

Towards a Vendor Neutral Kubernetes



kubernetes

About Me



Andrew Sy Kim

Kubernetes SIG Cloud Provider Co-Chair

Software Engineer at DigitalOcean




@a_sykim



@andrewsykim



What is this talk NOT about?





Multi-Cloud Kubernetes



How-to on supporting new vendors/clouds




Comparing Vendors/Distributions

Customizing & Extending Kubernetes for any Cloud



Cloud Providers





Bridging together Cloud & Kubernetes Resources

Cloud Provider



```
// Interface is an abstract, pluggable interface for cloud providers.
type Interface interface {
    // Initialize provides the cloud with a kubernetes client builder and may spawn goroutines
    // to perform housekeeping or run custom controllers specific to the cloud provider.
    // Any tasks started here should be cleaned up when the stop channel closes.
    Initialize(clientBuilder ControllerClientBuilder, stop <-chan struct{})
    // LoadBalancer returns a balancer interface. Also returns true if the interface is supported, false otherwise.
    LoadBalancer() (LoadBalancer, bool)
    // Instances returns an instances interface. Also returns true if the interface is supported, false otherwise.
    Instances() (Instances, bool)
    // Zones returns a zones interface. Also returns true if the interface is supported, false otherwise.
    Zones() (Zones, bool)
    // Clusters returns a clusters interface. Also returns true if the interface is supported, false otherwise.
    Clusters() (Clusters, bool)
    // Routes returns a routes interface along with whether the interface is supported.
    Routes() (Routes, bool)
    // ProviderName returns the cloud provider ID.
    ProviderName() string
    // HasClusterID returns true if a ClusterID is required and set
    HasClusterID() bool
}
```



Cloud Abstractions/APIs

- Nodes
 - instance types
 - regions/zones
 - addresses
 - provider IDs
- Load Balancers
- Routes/IPAM
- Volumes




Is Kubernetes Vendor Neutral?



