
Red Hat Ceph Storage Block Devices and OpenStack

Red Hat Ceph Storage integration with Red Hat OpenStack provides scalable and resilient block storage for cloud workloads, supporting virtual machine images, volumes, and guest disks.

Red Hat Ceph Storage integrates with OpenStack distributed block storage backend that scales horizontally and delivers high availability for cloud workloads. In an integrated environment, Ceph Block Devices are consumed by multiple OpenStack services, to store and manage images, volumes, persistent and virtual machine disks efficiently.

This section describes the supported integration models between Ceph and OpenStack, the OpenStack services that can use Ceph block storage, and key considerations to review before configuring the environment.

For detailed deployment guidance and service-specific configuration options, see the [Red Hat OpenStack Platform](#) documentation for your specific OpenStack release.

Ceph and OpenStack integration models

OpenStack can use Ceph Block Devices in the following deployment models:

OpenStack-managed Ceph cluster

In this model, OpenStack creates and manages the Ceph storage cluster as part of the deployment. OpenStack installs and configures Ceph components, including monitors and object storage daemons (OSDs). This approach simplifies deployment by allowing OpenStack to automate Ceph installation and configuration.

Existing Ceph cluster integration

OpenStack connects to a preconfigured Ceph storage cluster and uses existing Ceph monitors. Ceph block devices are configured as a storage backend for OpenStack services. This model is typically used when a Ceph cluster already exists or when storage and cloud infrastructure are managed separately.

Note: The procedures described in this guide assume that you are manually configuring OpenStack to use an existing Ceph cluster rather than using automated deployment tools such as an OpenStack director.

OpenStack services that use Ceph Block Devices

Ceph Block Devices can be used as storage backends for several OpenStack services:

Image storage (Glance)

OpenStack Glance manages virtual machine images, which are treated as immutable objects. When Glance is configured to use Ceph, images are stored in a Ceph Block Device pool and can be efficiently cloned using copy-on-write mechanisms.

Volume storage (Cinder and Cinder Backup)

OpenStack Cinder provides persistent block storage volumes that can be attached to running virtual machines, used for boot-from-volume instances, or backed up using Cinder Backup. Ceph can be used as a backend for both Cinder volumes and Cinder Backup storage.

Guest virtual machine disks (Nova)

Ceph block devices can store guest virtual machine disks directly. Instead of maintaining disk files on the hypervisor file system, Nova accesses disks from Ceph, improving scalability and enabling features such as live migration.

Image format and storage considerations

Ceph does not support the QCOW2 image format for hosting virtual machine disks. To boot virtual machines using Ceph block storage, images configured in Glance must use the RAW format.

This requirement applies to both ephemeral disks and boot-from-volume instances.

OpenStack can use Ceph for any combination of image storage, volume storage, or guest disk storage. There is no requirement to use Ceph for all three services.

Installing and configuring Ceph for OpenStack

Install and configure Red Hat Ceph Storage so that OpenStack can use Ceph Block Devices as a storage backend.

Before you begin, make sure that you have the following prerequisites in place:

- A new or existing Red Hat Ceph Storage cluster is installed and operational.
- Administrative access to the Ceph cluster.
- Access to the OpenStack environment that will use Ceph storage.

Complete all tasks in this section before proceeding to configure OpenStack services to use Ceph Block Devices. The following is a high-level workflow of the tasks required:

1. Install Red Hat Ceph Storage and verify that the cluster reports a healthy status.
For more information, see [“Installing the Ceph client on OpenStack” on page 2](#).
2. Create and configure Ceph storage pools for OpenStack services.
3. Configure authentication and access for OpenStack.
4. Validate connectivity between OpenStack and the Ceph cluster.

Creating Ceph pools for OpenStack

Create Ceph storage pools to store images, volumes, virtual machine disks, and backup data for OpenStack services.

PREREQUISITE

Before you begin, make sure that you have the prerequisites defined in [“Installing and configuring Ceph for OpenStack” on page 2](#) and access to a Ceph Monitor node.

Ceph pools are used by OpenStack services to store different types of data such as images, volumes, guest virtual machine disks, and backups. Creating separate pools helps isolate workloads and allows tuning based on usage patterns.

1. Verify that the Ceph cluster is running and in a healthy state.
Run the following command on a Ceph monitor node:

```
ceph -s
```

Ensure that the cluster reports a **HEALTH_OK** status.

