

Linux 10GbE NIC/iSCSI Performance

Chelsio T520-SO-CR vs. Intel X520 Latency and CPU Efficiency

Executive Summary

This paper presents NIC and iSCSI performance results comparing Chelsio's T520-SO-CR server adapter and Intel's Niantic series X520 adapter running at 10Gbps. The results show that Chelsio's adapter delivers lower latency over a range of I/O sizes. In addition, Chelsio's adapter in NIC mode shows improved CPU efficiency, freeing up processing resources for other memory intensive applications.

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, enabling offload storage, compute and networking traffic to run simultaneously. T5 provides extensive support for stateless offload operation for both IPv4 and IPv6 (IP, TCP and UDP checksums, Large Send Offload, Large Receive Offload, Receive Side Steering/Load Balancing, and flexible line rate Filtering). Furthermore, 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 Fibre Channel storage adapters and InfiniBand RDMA adapters. Nonetheless, they also excel at normal server adapter functionality, providing high packet processing rate, high throughput and low latency for common network applications.

This paper pits the cost effective T520-SO-CR against the standard Intel 10Gbps server adapter and compares the two in NIC and iSCSI performance.

NIC Test Results

The following graphs compare the single port unidirectional and bidirectional throughput numbers and CPU usage per Gbps, for the two adapters obtained by varying the I/O sizes using the **iperf** tool.

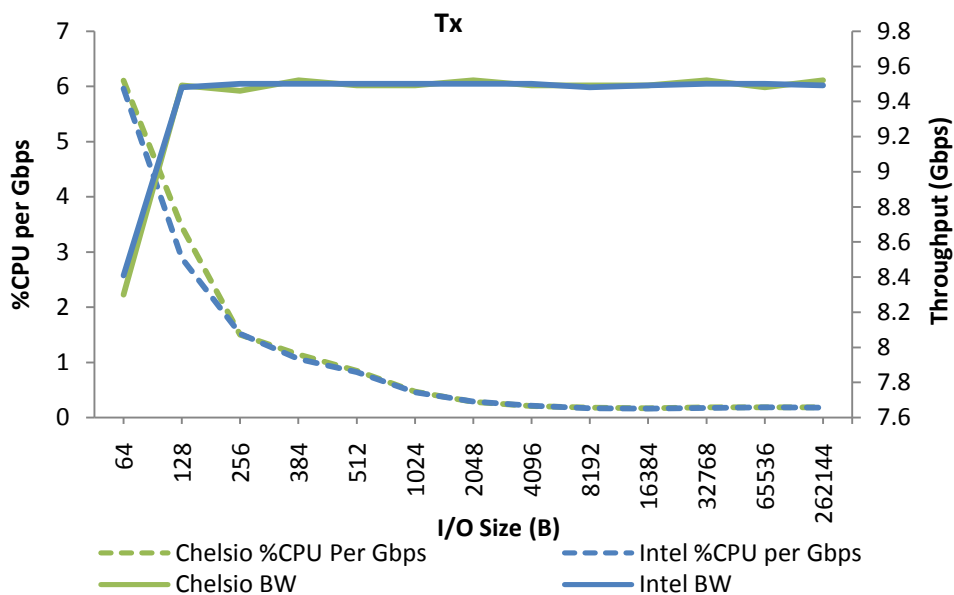


Figure 1 – Tx Throughput and %CPU/Gbps vs. I/O size

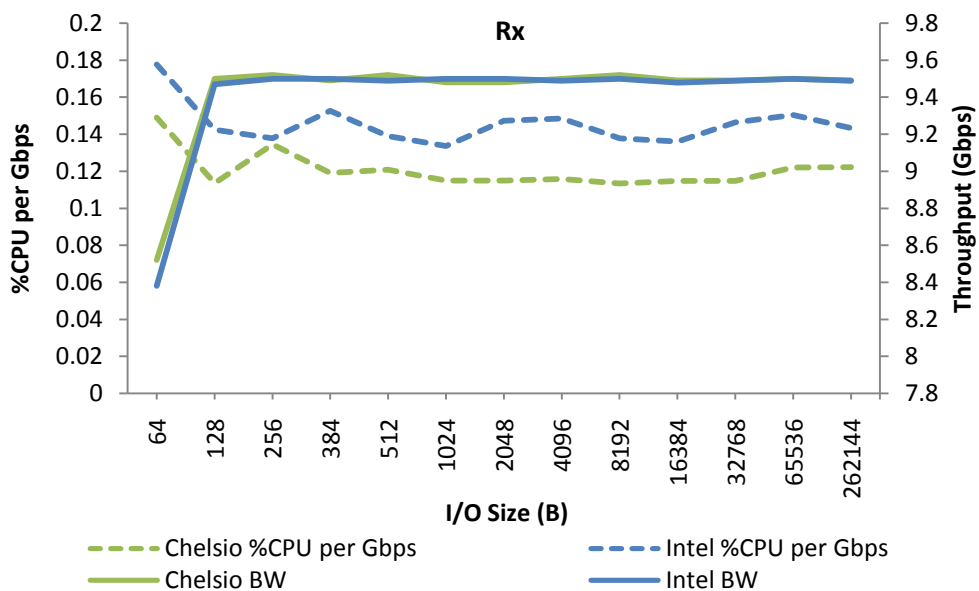


Figure 2 – Rx Throughput and %CPU/Gbps vs. I/O size




Figure 3 – Bidirectional Throughput and %CPU/Gbps vs. I/O size

The following graph compares the latency of the two adapters.

