Red Hat Summit

Connect

Exploring Latest Features in OpenShift Virtualization

Gökhan Göksu
Senior Solution Architect
gokhan@redhat.com
www.linkedin.com/in/ggoksu

Red Hat has a long history with Virtualization

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

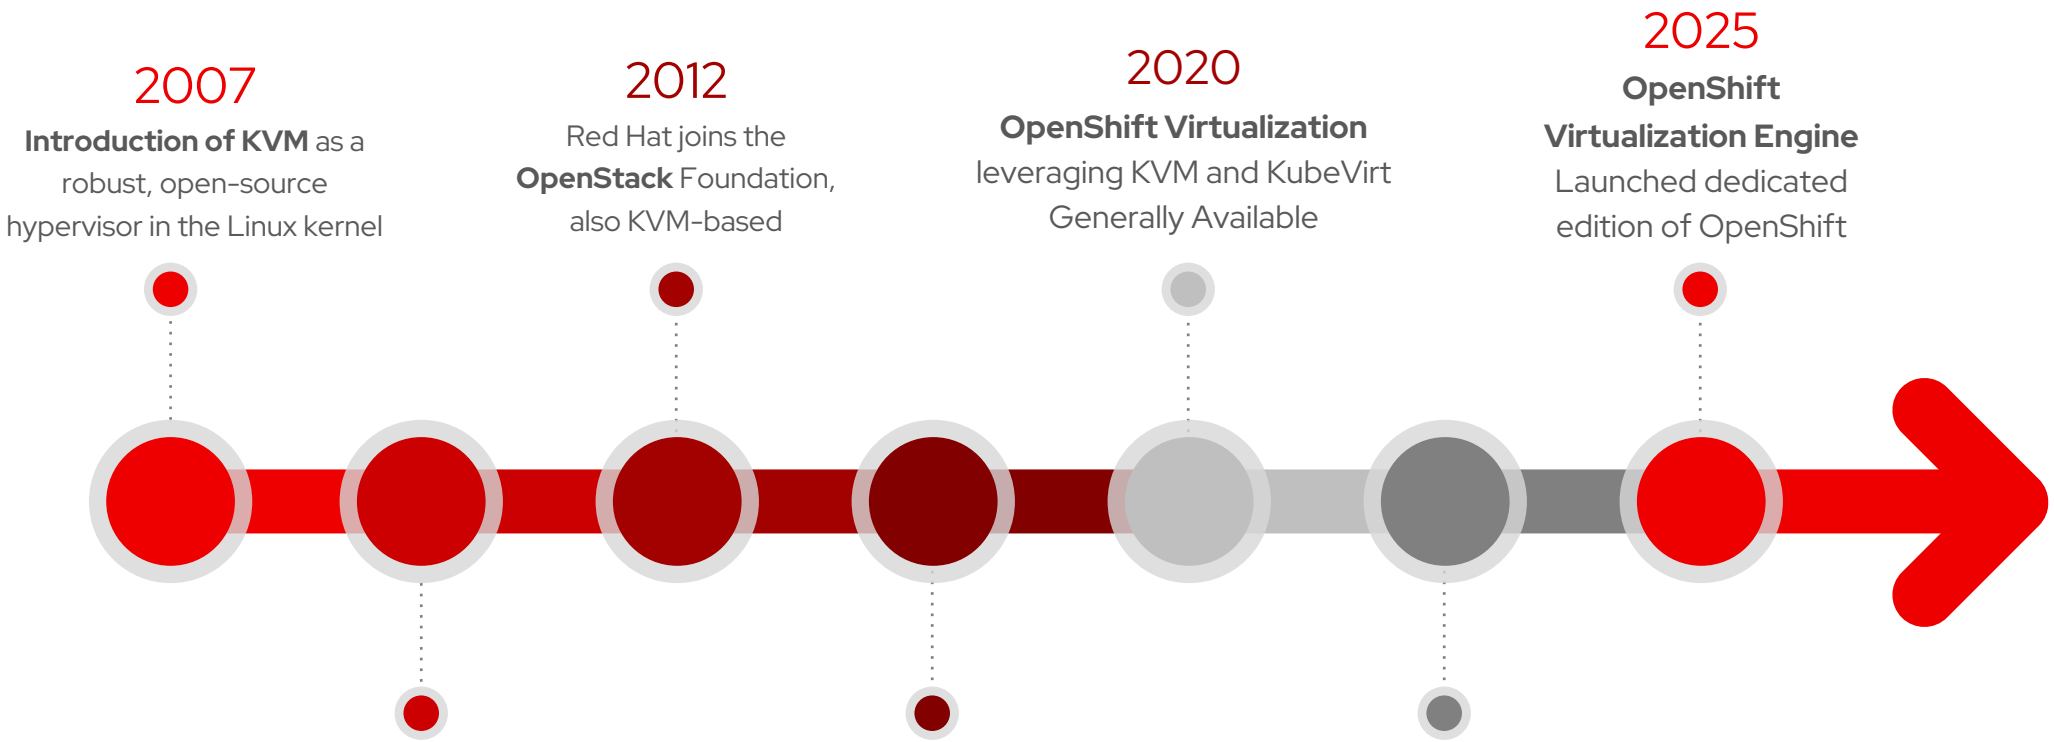

2007
Introduction of KVM as a robust, open-source hypervisor in the Linux kernel

2008
Acquisition of Qumranet, leading to the development of Red Hat Virtualization (RHV) based on KVM

2012
Red Hat joins the **OpenStack** Foundation, also KVM-based

2016
Modernize with KubeVirt VM management on Kubernetes

2020
OpenShift Virtualization leveraging KVM and KubeVirt Generally Available

2023
Migrate: Focus on Core IT Virtualization applications

2025
OpenShift Virtualization Engine Launched dedicated edition of OpenShift

What We Hear From the Field

"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 need 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

OpenShift Virtualization

- Virtual machines
  - Running in **containers**, managed as Pods
  - Using the **KVM** hypervisor
- Scheduled, deployed, and managed by **Kubernetes**
- Integrated with container orchestrator resources and services
  - Traditional Pod-like **network connectivity** and/or external VLAN
  - **Persistent storage** with PVC, PV, StorageClass

Red Hat OpenShift Virtualization

The modern option for general purpose virtualization

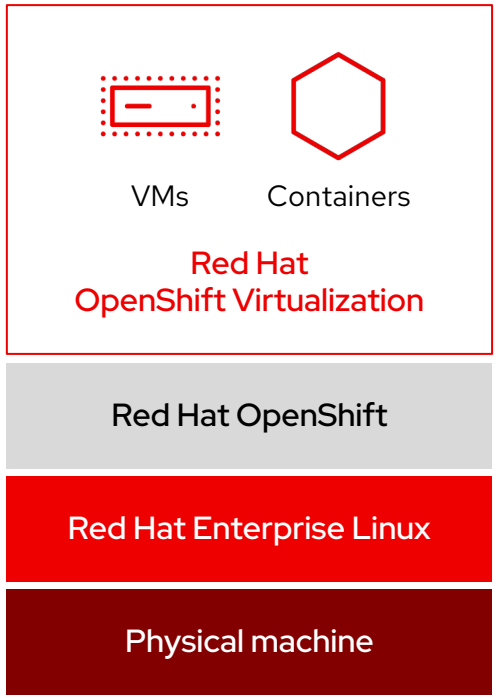

- **Unified platform**
  for virtual machines and containers
- **Consistent management**
  tools, interfaces, and APIs incl. ACM and AAP integrations
- **Performance and stability**
  of Linux, KVM, and qemu
- **Healthy open source community**
  the KubeVirt project is a top 10 CNCF active project, with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

