17:00 due to failure of public ethernet switch. Being followed by CMS sysadmin and IT-CS. When switch back re-established DB ops without problems. ATLAS: failure of streams propagation to IN2P3. Around 13:30. Caused by misconfig of one of parameters at destination (max # processes at DB level - set to 300 - discussing with DBA team on an eventual increase as this seems to be too low.) ATLAS: some activity with Panda people. Met Monday + Tuesday: migration of this application from MySQL to Oracle. Still some work before it can be declared 'production'. ATLAS will run on INTR (with backups!) and after this will decide when to move to production.

AOB:

- Daniele - some help from CMS contact at FZK to check CERN-PPS FTS server for delegation mode race-condition bug. After 1 week of tests confident that this was not seen. Nicolo will collect info and report back.

Thursday

Attendance: local(Gavin, Maria, Simone, Jamie, Miguel, Harry, Roberto, Kors, Andrea, Stephane, Ewan, Dirk, Diana, Julia, Ignacio);remote(Gonzalo, JT, Daniele, Michael, Jeremy).

Experiments round table:

- ATLAS (Simone) - PLANNING: next week the 2nd reprocessing of cosmic data will start. Data are at T1s - basically a single copy "on the Grid" - aka outside T0. Taipei will not be accessible; FZK has scheduled downtime to split SRM f/e. Taipei - reprocessing will be run at CERN; for FZK have to see - wait for a few days? Wait for FZK to be back online? (Not at CERN...) One day downtime for storage in Lyon (Tuesday). If anyone else plans ought holler! Brian- have now enabled pache which Graeme Stewart is testing. This may well solve replication problems. (Simone - pcache will download conditions DB file and hold it there as long as possible - useful for everybody). Gonzalo - pcache - ready to deploy? Simone - yes, ask Xavier or Graeme. Michael - when will validation jobs for reprocessing arrive? Simone - this is Rod - started tests today - within some hours...
- CMS reports (Daniele) - interested in reading a full post-mortem on CASTOR events from yesterday! Notified about problem with tape migration - not going so well until ~1 hour ago. Ticket last night about this. Something fixed earlier this morning? Ignacio - restarted things in the morning but realised that some files were not picked up. For post-mortem Olof has started to work on it. Daniele - will send some more info on what we experienced. T1s: closing tickets to IN2P3 x 2. Waiting for some transfers to T3 in FNAL - needed to validate some files - done. IN2P3 waiting for some transfers to move . CNAF: new ticket as of today - low priority - like to commission a link from France to CNAF - keep open until details resolved. Ticket open to CNAF since yesterday - check in detail some files not going to tape. FZK: ticket on transferring to CIEMAT - seems to be problem with tape connection. IN2P3: transfer errors FNAL->IN2P3. Nothing really waiting - seems to be ok now(?) Space increased in tape family by nnnTB - keep monitoring for a while. T2s: waiting for data deletion in Florida- now ok. London IC and Caltech. IC: slow transfers ->RAL (Sav 107383). Related to T2? Storage instabilities - >3TB backlog. Caltech: FNAL->Caltech transfer errors (security / permission issue?) Pisa - still to resinstall dCache headnode. MIT- no reply.
- ALICE -
- LHCb (Roberto) - data quality team still has to validate fest production. Massive production will take place then. Again problem at IN2P3 with wrong status returned by SRM. Freshly transferred data reported nearline. Lionel+ investigating. Patch from dcache doesn't work? Problems with some recons at NIKHEF. Access to conditions DB - seen also with SAM test jobs for this. Failure to contact LFC file catalog to access connection string for DB. Seems to be not happening systematically. Waiting jobs problem on WMS at CERN will not be fixed by the famous mega-patch - another one is needed. Gonzalo - any news on GGUS ticket on issue with monitoring jobs (LHCb SAM tests which make sqlite access via nfs - credentials of s/w manager - eventually hang NFS). Any news?? Roberto - didn't update GGUS - will move these tests to production role. Will inform when this will be done. LHCb manage prioritize jobs in their task q. Gonzalo - inter-VO issues.

