

Table of Contents

Minutes of the WLCG Tier1 Service Coordination Meeting of 29 July 2010.....	1
Agenda.....	1
Connected by phone.....	1
Local.....	1
Status of open GGUS tickets.....	1
LHCb.....	1
GGUS ALARM CHAIN.....	2
Review of recent / open SIRs.....	2
Conditions Data Access and related services.....	2
COOL, CORAL and POOL.....	2
Experiment Database Service Issues.....	3
AOB.....	4

Minutes of the WLCG Tier1 Service Coordination Meeting of 29 July 2010

Agenda

Connected by phone

- Gonzalo Merino (PIC), Michael Ernst (BNL), Alessandro Cavalli (CNAF), Carmine Cioffi (RAL), Onno Zweers (NL-T1), John DeStefano (ATLAS), Dario Barberis (ATLAS), Andrew (TRIUMF), BNL, Elizabeth Gallas (ATLAS), Felix Lee (ASGC), Alexander Verkooijen (NL-T1), Jhen-Wei Huang (ASGC), Elena Planas (PIC), Frederique Chollet (LCG-FR), Gareth Smith (RAL), Andreas Motzke (GridKA), Jon Bakken (FNAL), Carlos Fernando Gamboa(BNL), Jim Shank (ATLAS), Matt Hodges (RAL)

Local

- Helge Meinhard, Tim Bell, Gavin McCance, Maria Allandes, Jamie Shiers, Ian Fisk (CMS), Simone Campana, Andrea Valassi, Andrea Sciaba, Alessandro di Girolamo, Patricia Mendez Lorenzo, Jean-Philippe Baud, Maarten Litmaath, Jacek Wojcieszuk, Dawik Wojcik, (others)

Status of open GGUS tickets

- No tickets to follow up for ATLAS, ALICE or CMS.
- The tickets listed below are being followed up and the tickets should have been updated with the status
- The issue of GGUS<->PRMS synchronization needs to be followed up elsewhere - action: SCOD team to ensure this happens

LHCb

There are currently 26 tickets open since last year. Excluding the ones open in the last week (5) and filtering GGUS tickets submitted to (less important) T2s, the following cases deserve a close investigation:

- GGUS:58686 (ALARM-test) wasn't closed. PRMS: CT000000687493. Why the two systems did not talk properly?
- GGUS:58975 (problem with CRL update on AFS UI) was reopened; the ticket stayed in an open status while the PRMS associated (CT000000690572) was closed long before. I closed today this ticket (no longer relevant for LHCb) but it should be interesting to know (again) why PRMS and GGUS were not in sync.
- GGUS:59247 about LSF not properly reporting CPUTime at CERN. Ticket updated the 22/06 by Graciani Ricardo. No more news since then. This is has to be followed closely being an important matter for LHCb.
- GGUS:59880 at IN2P3-CC. People in Lyon are looking very actively at the problem but the issue seems to be still there (documented by a not negligible failure rate directly and indirectly due to that and SAM failures/degradation). May be some extra support from WLCG is needed?
- GGUS:59422 transfers from UK (and not only) to CERN. It has not been properly updated since beginning on July. The understanding of the issue might involve other service people (like Network).

Ticket	Site	Issue
GGUS:58686, GGUS:58975	CERN	GGUS/Remedy i/f

GGUS:59247 ↗	CERN	LSF & CPUtime
GGUS:59880 ↗	IN2P3	Access to shared area (timeouts)
GGUS:59422 ↗	GridPP and?	Transfers

GGUS ALARM CHAIN

- There are too many failures of something in the system - it needs to be much more reliable. All changes in the chain must be reported at the daily meeting and (it would seem) tested. To be escalated within WLCG and EGI (Jamie)

Review of recent / open SIRs

- In general, several of these SIRs have pending actions. The CERN vault cooling SIR does not follow the FIO(!) template.
- SIRs with out-standing or pending actions need to be followed up!
- SIR received for NDGF SRM outage on 2010/0714
- SIR received for GridKa cooling system failure incident of 2010/07/10.
- SIR received for reduced availability caused by data corruption at NL-T1 on 2010/07/05
- SIR being prepared from GGUS/OSG about notification issues
- SIR being prepared for CERN vault cooling issues

Conditions Data Access and related services

COOL, CORAL and POOL

- Fixes have been prepared in the CORAL FrontierAccess plugin for some bugs which prevented the readback of the ATLAS Detector Description via frontier (bug #70208 [↗](#)). These fixes have not been included in any CORAL release yet. The ATLAS Detector Description is normally not read back via Frontier, anyway.
- LHCb has reported problems in the CORAL LFCReplicaSvc component from the last release LCGCMT_58d (bug #70641 [↗](#)).
 - ◆ The problem has been tracked down to a well-known feature of Globus, which uses an own version of the GSSAPI library that is incompatible to that provided by the system. Whenever the system version of GSSAPI is loaded into application memory before the Globus version, software components which need the Globus GSSAPI extensions (like LFC) cease to work correctly.
 - ◆ The problem has been observed for the first time in the LCGCMT_58d release because of the upgrade to Xerces version 3.1.1, which is now linked to the libgssapi_krb5.so.2 system library because of its new 'network support' feature. The Xerces libraries are generally loaded into CORAL applications before the LFC plugin, triggering the load of the system GSSAPI before the Globus GSSAPI. The problem is being fixed for LHCb with a workaround that consists in a '3.1.1p1' rebuild of Xerces without network support, i.e. with no GSSAPI link dependency. The new Xerces version will be used to prepare a new LCGCMT_58e release for LHCb and also a modification of the LCGCMT_59 release already prepared for ATLAS.
 - ◆ The link dependency of Xerces 3.1.1 on libgssapi_krb5.so.2 has been found also in the frontier_client library. No action is planned for frontier_client because a similar configuration exists since a long time and no problems related to a GSSAPI mismatch with LFC have been reported so far.
 - ◆ The most appropriate solution to the problem would probably consist in patching Globus

