

100G iSCSI Offload Performance

Using White Box JBOF Storage Platform & Chelsio T6 Adapter

Executive Summary

Chelsio Unified Wire Engine offers storage protocol offload capability for accelerating both block (iSCSI, FCoE) and file (SMB, NFS, Object) level storage traffic. Chelsio’s iSCSI Offload solution runs at 100Gb and beyond, and will scale consistently with Ethernet evolution. Furthermore, iSCSI support is part of a complete, fully virtualized unified wire offload suite that includes FCoE, RDMA over Ethernet, TCP and UDP sockets and user space I/O. This paper demonstrates the capabilities of Chelsio’s iSCSI offload solution in a White Box JBOF Storage Platform setup with a Microsemi PCIe Switch. Chelsio’s iSCSI offload solution delivered line-rate throughput of 98 Gbps and 2M IOPS.

Test Results

The following graph presents READ, WRITE IOPS and throughput results of Chelsio’s T6 iSCSI solution using ramdisk (nullio) as storage array, for offloaded and non-offloaded traffic. The results are collected using **fiio** tool with I/O size varying from 4k to 256k bytes, with an access pattern of random READs and WRITES.

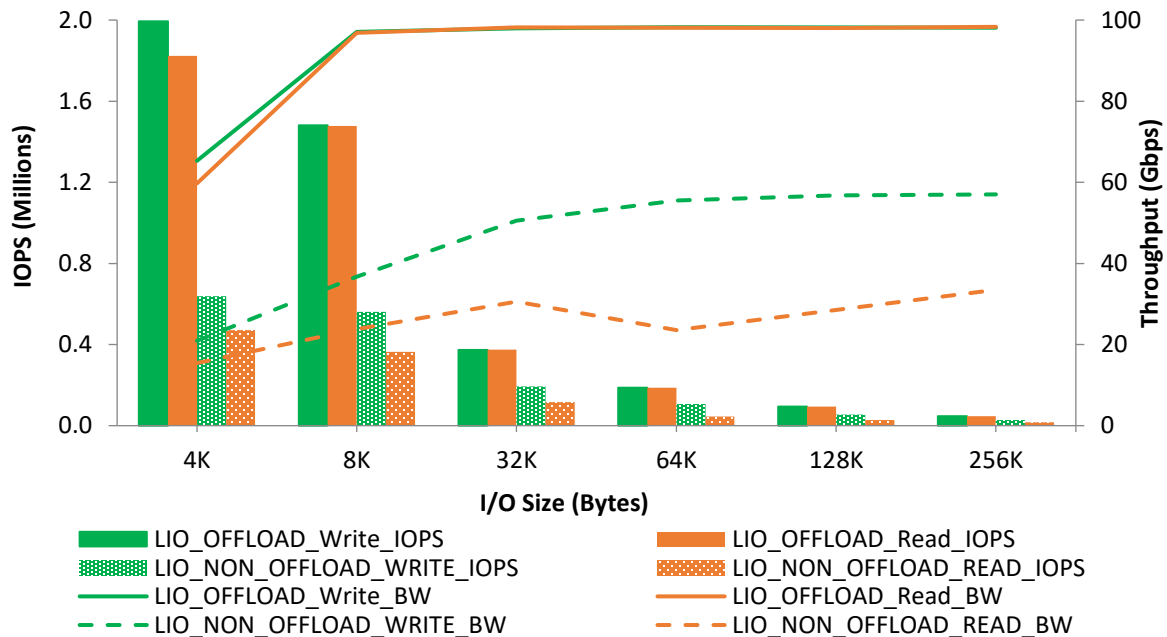


Figure 1 – ramdisk (nullio) READ, WRITE Throughput & IOPS vs. I/O size

T6 iSCSI solution delivers up to 98 Gbps line-rate throughput for offloaded traffic, for both READ and WRITE operations. Furthermore, it delivers a WRITE IOPS of approximately 2M at 4K I/O size for offloaded traffic. The iSCSI offload solution delivers upto 3x the throughput and IOPS of non-offload case, indicative of an efficient processing path.

The following graph presents READ, WRITE IOPS and throughput results of Chelsio iSCSI solution using SSD as storage array, for both offloaded and non-offloaded traffic. Please note that WRITE throughput and IOPS numbers are limited by SSD performance.




Figure 2 - SSD READ, WRITE Throughput & IOPS vs. I/O size

Chelsio’s T6 iSCSI solution delivers READ Throughput of 95.65 Gbps, and READ and WRITE IOPS of 1M for offloaded traffic. Similar to ramdisk (nullio) case, the iSCSI Offload solution exhibits distinct advantage compared with non-offload case. Further performance tuning is in progress.

