



Contribution ID: 75

Type: **Poster presentation**

Optimising query execution time in LHCb Bookkeeping System using partition pruning and partition wise joins

Monday, 14 October 2013 15:00 (45 minutes)

The LHCb experiment produces a huge amount of data which has associated metadata such as run number, data taking condition (detector status when the data was taken), simulation condition, etc. The data are stored in files, replicated on the Computing Grid around the world. The LHCb Bookkeeping System provides methods for retrieving datasets based on their metadata. The metadata is stored in a hybrid database model, which is a mixture of Relational and Hierarchical database models and is based on the Oracle Relational Database Management System (RDBMS). The database access has to be reliable and fast. In order to achieve a high timing performance, the tables are partitioned and the queries are executed in parallel. When we store large amounts of data the partition pruning is essential for database performance, because it reduces the amount of data retrieved from the disk and optimises the resource utilisation. This research presented here is focusing on the extended composite partitioning strategy such as range-hash partition, partition pruning and usage of the partition wise joins. The system has to serve thousands of queries per minute, the performance and capability of the system is measured when the above performance optimisation techniques are used.

Summary

Primary author: MATHE, Zoltan (CERN)

Co-author: CHARPENTIER, Philippe (CERN)

Presenter: MATHE, Zoltan (CERN)

Session Classification: Poster presentations

Track Classification: Data Stores, Data Bases, and Storage Systems