

IBM Storage Ceph 7

Installation Guide 2023-11-02



Important

DRAFT BETA 1 PROVIDED WITH NO SUPPORT: Do not upgrade a production installation to a beta release. Upgrade support is not provided between beta releases. No automated upgrade path is provided from a beta release to the final GA release. The cluster will need to be rebuilt for production use.

Last updated: 2023-11-02

Contents

Important.....	i
Figures.....	v
Tables.....	vii
Installing.....	9
Initial installation.....	9
Prerequisites.....	9
cephadm utility.....	9
Configuring NTP using chrony.....	12
Registering the IBM Storage Ceph nodes.....	12
Configuring Ansible inventory location.....	14
Creating an Ansible user with sudo access.....	15
Configuring SSH.....	16
Enabling SSH login as root user on Red Hat Enterprise Linux 9.....	18
Running the preflight playbook.....	19
Bootstrapping a new storage cluster.....	20
Disconnected installation.....	30
Launching the cephadm shell.....	39
cephadm commands.....	40
Verifying the cluster installation.....	44
Adding hosts.....	45
Removing hosts.....	49
Labeling hosts.....	50
Adding Monitor service.....	54
Setting up the admin node.....	56
Adding Manager service.....	58
Adding OSDs.....	58
Purging the Ceph storage cluster.....	59
Deploying client nodes.....	60
Managing an IBM Storage Ceph cluster using cephadm-ansible modules.....	62
cephadm-ansible modules.....	62
cephadm-ansible modules options.....	62
Bootstrapping a storage cluster using the cephadm_ansible modules.....	66
Adding or removing hosts using the ceph_orch_host module.....	68
Setting configuration options using the ceph_config module.....	72
Applying a service specification using the ceph_orch_apply module.....	73
Managing Ceph daemon states using the ceph_orch_daemon module.....	74
Comparison between Ceph Ansible and Cephadm.....	76
What to do next? Day 2.....	77

Figures

1. Ceph storage cluster deployment.....11

Tables

1. Custom image details for monitoring stack.....	33
2. The cephadm commands.....	40
3. Available options for the cephadm_bootstrap module.....	63
4. Available options for the ceph_orch_host module.....	64
5. Available options for the ceph_config module.....	65
6. Available options for the ceph_orch_apply module.....	65
7. Available options for the ceph_orch_daemon module.....	65
8. Available options for the cephadm_registry_login module.....	66
9. Day one operations.....	76
10. Day two operations.....	77

Installing

This information provides instructions on installing IBM Storage Ceph on Red Hat Enterprise Linux running on AMD64 and Intel 64 architectures.

Initial installation

As a storage administrator, you can use the `cephadm` utility to deploy new IBM Storage Ceph clusters.

The `cephadm` utility manages the entire life cycle of a Ceph cluster. Installation and management tasks comprise two types of operations:

- Day One operations involve installing and bootstrapping a bare-minimum, containerized Ceph storage cluster, running on a single node. Day One also includes deploying the Monitor and Manager daemons and adding Ceph OSDs.
- Day Two operations use the Ceph orchestration interface, `cephadm orch`, or the IBM Storage Ceph Dashboard to expand the storage cluster by adding other Ceph services to the storage cluster.

Prerequisites

- At least one running virtual machine (VM) or bare-metal server with an active internet connection.
- Ansible 2.9 or later.
- Server and Ceph repositories enabled.
- Root-level access to all nodes.
- An active IBM Network or service account to access the IBM Registry.
- Remove troubling configurations in iptables so that refresh of iptables services does not cause issues to the cluster. For an example, see [Verifying firewall rules are configured for default Ceph ports](#).

For the latest supported Red Hat Enterprise Linux versions, see [Compatibility matrix](#).

cephadm utility

The `cephadm` utility deploys and manages a Ceph storage cluster. It is tightly integrated with both the command-line interface (CLI) and the IBM Storage Ceph Dashboard web interface, so that you can manage storage clusters from either environment. `cephadm` uses SSH to connect to hosts from the manager daemon to add, remove, or update Ceph daemon containers. It does not rely on external configuration or orchestration tools such as Ansible or Rook.

Note: The `cephadm` utility is available after running the preflight playbook on a host.

The `cephadm` utility consists of two main components:

- The `cephadm` shell.
- The `cephadm` orchestrator.

The cephadm shell

The `cephadm` shell launches a bash shell within a container. This enables you to perform “Day One” cluster setup tasks, such as installation and bootstrapping, and to invoke `ceph` commands.

There are two ways to invoke the `cephadm` shell:

- Enter `cephadm shell` at the system prompt:

Example

```
[root@host01 ~]# cephadm shell  
[ceph: root@host01 /]# ceph -s
```

- At the system prompt, type `cephadm shell` and the command you want to run:

Example

```
[root@host01 ~]# cephadm shell ceph -s
```

Note: If the node contains configuration and keyring files in `/etc/ceph/`, the container environment uses the values in those files as defaults for the `cephadm shell`. However, if you execute the `cephadm shell` on a Ceph Monitor node, the `cephadm shell` inherits its default configuration from the Ceph Monitor container, instead of using the default configuration.

The cephadm orchestrator

The `cephadm` orchestrator enables you to perform “Day Two” Ceph functions, such as expanding the storage cluster and provisioning Ceph daemons and services. You can use the `cephadm` orchestrator through either the command-line interface (CLI) or the web-based IBM Storage Ceph Dashboard. Orchestrator commands take the form `ceph orch`.

The `cephadm` script interacts with the Ceph orchestration module used by the Ceph Manager.

How cephadm works

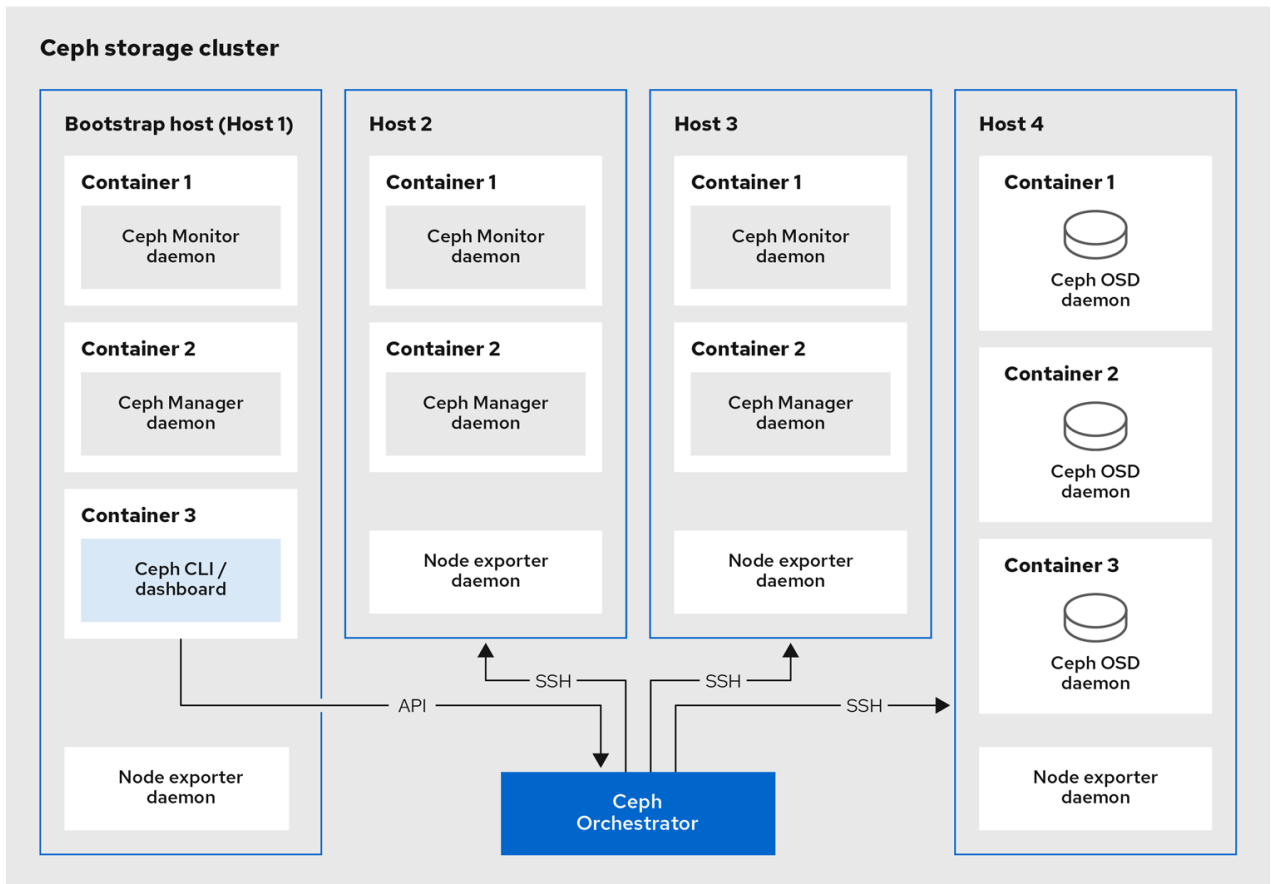
Use the `cephadm` command to manage the full lifecycle of an IBM Storage Ceph cluster.

The `cephadm` command can perform the following operations:

- Bootstrap a new IBM Storage Ceph cluster.
- Launch a containerized shell that works with the IBM Storage Ceph command-line interface (CLI).
- Aid in debugging containerized daemons.

The `cephadm` command uses `ssh` to communicate with the nodes in the storage cluster. This allows you to add, remove, or update IBM Storage Ceph containers without using external tools. Generate the `ssh` key pair during the bootstrapping process, or use your own `ssh` key.

The `cephadm` bootstrapping process creates a small storage cluster on a single node, consisting of one Ceph Monitor and one Ceph Manager, as well as any required dependencies. You then use the orchestrator CLI or the IBM Storage Ceph Dashboard to expand the storage cluster to include nodes, and to provision all of the IBM Storage Ceph daemons and services. You can perform management functions through the CLI or from the IBM Storage Ceph Dashboard web interface.



252_Ceph_0522

Figure 1. Ceph storage cluster deployment

cephadm-ansible playbooks

The `cephadm-ansible` package is a collection of Ansible playbooks to simplify workflows that are not covered by `cephadm`. After installation, the playbooks are located in `/usr/share/cephadm-ansible/`.

The `cephadm-ansible` package includes the following playbooks:

- `cephadm-preflight.yml`
- `cephadm-clients.yml`
- `cephadm-purge-cluster.yml`

The `cephadm-preflight` playbook

Use the `cephadm-preflight` playbook to initially setup hosts before bootstrapping the storage cluster and before adding new nodes or clients to your storage cluster. This playbook configures the Ceph repository and installs some prerequisites such as `podman`, `lvm2`, `chronyd`, and `cephadm`.

For more information, see [Running the preflight playbook](#).

The `cephadm-clients` playbook

Use the `cephadm-clients` playbook to set up client hosts. This playbook handles the distribution of configuration and keyring files to a group of Ceph clients.

For more information, see [“Deploying client nodes” on page 60](#).

The `cephadm-purge-cluster` playbook

Use the `cephadm-purge-cluster` playbook to remove a Ceph cluster. This playbook purges a Ceph cluster that is managed with `cephadm`.

For more information, see [Purging the Ceph storage cluster](#).

Configuring NTP using chrony

Accurate time keeping is important for a number of reasons. In networking for example, accurate time stamps in packets and logs are required. Chrony is an implementation of the Network Time Protocol (NTP).

Before you begin

1. You must have a running Red Hat Ceph Storage cluster.
2. You must have root-level access to the Ceph Monitor node.

About this task

This section explains how to configure NTP using chrony.

Procedure

1. Install chrony by running the following command as root:

```
# yum install chrony*
```

2. Edit the configuration file as per your requirements.

The default configuration file for chronyd is `/etc/chrony.conf`. The `-f` option can be used to specify an alternate configuration file path. It must be altered accordingly to include new NTP servers list and other directives as needed.

```
# vim /etc/chrony.conf
```

3. Restart chronyd service to load the changes:

```
# systemctl restart chronyd.service
```

4. Check if time has been synchronized:

```
# timedatectl
```

5. Display system time information:

```
# chronyc tracking
```

Example

The following `chrony.conf` is an example configuration file with different NTP servers specified with the `pool` directive instead of `multiplierserver` directive. The configuration file also shows some other directives that can be used.

```
pool pool.ntp.org iburst
driftfile /var/lib/chrony/drift
makestep 10 3
rtcsync
```

Registering the IBM Storage Ceph nodes

Register the IBM Storage Ceph nodes on Red Hat Enterprise Linux.

Before you begin

Before registering the IBM Storage Ceph nodes, be sure that you have:

- At least one running virtual machine (VM) or bare-metal server with an active internet connection.

- A valid IBM subscription with the appropriate entitlements.
- Root-level access to all nodes.

For the latest supported Red Hat Enterprise Linux versions, see [Compatibility matrix](#).

Procedure

1. Register the system, and when prompted, enter your Red Hat customer portal credentials.

Example

```
[root@admin ~]# subscription-manager register
```

2. Pull the latest subscription.

Syntax

```
subscription-manager refresh
```

3. Identify the appropriate subscription and retrieve its pool ID.
4. Attach the pool ID to gain access to the software entitlements.

Use the pool ID you identified in the previous step.

Example

```
[root@admin ~]# subscription-manager attach --pool=POOL_ID
```

5. Disable the software repositories.

Example

```
[root@admin ~]# subscription-manager repos --disable=*
```

6. Enable the Red Hat Enterprise Linux baseos and appstream repositories.

Example

```
[root@admin ~]# subscription-manager repos --enable=rhel-9-for-x86_64-baseos-rpms
[root@admin ~]# subscription-manager repos --enable=rhel-9-for-x86_64-appstream-rpms
```

7. Update the system.

Example

```
[root@admin ~]# dnf update
```

8. Enable the ceph-tools repository.

Example

```
[root@admin ~]# curl https://public.dhe.ibm.com/ibmdl/export/pub/storage/ceph/beta/ibm-storage-ceph-7-beta-rhel-9.repo | sudo tee /etc/yum.repos.d/ibm-storage-ceph-7-rhel-9.repo
```

9. Repeat steps “1” on page 13 through “8” on page 13 on all the nodes of the storage cluster.
10. Add license to install IBM Storage Ceph and click **Accept**.

Example

```
[root@admin ~]# dnf install ibm-storage-ceph-license
```

11. Accept the provisions.

Example

```
[root@admin ~]# sudo touch /usr/share/ibm-storage-ceph-license/accept
```

12. Install cephadm-ansible.

```
dnf install cephadm-ansible
```

Configuring Ansible inventory location

You can configure inventory location files for the `cephadm-ansible` staging and production environments. The Ansible inventory hosts file contains all the hosts that are part of the storage cluster.

You can list nodes individually in the inventory hosts file or you can create groups such as `[mons]`, `[osds]`, and `[rgws]` to provide clarity to your inventory and ease the usage of the `--limit` option to target a group or node when running a playbook.

Note: If deploying clients, client nodes must be defined in a dedicated `[clients]` group.

Prerequisites

- An Ansible administration node.
- Root-level access to the Ansible administration node.
- The `cephadm-ansible` package is installed on the node.

Procedure

1. Navigate to the `/usr/share/cephadm-ansible/` directory:

```
[root@admin ~]# cd /usr/share/cephadm-ansible
```

2. Optional: Create subdirectories for staging and production:

```
[root@admin cephadm-ansible]# mkdir -p inventory/staging inventory/production
```

3. Optional: Edit the `ansible.cfg` file and add the following line to assign a default inventory location:

```
[defaults]
inventory = ./inventory/staging
```

4. Optional: Create an inventory hosts file for each environment:

```
[root@admin cephadm-ansible]# touch inventory/staging/hosts
[root@admin cephadm-ansible]# touch inventory/production/hosts
```

5. Open and edit each hosts file and add the nodes and `[admin]` group:

Syntax

```
NODE_NAME_1
NODE_NAME_2

[admin]
ADMIN_NODE_NAME_1
```

- Replace `NODE_NAME_1` and `NODE_NAME_2` with the Ceph nodes such as monitors, OSDs, MDSs, and gateway nodes.
- Replace `ADMIN_NODE_NAME_1` with the name of the node where the admin keyring is stored.

Example

```
host02
host03
host04

[admin]
host01
```

Note: If you set the inventory location in the `ansible.cfg` file to staging, you need to run the playbooks in the staging environment as follows:

Syntax

```
ansible-playbook -i inventory/staging/hosts PLAYBOOK.yml
```

To run the playbooks in the production environment:

Syntax

```
ansible-playbook -i inventory/production/hosts PLAYBOOK.yml
```

Creating an Ansible user with sudo access

You can create an Ansible user with password-less `root` access on all nodes in the storage cluster to run the `cephadm-ansible` playbooks.

About this task

The Ansible user must be able to log into all the IBM Storage Ceph nodes as a user that has `root` privileges to install software and create configuration files without prompting for a password.

Important: If you are using Red Hat Enterprise Linux 9 only use these steps if you are a non-root user. If you are a root user, go to [“Bootstrapping a new storage cluster”](#) on page 20.

For more information about creating user accounts, see **Configuring basic system settings** > **Getting started with managing user accounts** within the *Red Hat Enterprise Linux* guide.

Note:

Before you begin

- Root-level access to all nodes.
- For Red Hat Enterprise 9, to log in as a root user, see [“Enabling SSH login as root user on Red Hat Enterprise Linux 9”](#) on page 18

Procedure

Complete these steps on each node in the storage cluster.

1. Log in to the node as the `root` user.

```
ssh root@HOST_NAME
```

Replace `HOST_NAME` with the host name of the Ceph node.

For example,

```
[root@admin ~]# ssh root@host01
```

Enter the `root` password when prompted.

2. Create a new Ansible user.

```
adduser USER_NAME
```

Replace `USER_NAME` with the new user name for the Ansible user.

Important: Do not use `ceph` as the user name. The `ceph` user name is reserved for the Ceph daemons. A uniform user name across the cluster can improve ease of use, but avoid using obvious user names, because intruders typically use them for brute-force attacks.

For example,

```
[root@host01 ~]# adduser ceph-admin
```

3. Set a new password for this user.

```
passwd USER_NAME
```

Replace *USER_NAME* with the new user name for the Ansible user.

For example,

```
[root@host01 ~]# passwd ceph-admin
```

Enter the new password twice when prompted.

4. Configure sudo access for the newly created user.

```
cat << EOF >/etc/sudoers.d/USER_NAME
$USER_NAME ALL = (root) NOPASSWD:ALL
EOF
```

Replace *USER_NAME* with the user name for the Ansible user, created in step “3” on page 16.

