



# OpenShift Container Platform 4.1

## Machine management

Adding and maintaining cluster machines



# OpenShift Container Platform 4.1 Machine management

---

Adding and maintaining cluster machines

## Legal Notice

Copyright © 2019 Red Hat, Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at

<http://creativecommons.org/licenses/by-sa/3.0/>

. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, the Red Hat logo, JBoss, OpenShift, Fedora, the Infinity logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux<sup>®</sup> is the registered trademark of Linus Torvalds in the United States and other countries.

Java<sup>®</sup> is a registered trademark of Oracle and/or its affiliates.

XFS<sup>®</sup> is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

MySQL<sup>®</sup> is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js<sup>®</sup> is an official trademark of Joyent. Red Hat is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack<sup>®</sup> Word Mark and OpenStack logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

## Abstract

This document provides instructions for managing the machines that make up an OpenShift Container Platform 4.1 cluster. Some tasks make use of the enhanced automatic machine management functions of an OpenShift Container Platform 4.1 cluster and some tasks are manual. Not all tasks that are described in this document are available in all installation types.

## Table of Contents

<b>CHAPTER 1. CREATING A MACHINESET</b> .....	<b>4</b>
1.1. MACHINE API OVERVIEW	4
1.2. SAMPLE YAML FOR A MACHINESET CUSTOM RESOURCE	5
1.3. CREATING A MACHINESET	6
<b>CHAPTER 2. MANUALLY SCALING A MACHINESET</b> .....	<b>9</b>
2.1. SCALING A MACHINESET MANUALLY	9
<b>CHAPTER 3. MODIFYING A MACHINESET</b> .....	<b>10</b>
3.1. MODIFYING A MACHINESET	10
<b>CHAPTER 4. APPLYING AUTOSCALING TO AN OPENSIFT CONTAINER PLATFORM CLUSTER</b> .....	<b>12</b>
4.1. ABOUT THE CLUSTERAUTOSCALER	12
4.2. ABOUT THE MACHINEAUTOSCALER	13
4.3. CONFIGURING THE CLUSTERAUTOSCALER	13
4.3.1. ClusterAutoscaler resource definition	13
4.3.2. Deploying the ClusterAutoscaler	15
4.4. CONFIGURING THE MACHINEAUTOSCALERS	15
4.4.1. MachineAutoscaler resource definition	15
4.4.2. Deploying the MachineAutoscaler	16
4.5. ADDITIONAL RESOURCES	16
<b>CHAPTER 5. CREATING INFRASTRUCTURE MACHINESETS</b> .....	<b>17</b>
5.1. OPENSIFT CONTAINER PLATFORM INFRASTRUCTURE COMPONENTS	17
5.2. CREATING INFRASTRUCTURE MACHINESETS FOR PRODUCTION ENVIRONMENTS	17
5.2.1. Sample YAML for a MachineSet Custom Resource	17
5.2.2. Creating a MachineSet	19
5.3. MOVING RESOURCES TO INFRASTRUCTURE MACHINESETS	21
5.3.1. Moving the router	21
5.3.2. Moving the default registry	23
5.3.3. Moving the monitoring solution	24
5.3.4. Moving the cluster logging resources	24
<b>CHAPTER 6. ADDING RHEL COMPUTE MACHINES TO AN OPENSIFT CONTAINER PLATFORM CLUSTER</b> .	<b>27</b>
6.1. ABOUT ADDING RHEL COMPUTE NODES TO A CLUSTER	27
6.2. SYSTEM REQUIREMENTS FOR RHEL COMPUTE NODES	27
6.3. PREPARING THE MACHINE TO RUN THE PLAYBOOK	28
6.4. PREPARING A RHEL COMPUTE NODE	29
6.5. ADDING A RHEL COMPUTE MACHINE TO YOUR CLUSTER	30
6.6. APPROVING THE CSRS FOR YOUR MACHINES	31
6.7. REQUIRED PARAMETERS FOR THE ANSIBLE HOSTS FILE	33
6.7.1. Removing RHCOS compute machines from a cluster	33
<b>CHAPTER 7. ADDING MORE RHEL COMPUTE MACHINES TO AN OPENSIFT CONTAINER PLATFORM CLUSTER</b> .....	<b>35</b>
7.1. ABOUT ADDING RHEL COMPUTE NODES TO A CLUSTER	35
7.2. SYSTEM REQUIREMENTS FOR RHEL COMPUTE NODES	35
7.3. PREPARING A RHEL COMPUTE NODE	36
7.4. ADDING MORE RHEL COMPUTE MACHINES TO YOUR CLUSTER	37
7.5. APPROVING THE CSRS FOR YOUR MACHINES	38
7.6. REQUIRED PARAMETERS FOR THE ANSIBLE HOSTS FILE	39

<b>CHAPTER 8. DEPLOYING MACHINE HEALTH CHECKS</b> .....	<b>41</b>
8.1. ABOUT MACHINEHEALTHCHECKS	41
8.2. SAMPLE MACHINEHEALTHCHECK RESOURCE	41
8.3. CREATING A MACHINEHEALTHCHECK RESOURCE	42



# CHAPTER 1. CREATING A MACHINESSET

## 1.1. MACHINE API OVERVIEW

The Machine API is a combination of primary resources that are based on the upstream Cluster API project and custom OpenShift Container Platform resources.

For OpenShift Container Platform 4.1 clusters, the Machine API performs all node host provisioning management actions after the cluster installation finishes. Because of this system, OpenShift Container Platform 4.1 offers an elastic, dynamic provisioning method on top of public or private cloud infrastructure.

The two primary resources are:

### Machines

A fundamental unit that describes the host for a Node. A machine has a `providerSpec`, which describes the types of compute nodes that are offered for different cloud platforms. For example, a machine type for a worker node on Amazon Web Services (AWS) might define a specific machine type and required metadata.

### MachineSets

Groups of machines. MachineSets are to machines as ReplicaSets are to Pods. If you need more machines or must scale them down, you change the `replicas` field on the MachineSet to meet your compute need.

The following custom resources add more capabilities to your cluster:

### MachineAutoscaler

This resource automatically scales machines in a cloud. You can set the minimum and maximum scaling boundaries for nodes in a specified MachineSet, and the MachineAutoscaler maintains that range of nodes. The MachineAutoscaler object takes effect after a ClusterAutoscaler object exists. Both ClusterAutoscaler and MachineAutoscaler resources are made available by the ClusterAutoscalerOperator.

### ClusterAutoscaler

This resource is based on the upstream ClusterAutoscaler project. In the OpenShift Container Platform implementation, it is integrated with the Machine API by extending the MachineSet API. You can set cluster-wide scaling limits for resources such as cores, nodes, memory, GPU, and so on. You can set the priority so that the cluster prioritizes pods so that new nodes are not brought online for less important pods. You can also set the ScalingPolicy so you can scale up nodes but not scale them down.

### MachineHealthCheck

This resource detects when a machine is unhealthy, deletes it, and, on supported platforms, makes a new machine.



#### NOTE

In version 4.1, MachineHealthChecks is a Technology Preview feature

In OpenShift Container Platform version 3.11, you could not roll out a multi-zone architecture easily because the cluster did not manage machine provisioning. Beginning with 4.1 this process is easier. Each MachineSet is scoped to a single zone, so the installation program sends out MachineSets across



availability zones on your behalf. And then because your compute is dynamic, and in the face of a zone failure, you always have a zone for when you must rebalance your machines. The autoscaler provides best-effort balancing over the life of a cluster.

