

Table of Contents

WLCG Tier1 Service Coordination Minutes - 11 March 2010.....	1
Data Management services & Other Tier1 Service Issue.....	1
Status of FTS 2.2.3 deployment.....	1
Storage systems: status, recent and planned changes.....	1
Other information.....	2
CASTOR.....	3
dCache.....	3
StoRM.....	3
LFC.....	3
Experiment issues.....	3
Service Update: support, problem handling and alarm escalation.....	3
Conditions data access and related services.....	4
Experiment Database Service Issues.....	4
WLCG Baseline Versions.....	5
AOB.....	5

WLCG Tier1 Service Coordination Minutes - 11 March 2010

Data Management services & Other Tier1 Service Issue

Status of FTS 2.2.3 deployment

- CERN: FTS 2.2.3 installed only on FTS22-T0-EXPORT but not exactly matching a gLite release (patches were applied individually).
Next steps:
 - ◆ March 16th: synchronize FTS-T0-EXPORT/FTS-T1-IMPORT to the latest gLite release;
 - ◆ March 17th: reinstall FTS-T2-SERVICE with the latest gLite release.
- ASGC: **UPDATE 22-03-10**: FTS 2.2.3 installed, it will be put in production on 24-03-10.
- BNL: FTS 2.2.3 installed on February 23rd.
- CNAF: FTS 2.2.3 installed on March 9th.
- FNAL: FTS 2.2.3 installed.
- IN2P3: Upgrade scheduled for March 18th (depending on some final tests). **UPDATE 22-03-10**: FTS 2.2.3 installed on 18-03-10.
- KIT: FTS 2.2.3 installed on February 18th.
- NDGF: Upgrade to be scheduled "soon". **UPDATE 24-03-10**: FTS 2.2.3 installed on March 16th.
- NL-T1: Upgrade ongoing. FTS went fine, but we observe problems starting the FTS agents. In contact with the developers. **UPDATE 22-03-10**: FTS 2.2.3 installed on 22-03-10.
- PIC: Upgrade scheduled for March 18th. **UPDATE 22-03-10**: FTS 2.2.3 installed on 19-03-10.
- RAL: Upgrade to be scheduled within the next two weeks. We completed the testing, are reviewing the deployment plan and discussing with the experiments to decide a date (17/3 will be proposed). **UPDATE 22-03-10**: FTS 2.2.3 installed on 17-03-10.
- TRIUMF: Upgrade scheduled for March 24th (if the results of the ongoing tests are positive). **UPDATE 25-03-10**: FTS 2.2.3 installed on 24-03-10.

Summary:

- 5 sites upgraded
- 1 site is upgrading
- 3 sites scheduled the upgrade
- 2 sites have not scheduled the upgrade

UPDATE 25-03-10

Summary:

- 12 (all) sites upgraded

Storage systems: status, recent and planned changes

Site	Status	Recent changes	Planned changes
CERN	CASTOR 2.1.9-4 (all) SRM 2.8-6 (ALICE, CMS, LHCb) SRM 2.9-2 (ATLAS)	8/3: ALICE Castor and SRM upgrade 8/3: LHCb Castor and SRM upgrade 9/3: ATLAS SRM upgrade 11/3: Castor xroot SSL plugin on castoratlas-xrdssl	

