

Red Hat  
**Summit**

**Connect**

# What you didn't know about OpenShift

Pilar Bravo

OpenShift Specialist Solution Architect Role Leader



- ✓ Computer Engineer
- ✓ Started as J2EE developer
- ✓ Red Hatter since 2008
- ✓ Software engineer, Consultant and Solution Architect
- ✓ Based in Madrid, Spain

**Pilar Bravo**

OpenShift Specialist Solution Architect Role Leader



I have OpenShift.  
Can I ... ?

# Administration

## Assess and monitor health

I want to be aware of my exposure to issues that can affect service availability, fault tolerance, performance or security.

## Recommendations

I want to receive personalised recommendations to avoid risks and get resolutions tailored to my individual clusters



## Support

I want to have proactive support so I can avoid problems before they happen.

## Knowledge

I want to take advantage of all the knowledge Red Hat has from the thousands of customers using OpenShift.

# Administration: Insights Advisor for OpenShift

- ▶ **Free service leveraging Red Hat experience with supporting and operating OpenShift**
- ▶ Insights Advisor UI adds **Upgrade risks (Preview)** - ML powered technology to identify potential blockers that could interrupt OCP upgrade flow
- ▶ **New Insights recommendations** focusing on preventing issues with OpenShift Data Foundation (ODF), OpenShift Cluster Version Operator (CVO) and OpenShift Cluster Autoscaler Operator (CAO)
- ▶ Insights recommendations available in Hybrid Cloud Console -> **Cluster History** page

<https://console.redhat.com/openshift/advisor>  
<https://console.redhat.com/settings/notifications/openshift>

The screenshot displays the OpenShift Insights Advisor interface. On the left is a navigation sidebar with options like Clusters, Overview, Releases, Downloads, Insights, Advisor, Subscriptions, Cost Management, Support Cases, Cluster Manager Feedback, Red Hat Marketplace, and Documentation. The main area shows 'Advisor recommendations' with a search bar and a table of recommendations. A red box highlights the '1 - 14 of 14' pagination control. Below the table, a detailed view for a specific recommendation (UID: 60e1fbfb-308f-4f1f-a212-c21a5097d05f) is shown, featuring a yellow warning banner for 'Resolve upgrade risks' and a table of 'Alerts firing' and 'Cluster operators'.

Name	Added	Category	Total risk	Clusters
Cluster upgrade will fail when default SCC gets changed	2 years ago	Service Availability	Important	1
Workloads are still using the deprecated APIs which will be removed in the next release	5 months ago	Service Availability	Important	6
SystemMemoryExceedsReservation alerts when the system daemons memory usage on nodes exceeds 90% of the reservation for them	4 months ago	Service Availability	Important	1
Workloads are using the deprecated PodSecurityPolicy API	5 months ago	Performance	Moderate	3
CVE-2021-30465: runc vulnerable to privilege escalation	9 months ago	Security	Moderate	1
Nodes will become Not Ready due to a CRI-O PID leak in the running OpenShift Container Platform version	5 months ago	Service Availability	Moderate	13
The running OpenShift version has reached its End of Life	2 years ago	Service Availability	Moderate	1
Pods could fail to start if openshift-samples is degraded due to FailedImageImport which is caused by a hiccup while talking to the Red Hat registries	2 years ago	Service Availability	Moderate	1

Name	Status	Namespace
ClusterOperatorDown	Critical	openshift-cluster-version
ClusterOperatorDown	Critical	openshift-cluster-version
etcdMembersDown	Critical	openshift-etcd

Name	Status	Message
authentication	Available	WellKnown_NotReady
machine-config	Available	-
authentication	Available	APIServerDeployment_UnavailablePod:WellKnownReadyController_SyncError

# Security

## Sandboxing for testing software

I want to run a containerized workload with known vulnerabilities or to handle an issue in a legacy application.

I want to have administrative control over pods or load custom kernel modules.

## Kernel isolation

I want to securely run workloads that require custom kernel tuning and being able to create custom kernel modules.



