INDITEX Talks

10 December 2025

Red Hat OpenShift Virtualization BootCamp

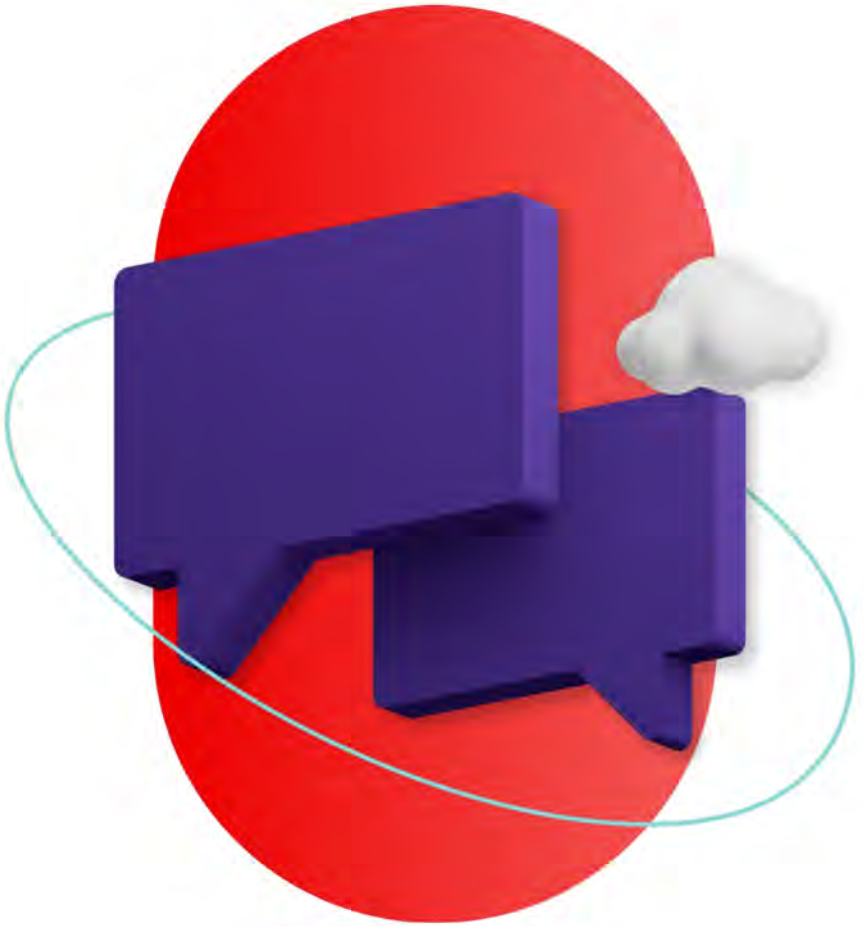

Daniel Dominguez
Specialist Solution Architect

Manuel Valle
Specialist Solution Architect

Red Hat | INDITEX

Image Resources

01
Sesión 1 | Presencial
10 de diciembre, 2025 | 10:00 - 14:00
OpenShift Virtualization: Consolidando Entornos

02
Sesión 2 | Virtual
11 de febrero, 2026 | 10:00 - 14:00
Poniendo en práctica OpenShift Virtualization

03
Sesión 3 | Presencial
24 de marzo, 2026 | 10:00 - 14:00
Ansible Automation Platform 2.6, una nueva experiencia

04
Sesión 4 | Virtual
TBC | 10:00 - 14:00
Tejiendo Patrones de Automatización: Laboratorio Práctico con Ansible:

What we'll discuss today

- Welcome and introduction
  - 10:00 - 10:30
- High Level Overview: OpenShift Virtualization and Live Demo
  - 10:30 - 11:45
- Break for coffee and checking emails
  - 11:45 - 12:15
- OpenShift Virt in Detail: Networking, Storage, Backup, DR
  - 12:15 - 14:00

Introduction

By the end of of the day...

- Consider an alternative solution for your virtual infrastructure
- Become familiar with Red Hat OpenShift Virtualization
- Understand that your migration and infrastructure modernization journey starts here
- Familiarize yourself with resources to support you along your journey

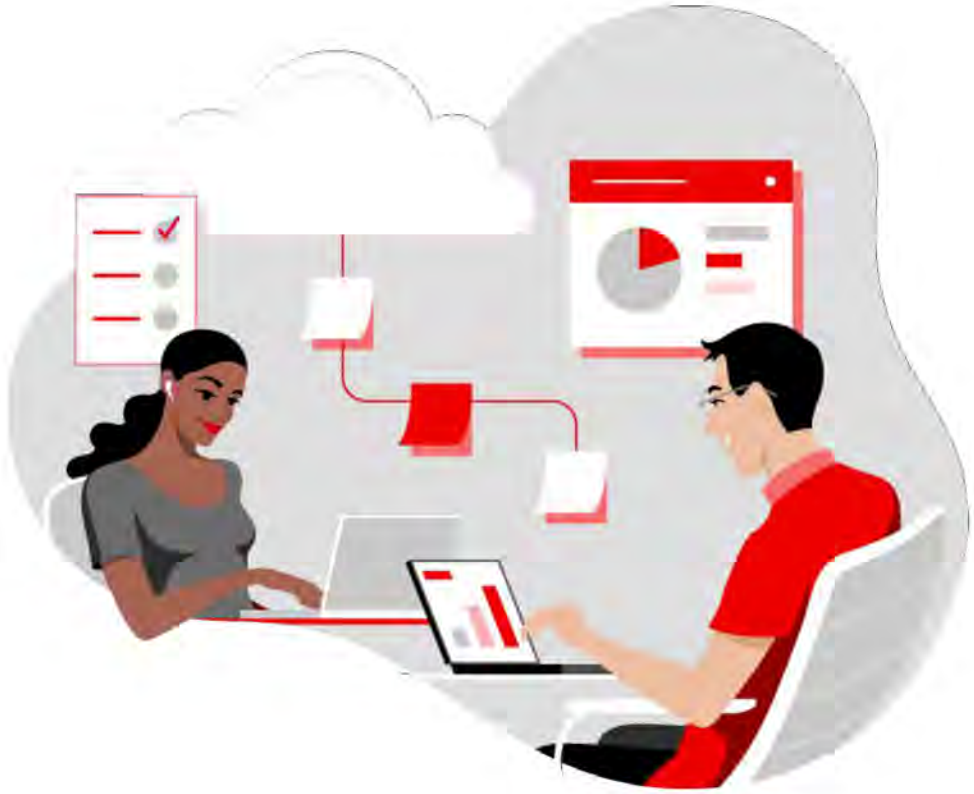

Virtualization @ Red Hat

Open Source driving KVM, RHEL, OpenStack and now OpenShift Virtualization

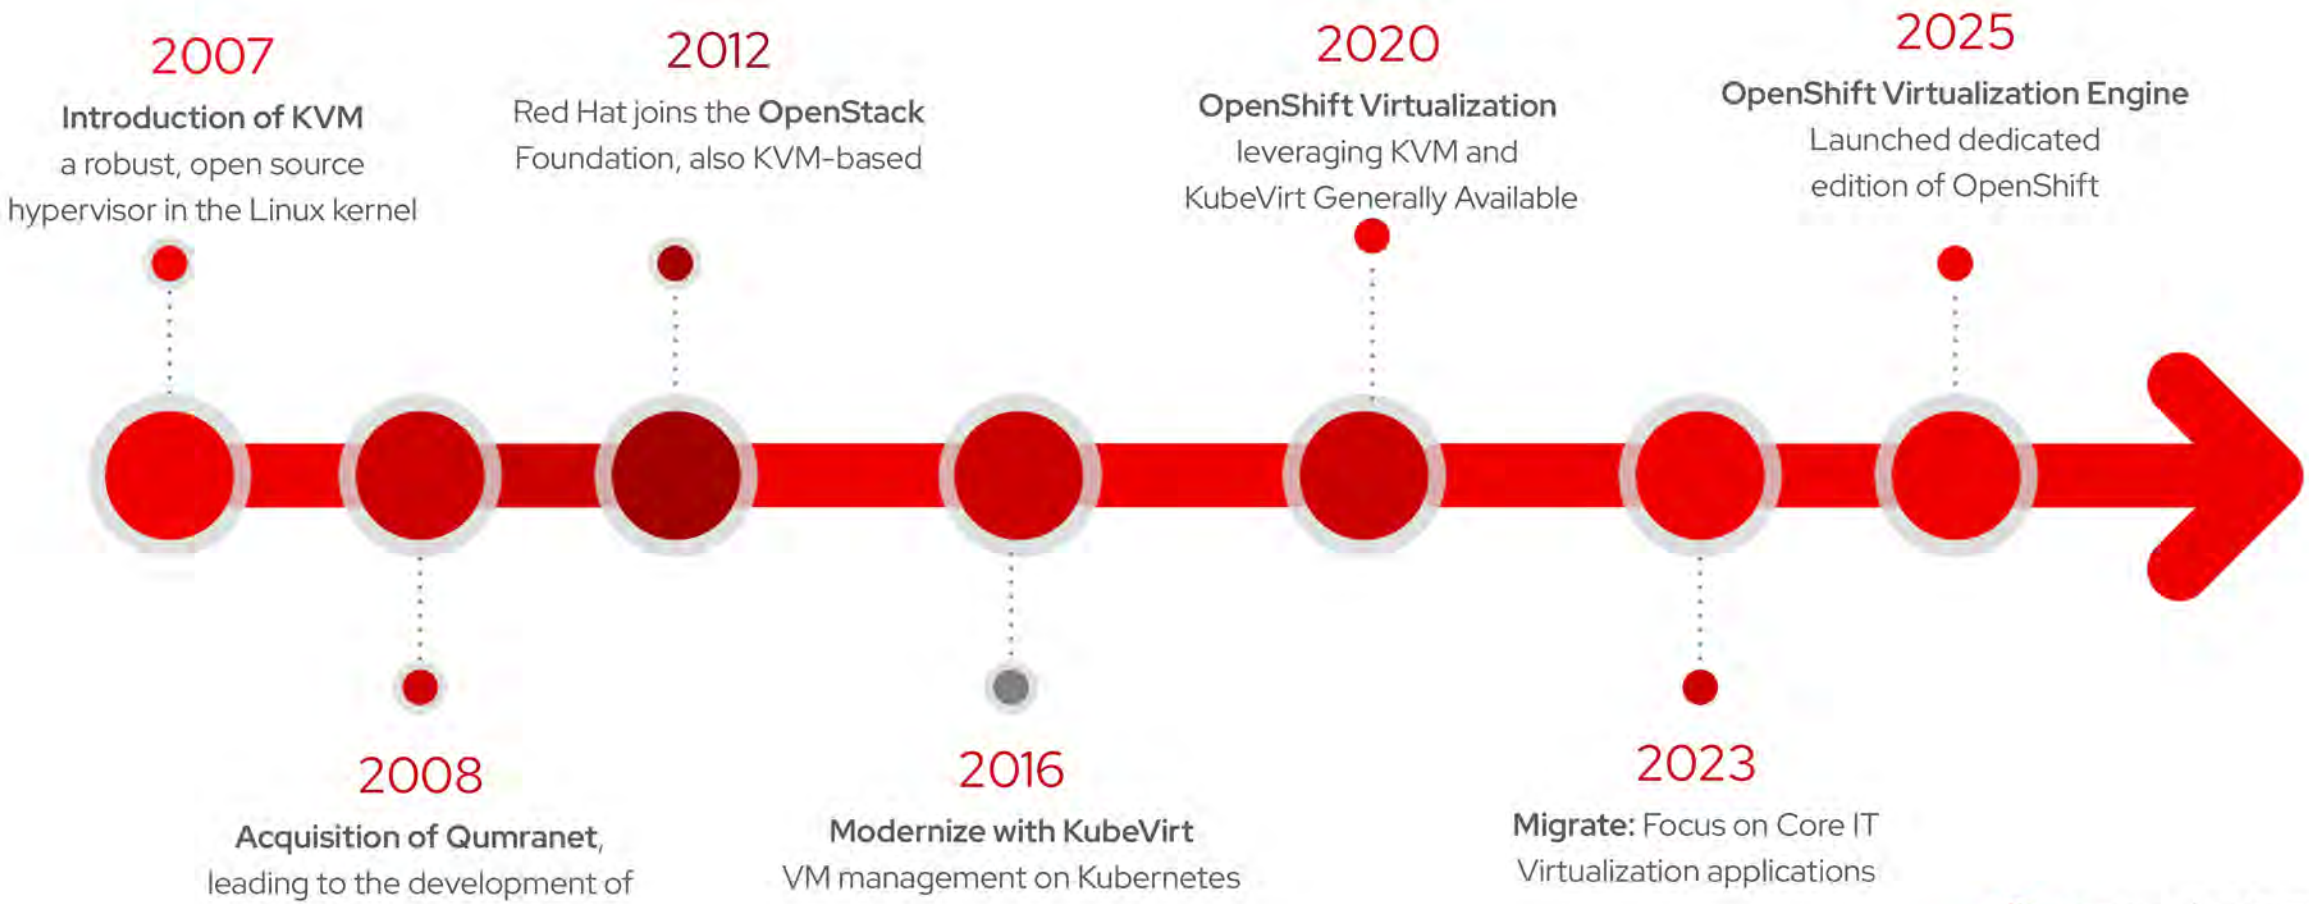

Red Hat | INDITEX

OpenShift Virtualization

Modernize workloads and support mixed applications

Start bringing VMs to OpenShift now

Support Linux and Windows apps and services in OpenShift as virtual machines with native Kubernetes tools and the security of the Red Hat platform

Deliver mixed applications on one platform

Add VM-based services such as databases to new and existing applications consisting of VMs, containers, and serverless

Modernize VMs to containers over time, or not

Refactor VMs to containerized services, or maintain as VMs. Your choice.

Red Hat OpenShift Virtualization

The modern option for general purpose virtualization

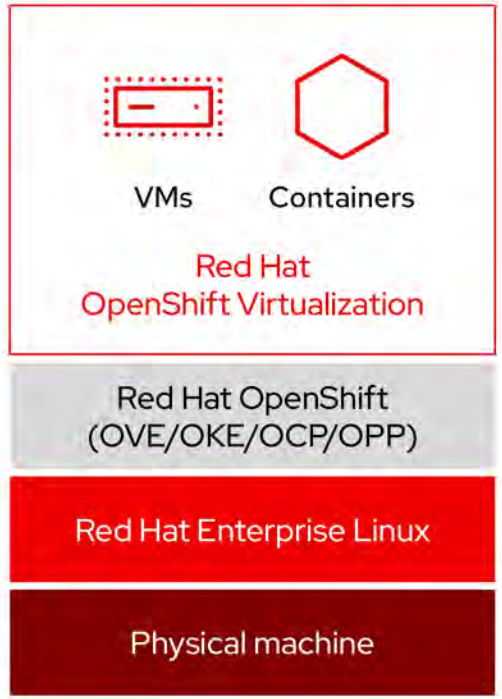

- **Unified platform**
  for virtual machines and containers*
- **Consistent management**
  tools, interfaces, and APIs incl. ACM and AAP integrations
- **Performance and stability**
  of Linux, KVM, and qemu
- **Healthy open source community**
  the KubeVirt project is a top 10 CNCF active project, with 200+ contributing companies
- **Diverse ecosystem**
  of Red Hat & partner operators
- **Included feature**
  of all OpenShift subscriptions (OVE/OKE/OCP/OPP)
- **Includes Red Hat Enterprise Linux**
  guest entitlements*
- **Supports Microsoft Windows**
  guests through Microsoft SVVP
- **Inbound guest migration**
  using Ansible Automation Platform + Migration Toolkit for Virtualization, Training and Consulting
- **Virt admin focused training**
  DO316, EX316

*excluding OVE which is virtual machines only and includes no RHEL guest entitlements

Migrate your traditional virtual machines

What we hear from customers ...

"I want to modernize"

- Wants to modernize to containers, but also run VMs in a more modern way
- Stand up a secondary virtualization platform for new workloads
- Legacy and next-gen virtualization platforms co-exist

"I want to migrate"

- Migrate off their current traditional virtualization platform completely, as quickly and as safely as possible
- Modernization is subordinate to migration; containers and Kubernetes are implementation details
- Willing to take calculated risk with their production workloads

Options for VMs with Red Hat in the datacenter

Building virtualized environments however our customers want them

Red Hat OpenShift Virtualization Engine

A virtualization-only platform that simplifies VM migration and offers a cost-effective alternative for organizations seeking to run and manage VMs

Red Hat OpenShift

OpenShift Virtualization is functionality found within all OpenShift editions, enabling the management of VMs in a hybrid cloud environment

OpenShift Virtualization Engine

Opening the door to virtualization and modernization

Unlimited VMs

Run as many VMs as you need, maximizing the value of your hardware. Purchase RHEL subscriptions, virtualized OpenShift for container-based applications, or upgrade to other bare metal OpenShift editions if needed.

128 core bare metal scale

Get bare metal scale with 128 cores per subscription - run more VMs on less hardware, optimizing your infrastructure efficiency.

Optional Advanced Cluster Management for Virtualization

Scale as big as you can; add Advanced Cluster Management for Virtualization to make management of thousands of nodes as easy as managing a single rack.

Workload monitoring and platform logging

Keep tabs on and track your environment with a preconfigured, preinstalled, and self-updating stack then stay in command with the included OpenShift GitOps operator to leverage Kubernetes-powered orchestration for VMs.

OpenShift Virtualization is available across major cloud providers

Migrate VMs to the cloud

Modernize and build new apps with cloud-native tech

Generally Available on Red Hat OpenShift Service on AWS and self-managed OpenShift on AWS

Public Preview of Azure Red Hat OpenShift

Technical Preview on OpenShift Dedicated and self-managed OpenShift on Google Cloud Platform

Generally Available on Red Hat OpenShift on IBM Cloud

Technical Preview on OpenShift on OCI

Technical Overview

Virtual Machines are not containers!

- Typically containers run a single app process
- Containers don’t mind to be terminated and restarted
- Containers are often running as multiple replicas
- Containers don’t need a hypervisor (just a host)
- There is nothing unusual about being ephemeral and immutable when you’re a container

But Virtual Machine can be run INSIDE a container

- KVM virtual machine is a process
- Containers are meant to encapsulate processes
- Both have the same underlying resource needs
  - CPU, Memory, Storage, Network, etc
- Orchestration (K8s) layer is transparent for VM
- KVM has 10+ years of production use

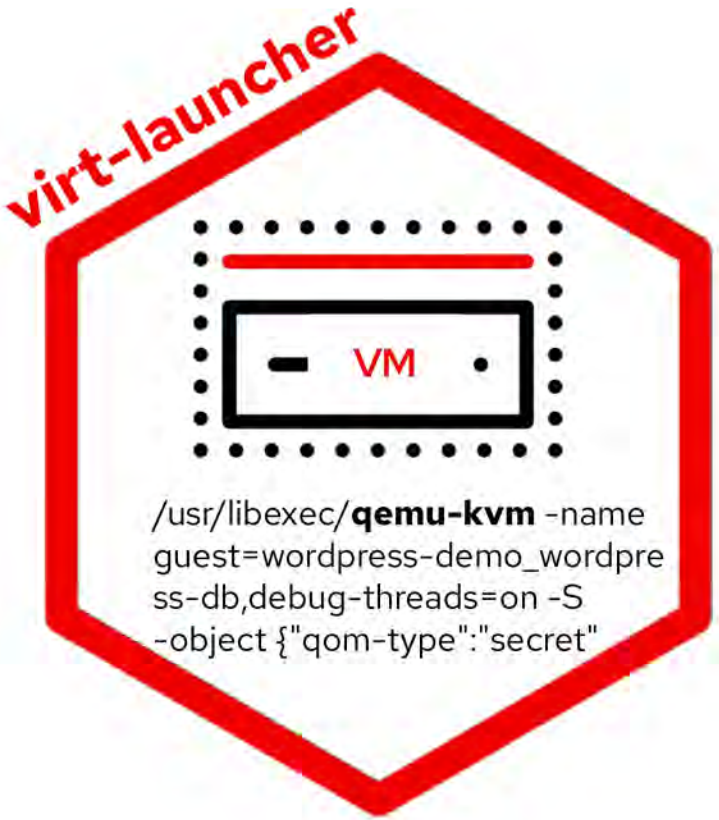

Containerizing KVM

Trusted, mature KVM wrapped in modern management and automation

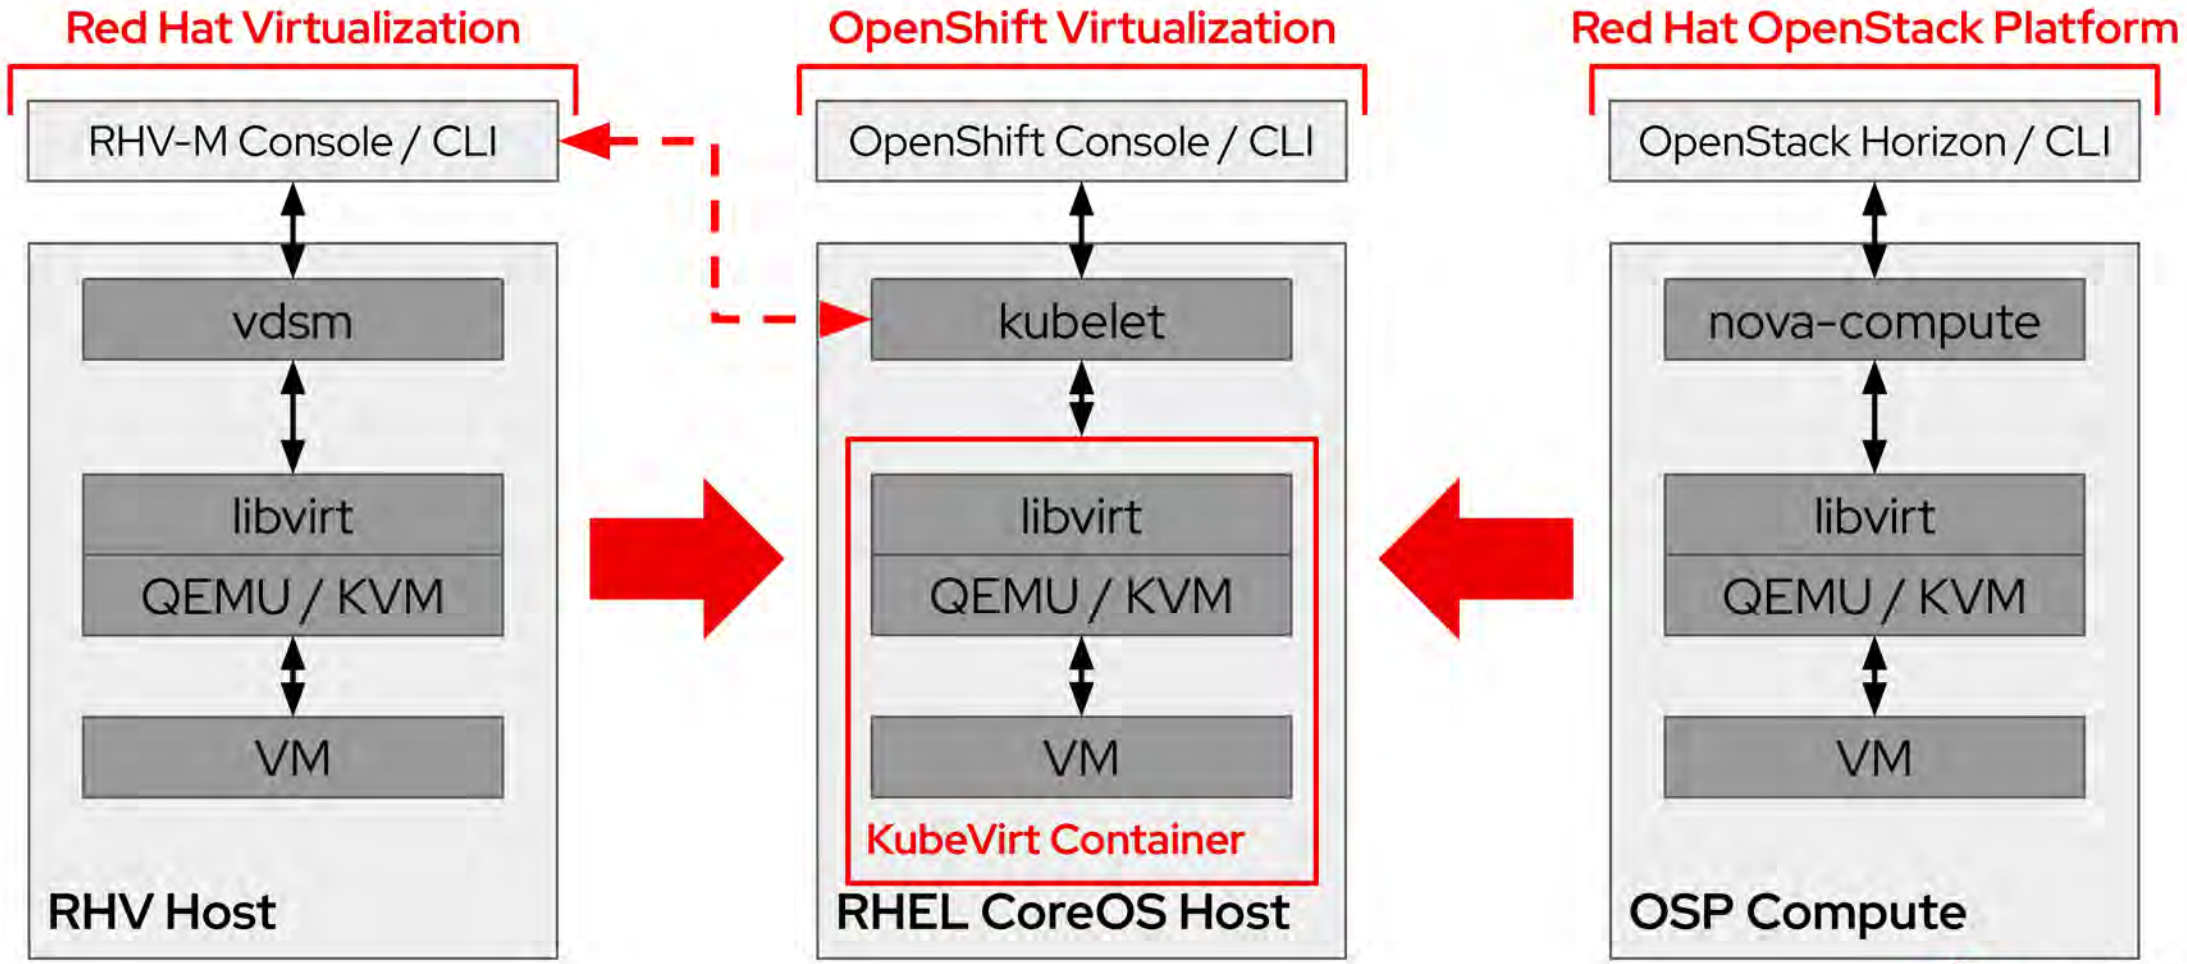

Using VMs and containers together

- Virtual machines connected to pod networks are accessible using standard Kubernetes methods:
  - Service
  - Route
  - Ingress
- 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

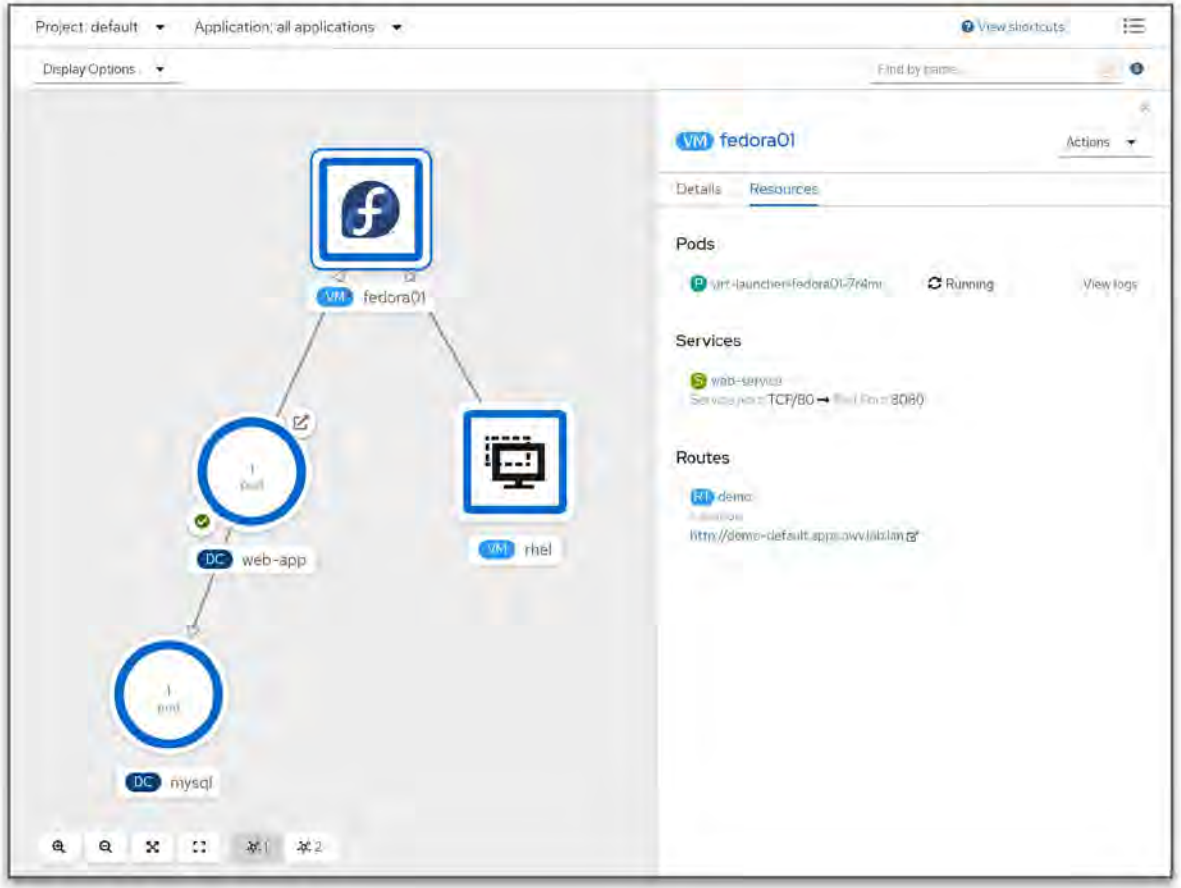

Red Hat | INDITEX

Advantages of running VMs in OpenShift

Networking

- integrated SDN & VM secondary networks
- network policies for segregation and security
- integrated DNS/IPAM
- flexible Load Balancer

Platform support

- live migration & HA
- live changes of disks and NICs
- CPU & memory overcommit
- access control & tenant separation

Automation

- integrated automation & 100% API
- integrated health checks
- storage provisioning
- multi-cluster & hybrid cloud

Application & Management support

- metrics and observability
- Service Mesh
- cost management
- tenant separation & AD/LDAP/oauth

... modern way of systems and applications management & huge ecosystem

Featuritis: 15+ years of enterprise virtualization experience

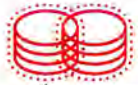

Enterprise virtualization capabilities

- **Live migration**
- Infrastructure fencing
- Application fencing
- **Hot pluggable disks**
- Non-disruptive upgrades
- Deduplicating memory with freePageReporting and KSM support
- **Performance, Scale, Stability of KVM**
- **CPU/MEM overcommit**
- VM disk resize
- **GPU passthrough, vGPU**
- UI for VM admins
- **VM export**
- **Microsoft Windows,** UEFI and Secure boot, Persistent vTPM
- Microsoft Windows Server Failover Cluster (WSFC)

Networking

- DPDK, **SRI-IOV,** IPv6
- Flat L2 Overlay network without the need to configure host networking
- **OVN Kubernetes** localnet as an alternative to the Bridge CNI
- Secondary network ipBlock policies and **micro segmentation**

Hybrid Cloud and managed services

- x86 and **ARM** based systems
- Single Node OpenShift
- 3- node compact cluster
- **ROSA, AWS Public Cloud**
- IBM Cloud (Tech Preview)

Storage, backups, and DR

- **Storage profiles for all major storage providers - ODF, Netapp, Pure/Portworx, Dell, Hitachi, HPE, IBM**
- **Backup / restore** with OADP
- **Portworx support for Metro-DR and Async-DR**
- GitOps-based VMs disaster recovery in **ODF Metro-DR**
- Kasten K10 by Veeam
- Trilio TVK
- Storware vProtect

Deployment Methods

Deploying on Bare Metal - Day 1

Interactive

★ Recommended | Web-based

Runs Assisted Installer with standard configuration settings to create your cluster.

PROS

- Easiest & Quickest
- Pre-deployment validations
- Best for customer driven PoCs

CONS

- Requires internet connection

Docs

Local Agent-based

CLI-based

Runs Assisted Installer securely and locally to create your cluster.

PROS

- Best of "both worlds" (A-I/IPI)
- Easiest to move to source control
- Best for disconnected

CONS

- Manual YAML file creation/editing

Docs

Automated

CLI-based

Auto-provision your infrastructure with minimal configuration to create your cluster.

PROS

- Great for repeat and large scale

CONS

- Requires Redfish compatibility or additional networking

Full control

CLI-based

Make all of the decisions when you create your cluster.

PROS

- Most customizable of all options
- Most versatile for underlying infrastructure

CONS

- Manual "everything" deployment

Complexity Increases

OpenShift 4.20 Supported Providers

Installation Experiences

aws
Outposts
Wavelength
Local Zones

Azure
Azure Stack Hub
Alibaba Cloud (Tech Preview)

IBM Cloud
IBM Power Systems
IBM Z and
IBM LinuxONE

Bare Metal

ORACLE
Cloud Infrastructure

Google Cloud

NUTANIX

vmware vSphere

Red Hat
OpenStack Services
on OpenShift

RED HAT
OPENSTACK
PLATFORM

Automated

Installer Provisioned Infrastructure

- Auto-provisions infrastructure
- *KS like
- Enables self-service

Full Control

User Provisioned Infrastructure

- Bring your own hosts
- You choose infrastructure automation
- Full flexibility
- Integrate ISV solutions

Interactive – Connected

Assisted Installer

- Hosted web-based guided experience
- Agnostic, bare metal, vSphere and Nutanix
- ISO driven

Local – Disconnected

Agent-based Installer

- Restricted network (disconnected / air -gapped)
- Automatable installations via CLI
- Bare metal, vSphere, SNO
- ISO driven

High-Level Bare Metal IPI

Installation Experiences

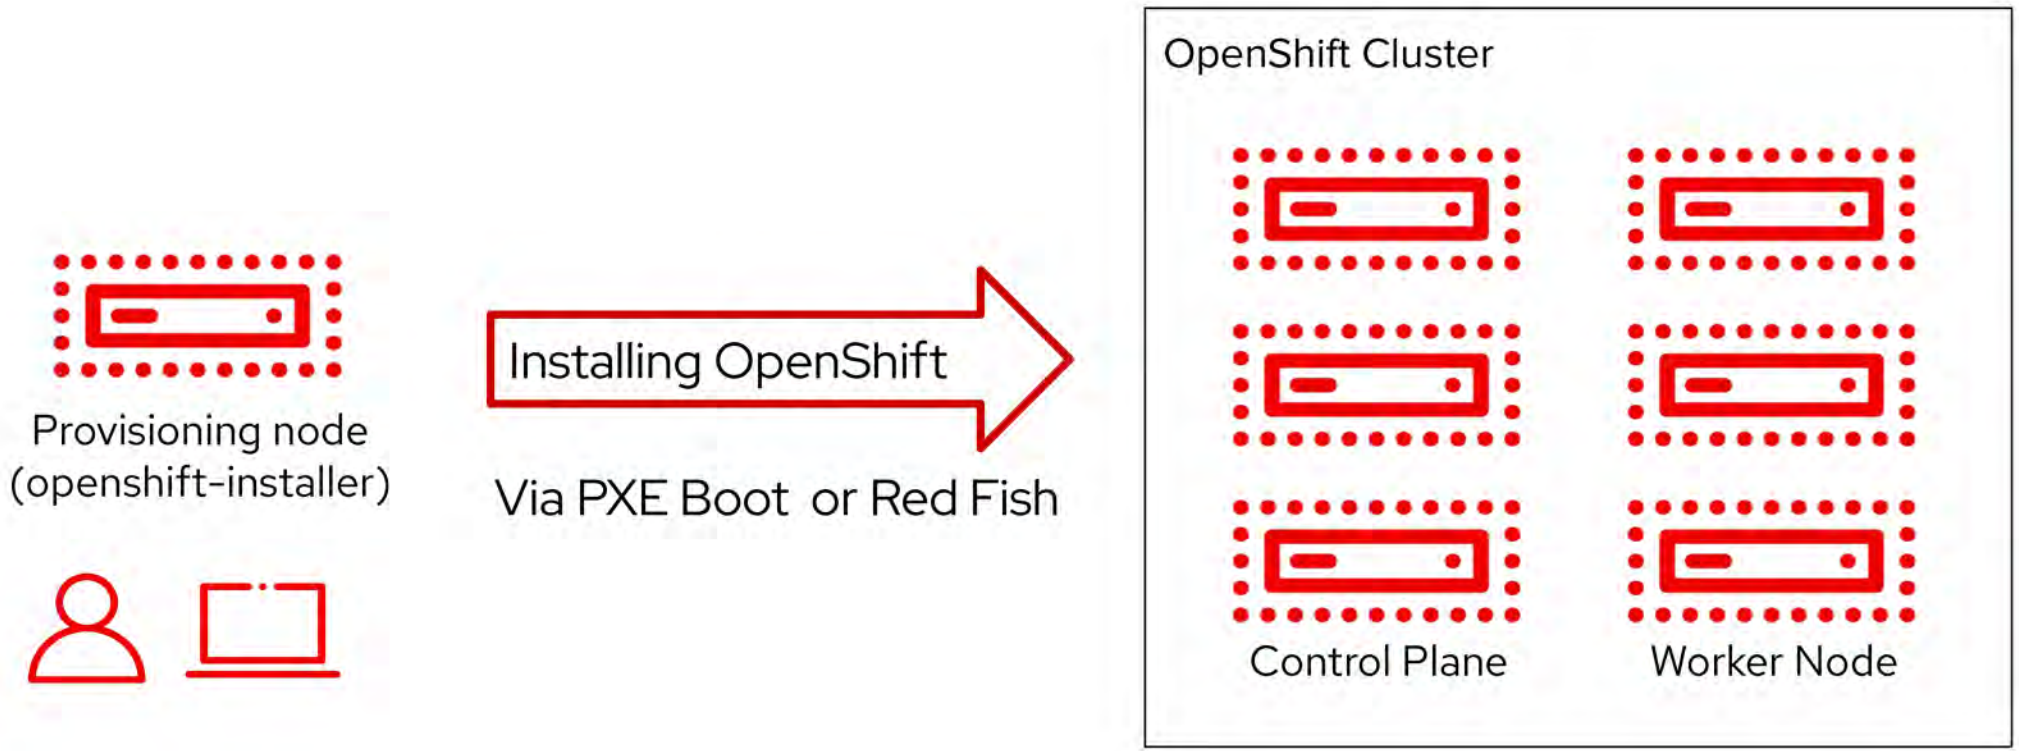

High-Level Agent-based Installer

Installation Experiences

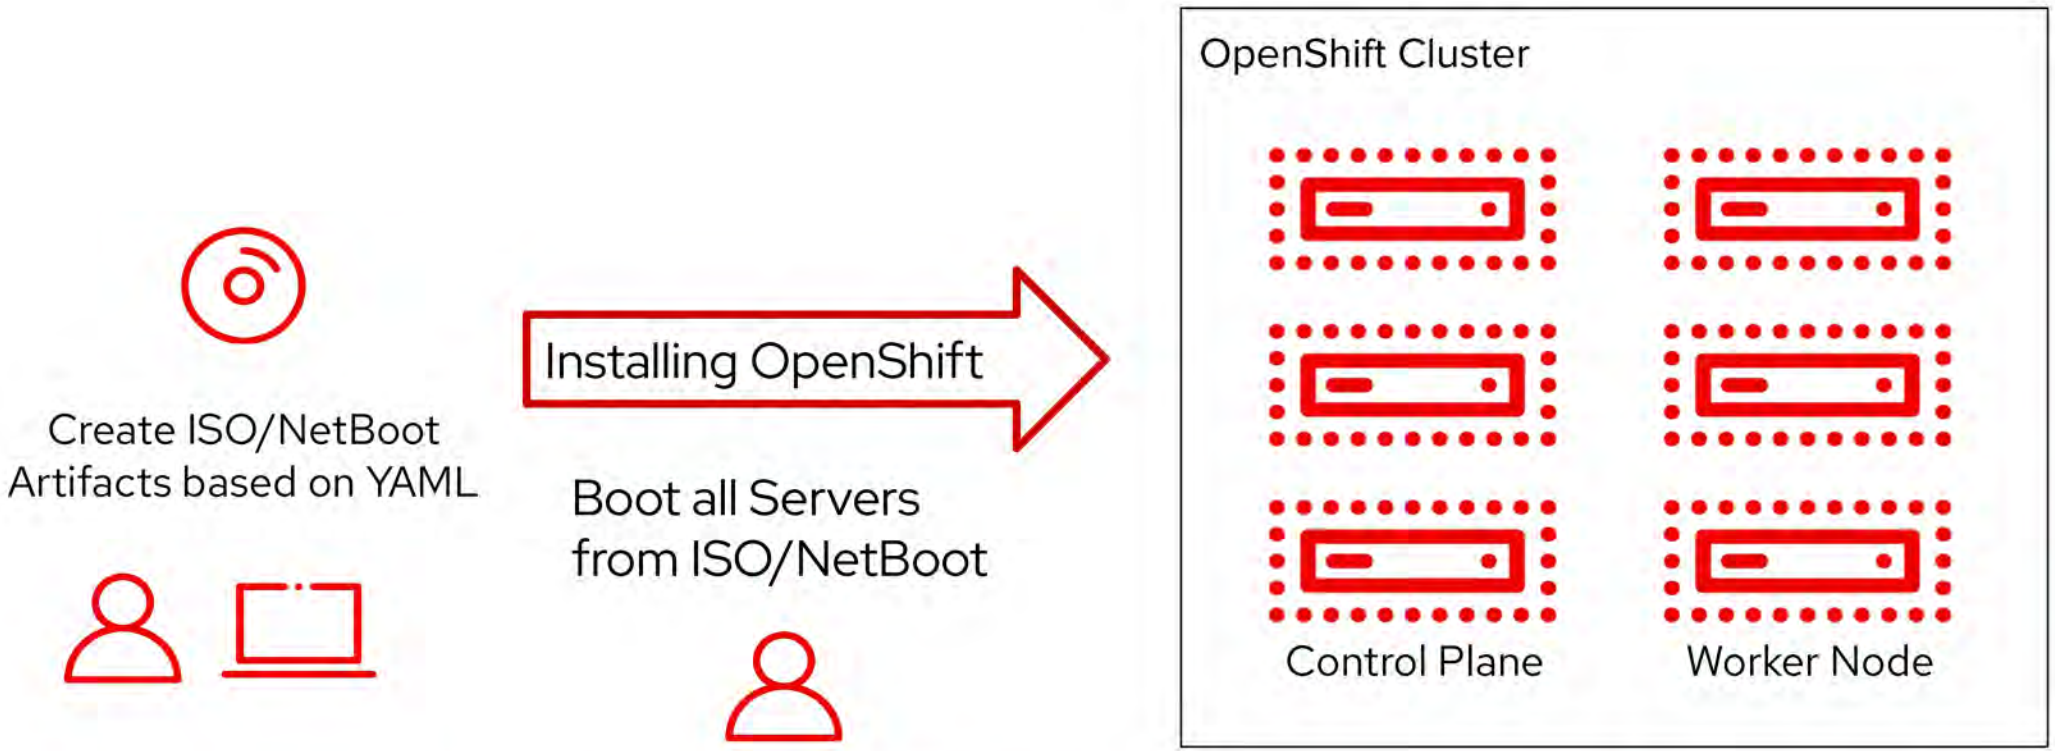

Commitment to API-First and Automation

High-Level Assisted Installer

Installation Experiences

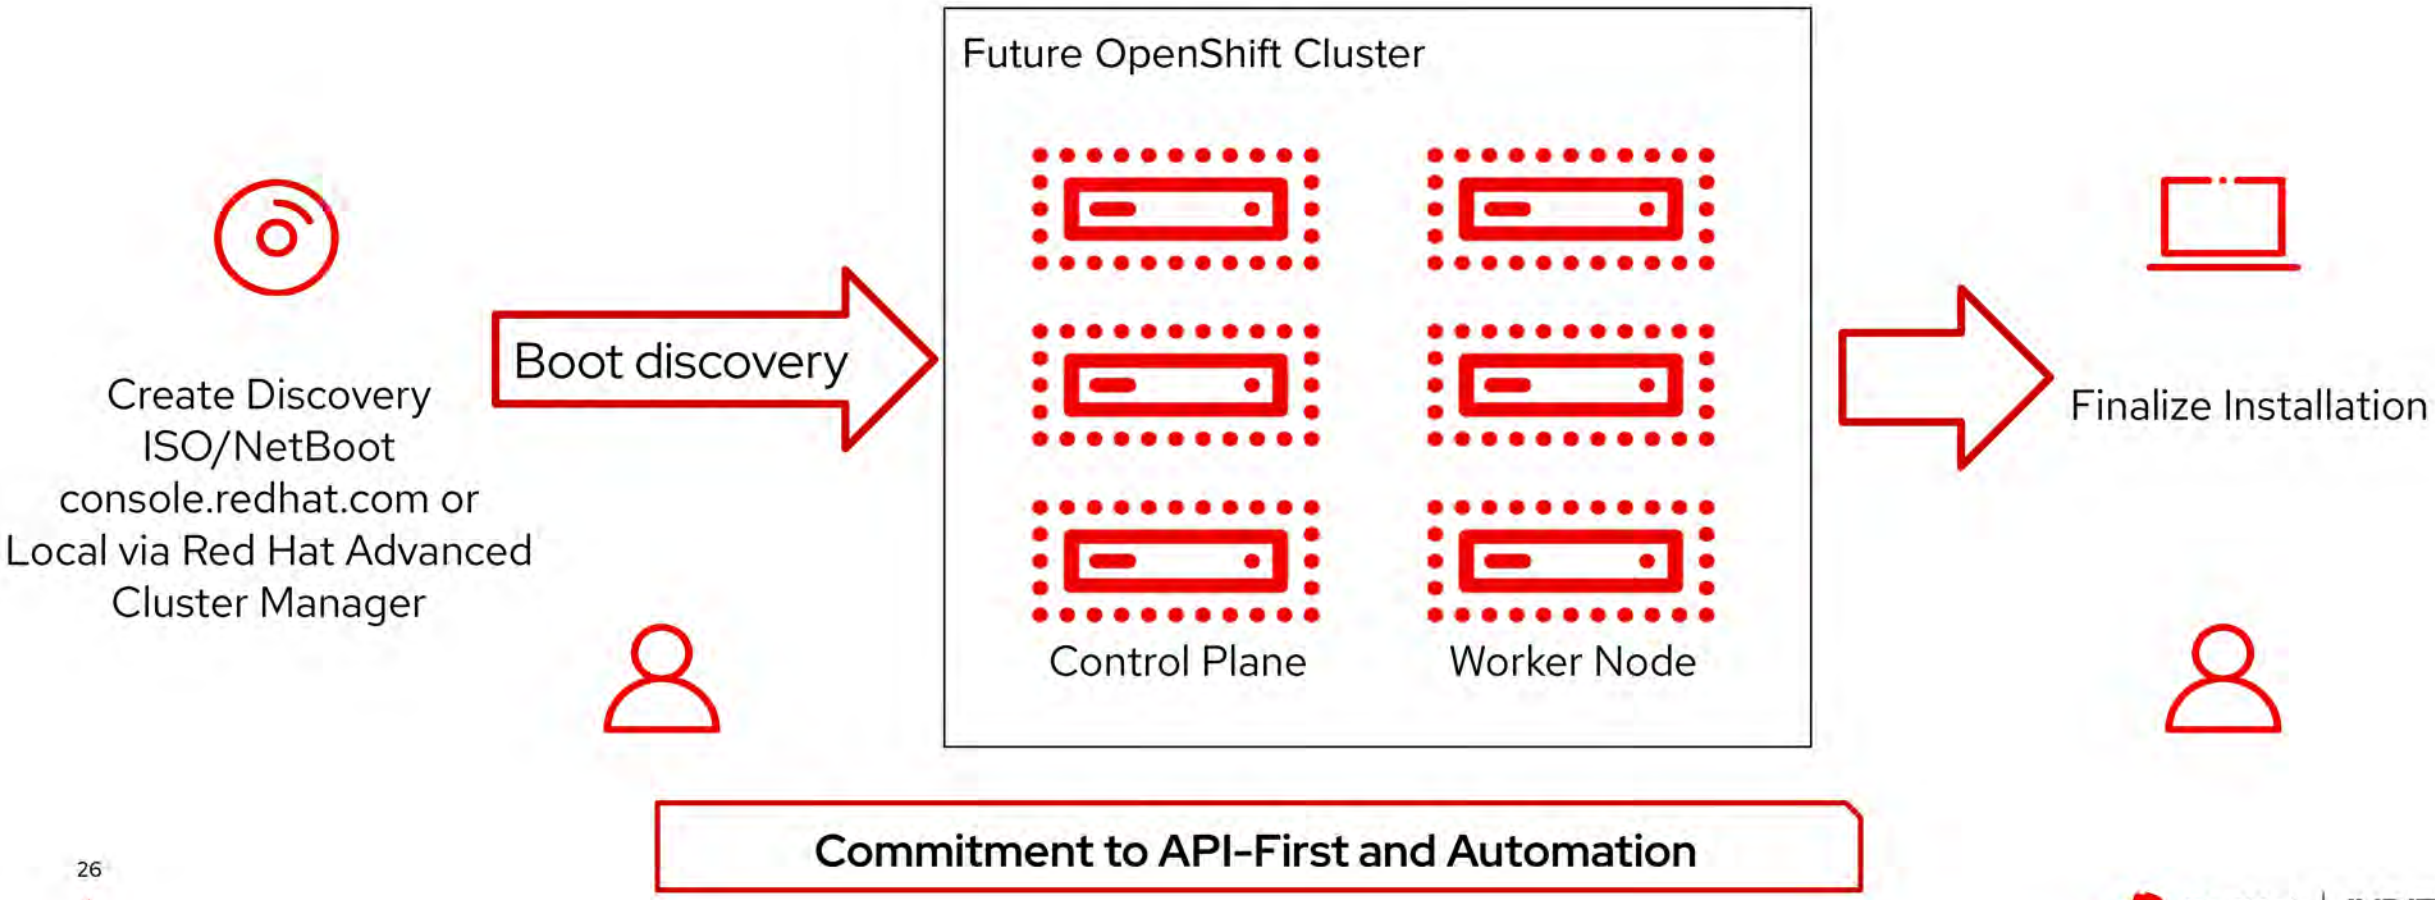

Commitment to API-First and Automation

Red Hat | INDITEX

OpenShift Operators

Operators

Operators act like an extension of the software vendor's engineering team, watching over a Kubernetes environment and using its current state to make decisions in real time

Ease the operational complexity

A cluster is managed by operators

```yaml
apiVersion:
apiextensions.k8s.io/v1beta1
kind: CustomResourceDefinition
metadata:
  name:
crontabs.stable.example.com
spec:
  group: stable.example.com
  version: v1
  scope: Namespaced
  Names:
    plural: crontabs
    singular: crontab
    kind: CronTab
    - ct
```

Operators in OpenShift

Operator Hub - Allows administrators to selectively make operators available from curated sources to users in the cluster.

