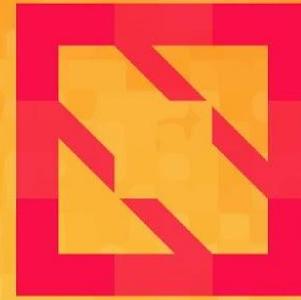




KubeCon



CloudNativeCon

North America 2019





kubeadm

Deep Dive



KubeCon



CloudNativeCon

North America 2019





Yago Nobre

Software Engineer, Nubank

@yagoyns



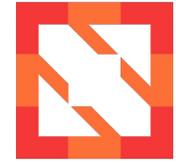
Fabrizio Pandini

Staff Member, VMware

@fabriziopandini



KubeCon



CloudNativeCon

North America 2019

Agenda

- Intro
- 2019 Highlights
- 2020 Roadmap



KubeCon



CloudNativeCon

North America 2019



Kubeadm is a node bootstrapper



KubeCon



CloudNativeCon

North America 2019

Someone or something should provide the machines

kubeadm creates a Kubernetes node on the machine

How kubeadm does node bootstrap?



KubeCon



CloudNativeCon

North America 2019

kubeadm's task is to set up a ***best-practice, conformant cluster***.

The user experience should be ***simple***.

kubeadm init, join ... done!

The ***scope is intentionally limited***.

Agnostic to infrastructure, CRI, CNI & CSI.

A SIG Cluster lifecycle project



KubeCon



CloudNativeCon

North America 2019

SIG Cluster Lifecycle's objective is to simplify creation, configuration, upgrade, downgrade, and teardown of Kubernetes clusters and their components.

