okd

OKD Virtualization

OKD Virtualization: *what's new, what's next*

Simone Tiraboschi

Principal Software Engineer - Red Hat

Red Hat

OKD

Virtualization

Agenda

- Intro: what are we talking about?
- CRC + OKD Virtualization: the simplest way to try OKD Virtualization @home
- New features:
  - Golden images
  - Kubevirt Tekton tasks
- What's next

What’s OKD?

OKD is a distribution of Kubernetes

OKD embeds Kubernetes and extends it with security and other integrated concepts

OKD adds developer and operations-centric tools

OKD is a sibling Kubernetes distribution to Red Hat OpenShift Container Platform (OCP for short)

https://flickr.com/photos/64711971@N08/6918935136/
“OKD & Kubernetes”

Governance:

Red Hat

THE UNIVERSITY OF MICHIGAN
ARTES · SCIENTIA · VERITAS
1817

What's OKD?

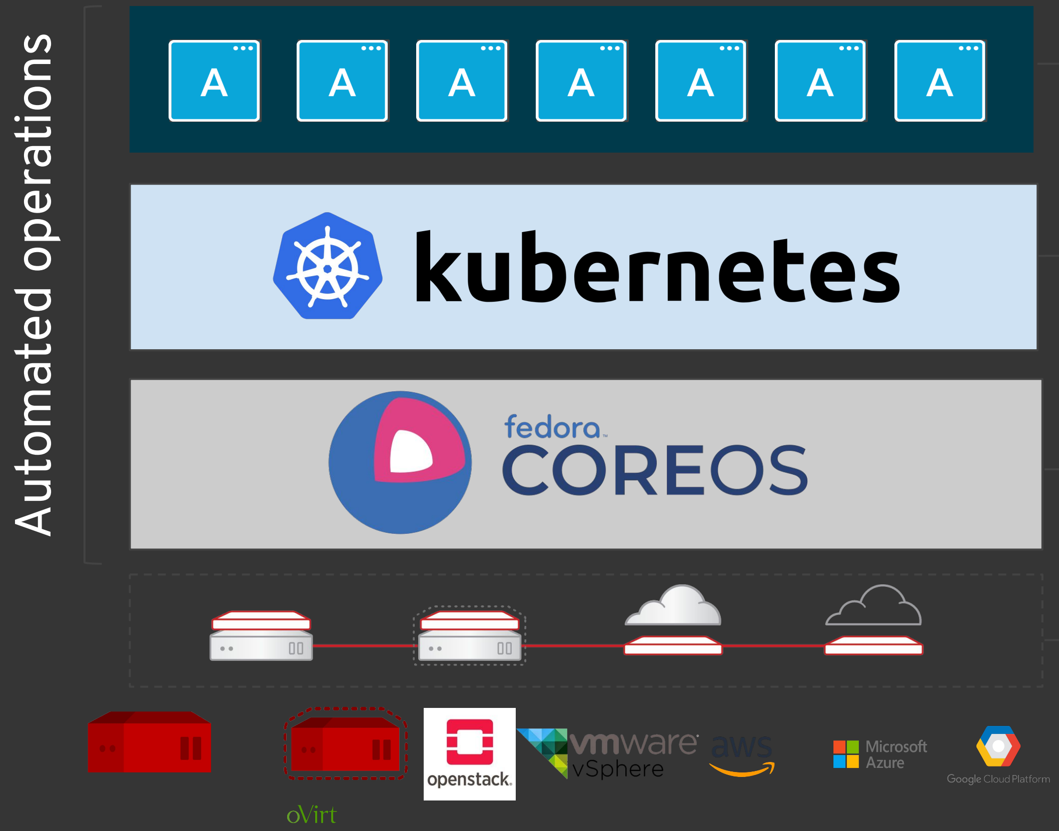

APPLICATIONS AND SERVICES
from Red Hat and community operators

PLATFORM AND CLUSTER MANAGEMENT
Kubernetes, security, monitoring, registry. etc

LINUX HOST with Fedora CoreOS*

FOR HYBRID / MULTI-CLOUD DEPLOYMENTS

What’s KubeVirt?

Kubernetes Virtualization API and runtime in order to define and manage virtual machines:

- Virtual machines
  - Running in containers
  - Using the KVM hypervisor
- Scheduled, deployed, and managed by Kubernetes
- Integrated with container orchestrator resources and services
  - Traditional Pod-like SDN connectivity and/or connectivity to external VLAN and other networks via multus
  - Persistent storage paradigm (PVC, PV, StorageClass)

https://flickr.com/photos/linein/2946303389/

“A Container & VM”

Why KubeVirt? Using VMs and containers together

- Follows Kubernetes paradigms:
  - Container Networking Interface (CNI)
  - Container Storage Interface (CSI)
  - Custom Resource Definitions (CRD, CR)
- Schedule, connect, and consume VM resources as container-native
- Virtual Machines connected to pod networks are accessible using standard Kubernetes methods:
  - Service
  - Route
  - Ingress
- VM-to-pod, and vice-versa, communication happens over SDN or ingress depending on network connectivity

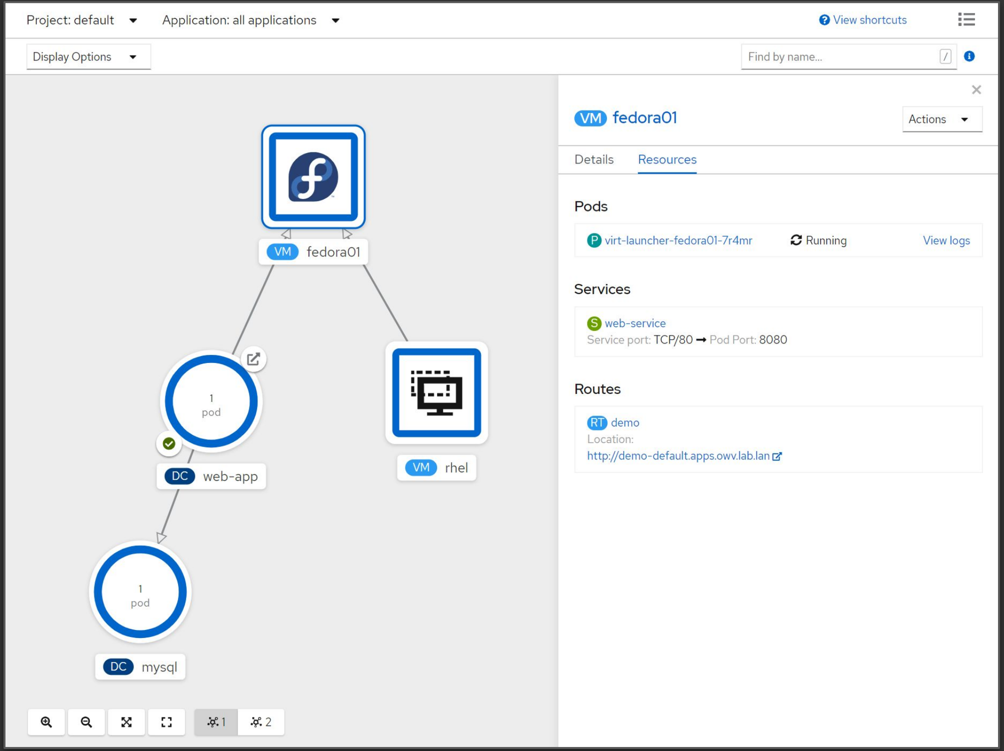

Project: default

Application: all applications

VM fedora01

Pods

P virt-launcher-fedora01-7r4mr Running View logs

Services

S web-service

Service port: TCP/80 ➞ Pod Port: 8080

Routes

RT demo

Location:

http://demo-default.apps.owv.lab.lan

CRC +
OKD
Virtualization

How can I try it @home? CRC + OKD Virtualization

CRC is the quickest way to get started building OpenShift (OCP/OKD) clusters. It is designed to run on a local computer to simplify setup and testing, and to emulate the cloud development environment locally with all of the tools needed to develop container-based applications.

It's not intended for production use!!!

- Single node cluster which behaves as both a control plane and worker node
- It's an ephemeral cluster inside a VM
- The OpenShift cluster runs in a virtual machine known as an instance on your laptop => you have to use nested virtualization to play with OKD Virtualization

How to: 1. Setup CRC

Enable nested virtualization on your laptop:

```
$ sudo modprobe -r kvm_intel
$ sudo modprobe kvm_intel nested=1
```

sed 's/intel/amd/g' on AMD machines

add 'options kvm_intel nested=1' to /etc/modprobe.d/kvm.conf to make it persistent

Download CRC binaries from https://github.com/crc-org/crc

Tune CRC configuration:

```
$ crc config set disk-size 64
$ crc config set memory 20480
$ crc config set enable-cluster-monitoring true
$ crc config set preset okd
```

Optionally enable the monitoring stack

Switch from OCP -> OKD

Configure prerequisites and download the bundle:

```
$ crc setup
```

Start your CRC instance:

```
$ crc start
```

CRC is already pre-configured to use kubevirt-hostpath-provisioner as a dynamic provisioner for PVs backed by CRC VM's filesystem.

No other manual actions are required on the CRC VM in order to be able to store and execute nested VMs inside the CRC VM.

After a few minutes...

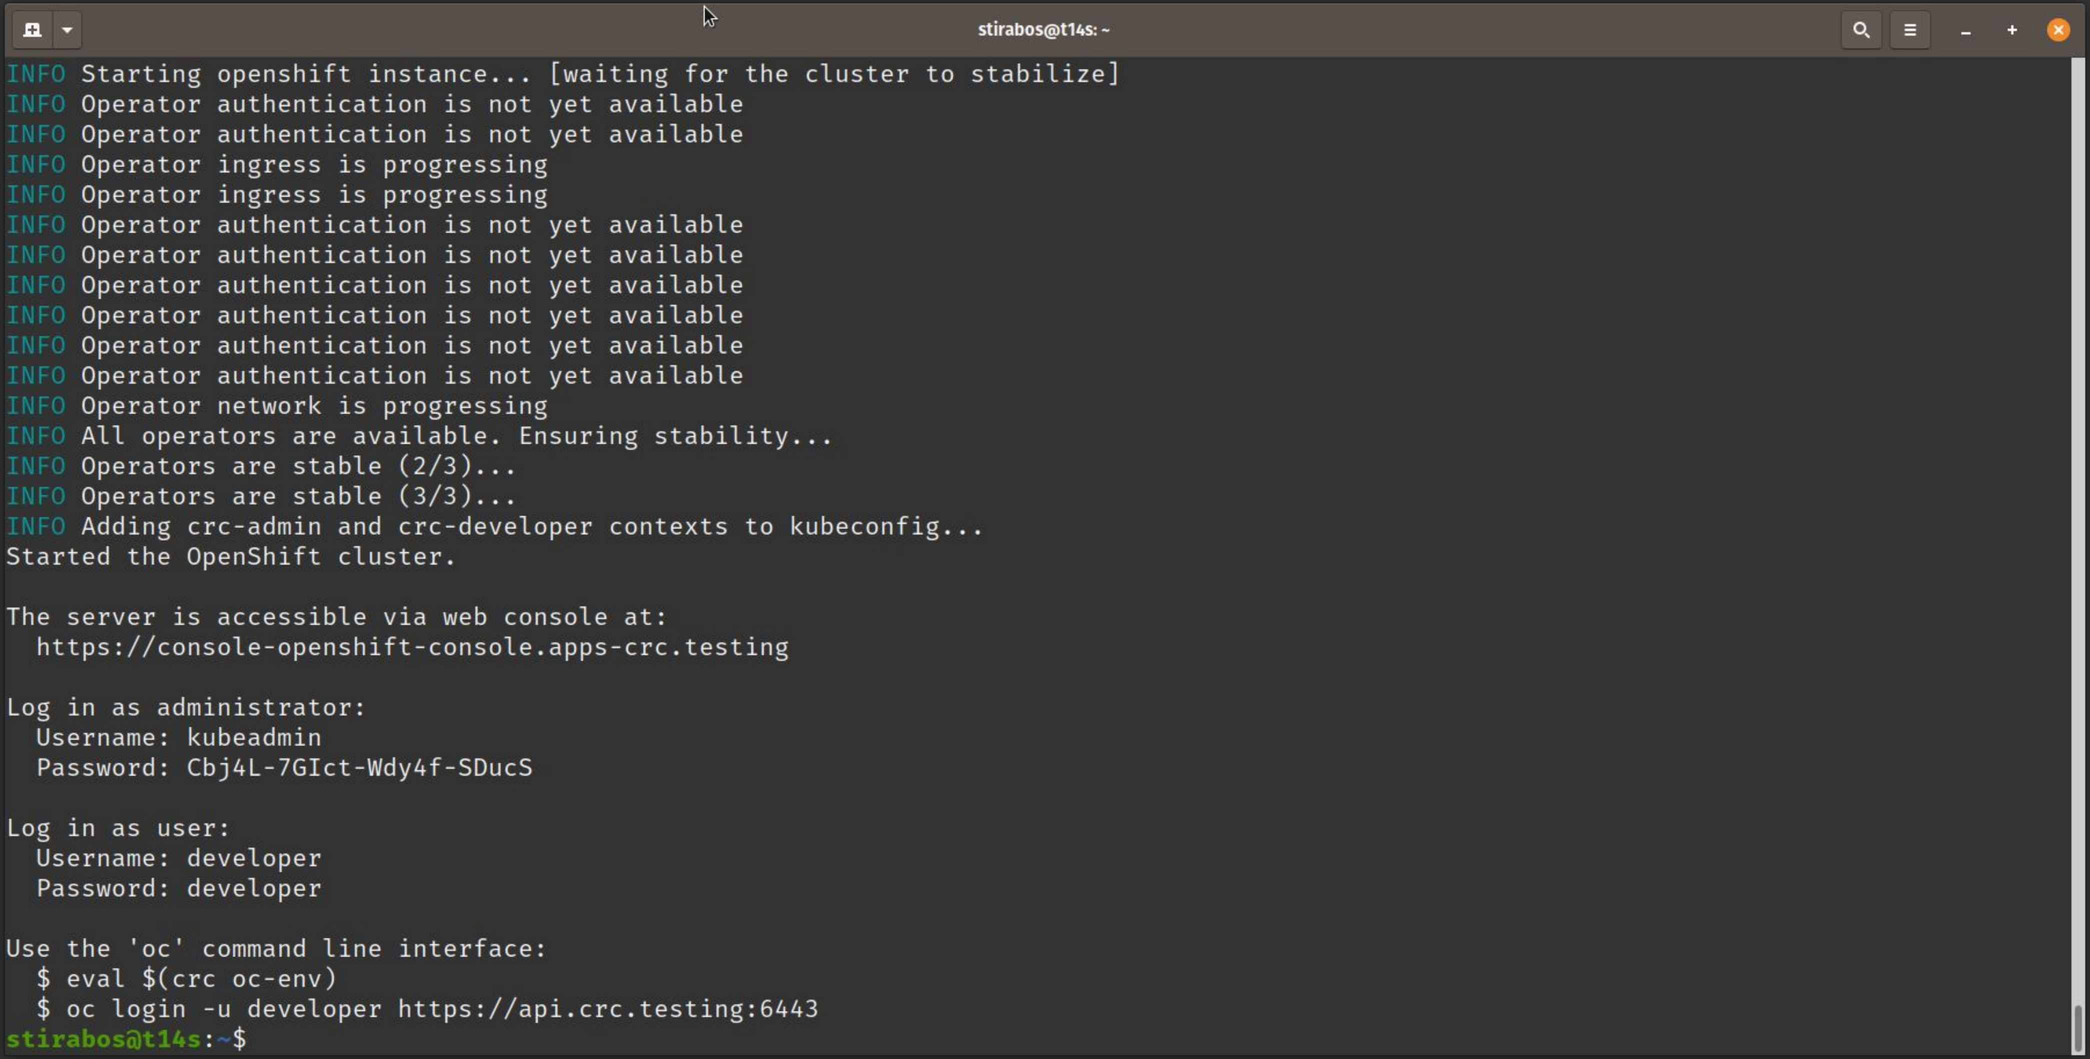

```
INFO Starting openshift instance... [waiting for the cluster to stabilize]
INFO Operator authentication is not yet available
INFO Operator authentication is not yet available
INFO Operator ingress is progressing
INFO Operator ingress is progressing
INFO Operator authentication is not yet available
INFO Operator authentication is not yet available
INFO Operator authentication is not yet available
INFO Operator authentication is not yet available
INFO Operator authentication is not yet available
INFO Operator authentication is not yet available
INFO Operator network is progressing
INFO All operators are available. Ensuring stability...
INFO Operators are stable (2/3)...
INFO Operators are stable (3/3)...
INFO Adding crc-admin and crc-developer contexts to kubeconfig...
Started the OpenShift cluster.

The server is accessible via web console at:
  https://console-openshift-console.apps-crc.testing

Log in as administrator:
  Username: kubeadmin
  Password: Cbj4L-7GIct-Wdy4f-SDucS

Log in as user:
  Username: developer
  Password: developer

Use the 'oc' command line interface:
  $ eval $(crc oc-env)
  $ oc login -u developer https://api.crc.testing:6443
stirabos@t14s:~$
```

Deploy **KubeVirt HyperConverged Cluster Operator** from the OperatorHub

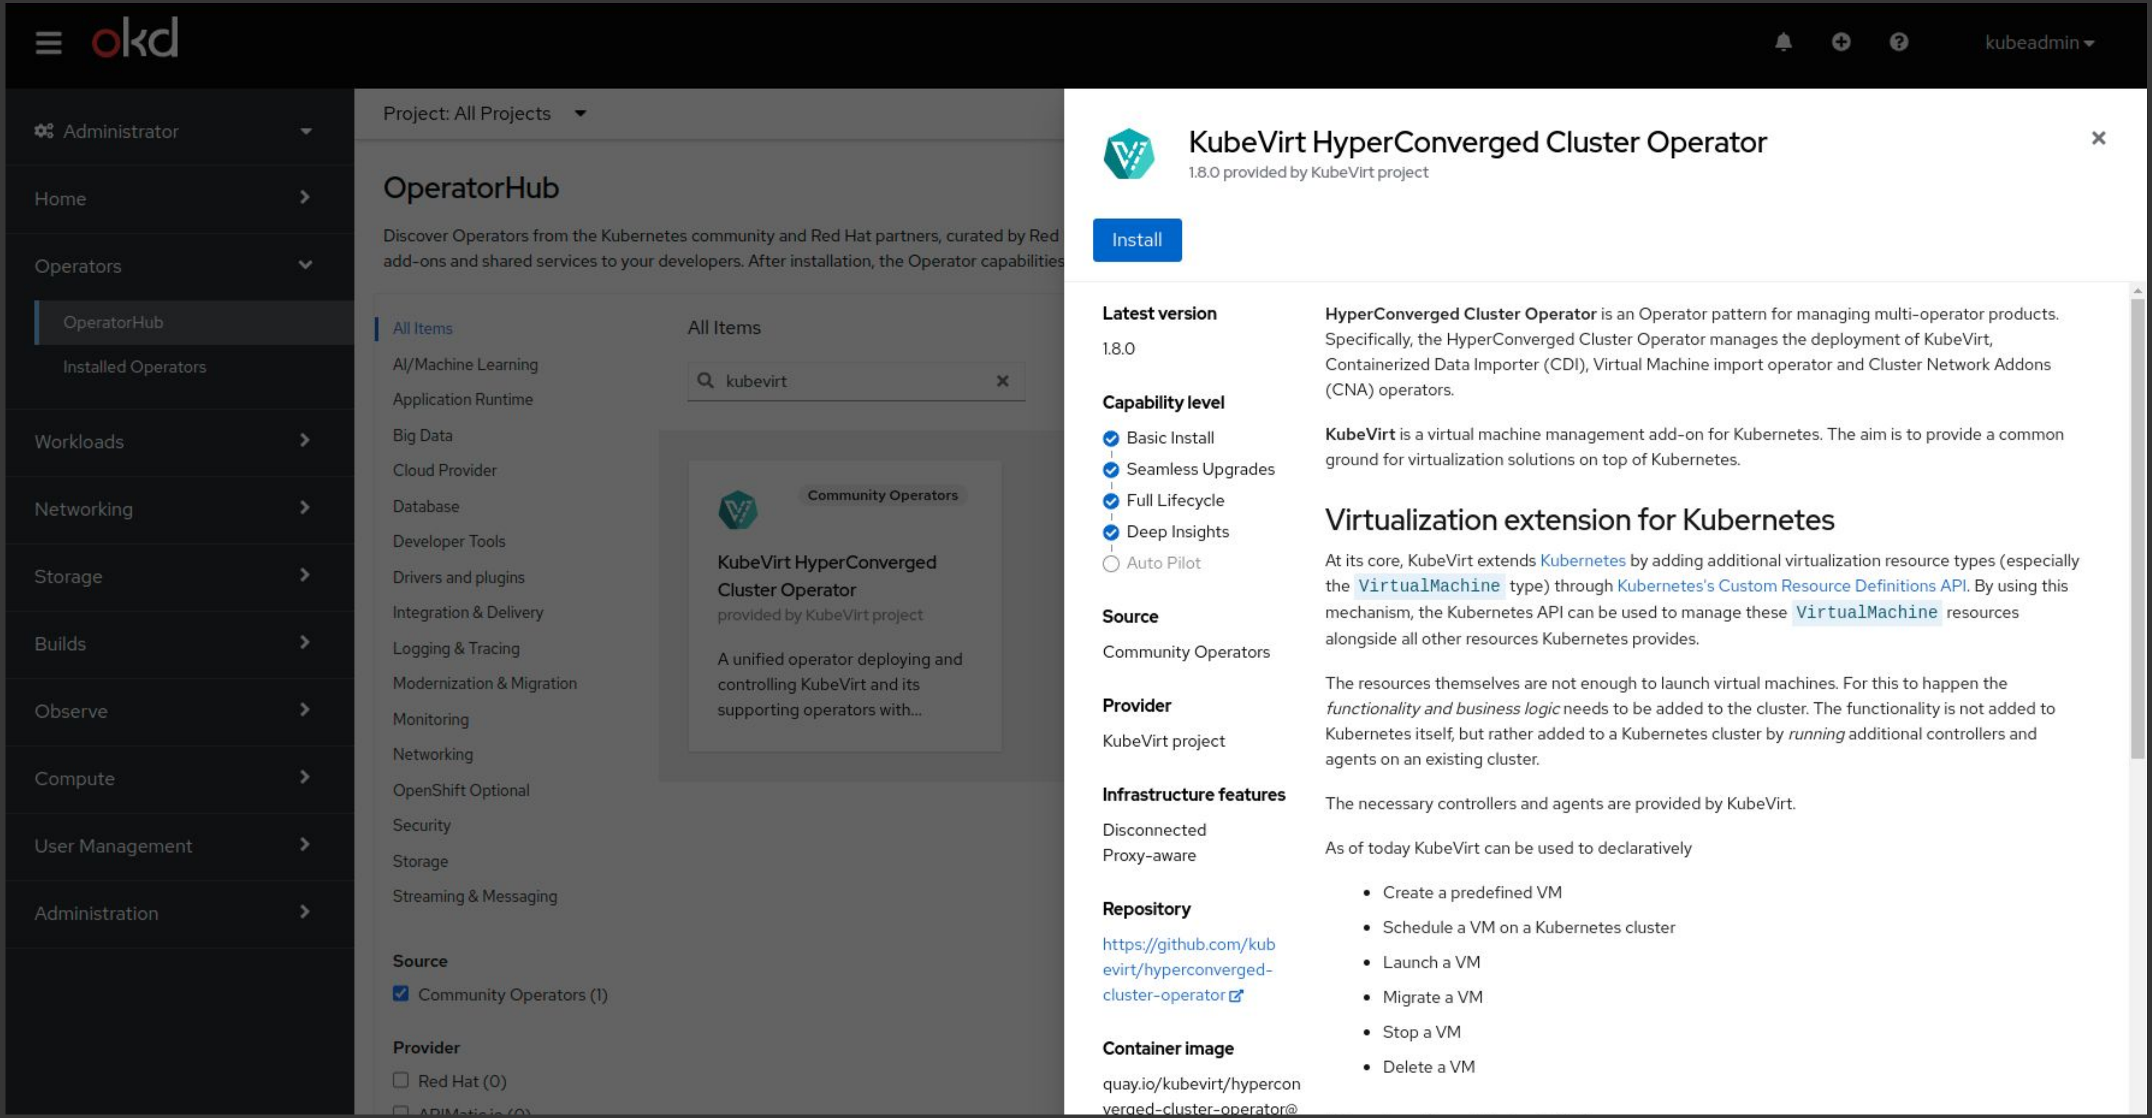

OperatorHub

Discover Operators from the Kubernetes community and Red Hat partners, curated by Red add-ons and shared services to your developers. After installation, the Operator capabilities

KubeVirt HyperConverged Cluster Operator
provided by KubeVirt project

A unified operator deploying and controlling KubeVirt and its supporting operators with...