2. Create Ceph pools for OpenStack workloads.
The following example creates pools used by OpenStack services, where 128 represents the number of placement groups (PGs) configured for each pool.

```
[root@rbd-client ~]# ceph osd pool create volumes 128
[root@rbd-client ~]# ceph osd pool create backups 128
[root@rbd-client ~]# ceph osd pool create images 128
[root@rbd-client ~]# ceph osd pool create vms 128
```

Result

Ceph storage pools required for OpenStack services are created and available. These pools can now be configured as backends for Glance, Cinder, and Nova.

AFTER COMPLETING THE TASK

Red Hat recommends calculating an appropriate number of placement groups based on the size, number of OSDs, and workload characteristics of the Ceph cluster.

Installing the Ceph client on OpenStack

Install the Ceph client packages on OpenStack nodes so that OpenStack services can communicate with the Ceph cluster.

PREREQUISITE

Before you begin, make sure that you have the following prerequisites in place:

1. Prerequisites defined in [“Installing and configuring Ceph for OpenStack”](#) on page 2.
2. Access to the Ceph software repositories.
3. Root access to the OpenStack nodes.

OpenStack services require Ceph client libraries to communicate with the Ceph cluster. These packages enable OpenStack components to access and manage Ceph Block Devices by using the Ceph Block Device (RBD) interface.

1. Install the required Ceph client packages on each of the Nova, Cinder, and Cinder Backup OpenStack nodes.

```
dnf install ceph-common python-rbd
```

2. Install the Ceph RBD Python package on the Glance node.

```
dnf install python-rbd
```

Result

The required Ceph client packages are installed on the OpenStack nodes. OpenStack services can now communicate with the Ceph cluster.

Copying the Ceph configuration file to OpenStack

Copy the Ceph configuration file to OpenStack nodes so that they can locate and connect to the Ceph cluster.

PREREQUISITE

Before you begin, make sure that you have the following prerequisites in place:

1. Prerequisites defined in [Installing and configuring Ceph for OpenStack](#).
2. Root access to the Ceph Monitor node.
3. Root access to the OpenStack nodes.

OpenStack services require the Ceph configuration file to identify the Ceph cluster monitors and connection settings. This file must be present on each OpenStack node that communicates with Ceph.

1. Copy the Ceph configuration file from a Ceph monitor node to all OpenStack nodes.

```
scp /etc/ceph/ceph.conf OPENSTACK_NODES:/etc/ceph/
```

Replace OPENSTACK_NODES with the hostnames or IP addresses of the OpenStack nodes.

2. Verify that the Ceph configuration file exists on each OpenStack node.
Confirm that the `/etc/ceph/ceph.conf` file is present and readable.

Result

The Ceph configuration file is available on all required OpenStack nodes. OpenStack services can now locate and connect to the Ceph cluster.

Configuring Ceph client authentication

Configure Ceph authentication so that OpenStack services can securely access Ceph Block Devices.

PREREQUISITE

Before you begin, make sure that you have the following prerequisites in place:

1. Prerequisites defined in [Installing and configuring Ceph for OpenStack](#).

2. Root access to a Ceph Monitor node.
3. Root access to specific OpenStack service nodes (Nova, Cinder, Cinder Backup, Glance).

Ceph uses authenticated clients to control access to storage pools. Before OpenStack services can use Ceph Block Devices, you must create Ceph users, distribute keyrings, and configure the libvirt secret used by Nova.

1. Create Ceph users for OpenStack services.
Run the following commands on a Ceph Monitor node:

```
ceph auth get-or-create client.cinder \
  mon 'allow r' \
  osd 'allow class-read object_prefix rbd_children, allow rwx pool=volumes, allow rwx pool=vms, allow rx pool=images'
```

```
ceph auth get-or-create client.cinder-backup \
  mon 'allow r' \
  osd 'allow class-read object_prefix rbd_children, allow rwx pool=backups'
```

```
ceph auth get-or-create client.glance \
  mon 'allow r' \
  osd 'allow class-read object_prefix rbd_children, allow rwx pool=images'
```

2. Copy the keyrings to the appropriate OpenStack nodes and set ownership.
For the Cinder volume service, run

```
ceph auth get-or-create client.cinder | \
  ssh CINDER_VOLUME_NODE sudo tee /etc/ceph/ceph.client.cinder.keyring

ssh CINDER_VOLUME_NODE sudo chown cinder:cinder \
  /etc/ceph/ceph.client.cinder.keyring
```

Repeat this step for `client.cinder-backup` and `client.glance` on their respective OpenStack nodes.

