

Analyzing Server-side Scalability of Image Filesystems & Attachment Technologies
Authors: Timothy Bargo, Aedan Wells, Michelle Yoon, J. Lowell Wofford, Travis Cotton
LA-UR-21-27857

High Performance Computing often deals in managing hundreds to thousands of compute nodes to solve large, complex problems. As we push the boundaries of compute, we continue to optimize the performance of all components of the cluster. A common method of compute cluster deployment is to utilize a master server to provide operating system images to the compute nodes. On large compute clusters this deployment method can lead to large workloads on the master server. In our work we compared different image filesystem types and attachment technologies to determine the most performant and scalable method to deploy a compute cluster. Our research demonstrated that the SquashFS filesystem and the Ceph RBD image attachment technology produce lower CPU and network loads in comparison to other combinations of SquashFS, Ext4, ISOFS, and XFS filesystems with Ceph RBD, NFS, NFS loopback, and iSCSI attachment technologies. Our results provide guidance in selecting the most scalable combination of technologies to deploy compute clusters.