EGI Community Forum 2012



Contribution ID: 17

Type: not specified

Handling of network and database instabilities in CORAL

Tuesday, 27 March 2012 14:00 (30 minutes)

Description of the Work

CORAL is a software package that was designed to simplify the development of applications, by screening individual users from the database-specific C++ APIs and SQL flavours.

It provides a C++ abstraction layer that supports data persistency for several backends and deployment models, including local access to SQLite files, direct client access to Oracle and MySQL servers, and read-only access to Oracle through the FroNTier/Squid and CoralServer/CoralServerProxy server/cache systems.

During 2010, several problems were reported by the LHC experiments using CORAL, involving application hangs or crashes after the network or the database servers became temporarily unavailable. CORAL already provided some level of handling of these instabilities, which are due to external causes and cannot be avoided, but this proved to be insufficient in some cases and to be itself the cause of other problems, such as the hangs mentioned before, in other cases. As a consequence, a major redesign of the CORAL plugins was implemented, with the aim of making the software more robust against these network glitches.

Conclusions

The CORAL software is widely used for accessing from C++ and python applications the data stored by the LHC experiments using a variety of relational database technologies (including Oracle, MySQL and SQLite). The new feature, implemented to cope with the network glitches, allow an automatic reconnection to the database making CORAL more robust and reliable in a transparent and safe way for the users.

Impact

The new implementation ensures that CORAL automatically reconnects to the database in a transparent way whenever possible and gently terminates the application when this is not possible. Internally, it takes care of resetting all relevant parameters of the underlying backend technology (such as OCI, the Oracle Call Interface).

Overview (For the conference guide)

The Large Hadron Collider (LHC), the world's largest and highest-energy particle accelerator, started its operations in September 2008 at CERN, Switzerland. Huge amounts of data are generated by the four experiments installed at different collision points along the LHC ring. The largest data volumes come from the 'event data' that record the signals left in the detectors by the particles generated in the LHC beam collisions and are generally stored on files. Relational database systems are commonly used instead to store the 'conditions data' that record the geometry, configuration and other working parameters of the detectors at the time the event Primary author: Dr TRENTADUE, Raffaello (CERN)
Co-authors: Mr LOTH, Alexander (CERN); Dr VALASSI, Andrea (CERN)
Presenters: Mr LOTH, Alexander (CERN); Dr TRENTADUE, Raffaello (CERN)
Session Classification: Community-tailored Services

Track Classification: Operational services and infrastructure