Chelsio T7 Performance Analysis on Linux for 400 GbE Networking in a Multi-Client Setup

Benchmarking 400 GbE Throughput and Scalability

Executive Summary

The Chelsio T7 adapters deliver exceptional networking performance on Linux, consistently achieving top-tier throughput. Engineered for high-speed data-center and enterprise environments, the T7 utilizes advanced offloads and a high-bandwidth PCIe Gen5 architecture to accelerate network I/O and minimize bottlenecks. With strong Linux driver support and optimized TCP/IP processing, it offers consistent low latency, excellent scalability, and high efficiency for modern applications requiring 400 GbE connectivity.

Test Results

The below graph shows that the Chelsio T7 adapter consistently delivers near line-rate throughput as IO size increases, with Tx and Rx performance efficiently scaling and reaching ~378 Gbps per direction. This demonstrates that a **single 400 Gbps port is capable of delivering performance close to line rate**, even as connection counts grow from 48 to 96. Overall, the Chelsio T7 provides stable, high-throughput performance regardless of workload intensity or the number of active connections.

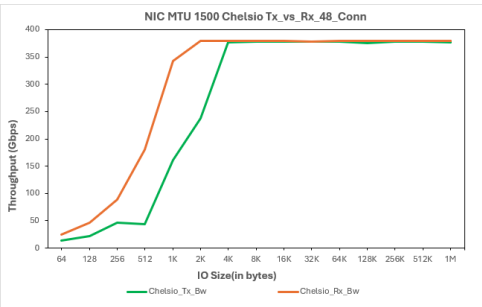

Figure 1 – NIC MTU 1500 Chelsio Tx and Rx with 48 Connections

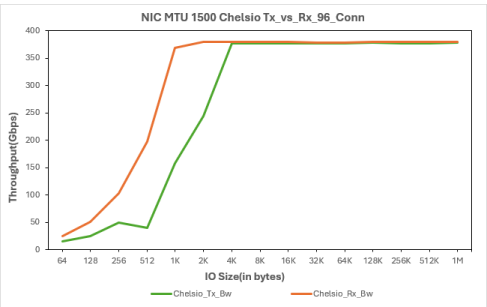

Figure 2 – NIC MTU 1500 Chelsio Tx and Rx with 96 Connections

The below graph shows the Chelsio T7 adapter sustains high throughput—reaching up to **~744 Gbps**—across both 48 and 96-connection configurations as IO sizes increase, with throughput scaling rapidly and remaining consistently high at larger IO sizes. Overall, the MTU 1500 BIDI results highlight the Chelsio T7's ability to sustain high, predictable bidirectional throughput across varying connection counts.

Copyright 2025. Chelsio Communications Inc. All rights reserved.

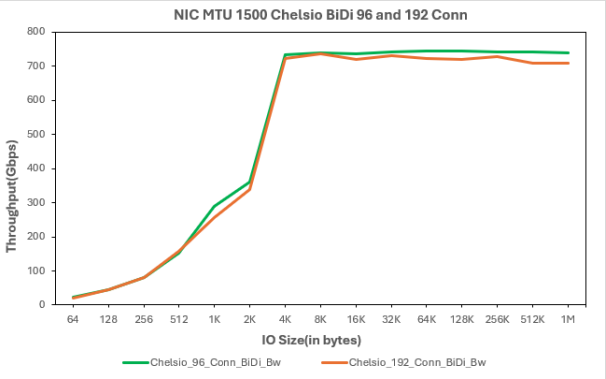

Figure 3 – NIC MTU 1500 Chelsio bidirectional with 96 and 192 Connections

Test Configuration

The setup consists of a server connected to three clients through a switch. The server is configured with the latest Chelsio Unified Wire drivers for Linux and equipped with a Chelsio S71400 adapter using a standard MTU of 1500 bytes. On the client side, ConnectX-7 MT2910 (MCX755106AC-HEAT), Chelsio T72200-FH (PT43250020), and BCM57608 (BCM957608-P2200GQF00) adapters are installed in the three client systems respectively, with Port 0 on each adapter configured with the standard MTU of 1500 bytes.

