

Windows Server 2012 R2 Network Direct Chelsio 40GbE T580-LP-CR vs. Mellanox MCX353A ConnectX-3 VPI 56G IB

Latency and Bandwidth Benchmark Results

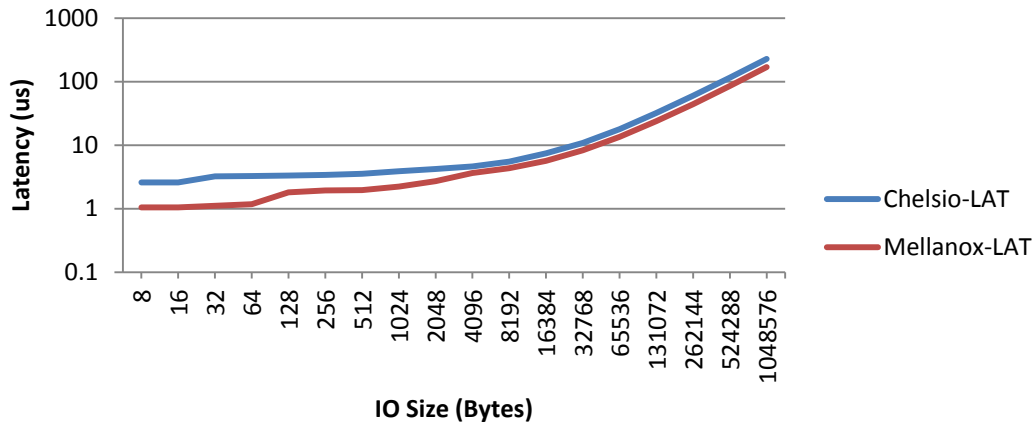
Chelsio's iWARP RDMA over Ethernet capability enables a user process on one system to transfer data directly between its virtual memory and the virtual memory of a process on another system without operating system intervention on either side of the communication. RDMA accomplishes this by offloading onto the "channel adapter" interface card the tasks traditionally performed by the operating system during network transfers. The result is high throughput, high message rate, low latency, and low CPU utilization message transfer.

RDMA is a protocol that provides a user-space I/O interface, which thanks to polling, enables low latency communications, and zero copy data transfer. In the RDMA paradigm, the network interface handles all protocol processing, and the CPU is practically bypassed and therefore fully offloaded. This translates to substantial reduction in CPU utilization, allowing power efficient and cost effective processors to be used to provide the same performance as more costly models. Chelsio's adapters implement the RDMA over Ethernet standard (iWARP) in hardware.

The Network Direct Service Provider Interface in Microsoft Windows Server 2012 R2 allows access to the RDMA capabilities of iWARP and InfiniBand adapters. This paper presents micro-benchmark results comparing 40GbE iWARP RDMA and IB FDR performance in raw bandwidth and latency using Windows Network Direct. Although such simple metrics are not necessarily always reflected in actual application performance, this paper shows that today, more than ever before, the two technologies are on a level footing.

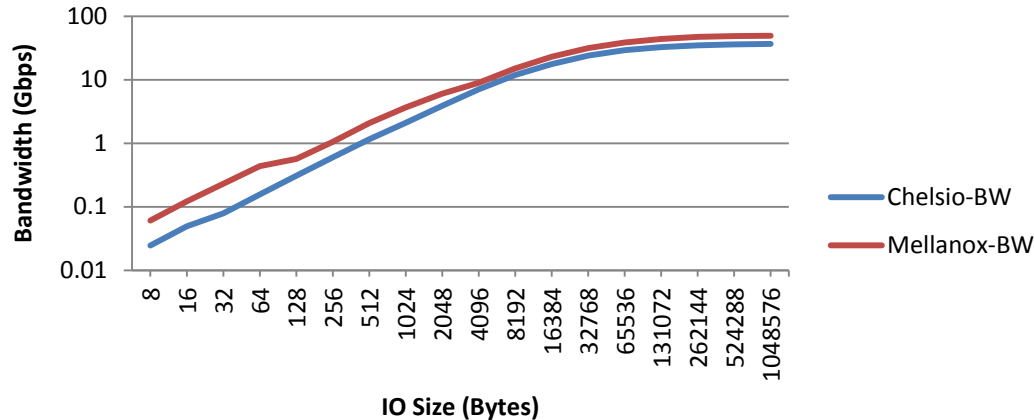
The first graph below shows latency numbers that are similar between Chelsio's 40GbE RDMA performance and Mellanox's 56G IB, especially with IO sizes of 128B or larger.

Windows RDMA 40GbE vs FDR IB Latency



Similarly, the graph below shows comparable bandwidth between the Chelsio and Mellanox adapters as the IO size increases beyond 1KB.

Windows RDMA 40GbE vs. FDR IB Bandwidth



Conclusion

This paper presents raw performance data comparing the Chelsio T580-LP-CR 40GbE iWARP adapter and Mellanox ConnectX-3 FDR InfiniBand adapter. Traditionally, InfiniBand had enjoyed a performance advantage in raw bandwidth and latency micro-benchmarks. However, with the latest 40GbE and 100GbE standards, Ethernet now shares the basic physical layer with IB and has essentially caught up on these basic metrics. In addition, Chelsio's adapter offers a high performance implementation of the iWARP standard which eliminates any perceived advantage to IB, obviating the need for using an esoteric technology to get high performance and efficient communication.