

# Cloud Native Runtimes for VMware Tanzu 2.0

Cloud Native Runtimes for VMware Tanzu 2.0

You can find the most up-to-date technical documentation on the VMware by Broadcom website at:

<https://docs.vmware.com/>

**VMware by Broadcom**  
3401 Hillview Ave.  
Palo Alto, CA 94304  
[www.vmware.com](http://www.vmware.com)

Copyright © 2024 Broadcom. All Rights Reserved. The term “Broadcom” refers to Broadcom Inc. and/or its subsidiaries. For more information, go to <https://www.broadcom.com>. All trademarks, trade names, service marks, and logos referenced herein belong to their respective companies. [Copyright and trademark information](#).

# Contents

Cloud Native Runtimes Overview	7
Overview	7
Warnings	8
Cloud Native Runtimes Release Notes	9
v2.0.5	9
Security Fixes	9
Components	9
v2.0.4	9
Security Fixes	10
Components	10
v2.0.2	10
Resolved Issues	10
Components	11
v2.0.1	11
Breaking Changes	11
New Features	12
Known Issues	12
Deprecation Notice	12
Components	13
Integrations you can use with Cloud Native Runtimes	14
CNR Integrations	14
Install Cloud Native Runtimes	15
Prerequisites	15
Install	15
Administrator Guide for Cloud Native Runtimes	18
Configure your External DNS with Cloud Native Runtimes	18
Overview	18
Configure custom domain	18
Configure Knative Service Domain Template	19
Use your existing TLS Certificate for Cloud Native Runtimes	20

Prerequisites	20
Installing Cloud Native Runtimes with your Existing Contour Installation	21
About Using Contour with Cloud Native Runtimes	21
Prerequisites	22
Identify Your Contour Version	22
Install Cloud Native Runtimes on a Cluster with Your Existing Contour Instances	22
Securing Your Web Workloads in Cloud Native Runtimes	23
Overview	23
Prerequisites	23
Enable Auto TLS Using an HTTP01 Challenge	23
Enable Auto TLS Using a DNS01 Challenge	25
Configuring Eventing with RabbitMQ for Cloud Native Runtimes	27
Overview	27
Install VMware Tanzu RabbitMQ for Kubernetes	27
Next Steps	28
Configuring Observability for Cloud Native Runtimes	28
Overview	28
Logging	28
Configure Logging with Fluent Bit	28
Forward Logs to vRealize	30
Metrics	30
Tracing	30
Configuring Tracing	30
Forwarding Trace Data to a Data Visualization Tool	31
Sending Trace Data to an Observability Platform	31
Use Wavefront Dashboards	32
Import Wavefront Dashboards	32
Import Dashboards with the Wavefront API	32
Import with the Ruby Wavefront CLI	33
Configuring Cloud Native Runtimes with Avi Vantage	33
Overview	33
Integrate Avi Vantage with Cloud Native Runtimes	33
About Routing with Avi Vantage and Cloud Native Runtimes	34
Configuring Cloud Native Runtimes with Tanzu Service Mesh	35
Overview	36

Run Cloud Native Runtimes on a Cluster Attached to Tanzu Service Mesh	36
Next Steps	36
<b>Troubleshooting Cloud Native Runtimes</b>	<b>37</b>
Cannot connect to app on AWS	37
Symptom	37
Solution	37
minikube Pods Fail to Start	37
Symptom	37
Solution	37
Pulling an image with imgpkg overwrites the cloud-native-runtimes directory	38
Symptom	38
Solution	38
Installation fails to reconcile app/cloud-native-runtimes	38
Symptom	38
Explanation	38
Solution	38
Example 1: The Cloud Provider does not support the creation of Service type LoadBalancer	39
Example 2: The webhook deployment failed	39
Cloud Native Runtimes Installation Fails with Existing Contour Installation	40
Symptom	40
Solution	40
Knative Serving Workloads fail to deploy with Forbidden errors on Run Cluster when using Multi Cluster setup on Openshift	41
Symptom	41
Solution	41
<b>Verifying Your Installation</b>	<b>43</b>
Prerequisites	43
Verify Installation of Knative Serving, Knative Eventing, and TriggerMesh SAWS	44
<b>Verifying Knative Serving for Cloud Native Runtimes</b>	<b>44</b>
About Verifying Knative Serving	44
Prerequisites	45
Test Knative Serving	45
Delete the Example Knative Service	46
<b>Verify Knative Eventing with Cloud Native Runtimes</b>	<b>46</b>
About Verifying Knative Eventing	47
Prerequisites	47

Prepare the RabbitMQ Environment	47
Verify Knative Eventing	48
Setup RabbitMQ Broker as the default in the cluster (optional)	51
Setup RabbitMQ Broker as the default in a namespace (optional)	52
Delete the Eventing Resources	54
<b>Verifying TriggerMesh SAWS for Cloud Native Runtimes</b>	<b>54</b>
Overview	55
Prerequisites	55
Verify TriggerMesh SAWS	55
<b>Upgrading Cloud Native Runtimes</b>	<b>58</b>
Prerequisites	58
Upgrade from a previous version to Cloud Native Runtimes v2.0.1	58
Known issue in Eventing 2.0 during the upgrade of a profile-based installation	59
Workaround	59
Upgrade Cloud Native Runtimes	59
<b>Uninstalling Cloud Native Runtimes</b>	<b>60</b>
Overview	60
Uninstall	60

# Cloud Native Runtimes Overview

This topic gives you an overview of Cloud Native Runtimes, commonly known as CNR.

## Overview

Cloud Native Runtimes (CNRs) is enterprise supported Knative, with the Carvel tools suite for deployment and Contour for networking. CNRs offers everything Knative does and some extras that make it ideal for cloud native application development. Cloud Native Runtimes gives developers environmental simplicity and administrators deployment control and it works on any single Kubernetes cluster running Kubernetes v1.22 and later.

CNR utilizes Knative’s main features of Serving and Eventing to provide:

- Automatic pod scaling.
- Traffic splitting by code release version.
- Event-triggered workloads.

Cloud Native Runtimes simplifies the Developer experience.

Kubernetes Developers need to know:	Cloud Native Runtimes Developers need to know:
Pods	Pods
Deployment & Rollout Progress	Knative Service
Service (networking model)	
Ingress	
Labels and selectors	

Cloud Native Runtimes increases Administrator control and support.

Administrators can:
Manage infrastructure costs with request driven autoscaling
Test deployments with traffic splitting by code version
Use Carvel command tools to simplify deployment
Receive Enterprise Support when they need it

Cloud Native Runtimes works well with these use cases:

- Batch Jobs Processing
- AI/ML

- Application or Network Monitoring
- IOT
- Event driven and serverless application architectures

For more information on the software that makes Cloud Native Runtimes see:

- Knative Documentation [Home - Knative](#)
- Carvel Tools Suite Documentation [Carvel - Home](#)
- Contour Networking Documentation [Contour](#)

## Warnings

Eventing in Tanzu Application Platform is deprecated and marked for removal in Tanzu Application Platform v1.7.0.



# Cloud Native Runtimes Release Notes

This topic contains release notes for Cloud Native Runtimes (CNR) for Tanzu v2.0.

## v2.0.5

**Release Date:** April 11, 2023

### Security Fixes

The following updates were done to the Eventing package:

- Update Knative dependencies, which contains bumps with fixes for:
  - ◊ [CVE-2022-32149](#)
  - ◊ [CVE-2022-27664](#)

**NOTE:** Knative 1.6 is out of support in Open Source. Therefore, the changes included in these bumps are maintained in VMware's proprietary copies of the Knative repositories.

### Components

Cloud Native Runtimes v2.0.5 uses the following component versions:

	Release	Details
Version		2.0.5
Release date		April 11, 2023
	Component	Version
	Knative Serving	1.6.3
	Knative net-certmanager	1.6.1
	Knative net-contour	1.6.1
	Knative Eventing	1.6.3
	Knative eventing-rabbitmq	1.6.1
	VMware sources-for-knative	1.6.1
	Triggermesh aws-event-sources	1.6.1

## v2.0.4

**Release Date:** March 6, 2023

## Security Fixes

The following updates were done to the CNRs and Eventing packages:

- Update the Ubuntu Bionic base image to [Bionic Tiny 1.3.109](#), which contains fixes for:
  - ◊ [CVE-2023-0286](#)
- Update Knative dependencies, which contains bumps with fixes for:
  - ◊ [CVE-2022-27191](#)

The following updates were done to the CNRs package only:

- Update Knative dependencies, which contains bumps with fixes for:
  - ◊ [CVE-2022-32149](#)
  - ◊ [CVE-2022-27664](#)

**NOTE:** Knative 1.6 is out of support in Open Source. Therefore, the changes included in these bumps are maintained in VMware's proprietary copies of the Knative repositories.

## Components

Cloud Native Runtimes v2.0.4 uses the following component versions:

	Release	Details
Version		2.0.4
Release date		March 6, 2023
	Component	Version
	Knative Serving	1.6.3
	Knative net-certmanager	1.6.1
	Knative net-contour	1.6.1
	Knative Eventing	1.6.0
	Knative eventing-rabbitmq	1.6.1
	VMware sources-for-knative	1.6.1
	Triggernesh aws-event-sources	1.6.0

## v2.0.2

**Release Date:** November 15, 2022

## Resolved Issues

This release has resolved the following issues:

- **Serving package:** This release has resolved the issue where on Multi Cluster setup with Openshift, Knative workloads may fail to come up on the `run` cluster due to SecurityContextConstraint restrictions. The steps outlined in the CNR [Troubleshooting section](#) only apply to CNR 2.0.1.

## Components

Cloud Native Runtimes v2.0.2 uses the following component versions:

Release		Details
Version		v2.0.2
Release date		October 11, 2022
Component		Version
Knative Serving		1.6.0
Knative cert-manager Integration		1.6.0
Knative Serving Contour Integration		1.6.0

## v2.0.1

**Release Date:** October 11, 2022

## Breaking Changes

- **Removed Knative Eventing:** Cloud Native Runtimes v2.0.1 no longer includes the following components:
  - ◊ Knative Eventing
  - ◊ Knative Eventing RabbitMQ Integration
  - ◊ VMware Tanzu Sources for Knative
  - ◊ TriggerMesh Sources from Amazon Web Services (SAWS)



### Note

If you previously installed Cloud Native Runtimes v1.3.0 or later, see [Upgrading Cloud Native Runtimes](#) before proceeding with the upgrade.

- **Removed CNRs Contour:** Cloud Native Runtimes no longer ships with a Contour instance and it will default to [TAP Contour](#).



### Note

Please note that a valid Contour instance must be present in a cluster where CNRs will be installed. If you are currently on a version that uses CNR's Contour and wish to upgrade, see [\[Migrating from CNR Contour to TAP Contour\]](#).

- **Removed config options:**
  - ◊ `reuse_crds`: CNRs will now always use the Contour CRDs present on the cluster. If you wish to use CNRs with an existing Contour installation, see [Installing Cloud Native](#)

### Runtimes with an Existing Contour Installation.

- ✦ `nodeport`: Due to the CNRs Contour removal, CNRs no longer has Envoy instances to configure, making this config no longer used.
- ✦ `provider: tkgs`: Due to CNRs Contour removal, it is no longer possible to set `provider: tkgs` since this config updated the Envoy instances that are no longer controlled by CNRs.

## New Features

New features in this release:

- **Golang Bump**: Built with [Golang 1.18.3](#).
- **Knative 1.6.0**: See the Release Notes for [Knative Serving](#).
- **Released Eventing v2.0.1**: [Knative Eventing](#)'s features are now shipped as a separate package called `eventing.tanzu.vmware.com`. Both the `cnrs` and `eventing` packages are installed by default on `full` and `light` TAP profiles. However, it is possible to exclude any of them, if desired.
- **Added Support for Openshift 4.10 on vSphere and Baremetal**

## Known Issues

This release has the following issues:

- **Eventing package**: There is an issue that prevents Knative eventing Sugar Controller to work with brokers other than the `MTChannelBasedBroker`. More information can be found at on [this github issue](#).
- **Serving package**: There is an issue where on Multi Cluster setup with Openshift, Knative workloads may fail to come up on the `run` cluster due to `SecurityContextConstraint` restrictions. For more information, see the [Troubleshooting section](#).
- **Upgrading a TAP Profile-based installation**: There is an issue that affects the Eventing package during the upgrade of a TAP profile installation. See [Upgrade instructions](#) for the recommended workaround.

