



## NVMe over TCP Rack-Scale Shared Flash

### Rack-Scale Storage

Most cloud data centers today rely on scaleout architectures that use direct-attached storage (DAS). While DAS based architectures are simple, deploying compute & storage in fixed ratios is an inefficient model for scaling cloud infrastructure. Historically, low cost hard disk drives enabled cheap overprovisioning of storage, which made this DAS approach popular. However, as data centers are transitioning to flash, the higher cost and dynamic nature of the SSD market require an architecture that emphasizes independent scaling of compute and storage, while using standard servers and SSDs.

Lightbits provides a rack-scale shared flash model where compute and storage are scaled independently to achieve maximum resource utilization and drive down TCO. Lightbits enables this disaggregation of compute and storage without compromising performance and without requiring changes to compute clients or networking infrastructure. Key attributes of the Lightbits solution include:

- Physically separating compute servers and their storage, enabling independent deployment, scalability, and upgrades for compute and storage
- Maximizing resource utilization
- Allowing for independent life-cycle management of compute and storage hardware, simplifying operations and driving down TCO
- Reducing or eliminating shuffling of data among nodes that can cause network congestion

### NVMe Over TCP


NVMe has emerged as the state of the art protocol for high performance SSDs. The NVMe over Fabrics (NVMeoF) specification extended NVMe to share PCIe SSDs over the network, with the initial implementations using an RDMA fabric.

Lightbits is collaborating with Facebook, Intel and other industry leaders to extend the NVMeoF standard to also support TCP/IP transport, which is complementary to RDMA fabrics. Disaggregation with NVMe-over-TCP has the distinct advantages of being simple and highly efficient. TCP/IP is ubiquitous, scalable, reliable, and ideal for short-lived connections. Also, migrating to shared flash storage with NVMe-over-TCP doesn't require any changes to the data center network infrastructure. Additionally, the Lightbits pre-standard NVMe-over-TCP solution, first in the industry, provides high performance and delivers consistently low latency.

## NVMe Over TCP Demo Setup

Flash Memory Summit 2017 - Intel Partners Pavilion

Facebook & Lightbits demo the industry's first pre-standard NVMe over TCP solution



## NVMe-over-TCP Latency Results

Random Read ( $\mu$ s)			Random Write ( $\mu$ s)		
Average	99%	99.9%	Average	99%	99.9%
120	167	212	47	71	95

## Contact Us

Lightbits, together with Facebook, Intel and other industry leaders, are actively working on defining and implementing NVMe-over-TCP as part of the NVMe.org working group. Join us to learn more and contribute.

For more information Email [info@lightbitslabs.com](mailto:info@lightbitslabs.com)