

FreeBSD 40GbE TOE Performance

High Efficiency with Direct Data Placement (DDP)

Chelsio is the leading provider of network protocol offloading technologies, and Chelsio's TCP Offload Engine (TOE) is the first and currently only engine capable of full TCP/IP at 10/40Gbps. Direct Data Placement (DDP) addresses the memory subsystem bottleneck issue on receive, by defining how data can be directly placed into an upper layer protocol's (ULP) receive buffer without intermediate buffers, thus enabling data transfers with minimal CPU utilization.

This paper presents FreeBSD TOE DDP Performance results using Chelsio's T580-CR Unified Wire Adapter. The results show T5 delivering line rate throughput even with a single connection. Furthermore, the results show remarkable improvements in CPU usage when DDP is in use, compared to regular NIC and TOE traffic. Thanks to an inbox 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 compares the single port unidirectional receive throughput results collected with **32K I/O** size over a range of network connections, using the **netperf** tool. The results show that Chelsio T5 delivers high and consistent performance in all the three modes (NIC, TOE, TOE with DDP) delivering line rate throughput, even with a single connection.

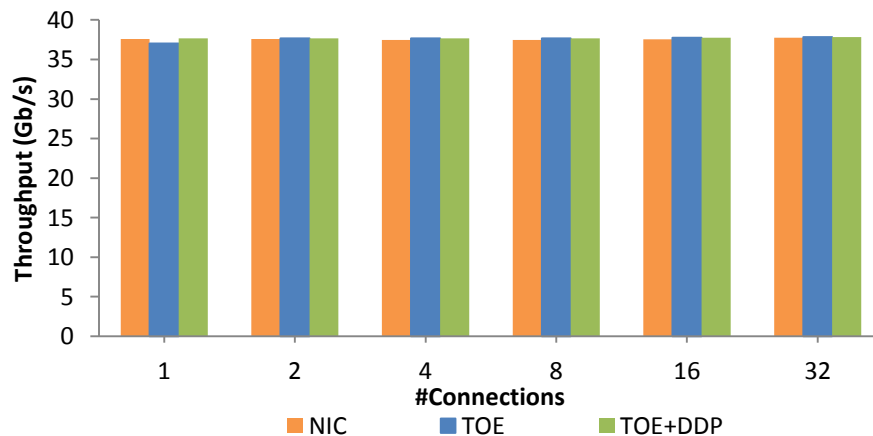


Figure 1 – Unidirectional Throughput vs. #Connections

The below graph compares the CPU Utilization results collected with **32K I/O** size over a range of network connections. The CPU Utilization in NIC mode increases dramatically as the number of connection increases. However, with DDP enabled, CPU utilization in TOE mode is negligible and remains constant even as the number of connection increases. The CPU Usage is up to 65% and 18% lower than regular NIC and TOE respectively, indicative of a more efficient data processing path. In addition, DDP does not pollute the CPU caches unnecessarily and thus saves bandwidth on the memory bus.




Figure 2 – CPU Utilization vs. #Connections

Test Setup

The test configuration consists of a Server connected back-to-back to a Client, using single 40Gbps link and standard MTU of 1500B. Both machines are configured with 1 Intel Xeon CPU E5-2687W v2 8-core processor running at 3.40GHz, 64 GB RAM and 1 Chelsio T580-CR adapter each. On the Server, FreeBSD Kernel is installed from <https://github.com/bsdjhb/freebsd.git>, ddp_bufid branch (commit d942f7638fa4c0f62900630b7846a76452030eaf). On the client, FreeBSD kernel is installed from the head repo.

Commands Used

Netperf is installed on both the machines from <https://github.com/bsdjhb/netperf2>, ddp-bufid-2.6.0 branch (commit 21df0c0f15e3ef6b5712d062b13d77dc0e15e403).

On the Server:

```
root@host:~# netserver -4
```

On the Client:

```
root@host:~# netperf -H <server_ip> -cC -t TCP_STREAM -l 30 -- -d -m 32768
```

Conclusion

This paper illustrated the benefits of using Chelsio's T580-CR unified wire adapter to provide exceptional bandwidth and CPU savings in FreeBSD. In NIC, TOE, and TOE with DDP mode, T5 maintains a superior profile across the board, reaching line rate for single as well as multiple network connections. Also, when DDP is in use, T5 in TOE mode delivers up to 65% reduction in CPU usage compared to NIC, and 18% less compared to regular TOE. The freed up CPU processing cycles can be used by other applications.

Related Links

[The Chelsio Terminator 5 ASIC](#)

[40Gb TOE vs NIC Performance](#)

[Linux 10GbE NIC/TOE Performance](#)

[OpenIndiana 40Gbps TOE vs. NIC Performance](#)

[TCP Offload at 40Gbps](#)