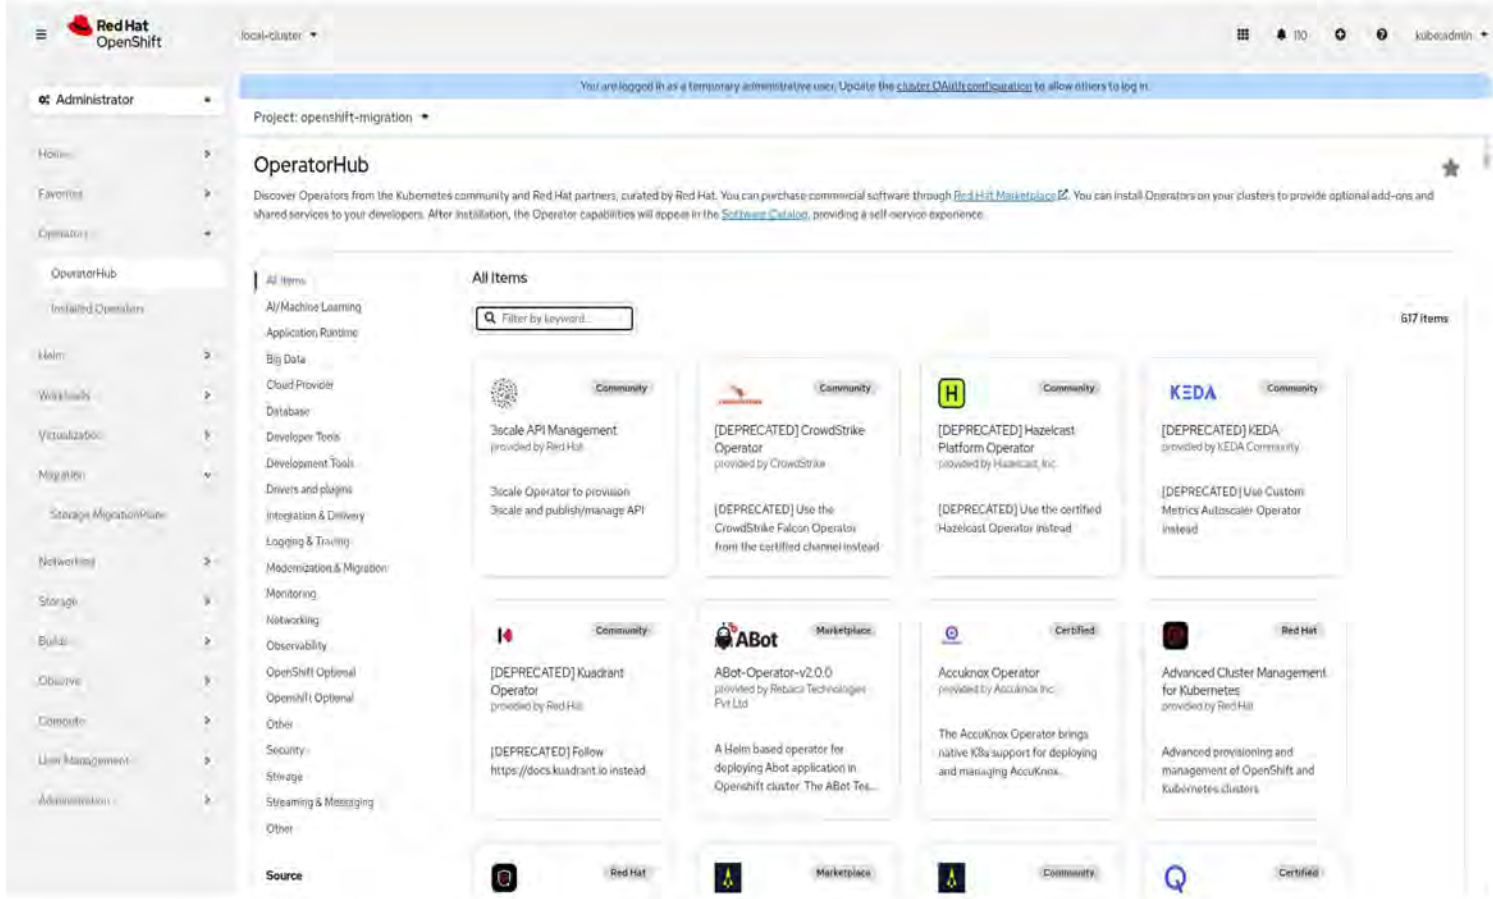

Red Hat | INDITEX

Marketplace

OperatorHub.io is designed to be the public registry for finding Kubernetes Operator backed services. It is a community project of Red Hat in collaboration with AWS, Google Cloud and Microsoft.

> The creation of OperatorHub.io provides a centralized repository that helps users and the community organize around Operators. We look forward to seeing growth and adoption of OperatorHub.io as an extension of the Kubernetes community

Aparna Sinha
Cloud Product Manager, Google Group

> Use of Kubernetes Operators is growing both inside Microsoft and amongst our customers, and we look forward to working with Red Hat and the broader community on this important technology

Gabe Monroy
Lead Program Manager, Containers, Microsoft Azure

Source: https://www.redhat.com/en/blog/introducing-operatorhubio-place-finding-kubernetes-native-services

Operators

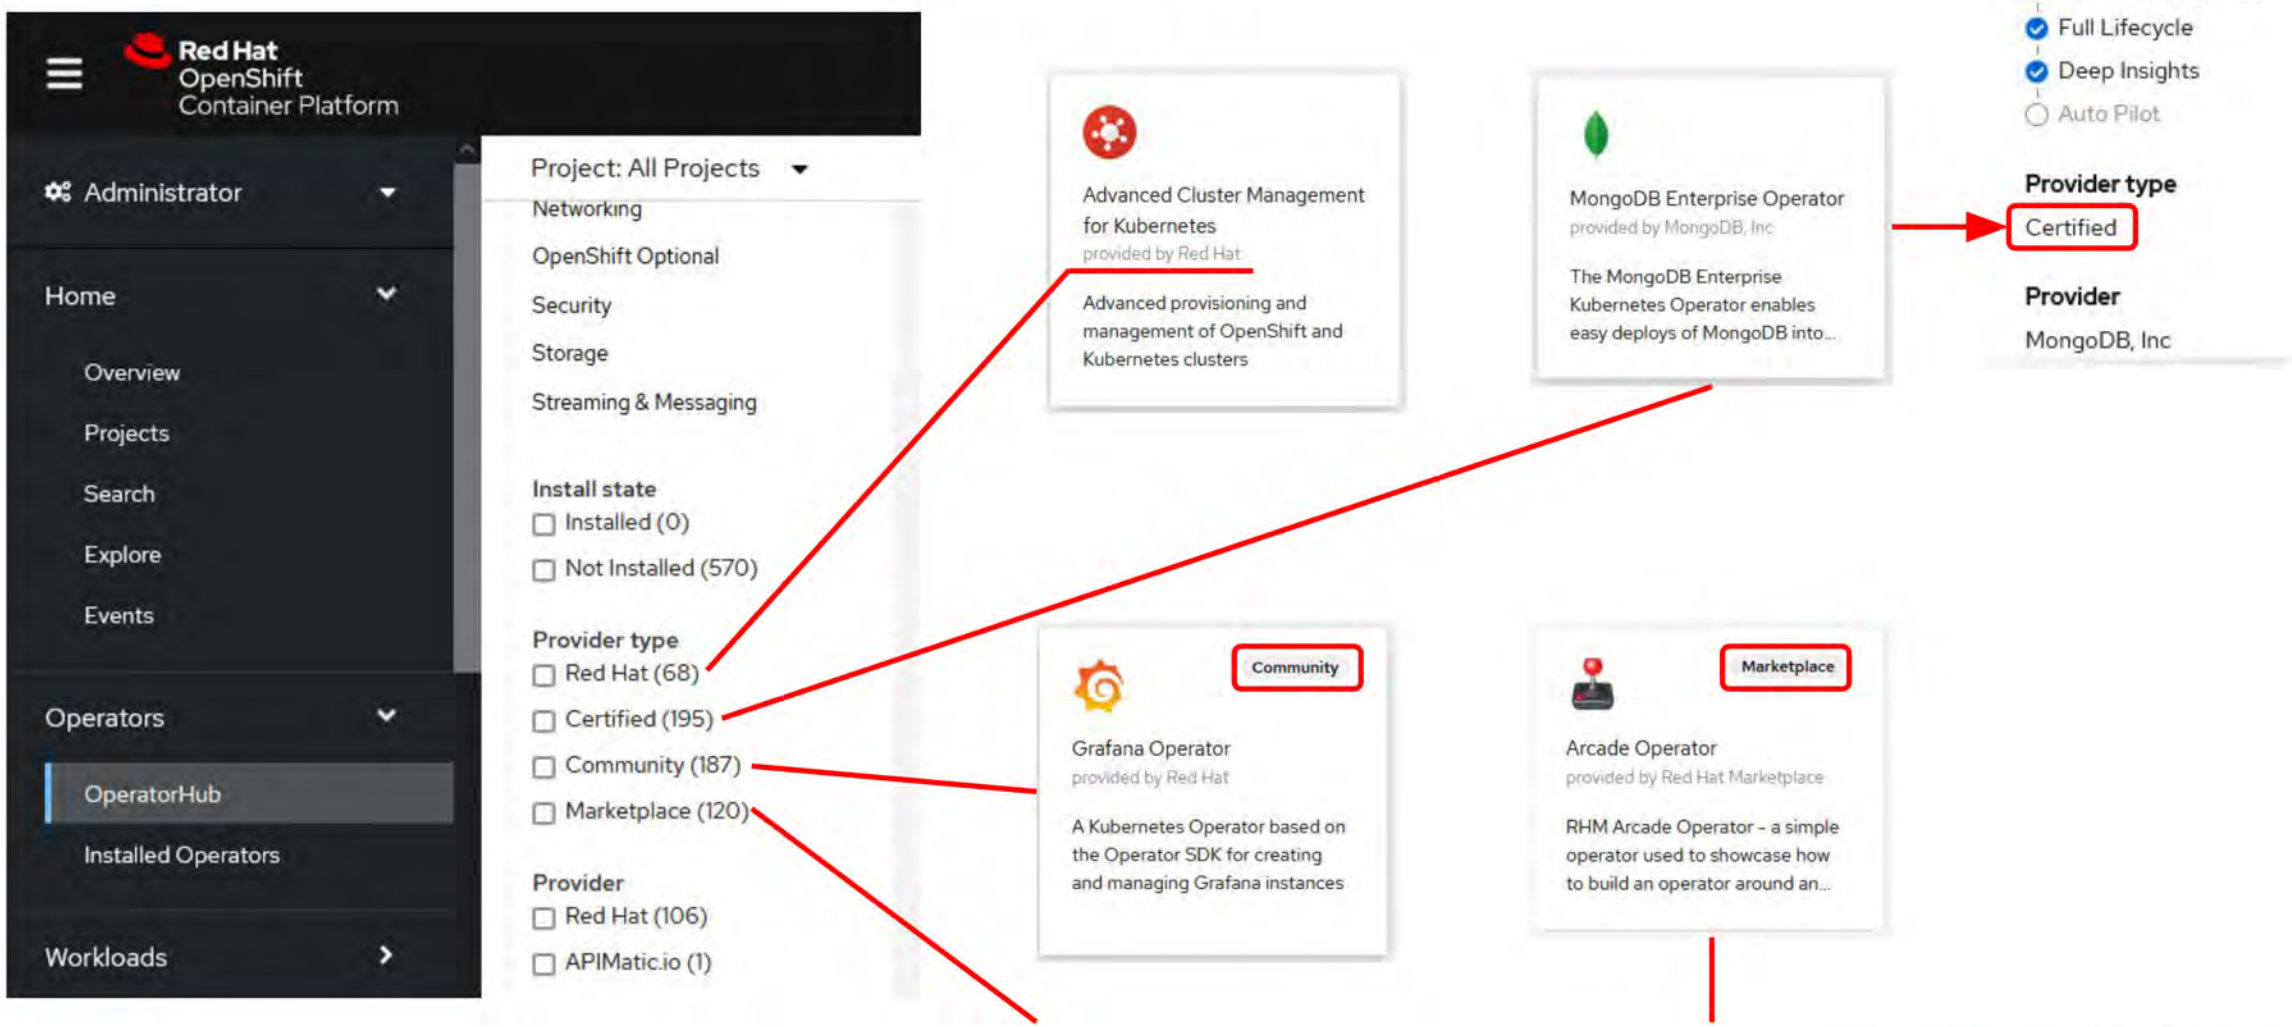
Red Hat
OpenShift
Container Platform
Administrator
Home
Overview
Projects
Search
Explore
Events
Operators
OperatorHub
Installed Operators
Workloads
Project: All Projects
Networking
OpenShift Optional
Security
Storage
Streaming & Messaging
Install state
Installed (0)
Not Installed (570)
Provider type
Red Hat (68)
Certified (195)
Community (187)
Marketplace (120)
Provider
Red Hat (106)
APIMatic.io (1)
Advanced Cluster Management for Kubernetes
provided by Red Hat
Advanced provisioning and management of OpenShift and Kubernetes clusters
MongoDB Enterprise Operator
provided by MongoDB, Inc
The MongoDB Enterprise Kubernetes Operator enables easy deploys of MongoDB into...
Community
Grafana Operator
provided by Red Hat
A Kubernetes Operator based on the Operator SDK for creating and managing Grafana instances
Marketplace
Arcade Operator
provided by Red Hat Marketplace
RHM Arcade Operator - a simple operator used to showcase how to build an operator around an...
Latest version
1.12.0
Capability level
Basic Install
Seamless Upgrades
Full Lifecycle
Deep Insights
Auto Pilot
Provider type
Certified
Provider
MongoDB, Inc

Industry behind operators

90+ Certified ISV Operators in Red Hat Program

Getting OpenShift Virtualization

OpenShift Virtualization Operator

- Operators - packaging, deploying and managing a Kubernetes native applications
- Certified and supported way to extend platform capabilities (virtualization and many more, including 3rd party components)
- Easy lifecycle management

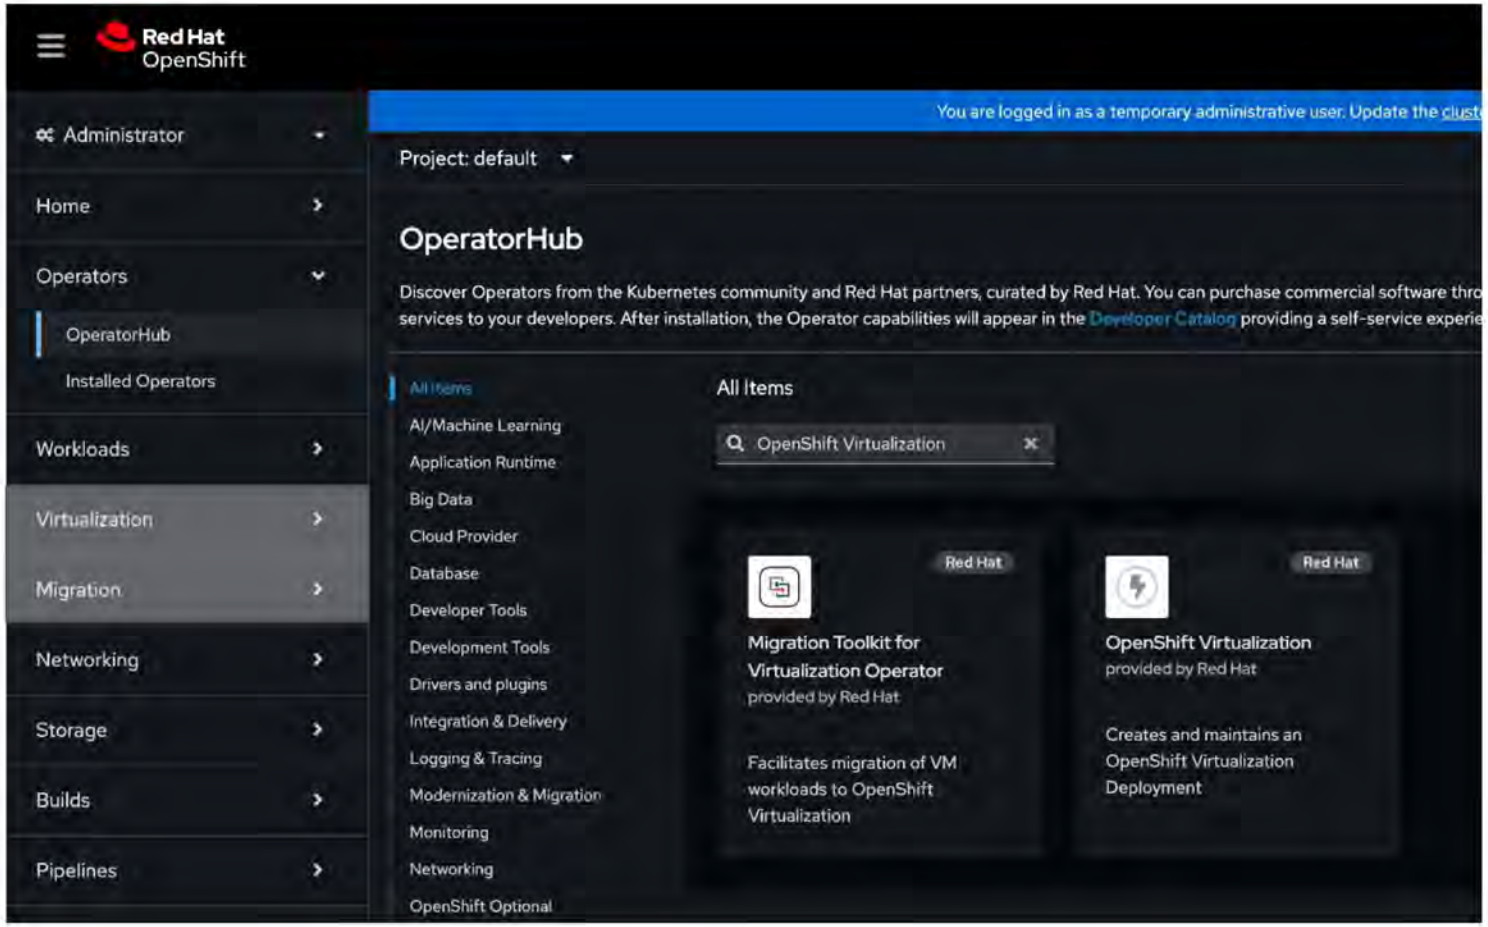

OpenShift Virtualization View

- Simplified view for virtualization admins
- The whole virtualization context in one place
- Reduces clatter made by noisy containers ;-)

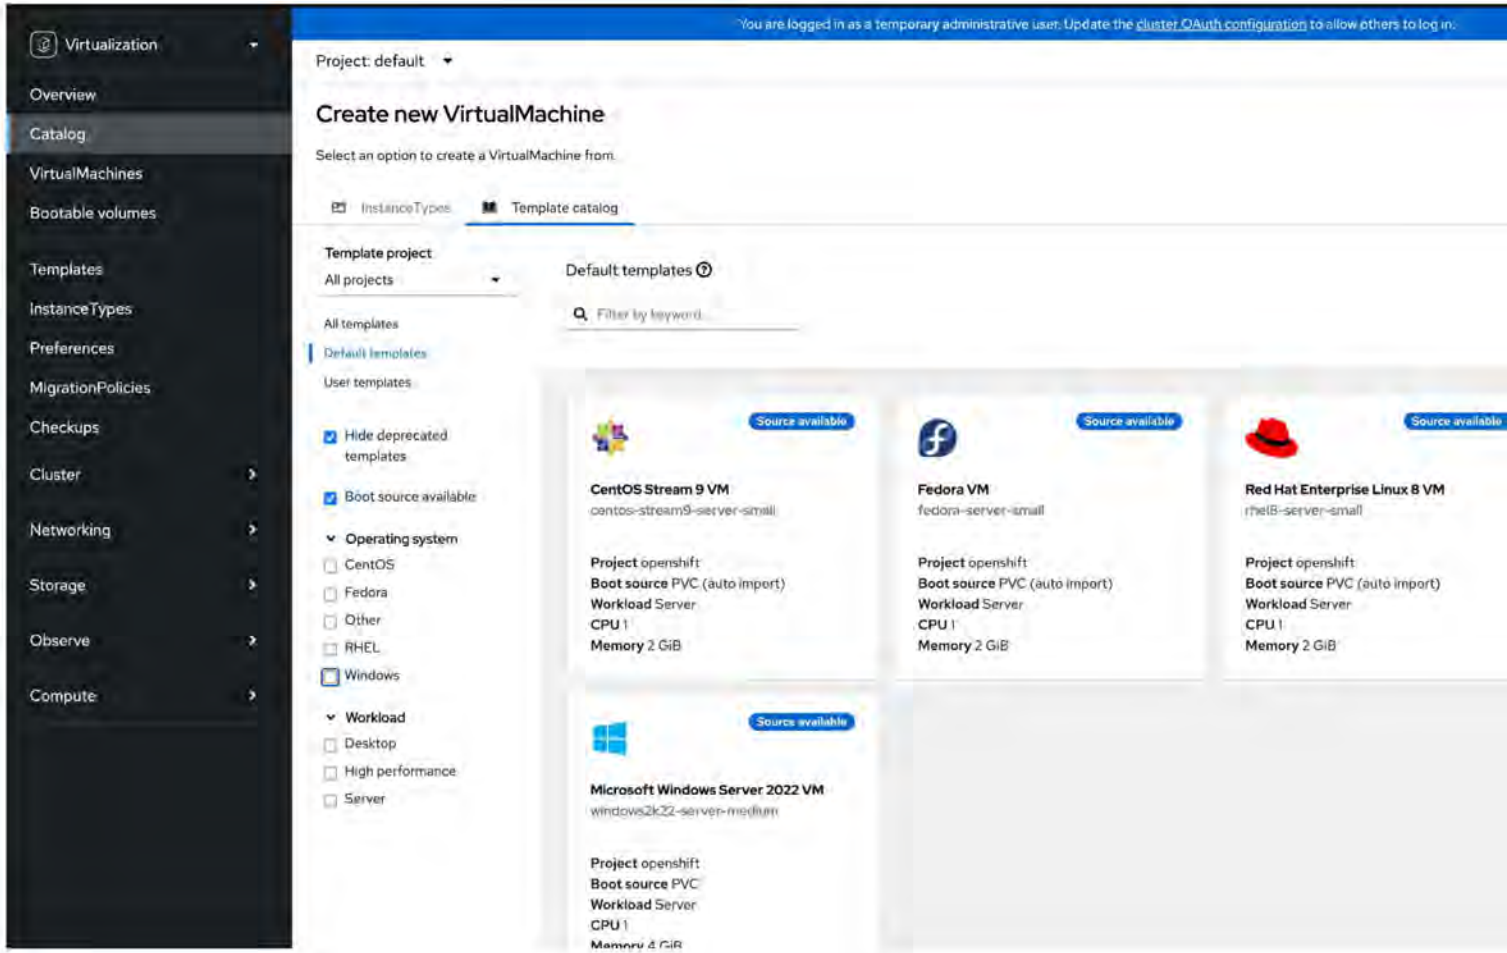

Red Hat | INDITEX

Workload mobility

Compute Live Migration

Virtual Machine Live Migration

- CPU and Memory changes
- Adding/removing devices
- Infrastructure maintenance
- Workload rebalance (descheduling)

