

Red Hat

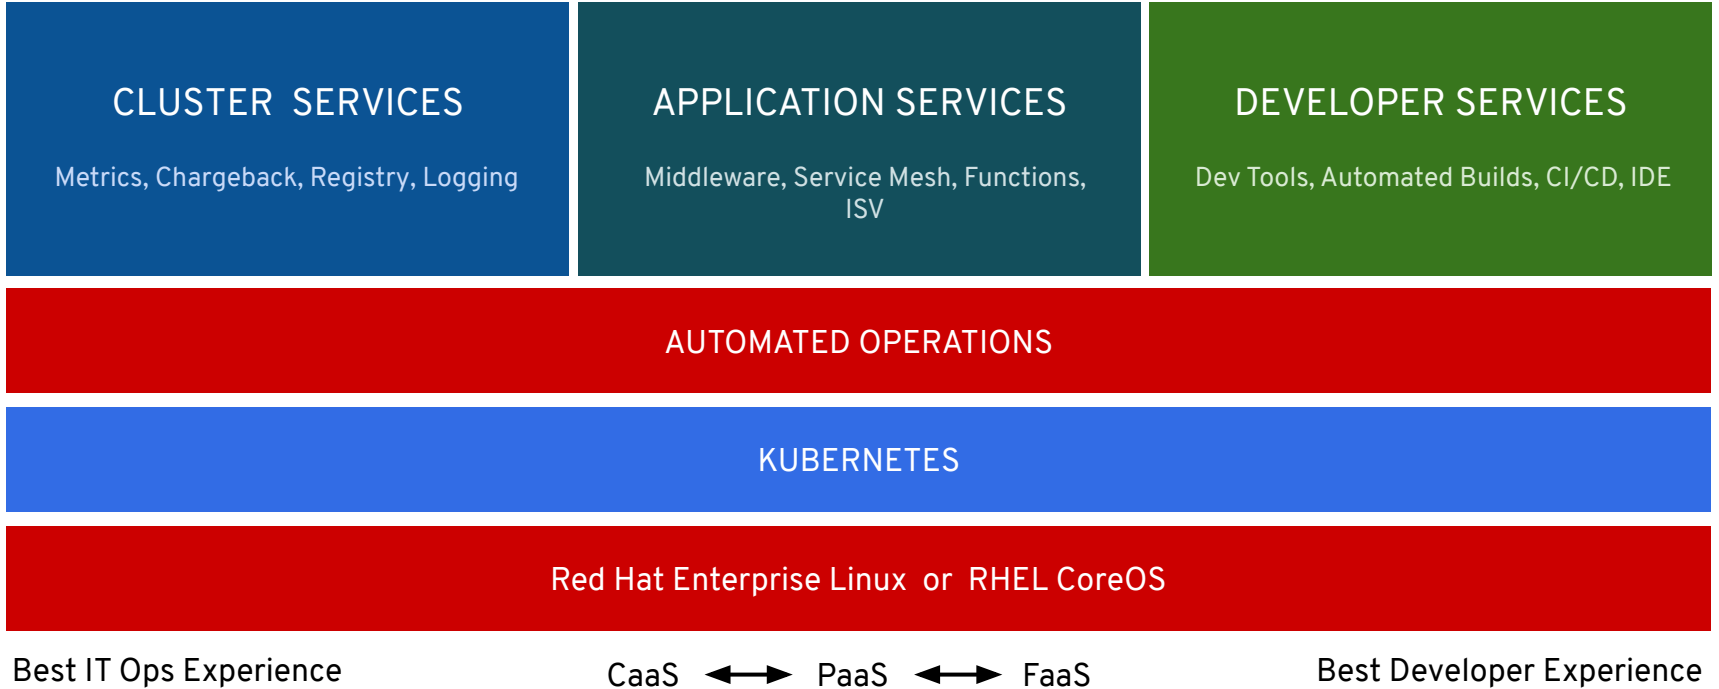


OpenShift 4 for Ops

Nick Schuetz
Principal Specialist Solution
Architect



OpenShift 4 Platform



OpenShift 4 Themes

DAY 2 OPERATIONS



Integrate CoreOS technology for a better install, re-config and upgrade experience.

Bring over-the-air upgrades to the platform.

IMMUTABLE INFRASTRUCTURE



Introduce Red Hat CoreOS as an immutable OS option. Enhance “infrastructure as code” throughout the platform.

OPERATOR FRAMEWORK




Provide tools, guidance and automation for customers and partners to deliver smart software on top of OpenShift


OpenShift 4.1 Workstreams Lifecycle



Installer + bootstrapping
Autoscale out of the box
MachineSet node pools




Red Hat Enterprise Linux CoreOS
Discourage SSH/node mutation
Ignition for Machine config



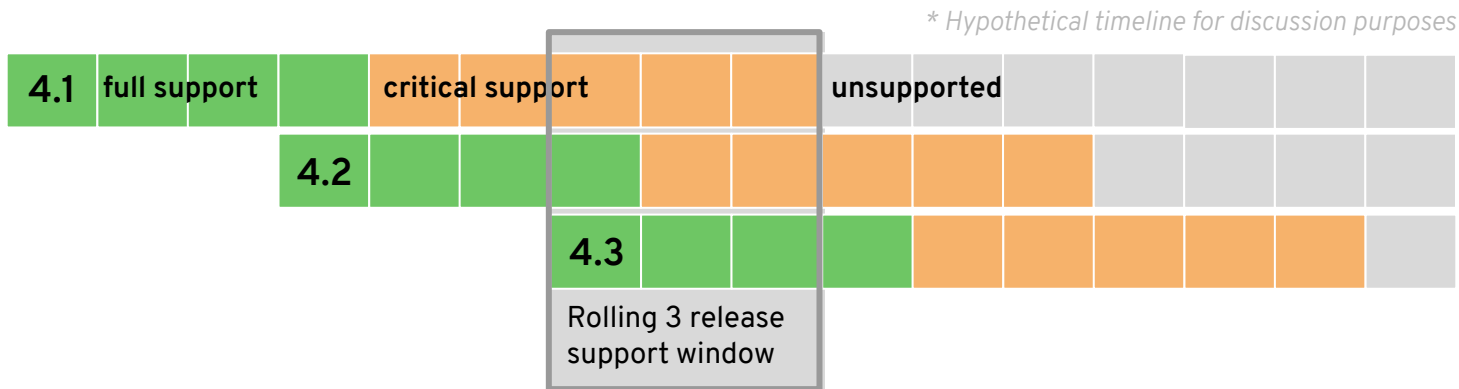
SDK & testing tools
OperatorHub for discovery
OLM delivers upper stack services

The New Platform Boundary

OpenShift 4 is aware of the entire infrastructure and brings the Operating System under management



OpenShift 4 Lifecycle



New model

Release based, not date based. Rolling three release window for support.

The overall 4 series will be supported for at least three years


- Minimum two years full support (likely more)
- One year maintenance past the end of full support

EUS release planned

Supported for 14 months of critical bug and critical security fixes instead of the normal 5 months. If you stay on the EUS for its entire life, you must use the application migration tooling to move to a new cluster

OpenShift 4 Upgrades

** Hypothetical timeline for discussion purposes*



OTA Upgrades

Works between two minor releases in a serial manner.

Happy path = migrate through each version

On a regular cadence, migrate to the next supported version.

Optional path = migration tooling

If you fall more than two releases behind, you must use the application migration tooling to move to a new cluster.

Current minor release

Full support for all bugs and security issues
1 month full support overlap with next release to aid migrations

Previous minor release

Fixes for critical bugs and security issues for 5 months

Installation Experiences

OPENSIFT CONTAINER PLATFORM

Full Stack Automated

Simplified opinionated “Best Practices” for cluster provisioning

Fully automated installation and updates including host container OS.



Pre-existing Infrastructure

Customer managed resources & infrastructure provisioning

Plug into existing DNS and security boundaries



HOSTED OPENSIFT

Azure Red Hat OpenShift

Deploy directly from the Azure console. Jointly managed by Red Hat and Microsoft Azure engineers.

OpenShift Dedicated

Get a powerful cluster, fully Managed by Red Hat engineers and support.

OpenShift 4 Supported Providers*

Full Stack Automated



Pre-existing Infrastructure




- * Requires Internet connectivity; support for cluster-wide proxy & disconnected installation/updating tentatively planned for 4.2.
- ** Automated Installer for Azure, OSP, GCP tentatively planned for 4.2.
- *** Azure UPI is currently in Dev Preview.

