

Parallel NFS over RDMA for HPC

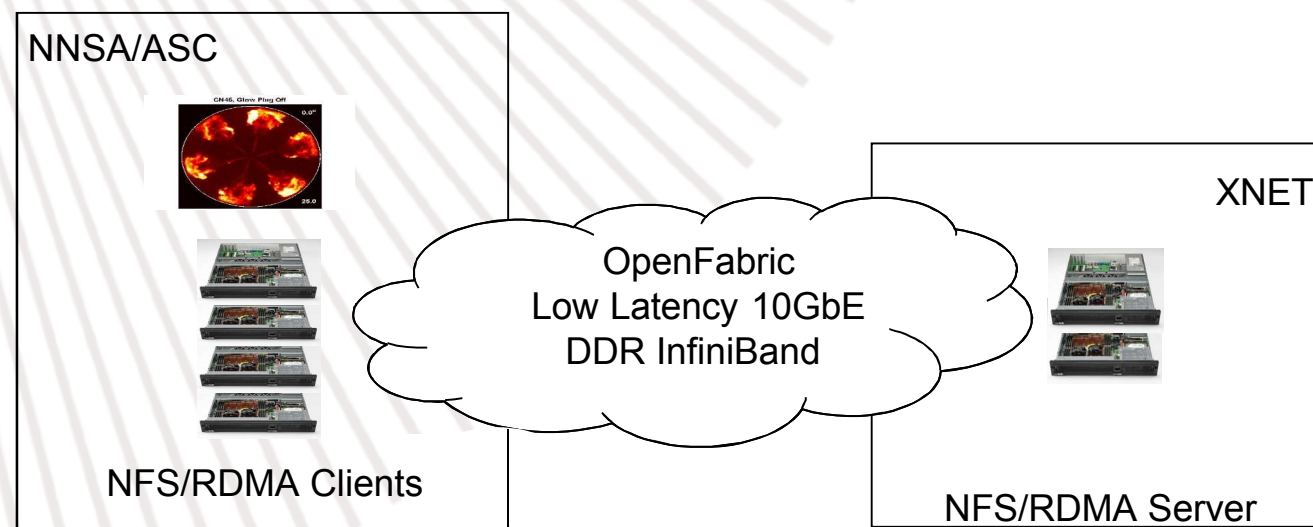
SAND2006-6941C



The latest interconnect technologies have enabled commodity High Performance Computing. However, commodity I/O storage system technologies lagged in similar advancement. This performance gap is of major concern because ASC applications are becoming more data intensive.

The ASC Scalable Computing team at Sandia National Laboratories is conducting a study to address this issue. Through research and industry partnerships, this study explores Parallel and RDMA technologies for I/O scalability. This effort also promotes a standardized common filesystem interface such as NFS to facilitate data sharing.

In this poster we demonstrate an early implementation of NFS over the emerging InfiniBand and iWARP technologies. We will be encoding and visualizing the isosurface of a medical data set stored on an RDMA NFS server, located at the XNET booth across the SCinet OpenFabric Low Latency 10GbE and InfiniBand networks. Parallel NFS over RDMA is still work-in-progress.



RDMA NFS achieved 928 MB/s throughput:

- o > 3 x NFS performance
- o 1/3 CPU overhead
- o Much improved scalability

Better application performance !!!



UNIVERSITY OF MICHIGAN