Red Hat

OpenShift Virtualization Features

Dedicated VM Admin View

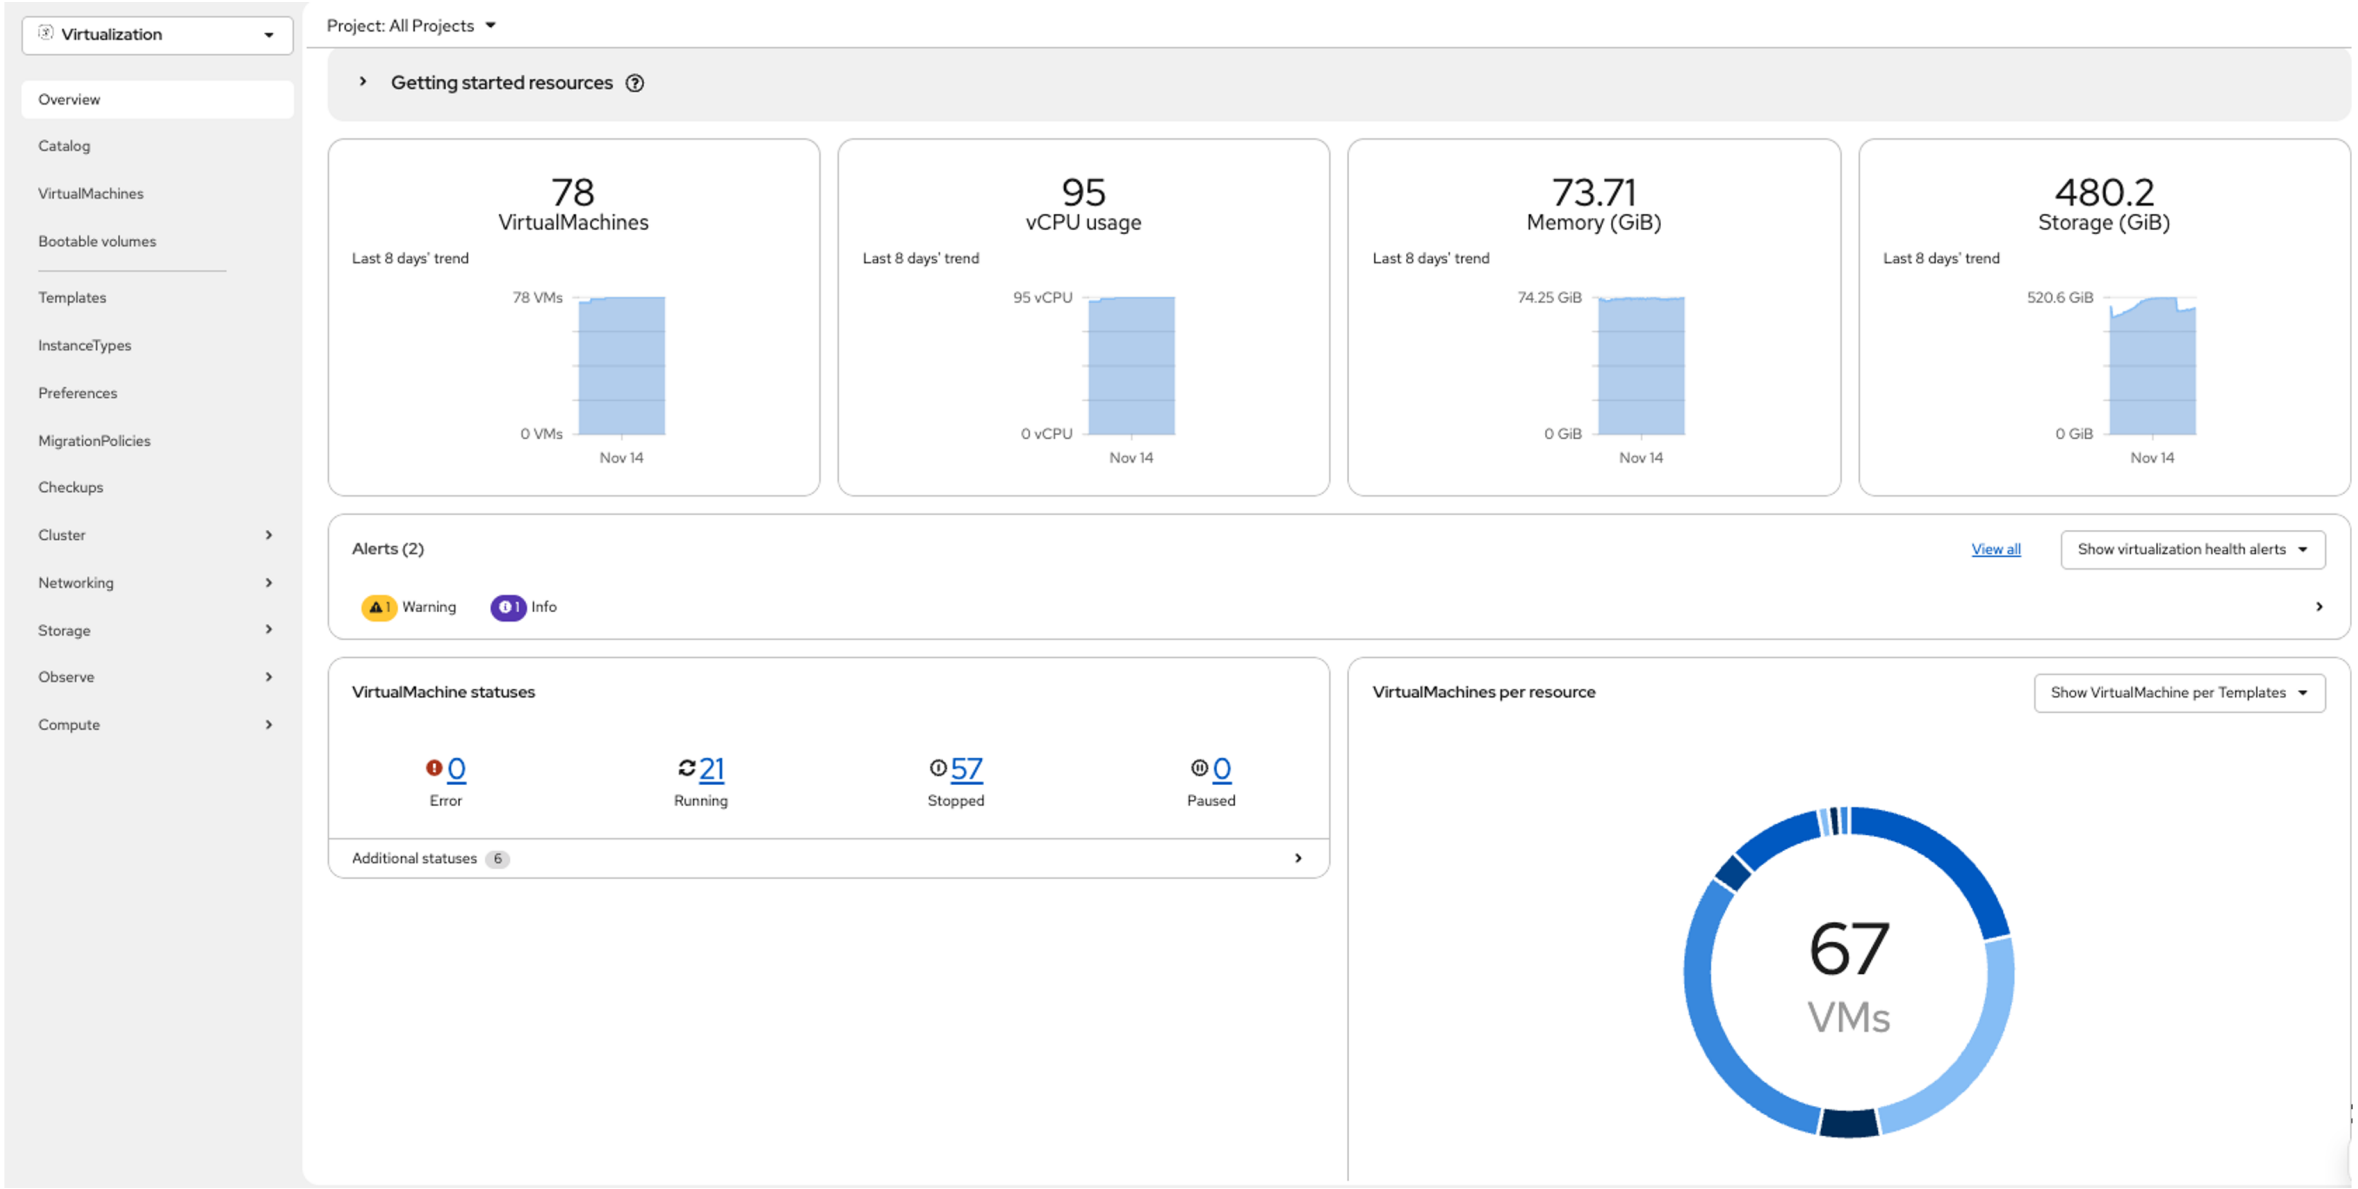

Designed for the Virtual Infrastructure Admin

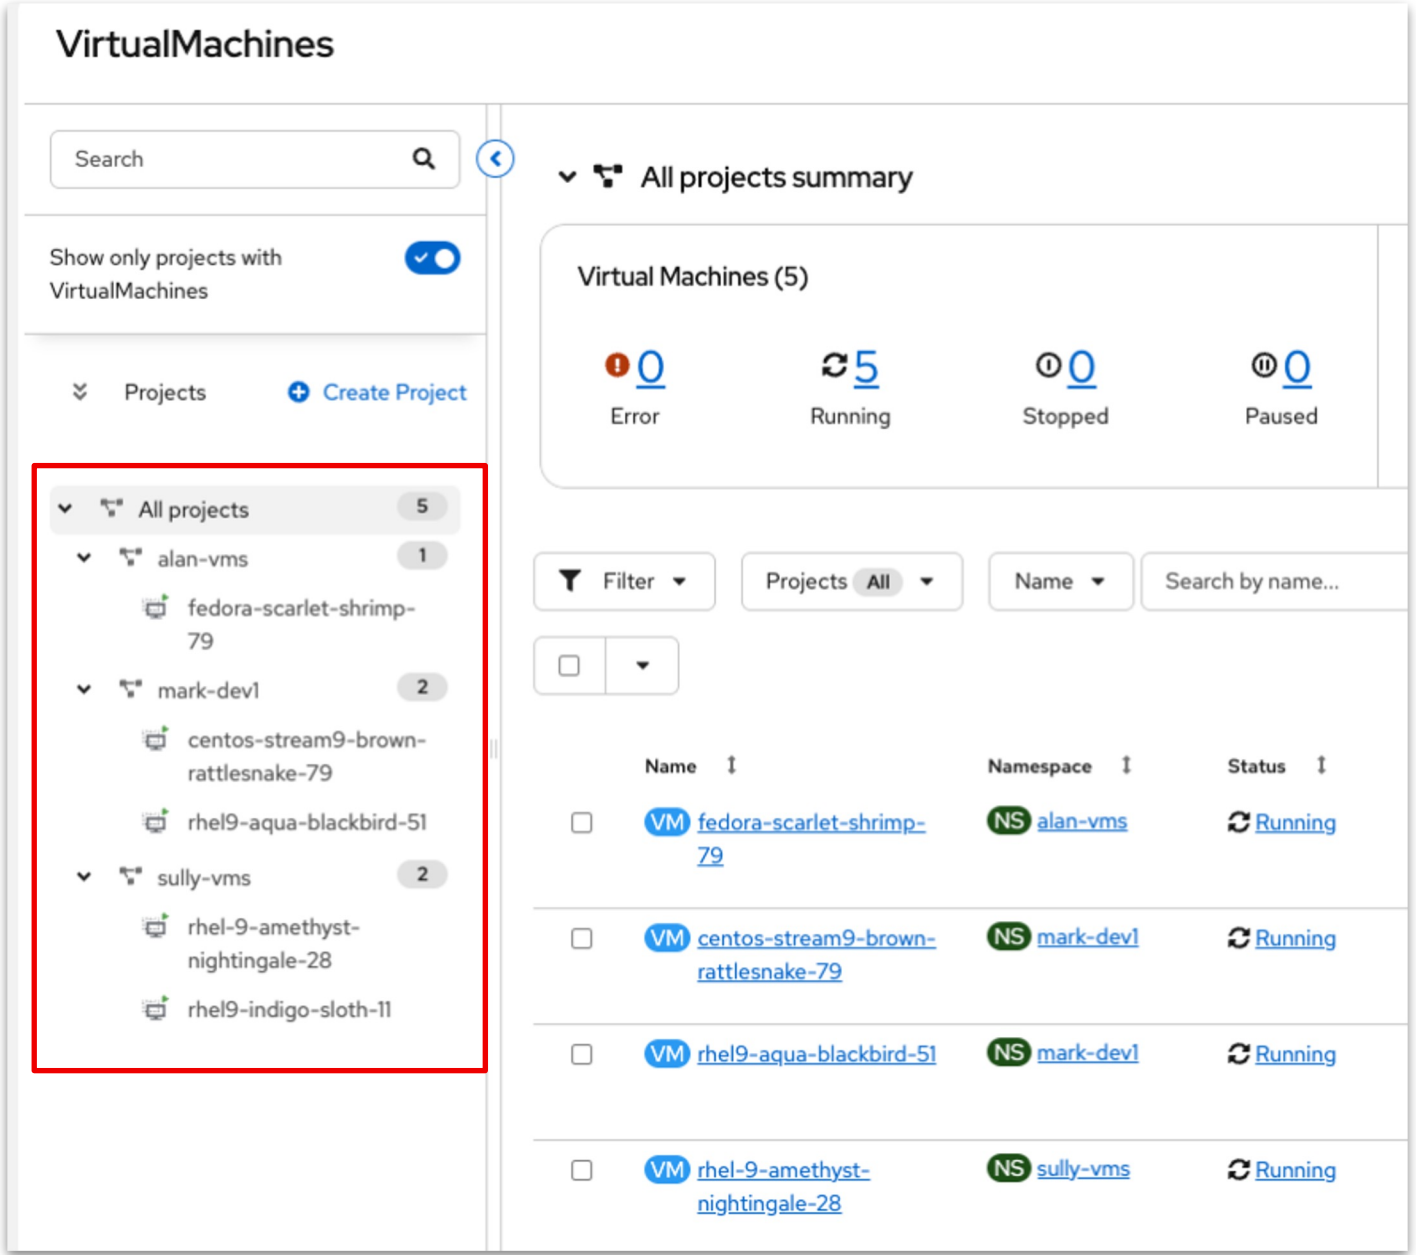

Red Hat

Simplified Management

- Multi-cluster Tree-view
- Optimize cluster for Virtualization with recommended operators
- OpenShift Virtualization installation in Disconnected and Registry-less (TP)
- New VM metrics and alerts for CPU and storage latency

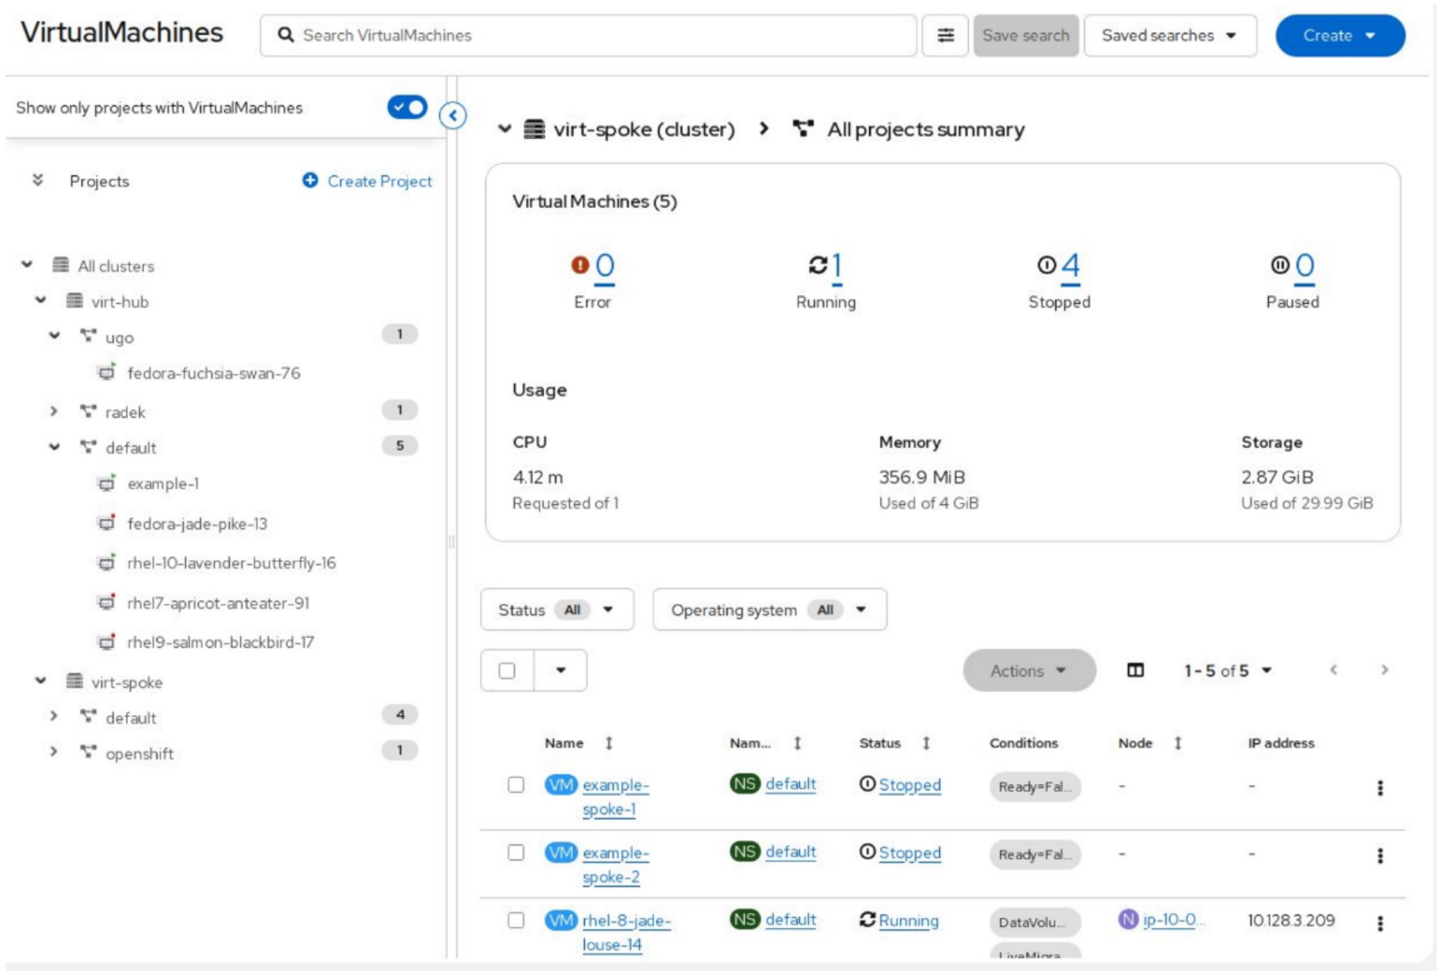

Red Hat

Virtual Machine Details

- Control the state and status of the virtual machine
- Actions menu allows quick access to common VM tasks
  - Start/stop/restart
  - Live migration
  - Clone
  - VM Console
  - Delete
- Change VM Settings like CPU, Memory, Network and Disks

