

# Table of Contents

|                                    |          |
|------------------------------------|----------|
| <b>Index.....</b>                  | <b>1</b> |
| <b>Introduction.....</b>           | <b>2</b> |
| <b>Database Structure.....</b>     | <b>3</b> |
| <b>Plan.....</b>                   | <b>4</b> |
| Schedule.....                      | 4        |
| Actual Achievements.....           | 4        |
| Streaming from CERN to Tier-1..... | 5        |
| <b>Links.....</b>                  | <b>7</b> |
| <b>Contacts.....</b>               | <b>8</b> |

# Index

# Introduction

# Database Structure

The database will contain

- conditions
- detector description

The "conditions" part will contain whatever is in the package Det/XMLConditions, while the "detector description" will contain Det/XMLDDDB.

The database structure will be modelled on the filesystem structure of XmlDDDB (v30r9) and XmlConditions (v2r0).

New conditions will come in two flavors: on-line and off-line. The on-line conditions will be written in the Oracle server in the PIT, the off-line one will be added to the Oracle server at IT. The changes occurring on one server need to be replicated to the other one (from PIT to IT and vice versa), and it means, since we use Oracle streams, that the two flavors have to be in different Oracle schemas (it is handled through a functionality added to Det/DetCond in LHCb v21r3).

Before the introduction of the support for multiple database schemas, it was assumed that the detector description part of the database would go to the same database schema as the off-line conditions. Now it looks more sensible to put it in a different schema to keep the separation we have in the XML files. We can also extract from the current XmlDDDB the part used only in simulation or visualization and put them in their own schema.

# Plan

I presented our plan at the 3D Meeting on the 22nd of June (see the slides [?](#)).

The main points are:

- set up streaming from CERN to Tier-1s
- test replication CERN -> Tier-1
- set up streaming with an LHCb managed RAC (PIT)
- test replication
  - ◆ CERN -> PIT
  - ◆ PIT -> CERN -> Tier-1

## Schedule

The schedule we aim to follow is:

- 31 Jul 2006 - 04 Aug 2006 -
  - ◆ Setup streaming to one Tier-1 (RAL?) (two schemas to be replicated).
  - ◆ Data integrity tests.
  - ◆ Test privileges transmission through streams.
- 07 Aug 2006 - 11 Aug 2006 -
  - ◆ Repeat the tests of the previous week with 2 Tier-1s (RAL and CNAF?)
- 14 Aug 2006 - 18 Aug 2006 -
  - ◆ Pause
- 21 Aug 2006 - 25 Aug 2006 -
  - ◆ Add another Tier-1.
  - ◆ Recovery Tests.
- 28 Aug 2006 - 01 Sep 2006 -
  - ◆ Test access to the DBs from the GRID (hopefully with the CORAL GRID enabled functionalities)
- 04 Sep 2006 - 08 Sep 2006 -
  - ◆ Set up a LHCb managed RAC to simulate the PIT one.
  - ◆ Test cross replication PIT <-> CERN
  - ◆ Test the two step replication PIT -> CERN -> Tier-1.
- 11 Sep 2006 - 15 Sep 2006 -
  - ◆ Pause (LHCb week in Heidelberg)
- 18 Sep 2006 - 22 Sep 2006 -
  - ◆ Include the 3 missing Tier-1s and test from the GRID (**with** the CORAL GRID enabled functionalities).
- Oct 2006 -
  - ◆ Conditions Database used by production jobs.

## Actual Achievements

- 07 Aug 2006
  - ◆ **Test of the software on the Grid:** Nicolas submitted few jobs on the grid to verify the installation of the libraries needed for CondDB (COOL, Oracle, etc.). The result is that the libraries we need are correctly installed.
- 18 Aug 2006
  - ◆ **Replication activated from LHCb integration RAC to RAL and GridKa:** Eva prepared the configuration and tested it.
- 22 Aug 2006

