Deep Dive: Cluster API

Cecile Robert-Michon & Naadir Jeewa

About us

KubeCon CloudNativeCon Europe 2020 Virtual

Cecile Robert-Michon
Software Engineer @ Microsoft
Cluster API Provider Azure Maintainer

@cecile @cecilerobertm

@CecileRobertMichon

Naadir Jeewa
Member of Technical Staff @ VMWare
Cluster API Provider AWS Maintainer

@naadir

@randomvariable

What is Cluster API?

"Cluster API is a Kubernetes project to bring declarative, Kubernetes-style APIs to cluster ***creation****,* ***configuration****, and* ***management****."*

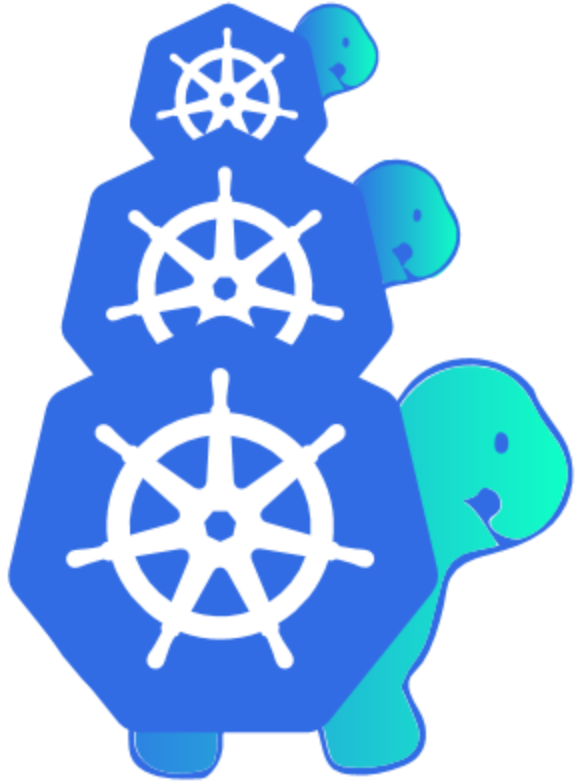

Cluster API is a building block

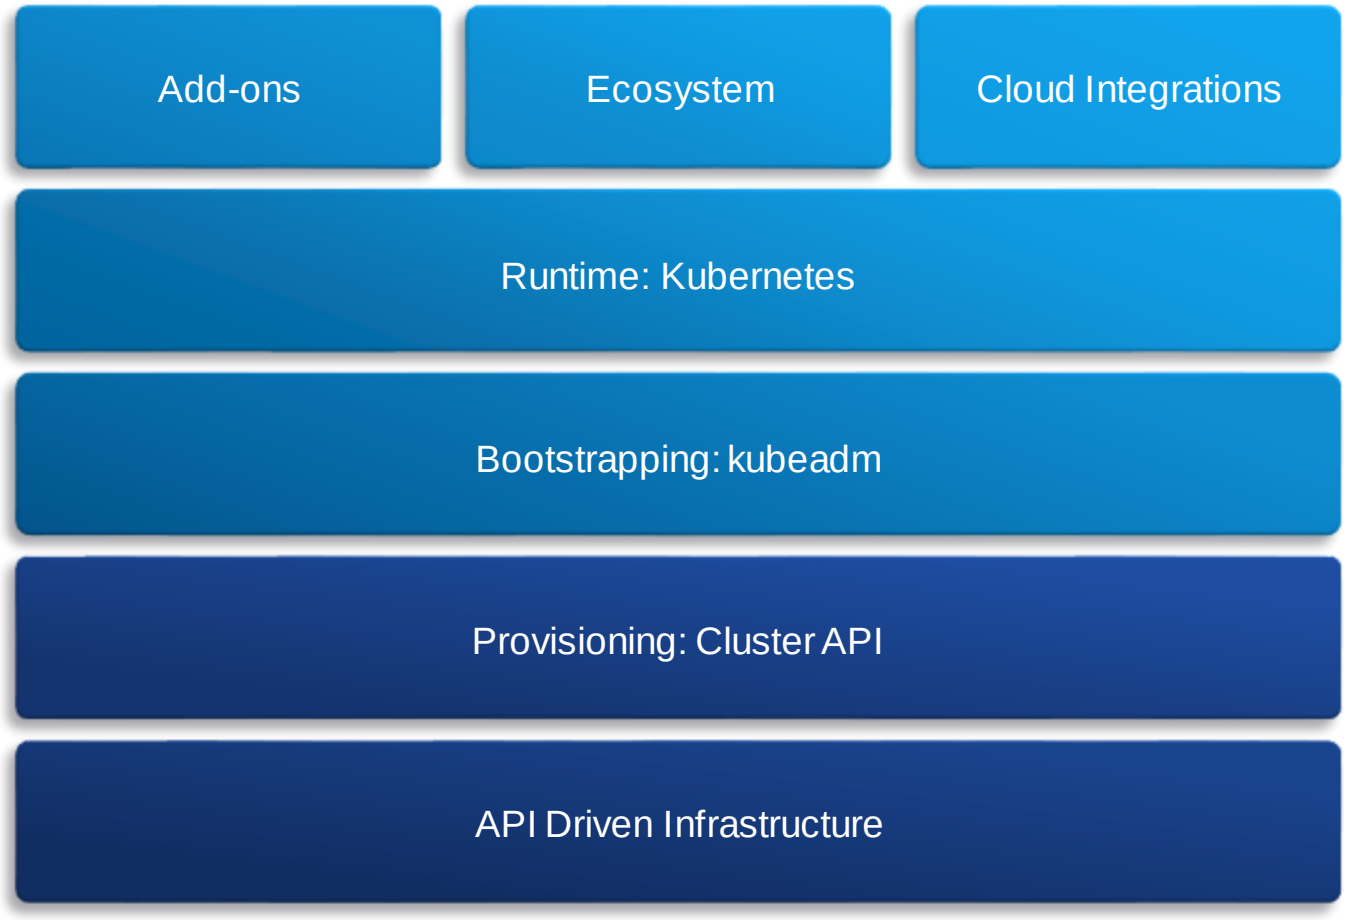

Cluster API is a building block

cluster-api.sigs.k8s.io/reference/providers.html

How does it work?