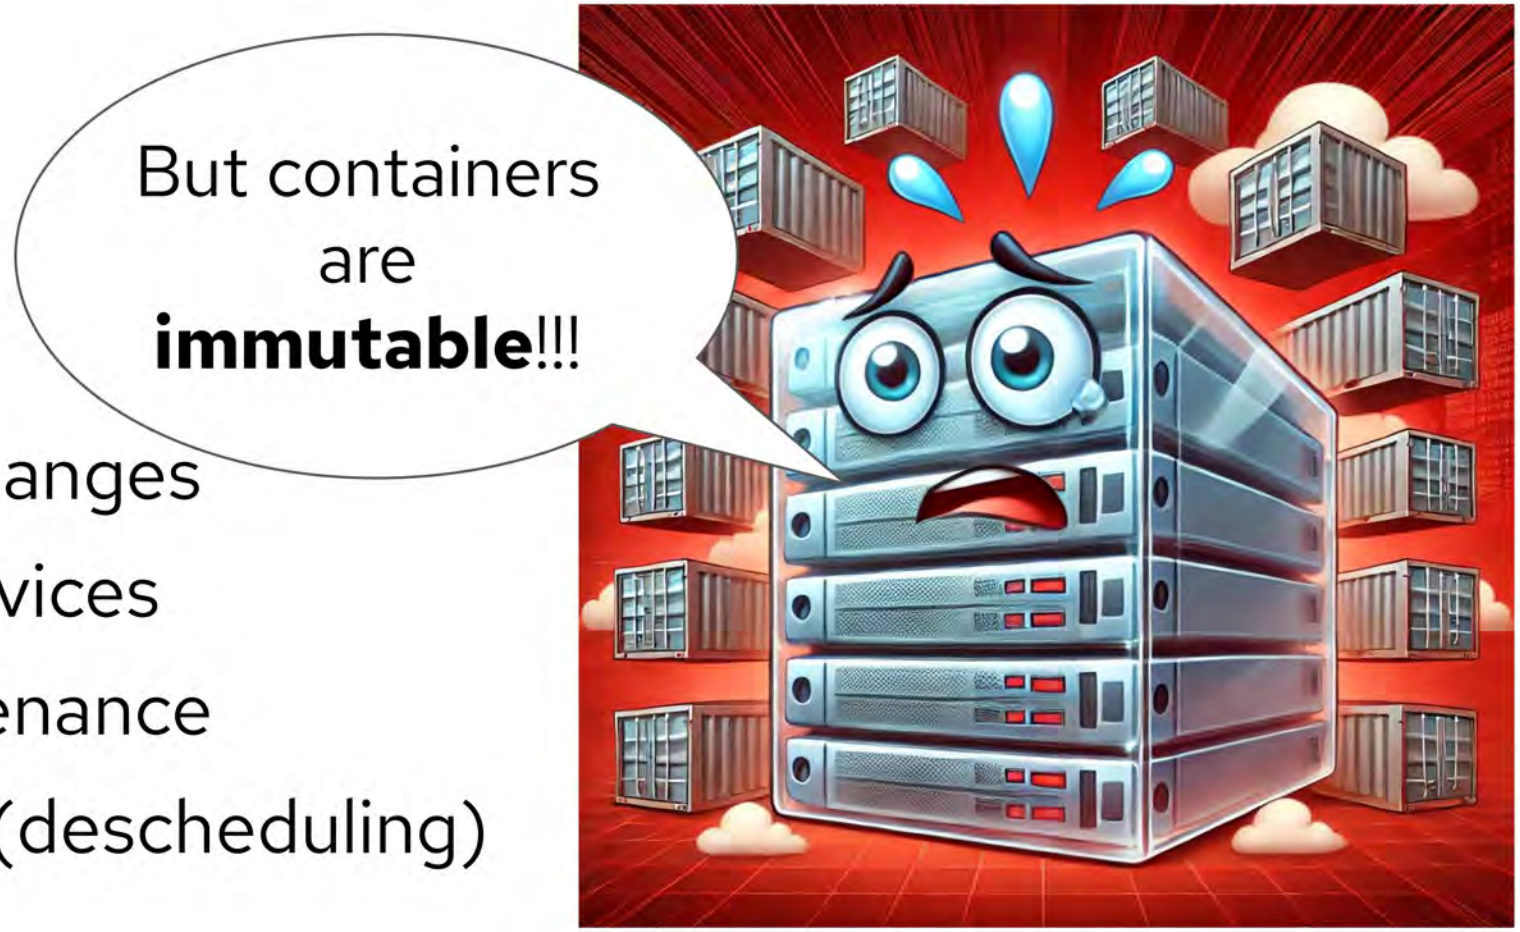

Virtual Machine Live Migration

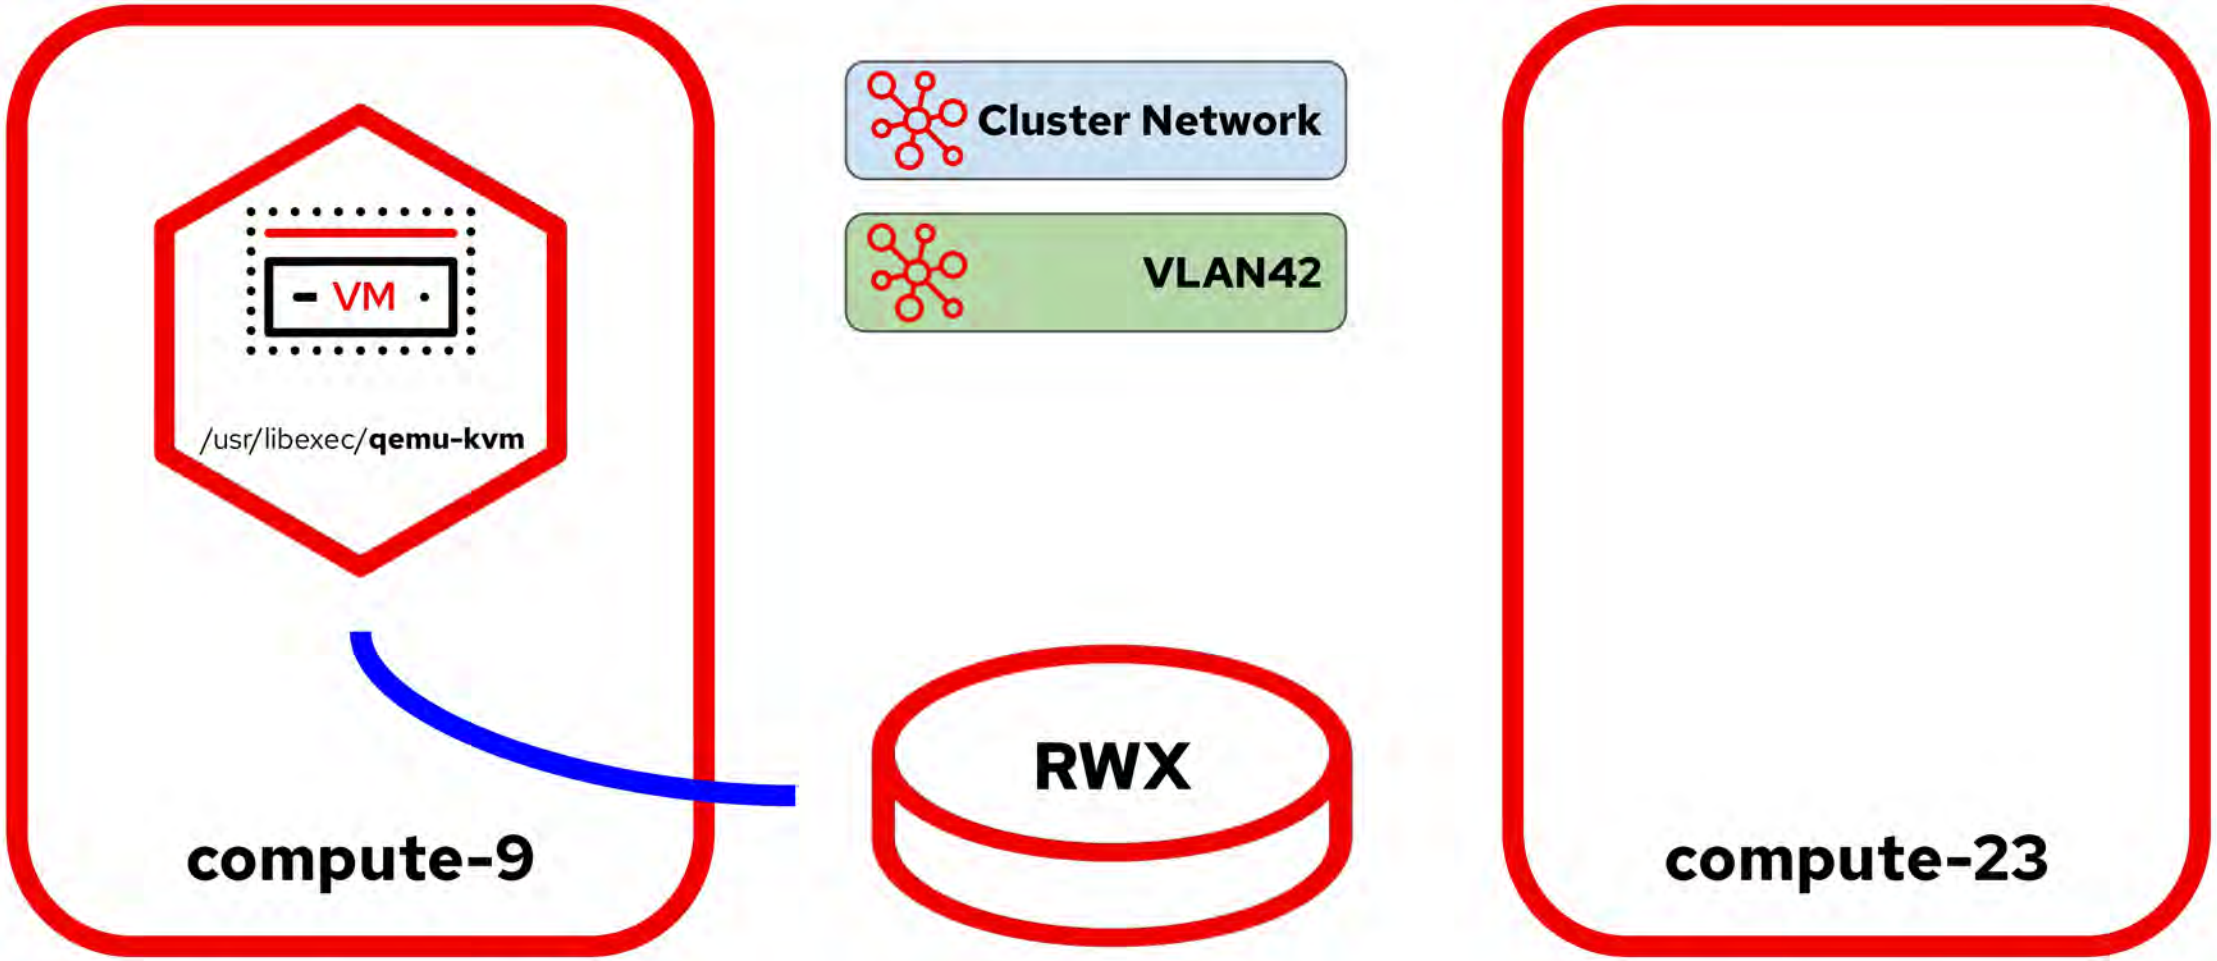

Virtual Machine Live Migration

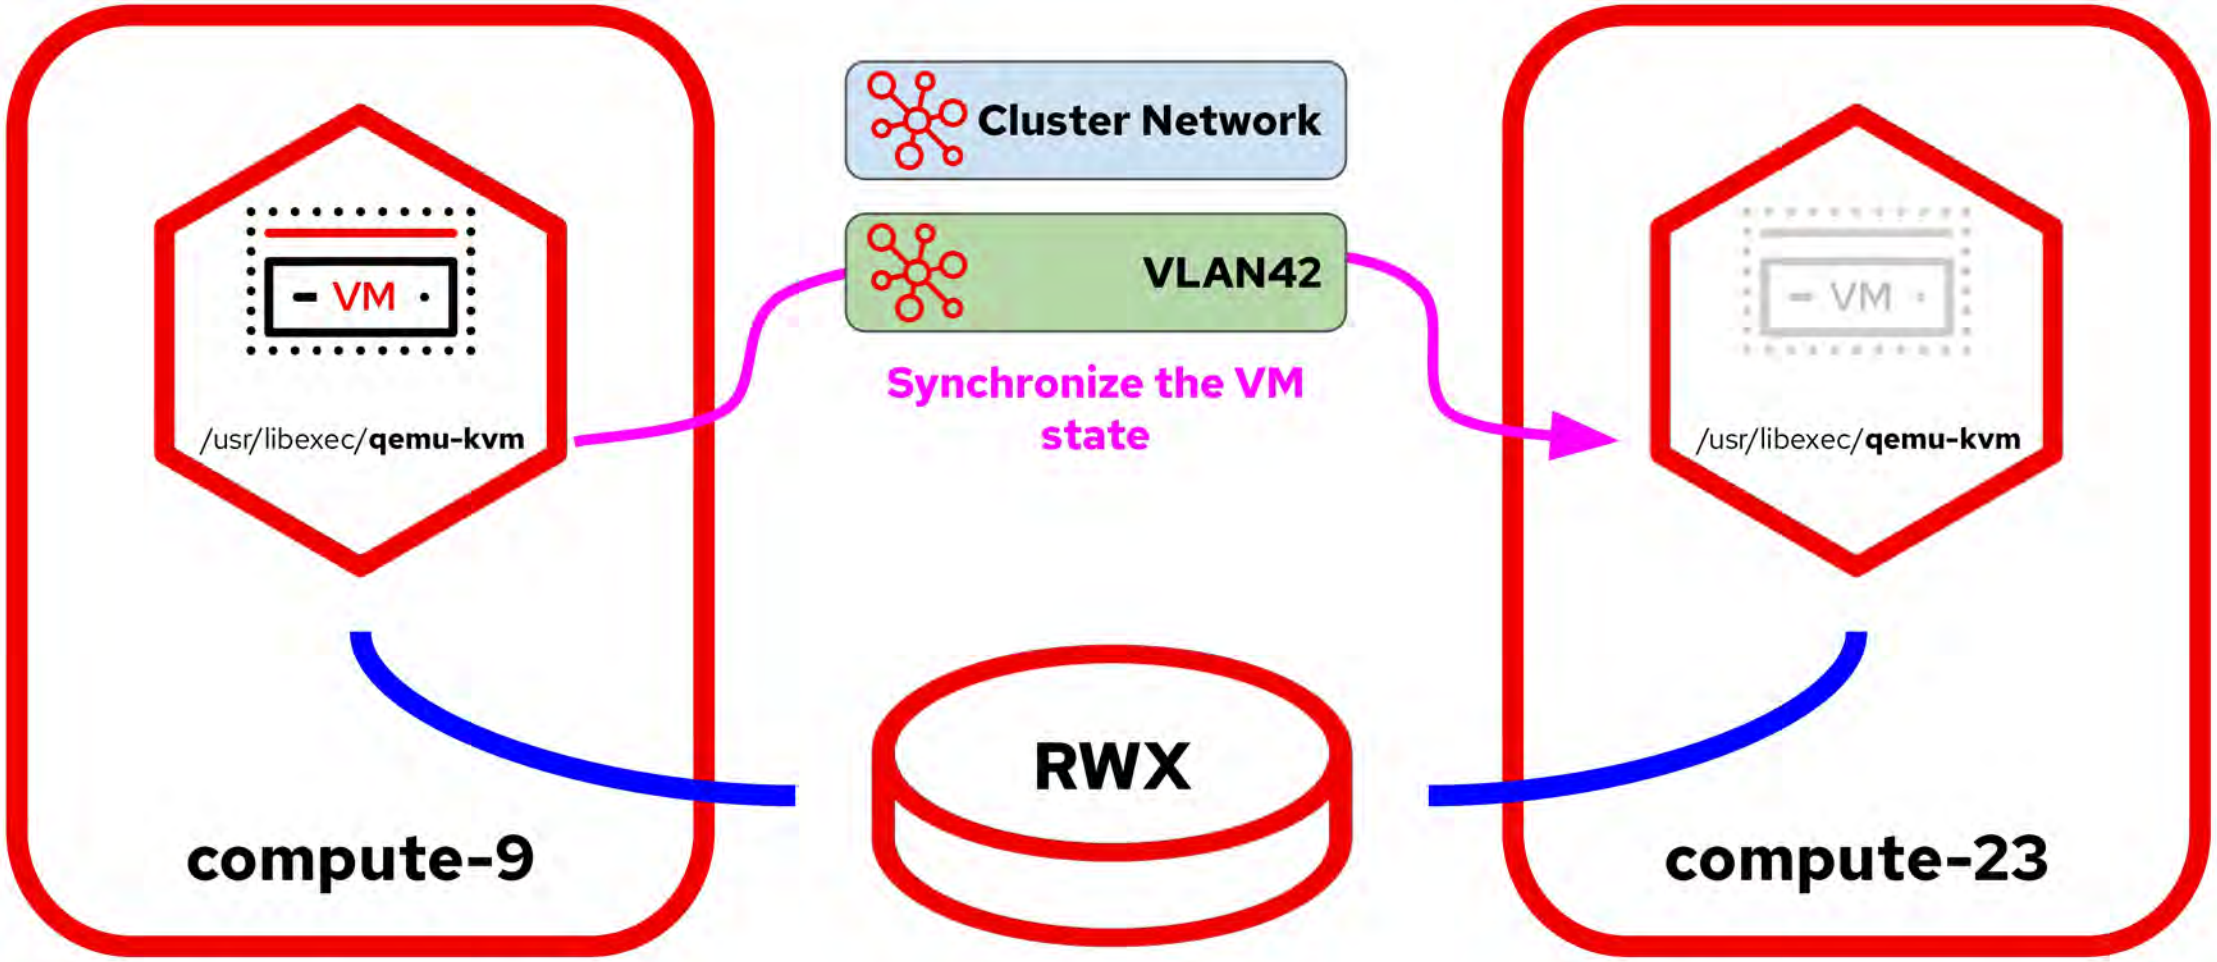

Virtual Machine Live Migration

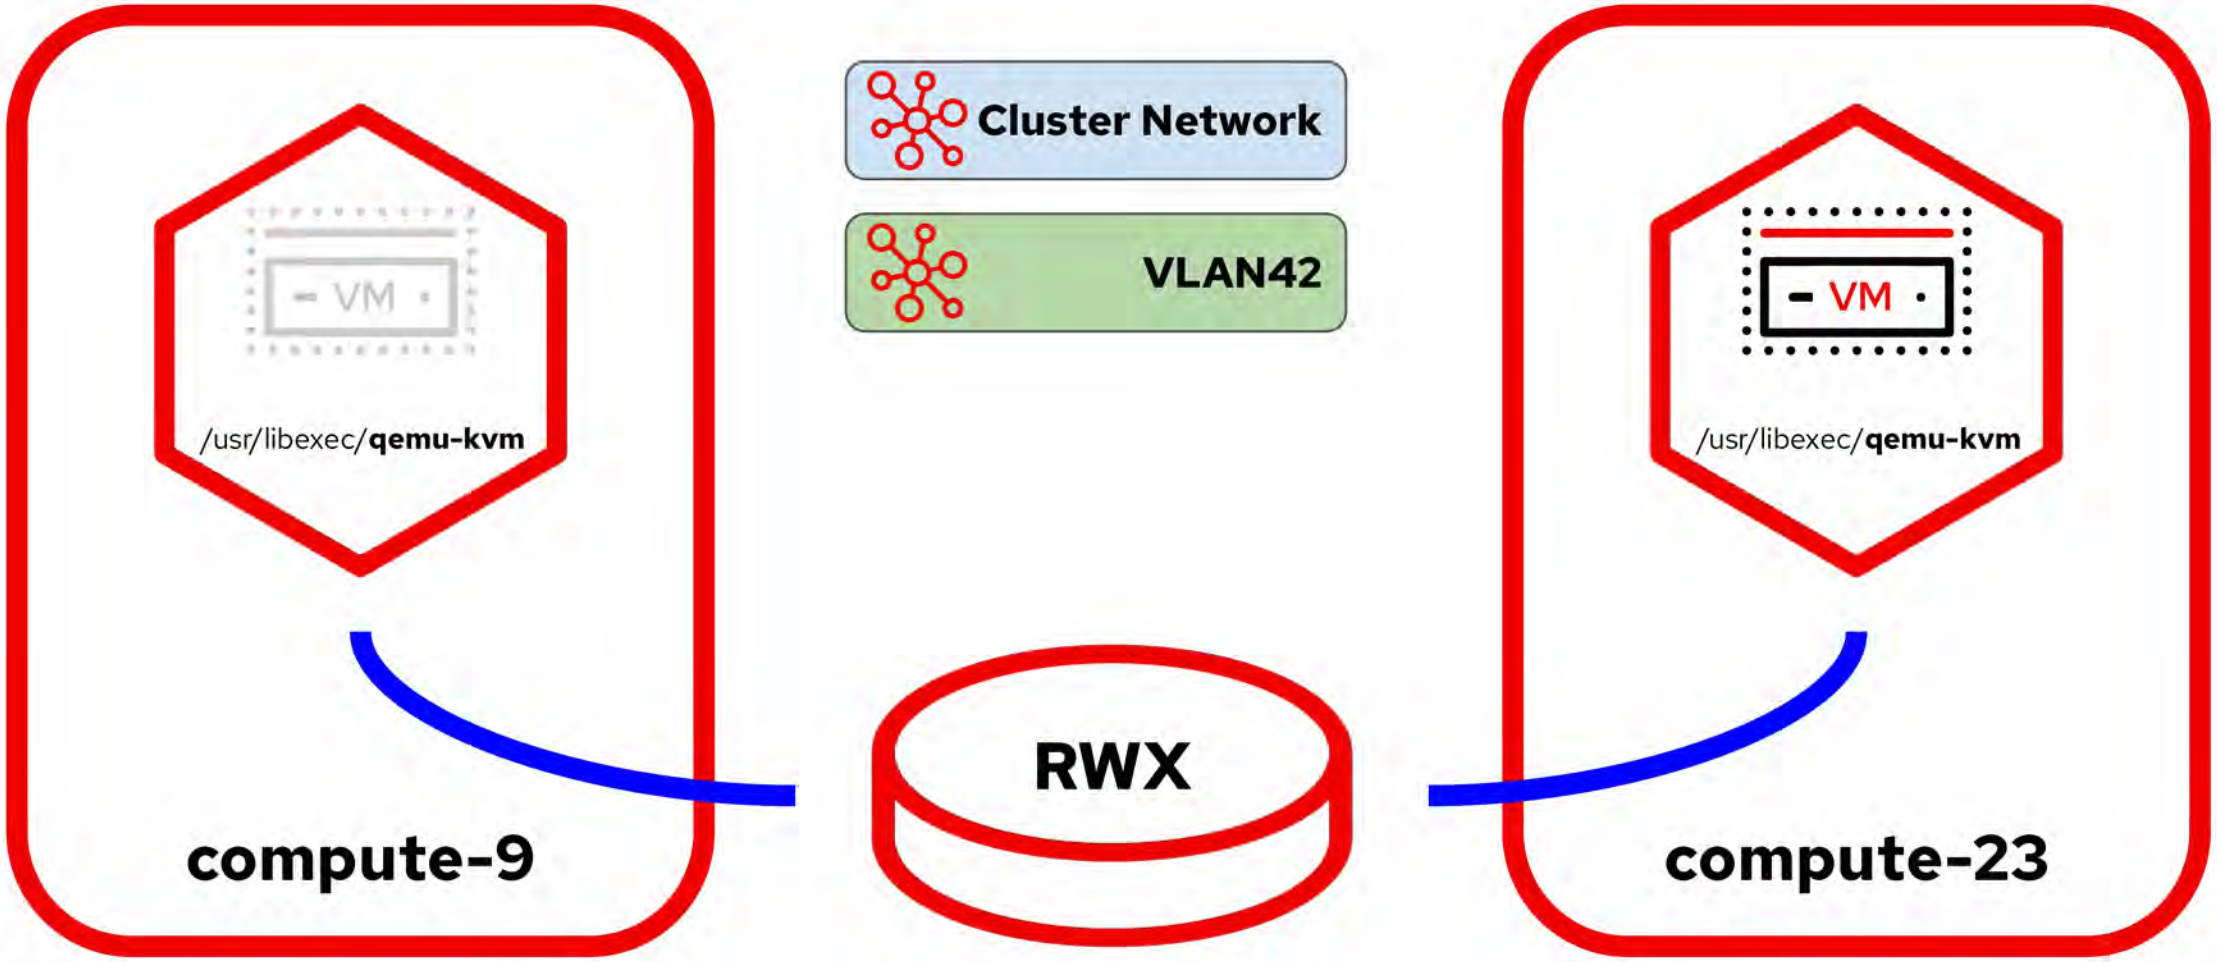

Virtual Machine Live Migration

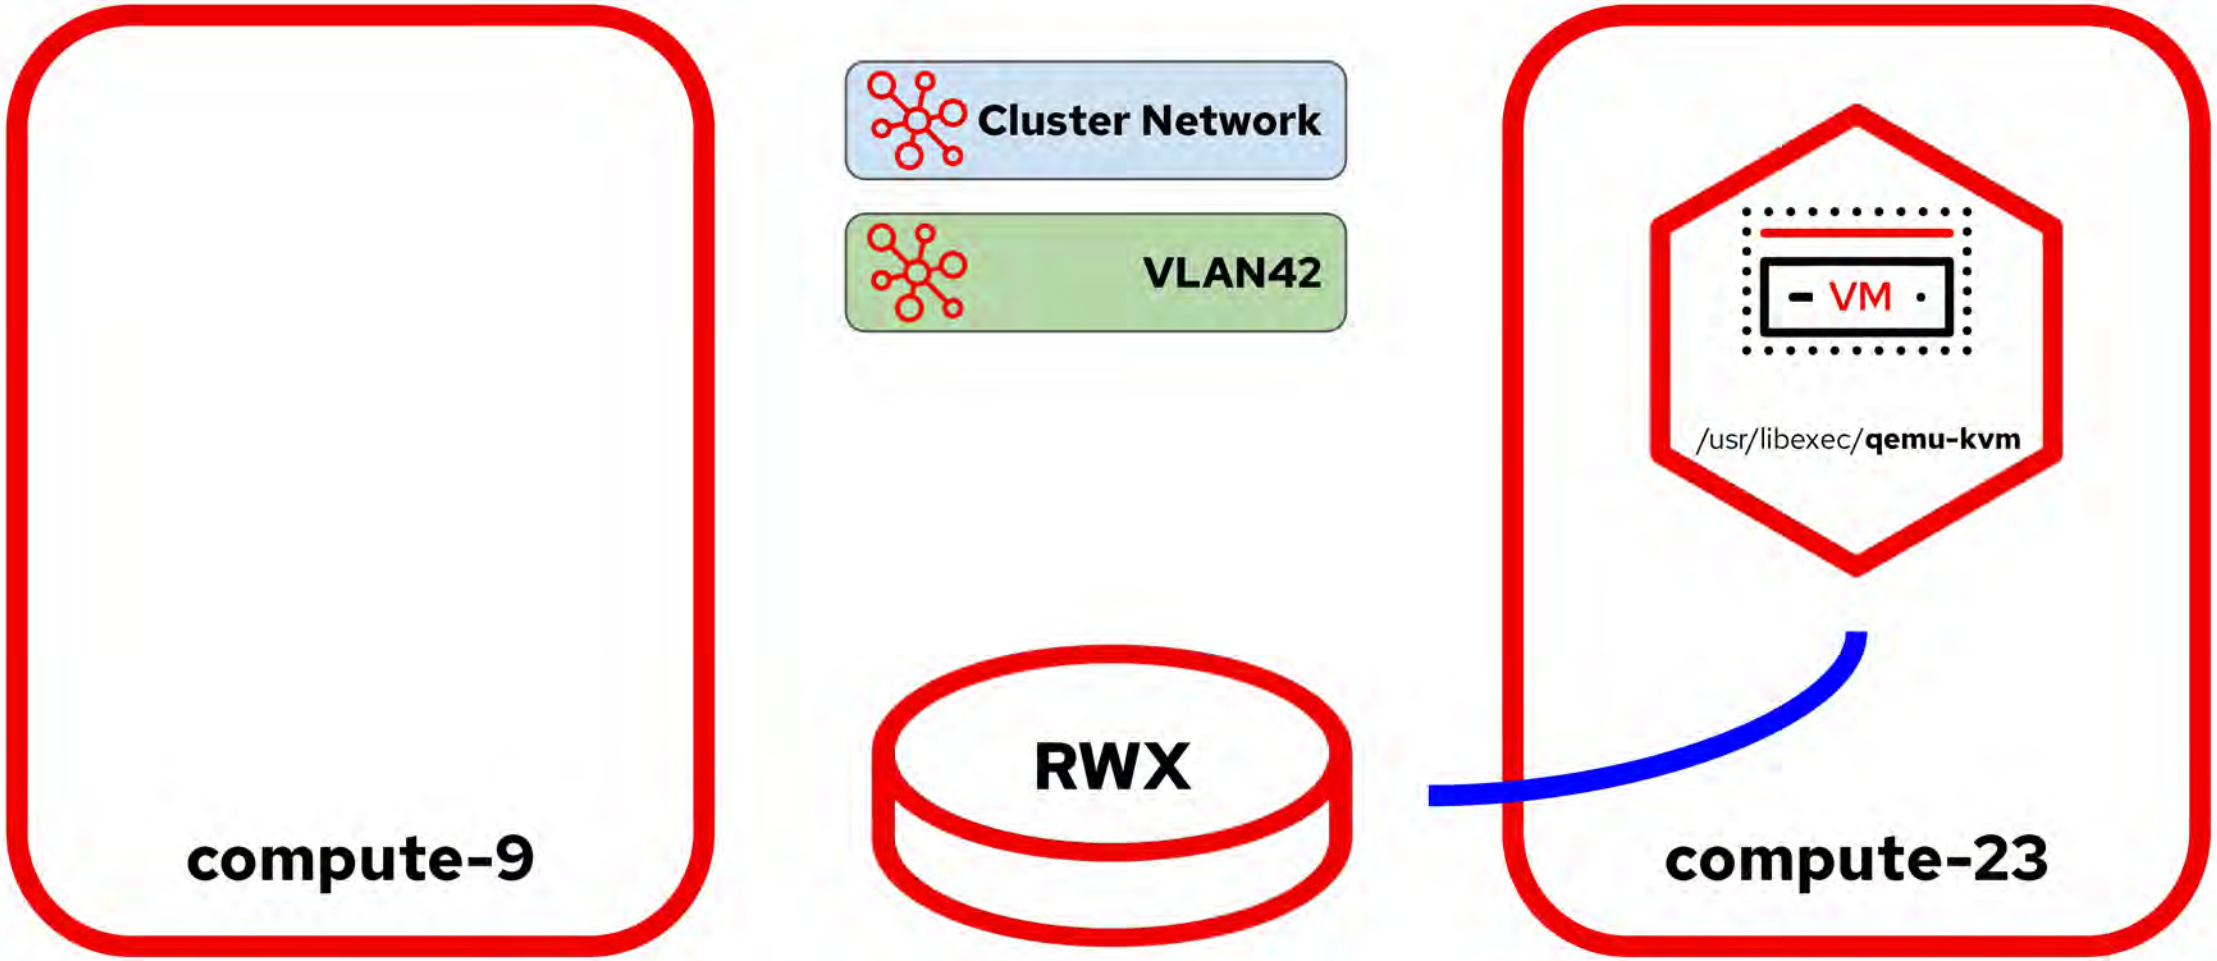

Cross-cluster Live Migration (Tech Preview 4.20)

