



**KubeCon**



**CloudNativeCon**

— North America 2018 —

## Rightsize your Pods with Vertical Pod Autoscaling

Beata Skiba <[bskiba@google.com](mailto:bskiba@google.com)>  
Software Engineer @ Google, working on Kubernetes Autoscaling



# Life of a resource request



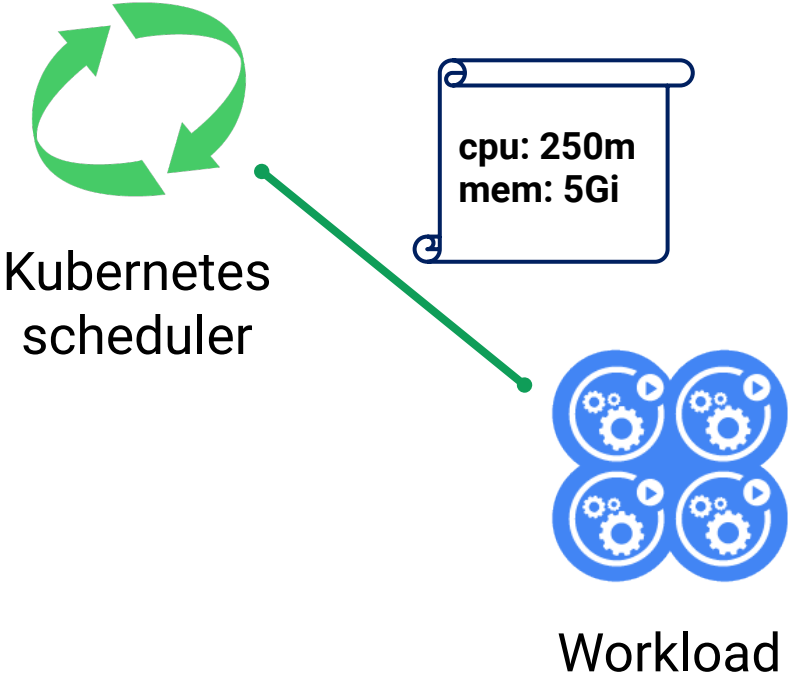
# Deploying is easy, right?

```
apiVersion: apps/v1
kind: Deployment
...
template:
  metadata:
    labels:
      app: worker
  spec:
    containers:
      - name: worker
        resources:
          requests:
            memory: "100M"
            cpu: "250m"
```

# Well...

```
apiVersion: apps/v1
kind: Deployment
...
template:
  metadata:
    labels:
      app: worker
  spec:
    containers:
      - name: worker
        resources:
          requests:
            memory: ??
            cpu: ??
```

# Resource request is a contract



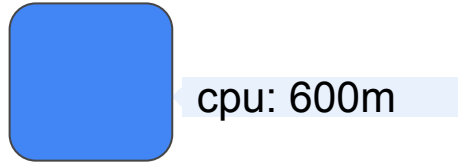
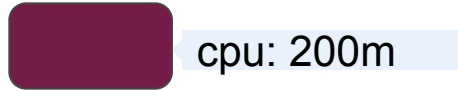
# Requests in action: Case 1