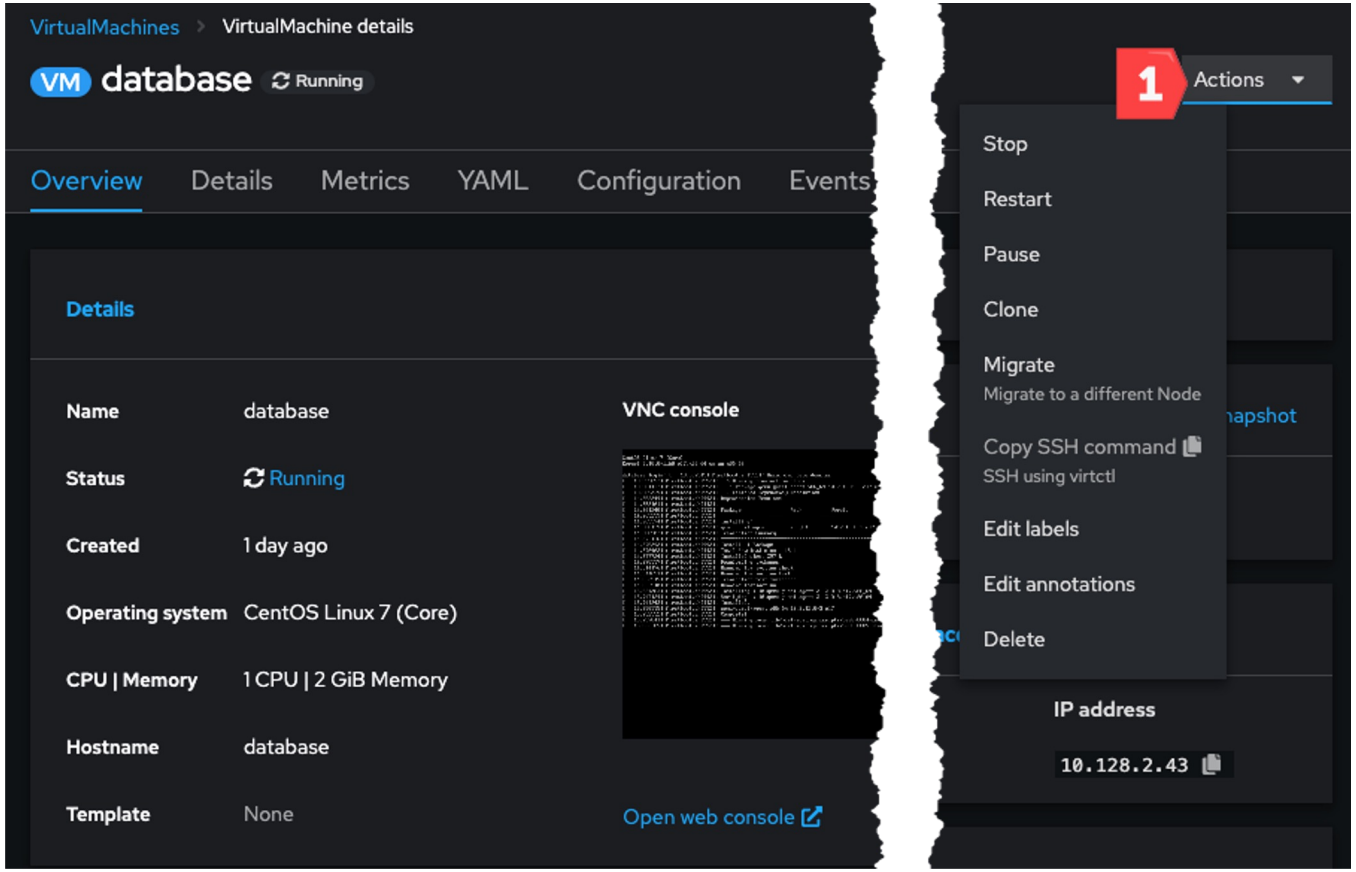

Red Hat

Example Host Network Configuration

- Pod, service, and machine network are configured by OpenShift automatically
- Use the NMstate to configure additional host network interfaces
  - `bond1` and `br1` in the example to the right
- VMs and Pods connect to one or more networks simultaneously

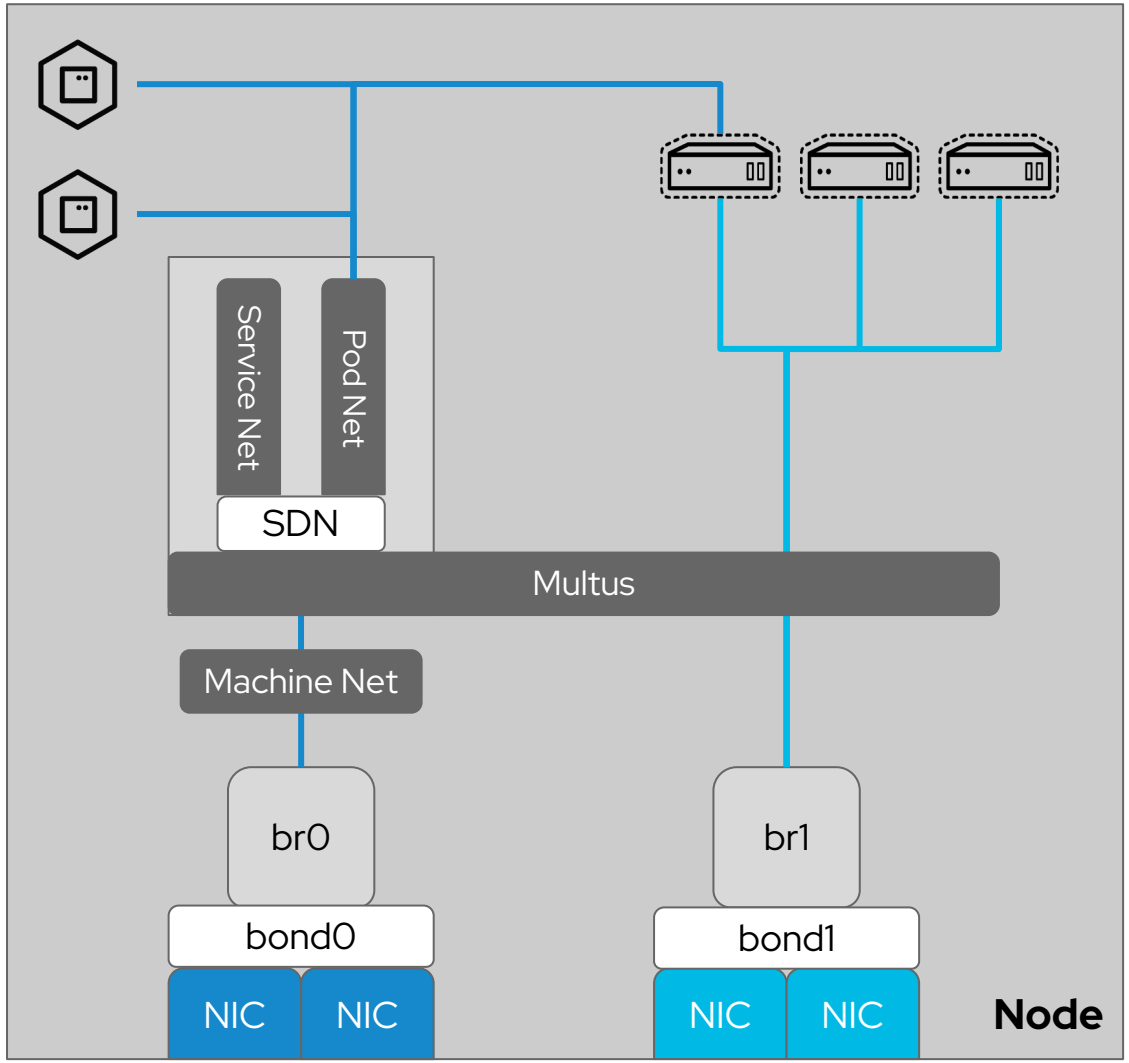

Red Hat

GUI-based host network configuration

- Apply NMstate configuration using a **form** in the OpenShift admin console
- Create and configure
  - Ethernet interface IP (static, DHCP)
  - Bonds - mode 1-6 bonds with options, including IP configuration
  - Linux bridge configuration utilizing ethernet and/or bonds for “uplinks”
- Specify **node selectors** to have configuration automatically applied to matching nodes

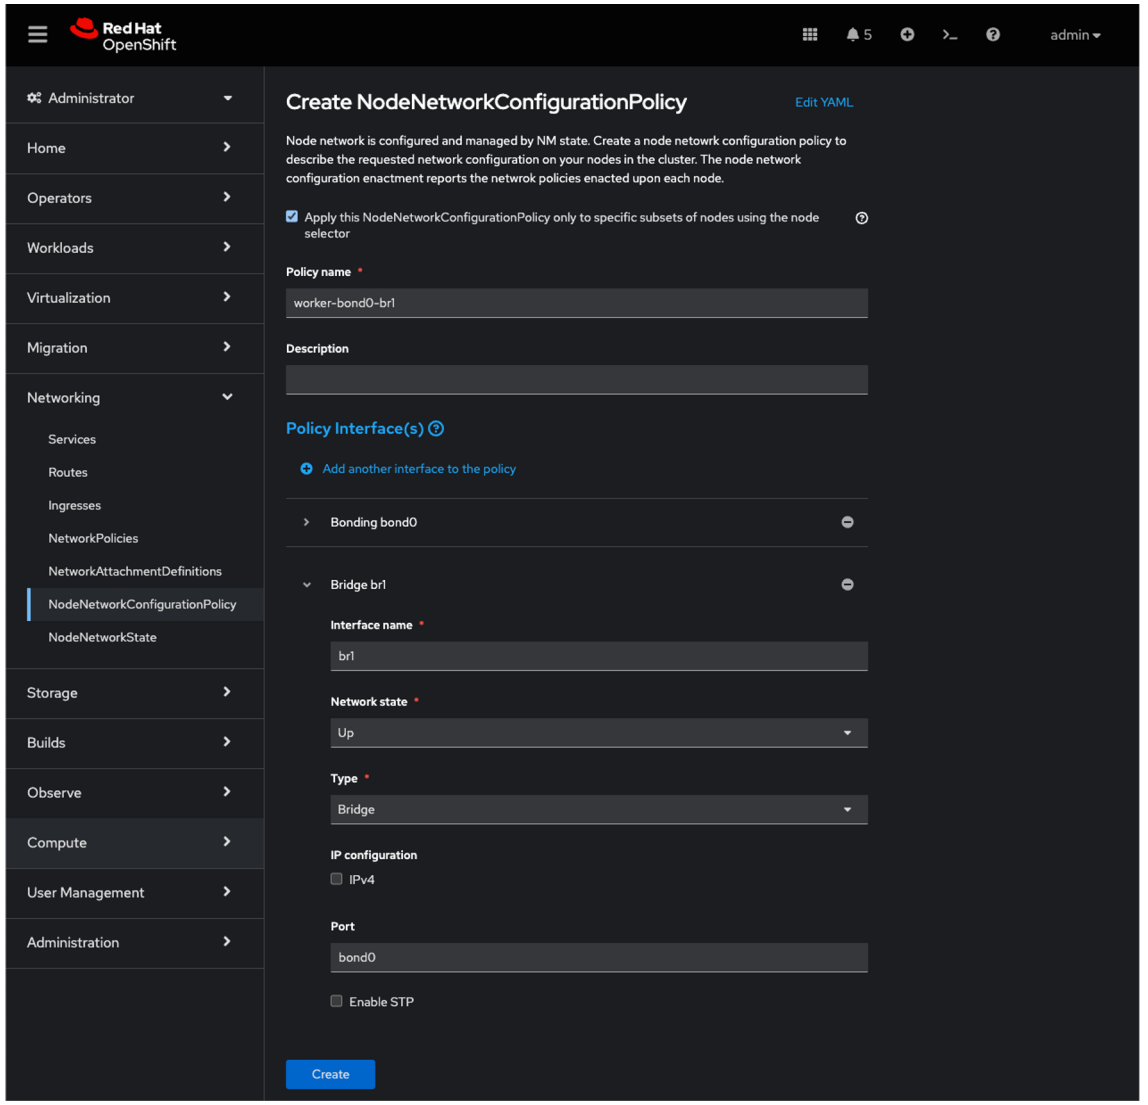

Red Hat

Host bond configuration

- NodeNetworkConfiguration-Policy (NNCP)
  - Nmstate operator CRD
  - Configure host network using declarative language
- Applies to all nodes specified in the `nodeSelector`, including newly added nodes automatically
- Update or add new NNCPs for additional host configs

```yaml
apiVersion: nmstate.io/v1alpha1
kind: NodeNetworkConfigurationPolicy
metadata:
  name: worker-bond1
spec:
  nodeSelector:
    node-role.kubernetes.io/worker: ""
  desiredState:
    interfaces:
    - name: bond1
      type: bond
      state: up
      ipv4:
        enabled: false
      link-aggregation:
        mode: balance-alb
        options:
          miimon: '100'
        slaves:
        - eth2
        - eth3
      mtu: 1450
```

