

RED HAT  
**SUMMIT**

# Rook: Automating Ceph for Kubernetes via Operators

Travis Nielsen  
Rook Maintainer  
May 7, 2019

# Platform Revolution: Kubernetes

- Kubernetes is the new platform for applications
- Based on containers
- <https://kubernetes.io/>

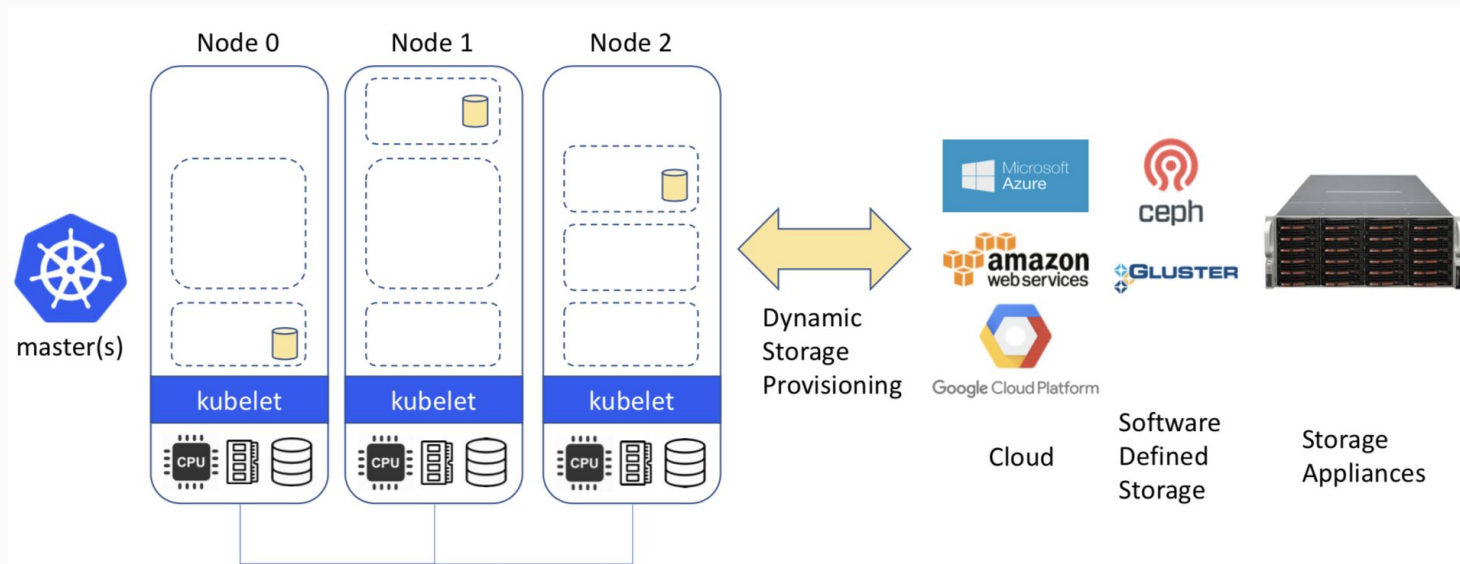
# Platform Revolution: OpenShift

- Red Hat® OpenShift® is a comprehensive enterprise-grade application platform, built for containers with Kubernetes.
- OpenShift is Red Hat's distro of Kubernetes
- <https://www.openshift.com/>



# Storage for Kubernetes

- Storage is traditionally external
- Volume plugins allow storage solutions to provide storage to your apps




# Traditional Storage Limitations

- Not portable: requires these services to be accessible
- Deployment burden of external solutions
- Vendor lock-in due to using provider managed services

# Storage ON Kubernetes

- Kubernetes can manage our storage solution
- Highly portable applications (including storage dependencies)
- Dedicated K8s storage cluster also possible



