

Containers and distributed applications

Xiang Li

Head of distributed system @ CoreOS

Distributed system at CoreOS

We solve hardest distributed system problems

Mission

Secure the Internet

STRATEGY

Make servers easy to upgrade

CoreOS Linux

Self update

Separate OS and the Apps

- enforce container oriented deployment

CoreOS Linux

1000+ releases

STRATEGY

Simplify application management

etcd

Configuration management
Distributed coordination

fleet

Distributed init system

Kubernetes

The “Replacement” for fleet

Manage containerized applications

STRATEGY

On any infrastructure

CoreOS Linux

Public cloud

- AWS, Azure, GCE, Packet...

Private cloud

- OpenStack

Bare metal

- iPXE

flannel

Enabling IP per container on any infrastructure

CoreOS

Accelerate the industry adoption of containers

Container

Run applications anywhere

Container

Package

Container

Push

Pull

Container

Run

Docker

```
$ docker build
```

```
$ docker push
```

```
$ docker run
```

Container Engine

Manage the lifecycle of a *single* application

Why CoreOS uses container

Make OS secure

Why CoreOS uses container

Deploy distributed
applications

Why Google uses container

Run distributed applications
efficiently


Container management

Managing the lifecycle of *one* application instance is not enough

Running containers on one node



Running containers on one node



Running containers on two nodes



Running containers on two nodes



Running containers on a cluster

Where to put my containers?



Running containers on a cluster

How to find my containers?



Running containers on a cluster


How to connect to my containers?



Running containers on a cluster




Scheduler



Running containers on a cluster




Scheduler




Running containers on a cluster

Scheduler




Running containers on a cluster

 color=yellow




Running containers on a cluster

 `color=yellow` Select color = yellow



Running containers on a cluster

yellow.mycluster
Select color = yellow



Running containers on a cluster

- Application deployment
 - Replication controller
 - Rolling upgrades
 - Auto scaling
- Configuration and secret
- Resource management
 - Storage
 - Network

GIFEE

Google infrastructure for everyone else - A flexible, secure, reliable infrastructure for running distributed applications.

- CoreOS

Application developers

I hate dealing with operations!

Application developers

Run my application on a “cluster manager”
now!

System builders

Trust us!

Operators

Here is Kubernetes!

Operators

Here is ~~Kubernetes!~~
Docker Swarm
DC/OS

But ...

Can your applications run
smoothly on the cluster
manager?

