

Industry's First 100G Offload with FreeBSD

Zero copy TCP @100Gbps with Less than 1% CPU usage using DDP

Executive Summary

Chelsio is the leading provider of network protocol offloading technologies, and with the new Terminator 6 (T6), Chelsio's TCP Offload Engine (TOE) is the first and currently only engine capable of full TCP/IP offload at 100Gbps. The Terminator series adapters provide a powerful zero copy capability for regular TCP connections, requiring no changes to sender or receiver applications, to deliver line rate performance at minimal CPU utilization, interrupts and context switches.

This paper presents benchmark results of Chelsio's FreeBSD TOE driver using Terminator 6 (T6) adapters running at 100Gbps. The results provide a preview of Chelsio's TCP offload technology using Direct Data Placement (DDP), with a superior throughput and minimized CPU processing cycles. Furthermore, T6 delivers line rate throughput even with a single connection.

Overview

The Terminator 6 (T6) ASIC from Chelsio Communications, Inc., is a highly integrated, hyper-virtualized 1/10/25/40/50/100GbE controller with full offload support of a complete Unified Wire solution. Furthermore, it scales to true 100 Gigabit line rate operation, from a single TCP connection to thousands of connections, and allows simultaneous low latency and high bandwidth operation thanks to multiple physical channels through the ASIC. T6 also has integrated TLS/SSL, DTLS, IPsec and SMB 3.X crypto offload support, thus enabling concurrent secure communication and secure storage.

Chelsio's TCP Offload Engine enables reliable data transfers with minimal CPU utilization, freeing up CPU resources to be used for application processing. This will enable savings in host CPU acquisition, power and operational costs. T6 ASIC 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. Thanks to an inbox driver in the FreeBSD kernel, T6-based adapters are plug-and-play solutions for extreme networking performance. T6 adapters, in addition, support advanced features such as traffic management, security and filtering, RDMA and iSCSI.

The Demonstration




Figure 1 - 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 1 Intel Xeon CPU E5-1620 v4 4-core processor (HT disabled) clocked at 3.50GHz, 32GB of RAM and FreeBSD head repo. Chelsio's T62100-LP-CR is installed in each machine and configured with inbox drivers. netperf tool is installed with AIO patches that were accepted upstream. MTU of 9000B is used.

Test Results

netperf tool is used to measure the TCP throughput.

Commands Used

The following sysctl variables are set on both server and client to enable DDP and TX Zero Copy:

```
[root@host~]# sysctl dev.t6nex.0.toe.tx_zcopy=1
[root@host~]# sysctl dev.t6nex.0.toe.ddp=1
[root@host~]# sysctl kern.ipc.maxsockbuf=4194304
```

On the Server:

```
[root@host~]# netserver -4
```

On The Client:

```
[root@host~]# netperf -cC -H <Server_IP> -D 10 -l 30 -- -aA -m 512k -s 2M -S 2M
Recv  Send  Send  Utilization  Service Demand
Socket Socket Message Elapsed  Send  Recv  Send  Recv
Size  Size  Size  Time      Throughput  local  remote  local  remote
bytes bytes bytes  secs.    10^6bits/s  % C   % C    us/KB  us/KB

2097152 2097152 524288 30.00    95451.08  8.43  0.98  0.033  0.003
```

As observed, T6 delivers line rate throughput even with a single connection. Also, the CPU utilization numbers for receive (only 1%) and send (8%) are indicative of a more efficient data processing path. In addition, DDP does not pollute the CPU cache unnecessarily and thus saves bandwidth on the memory bus.

Conclusion

This paper illustrated the benefits of using Chelsio's T6 TOE solution in providing exceptional bandwidth and CPU savings in FreeBSD. The new T6 adapter, T62100-LP-CR, performs exceptionally well, with a consistent line-rate bandwidth. Also apparent from the results is noticeably improved CPU efficiency, freeing up resources for useful application processing. TCP Offload is needed to realize the current 100Gbps and the forthcoming 200Gbps and 400Gbps bandwidths.

Related Links

[The Chelsio Terminator 6 ASIC](#)

[FreeBSD 100Gb demonstration using single connection](#)

[FreeBSD 40GbE TOE Performance](#)

[High Performance iSCSI at 100GbE](#)