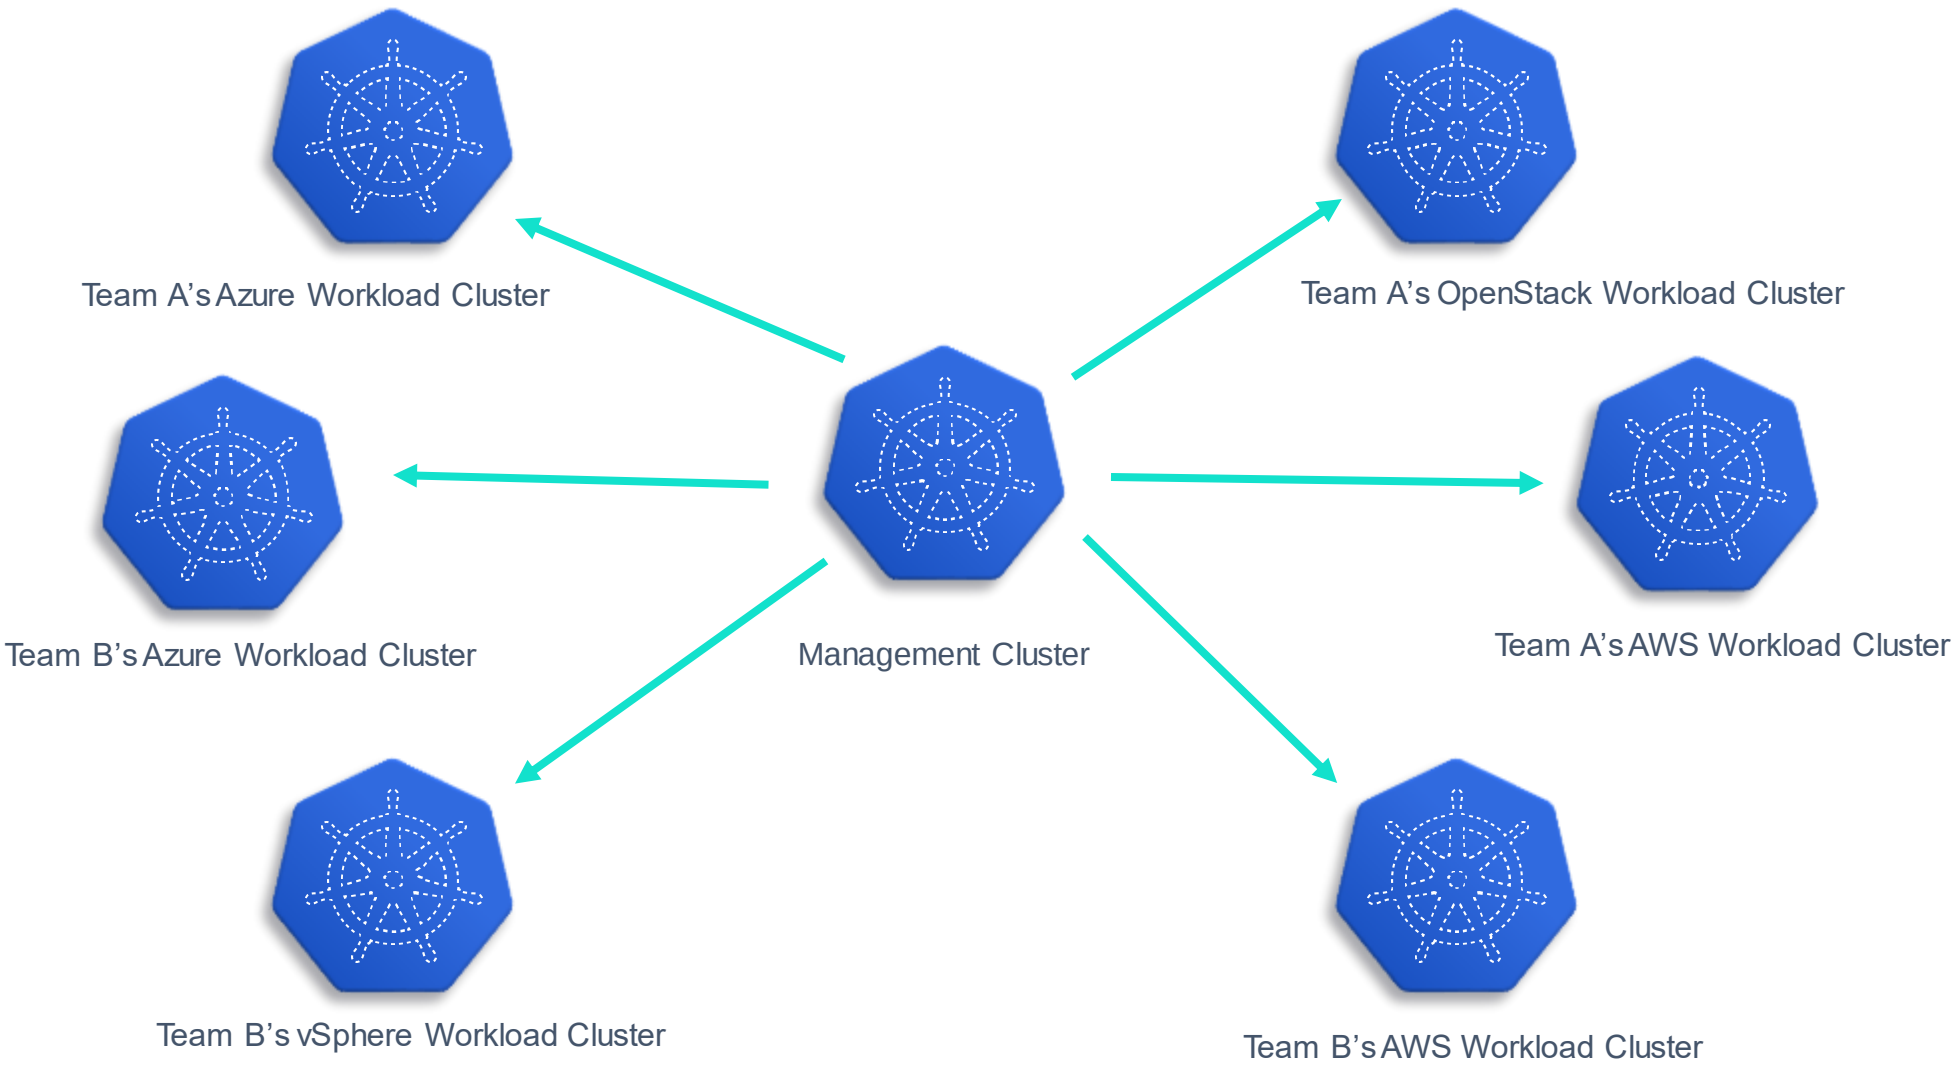

How does it work?

