



KubeCon



CloudNativeCon

— North America 2018 —

Deep Dive: Cluster API



Who are we?



KubeCon



CloudNativeCon

North America 2018



Robert Bailey

Staff Software Engineer @Google
@roberthbailey



David Watson

Software Engineer @Samsung
@davidewatson

Agenda



KubeCon



CloudNativeCon

North America 2018

- Motivation
- What is Cluster API?
- Bootstrapping / Pivoting
- API Specification
- Roadmap
- Getting Involved

Motivation



KubeCon



CloudNativeCon

North America 2018

- Cluster Management is difficult
 - kubeadm helps, but doesn't provide infrastructure management
- Tooling and ecosystem is fragmented
 - Too many tools to choose from
- Difficult to build higher order functionality
 - Additional automation (autoscaling)
 - Hybrid Cloud
 - Managed Control Plane

What is Cluster API?



KubeCon



CloudNativeCon

North America 2018

Declarative API

- Cluster
- Machine
- Machine Set + Machine Deployment
- MachineClass

Common Logic

- Machine Lifecycle (incl. Provisioning)
- Machine Upgrade

Pluggable Architecture

- Infrastructure platform (AWS, GCP, vSphere, etc.)
- Support for various Operating Systems

Tooling, Services

- Cluster Bootstrapping, Upgrade
