Red Hat Summit

Connect

# OpenShift Hosted Control Planes

21 November 2024

Gokhan Goksu
Senior Solution Architect, App Platforms
gokhan@redhat.com
www.linkedin.com/in/ggoksu

Red Hat

# Agenda

- What is Hosted Control Planes?
- Networking
- High Availability
- Upgrades
- KubeVirt Provider
- KubeVirt FAQ

# Standalone Openshift

- Control Plane hosted across 3 machines

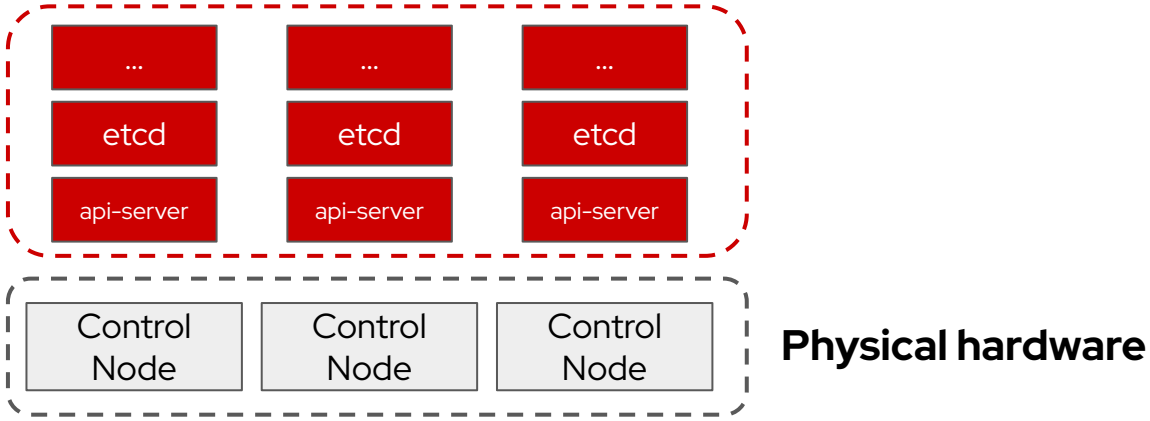




Red Hat

# Standalone Openshift

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

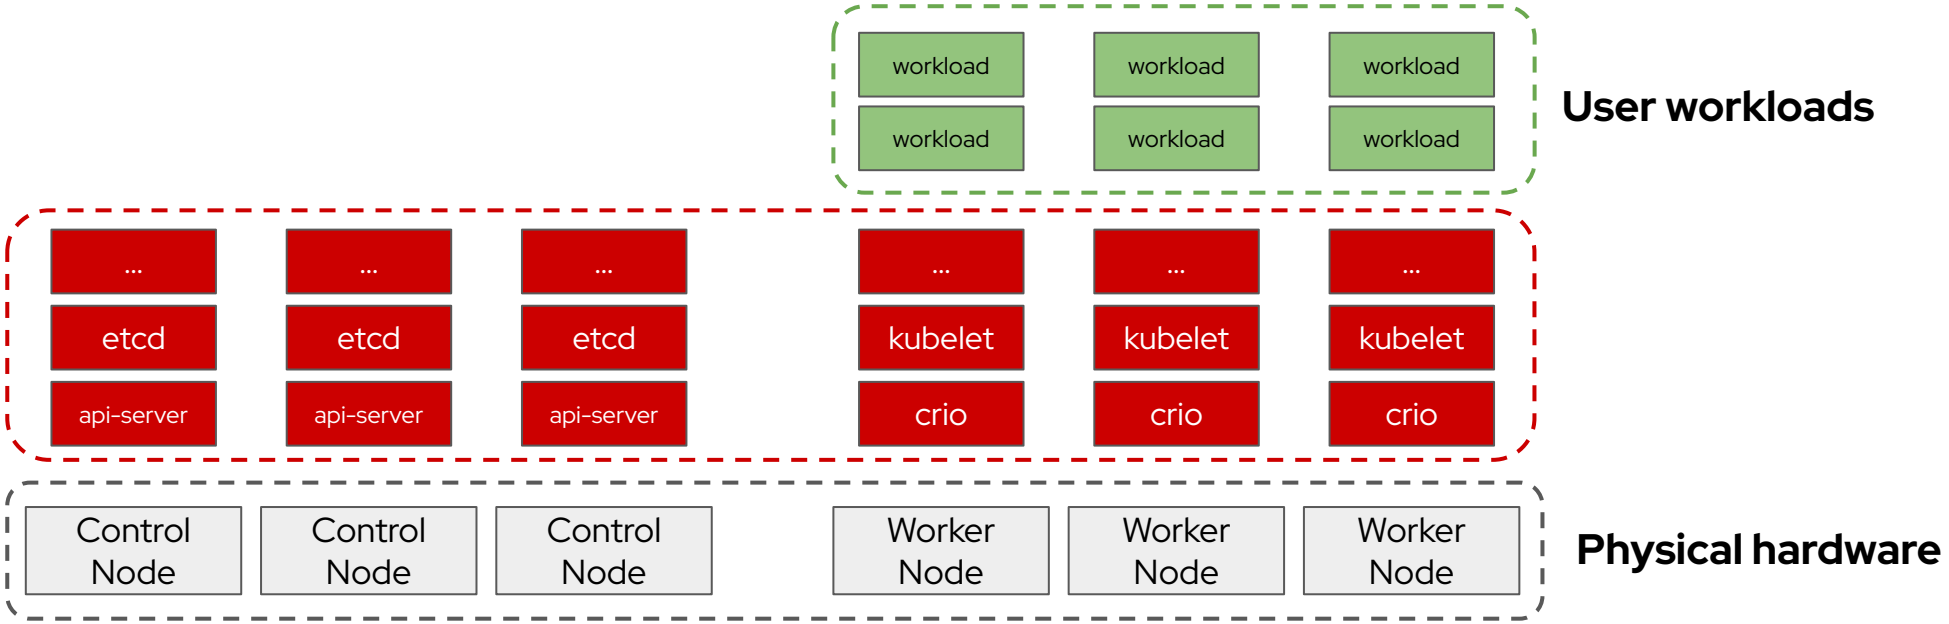




Red Hat

# Compact Openshift

- Self-contained
- Highly available control-plane
- Less room for actual workloads
- Ideal for resource-constrained environments

