Mirantis Technical Bulletin 2019-004

March 6, 2019

Add `osd blacklist` to the default `mon` permissions in Ceph

ISSUE

When a Ceph client fails, for example, due to power outage or segfault, the client cannot reconnect and write to any devices on reboot.

AFFECTS

All deployments without the `osd blacklist` command in the default permissions for the Ceph clients

FIX AVAILABLE FOR

This is a safe permission for all supported Ceph releases

SECURITY IMPACT

None

HOW TO DETERMINE IF YOU ARE AFFECTED

1. Log in to any `cmn` node.
2. Identify the existing caps for the Ceph client users. For example:

```
ceph auth list | egrep -A 3 'client.cinder|client.glance|client.nova'
```

Example of system response:

```
client.cinder
key: AQAtVzNcGOe5LBAAW437TQA94yBURi7c9Hx8UA==
caps: [mon] allow r
```



```
caps: [osd] profile rbd pool=volumes, profile rbd-read-only pool=images,
profile rbd pool=backups
client.glance
key: AQAuVzNcLL6vARAAqUS4Klq0ArDQvYlcMEEd6g==
caps: [mon] allow r
caps: [osd] profile rbd pool=images
client.nova
key: AQDAXTNcDV4QEBAA4/YaL+Yu/n6Ohk48Z3PtzA==
caps: [mon] allow r
caps: [osd] profile rbd pool=vms, profile rbd-read-only pool=images
```

All entries with `caps: [mon]` must have the `allow r, allow command "osd blacklist"` value set. The above example of system response does not include the `osd blacklist` command in the default permissions for the monitor caps and must be updated.

PREREQUISITES

Before applying the changes described below to a production environment:

1. Plan a maintenance window for the correction procedure.
2. Back up your deployment.
3. Test the correction procedure on your staging environment.

STEPS TO CORRECT

You can apply the changes either using Salt or manually.

To update caps using Salt:

Since the Salt formula currently does not support updating of client permissions, remove old keys and generate new keyrings.

1. Log in to the Salt Master node.
2. Back up `/srv/salt/reclass/classes/cluster/cluster_name/ceph/common.yml`.
3. In `/srv/salt/reclass/classes/cluster/cluster_name/ceph/common.yml`, update the Ceph-related client permissions, for example, for Nova, Glance, Cinder, by replacing

```
mon: "allow r"
```

with

```
mon: 'allow r, allow command \"osd blacklist\"'
```

4. Verify that the changes are applied successfully:

```
cat /srv/salt/reclass/classes/cluster/cluster_name/ceph/common.yml
```

Example of system response:

```
classes:
- system.linux.system.repo.mcp.apt_mirantis.ceph
- cluster.cookied-mcp-ocata-dvr-ceph.infra
parameters:
  ceph:
    common:
      public_network: 172.16.10.0/24
      cluster network: 172.16.10.0/24
      keyring:
        glance:
          name: ${_param:glance_storage_user}
          caps:
            mon: 'allow r, allow command \"osd blacklist\"'
            osd: "profile rbd pool=images"
        cinder:
          name: ${_param:cinder_storage_user}
          caps:
            mon: 'allow r, allow command \"osd blacklist\"'
            osd: "profile rbd pool=volumes, profile rbd-read-only
pool=images"
        nova:
          name: ${_param:nova_storage_user}
          caps:
            mon: 'allow r, allow command \"osd blacklist\"'
            osd: "profile rbd pool=vms, profile rbd-read-only
pool=images"
```

5. Remove the Ceph users from Ceph authentication. For example:

```
salt 'cmn01*' cmd.run 'ceph auth rm client.nova'
salt 'cmn01*' cmd.run 'ceph auth rm client.glance'
salt 'cmn01*' cmd.run 'ceph auth rm client.cinder'
```

6. Remove the keyring files from all Ceph nodes. For example:

```
salt -C "I@ceph:common" cmd.run 'rm /etc/ceph/ceph.client.nova.keyring'
salt -C "I@ceph:common" cmd.run 'rm /etc/ceph/ceph.client.cinder.keyring'
salt -C "I@ceph:common" cmd.run 'rm /etc/ceph/ceph.client.glance.keyring'
```

7. Regenerate keyrings:

```
salt 'cmn01*' state.sls ceph.setup.keyring
salt -C "I@ceph:common" state.sls ceph.setup.keyring
```