Cloud native application

Distributed
Stateless

Distributed and stateless

Scale

Distributed and stateless

Fault tolerance

Distributed and stateless

Easy to operate

Stateless

No dependency on local resources

Stateless

Treat storage as services


Stateless

Store configuration in the environment

Managed stateless application

Spec


- Container: `example.com/myApp?version=1`
- Replica: 3
- Restart: Always



Distributed and stateless

Spec


- Container: `example.com/myApp?version=1`
- Replica: 3
- Restart: Always



Distributed and stateless

Spec


- Container: `example.com/myApp?version=1`
- Replica: 3
- Restart: Always




Distributed and stateless

Spec

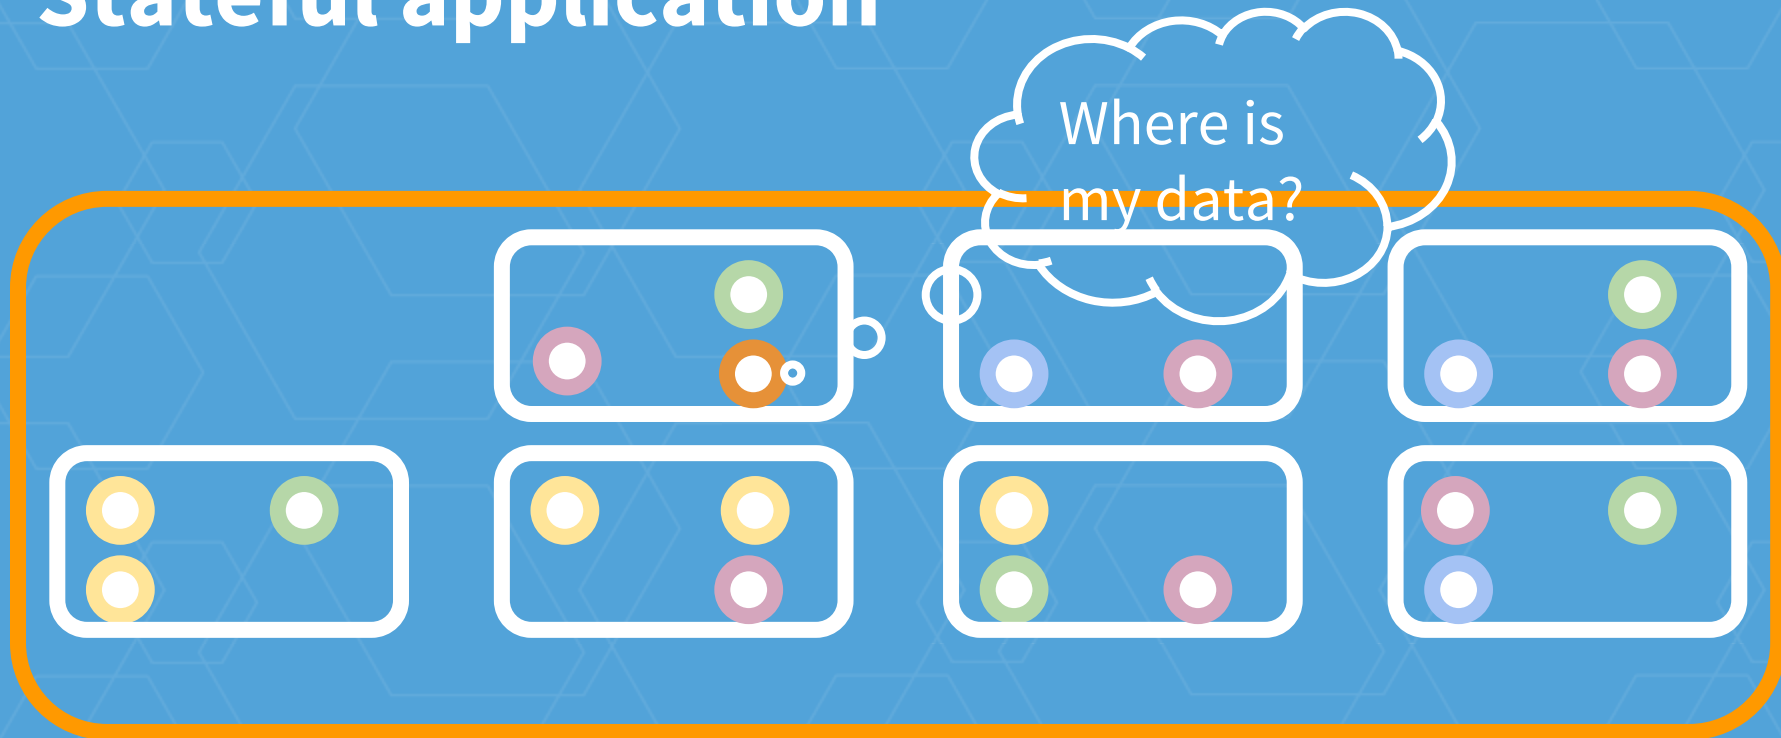
- Container: `example.com/myApp?version=1`
- Replica: 3
- Restart: Always



Stateful application



Stateful application



Stateful application

Fully managed by human

Or semi-fully managed by human

Container 2.0

Container 2.0 is the ability to run (and orchestrate) both stateless and stateful services on the same set of resources.

- Mesosphere


Example: manage etcd clusters

etcd

- Clustered key-value store
- Writes are persisted on disk
- Writes are replicated to all nodes

