

DC06 Activities

DC06 Aims

Challenge (using the LCG production services):

- Distribution of RAW data from CERN to Tier-1's
- Reconstruction/stripping at Tier-1's including CERN
- DST distribution to CERN & other Tier-1's
- Expected to last: End of September or once 25M events high and 25M events low luminosity are fully reprocessed for further physics analysis.

Status of LHCb SE's

CERN

2007/03/12 11:42: **CERN-disk OK**

2007/03/12 11:42: **CERN-tape OK**

CNAF

2007/03/12 11:42: CNAF-disk DOWN since 2007/03/11 13:00

2007/03/14 05:30: CNAF-disk FULL

2007/03/14 09:45 **CNAF-disk OK** 9 TB added

2007/03/12 11:42: CNAF-tape DOWN since 2007/03/11 13:00

2007/03/14 05:30: **CNAF-tape OK**

2007/06/30 10:30: CNAF-disk FULL

2007/03/12 11:42: IN2P3-disk DOWN since 2007/03/09

2007/03/13 00:00: IN2P3-disk OK

2007/03/14 11:00: IN2P3-disk DOWN GGUS 19677, SRM response very slow

2007/03/14 20:30: **IN2P3-disk OK**

2007/03/12 11:42: IN2P3-tape DOWN since 2007/03/09

2007/03/13 00:00: IN2P3-tape OK

2007/03/14 11:00: IN2P3-tape DOWN GGUS 19677, SRM response very slow

2007/03/14 20:30: **IN2P3-tape OK**

GRIDKA

2007/03/13 00:00: GRIDKA-disk FULL

2007/03/14 01:00: GRIDKA-disk DOWNGGUS 19591, SRM response very slow

2007/03/14 13:00: **GRIDKA-disk FULL** 19 TB added

2007/03/12 11:42: GRIDKA-tape OK

2007/03/14 01:00: GRIDKA-tape DOWNGGUS 19591, SRM response very slow

2007/03/14 20:30: **GRIDKA-tape OK**

2007/06/27 05:00 to 18:00: **GRIDKA-disk and tape DOWN Access to gateway nodes is possible again after 27/6 18:00 UTC but some services (shared disks, batch queues, SE, SRM,CE) will not be or are limited available during the following day (28/6).**

2007/06/30 10:00: **GRIDKA-disk FULL**

NIKHEF

2007/03/12 11:42: NIKHEF-disk OK

2007-03-12 13:41: NIKHEF-disk DOWN GGUS 19594, SRM response very slow

2007/03/14 09:30: **NIKHEF-disk OK**

2007/03/12 11:42: NIKHEF-tape OK

2007/03/13 00:00: **NIKHEF-tape DOWN**

2007/03/14 09:30: **NIKHEF-tape OK**

PIC

2007/03/12 11:42: PIC-disk FULL, 28 TB, new servers are expected in 3 weeks and should cover the 56 TB requested.

2007/03/12 15:30: PIC-disk DOWN all trasnfers from SE failing with java exception (networking problem)

2007/03/13 00:00: **PIC-disk FULL**

2007/03/12 11:42: **PIC-tape OK**

2007/06/20-21 09:00: **PIC-disk DOWN Upgrade of dcache infrastructure to 1.7. End of intervention on June, 22 19:00**

2007/06/22 11:42: **PIC-disk OK**

RAL

2007/03/12 11:42: RAL-disk OK

2007/03/14 00:00: RAL-disk DOWNProtocol not supported

2007/03/14 13:30: **RAL-disk OK**

2007/03/12 11:42: RAL-tape OK

2007/03/14 00:00: RAL-tape DOWN Protocol not supported

2007/03/14 13:30: **RAL-tape OK**

T1 status

- **CERN** :
- **CNAF** : "Scheduled downtime starting 9 July 09:00 and ending 11 July 18:00 in the CASTOR instance at CNAF for upgrade."
- **GRIDKA** :
- **IN2P3** :
- **PIC** :
- **NIKHEF** :
- **RAL** :

