

原生态 CEPH 安装步骤

存储软件 孔樱【技术大咖】 2024-03-31 发表

组网及说明

为了学习 X10000 以及 OneStor, 先要搞明白 CEPH 架构, 以及里面的各种概念。

再次基础上自行搭建一套原生态的 CEPH 软件定义存储,

在 VMWARE 环境下配置了三个虚机作为 (Mon, MDS, MGR) 节点;

后端用 3PAR 磁盘作为 OSD 磁盘。

配置步骤

```
[root@ceph-master cephcluster]# ceph-deploy new --cluster-network 10.1.1.0/24 --public-network  
10.12.180.0/22 ceph-master ceph-node1 ceph-node2  
[ceph_deploy.conf][DEBUG ] found configuration file at: /root/.cephdeploy.conf  
[ceph_deploy.cli][INFO ] Invoked (2.0.1): /usr/bin/ceph-deploy new --cluster-network 10.1.1.0/24 --pu  
blic-network 10.12.180.0/22 ceph-master ceph-node1 ceph-node2  
[ceph_deploy.cli][INFO ] ceph-deploy options:  
[ceph_deploy.cli][INFO ] username : None  
[ceph_deploy.cli][INFO ] func : <function new at 0x7f4998c40230>  
[ceph_deploy.cli][INFO ] verbose : False  
[ceph_deploy.cli][INFO ] overwrite_conf : False  
[ceph_deploy.cli][INFO ] quiet : False  
[ceph_deploy.cli][INFO ] cd_conf : <ceph_deploy.conf.cephdeploy.Conf instance at 0x  
7f49983a8cf8>  
[ceph_deploy.cli][INFO ] cluster : ceph  
[ceph_deploy.cli][INFO ] ssh_copykey : True  
[ceph_deploy.cli][INFO ] mon : ['ceph-master', 'ceph-node1', 'ceph-node2']  
[ceph_deploy.cli][INFO ] public_network : 10.12.180.0/22  
[ceph_deploy.cli][INFO ] ceph_conf : None  
[ceph_deploy.cli][INFO ] cluster_network : 10.1.1.0/24  
[ceph_deploy.cli][INFO ] default_release : False  
[ceph_deploy.cli][INFO ] fsid : None  
[ceph_deploy.new][DEBUG ] Creating new cluster named ceph  
[ceph_deploy.new][INFO ] making sure passwordless SSH succeeds  
[ceph-master][DEBUG ] connected to host: ceph-master  
[ceph-master][DEBUG ] detect platform information from remote host  
[ceph-master][DEBUG ] detect machine type  
[ceph-master][DEBUG ] find the location of an executable  
[ceph-master][INFO ] Running command: /usr/sbin/ip link show  
[ceph-master][INFO ] Running command: /usr/sbin/ip addr show  
[ceph-master][DEBUG ] IP addresses found: [u'10.12.180.122', u'192.168.122.1', u'10.1.1.122']  
[ceph_deploy.new][DEBUG ] Resolving host ceph-master  
[ceph_deploy.new][DEBUG ] Monitor ceph-master at 10.12.180.122  
[ceph_deploy.new][INFO ] making sure passwordless SSH succeeds  
[ceph-node1][DEBUG ] connected to host: ceph-master.h3cu.com  
[ceph-node1][INFO ] Running command: ssh -CT -o BatchMode=yes ceph-node1  
[ceph-node1][DEBUG ] connected to host: ceph-node1  
[ceph-node1][DEBUG ] detect platform information from remote host  
[ceph-node1][DEBUG ] detect machine type  
[ceph-node1][DEBUG ] find the location of an executable  
[ceph-node1][INFO ] Running command: /usr/sbin/ip link show  
[ceph-node1][INFO ] Running command: /usr/sbin/ip addr show  
[ceph-node1][DEBUG ] IP addresses found: [u'10.12.180.123', u'192.168.122.1', u'10.1.1.123']  
[ceph_deploy.new][DEBUG ] Resolving host ceph-node1  
[ceph_deploy.new][DEBUG ] Monitor ceph-node1 at 10.12.180.123  
[ceph_deploy.new][INFO ] making sure passwordless SSH succeeds  
[ceph-node2][DEBUG ] connected to host: ceph-master.h3cu.com  
[ceph-node2][INFO ] Running command: ssh -CT -o BatchMode=yes ceph-node2  
[ceph-node2][DEBUG ] connected to host: ceph-node2  
[ceph-node2][DEBUG ] detect platform information from remote host  
[ceph-node2][DEBUG ] detect machine type  
[ceph-node2][DEBUG ] find the location of an executable  
[ceph-node2][INFO ] Running command: /usr/sbin/ip link show
```

```
[ceph-node2][INFO ] Running command: /usr/sbin/ip addr show
[ceph-node2][DEBUG ] IP addresses found: [u'10.1.1.124', u'192.168.122.1', u'10.12.180.124']
[ceph_deploy.new][DEBUG ] Resolving host ceph-node2
[ceph_deploy.new][DEBUG ] Monitor ceph-node2 at 10.12.180.124
[ceph_deploy.new][DEBUG ] Monitor initial members are ['ceph-master', 'ceph-node1', 'ceph-node2']
[ceph_deploy.new][DEBUG ] Monitor addrs are [u'10.12.180.122', u'10.12.180.123', u'10.12.180.124']
[ceph_deploy.new][DEBUG ] Creating a random mon key...
[ceph_deploy.new][DEBUG ] Writing monitor keyring to ceph.mon.keyring...
[ceph_deploy.new][DEBUG ] Writing initial config to ceph.conf...
[root@ceph-master cephcluster]#
```

Each node install ceph package

```
yum install ceph ceph-mon ceph-mgr ceph-radosgw ceph-mds -y
```

Install monitor daemon

```
[root@ceph-master cephcluster]# ceph-deploy mon create-initial
[ceph_deploy.conf][DEBUG ] found configuration file at: /root/.cephdeploy.conf
[ceph_deploy.cli][INFO ] Invoked (2.0.1): /usr/bin/ceph-deploy mon create-initial
[ceph_deploy.cli][INFO ] ceph-deploy options:
[ceph_deploy.cli][INFO ] username          : None
[ceph_deploy.cli][INFO ] verbose           : False
[ceph_deploy.cli][INFO ] overwrite_conf    : False
[ceph_deploy.cli][INFO ] subcommand        : create-initial
[ceph_deploy.cli][INFO ] quiet             : False
[ceph_deploy.cli][INFO ] cd_conf           : <ceph_deploy.conf.cephdeploy.Conf instance at 0x
7fdebc0cba28>
[ceph_deploy.cli][INFO ] cluster          : ceph
[ceph_deploy.cli][INFO ] func              : <function mon at 0x7fdebc1197d0>
[ceph_deploy.cli][INFO ] ceph_conf         : None
[ceph_deploy.cli][INFO ] default_release   : False
[ceph_deploy.cli][INFO ] keyrings          : None
[ceph_deploy.mon][DEBUG ] Deploying mon, cluster ceph hosts ceph-master ceph-node1 ceph-node
2
[ceph_deploy.mon][DEBUG ] detecting platform for host ceph-master ...
[ceph-master][DEBUG ] connected to host: ceph-master
[ceph-master][DEBUG ] detect platform information from remote host
[ceph-master][DEBUG ] detect machine type
[ceph-master][DEBUG ] find the location of an executable
[ceph_deploy.mon][INFO ] distro info: CentOS Linux 7.9.2009 Core
[ceph-master][DEBUG ] determining if provided host has same hostname in remote
[ceph-master][DEBUG ] get remote short hostname
[ceph-master][DEBUG ] deploying mon to ceph-master
[ceph-master][DEBUG ] get remote short hostname
[ceph-master][DEBUG ] remote hostname: ceph-master
[ceph-master][DEBUG ] write cluster configuration to /etc/ceph/{cluster}.conf
[ceph-master][DEBUG ] create the mon path if it does not exist
[ceph-master][DEBUG ] checking for done path: /var/lib/ceph/mon/ceph-ceph-master/done
[ceph-master][DEBUG ] done path does not exist: /var/lib/ceph/mon/ceph-ceph-master/done
[ceph-master][INFO ] creating keyring file: /var/lib/ceph/tmp/ceph-ceph-master.mon.keyring
[ceph-master][DEBUG ] create the monitor keyring file
[ceph-master][INFO ] Running command: ceph-mon --cluster ceph --mkfs -i ceph-master --keyring /v
ar/lib/ceph/tmp/ceph-ceph-master.mon.keyring --setuser 167 --setgroup 167
[ceph-master][INFO ] unlinking keyring file /var/lib/ceph/tmp/ceph-ceph-master.mon.keyring
[ceph-master][DEBUG ] create a done file to avoid re-doing the mon deployment
[ceph-master][DEBUG ] create the init path if it does not exist
[ceph-master][INFO ] Running command: systemctl enable ceph.target
[ceph-master][INFO ] Running command: systemctl enable ceph-mon@ceph-master
[ceph-master][WARNIN] Created symlink from /etc/systemd/system/ceph-mon.target.wants/ceph-mo
n@ceph-master.service to /usr/lib/systemd/system/ceph-mon@.service.
[ceph-master][INFO ] Running command: systemctl start ceph-mon@ceph-master
[ceph-master][INFO ] Running command: ceph --cluster=ceph --admin-daemon /var/run/ceph/ceph-
mon.ceph-master.asok mon_status
[ceph-master][DEBUG ] ****
[ceph-master][DEBUG ] status for monitor: mon.ceph-master
```

```
[ceph-master][DEBUG ] {
[ceph-master][DEBUG ]   "election_epoch": 0,
[ceph-master][DEBUG ]   "extra_probe_peers": [
[ceph-master][DEBUG ]     {
[ceph-master][DEBUG ]       "addrvec": [
[ceph-master][DEBUG ]         {
[ceph-master][DEBUG ]           "addr": "10.12.180.123:3300",
[ceph-master][DEBUG ]           "nonce": 0,
[ceph-master][DEBUG ]           "type": "v2"
[ceph-master][DEBUG ]         },
[ceph-master][DEBUG ]         {
[ceph-master][DEBUG ]           "addr": "10.12.180.123:6789",
[ceph-master][DEBUG ]           "nonce": 0,
[ceph-master][DEBUG ]           "type": "v1"
[ceph-master][DEBUG ]         }
[ceph-master][DEBUG ]       ],
[ceph-master][DEBUG ]     },
[ceph-master][DEBUG ]   {
[ceph-master][DEBUG ]     "addrvec": [
[ceph-master][DEBUG ]       {
[ceph-master][DEBUG ]         "addr": "10.12.180.124:3300",
[ceph-master][DEBUG ]         "nonce": 0,
[ceph-master][DEBUG ]         "type": "v2"
[ceph-master][DEBUG ]       },
[ceph-master][DEBUG ]       {
[ceph-master][DEBUG ]         "addr": "10.12.180.124:6789",
[ceph-master][DEBUG ]         "nonce": 0,
[ceph-master][DEBUG ]         "type": "v1"
[ceph-master][DEBUG ]       }
[ceph-master][DEBUG ]     ],
[ceph-master][DEBUG ]   ],
[ceph-master][DEBUG ] ],
[ceph-master][DEBUG ]   "feature_map": {
[ceph-master][DEBUG ]     "mon": [
[ceph-master][DEBUG ]       {
[ceph-master][DEBUG ]         "features": "0x3ffddff8ffecffff",
[ceph-master][DEBUG ]         "num": 1,
[ceph-master][DEBUG ]         "release": "luminous"
[ceph-master][DEBUG ]       }
[ceph-master][DEBUG ]     ],
[ceph-master][DEBUG ]   },
[ceph-master][DEBUG ]   "features": {
[ceph-master][DEBUG ]     "quorum_con": "0",
[ceph-master][DEBUG ]     "quorum_mon": [],
[ceph-master][DEBUG ]     "required_con": "0",
[ceph-master][DEBUG ]     "required_mon": []
[ceph-master][DEBUG ]   },
[ceph-master][DEBUG ]   "monmap": {
[ceph-master][DEBUG ]     "created": "2024-02-26 10:25:25.099139",
[ceph-master][DEBUG ]     "epoch": 0,
[ceph-master][DEBUG ]     "features": {
[ceph-master][DEBUG ]       "optional": [],
[ceph-master][DEBUG ]       "persistent": []
[ceph-master][DEBUG ]     },
[ceph-master][DEBUG ]     "fsid": "c75bd276-c8d6-4b2c-b2a9-9062e2bf66a2",
[ceph-master][DEBUG ]     "min_mon_release": 0,
[ceph-master][DEBUG ]     "min_mon_release_name": "unknown",
[ceph-master][DEBUG ]     "modified": "2024-02-26 10:25:25.099139",
[ceph-master][DEBUG ]     "mons": [
[ceph-master][DEBUG ]       {
[ceph-master][DEBUG ]         "addr": "10.12.180.122:6789/0",
[ceph-master][DEBUG ]         "name": "ceph-master",
[ceph-master][DEBUG ]         "public_addr": "10.12.180.122:6789/0",
[ceph-master][DEBUG ]         "public_addrs": {

