

Red Hat
Summit

Connect

Introduction to Hosted Control Planes

a.k.a. HyperShift

Koray Şeremet

Senior Manager, Solution Architecture

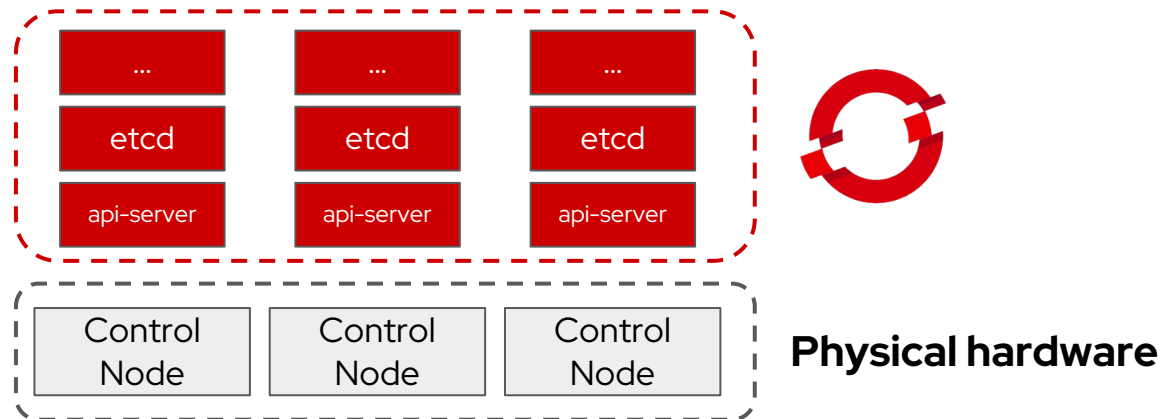
Overview

- ▶ What is Hypershift?
- ▶ Networking
- ▶ High availability
- ▶ Hypershift KubeVirt Provider
- ▶ Subscription Requirements

What Is Hypershift?

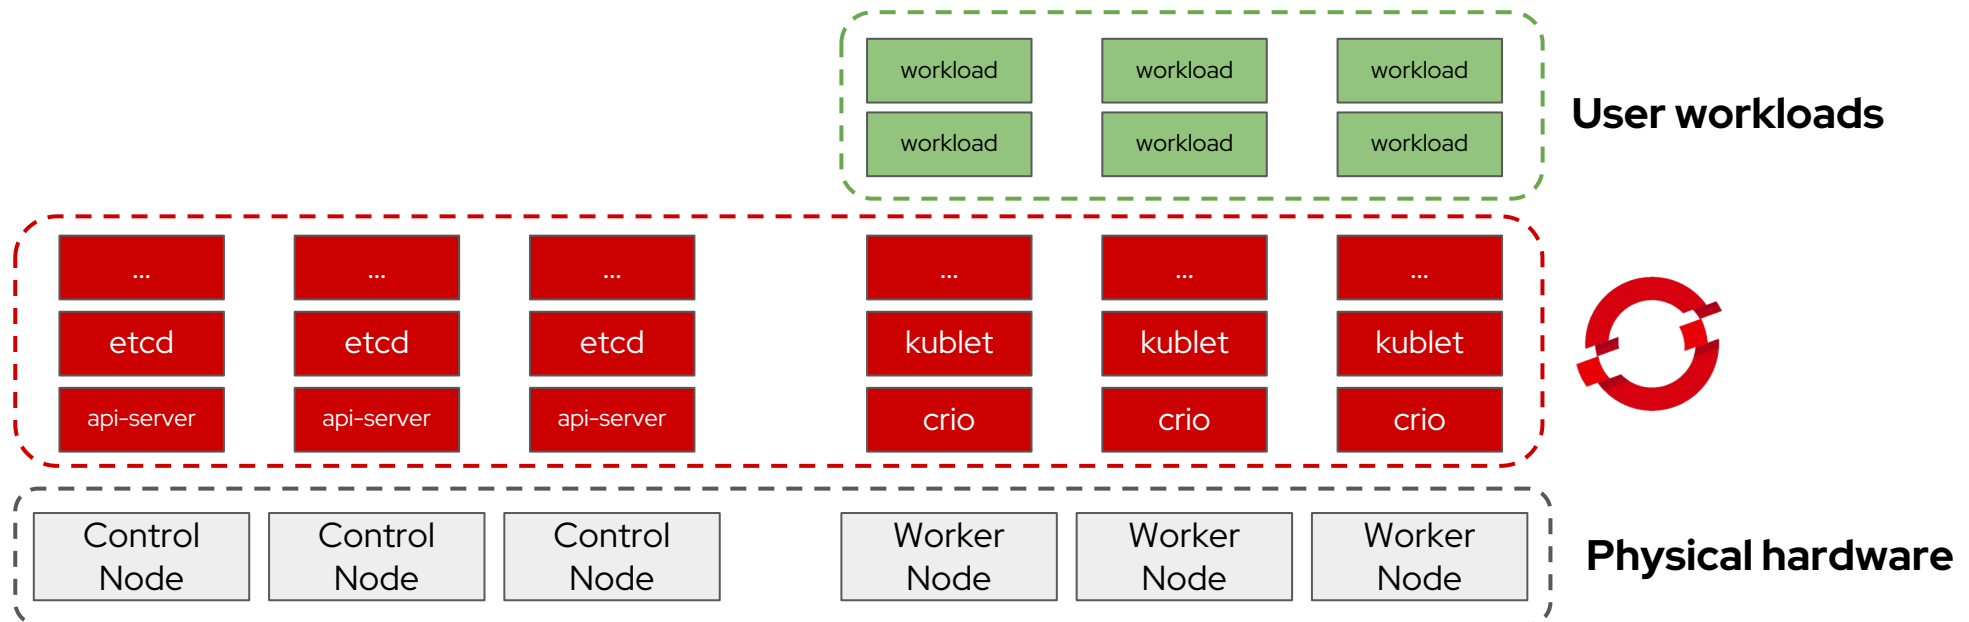
Standalone OpenShift

- ▶ Control Plane hosted across 3 machines



Standalone OpenShift

- ▶ Control Plane hosted across 3 machines
- ▶ Worker Nodes
- ▶ User Workloads



Use Cases



Ephemeral Clusters

Quickly (< 10 min) spin up/destroy clusters for CI and developers.



Clusters as a Service

On demand clusters driven by a declarative API



Cheaper Control Planes

Multiple Control Planes per node vs. 3 nodes for 1 Control Plane



Decoupled Life Cycle Management

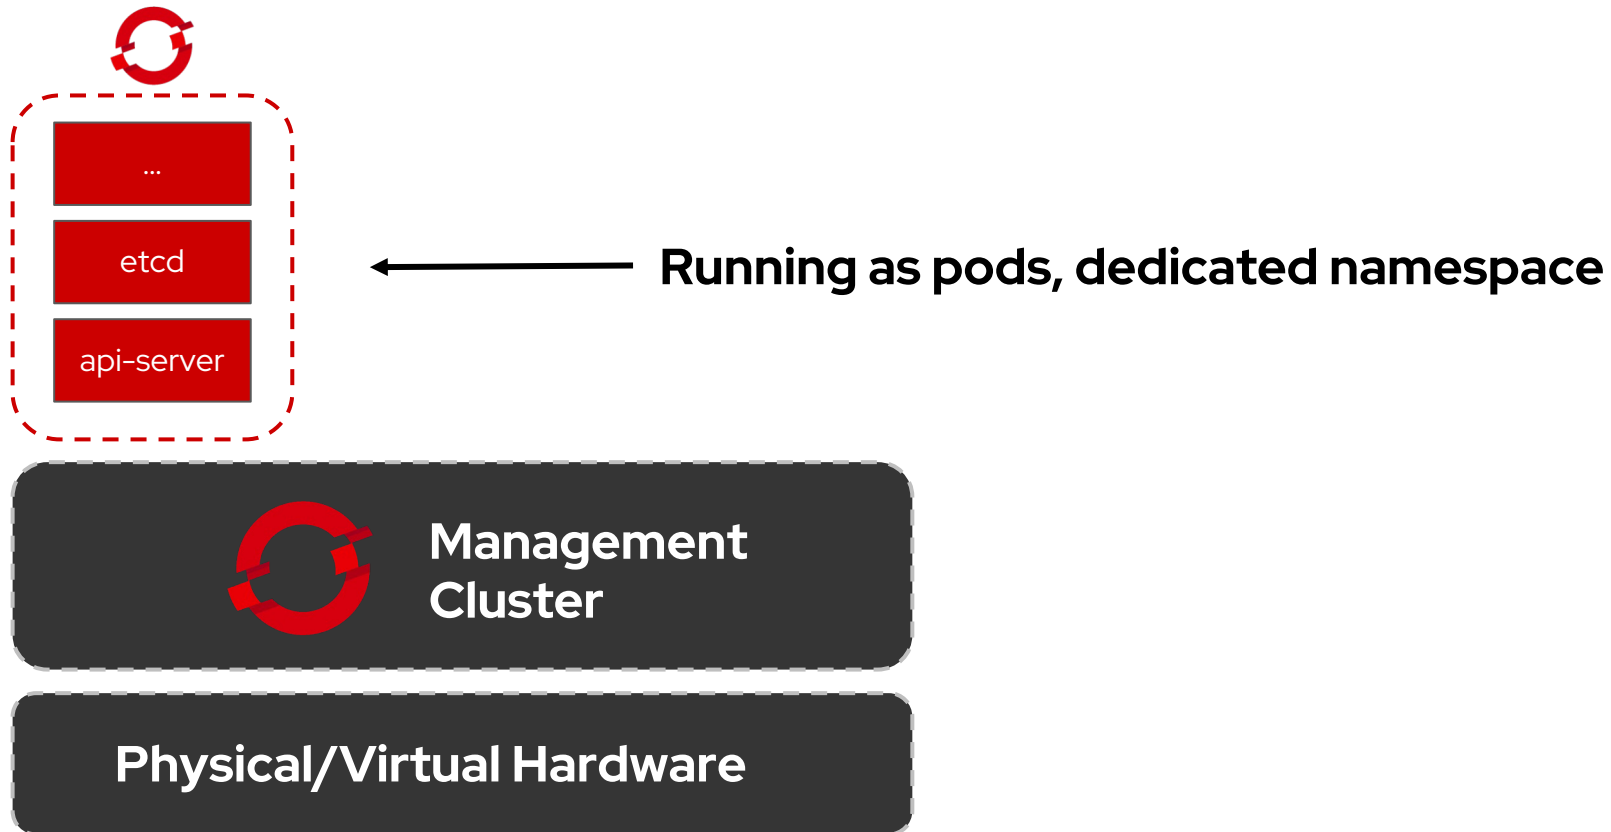
Upgrade the consolidated control planes out of cycle from the segmented worker nodes




Hosted Control Planes (HCP)

“Control Planes”

