Windows NVMe over Fabrics Performance

Line Rate Performance with Chelsio T6 100GbE iWARP RDMA

Executive Summary

NVMe over Fabrics specification extends the benefits of NVMe to large fabrics, beyond the reach and scalability of PCIe. NVMe enables deployments with hundreds or thousands of SSDs using a network interconnect, such as RDMA over Ethernet. T6 iWARP RDMA provides a low latency, high throughput, plug-and-play Ethernet solution for connecting high performance NVMe SSDs over a scalable, congestion controlled and traffic managed fabric, with no special configuration needed.

Chelsio iWARP enabled adapters power a range of Windows Server capabilities, including Storage Spaces Direct software-defined storage, Storage Replica for disaster recovery, SMB Direct for high performance file access, Network Direct for Windows HPC deployments, hardware offloaded iSCSI initiator for SAN applications. In addition, Chelsio's concurrent certification for Windows 10 Enterprise enables Client RDMA functions and allows the iWARP wire protocol be used on both ends of the wire. Chelsio's recent NVMe-oF initiator driver presents significant performance benefits in Windows Server environments. With 98 Gbps line-rate throughput, Chelsio's solution proves to be the best-in-breed in providing the next generation, scalable storage network over standard and cost effective Ethernet infrastructure with an efficient processing path.

Test Results

The following graph presents READ, WRITE IOPS and throughput results of Chelsio NVMe-oF initiator. The results are collected using the **diskspd** v2.0.17 tool with I/O size varying from 512 bytes to 512 Kbytes with an access pattern of random READs and WRITEs.

Figure 1 – IOPS and Throughput vs. I/O size

As evident from the graphs above, T6 solution delivers line-rate throughput of 98 Gbps for READ and WRITE. READ IOPS exceeds 1 Million and WRITE IOPS reach 820K. *Please note these are preliminary numbers and further IOPS tuning is in progress.*

Test Setup

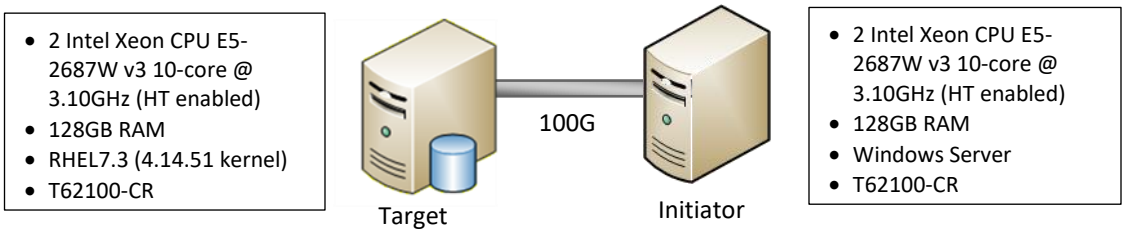

Figure 2 – Test setup

The setup consists of an NVMe target machine connected to an initiator machine with a single 100G port. MTU of 9000B is used. Latest Chelsio Unified Wire drivers for Linux and Windows are installed on target and initiator machines respectively.

Storage configuration

The target is configured with 2 Null Block devices, each of 1GB size.

Commands Used

C:\Users\Administrator> diskspd.exe -b<IO Size> -w<read/write> -r -o8 -W5 -C5 -d30 -t16 -Sh <disk #s>

Conclusion

This paper showcases the remote storage access performance capabilities of Chelsio T6 NVMe-oF over 100GbE iWARP fabric solution in Windows setup. Using iWARP RDMA enables the NVMe storage devices to be shared, pooled and managed more effectively across a low latency, high performance network. The results show that Chelsio's iWARP RDMA delivers line-rate throughput performance for both READ and WRITE operations. With concurrent support for iSCSI, iSER, FCoE and NVMe-oF initiators, Chelsio Converged Network Adapters are the best in class and well suited for Windows environments.

Related Links

100G NVMe over Fabrics JBOF
NVMe over Fabrics Performance for AMD EPYC
High Performance NVMe-oF with T6 100G iWARP RDMA
NVMe Over Fabrics Performance for Qualcomm ARM
NVMe over Fabrics iWARP Performance