Evaluating Hardware Compression Offload in a Lustre File System

Author: Mariana Hernandez

LA-UR-20-25808

Advancements in solid-state storage media, network, and PCIe performance, has allowed HPC storage system designers to begin practically designing an efficient bandwidth tier. Although SSDs have become affordable, this increased speed often means a reduction in capacity per device. When designing a storage system with an optimal performance level and limited capacity, compression becomes a vital component of the design. In measuring existing compression algorithms in ZFS, we find that these compression techniques drastically slow performance while providing high levels of compression or minimally impact performance while providing little compression benefit. To attain both fast and efficient compression of scientific datasets, compression can be offloaded to specialty hardware, such as Eideticom's NoLoad FPGA. We benchmarked two types of compression, GZIP and LZ4, on an in-memory ZFS based Lustre file system. Our results show that offloading compression to the FPGA proved to both minimize impact on performance and provide high compression efficiency. This talk will present challenges in storage node efficiency and provide insight into the potential efficiency gains using computational storage, while also describing the difficulties around configuring such a system.