

# A Brief History Of Containers


## Kubernetes In A Nutshell

---

Deven Phillips

Red Hat AppDev CoE

# It all started with CGroups



**cgroups** (abbreviated from **control groups**) is a Linux kernel feature that limits, accounts for, and isolates the resource usage (CPU, memory, disk I/O, network, etc.) of a collection of processes.

# CGroups Allow For Granular Control



## Resource limiting

groups can be set to not exceed a configured memory limit, which also includes the file system cache

## Prioritization

some groups may get a larger share of CPU utilization[10] or disk I/O throughput


## Accounting

measures a group's resource usage, which may be used, for example, for billing purposes

## Control

freezing groups of processes, their checkpointing and restarting


# LXC - Linux Containers



## System Level Virtualization

Allows for a single kernel to control multiple operating system “instances”


# Docker



## CGroups with A Packaged Filesystem

No shared filesystem unless explicitly configured

# OCI Containers



## Like Docker, But Less Vulnerable

Runs without 'root' privileges.


## Podman

Runs containers (including Docker containers)

## Buildah


Builds containers, like ``docker build ...``

# Containers Are Great ... UNTIL





## Docker “run”



### One-Offs

Each container needs to be started using the CLI or the Docker API

### Storage

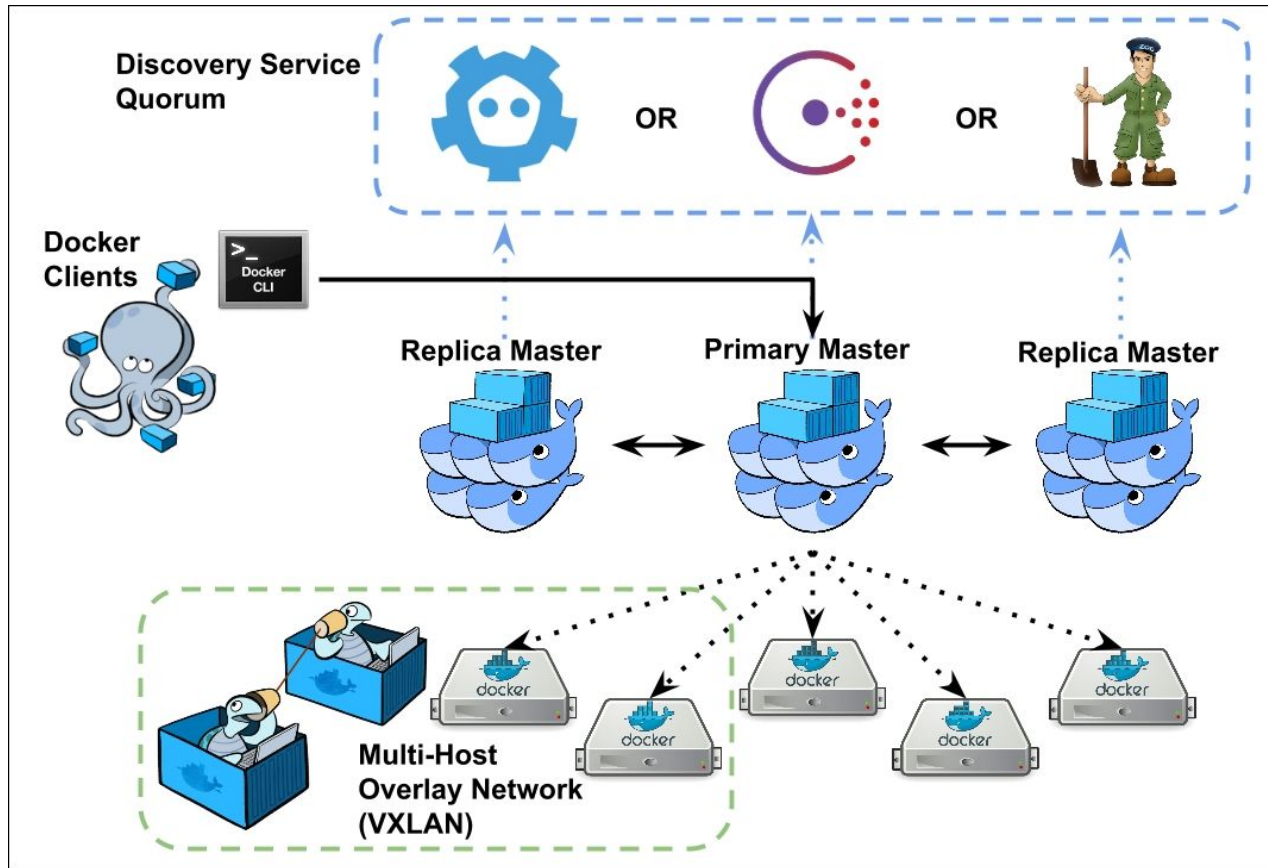
Storage is passed through to the underlying OS