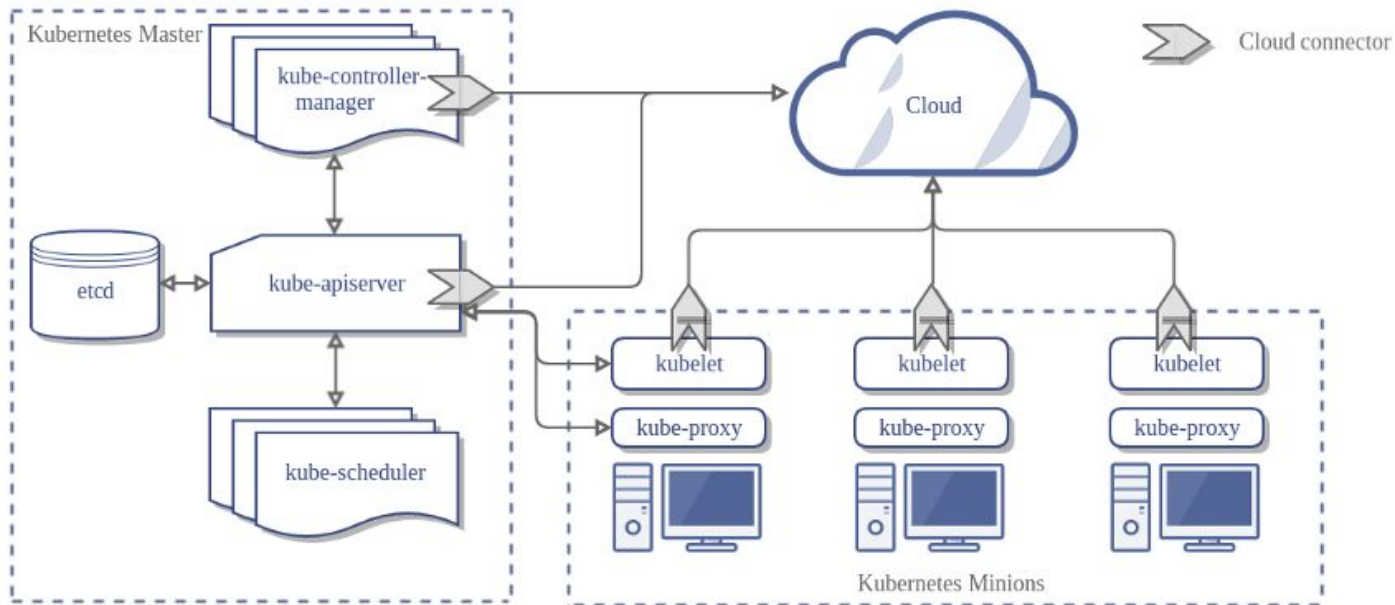
In-Tree & Out-of-Tree Cloud Providers



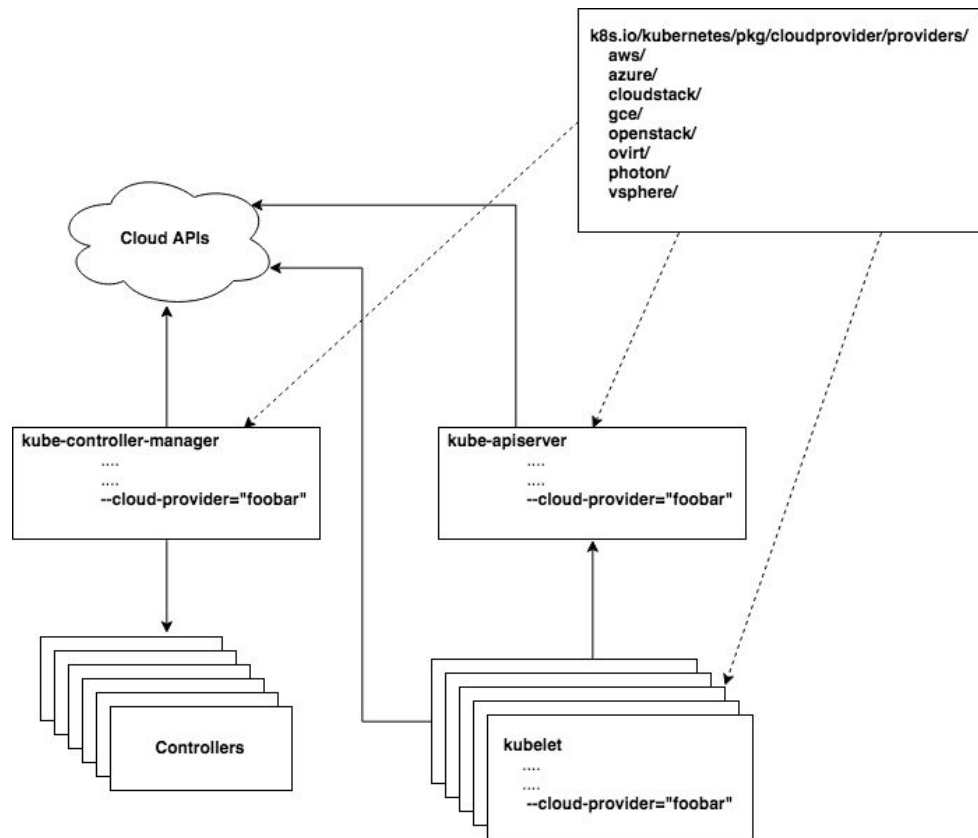
History of Kubernetes Cloud Providers



In-Tree



In-Tree





thockin commented on 29 Sep 2016

Member



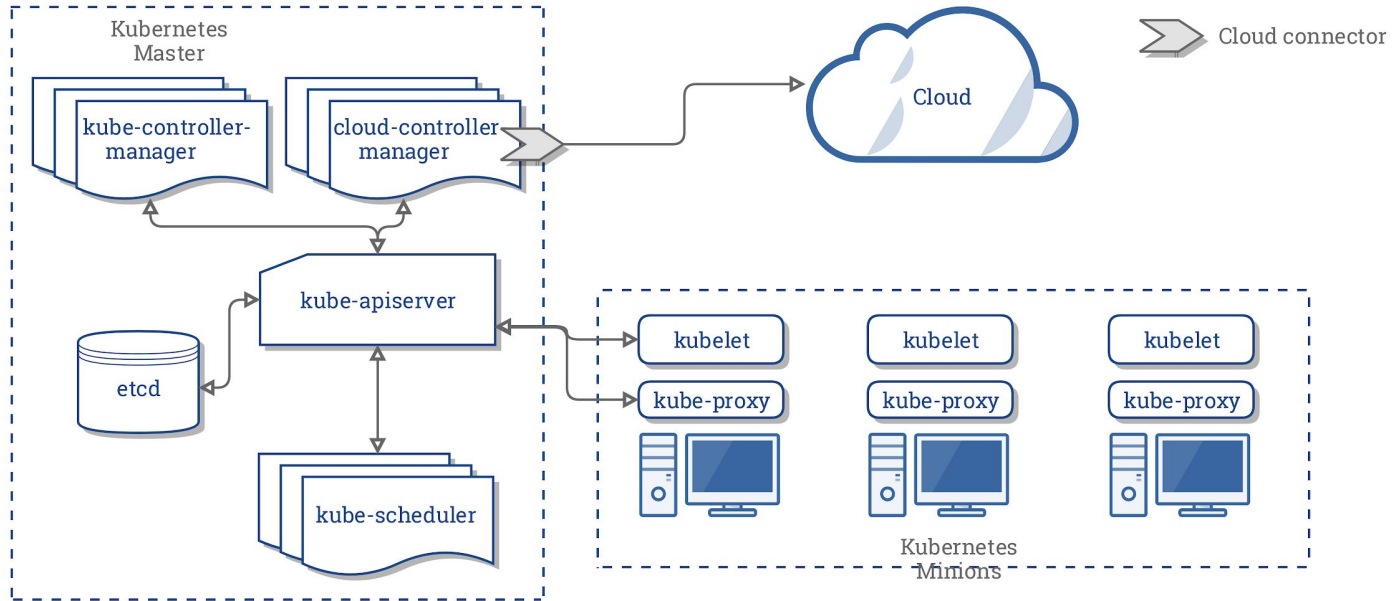
So this is a tough phase to be in. We know we want to change the way CloudProvider works, but that work hasn't really started yet. Are you guys willing to be a bit of a guinea pig? In broad strokes, here's what I am thinking. Don't build this as a Cloud Provider. Instead, let's make ControllerManager understand that "external" is a valid cloud provider, and write this as a controller. Simply run it somewhere in the cluster and let it watch the various APIs it needs.

While you do this, you can start to catalog all the places that assume CloudProvider is populated, and your controller can be a starting point for the discussion of how to evict google and amazon and openstack into their own modules.