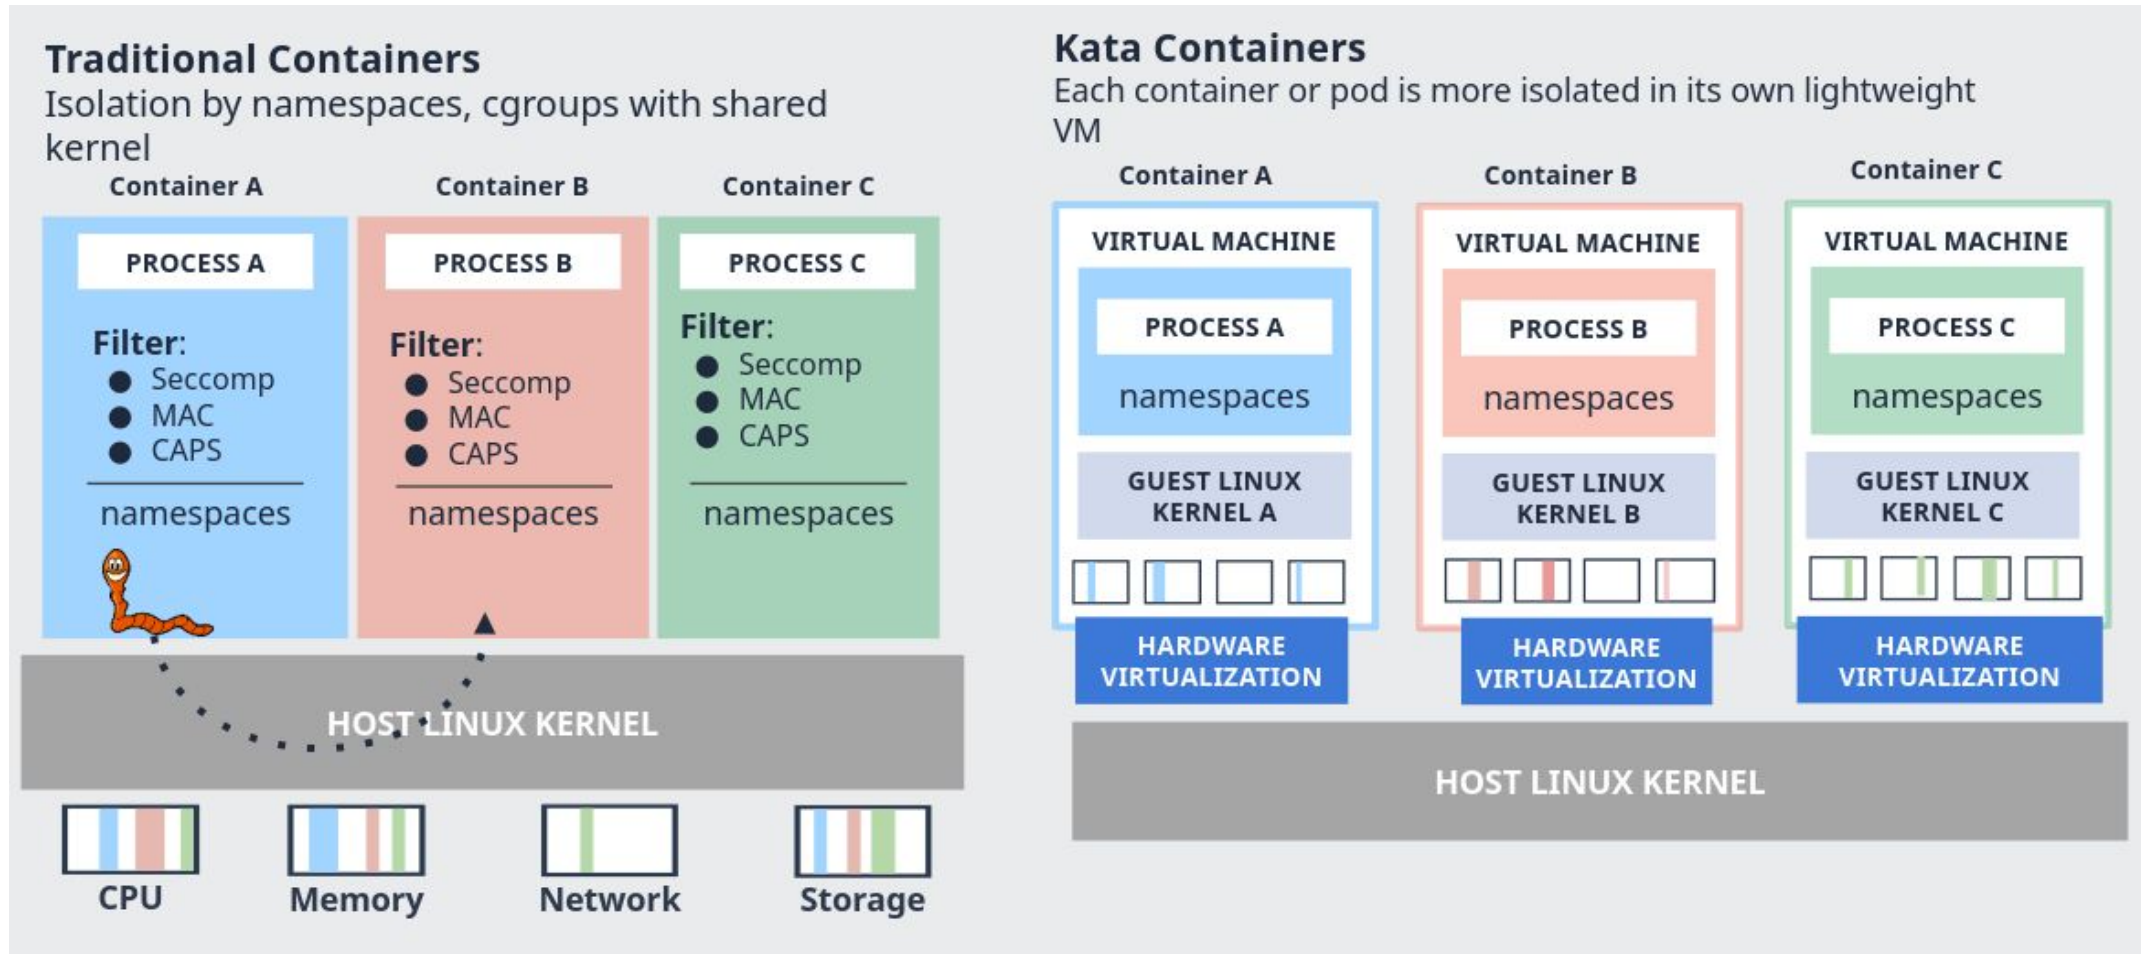
## Noisy neighbour (interference)

