

# Table of Contents

<b>WLCG Tier1 Service Coordination Minutes - 29 September 2011.....</b>	<b>1</b>
Attendance.....	1
Action list review.....	1
Release update.....	1
Data Management & Other Tier1 Service Issues.....	1
Other site news.....	2
CASTOR news.....	2
CERN operations and development.....	2
EOS news.....	2
EOS as a production service.....	2
xrootd news.....	2
dCache news.....	2
StoRM news.....	2
FTS news.....	2
DPM news.....	3
LFC news.....	3
LFC deployment.....	3
Experiment issues.....	3
WLCG Baseline Versions.....	4
Status of open GGUS tickets.....	4
Review of recent / open SIRs and other open service issues.....	4
Conditions data access and related services.....	4
Database services.....	4
AOB.....	5

# WLCG Tier1 Service Coordination Minutes - 29 September 2011

## Attendance

## Action list review

## Release update

## Data Management & Other Tier1 Service Issues

Site	Status	Recent changes	Planned changes
CERN	CASTOR 2.1.11-5 (SL5) for CMS and PUBLIC, others are on 2.1.11-2; SRM 2.10-x (SL4); xrootd: 2.1.11-1 FTS: 5 nodes in SLC5 3.7.0-3; 7 nodes in SLC4 3.2.1 EOS -0.1.0/xrootd-3.0.4		CASTOR 2.1.11-6 has been officially released (maintenance release addressing some issues with the transfermanager and the tapegateway components). This is scheduled for CASTORCMS on Monday (Oct 3rd). Soon this will be deployed at least on PUBLIC.
ASGC	CASTOR 2.1.11-5 SRM 2.11-0 DPM 1.8.0-1	22/09: CASTOR upgrade, no issues encountered	None
BNL	dCache 1.9.5-23 (PNFS, Postgres 9)	None	Transition from PNFS to Chimera during next LHC TS
CNAF	StoRM 1.7.0 (Atlas) Storm 1.5.0 (other endpoints)	the present version contains various patches which will be included in the new StoRM release, currently under certification	
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		
IN2P3	dCache 1.9.5-26 (Chimera) on core servers. Mix of 1.9.5-24 to 1.9.5-28 on pool nodes		Increase of RAM on Chimera node next week (site downtime on 2011-10-04)
KIT	dCache (admin nodes): 1.9.5-27 (ATLAS, Chimera), 1.9.5-26 (CMS, Chimera) 1.9.5-26 (LHCb, PNFS) dCache (pool nodes): 1.9.5-6 through 1.9.5-27		
NDGF	dCache 1.9.14 (Chimera) on core servers. Mix of 1.9.13 and 2.0.0 on pool nodes.		
NL-T1	dCache 1.9.12-? (Chimera) (SARA), DPM 1.7.3 (NIKHEF)	Upgrade to 1.9.12 solved problem with crashing pools	
PIC	dCache 1.9.12-10 (last upgrade to patch release on 13-Sep);		

	PNFS on Postgres 9.0		
RAL	CASTOR 2.1.10-2 2.1.10-0 (tape servers) SRM 2.10-0	None	None
TRIUMF	dCache 1.9.5-21 with Chimera namespace	None	Upgrade dCache to 1.9.5-28 and FTS to SL5 3.7.0-3 next Wednesday

## Other site news

### CASTOR news

#### CERN operations and development

- CASTOR 2.1.11-6 has been officially released (Release notes [↗](#)). This is a maintenance release addressing some issues with the transfermanager and the tapegateway components. This is scheduled for CASTORCMS on Monday (Oct 3rd). Soon this will be deployed at least on PUBLIC.

### EOS news

Dirk: informed that Bestman support will close at end of year. The plan is to replace SRM with direct GridFTP access. Compatible versions of FTS and lcg-util being prepared.

Simone: for ATLAS we need the FTS fix for checksum. Writes to SE rely on SRM, can be changed. Removal of files relies on SRM. Will check other bits and pieces.

Dirk: need to find a realistic plan. As soon as certification is done for FTS and lcg-utils we should start testing.

Maria: this item will be added as a regular point of discussion

Ian: the loss of Bestman affects many SEs around, for example Hadoop sites.

#### EOS as a production service

Maria: in terms of production quality services, could you clarify if there are changes in support?

Massimo: want to support CASTOR and EOS in the same way. Only difference: CASTOR has a formal piquet but EOS is best effort.

Dirk: we'd like to show that we don't need a piquet.

### xrootd news

### dCache news

### 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 released in EMI-1 Update 6 on Sep 1
  - ◆ restart/partial resume of failed transfers
- FTS 2.2.7 being prepared for certification: FTS 2.2.7 patch [↗](#) (see list of bugs at the end)
  - ◆ includes new overwrite logic
  - ◆ to be released for gLite and EMI

## DPM news

- DPM 1.8.2-2 - a problem was found in certification, fixed, and the release rebuilt
- DPM 1.8.2-3 ready for final certification (code already validated extensively and in use at some sites)
  - ◆ fast dpm-drain
  - ◆ filesystem selection algorithm configurable by admin
  - ◆ support for central banning (Argus)
  - ◆ Allow definition of the number of threads for DPM and SRM servers at startup time
  - ◆ <https://savannah.cern.ch/patch/?5005>
  - ◆ <https://savannah.cern.ch/patch/?5006>
- Monthly releases of new unstable components can be followed on the blog:
  - <https://svnweb.cern.ch/trac/lcgdm/blog>
  - ◆ This covers NFSv4.1, WebDAV, Nagios, Catalogue synchronisation & 'perfsuite'.

