

Performance of Hadoop file system on a HEPiX storage workgroup testbed

Monday, October 26, 2009 2:45 PM (30 minutes)

This work is continuation of storage solution testing performed by HEPiX storage workgroup on it's testbed at FZK. Hadoop, an Apache project, offers a cluster file system called HDFS inspired by Google File System and designed to run on commodity hardware. It has gained some popularity in OSG, where is has become a supported storage solution, and is currently in production at a few T2 sites.

In this series of tests we used HEPiX testbed worker nodes' hard drives as a underlying storage, without using external file servers or storage arrays. We used a standard HEPiX storage application suite to evaluate performance of this solution. We present obtained results in this paper.

Primary author: TRUNOV, Artem (Karlsruhe Institute of Technology)

Presenter: TRUNOV, Artem (Karlsruhe Institute of Technology)

Session Classification: Storage I

Track Classification: Storage