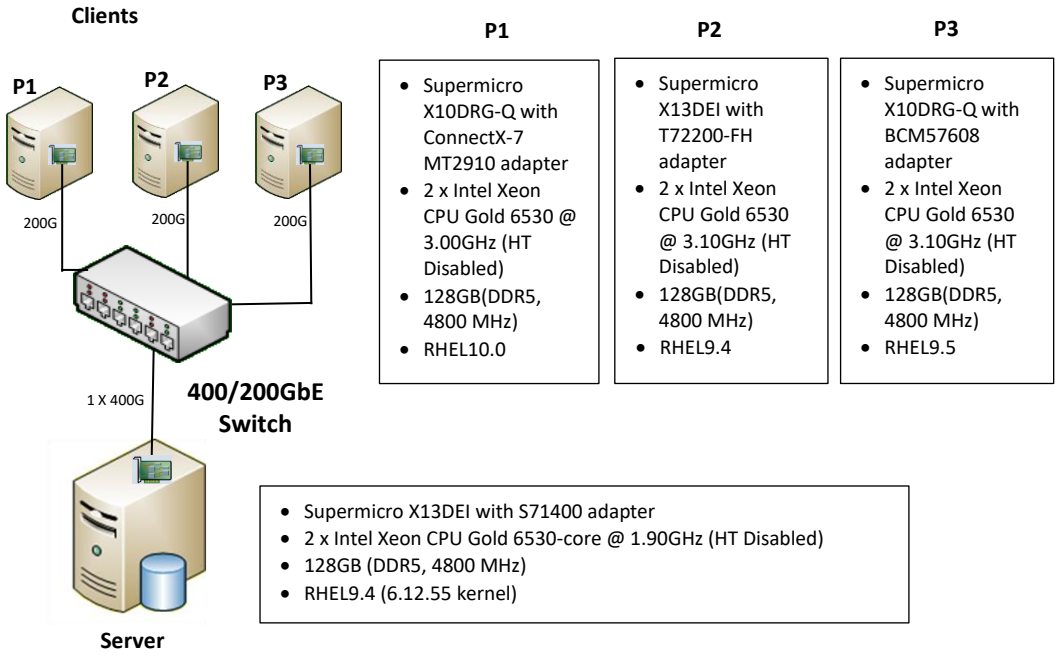

Figure 4 – Test Setup

Copyright 2025. Chelsio Communications Inc. All rights reserved.

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. Tune profile in the TuneD system tuning service:

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

iii. Set CPU Affinity as below:

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

Clients

BIOS Settings:

- Same as server BIOS options.

i. Install RHEL 10.0, RHEL 9.4, and RHEL 9.5 on the three client systems respectively.

ii. Tune profile in the TuneD system tuning service:

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

iii. Set IRQ Affinity as: Port mapped to CPU 0-31

Commands Used

1. Start iperf3 Server on Client Machines.

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

```
[root@client ~]# numa_node=$(cat
/sys/class/net/<intf_$X>/device/numa_node);
   numactl -N ${numa_node} -m ${numa_node} iperf3 -s -D -p 6001
done
```

   Notes:
   Replace "<intf_$X>" with the actual network interface name (Example: ens1f0).
   The numactl binding ensures the iperf3 process runs on the same NUMA node as its network interface, optimizing performance.

2. Run tests from the Server.

   Notes:
   The Server initiates all tests.
   Assume that the Server interface is attached to NUMA node 0 (Example: S71400 adapter in CPU1 PCIe slot).
   Client interfaces have IPs assigned as "<Client${client}_IP>".

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

```
[root@server ~]# for client in 1 2 3; do
       numactl -N 0 -m 0 iperf3 -c <Client${client}_IP> -Z --skip-
rx-copy -N -t 30 -i 5 -p 6001 -l <IO_Size> -P <Num_of_Connections>
-f g -O 10 &
done
wait
```

Copyright 2025. Chelsio Communications Inc. All rights reserved.

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

```
[root@server ~]# for client in 1 2 3; do
       numactl -N 0 -m 0 iperf3 -c <Client${client}_IP> -Z --skip-
rx-copy -N -t 30 -i 5 -p 6001 -l <IO_Size> -P <Num_of_Connections>
-f g -O 10 -R &
done
wait
```

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

```
[root@server ~]# for client in 1 2 3; do
       numactl -N 0 -m 0 iperf3 -c <Client${client}_IP> -Z --skip-
rx-copy -N -t 30 -i 5 -p 6001 -l <IO_Size> -P <Num_of_Connections>
-f g -O 10 --bidir &
done
wait
```

Conclusion

The Chelsio T7 adapters consistently deliver exceptional 400 GbE networking performance on Linux, providing high throughput across diverse workloads and connection scales. The results show stable Tx/Rx performance of **~378 Gbps** per direction and bidirectional throughput scaling up to **~744 Gbps** with MTU 1500, even as connection counts increase from 48 to 96. This sustained performance, combined with low latency and strong scalability, highlights the effectiveness of T7's advanced offload architecture and optimized Linux driver stack. Overall, the T7 delivers reliable, high-throughput, making it an ideal solution for modern data-center environments demanding predictable, high-performance 400 GbE connectivity.

Related Links

High-Performance NIC Benchmarking on Linux with Chelsio T7
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

Copyright 2025. Chelsio Communications Inc. All rights reserved.