## Deprecation Notice

This following features are deprecated in this release and will be removed in a future release of CNRs.

- **provider config option**: After the recent changes in this release, the `provider` option results in the same changes as the `lite` and `pdb` options combined. As a result, if you are on a development environment and wish to lower CPU and memory requests for resources, please use `lite.enable`. In case you also want to deactivate pod disruption budgets on Knative Serving and high availability is not indispensable in your development environment, please set `pdb.enable` to `false` from now on as `provider` will be removed in a future release of CNRs and Eventing.



### Note



`provider` is being deprecated on both CNRs and Eventing packages.

## Components

Cloud Native Runtimes v2.0.1 uses the following component versions:

	Release	Details
Version		v2.0.1
Release date		October 11, 2022
	Component	Version
	Knative Serving	1.6.0
	Knative cert-manager Integration	1.6.0
	Knative Serving Contour Integration	1.6.0

# Integrations you can use with Cloud Native Runtimes

This topic tells you the supported integrations for Cloud Native Runtimes. For more information regarding these integrations, see the [CNR Administrator Guide](#).

## CNR Integrations

CNR Integration	Version	Documentation
VMware Tanzu Observability	Supported	<a href="#">Configuring Observability for CNR</a>
Avi Vantage	Supported	<a href="#">Configuring CNR with Avi Vantage</a>
Rabbit MQ	Supported	<a href="#">Configuring CNR with RabbitMQ</a>
Tanzu Service Mesh	Supported	<a href="#">Configuring CNR with Tanzu Service Mesh</a>

For tools and software compatibility, please refer to [Tanzu Application Platform's requirements](#).

# Install Cloud Native Runtimes

This document describes how you can install Cloud Native Runtimes, commonly known as CNR, from the Tanzu Application Platform package repository.

**Note:** Use the instructions on this page if you do not want to use a profile to install packages. Both the full and light profiles include Cloud Native Runtimes. For more information about profiles, see [Installing the Tanzu Application Platform Package and Profiles](#).

## Prerequisites

Before installing Cloud Native Runtimes:

- Complete all prerequisites to install Tanzu Application Platform. For more information, see [Prerequisites](#).
- Contour is installed in the cluster. Contour can be installed from the [Tanzu Application package repository](#). If you have an existing Contour installation, see [Installing Cloud Native Runtimes with an Existing Contour Installation](#).

## Install

To install Cloud Native Runtimes:

1. List version information for the package by running:

```
tanzu package available list cnrs.tanzu.vmware.com --namespace tap-install
```

For example:

```
$ tanzu package available list cnrs.tanzu.vmware.com --namespace tap-install
- Retrieving package versions for cnrs.tanzu.vmware.com...
NAME                VERSION  RELEASED-AT
cnrs.tanzu.vmware.com 2.0.1    2022-09-19 00:00:00 -0800 PST
```

2. (Optional) Make changes to the default installation settings:

1. Gather values schema.

```
tanzu package available get cnrs.tanzu.vmware.com/2.0.1 --values-schema -n tap-install
```

For example:

```
$ tanzu package available get cnrs.tanzu.vmware.com/2.0.1 --values-schema -n tap-install
```

```
Retrieving package details for cnrs.tanzu.vmware.com/2.0.1...
KEY                                DEFAULT                                TYPE DESCRIPTION
lite.enable                         false                                boolean
Optional: Not recommended for production. Set to "true" to reduce CPU and Memory resource requests for all CNR Deployments, Daemonsets, and Statefulsets by half. On by default when "provider" is set to "local".
pdb.enable                          true                                 boolean
Optional: Set to true to enable Pod Disruption Budget. If provider local is set to "local", the PDB will be disabled automatically.
provider                            <nil>                               string
Optional: Kubernetes cluster provider. To be specified if deploying CNR on a local Kubernetes cluster provider.
default_tls_secret                  <nil>                               string
Optional: Overrides the config-contour configmap in namespace knative-serving.
domain_config                       <nil>                               object
Optional: Overrides the config-domain configmap in namespace knative-serving. Must be valid YAML.
domain_name                         <nil>                               string
Optional: Default domain name for Knative Services.
domain_template                     {{.Name}}.{{.Namespace}}.{{.Domain}} string
Optional: specifies the goolang text template string to use when constructing the Knative service's DNS name.
ingress.external.namespace          tanzu-system-ingress                string
Required: Specify a namespace where an existing Contour is installed on your cluster. CNR will use this Contour instance for external services.
ingress.internal.namespace          tanzu-system-ingress                string
Required: Specify a namespace where an existing Contour is installed on your cluster. CNR will use this Contour instance for internal services.
```

2. Create a `cnr-values.yaml` file by using the following sample as a guide to configure Cloud Native Runtimes:

**Note:** For most installations, you can leave the `cnr-values.yaml` empty, and use the default values.

```
---
# Configures the domain that Knative Services will use
domain_name: "mydomain.com"
```

**Configuration Notes:** \* If you are running on a single-node cluster, such as minikube, set the `provider: local` option. This option reduces resource requirements by using a NodePort service instead of a LoadBalancer and reduces the number of replicas.

- Cloud Native Runtimes reuses the existing `tanzu-system-ingress` Contour installation for external and internal access when installed in the `light` or `full` profile. If you want to use a separate Contour installation for system-internal traffic, set `cnrs.ingress.internal.namespace` to the empty string (`""`).
- If you are running on a multinode cluster, do not set `provider`.
- If your environment already has Contour installed, the installed Contour might conflict with the Cloud Native Runtimes installation. For information about how to prevent conflicts, see [Installing Cloud Native Runtimes with an Existing](#)



### Contour Installation.

3. Install the package by running:

```
tanzu package install cloud-native-runtimes -p cnrs.tanzu.vmware.com -v 2.0.1 -n tap-install --values-file cnr-values.yaml --poll-timeout 30m
```

For example:

```
$ tanzu package install cloud-native-runtimes -p cnrs.tanzu.vmware.com -v 2.0.1 -n tap-install --values-file cnr-values.yaml --poll-timeout 30m
- Installing package 'cnrs.tanzu.vmware.com'
| Getting package metadata for 'cnrs.tanzu.vmware.com'
| Creating service account 'cloud-native-runtimes-tap-install-sa'
| Creating cluster admin role 'cloud-native-runtimes-tap-install-cluster-role'
| Creating cluster role binding 'cloud-native-runtimes-tap-install-cluster-role binding'
- Creating package resource
- Package install status: Reconciling

Added installed package 'cloud-native-runtimes' in namespace 'tap-install'
```

4. Verify the package install by running:

```
tanzu package installed get cloud-native-runtimes -n tap-install
```

For example:

```
tanzu package installed get cloud-native-runtimes -n tap-install
| Retrieving installation details for cc...
NAME:                cloud-native-runtimes
PACKAGE-NAME:        cnrs.tanzu.vmware.com
PACKAGE-VERSION:     2.0.1
STATUS:              Reconcile succeeded
CONDITIONS:          [{ReconcileSucceeded True  }]
USEFUL-ERROR-MESSAGE:
```

Verify that **STATUS** is **Reconcile succeeded**.

# Administrator Guide for Cloud Native Runtimes

The next several pages show you how to use, troubleshoot, integrate, upgrade and uninstall Cloud Native Runtimes, commonly known as CNR.

## Configure your External DNS with Cloud Native Runtimes

This topic describes how you can configure your external DNS with Cloud Native Runtimes, commonly known as CNR.

### Overview

Knative uses `example.com` as the default domain.

**Note:** If you are setting up Cloud Native Runtimes for development or testing, you do not have to set up an external DNS. However, if you want to access your workloads (apps) over the internet, then you do need to set up a custom domain and an external DNS.

### Configure custom domain

To set up the custom domain and its external DNS record:

1. Configure your custom domain:

When your workloads are created, Knative will automatically create URLs for each workload based on the configuration in the domain ConfigMap.

- To set a default custom domain, edit your `cnr-values.yml` file to contain the following:

```
---
domain_name: "mydomain.com"
```

This will modify the Knative domain ConfigMap to use `domain_name` as the default domain.

**Note:** `domain_name` must be a valid DNS subdomain.

- **Advanced:** To overwrite the domain ConfigMap entirely, edit your `cnr-values.yml` file to contain your desired config-domain options, similar to the following:

```
---
domain_config: |
  ---
  mydomain.com: |
```

```
mydomain.org: |
  selector:
    app: nonprofit
```

This will replace the body of the Knative domain ConfigMap with `domain_config`. This will allow you to configure multiple custom domains, and configure a custom domain for a service depending on its labels.

See [Changing the default domain](#) for more information about the structure of the domain ConfigMap.

**Note:** `domain_config` must be valid YAML and a valid domain ConfigMap.

**Note:** You can only use one of `domain_config` or `domain_name` at a time. You may not use both.

2. Get the address of the cluster load balancer:

```
kubectl get service envoy -n EXTERNAL-CONTOUR-NS --output 'jsonpath={.status.loadBalancer.ingress}'
```

Where `EXTERNAL-CONTOUR-NS` is the namespace where a Contour serving external traffic is installed. If Cloud Native Runtimes was installed as part of a Tanzu Application Profile, this value will likely be `tanzu-system-ingress`.

If this command returns a URL instead of an IP address, then `ping` the URL to get the load balancer IP address.

3. Create a wildcard DNS `A` record that assigns the custom domain to the load balancer IP. Follow the instructions provided by your domain name registrar for creating records.

The record created looks like:

```
*.DOMAIN IN A TTL LOADBALANCER-IP
```

Where:

- ◆ `DOMAIN` is the custom domain.
- ◆ `TTL` is the time-to-live.
- ◆ `LOADBALANCER-IP` is the load balancer IP.

For example:

```
*.mydomain.com IN A 3600 198.51.100.6
```

If you chose to configure multiple custom domains above, you will need to create a wildcard DNS record for each domain.

## Configure Knative Service Domain Template

Knative uses domain template which specifies the goolang text template string to use when constructing the Knative service's DNS name. The default value is `{{.Name}}.{{.Namespace}}.{{.Domain}}`. Valid variables defined in the template include Name, Namespace, Domain, Labels,

and Annotations.

To configure domain template for the created Knative Services, edit your `cnr-values.yml` file to contain the following:

```
---
domain_template: "{{.Name}}-{{.Namespace}}.{{.Domain}}"
```

This will modify the Knative `domain-template` ConfigMap to use `domain_template` as the default domain template.

Changing this value might be necessary when the extra levels in the domain name generated are problematic for wildcard certificates that only support a single level of domain name added to the certificate's domain. In those cases you might consider using a value of `{{.Name}}-{{.Namespace}}.{{.Domain}}`, or removing the Namespace entirely from the template.

When choosing a new value, be thoughtful of the potential for conflicts, such as when users the use of characters like `-` in their service or namespace names.

`{{.Annotations}}` or `{{.Labels}}` can be used for any customization in the go template if needed.

It is strongly recommended to keep namespace part of the template to avoid domain name clashes: eg. `{{.Name}}-{{.Namespace}}.{{ index .Annotations "sub" }}.{{.Domain}}` and you have an annotation `{"sub": "foo"}`, then the generated template would be `{Name}-{Namespace}.foo.{Domain}`.

## Use your existing TLS Certificate for Cloud Native Runtimes

This topic tells you how to use your existing TLS Certificate for Cloud Native Runtimes, commonly known as CNR.

### Prerequisites

In order to configure TLS for Cloud Native Runtimes, you must first configure a Service Domain. For more information, see [Configuring External DNS with CNR](#).

To configure your TLS certificate for the created Knative Services, follow the steps:

- Create a Kubernetes Secret to hold your TLS Certificate

```
kubectl create -n <Enter developer namespace here> secret tls <Enter TLS_NAME name here> \
--key key.pem \
--cert cert.pem
```

- Create a delegation. To do so, create a `"tlscertdelegation.yaml"` file with following contents

```
apiVersion: projectcontour.io/v1
kind: TLSCertificateDelegation
metadata:
  name: default-delegation
  namespace: <Enter developer namespace here>
spec:
  delegations:
```

```
- secretName: <Enter the above TLS_NAME here>
  targetNamespaces:
  - "<Enter developer namespace here>"
```

Apply the above yaml file by running

```
k apply -f tlscertdelegation.yaml
```

- Include the following config in your `tap-values.yaml` file and redeploy:

```
---
default_tls_secret: "<Enter developer namespace here>/<Enter the above TLS_NAME here>"
```

To redeploy run -

```
tanzu package installed update tap -p tap.tanzu.vmware.com --values-file "tap-values.yaml" -n tap-install
```

This will modify the Knative `default_tls_secret` ConfigMap to use `default_tls_secret` as the default tls certificate

## Installing Cloud Native Runtimes with your Existing Contour Installation

This topic describes how you can configure Cloud Native Runtimes, commonly known as CNR, with your existing Contour instance. Cloud Native Runtimes uses Contour to manage internal and external access to the services in a cluster.