## 1.2. SAMPLE YAML FOR A MACHINESET CUSTOM RESOURCE

This sample YAML defines a MachineSet that runs in the **us-east-1a** Amazon Web Services (AWS) region and creates nodes that are labeled with **node-role.kubernetes.io/<role>**: ""

In this sample, **<clusterID>** is the cluster ID that you set when you provisioned the cluster and **<role>** is the node label to add.

```

apiVersion: machine.openshift.io/v1beta1
kind: MachineSet
metadata:
  labels:
    machine.openshift.io/cluster-api-cluster: <clusterID> 1
  name: <clusterID>-<role>-us-east-1a 2
  namespace: openshift-machine-api
spec:
  replicas: 1
  selector:
    matchLabels:
      machine.openshift.io/cluster-api-cluster: <clusterID> 3
      machine.openshift.io/cluster-api-machineset: <clusterID>-<role>-us-east-1a 4
  template:
    metadata:
      labels:
        machine.openshift.io/cluster-api-cluster: <clusterID> 5
        machine.openshift.io/cluster-api-machine-role: <role> 6
        machine.openshift.io/cluster-api-machine-type: <role> 7
        machine.openshift.io/cluster-api-machineset: <clusterID>-<role>-us-east-1a 8
    spec:
      metadata:
        labels:
          node-role.kubernetes.io/<role>: "" 9
      providerSpec:
        value:
          ami:
            id: ami-046fe691f52a953f9 10
          apiVersion: awsproviderconfig.openshift.io/v1beta1
          blockDevices:
            - ebs:
                iops: 0
                volumeSize: 120
                volumeType: gp2
          credentialsSecret:
            name: aws-cloud-credentials
          deviceIndex: 0
          iamInstanceProfile:
            id: <clusterID>-worker-profile 11
          instanceType: m4.large
          kind: AWSMachineProviderConfig
          placement:

```

```

availabilityZone: us-east-1a
region: us-east-1
securityGroups:
- filters:
  - name: tag:Name
    values:
      - <clusterID>-worker-sg 12
subnet:
filters:
- name: tag:Name
  values:
    - <clusterID>-private-us-east-1a 13
tags:
- name: kubernetes.io/cluster/<clusterID> 14
  value: owned
userDataSecret:
name: worker-user-data

```

**1 3 5 11 12 13 14** Specify the the cluster ID that you set when you provisioned the cluster.

**2 4 8** Specify the cluster ID and node label.

**6 7 9** Specify the node label to add.

**10** Specify a valid Red Hat Enterprise Linux CoreOS (RHCOS) AMI for your Amazon Web Services (AWS) zone for your OpenShift Container Platform nodes.

### 1.3. CREATING A MACHINESSET

In addition to to the ones created by the installation program, you can create your own MachineSets to dynamically manage the machine compute resources for specific workloads of your choice.

#### Prerequisites

- Deploy an OpenShift Container Platform cluster.
- Install the OpenShift Command-line Interface (CLI), commonly known as **oc**
- Log in to **oc** as a user with **cluster-admin** permission.

#### Procedure

1. Create a new YAML file that contains the MachineSet Custom Resource sample, as shown, and is named **<file\_name>.yaml**.

Ensure that you set the **<clusterID>** and **<role>** parameter values.

- a. If you are not sure about which value to set for an specific field, you can check an existing MachineSet from your cluster.

```
$ oc get machinesets -n openshift-machine-api
```

NAME	DESIRED	CURRENT	READY	AVAILABLE	AGE
agl030519-vplxk-worker-us-east-1a	1	1	1	1	55m
agl030519-vplxk-worker-us-east-1b	1	1	1	1	55m

```

agl030519-vplxk-worker-us-east-1c 1 1 1 1 55m
agl030519-vplxk-worker-us-east-1d 0 0 55m
agl030519-vplxk-worker-us-east-1e 0 0 55m
agl030519-vplxk-worker-us-east-1f 0 0 55m

```

- b. Check values of an specific MachineSet:

```

$ oc get machineset <machineset_name> -n \
  openshift-machine-api -o yaml

...

template:
  metadata:
    labels:
      machine.openshift.io/cluster-api-cluster: agl030519-vplxk 1
      machine.openshift.io/cluster-api-machine-role: worker 2
      machine.openshift.io/cluster-api-machine-type: worker
      machine.openshift.io/cluster-api-machineset: agl030519-vplxk-worker-us-east-1a

```

- 1** The cluster ID.
- 2** A default node label.

2. Create the new **MachineSet**:

```
$ oc create -f <file_name>.yaml
```

3. View the list of MachineSets:

```
$ oc get machineset -n openshift-machine-api
```

NAME	DESIRED	CURRENT	READY	AVAILABLE	AGE
agl030519-vplxk-infra-us-east-1a	1	1	1	1	11m
agl030519-vplxk-worker-us-east-1a	1	1	1	1	55m
agl030519-vplxk-worker-us-east-1b	1	1	1	1	55m
agl030519-vplxk-worker-us-east-1c	1	1	1	1	55m
agl030519-vplxk-worker-us-east-1d	0	0			55m
agl030519-vplxk-worker-us-east-1e	0	0			55m
agl030519-vplxk-worker-us-east-1f	0	0			55m

When the new MachineSet is available, the **DESIRED** and **CURRENT** values match. If the MachineSet is not available, wait a few minutes and run the command again.

4. After the new MachineSet is available, check status of the machine and the node that it references:

```
$ oc get machine -n openshift-machine-api

status:
  addresses:
    - address: 10.0.133.18
      type: InternalIP

```

```

- address: ""
  type: ExternalDNS
- address: ip-10-0-133-18.ec2.internal
  type: InternalDNS
lastUpdated: "2019-05-03T10:38:17Z"
nodeRef:
  kind: Node
  name: ip-10-0-133-18.ec2.internal
  uid: 71fb8d75-6d8f-11e9-9ff3-0e3f103c7cd8
providerStatus:
  apiVersion: awsproviderconfig.openshift.io/v1beta1
  conditions:
  - lastProbeTime: "2019-05-03T10:34:31Z"
    lastTransitionTime: "2019-05-03T10:34:31Z"
    message: machine successfully created
    reason: MachineCreationSucceeded
    status: "True"
    type: MachineCreation
  instanceId: i-09ca0701454124294
  instanceState: running
  kind: AWSMachineProviderStatus

```

- View the new node and confirm that the new node has the label that you specified:

```
$ oc get node <node_name> --show-labels
```

Review the command output and confirm that **node-role.kubernetes.io/<your\_label>** is in the **LABELS** list.



#### NOTE

Any change to a MachineSet is not applied to existing machines owned by the MachineSet. For example, labels edited or added to an existing MachineSet are not propagated to existing machines and Nodes associated with the MachineSet.

#### Next steps

If you need MachineSets in other availability zones, repeat this process to create more MachineSets.

## CHAPTER 2. MANUALLY SCALING A MACHINESET

You can add or remove an instance of a machine in a MachineSet.



### NOTE

If you need to modify aspects of a MachineSet outside of scaling, see [Modifying a MachineSet](#).

## 2.1. SCALING A MACHINESET MANUALLY

If you must add or remove an instance of a machine in a MachineSet, you can manually scale the MachineSet.

### Prerequisites

- Install an OpenShift Container Platform cluster and the **oc** command line.
- Log into **oc** as a user with **cluster-admin** permission.

### Procedure

1. View the MachineSets that are in the cluster:

```
$ oc get machinesets -n openshift-machine-api
```

The MachineSets are listed in the form of **<clusterid>-worker-<aws-region-az>**.

2. Scale the MachineSet:

```
$ oc scale --replicas=2 machineset <machineset> -n openshift-machine-api
```

Or:

```
$ oc edit machineset <machineset> -n openshift-machine-api
```

You can scale the MachineSet up or down. It takes several minutes for the new machines to be available.



### IMPORTANT

By default, the OpenShift Container Platform router pods are deployed on workers. Because the router is required to access some cluster resources, including the web console, do not scale the worker MachineSet to **1** or **0** unless you first relocate the router pods.

## CHAPTER 3. MODIFYING A MACHINESSET

You can make changes to a MachineSet, such as adding labels, changing the instance type, or changing block storage.



### NOTE

If you need to scale a MachineSet without making other changes, see [Manually scaling a MachineSet](#).

### 3.1. MODIFYING A MACHINESSET

To make changes to a MachineSet, edit the MachineSet YAML. Then remove all machines associated with the MachineSet by deleting or scaling down the machines to **0**. Then scale the replicas back to the desired number. Changes you make to a MachineSet do not affect existing machines.

If you need to scale a MachineSet without making other changes, you do not need to delete the machines.



### NOTE

By default, the OpenShift Container Platform router pods are deployed on workers. Because the router is required to access some cluster resources, including the web console, do not scale the worker MachineSet to **0** unless you first relocate the router pods.

#### Prerequisites

- Install an OpenShift Container Platform cluster and the `oc` command line.
- Log into `oc` as a user with `cluster-admin` permission.

#### Procedure

1. Edit the MachineSet:

```
$ oc edit machineset <machineset> -n openshift-machine-api
```

2. Scale down the MachineSet to **0**:

```
$ oc scale --replicas=0 machineset <machineset> -n openshift-machine-api
```

Or:

```
$ oc edit machineset <machineset> -n openshift-machine-api
```

Wait for the machines to be removed.

3. Scale up the MachineSet as needed:

```
$ oc scale --replicas=2 machineset <machineset> -n openshift-machine-api
```

Or:

```
$ oc edit machineset <machineset> -n openshift-machine-api
```

Wait for the machines to start. The new Machines contain changes you made to the Machineset.

## CHAPTER 4. APPLYING AUTOSCALING TO AN OPENSIFT CONTAINER PLATFORM CLUSTER

Applying autoscaling to an OpenShift Container Platform cluster involves deploying a ClusterAutoscaler and then deploying MachineAutoscalers for each Machine type in your cluster.

