

iWARP/RDMA Benefits in Windows 10

Chelsio T5 Client RDMA w/ Windows 10 Enterprise

Executive Summary

In this era of growing data, host systems are faced with unprecedented demand to handle complex and larger size of data sets, requiring expensive CPUs with bigger pools of memory DIMMS. This deluge of data is being fueled by the exponential growth in the parallel processing and capabilities adopted by many datacenter applications to address the growing demand of networking and storage requirements. This also presses the need for a truly converged and an all in-boxed I/O solution to meet today's demand of growing enterprise and cloud implementations.

To leverage a consistent performance enhancement on the client/host side, there is a need to have a network fabric that incorporates RDMA (Remote Data Memory Access) to create a large data pipe that can be used to move data quickly with low latency, high bandwidth and minimum CPU resources.

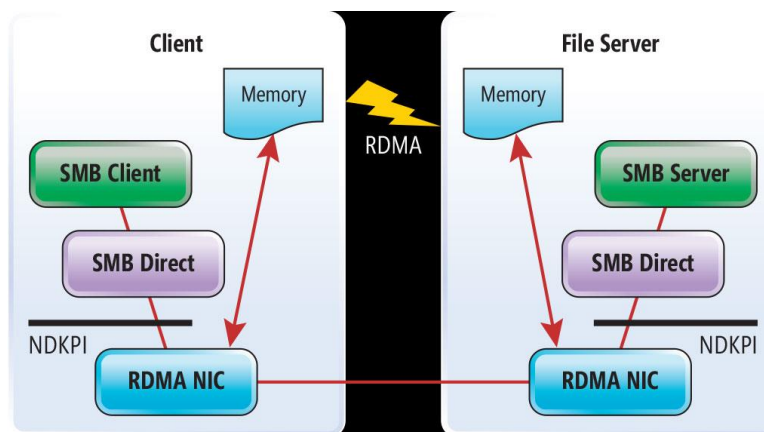


Figure 1 – RDMA NIC End-to-End Topology

Windows 10 Enterprise provides built-in support for using RDMA/SMB Direct client with Ethernet iWARP (RDMA/TCP) NICs, to support high-speed data transfers. iWARP based RNICs implement RDMA in hardware so as a result, data transfer is extremely fast with client-to-file server and performance almost equaling that of using local storage. RDMA NICs offload the server CPU, resulting in a high performance and more efficient solution.

Chelsio's Terminator 5 and recently released T6 ASIC solution offer a high performance, robust implementation of iWARP RDMA over 10/25/40/50/100Gb Ethernet Unified Wire adapters, delivering end-to-end RDMA latency that is comparable to InfiniBand, using a standard Ethernet infrastructure. Chelsio's iWARP is in production today and drivers are in-boxed. It is a high performance, robust, reliable and mature protocol that enables direct data placement, CPU savings, and RDMA functionality over TCP/IP and legacy Ethernet switches and internet with no performance penalties.

The Demonstration

The end-to-end RDMA solution setup consisted of a Windows Server 2016 Data Center Edition server machine connected to Windows 10 Enterprise client machine, using Chelsio 10Gb and 40Gb Ethernet Unified Wire adapters. The results show how iWARP RDMA enabled NIC can save up to 50% of CPU cycles/byte and also consume less %CPU compared to a standard NIC.

Test Configuration

The following sections provide the test setup and configuration details.

Topology

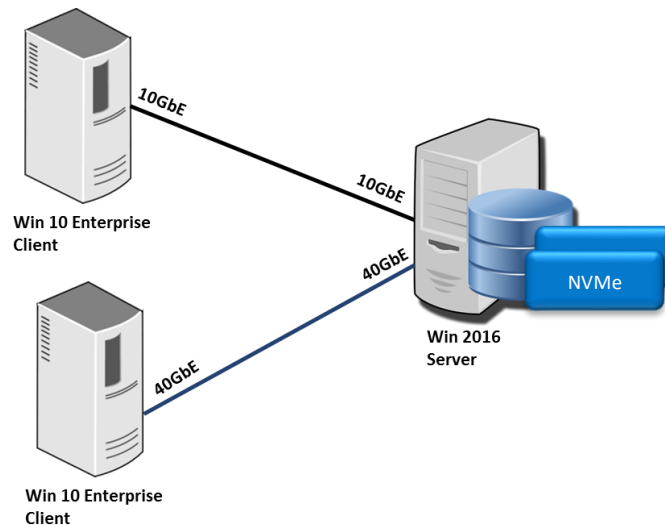


Figure 2 – Test Setup

Network Configuration

Windows Server 2016 node: Server node is installed with 2 Intel Xeon CPU E5-2687W v2 processors clocked at 3.40GHz (HT disabled), 64 GB of RAM, 4 SSD block devices and Microsoft Windows Server 2016 Datacenter operating system. Chelsio T520-LL-CR and T580-CR adapters are installed to provide 10GbE and 40GbE links speeds.

