



# OpenShift Container Platform 4.2

## Serverless applications

OpenShift Serverless installation, usage, and release notes



# OpenShift Container Platform 4.2 Serverless applications

---

OpenShift Serverless installation, usage, and release notes

## Legal Notice

Copyright © 2020 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 information on how to use OpenShift Serverless in OpenShift Container Platform 4.2

## Table of Contents

<b>CHAPTER 1. GETTING STARTED WITH OPENSIFT SERVERLESS</b> .....	<b>4</b>
1.1. HOW OPENSIFT SERVERLESS WORKS	4
1.2. APPLICATIONS ON OPENSIFT SERVERLESS	4
<b>CHAPTER 2. OPENSIFT SERVERLESS PRODUCT ARCHITECTURE</b> .....	<b>5</b>
2.1. KNATIVE SERVING	5
2.1.1. Knative Serving components	5
2.2. KNATIVE CLIENT	5
2.3. KNATIVE EVENTING	6
<b>CHAPTER 3. INSTALLING OPENSIFT SERVERLESS</b> .....	<b>7</b>
3.1. CLUSTER SIZE REQUIREMENTS	7
3.1.1. Scaling a MachineSet manually	7
3.2. INSTALLING THE OPENSIFT SERVERLESS OPERATOR	8
3.3. INSTALLING KNATIVE SERVING	8
3.4. UNINSTALLING KNATIVE SERVING	9
3.5. DELETING THE OPENSIFT SERVERLESS OPERATOR	10
3.6. DELETING KNATIVE SERVING CRDS FROM THE OPERATOR	10
<b>CHAPTER 4. GETTING STARTED WITH KNATIVE SERVICES</b> .....	<b>11</b>
4.1. CREATING A KNATIVE SERVICE	11
4.2. DEPLOYING A SERVERLESS APPLICATION	11
4.3. CONNECTING KNATIVE SERVICES TO EXISTING KUBERNETES DEPLOYMENTS	12
<b>CHAPTER 5. MONITORING OPENSIFT SERVERLESS COMPONENTS</b> .....	<b>13</b>
5.1. CONFIGURING CLUSTER FOR APPLICATION MONITORING	13
5.2. VERIFYING AN OPENSIFT CONTAINER PLATFORM MONITORING INSTALLATION FOR USE WITH KNATIVE SERVING	13
5.3. MONITORING KNATIVE SERVING USING THE OPENSIFT CONTAINER PLATFORM MONITORING STACK	14
<b>CHAPTER 6. USING METERING WITH OPENSIFT SERVERLESS</b> .....	<b>15</b>
6.1. INSTALLING METERING	15
6.2. DATASOURCES FOR KNATIVE SERVING METERING	15
6.2.1. Datasource for CPU usage in Knative Serving	15
6.2.2. Datasource for memory usage in Knative Serving	15
6.2.3. Applying Datasources for Knative Serving metering	16
6.3. QUERIES FOR KNATIVE SERVING METERING	16
6.3.1. Query for CPU usage in Knative Serving	16
6.3.2. Query for memory usage in Knative Serving	17
6.3.3. Applying Queries for Knative Serving metering	18
6.4. METERING REPORTS FOR KNATIVE SERVING	19
6.4.1. Running a metering report	19
<b>CHAPTER 7. CLUSTER LOGGING WITH OPENSIFT SERVERLESS</b> .....	<b>20</b>
7.1. ABOUT CLUSTER LOGGING	20
7.2. ABOUT DEPLOYING AND CONFIGURING CLUSTER LOGGING	20
7.2.1. Configuring and Tuning Cluster Logging	20
7.2.2. Sample modified Cluster Logging Custom Resource	22
7.3. USING CLUSTER LOGGING TO FIND LOGS FOR KNATIVE SERVING COMPONENTS	23
7.4. USING CLUSTER LOGGING TO FIND LOGS FOR SERVICES DEPLOYED WITH KNATIVE SERVING	24
<b>CHAPTER 8. CONFIGURING KNATIVE SERVING AUTOSCALING</b> .....	<b>25</b>

8.1. CONFIGURING CONCURRENT REQUESTS FOR KNATIVE SERVING AUTOSCALING	25
8.1.1. Configuring concurrent requests using the target annotation	26
8.1.2. Configuring concurrent requests using the containerConcurrency field	26
8.2. CONFIGURING SCALE BOUNDS KNATIVE SERVING AUTOSCALING	26
<b>CHAPTER 9. GETTING STARTED WITH KNATIVE CLIENT</b> .....	<b>28</b>
9.1. BEFORE YOU BEGIN	28
9.2. INSTALLING KNATIVE CLIENT	28
9.2.1. Installing the kn CLI using the OpenShift Container Platform web console	28
9.2.2. Installing the kn CLI for Linux using an RPM	29
9.2.3. Installing the kn CLI for Linux	29
9.2.4. Installing the kn CLI for macOS	30
9.2.5. Installing the kn CLI for Windows	30
9.3. BASIC WORKFLOW USING KNATIVE CLIENT	30
9.4. AUTOSCALING WORKFLOW USING KNATIVE CLIENT	32
9.5. TRAFFIC SPLITTING USING KNATIVE CLIENT	32
9.5.1. Assigning tag revisions	33
9.5.2. Unassigning tag revisions	34
9.5.3. Traffic flag operation precedence	34
9.5.4. Traffic splitting flags	35
<b>CHAPTER 10. OPENSIFT SERVERLESS RELEASE NOTES</b> .....	<b>36</b>
10.1. GETTING SUPPORT	36
10.2. RELEASE NOTES FOR RED HAT OPENSIFT SERVERLESS TECHNOLOGY PREVIEW 1.4.0	36
10.2.1. New features	36
10.2.2. Fixed issues	37
10.2.3. Known issues	37
10.3. RELEASE NOTES FOR RED HAT OPENSIFT SERVERLESS TECHNOLOGY PREVIEW 1.3.0	37
10.3.1. New features	37
10.3.2. Fixed issues	37
10.3.3. Known issues	37
10.4. RELEASE NOTES FOR RED HAT OPENSIFT SERVERLESS TECHNOLOGY PREVIEW 1.2.0	38
10.4.1. New features	38
10.4.2. Fixed issues	38
10.4.3. Known issues	38
10.5. RELEASE NOTES FOR RED HAT OPENSIFT SERVERLESS TECHNOLOGY PREVIEW 1.1.0	39
10.5.1. New features	39
10.5.2. Fixed issues	39
10.5.3. Known issues	39
10.6. RELEASE NOTES FOR RED HAT OPENSIFT SERVERLESS TECHNOLOGY PREVIEW 1.0.0	40
10.6.1. New features	40
10.6.2. Known issues	40
10.7. ADDITIONAL RESOURCES	40



# CHAPTER 1. GETTING STARTED WITH OPENSHIFT SERVERLESS



## IMPORTANT

You are viewing documentation for a release of Red Hat OpenShift Serverless that is no longer supported. Red Hat OpenShift Serverless is currently supported on OpenShift Container Platform 4.3 and newer.



## IMPORTANT

OpenShift Serverless 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/>.

OpenShift Serverless simplifies the process of delivering code from development into production by reducing the need for infrastructure set up or back-end development by developers.

## 1.1. HOW OPENSHIFT SERVERLESS WORKS

Developers on OpenShift Serverless can use the provided Kubernetes-native APIs, as well as familiar languages and frameworks, to deploy applications and container workloads. For information about installing OpenShift Serverless, see [Installing OpenShift Serverless](#).

OpenShift Serverless on OpenShift Container Platform enables stateful, stateless, and serverless workloads to all run on a single multi-cloud container platform with automated operations. Developers can use a single platform for hosting their microservices, legacy, and serverless applications.

OpenShift Serverless is based on the open source Knative project, which provides portability and consistency across hybrid and multi-cloud environments by enabling an enterprise-grade serverless platform.

