

Keep Kubernetes Caffeinated

Erin A. Boyd

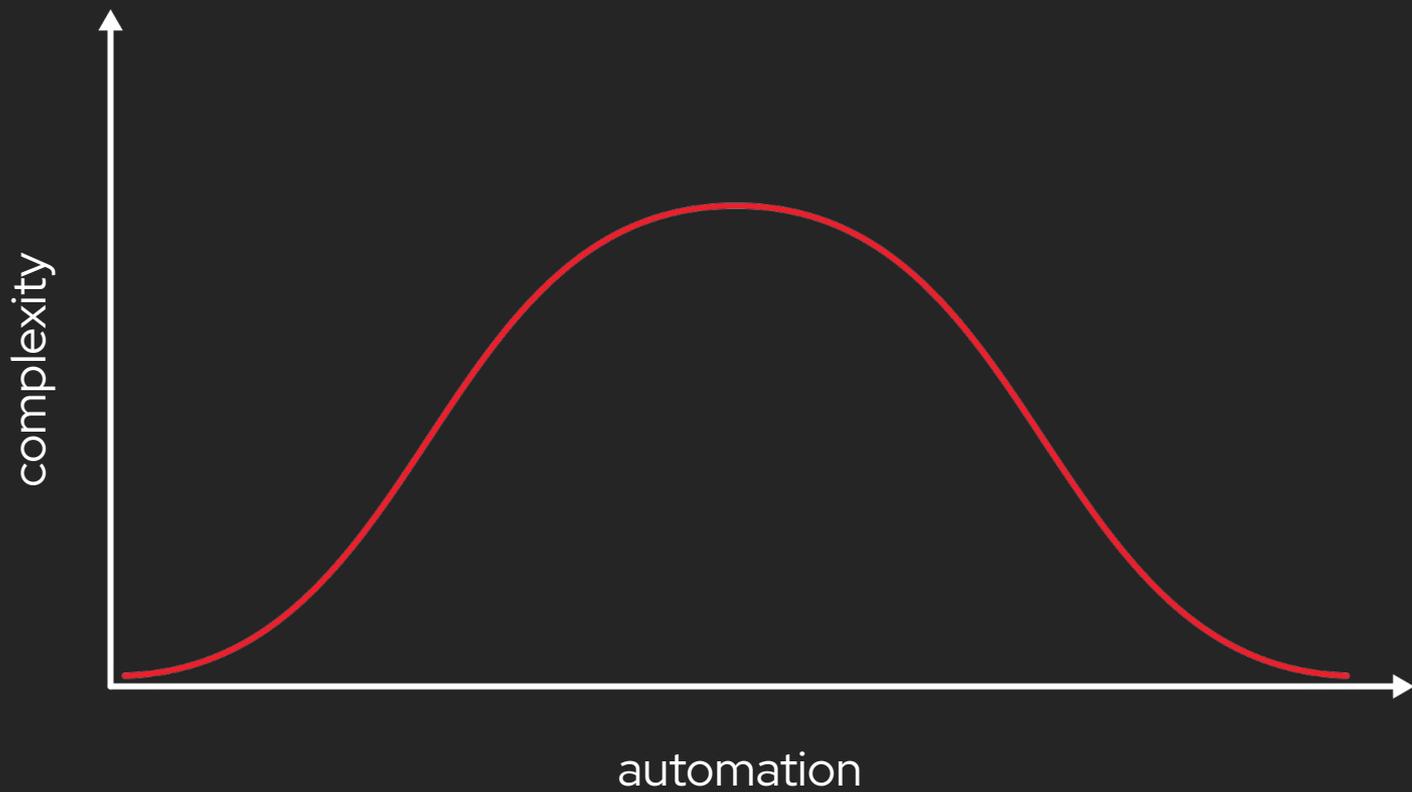
Senior Principal Software Engineer, Red Hat



