

This section documents the overall process for implementing KumoScale software in your environment using KumoScale custom resource files.

- [Overview of a KumoScale deployment](#)
- [KumoScale Operators](#)
- [KumoScale Custom Resource Definitions \(CRD\)](#)
- [Deployment Methodology](#)

As explained in [KumoScale Software Overview](#), the primary interface for KumoScale storage provisioning is the KumoScale Provisioner REST API. KumoScale also offers plug-ins to the Provisioner API for Kubernetes™ CSI, OpenStack™ Cinder and Ansible™ for Bare-Metal environments.

[illegible]

Supporting the Provisioner is the KumoScale control plane that is a set of Kubernetes Operators that control the storage system (whether there is one node or multiple nodes). These are deployed as a set of Kubernetes containers running in a cluster spread across the available storage nodes. KumoScale deploys several Kubernetes operators, detailed throughout this document, that control storage system configuration and operation.

The KumoScale Software Provisioner REST API is the primary interface for storage management. The KumoScale Provisioner REST API can be accessed via SSH or terminal simply by entering valid REST API credentials. Details on how to access the Provisioner service REST API are available in [KumoScale Provisioner REST API](#).

The KumoScale Control Plane supports two operators:

- There are multiple Custom Resource Definitions (CRD) used to manage features particular to each resource type.

For security reasons, users are not given direct access to the KumoScale storage cluster implemented in Appliance mode and must use a remote host to issue `kubectl` commands. `kubectl` is a command line tool for controlling Kubernetes clusters. For general information on `kubectl`, see <https://kubernetes.io/docs/reference/kubectl/overview>.

Configuring a remote host (based on Windows™, iOS™, or Linux™ operating systems) to enable access to them KumoScale storage cluster is documented in [Step 5 of the installation instructions for Appliance mode](#).

- You must use the version of kubectl compatible with your version of KumoScale (see the [Release Notes](#) for your environment). Users can use Secure Shell (SSH) or terminal access to storage nodes with the Cluster Manager CLI or the REST API.
- Storage administrators should use kubectl to tell KumoScale operators how to configure nodes, logging, and other components.

KumoScale Custom Resource Definitions

KumoScale software defines several Custom Resource Definitions (CRD) that tell KumoScale Operators exactly how to manage the storage system. CRDs are implemented as **yaml** files to support the deployment environment and KumoScale components.

This section summarizes the KumoScale CRD used to customize and support a deployment. Examples of each may be found in the **KumoScale_Operator/ks-config-operator-samples** directory.

- [KumoScale Master CRD](#) (Appliance Mode) are used to define the masters on the KumoScale storage cluster for Appliance mode. These CRD are configured and deployed during installation.
- [Kubernetes Storage Cluster Master CRD](#) (Managed Mode) are used to define the masters on the Kubernetes storage cluster for Managed mode.
- [Telemetry](#) and [Syslog CRD](#) are used to collect data for planning and support purposes and usually set up after the KumoScale Storage Cluster is configured.
- [Storage Node CRD](#) are one of several CRD used to configure storage nodes. You may set up different storage node CRD yaml files depending on different application needs and environmental settings. Storage nodes can be configured any time after installation.

In addition, the following CRD are described elsewhere in the links shown.

- [SSD Operations CRD](#) may be used to manage SSD.
- [Volume Operaton CRD](#) is used to perform volume migrations from one storage node to another.

NOTE: Not all parameters of KumoScale CRD are available in Managed mode; differences are noted where needed.

KumoScale Storage Cluster Master CRD for Appliance Mode

There are several CRDs used to configure masters on the KumoScale storage cluster. The only time that you will use the CRD files bulleted below is for installing, upgrading, or updating KumoScale software on the master nodes. See the [Installation Guide for Appliance Mode](#) for actions that require using these CRD and their specifications. A high-level summary of these actions follows:

- **Master CRD** may be configured at install time with the number of masters and affinity/anti-affinity parameters. There is only one master CRD file per KumoScale storage cluster. A skeleton of the CRD file is provided with the KumoScale software for appliance mode. You can modify the VIP for the master at a later time, only if the FQDN is used as explained in [Modifying the VIP for the KumoScale Master in Appliance Mode](#).
- **License CRD** is edited and deployed at install time with the license key provided to you by KIOXIA as explained in the [Installation Guide for Appliance Mode](#). You will not be able to provision or allocate volumes without a valid license key. If you need to update the key, for example if moving from a trial to a production license, see Replacing the License in [Maintenance, Troubleshooting, and Support](#).
- **Internal Services CRD** are CRD associated with Pmetheus, Grafana, Loki, and Fluentd. These CRD are edited at install time with the KumoScale storage cluster VIP and replication information. There is only one CRD file for each service and it has the form: <service name>.kumoscale.kioxia.com_v1_<service name>_cr.yaml. For example, `fluentd.kumoscale.kioxia.com_v1_fluentdservice_cr.yaml`

Modifying the VIP for the KumoScale Master in Appliance Mode

The master node is configured at install time, but you may find it necessary to modify the VIP for the node at a later time. This is only possible if a FQDN was used to define the VIP, and that your DNS server is available. Below are the steps to follow to modify the master VIP:

1. Confirm that the new IP and the Storage Cluster VIP are on the same subnet.
2. On your DNS server, assign the new VIP to your FQDN.

As long as the DNS server is available, your KumoScale appliance will use the new VIP.