- Autoscaling, Repair, Node Auto-provisioning



KubeCon



CloudNativeCon

North America 2018

Getting up and running

Bootstrapping

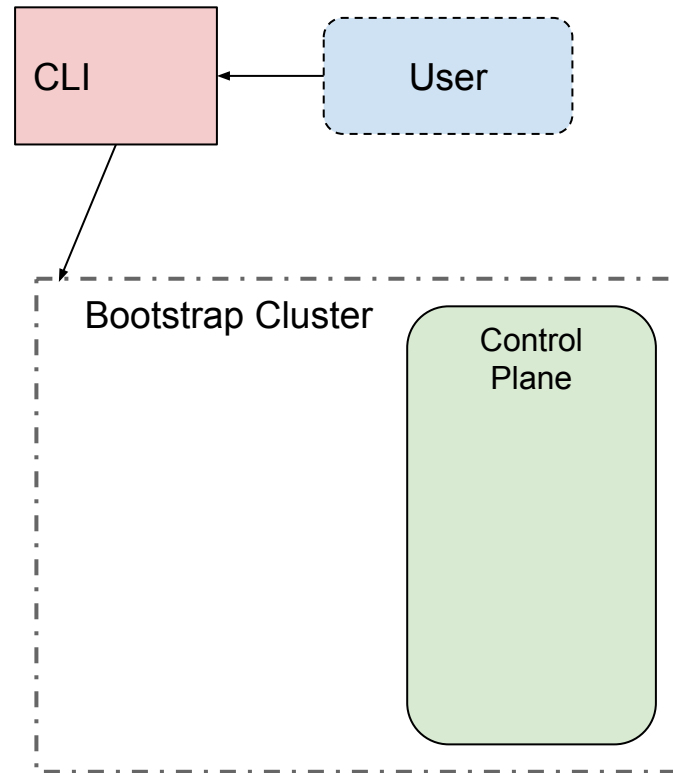


KubeCon



CloudNativeCon

North America 2018



Bootstrapping

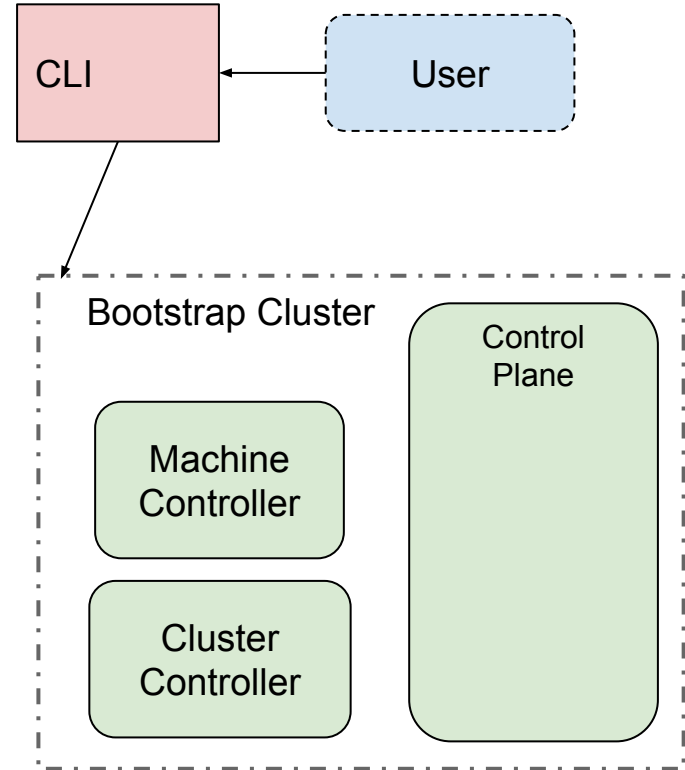


KubeCon



CloudNativeCon

North America 2018



Bootstrapping

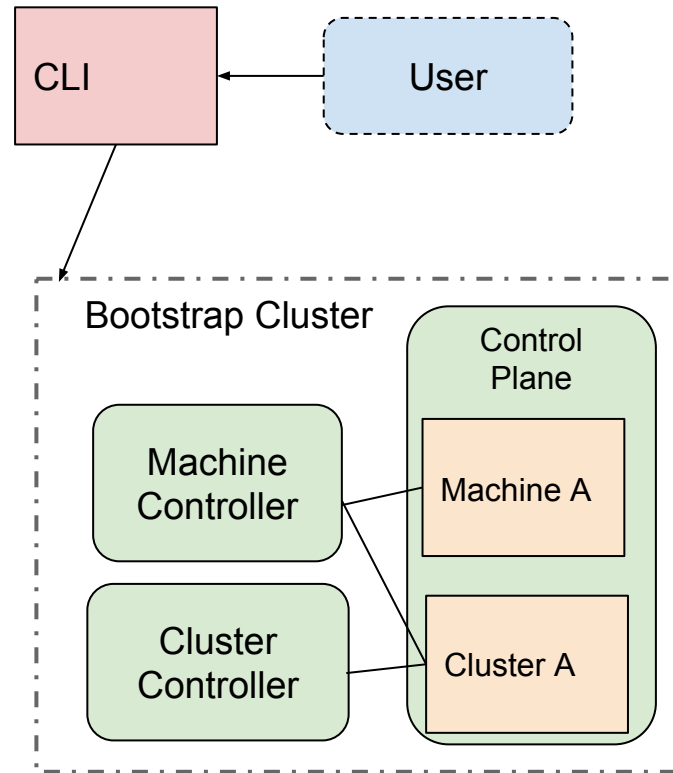


KubeCon



CloudNativeCon

North America 2018



Bootstrapping

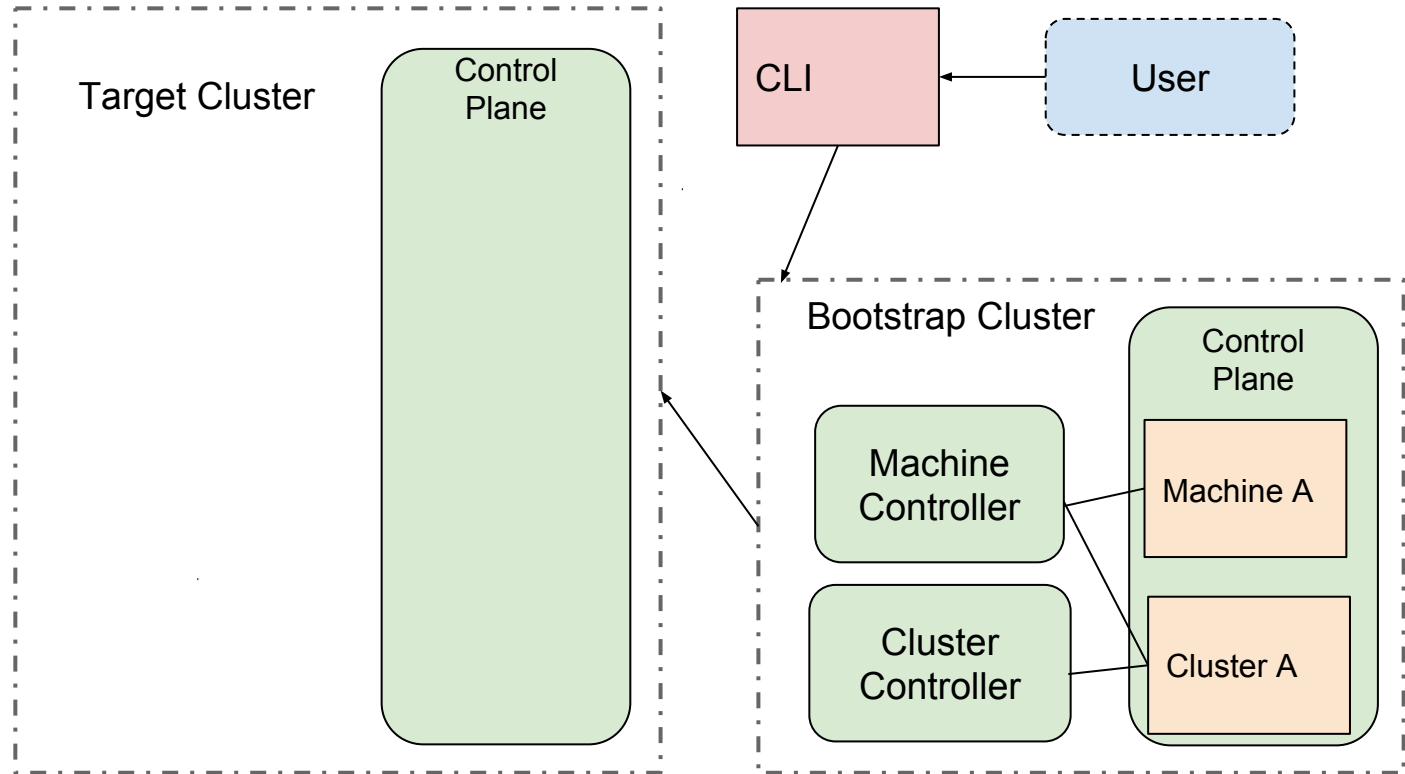


KubeCon



CloudNativeCon

North America 2018





KubeCon



CloudNativeCon

North America 2018

To pivot or not to pivot?

