



Contribution ID: 284

Type: oral presentation

Efficient Access to Remote Data in High Energy Physics

Monday 3 September 2007 16:30 (20 minutes)

Particle accelerators produce huge amounts of information in every experiment and such quantity cannot be stored easily in a personal computer. For that reason, most of the analysis is done using remote storage servers (this will be particularly true when the Large Hadron Collider starts its operation in 2007). Seeing how the bandwidth has increased in the last few years, the biggest problem of this approach at the moment is latency, which hurts considerably the performance of the analysis process.

Fortunately, particle events are independent of each other, which allows us to transfer the information that must be processed in the future while analyzing the data at hand. The independent nature also allows us to transfer many events instead of the single one needed at a given time. Such ideas are implemented in the data analysis framework ROOT, and its file servers (rootd, xrootd and a http plugin). Among the techniques used, we have pre-reads, pre-fetching, parallel streams and atomic readings for multiple requests. All these strategies present an enormous advantage and will facilitate processing remote files at almost the same speed as local ones, as long as the bandwidth does not present any limitations.

Primary authors: RADEMAKERS, Fons (CERN); FRANCO, Leandro (CERN); BRUN, Rene (CERN)

Presenter: FRANCO, Leandro (CERN)

Session Classification: Distributed data analysis and information management

Track Classification: Distributed data analysis and information management