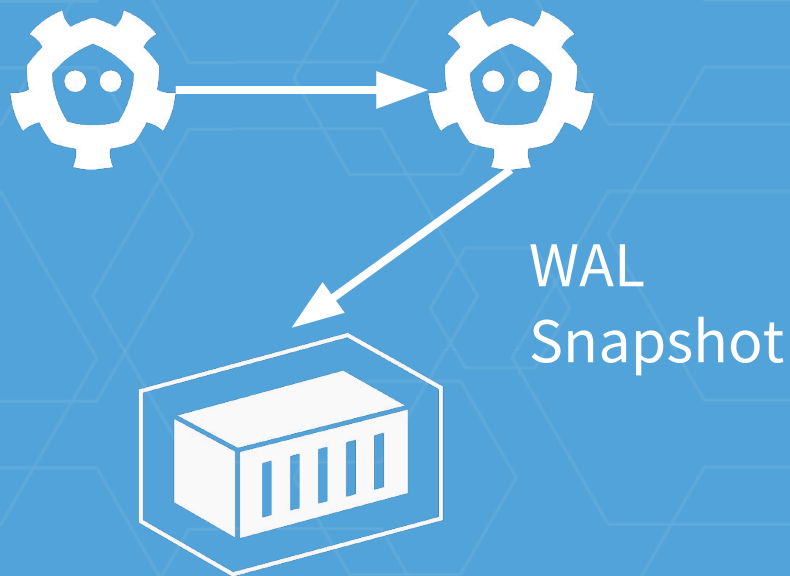
Example: manage etcd clusters

Tightly coupled with local storage




Example: manage etcd clusters

Indirectly coupled with local storage of its peers




Example: manage etcd clusters

Strong membership



Example: manage etcd clusters

Strong membership



Example: manage etcd clusters

External controller

- A big for loop to simulate human operator
 - Manage data migration
 - Manage membership changes
 - Manage configuration changes