# Will it fit?





Kubernetes  
scheduler

## Pods



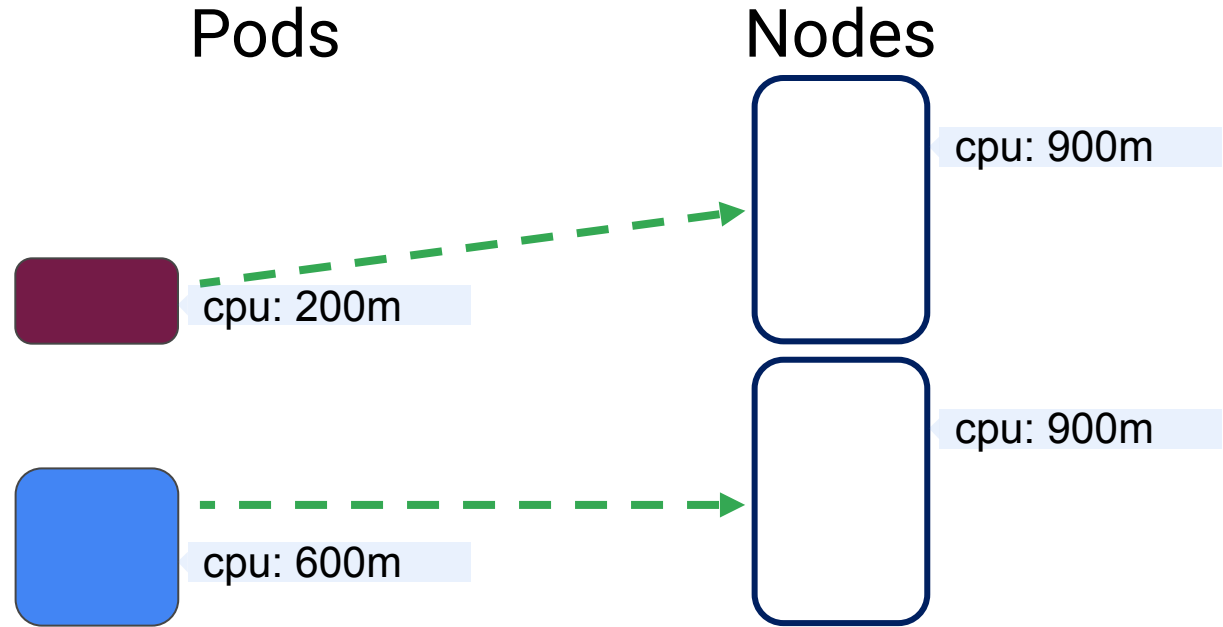
## Nodes



# Will it fit?



Kubernetes  
scheduler






# Will it fit?




Kubernetes  
scheduler

Pods




usage: 500m

Nodes



cpu: 900m



cpu: 900m

# Will it fit?




Kubernetes  
scheduler

Pods




cpu: 600m


Nodes



cpu: 900m



cpu: 900m




# Will it fit?



Kubernetes  
scheduler

Pods

Nodes

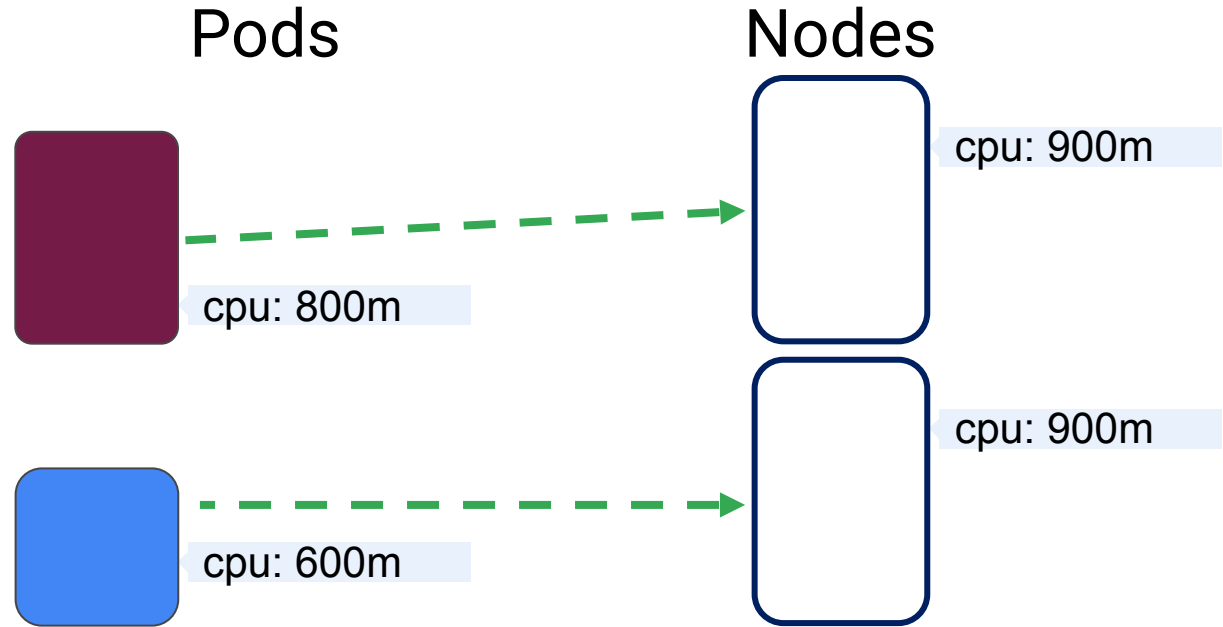


## Requests in action: Case 2

# Will it fit?



Kubernetes  
scheduler




# Will it fit?



Kubernetes scheduler

Pods

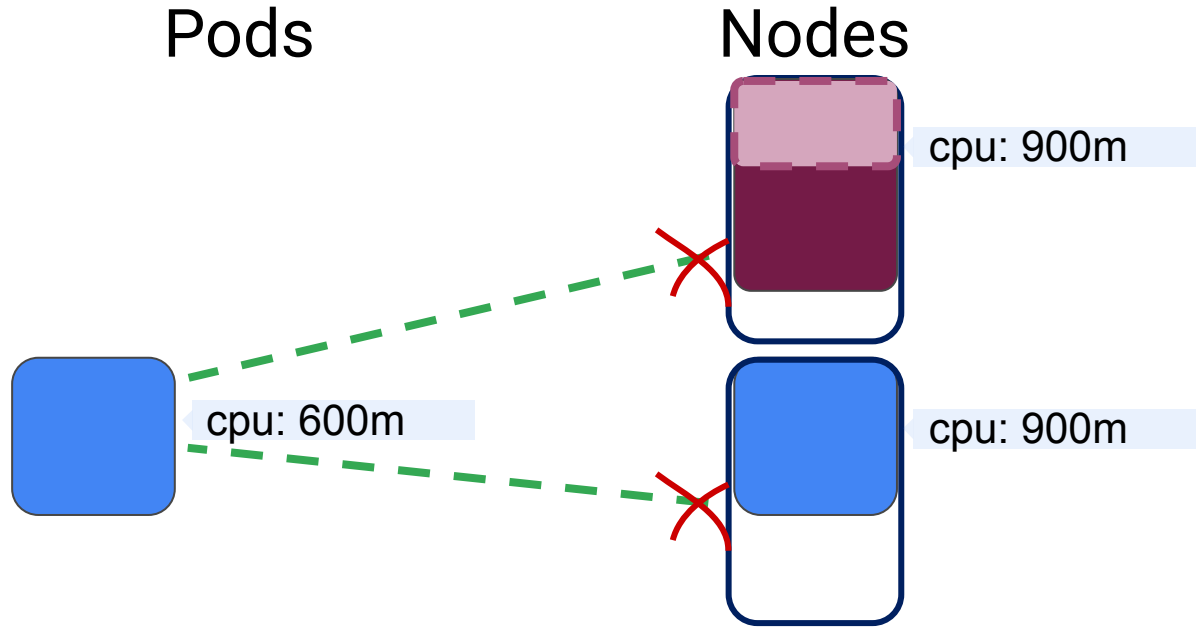
Nodes



# Will it fit?



Kubernetes  
scheduler