Windows 10 Enterprise node with 10GbE connectivity: This client node in the setup is installed with Intel Xeon CPU E5-1660 v2 @ 3.70GHz, 64 GB of RAM and Microsoft Windows 10 Enterprise operating system. Server and client nodes are connected back-to-back using Chelsio T520-LL-CR 10Gb Ethernet adapters on both ends and using Chelsio inbox SMB Direct and NIC drivers.

Windows 10 Enterprise node with 40GbE connectivity: This client node in the setup is installed with two Intel(R) Xeon(R) CPU E5-2687W v2 @ 3.40GHz, 64 GB of RAM and Microsoft Windows 10 Enterprise operating system. Server and client nodes are connected back-to-back using Chelsio T580-CR 40Gb Ethernet adapters on both ends and using Chelsio inbox SMB Direct and NIC drivers.

Command used

WRITE: diskspd.exe -b<bs> -w100 -r -o8 -t8 -W10 -d30 -C10 -n -Sh
Z:\testfile1.dat Y:\testfile1.dat X:\testfile1.dat W:\testfile1.dat

READ: diskspd.exe -b<bs> -w0 -r -o8 -t8 -W10 -d30 -C10 -n -Sh
Z:\testfile1.dat Y:\testfile1.dat X:\testfile1.dat W:\testfile1.dat

Test Results:

Server and Client connected at 10G link speed:

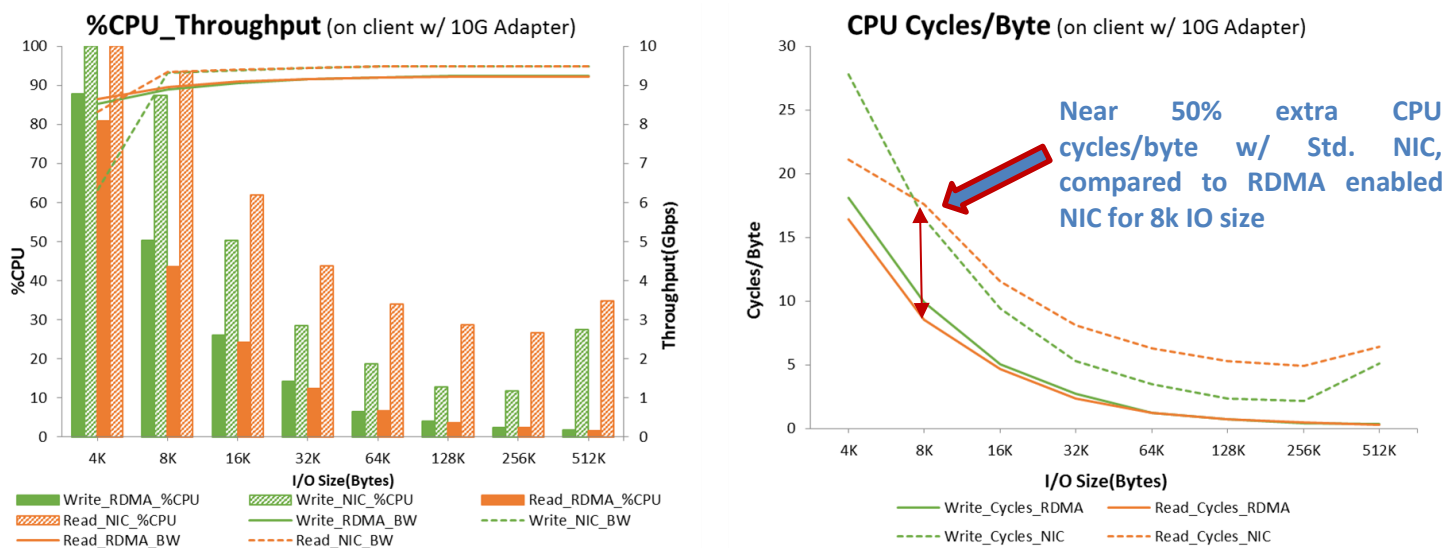


Figure 3 – Throughput, %CPU and CPU Cycles/Byte with Chelsio 10G adapters

Server and Client connected at 40G link speed:

