



KubeCon



CloudNativeCon

North America 2017

Cinder for Baremetal Kubernetes Clusters

John Griffith, *NetApp*

Huamin Chen, *Redhat*

Who are these two?

- John Griffith
 - Led the start of the Cinder project back in the day
 - Works on “Cloudy things”, OpenStack, SIG Storage, CSI etc
 - Lead Open Source development and strategy for SolidFire
 - Twitter: [@jdg_8](#), IRC: [jgriffith](#), Slack/Github: [j-griffith](#)
- Huamin Chen
 - Wrote many Kubernetes volume plugins back in the day
 - Member of Kubernetes SIG Storage
 - Twitter: [@root_fs](#), Slack/Github: [rootfs](#)

What's Cinder

- Block Storage as a Service
- Relatively mature, now around 6 years old
- Provides an abstraction for 70+ block storage backends
- Similar model/philosophy as Amazons EBS
 - with pluggable backends
 - or reference LVM provider
- Implements a standard block service for things like create, attach, snapshot and backup... (the list goes on)
- All drivers are in tree and require CI system
- > 500 community contributors

It's not just for OpenStack

- Initially was to provide Block Storage for OpenStack
- But always intended to be able to be used by itself
- Can be deployed in Containers as a standalone service
- True standalone (no-auth) or *sorta* standalone w/keystone
- Attach/Detach to bare-metal
- Same backend device maturity and support, just attach to bare-metal and consume in containers
- Cinder's one of the *simpler* services to deploy and maintain

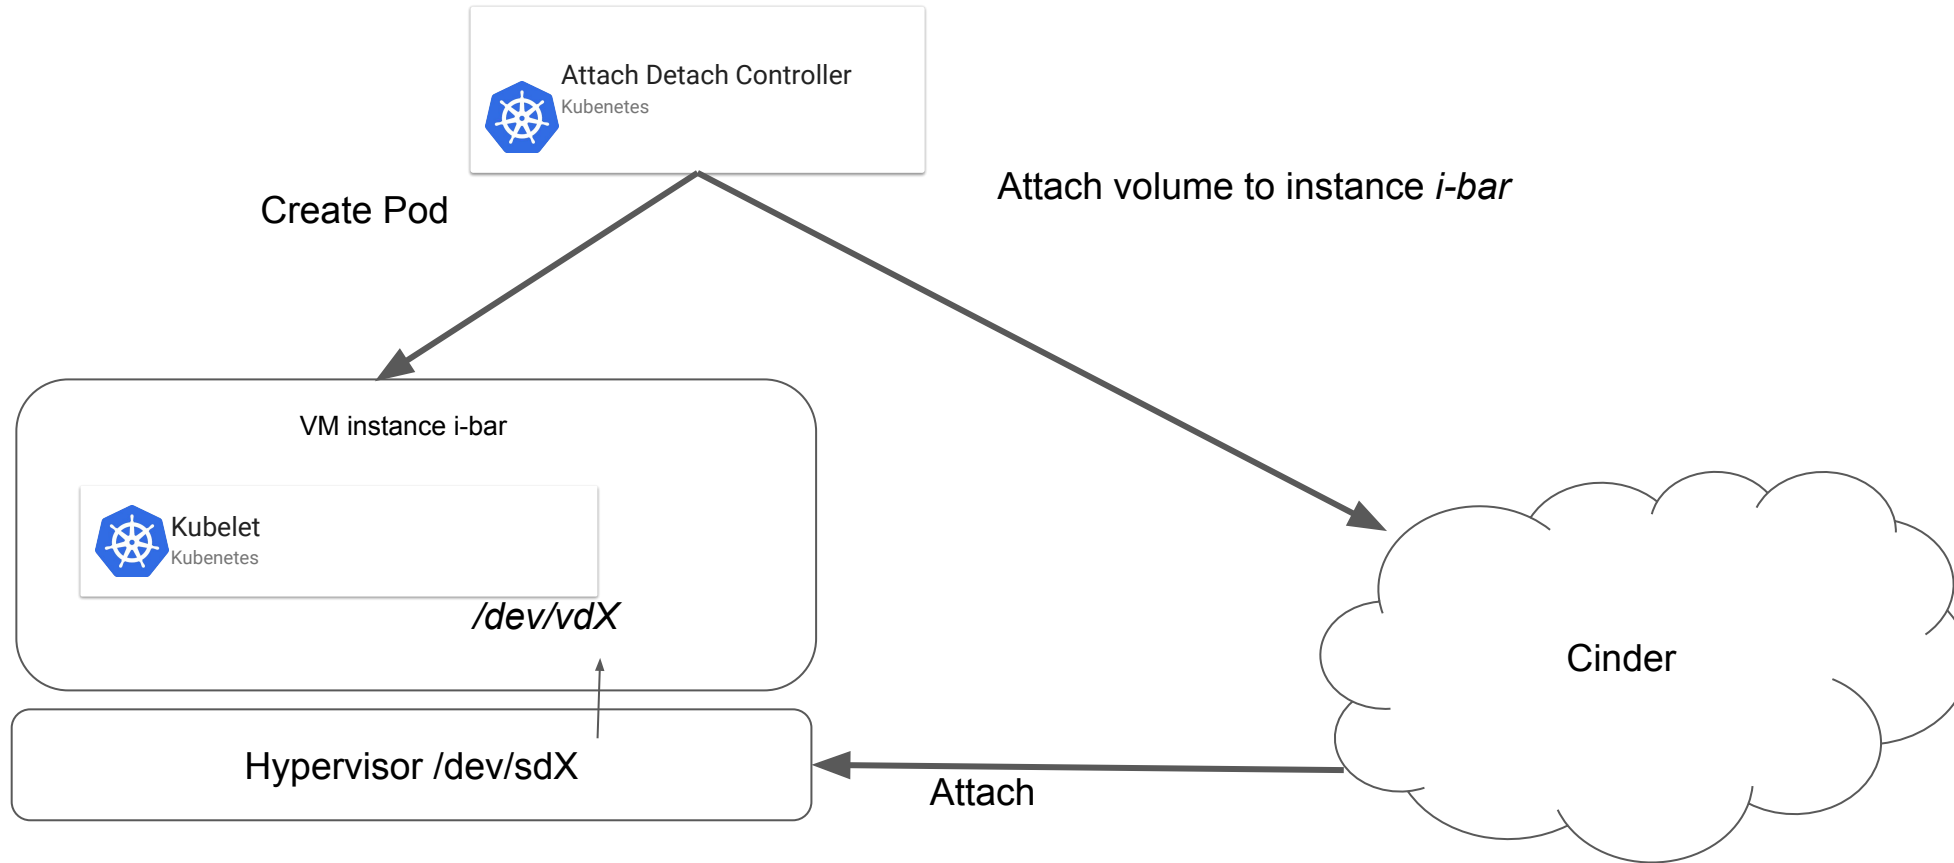
3 step deploy (*BlockBox*)

