



## RDMA over Ethernet for Windows Server 2012

---

The release of Microsoft Windows 8 and Windows Server 2012 is expected to be one of the largest, most feature rich Operating System releases for Microsoft. Besides a new user interface, Windows Server 2012 introduces a comprehensive set of solutions for virtualization, networking, and high availability storage, bringing to mainstream technologies that were to date confined to the High Performance Computing arena.

### SMB 3.0 over iWARP

Remote Direct Memory Access (RDMA) is one of the technologies that are making the transition from specialized HPC use to mainstream networking and storage applications. In Windows Server 2012, RDMA underlies the new SMB Direct protocol, a major enhancement to the SMB 3.0 version shipping with this release. By utilizing RDMA as a transport for SMB, unprecedented levels of performance and efficiency can be achieved. Chelsio worked closely with Microsoft to enable this functionality on its Terminator 4 (T4) powered network adapters, making use of Chelsio's leading implementation of iWARP, the standard for RDMA over Ethernet.

The Internet Wide Area RDMA Protocol – iWARP – is layered above the tried and tested TCP/IP protocol stack. As a result, iWARP enjoys all the stack's known benefits: routability across subnets, scalability, robustness, and reliability. While iWARP can utilize the new Data Center Bridging Ethernet mechanisms, it does not require them. Instead, it functions well with cost-effective legacy switching technology, and allows decoupling the server and switch refresh cycles.

Chelsio's implementation of iWARP benefits from the high performance, low latency Terminator Architecture Offload Engine, which fully handles data transfer in hardware, with virtually no host CPU involvement. This frees up the CPU to do more of the real application work. In addition to RDMA, Chelsio's NICs simultaneously support iSCSI, FCoE and TCP offload, as well as high performance networking and virtualization functions.

Noteworthy of the offload support infrastructure in Windows Server 2012 is that once the network adapter driver is installed, all the adapter's features and protocols are automatically enabled. Furthermore, with the new multi-channel SMB technology, Windows can choose the best protocol to use on each port at any time, as well as spread traffic over multiple different links using different protocols. Thus, the combination of Chelsio's T4 technology and Microsoft's SMB 3.0 results in a solution of high availability, performance and efficiency.

### Comprehensive Solution

In addition to the iWARP RDMA functionality utilized by SMB 3.0, Chelsio's adapters support:

- Stateless offload for both IPv4 and IPv6 with Scalable Networking features
- Hyper-V virtualization

- TCP offload (Chimney) implemented in hardware
- Network Direct RDMA for HPC clustering
- iSCSI Initiator with Boot capabilities
- FCoE Initiator with Boot capabilities

Chelsio’s comprehensive solution releases the full value and feature set of Windows Server 2012 in a single driver package. Chelsio further offers one of the most complete sets of adapters in the industry. These include adapters with 2x10G ports, 4x10G ports, combination of 2x10G ports and 2x1G ports, and a 4x1G port card. Various physical interfaces are also supported such as SFP+, QSFP+, CX4, and RJ-45.

### Benchmark Results

The following are sample benchmark results, which illustrate the performance levels seen with SMB3.0 over RDMA. More comprehensive results are available at [www.chelsio.com](http://www.chelsio.com).

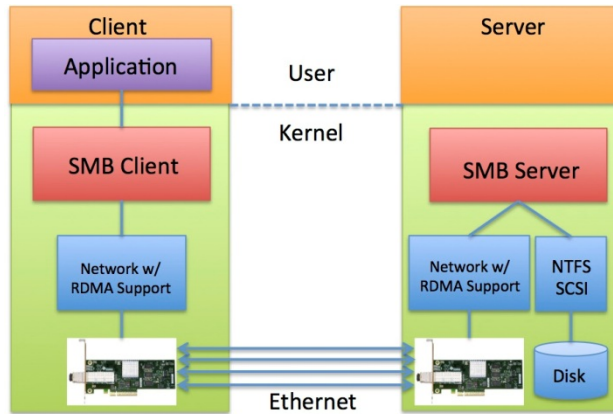


Figure 1 – System Diagrams




Figure 2 – Test Setup

Block Size	GB/sec	IOPS	Server CPU Utilization
8KB	2.6 GB/sec	334,827	45%
512KB	3.8 GB/sec	7,526	5%

Table 1 – Bidirectional sqlio Test Results

Tests run with the Chelsio T420-LP-CR card, driver version 4.0.0.14.

Test utility used is Microsoft’s sqlio program, with four 3GB RAM Disks. Command lines used are:

For 8K blocks: `sqlio.exe -s30 -T50 -t16 -o16 -b8 -BN -LS -fsequential -duvw testfile.dat`

For 512K blocks: `sqlio.exe -s30 -T50 -t2 -o16 -b512 -BN -LS -fsequential -duvw testfile.dat`