

Table of Contents

WLCG Tier1 Service Coordination Minutes - 23 September 2010.....	1
Present.....	1
Connected.....	1
Release update.....	1
WLCG Baseline Versions.....	1
Preparing for the Heavy Ion run - information for sites.....	1
Status of open GGUS tickets.....	1
Status of recent / open SIRs.....	2
Kernel Upgrade Status.....	2
Prolonged Site Downtimes.....	2
Conditions data access and related services.....	2
COOL, CORAL and POOL.....	2
Frontier.....	3
Database services.....	3
AOB.....	5

WLCG Tier1 Service Coordination Minutes - 23 September 2010

Present

Manuel, Nicolo, Patricia, Maria Alandes, Romain, Julia, Kate, Steve Gowdy, Oliver, Maarten, Miguel, Jamie, Maria, Roberto, Massimo, Andrea V, Stephane, Zsolt, MariaDZ, Harry, Kors, Dawid, Raffaello, Ian F, Alexei, Luca

Connected

Michael Ernst, Xavier Mol (KIT), Eter Pani, Gonzalo Merino (PIC), Jon Bakken, John DeStefano, Gunter Grein (GGUS), Carlos Fernando Gamboa, Dave Dykstra, Jhen-Wei Huang (ASGC), 00886922923533, Felix Lee (ASGC), Andrew (TRIUMF), Jon (NDGF), Alexander Verkooijen (NL-T1)

Release update

Maria Alandes presented the release status and in particular the plans for gLite 3.1 retirement. Please send your feedback to glite31-retirement@cernSPAMNOTNOSPAMPLEASE.ch.

WLCG Baseline Versions

- Release report: deployment status wiki page
- WLCG Baseline versions: table

Preparing for the Heavy Ion run - information for sites

See slides attached to the agenda [↗](#). (Ian Fisk for CMS, Massimo Lamanna for CASTOR).

In summary:

- LHCb does not take Heavy Ion (HI) data
- For CMS, HI data will be processed at CERN and reprocessed at FNAL.
- ATLAS HI data will be sent to all available ATLAS Tier1s.
- ALICE HI data will be distributed in the period following the HI run.
- In addition, pp activity (re-processing, MC, analysis) will continue (all experiments, all sites).

Status of open GGUS tickets

Mostly ATLAS network problems. The experiment requires a body related to WLCG and to the network providers between sites to take ownership of such tickets. Our current practice, i.e. notifying one of the 2 sites involved in a case of network problem, is correct, especially because it is not clear from the beginning this is a network problem indeed. The issue is who feels responsible for finding a solution to network problems between sites. A network authority should be involved in all relevant tickets, no matter which physical infrastructure is used. There is a GGUS Support Unit *Network Operations* but seems in-effective in its current incarnation. No traffic and no WLCG supporter involved. Details in <http://savannah.cern.ch/support/?115213> [↗](#). Jamie will make a presentation at the next OPN meeting asking for such a network authority to be involved from the beginning of network tickets' life. Alternatively, a WLCG supporter will be designated on a case by case basis.

Guenter presented the T0, T1 answers on local handling of ALARM tickets[?]. NL_T1 don't have experts on call or operators around the clock, so ALARMS outside working hours.

Status of recent / open SIRs

SIRs received for:

- NL-T1 DB problem - report[?]
- ASGC DB problem - report

SIRs pending for:

- Network issues - RAL/NDGF (GGUS:61306[?], finally solved after 29 days), BNL/CNAF (GGUS:61440[?]) [N.B. still to be clarified - who should produce each of these SIRs]
- LHCb LFC replication problem

Kernel Upgrade Status

An official release has been made now by RH. Sites should upgrade in the next 7 days. The recommended strategy in case of problems of such severity are:

1. apply upgrade (patch)
2. reinstall

That being said, it is clearly understood that it is a site responsibility to perform its own risk assessment and decide accordingly (and inform).

