

# FCoE at 40Gbps with FC-BB-6

## Line Rate Throughput and 2M+ IOPS over Terminator 5

---

### Executive Summary

This paper presents FCoE performance results for Chelsio's latest Terminator 5 (T5) ASIC running at 40Gbps. T5 is the industry's first high performance full FCoE offload solution with SAN management software support.

Performance results using a single FCoE target running over a Chelsio T580-CR Unified Wire Network adapter connected through a switch using one 40Gb port to 4 FCoE initiators, each connected at 40Gb, show I/O numbers in excess of 2M per second, and line rate throughput starting at IO sizes close to 2KB.

### 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 FCoE offload 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 eliminates the need for an expensive Fibre Channel Forwarding (FCF) enabled switch.

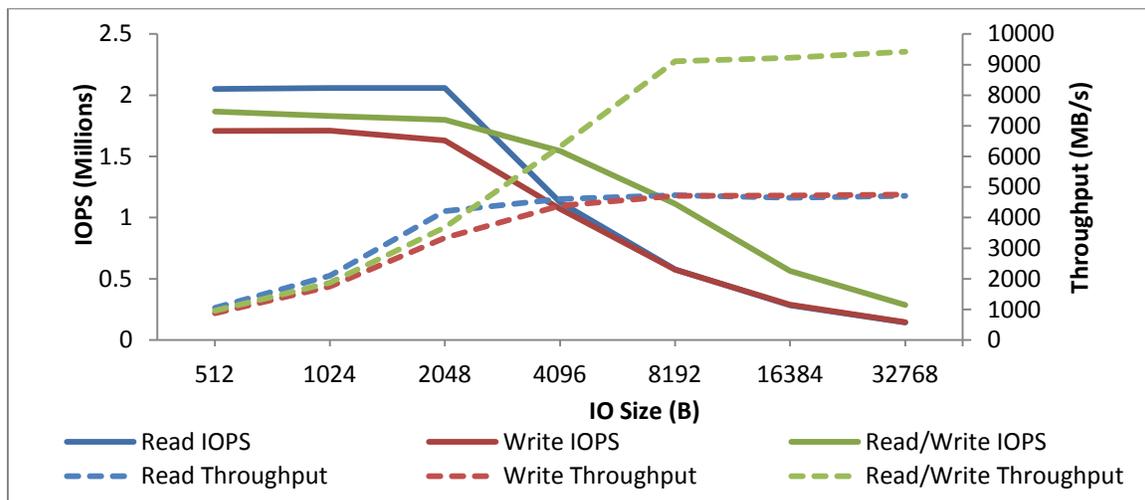
In addition to the comprehensive FCoE capability list, the performance results shown in the following section demonstrate that Chelsio's implementation is the highest performance available today.

## Test Results

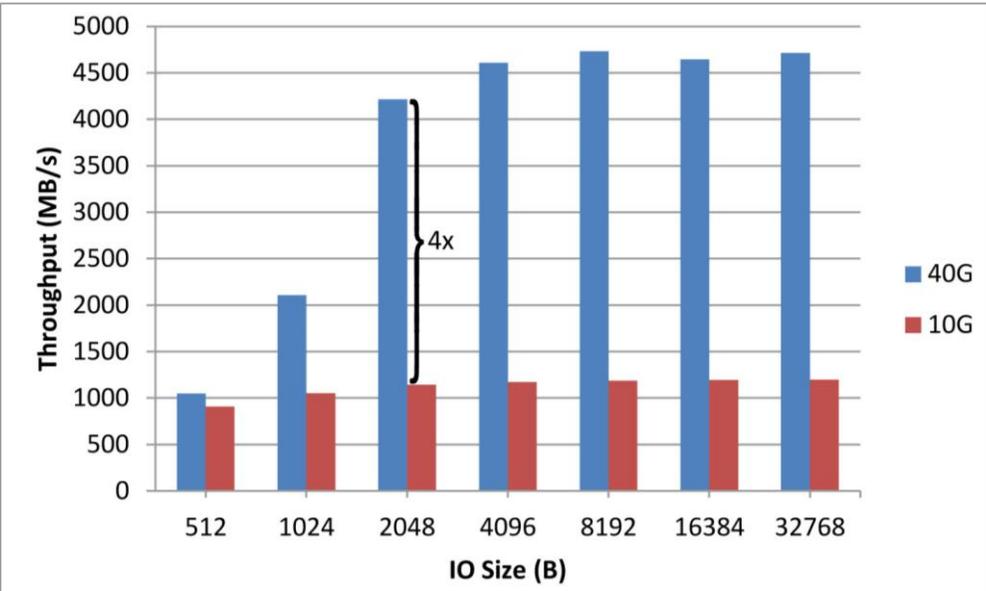
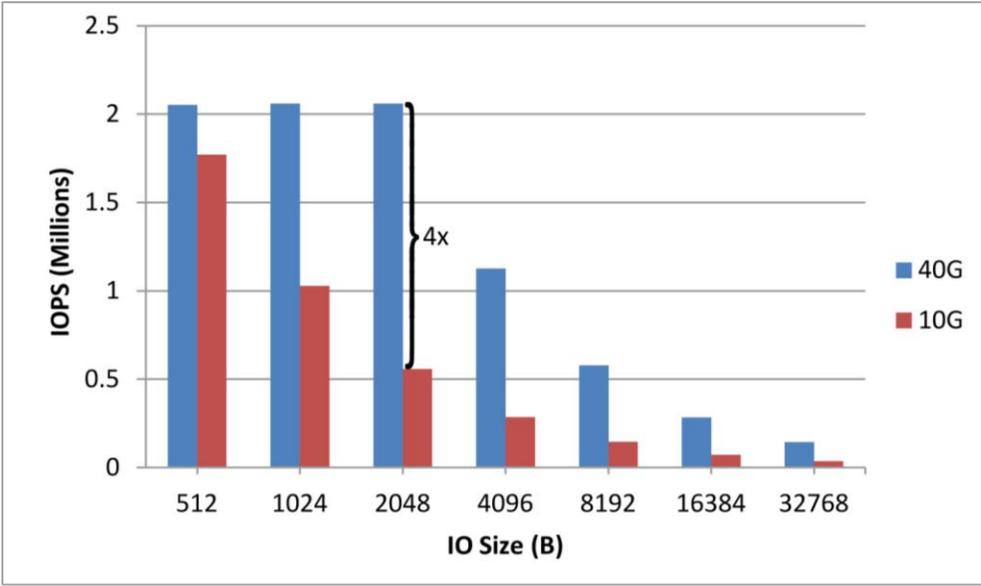
The following table summarizes the READ, WRITE and READ/WRITE throughput and IOPS numbers obtained varying the I/O sizes using the **xdd** tool. The maximum IOPS numbers consistently exceed 2M.