I want to run third-party workloads from multiple vendors, (like CNFs) without their custom settings interfering with packet tuning or with `sysctl` variables set by other apps.

## Privileged workloads

I want to securely run workloads that require elevated root privileges, like access to a specific physical device.

# Security: OpenShift Sandboxed Containers



# Virtualization

## Modernization

I want to modernize my virtualization workloads, but transforming everything in containers is slow and expensive.

## Single Architecture

I want to have a single cloud native platform to be able to host all my workloads.



## Vintage Applications

I want to keep my vintage applications as they are. I do not want to migrate them to containers.

I want to keep my Windows VMs as they are for the moment.

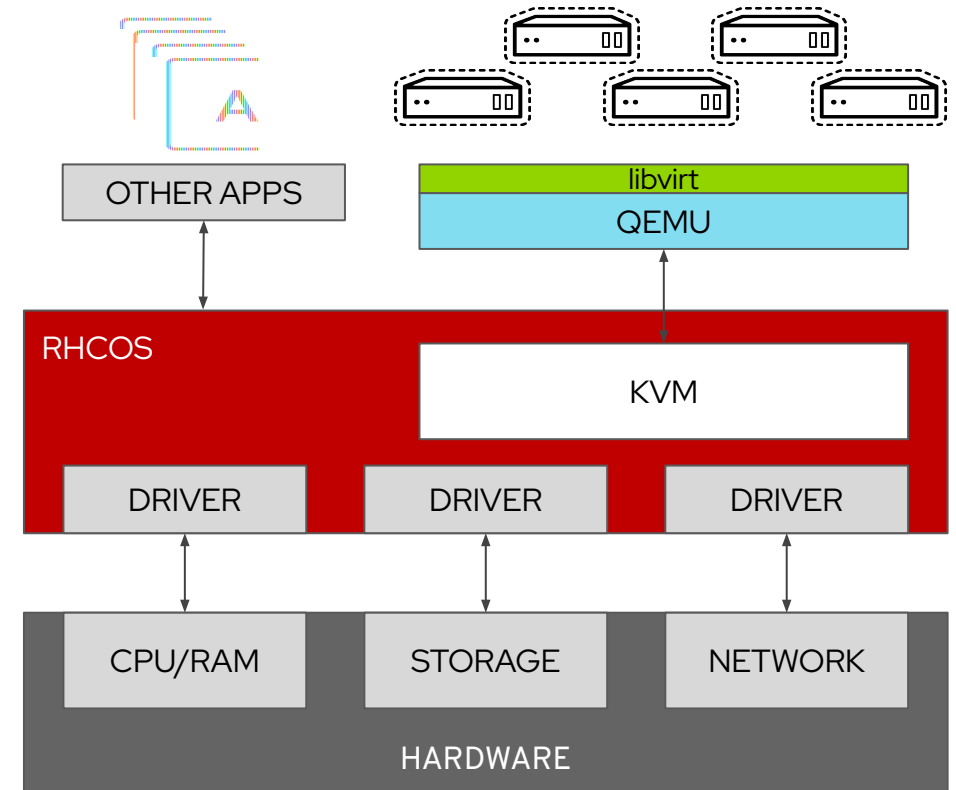
## Consistent Developer Experience

I want to have the same tools, processes and pipelines for all my applications. I want an unified control and DevOps pipeline.



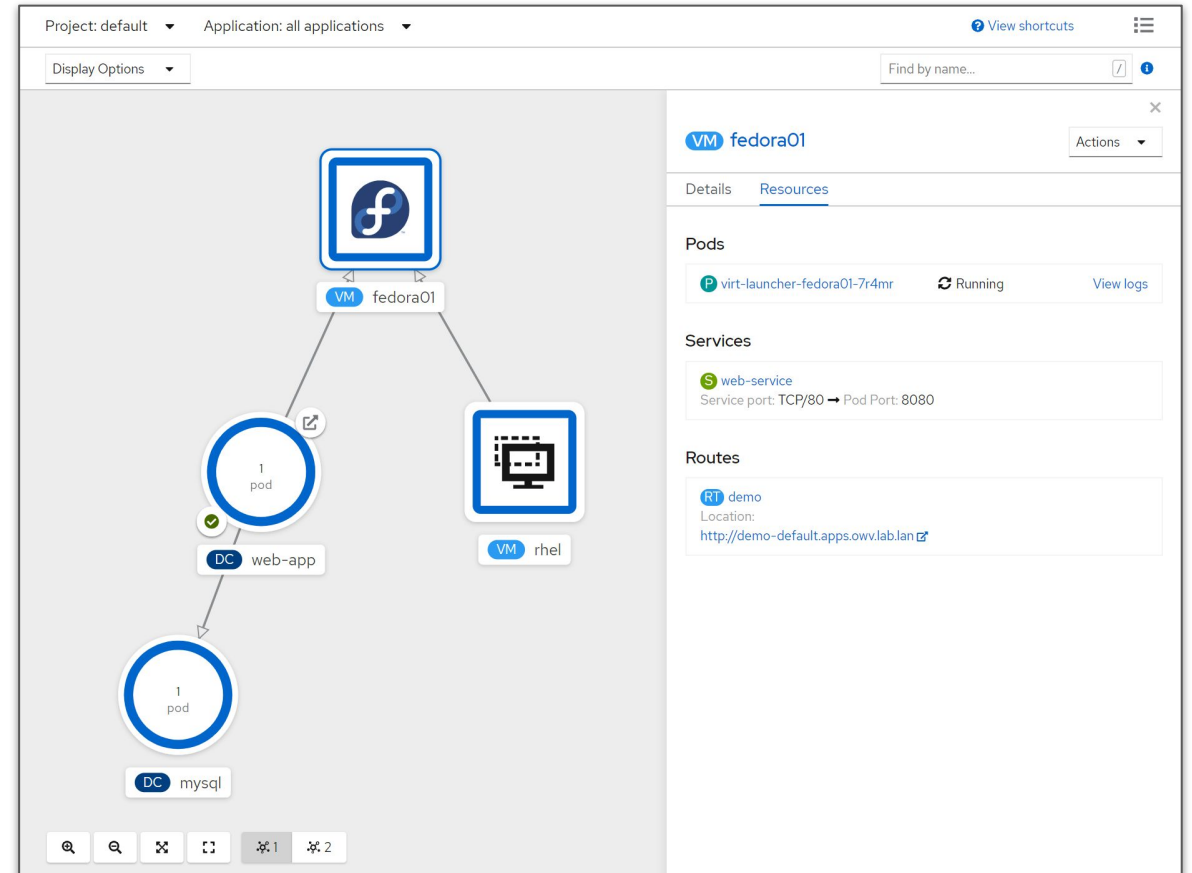
# Virtualization: OpenShift Virtualization