### 4.1. ABOUT THE CLUSTERAUTOSCALER

The ClusterAutoscaler adjusts the size of an OpenShift Container Platform cluster to meet its current deployment needs. It uses declarative, Kubernetes-style arguments to provide infrastructure management that does not rely on objects of a specific cloud provider. The ClusterAutoscaler has a cluster scope, and is not associated with a particular namespace.

The ClusterAutoscaler increases the size of the cluster when there are pods that failed to schedule on any of the current nodes due to insufficient resources or when another node is necessary to meet deployment needs. The ClusterAutoscaler does not increase the cluster resources beyond the limits that you specify.

The ClusterAutoscaler decreases the size of the cluster when some nodes are consistently not needed for a significant period, such as when it has low resource use and all of its important pods can fit on other nodes.

If the following types of pods are present on a node, the ClusterAutoscaler will not remove the node:

- Pods with restrictive PodDisruptionBudgets (PDBs).
- Kube-system pods that do not run on the node by default.
- Kube-system pods that do not have a PDB or have a PDB that is too restrictive.
- Pods that are not backed by a controller object such as a Deployment, ReplicaSet, or StatefulSet.
- Pods with local storage.
- Pods that cannot be moved elsewhere because of a lack of resources, incompatible node selectors or affinity, matching anti-affinity, and so on.
- Unless they also have a **"cluster-autoscaler.kubernetes.io/safe-to-evict": "true"** annotation, pods that have a **"cluster-autoscaler.kubernetes.io/safe-to-evict": "false"** annotation.

If you configure the ClusterAutoscaler, additional usage restrictions apply:

- Do not modify the nodes that are in autoscaled node groups directly. All nodes within the same node group have the same capacity and labels and run the same system pods.
- Specify requests for your pods.
- If you have to prevent pods from being deleted too quickly, configure appropriate PDBs.
- Confirm that your cloud provider quota is large enough to support the maximum node pools that you configure.
- Do not run additional node group autoscalers, especially the ones offered by your cloud provider.



The Horizontal Pod Autoscaler (HPA) and the ClusterAutoscaler modify cluster resources in different ways. The HPA changes the deployment's or ReplicaSet's number of replicas based on the current CPU load. If the load increases, the HPA creates new replicas, regardless of the amount of resources available to the cluster. If there are not enough resources, the ClusterAutoscaler adds resources so that the HPA-created pods can run. If the load decreases, the HPA stops some replicas. If this action causes some nodes to be underutilized or completely empty, the ClusterAutoscaler deletes the unnecessary nodes.

The ClusterAutoscaler takes pod priorities into account. The Pod Priority and Preemption feature enables scheduling pods based on priorities if the cluster does not have enough resources, but the ClusterAutoscaler ensures that the cluster has resources to run all pods. To honor the intention of both features, the ClusterAutoscaler includes a priority cutoff function. You can use this cutoff to schedule "best-effort" pods, which do not cause the ClusterAutoscaler to increase resources but instead run only when spare resources are available.

Pods with priority lower than the cutoff value do not cause the cluster to scale up or prevent the cluster from scaling down. No new nodes are added to run the pods, and nodes running these pods might be deleted to free resources.

## 4.2. ABOUT THE MACHINEAUTOSCALER

The MachineAutoscaler adjusts the number of Machines in the MachineSets that you deploy in an OpenShift Container Platform cluster. You can scale both the default **worker** MachineSet and any other MachineSets that you create. The MachineAutoscaler makes more Machines when the cluster runs out of resources to support more deployments. Any changes to the values in MachineAutoscaler resources, such as the minimum or maximum number of instances, are immediately applied to the MachineSet they target.



### IMPORTANT

You must deploy a MachineAutoscaler for the ClusterAutoscaler to scale your machines. The ClusterAutoscaler uses the annotations on MachineSets that the MachineAutoscaler sets to determine the resources that it can scale. If you define a ClusterAutoscaler without also defining MachineAutoscalers, the ClusterAutoscaler will never scale your cluster.

## 4.3. CONFIGURING THE CLUSTERAUTOSCALER

First, deploy the ClusterAutoscaler to manage automatic resource scaling in your OpenShift Container Platform cluster.



### NOTE

Because the ClusterAutoscaler is scoped to the entire cluster, you can make only one ClusterAutoscaler for the cluster.

### 4.3.1. ClusterAutoscaler resource definition

This **ClusterAutoscaler** resource definition shows the parameters and sample values for the ClusterAutoscaler.

```
apiVersion: "autoscaling.openshift.io/v1"
kind: "ClusterAutoscaler"
metadata:
  name: "default"
```

```

spec:
  podPriorityThreshold: -10 1
  resourceLimits:
    maxNodesTotal: 24 2
  cores:
    min: 8 3
    max: 128 4
  memory:
    min: 4 5
    max: 256 6
  gpus:
    - type: nvidia.com/gpu 7
      min: 0 8
      max: 16 9
    - type: amd.com/gpu 10
      min: 0 11
      max: 4 12
  scaleDown: 13
    enabled: true 14
    delayAfterAdd: 10m 15
    delayAfterDelete: 5m 16
    delayAfterFailure: 30s 17
    unneededTime: 60s 18

```

- 1 Specify the priority that a pod must exceed to cause the ClusterAutoscaler to deploy additional nodes. Enter a 32-bit integer value. The **podPriorityThreshold** value is compared to the value of the **PriorityClass** that you assign to each pod.
- 2 Specify the maximum number of nodes to deploy.
- 3 Specify the minimum number of cores to deploy.
- 4 Specify the maximum number of cores to deploy.
- 5 Specify the minimum amount of memory, in GiB, per node.
- 6 Specify the maximum amount of memory, in GiB, per node.
- 7 10 Optionally, specify the type of GPU node to deploy. Only **nvidia.com/gpu** and **amd.com/gpu** are valid types.
- 8 11 Specify the minimum number of GPUs to deploy.
- 9 12 Specify the maximum number of GPUs to deploy.
- 13 In this section, you can specify the period to wait for each action by using any valid [ParseDuration](#) interval, including **ns**, **us**, **ms**, **s**, **m**, and **h**.
- 14 Specify whether the ClusterAutoscaler can remove unnecessary nodes.
- 15 Optionally, specify the period to wait before deleting a node after a node has recently been *added*. If you do not specify a value, the default value of **10m** is used.
- 16

Specify the period to wait before deleting a node after a node has recently been *deleted*. If you do not specify a value, the default value of **10s** is used.

- 17 Specify the period to wait before deleting a node after a scale down failure occurred. If you do not specify a value, the default value of **3m** is used.
- 18 Specify the period before an unnecessary node is eligible for deletion. If you do not specify a value, the default value of **10m** is used.

### 4.3.2. Deploying the ClusterAutoscaler

To deploy the ClusterAutoscaler, you create an instance of the **ClusterAutoscaler** resource.

#### Procedure

1. Create a YAML file for the **ClusterAutoscaler** resource that contains the customized resource definition.
2. Create the resource in the cluster:

```
$ oc create -f <filename>.yaml 1
```

1 **<filename>** is the name of the resource file that you customized.

#### Next steps

- After you configure the ClusterAutoscaler, you must configure at least one MachineAutoscaler.

## 4.4. CONFIGURING THE MACHINEAUTOSCALERS

After you deploy the ClusterAutoscaler, deploy MachineAutoscaler resources that reference the MachineSets that are used to scale the cluster.



#### IMPORTANT

You must deploy at least one MachineAutoscaler resource after you deploy the ClusterAutoscaler resource.



#### NOTE

You must configure separate resources for each MachineSet. Remember that MachineSets are different in each AWS region, so consider whether you want to enable machine scaling in multiple regions.

### 4.4.1. MachineAutoscaler resource definition

This MachineAutoscaler resource definition shows the parameters and sample values for the MachineAutoscaler.

```
apiVersion: "autoscaling.openshift.io/v1beta1"
kind: "MachineAutoscaler"
```

```

metadata:
  name: "worker-us-east-1a" 1
  namespace: "openshift-machine-api"
spec:
  minReplicas: 1 2
  maxReplicas: 12 3
  scaleTargetRef: 4
    apiVersion: machine.openshift.io/v1beta1
    kind: MachineSet 5
    name: worker-us-east-1a 6

```

- 1 Specify the **MachineAutoscaler** name. To make it easier to identify which MachineSet this MachineAutoscaler scales, specify or include the name of the MachineSet to scale. The MachineSet name takes the following form: **<clusterid>-<machineset>-<aws-region-az>**
- 2 Specify the minimum number Machines of the specified type to deploy in the specified AWS zone.
- 3 Specify the maximum number Machines of the specified type to deploy in the specified AWS zone.
- 4 In this section, provide values that describe the existing MachineSet to scale.
- 5 The **kind** parameter value is always **MachineSet**.
- 6 The **name** value must match the name of an existing MachineSet, as shown in the **metadata.name** parameter value.