KubeVirt HyperConverged Cluster Operator
1.8.0 provided by KubeVirt project

Latest version
1.8.0

Capability level
- ✓ Basic Install
- ✓ Seamless Upgrades
- ✓ Full Lifecycle
- ✓ Deep Insights
- ○ Auto Pilot

Source
Community Operators

Provider
KubeVirt project

Infrastructure features
Disconnected
Proxy-aware

Repository
https://github.com/kubevirt/hyperconverged-cluster-operator

Container image
quay.io/kubevirt/hyperconverged-cluster-operator@

HyperConverged Cluster Operator is an Operator pattern for managing multi-operator products. Specifically, the HyperConverged Cluster Operator manages the deployment of KubeVirt, Containerized Data Importer (CDI), Virtual Machine import operator and Cluster Network Addons (CNA) operators.

KubeVirt is a virtual machine management add-on for Kubernetes. The aim is to provide a common ground for virtualization solutions on top of Kubernetes.

Virtualization extension for Kubernetes

At its core, KubeVirt extends Kubernetes by adding additional virtualization resource types (especially the `VirtualMachine` type) through Kubernetes's Custom Resource Definitions API. By using this mechanism, the Kubernetes API can be used to manage these `VirtualMachine` resources alongside all other resources Kubernetes provides.

The resources themselves are not enough to launch virtual machines. For this to happen the *functionality and business logic* needs to be added to the cluster. The functionality is not added to Kubernetes itself, but rather added to a Kubernetes cluster by *running* additional controllers and agents on an existing cluster.

The necessary controllers and agents are provided by KubeVirt.

As of today KubeVirt can be used to declaratively

- Create a predefined VM
- Schedule a VM on a Kubernetes cluster
- Launch a VM
- Migrate a VM
- Stop a VM
- Delete a VM

Deploy KubeVirt HyperConverged Cluster Operator

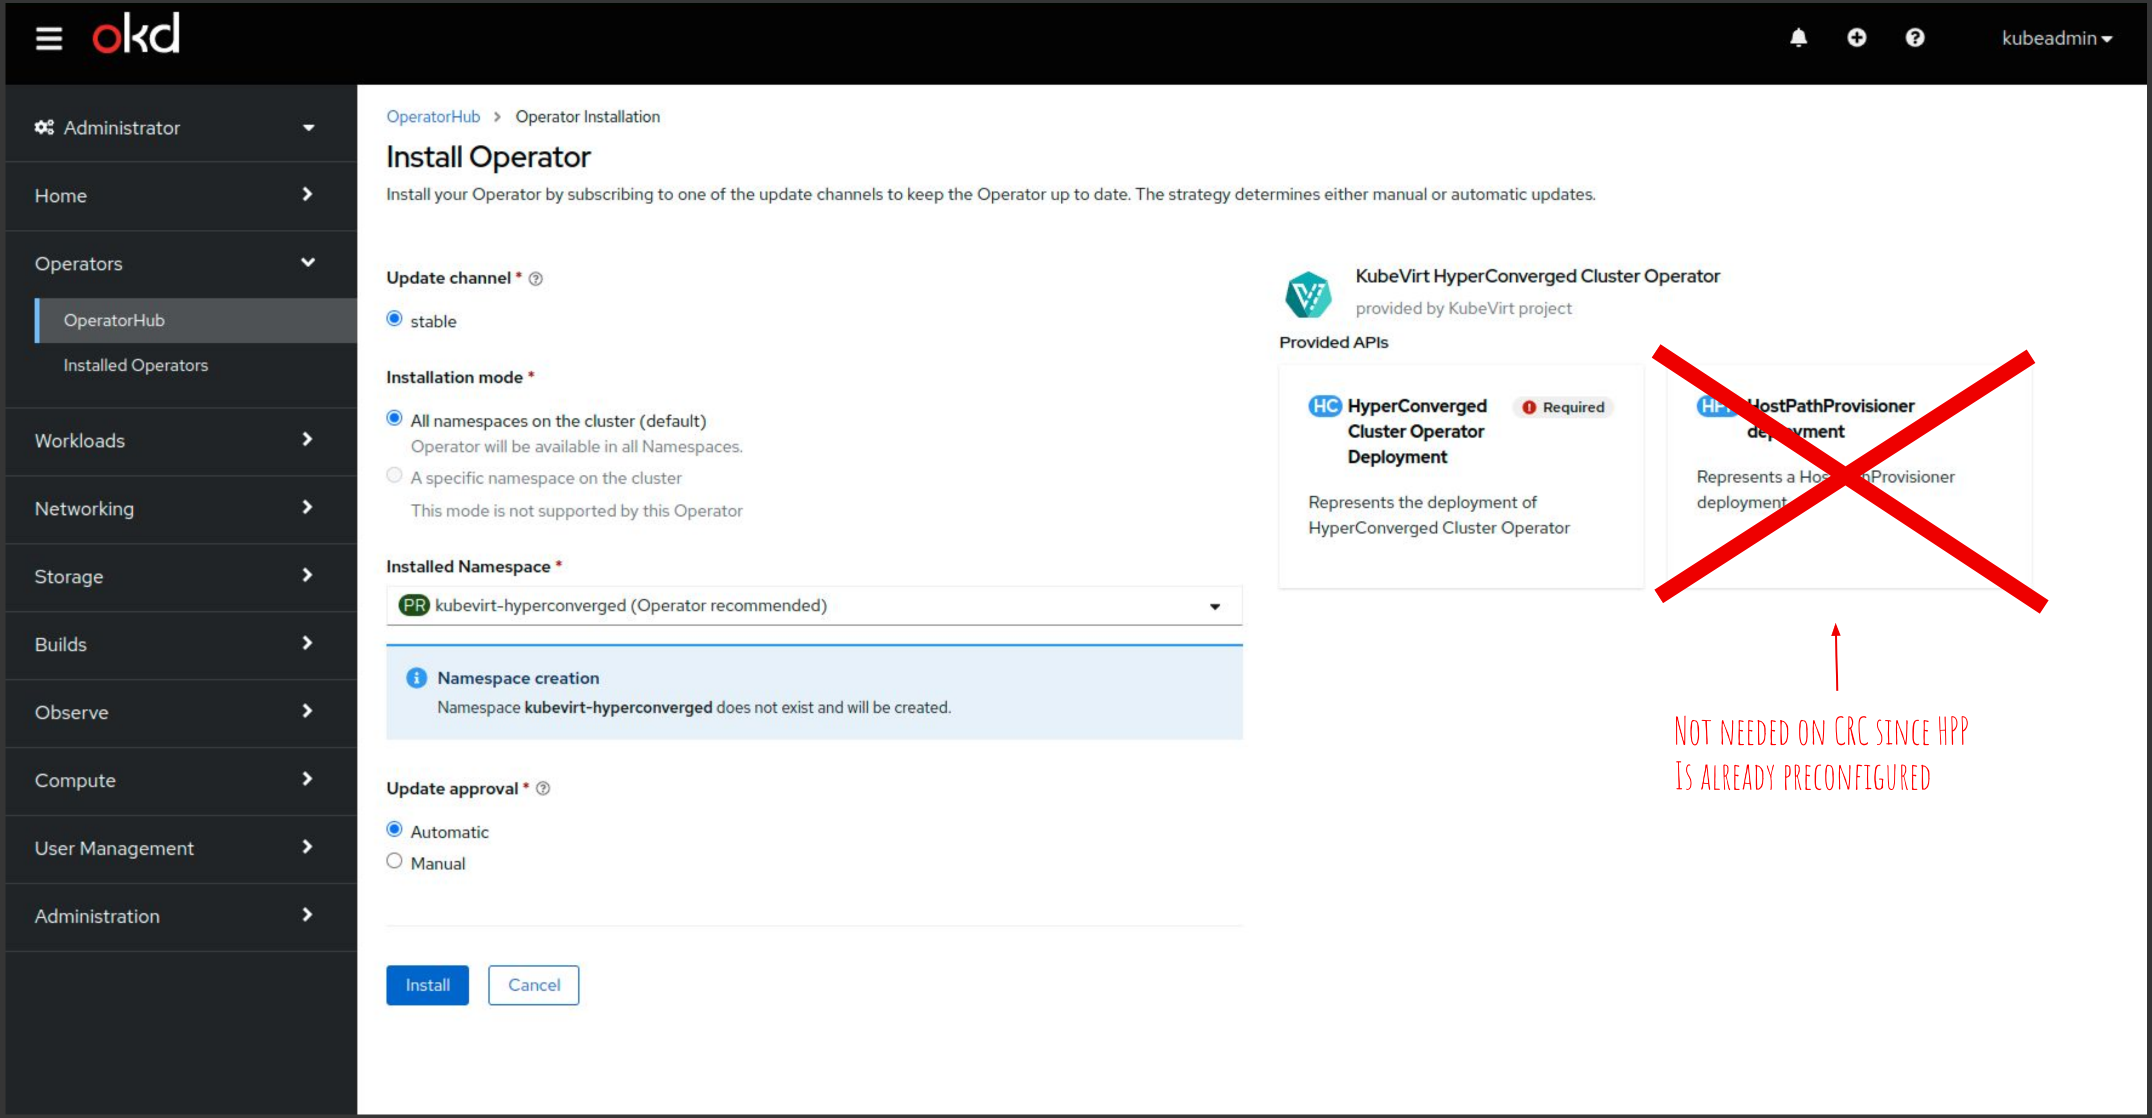

OperatorHub > Operator Installation

Install Operator

Install your Operator by subscribing to one of the update channels to keep the Operator up to date. The strategy determines either manual or automatic updates.

Update channel *

● stable

Installation mode *

● All namespaces on the cluster (default)
Operator will be available in all Namespaces.

○ A specific namespace on the cluster
This mode is not supported by this Operator

Installed Namespace *

PR kubevirt-hyperconverged (Operator recommended)

Namespace creation
Namespace **kubevirt-hyperconverged** does not exist and will be created.

Update approval *

● Automatic

○ Manual

Install | Cancel

KubeVirt HyperConverged Cluster Operator
provided by KubeVirt project

Provided APIs

HC **HyperConverged Cluster Operator Deployment** Required
Represents the deployment of HyperConverged Cluster Operator

~~HP HostPathProvisioner deployment~~
~~Represents a HostPathProvisioner deployment~~

NOT NEEDED ON CRC SINCE HPP
IS ALREADY PRECONFIGURED

HCO got deployed

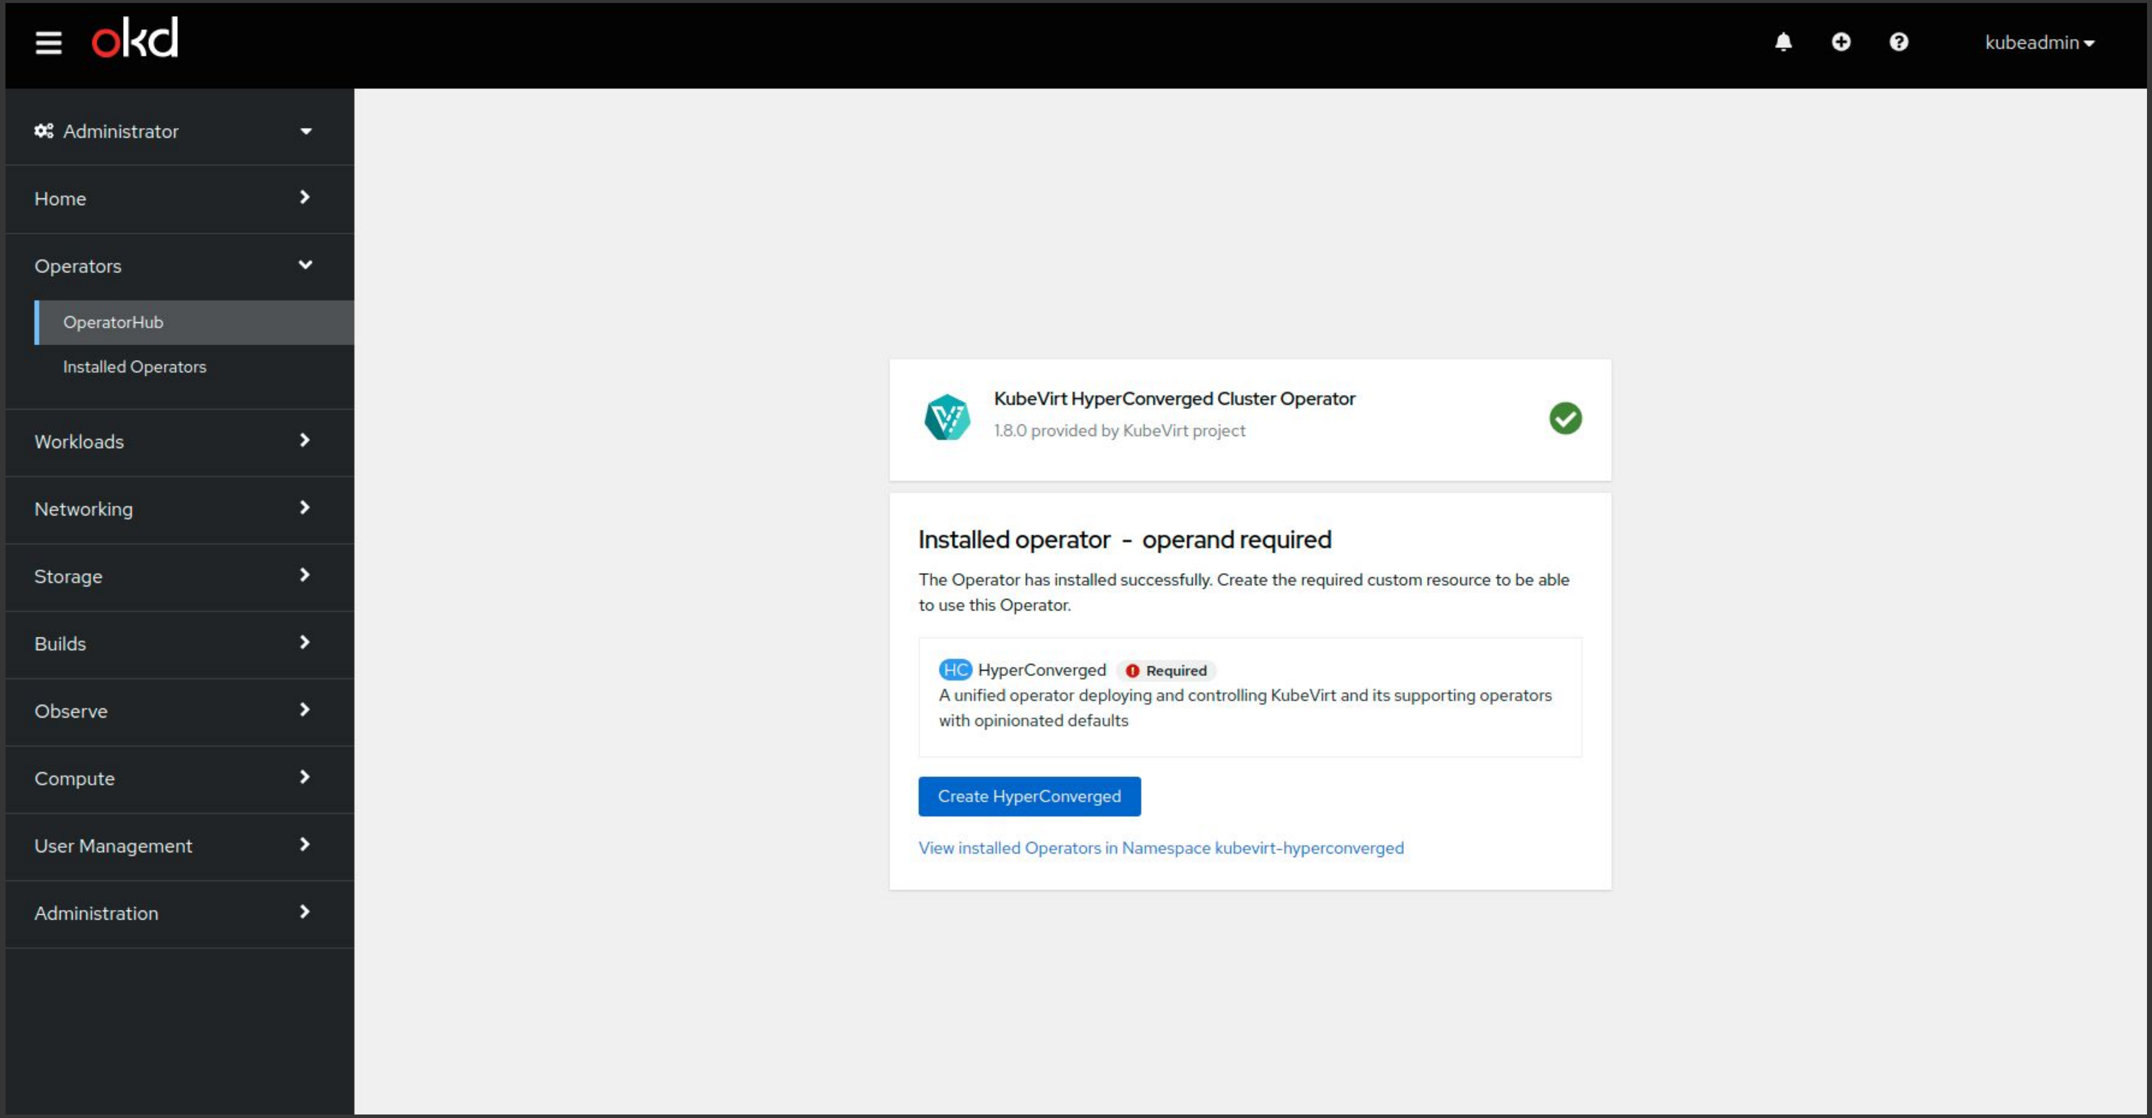

KubeVirt HyperConverged Cluster Operator

1.8.0 provided by KubeVirt project

Installed operator - operand required

The Operator has installed successfully. Create the required custom resource to be able to use this Operator.

