
OpenStack-Ansible Documentation: ceph_client role

Release 18.1.0.dev153

OpenStack-Ansible Contributors

Apr 17, 2024

CONTENTS

1	Configuring the Ceph client (optional)	1
1.1	Authentication	1
1.2	Configuration file overrides	1
1.3	Extra client configuration files	2
1.4	Monitors	3
1.5	Configure os_gnocchi with ceph_client	3
2	Ceph keyring from file example	5
2.1	Network configuration assumptions	5
2.2	Configuration	6
3	Default variables	9
4	Required variables	13
5	Dependencies	15
6	Example playbook	17

CONFIGURING THE CEPH CLIENT (OPTIONAL)

Ceph is a massively scalable, open source, distributed storage system.

These links provide details on how to use Ceph with OpenStack:

- [Ceph Block Devices and OpenStack](#)
- [Ceph - The De Facto Storage Backend for OpenStack \(Hong Kong Summit talk\)](#)
- [OpenStack Config Reference - Ceph RADOS Block Device \(RBD\)](#)
- [OpenStack-Ansible and Ceph Working Example](#)

Note: Configuring Ceph storage servers is outside the scope of this documentation.

1.1 Authentication

We recommend the `cephx` authentication method in the [Ceph config reference](#). OpenStack-Ansible enables `cephx` by default for the Ceph client. You can choose to override this setting by using the `cephx` Ansible variable:

```
cephx: False
```

Deploy Ceph on a trusted network if disabling `cephx`.

1.2 Configuration file overrides

OpenStack-Ansible provides the `ceph_conf_file` variable. This allows you to specify configuration file options to override the default Ceph configuration:

```
ceph_conf_file: |
  [global]
  fsid = 4037aa5f-abde-4378-9470-f73dbd6ceaba
  mon_initial_members = mon1.example.local,mon2.example.local,mon3.example.
  ↪local
  mon_host = 172.29.244.151,172.29.244.152,172.29.244.153
  auth_cluster_required = cephx
  auth_service_required = cephx
  auth_client_required = cephx
```

The use of the `ceph_conf_file` variable is optional. By default, OpenStack-Ansible obtains a copy of `ceph.conf` from one of your Ceph monitors. This transfer of `ceph.conf` requires the OpenStack-Ansible deployment host public key to be deployed to all of the Ceph monitors. More details are available here: [Deploying SSH Keys](#).

The following minimal example configuration sets `nova` and `glance` to use `ceph` pools: `ephemeral-vms` and `images` respectively. The example uses `cephx` authentication, and requires existing `glance` and `cinder` accounts for `images` and `ephemeral-vms` pools.

```
glance_default_store: rbd
nova_libvirt_images_rbd_pool: ephemeral-vms
```

For a complete example how to provide the necessary configuration for a Ceph backend without necessary access to Ceph monitors via SSH please see [Ceph keyring from file example](#).

1.3 Extra client configuration files

Deployers can specify extra Ceph configuration files to support multiple Ceph cluster backends via the `ceph_extra_confs` variable.

```
ceph_extra_confs:
- src: "/opt/rdb-1.conf"
  dest: "/etc/ceph/rdb-1.conf"
- src: "/opt/rdb-2.conf"
  dest: "/etc/ceph/rdb-2.conf"
```

These config file sources must be present on the deployment host.

Alternatively, deployers can specify more options in `ceph_extra_confs` to deploy keyrings, `ceph.conf` files, and configure `libvirt` secrets.

```
ceph_extra_confs:
- src: "/etc/openstack_deploy/ceph2.conf"
  dest: "/etc/ceph/ceph2.conf"
  mon_host: 192.168.1.2
  client_name: cinder2
  keyring_src: /etc/openstack_deploy/ceph2.client.cinder2.keyring
  keyring_dest: /etc/ceph/ceph2.client.cinder2.keyring
  secret_uuid: '{{ cinder_ceph_client_uuid2 }}'
- src: "/etc/openstack_deploy/ceph3.conf"
  dest: "/etc/ceph/ceph3.conf"
  mon_host: 192.168.1.3
  client_name: cinder3
  keyring_src: /etc/openstack_deploy/ceph3.client.cinder3.keyring
  keyring_dest: /etc/ceph/ceph3.client.cinder3.keyring
  secret_uuid: '{{ cinder_ceph_client_uuid3 }}'
```

The primary aim of this feature is to deploy multiple `ceph` clusters as `cinder` backends and enable `nova/libvirt` to mount block volumes from those backends. These settings do not override the normal deployment of `ceph` client and associated setup tasks.

Deploying multiple `ceph` clusters as `cinder` backends requires the following adjustments to each backend in `cinder_backends`

```
rbd_ceph_conf: /etc/ceph/ceph2.conf
rbd_pool: cinder_volumes_2
rbd_user: cinder2
rbd_secret_uuid: '{{ cinder_ceph_client_uuid2 }}'
volume_backend_name: volumes2
```

The dictionary keys `rbd_ceph_conf`, `rbd_user`, and `rbd_secret_uuid` must be unique for each ceph cluster to be used as a `cinder_backend`.