Pivoting

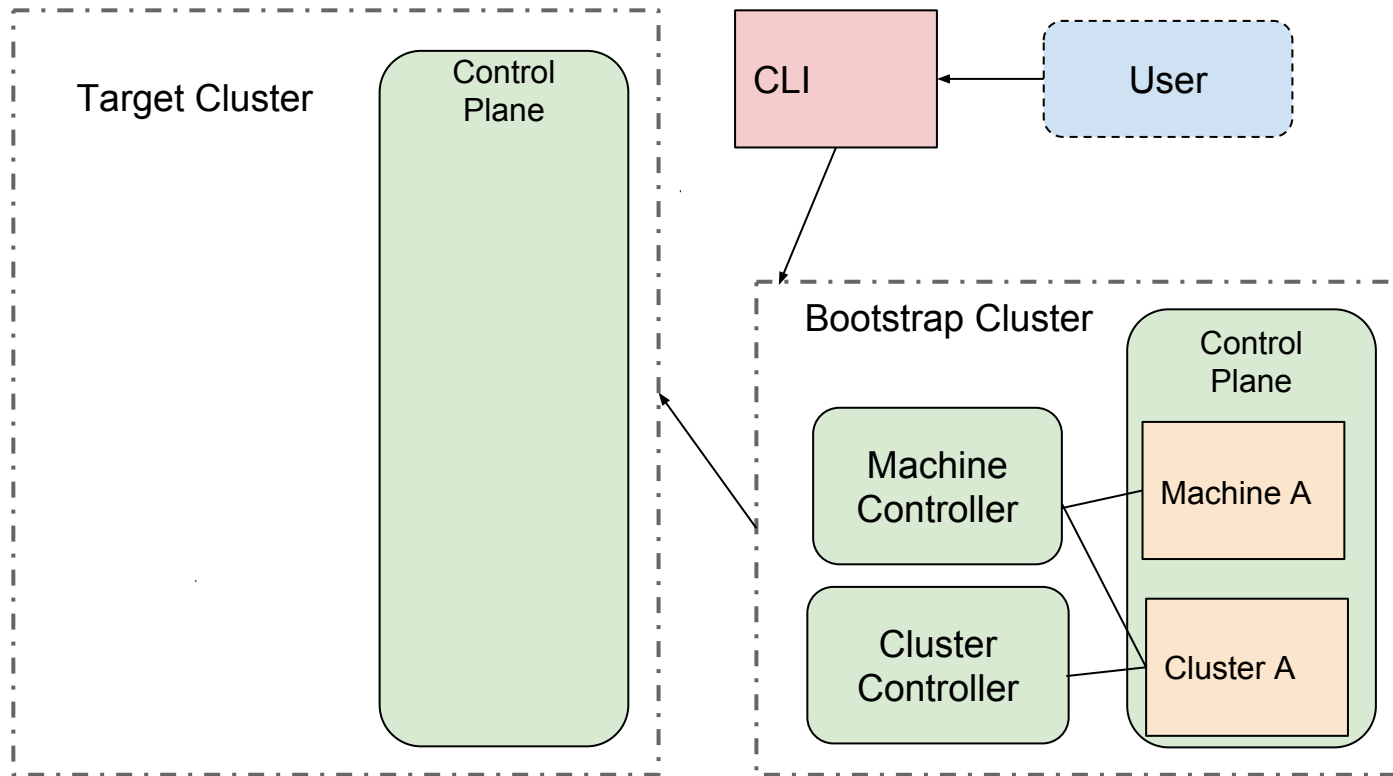


KubeCon



CloudNativeCon

North America 2018



Pivoting: Copy CRDs

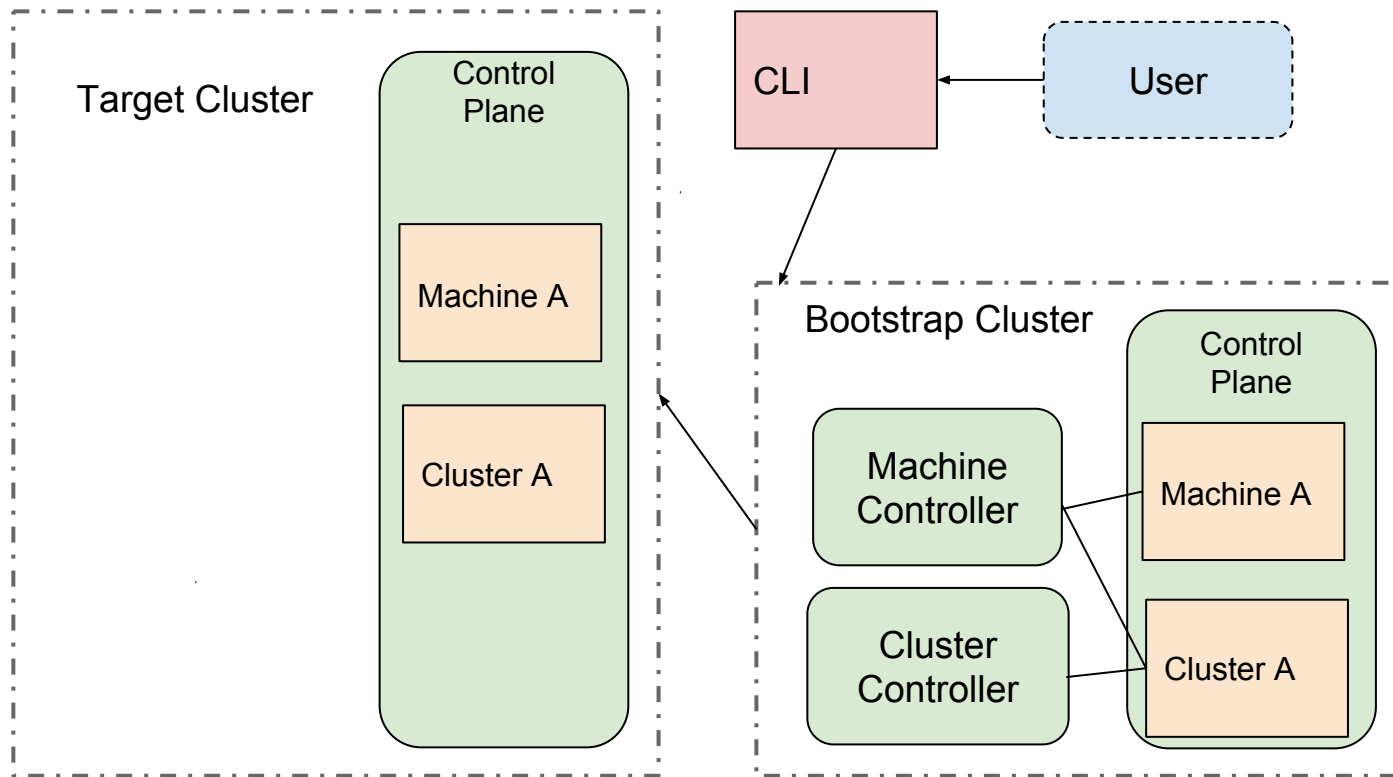


KubeCon



CloudNativeCon

North America 2018



Pivoting: Start Controllers

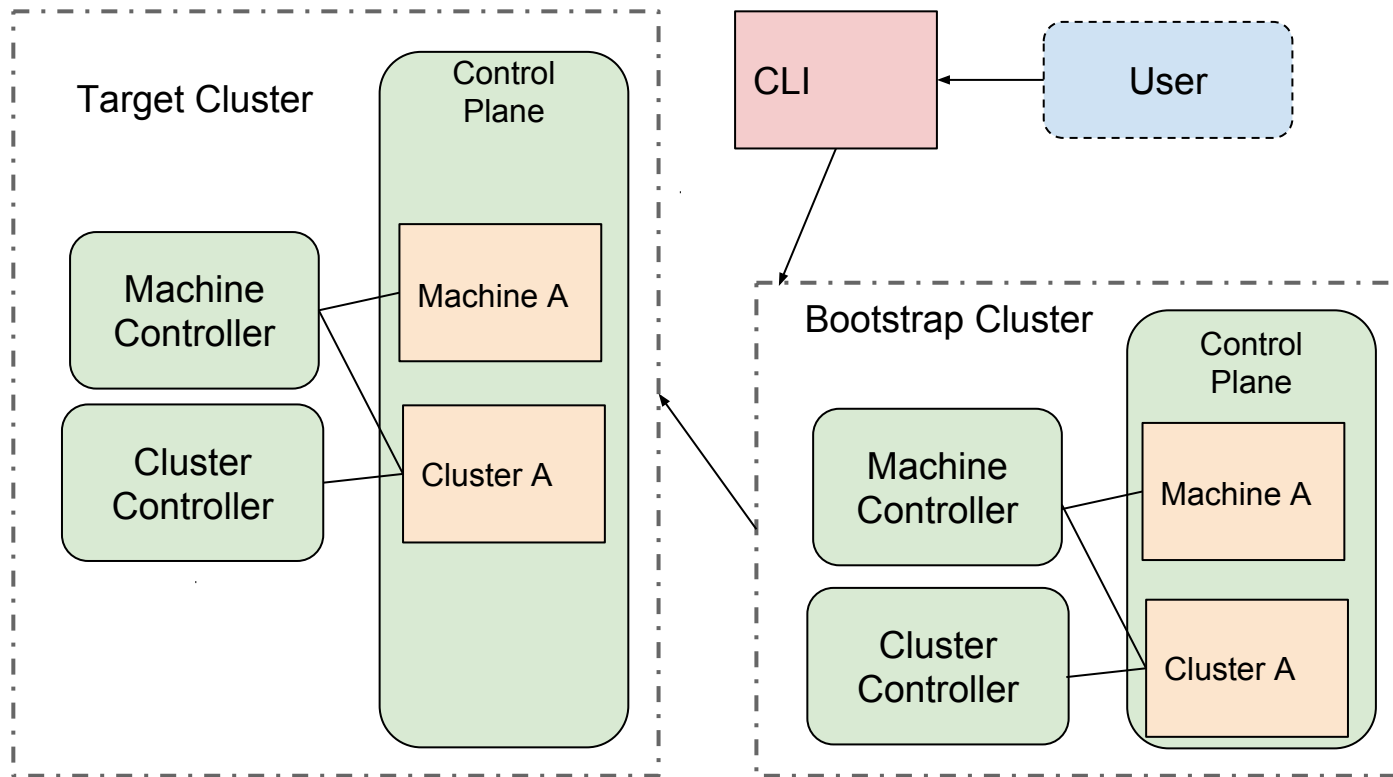


KubeCon



CloudNativeCon

North America 2018



Pivoting: Delete Bootstrap Cluster


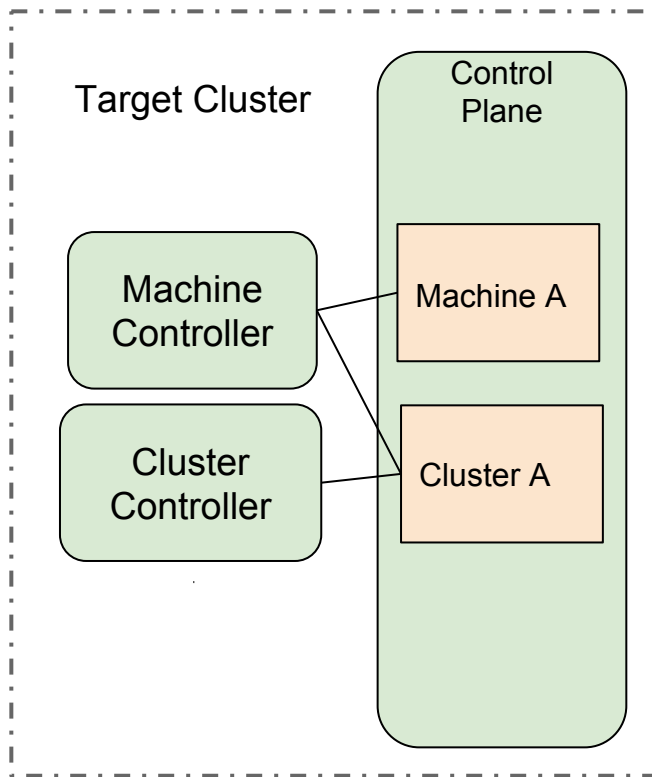


KubeCon



CloudNativeCon

North America 2018





KubeCon



CloudNativeCon

North America 2018

What's in the API?

Cluster



KubeCon



CloudNativeCon

North America 2018

Cluster-wide configuration.

Generic networking concepts:

- Pod and service ranges
- DNS domain
- [in status] Cluster endpoint

Providers can hook in where specificity is needed.

```
apiVersion: "cluster-api.k8s.io/v1alpha1"
kind: Cluster
metadata:
  name: my-first-cluster
spec:
  providerSpec:
    ...
  clusterNetwork:
    services:
      cidrBlocks: ["10.96.0.0/12"]
    pods:
      cidrBlocks: ["192.168.0.0/16"]
    serviceDomain: "cluster.local"
```

Machine



KubeCon



CloudNativeCon

North America 2018

Configuration for a specific machine.

Spec has the desired kubelet version.

Providers are responsible for defining the shape of the machine (cpu, ram, gpu, etc.).

```
apiVersion: "cluster.k8s.io/v1alpha1"
kind: Machine
metadata:
  name: my-first-machine
spec:
  providerSpec:
    ...
  versions:
    kubelet: 1.12.0
```

MachineSet



KubeCon



CloudNativeCon

North America 2018

Configuration for a group of homogeneous machines.

