

# FCoE at 40Gbps with FC-BB-6

## Line Rate Throughput and 4M IOPS over Terminator 5

---

### Executive Summary

This paper reports FCoE performance results for Chelsio's Terminator 5 (T5) ASIC running at 40Gbps. T5 is the industry's first high performance full FCoE offload solution with FC-BB-6 VN-to-VN and SAN management software support.

The results using a single FCoE target running over a Chelsio T580-CR Unified Wire Network adapter connected through a 40Gb switch to 5 FCoE initiators, show I/O numbers reaching nearly 4M per second, and line rate throughput starting at I/O sizes as low as 2KB. Further improvements are expected to be delivered by ongoing performance tuning.

### Overview

The Terminator 5 (T5) ASIC from Chelsio Communications, Inc. is a fifth generation, high-performance 2x40Gbps/4x10Gbps, unified wire engine with an SR-IOV capable PCIe Gen3 x8 interface.

T5 supports full FCoE offload, with hardware support for FC CRC, Direct Data Placement, T10-DIX enhanced data integrity, and FCoE RSS for CPU load balancing. T5 also supports the DCB protocol suite, including PFC, ETS and QCN, with a native DCBX protocol implementation for automatic configuration.

Furthermore, FCoE support in T5 is part of a complete, fully virtualized unified wire offload suite that includes iSCSI, RDMA, TCP/UDP sockets and user space I/O capabilities. All traffic types and protocols benefit from comprehensive traffic management functionality, and enhanced data integrity and reliability features.

Chelsio's FCoE solution is not only fully compliant with the FC-BB-5 standard, but is also ready for the new **FC-BB-6** standard. In fact, this paper provides results using the FC-BB-6 VN Port-to-VN Port mode (VN2VN), which improves performance and eliminates the need for an expensive Fibre Channel Forwarding (FCF) enabled switch.

In addition to T5's comprehensive FCoE capability list, the performance results shown in the below demonstrate that Chelsio's FCoE implementation delivers the highest performance available today.

### Test Results

The following graph plots the performance results, showing how line rate 40Gbps is achieved at I/O sizes as small as 2KB, with peak IOPS nearing 4M/sec for READ and 2.5M/sec for WRITE.

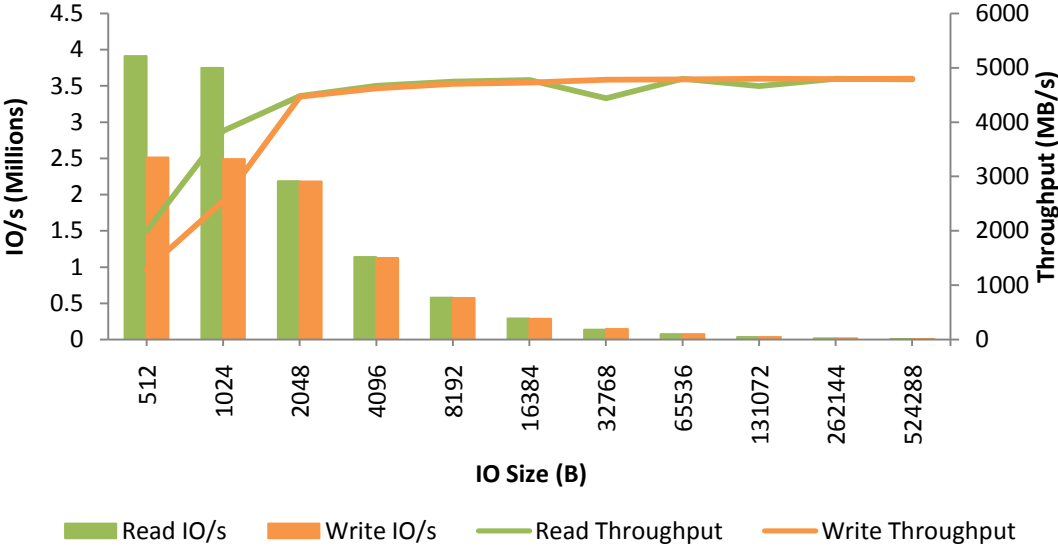


Figure 1 – Throughput and IOPS vs. IO size

The following graphs compares READ IOPS results at 40Gbps and 10Gbps, varying the READ I/O size. The results show that, starting at 2KB I/O size, the 40Gbps IOPS are 4 times those of 10Gbps, indicating perfect performance scaling.

