

NVMe over iWARP Fabrics Performance

T6 Based High Performance and Low Latency NVMe-oF Solution

Executive Summary

This paper demonstrates an efficient approach to access remote storage, made possible by iWARP, which enables the next generation scalable storage network over standard and cost effective Ethernet infrastructure.

Test Results

The following table presents the performance results using the Samsung PM1725 NVMe 1.6 TB SSDs over T6 100G iWARP fabric.

Metric		Local	Remote
Bandwidth @ 128KB	Read	5319 MB/s	5247 MB/s
	Write	1857 MB/s	1854 MB/s
Latency @ 4KB	Read	81.64 μ s	93.00 μ s
	Write	14.74 μ s	24.45 μ s
IOPs @ 4KB	Read	927609	919538
	Write	328376	309825

The results indicate that:

- Remote NVMe SSD IOPS and Bandwidth are same as local NVMe SSDs.
- Remote versus local NVMe device access adds less than 12 μ s latency.

Commands Used

Bandwidth: `fio --name=random --rw=rand<write/read> --size=400m --invalidate=1 --direct=1 --runtime=30 --time_based --ioengine=libaio --fsync_on_close=1 --group_reporting --filename=/dev/nvme0n1 --iodepth=32 --numjobs=12 --bs=128K`

Latency: `taskset -c 1 fio --name=random --rw=rand<read/write> --size=400m --invalidate=1 --direct=1 --time_based --ioengine=libaio --fsync_on_close=1 --group_reporting --filename=/dev/nvme0n1 --iodepth=1 --numjobs=1 --bs=4k --runtime=30`

IOPS: `fio --name=random --rw=rand<write/read> --size=400m --invalidate=1 --direct=1 --runtime=30 --time_based --ioengine=libaio --fsync_on_close=1 --group_reporting --filename=/dev/nvme0n1 --iodepth=32 --numjobs=12 --bs=4K`

Related Links

[Introducing NVMe over 100GbE iWARP Fabrics](#)

[High Performance NVMe over 40GbE iWARP](#)

[T6 100G NVMe-oF demonstration](#)

[Concurrent Support of NVMe-oF and Established Networked Block and File Storage](#)