External Rook Ceph cluster

Speaker

Parth Arora

Software Engineer @IBM Storage

Who Needs Storage for Kubernetes?

- Storage is commonly provided by cloud providers
- What about storage in your datacenter?
- Storage is traditionally not part of the cluster
  - Why should storage be external to K8s?
- Why not manage storage as any other K8s application?

What is Ceph?

- Distributed Software-Defined Storage solution
  - Block (RWO)
  - Shared File System (RWX)
  - Object (S3 Buckets)

- https://ceph.io/

- Open source

Why Ceph?

- Ceph designed to be consistent, not eventually consistent
- Data sharded across partitions (AZs), racks, nodes, disks
- Shard replication is configurable
- Proven highly durable
- Even in extreme disasters, data can be recovered manually

What is Rook?

- Rook is the management layer
  - Similar to cephadm (for non-K8s deployments)
- Operator
  - Automates the deployment of Ceph
- Custom Resource Definitions (CRDs)
  - Allows admin to control the Ceph deployment
  - Allows apps to consume storage like any other K8s storage
  - Storage Classes, Persistent Volume Claims

Architecture Layers

- Rook
  - Operator owns the management of Ceph
- Ceph-CSI
  - Driver dynamically provisioning Ceph volumes and attaching them to user workloads
- Ceph
  - Data layer

Rook Modes

- Converged clusters, is recommended if you just have a single cluster to manage the storage locally

- External Clusters

What is an External Cluster?

Ceph is running in an "external" cluster from the clients

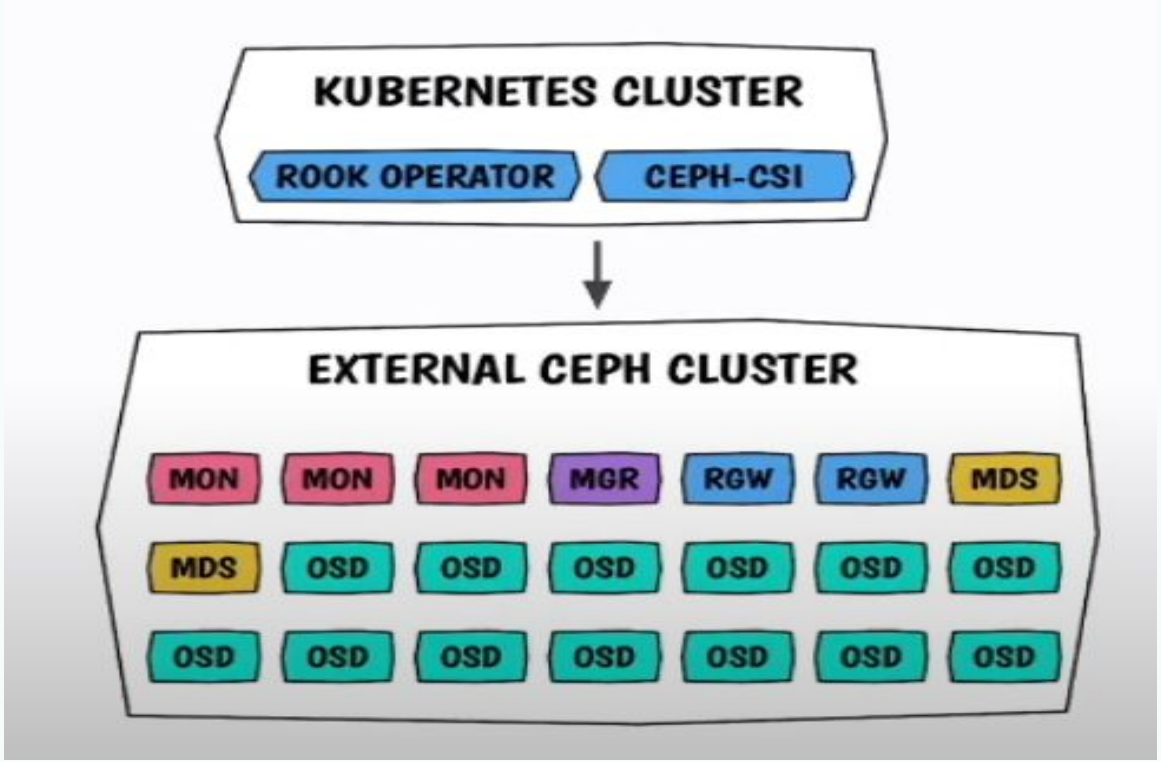

Why External Cluster?