Test Setup

Topology

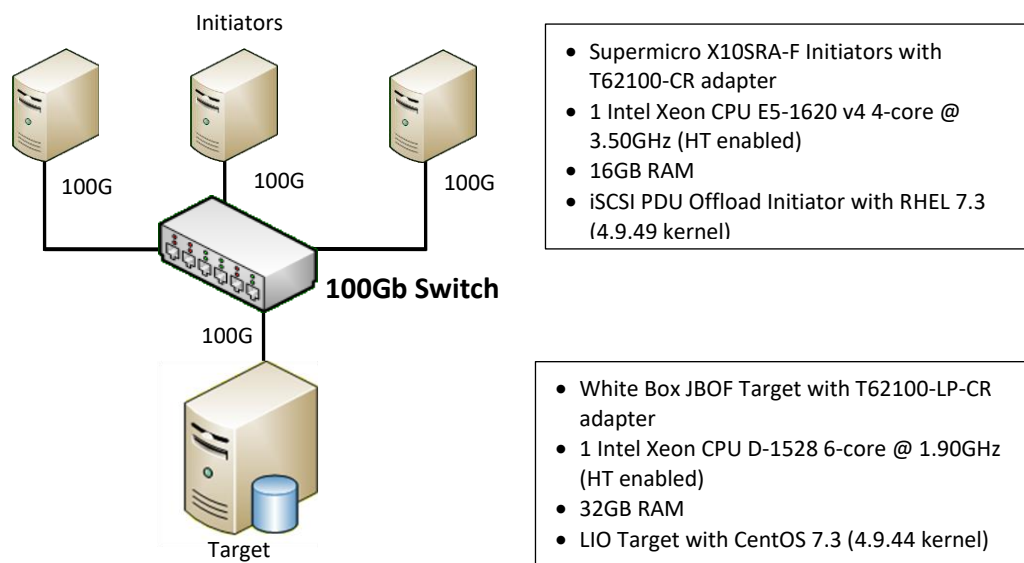


Figure 3 – Test Setup

The setup consists of a target machine connected to 3 initiator machines through a 100GbE switch using single port on each system. MTU of 9000B is used. Chelsio Unified Wire v3.6.0.3 is installed on each machine.

Storage configuration

For the ramdisk (nullio) setup, 24 targets are configured, each with 1 ramdisk (nullio) of 600MB size. In SSD scenario, the target is configured with 24 LUNs from 8 Samsung NVMe PCIe MZWLK800HCHP 800GB SSDs, each of 1GB size using Logical Volume Management.

3 Initiators are used and each of them connects to 8 targets.

Commands used

Ramdisk (nullio):

```
[root@host~]# fio --name=random --iodepth=64 --rw=randwrite/randread --size=400m --invalidate=1 --direct=1 --numjobs=6 --bs=<blocksize> --runtime=30 --time_based --ioengine=libaio --fsync_on_close=1 --group_reporting --filename=/dev/sdb:...:/dev/sdi
```

SSD:

```
[root@host~]# fio --name=random --iodepth=64 --rw=randwrite/randread --size=900m --invalidate=1 --direct=1 --numjobs=6 --bs=<blocksize> --runtime=30 --time_based --ioengine=libaio --fsync_on_close=1 --group_reporting --filename=/dev/sdb:...:/dev/sdi
```

Conclusion

This paper presented Chelsio T6 iSCSI solution's performance in a White Box JBOF Storage Platform setup using a T62100-LP-CR Unified Wire adapter and Microsemi PCIe Switch. Numbers were collected for offloaded and non-offloaded traffic, for both READ and WRITE operations. Using NVMe SSD and ramdisk (nullio) as storage, the results show that Chelsio's iSCSI Offload solution delivers:

- Line-rate throughput of up to 98 Gbps for both READ and WRITE using ramdisk (nullio) and 95.65 Gbps READ performance using SSDs.
- IOPS of ~2M using ramdisk (nullio) and more than 1M using SSDs.
- Improved bandwidth and IOPs compared to non-offload making it an ideal fit for applications looking for boost in performance or using small I/O sizes.

Related Links

[Demartek Evaluation: Chelsio Terminator 6 \(T6\) Unified Wire Adapter iSCSI Offload](#)

[Dell EMC Storage Center & Chelsio](#)

[100G iSCSI Performance for AMD EPYC](#)

[High performance 100G iSCSI Storage Solution](#)

[High Performance iSCSI at 100GbE](#)

[100G iSCSI – A Bright Future for Ethernet Storage](#)