#### 4.4.2. Deploying the MachineAutoscaler

To deploy the MachineAutoscaler, you create an instance of the **MachineAutoscaler** resource.

##### Procedure

1. Create a YAML file for the **MachineAutoscaler** resource that contains the customized resource definition.
2. Create the resource in the cluster:

```
$ oc create -f <filename>.yaml 1
```

- 1 **<filename>** is the name of the resource file that you customized.

## 4.5. ADDITIONAL RESOURCES

- For more information about pod priority, see [Including pod priority in pod scheduling decisions in OpenShift Container Platform](#).

## CHAPTER 5. CREATING INFRASTRUCTURE MACHINESETS

You can create a MachineSet to host only infrastructure components. You apply specific Kubernetes labels to these Machines and then update the infrastructure components to run on only those Machines. These infrastructure nodes are not counted toward the total number of subscriptions that are required to run the environment.



### IMPORTANT

Unlike earlier versions of OpenShift Container Platform, you cannot move the infrastructure components to the master Machines. To move the components, you must create a new MachineSet.

## 5.1. OPENSIFT CONTAINER PLATFORM INFRASTRUCTURE COMPONENTS

The following OpenShift Container Platform components are infrastructure components:

- Kubernetes and OpenShift Container Platform control plane services that run on masters
- The default router
- The container image registry
- The cluster metrics collection, or monitoring service
- Cluster aggregated logging
- Service brokers

Any node that runs any other container, pod, or component is a worker node that your subscription must cover.

## 5.2. CREATING INFRASTRUCTURE MACHINESETS FOR PRODUCTION ENVIRONMENTS

In a production deployment, deploy at least three MachineSets to hold infrastructure components. Both the logging aggregation solution and the service mesh deploy ElasticSearch, and ElasticSearch requires three instances that are installed on different nodes. For high availability, install deploy these nodes to different availability zones. Since you need different MachineSets for each availability zone, create at least three MachineSets.

### 5.2.1. Sample YAML for a MachineSet Custom Resource

This sample YAML defines a MachineSet that runs in the **us-east-1a** Amazon Web Services (AWS) region and creates nodes that are labeled with **node-role.kubernetes.io/<role>**: ""

In this sample, **<clusterID>** is the cluster ID that you set when you provisioned the cluster and **<role>** is the node label to add.

```
apiVersion: machine.openshift.io/v1beta1
kind: MachineSet
metadata:
```

```

labels:
  machine.openshift.io/cluster-api-cluster: <clusterID> 1
name: <clusterID>-<role>-us-east-1a 2
namespace: openshift-machine-api
spec:
  replicas: 1
  selector:
    matchLabels:
      machine.openshift.io/cluster-api-cluster: <clusterID> 3
      machine.openshift.io/cluster-api-machineset: <clusterID>-<role>-us-east-1a 4
template:
  metadata:
    labels:
      machine.openshift.io/cluster-api-cluster: <clusterID> 5
      machine.openshift.io/cluster-api-machine-role: <role> 6
      machine.openshift.io/cluster-api-machine-type: <role> 7
      machine.openshift.io/cluster-api-machineset: <clusterID>-<role>-us-east-1a 8
  spec:
    metadata:
      labels:
        node-role.kubernetes.io/<role>: "" 9
    providerSpec:
      value:
        ami:
          id: ami-046fe691f52a953f9 10
        apiVersion: awsproviderconfig.openshift.io/v1beta1
        blockDevices:
          - ebs:
              iops: 0
              volumeSize: 120
              volumeType: gp2
        credentialsSecret:
          name: aws-cloud-credentials
        deviceIndex: 0
        iamInstanceProfile:
          id: <clusterID>-worker-profile 11
        instanceType: m4.large
        kind: AWSMachineProviderConfig
        placement:
          availabilityZone: us-east-1a
          region: us-east-1
        securityGroups:
          - filters:
              - name: tag:Name
                values:
                  - <clusterID>-worker-sg 12
        subnet:
          filters:
            - name: tag:Name
              values:
                - <clusterID>-private-us-east-1a 13
        tags:
          - name: kubernetes.io/cluster/<clusterID> 14

```

```

value: owned
userDataSecret:
name: worker-user-data

```

**1 3 5 11 12 13 14** Specify the the cluster ID that you set when you provisioned the cluster.

**2 4 8** Specify the cluster ID and node label.

**6 7 9** Specify the node label to add.

**10** Specify a valid Red Hat Enterprise Linux CoreOS (RHCOS) AMI for your Amazon Web Services (AWS) zone for your OpenShift Container Platform nodes.

## 5.2.2. Creating a MachineSet

In addition to to the ones created by the installation program, you can create your own MachineSets to dynamically manage the machine compute resources for specific workloads of your choice.

### Prerequisites

- Deploy an OpenShift Container Platform cluster.
- Install the OpenShift Command-line Interface (CLI), commonly known as **oc**
- Log in to **oc** as a user with **cluster-admin** permission.

### Procedure

1. Create a new YAML file that contains the MachineSet Custom Resource sample, as shown, and is named **<file\_name>.yaml**.

Ensure that you set the **<clusterID>** and **<role>** parameter values.

- a. If you are not sure about which value to set for an specific field, you can check an existing MachineSet from your cluster.

```

$ oc get machinesets -n openshift-machine-api

NAME                                DESIRED  CURRENT  READY  AVAILABLE  AGE
agl030519-vplxk-worker-us-east-1a  1        1        1      1          55m
agl030519-vplxk-worker-us-east-1b  1        1        1      1          55m
agl030519-vplxk-worker-us-east-1c  1        1        1      1          55m
agl030519-vplxk-worker-us-east-1d  0        0        0      0          55m
agl030519-vplxk-worker-us-east-1e  0        0        0      0          55m
agl030519-vplxk-worker-us-east-1f  0        0        0      0          55m

```

- b. Check values of an specific MachineSet:

```

$ oc get machineset <machineset_name> -n \
  openshift-machine-api -o yaml

...

template:
  metadata:

```

labels:

```
machine.openshift.io/cluster-api-cluster: agl030519-vplxk 1
machine.openshift.io/cluster-api-machine-role: worker 2
machine.openshift.io/cluster-api-machine-type: worker
machine.openshift.io/cluster-api-machineset: agl030519-vplxk-worker-us-east-1a
```

- 1** The cluster ID.
- 2** A default node label.

## 2. Create the new **MachineSet**:

```
$ oc create -f <file_name>.yaml
```

## 3. View the list of MachineSets:

```
$ oc get machineset -n openshift-machine-api
```

NAME	DESIRED	CURRENT	READY	AVAILABLE	AGE
agl030519-vplxk-infra-us-east-1a	1	1	1	1	11m
agl030519-vplxk-worker-us-east-1a	1	1	1	1	55m
agl030519-vplxk-worker-us-east-1b	1	1	1	1	55m
agl030519-vplxk-worker-us-east-1c	1	1	1	1	55m
agl030519-vplxk-worker-us-east-1d	0	0			55m
agl030519-vplxk-worker-us-east-1e	0	0			55m
agl030519-vplxk-worker-us-east-1f	0	0			55m

When the new MachineSet is available, the **DESIRED** and **CURRENT** values match. If the MachineSet is not available, wait a few minutes and run the command again.

## 4. After the new MachineSet is available, check status of the machine and the node that it references:

```
$ oc get machine -n openshift-machine-api
```

status:

addresses:

```
- address: 10.0.133.18
  type: InternalIP
```

```
- address: ""
```

type: ExternalDNS

```
- address: ip-10-0-133-18.ec2.internal
```

type: InternalDNS

```
lastUpdated: "2019-05-03T10:38:17Z"
```

nodeRef:

kind: Node

name: ip-10-0-133-18.ec2.internal

uid: 71fb8d75-6d8f-11e9-9ff3-0e3f103c7cd8

providerStatus:

apiVersion: awsproviderconfig.openshift.io/v1beta1

conditions:

```
- lastProbeTime: "2019-05-03T10:34:31Z"
```

```
  lastTransitionTime: "2019-05-03T10:34:31Z"
```



```

message: machine successfully created
reason: MachineCreationSucceeded
status: "True"
type: MachineCreation
instanceId: i-09ca0701454124294
instanceState: running
kind: AWSMachineProviderStatus

```

- View the new node and confirm that the new node has the label that you specified:

```
$ oc get node <node_name> --show-labels
```

Review the command output and confirm that **node-role.kubernetes.io/<your\_label>** is in the **LABELS** list.