Non-disruption workload migration in a “shared nothing” architecture

Operational workflows

- Promote a workload from **`dev`** to **`test`** to **`prod`**

Workload Isolation

- Move “noisy neighbors” away from production workloads

Resource or Capacity Balancing

- Move workload amongst a fleet of clusters to balance capacity

Lifecycle management

- To evacuate workloads for rack maintenance / lifecycle

Storage Live Migration

VM Storage Mobility

Non-disruptive VM storage live migration

Performance Optimization

- Migration VM disk to a faster storage tier

Workload Isolation

- Move “noisy neighbors” away from production workloads

Resource or Capacity Balancing

- Move VMs to free storage space for growth

Storage lifecycle management

- Evacuate aging storage for migration to new storage system

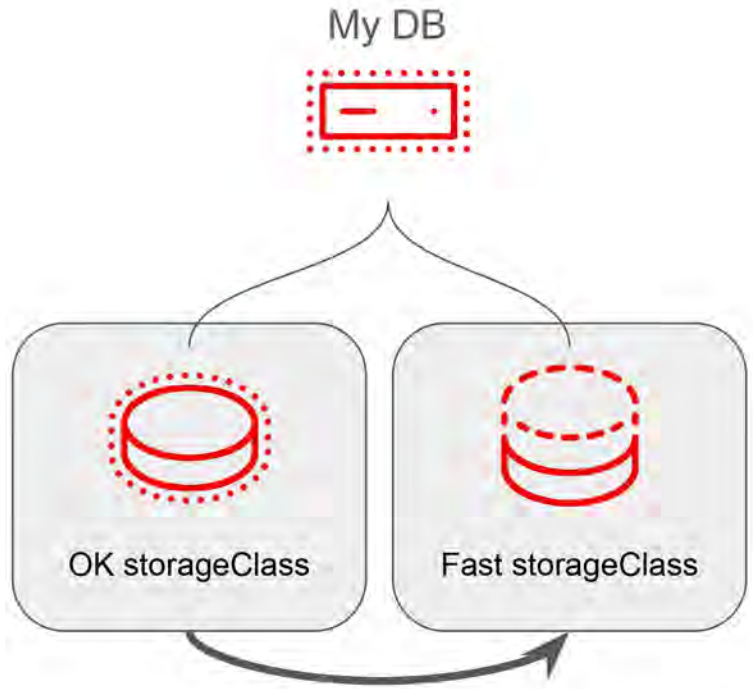

Initial Workloads

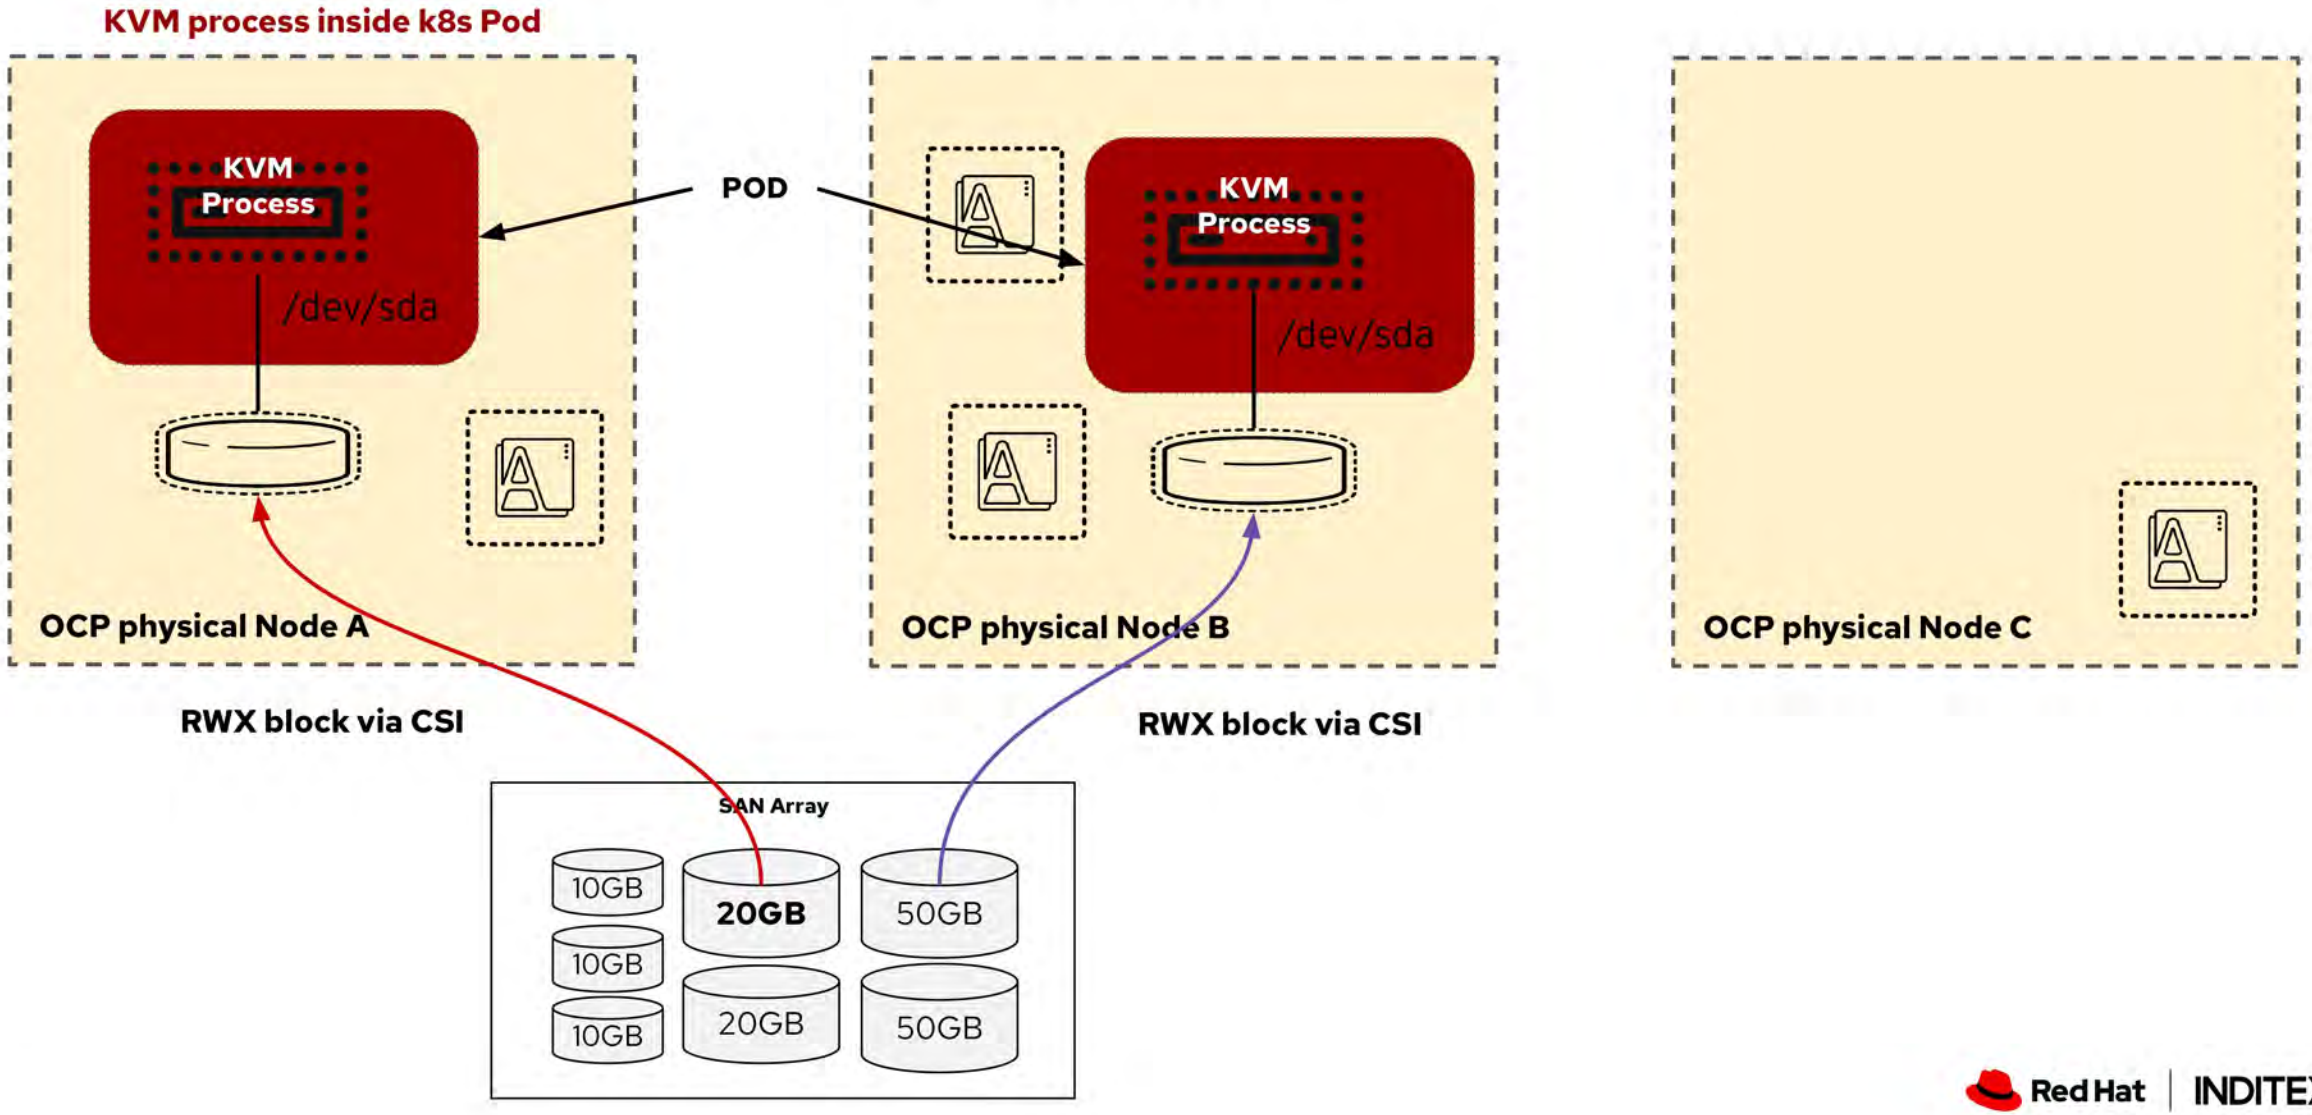

New Storage Connected to new target node

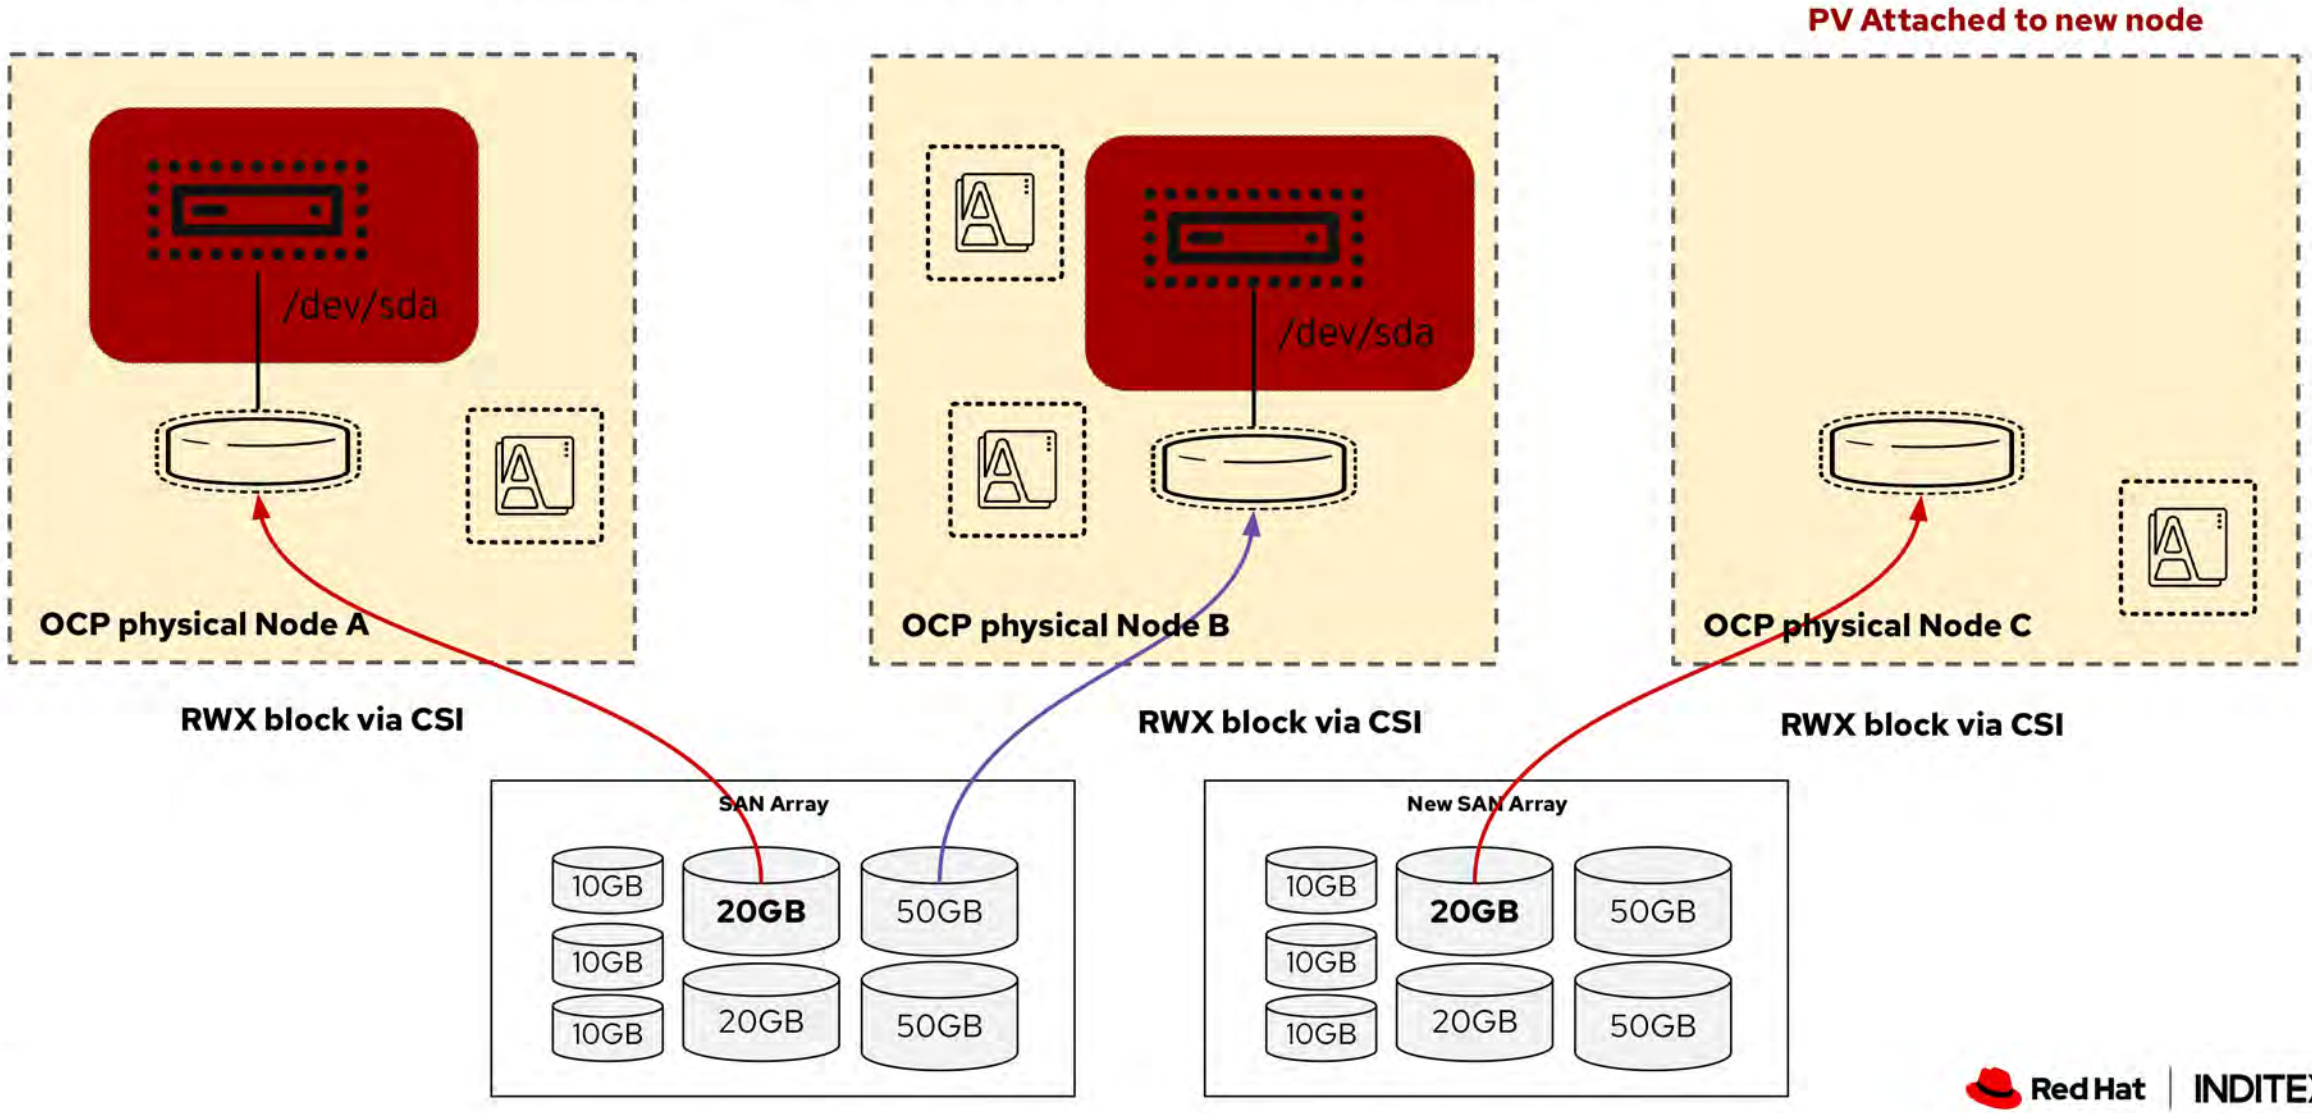

Red Hat | INDITEX

New Pod Created

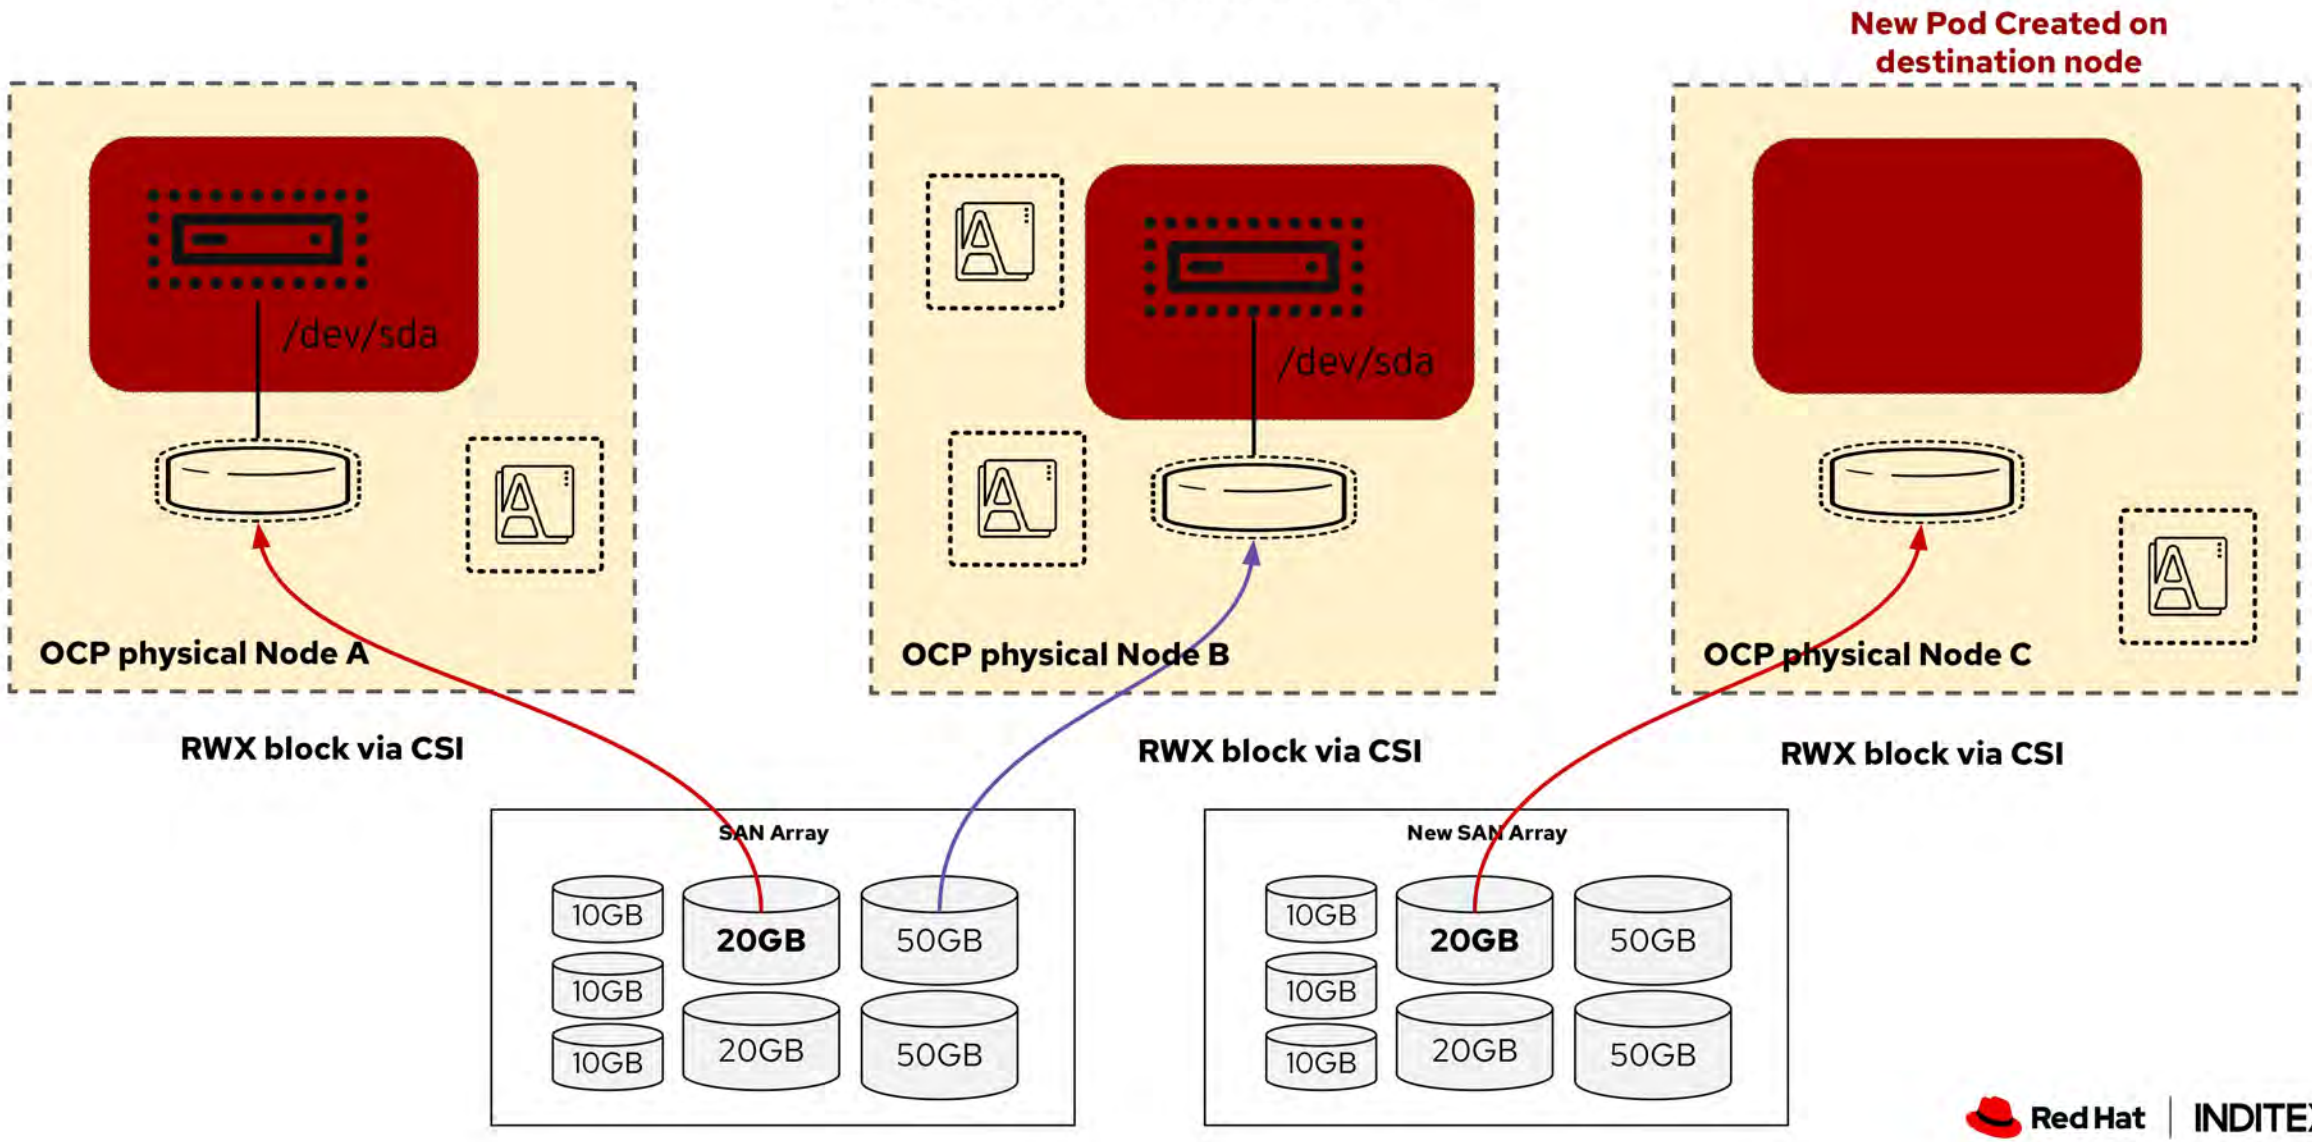

Red Hat | INDITEX

KVM process created

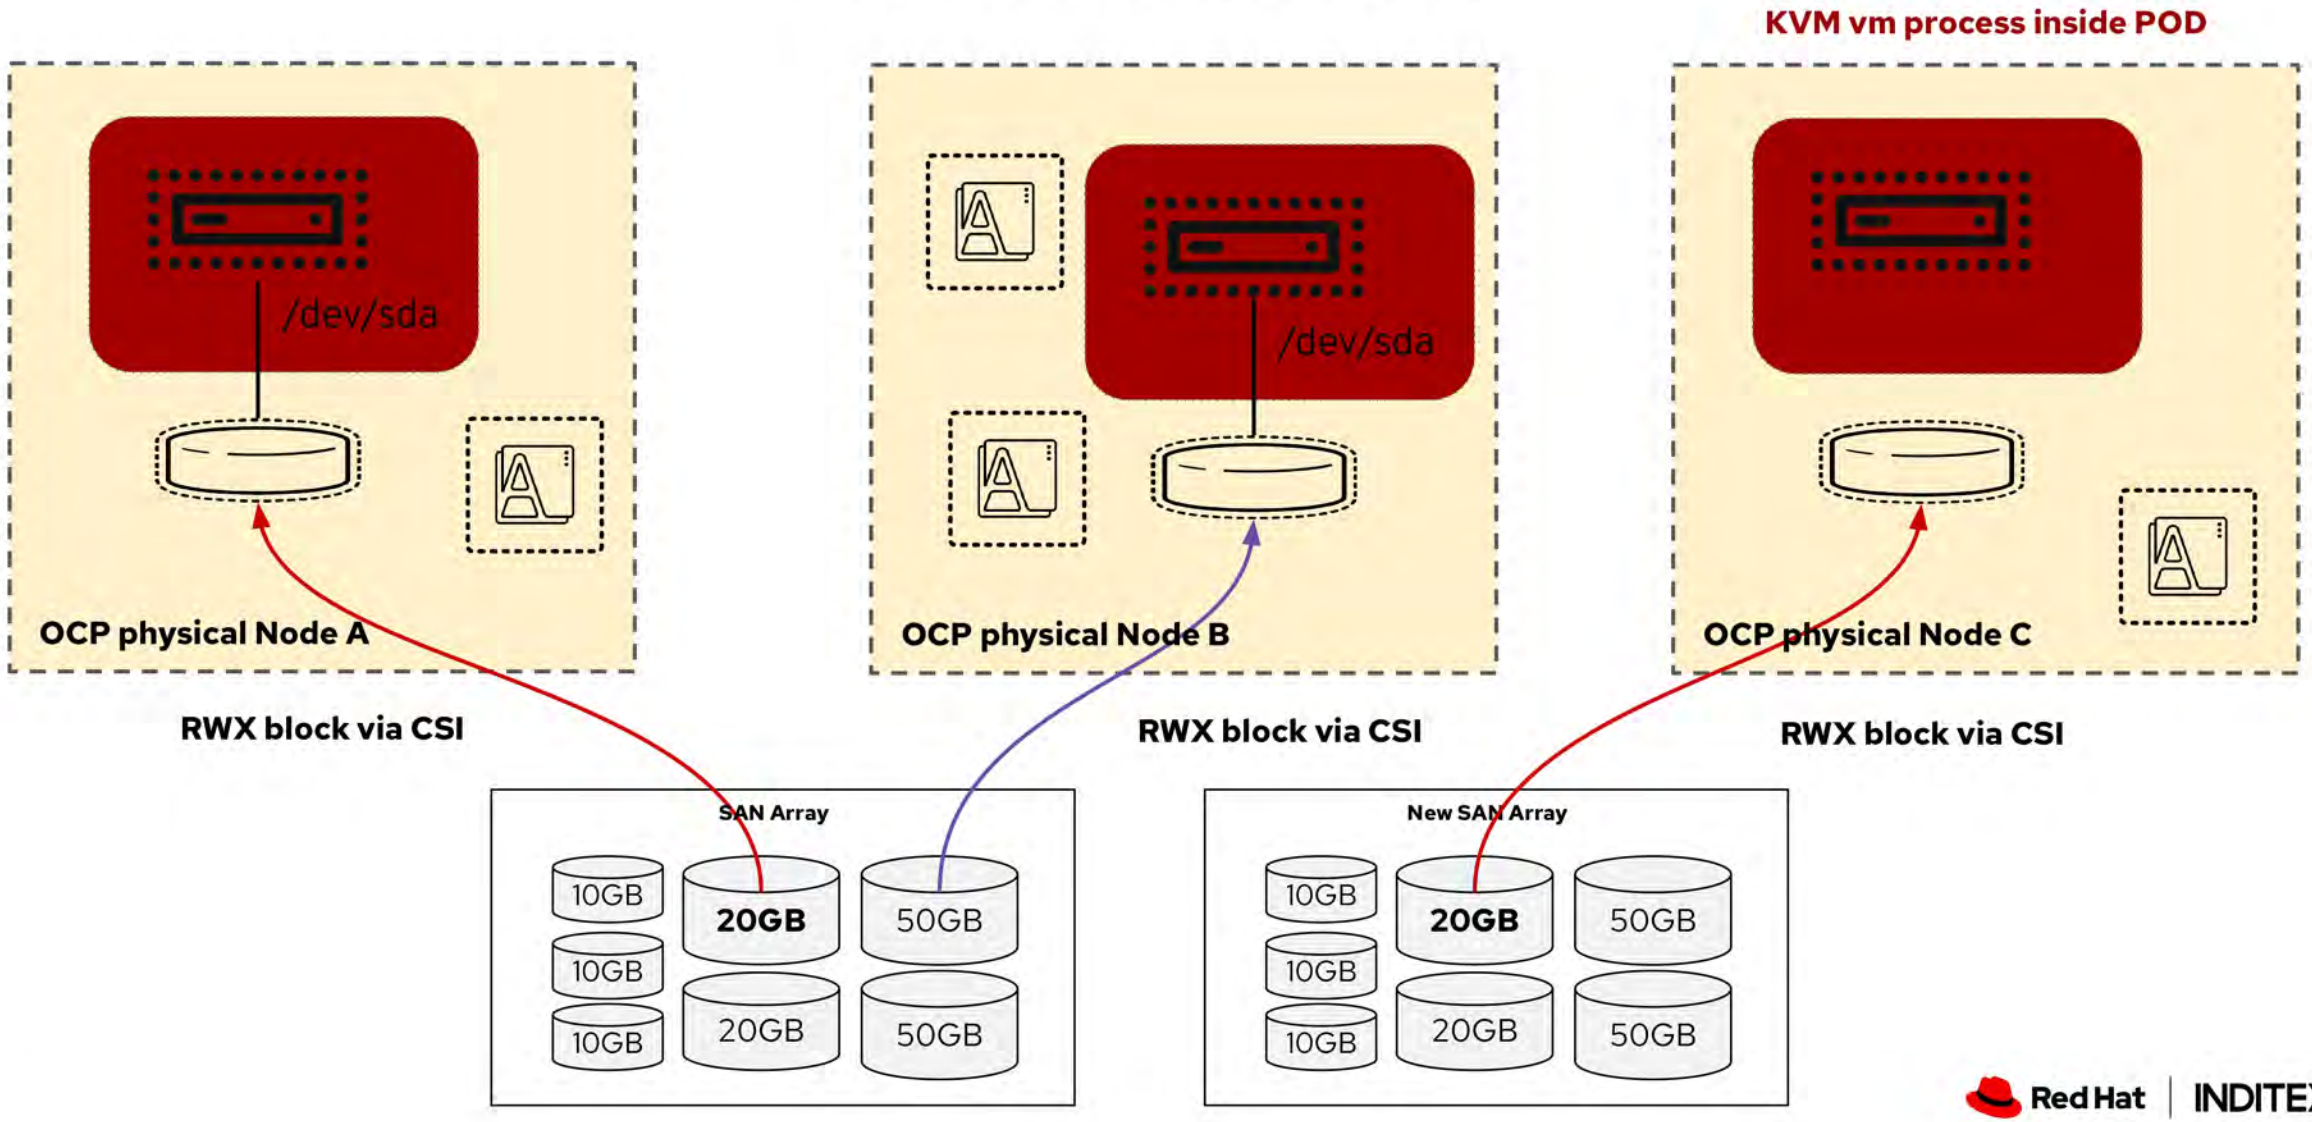

Red Hat | INDITEX

Memory/Disk synchronization

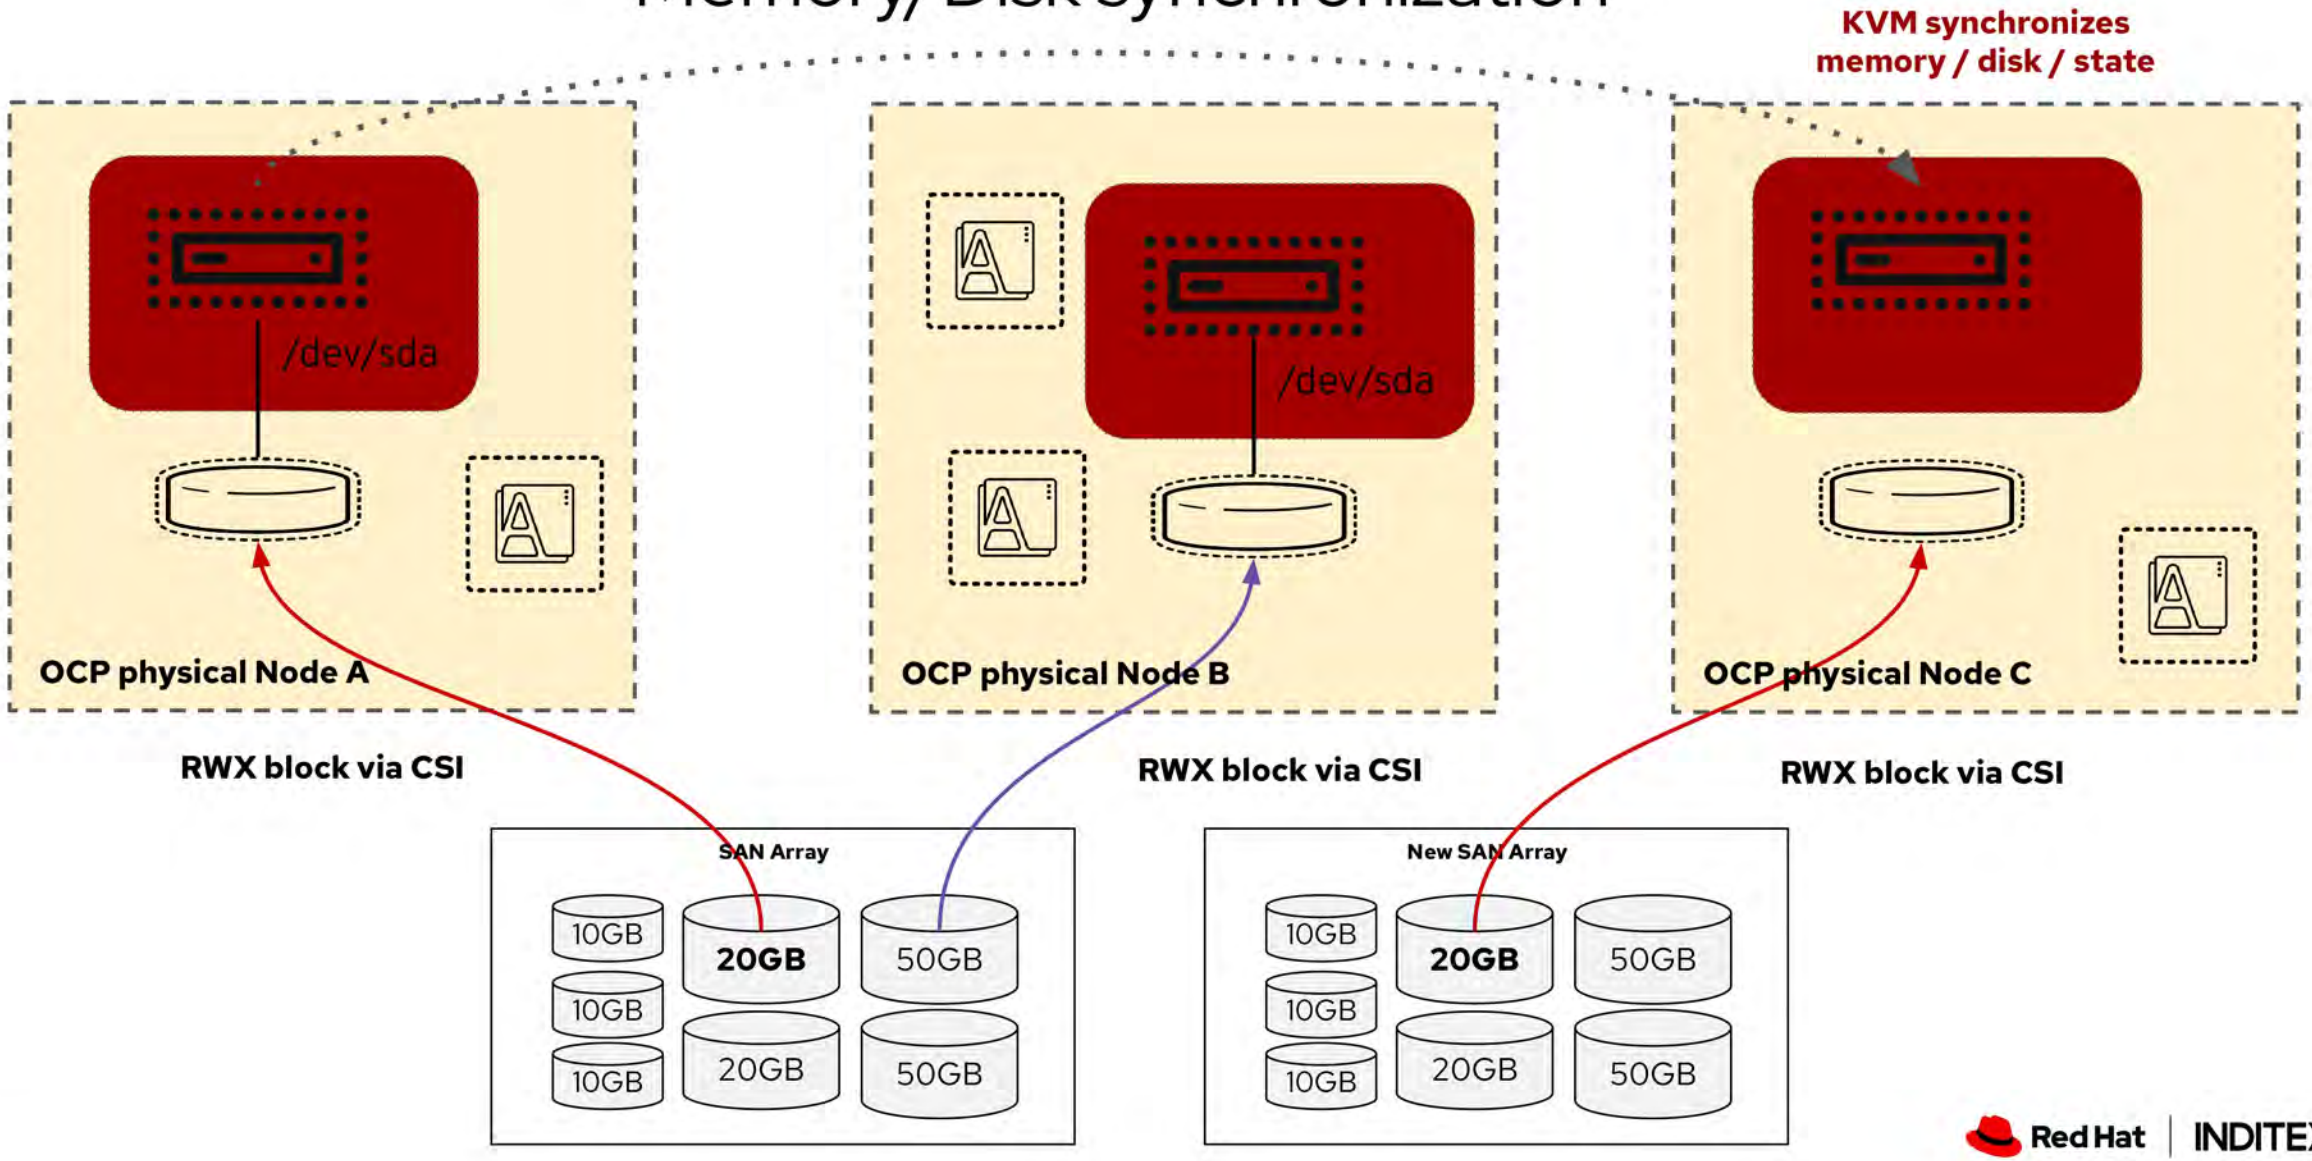

Red Hat | INDITEX

Final state and memory/disk synchronization

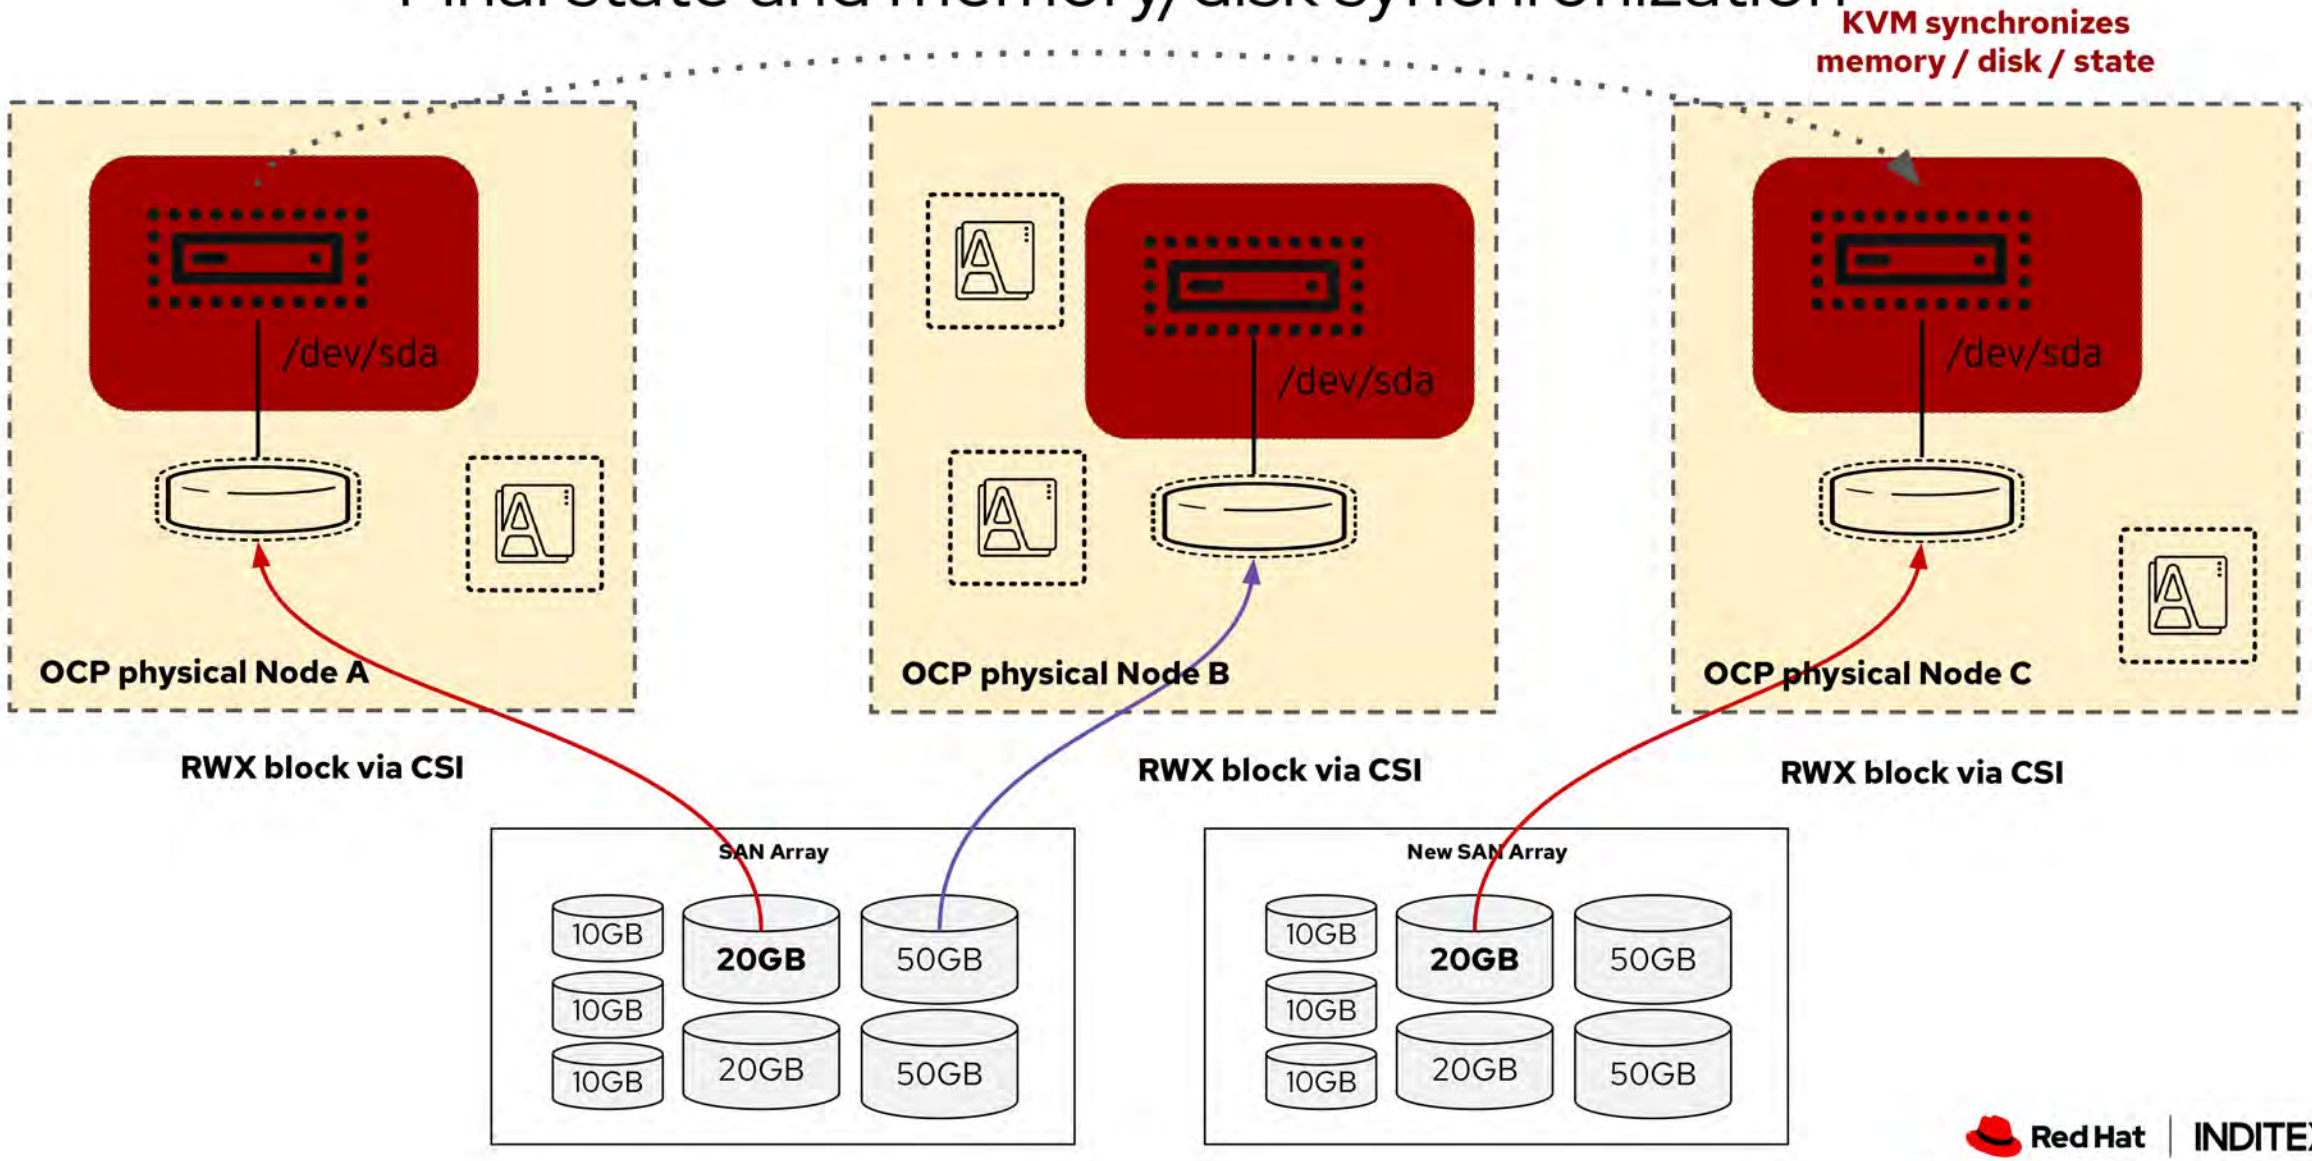

Migration Finalised

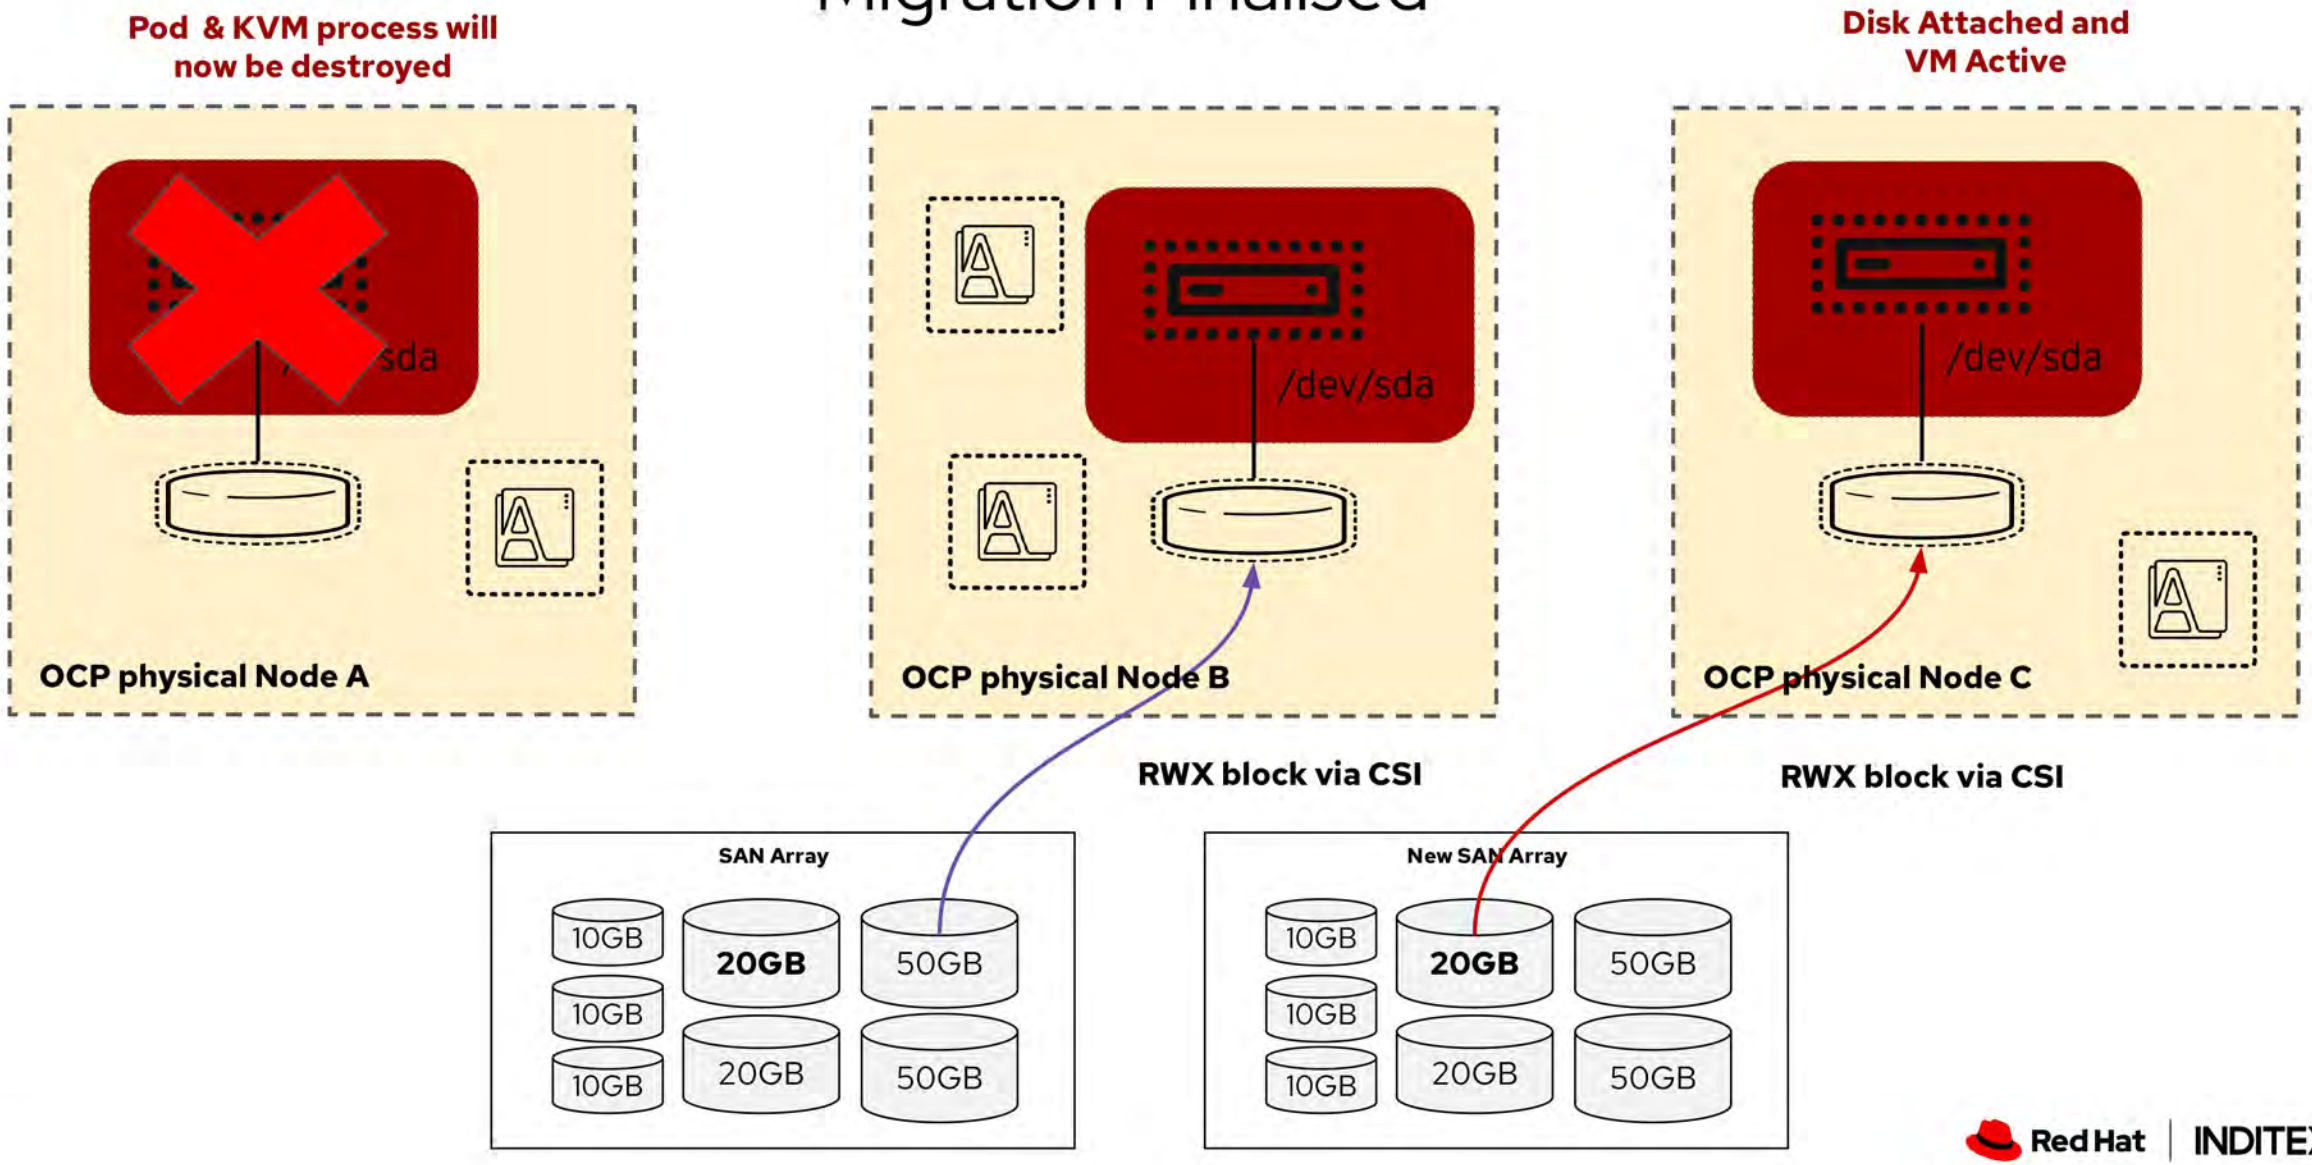

Red Hat | INDITEX

Workload migrated

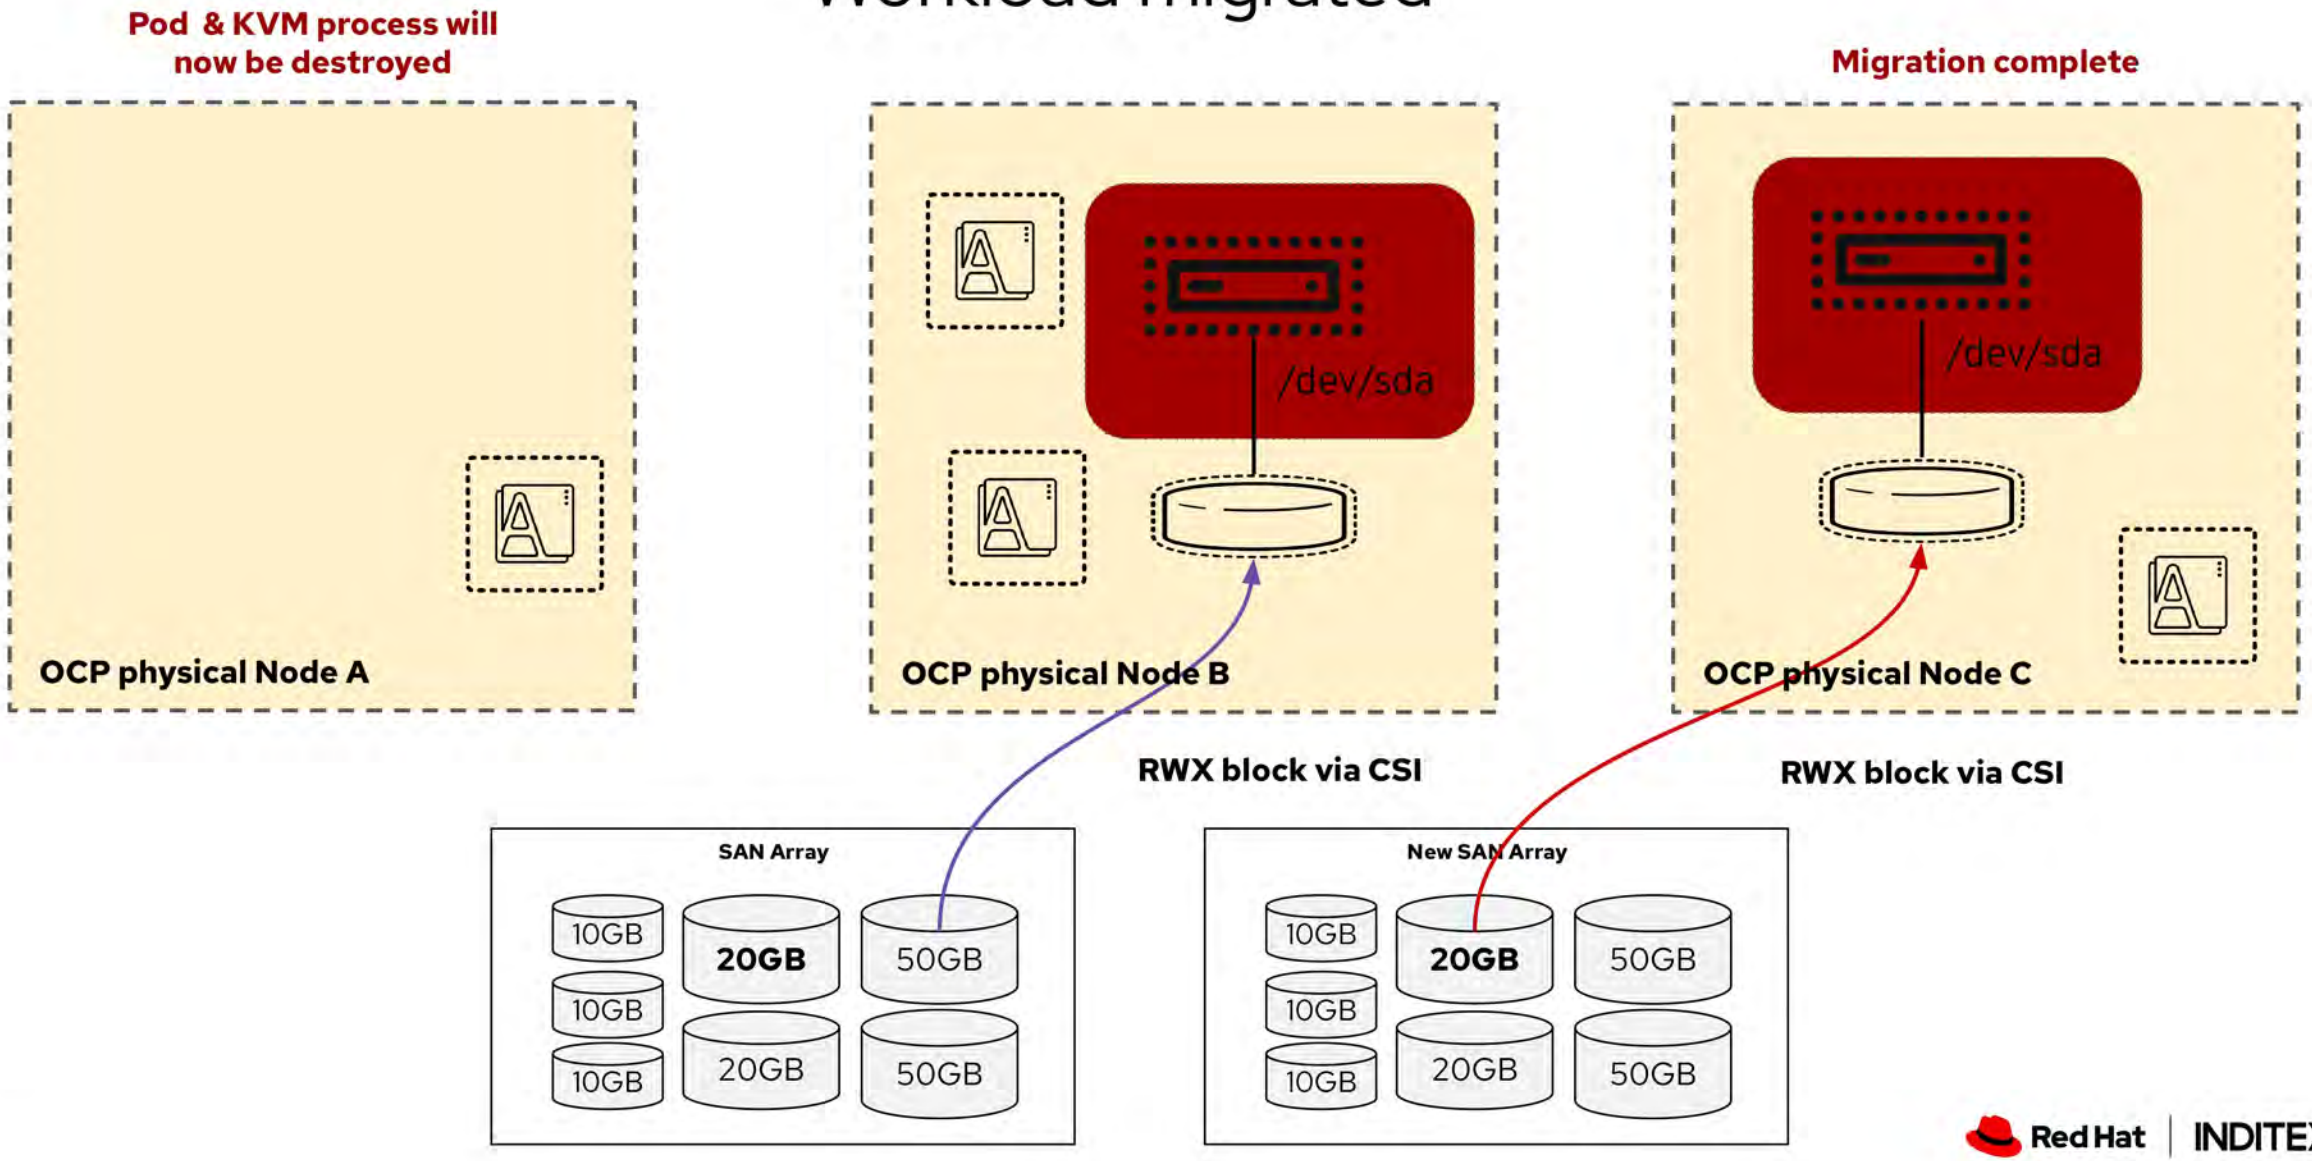

Red Hat | INDITEX

Migration Toolkit for Virtualization (MTV)

Migrating VM-based applications with minimal disruption

Migration toolkit for virtualization (MTV) included with OpenShift

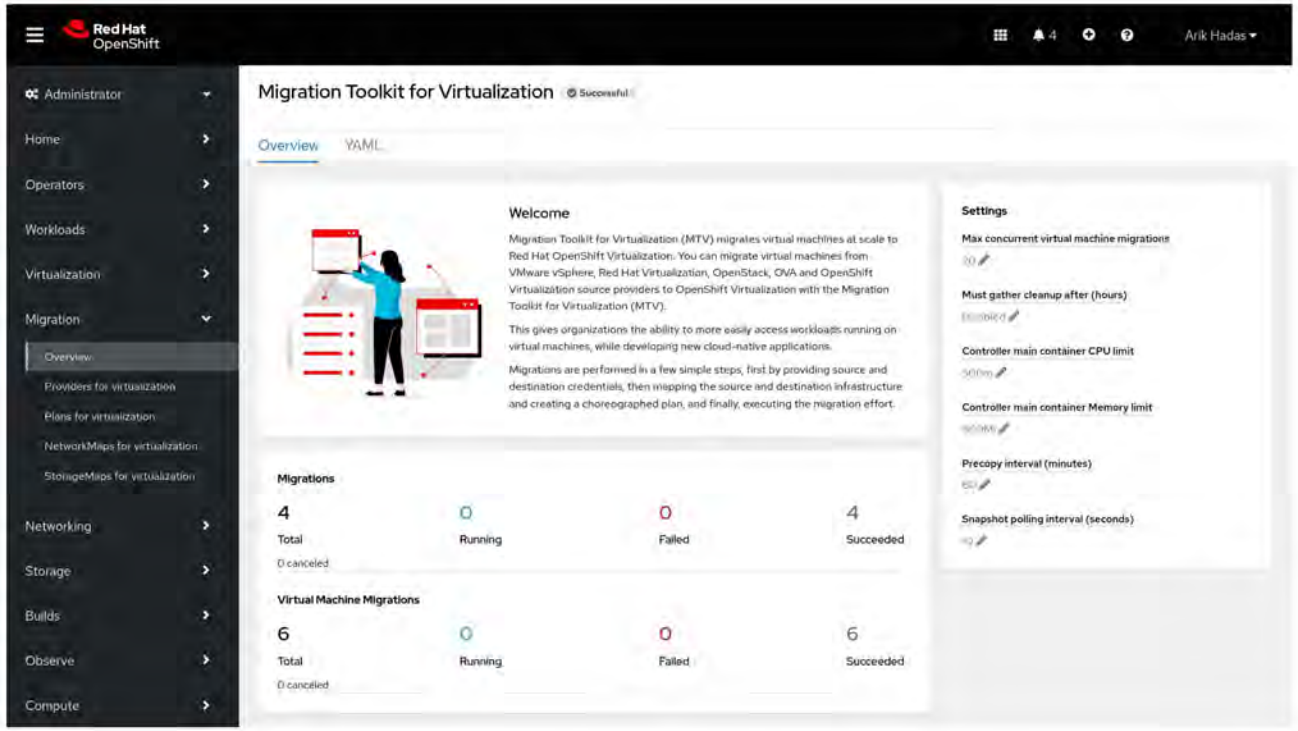

Easy migration of virtual machines

- Migrate virtual machines to OpenShift Virtualization in a few simple steps
- Provide source and destination credentials, map infrastructure, and create migration plans

Red Hat | INDITEX

Migration toolkit for virtualization

Making mass migration possible

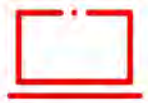

Easy to use UI

Simple, easy to use interface

Mass migration to OpenShift

Easily migrate VMs from VMware, Red Hat Virtualization and OpenStack to OpenShift and between OpenShift clusters

Pre-copied data

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

Migration toolkit for virtualization

Instant access to virtual machine migration at scale

Red Hat Virtualization

vmware

Red Hat OpenStack Services on OpenShift

Red Hat OpenShift

OpenShift Operator Hub

MTV is available as a one-click download in the OpenShift Operator Hub and integrates into the OpenShift menu structure.

Reduce Complexity

MTV offers all the tools needed to quickly and easily migrate virtual machines from alternative hypervisors without disrupting existing VM investments.

Provider-Plugin availability

Migration sources are defined by “Provider-Plugins”

Currently available Sources: RHV, VMware, OpenStack, vSphere-compatible OVA, OpenShift

Highlight features in 2.9

Storage Offloading

- Storage Offloading in MTV allows the migration process to delegate data transfer tasks to underlying storage systems.
- Tests on Hitachi Storage showing 10+ times faster copy processes.
- This significantly improves performance, reduces network load, and enables faster, more scalable VM migrations—especially in environments with CSI-compatible or advanced storage backends.

Raw Copy Mode

- The goal is to allow import of virtual machines (VMs) that fail during conversion—especially those with unsupported guest operating systems—by performing a raw disk copy without guest OS conversion.
- This feature is especially useful in environments where strict compatibility with guest OS types is not guaranteed, offering a fallback strategy to preserve and access VM content even if full automated migration fails.
- Please note that RCM is NOT what we want to have as a standard option, there is a reason why it has been called rescue mode as a working title.

Storage Offload

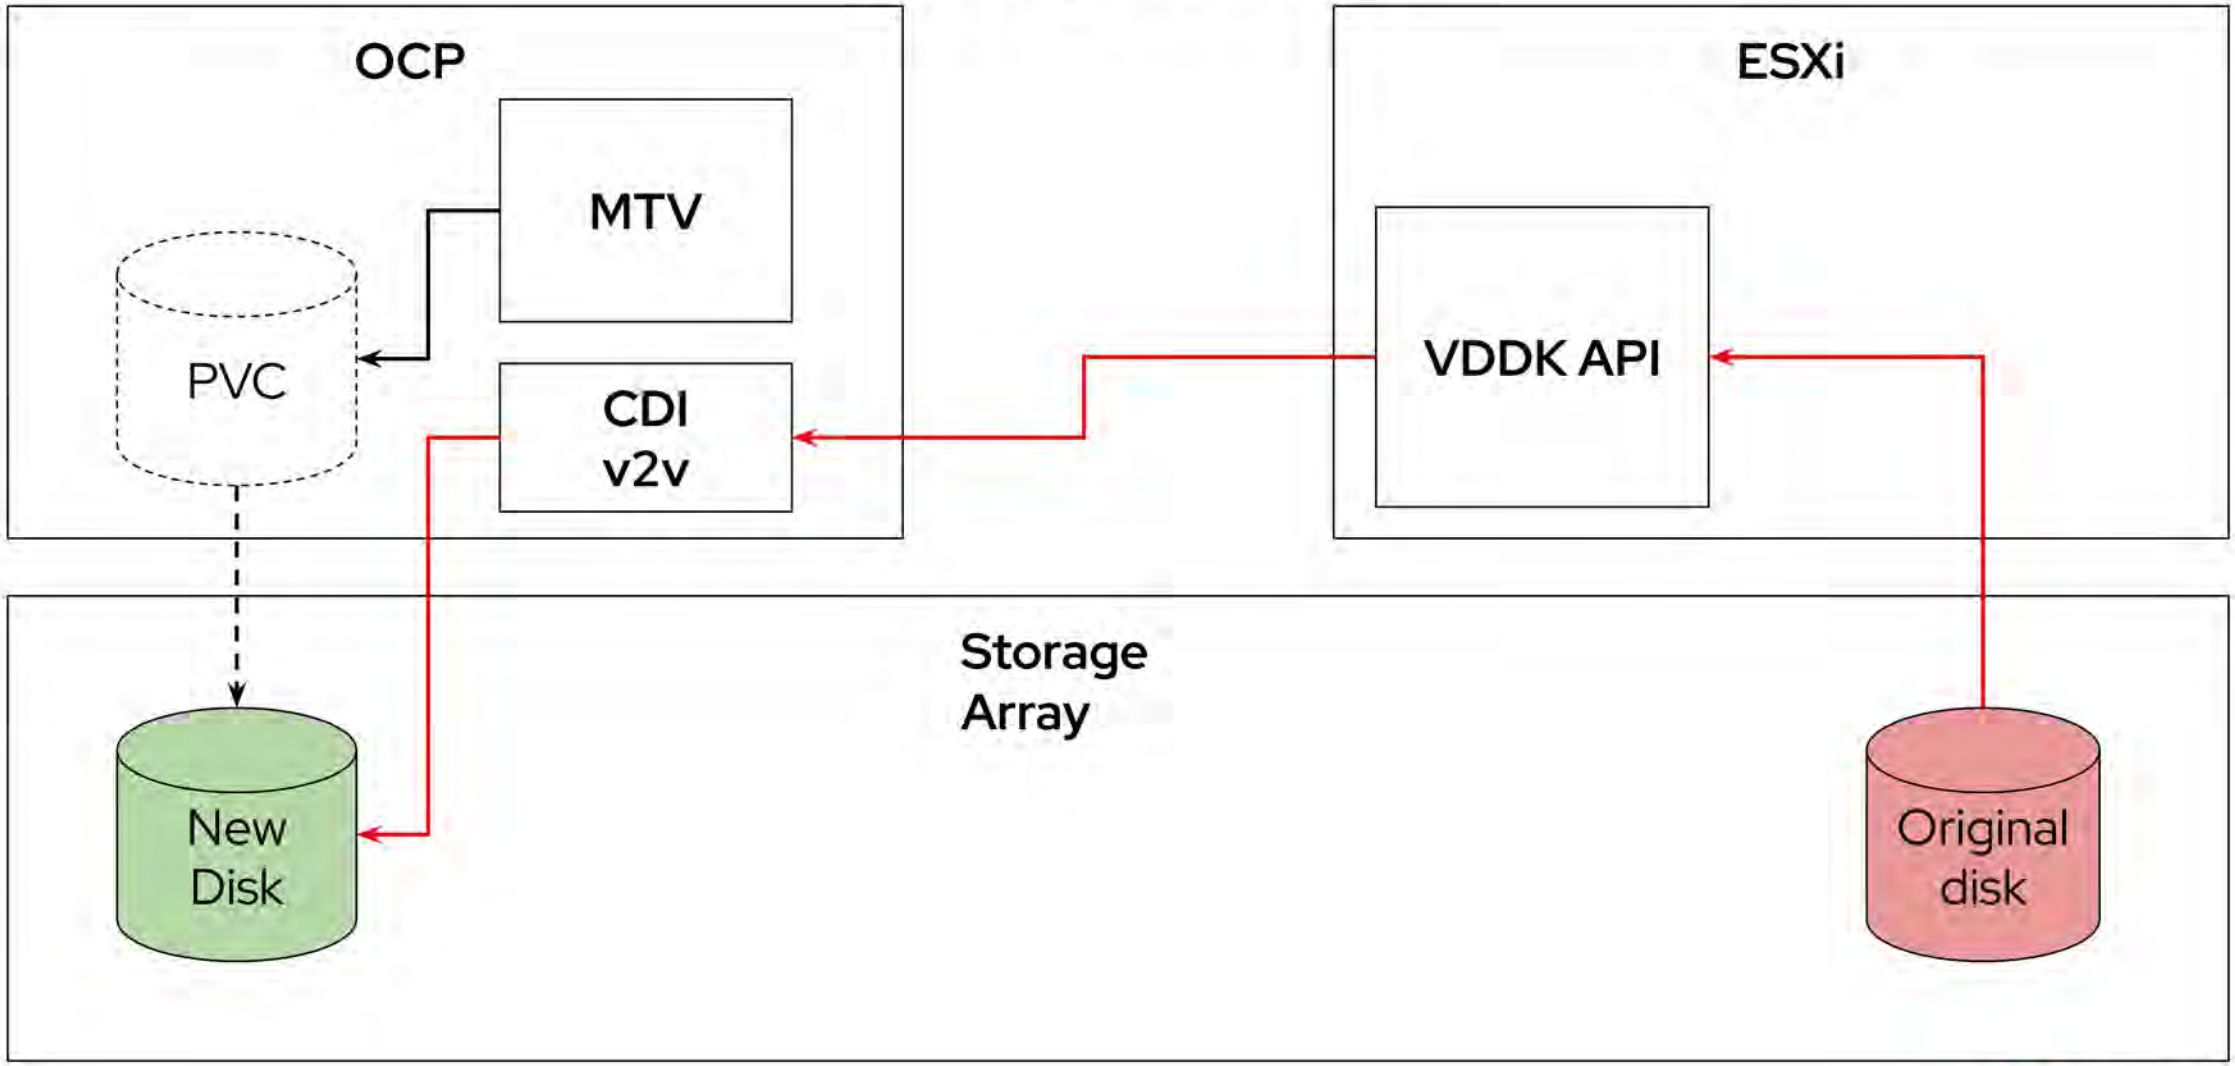

Storage Offload

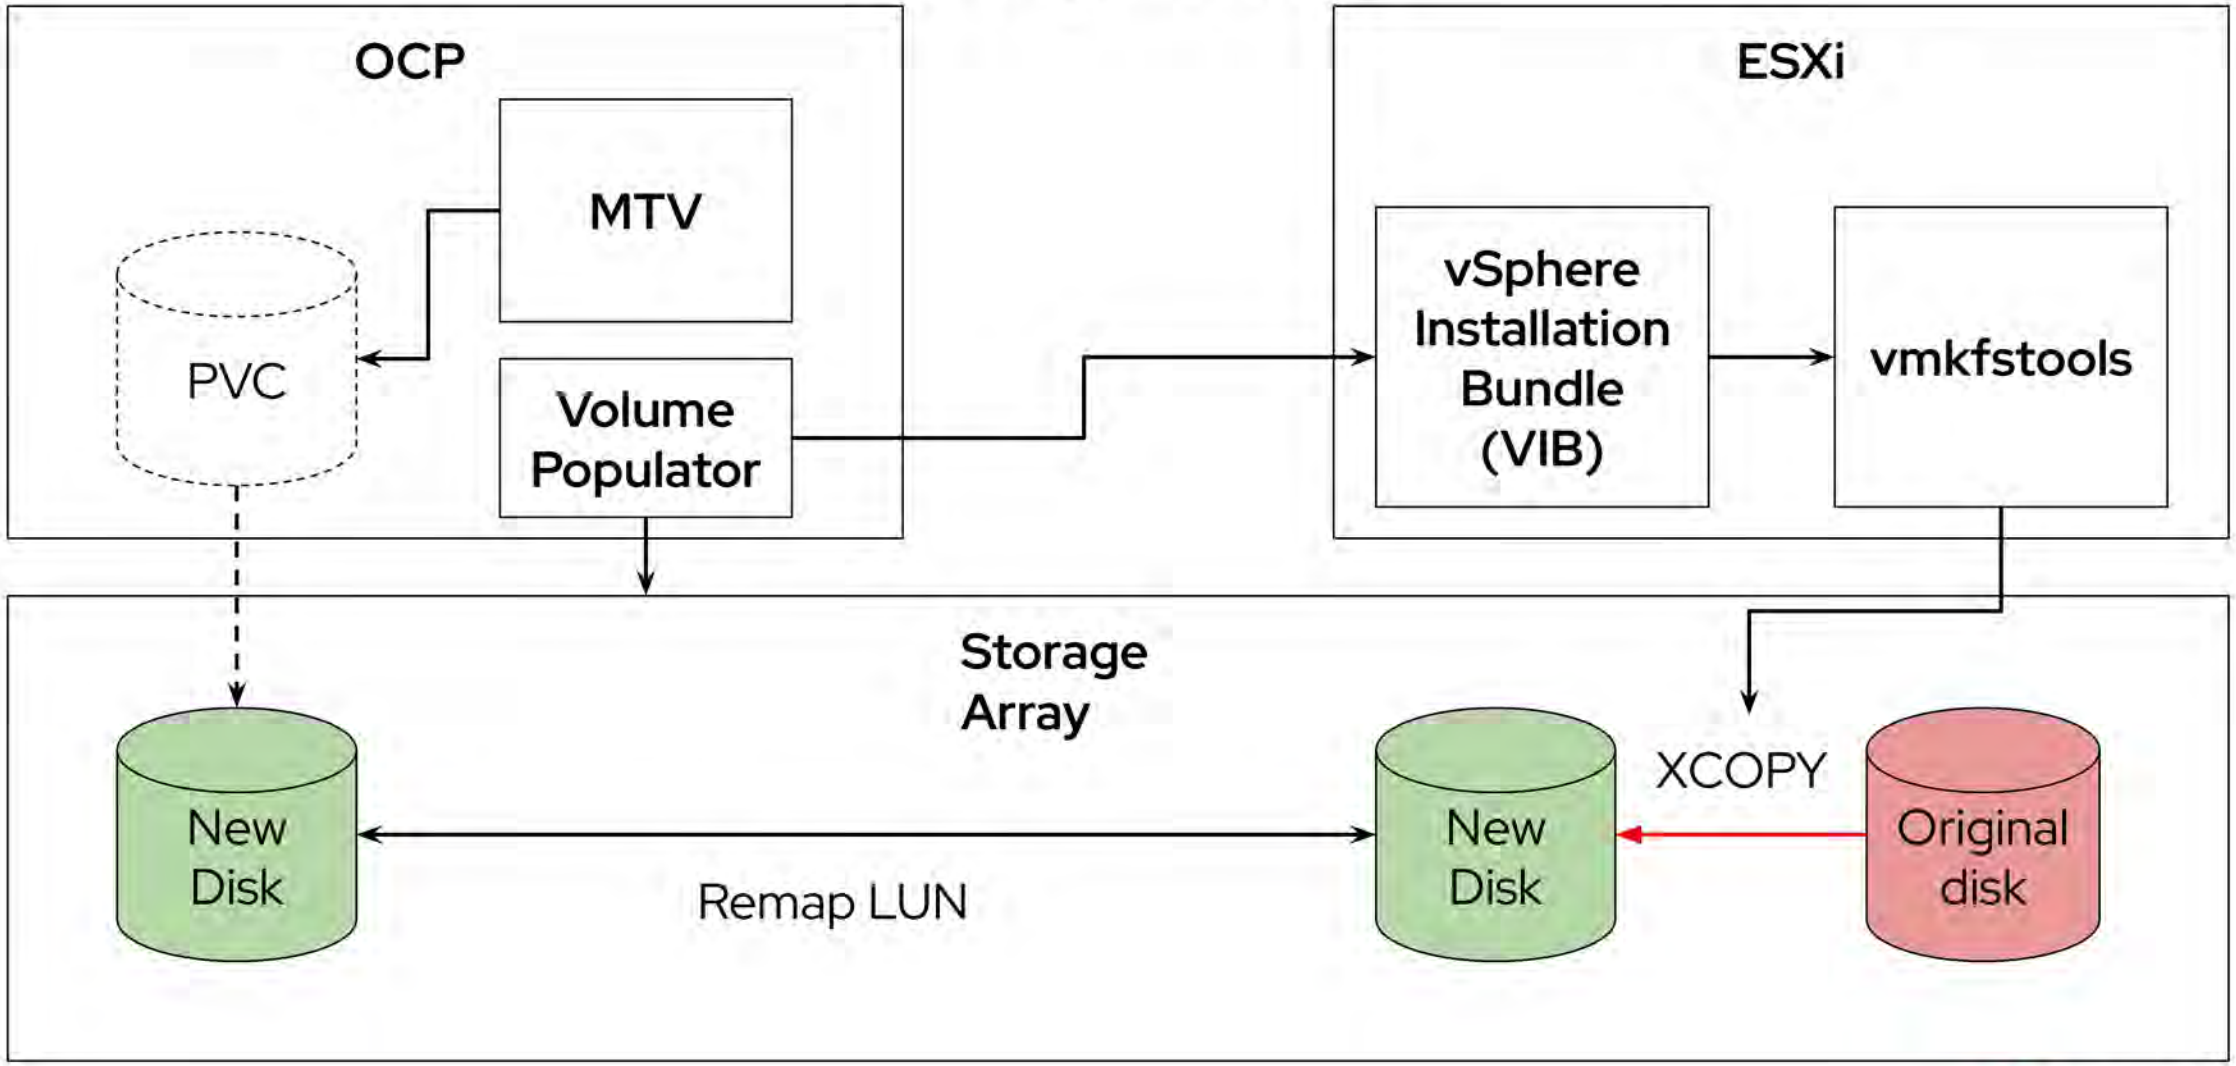

Demo
→ OpenShift Virtualization

Terminology comparison

Feature	RHV	OpenShift Virtualization	vSphere
Where VM disks are stored	Storage Domain	PVC / PV	Datastore
Policy based storage	None	StorageClass	Storage Policy Based Management (SPBM)
Non-disruptive VM migration	Live migration	Live migration	vMotion
Non-disruptive VM storage migration	Storage live migration	Storage class migration	Storage vMotion
Active resource balancing	Cluster scheduling policy	Pod eviction policy, descheduler	Dynamic Resource Scheduling (DRS)
Physical network configuration	Host network config (via nmstate w/4.4)	NMState Operator, Multus	vSwitch / DvSwitch
Overlay network configuration	OVN	OCP SDN (OpenShift-SDN, OVN-Kubernetes, CNI partners), Multus	NSX-T
Host / VM metrics	Data warehouse + Grafana (RHV 4.4)	OpenShift Metrics and Monitoring	vCenter, vRealize Operations