## 1.2. APPLICATIONS ON OPENSHIFT SERVERLESS

Applications are created using Custom Resource Definitions (CRDs) and associated controllers in Kubernetes, and are packaged as OCI compliant Linux containers that can be run anywhere.

To deploy applications in OpenShift Serverless, you must create Knative Services. For more information see [Getting started with Knative Services](#).



## CHAPTER 2. OPENSIFT SERVERLESS PRODUCT ARCHITECTURE



### IMPORTANT

You are viewing documentation for a release of Red Hat OpenShift Serverless that is no longer supported. Red Hat OpenShift Serverless is currently supported on OpenShift Container Platform 4.3 and newer.

## 2.1. KNATIVE SERVING

Knative Serving on OpenShift Container Platform builds on Kubernetes and Istio to support deploying and serving serverless applications.

It creates a set of Kubernetes Custom Resource Definitions (CRDs) that are used to define and control the behavior of serverless workloads on an OpenShift Container Platform cluster.

These CRDs can be used as building blocks to address complex use cases, such as rapid deployment of serverless containers, automatic scaling of Pods, routing and network programming for Istio components, or viewing point-in-time snapshots of deployed code and configurations.

### 2.1.1. Knative Serving components

The components described in this section are the resources that Knative Serving requires to be configured and run correctly.

#### Knative service resource

The **service.serving.knative.dev** resource automatically manages the whole lifecycle of a serverless workload on a cluster. It controls the creation of other objects to ensure that an app has a route, a configuration, and a new revision for each update of the service. Services can be defined to always route traffic to the latest revision or to a pinned revision.

#### Knative route resource

The **route.serving.knative.dev** resource maps a network endpoint to one or more Knative revisions. You can manage the traffic in several ways, including fractional traffic and named routes.

#### Knative configuration resource

The **configuration.serving.knative.dev** resource maintains the required state for your deployment. Modifying a configuration creates a new revision.

#### Knative revision resource

The **revision.serving.knative.dev** resource is a point-in-time snapshot of the code and configuration for each modification made to the workload. Revisions are immutable objects and can be retained for as long as needed. Cluster administrators can modify the **revision.serving.knative.dev** resource to enable automatic scaling of Pods in your OpenShift Container Platform cluster.

## 2.2. KNATIVE CLIENT

The Knative Client (**kn**) extends the functionality of the **oc** or **kubectl** tools to enable interaction with Knative components on OpenShift Container Platform. **kn** allows developers to deploy and manage applications without editing YAML files directly.

## 2.3. KNATIVE EVENTING

A developer preview version of Knative Eventing is available for use with OpenShift Serverless. However, this is not included in the OpenShift Serverless Operator and is not currently supported as part of this Technology Preview. For more information about Knative Eventing, including installation instructions and samples, see the [Knative Eventing on OpenShift Container Platform](#) documentation.

## CHAPTER 3. INSTALLING OPENSIFT SERVERLESS



### IMPORTANT

You are viewing documentation for a release of Red Hat OpenShift Serverless that is no longer supported. Red Hat OpenShift Serverless is currently supported on OpenShift Container Platform 4.3 and newer.



### IMPORTANT

OpenShift Serverless is not tested or supported for installation in a restricted network environment.

### 3.1. CLUSTER SIZE REQUIREMENTS

The cluster must be sized appropriately to ensure that OpenShift Serverless can run correctly. You can use the MachineSet API to manually scale your cluster up to the desired size.

An OpenShift cluster with 10 CPUs and 40 GB memory is the minimum requirement for getting started with your first serverless application. This usually means you must scale up one of the default MachineSets by two additional machines.



### NOTE

For this configuration, the requirements depend on the deployed applications. By default, each pod requests ~400m of CPU and recommendations are based on this value. In the given recommendation, an application can scale up to 10 replicas. Lowering the actual CPU request of the application further pushes the boundary.



### NOTE

The numbers given only relate to the pool of worker machines of the OpenShift cluster. Master nodes are not used for general scheduling and are omitted.

For more advanced use-cases, such as using OpenShift logging, monitoring, metering, and tracing, you must deploy more resources. Recommended requirements for such use-cases are 24 vCPUs and 96GB of memory.

#### Additional resources

For more information on using the MachineSet API, see [Creating MachineSets](#).

#### 3.1.1. 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 in to **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 **0** unless you first relocate the router pods.

## 3.2. INSTALLING THE OPENSIFT SERVERLESS OPERATOR

The OpenShift Serverless Operator can be installed using the OpenShift Container Platform instructions for installing Operators.

You can install the OpenShift Serverless Operator in the host cluster by following the OpenShift Container Platform instructions on installing an Operator.



### NOTE

The OpenShift Serverless Operator only works for OpenShift Container Platform versions 4.1.13 and later.

For details, see the OpenShift Container Platform documentation on [adding Operators to a cluster](#).



### IMPORTANT

The OpenShift Serverless Operator automatically installs the Service Mesh Operator. If you already have a community version of Maistra installed, this will cause a conflict with the OpenShift Serverless Operator Service Mesh auto-install. In this case, the already existing community version of Maistra will be used instead.

## 3.3. INSTALLING KNATIVE SERVING

You must create a **KnativeServing** object to install Knative Serving using the OpenShift Serverless Operator.



## IMPORTANT

You must create the **KnativeService** object in the **knative-serving** namespace, as shown in the sample YAML, or it is ignored.

### Sample serving.yaml

```
apiVersion: v1
kind: Namespace
metadata:
  name: knative-serving
---
apiVersion: operator.knative.dev/v1alpha1
kind: KnativeService
metadata:
  name: knative-serving
  namespace: knative-serving
```

### Prerequisite

- An account with cluster administrator access.
- Installed OpenShift Serverless Operator.

### Procedure

1. Copy the sample YAML file into **serving.yaml** and apply it using:

```
$ oc apply -f serving.yaml
```

2. Verify the installation is complete by using the command:

```
$ oc get knativeserving.operator.knative.dev/knative-serving -n knative-serving --
template='{{range .status.conditions}}{{printf "%s=%s\n" .type .status}}{{end}}'
```

Results should be similar to:

```
DeploymentsAvailable=True
InstallSucceeded=True
Ready=True
```

## 3.4. UNINSTALLING KNATIVE SERVING

To uninstall Knative Serving, you must remove its custom resource and delete the **knative-serving** namespace.

### Prerequisite

- Installed Knative Serving

### Procedure

1. To remove Knative Serving, use the following command:

```
$ oc delete knativeserving.operator.knative.dev knative-serving -n knative-serving
```

2. After the command has completed and all pods have been removed from the **knative-serving** namespace, delete the namespace by using the command:

```
$ oc delete namespace knative-serving
```

### 3.5. DELETING THE OPENSIFT SERVERLESS OPERATOR

You can remove the OpenShift Serverless Operator from the host cluster by following the OpenShift Container Platform instructions on deleting an Operator.

For details, see the OpenShift Container Platform documentation on [deleting Operators from a cluster](#).

### 3.6. DELETING KNATIVE SERVING CRDS FROM THE OPERATOR

After uninstalling the OpenShift Serverless Operator, the Operator CRDs and API services remain on the cluster. Use this procedure to completely uninstall the remaining components.

#### Prerequisite

- You have uninstalled Knative Serving and removed the OpenShift Serverless Operator using the previous procedure.

#### Procedure

1. Run the following command to delete the remaining Knative Serving CRDs:

```
$ oc delete crd knativeservings.operator.knative.dev
```

## CHAPTER 4. GETTING STARTED WITH KNATIVE SERVICES



### IMPORTANT

You are viewing documentation for a release of Red Hat OpenShift Serverless that is no longer supported. Red Hat OpenShift Serverless is currently supported on OpenShift Container Platform 4.3 and newer.

Knative services are Kubernetes services that a user creates to deploy a serverless application. Each Knative service is defined by a route and a configuration, contained in a **.yaml** file.