## NOTE

Any change to a MachineSet is not applied to existing machines owned by the MachineSet. For example, labels edited or added to an existing MachineSet are not propagated to existing machines and Nodes associated with the MachineSet.

## Next steps

If you need MachineSets in other availability zones, repeat this process to create more MachineSets.

## 5.3. MOVING RESOURCES TO INFRASTRUCTURE MACHINESSETS

Some of the infrastructure resources are deployed in your cluster by default. You can move them to the infrastructure MachineSets that you created.

### 5.3.1. Moving the router

You can deploy the router Pod to a different MachineSet. By default, the Pod is displayed to a worker node.

#### Prerequisites

- Configure additional MachineSets in your OpenShift Container Platform cluster.

#### Procedure

- View the **IngressController** Custom Resource for the router Operator:

```
$ oc get ingresscontroller default -n openshift-ingress-operator -o yaml
```

The command output resembles the following text:

```

apiVersion: operator.openshift.io/v1
kind: IngressController
metadata:
  creationTimestamp: 2019-04-18T12:35:39Z
finalizers:
  - ingresscontroller.operator.openshift.io/finalizer-ingresscontroller

```

```

generation: 1
name: default
namespace: openshift-ingress-operator
resourceVersion: "11341"
selfLink: /apis/operator.openshift.io/v1/namespaces/openshift-ingress-
operator/ingresscontrollers/default
uid: 79509e05-61d6-11e9-bc55-02ce4781844a
spec: {}
status:
  availableReplicas: 2
  conditions:
  - lastTransitionTime: 2019-04-18T12:36:15Z
    status: "True"
    type: Available
  domain: apps.<cluster>.example.com
  endpointPublishingStrategy:
    type: LoadBalancerService
  selector: ingresscontroller.operator.openshift.io/deployment-ingresscontroller=default

```

2. Edit the **ingresscontroller** resource and change the **nodeSelector** to use the **infra** label:

```
$ oc edit ingresscontroller default -n openshift-ingress-operator -o yaml
```

Add the **nodeSelector** stanza that references the **infra** label to the **spec** section, as shown:

```

spec:
  nodePlacement:
  nodeSelector:
    matchLabels:
      node-role.kubernetes.io/infra: ""

```

3. Confirm that the router pod is running on the **infra** node.
  - a. View the list of router pods and note the node name of the running pod:

```
$ oc get pod -n openshift-ingress -o wide
```

AME	READY	STATUS	RESTARTS	AGE	IP	NODE
NOMINATED	NODE	READINESS	GATES			
router-default-86798b4b5d-bdlvd	1/1	Running	0	28s	10.130.2.4	ip-10-0-217-226.ec2.internal
router-default-955d875f4-255g8	0/1	Terminating	0	19h	10.129.2.4	ip-10-0-148-172.ec2.internal

In this example, the running pod is on the **ip-10-0-217-226.ec2.internal** node.

- b. View the node status of the running pod:

```
$ oc get node <node_name> ❶
```

NAME	STATUS	ROLES	AGE	VERSION
ip-10-0-217-226.ec2.internal	Ready	infra,worker	17h	v1.11.0+406fc897d8

- ❶ Specify the **<node\_name>** that you obtained from the pod list.

Because the role list includes **infra**, the pod is running on the correct node.

### 5.3.2. Moving the default registry

You configure the registry Operator to deploy its pods to different nodes.

#### Prerequisites

- Configure additional MachineSets in your OpenShift Container Platform cluster.

#### Procedure

1. View the **config/instance** object:

```
$ oc get config/cluster -n openshift-ingress -o yaml
```

The output resembles the following text:

```
apiVersion: imageregistry.operator.openshift.io/v1
kind: Config
metadata:
  creationTimestamp: 2019-02-05T13:52:05Z
  finalizers:
  - imageregistry.operator.openshift.io/finalizer
  generation: 1
  name: cluster
  resourceVersion: "56174"
  selfLink: /apis/imageregistry.operator.openshift.io/v1/configs/cluster
  uid: 36fd3724-294d-11e9-a524-12fjee2931b
spec:
  httpSecret: d9a012ccd117b1e6616ceccb2c3bb66a5fed1b5e481623
  logging: 2
  managementState: Managed
  proxy: {}
  replicas: 1
  requests:
    read: {}
    write: {}
  storage:
    s3:
      bucket: image-registry-us-east-1-c92e88cad85b48ec8b312344dff03c82-392c
      region: us-east-1
status:
...
```

2. Edit the **config/instance** object:

```
$ oc edit config/cluster -n openshift-ingress
```

3. Add the following lines of text the **spec** section of the object:

```
nodeSelector:
  node-role.kubernetes.io/infra: ""
```

After you save and exit you can see the registry pod being moving to the infrastructure node.

### 5.3.3. Moving the monitoring solution

By default, the Prometheus Cluster Monitoring stack, which contains Prometheus, Grafana, and AlertManager, is deployed to provide cluster monitoring. It is managed by the Cluster Monitoring Operator. To move its components to different machines, you create and apply a custom ConfigMap.

#### Procedure

1. Save the following ConfigMap definition as the **cluster-monitoring-configmap.yaml** file:

```
apiVersion: v1
kind: ConfigMap
metadata:
  name: cluster-monitoring-config
  namespace: openshift-monitoring
data:
  config.yaml: |+
    alertmanagerMain:
      nodeSelector:
        node-role.kubernetes.io/infra: ""
    prometheusK8s:
      nodeSelector:
        node-role.kubernetes.io/infra: ""
    prometheusOperator:
      nodeSelector:
        node-role.kubernetes.io/infra: ""
    grafana:
      nodeSelector:
        node-role.kubernetes.io/infra: ""
    k8sPrometheusAdapter:
      nodeSelector:
        node-role.kubernetes.io/infra: ""
    kubeStateMetrics:
      nodeSelector:
        node-role.kubernetes.io/infra: ""
    telemeterClient:
      nodeSelector:
        node-role.kubernetes.io/infra: ""
```

Running this ConfigMap forces the components of the monitoring stack to redeploy to infrastructure nodes.

2. Apply the new ConfigMap:

```
$ oc create -f cluster-monitoring-configmap.yaml
```


3. Watch the monitoring Pods move to the new machines:

```
$ watch 'oc get pod -n openshift-monitoring -o wide'
```

### 5.3.4. Moving the cluster logging resources

You can configure the Cluster Logging Operator to deploy the pods for any or all of the Cluster Logging components, Elasticsearch, Kibana, and Curator to different nodes. You cannot move the Cluster Logging Operator pod from its installed location.

For example, you can move the Elasticsearch pods to a separate node because of high CPU, memory, and disk requirements.



## NOTE

You should set your MachineSet to use at least 6 replicas.

## Prerequisites

- Cluster logging and Elasticsearch must be installed. These features are not installed by default.

## Procedure

1. Edit the Cluster Logging Custom Resource in the **openshift-logging** project:

```
$ oc edit ClusterLogging instance

apiVersion: logging.openshift.io/v1
kind: ClusterLogging

....

spec:
  collection:
    logs:
      fluentd:
        resources: null
      rsyslog:
        resources: null
      type: fluentd
  curator:
    curator:
      nodeSelector: 1
        node-role.kubernetes.io/infra: "
      resources: null
      schedule: 30 3 * * *
      type: curator
  logStore:
    elasticsearch:
      nodeCount: 3
      nodeSelector: 2
        node-role.kubernetes.io/infra: "
      redundancyPolicy: SingleRedundancy
    resources:
      limits:
        cpu: 500m
        memory: 4Gi
      requests:
        cpu: 500m
        memory: 4Gi
    storage: {}
```

```
type: elasticsearch
managementState: Managed
visualization:
  kibana:
    nodeSelector: 3
      node-role.kubernetes.io/infra: " 4
    proxy:
      resources: null
    replicas: 1
    resources: null
    type: kibana
....
```

- 1 2 3 4 Add a **nodeSelector** parameter with the appropriate value to the component you want to move. You can use a **nodeSelector** in the format shown or use **<key>: <value>** pairs, based on the value specified for the node.

## CHAPTER 6. ADDING RHEL COMPUTE MACHINES TO AN OPENSIFT CONTAINER PLATFORM CLUSTER

In OpenShift Container Platform, you can add Red Hat Enterprise Linux (RHEL) compute, or worker, machines to a user-provisioned infrastructure cluster. You can use RHEL as the operating system on only compute machines.