Sites / Services round table:

- RAL (Gareth):
 1. There was a reference in the meeting on Tuesday to a loss of LHCb files on a disk server at RAL. To add more information on this: On a particular disk server the 'fsprobe' test showed evidence of data corruption. All data on the server has been declared as lost, given that we can have no guarantee of its integrity. This has been followed up via LHCb contacts.
 2. Alarm ticket tests. We have, so far, received a couple of these. From the documentation I was expecting that these would contain sample scenarios and we should say what we would do in

that case. Neither ticket received so far (Atlas, CMS) has had a scenario in it. For Atlas there are scenarios on the web page - is it up to us to pick one? For CMS I could not find sample scenarios on the web. Of course, the test so far has verified delivery of, and response to, the alarm tickets which is a good thing. I am assuming it is not for as the recipients of the tickets to pick (or invent ???) a scenario - although that would be an interesting and original test regime.... Diana - "sample scenarios" - request from USAG was that tickets should have some pretend problem - not just "test ticket" - e.g. a problem like "ALICE has a problem with VO box" or "problem with LFC". CMS did this but ATLAS not?? Stephane - we opened an alarm saying this is a test. Ignacio - we didn't see ticket. Diana - person on duty at ROC CERN did not assign PRMS ticket... Stephane - one site did not respond. (Taiwan!)

3. Brian - in addition to Gareth notes we finished upgrades for ATLAS and CMS to CASTOR to 2.1.7 release and ALICE is underway. Should be out of downtime on schedule at 18:00 CERN time.

- CERN (Ewan) - megapatch on all WMS nodes. Report any problems as normal. (Gavin) SRM LHCb at CERN - Andrew reported some security errors (background) - contacted developers who proposed config change (cgsisoap errors) - propose to make this change this afternoon! Will increase load on DB. Also # threads in daemon and # connections to DB... FTS service - firmware upgrades to boxes which should be "totally transparent"
- PIC (Gonzalo) - FTS server config at T1s. Contacted by FTS admins of FZK asking to configure FZK-PIC channel in FTS server - want to check status. Reasonable that this is a general config if so can it be made more efficient? Gavin - both sites should have admin control of the channel. No general solution of how to pass DNs around. Gonzalo - VOMS role for dteam VO triggered but these kind of issues - quiet automatised way. Gavin - will ask Akos to broadcast to FTS users how to do this...
- DB (Maria) - would like to announce downtime of DB services on March 18 related to network changes. Ewan - CERN batch queues already draining - 3 week q turned off last week. Ignacio - CMS have asked that we let jobs crash during intervention.

AOB:

Friday

Attendance: local(Gavin, Harry, Olof, Simone);remote(Angela, Michael, Gareth, Jeremy, Jeff). Apologies: Daniele (CMS) - cannot attend today (due to a delay in a meeting previous to this one).

Experiments round table:

- ATLAS (by email from Kors) - ATLAS reprocessing campaign has started!

the validation tasks for the re-processing have been submitted or are about to be submitted. These are very much like the tasks that were run during the Christmas re-processing campaign. Last time these tests revealed a problem in FZK but that is now understood and fixed. So we don't expect any site to fail this time. However Taipei is down because of the fire and won't be back up in time to participate. These data have to be re-processed at CERN now.

Tests have been performed at PIC and RAL to re-process data from tape. This still uses the old release of the reconstruction software but that is not important for this test. At PIC this went very well except that the task didn't finish because we seem to have a broken tape. It is good that this happens now because it gives us the opportunity to test how to fix this. We can be sure that this will happen again. Also at RAL this test is going well although a bit slower (on purpose).

All cosmic data has been cleaned from the buffers so the same tests can start in the other T1's. Lyon will bring the files on-line manually because there is still one component missing to have pre-staging been done by the site services. These tests tell us if pre-staging works and if the buffer turn over is more or less optimal for the jobs.

The plan still is to start the real re-processing of all the cosmic data next week. We know that there is a few day shut down at FZK so they will probably start a little bit later. We don't have to remove the RAW data from the disks because Panda can now distinguish between the copy on disk and the copy on tape and be made to choose the tape copy. This gives us a fall-back in case we do have problems with reading from tape. We should hope not to need this fall back solution because with real data we won't have an extra copy on disk.

We will have 2 measures against the "hot file" problem. There will be a conditions data tar ball per run and not one for all 100 runs together as we had over Christmas. So there will be fewer jobs at the same time trying to access these data. Secondly at a few sites we will test the "pcache solution" where the conditions data will be left on the worker node after the job has finished. If the next job on that node needs the same data it will just use it and not bring a fresh copy in.

During Christmas and New Year few people were available in the sites. This time we hope you will monitor closely this effort and report any irregularity to us. We need to measure how efficiently we can do re-processing, how many cpu's we use and how long it takes. We need to know if the stage buffer matches the number of tape drives to fill it and the number of cpu's to use it. And then there are the exceptions like broken tapes, files that seem to be missing for other reasons, job crashes and so on. This may be one of the last chances to test before we need it all working for real.

ATLAS reprocessing data volumes (by email from A.Klimentov): Dear Tier-1 and Cloud Reps,

March09 reprocessing will be started next week, most probably it will be Wed/Thu. It will be announced. Right now Operations is running tasks to validate sites and tests.

