

# TCP Offload at 40Gbps

*Reclaim CPU Cores with TCP/IP Full Offload*

---

## Overview

Chelsio is the leading provider of network protocol offloading technologies, and Chelsio's Terminator *TCP Offload Engine (TOE)* is the first and currently only engine capable of full TCP/IP at 40Gbps.

The unique ability of a *TOE* to perform the full transport layer functionality in hardware is essential to obtaining tangible benefits. The vital aspect of the transport layer is that it provides process-to-process communication, i.e. the data passed to the *TOE* comes straight from the application process, and the data delivered by the *TOE* goes straight to the application process. In contrast, lower layers only provide unreliable delivery functionality, which limits the usability of the data at these levels. This opens the way for very powerful extensions to pure protocol offload, including:


1. Direct Data Placement (DDP), which addresses the memory subsystem bottleneck problem on receive
2. Direct Data Sourcing (DDS), which addresses the memory subsystem bottleneck problem on send
3. Application layer data integrity check (CRC) offload, typically used in data critical applications, which are not satisfied with the relatively weak Internet checksum protection (e.g., iSCSI header and payload digests)
4. Reliable remote direct memory to memory access with RDMA
5. Further application layer offload, such as application layer payload recovery for end-to-end security protocol offload
6. Per connection TCP level traffic management and quality of service

The Terminator series – Terminator 4 (T4) and Terminator 5 (T5) – adapters can flexibly offload TCP/IP processing per connection, per-server or per-interface, while selectively and simultaneously tunnel traffic from non-offloaded connections to the host processor for the native TCP/IP stack to process. The Terminator series adapters provide a powerful zero copy capability for regular TCP connections, requiring no changes to the sender or receiver applications, to deliver line rate performance at minimal CPU utilization, interrupts and context switches.

## Test Results

To illustrate the benefits of TCP/IP socket offloads, tests were run using 8-streams unidirectional and bidirectional *iperf* using the Chelsio T580-CR Dual Port 40G adapter, comparing the *TOE* enabled and disabled cases.

Test setup: 1 socket x4 core Hyperthreaded (8 core) Intel(R) Xeon(R) CPU E5-1620 0 @ 3.60GHz connected back to back using RHEL 6.2 64 bit, T5 uwire config file and modprobe t4\_tom.



The graph above shows both the bandwidth achieved and CPU utilization for the offload (*TOE*) and NIC cases. While both can achieve line rate bidirectional transfer performance, the *TOE* frees up 3 cores that otherwise would have been wasted in network protocol stack processing. This effectively allows 75% more of the investment in processing to be reclaimed for useful applications.