### 6.1. ABOUT ADDING RHEL COMPUTE NODES TO A CLUSTER

In OpenShift Container Platform 4.1, you have the option of using Red Hat Enterprise Linux (RHEL) machines as compute, or worker, machines in your cluster if you use a user-provisioned infrastructure installation. You must use Red Hat Enterprise Linux CoreOS (RHCOS) machines for the control plane, or master, machines in your cluster.

As with all installations that use user-provisioned infrastructure, if you choose to use RHEL compute machines in your cluster, you take responsibility for all operating system life cycle management and maintenance, including performing system updates, applying patches, and completing all other required tasks.



#### IMPORTANT

Because removing OpenShift Container Platform from a machine in the cluster requires destroying the operating system, you must use dedicated hardware for any RHEL machines that you add to the cluster.



#### IMPORTANT

Swap memory is disabled on all RHEL machines that you add to your OpenShift Container Platform cluster. You cannot enable swap memory on these machines.

You must add RHEL compute machines to the cluster after you initialize the control plane.

### 6.2. SYSTEM REQUIREMENTS FOR RHEL COMPUTE NODES

The Red Hat Enterprise Linux (RHEL) compute, or worker, machine hosts in your OpenShift Container Platform environment must meet the following minimum hardware specifications and system-level requirements.

- You must have an active OpenShift Container Platform subscription on your Red Hat account. If you do not, contact your sales representative for more information.
- Production environments must provide compute machines to support your expected workloads. As an OpenShift Container Platform cluster administrator, you must calculate the expected workload and add about 10 percent for overhead. For production environments, allocate enough resources so that a node host failure does not affect your maximum capacity.
- Each system must meet the following hardware requirements:
  - Physical or virtual system, or an instance running on a public or private IaaS.
  - Base OS: [RHEL 7.6](#) with "Minimal" installation option.

**IMPORTANT**

Only RHEL 7.6 is supported in OpenShift Container Platform 4.1. You must not upgrade your compute machines to RHEL 8.

- NetworkManager 1.0 or later.
- 1 vCPU.
- Minimum 8 GB RAM.
- Minimum 15 GB hard disk space for the file system containing **/var/**.
- Minimum 1 GB hard disk space for the file system containing **/usr/local/bin/**.
- Minimum 1 GB hard disk space for the file system containing the system's temporary directory. The system's temporary directory is determined according to the rules defined in the tempfile module in Python's standard library.
- Each system must meet any additional requirements for your system provider. For example, if you installed your cluster on VMware vSphere, your disks must be configured according to its [storage guidelines](#) and the **disk.enableUUID=true** attribute must be set.

## 6.3. PREPARING THE MACHINE TO RUN THE PLAYBOOK


Before you can add compute machines that use Red Hat Enterprise Linux as the operating system to an OpenShift Container Platform 4.1 cluster, you must prepare a machine to run the playbook from. This machine is not part of the cluster but must be able to access it.

### Prerequisites

- Install the OpenShift Command-line Interface (CLI), commonly known as **oc**, on the machine that you run the playbook on.
- Log in as a user with **cluster-admin** permission.

### Procedure

1. Ensure that the **kubeconfig** file for the cluster and the installation program that you used to install the cluster are on the machine. One way to accomplish this is to use the same machine that you used to install the cluster.
2. Configure the machine to access all of the RHEL hosts that you plan to use as compute machines. You can use any method that your company allows, including a bastion with an SSH proxy or a VPN.
3. Configure a user on the machine that you run the playbook on that has SSH access to all of the RHEL hosts.

**IMPORTANT**

If you use SSH key-based authentication, you must manage the key with an SSH agent.



4. If you have not already done so, register the machine with RHSM and attach a pool with an **OpenShift** subscription to it:

- a. Register the machine with RHSM:

```
# subscription-manager register --username=<user_name> --password=<password>
```

- b. Pull the latest subscription data from RHSM:

```
# subscription-manager refresh
```

- c. List the available subscriptions:

```
# subscription-manager list --available --matches '*OpenShift*'
```

- d. In the output for the previous command, find the pool ID for an OpenShift Container Platform subscription and attach it:

```
# subscription-manager attach --pool=<pool_id>
```

5. Enable the repositories required by OpenShift Container Platform 4.1:

```
# subscription-manager repos \
  --enable="rhel-7-server-rpms" \
  --enable="rhel-7-server-extras-rpms" \
  --enable="rhel-7-server-ansible-2.7-rpms" \
  --enable="rhel-7-server-ose-4.1-rpms"
```

6. Install the required packages, including **Openshift-Ansible**:

```
# yum install openshift-ansible openshift-clients jq
```

The **openshift-ansible** package provides installation program utilities and pulls in other packages that you require to add a RHEL compute node to your cluster, such as Ansible, playbooks, and related configuration files. The **openshift-clients** provides the **oc** CLI, and the **jq** package improves the display of JSON output on your command line.

## 6.4. PREPARING A RHEL COMPUTE NODE

Before you add a Red Hat Enterprise Linux (RHEL) machine to your OpenShift Container Platform cluster, you must register each host with Red Hat Subscription Manager (RHSM), attach an active OpenShift Container Platform subscription, and enable the required repositories.

1. On each host, register with RHSM:

```
# subscription-manager register --username=<user_name> --password=<password>
```

2. Pull the latest subscription data from RHSM:

```
# subscription-manager refresh
```

3. List the available subscriptions:

```
# subscription-manager list --available --matches '*OpenShift*'
```

- In the output for the previous command, find the pool ID for an OpenShift Container Platform subscription and attach it:

```
# subscription-manager attach --pool=<pool_id>
```

- Disable all yum repositories:

- Disable all the enabled RHSM repositories:

```
# subscription-manager repos --disable="**"
```

- List the remaining yum repositories and note their names under **repo id**, if any:

```
# yum repolist
```

- Use **yum-config-manager** to disable the remaining yum repositories:

```
# yum-config-manager --disable <repo_id>
```

Alternatively, disable all repositories:

```
yum-config-manager --disable \*
```

Note that this might take a few minutes if you have a large number of available repositories

- Enable only the repositories required by OpenShift Container Platform 4.1:

```
# subscription-manager repos \
  --enable="rhel-7-server-rpms" \
  --enable="rhel-7-server-extras-rpms" \
  --enable="rhel-7-server-ose-4.1-rpms"
```

## 6.5. ADDING A RHEL COMPUTE MACHINE TO YOUR CLUSTER

You can add compute machines that use Red Hat Enterprise Linux as the operating system to an OpenShift Container Platform 4.1 cluster.

### Prerequisites

- You installed the required packages and performed the necessary configuration on the machine that you run the playbook on.
- You prepared the RHEL hosts for installation.

### Procedure

Perform the following steps on the machine that you prepared to run the playbook:

- Extract the pull secret for your cluster:

```
$ oc -n openshift-config get -o jsonpath='{.data.\.dockerconfigjson}' secret pull-secret |
base64 -d | jq .
```

2. Save the pull secret in a file that is named **pull-secret.txt**.
3. Create an Ansible inventory file that is named **<path>/inventory/hosts** that defines your compute machine hosts and required variables:

```
[all:vars]
ansible_user=root 1
#ansible_become=True 2

openshift_kubeconfig_path=~/.kube/config" 3
openshift_pull_secret_path=~/.pull-secret.txt" 4

[new_workers] 5
mycluster-worker-0.example.com
mycluster-worker-1.example.com
```

- 1 Specify the user name that runs the Ansible tasks on the remote compute machines.
- 2 If you do not specify **root** for the **ansible\_user**, you must set **ansible\_become** to **True** and assign the user sudo permissions.
- 3 Specify the path to the **kubeconfig** file for your cluster.
- 4 Specify the path to the file that contains the pull secret for the image registry for your cluster.
- 5 List each RHEL machine to add to your cluster. You must provide the fully-qualified domain name for each host. This name is the host name that the cluster uses to access the machine, so set the correct public or private name to access the machine.

4. Run the playbook:

```
$ cd /usr/share/ansible/openshift-ansible
$ ansible-playbook -i <path>/inventory/hosts playbooks/scaleup.yml 1
```

- 1 For **<path>**, specify the path to the Ansible inventory file that you created.

## 6.6. APPROVING THE CSRS FOR YOUR MACHINES

When you add machines to a cluster, two pending certificates signing request (CSRs) are generated for each machine that you added. You must confirm that these CSRs are approved or, if necessary, approve them yourself.

### Prerequisites

- You added machines to your cluster.
- Install the **jq** package.

