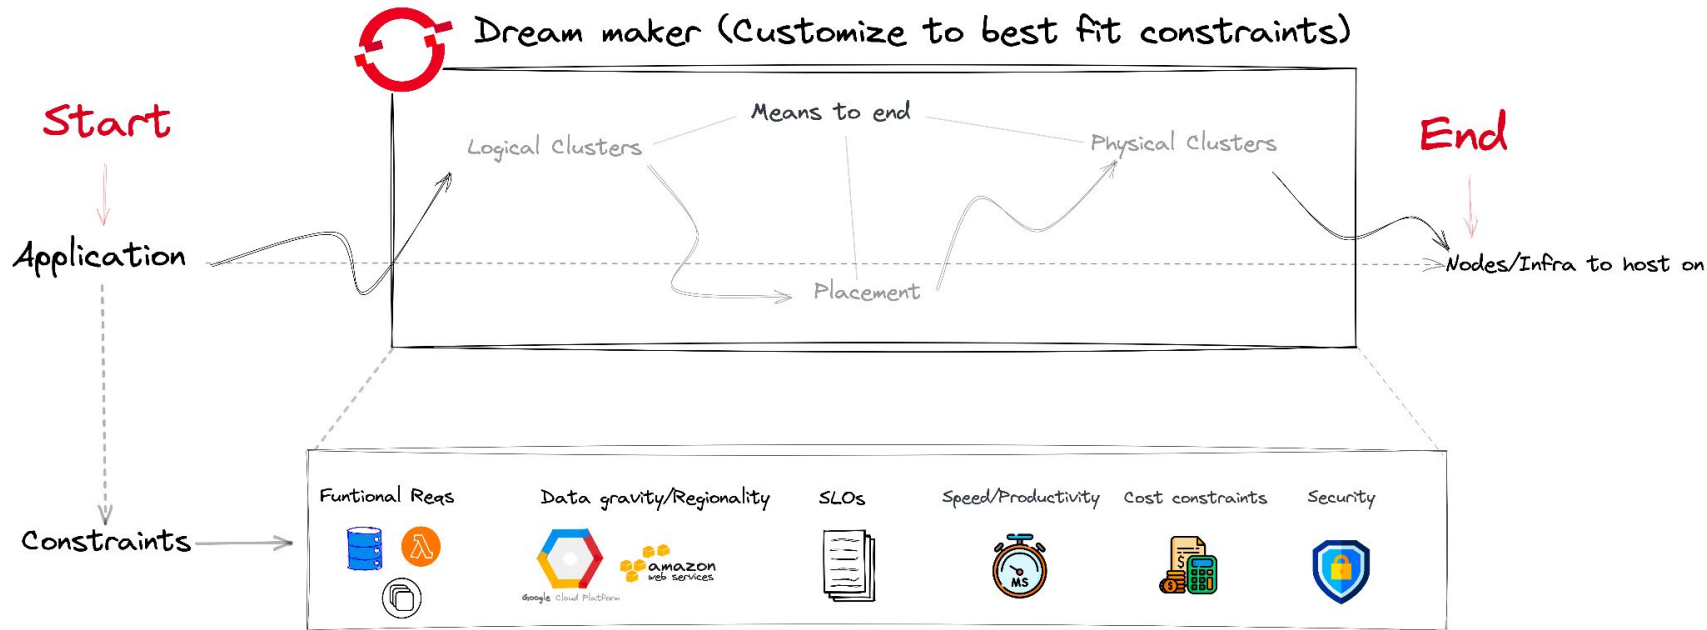


Hosted Control Planes

Patrick Ladd <pladd@redhat.com>
NYRHUG February 2024



The Big Picture



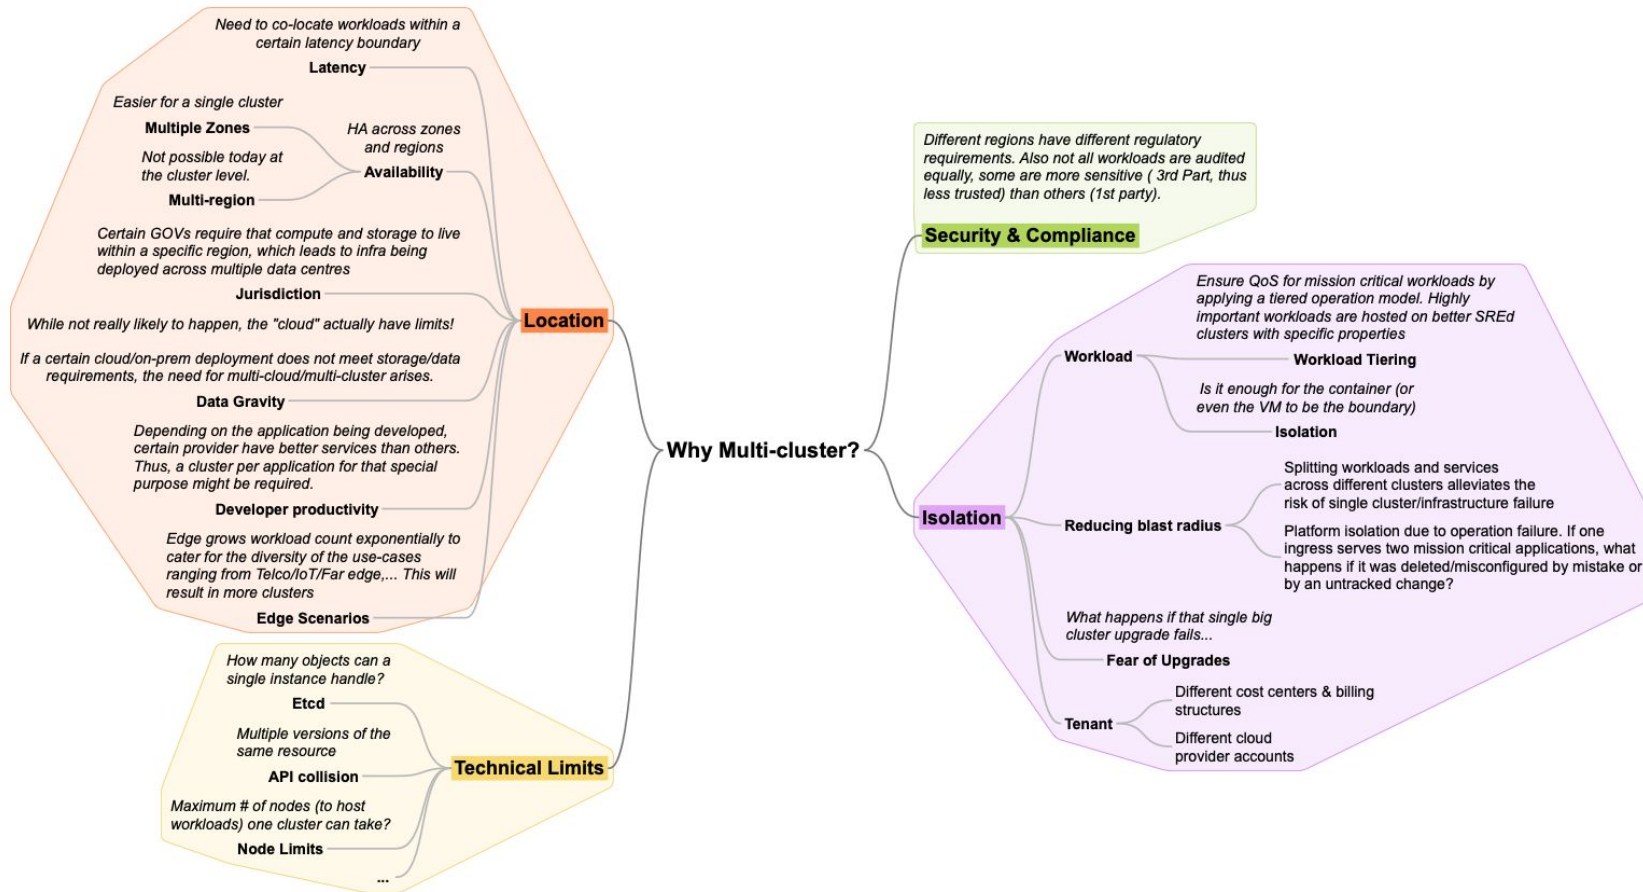
Users expect applications and services to be available 24/7

Companies desire more efficient use of cloud resources

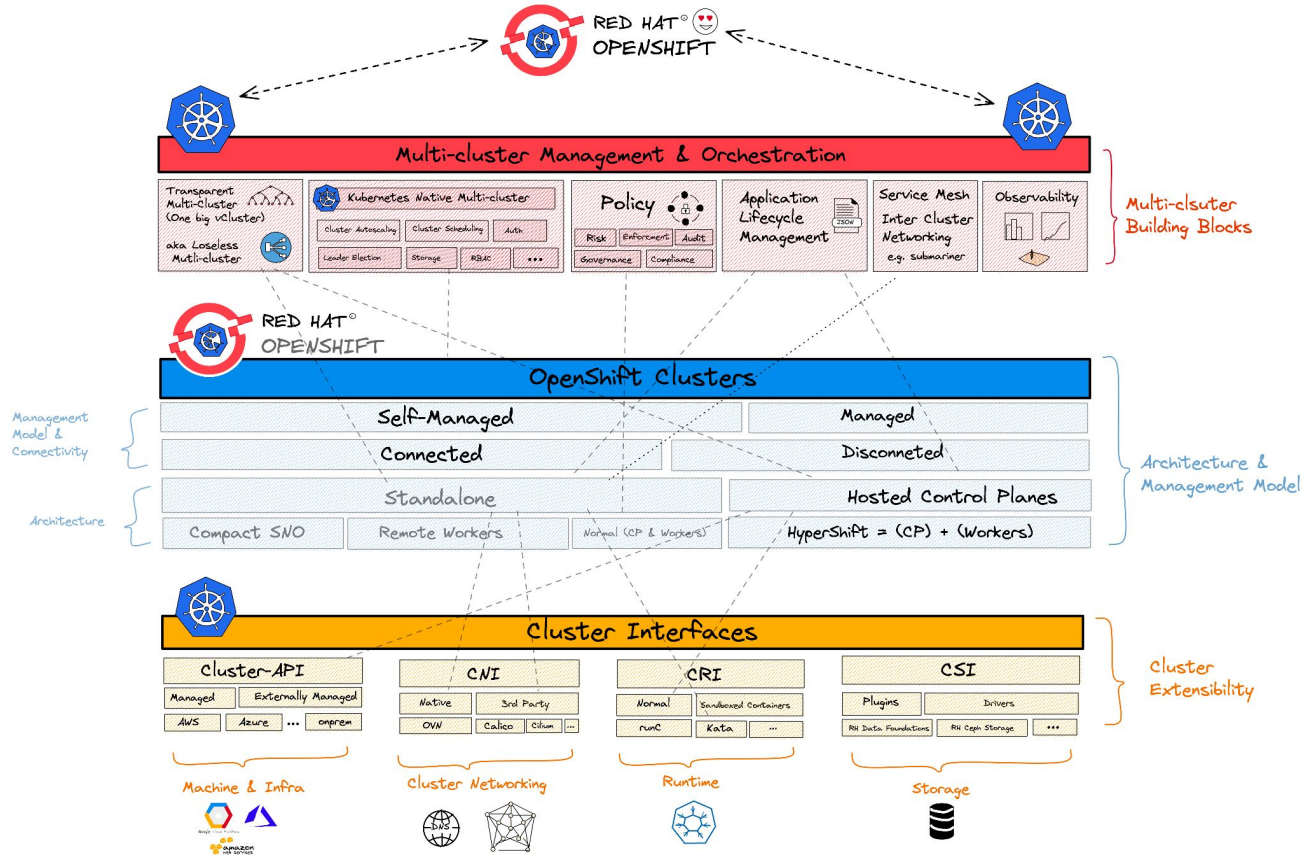
Developers expect to deploy code multiple times a day with no downtime