Centralized Ceph management in a single cluster with multiple K8s clusters

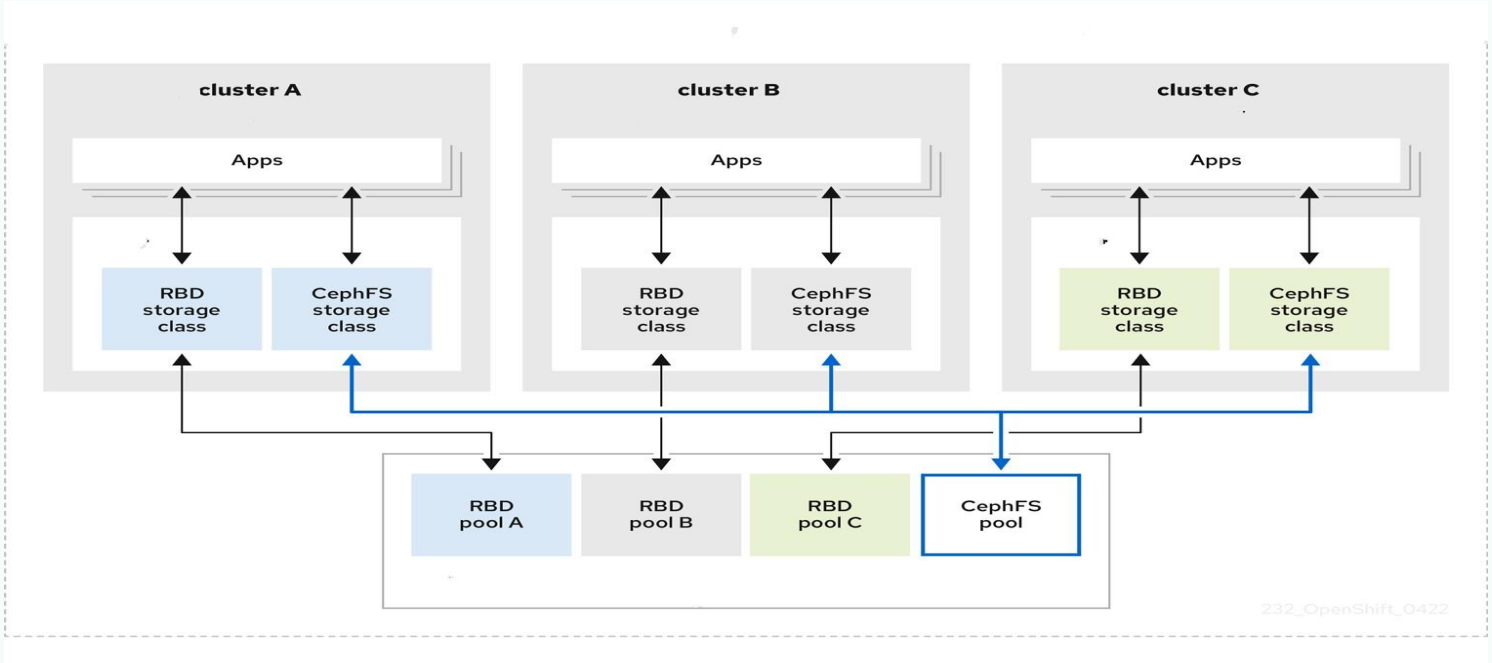

Why External Cluster?

- Already deployed Ceph cluster running not in a K8s environment

- Fully independent storage for another level of isolation from their K8 compute nodes

Dynamically create Block/File/object storage consumable by k8s applications

How to install External cluster?

1. Collect info on the producer cluster
2. Import info to the consumer cluster

[Installation Guide]

