

🏰 ROOK 车

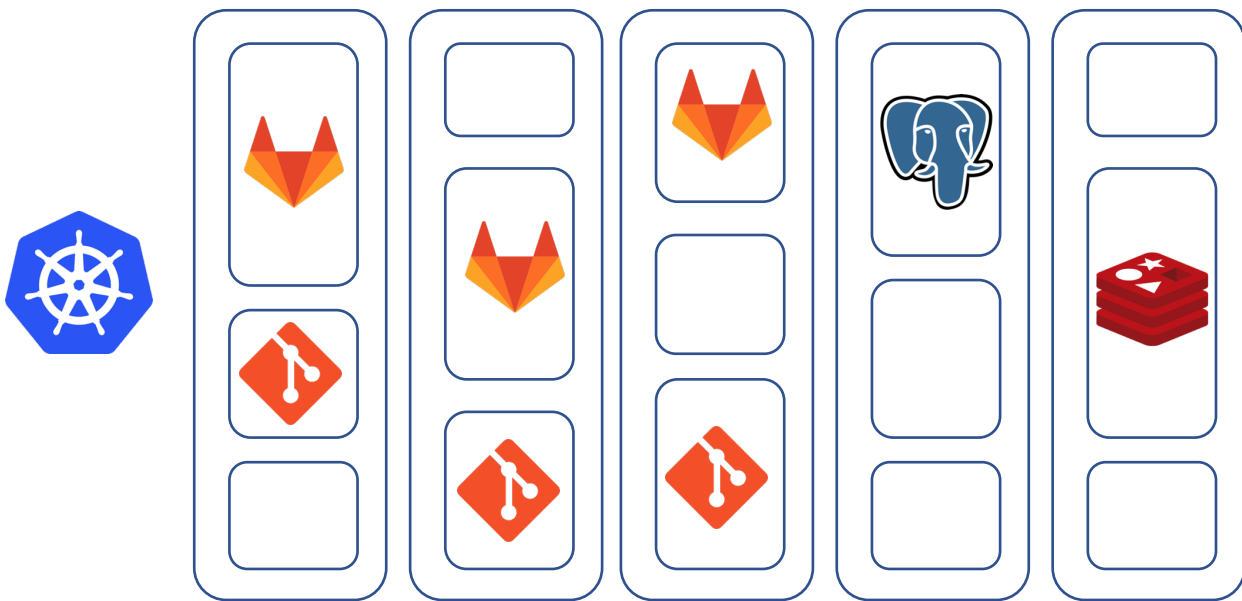
Making Ceph Awesome on Kubernetes
让 Ceph 在 Kubernetes 上光芒闪耀

Bassam Tabbara
Founder & CEO
Upbound, Inc.



kubernetes

Kubernetes – Cluster Operating System



- Container Packaging
- Resource Scheduling
- LifeCycle Management
- Horizontal Scaling
- Automated Rollout and Rollback
- Service Discovery
- Load Balancing
- Declarative Management


Kubernetes is everywhere



Google Cloud Platform





Kubernetes and Persistent Storage





Ceph Cluster

Ansible, Systemd, Ceph-Deploy, SSH, SaltStack, Scripts, etc.

Why not run Ceph on Kubernetes?



Rook makes Ceph on Kubernetes Awesome




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
- Open Source (Apache 2.0)
- Hosted by the Cloud-Native Compute Foundation (CNCF)


RADOS on Kubernetes with Rook

```
apiVersion: rook.io/v1alpha1
kind: Cluster
metadata:
  name: rook
  namespace: rook
spec:
  backend: ceph
  dataDirHostPath: /var/lib/rook
  hostNetwork: false
  storage:
    useAllNodes: true
    useAllDevices: false
    deviceFilter:
    metadataDevice:
    location:
    storeConfig:
      storeType: bluestore
```