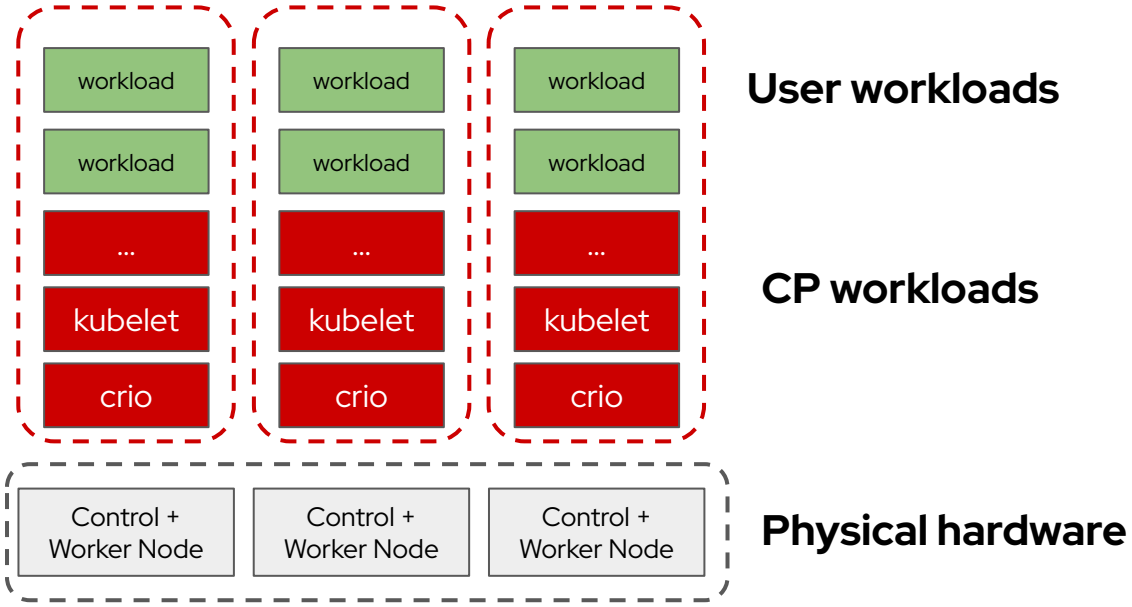




Red Hat

# Single-Node Openshift (SNO)

- Self-contained
- No High Availability
- Resource Constraints
- Edge Locations

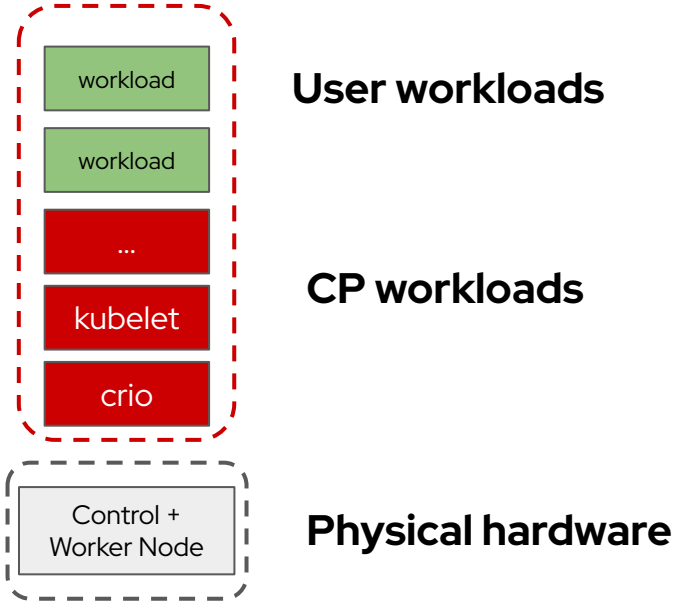




Red Hat

# Hosted Control Planes (HCP)

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

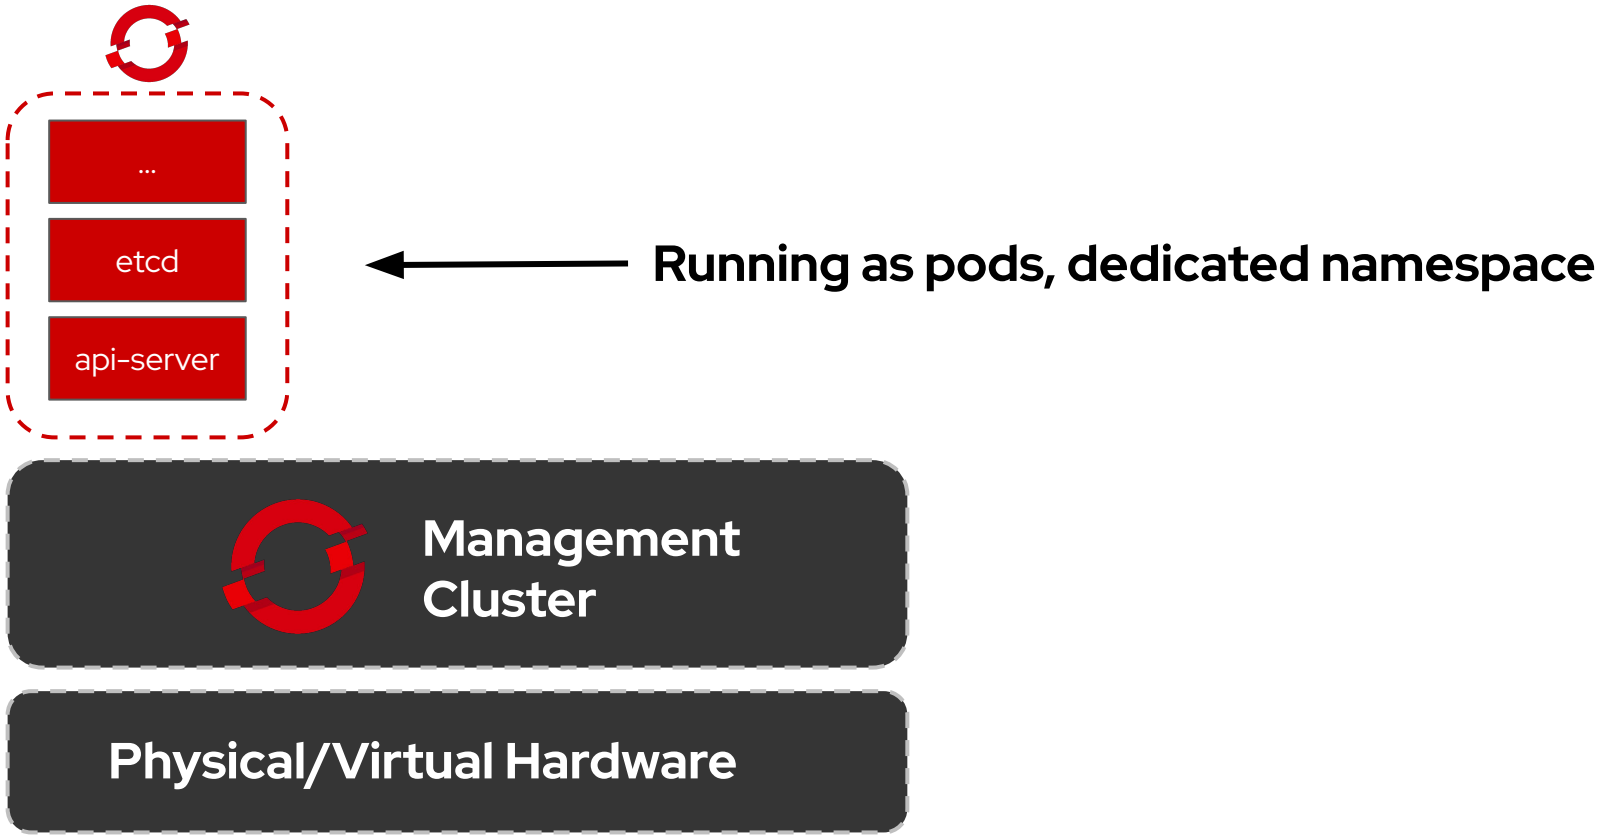




Red Hat

# Nodes Register with HCP

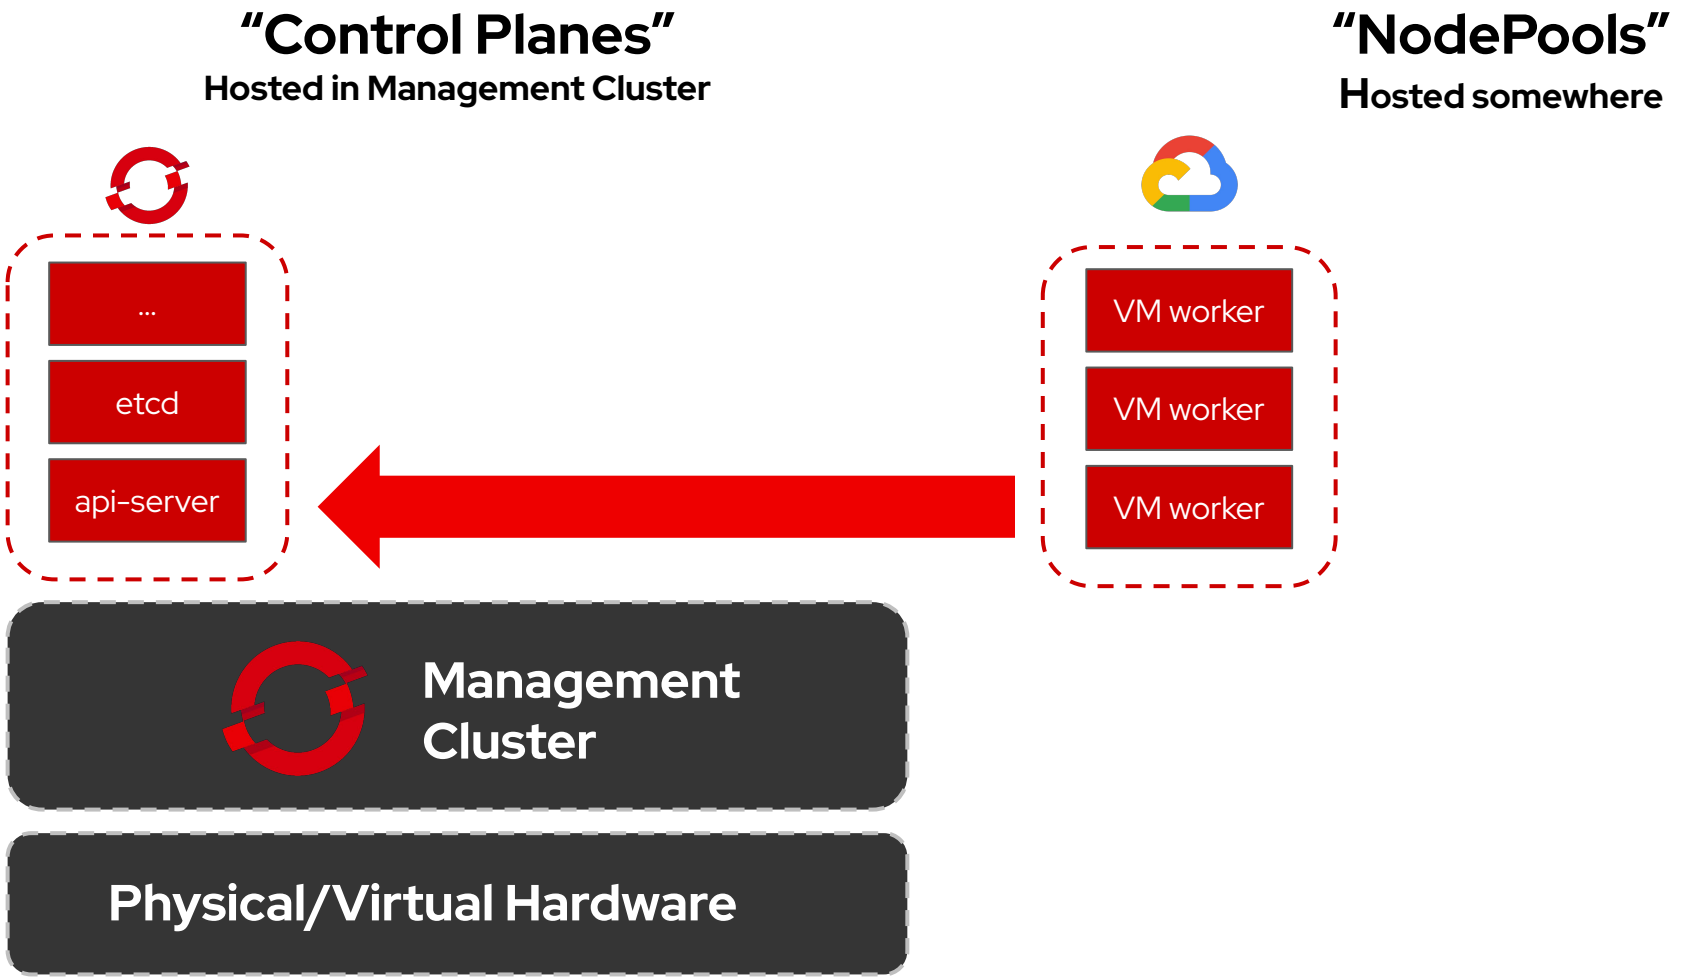




Red Hat

# HCP

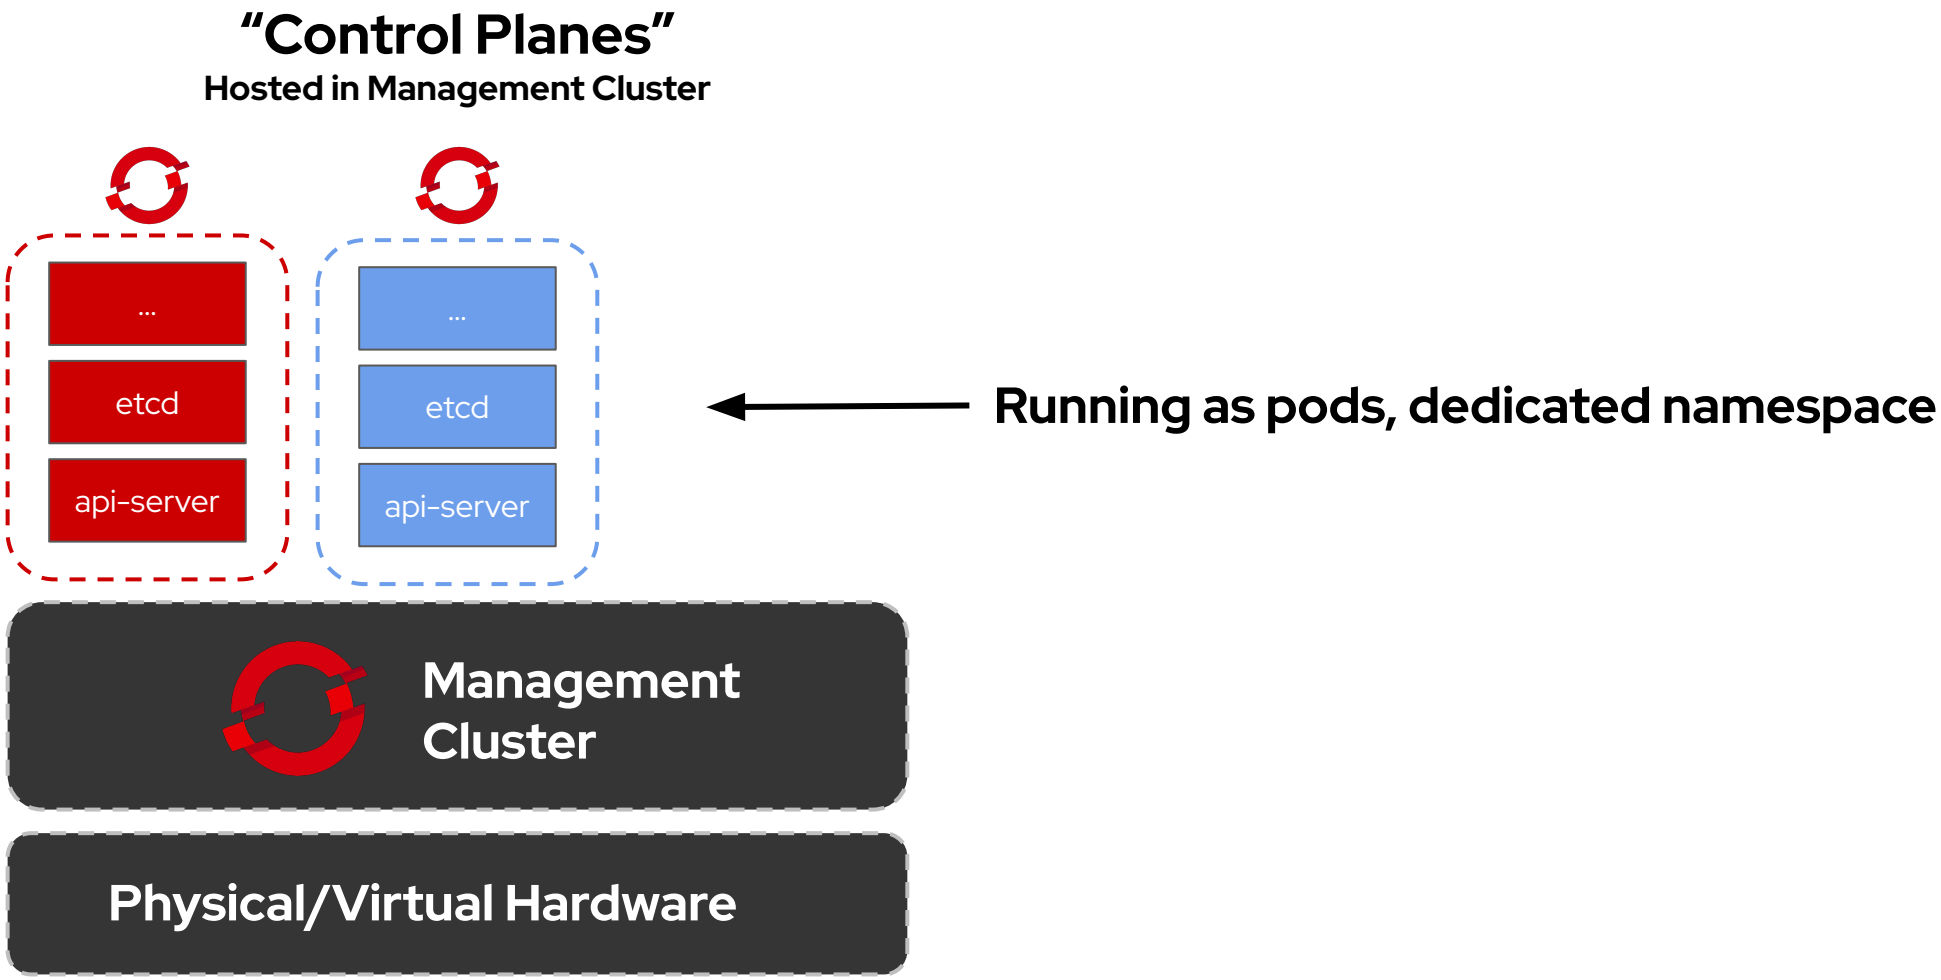
"Control Planes"
Hosted in Management Cluster
...
etcd
api-server
...
etcd
api-server
Running as pods, dedicated namespace
Management Cluster
Physical/Virtual Hardware