- *git clone https://github.com/openstack/cinder*
- *cd cinder/contrib/blockbox && Make*
 - Builds images (or build your own, download from dockerhub and skip)
- *docker-compose up*
 - Start all the Cinder services you need in containers
- DONE!
- Ok, Ok.. in reality you'll probably want to adjust/tweak a bit, but that's really all there is to it

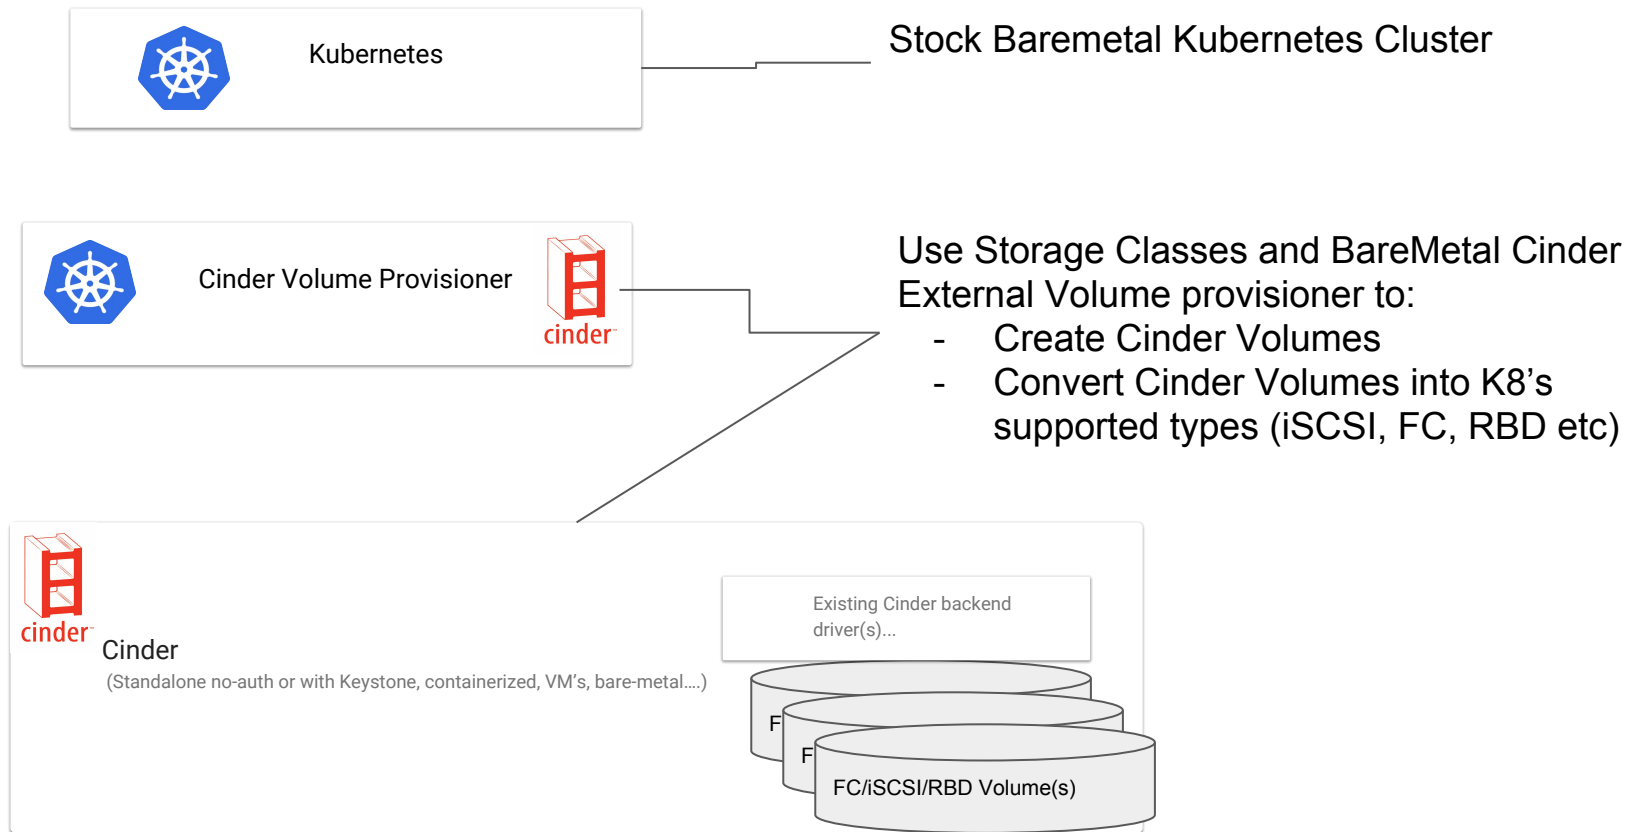
Kubernetes Storage Volumes

- Persistent Volume Sources:
 - Cloud Storage
 - AWS/EBS, Azure File/Disk, GCE PD, Cinder, vSphere
 - Block
 - iSCSI, FC, RBD, Scaleio, StorageOS, Portworx...
 - File Storage:
 - NFS, Gluster, CephFS, Quobyte....
- Features
 - Dynamic provisioning
 - Attach/Detach
 - Resize
 - Snapshot
 -