and/or the Grid client software that depends on Globus (e.g. the LFC client), either to use the system provided GSSAPI or to allow the Globus and system GSSAPI to coexist in the same application. A patch implementing the latter option seems to have been developed in October 2008 to address this issue, filed as Globus bug 6415. The possibility and timescales for fixing the problem in Globus will be followed up with the Grid middleware developers.

Experiment Database Service Issues

- Experiment reports:
 - ◆ ALICE:
 - ◇ A reboot of ALIONR DB and its disk arrays was performed on Monday 19th July. The intervention is supposed to resolve false warnings occasionally reported by one of the disk arrays.
 - ◆ ATLAS:
 - ◇ Oracle Listener on 2nd instance of Atlas offline database (ATLR) got stuck on 21th of July at around 8 AM. The issue was handled immediately by a DBA on-shift and the listener was restarted. Due to the problem big part of new database connections was being rejected for several minutes. Existing database sessions were not affected.
 - ◇ High CPU utilization was observed on 3rd node of Atlas offline database (ATLR) on Saturday 14th July starting from 3PM. The issue was caused by a spike of activity from Atlas Panda application. Issue was resolved together with Atlas.
 - ◇ Since Monday 26th July we are running series of stress tests aiming in reproduction of issues with April PSU patches from Oracle. So far the issues have not been reproduced and tests are continuing. A meeting with Atlas to discuss tests status and further plans is scheduled for Monday 2-8. Many thanks to ATLAS for their help.
 - ◆ CMS:
 - ◇ On Friday 23rd July a schema used by the application configuring the CMS Traker detector got corrupted due to a human error. The issue was fixed manually by developers.
 - ◇ On Monday 26th July another production schema (CMS_BEAM_COND) got corrupted due to a user error. The schema was successfully recovered from a backup.
 - ◆ LHCb:
 - ◇ One of the disk arrays hosting LHCb online production database rebooted unexpectedly during reconfiguration on Tuesday morning. The incident was transparent for the database users but some clean-up will be necessary. The root cause of the reboot is being investigated by LHCb sysadmins.
- Site reports:

Site	Status, recent changes, incidents, ...	Planned interventions
ASGC	* Patch April 2010/July 2010: - Testbed Single instance DB is patched to 10.2.0.4 - RAC Testbeds are ready to patch to 10.2.0.4, then, --> April 2010--> July 2010.	None
BNL	Nothing to report	Rolling back PSU April 2010 prior to deployment of PSU JULY 2010 in the development/pre-production DB.
CNAF	26th of July - problem with one of the LHCb clusters - connection timeout resulted in streams timeout - network problem - under investigation.	
KIT	Nothing to report	None
IN2P3		

NDGF	Nothing to report	None
PIC	* On 20th of July during scheduled downtime, we upgraded firmware for all the systems. Also, the HW problem with a LAN card of one node was finally solved.	No interventions
RAL	* On 20th of July - when trying to change multipath configuration some problems occurred, configuration roll-back was required. As a consequence ATLAS instance rebooted and there were problems afterwards with apply process. Few logical change records for ATLAS conditions replication between CERN and RAL were lost causing inconsistency of data and streaming failure. On Wednesday afternoon CERN DBAs resend missing data changes using additional temporary Streams setup. Consistency of data was restored and replication lag was recovered over Wednesday night. RAL DBAs opened a Service Request to understand the root cause but it is quite likely that this was another instance of bug 9232517 which we observed few months ago and which was not fixed by Oracle. Additionally Streams project manager was contacted to escalate the issue. Due to the issue RAL OGMA database was out of sync between Tuesday 20th July 12 AM and Thursday 22nd July 1AM.	None
SARA	Nothing to report	No interventions
TRIUMF	* Service incident - failed duplicating of our 3D Oracle database to new upgraded server (RH Linux 5 & SAS storage) * Data guard has been prepared now to switchover to new upgraded server (see planned interventions) * FTS Oracle database moved to new Oracle RAC servers.	* 2 hour outage planned for Thursday July 29 17:00hrs (UTC) to switch to TRAC standby DB running on new hardware

AOB

- Topics for future meetings:
 - ◆ The Site Status Board - Feedback from Sites
 - ◆ Prolonged site & service downtimes - strategies
 - ◆ Certification of the April Oracle PSU - follow-up
- Next meeting: in 1 month (stick to two week schedule but skip one meeting).

-- JamieShiers - 28-Jul-2010

This topic: LCG > WLCGTier1ServiceCoordinationMinutes100729

Topic revision: r11 - 2010-08-02 - JamieShiers



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