

Step 1 for All Nodes: Determine Installation Settings

This page explains Step 1 of installing KumoScale for Appliance Mode in a multi-node environment.

•Install and Configure the Node (Control Plane Endpoint/Additional Master/Worker Node)

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 Node

Note:

The first server to be installed and configured must be the first master of the storage cluster, also known as the Control Plane Endpoint.

You must provide the virtual IP or FQDN of the KumoScale storage cluster for this first node. It is bound to the management interface and must be on the same subnet as the KumoScale management IP. Do NOT provide it for any other node.

Several installation settings can be changed via boot parameters passed in through the installation menu. You will be prompted for values for these settings when installation begins. They are listed here so you can collect the required information in advance, thus enabling to go more smoothly. Several examples are provided in this section to help you determine how to use these parameters.

Installation Settings

Menu Option (Parameter)	Options (default in boldface)	How Used
Mode (kx_mode)	Deploy Purge	<ul style="list-style-type: none">Deploy (kx_mode=deploy)(default) will delete the contents of the system disk and deploy a new image.Purge (kx_mode=purge) will completely wipe out the contents of the devices specified in the kx_devices parameter. This is useful for wiping previous installations. For command line only: <ul style="list-style-type: none">kx_mode=upgrade will upgrade the existing deployment, without modifying the data partition. When specifying upgrade, do NOT pass any other kx_ parameter other than kx_debug; passing any other kx parameter will cause the upgrade to fail.
Auto Select Devices (kx_raid)	Auto Raid Single	<ul style="list-style-type: none">Auto (kx_raid=auto) (default): automatically discover the best RAID deployment and create a RAID over them.Raid(kx_raid=raid): install over two drives (either detected or specified). Will fail if not enough devices were detected or specified in kx_devices.Single (kx_raid=single): install over a single drive. Will fail if more than one drive was specified in kx_devices.
Device Type (kx_blocktype)	ATA NVMe	If specified, this indicates the type of device for deployment: It is ignored if Purge is selected, and purging will be determined according to the value of Boot Devices (kx_devices).. <ul style="list-style-type: none">ATA (kx_blocktype=ata) (default): Use only ATA devices for the deployment process, deploying to nvme, usb or any other device type will fail.NVMe (kx_blocktype=nvme): Use only NVMe devices for the deployment process, deploying to ata, usb or any other device type will fail.

Menu Option (Parameter)	Options (default in boldface)	How Used
Root Partition Size (kx_rootsize)	Integer with units (string) 16 GiB [1]	Modify the size of the root partition to this value. If units are not specified, it will assume MiB. If device/s were selected, validate that the selected size can fit the device/s capacity Note: You must enter the same value during an upgrade, or the upgrade will fail.
Boot Devices (kx_devices)	No default value.	<p>Allows you to select one or two deployment BOOT device/s (depending on the value of RAID mode) that matches the selected device type. You have the option of selecting the devices from a list showing Name, Size, Vendor, Model, S/N (if applicable).</p> <p>Required when Mode=Purge (kx_mode=purge), otherwise optional.</p> <p>This specifies the PCI slots of the drives that KumoScale software will be deployed to, as a comma-separated list of their slot numbers. The slot numbers should be provided in their full form, as they appear in the output of lspci -D.</p> <p>You can specify slots that belong to either an ATA or NVMe device, and your selected devices will be verified against the specified Block Type (kx_blocktype). For example:</p> <ul style="list-style-type: none">· If you specify two (2) devices and at the same time specify Raid Mode is single (kx_raid=single) the process will fail.· If you specify one (1) device and at the same time specify Raid Mode is Raid (kx_raid=raid) the process will fail.· If you specify a PCI slot that matches more than one (1) disk - such as `0000`, the process will fail. <p>Examples:</p> <ul style="list-style-type: none">· kx_devices=0000:1a:00.0,0000:5e:00.0 : Use two (2) devices - 0000:1a:00.0 and 0000:5e:00.0 and do not automatically detect which devices to use.· kx_devices=0000:18:00.0: Use one (1) device - 0000:18:00.0 and do not automatically detect which device to use.· kx_devices=0000:00:1f.2/ata9,0000:00:1f.2/ata10: Use two (2) SATA devices - 0000:00:1f.2/ata9 and 0000:00:1f.2/ata10 and do not automatically detect which devices to use.
IP Mode	DHCP Static	Sets the mode for the management IP. When Static is selected, the system prompts for additional information.
Static IP (kx_static_ip)		<p>Static IPV4 address for the Management IP address. The static IP and VIP must reside on the same subnet. When a value for this parameter is specified, the Static IP will be automatically configured based on the following algorithm for choosing the management interface:</p> <ol style="list-style-type: none">1. Use the first online physical network interface (OPERSTATE= up and CARRIER=1) found by sorting by minimum speed, PCI Bus, and MAC address.2. Configure the Static IP on the found interface: <p>Configure the Static IP on the found interface: Check the connectivity to the server provided by the kx_gateway parameter. If the server unreachable, then clear the current interface configuration, exclude the current interface from the list, and return to step 1.</p> <p>Note: If the Static IP is configured, the Dynamic Host Configuration Protocol (DHCP) network mode will be disabled.</p>
Subnet Mask (kx_subnet_mask)	255.255.0.0	Subnet mask to be configured with the Static IP in IPV4 IP address format xxx.xxx.xxx.xxx. This parameter is required when specifying a value for Static IP (kx_static_ip).
Default Gateway (kx_gateway)		Default gateway in IPV4 IP address format xxx.xxx.xxx.xxx. This parameter should be used only when specifying a value for Static IP (kx_static_ip).