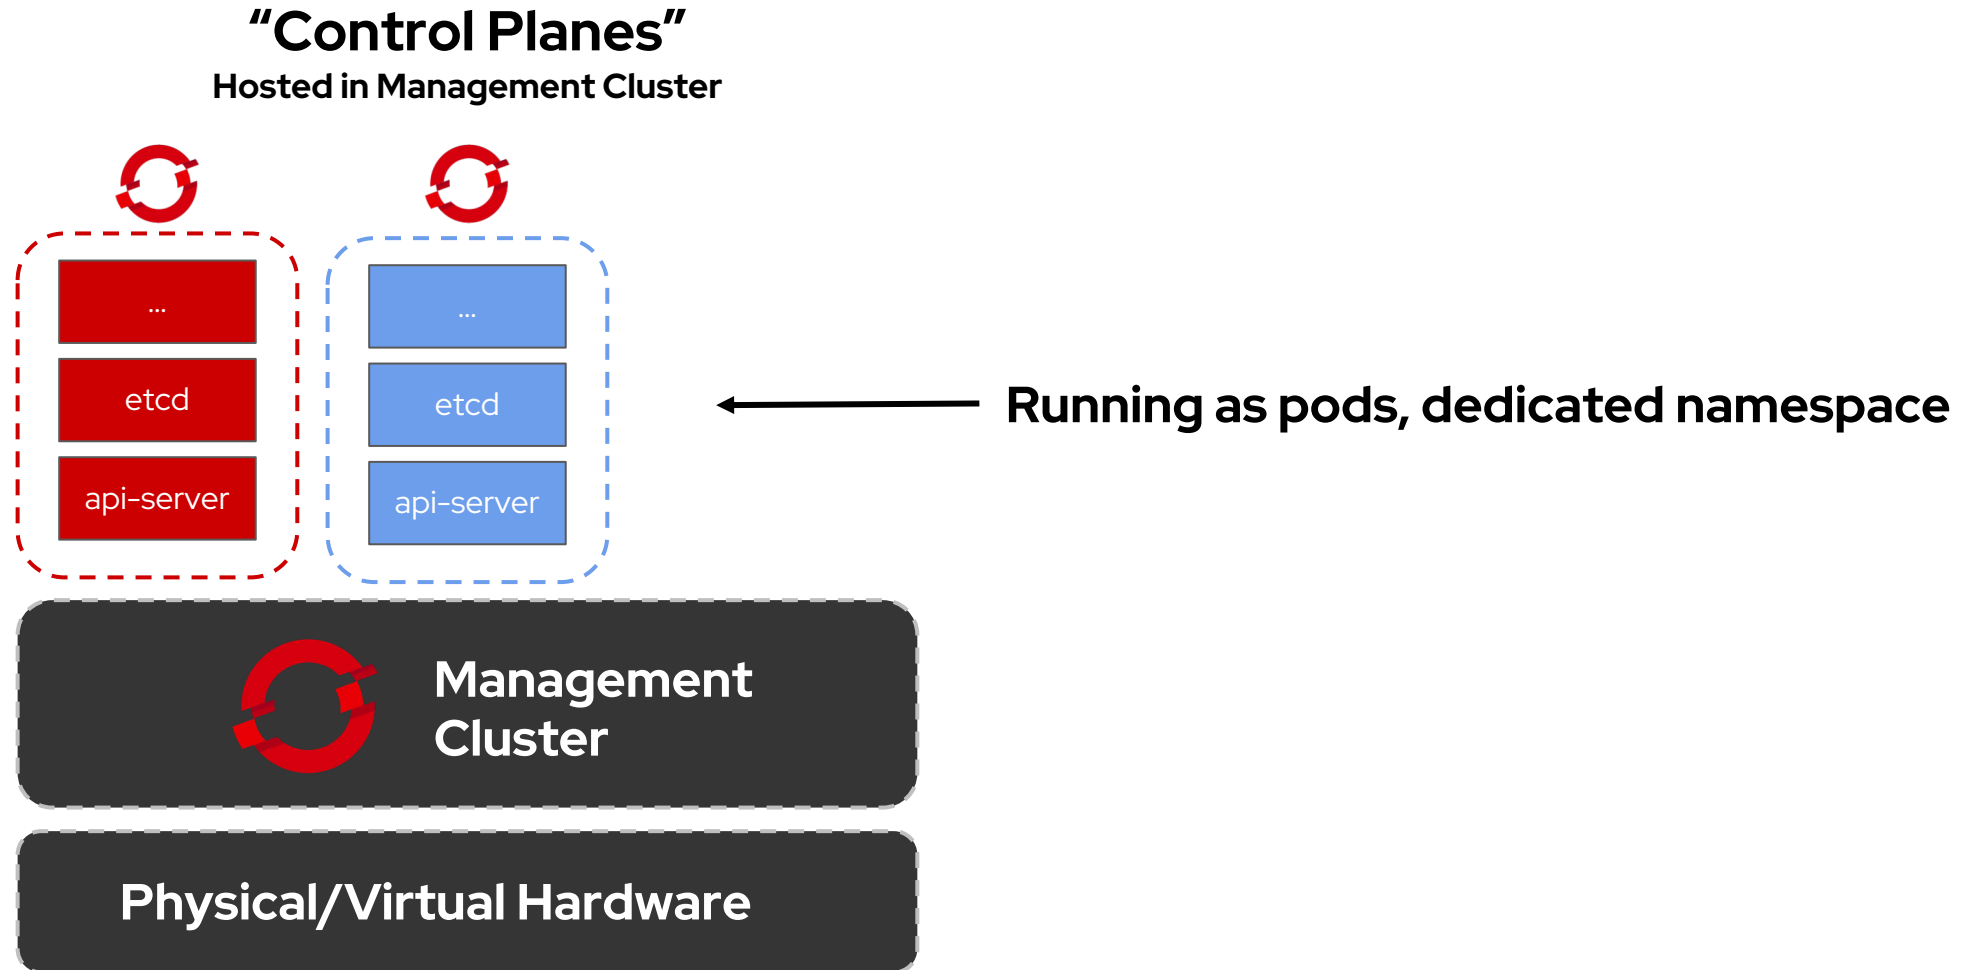
Hosted in Management Cluster



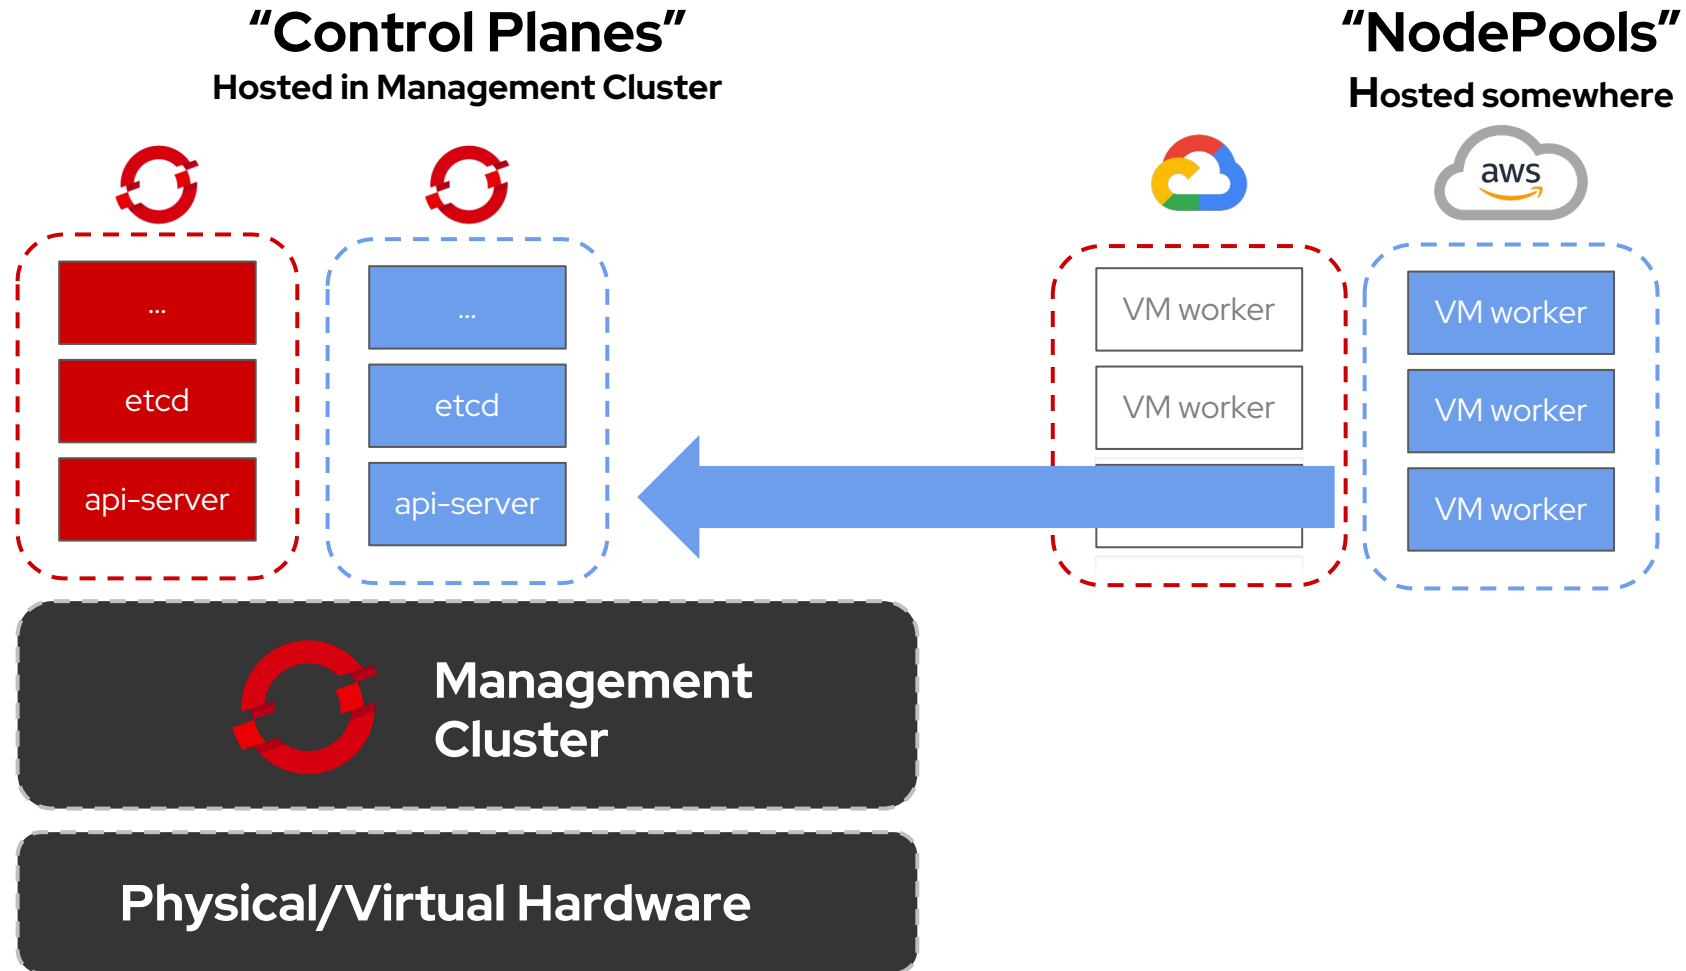
