

Ceph storage with Rook

Running Ceph on Kubernetes

Alexander Trost, Rook Maintainer and DevOps Engineer at @Cloudibility

Agenda

- What is Rook?
- Architecture of Rook
- Kubernetes Native Integration
- What can Rook help you do (better) with Ceph?
- Demo of Rook's capabilities
 - Creating a Ceph cluster
 - Showing ease of consuming storage using an example application
 - Adding and removing a new (Ceph) cluster node
- Why Rook?

What is Rook?

What is Rook?

- Cloud-Native Storage Orchestrator
- Extends Kubernetes with custom types and controllers
- Automates deployment, bootstrapping, configuration, provisioning, scaling, upgrading, migration, disaster recovery, monitoring, and resource management

What is Rook?

- Framework for many storage providers and solutions
- Open Source (Apache 2.0)
- Hosted by the Cloud-Native Computing Foundation (CNCF)

Rook Framework for Storage Solutions

- Rook is more than just a collection of Operators and CRDs
- **Framework** for storage providers to integrate their solutions into cloud-native environments
 - Storage resource normalization
 - Operator patterns/plumbing
 - Common policies, specs, logic
 - Testing effort
- Ceph, CockroachDB, Minio, Nexenta, and more...

Architecture of Rook

Architecture



Ceph on Kubernetes with Rook



Ceph on Kubernetes with Rook

```
apiVersion: ceph.rook.io/v1
kind: CephCluster
metadata:
                                                                            (?)
                                                          (\mathbf{Q})
                                                                   \mathbf{Q}
  name: rook-ceph
                                                                           OSD
                                                         MGR
                                                                 MON
                                                                                   OSD
  namespace: rook-ceph
spec:
  cephVersion:
    image: ceph/ceph:v13.2.4-20190109
                                                                          MON
  dataDirHostPath: /var/lib/rook
                                                          Δ
  dashboard:
    enabled: true
  mon:
    count: 3
                                                         MON
    allowMultiplePerNode: true
  storage:
    useAllNodes: true
    useAllDevices: false
```

OSD

OSD

Kubernetes Native Integration

StorageClass

```
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
    name: rook-block
    namespace: rook
provisioner: rook.io/block
parameters:
    pool: replicapool
    clusterName: rook
```

PersistentVolumeClaim

```
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
   name: wp-pv-claim
   labels:
      app: wordpress
spec:
   storageClassName: rook-block
   accessModes:
      - ReadWriteOnce
   resources:
      requests:
      storage: 20Gi
```

Results in..

PersistentVolume 20 Gigabyte

Rook Cluster CRD

```
apiVersion: ceph.rook.io/v1
kind: CephCluster
metadata:
  name: rook-ceph
  namespace: rook-ceph
spec:
  cephVersion:
    image: ceph/ceph:v13.2.4-20190109
  dataDirHostPath: /var/lib/rook
  dashboard:
    enabled: true
  mon:
    count: 3
    allowMultiplePerNode: true
  storage:
    useAllNodes: true
    useAllDevices: false
```

What can Rook help you do (better) with Ceph?

What can Rook help you do (better) with Ceph?

- Health checking for MONs with automatic failover
- Simple management of Ceph Cluster, Pools, Filesystem and RGW through Kubernetes objects.
- Storage Selection in one central place.

Storage Selection

```
storage:
   useAllNodes: true
   useAllDevices: false
   deviceFilter:
   config:
     # storeType: bluestore
     osdsPerDevice: "1"
   directories:
   - path: /rook/storage-dir
   nodes:
```

Storage Selection

nodes:

- name: "172.17.4.101" directories: - path: "/rook/storage-dir" resources: limits: cpu: "500m" memory: "1024Mi" requests: cpu: "500m" memory: "1024Mi" - name: "172.17.4.201" devices: - name: "sdb" - name: "nvme01" config: osdsPerDevice: "5"

Why Rook?

Why Rook?

- Simplifies running a storage backend by orchestrating it
 - Self managing (Health checking)
 - Dynamic provisioning of the storage
- Abstraction (coming) to reduce developer "overhead"
 - Object Store Bucket, Database Custom Resource Definition
 - (Less) vendor lock-in through running services your own

More than Ceph...

More than Ceph...

- Minio Object Storage
- CockroachDB Database
- EdgeFS
- NFS Server

All thanks to the community!



Demo

- Creating a Ceph cluster
- Showing ease of consuming storage using an example application
- Adding and removing a new (Ceph) cluster node

Demo Files: <u>GitHub galexrt/presentation-distributed-storage-with-rook</u>

How to get involved?

- Contribute to Rook
 - o <u>https://github.com/rook/rook</u>
 - o <u>https://rook.io/</u>
- Slack https://rook-io.slack.com/
 - **#conferences**
- Twitter <u>@rook_io</u>
- Forums <u>https://groups.google.com/forum/#!forum/rook-dev</u>
- Community Meetings

Questions?

https://github.com/rook/rook

https://rook.io/

Thank you!

https://github.com/rook/rook

https://rook.io/

The End

Twitter: @galexrt
GitHub: @galexrt
Blog: <u>https://edenmal.moe/</u>
Xing/LinkedIn: Alexander Trost
The Cloud Report: <u>the-report.cloud/?s=Rook</u>