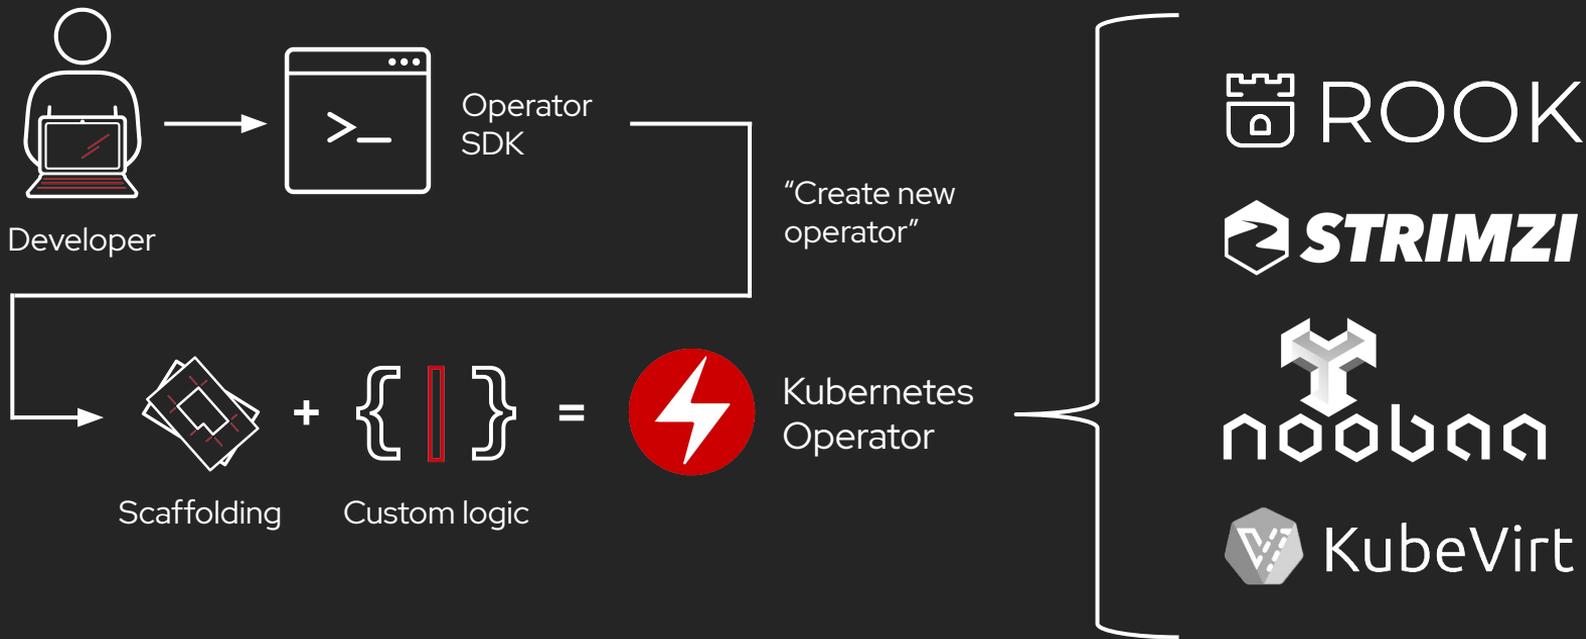
Coffee is in our **DNA**

Coffee is in our DNA

But it's a **process**



Simplicity is *key*



Single Button Automation

Foam

Milk

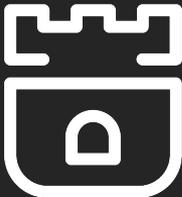
Espresso



Upgrade

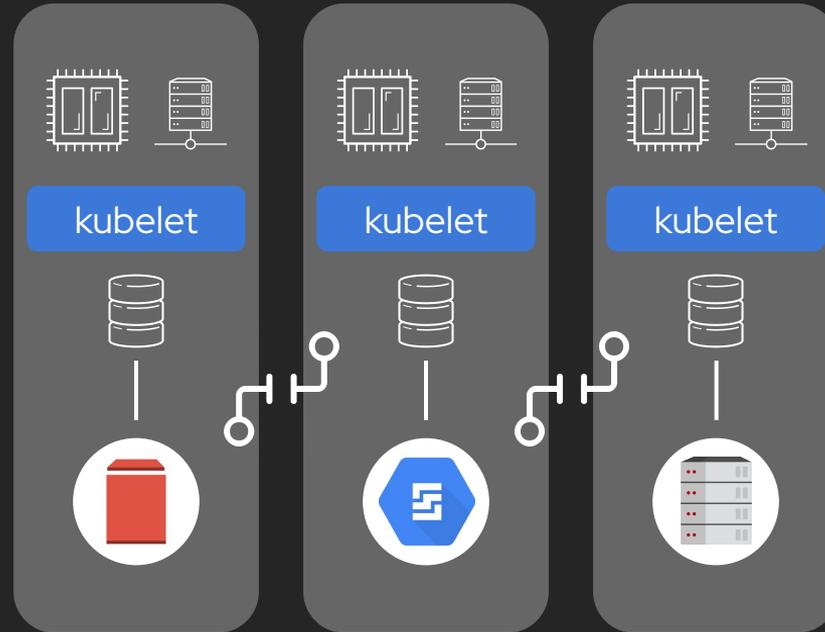
Manage

Deploy

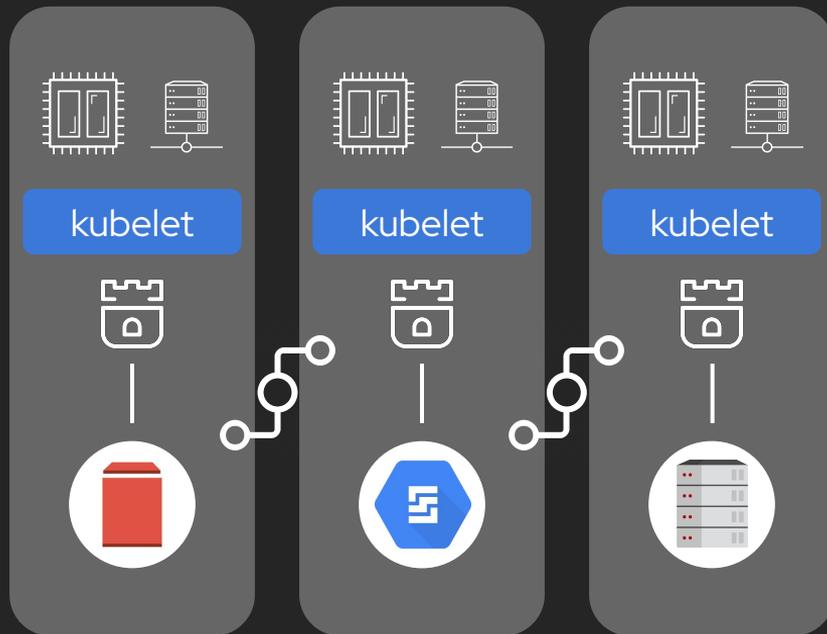
The icon is a white outline of a chess rook, which is a castle tower. It has a square base with a small archway in the center, and a crenellated top.

ROOK

Disconnected storage in Kubernetes



Universal backplane for storage



01

Select

"What storage should I use?"

02

Deploy

"How do I deploy and manage my storage?"

03

Integrate

"How do I make my deployed storage available in my cluster?"

04

Consume

"How does my stateful application provision and use available storage?"

01

Select

“What storage should I use?”

02

Deploy

“How do I deploy and manage my storage?”

03

Integrate

“How do I make my deployed storage available in my cluster?”

04

Consume

“How does my stateful application provision and use available storage?”

05

Scale

“How do I scale and manage my storage from cluster to cluster?”

