

Table of Contents

WLCG Tier1 Service Coordination Minutes - 21 January 2010.....	1
Database services.....	1
Data Management services.....	1
Conditions data access and related services.....	2
Workload Management services.....	3
Security related issues.....	3
Other VO services.....	3
WLCG Baseline Versions.....	3
AOB.....	3

WLCG Tier1 Service Coordination Minutes - 21 January 2010

Database services

- Tier0:
 - ◆ Oracle January Security patches. Wiki page updated by Luca.
 - ◆ Streams bug related to apply failover at destination sites: when apply is moved automatically to a different node, it stops applying LCRs. New alarm implemented in order to detect this. Site needs to be re-synchronize.
 - ◇ When receiving the alarm, contact Tier0 site, but never re-start the apply process.
 - ◇ Oracle support has reproduce the problem. More investigations ongoing.
 - ◇ One occurrence during Xmas at PIC, one more last week at RAL, both affecting LHCb conditions replication.
 - ◆ LFC replication from CERN to the 6 Tier1 sites was affected by a delay of 90 min on Tuesday. Capture process got stuck trying to register one archived log file from the source database on the downstream system. Problem was fixed immediately (capture process re-start) after getting the alarm from the monitoring system.
 - ◇ Roberto asked to shrink the alarm delay.
 - ◆ CMS PVSS test replication disabled while PVSS schemas are reorganized. Move to production afterwards.
- RAL (Carmine): New storage being installed. Tests will be done using the 3d databases but should be transparent. 1 hour downtime needed to relocate the voting disks.
- IN2P3 (Osman): January PSU to be applied in 2 weeks.
- PIC (Gonzalo): Nothing to report.
- NDGF (Jon): Security upgrade to be applied next week.
- Gridka(Andreas): January PSU patch to be scheduled.
- TRIUMF (Andrew): Nothing to report.
- BNL (Carlos): Memory on conditions database upgraded. Security patch deployed in development.
- CNAF (Alessandro): New 2-node console installed. New clusters being installed on new hardware. Migrations have to be scheduled, should use Data Guard.
- SARA (Alexander): Nothing to report.
- ASGC: Nothing to report.

- ATLAS: Nothing to report.
- LHCb: Problem with replication of LHCb integration schema, primary key removed causing apply abortions. Being investigated by Nico and Dawid.

Data Management services

- CASTOR: developers looking after the improvements agreed with ATLAS on the SRM-stager interaction and monitoring. Work is going on, no definite timescale for release. The intention is to provide a test setup to be given to ATLAS to check with an artificial load if problems are fixed.

- dCache
 - ◆ News: recommended version is 1.9.5 (not surprisingly). For the exact version, check the dCache download page and look for "green" versions. SARA upgraded to Chimera, KIT preparing to upgrade.
 - ◆ Issues: a gsidcap bug has been fixed and a patch given to IN2P3 and SARA; LHCb is going to test the patch with a test suite to run stress tests next week in Lyon.

- FTS: two serious bugs in FTS 2.2, one related to corruption of delegate proxies (#60095) and one related to agents crashing (59955) are ready but need to be certified. Normally the time scale to production would be ~2 months, but it can be speeded up. The idea is to deploy a test instance at CERN and have it tested by the experiments in about 2 weeks.
- LFC: a Persistency patch has been prepared to fix the problem in the Persistency LFC component, which caused a user to generate a very high load on an LFC instance for LHCb at CERN. The patch will be tested by LHCb. This component is used only by individual LHCb users, not by LHCb production jobs. Eventually this component will not be used by LHCb any more when the equivalent functionality (secure database access with Grid certificates) is provided by the Persistency server software.
- Site reports
 - ◆ IN2P3: reminded of transfer errors with FNAL due to a few 35 GB files.
 - ◆ RAL: reminded of a hardware problem with the node running the FTS agents
 - ◆ Other: nothing to report
- Experiments
 - ◆ ATLAS: reprocessing foreseen in February at Tier-1 sites. Load on databases is expected to be small because this will use the "dbrelease" mechanism. Frontier testing will continue in parallel.
 - ◆ LHCb: scheduled intervention on 26/1 to upgrade CASTORLHCb at CERN to 2.1.9 and to modify some stager privileges for LHCb administrators.
 - ◆ Other: nothing to report.