## Procedure

1. Confirm that the cluster recognizes the machines:

```
$ oc get nodes
```

```
NAME      STATUS    ROLES    AGE   VERSION
master-0  Ready    master   63m   v1.13.4+b626c2fe1
master-1  Ready    master   63m   v1.13.4+b626c2fe1
master-2  Ready    master   64m   v1.13.4+b626c2fe1
worker-0  NotReady worker   76s   v1.13.4+b626c2fe1
worker-1  NotReady worker   70s   v1.13.4+b626c2fe1
```

The output lists all of the machines that you created.

2. Review the pending certificate signing requests (CSRs) and ensure that you see a client and server request with **Pending** or **Approved** status for each machine that you added to the cluster:


```
$ oc get csr
```

```
NAME      AGE   REQUESTOR                                     CONDITION
csr-8b2br  15m   system:serviceaccount:openshift-machine-config-operator:node-
bootstrapper Pending 1
csr-8vnps  15m   system:serviceaccount:openshift-machine-config-operator:node-
bootstrapper Pending
csr-bfd72  5m26s system:node:ip-10-0-50-126.us-east-2.compute.internal
Pending 2
csr-c57lv  5m26s system:node:ip-10-0-95-157.us-east-2.compute.internal
Pending
...
```

- 1** A client request CSR.
- 2** A server request CSR.

In this example, two machines are joining the cluster. You might see more approved CSRs in the list.

3. If the CSRs were not approved, after all of the pending CSRs for the machines you added are in **Pending** status, approve the CSRs for your cluster machines:



### NOTE

Because the CSRs rotate automatically, approve your CSRs within an hour of adding the machines to the cluster. If you do not approve them within an hour, the certificates will rotate, and more than two certificates will be present for each node. You must approve all of these certificates. After you approve the initial CSRs, the subsequent node client CSRs are automatically approved by the cluster **kube-controller-manager**. You must implement a method of automatically approving the kubelet serving certificate requests.

- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

**1** **<csr\_name>** is the name of a CSR from the list of current CSRs.

- If all the CSRs are valid, approve them all by running the following command:

```
$ oc get csr -ojson | jq -r '.items[] | select(.status == {}) | .metadata.name' | xargs oc adm certificate approve
```

## 6.7. REQUIRED PARAMETERS FOR THE ANSIBLE HOSTS FILE

You must define the following parameters in the Ansible hosts file before you add Red Hat Enterprise Linux (RHEL) compute machines to your cluster.

Parameter	Description	Values
<b>ansible_user</b>	The SSH user that allows SSH-based authentication without requiring a password. If you use SSH key-based authentication, then you must manage the key with an SSH agent.	A user name on the system. The default value is <b>root</b> .
<b>ansible_become</b>	If the values of <b>ansible_user</b> is not root, you must set <b>ansible_become</b> to <b>True</b> , and the user that you specify as the <b>ansible_user</b> must be configured for passwordless sudo access.	<b>True</b> . If the value is not <b>True</b> , do not specify and define this parameter.
<b>openshift_kubeconfig_path</b>	Specifies a path to a local directory that contains the <b>kubeconfig</b> file for your cluster.	The path and name of the configuration file.
<b>openshift_pull_secret_path</b>	Specifies a path to the text file that contains the pull secret to the image registry for your cluster. Use the pull secret that you obtained from the <a href="#">OpenShift Infrastructure Providers</a> page. This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.	The path and name of the pull secret file.

### 6.7.1. Removing RHCOS compute machines from a cluster

After you add the Red Hat Enterprise Linux (RHEL) compute machines to your cluster, you can remove the the Red Hat Enterprise Linux CoreOS (RHCOS) compute machines.

#### Prerequisites

- You have added RHEL compute machines to your cluster.

## Procedure

1. View the list of machines and record the node names of the RHCOS compute machines:

```
$ oc get nodes -o wide
```

2. For each RHCOS compute machine, delete the node:

- a. Mark the node as unschedulable by running the **oc adm cordon** command:

```
$ oc adm cordon <node_name> 1
```

- 1 Specify the node name of one of the RHCOS compute machines.

- b. Drain all the pods from the node:

```
$ oc adm drain <node_name> --force --delete-local-data --ignore-daemonsets 1
```

- 1 Specify the node name of the RHCOS compute machine that you isolated.

- c. Delete the node:

```
$ oc delete nodes <node_name> 1
```

- 1 Specify the node name of the RHCOS compute machine that you drained.

3. Review the list of compute machines to ensure that only the RHEL nodes remain:

```
$ oc get nodes -o wide
```

4. Remove the RHCOS machines from the load balancer for your cluster's compute machines. You can delete the Virtual Machines or reimagine the physical hardware for the RHCOS compute machines.

## CHAPTER 7. ADDING MORE RHEL COMPUTE MACHINES TO AN OPENSIFT CONTAINER PLATFORM CLUSTER

If your OpenShift Container Platform cluster already includes Red Hat Enterprise Linux (RHEL) compute, or worker, machines, you can add more RHEL compute machines to it.

### 7.1. ABOUT ADDING RHEL COMPUTE NODES TO A CLUSTER

In OpenShift Container Platform 4.1, you have the option of using Red Hat Enterprise Linux (RHEL) machines as compute, or worker, machines in your cluster if you use a user-provisioned infrastructure installation. You must use Red Hat Enterprise Linux CoreOS (RHCOS) machines for the control plane, or master, machines in your cluster.

As with all installations that use user-provisioned infrastructure, if you choose to use RHEL compute machines in your cluster, you take responsibility for all operating system life cycle management and maintenance, including performing system updates, applying patches, and completing all other required tasks.



#### IMPORTANT

Because removing OpenShift Container Platform from a machine in the cluster requires destroying the operating system, you must use dedicated hardware for any RHEL machines that you add to the cluster.



#### IMPORTANT

Swap memory is disabled on all RHEL machines that you add to your OpenShift Container Platform cluster. You cannot enable swap memory on these machines.

You must add RHEL compute machines to the cluster after you initialize the control plane.

### 7.2. SYSTEM REQUIREMENTS FOR RHEL COMPUTE NODES

The Red Hat Enterprise Linux (RHEL) compute, or worker, machine hosts in your OpenShift Container Platform environment must meet the following minimum hardware specifications and system-level requirements.

- You must have an active OpenShift Container Platform subscription on your Red Hat account. If you do not, contact your sales representative for more information.
- Production environments must provide compute machines to support your expected workloads. As an OpenShift Container Platform cluster administrator, you must calculate the expected workload and add about 10 percent for overhead. For production environments, allocate enough resources so that a node host failure does not affect your maximum capacity.
- Each system must meet the following hardware requirements:
  - Physical or virtual system, or an instance running on a public or private IaaS.
  - Base OS: [RHEL 7.6](#) with "Minimal" installation option.



## IMPORTANT

Only RHEL 7.6 is supported in OpenShift Container Platform 4.1. You must not upgrade your compute machines to RHEL 8.

- NetworkManager 1.0 or later.
- 1 vCPU.
- Minimum 8 GB RAM.
- Minimum 15 GB hard disk space for the file system containing **/var/**.
- Minimum 1 GB hard disk space for the file system containing **/usr/local/bin/**.
- Minimum 1 GB hard disk space for the file system containing the system's temporary directory. The system's temporary directory is determined according to the rules defined in the `tempfile` module in Python's standard library.
- Each system must meet any additional requirements for your system provider. For example, if you installed your cluster on VMware vSphere, your disks must be configured according to its [storage guidelines](#) and the **disk.enableUUID=true** attribute must be set.

## 7.3. PREPARING A RHEL COMPUTE NODE

Before you add a Red Hat Enterprise Linux (RHEL) machine to your OpenShift Container Platform cluster, you must register each host with Red Hat Subscription Manager (RHSM), attach an active OpenShift Container Platform subscription, and enable the required repositories.

1. On each host, register with RHSM:

```
# subscription-manager register --username=<user_name> --password=<password>
```

2. Pull the latest subscription data from RHSM:

```
# subscription-manager refresh
```

3. List the available subscriptions:

```
# subscription-manager list --available --matches '*OpenShift*'
```

4. In the output for the previous command, find the pool ID for an OpenShift Container Platform subscription and attach it:

```
# subscription-manager attach --pool=<pool_id>
```

5. Disable all yum repositories:

- a. Disable all the enabled RHSM repositories:

```
# subscription-manager repos --disable="**"
```

- b. List the remaining yum repositories and note their names under **repo id**, if any:



```
# yum repolist
```

- c. Use **yum-config-manager** to disable the remaining yum repositories:

```
# yum-config-manager --disable <repo_id>
```

Alternatively, disable all repositories:

```
yum-config-manager --disable \*
```

Note that this might take a few minutes if you have a large number of available repositories

6. Enable only the repositories required by OpenShift Container Platform 4.1:

```
# subscription-manager repos \
  --enable="rhel-7-server-rpms" \
  --enable="rhel-7-server-extras-rpms" \
  --enable="rhel-7-server-ose-4.1-rpms"
```

## 7.4. ADDING MORE RHEL COMPUTE MACHINES TO YOUR CLUSTER

You can add more compute machines that use Red Hat Enterprise Linux as the operating system to an OpenShift Container Platform 4.1 cluster.

