High-Performance NIC Benchmarking on Linux with Chelsio T7

Delivering Scalable 400 GbE Connectivity and CPU Efficiency on Linux

Executive Summary

The Chelsio T7 adapters deliver exceptional networking performance on Linux, achieving near line-rate throughput with minimal CPU utilization. Designed for high-speed data center and enterprise workloads, the T7 leverages advanced offload capabilities and efficient PCIe Gen5 architecture to accelerate network I/O and reduce system bottlenecks. With robust Linux driver support and optimized TCP/IP performance, it ensures consistent low latency, scalability, and superior efficiency for modern applications demanding 400 GbE connectivity.

Test Results

Note: These results are preliminary, and further improvements are currently in progress.

The graph compares Chelsio NIC throughput for 32 and 64 connections at MTU 1500 for both TX and RX, showing that throughput scales consistently as connections increase and quickly reaches near-line-rate performance **~378 Gbps**. Overall, the adapter is able to reach line-rate or close-to-line-rate performance with a single peer machine, confirming its exceptional efficiency across varying I/O sizes.

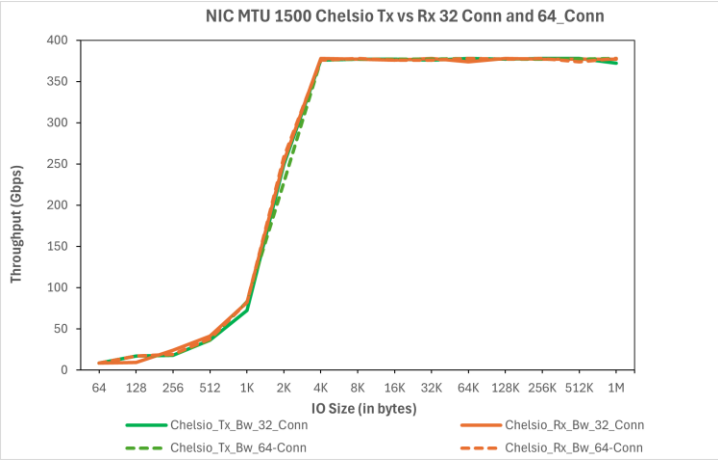

Figure 1 – Chelsio NIC MTU 1500: Tx and Rx Performance Scaling Across 32 and 64 Connections

The below graph compares Chelsio's bidirectional (BIDI) throughput performance for 32 and 64 connections at MTU 1500, where both configurations scale smoothly as load increases and quickly

reach near–line-rate performance of **732 Gbps** for both configurations. Overall, both 32 and 64-connection tests deliver high BIDI throughput.

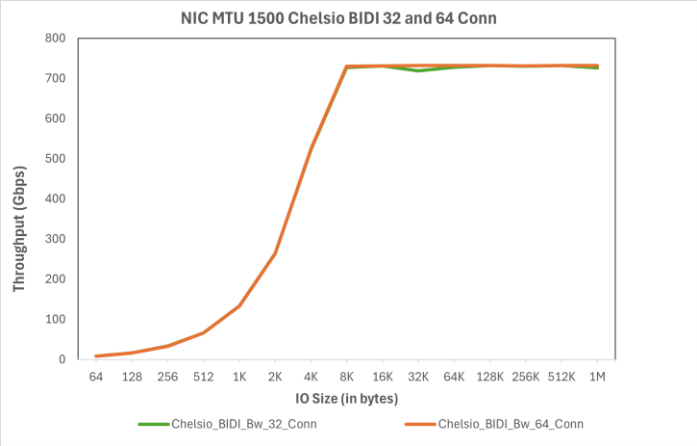

Figure 2 – Chelsio NIC MTU 1500: Bidirectional Performance Scaling Across 32 and 64 Connections

Test Configuration

The setup consists of a server connected to a client, both configured with the latest Chelsio Unified Wire drivers for Linux, which can be downloaded from the Chelsio Download Center. A Chelsio S71400 adapter is installed in the server, and another S71400 adapter is installed in the client. Port 0 on both adapters is configured, and both the server and client are set to the standard MTU of 1500 bytes.