### About Using Contour with Cloud Native Runtimes

The instructions on this page assume that you have an existing Contour installation on your cluster.

Follow the instructions on this page if:

- You have installed Contour as part of TAP and wish to configure Cloud Native Runtimes to use it.
- You see an error about [an existing Contour installation](#) when you install the Cloud Native Runtimes package.

Cloud Native Runtimes needs two instances of Contour: one instance for exposing services outside the cluster, and another instance for services that are private in your network. If installed as part of a Tanzu Application Platform profile, by default Cloud Native Runtimes will use the Contour instance installed in the namespace `tanzu-system-ingress` for both internal and external traffic.

If you already use a Contour instance to route requests from clients outside and inside the cluster, you may use your own Contour if it matches the Contour version used by [Tanzu Application's Platform](#).

You may use the same single instance of Contour for both internal and external traffic. However, this will cause internal and external traffic will be handled the same way. For example, if the Contour instance is configured to be accessible from clients outside the cluster, then any internal traffic will also be accessible from outside the cluster. Note that currently Tanzu Application Platform only

supports using a single Contour instance for both internal and external traffic.

In all of the above cases, Cloud Native Runtimes will use the Tanzu Application Platform's Contour `CustomResourceDefinitionS`.

## Prerequisites

The following prerequisites are required to configure Cloud Native Runtimes with an existing Contour installation:

- Contour version v1.22.0 (Contour version that is installed as part of TAP). To identify your cluster's Contour version, see [Identify Your Contour Version](#) below.
- Contour `CustomResourceDefinitionS` versions:

Resource Name	Version
<code>contourdeployments.projectcontour.io</code>	v1alpha1
<code>contourconfigurations.projectcontour.io</code>	v1alpha1
<code>extensionservices.projectcontour.io</code>	v1alpha1
<code>httpproxies.projectcontour.io</code>	v1
<code>tlscertificatedelegations.projectcontour.io</code>	v1

## Identify Your Contour Version

To identify your cluster's Contour version, run:

```
export CONTOUR_NAMESPACE=CONTOUR-NAMESPACE
export CONTOUR_DEPLOYMENT=$(kubectl get deployment --namespace $CONTOUR_NAMESPACE --output name)
kubectl get $CONTOUR_DEPLOYMENT --namespace $CONTOUR_NAMESPACE --output jsonpath="{.spec.template.spec.containers[].image}"
kubectl get crds extensionservices.projectcontour.io --output jsonpath="{.status.storedVersions}"
kubectl get crds httpproxies.projectcontour.io --output jsonpath="{.status.storedVersions}"
kubectl get crds tlscertificatedelegations.projectcontour.io --output jsonpath="{.status.storedVersions}"
```

Where `CONTOUR-NAMESPACE` is the namespace where Contour is installed on your Kubernetes cluster.

## Install Cloud Native Runtimes on a Cluster with Your Existing Contour Instances

To install Cloud Native Runtimes on a cluster with an existing Contour instance, you can add values to your `cnr-values.yml` file so that Cloud Native Runtimes will use your Contour instance.

An example of a `cnr-values.yml` file where you wish Cloud Native Runtimes to use the Contour version in a different namespace would look like this:

```
---
```

```

ingress:
  external:
    namespace: "tanzu-system-ingress"
  internal:
    namespace: "tanzu-system-ingress"

```

**Note:** If your Contour instance is removed or configured incorrectly, apps running on Cloud Native Runtimes will lose connectivity.

## Securing Your Web Workloads in Cloud Native Runtimes

This topic give you an overview of securing HTTP connections using TLS certificates in Cloud Native Runtimes, commonly known as CNR, for VMware Tanzu Application Platform and helps you configure TLS (Transport Layer Security).

### Overview

Cloud Native Runtimes supports both [HTTP01](#) and [DNS01](#) cert-manager challenge types. For more information about [cert-manager](#) challenge types, see [ACME](#) in the cert-manager documentation.

VMware recommends using Let's Encrypt as your certificate authority. However, you can integrate Cloud Native Runtimes with any ACME compatible certificate authority.

### Prerequisites

You can enable HTTPS with Automatic TLS certificate provisioning for Cloud Native Runtimes.

You need the following prerequisites to use secure HTTPS connections with automatic TLS certificate provisioning:

- A cluster configured to use a custom domain. See [Setting Up a Custom Domain](#) in the Knative documentation.
- A DNS provider configured with your domain name.
- cert-manager version 1.0.0 or later. See [Installing cert-manager for TLS certificates](#) in the Knative documentation.
- [HTTP01 challenges](#): An internet-reachable cluster.
- [DNS01 challenges](#): API access to set DNS records.

### Enable Auto TLS Using an HTTP01 Challenge

You can use the [HTTP01](#) challenge type to validate a domain with Cloud Native Runtimes. The [HTTP01](#) challenge requires that your load balancer be reachable from the internet via HTTP.



#### Note

With the [HTTP01](#) challenge type, you provision a certificate for each service.

To enable automatic TLS certificate provisioning using a [HTTP01](#) challenge, do the following:

1. Create a cert-manager `Issuer` or `ClusterIssuer` for the `HTTP01` challenge. See `Issuer` in the cert-manager documentation. The following example creates a `ClusterIssuer` using the Let's Encrypt Certificate Authority. See `Let's Encrypt`. To use a `ClusterIssuer` for the `HTTP01` challenge, run:

```
kubectl apply -f - <<EOF
apiVersion: cert-manager.io/v1
kind: ClusterIssuer
metadata:
  name: letsencrypt-http01-issuer
spec:
  acme:
    privateKeySecretRef:
      name: letsencrypt
    server: https://acme-v02.api.letsencrypt.org/directory
    solvers:
      - http01:
          ingress:
            class: contour
EOF
```

2. To validate that your `ClusterIssuer` was created successfully, run:

```
kubectl get clusterissuer letsencrypt-http01-issuer --output yaml
```

3. Edit your `config-certmanager` ConfigMap in the `knative-serving` namespace to reference the `ClusterIssuer` you created. Run:

```
kubectl edit configmap config-certmanager --namespace knative-serving
```

4. To define which `ClusterIssuer` will be used by Knative to issue certificates, add the following `issuerRef` block under the `data` section of the `config-certmanager` ConfigMap:

```
...
data:
  ...
  issuerRef: |
    kind: ClusterIssuer
    name: letsencrypt-http01-issuer
```

5. To validate that your ConfigMap was updated successfully, run:

```
kubectl get configmap config-certmanager --namespace knative-serving --output jsonpath="{.data.issuerRef}"
```

6. Edit the `config-network` ConfigMap in the `knative-serving` namespace to enable automatic TLS certificate provisioning and specify how HTTP requests are handled. Run:

```
kubectl edit configmap config-network --namespace knative-serving
```



#### Note

For `HTTP01` challenges, the `http-protocol` field must be set to `Enabled` for the



cluster to accept `HTTP01` challenge requests.

```
apiVersion: v1
kind: ConfigMap
metadata:
  name: config-network
  namespace: knative-serving
data:
  ...
  auto-tls: Enabled
  ...
  http-protocol: Enabled
  ...
```

- To validate that your ConfigMap was updated successfully, run:

```
kubectl get configmap config-network --namespace knative-serving --output jsonpath="{.data.auto-tls}"
kubectl get configmap config-network --namespace knative-serving --output jsonpath="{.data.http-protocol}"
```

- Verify that your automatic TLS certificate instance is configured correctly by deploying a sample app. See [Verify Auto TLS](#) in the Knative documentation.

## Enable Auto TLS Using a DNS01 Challenge

The `DNS01` challenge validates that you control your domain's DNS by accessing and updating your domain's TXT record. You need to provide a `cert-manager` with your DNS API credentials. For a list of DNS01 providers supported for the ACME `Issuer`, see the [cert-manager documentation](#).



### Note

You can provision certificates **per service** only.

To enable automatic TLS certificate provisioning using a `DNS01` challenge, do the following:

- Set up credentials for `cert-manager` to access your DNS records. For information about setting up credentials for your ACME `Issuer` supported DNS provider, see [Supported DNS01 providers](#) in the `cert-manager` documentation. In the next step, you create an `Issuer` on `cert-manager` with the configuration you set up.
- Create a `cert-manager Issuer` or `ClusterIssuer` for DNS01 challenge on the `cert-manager Issuer` you set up in the previous step. The following example uses Let's Encrypt and Google Cloud DNS. For information about other DNS providers supported by `cert-manager`, see the [cert-manager documentation](#). The `Issuer` assumes that your Kubernetes secret holds credentials for the service account created. Run the following command to apply the `ClusterIssuer`:

```
kubectl apply --filename - <<EOF
apiVersion: cert-manager.io/v1
kind: ClusterIssuer
```

```

metadata:
  name: letsencrypt-dns-issuer
spec:
  acme:
    server: https://acme-v02.api.letsencrypt.org/directory
    # This will register an issuer with LetsEncrypt.
    email: MY-EMAIL
    privateKeySecretRef:
      # Set privateKeySecretRef to any unused secret name.
      name: letsencrypt-dns-issuer
    solvers:
      - dns01:
          cloudDNS:
            project: $PROJECT_ID
            # Set this to the secret that we publish our service account key
            # in the previous step.
            serviceAccountSecretRef:
              name: cloud-dns-key
              key: key.json
EOF

```

Where `MY-EMAIL` is your email address.

- To verify that your ClusterIssuer is created successfully, run:

```
kubectl get clusterissuer letsencrypt-dns-issuer --output yaml
```

- Edit your `config-certmanager` ConfigMap in the `knative-serving` namespace to reference the ClusterIssuer created in the previous step. Run:

```
kubectl edit configmap config-certmanager --namespace knative-serving
```

- Add an `issuerRef` block under the `data` section of your ConfigMap. This defines the ClusterIssuer Knative uses to issue certificates. Run:

```

...
data:
...
  issuerRef: |
    kind: ClusterIssuer
    name: letsencrypt-dns-issuer

```

- To validate that your file was updated successfully, run:

```
kubectl get configmap config-certmanager --namespace knative-serving --output jsonpath="{.data.issuerRef}"
```

- To enable automatic TLS certificate provisioning and specify how HTTP requests are handled, edit your `config-network` ConfigMap in the `knative-serving` namespace:

```
kubectl edit configmap config-network --namespace knative-serving
```



#### Note

When using the DNS01 challenge type, the `http-protocol` field must be set

to `Enabled`.

```
apiVersion: v1
kind: ConfigMap
metadata:
  name: config-network
  namespace: knative-serving
data:
  ...
  auto-tls: Enabled
  ...
  http-protocol: Enabled
  ...
```

8. Validate that your file was updated successfully. Run:

```
kubectl get configmap config-network --namespace knative-serving --output jsonpath="{.data.auto-tls}"
kubectl get configmap config-network --namespace knative-serving --output jsonpath="{.data.http-protocol}"
```

9. Verify that your automatic TLS certificate instance is functioning correctly by deploying a sample app. See [Verify Auto TLS](#) in the Knative documentation.

## Configuring Eventing with RabbitMQ for Cloud Native Runtimes

This topic tells you how to use RabbitMQ as an event source to react to messages sent to a RabbitMQ exchange or as an event broker to distribute events within your app for Cloud Native Runtimes, commonly known as CNR.

### Overview

The integration allows you to create:

- **A RabbitMQ broker:** A Knative Eventing broker backed by RabbitMQ. This broker uses RabbitMQ exchanges to store CloudEvents that are then routed from one component to another.
- **A RabbitMQ source:** An event source that translates external messages on a RabbitMQ exchange to CloudEvents, which can then be used with Knative Serving or Knative Eventing over HTTP.

## Install VMware Tanzu RabbitMQ for Kubernetes

Before you can use or test RabbitMQ eventing on Cloud Native Runtimes, you need to install VMware Tanzu RabbitMQ for Kubernetes. Follow below steps to complete the installation -

1. [Accept End User License Agreement](#)
2. [Prerequisites:](#) This step installs kapp-controller and secretgen-controller. Skip this step if kapp-controller and secretgen-controller are already installed on your cluster.