For example,

```
[root@host01 ~]# cat << EOF >/etc/sudoers.d/ceph-admin
ceph-admin ALL = (root) NOPASSWD:ALL
EOF
```

5. Assign the correct file permissions to the new file.

```
chmod 0440 /etc/sudoers.d/USER_NAME
```

Replace *USER_NAME* with the user name for the Ansible user, created in step “3” on page 16.

For example,

```
[root@host01 ~]# chmod 0440 /etc/sudoers.d/ceph-admin
```

Configuring SSH

As a storage administrator, with Cephadm, you can use an SSH key to securely authenticate with remote hosts. The SSH key is stored in the monitor to connect to remote hosts.

Before you begin

- A running IBM Storage Ceph cluster.
- An Ansible administration node.
- Root-level access to the Ansible administration node.
- The cephadm-ansible package is installed on the node.

Procedure

1. Navigate to the cephadm-ansible directory.
2. Generate a new SSH key.

```
ceph cephadm generate-key
```

For example,

```
[ceph-admin@admin cephadm-ansible]$ ceph cephadm generate-key
```

3. Retrieve the public portion of the SSH key.

```
ceph cephadm get-pub-key
```

For example,

```
[ceph-admin@admin cephadm-ansible]$ ceph cephadm get-pub-key
```

4. Delete the currently stored SSH key.

```
ceph cephadm clear-key
```

For example,

```
[ceph-admin@admin cephadm-ansible]$ ceph cephadm clear-key
```

5. Restart the mgr daemon to reload the configuration.

```
ceph mgr fail
```

For example,

```
[ceph-admin@admin cephadm-ansible]$ ceph mgr fail
```

Configuring a different SSH user

As a storage administrator, you can configure a non-root SSH user who can log in to all the Ceph cluster nodes with enough privileges to download container images, start containers, and run commands without prompting for a password.

Important: Before configuring a non-root SSH user, the cluster SSH key needs to be added to the user's `authorized_keys` file and non-root users must have *passwordless* sudo access.

Prerequisites

- A running IBM Storage Ceph cluster.
- An Ansible administration node.
- Root-level access to the Ansible administration node.
- The `cephadm-ansible` package is installed on the node.
- Add the cluster SSH keys to the user's `authorized_keys`.
- Enable *passwordless* sudo access for the non-root users.

Procedure

1. Go to the `cephadm-ansible` directory.
2. Provide Cephadm the name of the user who is going to perform all the Cephadm operations.

Syntax

```
ceph cephadm set-user USER
```

Example

```
[ceph-admin@admin cephadm-ansible]$ ceph cephadm set-user user
```

3. Retrieve the SSH public key.

Syntax

```
ceph cephadm get-pub-key > ~/ceph.pub
```

Example

```
[ceph-admin@admin cephadm-ansible]$ ceph cephadm get-pub-key > ~/ceph.pub
```

4. Copy the SSH keys to all the hosts.

Syntax

```
ssh-copy-id -f -i ~/ceph.pub USER@HOST
```

Example

```
[ceph-admin@admin cephadm-ansible]$ ssh-copy-id ceph-admin@host01
```

Enabling SSH login as root user on Red Hat Enterprise Linux 9

Enabling SSH login as root user on Red Hat Enterprise Linux 9, as it is not automatically supported.

Red Hat Enterprise Linux 9 does not support SSH login as a root user even if `PermitRootLogin` parameter is set to `yes` in the `/etc/ssh/sshd_config` file. You get the following error:

Example

```
[root@host01 ~]# ssh root@myhostname
root@myhostname password:
Permission denied, please try again.
```

You can run one of the following methods to enable login as a root user:

- Use the `Allow root SSH login with password` flag while setting the root password during installation of Red Hat Enterprise Linux 9.
- Manually set the `PermitRootLogin` parameter after Red Hat Enterprise Linux 9 installation.

This section describes manual setting of the `PermitRootLogin` parameter.

Prerequisites

- Root-level access to all nodes.

Procedure

1. Open the `etc/ssh/sshd_config` file and set the `PermitRootLogin` to `yes`:

Example

```
[root@admin ~]# echo 'PermitRootLogin yes' >> /etc/ssh/sshd_config.d/01-permitrootlogin.conf
```

2. Restart the SSH service:

Example

```
[root@admin ~]# systemctl restart sshd.service
```

3. Login to the node as the root user:

Syntax

```
ssh root@HOST_NAME
```

Replace `HOST_NAME` with the host name of the Ceph node.

Example

```
[root@admin ~]# ssh root@host01
```

Enter the root password when prompted.

Reference

- For more information, see [Not able to login as root user via ssh in RHEL 9 server](#)

Running the preflight playbook

This Ansible playbook configures the Ceph repository and prepares the storage cluster for bootstrapping. It also installs some prerequisites, such as podman, lvm2, chronyd, and cephadm. The default location for cephadm-ansible and cephadm-preflight.yml is /usr/share/cephadm-ansible.

About this task

The preflight playbook uses the cephadm-ansible inventory file to identify all the admin and nodes in the storage cluster.

The default location for the inventory file is /usr/share/cephadm-ansible/hosts.

The following example shows the structure of a typical inventory file:

```
host02
host03
host04

[admin]
host01
```

The [admin] group in the inventory file contains the name of the node where the admin keyring is stored. On a new storage cluster, the node in the [admin] group will be the bootstrap node. To add additional admin hosts after bootstrapping the cluster see [“Setting up the admin node” on page 56](#).

Important: If you are performing a disconnected installation, see [“Running the preflight playbook for a disconnected installation” on page 34](#).

Note: Run the preflight playbook before you bootstrap the initial host.

Before you begin

- Root-level access to the Ansible administration node.
- Ansible user with sudo and passwordless ssh access to all nodes in the storage cluster.

Procedure

Note: In the following procedure, host01 is the bootstrap node.

1. Navigate to the /usr/share/cephadm-ansible directory.
2. Open and edit the hosts file and add your nodes.
For example,

```
host02
host03
host04

[admin]
host01
```

3. Add a license to install IBM Storage Ceph and click **Accept** on all nodes.

```
dnf install ibm-storage-ceph-license
```

For example,

```
[root@admin ~]# dnf install ibm-storage-ceph-license
```

4. Accept all provisions.

```
sudo touch /usr/share/ibm-storage-ceph-license/accept
```

For example,

```
[root@admin ~]# sudo touch /usr/share/ibm-storage-ceph-license/accept
```

5. Run the preflight playbook.

Note: When running the preflight playbook, `cephadm-ansible` automatically installs `chronyd` and `ceph-common` on the client nodes.

Important: Since this is a BETA release, you need to use the `ceph_origin` parameter set to custom to use the repository.

```
ansible-playbook -i INVENTORY_FILE cephadm-preflight.yml --extra-vars "ceph_origin=custom" -e "custom_repo_url=BETA_URL"
```

For example,

```
[ansible@admin cephadm-ansible]$ ansible-playbook -i hosts cephadm-preflight.yml --extra-vars "ceph_origin=custom" -e "custom_repo_url=https://public.dhe.ibm.com/ibmdl/export/pub/storage/ceph/beta/rhel9/x86_64/"
```

After installation is complete, `cephadm` resides in the `/usr/sbin/` directory.

- To run the preflight playbook on a selected set of hosts in a storage cluster, use the `--limit` option.

```
ansible-playbook -i INVENTORY_FILE cephadm-preflight.yml --extra-vars "ceph_origin=custom" -e "custom_repo_url=BETA_URL" --limit GROUP_NAME|NODE_NAME
```

Replace `GROUP_NAME` with a group name from your inventory file and `NODE_NAME` with a specific node name from your inventory file.

For example,

```
[ansible@admin cephadm-ansible]$ ansible-playbook -i hosts cephadm-preflight.yml --extra-vars "ceph_origin=custom" -e "custom_repo_url=https://public.dhe.ibm.com/ibmdl/export/pub/storage/ceph/beta/rhel9/x86_64/" --limit clients
[ansible@admin cephadm-ansible]$ ansible-playbook -i hosts cephadm-preflight.yml --extra-vars "ceph_origin=custom" -e "custom_repo_url=https://public.dhe.ibm.com/ibmdl/export/pub/storage/ceph/beta/rhel9/x86_64/" --limit host01
```

Bootstrapping a new storage cluster

Use the `cephadm` utility to bootstrap a new storage cluster.

About this task

The `cephadm` utility performs the following tasks during the bootstrap process:

- Installs and starts a Ceph Monitor daemon and a Ceph Manager daemon for a new IBM Storage Ceph cluster on the local node as containers.
- Creates the `/etc/ceph` directory.
- Writes a copy of the public key to `/etc/ceph/ceph.pub` for the IBM Storage Ceph cluster and adds the SSH key to the root user's `/root/.ssh/authorized_keys` file.
- Applies the `_admin` label to the bootstrap node.
- Writes a minimal configuration file needed to communicate with the new cluster to `/etc/ceph/ceph.conf`.
- Writes a copy of the `client.admin` administrative secret key to `/etc/ceph/ceph.client.admin.keyring`.

- Deploys a basic monitoring stack with prometheus, grafana, and other tools such as node-exporter and alert-manager.

Important:

- If you are performing a disconnected installation, see [“Performing a disconnected installation”](#) on page 36.
- If you are deploying a monitoring stack, see [Deploying the monitoring stack using the Ceph Orchestrator](#).
- Bootstrapping provides the default user name and password for the initial login to the Dashboard. Bootstrap requires you to change the password after you log in.
- Before you begin the bootstrapping process, make sure that the container image that you want to use has the same version of IBM Storage Ceph as cephadm. If the two versions do not match, bootstrapping fails at the Creating initial admin user stage.

Note:

- If you have an existing Prometheus services that you want to run with the new storage cluster, or if you are running Ceph with Rook, use the `--skip-monitoring-stack` option with the `cephadm bootstrap` command. This option bypasses the basic monitoring stack so that you can manually configure it later.
- Before you begin the bootstrapping process, you must create a username and password for the `cp.icr.io/cp` container registry.

Note: Enable Call Home and Storage Insights during installation. For bootstrap command options for Call Home and IBM Storage Insights, see [bootstrap command options](#).

For more information, see:

- [Recommended cephadm bootstrap command options](#)
- [Bootstrap command options](#)
- [Using a JSON file to protect login information](#)

Before you begin

- An IP address for the first Ceph Monitor container, which is also the IP address for the first node in the storage cluster.
- Login access to `cp.icr.io/cp`. For information about obtaining credentials for `cp.icr.io/cp`, see [“Obtaining entitlement key”](#) on page 29.
- A minimum of 10 GB of free space for `/var/lib/containers/`.
- Root-level access to all nodes.

Procedure

Important:

- Run **cephadm bootstrap** on the node that you want to be the initial Monitor node in the cluster. The `IP_ADDRESS` option should be the IP address of the node you are using to run **cephadm bootstrap**.

Note:

- If the storage cluster includes multiple networks and interfaces, be sure to choose a network that is accessible by any node that uses the storage cluster.
- If the local node uses fully-qualified domain names (FQDN), then add the `--allow-fqdn-hostname` option to the **cephadm bootstrap** command.
- To deploy a storage cluster using IPV6 addresses, use the IPV6 address format for the `--mon-ip IP_ADDRESS` option. For example: **cephadm bootstrap --mon-ip 2620:52:0:880:225:90ff:fefc:2536 --registry-json /etc/mylogin.json**.

- To route the internal cluster traffic over the public network, omit the `--cluster-network NETWORK_CIDR` option.
- Bootstrap a storage cluster.

```
cephadm bootstrap --cluster-network NETWORK_CIDR --mon-ip IP_ADDRESS --registry-url
cp.icr.io/cp --registry-username USER_NAME --registry-password PASSWORD --yes-i-know
```

For example,

```
[root@host01 ~]# cephadm bootstrap --cluster-network 10.10.128.0/24 --mon-ip 10.10.128.68 --
registry-url cp.icr.io/cp --registry-username myuser1 --registry-password mypassword1 --yes-
i-know
```

The script takes a few minutes to complete. Once the script completes, it provides the credentials to the IBM Storage Ceph Dashboard URL, a command to access the Ceph command-line interface (CLI), and a request to enable telemetry.

Ceph Dashboard is now available at:

```
URL: https://host01:8443/
User: admin
Password: i8nhu7zham
```

Enabling client.admin keyring and conf on hosts with "admin" label
You can access the Ceph CLI with:

```
sudo /usr/sbin/cephadm shell --fsid 266ee7a8-2a05-11eb-b846-5254002d4916 -c /etc/ceph/
ceph.conf -k /etc/ceph/ceph.client.admin.keyring
```

Please consider enabling telemetry to help improve Ceph:

```
ceph telemetry on
```

For more information see:

```
https://docs.ceph.com/docs/master/mgr/telemetry/
```

Bootstrap complete.

Recommended cephadm bootstrap command options

The **cephadm bootstrap** command has multiple options that allow you to specify file locations, configure SSH settings, set passwords, and perform other initial configuration tasks.

Use a basic set of command options for `cephadm bootstrap`. You can configure additional options after your initial cluster is up and running.

The following examples show how to specify the recommended options.

Syntax

```
cephadm bootstrap --ssh-user USER_NAME --mon-ip IP_ADDRESS --allow-fqdn-hostname --registry-
json REGISTRY_JSON
```

Example

```
[root@host01 ~]# cephadm bootstrap --ssh-user ceph --mon-ip 10.10.128.68 --allow-fqdn-hostname
--registry-json /etc/mylogin.json
```

For non-root users, see [Creating an Ansible user with sudo access](#) and [Enabling password-less SSH for Ansible](#) for more details.

Reference

- For more information about the `--registry-json` option, see [“Using a JSON file to protect login information”](#) on page 23.

- For more information about all available `cephadm bootstrap` options, see [“Bootstrap command options”](#) on page 26.
- For more information about bootstrapping the storage cluster as a non-root user, see [“Bootstrapping the storage cluster as a non-root user”](#) on page 25.

Using a JSON file to protect login information

As a storage administrator, you might choose to add login and password information to a JSON file, and then refer to the JSON file for bootstrapping. This protects the login credentials from exposure.

Note: A JSON file can also be used with the `cephadm --registry-login` command.

Prerequisites

- An IP address for the first Ceph Monitor container, which is also the IP address for the first node in the storage cluster.
- Login access to `cp.icr.io/cp`.
- A minimum of 10 GB of free space for `/var/lib/containers/`.
- Root-level access to all nodes.

Procedure

1. Create the JSON file. In this example, the file is named `mylogin.json`.

Syntax

```
{
  "url": "REGISTRY_URL",
  "username": "USER_NAME",
  "password": "PASSWORD"
}
```

Example

```
{
  "url": "cp.icr.io/cp",
  "username": "myuser1",
  "password": "mypassword1"
}
```

2. Bootstrap a storage cluster:

Syntax

```
cephadm bootstrap --mon-ip IP_ADDRESS --registry-json /etc/mylogin.json
```

Example

```
[root@host01 ~]# cephadm bootstrap --mon-ip 10.10.128.68 --registry-json /etc/mylogin.json
```

Bootstrapping a storage cluster using a service configuration file

Use the `cephadm bootstrap` command with the `--apply-spec` option to bootstrap the storage cluster and configure additional hosts and daemons using a service configuration file. The configuration file is a `.yaml` file that contains the service type, placement, and designated nodes for services that you want to deploy.

Note: If you want to use a non-default realm or zone for applications such as multi-site, configure your Ceph Object Gateway daemons after you bootstrap the storage cluster, instead of adding them to the configuration file and using the `--apply-spec` option. This gives you the opportunity to create the realm or zone you need for the Ceph Object Gateway daemons before deploying them.

Note: If deploying a NFS-Ganesha gateway, or Metadata Server (MDS) service, configure them after bootstrapping the storage cluster.

- To deploy a Ceph NFS-Ganesha gateway, you must create a RADOS pool first.
- To deploy the MDS service, you must create a CephFS volume first.

Note: If you run the bootstrap command with `--apply-spec` option, ensure to include the IP address of the bootstrap host in the specification file. This prevents resolving the IP address to loopback address while re-adding the bootstrap host where active Ceph Manager is already running. If you do not use the `--apply-spec` option during bootstrap and instead use `ceph orch apply` command with another specification file which includes re-adding the host and contains an active Ceph Manager running, then ensure to explicitly provide the `addr` field. This is applicable for applying any specification file after bootstrapping.

For more information, see [Operations](#).

Prerequisites

- At least one running virtual machine (VM) or server.
- Root-level access to all nodes.
- Login access to `cp.icr.io/cp`.
- Passwordless ssh is set up on all hosts in the storage cluster.
- `cephadm` is installed on the node that you want to be the initial Monitor node in the storage cluster.

For the latest supported Red Hat Enterprise Linux versions, see [Compatibility matrix](#).

Procedure

1. Log in to the bootstrap host.
2. Create the service configuration `.yaml` file for your storage cluster. The example file directs `cephadm bootstrap` to configure the initial host and two additional hosts, and it specifies that OSDs be created on all available disks.

Example

```
service_type: host
addr: host01
hostname: host01
---
service_type: host
addr: host02
hostname: host02
---
service_type: host
addr: host03
hostname: host03
---
service_type: host
addr: host04
hostname: host04
---
service_type: mon
placement:
  host_pattern: "host[0-2]"
---
service_type: osd
service_id: my_osds
placement:
  host_pattern: "host[1-3]"
data_devices:
  all: true
```

3. Bootstrap the storage cluster with the `--apply-spec` option:

Syntax

```
cephadm bootstrap --apply-spec CONFIGURATION_FILE_NAME --mon-ip MONITOR_IP_ADDRESS --registry-url cp.icr.io/cp --registry-username USER_NAME --registry-password PASSWORD
```

Example

```
[root@host01 ~]# cephadm bootstrap --apply-spec initial-config.yaml --mon-ip 10.10.128.68 --registry-url cp.icr.io/cp --registry-username myuser1 --registry-password mypassword1
```

The script takes a few minutes to complete. Once the script completes, it provides the credentials to the IBM Storage Ceph Dashboard URL, a command to access the Ceph command-line interface (CLI), and a request to enable telemetry.

Once your storage cluster is up and running, see [Operations](#) for more information about configuring additional daemons and services.

Reference

- [“Bootstrap command options” on page 26](#)

Bootstrapping the storage cluster as a non-root user

Use the **cephadm bootstrap** command with the **--ssh-user** option to bootstrap the IBM Storage Ceph cluster as a non-root user on the bootstrap node. The **--ssh-user** specifies a user for SSH connections to cluster nodes.

Non-root users must have passwordless sudo access. For more information, see [Creating an Ansible user with sudo access](#) and [Enabling password-less SSH for Ansible](#) for more details.