### 4.1. CREATING A KNATIVE SERVICE

To create a service, you must create the **service.yaml** file.

You can copy the sample below. This sample will create a sample go-lang application called **helloworld-go** and allows you to specify the image for that application.

```
apiVersion: serving.knative.dev/v1alpha1 1
kind: Service
metadata:
  name: helloworld-go 2
  namespace: default 3
spec:
  template:
    spec:
      containers:
        - image: gcr.io/knative-samples/helloworld-go 4
          env:
            - name: TARGET 5
              value: "Go Sample v1"
```

- 1** Current version of Knative
- 2** The name of the application
- 3** The namespace the application will use
- 4** The URL to the image of the application
- 5** The environment variable printed out by the sample application

### 4.2. DEPLOYING A SERVERLESS APPLICATION

To deploy a serverless application, you must apply the **service.yaml** file.

#### Procedure

1. Navigate to the directory where the **service.yaml** file is contained.
2. Deploy the application by applying the **service.yaml** file.

```
$ oc apply --filename service.yaml
```

Now that service has been created and the application has been deployed, Knative will create a new immutable revision for this version of the application.

Knative will also perform network programming to create a route, ingress, service, and load balancer for your application, and will automatically scale your pods up and down based on traffic, including inactive pods.



#### NOTE

The first time that a Knative service is created in a namespace, that namespace will automatically receive a new networking configuration. This might cause the initial service to take longer than is usually required for a service to become ready.

If the namespace has no existing NetworkPolicy configuration, an "allow all" type policy will be applied automatically. This policy will be removed automatically if all Knative Services are removed from that namespace and no other NetworkPolicy configurations have been applied.

## 4.3. CONNECTING KNATIVE SERVICES TO EXISTING KUBERNETES DEPLOYMENTS

Knative Services can call a Kubernetes deployment in any namespace, provided that there are no existing additional network barriers.

A Kubernetes deployment can call a Knative Service if:

- The Kubernetes deployment is in the same namespace as the target Knative Service.
- The Kubernetes deployment is in a namespace that was manually added to the ServiceMeshMemberRoll in **knative-serving-ingress**.
- The Kubernetes deployment uses the target Knative Service's public URL.



#### NOTE

Knative Services are accessed using a public URL by default. The target Knative Service must not be configured as a private, **cluster-local** visibility service if you want to connect it to your existing Kubernetes deployment using a public URL.



## CHAPTER 5. MONITORING OPENSIFT SERVERLESS COMPONENTS



### IMPORTANT

You are viewing documentation for a release of Red Hat OpenShift Serverless that is no longer supported. Red Hat OpenShift Serverless is currently supported on OpenShift Container Platform 4.3 and newer.

As a cluster administrator, you can deploy the OpenShift Container Platform monitoring stack and monitor the metrics of OpenShift Serverless components.

When using the OpenShift Serverless Operator, the required ServiceMonitor objects are created automatically for monitoring the deployed components.

OpenShift Serverless components, such as Knative Serving, expose metrics data. Administrators can monitor this data by using the OpenShift Container Platform web console.

### 5.1. CONFIGURING CLUSTER FOR APPLICATION MONITORING

Before application developers can monitor their applications, the human operator of the cluster needs to configure the cluster accordingly. This procedure shows how to.

#### Prerequisites

- You must log in as a user that belongs to a role with administrative privileges for the cluster.

#### Procedure

- In the OpenShift Container Platform web console, navigate to the **Operators → OperatorHub** page and install the Prometheus Operator in the namespace where your application is.
- Navigate to the **Operators → Installed Operators** page and install Prometheus, Alertmanager, Prometheus Rule, and Service Monitor in the same namespace.

### 5.2. VERIFYING AN OPENSIFT CONTAINER PLATFORM MONITORING INSTALLATION FOR USE WITH KNATIVE SERVING

Manual configuration for monitoring by an administrator is not required, but you can carry out these steps to verify that monitoring is installed correctly.

#### Procedure

- Verify that the ServiceMonitor objects are deployed.

```
$ oc get servicemonitor -n knative-serving
NAME      AGE
activator 11m
autoscaler 11m
controller 11m
```

2. Verify that the **openshift.io/cluster-monitoring=true** label has been added to the Knative Serving namespace:

```
$ oc get namespace knative-serving --show-labels
NAME          STATUS AGE LABELS
knative-serving Active 4d  istio-injection=enabled,openshift.io/cluster-
monitoring=true,serving.knative.dev/release=v0.7.0
```

### 5.3. MONITORING KNATIVE SERVING USING THE OPENSIFT CONTAINER PLATFORM MONITORING STACK

This section provides example instructions for the visualization of Knative Serving Pod autoscaling metrics by using the OpenShift Container Platform monitoring tools.

#### Prerequisites

- You must have the OpenShift Container Platform monitoring stack installed.

#### Procedure

1. Navigate to the OpenShift Container Platform web console and authenticate.
2. Navigate to **Monitoring** → **Metrics**.
3. Enter the **Expression** and select **Run queries**. To monitor Knative Serving autoscaler Pods, use this example expression.

```
autoscaler_actual_pods
```

You will now see monitoring information for the Knative Serving autoscaler Pods in the console.

# CHAPTER 6. USING METERING WITH OPENSIFT SERVERLESS



## IMPORTANT

You are viewing documentation for a release of Red Hat OpenShift Serverless that is no longer supported. Red Hat OpenShift Serverless is currently supported on OpenShift Container Platform 4.3 and newer.

As a cluster administrator, you can use metering to analyze what is happening in your OpenShift Serverless cluster.

For more information about metering on OpenShift Container Platform, see [About metering](#).

## 6.1. INSTALLING METERING

For information about installing metering on OpenShift Container Platform, see [Installing Metering](#).

## 6.2. DATASOURCES FOR KNATIVE SERVING METERING

The following **ReportDataSources** are examples of how Knative Serving can be used with OpenShift Container Platform metering.

### 6.2.1. Datasource for CPU usage in Knative Serving

This datasource provides the accumulated CPU seconds used per Knative service over the report time period.

#### YAML file

```
apiVersion: metering.openshift.io/v1
kind: ReportDataSource
metadata:
  name: knative-service-cpu-usage
spec:
  prometheusMetricsImporter:
    query: >
      sum
      by(namespace,
        label_serving_knative_dev_service,
        label_serving_knative_dev_revision)
      (
label_replace(rate(container_cpu_usage_seconds_total{container_name!="POD",container_name!="",pod_name!=""}[1m]), "pod", "$1", "pod_name", "(.*)")
*
      on(pod, namespace)
      group_left(label_serving_knative_dev_service, label_serving_knative_dev_revision)
      kube_pod_labels{label_serving_knative_dev_service!=""}
      )
```

### 6.2.2. Datasource for memory usage in Knative Serving

This datasource provides the average memory consumption per Knative service over the report time period.

### YAML file

```
apiVersion: metering.openshift.io/v1
kind: ReportDataSource
metadata:
  name: knative-service-memory-usage
spec:
  prometheusMetricsImporter:
    query: >
      sum
        by(namespace,
          label_serving_knative_dev_service,
          label_serving_knative_dev_revision)
        (
          label_replace(container_memory_usage_bytes{container_name!="POD",
container_name!="",pod_name!="}, "pod", "$1", "pod_name", "(.*)")
          *
          on(pod, namespace)
          group_left(label_serving_knative_dev_service, label_serving_knative_dev_revision)
          kube_pod_labels{label_serving_knative_dev_service!=""}
        )
```

### 6.2.3. Applying Datasources for Knative Serving metering

You can apply the **ReportDataSources** by using the following command:

```
$ oc apply -f <datasource-name>.yaml
```

#### Example

```
$ oc apply -f knative-service-memory-usage.yaml
```

## 6.3. QUERIES FOR KNATIVE SERVING METERING

The following **ReportQuery** resources reference the example **DataSources** provided.