- ◆ **Creation of databases LHCBCOND and DDDDB:** started at 14:00 and ended at ~18:30. LHCBCOND was replicate quickly (tag v2r0 applied), DDDDB replication at RAL ended at 23:50, while at GridKa it failed because of memory problems.
- 24 Aug 2006
  - ◆ **Testing on the GRID:** Prepared a package to run DetDescChecks on the GRID (accessing the master copy) and submitted (Ganga+DIRAC)
  - ◆ **Replication to GridKa:** The server has been recovered and the stream re-connected (Eva ). The copy of DDDDB is OK.
  - ◆ **Replication to IN2P3:** Eva configured the replication to IN2P3 and started to export the schema from the master to the slave (we agreed to not write the the master until the end of the export).
- 25 Aug 2006
  - ◆ **Test on the GRID:** The test failed because of various problems: wrong version of Brunel used (v30r7 instead of v30r6), not supported use-case, Oracle not installed on the CEs. At the end Marco adopted the solution of a custom wrapper script (based on the Ganga one) and a tarball with all the needed files that is downloaded with gridFTP (successfull).
  - ◆ **Replication to IN2P3:** IN2P3 copy ready and stream connected around 16:00 (Eva). The copies of both LHCBCOND and DDDDB are OK.
  - ◆ **Replication of OnLine database:** lhcb\_online\_conddb accounts created on the 3 slaves and streams connected (Eva).
- 28 Aug 2006
  - ◆ **Replication to Tier-1s:** Applied the tag v30r9 to DDDDB. The replication was smooth. *The problem is that the tagging of the whole DB is not atomic and the "global" tag appears before the end of the entire operation.* (we cannot use it to check if the replication is completed) Modified LHCBCOND too (tag v2r0p1)
- 30 Aug 2006
  - ◆ **Privileges replication:** Smooth, only a bit long (few hours).
- 05-28 Sep 2006
  - ◆ **Fake PIT:**
    - ◇ Non trivial!
      - Followed instructions at PSSGroup.ThroughputTestsPreparation, but it was not enough...
      - Few extra roles/functions were needed (Eva provided the scripts for them, see the attachments)
      - Set the SGA memory limit to 700M (import phase was failing)
      - Set the DB in ARCHIVELOG mode (see this article [↗](#))
    - ◇ Finally the 2 schemas have been successfully imported.
    - ◇ Eva enabled the two streams (CERN to PIT, PIT to CERN)
- 29 Sep 2006
  - ◆ **2 step replication:** I managed to insert conditions at ~1Hz without problems for the whole night and the 5 databases are consistent 😊

## Streaming from CERN to Tier-1

The first thing is the definition of the accounts on the Master copy. Currently we have (on the integration rack int4r\_lb):

lhcb\_conddb

schema owner for the off-line conditions (alignments, etc.), 10GB tablespace (LHCB\_CONDDDB\_DATA01)

lhcb\_online\_conddb

schema owner for on-line conditions (those written in the PIT), 10GB tablespace (LHCB\_ONLINE\_CONDDDB\_DATA01)

lhcb\_conddb\_manager

## CondDBCommissioning < LHCb < TWiki

account with read/write permission to the other schemas (he can store/tag data, but not add folders),  
no tablespace

lhcb\_conddb\_reader

account with read only permission on the main schemas, no tablespace

Due to absence for vacation of relevant people, Eva could set up the streaming from the LHCb integration RAC to RAL and GridKa only on Friday 18/08/2006.

# Links

- LCG 3D Project
- COOL [↗](#)



# Contacts

## CERN

Miguel Angio (integration RAC)  
Maria Girone (production RAC)  
Dirk Duellmann (3D)  
Eva Dafonte Perez (Oracle streaming)

## CNAF

Angelo Carbone

## GridKa

*waiting for reply*

## IN2P3

David Bouvet

## NIKHEF

*waiting for reply*

## PIC

Adria Casajus

## RAL

Raja Nandakumar

-- MarcoClemencic - 29 Aug 2006

---

This topic: LHCb > CondDBCommissioning

Topic revision: r16 - 2010-06-11 - PeterJones



Copyright &© 2008-2019 by the contributing authors. All material on this collaboration platform is the property of the contributing authors. Ideas, requests, problems regarding TWiki? Send feedback