

# Table of Contents

<b>WLCG Tier1 Service Coordination Minutes - 16 June 2011.....</b>	<b>1</b>
Attendance.....	1
Action list review.....	1
Release update.....	1
Data Management & Other Tier1 Service Issues.....	1
CASTOR news.....	2
CERN operations.....	2
Development.....	2
EOS news.....	2
xrootd news.....	2
dCache news.....	2
StoRM news.....	2
FTS news.....	2
DPM news.....	2
LFC news.....	2
LFC deployment.....	2
Experiment issues.....	3
WLCG Baseline Versions.....	3
GGUS issues.....	3
Review of recent / open SIRs and other open service issues.....	3
Conditions data access and related services.....	4
Database services.....	4
AOB.....	6

# WLCG Tier1 Service Coordination Minutes - 16 June 2011

## Attendance

- Local: Simone, Stephane, Alessandro, Dirk, Maite, MariaDZ, Andrea V, Maria, Jamie, Nicolo, Lawrence, Stefan, Andrea S, Maarten, Zsolt, Massimo
- Remote: Jon, Gonzalo, Felix, Andrew - TRIUMF, Patrick, Carlos, Andreas - KIT, Pierre - IN2P3, Roberto, Ken Bloom, Thomas - NDGF, Gareth, Daniele - CNAF, Jhen-Wei - ASGC

## Action list review

## Release update

## Data Management & Other Tier1 Service Issues

Site	Status	Recent changes	Planned changes
CERN	CASTOR 2.1.10 SRM 2.10-x xrootd: all 2.1.10 except CERNT3 (2.1.11)	2.1.11 on CERNT3, SL5 on all stager headnodes (disk servers migrated long ago)	2.1.11 for all (starting with ATLAS and CMS in the next MD/technical stop)
ASGC	CASTOR 2.1.10-0 SRM 2.10-2 DPM 1.8.0-1	None	None
BNL	dCache 1.9.5-23 (PNFS, Postgres 9)	None	Migration from pnfs to Chimera in summer 2011
CNAF	StoRM 1.5.6-3 SL4 (CMS, LHCb,ALICE) StoRM 1.6 SL5 (ATLAS)		Storm release 1.7.0-7 under test, will be on staged rollout shortly on Atlas
FNAL	dCache 1.9.5-23 (PNFS) httpd=1.9.5.-25 Scalla xrootd 2.9.1/1.4.2-4 Oracle Lustre 1.8.3	None	None
IN2P3	dCache 1.9.5-26 (Chimera) on core servers. Mix of 1.9.5-24 and 1.9.5-26 on pool nodes	Upgrade to version 1.9.5-26 on 2011-05-24	
KIT	dCache (admin nodes): 1.9.5-15 (Chimera), 1.9.5-24 (PNFS) dCache (pool nodes): 1.9.5-9 through 1.9.5-24	none	none
NDGF	dCache 1.9.12		
NL-T1	dCache 1.9.5-23 (Chimera) (SARA), DPM 1.7.3 (NIKHEF)		
PIC	dCache 1.9.5-26; PNFS, Postgres 9	Migrated from 25 to 26 last SD on June 8th	Migration to 1.9.12-x planned for August
RAL	CASTOR 2.1.10-0 2.1.9-1 (tape servers) SRM 2.10-2,2.8-6	none	upgrade CASTOR clients on farm to 2.1.10-0 next week and upgrade CASTOR to 2.1.10-1 during next technical stop

TRIUMF	dCache 1.9.5-21 with Chimera namespace	none	none
--------	--	------	------

## CASTOR news

- CASTOR 2.1.11-0 Release: [release notes](#)
- CASTOR SRM 2.11-0 Release: [release notes](#)
- CASTOR xrootd 2.1.11 Release: [release notes](#)

## CERN operations

### Development

## EOS news

ATLAS data migration to EOS going as expected, trying to have the maximum possible throughput. Installed version is 0.1.0. 1.2 M files (out of 5 M) were migrated in 20 days. Investigating the possibility to use GridFTP to bypass SRM slowness.

CMS has created a link to EOS in PhEDEx, transfers done via xrdcp, evaluating which protocol is best.

## xrootd news

## dCache news

Now 1.9.12.4 is the recommended "new" golden release.

## StoRM news

## FTS news

- FTS 2.2.5 in gLite Staged Rollout: [http://glite.cern.ch/staged\\_rollout](http://glite.cern.ch/staged_rollout)
- FTS 2.2.6 in development: FTS 2.2.6 patch  (see list of bugs at the end)

## DPM news

- DPM 1.8.1 in EGI UMD Staged Rollout: [changes](#)

