

Table of Contents

WLCG Tier1 Service Coordination Minutes - 11 February 2010.....	1
Database services.....	1
Data Management services.....	2
FTS 2.2.3 deployment.....	2
CASTOR.....	2
dCache.....	2
LFC, GFAL, lcg_util.....	2
Experiments.....	2
Conditions data access and related services.....	3
.Persistency and COOL.....	3
Workload Management services.....	4
Security related issues.....	4
Other VO services.....	4
WLCG Baseline Versions.....	4
AOB.....	4

WLCG Tier1 Service Coordination Minutes - 11 February 2010

Database services

- ALICE: nothing to report
- ATLAS:
 - ◆ Inconsistency at BNL: transaction is partially applied. Tracing enabled. There are 3 schemas, 9 tables affected. Since 3rd December (hw migration at BNL).
 - ◇ Data inconsistency is tolerable from the experiment point of view as these tables are not used by reprocessing. Other 9 replica sites can be used.
 - ◇ Consistency checks on Streams? Only for updates and deletes, checks on destination, not between source and destination. No errors on the Streams processes. Same number of transactions captured, propagated and applied for all replicas. Problem is inside the transaction.
 - ◇ In the future some consistency checks should be implemented. To be followed up on the ATLAS experiment.
 - ◇ New feature on 11g allows to compare objects between 2 databases. Consistency check is not automatic, still necessary to implement a script to do the checks. No definitive data for the 11g upgrade.
 - ◇ Why so long to make site consistent again? Site can be made consistent immediately. However, we need to gather information to understand the problem and get a fix for it. Otherwise, problem may happen again.
 - ◇ Time given to Oracle must be limited. Investigations are progressing, TRIUMF is collaborating to allow comparisons between replicas.
- CMS: PVSS Streams setup between online and offline progressing. Last schemas to be added next week.
- LHCb: Replication to RAL stopped after network intervention at RAL. Carmine used private email to notify the problem. Encourage people to use mailing lists so problems can be correctly followed up.
- Site reports:
 - ◆ PIC : Jan10 security upgrade applied. Memory upgraded.
 - ◆ CNAF: nothing to report.
 - ◆ SARA: nothing to report.
 - ◆ NDGF: Jan10 security upgrade applied. SGA increased.
 - ◆ GRIDKA: Jan10 security upgrade scheduled for 2 March.
 - ◆ TRIUMF: Jan10 security upgrade foreseen next or 2 weeks.
 - ◆ IN2P3: 3 new schemas added to AMI replication. Jan10 security upgrade pending while waiting for some PSEs.
 - ◆ RAL: nothing to report.
- Client 11g: Andrea opened Oracle SR. Patch available but does not work as expected. It introduces a new crash. Eric suggested to escalate SR. Wait until Monday and then decision will be taken. If no fixed go back to 10g.
- CASTOR CERN:
 - ◆ Security issue related to JAVA on Oracle database found recently. Privileges on affected package have been revoked on all databases at CERN.
- Tier0 - Streams:
 - ◆ Alerts on capture latency changed from 90min to 65min as requested by LHCb.
 - ◆ Alerts on time since the last LCR shipment in the Streams environment changed from 12h to 6h.
 - ◆ ASGC have not announced recent planned intervention to the 3d list. After intervention, propagation not re-enabled. New dba, he needs to review all the procedures.
- ATLAS Oracle Access in Reprocessing

Data Management services

FTS 2.2.3 deployment

FTS 2.2.3 has just been certified, is in "staged rollout" and will be released next Monday (to be confirmed). Right now it can be downloaded from the ETICS repository.

ATLAS and CMS would like all sites to upgrade to 2.2.3 as soon as possible; LHCb agrees with their requests, while ALICE has not yet a defined position on the matter.

All sites have agreed to upgrade by the end of February, but expressed concern about the availability of well defined upgrade procedures from 2.2, from 2.1 and from 2.0. A. Sciabà is going to follow up on this point.

CASTOR

- Versions deployed at CERN
 - ◆ CASTOR: 2.1.9-3 (ALICE, CMS, LHCb), 2.1.9-4 (ATLAS)
 - ◆ SRM: 2.8.5 (ALICE, CMS, LHCb), 2.8.6 (ATLAS), 2.9 being tested (should cure performance problems seen by ATLAS)
- No new releases are planned in the near future. Dirk invited all experiments to contact him to discuss data access via xrootd.
- CERN: discussing the upgrade to 2.1.9-4 for CMS next week for CMS. After that, the public instance and the ALICE and LHCb instances should be upgraded. Negotiations are to begin next week.
- No other CASTOR sites have planned any upgrade.

dCache

- PIC has performed today an upgrade of dCache.
- NDGF plans an upgrade for next week that should largely improve the xrootd access.

LFC, GFAL, lcg_util

- No news.

Experiments

- LHCb: during a recent and intense period of activity, srm-lhcb.cern.ch started giving errors to SAM tests. The cause is not fully understood (it was thought to be due to misconfigured new servers) but now the errors are gone. Documented in GGUS #55303 [↗](#). Following the meeting, this point has been clarified as follows. As mentioned in the GGUS ticket, the post-mortem for the LHCb lhcbdata incident is at <https://twiki.cern.ch/twiki/bin/view/CASTORService/IncidentsLhcbdata05Feb2010> with the follow up actions at the end. The documentation has been improved to avoid the root cause of the problem encountered. The SRM enhancement to reduce the impact of such a misconfiguration will be coming as part of SRM 2.9 (Castor bug <https://savannah.cern.ch/bugs/?45082> [↗](#)). The current plan is to produce a release for this during February and then discuss with the experiments on the deployment schedule via the castor change management process.
- ATLAS: K. Bos had asked how to make sure that user files in CASTORATLAST3 will be guaranteed not to be lost due to disk failures. A possibility could be to do backups to tape (without any constraint on the file size). D. Duellmann commented that CMS will request something similar and different options will be investigated and discussed with the experiments. I. Fisk mentioned that this is not different from protecting user files at Tier-2 sites, which are supposed to provide resilient enough disk storage.