Full Stack Automated Deployments

Day 1: openshift-install - Day 2: Operators

- User managed
- Operator managed



Generally Available

Full Stack Automated Deployments

Simplified Cluster Creation

Designed to easily provision a “best practices” OpenShift cluster

- New CLI-based installer with interactive guided workflow that allows for customization at each step
- Installer takes care of provisioning the underlying Infrastructure significantly reducing deployment complexity
- Leverages RHEL CoreOS for all node types enabling full stack automation of installation and updates of both platform and host OS content

Faster Install

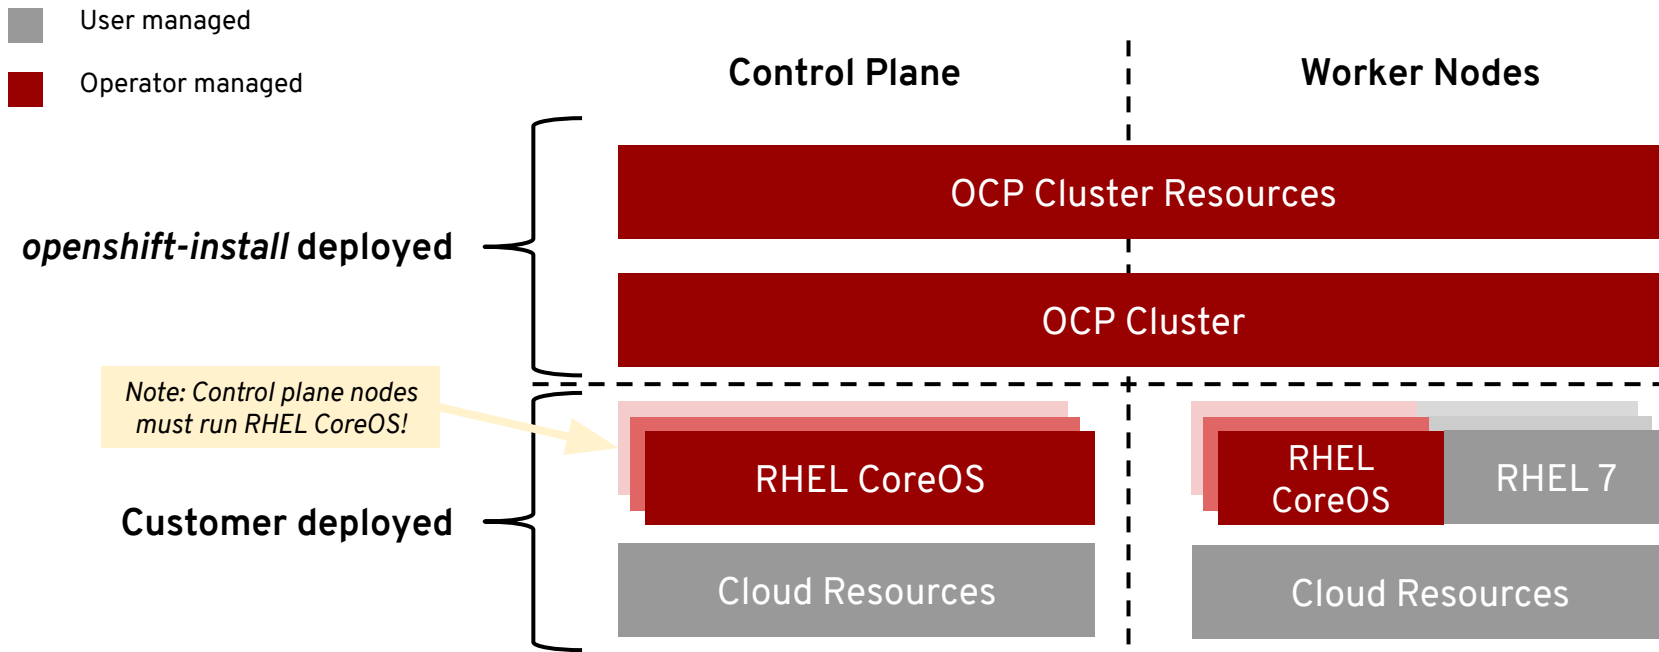
The installer typically finishes within 30 minutes

- Only minimal user input needed with all non-essential install config options now handled by component operator CRD's
- 4.1 provides support for AWS deployments with additional provider support planned in future releases
- [See the OpenShift documentation for more details](#)

```
$ ./openshift-install --dir ./demo create cluster
? SSH Public Key /Users/demo/.ssh/id_rsa.pub
? Platform aws
? Region us-west-2
? Base Domain example.com
? Cluster Name demo
? Pull Secret [?] for help
*****
INFO Creating cluster...
INFO Waiting up to 30m0s for the Kubernetes API...
INFO API v1.11.0+c69f926354 up
INFO Waiting up to 30m0s for the bootstrap-complete event...
INFO Destroying the bootstrap resources...
INFO Waiting up to 10m0s for the openshift-console route to be created...
INFO Install complete!
INFO Run 'export KUBECONFIG=<your working directory>/auth/kubeconfig' to
manage the cluster with 'oc', the OpenShift CLI.
INFO The cluster is ready when 'oc login -u kubeadmin -p <provided>'
succeeds (wait a few minutes).
INFO Access the OpenShift web-console here:
https://console-openshift-console.apps.demo.example.com
INFO Login to the console with user: kubeadmin, password: <provided>
```

Deploying to Pre-existing Infrastructure

Day 1: openshift-install - Day 2: Operators + Customer Managed Infra & Workers



Deploying to Pre-existing Infrastructure

Customized OpenShift Deployments

Enables OpenShift to be deployed to user managed resources and pre-existing infrastructure.

- Customers are responsible for provisioning all infrastructure objects including networks, load balancers, DNS, hardware/VMs and performing host OS installation
- Deployments can be performed both on-premise and to the public cloud
- OpenShift installer handles generating cluster assets (such as node ignition configs and kubeconfig) and aids with cluster bring-up by monitoring for bootstrap-complete and cluster-ready events
- While RHEL CoreOS is mandatory for the control plane, either RHEL CoreOS or RHEL 7 can be used for the worker/infra nodes
- Node auto-scaling can be setup for providers with OpenShift Machine API support
- [See the OpenShift documentation for more details](#)

```
$ cat ./demo/install-config.yaml
apiVersion: v1
baseDomain: example.com
compute:
- name: worker
  replicas: 0
controlPlane:
  name: master
...

$ ./openshift-install --dir ./demo create ignition-config
INFO Consuming "Install Config" from target directory

$ ./openshift-install --dir ./demo wait-for bootstrap-complete
INFO Waiting up to 30m0s for the Kubernetes API at
https://api.demo.example.com:6443...
INFO API v1.11.0+c69f926354 up
INFO Waiting up to 30m0s for the bootstrap-complete event...
$ ./openshift-install --dir ./demo wait-for cluster-ready

INFO Waiting up to 30m0s for the cluster at
https://api.demo.example.com:6443 to initialize...
INFO Install complete!
```

OpenShift Bootstrap Process: Self-Managed Kubernetes

How to boot a self-managed cluster:

- OpenShift 4 is unique in that management extends all the way down to the operating system
- Every machine boots with a configuration that references resources hosted in the cluster it joins enabling cluster to manage itself
- Downside is that every machine looking to join the cluster is waiting on the cluster to be created
- Dependency loop is broken using a bootstrap machine, which acts as a temporary control plane whose sole purpose is bringing up the permanent control plane nodes
- Permanent control plane nodes get booted and join the cluster leveraging the control plane on the bootstrap machine
- Once the pivot to the permanent control plane takes place, the remaining worker nodes can be booted and join the cluster

Bootstrapping process step by step:

1. Bootstrap machine boots and starts hosting the remote resources required for master machines to boot.
2. Master machines fetch the remote resources from the bootstrap machine and finish booting.
3. Master machines use the bootstrap node to form an etcd cluster.
4. Bootstrap node starts a temporary Kubernetes control plane using the newly-created etcd cluster.
5. Temporary control plane schedules the production control plane to the master machines.
6. Temporary control plane shuts down, yielding to the production control plane.
7. Bootstrap node injects OpenShift-specific components into the newly formed control plane.
8. Installer then tears down the bootstrap node or if user-provisioned, this needs to be performed by the administrator.

Comparison between deployments methods

	Full Stack Automation	Pre-existing Infrastructure
Build Network	Installer	User
Setup Load Balancers	Installer	User
Configure DNS	Installer	User
Hardware/VM Provisioning	Installer	User
OS Installation	Installer	User
Generate Ignition Configs	Installer	Installer
OS Support	Installer: RHEL CoreOS	User: RHEL CoreOS + RHEL 7
Node Provisioning / Autoscaling	Yes	Only for providers with OpenShift Machine API support
Customization & Provider Support	Best Practices: AWS	Yes: AWS, Bare Metal, & VMware

Generally Available

RHEL 7 Worker/Infra Nodes

Add RHEL 7.6 machines with Ansible

Use openshift-ansible to prepare, configure and join your RHEL 7 nodes to the cluster. After you have a functional control plane, nodes can be added to the cluster with the `scaleup` playbook.

- Pending certificates signing request (CSRs) for each RHEL machine must be approved before joining cluster
- [See the OpenShift documentation for more details](#)

Mixed clusters of RHEL 7 and RHEL CoreOS are ok

RHEL CoreOS is mandatory for the control plane, but mixed clusters of RHEL 7.6 and RHEL CoreOS are supported for any other node pools.

RHEL 8 is not yet supported for worker/infra nodes

Support will come in a later 4.x release

Upgrading OpenShift node components

With the upgrade playbook, RHEL 7 OpenShift components can be upgraded. Optionally, pre/post hooks are also available for performing custom tasks.

