

Multi-Cluster Experience at CERN

Thomas Hartland IT-CM-RPS



Existing projects

kubefed

- Formerly federation-v2
- Implemented using custom resource definitions
- Not a drop in replacement for a single cluster, have to rewrite manifests using CRDs

```
kind: Deployment
spec:
  replicas: 10
  template:          becomes...
    spec:
      containers:
        - image: nginx
```

```
kind: FederatedDeployment
spec:
  template:
    spec:
      replicas: 10
      template:
        spec:
          containers:
            - image: nginx
  placement:
    clusters:
      - name: cluster1
      - name: cluster2
  overrides:
    - clusterName: cluster2
      clusterOverrides:
        - path: "/spec/replicas"
          value: 5
```

kubefed

- Can use label selectors for the placement, but the person writing the manifest still needs to know the topology of the multi-cluster setup
- Anything else you might use (e.g workflow engine) does not create these federated types

kubefed

- Statuses of resources created in other clusters are not reported back
- Have to check them directly in the cluster they are running in

admiralty multi-cluster



- Admission controller intercepts pods that are being created and replaces them with “proxy pods”
- Creates “delegate pods” in other clusters using the original pod definitions

admiralty multi-cluster

No github activity for 8 months and then...

v0.4.0 see CHANGELOG.md

Signed-off-by: adrienjt <adrienjt@users.noreply.github.com>

 master  v0.4.0



adrienjt committed 5 days ago

 Showing **159 changed files** with **4,234 additions** and **3,766 deletions**.

admiralty multi-cluster

- Can annotate a pod to give a single cluster name to delegate to
- Pod status is synced back to the proxy pod
- Delegation and status sync were a bit unreliable in my testing

What I'm trying now

Use virtual-kubelet

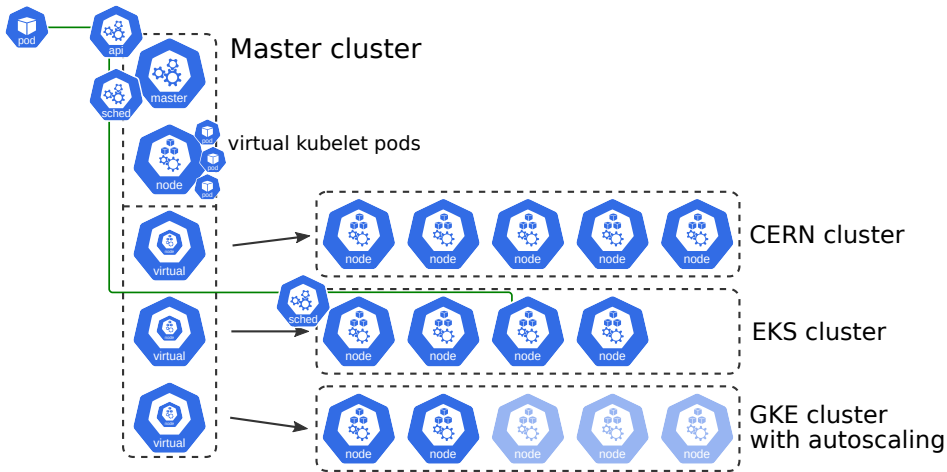
- Pod implementing the kubelet API and registering as a node
- Various providers to convert pods to other services

Use virtual kubelet

“It should be noted that VK is explicitly not intended to be an alternative to Kubernetes federation.”¹

¹virtual-kubelet CNCF proposal https://docs.google.com/document/d/1B0ZJ9CpNt6Q5TWoMCV_wgEhvrLA2WZz_vEjU4Q1eUQE

Use virtual kubelet



Use virtual kubelet

```
$ kubectl get nodes
```

NAME	STATUS	ROLES	AGE	VERSION
multicluster-cern-000-master-0	Ready	master	39d	v1.15.3
multicluster-cern-000-minion-0	Ready	<none>	39d	v1.15.3
multicluster-cern-001	Ready	agent	39d	v1.15.3
multicluster-cern-002	Ready	agent	39d	v1.15.3
multicluster-eks-001	Ready	agent	38d	v1.14.6-eks-5047ed
multicluster-gke-001	Ready	agent	39d	v1.13.11-gke.9

<https://gitlab.cern.ch/thartlan/virtual-cluster>

Restrictions

- Pod has to be independent of other cluster resources
- Scaling limit of number of nodes in the master cluster is no longer a problem, but scaling limit of pods might be

Others

- G-Research have been working on “armada”
<https://github.com/G-Research/armada>