Bare Metal cluster deployment options

Typical Cluster Layout

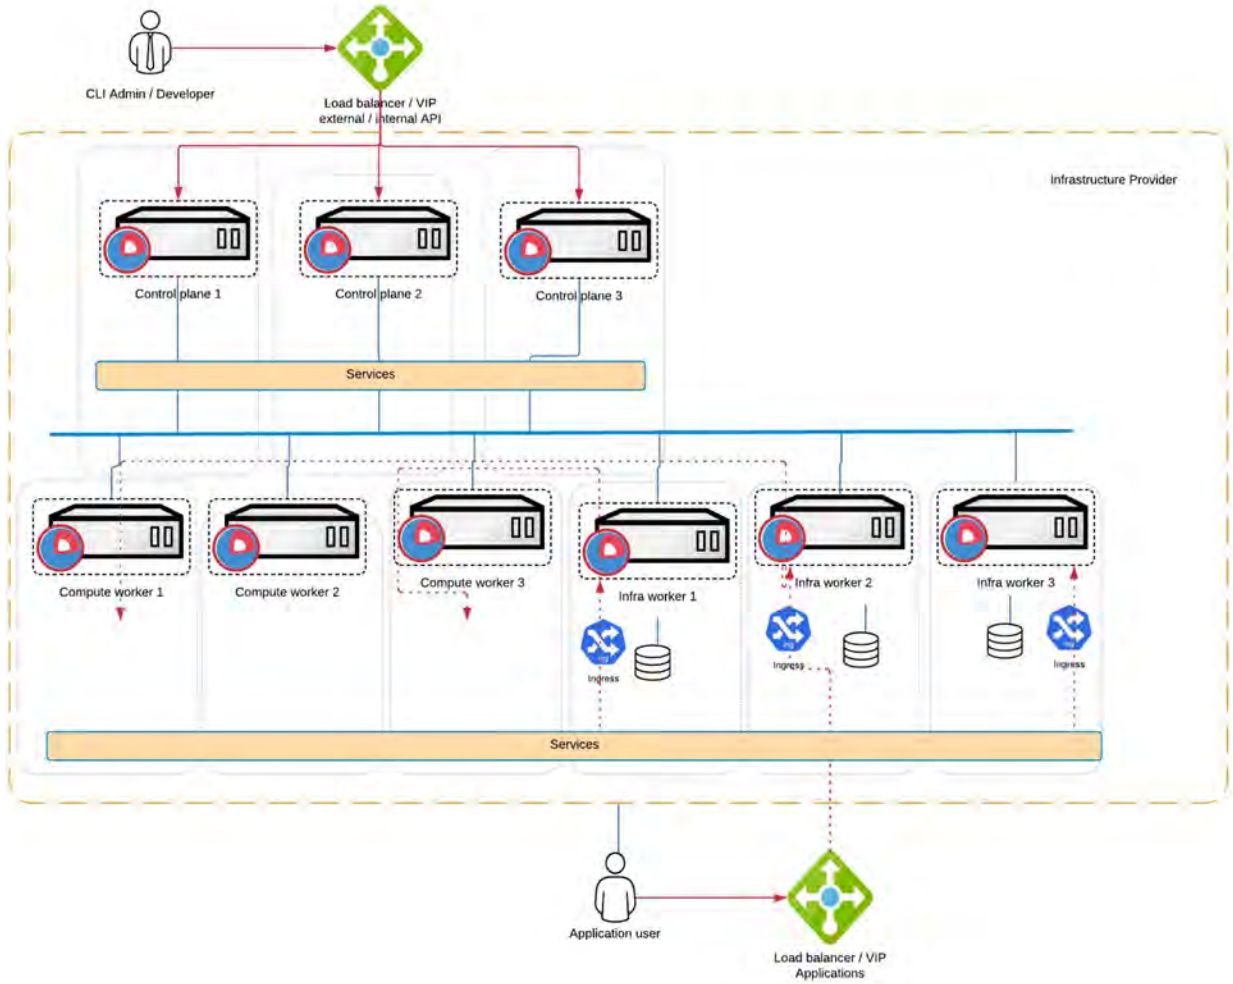

- **Control Plane (ex Master)** - No workloads, critical for API and etcd
- **Infra** - Hosts Ingress, monitoring, logging, registry etc.
- **Worker** - user workloads

Single node OpenShift - SNO

All-in-one - single node OpenShift deployment

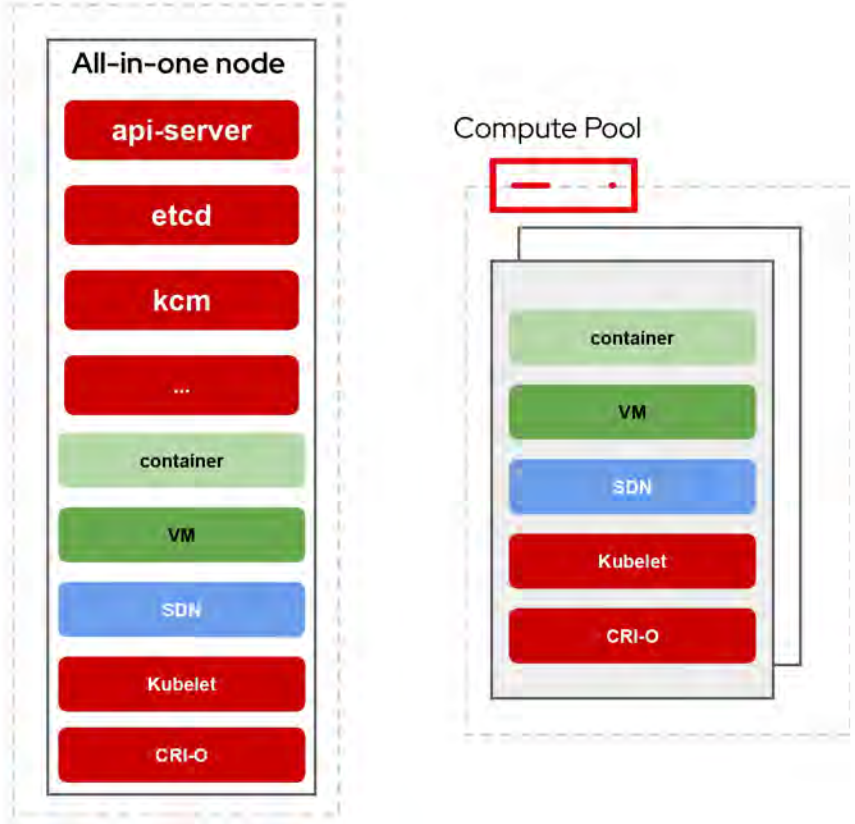

Two Node OpenShift with Arbiter (TNA)

Approach:

- Two node solution for cost sensitive customers
- Small arbiter node, running only 3d etcd instance
- Technically a three node cluster
- Arbiter Node is a regular node and could be used to run additional components/workload
- Arbiter node can be co-located (witness sled)
- Arbiter node has to be within <500msec max effective end to end latency (incl. Disc io)
- OCP Virtualization fully supported
- Hyperconverged Storage / SDS via Partners

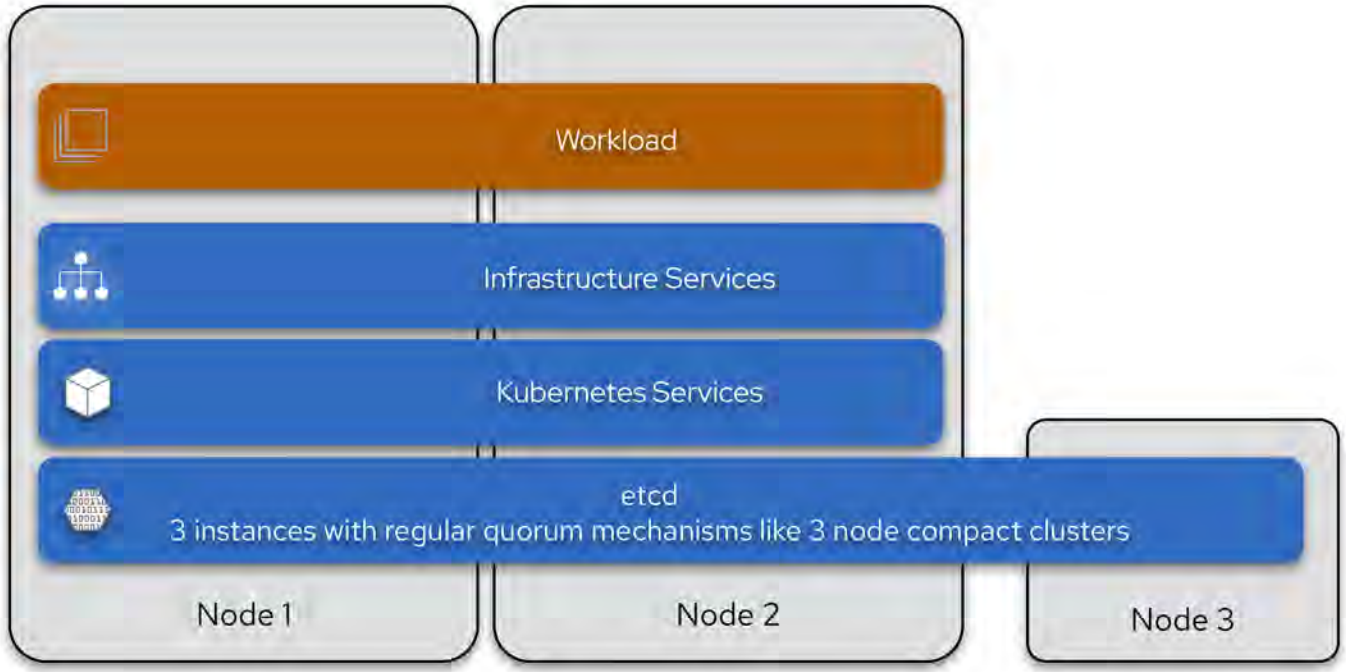

Timeline Targets:

- **V4.19** Technology Preview, X86+Arm Bare Metal only

Two Node OpenShift with Fencing (TNF)

Approach:

- True two node solution for cost sensitive customers
- Relies on proven RHEL-HA technologies (corosync, pacemaker) to provide etcd HA
- Uses fencing to protect against split brain situations: the surviving node power downs the failed node to guarantee consistency
- Requires a Base Management Controller (BMC) that supporters RedFish for fencing
- Node local storage supported (e.g. LVMS)

Timeline Targets:

- **V4.19** Technology Preview X86_64 only

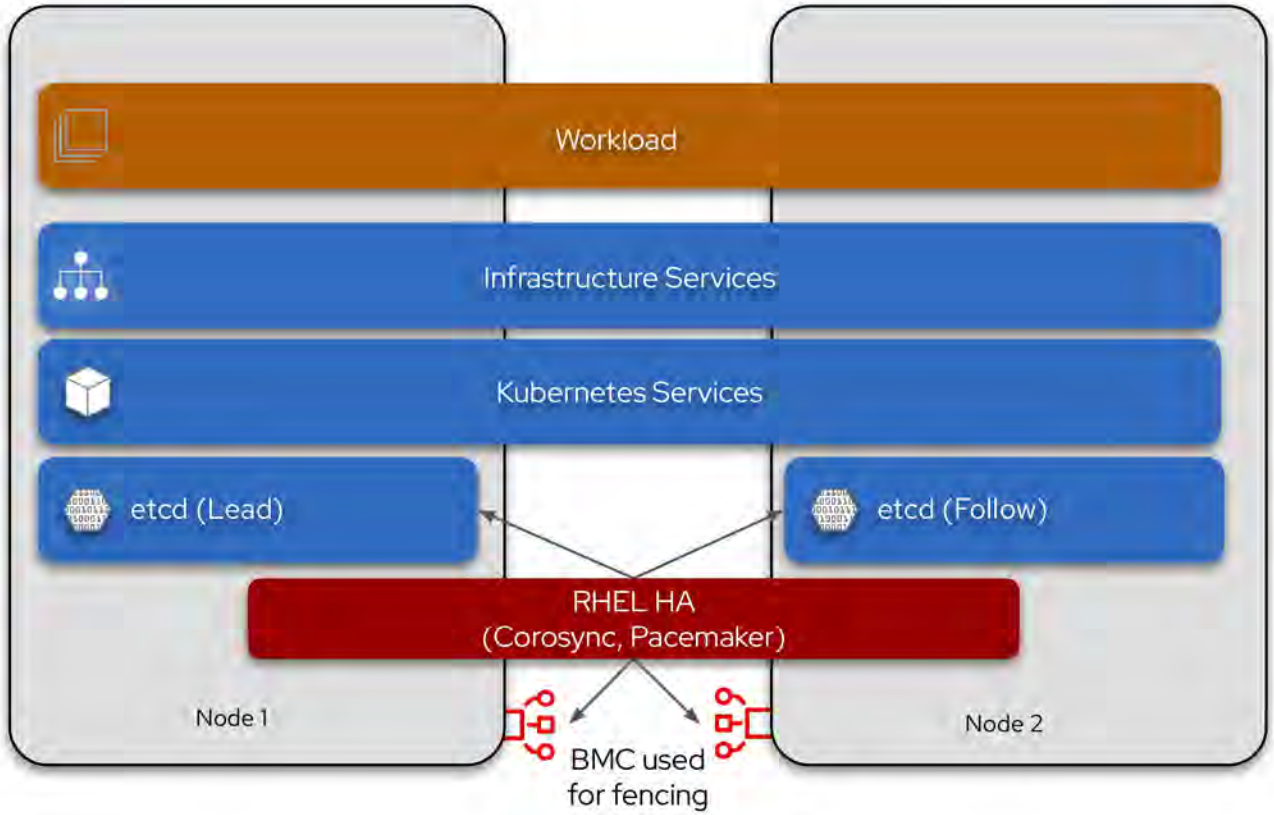

OpenShift Compact Cluster

Three nodes cluster - simple option for a HA deployment

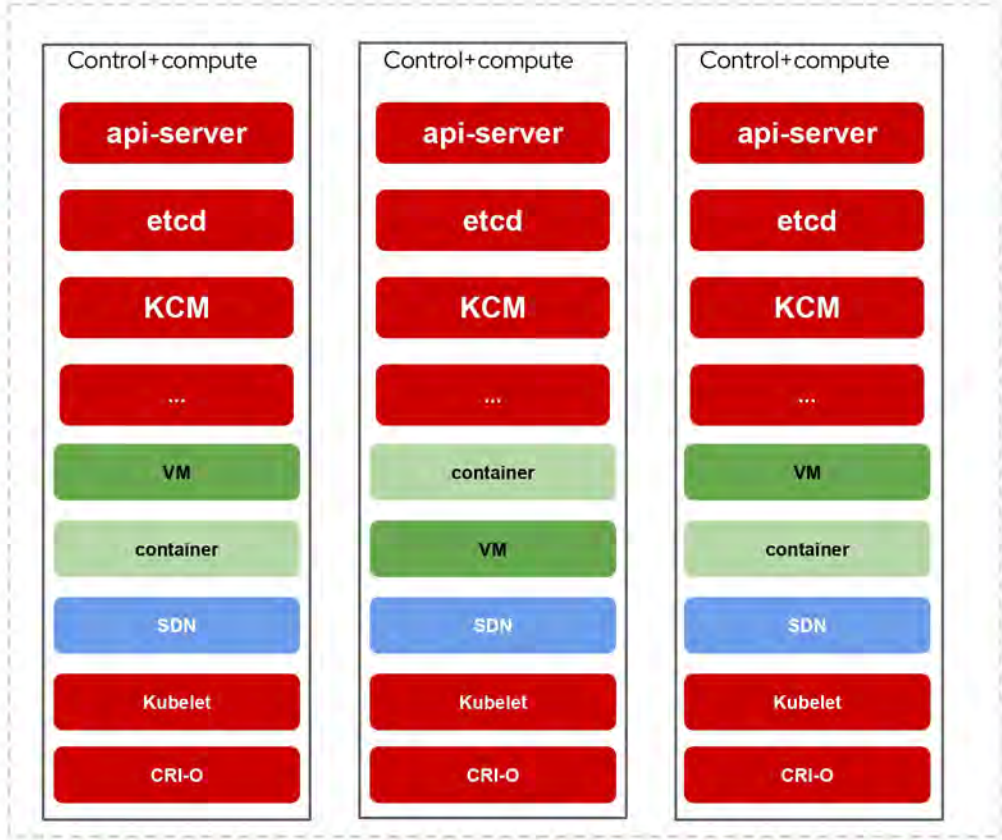

OpenShift Compact Cluster

Control-Plane (CP) incl. computes + dedicated computes

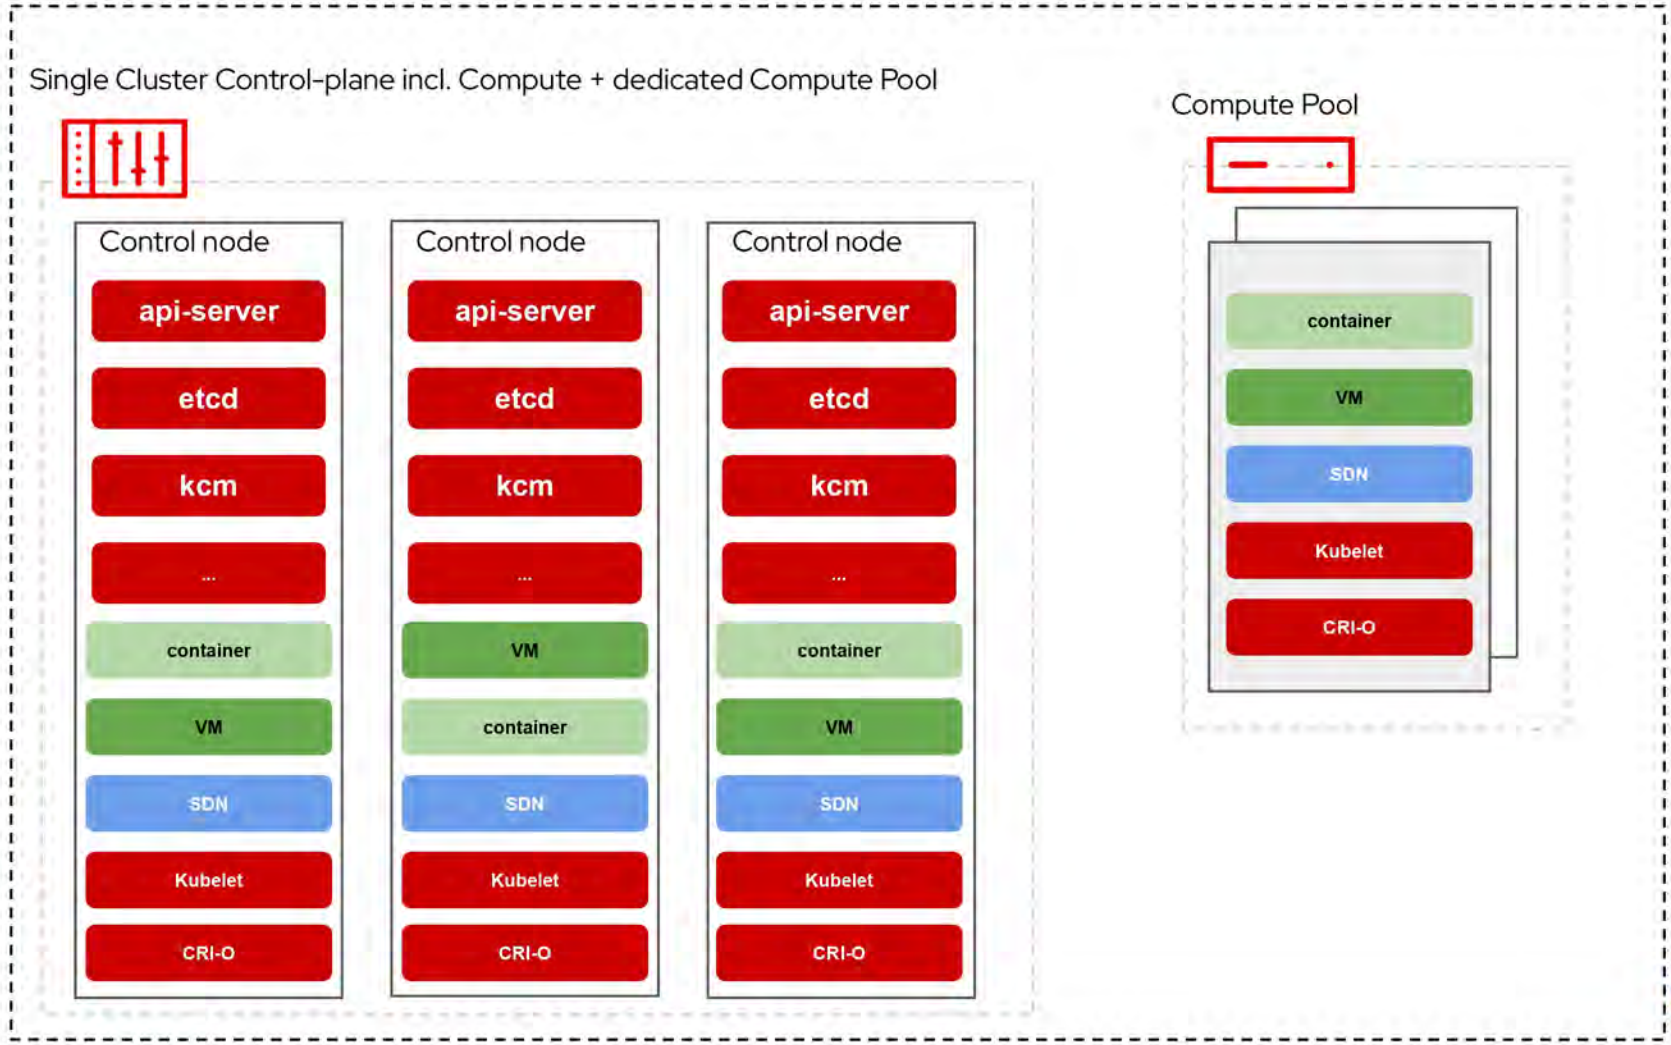

Red Hat | INDITEX

Standalone OpenShift

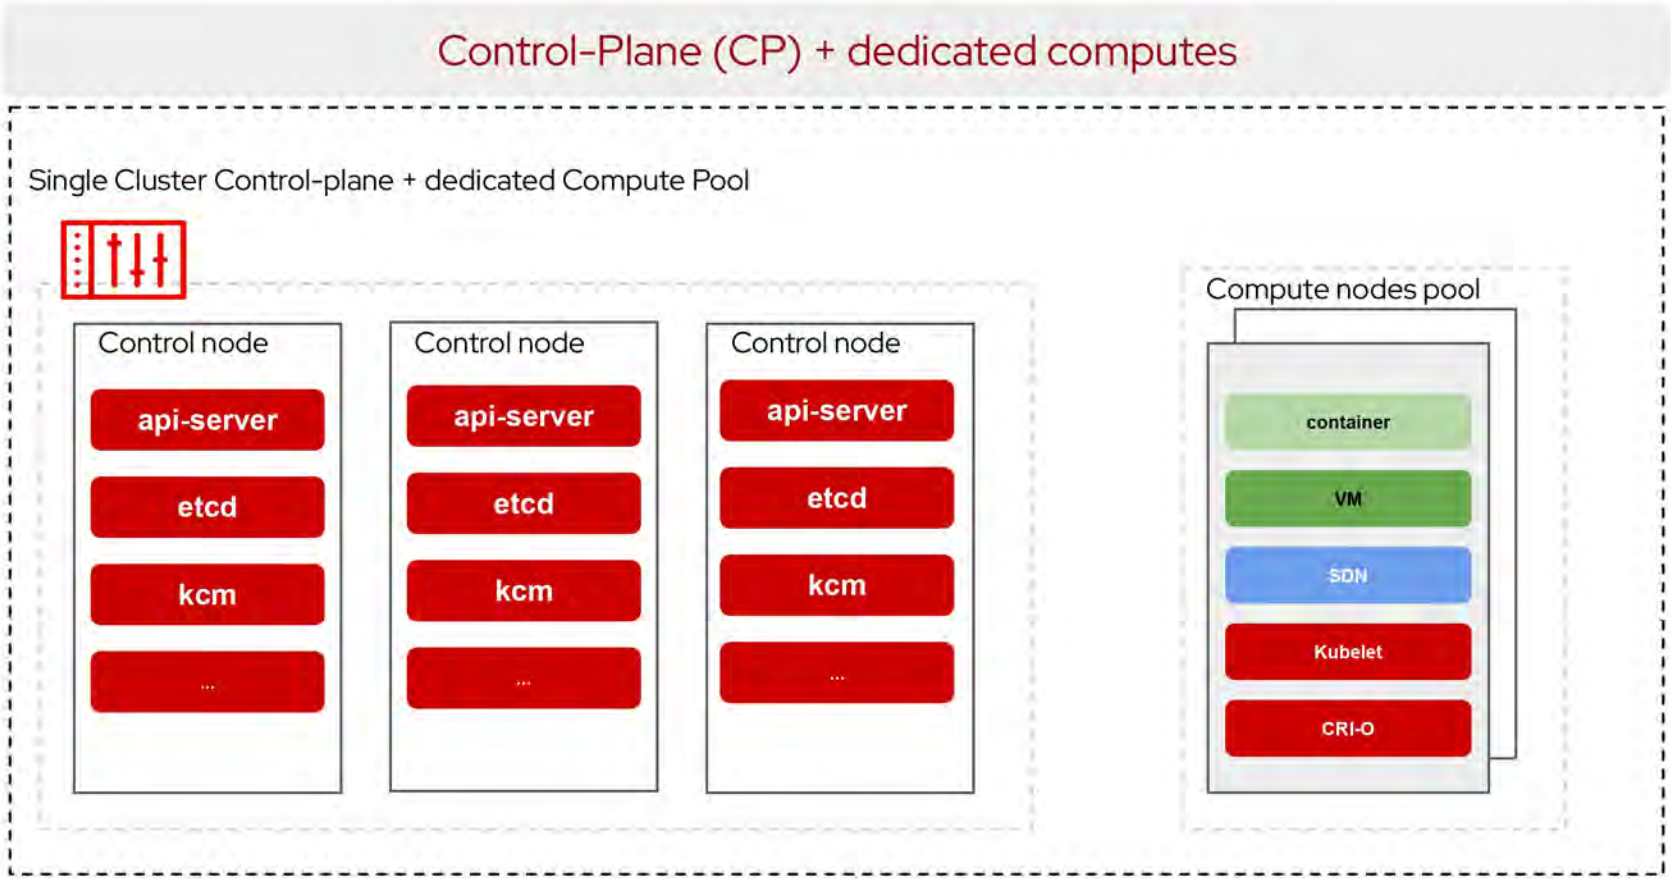

Utilizing OpenShift Virtualization to Consolidate OpenShift Clusters

Hosted Control Planes with KubeVirt provider

Increase Utilization of Infrastructure

- Consolidate multiple control planes to reduce unused and underutilized infrastructure
- Increase bare metal node utilization by hosting virtual worker nodes for multiple clusters

Reduce Dependency on Legacy Virtualization

- Eliminate the need to have legacy hypervisor layer to host your containerized infrastructure
- Underlying virtualization layer is included with hosted OpenShift cluster entitlements (no separate licensing needed)

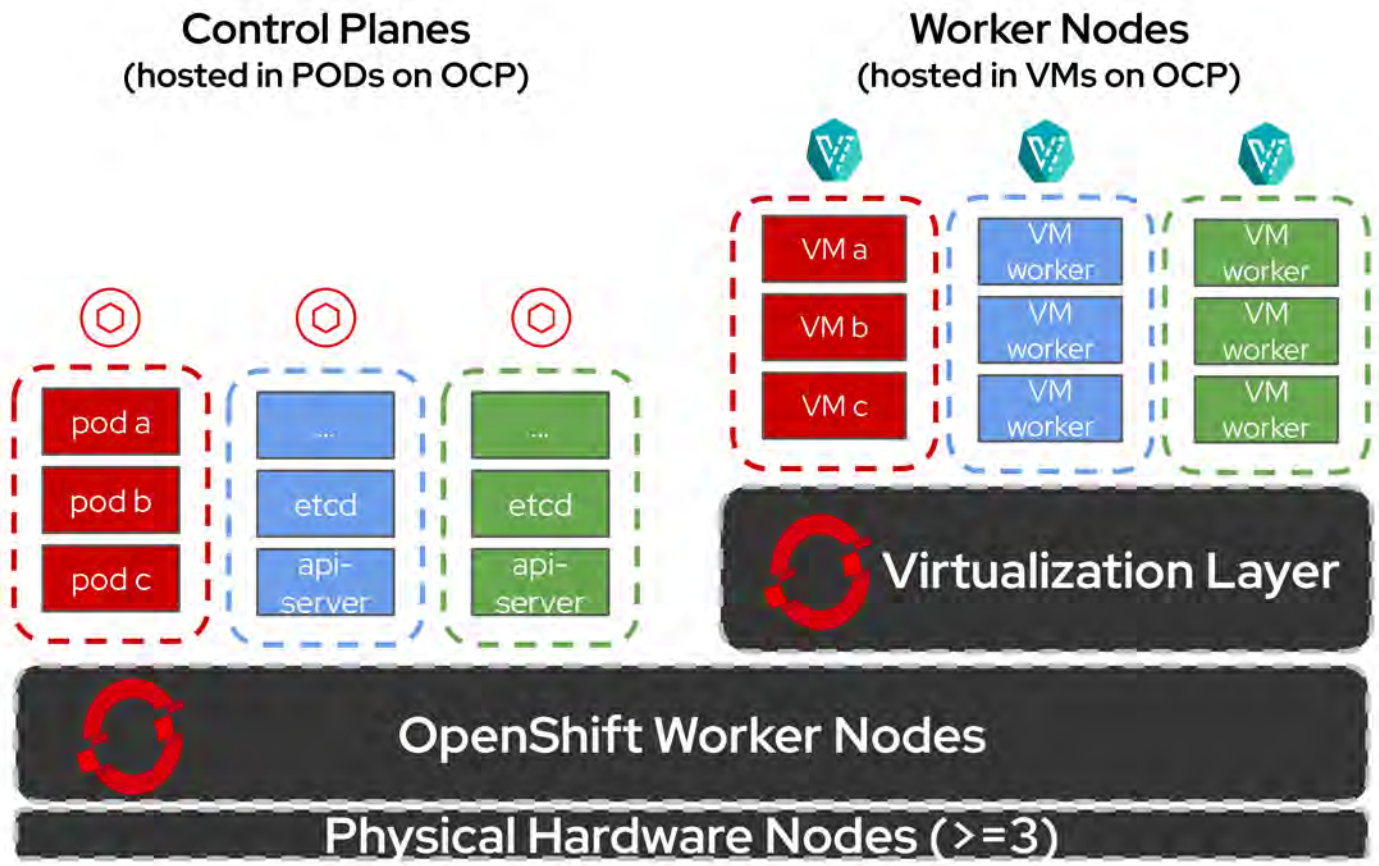

Red Hat | INDITEX

Utilizing OpenShift Virtualization to Consolidate OpenShift Clusters

Hosted Control Planes with KubeVirt provider

? Are there other ways of deployment?

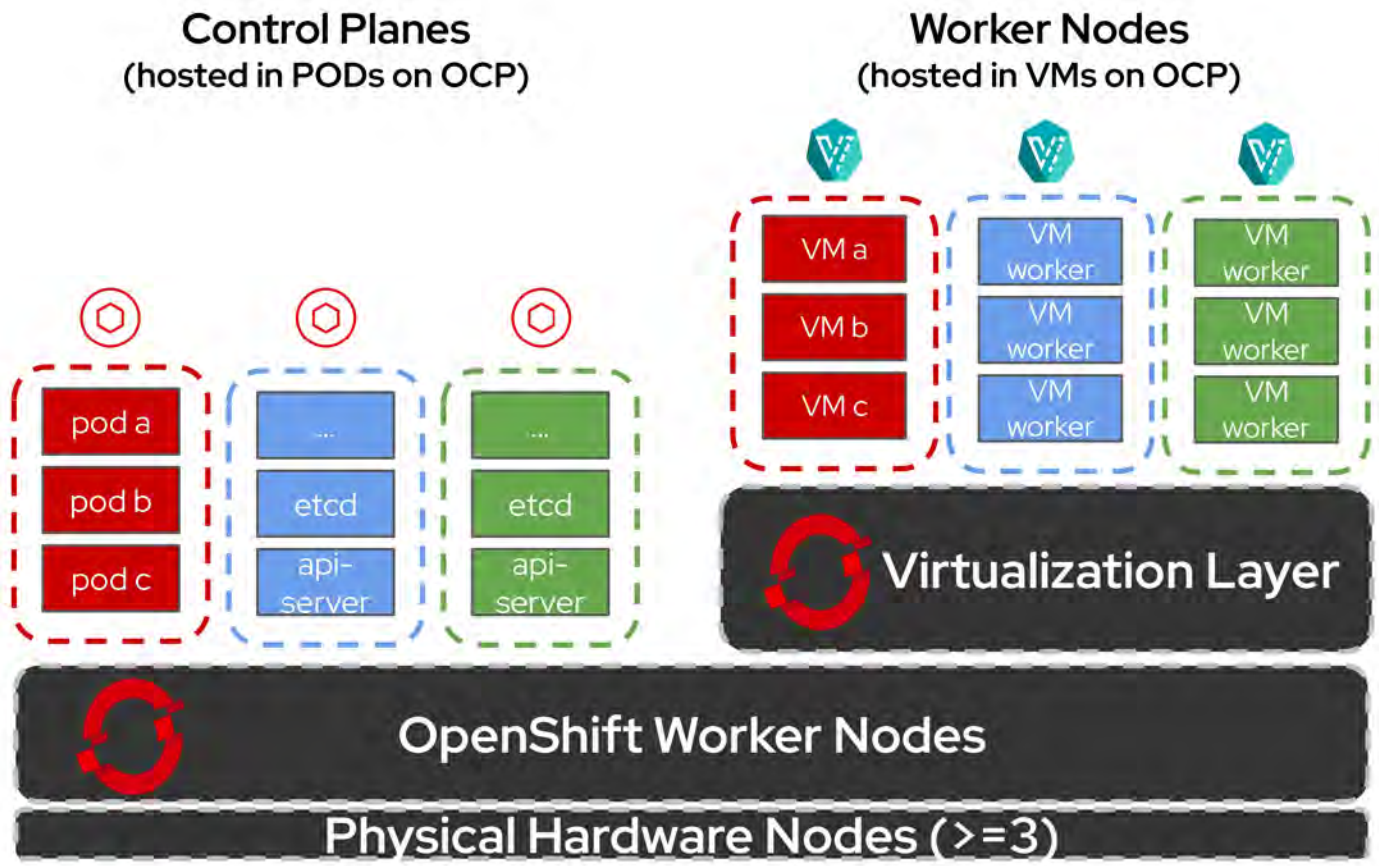

Utilizing OpenShift Virtualization to Consolidate OpenShift Clusters

Hosted Control Planes with KubeVirt provider

? Are there other ways of deployment?

- CP PODs and Worker VMs do not need to reside on the same cluster

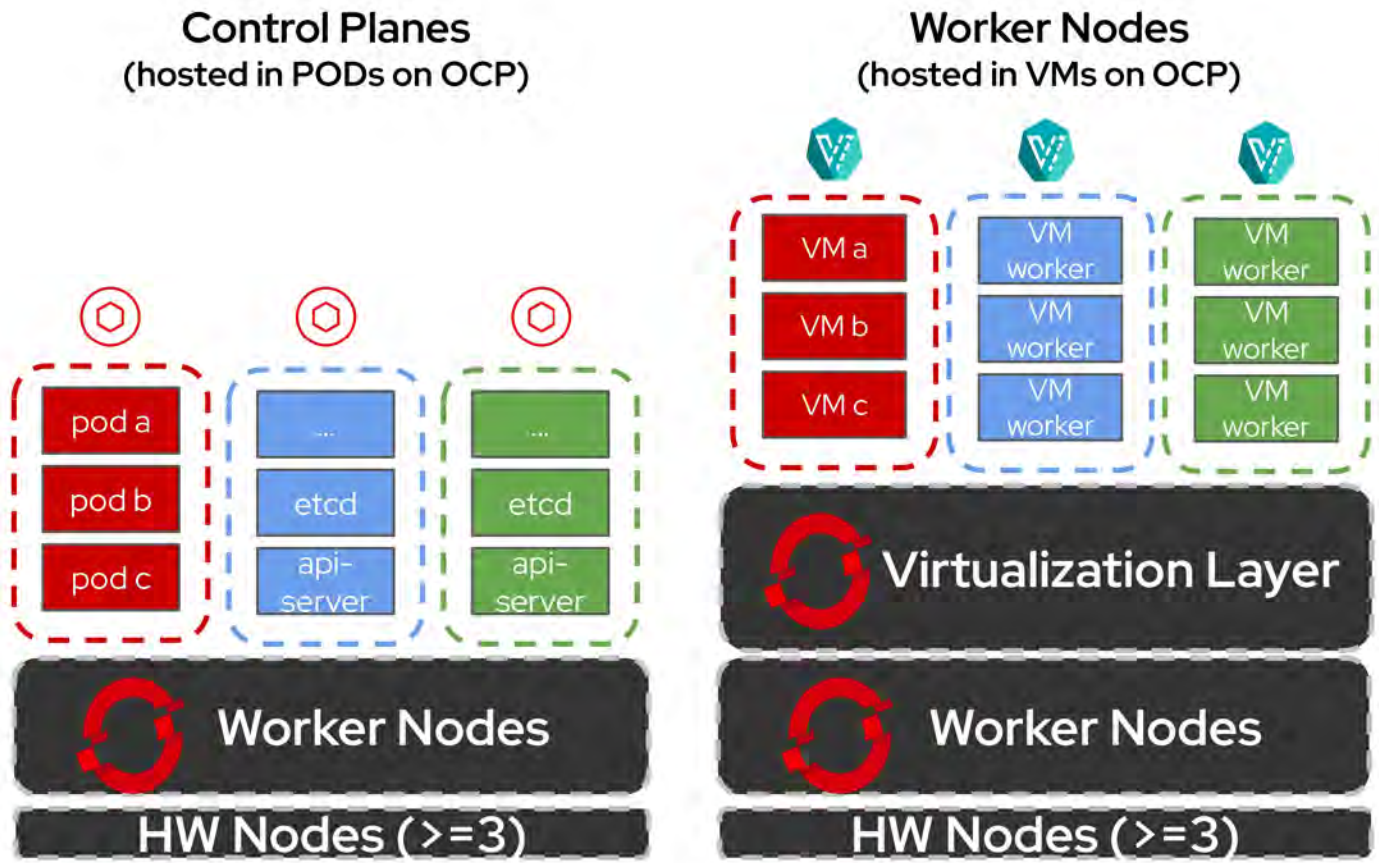

Utilizing OpenShift Virtualization to Consolidate OpenShift Clusters

Hosted Control Planes with KubeVirt provider

? Are there other ways of deployment?

- CP PODs and Worker VMs do not need to reside on the same cluster
- Worker VMs on Hyperscalers

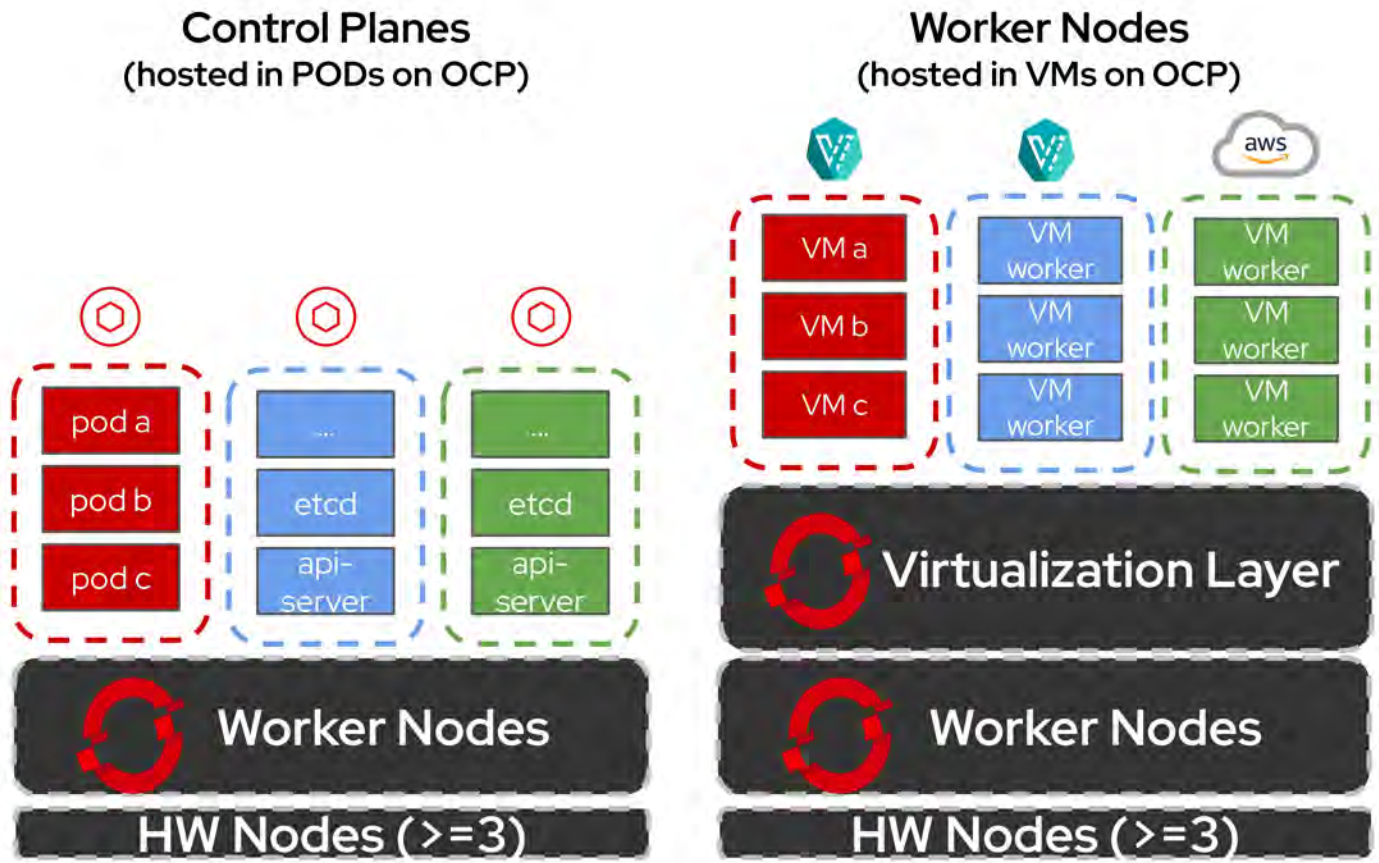

Red Hat | INDITEX

Utilizing OpenShift Virtualization to Consolidate OpenShift Clusters

Hosted Control Planes with KubeVirt provider

? Are there other ways of deployment?