RHEL admins are responsible for:

Keeping host inventory up to date

Refresh your list of hosts before an upgrade, to make sure no machines are missed

Defining Ansible hooks

Run playbooks that will cordon/uncordon machines, along with any pre/post upgrade actions

Updating RHEL RPM content

Security, performance and regular updates from Red Hat

Partitioning disks

Configure, maintain and health check your disks

Configuring network interfaces

Configure, secure and maintain settings within your data center's specifications

Adding & Updating RHEL 7 Worker Nodes

System Requirements

- OpenShift 4 cluster deployed to pre-existing infrastructure (not supported on installer provisioned clusters)
- Latest release of RHEL 7 Server installed on every machine joining the cluster
 - Ansible >= 2.7.8, OpenShift Client (oc)
 - Necessary entitlements and YUM repo access pre-configured for every machine

Create an Ansible Inventory

- Inventory file with the appropriate groups and variables defined
 - An example inventory can be found in [inventory/hosts.example](#)
- Required variables include:
 - `openshift_kubeconfig_path`: Path to the kubeconfig for the cluster
 - `openshift_pull_secret_path`: Path to the pull secret to the image registry

Run RHEL 7 node 'scaleup' playbook

```
$ cd openshift-ansible; ansible-playbook -i inventory/hosts playbooks/scaleup.yml
```

- Pending certificates signing request (CSRs) for each RHEL machine must be approved before joining cluster

```
$ oc adm certificate approve <csr_name>
```

Upgrading RHEL 7 OpenShift node components

- Leverages upgrade section of Ansible Inventory to specify nodes
- Custom tasks can be performed during upgrades at different stages of the upgrade; refer to 'hooks' in the documentation for more information.

```
$ cd openshift-ansible; ansible-playbook -i inventory/hosts playbooks/upgrade.yml
```

Red Hat Enterprise Linux

RED HAT® ENTERPRISE LINUX®

General Purpose OS

BENEFITS

- 10+ year enterprise life cycle
- Industry standard security
- High performance on any infrastructure
- Customizable and compatible with wide ecosystem of partner solutions

WHEN TO USE

When customization and integration with additional solutions is required

RED HAT® ENTERPRISE LINUX CoreOS

Immutable container host

- Self-managing, over-the-air updates
- Immutable and tightly integrated with OpenShift
- Host isolation is enforced via Containers
- Optimized performance on popular infrastructure

When cloud-native, hands-free operations are a top priority

Red Hat Enterprise Linux CoreOS

4.1 Image Availability:

- Amazon: AMIs
- vSphere: OVA
- Bare Metal: UEFI & BIOS

Installation Requirements:

- RHCOS image + ignition config (installer generated)

RHCOS Details

- RHEL 8 bits (4.18 kernel)
- Includes all packages required for OpenShift
- Over-The-Air updates encompass OCP & RHCOS
- Transactional host updates
- Read-only OS binaries
- Preconfigured for most environments

Bare Metal Installer (ISO or PXE):

```
coreos.inst=yes
coreos.inst.install_dev=sda
coreos.inst.image_url=http://10.10.10.1/rhcos-metal-uefi.raw.gz
coreos.inst.ignition_url=http://10.10.10.1/master.ign
```

Immutable Operating System

Red Hat Enterprise Linux CoreOS is versioned with OpenShift

CoreOS is tested and shipped in conjunction with the platform. Red Hat runs thousands of tests against these configurations.


Red Hat Enterprise Linux CoreOS is managed by the cluster

The Operating system is operated as part of the cluster, with the config for components managed by Machine Config Operator:

- CRI-O config
- Kubelet config
- Authorized registries
- SSH config


RHEL CoreOS admins are responsible for:

Nothing. 😊🙌



CRI-O Support in OpenShift

CRI-O tracks and versions identical to Kubernetes, simplifying support permutations



Generally Available

Machine Config Operator (MCO)

Example with CRI-O:

- Simplified management across entire cluster
- Container Runtime Config (CRC) is specialized Machine Config (MC) for CRI-O
- Container Runtime Config for CRI-O exposes configuration knobs

Leverages

- Custom Resource Definitions
- Machine Config Pools
- Machine Config Pool Selector
- Machine Configs
- Container Runtime Config (specialized MC)

```
apiVersion: machineconfiguration.openshift.io/v1
kind: ContainerRuntimeConfig
metadata:
  name: set-log-and-pid
spec:
  machineConfigPoolSelector:
    matchLabels:
      debug-crio: config-log-and-pid
  containerRuntimeConfig:
    pidsLimit: 2048
    logLevel: debug
```

```
$ oc get ContainerRuntimeConfig
```

NAME	AGE
set-log-and-pid	22h

Discovering Configuration Options

Cluster config objects are the human/machine APIs

These can be discovered with the commands to the right.

Configuration set via this mechanism will be carried forward through cluster updates.

Configuration Resources

Apiserver	Authentication
Build	Clusterversion
Console	Dns
Featuregate	Image
Infrastructure	Ingress
Network	Oauth
Project	Scheduler

```
$ oc api-resources --api-group config.openshift.io -o
name | cut -d . -f 1
apiservers
authentications
...
scheduler
```

```
$ oc explain --api-version=config.openshift.io/v1
scheduler.spec
KIND:      Scheduler
VERSION:   config.openshift.io/v1

FIELDS:
  defaultNodeSelector      <string>
    defaultNodeSelector helps set the cluster-wide
    default
    node selector to restrict pod placement to
    specific
```

```
$ oc explain --api-version=config.openshift.io/v1
scheduler.spec.policy
```


Graphical Re-configuration

Global Configuration


You complete most of the cluster configuration and customization after you deploy your OpenShift Container Platform cluster.

Change via Cluster Settings screen

Once you have discovered your desired settings (prev. slide), changes can be made via Console or CLI.

Operators apply these updates

One or more Operators are responsible for propagating these settings through the infrastructure



Network Configuration

Example #1: Operator-Assisted Ingress Ctrlr “Sharding”

In 4.1, the way you create a router to work with a shard is different (API call versus ‘oc adm’ command). A simple config (example to right) acted upon by the ingress operator automatically integrates sharding with the external (front-end) DNS/LB configured at install-time,.

```
apiVersion: operator.openshift.io/v1
kind: IngressController
metadata:
  namespace: openshift-ingress-operator
  name: internal-apps
spec:
  domain: internal-apps.dmace.devcluster.openshift.com
  routeSelector:
    matchLabels:
      environment: internal
```

Example #2: Create a Second Router

Ingress controller configuration is now a first-class object, meaning additional Ingress controllers can be created by making multiple Ingress objects. This is the preferred method for giving teams their own subdomains, replacing the ‘oc adm’ method (see right).

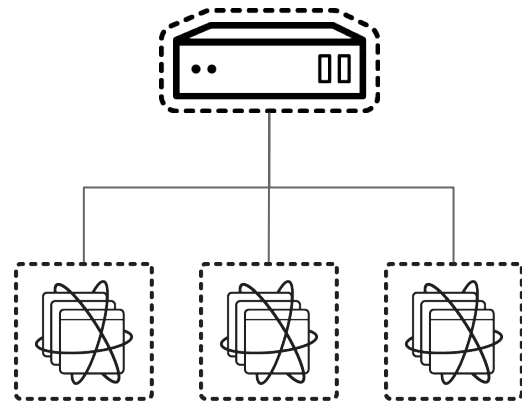
```
$ cat <<EOF | oc create -f -
apiVersion: operator.openshift.io/v1
kind: IngressController
metadata:
  namespace: openshift-ingress-operator
  name: finance-apps
spec:
  domain: finance-apps.openshift.example.com
EOF
```

HAProxy Optimization

HAProxy Large Performance Increase

OCP 4.1 deploys 2 router replicas by default (improving the observed performance $\sim 2x$), and sets 4 router threads by default (also improving observed performance another $\sim 2x$).

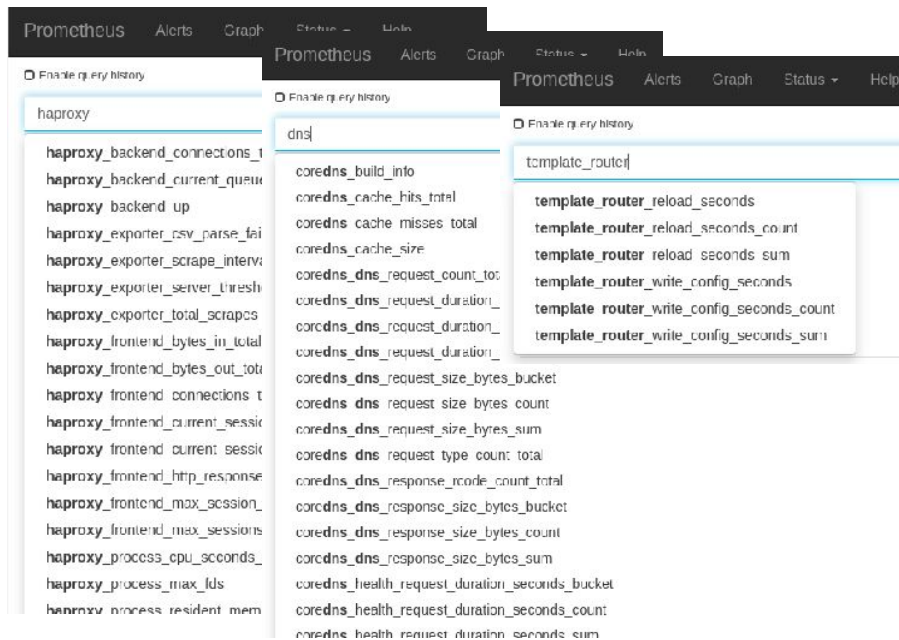
We increased these defaults for HAProxy so that the cluster can serve more routes.



Networking Metrics

Networking Metrics Visibility in Prometheus

Networking metrics (HAproxy/Router, CoreDNS, etc.) for consumption by Prometheus are GA at 4.1.



Networking Re-configuration

Networking Advanced Settings

These are the OpenShift SDN settings that can be tweaked at install-time:

- Mode: NetworkPolicy, Multitenant, Subnet
- VXLAN Port Number
- MTU (autodetected, once)
- External OpenVSwitch

[How to Modify Advanced Network Configuration Parameters](#)

```
spec:  
  defaultNetwork:  
    type: OpenShiftSDN  
  openshiftSDNConfig:  
    mode: NetworkPolicy  
    vxlanPort: 4789  
    mtu: 1450  
    useExternalOpenvswitch: false
```

NOTE: Most network settings cannot be changed safely and affect the entire cluster. The operator will prevent unsafe changes. If you need to force a change to a *non-production* cluster, see the operator README for the command, but a cluster re-install is likely to be the better choice.

kube-proxy Re-configuration

kube-proxy Advanced Settings

These are the kube-proxy settings that can be tweaked at install-time:

- iptablesSyncPeriod
- bindAddress
- proxyArguments (a list of kube-proxy command-line flags)

[How to Modify Advanced Network Configuration Parameters](#)

```
spec:  
  kubeProxyConfig:  
    iptablesSyncPeriod: 30s  
    bindAddress: 0.0.0.0  
    proxyArguments:  
      iptables-min-sync-period: ["30s"]
```

NOTE: Most network settings cannot be changed safely and affect the entire cluster. The operator will prevent unsafe changes. If you need to force a change to a *non-production* cluster, see the operator README for the command, but a cluster re-install is likely to be the better choice.

