

Table of Contents

Week of 140203.....	1
WLCG Operations Call details.....	1
General Information.....	1
Monday.....	1
Thursday.....	2

Week of 140203

WLCG Operations Call details

- At CERN the meeting room is 513 [R-068](#).
- For remote participation we use the Alcatel system. At 15.00 CE(S)T on Monday and Thursday (by default) do one of the following:
 1. Dial +41227676000 (just 76000 from a CERN office) and enter access code 0119168, or
 2. To have the system call you, click here
- In case of problems with Alcatel, we will use Vidyo as backup. Instructions can be found here . The SCOD will email the WLCG operations list in case the Vidyo backup should be used.

General Information

- The SCOD rota for the next few weeks is at ScodRota
- General information about the WLCG Service can be accessed from the Operations Web

Monday

Attendance:

- local: MariaD (SCOD), Maarten (ALICE), Massimo (CERN Data Mgmt), Vitor (CERN Grid Services), Felix (ASGC).
- remote: Roger (NDGF), Sang-Un (KISTI), Michael (BNL), Matteo (CNAF), Elena (ATLAS), Eric (CMS), Onno (NL_T1), Kyle (OSG), Tiju (RAL), Alexei (LHCb), Lisa (FNAL), Pepe (PIC).

Experiments round table:

- ATLAS reports (raw view) -
 - ◆ Central services/T0
 - ◇ CERN_PROD: Transfers were failing with permission denied errors on Monday morning. Noticed and fixed by CERN team. Thanks.
 - ◆ T1
 - ◇ TAIWAN: heavy SRM load caused transfer failures on Sunday (GGUS:100904). Fixed.
 - ◇ FZK: staging errors for DATATAPE on Friday (GGUS:100885). Fixed by issuing a retry for all outstanding stage requests for ATLAS and restarting tape storage software.
 - ◇ PIC: problem with one disk pool, which caused transfers to failed on Friday (GGUS:100874), dCache pool restarted.
- CMS reports (raw view) -
 - ◆ T1/T2/Others: Business as usual. Smooth running.
 - ◆ Preparing for DBS (data catalog) upgrade on Feb 10. That week will see little to no central processing
 - ◆ One problem: ARGUS cluster issue(s) (DNS? and then a new, uninitialized node in the cluster) caused problems with analysis jobs running.
 - ◇ Debugged by CMS analysis operations. Better would be to have SLS monitoring of the ARGUS cluster. Ticket is GGUS:100870

- ALICE -
 - ◆ sites please take note of the necessary WLCG VOBOX update announced last Fri
 - ◇ see details below
 - ◆ KIT
 - ◇ the number of corrupted files has *shrunk* by 45% to 26126
 - ◇ 21k files have been salvaged after all, thanks very much!
- LHCb reports (raw view) -
 - ◆ Mostly simulation and user jobs. Smooth running over most of the grid.
 - ◇ T0: Pilots aborted at ce202.cern.ch today. Ticket is GGUS:100902
 - ◇ T1: NTR
 - ◇ T2: NTR

Sites / Services round table:

- ASGC: ntr
- BNL: ntr
- FNAL: ntr
- OSG: ntr
- KISTI: ntr
- NL_T1: ntr
- CNAF: ntr
- PIC: ntr
- NDGF: ntr
- IN2P3: ntr (sent be email)
- RAL: Tomorrow, between 8-10hrs am UK time, tape system intervention. Site set at risk in GOCDB.
- CERN:
 - ◆ Grid Services: ntr
 - ◆ Data Mgmt:
 - ◇ Problem to access EOS from outside CERN. Now solved. Lasted for 1h 15'.
 - ◇ ROOT access to CASTOR is now switched off. Hardly 10 users concerned. They have been informed about alternative access methods.

AOB:

- WLCG VOBOX
 - ◆ as announced on the wlcg-operations list last Fri, please ensure your WLCG VOBOX instances generate host proxies with **1024-bit** keys!
 - ◆ preferably update Globus; correct minimal versions of the affected rpm:
 - ◇ globus-proxy-utils-5.0-6 (Globus 5.0)
 - UMD-3 SL6 rpm (can be installed manually also on UMD-2 machines)
 - UMD-3 SL5 rpm (ditto)
 - ◇ globus-proxy-utils-5.2-1 (Globus 5.2)
 - from EPEL for EMI-3 and EMI-2
 - ◆ otherwise one can apply this quick hack:

```
perl -pi.bak -e 's/ -q / -bits 1024 $&/' \  
/etc/vobox/templates/voname-box-proxyrenewal \  
/etc/init.d/*-box-proxyrenewal
```

Thursday

Attendance:

Monday

- local: MariaD (SCOD), Maarten (ALICE), Massimo (CERN Data Mgnt), Vitor (CERN Grid Services), Felix (ASGC), Pablo (GGUS), Przemek (DB), Alexandre (Dashboards).
- remote: Roger (NDGF), Michael (BNL), Saverio (CNAF), Eric (CMS), Dennis (NL_T1), Kyle (OSG), Gareth (RAL), Alexei (LHCb), Lisa (FNAL), Pepe (PIC), Jeremy (GridPP), Rolf (IN2P3), Pavel (KIT).