```

```
[ceph-master][DEBUG ]    "addrvec": [
[ceph-master][DEBUG ]        {
[ceph-master][DEBUG ]            "addr": "10.12.180.122:3300",
[ceph-master][DEBUG ]            "nonce": 0,
[ceph-master][DEBUG ]            "type": "v2"
[ceph-master][DEBUG ]        },
[ceph-master][DEBUG ]        {
[ceph-master][DEBUG ]            "addr": "10.12.180.122:6789",
[ceph-master][DEBUG ]            "nonce": 0,
[ceph-master][DEBUG ]            "type": "v1"
[ceph-master][DEBUG ]        }
[ceph-master][DEBUG ]    ],
[ceph-master][DEBUG ]},
[ceph-master][DEBUG ]    "rank": 0
[ceph-master][DEBUG ]},
[ceph-master][DEBUG ]    {
[ceph-master][DEBUG ]        "addr": "0.0.0.0:0/1",
[ceph-master][DEBUG ]        "name": "ceph-node1",
[ceph-master][DEBUG ]        "public_addr": "0.0.0.0:0/1",
[ceph-master][DEBUG ]        "public_addrs": {
[ceph-master][DEBUG ]            "addrvec": [
[ceph-master][DEBUG ]                {
[ceph-master][DEBUG ]                    "addr": "0.0.0.0:0",
[ceph-master][DEBUG ]                    "nonce": 1,
[ceph-master][DEBUG ]                    "type": "v1"
[ceph-master][DEBUG ]                }
[ceph-master][DEBUG ]            ],
[ceph-master][DEBUG ]        },
[ceph-master][DEBUG ]        "rank": 1
[ceph-master][DEBUG ]},
[ceph-master][DEBUG ]    {
[ceph-master][DEBUG ]        "addr": "0.0.0.0:0/2",
[ceph-master][DEBUG ]        "name": "ceph-node2",
[ceph-master][DEBUG ]        "public_addr": "0.0.0.0:0/2",
[ceph-master][DEBUG ]        "public_addrs": {
[ceph-master][DEBUG ]            "addrvec": [
[ceph-master][DEBUG ]                {
[ceph-master][DEBUG ]                    "addr": "0.0.0.0:0",
[ceph-master][DEBUG ]                    "nonce": 2,
[ceph-master][DEBUG ]                    "type": "v1"
[ceph-master][DEBUG ]                }
[ceph-master][DEBUG ]            ],
[ceph-master][DEBUG ]        },
[ceph-master][DEBUG ]        "rank": 2
[ceph-master][DEBUG ]},
[ceph-master][DEBUG ]    ],
[ceph-master][DEBUG ]},
[ceph-master][DEBUG ]    "name": "ceph-master",
[ceph-master][DEBUG ]    "outside_quorum": [
[ceph-master][DEBUG ]        "ceph-master"
[ceph-master][DEBUG ]    ],
[ceph-master][DEBUG ]    "quorum": [],
[ceph-master][DEBUG ]    "rank": 0,
[ceph-master][DEBUG ]    "state": "probing",
[ceph-master][DEBUG ]    "sync_provider": []
[ceph-master][DEBUG ]}
[ceph-master][DEBUG ] ****
[ceph-master][INFO ] monitor: mon.ceph-master is running
[ceph-master][INFO ] Running command: ceph --cluster=ceph --admin-daemon /var/run/ceph/ceph-mon.ceph-master.asok mon_status
[ceph_deploy.mon][DEBUG ] detecting platform for host ceph-node1 ...
[ceph-node1][DEBUG ] connected to host: ceph-node1
[ceph-node1][DEBUG ] detect platform information from remote host
[ceph-node1][DEBUG ] detect machine type
```

```
[ceph-node1][DEBUG ] find the location of an executable
[ceph_deploy.mon][INFO ] distro info: CentOS Linux 7.9.2009 Core
[ceph-node1][DEBUG ] determining if provided host has same hostname in remote
[ceph-node1][DEBUG ] get remote short hostname
[ceph-node1][DEBUG ] deploying mon to ceph-node1
[ceph-node1][DEBUG ] get remote short hostname
[ceph-node1][DEBUG ] remote hostname: ceph-node1
[ceph-node1][DEBUG ] write cluster configuration to /etc/ceph/{cluster}.conf
[ceph-node1][DEBUG ] create the mon path if it does not exist
[ceph-node1][DEBUG ] checking for done path: /var/lib/ceph/mon/ceph-ceph-node1/done
[ceph-node1][DEBUG ] done path does not exist: /var/lib/ceph/mon/ceph-ceph-node1/done
[ceph-node1][INFO ] creating keyring file: /var/lib/ceph/tmp/ceph-ceph-node1.mon.keyring
[ceph-node1][DEBUG ] create the monitor keyring file
[ceph-node1][INFO ] Running command: ceph-mon --cluster ceph --mkfs -i ceph-node1 --keyring /var/lib/ceph/tmp/ceph-ceph-node1.mon.keyring --setuser 167 --setgroup 167
[ceph-node1][INFO ] unlinking keyring file /var/lib/ceph/tmp/ceph-ceph-node1.mon.keyring
[ceph-node1][DEBUG ] create a done file to avoid re-doing the mon deployment
[ceph-node1][DEBUG ] create the init path if it does not exist
[ceph-node1][INFO ] Running command: systemctl enable ceph.target
[ceph-node1][INFO ] Running command: systemctl enable ceph-mon@ceph-node1
[ceph-node1][WARNIN] Created symlink from /etc/systemd/system/ceph-mon.target.wants/ceph-mon@ceph-node1.service to /usr/lib/systemd/system/ceph-mon@.service.
[ceph-node1][INFO ] Running command: systemctl start ceph-mon@ceph-node1
[ceph-node1][INFO ] Running command: ceph --cluster=ceph --admin-daemon /var/run/ceph/ceph-mon.ceph-node1.asok mon_status
[ceph-node1][DEBUG ] *****
[ceph-node1][DEBUG ] status for monitor: mon.ceph-node1
[ceph-node1][DEBUG ] {
[ceph-node1][DEBUG ]   "election_epoch": 0,
[ceph-node1][DEBUG ]   "extra_probe_peers": [
[ceph-node1][DEBUG ]     {
[ceph-node1][DEBUG ]       "addrvec": [
[ceph-node1][DEBUG ]         {
[ceph-node1][DEBUG ]           "addr": "10.12.180.122:3300",
[ceph-node1][DEBUG ]           "nonce": 0,
[ceph-node1][DEBUG ]           "type": "v2"
[ceph-node1][DEBUG ]         },
[ceph-node1][DEBUG ]         {
[ceph-node1][DEBUG ]           "addr": "10.12.180.122:6789",
[ceph-node1][DEBUG ]           "nonce": 0,
[ceph-node1][DEBUG ]           "type": "v1"
[ceph-node1][DEBUG ]         }
[ceph-node1][DEBUG ]       ],
[ceph-node1][DEBUG ]     },
[ceph-node1][DEBUG ]   {
[ceph-node1][DEBUG ]     "addrvec": [
[ceph-node1][DEBUG ]       {
[ceph-node1][DEBUG ]         "addr": "10.12.180.124:3300",
[ceph-node1][DEBUG ]         "nonce": 0,
[ceph-node1][DEBUG ]         "type": "v2"
[ceph-node1][DEBUG ]       },
[ceph-node1][DEBUG ]       {
[ceph-node1][DEBUG ]         "addr": "10.12.180.124:6789",
[ceph-node1][DEBUG ]         "nonce": 0,
[ceph-node1][DEBUG ]         "type": "v1"
[ceph-node1][DEBUG ]       }
[ceph-node1][DEBUG ]     ],
[ceph-node1][DEBUG ]   }
[ceph-node1][DEBUG ] ],
[ceph-node1][DEBUG ]   "feature_map": {
[ceph-node1][DEBUG ]     "mon": [
[ceph-node1][DEBUG ]       {
[ceph-node1][DEBUG ]         "features": "0x3fffff8ffecffff",
[ceph-node1][DEBUG ]         "num": 1,

```

```
[ceph-node1][DEBUG ]      "release": "luminous"
[ceph-node1][DEBUG ]    }
[ceph-node1][DEBUG ]  ]
[ceph-node1][DEBUG ] },
[ceph-node1][DEBUG ]  "features": {
[ceph-node1][DEBUG ]    "quorum_con": "0",
[ceph-node1][DEBUG ]    "quorum_mon": [],
[ceph-node1][DEBUG ]    "required_con": "0",
[ceph-node1][DEBUG ]    "required_mon": []
[ceph-node1][DEBUG ] },
[ceph-node1][DEBUG ]  "monmap": {
[ceph-node1][DEBUG ]    "created": "2024-02-26 10:25:28.274480",
[ceph-node1][DEBUG ]    "epoch": 0,
[ceph-node1][DEBUG ]    "features": {
[ceph-node1][DEBUG ]      "optional": [],
[ceph-node1][DEBUG ]      "persistent": []
[ceph-node1][DEBUG ] },
[ceph-node1][DEBUG ]    "fsid": "c75bd276-c8d6-4b2c-b2a9-9062e2bf66a2",
[ceph-node1][DEBUG ]    "min_mon_release": 0,
[ceph-node1][DEBUG ]    "min_mon_release_name": "unknown",
[ceph-node1][DEBUG ]    "modified": "2024-02-26 10:25:28.274480",
[ceph-node1][DEBUG ]    "mons": [
[ceph-node1][DEBUG ]      {
[ceph-node1][DEBUG ]        "addr": "10.12.180.123:6789/0",
[ceph-node1][DEBUG ]        "name": "ceph-node1",
[ceph-node1][DEBUG ]        "public_addr": "10.12.180.123:6789/0",
[ceph-node1][DEBUG ]        "public_addrs": {
[ceph-node1][DEBUG ]          "addrvec": [
[ceph-node1][DEBUG ]            {
[ceph-node1][DEBUG ]              "addr": "10.12.180.123:3300",
[ceph-node1][DEBUG ]              "nonce": 0,
[ceph-node1][DEBUG ]              "type": "v2"
[ceph-node1][DEBUG ]            },
[ceph-node1][DEBUG ]            {
[ceph-node1][DEBUG ]              "addr": "10.12.180.123:6789",
[ceph-node1][DEBUG ]              "nonce": 0,
[ceph-node1][DEBUG ]              "type": "v1"
[ceph-node1][DEBUG ]            }
[ceph-node1][DEBUG ]          ]
[ceph-node1][DEBUG ]        },
[ceph-node1][DEBUG ]        "rank": 0
[ceph-node1][DEBUG ]      },
[ceph-node1][DEBUG ]      {
[ceph-node1][DEBUG ]        "addr": "0.0.0.0:0/1",
[ceph-node1][DEBUG ]        "name": "ceph-master",
[ceph-node1][DEBUG ]        "public_addr": "0.0.0.0:0/1",
[ceph-node1][DEBUG ]        "public_addrs": {
[ceph-node1][DEBUG ]          "addrvec": [
[ceph-node1][DEBUG ]            {
[ceph-node1][DEBUG ]              "addr": "0.0.0.0:0",
[ceph-node1][DEBUG ]              "nonce": 1,
[ceph-node1][DEBUG ]              "type": "v1"
[ceph-node1][DEBUG ]            }
[ceph-node1][DEBUG ]          ]
[ceph-node1][DEBUG ]        },
[ceph-node1][DEBUG ]        "rank": 1
[ceph-node1][DEBUG ]      },
[ceph-node1][DEBUG ]      {
[ceph-node1][DEBUG ]        "addr": "0.0.0.0:0/2",
[ceph-node1][DEBUG ]        "name": "ceph-node2",
[ceph-node1][DEBUG ]        "public_addr": "0.0.0.0:0/2",
[ceph-node1][DEBUG ]        "public_addrs": {
[ceph-node1][DEBUG ]          "addrvec": [
[ceph-node1][DEBUG ]            {