HC HyperConverged Required

A unified operator deploying and controlling KubeVirt and its supporting operators with opinionated defaults

Create HyperConverged

View installed Operators in Namespace kubevirt-hyperconverged

Create a CR to trigger HCO

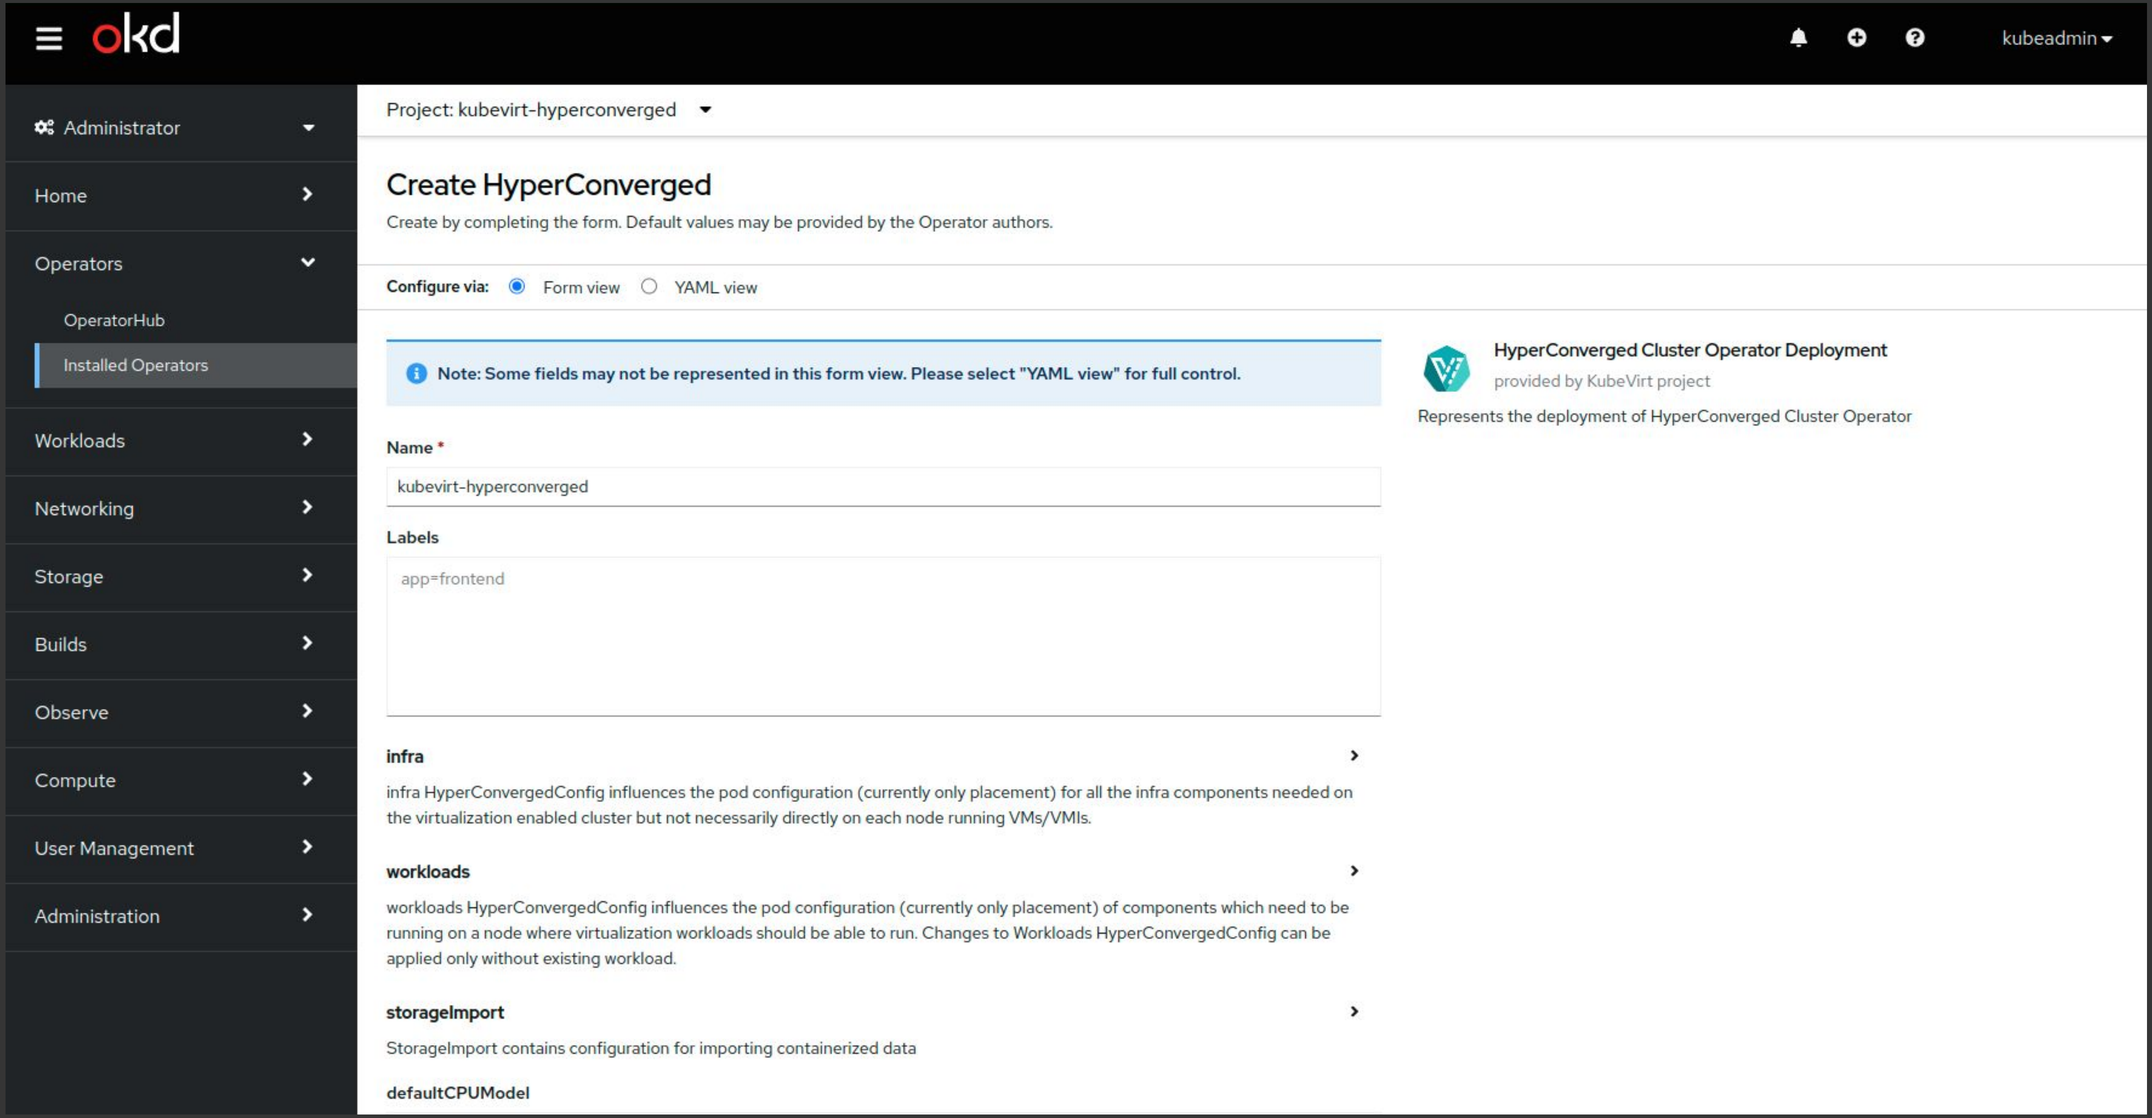

Project: kubevirt-hyperconverged

Create HyperConverged

Create by completing the form. Default values may be provided by the Operator authors.

Configure via: ● Form view ○ YAML view

Note: Some fields may not be represented in this form view. Please select "YAML view" for full control.

Name *

kubevirt-hyperconverged

Labels

app=frontend

infra

infra HyperConvergedConfig influences the pod configuration (currently only placement) for all the infra components needed on the virtualization enabled cluster but not necessarily directly on each node running VMs/VMIs.

workloads

workloads HyperConvergedConfig influences the pod configuration (currently only placement) of components which need to be running on a node where virtualization workloads should be able to run. Changes to Workloads HyperConvergedConfig can be applied only without existing workload.

storageImport

StorageImport contains configuration for importing containerized data

defaultCPUModel

HyperConverged Cluster Operator Deployment

provided by KubeVirt project

Represents the deployment of HyperConverged Cluster Operator

Deployment options

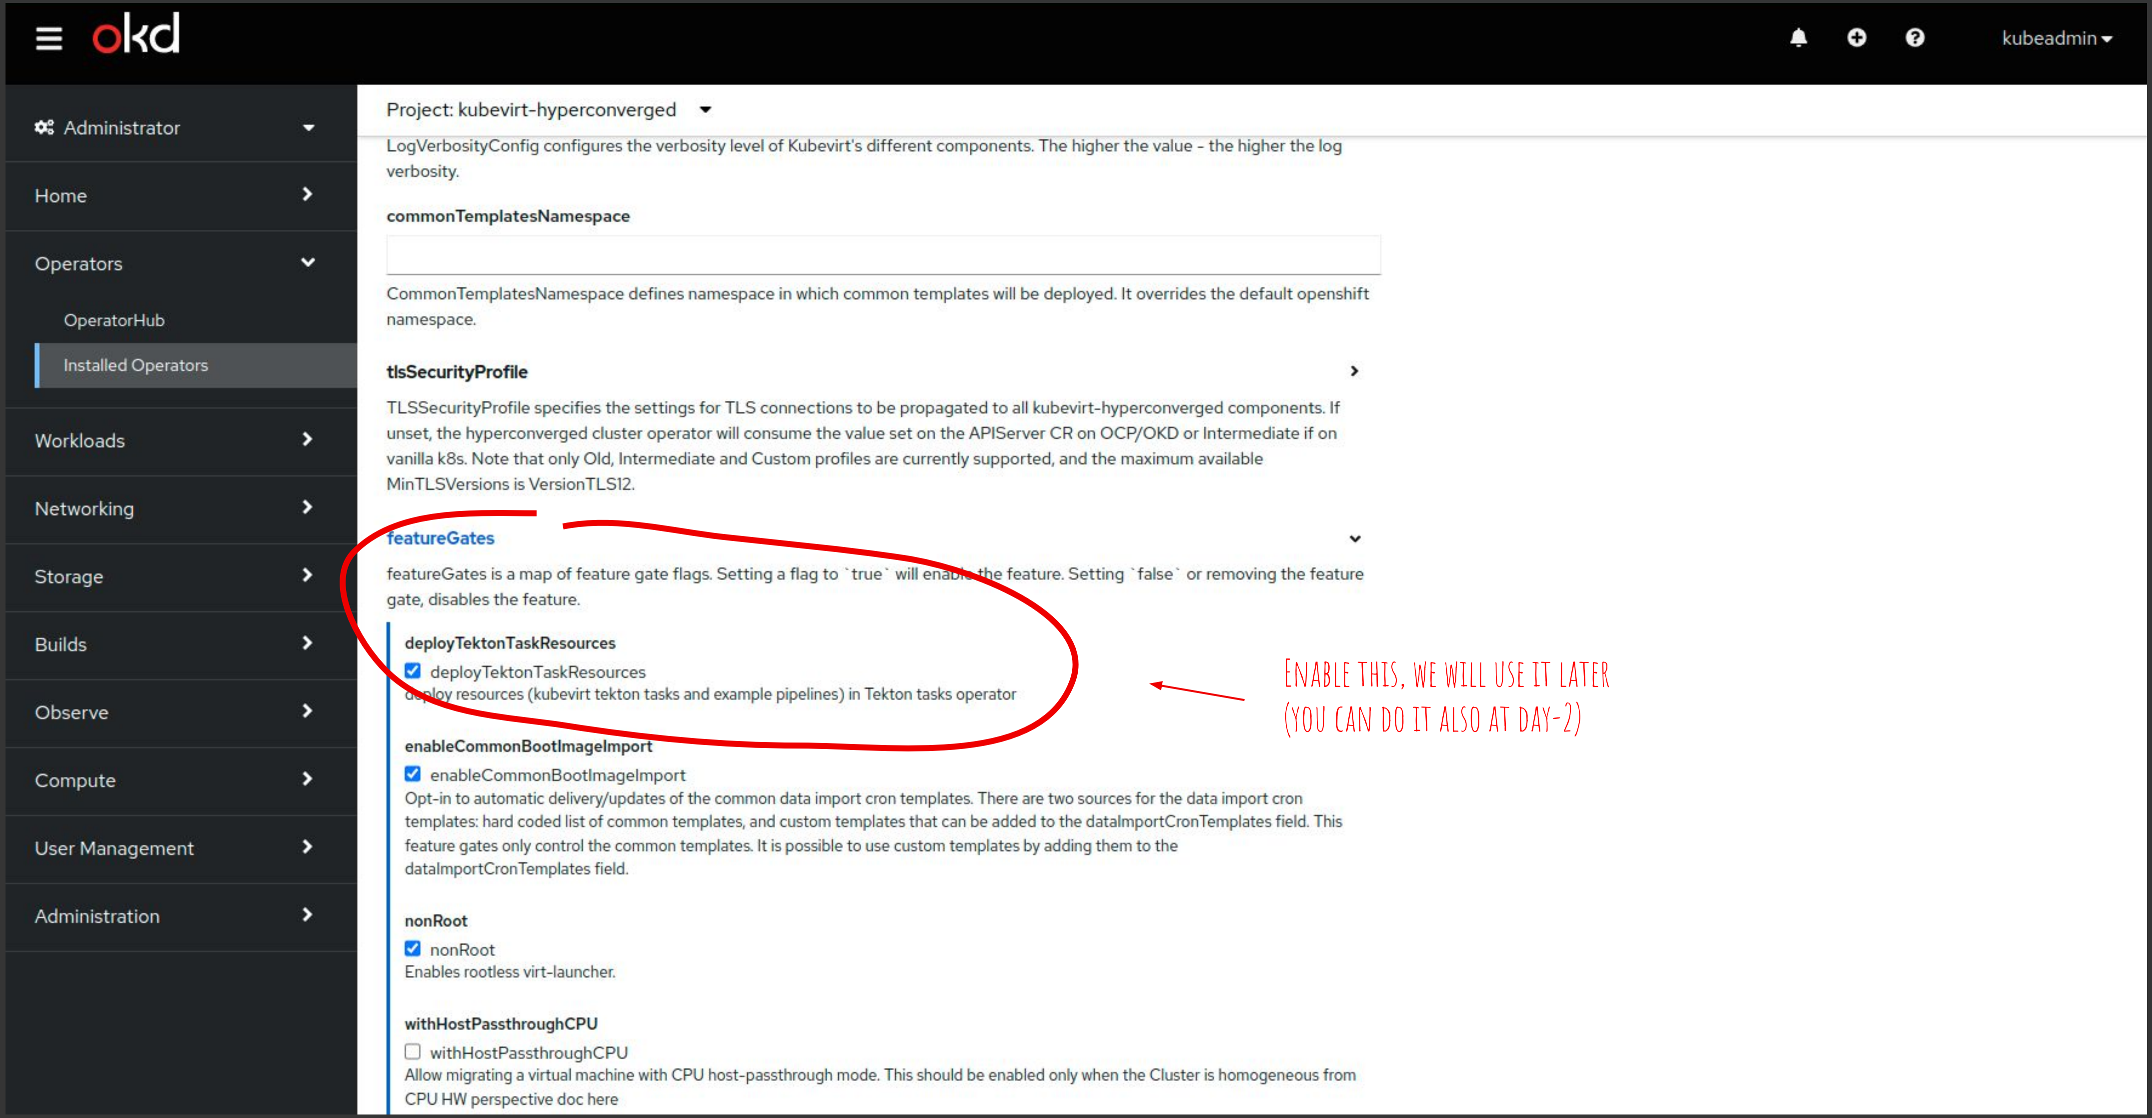

Project: kubevirt-hyperconverged

LogVerbosityConfig configures the verbosity level of Kubevirt's different components. The higher the value - the higher the log verbosity.

commonTemplatesNamespace

CommonTemplatesNamespace defines namespace in which common templates will be deployed. It overrides the default openshift namespace.

tlsSecurityProfile

TLSSecurityProfile specifies the settings for TLS connections to be propagated to all kubevirt-hyperconverged components. If unset, the hyperconverged cluster operator will consume the value set on the APIServer CR on OCP/OKD or Intermediate if on vanilla k8s. Note that only Old, Intermediate and Custom profiles are currently supported, and the maximum available MinTLSVersions is VersionTLS12.

featureGates

featureGates is a map of feature gate flags. Setting a flag to `true` will enable the feature. Setting `false` or removing the feature gate, disables the feature.

deployTektonTaskResources

☑ deployTektonTaskResources

deploy resources (kubevirt tekton tasks and example pipelines) in Tekton tasks operator

enableCommonBootImageImport

☑ enableCommonBootImageImport

Opt-in to automatic delivery/updates of the common data import cron templates. There are two sources for the data import cron templates: hard coded list of common templates, and custom templates that can be added to the dataImportCronTemplates field. This feature gates only control the common templates. It is possible to use custom templates by adding them to the dataImportCronTemplates field.

nonRoot

☑ nonRoot

Enables rootless virt-launcher.

withHostPassthroughCPU

☐ withHostPassthroughCPU

Allow migrating a virtual machine with CPU host-passthrough mode. This should be enabled only when the Cluster is homogeneous from CPU HW perspective doc here

ENABLE THIS, WE WILL USE IT LATER
(YOU CAN DO IT ALSO AT DAY-2)

After a few minutes...

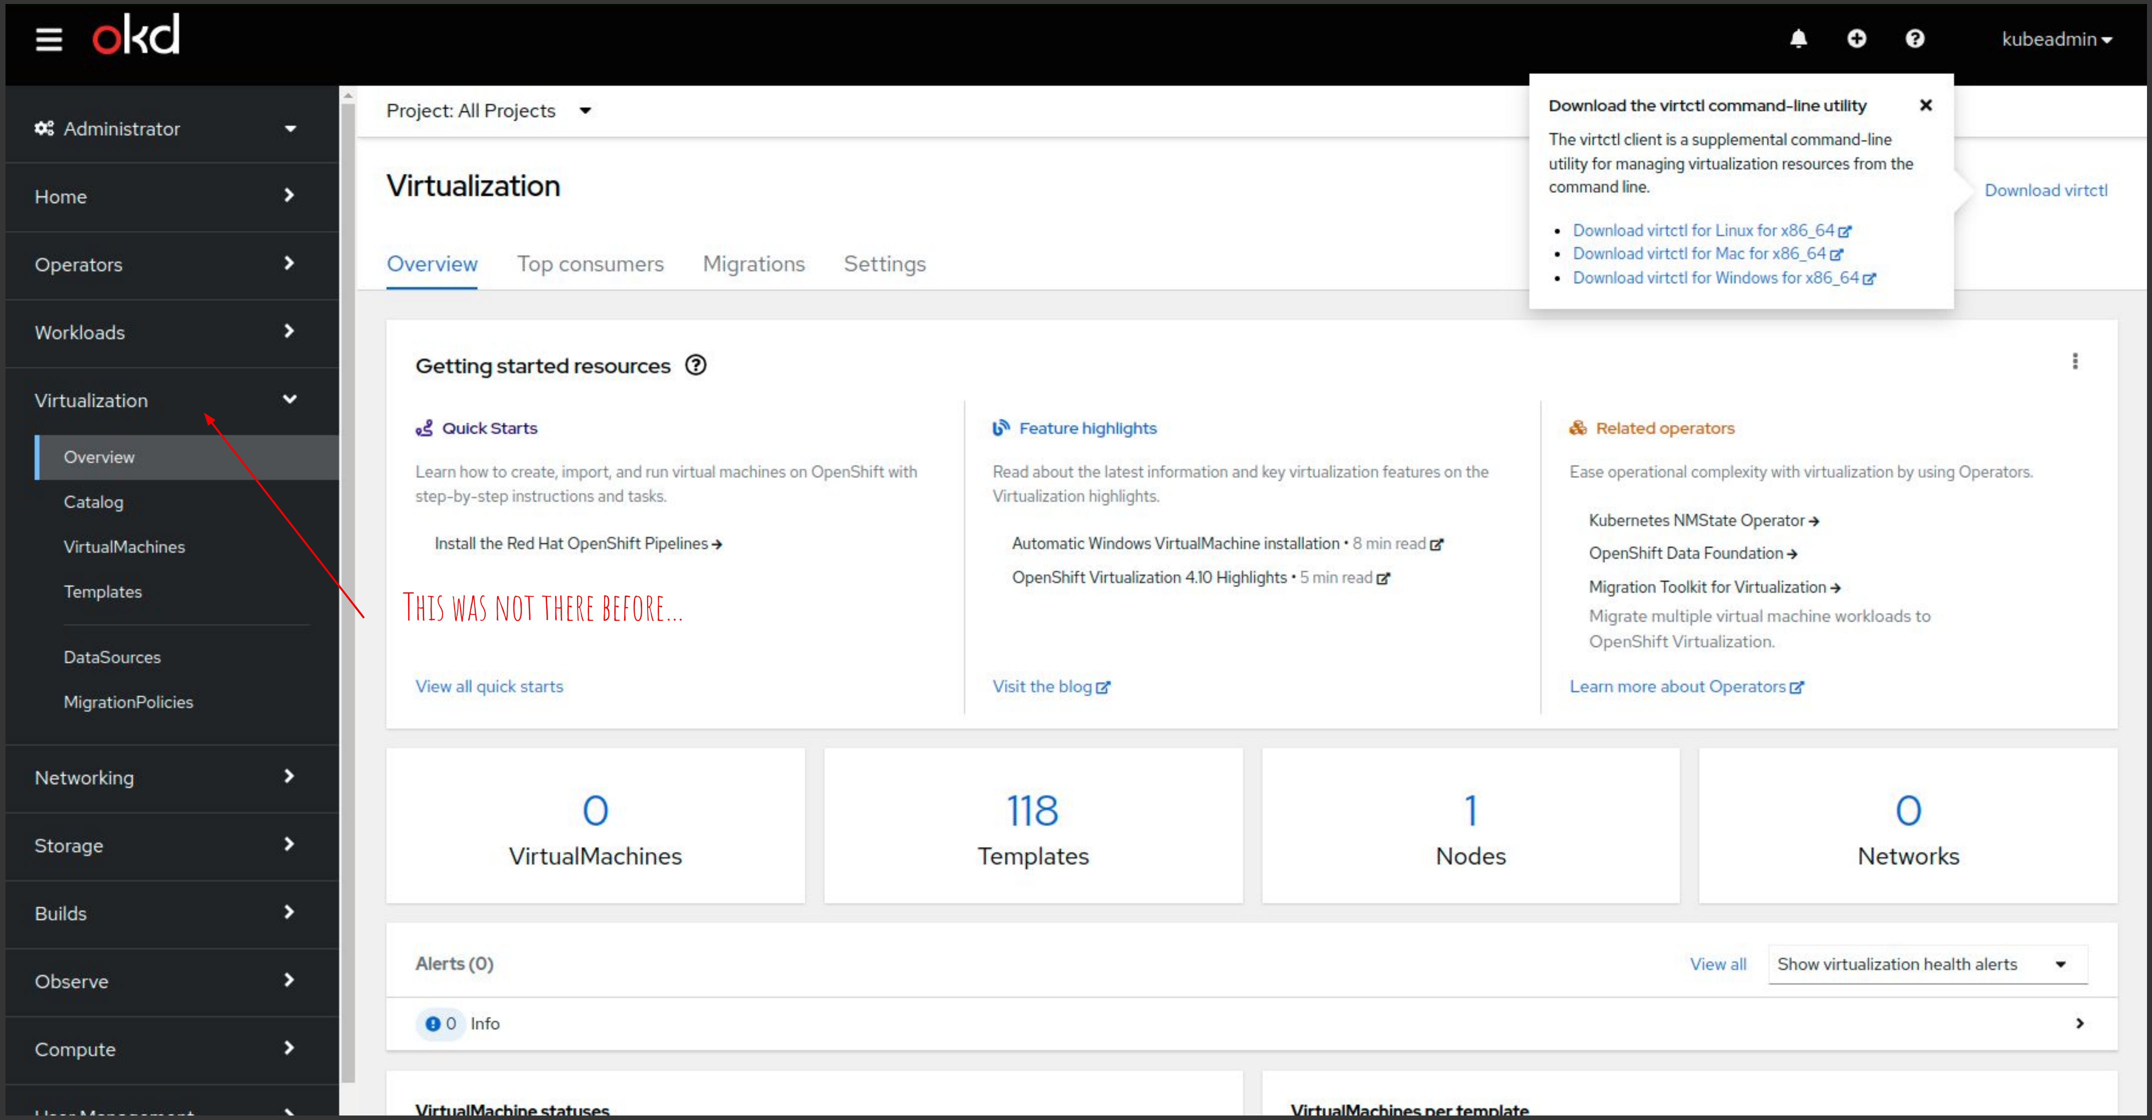

Virtualization

Download the virtctl command-line utility

The virtctl client is a supplemental command-line utility for managing virtualization resources from the command line.