Example: manage etcd clusters

Bootstrap a 3 member cluster



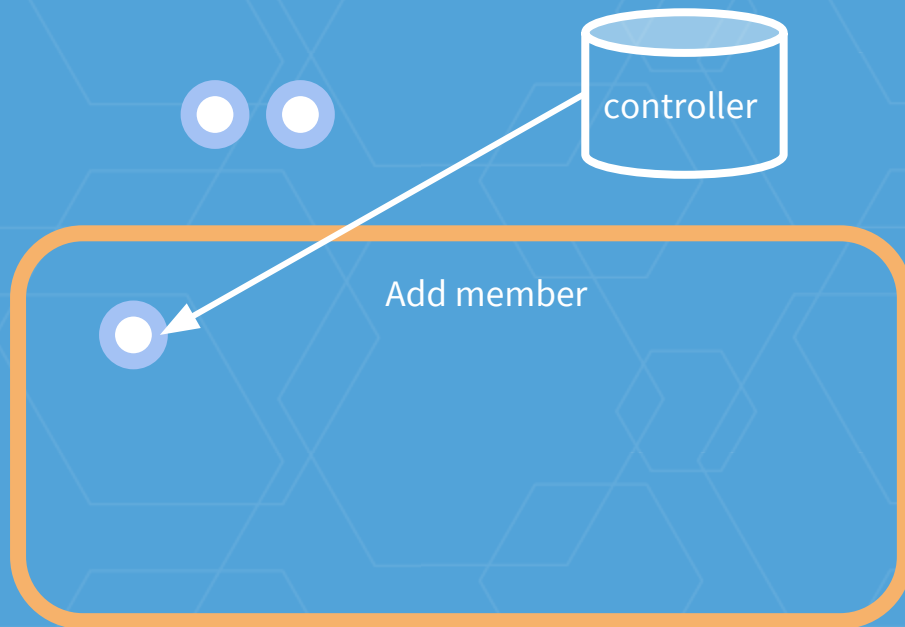
Example: manage etcd clusters

Bootstrap a 3 member cluster




Example: manage etcd clusters

Bootstrap a 3 member cluster



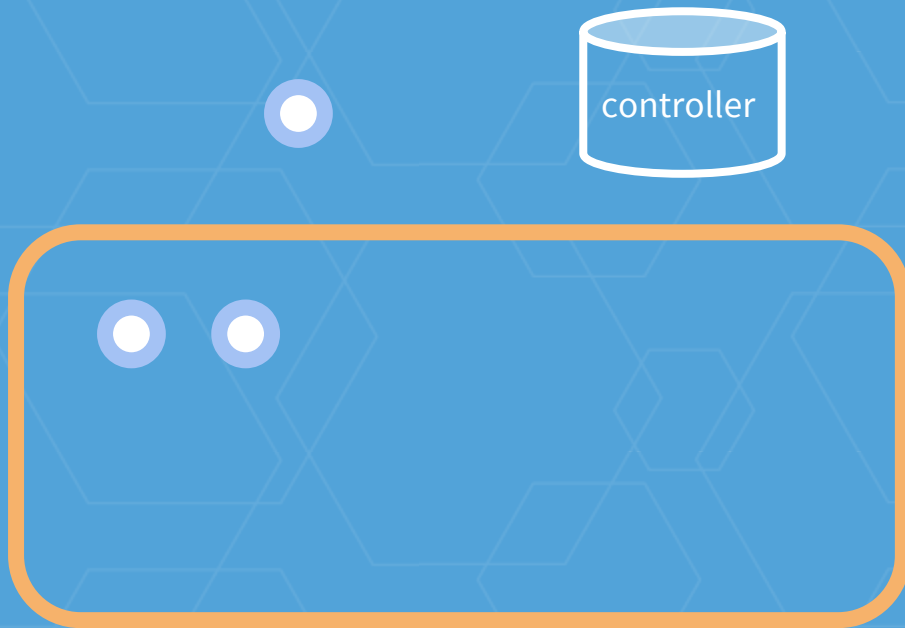
Example: manage etcd clusters

Bootstrap a 3 member cluster




Example: manage etcd clusters

Bootstrap a 3 member cluster



Example: manage etcd clusters

Bootstrap a 3 member cluster




Example: manage etcd clusters

Resize from 3 to 5

Select app = etcd


Add member if size < 5



Example: manage etcd clusters

Resize from 3 to 5

Select app = etcd
Add member if size < 5




Example: manage etcd clusters

Resize from 3 to 5

Select app = etcd


Add member if size < 5



Example: manage etcd clusters

Failure recovery


Select app = etcd



Example: manage etcd clusters

Failure recovery


Select app = etcd
Remove dead one



Example: manage etcd clusters

Failure recovery


Goto resize



Example: manage etcd clusters

Failure recovery

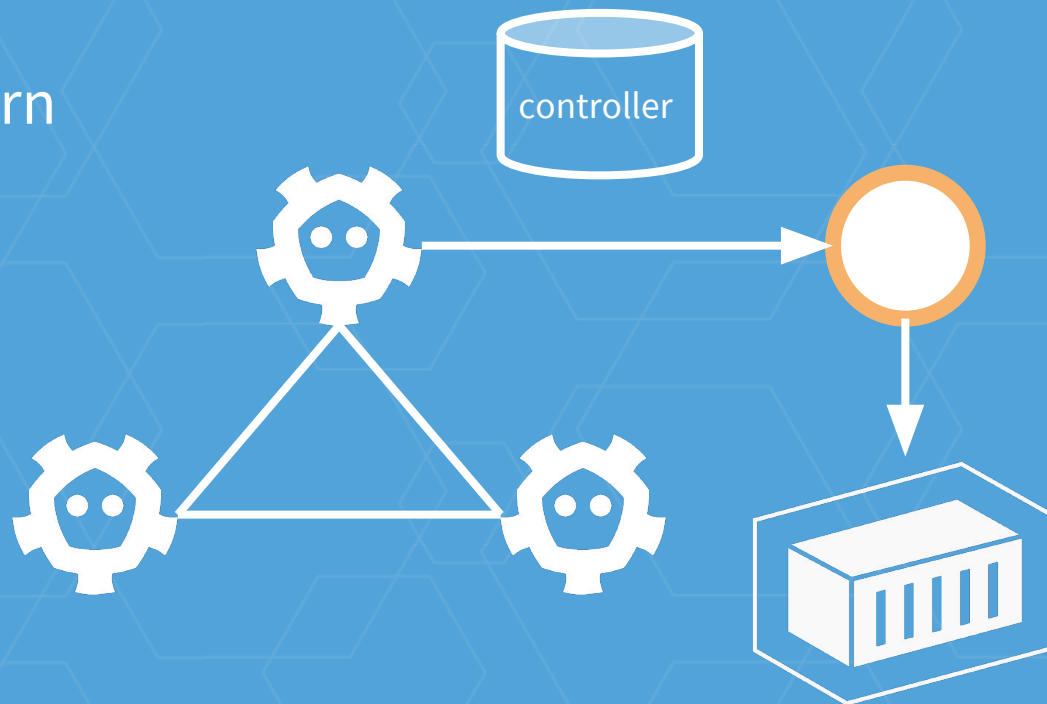
Select app = etcd



Example: manage etcd clusters

Backup

- Sidecar pattern



Stateful distributed application

More controllers to replace human

Tips for building distributed application

Use coordination libraries/software

- Leader election
- Locking
- Queue
- Barrier

Tips for building distributed application

Use a RPC framework

- strict contract between components
- across languages support
- client side load-balancing
- client side naming resolution
- auth

Tips for building distributed application

Metrics

- Request rate/duration
- Error rate
- Internal state

Tips for building distributed application

Logging

- Human actionable events
- Critical State changes

Future is bright