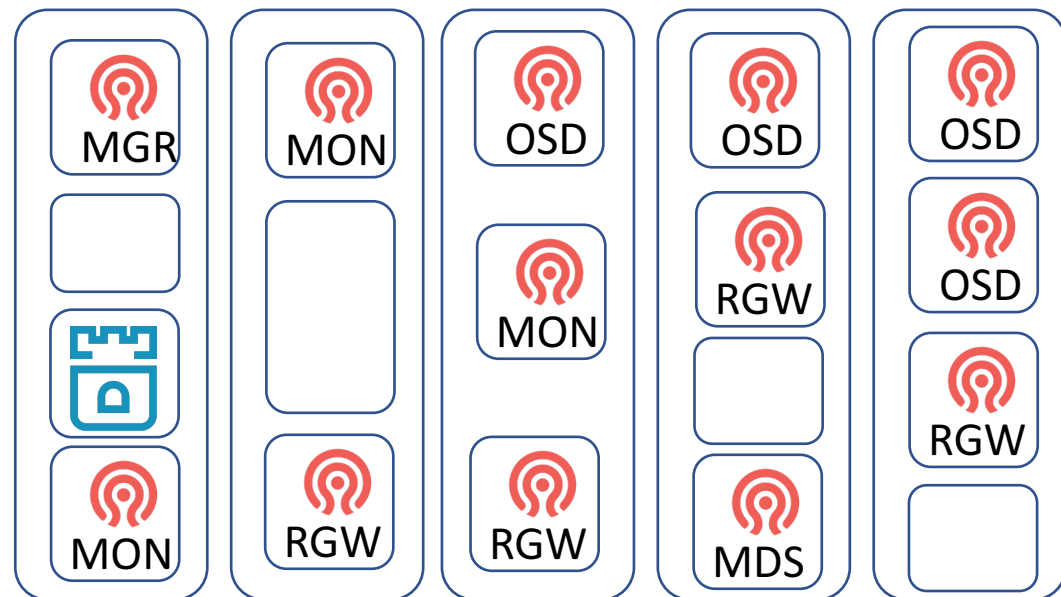
CephFS on Kubernetes with Rook

```
apiVersion: rook.io/v1alpha1
kind: Filesystem
metadata:
  name: myfs
  namespace: rook
spec:
  metadataPool:
    replicated:
      size: 1
  dataPools:
    - failureDomain: osd
      replicated:
        size: 1
  metadataServer:
    activeCount: 1
    activeStandby: true
```




RGW on Kubernetes with Rook

```
apiVersion: rook.io/v1alpha1
kind: ObjectStore
metadata:
  name: my-store
  namespace: rook
spec:
  gateway:
    type: s3
    sslCertificateRef:
    port: 80
    securePort: 443
    instances: 1
    allNodes: false
```