- Download virtctl for Linux for x86_64
- Download virtctl for Mac for x86_64
- Download virtctl for Windows for x86_64

Getting started resources

Quick Starts

Learn how to create, import, and run virtual machines on OpenShift with step-by-step instructions and tasks.

Install the Red Hat OpenShift Pipelines →

View all quick starts

Feature highlights

Read about the latest information and key virtualization features on the Virtualization highlights.

Automatic Windows VirtualMachine installation • 8 min read

OpenShift Virtualization 4.10 Highlights • 5 min read

Visit the blog

Related operators

Ease operational complexity with virtualization by using Operators.

Kubernetes NMState Operator →

OpenShift Data Foundation →

Migration Toolkit for Virtualization →

Migrate multiple virtual machine workloads to OpenShift Virtualization.

Learn more about Operators

0	118	1	0
VirtualMachines	Templates	Nodes	Networks

Alerts (0) View all Show virtualization health alerts

0 Info

VirtualMachine statuses

VirtualMachines per template

THIS WAS NOT THERE BEFORE...

Golden images

Golden Images: motivation

Hyperscalers (like AWS, GCP, Azure, IBM Cloud…) provide:

- root disk images for commonly used operating systems
- continuous updates of those images

Starting a fresh VM is really really quick and simple!

the KubeVirt ecosystem should provide the tools necessary for supporting a similar pattern of disk image availability within the cluster

Golden Images: design

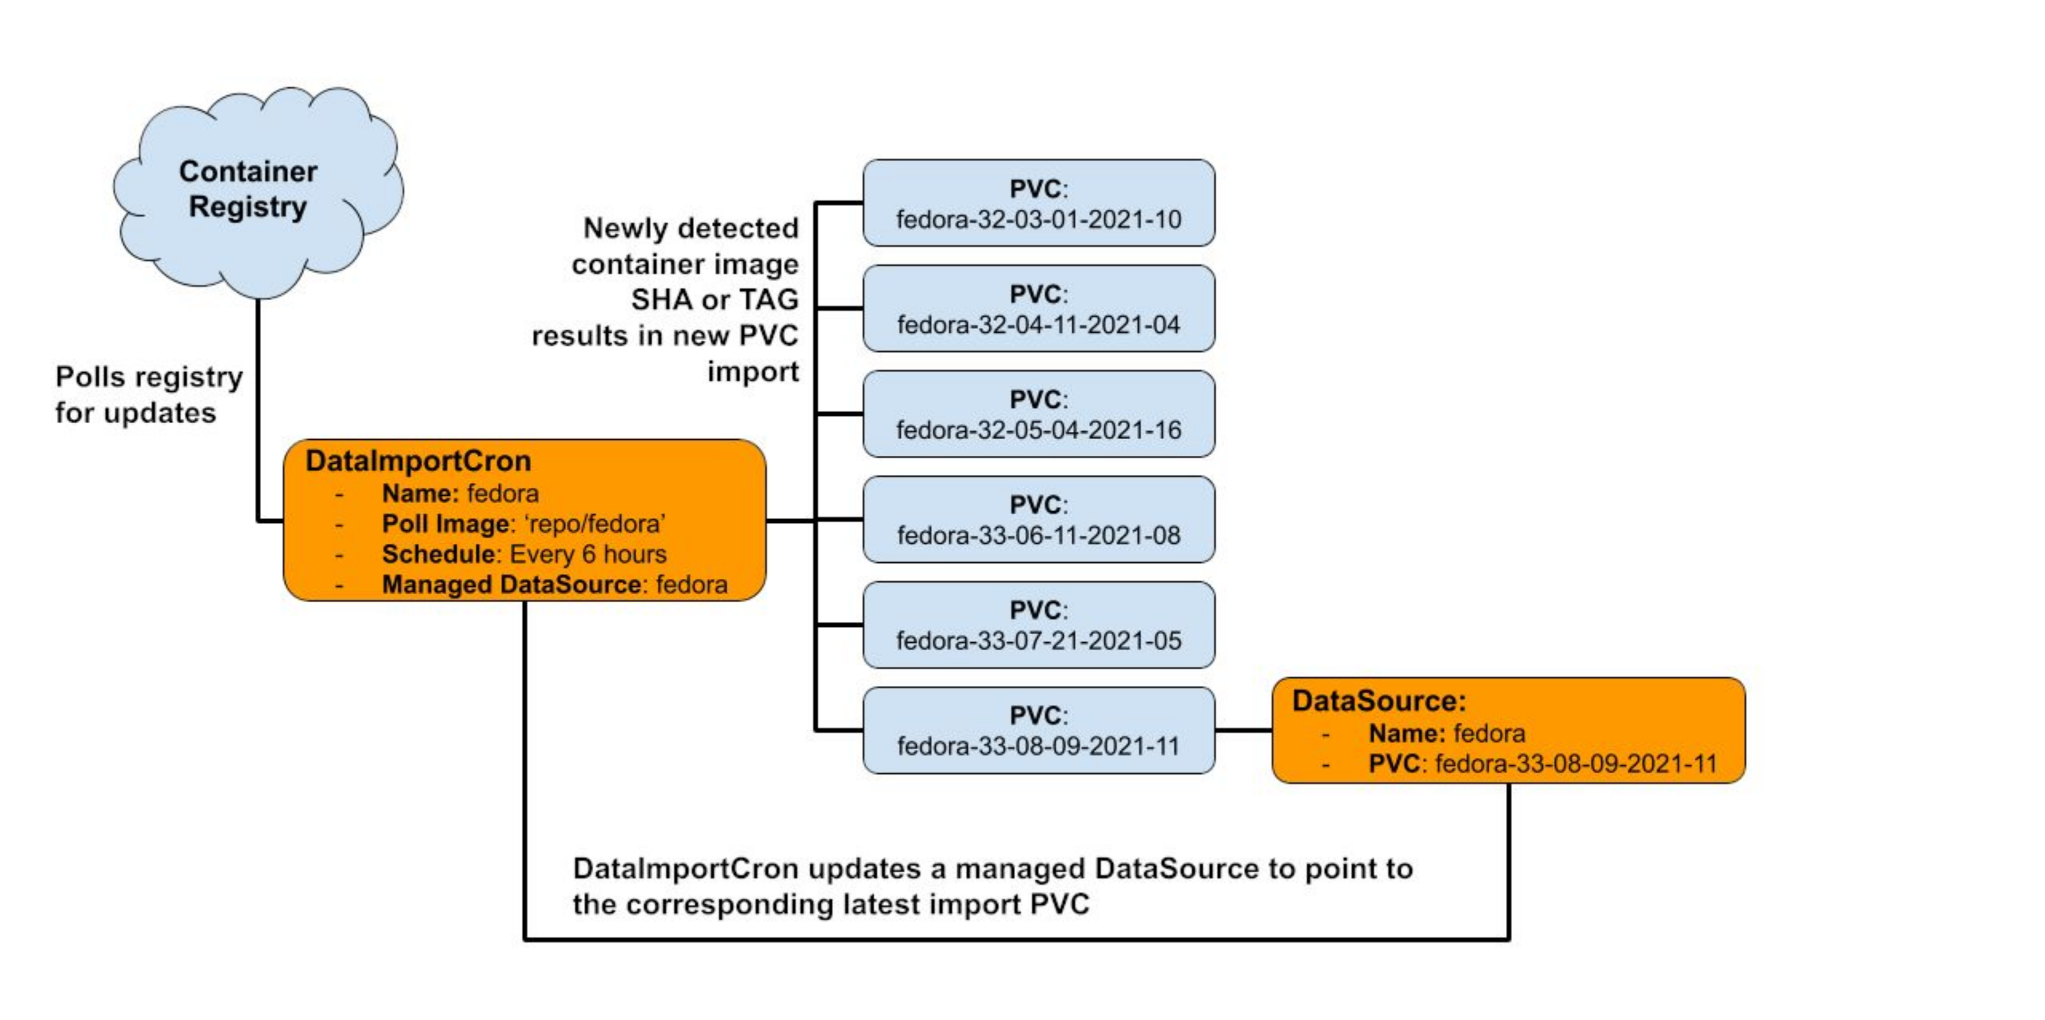

Container Registry
Polls registry for updates
Newly detected container image SHA or TAG results in new PVC import
DataImportCron
- Name: fedora
- Poll Image: 'repo/fedora'
- Schedule: Every 6 hours
- Managed DataSource: fedora
PVC: fedora-32-03-01-2021-10
PVC: fedora-32-04-11-2021-04
PVC: fedora-32-05-04-2021-16
PVC: fedora-33-06-11-2021-08
PVC: fedora-33-07-21-2021-05
PVC: fedora-33-08-09-2021-11
DataSource:
- Name: fedora
- PVC: fedora-33-08-09-2021-11
DataImportCron updates a managed DataSource to point to the corresponding latest import PVC

How they look like for cluster admins

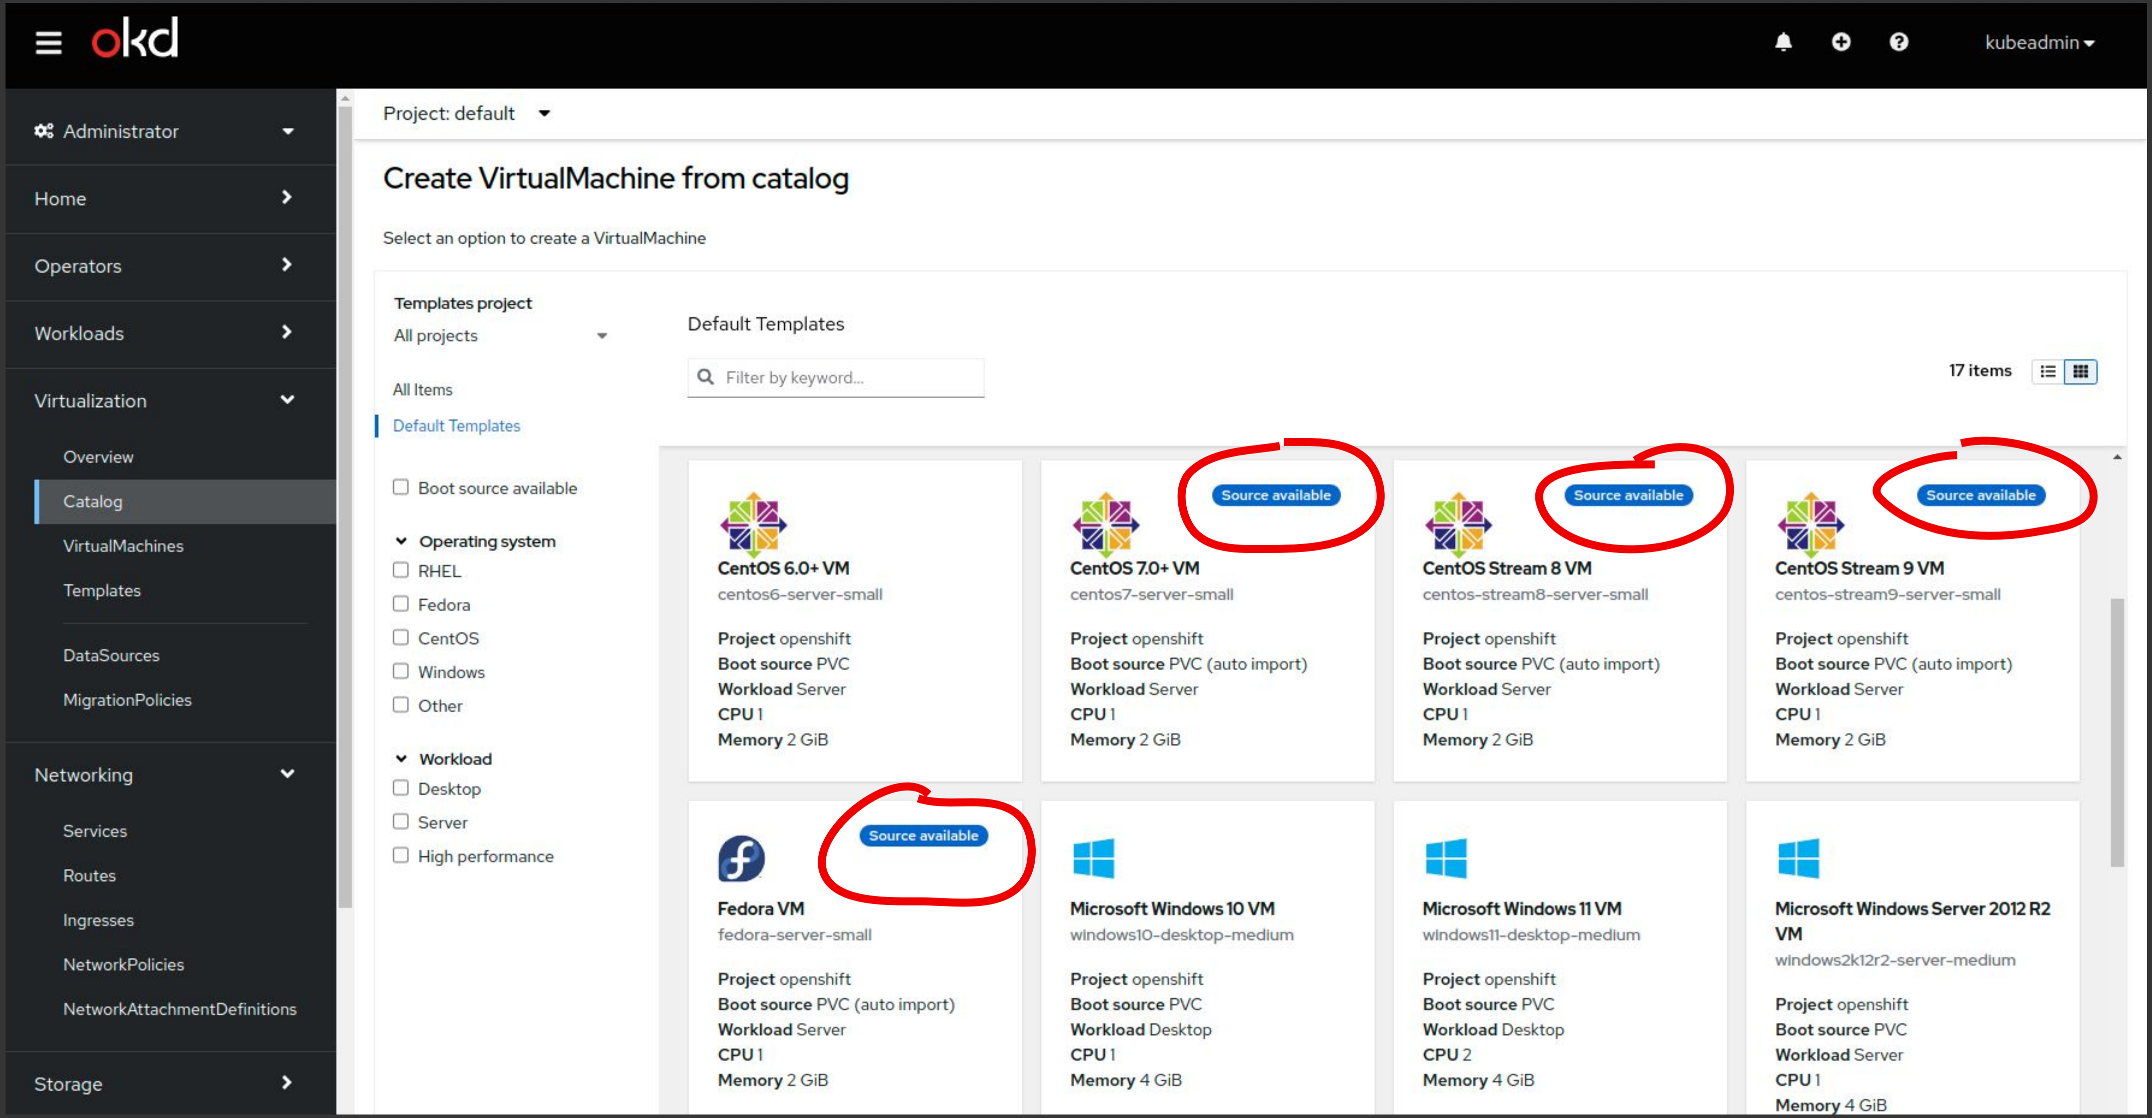

Project: default

Create VirtualMachine from catalog

Select an option to create a VirtualMachine

Templates project
All projects

All Items
Default Templates

□ Boot source available

Operating system
□ RHEL
□ Fedora
□ CentOS
□ Windows
□ Other

Workload
□ Desktop
□ Server
□ High performance

Default Templates

Filter by keyword...