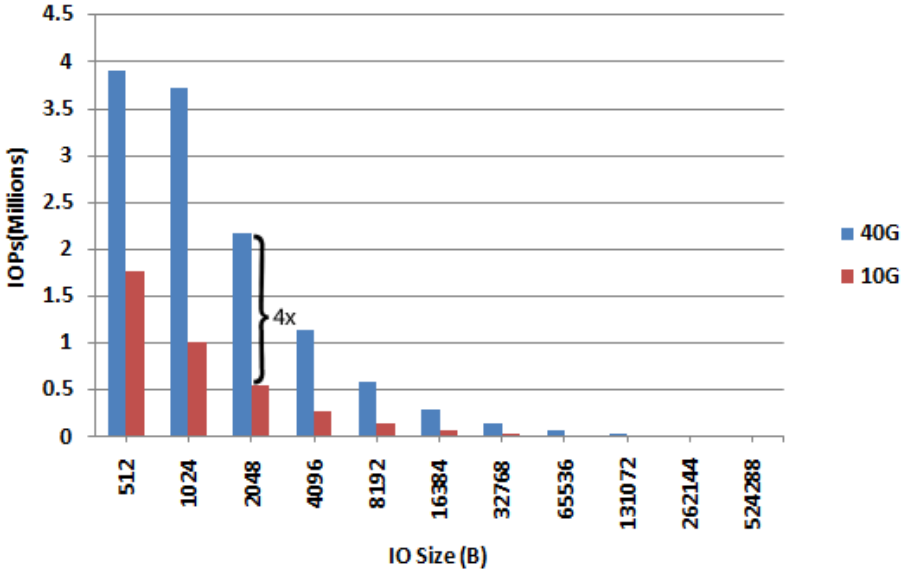


Figure 2 – READ IOPS at 40G and 10G

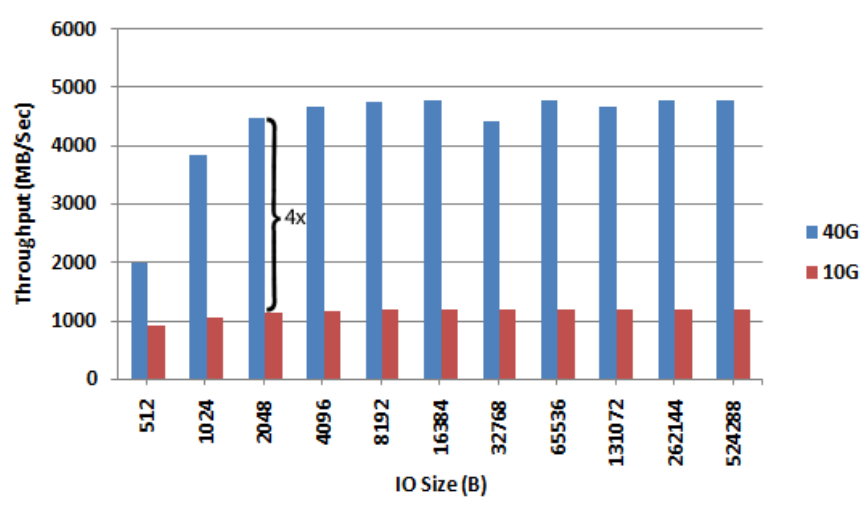


Figure 3 – READ Throughput at 40G and 10G

The throughput results above mirror those of the IOPS, and similar results were observed with WRITE testing.