IO Size (B)	T580-CR FCoE (BB-6)								
	READ			WRITE			READ/WRITE		
	MB/s	IOPS	CPU (%)	MB/s	IOPS	CPU (%)	MB/s	IOPS	CPU (%)
512	1049	2050764	91	874	1707264	94	956	1868071	94
1024	2107	2058485	91	1751	1710061	94	1875	1831542	94
2048	4216	2058875	91	3337	1629621	92	3684	1798989	94
4096	4610	1125673	78	4388	1071334	90	6325	1544316	93
8192	4731	577552	66	4710	575010	74	9107	1111695	91
16384	4647	283630	59	4722	288231	67	9228	563280	79
32768	4713	143857	59	4755	145138	64	9423	287579	78

The following graph plots the performance results, showing how line rate 40Gbps is achieved even at I/O sizes as small as 2KB.

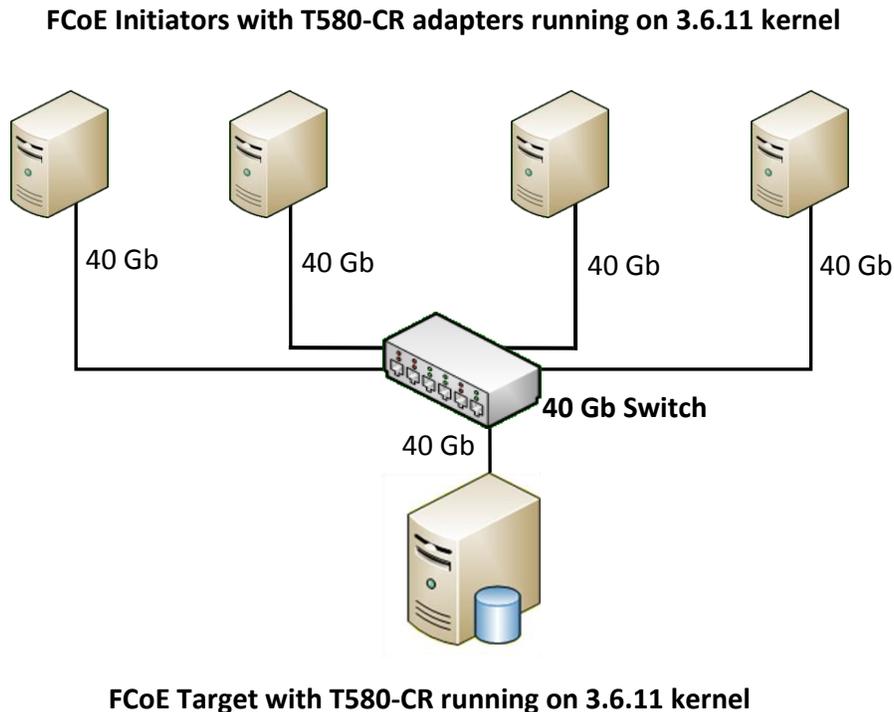


The following graphs plot the benchmarking data obtained varying I/O sizes with READ at 40Gbps and 10Gbps. The results show that, starting at 2KB I/O size, IOPS and throughput at 40Gbps are 4 times those of 10Gbps, indicating perfect performance scaling. Similar results were observed with WRITE and READ/WRITE.



## Test Topology

The following diagram shows the test setup and topology.



## Test Configuration

The following sections provide the test configuration details.

### Network Configuration

The network configuration consists of an FCoE target storage array connected to 4 FCoE initiator machines through a 40Gb switch. An FCoE compatible MTU of 2180B was used.

- **The storage array** was configured with an Intel Xeon CPU E5-2687W v2 processor 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.
- **The initiator machines** were each setup with an Intel Xeon CPU E5-1660 v2 processor running at 3.70GHz. 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 48 FCoE *ramdisk null-rw* targets. Each of the 4 initiators connects to 12 targets. The I/O sizes used varied from 512B to 32KB with an I/O access pattern of sequential READs and sequential WRITES.

### I/O Benchmarking Configuration

**xdd** is used to assess the I/O capacity of a configuration. This test used the following sample block sizes: 512B, 1KB, 2KB, 4KB, 8KB, 16KB and 32KB. Buffering is set to *none*, and the I/O access pattern is sequential.

### Command Used

```
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 provided 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 IO size
- IOPS **exceed 2M** for 512B IO size and reach the theoretical limit from 2KB IO size

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

## Related Links

[The Chelsio Terminator 5 ASIC](#)

[iSCSI over 40Gb Ethernet](#)

[TCP Offload at 40Gbps](#)