Experiments round table:

- ATLAS reports (raw view) -
 - ◆ Central services/T0
 - ◇ IT and DE clouds moved to FTS3
 - ◆ T1
 - ◇ CERN-PROD CVMFS inside CERN faulty GGUS:100928 <https://cern.service-now.com/service-portal/view-outage.do?from=CSP-Service-Status-Board>
- CMS reports (raw view) -
 - ◆ T1/T2/Others: Bussiness as usual. Smooth running.
 - ◆ Preparing for DBS (data catalog) upgrade on Feb 10. That week will see little to no central processing. Ramp-down has begun.
 - ◆ We are encouraging all our sites to switch to FTS3 server at RAL for load testing. Begins in a week or so.
- ALICE -
 - ◆ CNAF
 - ◇ tape SE updated to xrootd v3.3.4 (on Jan 28) with new checksum plugin successfully validated (Feb 5) with test transfers, thanks!
 - ◆ KIT
 - ◇ investigating why many jobs read a lot of data remotely from CERN
 - ◆ RRC-KI-T1
 - ◇ memory tuning for jobs ongoing, thanks!
- LHCb reports (raw view) -
 - ◆ Mostly simulation and user jobs. Smooth running over most of the grid.
 - ◆ T0: CVMFS caused 50%+ jobs failures Mo and Tu, back to normal since Wed
 - ◆ T1: NTR

Sites / Services round table:

- ASGC: ntr
- BNL: ntr
- FNAL: ntr
- OSG: ntr
- KISTI: not connected or I couldn't hear... sorry!
- NL_T1: ntr
- CNAF: ntr
- PIC: ntr
- NDGF: ntr
- IN2P3: ntr
- RAL: Now investigating CASTOR failing tests. Scheduling FTS3 tests.
- CERN:
 - ◆ Grid Services: ntr
 - ◆ Data Mgnt: ntr
 - ◆ Dashboards: The following FTS servers do not report information properly.

- ◇ These do not appear in the year log (is it possible they still have the old broker name hardcoded ?) : fts02.usatlas.bnl.gov, w-fts001.grid.sinica.edu.tw, fts-kit.gridka.de, fts-fzk.gridka.de
- ◇ These have authentication error in the broker: fts00.grid.hep.ph.ic.ac.uk, fts3.grid.sara.nl (empty or bogus username).

- ◆ Databases: Here is a short description of last Thursday's 2014/01/30 LCGR problems: *As a part of preparation work to migration and upgrade of LCGR database, the replication of the database has been established. Because of a misconfiguration of database archived log files deletion policy, the space on production database server has been exhausted and the database got stuck at around 7.20AM. The problem has been corrected at 9.30AM and the database came back. After the database went up, all applications, which cached their data during the database outage, started to write to the database the content of their cache. The LCGR database was not able to handle such strike of traffic in one moment, so it got stuck again. Another reboot of the database was required. The manually synchronized restart of applications allowed the database to come back to normal operation. The preparation work for the migration was continued during the day and around 5PM we hit an Oracle bug, which caused the database not to accept new connections. The existing ones were working properly. Around 6PM, restart of one of database nodes and cut of connection between the production and replicated database helped to solve the problem.*

- GGUS:

- ◆ Suggestion to remove three fields from the 'Ticket Submission Form' (see attachment). Those fields are hardly ever used, and they are anyways concatenated to the body of the issue. The meeting decided these fields can be deleted.

AOB:

- OpenSSL issue

- ◆ EGI broadcast [sent](#) Feb 4 describing current state of affairs and recipes for cures
- ◆ Sites using **HTCondor as batch system** may need to apply one of these configuration changes for now:
 - ◇ DELEGATE_JOB_GSI_CREDENTIALS = False
 - ◇ GSI_DELEGATION_KEYBITS = 1024
- ◆ HTCondor v8.0.6 will have the default increased to 1024

This topic: LCG > WLCGDailyMeetingsWeek140203

Topic revision: r15 - 2014-02-06 - MariaDimou



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