Nodes Register with HCP



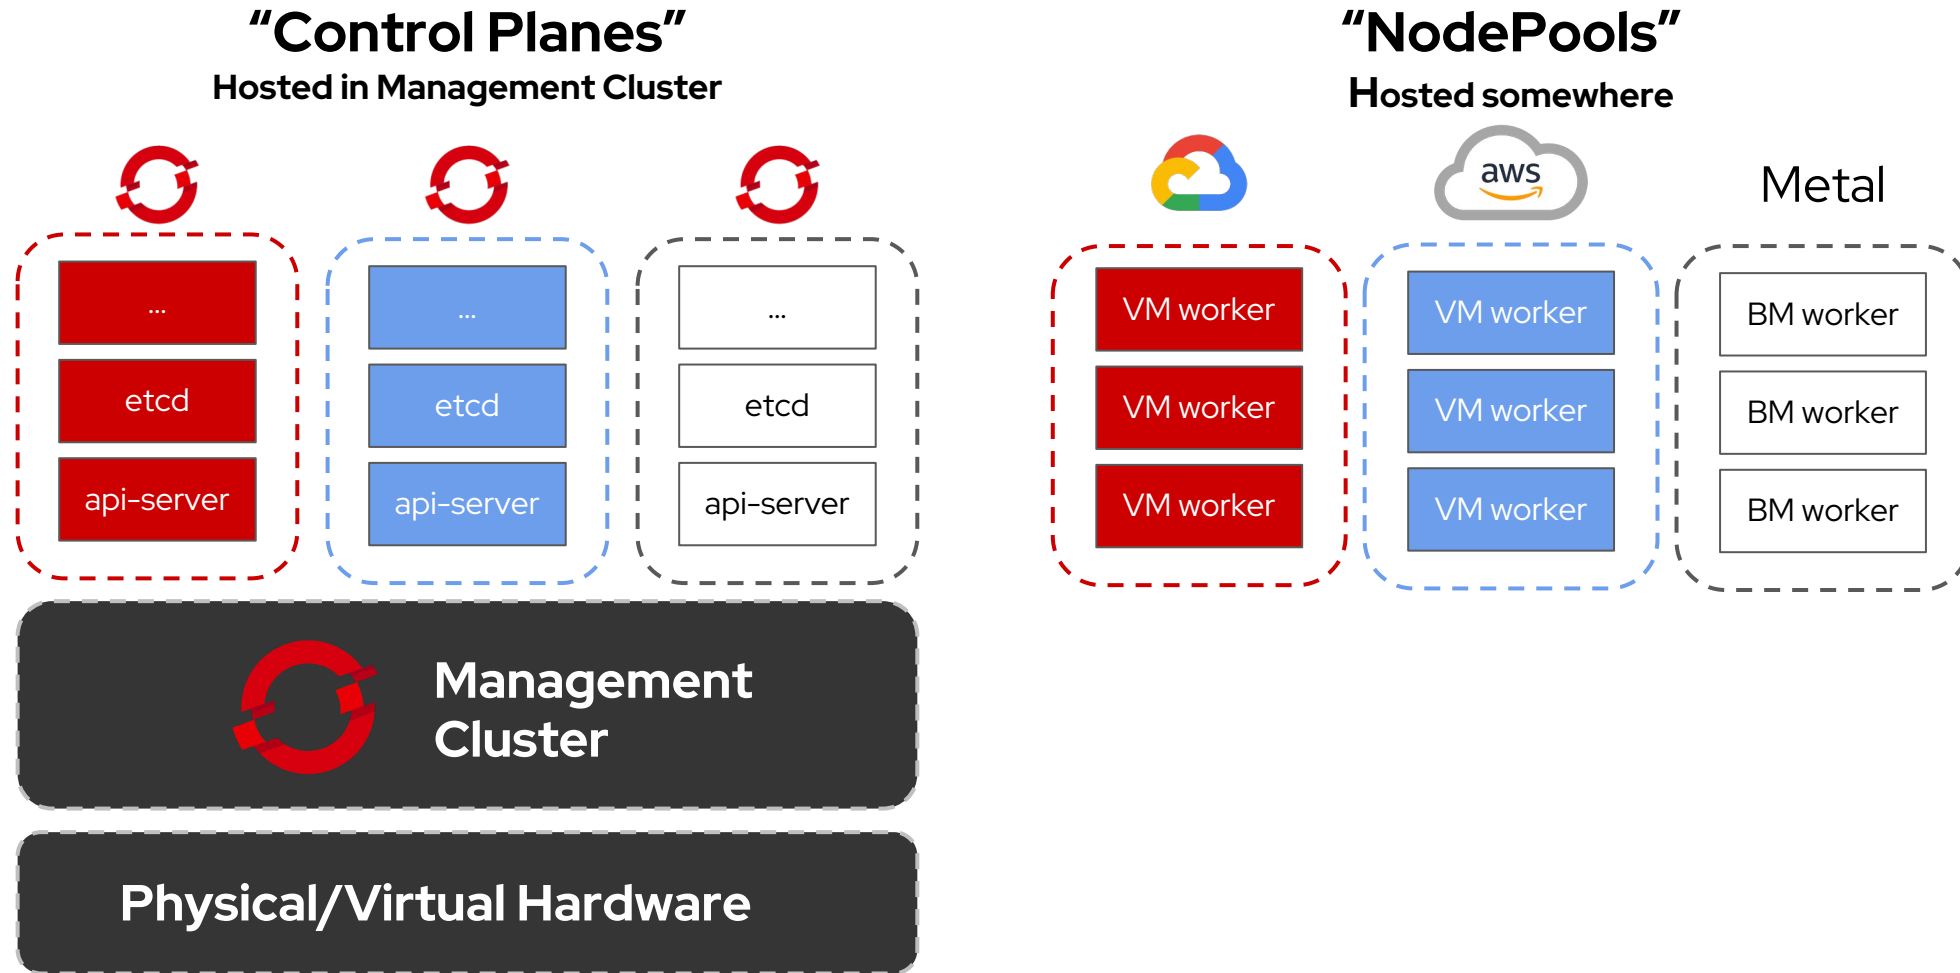
Hosted Control Planes (HCP)



