



Contribution ID: 402

Type: poster

Exploring Capable Systems for Fast Data Distribution over End-to-End 10Gb/s Paths

Wednesday 5 September 2007 08:00 (20 minutes)

A primary goal of the NSF-funded UltraLight Project is to expand existing data-intensive grid computing infrastructures to the next level by enabling a managed network that provides dynamically constructed end-to-end paths (optically or virtually, in whole or in part). Network bandwidth used to be the primary limiting factor, but with the recent advent of 10Gb/s network paths end-to-end, the end system has become the bottleneck for fast data distribution. As an additional goal of UltraLight we have been exploring tradeoff issues with relatively inexpensive solutions for capable end systems. The candidate system should be capable of driving 10 Gb/s WAN links in order to provide fast data cache capabilities for LHC(Large Hadron Collider) computing model.

In this paper, we perform various synthetic and application benchmarks on each architectural component in the data path of a disk-to-disk transfer. First, we find that disk subsystems are usually the main limiting factor for fast data distribution. Therefore, the candidate platforms are required to provide wide I/O paths to accommodate high-performance I/O devices. Second, we investigate a broad range of tunable parameters in the operating system and their impact on the throughput of disk-to-disk transfer. We observe improvement exceeding a factor of 2 in disk I/O throughput which directly impacts the throughput of large disk-to-disk transfers. We also find that disk and memory access patterns together with the size of various buffers play critical roles in maximizing the throughput of disk-to-disk transfers.

Submitted on behalf of Collaboration (ex, BaBar, ATLAS)

UltraLight

Author: Mr PARK, Kyu (Department of Electrical and Computer Engineering, University of Florida)

Co-authors: Dr GEORGE, Alan (Department of Electrical and Computer Engineering, University of Florida); Dr BOURILKOV, Dimitri (Department of Physics, University of Florida); Dr AVERY, Paul (Department of Physics, University of Florida); Dr CAVANAUGH, Richard (Department of Physics, University of Florida); Dr MCKEE, Shawn (Department of Physics, University of Michigan)

Presenter: Mr PARK, Kyu (Department of Electrical and Computer Engineering, University of Florida)

Session Classification: Poster 2

Track Classification: Distributed data analysis and information management