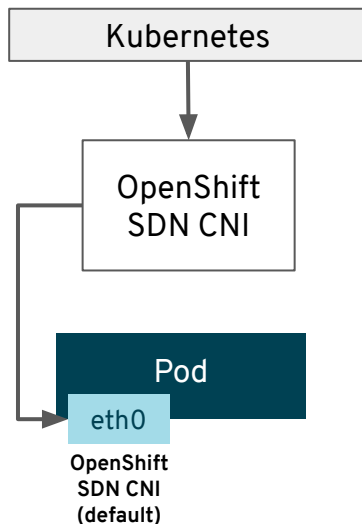
Networking Plug-ins

Multus Enables Multiple Networks & New Functionality to Existing Networking

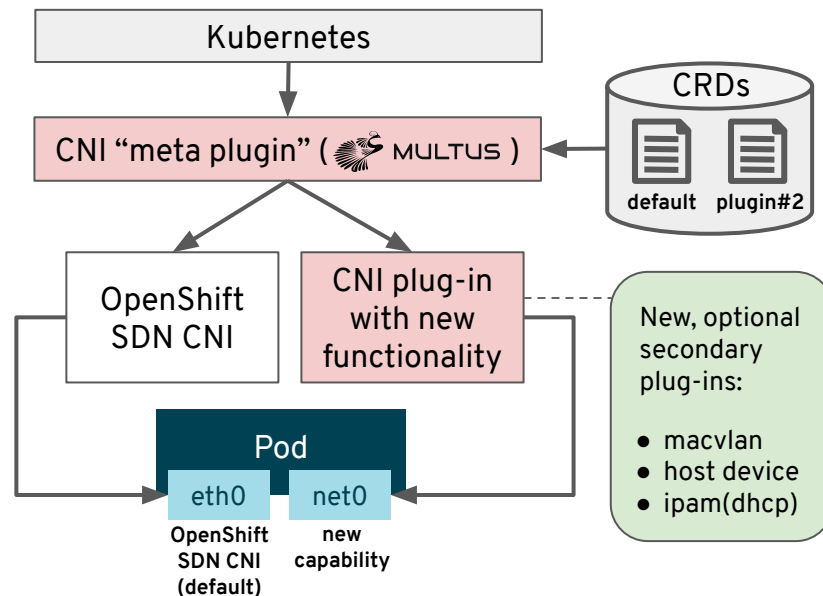
The Multus CNI “meta plugin” for Kubernetes enables one to create multiple network interfaces per pod, and assign a CNI plugin to each interface created.

1. Create pod annotation(s) to call out a list of intended network attachments...
2. ...each pointing to CNI network configurations packed inside CRD objects

3.x Capability...



4.x Capability...



New, optional secondary plug-ins:

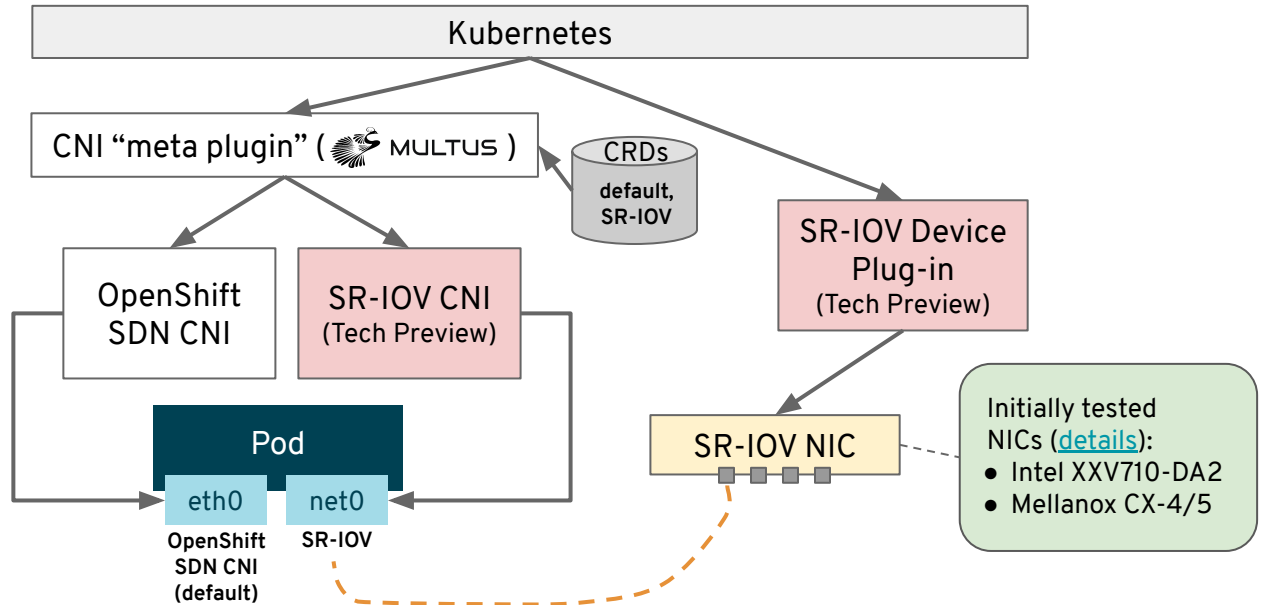
- macvlan
- host device
- ipam(dhcp)

High-Performance Networking (Tech Preview)

Approach Line-Rate Performance to the Pod

OCP 4.1 includes the capability to use specific SR-IOV hardware on cluster nodes to bring near-line rate network performance to cluster pods.

[Configuring SR-IOV](#)



Storage

Storage Focus

- Cluster Storage Operator
 - Sets up the default storage class
 - Looks through cloud provider and sets up the correct storage class
- Drivers themselves remain in-tree for now.
- CSI coming in the future.

Supported	Dev Preview
AWS EBS	Snapshot*
VMware vSphere Disk	EFS*
NFS	Local Volume*
iSCSI	Raw Block
Fibre Channel	
HostPath	

* via external provisioner

Configuring an Identity Provider

The Cluster Authentication Operator

- Use the *cluster-authentication-operator* to configure an Identity Provider. The configuration is stored in the *oauth/cluster* custom resource object inside the cluster.
- Once that's done, you may choose to remove kubeadmin (warning: there's no way to add it back).
- All the identity providers supported in 3.11 are supported in 4.1: LDAP, GitHub, GitHub Enterprise, GitLab, Google; OpenID Connect, HTTP request headers (for SSO), Keystone, Basic authentication.
- For more information:
[Understanding identity provider configuration cluster-authentication-operator](#)

Sample identity provider CR

```
apiVersion: config.openshift.io/v1
kind: OAuth
metadata:
  name: cluster
spec:
  identityProviders:
  - name: my_identity_provider 1
    mappingMethod: claim 2
    type: HTPasswd
    htpasswd:
      fileData:
        name: htpass-secret 3
```

- 1 This provider name is prefixed to provider user names to form an identity name.
- 2 Controls how mappings are established between this provider's identities and user objects.
- 3 An existing secret containing a file generated using [htpasswd](#).

Service Certificates and the Service CA

The Service CA Operator

- The Service Certificate Authority (CA) can sign server certificates for services running in the OpenShift cluster.
- The **service-ca operator** manages 3 controllers:
service-ca controller
configmap-cabundle-injector controller
apiservice-cabundle-injector controller
- The Service CA is valid for one year after installation. Manual steps to refresh the Service CA are documented. Automation is targeted for 4.2.
- For more information:

[Service Serving Certificate Secrets](#)
[openshift-service-ca-operator](#)

```
$ oc create configmap foobar --from-literal=key1=foo
configmap/foobar created
$ oc get configmap/foobar -o yaml
apiVersion: v1
data:
  key1: foo
kind: ConfigMap
metadata:
  creationTimestamp: 2018-09-11T23:44:56Z
  name: foobar
  namespace: myproject
  resourceVersion: "56490"
  selfLink: /api/v1/namespaces/myproject/configmaps/foobar
  uid: afee501b-b61c-11e8-833b-c85b762603b0
$ oc annotate configmap foobar service.beta.openshift.io/inject-cabundle="true"
configmap/foobar annotated
$ oc get configmap/foobar -o yaml
apiVersion: v1
data:
  key1: foo
  service-ca.crt: |
    -----BEGIN CERTIFICATE-----
    MIIDCjCCAfKgAwIBAgIBATANBgkqhkiG9w0BAQsFADA2MTQwMgYDVQDDctvcGvu
    c2hpZnQtY2VydmluZS1zZXJ2YW5nLXNpZ25lcAxtNTM2Njk1NTIxMB4XDTE4MDkx
    MTE5NTIwMVoXDTIzMDkxMDE5NTIwMlowNjE0MDIGA1UEAwrb3BlnNoaWZ0LXNl
    cnZpY2Utc2VydmluZy1zaWduZXJAMTUzNjY5NTUyMTCCASiWdQYJKoZIhvcNAQEB
    BQADggEPADCCAQoCggEBANP9Asc657SkWVP0ohmMlRlXQirL7taaarmM5L3/pNqeo
```

Custom Certificates for External Endpoints

Custom certificates for Ingress Controllers

- In OpenShift 4 adding custom certificates for external endpoints is a post-installation task.
- Cluster ingress is managed by the Ingress Operator which configures Ingress Controllers to route traffic as specified by OpenShift [Route](#) and Kubernetes [Ingress](#) resources.
- Updating the certificate for the ingress controller(s) covers the Web console, the API endpoint, the internal Registry and custom applications. For the Registry, you can choose to [define a separate route with certificate](#).
- For example: [Requesting and Installing Let's Encrypt Certificates for OpenShift 4](#)

Requesting and Installing Let's Encrypt Certificates for OpenShift 4

The Router expects the certificates in a Secret. This secret needs to be created in the project openshift-ingress.