- CP PODs and Worker VMs do not need to reside on the same cluster
- Worker VMs on Hyperscalers
- Workers can even be Baremetal

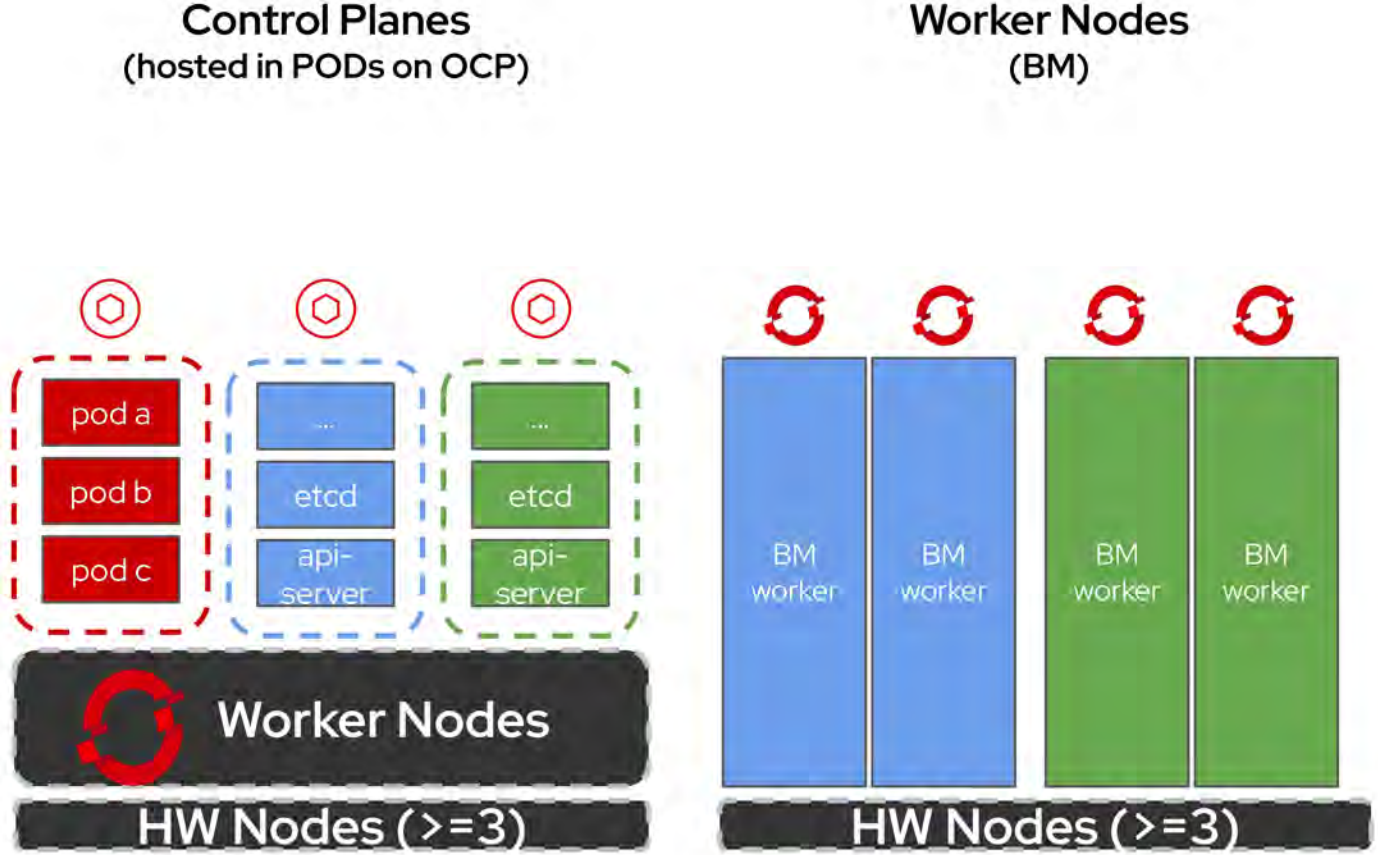

Red Hat | INDITEX

Utilizing OpenShift Virtualization to Consolidate OpenShift Clusters

Hosted Control Planes with KubeVirt provider

? Are there other ways of deployment?

- CP PODs and Worker VMs do not need to reside on the same cluster
- Worker VMs on Hyperscalers
- Workers can even be Baremetal
- Or even mixed

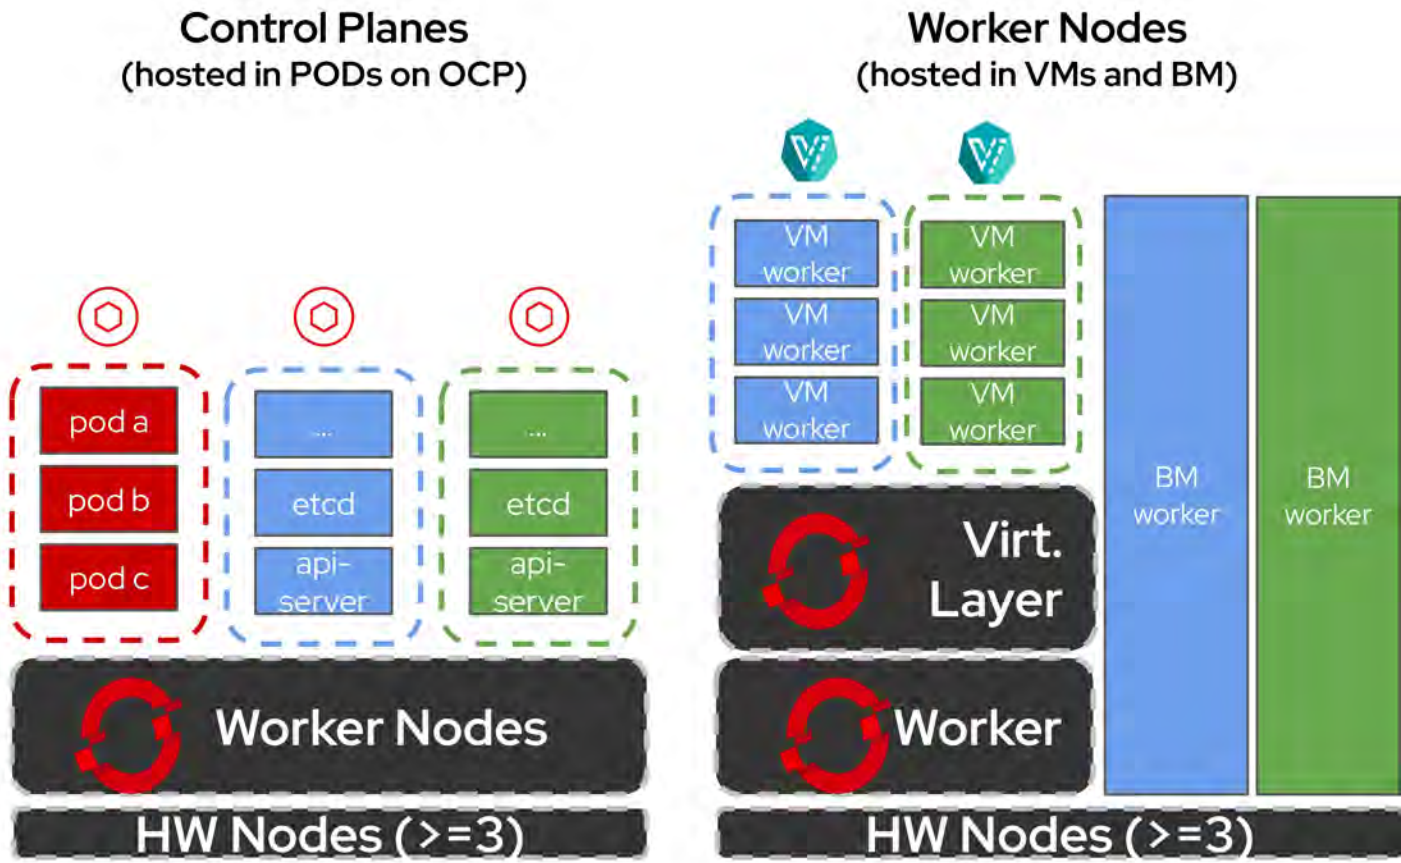

Why HCP is a game changer?

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

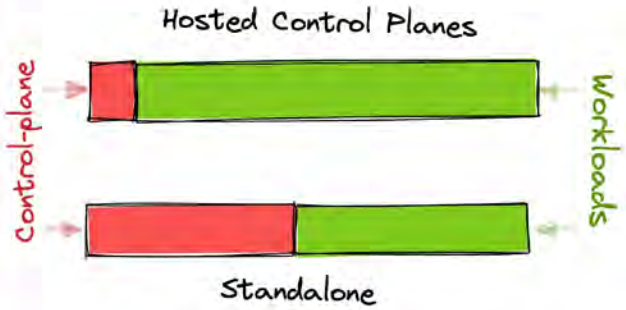

Networking

Levels of Network Integration

Can we use our VMware/RHV/OpenStack network layout?

- integrated OpenShift SDN (aka Pod Network)
- secondary SDN (aka OVN on secondary Networks)
  - This is the network type you will use when coming from OpenStack.
- traditional direct bridged VLAN networks (aka Multus)
  - This is the network type you use when coming from VMware/RHV.
- User Defined Network (UDN primary or secondary)
  - This is the network type you use when coming from VMware-NSX.

All network types can be used together on a single VM.

Note: This is not an exhaustive list of ISV partners, with new partners being added all the time

Virtual Machine Networking

- Virtual machines optionally connect to the standard pod network
  - OVNKubernetes (default since 4.12)
  - Partners, such as Calico, are also supported
- Additional network interfaces accessible via Multus:
  - Bridge, SR-IOV, OVN secondary networks
  - VLAN and other networks can be created at the host level using nmstate
- When using at least one interface on the default SDN, Service, Route, and Ingress configuration applies to VM pods the same as others

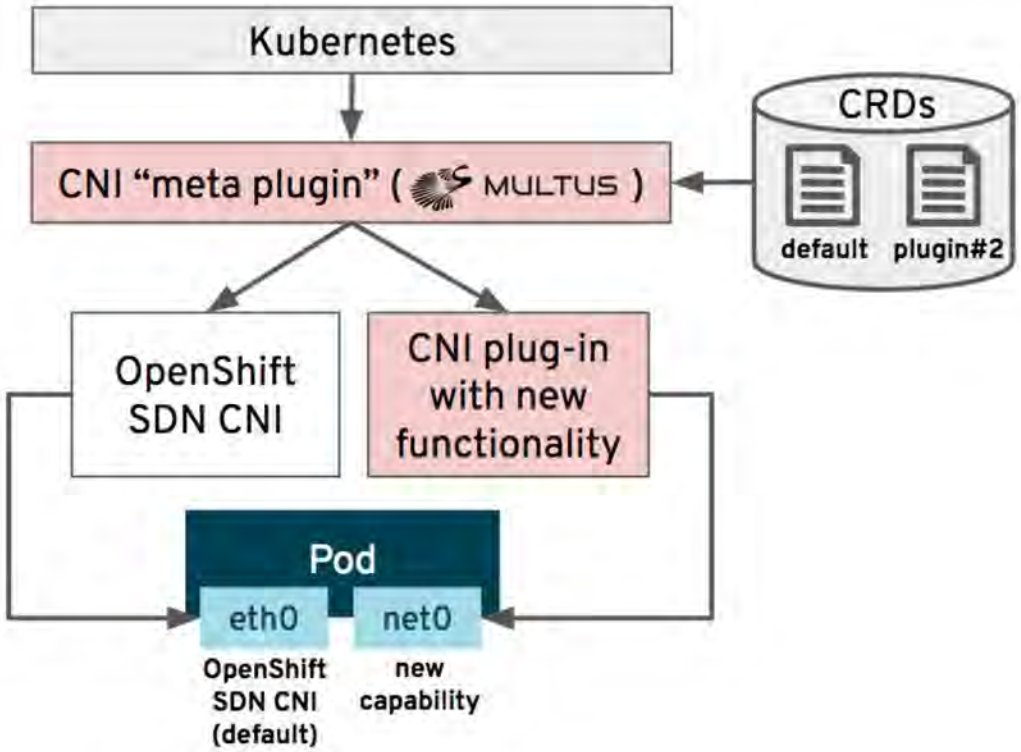

Default host network configuration

- Pod, service, and machine network are configured by OpenShift automatically
  - Use installer parameters (nmstate) for configuration of the NIC at install - `eth0` in the example to the right. NIC config options:
    - Static
    - DHCP
  - No redundancy
  - No high availability
  - No link aggregation

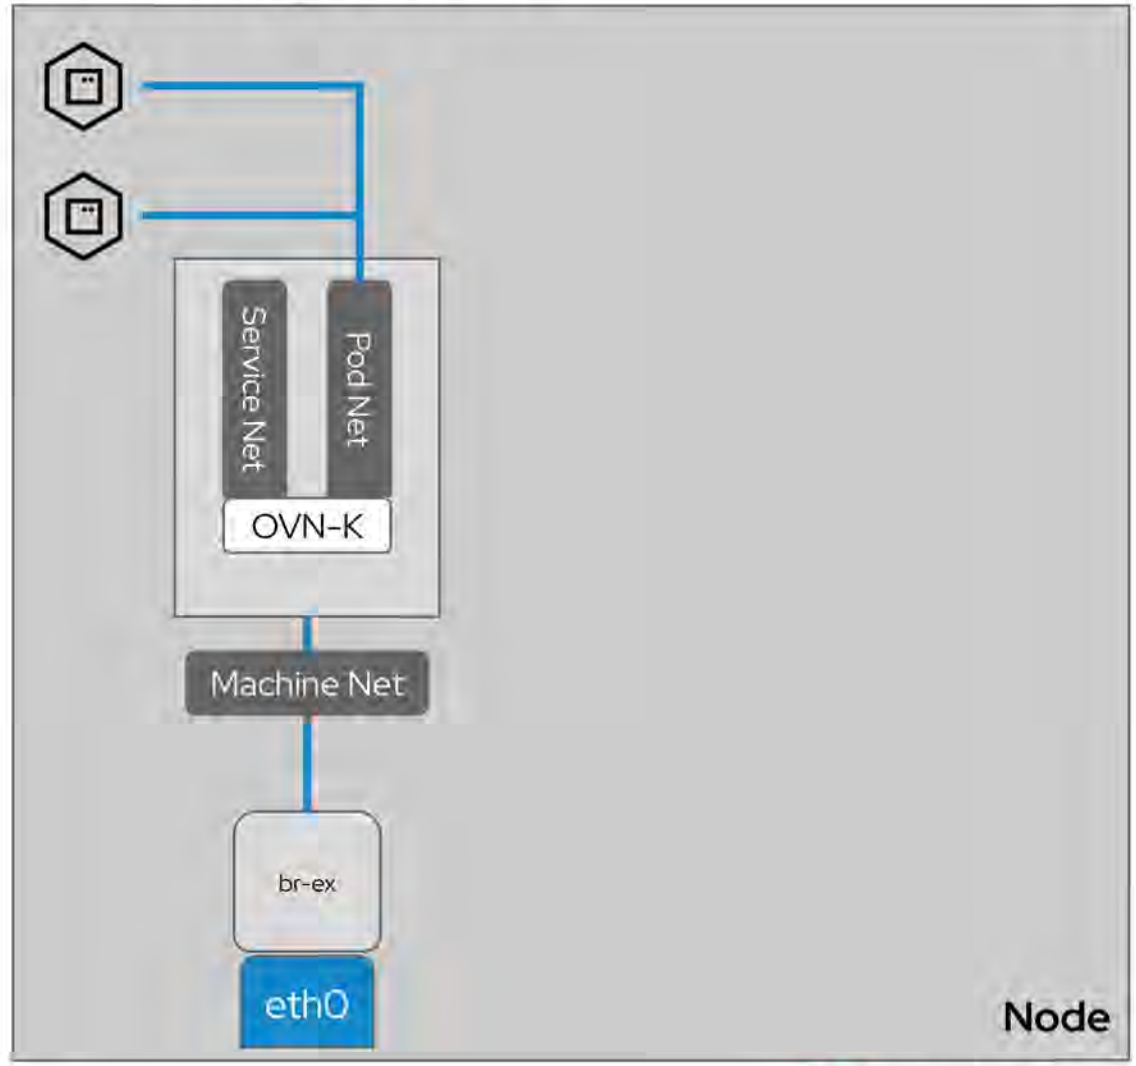

HA for Machine Net

- Use installer parameters (nmstate) for configuration of the NIC at install - `bond0` in the example to the right. Supported bonding modes: *
  - `1 - active-backup` (No configuration required on the switch)
  - `2 - balance-xor` (Requires static EtherChannel enabled)
  - `4 - 802.3ad` (Requires LACP-negotiated EtherChannel enabled)

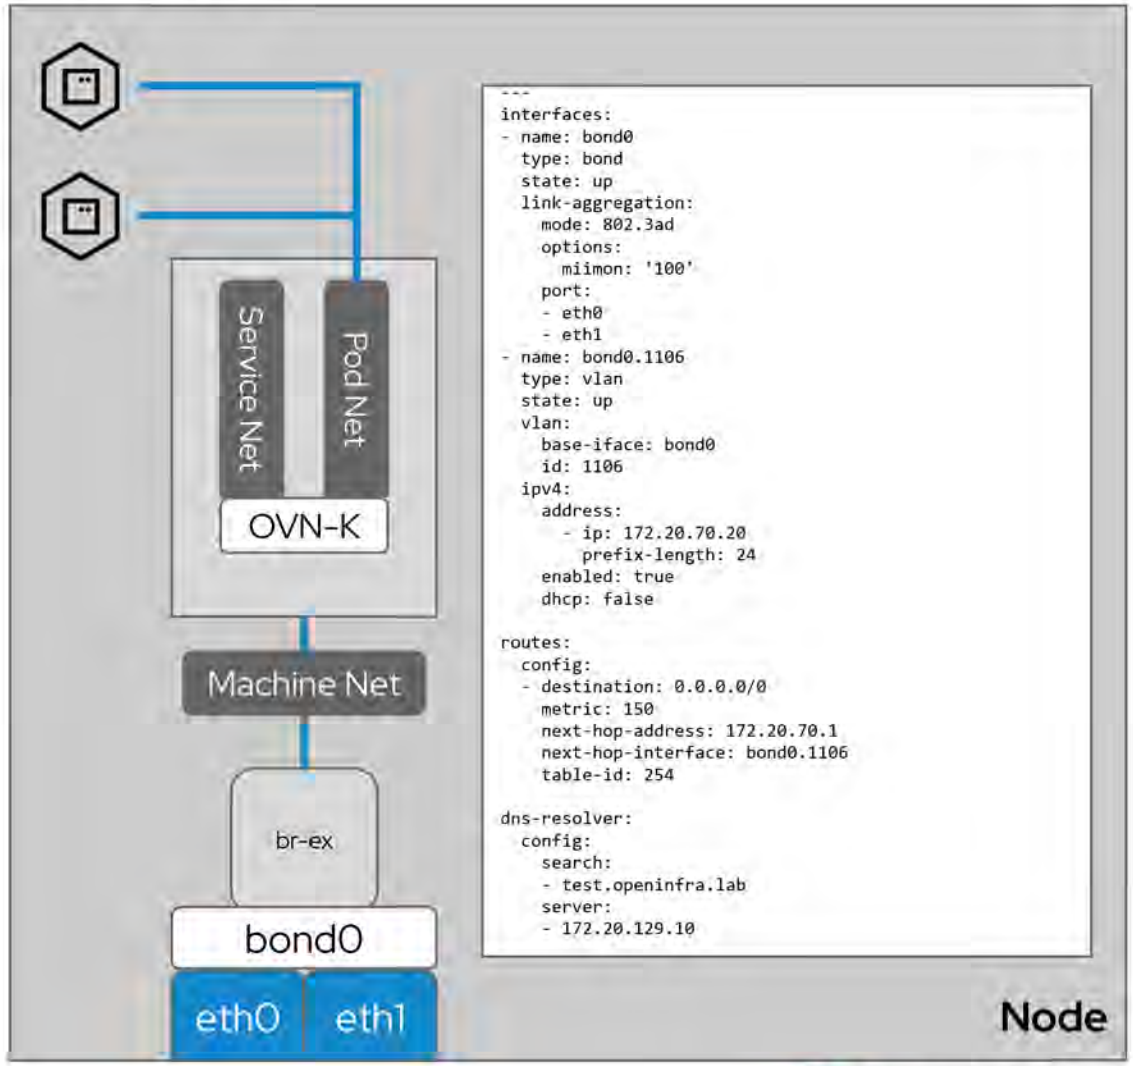

* other modes and why to avoid:

0 - balance-rr, 3 - broadcast: not suitable for TCP, packet order not guaranteed

5 - balance-tlb, 6 - balance-alb: both do not work with bridges - because it re-writes the source MAC

Operator Prerequisites

- In order to configure additional NICs (or bonds or bridges) on the host you will need to install the `nmstate` operator
  - `nmstate` is a declarative API for host network management
  - Uses `NodeNetworkConfigurationPolicy (NNCP)` CRD to accomplish its declarative state
  - `NNCP` embeds `nmstate` yaml configuration which declares the end state network configuration for network interfaces

VLAN tagged Networks for VMs - NAD

- Use the Network Attachment Definition CRD to configure additional VLAN tagged VM networks
  - `vlan-1124-br-data` in the example on the right
    - `vlanID:` this VLAN is expected to be tagged on the uplink switch trunk port
    - `subnets:` (only supported in OCP 4.16+) network subnet to be used by the NAD - implies port security and spoofing is on
    - `excludeSubnets:` optional but useful to exclude infrastructure IPs like gateways
    - `topology:` localnet is used for networks with a egress port via OVS bridge

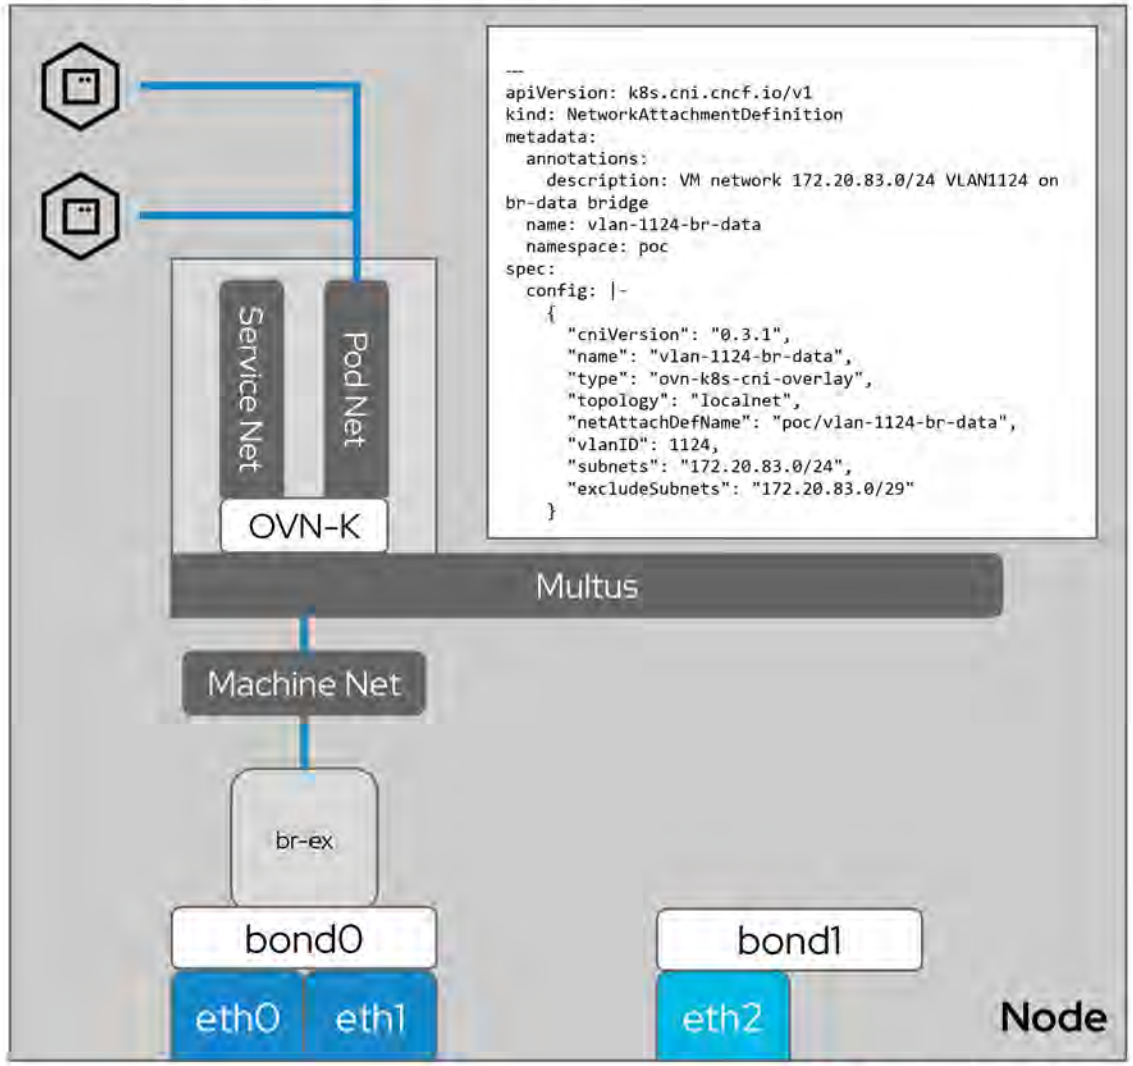


```
---
apiVersion: k8s.cni.cncf.io/v1
kind: NetworkAttachmentDefinition
metadata:
  annotations:
    description: VM network 172.20.83.0/24 VLAN1124 on
br-data bridge
  name: vlan-1124-br-data
  namespace: poc
spec:
  config: |-
    {
      "cniVersion": "0.3.1",
      "name": "vlan-1124-br-data",
      "type": "ovn-k8s-cni-overlay",
      "topology": "localnet",
      "netAttachDefName": "poc/vlan-1124-br-data",
      "vlanID": 1124,
      "subnets": "172.20.83.0/24",
      "excludeSubnets": "172.20.83.0/29"
    }
```

User Defined Networks: Introduction

Flexible Network Configurations for the users

Workload/Tenant Isolation

Ability to group different types of applications in different isolated networks that cannot talk to each other

Flexible Network Topologies

Ability to create **layer3** or **layer2** or **localnet** type networks which can act as either **primary** and **secondary** networks for your pods

Overlapping podIPs

Ability to create Multiple Networks in your cluster with same pod Subnet range thereby possible to have copies of setups!

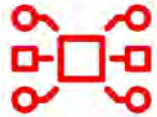

Kubernetes APIs supported!

Primary UDNs will have full support for **services, network policies, admin network policies, egressIPs**, and **ocp-routes**

User Defined Networks Use Case

Native Namespace Isolation: Each namespace can be 1 unique UDN

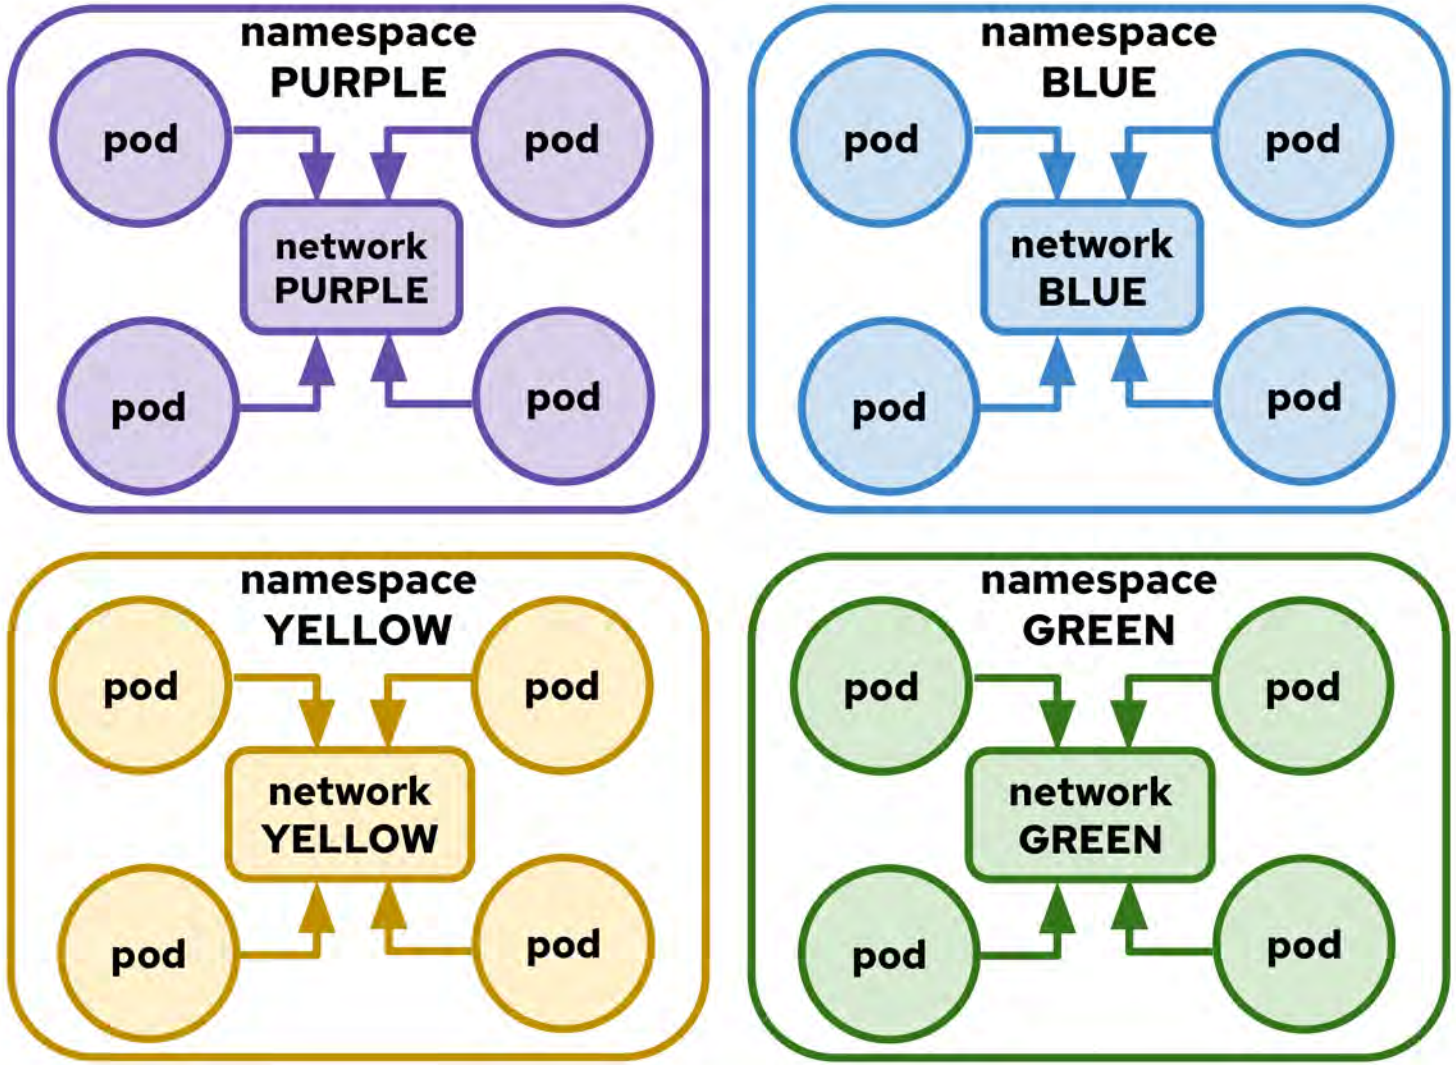

Networks purple, blue, yellow and green are disconnected networks (logically isolated islands) which guarantees native isolation of these namespaces

Red Hat | INDITEX

User Defined Networks Use Case

Native Tenant Isolation: Multiple namespaces can be connected to be part of the same UDN

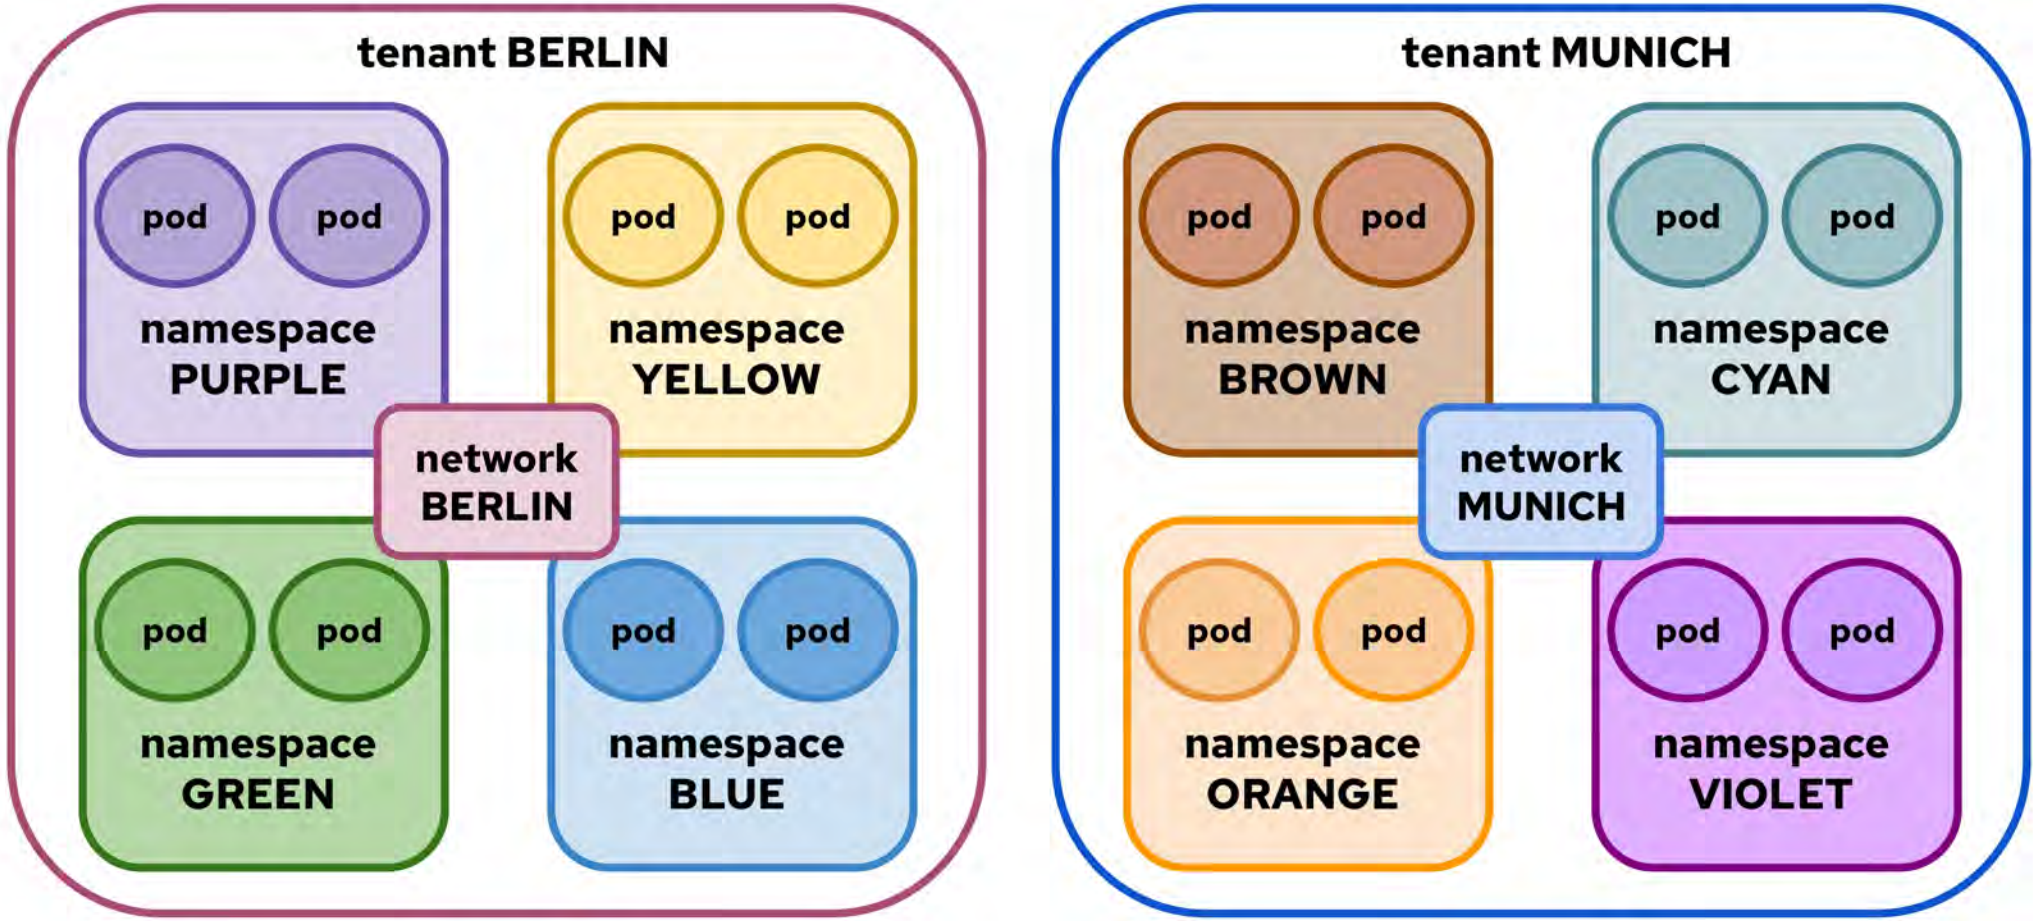

Tenants BERLIN and MUNICH are disconnected networks that guarantees native isolation

Namespaces purple, yellow, green and blue can talk to each other since they are all connected to BERLIN network but cannot talk with tenant MUNICH's namespaces

Network Policy

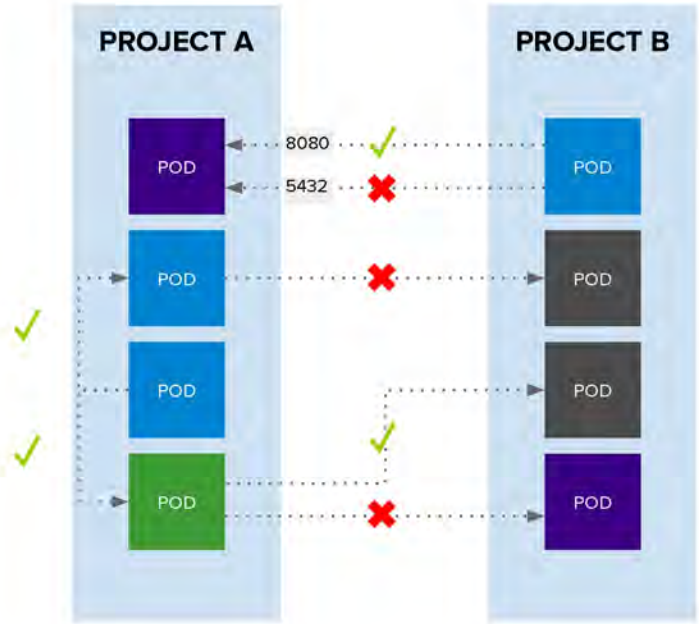

Example Policies

- Allow all traffic inside the project
- Allow traffic from green to gray
- Allow traffic to purple on 8080

```
apiVersion: extensions/v1beta1
kind: NetworkPolicy
metadata:
  name: allow-to-purple-on-8080
spec:
  podSelector:
    matchLabels:
      color: purple
  ingress:
  - ports:
    - protocol: tcp
      port: 8080
```

Micro-Segmentation for Virtual Machines made easy!

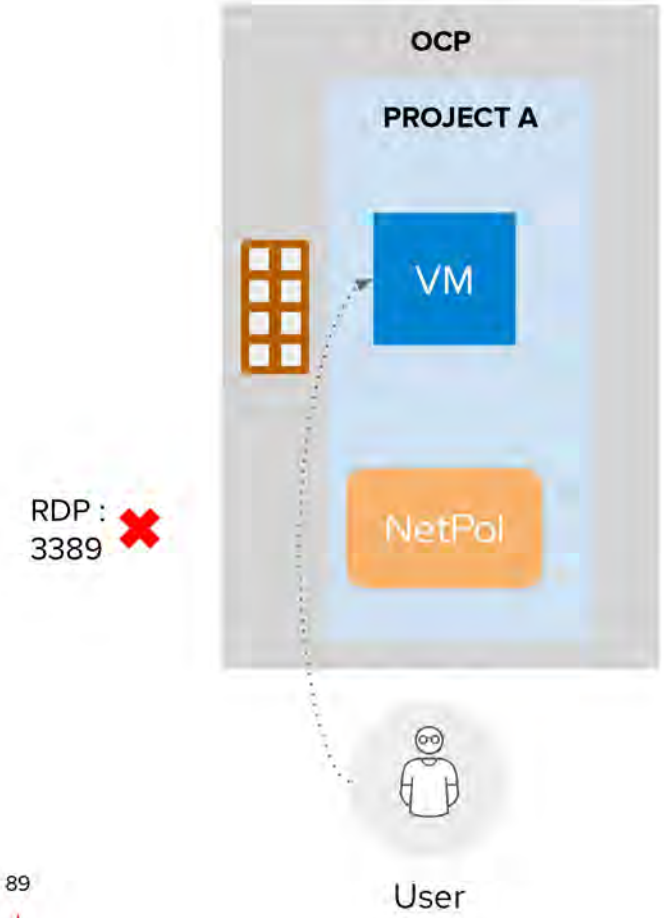


```yaml
kind: NetworkPolicy
apiVersion: networking.k8s.io/v1
metadata:
  name: deny-rdp-from-outside
spec:
  podSelector: {}
  ingress:
    - ports:
        - protocol: TCP
          port: 3389
      from:
        - podSelector: {}
```

Example Policies

- Applies to all pods in rguske-demo (because of `podSelector: {}`).
- Ingress traffic on port 3389 is only allowed if it comes from pods within the same namespace.
- External sources (other namespaces, nodes, or outside cluster) are blocked on port 3389.

Storage & Backup

Storage Integration - external Storage

Can we use our old storage system?

OpenShift relies on CSI (Container Storage Interface) drivers:

- come often as operators
- different ones can be mixed and even stacked
- need to support RWX feature (aka multi-attach) for live migration
- SAN sometimes need an additional layer for RWX (e.g. IBM Fusion Access for SAN, Arctera Infoscale, ~~ODF~~)

DELL Technologies

Red Hat OpenShift Data Foundation

HITACHI Inspire the Next

INFINIDAT

NetApp

portworx by Pure Storage

IBM

Hewlett Packard Enterprise

Note: This is not an exhaustive list of ISV partners, with new partners being added all the time

Storage Integration - SDS

- ODF (OpenShift Data Foundation)
- IBM FDF (Fusion Data Foundation)
- Portworx (by Pure Storage)
- Dell Powerflex

Note: This is not an exhaustive list of ISV partners, with new partners being added all the time

What does storage for VMware look like?

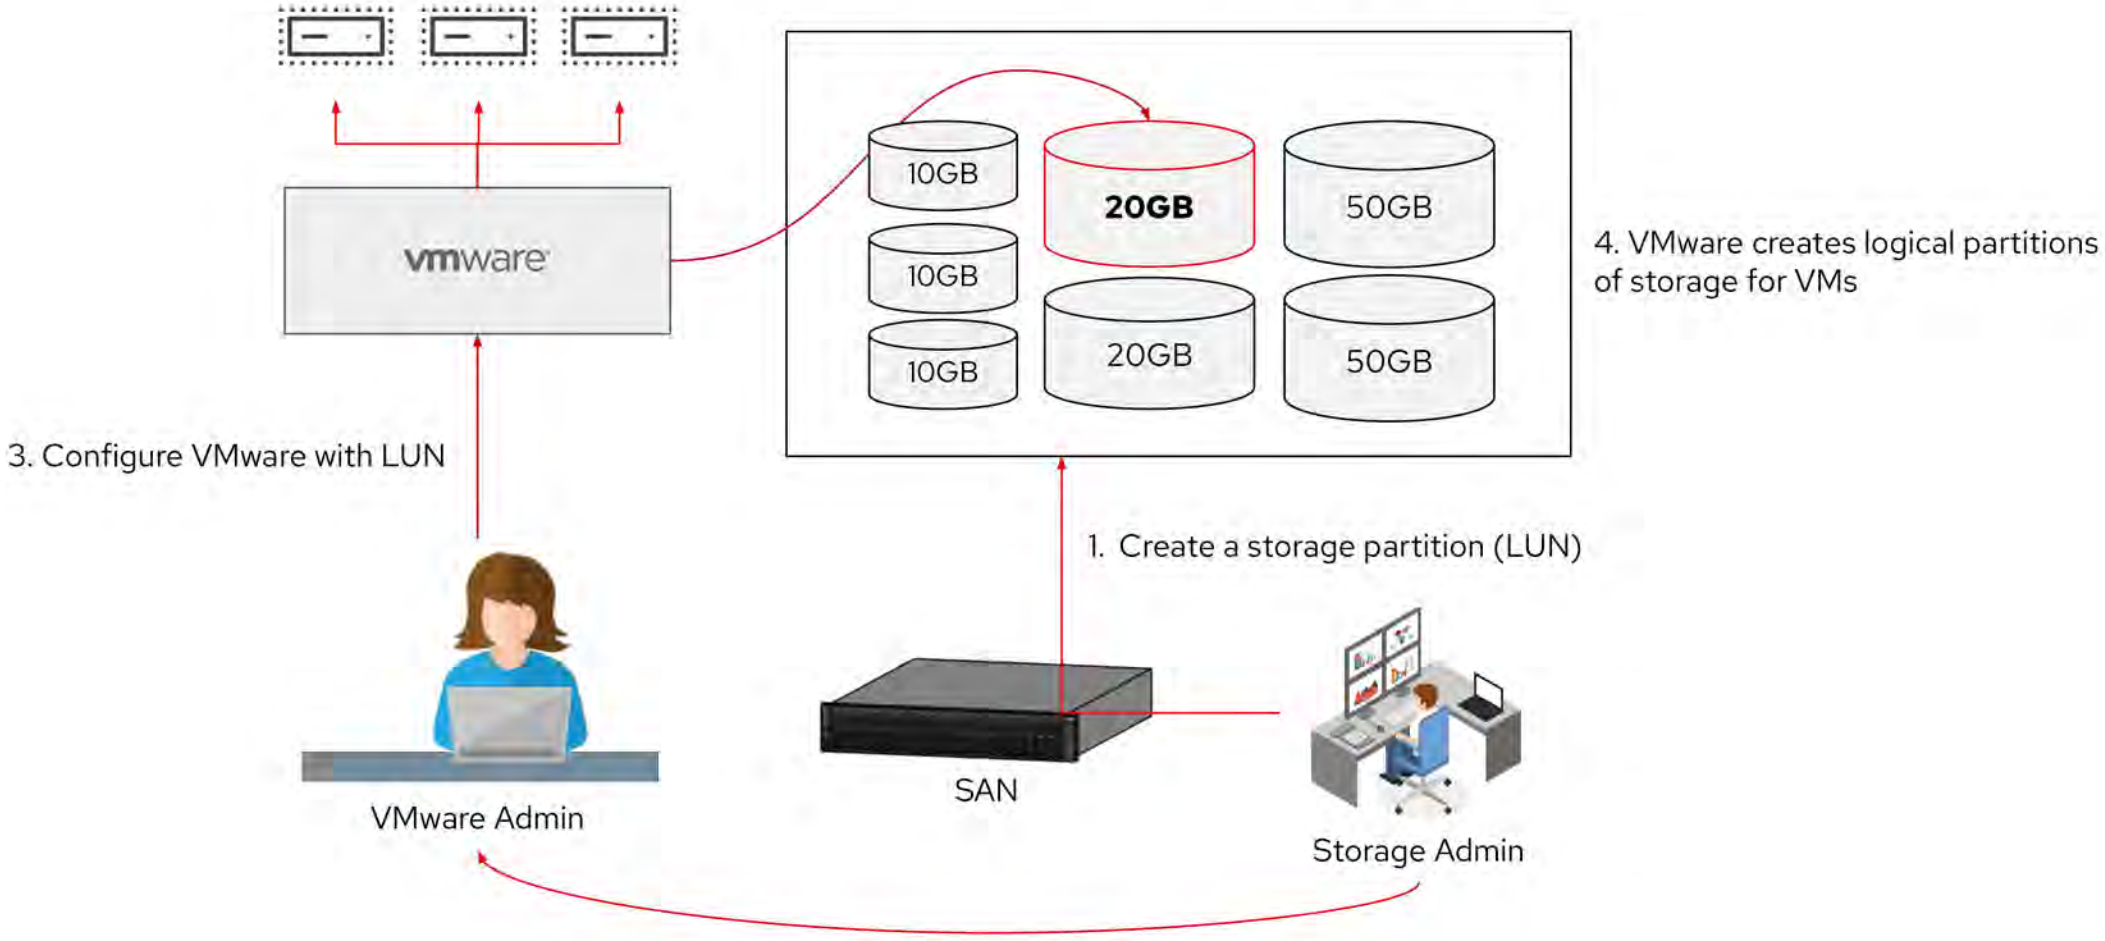

How does storage for Kubernetes differ?