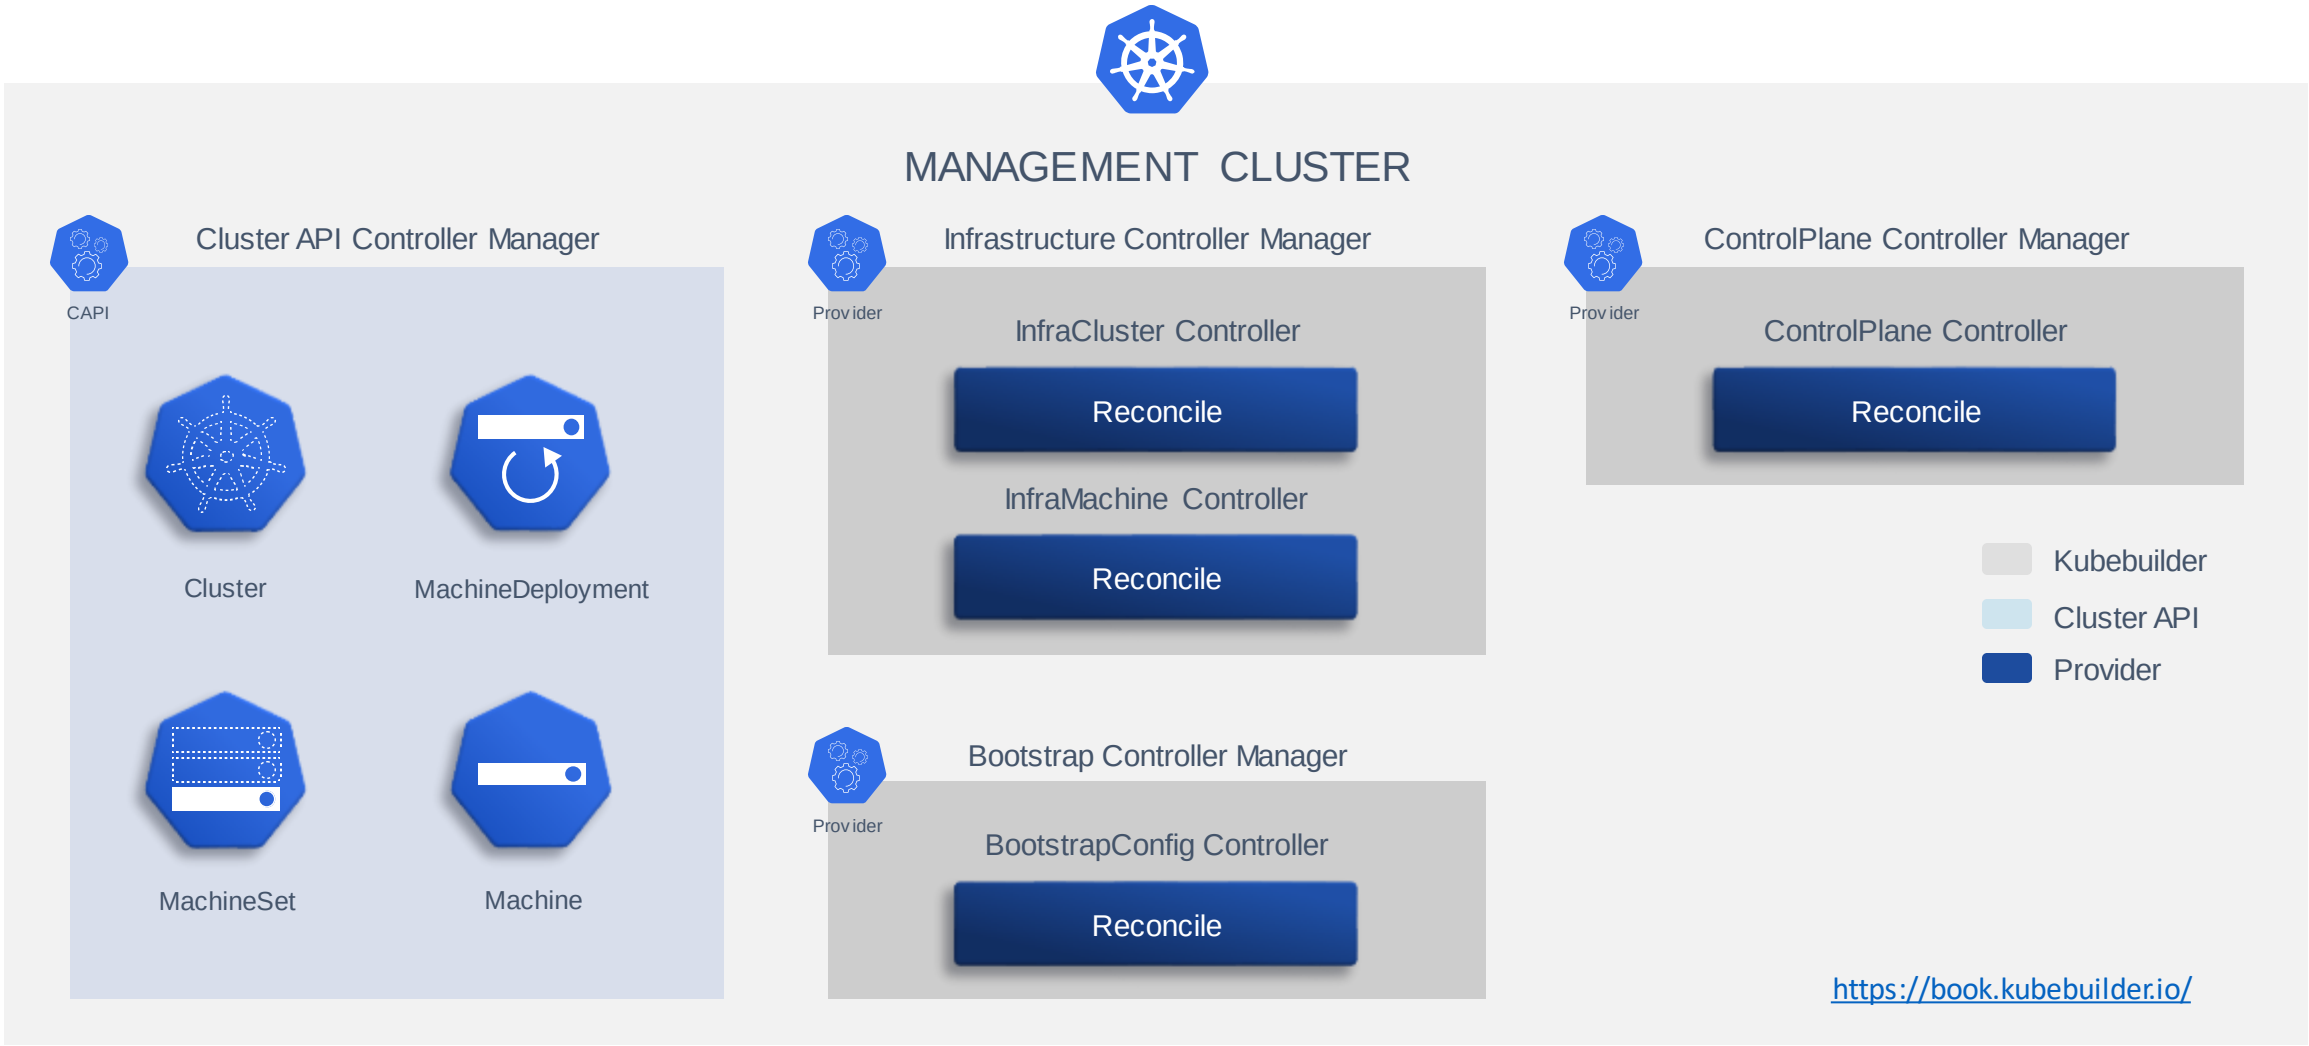

Self-managed clusters

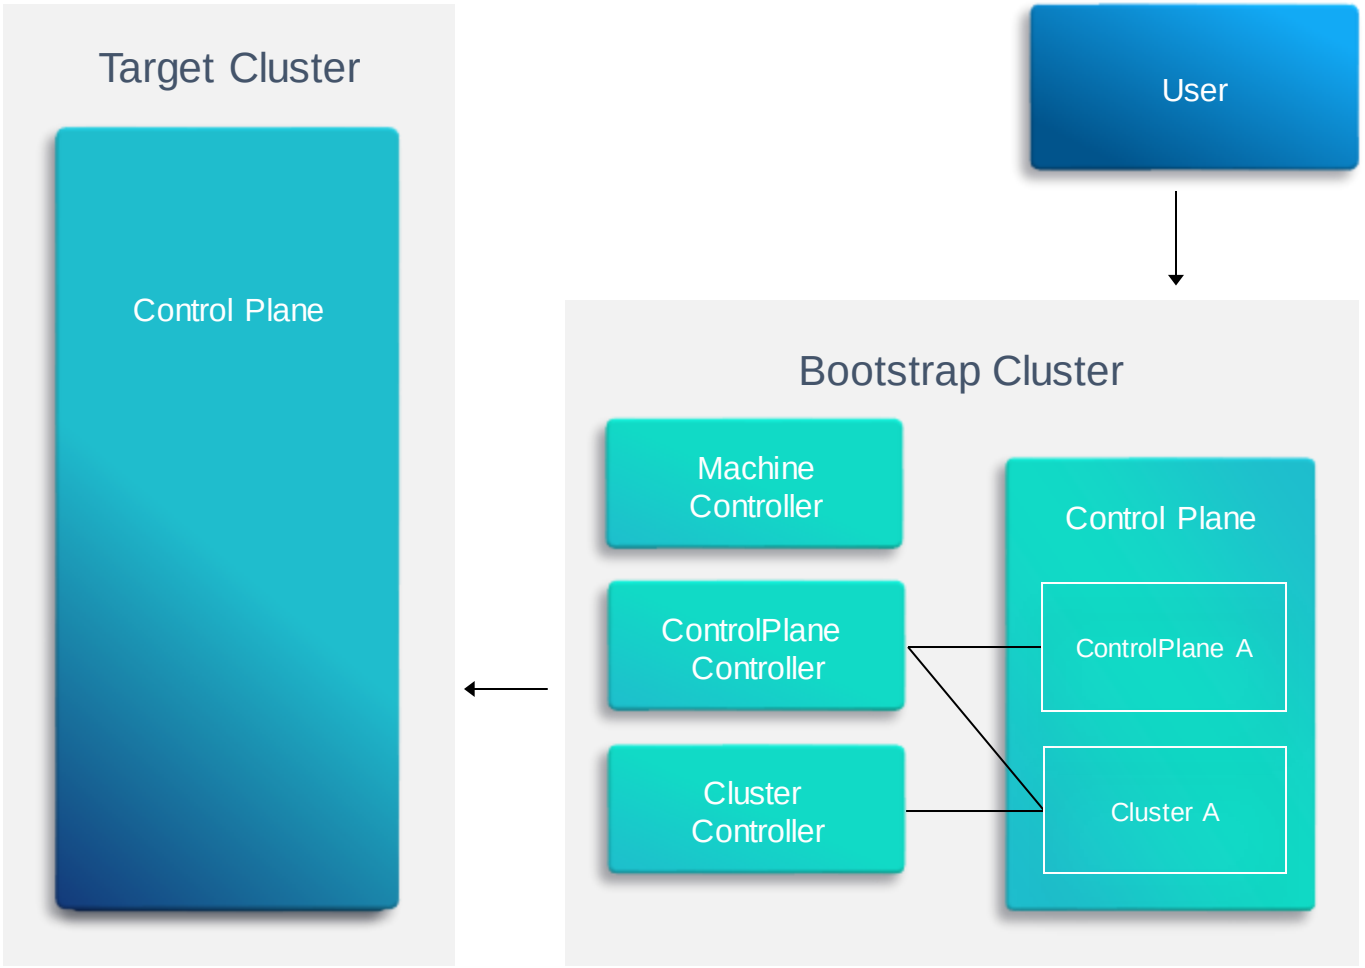

Self-managed clusters

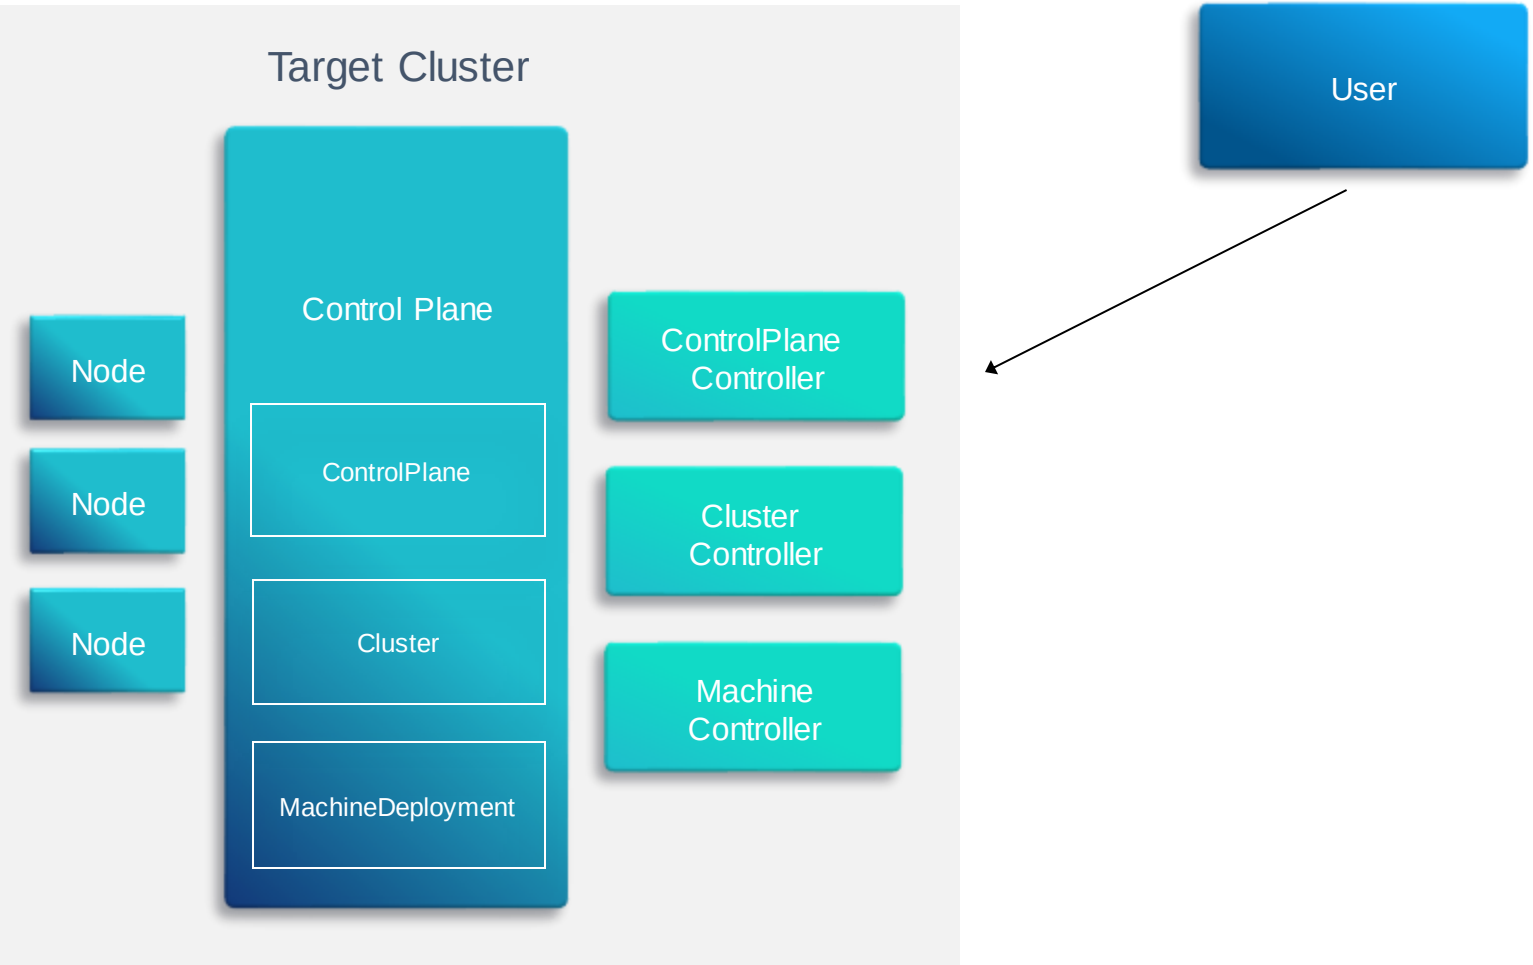

The Spec

Cluster

Cluster-wide configuration

Generic networking concepts like pod and service ranges or DNS domain

Providers can modify and override behavior where needed

```yaml
apiVersion: cluster.x-k8s.io/v1alpha3
kind: Cluster
metadata:
  name: cluster-api-demo
spec:
  clusterNetwork:
    services:
      cidrBlocks: ["10.96.0.0/12"]
    pods:
      cidrBlocks: ["192.168.0.0/16"]
    serviceDomain: "cluster.local”
  infrastructureRef:
    kind: AWSCluster
    apiVersion: infrastructure.cluster.x-
k8s.io/v1alpha3
    name: cluster-api-demo