Red Hat

# Nodes Register with HCP

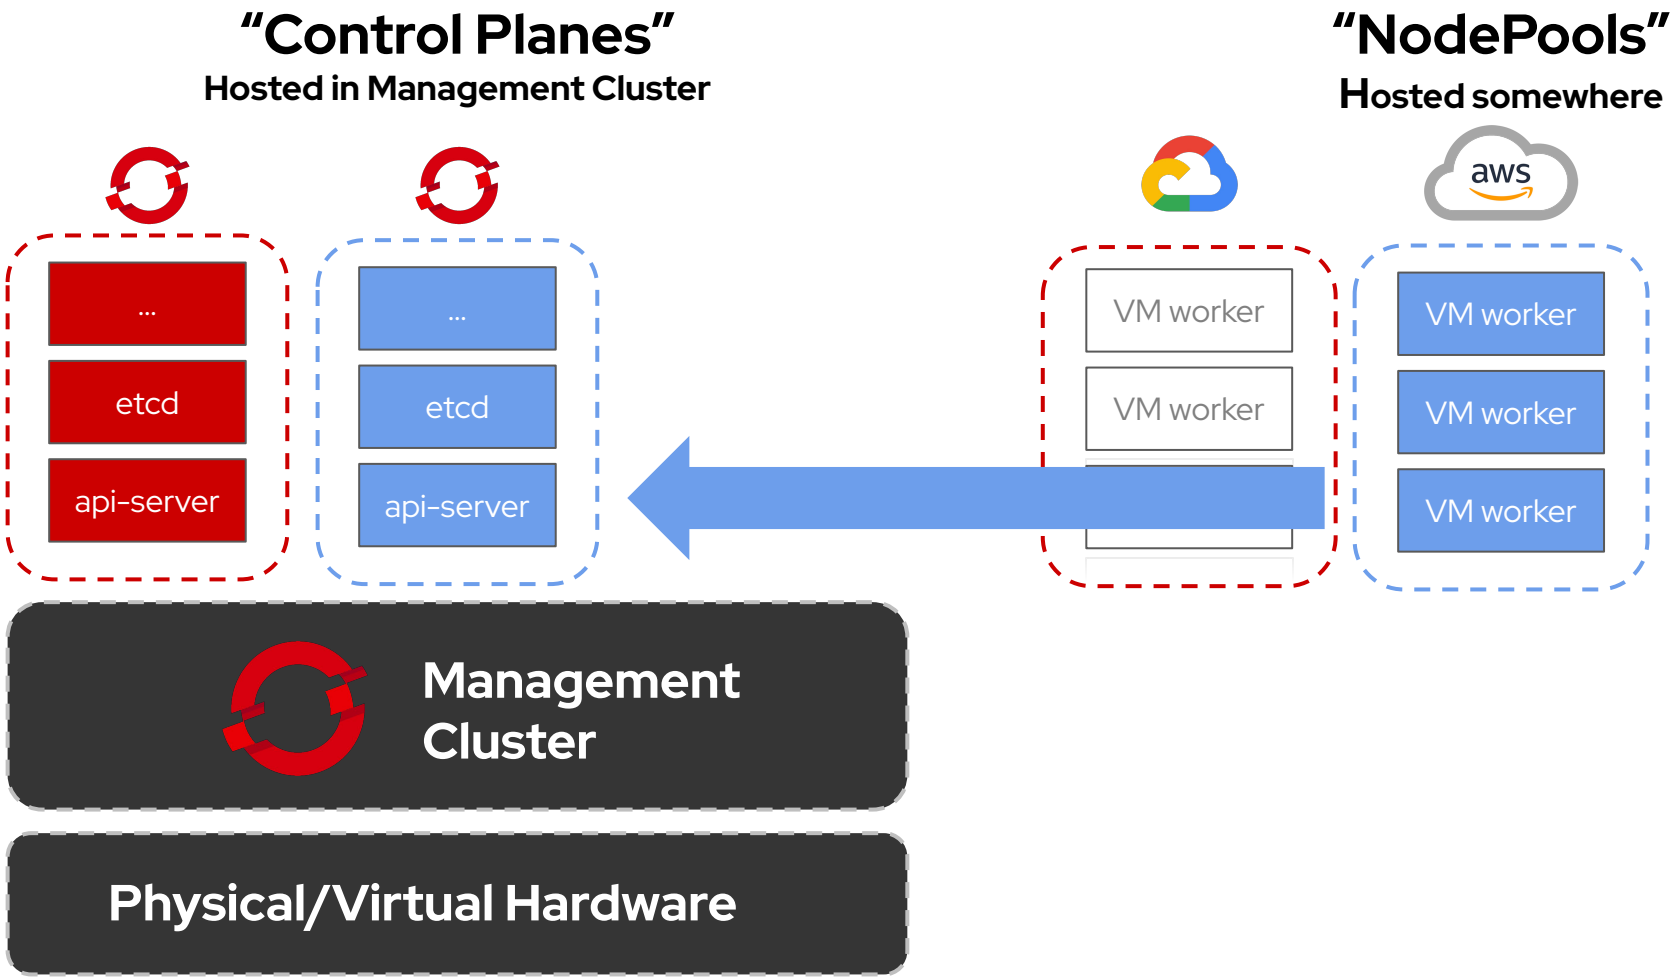




Red Hat

# HCP

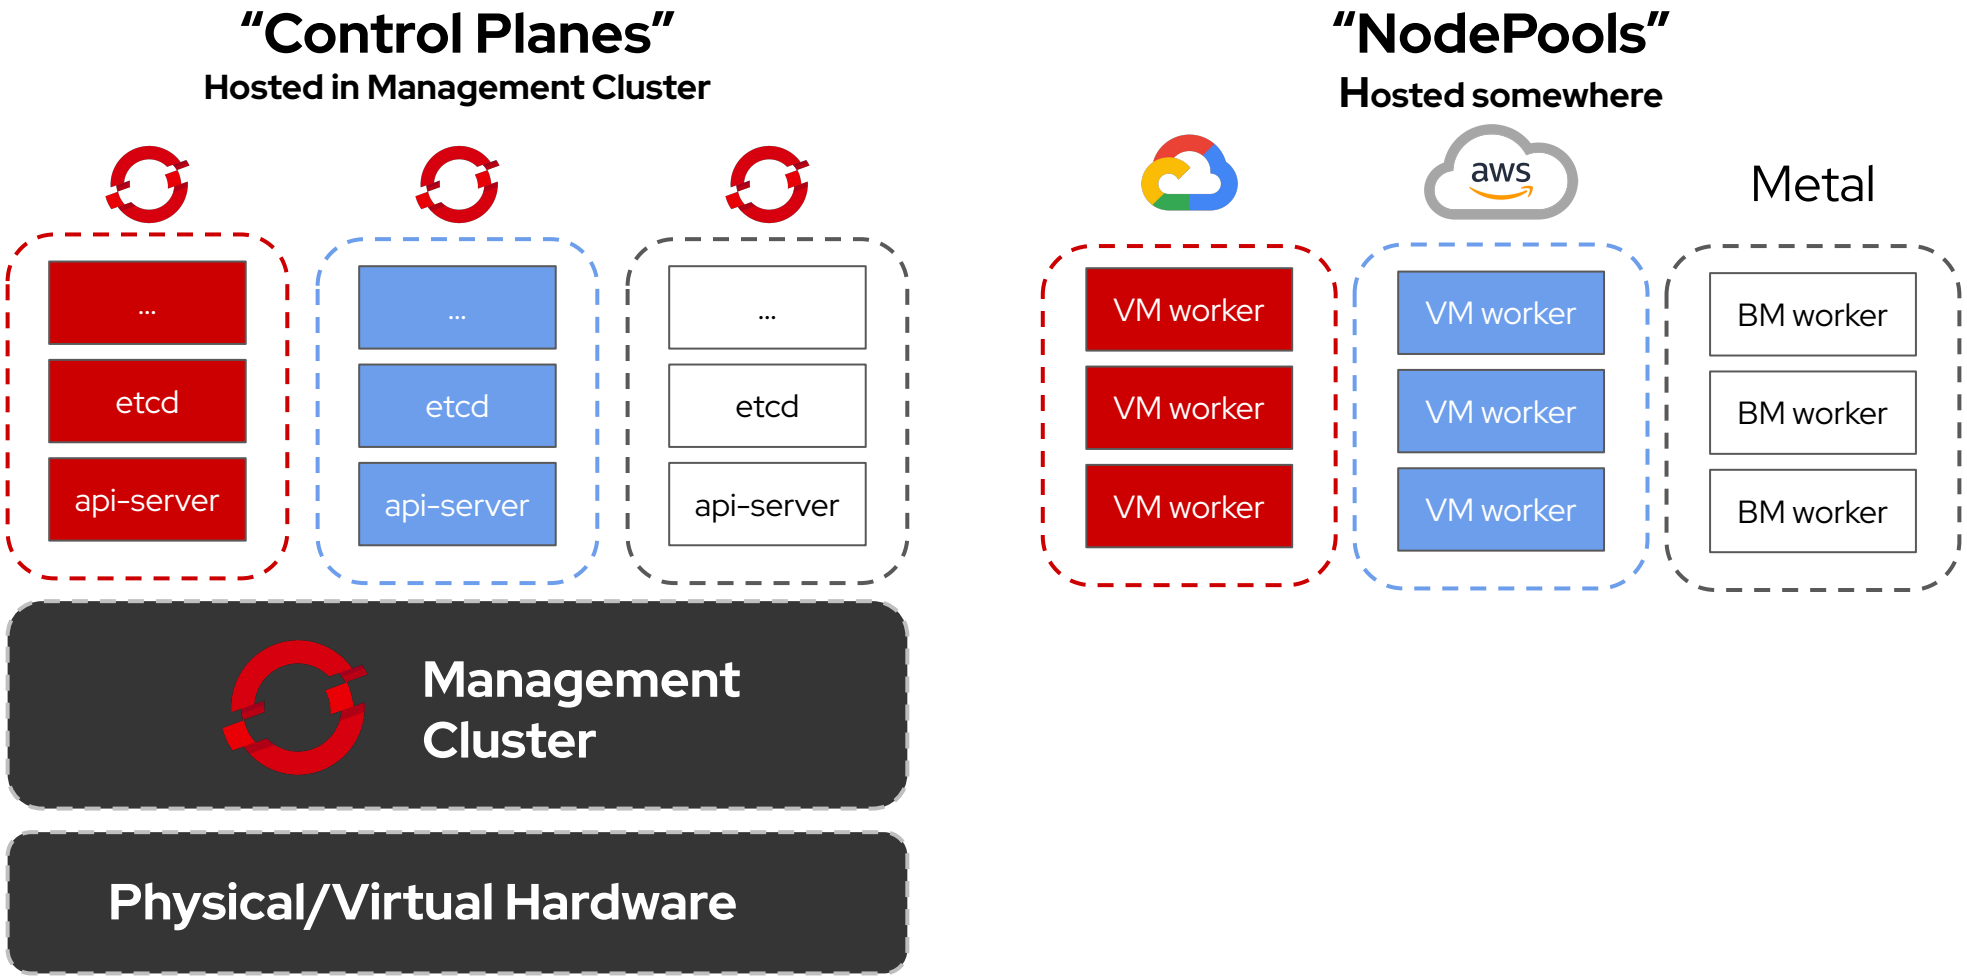
"Control Planes"
Hosted in Management Cluster
...
etcd
api-server
...
etcd
api-server
...
etcd
api-server
Management Cluster
Physical/Virtual Hardware
"NodePools"
Hosted somewhere
aws
Metal
VM worker
VM worker
VM worker
VM worker
VM worker
VM worker
BM worker
BM worker
BM worker

# HCP

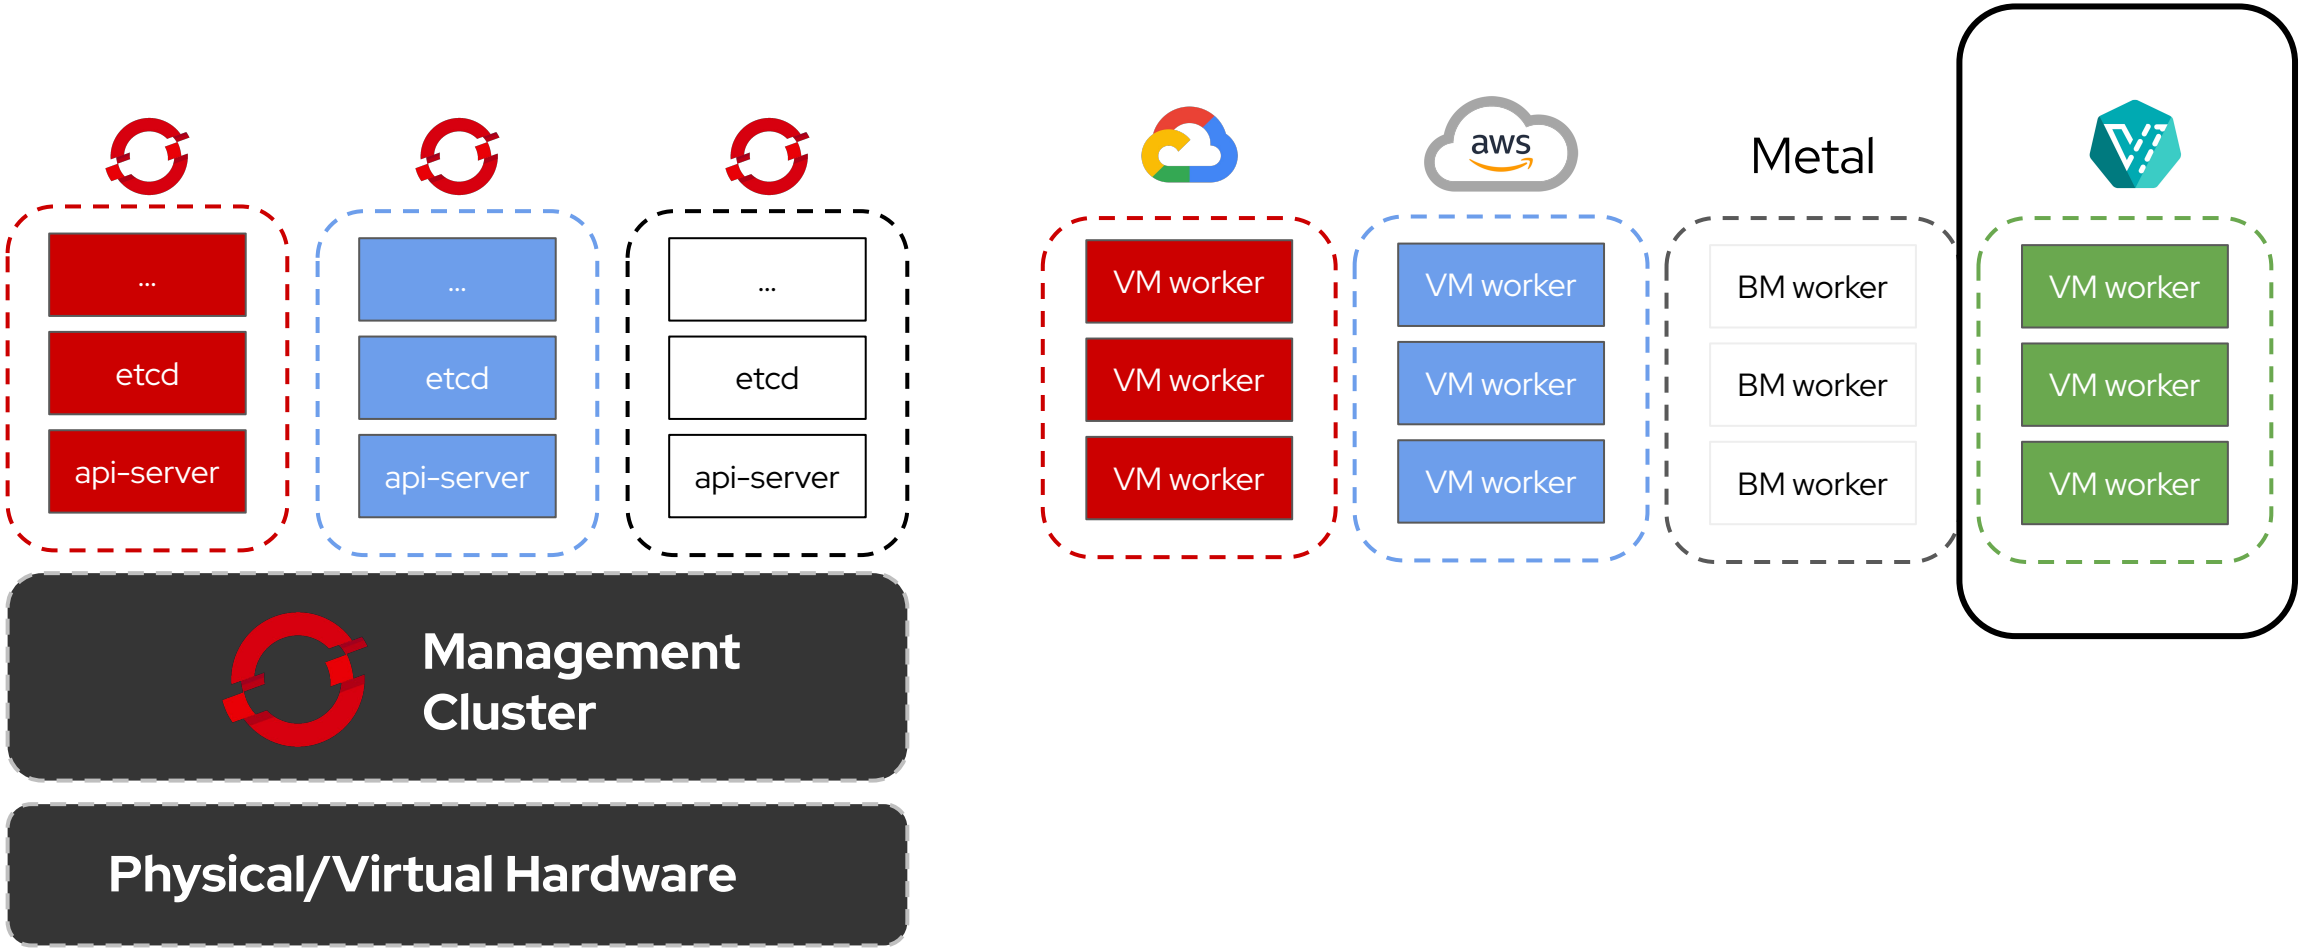
...
etcd
api-server
...
etcd
api-server
...
etcd
api-server
aws
Metal
VM worker
VM worker
VM worker
VM worker
VM worker
VM worker
BM worker
BM worker
BM worker
VM worker
VM worker
VM worker
Management Cluster
Physical/Virtual Hardware