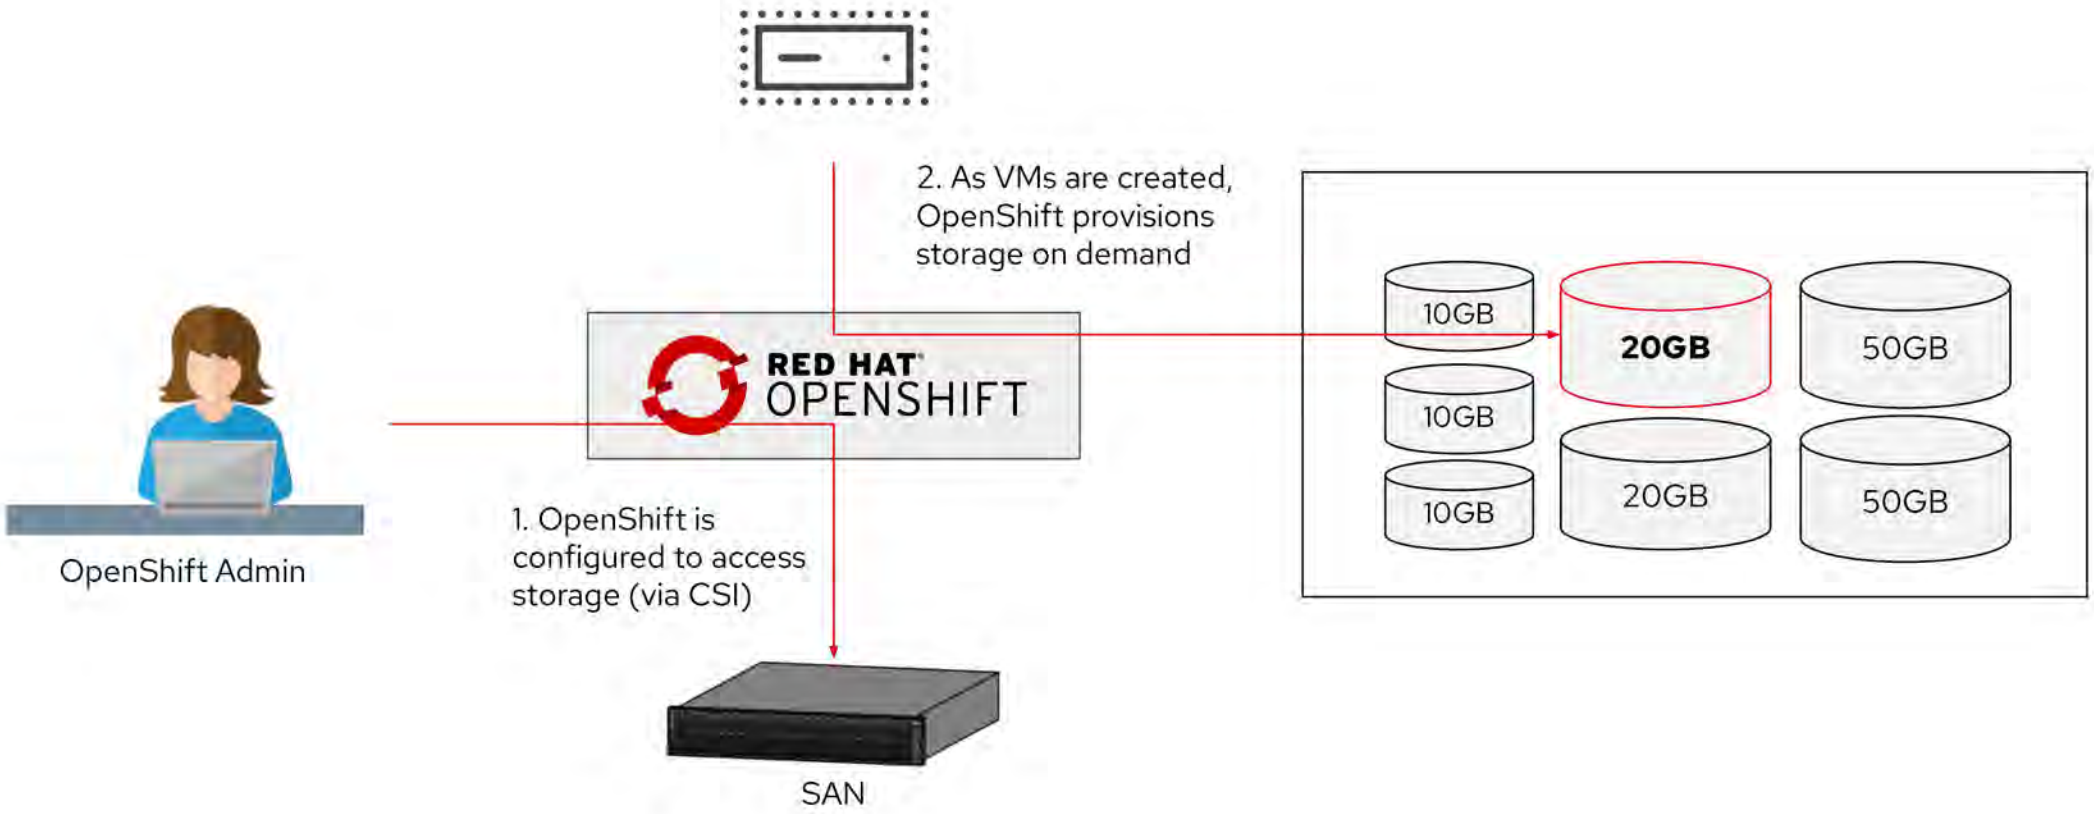

Kubernetes Storage Terminology

- STORAGE CLASS: "class" of storage service
  - Abstraction of back-end storage systems
  - OCP Clusters can have multiple storage classes

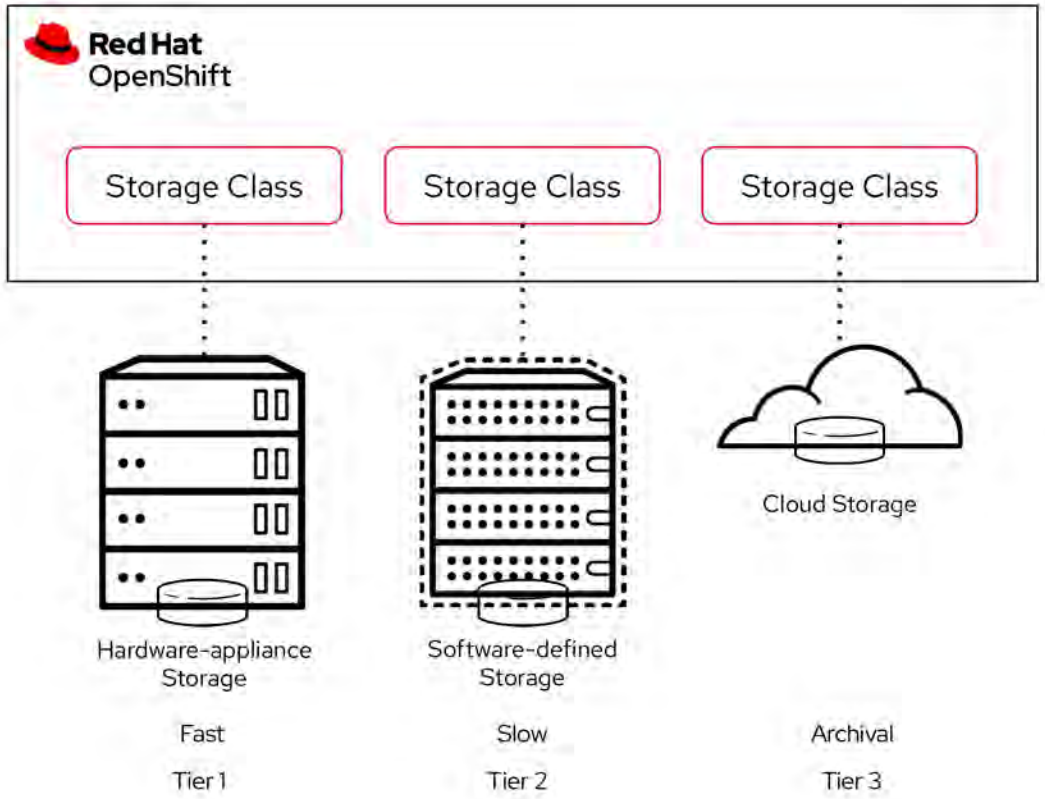

Kubernetes Storage Terminology

- STORAGE CLASS: “class” of storage service
  - Abstraction of back-end storage systems
  - OCP Clusters can have multiple storage classes

- PVC: Persistent Volume Claim,
  - request and attach persistent storage to a pod
  - specifies workload requirements eg resize, snapshot, read-only

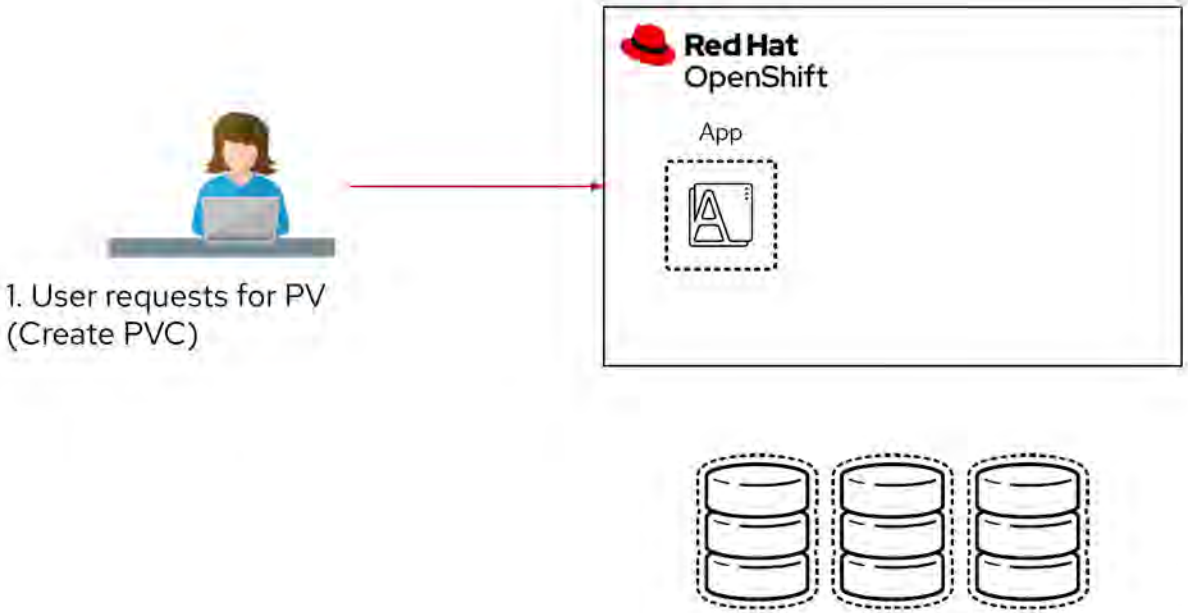

Kubernetes Storage Terminology

- STORAGE CLASS: "class" of storage service
  - Abstraction of back-end storage systems
  - OCP Clusters can have multiple storage classes

- PVC: Persistent Volume Claim,
  - request and attach persistent storage to a pod
  - specifies workload requirements eg resize, snapshot, read-only

- PV: Persistent Volume
  - reference to actual storage on SAN / NAS or local disk
  - can be static or dynamic
  - Usually connected to containers as a filesystem
  - Usually connected to VMs as a virtual disk

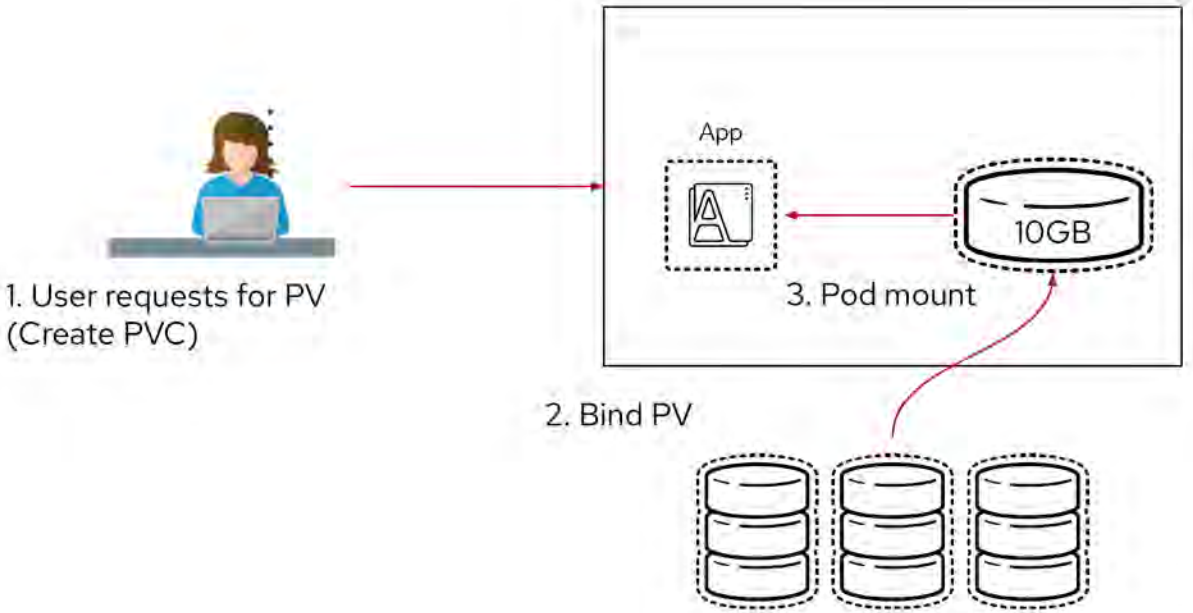

Kubernetes Storage Terminology

- STORAGE CLASS: “class” of storage service
  - Abstraction of back-end storage systems
  - OCP Clusters can have multiple storage classes

- PVC: Persistent Volume Claim,
  - request and attach persistent storage to a pod
  - specifies workload requirements eg resize, snapshot, read-only

- PV: Persistent Volume
  - reference to actual storage on SAN / NAS or local disk
  - can be static or dynamic
  - Usually connected to containers as a filesystem
  - Usually connected to VMs as a virtual disk

- CSI: Container Storage Interface
  - Vendor agnostic layer consume storage from storage backends
  - Protocol agnostic
    - FC /iSCSI/NVMe-oF/NFS/CIFS/CephFS/CephRBD etc

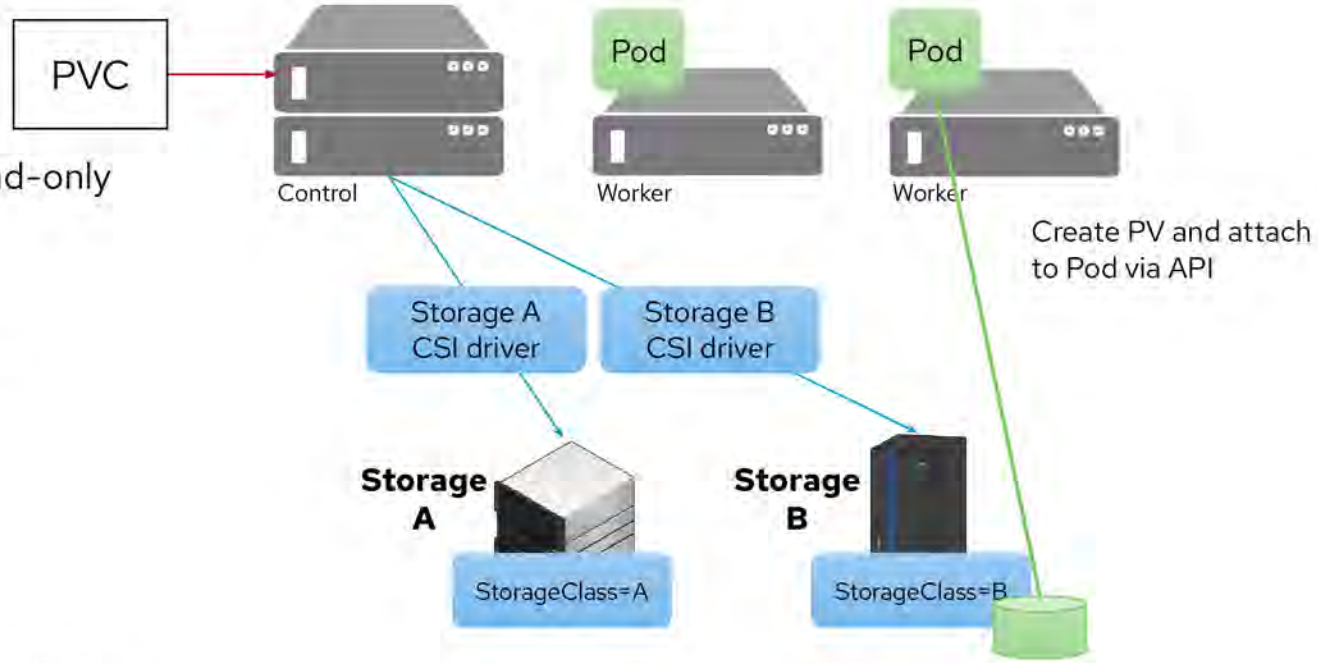

Kubernetes Storage Modes

- Modes of access control mode
- PV is only used in one mode; never used in more than one mode
- Available access modes vary depending on back-end storage systems
  - If an unavailable mode is specified in PVC, the PVC will fail and any PVs will not be bound.

Read Write Once (RWO)

- Read/Write on single node
- Available for almost all storages
- Most frequently used

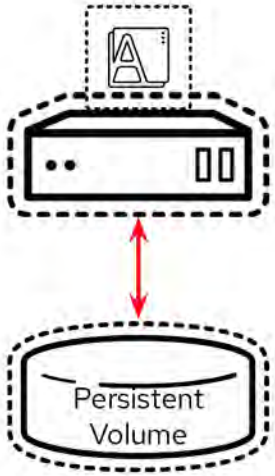

Read Only Many (ROX)

- Read-only on multiple nodes simultaneously
- Not often used and of limited use.

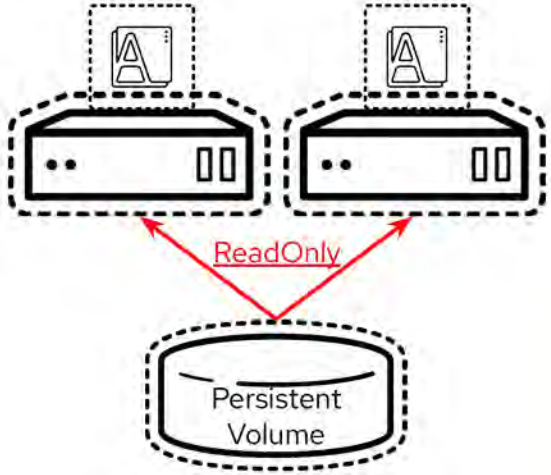

Read Write Many (RWX)

- Read/Write on multiple nodes simultaneously
- Basically available with file storages (e.g. NAS, NFS server)

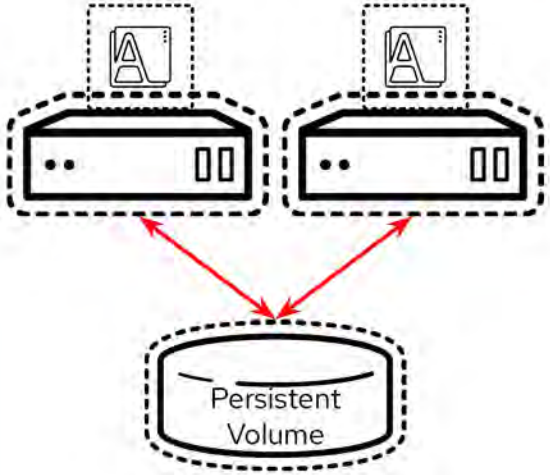

ReadWriteOncePod is now GA as of k8s 1.29 ReadWriteOncePod is now GA as of k8s 1.29
For more details see https://docs.openshift.com/container-platform/4.16/storage/understanding-persistent-storage.html

Storage layer functionality recommended

CSI plugin features

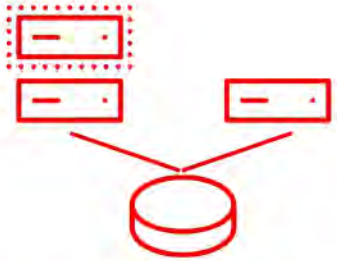

Access attribute: RWX Block

VM live migration &

VM restart

Volume snapshots

CSI:

VM cloning; VM backup

Volume cloning

VM cloning & VM templates

(aka golden images)

Availability

Data redundancy across

failure domains

Backup

Backup Integration

What backup system is compatible?

OpenShift provides OADP (OpenShift API for Data Protection) operator:

- Can be used as a very simple standalone solution
- Used by most Backup Vendors
- Supported by Trilio, Velero, StorWare, IBM Fusion, Kasten by Veeam, Rubrik, Veritas, Cohesty, Dell PowerProtect Data Manager, Commvault etc.
- StorWare can be used as proxy for IBM Spectrum Protect, et al

Note: This is not an exhaustive list of ISV partners, with new partners being added all the time

Snapshots are not backups!

- Snapshots are not stored on a separate device
  - If the storage fails, e.g. RAID group failure, snapshots are also lost. There is no protection against the storage device failing
- Snapshots are a point-in-time recovery tool
  - Useful for data corruption protection, e.g. ransomware, if a snapshot exists before corruption happened
  - Useful for recovery/roll back due to error
- Snapshot implementations depend on, and therefore vary by, the storage provider
  - Performance and capacity impact can be dramatically different
  - Snapshot creation and recovery is dependent on the CSI driver
    - Reverting to a snapshot is done by cloning the snapshot to a new disk and attaching it to the VM

Changed Block Tracking (CBT)

- Changed block tracking relies on the underlying storage to identify which blocks have been modified since the last backup operation
  - Avoids the backup suite having to scan the disk to detect changed blocks

Backup flow with CBT

1. Copy all blocks
2. Track changes in the map
3. Using the change map, copy only changes
4. Reset change map

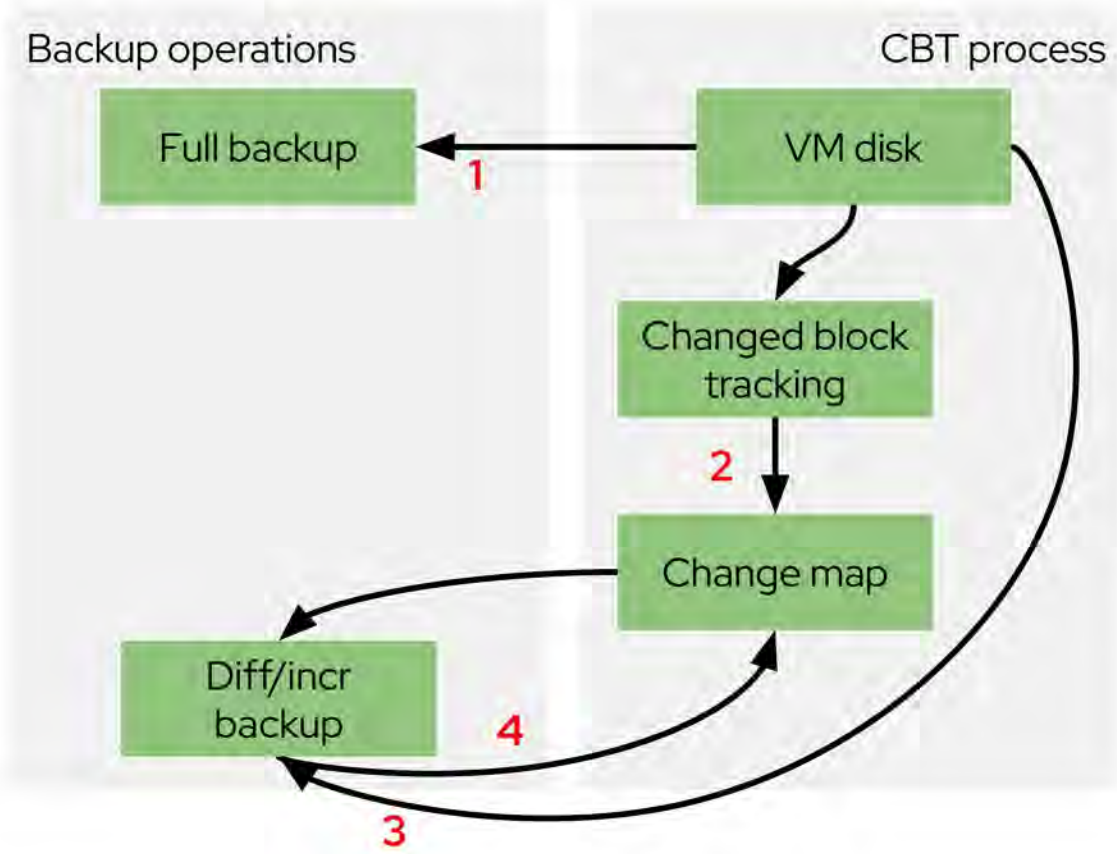

Protecting VMs with multiple disks - VolumeGroupSnapshot

- Kubernetes snapshots are non-deterministic, meaning that the time that the snapshot is created is not exact
- For VMs with multiple disks requiring all volumes to be quiesced, snapshotted, and backed up at the same time, this can be a challenge
  - Disks being snapshotted at different times may lead to application inconsistency
- Upstream work is happening here, KEP-3476

High Availability and Disaster Recovery

Typical Cluster Layout

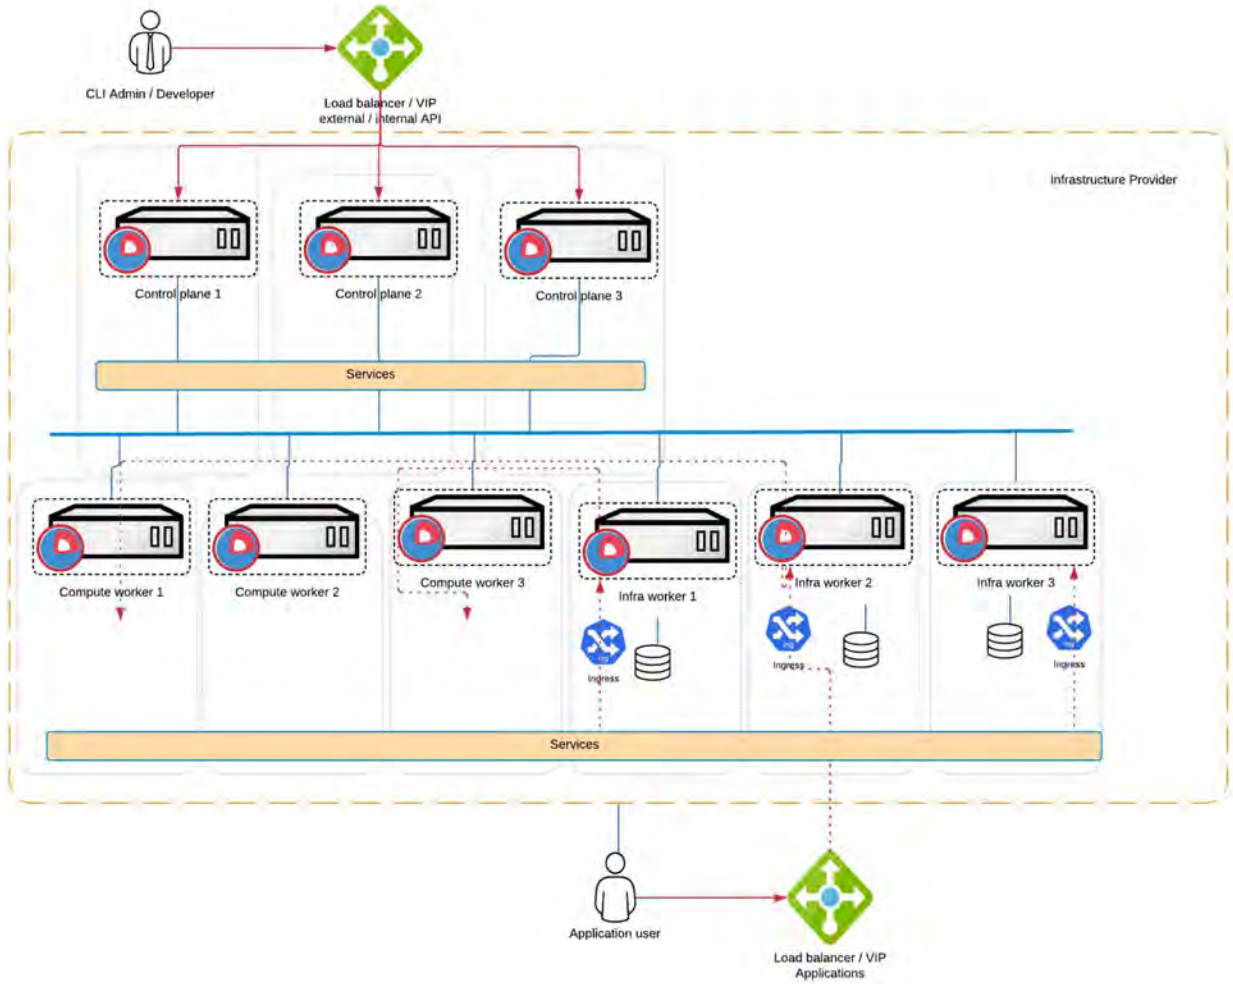

- **Control Plane (ex Master)** - No workloads, critical for API and etcd
- **Infra** - Hosts Ingress, monitoring, logging, registry etc.
- **Worker** - user workloads

Single Control Plane Loss

A single Control Plane is disabled

Impact:

- Control plane fully functional but no longer HA
- Infrastructure Pods: All pods are still up and fully running and still HA
- Customer Applications: All pods are still up and fully running and still HA

Solution:

- Recover ETCD member
- Create a new control plane node with the same settings as the previous node.

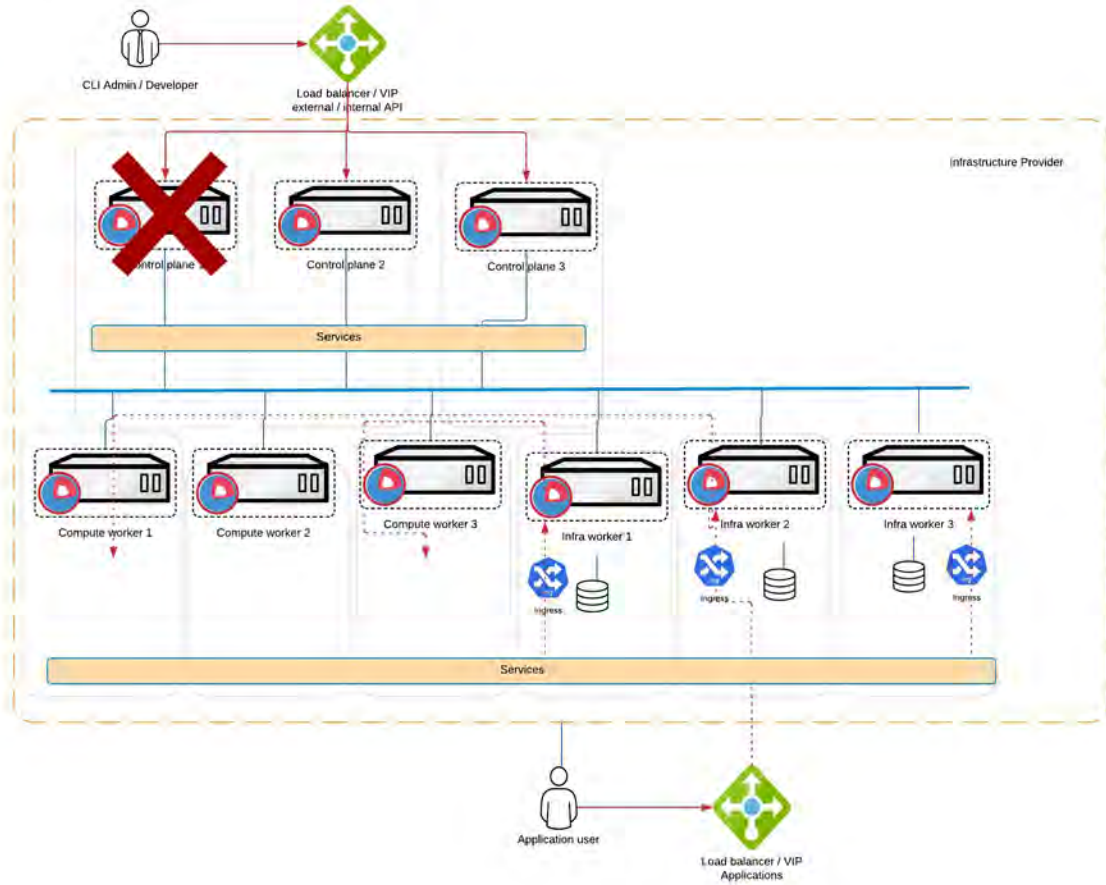

Red Hat | INDITEX

Double Control Plane Loss

Two Control Planes are disabled

Impact:

- Control plane fails and stays in **read-only mode**
- Infrastructure Pods: All pods are still up and fully running and still HA
- Customer Applications: All pods are still up and fully running and still HA
- **No changes possible to the cluster**
- **No recovery mechanisms will be executed**

Solution:

- Deploy new control planes
- Manually resync etcd

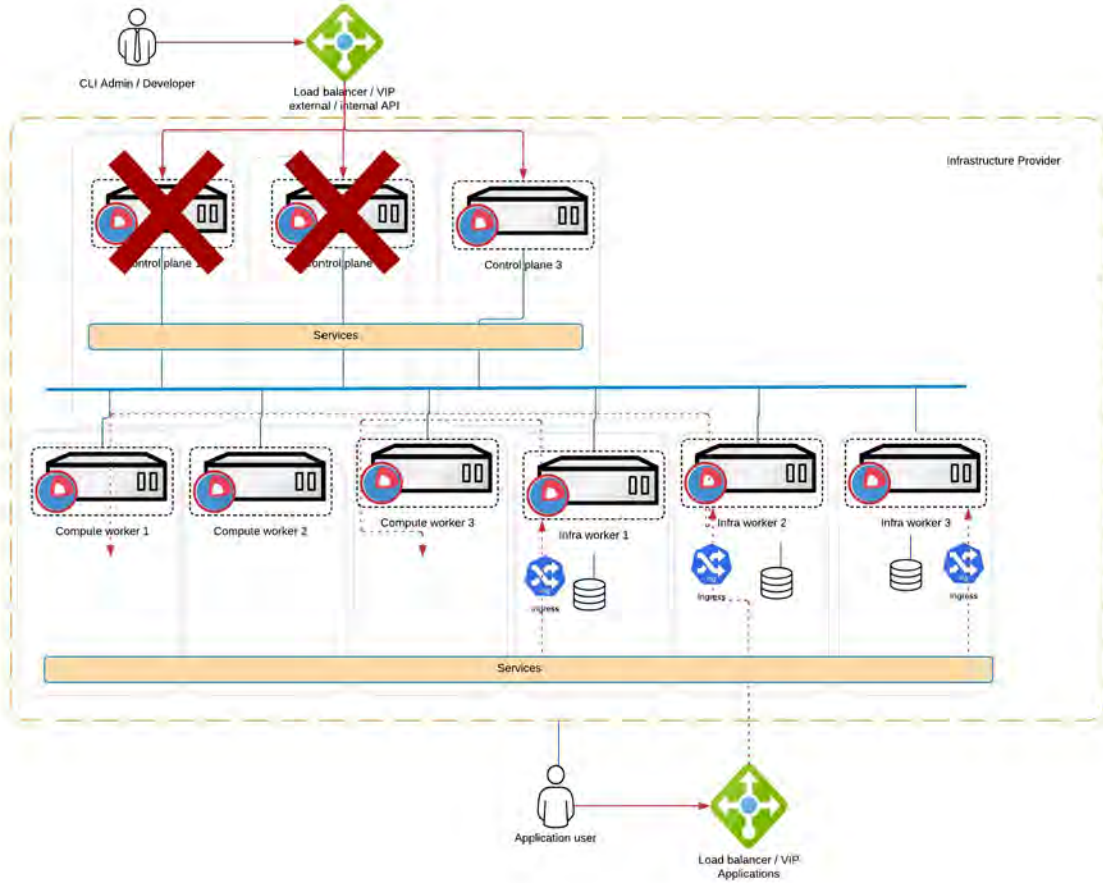

Red Hat | INDITEX

Single Infrastructure Node Loss

An Infrastructure Node is disabled

Impact:

- Control plane: fully functional and still HA
- Infrastructure Pods:
  - Stateless pods (Ingress): fully functional and still HA
  - Stateful pods (Elasticsearch/Registry/monitoring) or storage : fully functional.
  - Elasticsearch is no longer HA
- Customer Applications: All pods are still up and fully running and still HA

Solution:

- Create a new infrastructure node. With the Machine API, simply use the scale up procedure

Make sure the available nodes have the resources so the pods from the unavailable node can be rescheduled. Otherwise, use only two elasticsearch and registry pods

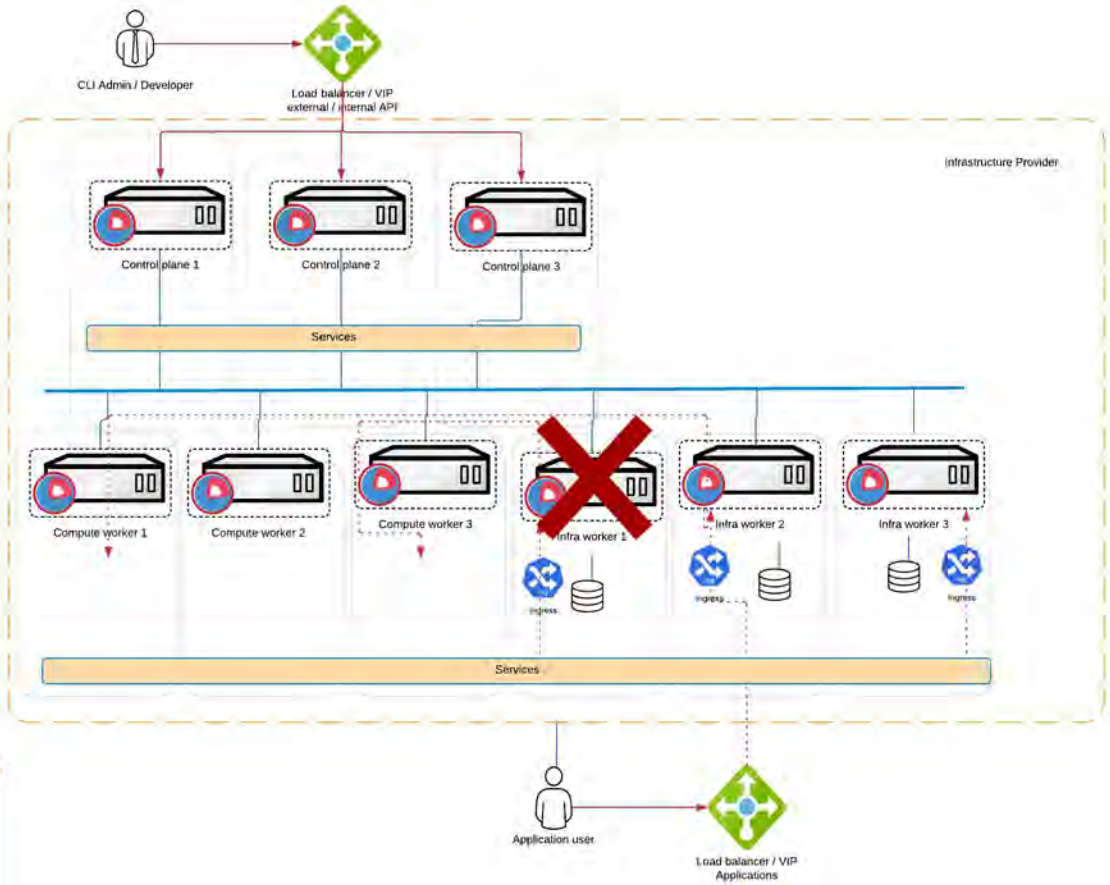

Red Hat | INDITEX

Double Infrastructure Node Loss

2 Infrastructure Node are disabled

Impact:

- Control plane: fully functional and still HA
- Infrastructure Pods:
  - Stateless pods (Ingress): fully functional and still HA
  - Stateful pods (Elasticsearch) : non-functional anymore as quorum is lost!
  - No logs can be written anymore, no s
- Customer Applications: All pods are still up and fully running and still HA

Solution:

- Create a new infrastructure node. With the Machine API, simply use the scale up procedure

Make sure the available nodes have the resources so the pods from the unavailable node can be rescheduled.
Manual intervention might be required to restore services!

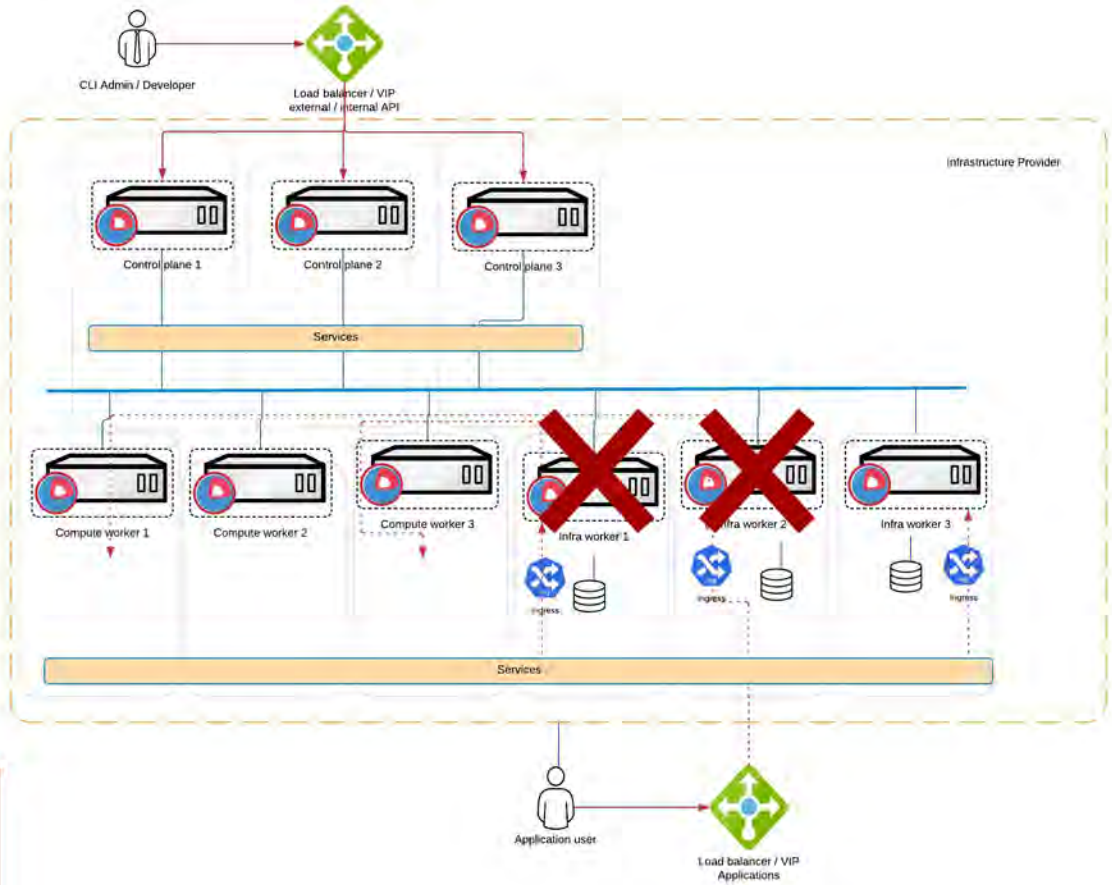

Red Hat | INDITEX

Compute Node Loss

A Compute Node is disabled

Impact:

- Control plane: fully functional and stillHA
- Infrastructure Pods: All pods are still up and fully running and still HA
- Customer Applications: All pods are still up and fully running and still HA depending on the number of compute nodes available.

Solution:

- Create a new node. WIth the Machine API, simply use the scale up procedure.

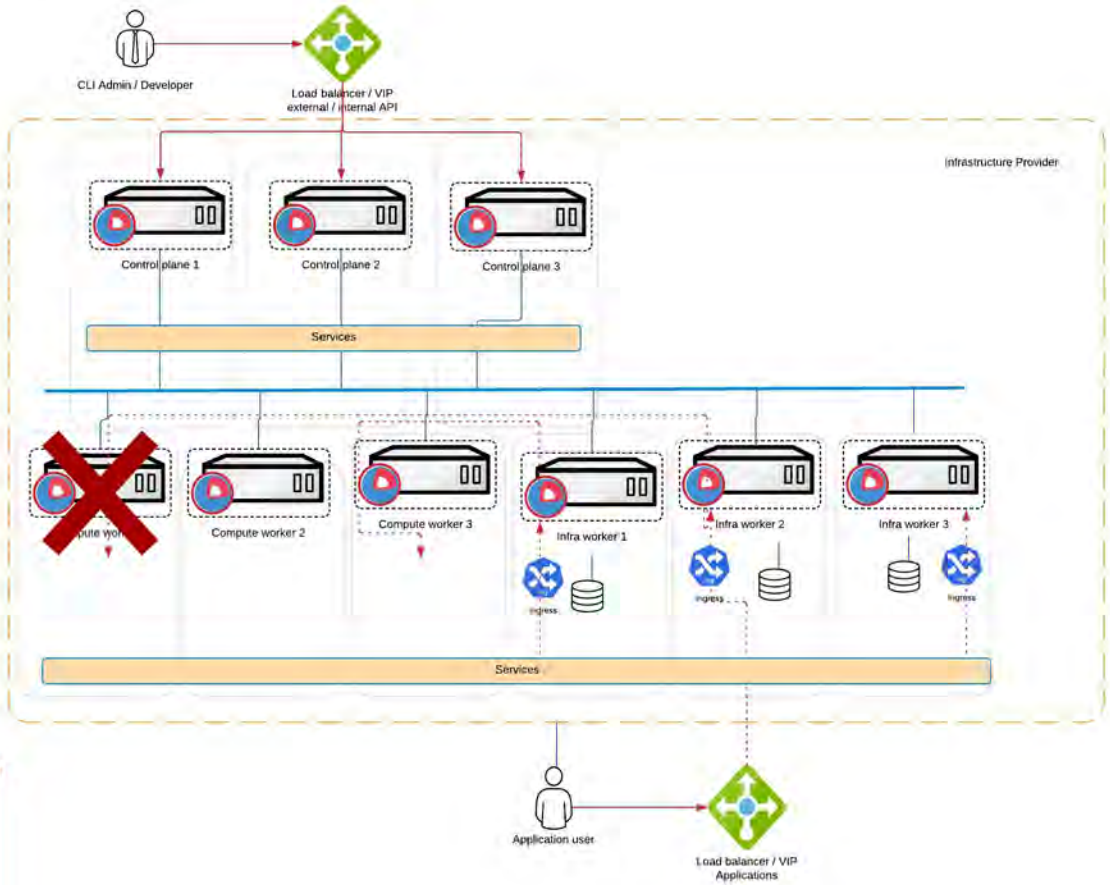

Make sure the available nodes have the resources so the pods from the unavailable node can be rescheduled.

Double Compute Node Loss

Multiple Compute Nodes are lost

Impact:

- Control plane: fully functional and stillHA
- Infrastructure Pods: All pods are still up and fully running and still HA
- Customer Applications: All stateless pods are still up and fully running, are rescheduled and still HA depending on the number of compute nodes available. All stateful applications: depending on where they lived, their quorum might be lost and needs to be restored.
  RWO Storage must be considered

Solution:

- Create a new node. With the Machine API, simply use the scale up procedure.
- Manually recover lost quorum

Make sure the available nodes have the resources so the pods from the unavailable node can be rescheduled.
Make sure your application pods are properly distributed (e.g. different failure domains), so that the possibility of multiple nodes failing might not impact your quorum!