Part of a composable solution

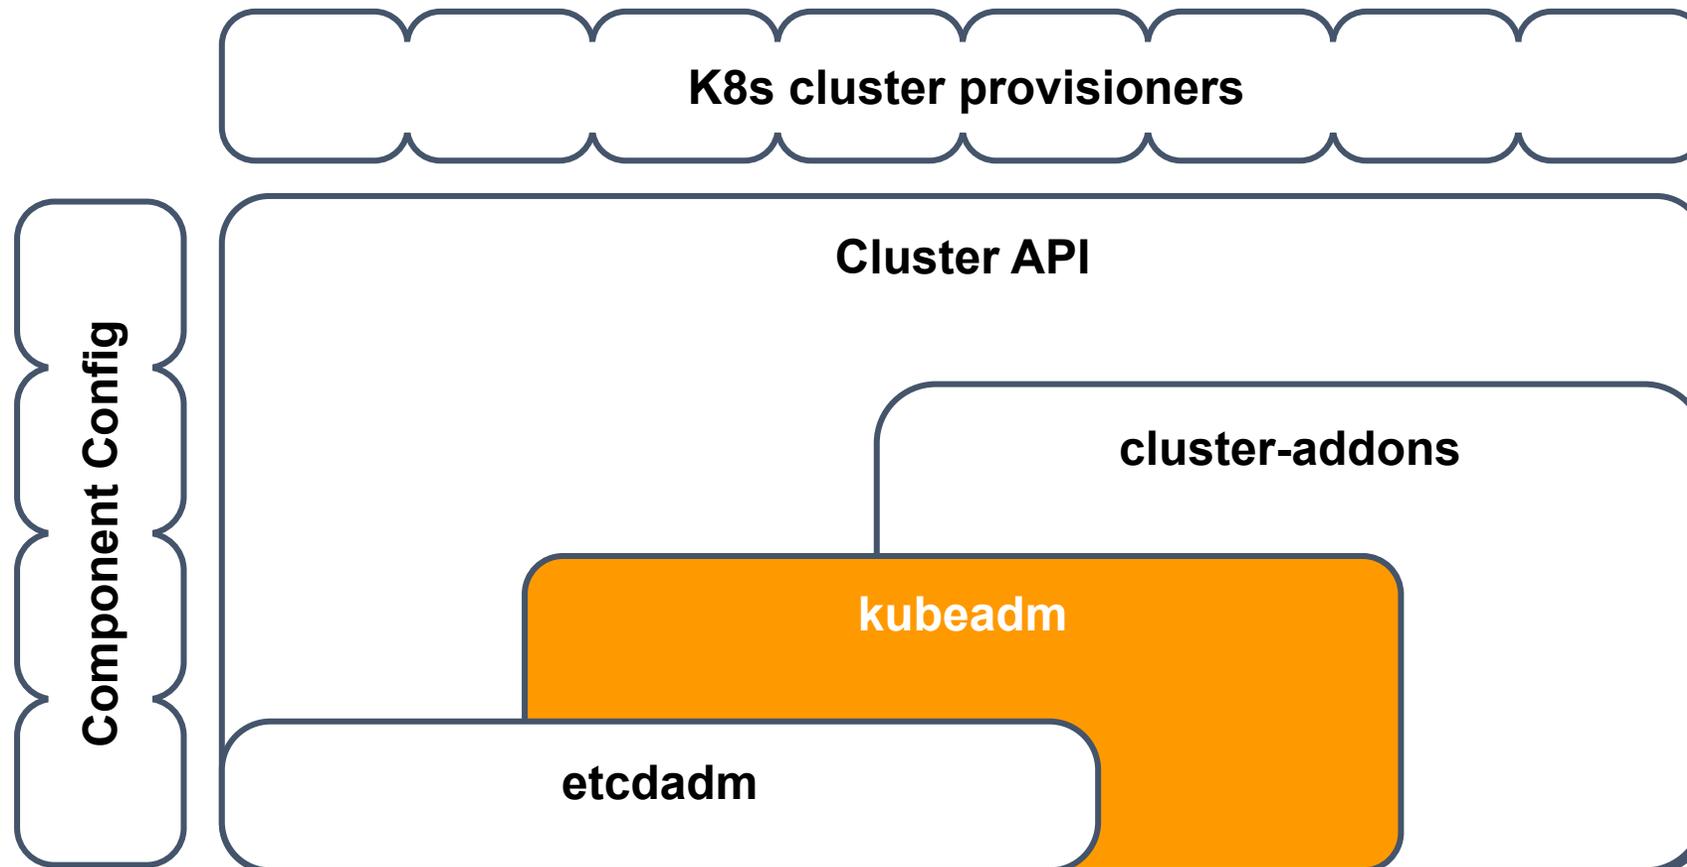


KubeCon



CloudNativeCon

North America 2019





KubeCon



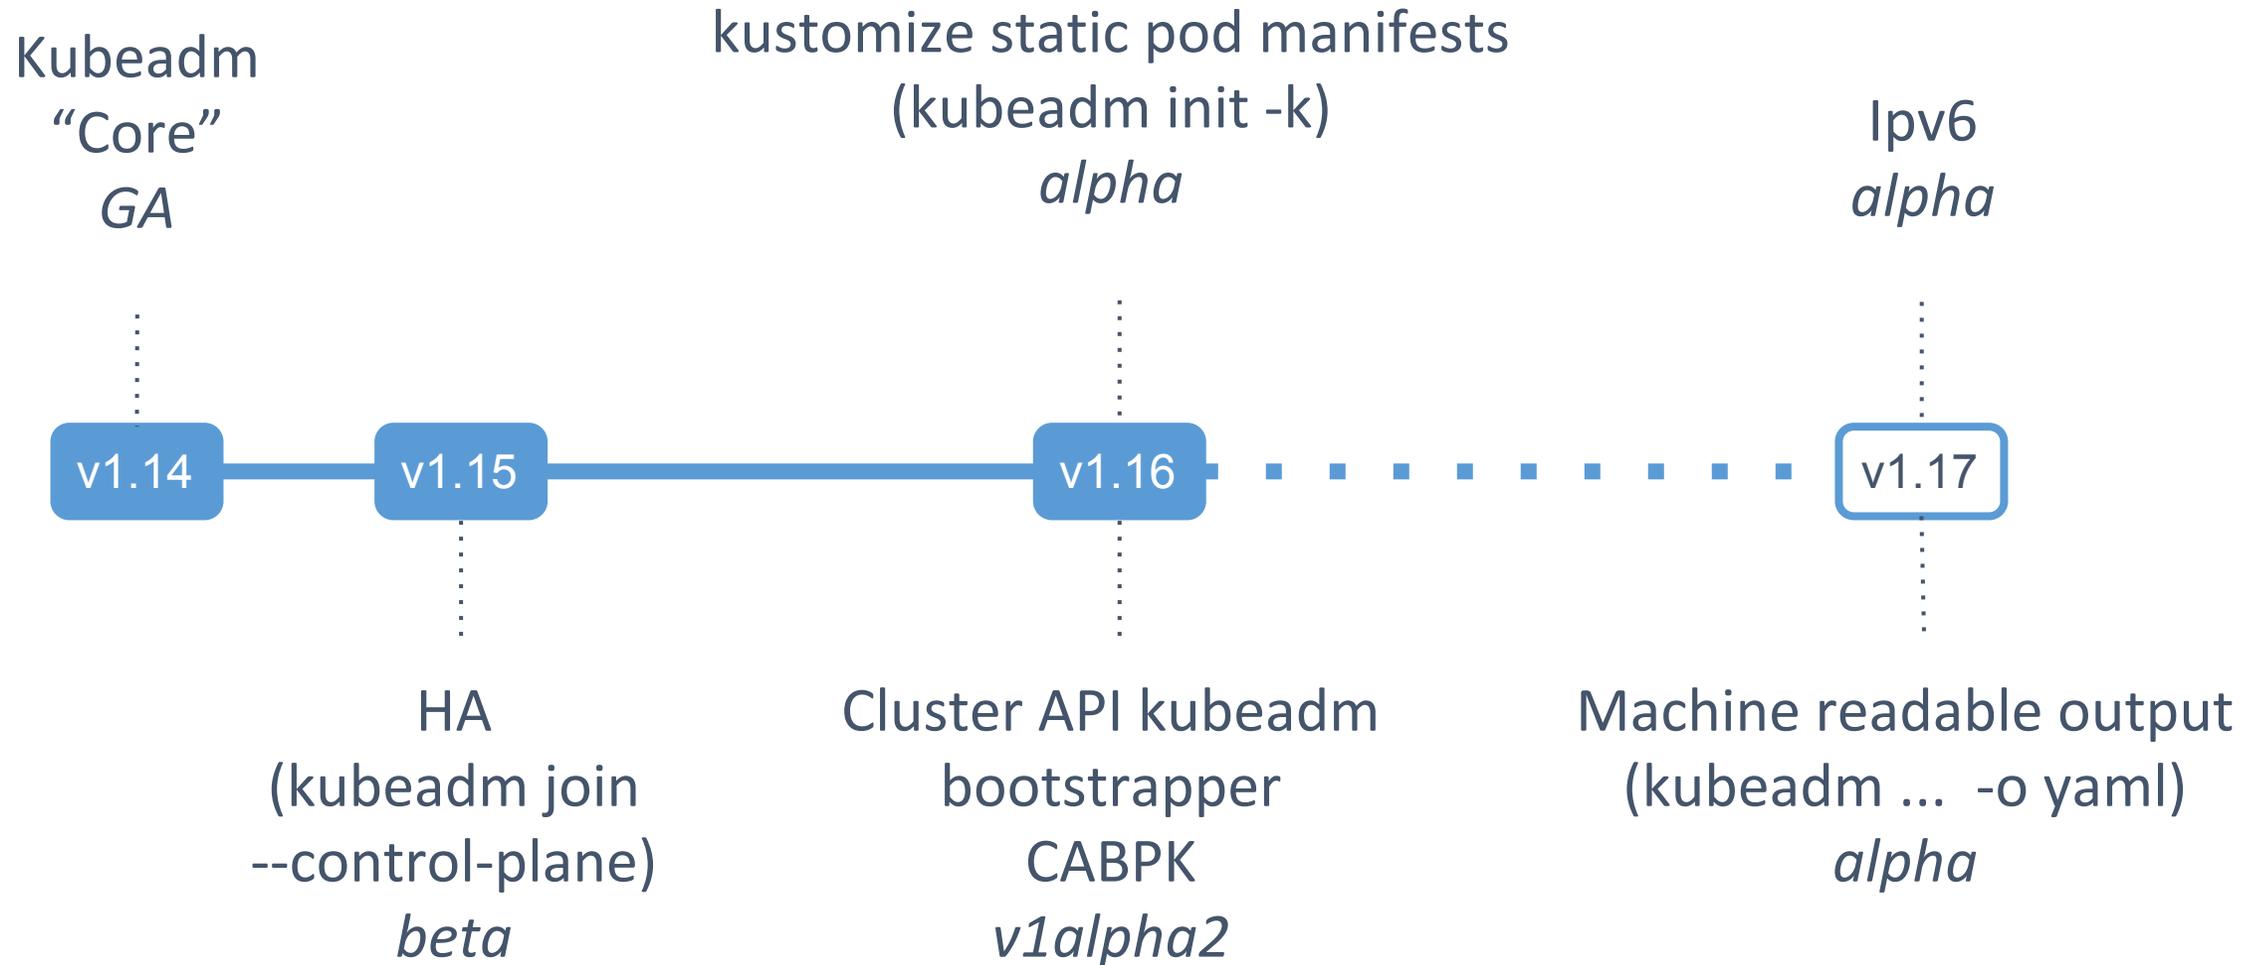
CloudNativeCon

North America 2019

2019 Highlights



2019 Main Achievements and WIPs



kustomize support



KubeCon



CloudNativeCon

North America 2019

Thanks to kustomize support, no more need of manual changes in static pod manifests!

Create a kustomize patch for static pod manifests:

```
$ cat <<EOF > /tmp/patch1.yaml
apiVersion: v1
kind: Pod
metadata:
  name: kube-apiserver
  namespace: kube-system
spec:
  containers:
  - name: kube-apiserver
    livenessProbe:
      initialDelaySeconds: 30
      timeoutSeconds: 30
EOF
```



Make kubeadm init, join or upgrade to apply the patch:

```
$ kubeadm init -k /kinder/kustomize