# What is Rook?

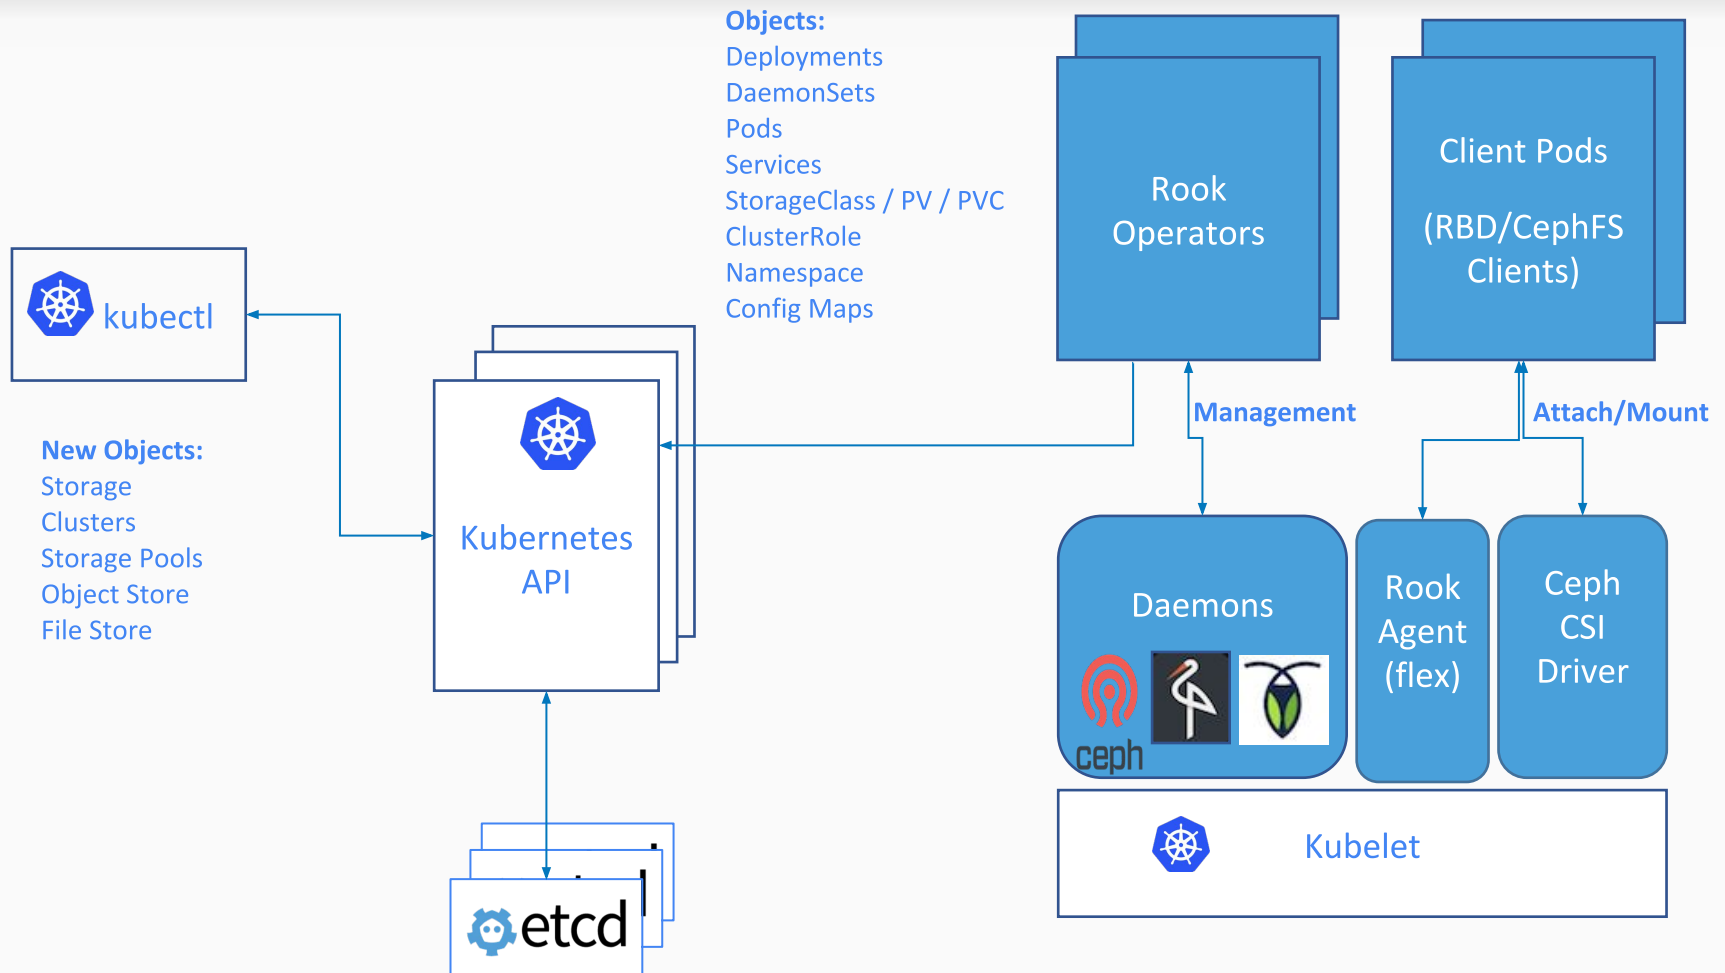
- <https://rook.io>
- Storage Provider for Kubernetes
- Uses Kubernetes patterns (custom types and controllers)
- Automates deployment, bootstrapping, configuration, provisioning, scaling, upgrading, migration, disaster recovery, monitoring, and resource management
- Open Source (Apache 2.0)
- Hosted by the Cloud-Native Computing Foundation (CNCF)
- Framework for several storage providers
  - Ceph: <https://www.ceph.com>

