

CephFS

- Ceph ceph-deploy
- Ceph



- Astra Linux Special Edition .10015-01 (1.7), .10015-10
- Astra Linux Special Edition .10015-17
- Astra Linux Special Edition .10015-37 (7.7)
- Astra Linux Special Edition .10015-03 (7.6)
- Astra Linux Special Edition .10152-02 (4.7)
- Astra Linux Special Edition .10015-01 (1.6)
- Astra Linux Special Edition .10015-16 . 1
- Astra Linux Special Edition .10015-16 . 2
- Astra Linux Special Edition .10265-01 (8.1)
- Astra Linux Common Edition 2.12

, . Astra Linux Ceph.

Ceph --- , Ceph

Ceph :

- cephfs;
- rbd;
- s3- api.

Ceph ceph-deploy

Ceph 3:

- astra-ceph1;
- astra-ceph2;
- astra-ceph3.

astra-ceph-admin.

- (MON);
- (OSD).



Ceph OSD .

- /dev/sda Astra Linux;
- /dev/sdb OSD.

, astra-ceph1 ceph-mgr (Ceph Manager Daemon).

, Ceph astra-ceph1,astra-ceph2 astra-ceph3 :

- IP-, DNS IP- astra-ceph-admin;
- . astra-ceph-admin.



Ceph .

ceph-deploy astra-ceph-admin. , /etc/hosts :

```
10.0.0.100      astra-ceph-admin
10.0.0.171      astra-ceph1
10.0.0.172      astra-ceph2
10.0.0.173      astra-ceph3
```

ceph-deploy Ceph , ssh, sudo Python.

1. :

a. python ssh, :

```
sudo apt install python ssh
```

b. ssh, :

```
sudo systemctl enable --now ssh
```

 , SSH ;

c. , , ceph-adm:

```
sudo adduser ceph-adm
```

, . "y" ("");

d. (sudo) , :

```
echo "ceph-adm ALL = (root) NOPASSWD:ALL" | sudo tee /etc/sudoers.d/ceph-adm
sudo chmod 0440 /etc/sudoers.d/ceph-adm
```

e. ceph-adm:

```
sudo pdpl-user -i 63 ceph-adm
```



Ceph ceph-deploy astra-ceph-admin ceph-adm.
ceph-deploy (sudo root).

Ceph astra-ceph-admin :

1. ceph-deploy:

```
sudo apt install ceph-deploy python
```

2. ceph-adm:

```
su - ceph-adm
```

3. ceph-adm :

a. ssh- :

```
ssh-keygen
for N in $(seq 1 3); do ssh-copy-id ceph-adm@astra-ceph$N; done
```

```
<<yes>> (<<>>) ceph-adm;
```

b. ssh- :

```
ssh-copy-id ceph-adm@astra-ceph-admin
```

c. Ceph :

```
ceph-deploy --username ceph-adm install --mon --osd astra-ceph1
astra-ceph2 astra-ceph3
```

```
--mon --osd Ceph, . Ceph;
```

d. :

```
for N in $(seq 1 3); do ssh ceph-adm@astra-ceph$N sudo reboot; done
```

e. astra-cephl ceph-mgr():

```
ceph-deploy --username ceph-adm install --mgr astra-ceph1
```

f. Ceph, , :

```
ceph-deploy --username ceph-adm new astra-ceph1 astra-ceph2 astra-
ceph3
```

```
ceph.conf keyring- ;
```

g. :

```
ceph-deploy --username ceph-adm mon create-initial
```

h. mgr astra-ceph1 :

```
ceph-deploy --username ceph-adm mgr create astra-ceph1
```

i. OSD /dev/sdb :

```
ceph-deploy --username ceph-adm osd create --data /dev/sdb astra-
ceph1
ceph-deploy --username ceph-adm osd create --data /dev/sdb astra-
ceph2
ceph-deploy --username ceph-adm osd create --data /dev/sdb astra-
ceph3
```

j. Ceph astra-ceph-admin :

```
ceph-deploy --username ceph-adm install --cli astra-ceph-admin
```