Red Hat

# 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

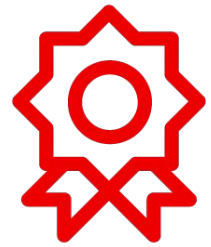

## Decoupled Life Cycle Management

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

# Networking

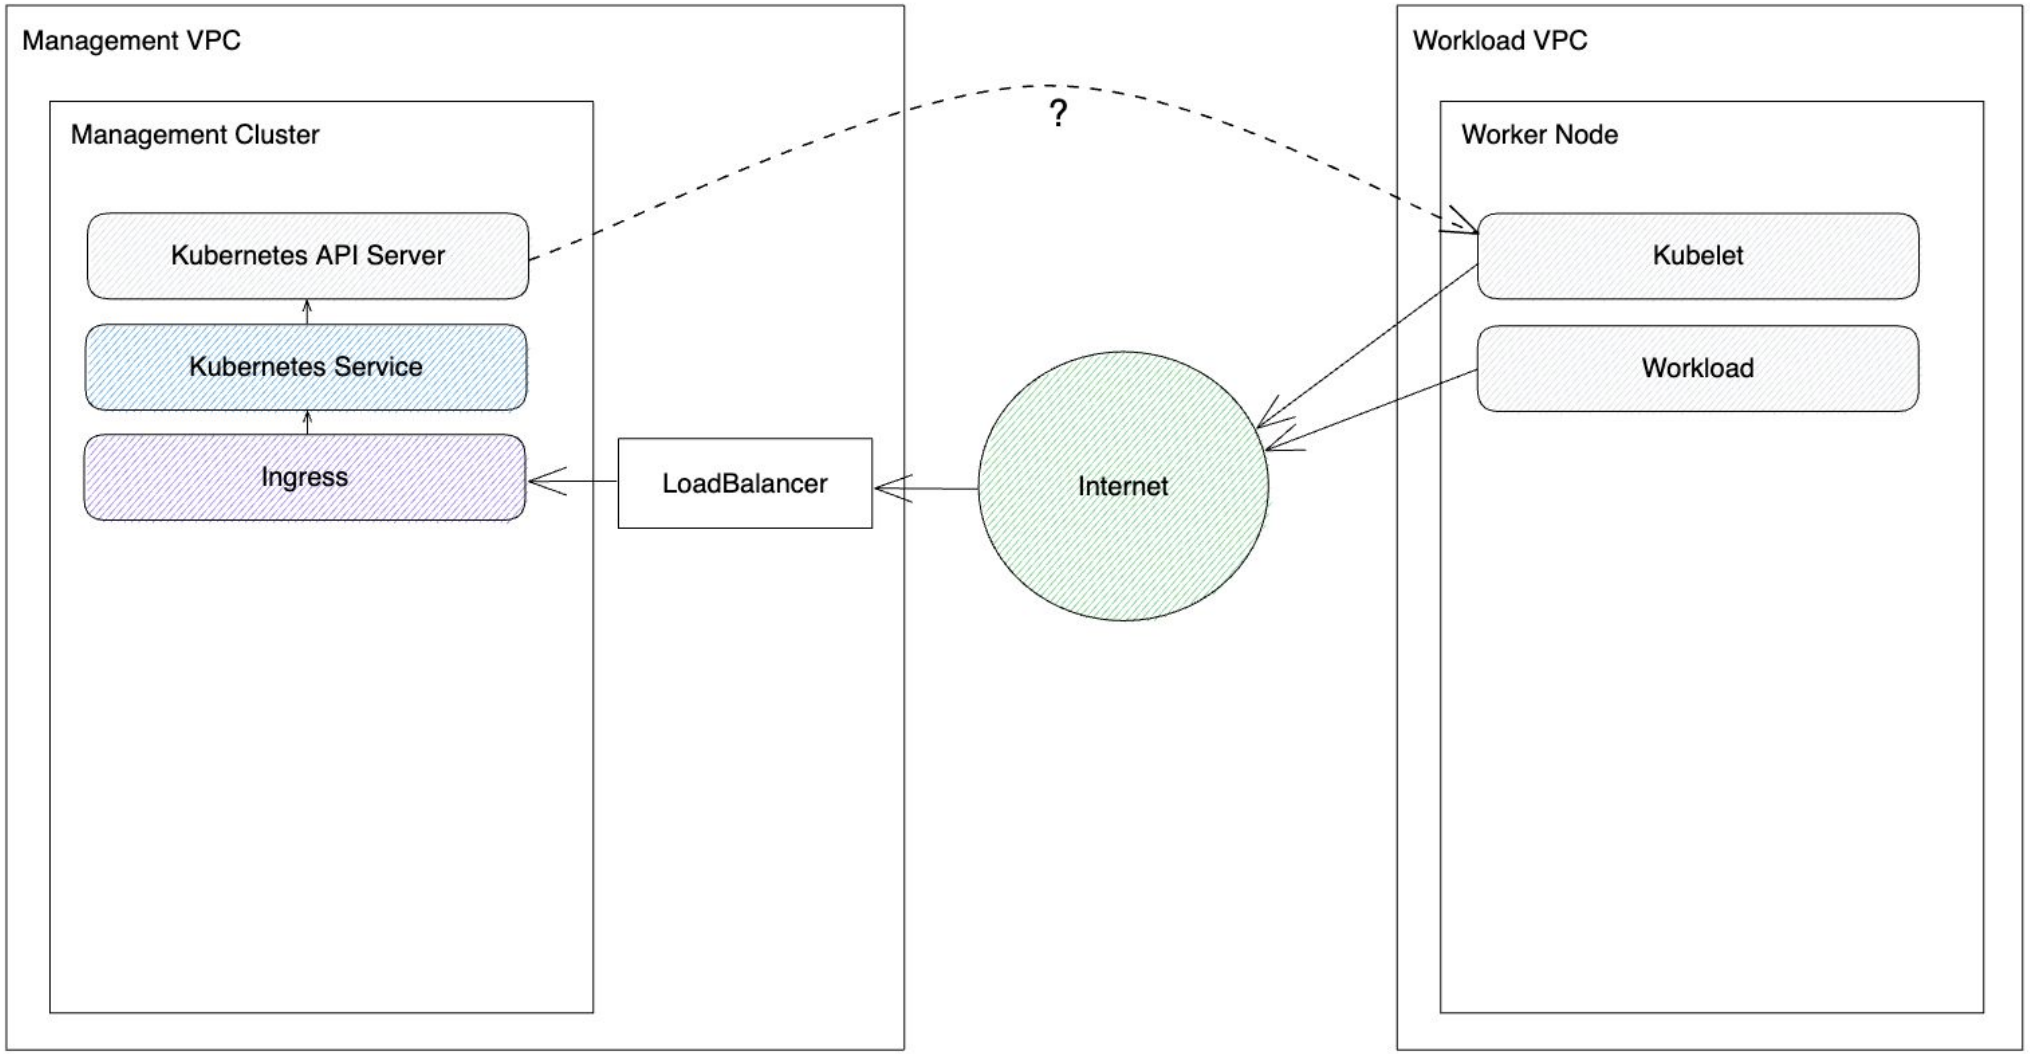

# Networking

## Networking between the management cluster and the hosted clusters

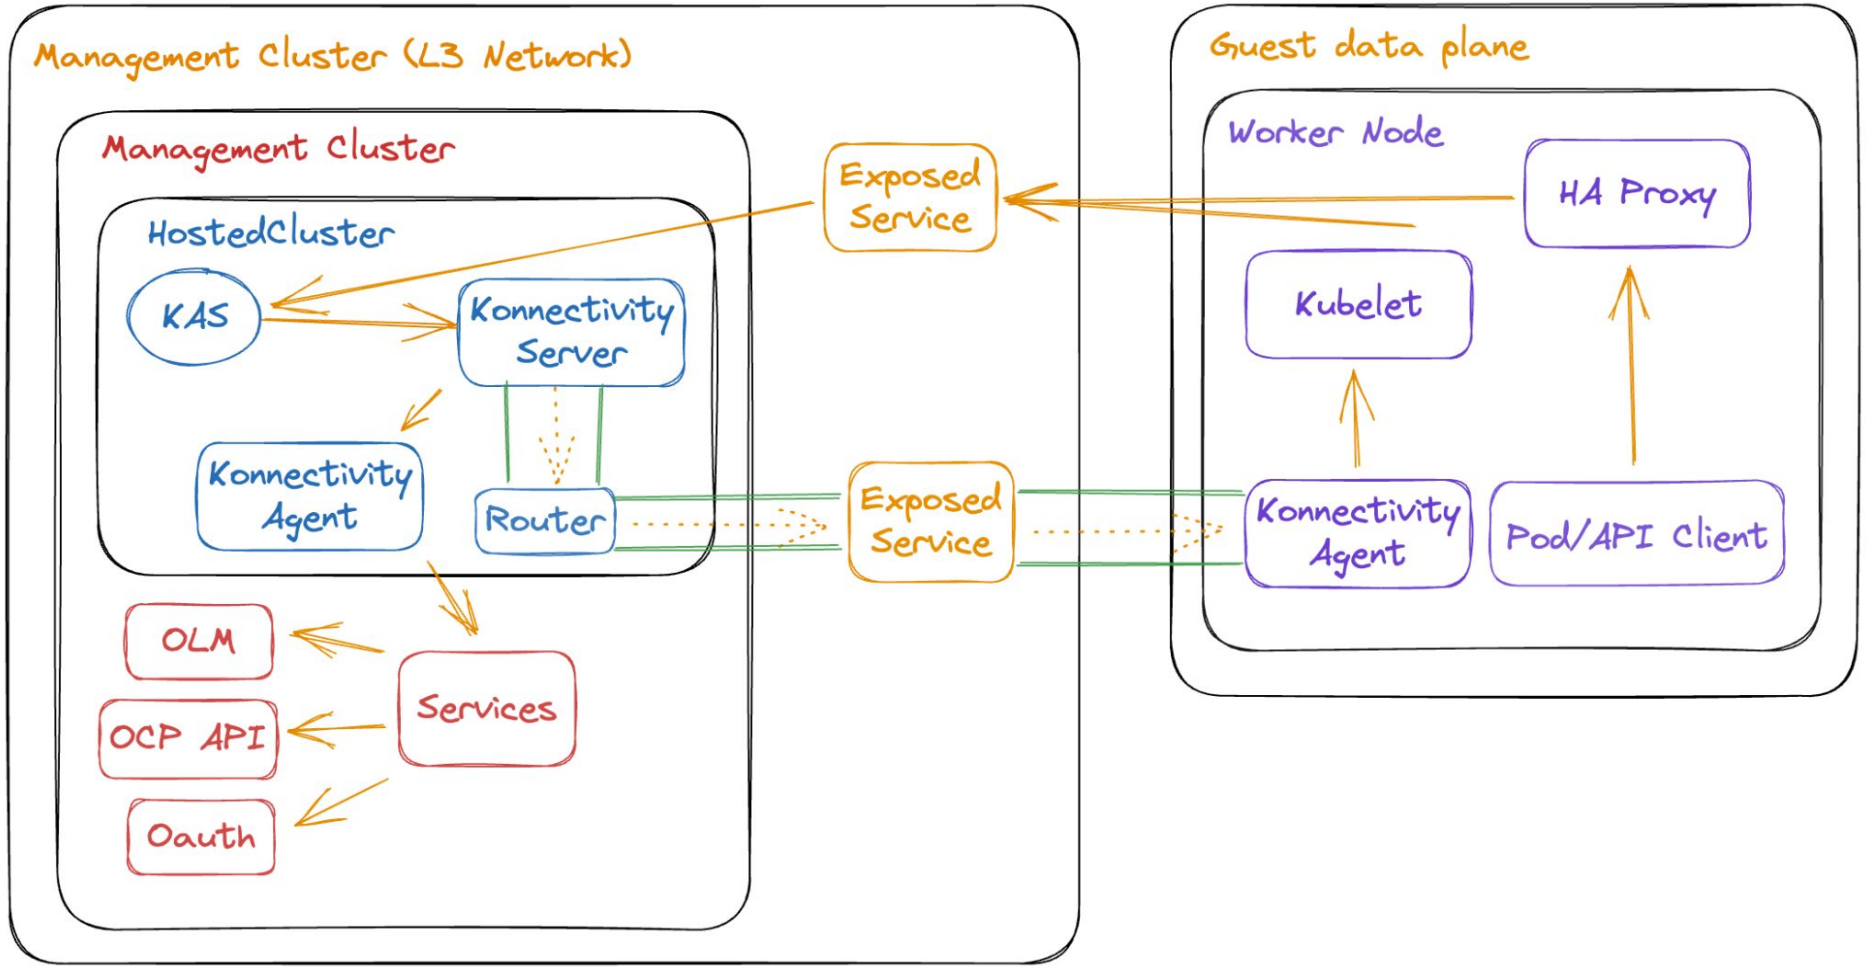




Red Hat

# High Availability

## High-level summary of different failure scenarios

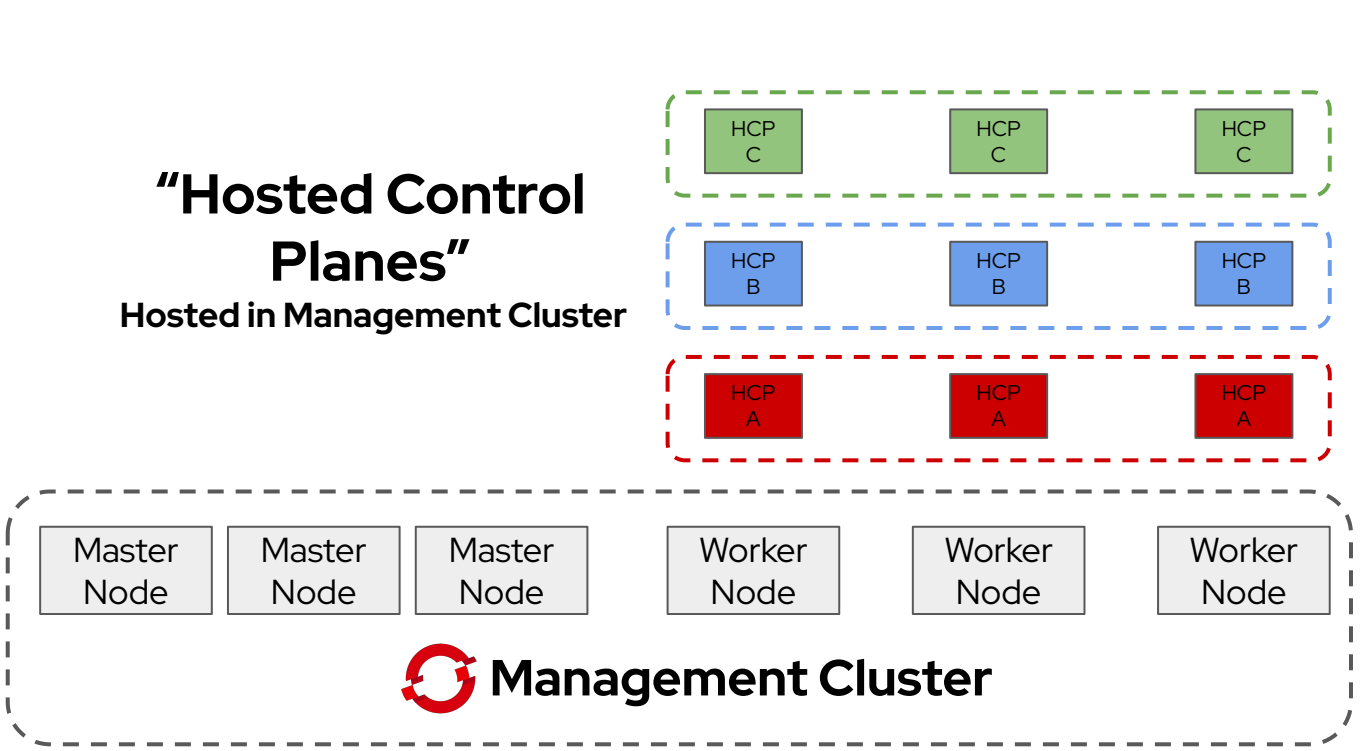


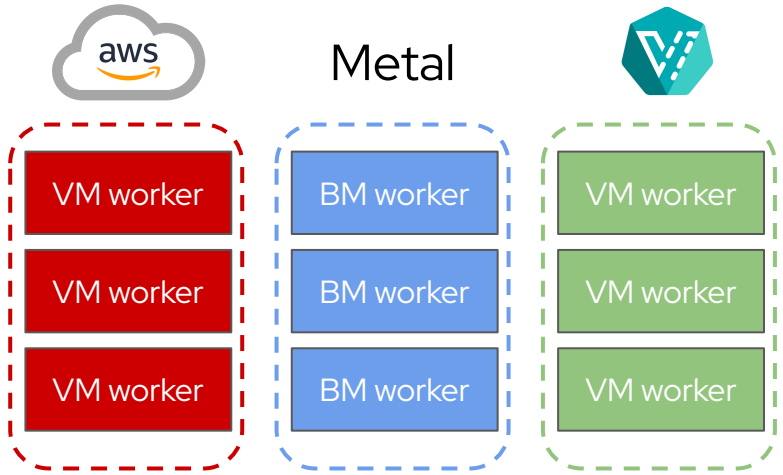




Red Hat