```

```
[ceph-node1][DEBUG ]      "addr": "0.0.0.0:0",
[ceph-node1][DEBUG ]      "nonce": 2,
[ceph-node1][DEBUG ]      "type": "v1"
[ceph-node1][DEBUG ]    }
[ceph-node1][DEBUG ]  ]
[ceph-node1][DEBUG ] },
[ceph-node1][DEBUG ] "rank": 2
[ceph-node1][DEBUG ] }
[ceph-node1][DEBUG ] ]
[ceph-node1][DEBUG ] },
[ceph-node1][DEBUG ] "name": "ceph-node1",
[ceph-node1][DEBUG ] "outside_quorum": [
[ceph-node1][DEBUG ] "ceph-node1"
[ceph-node1][DEBUG ] ],
[ceph-node1][DEBUG ] "quorum": [],
[ceph-node1][DEBUG ] "rank": 0,
[ceph-node1][DEBUG ] "state": "probing",
[ceph-node1][DEBUG ] "sync_provider": []
[ceph-node1][DEBUG ] }

[ceph-node1][DEBUG ] ****
[ceph-node1][INFO ] monitor: mon.ceph-node1 is running
[ceph-node1][INFO ] Running command: ceph --cluster=ceph --admin-daemon /var/run/ceph/ceph-mon.ceph-node1.asok mon_status
[ceph_deploy.mon][DEBUG ] detecting platform for host ceph-node2 ...
[ceph-node2][DEBUG ] connected to host: ceph-node2
[ceph-node2][DEBUG ] detect platform information from remote host
[ceph-node2][DEBUG ] detect machine type
[ceph-node2][DEBUG ] find the location of an executable
[ceph_deploy.mon][INFO ] distro info: CentOS Linux 7.9.2009 Core
[ceph-node2][DEBUG ] determining if provided host has same hostname in remote
[ceph-node2][DEBUG ] get remote short hostname
[ceph-node2][DEBUG ] deploying mon to ceph-node2
[ceph-node2][DEBUG ] get remote short hostname
[ceph-node2][DEBUG ] remote hostname: ceph-node2
[ceph-node2][DEBUG ] write cluster configuration to /etc/ceph/{cluster}.conf
[ceph-node2][DEBUG ] create the mon path if it does not exist
[ceph-node2][DEBUG ] checking for done path: /var/lib/ceph/mon/ceph-ceph-node2/done
[ceph-node2][DEBUG ] done path does not exist: /var/lib/ceph/mon/ceph-ceph-node2/done
[ceph-node2][INFO ] creating keyring file: /var/lib/ceph/tmp/ceph-ceph-node2.mon.keyring
[ceph-node2][DEBUG ] create the monitor keyring file
[ceph-node2][INFO ] Running command: ceph-mon --cluster ceph --mkfs -i ceph-node2 --keyring /var/lib/ceph/tmp/ceph-ceph-node2.mon.keyring --setuser 167 --setgroup 167
[ceph-node2][INFO ] unlinking keyring file /var/lib/ceph/tmp/ceph-ceph-node2.mon.keyring
[ceph-node2][DEBUG ] create a done file to avoid re-doing the mon deployment
[ceph-node2][DEBUG ] create the init path if it does not exist
[ceph-node2][INFO ] Running command: systemctl enable ceph.target
[ceph-node2][INFO ] Running command: systemctl enable ceph-mon@ceph-node2
[ceph-node2][WARNIN] Created symlink from /etc/systemd/system/ceph-mon.target.wants/ceph-mon@ceph-node2.service to /usr/lib/systemd/system/ceph-mon@.service.
[ceph-node2][INFO ] Running command: systemctl start ceph-mon@ceph-node2
[ceph-node2][INFO ] Running command: ceph --cluster=ceph --admin-daemon /var/run/ceph/ceph-mon.ceph-node2.asok mon_status
[ceph-node2][DEBUG ] ****
[ceph-node2][DEBUG ] status for monitor: mon.ceph-node2
[ceph-node2][DEBUG ] {
[ceph-node2][DEBUG ] "election_epoch": 1,
[ceph-node2][DEBUG ] "extra_probe_peers": [
[ceph-node2][DEBUG ] {
[ceph-node2][DEBUG ] "addrvec": [
[ceph-node2][DEBUG ] {
[ceph-node2][DEBUG ] "addr": "10.12.180.122:3300",
[ceph-node2][DEBUG ] "nonce": 0,
[ceph-node2][DEBUG ] "type": "v2"
[ceph-node2][DEBUG ] },
```

```
[ceph-node2][DEBUG ]    {
[ceph-node2][DEBUG ]      "addr": "10.12.180.122:6789",
[ceph-node2][DEBUG ]      "nonce": 0,
[ceph-node2][DEBUG ]      "type": "v1"
[ceph-node2][DEBUG ]    }
[ceph-node2][DEBUG ]  ],
[ceph-node2][DEBUG ]  },
[ceph-node2][DEBUG ]  {
[ceph-node2][DEBUG ]    "addrvec": [
[ceph-node2][DEBUG ]      {
[ceph-node2][DEBUG ]        "addr": "10.12.180.123:3300",
[ceph-node2][DEBUG ]        "nonce": 0,
[ceph-node2][DEBUG ]        "type": "v2"
[ceph-node2][DEBUG ]      },
[ceph-node2][DEBUG ]      {
[ceph-node2][DEBUG ]        "addr": "10.12.180.123:6789",
[ceph-node2][DEBUG ]        "nonce": 0,
[ceph-node2][DEBUG ]        "type": "v1"
[ceph-node2][DEBUG ]      }
[ceph-node2][DEBUG ]    ],
[ceph-node2][DEBUG ]  ],
[ceph-node2][DEBUG ]  }
[ceph-node2][DEBUG ] ],
[ceph-node2][DEBUG ] "feature_map": {
[ceph-node2][DEBUG ]   "mon": [
[ceph-node2][DEBUG ]     {
[ceph-node2][DEBUG ]       "features": "0x3fffff8ffecffff",
[ceph-node2][DEBUG ]       "num": 1,
[ceph-node2][DEBUG ]       "release": "luminous"
[ceph-node2][DEBUG ]     }
[ceph-node2][DEBUG ]   ],
[ceph-node2][DEBUG ] },
[ceph-node2][DEBUG ] "features": {
[ceph-node2][DEBUG ]   "quorum_con": "0",
[ceph-node2][DEBUG ]   "quorum_mon": [],
[ceph-node2][DEBUG ]   "required_con": "0",
[ceph-node2][DEBUG ]   "required_mon": []
[ceph-node2][DEBUG ] },
[ceph-node2][DEBUG ] "monmap": {
[ceph-node2][DEBUG ]   "created": "2024-02-26 10:25:31.068955",
[ceph-node2][DEBUG ]   "epoch": 0,
[ceph-node2][DEBUG ]   "features": {
[ceph-node2][DEBUG ]     "optional": [],
[ceph-node2][DEBUG ]     "persistent": []
[ceph-node2][DEBUG ]   },
[ceph-node2][DEBUG ]   "fsid": "c75bd276-c8d6-4b2c-b2a9-9062e2bf66a2",
[ceph-node2][DEBUG ]   "min_mon_release": 0,
[ceph-node2][DEBUG ]   "min_mon_release_name": "unknown",
[ceph-node2][DEBUG ]   "modified": "2024-02-26 10:25:31.068955",
[ceph-node2][DEBUG ]   "mons": [
[ceph-node2][DEBUG ]     {
[ceph-node2][DEBUG ]       "addr": "10.12.180.123:6789/0",
[ceph-node2][DEBUG ]       "name": "ceph-node1",
[ceph-node2][DEBUG ]       "public_addr": "10.12.180.123:6789/0",
[ceph-node2][DEBUG ]       "public_addrs": {
[ceph-node2][DEBUG ]         "addrvec": [
[ceph-node2][DEBUG ]           {
[ceph-node2][DEBUG ]             "addr": "10.12.180.123:3300",
[ceph-node2][DEBUG ]             "nonce": 0,
[ceph-node2][DEBUG ]             "type": "v2"
[ceph-node2][DEBUG ]           },
[ceph-node2][DEBUG ]           {
[ceph-node2][DEBUG ]             "addr": "10.12.180.123:6789",
[ceph-node2][DEBUG ]             "nonce": 0,
[ceph-node2][DEBUG ]             "type": "v1"
[ceph-node2][DEBUG ]           }
[ceph-node2][DEBUG ]         ]
[ceph-node2][DEBUG ]       }
[ceph-node2][DEBUG ]     }
[ceph-node2][DEBUG ]   ]
[ceph-node2][DEBUG ] }
```

```
[ceph-node2][DEBUG ]      }
[ceph-node2][DEBUG ]      ]
[ceph-node2][DEBUG ]      },
[ceph-node2][DEBUG ]      "rank": 0
[ceph-node2][DEBUG ]      },
[ceph-node2][DEBUG ]      {
[ceph-node2][DEBUG ]      "addr": "10.12.180.124:6789/0",
[ceph-node2][DEBUG ]      "name": "ceph-node2",
[ceph-node2][DEBUG ]      "public_addr": "10.12.180.124:6789/0",
[ceph-node2][DEBUG ]      "public_addrs": {
[ceph-node2][DEBUG ]      "addrvec": [
[ceph-node2][DEBUG ]      {
[ceph-node2][DEBUG ]      "addr": "10.12.180.124:3300",
[ceph-node2][DEBUG ]      "nonce": 0,
[ceph-node2][DEBUG ]      "type": "v2"
[ceph-node2][DEBUG ]      },
[ceph-node2][DEBUG ]      {
[ceph-node2][DEBUG ]      "addr": "10.12.180.124:6789",
[ceph-node2][DEBUG ]      "nonce": 0,
[ceph-node2][DEBUG ]      "type": "v1"
[ceph-node2][DEBUG ]      }
[ceph-node2][DEBUG ]      ],
[ceph-node2][DEBUG ]      "rank": 1
[ceph-node2][DEBUG ]      },
[ceph-node2][DEBUG ]      {
[ceph-node2][DEBUG ]      "addr": "0.0.0.0:0/1",
[ceph-node2][DEBUG ]      "name": "ceph-master",
[ceph-node2][DEBUG ]      "public_addr": "0.0.0.0:0/1",
[ceph-node2][DEBUG ]      "public_addrs": {
[ceph-node2][DEBUG ]      "addrvec": [
[ceph-node2][DEBUG ]      {
[ceph-node2][DEBUG ]      "addr": "0.0.0.0:0",
[ceph-node2][DEBUG ]      "nonce": 1,
[ceph-node2][DEBUG ]      "type": "v1"
[ceph-node2][DEBUG ]      }
[ceph-node2][DEBUG ]      ],
[ceph-node2][DEBUG ]      "rank": 2
[ceph-node2][DEBUG ]      }
[ceph-node2][DEBUG ]      ],
[ceph-node2][DEBUG ]      },
[ceph-node2][DEBUG ]      "name": "ceph-node2",
[ceph-node2][DEBUG ]      "outside_quorum": [],
[ceph-node2][DEBUG ]      "quorum": [],
[ceph-node2][DEBUG ]      "rank": 1,
[ceph-node2][DEBUG ]      "state": "electing",
[ceph-node2][DEBUG ]      "sync_provider": []
[ceph-node2][DEBUG ] }

[ceph-node2][DEBUG ] ****
[ceph-node2][INFO ] monitor: mon.ceph-node2 is running
[ceph-node2][INFO ] Running command: ceph --cluster=ceph --admin-daemon /var/run/ceph/ceph-mon.ceph-node2.asok mon_status
[ceph_deploy.mon][INFO ] processing monitor mon.ceph-master
[ceph-master][DEBUG ] connected to host: ceph-master
[ceph-master][DEBUG ] detect platform information from remote host
[ceph-master][DEBUG ] detect machine type
[ceph-master][DEBUG ] find the location of an executable
[ceph-master][INFO ] Running command: ceph --cluster=ceph --admin-daemon /var/run/ceph/ceph-mon.ceph-master.asok mon_status
[ceph_deploy.mon][WARNIN] mon.ceph-master monitor is not yet in quorum, tries left: 5
[ceph_deploy.mon][WARNIN] waiting 5 seconds before retrying
[ceph-master][INFO ] Running command: ceph --cluster=ceph --admin-daemon /var/run/ceph/ceph-mon.ceph-master.asok mon_status
```

```
[ceph_deploy.mon][WARNIN] mon.ceph-master monitor is not yet in quorum, tries left: 4
[ceph_deploy.mon][WARNIN] waiting 10 seconds before retrying
[ceph-master][INFO ] Running command: ceph --cluster=ceph --admin-daemon /var/run/ceph/ceph-mon.ceph-master.asok mon_status
[ceph_deploy.mon][INFO ] mon.ceph-master monitor has reached quorum!
[ceph_deploy.mon][INFO ] processing monitor mon.ceph-node1
[ceph-node1][DEBUG ] connected to host: ceph-node1
[ceph-node1][DEBUG ] detect platform information from remote host
[ceph-node1][DEBUG ] detect machine type
[ceph-node1][DEBUG ] find the location of an executable
[ceph-node1][INFO ] Running command: ceph --cluster=ceph --admin-daemon /var/run/ceph/ceph-mon.ceph-node1.asok mon_status
[ceph_deploy.mon][INFO ] mon.ceph-node1 monitor has reached quorum!
[ceph_deploy.mon][INFO ] processing monitor mon.ceph-node2
[ceph-node2][DEBUG ] connected to host: ceph-node2
[ceph-node2][DEBUG ] detect platform information from remote host
[ceph-node2][DEBUG ] detect machine type
[ceph-node2][DEBUG ] find the location of an executable
[ceph-node2][INFO ] Running command: ceph --cluster=ceph --admin-daemon /var/run/ceph/ceph-mon.ceph-node2.asok mon_status
[ceph_deploy.mon][INFO ] mon.ceph-node2 monitor has reached quorum!
[ceph_deploy.mon][INFO ] all initial monitors are running and have formed quorum
[ceph_deploy.mon][INFO ] Running gatherkeys...
[ceph_deploy.gatherkeys][INFO ] Storing keys in temp directory /tmp/tmpelzGR7
[ceph-master][DEBUG ] connected to host: ceph-master
[ceph-master][DEBUG ] detect platform information from remote host
[ceph-master][DEBUG ] detect machine type
[ceph-master][DEBUG ] get remote short hostname
[ceph-master][DEBUG ] fetch remote file
[ceph-master][INFO ] Running command: /usr/bin/ceph --connect-timeout=25 --cluster=ceph --admin-daemon=/var/run/ceph/ceph-mon.ceph-master.asok mon_status
[ceph-master][INFO ] Running command: /usr/bin/ceph --connect-timeout=25 --cluster=ceph --name mon. --keyring=/var/lib/ceph/mon/ceph-ceph-master/keyring auth get client.admin
[ceph-master][INFO ] Running command: /usr/bin/ceph --connect-timeout=25 --cluster=ceph --name mon. --keyring=/var/lib/ceph/mon/ceph-ceph-master/keyring auth get client.bootstrap-mds
[ceph-master][INFO ] Running command: /usr/bin/ceph --connect-timeout=25 --cluster=ceph --name mon. --keyring=/var/lib/ceph/mon/ceph-ceph-master/keyring auth get client.bootstrap-mgr
[ceph-master][INFO ] Running command: /usr/bin/ceph --connect-timeout=25 --cluster=ceph --name mon. --keyring=/var/lib/ceph/mon/ceph-ceph-master/keyring auth get client.bootstrap-osd
[ceph-master][INFO ] Running command: /usr/bin/ceph --connect-timeout=25 --cluster=ceph --name mon. --keyring=/var/lib/ceph/mon/ceph-ceph-master/keyring auth get client.bootstrap-rgw
[ceph_deploy.gatherkeys][INFO ] Storing ceph.client.admin.keyring
[ceph_deploy.gatherkeys][INFO ] Storing ceph.bootstrap-mds.keyring
[ceph_deploy.gatherkeys][INFO ] Storing ceph.bootstrap-mgr.keyring
[ceph_deploy.gatherkeys][INFO ] keyring 'ceph.mon.keyring' already exists
[ceph_deploy.gatherkeys][INFO ] Storing ceph.bootstrap-osd.keyring
[ceph_deploy.gatherkeys][INFO ] Storing ceph.bootstrap-rgw.keyring
[ceph_deploy.gatherkeys][INFO ] Destroy temp directory /tmp/tmpelzGR7
[root@ceph-master cephcluster]#
```

```
[root@ceph-master cephcluster]# ceph -s
2024-02-26 10:35:44.849 7f54ed84d700 -1 auth: unable to find a keyring on
/etc/ceph/ceph.client.admin.keyring,/etc/ceph/ceph.keyring,/etc/ceph/keyring,/etc/ceph/keyring.bin,: (2
) No such file      or directory
2024-02-26 10:35:44.849 7f54ed84d700 -1 AuthRegistry(0x7f54e80662b8) no keyring found at /etc/ceph/ceph.client.admin.keyring,/etc/ceph/ceph.keyring,/etc/ceph/keyring,/etc/ceph/keyring.bin,, d      i
sabling cephx
2024-02-26 10:35:44.879 7f54ed84d700 -1 auth: unable to find a keyring on
/etc/ceph/ceph.client.admin.keyring,/etc/ceph/ceph.keyring,/etc/ceph/keyring,/etc/ceph/keyring.bin,: (2
) No such file      or directory
2024-02-26 10:35:44.879 7f54ed84d700 -1 AuthRegistry(0x7f54e80c87b8) no keyring found at /etc/ceph/ceph.client.admin.keyring,/etc/ceph/ceph.keyring,/etc/ceph/keyring,/etc/ceph/keyring.bin,, d      i
sabling cephx
2024-02-26 10:35:44.880 7f54ed84d700 -1 auth: unable to find a keyring on
```

```
/etc/ceph/ceph.client.admin.keyring,/etc/ceph/ceph.keyring,/etc/ceph/keyring,/etc/ceph/keyring.bin,: (2
) No such file      or directory
2024-02-26 10:35:44.880 7f54ed84d700 -1 AuthRegistry(0x7f54ed84be78) no keyring found at /etc/c
eph/ceph.client.admin.keyring,/etc/ceph/ceph.keyring,/etc/ceph/keyring,/etc/ceph/keyring.bin,, d      i
sabling cephx
[errno 2] error connecting to the cluster
[root@ceph-master cephcluster]#
```