# Rook Project

- Upstream
  - v1.0 was just released!
  - Operator Hub integration in progress  
<https://www.operatorhub.io/>
- Downstream
  - Integration in progress: RHHI Next and OCS 4

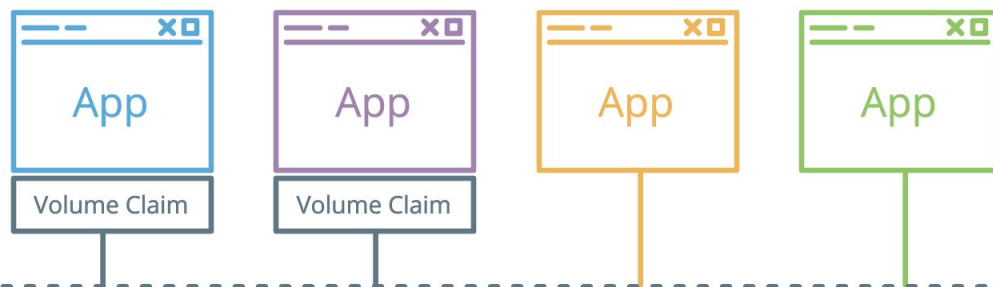


# Rook Architecture

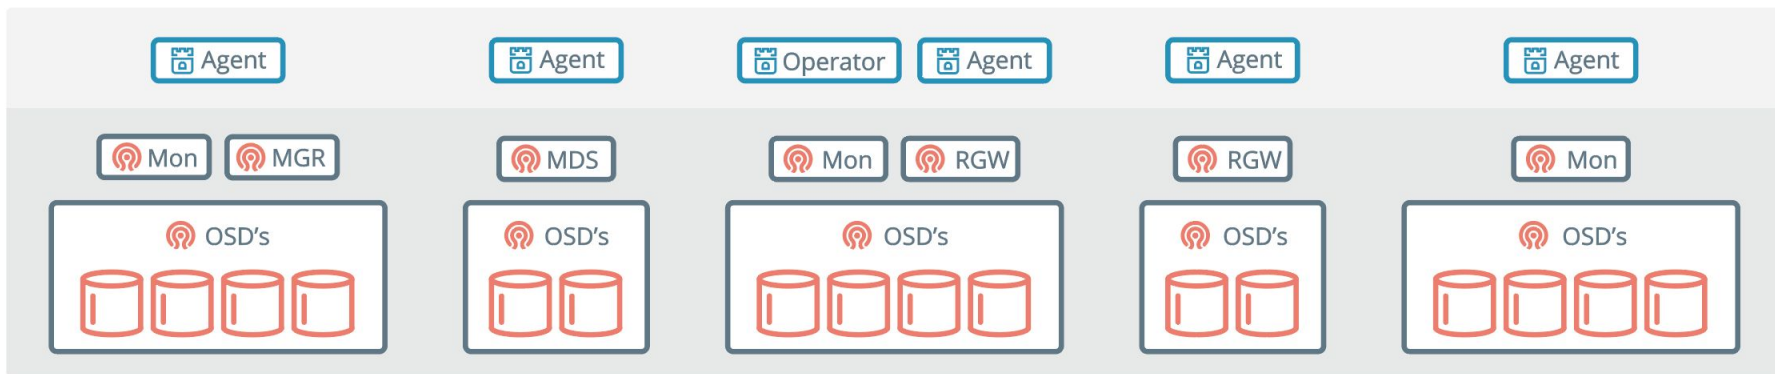


# Ceph on Kubernetes with Rook

## Rook Architecture



### ROOK pods



# Operator Pattern

- OpenShift strategy: Operators are the future of automation
- Codifies domain expertise to deploy and manage an application
  - Automates actions a human would normally do
- Apply user's desired state
  - Observe - discover current actual state of cluster
  - Analyze - determine differences from desired state
  - Act - perform operations to drive actual towards desired