$ kubeadm join -k /kinder/kustomize
$ kubeadm upgrade apply -k /kinder/kustomize
$ kubeadm upgrade node -k /kinder/kustomize
```

Machine readable output



KubeCon



CloudNativeCon

North America 2019

Thanks to machine readable output, integrating kubeadm with automation tools is now simpler for an initial set of commands!

Before, only:

```
$ kubeadm token list
```

```
TOKEN          TTL          EXPIRES          USAGES          ...  
7vg8cr.pks5g06s84aisb27 <invalid> 2019-06-05T17:13:55+03:00 authentication,signing ...
```

Now also:

```
$ kubeadm token list -o json
```

```
$ kubeadm token list -o yaml
```

```
$ kubeadm token list -o "go-template={{println .token .description .groups}}"
```

```
$ kubeadm token list -o "jsonpath={.token} {.description}{'\n'}}"
```

help wanted

Cluster API kubeadm bootstrapper



KubeCon



CloudNativeCon

North America 2019

The Cluster API is a Kubernetes project to bring declarative, Kubernetes-style APIs to cluster creation, configuration, and management. *And it uses kubeadm!*

The entire kubeadm community can benefit from this, because kubeadm now is battle tested on several cloud providers and on clusters with a much higher number of nodes.

Behind the scenes, every day ...



KubeCon



CloudNativeCon

North America 2019

104

E2E Tests

520

Nodes

768

init/join/upgrade etc.

15.120

Conformance
unit tests

Many, many, E2E tests
for higher level tools using
kubeadm:

- minikube
- kind
- kubespray
- Cluster API

A lot of priceless
users feedback

help wanted

good first issue



KubeCon



CloudNativeCon

North America 2019

2020 Roadmap



Your opinion DOES MATTER.