将admin秘钥配置文件copy到各节点

```
[root@ceph-master cephcluster]# ceph-deploy admin ceph-master ceph-node1 ceph-node2
[ceph_deploy.conf][DEBUG ] found configuration file at: /root/.cephdeploy.conf
[ceph_deploy.cli][INFO ] Invoked (2.0.1): /usr/bin/ceph-deploy admin ceph-master ceph-node1 ceph-
node2
[ceph_deploy.cli][INFO ] ceph-deploy options:
[ceph_deploy.cli][INFO ] username          : None
[ceph_deploy.cli][INFO ] verbose           : False
[ceph_deploy.cli][INFO ] overwrite_conf    : False
[ceph_deploy.cli][INFO ] quiet              : False
[ceph_deploy.cli][INFO ] cd_conf            : <ceph_deploy.conf.cephdeploy.Conf instance at 0x
7fe892f5c128>
[ceph_deploy.cli][INFO ] cluster            : ceph
[ceph_deploy.cli][INFO ] client             : ['ceph-master', 'ceph-node1', 'ceph-node2']
[ceph_deploy.cli][INFO ] func               : <function admin at 0x7fe893c695f0>
[ceph_deploy.cli][INFO ] ceph_conf          : None
[ceph_deploy.cli][INFO ] default_release    : False
[ceph_deploy.admin][DEBUG ] Pushing admin keys and conf to ceph-master
[ceph-master][DEBUG ] connected to host: ceph-master
[ceph-master][DEBUG ] detect platform information from remote host
[ceph-master][DEBUG ] detect machine type
[ceph-master][DEBUG ] write cluster configuration to /etc/ceph/{cluster}.conf
[ceph_deploy.admin][DEBUG ] Pushing admin keys and conf to ceph-node1
[ceph-node1][DEBUG ] connected to host: ceph-node1
[ceph-node1][DEBUG ] detect platform information from remote host
[ceph-node1][DEBUG ] detect machine type
[ceph-node1][DEBUG ] write cluster configuration to /etc/ceph/{cluster}.conf
[ceph_deploy.admin][DEBUG ] Pushing admin keys and conf to ceph-node2
[ceph-node2][DEBUG ] connected to host: ceph-node2
[ceph-node2][DEBUG ] detect platform information from remote host
[ceph-node2][DEBUG ] detect machine type
[ceph-node2][DEBUG ] write cluster configuration to /etc/ceph/{cluster}.conf
[root@ceph-master cephcluster]#
```

```
[root@ceph-master cephcluster]# ceph -s
cluster:
  id:  c75bd276-c8d6-4b2c-b2a9-9062e2bf66a2
  health: HEALTH_WARN
    mons are allowing insecure global_id reclaim
    clock skew detected on mon.ceph-node2

  services:
    mon: 3 daemons, quorum ceph-master,ceph-node1,ceph-node2 (age 11m)
    mgr: no daemons active
    osd: 0 osds: 0 up, 0 in

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

