

Windows Server 2012 R2 SMB Performance

RDMA vs. NIC Throughput, IOPS and CPB

Executive Summary

A most notable feature of Windows Server 2012 R2 is the release of SMB 3.0, which seamlessly leverages RDMA-enabled adapters for improved performance and efficiency, when available. This paper demonstrates the benefits of SMBDirect (SMB over RDMA) through a performance comparison of Chelsio's T520-LL-CR RDMA-enabled adapter and Intel's X520-DA2 server adapter. While the T520-LL-CR shows a better performance profile across the board, the results clearly show that when RDMA is in use, significant performance and efficiency benefits are obtained.

Overview

Remote Direct Memory Access (RDMA) technology, which underlies the new SMB Direct protocol (part of SMB version 3.0) is a transport medium for SMB, which can be utilized to achieve unprecedented levels of performance and efficiency. Chelsio worked closely with Microsoft to enable this functionality via its high performance implementation of RDMA over Ethernet – the Internet Wide Area RDMA Protocol (iWARP).

One of the main advantages of the new SMB 3.0 implementation is that once the network adapter driver is installed, all its features are automatically enabled. Furthermore, with the new multi-channel SMB technology, Windows can choose the best protocol to use at any time, as well as aggregate traffic over multiple different links using different protocols.

The combination of Chelsio's T5 technology and Microsoft's SMB 3.0 therefore results in a highly efficient plug-and-play solution that can move large amount of data at high speed with minimal CPU utilization.

This paper demonstrates these benefits through a performance characterization of an RDMA enabled adapter, Chelsio's 520-LL-CR, compared to the high end Intel X520-DA2 server adapter, which lacks this capability. The two transport modes are compared in terms of I/O per second (IOPS), throughput, and CPU efficiency (expressed in cycles per byte – CPB). A lower CPB value indicates a more efficient data transfer.

In addition, SMB over iWARP also benefits from greatly improved data integrity protection, thanks to iWARP's end-to-end payload CRC (in lieu of simple checksums for the standard NIC). Being especially designed for storage networking, T5 incorporates additional reliability features, including internal datapath CRC and ECC-protected memory.

Test Results

The following graphs compare the throughput and IOPS performance results, for the two adapters at different I/O size.




Figure 1 – RDMA and NIC Throughput Comparison

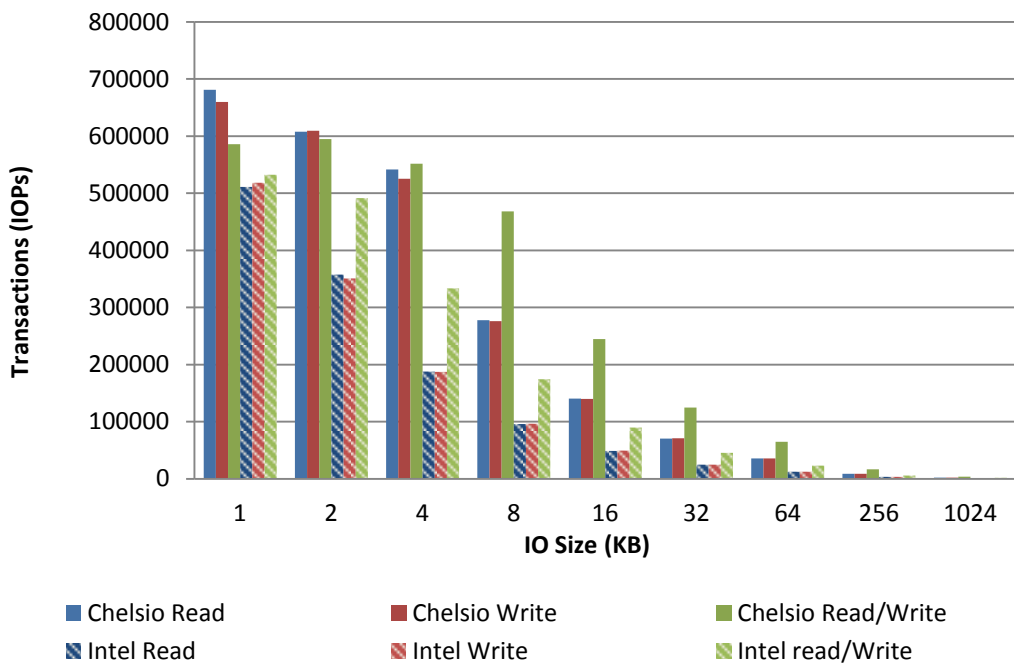


Figure 2 – RDMA and NIC IOPS Comparison

The results above reveal that Chelsio’s adapter enjoys superior performance throughout, reaching line rate unidirectional throughput at 4KB I/O size, more than 2x the Intel adapter. However, when RDMA kicks in as the I/O size exceeds 4KB, performance shoots up to near line rate in bidirectional performance, about 3x the results of Intel’s adapter. The graph below further validates how, at that point, data transfer costs significantly decrease as well, reaching close to 50% those of the regular NIC.