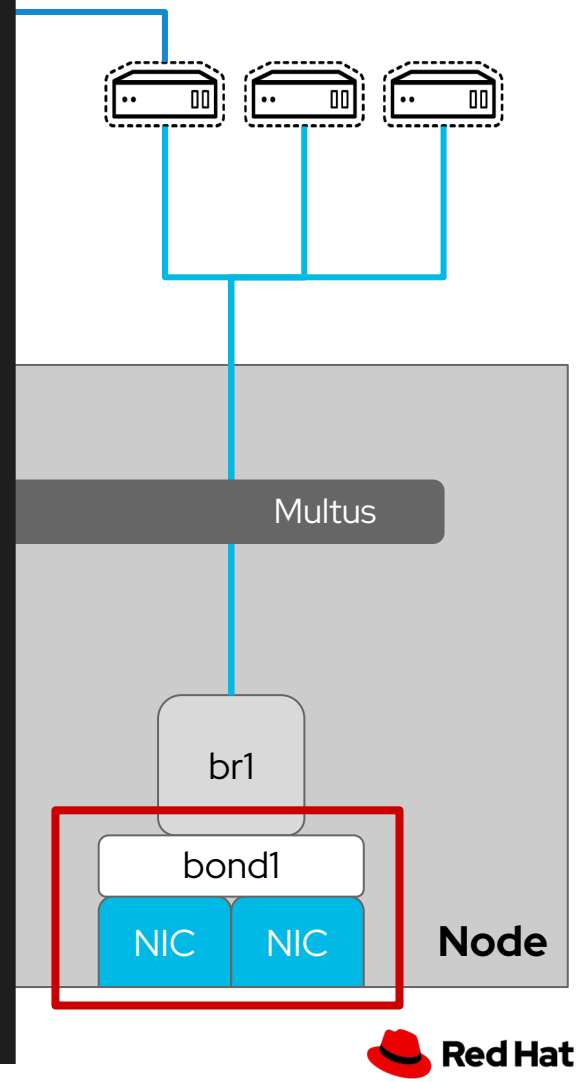

Red Hat

Host bridge configuration

```yaml
apiVersion: nmstate.io/v1alpha1
kind: NodeNetworkConfigurationPolicy
metadata:
  name: worker-bond1-br1
spec:
  nodeSelector:
    node-role.kubernetes.io/worker: ""
  desiredState:
    interfaces:
      - name: br1
        description: br1 with bond1
        type: linux-bridge
        state: up
        ipv4:
          enabled: false
        bridge:
          options:
            stp:
              enabled: false
          port:
            - name: bond1
```

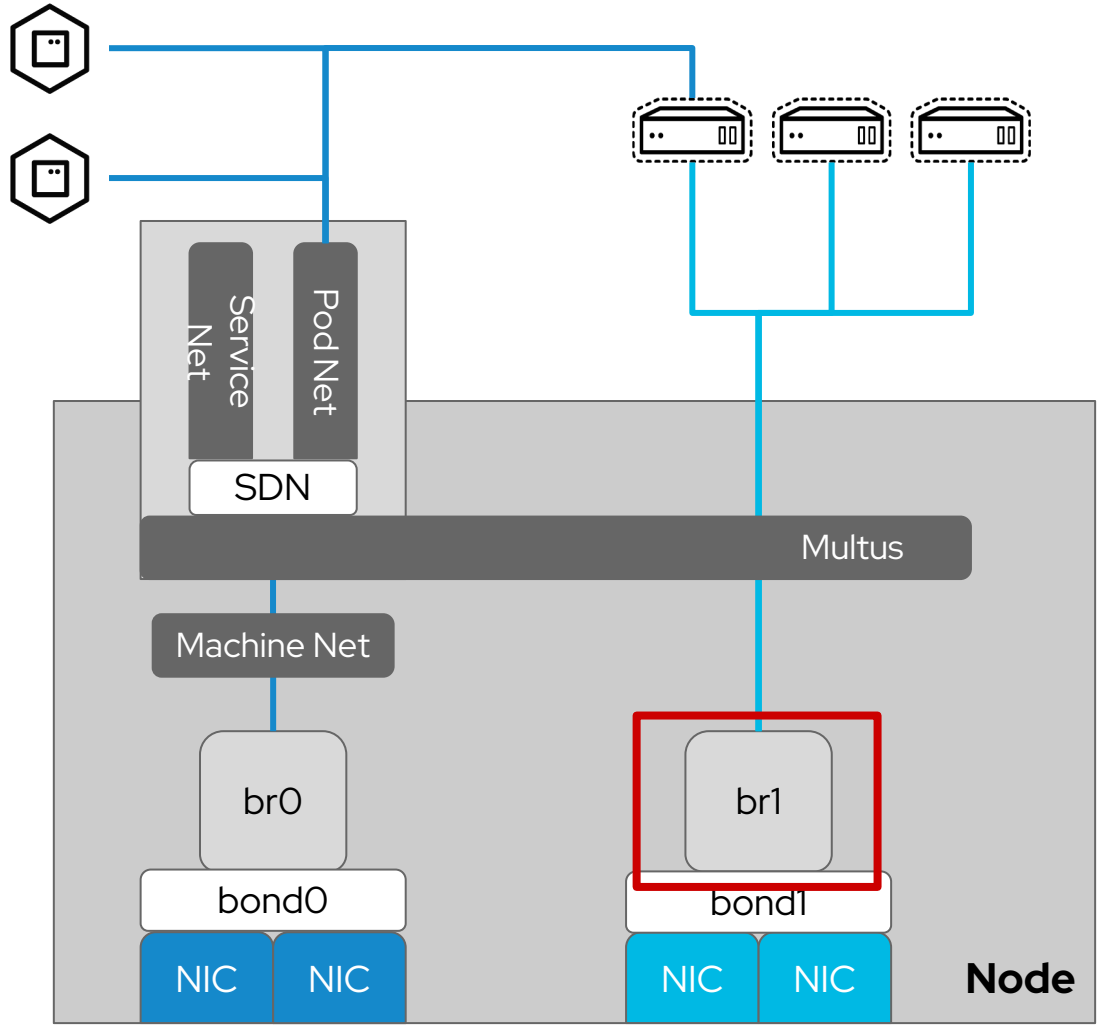

Red Hat

GUI-based Network Attachment Definition

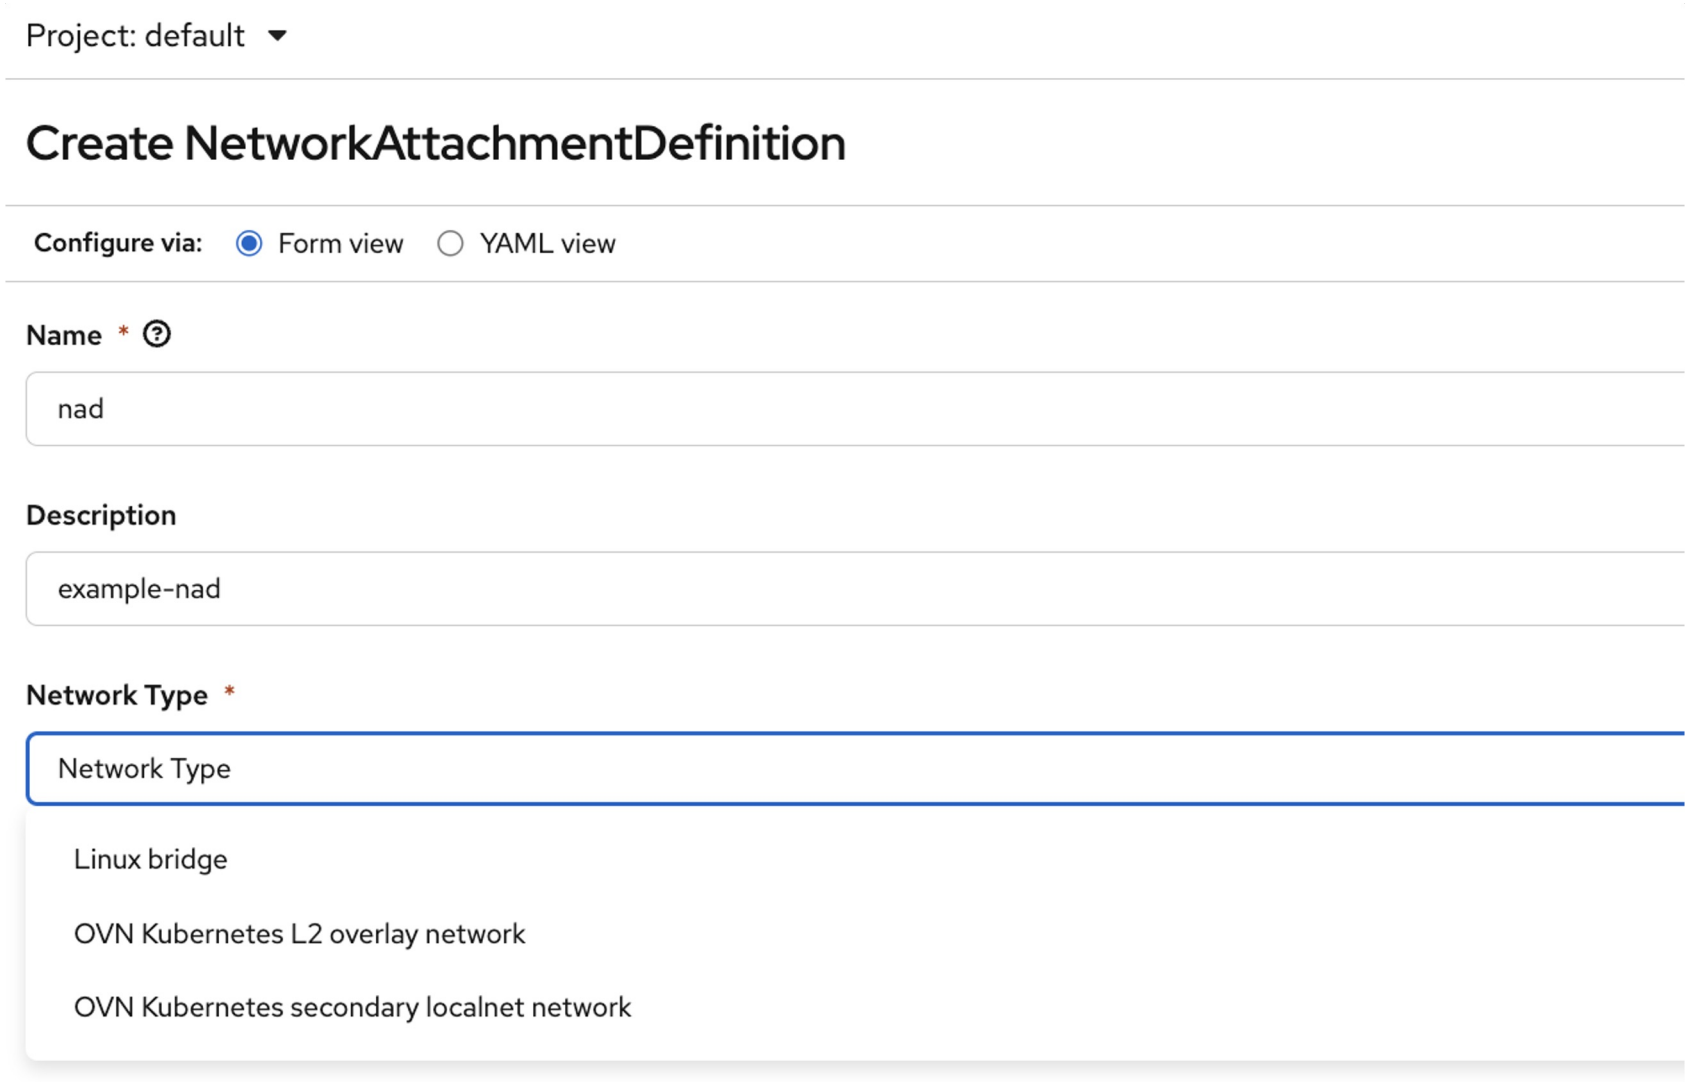

Red Hat

Network Attachment Definition configuration

- `net-attach-def` configures multus to allow the VM to access an underlying resource
  - Optionally define VLAN tags
- Limited to the namespace it's created in
  - Except the `default` namespace, which is available to all

```yaml
apiVersion: k8s.cni.cncf.io/v1
kind: NetworkAttachmentDefinition
metadata:
  annotations:
    k8s.v1.cni.cncf.io/resourceName: bridge.network.kubevirt.io/br1
  name: vlan-93
  namespace: default
spec:
  config: >-
    {
      "name":"vlan-93"
      "type":"cnv-bridge",
      "cniVersion":"0.3.1",
      "bridge":"br1",
      "vlan":93,"
      macspoofchk":true,
      "ipam":{},
      "preserveDefaultVlan":false
    }
```