## LFC news

- LFC 1.8.1 in EGI UMD Verification

## LFC deployment

Site	Version	OS, n-bit	Backend	Upgrade plans
ASGC	1.8.0-1	SLC5 64-bit	Oracle	None
BNL	1.8.0-1	SL5, 64-bit	Oracle	None
CERN	1.7.3 64-bit	SLC4	Oracle	Upgrade to SLC5 64-bit pending
CNAF	1.7.4-7	SL5 64-bit	Oracle	
FNAL	N/A			Not deployed at Fermilab
IN2P3	1.8.0-1	SL5 64-bit	Oracle 11g	Oracle DB migrated to 11g on Feb. 8th
KIT	1.7.4-7	SL5 64-bit	Oracle	Oracle backend migration pending
NDGF	1.7.4.7-1	Ubuntu 9.10 64-bit	MySQL	None
NL-T1	1.7.4-7	CentOS5 64-bit	Oracle	

PIC	1.7.4-7	SL5 64-bit	Oracle	
RAL	1.7.4-7	SL5 64-bit	Oracle	
TRIUMF	1.7.3-1	SL5 64-bit	MySQL	

## Experiment issues

- ATLAS LFC consolidation has started at CERN and SARA, done with an ad-hoc tool. The first migration tests took 7 days, went down to 3 after improving the script, but it needs to become faster; now it works at 40 Hz. The intervention is rolling, no there are no service outages. ATLAS is in touch with PES to understand what's the best LFC frontend configuration. The first real migration will happen at CERN and SARA and should be completed by the beginning of July; the goal for the other sites is the end of the year.

## WLCG Baseline Versions

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

## GGUS issues

- The tickets of concern to the experiments were presented as per slides on the agenda page <https://indico.cern.ch/getFile.py/access?contribId=4&resId=2&materialId=slides&confId=143633> Comments at the meeting about GGUS:69739 were recorded in the ticket (see Public Diary of 2011/06/20).
- The final version of "Type of Problem" values' proposal was presented as per slides on the agenda page and Savannah:117206. Value "VO Specific software" should be changed into "Middleware". Future development will be recorded in savannah as always.

## Review of recent / open SIRs and other open service issues

- Report on the KDC problem: <https://savannah.cern.ch/bugs/?82793>
  - ◆ *Reminder:* Since May 26th the Kerberos (KDC) service at CERN observes peaks of very high load originating from batch jobs run by ATLAS users. These jobs were issuing bunch of concurrent file access requests from 'castoratlas', typically via the 'xrscp' copy command. The investigation of the problem involved several server-side (batch, KDC, xrootd) and client-side (ROOT, POOL, ATLAS) components, and required the collaboration of many groups in IT (ES, OIS, PES, DSS) and PH (SFT/ROOT, ATLAS).
  - ◆ *Origine of the problem:* The problem is caused by a fake Kerberos authentication error followed by an incorrect handling of the error by the xrootd-client, resulting in an infinite loop of attempts to reinitialize the credentials with the KDC. The xrootd-client error handling bug had already being fixed mid 2010 but not yet available in version used by ATLAS. The Kerberos authentication error is instead not yet fully understood and is still under investigation; the current hypothesis is that it may be due to a bug in the Kerberos libraries.
  - ◆ *ROOT versions affected:* The ROOT versions affected by the xrootd-client bug are all the versions prior to 5-28-00a, therefore including the version 5.26/00e - used in production by ATLAS - and the version 5.27/06 used by CMS . The fix has been already included in the patch branches of these versions, 5-26-00-patches and 5-27-06-patches . LHCb (using version 5.28/00b) and ALICE (version 5.28/00d) are not affected by the problem.
  - ◆ *Deployment of the client-side fix:* the deployment of the fix for ATLAS and CMS has been discussed during the Architect Forum meeting on June 16th, 2011.
    - ◇ *ATLAS:* The ROOT team is going to provide a new tag 5-26-00f and the associated binaries in the standard LCG area under AFS; the binaries will be backward-compatible with 5-26-00e. A patch with only the fixed version of the

affected executable (xrdcp) and of the affected libraries (libXrdClient.so, libXrdSeckrb5.so) will be also made available for those users who cannot move to 5.26/00f .

- ◇ *CMS*: Binaries for the CMS version of ROOT 5-27-06 will; be re-built with the relevant patch (<http://root.cern.ch/viewcvs?view=rev&revision=39740>) included.
- ◇ *Update of the default 'xrdcp' on lxplus/lxbatch*: It has also been remarked that an old version of 'xrdcp' is available in the system standard paths of lxplus/lxbatch. The IT/DSS team will take care to upgrade this version to a bug-free version to avoid any fortuitous occurrence of the problem.
- ◆ *Remaining actions*: the full solution of this problem requires fixing the fake authentication failures which trigger the KDC floods. This is under investigation.

