

Linux NIC and iSCSI Performance over 10GbE

Chelsio T520-LL-CR vs. Intel Fortville XL710

Executive Summary

This paper presents NIC and iSCSI performance results comparing Chelsio's T520-LL-CR and Intel's latest XL710 "Fortville" server adapter running at 10Gbps. The results demonstrate that Chelsio's NIC provides drastically lower latency.

In iSCSI tests, the Chelsio's adapter provides significantly higher throughput and IOPs with outstanding small I/O performance and noticeable CPU savings. The T520-LL-CR notably delivers line rate 10Gbps at the I/O sizes that are more representative of actual application use. These advantages make the T520-LL-CR the best performing unified 10Gbps Ethernet adapter in the market.

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. Unlike other converged Ethernet adapters, the Chelsio T5 based NICs 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 T520-LL-CR against the latest Intel 10Gbps server adapter. The following sections start by comparing the two in server adapter (NIC) latency benchmarks, followed by iSCSI storage performance, comparing the full offload iSCSI support of the Chelsio adapter to the Intel adapter.

NIC Test Results

The following graph compares the single port latency of the two adapters, obtained by varying the I/O sizes using the **netperf** tool.

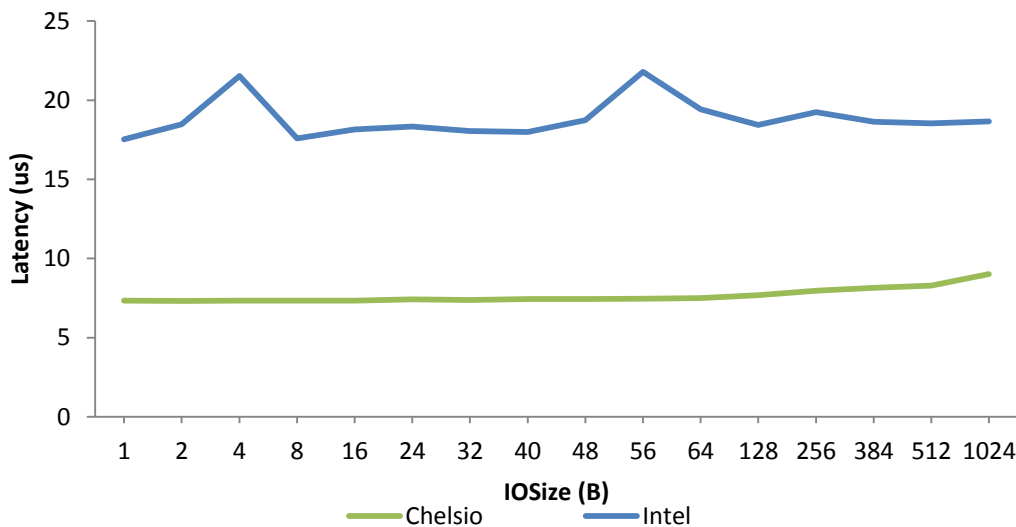


Figure 1 – Latency vs. I/O size

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 significantly higher, placing it outside of the range of interest for low latency applications.

