

OpenStack: How do I integrate Bright OpenStack with an external Ceph storage?

How do I add External Ceph to Bright OpenStack?

A local Ceph can be set up on Bright clusters with the `cm-ceph-setup` script and used as a backend storage for Bright OpenStack.

An external Ceph, on the other hand, can be integrated with Bright OpenStack, by first deploying Bright OpenStack, and then changing the backend storage for different OpenStack services. The following steps can be followed:

The first step in external Ceph integration is to make sure that the external Ceph storage can be reached and queried from the Bright cluster. This can be done by installing the required Ceph packages and copying the required configuration and key files as follows:

1. Install the Ceph packages on the head node and in the software images:

```
# yum install ceph-common python-rbd
```

```
# yum install ceph-common python-rbd --installroot=/cm/images/default-image
```

** Note: the `default-image` can be substituted by the relevant image name.

2. Copy the `ceph.conf` configuration file and the admin keyring from the external Ceph cluster to the Bright cluster head node and software images. The "`/etc/ceph/ceph.conf`" and "`/etc/ceph/ceph.client.admin.keyring`" will be required.

3. Verify that the Ceph cluster can be queried from the Bright cluster:

```
# ceph -s
```

```
# ceph osd dump
```

OpenStack: How do I integrate Bright OpenStack with an external Ceph storage?

After connecting the Bright cluster with the Ceph cluster, the required OpenStack pools for Bright OpenStack need to be created. This can be done as follows:

Note: The shorewall rules on the head node may need to be adjusted to allow traffic from the external Ceph storage.

1. Cinder

Generate Ceph UUID

```
[root@ma-b72-c7 ~]# ceph
```

```
[ma-b72-c7]# openstack settings
```

```
[ma-b72-c7->openstack[default]->settings]# advanced
```

```
[ma-b72-c7->openstack[default]->settings->advanced]# get rbdsecretuuid
```

```
73db5161-4a91-4be0-8cc0-41049578b7a1
```

```
[ma-b72-c7->openstack[default]->settings->advanced]#
```

Setup Ceph pool

a. Compute number of PGS: 16

This has to be done according to the number of placement groups in the external Ceph. The “ceph -s” command should show the number of available PGS. The admin can then decide the correct number of PGS to be used.

b. Delete current osd pool if it exists:

```
[root@ma-b72-c7 ~]# ceph osd pool delete openstack_volumes openstack_volumes --yes-i-really-really-mean-it
```

```
pool 'openstack_volumes' does not exist
```

```
[root@ma-b72-c7 ~]# ceph osd pool delete openstack_volume_backups openstack_volume_backups --yes-i-really-really-mean-it
```

OpenStack: How do I integrate Bright OpenStack with an external Ceph storage?

```
pool 'openstack_volume_backups' does not exist
```

c. Create osd pool

```
[root@ma-b72-c7 ~]# ceph osd pool create openstack_volumes 16
```

```
pool 'openstack_volumes' created
```

```
[root@ma-b72-c7 ~]# ceph osd pool create openstack_volume_backups 16
```

```
pool 'openstack_volume_backups' created
```

```
[root@ma-b72-c7 ~]# ceph auth del client.cinder
```

```
entity client.cinder does not exist
```

```
[root@ma-b72-c7 ~]# ceph auth get-or-create client.cinder mon 'allow r' osd 'allow class-read object_prefix rbd_children, allow rwx pool=openstack_volumes, allow rwx pool=openstack_volume_backups, allow rwx pool=openstack_vms, allow rwx pool=openstack_images'
```

```
[client.cinder]
```

```
key = AQC3GatXs/xSFxAAbMLvvdmqZqIAUfUTuzDirA==
```

```
[root@ma-b72-c7 ~]# cat > /etc/ceph/ceph.cinder.keyring
```

```
[client.cinder]
```

```
key = AQC3GatXs/xSFxAAbMLvvdmqZqIAUfUTuzDirA==
```

d. Create volume backend:

```
[root@ma-b72-c7 ~]# ceph
```

```
[ma-b72-c7]# configurationoverlay
```

```
[ma-b72-c7->configurationoverlay]# use openstackcontrollers
```

```
[ma-b72-c7->configurationoverlay[OpenStackControllers]]# roles
```

```
[ma-b72-c7->configurationoverlay[OpenStackControllers]->roles]# use openstack::volume
```