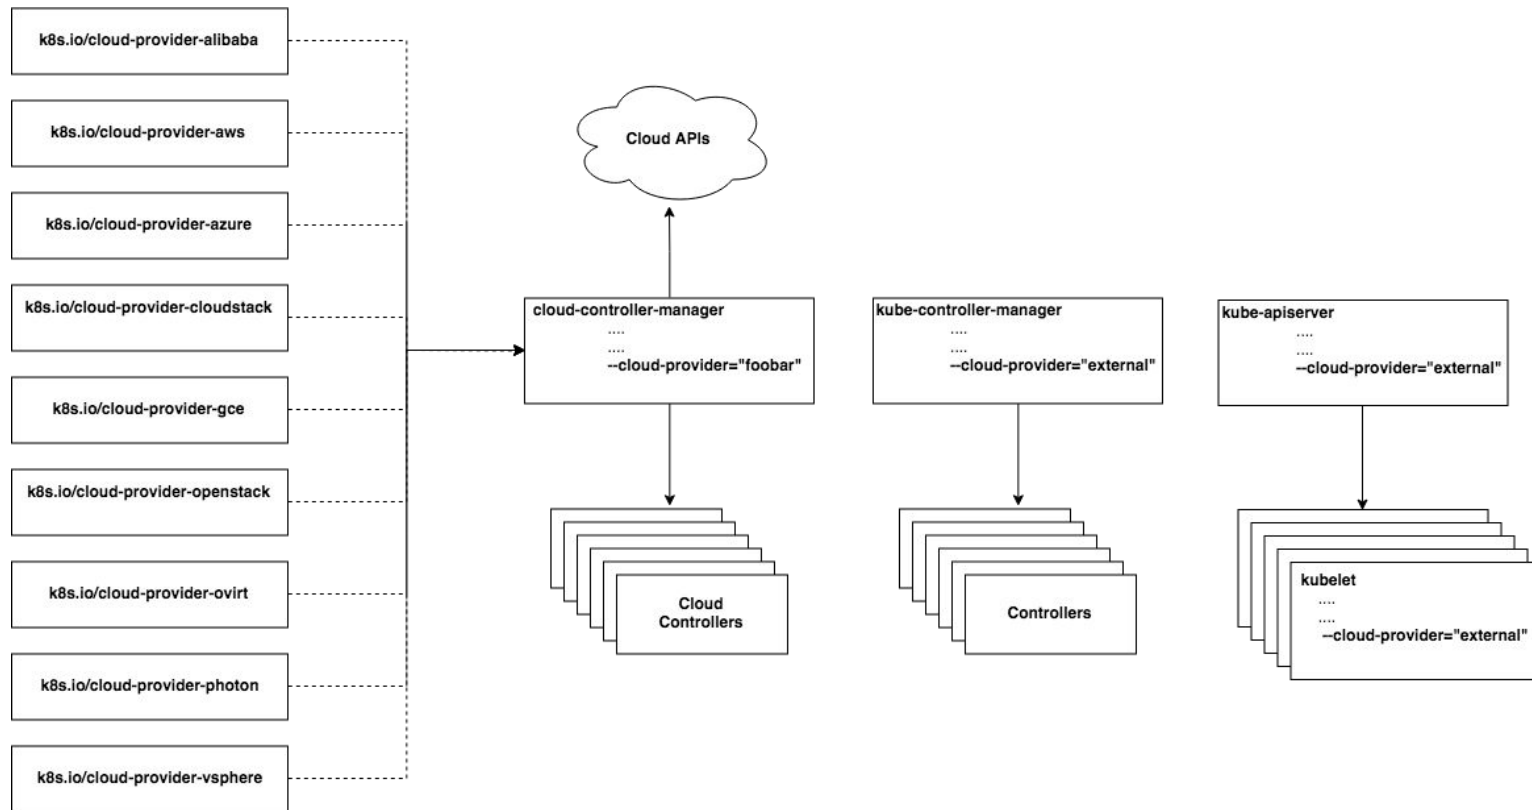
@errordeveloper




Out-of-Tree



Out-of-Tree



Is Kubernetes Cloud Provider Neutral?



- _ (ツ) _ / -





Advantage In-Tree: Visibility



Advantage Out-of-Tree: Faster Releases



Advantage In-Tree: Using Upstream Tools



Advantage Out-of-Tree: **NOT** Using Upstream Tools




Advantage In-Tree: Community Contribution




Advantage Out-of-Tree: Flexible Ownership


Inconsistent Integration



Community Acknowledgement



Converging Towards One Model





Where Are We Today?

In-Tree Cloud Providers

```
k8s.io/kubernetes/pkg/cloudprovider/  
└─ providers  
    ├── aws  
    ├── azure  
    ├── cloudstack  
    ├── gce  
    ├── openstack  
    ├── ovirt  
    ├── photon  
    └── vsphere
```



Migrate Out-of-Tree

k8s.io/

- cloud-provider
- cloud-provider-alibaba
- cloud-provider-aws
- cloud-provider-azure
- cloud-provider-cloud-stack
- cloud-provider-gcp
- cloud-provider-openstack
- cloud-provider-ovirt
- cloud-provider-cloudstack
- cloud-provider-photon
- cloud-provider-vmware