Prerequisites

- An IP address for the first Ceph Monitor container, which is also the IP address for the initial Monitor node in the storage cluster.
- Login access to cp.icr.io/cp.
- A minimum of 10 GB of free space for /var/lib/containers/.
- SSH public and private keys.
- Passwordless sudo access to the bootstrap node.

Procedure

1. Change to sudo on the bootstrap node:

Syntax

```
su - SSH_USER_NAME
```

Example

```
[root@host01 ~]# su - ceph  
Last login: Tue Sep 14 12:00:29 EST 2021 on pts/0
```

2. Establish the SSH connection to the bootstrap node:

Example

```
[ceph@host01 ~]# ssh host01  
Last login: Tue Sep 14 12:03:29 EST 2021 on pts/0
```

3. Optional: Invoke the cephadm bootstrap command.

Note: Using private and public keys is optional.

If SSH keys have not previously been created, these can be created during this step.

Include the `--ssh-private-key` and `--ssh-public-key` options:

Syntax

```
cephadm bootstrap --ssh-user USER_NAME --mon-ip IP_ADDRESS --ssh-private-key PRIVATE_KEY  
--ssh-public-key PUBLIC_KEY --registry-url cp.icr.io/cp --registry-username USER_NAME --  
registry-password PASSWORD
```

Example

```
cephadm bootstrap --ssh-user ceph --mon-ip 10.10.128.68 --ssh-private-key /home/ceph/.ssh/  
id_rsa --ssh-public-key /home/ceph/.ssh/id_rsa.pub --registry-url cp.icr.io/cp --registry-  
username myuser1 --registry-password mypassword1
```

Reference

- [“Bootstrap command options” on page 26](#)
- For more information about utilizing Ansible to automate bootstrapping a rootless cluster, see the knowledge base article [Red Hat Ceph Storage 5.3 rootless deployment utilizing ansible ad-hoc commands](#).

Bootstrap command options

The **cephadm bootstrap** command bootstraps a IBM Storage Ceph cluster on the local host. It deploys a MON daemon and a MGR daemon on the bootstrap node, automatically deploys the monitoring stack on the local host, and calls **ceph orch host add HOSTNAME**.

The following table lists the available options for **cephadm bootstrap**.

cephadm bootstrap option	Description
<code>--config CONFIG_FILE, -c CONFIG_FILE</code>	<i>CONFIG_FILE</i> is the <code>ceph.conf</code> file to use with the bootstrap command
<code>--cluster-network NETWORK_CIDR</code>	Use the subnet defined by <i>NETWORK_CIDR</i> for internal cluster traffic. This is specified in CIDR notation. For example: <code>10.10.128.0/24</code> .
<code>--mon-id MON_ID</code>	Bootstraps on the host named <i>MON_ID</i> . Default value is the local host.
<code>--mon-addrv MON_ADDRV</code>	mon IPs (e.g., <code>[v2:localipaddr:3300,v1:localipaddr:6789]</code>)
<code>--mon-ip IP_ADDRESS</code>	IP address of the node you are using to run <code>cephadm bootstrap</code> .
<code>--mgr-id MGR_ID</code>	Host ID where a MGR node should be installed. Default: randomly generated.
<code>--fsid FSID</code>	Cluster FSID.
<code>--output-dir OUTPUT_DIR</code>	Use this directory to write config, keyring, and pub key files.
<code>--output-keyring OUTPUT_KEYRING</code>	Use this location to write the keyring file with the new cluster admin and mon keys.

cephadm bootstrap option	Description
--output-config <i>OUTPUT_CONFIG</i>	Use this location to write the configuration file to connect to the new cluster.
--output-pub-ssh-key <i>OUTPUT_PUB_SSH_KEY</i>	Use this location to write the public SSH key for the cluster.
--skip-ssh	Skip the setup of the ssh key on the local host.
--initial-dashboard-user <i>INITIAL_DASHBOARD_USER</i>	Initial user for the dashboard.
--initial-dashboard-password <i>INITIAL_DASHBOARD_PASSWORD</i>	Initial password for the initial dashboard user.
--ssl-dashboard-port <i>SSL_DASHBOARD_PORT</i>	Port number used to connect with the dashboard using SSL.
--dashboard-key <i>DASHBOARD_KEY</i>	Dashboard key.
--dashboard-crt <i>DASHBOARD_CRT</i>	Dashboard certificate.
--ssh-config <i>SSH_CONFIG</i>	SSH config.
--ssh-private-key <i>SSH_PRIVATE_KEY</i>	SSH private key.
--ssh-public-key <i>SSH_PUBLIC_KEY</i>	SSH public key.
--ssh-user <i>SSH_USER</i>	Sets the user for SSH connections to cluster hosts. Passwordless sudo is needed for non-root users.
--skip-mon-network	Sets mon public_network based on the bootstrap mon ip.
--skip-dashboard	Do not enable the Ceph Dashboard.
--dashboard-password-noupdate	Disable forced dashboard password change.
--no-minimize-config	Do not assimilate and minimize the configuration file.
--skip-ping-check	Do not verify that the mon IP is pingable.
--skip-pull	Do not pull the latest image before bootstrapping.
--skip-firewalld	Do not configure firewalld.
--allow-overwrite	Allow the overwrite of existing --output-* config/keyring/ssh files.
--allow-fqdn-hostname	Allow fully qualified host name.
--skip-prepare-host	Do not prepare host.

cephadm bootstrap option	Description
--orphan-initial-daemons	Do not create initial mon, mgr, and crash service specs.
--skip-monitoring-stack	Do not automatically provision the monitoring stack] (prometheus, grafana, alertmanager, node-exporter).
--apply-spec <i>APPLY_SPEC</i>	Apply cluster spec file after bootstrap (copy ssh key, add hosts and apply services).
--registry-url <i>REGISTRY_URL</i>	Specifies the URL of the custom registry to log into. For example: cp.icr.io/cp.
--registry-username <i>REGISTRY_USERNAME</i>	User name of the login account to the custom registry.
--registry-password <i>REGISTRY_PASSWORD</i>	Password of the login account to the custom registry.
--registry-json <i>REGISTRY_JSON</i>	JSON file containing registry login information.

The following table lists the available options that are required to enable Call home via the bootstrap command -

cephadm bootstrap option	Description
--enable-ibm-call-home	Flag to opt-in to IBM Call Home. Syntax: <pre>cephadm -v --image cp.stg.icr.io/cp/ibm-ceph/ceph-7-rhel9:7-2 bootstrap --mon-ip 10.0.210.25 --registry-json registry.json --enable-ibm-call-home --call-home-config call_home.json --enable-storage-insights --storage-insights-config si_opt.json</pre>
--call-home-icn	IBM Customer Number (7 alphanumeric characters long).
--ceph-call-home-contact-first-name	Storage help desk contact first name.
--ceph-call-home-contact-last-name	Storage help desk contact last name.
--ceph-call-home-contact-email	Storage help desk email. For example, <i>storage_helpdesk@ibm.com</i> .
--ceph-call-home-contact-phone	Storage help desk phone number.
--ceph-call-home-country-code	2-letter country code.

cephadm bootstrap option	Description
--call-home-config	<p>File path to a JSON config file which includes the following keys:</p> <pre>{ "icn": "<IBM_CUSTOMER_NUMBER>", "email": "<CALL_HOME_CONTACT_EMAIL_ADDRESS>", "phone": "<CALL_HOME_CONTACT_PHONE_NUMBER>", "first_name": "<CALL_HOME_CONTACT_FIRST_NAME>", "last_name": "<CALL_HOME_CONTACT_LAST_NAME>", "country_code": "<CUSTOMER_COUNTRY_CODE>" }</pre> <p>Use either the direct parameters with their options, or the --call-home-config flag with the file path.</p>

Enable IBM Storage Insights too by providing the following flags and parameters:

cephadm bootstrap option	Description
--enable-storage-insights	<p>Flag to opt-in to IBM Storage Insights.</p> <p>Syntax:</p> <pre>cephadm -v --image cp.stg.icr.io/cp/ibm-ceph/ceph-7-rhel9:7-2 bootstrap --mon-ip 10.0.210.25 --registry-json registry.json --enable-ibm-call-home --call-home-config call_home.json --enable-storage-insights --storage-insights-config si_opt.json</pre>
--storage-insights-tenant-id	A Storage Insights tenant ID to be associated with this deployment.
--storage-insights-config	<p>File path to a JSON config file which includes the following key:</p> <pre>{"tenant_id": "<STORAGE_INSIGHTS_TENANT_ID>"}</pre> <p>Use either the direct parameters with their options, or the --storage-insights-config flag and specify its file path.</p>

Reference

- For more information about the --skip-monitoring-stack option, see [“Adding hosts” on page 45](#).
- For more information about logging into the registry with the registry-json option, see help for the registry-login command.
- For more information about cephadm options, see help for cephadm.

Obtaining entitlement key

Entitlement keys determine whether IBM Storage Ceph can automatically pull the required container default images. During installation, image pull failures can occur due to an invalid entitlement key or a key belonging to an account that does not have entitlement to IBM Storage Ceph.

Procedure

1. Log in to the IBM container software library with the IBMid and password that is associated with the entitled IBM Storage Ceph software.

2. In the navigation bar, click **Get entitlement key**.
3. On the **Access your container software** page, click **Copy key** to copy the generated entitlement key.
4. Save the key to a secure location for future use.

Disconnected installation

As a storage administrator, you can use a disconnected installation procedure to install cephadm and bootstrap your storage cluster on a private network. A disconnected installation uses a private registry for installation.

Configuring a private registry for a disconnected installation

Use this procedure when the IBM Storage Ceph nodes do NOT have access to the Internet during deployment.

Before you begin

- At least one running virtual machine (VM) or server with an active internet connection.
- Login access to `cp.icr.io/cp`.
- Root-level access to all nodes.

For the latest supported Red Hat Enterprise Linux versions, see [Compatibility matrix](#).

About this task

Follow this procedure to set up a secure private registry using authentication and a self-signed certificate. Perform these steps on a node that has both Internet access and access to the local cluster.

Note: Do not use an insecure registry for production.

Procedure

1. Register the hosts, and when prompted, enter the appropriate Red Hat customer portal credentials:

Example

```
[root@admin ~]# subscription-manager register
```

2. Pull the latest subscription:

Example

```
subscription-manager refresh
```

3. Identify the appropriate subscription and retrieve its Pool ID.
4. Attach the pool ID to gain access to the software entitlements. Use the Pool ID you identified in the previous step.

Example

```
[root@admin ~]# subscription-manager attach --pool=POOL_ID
```

5. Disable the software repositories:

Example

```
[root@admin ~]# subscription-manager repos --disable=*
```

6. Enable the Red Hat Enterprise Linux baseos and appstream repositories.

Example

```
[root@admin ~]# subscription-manager repos --enable=rhel-9-for-x86_64-baseos-rpms
[root@admin ~]# subscription-manager repos --enable=rhel-9-for-x86_64-appstream-rpms
```

7. Update the system:

Example

```
[root@admin ~]# dnf update
```

8. Enable the ceph-tools repository.

Example

```
[root@admin ~]# curl https://public.dhe.ibm.com/ibmdl/export/pub/storage/ceph/beta/ibm-storage-ceph-7-beta-rhel-9.repo | sudo tee /etc/yum.repos.d/ibm-storage-ceph-7-rhel-9.repo
```

9. Install the podman and httpd-tools packages.

```
[root@admin ~]# dnf install -y podman httpd-tools
```

10. Create folders for the private registry.

```
[root@admin ~]# mkdir -p /opt/registry/{auth,certs,data}
```

The registry is stored in `/opt/registry` and the directories are mounted in the container running the registry.

- The `auth` directory stores the `htpasswd` file the registry uses for authentication.
- The `certs` directory stores the certificates the registry uses for authentication.
- The `data` directory stores the registry images.

11. Create credentials for accessing the private registry.

```
htpasswd -bBc /opt/registry/auth/htpasswd PRIVATE_REGISTRY_USERNAME
PRIVATE_REGISTRY_PASSWORD
```

- The `b` option provides the password from the command line.
- The `B` option stores the password using `Bcrypt` encryption.
- The `c` option creates the `htpasswd` file.
- Replace `PRIVATE_REGISTRY_USERNAME` with the username to create for the private registry.
- Replace `PRIVATE_REGISTRY_PASSWORD` with the password to create for the private registry username.

For example:

```
[root@admin ~]# htpasswd -bBc /opt/registry/auth/htpasswd myregistryusername
myregistrypassword1
```

12. Create a self-signed certificate.

```
openssl req -newkey rsa:4096 -nodes -sha256 -keyout /opt/registry/certs/domain.key -x509
-days 365 -out /opt/registry/certs/domain.crt -addext "subjectAltName = DNS:LOCAL_NODE_FQDN"
```

Replace `LOCAL_NODE_FQDN` with the fully qualified host name of the private registry node.

There is a prompt for the respective options for your certificate. The `CN=` value is the host name of your node and should be resolvable by DNS or the `/etc/hosts` file.

For example:

```
[root@admin ~] # openssl req -newkey rsa:4096 -nodes -sha256 -keyout /opt/registry/certs/
domain.key -x509 -days 365 -out /opt/registry/certs/domain.crt -addext "subjectAltName =
DNS:admin.lab.ibm.com"
```

Note: When creating a self-signed certificate, be sure to create a certificate with a proper Subject Alternative Name (SAN). Podman commands that require TLS verification for certificates that do not include a proper SAN, return the following error:

```
x509: certificate relies on legacy Common Name field, use SANs or temporarily enable Common Name matching with GODEBUG=x509ignoreCN=0
```

13. Create a symbolic link to `domain.cert`.

This allows skopeo to locate the certificate with the file extension `.cert`.

For example:

```
[root@admin ~]# ln -s /opt/registry/certs/domain.crt /opt/registry/certs/domain.cert
```

14. Add the certificate to the trusted list on the private registry node.

```
cp /opt/registry/certs/domain.crt /etc/pki/ca-trust/source/anchors/  
update-ca-trust  
trust list | grep -i "LOCAL_NODE_FQDN"
```

Replace `LOCAL_NODE_FQDN` with the FQDN of the private registry node.

For example:

```
[root@admin ~]# cp /opt/registry/certs/domain.crt /etc/pki/ca-trust/source/anchors/  
[root@admin ~]# update-ca-trust  
[root@admin ~]# trust list | grep -i "admin.lab.ibm.com"  
  
label: admin.lab.ibm.com
```

15. Copy the certificate to any nodes that will access the private registry for installation and update the trusted list.

For example:

```
[root@admin ~]# scp /opt/registry/certs/domain.crt root@host01:/etc/pki/ca-trust/source/  
anchors/  
[root@admin ~]# ssh root@host01  
[root@host01 ~]# update-ca-trust  
[root@host01 ~]# trust list | grep -i "admin.lab.ibm.com"  
  
label: admin.lab.ibm.com
```

16. Start the local secure private registry.

```
[root@admin ~]# podman run --restart=always --name NAME_OF_CONTAINER \  
-p 5000:5000 -v /opt/registry/data:/var/lib/registry:z \  
-v /opt/registry/auth:/auth:z \  
-v /opt/registry/certs:/certs:z \  
-e "REGISTRY_AUTH=htpasswd" \  
-e "REGISTRY_AUTH_HTPASSWD_REALM=Registry Realm" \  
-e REGISTRY_AUTH_HTPASSWD_PATH=/auth/htpasswd \  
-e "REGISTRY_HTTP_TLS_CERTIFICATE=/certs/domain.crt" \  
-e "REGISTRY_HTTP_TLS_KEY=/certs/domain.key" \  
-e REGISTRY_COMPATIBILITY_SCHEMA1_ENABLED=true \  
-d registry:2
```

Replace `NAME_OF_CONTAINER` with a name to assign to the container.

This starts the private registry on port 5000 and mounts the volumes of the registry directories in the container running the registry.

For example:

```
[root@admin ~]# podman run --restart=always --name myprivateregistry \
-p 5000:5000 -v /opt/registry/data:/var/lib/registry:z \
-v /opt/registry/auth:/auth:z \
-v /opt/registry/certs:/certs:z \
-e "REGISTRY_AUTH=htpasswd" \
-e "REGISTRY_AUTH_HTPASSWD_REALM=Registry Realm" \
-e REGISTRY_AUTH_HTPASSWD_PATH=/auth/htpasswd \
-e "REGISTRY_HTTP_TLS_CERTIFICATE=/certs/domain.crt" \
-e "REGISTRY_HTTP_TLS_KEY=/certs/domain.key" \
-e REGISTRY_COMPATIBILITY_SCHEMA1_ENABLED=true \
-d registry:2
```

17. On the local registry node, verify that `cp.icr.io/cp` is in the container registry search path.
 - a) Open `/etc/containers/registries.conf` for editing.
 - b) Optional: If needed, add `cp.icr.io/cp` to the `unqualified-search-registries` list.

```
unqualified-search-registries = ["cp.icr.io/cp"]
```

18. With your IBM Customer Portal credentials, login to `cp.icr.io/cp`.

```
podman login cp.icr.io/cp
```

19. Copy the following images from the IBM Customer Portal to the private registry.

<i>Table 1. Custom image details for monitoring stack</i>	
Monitoring stack component	Image details
Prometheus	cp.icr.io/cp/ibm-ceph/prometheus:v4.12
Grafana	cp.icr.io/cp/ibm-ceph/ceph-7-dashboard-rhel9:latest
Node-exporter	cp.icr.io/cp/ibm-ceph/prometheus-node-exporter:v4.12
AlertManager	cp.icr.io/cp/ibm-ceph/prometheus-alertmanager:v4.12
HAProxy	cp.icr.io/cp/ibm-ceph/haproxy-rhel9:latest
Keepalived	cp.icr.io/cp/ibm-ceph/keepalived-rhel9:latest
SNMP Gateway	cp.icr.io/cp/ibm-ceph/snmp-notifier-rhel9:latest

```
podman run -v /CERTIFICATE_DIRECTORY_PATH:/certs:Z -v /CERTIFICATE_DIRECTORY_PATH/
domain.cert:/certs/domain.cert:Z --rm registry.redhat.io/rhel9/skopeo skopeo copy --
remove-signatures --src-creds IBM_CUSTOMER_PORTAL_LOGIN:IBM_CUSTOMER_PORTAL_PASSWORD --dest-
cert-dir=./certs/ --dest-creds PRIVATE_REGISTRY_USERNAME:PRIVATE_REGISTRY_PASSWORD docker://
cp.icr.io/cp/SRC_IMAGE:SRC_TAG docker://LOCAL_NODE_FQDN:5000/DST_IMAGE:pDST_TAG
```

- Replace `CERTIFICATE_DIRECTORY_PATH` with the directory path to the self-signed certificates.
- Replace `CERTIFICATE_DIRECTORY_PATH` and `IBM_CUSTOMER_PORTAL_PASSWORD` with your IBM Customer Portal credentials.
- Replace `PRIVATE_REGISTRY_USERNAME` and `PRIVATE_REGISTRY_PASSWORD` with the private registry credentials.
- Replace `SRC_IMAGE` and `SRC_TAG` with the name and tag of the image to copy from `cp.icr.io/cp`.
- Replace `DST_IMAGE` and `DST_TAG` with the name and tag of the image to copy to the private registry.
- Replace `LOCAL_NODE_FQDN` with the FQDN of the private registry.