OpenStack: How do I integrate Bright OpenStack with an external Ceph storage?

```
[ma-b72-c7->configurationoverlay[OpenStackControllers]->roles[OpenStack::Volume]->volumebackends]% add ceph ceph
```

```
[ma-b72-c7->configurationoverlay*[OpenStackControllers*]->roles*[OpenStack::Volume*]->volumebackends*[ceph*]]% set rbdpool openstack_volumes
```

```
[root@ma-b72-c7 ~]# cmsk [ma-b72-c7]% configurationoverlay [ma-b72-c7->configurationoverlay]% use openstackcontrollers
```

```
[ma-b72-c7->configurationoverlay[OpenStackControllers]]% roles  
[ma-b72-c7->configurationoverlay[OpenStackControllers]->roles]% use openstack::volumebackup
```

```
[ma-b72-c7->configurationoverlay[OpenStackControllers]->roles[OpenStack::VolumeBackup]]% backupbackends
```

```
[ma-b72-c7->configurationoverlay[OpenStackControllers]->roles[OpenStack::VolumeBackup]->backupbackends]% add ceph ceph
```

e. Store Ceph UUID and Setup Ceph Secret

```
[root@ma-b72-c7 ~]# ceph auth get-key client.cinder >  
/etc/libvirt/secrets/73db5161-4a91-4be0-8cc0-41049578b7a1.base64.tmp
```

```
[root@ma-b72-c7 ~]# ceph auth get-key client.cinder >  
/cm/images/default-image/etc/libvirt/secrets/73db5161-4a91-4be0-8cc0-41049578b7a1.base64.tmp
```

```
[root@ma-b72-c7 ~]# perl -pe 'chomp if eof' /etc/libvirt/secrets/73db5161-4a91-4be0-8cc0-41049578b7a1.base64.tmp >  
/etc/libvirt/secrets/73db5161-4a91-4be0-8cc0-41049578b7a1.base64
```

```
[root@ma-b72-c7 ~]# perl -pe 'chomp if eof'  
/cm/images/default-image/etc/libvirt/secrets/73db5161-4a91-4be0-8cc0-41049578b7a1.base64.tmp  
/cm/images/default-image/etc/libvirt/secrets/73db5161-4a91-4be0-8cc0-41049578b7a1.base64
```

```
[root@ma-b72-c7 ~]# chmod 600 /etc/libvirt/secrets/73db5161-4a91-4be0-8cc0-41049578b7a1.xml
```

```
[root@ma-b72-c7 ~]# chmod 600  
/cm/images/default-image/etc/libvirt/secrets/73db5161-4a91-4be0-8cc0-41049578b7a1.xml
```

f. If you are using multiple cinder backends, you will need to create volume types and assign it to a backend :

```
#openstack volume type create ceph
```

```
#openstack volume type set --property volume_backend_name=ceph ceph
```

OpenStack: How do I integrate Bright OpenStack with an external Ceph storage?

If the second storage backend is NFS :

```
#openstack volume type create <nfs>
```

```
#openstack volume type set --property volume_backend_name=nfs nfs
```

To create a ceph backed volume :

```
#openstack volume create cephVolume --size 100 --type ceph --description " A ceph volume "
```

To validate the new volume :

```
#rbd ls -p openstack_volumes
```

2. Glance

Setup Ceph Pool

a. Compute number of PGS

This has to be done according to the number of placement groups in the external Ceph. The “`ceph -s`” command should show the number of available PGS. The admin can then decide the correct number of PGS to be used.

b. Delete current osd pool if it exists:

```
[root@ma-b72-c7 ~]# ceph osd pool delete openstack_images openstack_images --yes-i-really-really-mean-it pool 'openstack_images' does not exist
```

c. Create osd pool

```
[root@ma-b72-c7 ~]# ceph osd pool create openstack_images 16
```