Companies Must comply with Security standards / Regulations

The Big Picture - Why Multi-cluster




The Big Picture - Dream Maker (aka OpenShift) Tech Stack



Standalone OpenShift



Management
Model &
Connectivity



Architecture &
Management Model

Personas

Cluster Admin



oc/GitOps/ ...
Managed deployments/
RBAC/Policies

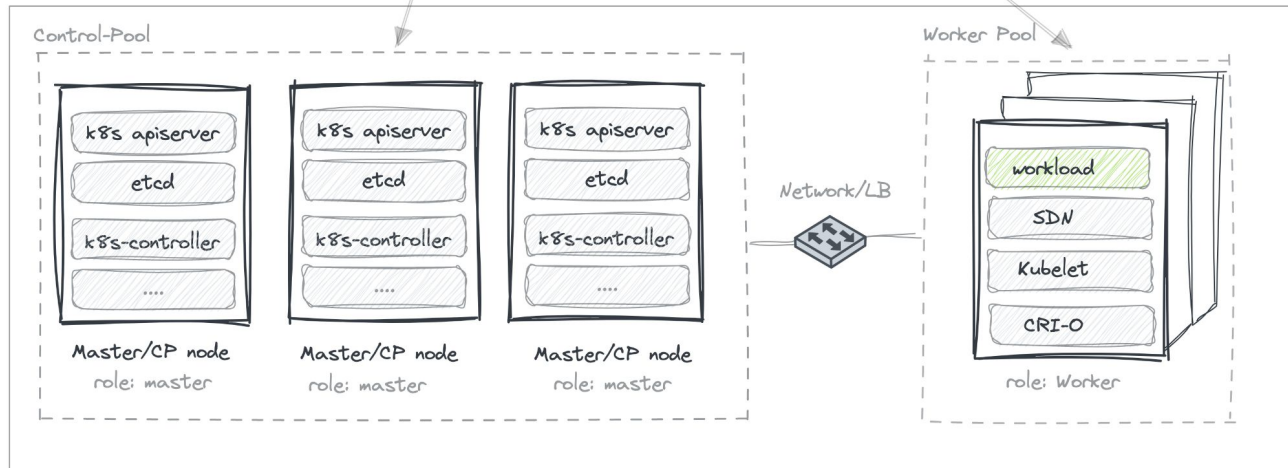
Developer




Run workloads
(IDE/Odo/ ...)



OpenShift cluster



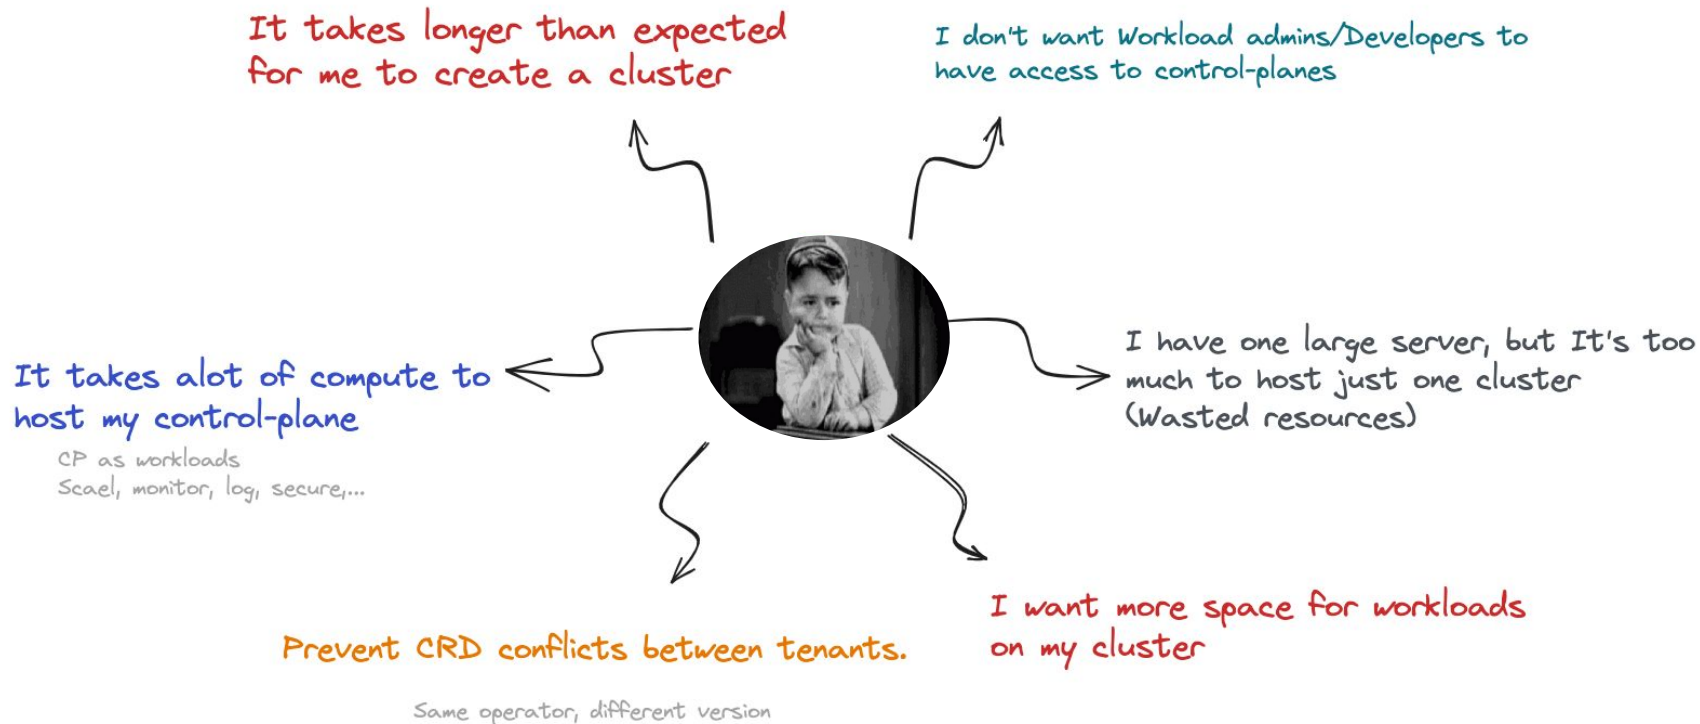
Where is the majority of your Kubernetes spend?



Source
CNCF
© Statista 2023


Additional Information:
Worldwide; April and May 2021; 178 respondents; Cloud native community*

Short Stories / Use-cases



Short Stories / Use-cases


Percentage (%) of Customers per Theme/Need for HCP



Hosted Control Planes (HyperShift)

RED HAT [®] OPENSHPIFT			
OpenShift Clusters			
Self-Managed		Managed	
Connected		Disconnected	
Standalone		Hosted Control Planes	
Compact S/O	Remote Workers	Normal CP + Workers	HyperShift = (CP) + (Workers)

- An **OpenShift** Topology
- Service for **hosting OpenShift control planes** at **scale**
- Solves for **cost** and **time to provision**
- Portable **across clouds**
- Provides **strong separation of concerns** between management and workloads.



Why HCP?

HCP



Supported OpenShift topology



Reduced infrastructure costs / densification



Faster cluster creation




Strong separation between control and workload




Support multi-arch / multi-env



Centralized management in a "Managed" model



Cluster Sizes Trending Down, Cluster Count UP!





Costs savings with HCP Relative to Cluster Size

Large cluster

~~3 Master nodes~~

100 Worker nodes

= Saving : 2.3%

Medium cluster

~~3 Master nodes~~

10 Worker nodes

= Saving : 23%

Small cluster



~~3 Master nodes~~

3 Worker nodes

= Saving : 50%

HCP Architecture & Support

Management cluster



Node Pool (s)



Node Pool (s)



Node Pool (s)

Service Provider



Provide infra for hosting
SRE ... SLOs + uptime

Service Consumer



Can request cluster control planes
can request workers

Roles/Personas

Instance Admin



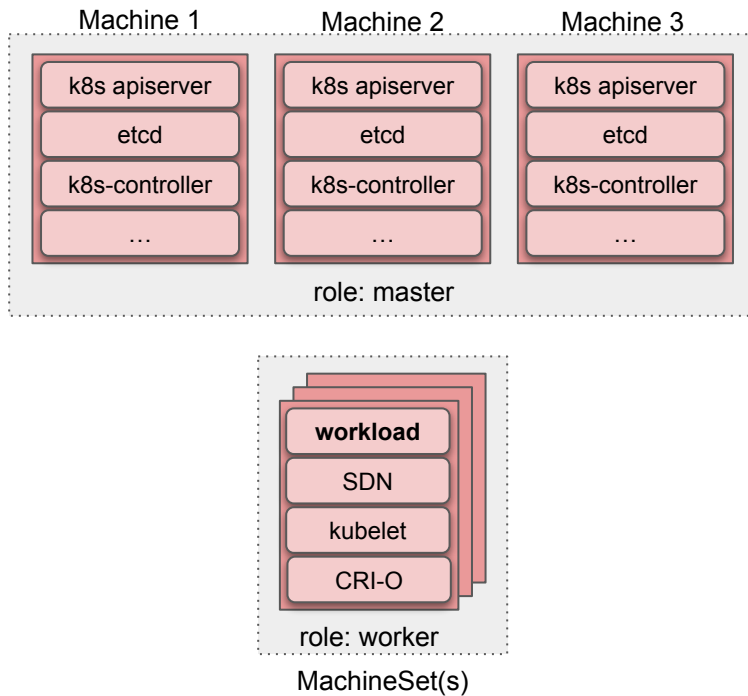
Cluster-admin, Managed deployments/
RBAC/Policies

Instance Developer





Run workloads
(IDE/Odo/...)