### Test Configuration

The following sections provide the test setup and configuration details.

#### Topology

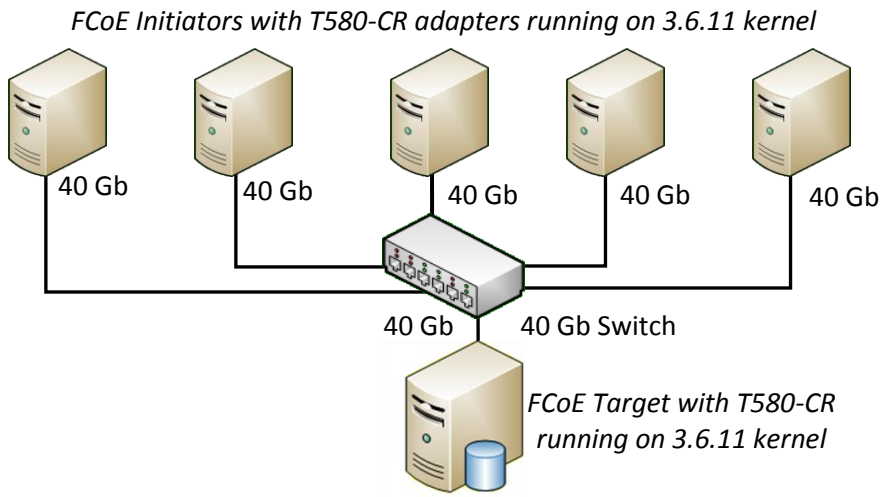


Figure 4 – Test Setup

#### Network Configuration

The network configuration consists of an FCoE target storage array connected through a single 40Gb port of a 40Gb switch to 5 FCoE initiator machines all connected with 40Gb links. An FCoE compatible MTU of 2180B was used.

- **The storage array** was configured with 2 Intel Xeon CPU E5-2687W v2 8-core processors (HT enabled) running at 3.40GHz with 64 GB of RAM. Chelsio T580-CR adapter was installed in the system with Chelsio's FCoE target driver and RHEL 6.4 (3.6.11 Kernel) operating system, taking advantage of the T5 ASIC's FCoE BB-6 VN2VN technology at 40Gb.
- Each of the **initiator machines** was setup with an Intel Xeon CPU E5-1660 v2 6-core processor (HT enabled) running at 3.70GHz with 64 GB of RAM. Chelsio T580-CR adapter was installed in each system with Open-FCoE and RHEL 6.4(3.6.11 kernel) operating system.

### Storage Topology and Configuration

The storage array contains 60 FCoE *ramdisk null-rw* targets. Each of the 5 initiators connects to 12 targets.

### I/O Benchmarking Configuration

**xdd** is used to assess the I/O capacity of a configuration. This test uses sample block sizes in the range of 512B-512KB. Buffering is set to *none*, and the I/O access pattern used is sequential READs and WRITEs.

### Command Used

```
[root@host]# xdd -sgio -op <read/write> -targets <#targets> -rwratio  
<0/50/100> <Sg luns e.g /dev/sg1> -queuedepth 16 -reqsize <IOSize> -blocksize  
512 -numreqs 2097152 -runtime <time> -datapattern sequenced -maxall
```

## Conclusion

This paper reported performance results for Chelsio's offloaded FCoE solution running the T580-CR Unified Wire Network adapter. The results demonstrate that:

- Chelsio's T5 delivers line rate 40Gbps FCoE performance from 2KB I/O size, with perfect scaling from 10Gbps
- IOPS **reach** nearly **4M** at 512B I/O size, and line rate starting at 2KB I/O size

Part of Chelsio's Unified Wire Ethernet solution, the T5 FCoE implementation provides the same unique combination of uncompromising performance and rich feature set as the rest of the offloaded protocols.

## Related Links

[The Chelsio Terminator 5 ASIC](#)

[iSCSI at 40Gbps](#)

[TCP Offload at 40Gbps](#)