# Will it fit?



Kubernetes  
scheduler

## Pods




cpu: 600m

## Nodes



cpu: 900m



cpu: 900m



# Will it fit?




Kubernetes  
scheduler




Cluster  
Autoscaler

## Pods




cpu: 600m

## Nodes

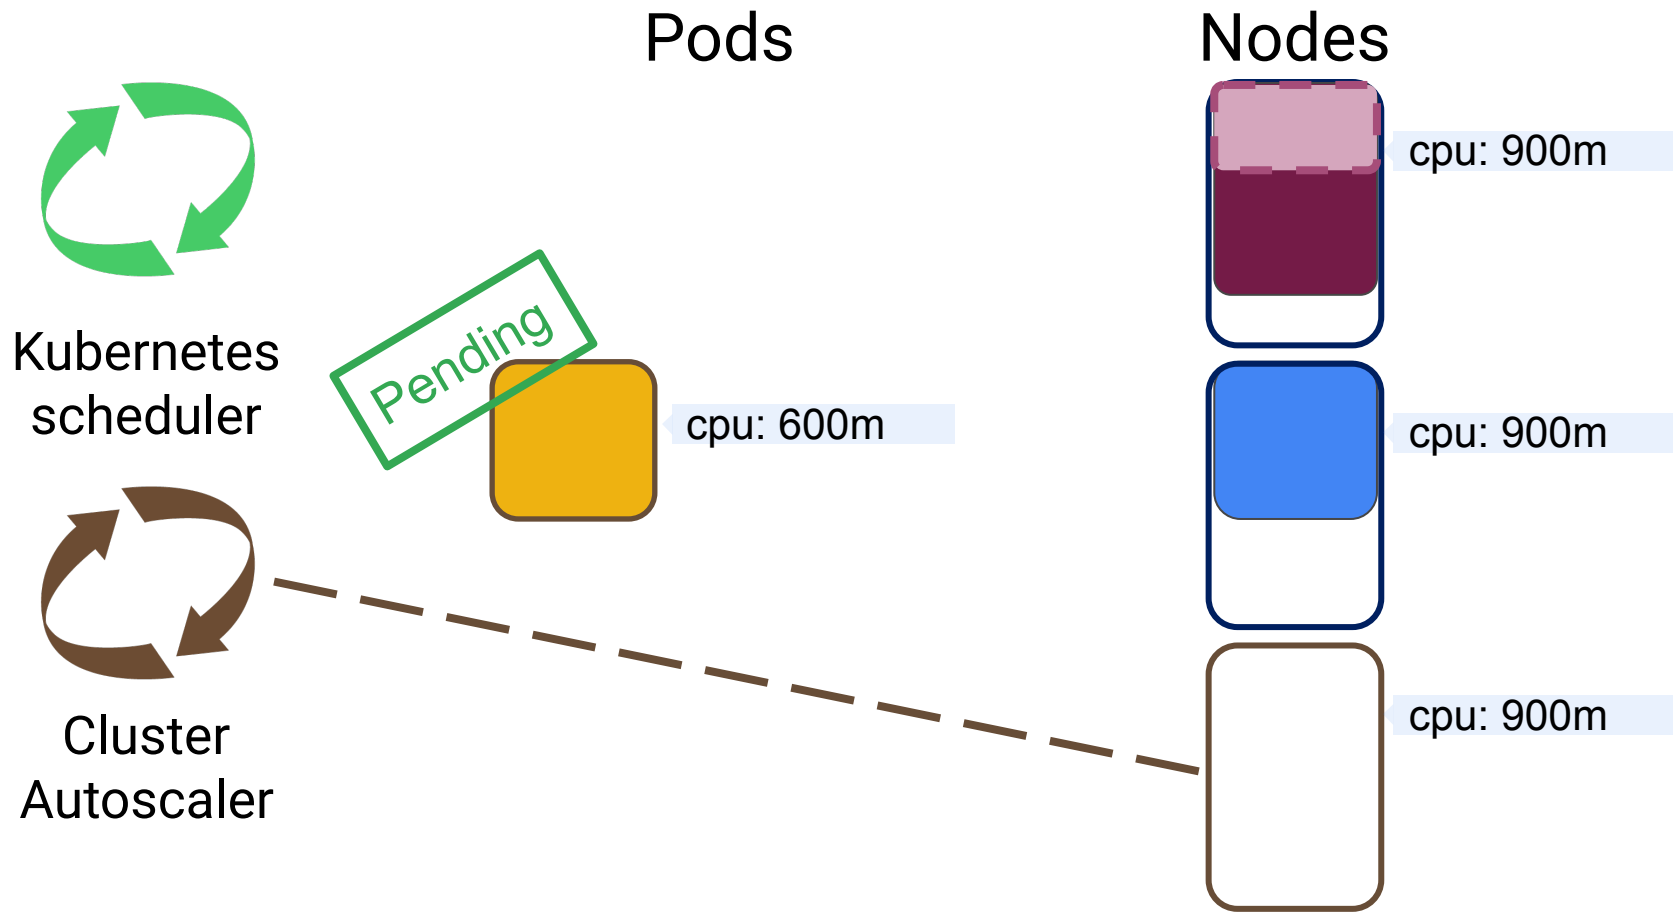


cpu: 900m



cpu: 900m


# Will it fit?



# Will it fit?



Kubernetes  
scheduler



Cluster  
Autoscaler

Pods




cpu: 600m

Nodes




cpu: 900m



cpu: 900m




cpu: 900m



# HELP! My cluster is on fire

## Too little

- CPU starvation
- Out Of Memory errors
- Workload eviction



## Too much

- Resource waste



# How to get it right?


- ? Canarying
- ? Trial and error
- ? Tea leaves

But....