Menu Option (Parameter)	Options (default in boldface)	How Used
First Master Node	Yes No	Select whether the deployment is set for a first master node (Yes) or not. The default value is No.
Control Plane Endpoint VIP/Name (kx_cluster_vip)	Required for the first node only. Do NOT provide for any other node.	<p>The Control Plane Endpoint is also referred to as the First Master. The virtual IP or FQDN of the KumoScale storage cluster must be provided during the installation of the first node for the cluster. It is bound to the management interface and must be on the same subnet as the KumoScale management IP. Do NOT provide it for any other node.</p> <p>Note that 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 using your DNS server and KumoScale will use the updated VIP as long as the DNS server is available.</p> <p>Example values: 192.0.2.0 or myvip.domain.com</p>
Pods CIDR (kx_pods_cidr)	Address in CIDR notation	<p>This is an optional parameter that allows you to override the IP range used for communication. It can be defined for the first master node installation only. It allows the user to manually select the value of \$POD_NETWORK_CIDR. When this parameter is not defined by the user, a /16 range is selected out of the free private range. The value must be in CIDR notation and must define a free private address range. The free private range is within:</p> <ol style="list-style-type: none">192.168.0.0/16 - entire range10.0.0.0/8 - excluding 10.96.0.0/12 used for Kubernetes services (10.0.0.0/16 - 10.95.0.0/16, 10.112.0.0/16 - 10.255.0.0/16)172.16.0.0/12 - excluding 172.16.0.0/16 used by docker (172.17.0.0/16 - 172.31.0.0/16) <p>· All the above excluding subnets configured on the machine (management: static/dhcp).</p>

Bypassing the Installation Menu

If you prefer to input the parameters through the command line rather than a menu, you must specify the boot parameter **kx_nomenu=true** (default value is false) and provide the values for the above parameters according to your install approach:

- Network (PXE, iPXE) Installation:** Add these parameters to the pxelinux menu entry of the cfg menu.
- Virtual CD:** Supply the values through the grub menu.
- USB Installation:** Choose the installation option during system installation and press Tab to enter each value.

Debugging Installation

By default, a limited amount of information is displayed during installation. If you are having issues with installing KumoScale software, you may find it helpful to disable quiet mode and print extended information during deployment. To do this, specify the boot parameter **kx_debug=true** (default value is false) according to your install approach:

- Network (PXE, iPXE) Installation:** Add the parameter to the pxelinux menu entry of the cfg menu.
- Virtual CD:** Supply the values through the grub menu.
- USB Installation:** Choose the installation option during system installation and press Tab to enter each value.

Examples

Below are examples of a network installation where a new installation of KumoScale software is deployed. The content of each menu entry sets up a node that will be configured as a master or worker node:

Install KumoScale on a Master with a Static IP: Delete the contents of the system disk and install a new image of KumoScale software on a server that will be used as the master node of the storage cluster with VIP 192.0.2.0. Deploy across a RAID created based on KumoScale software discovery.

Installation Menu options:

Mode=Deploy
Raid Mode=Auto
IP Mode=Static
Static IP= 192.0. ###.##
Subnet Mask=192.0. ##.##
Default Gateway =###.##.##
Control Plane Endpoint=192.0.2.0

Bypassing Install Menu and using Boot Menu:

```
menuentry 'Install-KumoScale-static-master' {
linuxefi ks/vmlinuz-3.22-15549 ro copytoram=y panic=30 ip=:::::dhcp ipv6.disable=1 biosdevname=0 net.ifnames=0
quiet root=live: http://###.##.##/files/Install/KumoScale/3.22_golden/kumoscale-3.22-15757-img.raw
rd.driver.pre=loaop rd.writable.fsimg=1 rd.neednet=1 rd.live.image rd.live.ram=0 kx_static_ip= 192.0.x.xx
kx_subnet_mask= 192.0.x.xx kx_gateway= ###.##.## kx_cluster_vip=192.0.2.0 kx_mode=deploy kx_raid=auto
initrdefi ks/initrd-3.22-15549.img
}
```

Install KumoScale on a Single Storage Node: Delete the contents of the system disk and install a new image of KumoScale software on a storage node. Deploy over a single drive.