This is a Roadmap from the community to the community.

File issues, provide feedback, request new features.



New roadmap, old mantra...



KubeCon



CloudNativeCon

North America 2019

Be Simple

for the 80% of the use cases

Be Extensible

for the remaining 20% of the use cases

P0 - Simplify Day 2 Operations



KubeCon



CloudNativeCon

North America 2019

As of today, all the day 2 operations require users to execute kubectl commands on different machines, like e.g. upgrades

This approach worked well so far, but now it is showing its limits.
It is time to rethink day 2 operations in kubectl

RFE: Rethinking day 2 with kubeadm



KubeCon



CloudNativeCon

North America 2019

```
kubectl apply -f cluster-desired-state.yaml
```

The kubeadm operator prototype aims to encode operational best practices in a controller and bring declarative, Kubernetes-style APIs to kubeadm day 2 operations

priority/critical-urgent

help wanted

P1 - Extensibility for addons



KubeCon



CloudNativeCon

North America 2019

kubeadm installs kube-proxy and CoreDNS in order to get a working, conformant cluster.

Today a growing number of users have more sophisticated needs, and so we want to make it easier for the kubeadm users to get access to a richer set of addons (e.g. local DNS cache, kube-router)

RFE: The Cluster addons project



KubeCon



CloudNativeCon

North America 2019

kubeadm is not going to solve this problem by its own.

We are going to integrate with the Cluster Addons project, that is a SCL effort to create a curated list of addons, that, due to their role, are tied to or dependent on the configuration of other cluster components.

priority/important-soon

help wanted

P3 - Extend etcd support



KubeCon



CloudNativeCon

North America 2019

kubeadm by default installs etcd in order to get a working cluster.

kubeadm offers very limited features for etcd management, so we
we want to make it easier for kubeadm users to get access to
better etcd tooling.

RFE: The etcdadm addon project



KubeCon



CloudNativeCon

North America 2019

Also in this case,

kubeadm is not going to solve this problem by its own.

We are going to integrate with the etcdadm project, that is a SCL effort to create a command-line tool for operating an etcd cluster.

priority/important-soon

help wanted

... and some housekeeping



KubeCon



CloudNativeCon

North America 2019

Move kubeadm out of tree

As of today kubeadm code lives in a folder under the main Kubernetes GitHub repository, but we are now requested to move kubeadm out of tree

This step enables a whole new set of scenarios for kubeadm (e.g. vendoring)

priority/critical-urgent

help wanted

Code cleanup/refactor

Current Kubeadm code organization reflects GA goals and objectives, but it shows limits for the new roadmap

It is time to pay down some kubeadm technical debts

priority/backlog

help wanted

Reach out if you are interested !

Any skill, any expertise.

feature request, feedback, issue, code, test, documentation...



New contributors onboarding sessions soon !

Code walk-through, developing kubeadm 101,
backlog grooming,
office hours ...



The “kubeadm” team

Smaller core group of active maintainers

Tim, Lubomir, Ross, Jason, Fabrizio, Yassine, Chuck (VMWare)

Marek, Rafael (SUSE), Alex, Ed (Intel)

Yago (NuBank), Lucas (Independent/Weaveworks)

SIG Cluster Lifecycle

100s of contributors across several companies

We’re working on growing the contributor/reviewers pool



Look for “good first issue”, “help wanted” and “sig/cluster-lifecycle” labeled issues in our repositories

[Contributing to kubeadm](#)

[SIG Cluster Lifecycle New Contributor Onboarding](#)

Attend our Zoom meetings / be around on Slack

Join [#sig-cluster-lifecycle](#), [#kubeadm](#) channels

Weekly “Office Hours” meetings for kubeadm and Cluster API, bi-weekly for other projects

Full list of SIG meetings and links to minutes and recordings can be found on [SIG page](#)

Meetings are recorded: follow the [SIG Cluster Lifecycle YouTube playlist](#)





KubeCon



CloudNativeCon

North America 2019



kubeadm

Deep Dive

*Thank you for attending!
Q & A time!*

