Deploy Ceph Cluster using Cephadm on CentOS 8

Ach.Chuanuk Chikam · Follow
Jun 14, 2020 · 10 min read

Ceph is an open source software (software-defined storage) storage platform, implements object storage on a single distributed computer cluster, and provides it in-s interfaces for object-, block and file-level storage. Ceph aims primarily for completely distributed operation without a single point of failure, scalable to the exabyte level, and freely available.

There are several different ways to install Ceph. Choose the method that best suits your needs. For recommendation on ceph documentation is used cephadm. Cephadm installs and manages a Ceph cluster using containers and systemd, with tight integration with the CLI and dashboard GUI. Cephadm only supports Octopus and newer releases. Cephadm is fully integrated with the new orchestration API and fully supports the new CLI and dashboard features to manage cluster deployment. Cephadm requires container support (podman or docker) and Python 3.

Prerequisites:

- OS CentOS 8
- 3 Nodes with 3 additional disk (sdb, sdc, sdd)

Lab Environment:

Server	OS	Specs	IP Address	Additional Disk
Ceph-01	CentOS 8	4GB RAM, 2vCPU	192.168.1.21	2xlb, sdb, sdd
Ceph-02	CentOS 8	4GB RAM, 2vCPU	192.168.1.22	2xlb, sdb, sdd
Ceph-03	CentOS 8	4GB RAM, 2vCPU	192.168.1.23	2xlb, sdb, sdd

Setup Ceph Cluster:

Let's take a look at the steps required to set up Ceph Cluster using kubeadm on CentOS 8. The steps look something like the following:

1. Login as root and add host on /etc/hosts

```
$ sudo -i
$ sudo tee -a /etc/hosts<<EOF
192.168.1.21 ceph-01 ceph-01.localdomain
192.168.1.222 ceph-02 ceph-02.localdomain
192.168.1.223 ceph-03 ceph-03.localdomain
EOF

Set hostnam for 3 node ceph cluster
# hostnamc ceph-01
# sudo ssh root@192.168.1.222 hostnamc ceph-02
# sub ssh-01 hostnamc ceph-03
```

2. Set passwordless

```
# ssh-keygen
# ssh-copy-id ceph-01
# ssh-copy-id ceph-02
# ssh-copy-id ceph-03
```

3. Update dan Upgrade package

```
# dnf update -y; dnf upgrade -y
```

4. Install python3, lvm2, and podman

```
# dnf install -y python3 lvm2 podman
```

5. Install Cephadm using curl-based installation method

```
Use curl to fetch the most recent version of the standalone script
[root@ceph-01 ~]# curl --silent --remote-name --location
https://github.com/ceph/ceph/raw/octopus/src/cephadm/cephadm

Make the cephadm script executable:
[root@ceph-01 ~]# chmod +x cephadm

To install the packages that provide the cephadm command, run the
following commands:
[root@ceph-01 ~]# ./cephadm add-repo --release octopus
[root@ceph-01 ~]# ./cephadm install

Install ceph-common and confirm that cephadm is now in your PATH
by running which:
[root@ceph-01 ~]# dnf install -y ceph-common
[root@ceph-01 ~]# which cephadm
```

6. Bootstrap new Ceph cluster

The first step in creating a new Ceph cluster is running the `cephadm bootstrap` command on the Ceph cluster's first host. The act of running the `cephadm bootstrap` command on the Ceph cluster's first host creates the Ceph cluster's first "monitor daemon", and that monitor daemon needs an IP address. You must pass the IP address of the Ceph cluster's first host to the `ceph bootstrap` command, so you'll need to know the IP address of that host.

```
[root@ceph-01 ~]# cephadm bootstrap --mon-ip 192.168.1.231
```

Check Ceph dashboard, access IP address of ceph-01 https://192.168.1.231:8443/ and use credentials from the `cephadm bootstrap` output then set a new password

