

Overview to Installing KumoScale for Appliance Mode

An overview to the software, requirements, and process for installing KumoScale Storage Software in Appliance Mode. At the end of this document you will be guided to the installation steps appropriate for your environment.

What is KumoScale for Appliance Mode

The installation process consists of installing and configuring the following software components:

KumoScale Storage Cluster: KumoScale management services run on a Kubernetes™ orchestration cluster hosted on the storage nodes themselves. The services that run on this cluster are configured for native Kubernetes environments' data resiliency and high availability (HA). The installation automatically installs the following essential services on the cluster:

- **KumoScale Provisioner**, a service for provisioning, analyzing, and managing storage nodes. The Provisioner is installed on the master node of the storage cluster and acts as a load balancer, using storage requirements and analytics to determine the best location for allocating requested volumes.
- **KumoScale Control Operators** to manage KumoScale services. These control operators, also known as the Control Plane, run as services on the Kubernetes cluster.
- **The Container Storage Interface (CSI) driver for Kubernetes** to support implementation of the storage cluster and storage nodes
- **External Services (optional)** that provide internal logging and analysis can be added to the storage cluster.

KumoScale Storage Nodes: KumoScale software is installed on servers equipped with NVMe™ SSDs and configured as master or worker nodes.

More detailed information, including a system overview describing all components and interfaces, is available in the Overview section of the [KumoScale User Guide](#).

Once you have complete installing and configuring KumoScale software, your environment will be ready for provisioning as shown in Figure 1. Example of a KumoScale Software Deployment.

NOTE: If you want to install KumoScale to work with your own Kubernetes cluster, see the Installation instructions for [KumoScale in Managed Mode with Kubernetes](#).

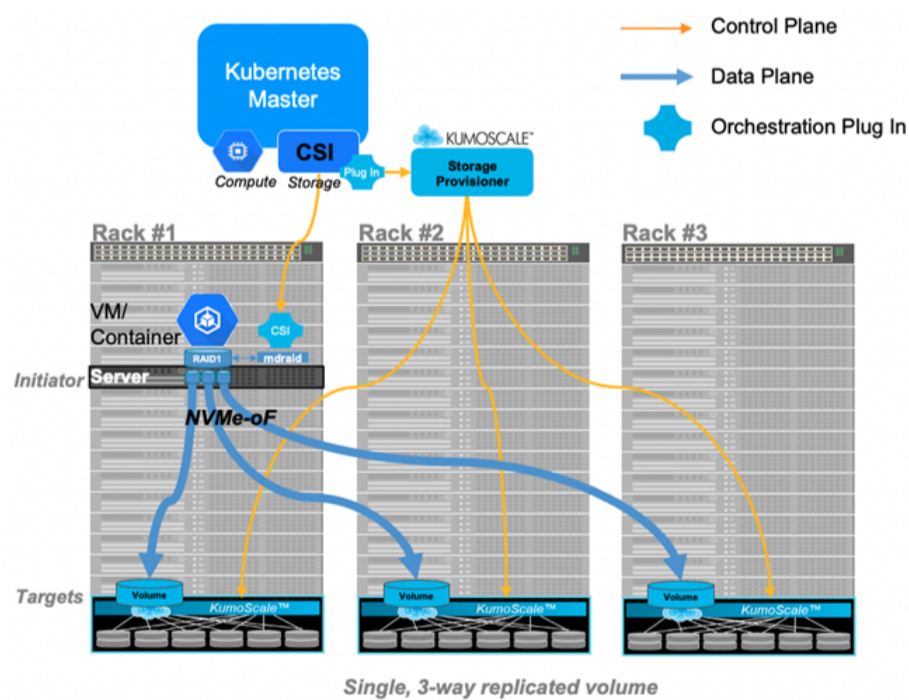


Figure 1. Example of a KumoScale Software Deployment

Intended Audience

This document is written for storage administrators. It is assumed that the reader has a working knowledge of storage, networking, and Kubernetes orchestration concepts.

Release Notes

For known issues and changes from the previous release, see the [Release Notes](#).

Preparing for Installation

This section details prerequisites for a successful implementation of KumoScale in Appliance mode. You should not proceed with installation and configuration until you have verified that your environment supports these requirements.

1. **Hardware Requirements:** KumoScale software may be installed on any of the server and network configurations defined in the [KumoScale Hardware Compatibility List \(HCL\)](#). If you wish to add RDMA NICs not included in the HCL, contact your KIOXIA representative.

2. Installation Device: KumoScale software supports deployment on platforms with either SATA drives or NVMe drives. By default, SATA drive(s) is/are used for installation. NVMe drive(s) will be used if no SATA drive is available. The installation devices are formatted as mirrored (RAID-1) devices if at least two drives are available. Otherwise, the installation is done on the single drive. More installation device options can be found under [Step 1. Determine Installation Parameters](#).

NOTE: We recommend using two SATA SSDs as installation media.

3. Network and Environment Requirements: There are three ways you can boot into the image depending on the device you choose.

NOTE: KumoScale storage node installation is supported only on platforms that are configured for UEFI.

