

Table of Contents

WLCG Tier1 Service Coordination Minutes - 2nd December 2010.....	1
Attendance.....	1
LHC machine - shutdown and 2011 startup plans.....	1
Security updates.....	1
GGUS news.....	1
CERNVM-FS.....	1
Release Update.....	1
WLCG Baseline Versions.....	2
Data Management & Other Tier1 Service Issues.....	2
Other site news.....	2
CASTOR news.....	2
CERN operations.....	2
Development.....	3
xrootd news.....	3
dCache news.....	3
StoRM news.....	3
FTS news.....	3
DPM news.....	3
LFC news.....	3
LFC deployment.....	3
Experiment issues.....	3
BDII deployment plan.....	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
Experiment Database Service Issues.....	4
Dates & topics for future meetings.....	5
AOB.....	5

WLCG Tier1 Service Coordination Minutes - 2nd December 2010

Attendance

LHC machine - shutdown and 2011 startup plans

Talk postponed.

Security updates

Romain reported about two incidents (details not given on purpose).

GGUS news

Thorsten reported about two problems. On November 16 a SOAP component caused the web services to block. A guess was that it could be due to the simultaneous submission of several tickets to Spain but an attempt to reproduce it did not cause any problem. The logs did not contain any useful information and the problem is not yet understood.

On November 26 the GGUS Oracle database was unavailable for 1.5 hours, due to moving to a new database with high availability and a newer version of Oracle (11). There was a misconfiguration of this HA cluster which was fixed this morning during a short "at risk" downtime.

CERNVM-FS

Apart from what is said in the slides, it was clarified that the stress test foreseen at RAL will involve all the affected parties, including the central repository at CERN.

In general this is still an experimental service because support from CERN is not yet full (maybe it will be so after January). RAL will set up a mirror web repository but the release of software will still happen at CERN.

LHCb is starting doing tests at NIKHEF where the site will change the environment variable pointing to the software area (NFS or CVM-FS) as requested. It should be possible to do the same at RAL.

Joel asked about the plans of other Tier-1 sites: Pierre said that IN2P3 is interested but they must give priority to solving their AFS problems. KIT was not available to answer during the meeting.

Ian F. said that CMS plans to start using it at Tier-3 sites and possibly extending it to other sites later.

Stephane said that the first ATLAS tests are encouraging (results will be shown tomorrow at an ATLAS meeting).

Ron said that for SARA they don't have tests planned but will discuss with NIFHEF.

Release Update

The main point was the new version of the WMS, which allows to use VOMS from gLite 3.2. Sites are urged to upgrade ASAP to be able to upgrade all VOMS servers to 3.2.

Many patches in staged rollout (DPF, LFC, glexec, etc.). A CREAM patch had to be rejected.

WLCG Baseline Versions

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

Data Management & Other Tier1 Service Issues

Site	Status	Recent changes	Planned changes
CERN	CASTOR 2.1.9-8 (ATLAS) CASTOR 2.1.9-9 (ALICE, CMS and LHCb) SRM 2.9-4 (all) xrootd 2.1.9-7		
ASGC	CASTOR 2.1.7-19 (stager, nameserver) CASTOR 2.1.8-14 (tapeserver) SRM 2.8-2	29/11: network maintenance ↗ , storage services stopped	None
BNL	dCache 1.9.4-3 (PNFS)	None	None
CNAF	StoRM 1.5.4-5 (ATLAS, CMS, LHCb,ALICE)		
FNAL	dCache 1.9.5-23 (PNFS) Scalla xrootd 2.9.1/1.4.2-4	None	None
IN2P3	dCache 1.9.5-22 (Chimera)		
KIT	dCache 1.9.5-15 (admin nodes) (Chimera) dCache 1.9.5-5 - 1.9.5-15 (pool nodes)		
NDGF	dCache 1.9.7 (head nodes) (Chimera) dCache 1.9.5, 1.9.6 (pool nodes)		
NL-T1	dCache 1.9.5-23 (Chimera) (SARA), DPM 1.7.3 (NIKHEF)		
PIC	dCache 1.9.5-23 (PNFS)		
RAL	CASTOR 2.1.7-27 and 2.1.9-6 (stagers) 2.1.9-1 (tape servers) SRM 2.8-2 and SRM 2.8-6	Added 2 new SRM backends for ATLAS	ATLAS upgrade to 2.1.9-6 on 6-8/12/10
TRIUMF	dCache 1.9.5-21 with Chimera namespace		

Other site news

The FTS channels to TW-FTT were created at all relevant sites.

CASTOR news

CERN operations

There will be a deployment campaign in January. Now busy with closing the new release and with testing and planning.

[ACTION] It would be good to have from the experiments information about low and high points of activity foreseen for January.

Development

No significant news.

xrootd news

dCache news

No significant news.

StoRM news

FTS news

FTS 2.2.5 still in certification.