1. Use the following command to create the secret – and if you have existing certificates, make sure to provide the path to your certificates instead.

```
oc create secret tls router-certs --cert=${CERTDIR}/fullchain.pem --key=${CERTDIR}/key.pe
```

2. Now update the Custom Resource for your router. The default custom resource is of type IngressController, is named default and is located in the openshift-ingress-operator project. Note that this project is different from where you created the secret earlier.

```
oc patch ingresscontroller default -n openshift-ingress-operator --type=merge --patch='{'
```

3. This is all you need to do. After you update the IngressController object the OpenShift ingress operator notices that the custom resource has changed and therefore re-deploys the router.

You now have proper certificates on the router – and this includes custom applications, the Web Console for your OpenShift Cluster and the API Endpoint.

Registry Configuration

Image Registry Operator Configuration Parameters

Several parameters support changing the registry configuration. Those are defined in the `configs.imageregistry.operator.openshift.io` resource

Additional configuration by ConfigMap and Secret resources located within the `openshift-image-registry` namespace

Documentation:

https://docs.openshift.com/container-platform/4.1/registry/configuring-registry-operator.html#registry-operator-configuration-resource-overview_configuring-registry-operator

Image Registry Operator configuration parameters

The `configs.imageregistry.operator.openshift.io` resource offers the following configuration parameters.

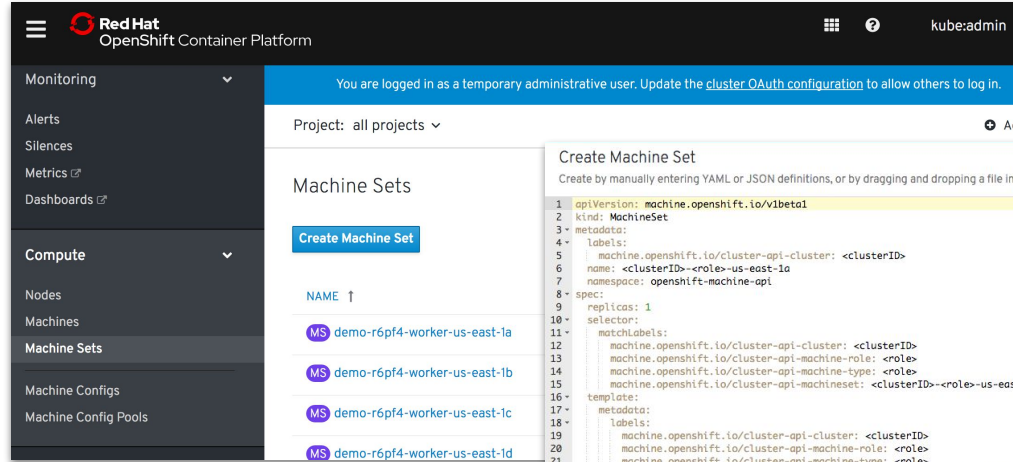
Parameter	Description
ManagementState	Managed: The Operator updates the registry as configuration resources are updated. Unmanaged: The Operator ignores changes to the configuration resources.
Removed	The Operator removes the registry instance and tear down any storage that the Operator provisioned.
Logging	Sets LogLevel of the registry instance.
HTTPSecret	Value needed by the registry to secure uploads, generated by default.
Proxy	Defines the Proxy to be used when calling master API and upstream registries.
Storage	StorageType: Details for configuring registry storage, for example S3 bucket coordinates. Normally configured by default.
Requests	API Request Limit details. Controls how many parallel requests a given registry instance will handle before queuing additional requests.
DefaultRoute	Determines whether or not an external route is defined using the default hostname. If enabled, the route uses re-encrypt encryption. Defaults to false.
Routes	Array of additional routes to create. You provide the hostname and certificate for the route.
Replicas	Replica count for the registry.

Infrastructure MachineSets

MachineSets that are purpose built for Infrastructure Services

- Elasticsearch, Prometheus, Router, Registry
- Out-of-box installer does not create a MachineSets dedicated for Infra services
- Create a MachineSet via console or cli and label them with desired roles
- Redeploy Infra Services with nodeSelector set to the designated role

[Documentation: Creating Infrastructure MachineSets](#)



Moving components to dedicated Node pools

Create a new Node pool

First, make a new MachineSet with a template that contains a custom label, like `role=logging`. Optionally customize the resources or security groups on this set of Nodes.

[Documentation: Roles in OpenShift](#)

Update Node selectors

Change the node selectors to your new Node pool labels. This process uses different CRDs based on the component:

- [Cluster Monitoring](#) (docs)
- [Router](#) (lab)
- [Registry](#) (lab)
- [Logging](#) (docs)

```
# new Node pools added
$ oc get nodes
NAME
ip-10-0-132-151.us-east-2.compute.internal worker
Ip-10-0-137-86.us-east-2.compute.internal worker
ip-10-0-138-38.us-east-2.compute.internal worker
ip-10-0-139-204.us-east-2.compute.internal worker
ip-10-0-139-249.us-east-2.compute.internal master
ip-10-0-144-70.us-east-2.compute.internal master
ip-10-0-145-199.us-east-2.compute.internal master
internal logging
internal logging
internal monitoring
internal monitoring
internal router
internal router
```

NAME	NAMESPACE
demo-r6pf4-worker-us-east-1a	openshift-machine-api
demo-r6pf4-worker-us-east-1b	openshift-machine-api
demo-r6pf4-worker-us-east-1c	openshift-machine-api
demo-r6pf4-worker-us-east-1d	openshift-machine-api

Generally Available

Additional Build Configurations

Default build configurations for buildconfigs


- Additional CAs to be trusted for image pull/push
- Proxy setting for image pull/push and source download
- Proxy settings for git commands
- Environment variables to set on all builds
- Docker labels to apply to resulting images
- Build resource requirements
- Default values to override on builds even if user has provided values on the buildconfig

```
apiVersion: config.openshift.io/v1
kind: Build
metadata:
  name: cluster
spec:
  additionalTrustedCA:
    name: trustedCAsConfigMap
  buildDefaults:
    defaultProxy: # http, https and no proxy
    gitProxy: # http, https and no proxy
    env: # key-values
    imageLabels:
    resources:
      limits: # cpu, memory
      requests: # cpu, memory
  buildOverrides:
    imageLabels:
    nodeSelector:
    tolerations:
```

```
oc edit build.config.openshift.io/cluster
```


Cloud-like Simplicity, Everywhere

Full-stack automated operations across any on-premises,
cloud, or hybrid infrastructure



OpenShift Cluster Manager on cloud.redhat.com

Automatic registration of OpenShift clusters

View cluster versions and capacity in one place, no matter what infrastructure you are running on. Integrated with RHSM.

OpenShift Dedicated cluster management

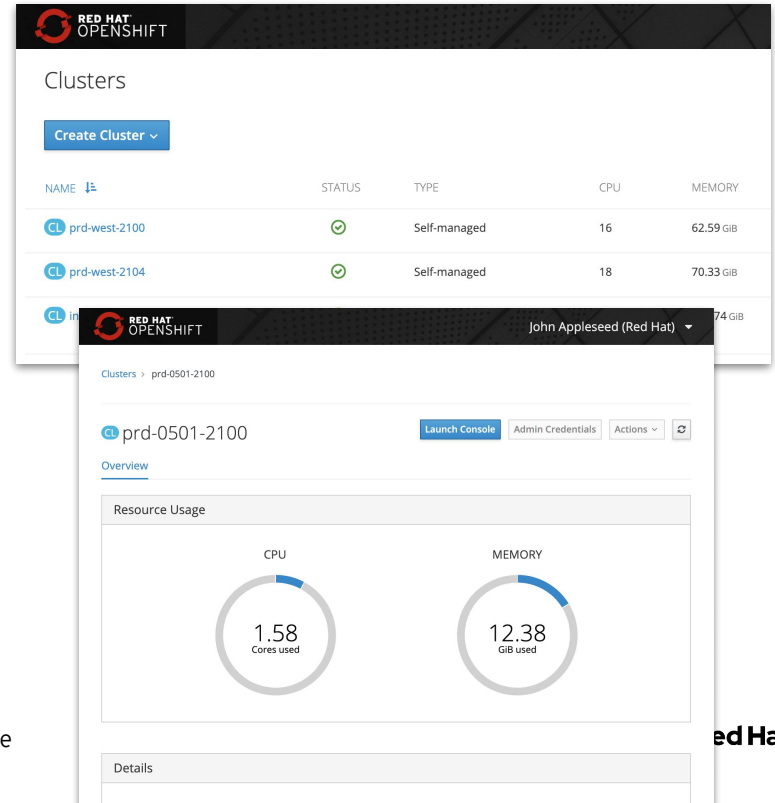
Self-service cluster deployment, scaling, and management for OpenShift Dedicated coming soon.

Azure Red Hat OpenShift

Information about these clusters will be coming at a later date.