05

Scale

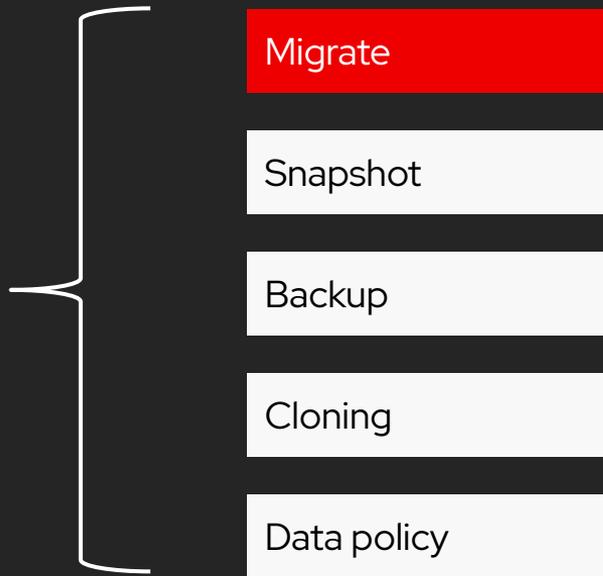
"How do I scale and manage
my storage from cluster to
cluster?"

==

Hard

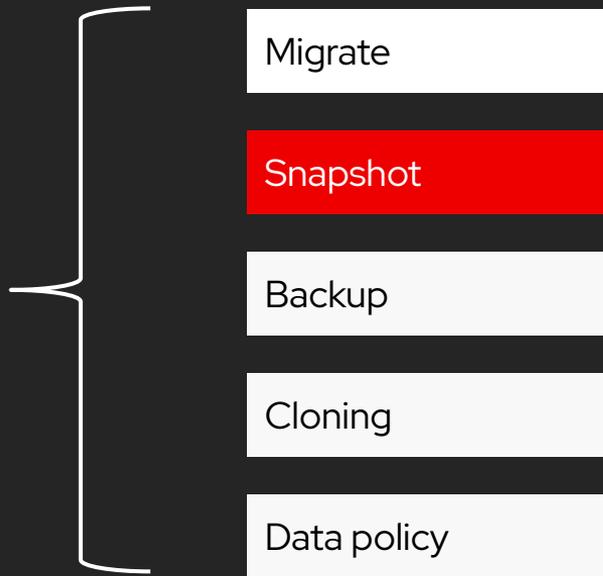
05 Scale

"How do I scale and manage my storage from cluster to cluster?"



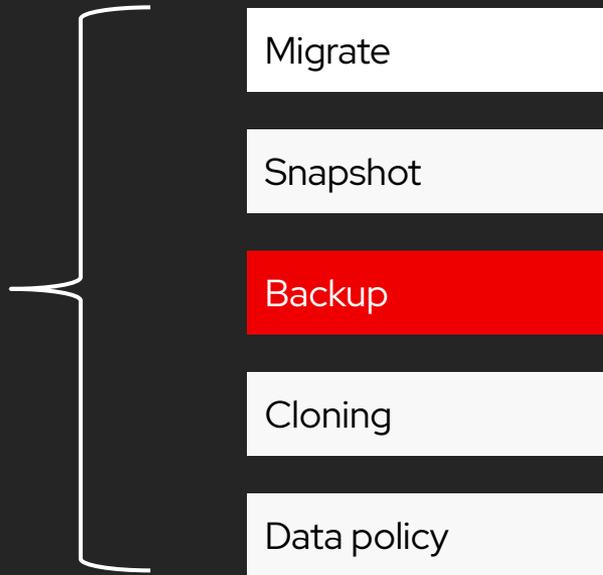
05 Scale

"How do I scale and manage my storage from cluster to cluster?"