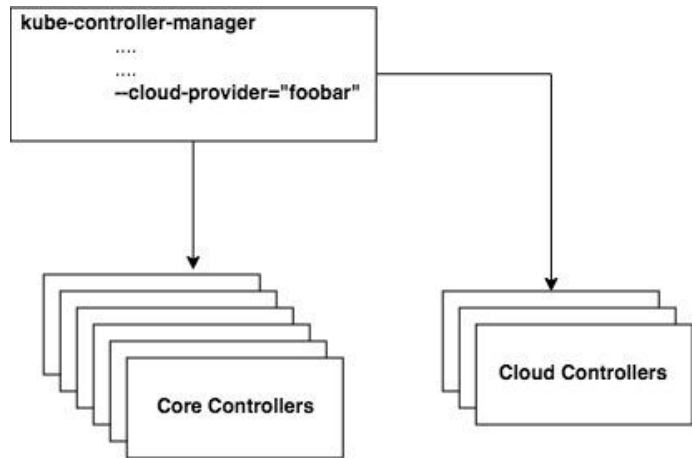




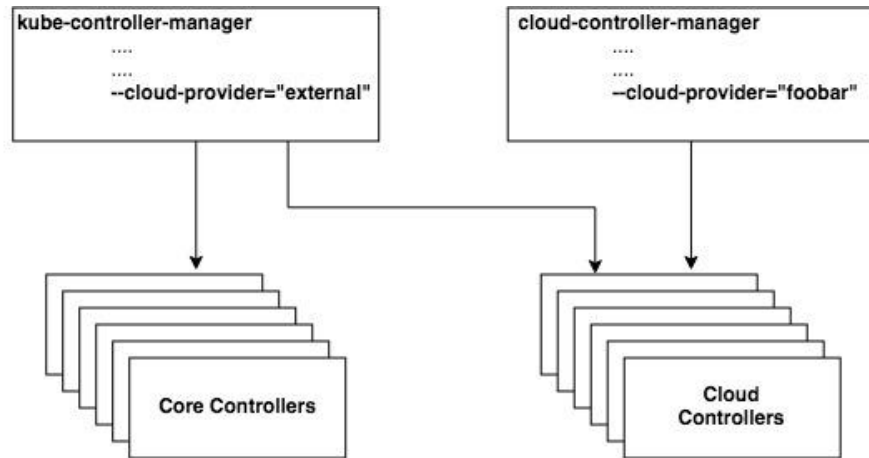

What's taking so long?