# High Availability

## Loss of management cluster worker

Hosted control plane API is still available.

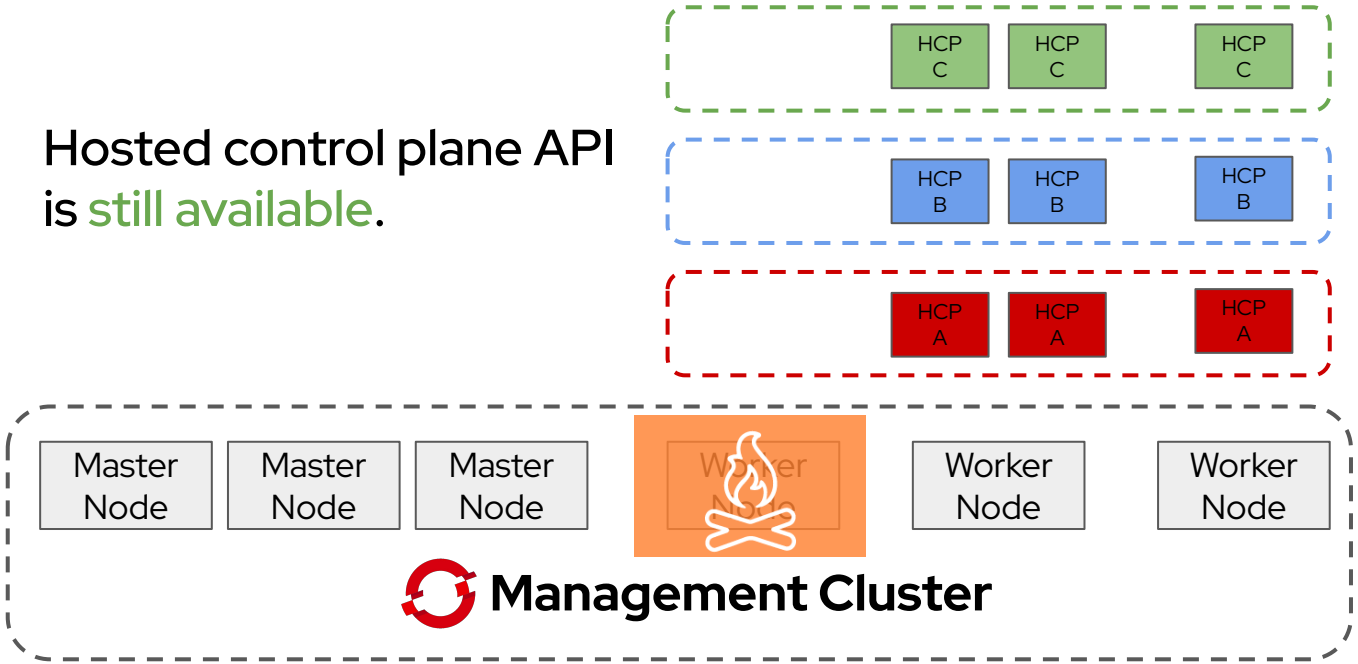


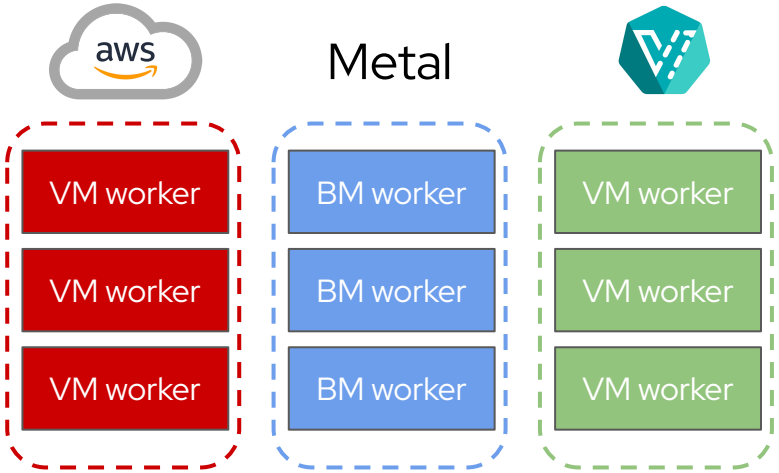


Hosted cluster data plane is still available.

Impacted hosted control plane components are rescheduled.



Red Hat

# High Availability

## Loss of management cluster availability zone

Hosted control plane API is still available.

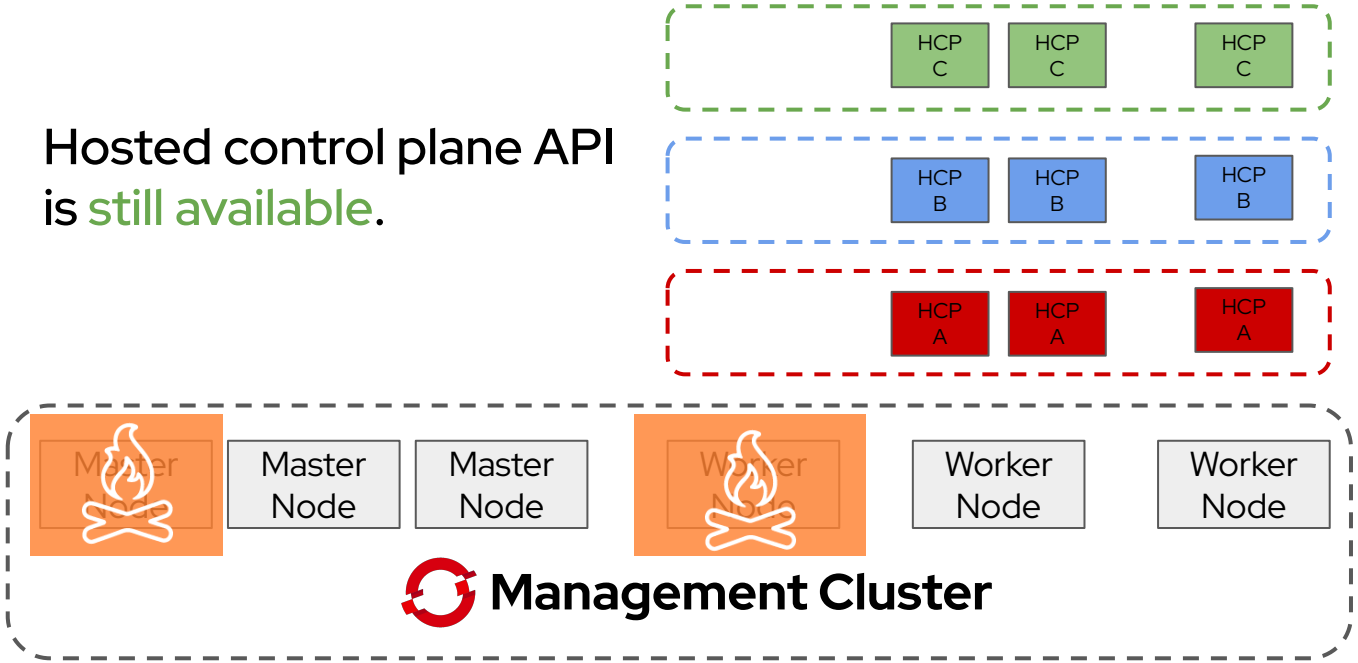


Impacted hosted control planes maintain quorum.

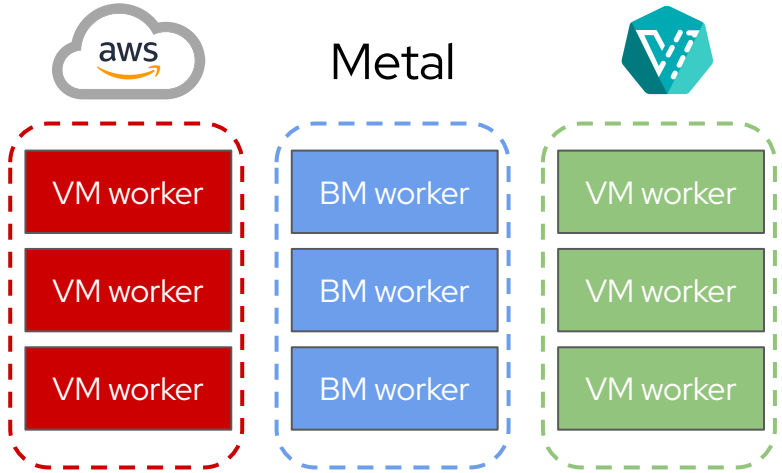


Hosted cluster data plane is still available.



