

Linux OpenStack iSCSI at 10GbE

Throughput Benchmark Results

Executive Summary

OpenStack is an open source cloud computing software platform that provides tools to build and manage public or private clouds. Completely scalable and robust, OpenStack can cater to the cloud requirement of any scale: whether it is large enterprises looking for a global cloud solution or small companies wanting to install a private cloud. iSCSI, the Internet SCSI standard, is the leading Ethernet SAN protocol with native support integrated into OpenStack.

This paper presents 10GbE benchmark results of Chelsio's iSCSI solution integrated in OpenStack using the latest Terminator 5 ASIC. With a superior throughput curve, the results provide a preview of the benefits of Chelsio's iSCSI solution.

Overview

The Terminator 5 (T5) ASIC from Chelsio Communications, Inc. is a fifth generation, high-performance 2x40Gbps/4x10Gbps, unified wire engine which offers PDU iSCSI offload capability in protocol acceleration for both file and block-level storage (iSCSI) traffic. 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.

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. In fact, T5 includes enhanced data integrity protection for all protocols, and particularly so for storage traffic, including full end-to-end T10-DIF support for both iSCSI and FCoE, as well as internal datapath CRC, and ECC-protected memory. 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 provides benchmark results that demonstrate T5's performance within the OpenStack virtualized environment using the T520-LL-CR 2x10Gbps server adapter.

Test Results

The following graph shows the READ and WRITE throughput numbers, obtained by varying the I/O sizes using the **fiio** tool.

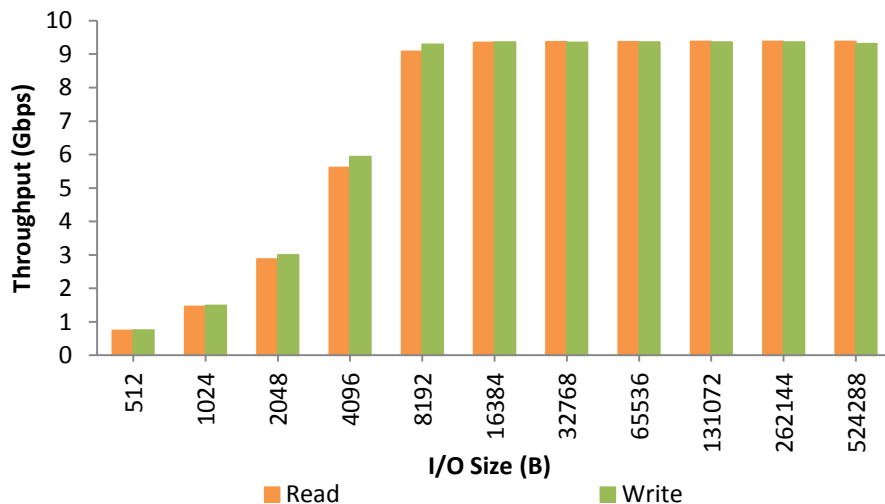


Figure 1 – READ and WRITE Throughput vs. I/O size

The above results reveal that Chelsio’s adapter achieves superior throughput, reaching line rate at I/O size as small as 8KB for both READ and WRITE.

Test Configuration

The following sections provide the test setup and configuration details.

Topology

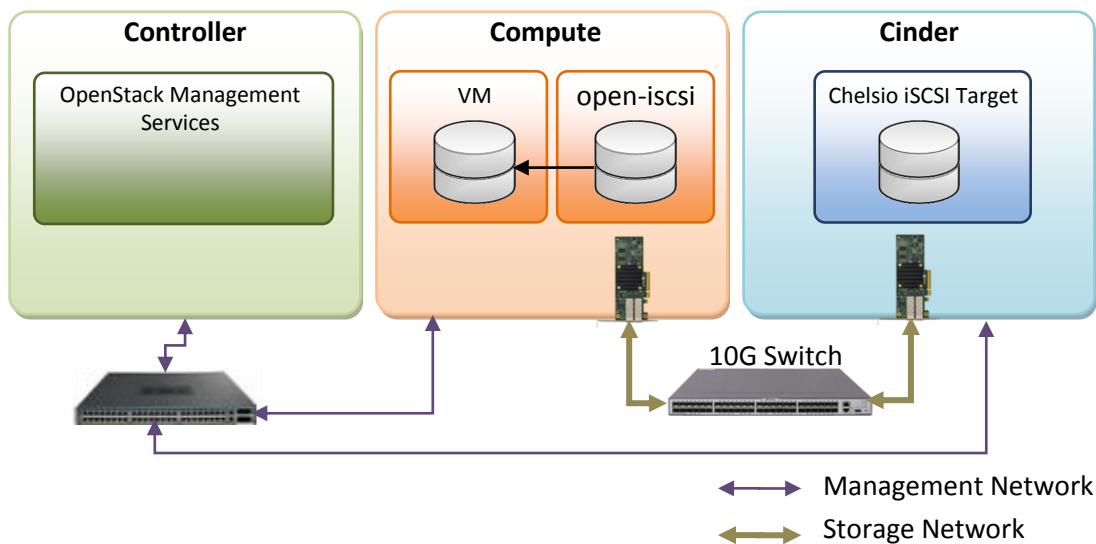


Figure 2 –Test Setup

Storage Topology and Configuration

The setup consists of 3 identical machines: a Compute node, Cinder node and Controller, each configured with Intel Xeon CPU E5-1660 v2 processor clocked at 3.70GHz (HT enabled), 64GB of RAM, CentOS operating system and OpenStack snapshot Kilo-3. 1 T520-LL-CR adapter is installed on each of the Compute and Cinder node with Chelsio Unified Wire driver package v2.11.0.0. Standard MTU of 1500B is used.

Using the Controller, a VM instance is created with custom flavor having 2 vCPUs, 4GB of RAM and 10GB disk. CentOS-6-x86_64-GenericCloud raw image is used for operating system. Additionally, 2 volumes are created and attached to the instance. On the Cinder node, *LVMVolumeDriver* is used as the volume driver with *chelsio iscsi_helper* (iscsictl). A ramdisk is used as the backing storage for the LVM volume group. The Cinder node volumes are attached to the Compute node instances using open-iscsi Initiator.

I/O Benchmarking Configuration

fiio is used to assess the storage capacity of the configuration. The I/O sizes used varied from 512B to 512KB with an I/O access pattern of random READs and WRITEs.

Command Used

On the VM instance:

```
[root@host]# fio --rw=randread --bs=<IO Size> --numjobs=4 --iodepth=32 --  
runtime=20 --time_based --size=8388608k --ioengine=libaio --direct=1 --  
invalidate=1 --fsync_on_close=1 --norandommap --group_reporting --exitall --  
name disk1 --filename=/dev/vdb
```

Conclusions

This paper showcases the significant performance benefits of Chelsio T5 iSCSI solution for the OpenStack platform. The benchmark results demonstrate that Chelsio's T520-LL-CR adapter delivers line rate throughput and performs consistently under heavy load, as shown by steady READ and WRITE numbers starting at I/O sizes as low as 8K.

Related Links

[The Chelsio Terminator 5 ASIC](#)

[iSCSI at 40Gbps](#)

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

[High Performance iSCSI for Virtual Machines](#)

[iSCSI Heritage and Future](#)