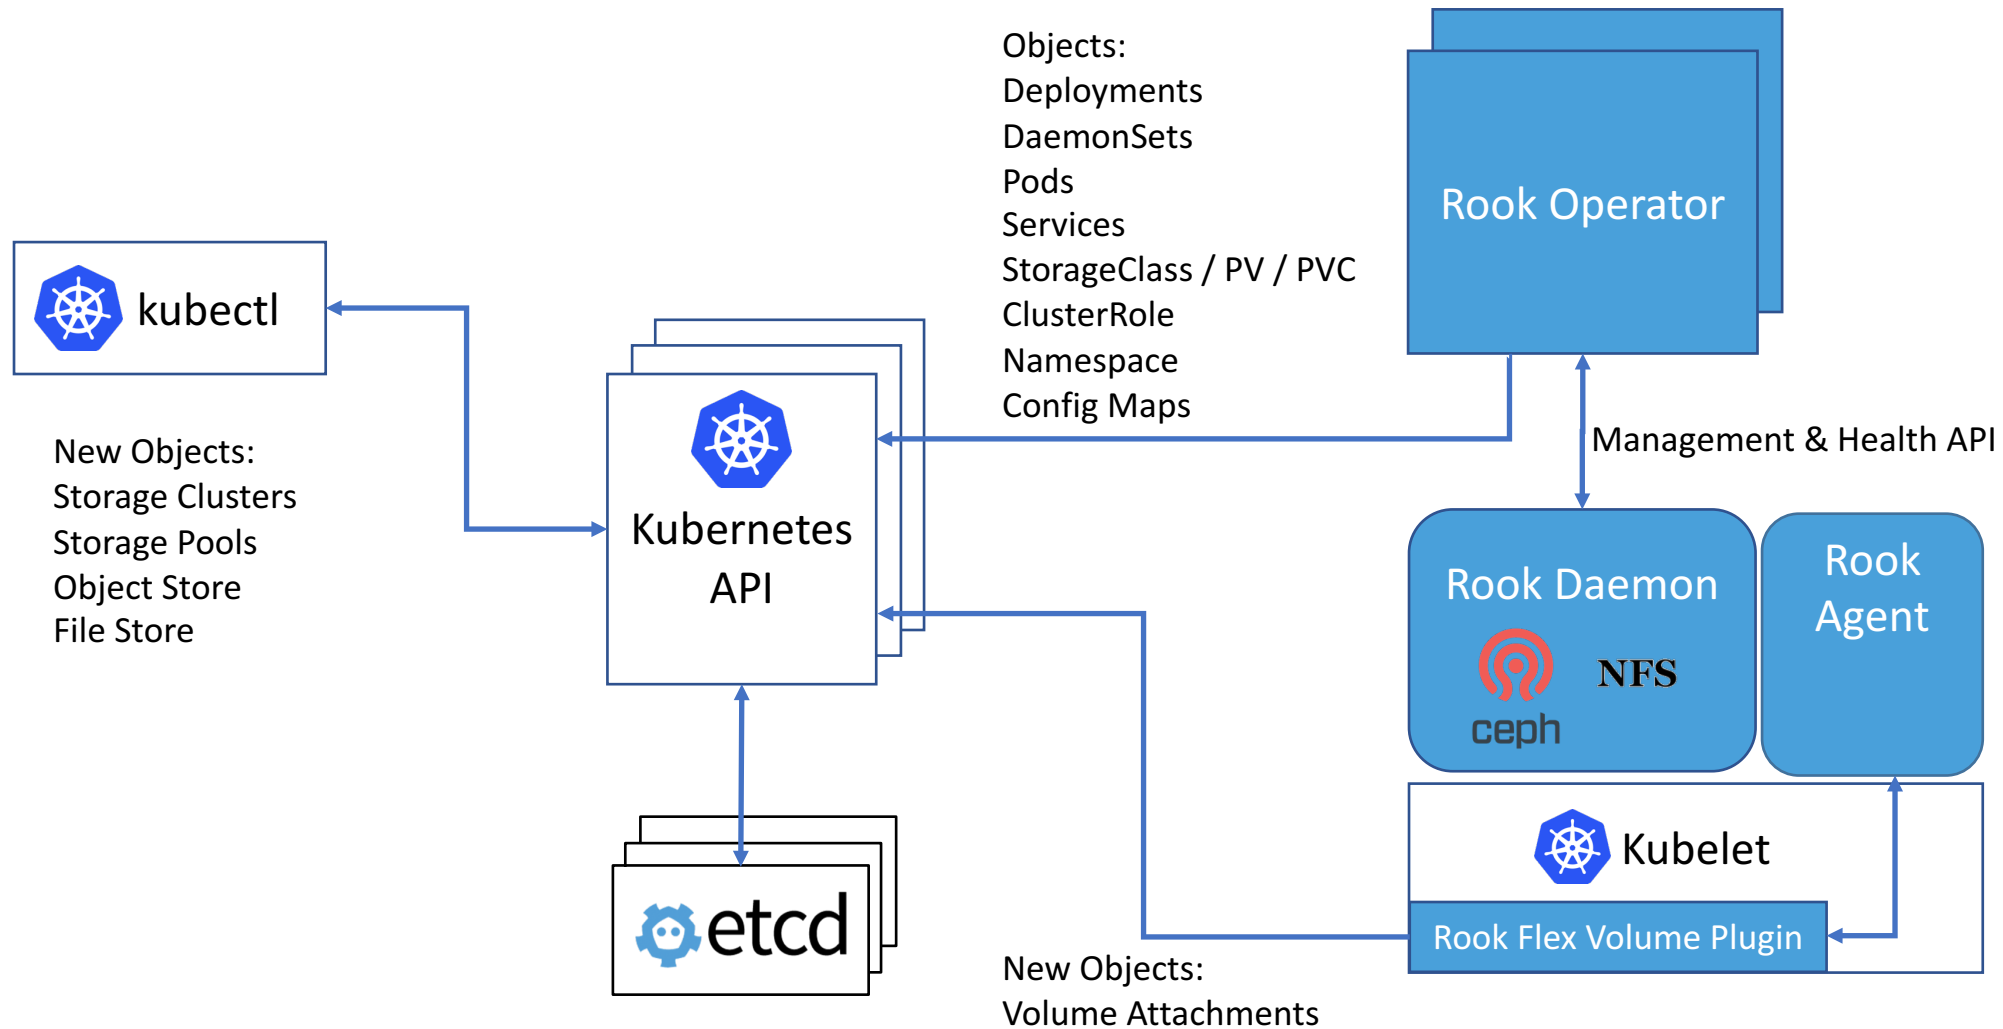


Rook Operator

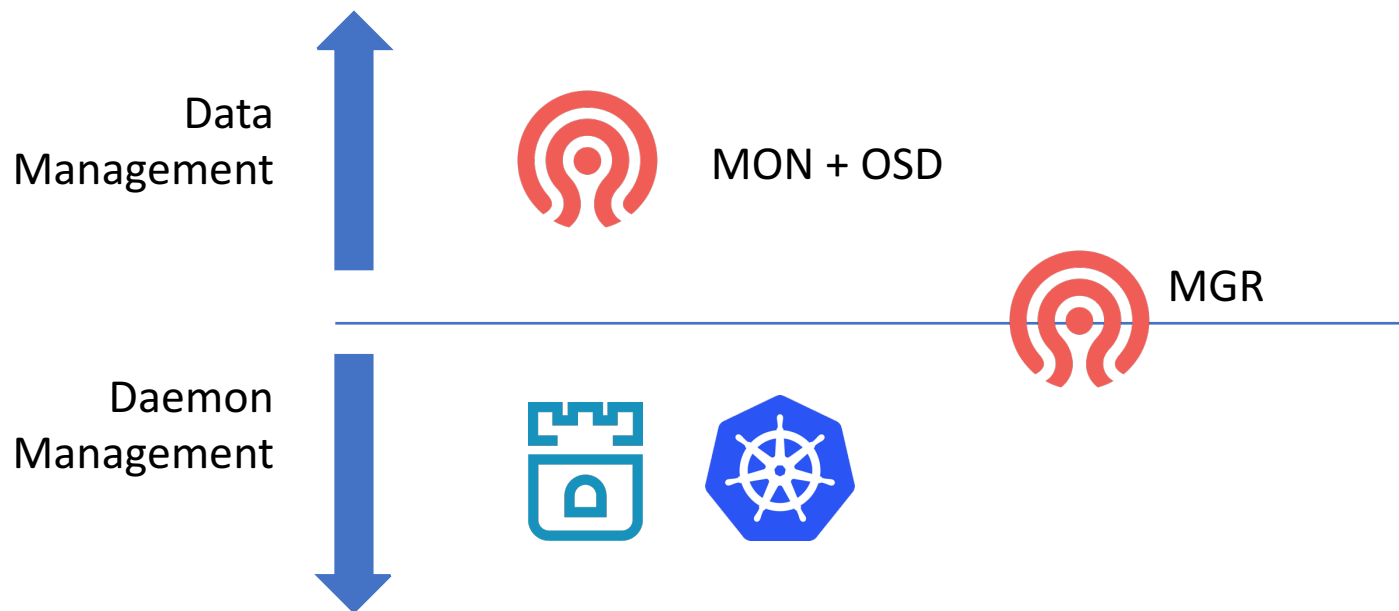
- Defines *desired state* for the storage cluster
 - Storage Cluster, Pool, Object Store, etc.
- The Operator runs reconciliation loops
 - Watches for changes in desired state
 - Watches for changes in the cluster
 - Applies changes to the cluster to make it match desired
- The Operator leverages the full power of K8S
 - Services, ReplicaSets, DaemonSets, Secrets, ...
 - Contains all the logic to manage Ceph at scale
 - Handles stateful upgrades
 - Handles rebalancing the cluster
- Not on the data path – can be offline for minutes