Hosted in the United States

Other geographies may come later. You can [opt-out](#) too.



Generally Available

OpenShift Subscription Management

Moves from node management to cluster management

Entitle clusters and not nodes. Nodes too dynamic. We do not block on usage. Requires telemeter Opt-In.

Dynamically adds and removes nodes

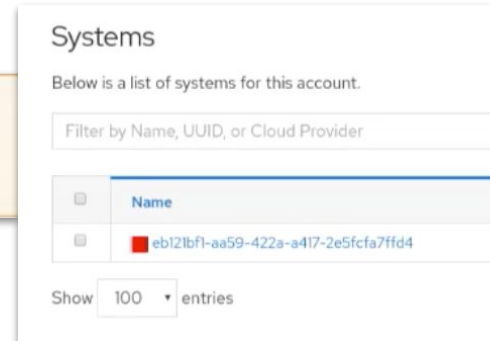
UHC will dynamically add and remove nodes from your subscription allocations to the cluster in 24 hour intervals. This will move to instantaneous across the next several releases.

Connected to the same backend as Subscription Portal and Satellite

Allocation numbers you see at cloud.redhat.com for OCP can be also seen on the subscription portal at access.redhat.com

Removes OCP Infrastructure from the count

UHC will figure out which pods are your OCP infra pods and subtract out their usage from your core count so you are not charged.



Generally Available

Automated Container Operations

Fully automated day-1 and day-2 operations

INSTALL	DEPLOY	HARDEN	OPERATE
AUTOMATED OPERATIONS			
Infra provisioning	Full-stack deployment	Secure defaults	Multi-cluster aware
Embedded OS	On-premises and cloud	Network isolation	Monitoring and alerts
	Unified experience	Audit and logs	Full-stack patch & upgrade
		Signing and policies	Zero downtime upgrades
			Vulnerability scanning

Smarter Software Updates

No downtime for well behaving apps

Applications with multiple replicas, using liveness probes, health checks and taints/tolerations

Node Pools with more than one worker and slack resources

Maintenance window for entire cluster


No need for separate windows for each component

Upgrade runs completely on the cluster

No more long running processes on a workstation

Constant health checking from each Operator

Operators are constantly looking for incompatibilities and issues that might arise



Rolling Machine Updates

Single-click updates

- RHEL CoreOS version & config
- Kubernetes core components
- OpenShift cluster components

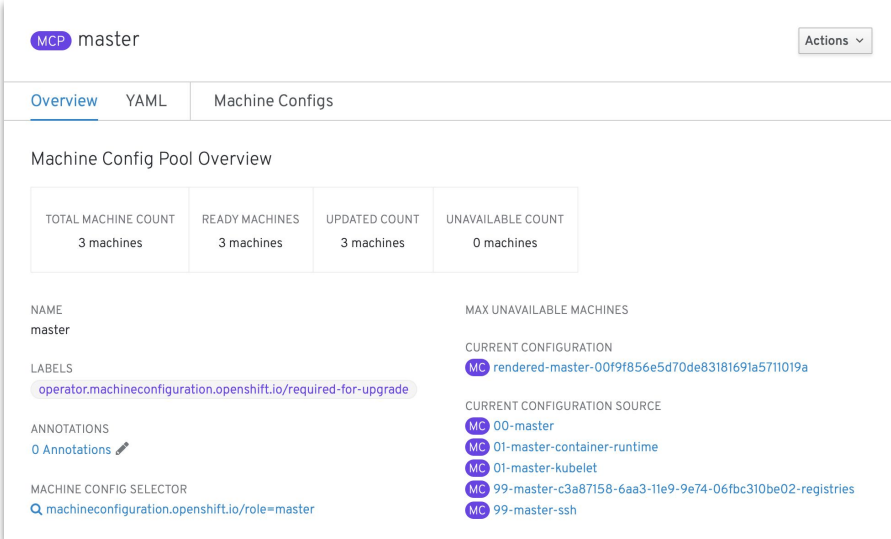
Configure how many machines can be unavailable

Set the “maxUnavailable” setting in the MachineConfigPool to maintain high availability while rolling out updates.

The default is 1.

Machine Config Operator (MCO) controls updates

This is a DaemonSet that runs on all Nodes in the cluster. When you upgrade with `oc adm upgrade`, the MCO executes these changes.



The screenshot displays the OpenShift console interface for the 'master' Machine Config Pool (MCP). The page is titled 'MCP master' and includes an 'Actions' dropdown menu. Below the title, there are tabs for 'Overview', 'YAML', and 'Machine Configs'. The 'Overview' tab is selected, showing a 'Machine Config Pool Overview' section with a table of statistics:

TOTAL MACHINE COUNT	READY MACHINES	UPDATED COUNT	UNAVAILABLE COUNT
3 machines	3 machines	3 machines	0 machines

Below the table, there are several sections of metadata:

- NAME:** master
- MAX UNAVAILABLE MACHINES:** (empty)
- LABELS:** operator.machineconfiguration.openshift.io/required-for-upgrade
- CURRENT CONFIGURATION:** MC rendered-master-00f9f856e5d70de83181691a5711019a
- ANNOTATIONS:** 0 Annotations
- CURRENT CONFIGURATION SOURCE:** MC 00-master, MC 01-master-container-runtime, MC 01-master-kubelet, MC 99-master-c3a87158-6aa3-11e9-9e74-06fbc310be02-registries, MC 99-master-ssh
- MACHINE CONFIG SELECTOR:** Q machineconfiguration.openshift.io/role=master

Pod Eviction Behavior

Eviction has been moved from Node Conditions to Taints

A new `DefaultTolerationSeconds` mutating admission plugin will add a 5 min eviction timeout unless specifically set by the Pod. This gives users more control vs the previous cluster-wide setting.

Don't tolerate the `node.kubernetes.io/unreachable` taint

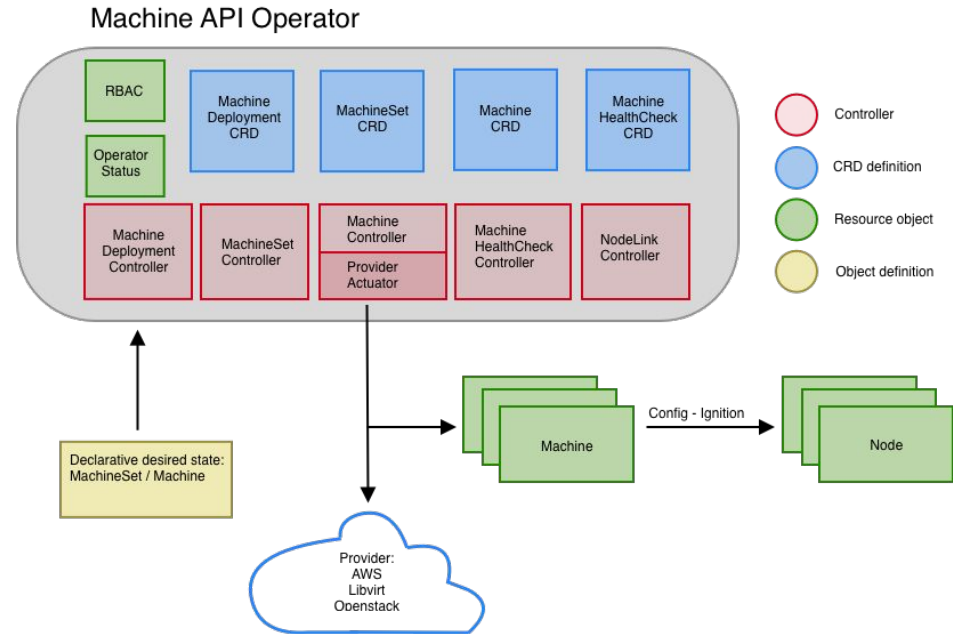
This taint is used by the cluster to evict and reschedule Nodes. If your Pod tolerates this, it will not be rescheduled on Node failure.

The image shows a screenshot of a Kubernetes pod configuration. On the left, a summary panel lists '2 Tolerations' with a red box around it. A red arrow points from this box to a detailed view of the pod's configuration on the right. The detailed view shows the following configuration:

```
198 secret:
199   secretName: default-token-4w4jp
200   defaultMode: 420
201 dnsPolicy: ClusterFirst
202 tolerations:
203 - key: node.kubernetes.io/not-ready
204   operator: Exists
205   effect: NoExecute
206   tolerationSeconds: 300
207 - key: node.kubernetes.io/unreachable
208   operator: Exists
209   effect: NoExecute
210   tolerationSeconds: 300
211 status:
212   phase: Running
213   conditions:
214 - type: Initialized
215   status: 'True'
216   lastProbeTime: null
```

Cloud API

- Provide a single view and control across multiple cluster types
- *Machine API*:
 - Set up definitions via CRDs
 - Machine: a node
 - MachineSet: think ReplicaSet
 - Actuators roll definitions across clusters
 - Nodes are drained before deletion
- *Cluster Autoscaler*: provide/remove additional nodes on demand
- AWS (4.1), Azure/GCP (target 4.2), VMWare (Future)



Cluster Monitoring

Cluster monitoring is installed by default

- Exposes resource metrics for Horizontal Pod Autoscaling (HPA) by default
 - HPA based on custom metric is tech preview
- No manual etcd monitoring configuration anymore
- New screens for managing Alerts & Silences
- More metrics available for troubleshooting purposes (e.g. HAproxy)
- Configuration via ConfigMaps and Secrets

Alerts | Alertmanager UI

Alerts help notify you when certain conditions in your environment are met. [Learn more about how alerts are](#)

12 Firing | 0 Silenced | 0 Pending | 77 Not Firing | [Select All Filters](#)

NAME ↑	STATE
AL CPUThrottlingHigh 39% throttling of CPU in namespace metering-demo for container tiller in pod metering-operator-5c9c754b85-19ds2.	Firing Since Apr 29, 11:52
AL CPUThrottlingHigh 28% throttling of CPU in namespace metering-demo for container reporting-operator in pod reporting-operator-6c666b88db-qvbb5.	Firing Since May 2, 6:47 a
AL CPUThrottlingHigh 81% throttling of CPU in namespace metering-demo for container metering-operator in pod metering-operator-5c9c754b85-19ds2.	Firing Since Apr 29, 11:52
AL KubeDeploymentReplicasMismatch Deployment openshift-operators/mongodb-enterprise-operator has not matched the expected number of replicas for longer than an hour.	Firing Since May 2, 1:34 p
AL KubePodCrashLooping Pod openshift-operators/mongodb-enterprise-operator-7b6954d84d-g69b4 (mongodb-enterprise-operator) is crashing 0.92 times / 1.5 minutes	Firing Since Apr 29, 2:52


Telemetry

Collects anonymized data from any OpenShift 4 cluster deployment

- Red Hat gains quality assurance with anonymous data reporting faults encountered during upgrade
- Show utilization of all your clusters at cloud.redhat.com
- Perform subscription management at cloud.redhat.com

Opt-out is only available for self-managed OpenShift clusters but we strongly discourage that as you will lose all of the features described above.