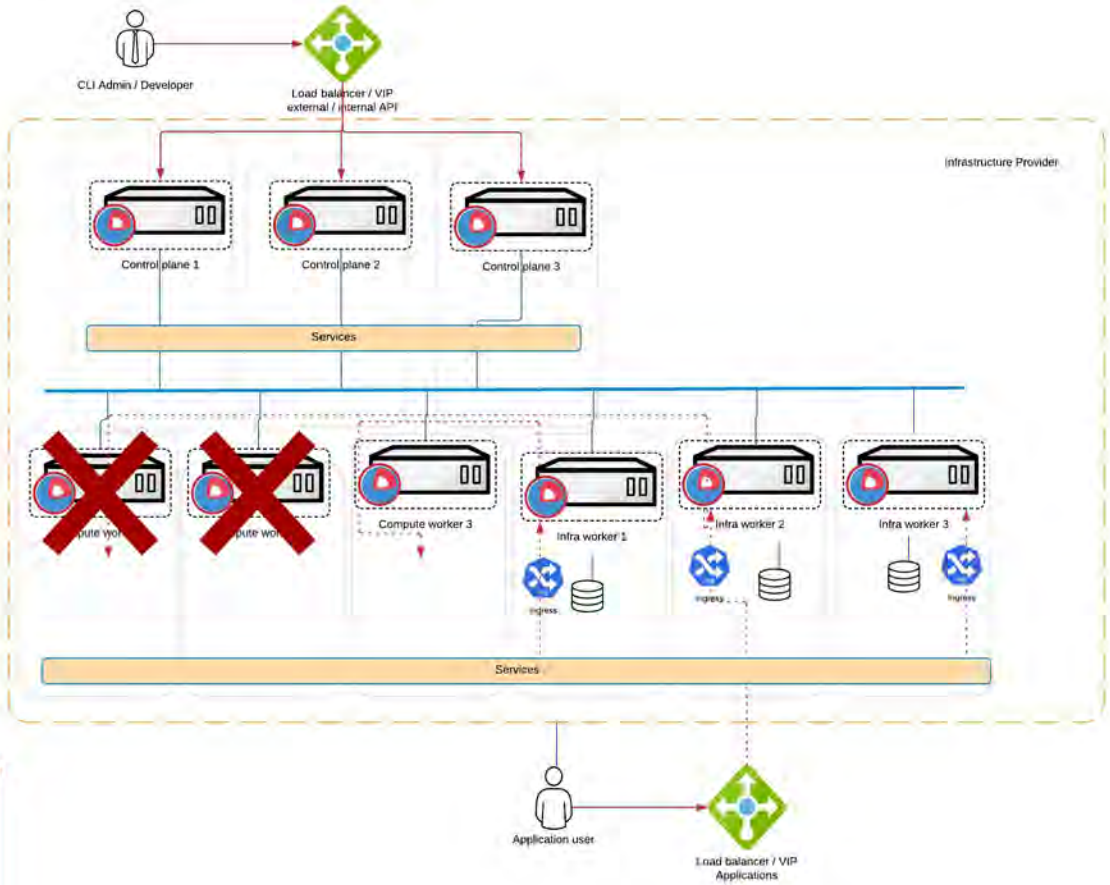

Red Hat | INDITEX

Disaster Recovery

First things first

- **Recovery Point Objective** (RPO) - the maximum acceptable amount of data that can be lost after recovery from disaster - expressed in time
- **Recovery Time Objective** (RTO) - the maximum acceptable time than service can be down after disaster event before being recovered

Application HA

Multi-Zone spanning Cluster for local HA

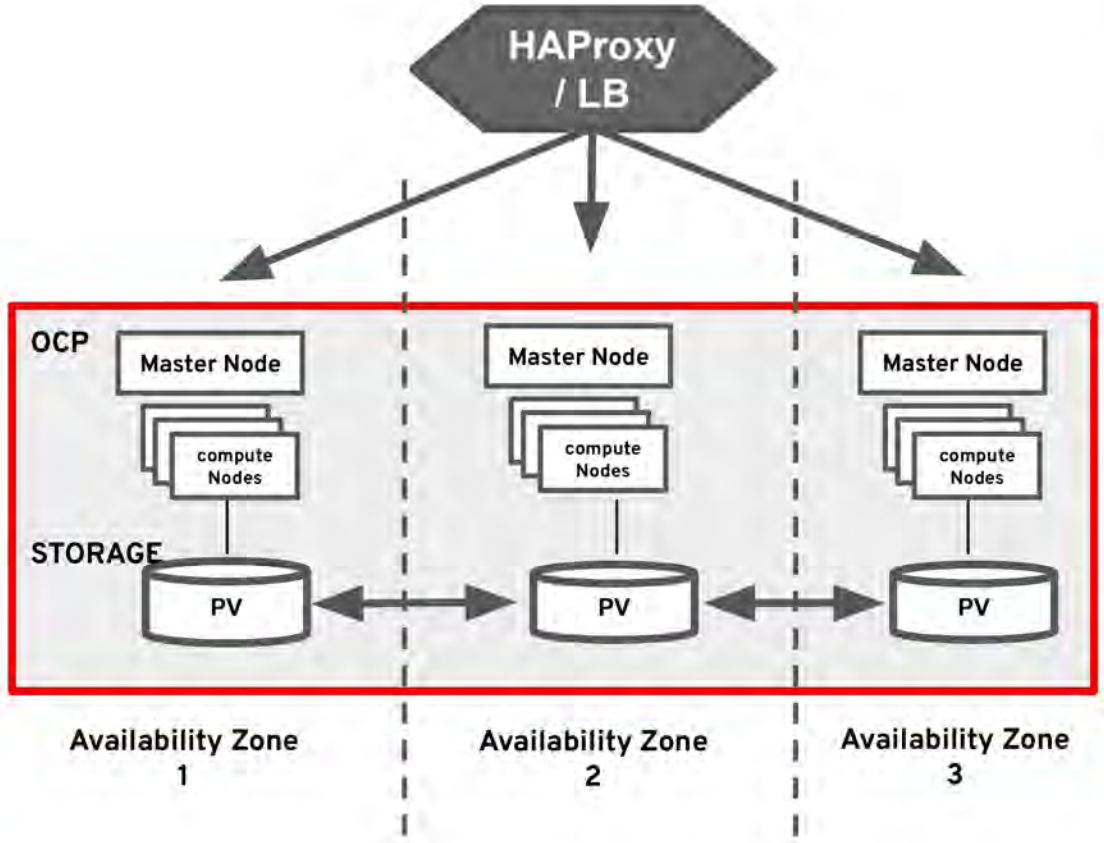

- HA for Stateful Applications deployed on cluster that is stretched across **Availability Zones within a region**
- Installer ensures that resources are deployed across all AZs making the cluster resilient against **failures of any single AZ**
- Storage backend provides synchronous consistent copies in all AZs ensuring **no data loss during zone failure**
- Suitable for public cloud platforms with Regions supporting 3 or more AZs
- Can be **deployed on-prem** when AZs are connected by networks with **<10ms latency**

RTO ~ seconds
RPO = 0

Stretch Cluster-DR

No Data loss Data Mirroring, with single stretched OCP cluster

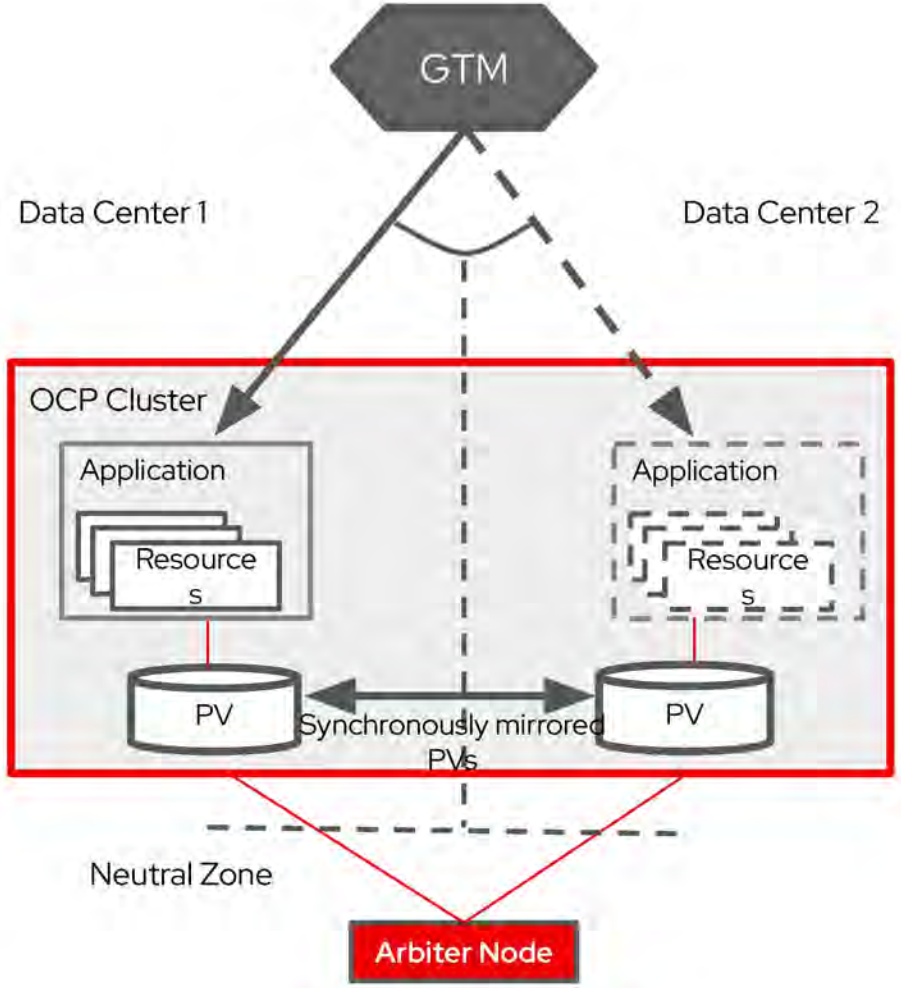

- A **single OCP and hyperconverged STORAGE** are deployed in different AZs or data centers to provide a complete fault isolated configuration
- Storage cluster provides **persistent synchronous mirrored volumes** across two data centers enabling **zero RPO**
- Applications recover automatically on surviving nodes in the other data center ensuring **very low RTO**
- Requires **arbiter node in a third site** for storage cluster and OCP Master node
- **Low latency high throughput link** between DC1 and DC2

RTO ~ seconds
RPO = 0

Red Hat | INDITEX

Metro-DR

No Data loss Data Mirroring, across multiple OCP clusters

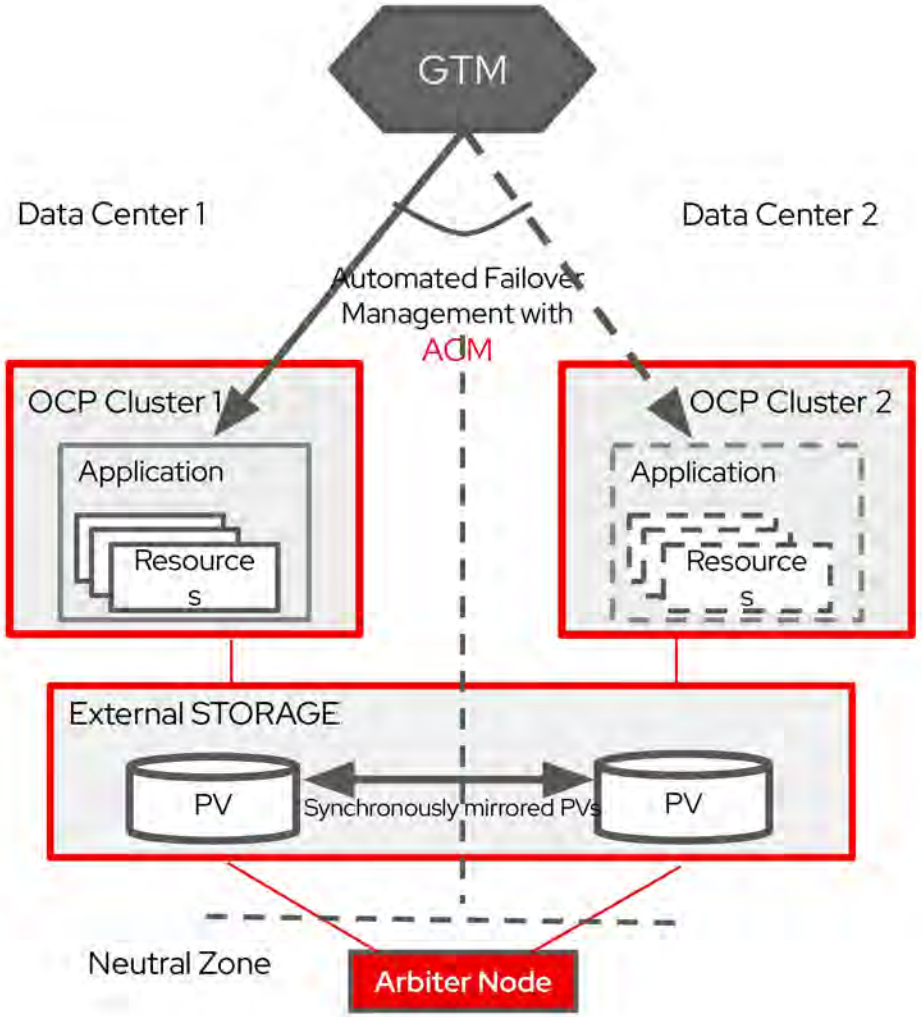

- **Multiple OCP clusters** deployed in different AZs provide a complete fault isolated configuration
- **External** storage appliance provides persistent synchronous mirrored volumes across multiple OCP clusters enabling **zero RPO**
- ACM managed automated Application failover across clusters **reduces RTO**
- Requires **arbiter node in a third site** for storage cluster
- **Low latency high throughput link** between DC1 and DC2
- Arbiter node can be deployed over higher latency networks provided by public clouds

RTO ~ minutes
RPO = 0

Red Hat | INDITEX

Regional-DR with Failover Automation

Protection against Geographic Scale Disasters

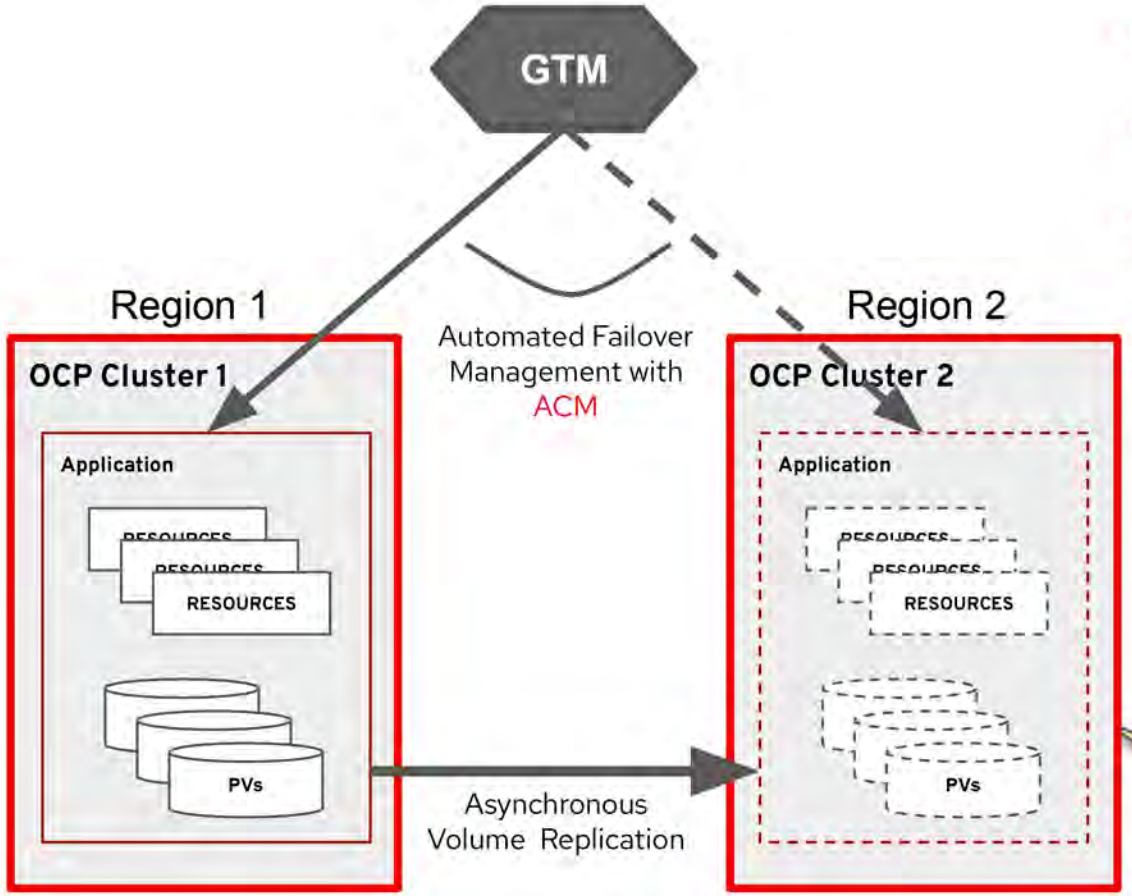

- Asynchronous Volume Replication → **low RPO**
- Automated Failover Management → **low RTO**
- ACM Multi-Cluster manager enables failover and failback automation at application granularity
- **Both clusters remain active** with Apps distributed and protected among them

RTO ~ minutes
RPO ~ minutes

Red Hat | INDITEX

Comprehensive & Flexible Data Protection for the desired SLO

RPO+RTO

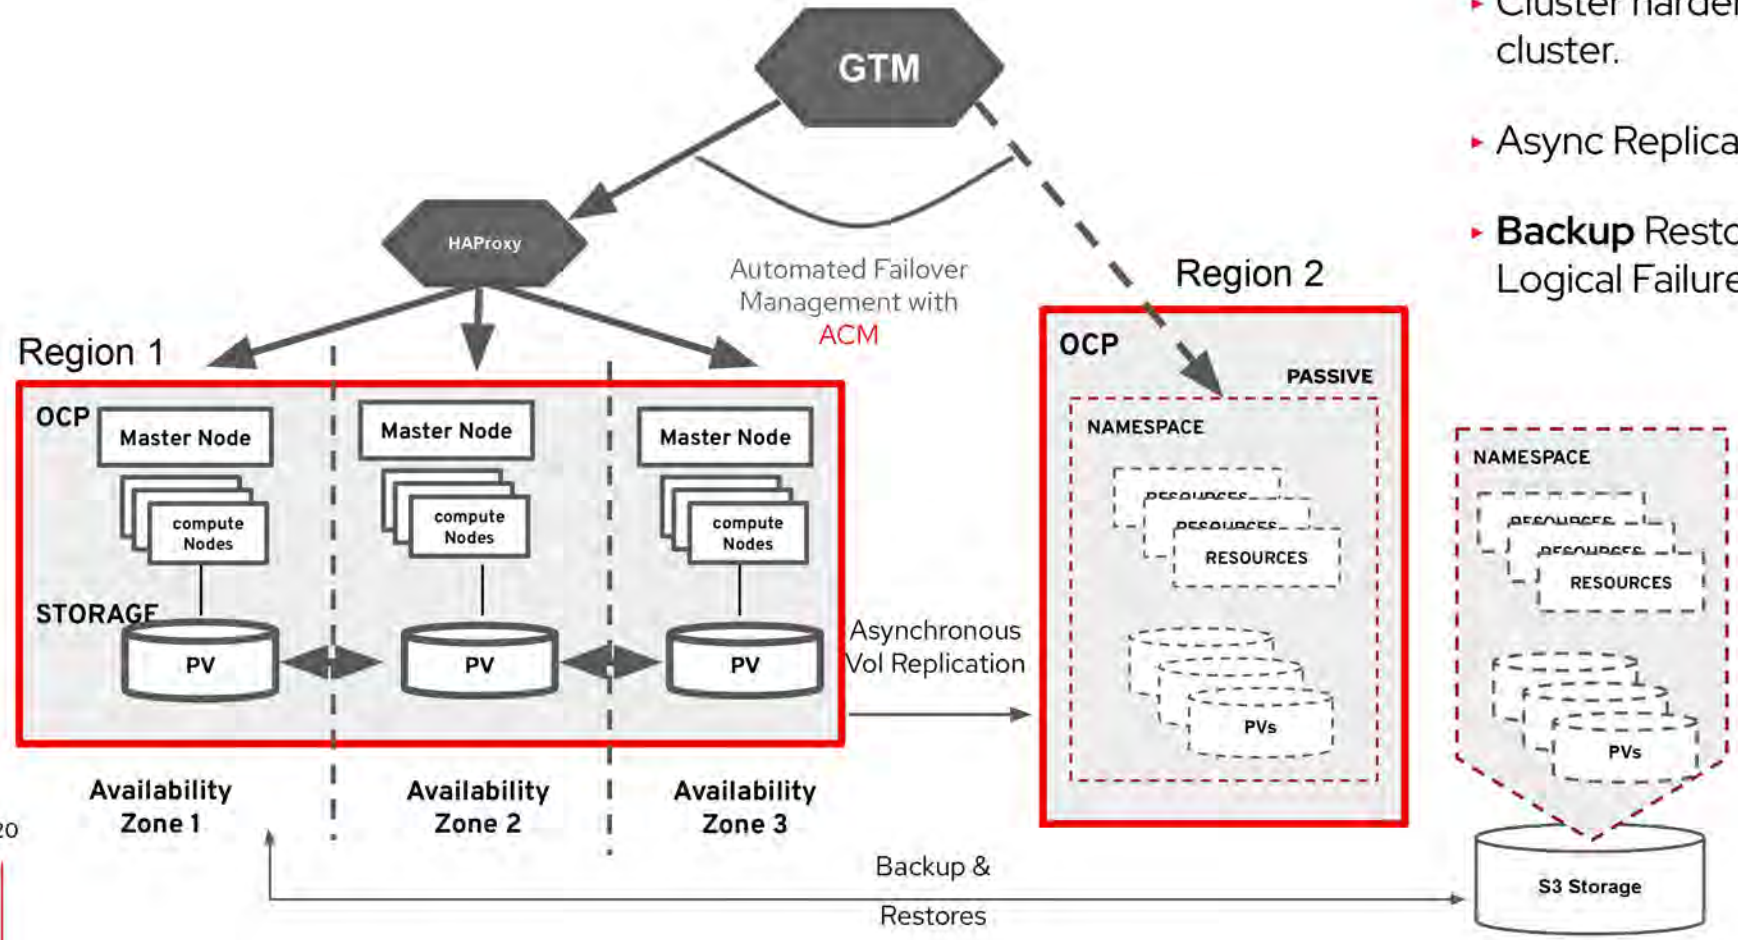

Multi-tier protection against various failure conditions:

- Cluster hardening with Multi-Zone spanning OCP cluster.
- Async Replication for HW and Data Center Failures
- **Backup** Restores from Snapshots for Software & Logical Failures

Red Hat | INDITEX

Introducing Advanced Cluster Management (ACM)

Red Hat Advanced Cluster Management

For Kubernetes

Simplified operation and maintenance

View, manage, operate and solve issues all through a single console.

Runs on OpenShift

Like any other Kubernetes app, easily run and manage it on top of a OpenShift cluster.

Hub-Spoke architecture

Have all configurations managed by the Hub cluster component and seamlessly add Spoke Kubernetes clusters to the central hub.

Tight Integration

RHACM comes with a rich API, add-ons and it can integrate with some key other enterprise tools.

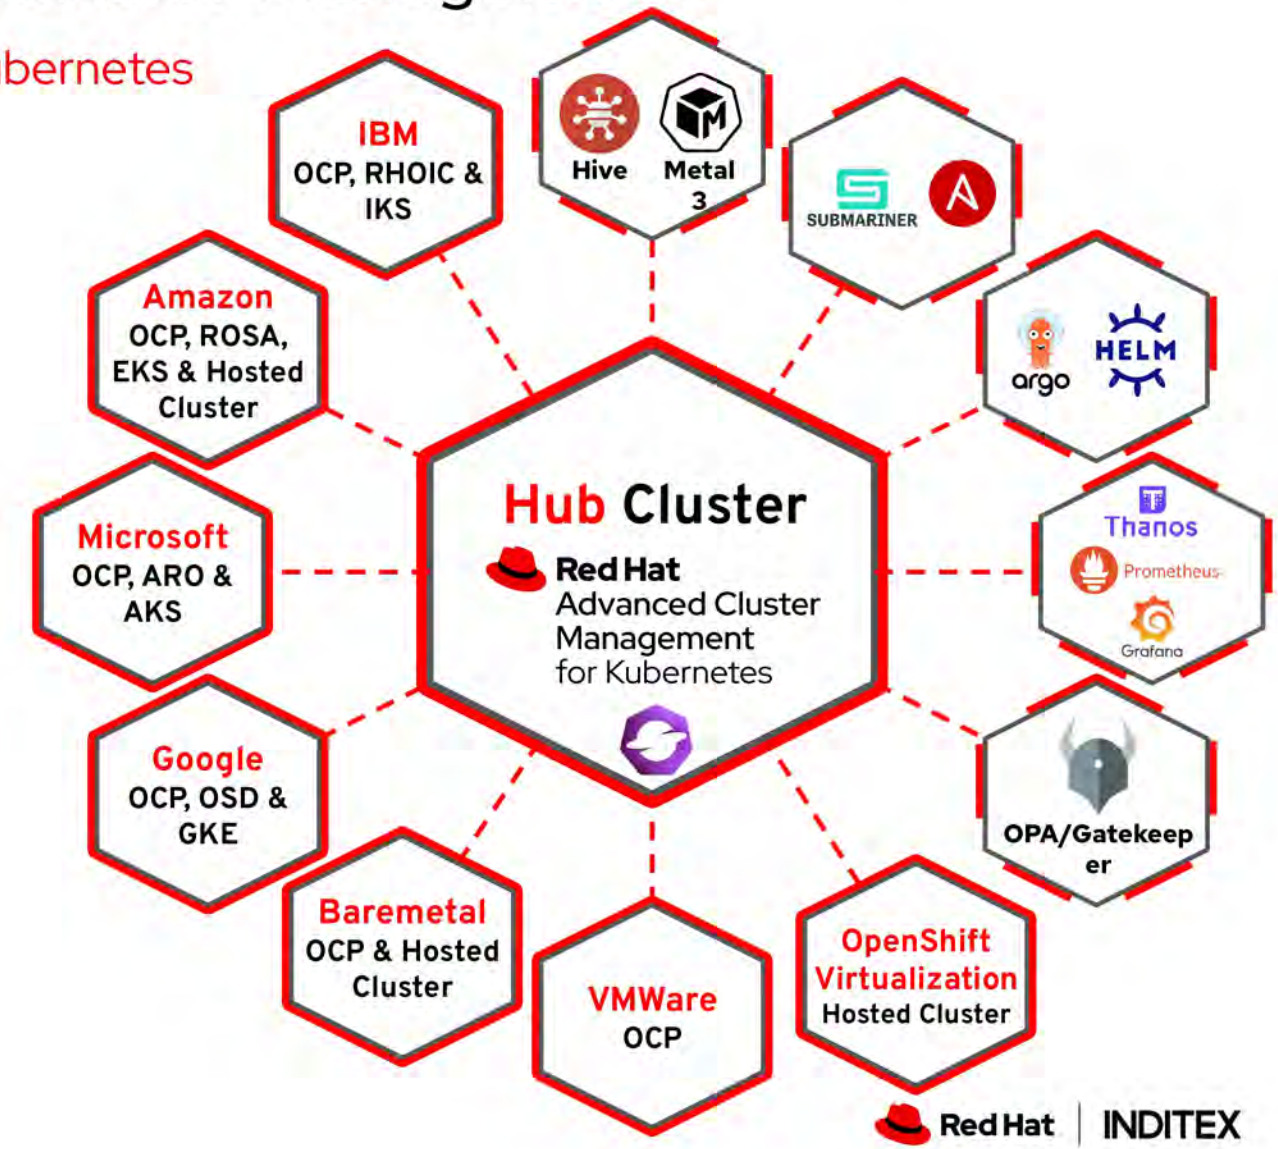

Rules of the game

OpenShift Virtualization encourages you to change

- it brings with it a new view on operational models, including adaptation to kubernetes principles and encouraging for e.g. automation and DevOps.

Sessions

01

Sesión 1 | Presencial
10 de diciembre, 2025 | 10:00 - 14:00
OpenShift Virtualization: Consolidando Entornos

02

Sesión 2 | Virtual
11 de febrero, 2026 | 10:00 - 14:00 **Poniendo en práctica OpenShift Virtualization**

03

Sesión 3 | Presencial
24 de marzo, 2026 | 10:00 - 14:00
Ansible Automation Platform 2.6, una nueva experiencia

04

Sesión 4 | Virtual
TBC | 10:00 - 14:00
Tejiendo Patrones de Automatización: Laboratorio Práctico con Ansible:

Thank you

Red Hat is the world's leading provider of enterprise open source software solutions. Award-winning support, training, and consulting services make Red Hat a trusted adviser to the Fortune 500.

linkedin.com/company/red-hat

youtube.com/user/RedHatVideos

facebook.com/redhatinc

twitter.com/RedHat

Red Hat | INDITEX

OpenShift Virtualization LAB

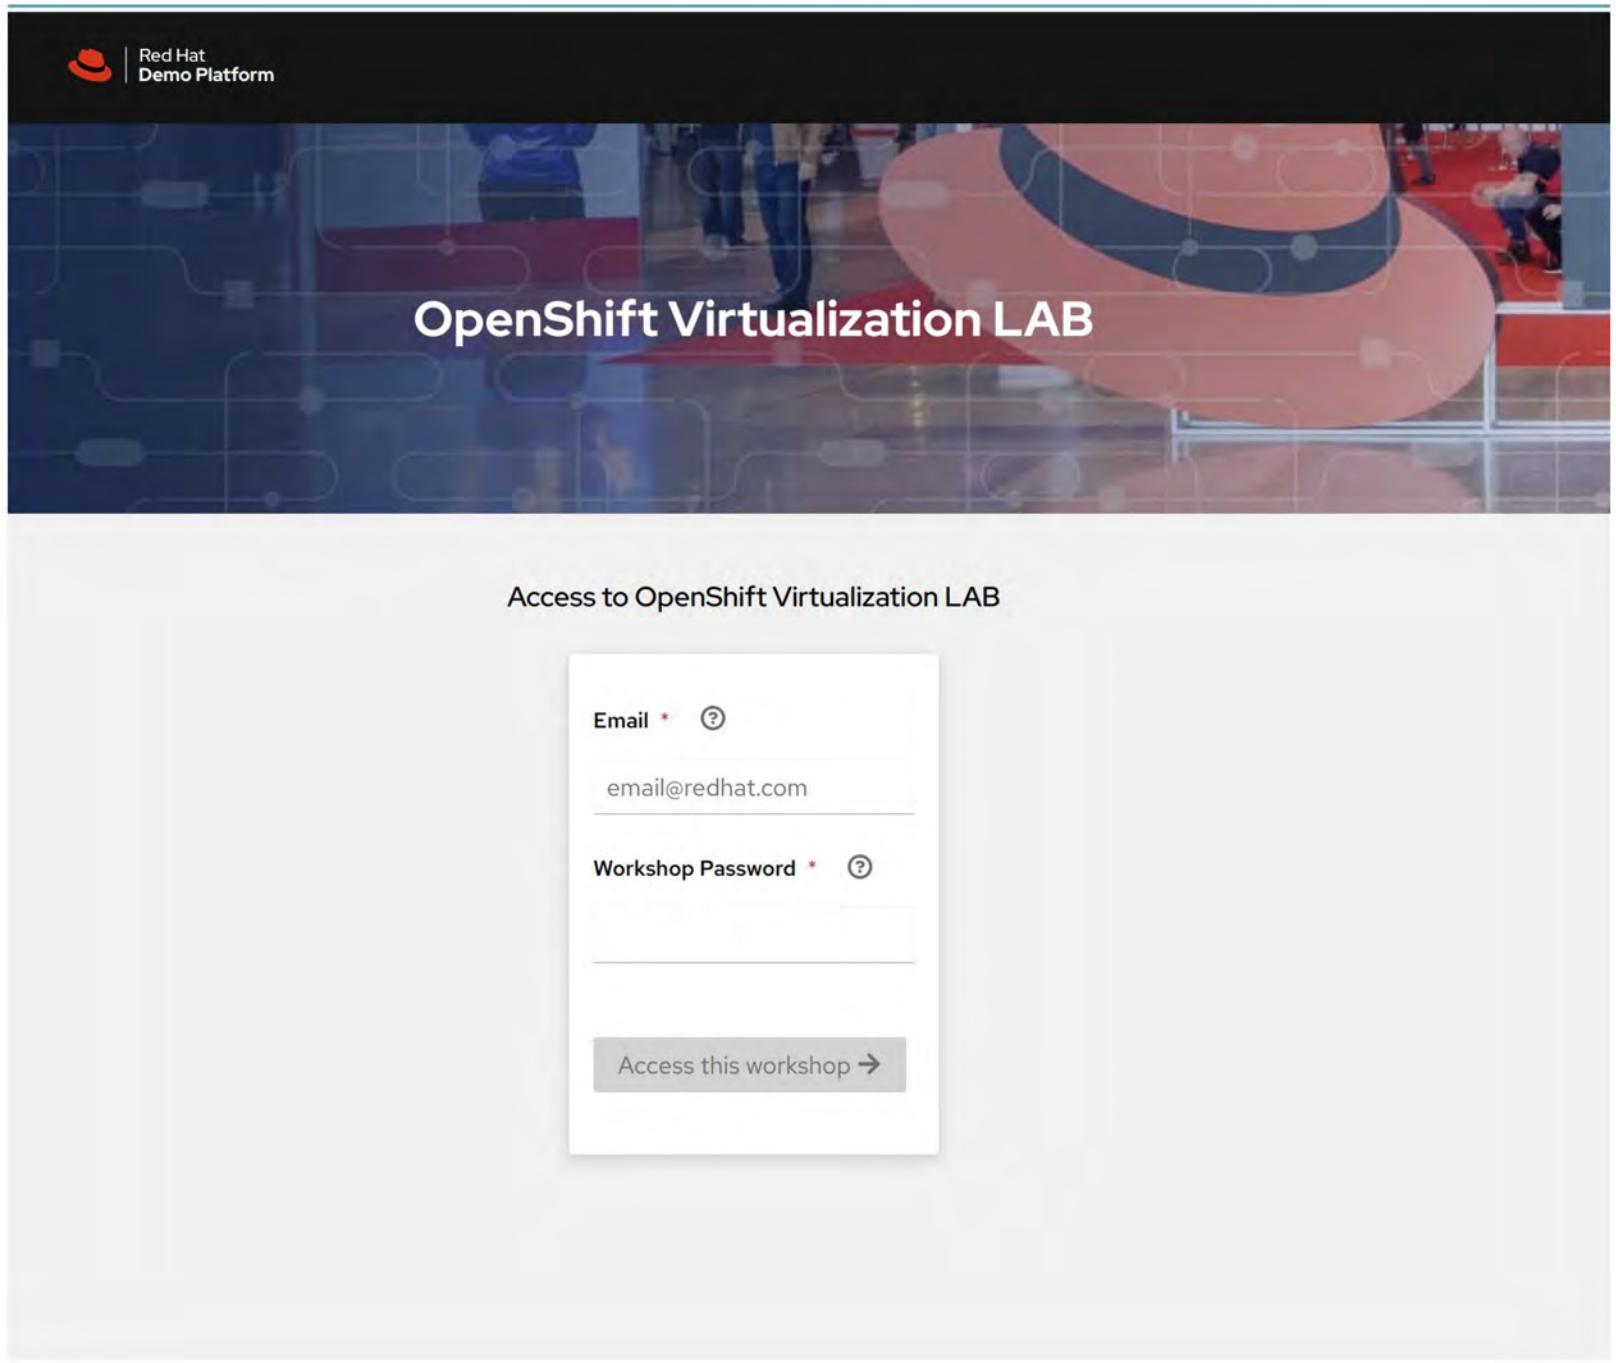
Red Hat
Demo Platform
OpenShift Virtualization LAB
Access to OpenShift Virtualization LAB
Email *
email@redhat.com
Workshop Password *
Access this workshop

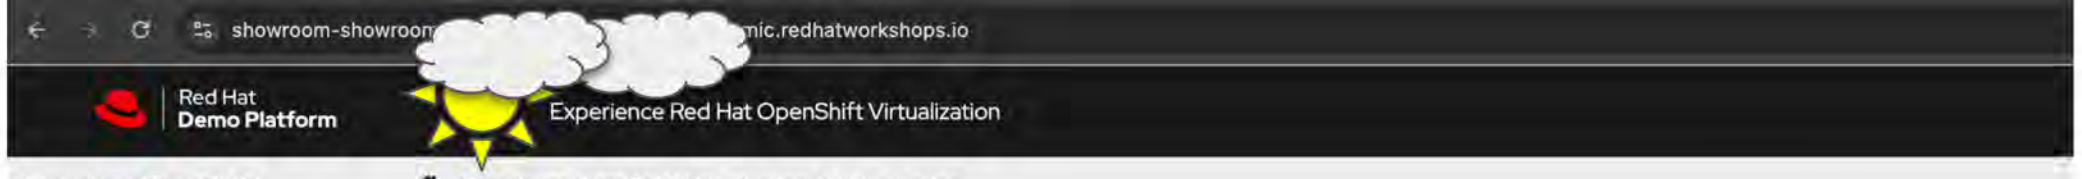

Welcome to Experience OpenShift Virtualization!

Introduction

OpenShift Virtualization enables you to bring virtual machines onto a modern, Kubernetes-based infrastructure. It enables the development and delivery of new applications as well as the modernization of existing ones and can create applications that consist of virtual machines, containers, and serverless functions - all managed together using Kubernetes-native tools and paradigms.

This roadshow event is organized to allow you to have a hands-on experience with Red Hat OpenShift Virtualization.

In this event we will explore many common management activities that virtualization administrators often encounter in their day to day workflows.

Who Will Benefit Most from attending a Roadshow?

Virtual Machine Administrators – Those responsible for day to day management of virtual guests in OpenShift Virtualization. These users will often find themselves responsible for provisioning virtual guests, and day to day management of the guests, and the applications running within.

Virtual Infrastructure Administrators – Those responsible for the physical infrastructure hosting the OpenShift Virtualization solution. These users will be responsible for physical hardware, storage, and networking changes to the environment, that will affect the day to day operations of the running virtual machines.

Credentials for the OpenShift Console

Your OpenShift cluster console is available here.

Your login is available with:

- **User:** use
- **Password:**

vCenter Access

In the migration chapter of the lab, you will be asked to login and examine a VMware vSphere environment.

For access, please use the following credentials:

- **vcenter_user:** ocpvi
- **vcenter_password:** e9RiS

Version Information

This edition of the OpenShift Virtualization Roadshow has been developed using the following software versions:

- Red Hat OpenShift 4.18.3
- Red Hat OpenShift Virtualization 4.18.0
- Red Hat OpenShift Data Foundation 4.18.1
- Red Hat OADP 1.4.4
- Red Hat Migration Toolkit for Virtual Machines 2.7.2

(Accurate as of March 24, 2025)

Next
Virtual Machine Management >

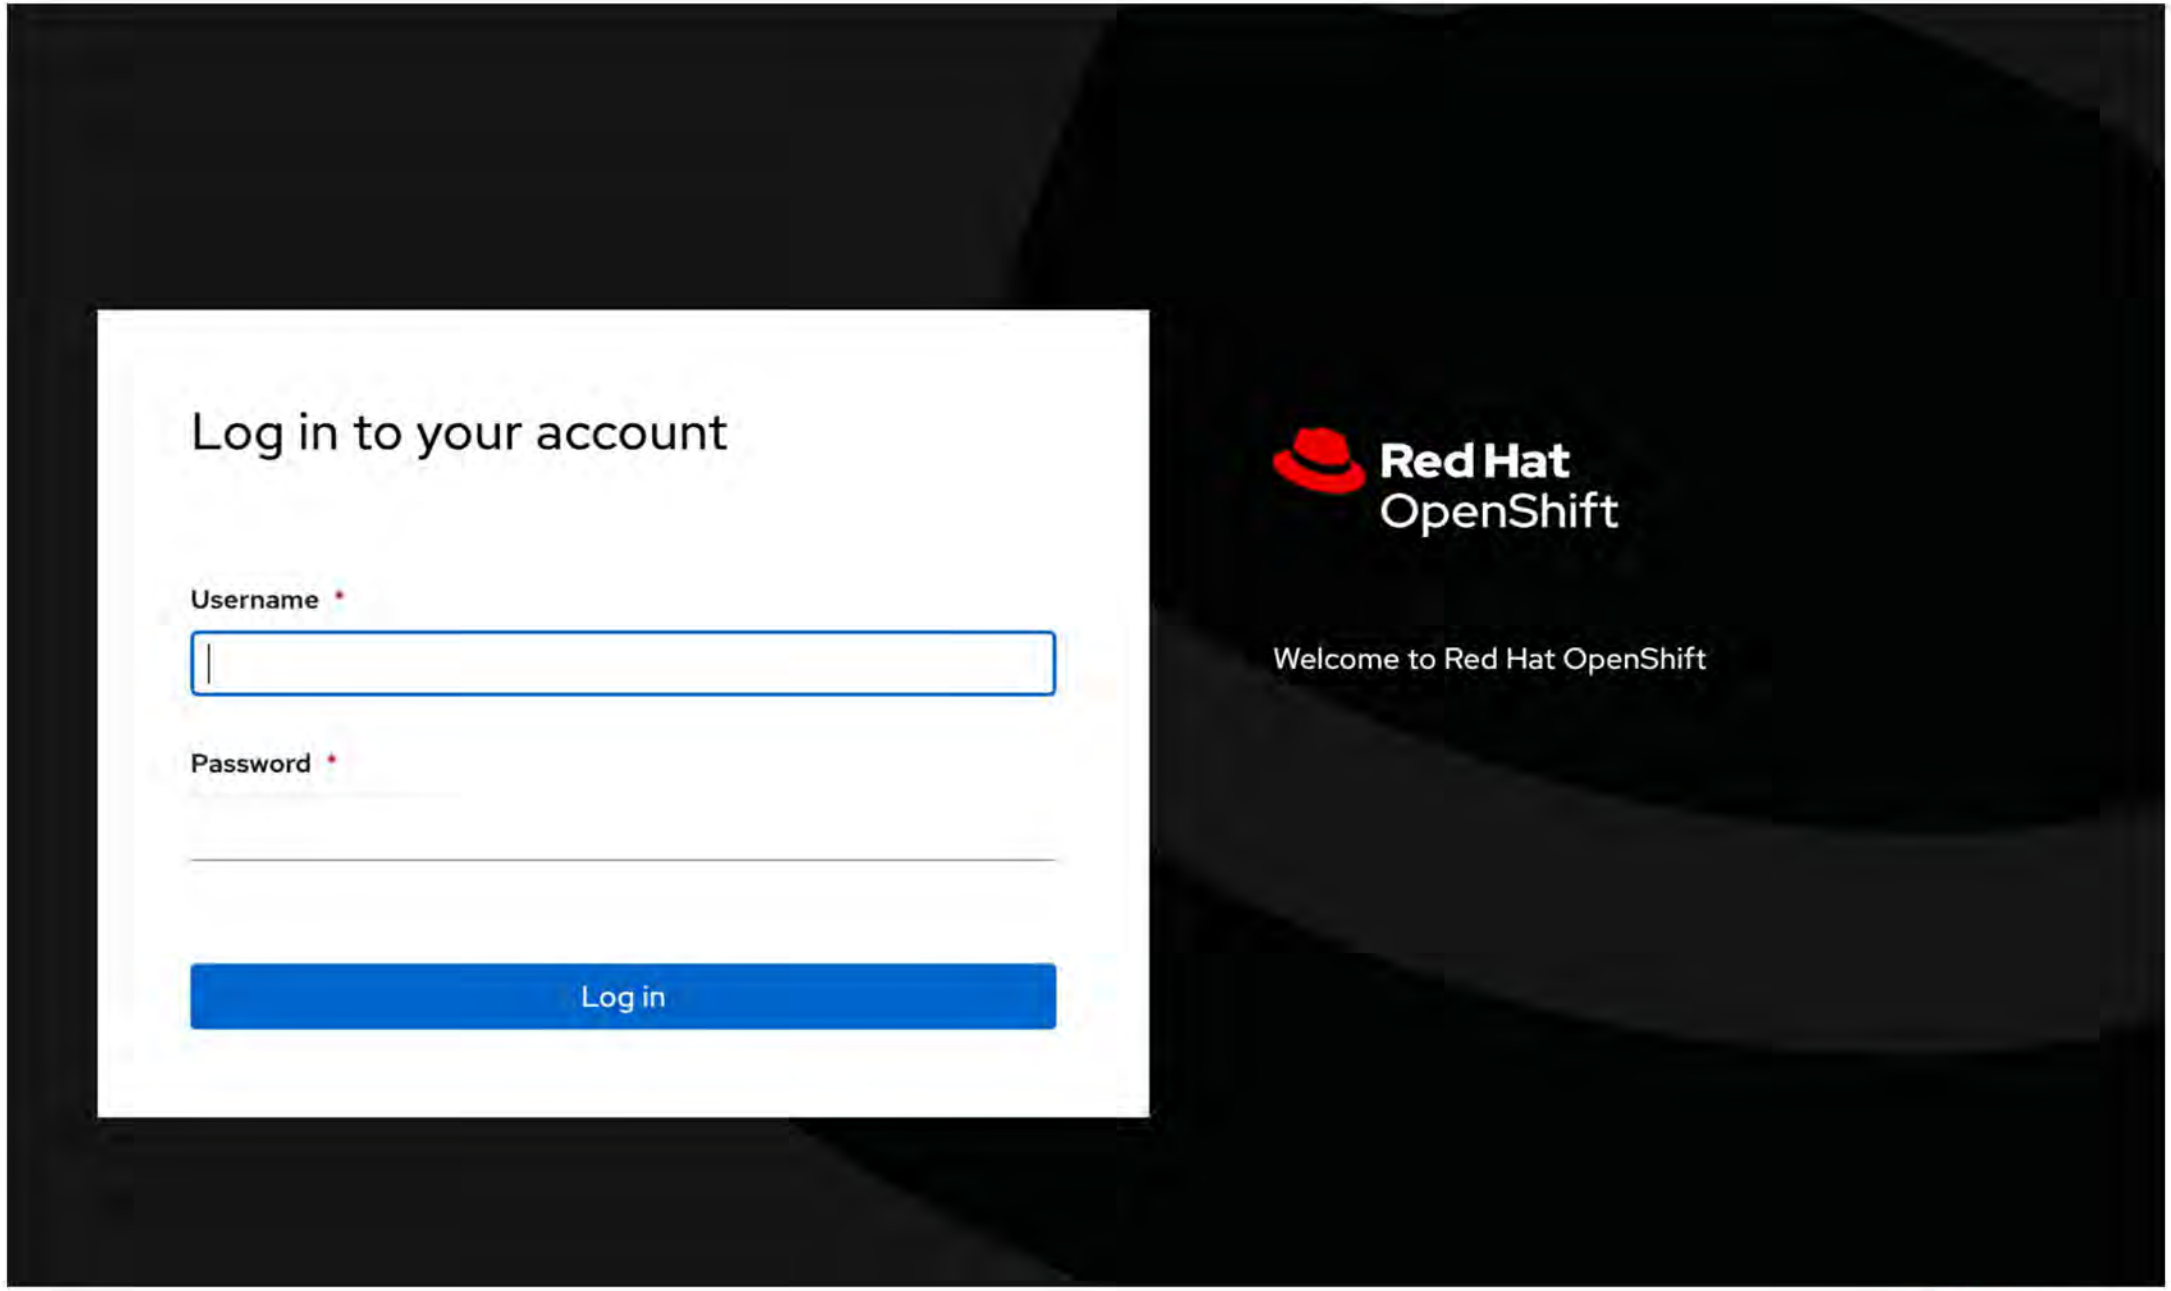
Log in to your account
Username *
Password *
Log in
Red Hat
OpenShift
Welcome to Red Hat OpenShift

IMPORTANT

Having many participants performing the same task in parallel in a simulated lab environment can cause this task to perform much slower than in a real environment. For this lab instance we have limited the number of in-flight VMs to 1 at a time. Please be patient with this process as it completes. You may continue with other sections in the roadshow as the migrations complete.

References & Links

Resources to explore further

- Demo environment for partners
  - ✓ Name of the lab we completed: Experience OpenShift Virtualization Roadshow
- OpenShift Virtualization catalog
  - ✓ List of certified partners (storage, networking, backup, ...)
- OpenShift Virtualization: Reference implementation guide
- Training and certification for Red Hat OpenShift Virtualization
- Configuring USB host passthrough
- Certified Guest Operating Systems in OpenShift Virtualization

Thank you

Red Hat is the world's leading provider of enterprise open source software solutions. Award-winning support, training, and consulting services make Red Hat a trusted adviser to the Fortune 500.

linkedin.com/company/red-hat

youtube.com/user/RedHatVideos

facebook.com/redhatinc

twitter.com/RedHat

Red Hat | INDITEX