Installation Menu options:

Mode=Deploy
Raid=Single
(take default values for all else)

Bypassing Installation Menu and using Boot Menu:

```
menuentry 'Install-KumoScale-single' {
linuxefi ks/vmlinuz-3.22-15757 ro copytoram=y panic=30 ip=:::::dhcp ipv6.disable=1 biosdevname=0 net.ifnames=0
quiet root=live: http://###.##.##/files/Install/KumoScale/3.22_golden/kumoscale-3.22-15757-img.raw
rd.driver.pre=loaop rd.writable.fsimg=1 rd.neednet=1 rd.live.image rd.live.ram=0 kx_mode=deploy kx_raid=single
initrdefi ks/initrd-3.22-15757.img
}
```

Install KumoScale on a storage node using Autodiscovery: Delete the contents of the system disk and install a new image of KumoScale software on a storage node over https. Deploy across a RAID created based on KumoScale software discovery.

Installation Menu options:

Mode=Deploy
Raid Mode=Auto
(take default values for all else)

Bypassing Install Menu and using Boot Menu:

```
menuentry 'Install-KumoScale-auto' {
linuxefi ks/vmlinuz-3.22-15757 ro copytoram=y panic=30 ip=:::::dhcp ipv6.disable=1 biosdevname=0 net.ifnames=0
quiet root=live: http://###.##.##/files/Install/KumoScale/3.22_golden/kumoscale-3.22-15757-img.raw
rd.driver.pre=loaop rd.writable.fsimg=1 rd.neednet=1 rd.live.image rd.live.ram=0 --noverifyssl kx_mode=deploy
kx_raid=auto initrdefi ks/initrd-3.22-15757.img
}
```

Install KumoScale across a RAID using NVMe: Delete the contents of the system disk and install a new image of KumoScale software on a storage node. Deploy across a RAID using only NVMe, based on KumoScale software discovery.

Installation Menu options:

Mode=Deploy
Raid Mode=Auto
Device Type=Nvme
(take default values for all else)

Bypassing Install Menu and using Boot Menu:

```
menuentry 'Install-KumoScale-auto-NVMe )' {
linuxefi ks/vmlinuz-3.22-15757 ro copytoram=y panic=30 ip=:::::dhcp ipv6.disable=1 biosdevname=0 net.ifnames=0
quiet root=live: http://###.##.##/files/Install/KumoScale/3.22_golden/kumoscale-3.22-15757-img.raw
rd.driver.pre=loaop rd.writable.fsimg=1 rd.neednet=1 rd.live.image rd.live.ram=0 kx_mode=deploy kx_raid=auto
kx_blocktype=nvme initrdefi ks/initrd-3.22-15757.img
}
```

Install KumoScale and Purge NVMe: Delete the contents of the drives associated with the PCI slots
0000:40:01.1/0000:41:00.0/0000:42:02.0/0000:45:00.0.

Installation Menu options:

Mode=Purge
Device Type=Nvme
Boot Devices= 0000:40:01.1/0000:41:00.0/0000:42:02.0/0000:45:00.0
(take default values for all else)

Bypassing Install Menu and using Boot Menu:

```
menuentry 'Install- KumoScale-purge-NVMe' {
linuxefi ks/vmlinuz-3.22-15757 ro copytoram=y panic=30 ip=:::::dhcp ipv6.disable=1 biosdevname=0 net.ifnames=0
quiet root=live: http://###.##.##/files/Install/KumoScale/3.22_golden/kumoscale-3.22-15757-img.raw
rd.driver.pre=loaop rd.writable.fsimg=1 rd.neednet=1 rd.live.image rd.live.ram=0 kx_mode=purge kx_blocktype=nvme
kx_devices=0000:40:01.1/0000:41:00.0/0000:42:02.0/0000:45:00.0
initrdefi ks/initrd-3.22-15757.img
}
```

Install KumoScale and Purge SATA: Delete the contents of SATA SSDs.

Installation Menu options:

Mode=Purge

Boot Devices= host9

(take default values for all else)

Bypassing Install Menu and using Boot Menu:

```
menuentry 'Install- KumoScale-purge-SATA' {
linuxefi ks/vmlinuz-3.22-15757 ro copytoram=y panic=30 ip=:::::dhcp ipv6.disable=1 biosdevname=0 net.ifnames=0
quiet root=live: http://###.##.##/files/Install/KumoScale/3.22_golden/kumoscale-3.22-15757-img.raw
rd.driver.pre=loaop rd.writable.fsimg=1 rd.neednet=1 rd.live.image rd.live.ram=0 kx_mode=purge kx_devices=host9
initrdefi ks/initrd-3.22-15757.img
}
```

Next: [Step 2: Configure the boot order.](#)

[1] Definition of capacity - KIOXIA Corporation defines a megabyte (MB) as 1,000,000 bytes, a gigabyte (GB) as 1,000,000,000 bytes and a terabyte (TB) as 1,000,000,000,000 bytes. A computer operating system, however, reports storage capacity using powers of 2 for the definition of 1Gbit = 230 bits = 1,073,741,824 bits, 1GB = 230 bytes = 1,073,741,824 bytes and 1TB = 240 bytes = 1,099,511,627,776 bytes and therefore shows less storage capacity. Available storage capacity (including examples of various media files) will vary based on file size, formatting, settings, software and operating system, and/or pre-installed software applications, or media content. Actual formatted capacity may vary.
