

Windows iSER Performance at 100Gbps

RDMA Performance and Efficiency without Fabric Overhaul

The iSCSI Extensions for RDMA (iSER) protocol is a translation layer for operating iSCSI over RDMA transports, such as iWARP/Ethernet or InfiniBand. Thanks to its hardware offloaded TCP/IP foundation, iWARP provides the high performance, low latency and efficiency benefits of RDMA and runs over standard Ethernet gear, without the need for special configuration or additional management costs.

This paper presents 100GbE iSER performance over iWARP RDMA using Chelsio T62100-CR Unified Wire adapter in Windows Server 2019 setup. The results demonstrate that Chelsio T6 adapter achieves line rate performance of 98 Gbps and 500K IOPS. Chelsio’s iSER solution provides high performance and efficiency, along with routability to scale to large datacenters, clouds and long distances.

Test Results

The following graph plots the iSER READ and WRITE throughput and IOPS numbers using the **lometer** tool v1.1.0. The I/O size used varies from 512B to 512KB with an access pattern of random READs and WRITES.




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

The results above reveal that iSER over iWARP RDMA delivers high and consistent performance throughout, reaching line rate of 98 Gbps for both READ and WRITE. WRITE IOPS reaches a maximum of 500K. iWARP RDMA uses a hardware TCP/IP stack that runs in the adapter, completely bypassing the host software stack, thus eliminating any inefficiencies due to software processing. It also provides benefits of CPU bypass and zero copy, resulting in significant CPU savings.

Test Configuration

The setup consists of a Windows iSER Initiator machine connected to 2 LIO iSER target machines through a 100GbE switch using single port on each system. MTU of 9000B is used. Latest Unified Wire drivers for Windows and Linux and installed in Initiator and target machines respectively.




Figure 2- Test Setup

Storage Topology and Configuration

Each LIO target is configured in offload mode with 8 Ramdisk LUNs, each of 600MB size. The iSER initiator connects to the targets and lomometer is run with single worker thread per target device and 8 outstanding I/Os.

Conclusion

This paper provided iSER performance results of Chelsio's T62100-CR Unified Wire adapter in Windows Server 2019 setup. The benchmark results show that T6 delivers superior bandwidth, reaching line rate for both READ and WRITE. As clear from the results, iSER over Chelsio's iWARP RDMA can achieve high performance without the need for a new fabric that is not compatible with the large Ethernet installed base.

Related Links

- [Windows NVMe over Fabrics Performance](#)
- [Secure Data Replication at 100GbE](#)
- [High Performance S2D with Chelsio 100GbE](#)
- [Windows SMB 3.1.1 Performance at 100Gbps](#)