Standalone Clusters



Standalone Clusters



Standalone Clusters

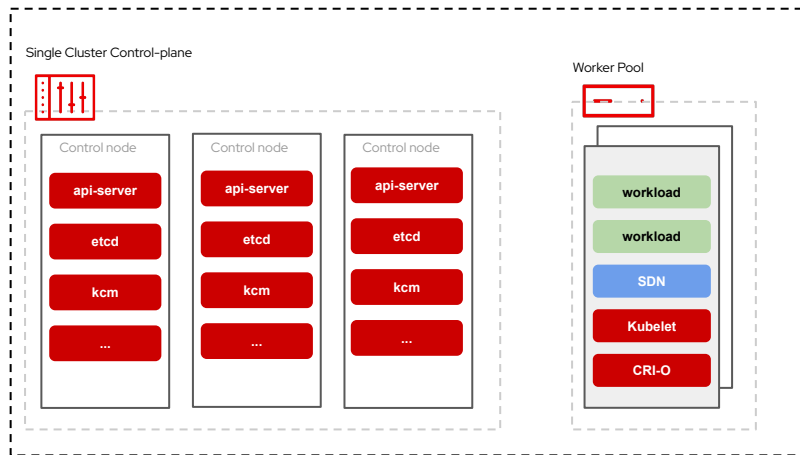


Hypershift Brings Externally Managed Control-Planes

Standalone OpenShift

Control-Plane (CP) + Workers

Standalone OpenShift **Cluster** (dedicated CP nodes)



Low CAPEX and OPEX costs
(bundling of CPs + CP as pods)



Central Management of CPs
(easy operation & maintenance)



Multi-arch support
(e.g. CP x86, workers ARM)



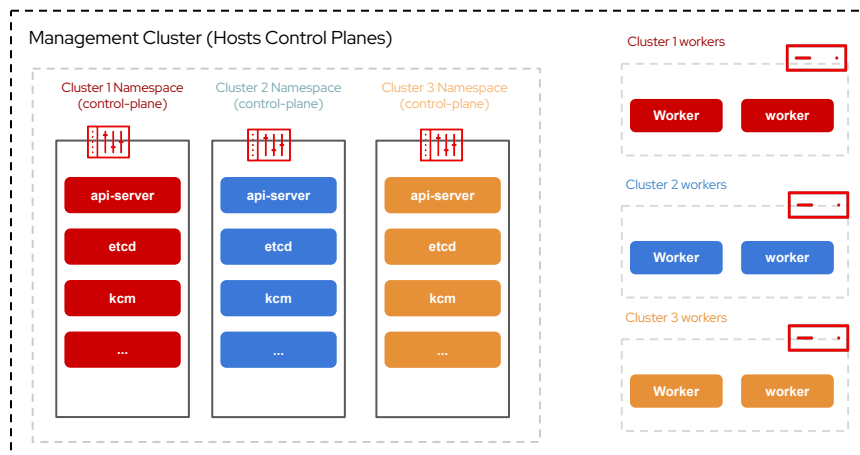
HyperShift

Control-Plane (CP)



Workers

HyperShift **Clusters** (decoupled CP and workers)



Network & Trust
segmentation




Mixed Iaas For CP and
Workers



Fast cluster bootstrapping
(CP as Pods)

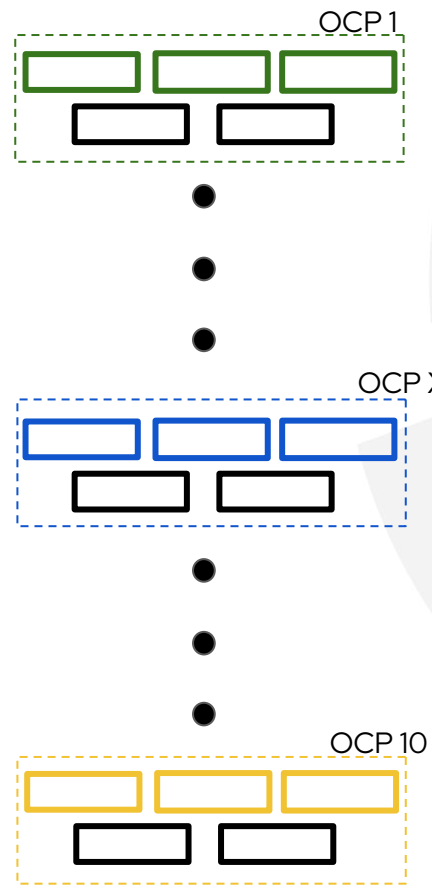



Management cluster




 **Red Hat**
Advanced Cluster
Management
for Kubernetes

Managed clusters





Without HCP

Management cluster




 **Red Hat**
Advanced Cluster
Management
for Kubernetes


Managed clusters



•
•
•

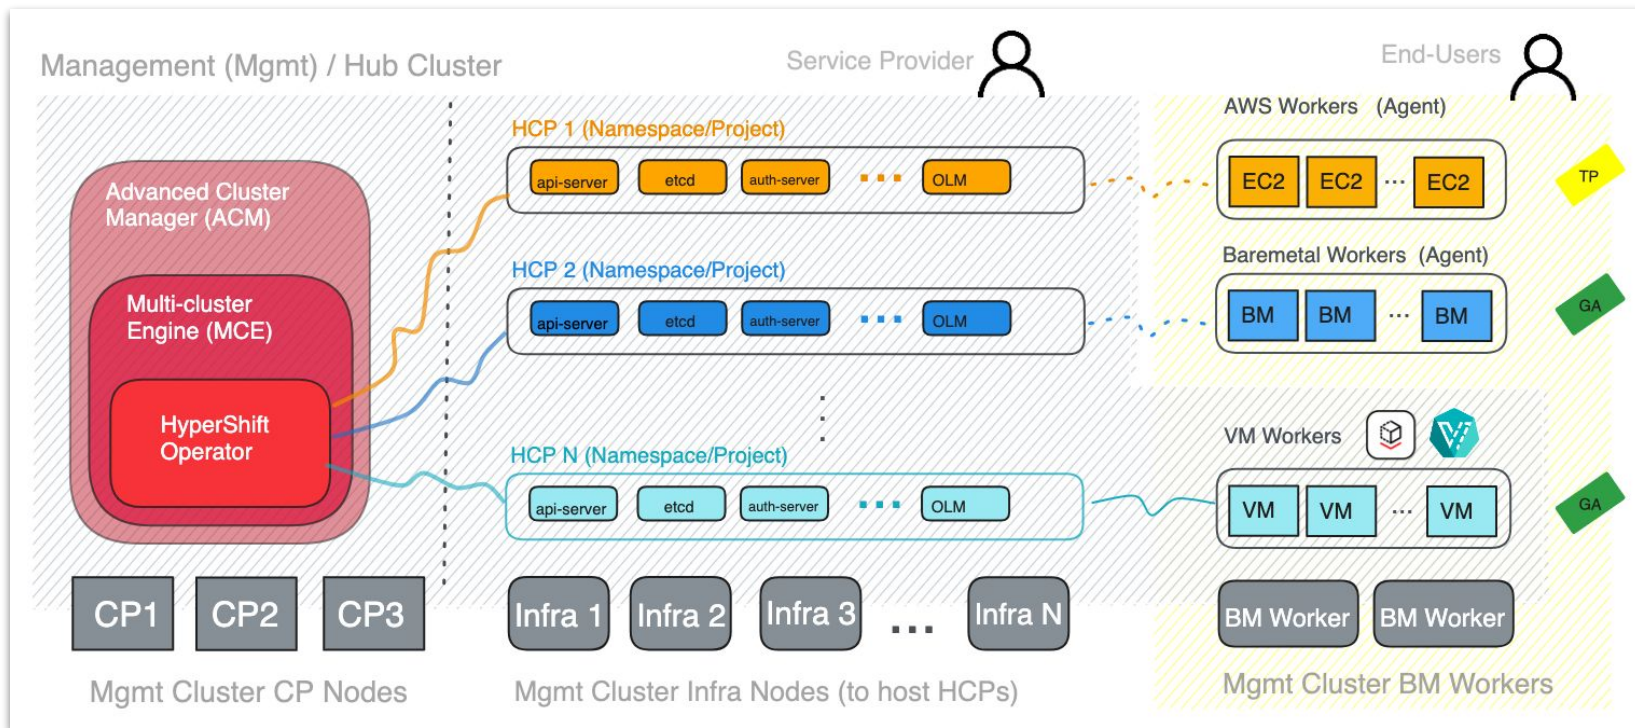


•
•
•



With HCP

Architecture Overview





Product

Project
Hypershift



Hosted Control Plane (HCP)



