

Latencies and data access. Boosting the performance of distributed applications.

Wednesday 15 February 2006 16:20 (20 minutes)

The latencies induced by network communication often play a big role in reducing the performances of systems which access big amounts of data in a distributed environment. The problem is present in Local Area Networks, but in Wide Area Networks is much more evident. It is generally perceived as a critical problem which makes very difficult to get access to remote data. However, a more detailed analysis on the access pattern of the involved applications can be used to understand the characteristics of the stream of the data requests, and develop techniques to optimize it. This work started from the analysis of the access patterns of the BaBar experiment's physics analysis data, but the methods and the results can be applied in other computing environments as well. We show how the exploit of caching and asynchronous prefetching techniques is able to enhance the performances of such kind of applications in Local Area Networks, and is able to lower the total latencies for Wide Area Networks data access of an order of magnitude. Moreover, the ability to process file open requests in parallel can be a very interesting performance enhancement for applications which need access to many files at once. Such general techniques have been implemented in the client side of the xrootd data access system, which, for its performances and its fault tolerant architecture, showed itself as an ideal testbed for such kind of enhancements.

Primary author: Mr FURANO, Fabrizio (INFN sez. di Padova)

Presenter: Mr FURANO, Fabrizio (INFN sez. di Padova)

Session Classification: Distributed Data Analysis

Track Classification: Distributed Data Analysis