



Contribution ID: 432

Type: poster

Distributed database services in PHENIX - what it takes to support a Petabyte experiment.

Monday, 3 September 2007 08:00 (20 minutes)

After seven years of running and collecting 2 Petabytes of physics data, PHENIX experiment at the Relativistic Heavy Ion Collider (RHIC) has gained a lot of experience with database management systems (DBMS).

Serving all of the experiment's operations - data taking, production and analysis - databases provide 24/7 access to calibrations and book-keeping information for hundreds of users at several computing centers worldwide and face the following challenges:

- Simultaneous data taking, production and analysis result in hundreds of concurrent database connections and heavy server I/O load.
- Online data production at remote sites requires a high degree of Master-Slave server synchronization.
- Database size (presently 100GB with half of data added in the last few months) raises scalability concerns.
- Long life of modern HENP experiments and fast development of database technologies make prediction of the best DBMS provider 5-10 years down the road difficult and require investments in design and support of good APIs.

In this talk PHENIX solutions to the above problems will be presented and the trade-offs discussed.

Submitted on behalf of Collaboration (ex, BaBar, ATLAS)

PHENIX

Primary author: SOURIKOVA, Irina (BROOKHAVEN NATIONAL LABORATORY)

Presenter: SOURIKOVA, Irina (BROOKHAVEN NATIONAL LABORATORY)

Session Classification: Poster 1

Track Classification: Online Computing