

ESXi iSCSI Performance at 40Gbps

Comprehensive Storage Virtualization with Terminator 5

iSCSI, the Internet SCSI standard, is the leading Ethernet SAN protocol, with native initiator support integrated into all the major operating systems and built-in RDMA for high performance offloaded implementations. Storage virtualization using iSCSI is rapidly becoming the most preferred method of storage management in virtualized IT infrastructures as they offer simplicity, scalability and resilience; coupled with lower operational costs given iSCSI does not have to have dedicated hardware support and can indeed run on any server platform.

The Terminator 5 (T5) from Chelsio Communications, Inc. is a fifth generation, hyper-virtualized, high-performance unified wire engine with full iSCSI offload support, among other networking, clustering and storage protocols. By leveraging Chelsio’s proven TCP Offload Engine (TOE), offloaded iSCSI over T5 enjoys a distinct performance advantage over regular NIC, as well as superior data integrity protection. iSCSI preserves existing equipment without requiring a fabric overhaul, additional acquisition and management costs, with a high performance option thanks to hardware offload, and a strong roadmap to 100Gbps and beyond.

Test Results

The following graph plots the throughput results obtained from a virtual machine at different I/O sizes, using the **fiio** tool, with an access pattern of random READs and WRITEs.

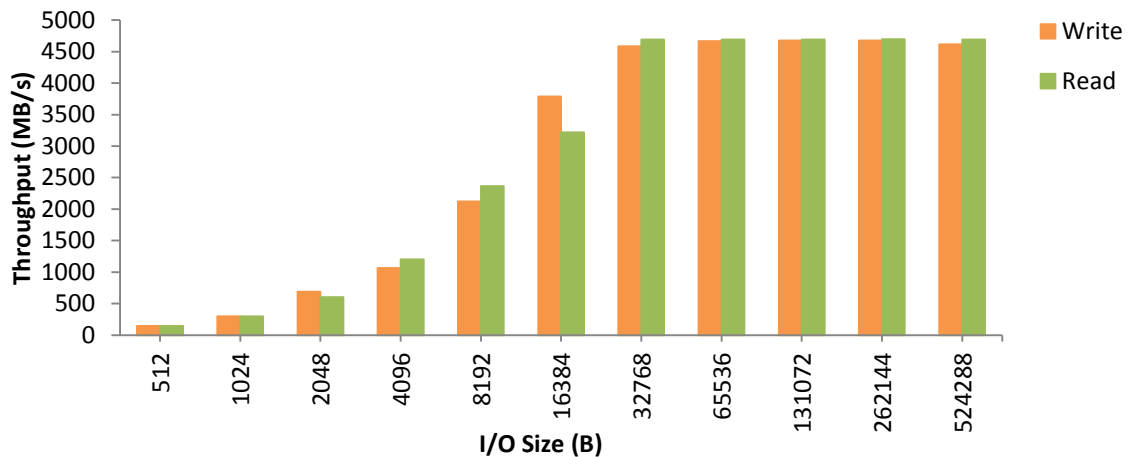


Figure 1 – READ and WRITE Throughput vs. I/O size

As clear from the graph above, Chelsio’s adapter delivers consistent throughput, reaching line rate at 32KB I/O size for both READ and WRITE. The virtual machines can experience large improvements in storage I/O performance using the Chelsio’s hardware offloaded iSCSI solution.

Test Configuration

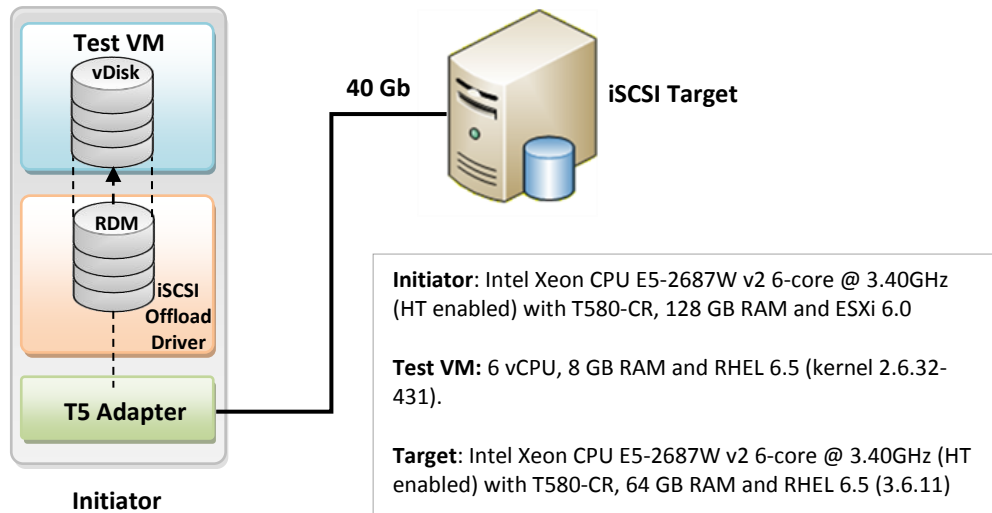


Figure 2 – Test Setup

Storage Topology and Configuration

The setup consists of a target storage array connected back-to-back to an initiator using a single 40Gbps link, with MTU of 9000B. Chelsio Unified Wire driver package v2.11.1.0 is installed on the target machine, whereas on the initiator machine, Chelsio iSCSI Initiator driver v1.1.x.x is installed. The initiator connects to the target having 4 *ramdisk* block devices (LUNs) each of 1GB size. A Test VM is configured on the initiator machine. Para-Virtualized SCSI controller is used for attaching the discovered LUNs to the Test VM as Raw Device Mapping (RDM) disks, which appear as local disks to the VM.

Command Used

```
[root@host~]# fio --name=rand<read/write> --iodepth=16 --rw=rand<read/write> --
size=800m --direct=1 --invalidate=1 --fsync_on_close=1 --norandommap --
group_reporting --numjobs=2 --bs=<io_size> --runtime=30 --time_based --
filename=<disk_to_use>
```

Conclusion

This paper provided 40GbE iSCSI performance results for Chelsio's T580-CR Unified Wire adapter with VMWare's ESXi hypervisor. The adapter performs consistently under load and achieves line rate throughput for both READ and WRITE at 32KB I/O size. Chelsio T5's iSCSI solution provides an all-round SAN solution for virtualized environments with exceptional I/O performance and efficiency.

Related Links

[The Chelsio Terminator 5 ASIC](#)
[iSCSI Heritage and Future](#)
[iSCSI at 40Gbps](#)