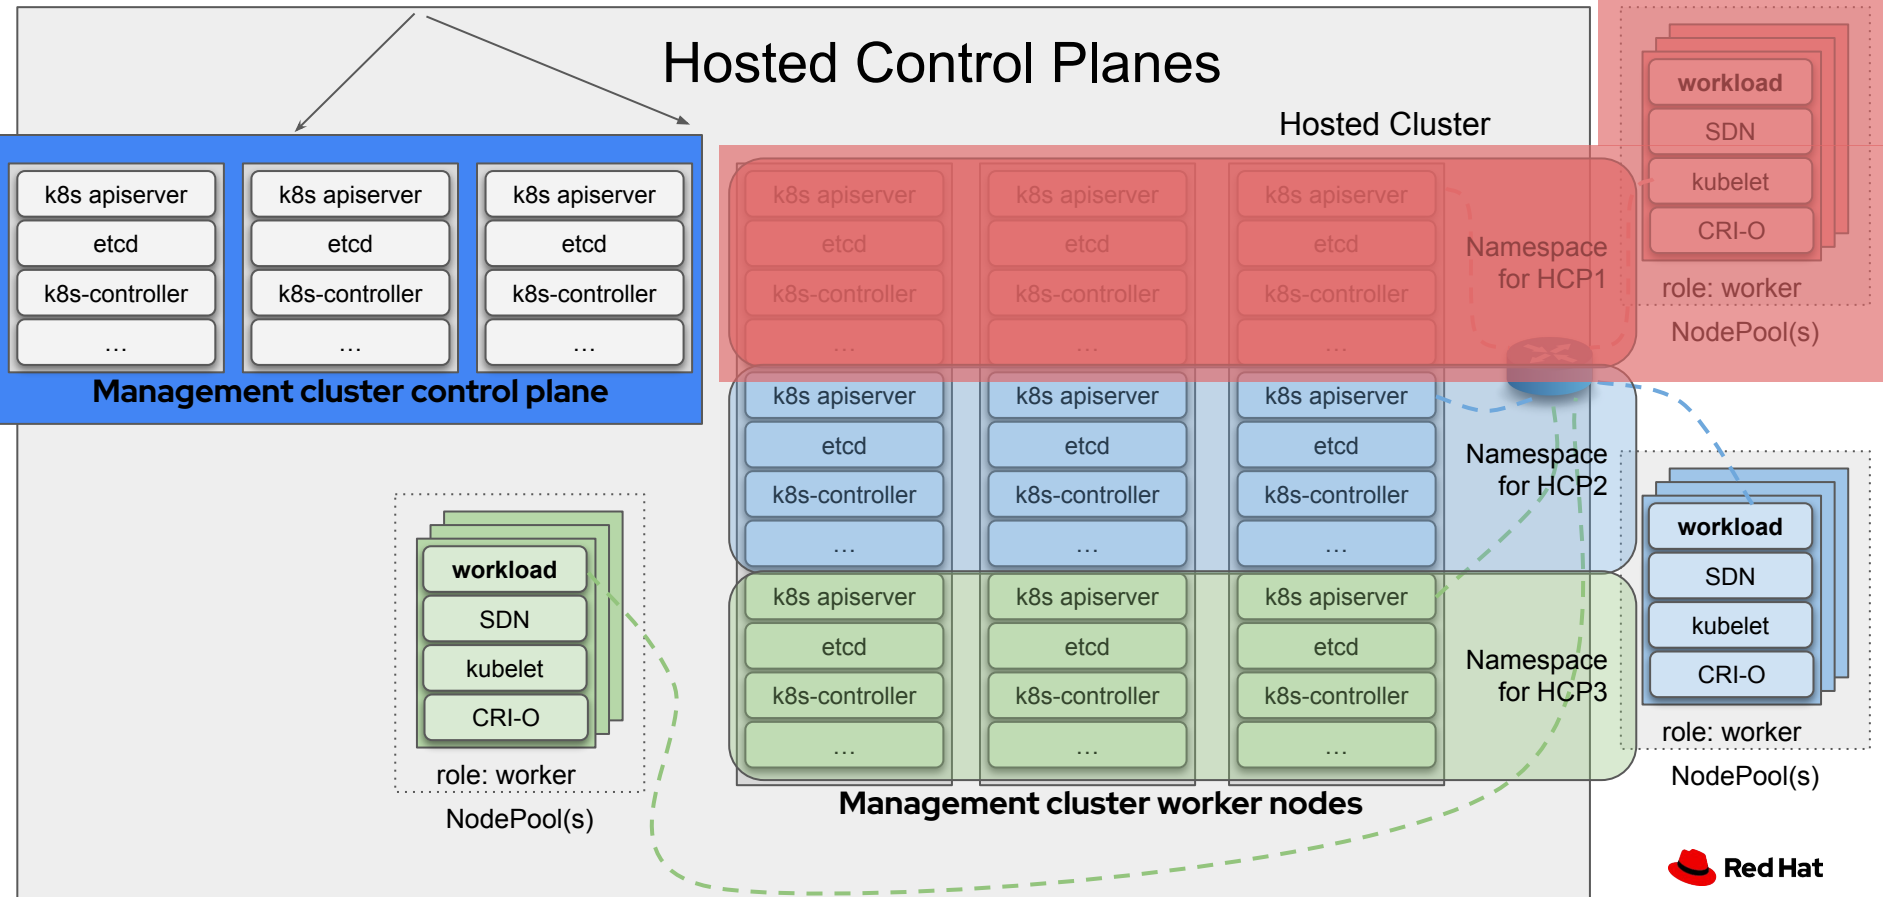
OCP 4.14 / ACM 2.9 / MCE 2.4

Release date : 11/2023



Standalone OpenShift Cluster

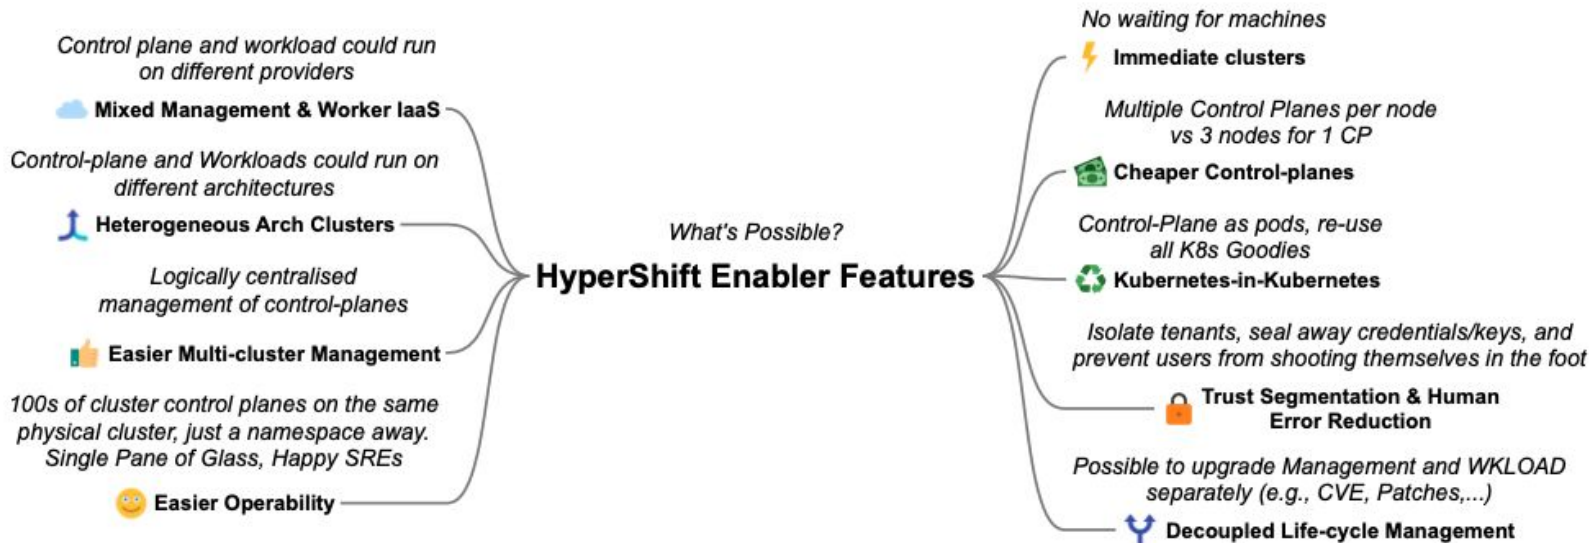
CP = Min(x=physical limit of node , y=max pods, z=etcd object sizes, ...)



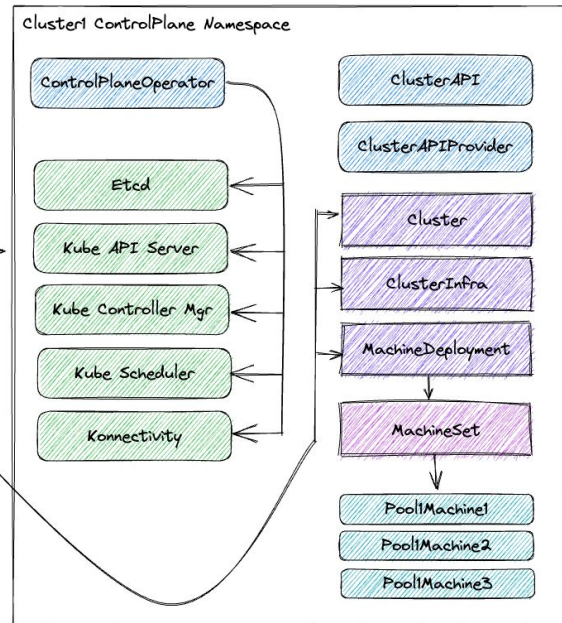


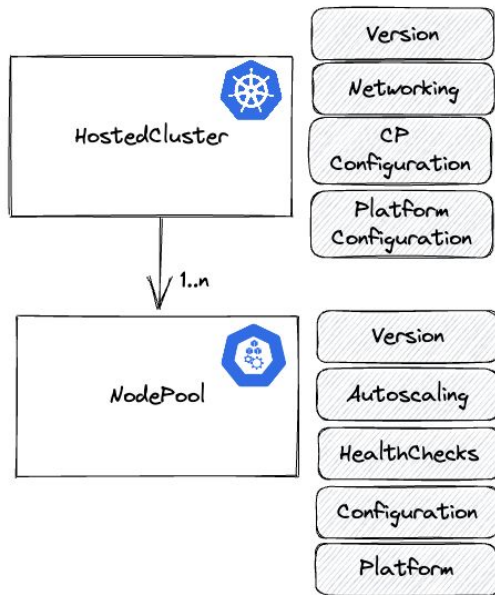
Why HyperShift?



OpenShift Clusters			
Self-Managed		Managed	
Connected		Disconnected	
Standalone			External CP
Compact SWO	Remote Workers	Normal K8s & Workers	HyperShift = (CP) + (Workers)

