

FreeBSD NIC at 40Gbps

Extreme Performance with Terminator 5

Executive Summary

This paper presents NIC performance results for FreeBSD using Chelsio's latest Terminator 5 (T5) ASIC running at 40Gbps. The results show 40Gbps line rate throughput with I/O sizes as small as 2KB for one-way traffic and 4KB for bidirectional traffic, using standard 1500B Ethernet frame sizes.

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. T5 enables offloaded storage, compute and networking traffic to run simultaneously. T5 also provides extensive support for stateless offload operation for both IPv4 and IPv6 (IP, TCP and UDP checksum offload, Large Send Offload, Larger Receive Offload, Receive Side Steering/Load Balancing, and flexible line rate Filtering). In addition, T5 is 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. However, they also excel at normal server adapter functionality, providing high packet processing rate, high throughput and low latency for common network applications. This paper demonstrates T5's FreeBSD performance using the T580-LP-CR adapter, showing line rate throughput even at small I/O sizes.

Thanks to an in-box driver in the FreeBSD kernel, T5-based adapters are plug-and-play solutions for extreme networking performance, with current or forthcoming support of feature rich capabilities such as traffic management, security and filtering, RDMA and iSCSI.

Test Results

The following graph summarizes the TX, Rx and BiDi throughput numbers along with CPU utilization seen while varying the I/O sizes using the **iperf** tool. A standard MTU size of 1500B was used.

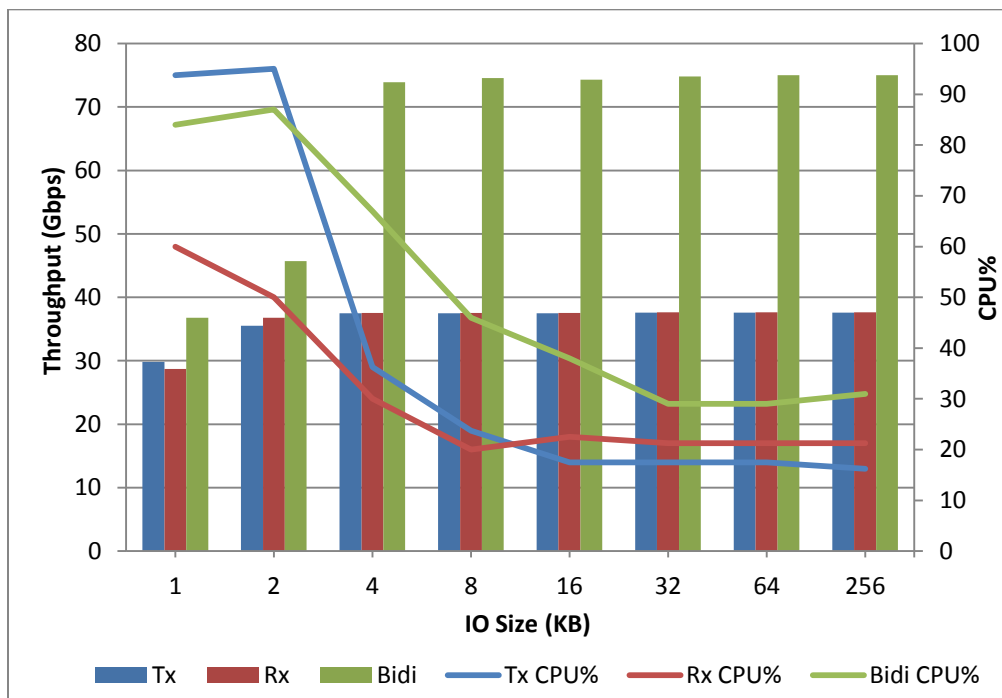


Figure 1 – Throughput and CPU Utilization

The maximum throughput numbers consistently reach line rate (throughput exceeds 37G) even at I/O sizes as small as 2KB.

Test Setup

The following diagram shows the test setup and topology.




Figure 2 – Simple Back-to-Back Test Topology

Test Configuration

The following sections provide the test configuration details.

Network Configuration

The test configuration consisted of 2 identical machines; each with Intel Xeon CPU E5-1660v2 8-core processor running at 3.70GHz, connected back-to-back using a single 40Gbps link. Standard MTU of 1500B was used.

One Chelsio T580-LP-CR adapter was installed in each system with FreeBSD 10.0-RELEASE operating system. The Chelsio adapters utilized the in-box driver.

I/O Benchmarking Configuration

iperf is used to measure network throughput. This test used sample IO sizes varying from 1KB to 256KB.

Commands Used

On the Server:

```
root@host:~# iperf -s -p <port> -w 512k
```

On the Client:

```
root@host:~# iperf -c <Server IP> -p <port> -l <IO Size> -t 30 -P 8 -w 512k
```

Conclusion

This paper provided performance results for Chelsio's T5 ASIC in FreeBSD, using Chelsio's T580-LP-CR server adapter. T5 delivers line rate 40Gbps NIC performance starting at 2KB IO size.

Related Links

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