100G Network Performance for AMD EPYC

Using AMD EPYC 7551 Platform & Chelsio T6 Adapter

Executive Summary

AMD EPYC, industry's first hardware-embedded x86 server security solution, is a system on chip (SoC) which provides exceptional processing power coupled with high-end memory and I/O resources to meet workload demands of any scale, from virtualized infrastructures to cloud-era datacenters. The combination of the AMD EPYC 7551 server with Chelsio's industry-leading Unified Wire adapter solution delivers compelling performance, power and total cost of ownership (TCO) advantages. This enables innovative topologies and networked computing models to address the most demanding processing needs.

Chelsio adapters provide 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). Designed for industry-leading performance, efficiency, and with the unique ability to fully offload TCP/IP, iSCSI and iWARP protocols using a single ASIC and firmware, Chelsio adapters unburden communication responsibilities and processing overhead from servers and storage systems by enabling a true converged Unified Wire solution resulting in a dramatic increase in application performance with a minimum of CPU cycles. This paper shows Chelsio 100G Network adapter T62100-CR delivering line-rate 99 Gbps throughput for both Transmit and Receive directions in an AMD EPYC Server environment.

Test Results

The following graphs plot the Tx and Rx performance results varying the number of connections. The results are obtained using Iperf tool across I/O sizes ranging from 64 bytes to 512 Kbytes.

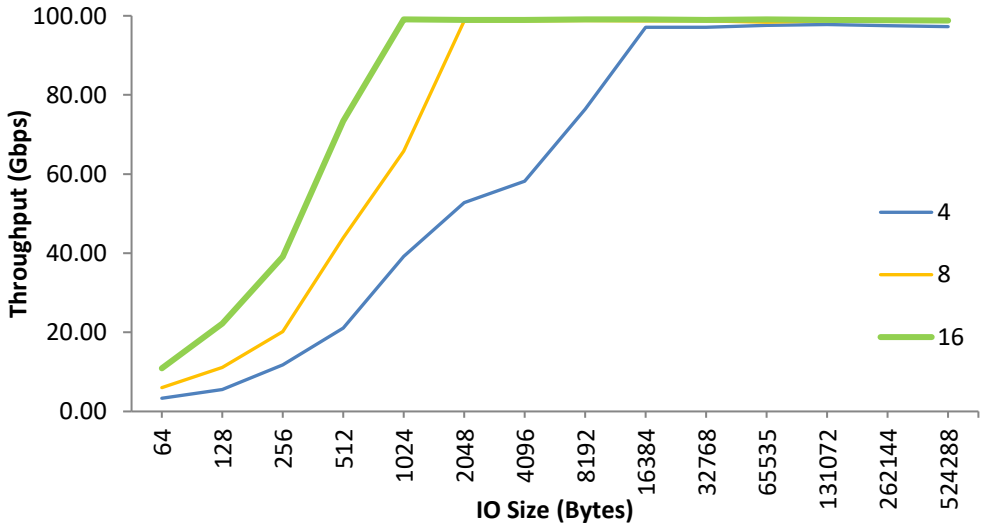

Figure 1 – Tx Performance vs. I/O size

Chelsio adapter delivers line-rate throughput of 99 Gbps even with multiple connections.

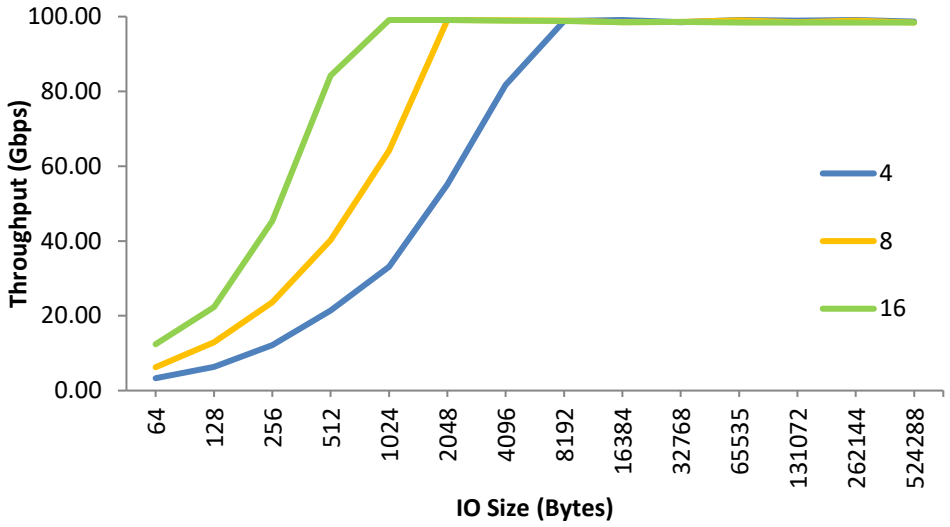

Figure 2 – Rx Performance vs. I/O size

Rx Performance graph reflects similar results, with line-rate throughput of 99 Gbps.