Nodes Register with HCP



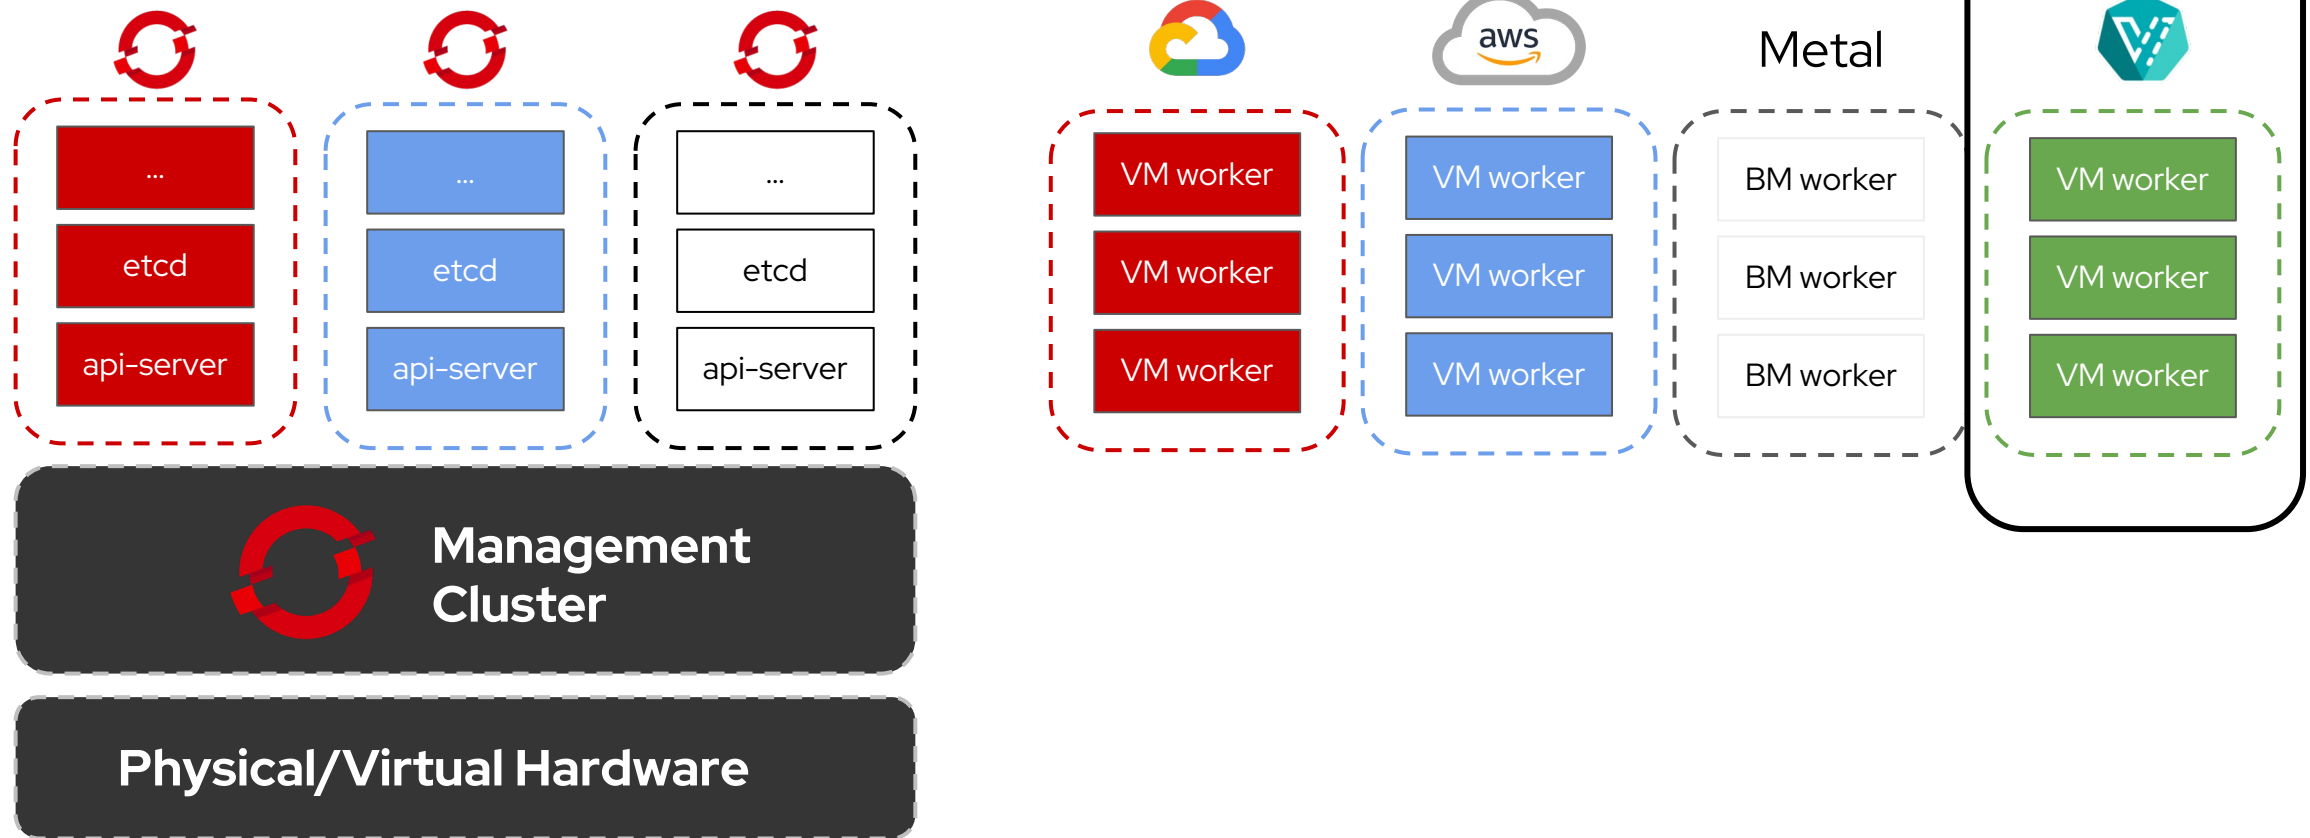
HyperShift



No Cloud Provider? No problem

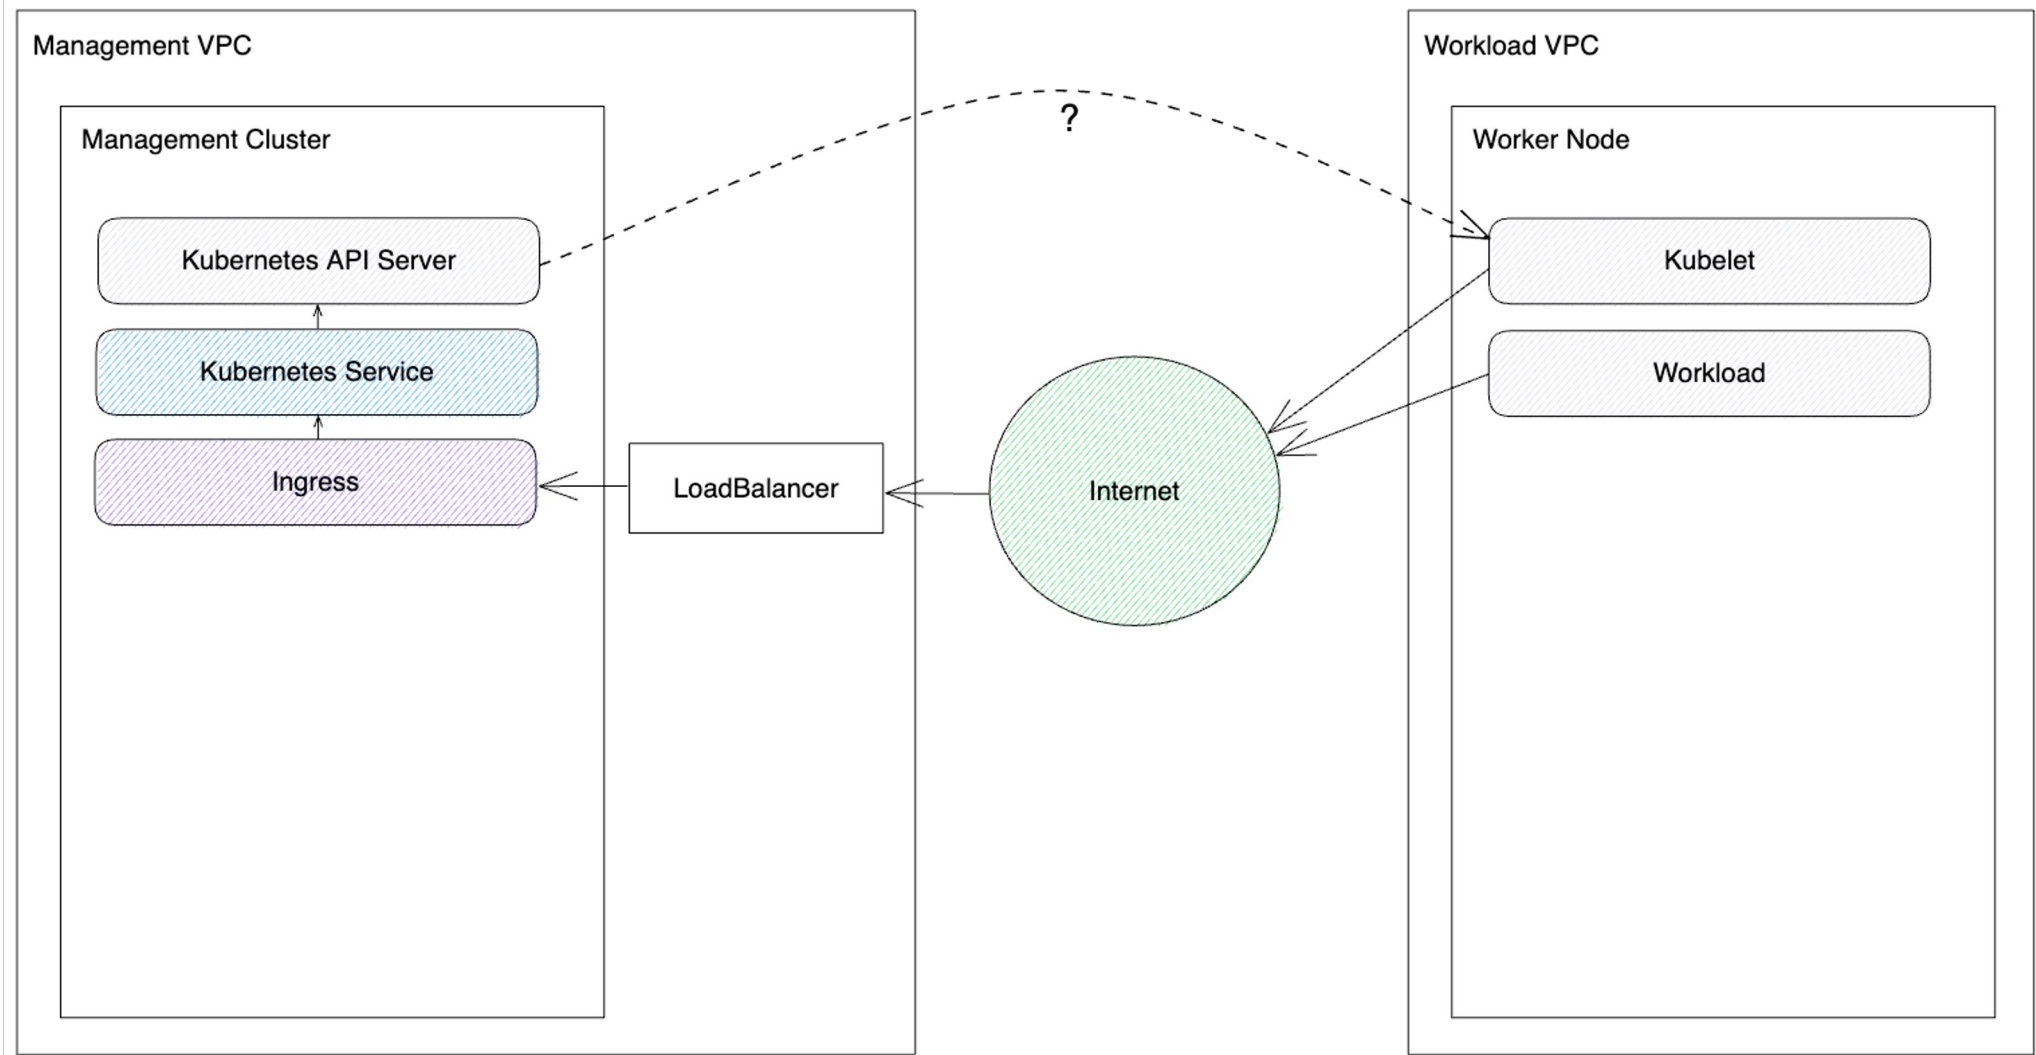
HyperShift

KubeVirt is just another provider



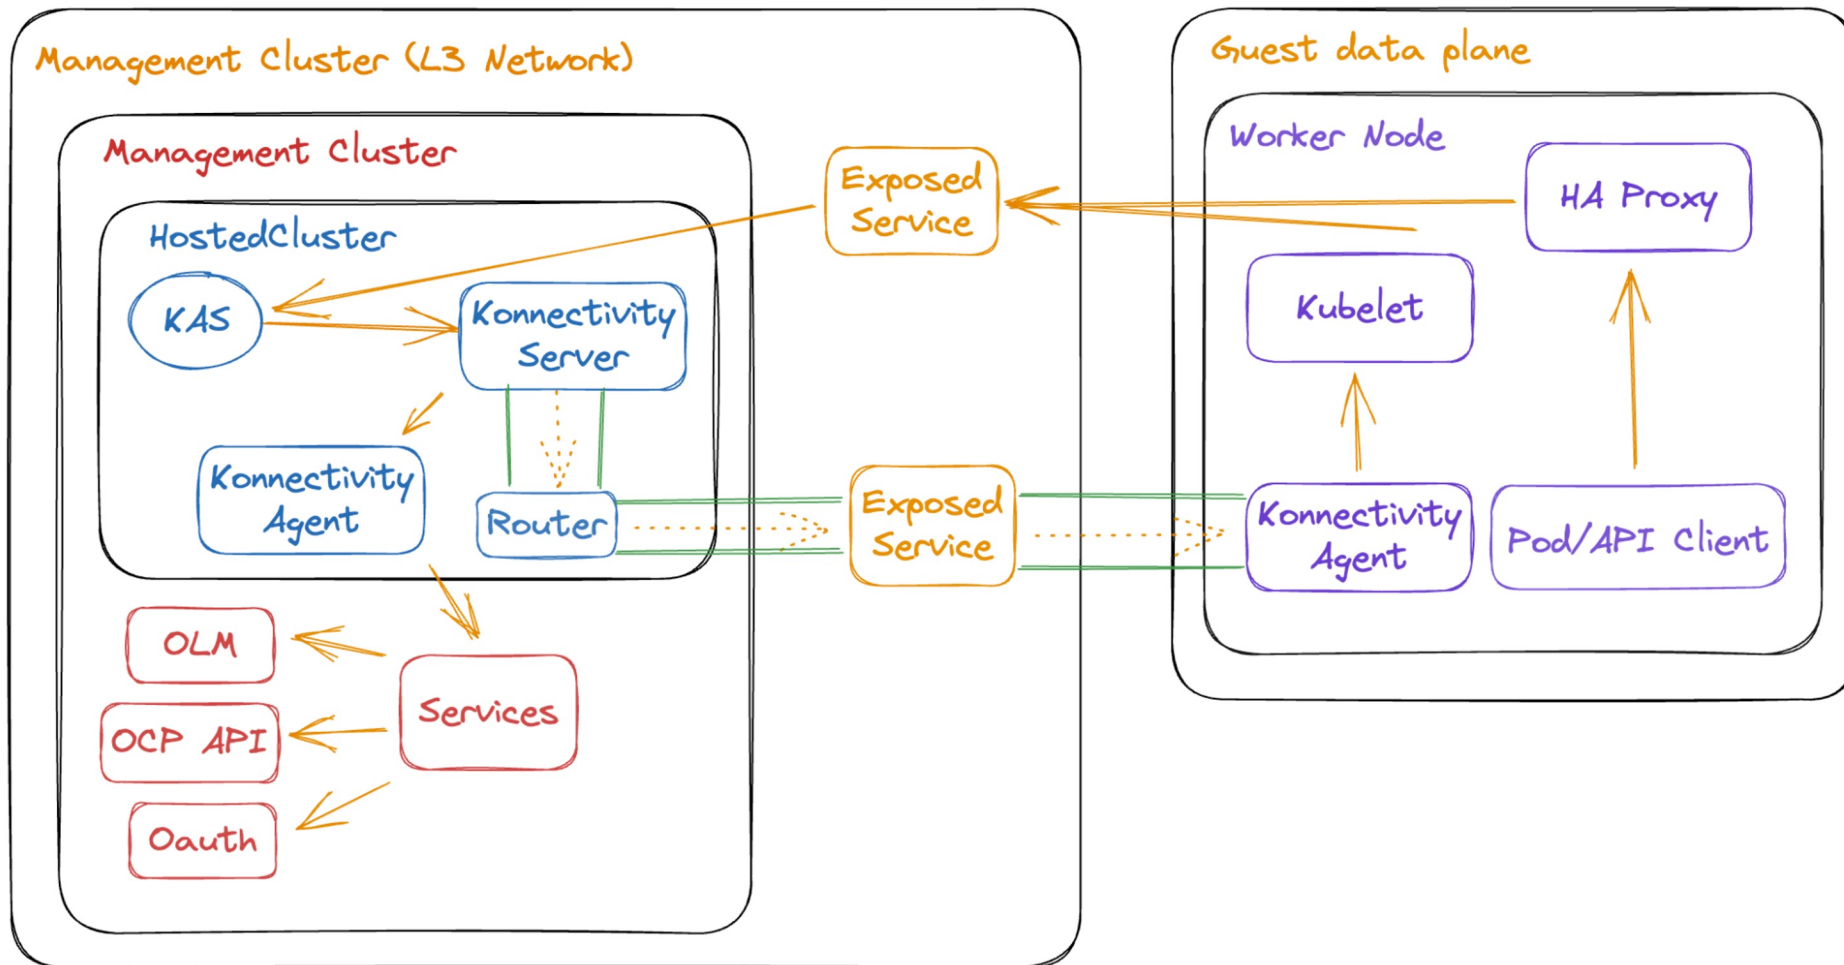
Networking

Networking



Networking

Networking between the management cluster and the hosted clusters


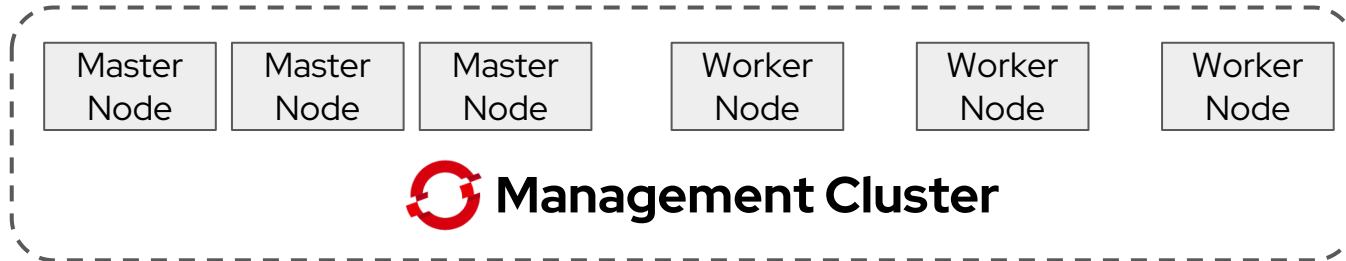
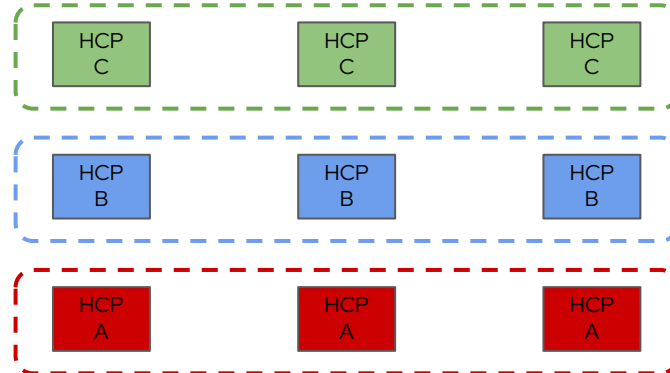


High Availability

High Availability

High-level summary of different failure scenarios

**“Hosted Control
Planes”**