Verify Ceph CLI

```
Confirm that the ceph command is accessible with:
[root@ceph-01 ~]# ./cephadm shell -- ceph -v
ceph version 15.2.0 (34ce481bca7a854ca3ed8a8firc1e44e1e29d011)
octopus (stable)

Check status of ceph cluster. OK for [HEALTH_WARN] because OSDs are
not added yet

[root@ceph-01 ~]# ceph -s
  cluster:
    id:     3f6a2f1b-22dc-11ec-b90d-0050568ee762
    health: HEALTH_WARN
            OSD count 0 < osd_pool_default_size 3

  services:
    mon: 1 daemons, quorum ceph-01 (age 14m)
    mgr: ceph-01.ugfjcnrtfx(active, since 12m)
    osd: 0 osds: 0 up, 0 in

  data:
    pools:   0 pools, 0 pools
    objects: 0 objects, 0 B
    usage:   0 B used, 0 B / 0 B avail
    pgs:

[root@ceph-01 ~]# ceph -w
  cluster:
    id:     3f6a2f1b-22dc-11ec-b90d-0050568ee762
    health: HEALTH_WARN
            OSD count 0 < osd_pool_default_size 3

  services:
    mon: 1 daemons, quorum ceph-01 (age 14m)
    mgr: ceph-01.ugfjcnrtfx(active, since 14m), standbys: ceph-
02.ldrzqq
    osd: 0 osds: 0 up, 0 in (since 14s)

  data:
    pools:   0 pools, 0 pools
    objects: 0 objects, 0 B
    usage:   0 B used, 0 B / 0 B avail
    pgs:

[root@ceph-01 ~]# ceph -v
ceph version 15.2.0 (34ce481bca7a854ca3ed8a8firc1e44e1e29d011)
octopus (stable)
[root@ceph-01 ~]# ceph osd tree
ID  CLASS  WEIGHT  TYPE NAME       STATUS  REWEIGHT  PRI-AFF
-1         0.02930 root default
-3         0.01953     host ceph-01
 1    hdd  0.00980         osd.1       up   1.00000   1.00000
 3    hdd  0.00980         osd.3       up   1.00000   1.00000
-7         0.00977     host ceph-02
 0    hdd  0.00980         osd.0       up   1.00000   1.00000
 2    hdd  0.00977         osd.2       up   1.00000   1.00000
-5         0.02930     host ceph-03
 4    hdd  0.00980         osd.4       up   1.00000   1.00000
 5    hdd  0.00980         osd.5       up   1.00000   1.00000
 6    hdd  0.00980         osd.6       up   1.00000   1.00000
```

7. Add hosts to the cluster

To add each new host to the cluster, perform two steps:

```
Install the cluster's public SSH key in the new host's root user's
authorized_keys file. For example, run the following commands:
[root@ceph-01 ~]# ./cephadm shell -- ceph cephadm get-pub-key
>~/ceph.pub
[root@ceph-01 ~]# ./cephadm shell -- ceph-volume lvm list
Tell Ceph that the new node is part of the cluster, make sure
the new host is installed and available on our node
[root@ceph-01 ~]# ceph orch host add ceph-02 192.168.1.222
Added host 'ceph-02' with addr '192.168.1.222'
[root@ceph-01 ~]#
[root@ceph-01 ~]# ceph orch host add ceph-03 192.168.1.223
Added host 'ceph-03' with addr '192.168.1.223'
[root@ceph-01 ~]#
[root@ceph-01 ~]# ceph orch host ls
HOST     ADDR         LABELS  STATUS
ceph-01  192.168.1.21  _admin
ceph-02  192.168.1.222
ceph-03  192.168.1.223
[root@ceph-01 ~]#
```

8. Deploy OSDs to the cluster

Run this command to display an inventory of storage devices on all cluster hosts:

```
[root@ceph-01 ~]# ceph orch device ls
Hostname  Path      Type  Size  Health   Ident  Fault
Available
ceph-01   /dev/sdb  ssd    10.7G  Unknown  N/A    N/A
Yes
ceph-01   /dev/sdc  ssd    10.7G  Unknown  N/A    N/A
Yes
ceph-01   /dev/sdd  ssd    10.7G  Unknown  N/A    N/A
Yes
ceph-02   /dev/sdb  ssd    10.7G  Unknown  N/A    N/A
Yes
ceph-02   /dev/sdc  ssd    10.7G  Unknown  N/A    N/A
Yes
ceph-02   /dev/sdd  ssd    10.7G  Unknown  N/A    N/A
Yes
ceph-03   /dev/sdb  ssd    10.7G  Unknown  N/A    N/A
Yes
ceph-03   /dev/sdc  ssd    10.7G  Unknown  N/A    N/A
Yes
ceph-03   /dev/sdd  ssd    10.7G  Unknown  N/A    N/A
Yes
```