iSCSI Test Results

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

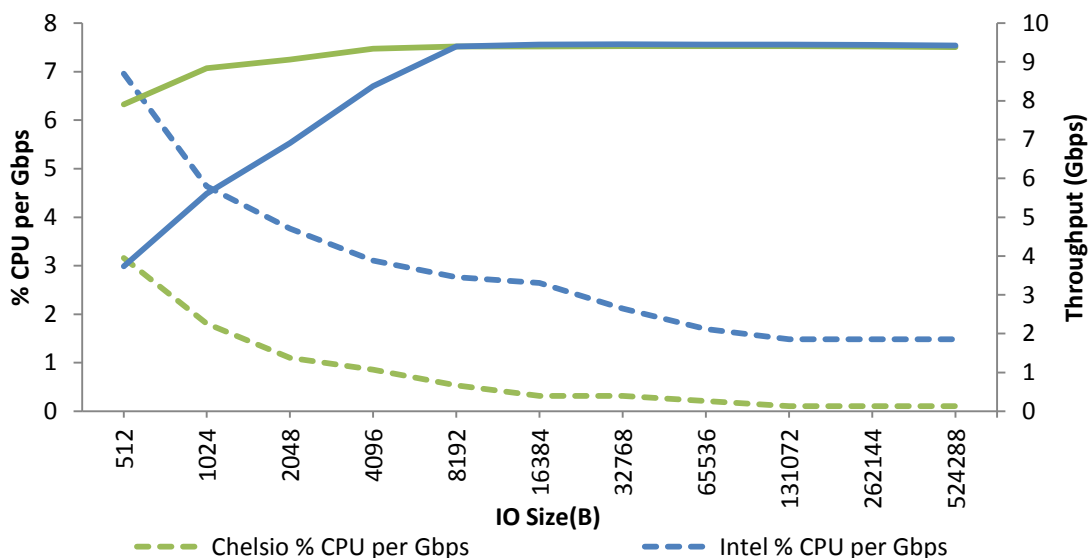


Figure 2 – READ Throughput and %CPU/Gbps vs. I/O size

The READ graph above shows Chelsio’s T520-LL-CR performance to be superior in both CPU utilization and throughput, reaching line rate at ¼ the I/O size needed for Intel’s Fortville XL 710, and consuming far fewer CPU cycles per Gbps.

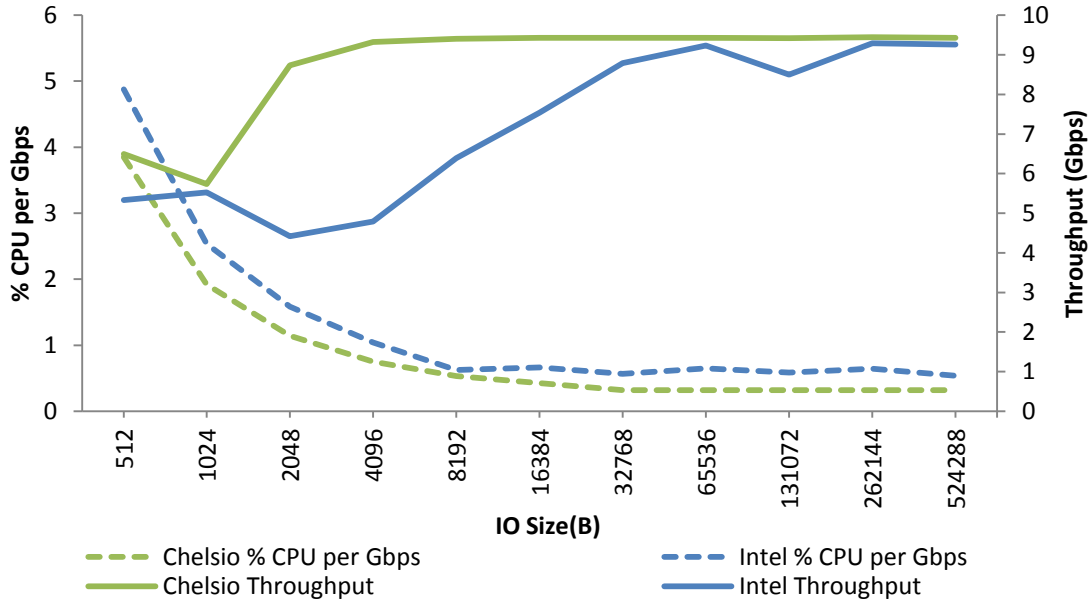


Figure 3 – WRITE Throughput and %CPU/Gbps vs. I/O size

The WRITE results above clearly show that Chelsio’s adapter provides more efficient and consistent performance than Intel’s, which struggles to reach line rate even at large I/O sizes.