Hosted in Management
Cluster

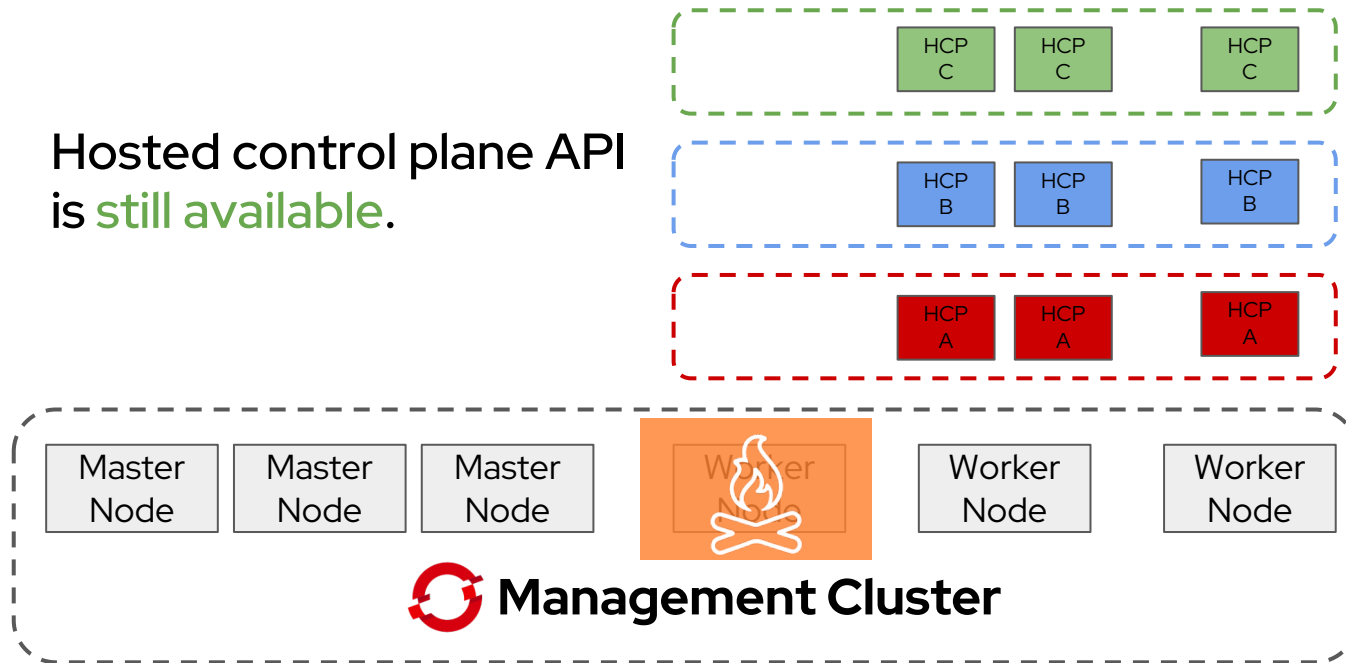


“NodePools”
Hosted somewhere


High Availability

Loss of management cluster worker

Hosted control plane API is **still available**.



Impacted hosted control plane components are **rescheduled**.

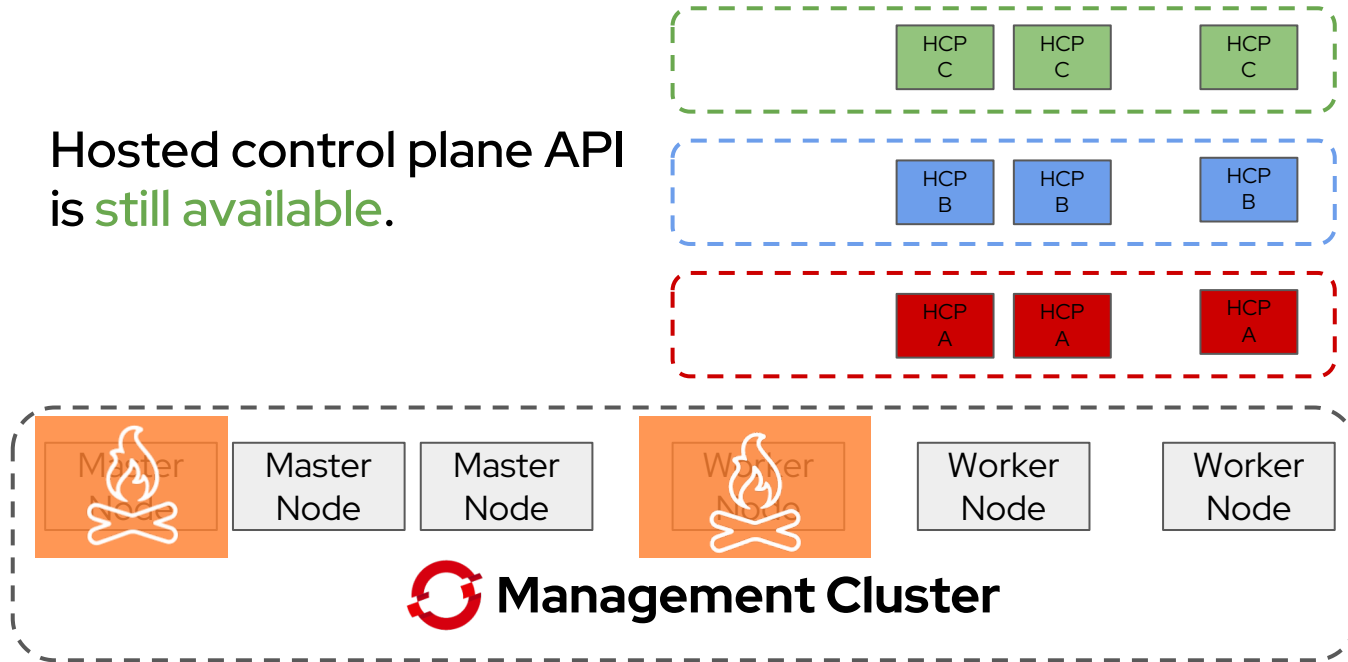


Hosted cluster data plane is **still available**.


High Availability

Loss of management cluster availability zone

Hosted control plane API is **still available**.



Impacted hosted control planes **maintain quorum**.


