Chelsio 100G DPDK Performance

High Speed User Space Packet Processing with Terminator 6

Overview

The Terminator 6 (T6) is Chelsio's next generation of highly integrated, hyper-virtualized 1/10/25/40/50/100Gb Ethernet ASIC Solution. T6 is built around a programmable protocol-processing engine, with full offload of a complete Unified Wire solution comprising NIC, TOE, iWARP RDMA, iSCSI, OvS, DPDK, NAT, and Crypto functionalities. T6 based Ethernet adapters scale to a true 100Gb line rate operation from a single TCP connection to thousands of connections, and allows simultaneous low latency and high bandwidth operation, thanks to programmable multiple physical channels through the ASIC.

This paper compares Data Plane Development Kit (DPDK) packet rate performance results using Chelsio T6 and Mellanox ConnectX-4 Ethernet Adapters running at 100Gb, measuring Tx and Rx rate for each adapter. The MPPS numbers reflect the performance advantages of T6 adapter, particularly at the challenging small I/O sizes that are more representative of actual application requirements.

Chelsio DPDK – Poll Mode Driver (PMD)

The Chelsio T6 DPDK driver packages for Linux and FreeBSD platforms are a collection of data plane libraries and NIC drivers optimized for running in user space to boost packet per second performance. Chelsio's 1/10/25/40/50/100GbE CXGBE PMD (Poll Mode Driver) supports the following key features:

- Multiple queues for Tx and Rx
- Receive Side Scaling (RSS)
- VLAN filtering
- Checksum offload
- Promiscuous mode
- All multicast mode
- Port Hardware/Traffic statistics
- Jumbo frames
- Classification and Filtering

Chelsio T6 adapters also support an offload solution for Traffic Management, Classification, Filtering, and Crypto functions. Once combined and configured with the DPDK interface, these functions further enhance the DPDK functionalities by enabling network security, monitoring and QoS for both incoming/outgoing traffic.

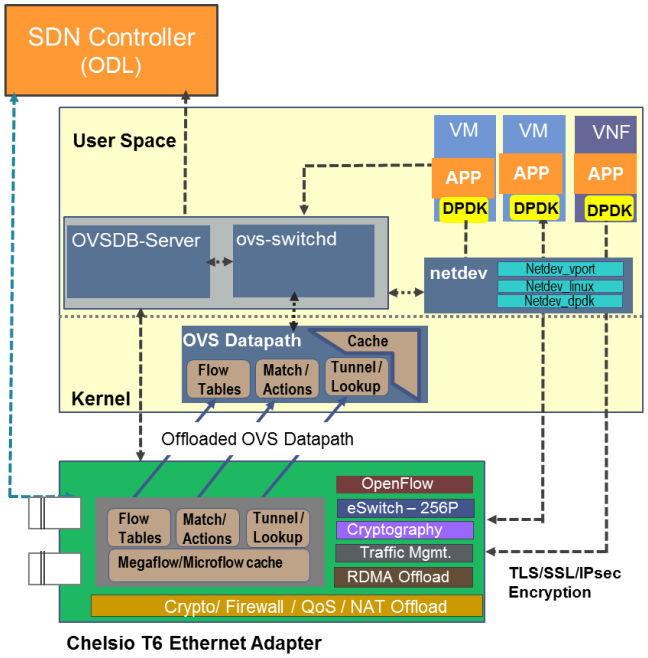

Figure 1 – Chelsio DPDK and OVS Offload Solution

The Chelsio solution provides a command line interface for packet classification and filtering features available in hardware. SDK/API are also available to integrate third party management applications.

Chelsio T6 adapters support OVS Kernel data path offload by offloading flow match entries and action processing onto the Chelsio adapter (at port level) for hardware acceleration of OVS datapath processing. This is achieved by adding, removing and synchronizing the OVS flow table entries on the adapter. Chelsio CLI/GUI tools can be used to enable and configure these features.

Test Results

The following graphs compare the single port Tx and Rx packet processing rate (MPPS) achieved by the T6 and ConnectX-4 adapters, using DPDK-Pktgen tool at different I/O sizes from 64B to 1024B.