A template for creating machines and a desired number of machines.

```
apiVersion: "cluster.k8s.io/v1alpha1"
kind: MachineSet
metadata:
  name: my-first-machine-set
spec:
  replicas: 3
  template:
    metadata:
    spec:
      ...
```

MachineDeployment



KubeCon



CloudNativeCon

North America 2018

Declarative updates for Machines and MachineSets.

Update strategy allows control of the rate at which a change is applied.

```
apiVersion: "cluster.k8s.io/v1alpha1"
kind: MachineDeployment
metadata:
  name: my-first-machine-deployment
spec:
  replicas: 3
  template:
    ...
  strategy:
    type: RollingUpdate
    rollingUpdate:
      maxUnavailable: 0
      maxSurge: 1
```

Analogy to existing resources



KubeCon



CloudNativeCon

North America 2018

Pod



RepliaSet



Deployment



Machine



MachineSet



MachineDeployment



Provider Spec



KubeCon



CloudNativeCon

North America 2018

Platform-specific configuration
for Machines and Clusters.

The embedded type is expected
to itself be a versioned
Kubernetes-style type.

```
apiVersion: "cluster.k8s.io/v1alpha1"
kind: Machine
metadata:
  name: my-first-machine
spec:
  providerSpec:
    value:
      apiVersion: "gceproviderconfig/v1alpha1"
      kind: "GCEMachineProviderConfig"
      project: "gke-kubecon"
      zone: "us-central1-f"
      machineType: "n1-standard-2"
      os: "ubuntu-1604-lts"
```

Sharing Provider Specs



KubeCon



CloudNativeCon

North America 2018

Repeating the provider specific fields in every Machine or Machine template gets cumbersome.

How can we share the data to reduce duplication?

```
apiVersion: "cluster.k8s.io/v1alpha1"
kind: Machine
metadata:
  name: my-first-machine
spec:
  providerSpec:
    valueFrom:
      machineClass:
        kind: MachineClass
        name: my-first-machine-class
```

MachineClass



KubeCon



CloudNativeCon

North America 2018

Platform-specific configuration for Machines.

The embedded type is expected to itself be a versioned Kubernetes-style type.

```
apiVersion: "cluster-api.k8s.io/v1alpha1"
kind: MachineClass
metadata:
  name: my-first-machine-class
spec:
  providerSpec:
    apiVersion: "gceproviderconfig/v1alpha1"
    kind: "GCEMachineProviderConfig"
    project: "gke-kubecon"
    zone: "us-central1-f"
    machineType: "n1-standard-2"
    os: "ubuntu-1604-lts"
```


Analogy to existing resources



KubeCon



CloudNativeCon

North America 2018

StorageClass



MachineClass



Demo



KubeCon



CloudNativeCon

North America 2018

Creating a new provider

Future Roadmap



KubeCon



CloudNativeCon

North America 2018

- Alpha and beta releases
 - Alpha in 2-3 months
 - Beta later this year (may be blocked on CRD features being released)
- Multi-master
 - Possible today on AWS with clusterctl phases and complex docs
- Prow integration
 - Both for cluster api and k/k

Future Roadmap

- Baremetal
 - Looking into a generic provider using webhooks
- Documentation
 - gitbook
- Adopt horizontal SIG tooling
 - etcdadm
 - Addons as operators
 - Bundle spec

Getting Involved



KubeCon



CloudNativeCon

North America 2018

- Help with documentation / project management
 - No coding skills necessary!
- Look for issues with "help wanted" or "good first issue"
 - Start with small changes, work your way up to larger changes
- Contribute to the provider for your environment
 - Or create one if it doesn't exist

Getting Involved



KubeCon



CloudNativeCon

North America 2018

- <https://github.com/kubernetes-sigs/cluster-api>
- Join kubernetes-sig-cluster-lifecycle
- Weekly meeting on Wed @ 10:00 PT
- Office Hours: Weekly on Wed @ 06:00 AM and Tues @ 12:00 PT
 - Meeting notes - <https://goo.gl/KrjNgZ>
 - Recordings - <https://goo.gl/nD5HEd>
- Slack: #cluster-api



KubeCon



CloudNativeCon

North America 2018

Questions?



KubeCon

CloudNativeCon

————— North America 2018 —————

Thank You!