- Virtual machines
  - ◆ Running in containers, managed as Pods
  - ◆ Using the KVM hypervisor
- Scheduled, deployed, and managed by Kubernetes
- OpenShift Virtualization uses KVM, the Linux kernel hypervisor
- KVM is a core component of the Red Hat Enterprise Linux kernel
  - ◆ KVM has 10+ years of production use: Red Hat Virtualization, Red Hat OpenStack Platform, and RHEL all leverage KVM, QEMU, and libvirt
- QEMU uses KVM to execute virtual machines
- libvirt provides a management abstraction layer



# Virtualization: OpenShift Virtualization

- Virtual Machines connected to pod networks are accessible using standard Kubernetes methods:
  - ◆ Service
  - ◆ Route
  - ◆ Pipelines
  - ◆ etc.
- Network policies apply to VM pods the same as application pods
- VM-to-pod, and vice-versa, communication happens over SDN or ingress depending on network connectivity



# Virtualization: OpenShift Virtualization

## Migration Toolkit for Virtualization

- Easy to use UI
- Mass migration of VMs from VMware, Red Hat Virtualization and OpenStack to OpenShift (Tech Preview)
- VM data pre-copied before shutdown (Warm Migration) for VMware and RHV migrations
- VM validation service: Run checks on VM configuration to avoid migration issues
- Parallelized VM conversion
  - ◆ Maximize throughput
- Migration Network Selection
  - ◆ Avoid impact on other running workloads

The top screenshot displays the 'Providers' page in the OpenShift Migration Toolkit. The left sidebar shows the navigation menu with 'Migration' expanded to 'Providers for virtualization'. The main content area shows a table of providers:

Name	Status	Type	Endpoint	VMs	Net...	Hosts
vsphere	Ready	VMware	https://rhev-node-05.rdu2.scalelab.redhat.com/sdk	7736	2	5
rhv	Ready	oVirt	https://hosted-engine-08.lab.eng.tlv2.redhat.com/ovirt-engine/api	8	2	3
ocp	Ready	KubeVirt		5	1	

The bottom screenshot displays the 'NetworkMaps' page. The left sidebar shows the navigation menu with 'Migration' expanded to 'NetworkMaps for virtualization'. The main content area shows a table of network maps and a diagram illustrating network mapping:

Name	Source...	Target...	From	Status	To
vsphere-map	vsphere	ocp	1	Ready	Pod network

The diagram shows a box labeled 'Mgmt Network /Datacenter/network/Mgmt Network' with an arrow pointing to a box labeled 'Pod network'.

# Flexibility

## Agility

It takes longer than expected to create a cluster. I want to be able to create clusters quickly, even as part of a pipeline as temporary clusters.

## Restrict administrator access

As an administrator, I don't want the cluster users to have access to the control plane.



## Resource Usage


It takes a lot of compute to run my control plane, and I am wasting resources most of the time. I want to optimize resource usage.

## Lack of space

I want more space for the workloads in my cluster, instead of provisioning more hardware.