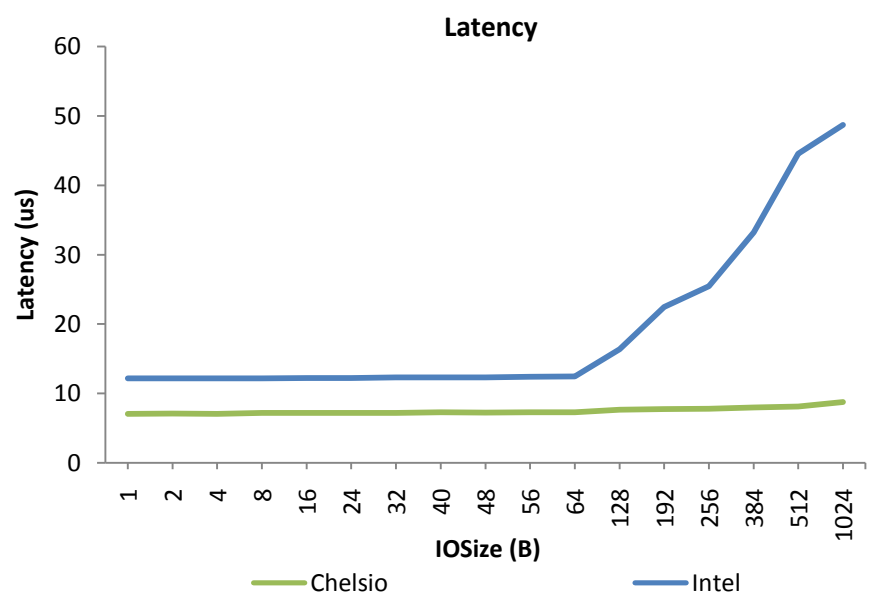


Figure 4 – Latency vs. I/O size

The results clearly show Chelsio’s advantage in latency, with a superior profile that remains flat across the range of study, whereas Intel’s latency is both significantly higher and increases sharply with I/O size.

iSCSI Test Results

The following graphs compare the single port READ, WRITE and READ/WRITE IOPS and Throughput numbers for the two adapters obtained by varying the I/O sizes using the **iometer** tool.




Figure 5 – READ,WRITE and READ/WRITE IOPS vs. I/O size

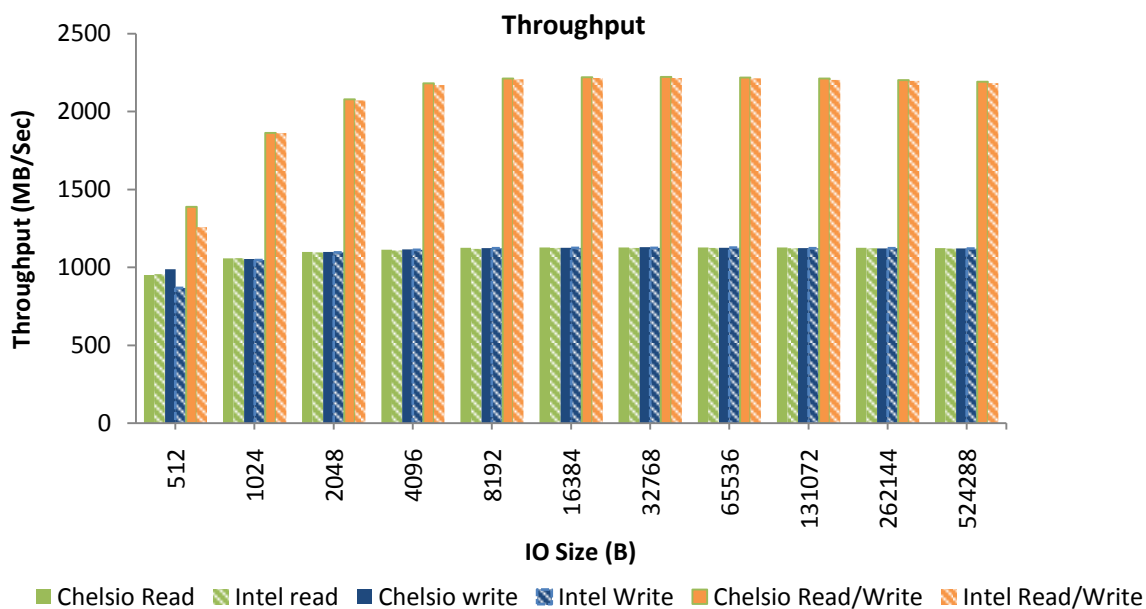


Figure 6 – READ,WRITE and READ/WRITE Throughput vs. I/O size

Test Configuration

The following sections provide the test setup and configuration details.

