QuickSAN: A Storage Area Network for Fast, Distributed, Solid State Disks

Adrian M. Caulfield
acaulfie@cs.ucsd.edu

Steven Swanson
swanson@cs.ucsd.edu

Computer Science and Engineering Department
University of California, San Diego

Abstract

Solid State Disks (SSDs) based on flash and other non-volatile memory technologies reduce storage latencies from 10s of milliseconds to 10s or 100s of microseconds, transforming previously inconsequential storage overheads into performance bottlenecks. This problem is especially acute in storage area network (SAN) environments where complex hardware and software layers (distributed file systems, block severs, network stacks, etc.) lie between applications and remote data. These layers can add hundreds of microseconds to requests, obscuring the performance of both flash memory and faster, emerging non-volatile memory technologies.

We describe QuickSAN, a SAN prototype that eliminates most software overheads and significantly reduces hardware overheads in SANs. QuickSAN integrates a network adapter directly into SSDs, so the SSDs can communicate directly with one another to service storage accesses as quickly as possible. QuickSAN can also give applications direct access to both local and remote data without operating system intervention, further reducing software costs. Our evaluation of QuickSAN demonstrates remote access latencies of 20 μs for 4 KB requests, bandwidth improvements of as much as 163$\times$ for small accesses compared with an equivalent iSCSI implementation, and 2.3-3.0$\times$ application level speedup for distributed sorting. We also show that QuickSAN improves energy efficiency by up to 96% and that QuickSAN's networking connectivity allows for improved cluster-level energy efficiency under varying load.

A copy of this technical report can be obtained by sending a request to swanson@cs.ucsd.edu