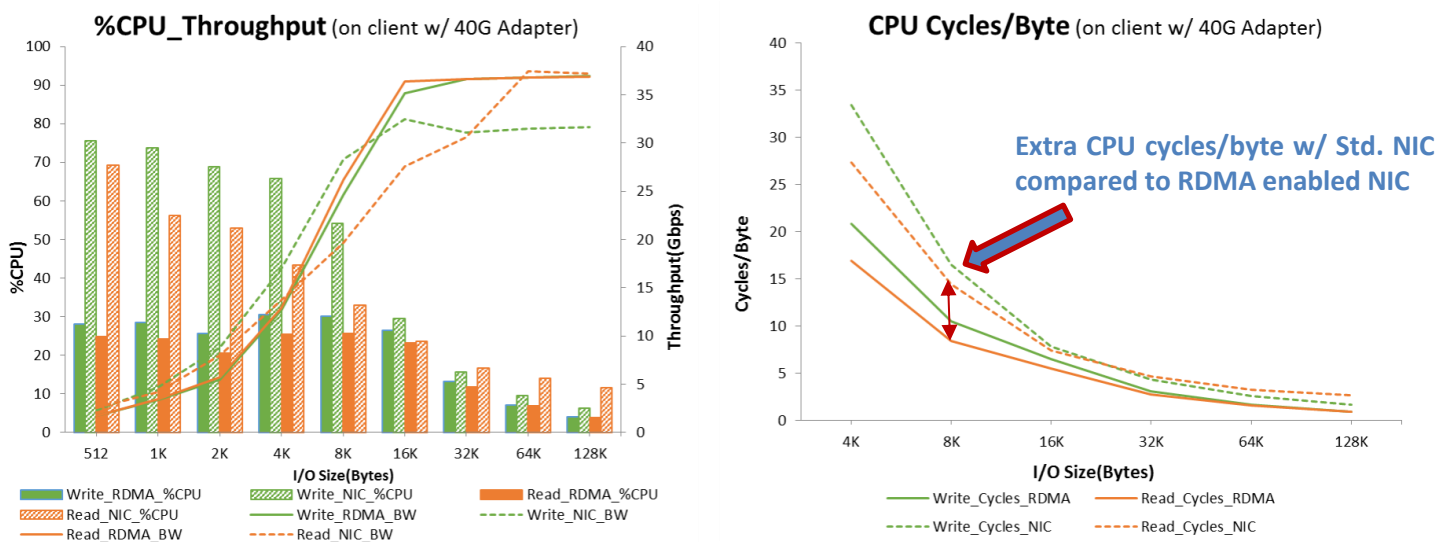


Figure 4 – Throughput, %CPU and CPU Cycles/Byte with Chelsio 40G adapters

Conclusion

This paper showcases the significant performance benefits of Chelsio T5 iWARP RDMA enabled 10GbE and 40GbE adapters with Windows 10 Enterprise. Chelsio Unified Wire 10/40Gb Ethernet adapters can save %CPU and also consume much less CPU cycles/byte compared to a standard NIC. Above test results shows:

- **CPU Cycles per Byte:** Standard NIC consumed almost 50% extra CPU cycles/byte compared to RDMA enabled NIC for 8k IO size. RDMA enabled NIC uses less number of processor cycles to move each Byte and keeps the CPU free for the other tasks.
 - Example:
 - **8K Read @10G - RDMA: 8.52, NIC: 17.61, Write - RDMA: 9.91, NIC: 16.48**
 - **8K Read @40G - RDMA: 8.43, NIC: 14.45, Write - RDMA: 10.51, NIC: 16.50**
- **Throughput and %CPU:** Chelsio 10GbE and 40GbE adapters not only deliver line rate performance for both RDMA enabled and standard NIC interfaces but uses much less %CPU for iWARP/RDMA interface by leveraging direct data placement capabilities and keeps the CPU free for the other applications.

Overall a 40G NIC utilizes 55% CPU on writes, it indicates that chances are high to exhaust a CPU and/or run out of memory BW in a system with a 100G NIC adapter and application will never be able to experience the true 100G capacity, while a RDMA enabled NIC utilizes 50% less CPU for the same amount of data transfer at the same IO size and showcases, how Chelsio offload capabilities enable a high performance I/O without worrying about the CPU cost.

Related Links

[Storage IOPS Update with S2D - Microsoft Blog](#)

[High Performance S2D with Chelsio 40GbE](#)

[The Chelsio Terminator 5 ASIC](#)

[Windows Server 2012 R2 SMB Performance](#)

[iWARP: Ready for Data Center and Cloud Applications](#)