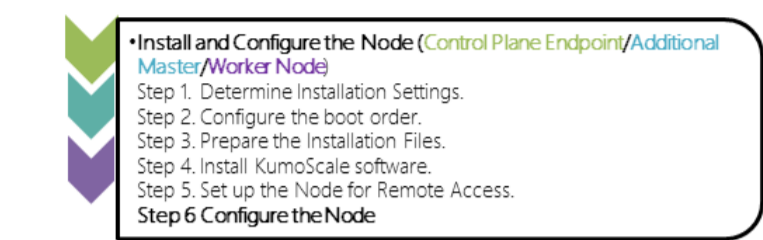


Step 6 for All Nodes. Configuring the Node

This page explains Step 6 of installing KumoScale for Appliance Mode in a multi-node environment.



This section explains how to configure master and worker nodes of the KumoScale Storage Cluster. The next step depends on what type of node you need to configure. Remember you must configure all masters before configuring storage nodes. If you are configuring the:

- **First (master) node**, follow the instructions in [Step 6: Configure the Control Plane Endpoint](#).
- **Additional Masters**, follow the instructions in [Step 6: Configure Additional Masters](#).
- **Worker nodes**, follow the instructions in [Step 6: Configure Worker Nodes for Storage](#).

First Node Step 6: Configure the Control Plane Endpoint

You are now at the last step for installation and configuration of the first master as shown in the diagram.



You need to set up the first master 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 Endpoint**. Once completed, this step is not needed for any other masters added to the KumoScale storage cluster, nor for any other storage nodes.

From the remote host, complete the steps below; you should be at the **CLI**> prompt. The sample custom resource files referenced in this section are included with the KumoScale software and are in the directory `operators/ks-config-operator/samples`.

1. **Set the secret for KumoScale software.**

`kumoscale-secret.yaml` should contain the desired admin password base64 encoded. Password requirements are defined according to the current Linux OS password policy.

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-1a>
```

c. Set the secret with:

```
kubectl create -f kumoscale-secret.yaml
```

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

2. **(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.

a. Edit the license Custom Resource, `kioxia.com_v1_license_cr.yaml`, and replace the value of license with the license key provided by KIOXIA.

b. Save the file then run this command to install the license:

```
kubectl create -f kumoscale.kioxia.com_v1_license_cr.yaml
```

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

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

c. You can also validate that the KumoScale software license was installed with the following command:

```
kubect1 describe licenses provisioner-license
```

3. Edit the Master CR, yaml, and set the value of:
- a. **numberOfMasters** to the number of masters you want to be in the KumoScale storage cluster. See [Number of Masters](#) for cluster requirements.
 - b. **affinity** or **anti-affinity** to specify any constraints on when a node can be added to the storage cluster as explained in [Custom Requirements for the Masters](#). Both entries are ‘commented’ which mean they do not apply. If you want to specify custom constraints you will need to uncomment these values.

For example, the Master CR below specifies that the storage cluster has the following requirements:

- It has three (3) masters.
- To join the cluster as a master, the node must have the name: node1, node3, or node9.
- To join the cluster as a master, the node must not be in the same rack or region as existing masters in the KumoScale storage cluster.

```
apiVersion: kumospace.kioxia.com/v1
kind: Master
metadata:
  name: master
spec:
  numberOfMasters: 3
  affinity:
    matchExpressions:
      -key: Kubernetes.io/hostname
      operator: In
      values:
        - node1
        - node3
        - node9
  antiAffinity:
    - topologyKey: "topology.kubernetes.io/rack"
    - topologyKey: "topology.kubernetes.io/region"
```

4. Configure the first master node for the cluster with:

```
kubect1 create -f kumospace_v1_master_cr.yaml
```

You will receive a message that the first master node was configured.

5. You now need to add this as a storage node on the KumoScale storage cluster. Create a new CR to define the node. Edit **kumospace_v1_storagenode_CR.yaml** and save with a new name such as **kumospace_v1_storagenode1_CR.yaml**.
6. Your new CR should have new values specific to the server. A complete list of possible parameters is available in the [KumoScale User Guide](#). You should review all the parameters before creating the node. The following parameters are required:
- **name:** Host name.
 - **initMgmtIp:** IP address of the server. If you specified the static IP during installation, you must provide the same address here. Otherwise, provide the KumoScale software IP returned using DHCP.
 - **adminSecretName:** The name of the secret created in step 1 above.
 - **timeSettings:** time zone ID.
 - **mode:** NTP.
 - **ntpServer:**NTP server FQDN or IP address.
 - **topology:** This information will be verified against any affinity or anti-affinity information in the Master CR. The server will only join as a master if it meets the requirements.
 - **portals:** The network ports being used for storage data.
 - **transportType:** TCP_IP or RoCEv2.

An example CR for a server deployed as a master to the storage cluster is shown below. Note that we do not show the IP addresses and subnet for security reasons. You will need to specify addresses.

```
apiVersion: kumospace.kioxia.com/v1
kind: StorageNode
metadata:
  name: ks-node1-000c298c715f
spec:
  initMgmtIp: ###.##.##.###
  adminSecretName: kumospace-secret
  groupName: group1
  timeSettings:
    timeZoneID: Greenwich

  mode: NTP
  ntpServer: ###.###.###.###
  network:
    portals:
      - ip: ###.###.###.###
        name: portal1
        subnet: ###.###.##.#
        interface: kx0
        port: 4420
        transportType: TCP_IP
  topology:
    - name: topology.kubernetes.io/rack
```

```
value: "RACK1"
- name: topology.kubernetes.io/zone
value: "LAB"
```

Once the above configuration file has been prepared, run following command:

```
kubectl create -f kumoscale_v1_storagenode1_CR.yaml
```

7. You should receive confirmation that the master node was created. In addition, you can verify that the node was created as a master with the name provided with any of the commands below.

