

# Windows SMB 3.0 Performance at 40Gbps

## RDMA vs. NIC Throughput, IOPS and CPU Utilization

---

### Executive Summary

A notable feature of Windows Server 2012 R2 is the release of SMB 3.0 with SMB Direct (SMB over RDMA), which seamlessly leverages RDMA-enabled network adapters for improved storage bandwidth, latency and efficiency. This paper demonstrates the benefits of SMB Direct through a performance comparison of Chelsio's T580-CR RDMA-enabled adapter and Intel's latest XL710 40GbE server adapter. While the T580-CR exhibits a better performance profile in both NIC and RDMA operation, the results clearly show that when RDMA is in use, significant performance and efficiency gains 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 that can be utilized to achieve unprecedented levels of performance and efficiency. Chelsio worked closely with Microsoft to enable this functionality via Chelsio's high performance hardware implementation of the iWARP protocol, the open IETF standard for RDMA over Ethernet.

One of the main advantages of the SMB 3.0 implementation is that once the network adapter driver is installed, all its features are automatically enabled and made available to the SMB application. 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 Chelsio's T580-CR RDMA enabled adapter, comparing it to the high end Intel XL710 40GbE server adapter, which lacks RDMA support but provides stateless offload NIC capabilities. The RDMA and NIC transport modes are compared in terms of I/O per second (IOPS), throughput, and CPU efficiency (expressed in cycles per byte – CPB).

In addition to performance and efficiency gains, SMB over iWARP 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). Furthermore, being especially designed for storage networking, T5 incorporates additional reliability features, including internal datapath CRC and ECC-protected memory. For these reasons, T5 adapters have been selected for Microsoft's Cloud Platform System, a scalable, turnkey private cloud solution.

## Test Results

The following graphs compare the throughput and IOPS benchmark results for the two adapters at different I/O sizes.




Figure 1 – RDMA and NIC Throughput Comparison


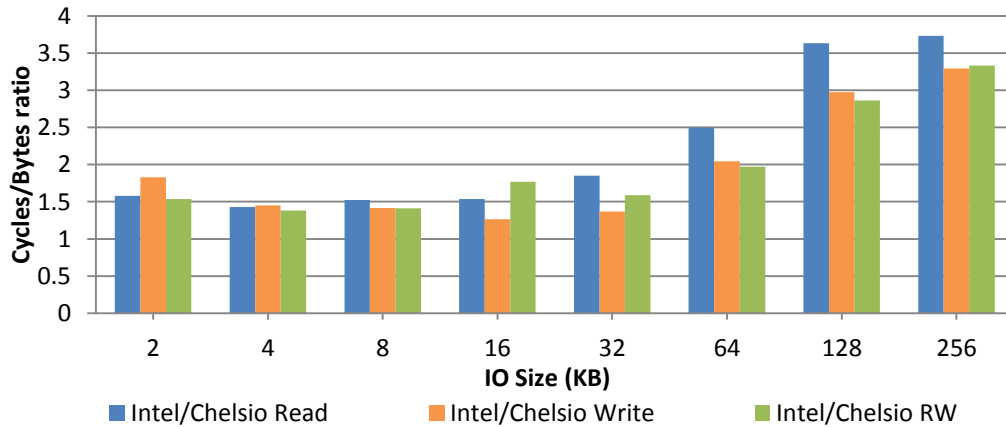


Figure 2 – RDMA and NIC IOPS Comparison

The above results reveal that Chelsio’s adapter provides significantly higher and more consistent performance, reaching line rate unidirectional throughput at 8KB I/O size. The difference is particularly clear when RDMA kicks in as the I/O size exceeds 4KB, with up to 2x the performance of the Intel NIC in bidirectional transfers. The Intel adapter also exhibits large variability in the same test environment, symptomatic of performance corners.

The graph below plots the CPB ratio for the two adapters, showing the Intel adapter at 1.5x the CPU utilization, rising to 3.5x as the I/O size increases and RDMA is in use.




**Figure 3 – RDMA and NIC Cycles/Bytes Ratio**

The tests results therefore 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 users management and configuration burden.

## Test Configuration

The following sections provide the test setup and configuration details.

### Topology



**Figure 4 – Simple Back-to-Back Test Topology**

### Network Configuration

The test configuration consists of 2 machines connected back-to-back using a single port: a Server and Client, each with 2 Intel Xeon CPU E5-2687W v2 8-core processors clocked at 3.40GHz and with 64 GB of RAM. Windows Server 2012 R2 operating system is installed on both the machines. Standard MTU of 1500B is configured.

The Chelsio setup uses one T580-CR adapter installed in each system, with Windows Unified Wire software v5.0.0.36, whereas the Intel setup uses one XL710 adapter installed in each system, with driver v1.0.114.

### Storage Topology and Configuration

The storage array is configured with 20 test files, each of 1 GB size.

### I/O Benchmarking Configuration

Microsoft's **Diskspd v2.0.0** is used to assess the I/O performance of the configurations. This test used I/O sizes varying from 2KB to 256KB.

### Commands Used

```
C:\Users\Administrator> diskspd.exe -b<IO Size> -w<0/100/50> -l -r -o10 -W20 -d20 -t2 -h  
\\<server_ip>\<share_name>\<testfile1>.dat ... \\<server_ip>\<share_name>\<testfile20>.dat
```

### Conclusion

This paper provided performance results for SMB 3.0 running over Chelsio's T5 RDMA enabled Ethernet adapter T580-CR comparing it to Intel's XL710 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 up to 2x the throughput of Intel for READ/WRITE when RDMA is used.
- Intel's CPU utilization is 1.5x to 3.5x that of Chelsio's T5 adapter 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 2012 R2 storage networking.

### Related Links

[The Chelsio Terminator 5 ASIC](#)

[SMBDirect Latency on Windows Server 2012 R2](#)

[SMBDirect 40 GbE iWARP vs 56G InfiniBand](#)

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