3. Prepare for the installation:
  - ◆ [Provide imagePullSecrets](#)
  - ◆ [Install the PackageRepository](#)
  - ◆ [Create a Service Account](#)
  - ◆ [Install Cert-Manager](#): Skip this step if cert-manager is already installed on your cluster.
4. [Install the Tanzu RabbitMQ Package](#): This step will install the Tanzu RabbitMQ Cluster Operator, Message Topology Operator, and Standby Replication Operator on your cluster.

## Next Steps

After completing these installations, you can:

- Verify your Knative Eventing installation using an example RabbitMQ broker. For instructions, see [Verify Knative Eventing](#).
- Create a broker, producer, and a consumer to use RabbitMQ and Knative Eventing with your own app.

## Configuring Observability for Cloud Native Runtimes

This topic tells you how to configure observability for Cloud Native Runtimes, commonly known as CNR.

### Overview

You can set up integrations with third-party observability tools to use logging, metrics, and tracing with Cloud Native Runtimes. These observability integrations allow you to monitor and collect detailed metrics from your clusters on Cloud Native Runtimes. You can collect logs and metrics for all workloads running on a cluster. This includes Cloud Native Runtimes components or any apps running on Cloud Native Runtimes. The integrations in this topic are recommended by VMware, however you can use any Kubernetes compatible logging, metrics, and tracing platforms to monitor your cluster workload.

### Logging

You can collect and forward logs for all workloads on a cluster, including Cloud Native Runtimes components or any apps running on Cloud Native Runtimes. You can use any logging platform that is compatible with Kubernetes to collect and forward logs for Cloud Native Runtimes workloads. VMware recommends using Fluent Bit to collect logs and then forward logs to vRealize. The following sections describe configuring logging for Cloud Native Runtimes with Fluent Bit and vRealize as an example.

### Configure Logging with Fluent Bit

You can use Fluent Bit to collect logs for all workloads on a cluster, including Cloud Native Runtimes components or any apps running on Cloud Native Runtimes. For more information about using

Fluent Bit logs, see [Fluent Bit Kubernetes Logging](#) in the Fluent Bit documentation.

Fluent Bit lets you collect logs from Kubernetes containers, add Kubernetes metadata to these logs, and forward logs to third-party log storage services. For more information about collecting logs, see [Logging](#) in the Knative documentation.

If you are using Tanzu Mission Control (TMC), vSphere 7.0 with Tanzu, or Tanzu Kubernetes Cluster to manage your cloud native environment, you must set up a role binding that grants required permissions to Fluent Bit containers in order to configure logging with any integration. Then, you can follow the instructions in the Fluent Bit documentation to complete the logging configuration. For more information about configuring Fluent Bit logging, see [Installation](#) in the Fluent Bit documentation.

To configure logging with Fluent Bit for your Cloud Native Runtimes environment:

1. VMware recommends that you add any integrations to the [ConfigMap](#) in both your Knative Serving and Knative Eventing namespaces. Follow the logging configuration steps in the Fluent Bit documentation to create the [Namespace](#), [ServiceAccount](#), [Role](#), [RoleBinding](#), and [ConfigMap](#). To view these steps, see [Installation](#) in the Fluent Bit documentation.
2. If you are using TMC, vSphere with Tanzu, or Tanzu Kubernetes Cluster to manage your cloud native environment, create a role binding in the Kubernetes namespace where your integration will be deployed to grant permission for privileged Fluent Bit containers. For information about creating a role binding on a Tanzu platform, see [Add a Role Binding](#) in the TMC documentation. For information about viewing your Kubernetes namespaces, see [Viewing Namespaces](#) in the Kubernetes documentation. Create the following role binding:

```
apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
  name: fluentbit-psp-rolebinding
  namespace: FLUENTBIT-NAMESPACE
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: PRIVILEGED-CLUSTERROLE
subjects:
- kind: ServiceAccount
  name: FLUENTBIT-SERVICEACCOUNT
  namespace: FLUENTBIT-NAMESPACE
```

Where:

- ◆ [FLUENTBIT-NAMESPACE](#) is your Fluent Bit namespace.
- ◆ [PRIVILEGED-CLUSTERROLE](#) is the name of your privileged cluster role.
- ◆ [FLUENTBIT-SERVICEACCOUNT](#) is your Fluent Bit service account.

3. To verify that you have configured logging successfully, run the following to access logs through your web browser:

```
kubectl port-forward --namespace logging service/log-collector 8080:80
```

For more information about accessing Fluent Bit logs, see [Logging](#) in the Knative documentation.

## Forward Logs to vRealize

After you configure log collection, you can forward logs to log management services. vRealize Log Insight is one service you can use with Cloud Native Runtimes. vRealize Log Insight is a scalable log management solution that provides log management, dashboards, analytics, and third-party extensibility for infrastructure and apps. For more information about vRealize Log Insight, see the [VMware vRealize Log Insight Documentation](#).

To forward logs from your Cloud Native Runtimes environment to vRealize, you can use a new or existing instance of Tanzu Kubernetes Cluster. For information about how to configure log forwarding to vRealize from Tanzu Kubernetes Cluster, see the [Configure Log forwarding from VMware Tanzu Kubernetes Cluster to vRealize Log Insight Cloud](#) blog.

## Metrics

Cloud Native Runtimes integrates with Prometheus and Tanzu Observability by Wavefront to collect metrics on components or apps. For more information about integrating with Prometheus, see [Overview](#) in the Prometheus documentation and [Kubernetes Integration](#) in the Wavefront documentation.

You can configure Prometheus endpoints on Cloud Native Runtimes components in order to be able to collect metrics on your components or apps. For information about configuring this, see the [Prometheus documentation](#).

You can use annotation based discovery with Prometheus to define which Kubernetes objects in your Cloud Native Runtimes environment to add metadata and collect metrics in a more automated way. For more information about using annotation based discovery, see [Annotation based discovery](#) in GitHub.

You can then use the Wavefront Collector for Kubernetes collector to dynamically discover and scrape pods with the `prometheus.io/scrape` annotation prefix. For information about the Kubernetes collector, see [Wavefront Collector for Kubernetes](#) in GitHub.



### Note

: All Cloud Native Runtimes related metrics are emitted with the prefix `tanzu.vmware.com/cloud-native-runtimes.*`.

## Tracing

Tracing is a method for understanding the performance of specific code paths in apps as they handle requests. You can configure tracing to collect performance metrics for your apps or Cloud Native Runtimes components. You can trace which aspects of Cloud Native Runtimes and workloads running on Cloud Native Runtimes are performing poorly.

## Configuring Tracing

You can configure tracing for your apps on Cloud Native Runtimes. To do this, you configure tracing for both Knative Serving and Eventing by editing the ConfigMap for your Knative namespace.

To configure tracing, do the following:

1. Configure the `config-tracing` ConfigMap in your Knative component namespace. VMware recommends that you add any integrations to the ConfigMap in both your Serving and Eventing namespaces. For information on how to enable request traces in each component, see the following Knative documentation:
  - ◊ Serving. See [Accessing request traces](#).
  - ◊ Eventing. See [Accessing CloudEvent traces](#).

## Forwarding Trace Data to a Data Visualization Tool

You can use the OpenTelemetry integration with Tanzu Observability by Wavefront to forward trace data to Tanzu Observability by Wavefront. For information about forwarding trace data, see [Sending Metrics Data to Wavefront](#) in the Wavefront documentation.

To configure to send trace data to Cloud Native Runtimes tracing with Tanzu Observability by Wavefront and the OpenTelemetry integration, do the following:

1. Use the following documentation to configure the OpenTelemetry Integration to send trace data to with Cloud Native Runtimes. For more information about sending trace data with OpenTelemetry, see [OpenTelemetry Integration](#) in the Wavefront documentation.
2. Deploy the Wavefront Proxy. For more information about wavefront proxies, see [Deploy a Wavefront Proxy in Kubernetes](#) in the Wavefront documentation.
  - ◊ Use the following .yaml file to install the Wavefront proxy in your Kubernetes cluster: [wavefront.yaml](#).
  - ◊ Provide the URL of your Wavefront instance and a [Wavefront token](#).
  - ◊ Uncomment the lines indicated in the yaml file to enable consumption of Zipkin traces.

## Sending Trace Data to an Observability Platform

You can send trace data to an observability and analytics platform to view and monitor your trace data in dashboards.

One way to do this is to integrate Tanzu Observability by Wavefront with your Cloud Native Runtimes environment. To view your trace data in Wavefront, you configure Cloud Native Runtimes to send traces to the Wavefront proxy and then configure the Wavefront proxy to consume Zipkin spans.

For more information about using Zipkin for tracing, see the [Zipkin](#) documentation.

You can send trace data from Cloud Native Runtimes to Wavefront by using Zipkin as the backend and defining the Zipkin endpoint as the Wavefront proxy URL listening over port `9411`. You configure Cloud Native Runtimes to send traces directly to the Wavefront proxy by editing the `zipkin-endpoint` property in the ConfigMap to point to the Wavefront proxy URL. You can configure the Wavefront proxy to consume Zipkin spans by listening to port `9411`.

To send trace data to Tanzu Observability by Wavefront:

1. Edit the ConfigMap to enable the Zipkin tracing integration. VMware recommends that you add any integrations to the ConfigMap in both your Serving and Eventing namespaces. Edit the Knative `config-tracing` ConfigMap to set `backend` to `zipkin` and pass the Wavefront proxy

URL in the zipkin-endpoint field:

```
Kubectl edit configmap config-tracing --namespace knative-serving apiVersion: v1
kind: ConfigMap
metadata:
  name: config-tracing
  ...
data:
  backend: "zipkin"
  zipkin-endpoint: "http://wavefront-proxy.default.svc.cluster.local:9411/api/v2/
spans"  ...
```

## Use Wavefront Dashboards

Cloud Native Runtimes provides two Wavefront dashboards in JSON format. You can use these dashboard to monitor your apps and investigate performance issues. For information about configuring dashboards, see [Create and Customize Dashboards](#) in the Wavefront documentation.

The following Wavefront dashboards are compatible with Cloud Native Runtimes: - Application Operator Service View. See [app-operator-service-view.json](#) in the Cloud Native Runtimes installation .tar file. - Application Operator Revision View. See [app-operator-revision-view.json](#) in the Cloud Native Runtimes installation .tar file.

To import a dashboard JSON file, use one of the following methods: - [Wavefront REST API](#) - [Wavefront CLIs](#).

You must provide the URL of your Wavefront instance and a Wavefront token. For more information about Wavefront tokens, see [Generating an API Token](#) in the Wavefront documentation.

## Import Wavefront Dashboards

You can import the Wavefront dashboards using either the Wavefront API or the Ruby Wavefront CLI. For more information about Wavefront dashboard, see [Import Dashboards with the Wavefront API](#) or [Import with a Ruby Wavefront CLI](#) below.

### Import Dashboards with the Wavefront API

To import a Wavefront dashboard with the Wavefront API, run:

```
curl -H "Content-Type: application/json" -H 'Authorization: Bearer <wavefront-token>' \
  https://<wavefront-instance>.wavefront.com/api/v2/dashboard -d @observability/wavefront/app-operator-service-view.json

curl -H "Content-Type: application/json" -H 'Authorization: Bearer <wavefront-token>' \
  https://<wavefront-instance>.wavefront.com/api/v2/dashboard -d @dashboards/wavefront/app-operator-revision-view.json
```

After you run the import code, the Wavefront API creates two dashboards with the following names and URLs:

- **Title:** [Cloud Native Runtimes App Operator - Service View](#)  
**URL:** [https://<wavefront-instance>.wavefront.com/dashboards/App-Operator-Service-](#)



Level

- **Title:** Cloud Native Runtimes App Operator - Revision View

**URL:** `https://<wavefront-instance>.wavefront.com/dashboards/App-Operator-Revision-Level`

## Import with the Ruby Wavefront CLI

To import a Wavefront dashboard with the Ruby Wavefront CLI, run:

```
export WAVEFRONT_TOKEN=<wavefront-token>
export WAVEFRONT_ENDPOINT=<wavefront-instance>.wavefront.com

wf config envvars
wf dashboard import observability/wavefront/app-operator-service-view.json
wf dashboard import dashboards/wavefront/app-operator-revision-view.json
```

After you run the import code, the Ruby Wavefront CLI creates two dashboards with a name and URL.