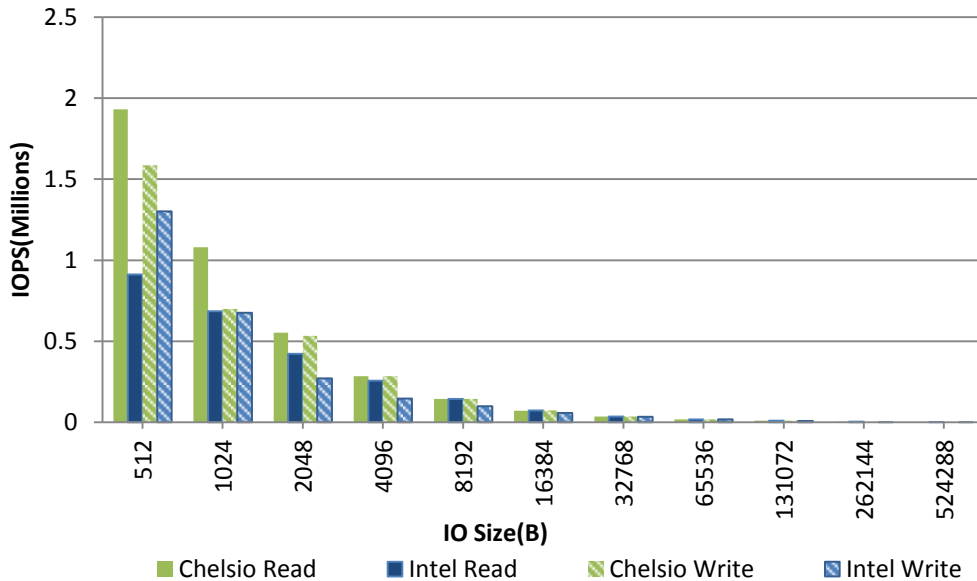


Figure 4 – READ and WRITE IOPs vs. I/O size


The IOPS numbers reflect the performance advantages of Chelsio’s adapter, particularly at the challenging small I/O sizes that are more representative of actual application requirements.

Test Configuration

The following sections provide the test setup and configuration details.

NIC Topology

*Server running
 RHEL 6.5 (3.6.11)*



*Client running
 RHEL 6.5 (3.6.11)*

Figure 5 – Simple Back-to-Back Test Topology

Network Configuration

The NIC setup consists of 2 machines connected back-to-back using single port on each system: a Server and Client, each configured with Intel Xeon CPU E5-2687W v2 8-core processor 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 a T520-LL-CR adapter in each system with Chelsio network driver v2.11.0.0 whereas the Intel setup used a XL710 adapter in each system with Intel network driver v1.0.15.

iSCSI Topology

*iSCSI Initiators
 with T520-LL-CR
 adapters running
 on Windows
 Server 2012 R2*




Figure 6 – 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 2 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 v5.2.0-0850 was installed with RHEL 6.5 (3.6.11 Kernel) operating system.

The Chelsio setup was configured in ULP mode with CRC enabled using T520-LL-CR adapter. The Intel setup was configured in AUTO mode with CRC enabled using XL710 adapter.

- **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. T520-LL-CR adapter was installed in each system with Windows MS Initiator, Unified Wire Software v5.0.0.33 and Windows 2012 R2 operating system.

The storage array contains 8 iSCSI *ramdisk null-rw* targets. Each of the 4 initiators connects to 2 targets.

I/O Benchmarking Configuration

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

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 and WRITES.

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

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>
```

Conclusions

This paper compared performance results of Chelsio's T520-LL-CR and Intel's XL710 server adapter in Linux. The benchmark results demonstrate that Chelsio's T5-based adapter delivers:

- Drastically lower latency compared to the Intel adapter making it the ideal choice for low latency applications.

- Higher performance and higher efficiency networking using the iSCSI protocol, thanks to iSCSI offload in hardware.

The results thus show that the T5-based adapters provide a unique combination of a complete suite of networking and storage protocols with high performance and high efficiency operation, making them great all-around unified wire adapters.

Related Links

[The Chelsio Terminator 5 ASIC](#)

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

[High Frequency Trading Report](#)

[Packet Rate Performance Report](#)