

Table of Contents

WLCG Tier1 Service Coordination Minutes - 1 July 2010.....	1
Attendance.....	1
Minutes of last meeting and matters arising.....	1
Alarm chain tests and Alarm Issues.....	1
Status of Open GGUS Tickets.....	1
Preparation for WLCG Collaboration Workshop 7-9 July.....	1
Deployment / Rollout Issues.....	1
glexec.....	2
Data Management & Other Tier1 Service Issues.....	2
CASTOR news.....	3
dCache news.....	3
StoRM news.....	3
DPM news.....	3
LFC news.....	3
FTS news.....	3
WLCG Baseline Versions.....	3
Conditions data access and related services.....	4
.COOL, CORAL and POOL.....	4
Database services.....	4
AOB.....	6

WLCG Tier1 Service Coordination Minutes - 1 July 2010

Attendance

Local: Jacek, Lola, Jean-Phillipe, Manuel, Patricia, Tim, Andrea S, Andrea V, Flavia, Jamie, Maria G, Maria A, Maria DZ, Marie-Christine, Dawid, Harry, Dario, Alexei, Simone, Stephane, Maarten, Alessandro, Julia, IanF, Helge

Remote: Federico & Roberto (LHCb), Carmine Cioffi, Ron Trompert (NL-T1), Gonzalo Merino (PIC Tier1), Jon Bakken, John DeStefano, Angela Poschlad (KIT), Elizabeth Gallas, dave.dykstra@cernNOSPAMPLEASE.ch, Carlos Fernando Gamboa, Andreas Motzke (KIT), patrick, FelixLee - ASGC), Elena Planas (PIC), Alessandro (CNAF), Alexander Verkooijen (SARA)

Minutes of last meeting and matters arising

Alarm chain tests and Alarm Issues

Three problems met at the regular ALARM test of this GGUS Release (June 23rd 2010):

1. Test ALARM email notifications were signed with a wrong certificate and had to be repeated.
2. NDGF took a week to realise they received an ALARM.
3. Some email notifications were not received. The investigation is now at the level of KIT,SARA,CERN mailserver logs.

Status of Open GGUS Tickets

Experiments complain at times they don't get the support they need. To dig into such cases a new procedure starts today:

1. Experiments send to Maria problematic GGUS ticket numbers every Wednesday no later than 6pm.
2. Maria presents analysis results on Thursdays at 3:30pm.

Preparation for WLCG Collaboration Workshop 7-9 July

- Things basically work. Some topics raised at June GDB but no show-stoppers per se. Things working ok but some specific points that need to be addressed.
- Data and work flows have been quite variable. If there are any problems plenty of time to catch up - no big backlogs.(Can keep up on average due to this - still wait for sustained data rates...)
- Other Tier1s??

- FNAL - I agree (Jon)
- PIC - agree (Gonzalo)
- KIT - same
- NL- T1 same

Deployment / Rollout Issues

glexec

- LCG-CE test results [↗](#)
- CREAM test results [↗](#)

Site	"/ops/Role=pilot" job + glexec test	glexec capability in BDII
ASGC	LCG-CE OK, CREAM AuthZ/LCMAPS error	
BNL		
CERN	OK	OK
CNAF	LCMAPS error	
FNAL	OK for CMS	OK
IN2P3CC	CE AuthZ error or LCMAPS error	OK
KIT	OK	OK
NDGF	n/a	n/a
NIKHEF	LCG-CE OK, CREAM connection error	OK
PIC	LCMAPS error	OK
RAL	OK	OK
SARA	OK	OK
TRIUMF	OK	

Data Management & Other Tier1 Service Issues

Site	Status	Recent changes	Planned changes
CERN	CASTOR 2.1.9-5 (All) SRM 2.9-3 (all)	None	None planned for the LHC instances
ASGC	CASTOR 2.1.7-19 (stager, nameserver) CASTOR 2.1.8-14 (tapeserver) SRM 2.8-2	22/6 ↗ : Re-cabling stage 1/3 for CASTOR and SRM servers. 24/6 ↗ : Re-cabling stage 2/3 for about 50 disk servers and Oracle DB. 29/6 ↗ : Re-cabling stage 3/3 for 50 disk servers.	None
BNL	dCache 1.9.4-3	None	None
CNAF	CASTOR 2.1.7-27 (ALICE) SRM 2.8-5 (ALICE) StoRM 1.5.1-3 (ATLAS, CMS, LHCb,ALICE)	Put in production 1.2 PB of new disks, 8 disk servers (10 Gb) and 4 GridFTP servers for ATLAS. Taken out of production 430 TB of disk, 18 disk servers, 6 GridFTP servers (all 2x1 Gb). Migrated 250 TB of ATLAS data without service interruption	Alice disk space expansion and data migration (transparent for users)
FNAL	dCache 1.9.5-10 (admin nodes) dCache 1.9.5-12 (pool nodes)	Added 2 PB dCache FY10 disk, retiring old disk now	Buying 1 PB dCache disk for LPCCAF
IN2P3	dCache 1.9.5-11 with Chimera		
KIT	dCache 1.9.5-15 (admin nodes) dCache 1.9.5-5 - 1.9.5-15 (pool nodes)		6/7: short SRM restart, 1 hour of downtime
NDGF			