The Service View of the Cloud Native Runtimes App Operator dashboard will have the following title and URL:

- **Title:** Cloud Native Runtimes App Operator - Service View

**URL:** `https://<wavefront-instance>.wavefront.com/dashboards/App-Operator-Service-Level`

The Revision View of the Cloud Native Runtimes App Operator dashboard will have the following title and URL:

- **Title:** Cloud Native Runtimes App Operator - Revision View

**URL:** `https://<wavefront-instance>.wavefront.com/dashboards/App-Operator-Revision-Level`

## Configuring Cloud Native Runtimes with Avi Vantage

This topic tells you how to configure Cloud Native Runtimes, commonly known as CNR, with Avi Vantage.

### Overview

You can configure Cloud Native Runtimes to integrate with Avi Vantage. Avi Vantage is a multi-cloud platform that delivers features such as load balancing, security, and container ingress services. The Avi Controller provides a control plane while the Avi Service Engines provide a data plane. Once set up, the Avi Service Engines forward incoming traffic to your Kubernetes cluster's Envoy pods, which are created and managed by Contour.

For information about Avi Vantage, see [Avi Documentation](#).

## Integrate Avi Vantage with Cloud Native Runtimes

This procedure assumes that you have already installed Cloud Native Runtimes.

If you have not already installed Cloud Native Runtimes, see [Installing Cloud Native Runtimes](#). If you already have a Contour installation on your cluster, see [Installing Cloud Native Runtimes with an Existing Contour Installation](#).

To configure Cloud Native Runtimes with Avi Vantage, do the following:

1. Deploy the Avi Controller on any Avi supported infrastructure providers. For a list of Avi supported providers, see [Avi Installation Guides](#). For more information about deploying an Avi Controller, see [Install Avi Kubernetes Operator](#) in the Avi Vantage documentation.
2. Deploy the Avi Kubernetes Operator to your Kubernetes cluster where Cloud Native Runtimes is hosted. See [Install AKO for Kubernetes](#) in the Avi Vantage documentation.
3. Connect to a test app and verify that it is reachable. Run:

```
"curl -H KNATIVE-SERVICE-DOMAIN" ENVOY-IP
```

Where:

- ◆ `KNATIVE-SERVICE-DOMAIN` is the name of your domain.
- ◆ `ENVOY-IP` is the IP address of your Envoy instance.

For more information about deploy a sample application and connect to the application, see [Test Knative Serving](#).

4. (Optional) Create a DNS record that will configure your KService URL to point to the Avi Service Engines, and resolve to the external IP of the Envoy. You can create a DNS record on any platform that supports DNS services. Refer to the documentation for your DNS service platform for more information.

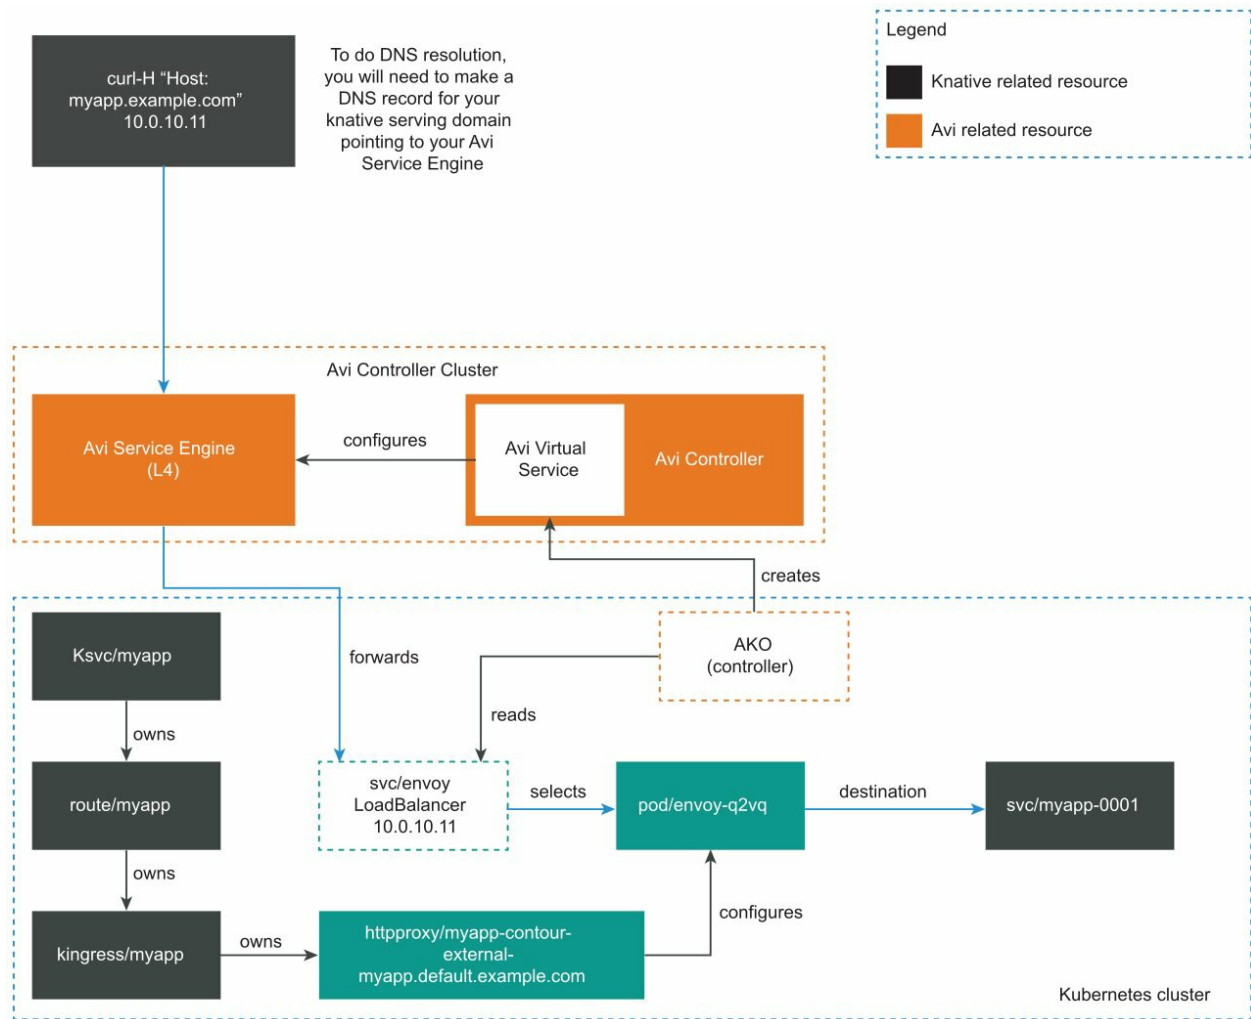
To get the KService URL, run:

```
kn route describe APP-NAME | grep "URL"
```

To get Envoy's external IP, follow step 3 in [Test Knative Serving](#) in *Verifying your Serving Installation*.

## About Routing with Avi Vantage and Cloud Native Runtimes

The following diagram shows how Avi Vantage integrates with Cloud Native Runtimes:



When Contour creates a Kubernetes LoadBalancer service for Envoy, the Avi Kubernetes Operator (AKO) sees the new LoadBalancer service. Then Avi Controller creates a Virtual Service. For information about LoadBalancer services, see [Type LoadBalancer](#) in the Kubernetes documentation.

For each Envoy service, Avi Controller creates a corresponding Virtual Service. See [Virtual Services](#) in the Avi Vantage documentation.

After Avi Controller creates a Virtual Service, the Controller configures the Avi Service Engines to forward traffic to the Envoy pods. The Envoy pods route traffic based on incoming requests, including traffic splitting and path based routing.

The Avi Controller provides Envoy with an external IP address so that apps are reachable by the app developer.

✍

**Note**

Avi does not interact directly with any Cloud Native Runtimes resources. Avi Vantage forwards all incoming traffic to Envoy.

## Configuring Cloud Native Runtimes with Tanzu Service Mesh

This topic tells you how to configure Cloud Native Runtimes, commonly known as CNR, with Tanzu Service Mesh.

## Overview

You cannot install Cloud Native Runtimes on a cluster that has Tanzu Service Mesh attached.

This workaround describes how Tanzu Service Mesh can be configured to ignore the Cloud Native Runtimes. This allows Contour to provide ingress routing for the Knative workloads, while Tanzu Service Mesh continues to satisfy other connectivity concerns.

**Note:** Cloud Native Runtimes workloads are unable to use Tanzu Service Mesh features like Global Namespace, Mutual Transport Layer Security authentication (mTLS), retries, and timeouts.

For information about Tanzu Service Mesh, see [Tanzu Service Mesh Documentation](#).

## Run Cloud Native Runtimes on a Cluster Attached to Tanzu Service Mesh

This procedure assumes you have a cluster attached to Tanzu Service Mesh, and that you have not yet installed Cloud Native Runtimes.

**Note:** If you installed Cloud Native Runtimes on a cluster that has Tanzu Service Mesh attached before doing the procedure below, pods fail to start. To fix this problem, follow the procedure below and then delete all pods in the excluded namespaces.

Configure Tanzu Service Mesh to ignore namespaces related to Cloud Native Runtimes:

1. Navigate to the **Cluster Overview** tab in the Tanzu Service Mesh UI.
2. On the cluster where you want to install Cloud Native Runtimes, click **...**, then select **Edit Cluster...**
3. Create an Is Exactly rule for each of the following namespaces:
  - ◆ CONTOUR-NS
  - ◆ knative-serving
  - ◆ knative-eventing
  - ◆ knative-sources
  - ◆ triggermesh
  - ◆ vmware-sources
  - ◆ tap-install
  - ◆ rabbitmq-system
  - ◆ kapp-controller
  - ◆ The namespace or namespaces where you plan to run Knative workloads.

Where CONTOUR-NS is the namespace(s) where Contour is installed on your cluster. If Cloud Native Runtimes was installed as part of a Tanzu Application Profile, this value will likely be `tanzu-system-ingress`.

## Next Steps

After configuring Tanzu Service Mesh, install Cloud Native Runtimes and verify your installation:

1. Install Cloud Native Runtimes. See [Installing Cloud Native Runtimes](#).
2. Verify your installation. See [Verifying Your Installation](#).

**Note:** You must create all Knative workloads in the namespace or namespaces where you plan to run these Knative workloads. If you do not, your pods fail to start.

## Troubleshooting Cloud Native Runtimes

This topic tells you how to troubleshoot Cloud Native Runtimes, commonly known as CNR, installation or configuration.

### Cannot connect to app on AWS

#### Symptom

On AWS, you see the following error when connecting to your app:

```
curl: (6) Could not resolve host: a*****7.us-west-2.elb.amazonaws.com
```

#### Solution

Try connecting to your app again after 5 minutes. The AWS LoadBalancer name resolution takes several minutes to propagate.

### minikube Pods Fail to Start

#### Symptom

On minikube, you see the following error when installing Cloud Native Runtimes:

```
3:03:59PM: error: reconcile job/contour-certgen-v1.10.0 (batch/v1) namespace: contour-internal
Pod watching error: Creating Pod watcher: Get "https://192.168.64.17:8443/api/v1/pods?labelSelector=kapp.k14s.io%2Fapp%3D1618232545704878000&watch=true": dial tcp 192.168.64.17:8443: connect: connection refused
kapp: Error: waiting on reconcile job/contour-certgen-v1.10.0 (batch/v1) namespace: CONTOUR-NS:
  Errored:
    Listing schema.GroupVersionResource{Group:"", Version:"v1", Resource:"pods"}, namespace: true:
      Get "https://192.168.64.17:8443/api/v1/pods?labelSelector=kapp.k14s.io%2Fassociation%3Dv1.572a543d96e0723f858367fcf8c6af4e": unexpected EOF
```

Where CONTOUR-NS is the namespace where Contour is installed on your cluster. If Cloud Native Runtimes was installed as part of a Tanzu Application Profile, this value will likely be `tanzu-system-ingress`.

## Solution

Increase your available system RAM to at least 4 GB.

## Pulling an image with `imgpkg` overwrites the `cloud-native-runtimes` directory

### Symptom

When relocating an image to a private registry and later pulling that image with `imgpkg pull --lock LOCK-OUTPUT -o ./cloud-native-runtimes`, the contents of the `cloud-native-runtimes` are overwritten.

### Solution

Upgrade the `imgpkg` version to v0.13.0 or later.