Network Installation: This is the preferred installation method when installing KumoScale on multiple nodes. It works when installing remotely and requires a TFTP/DHCP PXE server that supports DNS, HTTP / HTTPS / FTP services on the installation site. You can review details on this installation at [Network Installation using a PXE Server](#) and [Network Installation using a iPXE Server](#).

Virtual CD Installation: The Virtual CD is a remote server containing the image and must be accessible from either the Baseboard Management Controller (BMC), integrated Lights Out (iLO) or iDRAC of the server. This option should be used when a PXE server is not available, and only remote access is available. You will need a way to transfer files over the network such as a WebDAV server, TFTP server et. al. You can review details on this installation at Virtual CD Installation.

USB Installation: This method installs the image directly on a server using a USB drive. This option is typically chosen when installing on a small number of nodes, or when a PXE server is not available. This method does not assume any special conditions.You can review details on this installation at USB Installation.

WARNING: Do not connect the bootable USB drive to a production server since it might delete its content without a warning.

Configuration Requirements

This section details the requirements needed to properly configure the KumoScale storage cluster and your storage nodes. You should not proceed with installation until confirming you have the information listed below.

1. License Key. The 3.22 download package comes with a license key that is good to set up up to five (5) storage nodes and lasts for three (3) months. If you need to use KumoScale outside these conditions, you will need to install the license key file provided to you by KIOXIA to configure the storage nodes. You will not be able to use KumoScale without a valid license key.

2. Administrative Host with Support for Kubernetes kubectl 1.17.5. For security reasons, users are not given direct access to the KumoScale storage cluster or storage nodes. All post-install steps are executed on a Windows™/OS X/Linux™ host that can connect to the KumoScale master node over the network and supports the Kubernetes command-line tool, [kubectl](#). Ensure that you have access to a host with this tool before starting the installation. This is your administrative host. You will need to use the version of kubectl compatible with Kubernetes v1.17.5, the version of Kubernetes implemented in KumoScale 3.22. To install kubectl on your administrative host, please follow the steps [here](#).

3. KumoScale Management IP and Storage Cluster VIP must be on the same subnet. The virtual IP or FQDN of the KumoScale storage cluster must be provided during the installation of the first master node for the cluster. It is bound to the management interface and must be on the same subnet as the KumoScale management IP.

4. Master Nodes are on the same subnet. The master nodes of the KumoScale storage cluster must all be on the same subnet.

5. Number of Master Nodes. A KumoScale storage cluster is optimally configured with an odd number of masters on the same subnet. For single node deployments, one (1) master is sufficient. High Availability (HA) systems should contain at least three (3) masters. Table KumoScale Cluster Configurations shows possible cluster configurations.

KumoScale Cluster Configurations			
Cluster size	Number of Master Nodes	Number of Worker Nodes	Deployment Environment
1	1	0	No resiliency
3	3	0	
4	3	1	Odd number of masters
5	5	0	Best practice
More than 5	5	1+	

The number of master nodes is specified when you install the first node. If the number of master nodes is specified as, say, three, then the first three nodes installed are master nodes. All nodes added after that are configured as worker nodes. You may add nodes to your KumoScale cluster at any time. Note that it is possible to prevent a node from joining as a master node (see [Step 6. Configure the Node](#)).

6. VIP for the First Master aka Control Plane Endpoint. You need to set up one server in the KumoScale storage cluster with licensing and configuration details that will be used to deploy the other nodes. This is referred to as the Control Plane End Point or First Master. It will also serve as the KumoScale storage cluster Virtual IP Address (VIP). You must provide it when you install KumoScale software on the very first node (details are in *Installation Command Line Parameters*). All master nodes must be on the same subnet, so a node will join as a master only if it is on the same subnet as the first master.

NOTE: In KumoScale 3.22 if you specify a FQDN rather than an IP address for the VIP, you will be able to modify the VIP at a later time. This is explained in the KumoScale User Guide.

7. Custom Requirements for Masters (optional)

. You may choose to specify requirements for when a node can join as a master of the KumoScale storage cluster. These restrictions need to be provided when configuring the Master CR and are referred to as affinity parameters.

- **Affinity:** the node meets a certain set of requirements such as it is in a specific rack, or it has a particular name
- **Antiaffinity:** the node DOES NOT meet a certain set of requirements such as it is NOT located in a specific rack.

You will need to select whether you want a node to join only when there is affinity or antiaffinity or neither; you cannot have both affinity and antiaffinity. The default setting is neither. You will express this in the Master CR at [Step 6. Configure the Control Plane Endpoint \(First Master\)](#). Examples of syntax are provided in that section.

8. Single Node vs. Multiple Node Deployments

A **single node** environment consists of one master node only; no workers. This is useful for testing or For Proof of Concept (PoC).

A multiple node deployment consists of at least two nodes: more than one master node or at least one worker node. This is common with production deployments.

Once you have confirmed your environment meets installation requirements and you have collected all the information needed for configuring KumoScale, proceed with installation based on your deployment.

Next: [Download KumoScale software](#)