17 items

CentOS 6.0+ VM
centos6-server-small
Project openshift
Boot source PVC
Workload Server
CPU 1
Memory 2 GiB

Source available
CentOS 7.0+ VM
centos7-server-small
Project openshift
Boot source PVC (auto import)
Workload Server
CPU 1
Memory 2 GiB

Source available
CentOS Stream 8 VM
centos-stream8-server-small
Project openshift
Boot source PVC (auto import)
Workload Server
CPU 1
Memory 2 GiB

Source available
CentOS Stream 9 VM
centos-stream9-server-small
Project openshift
Boot source PVC (auto import)
Workload Server
CPU 1
Memory 2 GiB

Source available
Fedora VM
fedora-server-small
Project openshift
Boot source PVC (auto import)
Workload Server
CPU 1
Memory 2 GiB

Microsoft Windows 10 VM
windows10-desktop-medium
Project openshift
Boot source PVC
Workload Desktop
CPU 1
Memory 4 GiB

Microsoft Windows 11 VM
windows11-desktop-medium
Project openshift
Boot source PVC
Workload Desktop
CPU 2
Memory 4 GiB

Microsoft Windows Server 2012 R2 VM
windows2k12r2-server-medium
Project openshift
Boot source PVC
Workload Server
CPU 1
Memory 4 GiB

How can I use them

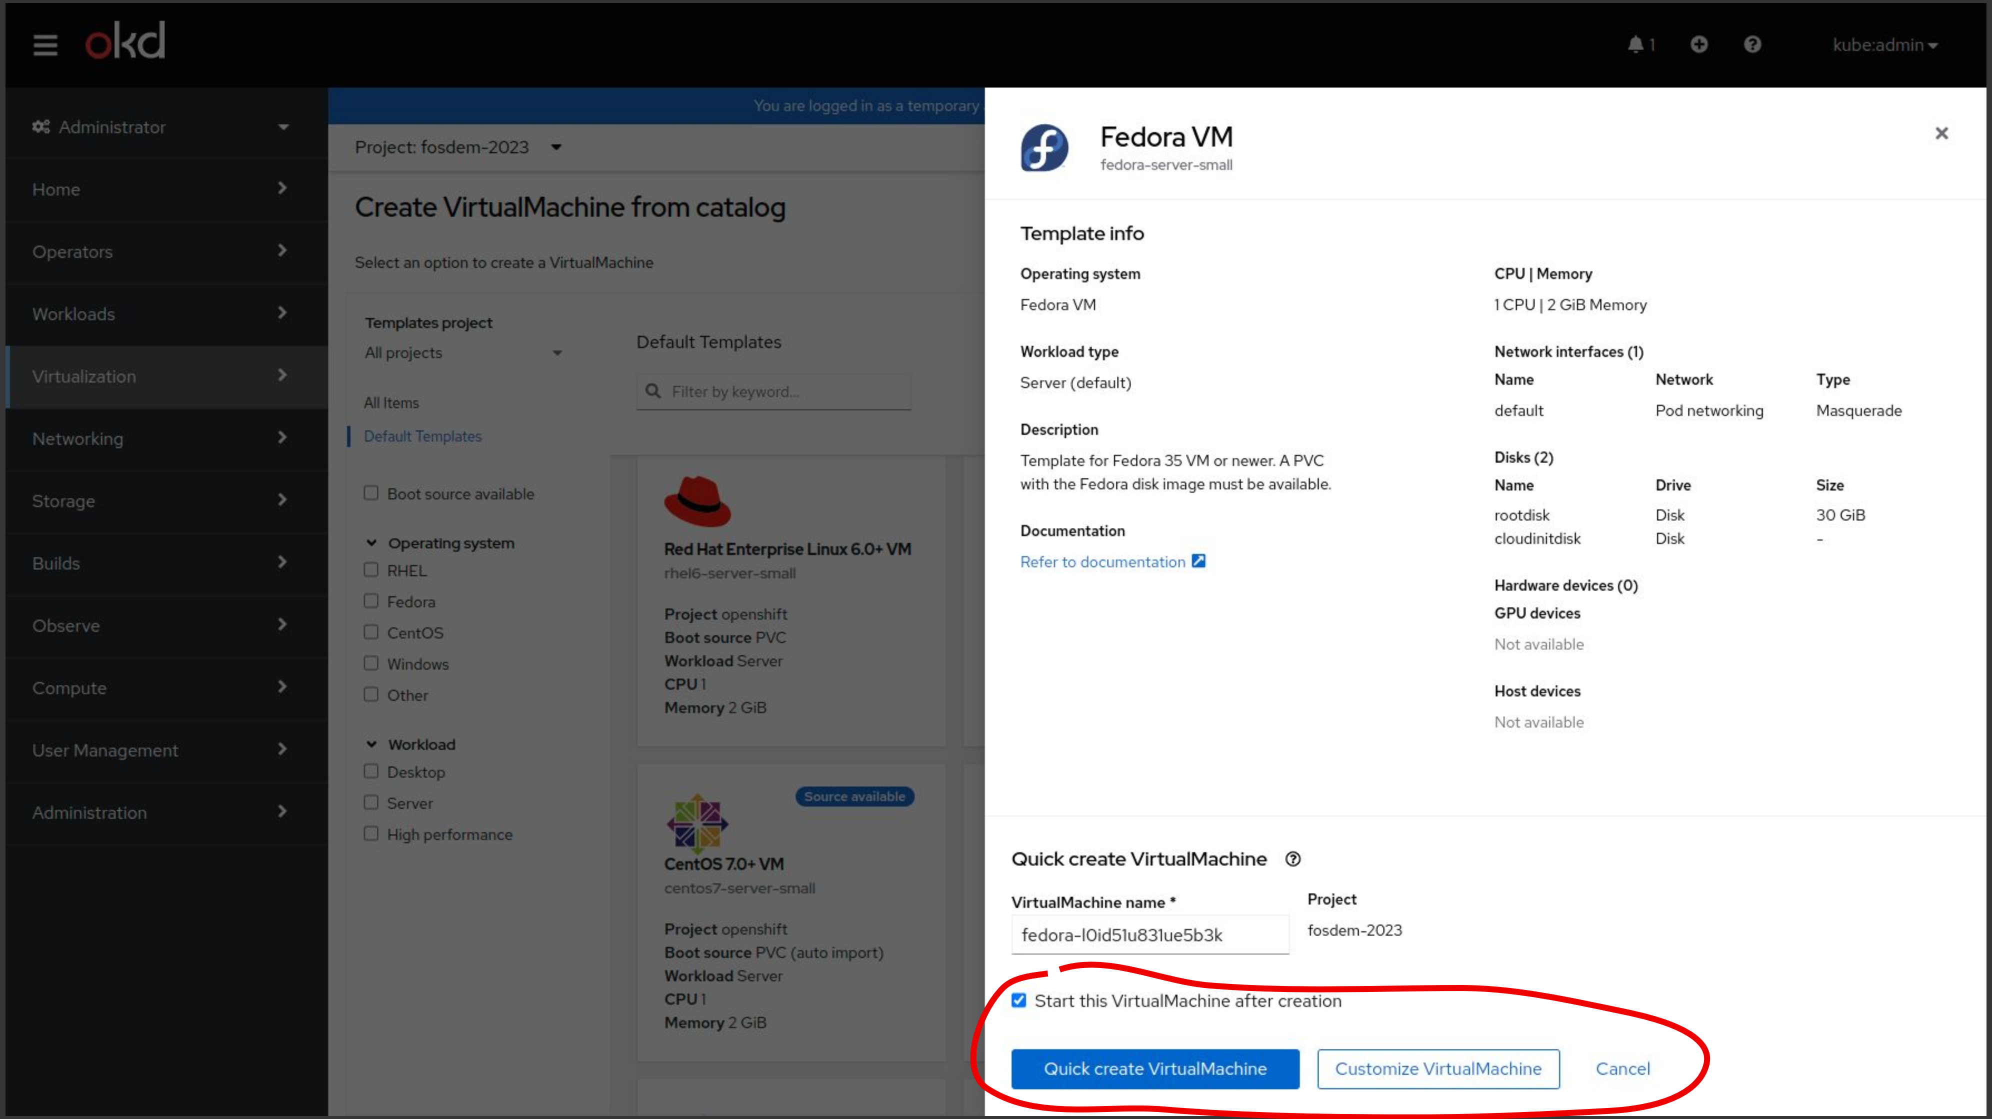

Project: fosdem-2023

Create VirtualMachine from catalog

Select an option to create a VirtualMachine

Templates project: All projects

All Items

Default Templates

□ Boot source available

Operating system
- □ RHEL
- □ Fedora
- □ CentOS
- □ Windows
- □ Other

Workload
- □ Desktop
- □ Server
- □ High performance

Default Templates

Filter by keyword...

Red Hat Enterprise Linux 6.0+ VM
rhel6-server-small

Project openshift
Boot source PVC
Workload Server
CPU 1
Memory 2 GiB

Source available

CentOS 7.0+ VM
centos7-server-small

Project openshift
Boot source PVC (auto import)
Workload Server
CPU 1
Memory 2 GiB

Fedora VM
fedora-server-small

Template info

Operating system
Fedora VM

Workload type
Server (default)

Description
Template for Fedora 35 VM or newer. A PVC with the Fedora disk image must be available.

Documentation
Refer to documentation

CPU | Memory
1 CPU | 2 GiB Memory

Network interfaces (1)

Name	Network	Type
default	Pod networking	Masquerade

Disks (2)

Name	Drive	Size
rootdisk	Disk	30 GiB
cloudinitdisk	Disk	-

Hardware devices (0)

GPU devices
Not available

Host devices
Not available

Quick create VirtualMachine

VirtualMachine name *: fedora-l0id51u831ue5b3k

Project: fosdem-2023

☑ Start this VirtualMachine after creation

Quick create VirtualMachine | Customize VirtualMachine | Cancel

On the storage side...

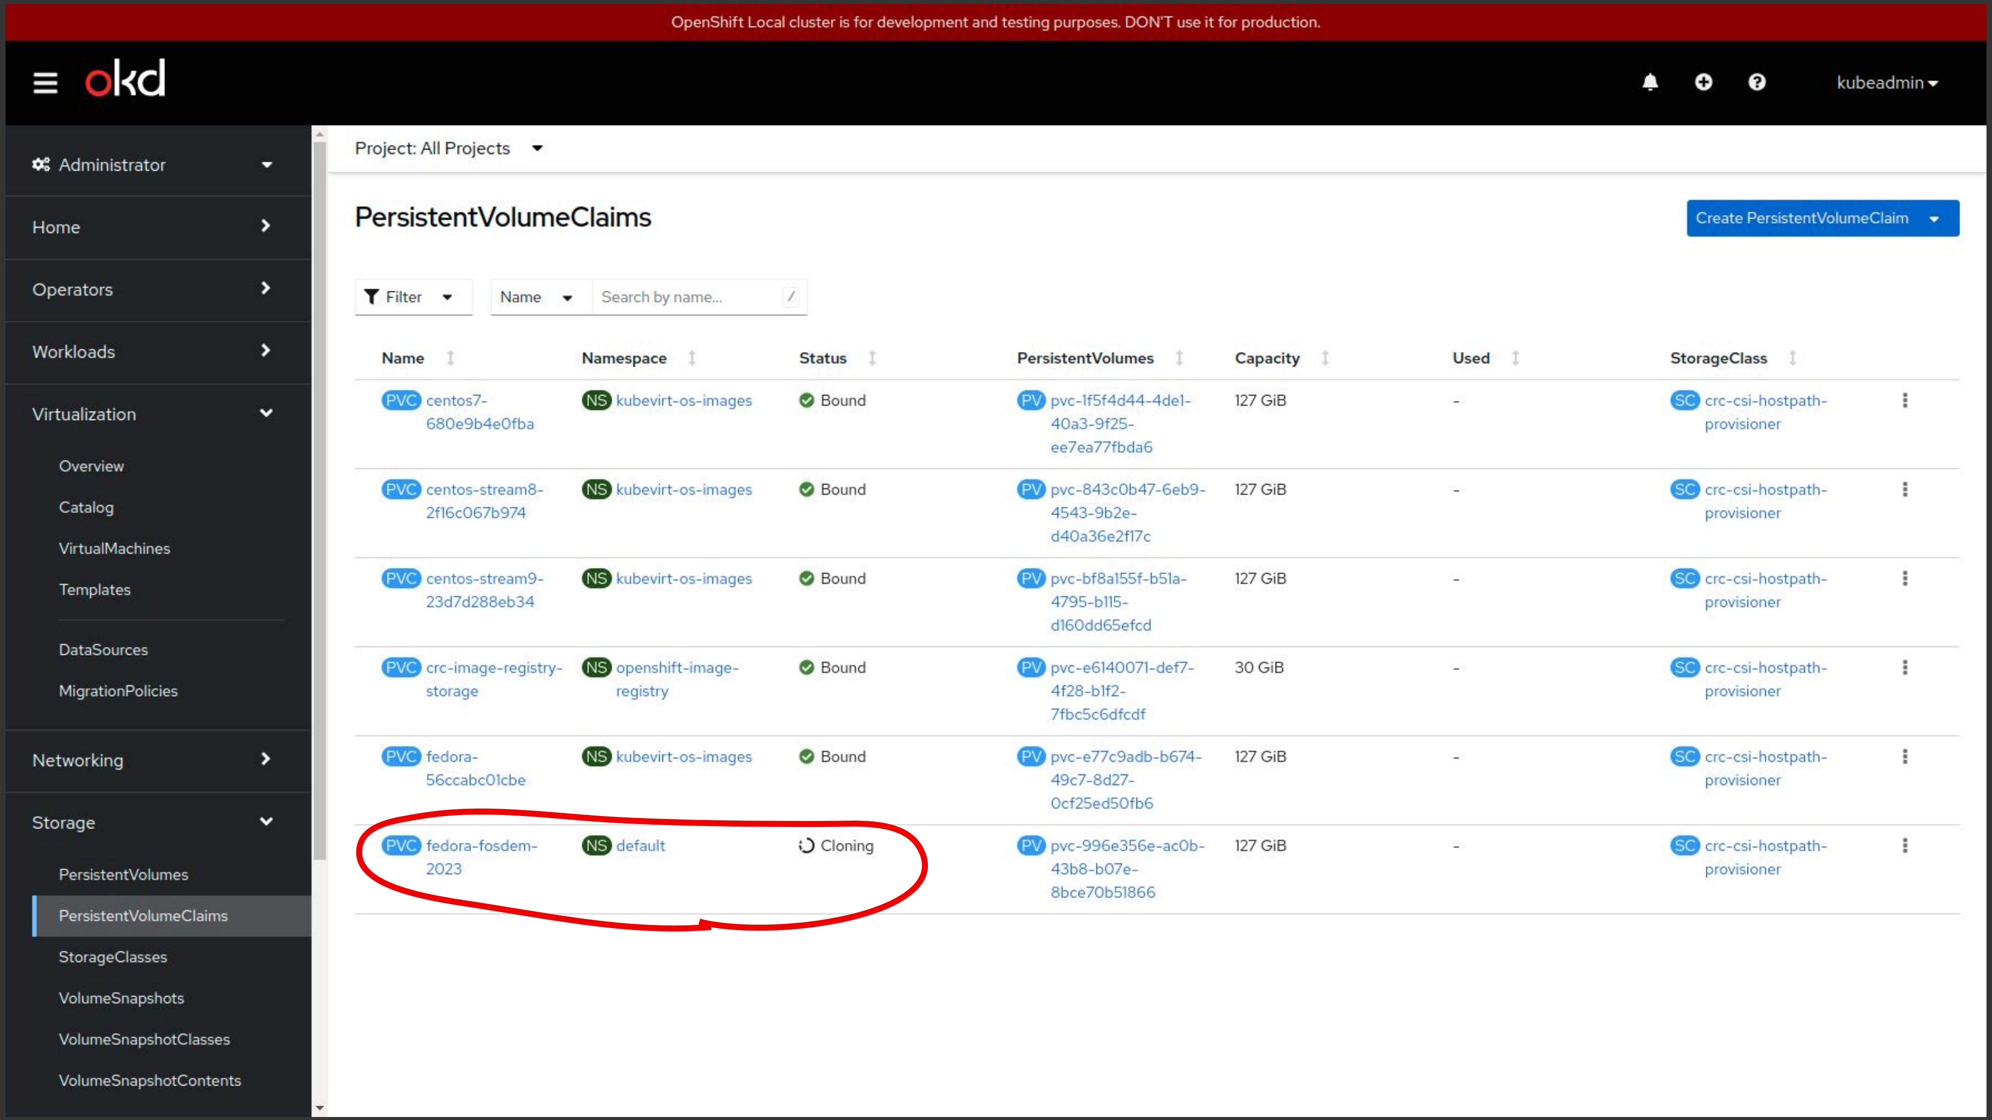

OpenShift Local cluster is for development and testing purposes. DON'T use it for production.

PersistentVolumeClaims

Name	Namespace	Status	PersistentVolumes	Capacity	Used	StorageClass
PVC centos7-680e9b4e0fba	NS kubevirt-os-images	✓ Bound	PV pvc-1f5f4d44-4de1-40a3-9f25-ee7ea77fbda6	127 GiB	-	SC crc-csi-hostpath-provisioner
PVC centos-stream8-2f16c067b974	NS kubevirt-os-images	✓ Bound	PV pvc-843c0b47-6eb9-4543-9b2e-d40a36e2f17c	127 GiB	-	SC crc-csi-hostpath-provisioner
PVC centos-stream9-23d7d288eb34	NS kubevirt-os-images	✓ Bound	PV pvc-bf8a155f-b51a-4795-b115-d160dd65efcd	127 GiB	-	SC crc-csi-hostpath-provisioner
PVC crc-image-registry-storage	NS openshift-image-registry	✓ Bound	PV pvc-e6140071-def7-4f28-b1f2-7fbc5c6dfcdf	30 GiB	-	SC crc-csi-hostpath-provisioner
PVC fedora-56ccabc01cbe	NS kubevirt-os-images	✓ Bound	PV pvc-e77c9adb-b674-49c7-8d27-0cf25ed50fb6	127 GiB	-	SC crc-csi-hostpath-provisioner
PVC fedora-fosdem-2023	NS default	Cloning	PV pvc-996e356e-ac0b-43b8-b07e-8bce70b51866	127 GiB	-	SC crc-csi-hostpath-provisioner

My nice VM is there...

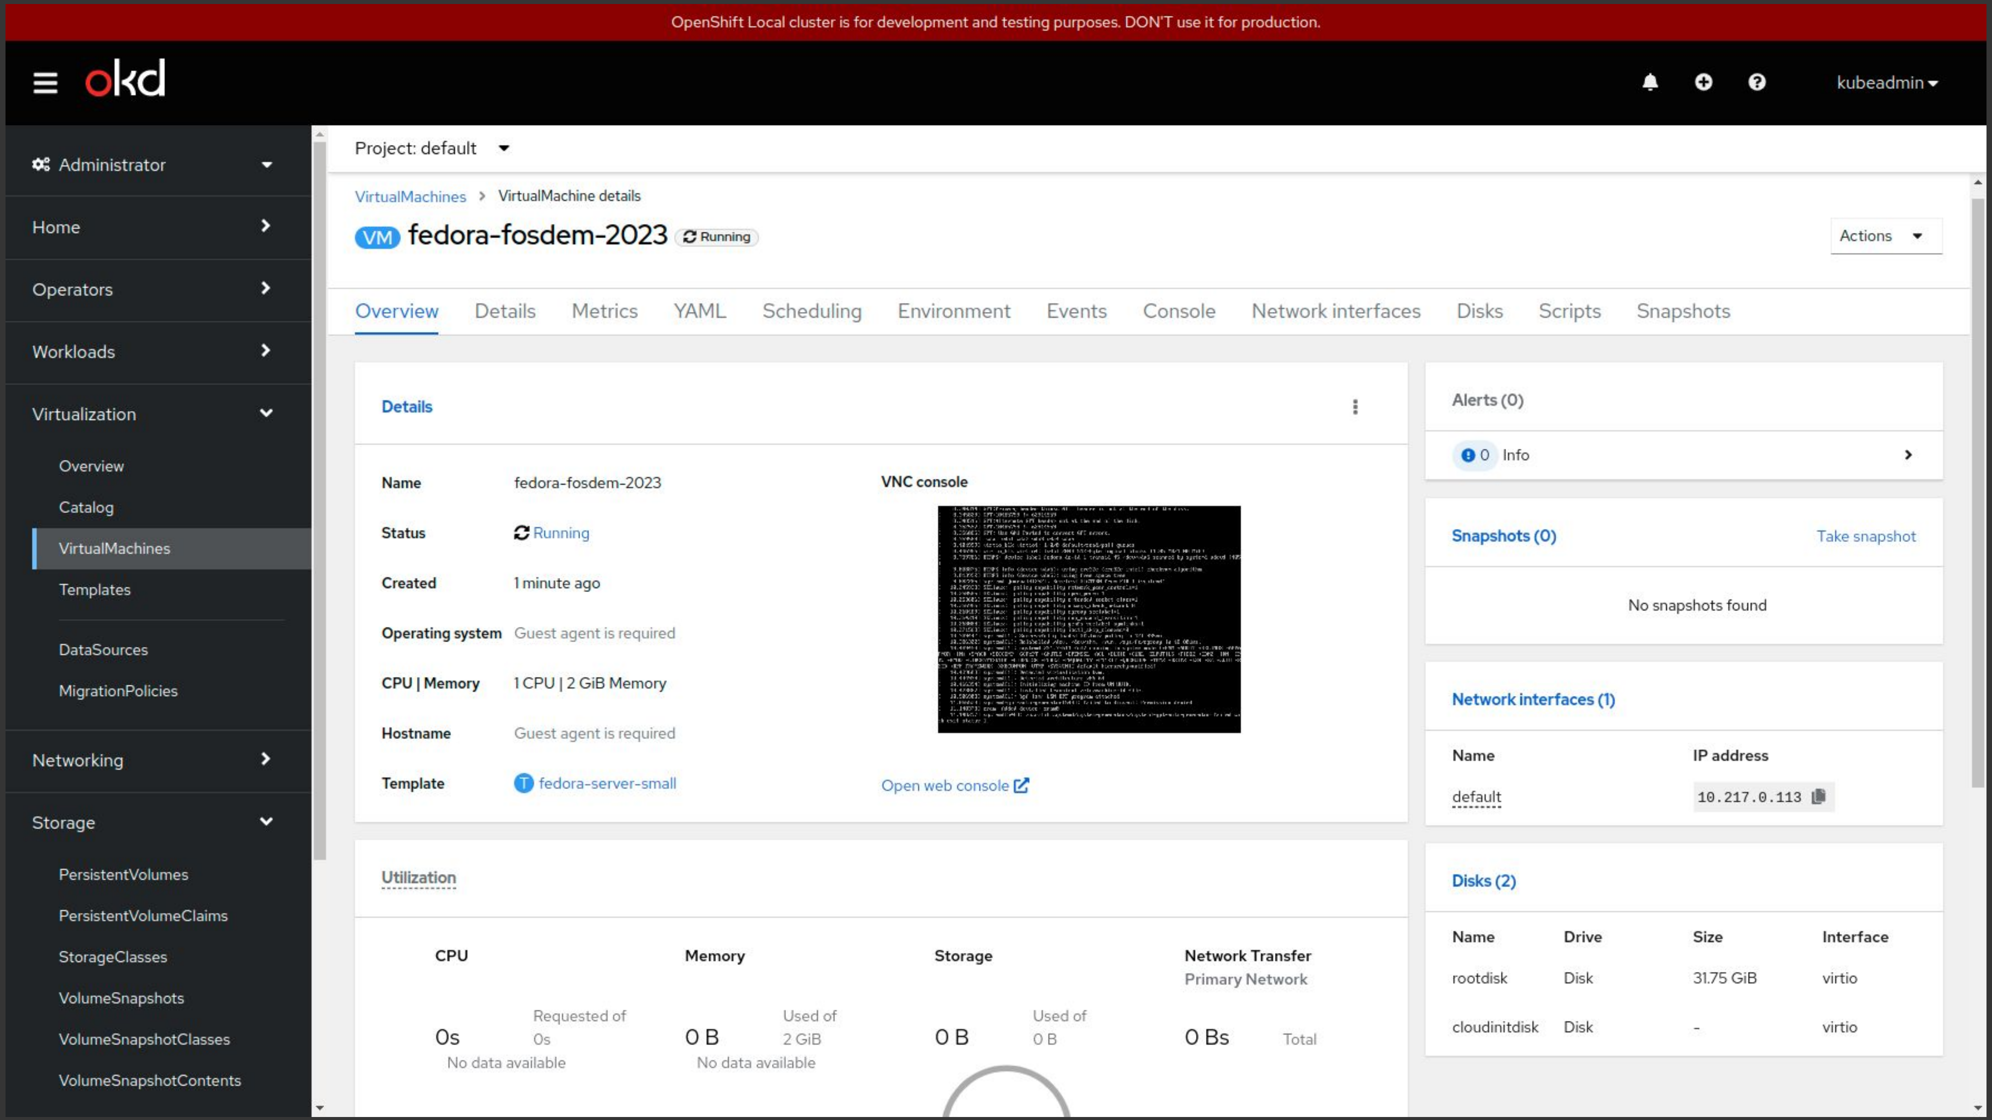

VM fedora-fosdem-2023 Running

Details

Name	fedora-fosdem-2023
Status	Running
Created	1 minute ago
Operating system	Guest agent is required
CPU | Memory	1 CPU | 2 GiB Memory
Hostname	Guest agent is required
Template	fedora-server-small

