



Contribution ID: 16

Type: **not specified**

Monitoring in CORAL

Tuesday, 27 March 2012 17:05 (25 minutes)

Description of the Work

CORAL 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.

Given the huge amount of operations executed by several CORAL clients at the same time on several database servers, it was crucial to develop a monitoring system with two main goals: first, to allow individual CORAL users to study and optimize the performance of the relational operations executed by their applications; second, to check whether the whole system is properly working and well configured. Client-level monitoring functionalities already existed in CORAL, but they have recently been reviewed and significantly improved, especially for the Oracle and Frontier plugins, and the same functionality are also being integrated into the CORAL server component (itself a CORAL-based application) and its client plugin. Work is in progress also on the monitoring of the CoralServerProxy components and on the aggregation of the monitoring information these proxies provide when deployed in a hierarchical structure, such as that used by the ATLAS High Level Trigger system. This presentation will report on the status of this work at the time of the 2012 EGI Community Forum, covering the design and implementation of these new features and the results from the first experience with their use.

Conclusions

The CORAL software is widely used by the LHC experiments for storing and accessing conditions data using relational database technologies. It provides a C++ abstraction layer that supports data persistency for several backends and deployment models. CORAL is appreciated for its reliability and performance.

Impact

CORAL had proven that it is a fundamental framework for database abstraction. It is used by various experiments and software components. It is appreciated for its reliability and performance.

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

data were collected. The CORAL software is widely used by the LHC experiments for storing and accessing conditions data using relational database technologies.

Primary author: Mr LOTH, Alexander (CERN)

Co-authors: Dr VALASSI, Andrea (CERN); Dr TRENTADUE, Raffaello (CERN)

Presenter: Mr LOTH, Alexander (CERN)

Session Classification: Quality Assurance - Infrastructure

Track Classification: Operational services and infrastructure