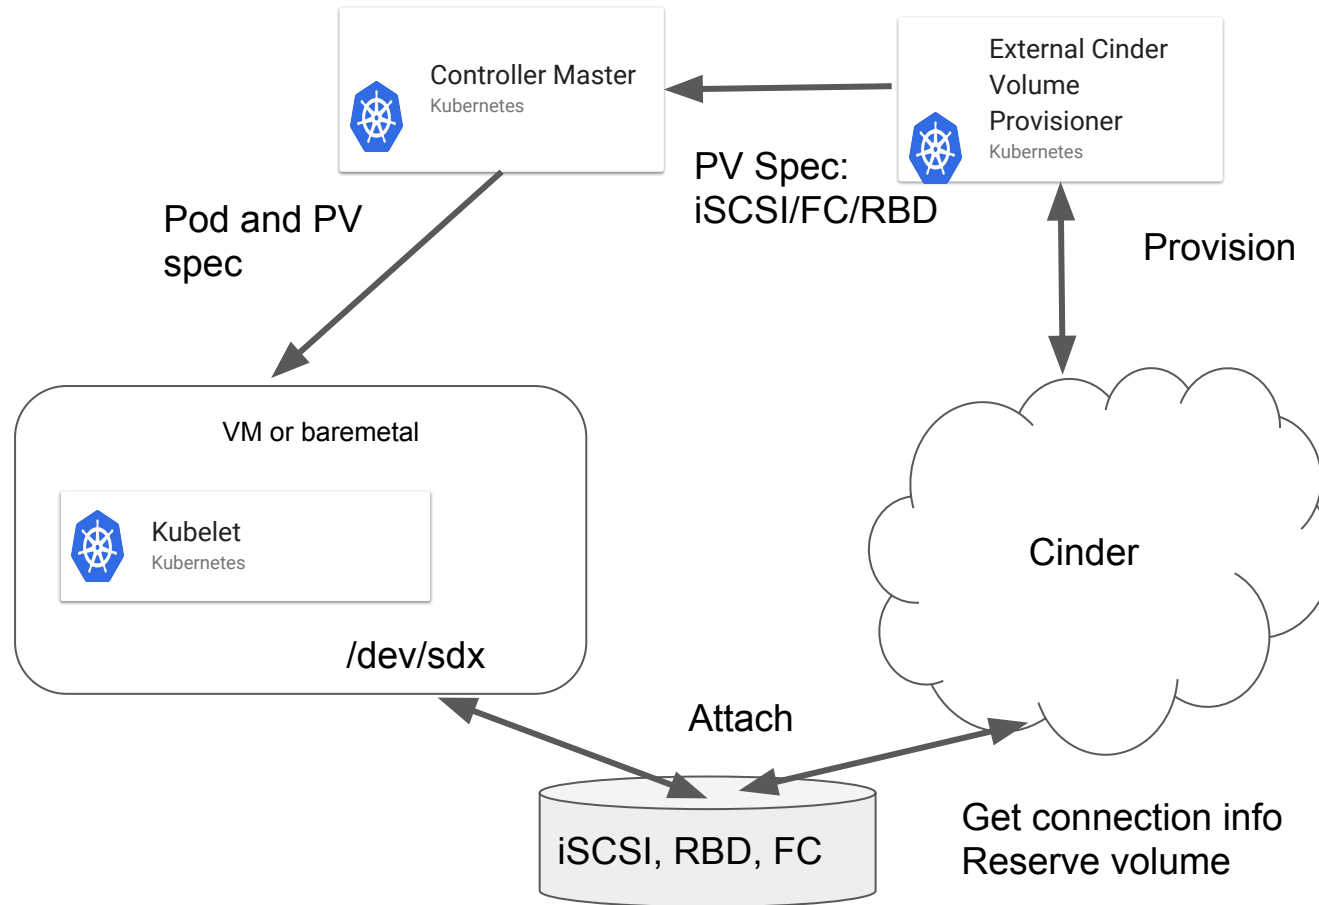
Cloud Storage in Kubernetes (OpenStack)