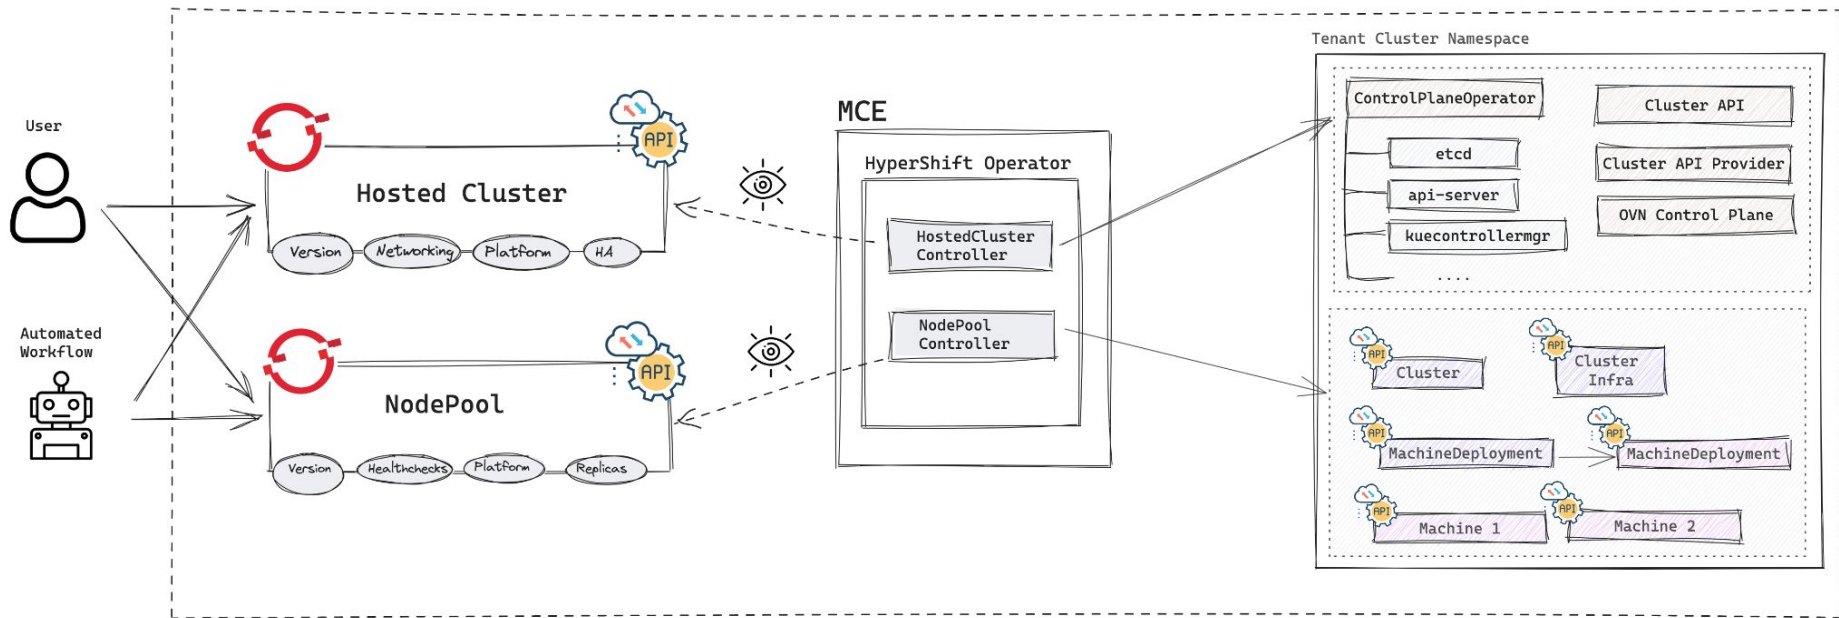


Hosted Control Planes APIs (Zoom-In to APIs)



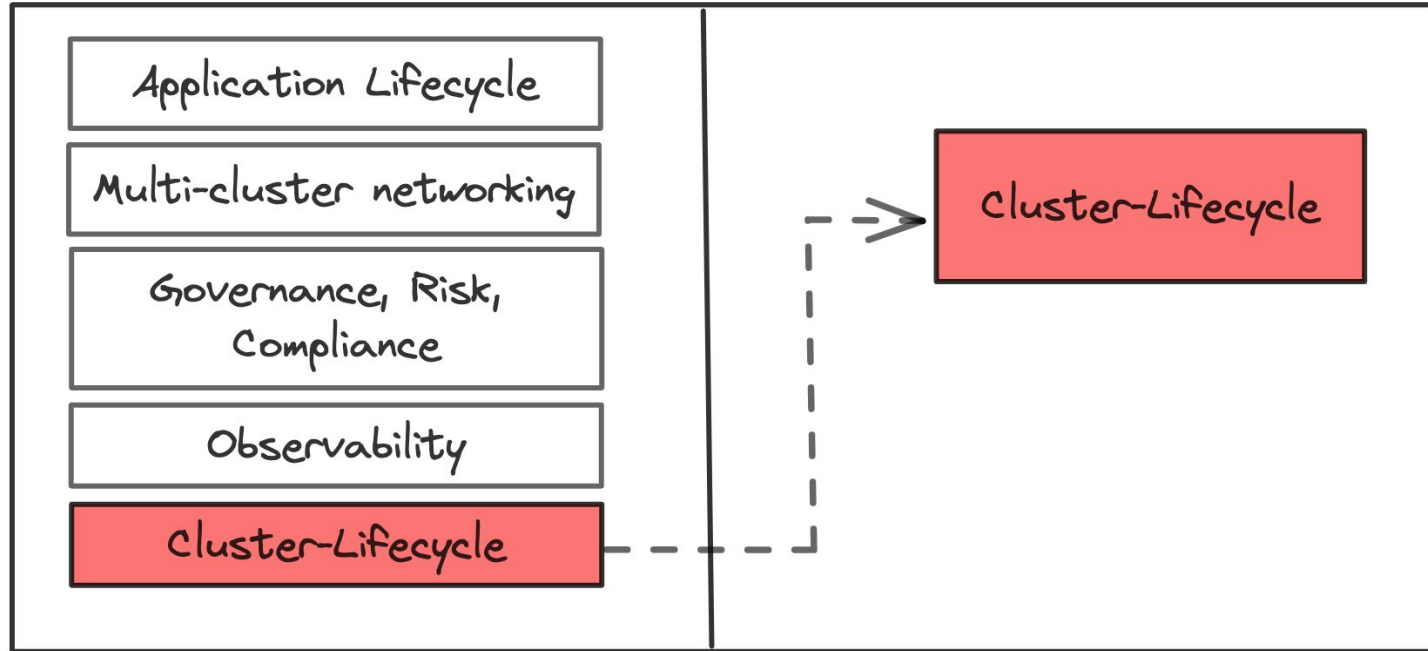
Hosted Control Planes APIs (Zoom-In to APIs)

Management Cluster

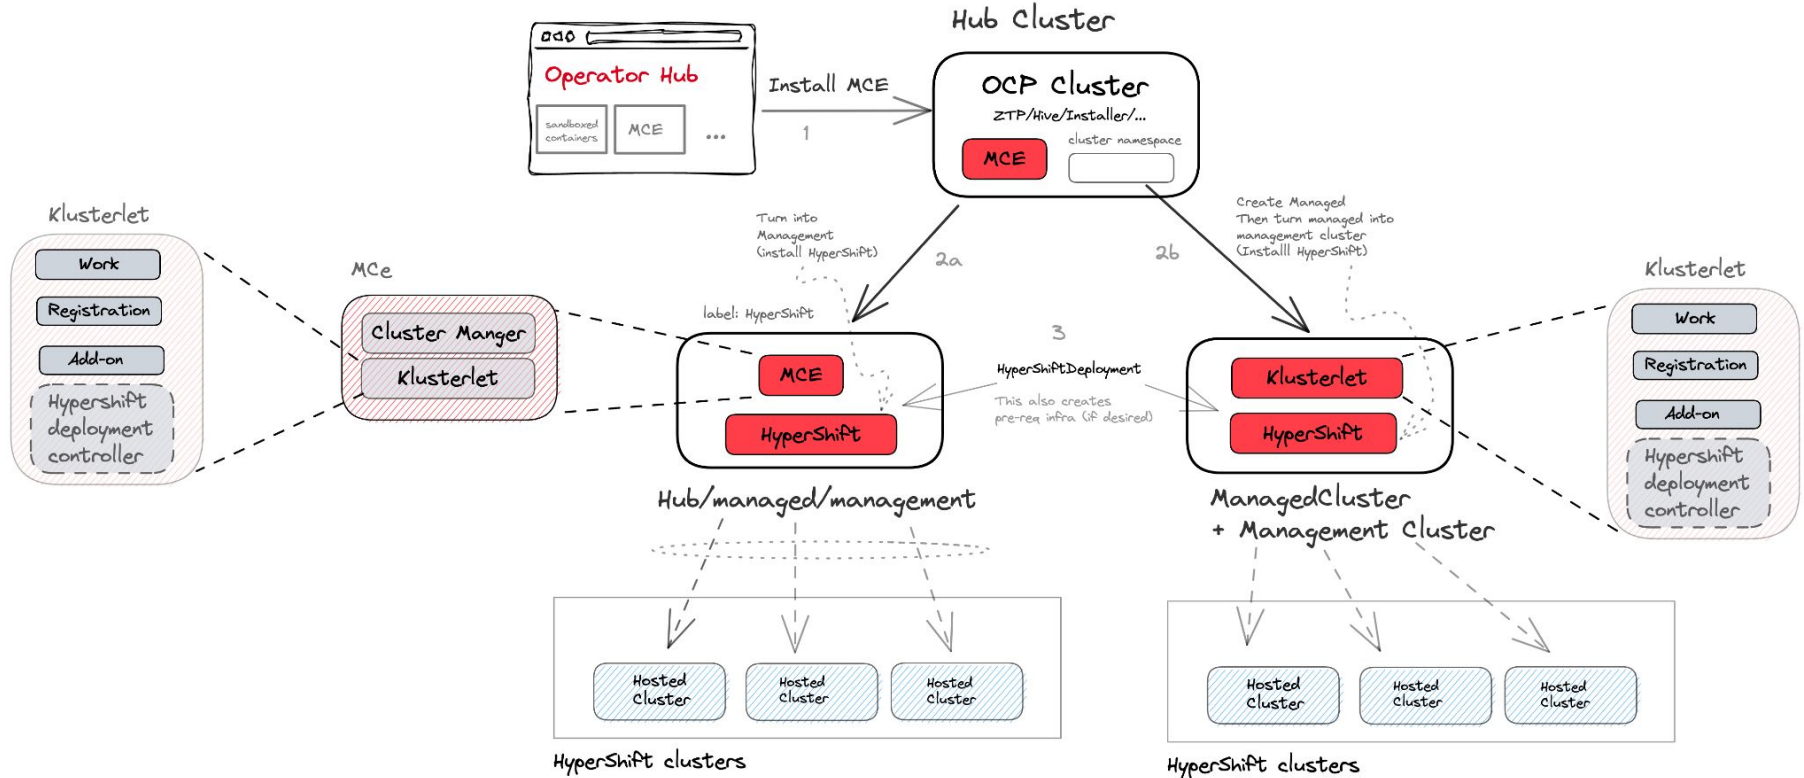


HCP is Available via ACM (MCE)