VNC console

Open web console

Utilization

CPU	Memory	Storage	Network Transfer Primary Network
0s Requested of 0s No data available	0 B Used of 2 GiB No data available	0 B Used of 0 B	0 Bs Total

Alerts (0)

0 Info

Snapshots (0) Take snapshot

No snapshots found

Network interfaces (1)

Name	IP address
default	10.217.0.113

Disks (2)

Name	Drive	Size	Interface
rootdisk	Disk	31.75 GiB	virtio
cloudinitdisk	Disk	-	virtio

And it has been customized from the template

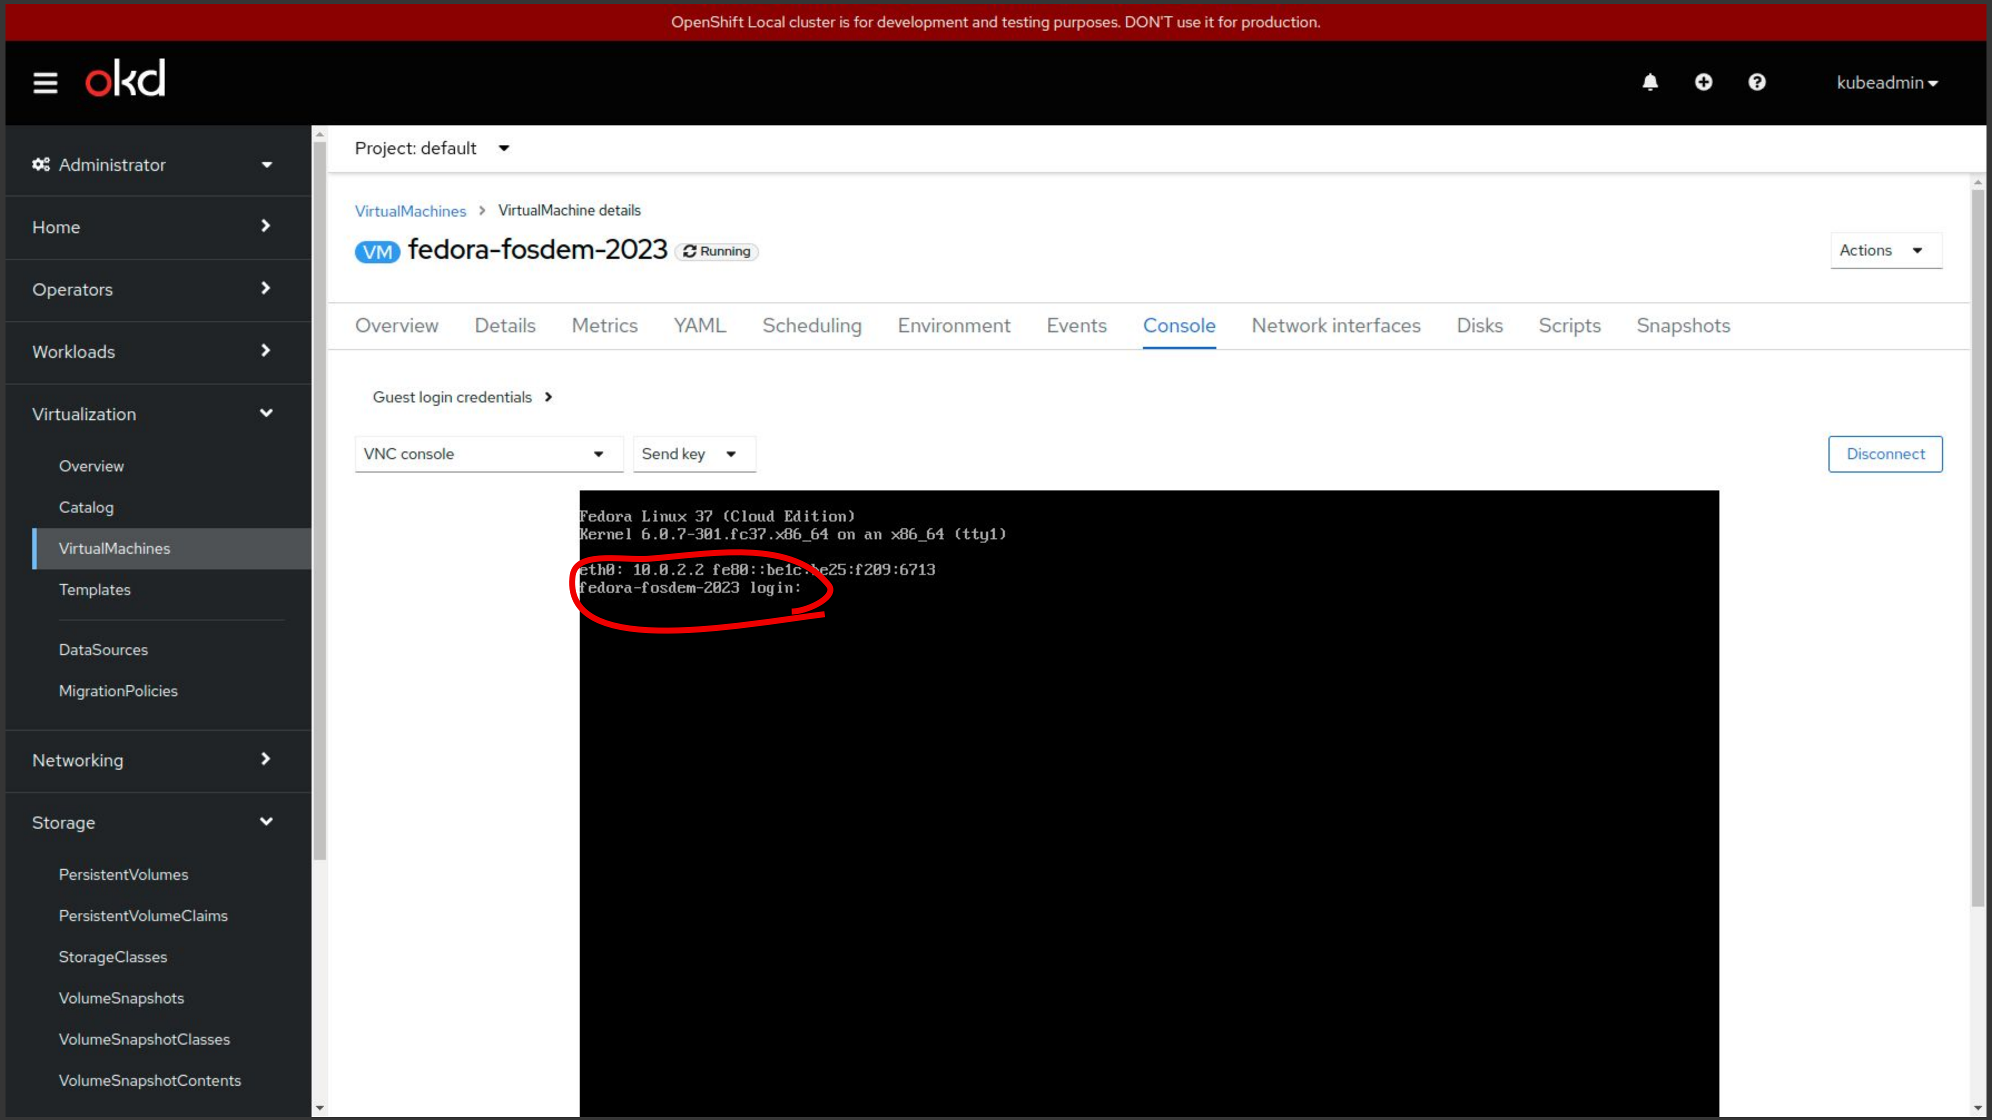

OpenShift Local cluster is for development and testing purposes. DON'T use it for production.

```
Fedora Linux 37 (Cloud Edition)
Kernel 6.0.7-301.fc37.x86_64 on an x86_64 (tty1)

eth0: 10.0.2.2 fe80::be1c:be25:f209:6713
fedora-fosdem-2023 login:
```

DataImportCrons

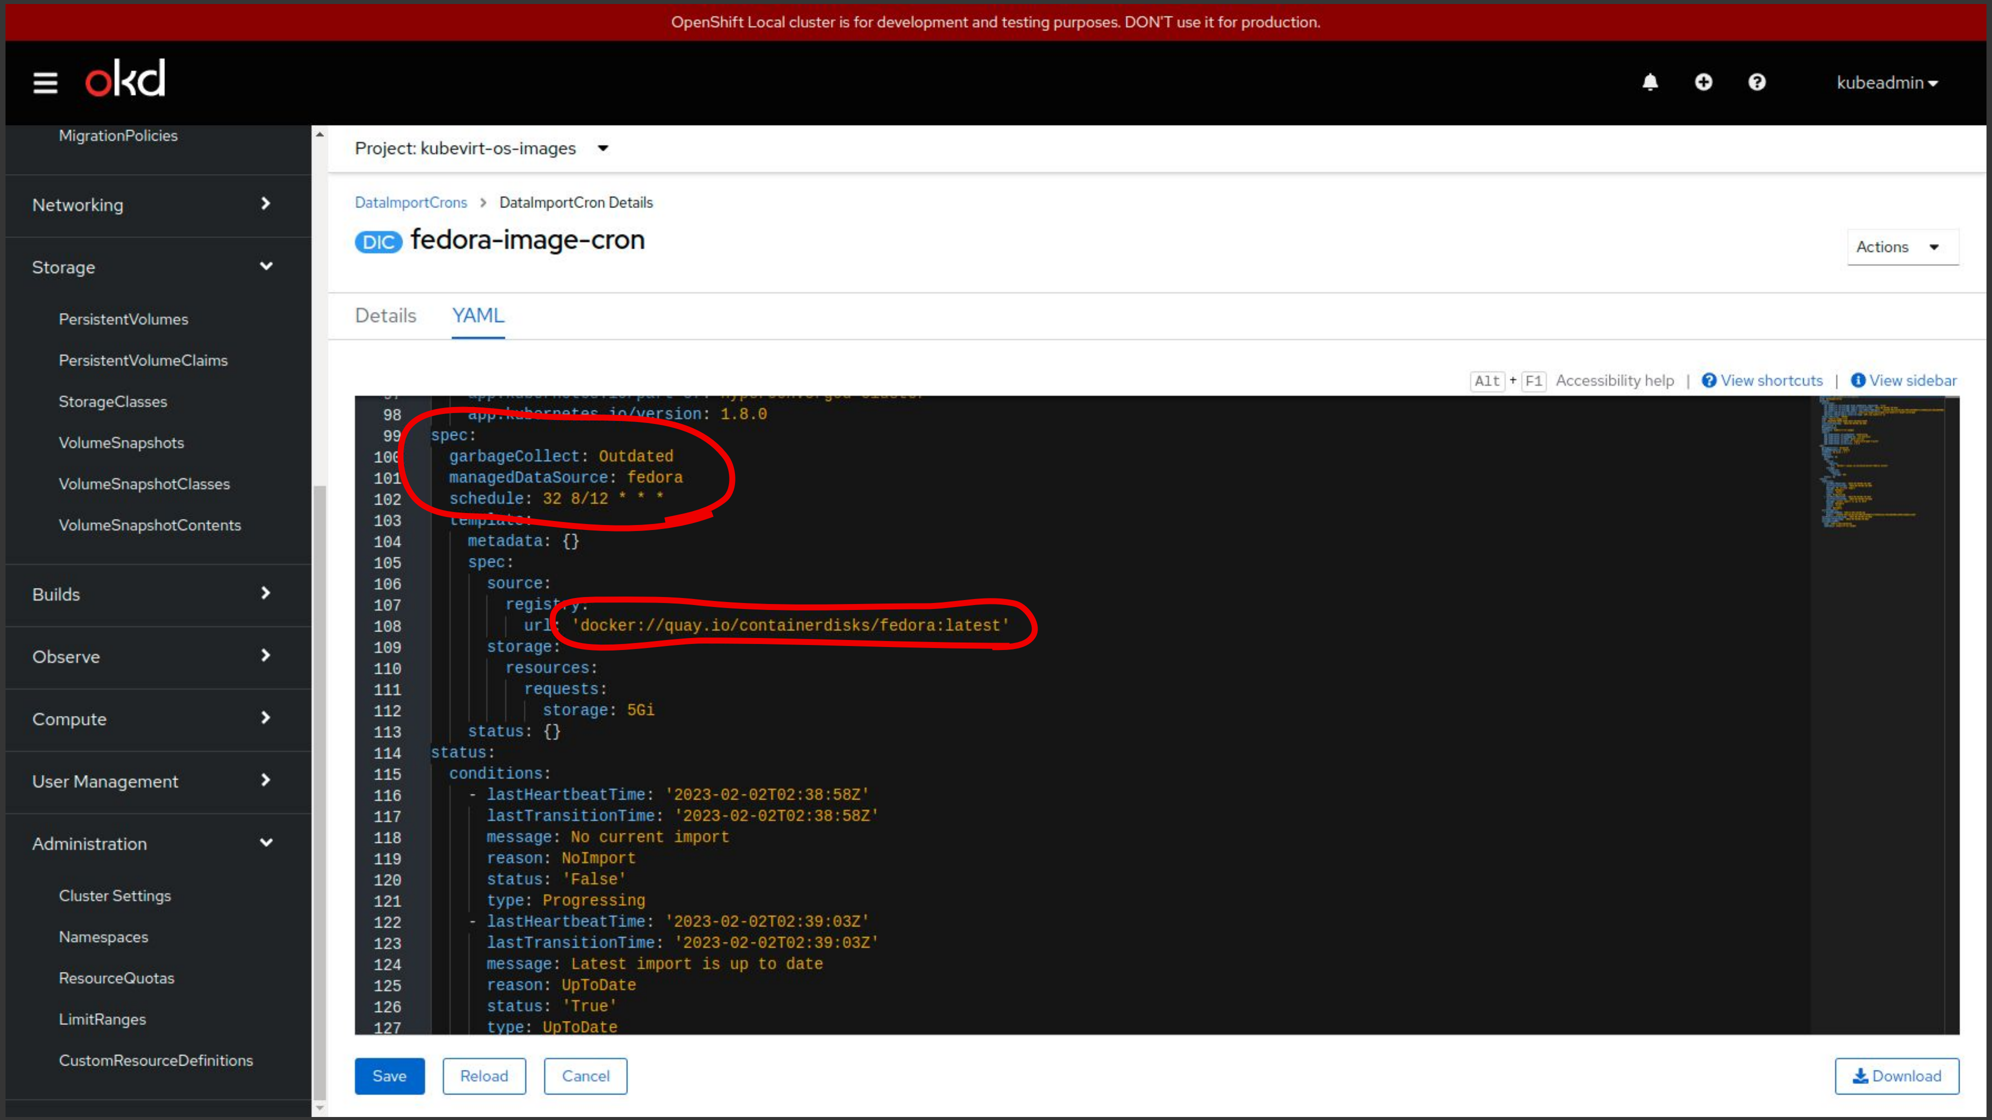

OpenShift Local cluster is for development and testing purposes. DON'T use it for production.

Project: kubevirt-os-images

fedora-image-cron

```yaml
    app.kubernetes.io/part-of: hyperconverged-cluster
    app.kubernetes.io/version: 1.8.0
spec:
  garbageCollect: Outdated
  managedDataSource: fedora
  schedule: 32 8/12 * * *
  template:
    metadata: {}
    spec:
      source:
        registry:
          url: 'docker://quay.io/containerdisks/fedora:latest'
      storage:
        resources:
          requests:
            storage: 5Gi
    status: {}
status:
  conditions:
    - lastHeartbeatTime: '2023-02-02T02:38:58Z'
      lastTransitionTime: '2023-02-02T02:38:58Z'
      message: No current import
      reason: NoImport
      status: 'False'
      type: Progressing
    - lastHeartbeatTime: '2023-02-02T02:39:03Z'
      lastTransitionTime: '2023-02-02T02:39:03Z'
      message: Latest import is up to date
      reason: UpToDate
      status: 'True'
      type: UpToDate
```

Custom DataImportCrons

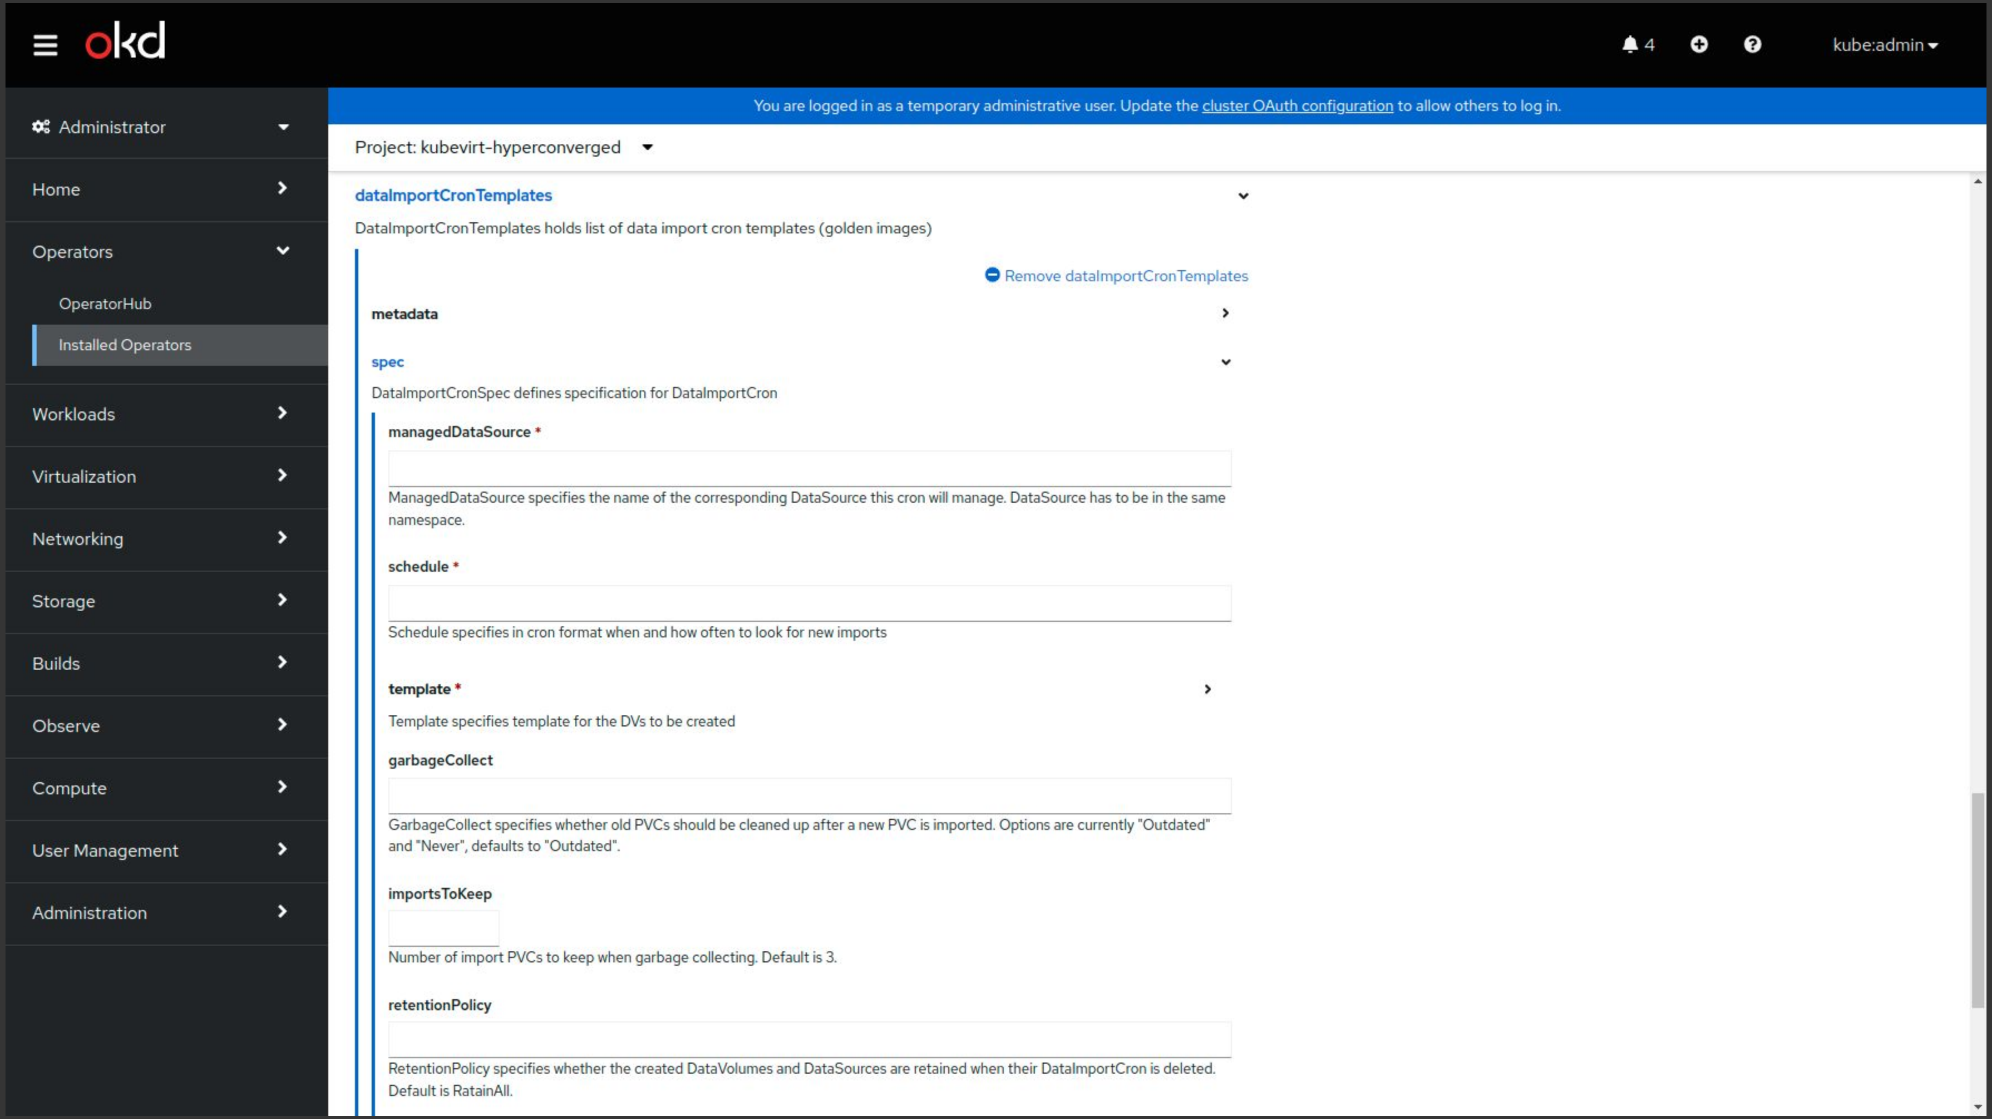

You are logged in as a temporary administrative user. Update the cluster OAuth configuration to allow others to log in.