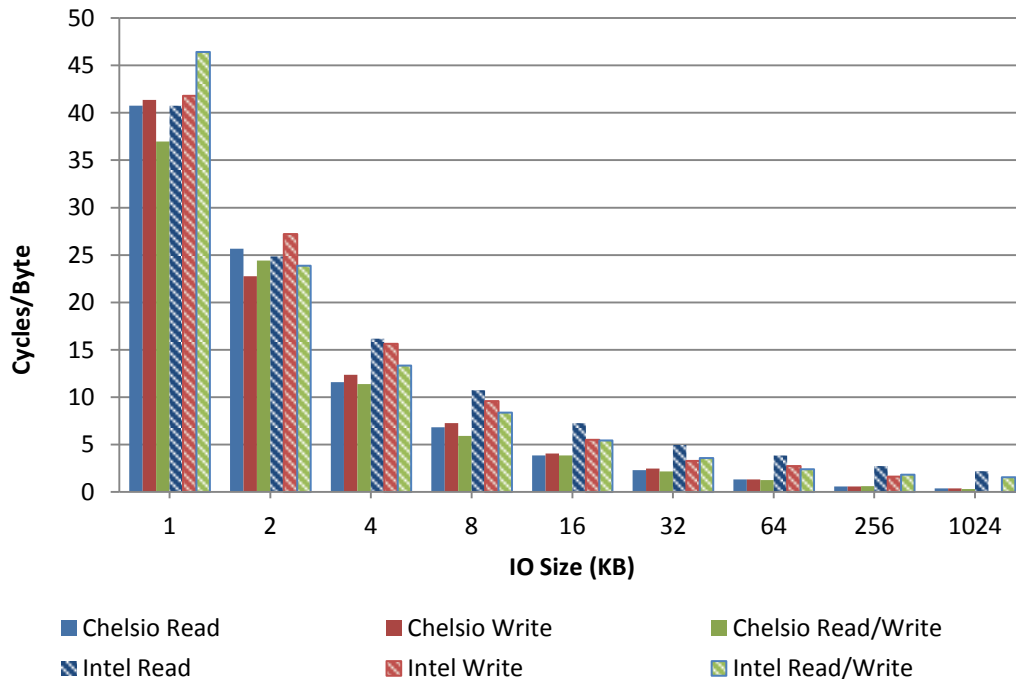



Figure 3 – RDMA and NIC CPU Utilization Comparison (Lower is Better)

The tests results therefore clearly establish that SMB significantly benefits from an RDMA transport in increased performance and efficiency compared to regular NICs. These benefits are automatically and transparently enabled when using Windows Server 2012 R2, minimizing the user’s management and configuration burden.

Test Topology

The following diagram shows the test setup and topology.



Test Configuration

The test configuration consists of 2 machines connected back-to-back. A Server and Client, each with Intel Xeon CPU E5-1660 v2 processor clocked at 3.70GHz, with 64 GB of RAM, run the Windows Server 2012 R2 operating system. Standard MTU of 1500B is configured.

The Chelsio setup uses one T520-LL-CR adapter installed in each system, with driver version 5.2.1.0.

The Intel setup uses one X520-DA2 adapter installed in each system, with inbox driver.

Benchmarking Tool

Microsoft's **sqlio v2.15** is used to assess the I/O performance of the configurations. This test used I/O sizes varying from 1KB to 1MB.

Command Used

```
sqlio2.exe -s<time in sec> -t<#threads> -o<#outstanding> -b<io size in KB> -T<percentage of RW mix> -BN -LS -frandom -d<drive> testfile.dat
```

Conclusion

This paper provided performance results for SMB 3.0 running over Chelsio's T5 RDMA enabled Ethernet adapter and compared it to Intel's X520-DA2 non-RDMA server adapter. The results demonstrate the benefits of RDMA in improved performance and efficiency. The following are the takeaway data points:

- Chelsio's T5 reaches 3x the performance of Intel X520-DA2 for READ, WRITE and READ/WRITE when RDMA is used
- Chelsio T5's cycles per byte results are 25% to 80% lower than those of Intel X520-DA2 when RDMA is used

With plug-and-play operation, enhanced reliability, higher efficiency and line rate performance, SMB over Chelsio's T5 series of RDMA enabled adapters is a very compelling solution for Windows Server R2 storage networking.

Related Links

[The Chelsio Terminator 5 ASIC](#)

[SMBDirect over Ethernet using iWARP on Windows Server 2012 R2](#)

[Delivering Superior Performance with SMB over Ethernet on Windows Server 2012](#)