High Level component view

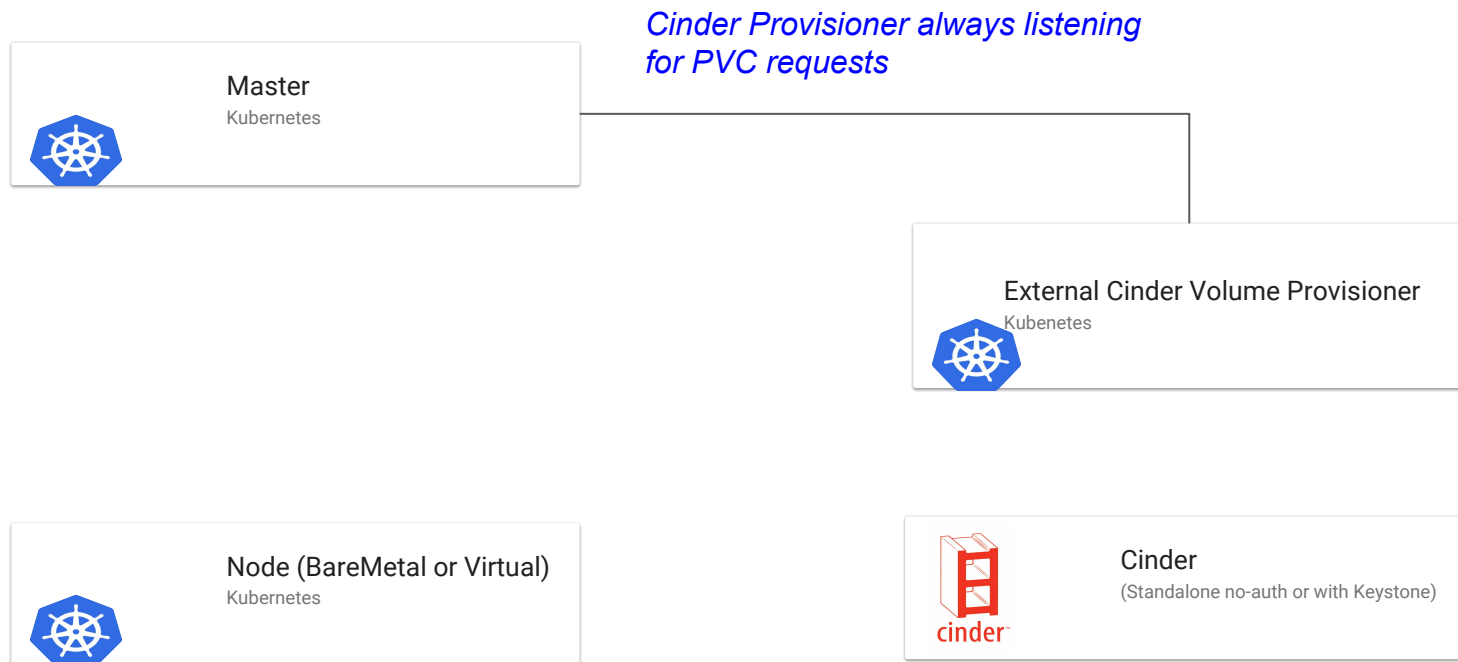


External provisioner process flow



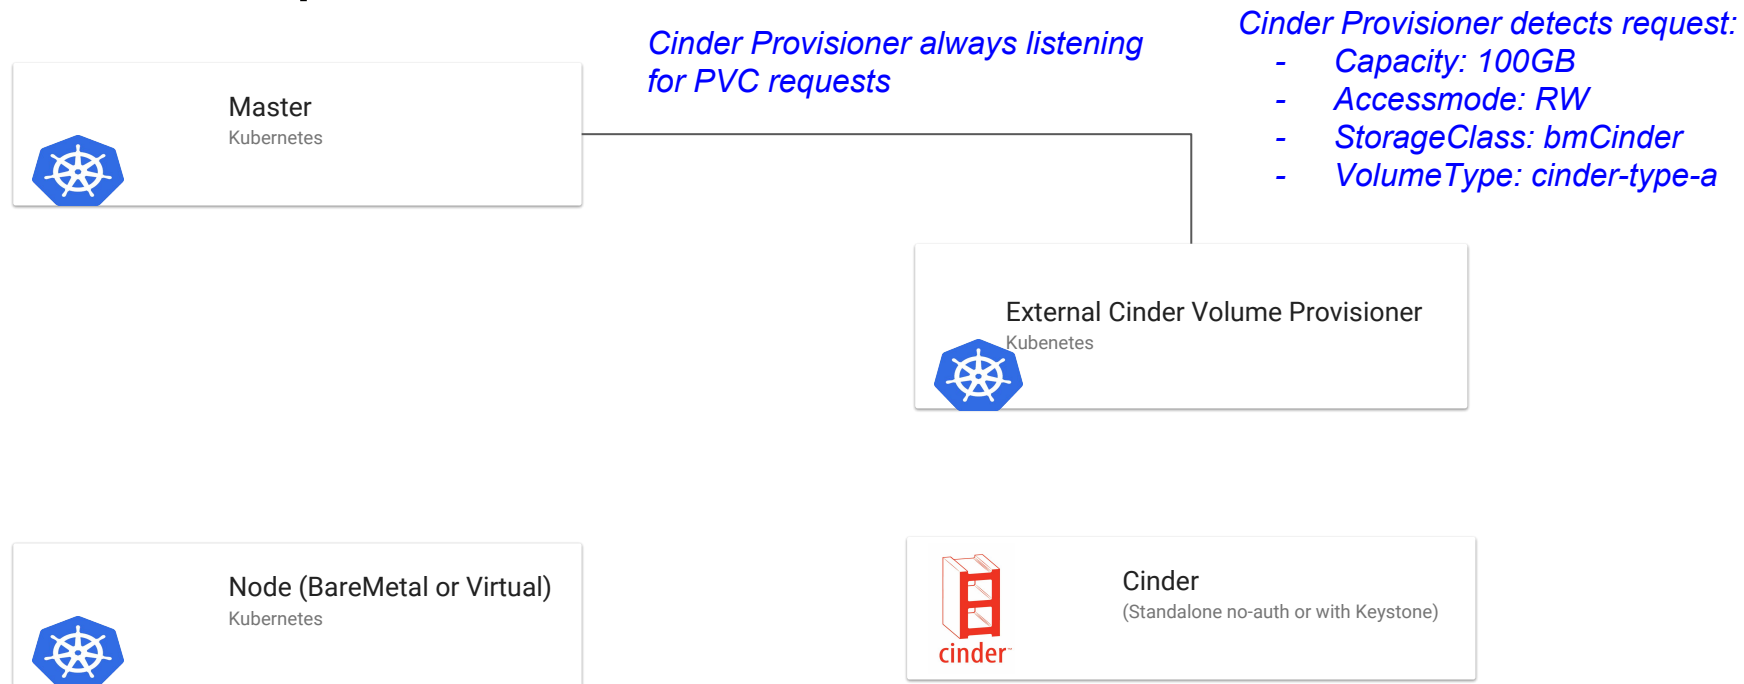
Example process/flow

