Contribution ID: **431** Type: **oral**

Fair-share scheduling algorithm for a tertiary storage system

Monday, 23 March 2009 17:50 (20 minutes)

Any experiment facing Peta bytes scale problems is in need for a highly scalable mass storage system (MSS) to keep a permanent copy of their valuable data. But beyond the permanent storage aspects, the sheer amount of data makes complete dataset availability onto "live storage"(centralized or aggregated space such as the one provided by Scala/Xrootd) cost prohibitive implying that a dynamic population from MSS to faster storage is needed. One of the most difficult aspects of dealing with MSS is the robotic tape component and its intrinsically long access times (latencies) that can dramatically affect the overall performance of any data access systems having MSS as their primary data storage.

To speed the retrieval of such data, one could "organize" the requests according to criterion with an aim to deliver maximal data throughput. However, such approaches are often orthogonal to the fairness and a tradeoff between quality of service (responsiveness) and throughput is necessary for an optimal and practical implementation of a truly faire-share oriented file restore policy. Starting from explaining the key criterion used to build such policy, we will present an evaluation and comparisons of three different algorithms, offering fairshare file restoration from MSS and discuss their respective merits. We will further quantify their use impact on a typical file restoration for the RHIC/STAR experimental setup and this, within a development, analysis and production environment relying on a shared MSS service.

Presentation type (oral | poster)

oral

Authors: Dr LAURET, Jerome (BROOKHAVEN NATIONAL LABORATORY); Mr JAKL, Pavel (Nuclear Physics Inst., Academy of Sciences, Praha)

Presenter: Mr JAKL, Pavel (Nuclear Physics Inst., Academy of Sciences, Praha)

Session Classification: Hardware and Computing Fabrics

Track Classification: Hardware and Computing Fabrics