k. keyring- admin astra-ceph-admin :

```
ceph-deploy admin astra-ceph-admin
```

I. RBD(), Metadata Server Daemon - astra-ceph1:

```
ceph-deploy --username ceph-adm install --mds astra-ceph1  
ceph-deploy --username ceph-adm mds create astra-ceph1
```

m. ceph-adm:

```
exit
```

4. Ceph :

```
sudo ceph -s
```

```
health HEALTH_OK:
```

```
cluster:  
  id: 03ff5b8a-a453-4da3-8296-2d473649bcc4  
  health: HEALTH_OK  
services:  
  mon: 3 daemons, quorum astra-ceph1,astra-ceph2,astra-ceph3 (age 3h)  
    mgr: astra-ceph1(active, since 3h)  
    osd: 3 osds: 3 up (since 3h), 3 in (since 25h)  
data:  
  pools: 0 pools, 0 pgs  
  objects: 0 objects, 0 bytes  
  usage: 3.0 GiB used, 45 GiB / 48 GiB avail  
  pgs:
```

Ceph

cephfs. cephfs , . , Ceph , (MDS).
cephfs astra-ceph-admin:

1. Metadata Server Daemon - astra-ceph1:

```
ceph-deploy --username ceph-adm install --mds astra-ceph1  
ceph-deploy --username ceph-adm mds create astra-ceph1
```

2. cephfs_data cephfs_metadata, - (PG) 64:

```
sudo ceph osd pool create cephfs_data 64  
sudo ceph osd pool create cephfs_metadata 64  
sudo ceph osd pool application enable cephfs_metadata cephfs
```

PG , 2 (64, 128, 256...). OSD . PG OSD, . PG, . PG:

```
Total PGs = (Number OSD * 100) / max_replication_count
```

Ceph: <https://docs.ceph.com/en/latest/rados/operations/placement-groups/#choosing-the-number-of-placement-groups>.

3. cephfs:

```
sudo ceph fs new cephfs cephfs_metadata cephfs_data
```

4. cephfs astra-ceph-admin:
a. , :

```
cat ceph.client.admin.keyring | grep key
```

key. :

```
key = AQBBVX1g12oJJBAh40D+1Kphz/0QA/Gbkz1sw==
```

b. ():

```
echo "AQBBVX1g12oJJBAh40D+1Kphz/0QA/Gbkz1sw==" > admin.secret
```

c. , cephfs:

```
sudo mkdir /mnt/cephfs
```

d. cephfs:

```
sudo mount -t ceph 10.0.0.171:6789:/ /mnt/cephfs -o name=admin,  
secretfile=admin.secret
```

10.0.0.171 --- . , , , , , , , :

```
sudo mount -t ceph 10.0.0.171,10.0.0.172,10.0.0.173:/ /mnt  
/mycephfs -o name=admin,secretfile=admin.secret
```

e. , :

```
df -h | grep cephfs
```

:

```
10.0.0.171:6789:/ 15G 0 15G 0% /mnt/cephfs
```

15 (3 16, 3 1). , 15 ceph. .

. cephfs :

1. astra-ceph-admin cephfs data1, :

```
sudo mkdir /mnt/cephfs/data1
```

2. client.data1 data1, :

```
sudo ceph auth get-or-create client.data1 mon 'allow r' mds 'allow r,  
allow rw path=/data1' osd 'allow rw pool=cephfs_data'
```

key. :

```
[client.data1]  
key = AQDh335g/MDeKBAAOxnXO/H4W7g2snP0pq+lCA==
```

:

```
sudo ceph auth get-key client.data1
```

3. , cephfs, data1, IP- :

```
sudo mount -t ceph 10.0.0.171,10.0.0.172,10.0.0.173:/ /mnt -o  
name=data1,secret='AQDh335g/MDeKBAAOxnXO/H4W7g2snPOpq+lCA=='
```

10.0.0.171,10.0.0.172,10.0.0.173 --- IP- , secret --- , .

4. , :

```
df -h | grep mnt
```

:

10.0.0.171,10.0.0.172,10.0.0.173:/	15G	0	15G	0% /mnt
------------------------------------	-----	---	-----	---------

, data1 cephfs. , , .