Conditions data access and related services

- Frontier
 - ◆ General issues
 - ◇ Maria proposes that Frontier is listed in the baseline services twiki.
 - ◇ John: should we discontinue the weekly ATLAS Frontier meetings? Maria: the T1 Service Coordination Meeting is a general meeting dealing with news and just summaries of activities from ATLAS and CMS. Something internal in ATLAS should be discussed in ATLAS dedicated meetings.
 - ◆ ATLAS
 - ◇ BNL: no specific issue to report. BNL was recently in contact with RAL because of some configuration problem with Frontier servers. RAL updated the frontier server to the latest version: 3.22. Today they are running Hammercloud to check that everything is OK.
 - ◇ Need to understand how conditions are updated and how the information is used at Tier-2 in order to setup minimum expiration time for the cache. This needs to be discussed within ATLAS by John, Douglas and other experts.
 - ◇ Need to understand how to structure the monitoring for data servers. A proposal came from Dave Dykstra. Probably the coordination of this activity should be done not by BNL but by somebody within ATLAS with a better understanding of what needs to be monitored.
 - ◇ CMS developers proposed that the number of Frontier servers in Europe could be reduced. The possible use of Frontier for T1 data access also needs to be understood. Alexei: in ATLAS there are different types of jobs, it is not clear which ones should use Frontier. Dario would like to see performance figures before taking a decision. This discussion should happen in ATLAS. IN2P3: there should be a clear statement about what is required from a site.
 - ◆ CMS
 - ◇ Andrea was contacted by Stephen to discuss the Frontier cache consistency problems observed in T0 processing for CMS. After a change of run, the first jobs processing

events from the new run fail because they use stale conditions from the previous run. The problem is that the same SQL text is used to retrieve queries from all runs, so that wrong results are used until the stale cache is refreshed, which typically takes 15 minutes. CMS will modify the SQL queries by adding the run number explicitly, so that the risk of using a stale cache is avoided. This solution will be tested to understand if it leads to any performance penalty (too frequent reloads of valid data already in the cache).

- Persistency server and COOL
 - ◆ The Persistency server software is used in production in the ATLAS online system since November 2009. It is used for read-only access with a data cache to configure the HLT system. Developments are ongoing to improve the Persistency server for ATLAS online, by adding monitoring features and improving performance. Developments are also ongoing to prepare a Persistency server version which can be used for offline users at CERN and remote sites, in particular providing secure access with Grid certificates (replacing the LFC component presently used by LHCb) and read-write access.
 - ◆ A new version of the Persistency Framework software (Persistency 2.3.6, COOL 2.8.5, POOL 2.9.5), including several fixes and improvements, has been released for the LCGCMT_58 configuration. This new configuration was requested by LHCb and involves the upgrade to the latest ROOT 5.26.00a and other externals, such as frontier_client 2.7.12.

Workload Management services

Security related issues

Other VO services

CMS would like to see a place where the deployment status of SCAS (and in future ARGUS) is published, and where to see which sites have enabled and tested glxexec. To be addressed by finding a volunteer to maintain the information.

WLCG Baseline Versions

Flavia pointed out that we should also indicate if the information providers of the baseline versions of storage systems are compliant with the installed capacity document.

- Table

AOB

CMS is working on risk analysis and mitigation in case of T1 service downtimes: policies for reallocation of workflows and resources. Jamie pointed out that it should be coordinated at WLCG level, follow up offline producing a WLCG document. ATLAS also supports coordination at WLCG level to avoid conflicts between VO actions, has produced an ATLAS site exclusion policy document with metrics and is willing to provide it for extension to all WLCG. Comment from RAL: inform sites about any such action.

-- JamieShiers - 25-Jan-2010

This topic: LCG > WLCGTier1ServiceCoordinationMinutes100121
Topic revision: r3 - 2010-06-11 - PeterJones



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