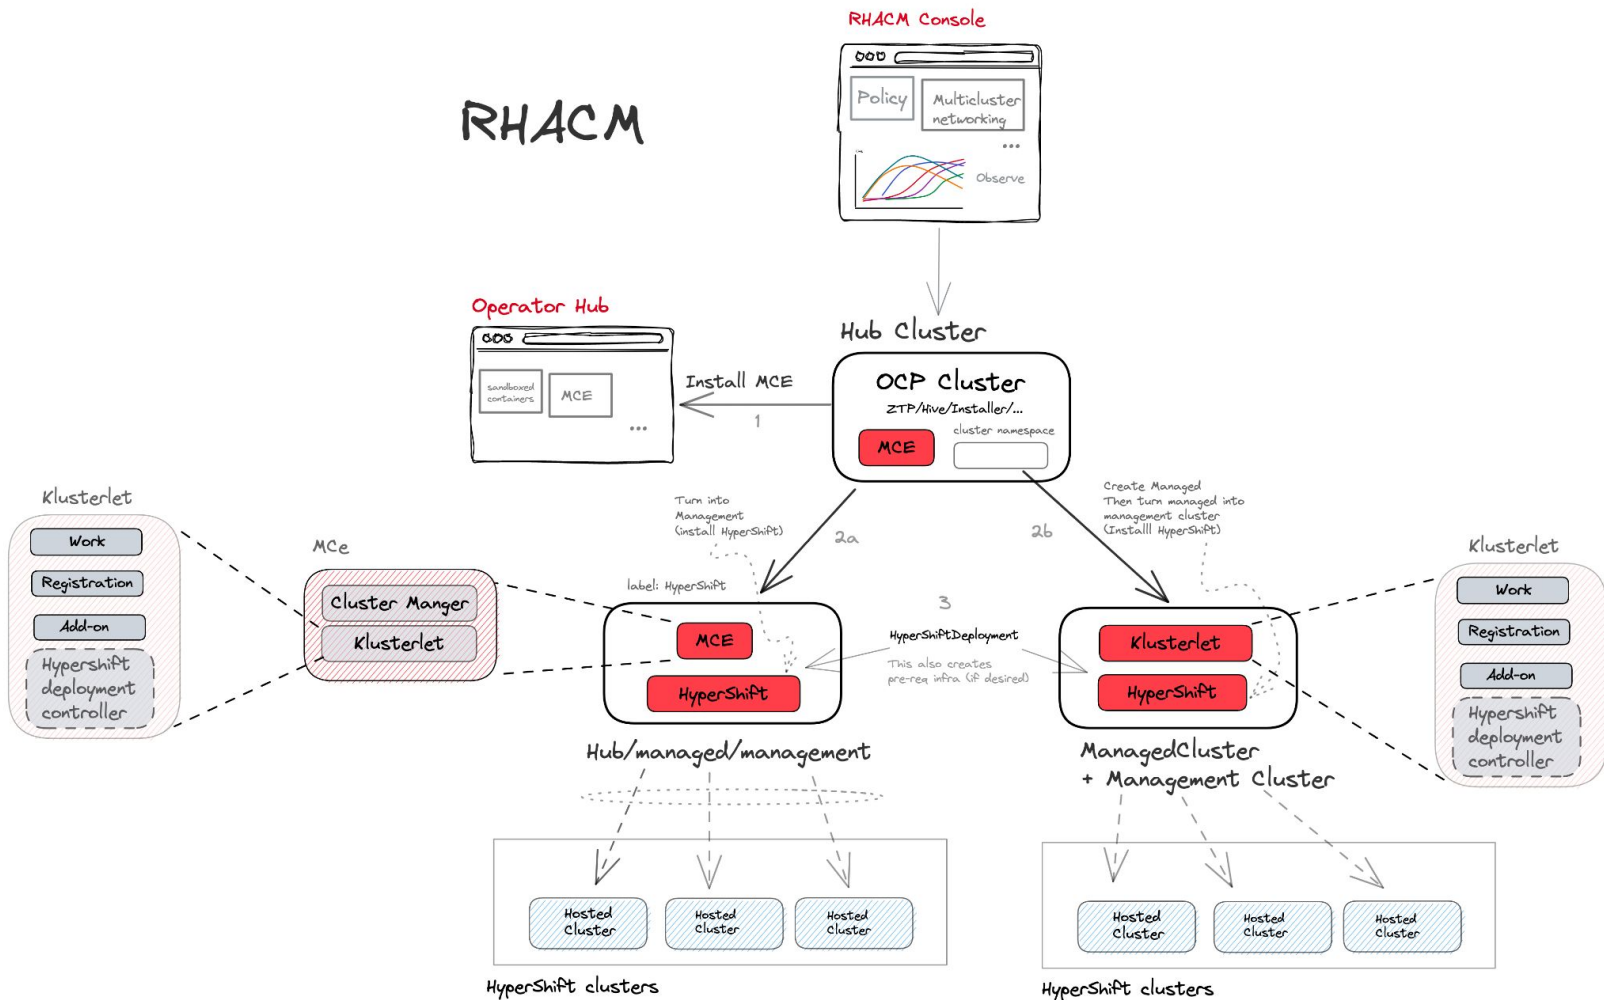
ACM vs. MCE



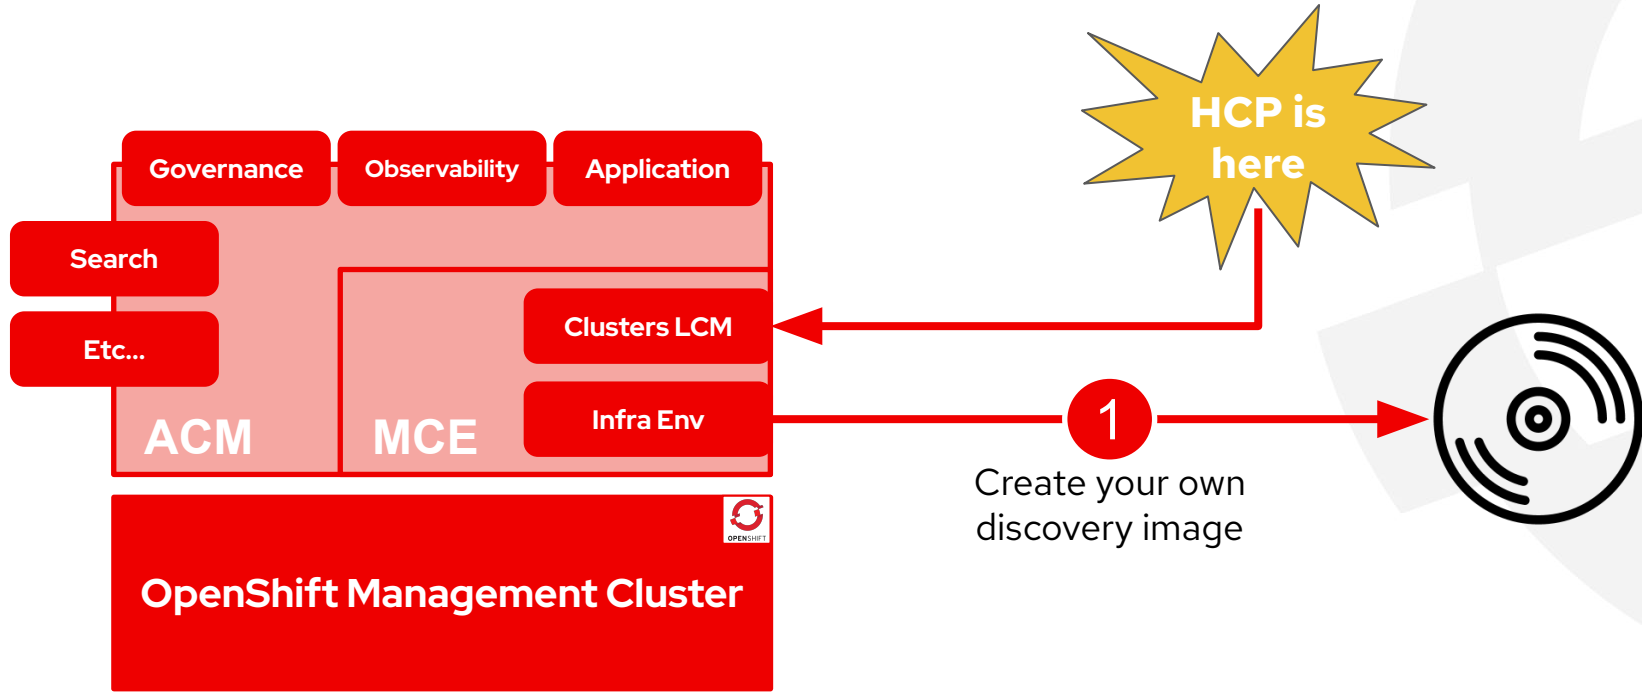
OpenShift/MCE



RHACM



Bare Metal provider is a **GA provider** for **HCP** and use existing method based on **Assisted Installer Service** included in ACM


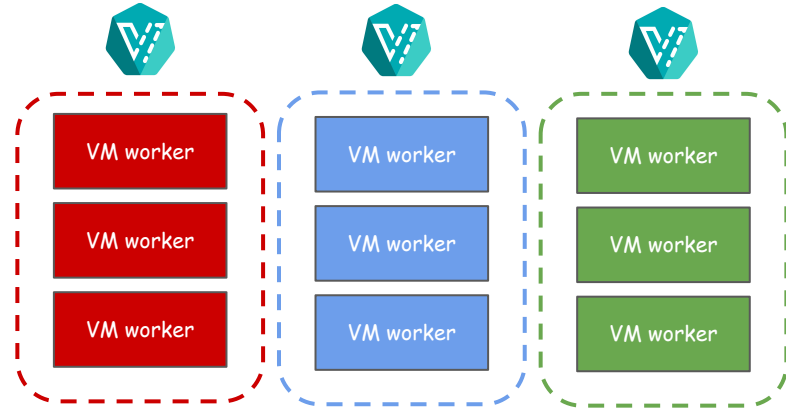
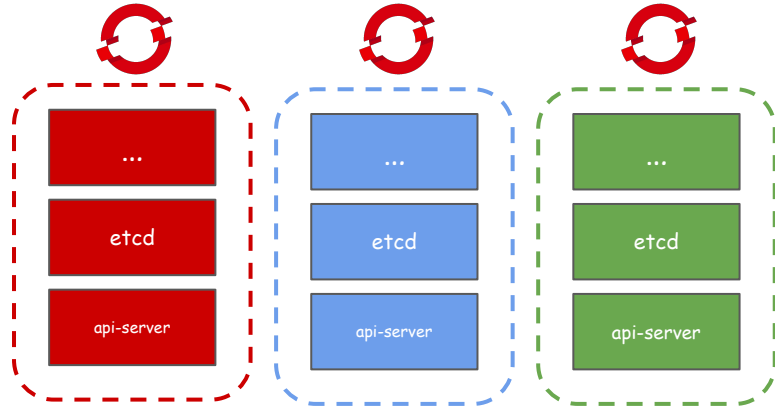