# Flexibility

## Where is the majority of your Kubernetes spend?

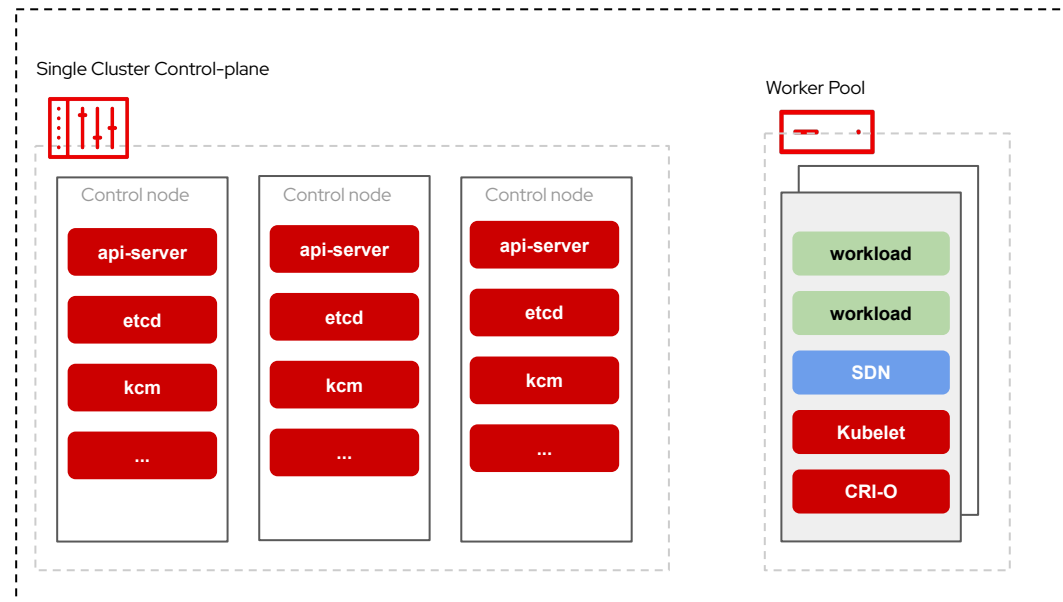


# Flexibility: OpenShift Hosted Control Planes

## 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)\*



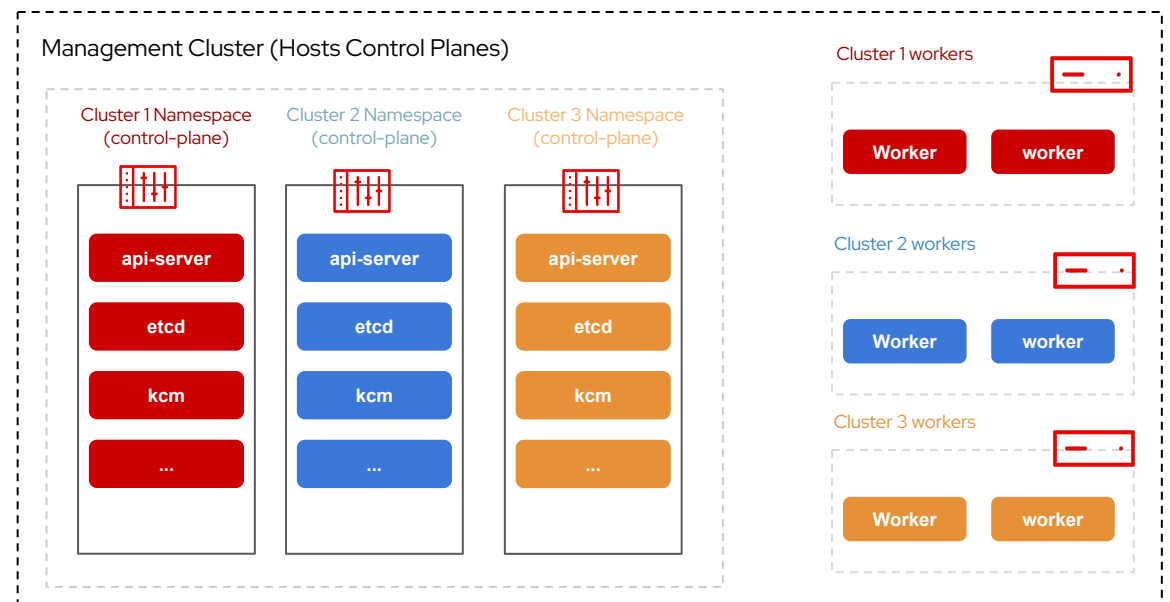
## HyperShift

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)