Red Hat

# High Availability

## Loss of management cluster control plane

Hosted control plane API is still available.

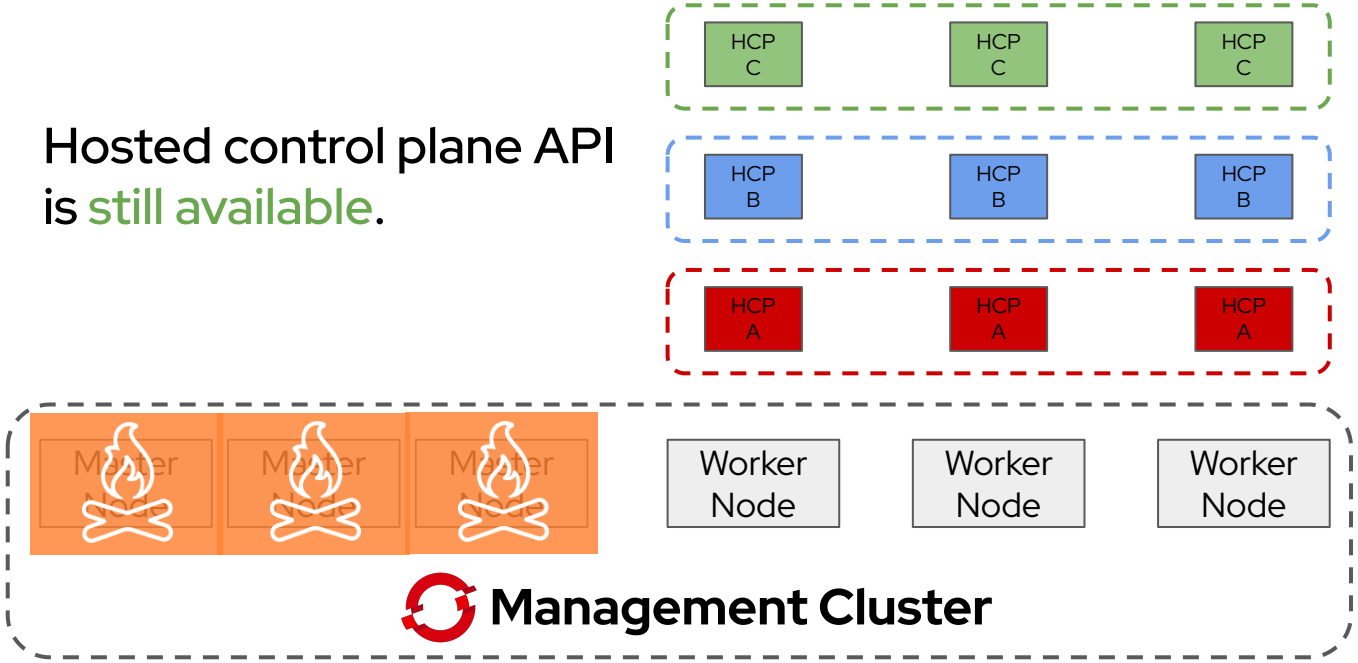


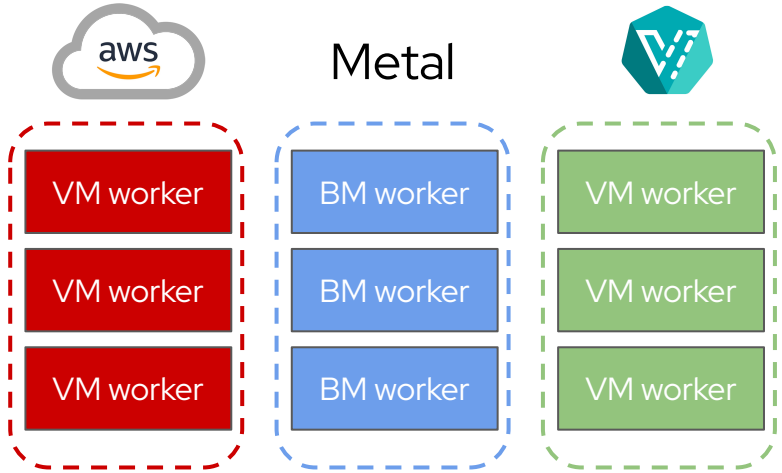


Hosted cluster data plane is still available.



Red Hat

# High Availability

## Loss of management cluster control plane and workers

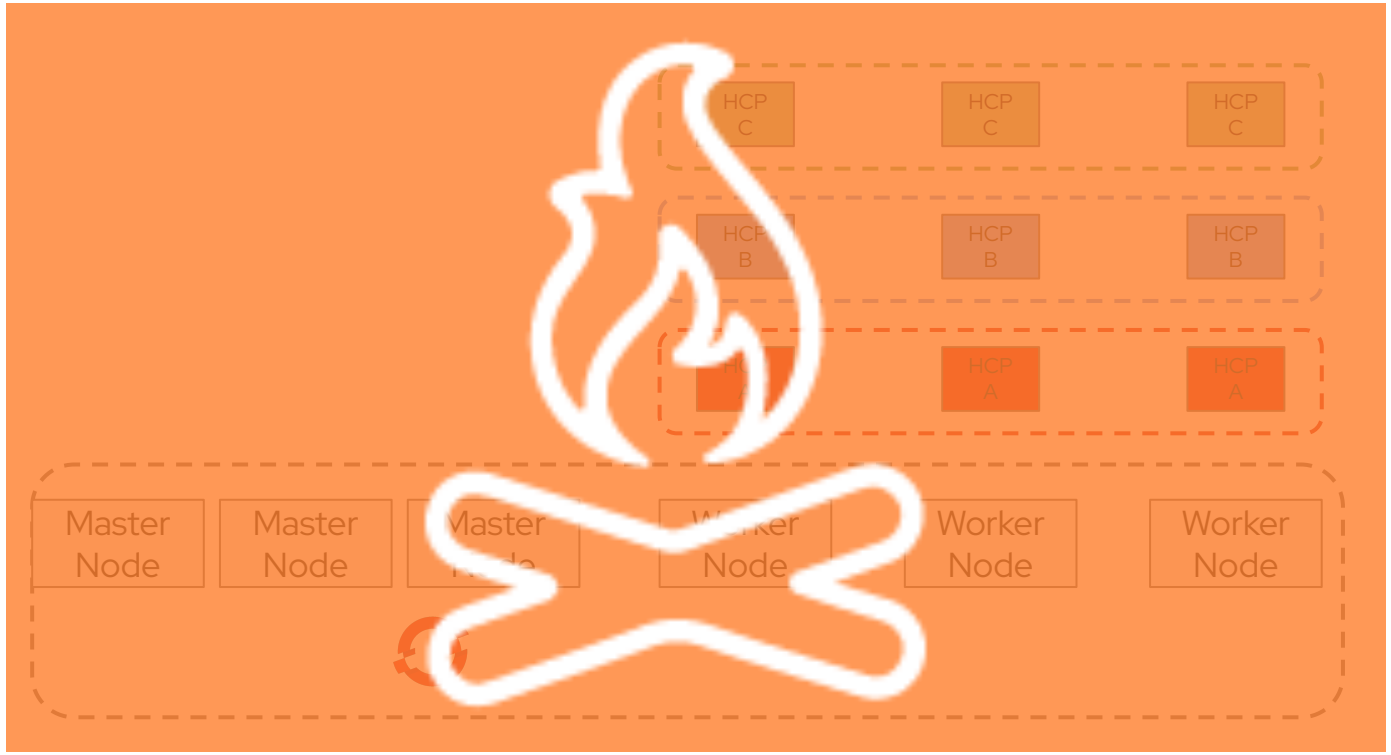


Hosted control plane API is not available.

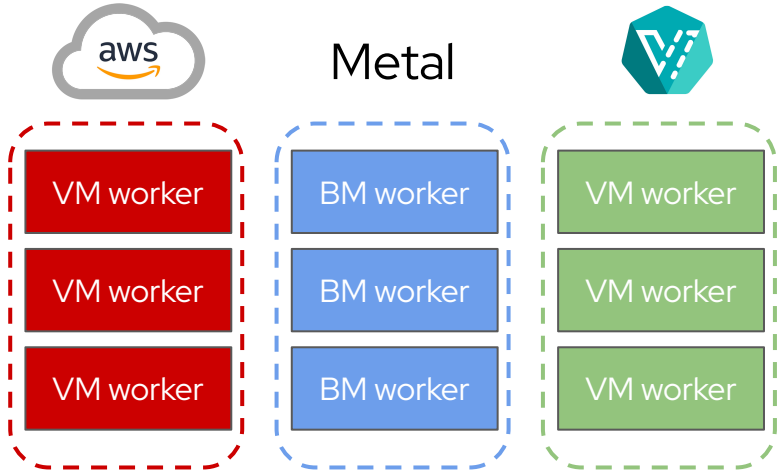


Hosted cluster data plane is still available.



Red Hat

# Upgrades are decoupled for CP and node pools

- Management cluster can host different versions of control planes
- Hosted cluster handles control plane updates, and node pools handle node upgrades
- Node pools support **replace** updates and **in-place** updates
- Supported versions config map

# OpenShift Virtualization

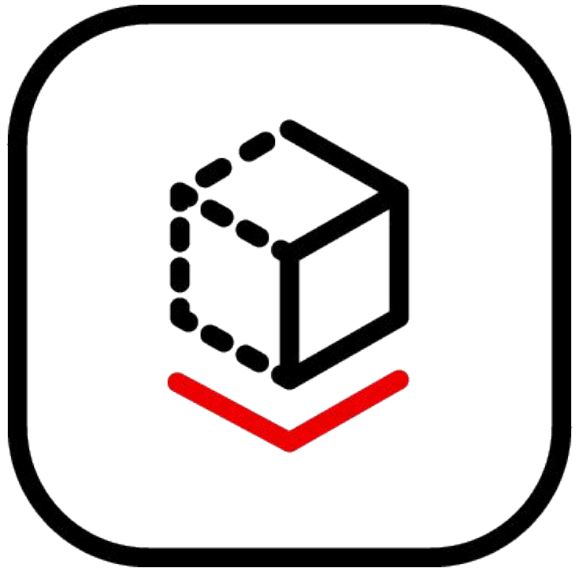

A virtualization API and runtime for OpenShift, built on KubeVirt, to run and manage virtual machines using a Kubernetes-native way