### Networking

Has to be configured manually for each container



# Docker Swarm



## Discovery Server

Raft/Zookeeper/etcd

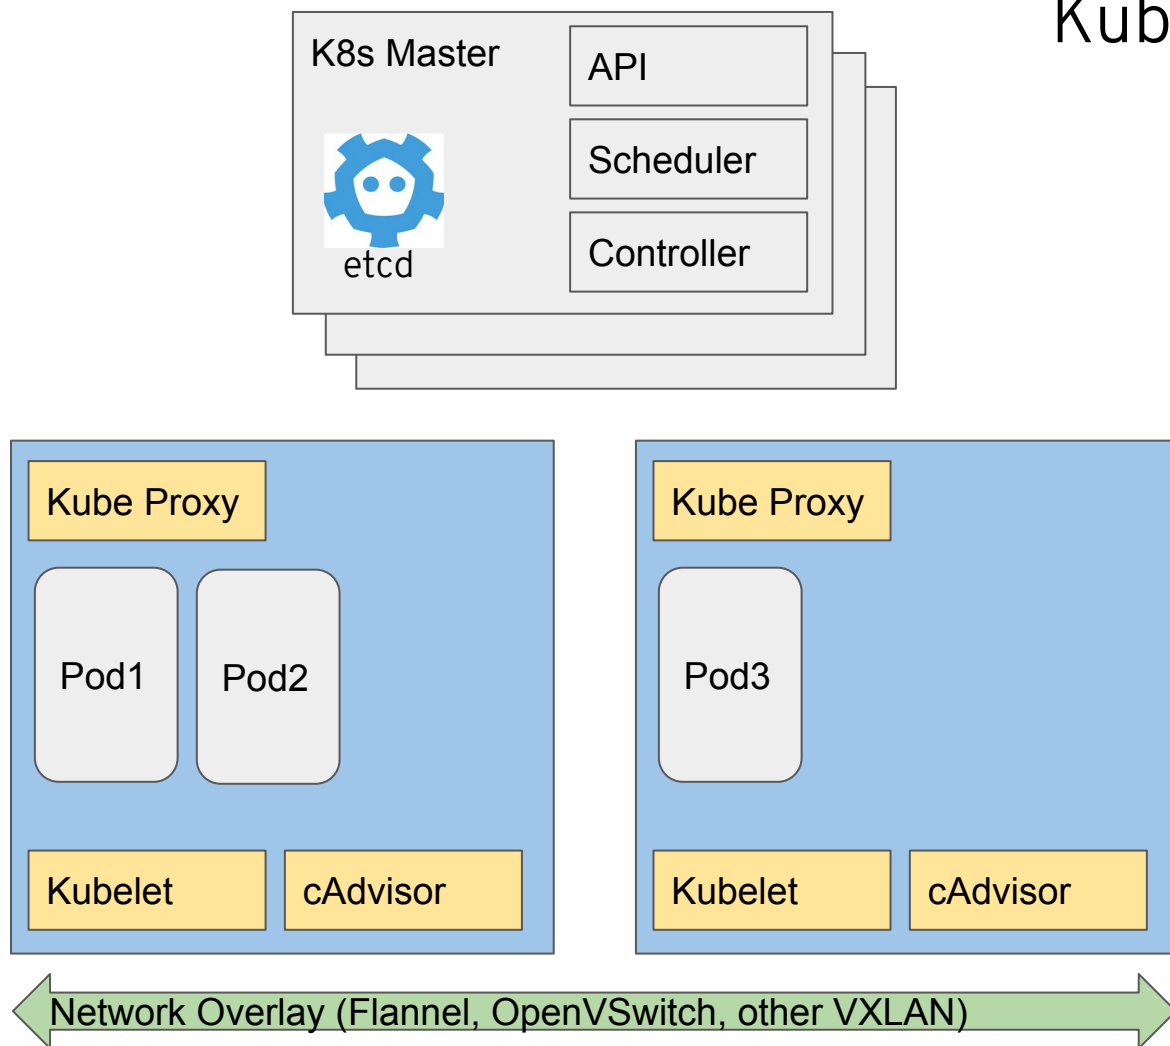
## Networking

Overlay Network Using VXLANs

## Storage

Not really addressed by Swarm, but you could use other distributed filesystems

# Kubernetes



## Kube Master

etcd, API, Controller; written in GoLang

## Networking

Overlay Network Using VXLANs

## Storage

Several cloud-native storage options now available, or use external Gluster/NFS/Ceph/S3/etc...