## Installation fails to reconcile `app/cloud-native-runtimes`

### Symptom

When installing Cloud Native Runtimes, you see one of the following errors:

```
11:41:16AM: ongoing: reconcile app/cloud-native-runtimes (kappctrl.k14s.io/v1alpha1) namespace: cloud-native-runtime
11:41:16AM: ^ Waiting for generation 1 to be observed
kapp: Error: Timed out waiting after 15m0s
```

Or,

```
3:15:34PM: ^ Reconciling
3:16:09PM: fail: reconcile app/cloud-native-runtimes (kappctrl.k14s.io/v1alpha1) namespace: cloud-native-runtimes
3:16:09PM: ^ Reconcile failed: (message: Deploying: Error (see .status.usefulErrorMessage for details))

kapp: Error: waiting on reconcile app/cloud-native-runtimes (kappctrl.k14s.io/v1alpha1) namespace: cloud-native-runtimes:
  Finished unsuccessfully (Reconcile failed: (message: Deploying: Error (see .status.usefulErrorMessage for details)))
```

## Explanation

The `cloud-native-runtimes` deployment app installs the subcomponents of Cloud Native Runtimes. Error messages about reconciling indicate that one or more subcomponents have failed to install.

## Solution

Use the following procedure to examine logs:

1. Get the logs from the `cloud-native-runtimes` app by running:

```
kubectl get app/cloud-native-runtimes -n cloud-native-runtimes -o jsonpath="{.s
```

```
tatus.deploy.stdout}"
```

**Note:** If the command does not return log messages, then kapp-controller is not installed or is not running correctly.

2. Review the output for sub component deployments that have failed or are still ongoing. See the examples below for suggestions on resolving common problems.

### Example 1: The Cloud Provider does not support the creation of Service type LoadBalancer

Follow these steps to identify and resolve the problem of the cloud provider not supporting services of type `LoadBalancer`:

1. Search the log output for `Load balancer`, for example by running:

```
kubectl -n cloud-native-runtimes get app cloud-native-runtimes -ojsonpath="{.status.deploy.stdout}" | grep "Load balancer" -C 1
```

2. If the output looks similar to the following, ensure that your cloud provider supports services of type `LoadBalancer`. For more information, see [Prerequisites](#).

```
6:30:22PM: ongoing: reconcile service/envoy (v1) namespace: CONTOUR-NS
6:30:22PM: ^ Load balancer ingress is empty
6:30:29PM: ---- waiting on 1 changes [322/323 done] ----
```

Where `CONTOUR-NS` is the namespace where Contour is installed on your cluster. If Cloud Native Runtimes was installed as part of a Tanzu Application Profile, this value will likely be `tanzu-system-ingress`.

### Example 2: The webhook deployment failed

Follow these steps to identify and resolve the problem of the `webhook` deployment failing in the `vmware-sources` namespace:

1. Review the logs for output similar to the following:

```
10:51:58PM: ok: reconcile customresourcedefinition/httpproxies.projectcontour.io (apiextensions.k8s.io/v1) cluster
10:51:58PM: fail: reconcile deployment/webhook (apps/v1) namespace: vmware-sources
10:51:58PM: ^ Deployment is not progressing: ProgressDeadlineExceeded (message: ReplicaSet "webhook-6f5d979b7d" has timed out progressing.)
```

2. Run `kubectl get pods` to find the name of the pod:

```
kubectl get pods --show-labels -n NAMESPACE
```

Where `NAMESPACE` is the namespace associated with the reconcile error, for example, `vmware-sources`.

For example,

```
$ kubectl get pods --show-labels -n vmware-sources
NAME                                READY   STATUS    RESTARTS   AGE   LABELS
webhook-6f5d979b7d-cxr9k           0/1    Pending   0          44h   app=webhook,kapp.k14s.io/app=1626302357703846007,kapp.k14s.io/association=v1.9621e0a793b4e925077dd557acedbcfe,pod-template-hash=6f5d979b7d,role=webhook,sources.tanzu.vmware.com/release=v0.23.0
```

### 3. Run `kubectl logs` and `kubectl describe`:

```
kubectl logs PODNAME -n NAMESPACE
kubectl describe pod PODNAME -n NAMESPACE
```

Where:

- ◆ `PODNAME` is found in the output of step 3, for example `webhook-6f5d979b7d-cxr9k`.
- ◆ `NAMESPACE` is the namespace associated with the reconcile error, for example, `vmware-sources`.

For example:

```
$ kubectl logs webhook-6f5d979b7d-cxr9k -n vmware-sources

$ kubectl describe pod webhook-6f5d979b7d-cxr9k -n vmware-sources
Events:
  Type            Reason             Age          From              Message
  ----            -
Warning          FailedScheduling   80s (x14 over 14m)  default-scheduler  0/1 nodes are available: 1 Insufficient cpu.
```

### 4. Review the output from the `kubectl logs` and `kubectl describe` commands and take further action.

For this example of the webhook deployment, the output indicates that the scheduler does not have enough CPU to run the pod. In this case, the solution is to add nodes or CPU cores to the cluster. If you are using Tanzu Mission Control (TMC), increase the number of workers in the node pool to three or more through the TMC UI. See [Edit a Node Pool](#), in the TMC documentation.

## Cloud Native Runtimes Installation Fails with Existing Contour Installation

### Symptom

You see the following error message when you run the install script:

```
Could not proceed with installation. Refer to Cloud Native Runtimes documentation for details on how to utilize an existing Contour installation. Another app owns the custom resource definitions listed below.
```

### Solution

Follow the procedure in [Install Cloud Native Runtimes on a Cluster with Your Existing Contour Instances](#) to resolve the issue.



# Knative Serving Workloads fail to deploy with Forbidden errors on Run Cluster when using Multi Cluster setup on Openshift

This only applies to CNR 2.0.1

## Symptom

You see the following error message when you describe a workload:

```
...
pods "<pod name>" is forbidden: unable to validate against any security context constraint: [provider "anyuid": Forbidden: not usable by user or serviceaccount, spec.containers[0].securityContext.runAsUser: Invalid value: 1000: must be in the ranges: [1000740000, 1000749999]]
...
```

## Solution

The service account (or user) used to create the workload is bound to the “restricted” SecurityContextConstraint, and thus is not able to run a pod with a UserID outside the set range.

You must bind the service accounts on your cluster to the “nonroot” SecurityContextConstraint to allow running as any nonroot UserID.

You can apply the following YAML to the `run` cluster to achieve this:

```
---
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRole
metadata:
  name: cnrs-nonroot
rules:
- apiGroups:
  - security.openshift.io
  resourceNames:
  - nonroot
  resources:
  - securitycontextconstraints
  verbs:
  - use
---
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  name: cnrs-nonroot-crb
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: cnrs-nonroot
subjects:
- apiGroup: rbac.authorization.k8s.io
  kind: Group
  name: system:serviceaccounts
```

This will associate the `nonroot` SecurityContextConstraint to all ServiceAccounts. If you know the specific ServiceAccounts used to deploy workloads, you can modify the `.subjects` section of the ClusterRoleBinding above with specific subjects:

```
subjects:  
- kind: ServiceAccount  
  name: <sa name>  
  namespace: <workload namespace>
```

# Verifying Your Installation

This topic tells you how to verify your Cloud Native Runtimes, commonly known as CNR, installation. You can verify that your Cloud Native Runtimes installation was successful by testing Knative Serving, Knative Eventing, and TriggerMesh Sources for Amazon Web Services (SAWS).

## Prerequisites

1. Create a namespace and environment variable where you want to create Knative services.  
Run:

**Note:** This step covers configuring a namespace to run Knative services. If you rely on a SupplyChain to deploy Knative services into your cluster, skip this step because namespace configuration is covered in [Set up developer namespaces to use installed packages](#). Otherwise, you must complete the following steps for each namespace where you create Knative services.

```
export WORKLOAD_NAMESPACE='cnr-demo'
kubectl create namespace ${WORKLOAD_NAMESPACE}
```

2. Configure a namespace to use Cloud Native Runtimes. If during the TAP installation you relocated images to another registry, you must grant service accounts that run Knative services using Cloud Native Runtimes access to the image pull secrets. This includes the `default` service account in a namespace, which is created automatically but not associated with any image pull secrets. Without these credentials, attempts to start a service fail with a timeout and the pods report that they are unable to pull the `queue-proxy` image.
  1. Create an image pull secret in the namespace Knative services will run and fill it from the `tap-registry` secret mentioned in [Add the Tanzu Application Platform package repository](#). Run the following commands to create an empty secret and annotate it as a target of the secretgen controller:

```
kubectl create secret generic pull-secret --from-literal=.dockerconfigjso
n={} --type=kubernetes.io/dockerconfigjson -n ${WORKLOAD_NAMESPACE}

kubectl annotate secret pull-secret secretgen.carvel.dev/image-pull-secre
t="" -n ${WORKLOAD_NAMESPACE}
```

2. After you create a `pull-secret` secret in the same namespace as the service account, run the following command to add the secret to the service account:

```
kubectl patch serviceaccount default -p '{"imagePullSecrets": [{"name": "
pull-secret"}]}' -n ${WORKLOAD_NAMESPACE}
```

- Verify that a service account is correctly configured by running:

```
kubectl describe serviceaccount default -n ${WORKLOAD_NAMESPACE}
```

For example:

```
kubectl describe sa default -n cnr-demo
Name:                default
Namespace:           cnr-demo
Labels:              <none>
Annotations:         <none>
Image pull secrets:  pull-secret
Mountable secrets:   default-token-xh6p4
Tokens:              default-token-xh6p4
Events:              <none>
```

The service account has access to the `pull-secret` image pull secret.

Verify that `STATUS` is `Reconcile succeeded`.

## Verify Installation of Knative Serving, Knative Eventing, and TriggerMesh SAWS

To verify the installation of Knative Serving, Knative Eventing, and Triggermesh SAWS:

- Create a namespace and environment variable for the test. Run:

```
export WORKLOAD_NAMESPACE='cnr-demo'
kubectl create namespace ${WORKLOAD_NAMESPACE}
```

- Verify installation of the components that you intend to use:

To test...	Create...	For instructions, see...
Knative Serving	a test service	<a href="#">Verifying Knative Serving</a>
Knative Eventing	a broker, a producer, and a consumer	<a href="#">Verifying Knative Eventing</a>
TriggerMesh SAWS	an AWS source and trigger it	<a href="#">Verifying TriggerMesh SAWS</a>

- Delete the namespace that you created for the demo. Run:

```
kubectl delete namespaces ${WORKLOAD_NAMESPACE}
unset WORKLOAD_NAMESPACE
```

## Verifying Knative Serving for Cloud Native Runtimes

This topic tells you how to verify that Knative Serving was successfully installed for Cloud Native Runtimes, commonly known as CNR.

### About Verifying Knative Serving

To verify that Knative Serving was successfully installed, create an example Knative service and test it.

The procedure below shows you how to create an example Knative service using the Cloud Native Runtimes sample app, `hello-yeti`. This sample is custom built for Cloud Native Runtimes and is stored in the VMware Harbor registry.

**Note:** If you do not have access to the Harbor registry, you can use the [Hello World - Go](#) sample app in the Knative documentation.

## Prerequisites

Before you verify Knative Serving, you must have a namespace where you want to deploy Knative services. This namespace will be referred as `_${WORKLOAD_NAMESPACE}` in this tutorial. See step 1 of [Verifying Your Installation](#) for more information.

## Test Knative Serving

To create an example Knative service and use it to test Knative Serving:

1. If you are verifying on Tanzu Mission Control or vSphere 7.0 with Tanzu, then create a role binding in the `_${WORKLOAD_NAMESPACE}` namespace. Run:

```
kubectl apply -n "$_${WORKLOAD_NAMESPACE}" -f - << EOF
---
apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
  name: $_${WORKLOAD_NAMESPACE}-psp
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: cnr-restricted
subjects:
- kind: Group
  name: system:serviceaccounts:$_${WORKLOAD_NAMESPACE}
EOF
```

2. Deploy the sample app using the `kn` CLI. Run:

```
kn service create hello-yeti -n $_${WORKLOAD_NAMESPACE} \
  --image projects.registry.vmware.com/tanzu_serverless/hello-yeti@sha256:17d64
0edc48776cfc604a14fbabf1b4f88443acc580052eef3a753751ee31652 --env TARGET='hello
-yeti'
```

If you are verifying on Tanzu Mission Control or vSphere 7.0 with Tanzu, then add `--user 1001` to the command above to run it as a non-root user.