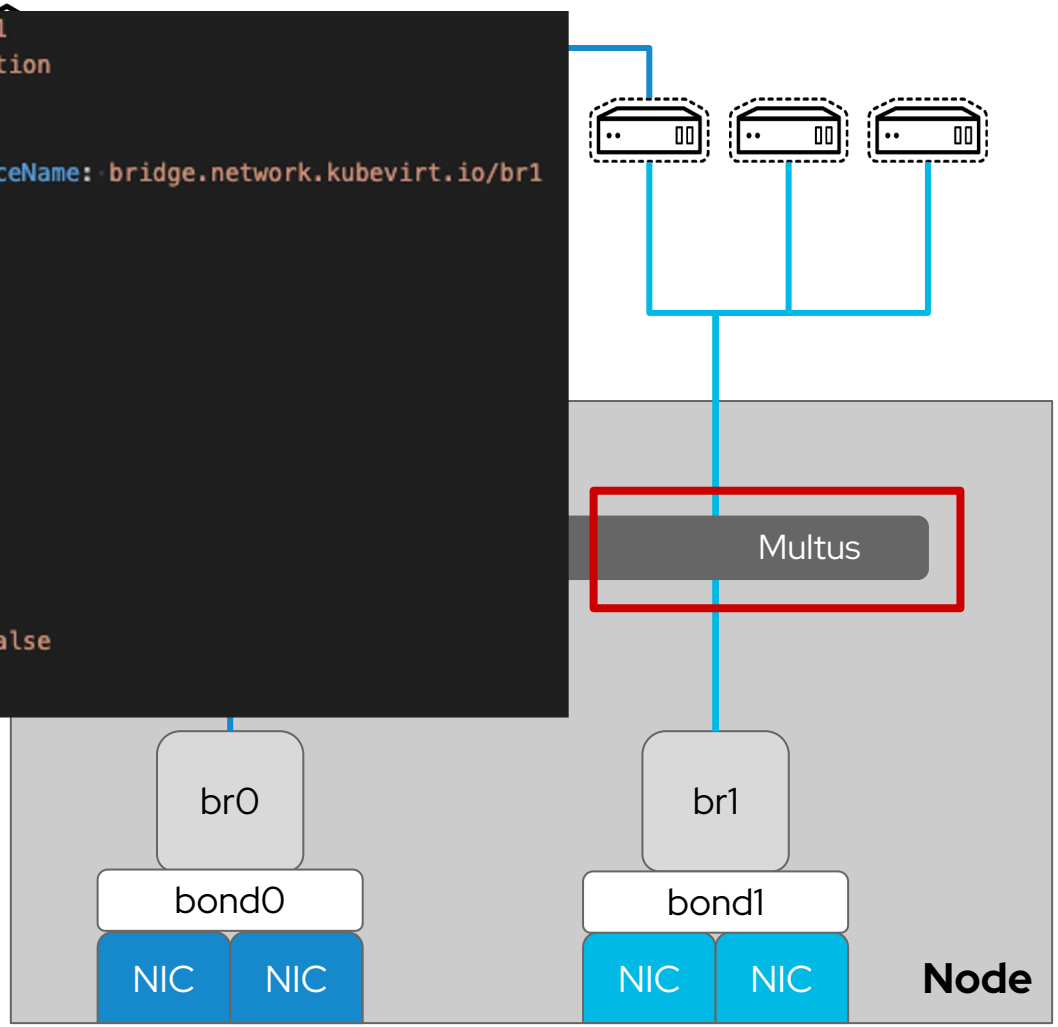

Red Hat

Host network configuration status

- Use the admin console to view the **NodeNetworkState**
- Detailed configuration information for host networking including
  - IP and MAC addresses
  - Bond configuration
  - Bridge configuration
- **Review** and **troubleshoot** host network configuration

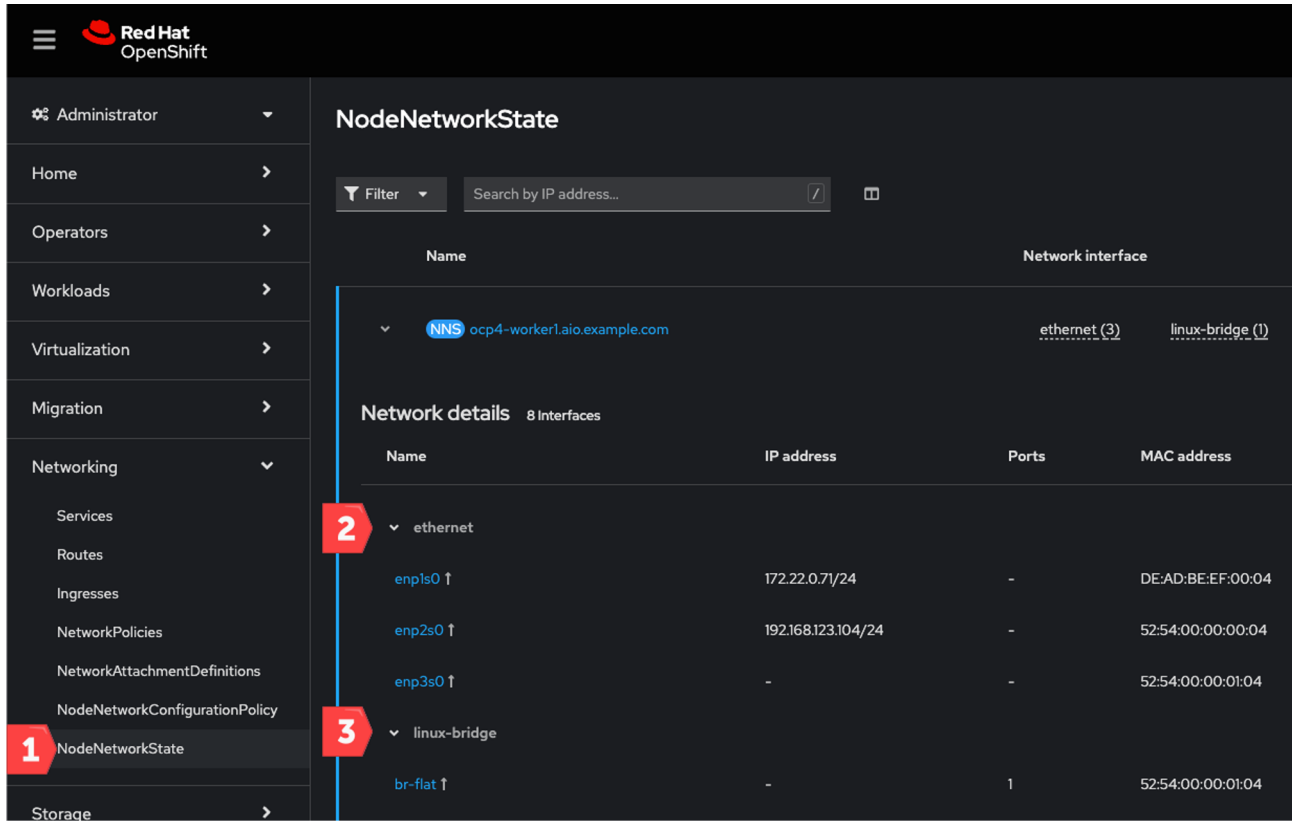

Red Hat

Node Network State (NNS) Topology View

Visualize node network configuration

A graphical representation of the nodes network configuration to provide admins a visual way to see and search the nodes configurations.

Also provides detailed view for each of the network devices

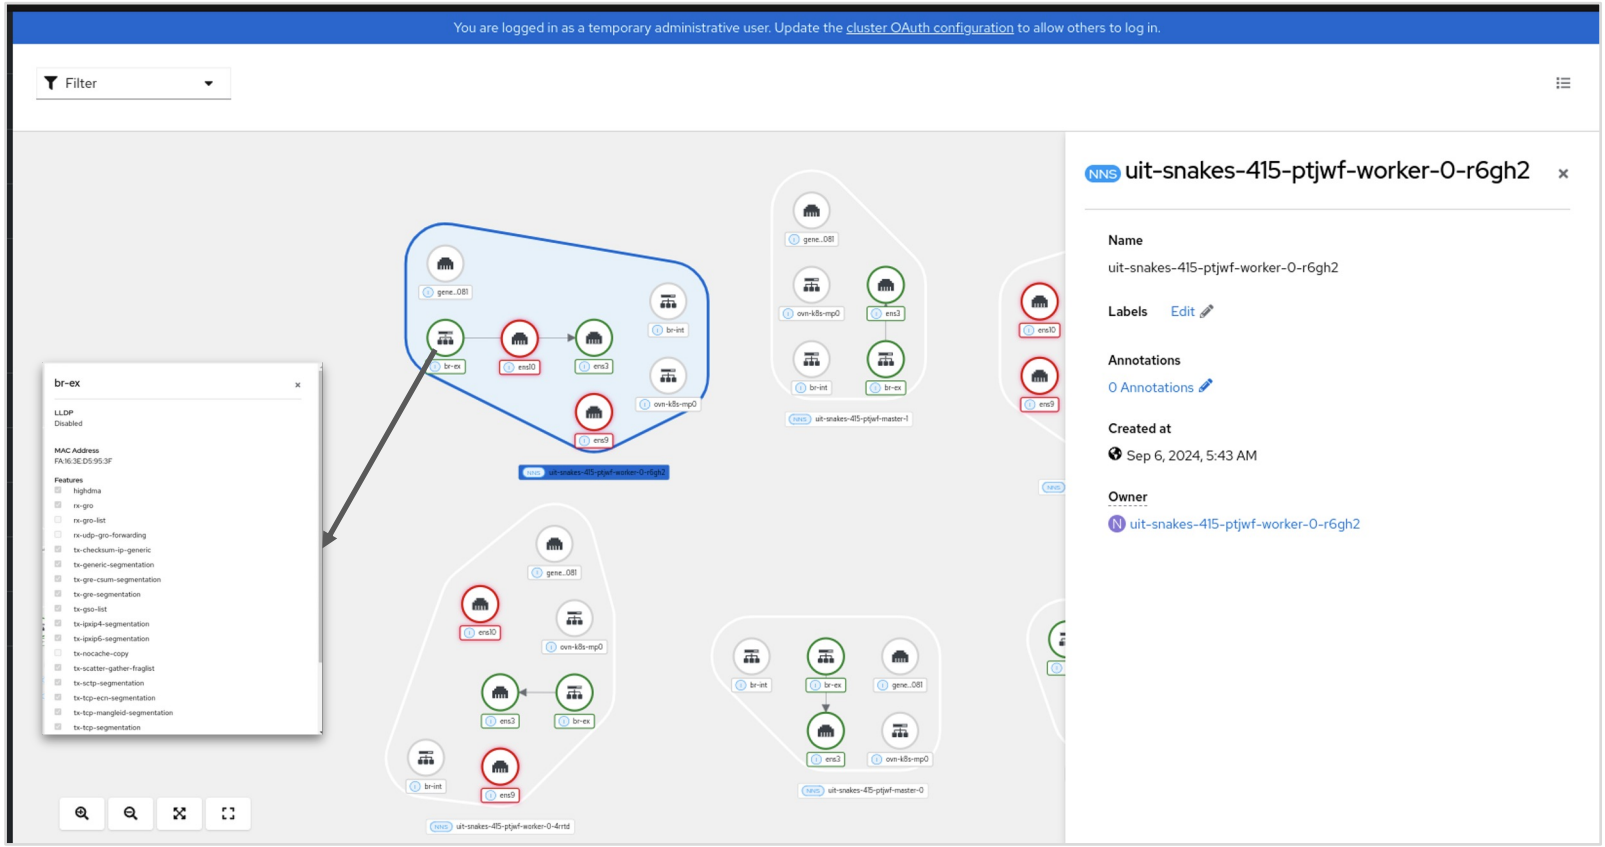

Product Manager: Ronen Sde'Or

Red Hat

Native Network Isolation for Namespaces

A better solution for the monolithic layer 3 Kubernetes pod network

- **User Defined Network (UDN)** support in OVN-Kubernetes
- One or more namespaces in each UDN (tenant)
- Support for:
  - **OpenShift Virtualization**
  - static IP assignments for the life of VMs (for OCP Virt)
  - L2, L3 & localnet UDN topologies
  - overlapping pod IPs across UDNs
  - Kubernetes Network Policy
  - clusterIP services and external services