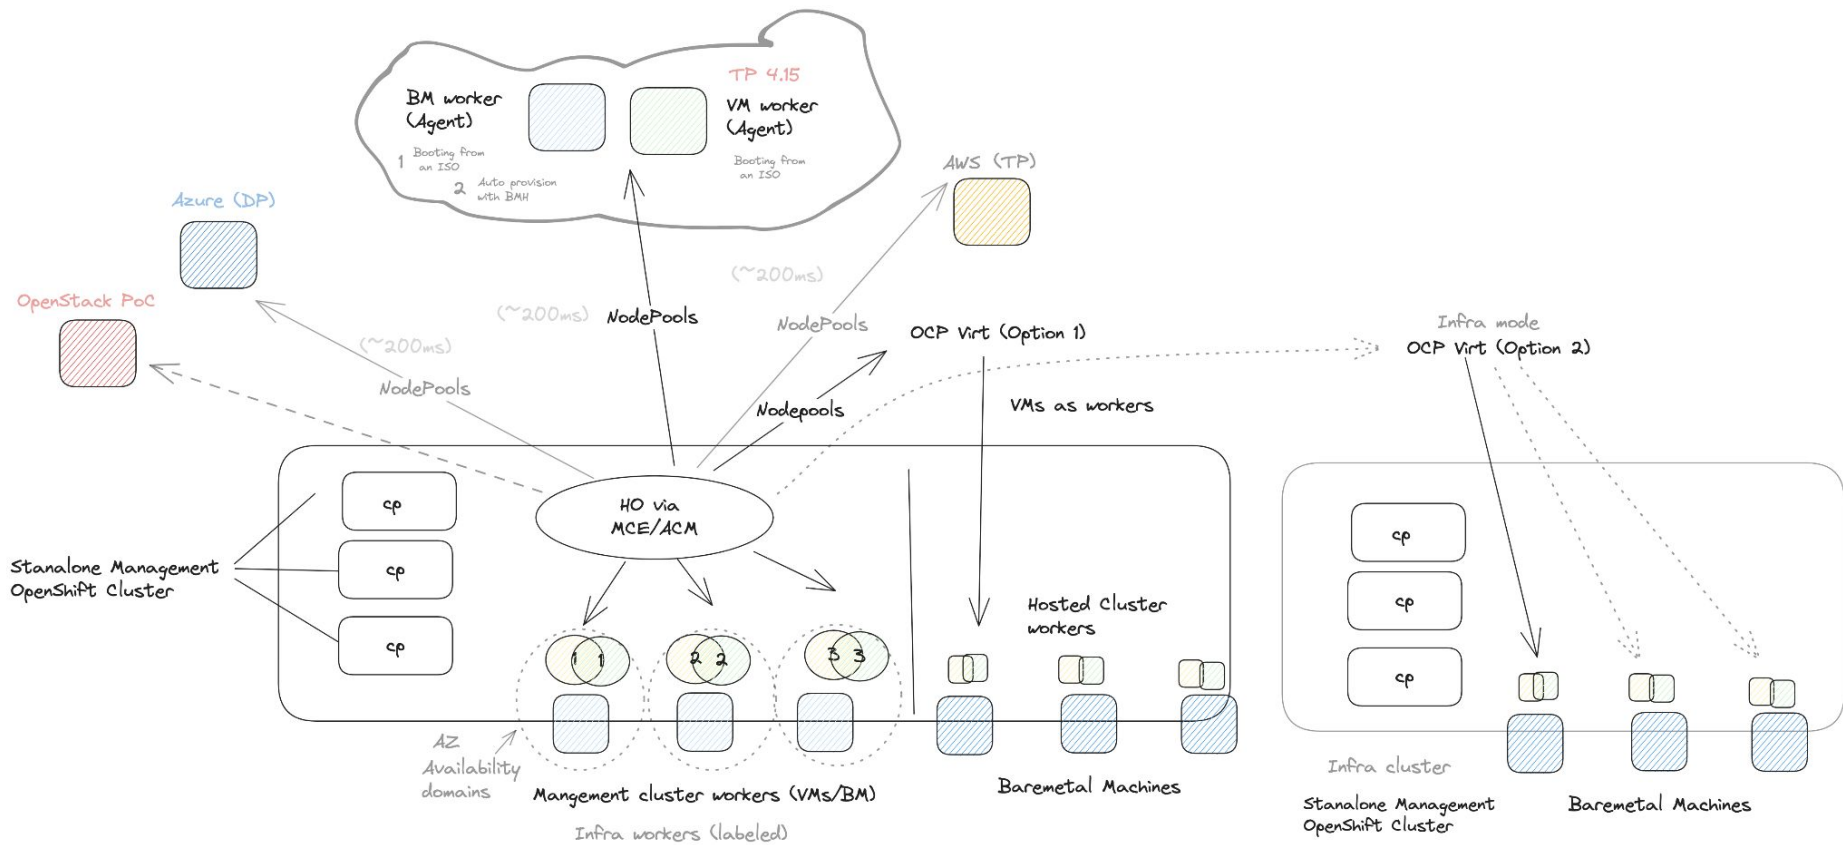
The OpenShift Virtualization Provider

Worker Nodes (hosted in VMs on OCP)

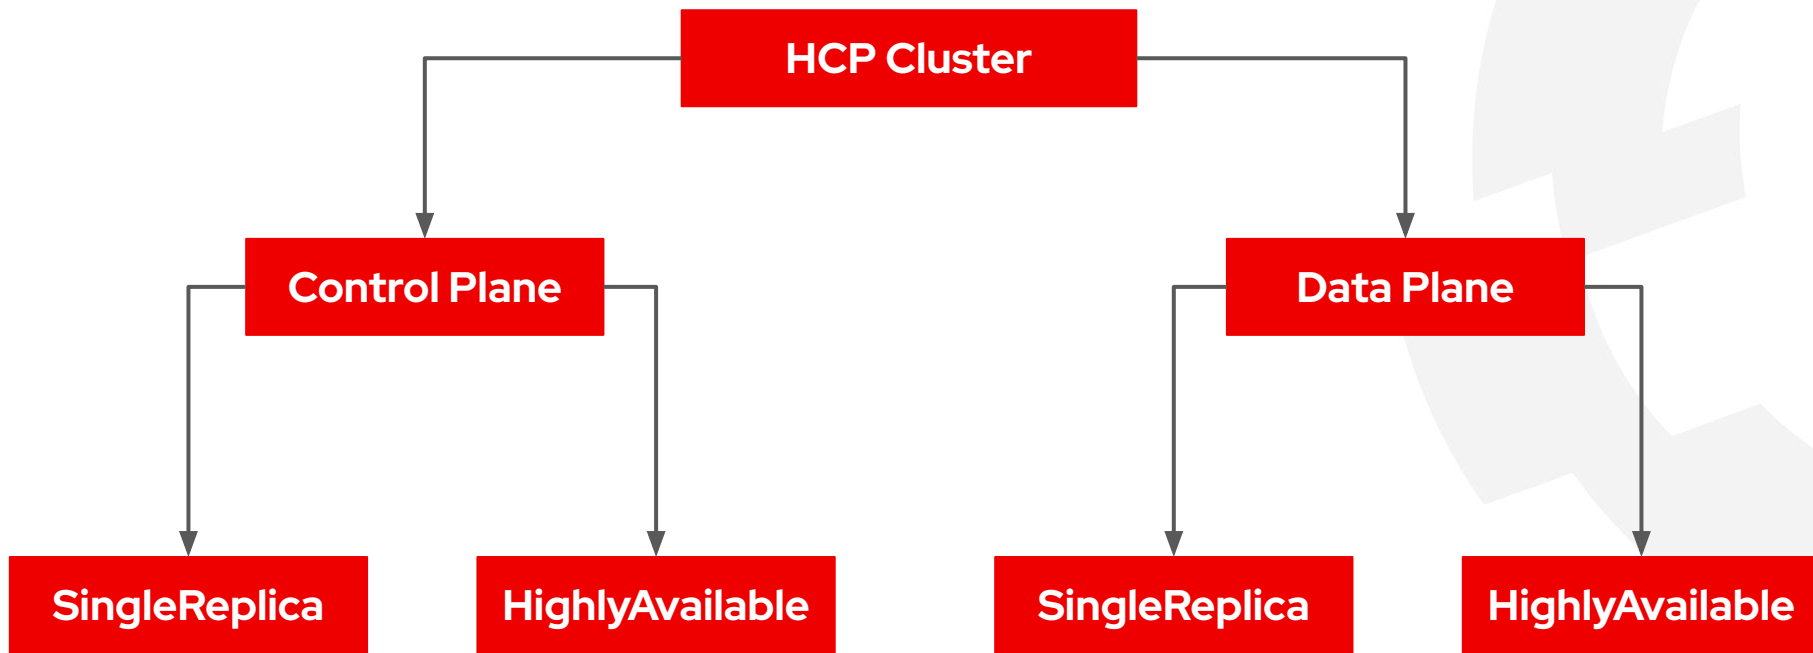


Control Planes (hosted in OCP)





High Availability



Create HCP cluster

Via ACM WebUI



Hosted

Run an OpenShift cluster where the control plane is decoupled from the data plane, and is treated like a multi-tenant workload on a hosting service cluster. The data plane is on a separate network domain that allows segmentation between management and workload traffic.

- ✓ Reduces costs by efficiently reusing an OpenShift cluster to host multiple control planes.
- ✓ Quickly provisions clusters.

Dedicated CLI



hcp - Hosted Control Plane Command Line Interface (CLI)

With the Hosted Control Plane command line interface, you can create and manage OpenShift hosted clusters.

- [Download hcp CLI for Linux for x86_64](#)
- [Download hcp CLI for Mac for x86_64](#)
- [Download hcp CLI for Windows for x86_64](#)
- [Download hcp CLI for Linux for ARM 64](#)
- [Download hcp CLI for Mac for ARM 64](#)
- [Download hcp CLI for Linux for IBM Power](#)
- [Download hcp CLI for Linux for IBM Power, little endian](#)
- [Download hcp CLI for Linux for IBM Z](#)

Create HCP cluster

Time Provisioning = ~10 min

```
export CLUSTER_NAME=hcp01
export PULL_SECRET="./pull-secret"
export SSH_KEY="./dm_key.pub"
export MEM="8Gi"
export CPU="4"
export WORKER_COUNT="3"
export BASE_DOMAIN=drkspace.fr
export CP_DEPLOYMENT_MODE="SingleReplica"
export INFRA_DEPLOYMENT_MODE="SingleReplica"
```

Variables definition

```
hcp create cluster kubevirt \
--name $CLUSTER_NAME \
--release-image $RELEASE_IMAGE \
--node-pool-replicas $WORKER_COUNT \
--pull-secret $PULL_SECRET \
--ssh-key $SSH_KEY \
--memory $MEM \
--cores $CPU \
--control-plane-availability-policy $CP_DEPLOYMENT_MODE \
--infra-availability-policy $INFRA_DEPLOYMENT_MODE
```

Cluster creation

<https://github.com/davmartini/redhat-techs/tree/main/openshift/hcp>

Create HCP cluster

The screenshot shows the Red Hat OpenShift Clusters management interface. The top navigation bar includes the Red Hat OpenShift logo, 'All Clusters' dropdown, and user information 'David Martini'. The left sidebar shows 'Infrastructure' and 'Credentials' sections. The main content area is titled 'Clusters' and contains a table of cluster information.

