



FOR IMMEDIATE RELEASE

Media Contact:
media@chelsio.com
Chelsio Communications
1-408-962-3677

**CHELSIO TO DEMONSTRATE HIGH PERFORMANCE CUDA CLUSTERING WITH
40G ETHERNET AND NVIDIA GPUDIRECT**

SUNNYVALE, CA – November 11, 2015 – Chelsio Communications, Inc., a leading provider of Ethernet adapters for storage networking, virtualized enterprise data centers, cloud service installations, and cluster computing environments, today announced that it will demonstrate at SC15 (in Austin TX, November 16-19) high performance direct access to NVIDIA GPU processing at 40G speeds over Ethernet, showing significantly higher overall cluster throughput and much lower latency as compared to standard 40G server NICs.

The SC15 demo shows throughput increases up to 400% when Chelsio GPUDirect RDMA is enabled, with latency reduced by half, from more than 20us to less than 10us over most of the byte range.

With iWARP RDMA, network access to the GPU is achieved with both high performance and high efficiency. Since the host CPU and memory are completely bypassed, communication overhead and bottlenecks are eliminated, resulting in minimal impact on host resources, and translating to significantly higher overall cluster performance.

“We are pleased to be able to show the dramatic performance advantages of using iWARP RDMA with the leading accelerated computing platform from NVIDIA,” said Jim Johnston, senior director of marketing for Chelsio. “Our iWARP RDMA solution leverages existing Ethernet infrastructure and requires no new protocols, interoperability, or long maturity period as the no-risk path for Ethernet-based large-scale GPU clustering. We look forward to the continued increase in adoption of Ethernet in high performance applications.”



Demo Configuration: The demonstration consists of two servers connected back-to-back using a single 40G port. Each server uses an Intel Xeon CPU E5-160 v2 6-core processor clocked at 3.7GHz, with 64GB of RAM and RHEL6.3 operating system. Standard MTU of 1500B is configured. One Chelsio T580-CR and one Tesla K80 GPU adapter is installed in each system with Chelsio GPUDirect RDMA driver, CUDA v6.5, OpenMPI v1.8.4 9 with CUDA support), OSU micro benchmarking tools v4.4.1 and NVIDIA peer memory driver. openmpi was used to measure the throughput and latency, and the I/O sizes used varied from 1B to 8KB. The complete benchmark paper is available here: [GPUDirect RDMA over 40Gbps Ethernet](#).

About Chelsio iWARP

Chelsio's Terminator 5 ASIC offers a high performance, robust third generation implementation of RDMA (Remote Direct Memory Access) over 40G Ethernet, the Internet Wide Area RDMA Protocol – iWARP. T5 delivers end-to-end RDMA latency that is comparable to InfiniBand, using standard Ethernet infrastructure. Chelsio's iWARP is in production today in GPU applications, in storage applications as a fabric for clustered storage, for Lustre and other storage applications, for Microsoft Azure Stack storage, for HPC applications, and for remote replication and disaster recovery. 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.

About Chelsio Communications

Chelsio is a recognized leader in high performance (10G/25G/40G/50G/100G) Ethernet adapters for networking and storage within virtualized enterprise data centers, public and private hyperscale clouds, and cluster computing environments. With a clear emphasis on performance and delivering the only robust offload solution, as opposed to simple speeds and feeds, Chelsio has set itself apart from the competition. The Chelsio Unified Wire fully offloads all protocol traffic, providing no-compromise performance with high packet processing



capacity, sub-microsecond hardware latency and high bandwidth. Visit the company at www.chelsio.com.

###