Project: kubevirt-hyperconverged

dataImportCronTemplates

DataImportCronTemplates holds list of data import cron templates (golden images)

Remove dataImportCronTemplates

metadata

spec

DataImportCronSpec defines specification for DataImportCron

managedDataSource *

ManagedDataSource specifies the name of the corresponding DataSource this cron will manage. DataSource has to be in the same namespace.

schedule *

Schedule specifies in cron format when and how often to look for new imports

template *

Template specifies template for the DVs to be created

garbageCollect

GarbageCollect specifies whether old PVCs should be cleaned up after a new PVC is imported. Options are currently "Outdated" and "Never", defaults to "Outdated".

importsToKeep

Number of import PVCs to keep when garbage collecting. Default is 3.

retentionPolicy

RetentionPolicy specifies whether the created DataVolumes and DataSources are retained when their DataImportCron is deleted. Default is RatainAll.

Kubevirt
Tekton
Pipelines

Kubevirt Tekton Pipelines: motivation

How can I automate the execution of complex, long and error prone tasks?

Like creating a custom golden image and installing the OS,
Configuring it,
Updating...

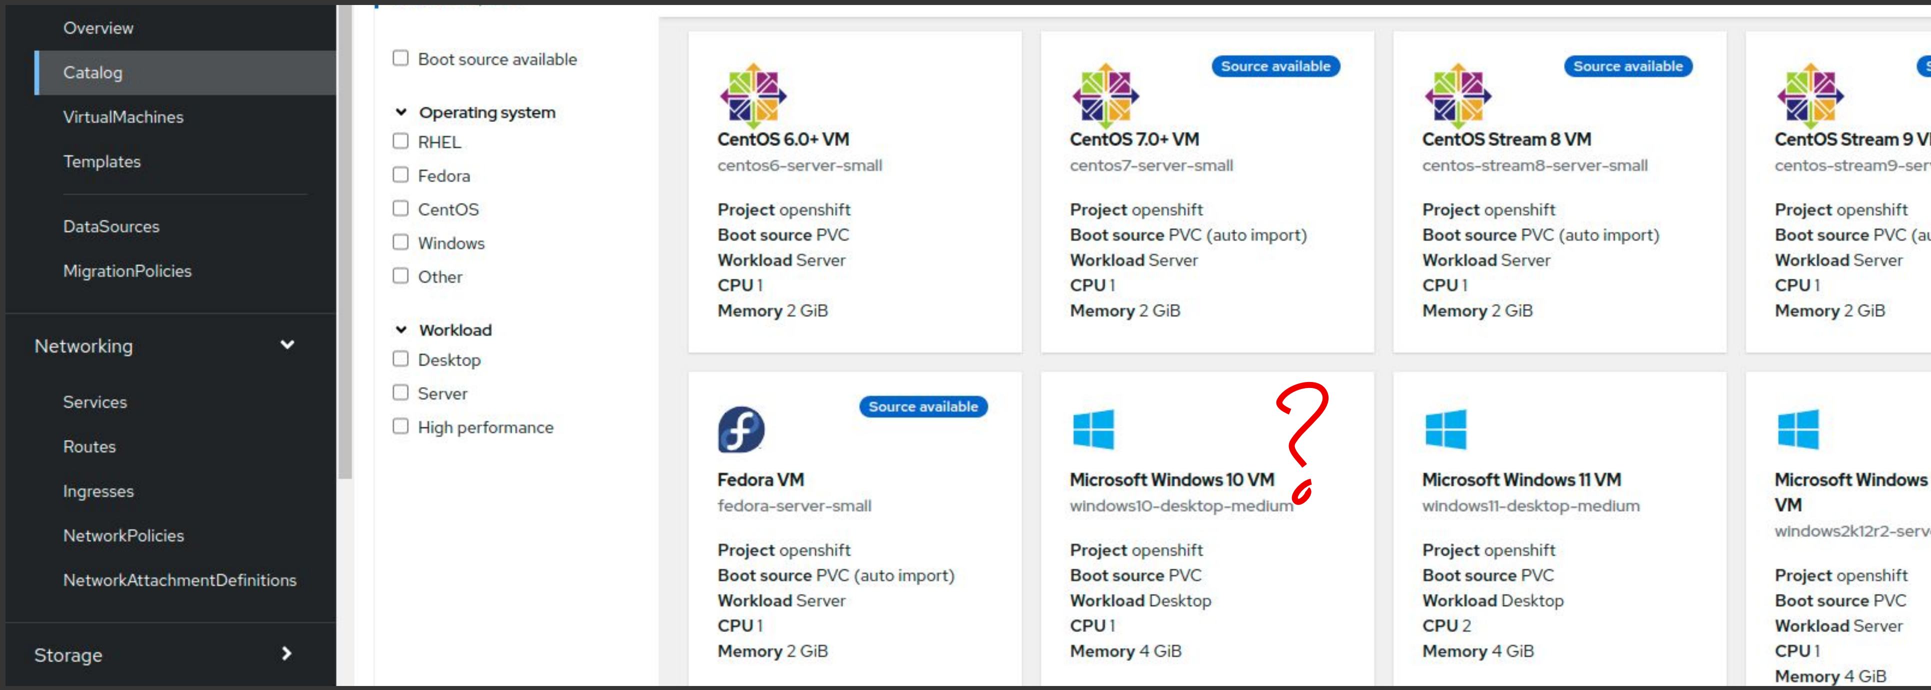

CentOS 6.0+ VM
centos6-server-small
Project openshift
Boot source PVC
Workload Server
CPU 1
Memory 2 GiB

CentOS 7.0+ VM
centos7-server-small
Project openshift
Boot source PVC (auto import)
Workload Server
CPU 1
Memory 2 GiB

CentOS Stream 8 VM
centos-stream8-server-small
Project openshift
Boot source PVC (auto import)
Workload Server
CPU 1
Memory 2 GiB

CentOS Stream 9 V
centos-stream9-ser
Project openshift
Boot source PVC (a
Workload Server
CPU 1
Memory 2 GiB

Fedora VM
fedora-server-small
Project openshift
Boot source PVC (auto import)
Workload Server
CPU 1
Memory 2 GiB

Microsoft Windows 10 VM
windows10-desktop-medium
Project openshift
Boot source PVC
Workload Desktop
CPU 1
Memory 4 GiB

Microsoft Windows 11 VM
windows11-desktop-medium
Project openshift
Boot source PVC
Workload Desktop
CPU 2
Memory 4 GiB

Microsoft Windows VM
windows2k12r2-serv
Project openshift
Boot source PVC
Workload Server
CPU 1
Memory 4 GiB

Tekton

- Tekton AKA OpenShift Pipelines is a cloud-native, continuous integration and continuous delivery (CI/CD) solution based on Kubernetes resources

- It uses Tekton building blocks to automate deployments across multiple platforms by abstracting away the underlying implementation details

- Tekton introduces a number of standard custom resource definitions (CRDs) for defining CI/CD pipelines that are portable across Kubernetes distributions.

Still not available out of the box in community-operators

You can install it with:

```
$ export TEKTON_VERSION=v0.64.0
$ oc apply -f
"https://github.com/tektoncd/operator/releases/download
/${TEKTON_VERSION}/openshift-release.yaml"
```

How does Tekton works?

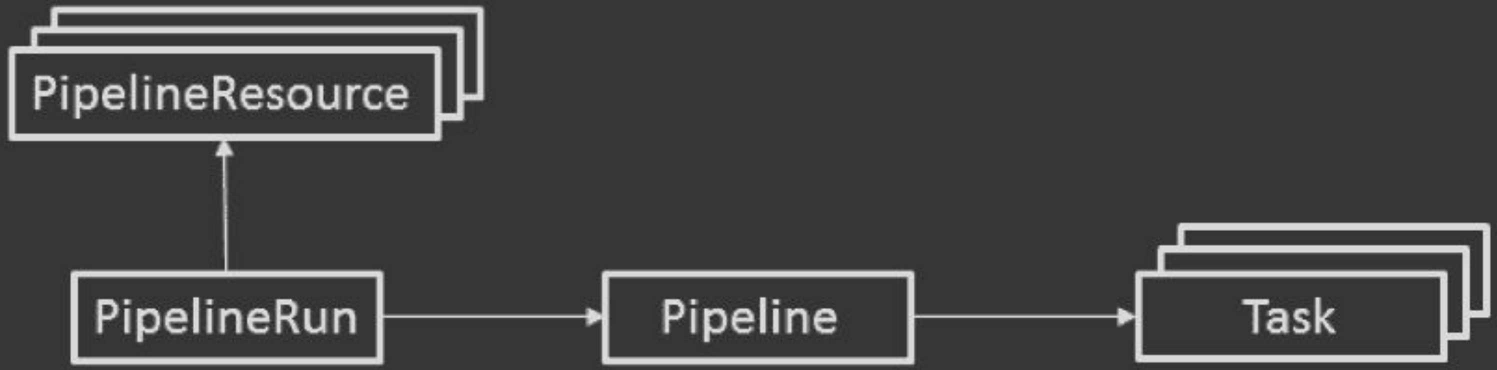

- Task: Defines a set of build steps, such as compiling code, running tests, and building and deploying images.
- Pipeline: Defines the set of tasks that compose a pipeline.
- PipelineResource: Defines an object that is an input (such as a Git repository) or an output (such as a Docker image) of the pipeline.
- PipelineRun: Instantiates a Pipeline for execution with specific inputs, outputs, and execution parameters.

Kubevirt Tekton tasks

KubeVirt Tekton tasks provide OpenShift Virtualization - specific Tekton tasks, which focus on:

- Creating, updating and managing resources of KubeVirt (VMs, DataVolumes, DataSources, Templates)
- Manipulating disk images with libguestfs tools

https://github.com/kubevirt/kubevirt-tekton-tasks

Kubevirt Tekton tasks

- Create Virtual Machines
  - create-vm-from-manifest - create VM from yaml manifest
  - create-vm-from-template - create VM from template
- Utilize Templates
  - copy-template - copy template
  - modify-vm-template - update template metadata
- Create DataVolumes/DataSources
  - create-datavolume-from-manifest (4.11) - create dataVolume
  - modify-data-object (>= 4.12) - create/delete dataVolume/dataSource
- Generate SSH Keys
  - generate-ssh-keys - generate SSH keys and store them in cluster
- Execute Commands in Virtual Machines
  - execute-in-vm: execute commands over SSH
  - cleanup-vm: execute commands and/or stop/delete VMs
- Manipulate PVCs with libguestfs tools
  - disk-virt-customize: execute virt-customize commands in PVCs
  - disk-virt-sysprep: execute virt-sysprep commands in PVCs
- Wait for Virtual Machine Instance Status
  - wait-for-vmi-status - wait for VM status

Tekton Task operator

- The Tekton Tasks Operator (TTO) is a operator that takes care of deploying Kubevirt tasks and example pipelines.

- Starting from OKD Virtualization 4.11, TTO is deployed by default, but does not deploy any resources until explicitly enabled.

- To enable resource creation - spec.featureGates.deployTektonTaskResources needs to be updated in the HCO CR:

```
$ oc patch hco kubevirt-hyperconverged  --type=merge -p '{"spec":{"featureGates":
{"deployTektonTaskResources": true}}}'
```

https://github.com/kubevirt/tekton-tasks-operator

Tekton Task operator - Example pipelines

- windows10-installer - Populates golden windows 10 image in openshift-virtualization-os-images namespace
- windows10-customize - Installs sql server in golden windows image and creates new image and template

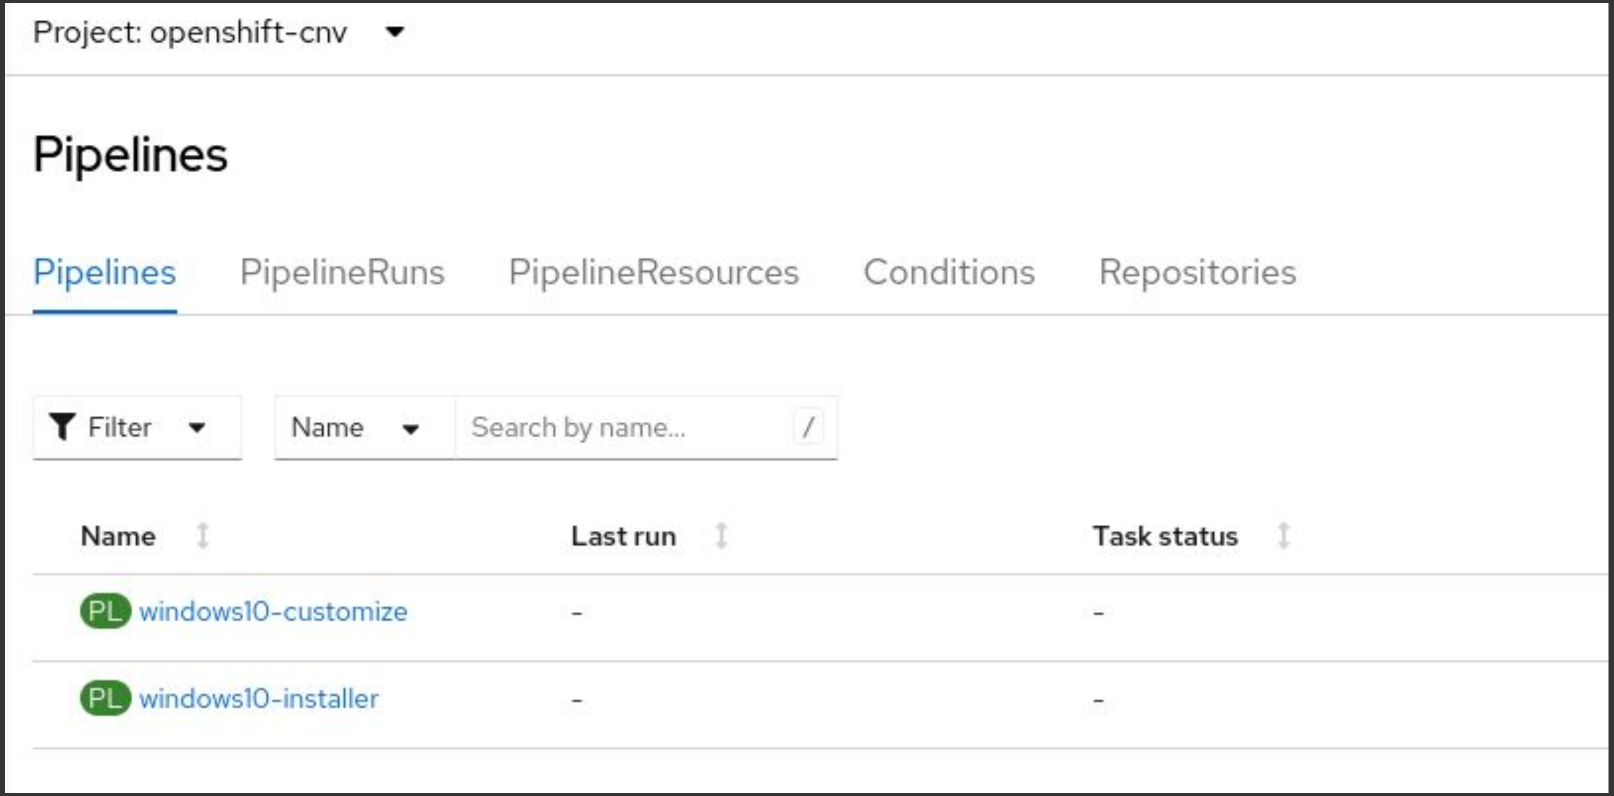

Project: openshift-cnv

Pipelines

Name	Last run	Task status
PL windows10-customize	-	-
PL windows10-installer	-	-

Example pipelines - windows10-installer

- Populates golden windows image in kubevirt-os-images namespace
- Installs virtio drivers
- The minimal required input is the URL to the Windows 10 ISO
- The default Windows 10 template will detect the golden image as default disk image automatically
- Pipeline definition: https://github.com/kubevirt/tekton-tasks-operator/blob/main/data/tekton-pipelines/okd/windows10-installer.yaml

