

## xrootd Server Clustering

*Wednesday, 15 February 2006 09:00 (20 minutes)*

Server clustering is an effective method in increasing the pool of resources available to applications. Many clustering mechanisms exist; each with its own strengths as well as weaknesses. This paper describes the mechanism used by xrootd to provide a uniform data access space consisting of an unbounded number of independent distributed servers. We show how the mechanism is especially effective in reducing data request routing latency as well as eliminating most of the cluster definition details to allow super large clusters to be constructed on-the-fly with a minimum amount of administration.

**Primary author:** HANUSHEVSKY, Andrew (Stanford Linear Accelerator Center)

**Co-authors:** WEEKS, Bill (Stanford Linear Accelerator Center); LAURET, Jerome (BNL); KROEGER, Wilko (Stanford Linear Accelerator Center)

**Presenter:** HANUSHEVSKY, Andrew (Stanford Linear Accelerator Center)

**Session Classification:** Poster

**Track Classification:** Distributed Event production and processing