External provisioner “listening” for PVC requests



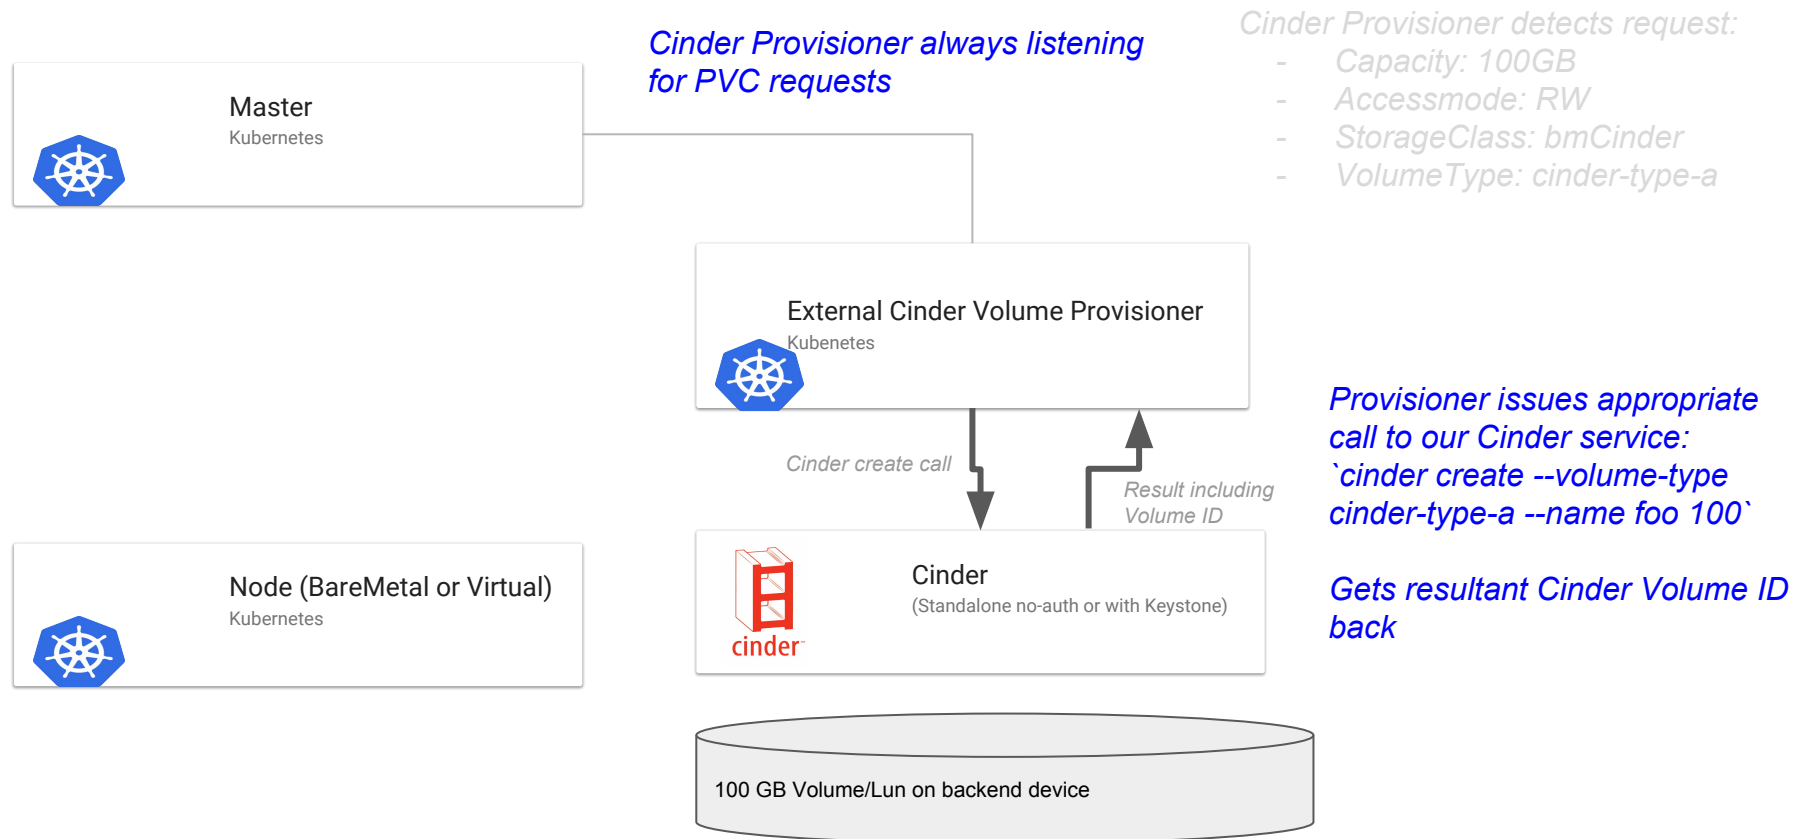
Example process/flow

PVC Request via Kubernetes



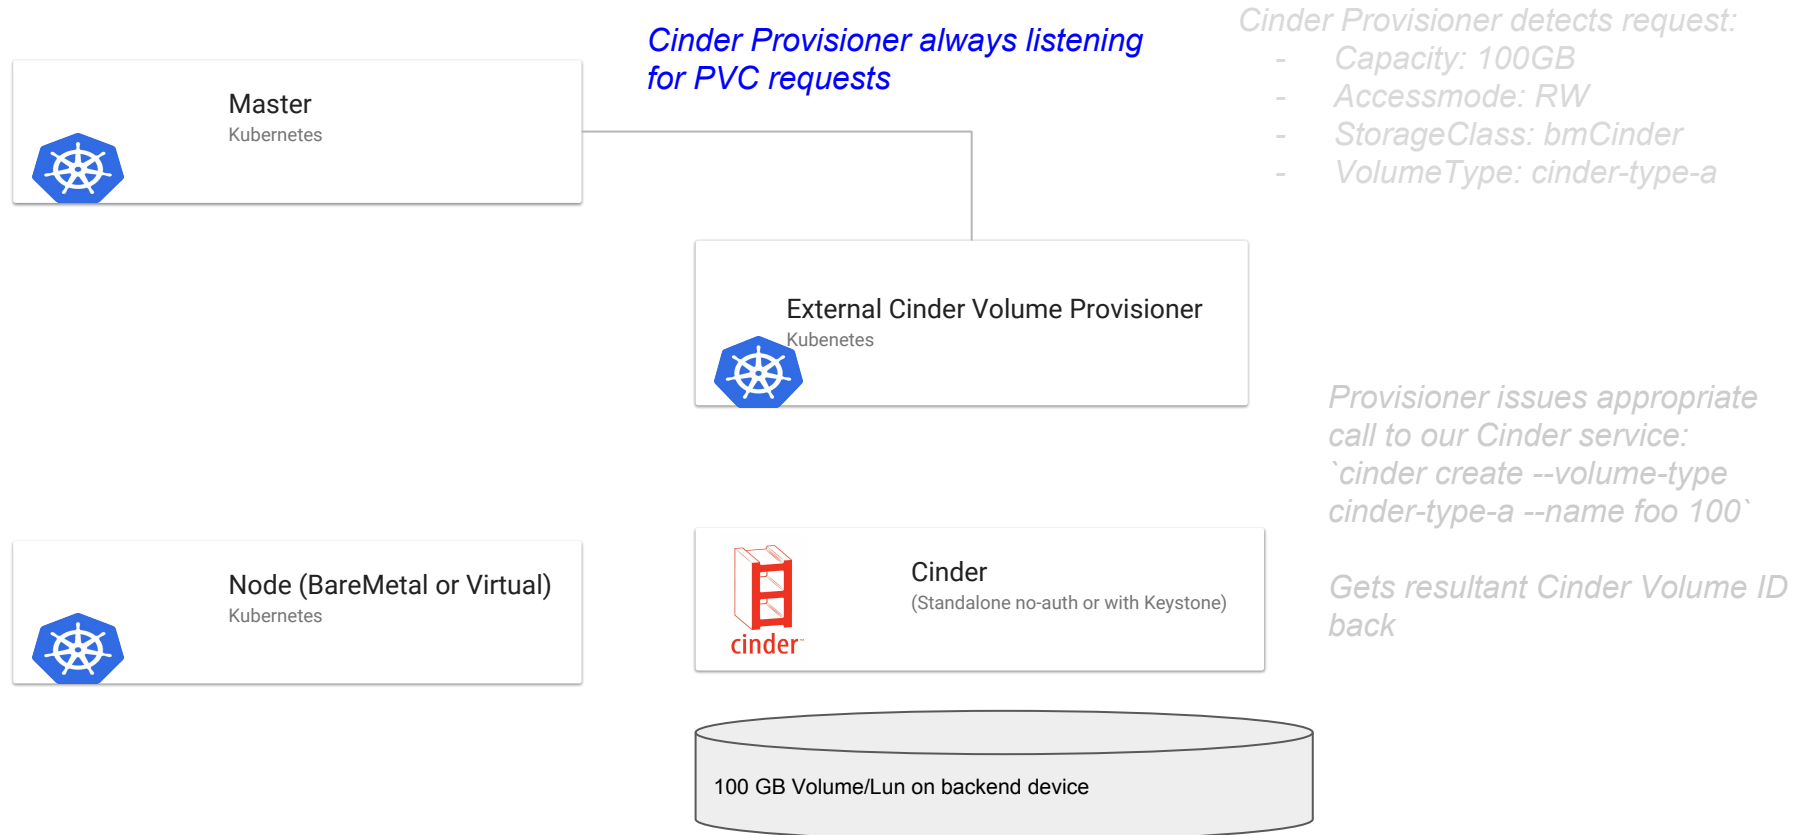