# HCP + Kubevirt

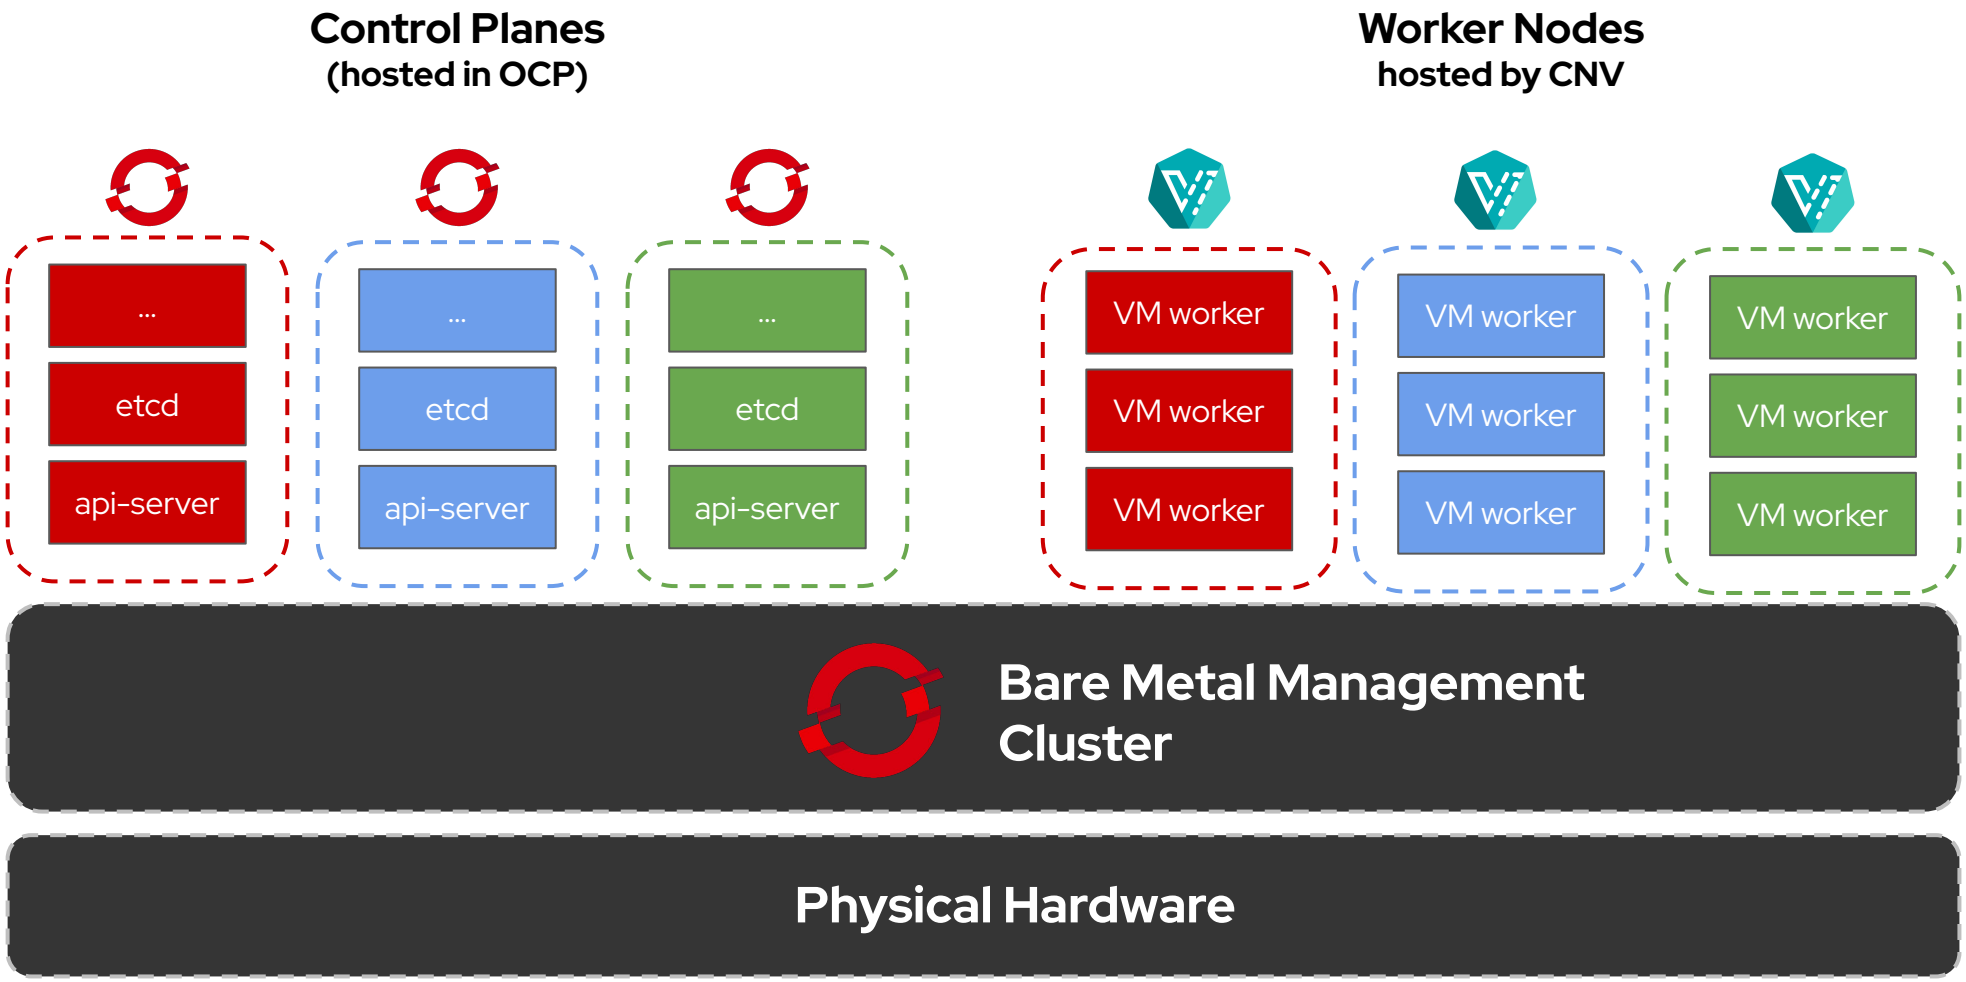




Red Hat

# 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

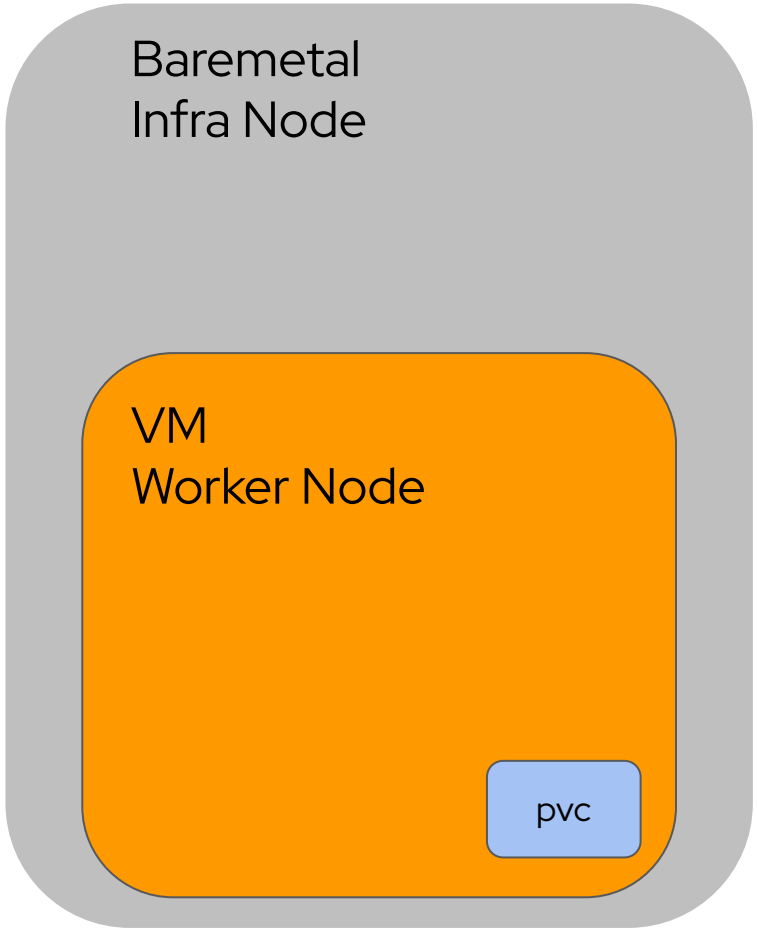


- Extends infra StorageClass into the guest clusters hosted by KubeVirt
- Utilizes HotPlug to make infra PVCs available within guest clusters
- Flow example...
  - User within guest cluster creates a PVC

# KubeVirt CSI Driver

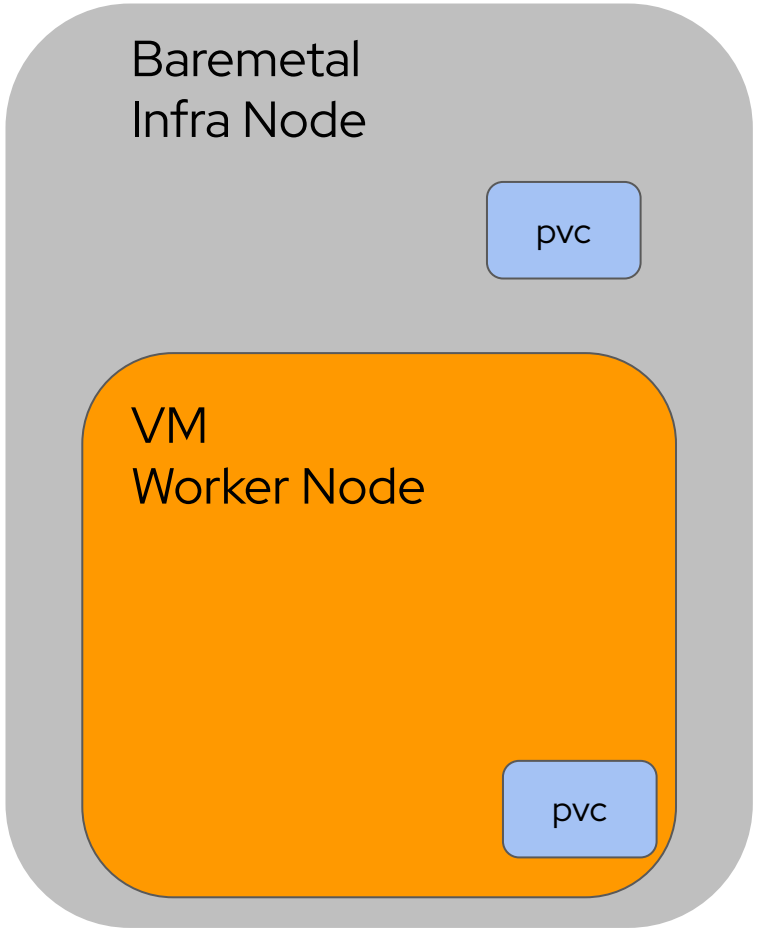


- Extends infra StorageClass into the guest clusters hosted by KubeVirt
- Utilizes HotPlug to make infra PVCs available within guest clusters
- Flow example...
  - User within guest cluster creates a PVC
  - KubeVirt CSI driver mirrors this PVC to the infra cluster

# KubeVirt CSI Driver

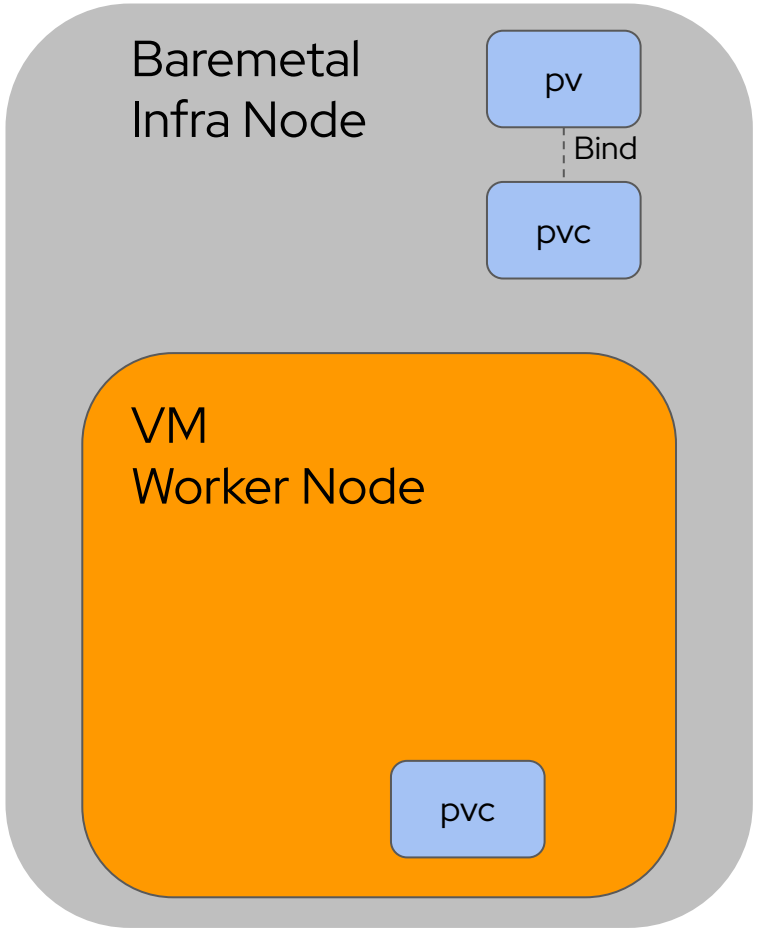


- Extends infra StorageClass into the guest clusters hosted by KubeVirt
- Utilizes HotPlug to make infra PVCs available within guest clusters
- Flow example...
  - User within guest cluster creates a PVC
  - KubeVirt CSI driver mirrors this PVC to the infra cluster
  - Infra cluster's dynamic storage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

# KubeVirt CSI Driver

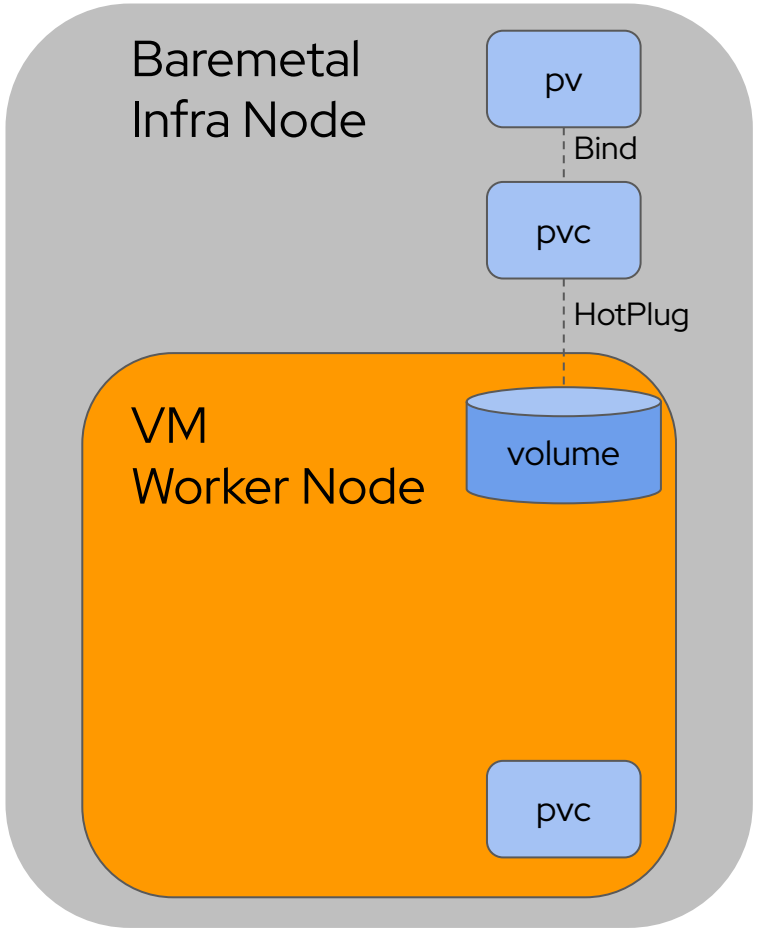


- Extends infra StorageClass into the guest clusters hosted by KubeVirt
- Utilizes HotPlug to make infra PVCs available within guest clusters
- Flow example…
  - User within guest cluster creates a PVC
  - KubeVirt CSI driver mirrors this PVC to the infra cluster
  - Infra cluster's dynamic storage provisioner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

