

# OpenStack: How do I configure CephFS?

This article explains how to configure CephFS on a Bright Cluster. This procedure has the following requirements:

- Bright Cluster Manager 8.0 or higher version.
- Ceph already deployed as described in the Bright Openstack Deployment Manual.
- There are enough placement groups to create the necessary Ceph pools.

The examples about how to mount the filesystem assume that the server that going to have the filesystem mounted on it is in the Ceph public network (this is true for the head node and all compute nodes in a typical installation).

## 1 - Deploy Ceph Metadata Server

Choose which server will be used to run the Ceph Metadata Server. Then run the following commands on the head node.

```
# cmsb
```

```
% device use <MDSSERVER>
```

```
% roles
```

```
% assign cephmbs
```

```
% commit
```

In the above commands, replace <MDSSERVER> with the name of the node which will run the Ceph Metadata Server.

## 2 - Create Ceph pools

Create two Ceph pools. One will be used for the filesystem itself and the other one will be used for the

# OpenStack: How do I configure CephFS?

metadata. The pools can be created by running the following commands on the head node:

```
# ceph
```

```
% ceph
```

```
% pools
```

```
% add cephfs_metadata
```

```
% set pgnum 32; set replicas 3
```

```
% commit
```

```
% add cephfs_data
```

```
% set pgnum 32; set replicas 3
```

```
% commit
```

The above commands can be modified to specify the desired number of placement groups and replicas.

## 3 - Create Ceph filesystem

Run the following command on the head node:

```
# ceph fs new cephfs cephfs_metadata cephfs_data
```

In the preceding command, cephfs is a name which will identify the created filesystem, while cephfs\_metadata and cephfs\_data are the pools created in step 2.

## 4 - Mounting the filesystem (optional)

# OpenStack: How do I configure CephFS?

The documentation at <http://docs.ceph.com> explains different ways to mount the filesystem (at <http://docs.ceph.com/docs/jewel/cephfs/> for Ceph Jewel, at <http://docs.ceph.com/docs/luminous/cephfs/> for Ceph Luminous, and so on).

The following examples assume that the filesystem is going to be mounted on the head node and that the client.admin Ceph identity will be used for authentication. Please note that additional steps will be needed for other cases.

CephFS can be mounted either by using a FUSE (user space) client or a kernel client. It is recommended to go over the CephFS documentation (and the <http://docs.ceph.com/docs/master/cephfs/best-practices> section in particular) to decide which one is better suited for the particular use case.

## 4.1 - Mounting the filesystem using FUSE

In order to mount the Ceph filesystem created in the previous steps, first install the Ceph FUSE client:

```
# yum install ceph-fuse
```

The head node already has the `/etc/ceph/ceph.conf` file with the Ceph configuration and the `/etc/ceph/ceph.client.admin.keyring` with the keyring for the client.admin

identity, so the filesystem can be mounted by running this command (assuming `/mnt/cephfs` exists):

```
# ceph-fuse /mnt/cephfs/
```

# OpenStack: How do I configure CephFS?

This configuration can be made persistent via an entry in `/etc/fstab`. This requires the version of the `util-linux` package which was included on RHEL 7.3. So if the following procedure doesn't work, the administrator should try updating the `util-linux` package.

For Ceph Jewel (the version supported by Bright 8.0) the following entry has to be added to `/etc/fstab`:

```
id=admin,conf=/etc/ceph/ceph.conf /mnt/cephfs fuse.ceph _netdev,defaults 0 0
```

For newer versions of Ceph, a different syntax can be used:

```
none /mnt/cephfs fuse.ceph ceph.id=admin,ceph.conf=/etc/ceph/ceph.conf,_netdev,defaults 0 0
```

For the preceding examples, the mount points can also be added to `/etc/fstab` by `CMDaemon`. Thus, the example for Ceph Jewel can be configured by running the following commands on the head node:

```
# cmsg
```

```
% device use master
```

```
% fsmounts
```

```
% add /mnt/cephfs
```

```
% set device id=admin,conf=/etc/ceph/ceph.conf
```

```
% set filesystem fuse.ceph
```

```
% set mountoptions _netdev,defaults
```

```
% commit
```

The example for newer versions of Ceph can be configured by running the following commands on the head node:

# OpenStack: How do I configure CephFS?

```
# cmsg
```

```
% device use master
```

```
% fsmounts
```

```
% add /mnt/cephfs
```

```
% set device none
```

```
% set filesystem fuse.ceph
```

```
% set mountoptions ceph.id=admin,ceph.conf=/etc/ceph/ceph.conf,_netdev,defaults
```

```
% commit
```

## 4.2 - Mounting the filesystem using FUSE (another method, using systemd unit)

Starting with Ceph Luminous (the version supported in Bright 8.1) and when using RHEL 7 or CentOS 7, the ceph-fuse package includes a systemd unit that can be used to mount the filesystem. For example, to mount the filesystem on /mnt/cephfs, the following command can be run (assuming /mnt/cephfs exists):

```
# systemctl start ceph-fuse@/mnt/cephfs
```

It is possible to make this configuration persistent by running the following command:

```
# systemctl enable ceph-fuse@-mnt-cephfs
```

In the above command, please note that even though the mount point is "/mnt/cephfs" the "/" in the path has to be substituted by a "-".

# OpenStack: How do I configure CephFS?

## 4.3 - Mounting the filesystem using the kernel driver

Run the following command on the head node to generate a file with the client.admin secret key:

```
# ceph auth get-key client.admin > /etc/ceph/ceph.admin.secret
```

Then run the following command (assuming /mnt/cephfs exists):

```
# mount -t ceph <monitor>:6789:/ /mnt/cephfs -o name=admin,secretfile=/etc/ceph/ceph.admin.secret
```

In the above command, substitute <monitor> for the IP address of one of the Ceph monitors.

To make this configuration persistent, the following line has to be added to /etc/fstab:

```
<monitor>:6789          /mnt/cephfs          ceph
name=admin,secretfile=/etc/ceph/ceph.admin.secret 0 0
```

If the client node is a node managed by CMDaemon, like the head node, then the mountpoint can be added by running the following commands on the head node:

```
# cmsg
```

```
% device use master
```

```
% fsmounts
```

```
% add /mnt/cephfs
```

# OpenStack: How do I configure CephFS?

```
% set device <monitor>:6789
```

```
% set filesystem ceph
```

```
% set mountoptions name=admin,secretfile=/etc/ceph/ceph.admin.secret
```

```
% commit
```

Unique solution ID: #1411

Author: Carlos Pintado

Last update: 2018-10-16 12:09