1.4 Monitors

The [Ceph Monitor](#) maintains a master copy of the cluster map. OpenStack-Ansible provides the `ceph_mons` variable and expects a list of IP addresses for the Ceph Monitor servers in the deployment:

```
ceph_mons:
- 172.29.244.151
- 172.29.244.152
- 172.29.244.153
```

1.5 Configure os_gnocchi with ceph_client

If the `os_gnocchi` role is going to utilize the `ceph_client` role, the following configurations need to be added to the user variable file:

```
ceph_extra_components:
- component: gnocchi_api
  package: "{{ python_ceph_packages }}"
  client:
    - '{{ gnocchi_ceph_client }}'
  service: '{{ ceph_gnocchi_service_names }}'
```


CEPH KEYRING FROM FILE EXAMPLE

OpenStack-Ansible (OSA) allows to deploy an OpenStack environment that uses an existing Ceph cluster for block storage for images, volumes and instances. Interaction with the Ceph cluster is normally done using SSH to Ceph MONs. To avoid the SSH access to the Ceph cluster nodes all necessary client configurations can be read from files. This example describes what these files need to contain.

This example has just a single main requirement. You need to configure a storage network in your OpenStack environment. Both Ceph services - the MONs and the OSDs - need to be connected to this storage network, too. On the OpenStack side you need to connect the affected services to the storage network. Glance to store images in Ceph, Cinder to create volumes in Ceph and in most cases the compute nodes to use volumes and maybe store ephemeral discs in Ceph.

2.1 Network configuration assumptions

The following CIDR assignments are used for this environment.

Network	CIDR
Storage Network	172.29.244.0/22

2.1.1 IP assignments

The following host name and IP address assignments are used for this environment.

Host name	Storage IP
ceph1	172.29.244.18
ceph2	172.29.244.19
ceph3	172.29.244.20

2.2 Configuration

2.2.1 Environment customizations

For a ceph environment, you can run the `cinder-volume` in a container. By default `cinder-volume` runs on the host. See [here](#) an example how to a service in a container.

2.2.2 User variables

The `/etc/openstack_deploy/user_variables.yml` file defines the global overrides for the default variables.

For this example environment, we configure an existing Ceph cluster, that we want the OpenStack environment to connect to. Your `/etc/openstack_deploy/user_variables.yml` must have the following content to configure ceph for images, volumes and instances. If not all necessary block storages should be provided from the Ceph backend, do only include the block storage you want to store in Ceph:

```
---
# OSA options for using an existing Ceph deployment. This example can be
# used
# if all configuration needs to come from OSA configuration files instead
# of
# the Ceph MONs.

# Directory containing the Ceph keyring files with access credentials.
ceph_keyrings_dir: /etc/openstack_deploy/ceph-keyrings

# General Ceph configuration file containing the information for Ceph
# clients
# to connect to the Ceph cluster.
ceph_conf_file: |
    [global]
    mon initial members = ceph1,ceph2,ceph3
    ## Ceph clusters starting with the Nautilus release can support the v2
    # wire protocol
    mon host = [v2:172.29.244.18:3300,v1:172.29.244.18:6789],[v2:172.29.244.
    # 19:3300,v1:172.29.244.19:6789],[v2:172.29.244.20:3300,v1:172.29.244.
    # 20:6789]
    ## for a Ceph cluster not supporting the v2 wire protocol (before
    # Nautilus release)
    # mon host = [v1:172.29.244.18:6789],[v1:172.29.244.19:6789],[v1:172.29.
    # 244.20:6789]

# For configuring the Ceph backend for Glance to store images in Ceph.
glance_ceph_client: glance
glance_default_store: rbd
glance_rbd_store_pool: images

# For configuring a backend in Cinder to store volumes in Ceph. This
# configuration will be used for Nova compute and libvirt to access
# volumes.
cinder_ceph_client: cinder
```

(continues on next page)

(continued from previous page)

```
cinder_backends:
  rbd:
    volume_driver: cinder.volume.drivers.rbd.RBDDriver
    rbd_pool: volumes
    rbd_ceph_conf: /etc/ceph/ceph.conf
    rbd_store_chunk_size: 8
    volume_backend_name: rbd
    rbd_user: "{{ cinder_ceph_client }}"
    rbd_secret_uuid: "{{ cinder_ceph_client_uuid }}"
    report_discard_supported: true

# Configuration for Nova compute and libvirt to store ephemeral discs in_
↳Ceph.
nova_libvirt_images_rbd_pool: vms
```

2.2.3 Ceph keyrings

With the above settings in the `/etc/openstack_deploy/user_variables.yml` we configured to read the credentials for accessing the Ceph cluster in the `/etc/openstack_deploy/ceph-keyrings/` directory. We need to place now the keyring files for Ceph credentials into this directory. They need to be named according to the ceph client names, e.g. `glance.keyring` according to `glance_ceph_client: glance`. See the following example for the file contents:

```
[client.glance]
key = AQC93h9fAAAAABAAUrAlQF+xJnjD6E8ChZkTaQ==
```

This Ansible role installs the Ceph operating system packages used to interact with a Ceph cluster.