# Where to get the numbers from?

- ⊘ Time, effort & resources
- ⊘ Canary may differ from prod
- ⊘ Resource needs will change over time



# VPA to the rescue



# VPA to the rescue

Observes  
usage




Recommends  
resources



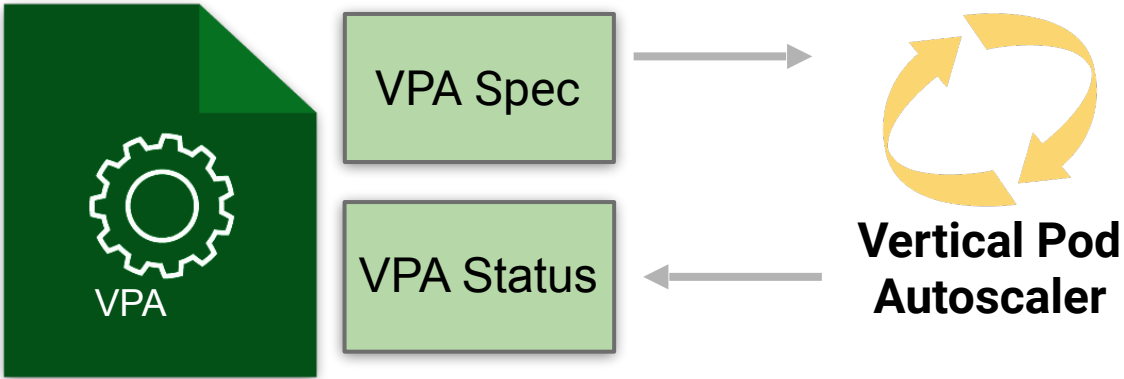
Updates  
resources

EXPERIMENTAL





# It's VPA objects all the way down



VPA Object



# What the Spec?

```
apiVersion: "autoscaling.k8s.io/v1beta1"
kind: VerticalPodAutoscaler
metadata:
  name: worker-vpa
spec:
  selector:
    matchLabels:
      app: worker
  updatePolicy:
    updateMode: "Off"
  resourcePolicy:
    containerPolicies:
      - containerName: "*"
        maxAllowed:
          memory: 5Gi
```