# Custom Resource Definitions (CRDs)

- Custom Resource Definition (CRDs) are arbitrary types that extend the Kubernetes API
  - look just like any other built-in object (e.g. Pod)
  - Enabled native `oc/kubectl` experience
- A means for user to describe their desired state
- Standard yaml files

```
apiVersion: ceph.rook.io/v1
```

```
kind: CephCluster
```

```
metadata:
```

```
  name: my-cluster
```

```
spec:
```

```
  dashboard:
```

```
    enabled: true
```

```
  network:
```

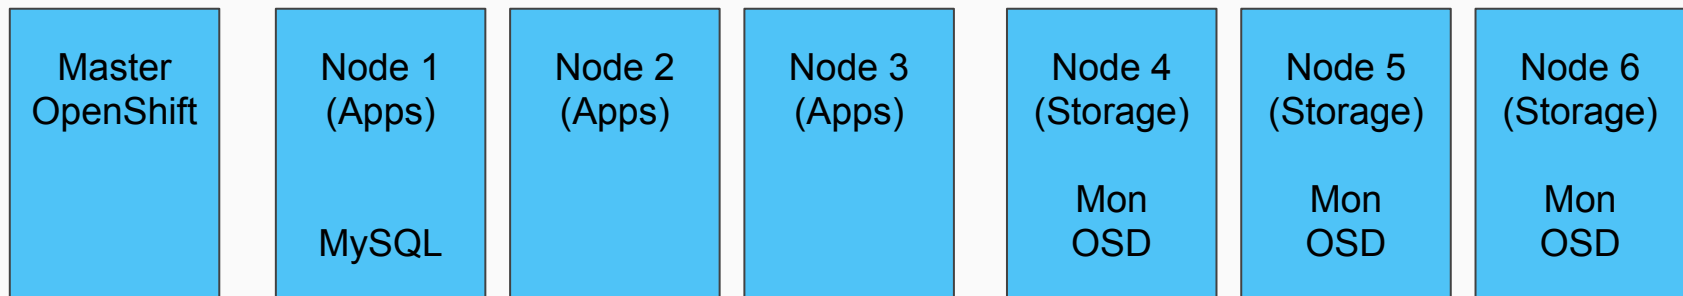
```
    hostNetwork: true
```

# Rook Ceph Operator

- The Operator leverages the full power of K8S
  - Services, ReplicaSets, DaemonSets, Secrets, ...
- Contains all the logic to manage storage systems at scale
  - Handle stateful upgrades
  - Handle rebalancing the cluster
  - Handle health and monitoring tasks
- Not on the data path – can be offline for minutes

# DEMO: Rook Walkthrough

- Start a Rook Ceph cluster on OpenShift



# CephCluster CRD

- Ceph CRDs are v1
- Host path
  - Required for persistence of the data
- Dashboard
  - Web UI to view and manage the Ceph cluster
- Network

```
apiVersion: ceph.rook.io/v1
kind: CephCluster
metadata:
  name: my-cluster
spec:
  dataDirHostPath: /var/lib/rook
  cephVersion:
    image: ceph/ceph:v13.2.2-20181023
  dashboard:
    enabled: true
  network:
    hostNetwork: true
  ...
```

# How to get involved?

- <https://rook.io/>
- Try out the Rook v1.0 release!
- Contribute to Rook: <https://github.com/rook/rook>
- Slack - <https://rook-io.slack.com/>
- Twitter - @rook\_io
- Forums - <https://groups.google.com/forum/#!forum/rook-dev>
- Community Meetings



**RED HAT  
SUMMIT**

THANK YOU



[linkedin.com/company/Red-Hat](https://www.linkedin.com/company/Red-Hat)



[facebook.com/RedHatinc](https://www.facebook.com/RedHatinc)



[youtube.com/user/RedHatVideos](https://www.youtube.com/user/RedHatVideos)



[twitter.com/RedHat](https://twitter.com/RedHat)