```
kubectl get nodes -A -o wide
kubectl get storagenodes -A -o wide
kubectl cluster-info
```

For detailed information on the node enter:

```
kubectl describe storagenodes
```

Next: You are now done setting up the Control Plane Endpoint. You should now either

- Return to [Step 1](#) to install KumoScale on Additional masters, if the number of masters is greater than one (1), or
- Proceed to [Step 7](#) to verify the KumoScale storage cluster.

Additional Masters Step 6: Configure Additional Masters

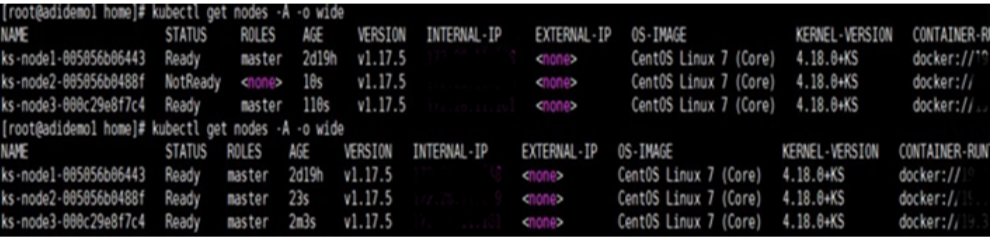
•**Install and Configure Additional Masters if Number of Masters > 1**
Repeat steps 1 through 6 for any additional masters.
Step 1. Determine Installation Settings.
Step 2. Configure the boot order.
Step 3. Prepare the Installation Files.
Step 4. Install KumoScale software.
Step 5. Set up the Node for Remote Access.
Step 6. Configure the master.

Note: To proceed with adding nodes, you must confirm that

- The storage node is an uninitialized (not already configured) node with all available SSDs attached.
- You have both a KumoScale secret and a valid license file.

As noted earlier, you cannot add nodes as storage nodes until the number of master servers in the KumoScale storage cluster is equal to the value of **numberOfMasters** specified in the Master CR. Servers will stay in a *Not Ready* state until all required masters are created.

For example, the screen shot below shows the results of **kubectl get nodes -A -o wide** when only two (2) masters have been configured for a cluster of three (3), that is **numberOfMasters** is 3, and when all three (3) masters have been configured.



To configure a node as a master on the KumoScale cluster:

1. Edit the storage node CR created in [Configure the Control Plane Endpoint \(First Master\)](#). For example **kumoscale_v1_storagenode1_CR.yaml**. Give it values specific for this server:

name
IP

Other parameters, such as the **adminSecretName** will not change. Topology information will be verified against any affinity or anti-affinity information in the Master CR. The server will only join the cluster as a master if it meets the requirements.

2. Deploy the node as a master on the KumoScale cluster using your storage node CR. For example:

```
kubectl create -f kumoscale_v1_storagenode1_cr.yaml
```

Note: When adding an additional second node as a master, if you issue the kubectl describe command, you may get the message **Cannot add a master while there are not enough masters pending**. This is not an error but only reflects that since the number of master nodes needs be an odd number, the storage node is in a pending state until a third storage node is added as a master. You will see this again with an even numbered master node.

Next:

- If there are additional masters to be set up (i.e. numMasters > the number of masters installed and configured with KumoScale), then return to [Step 1](#) to repeat steps 1 through 6 for any additional masters.
- If all masters have been set up, then proceed to [Step 7](#) to complete the setup of the KumoScale Storage Cluster Configuration.

Worker Nodes Step 6: Configure Storage Nodes on the KumoScale Cluster

•Install and Configure Storage Nodes if Number of Workers >0
For each storage node perform the following:

Step 1. Determine Installation Settings.
Step 2. Configure the boot order.
Step 3. Prepare the Installation Files.
Step 4. Install KumoScale software.
Step 5. Set up the Node for Remote Access.
Step 6. **Configure the storage node for your application requirements.**

Notes:

- 1) Storage nodes are configured by creating nodes from customizations you make to the Storage Node CR. The KumoScale User Guide describes the many parameters that can be specified so that you can create your own templates for different types of storage nodes. You should review the possibilities before creating the nodes as indicated in Step (2) below.
- 2) It is very important that all components in a cluster are synchronized. Some sites may use external worldwide NTP servers while others might use an internal NTP server. We strongly recommend that you validate cluster wide timing synchronization between all components to reduce the risk of issues related to nodes probe timers.

To configure KumoScale software on each server designated for storage, complete the following steps:

1. Create the CR for your storage node using `kumoscale_v1_storagenode_CR.yaml` as a template. Your new CR, we will refer to it as `kumoscale_v1_storagenode2_cr.yaml`, should have values specific to the node such as name, IP, and network information. A description of all possible settings with examples is given in the KumoScale User Manual. You may defer creating CRs and deploying storage nodes as workers until you have fully explored all possible settings for the nodes.
2. Deploy the node using the settings specified in yaml with:

```
kubectl create -f kumoscale_v1_storagenode2.yaml
```

Next:

- If there are additional workers to configure, return to [Step 1](#) to repeat steps 1 through 6 on the additional worker nodes.
- If all workers and masters are configured, proceed to [Step 7](#) to verify the storage node configuration.