# Scale WHAT?

```
apiVersion: "autoscaling.k8s.io/v1beta1"
kind: VerticalPodAutoscaler
metadata:
  name: worker-vpa
spec:
  selector:
    matchLabels:
      app: worker
  updatePolicy:
    updateMode: "Off"
  resourcePolicy:
    containerPolicies:
      - containerName: "*"
      maxAllowed:
        memory: 5Gi
```

# Scale HOW?

```
apiVersion: "autoscaling.k8s.io/v1beta1"
kind: VerticalPodAutoscaler
metadata:
  name: worker-vpa
spec:
  selector:
    matchLabels:
      app: worker
  updatePolicy:
    updateMode: "Off"
  resourcePolicy:
    containerPolicies:
      - containerName: "*"
        maxAllowed:
          memory: 5Gi
```

Off

Initial

Auto

EXPERIMENTAL

# Scale HOW?

```
apiVersion: "autoscaling.k8s.io/v1beta1"
kind: VerticalPodAutoscaler
metadata:
  name: worker-vpa
spec:
  selector:
    matchLabels:
      app: worker
  updatePolicy:
    updateMode: "Off"
  resourcePolicy:
    containerPolicies:
      - containerName: "*"
        maxAllowed:
          memory: 5Gi
```

# VPA status

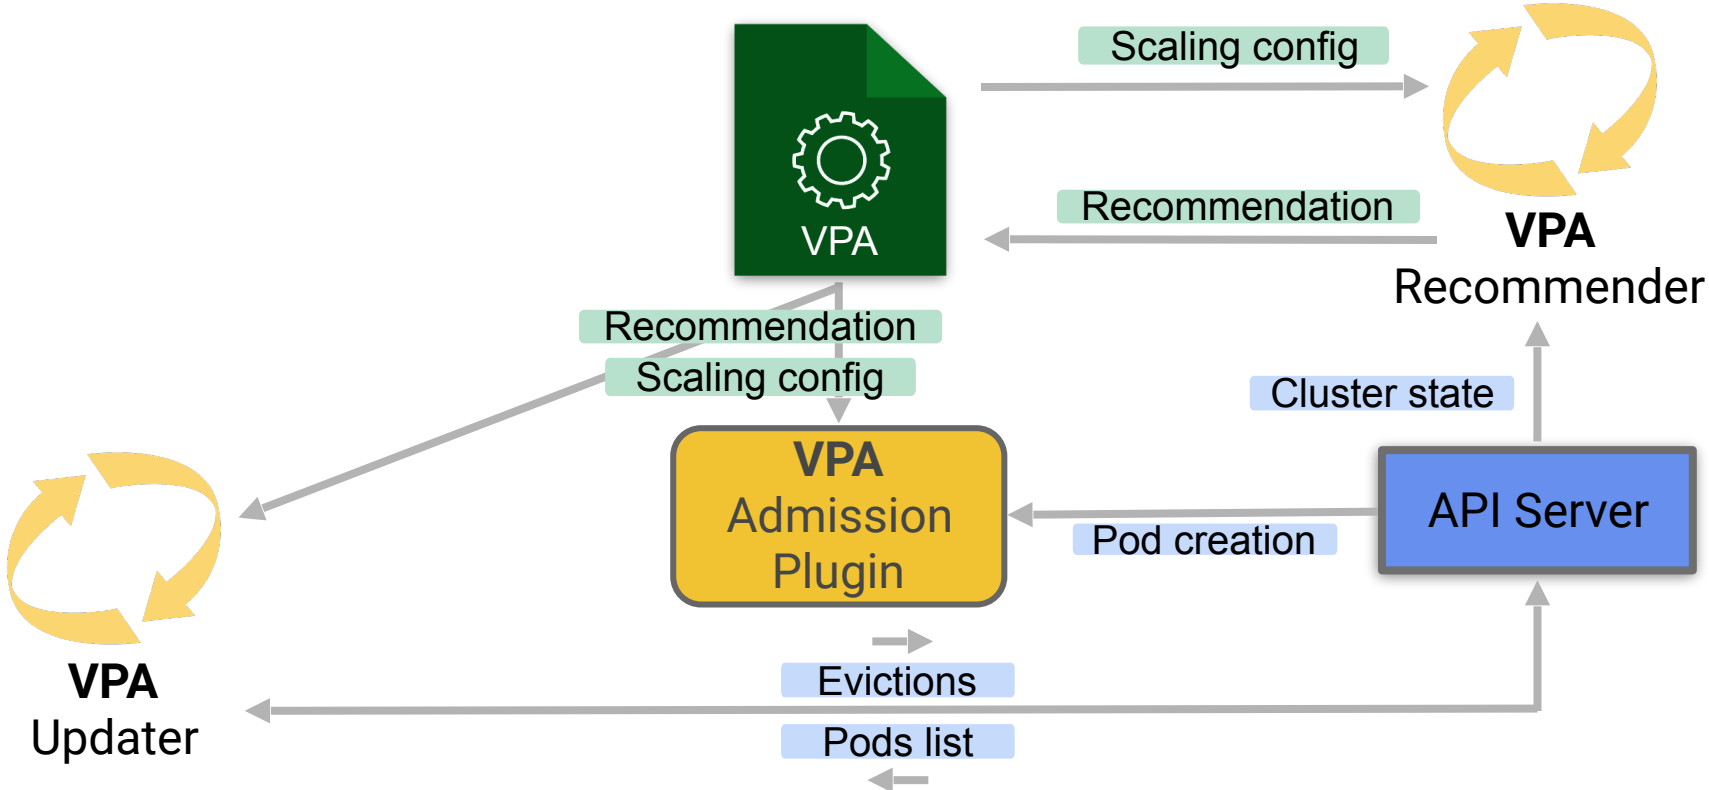
```
status:
  recommendation:
    containerRecommendations:
      - containerName: worker
        lowerBound:
          cpu: 584m
          memory: 218M
          target:
            cpu: 587m
            memory: 262M
        upperBound:
          cpu: 2088m
          memory: 387M
```