### 6.3.1. Query for CPU usage in Knative Serving

#### YAML file

```
apiVersion: metering.openshift.io/v1
kind: ReportQuery
metadata:
  name: knative-service-cpu-usage
spec:
  inputs:
    - name: ReportingStart
      type: time
    - name: ReportingEnd
```

```

  type: time
- default: knative-service-cpu-usage
  name: KnativeServiceCpuUsageDataSource
  type: ReportDataSource
columns:
- name: period_start
  type: timestamp
  unit: date
- name: period_end
  type: timestamp
  unit: date
- name: namespace
  type: varchar
  unit: kubernetes_namespace
- name: service
  type: varchar
- name: data_start
  type: timestamp
  unit: date
- name: data_end
  type: timestamp
  unit: date
- name: service_cpu_seconds
  type: double
  unit: cpu_core_seconds
query: |
  SELECT
    timestamp '{| default .Report.ReportingStart .Report.Inputs.ReportingStart | prestoTimestamp |}'
AS period_start,
    timestamp '{| default .Report.ReportingEnd .Report.Inputs.ReportingEnd | prestoTimestamp |}' AS
period_end,
    labels['namespace'] as project,
    labels['label_serving_knative_dev_service'] as service,
    min("timestamp") as data_start,
    max("timestamp") as data_end,
    sum(amount * "timeprecision") AS service_cpu_seconds
FROM {| dataSourceTableName .Report.Inputs.KnativeServiceCpuUsageDataSource |}
WHERE "timestamp" >= timestamp '{| default .Report.ReportingStart .Report.Inputs.ReportingStart
| prestoTimestamp |}'
  AND "timestamp" < timestamp '{| default .Report.ReportingEnd .Report.Inputs.ReportingEnd |
prestoTimestamp |}'
  GROUP BY labels['namespace'],labels['label_serving_knative_dev_service']

```

### 6.3.2. Query for memory usage in Knative Serving

#### YAML file

```

apiVersion: metering.openshift.io/v1
kind: ReportQuery
metadata:
  name: knative-service-memory-usage
spec:
  inputs:
  - name: ReportingStart
    type: time

```

```

- name: ReportingEnd
  type: time
- default: knative-service-memory-usage
  name: KnativeServiceMemoryUsageDataSource
  type: ReportDataSource
columns:
- name: period_start
  type: timestamp
  unit: date
- name: period_end
  type: timestamp
  unit: date
- name: namespace
  type: varchar
  unit: kubernetes_namespace
- name: service
  type: varchar
- name: data_start
  type: timestamp
  unit: date
- name: data_end
  type: timestamp
  unit: date
- name: service_usage_memory_byte_seconds
  type: double
  unit: byte_seconds
query: |
  SELECT
    timestamp '{| default .Report.ReportingStart .Report.Inputs.ReportingStart| prestoTimestamp |}'
AS period_start,
    timestamp '{| default .Report.ReportingEnd .Report.Inputs.ReportingEnd | prestoTimestamp |}' AS
period_end,
    labels['namespace'] as project,
    labels['label_serving_knative_dev_service'] as service,
    min("timestamp") as data_start,
    max("timestamp") as data_end,
    sum(amount * "timeprecision") AS service_usage_memory_byte_seconds
FROM {| dataSourceTableName .Report.Inputs.KnativeServiceMemoryUsageDataSource |}
WHERE "timestamp" >= timestamp '{| default .Report.ReportingStart .Report.Inputs.ReportingStart
| prestoTimestamp |}'
AND "timestamp" < timestamp '{| default .Report.ReportingEnd .Report.Inputs.ReportingEnd |
prestoTimestamp |}'
GROUP BY labels['namespace'],labels['label_serving_knative_dev_service']

```

### 6.3.3. Applying Queries for Knative Serving metering

You can apply the **ReportQuery** by using the following command:

```
$ oc apply -f <query-name>.yaml
```

#### Example

```
$ oc apply -f knative-service-memory-usage.yaml
```

## 6.4. METERING REPORTS FOR KNATIVE SERVING

You can run metering reports against Knative Serving by creating **Report** resources. Before you run a report, you must modify the input parameter within the **Report** resource to specify the start and end dates of the reporting period.

### YAML file

```
apiVersion: metering.openshift.io/v1
kind: Report
metadata:
  name: knative-service-cpu-usage
spec:
  reportingStart: '2019-06-01T00:00:00Z' 1
  reportingEnd: '2019-06-30T23:59:59Z' 2
  query: knative-service-cpu-usage 3
runImmediately: true
```

- 1 Start date of the report, in ISO 8601 format.
- 2 End date of the report, in ISO 8601 format.
- 3 Either **knative-service-cpu-usage** for CPU usage report or **knative-service-memory-usage** for a memory usage report.

### 6.4.1. Running a metering report

Once you have provided the input parameters, you can run the report using the command:


```
$ oc apply -f <report-name>.yaml
```

You can then check the report as shown in the following example:

```
$ kubectl get report
```

NAME	QUERY	SCHEDULE	RUNNING	FAILED	LAST REPORT
knative-service-cpu-usage	knative-service-cpu-usage		Finished		2019-06-30T23:59:59Z
					10h

## CHAPTER 7. CLUSTER LOGGING WITH OPENSIFT SERVERLESS



### IMPORTANT

You are viewing documentation for a release of Red Hat OpenShift Serverless that is no longer supported. Red Hat OpenShift Serverless is currently supported on OpenShift Container Platform 4.3 and newer.

### 7.1. ABOUT CLUSTER LOGGING

OpenShift Container Platform cluster administrators can deploy cluster logging using a few CLI commands and the OpenShift Container Platform web console to install the Elasticsearch Operator and Cluster Logging Operator. When the operators are installed, create a Cluster Logging Custom Resource (CR) to schedule cluster logging pods and other resources necessary to support cluster logging. The operators are responsible for deploying, upgrading, and maintaining cluster logging.

You can configure cluster logging by modifying the Cluster Logging Custom Resource (CR), named **instance**. The CR defines a complete cluster logging deployment that includes all the components of the logging stack to collect, store and visualize logs. The Cluster Logging Operator watches the **ClusterLogging** Custom Resource and adjusts the logging deployment accordingly.

Administrators and application developers can view the logs of the projects for which they have view access.

### 7.2. ABOUT DEPLOYING AND CONFIGURING CLUSTER LOGGING

OpenShift Container Platform cluster logging is designed to be used with the default configuration, which is tuned for small to medium sized OpenShift Container Platform clusters.

The installation instructions that follow include a sample Cluster Logging Custom Resource (CR), which you can use to create a cluster logging instance and configure your cluster logging deployment.

If you want to use the default cluster logging install, you can use the sample CR directly.

If you want to customize your deployment, make changes to the sample CR as needed. The following describes the configurations you can make when installing your cluster logging instance or modify after installation. See the Configuring sections for more information on working with each component, including modifications you can make outside of the Cluster Logging Custom Resource.

#### 7.2.1. Configuring and Tuning Cluster Logging

You can configure your cluster logging environment by modifying the Cluster Logging Custom Resource deployed in the **openshift-logging** project.

You can modify any of the following components upon install or after install:

##### Memory and CPU

You can adjust both the CPU and memory limits for each component by modifying the **resources** block with valid memory and CPU values:

```
spec:
  logStore:
```



```

elasticsearch:
  resources:
    limits:
      cpu:
      memory:
    requests:
      cpu: 1
      memory: 16Gi
  type: "elasticsearch"
collection:
  logs:
    fluentd:
      resources:
        limits:
          cpu:
          memory:
        requests:
          cpu:
          memory:
      type: "fluentd"
  visualization:
    kibana:
      resources:
        limits:
          cpu:
          memory:
        requests:
          cpu:
          memory:
      type: kibana
  curation:
    curator:
      resources:
        limits:
          memory: 200Mi
        requests:
          cpu: 200m
          memory: 200Mi
      type: "curator"

```

### Elasticsearch storage

You can configure a persistent storage class and size for the Elasticsearch cluster using the **storageClass name** and **size** parameters. The Cluster Logging Operator creates a **PersistentVolumeClaim** for each data node in the Elasticsearch cluster based on these parameters.

```

spec:
  logStore:
    type: "elasticsearch"
  elasticsearch:
    nodeCount: 3
    storage:
      storageClassName: "gp2"
      size: "200G"

```

This example specifies each data node in the cluster will be bound to a **PersistentVolumeClaim** that requests "200G" of "gp2" storage. Each primary shard will be backed by a single replica.