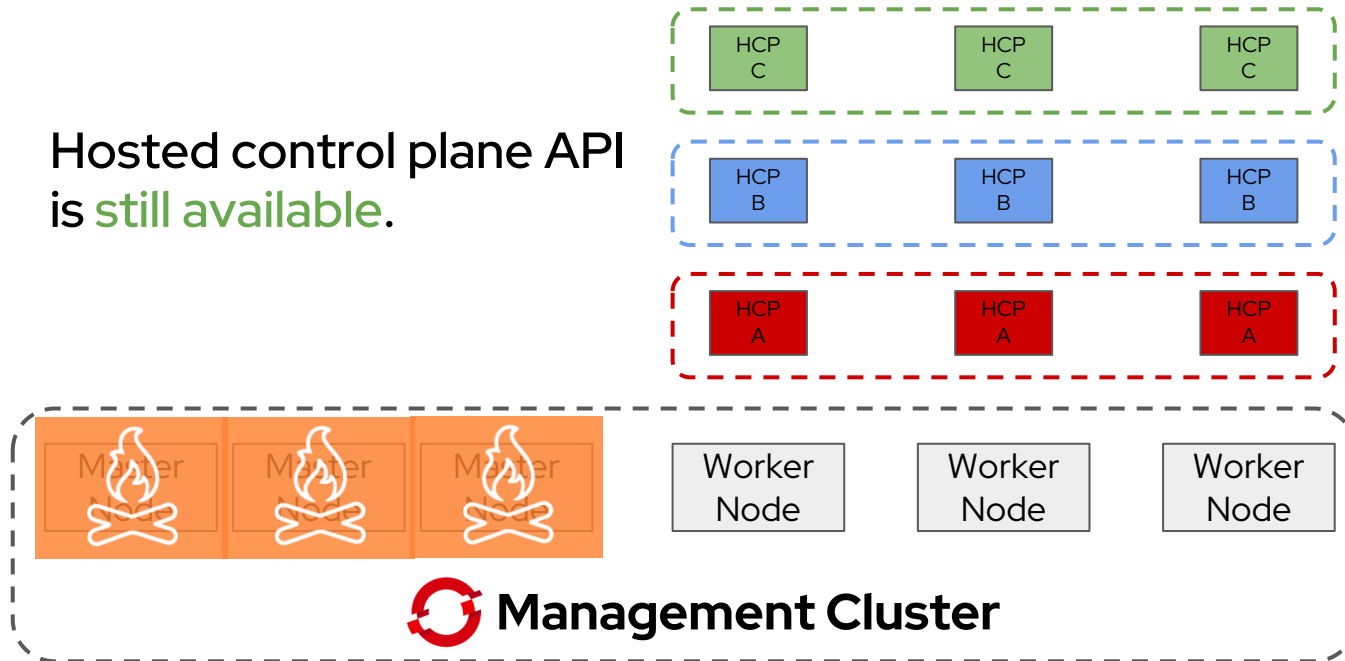


Hosted cluster data plane is **still available**.

High Availability

Loss of management cluster control plane


Hosted control plane API is **still available**.




Hosted cluster data plane is **still available**.

High Availability

Loss of management cluster control plane and workers



Hosted control plane API is **not available.**




Hosted cluster data plane is **still available.**

Hypershift KubeVirt Provider

Hypershift + Kubevirt

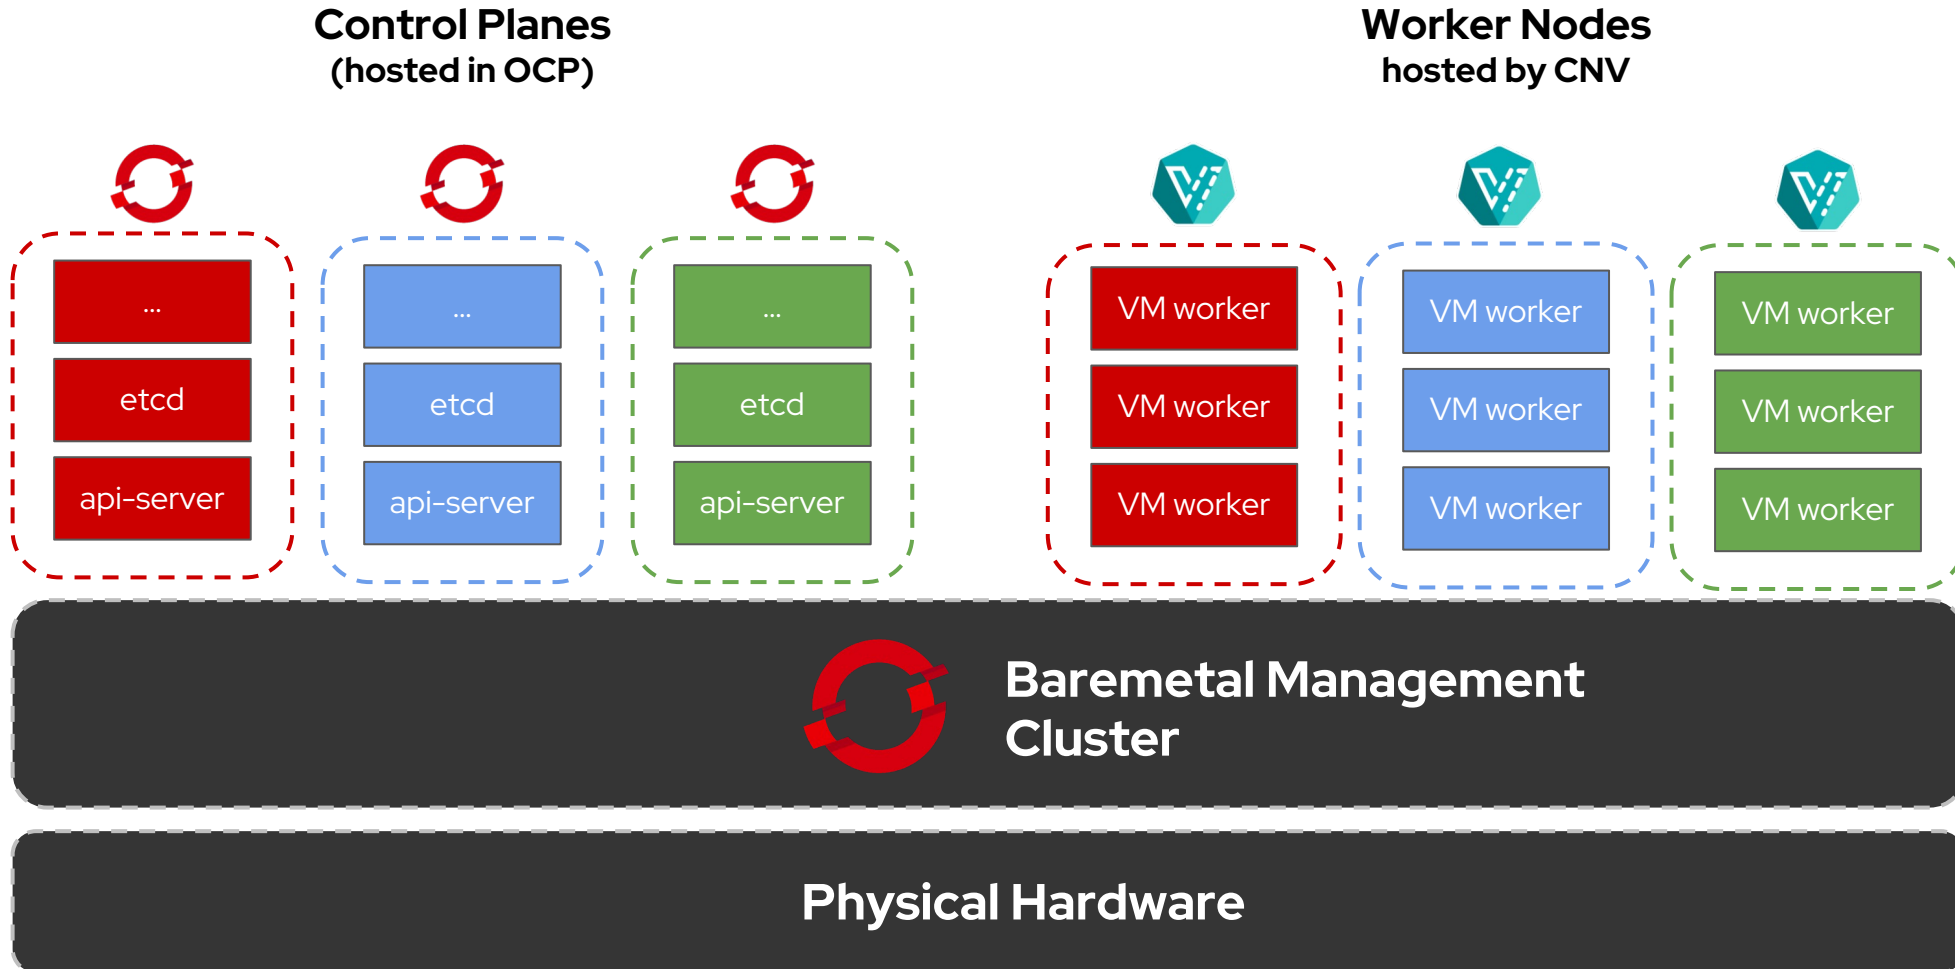
Control Planes
(hosted in OCP)



 **Baremetal Management Cluster**

Physical Hardware

Hypershift + Kubevirt



Why Hypershift with CNV

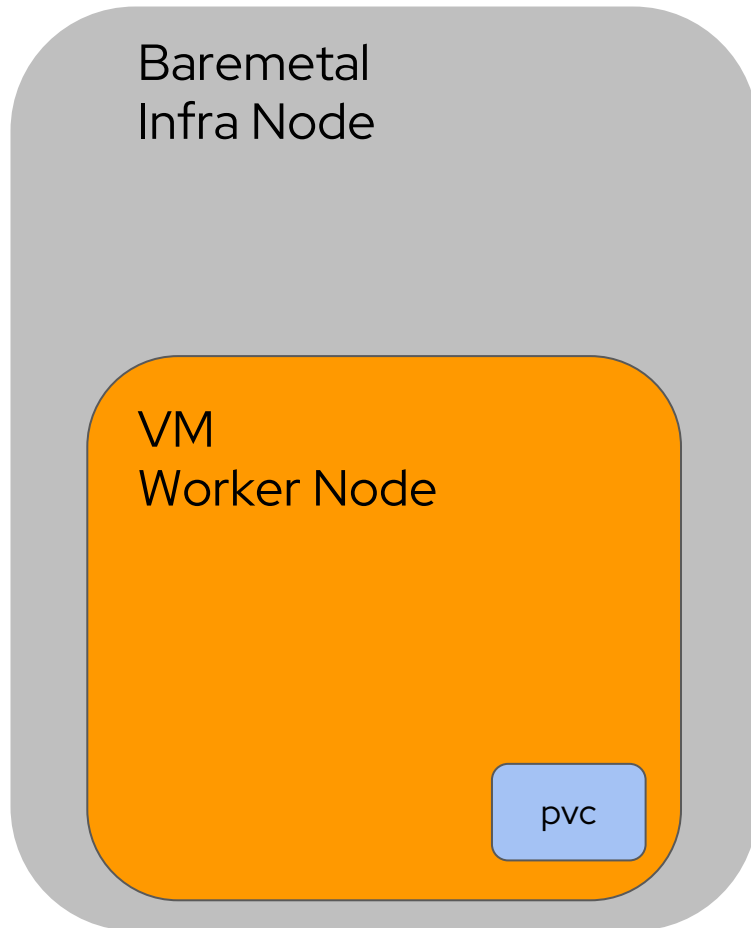
- ▶ **Multi Tenancy:** Give people entire clusters instead of namespaces.
- ▶ **API Isolation:** Hypershift solves the api isolation problem space.
- ▶ **Ephemeral Clusters:** quickly (< 10 min) spin up/destroy clusters for CI and developers
- ▶ **Clusters as a Service:** On demand clusters driven by a declarative API