For example:

```
[root@admin ~]# podman run -v /opt/registry/certs:/certs:Z -v /opt/registry/certs/
domain.cert:/certs/domain.cert:Z --rm --rm registry.redhat.io/rhel9/skopeo skopeo copy --
remove-signatures --src-creds myusername:mypassword1 --dest-cert-dir=./certs/ --dest-creds
myregistryusername:myregistrypassword1 docker://cp.icr.io/cp/ibm-ceph/ceph-7-rhel9:latest
docker://admin.lab.ibm.com:5000/ibm-ceph/ceph-7-rhel9:latest

[root@admin ~]# podman run -v /opt/registry/certs:/certs:Z -v /opt/registry/certs/
domain.cert:/certs/domain.cert:Z --rm registry.redhat.io/rhel9/skopeo skopeo copy --
remove-signatures --src-creds myusername:mypassword1 --dest-cert-dir=./certs/ --dest-
creds myregistryusername:myregistrypassword1 docker://cp.icr.io/cp/ibm-ceph/prometheus-node-
exporter:v4.12 docker://admin.lab.ibm.com:5000/ibm-ceph/prometheus-node-exporter:v4.12 docker

[root@admin ~]# podman run -v /opt/registry/certs:/certs:Z -v /opt/registry/certs/
domain.cert:/certs/domain.cert:Z --rm registry.redhat.io/rhel9/skopeo skopeo copy --
remove-signatures --src-creds myusername:mypassword1 --dest-cert-dir=./certs/ --dest-
creds myregistryusername:myregistrypassword1 docker://cp.icr.io/cp/ibm-ceph/ceph-7-dashboard-
rhel9:latest docker://admin.lab.ibm.com:5000/ibm-ceph/ceph-7-dashboard-rhel9:latest

[root@admin ~]# podman run -v /opt/registry/certs:/certs:Z -v /opt/registry/certs/
domain.cert:/certs/domain.cert:Z --rm registry.redhat.io/rhel9/skopeo skopeo copy --
remove-signatures --src-creds myusername:mypassword1 --dest-cert-dir=./certs/ --dest-
creds myregistryusername:myregistrypassword1 docker://cp.icr.io/cp/ibm-ceph/prometheus:v4.12
docker://admin.lab.ibm.com:5000/ibm-ceph/prometheus:v4.12

[root@admin ~]# podman run -v /opt/registry/certs:/certs:Z -v /opt/registry/certs/
domain.cert:/certs/domain.cert:Z --rm registry.redhat.io/rhel9/skopeo skopeo copy
--remove-signatures --src-creds myusername:mypassword1 --dest-cert-dir=./certs/ --
dest-creds myregistryusername:myregistrypassword1 docker://cp.icr.io/cp/ibm-ceph/prometheus-
alertmanager:v4.12 docker://admin.lab.ibm.com:5000/ibm-ceph/prometheus-alertmanager:v4.12
```

Note: For more image Ceph package versions, see [What are the Red Hat Ceph Storage releases and corresponding Ceph package versions?](#) within the Red Hat Customer Portal..

20. Verify that the images reside in the local registry.

Use the **curl** command.

```
curl -u PRIVATE_REGISTRY_USERNAME:PRIVATE_REGISTRY_PASSWORD https://LOCAL_NODE_FQDN:5000/v2/
catalog
```

For example:

```
[root@admin ~]# curl -u myregistryusername:myregistrypassword1 https://
admin.lab.ibm.com:5000/v2/_catalog

{"repositories":["ibm-ceph/prometheus","ibm-ceph/prometheus-alertmanager","ibm-ceph/
prometheus-node-exporter","ibm-ceph/ceph-7-dashboard-rhel9","ibm-ceph/ceph-7-rhel9"]}
```

Running the preflight playbook for a disconnected installation

Use the `cephadm-preflight.yml` Ansible playbook to configure the Ceph repository and prepare the storage cluster for bootstrapping. It also installs some prerequisites, such as `podman`, `lvm2`, `chronyd`, and `cephadm`.

About this task

The preflight playbook uses the `cephadm-ansible` inventory hosts file to identify all the nodes in the storage cluster. The default location for `cephadm-ansible`, `cephadm-preflight.yml`, and the inventory hosts file is `/usr/share/cephadm-ansible/`.

The following example shows the structure of a typical inventory file:

```
host02
host03
host04

[admin]
host01
```

The [admin] group in the inventory file contains the name of the node where the admin keyring is stored.

Note: Run the preflight playbook before you bootstrap the initial host.

Before you begin

- The cephadm-ansible package is installed on the Ansible administration node.
- Root-level access to all nodes in the storage cluster.
- Passwordless ssh is set up on all hosts in the storage cluster.
- Nodes configured to access a local YUM repository server with the following repositories enabled on respective Red Hat Enterprise Linux versions.
 - rhel-9-for-x86_64-baseos-rpms
 - rhel-9-for-x86_64-appstream-rpms
 - ```
curl https://public.dhe.ibm.com/ibmdl/export/pub/storage/ceph/beta/ibm-storage-ceph-7-beta-rhel-9.repo | sudo tee /etc/yum.repos.d/ibm-storage-ceph-7-rhel-9.repo
```

**Note:** For more information about setting up a local YUM repository, see the Red Hat knowledge base article [Creating a Local Repository and Sharing with Disconnected/Offline/Air-gapped Systems](#).

## Procedure

1. Navigate to the /usr/share/cephadm-ansible directory on the Ansible administration node.
2. Open and edit the hosts file and add your nodes.
3. Add license to install IBM Storage Ceph and click **Accept** on all nodes.

```
dnf install ibm-storage-ceph-license
```

For example,

```
[root@admin ~]# dnf install ibm-storage-ceph-license
```

4. Accept the provisions.

```
sudo touch /usr/share/ibm-storage-ceph-license/accept
```

For example,

```
[root@admin ~]# sudo touch /usr/share/ibm-storage-ceph-license/accept
```

5. Run the preflight playbook.

**Note:** When running the preflight playbook, cephadm-ansible automatically installs chronyd and ceph-common on the client nodes.

Use **ceph\_origin** parameter set to custom to use a local YUM repository.

```
ansible-playbook -i INVENTORY_FILE cephadm-preflight.yml --extra-vars "ceph_origin=custom" -e "custom_repo_url=CUSTOM_REPO_URL"
```

For example,

```
[ansible@admin cephadm-ansible]$ ansible-playbook -i hosts cephadm-preflight.yml --extra-vars "ceph_origin=custom" -e "custom_repo_url=http://mycustomrepo.lab.ibm.com/x86_64/os/"
```

After installation is complete, cephadm resides in the /usr/sbin/ directory.

- Use the **--limit** option to run the preflight playbook on a selected set of hosts in the storage cluster.

```
ansible-playbook -i INVENTORY_FILE cephadm-preflight.yml --extra-vars "ceph_origin=custom"
-e "custom_repo_url=CUSTOM_REPO_URL" --limit GROUP_NAME|NODE_NAME
```

Replace *GROUP\_NAME* with a group name from your inventory file and *NODE\_NAME* with a specific node name from your inventory file.

For example,

```
[ansible@admin cephadm-ansible]$ ansible-playbook -i hosts cephadm-preflight.yml --extra-vars
"ceph_origin=custom" -e "custom_repo_url=http://mycustomrepo.lab.ibm.com/x86_64/os/" --limit
clients
[ansible@admin cephadm-ansible]$ ansible-playbook -i hosts cephadm-preflight.yml --extra-vars
"ceph_origin=custom" -e "custom_repo_url=http://mycustomrepo.lab.ibm.com/x86_64/os/" --limit
host02
```

## Performing a disconnected installation

Before you can perform the installation, you must obtain an IBM Storage Ceph container image, either from a proxy host that has access to the IBM registry or by copying the image to your local registry.

**Important:** Before you begin the bootstrapping process, make sure that the container image that you want to use has the same version of IBM Storage Ceph as `cephadm`. If the two versions do not match, bootstrapping fails at the `Creating initial admin user` stage.

**Note:** If your local registry uses a self-signed certificate with a local registry, ensure you have added the trusted root certificate to the bootstrap host. For more information, see [“Configuring a private registry for a disconnected installation” on page 30](#).

### Prerequisites

- At least one running virtual machine (VM) or server.
- Root-level access to all nodes.
- Passwordless ssh is set up on all hosts in the storage cluster.
- The preflight playbook has been run on the bootstrap host in the storage cluster. For more information, see [“Running the preflight playbook for a disconnected installation” on page 34](#).
- A private registry has been configured and the bootstrap node has access to it. For more information, see [“Configuring a private registry for a disconnected installation” on page 30](#)
- An IBM Storage Ceph container image resides in the custom registry.

### Procedure

1. Log in to the bootstrap host.
2. Bootstrap the storage cluster:

Syntax

```
cephadm --image PRIVATE_REGISTRY_NODE_FQDN:5000/CUSTOM_IMAGE_NAME:IMAGE_TAG bootstrap --
mon-ip IP_ADDRESS --registry-url PRIVATE_REGISTRY_NODE_FQDN:5000 --registry-username
PRIVATE_REGISTRY_USERNAME --registry-password PRIVATE_REGISTRY_PASSWORD
```

- Replace *PRIVATE\_REGISTRY\_NODE\_FQDN* with the fully qualified domain name of your private registry.
- Replace *CUSTOM\_IMAGE\_NAME* and *IMAGE\_TAG* with the name and tag of the IBM Storage Ceph container image that resides in the private registry.
- Replace *IP\_ADDRESS* with the IP address of the node you are using to run `cephadm bootstrap`.
- Replace *PRIVATE\_REGISTRY\_USERNAME* with the username to create for the private registry.

- Replace `PRIVATE_REGISTRY_PASSWORD` with the password to create for the private registry username.

#### Example

```
[root@host01 ~]# cephadm --image admin.lab.ibm.com:5000/ibm-ceph/ceph-7-rhel9:latest
bootstrap --mon-ip 10.10.128.68 --registry-url admin.lab.ibm.com:5000 --registry-username
myregistryusername --registry-password myregistrypassword1
```

The script takes a few minutes to complete. Once the script completes, it provides the credentials to the IBM Storage Ceph Dashboard URL, a command to access the Ceph command-line interface (CLI), and a request to enable telemetry.

Ceph Dashboard is now available at:

```
URL: https://host01:8443/
User: admin
Password: i8nhu7zham
```

Enabling client.admin keyring and conf on hosts with "admin" label  
You can access the Ceph CLI with:

```
sudo /usr/sbin/cephadm shell --fsid 266ee7a8-2a05-11eb-b846-5254002d4916 -c /etc/ceph/
ceph.conf -k /etc/ceph/ceph.client.admin.keyring
```

Please consider enabling telemetry to help improve Ceph:

```
ceph telemetry on
```

For more information see:

```
https://docs.ceph.com/docs/master/mgr/telemetry/
```

Bootstrap complete.

After the bootstrap process is complete, configure the container images, as detailed in [“Changing configurations of custom container images for disconnected installations”](#) on page 37.

Once your storage cluster is up and running, configure additional daemons and services. For more information, see [Operations](#)

## Changing configurations of custom container images for disconnected installations

After you perform the initial bootstrap for disconnected nodes, you must specify custom container images for monitoring stack daemons. You can override the default container images for monitoring stack daemons, since the nodes do not have access to the default container registry.

**Note:** Make sure that the bootstrap process on the initial host is complete before making any configuration changes.

By default, the monitoring stack components are deployed based on the primary Ceph image. For disconnected environment of the storage cluster, you can use the latest available monitoring stack component images.

**Note:** When using a custom registry, be sure to log in to the custom registry on newly added nodes before adding any Ceph daemons.

### Syntax

```
ceph cephadm registry-login --registry-url CUSTOM_REGISTRY_NAME --registry_username
REGISTRY_USERNAME --registry_password REGISTRY_PASSWORD
```

### Example

```
ceph cephadm registry-login --registry-url myregistry --registry_username myregistryusername --
registry_password myregistrypassword1
```

For more information, see [“Performing a disconnected installation” on page 36](#).

## Prerequisites

- At least one running virtual machine (VM) or server.
- Root-level access to all nodes.
- Passwordless ssh is set up on all hosts in the storage cluster.

For the latest supported Red Hat Enterprise Linux versions, see [Compatibility matrix](#).

## Procedure

1. Set the custom container images with the `ceph config` command:

Syntax

```
ceph config set mgr mgr/cephadm/OPTION_NAME CUSTOM_REGISTRY_NAME/CONTAINER_NAME
```

Use the following options for `OPTION_NAME`:

```
container_image_prometheus
container_image_grafana
container_image_alertmanager
container_image_node_exporter
```

Example

```
[root@host01 ~]# ceph config set mgr mgr/cephadm/container_image_prometheus myregistry/
mycontainer
[root@host01 ~]# ceph config set mgr mgr/cephadm/container_image_grafana myregistry/
mycontainer
[root@host01 ~]# ceph config set mgr mgr/cephadm/container_image_alertmanager myregistry/
mycontainer
[root@host01 ~]# ceph config set mgr mgr/cephadm/container_image_node_exporter myregistry/
mycontainer
```

2. Redeploy node-exporter:

Syntax

```
ceph orch redeploy node-exporter
```

### Note:

- If any of the services do not deploy, you can redeploy them with the **ceph orch redeploy** command.
- By setting a custom image, the default values for the configuration image name and tag will be overridden, but not overwritten. The default values change when updates become available. By setting a custom image, you will not be able to configure the component for which you have set the custom image for automatic updates. You will need to manually update the configuration image name and tag to be able to install updates.
- If you choose to revert to using the default configuration, you can reset the custom container image. Use **ceph config rm** to reset the configuration option:

Syntax

```
ceph config rm mgr mgr/cephadm/OPTION_NAME
```

Example

```
ceph config rm mgr mgr/cephadm/container_image_prometheus
```

## Reference

- [“Performing a disconnected installation” on page 36](#)



## Adding hosts in disconnected deployments

If you are running a storage cluster on a private network and your host domain name server (DNS) cannot be reached through private IP, you must include both the host name and the IP address for each host you want to add to the storage cluster.

### Before you begin

- A running IBM Storage Ceph cluster.
- Root-level access to all hosts in the storage cluster.

### Procedure

1. Log into the cephadm shell.

Syntax

```
[root@host01 ~]# cephadm shell
```

2. Add the host.

Syntax

```
ceph orch host add HOST_NAME HOST_ADDRESS
```

Example

```
[ceph: root@host01 /]# ceph orch host add host03 10.10.128.70
```

## Launching the cephadm shell

The **cephadm shell** command launches a bash shell in a container with all of the Ceph packages installed. This enables you to perform “Day One” cluster setup tasks, such as installation and bootstrapping, and to invoke ceph commands.

### Prerequisites

- A running IBM Storage Ceph cluster.
- Root-level access to all nodes in the storage cluster.

### Procedure

There are two ways to launch the cephadm shell:

- Enter **cephadm shell** at the system prompt. This example invokes the **ceph -s** command from within the shell.

Example

```
[root@host01 ~]# cephadm shell
[ceph: root@host01 /]# ceph -s
```

- At the system prompt, type **cephadm shell** and the command you want to execute:

Example

```
[root@host01 ~]# cephadm shell ceph -s
cluster:
 id: f64f341c-655d-11eb-8778-fa163e914bcc
 health: HEALTH_OK

services:
 mon: 3 daemons, quorum host01,host02,host03 (age 94m)
 mgr: host01.lbnhug(active, since 59m), standbys: host02.rofgay, host03.ohipra
 mds: 1/1 daemons up, 1 standby
 osd: 18 osds: 18 up (since 10m), 18 in (since 10m)
```

```

 igw: 4 daemons active (2 hosts, 1 zones)

data:
 volumes: 1/1 healthy
 pools: 8 pools, 225 pgs
 objects: 230 objects, 9.9 KiB
 usage: 271 MiB used, 269 GiB / 270 GiB avail
 pgs: 225 active+clean

io:
 client: 85 B/s rd, 0 op/s rd, 0 op/s wr

```

**Note:** If the node contains configuration and keyring files in `/etc/ceph/`, the container environment uses the values in those files as defaults for the `cephadm` shell. If you execute the `cephadm` shell on a MON node, the `cephadm` shell inherits its default configuration from the MON container, instead of using the default configuration.

## cephadm commands

The `cephadm` is a command line tool to manage the local host for the Cephadm Orchestrator. It provides commands to investigate and modify the state of the current host.

Some of the commands are generally used for debugging.

