



Contribution ID: 307

Type: oral presentation

Latest Developments in the PROOF System

Thursday, September 6, 2007 3:20 PM (20 minutes)

The goal of PROOF (Parallel ROOt Facility) is to enable interactive analysis of large data sets in parallel on a distributed cluster or multi-core machine. PROOF represents a high-performance alternative to a traditional batch-oriented computing system.

The ALICE collaboration is planning to use PROOF at the CERN Analysis Facility (CAF) and has been stress testing the system since mid 2006 on a 40 machine pilot cluster. The ALICE CAF is expected to grow to around 500 machines.

The testing by ALICE has allowed us to identify missing functionality and to improve the system in many ways. Areas of significant development include: a dataset manager to optimally distribute data on the cluster; facilities to upload and manage the experiment software; a new “packetizer” which significantly reduces the end-of-query tails; a worker-level priority-based scheduling system to control the fraction of resources assigned to a group of users; improved error handling and user feedback mechanism; and much more.

The CMS collaboration is also actively investigating PROOF as Tier-2 analysis facility.

Current activities focus on the development of a central scheduling system that uses the OLBD/XROOTD control network as information routing system. This scheduler aims to improve resource sharing in a multi-user environment, taking per-query decisions based on the status of the farm, the query requirements and the history and priorities of the user.

In this paper we will describe in detail the recent developments, the status of the current activities, and outline the future plans to bring PROOF in production for LHC analysis.

Primary author: Dr RADEMAKERS, Fons (CERN)

Co-authors: Mr BELLENOT, Bertrand (CERN); Dr GANIS, Gerardo (CERN); Mr IWASZKIEWICZ, Jan (CERN); Dr BALLINTIJN, Maarten (MIT); Dr BRUN, Rene (CERN)

Presenter: Dr RADEMAKERS, Fons (CERN)

Session Classification: Distributed data analysis and information management

Track Classification: Distributed data analysis and information management