To clone or view the source code for this repository, visit the role repository for [ceph_client](#).

DEFAULT VARIABLES

```
# Set the package install state for distribution packages
# Options are 'present' and 'latest'
ceph_client_package_state: "{{ package_state | default('latest') }}"

# to use Ceph in OSA, you need to
# - have the needed pools and a client user (for glance, cinder and/or
  →nova)
#   pre-provisioned in your ceph cluster; OSA assumes to have root access
  →to
#   the monitor hosts
# - configure / overrules following defaults in osa's user config
# - some ceph specific vars are (also) part of other role defaults:
#   * glance
#   * nova
# - cinder gets configured with ceph if there are cinder backends defined
  →with
#   the rbd driver (see openstack_user_config.yml.example)

# The ceph_pkg_source variable controls the install source for the Ceph
  →packages.
# Valid values include:
#   * ceph This option installs Ceph from a ceph.com repo. Additional
  →variables to
#       adjust items such as Ceph release and regional download mirror
  →can be found
#       in vars/*.yml
#
#   * distro This options installs Ceph from the operating system's
  →default repository and
#       unlike the other options does not attempt to manage package
  →keys or add additional
#       package repositories.
ceph_pkg_source: ceph
ceph_stable_release: octopus
ceph_apt_pinned_packages: [{ package: "*", release: "ceph.com", priority:
  →1001 }]

# Ceph Authentication
cephx: true

# Ceph Monitors
# A list of the IP addresses for your Ceph monitors
ceph_mons: []
```

(continues on next page)

(continued from previous page)

```

# Path to local ceph.conf file
# Leave this commented to obtain a ceph.conf from one of the monitors,
↳ defined in ceph_mons
#ceph_conf_file: |
# [global]
# fsid = 4037aa5f-abde-4378-9470-f73dbd6ceaba
# mon_initial_members = mon1.example.local,mon2.example.local,mon3.
↳ example.local
# mon_host = 10.16.5.40,10.16.5.41,10.16.5.42
# auth_cluster_required = cephx
# auth_service_required = cephx
# auth_client_required = cephx

# Path to local keyrings directory
# If you want to provide keyrings from existing files, because you do not,
↳ have ssh access to the monitors
# set the path to the repository containing the keyrings files.
# ie : ceph_keyrings_dir: /etc/openstack_deploy/ceph-conf
# The filenames inside the keyring directory must be in the structure of,
↳ client-name.keyring
# ie: /etc/openstack_deploy/ceph-conf
# cinder.keyring
# glance.keyring
# etc..
#ceph_keyrings_dir: "/etc/openstack/ceph-keyrings"

# Ceph client usernames for glance, cinder+nova and gnocchi
glance_ceph_client: glance
cinder_ceph_client: cinder
manila_ceph_client: manila
cinder_backup_ceph_client: cinder-backup
gnocchi_ceph_client: gnocchi

# by default we assume you use rbd for both cinder and nova, and as libvirt
# needs to access both volumes (cinder) as boot disks (nova) we default to
# reuse the cinder_ceph_client
# only need to change this if you'd use ceph for boot disks and not for,
↳ volumes
nova_ceph_client: '{{ cinder_ceph_client }}'

# overruled in user_secrets:
# TODO(odyssey4me) - the uuid should be removed, there should be no,
↳ defaults for secrets
nova_ceph_client_uuid: 457eb676-33da-42ec-9a8c-9293d545c337
cephkeys_access_group: ceph

openstack_service_system_user: null

ceph_cinder_service_names:
- cinder-volume
- cinder-backup

ceph_nova_service_names:
- nova-compute

```

(continues on next page)

(continued from previous page)

```
ceph_manila_service_names:
  - manila-api
  - manila-data
  - manila-share

ceph_glance_service_names:
  - glance-api

ceph_gnocchi_service_names:
  - gnocchi-api
  - gnocchi-metricd

ceph_extra_auth_groups: "{{ ceph_extra_config_groups }}"
ceph_extra_config_groups:
  - cinder_backup
  - cinder_volume
ceph_extra_compute_group: nova_compute

ceph_client_ceph_conf_overrides: "{{ ceph_conf_overrides | default({}) }}"
```


REQUIRED VARIABLES

None.

DEPENDENCIES

None.

EXAMPLE PLAYBOOK

```
- name: Install Ceph client
  hosts: all
  user: root
  roles:
    - role: "openstack-ansible-ceph_client"
```