## NOTE

Omitting the **storage** block results in a deployment that includes ephemeral storage only.

```
spec:
  logStore:
    type: "elasticsearch"
  elasticsearch:
    nodeCount: 3
    storage: {}
```

## Elasticsearch replication policy

You can set the policy that defines how Elasticsearch shards are replicated across data nodes in the cluster:

- **FullRedundancy**. The shards for each index are fully replicated to every data node.
- **MultipleRedundancy**. The shards for each index are spread over half of the data nodes.
- **SingleRedundancy**. A single copy of each shard. Logs are always available and recoverable as long as at least two data nodes exist.
- **ZeroRedundancy**. No copies of any shards. Logs may be unavailable (or lost) in the event a node is down or fails.

## Curator schedule

You specify the schedule for Curator in the [cron format](#).

```
spec:
  curation:
    type: "curator"
  resources:
    curator:
      schedule: "30 3 * * *"
```

## 7.2.2. Sample modified Cluster Logging Custom Resource

The following is an example of a Cluster Logging Custom Resource modified using the options previously described.

### Sample modified Cluster Logging Custom Resource

```
apiVersion: "logging.openshift.io/v1"
kind: "ClusterLogging"
metadata:
  name: "instance"
  namespace: "openshift-logging"
spec:
  managementState: "Managed"
```

```

logStore:
  type: "elasticsearch"
  elasticsearch:
    nodeCount: 2
    resources:
      limits:
        memory: 2Gi
      requests:
        cpu: 200m
        memory: 2Gi
    storage: {}
    redundancyPolicy: "SingleRedundancy"
visualization:
  type: "kibana"
  kibana:
    resources:
      limits:
        memory: 1Gi
      requests:
        cpu: 500m
        memory: 1Gi
    replicas: 1
curation:
  type: "curator"
  curator:
    resources:
      limits:
        memory: 200Mi
      requests:
        cpu: 200m
        memory: 200Mi
    schedule: "*/5 * * * *"
collection:
  logs:
    type: "fluentd"
    fluentd:
      resources:
        limits:
          memory: 1Gi
        requests:
          cpu: 200m
          memory: 1Gi

```

## 7.3. USING CLUSTER LOGGING TO FIND LOGS FOR KNATIVE SERVING COMPONENTS


### Procedure

1. To open the Kibana UI, the visualization tool for Elasticsearch, use the following command to get the Kibana route:

```
$ oc -n openshift-logging get route kibana
```

2. Use the route's URL to navigate to the Kibana dashboard and log in.

3. Ensure the index is set to **.all**. If the index is not set to **.all**, only the OpenShift system logs will be listed.
4. You can filter the logs by using the **knative-serving** namespace. Enter **kubernetes.namespace\_name:knative-serving** in the search box to filter results.



#### NOTE

Knative Serving uses structured logging by default. You can enable the parsing of these logs by customizing the cluster logging Fluentd settings. This makes the logs more searchable and enables filtering on the log level to quickly identify issues.

## 7.4. USING CLUSTER LOGGING TO FIND LOGS FOR SERVICES DEPLOYED WITH KNATIVE SERVING


With OpenShift Cluster Logging, the logs that your applications write to the console are collected in Elasticsearch. The following procedure outlines how to apply these capabilities to applications deployed by using Knative Serving.

### Procedure

1. Use the following command to find the URL to Kibana:

```
$ oc -n cluster-logging get route kibana`
```

2. Enter the URL in your browser to open the Kibana UI.
3. Ensure the index is set to **.all**. If the index is not set to **.all**, only the OpenShift system logs will be listed.
4. Filter the logs by using the Kubernetes namespace your service is deployed in. Add a filter to identify the service itself: **kubernetes.namespace\_name:default AND kubernetes.labels.serving\_knative\_dev\service:{SERVICE\_NAME}**.



#### NOTE

You can also filter by using **/configuration** or **/revision**.

5. You can narrow your search by using **kubernetes.container\_name:<user-container>** to only display the logs generated by your application. Otherwise, you will see logs from the queue-proxy.



#### NOTE

Use JSON-based structured logging in your application to allow for the quick filtering of these logs in production environments.

## CHAPTER 8. CONFIGURING KNATIVE SERVING AUTOSCALING



### IMPORTANT

You are viewing documentation for a release of Red Hat OpenShift Serverless that is no longer supported. Red Hat OpenShift Serverless is currently supported on OpenShift Container Platform 4.3 and newer.

OpenShift Serverless provides capabilities for automatic Pod scaling, including scaling inactive Pods to zero, by enabling the Knative Serving autoscaling system in an OpenShift Container Platform cluster.

To enable autoscaling for Knative Serving, you must configure concurrency and scale bounds in the revision template.



### NOTE

Any limits or targets set in the revision template are measured against a single instance of your application. For example, setting the **target** annotation to **50** will configure the autoscaler to scale the application so that each instance of it will handle 50 requests at a time.

## 8.1. CONFIGURING CONCURRENT REQUESTS FOR KNATIVE SERVING AUTOSCALING

You can specify the number of concurrent requests that should be handled by each instance of an application (revision container) by adding the **target** annotation or the **containerConcurrency** field in the revision template.

Here is an example of **target** being used in a revision template:

```
apiVersion: serving.knative.dev/v1alpha1
kind: Service
metadata:
  name: myapp
spec:
  template:
    metadata:
      annotations:
        autoscaling.knative.dev/target: 50
    spec:
      containers:
        - image: myimage
```

Here is an example of **containerConcurrency** being used in a revision template:

```
apiVersion: serving.knative.dev/v1alpha1
kind: Service
metadata:
  name: myapp
spec:
  template:
```

```


metadata:
  annotations:
spec:
  containerConcurrency: 100
  containers:
  - image: myimage

```

Adding a value for both **target** and **containerConcurrency** will target the **target** number of concurrent requests, but impose a hard limit of the **containerConcurrency** number of requests.

For example, if the **target** value is 50 and the **containerConcurrency** value is 100, the targeted number of requests will be 50, but the hard limit will be 100.

If the **containerConcurrency** value is less than the **target** value, the **target** value will be tuned down, since there is no need to target more requests than the number that can actually be handled.



#### NOTE

**containerConcurrency** should only be used if there is a clear need to limit how many requests reach the application at a given time. Using **containerConcurrency** is only advised if the application needs to have an enforced constraint of concurrency.

### 8.1.1. Configuring concurrent requests using the target annotation

The default target for the number of concurrent requests is **100**, but you can override this value by adding or modifying the **autoscaling.knative.dev/target** annotation value in the revision template.

Here is an example of how this annotation is used in the revision template to set the target to **50**.

```
autoscaling.knative.dev/target: 50
```

### 8.1.2. Configuring concurrent requests using the containerConcurrency field

**containerConcurrency** sets a hard limit on the number of concurrent requests handled.

```
containerConcurrency: 0 | 1 | 2-N
```

0


allows unlimited concurrent requests.

1

guarantees that only one request is handled at a time by a given instance of the revision container.

2 or more

will limit request concurrency to that value.



#### NOTE

If there is no **target** annotation, autoscaling is configured as if **target** is equal to the value of **containerConcurrency**.