Prolonged Site Downtimes

Timelines were proposed. To be used with caution.

- First 24h - site internal - report to daily meeting and via any open GGUS tickets.
- Following 48h - 1st level escalation (WLCG MB informed), prepare for longer downtime
- Up to 2 weeks - 2nd level escalation (WLCG helps mediate choice of backup sites (inter-VO issues))
- Beyond 2 weeks - 3rd level escalation (WLCG OB & CB informed).

Obviously flexibility will be required - these are intended as a guideline.

Conditions data access and related services

COOL, CORAL and POOL

- A patch (tagged as CORAL_2_3_12) has been prepared for CMS to fix a crash in CORAL_2_3_10 with optimized gcc 4.3 builds, which is now understood to be due to a bug in the OracleAccess plugin.
 - ◆ The crash had already been observed last year, but had initially been attributed to a bug in gcc optimization and had been solved by the workaround of disabling optimization on one C++ file.
 - ◆ Following this issue, discussions are ongoing with CMS to set up nightly builds and tests of CORAL in their development framework (as is already done for ATLAS and LHCb), to ease and speed up the integration of new CORAL versions into the CMS software in the future.

Frontier

- ATLAS weekly FroNTier meetings [↗](#)

Database services

- Topics of general discussion
 - ◆ PSU Jul testing
 - ◇ Rollingness tests are in progress - we are able to reproduce the cluster-wide hang when patching with PSU APR with load from siwngbench and COOL. We will re-use those tests we are doing now to test Oct PSU when it comes out hoping to discover similar issues if they are introduced by the patch.
 - ◇ We support our recommendation of not installing the PSUs if they were not installed already, if they were please keep an eye on the systems.
 - ◆ CMS spontaneous reboots (see below) - all issues have been identified to be caused by excessive PGA memory consumption. We are identifying all the queries responsible and we will contact developers and application owners individually.
 - ◆ Open ports and network at Tier1s - discussion
 - ◆ Distributed Database Operations Workshop - proposal
- Experiment reports:
 - ◆ ALICE:
 - ◇ Nothing to report
 - ◆ ATLAS:
 - ◇ Atlas node 3 crashed on Monday (30th of Aug) night due to high load from PANDA. High load was caused by a particular activity from PANDA applications. PANDA development team has been contacted and they have disabled this activity in production and will develop an alternative mechanism.
 - ◇ Standby DB corruption problem is still under investigation by the team and Oracle Support.
 - ◆ CMS:
 - ◇ On Friday morning (27th of Aug) there were single reboots of nodes 3 and 4 of CMS offline database (CMSR). The problematic processes have been killed. Following those incidents Streams replication from online to offline got stuck and a manual intervention was needed to make it working again. Finally around the noon nodes 3 and 4 rebooted again, which blocked the replication again. The streaming issue was fixed around 13:00.
 - ◇ On Monday 30th August there were two issues affecting replication of PVSS data of CMS experiment. First issue seems to be caused by an unknown bug and is still being investigated. Streaming could be restarted with a manual intervention. Second issue occurred few hours later and was caused by a mistake of a developer who added a column with an unsupported datatype to a replicated table. The issue was fixed by dropping the offending column (done by the user).
 - ◇ On Monday evening (13th of Sep) CMSR database suffered again from spontaneous node reboots. Thanks to extra monitoring deployed recently it was possible to determine that the reboots were caused by excessive PGA memory consumption by sessions of CMS DataSet Bookkeeping application. Exact queries triggering the issue will still have to be identified.
 - ◆ LHCb:
 - ◇ LHCb requested an urgent rollback of conditions data one day behind due to a user error. In the end the the rollback on our side was not necessary as LHCb found a work-around to fix the data. (Recovery discussion requested by Marco).
 - ◇ On Sunday night (12th of Sep) we have observed deadlocks between streams propagation jobs at LHCb downstream capture database. Killing of blocking session

was required in order to resume the replication.