Problems / issues during the Christmas break of 2006

- 24/12/2006 : GridKa banned by Davide as many jobs were failing there
 1. 31/12/2006 update : GridKa unbanned as a test measure. Banned once again as jobs continued to fail there. Johan to follow it up.
 2. 5/1/07 update : Not yet understood but will try again after Monday when dCache is upgraded at GridKa
- 26/12/2006 : Bookkeeping down (or at least a part of it ... see Davide's email about it to lhcb-production). Update : fixed by Carmine on 2/1/2007
- 28/12/2006 : LCG.IHEP.SU banned by Davide (failing jobs - bus error)
 1. 5/1/07 update : Test unbanning of IHEP. Banned again after jobs still failed. See Raja's email to lhcb-production.
- 29/12/2006 : Problems with reconstruction jobs at RAL (production ID : 1536) - FID error
- 29/12/2006 : Problems with reconstruction jobs at PIC (production ID : 1536) - file does not exist
- 2/1/2007 : Jobs not running at QMUL and PIC since last week - GGUS ticket (id = 17013, 17014) submitted by Raja.
 1. 3/1/07 QMUL update : Power failure in the building hosting the machines. Waiting for problem to be resolved before machines are turned on.
 2. 5/1/07 QMUL update : System is partially back now according to GGUS ticket. Jobs still not going to QMUL.
 3. 6/1/07 PIC update : Jobs are running at PIC now
- 2/1/2007 : Problems with reconstruction jobs at RAL, PIC, CERN (production ID : 1536) - files donot exist. Due to overloads in the tape systems at the respective sites?
- 2/1/2007 : Problems with reconstruction jobs at Dirac.Lyon.fr (production ID : 1536) - Variable VO_LHCB_SW_DIR is not defined
- 3/1/2007 : SAM test results not available. Opened GGUS ticket 17032 (Raja, with copy to Roberto Santinelli). Possibly a LHCb problem, since other VO-s seem to return reasonable results.
 1. 5/1/07 : The problem is that the LHCb tests have not been submitted for some time. (Maybe someone's cron job has stopped working?)
- 3/1/2007 : Problem with lfc causing a large number of jobs to fail (primarily reconstruction jobs for now).
- 5/1/2007 : RAL tape robot going down on Wednesday, 10/1/2007 from 9AM-1PM (UK) estimated

Problems/issues in the last 24 hours

- JINR banned since it seems like they have only PIII machines there (Andrei)
- QMUL banned to be checked why LHCb not submitting jobs there (Raja).
- Problem related to "sse2" flag: it doesn't work for PIII architecture while it does on PIV. They are risking to ban sites with hundred CPUs just because a very few of them is PIII. A crude check of sse2 availability at the level of the job wrapper (like "no go" if not there would make it) should be foreseen.
 1. A check for sse2 flag is implemented in the pilot agent by Ricardo. If machine does not have the flag, the agent will sleep for 10 minutes (even before downloading DIRAC) before dying to prevent black hole effect
 2. Andrei is implementing a monitoring of these pilot agents to see how efficient we will be while running in this mode.
 3. A better algorithm will be implemented soon.
- Jobs stalling on GridKa, IN2P3 and RAL. LHCb local contacts at each site to follow up on it.

Resource Issues

1. SARA/ NIKHEF gsidcap issue (see outstanding problems)
2. CERN: disk server for Production UIs. (see Outstanding problems)
3. IN2P3 does use gsidcap for the tape storage. This also prevents LHCb using Lyon as expected since the production version of ROOT was not able to use gsidcap until Inext release of AA. The disk endpoint of IN2P3 will be used instead.

Outstanding Problems

- LFC replication to T1. **Top priority** Intervention supposed to take place on November 2nd. The intervention (Diana Bosio will coordinate everything) consists on:
 - ◆ - move to new mid-range servers for LFC front-end nodes at CERN;
 - ◆ - upgrade of LFC to 1.5.10 at CERN and CNAF;
 - ◆ - move of DB backend from LCG to LHCb RAC at CERN;
 - ◆ - opening of production R/O LFC service for LHCb at CNAF;
 - ◆ - production Oracle-streams based replication service from CERN to CNAF.
- LFC update of the ownership (and accordingly of ACLs) on all files and directories. **High priority** This activity's priority has been downgraded with respect the first point because the previous intervention already allow the mixed VOMS and no-VOMS production managers accessing/modifying newly created files.
- NIKHEF GSIDCAP ISSUE: **High priority** The center has been downgraded to T2 role because it's currently impossible accessing files stored in the WAN connected Storage at SARA from WN via dcache. Await patched version of the dCache client is gonna be released for test. This version wouldn't require Inbound connectivity on the WN because it wouldn't require calls client back. Until further news, NIKHEF sits out DC06 activity. New dcache libraries and clients are not yet officially released. There are stability and functionality tests undergoing from NIKHEF people. It seems that they will be glad to install on their sites these new clients (even before they get certificated by certification team) once they will prove to be running fine. On the other hands the certification and deployment process for these clients, once the experts in NIKHEF will give the green light, could be dramatically accelerated. These new libraries are backward compatible. The user has just to set a special environment variable to use them or the old ones. Experiment side, LHCb is also testing the GFAL plugin in the Root application. They didn't yet manage to access a file via gsidcap protocol (currently published only at NIKHEF and IN2P3-Tape) with the application crashing. The transfer share to NIKHEF/SARA is still set to 0. Application side (directly from Root **BUT** not within GAUDI) they manage to read a file using dcap:gsidcap:// as protocol and the latest libstunnel libraries

for dcap (against a dcache server in the CERN certification testbed). The GFAL plugin for root doesn't work. However some tests against IN2P3 and RAL (using a backdoor method in ROOT) and the latest version of dCache clients went successfully. Waiting for site managers of SARA (Ron Trompert) to upgrade their own dcache server (expected for the end of September). New hardware delivered, fully installation of the servers will happen on mid of October.