	dCache 1.9.7 (head nodes) dCache 1.9.5, 1.9.6 (pool nodes)		
NL-T1	dCache 1.9.5-21 with chimera (SARA), DPM 1.7.3 (NIKHEF)		26/7: one day of downtime to upgrade firmware on storage service and reboot with new kernel
PIC	dCache 1.9.5-20rc1		20/7: in the morning, intervention scheduled (no details)
RAL	CASTOR 2.1.7-27 and 2.1.9-6 (stagers) CASTOR 2.1.8-3 (nameserver central node) CASTOR 2.1.8-8, 2.1.8-14 and 2.1.9-1 (tape servers) SRM 2.8-2	LHCb, ALICE upgraded to 2.1.9	CMS, ATLAS to be upgraded to 2.1.9 during November
TRIUMF	dCache 1.9.5-17 with Chimera namespace		

CASTOR news

- We are upgrading CASTOR 2.1.9-7 on the pre-production and repack instances. This will not affect the experiment instances.

dCache news

The extended tape protection will be part of the golden release from next week with version 1.9.5-20. The new functionality has been thoroughly tested by PIC which is already running the release candidate.

dCache is planning for the next Golden Release Series, which will become available starting spring 2011. The current Golden Release (1.9.5) will nevertheless be supported until the end of the current LHC run period.

StoRM news

DPM news

NIKHEF should upgrade to 1.7.4-7 as soon as available as performance is going to improve, 1) because now old requests get automatically cleaned from the database and 2) because it works better with the ROOT TTreeCache mechanism.

LFC news

Sites are recommended to upgrade to 1.7.4-7 as soon as it is available.

FTS news

Testing of 2.2.5 should end by mid July. New features include support for sites without an SRM endpoint and VOMS server certificates not being needed any longer.

WLCG Baseline Versions

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

Conditions data access and related services

COOL, CORAL and POOL

- Two new releases of COOL, CORAL and POOL (LCGCMT_58d and LCGCMT_59) are being prepared for LHCb and ATLAS, including several bug fixes and improvements in all Persistency software projects. The two new releases use the same code base for COOL, CORAL and POOL: the only difference is that the ATLAS release is based on python 2.6, whereas the LHCb release is based on python 2.5. The same versions of all other externals are also used, including ROOT: as of this release, ATLAS will also use ROOT 5.26 (like LHCb), while it was previously using ROOT 5.22.
- ATLAS weekly FroNTier meetings [↗](#)

Database services

- Experiment reports:
 - ◆ ALICE:
 - ◇ Plans for new hardware acquisition to gradually replace the DB hardware in the pit (no dates yet).
 - ◆ ATLAS:
 - ◇ April PSU was rolled back on ATONR and ATLR. ATLARC still has the patch - no issues observed.
 - ◇ Integration DBs migrated to new hardware.
 - ◇ Few node reboots of ATLR cluster caused by short spikes of very high load (>180), problem under investigation.
 - ◇ Fourth node of Atlas Offline database (ATLR) was evicted from the cluster on June 26th at around 15:10. The eviction was caused by high load generated by Atlas Cool application. Unfortunately service failover mechanism, normally instantaneous, did not re-act properly and 9 Oracle services got stuck following the eviction. Manual intervention was needed to recover affected services. The recovery has been completed by a person on shift at around 16:00. The high load issue has been followed up with ATLAS and was traced down to batch jobs that didn't go through the Frontier server as they should have. Measures have been taken by ATLAS to avoid such issues in the future. The misbehaviour of Oracle service failover mechanism seems to be related to some known bugs, however we are still waiting for confirmation from Oracle Support Services.
 - ◆ CMS:
 - ◇ April PSU was rolled back on CMSONR and CMSR. CMSARC still has the patch - no issues observed.
 - ◇ Missing CMSR3 node added to the cluster on 2nd of June.
 - ◇ Integration DBs migrated to new hardware.
 - ◇ Few nodes reboots of CMSR cluster. The root cause not clear yet. In some cases it could be due to spikes of load but it is also possible that there is a kernel/hardware issue. The issue is followed up with Linux Support.
 - ◆ LHCb:
 - ◇ April PSU was rolled back on LHCBONR and LHCBR.
 - ◇ Some small issues with RUNDB application - high load caused by backup, fixed by relocating RUNDB services.
- Site reports:

Site	Status, recent changes, incidents, ...	Planned interventions
------	--	-----------------------

