



Flash-enabled Agility

KumoScale™ NVMe over Fabrics™ Shared Accelerated Storage Software

As business moves to private cloud and flash storage becomes the norm, data center architectures must be flexible with an elastic infrastructure able to share valuable compute and flash storage resources efficiently. Data center architectures can utilize NVMe over Fabrics™ (NVMe-oF™) with disaggregated storage and KumoScale software to meet these requirements and help maintain a business advantage. Cloud-based storage is not just about capacity and performance – it's about better efficiency, capacity utilization and storage availability. Toshiba continues to lead the industry by building reliable, high performance and low-latency innovative technologies, like KumoScale, supporting data intensive and performance hungry applications in cloud infrastructures.

Introduction

KumoScale, Toshiba's NVMe-oF shared accelerated storage software, enables data center architects to create more flexible and scalable infrastructures significantly lowering their CapEx and OpEx costs. KumoScale seamlessly integrates with multiple orchestration and provisioning frameworks, enabling dynamic management of compute and storage allocation matched to each application instance with high performance and low latency. In addition, KumoScale generates a complete view of workload and wear state across the networked solid-state drive (SSD) inventory with configurable event notifications and automatic remediation of common problems.

Business Benefits

- Save storage CapEx with better flash capacity utilization
- Reduce stranded compute resources by enabling server SKU consolidation
- Increase revenue agility by instantly adapting to changing workloads and peak demand

Infrastructure Benefits

- Greater IOPS and lower latency for the cloud
- PCIe® NVMe™ DAS-like performance benefits now shared across the network
- Lower per-node cost by sharing flash
- Reduce need for specialized compute-node SKUs
- Enable higher application uptime; live migrate vs reboot VMs/containers when deploying security updates
- Contain cluster growth by enabling stateful applications to be efficiently load balanced

Features

KumoScale software enables agile implementation and scaling of private cloud environments through the following features:

- Abstracts details of physical NVMe SSDs; provisions virtual volumes
- Secure ACL maintains and specifies access rights to NVMe resources
- Works with standard Linux and SPDK host drivers
- Compatible with NVMe-compliant SSDs
- RESTful API – integrates with popular provisioning, orchestration, and telemetry frameworks, including Kubernetes®, Intel® RSD and OpenStack®
- N x 100G Ethernet network ports; full network bandwidth at 4KB random I/O
- <20µS 4K incremental latency (QD=1)
- NVMe Express™ v1.3 and NVMe-oF v1.0 compliant
- University of New Hampshire InterOperability Lab certification

High-Performance Capabilities Across the Network

KumoScale shared accelerated storage technology converts industry-standard NVMe server platforms into ultra-high-performance block storage pools, with capacity and IOPS sharable by thousands of job instances over a data center network. KumoScale software enabled nodes deliver the right amount of capacity and IOPS needed to each application instance, outperforming direct-attached SATA SSDs performance and latency.

Toshiba Labs baseline KumoScale performance for latency, IOPS and throughput are provided in the chart below*.

KumoScale Software-enabled Performance

| Block Size | Δ Latency ^{1,2} (μS) | | IOPS ³ | | Throughput (GB/s) | |
|------------|-------------------------------|-------|-------------------|-------|-------------------|-------|
| | Read | Write | 50/50 | 100%W | 50/50 | 100%W |
| 4 KiB | 13 | 8 | 5.9M | 4.6M | 23 | 19 |
| 128 KiB | 29 | 34 | 218K | 155K | 27 | 19 |

¹ To provide results that are independent of SSD performance, latencies are given as a delta relative to the latency measured on the same drive in a direct-attached configuration

² Queue depth = 1

³ Full duplex

* KumoScale performance was measured using a mid-range, single CPU platform equipped with 24 Toshiba SSDs, and is dependent on the hardware platform and installed SSDs.

Hardware and Deployment Configuration

KumoScale software operates on the latest Linux® kernel release and is currently available on Intel™ x86 and AMD EPYC™ platforms offering single or dual-redundant controllers. KumoScale storage nodes typically include:

- Multiple 16-lane PCIe slots for RNICs
- Typical 2U, 24x2.5" bay implementation

Toshiba recommends Intel® Xeon® CPU E5-2690 v4 @2.30GHz or AMD EPYC™ 7XXX CPU @2.9GHz and 64GiB of DRAM.

Kumoscale software is integrated into a storage node - provisioned, scheduled and managed like compute nodes and scaled to supply the storage needs of one or a few racks. Toshiba's software enables server consolidation, addressing the flash storage needs of a rack of diskless compute nodes with a typical 2U storage server.

Essential KumoScale elements

- A datapath implemented within the Linux kernel
- A management module which runs in user space

Additional questions? Contact your Toshiba representative or visit: KumoScale.com.