3. Run one of the following commands to retrieve the external address for your ingress, depending on your IaaS:

IaaS	◆	Tanzu Kubernetes Grid on AWS
:	◆	Tanzu Mission Control on AWS
	◆	Amazon Elastic Kubernetes Service

```
Run :
export EXTERNAL_ADDRESS=$(kubectl get service envoy -n tanzu-system-ingre
ss \
-o jsonpath='{.status.loadBalancer.ingress[0].hostname}')
```

- iaaS** :
- ◆ vSphere 7.0 on Tanzu
  - ◆ Tanzu Kubernetes Grid on vSphere/Azure/GCP
  - ◆ Tanzu Kubernetes Grid Integrated Edition
  - ◆ Tanzu Mission Control on vSphere
  - ◆ Azure Kubernetes Service
  - ◆ Google Kubernetes Engine

```
Run :
export EXTERNAL_ADDRESS=$(kubectl get service envoy -n tanzu-system-ingre
ss \
-o jsonpath='{.status.loadBalancer.ingress[0].ip}')
```

- iaaS:** Local Kubernetes Cluster:
- ◆ Docker desktop
  - ◆ Minikube

```
Run:
export EXTERNAL_ADDRESS='localhost:8080'
```

And set up port-forwarding in a separate terminal:

```
kubectl -n tanzu-system-ingress port-forward svc/envoy 8080:80
```

4. Connect to the app. Run:

```
curl -H "Host: hello-yeti.${WORKLOAD_NAMESPACE}.example.com" ${EXTERNAL_ADDRESS}
```

If external DNS is correctly configured, you can also visit the URL in a web browser.

On success, you see a reply from our mascot, Carl the Yeti.

## Delete the Example Knative Service

After verifying your serving installation, delete the example Knative service and unset the environment variable:

1. Run:

```
kn service delete hello-yeti -n ${WORKLOAD_NAMESPACE}
unset EXTERNAL_ADDRESS
```

2. If you created port forwarding in step 3 above, then terminate that process.

## Verify Knative Eventing with Cloud Native Runtimes

Eventing in Tanzu Application Platform is deprecated and marked for removal in Tanzu Application Platform v1.7.0.

This topic tells you how to verify that Knative Eventing was successfully installed with Cloud Native Runtimes, commonly known as CNR.

### About Verifying Knative Eventing

You can verify Knative Eventing by setting up a broker, creating a producer, and creating a consumer. If your installation was successful, you can create a test eventing workflow and see that the events appear in the logs.

You can use either an in-memory broker or a RabbitMQ broker to verify Knative Eventing:

- **RabbitMQ broker:** Using a RabbitMQ broker to verify Knative Eventing is a scalable and reliable way to verify your installation. Verifying with RabbitMQ uses methods similar to production environments.
- **In-memory broker:** Using an in-memory broker is a fast and lightweight way to verify that the basic elements of Knative Eventing are installed. An in-memory broker is not meant for production environments or for use with apps that you intend to take to production.

### Prerequisites

Before you verify Knative Eventing, you must:

- Have a namespace where you want to deploy your Knative resources. This namespace will be referred as `${WORKLOAD_NAMESPACE}` in this tutorial. See step 1 of [Verifying Your Installation](#) for more information.
- Create the following role binding in the `${WORKLOAD_NAMESPACE}` namespace. Run:

```
cat <<EOF | kubectl apply -f -
apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
  name: ${WORKLOAD_NAMESPACE}-psp
  namespace: ${WORKLOAD_NAMESPACE}
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: eventing-restricted
subjects:
- kind: Group
  name: system:serviceaccounts:${WORKLOAD_NAMESPACE}
EOF
```

### Prepare the RabbitMQ Environment

If you are using a RabbitMQ broker to verify Knative Eventing, follow the procedure in this section. If you are verifying with the in-memory broker, skip to [Verify Knative Eventing](#).

To prepare the RabbitMQ environment before verifying Knative Eventing:

1. Set up the RabbitMQ integration as described in [Configuring Eventing with RabbitMQ](#).
2. On the Kubernetes cluster where Eventing is installed, deploy a RabbitMQ cluster using the RabbitMQ Cluster Operator by running:

```
cat <<EOF | kubectl apply -f -
apiVersion: rabbitmq.com/v1beta1
kind: RabbitmqCluster
metadata:
  name: my-rabbitmq
  namespace: ${WORKLOAD_NAMESPACE}
spec:
  replicas: 1
  override:
    statefulSet:
      spec:
        template:
          spec:
            securityContext: {}
            containers:
            - name: rabbitmq
              env:
                - name: ERL_MAX_PORTS
                  value: "4096"
            initContainers:
            - name: setup-container
              securityContext:
                runAsUser: 999
                runAsGroup: 999
EOF
```

**Note:** The `override` section can be omitted if your cluster allows containers to run as `root`.

3. Create a RabbitmqBrokerConfig

```
cat <<EOF | kubectl apply -f -
apiVersion: eventing.knative.dev/v1alpha1
kind: RabbitmqBrokerConfig
metadata:
  name: default-config
  namespace: ${WORKLOAD_NAMESPACE}
spec:
  rabbitmqClusterReference:
    name: my-rabbitmq
    namespace: ${WORKLOAD_NAMESPACE}
  queueType: quorum
EOF
```

## Verify Knative Eventing

To verify installation of Knative Eventing create and test a broker, procedure, and consumer in the `${WORKLOAD_NAMESPACE}` namespace:

1. Create a broker.



For the RabbitMQ broker. Run:

```
cat <<EOF | kubectl apply -f -
apiVersion: eventing.knative.dev/v1
kind: Broker
metadata:
  name: default
  namespace: ${WORKLOAD_NAMESPACE}
  annotations:
    eventing.knative.dev/broker.class: RabbitMQBroker
spec:
  config:
    apiVersion: eventing.knative.dev/v1alpha1
    kind: RabbitmqBrokerConfig
    name: default-config
    namespace: ${WORKLOAD_NAMESPACE}
EOF
```

For the in-memory broker. Run:

```
cat <<EOF | kubectl create -f -
apiVersion: eventing.knative.dev/v1
kind: Broker
metadata:
  name: default
  namespace: ${WORKLOAD_NAMESPACE}
EOF
```

## 2. Create a consumer for the events. Run:

```
cat <<EOF | kubectl create -f -
---
apiVersion: apps/v1
kind: Deployment
metadata:
  labels:
    app: event-display
  name: event-display
  namespace: ${WORKLOAD_NAMESPACE}
spec:
  replicas: 1
  selector:
    matchLabels:
      app: event-display
  template:
    metadata:
      labels:
        app: event-display
    spec:
      containers:
        - name: user-container
          image: gcr.io/knative-releases/knative.dev/eventing-contrib/cmd/event
_display
          ports:
            - containerPort: 8080
              name: user-port
              protocol: TCP
---
```

```

apiVersion: v1
kind: Service
metadata:
  labels:
    app: event-display
  name: event-display-service
  namespace: ${WORKLOAD_NAMESPACE}
spec:
  ports:
    - port: 80
      protocol: TCP
      targetPort: 8080
  selector:
    app: event-display
EOF

```

### 3. Create a trigger. Run:

```

cat <<EOF | kubectl apply -f -
apiVersion: eventing.knative.dev/v1
kind: Trigger
metadata:
  name: event-display
  namespace: ${WORKLOAD_NAMESPACE}
spec:
  broker: default
  subscriber:
    ref:
      apiVersion: v1
      kind: Service
      name: event-display-service
      namespace: ${WORKLOAD_NAMESPACE}
EOF

```

### 4. Create a producer. This will send a message every minute. Run:

```

cat <<EOF | kubectl create -f -
apiVersion: sources.knative.dev/v1
kind: PingSource
metadata:
  name: test-ping-source
  namespace: ${WORKLOAD_NAMESPACE}
spec:
  schedule: "*/1 * * * *"
  data: '{"message": "Hello Eventing!"}'
  sink:
    ref:
      apiVersion: eventing.knative.dev/v1
      kind: Broker
      name: default
      namespace: ${WORKLOAD_NAMESPACE}
EOF

```

### 5. Verify that the event appears in your consumer logs. Run:

```

kubectl logs deploy/event-display -n ${WORKLOAD_NAMESPACE} --since=10m --tail=5
0 -f

```

## Setup RabbitMQ Broker as the default in the cluster (optional)

Eventing provides a `config-br-defaults` ConfigMap that contains the configuration setting that govern default Broker creation.

This example configuration will set RabbitMQ as the default broker on the cluster:

```
apiVersion: v1
kind: ConfigMap
metadata:
  name: config-br-defaults
  namespace: knative-eventing
  labels:
    eventing.knative.dev/release: devel
data:
  default-br-config: |
    clusterDefault:
      brokerClass: RabbitMQBroker
      apiVersion: eventing.knative.dev/v1alpha1
      kind: RabbitmqBrokerConfig
      name: default-config
      namespace: ${WORKLOAD_NAMESPACE} # This should be the name of your namespace.
      delivery:
        retry: 3
        backoffDelay: PT0.2S
        backoffPolicy: exponential
```

To achieve this you can:

1. Run:

```
kubectl patch -n knative-eventing cm config-br-defaults --type merge --patch '{
  "data": {"default-br-config": "clusterDefault:\n      brokerClass: RabbitMQBrok
er\n      apiVersion: eventing.knative.dev/v1alpha1\n      kind: RabbitmqBroker
Config\n      name: default-config\n      namespace: \"${WORKLOAD_NAMESPACE}\"'\n
n      delivery:\n      retry: 3\n      backoffDelay: PT0.2S\n      backo
ffPolicy: exponential\n"}}'
```

2. Check that the ConfigMap looks as intended.

```
kubectl get -n knative-eventing cm config-br-defaults -o yaml
```

3. Now, to create a RabbitMQ broker you can run:

```
cat <<EOF | kubectl create -f -
apiVersion: eventing.knative.dev/v1
kind: Broker
metadata:
  name: broker-using-defaults
  namespace: ${WORKLOAD_NAMESPACE}
EOF
```

Eventing will automatically set the `brokerClass` to `RabbitMQBroker` and it will set up the `spec.config` to `RabbitmqBrokerConfig` with name `default-config`.

## Setup RabbitMQ Broker as the default in a namespace (optional)

You can also use the `config-br-defaults` ConfigMap to set up the default broker configuration for a given namespace.

Let us suppose you want to have the MTChannelBroker as the default for the cluster and the RabbitMQ Broker as the default for your workload namespace.

To do this, we want that our `config-br-defaults` ConfigMap looks like this:

```
apiVersion: v1
kind: ConfigMap
metadata:
  name: config-br-defaults
  namespace: knative-eventing
  labels:
    eventing.knative.dev/release: devel
data:
  default-br-config: |
    clusterDefault:
      brokerClass: MTChannelBasedBroker
      apiVersion: v1
      kind: ConfigMap
      name: config-br-default-channel
      namespace: knative-eventing
      delivery:
        retry: 10
        backoffDelay: PT0.2S
        backoffPolicy: exponential
    namespaceDefaults:
      ${WORKLOAD_NAMESPACE}: # This should be the name of your namespace.
        brokerClass: RabbitMQBroker
        apiVersion: eventing.knative.dev/v1alpha1
        kind: RabbitmqBrokerConfig
        name: default-config
        namespace: ${WORKLOAD_NAMESPACE} # This should be the name of your namespace.
        delivery:
          retry: 3
          backoffDelay: PT0.2S
          backoffPolicy: exponential
```

To achieve this you can:

1. Run:

```
kubectl patch -n knative-eventing cm config-br-defaults --type merge --patch '{
  "data": {"default-br-config": "clusterDefault:\n brokerClass: MTChannelBasedBr
  oker\n apiVersion: v1\n kind: ConfigMap\n name: config-br-default-channel\n
  namespace: knative-eventing\n delivery:\n   retry: 10\n   backoffDelay: PT0
  .2S\n   backoffPolicy: exponential\nnamespaceDefaults:\n   \"${WORKLOAD_NAMESPA
  CE}\"":\n     brokerClass: RabbitMQBroker\n     apiVersion: eventing.knative.dev/v
  1alpha1\n     kind: RabbitmqBrokerConfig\n     name: default-config\n     namespac
  e: \"${WORKLOAD_NAMESPACE}\"'\n     delivery:\n       retry: 3\n       backoffDelay
  : PT0.2S\n       backoffPolicy: exponential\n"}'}
```

2. Check that the ConfigMap looks as intended.

```
kubectl get -n knative-eventing cm config-br-defaults -o yaml
```

3. With this configuration when you create a Broker in the `default` namespace it will be a `MTChannelBasedBroker`.

```
cat <<EOF | kubectl create -f -
apiVersion: eventing.knative.dev/v1
kind: Broker
metadata:
  name: broker-in-default-ns
  namespace: default
EOF
```

4. Check the type of this broker like so:

```
kubectl get -n default broker broker-in-default-ns -o yaml
```

It will show something like this:

```
apiVersion: eventing.knative.dev/v1
kind: Broker
metadata:
  annotations:
    eventing.knative.dev/broker.class: MTChannelBasedBroker
  name: broker-in-default-ns
  namespace: default
spec:
  config:
    apiVersion: v1
    kind: ConfigMap
    name: config-br-default-channel
    namespace: knative-eventing
  delivery:
    backoffDelay: PT0.2S
    backoffPolicy: exponential
    retry: 10
```

Notice the `eventing.knative.dev/broker.class: MTChannelBasedBroker` annotation.

