

KumoScale Software Overview

This section describes the KumoScale software architecture, defines various software components, and outlines the terminology used throughout the document.

What is KumoScale?

KumoScale software is a clustered, scale-out NVMe-oF™ protocol storage system designed for a modern data center. KumoScale software is available in one of two modes both of which host services that provision and manage storage allocation within a data center orchestration framework. This includes the KumoScale Provisioner Service, interfaces, mapping, and monitoring tools for allocating, analyzing, and managing storage.

- **Appliance Mode** software is installed and configured on the **KumoScale Storage Cluster**, a Kubernetes cluster.
- **Managed Mode with Kubernetes** software is installed and configured for your own Kubernetes cluster. This is any healthy, high-availability cluster that is installed and configured per best practices and guidelines of the selected Kubernetes distribution.

KumoScale Components

This section describes the components of KumoScale when used in Appliance mode or in Managed mode with a Kubernetes Cluster.

- **KumoScale Storage Nodes** are storage appliances equipped with NVMe™ Solid State Drives (SSDs) and configured as worker nodes that manage the virtual volumes.
- **KumoScale Control Plane** consists of operators used to create and configure **Custom Resources** used in the storage environment. The control operators are installed as part of the KumoScale Kubernetes storage cluster when using appliance mode and are installed on top of your own Kubernetes storage cluster when using managed mode.
- **KumoScale Provisioner** is a service on the storage cluster that
 - Creates virtual clusters for different tenants.
 - Allocates volumes according to applications requirements, storage nodes status, and utilization.
 - Creates entities referred to as **targets** on KumoScale storage nodes to connect to the application initiators (hosts).
 - Adds the volumes as namespaces.
 - Maintains the connection status for application initiators.

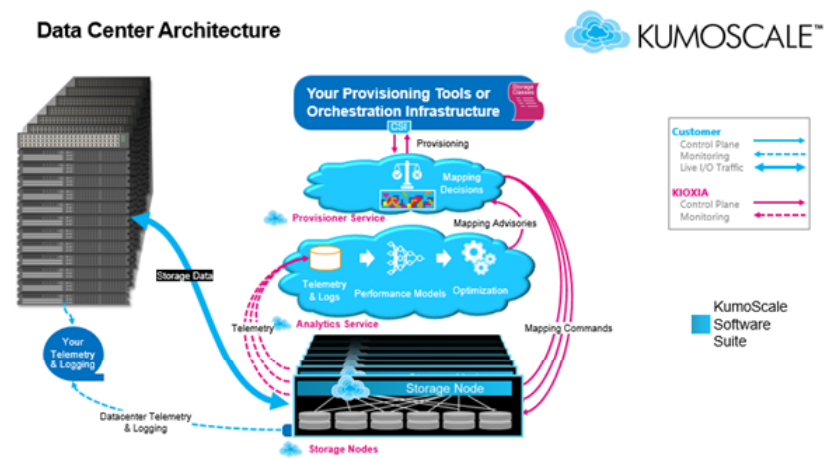
The Provisioner is discussed in more detail in [Deploying a KumoScale Storage Provisioning System](#) and throughout this document.

- **KumoScale Interfaces** are installed on the storage cluster to support integration with your current orchestration environment. The main KumoScale interface is a
 - **RESTful Application Programming Interface (API)** designed for proprietary orchestration environments.

For easier and faster deployment, KumoScale software also supports:

- **KumoScale Cluster Manager Command Line Interface (CLI)** for storage control, monitoring, and user management. The [KumoScale Cluster Manager CLI Guide](#) includes complete information and notes differences between appliance and managed mode.
- **Kubernetes CSI Driver** to interface between KumoScale and Kubernetes containerized environments. The [KumoScale CSI Driver User Guide](#) provides details on how to use this interface.
- **Ansible for Bare-Metal** for provisioning storage to support a cross-domain resiliency solution. The [KumoScale Ansible User Guide](#) provides details on how to use this interface.
- **OpenStack™ Cinder Driver** to integrate with an OpenStack platform via the Cinder driver plug-in and the NVMe-oF connector. Cinder is a block storage service for OpenStack platforms, and designed to present storage resources to end users that can be consumed by the OpenStack Nova Compute Project. The [KumoScale Guide for OpenStack](#) provides details on how to use this interface.
- **KumoScale Time Series Databases (TSDBs)** and **Syslog** servers collect the telemetry data for analysis. KumoScale software analyzes the data to produce recommendations on managing the environment.
- **Compute nodes** within **application initiators (hosts)** are nodes in clustered application servers that connect to their allocated volumes via the *nvme-connect* command. Telemetry and Syslog events are collected from them as well. An agent monitors the state of the replication layer and generates events to the system log.

When KumoScale is installed in appliance mode, these components work in a data center as shown in the high-level architecture below.



KumoScale Software High-Level Architecture for Appliance Mode

In both Appliance and Managed mode, the data center infrastructure environment requests volumes from KumoScale software according to the requirements of hosted applications. These requests are implemented using KumoScale interfaces.

KumoScale storage nodes maintain the persistence of the configuration and manage the volume virtualization. The nodes periodically send telemetry on the physical SSDs and virtual volumes.

KumoScale Software Provisioning

The KumoScale Provisioner service is a distributed, stateless, resilient service that accepts requests for volumes, along with a specification detailing the volume's requirements known as the [storage class](#). The KumoScale Provisioner service returns the logical identity and network location where the requested volume can be accessed.

Provisioning takes into account a variety of factors to arrive at an optimal placement decision, such as resilience and topology requirements, capacity, and node utilization. In addition, the Provisioner Service will take into account the desired Quality of Service (QoS) parameters. KumoScale software interfaces support additional functions for managing provisioning in specific orchestration environments.

Solid State Disk (SSD) Groups

Every SSD deployed within a KumoScale storage node is automatically sorted into a group. The SSD group serves as a pool of storage. You can expand a group by adding new SSDs to it. The new capacity becomes available for mapping new volumes. See [Adding an SSD to a Group](#) for additional information.

Virtual Volumes

KumoScale software maps virtual volumes to physical drives. In order to optimize the memory utilization within KumoScale software and to enable the easiest and most intuitive integration with orchestration frameworks, KumoScale software automatically places the volumes within the storage node, taking into account the storage class defined by the user.

Targets and Access Control List (ACL)

Once volumes are created, they are exposed over an Ethernet fabric. Application initiators (hosts) and compute nodes establish an NVMe-oF protocol connection to the volumes using a logical entity called a target.

Connectivity between application initiators (hosts) and targets is controlled in the Access Control List (ACL). When a volume is created, it is added to a target with the corresponding ACL as a namespace, thus enabling the host to connect to it. The ACL is set per the host and is associated with the initiator/target pair. The permitted connection types are 'Read Only' and 'Read/Write' access.

An additional access control feature is a discovery ON/OFF setting. When discovery is set to OFF, the NVMe-oF protocol resource **cannot** be discovered over the fabric using the NVMe Discover command. The only way to connect to it is to know ahead of time the unique worldwide NVMe Qualified Name (NQN) of that resource. This prevents the discovery of storage resources by application initiators (hosts), even when they are connected to a data center trusted network.

The target creation and ACL settings are managed by the KumoScale Provisioner service, derived from connecting host parameters, and do not require the user to configure.

Targets and the ACL are described further in [Target Management](#).

Next: [Deploying a KumoScale Storage Provisioning System](#)