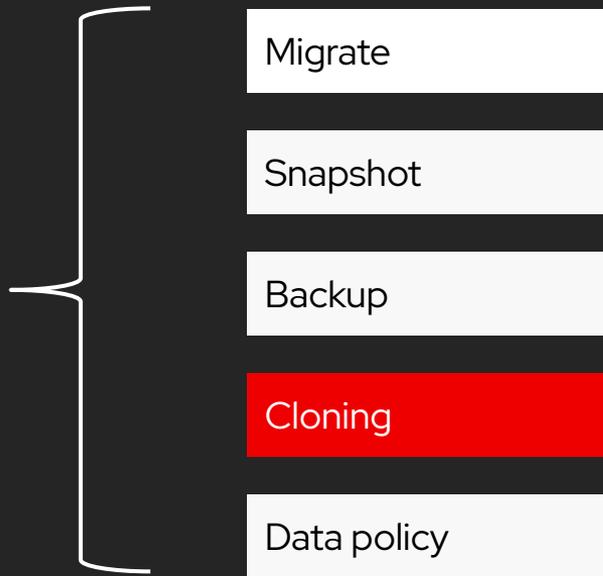
05 Scale

"How do I scale and manage
my storage from cluster to
cluster?"



05 Scale

"How do I scale and manage my storage from cluster to cluster?"



05 Scale

"How do I scale and manage my storage from cluster to cluster?"

Migrate

Snapshot

Backup

Cloning

Data policy

ROOK

Deploy

Integrate

Consume



Scale

Migrate

Policy

Storage is **critical**
to hybrid cloud

Thank you

Red Hat is the world's leading provider of enterprise open source software solutions.

Award-winning support, training, and consulting services make Red Hat a trusted adviser to the Fortune 500.



@erinaboyn



erinaboyn / erinboyn

Thank you



Twitter: @erinaboyd
Github: erinaboyd / erinboyd
Etc...

OperatorHub.io | The registry

operatorhub.io/?sort=descending

OperatorHub.io

Search OperatorHub... Contribute

Welcome to OperatorHub.io

OperatorHub.io is a new home for the Kubernetes community to share Operators. Find an existing Operator or list your own today.

78 ITEMS

VIEW SORT Z-A

CATEGORIES

- AI/Machine Learning
- Application Runtime
- Big Data
- Cloud Provider
- Database
- Developer Tools
- Integration & Delivery
- Logging & Tracing
- Monitoring
- Networking
- OpenShift Optional
- Security
- Storage
- Streaming & Messaging

PROVIDER

- Altinity (1)
- Amazon Web Services (1)
- Appody (1)
- Aqua Security (1)
- AtlasMap (1)
- Show 51 more

CAPABILITY LEVEL

- Basic Install (38)
- Seamless Upgrades (13)
- Full Lifecycle (16)



WildFly
provided by WildFly

Operator that creates and manages Java applications running on WildFly.



Vault Operator
provided by Banzai Cloud

Creates and maintain highly-available Vault clusters on Kubernetes



Twistlock Console Operator
provided by Twistlock

Deploy Twistlock cloud native security in Kubernetes.



Turbonomic Platform Operator
provided by Turbonomic, Inc.

Turbonomic Workload Automation for Multicloud



TiDB Operator
provided by PingCAP

TiDB is a popular NewSQL MySQL-compatible database. The TiDB Operator manages



Sysdig Agent Operator
provided by Sysdig

Sysdig is a unified platform for container and microservices monitoring, troubleshooting



Synopsys Operator
provided by Synopsys, Inc

Synopsys Operator is a cloud-native administration utility for Synopsys software.



Strimzi Apache Kafka Operator
provided by Red Hat

Strimzi provides a way to run an Apache Kafka cluster on



StorageOS
provided by StorageOS, Inc

Cloud-native, persistent storage for containers.