[Complete list of collected metrics](#)



Cluster Logging

Cluster Logging is lifecycle managed via Operator Lifecycle Management

- Install the Elasticsearch and Cluster Logging Operators from OperatorHub
- Create an instance of Cluster Logging. fluentd, Elasticsearch and Kibana (with Operators) are created
- Changing the out-of-box configuration:
 - CPU, memory requests and limits, PVC sizes etc can be changed by editing the Cluster Logging Operator YAML
- Direct Elasticsearch and Kibana Deployments to dedicated Nodes (recommended for production usage)

Create Operator Subscription

Keep your service up to date by selecting a channel and approval strategy. The strategy determines either manual or automatic updates.

Installation Mode *

All namespaces on the cluster
This mode is deprecated. Operator will be available in a single namespace only.

A specific namespace on the cluster
Operator will be available in a single namespace only.

Update Channel *

preview

Approval Strategy *


Automatic
 Manual

```
# configure via CRD
apiVersion: "logging.openshift.io/v1"
kind: "ClusterLogging"
metadata:
  name: "instance"
  namespace: "openshift-logging"
spec:
  managementState: "Managed"
  logStore:
    type: "elasticsearch"
    elasticsearch:
      nodeCount: 3
      resources:
        limits:
          cpu: 800m
          memory: 1Gi
          requests:
            cpu: 800m
            memory: 1Gi
      storage:
        storageClassName: gp2
        size: 100G
        redundancyPolicy: "SingleRedundancy"
    visualization:
      type: "kibana"
      kibana:
        replicas: 1
```

Generally Available

A broad ecosystem of workloads

Operator-backed services allow for a SaaS experience on your own infrastructure



Red Hat Certified Operators

DEVOPS



APM



INSTANA



DATA SERVICES



DATABASE



SECURITY



STORAGE



Generally Available

OperatorHub data sources

Requires an online cluster

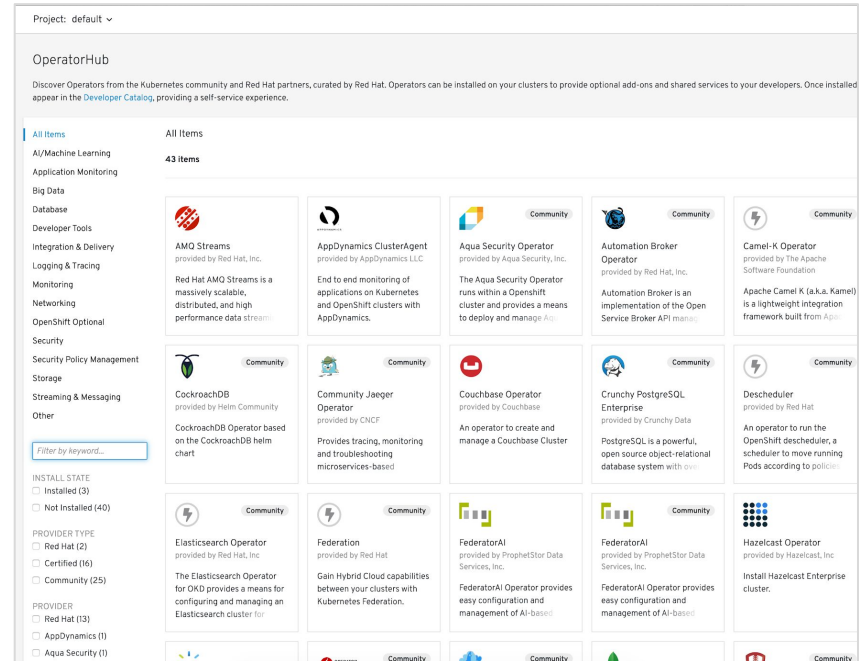
- For 4.1, the cluster must have connectivity to the internet
- Later 4.x releases will add offline capabilities

Operator Metadata

- Stored in quay.io
- Fetches channels and available versions for each Operator

Container Images

- Red Hat products and certified partners come from RHCC
- Community content comes from a variety of registries



Generally Available

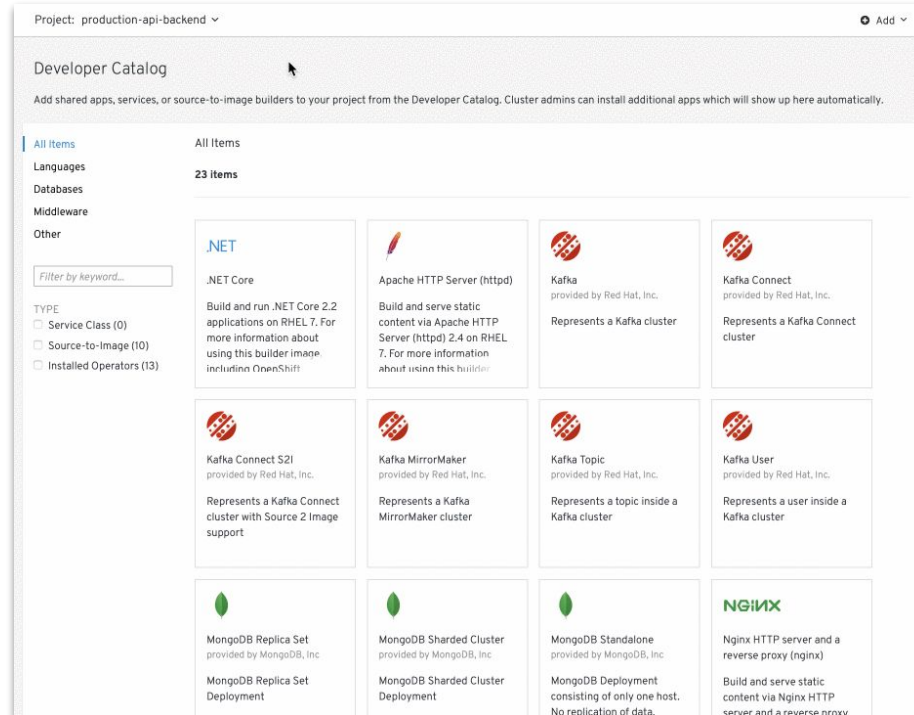
Services ready for your developers

New Developer Catalog aggregates apps

- Blended view of Operators, Templates and Broker backed services
- Operators can expose multiple CRDs. Example:
 - MongoDBReplicaSet
 - MongoDBSharded Cluster
 - MongoDBStandalone
- Developers can't see any of the admin screens

Self-service is key for productivity

- Developers with access can change settings and test out new services at any time



Operators as a First-Class Citizen

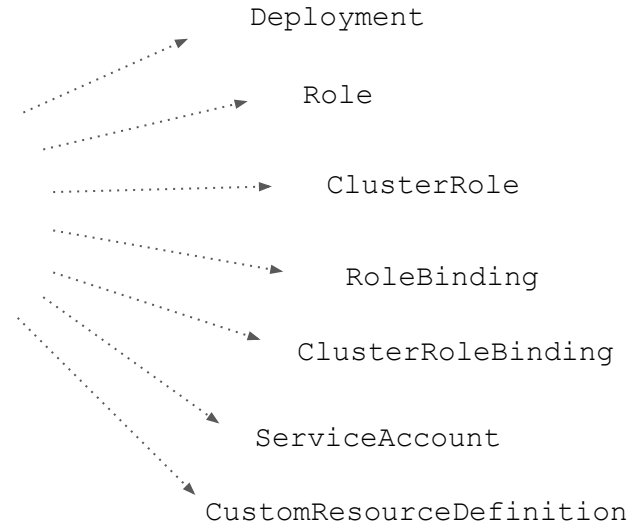


YourOperator v1.1.2
Bundle



**OPERATOR
LIFECYCLE MANAGER**

Operator Deployment
Custom Resource
Definitions
RBAC
API Dependencies
Update Path
Metadata