# VPA under the hood

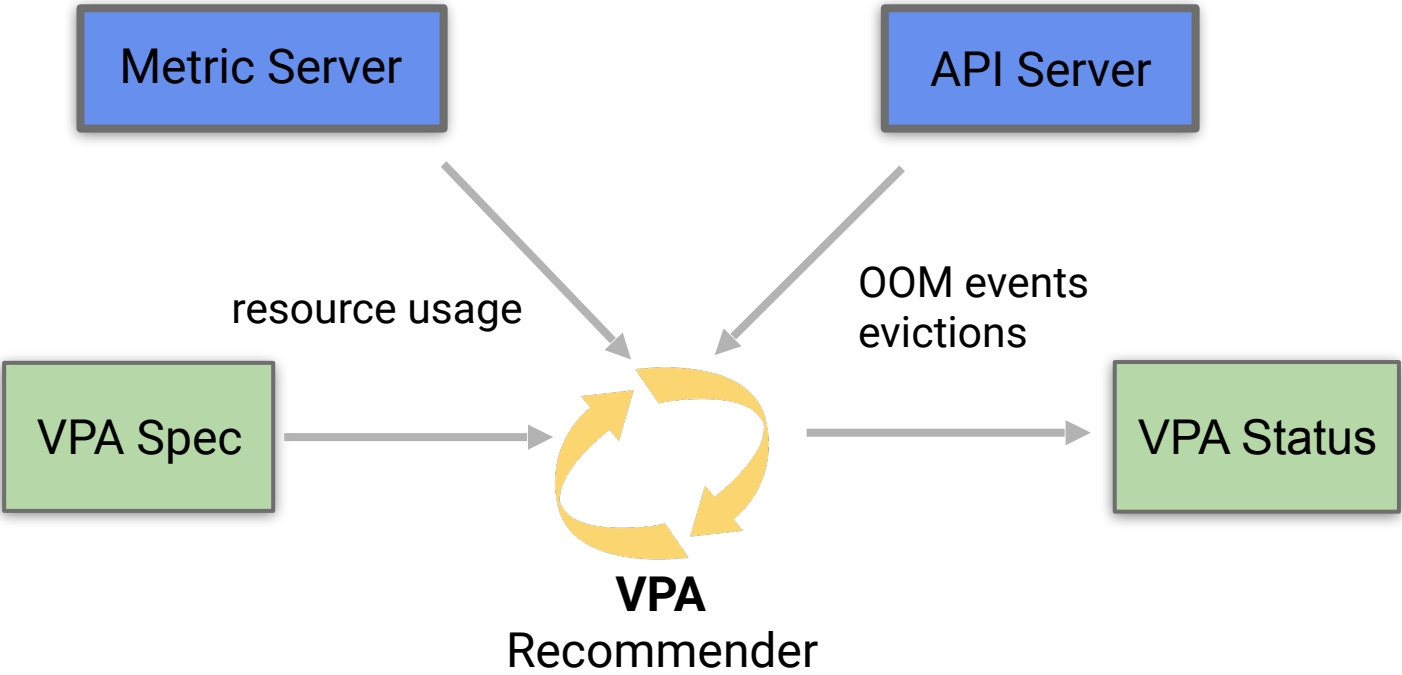




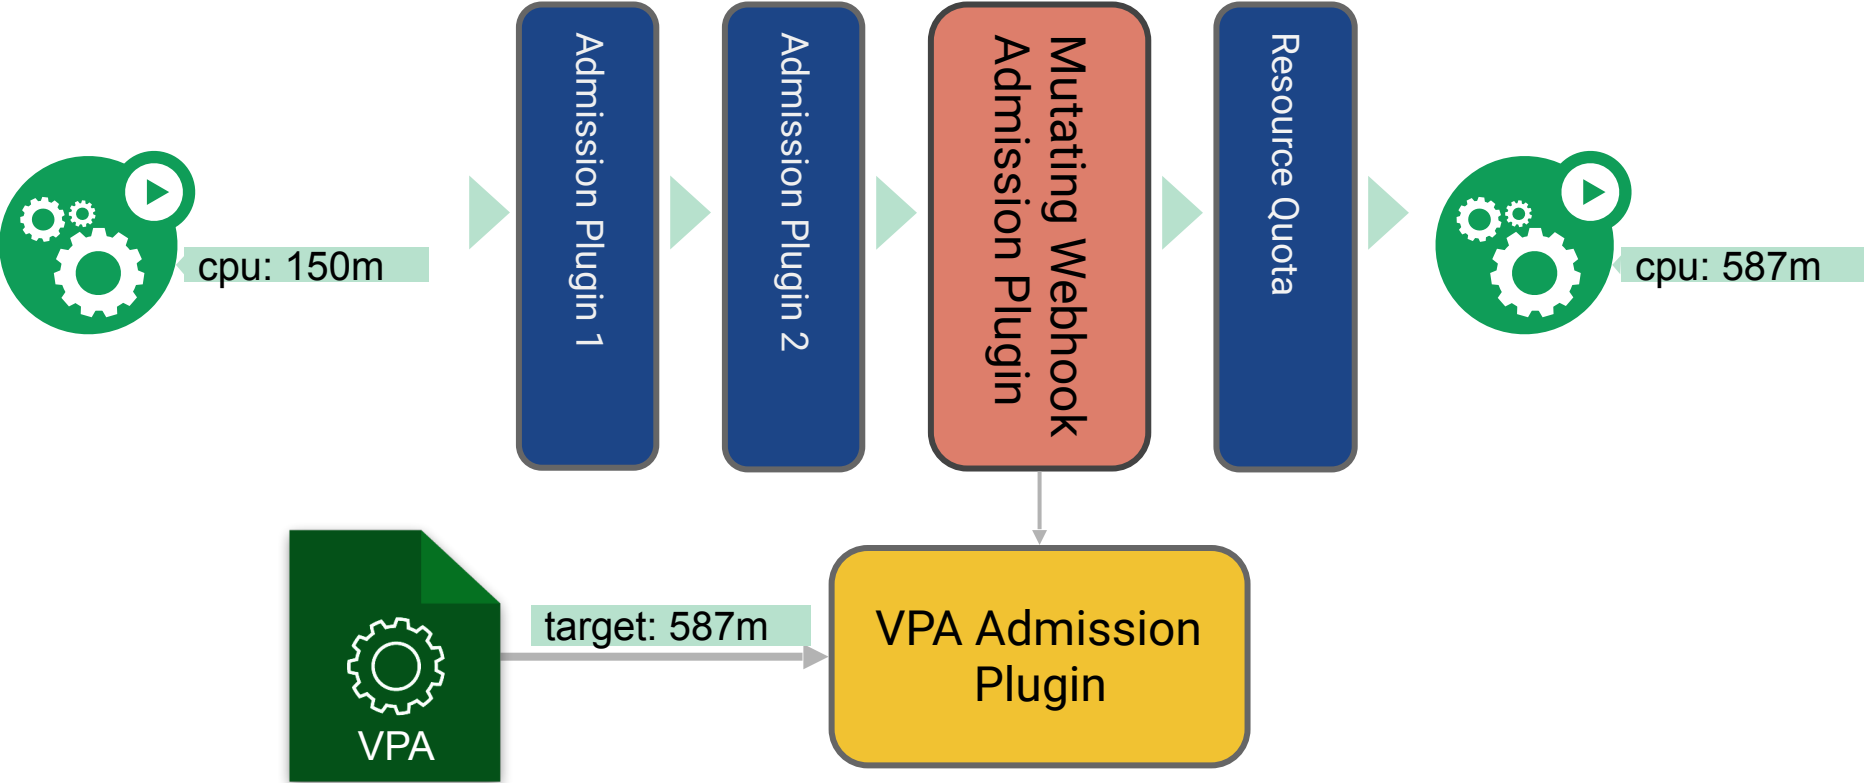
# Bits and pieces



# Watch and recommend

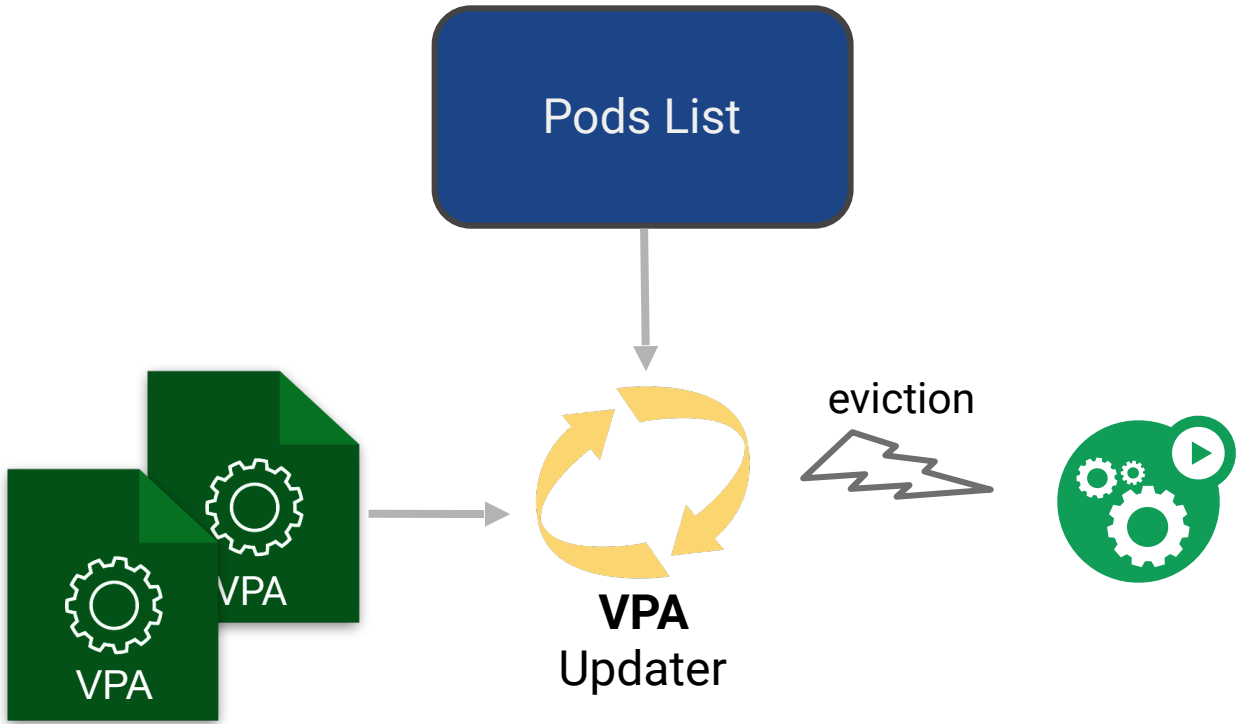


# Apply recommendations



# Restart running workloads

EXPERIMENTAL



# Project status

Status

&

Plans

2018/04 VPA Alpha

---

2018/12 VPA Beta

---

Waiting for your feedback!

Where is the limit?

Can you not restart my  
workloads?

Check it out!

▶ <https://github.com/kubernetes/autoscaler/> ◀

The screenshot shows the GitHub repository page for `kubernetes/autoscaler`. At the top, there is a search bar and navigation links for Pull requests, Issues, Marketplace, and Explore. Below the repository name, there are buttons