## 8.2. CONFIGURING SCALE BOUNDS KNATIVE SERVING AUTOSCALING

The **minScale** and **maxScale** annotations can be used to configure the minimum and maximum number of Pods that can serve applications. These annotations can be used to prevent cold starts or to help control computing costs.

### minScale

If the **minScale** annotation is not set, Pods will scale to zero (or to 1 if `enable-scale-to-zero` is false per the **ConfigMap**).

### maxScale

If the **maxScale** annotation is not set, there will be no upper limit for the number of Pods created.

**minScale** and **maxScale** can be configured as follows in the revision template:

```
spec:
  template:
    metadata:
      autoscaling.knative.dev/minScale: "2"
      autoscaling.knative.dev/maxScale: "10"
```

Using these annotations in the revision template will propagate this configuration to **PodAutoscaler** objects.



### NOTE

These annotations apply for the full lifetime of a revision. Even when a revision is not referenced by any route, the minimal Pod count specified by **minScale** will still be provided. Keep in mind that non-routeable revisions may be garbage collected, which enables Knative to reclaim the resources.

## CHAPTER 9. GETTING STARTED WITH KNATIVE CLIENT



### IMPORTANT

You are viewing documentation for a release of Red Hat OpenShift Serverless that is no longer supported. Red Hat OpenShift Serverless is currently supported on OpenShift Container Platform 4.3 and newer.

Knative Client (**kn**) is the Knative command line interface (CLI). The CLI exposes commands for managing your applications, as well as lower level tools to interact with components of OpenShift Container Platform. With **kn**, you can create applications and manage OpenShift Container Platform projects from the terminal.

### 9.1. BEFORE YOU BEGIN


Knative Client does not have its own log in mechanism. To log in to the cluster you must install the **oc** CLI and use **oc** login.

Installation options for the **oc** CLI will vary depending on your operating system. For more information on installing the **oc** CLI for your operating system and logging in with **oc**, see the [CLI getting started](#) documentation.

### 9.2. INSTALLING KNATIVE CLIENT

#### 9.2.1. Installing the kn CLI using the OpenShift Container Platform web console

Once the OpenShift Serverless Operator is installed, you will see a link to download the **kn** CLI for Linux, macOS and Windows from the **Command Line Tools** page in the OpenShift Container Platform web console.

You can access the **Command Line Tools** page by clicking the  icon in the top right corner of the web console and selecting **Command Line Tools** in the drop down menu.

#### Procedure


1. Download the **kn** CLI from the **Command Line Tools** page.
2. Unpack the archive:

```
$ tar -xf <file>
```

3. Move the **kn** binary to a directory on your PATH.
4. To check your path, run:

```
$ echo $PATH
```



**NOTE**

If you do not use RHEL or Fedora, ensure that **libc** is installed in a directory on your library path. If **libc** is not available, you might see the following error when you run CLI commands:

```
$ kn: No such file or directory
```

### 9.2.2. Installing the kn CLI for Linux using an RPM

For Red Hat Enterprise Linux (RHEL), you can install **kn** as an RPM if you have an active OpenShift Container Platform subscription on your Red Hat account.

**Procedure**

- Use the following command to install **kn**:

```
# subscription-manager register
# subscription-manager refresh
# subscription-manager attach --pool=<pool_id> 1
# subscription-manager repos --enable="openshift-serverless-1-for-rhel-8-x86_64-rpms"
# yum install openshift-serverless-clients
```

- 1** Pool ID for an active OpenShift Container Platform subscription

### 9.2.3. Installing the kn CLI for Linux


For Linux distributions, you can download the CLI directly as a **tar.gz** archive.

**Procedure**

1. Download the [CLI](#).
2. Unpack the archive:
 

```
$ tar -xf <file>
```
3. Move the **kn** binary to a directory on your PATH.
4. To check your path, run:
 

```
$ echo $PATH
```

**NOTE**

If you do not use RHEL or Fedora, ensure that **libc** is installed in a directory on your library path. If **libc** is not available, you might see the following error when you run CLI commands:

```
$ kn: No such file or directory
```

### 9.2.4. Installing the kn CLI for macOS

**kn** for macOS is provided as a **tar.gz** archive.

#### Procedure

1. Download the [CLI](#).
2. Unpack and unzip the archive.
3. Move the **kn** binary to a directory on your PATH.
4. To check your PATH, open a terminal window and run:

```
$ echo $PATH
```

### 9.2.5. Installing the kn CLI for Windows

The CLI for Windows is provided as a zip archive.

#### Procedure

1. Download the [CLI](#).
2. Unzip the archive with a ZIP program.
3. Move the **kn** binary to a directory on your PATH.
4. To check your PATH, open the Command Prompt and run the command:

```
C:\> path
```

## 9.3. BASIC WORKFLOW USING KNATIVE CLIENT

Use this basic workflow to create, read, update, delete (CRUD) operations on a service. The following example deploys a [simple Hello World service](#) that reads the environment variable **TARGET** and prints its output.

#### Procedure

1. Create a service in the **default** namespace from an image.

```
$ kn service create hello --image gcr.io/knative-samples/helloworld-go --env TARGET=Knative  
Creating service 'hello' in namespace 'default':
```

```
0.085s The Route is still working to reflect the latest desired specification.  
0.101s Configuration "hello" is waiting for a Revision to become ready.  
11.590s ...  
11.650s Ingress has not yet been reconciled.  
11.726s Ready to serve.
```

```
Service 'hello' created with latest revision 'hello-gsdks-1' and URL:  
http://hello.default.apps-crc.testing
```

- 2. List the service.

```
$ kn service list
NAME URL LATEST AGE CONDITIONS READY
REASON
hello http://hello.default.apps-crc.testing hello-gsdks-1 8m35s 3 OK / 3 True
```

- 3. Check if the service is working by using the **curl** service endpoint command:

```
$ curl http://hello.default.apps-crc.testing

Hello Knative!
```

- 4. Update the service.

```
$ kn service update hello --env TARGET=Kn
Updating Service 'hello' in namespace 'default':

10.136s Traffic is not yet migrated to the latest revision.
10.175s Ingress has not yet been reconciled.
10.348s Ready to serve.

Service 'hello' updated with latest revision 'hello-dghll-2' and URL:
http://hello.default.apps-crc.testing
```

The service's environment variable **TARGET** is now set to **Kn**.

- 5. Describe the service.

```
$ kn service describe hello
Name: hello
Namespace: default
Age: 13m
URL: http://hello.default.apps-crc.testing
Address: http://hello.default.svc.cluster.local

Revisions:
100% @latest (hello-dghll-2) [2] (1m)
Image: gcr.io/knative-samples/helloworld-go (pinned to 5ea96b)

Conditions:
OK TYPE AGE REASON
++ Ready 1m
++ ConfigurationsReady 1m
++ RoutesReady 1m
```

- 6. Delete the service.

```
$ kn service delete hello
Service 'hello' successfully deleted in namespace 'default'.
```

You can then verify that the **hello** service is deleted by attempting to **list** it.

```
$ kn service list hello
No services found.
```

## 9.4. AUTOSCALING WORKFLOW USING KNATIVE CLIENT

You can access autoscaling capabilities by using **kn** to modify Knative services without editing YAML files directly.

Use the **service create** and **service update** commands with the appropriate flags to configure the autoscaling behavior.

Flag	Description
<b>--concurrency-limit int</b>	Hard limit of concurrent requests to be processed by a single replica.
<b>--concurrency-target int</b>	Recommendation for when to scale up based on the concurrent number of incoming requests. Defaults to <b>--concurrency-limit</b> .
<b>--max-scale int</b>	Maximum number of replicas.
<b>--min-scale int</b>	Minimum number of replicas.

## 9.5. TRAFFIC SPLITTING USING KNATIVE CLIENT

**kn** helps you control which revisions get routed traffic on your Knative service.

Knative service allows for traffic mapping, which is the mapping of revisions of the service to an allocated portion of traffic. It offers the option to create unique URLs for particular revisions and has the ability to assign traffic to the latest revision.