    namespace: default
  controlPlaneRef:
    kind: KubeadmControlPlane
    apiVersion: controlplane.cluster.x-
k8s.io/v1alpha3
    name: capi-demo-control-plane
```

The Spec

InfraCluster
(eg. AWSCluster, AzureCluster, etc.)

Provider-specific cluster configuration

```
apiVersion: infrastructure.cluster.x-k8s.io/v1alpha3
kind: AWSCluster
metadata:
  name: cluster-api-demo
spec:
  region: us-east-2
  sshKeyName: default
```

The Spec

KubeadmControlPlane

Declarative control plane lifecycle management with Kubeadm

Replicas has the desired number of control plane machines

InfrastructureTemplate provides pluggable provider-specific machine definitions for control plane machines

KubeadmConfig provides means for configuring initialization, cluster and join configuration for control plane machines


```yaml
apiVersion: controlplane.cluster.x-k8s.io/v1alpha3
kind: KubeadmControlPlane
metadata:
  name: capi-demo-control-plane
spec:
  replicas: 3
  infrastructureTemplate:
    kind: AWSMachineTemplate
    apiVersion: infrastructure.cluster.x-k8s.io/v1alpha3
    name: capi-demo-control-plane
  kubeadmConfigSpec:
    initConfiguration:
      ...
    clusterConfiguration:
      ...
    joinConfiguration:
      ...
  version: 1.18.6
```

The Spec

MachineDeployment

Declarative updates for Machines via MachineSets

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

```yaml
apiVersion: cluster.x-k8s.io/v1alpha3
kind: MachineDeployment
metadata:
  name: nodepool-0
  labels: {cluster.k8s.io/cluster-name: cluster-api-demo}
spec:
  replicas: 3
  selector:
    matchLabels:
      cluster.x-k8s.io/cluster-name: cluster-api-demo
      nodepool: nodepool-0
  template:
    metadata:
      labels:
        cluster.x-k8s.io/cluster-name: cluster-api-demo
        nodepool: nodepool-0
    spec:
      version: v1.18.6
      bootstrap:
        configRef:
          name: nodepool-0
          apiVersion: bootstrap.cluster.x-k8s.io/v1alpha3
          kind: KubeadmConfigTemplate
      infrastructureRef:
        name: nodepool-0
        apiVersion: infrastructure.cluster.x-k8s.io/v1alpha3
        kind: AWSMachineTemplate
```