- CERN New diskserver for replacing the old lxfsrk524 serving the production UI and used for running *agent director and monitor* **Low Priority**. To discuss whether we will need it once the Sandbox& Logs storage element is operative.
- Backup of databases. **Medium priority** issue
- Ability of CNAF, GridKa, NIKHEF/SARA to take on output data from jobs on the grid during CERN downtime. **Medium priority** issue.
- Make the Bookkeeping standalone (independent of AFSLoginProblems or other network mounted drives).
 1. Make the machine into a user interface
 2. Local installation of lcg-utils (why? certificates?)
 3. Medium priority
- Endpoints name convention:
 1. The major issue here is that VO namespace should be sacrosanct. Jeff's proposal (site-related-path/VO-path) presented at last ops meeting is reasonable but should be enforced
 2. srm host-name should be generic enough that it doesn't change with time unless forced by special events (e.g. not be the machine cryptic name, not contain the local storage technology). Prefer srm. to dcache05srm. or castorsrm. SURLS are registered in the FC that necessitate lengthy updates when endpoint names are changed.
- Permission for the shift manager to write to data on the production area using their certificates. (Priority as needed)
 1. Shift manager needs to test out the system when they are assistant in the previous week
 2. Need to contact one of the responsables (Joel, Nick) if they do not have the permission
- Multiple jobs submissions (High priority)
 1. Job actually submitted but not marked as submitted by the server
 2. Temporarily fixed for now by Gennady - mark by ourself as submitted **before** submission
 3. Do not want the same job submitted multiple times!
 4. Needs to be understood in better detail especially on the server side
 5. Also gave rise to transient problems at GridKa when jobs failed to transfer output data while the actual data had already been uploaded by a previous job.
- How do we use the databases in DIRAC (Medium - High priority)
 1. The location of databases - centrally at CERN or on the DIRAC machines?
 2. The location of the Sandboxes - within the database or in a SE?
 3. Need / ability to restore the database after a crash.
 4. The need for a "live" backup of the database using a slave server
- Consistency check between LFC, Bookkeeping and what is actually on the SE. (Medium priority)
 1. Check both ways - whether what is on the SE is actually registered and vice versa
 2. Crucial for analysis
 3. A lot of DC04 stripped data missing at RAL (and a few elsewhere).
 4. Andrew Smith to follow this up and transfer the needed files.

Ongoing Tests

1. gLite WMS tests (Gianluca)
2. Stripping tests: jobs failed to upload output because wrongly defined the Local SE.
3. LFC RO replication activity. ([<http://indico.cern.ch/conferenceDisplay.py?confId=7429>])
4. Rotation of the Tier0 SE among CERN, RAL, PIC, Lyon
5. Submit test jobs to test system to test Andrei's new failover mechanism

LCG Middleware

1. glxexec on the T1 for exploiting Dirac prioritization mechanism.
2. lcg-utils: improvement/development. (HIGH Priority)
 1. Need a lcg-util for uploading files to grid without registering it. (lcg-cp that does accept SURL as destination or a lcg-cr that doesn't register). --> *Customized distribution of lcg-utils contains a special version lcg-cp supporting SURLs, written ad hoc for LHCb*
 2. lcg-gt should be able to analyze a list of protocols rather than invoking lcg-gt as many times as there are protocols to be checked
 3. lcg-cr should be able to register a file in the catalog with an additional option that allows to specify the host field.
 4. lcg-cr should delete the physical replica when the registration on the file catalog doesn't succeed. (for consistency of the file catalog). It should also delete the temporary entry in the FC when the transfer fails or the replica already exists. --> *An ad hoc lcg-utils written for LHCb called lcg-clean allows for removing physical replicas on the Storage without any check on LFC, as required*
 5. lcg-gt should return TURL that are compatible with ROOT application. Using the gfal_plugin there wouldn't be any problem, but the current incompatibility should be fixed.
 6. The deployed dCache client in gLite 3.0.0 has a bug in libgsiTunnel that prevents using gsidcap (only protocol available at some sites). This implies a LCG AA release with the current libgsiTunnel.so library.
 7. getbestfile method in GFAL flawed logic when checking on domain names (example is GRIDKA where SE and WNs sit on different domain) so doesn't match. Going through the .BrokerInfo and then through the IS would be the solution. Reported and to be further discussed with developers.

Minutes of Meetings

Previous Problems

Detailed Status of various centers

This topic: LHCb > DC06Activity

Topic revision: r87 - 2007-07-04 - unknown



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