



Contribution ID: 22

Type: **Parallel Talk**

Petaminer: Efficient Navigation to Petascale Data Using Event-Level Metadata

Tuesday 4 November 2008 14:25 (25 minutes)

HEP experiments at the LHC store petabytes of data in ROOT files described with TAG metadata. The LHC experiments have challenging goals for efficient access to this data. Physicists need to be able to compose a metadata query and rapidly retrieve the set of matching events. Such skimming operations will be the first step in the analysis of LHC data, and improved efficiency will facilitate the discovery process by permitting rapid iterations of data evaluation and retrieval. Furthermore, efficient selection of LHC data helps enable the tiered data distribution system adopted by LHC experiments, in which massive raw data resides at a few central sites, while higher quality, smaller scale skimmed data is replicated at many lower tier sites with more modest computational resources.

To address this problem, we are developing a custom MySQL storage engine to enable the MySQL query processor to directly access TAG data stored in ROOT TTrees. As ROOT TTrees are column-oriented, reading them directly will provide improved performance over traditional row-oriented TAG databases. In addition, to the efficient SQL query interface to the data stored in ROOT TTrees, the Petaminer technology will enable rich MySQL index-building capabilities to add indices to the data in ROOT TTrees, providing further optimization to TAG query performance.

Column-oriented databases are an emerging technique for achieving higher performance than traditional row-oriented databases, especially in large scale data-mining scenarios.

We will present first results of our feasibility studies of creating a column-oriented MySQL storage engine that enables MySQL to access TAG metadata directly from ROOT files.

Summary

We report on the development of a custom MySQL storage engine to directly access TAG metadata stored in ROOT TTrees. The ability to directly read column-oriented ROOT data via a MySQL front end, and to create MySQL indexes to optimize data access, offers the potential to optimize the speed of TAG metadata retrieval in LHC data analysis.

Authors: VANIACHINE, Alexandre (Argonne National Laboratory); MALON, David (Argonne National Laboratory); CRANSHAW, Jack (Argonne National Laboratory); HAMILL, Paul (Tech-X Corporation)

Presenter: VANIACHINE, Alexandre (Argonne National Laboratory)

Session Classification: Data Analysis - Algorithms and Tools

Track Classification: 2. Data Analysis