# Hosted Control Planes + Virtualization

## **I am running OpenShift on VMs**

I want to have a single cloud native platform to be able to host all my workloads without entitlements needed for my bare metal hardware.



## **I am running OpenShift on bare metal**

I want to host multiple OpenShift clusters and increase utilization.

## **I provide dedicated clusters**

I want to faster cluster provisioning, offload administration to cluster admins, improve cluster isolation and be able to run different OpenShift versions.

# Hosted Control Planes + Virtualization



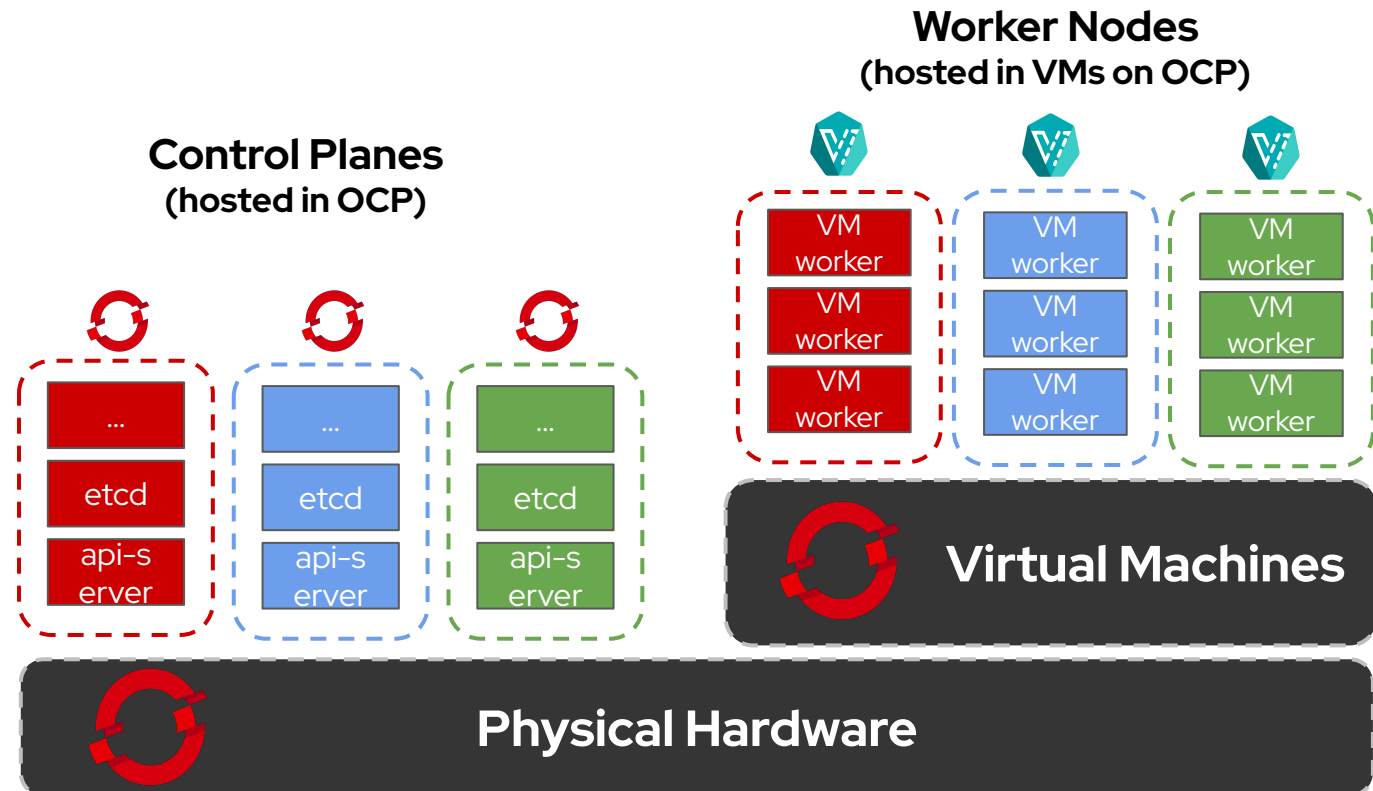
## Increase Utilization of Infrastructure