With every update to the configuration of the service, a new revision is created with the service route pointing all the traffic to the latest ready revision by default.

You can change this behavior by defining which revision gets a portion of the traffic.

### Procedure

- Use the **kn service update** command with the **--traffic** flag to update the traffic.

**NOTE**

**--traffic RevisionName=Percent** uses the following syntax:

- The **--traffic** flag requires two values separated by separated by an equals sign (=).
- The **RevisionName** string refers to the name of the revision.
- **Percent** integer denotes the traffic portion assigned to the revision.
- Use identifier **@latest** for the RevisionName to refer to the latest ready revision of the service. You can use this identifier only once with the **--traffic** flag.
- If the **service update** command updates the configuration values for the service along with traffic flags, the **@latest** reference will point to the created revision to which the updates are applied.
- **--traffic** flag can be specified multiple times and is valid only if the sum of the **Percent** values in all flags totals 100.

**NOTE**


For example, to route 10% of traffic to your new revision before putting all traffic on, use the following command:

```
$ kn service update svc --traffic @latest=10 --traffic svc-vwxyz=90
```

### 9.5.1. Assigning tag revisions

A tag in a traffic block of service creates a custom URL, which points to a referenced revision. A user can define a unique tag for an available revision of a service which creates a custom URL by using the format **http(s)://TAG-SERVICE.DOMAIN**.

A given tag must be unique to its traffic block of the service. **kn** supports assigning and unassigning custom tags for revisions of services as part of the **kn service update** command.

**NOTE**

If you have assigned a tag to a particular revision, a user can reference the revision by its tag in the **--traffic** flag as **--traffic Tag=Percent**.

#### Procedure

- Use the following command:

```
$ kn service update svc --tag @latest=candidate --tag svc-vwxyz=current
```

**NOTE**

**--tag RevisionName=Tag** uses the following syntax:

- **--tag** flag requires two values separated by a **=**.
- **RevisionName** string refers to name of the **Revision**.
- **Tag** string denotes the custom tag to be given for this Revision.
- Use the identifier **@latest** for the RevisionName to refer to the latest ready revision of the service. You can use this identifier only once with the **--tag** flag.
- If the **service update** command is updating the configuration values for the Service (along with tag flags), **@latest** reference will be pointed to the created Revision after applying the update.
- **--tag** flag can be specified multiple times.
- **--tag** flag may assign different tags to the same revision.

### 9.5.2. Unassigning tag revisions

Tags assigned to revisions in a traffic block can be unassigned. Unassigning tags removes the custom URLs.

**NOTE**

If a revision is untagged and it is assigned 0% of the traffic, it is removed from the traffic block entirely.

#### Procedure

- A user can unassign the tags for revisions using the **kn service update** command:

```
$ kn service update svc --untag candidate
```

**NOTE**

**--untag Tag** uses the following syntax:

- The **--untag** flag requires one value.
- The **tag** string denotes the unique tag in the traffic block of the service which needs to be unassigned. This also removes the respective custom URL.
- The **--untag** flag can be specified multiple times.

### 9.5.3. Traffic flag operation precedence

All traffic-related flags can be specified using a single **kn service update** command. **kn** defines the precedence of these flags. The order of the flags specified when using the command is not taken into account.

The precedence of the flags as they are evaluated by **kn** are:

1. **--untag**: All the referenced revisions with this flag are removed from the traffic block.
2. **--tag**: Revisions are tagged as specified in the traffic block.
3. **--traffic**: The referenced revisions are assigned a portion of the traffic split.


#### 9.5.4. Traffic splitting flags

**kn** supports traffic operations on the traffic block of a service as part of the **kn service update** command.

The following table displays a summary of traffic splitting flags, value formats, and the operation the flag performs. The "Repetition" column denotes whether repeating the particular value of flag is allowed in a **kn service update** command.

Flag	Value(s)	Operation	Repetition
<b>--traffic</b>	<b>RevisionName=Percent</b>	Gives <b>Percent</b> traffic to <b>RevisionName</b>	Yes
<b>--traffic</b>	<b>Tag=Percent</b>	Gives <b>Percent</b> traffic to the Revision having <b>Tag</b>	Yes
<b>--traffic</b>	<b>@latest=Percent</b>	Gives <b>Percent</b> traffic to the latest ready Revision	No
<b>--tag</b>	<b>RevisionName=Tag</b>	Gives <b>Tag</b> to <b>RevisionName</b>	Yes
<b>--tag</b>	<b>@latest=Tag</b>	Gives <b>Tag</b> to the latest ready Revision	No
<b>--untag</b>	<b>Tag</b>	Removes <b>Tag</b> from Revision	Yes

## CHAPTER 10. OPENSIFT SERVERLESS RELEASE NOTES



### IMPORTANT


You are viewing documentation for a release of Red Hat OpenShift Serverless that is no longer supported. Red Hat OpenShift Serverless is currently supported on OpenShift Container Platform 4.3 and newer.

For an overview of OpenShift Serverless functionality, see [Getting started with OpenShift Serverless](#).

### 10.1. GETTING SUPPORT

If you experience difficulty with a procedure described in this documentation, visit the [Customer Portal](#) to learn more about support for Technology Preview features.

### 10.2. RELEASE NOTES FOR RED HAT OPENSIFT SERVERLESS TECHNOLOGY PREVIEW 1.4.0



### IMPORTANT

OpenShift Serverless 1.4.0 contains a bad owner reference that causes the Kubernetes Garbage Collector to incorrectly remove the entire Knative control plane, including all of your services. You must install OpenShift Serverless 1.4.1 to fix this issue.

#### 10.2.1. New features

- OpenShift Serverless 1.4.0 is available on OpenShift Container Platform 4.2 and newer versions.
- OpenShift Serverless has been updated to use Knative Serving 0.11.1.
- OpenShift Serverless has been updated to use Knative Client (**kn** CLI) 0.11.0.
- OpenShift Serverless has been updated to use Knative Serving Operator 0.11.1.
- The **kn** CLI is now available for download through the **Command Line Tools** page in the OpenShift Container Platform web console.
- The **KnativeServing** object's API group has changed in this release from **servicing.knative.dev** to **operator.knative.dev**. You will need to adjust any of your scripts or applications that rely on the old API group to use the new API group. The OpenShift Serverless installation instructions have been updated to use the new API group.

When upgrading from OpenShift Serverless 1.3.0 to 1.4.0, the OpenShift Serverless Operator will create a **KnativeServing** custom resource (CR) in the new API group for you. This CR will be a mirror of the **KnativeServing** CR in the old group that was used in OpenShift Serverless 1.3.0.


If you need to keep using the old group temporarily, you can use the old CR as before. However, this CR is deprecated and will eventually be removed.

Once you have updated references to the new API group, you can remove any older CR versions and use the newly deployed **KnativeServing** CR instead. To safely do this without downtime, remove the owner reference from the newly deployed **KnativeServing** CR using:



```
$ oc edit knativeserving.operator.knative.dev knative-serving -n knative-serving
```

After the owner reference has been removed, you can safely remove any older CR versions and start using the new one.



### IMPORTANT

If a previous version of the CR exists, changes to the new CR will be overwritten by the OpenShift Serverless Operator. While the old CR is still active, all changes need to be made to that CR.

## 10.2.2. Fixed issues

- Connecting to a private, cluster local Knative Service from a namespace that was not part of the **knative-serving-ingress** Service Mesh was failing on **i/o timeout**. This issue is now fixed.
- The **container\_name** and **pod\_name** metric labels were removed in OpenShift Container Platform 4.3. The documentation has been updated to use the new **container** and **pod** metric labels instead. If you are using metering with Serverless on OpenShift Container Platform 4.3 or later, you must update your Prometheus queries according to the current version of the Serverless metering documentation.