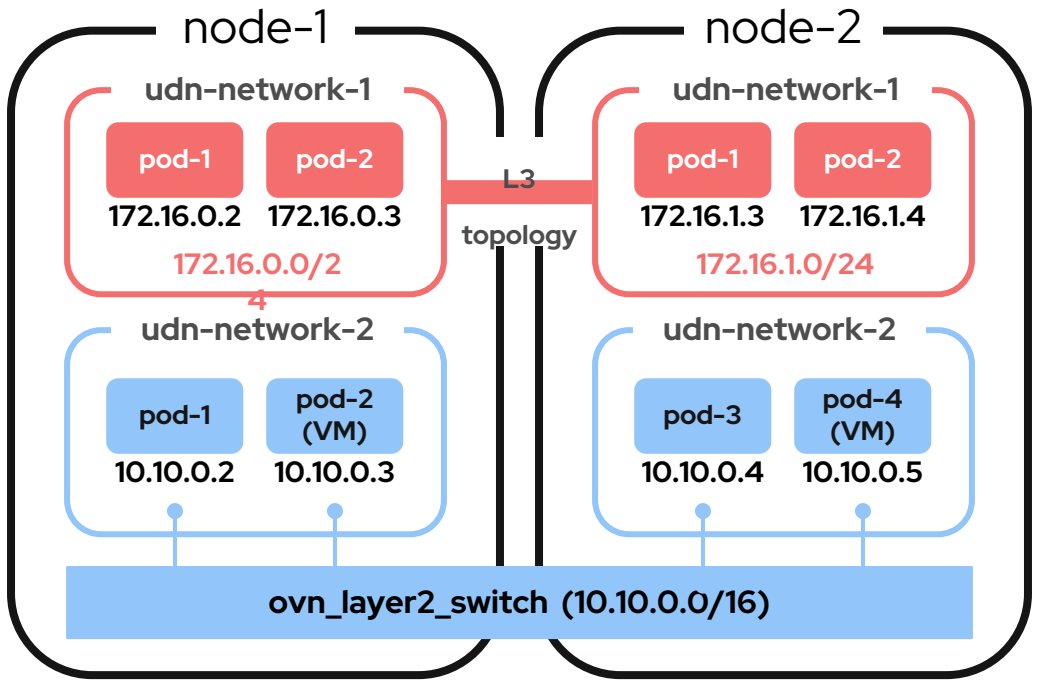

Red Hat

Networking Enhancements

- Routed ingress (BGP) for L2 User Defined Network
- Ability to change the virtual network interface link state of a running VM

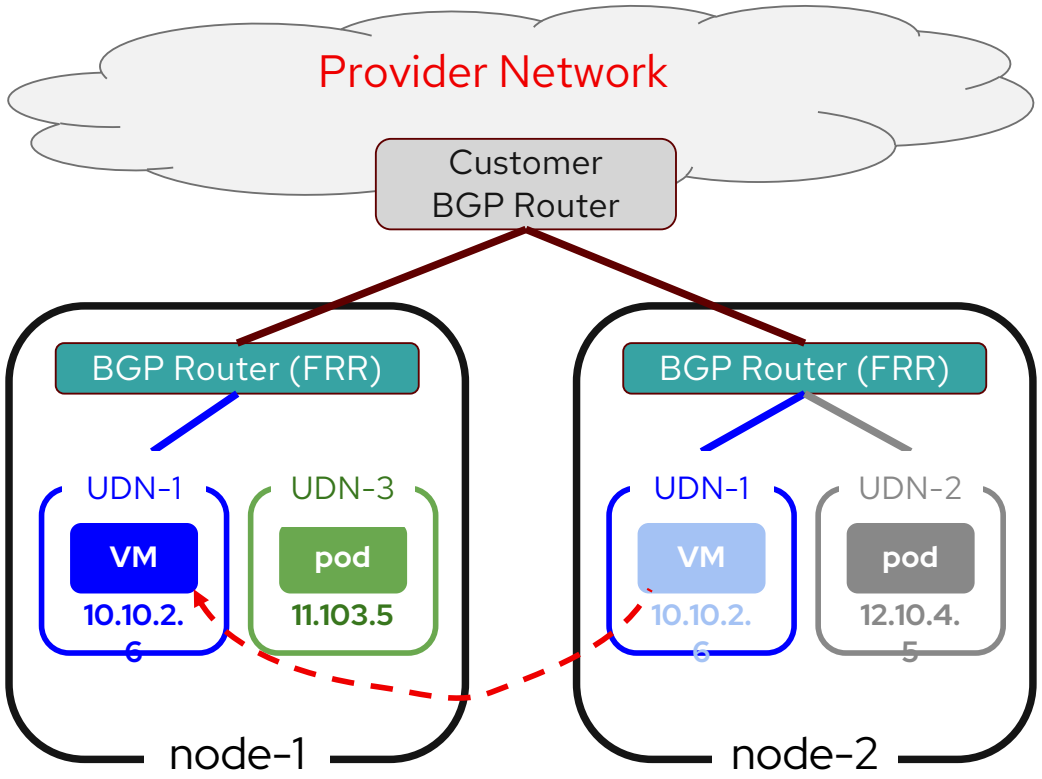

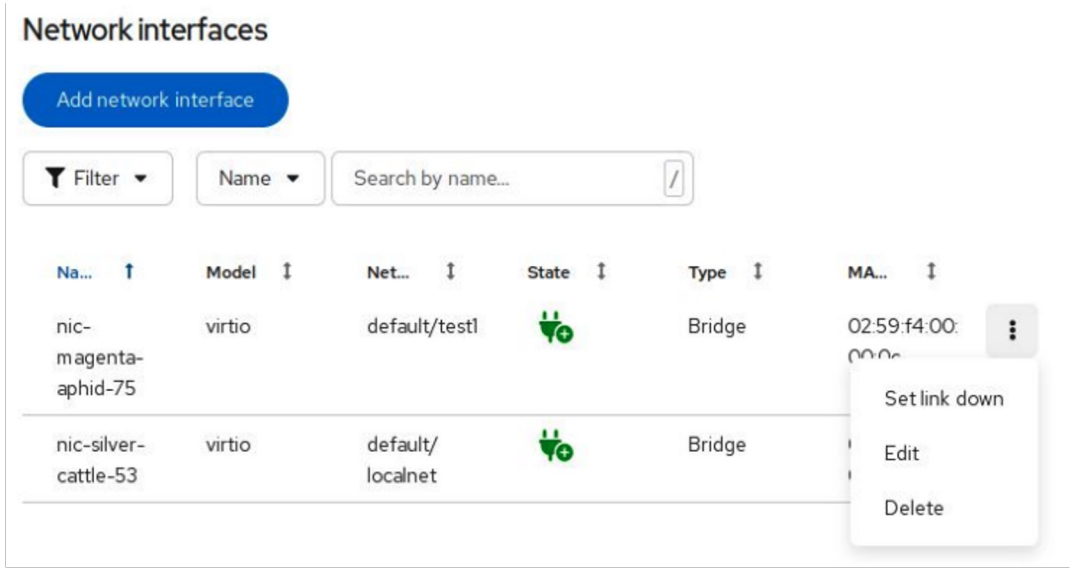

Red Hat

Live Migration

- Live migration moves a virtual machine from one node to another in the OpenShift cluster
- Can be triggered via GUI, CLI, API, or automatically
- Shared storage (RWX) is required
- New Features
  - Live migration to **specific node**
  - **Cross-cluster** VM live migration (TP)

Red Hat

Demo: Live Migration

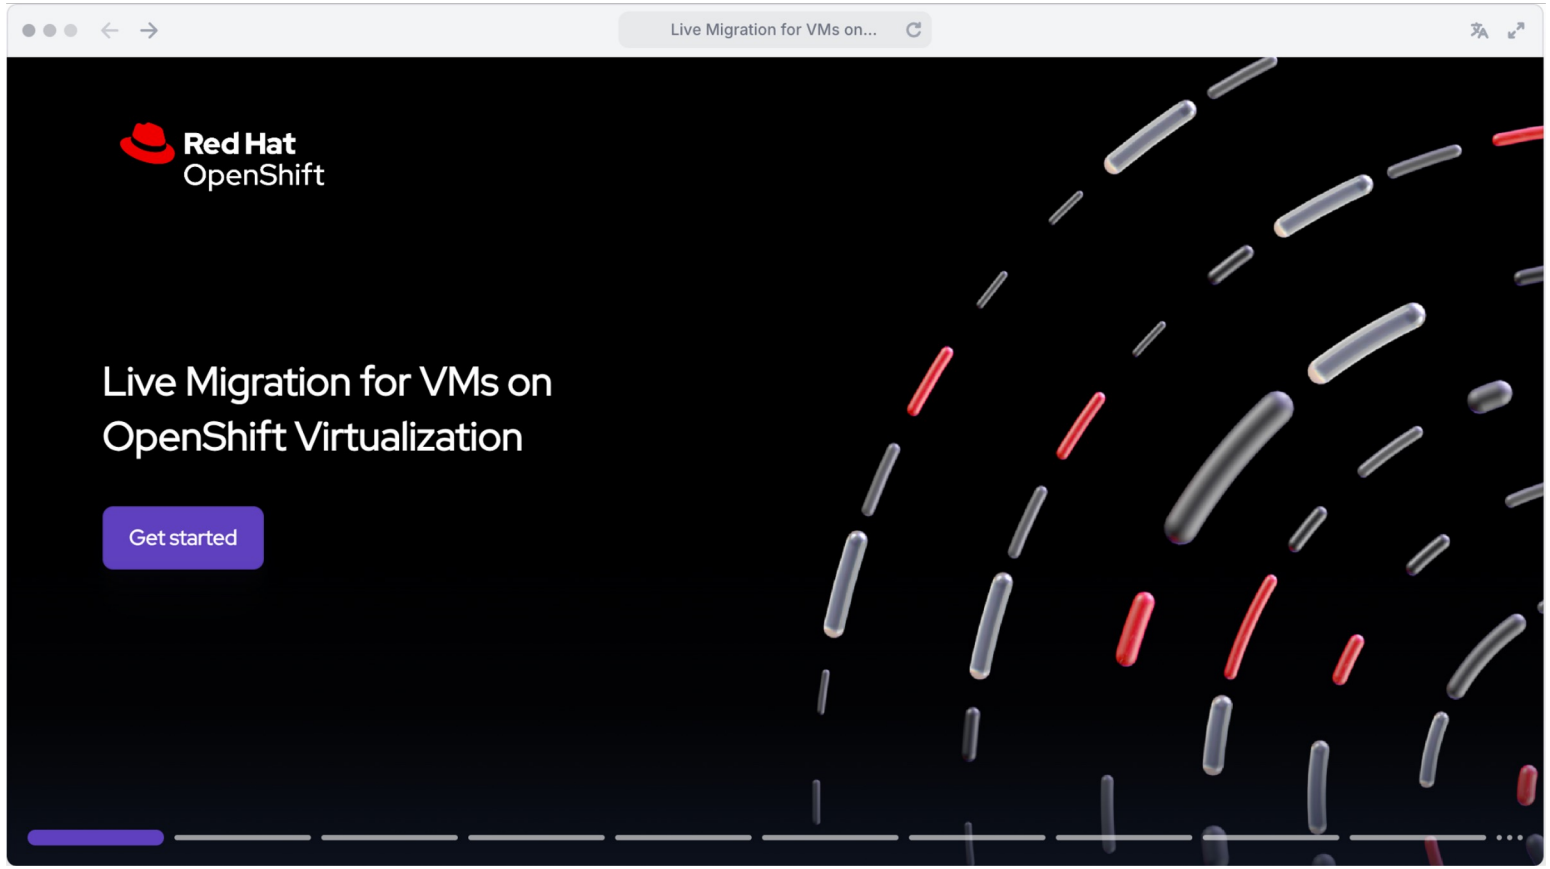

Red Hat

Load-aware Balancing (Descheduler)

