

OpenStack Integration Installation and Configuration

This section describes the installation and configuration requirements to bring up the KumoScale software. It is assumed that OpenStack is already installed.

Environment Requirements

1. Verify that your environment includes the following:
- A controller node.
 - A compute node.
 - A network node (optional).
 - All nodes run CentOS™ 8.4 and Ubuntu™ 20.04.
2. Verify that you have the following packages installed:
- `nvme-cli`
 - `mdadm`
 - `dmidecode`
 - `util-linux`
3. Verify that OpenStack platform version **Yoga** is installed.
4. Verify that KumoScale software is installed, the management cluster is configured, NVMe™ modules are loaded, and KumoScale storage nodes have been added to the KumoScale Provisioner service.

Installing the KumoScale Agent, NVMe-oF Connector and Cinder Driver

The KumoScale NVMe-oF connector and agent are required for both compute nodes (Nova) and volume hosts (Cinder). The NVMe-oF suite for OpenStack contains installation scripts for both as noted in the following sections. The archive file **install-kioxia-nvmeof-<version>.tar** contains the installation for the NVMe-oF connector and agent and **install-kioxia-cinder-<version>.tar** contains the Cinder driver.

Install and Configure the NVMe-OF Connector and Agent

1. Extract the tar file **install-kioxia-nvmeof-<version>.tar** available from the NVMe-oF suite for OpenStack.
2. Confirm **nvme-cli** is installed as documented in [Environment Requirements](#)
3. Run the extracted script: **install-kioxia-nvmeof-<version>.sh**

```
./install_kioxia_nvmeof_<version>.sh
```

You will be prompted for the value of **OS_BRICK_PATH**: Provide the location. For example,

RHEL/CentOS, **/usr/lib/python3.6/site-packages**

Ubuntu 20.04, **/usr/lib/python3/dist-packages**

4. Configure the agent by updating parameter values in the KumoScale Provisioner configuration file, **/etc/kioxia/provisioner.conf** as noted below.
- Provide values for **prov_ip**, **prov_port** **cert_file** .
 - For LOCAL authorization mode only, provide a value for
 - For Open ID Connection (OpenIDC or OIDC) authentication mode, provide values for **client_id**, **client_secret** , **client_scope** (optional), and **token_url**. In any other case, remove the entry or specify **None**. These values are case sensitive so confirm that you use **None** rather than **none** or **NONE**.

Below is an example of the contents of a valid **provisioner.conf** file for LOCAL authorization mode:

```
prov_ip=###.##.##
prov_port=30100
token=*****
cert_file='/etc/kioxia/ssdtoolbox.pem'
```

Install and Configure the Cinder Driver

1. Extract the tar file **install-kioxia-cinder-<version>.tar** available from the NVMe-oF suite for OpenStack.
2. Install the Cinder driver by running the extracted script **install-kioxia-cinder-<version>.sh**:

```
./install_kioxia_cinder_<version>.sh
```

You will be prompted for the value of **CINDER_PATH**. Provide the location. For example,

RHEL/CentOS: **/usr/lib/python3.6/site-packages**

Ubuntu: **/usr/lib/python3/dist-packages**

3. Configure the Cinder driver by editing the file **/etc/cinder/cinder.conf** For example, to configure the backend edit the file with the following information:

- a. Under the [DEFAULT] section, set the value of the **enabled_backends** For example:

```
[DEFAULT]

enabled_backends = kumoscale-1
```

- b. Add a section for the storage node group specified in the **enabled_backends**
c. In the newly created storage node group section, set the following configuration options:

```
[kumoscale-1]
# Storage node (backend) name
volume_backend_name=kumoscale-1
# The driver path
volume_driver=cinder.volume.drivers.kioxia.kumoscale.KumoScaleBaseVolumeDriver
# Kumoscale provisioner URL
kioxia_url=https://##.##.##.##:30100
# Kumoscale provisioner cert file
kioxia_cafile=/etc/kioxia/ssdtoolbox.pem
# Kumoscale provisioner token
kioxia_token=#####
```

- d. Restart the service cinder/nova and verify that the **pem** file exists in **/etc/kioxia/**
e. Run the following in the controller system to create a new type:

```
openstack volume type create kumoscale-1
```

- f. Run the following in the controller system:

```
openstack volume type set --property volume_backend_name=kumoscale-1 kumoscale-1
```

The following operations are supported:

- Create, list, delete, attach, and detach volumes
- Create, list, and delete volume snapshots
- Create a volume from a snapshot
- Copy an image to a volume.
- Copy a volume to an image.
- Create volume from snapshot
- Clone a volume
- Extend a volume

[KumoScale Configuration Options](#) lists the parameters that may be used to complete the above actions.