8. Refresh pillar data and the minion grains:

```
salt '*' saltutil.refresh_pillar
salt '*' saltutil.refresh_grains
```

9. Regenerate virsh secrets on the compute nodes:

```
salt 'cmp*' state.sls nova
```

10. Restart the services related to Ceph. For example, Nova, Cinder, Glance.

To update caps manually:

1. Log in to the Salt Master node.
2. Back up `/srv/salt/reclass/classes/cluster/<cluster_name>/ceph/common.yml`.
3. In `/srv/salt/reclass/classes/cluster/<cluster_name>/ceph/common.yml`, update the Ceph-related client permissions, for example, for Nova, Glance, Cinder, by replacing

```
mon: "allow r"
```

with

```
mon: 'allow r, allow command \"osd blacklist\"'
```

For example:

```
classes:
- system.linux.system.repo.mcp.apt_mirantis.ceph
- cluster.cookied-mcp-ocata-dvr-ceph.infra

parameters:
...
  ceph:
    common:
      public_network: 172.16.10.0/24
```

```
      cluster network: 172.16.10.0/24
      keyring:
        glance:
          name: ${_param:glance_storage_user}
          caps:
            mon: 'allow r, allow command \"osd blacklist\"'
            osd: "profile rbd pool=images"
        cinder:
          name: ${_param:cinder_storage_user}
          caps:
            mon: 'allow r, allow command \"osd blacklist\"'
            osd: "profile rbd pool=volumes, profile rbd-read-only
pool=images"
        nova:
          name: ${_param:nova_storage_user}
          caps:
            mon: 'allow r, allow command \"osd blacklist\"'
            osd: "profile rbd pool=vms, profile rbd-read-only
pool=images"
```

4. Log in to any `cmn` node.
5. Identify the existing caps for the Ceph client users. For example:

```
ceph auth list | egrep -A 3 'client.cinder|client.glance|client.nova'
```

Example of system response:

```
client.cinder
key: AQAtVzNcGOe5LBAAW437TQA94yBURi7c9Hx8UA==
caps: [mon] allow r
caps: [osd] profile rbd pool=volumes, profile rbd-read-only pool=images,
profile rbd pool=backups
client.glance
key: AQAuVzNcLL6vARAAqUS4Klq0ArDQvYlcMEEd6g==
caps: [mon] allow r
caps: [osd] profile rbd pool=images
client.nova
key: AQDAXTNcDV4QEBAA4/YaL+Yu/n6Ohk48Z3PtzA==
caps: [mon] allow r
caps: [osd] profile rbd pool=vms, profile rbd-read-only pool=images
```

6. Using the output of the command above, update the Ceph `rbd` client users for `[mon]` that have only `allow r` caps and do not modify the existing caps for `[osd]`. Use the following command:

```
ceph auth caps client.<ID> mon 'allow r, allow command "osd blacklist"'
osd '<existing_OSD_caps_for_user>'
```

WARNING: Before you change the permissions manually, double-check that you have explicitly defined all required permissions for the Ceph client users, since the commands above overwrite all existing permissions.

For example:

```
ceph auth caps client.nova mon 'allow r, allow command "osd blacklist"'
osd 'profile rbd pool=vms, profile rbd-read-only pool=images'

ceph auth caps client.glance mon 'allow r, allow command "osd
blacklist"' osd 'profile rbd pool=images'

ceph auth caps client.cinder mon 'allow r, allow command "osd
blacklist"' osd 'profile rbd pool=volumes, profile rbd-read-only
pool=images, profile rbd pool=backups'
```

STEPS TO REVERT

1. In `/srv/salt/reclass/classes/cluster/<cluster_name>/ceph/common.yml`, revert the previously applied changes for the Ceph client users using the backup `ceph/common.yml` file.

2. Revert to the original permissions of the Ceph client users. For example:

```
ceph auth caps client.nova mon 'allow r' osd 'profile rbd pool=vms,
profile rbd-read-only pool=images'

ceph auth caps client.glance mon 'allow r' osd 'profile rbd pool=images'

ceph auth caps client.cinder mon 'allow r' osd 'profile rbd
pool=volumes, profile rbd-read-only pool=images, profile rbd
pool=backups'
```