Zero Downtime Migration



Kubernetes v1.X

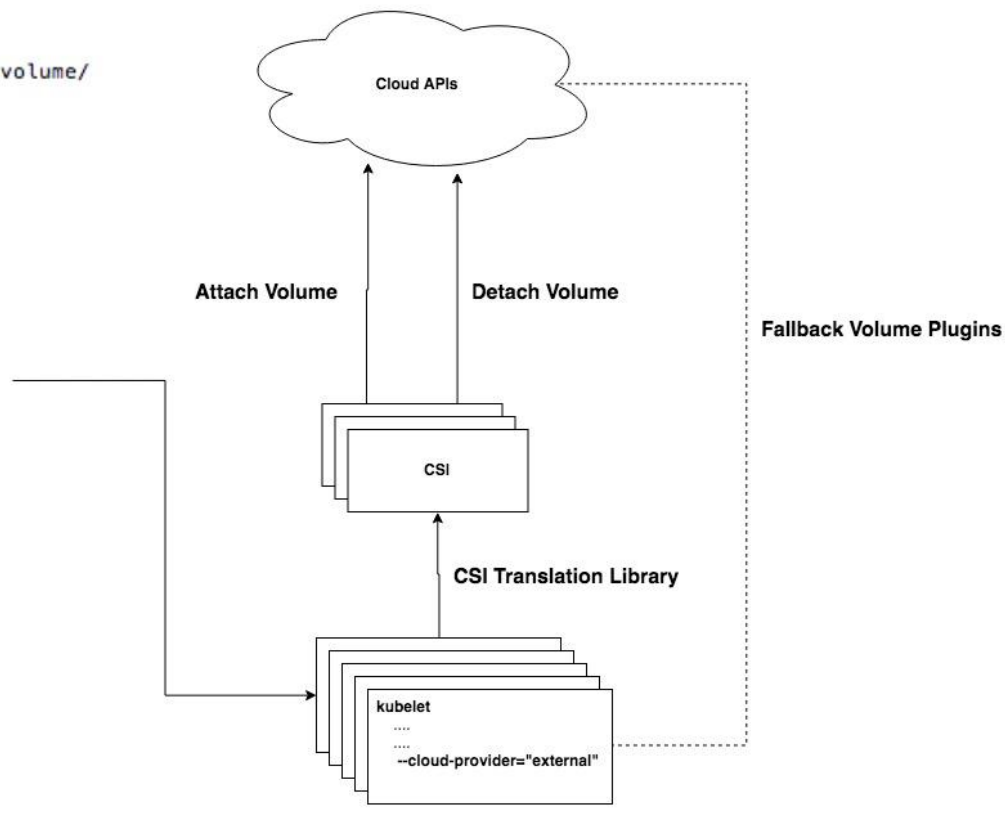


Kubernetes v1.Y




k8s.io/kubernetes/pkg/volume/

- awsebs
- azure-dd
- azure-file
- cephfs
- cinder
- configmap
- csi
- downwardapi
- emptydir
- fc
- flexvolume
- flocker
- gcepd
- git-repo
- glusterfs
- host-path
- iscsi
- local
- nfs
- photon-pd
- portworx
- projected
- quobyte
- rbd
- scaleio
- secret
- storageos
- testing
- util
- validation
- vsphere-volume



“Staging” In-Tree Cloud Providers

```
k8s.io/kubernetes/pkg/cloudprovider/  
└─ providers  
   └─ aws  
   └─ azure  
   └─ cloudstack  
   └─ gce  
   └─ openstack  
   └─ ovirt  
   └─ photon  
   └─ vsphere
```



```
k8s.io/kubernetes/staging/src/k8s.io/  
└─ cloud-provider  
   └─ cloud-provider-aws  
   └─ cloud-provider-azure  
   └─ cloud-provider-cloud-stack  
   └─ cloud-provider-gcp  
   └─ cloud-provider-openstack  
   └─ cloud-provider-ovirt  
   └─ cloud-provider-cloudstack  
   └─ cloud-provider-photon  
   └─ cloud-provider-vsphere
```



Staging Symlinks


```
~/go/src/k8s.io/kubernetes/vendor/k8s.io $ $ ls -l
lrwxr-xr-x  api -> ../../staging/src/k8s.io/api
lrwxr-xr-x  client-go -> ../../staging/src/k8s.io/client-go
...
...
...
lrwxr-xr-x  cloud-provider -> ../../staging/src/k8s.io/cloud-provider
lrwxr-xr-x  cloud-provider-aws -> ../../staging/src/k8s.io/cloud-provider-aws
lrwxr-xr-x  cloud-provider-azure -> ../../staging/src/k8s.io/cloud-provider-azure
lrwxr-xr-x  cloud-provider-cloudstack -> ../../staging/src/k8s.io/cloud-provider-cloudstack
lrwxr-xr-x  cloud-provider-gcp -> ../../staging/src/k8s.io/cloud-provider-gcp
lrwxr-xr-x  cloud-provider-openstack -> ../../staging/src/k8s.io/cloud-provider-openstack
lrwxr-xr-x  cloud-provider-ovirt -> ../../staging/src/k8s.io/cloud-provider-ovirt
lrwxr-xr-x  cloud-provider-photon -> ../../staging/src/k8s.io/cloud-provider-photon
lrwxr-xr-x  cloud-provider-vsphere -> ../../staging/src/k8s.io/cloud-provider-vsphere
```



“Staging” In-Tree Cloud Providers

k8s.io/kubernetes/staging/src/k8s.io/

- cloud-provider
- cloud-provider-aws
- cloud-provider-azure
- cloud-provider-cloud-stack
- cloud-provider-gcp
- cloud-provider-openstack
- cloud-provider-ovirt
- cloud-provider-cloudstack
- cloud-provider-photon
- cloud-provider-vmware