KumoScale Cinder Driver Configuration Options

The following table contains the configuration options supported by the KIOXIA KumoScale Cinder driver.

Description of KIOXIA KumoScale configuration options	
Configuration option = Default value	Description
kioxia_block_size = 4096	(Integer) Volume block size in bytes - 512 or 4096 (Default).
kioxia_cafile = None	(String) Cert for provisioner REST API SSL
kioxia_client_id	The ID for the client, which has a service account role of ADMIN. Relevant only for OpenIDC authentication mode
kioxia_client_secret	The client secret. Relevant only for OpenIDC authentication mode
kioxia_client_scope	The scope (role) of the client defined in the Provisioner. Relevant only for OpenIDC authentication mode under Active Directory Federation Services (ADFS).
kioxia_desired_bw_per_gb = 0	(Integer) Desired bandwidth in B/s per GB.
kioxia_desired_iops_per_gb = 0	(Integer) Desired IOPS/GB.
kioxia_max_bw_per_gb = 0	(Integer) Upper limit for bandwidth in B/s per GB.
kioxia_max_iops_per_gb = 0	(Integer) Upper limit for IOPS/GB.
kioxia_max_replica_down_time = 0	(Integer) Replicated volume max downtime for replica in minutes Default of 0 indicates forever. Otherwise, must be a value between 5 and 1440 (24 hours).

Configuration option = Default value	Description
kioxia_num_replicas = 1	(Integer) Number of volume replicas.
kioxia_provisioning_type = THICK	(String(choices=['THICK', 'THIN'])) Specify whether to use a thin or thick volume.
kioxia_same_rack_allowed = False	(Boolean) Can more than one replica be allocated to same rack.
kioxia_snap_reserved_space_percentage = 0	(Integer) Percentage of the parent volume to be used for log.
kioxia_snap_vol_reserved_space_percentage = 0	(Integer) Writable snapshot percentage of parent volume used for log.
kioxia_snap_vol_span_allowed = True	(Boolean) Allow span in snapshot volume - Default: True.
kioxia_span_allowed = True	(Boolean) Allow span – Default: True.
kioxia_token = None	(String) KumoScale Provisioner authorization token. Relevant only for OpenIDC authentication mode.
kioxia_url = None	(String) KumoScale Provisioner full URL. A valid URL must be provided.
kioxia_vol_reserved_space_percentage = 0	(Integer) Thin volume reserved capacity ^[1] allocation percentage.
kioxia_writable = False	(Boolean) Specify whether snapshot volumes from this class are writable or not.
volume_backend_name	The name of the KumoScale storage node to be identified by Cinder.
volume_driver	Full path and name of the KumoScale Volume Driver.

Multipath Configuration for OpenStack Yoga

KumoScale OpenStack supports MPIO as implemented in yoga. To add support for MPIO, you will need to edit the file `/etc/nova/nova-compute.conf` and set the `volume_use_multipath` parameter to `True` as follows:

```
[libvirt]
volume_use_multipath = True
```

For more information on MPIO support in yoga, see <https://specs.openstack.org/openstack/cinder-specs/specs/yoga/nvme-multipath.html>.

You should also confirm that your kernel supports Asymmetric Namespace Access (ANA) multipath by verifying that the value of multipath in `/sys/module/nvme_core/parameters/multipath` is `Y`.

Maintenance of NVMe-oF Connector and Cinder Driver

This section documents how to uninstall and upgrade KumoScale software for the Nova host and Cinder driver.