```
[root@ceph-master cephcluster]#
```

安装manager daemon监控节点

```
[root@ceph-master cephcluster]# ceph-deploy mgr create ceph-node1
[ceph_deploy.conf][DEBUG ] found configuration file at: /root/.cephdeploy.conf
[ceph_deploy.cli][INFO ] Invoked (2.0.1): /usr/bin/ceph-deploy mgr create ceph-node1
[ceph_deploy.cli][INFO ] ceph-deploy options:
[ceph_deploy.cli][INFO ] username : None
[ceph_deploy.cli][INFO ] verbose : False
[ceph_deploy.cli][INFO ] mgr : [('ceph-node1', 'ceph-node1')]
[ceph_deploy.cli][INFO ] overwrite_conf : False
[ceph_deploy.cli][INFO ] subcommand : create
[ceph_deploy.cli][INFO ] quiet : False
[ceph_deploy.cli][INFO ] cd_conf : <ceph_deploy.conf.cephdeploy Conf instance at 0x
7f6ecca75950>
[ceph_deploy.cli][INFO ] cluster : ceph
[ceph_deploy.cli][INFO ] func : <function mgr at 0x7f6ecd355500>
[ceph_deploy.cli][INFO ] ceph_conf : None
[ceph_deploy.cli][INFO ] default_release : False
[ceph_deploy.mgr][DEBUG ] Deploying mgr, cluster ceph hosts ceph-node1:ceph-node1
[ceph-node1][DEBUG ] connected to host: ceph-node1
[ceph-node1][DEBUG ] detect platform information from remote host
[ceph-node1][DEBUG ] detect machine type
[ceph_deploy.mgr][INFO ] Distro info: CentOS Linux 7.9.2009 Core
[ceph_deploy.mgr][DEBUG ] remote host will use systemd
[ceph_deploy.mgr][DEBUG ] deploying mgr bootstrap to ceph-node1
[ceph-node1][DEBUG ] write cluster configuration to /etc/ceph/{cluster}.conf
[ceph-node1][WARNIN] mgr keyring does not exist yet, creating one
[ceph-node1][DEBUG ] create a keyring file
[ceph-node1][DEBUG ] create path recursively if it doesn't exist
[ceph-node1][INFO ] Running command: ceph --cluster ceph --name client.bootstrap-mgr --keyring /v
ar/lib/ceph/bootstrap-mgr/ceph.keyring auth get-or-create mgr.ceph-node1 mon allow profile mgr osd
allow * mds allow * -o /var/lib/ceph/mgr/ceph-ceph-node1/keyring
[ceph-node1][INFO ] Running command: systemctl enable ceph-mgr@ceph-node1
[ceph-node1][WARNIN] Created symlink from /etc/systemd/system/ceph-mgr.target.wants/ceph-mgr
@ceph-node1.service to /usr/lib/systemd/system/ceph-mgr@.service.
[ceph-node1][INFO ] Running command: systemctl start ceph-mgr@ceph-node1
[ceph-node1][INFO ] Running command: systemctl enable ceph.target
[root@ceph-master cephcluster]#
```

```
[root@ceph-master cephcluster]# ceph -s
cluster:
  id: c75bd276-c8d6-4b2c-b2a9-9062e2bf66a2
  health: HEALTH_WARN
    OSD count 0 < osd_pool_default_size 3
    mons are allowing insecure global_id reclaim
    clock skew detected on mon.ceph-node2

services:
  mon: 3 daemons, quorum ceph-master,ceph-node1,ceph-node2 (age 13m)
  mgr: ceph-node1(active, since 32s)
  osd: 0 osds: 0 up, 0 in

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

[root@ceph-master cephcluster]#

```
[root@ceph-master cephcluster]# ceph-deploy mgr create ceph-node2
[ceph_deploy.conf][DEBUG ] found configuration file at: /root/.cephdeploy.conf
[ceph_deploy.cli][INFO ] Invoked (2.0.1): /usr/bin/ceph-deploy mgr create ceph-node2
[ceph_deploy.cli][INFO ] ceph-deploy options:
[ceph_deploy.cli][INFO ] username : None
[ceph_deploy.cli][INFO ] verbose : False
```

```

[ceph_deploy.cli][INFO ] mgr : [('ceph-node2', 'ceph-node2')]
[ceph_deploy.cli][INFO ] overwrite_conf : False
[ceph_deploy.cli][INFO ] subcommand : create
[ceph_deploy.cli][INFO ] quiet : False
[ceph_deploy.cli][INFO ] cd_conf : <ceph_deploy.conf.cephdeploy.Conf instance at 0x
7feadcf01950>
[ceph_deploy.cli][INFO ] cluster : ceph
[ceph_deploy.cli][INFO ] func : <function mgr at 0x7feadd7e1500>
[ceph_deploy.cli][INFO ] ceph_conf : None
[ceph_deploy.cli][INFO ] default_release : False
[ceph_deploy.mgr][DEBUG ] Deploying mgr, cluster ceph hosts ceph-node2:ceph-node2
[ceph-node2][DEBUG ] connected to host: ceph-node2
[ceph-node2][DEBUG ] detect platform information from remote host
[ceph-node2][DEBUG ] detect machine type
[ceph_deploy.mgr][INFO ] Distro info: CentOS Linux 7.9.2009 Core
[ceph_deploy.mgr][DEBUG ] remote host will use systemd
[ceph_deploy.mgr][DEBUG ] deploying mgr bootstrap to ceph-node2
[ceph-node2][DEBUG ] write cluster configuration to /etc/ceph/{cluster}.conf
[ceph-node2][WARNIN] mgr keyring does not exist yet, creating one
[ceph-node2][DEBUG ] create a keyring file
[ceph-node2][DEBUG ] create path recursively if it doesn't exist
[ceph-node2][INFO ] Running command: ceph --cluster ceph --name client.bootstrap-mgr --keyring /v
ar/lib/ceph/bootstrap-mgr/ceph.keyring auth get-or-create mgr.ceph-node2 mon allow profile mgr osd
allow * mds allow * -o /var/lib/ceph/mgr/ceph-ceph-node2/keyring
[ceph-node2][INFO ] Running command: systemctl enable ceph-mgr@ceph-node2
[ceph-node2][WARNIN] Created symlink from /etc/systemd/system/ceph-mgr.target.wants/ceph-mgr
@ceph-node2.service to /usr/lib/systemd/system/ceph-mgr@.service.
[ceph-node2][INFO ] Running command: systemctl start ceph-mgr@ceph-node2
[ceph-node2][INFO ] Running command: systemctl enable ceph.target
[root@ceph-master cephcluster]#

```

```

[root@ceph-master cephcluster]# ceph-deploy mgr create ceph-master
[ceph_deploy.conf][DEBUG ] found configuration file at: /root/.cephdeploy.conf
[ceph_deploy.cli][INFO ] Invoked (2.0.1): /usr/bin/ceph-deploy mgr create ceph-master
[ceph_deploy.cli][INFO ] ceph-deploy options:
[ceph_deploy.cli][INFO ] username : None
[ceph_deploy.cli][INFO ] verbose : False
[ceph_deploy.cli][INFO ] mgr : [('ceph-master', 'ceph-master')]
[ceph_deploy.cli][INFO ] overwrite_conf : False
[ceph_deploy.cli][INFO ] subcommand : create
[ceph_deploy.cli][INFO ] quiet : False
[ceph_deploy.cli][INFO ] cd_conf : <ceph_deploy.conf.cephdeploy.Conf instance at 0x
7ffb37ccf950>
[ceph_deploy.cli][INFO ] cluster : ceph
[ceph_deploy.cli][INFO ] func : <function mgr at 0x7ffb385af500>
[ceph_deploy.cli][INFO ] ceph_conf : None
[ceph_deploy.cli][INFO ] default_release : False
[ceph_deploy.mgr][DEBUG ] Deploying mgr, cluster ceph hosts ceph-master:ceph-master
[ceph-master][DEBUG ] connected to host: ceph-master
[ceph-master][DEBUG ] detect platform information from remote host
[ceph-master][DEBUG ] detect machine type
[ceph_deploy.mgr][INFO ] Distro info: CentOS Linux 7.9.2009 Core
[ceph_deploy.mgr][DEBUG ] remote host will use systemd
[ceph_deploy.mgr][DEBUG ] deploying mgr bootstrap to ceph-master
[ceph-master][DEBUG ] write cluster configuration to /etc/ceph/{cluster}.conf
[ceph-master][WARNIN] mgr keyring does not exist yet, creating one
[ceph-master][DEBUG ] create a keyring file
[ceph-master][DEBUG ] create path recursively if it doesn't exist
[ceph-master][INFO ] Running command: ceph --cluster ceph --name client.bootstrap-mgr --keyring /v
ar/lib/ceph/bootstrap-mgr/ceph.keyring auth get-or-create mgr.ceph-master mon allow profile mgr osd
allow * mds allow * -o /var/lib/ceph/mgr/ceph-ceph-master/keyring
[ceph-master][INFO ] Running command: systemctl enable ceph-mgr@ceph-master
[ceph-master][WARNIN] Created symlink from /etc/systemd/system/ceph-mgr.target.wants/ceph-mgr
@ceph-master.service to /usr/lib/systemd/system/ceph-mgr@.service.

```

```

[ceph-master][INFO ] Running command: systemctl start ceph-mgr@ceph-master
[ceph-master][INFO ] Running command: systemctl enable ceph.target
[root@ceph-master cephcluster]# ceph -s
cluster:
  id:  c75bd276-c8d6-4b2c-b2a9-9062e2bf66a2
  health: HEALTH_WARN
    OSD count 0 < osd_pool_default_size 3
    mons are allowing insecure global_id reclaim
    clock skew detected on mon.ceph-node2

services:
  mon: 3 daemons, quorum ceph-master,ceph-node1,ceph-node2 (age 15m)
  mgr: ceph-node1(active, since 118s), standbys: ceph-node2, ceph-master
  osd: 0 osds: 0 up, 0 in

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

[root@ceph-master cephcluster]#

```

部署OSD节点

```

[root@ceph-master cephcluster]# ceph-deploy osd create ceph-master --data /dev/sdb
[ceph_deploy.conf][DEBUG ] found configuration file at: /root/.cephdeploy.conf
[ceph_deploy.cli][INFO ] Invoked (2.0.1): /usr/bin/ceph-deploy osd create ceph-master --data /dev/sd
b
[ceph_deploy.cli][INFO ] ceph-deploy options:
[ceph_deploy.cli][INFO ] verbose          : False
[ceph_deploy.cli][INFO ] bluestore        : None
[ceph_deploy.cli][INFO ] cd_conf         : <ceph_deploy.conf.cephdeploy.Conf instance at 0x
7f374df98b48>
[ceph_deploy.cli][INFO ] cluster          : ceph
[ceph_deploy.cli][INFO ] fs_type          : xfs
[ceph_deploy.cli][INFO ] block_wal        : None
[ceph_deploy.cli][INFO ] default_release   : False
[ceph_deploy.cli][INFO ] username         : None
[ceph_deploy.cli][INFO ] journal          : None
[ceph_deploy.cli][INFO ] subcommand       : create
[ceph_deploy.cli][INFO ] host             : ceph-master
[ceph_deploy.cli][INFO ] filestore        : None
[ceph_deploy.cli][INFO ] func             : <function osd at 0x7f374dfc6c80>
[ceph_deploy.cli][INFO ] ceph_conf        : None
[ceph_deploy.cli][INFO ] zap_disk        : False
[ceph_deploy.cli][INFO ] data             : /dev/sdb
[ceph_deploy.cli][INFO ] block_db         : None
[ceph_deploy.cli][INFO ] dmcrypt          : False
[ceph_deploy.cli][INFO ] overwrite_conf   : False
[ceph_deploy.cli][INFO ] dmcrypt_key_dir  : /etc/ceph/dmcrypt-keys
[ceph_deploy.cli][INFO ] quiet            : False
[ceph_deploy.cli][INFO ] debug            : False
[ceph_deploy.osd][DEBUG ] Creating OSD on cluster ceph with data device /dev/sdb
[ceph-master][DEBUG ] connected to host: ceph-master
[ceph-master][DEBUG ] detect platform information from remote host
[ceph-master][DEBUG ] detect machine type
[ceph-master][DEBUG ] find the location of an executable
[ceph_deploy.osd][INFO ] Distro info: CentOS Linux 7.9.2009 Core
[ceph_deploy.osd][DEBUG ] Deploying osd to ceph-master
[ceph-master][DEBUG ] write cluster configuration to /etc/ceph/{cluster}.conf
[ceph-master][WARNIN] osd keyring does not exist yet, creating one
[ceph-master][DEBUG ] create a keyring file
[ceph-master][DEBUG ] find the location of an executable