- CPU utilization based Automatic VM workload balancing

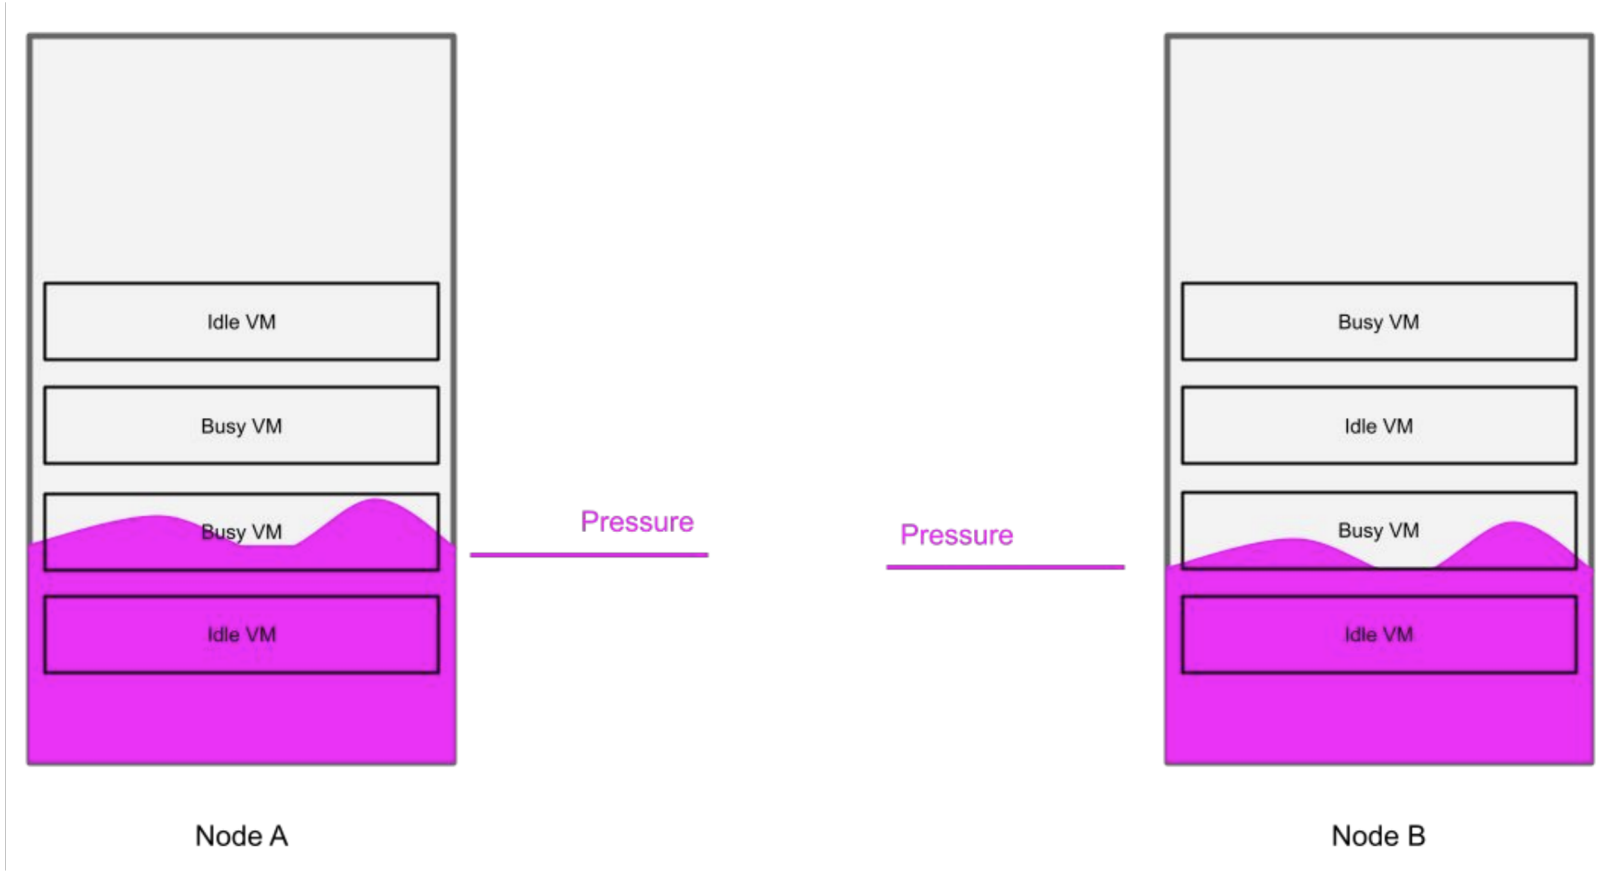

Red Hat

Demo: Load-aware Balancing (Descheduler)

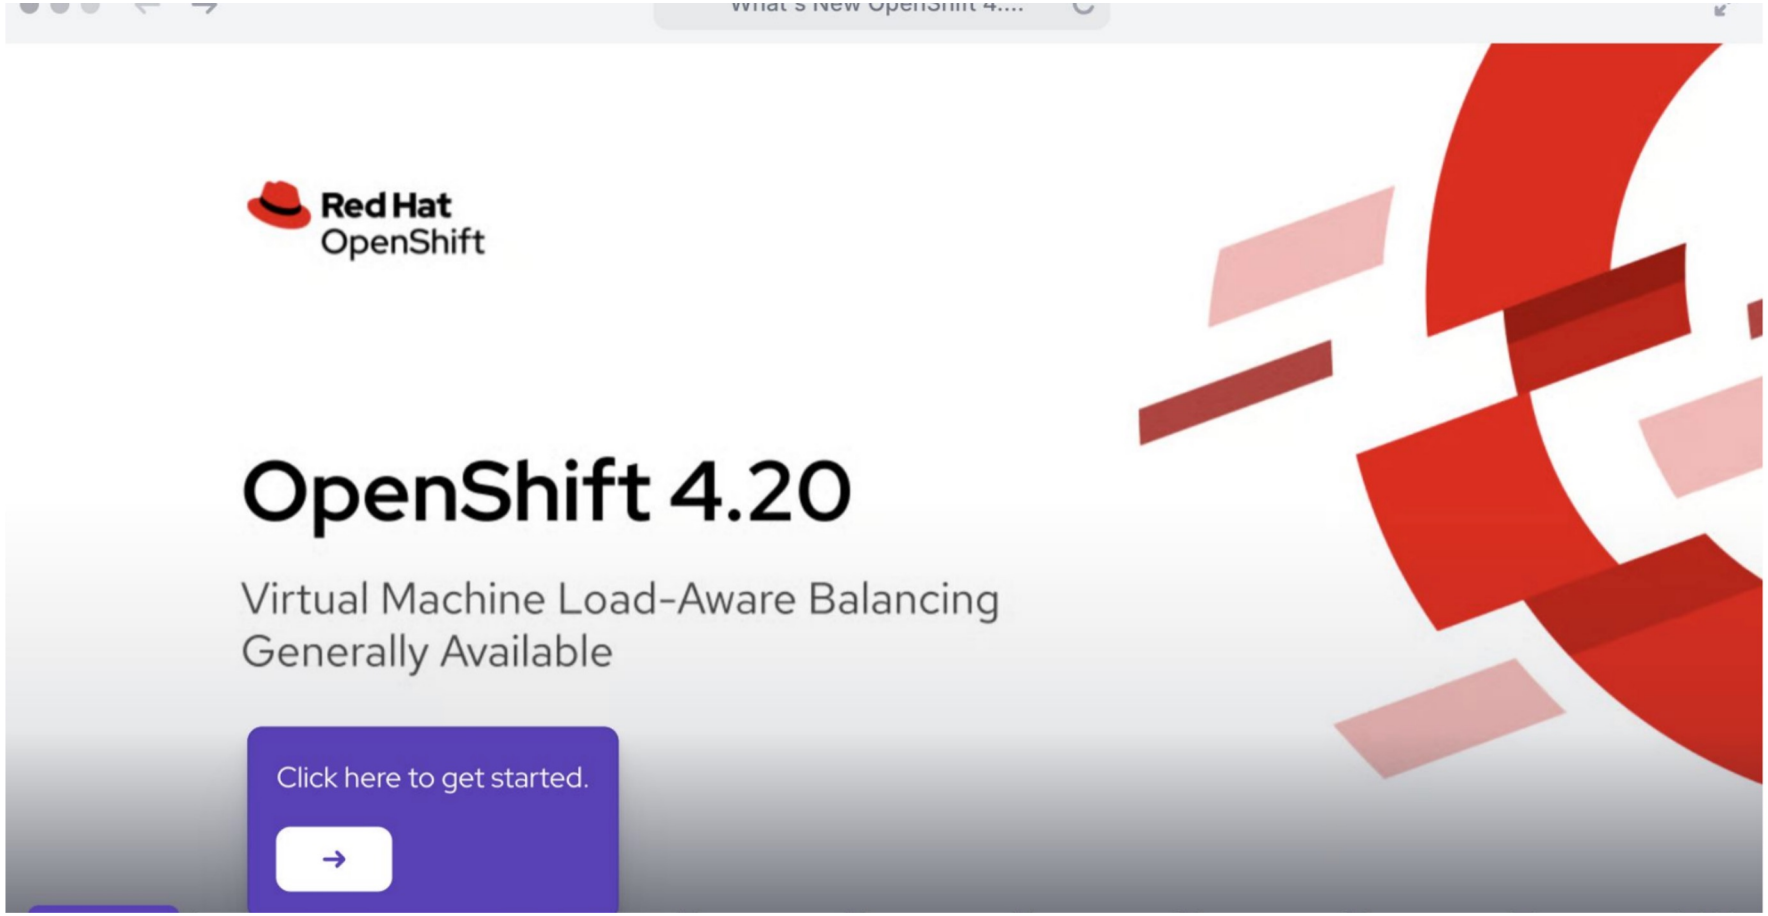

Red Hat

Migrate your virtual machines with minimal disruption

Migration toolkit for virtualization (MTV) included with OpenShift

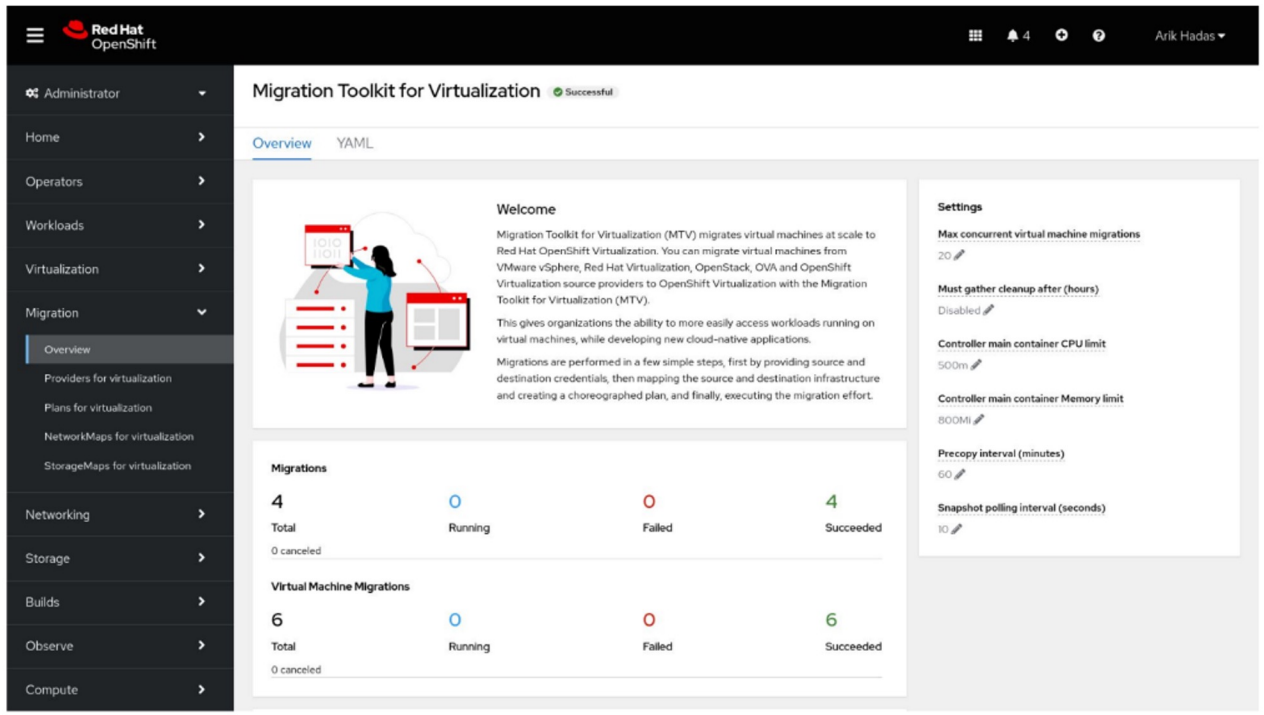

VM migrations made simple

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

Red Hat

Migration toolkit for virtualization

Making mass migration possible

Easy to use UI

Simple, easy to use interface

Concurrent migrations

Easily migrate VMs concurrently 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

Storage offloading

Accelerate migrations with underlying storage systems versus IP network (Tech Preview MTV 2.10)

Migration network selection

Avoid impact on other running workloads

Red Hat

Accelerate VM migrations with existing storage systems

Fast, low-risk migrations with storage offloading

Faster than traditional migrations

- Reduces IP network load
- Offloads compute resources to the storage system
- Migrations up to **10x faster***

Compatible with supported certified storage partners

- Partner list growing

Hitachi Vantara

NetApp®

Hewlett Packard Enterprise

DELL

PURESTORAGE®

*Based on Hitachi's internal testing

Red Hat

Two Node OpenShift with Arbiter (TNA)

New Topology Approach:

- **Two node** solution for cost sensitive customers who do not need a full third node for their workload
- Small arbiter node (2 vCPU, 8Gi), running only third etcd instance
- **Same HA characteristics** as a regular three node cluster - tolerates single node outage
- **OpenShift Virtualization** fully supported
- Hyperconverged Storage / Software Defined Storage (replica 2, disks only on node 1 and 2) via partners:
  - Pure / Portworx
  - Arctera / Infoscale
  - IBM / Fusion