**Note:** `cephadm` is not required on all hosts, however, it is useful when investigating a particular daemon. The `cephadm-ansible-preflight` playbook installs `cephadm` on all hosts and the `cephadm-ansible-purge` playbook requires `cephadm` be installed on all hosts to work properly.

| Command         | Description                                                                                                                                            | Syntax                                                                                                                                                                   | Example                                                                                                               |
|-----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------|
| <b>add-repo</b> | Configure the local package repository to also include the Ceph repository. This is mainly used for the disconnected installation of IBM Storage Ceph. | <code>cephadm add-repo [--release RELEASE] [--version VERSION] [--repo-url REPO_URL]</code>                                                                              | <pre>[root@host01 ~]# ./cephadm add-repo --version 15.2.1 [root@host01 ~]# ./cephadm add-repo --release pacific</pre> |
| <b>adopt</b>    | Convert an upgraded storage cluster daemon to run <code>cephadm</code> .                                                                               | <code>cephadm adopt [-h] --name DAEMON_NAME --style STYLE [--cluster CLUSTER] --legacy-dir [LEGACY_DIR] --config-json CONFIG_JSON [--skip-firewall] [--skip-pull]</code> | <pre>[root@host01 ~]# cephadm adopt --style=legacy --name prometheus.host02</pre>                                     |

Table 2. The **cephadm** commands (continued)

| Command              | Description                                                                                                                                                                                                                                                                                                              | Syntax                                                                                                                                                                                                                                                                                                                                                                 | Example                                                                                                                       |
|----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| <b>ceph-volume</b>   | This command is used to list all the devices on the particular host. Run the <code>ceph-volume</code> command inside a container Deploys OSDs with different device technologies like lvm or physical disks using pluggable tools and follows a predictable, and robust way of preparing, activating, and starting OSDs. | <code>cephadm ceph-volume inventory/simple/raw/lvm [-h] [--fsid FSID] [--config-json CONFIG_JSON] [--config CONFIG, -c CONFIG] [--keyring KEYRING, -k KEYRING]</code>                                                                                                                                                                                                  | <pre>[root@host01 ~]# cephadm ceph-volume inventory --fsid f64f341c-655d-11eb-8 778-fa163e914bcc</pre>                        |
| <b>check-host</b>    | Check the host configuration that is suitable for a Ceph cluster.                                                                                                                                                                                                                                                        | <code>cephadm check-host [--expect-hostname HOSTNAME]</code>                                                                                                                                                                                                                                                                                                           | <pre>[root@host01 ~]# cephadm check-host --expect-hostname host02</pre>                                                       |
| <b>deploy</b>        | Deploys a daemon on the local host.                                                                                                                                                                                                                                                                                      | <code>cephadm shell deploy DAEMON_TYPE [-h] [--name DAEMON_NAME] [--fsid FSID] [--config CONFIG, -c CONFIG] [--config-json CONFIG_JSON] [--keyring KEYRING] [--key KEY] [--osd-fsid OSD_FSID] [--skip-firewalld] [--tcp-ports TCP_PORTS] [--reconfig] [--allow-pttrace] [--memory-request MEMORY_REQUEST] [--memory-limit MEMORY_LIMIT] [--meta-json META_JSON]</code> | <pre>[root@host01 ~]# cephadm shell deploy mon --fsid f64f341c-655d-11eb-8 778-fa163e914bcc</pre>                             |
| <b>enter</b>         | Run an interactive shell inside a running daemon container.                                                                                                                                                                                                                                                              | <code>cephadm enter [-h] [--fsid FSID] --name NAME [command [command ...]]</code>                                                                                                                                                                                                                                                                                      | <pre>[root@host01 ~]# cephadm enter -- name 52c611f2b1d9</pre>                                                                |
| <b>help</b>          | View all the commands supported by <code>cephadm</code> .                                                                                                                                                                                                                                                                | <code>cephadm help</code>                                                                                                                                                                                                                                                                                                                                              | <pre>[root@host01 ~]# cephadm help</pre>                                                                                      |
| <b>install</b>       | Install the packages.                                                                                                                                                                                                                                                                                                    | <code>cephadm install PACKAGES</code>                                                                                                                                                                                                                                                                                                                                  | <pre>[root@host01 ~]# cephadm install ceph-common ceph-osd</pre>                                                              |
| <b>inspect-image</b> | Inspect the local Ceph container image.                                                                                                                                                                                                                                                                                  | <code>cephadm --image IMAGE_ID inspect-image</code>                                                                                                                                                                                                                                                                                                                    | <pre>[root@host01 ~]# cephadm --image 13ea90216d0be03003d1 2d7869f72ad9de5cec9e 54a27fd308e01e467c0d 4a0a inspect-image</pre> |

Table 2. The **cephadm** commands (continued)

| Command              | Description                                                                                                                                                                                                                                                                                                       | Syntax                                                                                                                                                                                                                                     | Example                                                                                                                                                                                                                                                                                                              |
|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>list-networks</b> | List the IP networks.                                                                                                                                                                                                                                                                                             | <code>cephadm list-networks</code>                                                                                                                                                                                                         | <pre>[root@host01 ~]# cephadm list-networks</pre>                                                                                                                                                                                                                                                                    |
| <b>ls</b>            | List daemon instances known to cephadm on the hosts. You can use <code>--no-detail</code> for the command to run faster, which gives details of the daemon name, fsid, style, and systemd unit per daemon. You can use <code>--legacy-dir</code> option to specify a legacy base directory to search for daemons. | <code>cephadm ls [--no-detail] [--legacy-dir LEGACY_DIR]</code>                                                                                                                                                                            | <pre>[root@host01 ~]# cephadm ls --no-detail</pre>                                                                                                                                                                                                                                                                   |
| <b>logs</b>          | Print journald logs for a daemon container. This is similar to the <code>journalctl</code> command.                                                                                                                                                                                                               | <code>cephadm logs [--fsid FSID] --name DAEMON_NAME</code><br><code>cephadm logs [--fsid FSID] --name DAEMON_NAME -- -n NUMBER # Last N lines</code><br><code>cephadm logs [--fsid FSID] --name DAEMON_NAME -- -f # Follow the logs</code> | <pre>[root@host01 ~]# cephadm logs --fsid 57bdbb48-ee04-11eb-9962-001a4a000672 --name osd.8</pre> <pre>[root@host01 ~]# cephadm logs --fsid 57bdbb48-ee04-11eb-9962-001a4a000672 --name osd.8 -- -n 20</pre> <pre>[root@host01 ~]# cephadm logs --fsid 57bdbb48-ee04-11eb-9962-001a4a000672 --name osd.8 -- -f</pre> |
| <b>prepare-host</b>  | Prepare a host for cephadm.                                                                                                                                                                                                                                                                                       | <code>cephadm prepare-host [--expect-hostname HOSTNAME]</code>                                                                                                                                                                             | <pre>[root@host01 ~]# cephadm prepare-host</pre> <pre>[root@host01 ~]# cephadm prepare-host --expect-hostname host01</pre>                                                                                                                                                                                           |
| <b>pull</b>          | Pull the Ceph image.                                                                                                                                                                                                                                                                                              | <code>cephadm [-h] [--image IMAGE_ID] pull</code>                                                                                                                                                                                          | <pre>[root@host01 ~]# cephadm --image 13ea90216d0be03003d12d7869f72ad9de5cec9e54a27fd308e01e467c0d4a0a pull</pre>                                                                                                                                                                                                    |

Table 2. The **cephadm** commands (continued)

| Command                      | Description                                                                                                                                                                                                                                                                                                                                          | Syntax                                                                                                                                                                                                                                                                                                                                                                                       | Example                                                                                                                                                                                                                                                                                                                                        |
|------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p><b>registry-login</b></p> | <p>Give cephadm login information for an authenticated registry. Cephadm attempts to log the calling host into that registry.</p>                                                                                                                                                                                                                    | <pre>cephadm registry- login --registry- url REGISTRY_URL --registry-username USERNAME --registry- password PASSWORD [--fsid FSID] [--registry-json JSON_FILE]</pre> <p>You can also use a JSON registry file containing the login info formatted as:</p> <pre>cat REGISTRY_FILE  {   "url": "REGISTRY_URL"   ,   "username": "REGISTRY_USERNAME",   "password": "REGISTRY_PASSWORD" }</pre> | <pre>[root@host01 ~]# cephadm registry- login --registry- url cp.icr.io/cp --registry-username myuser1 --registry- password mypassword1</pre> <pre>[root@host01 ~]# cat registry_file  {   "url": "cp.icr.io/ cp",   "username": "myuser",   "password": "mypass" }</pre> <pre>[root@host01 ~]# cephadm registry- login -i registry_file</pre> |
| <p><b>rm-daemon</b></p>      | <p>Remove a specific daemon instance. If you run the cephadm <code>rm-daemon</code> command on the host directly, although the command removes the daemon, the cephadm <code>mgr</code> module notices that the daemon is missing and redeploys it. This command is problematic and should be used only for experimental purposes and debugging.</p> | <pre>cephadm rm-daemon --fsid FSID -- name DAEMON_NAME [--force ] [--force- delete-data]</pre>                                                                                                                                                                                                                                                                                               | <pre>[root@host01 ~]# cephadm rm-daemon -- fsid f64f341c-655d-11eb-8 778-fa163e914bcc -- name osd.8</pre>                                                                                                                                                                                                                                      |

Table 2. The **cephadm** commands (continued)

| Command           | Description                                                                                                                                                                                                                                                                                                                                 | Syntax                                                                                                                                                                                   | Example                                                                                                 |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------|
| <b>rm-cluster</b> | Remove all the daemons from a storage cluster on that specific host where it is run. Similar to <code>rm-daemon</code> , if you remove a few daemons this way and the Ceph Orchestrator is not paused and some of those daemons belong to services that are not unmanaged, the <code>cephadm</code> orchestrator just redeploys them there. | <code>cephadm rm-cluster --fsid FSID [--force]</code>                                                                                                                                    | <pre>[root@host01 ~]# cephadm rm-cluster --fsid f64f341c-655d-11eb-8778-fa163e914bcc</pre>              |
| <b>rm-repo</b>    | Remove a package repository configuration. This is mainly used for the disconnected installation of IBM Storage Ceph.                                                                                                                                                                                                                       | <code>cephadm rm-repo [-h]</code>                                                                                                                                                        | <pre>[root@host01 ~]# cephadm rm-repo</pre>                                                             |
| <b>run</b>        | Run a Ceph daemon, in a container, in the foreground.                                                                                                                                                                                                                                                                                       | <code>cephadm run [--fsid FSID] --name DAEMON_NAME</code>                                                                                                                                | <pre>[root@host01 ~]# cephadm run --fsid f64f341c-655d-11eb-8778-fa163e914bcc --name osd.8</pre>        |
| <b>shell</b>      | Run an interactive shell with access to Ceph commands over the inferred or specified Ceph cluster. You can enter the shell using the <code>cephadm shell</code> command and run all the orchestrator commands within the shell.                                                                                                             | <code>cephadm shell [--fsid FSID] [--name DAEMON_NAME, -n DAEMON_NAME] [--config CONFIG, -c CONFIG] [--mount MOUNT, -m MOUNT] [--keyring KEYRING, -k KEYRING] [--env ENV, -e ENV]</code> | <pre>[root@host01 ~]# cephadm shell -- ceph orch ls [root@host01 ~]# cephadm shell</pre>                |
| <b>unit</b>       | Start, stop, restart, enable, and disable the daemons with this operation. This operates on the daemon's systemd unit.                                                                                                                                                                                                                      | <code>cephadm unit [--fsid FSID] --name DAEMON_NAME start/stop/restart/enable/disable</code>                                                                                             | <pre>[root@host01 ~]# cephadm unit --fsid f64f341c-655d-11eb-8778-fa163e914bcc --name osd.8 start</pre> |
| <b>version</b>    | Provides the version of the storage cluster.                                                                                                                                                                                                                                                                                                | <code>cephadm version</code>                                                                                                                                                             | <pre>[root@host01 ~]# cephadm version</pre>                                                             |

## Verifying the cluster installation

Once the cluster installation is complete, you can verify that the IBM Storage Ceph installation is running properly.

There are two ways of verifying the storage cluster installation as a root user:

- Run the **podman ps** command.
- Run the **cephadm shell ceph -s** command.

## Prerequisites

- Root-level access to all nodes in the storage cluster.

## Procedure

- Run the **podman ps** command:

Example

```
[root@host01 ~]# podman ps
```

**Note:** In the NAMES column, the unit files now include the FSID.

- Run the **cephadm shell ceph -s** command:

Example

```
[root@host01 ~]# cephadm shell ceph -s
```

```
cluster:
 id: f64f341c-655d-11eb-8778-fa163e914bcc
 health: HEALTH_OK

services:
 mon: 3 daemons, quorum host01,host02,host03 (age 94m)
 mgr: host01.lbnhug(active, since 59m), standbys: host02.rofgay, host03.ohipra
 mds: 1/1 daemons up, 1 standby
 osd: 18 osds: 18 up (since 10m), 18 in (since 10m)
 rgw: 4 daemons active (2 hosts, 1 zones)

data:
 volumes: 1/1 healthy
 pools: 8 pools, 225 pgs
 objects: 230 objects, 9.9 KiB
 usage: 271 MiB used, 269 GiB / 270 GiB avail
 pgs: 225 active+clean

io:
 client: 85 B/s rd, 0 op/s rd, 0 op/s wr
```

**Note:** The health of the storage cluster is in *HEALTH\_WARN* status as the hosts and the daemons are not added.

## Adding hosts

Bootstrapping the IBM Storage Ceph installation creates a working storage cluster, consisting of one Monitor daemon and one Manager daemon within the same container. As a storage administrator, you can add additional hosts to the storage cluster and configure them.

**Important:** For adding hosts with disconnected installations, see [“Adding hosts in disconnected deployments”](#) on page 39.

### Note:

- Running the preflight playbook installs podman, lvm2, chronyd, and cephadm on all hosts listed in the Ansible inventory file.
- When using a custom registry, be sure to log in to the custom registry on newly added nodes before adding any Ceph daemons.

### Syntax

```
ceph cephadm registry-login --registry-url CUSTOM_REGISTRY_NAME --registry_username
REGISTRY_USERNAME --registry_password REGISTRY_PASSWORD
```

## Example

```
ceph cephadm registry-login --registry-url myregistry --registry_username myregistryusername
--registry_password myregistrypassword1
```

## Prerequisites

- A running IBM Storage Ceph cluster.
- Root-level or user with sudo access to all nodes in the storage cluster.
- Register the nodes to IBM subscription.
- Ansible user with sudo and passwordless ssh access to all nodes in the storage cluster.

## Procedure

**Note:** In the following procedure, use either `root`, as indicated, or the username with which the user is bootstrapped.

1. From the node that contains the admin keyring, install the storage cluster's public SSH key in the root user's `authorized_keys` file on the new host:

Syntax

```
ssh-copy-id -f -i /etc/ceph/ceph.pub user@NEWHOST
```

Example

```
[root@host01 ~]# ssh-copy-id -f -i /etc/ceph/ceph.pub root@host02
[root@host01 ~]# ssh-copy-id -f -i /etc/ceph/ceph.pub root@host03
```

2. Navigate to the `/usr/share/cephadm-ansible` directory on the Ansible administration node.

Example

```
[ansible@admin ~]$ cd /usr/share/cephadm-ansible
```

3. From the Ansible administration node, add the new host to the Ansible inventory file. The default location for the file is `/usr/share/cephadm-ansible/hosts`. The following example shows the structure of a typical inventory file:

**Note:** If you have previously added the new host to the Ansible inventory file and run the preflight playbook on the host, skip to step 4.

Example

```
[ansible@admin ~]$ cat hosts

host02
host03
host04

[admin]
host01
```

4. Run the preflight playbook with the `--limit` option:

Syntax

```
ansible-playbook -i INVENTORY_FILE cephadm-preflight.yml --extra-vars "ceph_origin=ibm" --
limit NEWHOST
```

Example

```
[ansible@admin cephadm-ansible]$ ansible-playbook -i hosts cephadm-preflight.yml --extra-
vars "ceph_origin=ibm" --limit host02
```



The preflight playbook installs podman, lvm2, chronyd, and cephadm on the new host. After installation is complete, cephadm resides in the `/usr/sbin/` directory.

For Red Hat Enterprise Linux 9, install podman, lvm2, chronyd, and cephadm manually:

Example

```
[root@host01 ~]# dnf install podman lvm2 chronyd cephadm
```

5. From the bootstrap node, use the cephadm orchestrator to add the new host to the storage cluster:

Syntax

```
ceph orch host add NEWHOST
```

Example

```
[ceph: root@host01 /]# ceph orch host add host02
Added host 'host02' with addr '10.10.128.69'
[ceph: root@host01 /]# ceph orch host add host03
Added host 'host03' with addr '10.10.128.70'
```

6. Optional: You can also add nodes by IP address, before and after you run the preflight playbook. If you do not have DNS configured in your storage cluster environment, you can add the hosts by IP address, along with the host names.

Syntax

```
ceph orch host add HOSTNAME IP_ADDRESS
```

Example

```
[ceph: root@host01 /]# ceph orch host add host02 10.10.128.69
Added host 'host02' with addr '10.10.128.69'
```

- View the status of the storage cluster and verify that the new host has been added. The *STATUS* of the hosts is blank, in the output of the `ceph orch host ls` command.

Example

```
[ceph: root@host01 /]# ceph orch host ls
```

## Reference

- [“Registering the IBM Storage Ceph nodes” on page 12](#)
- [“Creating an Ansible user with sudo access” on page 15](#)

## Using the `addr` option to identify hosts

The `addr` option offers an additional way to contact a host. Add the IP address of the host to the `addr` option. If `ssh` cannot connect to the host by its hostname, then it uses the value stored in `addr` to reach the host by its IP address.

### Prerequisites

- A storage cluster that has been installed and bootstrapped.
- Root-level access to all nodes in the storage cluster.

### Procedure

1. Log in to the cephadm shell:

## Example

```
[root@host01 ~]# cephadm shell
```

## 2. Add the IP address:

### Syntax

```
ceph orch host add HOSTNAME IP_ADDR
```

## Example

```
[ceph: root@host01 /]# ceph orch host add host02 10.10.128.68
```

**Note:** If adding a host by hostname results in that host being added with an IPv6 address instead of an IPv4 address, use the **ceph orch host** command to specify the IP address of that host:

### Syntax

```
ceph orch host set-addr HOSTNAME IP_ADDR
```

To convert the IP address from IPv6 format to IPv4 format for a host you have added, use the following command:

```
ceph orch host set-addr HOSTNAME IPV4_ADDRESS
```

## Adding multiple hosts

Add multiple hosts at a time to the storage cluster, by using a YAML file.

**Note:** Be sure to create the `hosts.yaml` file within a host container, or create the file on the local host and then use the `cephadm shell` to mount the file within the container. The `cephadm shell` automatically places mounted files in `/mnt`. If you create the file directly on the local host and then apply the `hosts.yaml` file instead of mounting it, you might see a `File does not exist` error.

### Prerequisites

- A storage cluster that has been installed and bootstrapped.
- Root-level access to all nodes in the storage cluster.

### Procedure

1. Copy over the public ssh key to each of the hosts that you want to add.
2. Use a text editor to create a `hosts.yaml` file.
3. Add the host descriptions to the `hosts.yaml` file, as shown in the following example. Include the labels to identify placements for the daemons that you want to deploy on each host. Separate each host description with three dashes (`---`).

## Example

```
service_type: host
addr:
hostname: host02
labels:
- mon
- osd
- mgr

service_type: host
addr:
hostname: host03
labels:
- mon
- osd
- mgr

```

```
service_type: host
addr:
hostname: host04
labels:
- mon
- osd
```

4. If you created the `hosts.yaml` file directly on the local host, use the `cephadm` shell to mount the file:

Example

```
[root@host01 ~]# cephadm shell --mount hosts.yaml -- ceph orch apply -i /mnt/hosts.yaml
```

5. If you created the `hosts.yaml` file within the host container, invoke the `ceph orch apply` command:

Example

```
[ceph: root@host01 /]# ceph orch apply -i hosts.yaml
Added host 'host02' with addr '10.10.128.69'
Added host 'host03' with addr '10.10.128.70'
Added host 'host04' with addr '10.10.128.71'
```

6. View the list of hosts and their labels:

Example

```
[ceph: root@host01 /]# ceph orch host ls
HOST ADDR LABELS STATUS
host02 host02 mon osd mgr
host03 host03 mon osd mgr
host04 host04 mon osd
```

**Note:** If a host is online and operating normally, its status is blank. An offline host shows a status of OFFLINE, and a host in maintenance mode shows a status of MAINTENANCE.

