



Software that Delivers NVMe™ Flash as a Service

KumoScale™ storage software suite implements a fast, networked block storage service based on the NVMe-oF™ (NVMe™ over Fabrics) standard. Focused on speed, scale and very low cost, KumoScale software delivers the efficiency and flexibility of a shared, clustered storage pool at the performance of NVMe flash.

■ KumoScale Key Features

- **Born in the cloud** – Integrated with Kubernetes®, Ansible®, Graphite™ and other cloud-native infrastructure software
- **Open** – Works with standards-compliant NVMe-oF initiator and NVMe drives.
- **Network Fabrics** – Supports RDMA for highest network performance and TCP/IP for broadest network compatibility
- **Resilient** – Cross-domain volume replication for legacy applications; simple volumes for cloud-native applications
- **Analytics-driven provisioning** – Based on storage classes, QoS and telemetry data
- **Standard platform** – Runs on standard x86-based (including AMD-based) storage platforms
- **Trusted supplier** – Inventor and first to mass production of NAND flash memory.

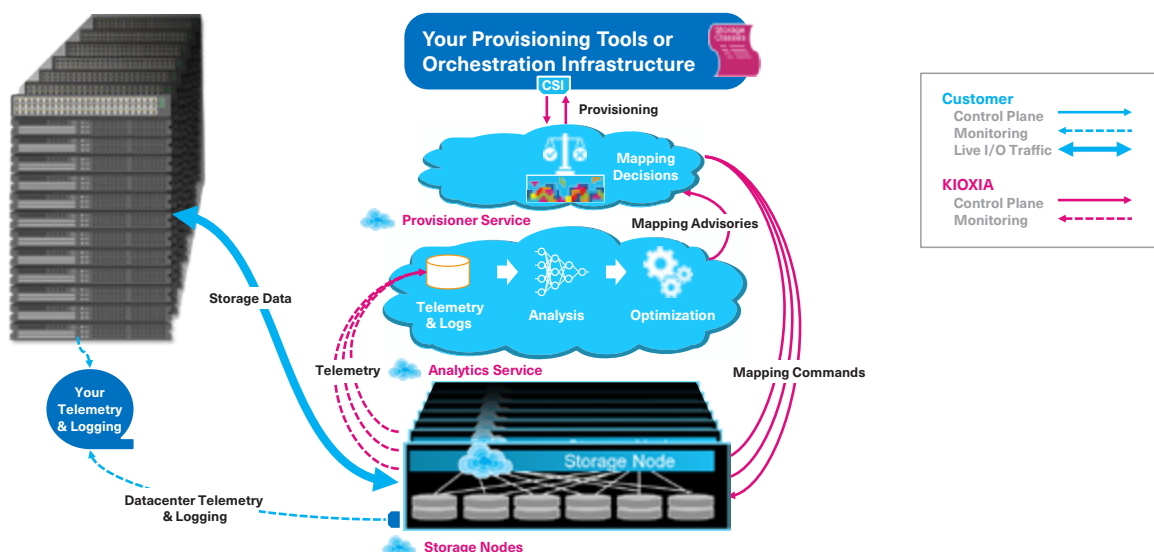
■ Application Benefits

- Customize storage performance for each application based on its specific requirements
- Significantly improves query performance for OLTP and analytics databases
- Provide instant deployment of container-based workloads
- Enable stateful applications to be efficiently load-balanced and elastically scaled

■ Business Benefits

- Save storage CapEx with better flash capacity utilization
- Perform the same work with fewer compute nodes by blending workloads and eliminating stranded resources
- Increase revenue agility by rapidly adapting stateful workloads during demand peaks
- Accelerate decision making through faster ingest and analysis
- Simplify procurement and maintenance by consolidating server SKUs

Data Center System Architecture



■ KumoScale™ Software Suite

The KumoScale software suite delivers performant NVMe™ flash storage as a networked service. It is designed for deployment at scale in medium to large scale on-premise clouds.

Typical applications include:

- Cloud services
- Consumer SaaS services
- Marketplaces, clearinghouses (travel, tickets, stock trades)
- Massively multi-player gaming

■ Share Fast Flash-based SSDs

Capacity and performance of today's solid state disks have exploded and are far more than a single compute node can consume. Sharing SSDs over the network is an urgent need. Until now, networked storage has been either slow or expensive, or both. No longer. KumoScale software delivers the efficiency and flexibility of a shared storage pool at the performance of NVMe SSD.