```

```
[ceph-master][INFO ] Running command: /usr/sbin/ceph-volume --cluster ceph lvm create --bluestore  
--data /dev/sdb  
[ceph-master][WARNIN] Running command: /usr/bin/ceph-authtool --gen-print-key  
[ceph-master][WARNIN] Running command: /usr/bin/ceph --cluster ceph --name client.bootstrap-osd  
--keyring /var/lib/ceph/bootstrap-osd/ceph.keyring -i - osd new 3329dba7-2c9f-48ac-aded-  
18d8692ee03d  
[ceph-master][WARNIN] Running command: /usr/sbin/vgcreate --force --yes ceph-1f4db3a1-f95d-4e6  
8-9878-ba6f19484119 /dev/sdb  
[ceph-master][WARNIN] stdout: Physical volume "/dev/sdb" successfully created.  
[ceph-master][WARNIN] stdout: Volume group "ceph-1f4db3a1-f95d-4e68-9878-ba6f19484119" succ  
essfully created  
[ceph-master][WARNIN] Running command: /usr/sbin/lvcreate --yes -l 25599 -n osd-block-3329dba7-  
2c9f-48ac-aded-18d8692ee03d ceph-1f4db3a1-f95d-4e68-9878-ba6f19484119  
[ceph-master][WARNIN] stdout: Logical volume "osd-block-3329dba7-2c9f-48ac-aded-18d8692ee03  
d" created.  
[ceph-master][WARNIN] Running command: /usr/bin/ceph-authtool --gen-print-key  
[ceph-master][WARNIN] Running command: /usr/bin/mount -t tmpfs tmpfs /var/lib/ceph/osd/ceph-0  
[ceph-master][WARNIN] Running command: /usr/bin/chown -h ceph:ceph /dev/ceph-1f4db3a1-f95d-4  
e68-9878-ba6f19484119/osd-block-3329dba7-2c9f-48ac-aded-18d8692ee03d  
[ceph-master][WARNIN] Running command: /usr/bin/chown -R ceph:ceph /dev/dm-3  
[ceph-master][WARNIN] Running command: /usr/bin/ln -s /dev/ceph-1f4db3a1-f95d-4e68-9878-ba6f1  
9484119/osd-block-3329dba7-2c9f-48ac-aded-18d8692ee03d /var/lib/ceph/osd/ceph-0/block  
[ceph-master][WARNIN] Running command: /usr/bin/ceph --cluster ceph --name client.bootstrap-osd  
--keyring /var/lib/ceph/bootstrap-osd/ceph.keyring mon getmap -o /var/lib/ceph/osd/ceph-0/activate.m  
onmap  
[ceph-master][WARNIN] stderr: 2024-02-26 10:43:55.133 7f0508965700 -1 auth: unable to find a key  
ring on /etc/ceph/ceph.client.bootstrap-  
osd.keyring,/etc/ceph/ceph.keyring,/etc/ceph/keyring,/etc/ceph/keyring.bin,: (2) No such file or  
directory  
[ceph-master][WARNIN] 2024-02-26 10:43:55.133 7f0508965700 -1 AuthRegistry(0x7f05040662f8) n  
o keyring found at /etc/ceph/ceph.client.bootstrap-osd.keyring,/etc/ceph/ceph.keyring,/etc/ceph/keyrin  
g,/etc/ceph/keyring.bin,, disabling cephx  
[ceph-master][WARNIN] stderr: got monmap epoch 1  
[ceph-master][WARNIN] Running command: /usr/bin/ceph-authtool /var/lib/ceph/osd/ceph-0/keyring -  
-create-keyring --name osd.0 --add-key AQDq+tlfpThERAA2avw8tkt7Eh00hLCTJt53Q==  
[ceph-master][WARNIN] stdout: creating /var/lib/ceph/osd/ceph-0/keyring  
[ceph-master][WARNIN] stdout: added entity osd.0  
auth(key=AQDq+tlfpThERAA2avw8tkt7Eh00hLCTJt53Q==)  
[ceph-master][WARNIN] Running command: /usr/bin/chown -R ceph:ceph /var/lib/ceph/osd/ceph-0/ke  
yring  
[ceph-master][WARNIN] Running command: /usr/bin/chown -R ceph:ceph /var/lib/ceph/osd/ceph-0/  
[ceph-master][WARNIN] Running command: /usr/bin/ceph-osd --cluster ceph --osd-objectstore bluest  
ore --mkfs -i 0 --monmap /var/lib/ceph/osd/ceph-0/activate.monmap --keyfile - --osd-data /var/lib/ceph/  
osd/ceph-0/ --osd-uuid 3329dba7-2c9f-48ac-aded-18d8692ee03d --setuser ceph --setgroup ceph  
[ceph-master][WARNIN] stderr: 2024-02-26 10:43:55.704 7fb0308b7a80 -1 bluestore(/var/lib/ceph/os  
d/ceph-0) _read_fsid unparsable uuid  
[ceph-master][WARNIN] --> ceph-volume lvm prepare successful for: /dev/sdb  
[ceph-master][WARNIN] Running command: /usr/bin/chown -R ceph:ceph /var/lib/ceph/osd/ceph-0  
[ceph-master][WARNIN] Running command: /usr/bin/ceph-bluestore-tool --cluster=ceph prime-osd-  
dir --dev /dev/ceph-1f4db3a1-f95d-4e68-9878-ba6f19484119/osd-block-3329dba7-2c9f-48ac-aded-18  
d8692ee03d --path /var/lib/ceph/osd/ceph-0 --no-mon-config  
[ceph-master][WARNIN] Running command: /usr/bin/ln -snf /dev/ceph-1f4db3a1-f95d-4e68-9878-  
ba6f19484119/osd-block-3329dba7-2c9f-48ac-aded-18d8692ee03d /var/lib/ceph/osd/ceph-0/block  
[ceph-master][WARNIN] Running command: /usr/bin/chown -h ceph:ceph /var/lib/ceph/osd/ceph-  
0/block  
[ceph-master][WARNIN] Running command: /usr/bin/chown -R ceph:ceph /dev/dm-3  
[ceph-master][WARNIN] Running command: /usr/bin/chown -R ceph:ceph /var/lib/ceph/osd/ceph-0  
[ceph-master][WARNIN] Running command: /usr/bin/systemctl enable ceph-volume@lvm-0-3329dba  
7-2c9f-48ac-aded-18d8692ee03d  
[ceph-master][WARNIN] stderr: Created symlink from /etc/systemd/system/multi-user.target.wants/ce  
ph-volume@lvm-0-3329dba7-2c9f-48ac-aded-18d8692ee03d.service to /usr/lib/systemd/system/ceph  
-volume@.service.  
[ceph-master][WARNIN] Running command: /usr/bin/systemctl enable --runtime ceph-osd@0  
[ceph-master][WARNIN] stderr: Created symlink from /run/systemd/system/ceph-osd.target.wants/ce
```

```
ph-osd@0.service to /usr/lib/systemd/system/ceph-osd@.service.  
[ceph-master][WARNIN] Running command: /usr/bin/systemctl start ceph-osd@0  
[ceph-master][WARNIN] -> ceph-volume lvm activate successful for osd ID: 0  
[ceph-master][WARNIN] --> ceph-volume lvm create successful for: /dev/sdb  
[ceph-master][INFO ] checking OSD status...  
[ceph-master][DEBUG ] find the location of an executable  
[ceph-master][INFO ] Running command: /bin/ceph --cluster=ceph osd stat --format=json  
[ceph_deploy.osd][DEBUG ] Host ceph-master is now ready for osd use.  
[root@ceph-master cephcluster]#
```

```
[root@ceph-master cephcluster]# ceph -s  
cluster:  
  id:  c75bd276-c8d6-4b2c-b2a9-9062e2bf66a2  
  health: HEALTH_WARN  
    OSD count 1 < osd_pool_default_size 3  
    mons are allowing insecure global_id reclaim  
    clock skew detected on mon.ceph-node2  
  
services:  
  mon: 3 daemons, quorum ceph-master,ceph-node1,ceph-node2 (age 18m)  
  mgr: ceph-node1(active, since 5m), standbys: ceph-node2, ceph-master  
  osd: 1 osds: 1 up (since 26s), 1 in (since 26s)  
  
data:  
  pools: 0 pools, 0 pgs  
  objects: 0 objects, 0 B  
  usage: 1.0 GiB used, 99 GiB / 100 GiB avail  
  pgs:
```

```
[root@ceph-master cephcluster]# ceph osd tree  
ID CLASS WEIGHT TYPE NAME      STATUS REWEIGHT PRI-AFF  
-1     0.09769 root default  
-3     0.09769  host ceph-master  
0     hdd 0.09769    osd.0      up  1.00000 1.00000  
[root@ceph-master cephcluster]#
```

Repeat for other osd create

```
[ceph_deploy.osd][DEBUG ] Host ceph-master is now ready for osd use.  
[root@ceph-master cephcluster]# ceph -s  
cluster:  
  id:  c75bd276-c8d6-4b2c-b2a9-9062e2bf66a2  
  health: HEALTH_WARN  
    mons are allowing insecure global_id reclaim  
    clock skew detected on mon.ceph-node2  
  
services:  
  mon: 3 daemons, quorum ceph-master,ceph-node1,ceph-node2 (age 24m)  
  mgr: ceph-node1(active, since 11m), standbys: ceph-node2, ceph-master  
  osd: 5 osds: 5 up (since 6s), 5 in (since 6s)  
  
task status:  
  
data:  
  pools: 0 pools, 0 pgs  
  objects: 0 objects, 0 B  
  usage: 5.0 GiB used, 495 GiB / 500 GiB avail  
  pgs:  
  
[root@ceph-master cephcluster]#
```

