



Contribution ID: 287

Type: **oral presentation**

Super scaling PROOF to very large clusters

Thursday, September 30, 2004 3:00 PM (20 minutes)

The Parallel ROOT Facility, PROOF, enables a physicist to analyze and understand very large data sets on an interactive time scale. It makes use of the inherent parallelism in event data and implements an architecture that optimizes I/O and CPU utilization in heterogeneous clusters with distributed storage. Scaling to many hundreds of servers is essential to process tens or hundreds of gigabytes of data interactively. This is supported by the industry trend to pack more CPU's into single systems and to create bigger clusters by increasing the number of systems per rack. We will describe the latest developments in PROOF and the development of a standardized benchmark for PROOF clusters. The benchmark is self contained and measures the network, the I/O and the processing characteristics of a cluster. We will present the comprehensive results of the benchmark for several clusters, demonstrating the performance and scalability of PROOF on very large clusters.

Primary authors: RADEMAKERS, F. (CERN); ROLAND, G. (MIT); GULBRANDSEN, K. (MIT); BALLINTIJN, M. (MIT); CANAL, P. (Fermilab); BRUN, R. (CERN)

Presenter: BALLINTIJN, M. (MIT)

Session Classification: Distributed Computing Systems and Experiences

Track Classification: Track 5 - Distributed Computing Systems and Experiences