Test Setup

The setup consists of an AMD EPYC machine connected to a PEER machine with a single 100G port. MTU of 9000B was used. Latest Chelsio Unified Wire drivers for Linux was installed on both machines.

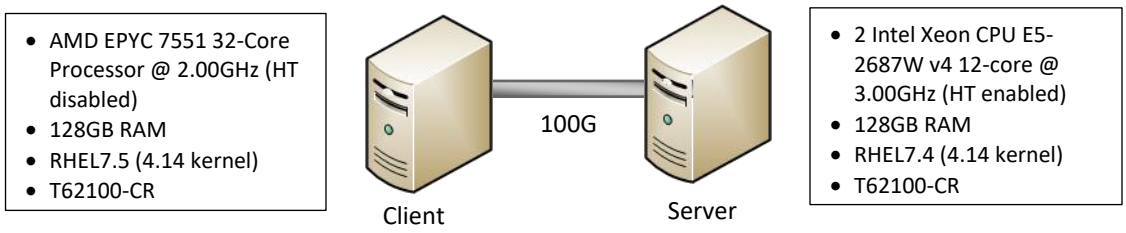

Figure 3 – Test Setup

Setup Configuration

Following performance tunings were done on AMD EPYC server:

i. Updated BIOS to latest version (v1.1b was used in this case).
ii. BIOS Settings:
SVM (Virtualization), Global C-state Control, Hyperthreading, IOMMU, SR-IOV and Core Performance Boost were *Disabled.*
Determinism Slider was set to *Performance*
Memory Interleaving was set to *Auto*
iii. All the memory channels were populated with maximum supported speed.
iv. Added *'iommu=pt cpuidle.off=1 processor.max_cstate=0'* to the kernel command line to disable c-states.
v. Following services were stopped.

```
[root@host~]# systemctl stop firewalld.service
[root@host~]# systemctl stop irqbalance.service
```

vi. Following power saving profiles were set.

```
[root@host~]# tuned-adm profile network-throughput
[root@host~]# cpupower frequency-set --governor performance
```

vii. Following Sysctls were set:

```
[root@host~]# sysctl -w net.ipv4.tcp_timestamps=0
[root@host~]# sysctl -w net.core.netdev_max_backlog=250000
[root@host~]# sysctl -w net.core.rmem_max=1048576
[root@host~]# sysctl -w net.core.wmem_max=1048576
[root@host~]# sysctl -w net.core.rmem_default=1048576
[root@host~]# sysctl -w net.core.wmem_default=1048576
[root@host~]# sysctl -w net.ipv4.tcp_rmem="4096      1048576 1048576"
[root@host~]# sysctl -w net.ipv4.tcp_wmem="4096      1048576 1048576"
```

viii. Installed Chelsio Unified Wire v3.9.0.0.

```
[root@host~]# make install
```

ix. Chelsio interface was assigned with IPv4 address, MTU 9000 and brought-up.

```
[root@host~]# modprobe cxgb4
[root@host~]# ifconfig ethX <IP address> mtu 9000 up
```

x. Enabled adaptive-rx for Chelsio interface.

```
[root@host~]# ethtool -C ethX adaptive-rx on
```

xi. Mapped the Chelsio Interface IRQ's to different CPU cores.

```
[root@host~]# t4_perftune.sh -n -Q nic
```

Commands Used

Server: *iperf -s -p <port>*
Client: *iperf -c <Server IP> -p <port> -l <IO Size> -t 30 -P <# Conn>*

Conclusion

This paper provided the performance results for Chelsio 100G Network adapter in AMD EPYC server environment. With line-rate throughput of 99 Gbps, Chelsio Unified Wire Ethernet adapters provide the optimal and scalable networking building block for datacenters with AMD EPYC servers.

Related Links

100G NVMe over Fabrics for AMD EPYC
100G iSCSI Performance for AMD EPYC
100G OVS Kernel Datapath Offload for AMD EPYC
FreeBSD 100G TOE Performance for AMD EPYC