对于其他节点的 osd 创建也是在主节点创建 (因为其他节点没有安装 ceph-deploy)

```
[root@ceph-master cephcluster]# ceph-deploy osd create ceph-node1 --data /dev/sdf  
[ceph_deploy.conf][DEBUG ] found configuration file at: /root/.cephdeploy.conf  
[ceph_deploy.cli][INFO ] Invoked (2.0.1): /usr/bin/ceph-deploy osd create ceph-node1 --data /dev/sdf
```

```
[ceph_deploy.cli][INFO ] ceph-deploy options:
[ceph_deploy.cli][INFO ] verbose           : False
[ceph_deploy.cli][INFO ] bluestore         : None
[ceph_deploy.cli][INFO ] cd_conf          : <ceph_deploy.conf.cephdeploy.Conf instance at 0x
7f738d37fb48>
[ceph_deploy.cli][INFO ] cluster          : ceph
[ceph_deploy.cli][INFO ] fs_type          : xfs
[ceph_deploy.cli][INFO ] block_wal        : None
[ceph_deploy.cli][INFO ] default_release   : False
[ceph_deploy.cli][INFO ] username         : None
[ceph_deploy.cli][INFO ] journal          : None
[ceph_deploy.cli][INFO ] subcommand       : create
[ceph_deploy.cli][INFO ] host             : ceph-node1
[ceph_deploy.cli][INFO ] filestore        : None
[ceph_deploy.cli][INFO ] func             : <function osd at 0x7f738d3adc80>
[ceph_deploy.cli][INFO ] ceph_conf        : None
[ceph_deploy.cli][INFO ] zap_disk         : False
[ceph_deploy.cli][INFO ] data             : /dev/sdf
[ceph_deploy.cli][INFO ] block_db         : None
[ceph_deploy.cli][INFO ] dmcrypt          : False
[ceph_deploy.cli][INFO ] overwrite_conf   : False
[ceph_deploy.cli][INFO ] dmcrypt_key_dir  : /etc/ceph/dmrypt-keys
[ceph_deploy.cli][INFO ] quiet            : False
[ceph_deploy.cli][INFO ] debug            : False
[ceph_deploy.osd][DEBUG ] Creating OSD on cluster ceph with data device /dev/sdf
[ceph-node1][DEBUG ] connected to host: ceph-node1
[ceph-node1][DEBUG ] detect platform information from remote host
[ceph-node1][DEBUG ] detect machine type
[ceph-node1][DEBUG ] find the location of an executable
[ceph_deploy.osd][INFO ] Distro info: CentOS Linux 7.9.2009 Core
[ceph_deploy.osd][DEBUG ] Deploying osd to ceph-node1
[ceph-node1][DEBUG ] write cluster configuration to /etc/ceph/{cluster}.conf
[ceph-node1][WARNIN] osd keyring does not exist yet, creating one
[ceph-node1][DEBUG ] create a keyring file
[ceph-node1][DEBUG ] find the location of an executable
[ceph-node1][INFO ] Running command: /usr/sbin/ceph-volume --cluster ceph lvm create --bluestore
--data /dev/sdf
[ceph-node1][WARNIN] Running command: /bin/ceph-authtool --gen-print-key
[ceph-node1][WARNIN] Running command: /bin/ceph --cluster ceph --name client.bootstrap-osd --ke
yring /var/lib/ceph/bootstrap-osd/ceph.keyring -i - osd new 12feeba3-6e93-479e-b64d-7dbeee50c4d2
[ceph-node1][WARNIN] Running command: /usr/sbin/vgcreate --force --yes ceph-eda4fc67-5bbb-471
7-95e7-5e99ab9f09f8 /dev/sdf
[ceph-node1][WARNIN] stdut: Physical volume "/dev/sdf" successfully created.
[ceph-node1][WARNIN] stdut: Volume group "ceph-eda4fc67-5bbb-4717-95e7-5e99ab9f09f8" succ
essfully created
[ceph-node1][WARNIN] Running command: /usr/sbin/lvcreate --yes -l 25599 -n osd-block-12feeba3-6
e93-479e-b64d-7dbeee50c4d2 ceph-eda4fc67-5bbb-4717-95e7-5e99ab9f09f8
[ceph-node1][WARNIN] stdut: Logical volume "osd-block-12feeba3-6e93-479e-b64d-7dbeee50c4d2" cre
ated.
[ceph-node1][WARNIN] Running command: /bin/ceph-authtool --gen-print-key
[ceph-node1][WARNIN] Running command: /bin/mount -t tmpfs tmpfs /var/lib/ceph/osd/ceph-5
[ceph-node1][WARNIN] Running command: /bin/chown -h ceph:ceph /dev/ceph-eda4fc67-5bbb-4717
-95e7-5e99ab9f09f8/ceph-block-12feeba3-6e93-479e-b64d-7dbeee50c4d2 /var/lib/ceph/osd/ceph-5/block
[ceph-node1][WARNIN] Running command: /bin/ceph --cluster ceph --name client.bootstrap-osd --ke
yring /var/lib/ceph/bootstrap-osd/ceph.keyring mon getmap -o /var/lib/ceph/osd/ceph-5/activate.monm
ap
[ceph-node1][WARNIN] stderr: 2024-02-26 10:51:16.499 7f8c7c007700 -1 auth: unable to find a keyr
ing on /etc/ceph/ceph.client.bootstrap-osd.keyring,/etc/ceph/ceph.keyring,/etc/ceph/keyring,/etc/ceph/
keyring.bin,: (2) No such file or directory
[ceph-node1][WARNIN] 2024-02-26 10:51:16.499 7f8c7c007700 -1 AuthRegistry(0x7f8c740662f8) no
keyring found at /etc/ceph/ceph.client.bootstrap-
```

```

osd.keyring,/etc/ceph/ceph.keyring,/etc/ceph/keyring,/etc/ceph/keyring.bin,, disabling cephx
[ceph-node1][WARNIN] stderr: got monmap epoch 1
[ceph-node1][WARNIN] Running command: /bin/ceph-authtool /var/lib/ceph/osd/ceph-5/keyring --create-keyring --name osd.5 --add-key AQCj/NtliVFHFhAAlejVmAdKBnCyBxP4x6mINw==
[ceph-node1][WARNIN] stdout: creating /var/lib/ceph/osd/ceph-5/keyring
[ceph-node1][WARNIN] stdout: added entity osd.5 auth(key=AQCj/NtliVFHFhAAlejVmAdKBnCyBxP4x6mINw==)
[ceph-node1][WARNIN] Running command: /bin/chown -R ceph:ceph /var/lib/ceph/osd/ceph-5/keyring
[ceph-node1][WARNIN] Running command: /bin/chown -R ceph:ceph /var/lib/ceph/osd/ceph-5/
[ceph-node1][WARNIN] Running command: /bin/ceph-osd --cluster ceph --osd-objectstore bluestore -mkfs -i 5 --monmap /var/lib/ceph/osd/ceph-5/activate.monmap --keyfile --osd-data /var/lib/ceph/osd/ceph-5 --osd-uuid 12feeba3-6e93-479e-b64d-7dbeee50c4d2 --setuser ceph --setgroup ceph
[ceph-node1][WARNIN] stderr: 2024-02-26 10:51:16.952 7f54dd1e7a80 -1 bluestore(/var/lib/ceph/osd/ceph-5/) _read_fsid unparsable uuid
[ceph-node1][WARNIN] --> ceph-volume lvm prepare successful for: /dev/sdf
[ceph-node1][WARNIN] Running command: /bin/chown -R ceph:ceph /var/lib/ceph/osd/ceph-5
[ceph-node1][WARNIN] Running command: /bin/ceph-bluestore-tool --cluster=ceph prime-osd-dir --dev /dev/ceph-eda4fc67-5bbb-4717-95e7-5e99ab9f09f8/osd-block-12feeba3-6e93-479e-b64d-7dbeee50c4d2 --path /var/lib/ceph/osd/ceph-5 --no-mon-config
[ceph-node1][WARNIN] Running command: /bin/ln -snf /dev/ceph-eda4fc67-5bbb-4717-95e7-5e99ab9f09f8/osd-block-12feeba3-6e93-479e-b64d-7dbeee50c4d2 /var/lib/ceph/osd/ceph-5/block
[ceph-node1][WARNIN] Running command: /bin/chown -h ceph:ceph /var/lib/ceph/osd/ceph-5/block
[ceph-node1][WARNIN] Running command: /bin/chown -R ceph:ceph /dev/dm-3
[ceph-node1][WARNIN] Running command: /bin/chown -R ceph:ceph /var/lib/ceph/osd/ceph-5
[ceph-node1][WARNIN] Running command: /bin/systemctl enable ceph-volume@lvm-5-12feeba3-6e93-479e-b64d-7dbeee50c4d2
[ceph-node1][WARNIN] stderr: Created symlink from /etc/systemd/system/multi-user.target.wants/ceph-volume@lvm-5-12feeba3-6e93-479e-b64d-7dbeee50c4d2.service to /usr/lib/systemd/system/ceph-volume@.service.
[ceph-node1][WARNIN] Running command: /bin/systemctl enable --runtime ceph-osd@5
[ceph-node1][WARNIN] stderr: Created symlink from /run/systemd/system/ceph-osd.target.wants/ceph-osd@5.service to /usr/lib/systemd/system/ceph-osd@.service.
[ceph-node1][WARNIN] Running command: /bin/systemctl start ceph-osd@5
[ceph-node1][WARNIN] --> ceph-volume lvm activate successful for osd ID: 5
[ceph-node1][WARNIN] --> ceph-volume lvm create successful for: /dev/sdf
[ceph-node1][INFO ] checking OSD status...
[ceph-node1][DEBUG ] find the location of an executable
[ceph-node1][INFO ] Running command: /bin/ceph --cluster=ceph osd stat --format=json
[ceph_deploy.osd][DEBUG ] Host ceph-node1 is now ready for osd use.

```

```

[root@ceph-master cephcluster]# ceph -s
cluster:
  id:  c75bd276-c8d6-4b2c-b2a9-9062e2bf66a2
  health: HEALTH_WARN
    mons are allowing insecure global_id reclaim
    clock skew detected on mon.ceph-node2

services:
  mon: 3 daemons, quorum ceph-master,ceph-node1,ceph-node2 (age 25m)
  mgr: ceph-node1(active, since 12m), standbys: ceph-node2, ceph-master
  osd: 6 osds: 6 up (since 5s), 6 in (since 5s)

data:
  pools: 0 pools, 0 pgs
  objects: 0 objects, 0 B
  usage: 6.0 GiB used, 594 GiB / 600 GiB avail
  pgs:

[root@ceph-master cephcluster]#

```

三个节点，每个节点5块100G的数据盘，所以有 15 个 OSD

```
[root@ceph-master cephcluster]# ceph -s
```

```

cluster:
id: c75bd276-c8d6-4b2c-b2a9-9062e2bf66a2
health: HEALTH_WARN
    mons are allowing insecure global_id reclaim     >>> 注此处还有告警
    clock skew detected on mon.ceph-node2

services:
mon: 3 daemons, quorum ceph-master,ceph-node1,ceph-node2 (age 29m)
mgr: ceph-node1(active, since 16m), standbys: ceph-node2, ceph-master
osd: 15 osds: 15 up (since 17s), 15 in (since 17s)

task status:

data:
pools: 0 pools, 0 pgs
objects: 0 objects, 0 B
usage: 15 GiB used, 1.5 TiB / 1.5 TiB avail
pgs:

[root@ceph-master cephcluster]#

```

```

[root@ceph-master cephcluster]# ceph osd tree
ID CLASS WEIGHT TYPE NAME      STATUS REWEIGHT PRI-AFF
-1   1.46530 root default
-3   0.48843  host ceph-master
  0 hdd 0.09769  osd.0      up 1.00000 1.00000
  1 hdd 0.09769  osd.1      up 1.00000 1.00000
  2 hdd 0.09769  osd.2      up 1.00000 1.00000
  3 hdd 0.09769  osd.3      up 1.00000 1.00000
  4 hdd 0.09769  osd.4      up 1.00000 1.00000
-5   0.48843  host ceph-node1
  5 hdd 0.09769  osd.5      up 1.00000 1.00000
  6 hdd 0.09769  osd.6      up 1.00000 1.00000
  7 hdd 0.09769  osd.7      up 1.00000 1.00000
  8 hdd 0.09769  osd.8      up 1.00000 1.00000
  9 hdd 0.09769  osd.9      up 1.00000 1.00000
-7   0.48843  host ceph-node2
 10 hdd 0.09769  osd.10     up 1.00000 1.00000
 11 hdd 0.09769  osd.11     up 1.00000 1.00000
 12 hdd 0.09769  osd.12     up 1.00000 1.00000
 13 hdd 0.09769  osd.13     up 1.00000 1.00000
 14 hdd 0.09769  osd.14     up 1.00000 1.00000

```

[root@ceph-master cephcluster]#

禁用不安全模式后，集群恢复正常

```

[root@ceph-master cephcluster]# ceph config set mon auth_allow_insecure_global_id_reclaim f
also
[root@ceph-master cephcluster]#

```

```

[root@ceph-master cephcluster]# ceph -s
cluster:
id: c75bd276-c8d6-4b2c-b2a9-9062e2bf66a2
health: HEALTH_OK

services:
mon: 3 daemons, quorum ceph-master,ceph-node1,ceph-node2 (age 38m)
mgr: ceph-node1(active, since 25m), standbys: ceph-node2, ceph-master
osd: 15 osds: 15 up (since 8m), 15 in (since 8m)

data:
pools: 0 pools, 0 pgs
objects: 0 objects, 0 B
usage: 15 GiB used, 1.5 TiB / 1.5 TiB avail
pgs:

[root@ceph-master cephcluster]#