# KubeVirt CSI Driver

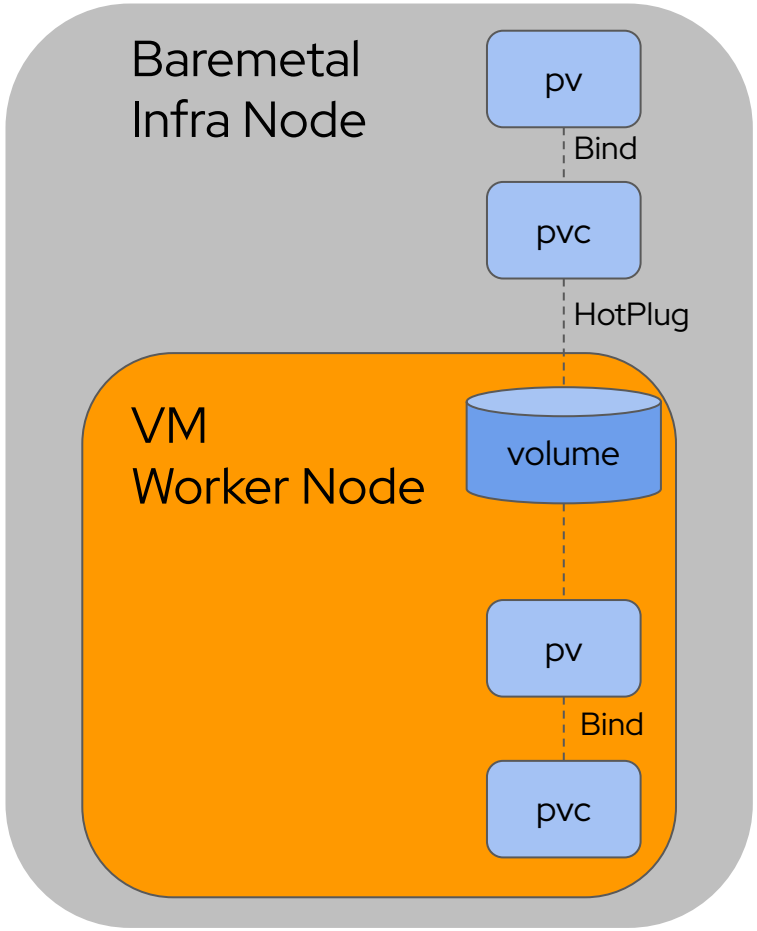


- Extends infra StorageClass into the guest clusters hosted by KubeVirt
- Utilizes HotPlug to make infra PVCs available within guest clusters
- Flow example…
  - User within guest cluster creates a PVC
  - KubeVirt CSI driver mirrors this PVC to the infra cluster
  - Infra cluster's dynamic storage provisioner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

# 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

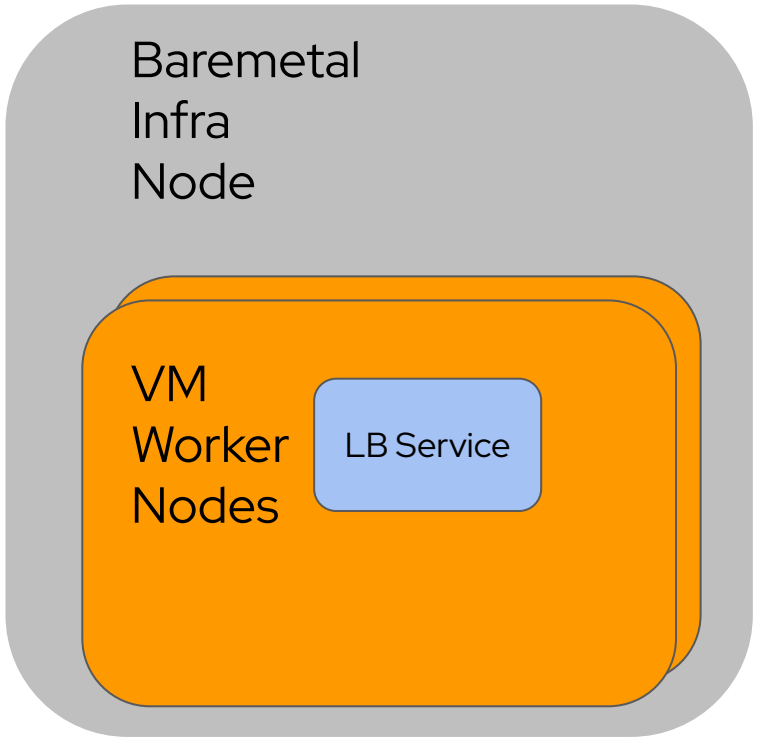


- Provides Load Balancer support to KubeVirt guest clusters
- Similar to KubeVirt CSI in that it is mirroring infra capabilities to guest clusters.
- Flow example...
  - User within guest cluster creates a LoadBalancer service

# Cloud Provider KubeVirt

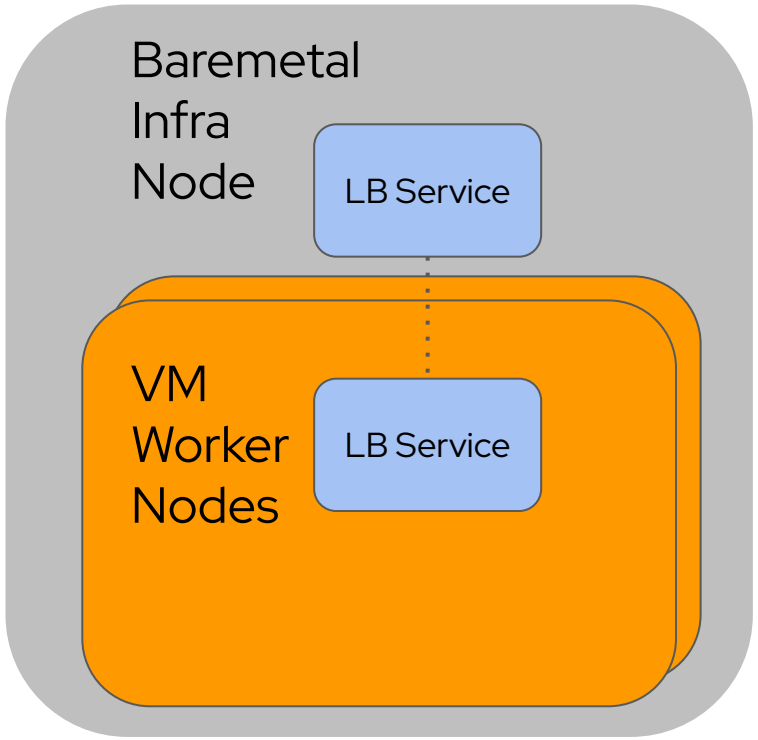


- Provides Load Balancer support to KubeVirt guest clusters
- Similar to KubeVirt CSI in that it is mirroring infra capabilities to guest clusters.
- Flow example...
  - User within guest cluster creates a LoadBalancer service
  - Cloud Provider Kubevirt controller creates corresponding LB on infra cluster

# Cloud Provider KubeVirt

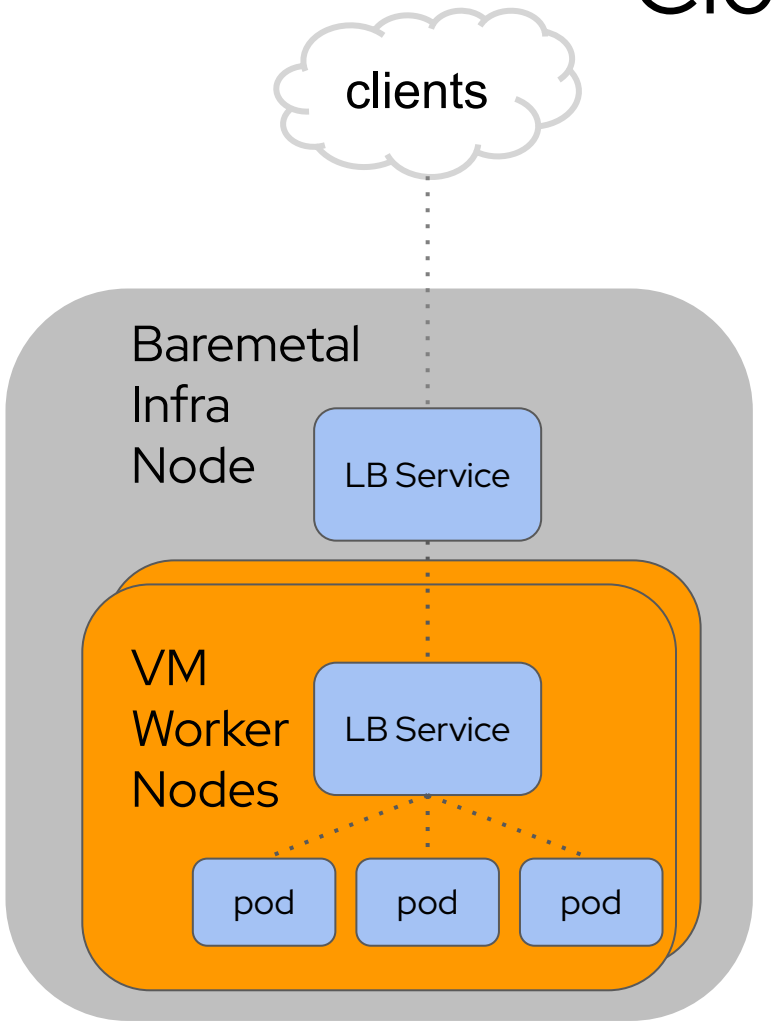


- Provides Load Balancer support to KubeVirt guest clusters
- Similar to KubeVirt CSI in that it is mirroring infra capabilities to guest clusters.
- Flow example...
  - User within guest cluster creates a LoadBalancer service
  - Cloud Provider Kubevirt controller creates corresponding LB on infra cluster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

# Hosted KubeVirt Clusters

- Pod network attachment can be disabled
- Additional networks can be specified for NodePools
- Boot image for worker nodes can be cached for faster boots
- PCI devices such as GPUs in management cluster can be utilized in worker nodes
- Management and guest clusters can have different DNS domains
- External Infrastructure can be used for guest clusters

# Advanced Cluster Management for Kubernetes

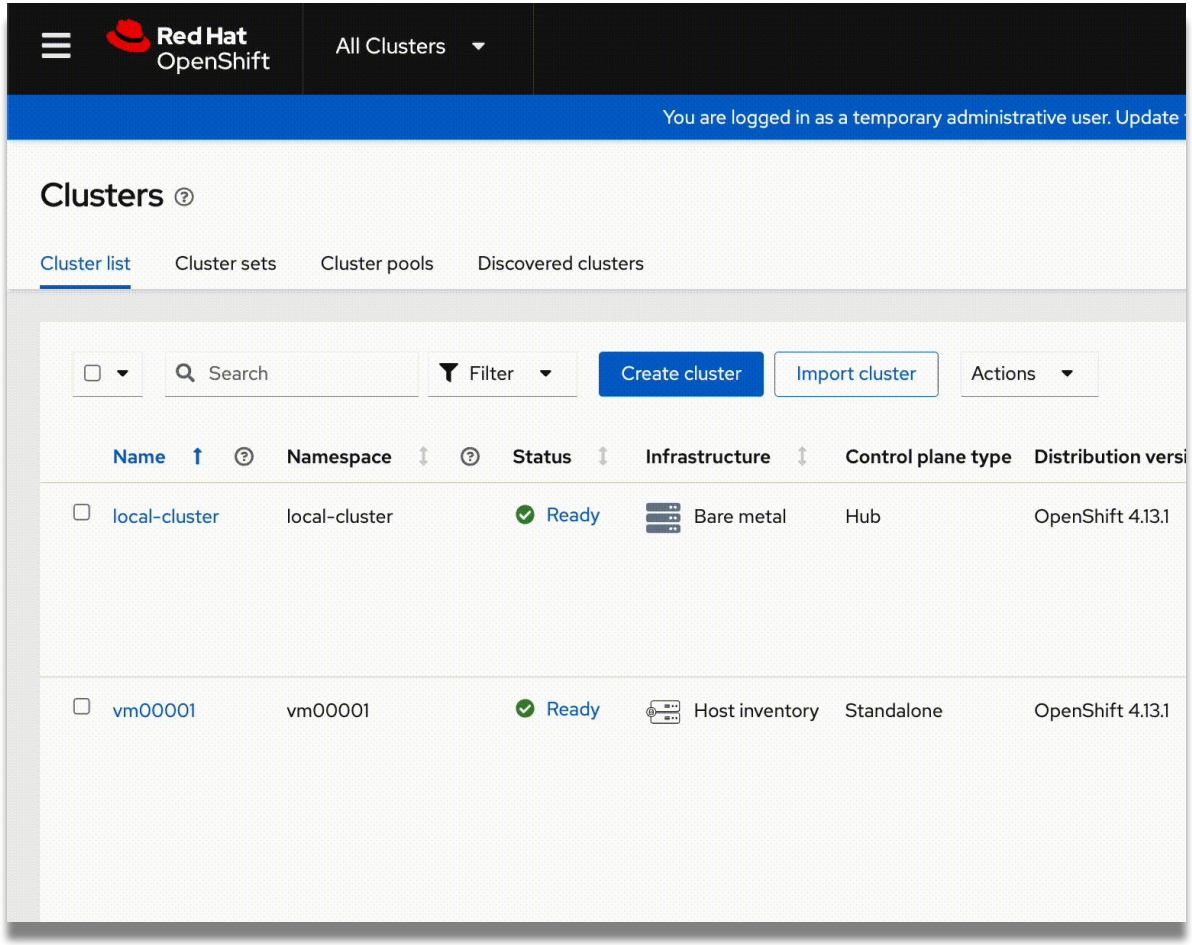


- **Hosted Control Planes Available on ACM**
  Hosted Control Planes is available on OpenShift Virtualization and Bare Metal
- ACM centrally manages standalone and hosted clusters
  - Create, Delete, Connect, Monitor Clusters
- Policy Governance
- Application Lifecycle



Red Hat

Red Hat Summit

Connect

# Thank you


linkedin.com/company/red-hat


facebook.com/redhatinc


youtube.com/user/RedHatVideos


twitter.com/RedHat

Red Hat