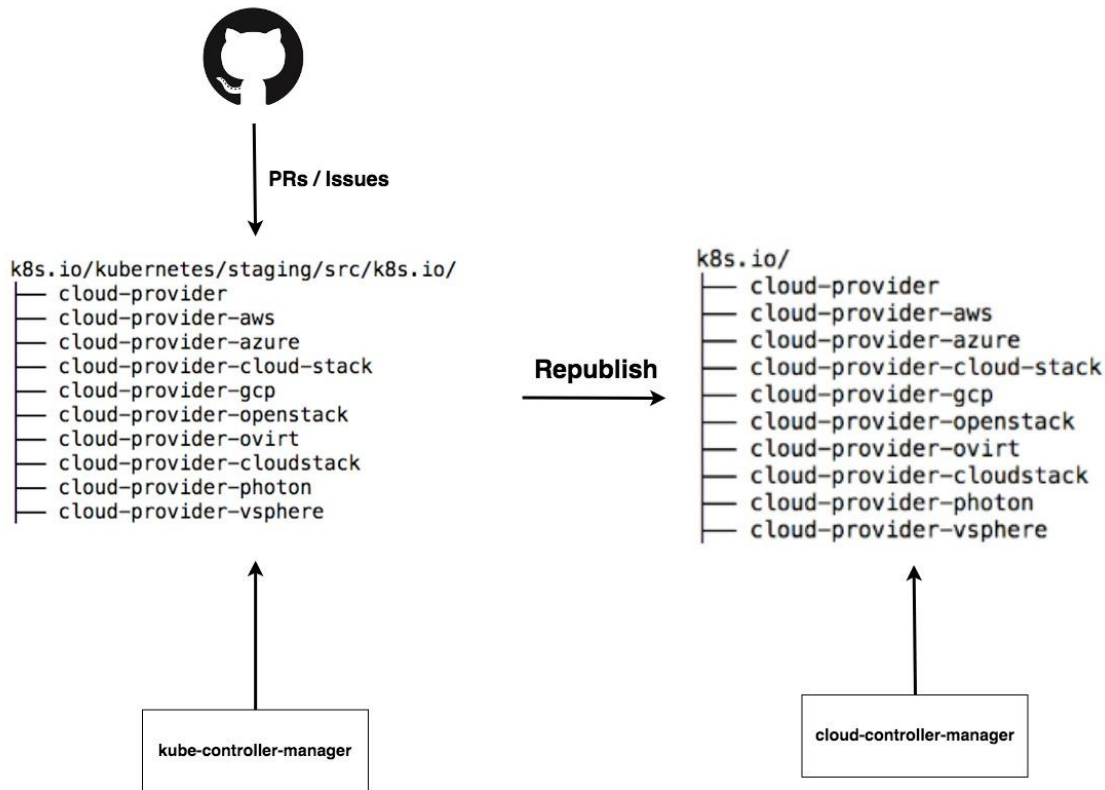


k8s.io/

- cloud-provider
- cloud-provider-aws
- cloud-provider-azure
- cloud-provider-cloud-stack
- cloud-provider-gcp
- cloud-provider-openstack
- cloud-provider-ovirt
- cloud-provider-cloudstack
- cloud-provider-photon
- cloud-provider-vmware



“Staging” In-Tree Cloud Providers



Breaking Circular Dependencies

k8s.io/kubernetes/.../providers

- aws
- azure
- cloudstack
- fake
- gce
- openstack
- ovirt
- photon
- vsphere



k8s.io/kubernetes



Removing cloud provider dependencies to kubernetes/kubernetes #69585

Open andrewsykim opened this issue on Oct 10 · 31 comments



andrewsykim commented on Oct 10 • edited

Member



Is this a **BUG REPORT** or **FEATURE REQUEST**?:

/kind feature

As part of a long running initiative to remove cloud providers out of kubernetes/kubernetes, it's required to remove dependencies to kubernetes/kubernetes so we can place them into a staging directory. The following dependencies need to be removed from

k8s.io/kubernetes/pkg/cloudprovider/providers :

