



Contribution ID: 328

Type: oral presentation

The Next Generation Root File Server

Monday, 27 September 2004 16:30 (20 minutes)

As the BaBar experiment shifted its computing model to a ROOT-based framework, we undertook the development of a high-performance file server as the basis for a fault-tolerant storage environment whose ultimate goal was to minimize job failures due to server failures. Capitalizing on our five years of experience with extending Objectivity's Advanced Multithreaded Server (AMS), elements were added to remove as many obstacles to server performance and fault-tolerance as possible. The final outcome was xrootd, upwardly and downwardly compatible with the current file server, rootd. This paper describes the essential protocol elements that make high performance and fault-tolerance possible; including asynchronous parallel requests, stream multiplexing, data pre-fetch, automatic data segmenting, and the framework for a structured peer-to-peer storage model that allows massive server scaling and client recovery from multiple failures. The internal architecture of the server is also described to explain how high performance was maintained and full compatibility was achieved. Now in production at Stanford Linear Accelerator Center, Rutherford Appleton Laboratory (RAL), INFN, and IN2P3; xrootd has shown that our design provides what we set out to achieve. The xrootd server is now part of the standard ROOT distribution so that other experiments can benefit from this data serving model within a standard HEP event analysis framework.

Primary authors: DORIGO, A. (INFN, Padova); HANUSHEVSKY, A. (SLAC); FURANO, F. (INFN, Padova)

Presenter: HANUSHEVSKY, A. (SLAC)

Session Classification: Distributed Computing Services

Track Classification: Track 4 - Distributed Computing Services