```

查看 mon 状态

方式一

```
[root@ceph-master cephcluster]# ceph mon stat
e1: 3 mons at [ceph-master=[v2:10.12.180.122:3300/0,v1:10.12.180.122:6789/0],ceph-node1=
[v2:10.12.180.123:3300/0,v1:10.12.180.123:6789/0],ceph-node2=
[v2:10.12.180.124:3300/0,v1:10.12.180.124:6789/0]}, election epoch 8, leader 0 ceph-master, quoru
m 0,1,2 ceph-master,ceph-node1,ceph-node2
```

方式二

```
[root@ceph-master cephcluster]# ceph mon dump
epoch 1
fsid c75bd276-c8d6-4b2c-b2a9-9062e2bf66a2
last_changed 2024-02-26 10:25:25.099139
created 2024-02-26 10:25:25.099139
min_mon_release 14 (nautilus)
0: [v2:10.12.180.122:3300/0,v1:10.12.180.122:6789/0] mon.ceph-master
1: [v2:10.12.180.123:3300/0,v1:10.12.180.123:6789/0] mon.ceph-node1
2: [v2:10.12.180.124:3300/0,v1:10.12.180.124:6789/0] mon.ceph-node2
dumped monmap epoch 1
[root@ceph-master cephcluster]#
```

方式三

```
[root@ceph-master cephcluster]# ceph quorum_status --format json-pretty
```

```
{
  "election_epoch": 8,
  "quorum": [
    0,
    1,
    2
  ],
  "quorum_names": [
    "ceph-master",
    "ceph-node1",
    "ceph-node2"
  ],
  "quorum_leader_name": "ceph-master",
  "quorum_age": 2492,
  "monmap": {
    "epoch": 1,
    "fsid": "c75bd276-c8d6-4b2c-b2a9-9062e2bf66a2",
    "modified": "2024-02-26 10:25:25.099139",
    "created": "2024-02-26 10:25:25.099139",
    "min_mon_release": 14,
    "min_mon_release_name": "nautilus",
    "features": {
      "persistent": [
        "kraken",
        "luminous",
        "mimic",
        "osdmap-prune",
        "nautilus"
      ],
      "optional": []
    },
    "mons": [
      {
        "rank": 0,
        "name": "ceph-master",
        "public_addrs": {
          "v2": "10.12.180.122:3300",
          "v1": "10.12.180.122:6789"
        }
      }
    ]
  }
}
```

```

"addrvec": [
    {
        "type": "v2",
        "addr": "10.12.180.122:3300",
        "nonce": 0
    },
    {
        "type": "v1",
        "addr": "10.12.180.122:6789",
        "nonce": 0
    }
]
},
{
    "addr": "10.12.180.122:6789/0",
    "public_addr": "10.12.180.122:6789/0"
},
{
    "rank": 1,
    "name": "ceph-node1",
    "public_addrs": {
        "addrvec": [
            {
                "type": "v2",
                "addr": "10.12.180.123:3300",
                "nonce": 0
            },
            {
                "type": "v1",
                "addr": "10.12.180.123:6789",
                "nonce": 0
            }
        ]
    },
    "addr": "10.12.180.123:6789/0",
    "public_addr": "10.12.180.123:6789/0"
},
{
    "rank": 2,
    "name": "ceph-node2",
    "public_addrs": {
        "addrvec": [
            {
                "type": "v2",
                "addr": "10.12.180.124:3300",
                "nonce": 0
            },
            {
                "type": "v1",
                "addr": "10.12.180.124:6789",
                "nonce": 0
            }
        ]
    },
    "addr": "10.12.180.124:6789/0",
    "public_addr": "10.12.180.124:6789/0"
}
]
}
}

[root@ceph-master cephcluster]#

```

配置 ceph-dashboard 图形化界面

https://blog.51cto.com/u_14035463/5585093

每个节点安装 ceph-dashboard

```
[root@ceph-master cephcluster]# yum install ceph-mgr-dashboard -y
Loaded plugins: fastestmirror, langpacks, priorities
Loading mirror speeds from cached hostfile
 * base: mirrors.jlu.edu.cn
 * epel: mirrors.bfsu.edu.cn
 * extras: mirrors.jlu.edu.cn
 * updates: mirrors.jlu.edu.cn
Resolving Dependencies
--> Running transaction check
---> Package ceph-mgr-dashboard.noarch 2:14.2.22-0.el7 will be installed
---> Processing Dependency: ceph-grafana-dashboards = 2:14.2.22-0.el7 for package: 2:ceph-mgr-da
shboard-14.2.22-0.el7.noarch
---> Processing Dependency: python-routes for package: 2:ceph-mgr-dashboard-14.2.22-0.el7.noarch
---> Running transaction check
---> Package ceph-grafana-dashboards.noarch 2:14.2.22-0.el7 will be installed
---> Package python-routes.noarch 0:1.13-2.el7 will be installed
---> Processing Dependency: python-repoze-lru for package: python-routes-1.13-2.el7.noarch
---> Running transaction check
---> Package python-repoze-lru.noarch 0:0.4-3.el7 will be installed
---> Finished Dependency Resolution
```

Dependencies Resolved

Transaction Summary

Install 1 Package (+3 Dependent packages)

Total download size: 4.6 M

Installed size: 19 M

Downloading packages:

(1/4): python-repoze-lru-0.4-

3.el7.noarch.rpm

13 kB 00:00:00

(2/4): python-routes-1.13-

2.el7.noarch.rpm

| 640 kB 00:00:00

(3/4): ceph-grafana-dashboards-14.2.22-

0.el7.noarch.rpm

00:00:02

(4/4): ceph-mgr-dashboard-14.2.22-

0.el7.noarch.rpm

MB 00:00:04

```
-----  
Total  
    1.1 MB/s | 4.6 MB 00:00:04  
Running transaction check  
Running transaction test  
Transaction test succeeded  
Running transaction  
  Installing : python-repoze-lru-0.4-  
3.el7.noarch  
1/4  
  Installing : python-routes-1.13-  
2.el7.noarch  
2/4  
  Installing : 2:ceph-grafana-dashboards-14.2.22-  
0.el7.noarch  
    Installing : 2:ceph-mgr-dashboard-14.2.22-  
0.el7.noarch  
3/4  
  Verifying : 2:ceph-grafana-dashboards-14.2.22-  
0.el7.noarch  
    Verifying : 2:ceph-mgr-dashboard-14.2.22-  
0.el7.noarch  
4/4  
  Verifying : python-routes-1.13-  
2.el7.noarch  
3/4  
  Verifying : python-repoze-lru-0.4-  
3.el7.noarch  
4/4  
  
Installed:  
  ceph-mgr-dashboard.noarch 2:14.2.22-0.el7  
  
Dependency Installed:  
  ceph-grafana-dashboards.noarch 2:14.2.22-0.el7           python-repoze-lru.noarch 0:0.4-3.  
el7          python-routes.noarch 0:1.13-2.el7  
  
Complete!  
[root@ceph-master cephcluster]#
```

开启 ceph mgr 功能

```
[root@ceph-master cephcluster]# ceph mgr module enable dashboard  
[root@ceph-master cephcluster]#
```

只需在一个节点开启就行，在其他节点开启会提示已经开启该功能

```
[root@ceph-node2 yum.repos.d]# ceph mgr module enable dashboard  
module 'dashboard' is already enabled  
[root@ceph-node2 yum.repos.d]#
```

生成并安装自签名的证书

```
[root@ceph-master cephcluster]# ceph dashboard create-self-signed-cert  
Self-signed certificate created  
[root@ceph-master cephcluster]#
```

```
[root@ceph-master cephcluster]# ceph -s  
cluster:  
  id:  c75bd276-c8d6-4b2c-b2a9-9062e2bf66a2  
  health: HEALTH_OK  
  
services:  
  mon: 3 daemons, quorum ceph-master,ceph-node1,ceph-node2 (age 3h)  
  mgr: ceph-node1(active, since 2h), standbys: ceph-master, ceph-node2  >>> GUI 会在 mgr 节点  
开启  
  osd: 15 osds: 15 up (since 2h), 15 in (since 2h)
```

```
data:  
pools: 0 pools, 0 pgs  
objects: 0 objects, 0 B  
usage: 15 GiB used, 1.5 TiB / 1.5 TiB avail  
pgs:
```

```
[root@ceph-master cephcluster]#
```

这时登录 ceph-node1, 可以访问图形界面

```
[root@ceph-master cephcluster]# ceph mgr services  
{  
    "dashboard": "https://ceph-node1.h3cu.com:8443/"  
}  
[root@ceph-master cephcluster]#
```

<http://10.12.180.123:8443>

接下来配置登录用户名和密码

```
[root@ceph-master cephcluster]# ceph dashboard ac-user-create admin administrator -i cephps  
w.txt  
{"username": "admin", "lastUpdate": 1708927490, "name": null, "roles": ["administrator"], "password":  
"$2b$12$X9PZ5QuGABnYeOYMaCLueeq5yyRj6FBecGp8al4jl3GPWMvfzb4xa", "email": null}  
[root@ceph-master cephcluster]#
```

设置的 admin 用户名的密码为 HPinside!

```
[root@ceph-master cephcluster]# more cephpsw.txt  
HPinside!  
[root@ceph-master cephcluster]#
```

用户名, 密码登录 GUI

<http://10.12.180.123:8443>

```
username: admin  
password: HPinside!
```

图形化界面的缺省端口为 8443 , 如果指定为其他端口

```
[root@ceph-master cephcluster]# ceph config set mgr mgr/dashboard/server_port xxxx
```

指定图形界面访问的地址

```
[root@ceph-master cephcluster]# ceph config set mgr mgr/dashboard/server_addr xx.xx.xx.xx
```

关闭HTTPS

```
[root@ceph-master cephcluster]# ceph config set mgr mgr/dashboard/ssl false
```

开启 Object Gateway 管理功能

登录图形界面, 选择 Object Gaway

Information

No RGW service is running.

Please consult the documentation on how to configure and enable the Object Gateway management functionality.

接下来创建一个带system选项的用户

```
[root@ceph-master cephcluster]# radosgw-admin user create --uid ky_rgw --display-name="Kong  
Ying RGW" --system  
{  
    "user_id": "ky_rgw",  
    "display_name": "KongYing RGW",  
    "email": "",  
    "suspended": 0,
```

```

"max_buckets": 1000,
"subusers": [],
"keys": [
  {
    "user": "ky_rgw",
    "access_key": "QXTLP6XLHXUK425QZ5JE",
    "secret_key": "0byyvelODLzP53KCAi0Ecj48QTtjYRnPiufG6Ygm"
  }
],
"swift_keys": [],
"caps": [],
"op_mask": "read, write, delete",
"system": "true",           >>> Enable system option
"default_placement": "",
"default_storage_class": "",
"placement_tags": [],
"bucket_quota": {
  "enabled": false,
  "check_on_raw": false,
  "max_size": -1,
  "max_size_kb": 0,
  "max_objects": -1
},
"user_quota": {
  "enabled": false,
  "check_on_raw": false,
  "max_size": -1,
  "max_size_kb": 0,
  "max_objects": -1
},
"temp_url_keys": [],
"type": "rgw",
"mfa_ids": []
}

```

[root@ceph-master cephcluster]#

```

[root@ceph-master cephcluster]# echo QXTLP6XLHXUK425QZ5JE >rgw_access.key
[root@ceph-master cephcluster]# echo 0byyvelODLzP53KCAi0Ecj48QTtjYRnPiufG6Ygm > rgw_
secret.key
[root@ceph-master cephcluster]# ceph dashboard set-rgw-api-access-key -i rgw_access.key
Option RGW_API_ACCESS_KEY updated
[root@ceph-master cephcluster]# ceph dashboard set-rgw-api-secret-key -i rgw_secret.key
Option RGW_API_SECRET_KEY updated
[root@ceph-master cephcluster]#

```

根据带 system 选项创建的用户，就可以正常打开 GUI -> Object Gateway 界面

配置关键点

希望能实现所有 X10000 和 OneStor的所有软件功能。