## Removing hosts

Remove hosts of a Ceph cluster with the Ceph Orchestrators.

All daemons are removed with the `drain` option which adds the `_no_schedule` label to ensure that you cannot deploy any daemons or a cluster till the operation is complete.

**Important:** If you are removing the bootstrap host, be sure to copy the admin keyring and the configuration file to another host in the storage cluster before you remove the host.

## Prerequisites

- A running IBM Storage Ceph cluster.
- Root-level access to all the nodes.
- Hosts are added to the storage cluster.
- All the services are deployed.
- Cephadm is deployed on the nodes where the services have to be removed.

## Procedure

1. Log into the Cephadm shell:

Example

```
[root@host01 ~]# cephadm shell
```

2. Fetch the host details:

### Example

```
[ceph: root@host01 /]# ceph orch host ls
```

3. Drain all the daemons from the host:

### Syntax

```
ceph orch host drain HOSTNAME
```

### Example

```
[ceph: root@host01 /]# ceph orch host drain host02
```

The `_no_schedule` label is automatically applied to the host which blocks deployment.

4. Check the status of OSD removal:

### Example

```
[ceph: root@host01 /]# ceph orch osd rm status
```

When no placement groups (PGs) remain on the OSD, the OSD is decommissioned and removed from the storage cluster.

5. Check if all the daemons are removed from the storage cluster:

### Syntax

```
ceph orch ps HOSTNAME
```

### Example

```
[ceph: root@host01 /]# ceph orch ps host02
```

6. Remove the host:

### Syntax

```
ceph orch rm HOSTNAME
```

### Example

```
[ceph: root@host01 /]# ceph orch rm host02
```

## Reference

- [Adding hosts using the Ceph Orchestrator](#)
- [Listing hosts using the Ceph Orchestrator](#)

## Labeling hosts

The Ceph Orchestrator supports assigning labels to hosts. Labels are free-form and have no specific meanings. This means that you can use `mon`, `monitor`, `mycluster_monitor`, or any other text string. Each host can have multiple labels.

For example, the `mon` label can be applied to all hosts on which you want to deploy Ceph Monitor daemons, `mgr` for all hosts on which you want to deploy Ceph Manager daemons, `rgw` for Ceph Object Gateway daemons, and so on.

Labeling all the hosts in the storage cluster helps to simplify system management tasks by allowing you to quickly identify the daemons running on each host. In addition, you can use the Ceph orchestrator or a YAML file to deploy or remove daemons on hosts that have specific host labels.

## Prerequisites

- A storage cluster that has been installed and bootstrapped.

## Adding a label to a host

Use the Ceph Orchestrator to add a label to a host. Labels can be used to specify placement of daemons. Each host can have multiple labels.

### Prerequisites

- A storage cluster that has been installed and bootstrapped.
- Root-level access to all nodes in the storage cluster.
- Hosts are added to the storage cluster.

### Procedure

1. Launch the cephadm shell:

```
[root@host01 ~]# cephadm shell
[ceph: root@host01 /]#
```

2. Add a label to a host:

Syntax

```
ceph orch host label add HOSTNAME LABEL
```

Example

```
[ceph: root@host01 /]# ceph orch host label add host02 mon
```

### Verification

- List the hosts:

Example

```
[ceph: root@host01 /]# ceph orch host ls
```

## Removing a label from a host

Use the Ceph Orchestrator to remove a label from a host.

### Prerequisites

- A storage cluster that has been installed and bootstrapped.
- Root-level access to all nodes in the storage cluster.

### Procedure

1. Launch the cephadm shell:

```
[root@host01 ~]# cephadm shell
[ceph: root@host01 /]#
```

2. Use the ceph orchestrator to remove a label from a host:

Syntax

```
ceph orch host label rm HOSTNAME LABEL
```

Example

```
[ceph: root@host01 /]# ceph orch host label rm host02 mon
```

## Verification

- List the hosts:

Example

```
[ceph: root@host01 /]# ceph orch host ls
```

## Using host labels to deploy daemons on specific hosts

Use host labels to deploy daemons to specific hosts.

### About this task

There are two ways to use host labels to deploy daemons on specific hosts:

- By using the **--placement** option from the command line.  
For example, **ceph orch apply DAEMON --placement="label:LABEL"**
- By using a YAML file.

### Before you begin

- A storage cluster that has been installed and bootstrapped.
- Root-level access to all nodes in the storage cluster.

### Procedure

1. Log into the Cephadm shell.

Example

```
[root@host01 ~]# cephadm shell
```

2. List all current hosts and labels.

Example

```
[ceph: root@host01 /]# ceph orch host ls
HOST ADDR LABELS STATUS
host01 _admin mon osd mgr
host02 mon osd mgr mylabel
```

3. Optional: Use the **--placement** option to deploy a daemon from the command line.

### Deploying daemons from the command line

#### Procedure

1. Log into the Cephadm shell.

Example

```
[root@host01 ~]# cephadm shell
```

2. List all current hosts and labels.

Example

```
[ceph: root@host01 /]# ceph orch host ls
```

| HOST   | ADDR | LABELS              | STATUS |
|--------|------|---------------------|--------|
| host01 |      | _admin mon osd mgr  |        |
| host02 |      | mon osd mgr mylabel |        |

- Use the **--placement** option to deploy a daemon from the command line.

Syntax

```
ceph orch apply DAEMON --placement="label:LABEL"
```

Example

```
[ceph: root@host01 /]# ceph orch apply prometheus --placement="label:mylabel"
```

## Deploying daemons with a YAML file

### Procedure

- Log into the Cephadm shell.

Example

```
[root@host01 ~]# cephadm shell
```

- List all current hosts and labels.

Example

```
[ceph: root@host01 /]# ceph orch host ls
```

| HOST   | ADDR | LABELS              | STATUS |
|--------|------|---------------------|--------|
| host01 |      | _admin mon osd mgr  |        |
| host02 |      | mon osd mgr mylabel |        |

- Assign the daemon to a specific host label in a YAML file.

Be sure to specify the service type and label.

- Create the `placement.yml` file.

```
[ceph: root@host01 /]# vi placement.yml
```

- Specify the service type and label within the YAML file.

Example

```
service_type: prometheus
placement:
 label: "mylabel"
```

- Apply the daemon placement file.

Syntax

```
ceph orch apply -i FILENAME
```

Example

```
[ceph: root@host01 /]# ceph orch apply -i placement.yml
Scheduled prometheus update...
```

### What to do next

Verify the status of the daemons.

## Syntax

```
ceph orch ps --daemon_type=DAEMON_NAME
```

## Example

```
[ceph: root@host01 /]# ceph orch ps --daemon_type=prometheus
NAME HOST PORTS STATUS REFRESHED AGE MEM USE MEM LIM VERSION
IMAGE ID CONTAINER ID
prometheus.host02 host02 *:9095 running (2h) 8m ago 2h 85.3M - 2.22.2
ac25aac5d567 ad8c7593d7c0
```

## Adding Monitor service

A typical IBM Storage Ceph storage cluster has three or five monitor daemons deployed on different hosts. If your storage cluster has five or more hosts, deploy five Monitor nodes.

### About this task

#### Note:

- In the case of a firewall, see [Firewall settings for Ceph Monitor node](#).
- The bootstrap node is the initial monitor of the storage cluster. Be sure to include the bootstrap node in the list of hosts to which you want to deploy.
- If you want to apply Monitor service to more than one specific host, be sure to specify all of the host names within the same **ceph orch apply** command. If you specify **ceph orch apply mon --placement host1** and then specify **ceph orch apply mon --placement host2**, the second command removes the Monitor service on host1 and applies a Monitor service to host2.

If your Monitor nodes or your entire cluster are located on a single subnet, then cephadm automatically adds up to five Monitor daemons as you add new hosts to the cluster. cephadm automatically configures the Monitor daemons on the new hosts. The new hosts reside on the same subnet as the first (bootstrap) host in the storage cluster. cephadm can also deploy and scale monitors to correspond to changes in the size of the storage cluster.

### Before you begin

- Root-level access to all hosts in the storage cluster.
- A running storage cluster.

### Procedure

1. Apply five Monitor daemons to five random hosts in the storage cluster.

```
ceph orch apply mon 5
```

2. Disable automatic Monitor deployment.

```
ceph orch apply mon --unmanaged
```

## Deploying Ceph monitor nodes using host labels

A typical IBM Storage Ceph storage cluster has three or five Ceph Monitor daemons deployed on different hosts. If your storage cluster has five or more hosts, deploy five Ceph Monitor nodes. Use host labels to identify the hosts that contain Ceph Monitor nodes.

If your Ceph Monitor nodes or your entire cluster are located on a single subnet, then cephadm automatically adds up to five Ceph Monitor daemons as you add new nodes to the cluster. cephadm automatically configures the Ceph Monitor daemons on the new nodes. The new nodes reside on the



same subnet as the first (bootstrap) node in the storage cluster. cephadm can also deploy and scale monitors to correspond to changes in the size of the storage cluster.

## Prerequisites

- Root-level access to all nodes in the storage cluster.
- A running storage cluster.

## Procedure

1. Assign the mon label to the host:

Syntax

```
ceph orch host label add HOSTNAME mon
```

Example

```
[ceph: root@host01 /]# ceph orch host label add host02 mon
[ceph: root@host01 /]# ceph orch host label add host03 mon
```

2. View the current hosts and labels:

Syntax

```
ceph orch host ls
```

Example

```
[ceph: root@host01 /]# ceph orch host ls
HOST ADDR LABELS STATUS
host01 mon, mgr, _admin
host02 mon
host03 mon
host04
host05
host06
```

- Deploy Ceph Monitor daemons based on the host label:

Syntax

```
ceph orch apply mon label:mon
```

- Deploy Ceph Monitor daemons on a specific set of hosts:

Syntax

```
ceph orch apply mon HOSTNAME1,HOSTNAME2,HOSTNAME3
```

Example

```
[ceph: root@host01 /]# ceph orch apply mon host01,host02,host03
```

**Note:** Be sure to include the bootstrap node in the list of hosts to which you want to deploy.

## Adding Ceph Monitor nodes by IP address or network name

A typical IBM Storage Ceph storage cluster has three or five Ceph Monitor daemons deployed on different hosts. If your storage cluster has five or more hosts, deploy five Ceph Monitor nodes. Add Ceph Monitor nodes by IP address or network name.

A typical IBM Storage Ceph storage cluster has three or five monitor daemons deployed on different hosts. If your storage cluster has five or more hosts, IBM recommends that you deploy five Monitor nodes.

If your Monitor nodes or your entire cluster are located on a single subnet, then `cephadm` automatically adds up to five Monitor daemons as you add new nodes to the cluster. You do not need to configure the Monitor daemons on the new nodes. The new nodes reside on the same subnet as the first node in the storage cluster. The first node in the storage cluster is the bootstrap node. `cephadm` can also deploy and scale monitors to correspond to changes in the size of the storage cluster.

### Prerequisites

- Root-level access to all nodes in the storage cluster.
- A running storage cluster.

### Procedure

- To deploy each additional Ceph Monitor node:

Syntax

```
ceph orch apply mon NODE:IP_ADDRESS_OR_NETWORK_NAME [NODE:IP_ADDRESS_OR_NETWORK_NAME...]
```

Example

```
[ceph: root@host01 /]# ceph orch apply mon host02:10.10.128.69 host03:mynetwork
```

## Setting up the admin node

Use an admin node to administer the storage cluster.

An admin node contains both the cluster configuration file and the admin keyring. Both of these files are stored in the directory `/etc/ceph` and use the name of the storage cluster as a prefix.

For example, the default ceph cluster name is `ceph`. In a cluster using the default name, the admin keyring is named `/etc/ceph/ceph.client.admin.keyring`. The corresponding cluster configuration file is named `/etc/ceph/ceph.conf`.

To set up additional hosts in the storage cluster as admin nodes, apply the `_admin` label to the host you want to designate as an administrator node.

**Note:** By default, after applying the `_admin` label to a node, `cephadm` copies the `ceph.conf` and `client.admin` keyring files to that node. The `_admin` label is automatically applied to the bootstrap node unless the `--skip-admin-label` option was specified with the `cephadm bootstrap` command.

### Prerequisites

- A running storage cluster with `cephadm` installed.
- The storage cluster has running Monitor and Manager nodes.
- Root-level access to all nodes in the cluster.

### Procedure

1. Use `ceph orch host ls` to view the hosts in your storage cluster:

Example

```
[root@host01 ~]# ceph orch host ls
HOST ADDR LABELS STATUS
host01 mon, mgr, _admin
host02 mon
host03 mon, mgr
host04
host05
host06
```

2. Use the `_admin` label to designate the admin host in your storage cluster. For best results, this host should have both Monitor and Manager daemons running.

## Syntax

```
ceph orch host label add HOSTNAME _admin
```

## Example

```
[root@host01 ~]# ceph orch host label add host03 _admin
```

3. Verify that the admin host has the `_admin` label.

## Example

```
[root@host01 ~]# ceph orch host ls
HOST ADDR LABELS STATUS
host01 mon,mgr,_admin
host02 mon
host03 mon,mgr,_admin
host04
host05
host06
```

4. Log in to the admin node to manage the storage cluster.

## Removing the admin label from a host

You can use the Ceph orchestrator to remove the admin label from a host.

### Prerequisites

- A running storage cluster with `cephadm` installed and bootstrapped.
- The storage cluster has running Monitor and Manager nodes.
- Root-level access to all nodes in the cluster.

### Procedure

1. Use the **`ceph orch host ls`** command to view the hosts and identify the admin host in your storage cluster:

## Example

```
[root@host01 ~]# ceph orch host ls
HOST ADDR LABELS STATUS
host01 mon,mgr,_admin
host02 mon
host03 mon,mgr,_admin
host04
host05
host06
```

2. Log into the Cephadm shell:

## Example

```
[root@host01 ~]# cephadm shell
```

3. Use the ceph orchestrator to remove the admin label from a host:

## Syntax

```
ceph orch host label rm HOSTNAME LABEL
```

## Example

```
[ceph: root@host01 /]# ceph orch host label rm host03 _admin
```

4. Verify that the admin host has the `_admin` label.

## Example

```
[root@host01 ~]# ceph orch host ls
HOST ADDR LABELS STATUS
host01 mon,mgr,_admin
host02 mon
host03 mon,mgr
host04
host05
host06
```

**Important:** After removing the admin label from a node, ensure you remove the `ceph.conf` and `client.admin` keyring files from that node. Also, the node must be removed from the admin Ansible inventory file.

## Adding Manager service

cephadm automatically installs a Manager daemon on the bootstrap node during the bootstrapping process. Use the Ceph orchestrator to deploy additional Manager daemons.

The Ceph orchestrator deploys two Manager daemons by default. To deploy a different number of Manager daemons, specify a different number. If you do not specify the hosts where the Manager daemons should be deployed, the Ceph orchestrator randomly selects the hosts and deploys the Manager daemons to them.

**Note:** If you want to apply Manager daemons to more than one specific host, be sure to specify all of the host names within the same `ceph orch apply mgr` command. If you specify `ceph orch apply mgr --placement host1` and then specify `ceph orch apply mgr --placement host2`, the second command removes the Manager daemon on host1 and applies a Manager daemon to host2.

Use the `--placement` option to deploy to specific hosts.

## Prerequisites

- A running storage cluster.

## Procedure

- To specify that you want to apply a certain number of Manager daemons to randomly selected hosts:

Syntax

```
ceph orch apply mgr NUMBER_OF_DAEMONS
```

Example

```
[ceph: root@host01 /]# ceph orch apply mgr 3
```

- To add Manager daemons to specific hosts in your storage cluster:

Syntax

```
ceph orch apply mgr --placement "HOSTNAME1 HOSTNAME2 HOSTNAME3"
```

Example

```
[ceph: root@host01 /]# ceph orch apply mgr --placement "host02 host03 host04"
```

## Adding OSDs

Use this information to deploy OSDs on a specific device or host.

Cephadm does not provision an OSD on a device that is not available. A storage device is considered available if it meets all of the following conditions:

- The device must have no partitions.
- The device must not be mounted.
- The device must not contain a file system.
- The device must not contain a Ceph BlueStore OSD.
- The device must be larger than 5 GB.

**Important:** By default, the `osd_memory_target_autotune` parameter is set to `true` in IBM Storage Ceph.

For information about tuning OSD memory automatically, see [Automatically tuning OSD memory](#).

## Prerequisites

- A running IBM Storage Ceph cluster.

## Procedure

1. List the available devices to deploy OSDs:

Syntax

```
ceph orch device ls [--hostname=HOSTNAME1 HOSTNAME2] [--wide] [--refresh]
```

Example

```
[ceph: root@host01 /]# ceph orch device ls --wide --refresh
```

2. You can either deploy the OSDs on specific hosts or on all the available devices:

- To create an OSD from a specific device on a specific host:

Syntax

```
ceph orch daemon add osd HOSTNAME:DEVICE_PATH
```

Example

```
[ceph: root@host01 /]# ceph orch daemon add osd host02:/dev/sdb
```

- To deploy OSDs on any available and unused devices, use the `--all-available-devices` option.

Example

```
[ceph: root@host01 /]# ceph orch apply osd --all-available-devices
```

**Note:** This command creates colocated WAL and DB daemons. If you want to create non-colocated daemons, do not use this command.

## Reference

- For more information about drive specifications for OSDs, see [Advanced service specifications and filters for deploying OSDs](#)
- For more information about zapping devices to clear data on devices, see [Zapping devices for Ceph OSD deployment](#)

## Purging the Ceph storage cluster

Purging the Ceph storage cluster clears any data or connections that remain from previous deployments on your server. For Red Hat Enterprise Linux 9, use the `cephadm rm-cluster` command, since Ansible is not supported.

## Purging using Red Hat Enterprise Linux 9

### Before you begin

- A running IBM Storage Ceph.