3. Copy the Cinder keyring to each Nova compute node.

```
ceph auth get-or-create client.cinder | \
  ssh NOVA_NODE sudo tee /etc/ceph/ceph.client.cinder.keyring
```

4. Retrieve the Cinder secret key on a Nova compute node.

```
ceph auth get-key client.cinder > client.cinder.key
```

5. Generate a UUID for the libvirt secret.

```
uuidgen > uuid-secret.txt
```

6. Define and set the Ceph secret in libvirt.

```
cat > secret.xml <<EOF
<secret ephemeral='no' private='no'>
  <uuid>$(cat uuid-secret.txt)</uuid>
  <usage type='ceph'>
    <name>client.cinder secret</name>
  </usage>
</secret>
EOF
```

```
virsh secret-define --file secret.xml

virsh secret-set-value \
  --secret $(cat uuid-secret.txt) \
  --base64 $(cat client.cinder.key)
```

7. Remove the temporary files.

```
rm -f client.cinder.key secret.xml
```

Result

Ceph client authentication is configured for OpenStack services. OpenStack components can now authenticate securely and access Ceph Block Devices.

Configuring OpenStack to use Ceph Block Devices

Configure OpenStack services to use Ceph block devices as a scalable and resilient storage backend after the Ceph cluster is prepared.

As a storage administrator, you must configure OpenStack services to use Ceph Block Devices after installing and preparing the Ceph cluster. This configuration enables OpenStack components to consume Ceph storage for images, volumes, backups, and virtual machine disks.

This section describes the tasks required to configure OpenStack components to integrate with Ceph block storage. Complete all tasks in this section only after Ceph is fully installed and configured for OpenStack.

Before you begin, make sure that you have the following prerequisites in place:

- Red Hat Ceph Storage is installed and configured for use with OpenStack.
- Ceph storage pools required for OpenStack service integration are created.
- Ceph client authentication is configured.
- A running OpenStack environment is available.

Configuring Cinder to use Ceph Block Devices

Configure OpenStack Cinder to use Ceph Block Devices as the backend for volume storage.

PREREQUISITE

Before you begin, make sure that you have the following prerequisites in place:

1. Prerequisites defined in [“Installing and configuring Ceph for OpenStack” on page 2](#).
2. Root access to the Cinder node.
3. A Ceph pool is created for storing Cinder volumes.
4. The Ceph user and UUID secret required for Cinder are configured.

Cinder uses a backend driver to provision and manage block storage volumes. When configured to use Ceph, Cinder stores volumes as Ceph Block Devices (RBDs) in a designated Ceph pool.

1. Edit the Cinder configuration file.
For example, to open the configuration file on the Cinder node:

```
vi /etc/cinder/cinder.conf
```

2. In the [DEFAULT] section, enable Ceph as the backend.
Add or update the following configuration options, as required.

```
enabled_backends = ceph  
glance_api_version = 2
```

3. Create a [ceph] section and configure the Ceph volume driver.

```
[ceph]  
volume_driver = cinder.volume.drivers.rbd.RBDDriver  
rbd_cluster_name = ceph  
rbd_pool = volumes  
rbd_user = cinder  
rbd_ceph_conf = /etc/ceph/ceph.conf  
rbd_secret_uuid = UUID  
rbd_flatten_volume_from_snapshot = false  
rbd_max_clone_depth = 5  
rbd_store_chunk_size = 4  
rados_connect_timeout = -1
```

Replace *UUID* with the UUID of the libvirt secret configured for the Cinder service.

4. Save the file and exit the editor.

Result

Cinder is configured to use Ceph Block Devices as its volume backend. Volumes can now be created and managed by using Ceph storage.

AFTER COMPLETING THE TASK

Restart the OpenStack service to complete the configuration. For more information, see [“Restarting the OpenStack services” on page 8](#).

Configuring Cinder Backup to use Ceph Block Devices

Configure OpenStack Cinder Backup to store volume backup data in Ceph Block Storage.

PREREQUISITE

Before you begin, make sure that you have the following prerequisites in place:

- Root access to the Cinder node.
- A Ceph pool is created for storing backup data.

Cinder Backup enables backup and restore of Cinder volumes. When configured to use Ceph, backup data is stored as RADOS objects in a designated Ceph pool.

1. Edit the Cinder configuration file.
For example, to open the configuration file on the Cinder node:

```
vi /etc/cinder/cinder.conf
```

2. Add the Cinder Backup configuration to the [ceph] section of the configuration file.

```
backup_driver = cinder.backup.drivers.ceph
backup_ceph_conf = /etc/ceph/ceph.conf
backup_ceph_pool = backups
backup_ceph_user = cinder-backup
backup_ceph_chunk_size = 134217728
backup_ceph_stripe_unit = 0
backup_ceph_stripe_count = 0
restore_discard_excess_bytes = true
```

3. Verify that the Cinder Backup service is enabled in the OpenStack dashboard configuration.

Result

Cinder Backup is configured to use Ceph Block Devices for storing backup data. Volume backups can now be created and managed by using the Ceph backend.

