



Contribution ID: 310

Type: oral presentation

Real-time Data Access Monitoring in Distributed, Multi-Petabyte Systems

Thursday, September 6, 2007 2:40 PM (20 minutes)

Petascale systems are in existence today and will become widespread in the next few years. Such systems are inevitably very complex, highly distributed and heterogeneous. Monitoring a petascale system in real time and understanding its status at any given moment without impacting its performance is a highly intricate task. Common approaches and off the shelf tools are either unusable, do not scale, or severely impact the performance of the servers that are monitored. This talk will describe an unobtrusive monitoring software developed at Stanford Linear Accelerator Center (SLAC) and currently deployed by the BaBar Experiment that uses the xrootd file access system to access its highly distributed petascale production data set. The system facilitates central monitoring of all BaBar Tier A centers at SLAC. The talk will describe the employed solutions, the lessons learned, and the issues still to be addressed, and discuss the advantages of such a system in predicting the storage needs and understanding data access patterns. It will further explain how the system can be deployed in other High Energy Physics centers where the data servers may be shared by many experiments and run under a different file access system.

Author: Dr AZEMOON, Tofigh (Stanford Linear Accelerator Center)

Co-authors: Mr HANUSHEVSKY, Andrew (Stanford Linear Accelerator Center); Mr BECLA, Jacek (Stanford Linear Accelerator Center); Mr MASSIMILIANO, Turri (Stanford Linear Accelerator Center)

Presenter: Dr AZEMOON, Tofigh (Stanford Linear Accelerator Center)

Session Classification: Distributed data analysis and information management

Track Classification: Distributed data analysis and information management