### Procedure

1. Disable cephadm to stop all the orchestration operations to avoid deploying new daemons.

Example

```
[ceph: root#host01 /]# ceph mgr module disable cephadm
```

2. Get the FSID of the cluster.

Example

```
[ceph: root#host01 /]# ceph fsid
```

3. Purge the Ceph daemons from all hosts in the cluster.

Syntax

```
cephadm rm-cluster --force --zap-osds --fsid FSID
```

Example

```
[ceph: root#host01 /]# cephadm rm-cluster --force --zap-osds --fsid a6ca415a-cde7-11eb-a41a-002590fc2544
```

## Deploying client nodes

As a storage administrator, you can deploy client nodes by running the `cephadm-preflight.yml` and `cephadm-clients.yml` playbooks. The `cephadm-preflight.yml` playbook configures the Ceph repository and prepares the storage cluster for bootstrapping. It also installs some prerequisites, such as `podman`, `lvm2`, `chronyd`, and `cephadm`.

The `cephadm-clients.yml` playbook handles the distribution of configuration and keyring files to a group of Ceph clients.

**Note:** If you are not using the `cephadm-ansible` playbooks, after upgrading your Ceph cluster, you must upgrade the `ceph-common` package and client libraries on your client nodes. For more information, see [Upgrading the IBM Storage Ceph cluster](#).

### Prerequisites

- Root-level access to the Ansible administration node.
- Ansible user with `sudo` and passwordless `ssh` access to all nodes in the storage cluster.
- The `cephadm-ansible` package is installed.
- The `preflight` playbook has been run on the initial host in the storage cluster. For more information, see [“Running the preflight playbook” on page 19](#).
- The `client_group` variable must be specified in the Ansible inventory file.
- The `[admin]` group is defined in the inventory file with a node where the admin keyring is present at `/etc/ceph/ceph.client.admin.keyring`.

### Procedure

1. As an Ansible user, navigate to the `/usr/share/cephadm-ansible` directory on the Ansible administration node.

## Example

```
[ceph-admin@admin ~]$ cd /usr/share/cephadm-ansible
```

2. Open and edit the hosts inventory file and add the [clients] group and clients to your inventory:

## Example

```
host02
host03
host04

[clients]
client01
client02
client03

[admin]
host01
```

3. Run the cephadm-preflight.yml playbook to install the prerequisites on the clients:

## Syntax

```
ansible-playbook -i INVENTORY_FILE cephadm-preflight.yml --limit CLIENT_GROUP_NAME|
CLIENT_NODE_NAME
```

## Example

```
[ceph-admin@admin cephadm-ansible]$ ansible-playbook -i hosts cephadm-preflight.yml --limit
clients
```

4. Run the cephadm-clients.yml playbook to distribute the keyring and Ceph configuration files to a set of clients.

- a. To copy the keyring with a custom destination keyring name:

## Syntax

```
ansible-playbook -i INVENTORY_FILE cephadm-clients.yml --extra-vars
'{"fsid":"FSID","keyring":"KEYRING_PATH","client_group":"CLIENT_GROUP_NAME","conf":"CEPH_C
ONFIGURATION_PATH","keyring_dest":"KEYRING_DESTINATION_PATH"}
```

- Replace INVENTORY\_FILE with the Ansible inventory file name.
- Replace FSID with the FSID of the cluster.
- Replace KEYRING\_PATH with the full path name to the keyring on the admin host that you want to copy to the client.
- Optional: Replace CLIENT\_GROUP\_NAME with the Ansible group name for the clients to set up.
- Optional: Replace CEPH\_CONFIGURATION\_PATH with the full path to the Ceph configuration file on the admin node.
- Optional: Replace KEYRING\_DESTINATION\_PATH with the full path name of the destination where the keyring will be copied.

**Note:** If you do not specify a configuration file with the conf option when you run the playbook, the playbook generates and distributes a minimal configuration file. By default, the generated file is located at /etc/ceph/ceph.conf.

## Example

```
[ceph-admin@host01 cephadm-ansible]$ ansible-playbook -i hosts cephadm-
clients.yml --extra-vars '{"fsid":"266ee7a8-2a05-11eb-b846-5254002d4916","keyring":"/etc/
ceph/ceph.client.admin.keyring","client_group":"clients","conf":"/etc/ceph/
ceph.conf","keyring_dest":"/etc/ceph/custom.name.ceph.keyring"}
```

- b. To copy a keyring with the default destination keyring name of ceph.keyring and using the default group of clients:

## Syntax

```
ansible-playbook -i INVENTORY_FILE cephadm-clients.yml --extra-vars
'{"fsid":"FSID","keyring":"KEYRING_PATH","conf":"CONF_PATH"}'
```

## Example

```
[ceph-admin@host01 cephadm-ansible]$ ansible-playbook -i hosts cephadm-
clients.yml --extra-vars '{"fsid":"266ee7a8-2a05-11eb-b846-5254002d4916","keyring":"/etc/
ceph/ceph.client.admin.keyring","conf":"/etc/ceph/ceph.conf"}'
```

## Verification

- Log into the client nodes and verify that the keyring and configuration files exist.

### Example

```
[user@client01 ~]# ls -l /etc/ceph/
-rw-----. 1 ceph ceph 151 Jul 11 12:23 custom.name.ceph.keyring
-rw-----. 1 ceph ceph 151 Jul 11 12:23 ceph.keyring
-rw-----. 1 ceph ceph 269 Jul 11 12:23 ceph.conf
```

## Reference

- For more information about admin keys, see [Ceph User Management](#).
- For more information about the `cephadm-preflight` playbook, see [“Running the preflight playbook” on page 19](#).

# Managing an IBM Storage Ceph cluster using `cephadm-ansible` modules

As a storage administrator, you can use `cephadm-ansible` modules in Ansible playbooks to administer your IBM Storage Ceph cluster. The `cephadm-ansible` package provides several modules that wrap `cephadm` calls to let you write your own unique Ansible playbooks to administer your cluster.

**Note:** `cephadm-ansible` modules currently only support the most important tasks. Any operation not covered by `cephadm-ansible` modules must be completed using either the `command` or `shell` Ansible modules in your playbooks.

## `cephadm-ansible` modules

The `cephadm-ansible` modules are a collection of modules that simplify writing Ansible playbooks by providing a wrapper around `cephadm` and `ceph orch` commands. You can use the modules to write your own unique Ansible playbooks to administer your cluster using one or more of the modules.

The `cephadm-ansible` package includes the following modules:

- `cephadm_bootstrap`
- `ceph_orch_host`
- `ceph_config`
- `ceph_orch_apply`
- `ceph_orch_daemon`
- `cephadm_registry_login`

## `cephadm-ansible` modules options

This section lists the available options for the `cephadm-ansible` modules.



Options listed as required need to be set when using the modules in your Ansible playbooks. Options listed with a default value of `true` indicate that the option is automatically set when using the modules and you do not need to specify it in your playbook. For example, for the `cephadm_bootstrap` module, the Ceph Dashboard is installed unless you set `dashboard: false`.

| <i>Table 3. Available options for the <code>cephadm_bootstrap</code> module.</i> |                                                                                                                                        |                 |                |
|----------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------|-----------------|----------------|
| <b>cephadm_bootstrap</b>                                                         | <b>Description</b>                                                                                                                     | <b>Required</b> | <b>Default</b> |
| <code>mon_ip</code>                                                              | Ceph Monitor IP address.                                                                                                               | true            |                |
| <code>image</code>                                                               | Ceph container image.                                                                                                                  | false           |                |
| <code>docker</code>                                                              | Use <code>docker</code> instead of <code>podman</code> .                                                                               | false           |                |
| <code>fsid</code>                                                                | Define the Ceph FSID.                                                                                                                  | false           |                |
| <code>pull</code>                                                                | Pull the Ceph container image.                                                                                                         | false           | true           |
| <code>dashboard</code>                                                           | Deploy the Ceph Dashboard.                                                                                                             | false           | true           |
| <code>dashboard_user</code>                                                      | Specify a specific Ceph Dashboard user.                                                                                                | false           |                |
| <code>dashboard_password</code>                                                  | Ceph Dashboard password.                                                                                                               | false           |                |
| <code>monitoring</code>                                                          | Deploy the monitoring stack.                                                                                                           | false           | true           |
| <code>firewalld</code>                                                           | Manage firewall rules with <code>firewalld</code> .                                                                                    | false           | true           |
| <code>allow_overwrite</code>                                                     | Allow overwrite of existing <code>--output-config</code> , <code>--output-keyring</code> , or <code>--output-pub-ssh-key</code> files. | false           | false          |
| <code>registry_url</code>                                                        | URL for custom registry.                                                                                                               | false           |                |
| <code>registry_username</code>                                                   | Username for custom registry.                                                                                                          | false           |                |
| <code>registry_password</code>                                                   | Password for custom registry.                                                                                                          | false           |                |
| <code>registry_json</code>                                                       | JSON file with custom registry login information.                                                                                      | false           |                |

Table 3. Available options for the `cephadm_bootstrap` module. (continued)

| <b>cephadm_bootstrap</b>         | <b>Description</b>                                              | <b>Required</b> | <b>Default</b> |
|----------------------------------|-----------------------------------------------------------------|-----------------|----------------|
| <code>ssh_user</code>            | SSH user to use for <code>cephadm ssh</code> to hosts.          | false           |                |
| <code>ssh_config</code>          | SSH config file path for <code>cephadm SSH</code> client.       | false           |                |
| <code>allow_fqdn_hostname</code> | Allow hostname that is a fully-qualified domain name (FQDN).    | false           | false          |
| <code>cluster_network</code>     | Subnet to use for cluster replication, recovery and heartbeats. | false           |                |

Table 4. Available options for the `ceph_orch_host` module.

| <b>ceph_orch_host</b>        | <b>Description</b>                                                                                                                                                                                                                                                                                    | <b>Required</b>             | <b>Default</b> |
|------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------|----------------|
| <code>fsid</code>            | The FSID of the Ceph cluster to interact with.                                                                                                                                                                                                                                                        | false                       |                |
| <code>image</code>           | The Ceph container image to use.                                                                                                                                                                                                                                                                      | false                       |                |
| <code>name</code>            | Name of the host to add, remove, or update.                                                                                                                                                                                                                                                           | true                        |                |
| <code>address</code>         | IP address of the host.                                                                                                                                                                                                                                                                               | true when state is present. |                |
| <code>set_admin_label</code> | Set the <code>_admin</code> label on the specified host.                                                                                                                                                                                                                                              | false                       | false          |
| <code>labels</code>          | The list of labels to apply to the host.                                                                                                                                                                                                                                                              | false                       | []             |
| <code>state</code>           | If set to <code>present</code> , it ensures the name specified in <code>name</code> is present. If set to <code>absent</code> , it removes the host specified in <code>name</code> . If set to <code>drain</code> , it schedules to remove all daemons from the host specified in <code>name</code> . | false                       | present        |

Table 5. Available options for the `ceph_config` module

| <b>ceph_config</b>  | <b>Description</b>                                       | <b>Required</b>       | <b>Default</b> |
|---------------------|----------------------------------------------------------|-----------------------|----------------|
| <code>fsid</code>   | The FSID of the Ceph cluster to interact with.           | false                 |                |
| <code>image</code>  | The Ceph container image to use.                         | false                 |                |
| <code>action</code> | Whether to set or get the parameter specified in option. | false                 | set            |
| <code>who</code>    | Which daemon to set the configuration to.                | true                  |                |
| <code>option</code> | Name of the parameter to set or get.                     | true                  |                |
| <code>value</code>  | Value of the parameter to set.                           | true if action is set |                |

Table 6. Available options for the `ceph_orch_apply` module.

| <b>ceph_orch_apply</b> | <b>Description</b>                             | <b>Required</b> |
|------------------------|------------------------------------------------|-----------------|
| <code>fsid</code>      | The FSID of the Ceph cluster to interact with. | false           |
| <code>image</code>     | The Ceph container image to use.               | false           |
| <code>spec</code>      | The service specification to apply.            | true            |

Table 7. Available options for the `ceph_orch_daemon` module.

| <b>ceph_orch_daemon</b> | <b>Description</b>                                  | <b>Required</b>                                                                                                                                        |
|-------------------------|-----------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>fsid</code>       | The FSID of the Ceph cluster to interact with.      | false                                                                                                                                                  |
| <code>image</code>      | The Ceph container image to use.                    | false                                                                                                                                                  |
| <code>state</code>      | The desired state of the service specified in name. | true<br>If started, it ensures the service is started.<br>If stopped, it ensures the service is stopped.<br>If restarted, it will restart the service. |
| <code>daemon_id</code>  | The ID of the service.                              | true                                                                                                                                                   |

| <i>Table 7. Available options for the ceph_orch_daemon module. (continued)</i> |                      |                 |
|--------------------------------------------------------------------------------|----------------------|-----------------|
| <b>ceph_orch_daemon</b>                                                        | <b>Description</b>   | <b>Required</b> |
| daemon_type                                                                    | The type of service. | true            |

| <i>Table 8. Available options for the cephadm_registry_login module</i> |                                                                                                                                        |                           |                |
|-------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------|---------------------------|----------------|
| <b>cephadm_registry_login</b>                                           | <b>Description</b>                                                                                                                     | <b>Required</b>           | <b>Default</b> |
| state                                                                   | Login or logout of a registry.                                                                                                         | false                     | login          |
| docker                                                                  | Use docker instead of podman.                                                                                                          | false                     |                |
| registry_url                                                            | The URL for custom registry.                                                                                                           | false                     |                |
| registry_username                                                       | Username for custom registry.                                                                                                          | true when state is login. |                |
| registry_password                                                       | Password for custom registry.                                                                                                          | true when state is login. |                |
| registry_json                                                           | The path to a JSON file. This file must be present on remote hosts prior to running this task. This option is currently not supported. |                           |                |

## Bootstrapping a storage cluster using the cephadm\_ansible modules

As a storage administrator, you can bootstrap a storage cluster using the cephadm-ansible modules such as cephadm\_bootstrap and cephadm\_registry\_login playbook.

### Prerequisites

- An IP address for the first Ceph Monitor container, which is also the IP address for the first node in the storage cluster.
- Login access to `cp.icr.io/cp`.
- A minimum of 10 GB of free space for `/var/lib/containers/`.
- Installation of the cephadm-ansible package on the Ansible administration node.
- Passwordless SSH is set up on all hosts in the storage cluster.
- Hosts are registered with CDN.

For the latest supported Red Hat Enterprise Linux versions, see [Compatibility matrix](#).

### Procedure

1. Log in to the Ansible administration node.
2. Navigate to the `/usr/share/cephadm-ansible` directory on the Ansible administration node:

## Example

```
[ansible@admin ~]$ cd /usr/share/cephadm-ansible
```

3. Create the hosts file and add hosts, labels, and monitor IP address of the first host in the storage cluster:

## Syntax

```
sudo vi INVENTORY_FILE

HOST1 labels="[LABEL1, LABEL2]"
HOST2 labels="[LABEL1, LABEL2]"
HOST3 labels="[LABEL1]"

[admin]
ADMIN_HOST monitor_address=MONITOR_IP_ADDRESS labels="[ADMIN_LABEL, LABEL1, LABEL2]"
```

## Example

```
[ansible@admin cephadm-ansible]$ sudo vi hosts

host02 labels="['mon', 'mgr']"
host03 labels="['mon', 'mgr']"
host04 labels="['osd']"
host05 labels="['osd']"
host06 labels="['osd']"

[admin]
host01 monitor_address=10.10.128.68 labels="['_admin', 'mon', 'mgr']"
```

4. Run the preflight playbook:

## Syntax

```
ansible-playbook -i INVENTORY_FILE cephadm-preflight.yml --extra-vars "ceph_origin=ibm"
```

## Example

```
[ansible@admin cephadm-ansible]$ ansible-playbook -i hosts cephadm-preflight.yml --extra-vars "ceph_origin=ibm"
```

5. Create a playbook to bootstrap your cluster:

## Syntax

```
sudo vi PLAYBOOK_FILENAME.yml

- name: NAME_OF_PLAY
 hosts: BOOTSTRAP_HOST
 become: USE_ELEVATED_PRIVILEGES
 gather_facts: GATHER_FACTS_ABOUT_REMOTE_HOSTS
 tasks:
 - name: NAME_OF_TASK
 cephadm_registry_login:
 state: STATE
 registry_url: REGISTRY_URL
 registry_username: REGISTRY_USER_NAME
 registry_password: REGISTRY_PASSWORD

 - name: NAME_OF_TASK
 cephadm_bootstrap:
 mon_ip: "{{ monitor_address }}"
 dashboard_user: DASHBOARD_USER
 dashboard_password: DASHBOARD_PASSWORD
 allow_fqdn_hostname: ALLOW_FQDN_HOSTNAME
 cluster_network: NETWORK_CIDR
```

## Example

```
[ansible@admin cephadm-ansible]$ sudo vi bootstrap.yml
```

```

- name: bootstrap the cluster
 hosts: host01
 become: true
 gather_facts: false
 tasks:
 - name: login to registry
 cephadm_registry_login:
 state: login
 registry_url: cp.icr.io/cp
 registry_username: user1
 registry_password: mypassword1

 - name: bootstrap initial cluster
 cephadm_bootstrap:
 mon_ip: "{{ monitor_address }}"
 dashboard_user: mydashboarduser
 dashboard_password: mydashboardpassword
 allow_fqdn_hostname: true
 cluster_network: 10.10.128.0/28

```

6. Run the playbook:

Syntax

```
ansible-playbook -i INVENTORY_FILE PLAYBOOK_FILENAME.yml -vvv
```

Example

```
ansible@admin cephadm-ansible]$ ansible-playbook -i hosts bootstrap.yml -vvv
```

- Review the Ansible output after running the playbook.

## Adding or removing hosts using the `ceph_orch_host` module

Add and remove hosts in your storage cluster by using the `ceph_orch_host` module in your Ansible playbook.

### Prerequisites

- A running IBM Storage Ceph cluster.
- Register the nodes to the CDN and attach subscriptions.
- Ansible user with `sudo` and passwordless SSH access to all nodes in the storage cluster.
- Installation of the `cephadm-ansible` package on the Ansible administration node.
- New hosts have the storage cluster's public SSH key. For more information about copying the storage cluster's public SSH keys to new hosts, see ["Adding hosts" on page 45](#).

### Procedure

1. Use the following procedure to add new hosts to the cluster:

- a. Log in to the Ansible administration node.
- b. Navigate to the `/usr/share/cephadm-ansible` directory on the Ansible administration node:

Example

```
[ansible@admin ~]$ cd /usr/share/cephadm-ansible
```

c. Add the new hosts and labels to the Ansible inventory file.