Generally Available

Operator Lifecycle Management


Operator Catalog



Subscription for YourOperator



Version



Time

Generally Available

Operator Lifecycle Management

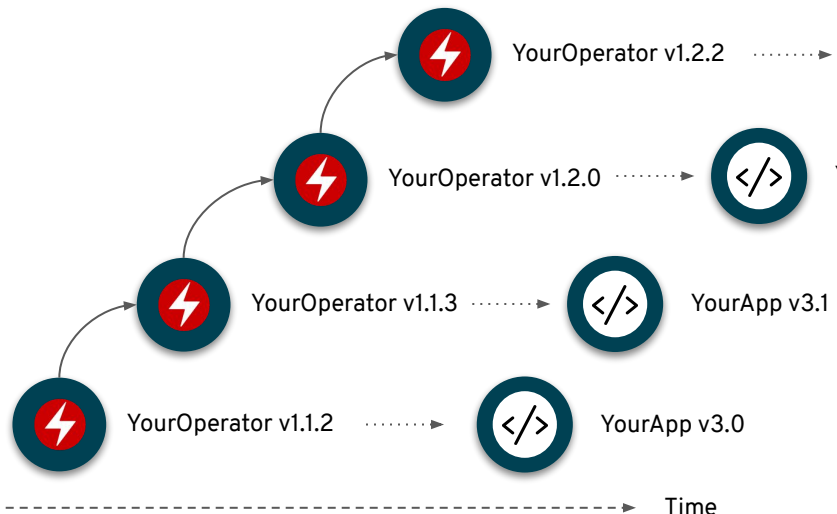
Operator Catalog



Subscription for YourOperator



Version



Time

Generally Available

Operator Upgrade in Detail

OperatorHub facilitates upgrades of installed Operators

- Manual or automatic modes can be chosen per Operator
- The Operator itself is upgraded by OLM via Deployment and a regular rolling upgrade
- The objects managed by the Operator use built in mechanisms to maintain HA
 - Deployments/StatefulSets
 - affinity/anti-affinity
 - taints/tolerations
 - PodDisruptionBudgets
- Behavior is dependent on the maturity of the Operator
- Optional cluster components like Cluster Logging are well behaved during upgrades


The screenshot displays the Red Hat OpenShift Operator Management interface. The left sidebar shows navigation options like Home, Catalog, Operator Management, and Workloads. The main content area shows 'Operator Subscriptions' with a table of installed operators. A tooltip is visible over the table, highlighting the 'STATUS', 'CHANNEL', and 'APPROVAL STRATEGY' columns.

NAME	NAMESPACE	STATUS	CHANNEL	APPROVAL STRATEGY
amq-streams	openshift-operators	Up to date	stable	Automatic
cockroachdb	openshift-operators	Up to date	stable	Automatic
codeready-workspaces	codeready	Up to date	final	Automatic
couchbase-enterprise-certified	robszumski-api-backend	Up to date	preview	Automatic


Tooltip details:

STATUS	CHANNEL	APPROVAL STRATEGY
Up to date	stable	Automatic
Up to date	stable	Automatic

Build Operators for your apps



Build operators from Helm chart, without any coding



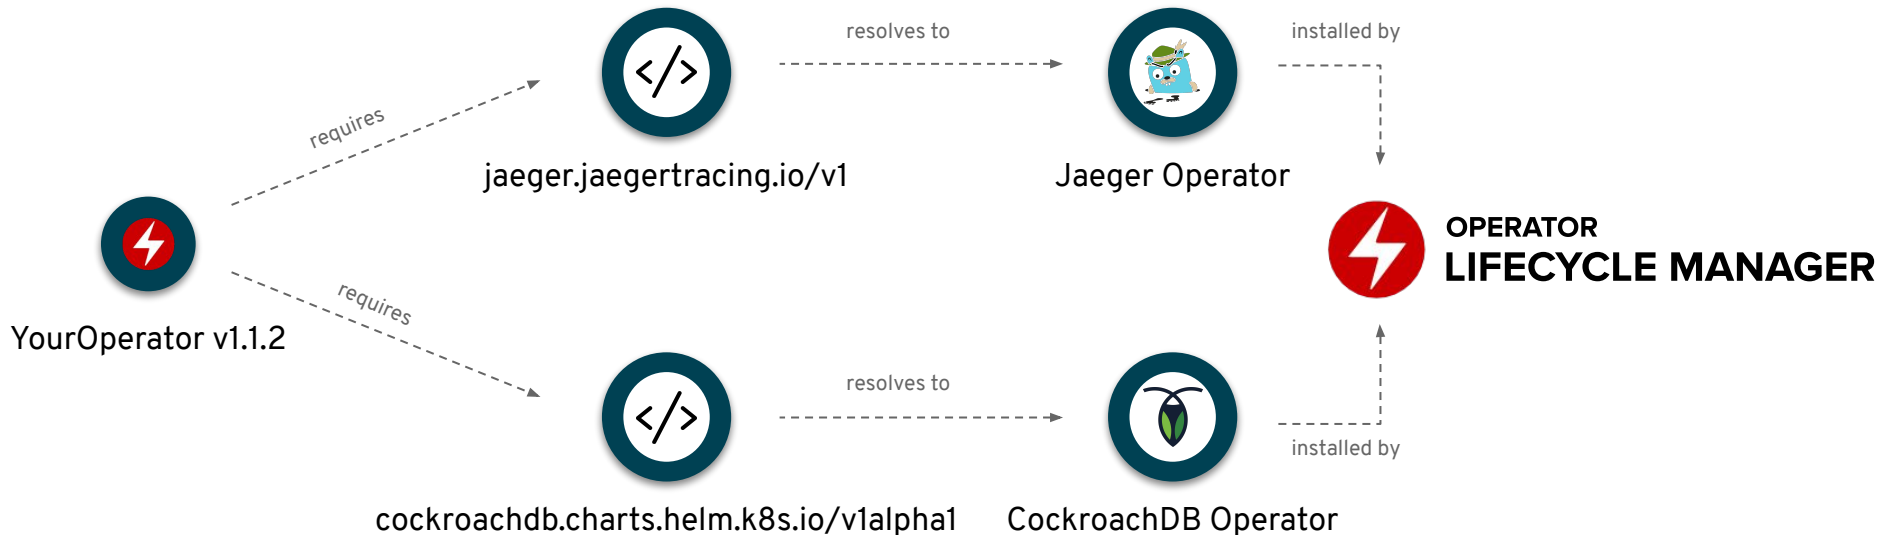
Build operators from Ansible playbooks and APBs



Build advanced operators for full lifecycle management

Depend on other Operators

Operator Framework Dependency Graphs




Red Hat Universal Base Image

Enable an ecosystem of freely distributable operators for Kubernetes/OpenShift



Base Images




Pre-Built Language Images




Package Subset


UBI and Host interactions



Red Hat Enterprise Linux 7



Red Hat Enterprise Linux 8



Like any community distro

Hosted OpenShift

Get the best of OpenShift without being on call



Hosted OpenShift Benefits

OPENSIFT CONTAINER PLATFORM

Full Stack Automated	Pre-existing Infrastructure
Skip the on-call rotation	Customer managed
Red Hat engineers keep you up to date	Provisioning
Expand capacity without hassle	Provisioning

HOSTED OPENSIFT

Azure Red Hat OpenShift	OpenShift Dedicated
Deploy directly from the Azure console.	Powerful cluster, no maintenance needed
Jointly managed by Red Hat and Azure engineers.	Managed by Red Hat engineers and support
Free your team from the distraction of ops	Free your team from the distraction of ops

Azure Red Hat OpenShift

Jointly engineered, operated, and supported by both Microsoft and Red Hat with an integrated support experience

Experience OpenShift as a native Microsoft Azure service.

- Create fully managed OpenShift clusters in minutes using `az openshift create`
- Add or remove compute nodes to match resource demand using `az openshift scale`
- 99.9% SLA
- Will soon inherit Azure regulatory compliance
- Pricing available at <https://azure.microsoft.com/en-us/pricing/details/openshift/>



OpenShift Dedicated

Dedicated with OpenShift 3

Available today, hosted on Amazon Web Services

Consumption based billing now available

Bring Your Own Cloud Account

Dedicated with OpenShift 4

Initial availability June 2019

Broader availability in fiscal Q2

OperatorHub

Red Hat products and certified Operators will be added in a curated catalog later in the year.


The Service Catalog and Brokers will not migrate to Dedicated due to their deprecation.

Connected to cloud.redhat.com

Clusters will appear beside other self-managed installs

Migrating to OpenShift 4

Tooling and advice for moving from OpenShift 3.x to 4.x



App migration experience

Using open source tooling based on Velero

Velero is an upstream project previously known as Ark. Check out [this video](#) if you are curious and want to get a sneak peek at our capabilities.

What's moved during a migration

- Namespaces
- Persistent Volumes (move or copy)
- All important resource objects (Deployments, StatefulSets, etc)

Available in OpenShift 4.2

Customers are anxious to get their hands on this, but we want to get it right. We would love to receive sample application workloads to test.

Name	Migrations	Source	Target	Repository	Persistent Volumes	Last Status
demo plan	2	Summit Demo Source Cluster	Target cluster	mydemobucket	2	Migrated Success
demo2	2	Summit Demo Source Cluster	Target cluster	mydemobucket	2	Migrated Success

Why did we choose this migration strategy?

Reducing risk


A ton of innovation went into OpenShift 4, and an in-place upgrade would have risk of failure in which there is no forwards or backwards remediation. It allows you to skip from 3.7/3.9/3.10/3.11 to 4.x. Skipping the need to install each one.

Useful for 4-to-4 migrations


A general migration tool is frequently requested and a better long term investment. Helps you build a foundation towards making your cluster investments less fragile.

Allows for staging


Stage a mock migration before doing it live, on a Project by Project basis. Extremely useful for success.



Questions?

 [linkedin.com/company/red-hat](https://www.linkedin.com/company/red-hat)

 [facebook.com/redhatinc](https://www.facebook.com/redhatinc)

 [youtube.com/user/RedHatVideos](https://www.youtube.com/user/RedHatVideos)

 twitter.com/RedHat