KubeVirt CSI Driver



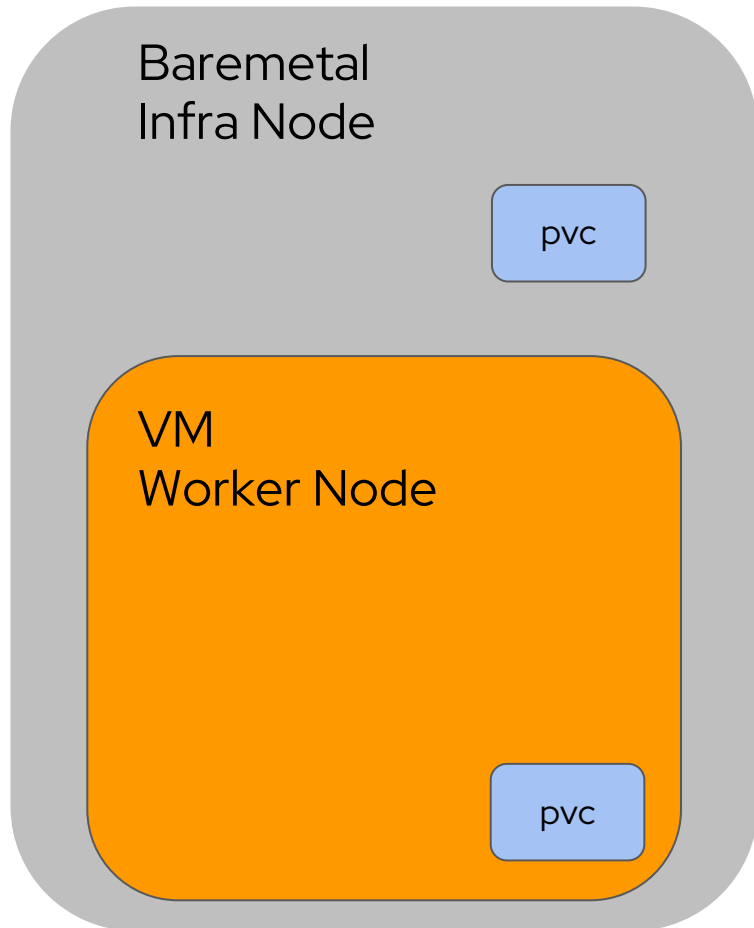
- ▶ Extends infra StorageClass into the guest clusters hosted by KubeVirt
- ▶ Utilizes HotPlug to make infra PVCs available within guest clusters
- ▶ Flow example...

KubeVirt CSI Driver



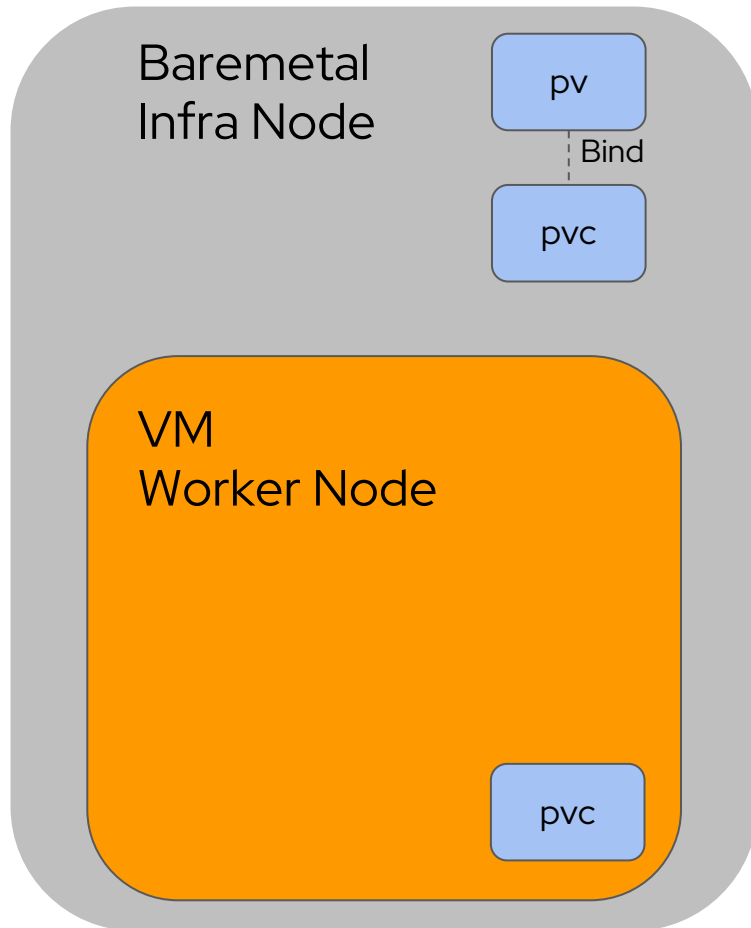
- ▶ Extends infra StorageClass into the guest clusters hosted by KubeVirt
- ▶ Utilizes HotPlug to make infra PVCs available within guest clusters
- ▶ Flow example...
 - User within guest cluster creates a PVC

KubeVirt CSI Driver



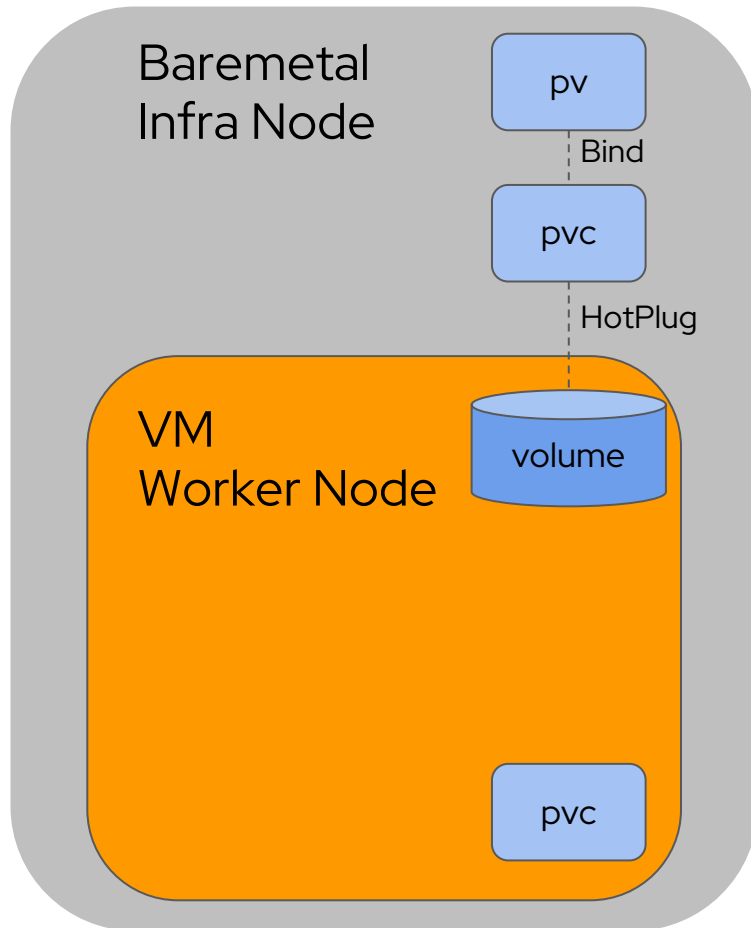
- ▶ Extends infra StorageClass into the guest clusters hosted by KubeVirt
- ▶ Utilizes HotPlug to make infra PVCs available within guest clusters
- ▶ Flow example...
 - User within guest cluster creates a PVC
 - KubeVirt CSI driver mirrors this PVC to the infra cluster

KubeVirt CSI Driver



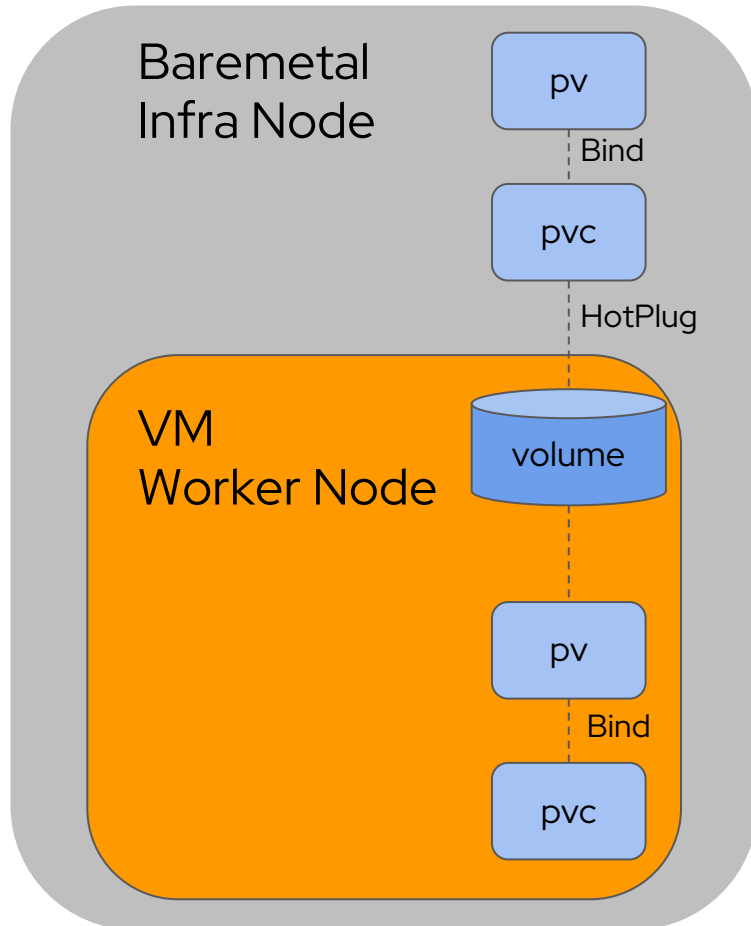
- ▶ Extends infra StorageClass into the guest clusters hosted by KubeVirt
- ▶ Utilizes HotPlug to make infra PVCs available within guest clusters
- ▶ Flow example...
 - User within guest cluster creates a PVC
 - KubeVirt CSI driver mirrors this PVC to the infra cluster
 - Infra cluster's dynamic storage provisioner creates the PV and binds it to PVC