Syntax

```

sudo vi INVENTORY_FILE

NEW_HOST1 labels="[LABEL1, LABEL2]"
NEW_HOST2 labels="[LABEL1, LABEL2]"
NEW_HOST3 labels="[LABEL1]"

```

```
[admin]
ADMIN_HOST monitor_address=MONITOR_IP_ADDRESS labels="[ADMIN_LABEL, LABEL1, LABEL2]"
```

### Example

```
[ansible@admin cephadm-ansible]$ sudo vi hosts

host02 labels="['mon', 'mgr']"
host03 labels="['mon', 'mgr']"
host04 labels="['osd']"
host05 labels="['osd']"
host06 labels="['osd']"

[admin]
host01 monitor_address= 10.10.128.68 labels="['_admin', 'mon', 'mgr']"
```

- d. Run the preflight playbook with the `--limit` option:

### Syntax

```
ansible-playbook -i INVENTORY_FILE cephadm-preflight.yml --extra-vars "ceph_origin=ibm"
--limit NEWHOST
```

### Example

```
[ansible@admin cephadm-ansible]$ ansible-playbook -i hosts cephadm-preflight.yml --extra-
vars "ceph_origin=ibm" --limit host02
```

The preflight playbook installs `podman`, `lvm2`, `chronyd`, and `cephadm` on the new host. After installation is complete, `cephadm` resides in the `/usr/sbin/` directory.

- e. Create a playbook to add the new hosts to the cluster:

### Syntax

```
sudo vi PLAYBOOK_FILENAME.yml

- name: PLAY_NAME
 hosts: HOSTS_OR_HOST_GROUPS
 become: USE_ELEVATED_PRIVILEGES
 gather_facts: GATHER_FACTS_ABOUT_REMOTE_HOSTS
 tasks:
 - name: NAME_OF_TASK
 ceph_orch_host:
 name: "{{ ansible_facts[hostname] }}"
 address: "{{ ansible_facts[default_ipv4][address] }}"
 labels: "{{ labels }}"
 delegate_to: HOST_TO_DELEGATE_TASK_TO

 - name: NAME_OF_TASK
 when: inventory_hostname in groups[admin]
 ansible.builtin.shell:
 cmd: CEPH_COMMAND_TO_RUN
 register: REGISTER_NAME

 - name: NAME_OF_TASK
 when: inventory_hostname in groups[admin]
 debug:
 msg: "{{ REGISTER_NAME.stdout }}"
```

**Note:** By default, Ansible executes all tasks on the host that matches the `hosts` line of your playbook. The `ceph_orch` commands must run on the host that contains the admin keyring and the Ceph configuration file. Use the `delegate_to` keyword to specify the admin host in your cluster.

### Example

```
[ansible@admin cephadm-ansible]$ sudo vi add-hosts.yml

- name: add additional hosts to the cluster
 hosts: all
```

```

become: true
gather_facts: true
tasks:
 - name: add hosts to the cluster
 ceph_orch_host:
 name: "{{ ansible_facts['hostname'] }}"
 address: "{{ ansible_facts['default_ipv4']['address'] }}"
 labels: "{{ labels }}"
 delegate_to: host01

 - name: list hosts in the cluster
 when: inventory_hostname in groups['admin']
 ansible.builtin.shell:
 cmd: ceph orch host ls
 register: host_list

 - name: print current list of hosts
 when: inventory_hostname in groups['admin']
 debug:
 msg: "{{ host_list.stdout }}"

```

In this example, the playbook adds the new hosts to the cluster and displays a current list of hosts.

- f. Run the playbook to add additional hosts to the cluster:

#### Syntax

```
ansible-playbook -i INVENTORY_FILE PLAYBOOK_FILENAME.yml
```

#### Example

```
[ansible@admin cephadm-ansible]$ ansible-playbook -i hosts add-hosts.yml
```

2. Use the following procedure to remove hosts from the cluster:

- a. Log in to the Ansible administration node.
- b. Navigate to the `/usr/share/cephadm-ansible` directory on the Ansible administration node:

#### Example

```
[ansible@admin ~]$ cd /usr/share/cephadm-ansible
```

- c. Create a playbook to remove a host or hosts from the cluster:

#### Syntax

```

sudo vi PLAYBOOK_FILENAME.yml

- name: NAME_OF_PLAY
 hosts: ADMIN_HOST
 become: USE_ELEVATED_PRIVILEGES
 gather_facts: GATHER_FACTS_ABOUT_REMOTE_HOSTS
 tasks:
 - name: NAME_OF_TASK
 ceph_orch_host:
 name: HOST_TO_REMOVE
 state: STATE

 - name: NAME_OF_TASK
 ceph_orch_host:
 name: HOST_TO_REMOVE
 state: STATE
 retries: NUMBER_OF_RETRIES
 delay: DELAY
 until: CONTINUE_UNTIL
 register: REGISTER_NAME

 - name: NAME_OF_TASK
 ansible.builtin.shell:
 cmd: ceph orch host ls
 register: REGISTER_NAME

 - name: NAME_OF_TASK

```



```
debug:
 msg: "{{ REGISTER_NAME.stdout }}"
```

## Example

```
[ansible@admin cephadm-ansible]$ sudo vi remove-hosts.yml

- name: remove host
 hosts: host01
 become: true
 gather_facts: true
 tasks:
 - name: drain host07
 ceph_orch_host:
 name: host07
 state: drain

 - name: remove host from the cluster
 ceph_orch_host:
 name: host07
 state: absent
 retries: 20
 delay: 1
 until: result is succeeded
 register: result

 - name: list hosts in the cluster
 ansible.builtin.shell:
 cmd: ceph orch host ls
 register: host_list

 - name: print current list of hosts
 debug:
 msg: "{{ host_list.stdout }}"
```

In this example, the playbook tasks drain all daemons on host07, removes the host from the cluster, and displays a current list of hosts.

d. Run the playbook to remove host from the cluster:

## Syntax

```
ansible-playbook -i INVENTORY_FILE PLAYBOOK_FILENAME.yml
```

## Example

```
[ansible@admin cephadm-ansible]$ ansible-playbook -i hosts remove-hosts.yml
```

## Verification

- Review the Ansible task output displaying the current list of hosts in the cluster:

## Example

```
TASK [print current hosts]

Friday 24 June 2022 14:52:40 -0400 (0:00:03.365) 0:02:31.702 *****
ok: [host01] =>
msg: |-
 HOST ADDR LABELS STATUS
 host01 10.10.128.68 _admin mon mgr
 host02 10.10.128.69 mon mgr
 host03 10.10.128.70 mon mgr
 host04 10.10.128.71 osd
 host05 10.10.128.72 osd
 host06 10.10.128.73 osd
```

## Setting configuration options using the ceph\_config module

As a storage administrator, you can set or get IBM Storage Ceph configuration options using the `ceph_config` module.

### Prerequisites

- A running IBM Storage Ceph cluster.
- Ansible user with `sudo` and passwordless SSH access to all nodes in the storage cluster.
- Installation of the `cephadm-ansible` package on the Ansible administration node.
- The Ansible inventory file contains the cluster and admin hosts. For more information about adding hosts to your storage cluster, see [“Adding or removing hosts using the `ceph\_orch\_host` module” on page 68.](#)

### Procedure

1. Log in to the Ansible administration node.
2. Navigate to the `/usr/share/cephadm-ansible` directory on the Ansible administration node:

Example

```
[ansible@admin ~]$ cd /usr/share/cephadm-ansible
```

3. Create a playbook with configuration changes:

Syntax

```
sudo vi PLAYBOOK_FILENAME.yml

- name: PLAY_NAME
 hosts: ADMIN_HOST
 become: USE_ELEVATED_PRIVILEGES
 gather_facts: GATHER_FACTS_ABOUT_REMOTE_HOSTS
 tasks:
 - name: NAME_OF_TASK
 ceph_config:
 action: GET_OR_SET
 who: DAEMON_TO_SET_CONFIGURATION_TO
 option: CEPH_CONFIGURATION_OPTION
 value: VALUE_OF_PARAMETER_TO_SET

 - name: NAME_OF_TASK
 ceph_config:
 action: GET_OR_SET
 who: DAEMON_TO_SET_CONFIGURATION_TO
 option: CEPH_CONFIGURATION_OPTION
 register: REGISTER_NAME

 - name: NAME_OF_TASK
 debug:
 msg: "MESSAGE_TO_DISPLAY {{ REGISTER_NAME.stdout }}"
```

Example

```
[ansible@admin cephadm-ansible]$ sudo vi change_configuration.yml

- name: set pool delete
 hosts: host01
 become: true
 gather_facts: false
 tasks:
 - name: set the allow pool delete option
 ceph_config:
 action: set
 who: mon
 option: mon_allow_pool_delete
 value: true
```

```

- name: get the allow pool delete setting
 ceph_config:
 action: get
 who: mon
 option: mon_allow_pool_delete
 register: verify_mon_allow_pool_delete

- name: print current mon_allow_pool_delete setting
 debug:
 msg: "the value of 'mon_allow_pool_delete' is
 {{ verify_mon_allow_pool_delete.stdout }}"

```

In this example, the playbook first sets the `mon_allow_pool_delete` option to `false`. The playbook then gets the current `mon_allow_pool_delete` setting and displays the value in the Ansible output.

#### 4. Run the playbook:

##### Syntax

```
ansible-playbook -i INVENTORY_FILE _PLAYBOOK_FILENAME.yml
```

##### Example

```
[ansible@admin cephadm-ansible]$ ansible-playbook -i hosts change_configuration.yml
```

## Verification

- Review the output from the playbook tasks.

##### Example

```

TASK [print current mon_allow_pool_delete setting]

Wednesday 29 June 2022 13:51:41 -0400 (0:00:05.523) 0:00:17.953 *****
ok: [host01] =>
 msg: the value of 'mon_allow_pool_delete' is true

```

## Reference

- For more information, see [Configuring](#).

## Applying a service specification using the `ceph_orch_apply` module

As a storage administrator, you can apply service specifications to your storage cluster using the `ceph_orch_apply` module in your Ansible playbooks. A service specification is a data structure to specify the service attributes and configuration settings that is used to deploy the Ceph service. You can use a service specification to deploy Ceph service types like `mon`, `crash`, `mgs`, `mgr`, `osd`, `rdb`, or `rbd-mirror`.

## Prerequisites

- A running IBM Storage Ceph cluster.
- Ansible user with `sudo` and passwordless SSH access to all nodes in the storage cluster.
- Installation of the `cephadm-ansible` package on the Ansible administration node.
- The Ansible inventory file contains the cluster and admin hosts. For more information about adding hosts to your storage cluster, see [“Adding or removing hosts using the `ceph\_orch\_host` module”](#) on page 68.

## Procedure

1. Log in to the Ansible administration node.
2. Navigate to the `/usr/share/cephadm-ansible` directory on the Ansible administration node:

## Example

```
[ansible@admin ~]$ cd /usr/share/cephadm-ansible
```

### 3. Create a playbook with the service specifications:

#### Syntax

```
sudo vi PLAYBOOK_FILENAME.yml

- name: PLAY_NAME
 hosts: HOSTS_OR_HOST_GROUPS
 become: USE_ELEVATED_PRIVILEGES
 gather_facts: GATHER_FACTS_ABOUT_REMOTE_HOSTS
 tasks:
 - name: NAME_OF_TASK
 ceph_orch_apply:
 spec: |
 service_type: SERVICE_TYPE
 service_id: UNIQUE_NAME_OF_SERVICE
 placement:
 host_pattern: HOST_PATTERN_TO_SELECT_HOSTS
 label: LABEL
 spec:
 SPECIFICATION_OPTIONS:
```

## Example

```
[ansible@admin cephadm-ansible]$ sudo vi deploy_osd_service.yml

- name: deploy osd service
 hosts: host01
 become: true
 gather_facts: true
 tasks:
 - name: apply osd spec
 ceph_orch_apply:
 spec: |
 service_type: osd
 service_id: osd
 placement:
 host_pattern: '*'
 label: osd
 spec:
 data_devices:
 all: true
```

In this example, the playbook deploys the Ceph OSD service on all hosts with the label `osd`.

### 4. Run the playbook:

#### Syntax

```
ansible-playbook -i INVENTORY_FILE _PLAYBOOK_FILENAME.yml
```

## Example

```
[ansible@admin cephadm-ansible]$ ansible-playbook -i hosts deploy_osd_service.yml
```

## Verification

- Review the output from the playbook tasks.

## Managing Ceph daemon states using the `ceph_orch_daemon` module

Start, stop, and restart Ceph daemons on hosts using the `ceph_orch_daemon` module in your Ansible playbooks.

## Prerequisites

- A running IBM Storage Ceph cluster.
- Ansible user with sudo and passwordless SSH access to all nodes in the storage cluster.
- Installation of the `cephadm-ansible` package on the Ansible administration node.
- The Ansible inventory file contains the cluster and admin hosts. For more information about adding hosts to your storage cluster, see [“Adding or removing hosts using the `ceph\_orch\_host` module” on page 68.](#)

## Procedure

1. Log in to the Ansible administration node.
2. Navigate to the `/usr/share/cephadm-ansible` directory on the Ansible administration node:

Example

```
[ansible@admin ~]$ cd /usr/share/cephadm-ansible
```

3. Create a playbook with daemon state changes:

Syntax

```
sudo vi PLAYBOOK_FILENAME.yml

- name: PLAY_NAME
 hosts: ADMIN_HOST
 become: USE_ELEVATED_PRIVILEGES
 gather_facts: GATHER_FACTS_ABOUT_REMOTE_HOSTS
 tasks:
 - name: NAME_OF_TASK
 ceph_orch_daemon:
 state: STATE_OF_SERVICE
 daemon_id: DAEMON_ID
 daemon_type: TYPE_OF_SERVICE
```

Example

```
[ansible@admin cephadm-ansible]$ sudo vi restart_services.yml

- name: start and stop services
 hosts: host01
 become: true
 gather_facts: false
 tasks:
 - name: start osd.0
 ceph_orch_daemon:
 state: started
 daemon_id: 0
 daemon_type: osd

 - name: stop mon.host02
 ceph_orch_daemon:
 state: stopped
 daemon_id: host02
 daemon_type: mon
```

In this example, the playbook starts the OSD with an ID of 0 and stops a Ceph Monitor with an id of host02.

4. Run the playbook:

Syntax

```
ansible-playbook -i INVENTORY_FILE _PLAYBOOK_FILENAME.yml
```

## Example

```
[ansible@admin cephadm-ansible]$ ansible-playbook -i hosts restart_services.yml
```

## Verification

- Review the output from the playbook tasks.

## Comparison between Ceph Ansible and Cephadm

Understand the differences between Cephadm and Ceph-Ansible playbooks for the containerized deployment of the storage cluster.

The tables compare Cephadm with Ceph-Ansible playbooks for managing the containerized deployment of a Ceph cluster for day one and day two operations.

| Description                                  | Ceph-Ansible                                                                                            | Cephadm                                                                                                                                                                       |
|----------------------------------------------|---------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Installation of the IBM Storage Ceph cluster | Run the <code>site-container.yml</code> playbook.                                                       | Run <code>cephadm bootstrap</code> command to bootstrap the cluster on the admin node.                                                                                        |
| Addition of hosts                            | Use the Ceph Ansible inventory.                                                                         | Run <code>ceph orch add host HOST_NAME</code> to add hosts to the cluster.                                                                                                    |
| Gathering Ceph logs                          | Run the <code>gather_ceph_logs</code> playbook.                                                         | Run the <code>journalctl</code> command.                                                                                                                                      |
| Addition of monitors                         | Run the <code>add-mon.yml</code> playbook.                                                              | Run the <code>ceph orch apply mon</code> command.                                                                                                                             |
| Addition of managers                         | Run the <code>site-container.yml</code> playbook.                                                       | Run the <code>ceph orch apply mgr</code> command.                                                                                                                             |
| Addition of OSDs                             | Run the <code>add-osd.yml</code> playbook.                                                              | Run the <code>ceph orch apply osd</code> command to add OSDs on all available devices or on specific hosts.                                                                   |
| Addition of OSDs on specific devices         | Select the devices in the <code>osd.yml</code> file and then run the <code>add-osd.yml</code> playbook. | Select the <code>paths</code> filter under the <code>data_devices</code> in the <code>osd.yml</code> file and then run <code>ceph orch apply -i FILE_NAME.yml</code> command. |
| Addition of MDS                              | Run the <code>site-container.yml</code> playbook.                                                       | Run the <code>ceph orch apply FILESYSTEM_NAME</code> command to add MDS.                                                                                                      |
| Addition of Ceph Object Gateway              | Run the <code>site-container.yml</code> playbook.                                                       | Run the <code>ceph orch apply rgw</code> commands to add Ceph Object Gateway.                                                                                                 |

| Description                               | Ceph-Ansible                                                               | Cephadm                                                                                    |
|-------------------------------------------|----------------------------------------------------------------------------|--------------------------------------------------------------------------------------------|
| Removing hosts                            | Use the Ansible inventory.                                                 | Run <code>ceph orch host rm HOST_NAME</code> to remove the hosts.                          |
| Removing monitors                         | Run the <code>shrink-mon.yml</code> playbook.                              | Run <code>ceph orch apply mon</code> to redeploy other monitors.                           |
| Removing managers                         | Run the <code>shrink-mon.yml</code> playbook.                              | Run <code>ceph orch apply mgr</code> to redeploy other managers.                           |
| Removing OSDs                             | Run the <code>shrink-osd.yml</code> playbook.                              | Run <code>ceph orch osd rm OSD_ID</code> to remove the OSDs.                               |
| Removing MDS                              | Run the <code>shrink-mds.yml</code> playbook.                              |                                                                                            |
| Deployment of Ceph Object Gateway         | Run the <code>site-container.yml</code> playbook.                          | Run <code>ceph orch apply rgw _SERVICE_NAME_</code> to deploy Ceph Object Gateway service. |
| Removing Ceph Object Gateway              | Run the <code>shrink-rgw.yml</code> playbook.                              | Run <code>ceph orch rm SERVICE_NAME</code> to remove the specific service.                 |
| Block device mirroring                    | Run the <code>site-container.yml</code> playbook.                          | Run <code>ceph orch apply rbd-mirror</code> command.                                       |
| Minor version upgrade of IBM Storage Ceph | Run the <code>infrastructure-playbooks/rolling_update.yml</code> playbook. | Run <code>ceph orch upgrade start</code> command.                                          |
| Deployment of monitoring stack            | Edit the <code>all.yml</code> file during installation.                    | Run the <code>ceph orch apply -i FILE.yml</code> after specifying the services.            |

## What to do next? Day 2

As a storage administrator, once you have installed and configured IBM Storage Ceph, you are ready to perform "Day Two" operations for your storage cluster. These operations include adding metadata servers (MDS) and object gateways (RGW), and configuring services.

For more information about how to use the cephadm orchestrator to perform "Day Two" operations, see [Operations](#).

To deploy, configure, and administer the Ceph Object Gateway on "Day Two" operations, see [Object Gateway](#).