Tell Ceph to consume any available and unused storage device execute `ceph orch apply osd --all-available-devices`

```
[root@ceph-01 ~]# ceph orch apply osd --all-available-devices
Scheduled osd.all-available-devices update...
[root@ceph-01 ~]# ceph -s
  cluster:
    id:     3f6a2f1b-22dc-11ec-b90d-0050568ee762
    health: HEALTH_OK

  services:
    mon: 1 daemons, quorum ceph-01(age5.ceph-01,ceph-03 (age 3m)
    mgr: ceph-01.ugfjcnrtfx(active, since 6m), standbys: ceph-
02.ldrzqq
    osd: 6 osds: 6 up, 6 in (since 24s)

  data:
    pools:   0 pools, 0 pools
    objects: 0 objects, 0 B
    usage:   6 B used, 6 B / 6 B avail
    pgs:

[root@ceph-01 ~]#
[root@ceph-01 ~]# ceph osd tree
ID  CLASS  WEIGHT  TYPE NAME       STATUS  REWEIGHT  PRI-AFF
-1         0.02930 root default
-3         0.01953     host ceph-01
 1    hdd  0.00980         osd.1       up   1.00000   1.00000
 3    hdd  0.00980         osd.3       up   1.00000   1.00000
-7         0.00977     host ceph-02
 0    hdd  0.00980         osd.0       up   1.00000   1.00000
 2    hdd  0.00977         osd.2       up   1.00000   1.00000
-5         0.02930     host ceph-03
 4    hdd  0.00980         osd.4       up   1.00000   1.00000
 5    hdd  0.00980         osd.5       up   1.00000   1.00000
```

9. Deploy ceph-mon (ceph monitor daemon)

Ceph-mon is the cluster monitor daemon for the Ceph distributed file system. One or more instances of **ceph-mon** form a Paxos part-time parliament cluster that provides extremely reliable and durable storage of cluster membership, configuration, and state. Add ceph-mon to all node using placement option

```
[root@ceph-01 ~]# ceph orch apply mon --placement="ceph-01,ceph-
02,ceph-03"
Scheduled mon update...
[root@ceph-01 ~]# ceph orch ps | grep mon
mon.ceph-01  ceph-01  running (14m)   8m
ago 92m    200M    2048M  16.2.0   81a7211be67a  f2a2c49ff9b
mon.ceph-02  ceph-02  running (2m)    2m
ago 2m     29M     2048M  16.2.0   81a7211be67a  3d77b6a2d9c
mon.ceph-03  ceph-03  running (2m)    2m
ago 2m     26M     2048M  16.2.0   81a7211be67a  3cc1d8
[root@ceph-01 ~]#
```

10. Deploy ceph-mgr (ceph manager daemon)

The Ceph Manager daemon (ceph-mgr) runs alongside monitor daemons, to provide additional monitoring and interfaces to external monitoring and management systems.

```
[root@ceph-01 ~]# ceph orch apply mgr --placement="ceph-01,ceph-
02,ceph-03"
Scheduled mgr update...
[root@ceph-01 ~]# ceph orch ps | grep mgr
mgr.ceph-01.ugfjcn  ceph-01  *:9283  running (14m)   8m
ago 94m    465M        -  16.2.0   81a7211be67a  c3a6ce1f9f8b
mgr.ceph-02.ldrzqq  ceph-02  *:8443  running (4m)    4m
ago 4m     429M    512M  16.2.0   81a7211be67a  1d77b6a2d9c
mgr.ceph-03.ccbzzr  ceph-03  *:8443  running (4m)    4m
ago 4m     415M    512M  16.2.0   81a7211be67a  c7e0f0e4a2de
[root@ceph-01 ~]#
```

11. Set label on all nodes

The orchestrator supports assigning labels to hosts. Labels are free form and have no particular meaning by itself and each host can have multiple labels. They can be used to specify placement of daemons.

```
Set label osd-node on all node
[root@ceph-01 ~]# ceph orch host label add ceph-01 osd-node
[root@ceph-01 ~]# ceph orch host label add ceph-02 osd-node
[root@ceph-01 ~]# ceph orch host label add ceph-03 osd-node

Set label mon on all node
[root@ceph-01 ~]# ceph orch host label add ceph-01 mon
[root@ceph-01 ~]# ceph orch host label add ceph-02 mon
[root@ceph-01 ~]# ceph orch host label add ceph-03 mon

Set label mgr on all node
[root@ceph-01 ~]# ceph orch host label add ceph-01 mgr
[root@ceph-01 ~]# ceph orch host label add ceph-02 mgr
[root@ceph-01 ~]# ceph orch host label add ceph-03 mgr
```

Now verify ceph cluster and see everything is great

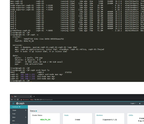

See other content

References :

- https://docs.ceph.com/en/latest/cephadm/install/
- https://www.server-world.info/

#CEPH #Storage #CentOS8 #SDS

Written by Ach.Chuanuk Chikam
DevOpsEngineer | SE/Sysadmin
Follow
Cloud Enthusiast | Linux Administrator, OpenStack, OpenShift, Ceph Certified | AWS SAA | Google Cloud ACE | Linkedin
https://www.linkedin.com/in/achchuanuk/0b80

No responses yet