ASGC	?	?	?
BNL	dCache 1.9.4-3	none	none
CNAF	CASTOR 2.1.7-27 (ALICE) SRM 2.8-5 (ALICE) StoRM 1.5.1-2 (ATLAS, LHCb) StoRM 1.4 (CMS)	StoRM upgraded for ATLAS and LHCb	15/3: StoRM upgrade for CMS
FNAL	dCache 1.9.5-10 (admin nodes) dCache 1.9.5-12 (pool nodes)	none	none
IN2P3	dCache 1.9.5-11 with Chimera	none	April: configuration update and tests on tape robot, integrate new drives Q3-Q4: one week downtime to upgrade the MSS
KIT	dCache 1.9.5-15 (admin nodes) dCache 1.9.5-5 - 1.9.5-15 (pool nodes)	1-4/2: Chimera migration for ATLAS 18/2: dCache update on admin nodes	none
NDGF	dCache 1.9.6 dCache 1.9.5 (some pools) dCache 1.9.7 (some pilot admin nodes)	none	none
NL-T1	dCache 1.9.5-16	none	next LHC stop: migrate three dCache admin nodes to new hardware
PIC	dCache 1.9.5-15	none	none
RAL	CASTOR 2.1.7-24 (stagers) CASTOR 2.1.8-3 (nameserver central node) CASTOR 2.1.8-17 (nameserver local node on SRM machines) CASTOR 2.1.8-8, 2.1.8-14 and 2.1.9-1 (tape servers) SRM 2.8-2	used a replicating svcclass for CMS "hot" files	none
TRIUMF	1.9.5-11	none	none

Other information

- BNL is testing dCache 1.9.6-2 (double SRM installation with Terracotta).
- FNAL will upgrade the admin nodes to get updates to info cells and pin manager, but not planned yet (may be months away).
- NDGF will upgrade to dCache 1.9.7 probably in a few weeks after having verified stability and updated the internal documentation.
- PIC will upgrade dCache once a couple of issues is addressed:
 - ◆ too much memory usage by pools at startup
 - ◆ the xrootd implementation does not implement read permissions, which is an issue for a multi-VO installation.
- RAL will upgrade to CASTOR 2.1.8/2.1.9 but there is no schedule for this. We experienced some problems with CASTOR:
 - ◆ recurrence of Oracle bigId problem seen in VDQM and SRM databases
 - ◆ Tape Migration problems due to misconfiguration
 - ◆ Memory usage for LSF mbatchd daemon
 - ◆ Draining of disk servers competing with production work

CASTOR

SRM 2.9 is being tested by ATLAS. This version brings some new features ([link1](#), [link2](#)).

dCache

dCache 1.9.7 was released on March 9th. The number of internal changes of dCache from the 1.9.6 series is small. Instead, the focus of 1.9.7 is to let go of some old baggage: many configuration files have changed, libraries have been upgraded and some deprecated components have been removed. See the release notes [here](#).

StoRM

Version 1.5 was released on 26/2 for SL4; the SL5 will be released soon. New features:

- support for hierarchical storage (storage classes T1D0 and T1D1) based on GPFS and TSM
- the checksummer service which acts on behalf of StoRM to perform the computation of checksum
- various enhancements and bug fixes

Note that from now on bugs will be fixed only in version 1.5.

LFC

LFC 1.7.3 has been released for SL5. The SL4 version will follow.

Experiment issues

Service Update: support, problem handling and alarm escalation

Following the reorganisation of IT department and some questions that had been raised at the daily WLCG operations meetings, the DB, DSS and PES groups of CERN IT were asked to summarize how alarm tickets in particular were handled.

The main points are:

- People within the experiments authorized to issue alarm tickets for critical services should continue to do so as and when they consider that this is required;
- As has been the case since the alarm mechanism was introduced (since when the number of real - as opposed to test - alarm tickets has been extremely low), each such ticket is analyzed at the subsequent WLCG Management Board in the Operations report against the targets (for the Tier0) that are summarized below:

Time Interval	Issue	Tier0 Target
30'	Operation response to alarm / phone call to +4122 767 5011	99%
1 hour	Operator response to alarm / call to x5011	100%
4 hours	Expert intervention in response to above	95%
8 hours	Problem resolved	90%
24 hours	Problem resolved	99%

Although the number of problems has been too low to establish reliable statistics in all real cases concerning the Tier0 these targets have been met.

The slides from the various groups can be found [here](#).