OpenStack: How do I integrate Bright OpenStack with an external Ceph storage?

```
pool 'openstack_images' created
```

```
[root@ma-b72-c7 ~]# ceph auth del client.glance
```

```
entity client.glance does not exist
```

```
[root@ma-b72-c7 ~]# ceph auth get-or-create client.glance mon 'allow r' osd 'allow class-read object_prefix rbd_children, allow rwx pool=openstack_images'
```

```
[client.glance]
```

```
key = AQD4GqtXLYZKAhAA/cEEweHueCEosuwQEcx9dQ==
```

```
[root@ma-b72-c7 ~]# cat > /etc/ceph/ceph.glance.keyring
```

```
[client.glance]
```

```
key = AQD4GqtXLYZKAhAA/cEEweHueCEosuwQEcx9dQ==
```

d. Create image backend

```
[root@ma-b72-c7 ~]# cmsh
```

```
[ma-b72-c7]% configurationoverlay
```

```
[ma-b72-c7->configurationoverlay]% use openstackcontrollers
```

```
[ma-b72-c7->configurationoverlay[OpenStackControllers]]% roles
```

```
[ma-b72-c7->configurationoverlay[OpenStackControllers]->roles]% use openstack::imageapi
```

```
[ma-b72-c7->configurationoverlay[OpenStackControllers]->roles[OpenStack::ImageApi]]% imagebackends
```

```
[ma-b72-c7->configurationoverlay[OpenStackControllers]->roles[OpenStack::ImageApi]->imagebackends]% add ceph  
ceph
```

```
[ma-b72-c7->configurationoverlay[OpenStackControllers]->roles[OpenStack::ImageApi]->imagebackends[ceph]]% set  
rbdstorepool openstack_images
```

```
[ma-b72-c7->configurationoverlay*[OpenStackControllers*]->roles*[OpenStack::ImageApi*]->imagebackends*[ceph*]]%  
commit
```

3. Nova

OpenStack: How do I integrate Bright OpenStack with an external Ceph storage?

Setup Ceph Pool

a. Compute number of PGS

This has to be done according to the number of placement groups in the external Ceph. The “ceph -s” command should show the number of available PGS. The admin can then decide the correct number of PGS to be used.

b. Delete current osd pool if it exists:

```
[root@ma-b72-c7 ~]# ceph osd pool delete openstack_vms openstack_vms --yes-i-really-really-mean-it
```

```
pool 'openstack_vms' does not exist
```

c. Create osd pool

```
[root@ma-b72-c7 ~]# ceph osd pool create openstack_vms 16
```

```
pool 'openstack_vms' created
```

d. Create image backend

```
[root@ma-b72-c7 ~]# cmsn
```

```
[ma-b72-c7]# configurationoverlay
```

```
[ma-b72-c7->configurationoverlay]# use openstackhypervisors
```

```
[ma-b72-c7->configurationoverlay[OpenStackHypervisors]]# roles
```

```
[ma-b72-c7->configurationoverlay[OpenStackHypervisors]->roles]# use openstack::compute
```

```
[ma-b72-c7->configurationoverlay[OpenStackHypervisors]->roles[OpenStack::Compute]]# imagebackends
```

```
[ma-b72-c7->configurationoverlay[OpenStackHypervisors]->roles[OpenStack::Compute]->imagebackends]# add ceph
```

```
[ma-b72-c7->configurationoverlay[OpenStackHypervisors]->roles[OpenStack::Compute]->imagebackends[ceph*]]# set rbdpool openstack_vms
```

```
[ma-b72-c7->configurationoverlay[OpenStackHypervisors]->roles[OpenStack::Compute]->imagebackends[ceph*]]#
```

OpenStack: How do I integrate Bright OpenStack with an external Ceph storage?

```
set rbdsecretuuid 73db5161-4a91-4be0-8cc0-41049578b7a1
```

```
[ma-b72-c7->configurationoverlay[OpenStackHypervisors]->roles[OpenStack::Compute]->imagebackends[ceph*]]%  
commit
```

```
[ma-b72-c7->configurationoverlay[OpenStackHypervisors]->roles[OpenStack::Compute]->imagebackends[ceph]]% ..
```

```
[ma-b72-c7->configurationoverlay[OpenStackHypervisors]->roles[OpenStack::Compute]->imagebackends]% ..
```

```
[ma-b72-c7->configurationoverlay[OpenStackHypervisors]->roles[OpenStack::Compute]->fsmounts]% remove  
/var/lib/nova/instances
```

```
[ma-b72-c7->configurationoverlay[OpenStackHypervisors]->roles[OpenStack::Compute]->fsmounts*]% commit
```

Unique solution ID: #1317

Author: mohamed adel

Last update: 2016-09-23 14:22