## LFC news

- LFC 1.8.2-2 has been certified
  - ◆ fix for read-only replica operation (LHCb)
  - ◆ support for central banning (Argus)
  - ◆ <https://savannah.cern.ch/patch/?5003>
  - ◆ <https://savannah.cern.ch/patch/?5004>

## 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	1.8.2-0 64-bit	SLC5	Oracle	Upgrade to SLC5 64-bit only pending for lfcshared1/2
CNAF	1.7.4-7 (ATLAS, to be dismissed) 1.8.0-1 (LHCb, recently updated)	SL5 64-bit	Oracle	
FNAL	N/A			Not deployed at Fermilab
IN2P3	1.8.0-1	SL5 64-bit	Oracle 11g	Plan to migrate to 1.8.0-2 asap
KIT	1.7.4-7	SL5 64-bit	Oracle	Oracle backend migration pending
NDGF	1.7.4.7-1	Ubuntu 10.04 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	None

## Experiment issues

(MariaDZ) There was an ATLAS presentation on ALARM response requirements that became a Critical Services discussion. The following email was sent to wlcg-service-coordination e-group for comments: *This contains Tier0 only but not very old so maybe it can be updated to cover all Tier1s as well and required response times:*

*[https://twiki.cern.ch/twiki/bin/view/LCG/WLCGCriticalServices#Tier0\\_Critical\\_Services\\_Generic](https://twiki.cern.ch/twiki/bin/view/LCG/WLCGCriticalServices#Tier0_Critical_Services_Generic) . I take the liberty to send it to the list because I was the most recent editor.*

## WLCG Baseline Versions

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

## Status of open GGUS tickets

- The open tickets of concern for ATLAS and CMS concerned US Tier2s. They were included in the presentation, nevertheless, to prompt the relevant Tier1s to help solving these long lasting issues.
- There was a short presentation on the Type of Problem values.

All slides available from

<https://indico.cern.ch/materialDisplay.py?contribId=5&materialId=slides&confId=156754>

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

### Conditions data access and related services

### Database services

- Experiment reports:
  - ◆ ATLAS:
    - ◇ On Monday (19.09) at 10AM ATLAS streams replication to T1s got stuck. The reasons of the problem were Oracle internal queuing processes which were preventing from accessing the queues. All blocking process had to be killed and affected database was restarted. Replication service was available at 2PM. The service request to the Oracle has been opened as the same problem was observed 3 weeks ago after applying CPU July on downstream capture database.
    - ◇ Gancho updated the TWiki (<https://twiki.cern.ch/twiki/bin/viewauth/Atlas/DatabaseVolumes>) with the latest projections to 2014 for ATLAS Conditions DB Volumes to Tier-1s. This was prompted by a request from Carlos Gamboa, who was doing hardware purchase planning. (Elizabeth)
  - ◆ CMS:
    - ◇ On Thursday (15.09) second node of CMSR was rebooted and on Wednesday (21.09) all nodes but first were rebooted by Clusterware. The only indication of cause is high load which is growing very fast about 2 minutes before reboot. Unfortunately existing logs and trace files do not allow for determining the root cause. Oracle OS Watcher software will be deployed today to gather additional diagnostic information in case the problem re-appears.
    - ◇ On Friday 23rd September around 5:30 in the morning 5 out of 6 nodes of CMS online production database went down due to a failure of the cluster interconnect switch. The switch has been fixed by CMS sysadmins around 9am and by 9:30 the database was fully available again. In order to limit the impact of similar issues in the future CMS deployed a secondary switch dedicated for the cluster interconnect.
    - ◇ On Tuesday 27th September at 14:00 CMS offline database (CMSR) hung completely following a vendor mistake during replacement of a broken disk in one of the disk arrays used by the database. Even though normally such a problem should be transparently handled by Oracle ASM software, this time, for a reason which is still not understood, it caused unavailability of the whole system. We suspect issues with the disk array's controller and plan to drain the disk array and examine it.

◇ As a side effect of the hang of CMSR on 27th September one of the tablespaces used by CMS Dataset Bookkeeping application was put offline by Oracle make the data stored in there unavailable. The problem has been reported by CMS at 20:30 on was fixed within half an hour. Additional monitoring is being deployed to discover such problems quicker.

◆ General:

◇ New procedure has been developed in order to crosscheck the content of the streams dictionary between source and replica databases. It has been deployed as a weekly database job in each LCG replication environment. This will provide low level validation of replication configuration and detect potential problems with data consistency (which we observed few times in the past).

• Site reports:

Site	Status, recent changes, incidents, ...	Planned interventions
BNL	-- Applied RHEL 5 Operating System (OS) kernel security patches. -- Applied quarterly Oracle Critical Update (CPU) 2011. -- Updated oracle Automatic Storage Management (ASM) file system libraries.	-- To apply proposed patch from Oracle (P6011045) in Conditions Database
CNAF	Applied RHEL 5 OS latest kernel/updates Applied JUL PSU on FTS and LHCb	
KIT		Plans to apply CPU July 2011 for ATLAS ConditionDB and LFC/FTS-DB around 18-19th of October. Intervention will not be transparent for ATLAS as a short shutdown will be required to fix some issues with spfile (after last migration with DG).
IN2P3		
PIC	July CPU patch was applied on 20th and 21th of September on all databases. ATLAS database was definitely stopped last week.	None
RAL	Incident on CASTOR on the 27th, resolved and root cause under investigation.	None
SARA	Nothing to report	No interventions
TRIUMF	Nothing to report (not attending today)	Plan to apply July 2011 CPU on Oct 5.

## AOB

-- AndreaSciaba - 28-Sep-2011

---

This topic: LCG > WLCGTier1ServiceCoordinationMinutes110929

Topic revision: r18 - 2011-09-29 - MariaDimou



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](#)