Rook Ceph plugin

Introduce

This plugin describe the deployment with containerized Ceph cluster by Rook operator and provide storage service in edge location node.

As a distributed backend storage, Ceph is widely used on Cloud and Edge solutions, it can provide Block, Filesystem and Object storage service. Traditional deployment for Ceph tools include: ceph-deploy, puppet, ansible etc.

Rook is a storage orchestrator for Kubernetes, which turns distributed storage system into self-managing, self-scaling, self-healing storage services, compared with traditional tools, there are several expectations by using Rook:

- Reduced deployment time for new clusters
- Simplified upgrades
- More agile horizontal scaling
- Better failure tolerance
- Reduced reliance on expert Ceph operators

CSI (Container Storage Interface) is a standard for exposing arbitrary block and file storage storage system to containerized workload on Container Orchestration Systems (COs) like Kubernetes. Also Ceph-CSI is supported by Rook from v1.0.

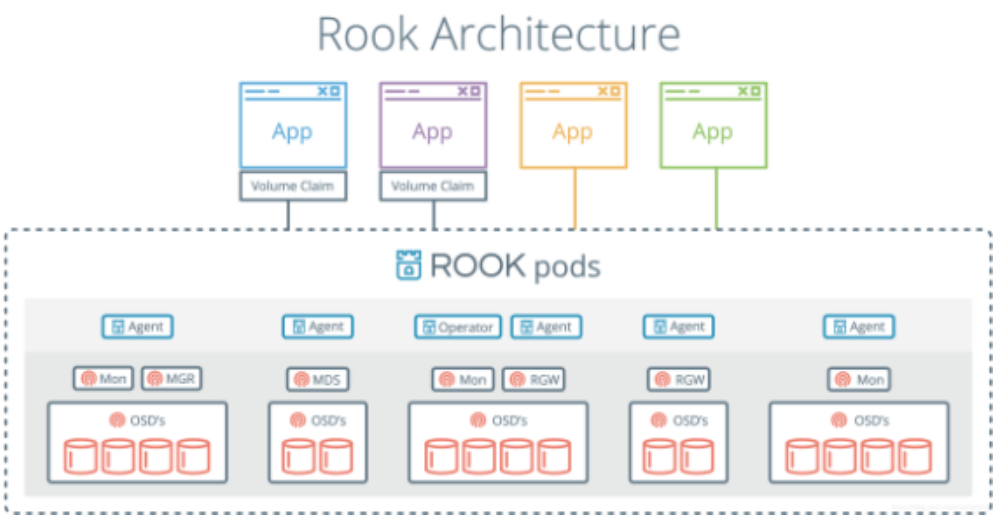

Implement

Rook implementation include two parts, Rook operator and Rook Ceph cluster.

Rook operator include CRDs and also including Ceph-CSI support for volume provisioning.

```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: rook-ceph-operator
  namespace: rook-ceph
labels:
  operator: rook
  storage-backend: ceph
spec:
  selector:
    matchLabels:
      app: rook-ceph-operator
  replicas: 1
  template:
    metadata:
      labels:
        app: rook-ceph-operator
    spec:
      serviceAccountName: rook-ceph-system
      containers:
      - name: rook-ceph-operator
        image: rook/ceph:v1.0.4
        args: ["ceph", "operator"]
        volumeMounts:
        - mountPath: /var/lib/rook
          name: rook-config
        - mountPath: /etc/ceph
          name: default-config-dir
        env:
        - name: ROOK_CURRENT_NAMESPACE_ONLY
          value: "true"
        # CSI enablement
        - name: ROOK_CSI_ENABLE_CEPHFS
          value: "true"
        - name: ROOK_CSI_CEPHFS_IMAGE
          value: "quay.io/cephcsi/cephfsplugin:v1.0.0"
        - name: ROOK_CSI_ENABLE_RBD
          value: "true"
        - name: ROOK_CSI_RBD_IMAGE
          value: "quay.io/cephcsi/rbdplugin:v1.0.0"
        - name: ROOK_CSI_REGISTRAR_IMAGE
          value: "quay.io/k8scsi/csi-node-driver-registrar:v1.0.2"
        - name: ROOK_CSI_PROVISIONER_IMAGE
          value: "quay.io/k8scsi/csi-provisioner:v1.0.1"
        - name: ROOK_CSI_SNAPSHOTTER_IMAGE
          value: "quay.io/k8scsi/csi-snapshotter:v1.0.1"
        - name: ROOK_CSI_ATTACHER_IMAGE
          value: "quay.io/k8scsi/csi-attacher:v1.0.1"
        # The name of the node to pass with the downward API
        - name: NODE_NAME
          valueFrom:
            fieldRef:
              fieldPath: spec.nodeName
        # The pod name to pass with the downward API
        - name: POD_NAME
          valueFrom:
            fieldRef:
              fieldPath: metadata.name
        # The pod namespace to pass with the downward API
        - name: POD_NAMESPACE
          valueFrom:
            fieldRef:
              fieldPath: metadata.namespace
      volumes:
      - name: rook-config
        emptyDir: {}
      - name: default-config-dir
        emptyDir: {}
```

Rook Ceph cluster include settings for Ceph monitor and storage, currently the setting as following:

```
apiVersion: ceph.rook.io/v1
kind: CephCluster
metadata:
name: rook-ceph
namespace: rook-ceph
spec:
cephVersion:
image: ceph/ceph:v13.2.2-20190410
allowUnsupported: false
dataDirHostPath: /var/lib/rook
mon:
count: 3
allowMultiplePerNode: true
dashboard:
enabled: true
network:
hostNetwork: false
rbdMirroring:
workers: 0
annotations:
resources:
storage: # cluster level storage configuration and selection
useAllNodes: true
useAllDevices: false
deviceFilter:
location:
config:
storeType: filestore
metadataDevice: # "md0" specify a non-rotational storage so ceph-volume will use it as block db device of bluestore.
databaseSizeMB: "10240" # uncomment if the disks are smaller than 100 GB
journalSizeMB: "10240" # uncomment if the disks are 20 GB or smaller
directories:
- path: "/var/lib/rook/storage-dir"
```