## 10.2.3. Known issues

- Unqualified usage of **knativeserving** in **oc** commands no longer works because of the migration to a new API group. For example, this command will not work:

```
$ oc get knativeserving -n knative-serving
```

Use the explicit fully-qualified format instead. For example:

```
$ oc get knativeserving.operator.knative.dev -n knative-serving
```

- OpenShift Container Platform scale from zero latency causes a delay of approximately 10 seconds when creating pods. This is a current OpenShift Container Platform limitation.

## 10.3. RELEASE NOTES FOR RED HAT OPENSIFT SERVERLESS TECHNOLOGY PREVIEW 1.3.0

### 10.3.1. New features

- OpenShift Serverless has been updated to use Knative Serving 0.10.1.
- OpenShift Serverless has been updated to use Knative Client (**kn** CLI) 0.10.0.
- OpenShift Serverless 1.3.0 is available on OpenShift Container Platform 4.2 and newer versions.

### 10.3.2. Fixed issues

- Fixed a bug which caused Routes to have incorrect cross-namespaced **OwnerReferences**.

### 10.3.3. Known issues

- Connecting to a private, cluster local Knative Service from a namespace that is not part of the **knative-serving-ingress** Service Mesh fails on **i/o timeout**.

## 10.4. RELEASE NOTES FOR RED HAT OPENSIFT SERVERLESS TECHNOLOGY PREVIEW 1.2.0

### 10.4.1. New features

- OpenShift Serverless has been updated to use Knative Serving 0.9.0.
- OpenShift Serverless has been updated to use Knative Client (**kn CLI**) 0.9.0.
- OpenShift Serverless on OpenShift Container Platform 4.2 now uses the Operator Lifecycle Manager (OLM) dependency resolution mechanism to install the ServiceMesh Operator automatically. The required ServiceMeshControlPlane and ServiceMeshMemberRoll are also installed and managed for the user.
- Access to the KnativeServing resource is now restricted to **cluster-admin** roles to prevent any user from blocking the resource. Only **cluster-admin** roles can create KnativeServing CRs.
- The OpenShift Serverless Operator can now be found in the OperatorHub by searching for "knative".
- The OpenShift Container Platform web console now shows status conditions for the KnativeServing resource.
- In version 1.2.0, the OpenShift Serverless Operator inspects network policies for namespaces. If no network policy exists, the Operator automatically creates a wide open policy, to ensure that traffic can flow in and out of the namespace and OpenShift routes can be used.

If there is an existing network policy, OpenShift Serverless will not create a new policy. The Operator expects the user to continue managing their own network policies as needed for their applications. For example, the user must set policies that allow traffic to flow in and out of the namespace, and allow OpenShift routes to still be used after the namespace is added to a ServiceMeshMemberRoll.

### 10.4.2. Fixed issues

- In previous releases, using the same services or routes in different namespaces caused services to not work properly and caused OpenShift Container Platform routes to be overridden. This issue has been fixed.
- In previous releases, different traffic split targets required a mandatory tag. A single traffic split can now be defined with **untagged** traffic targets.
- Existing Knative Services and Routes which had been created with public visibility in OpenShift Serverless Operator version 1.1.0 were not able to be updated to cluster-local visibility. This issue is now fixed.
- The **Unknown Uninitialized : Waiting for VirtualService** error has been fixed.
- Knative service no longer returns a 503 status code when the cluster is running for a long time.

### 10.4.3. Known issues

- Installing the OpenShift Serverless Operator on OpenShift Container Platform versions older than 4.2.4 using OLM may incorrectly use community versions of the required dependencies. As a workaround, on OpenShift Container Platform versions older than 4.2.4, explicitly install the Red Hat provided versions of the Elastic Search, Jaeger, Kiali and ServiceMesh Operators before installing the OpenShift Serverless Operator.
- If you are upgrading the OpenShift Serverless from version 1.1.0 to version 1.2.0 and you have set up a ServiceMeshControlPlane and ServiceMeshMemberRoll to work with your Knative Serving instance, you must remove the **knative-serving** namespace and any other namespaces that contain Knative Services from the ServiceMeshMemberRoll in **istio-system**. You can also delete the ServiceMeshControlPlane from the namespace entirely if it is not required for other applications.

Once the upgrade starts, existing services will continue to work as before, but new Services will never become ready. Once you unblock the release by removing the **knative-serving** and any other relevant namespaces from the ServiceMeshMemberRoll, there will be a brief outage to all active Services. This will fix itself. Make sure that you remove all namespaces containing Knative Services from the original ServiceMeshMemberRoll.

- gRPC and HTTP2 do not work against routes. This is a known limitation of OpenShift routes.

## 10.5. RELEASE NOTES FOR RED HAT OPENSIFT SERVERLESS TECHNOLOGY PREVIEW 1.1.0

### 10.5.1. New features

- OpenShift Serverless has been updated to use Knative Serving 0.8.1.
- Enhanced Operator metadata now includes more information regarding support state and a link to the official installation documentation.
- A developer preview version of Knative Eventing is now available for use with OpenShift Serverless, however this is not included in the OpenShift Serverless Operator and is not currently supported as part of this Technology Preview. For more information, see [Knative Eventing on OpenShift Container Platform](#).

### 10.5.2. Fixed issues

- Users who were not project administrators would previously see the following error when using OpenShift Serverless:

```
revisions.serving.knative.dev: User "sounds" cannot list resource "revisions"
```

This issue has now been fixed with the addition of new RBAC rules.

- A race condition was preventing Istio sidecar injection from working correctly. Istio did not consider the **knative-serving** namespace to be present in the ServiceMeshMemberRoll at the time of Pod creation. Istio now waits for status information from ServiceMeshMemberRoll which fixes this issue.

### 10.5.3. Known issues

- Users may see the error **Unknown Uninitialized : Waiting for VirtualService to be ready** while waiting for a service in a newly created namespace to be ready, which can take several minutes.

If a user allows enough time between the creation of a namespace and the creation of a service in the namespace (approximately one minute), this error may be avoided.

- Existing Knative Services and Routes which have been created with public visibility cannot be updated to cluster-local visibility. If you require cluster-local visibility on Knative Services and Routes, this must be configured at the time of creating these resources.
- Knative service returns a 503 status code when the cluster is running for a long time. The Knative Serving Pods do not show any errors. Restarting the **istio-pilot** Pod temporarily fixes the issue.
- gRPC and HTTP2 do not work against routes. This is a known limitation of OpenShift routes.

## 10.6. RELEASE NOTES FOR RED HAT OPENSIFT SERVERLESS TECHNOLOGY PREVIEW 1.0.0

### 10.6.1. New features

This release of OpenShift Serverless introduces the OpenShift Serverless Operator, which supports Knative Serving 0.7.1 and is tested for OpenShift Service Mesh 1.0.

### 10.6.2. Known issues

The following limitations exist in OpenShift Serverless at this time:

- The Knative Serving Operator should wait for ServiceMeshMemberRoll to include the **knative-serving** namespace. The installation procedure recommends creating the **knative-serving** namespace and then installing the operator. Istio does not consider the **knative-serving** namespace to be in the ServiceMeshMemberRoll when the Knative Serving Pods are being created. Consequently, the sidecars are not injected.
- Knative service returns a 503 status code when the cluster is running for a long time. The Knative Serving Pods do not show any errors. Restarting the **istio-pilot** Pod temporarily fixes the issue.
- The gRPC and HTTP2 do not work against routes. This is a known limitation of OpenShift routes.

## 10.7. ADDITIONAL RESOURCES

OpenShift Serverless is based on the open source Knative project.

- For details about the latest Knative Serving release, see the [Knative Serving releases page](#).
- For details about the latest Knative Client release, see the [Knative Client releases page](#).
- For details about the latest Knative Eventing release, see the [Knative Eventing releases page](#).



### NOTE

Knative Eventing is currently available as a Developer Preview on OpenShift Container Platform. See the upstream [Knative Eventing on OpenShift Container Platform documentation](#).

