DigitalOcean

Ceph operations

A DigitalOcean journey

Ceph Days Silicon Valley - 2025/03/25

Alex Marangone - Storage Systems

Contents

- Ceph at DO
- Operations journey
- Community involvement
- Q&A

Quick Stats

65 Cephs

57 Production clusters

8 Staging clusters

250+ PB

Total raw Ceph capacity

12+ PB in our biggest cluster

30,000+

OSDs in the fleet across **1,700+ nodes**

Containers!

- **Initial goal: detach OS release from Ceph version**
  - Removes upgrade headaches
- **No noticeable overhead**
- **Everything still works as if it was non-containerized**

```
#!/bin/bash
TOOLS_CONTAINER="ceph-tools"
if [ -e /usr/bin/ceph ] && [ ! -L /usr/bin/ceph ]; then
    /usr/bin/ceph "$@"
else
    docker exec "$TOOLS_CONTAINER" ceph "$@"
fi
```

```
ExecStartPre=/usr/bin/install -d -m0770 -o ceph -g ceph /var/run/ceph
ExecStartPre=/root/osd_mount_script.sh "%i"
ExecStart=/usr/bin/docker start -a ceph-osd-%i
ExecStop=/usr/bin/docker stop -t 300 ceph-osd-%i
```

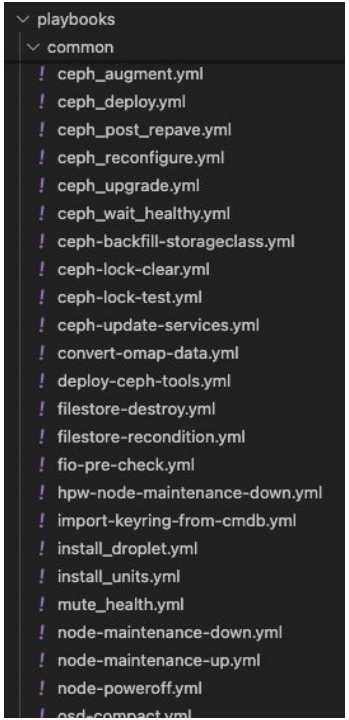

Deployment

- **Build your own vs use upstream? We built our own**
- **Our use cases are limited vs what Ceph has to support**
- **Allows for tight integration with our internal systems**
  - This allows us to deploy and release a cluster to prod within the same day with time to spare
- **Allows for quicker iteration and deployment**
- **Same tooling can be leveraged for daily operations**
- **Our stack is mostly ansible orchestrated by AWX**

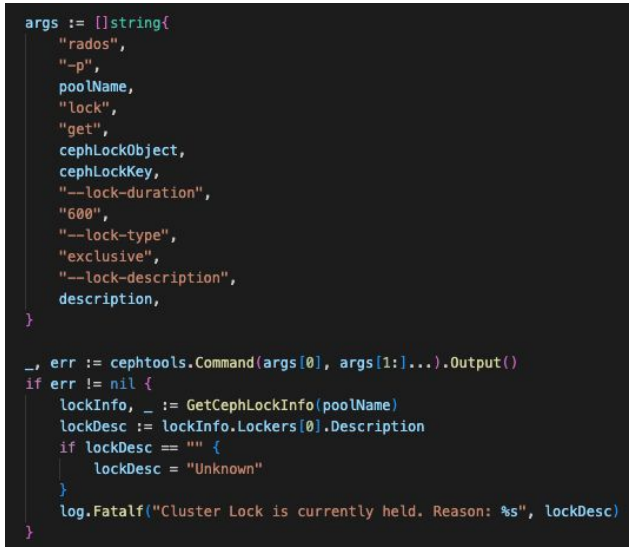

```
args := []string{
    "rados",
    "-p",
    poolName,
    "lock",
    "get",
    cephLockObject,
    cephLockKey,
    "--lock-duration",
    "600",
    "--lock-type",
    "exclusive",
    "--lock-description",
    description,
}

_, err := cephtools.Command(args[0], args[1:]...).Output()
if err != nil {
    lockInfo, _ := GetCephLockInfo(poolName)
    lockDesc := lockInfo.Lockers[0].Description
    if lockDesc == "" {
        lockDesc = "Unknown"
    }
    log.Fatalf("Cluster Lock is currently held. Reason: %s", lockDesc)
}
```

Operations goal: no ssh

- **How do you run any sort of create/update/delete operation safely?**
  - Two persons operating on the same cluster can conflict and create an outage
- **Create automation for all common tasks**
  - Automation check health of the cluster
  - If healthy attempts to grab a rados lock on a uniquely named object
  - Does its thing → release the lock
  - Other tasks will wait for the lock to release before operating

Operations: disks

- **With 30k+ drives proper automation becomes critical**
- **Storman (now a service) handles the LC of every OSD**
  - Wraps ceph-volume
  - Manages firmware updates
  - Does slack alerts
  - Diagnoses failed OSDs
- **Crawl → Walk → Run**
  - *At first* disk failure would require someone to follow a tedious process
  - *Then* disk failure would require someone to SSH and run a couple of commands + ticket creation
  - *Now* disk failure only require human intervention if the disk needs to be physically replaced

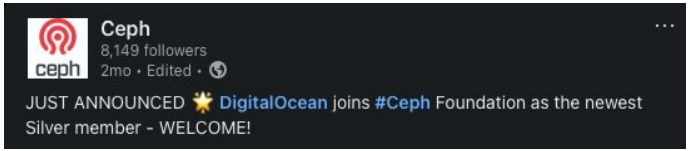

Community

- **We're getting more involved and want to do more**
  - We're BACK in the Ceph foundation
  - Ceph performance meetings
  - Crimson testing

- **On Github**
  - digitalocean/ceph_exporter : prom exporter for ceph
  - digitalocean/pgremapper : Make node operations drama-free
  - We're looking at more


DigitalOcean
Thank You!
Hiring do.co/jobs - amarangone@digitalocean.com
Q&A?