CERN	<p>* CMS PVSS replication affected between 10 a.m June 15th and 8 p.m. June 17th (peak delay reached 24h) due to huge transactions on PVSS replication that caused series of capture problems, we were forced to change memory parameters several times and restart capture to make the transaction go through streaming. Message passed via coordinators to remind users not to run big transactions by hand on streamed schemas.</p> <p>* Some issues with LCG_DASHBOARD application - some queries (using a specific date conversion function) get invalidated (one occurrence per day or less). Application encounters an exception: 'ORA-01003: no statement parsed'. Service request opened over two weeks ago, problem still under investigation.</p> <p>* On Thursday June 24th at around 21:00 the propagation process sending data to PIC in LHCb conditions data replication setup got stuck unexpectedly. In the morning of June 25th at around 10:20 the hanging process blocked whole replication of conditions data to Tier1 sites. Manual intervention was needed in order to recover the replication. The intervention was completed at around 11:00. The root cause of the hang could not be unambiguously determined by some symptoms indicate that it could be triggered by massive execution of DDL operations in the source database. Unfortunately there is not enough information to follow this case with Oracle Support Services.</p>	None
ASGC	<p>* 3D Database weekly reports deployment - installation completed, but audit_trail is disabled. It will be set during next scheduled intervention.</p> <p>* Incident: asgc3d shutdown/restart not completed due to session problems. Apply process was not working on ASGC and could not be started as the DB did not finish the shutdown. Fixed at 06.29.2010 16:10 UTC</p> <p>* Patch April 2010 status: we are still working on our Oracle RAC testbed verification.</p>	<p>asgc3d Schedule Downtime: 2010-06-29 02:00:00 ~ 2010-06-29 04:00:00 for enabling audit_trail</p>
BNL	<p>* One ORA-07445 entry observed in alert log after applying the PSU APRIL 2010. No problems observed during the deployment of this patch, no service interruption. Database service is working as usual. Created an SR (3-182235985) to follow up ORA-07445 issue. Oracle proposed some actions. Currently considering:</p> <ul style="list-style-type: none"> - We will not rollback the PSU APRIL. - Deploy the one-off patch if this problems appear again (on TOP of PSU APRIL/PSU JULY). 	None
CNAF	<p>New installations (?) for ATLAS and LHCb - 10.2.0.4.4 PSU (CRS and DB) and 10.2.0.4.2 PSU (CRS and DB) respectively. "Weekly reports" tool to be installed.</p>	
KIT	<p>* We decided to not apply the April Oracle security patch (thus cancelled the scheduled intervention on June 29) and wait for the July patch.</p>	<p>Scheduled downtime (severity: outage) on July 6, 7:00-8:00 UTC: network maintenance (switch reconfig). 3D services (LHCb RAC, ATLAS RAC, ATLAS squid, CMS squids) are affected.</p>
IN2P3	NTR	None
NDGF	<p>* The weekly reports have been installed in the Atlas database.</p>	None
PIC	<p>* Applied PSU April 2010 patches on all databases except FTS (there was an issue with a LAN card) last month with no problems.</p> <p>* Audit is turned off in our systems, but we are planning to turn it on</p>	<p>Scheduled interventions on 6th of July, 8th of July and downtime on 20th of July,</p>

	during next scheduled downtime (20th July). * To avoid problems we are going to plan rollback PSU patch on 6th of July (for LFC) and 8th of July (for ATLAS, LHC and TAGS) * The LAN issue is not solved yet, the system is up and running using standby LAN card. During the next scheduled downtime it's planned a hw intervention, that would cause a total stop of the databases.	see report.
RAL	* No problems on the 3D database even though they have been upgraded to 10.2.0.4.4 a month ago. * Castor - disk servers will be updated to SL5 and use the 10.2.0.3 instant client.	
SARA	* We will not install the April PSU because of the trouble it causes.	None
TRIUMF	* APRIL2010 PSU not rolled back. * New Oracle 11gR2 RAC (2 nodes) installed for TAGs. Florbela is in the process of testing & measuring TAGs uploads.	None

- PSU update:
 1. We have opened a service request (3-1826315781) to Oracle Support on the issue we observed on ATLAS databases after applying the PSU April 2010.
 2. According to Oracle Support the symptoms we observed are similar to symptoms of one of known bugs and that there is a patch (6196748) fixing this bug that can be applied on top of the PSU.
 3. The problem is that in order to verify if the aforementioned one-off patch really solves the issue, we need to be able to reproduce the issue in a test environment and this we haven't managed so far. Thus, there is no safe way to verify that the patch will resolve issue of ORA-7445 and spikes of load following.
- Taking into account the current situation, action plan and recommendations are the following:
 1. We keep un-patched all the production databases that experienced the issue after the PSU was applied. The list includes ATLR, ATONR and LHCb databases. Fortunately the security and functional issues that were supposed to be fixed by PSU April 2010 are not very critical so the risk associated with running without the PSU installed is not high.
 2. We keep trying to verify the one-off patch in a test environment.
 3. We will review the situation in the middle of July when new bunch of security patches will be released by Oracle.

AOB

-- JamieShiers - 30-Jun-2010

This topic: LCG > WLCGTier1ServiceCoordinationMinutes100701

Topic revision: r26 - 2010-10-27 - MatthewViljoen1



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