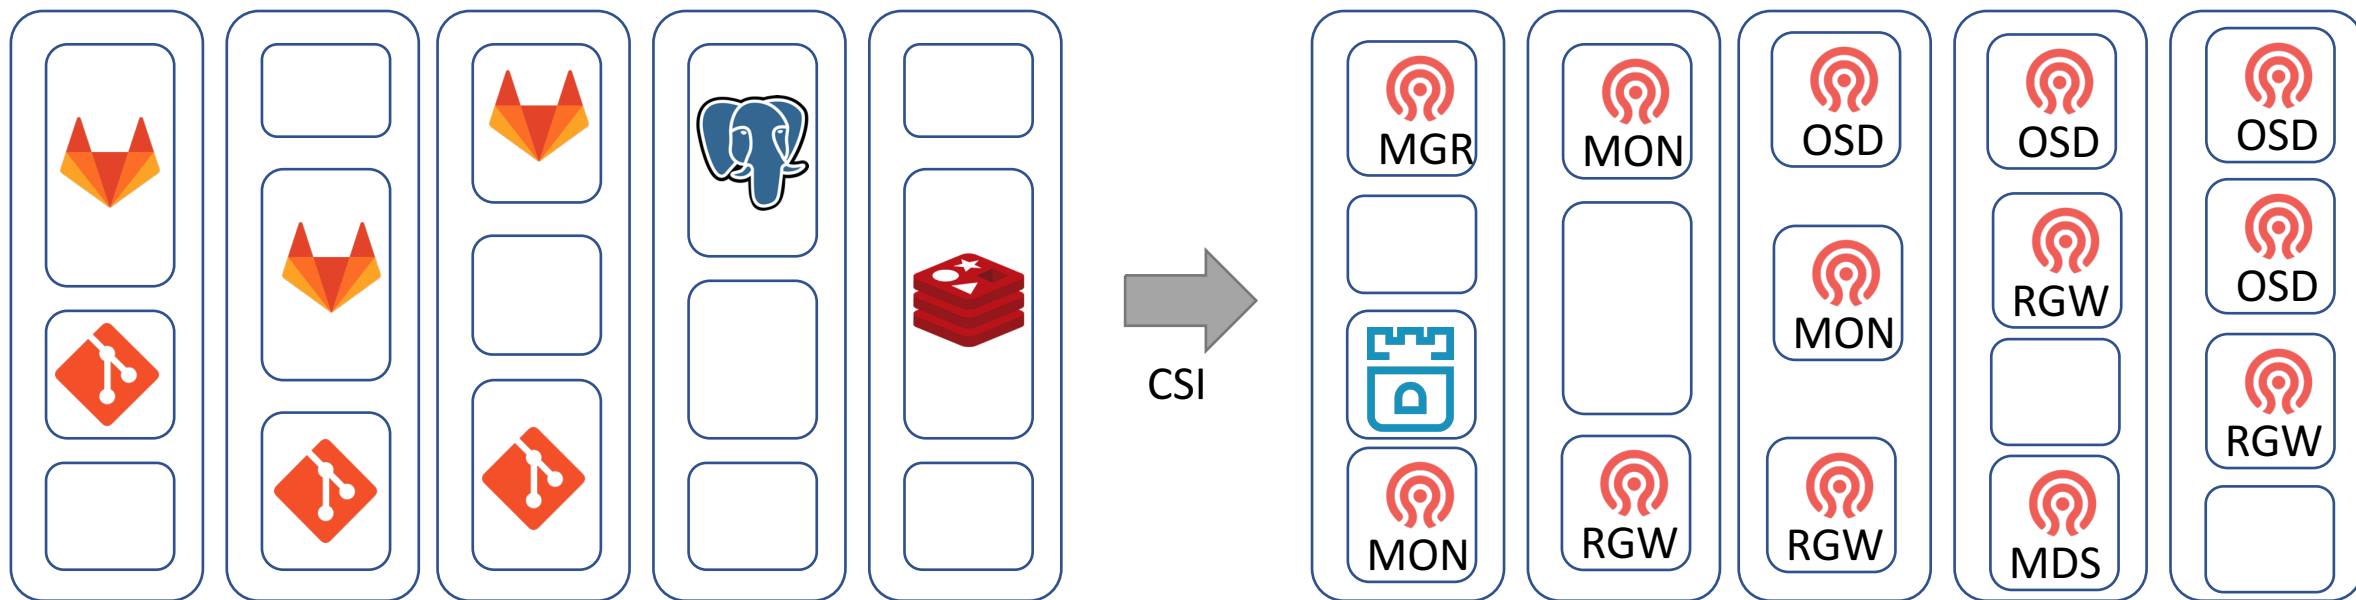
Rook Architecture



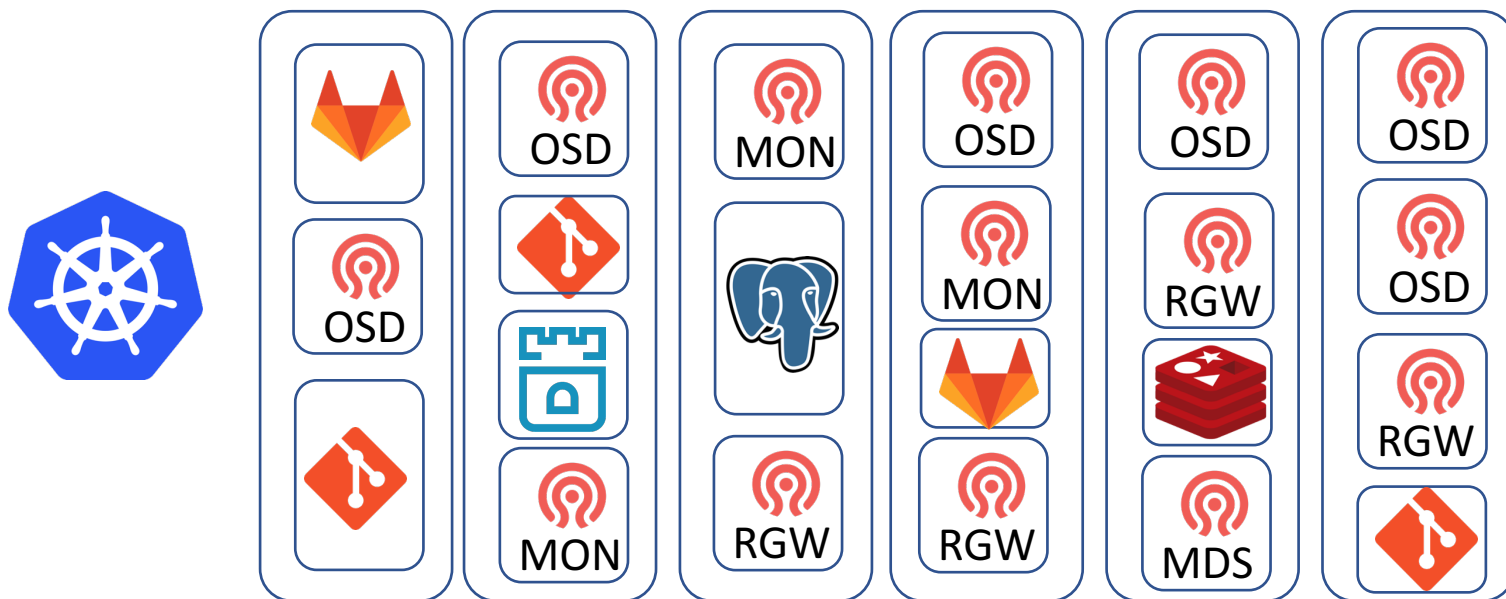
Management Division of Labor




Kubernetes with Dedicated Storage Cluster



Kubernetes with Converged Storage



Always deploy Ceph on Kubernetes



How to get involved?

- Contribute to Rook
 - <https://github.com/rook/rook>
 - <https://rook.io>
- Slack - <https://rook-slackin.herokuapp.com/>
- Twitter - @rook_io
- Forums - <https://groups.google.com/forum/#!forum/rook-dev>
- Community Meetings



Ceph中国社区

IT大咖说
知识共享平台

Thank you!
谢谢