Example process/flow

Volume Creation



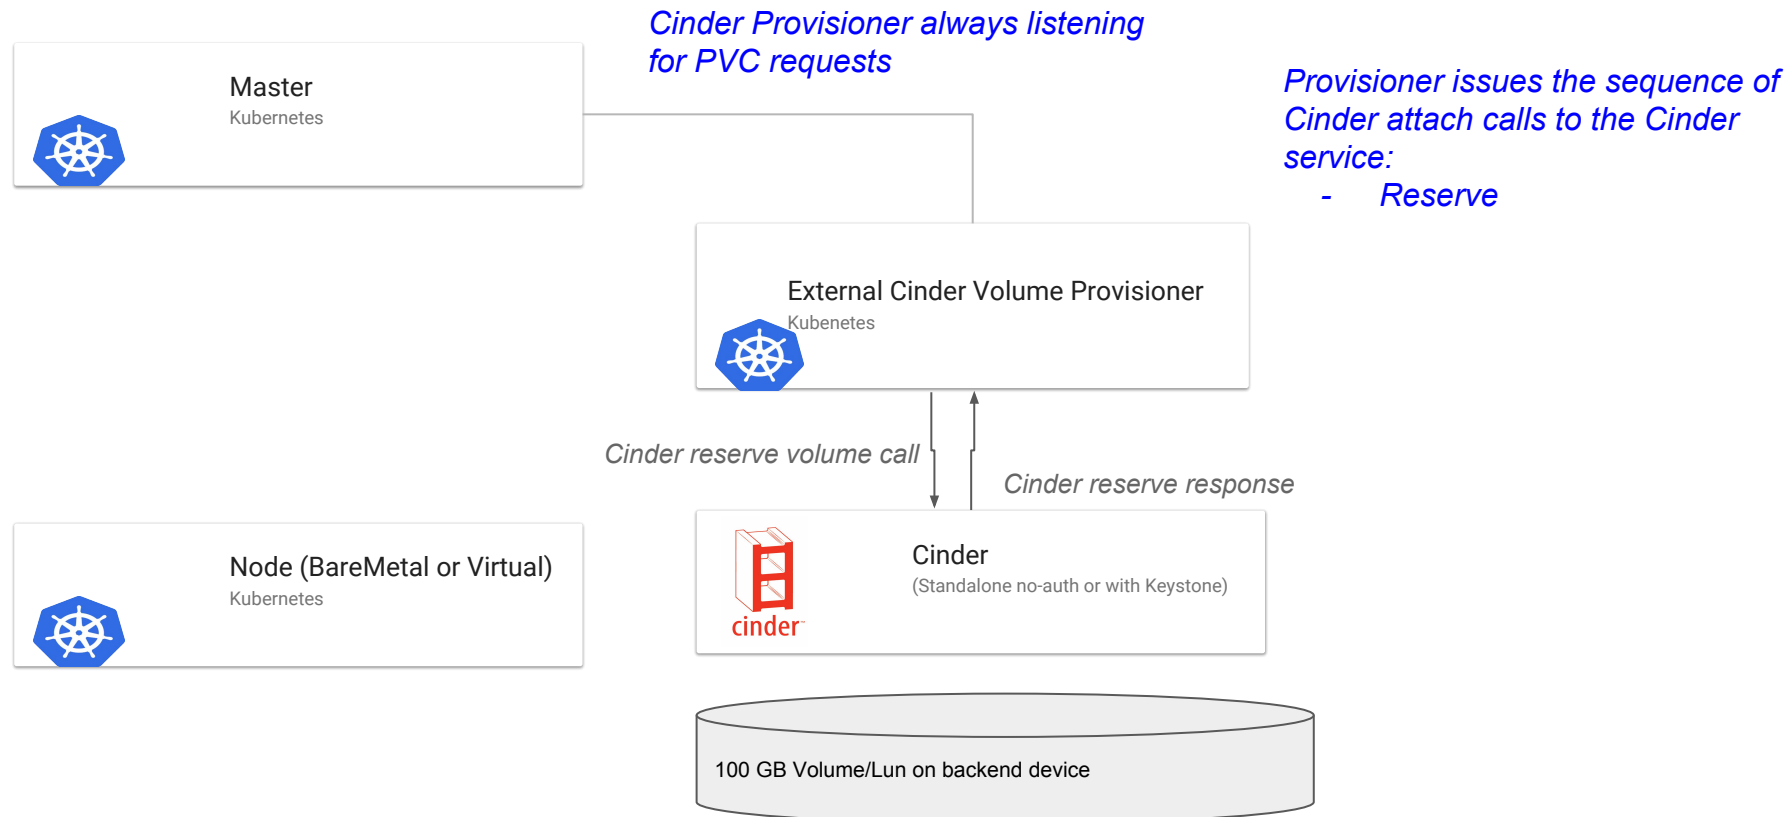
Example process/flow

Waiting for new PVC or Attach request



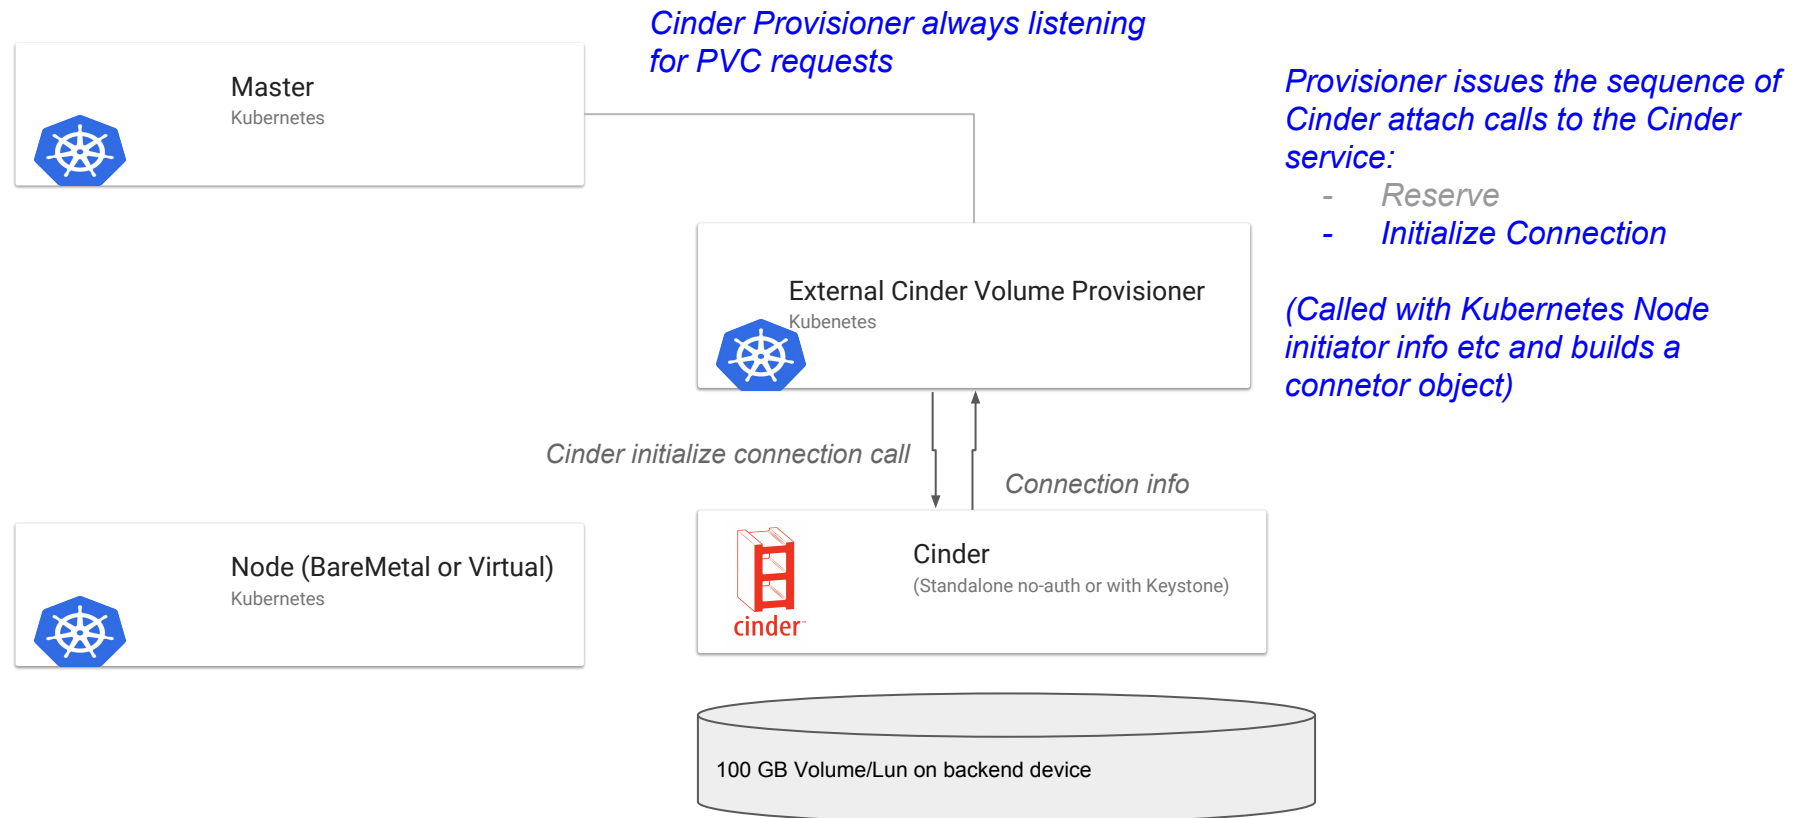