ATLAS pointed out a discrepancy between the support levels of CASTOR - for which an on-call ("piquet") service is provided - and services that it depends on, such as the CASTOR DBs for which an on-call service is not currently provided by IT-DB (although the non-CASTOR physics DBs are covered).

The Physics Database services have on-call experts with 24x7 coverage (online & offline databases), whereas Database replication has 8x5 coverage with best-effort outside working hours.

A presentation on Service Deployment and Support was made based on input from ATLAS. This included the following key points:

- Major upgrades should be carried out during LHC technical stops (it being understood that the current WLCG procedures for announcing downtimes have to be revised to cater for the needed service + experiment flexibility in rescheduling associated downtimes);
 - ◆ The recently established procedure for handling interventions, with a "change assessment" followed by testing of the functionality by the experiment + IT-ES being much appreciated
- The expectation that Critical Services - such as those identified by WLCG bodies already at the time of CCRC'08 (wiki page) - have 24x7 coverage
- It was noted that GGUS TEAM and ALARM tickets work well for ATLAS but that a GGUS template, handling the different information that various sites require for ALARM tickets, be provided (based on keywords).

Conditions data access and related services

- ATLAS weekly FroNTier meetings [☞](#)

Experiment Database Service Issues

- Experiments reports:
 - ◆ ALICE: nothing to report
 - ◆ ATLAS: incident on streaming from online to offline. Caused by an application misbehavior: several million rows inserted causing 14 hours delay in the replication. ATLAS is investigating the problem.
 - ◆ CMS: streaming issues with PVSS replication caused by users mistakes: views are created referencing tables which do not exist at the destination database.
 - ◇ SLS modification in order to distinguish between the PVSS and the conditions replication setups.
 - ◆ LHCb: last week LHCb conditions requested one schema to be restored. Operation was not properly replicated due to the default rules used on the apply processes in order to fix a Streams bug. Eva is reviewing the rules now that the bug has been fixed.
 - ◇ LHCb-LFC apply process aborted at RAL due to a inconsistency found on the destination database. It seems that a new user was added at RAL, however this should not be allowed (RAL must be a read-only replica).
- Tier0:
 - ◆ Hw resources: new hw being stress tested. Foreseen to finish in 2 weeks. Databases migration will start in April (if everything goes fine).
 - ◆ Streams monitoring improvements implemented:
 - ◇ auto restart of the capture processes under known failures
 - ◇ monitoring and notifications in case of long running transactions
 - ◇ working on apply handlers in order to avoid apply failures due to grant operations (very frequent problem in the CMS replication environments)
 - ◆ Tier 1 reports: deployed in test environment by Carlos at BNL, there will be a presentation in 2 weeks.

- Sites status:
 - ◆ RAL: January security patches scheduled for early next week
 - ◆ CNAF: migration of the LHCb and LFC clusters ongoing
 - ◇ Question about Frontier installation for ATLAS conditions: Barbara will contact Simone and Flavia.
 - ◇ Question about Streams future: no plan to replace Streams in the near future. IT-DB would like to start testing Golden Gate (recently acquired by Oracle) and analyze the benefits and implications on the current replication systems. With the use of Active DataGuard in Oracle 11gR2 we might avoid the use of Streams at Tier0 (replication between online and offline databases).
 - ◇ VOMS replication using Streams? It is up to voms people to decide if they need replication and which is the best technology to do that.
 - ◆ Gridka: January security patches applied last week.
 - ◆ PIC: TAGs database being prepared
 - ◆ SARA: January security patches scheduled for next week
 - ◆ BNL: January security patches scheduled for next week

WLCG Baseline Versions

- Release report: deployment status wiki page
- WLCG Baseline versions: table
 - ◆ Proposed change for STORM from 1.4 to 1.5 (support for 1.4 has stopped)
- Release update - wiki

AOB

This topic: LCG > WLCGTier1ServiceCoordinationMinutes100311

Topic revision: r11 - 2010-03-25 - AndreaSciaba



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