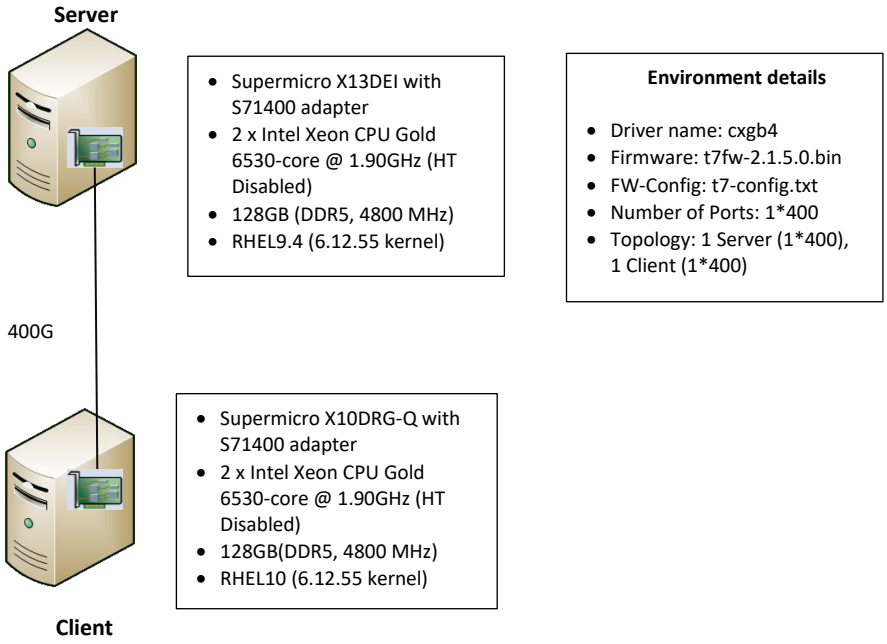

Figure 3 – Test Setup

*These performance numbers represent application-layer payload throughput and are slightly lower than the physical wire rate due to protocol and framing overhead on the link.

Setup Configuration

Server

BIOS settings:

- Disable Hyper-Threading Technology.
- Disable Sub-NUMA Clustering (SNC).
- Disable CPU Power Management Features including C-states, P-states, and Turbo Boost
- Configure System power management profile in Maximum Performance Mode.

i. Install RHEL 9.4 and compile 6.12.55 kernel.

ii. Set CPU Affinity as below:

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

iii. Tune profile in the TuneD system tuning service:

```
[root@host ~]# tuned-adm profile network-throughput
```

iv. Run the below command:

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

Clients

BIOS Settings:

- Same as server BIOS options.
- Set CPU Affinity as below:

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

i. Install RHEL 10.0 and compile 6.12.55 kernel.

ii. Tune profile in the TuneD system tuning service:

```
[root@host ~]# tuned-adm profile network-throughput
```

iii. Run the below command:

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

Commands Used

1. Start iperf3 Server on Server Machine.

   Run the following command on each client machine to start the iperf3 server on both interfaces:

   ```
   [root@client ~]# numactl -N0 -m0 iperf3 -s -D -p 6001
   ```

   The numactl binding ensures the iperf3 process runs on the same NUMA node as its network interface, optimizing performance.

2. Run tests from the Client.

   The Client initiates all tests.

   Assume that the DUT interface is attached to NUMA node 0 (Example: S71400 adapter in CPU1 PCIe slot).

   2.1. Transmit (TX) Test: To measure transmit performance from the DUT to clients,

   ```
   [root@server ~]# numactl -N0 -m0 iperf3 -c 101.1.1.15 -Z --skip-
   rx-copy -N -t 30 -P 32 -l 256K -p 6001 -O 10
   ```

2.2. Receive (RX) Test: To measure receive performance from clients to the DUT,

```
[root@server ~]# numactl -N0 -m0 iperf3 -c 101.1.1.15 -Z --skip-
rx-copy -N -t 30 -P 32 -l 256K -p 6001 -O 10 -R
```

2.3. Bi-directional (BiDi) Test: To measure simultaneous transmit and receive performance,

```
[root@server ~]# numactl -N0 -m0 iperf3 -c 101.1.1.15 -Z --skip-
rx-copy -N -t 30 -P 32 -l 256K -p 6001 -O 10 --bidir
```

Conclusion

The performance evaluation results clearly show that the Chelsio T7 adapter demonstrates strong, consistent networking performance on Linux, achieving near line-rate Tx/Rx throughput of **~378 Gbps** and bidirectional (BIDI) throughput of about **732 Gbps**. Performance remains stable as connection counts scale from 32 to 64, with both unidirectional and bidirectional tests showing nearly identical results across configurations. Overall, the T7 adapter provides a scalable and reliable solution for next-generation high-performance **400GbE** data-center deployments.

Related Links

400G kTLS/SSL Offload & Encryption with FreeBSD
Chelsio T7 NIC Performance Leadership Across Tx, Rx, and BiDi Workloads
High-Performance NIC Optimization on Linux With Chelsio T7
AI Networking Solution: Chelsio T7 DPU and S7/T6 SmartNICs
T7 Product Brief
AI Networking: The Role of DPUs
Offload Protocols with Inline IPsec demonstration on T7 Emulation Platform
iSCSI JBOF with T7
NVMe/TCP and iSCSI JBOF with T7