## Conditions data access and related services

### Database services

- Experiment reports (long list covering issues since 5th of May):
  - ◆ ALICE:
  - ◆ ATLAS:
    - ◇ First instance of Atlas offline database (ADCR) has crashed on Sunday (08.05). Issue has been caused by internal database error and is now being under investigation. Services were available on the surviving nodes while the instance restarted and were relocated back to instance one after it came back into operation.
    - ◇ We had three hangs of ATLAS offline DB (ADCR) during which the service was not available: Monday 16th between 16:25 and 17:10, Monday 16th between 21:50 and 23:30 and Tuesday 17th between 1:50 a.m. and 2:40 a.m. No data loss occurred except for all uncommitted data. All incidents were caused by unusual reaction of ASM on a broken disk (itstor737 disk 3). ASM did not properly initiated a rebalance operation during the first incident and was affected by some problems during second and third. After the incidents a normal rebalance has finished and we were trying to forcefully evict the problematic disk. SR has been opened on this issue.
    - ◇ On 18th of May around 11:20 ADCR DB experienced another disk failure during rebalancing operation which did not finish. Decision was taken to switchover the DB to the standby cluster. Switchover completed successfully after several minor issues and the DB was back operational at 13:05. IN2P3 reported that AMI applications are not able to reach the DB. It turned out that DB was not visible outside of CERN. We requested the port on the firewall to be opened and it was done the next day (19th of May in the morning).
    - ◇ The INT8R Atlas integration database has been migrated to Oracle 11g (version 11.2.0.2) on Monday (30.05).
    - ◇ Atlas offline (ADCR) database has been switched back to original hardware on Tuesday (31.05). Switchover operation required 1h database downtime.
    - ◇ ADCR database, hosting Atlas DQ2, PANDA and PRODSYS services has suffered full downtime on Monday (06.06) from 13:20 till 14:00. The cause of downtime has been a DB hang caused by the ASM layer following a double disk failure occurred in the morning. A full restart of the DB and ASM unblocked the hang. We are following up with Oracle support on the causes of the hang. We are also investigating on the rate of disk failures observed in this system recently (2 disks have failed over the long WE in addition to the 2 failed on Monday)
    - ◇ There has been a failure of one of redundant Fibre Channel Switches on 09.06 to which storages of ATLAS offline DB and ATLAS online DB are connected. Issue is now being investigated by FC-sunsupport. We are also discussing with ATLAS a schedule for replacing the swich.

- ◆ CMS:
  - ◇ 3rd node of the CMS offline production database (CMSR) rebooted on 14th of June at 14:45, most likely due to an issue with HBA's firmware or driver. Due to the reboot some sessions of several CMS applications failed. Fortunately none of the 3 main applications running on the DB were affected. The machine re-joined the cluster around 15:00.
- ◆ LHCb:
  - ◇ Streams issues - see below
- ◆ Streams:
  - ◇ Due to SARA's Advance Queuing configuration problems which were caused by January migration of the database back to RAC (consequence of the August 2010 storage failure), replication of LHCb conditions data was frozen from Saturday (22.05) midnight until Monday 11 o'clock. Problem was temporarily fixed by migration of buffered queue to different instance. Permanent solution is being discussed with Oracle Support.
  - ◇ After Tuesday (24.05 10:00) SARA's database upgrade to 10.2.0.5 both LHCb replication (conditions data and LFC) were hit by streams related Oracle bug. Solution: Patches fixing the bug has been applied.
  - ◇ After extended intervention on Tuesday (24.05 7:00) on database storage at IN2P3, LFC replica became partially inconsistent - at least one transaction was missing. After skipping some other transactions, streams managed to continue replication without further errors. Full LFC data consistency at IN2P3 has been restored on Wednesday afternoon (25.05).
  - ◇ Replication to SARA was down from Sunday 29.05 (around 9:45 am) to Monday 30.05 (around 10:00 am) due to a spanning tree problem in a part of SARA's network. The replica at SARA was out of sync during the period of the issue which could potentially cause some problems for the jobs using this replica.
  - ◇ The replication of ATLAS conditions data to BNL hanged twice: Monday night and Tuesday night for 6 hours. Root cause of the problem is gathering statistics job which due to cross cluster contention blocks processing of big transactions by streams components. The temporal workaround is to kill manually the locking job. This problem is being investigated by Oracle Support without any results so far.

- Site reports:

Site	Status, recent changes, incidents, ...	Planned interventions
ASGC	Restore FTS Database: Unexpected power cut.	None
BNL	- Contention observed between the apply process and the gather_stats_jobs (similar issue reported in 04/13/11) in the Conditions DB. This issue is now being followed up along with ORACLE via SR 3-3535183751, information about this recent problem and different logs has been collected and sent to Oracle. No database service was disrupted during this issue. - A test 11g2.0.2R two nodes RAC was deployed no ACFS used.	None
CNAF	Nothing to report	None
KIT	May 24th - Migration of FTS DB to new hardware (new 3-node RAC).	None
IN2P3	Nothing to report	None
NDGF	Nothing to report	None
PIC	- Two weeks ago (31th of may and 1st of June) we had applied latest CPU patch and kernel updates. - Last week (8th of June) during a SD we'd upgraded FW of Storage systems.	None
RAL	RAL is in the process of testing dataguard implementation for Castor and LFC/FTS. This will take the next couple of weeks.	None
SARA	Nothing to report	

		No interventions
TRIUMF	Applied APR2011 CPU	None

## AOB

-- AndreaSciaba - 15-Jun-2011

---

This topic: LCG > WLCGTier1ServiceCoordinationMinutes110616

Topic revision: r19 - 2011-06-30 - 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](#)