- ◇ On Tuesday (14th of Sep) we were notified by LHCb people about slow down of LFC replication. It was caused by self disabled real time downstream capture optimization which happened during recovering of SARA. We have re-enabled it for LFC and ATLAS replication.

- Site reports:

Site	Status, recent changes, incidents, ...	Planned interventions
ASGC	<ul style="list-style-type: none"> * DB status: - Scheduled downtime: stager and fts DB corrupted. Recovered on 3 September. - 3D Stream: Unable to extend space for tables and indexes in SYSAUX caused by ASM disk error. Drop bad disk and add new datafile to SYSAUX. Recovered it on 16 September. * Patch April 2010/July 2010: - RAC Testbeds were patched to 10.2.0.4. - Prepare for patch April 2010. 	None
BNL	<ul style="list-style-type: none"> -BNL OEM was enabled to work with DOE Grids Certificates. Underlying production agents were reconfigured to work with the new OEM BNL third party Certificate. * Deployed PSU JULY 2010 and 6196748Patch on development/pre-production cluster and TAGS tests cluster. Postponed deployment of PSU JULY patch on production Cond DB per Tier1 Service Coordination Meeting (WLCG T1SC100826) recommendation. * Observed 07445 entry on Cond DB error logs. 	<ul style="list-style-type: none"> * To deploy patch 6196748. * To test/deploy recent OS kernel on production clusters per asmlibs packages availability.
CNAF	<ul style="list-style-type: none"> - updated parameters for LHCB CONDDDB which reached session limits - opened access for LHCB cluster to "world" (CNAF policy is open-on request, preferably not world-opened)%BR- successfully tested TDP backup with TSM, now we have to buy licenses and put in prod - solved network problems on LHCB cluster with RH5 kernel upgrade 	- we have planned to install frontier for Atlas, not yet fully ready to say a "due date"
IN2P3	<ul style="list-style-type: none"> DBAMI - 16/09/2010 8:00 - 9:00 UTC - One year ago, AMI raised a bug on Oracle database : when NLS_SORT=BINARY_CI AND NLS_COMP=LINGUISTIC are altered, any select including a "like" operator returned a wrong result (bug 7522759) Oracle support provided us a patch we tested and installed. So, would it be possible to install this patch at CERN so that the AMI application has the same behavior at CERN and IN2P3 ? ALL 3D databases - 21/09/2010 17:00 - 21:30 UTC - All databases were completely unavailable due to a network maintenance. All streams processes were stopped before the intervention. DBLHCB - 03/09/2010 08:00 - 08-20 - The Lhcb experience reached the max number session authorized. We increased the session number to 550 (old 300) 	
KIT	<ul style="list-style-type: none"> On Sep 7, an ATLAS LFC DB tablespace reached its limit. The tablespace was extended. On Sep 8, LHCb LFC/3D streams had to be restarted due to a short network interruption. 	PSUs Jul/Oct to be discussed
NDGF	Nothing to report	None
PIC	Nothing to report	No interventions
RAL	Ongoing investigations into performance of disk access for LHCb. Failure of CMS disk server shortly after it was returned to use following earlier failure led to loss of 30 files.	Upgrade of LHCb Castor instance to version 2.1.9 on 27-29 September.
SARA		No interventions

	<p>SARA has uploaded a site incident report concerning the database problem that started at August 18th. It can be found at http://sirs.grid.sara.nl/docs/NL-T1_SIR-20100818.pdf</p> <p>Currently running DBs on a single server. In the mean time working with Sun and Oracle to find the cause of the data corruption problem on RAC cluster. Haven't found anything conclusive yet so it is hard to give an estimate of when we'll be moving the DB's back to the original hardware.</p>	
TRIUMF	NTR	Planned outage next week to update Linux kernel on our 3D Oracle RAC servers.

AOB

-- JamieShiers - 23-Sep-2010

This topic: LCG > WLCGTier1ServiceCoordinationMinutes100923

Topic revision: r14 - 2010-09-27 - 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