Rael Garcia from go.rael.dev/kubecap19-eu

The Spec

Machine

Configuration for a **specific machine**

Spec has the desired kubelet version

Providers can **modify and override behavior where needed**

```
apiVersion: cluster.x-k8s.io/v1alpha3
kind: Machine
metadata:
  name: capi-demo-machine
  labels:
    cluster.k8s.io/cluster-name: cluster-api-demo
spec:
 bootstrap:
    configRef:
      kind: KubeadmConfig
      apiVersion: bootstrap.cluster.x-k8s.io/v1alpha3
      namespace: default
      name: capi-demo-machine
  infrastructureRef:
    kind: AWSMachine
    apiVersion: infrastructure.cluster.x-k8s.io/v1alpha3
    namespace: default
    name: capi-demo-machine
  version: "v1.18.6"
```

The Spec

InfraMachine
(AWSMachine, AzureMachine)

Infrastructure provider-specific machine configuration

```
apiVersion: infrastructure.cluster.x-k8s.io/v1alpha3
kind: AWSMachine
metadata:
  name: capi-demo-machine
  labels:
    cluster.k8s.io/cluster-name: cluster-api-demo
spec:
  instanceType: m5.large
  iamInstanceProfile: "controllers.cluster-api-provider-
aws.sigs.k8s.io"
  sshKeyName: default
```

Demo

What's new in 0.3?

KubeCon CloudNativeCon Europe 2020 Virtual

Control Plane

Control Plane managed as a single entity instead of individual machines

Upgrade of the Control Plane nodes

Machine Pools

API to create and manage groups of machines, e.g. Azure VM Scale Sets, AWS ASGs

Machine Health Checks

Allows for remediation of machines which have become unhealthy, based on node conditions

Testing

Testing framework to enable infrastructure and bootstrap providers to validate Cluster API behaviors

Failure Domains

Allows spreading machines out to reduce the risk of a target cluster failing due to a domain outage

Clusterctl

Rewrite of initial CLI to manage getting started with Cluster API

Conditions

Provides more detailed information on cluster and machine state

Cluster Autoscaler

Initial implementation of a Cluster API provider for cluster-autoscaler

What's next?

KubeCon CloudNativeCon Europe 2020 Virtual

- Working towards stabilizing Cluster API to reach beta status
- UX improvements
- Better management of cluster addons
- Detection of machine failures prior to cluster join
- Load Balancer Providers: Pluggable load balancers for the Kubernetes API server for on-premise environments
- Move MachinePool out of experimental
- Windows nodes
- Ability to provision infrastructure across different accounts for different users (AWS and Azure)

https://cluster-api.sigs.k8s.io/roadmap.html

Want to get involved?

KubeCon CloudNativeCon Europe 2020 Virtual

- https://github.com/kubernetes-sigs/cluster-api
- Join **kubernetes-sig-cluster-lifecycle**
- Weekly meeting on **Wed @ 10:00AM** PT (5:00PM UTC)
- K8s Slack: **#cluster-api**

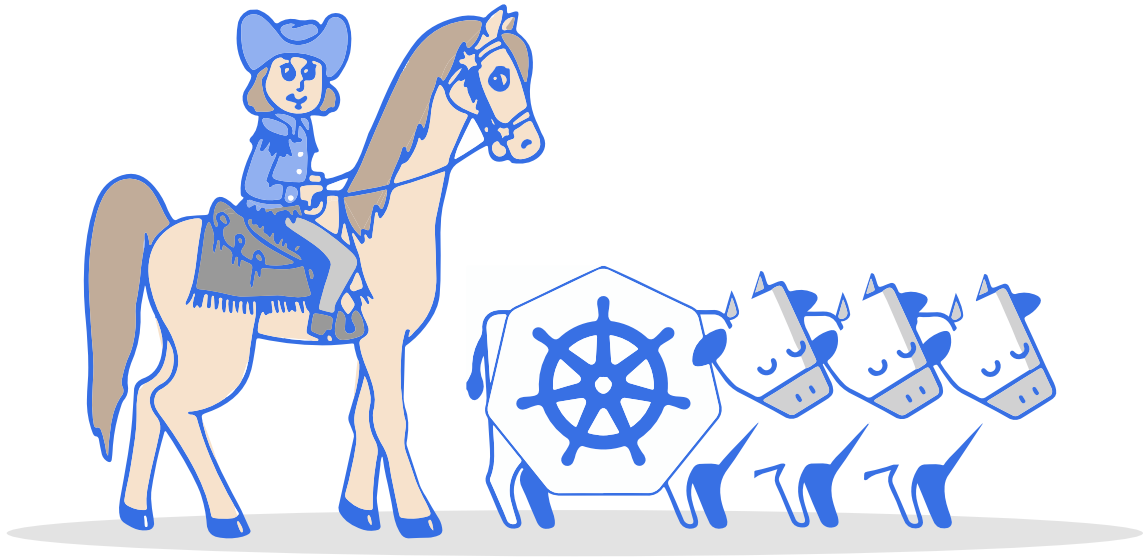


KubeCon
CloudNativeCon
Europe 2020
Virtual
HELM
KEEP CLOUD NATIVE
CONNECTED