20th International Conference on Computing in High Energy and Nuclear Physics (CHEP2013)

Contribution ID: **477** Type: **Poster presentation**

System level traffic shaping in diskservers with heterogeneous protocols

Monday 14 October 2013 15:00 (45 minutes)

Disk access and tape migrations compete for network bandwidth in CASTOR's disk servers, over various protocols: RFIO, Xroot, root and GridFTP. As there are a limited number of tape drives, it is important be keep them busy all the time, at their nominal speed. With potentially 100s of user read streams per server, the bandwidth for the tape migrations has to be guaranteed to a controlled level, and not the default fair share the system gives by default. Xroot provides a prioritization mechanism , but using it implies moving exclusively to the Xroot protocol, which is not possible in short to mid-term time frame, as users are equally using all. The greatest commonality of all those protocols is not more that usage of TCP/IP. We investigated the Linux kernel traffic shaper to control TCP/IP bandwidth. The performance and limitations of the traffic shaper have been understood in test environment, and satisfactory working point has been found for production. Notably, TCP offload engines' negative impact on traffic shaping, and the limitations of the length of the traffic shaping rules were discovered and measured. A suitable working point has been found and the traffic shaping is now successfully deployed in the CASTOR production systems at CERN. This system level approach could be transposed easily to other environments.

Author: CANO, Eric (CERN)

Co-author: KRUSE, Daniele Francesco (CERN)

Presenter: KRUSE, Daniele Francesco (CERN)

Session Classification: Poster presentations

Track Classification: Data Stores, Data Bases, and Storage Systems