## Scaling

Pods can be “replicated” to scale horizontally

## Management

Via CLI, Web, or API

## Centralized Logging

Elasticsearch/Fluentd/Kibana

# Kubernetes Operators



## Extend Kubernetes API

Custom Resource Definitions

## Integrate With Kubernetes Controller

Can watch for changes and respond automatically

## Add New Functionality

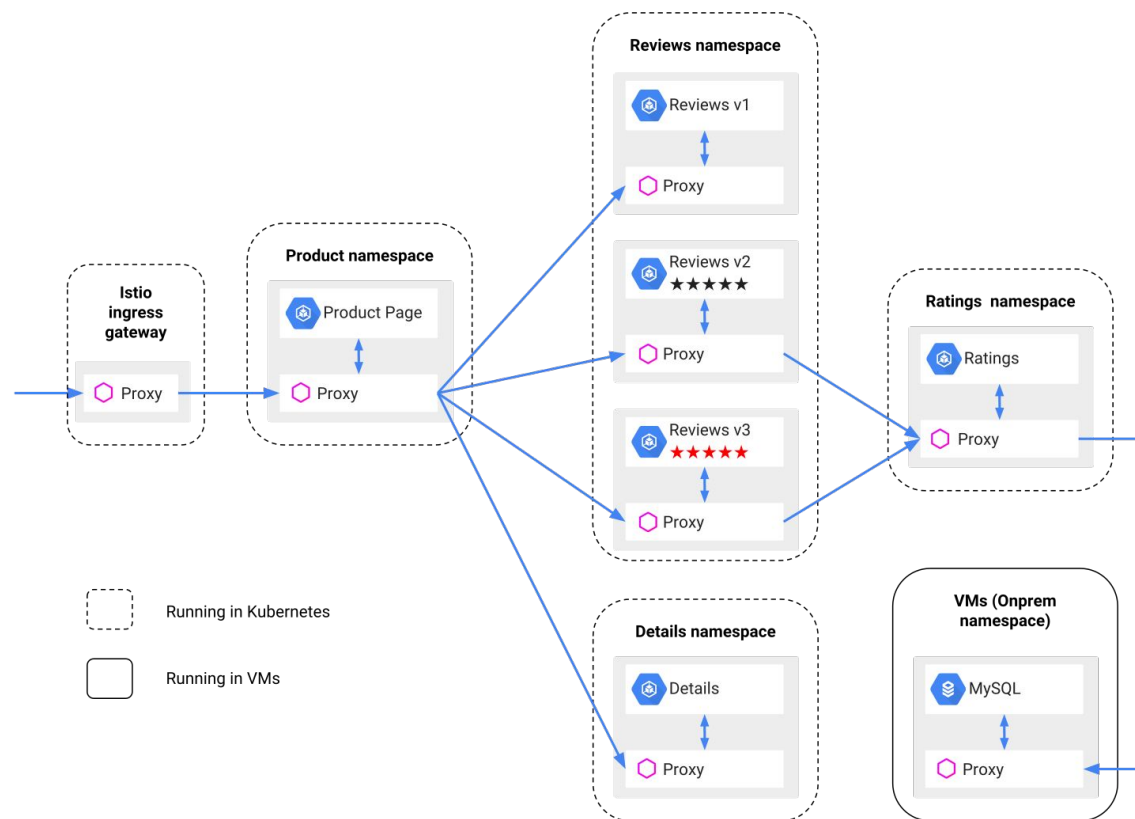
Operators for storage, databases, security, messaging, etc...