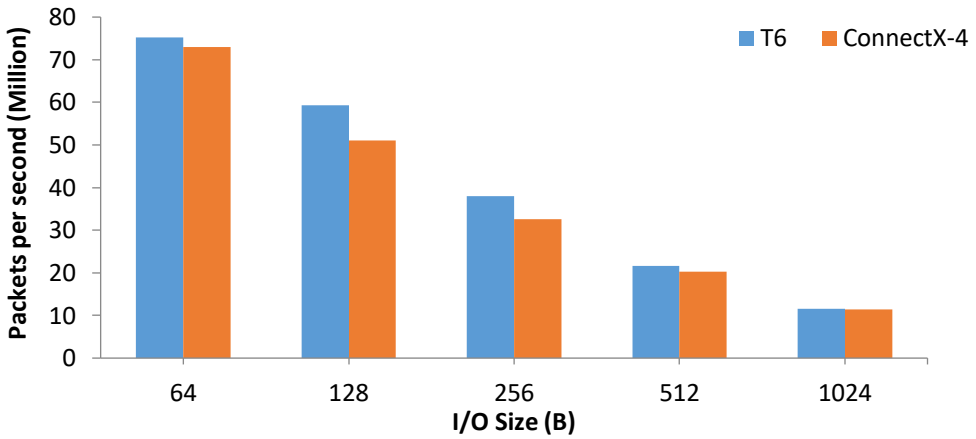

Figure 2 – Tx Traffic Performance

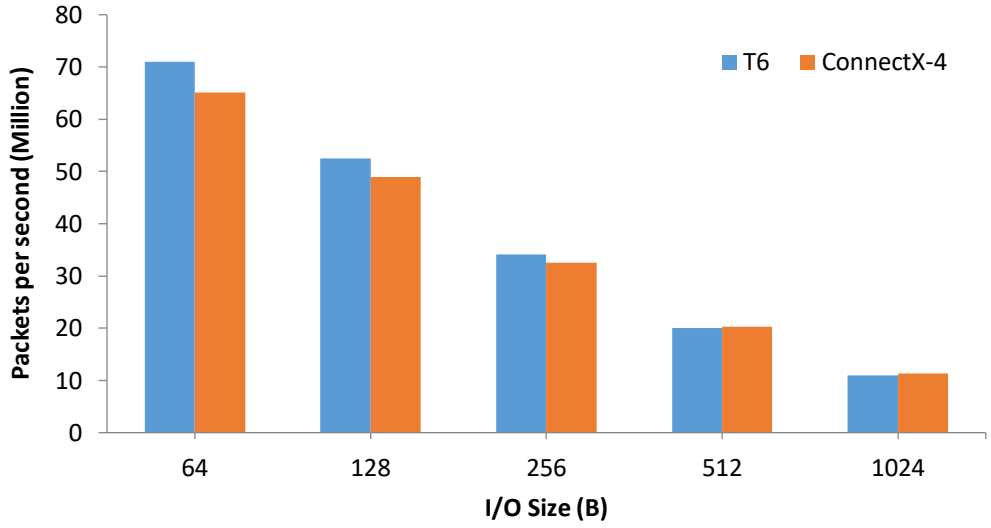

Figure 3 – Rx Traffic Performance

The Chelsio DPDK solution achieves higher packet rate for both transmit and receive directions across the range of I/O sizes. The Performance at lower I/O sizes upto 75 MPPS, makes T6 the ideal choice for high-performance applications, including wire analytics, traffic generation and monitoring applications that so far required specialized hardware.

Test Configuration

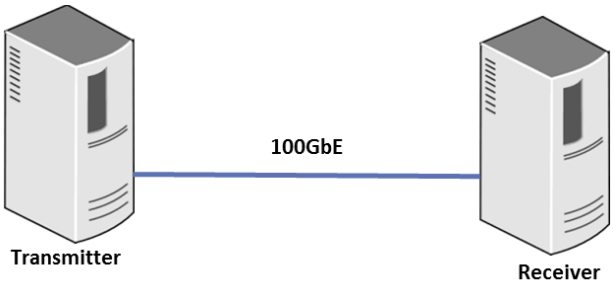

Figure 4 – Test Setup Topology

Network Configuration

The test configuration consists of 2 machines, Transmitter and Receiver, connected back-to-back using a T6 and ConnectX-4 adapters on both ends. The Receiver is configured with 1 Intel Xeon CPU E5-1660 v2 6-core processor clocked at 3.70GHz (HT enabled) and 64GB of RAM. The Transmitter is using same processor with 32GB of RAM. Both machines are using RHEL 7.3 operating system. Standard MTU of 1500B is used.

Mellanox driver: 4.0-1.0.1.0, Mellanox DPDK driver: 16.11_1.5 was installed on both the machines. Latest Chelsio DPDK software from http://service.chelsio.com/ was installed on both the machines.

Commands Used

The following commands are executed on both Transmitter and Receiver systems for Unidirectional packet rate performance:

Chelsio: pktgen -c fff -n 4 -- -T -P -m "[1-3:4-5/7-11].0" -N
Mellanox: pktgen -c fff -n 4 -- -T -P -m "[2-10].0" -N

Conclusion

This benchmark provides DPDK performance results for T6 and ConnectX-4 100GbE adapters. The results show Tx and Rx packet processing capacity for both the adapters. T6 clearly demonstrates exceptional packet processing performance upto 75 MPPS, especially for the for small size packets (64B and 128B). Such record-breaking performance and complimentary support for Crypto, OVS Offload, Traffic QoS, packet classification and filtering functions make T6's DPDK solution best in class, and an ideal fit for networking applications using DPDK and striving for the best in packet processing performance.

Related Links

Chelsio DPDK Solution Overview
Linux 40GbE DPDK Performance
The Chelsio Terminator 6 ASIC
The Chelsio Product Selector - 1/10/25/40/50/100GbE Unified Wire Adapters