Cluster list Cluster sets Cluster pools Discovered clusters [Get started with Multicuster Hub](#)

Cluster management controls: Search, Filter, Create cluster, Import cluster, Actions, 1 - 2 of 2

Name	Namespace	Status	Infrastructure	Control plane type	Distribution version	Labels	Nodes	Add-ons	Creation date
<input type="checkbox"/> hcp01	clusters	Ready	Red Hat OpenShift Virtualization	Hosted	OpenShift 4.14.3	openshiftVersion-m... openshiftVersion-m... 8 more	3	2	15/12/2023, 10:01:05
<input type="checkbox"/> local-cluster	local-cluster	Ready	Other	Hub	OpenShift 4.14.5	openshiftVersion-m... openshiftVersion-m... velero.io/exclude-fro... 10 more	1	3	13/12/2023, 17:51:23

1 - 2 of 2 items 1 of 1 page

Create HCP cluster

hcp01

[Download kubeconfig](#)

[Actions](#) ▾

[Overview](#)

[Nodes](#)

[Add-ons](#)

▼ **Control plane status**

✔ > [Control plane](#)

✔ ▼ [Cluster node pools](#)

🔍 Search

Add node pool

1 - 1 of 1 ▾

< >

Node pool ↑	Status ↓	Distribution version ↓	Root volume ↓	Compute ↓	Nodes ↓	Health check ↓	Upgrade type ↓	Autoscaling ↓	
hcp01	✔ Ready	OpenShift 4.14.3			3	False	Replace	False	⋮

1 - 1 of 1 items ▾

<< <

1 of 1 page

> >>

HCP cluster on Management Cluster

Project: clusters-hcp01

Pods

Filter Name Search by name...

Name ↑	Status ↓	Ready ↓
capi-provider-845fd9b4b5-hf4qz	Running	1/1
catalog-operator-5df44dd8bc-lhmcg	Running	2/2
certified-operators-catalog-7ddfd77c96-4shq4	Running	1/1
cluster-api-54b7fb46f-w5d4z	Running	1/1
cluster-autoscaler-5b89666595-mgl2q	Running	1/1
cluster-image-registry-operator-57b667d574-2xfnk	Running	2/2
cluster-network-operator-dc856477c-2xwm4	Running	2/2

Control Plane Pods

Project: clusters-hcp01

VirtualMachines

Filter Name Search by name...

Name ↑	Status ↓
hcp01-8b7a3cf7-cz5pv	Running Not migratable
hcp01-8b7a3cf7-dlbdl	Running Not migratable
hcp01-8b7a3cf7-fw2xp	Running Not migratable

Data Plane VMs

Scale UP / Scale DOWN

Extend existing NodePool

1

Control plane status

- Control plane
- Cluster node pools

Search Add node pool 1-1 of 1

Node pool	Status	Distribution version	Root volume	Compute	Nodes	Health check	Upgrade type	Autoscaling
hcp01	Ready	OpenShift 4.14.3			3	False	Replace	False

1-1 of 1 items Manage node pool Remove node pool

Manage node pool

Namespace clusters

Name hcp01

OpenShift version 4.14.3

Number of nodes *

- 5 +

Update

Cancel

Scale UP / Scale DOWN


Create new NodePool

```
hcp create nodepool kubevirt \  
--cluster-name=hcp01 \  
--name=new-nodepool \  
--node-count=3 \  
--cores=2 \  
--memory=8Gi \  
--root-volume-size=32 \  
--node-upgrade-type=InPlace \  
1
```


2

Node pool ↑	Status ↓	Distribution version ↓	Root volume ↓	Compute ↓	Nodes ↓	Health check ↓	Upgrade type ↓	Autoscaling ↓
hcp01	Ready	OpenShift 4.14.3			3	False	Replace	False
new-nodepool	Pending	OpenShift 4.14.3			3	False	InPlace	False

HCP Upgrade - General aspects



Control plane and Data plane upgrades are decoupled



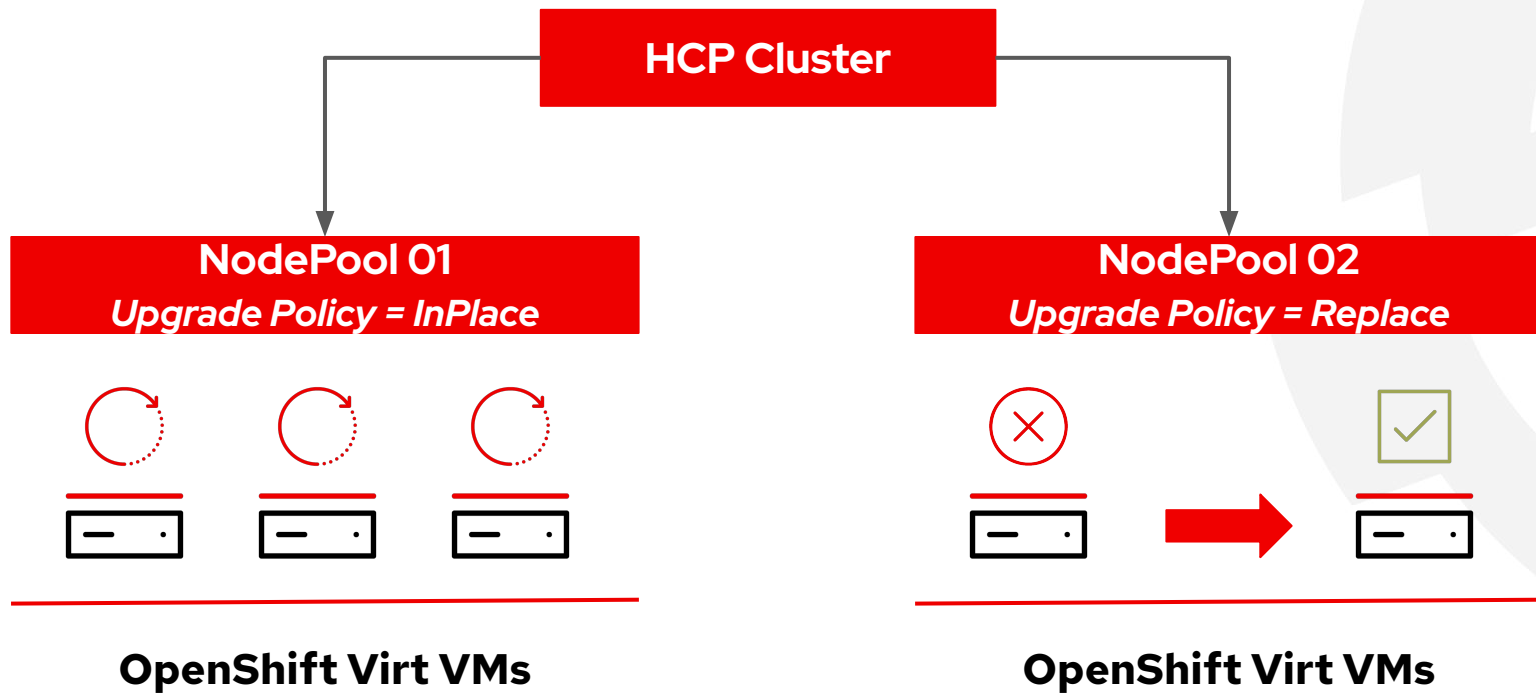
Each Node Pools can be upgraded separately



Different versions can be works in parallel following Version Skew Policy

Capture

HCP Upgrade - Methods



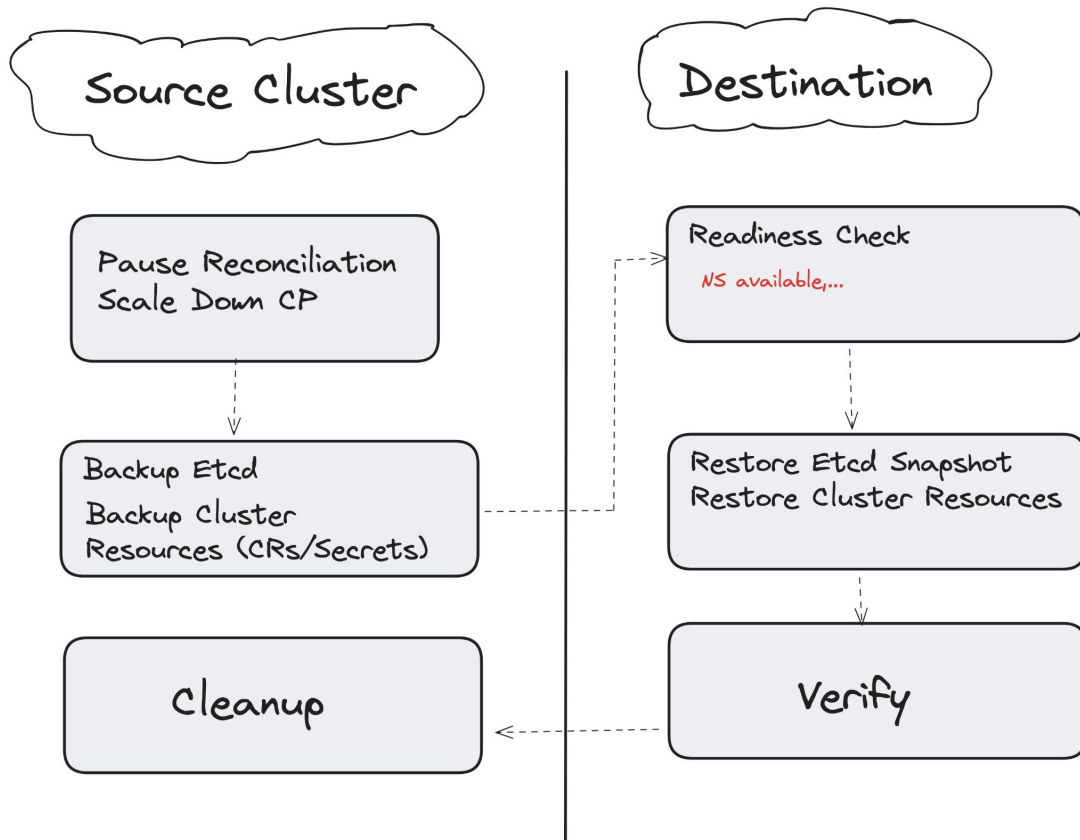
Failure Modes & DR

Failure Scenarios

Failure	Result
Loss of management cluster worker	Hosted control plane API is still available . Hosted cluster data plane is still available . Impacted hosted control plane member is rescheduled .
Loss of management cluster availability zone	Hosted control plane API is still available . Hosted cluster data plane is still available . Impacted hosted control plane maintains quorum .
Loss of management cluster control plane	Hosted control plane API is still available . Hosted cluster data plane is still available .
Loss of management cluster control plane and workers	Hosted control plane API is not available . Hosted cluster data plane is still available .

See a live [demo](#) of these situations!

Disaster Recovery & Migration of Control Planes




[Backup, restore, and disaster recovery for hosted control planes](#)

Performance & Sizing

Static Sizing (Request-Based)

Request-based Resource Consumption per control Plane		
72 pods per HA HCP		
CPU	Memory	Storage
5.5 cores	19 GO	12 Gig for Etcd 3 Gig for OVN



OpenShift Worker Nodes

BM | VM | Cloud | x86 | ARM

Demo

Show me how it works...

```
demo 1:cli# 2:hclusters 3:nodepoo> CPU:11.4% | MEM:12% 7.2G | 0.19 0.16 0.15 | cewong@cewong-rhel | Sep 13 21:37 100% ●  
→ ~
```