5. Now try to create a Broker in the `${WORKLOAD_NAMESPACE}`.

```
cat <<EOF | kubectl create -f -
apiVersion: eventing.knative.dev/v1
kind: Broker
metadata:
  name: broker-in-workload-ns
  namespace: ${WORKLOAD_NAMESPACE}
EOF
```

6. Check the type of this broker like so:

```
kubectl get -n ${WORKLOAD_NAMESPACE} broker broker-in-workload-ns -o yaml
```

It will show something like this:

```

apiVersion: eventing.knative.dev/v1
kind: Broker
metadata:
  annotations:
    eventing.knative.dev/broker.class: RabbitMQBroker
  name: broker-in-workload-ns
  namespace: ${WORKLOAD_NAMESPACE}
spec:
  config:
    apiVersion: eventing.knative.dev/v1alpha1
    kind: RabbitmqBrokerConfig
    name: default-config
    namespace: ${WORKLOAD_NAMESPACE}
  delivery:
    backoffDelay: PT0.2S
    backoffPolicy: exponential
    retry: 3

```

Notice the `eventing.knative.dev/broker.class: RabbitMQBroker` annotation.

## Delete the Eventing Resources

After verifying your Eventing installation, clean up by deleting the resources used for the test:

1. Delete the eventing resources:

```

kubectl delete pingsource/test-ping-source -n ${WORKLOAD_NAMESPACE}
kubectl delete trigger/event-display -n ${WORKLOAD_NAMESPACE}
kubectl delete service/event-display-service -n ${WORKLOAD_NAMESPACE}
kubectl delete deploy/event-display -n ${WORKLOAD_NAMESPACE}
kubectl delete broker/default -n ${WORKLOAD_NAMESPACE}

```

2. If you followed [Setup RabbitMQ Broker as the default in the cluster \(optional\)](#), delete the broker like so:

```

kubectl delete broker/broker-using-defaults -n ${WORKLOAD_NAMESPACE}

```

3. If you followed [Setup RabbitMQ Broker as the default in a namespace \(optional\)](#), delete the brokers like so:

```

kubectl delete broker/broker-in-default-ns -n default
kubectl delete broker/broker-in-workload-ns -n ${WORKLOAD_NAMESPACE}

```

4. If you created a RabbitMQ cluster:

```

kubectl delete rabbitmqbrokerconfig/default-config -n ${WORKLOAD_NAMESPACE}
kubectl delete rabbitmqcluster/my-rabbitmq -n ${WORKLOAD_NAMESPACE}

```

5. Delete the role binding:

```

kubectl delete rolebinding/${WORKLOAD_NAMESPACE}-psp -n ${WORKLOAD_NAMESPACE}

```

## Verifying TriggerMesh SAWS for Cloud Native Runtimes

This topic tells you how to verify that TriggerMesh Sources for Amazon Web Services (SAWS) was installed successfully for Cloud Native Runtimes, commonly known as CNR.

## Overview

TriggerMesh SAWS allows you to consume events from your AWS services and send them to workloads running in your cluster.

Cloud Native Runtimes includes an installation of the Triggermesh SAWS controller and CRDs. You can find the controller in the `triggermesh` namespace.

For general information about TriggerMesh SAWS, see [aws-event-sources](#) in GitHub.

The procedure below shows you how to test TriggerMesh SAWS using the example of an event source for Amazon CodeCommit. If you want to test using a different AWS service, see [samples](#) in GitHub. The basic steps are the same, regardless of the AWS service you choose: create a broker, trigger, and consumer and then test.

## Prerequisites

Before you verify TriggerMesh SAWS with AWS CodeCommit, you must have:

- An AWS service account
- An AWS CodeCommit repository with push and pull access
- Have a namespace where you want to deploy Knative services. This namespace will be referred as `${WORKLOAD_NAMESPACE}` in this tutorial. See step 1 of [Verifying Your Installation](#) for more information.

## Verify TriggerMesh SAWS

To verify TriggerMesh SAWS with AWS CodeCommit:

1. Create a broker:

```
kubectl apply -f - << EOF
apiVersion: eventing.knative.dev/v1
kind: Broker
metadata:
  name: broker
  namespace: ${WORKLOAD_NAMESPACE}
EOF
```

2. Create a trigger:

```
kubectl apply -f - << EOF
---
apiVersion: eventing.knative.dev/v1
kind: Trigger
metadata:
  name: trigger
  namespace: ${WORKLOAD_NAMESPACE}
spec:
  broker: broker
```

```

subscriber:
  ref:
    apiVersion: serving.knative.dev/v1
    kind: Service
    name: consumer
    namespace: ${WORKLOAD_NAMESPACE}
EOF

```

### 3. Create a consumer:

```

kubectl apply -f - << EOF
---
apiVersion: serving.knative.dev/v1
kind: Service
metadata:
  name: consumer
  namespace: ${WORKLOAD_NAMESPACE}
spec:
  template:
    spec:
      containers:
        - image: gcr.io/knative-releases/knative.dev/eventing-contrib/cmd/event
          _display
EOF

```

### 4. Add an AWS service account secret:

```

kubectl -n ${WORKLOAD_NAMESPACE} create secret generic awscreds \
--from-literal=aws_access_key_id=${AWS_ACCESS_KEY_ID} \
--from-literal=aws_secret_access_key=${AWS_SECRET_ACCESS_KEY}

```

Where:

- ◆ `AWS_ACCESS_KEY_ID` is the AWS access key ID for your AWS service account.
- ◆ `AWS_SECRET_ACCESS_KEY` is your AWS access key for your AWS service account.

### 5. Create the AWSCodeCommitSource:

```

kubectl apply -f - << EOF
apiVersion: sources.triggermesh.io/v1alpha1
kind: AWSCodeCommitSource
metadata:
  name: source
  namespace: ${WORKLOAD_NAMESPACE}
spec:
  arn: ARN
  branch: BRANCH

  eventTypes:
    - push
    - pull_request

  credentials:
    accessKeyID:
      valueFromSecret:
        name: awscreds
        key: aws_access_key_id

```



```

secretAccessKey:
  valueFromSecret:
    name: awscreds
    key: aws_secret_access_key

sink:
  ref:
    apiVersion: eventing.knative.dev/v1
    kind: Broker
    name: broker
    namespace: ${WORKLOAD_NAMESPACE}
EOF

```

Where:

- ◆ **ARN** is Amazon Resource Name (ARN) of your CodeCommit repository. For example, `arn:aws:codecommit:eu-central-1:123456789012:triggermeshtest`.
  - ◆ **BRANCH** is the branch of your CodeCommit repository that you want the trigger to watch. For example, `main`.
6. Patch the `awscodecommitsource-adapter` service account to pull images from the private registry using the `tap-registry` secret, created during the TAP installation. Note that the `awscodecommitsource-adapter` service account was created on the previous step during the creation of `AWSCodeCommitSource`.

```

kubectl patch serviceaccount -n ${WORKLOAD_NAMESPACE} awscodecommitsource-adapter -p '{"imagePullSecrets": [{"name": "tap-registry"}]}'

```

Note: It may be necessary to delete the current `awscodecommitsource-source` **Pod** so a new pod is created with the new `imagePullSecrets`.

7. Create an event by pushing a commit to your CodeCommit repository.
8. Watch the consumer logs to see that the event appears after a minute:

```

kubectl logs -l serving.knative.dev/service=consumer -c user-container -n ${WORKLOAD_NAMESPACE} --since=10m --tail=50

```

# Upgrading Cloud Native Runtimes

This topic tells you how to upgrade Cloud Native Runtimes for Tanzu to the latest version.

New versions of Cloud Native Runtimes are available from the Tanzu Application Platform package repository, and can be upgraded to as part of [upgrading Tanzu Application Platform as a whole](#).

## Prerequisites

The following prerequisites are required to upgrade Cloud Native Runtimes:

- An updated Tanzu Application Platform package repository with the version of Cloud Native Runtimes you wish to upgrade to. For more information, see the [documentation on adding a new package repository](#).

## Upgrade from a previous version to Cloud Native Runtimes v2.0.1

Cloud Native Runtimes v2.0.1 no longer packages Knative Eventing and related resources. All the Knative Eventing resources are now part of a new package called Eventing for Tanzu. See [official documentation](#) for details.

If you previously installed Cloud Native Runtimes and don't use any Eventing components, you can go ahead and upgrade to the newest version by following the [Upgrade Cloud Native Runtimes](#) instructions. As part of the upgrade process, all the Eventing components will be deleted from your cluster. If you wish to install the Eventing package later, you can do so by following the package's official installation instructions.

If you have previously installed Cloud Native Runtimes for Tanzu and wish to keep any Eventing components, you can follow the instructions below.

**Warning:** *There can be some downtime in your Eventing workloads during the upgrade, since related deployments will be deleted and recreated during the process.*

1. First, install the new Eventing package by running the following command.

```
tanzu package install eventing -p eventing.tanzu.vmware.com -v EVENTING-VERSION -n tap
--install --values-file eventing-values.yaml
```

Where `EVENTING-VERSION` is the latest version of Eventing for Tanzu available as part of the new Tanzu Application Platform package repository, starting from version 1.3.0.

As part of Eventing package installation, the new package will take ownership of the existing Eventing resources in your cluster.

1. Then, you can upgrade Cloud Native Runtimes to the newest version as usual:

```
tanzu package installed update cloud-native-runtimes -p cnrs.tanzu.vmware.com -v CNR-VERSION --values-file cnr-values.yaml -n tap-install
```

Where `CNR-VERSION` is the latest version of Cloud Native Runtimes available as part of the new Tanzu Application Platform package repository.

## Known issue in Eventing 2.0 during the upgrade of a profile-based installation

During the upgrade from TAP v1.2 to TAP v1.3 of a profile-based installation, the Eventing package may fail to reconcile, causing TAP also to not reconcile successfully. Some pods in the `knative-eventing` namespace could be in an error state (e.g. `CrashLoopBackOff`) and may not recover from it. The pods may fail with the errors described below.

The logs from `rabbitmq-broker-webhook` pods will show the error:

```
Error reading/parsing logging configuration: timed out waiting for the condition: Unauthorized
```

The command `kubectl get app eventing -n tap-install -oyaml` may show the following error in its output:

```
create configmap/default-ch-webhook (v1) namespace: knative-eventing
^ Retryable error: Creating resource configmap/default-ch-webhook (v1) namespace: knative-eventing: API server says: Internal error occurred: failed calling webhook "config.webhook.rabbitmq.sources.knative.dev": failed to call webhook: Post "https://rabbitmq-webhook.knative-sources.svc:443/config-validation?timeout=2s": no endpoints available for service "rabbitmq-webhook" (reason: InternalError)
```

### Workaround

The provided workaround is to delete the Eventing package using the command below, and wait for it to be re-created again.

```
tanzu package installed delete eventing.tanzu.vmware.com -n tap-install
```

## Upgrade Cloud Native Runtimes

To upgrade the Cloud Native Runtimes PackageInstall specifically, run:

```
tanzu package installed update cloud-native-runtimes -p cnrs.tanzu.vmware.com -v CNR-VERSION --values-file cnr-values.yaml -n tap-install
```

Where `CNR-VERSION` is the latest version of Cloud Native Runtimes available as part of the new Tanzu Application Platform package repository.

# Uninstalling Cloud Native Runtimes

This topic tells you how to uninstall Cloud Native Runtimes.

## Overview

Cloud Native Runtimes is part of the Tanzu Application Platform package repository. For information on uninstalling the entire Tanzu Application Platform package repository, see the [Tanzu Application Platform uninstall documentation](#).

## Uninstall

To uninstall Cloud Native Runtimes specifically:

1. Delete the installed package:

```
tanzu package installed delete cloud-native-runtimes --namespace tap-install
```