for View Repository, Unwatch (89), Unstar (1,279), and Fork. The main navigation tabs include Code, Issues (74), Pull requests (14), Projects (0), Wiki, and Insights. The repository description is "Autoscaling components for Kubernetes". Below this, there are statistics for 2,478 commits, 26 branches, 45 releases, 162 contributors, and Apache-2.0 license. There are buttons for Branch: master, New pull request, Create new file, Upload files, Find file, and Clone or download. A list of recent commits and files is shown, with the `vertical-pod-autoscaler` folder highlighted by a red box. The list includes folders like `addon-resizer`, `builder`, `cluster-autoscaler`, `hack`, and `vertical-pod-autoscaler`, along with files like `.gitignore` and `.travis.yml`.

File/Folder	Description	Time
<code>vertical-pod-autoscaler</code>	Fix gofmt errors	23 hours
<code>hack</code>	Update go version used from 1.10.2 to 1.11.2 to match one used by k8s	23 hours
<code>cluster-autoscaler</code>	Initialize klog	18 hours
<code>builder</code>	Merge pull request #1447 from aleksandra-malinowska/add-owners-builder	2 hours
<code>addon-resizer</code>	Merge pull request #1299 from Rajat-0/import_fix	a month
<code>.travis.yml</code>	Update go version used from 1.10.2 to 1.11.2 to match one used by k8s	23 hours
<code>.gitignore</code>	Build pod_nanny in a container.	a year



**KubeCon**

**CloudNativeCon**

————— **North America 2018** —————

**Thank You**