KubeVirt CSI Driver



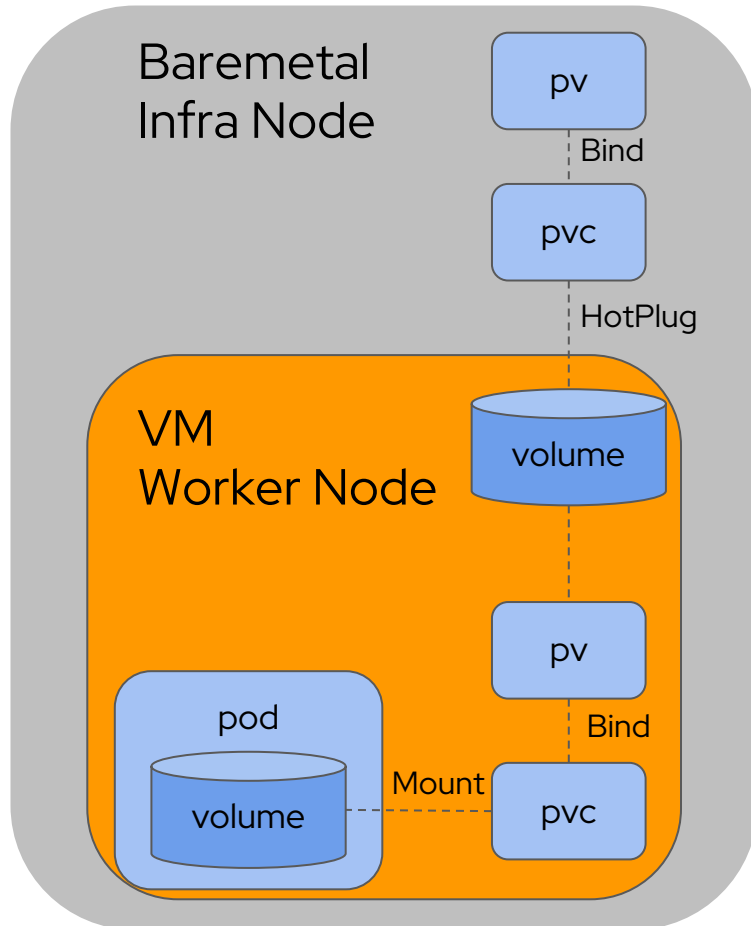
- ▶ Extends infra StorageClass into the guest clusters hosted by KubeVirt
- ▶ Utilizes HotPlug to make infra PVCs available within guest clusters
- ▶ Flow example...
 - User within guest cluster creates a PVC
 - KubeVirt CSI driver mirrors this PVC to the infra cluster
 - Infra cluster's dynamic storage provisioner creates the PV and binds it to PVC
 - KubeVirt CSI HotPlugs the PVC to the VM

KubeVirt CSI Driver



- ▶ Extends infra StorageClass into the guest clusters hosted by KubeVirt
- ▶ Utilizes HotPlug to make infra PVCs available within guest clusters
- ▶ Flow example...
 - User within guest cluster creates a PVC
 - KubeVirt CSI driver mirrors this PVC to the infra cluster
 - Infra cluster's dynamic storage provisioner creates the PV and binds it to PVC
 - KubeVirt CSI HotPlugs the PVC to the VM
 - Volume becomes a PV and is bound to PVC within Guest Cluster

KubeVirt CSI Driver




- ▶ Extends infra StorageClass into the guest clusters hosted by KubeVirt
- ▶ Utilizes HotPlug to make infra PVCs available within guest clusters
- ▶ Flow example...
 - User within guest cluster creates a PVC
 - KubeVirt CSI driver mirrors this PVC to the infra cluster
 - Infra cluster's dynamic storage provisioner creates the PV and binds it to PVC
 - KubeVirt CSI HotPlugs the PVC to the VM
 - Guest PVC is attached to a pod workload

Cloud Provider KubeVirt




- ▶ Provides Load Balancer support to KubeVirt guest clusters
- ▶ Similar to KubeVirt CSI in that it is mirroring infra capabilities to guest clusters.

Cloud Provider KubeVirt



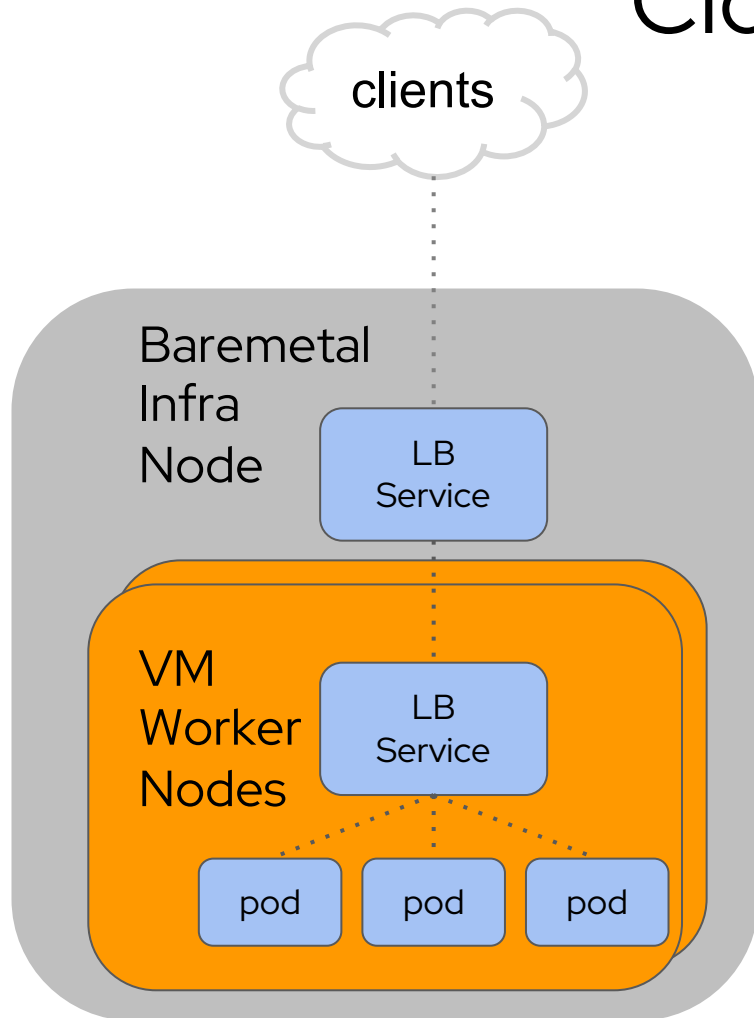
- ▶ Provides Load Balancer support to KubeVirt guest clusters
- ▶ Similar to KubeVirt CSI in that it is mirroring infra capabilities to guest clusters.
- ▶ Flow example...
 - User within guest cluster creates a LoadBalancer service

Cloud Provider KubeVirt



- ▶ Provides Load Balancer support to KubeVirt guest clusters
- ▶ Similar to KubeVirt CSI in that it is mirroring infra capabilities to guest clusters.
- ▶ Flow example...
 - User within guest cluster creates a LoadBalancer service
 - Cloud Provider Kubevirt controller creates corresponding LB on infra cluster

Cloud Provider KubeVirt



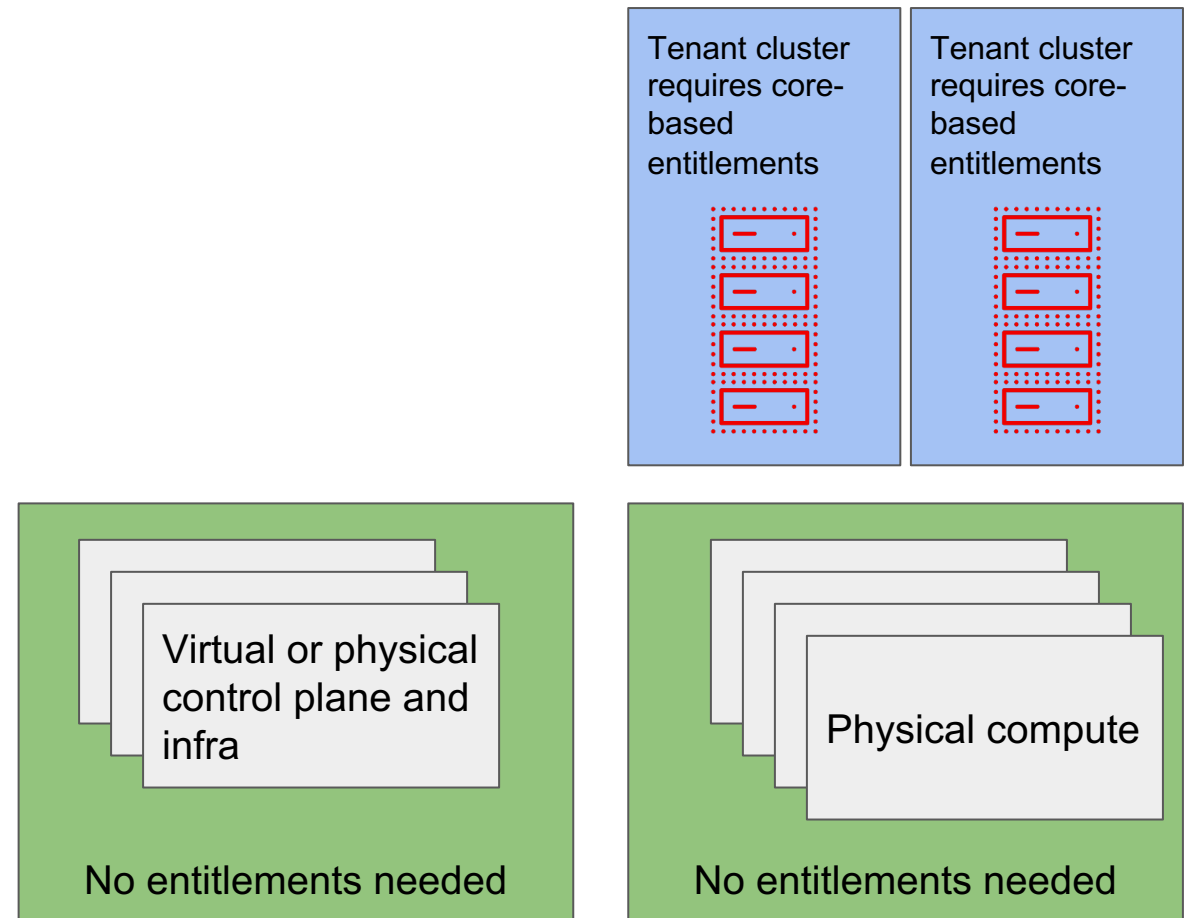
- ▶ Provides Load Balancer support to KubeVirt guest clusters
- ▶ Similar to KubeVirt CSI in that it is mirroring infra capabilities to guest clusters.
- ▶ Flow example...
 - User within guest cluster creates a LoadBalancer service
 - Cloud Provider Kubevirt controller creates corresponding LB on infra cluster
 - Infa LB maps to guest cluster VM pods to pass traffic to guest cluster LB

Subscription Requirements

Hosted Cluster Subscription Requirement

Only tenant clusters on Bare metal worker nodes

- Customers need **core-based entitlements for all virtualized** OpenShift compute nodes
- **No entitlement** needed for control plane or underlying physical nodes (as long as no other workload on the bare metal node)



Red Hat
Summit

Connect

Thank you



[linkedin.com/company/red-hat](https://www.linkedin.com/company/red-hat)



[facebook.com/redhatinc](https://www.facebook.com/redhatinc)



[youtube.com/user/RedHatVideos](https://www.youtube.com/user/RedHatVideos)



twitter.com/RedHat

Breakout Session Feedback



<https://forms.gle/2nCjKPgQBHP7jEoS9>