Provisioning detail

Provisioning our Volume to a POD (1 of 4)



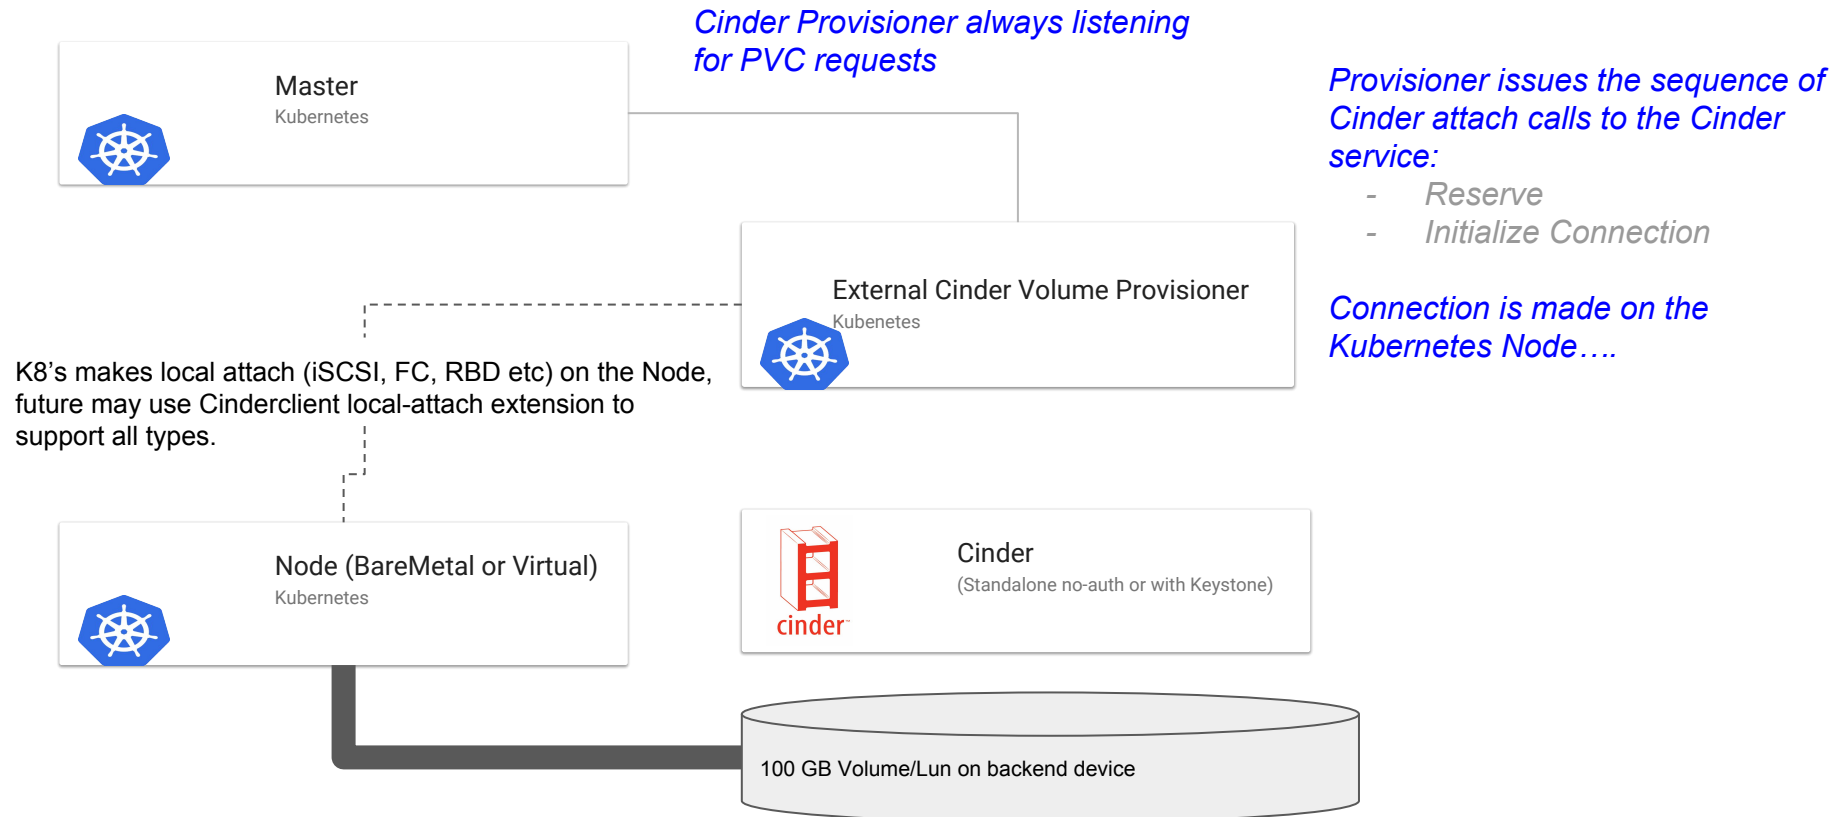
Provisioning detail

Provisioning our Volume to a POD (2 of 4)



