

OpenIndiana 40Gbps TOE vs. NIC Performance

Throughput Benchmark Results

Executive Summary

Chelsio is the leading provider of network protocol offloading technologies, and Chelsio's TCP Offload Engine (TOE) is the only engine capable of full TCP/IP at 40Gbps. Chelsio's latest Terminator 5 (T5) adapters can flexibly offload TCP/IP processing per connection, per-server or per-interface, while selectively and simultaneously acting as a stateless server adapter for traffic from non-offloaded connections, delivering full Ethernet frames 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.

This paper presents early benchmarks results of the TOE driver for OpenIndiana, an OpenSolaris derivate, using the T5 ASIC running at 40Gbps. The results provide a preview of the benefits of Chelsio's TCP offload technology over regular NIC adapters, showing that TCP offload provides a superior throughput curve, with significantly higher performance at small I/O sizes, while consistently utilizing a fraction of the CPU processing cycles of NIC mode. Performance tuning of the TOE driver is in progress and is expected to deliver additional gains for TOE mode.

Overview

The Terminator 5 (T5) ASIC from Chelsio Communications, Inc. is a fifth generation, high-performance 2x40Gbps/4x10Gbps server adapter engine with Unified Wire capability, allowing offload storage, compute and networking traffic to run simultaneously. In addition to the comprehensive full offload capability, T5 provides extensive support for stateless offload operation for both IPv4 and IPv6 (IP, TCP and UDP checksum offload, Large Send Offload, Large Receive Offload, Receive Side Steering/Load Balancing, and flexible line rate Filtering). T5 is also a fully virtualized NIC engine with separate configuration and traffic management for 128 virtual interfaces, and includes an on-board switch that offloads the hypervisor v-switch.

Thanks to integrated, standards based FCoE/iSCSI and RDMA offload, T5 based adapters are high performance drop in replacements for FibreChannel storage adapters and InfiniBand RDMA adapters. They also excel at normal server adapter functionality, providing high packet processing rate, high throughput and low latency for common network applications. This paper presents T5's OpenIndiana TCP offload performance using the T580-CR adapter, showing superior throughput at small I/O sizes, which are typical of many applications. Unlike RDMA, TOE performance benefits are obtained without any application changes, using the standard BSD socket interface. This makes it an attractive instant ON, zero effort solution for popular applications as well as infrastructure suites such as Cloud Stacks and Big Data installations.

Test Results

The following graph compares unidirectional throughput and CPU utilization results for offloaded (TOE) and non-offloaded (NIC) traffic at different I/O sizes, using the **netperf** tool.

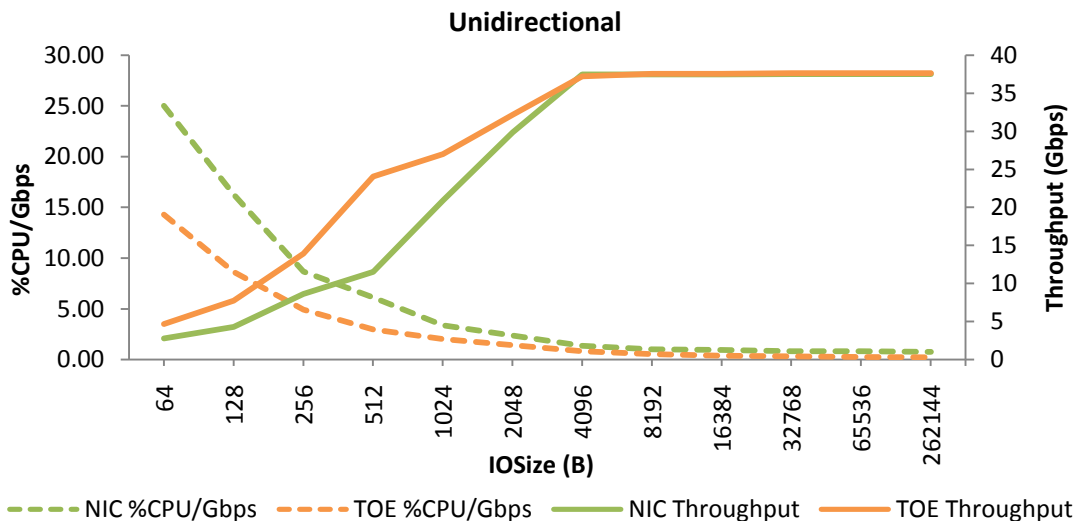


Figure 1 – Unidirectional Throughput and %CPU vs. I/O size

The results show a significantly higher throughput at small I/O sizes (below 4KB), with TOE achieving up to 2x the NIC performance. Furthermore, the TOE’s CPU usage numbers (in %CPU per Gbps) are again about half those of the NIC, resulting in savings that are made available for user applications.

