

Set up Storage Nodes for Provisioning

This section describes the process for setting up and adding storage nodes to your cluster when using KumoScale software in Managed Mode.

•Set up Storage Nodes for Provisioning

Step 1 Set up your storage cluster to use the Provisioner

Step 2 Add storage nodes.

Set up your Storage Cluster to use the KumoScale Provisioner

•Set up Storage Nodes for Provisioning

Step 1 Set up your storage cluster to use the Provisioner.

To configure the Provisioner for your existing Kubernetes cluster, you will need to complete the following in the order listed:

1. (OPTIONAL) Set the Provisioner license.
2. If using OpenID Connect (OIDC) for authentication, set up authorization CR files.
3. Set the KumoScale secret.

We will show you how to use custom resource files to complete these steps. See the [KumoScale Provisioner REST API](#) or [Cluster Manager CLI](#) if you wish to use those tools instead.

1. **(OPTIONAL) Specify the license key provided to you by KIOXIA.** KumoScale comes packaged with a license key that is valid for up to three (3) months and five (5) storage nodes. You can skip this step if your current deployment can support these restrictions and deploy the production license at a later date. If you prefer to set the license custom resource file with the production key provided to you by KIOXIA:

a. Edit the license Custom Resource located in Kumoscale Operator/deploy/crds/kumoscale.kioxia.com_v1_license_cr.yaml, and replace the value of license with the license key provided by KIOXIA.

b. Save the file then install the license:

```
$ kubectl apply -f kumoscale.kioxia.com_v1_license_cr.yaml
```

This will return a message about the KumoScale Provisioner service license being created like:

```
license.kumoscale.kioxia.com/provisioner-license created
```

c. You can validate that the KumoScale software license was installed with the KumoScale Cluster Manager command [license-show](#), or by using the following command:

```
$ kubectl describe license provisioner-license
```

2. **If you are using OpenID Connect (OIDC) for authentication**, you will need to set up the custom resources **authorization-secret.yaml** and **authorizationserver_cr.yaml** as described in the [Authentication](#) section of the User Guide.

3. **Create the KumoScale secret** with the file **kumoscale-secret.yaml**. An example file is in the directory **KumoScale Operator/ks-install-operator/deploy/crds**. The secret should contain the desired admin password base64 encoded. Password requirements are defined according to the current Linux OS password policy. You must change the default password; failure to do so will result in an error and you will not be able to add storage nodes successfully.

a. To encode the admin password, run the command:

```
$ echo -n 'YourPassword' | base64
```

b. Edit **kumoscale-secret.yaml** and copy the password returned above into the **password** field as shown below:

```
apiVersion: v1
kind: Secret
metadata:
  name: kumoscale-secret
  namespace: default
type: Opaque
data:
  password: <password-returned-from-step-4a>
```

c. Set the secret with:

```
$ kubectl apply -f kumoscale-secret.yaml
```

d. The system will return confirmation that the secret file was created.

Add Storage Nodes

•Set up Storage Nodes for Provisioning

Step 2 Add storage nodes.

You can now add storage nodes to your cluster using custom resource files and **kubect**l. A sample storage node Custom Resource (CR) file, **kumoscale_v1_storagenode_cr.yaml** is available in the **KumoScale Operator/ks-install-operator/deploy/crds** directory included with KumoScale software.

To configure and deploy a new storage node, follow these steps:

1. Confirm that your data network is configured and running.
2. Make a copy of the example file **kumoscale_v1_storagenode_cr.yaml** for editing, and save to a separate directory. For example, **myapp/deploy/crds/myapp_storagenode_cr.yaml**.
3. Update **myapp_storagenode_cr.yaml** with required values for the parameters that apply to your deployment, such as

Parameter	Description
name	Unique name for the storage node. It must comply with the Name rules (see KumoScale Field Types)
initMgmtIp	The initial management IP address.
adminSecretName	Name of the secret file created during installation.
groupName	Name of the group. It can remain with the default value.
AuthenticationMode	The authentication mode; select one: Local , Ldap , OPENIDC
portal:ip	IP address of the data network; typically, the same used for the initMgmtIP suffix.
portal:subnet	The IP address subnet mask.
portal:interface	The interface to host the IP address. A list of available interfaces may be obtained using the Cluster Manager CLI interface-show command.

The above parameters are the minimum you will need to specify; there are many others. See the [Creating and Managing Storage Nodes in KumoScale](#) for complete documentation on parameters to use when adding and modifying storage nodes .

You can add definitions of multiple nodes in the same file; just confirm that the name and IP addresses are different.

Note on Topology parameters: If you plan to use the KumoScale Kubernetes CSI driver, you need to also label each node with topology labels before installing the driver. If not you will need to redeploy the nodes after labeling these attributes:

```
topology.kubernetes.io/region
topology.kubernetes.io/zone
topology.kubernetes.io/rack
```

4. Create the node with:

```
$ kubectl apply -f myapp_storagenode_cr.yaml
```

Verify that the node was created with kubectl or using the KumoScale Cluster Manager CLI command [storagenode-show](#). For example

```
$ kubectl get storagenodes

NAME          AGE   STATE    VERSION    ALERTS
storagenode1  18h   Available  3.22-15989  0
```

5. Continue to add as many storage nodes as needed. You can use separate files for each node or use the same file for multiple nodes and run the apply command once . You will need to confirm that the name and IP are different.

Next Steps

(Optional) Using Managed Mode

For more information on using KumoScale in Managed Mode see:

- [Maintenance Activities for KumoScale in Managed Mode](#)
- [Using the KumoScale Provisioner in Managed Mode](#)
- [Using the KumoScale Engine in Managed Mode](#)

(Optional) Install Prometheus

Follow the instructions at [Using Prometheus with KumoScale in Managed Mode](#) to install Prometheus.

Implementing and Managing your Storage Environment

See the [KumoScale User Guide](#) for complete information on how to set up and use storage nodes, provisioning, logging, and analysis for your environment.

Setting up Initiators

To set up initiator nodes to request volumes, refer to KumoScale documentation on the interface applicable to your deployment:

- [CSI](#), to install and use the CSI driver for Kubernetes.
- [OpenStack](#), to integrate with OpenStack.
- [Ansible](#), for bare-metal integration.

You will also need to install the NVMe host module patch on these nodes as documented in [Installing NVMe Host Module Patch on Compute Nodes](#).