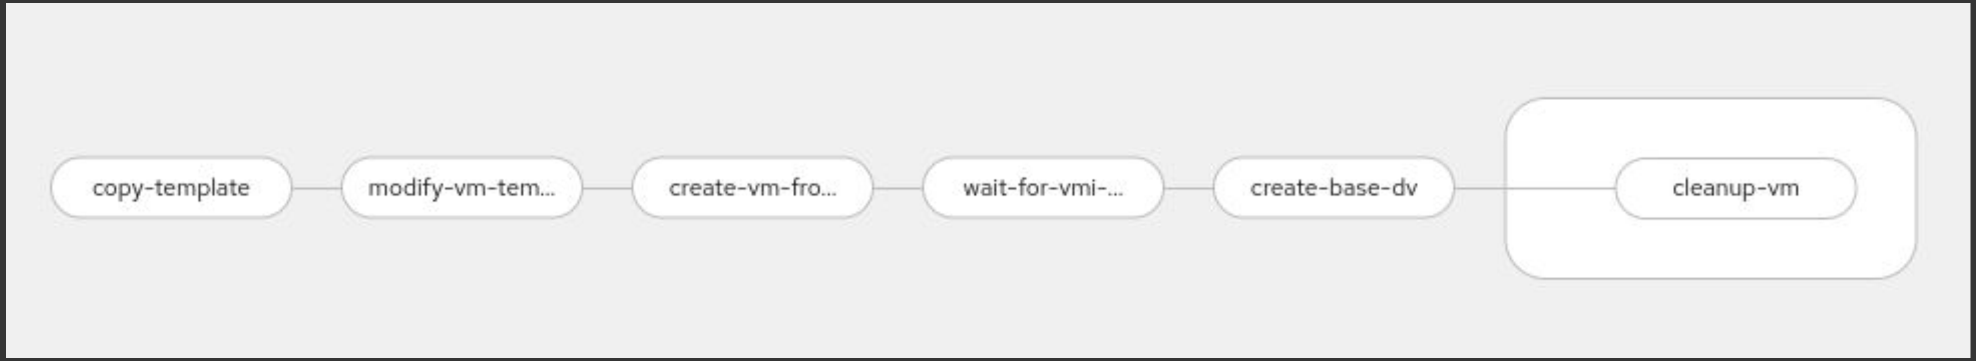

Example pipelines - windows10-installer demo

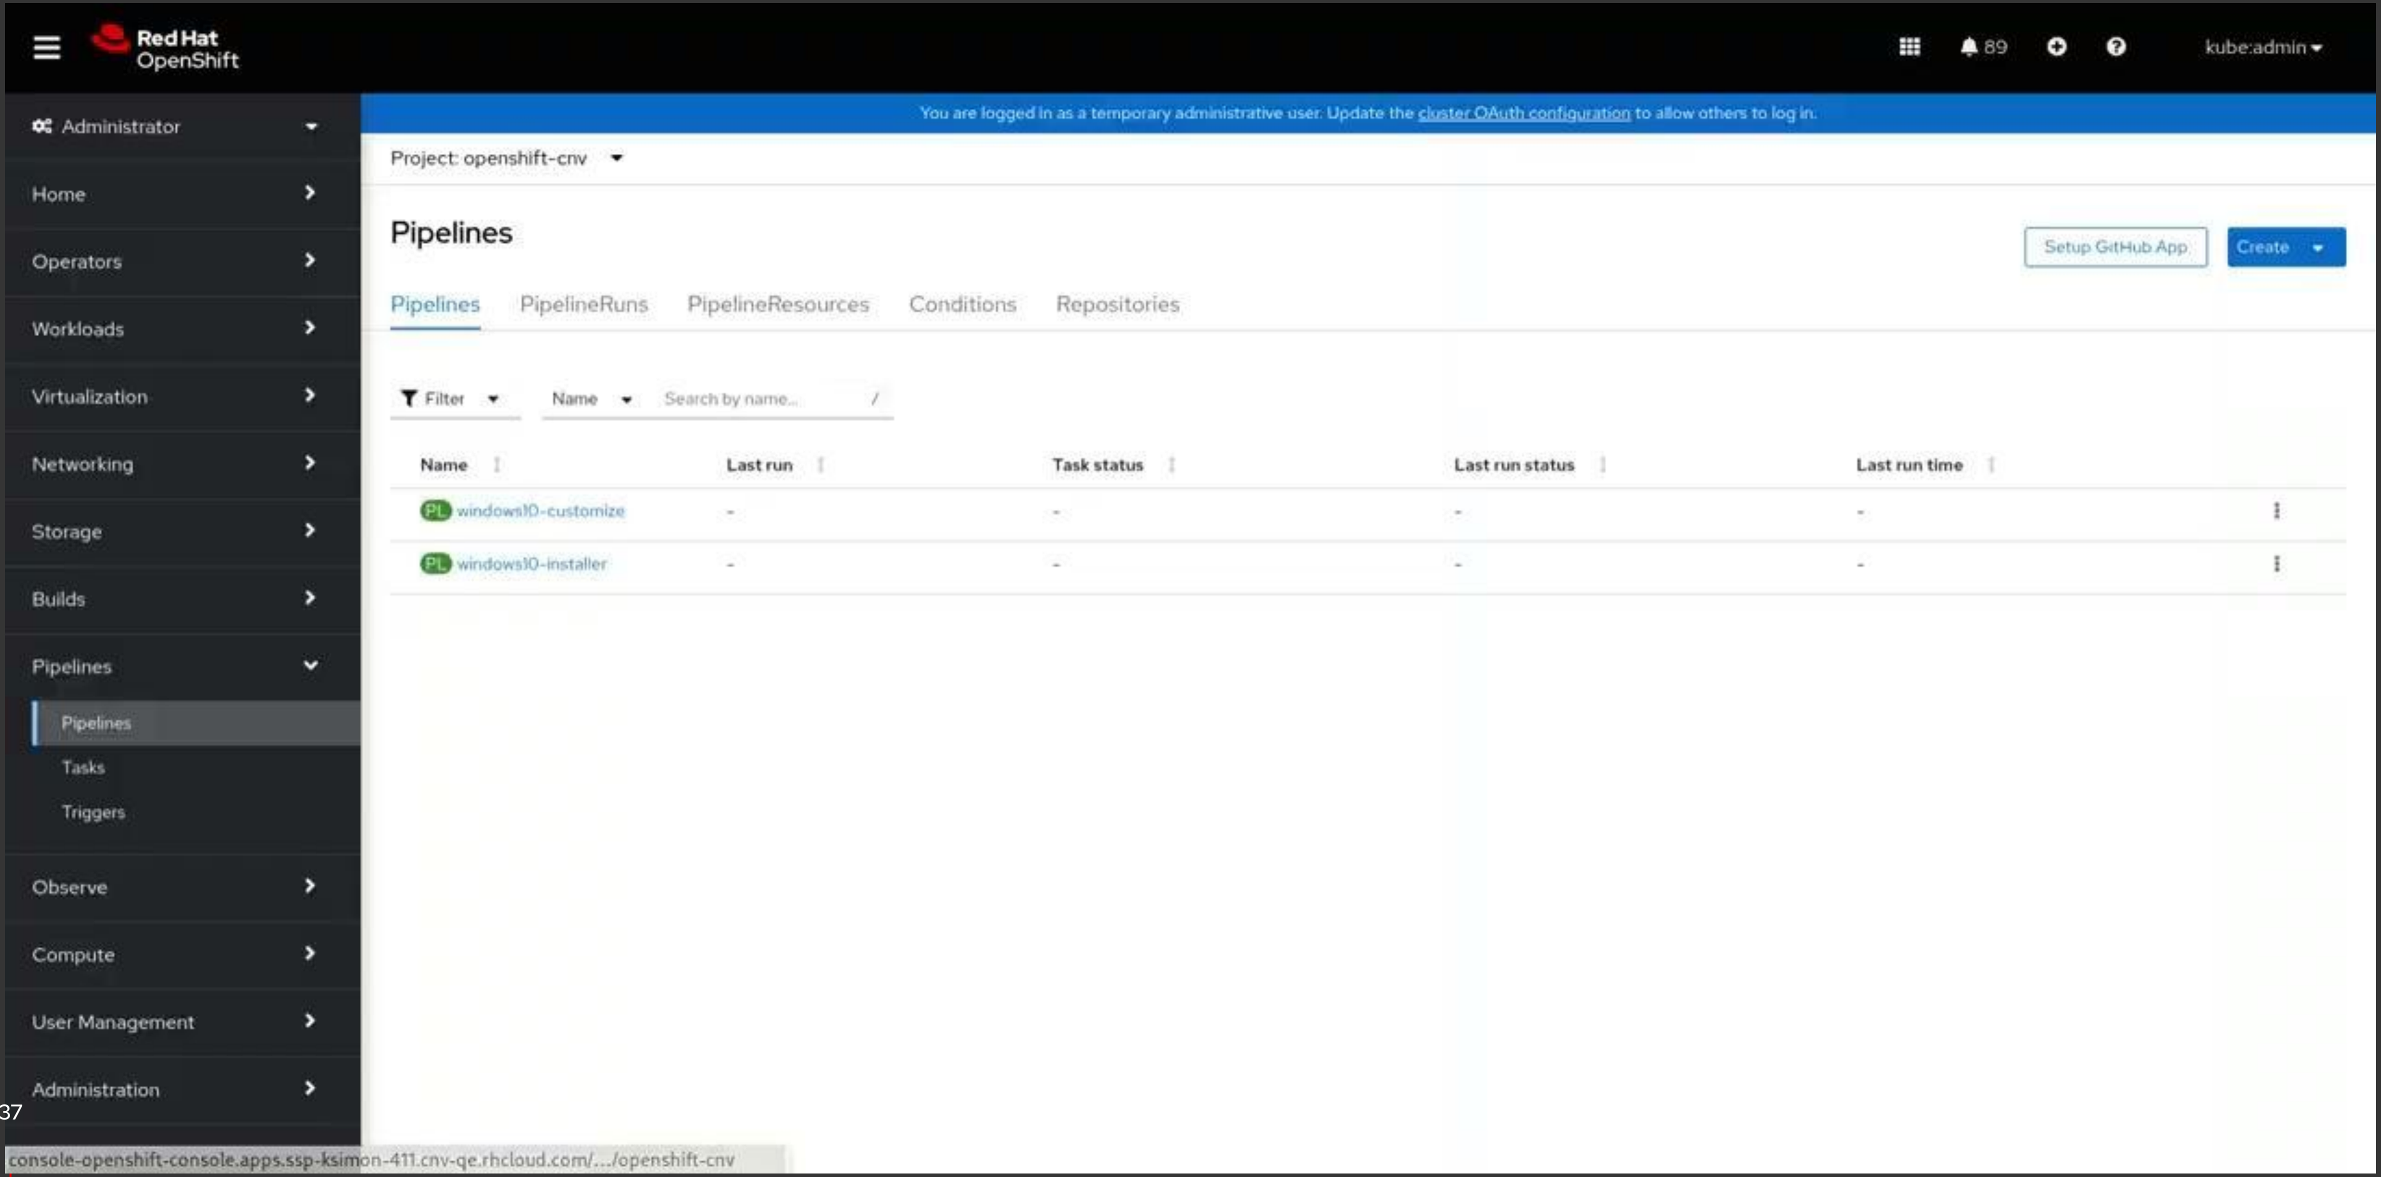

You are logged in as a temporary administrative user. Update the cluster OAuth configuration to allow others to log in.

Project: openshift-cnv

Pipelines

Name	Last run	Task status	Last run status	Last run time
PL windows10-customize	-	-	-	-
PL windows10-installer	-	-	-	-

console-openshift-console.apps.ssp-ksimon-411.cnv-qe.rhcloud.com/.../openshift-cnv

Example pipelines - windows10-customize

- Creates golden images with customizations applied on top of a basic Windows installation
- Uses image created by windows10-installer pipeline
- Example pipeline installs MS SQL Server in Windows 10

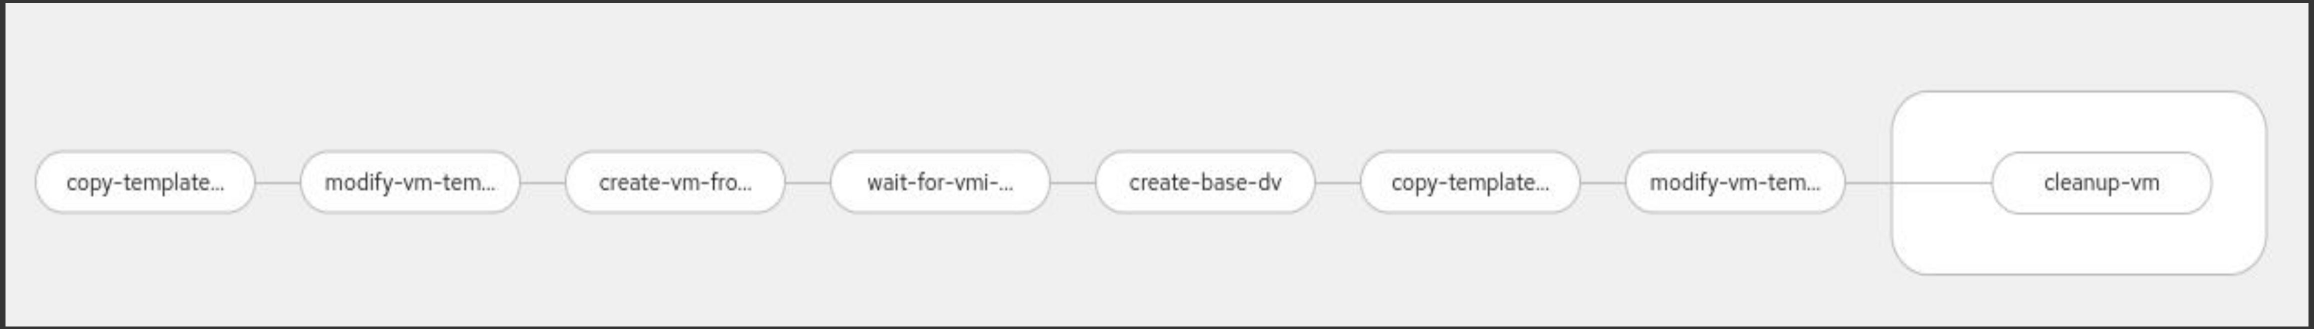

Example pipelines - windows10-customize demo

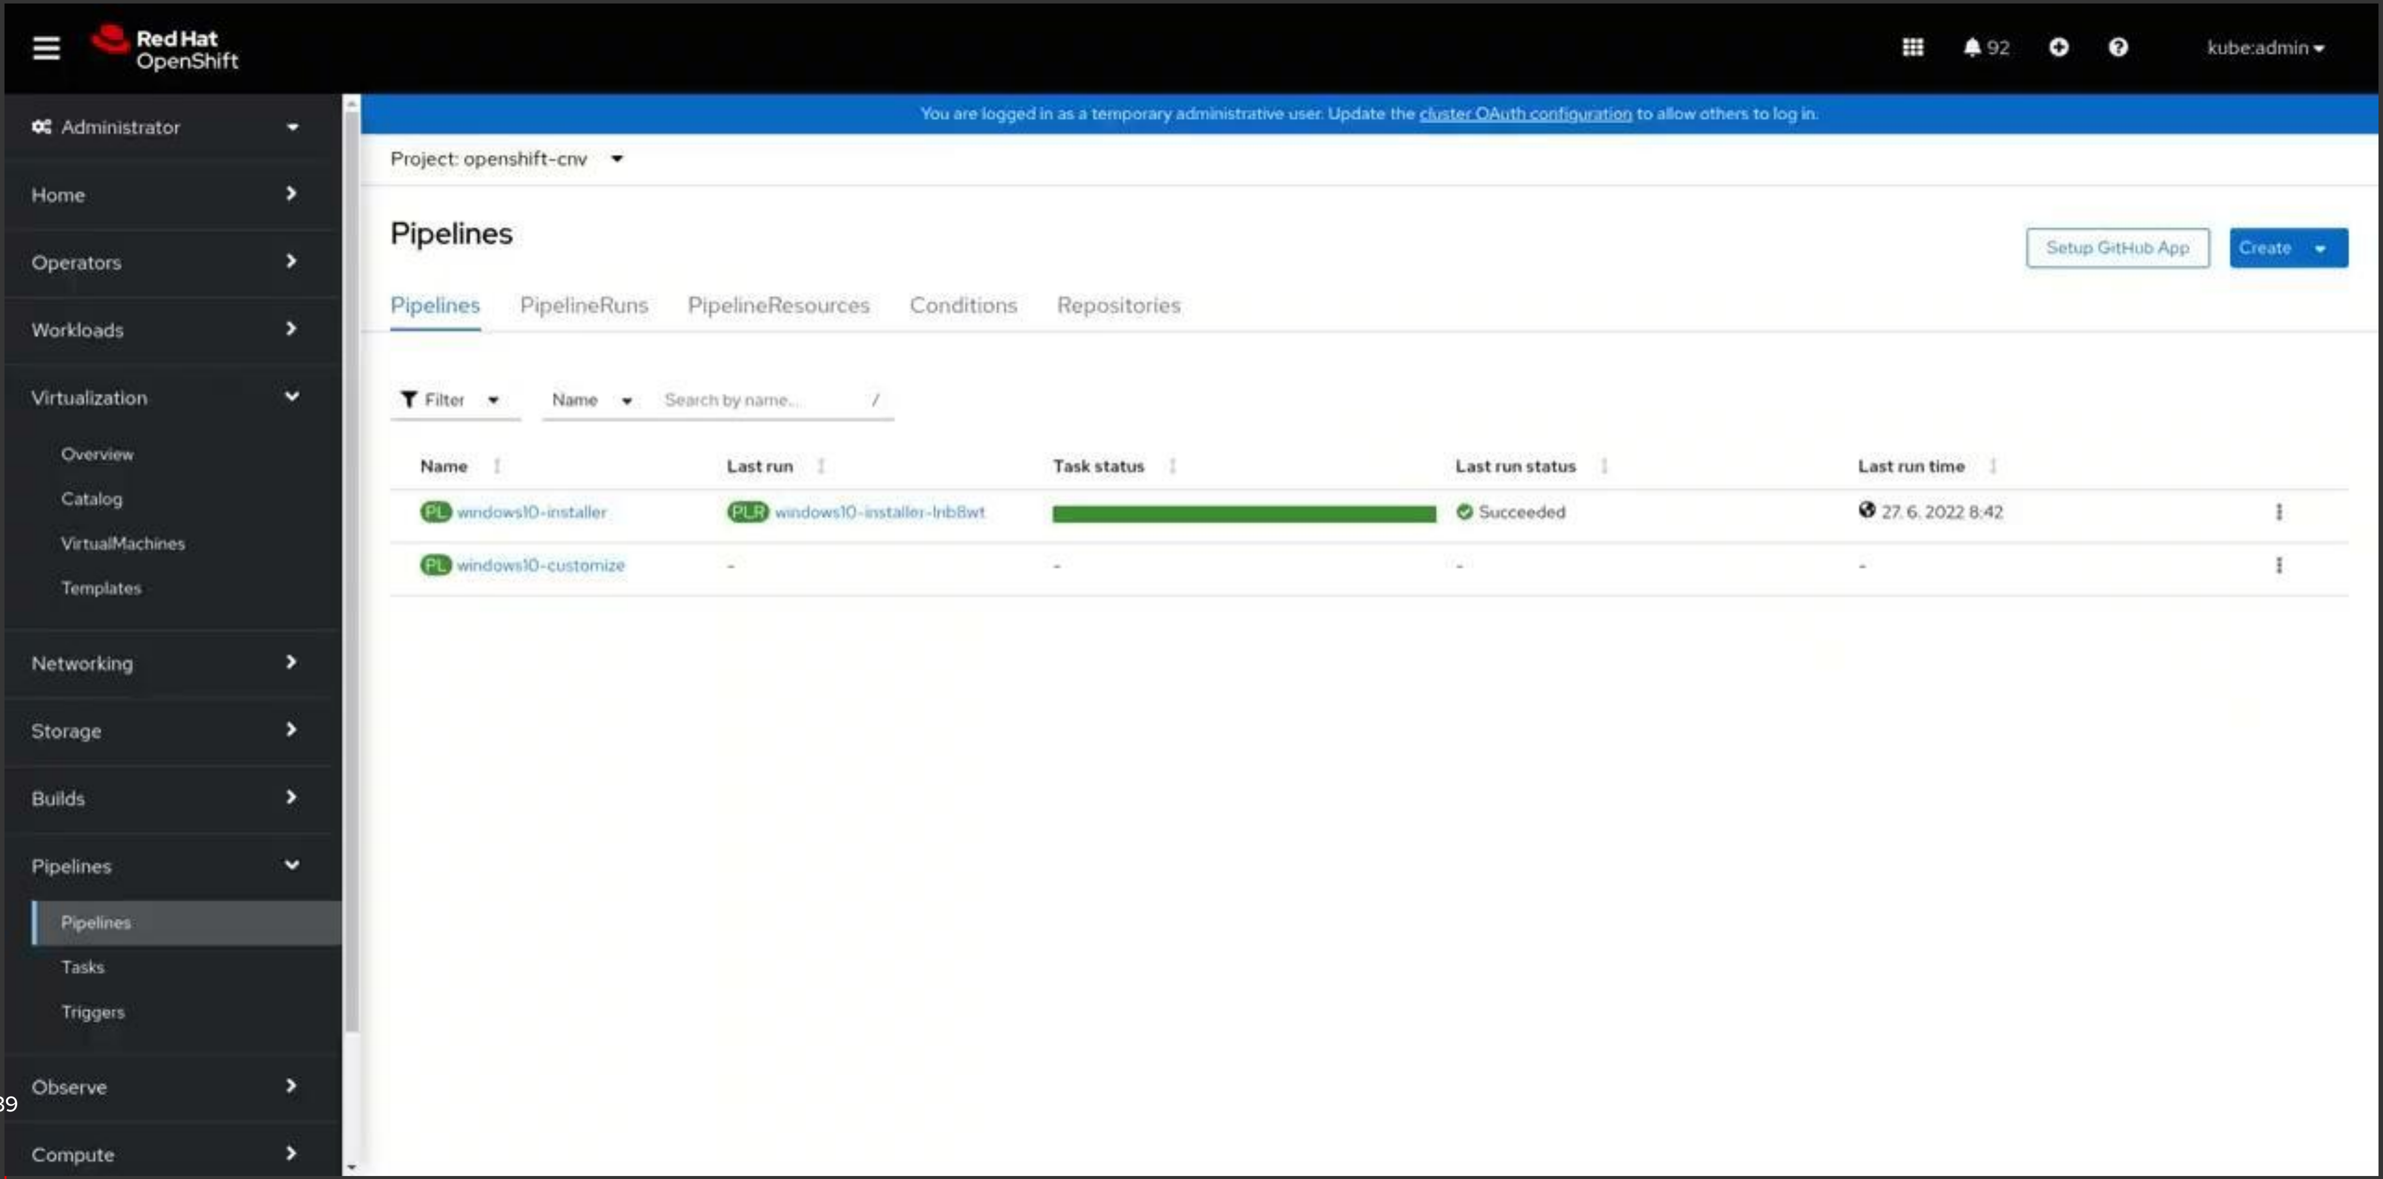

You are logged in as a temporary administrative user. Update the cluster OAuth configuration to allow others to log in.

Project: openshift-cnv

Pipelines

Name	Last run	Task status	Last run status	Last run time
PL windows10-installer	PLR windows10-installer-lnb8wt	🟩	✓ Succeeded	27. 6. 2022 8:42
PL windows10-customize	-	-	-	-

What's next

What's next

OKD: New Patterns, New CI/CD Pipelines and a new CoreOS

- What we call OKD is now "OKD running on the latest stable release of Fedora CoreOS (FCOS for short)"
- OKD Streams built using Tekton pipelines

→ OKD CentOS **S**treams **C**ore**OS** ('SCOS' for short) (a real upstream for OCP on RHEL9)

https://www.okd.io/blog/2022-10-25-OKD-Streams-Building-the-Next-Generation-of-OKD-together/

OKD Virtualization

- Enhance support for Tekton pipelines
- Additional metrics for monitoring and alerting
- ARM support
- Backup/Restore
- Non-privileged controller
- Real time virtual machine
- ....

OKD Working Group

Website
okd.io

Twitter
twitter.com/okd_io

Slack
#openshift-dev on kubernetes.slack.com

Google Group
groups.google.com/forum/#!forum/okd-wg

Bi-weekly Video Conference Meetings
apps.fedoraproject.org/calendar/okd

Repositories
github.com/openshift/community
github.com/openshift/okd

OKD Virtualization SIG

OKD
Virtualization

Reddit
www.reddit.com/r/OKD_Virtualization

Twitter
twitter.com/OKD_Virt_SIG

Slack
#virtualization on kubernetes.slack.com

Website
okd-virtualization.github.io

GitHub
github.com/okd-virtualization

Thank you

https://www.youtube.com/playlist?list=PLaR6Rq6Z4Iqc3WjZB-rUTPru8RKyOCnBo

https://twitter.com/okd_io

https://twitter.com/OKD_Virt_SIG

https://www.okd.io

https://okd-virtualization.github.io/