The following graph compares bidirectional throughput and CPU utilization results for offloaded (TOE) and non-offloaded (NIC) traffic at different I/O sizes, using the **netperf** tool.

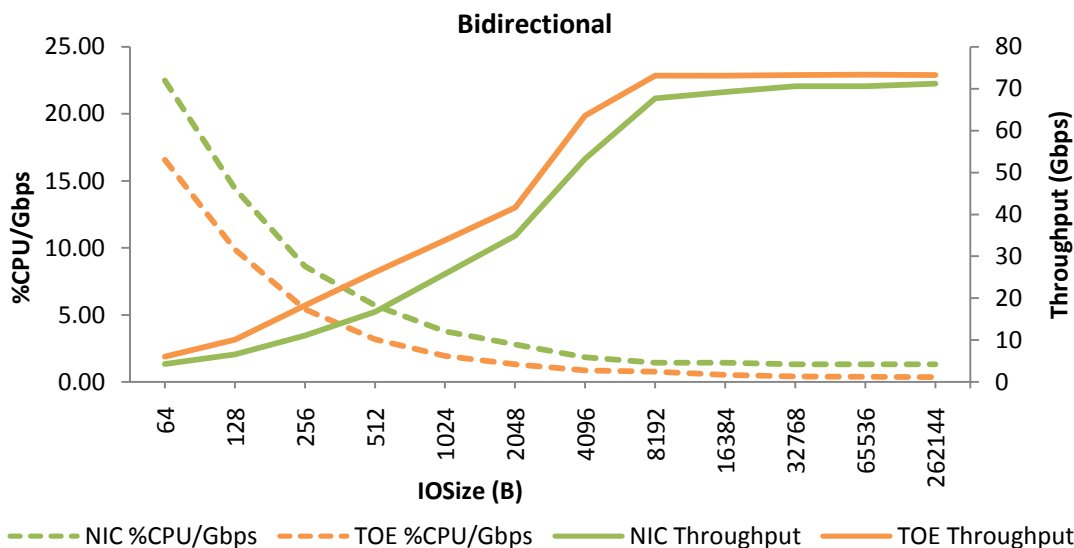


Figure 2 – Bidirectional Throughput and %CPU vs. I/O size

The bidirectional results mirror those of the unidirectional case, with both throughput and CPU efficiency gains transparently obtained through enabling TCP offload under the hood.

Test Configuration

The following sections provide the test setup and configuration details.

Topology




Figure 3 – Simple Back-to-Back Test Topology

Network Configuration

The test setup consists of 2 machines connected back-to-back using a single port, a Server and Client, each configured with Intel Xeon CPU E5-1660 v2 hexa-core processor clocked at 3.70GHz (HT enabled) and 64GB of RAM. Chelsio's Network driver v3.0.0.5 was installed with OpenIndiana Development oi_151.1.8 x86 operating system. Standard MTU of 1500B is used.

I/O Benchmarking Configuration

netperf was used to measure network throughput. This test used I/O sizes varying from 64B to 256KB.

Command Used

```
root@host:~# netperf -t TCP_STREAM -H <Server_IP> -l <time_duration> -- -P  
<local_port>,<Remote_port> -m <local message length> -M <Remote message length>
```

Conclusion

This paper provided performance results for Chelsio's T5 ASIC the OpenIndiana distribution, an OpenSolaris derivate, comparing Chelsio's T580-CR server adapter running in NIC and TOE modes. While the NIC performance is good, the results show Chelsio's adapter providing better throughput in TOE mode at small I/O sizes, reaching 40Gbps line rate at 4KB for unidirectional and 8KB for bidirectional traffic. Additionally, the TOE mode noticeably improves CPU efficiency, freeing up CPU resources for useful application processing.

Related Links

[The Chelsio Terminator 5 ASIC](#)
[40Gb TOE vs NIC Performance](#)
[Solaris/OpenIndiana at 40Gbps](#)