Spinnaker Operator
provided by PpsMx

Spinnaker is an Open Source, multi-cloud Continuous delivery platform to perform



Spark Operator
provided by GoogleCloudPlatform

Apache Spark is a unified analytics engine for large



Sematext Operator
provided by Sematext

Full stack observability with Sematext Monitoring Platform



Seldon Operator
provided by Seldon Technologies

The Seldon operator for management, deployment



Rook-EdgeFS
provided by The Rook Authors

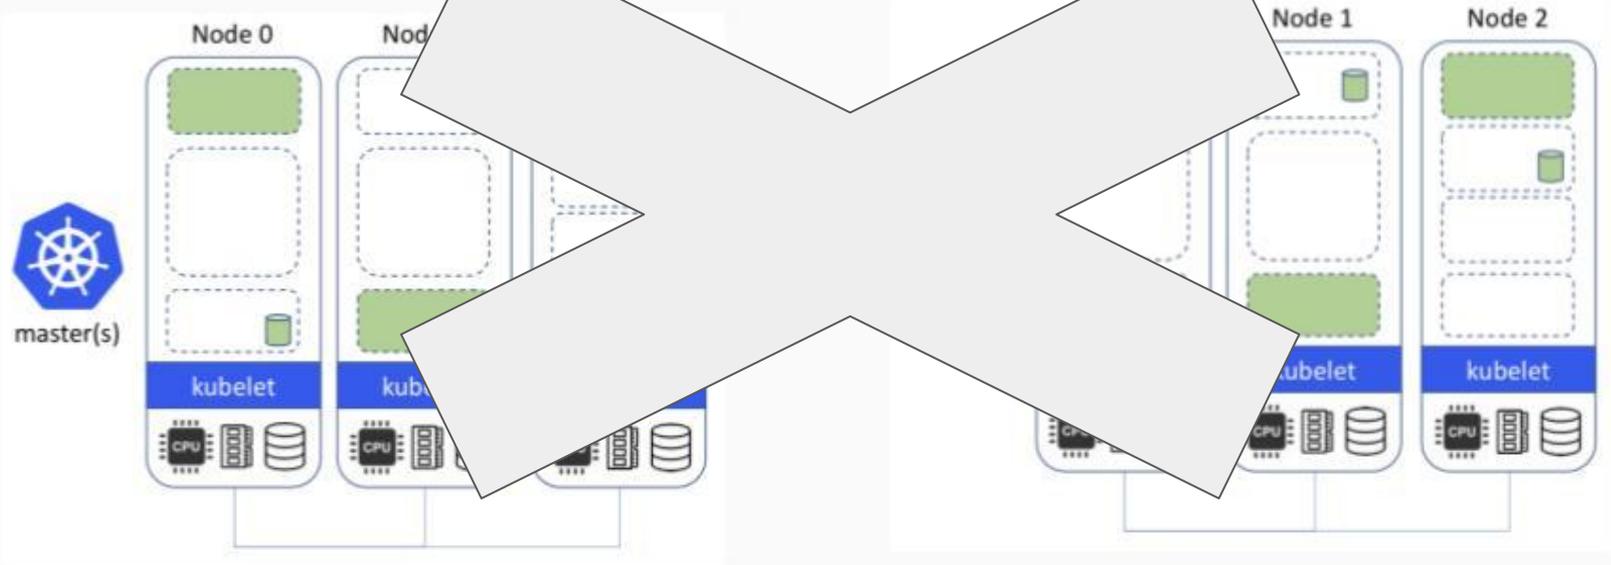
Install and maintain EdgeFS Storage cluster