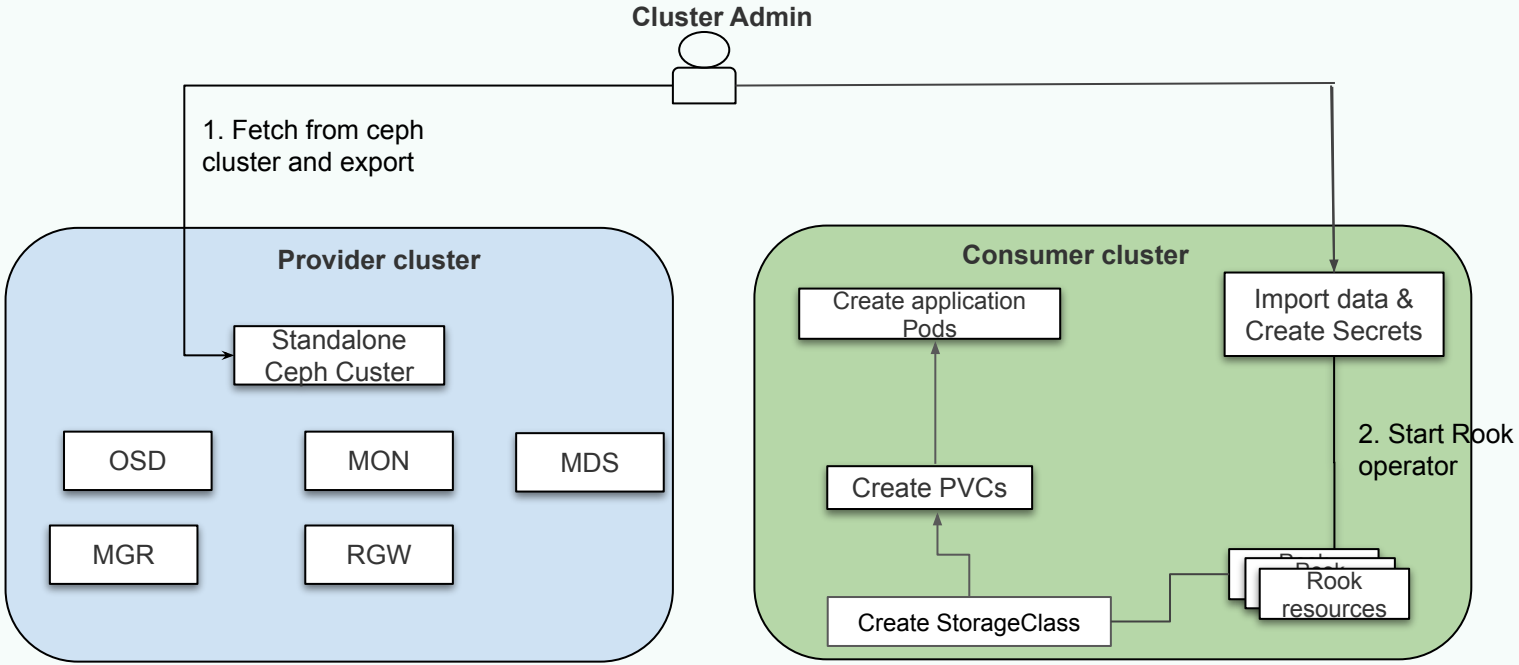

How to export from external cluster?

On Standalone Ceph Cluster:

- **Run the python script create-external-cluster-resources.py for creating all users and keys.**
  `CMD: # python3 ceph-external-cluster-details-exporter.py --rbd-data-pool-name ceph-rbd --rgw-endpoint xxx.xxx.xxx.xxx:xxxx`

- **Export the output**

How to import external cluster?

On the K8s consumer cluster

- Run the import script
  CMD: `. import-external-cluster.sh`

- Deploy manifests create common.yaml, crds.yaml, operator.yaml, common-external.yaml and cluster-external.yaml

Verify Connection

- Verify the consumer cluster is connected to the source ceph cluster:

```
$ kubectl -n rook-ceph-external get CephCluster

NAME                 DATADIRHOSTPATH   MONCOUNT   AGE    STATE
HEALTH
rook-ceph-external   /var/lib/rook                162m   Connected
HEALTH_OK
```

- Verify pool, StorageClass
- Create respective pvc from storageclass and application pods which uses the pvc for storage

Demo

Demo Link

New Features

- Rados namespaces
- subVolumeGroup
- IPV6 endpoints
- V2 mon port
- RGW Multisite

Future Works!

- Add replica-1 support
- Add support for topology awareness
- Make the connection steps easier
- Improve documentation

Community

Website	https://rook.io/
Documentation	https://rook.io/docs/rook/v1.13/
Slack	https://rook-io.slack.com/
Github	https://github.com/rook/rook
Twitter	@rook_io

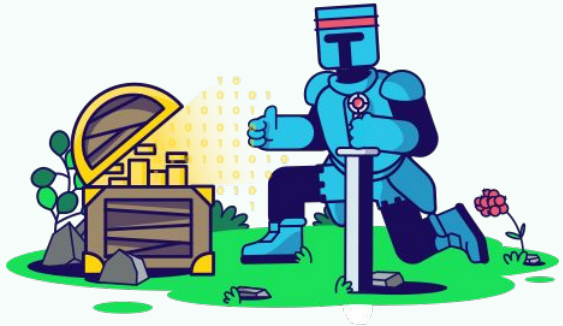

Thanks

Q&A