Conditions data access and related services

Persistence and COOL

- Oracle 11g client issues.
 - ◆ Last week ATLAS reported crashes of the Oracle 11.2 client used by CORAL and COOL on AMD Opteron quad-core nodes at NDGF ((ATLAS bug #62194). This is preventing the NDGF Tier-1 from participating in ATLAS reprocessing. Andrea reported the issue as a Service Request and obtained a possible patch from Oracle Support (Oracle bug 8670579). The patch however does not fix the issue and introduces instead new crashes for both Intel (tested at CERN) and AMD Opteron (tested in Ljubljana and Chicago). ATLAS will decide on Monday whether to go back to the 10.2 client, if no fix can be obtained. Oracle Support has already been notified that the fix is required by Monday at the latest. Eric suggested that the issue can be further escalated without going to priority 1.
 - ◆ Andrea confirmed that keeping the 11g client is desirable (because a few issues have been fixed in the 11.2 but not in the 10.2 client, and 10.2 will eventually be discontinued), but going back to the 10g client would still be possible:
 - ◇ The lightweight instant client can be used also with 10g. It is enough to replace libocii.so by libociicus.so in LD_LIBRARY_PATH. A rebuild is not necessary because this library is internally loaded at runtime.
 - ◇ A workaround for the SELinux issue (CORAL bug #45238) exists also for 10g. The Oracle libraries should be modified by 'chcon -t textrel_shlib_t *.so' if they are on local disk (this is not necessary on NFS or AFS). This workaround is in any case necessary also on 11g for the linux64 OCI (Oracle bug 9215184) and OCCI (Oracle bug 8979916) libraries because the bug has only been fixed by Oracle for linux32.
 - ◇ The OCIEnvCreate issue (CORAL bug #31554, aka Oracle bug 5521417) is fixed in 11g, but a backport to 10g was requested and is included in the "10.2.0.4p1" client that was built for the LCG AA projects.
 - ◇ The OCILobRead hang in multi-threaded mode (CORAL bug #47435, aka Oracle bug 6667800) is fixed in 11g, but a workaround that is valid also for 10g is applied in CORAL.
 - ◇ The OCILobRead crash in multi-threaded mode (CORAL bug #49662, aka Oracle bug 8590389) has been solved by a workaround in CORAL that is valid for both 11g and 10g (no fix is planned by Oracle as this is rated as being caused by invalid code).
 - ◇ The OCISmtRelease crash (CORAL bug #61090, aka Oracle bug 9342118), has been addressed by a bug fix in CORAL (committed to CVS but not yet included in any release). In parallel, Oracle will prepare a fix (which is only needed for the 11g client as the issue is absent in the 10g client).
 - ◇ While the 11gR2 32-bit *server* software is not supported by Oracle on 64-bit O/S, metalink reports that "the 32-bit Linux x86 Oracle Database *client* 11gR2 is supported on the Linux x86-64 AMD64/EM64T (64-bit OS)".
- Persistence Framework release
 - ◆ A new release of COOL, CORAL and POOL (LCGCMT_58a) is being prepared for LHCB. This is meant to upgrade to ROOT 5.26.00b to fix an issue observed by the dcache client at remote sites.
- Oracle partitioning for ATLAS/LHCB conditions in COOL and for CMS conditions
 - ◆ Discussions are ongoing between the COOL team and CMS about Oracle partitioning for CMS conditions, to share ideas and experience. New features may need to be added to CORAL to implement schema partitioning in both COOL and the CMS conditions software.
 - ◆ Tests with a prototype implementation of Oracle schema partitioning in COOL (see COOL task #8077) have confirmed that query performance is bad because of a bug in the Oracle 10.2.0.4 server version (Oracle bug 6029469), which is fixed in the 11.2 server (and should be

fixed also in the upcoming 10.2.0.5 server). Luca mentioned that the 10.2.0.5 server software should be available around Easter 2010.

- ATLAS weekly FroNTier meetings [↗](#)
- A new version of the frontier_client library (v. 2.7.13) has been released on the 5th of February 2010 to fix issues seen on the CMS Online cluster where jobs forks multiple processes after most of the conditions are loaded.
This version is distributed by CMS through rpms. However it is not included in the Application Area (AA) distribution since LHCb is not using it and ATLAS is not affected by the issues that this version cures.
- New squid rpms version 2.7.STABLE-7-7 are available. They introduce a new configuration methodology through the customize.sh script which allows to preserve configuration parameters through the upgrades.
- Support for CERNVM though frontier/squid is being discussed. Soon instructions will be distributed for the Tier-3s that would like to use CERNVM pointing to one of the available Tier-2 squids.

Workload Management services

Security related issues

- A discussion with CMS will take place in the near future to see how to tackle ATLAS and CMS common needs in terms of security and service management.

Other VO services

WLCG Baseline Versions

- Release report: deployment status wiki page
- WLCG Baseline versions: table
 - ◆ Versions were changed to match the oldest version supported by EGEE, or a later one if so recommended by the developers or if it contains a significant number of bug fixes. Being discussed how to improve the determination of the baseline versions.

AOB

-- JamieShiers - 11-Feb-2010

This topic: LCG > WLCGTier1ServiceCoordinationMinutes100211

Topic revision: r17 - 2010-06-11 - PeterJones



Copyright &© 2008-2022 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