■ Run Stateful Instances on Any Node

Elastic application scaling is table-stakes in the modern cloud. Per-node workload blending will be soon as well. Both require "run any instance on any node" flexibility. That is fine for web servers, but moving stateful instances under load is painful, slow and error prone, unless the storage is visible everywhere. KumoScale software allows even the most latency-sensitive applications to be scheduled with high flexibility.

■ How It Works

The KumoScale suite consists of multiple components:

- **KumoScale Storage Node** accepts NVMe-oF™ I/O commands directed at virtual volumes and translates them into NVMe commands directed at the physical SSDs installed in the storage node.
- **KumoScale Provisioner** is a distributed, resilient service that tracks the fleet of SSDs and KumoScale storage nodes. It handles the mapping of user volumes to nodes and to physical drives.
- **KumoScale CSI driver** is an implementation of the Container Storage Interface— a standards-based method for connecting storage to orchestrated containers. It provides an API translation between an orchestration system and KumoScale Provisioner for containerized deployments.
- **KumoScale Analytics Service** collects and analyzes telemetry information from the storage nodes for the purpose of optimized data placement.

■ KumoScale Platform Requirements

KumoScale is a software only storage fabric solution that supports both bare metal and Kubernetes® deployments on a wide range of industry standard servers. Storage nodes are typically one or two-socket x86-based servers with 8 to 24 NVMe bays and two 100G Ethernet NICs. It works with standards-compliant NVMe SSD. With the exception of the Kubernetes/CSI driver, KumoScale software does not require any proprietary agent or code resident on the Initiator (client). Any NVMe-oF compliant initiator may be used, including Linux® kernel version 4.8 or later.

■ How It's Deployed

KumoScale software supports both containerized and bare metal environments. A deployment consists of two parts: KumoScale storage nodes and control plane services, which comprise a set of provisioning, analysis, and management processes.

Deploying Storage Nodes:

KumoScale software uses PXE to boot from a centralized server. Deployments can be fully automated by integrating into a datacenter framework, such as Ansible®.

Deploying Control Plane Services:

KumoScale Provisioner, Analytics and Manager are typically deployed on different compute nodes in separate failure zones for resilience.

■ How Storage is Provisioned and Managed

In a Kubernetes environment, customer pods are connected to KumoScale storage by the Container Storage Interface. CSI consists of a driver, which is deployed by Kubernetes on each compute node, and a CSI controller running on the Kubernetes master. For bare-metal environments, Ansible playbooks or other popular automation tools may be used for deployment automation and provisioning of storage to compute nodes. For short-term testing and Proof-of-Concept activities, CLI and GUI interfaces are also available.

■ Monitoring and Alerting

KumoScale software provides telemetry (time series metrics) and logging (asynchronous events) to an existing telemetry and logging infrastructure. These interfaces are "pluggable," so adaptation to new monitoring frameworks is done easily and quickly.

■ About the Network

The NVMe over Fabrics standard offers several network transport layer options. The KumoScale solution runs on Ethernet, supporting RoCE v2, TCP/IP or a combination of both.

■ How it Performs

Two 100Gbps Ethernet channels can carry roughly 23 GB/s full duplex in an uncongested network. This is equivalent to 5.7 million 4kB I/O's. A KumoScale storage node is capable of filling this pipe with 4kB random I/Os over RoCE, with additional read latency of about 15us relative to a read from the same SSD attached to the initiator's PCIe® bus.

■ To Learn More

Contact your KIOXIA representative or visit www.kumoscale.com

© 2020 KIOXIA Corporation. All rights reserved. Information in this frequently asked questions document, including product specifications, tested content, and assessments are current and believed to be accurate as of the date that the document was published, but is subject to change without prior notice. Technical and application information contained here is subject to the most recent applicable KIOXIA product specifications.

NVMe, NVMe-oF and NVM Express are trademarks of NVM Express, Inc. PCIe is a registered trademark of PCI-SIG. Linux is a registered trademark of Linus Torvalds. All other trademarks or registered trademarks are the property of their respective owners.