### Prerequisites

- Your OpenShift Container Platform cluster already contains RHEL compute nodes.
- The **hosts** and **pull-secret.txt** files that you used to add the first RHEL compute machines to your cluster are on the machine that you use to run the playbook.
- The machine that you run the playbook on must be able to access all of the RHEL hosts. You can use any method that your company allows, including a bastion with an SSH proxy or a VPN.
- The **kubeconfig** file for the cluster and the installation program that you used to install the cluster are on the machine that you use to run the playbook.
- You must prepare the RHEL hosts for installation.
- Configure a user on the machine that you run the playbook on that has SSH access to all of the RHEL hosts.
- If you use SSH key-based authentication, you must manage the key with an SSH agent.
- Install the OpenShift Command-line Interface (CLI), commonly known as **oc**, on the machine that you run the playbook on.

### Procedure

1. Open the Ansible inventory file at `<path>/inventory/hosts` that defines your compute machine hosts and required variables.
2. Rename the **[new\_workers]** section of the file to **[workers]**.

3. Add a **[new\_workers]** section to the file and define the fully-qualified domain names for each new host. The file resembles the following example:

```
[all:vars]
ansible_user=root
#ansible_become=True

openshift_kubeconfig_path=~/.kube/config"
openshift_pull_secret_path=~/.pull-secret.txt"

[workers]
mycluster-worker-0.example.com
mycluster-worker-1.example.com

[new_workers]
mycluster-worker-2.example.com
mycluster-worker-3.example.com
```

In this example, the **mycluster-worker-0.example.com** and **mycluster-worker-1.example.com** machines are in the cluster and you add the **mycluster-worker-2.example.com** and **mycluster-worker-3.example.com** machines.

4. Run the scaleup playbook:

```
$ cd /usr/share/ansible/openshift-ansible
$ ansible-playbook -i /<path>/inventory/hosts playbooks/scaleup.yml 1
```

**1** For **<path>**, specify the path to the Ansible inventory file that you created.

## 7.5. APPROVING THE CSRS FOR YOUR MACHINES

When you add machines to a cluster, two pending certificates signing request (CSRs) are generated for each machine that you added. You must confirm that these CSRs are approved or, if necessary, approve them yourself.

### Prerequisites

- You added machines to your cluster.
- Install the **jq** package.

### Procedure

1. Confirm that the cluster recognizes the machines:

```
$ oc get nodes

NAME      STATUS   ROLES    AGE   VERSION
master-0  Ready    master   63m   v1.13.4+b626c2fe1
master-1  Ready    master   63m   v1.13.4+b626c2fe1
master-2  Ready    master   64m   v1.13.4+b626c2fe1
worker-0  NotReady worker   76s   v1.13.4+b626c2fe1
worker-1  NotReady worker   70s   v1.13.4+b626c2fe1
```

The output lists all of the machines that you created.

- Review the pending certificate signing requests (CSRs) and ensure that the you see a client and server request with **Pending** or **Approved** status for each machine that you added to the cluster:

```
$ oc get csr

NAME      AGE      REQUESTOR                                     CONDITION
csr-8b2br 15m      system:serviceaccount:openshift-machine-config-operator:node-
bootstrapper Pending 1
csr-8vnps 15m      system:serviceaccount:openshift-machine-config-operator:node-
bootstrapper Pending
csr-bfd72 5m26s    system:node:ip-10-0-50-126.us-east-2.compute.internal
Pending 2
csr-c57lv 5m26s    system:node:ip-10-0-95-157.us-east-2.compute.internal
Pending
...
```

- A client request CSR.
- A server request CSR.

In this example, two machines are joining the cluster. You might see more approved CSRs in the list.

- If the CSRs were not approved, after all of the pending CSRs for the machines you added are in **Pending** status, approve the CSRs for your cluster machines:



#### NOTE

Because the CSRs rotate automatically, approve your CSRs within an hour of adding the machines to the cluster. If you do not approve them within an hour, the certificates will rotate, and more than two certificates will be present for each node. You must approve all of these certificates. After you approve the initial CSRs, the subsequent node client CSRs are automatically approved by the cluster **kube-controller-manager**. You must implement a method of automatically approving the kubelet serving certificate requests.

- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

- <csr\_name>** is the name of a CSR from the list of current CSRs.

- If all the CSRs are valid, approve them all by running the following command:

```
$ oc get csr -ojson | jq -r '.items[] | select(.status == {} ) | .metadata.name' | xargs oc adm
certificate approve
```

## 7.6. REQUIRED PARAMETERS FOR THE ANSIBLE HOSTS FILE

You must define the following parameters in the Ansible hosts file before you add Red Hat Enterprise Linux (RHEL) compute machines to your cluster.

Parameter	Description	Values
<b>ansible_user</b>	The SSH user that allows SSH-based authentication without requiring a password. If you use SSH key-based authentication, then you must manage the key with an SSH agent.	A user name on the system. The default value is <b>root</b> .
<b>ansible_become</b>	If the values of <b>ansible_user</b> is not root, you must set <b>ansible_become</b> to <b>True</b> , and the user that you specify as the <b>ansible_user</b> must be configured for passwordless sudo access.	<b>True</b> . If the value is not <b>True</b> , do not specify and define this parameter.
<b>openshift_kubeconfig_path</b>	Specifies a path to a local directory that contains the <b>kubeconfig</b> file for your cluster.	The path and name of the configuration file.
<b>openshift_pull_secret_path</b>	Specifies a path to the text file that contains the pull secret to the image registry for your cluster. Use the pull secret that you obtained from the <a href="#">OpenShift Infrastructure Providers</a> page. This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.	The path and name of the pull secret file.

## CHAPTER 8. DEPLOYING MACHINE HEALTH CHECKS

You can configure and deploy a machine health check to automatically repair damaged machines in a machine pool.



### IMPORTANT

Machine health checks is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see <https://access.redhat.com/support/offerings/techpreview/>.

### Prerequisites

- [Enable a FeatureGate](#) so you can access Technology Preview features.



### NOTE

Turning on Technology Preview features cannot be undone and prevents upgrades.

## 8.1. ABOUT MACHINEHEALTHCHECKS

MachineHealthChecks automatically repairs unhealthy Machines in a particular MachinePool.

To monitor machine health, you create a resource to define the configuration for a controller. You set a condition to check for, such as staying in the **NotReady** status for 15 minutes or displaying a permanent condition in the node-problem-detector, and a label for the set of machines to monitor.



### NOTE

You cannot apply a MachineHealthCheck to a machine with the master role.

The controller that observes a MachineHealthCheck resource checks for the status that you defined. If a machine fails the health check, it is automatically deleted and a new one is created to take its place. When a machine is deleted, you see a **machine deleted** event. To limit disruptive impact of the machine deletion, the controller drains and deletes only one node at a time.

To stop the check, you remove the resource.

## 8.2. SAMPLE MACHINEHEALTHCHECK RESOURCE

The MachineHealthCheck resource resembles the following YAML file:

### MachineHealthCheck

```
apiVersion: healthchecking.openshift.io/v1alpha1
kind: MachineHealthCheck
metadata:
```

```

name: example 1
namespace: openshift-machine-api
Spec:
  Selector:
    matchLabels:
      machine.openshift.io/cluster-api-cluster: <cluster_name> 2
      machine.openshift.io/cluster-api-machine-role: <label> 3
      machine.openshift.io/cluster-api-machine-type: <label> 4
      machine.openshift.io/cluster-api-machineset: <cluster_name>-<label>-<AWS-zone> 5

```

- 1** Specify the name of the MachineHealthCheck to deploy. Include the name of the MachinePool to track.
- 2** Specify the name of your cluster.
- 3** **4** Specify a label for the MachinePool that you want to check.
- 5** Specify the MachineSet to track in **<cluster\_name>-<label>-<AWS-zone>** format. For example, **prod-node-us-east-1a**.

### 8.3. CREATING A MACHINEHEALTHCHECK RESOURCE

You can create a MachineHealthCheck resource for all MachinePools in your cluster except the **master** pool.

#### Prerequisites

- Install the **oc** command line interface.

#### Procedure

1. Create a **healthcheck.yml** file that contains the definition of your MachineHealthCheck.
2. Apply the **healthcheck.yml** file to your cluster:

```
$ oc apply -f healthcheck.yml
```