Computing management proposed to do reprocessing 'from TAPE'. Files staging will be done by Panda (no files pre-staging in advance), data volume per site is in the table below.

The above scenario is under testing at PIC. More tests are scheduled for other sites.

ASGC share will be redistributed between Tier-1s, it was reported today that reprocessing at IN2P3 will be done for disk resident files.

Please, inform ADC Operations before Mon Mar 9th evening CET, if there are any constraints to run reprocessing from 'TAPE' on your Tier-1. And please verify that disk buffer is cleaned, as it was requested earlier.

| Tier-1 | DS@tape | (files/TB) |
|-------------|---------|---------------|
| BNL | 648 | (89459/140.2) |
| FZK | 200 | (31516/57.6) |
| IN2P3-CC | 269 | (34956/61.3) |
| INFN-T1 | 85 | (13009/21.5) |
| NDGF-T1 | 93 | (9510/14.6) |
| PIC | 88 | (9664/17.1) |
| RAL | 181 | (41162/69.0) |
| SARA-MATRIX | 237 | (28315/47.4) |
| TAIWAN | 99 | (12145/20.4) |
| TRIUMF | 73 | (9740/14.6) |

*) Table shows how data was distributed initially, as you are aware some data redistribution was done during Dec08 reprocessing.

ATLAS Friday news (SC): First results from the reprocessing yesterday from PIC. 200 jobs failed with unavailable files due to a broken tape. The procedure for recovering data from CERN was applied and the jobs are running and should finish soon. Tests at other Tier 1 will start soon. ATLAS considered if these tests should also be used to produce data for physics in which case they would not insist on recalling input data from tape but have decided to maintain this as a computing exercise and insist on recall from tape. ATLAS have agreed that FZK migrate to FTS 2.1 during their downtime next week (see below). Also they have been testing the FTS patch in the PPS that fixes the proxy delegation problem and are happy with it so it should be released next week.

- CMS reports - See apologies above.
- ALICE -
- LHCb - FEST09 is now running with transfers and reconstruction jobs at T1.

LHCb had yesterday a severe problem uploading data to the USER space token both at CERN and RAL with error message "globus_xio: An end of file occurred".

Open a TEAM GGUS ticket (https://gus.fzk.de/ws/ticket_info.php?ticket=46946) (but also a direct call yesterday evening to CASTOR guys being a show stopper for all users uploading data there) the problem has gone today. Most likely the solution has been the same I saw for RAL and the same as of two weeks ago at CERN : i.e. by killing all gridftp pending processes. Please note that this top priority TEAM ticket is still at the CERN-ROC to be assigned.

At RAL (spotted by SAM) the same symptoms, I open a GGUS this morning https://gus.fzk.de/ws/ticket_info.php?ticket=46950, they recovered the rogue diskserver by cleaning gridftp processes. Problem solved and verified. A more robust and generic solution for CASTOR gridftp server should be implemented. A comment by Brian Davies was that the problem is due to a gridftp server limit of 100 processes.

IN2p3: wrong locality returned by SRM seems to be the right one because they have different disk pool for reading data, importing and exporting data. A BringOnline call should be issued to get the right status.

NIKHEF: issue with some reconstruction jobs accessing ConditionDatabase, experts still working to understand with Persistency devs also involved.

Sites / Services round table:

FZK ATLAS (by email from S.Nderitu): There will be an ATLAS downtime scheduled from 09-03-2009 morning till 13-03-2009 evening. This is to facilitate the dCache split to set up an ATLAS separate dCache instance. During this period, the following will be done:

- ATLAS pools have to be reconfigured and reconnected.
- dCache DBs need to be migrated to new ATLAS dCache instance and dCache service nodes for ATLAS have to be fully setup and started.
- In parallel LFC must be updated to reflect the new SRM endpoint atlassrm-fzk.gridka.de which requires downtime for LFC as well.

The downtime includes some safety margin and time for testing, but the migration will take at least until 11-03-2009.

Production and data management will be down for the whole DE cloud during the migration. Distributed analysis at T2s will be down during the GridKa-LFC update (downtime scheduled for 09-03-2009 till 11-03-2009).

RAL: Brian reported issues with bulk deletion of ATLAS data where there were some incorrectly owned files and directories - will probably be seen at other sites. Gareth reported the CASTOR ALICE upgrade to 2.1.7 release overran with the service being restarted overnight.

AOB:

-- JamieShiers - 26 Feb 2009

This topic: LCG > WLCGDailyMeetingsWeek090302

Topic revision: r21 - 2010-06-11 - PeterJones



Copyright &© 2008-2022 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

or Ideas, requests, problems regarding TWiki? use Discourse or Send feedback