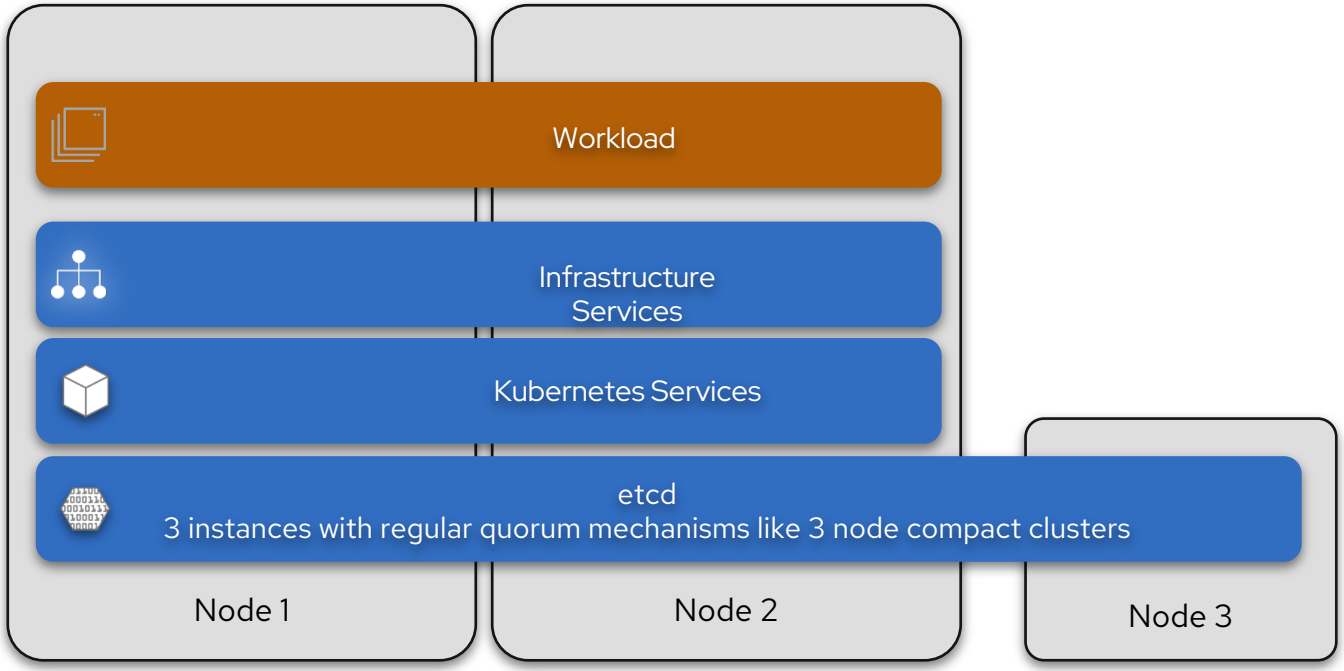


```
% oc get nodes
NAME  STATUS  ROLES
node1 Ready   master,worker
node2 Ready   master,worker
node3 Ready   arbiter
```

Operational Guidance with OpenShift Lightspeed

Generative AI based conversational assistant for OpenShift

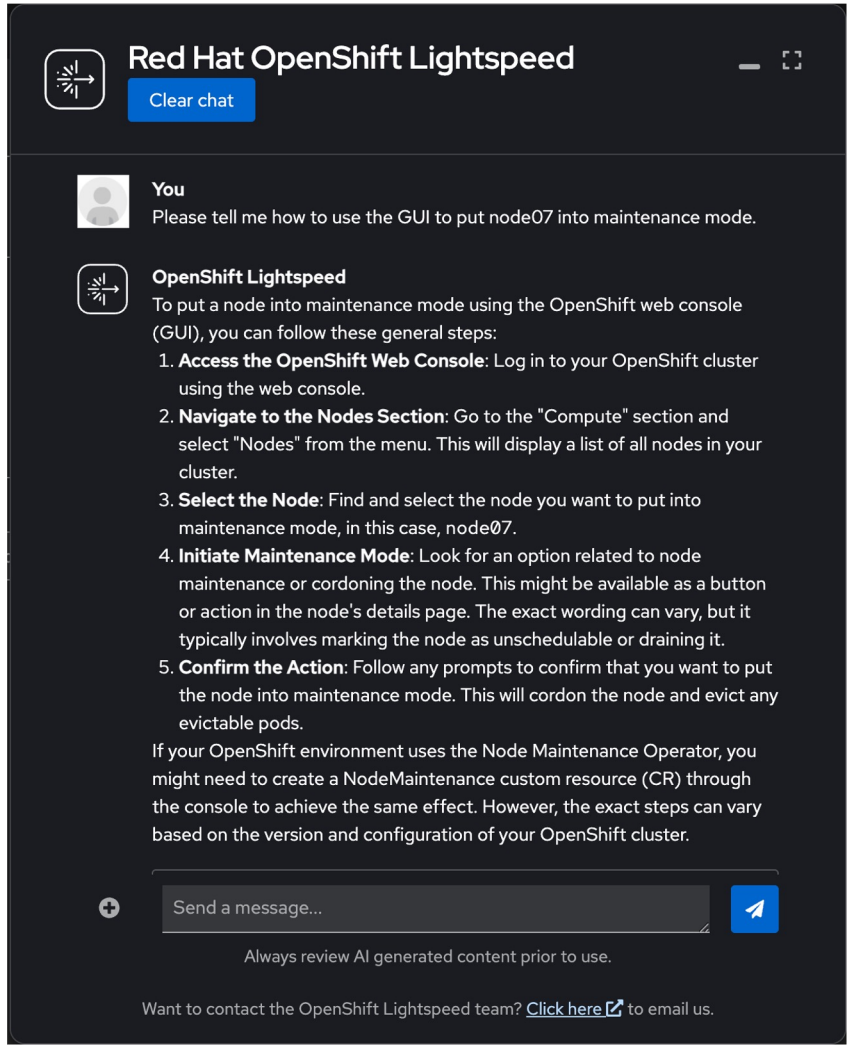

- Generative AI-based conversational assistant in OpenShift Console, accessible throughout the Management Console
- Assists traditional vAdmins transition to OpenShift Virtualization
- Get product help and helps guide daily tasks and troubleshooting
- Context Aware, and trained on VMware -> OpenShift product terms
- Bring your own LLM provider
- For both VMs and containers

Red Hat

Easy/Guided Install

Installation Wizard tuned for OpenShift Virtualization (Tech Preview)

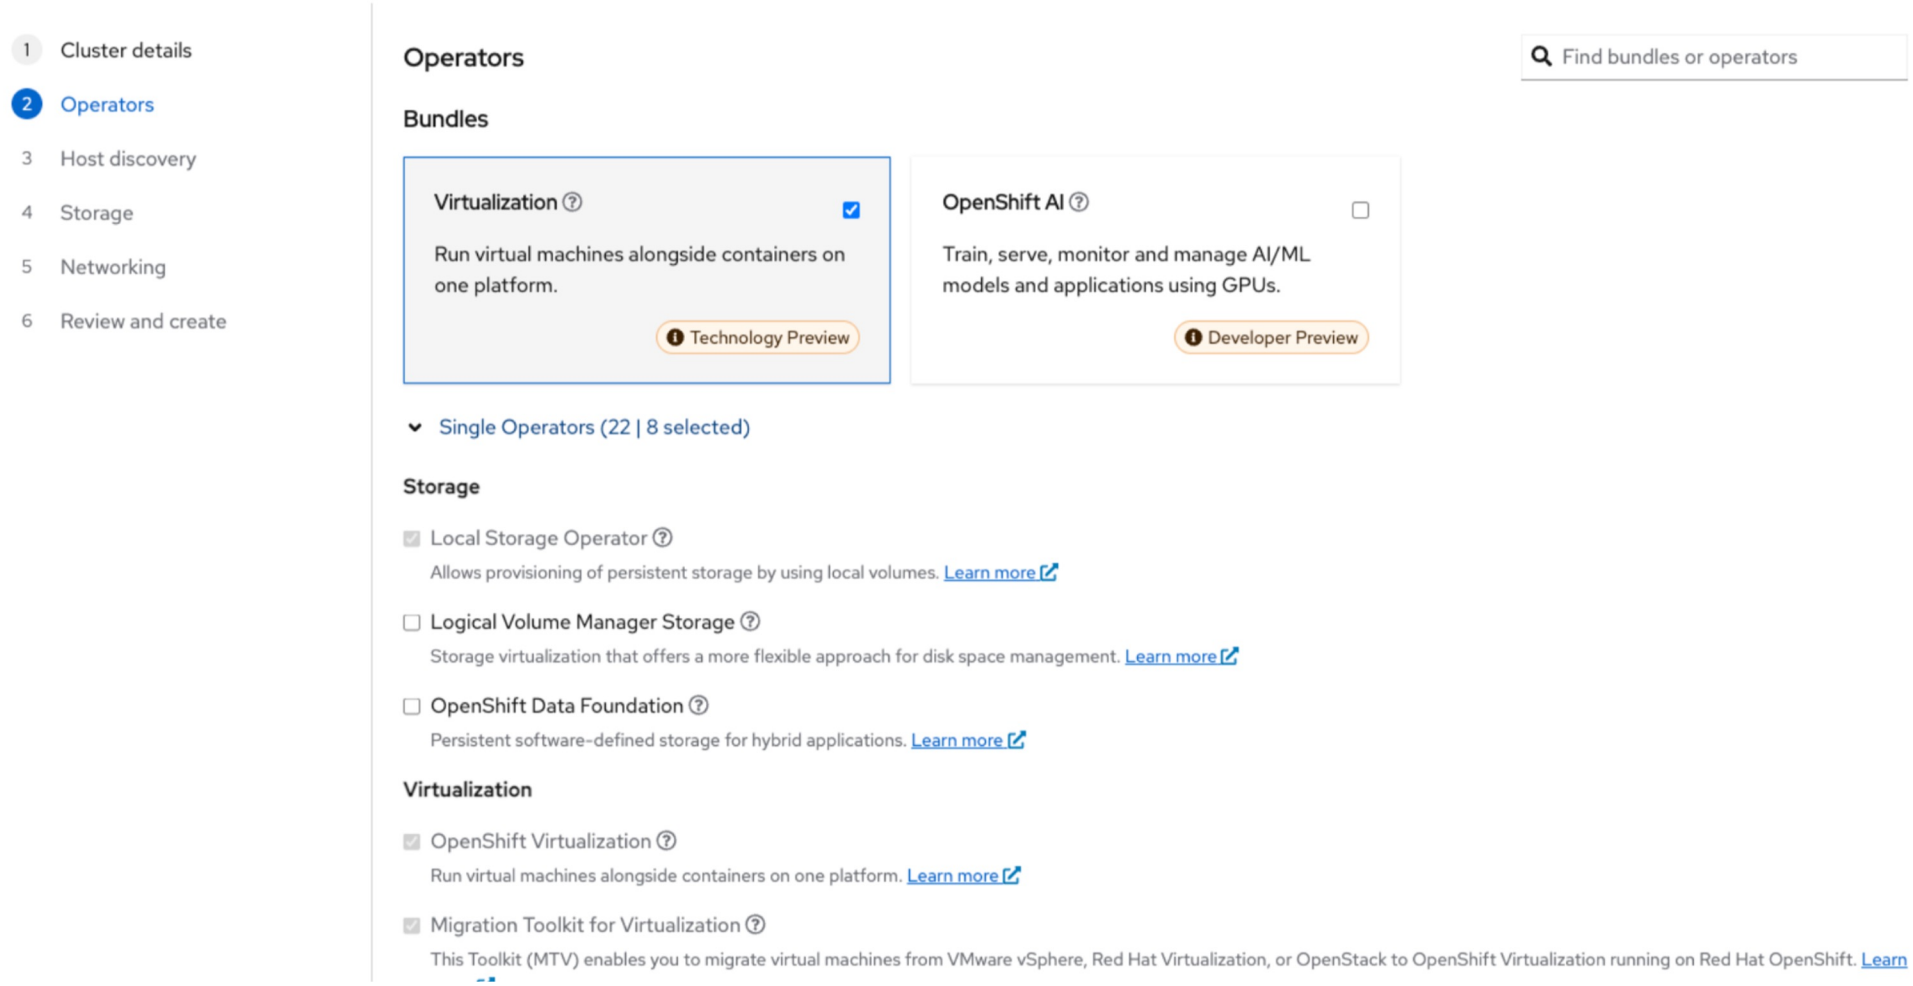

Red Hat

Red Hat OpenShift Virtualization
What's Next

Ease of Migration

- **Live migration of VMs and Storage across clusters**
- GA of Storage accelerated migration to OpenShift
- Migration from Hyper-V
- Self Guided Migration assistant

Networking

- **BGP and EVPN for stretching networks across clusters**

Simplified & Scaled Administration

- **Extend single cluster experience to multiple clusters**
- VM Right Sizing operations
- Guided networking configuration

Edge to Cloud

- **Two node HA Solution (no arbiter needed) for Edge deployments**
- GA support for ARM platforms
- GA of Oracle OCI, GCP, Azure ARO

Storage & Data Protection

- **Change block tracking for Incremental backup**

Security and Performance

- **Optimized defaults for multiple live migrations and databases**
- VM vulnerability reporting in RHACS
- Compliance Operator for Hardened OpenShift Virtualization

Red Hat

Red Hat Summit

Connect

Thank you

linkedin.com/company/red-hat

facebook.com/redhatinc

youtube.com/user/RedHatVideos

twitter.com/RedHat