- Reduce unused and underutilized infrastructure
- Increase bare metal node utilization by deploying multiple hosted clusters.



## Reduce Dependency on Legacy Virtualization

- Eliminate legacy hypervisor hosting your container platform.
- Underlying virtualization layer is included with hosted OpenShift cluster entitlements (no separate licensing needed)






# Edge

**I have different platforms with different resources and locations**

Can I work the same way, with the same deployment model, everywhere? Can I move my workloads wherever it makes more sense?



**My apps have various components**

Can I upgrade each one separately? Can I standardize the way I work in the datacenter across all my devices?


**I need to create real time applications**

Can I use OpenShift for that? Can I move my applications closer to the users? Can I enable new business models with OpenShift?

# OpenShift at the Edge



**OPENSIFT  
SINGLE NODE**



**OPENSIFT  
REMOTE WORKER NODES**



**OPENSIFT  
3-NODE**

Definition

**OpenShift** deployment on a single box (Control Plane + Worker) with fair amount of resources.

**OpenShift** Control Plane resides in a central location ("good enough" network required), whereas Workers are distributed sharing the control plane.

**OpenShift** Control Plane and Workers reside on the same node. HA setup with 3 servers.

Examples

In-vehicle field operations

Telco 5G sparsely populated areas

In-field single server operations

Disconnected environments

Telco 5G far edge - RAN

IoT / data collection gateways

Visual Inspection using Machine Vision

Critical Workload requiring High Availability

Telco 5G near edge & MEC

Edge AI & Data pipelining

Smart manufacturing

Remote office

Disconnected clusters

# Artificial Intelligence

## **Tested and supported AI/ML tools**

I want to use tools like JupyterHub, TensorFlow and PyTorch in OpenShift, fully upgraded, tested and supported.

## **Freedom of choice**

I want to extend OpenShift with more services of my choice. Including not only Red Hat technology but partners, like Anaconda, IBM, Seldon, etc...




## **I want to start fast and scale quickly**

As a data scientist, I want extra resources to start fast and scale, with things like integrated GPU support with NVIDIA and CUDA runtimes.

## **Develop best practices**

I want to use the same deployment model for all my applications, integrate regular apps with AI/ML models. Also control which data I am uploading to the cloud.



## AI for the Open Hybrid Cloud

Train, serve, monitor and manage the lifecycle of AI/ML models and applications, from experiments to production



### Red Hat OpenShift AI

Expands upon the proven capabilities of Red Hat OpenShift and Red Hat OpenShift Data Science, to:

- ▶ Provide a unified platform for data scientists, application developers and IT Ops
- ▶ Scale to handle workload demands of foundation models (volume of data, duration of training run, size of model, acceleration required, and scalability).
- ▶ Deliver consistency, ease-of-use, and cloud-to-edge deployment options.
- ▶ Underlying platform for training, serving and tuning foundation models for IBM watsonx.ai and Red Hat Ansible Lightspeed with Watson Code Assistant

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](https://twitter.com/RedHat)

# Resources

# Resources

[Insights Advisor for OpenShift](#) - Red Hat Hybrid Cloud Console  
[Red Hat Insights Upgrade Risks for Red Hat OpenShift](#) - Red Hat Blog

[Learn about OpenShift Sandboxed Containers](#) - Red Hat Blog  
[Kata Containers Project](#)

[Migrate your VMs to OpenShift Virtualization using Migration Toolkit for Virtualization 2.4](#) - Red Hat Blog

[Hosted Control Planes](#)

[Nested OpenShift using OpenShift Virtualization](#) - Red Hat Blog