AFTER COMPLETING THE TASK

Restart the OpenStack service to complete the configuration. For more information, see [“Restarting the OpenStack services” on page 8](#).

Configuring Glance to use Ceph Block Devices

Configure OpenStack Glance to store virtual machine images in Ceph block storage.

PREREQUISITE

Before you begin, make sure that you have the following prerequisites in place:

- Root access to the Glance API node.
- A Ceph pool is created for storing images.

Glance manages virtual machine images in OpenStack. When configured to use Ceph, images are stored as Ceph Block Devices (RBDs), which enables efficient cloning and improved performance through copy-on-write operations.

1. Edit the Glance API configuration file.
For example, to open the configuration file on the Glance API node:

```
vi /etc/glance/glance-api.conf
```

2. Configure Glance to use Ceph as the default image store.

```
stores = rbd
default_store = rbd
rbd_store_pool = images
rbd_store_user = glance
rbd_store_ceph_conf = /etc/ceph/ceph.conf
rbd_store_chunk_size = 8
```

3. Enable copy-on-write image cloning.

```
show_image_direct_url = True
```

4. If enabled, disable cache management.
Ensure the following option is set:

```
flavor = keystone
```

Result

Glance is configured to use Ceph Block Devices as the backend for virtual machine images. Images can now be stored, cloned, and managed efficiently by using Ceph.

AFTER COMPLETING THE TASK

Restart the OpenStack service to complete the configuration. For more information, see [“Restarting the OpenStack services”](#) on page 8.

Configuring Nova to use Ceph Block Devices

Configure OpenStack Nova to store virtual machine disks in Ceph block storage.

PREREQUISITE

Before you begin, make sure that you have the following prerequisites in place:

- Root access to all Nova compute nodes.
- Ceph authentication is configured for Nova.

Nova uses the libvirt driver to manage virtual machine disks. When configured to use Ceph, Nova stores instance disks as Ceph Block Devices (RBDs), which improves scalability and enables features such as live migration.

1. Edit the Ceph configuration file on the Nova compute nodes.
For example, to open the configuration file on the Nova compute nodes:

```
vi /etc/ceph/ceph.conf
```

- a. Add the required client configuration settings.

```
[client]
rbd cache = true
rbd cache writethrough until flush = true
rbd concurrent management ops = 20
admin socket = /var/run/ceph/guests/$cluster-$type.$id.$pid.$cctid.asok
log file = /var/log/ceph/qemu-guest-$pid.log
```

2. Create the required directories and set permissions.
Run the following commands on each Nova compute node:

```
mkdir -p /var/run/ceph/guests /var/log/ceph
chown qemu:libvirt /var/run/ceph/guests /var/log/ceph
```

3. Edit the Nova configuration file.
For example, to open the configuration file on the Nova compute nodes:

```
vi /etc/nova/nova.conf
```

- a. Configure the [libvirt] section to use Ceph.
Ensure the following options are present:

```
[libvirt]
images_type = rbd
images_rbd_pool = vms
images_rbd_ceph_conf = /etc/ceph/ceph.conf
rbd_user = cinder
rbd_secret_uuid = UUID
disk_cachemodes = network=writeback
hw_disk_discard = unmap
```

Replace *UUID* with the UUID of the Ceph secret configured for Nova.

Result

Nova is configured to store virtual machine disks in Ceph Block Devices. New instances can now use Ceph-backed disks, enabling scalable and resilient compute storage.

AFTER COMPLETING THE TASK

Restart the OpenStack service to complete the configuration. For more information, see [“Restarting the OpenStack services”](#) on page 8.

Restarting the OpenStack services

Restart OpenStack services to apply the Ceph configuration changes.

PREREQUISITE

Before you begin, make sure that you have completed the configuration of all OpenStack services that use Ceph, as detailed in [Installing and configuring Ceph for OpenStack](#).

1. You have completed configuration of all OpenStack services that use Ceph.

After configuring OpenStack services to use Ceph Block Devices, the services must be restarted so that the configuration changes take effect.

1. Restart the OpenStack services by running the following commands on each of the OpenStack nodes.

```
systemctl restart openstack-cinder-volume
systemctl restart openstack-cinder-backup
systemctl restart openstack-glance-api
systemctl restart openstack-nova-compute
```

2. Verify that all services start successfully.
Check the service status and logs to confirm that no errors are reported.

Result

All OpenStack services are restarted successfully, and the Ceph configuration changes are active.