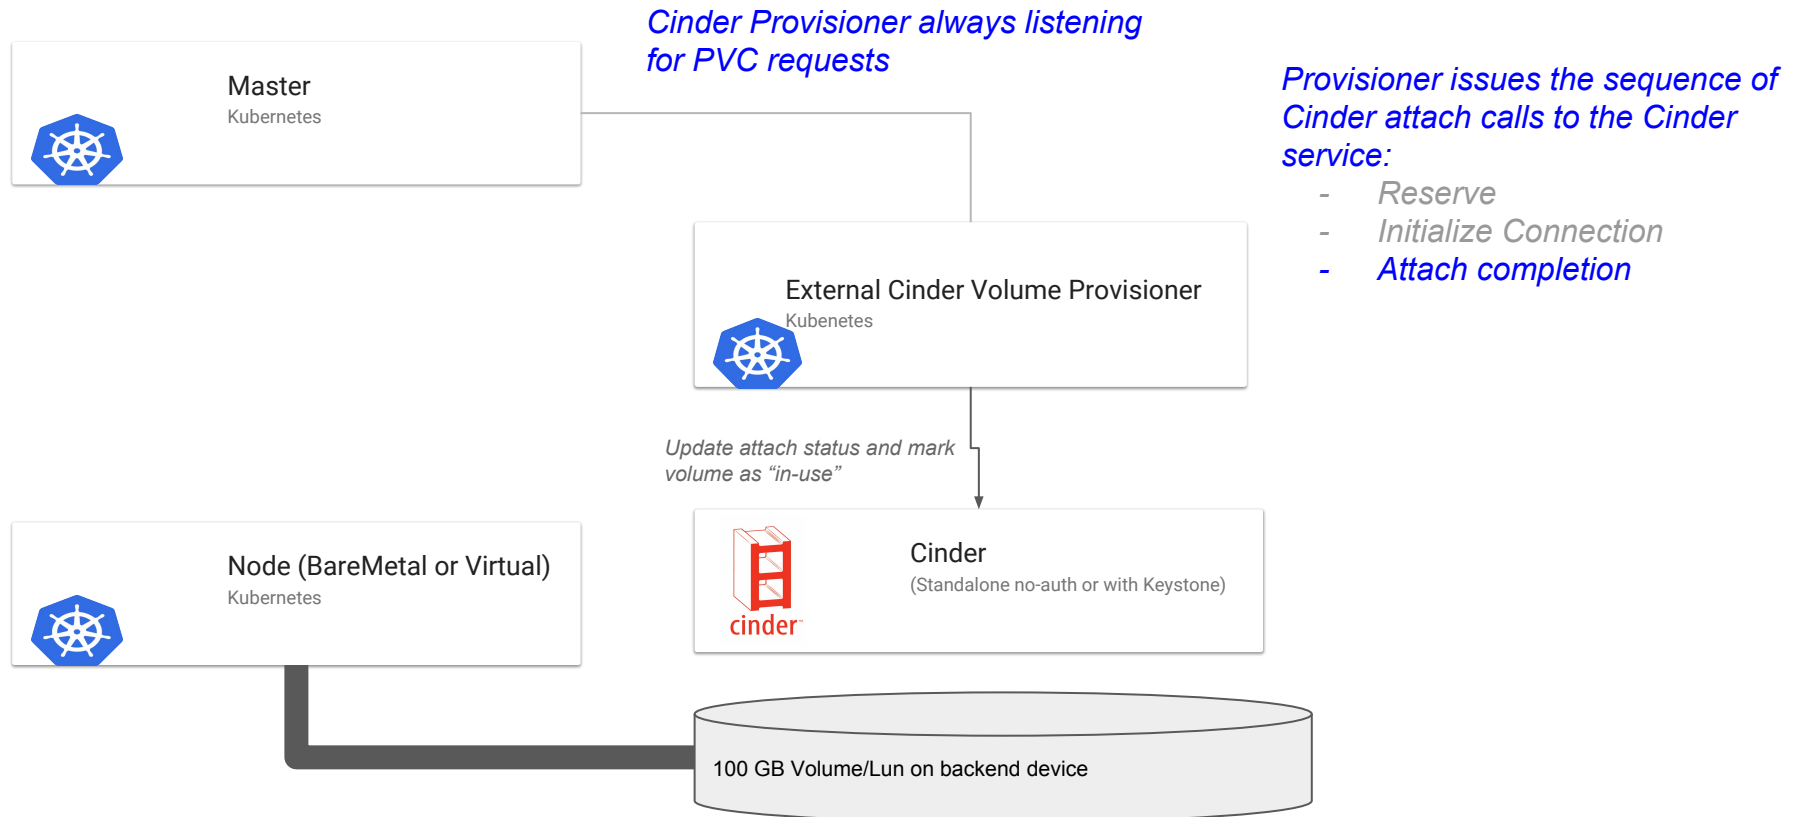
Provisioning detail

Provisioning our Volume to a POD (3 of 4)



Provisioning detail

Provisioning our Volume to a POD (4 of 4)



Current Status

- Merged into Kubernetes-Incubator/external-storage
 - <https://github.com/kubernetes-incubator/external-storage>
- Supports iSCSI and RBD
 - iSCSI is predominant Cinder data transport (49 out of 71)
- FC and NFS are likely next
 - Might be able to leverage the existing cinder extensions here?
 - Use tcmu-runner for NFS?

Kubernetes Storage Feature Integration

Kubernetes Storage Features	Links	Benefits
Volume Snapshot	https://github.com/kubernetes-incubator/external-storage/tree/master/snapshot	Use Cinder snapshot out of the box
Volume Resize	https://kubernetes.io/docs/concepts/storage/persistent-volumes/#expanding-persistent-volumes-claims	Use Cinder volume extend out of the box
Future features (replication, etc)	N/A	Cinder replication, etc

What's next?

- Container Storage Interface (CSI) Integration
 - Provisioner
 - <https://github.com/kubernetes-csi/external-provisioner>
 - Implements CSI CO CreateVolume/DeleteVolume API
 - Attacher
 - <https://github.com/kubernetes-csi/external-attacher>
 - Implements CSI CO Publish/Unpublish Volume
 - Driver
 - <https://github.com/kubernetes-csi/drivers>
 - Will implement CSI SP API and call Cinder Volume and Attachment API to work with Provisioner and Attacher

Additional CSI thoughts

- Number of ways to do this
- Shim layer on top of Cinder to make it “talk CSI”
- Maybe Cinder version that at some point accepts CSI drivers and protocol?