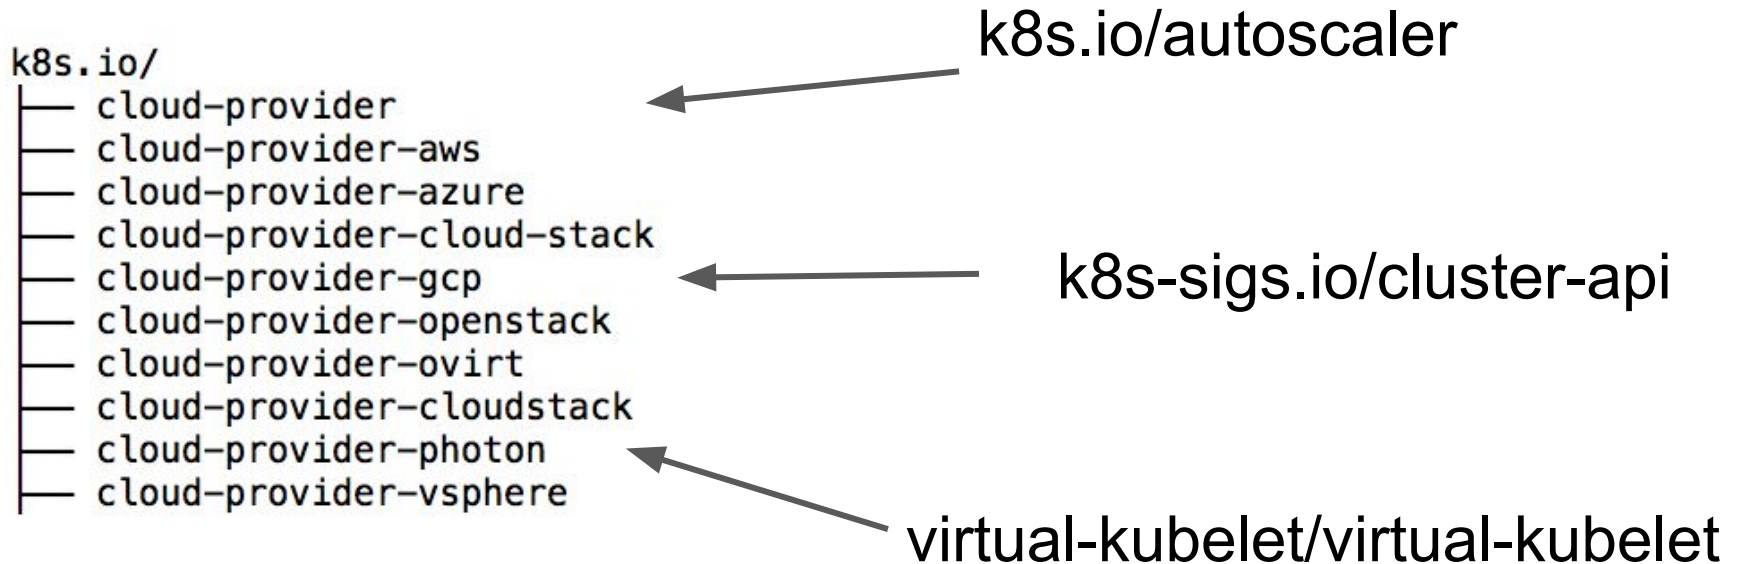
Dependency checklist:

- pkg/api/service (@mcrute)
- pkg/api/v1/service (@mcrute)
- pkg/apis/core/v1/helper (@islinwb)
- pkg/controller (@stewart-yu)
- pkg/credentialprovider/aws (@yastij)
- pkg/master/ports (@d-nishi)
- pkg/cloudprovider (@cheftako)
- pkg/features
- pkg/kubelet/apis (@islinwb)
- pkg/util/keymutex (@dims)
- pkg/util/mount (@mcrute)
- pkg/util/net/sets (@mcrute)
- pkg/util/node (@cheftako)
- pkg/util/parsers (@cheftako)
- pkg/util/strings (@dims)
- pkg/version (@islinwb)
- pkg/volume (@islinwb)
- pkg/volume/util (@islinwb)
- pkg/volume/util/volumepathhandler


/sig cloud-provider



Consumed for Other Repos



Benefits





Smaller Binaries

Smaller Binaries

Component	With Cloud Providers	Without Cloud Providers
kubelet	126 MiB	56 MiB
kube-controller-manager	146.3 MiB	89.9 MiB
kube-apiserver	147.5 MiB	123.1 MiB





Separation of Concern



Open Cloud



Frameworks Instead of Interfaces

Community & Governance





SIG per Cloud Provider is Not Sustainable



Cloud Providers as Subprojects

New Cloud Providers in Kubernetes

- Conformance certified
- Reporting conformance test results to testgrid
- Well documented
- User experience reports
- Well tested/used CCM
- Tech leads are members of the Kubernetes org



Intro: SIG Cloud Provider




- Dec 12th 1:45 pm - 2:20 pm
- Chris Hoge (Open Stack) & Jago Macleod (Google)

Deep Dive: SIG Cloud Provider



- Dec 13th 10:50 am - 11:25 am
- Chris Hoge (Open Stack) & Jago Macleod (Google)

Thank you!



Questions!