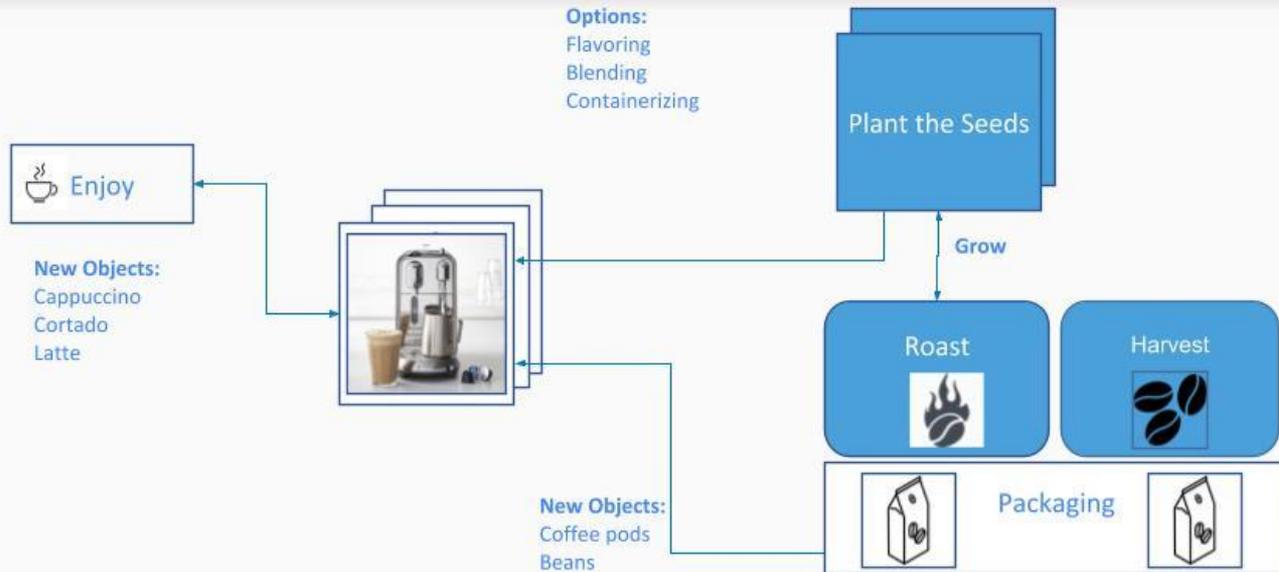
Rook-Ceph
provided by The Rook Authors

Install and maintain Ceph Storage cluster

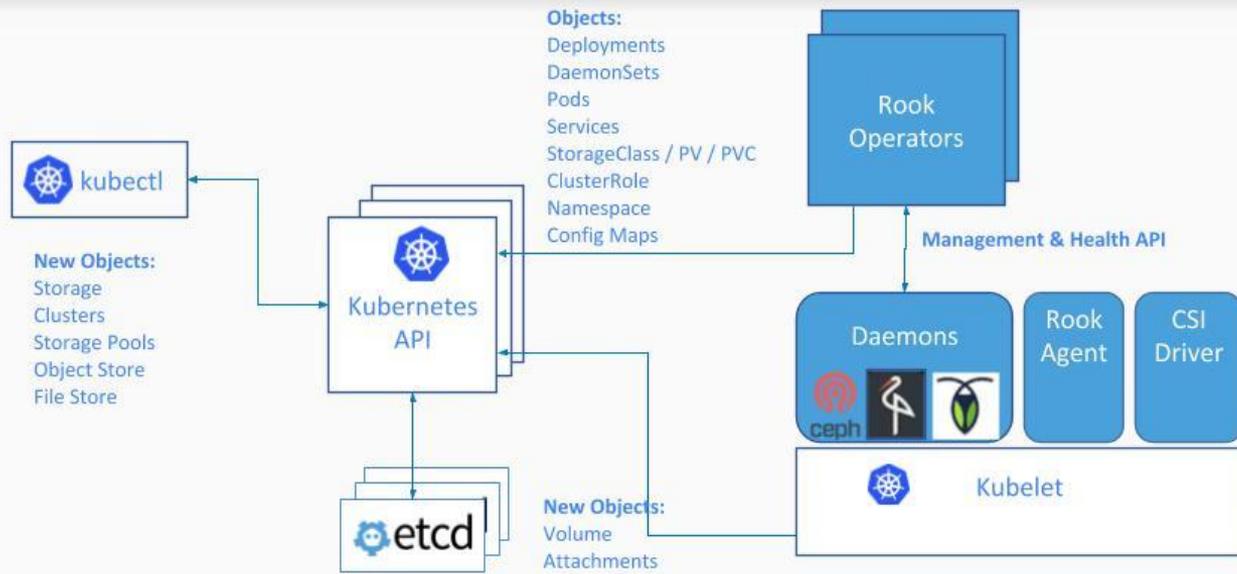
Backplane



Coffee Architecture



Rook Architecture



05

Multicluster scale

“How do I scale and manage
my storage from cluster to
cluster?”



Migrate

Snapshot

Backup

Cloning

Data Policy

Databases

Message
Queues

Object
Stores