Uninstalling the NVMe-oF Connector and Agent

RHEL/CentOS:

```
sudo rpm -ev nvmeof-agent-${RELEASE}-1.x86_64
```

Ubuntu:

```
sudo dpkg -r nvmeof-agent
```

Upgrading the NVMe-oF Connector and Agent

RHEL/CentOS:

```
rpm -Uvh kioxia-nvmeof-${RELEASE}-1.x86_64.rpm
```

Ubuntu:

```
sudo dpkg --install kioxia-nvmeof-${RELEASE}_amd64.deb
```

Uninstalling the Cinder Driver

RHEL/CentOS:

```
sudo rpm -ev kioxia-cinder-${RELEASE}-1.x86_64
```

Ubuntu:

```
sudo dpkg -r kioxia-cinder
```

Upgrading the Cinder Driver

RHEL/CentOS:

```
rpm -Uvh kioxia-cinder- $\${RELEASE}$ -1.x86_64.rpm
```

Ubuntu:

```
sudo dpkg --install kioxia-cinder- $\${RELEASE}$ _amd64.deb
```

Changing the Truststores

To change the truststore used with KumoScale for Open Stack, you will need to provide the location for the new Privacy Enhanced Mail (PEM) file for both the NVMe-OF agent and Cinder.

Set the OpenStack NVMe-oF Agent Truststore

To set the Truststore for the NVMe-OF agent:

1. Define the location of the certificate (pem) file in `/etc/kioxia/provisioner.conf` file:

```
cert_file= path to truststore pem file
```

2. Restart the OpenStack Nova Compute service:

```
systemctl restart nova-compute
```

Set the OpenStack Cinder Truststore

To set the Truststore for Cinder:

1. Define the below backends property in `/etc/cinder/cinder.conf` file:

```
kioxia_cafile= path to truststore pem file
```

2. Restart the OpenStack Cinder service:

```
systemctl restart cinder-volume
```

Live Migration Support

The following instructions are for deployments with QEMU™ or KVM hypervisors. Different deployments may require different configurations according to the hypervisor.

Note: We advise using physical servers as compute nodes, since if one or more of the compute nodes are VMs, the result of the live migration may be unknown and cause the cluster to crash.

1. Configure live migration according to the OpenStack platform user guide available at [OpenStack Yoga Project user guides](#).
2. If the compute initiator is a VMware ESX® hypervisor, enable "Expose hardware assisted virtualization to the guest OS" in the CPU settings.
3. Verify the instance you intend to migrate has only volumes (including the OS).
4. Edit the configuration file `/etc/nova/nova.conf` . Add the following entry to the [DEFAULT] section:

```
live_migration_flag=VIR_MIGRATE_UNDEFINE_SOURCE,VIR_MIGRATE_PEER2PEER,VIR_MIGRATE_LIVE
```

5. If the target initiator and source have different CPU features, set the following in the [libvirt] section:

```
cpu_mode=custom  
  
cpu_model=[baseline-model]
```

6. Compute the baseline-model according to the instructions in the following page: [Fedora 18 Virtualization Administration Guide](#).
7. Restart Nova.
8. For Ubuntu only, add the following line to `/etc/neutron/plugins/ml2/linuxbridge_agent.ini`: on both servers:

```
[linux_bridge]  
  
physical_interface_mappings = physnet1:<relevant interface name on the actual host>
```

9. Configure **libvirtd** on both servers. The following example configuration is a simple configuration for libvirtd↔libvirtd communication over TCP with no authentication. On both servers:

- a. Remove the remark symbol from the following line in `/etc/sysconfig/libvirtd`:

```
LIBVIRT_ARGS="--listen"
```

- b. For Ubuntu only, add the following line to `/etc/default/libvirtd`:

```
libvirtd_opts="-l"
```

c. For a non-production, vanilla Opentack environment, remove the remark symbol from the following lines in `/etc/libvirt/libvirtd.conf`:

```
listen_tls = 0
listen_tcp = 1
```

Note that these settings are only needed for a non-production, vanilla OpenStack environment. Production environments should already be set up for live migration.

d. Remove the remark symbol from the following line in `/etc/libvirt/libvirtd.conf` and set its value:

```
auth_tcp = "none"
```

10. Change **conf** as follows. Add

```
live_migration_uri=qemu+ssh://root@%s/system
```

Use **virsh** to check **libirtd** connectivity from both systems, for example:

```
virsh -c qemu+ssh://root@blade-1/system
```

Add the following:

```
[service_user]
send_service_user_token = True
auth_type = password
project_domain_name = Default
project_name = service
user_domain_name = Default
password = servicepassword
username = nova
auth_url = http://172.28.30.23:5000
```

11. Restart nova-compute:

```
systemctl restart nova-*
```

12. Restart **libvirtd**.

See Next: [OpenStack Storage Provisioning](#)

[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.