Think of it like AWS servers on your own cloud

## Codify Domain Expertise

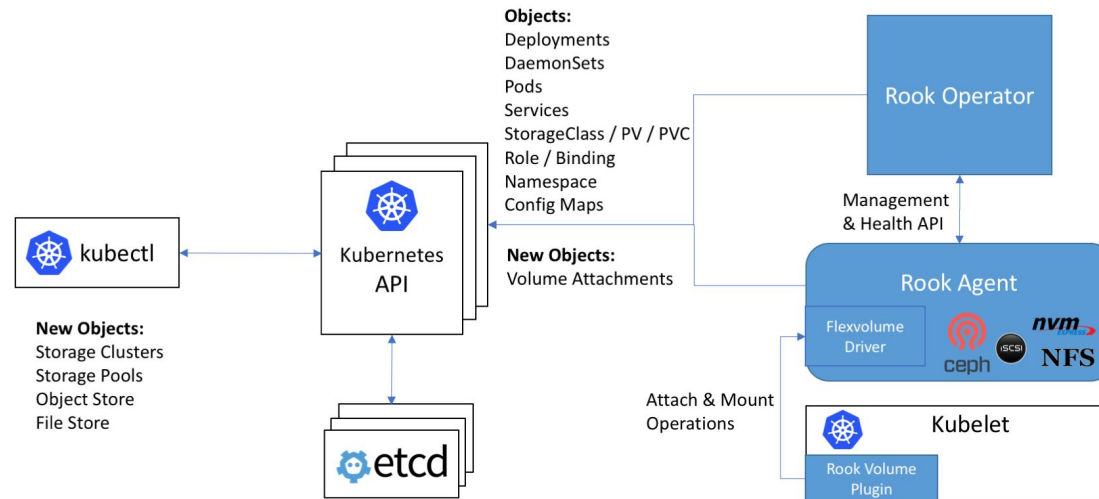
Make services self-managing (as much as possible)

# Istio Operator




- End-To-End Encryption
- Control Plane
- Tracing
- Traffic Management
- Metrics

# Rook Operator



- Scalable Storage
- Storage Management
- Portable Storage
- Replication
- Volumes
- Block Devices
- Object Stores

# Kubernetes Federation



**Sync Resources Across Kubernetes Clusters**

etcd Sync

**Migrate Workloads/Storage/Configuration**

Allow workloads to migrate from one DC to another

## Further Resources



[Rancher Kubernetes](#)



[OpenShift Kubernetes](#)

- [Operator Hub](#) - A Marketplace for K8s Operators
- [OpenVSwitch](#) - An Open Source virtual switch
- [Operator Framework](#) - How to create Operators
- Mini-Kubernetes For Testing/Experimentation
  - [K3s](#)
  - [MiniKube](#)
  - [MicroK8s](#)



# 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.



[linkedin.com/company/red-hat](https://linkedin.com/company/red-hat)



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



[facebook.com/redhatinc](https://facebook.com/redhatinc)



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