DPM news

No significant news.

LFC news

No significant news.

LFC deployment

Site	Version	OS, n-bit	Backend	Upgrade plans
ASGC	1.7.2-4	SLC4 64-bit	Oracle	Testing ongoing, upgrade by the end of the year
BNL	1.7.2-4	SL4	Oracle	1.7.4 on SL5 postponed to January
CERN	1.7.3 64-bit	SLC4	Oracle	Will upgrade to SLC5 64-bit by the end of the year
CNAF	1.7.2-4	SLC4 32-bit	Oracle	1.7.4 on SL5 64-bit in November
FNAL	N/A			Not deployed at Fermilab
IN2P3	1.7.4-7	SL5 - 64 bits	Oracle	
KIT	1.7.4	SL5 64-bit	Oracle	
NDGF				
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	

[NOTE]: BNL and CNAF should better upgrade to 1.8.0 because of the VOMS library memory leaks in 1.7.4.

Experiment issues

Simone reviewed the issues ATLAS has experienced with dCache at IN2P3. Pierre explained that the suggestion from dCache that it could be related to using Solaris was actually wrong (it mistakenly referred to another problem). There is no real evidence that the problems are the consequence of the dCache upgrade and they still need to be understood.

Jon reported as something potentially interesting for all dCache sites that FNAL had major process scheduling problems with the kernel coming with SL5 and they solved them by using the latest available kernel. dCache developers were not involved and it would be useful to let them aware of FNAL's findings.

BDII deployment plan

Some points were discussed during the talk. Highlights follow.

The MoU prescriptions for "other services" (like the BDII) require 98% availability at prime hours, 97% otherwise.

Published data should be not more than 15' old (1 hour was considered too old).

It was clarified that the quality of service of the top BDII at CERN should be no less than at Tier-1 sites and that best effort support does not imply a lower quality of service.

Finally it was stressed that best practices and requirements should be clearly separated in the document (the requirements must be associated to specific metrics).

Site	Plan
NL-T1	There are in total more than 5 top-level BDII's at the NL-T1. In LCG_GFAL_INFOSYS at both SARA and NIKHEF there are three top-level BDII's configured. At NIKHEF two BDII's from NIKHEF and one BDII at SARA configured. At SARA there are two SARA BDII's and one NIKHEF BDII in LCG_GFAL_INFOSYS
US ATLAS-T1	Working with OSG on the deployment of a resilient and performant top-level BDII infrastructure in the US

Status of open GGUS tickets

Review of recent / open SIRs and other open service issues

Conditions Data Access and related services

Dave reported an ATLAS Frontier server (and database) overload. The database server had to be rebooted. Alessandro offered a possible explanation, as the software used in the reprocessing campaign had a bug[?] and jobs were repeatedly connecting to the database instead than connecting to Frontier.

Experiment Database Service Issues

- Experiment reports:
 - ◆ ALICE:
 - ◇ Nothing to report
 - ◆ ATLAS:
 - ◇ Atlas offline database suffered from 4 instance reboots this week. Instance 4 rebooted on 28.11, 30.11, 02.12 morning around 4AM and instance 3 rebooted on 02.12 around 11:30AM. Initially high load caused by COOL application was suspected as rootcause however there have been corresponding I/O errors and spikes of physical writes observed on 02.12 which points out to disk or hardware related problems. DBAs are currently working on this problem to understand the root cause and provide a fix to the issue as soon as possible.
 - ◆ CMS:
 - ◇ On Wednesday (1st Dec) morning CMS PVSS streaming aborted once again for 30 minutes while executing modifications (adding new table partitions for 2011) on one

of the replicated tables. In fact all changes were already there manually applied by user job. That caused dictionary inconsistency and abort of apply process. Colliding changes have been marked to be skipped and apply process was restarted.

◇ On Thursday (2st Dec) CMS PVSS aborted several times due to missing tablespace on offline database - they were not created together with corresponding tablespaces on online database. All related streams errors were solved manually by creating proper tablespaces on the offline database.

- ◆ LHCb:
 - ◇ nothing

- Site reports:

Site	Status, recent changes, incidents, ...	Planned interventions
ASGC	Nothing to report	None
BNL	Validations for new hardware Working on improvements for Weekly reports	None
CNAF	Nothing to report	None
KIT	Nothing to report	None
IN2P3	Nothing to report	None
NDGF	Nothing to report	None
PIC	Nothing to report	None
RAL	Nothing to report	None
SARA	Nothing to report	Next Tuesday migration to the cluster
TRIUMF	Database was not accessible during last weekend due number of session exceeded because resource_limit parameter was set to FALSE profiles were not working	None

Dates & topics for future meetings

AOB

-- JamieShiers - 23-Nov-2010

This topic: LCG > WLCGTier1ServiceCoordinationMinutes101202
 Topic revision: r12 - 2010-12-02 - 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