Storage Cluster Master CRD for Managed Mode

There are several CRDs used to configure masters on the Kubernetes storage cluster. See [KumoScale installation guide for Managed Mode with Kubernetes](#) for actions that require using these CRD at install time. A high-level summary of these actions follows:

- **License CRD** is deployed at install time with a license key provided to you by KIOXIA as explained in Installation Guide for Managed Mode for Kubernetes. You will not be able to provision or allocate volumes without a valid license key. If you need to update the key, for example if moving from a trial to a production license, see [Replacing the License](#).
- **Internal Services CRD** are CRD associated with Grafana, Loki, Fluentd, and Prometheus. Prometheus may be used in both appliance mode and managed mode. Loki and Fluentd are only supported in appliance mode; contact KIOXIA support for information on how to use these applications with KumoScale in Managed Mode. There is a CRD file for Prometheus, Loki, and Fluentd. There is only one CRD file for each service and it has the form:
<service name>.kumoscale.kioxia.com_v1_<service name>_cr.yaml.
For example, `prometheus.kumoscale.kioxia.com_v1_prometheus service_cr.yaml`

Telemetry and Syslog CRD

In addition to the external tools above, there are several CRD that can be used to collect data for planning and support purposes. These CRD are defined in [Telemetry](#) and [Syslog](#) along with instructions and examples for creation.

Storage Node CRD

Storage node CRD files are used to configure masters and workers on the storage cluster. You can create templates for storage nodes intended to be used by particular applications. Storage node parameters are explained with examples in [Creating and Managing Storage Nodes](#).

Deployment Methodology

This section describes the high-level steps for implementing KumoScale software in your environment. Details on how to complete each step is provided in this and other documents as shown in the links.

Note that not all features will be available to users of Managed mode. Differences between appliance and managed mode are noted throughout the documentation.

1. Develop an implementation plan: Review your environment architecture and requirements with respect to KumoScale software features and requirements. You will need to determine:

- System, network, and topological requirements of the KumoScale storage cluster. [Creating and Managing Storage Nodes](#) explains how these parameters are set using KumoScale CRDs to deploy storage nodes. This will help you determine the parameters you will need to define in CRD files used for your implementation.
- The number of masters you will need are based on your need to scale up. This is explained in [Before you Begin using KumoScale](#)
- Logging and Telemetry needs, as outlined above, are implemented with KumoScale CRD. You will need to create and specify logging and telemetry parameters in KumoScale Custom Resource files described in [Creating and Managing Storage Nodes](#).
- A list of users who will need access your environment and what their roles will be in all phases of design, implementation, and ongoing support. You will need to determine the [Authentication](#) method that best applies in your environment.
- [KumoScale interfaces](#) you will use to integrate KumoScale software into your current orchestration framework.

At this point you are ready to begin implementing your design.

2. Install KumoScale software: Install KumoScale software on servers to be configured for the KumoScale Storage Cluster as documented in the *KumoScale Installation Guide* for either Managed or Appliance mode. Note that for this guide, we assume Managed mode is used with a healthy, high-availability Kubernetes cluster to support the features documented in this guide.

3. Define and deploy the KumoScale Storage Cluster: Create the storage nodes in the KumoScale Storage Cluster as instructed in the KumoScale Installation Guide. You will need to create and specify values for the:

- [Master Custom Resource](#) for Appliance mode used for defining requirements for KumoScale storage cluster masters, and
- [Storage Node Custom Resource](#) used to create storage nodes in the KumoScale storage cluster.

4. Install internal telemetry tools on the master nodes: Install tools used for analytics as documented in the KumoScale Installation Guide for [Appliance Mode](#) or [Managed Mode](#). Note that not all tools are available in Managed Mode. Contact KIOXIA support if you wish to use other tools.

5. Optionally configure and deploy telemetry and Syslog. This step should be done if you want to push data to an external tool such as Syslog:

- **Define and deploy telemetry:** Create and deploy the CRD for telemetry as described in [Telemetry](#).
- **Define and deploy Syslog:** Create and deploy the CRD for Syslog as described in [Syslog](#).

At this point, your implementation is equipped with masters, and you are ready to add storage nodes as workers.

6. Install KumoScale software on servers to be used as storage nodes as documented in the KumoScale Installation Guide for Appliance Mode or Managed Mode.

7. Create the storage node CRD files: Define the [Storage Node Custom Resource definitions](#) to be used for deploying storage nodes. You may elect to create several CRD definitions to reflect different application needs or different environment settings.

8. Configure the storage nodes for provisioning by the KumoScale Provisioner as described in [Creating and Managing Storage Nodes in KumoScale](#).

Your implementation is now equipped with storage nodes that will be automatically provisioned, monitored, and logged by KumoScale software as described in the [KumoScale Software Overview](#). You are ready to set up user access and interfaces to other orchestration environments.

9. Create additional users (Appliance mode only): [Configure RBAC Users](#) to fulfill other roles for supporting the deployment going forward.

10. Integrate with your orchestration environment:

Use the KumoScale REST API, Cluster Manager CLI, Ansible, OpenStack, or CSI interfaces to complete the remaining steps that may be required to successfully implement your design. Note that orchestration depends on whether you are using KumoScale in appliance or managed mode. [Setting Up Initiators](#) provides the initial requirements for your application initiators, also known as compute nodes. For complete details on how to use these interfaces see the KumoScale documentation page for the appropriate user guide (REST API, Ansible, OpenStack, Cluster Manager CLI, or CSI).

Next: [Creating and Managing Storage Nodes in KumoScale using Custom Resource Files](#)