NIC Topology




Figure 7 – Simple Back-to-Back Test Topology

Network Configuration

The NIC setup consisted of 2 machines connected back-to-back using single port: a Server and Client, each setup with Intel Xeon CPU E5-2687W v2 8-core processors running at 3.40GHz (HT enabled) and 128 GB of RAM. RHEL 6.5 (3.6.11 Kernel) operating system was installed on both machines. Standard MTU of 1500B was used.

The Chelsio setup used T520-SO-CR adapter in each system with Chelsio Network driver v2.10.1.0 whereas the Intel setup used X520 adapter in each system with inbox driver.

iSCSI Topology




Figure 8 – iSCSI Target Connected to 4 Initiators Using a 10Gb Switch

Storage Topology and Configuration

The iSCSI setup consisted of a target storage array connected to 4 iSCSI initiator machines through a 10Gb switch using single port on each system. Standard MTU of 1500B was used.

- **The storage array** was configured with two Intel Xeon CPU E5-2687W v2 8-core processors running at 3.40GHz (HT enabled) and 64 GB of RAM. Chelsio's iSCSI target driver v5.2.0-0813 was installed with RHEL 6.5 (3.6.11 Kernel) operating system.

The Chelsio setup used Chelsio T520-SO-CR adapter while the Intel setup used X520 adapter in the target system.

- **The initiator machines** were each setup with an Intel Xeon CPU E5-2687W v2 8-core processors running at 3.40GHz (HT enabled) and 128 GB of RAM. Chelsio T520-LL-CR adapter was installed in each system with Windows MS Initiator, Chelsio Network driver v5.3.14.0 and Windows 2012 R2 operating system.

The storage array contains 32 iSCSI *ramdisk null-rw* targets. Each of the 4 initiators connects to 8 targets.

I/O Benchmarking Configuration

iometer was used to assess the storage capacity of a configuration. The I/O sizes used varied from 512B to 512KB with an I/O access pattern of random READs, WRITEs and READ/WRITEs.

Iperf was used to measure network throughput. This test used sample IO sizes varying from 64B to 256KB.

Netperf was used to measure the network latency. This test used sample IO sizes varying from 1B to 1KB.

Parameters passed to iometer

- `dynamo.exe -l remote_iometer_ip -m localmachine ip //Add it for all initiators.`
- 30 outstanding IO per Target.
- 16 worker threads.

Commands Used

For all the tests, *adaptive-rx* was enabled on all Chelsio interfaces using the following command:

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

Additionally, the following system wide settings were made:

```
[root@host]# sysctl -w net.ipv4.tcp_timestamps=0
[root@host]# sysctl -w net.ipv4.tcp_sack=0
[root@host]# sysctl -w net.ipv4.tcp_low_latency=1
[root@host]# sysctl -w net.ipv4.tcp_window_scaling=0
[root@host]# sysctl -w net.ipv4.tcp_dsack=0
[root@host]# sysctl -w net.ipv4.tcp_tw_reuse=1
[root@host]# sysctl -w net.ipv4.tcp_tw_recycle=1
```

```
[root@host]# sysctl -w net.core.netdev_max_backlog=250000
[root@host]# sysctl -w net.core.rmem_max=16777216
[root@host]# sysctl -w net.core.wmem_max=16777216
[root@host]# sysctl -w net.core.rmem_default=16777216
[root@host]# sysctl -w net.core.wmem_default=16777216
[root@host]# sysctl -w net.core.optmem_max=16777216
[root@host]# sysctl -w net.ipv4.tcp_rmem='4096 87380 16777216'
[root@host]# sysctl -w net.ipv4.tcp_wmem='4096 65536 16777216'
```

Throughput test:

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
```

Latency test:

On the Server:

```
[root@host]# netserver
```

On the Client:

```
[root@host]# netperf -H <server IP> -t TCP_RR -l 30 -- -r <IO Size>,<IO Size>
```

Conclusion

This paper compared performance results of Chelsio's T520-SO-CR server adapter and Intel's X520 adapter in Linux. Chelsio's T5 delivers:

- Latency low compared to Intel adapter making it the ideal choice for low latency applications.
- Better CPU utilization than Intel adapter in NIC mode, resulting in savings that can be made available for user applications.

Related Links

[The Chelsio Terminator 5 ASIC](#)

[iSCSI over 40Gb Ethernet](#)

[STAC-N1 Benchmark For UDP Traffic](#)

[Packet Rate Performance Report](#)

[10Gb TOE vs NIC Performance](#)