

Deployment to Kubernetes

This document describes how to use the Carvel packages to deploy Spring Cloud Data Flow, and it also has a [section](#) on creating local development cluster using minikube or kind.

Deployment using Carvel

Deployment of a carvel package requires the installation of tools and specific Kubernetes controllers. Then you will add the package repository to the cluster and install the application.

For local minikube or kind cluster you can use: [Configure Kubernetes for local development or testing](#), and follow the instructions until the section *Deploy Spring Cloud Data Flow*

Required Tools

- `kubectl` - Kubernetes CLI (Install with `brew install kubectl`)
- `carvel` - Packaging and Deployment tools

Carvel CLI can be installed using:

```
wget -O- https://carvel.dev/install.sh | bash
# or with curl...
curl -L https://carvel.dev/install.sh | bash
```

Alternative following the instructions at the bottom of the home page at carvel.dev

The following tools are use by the scripts.

- `jq` - lightweight JSON parser
- `yq` - lightweight YAML parser
- `wget` - Invoke http requests.
- `dirname` provides the directory part of a filename.
- `readlink` provides absolute path of a relative link.

NOTE

Some of these utilities are not installed in macOS or *nix by default but will be available from MacPorts or HomeBrew.

Scripts

These scripts assume you are connected to a Kubernetes cluster and `kubectl` is available.

Name	Arguments	Descriptions
<code>start-deploy.sh</code>	<broker> [scdf-type] [namespace] [release snapshot]	Configures environmental variables needs for the rest of the scripts. <code>BROKER</code> , <code>NS</code> and <code>SCDF_TYPE</code> are set. The default <code>NS</code> is <code>scdf</code> . The namespace will be created if it doesn't exist by <code>setup-scdf-repo.sh</code> . The default <code>SCDF_TYPE</code> is <code>oss</code> . <code>release snapshot</code> and <code>scdf-type</code> will determine the value of <code>PACKAGE_VERSION</code> set.
<code>prepare-cluster.sh</code>	N/A	Installs cert-manager, secretgen-controller and kapp-controller
<code>carvel-use-template.sh</code>	[scdf-type] (oss, pro)	Creates <code>scdf-values.yml</code> in current directory based on <code>scdf-pro-values.yml</code> or <code>scdf-oss-values.yml</code>
<code>setup-scdf-repo.sh</code>	[scdf-type] (oss, pro)	Creates the namespace and installs the relevant Carvel package and credentials. If the optional <code>scdf-type</code> is not provided the environmental variable <code>SCDF_TYPE</code> will be used.
<code>configure-prometheus-proxy.sh</code>	<host> <port> [step]	Configures Spring Boot Actuator properties for Data Flow, Skipper, Streams and Tasks. Default <code>step</code> is 10s
<code>configure-database.sh</code>	<app> <database> <url> <username/secret-name> [password/secret-username-key] [secret-password-key]	If only <code>secret-name</code> is provided then <code>secret-username-key</code> defaults to <code>username</code> and <code>secret-password-key</code> defaults to <code>password</code> . The following 3 combinations are allowed after the <code>url</code> : <ul style="list-style-type: none">• <secret-name>• <secret-name> <username-key> <password-key>• <username> <password>
<code>deploy-scdf.sh</code>	[app-name]	Deploys the application using the package and <code>scdf-values.yml</code> in the current directory. The default <code>app-name</code> is <code>scdf-\${SCDF_TYPE}</code> .
<code>update-scdf.sh</code>	[app-name]	Updated the deployed application using a modified values file. The default <code>app-name</code> is <code>scdf-\${SCDF_TYPE}</code> .
<code>export-dataflow-ip.sh</code>	N/A	Will print the URL to access dataflow. If you use <code>source ./export-dataflow-ip.sh</code> it will export <code>DATAFLOW_URL</code> to be used by <code>register-apps.sh</code>
<code>register-apps.sh</code>	<broker> [stream-application-version]	<code>broker</code> must be one of rabbit or kafka. <code>stream-application-version</code> is optional and will install the latest version. The latest version is 2021.1.2

NOTE

Take note that the registration of application in the *pro* version can take a few minutes since it retrieves all version information and metadata upfront.

Preparation

You will need to prepare a values file named `scdf-values.yml`. The following steps will provide help.

Prepare Configuration parameters

Executing the following script will configure the environmental variables needed.

```
source ./carvel/start-deploy.sh <broker> <namespace> [scdf-type] [release|snapshot]
```

Where:

- **broker** is one of `rabbitmq` or `kafka`
- **namespace** A valid Kubernetes namespace other than `default`
- **scdf-type** One of `oss` or `pro`. `oss` is the default.
- **release|snapshot** and **scdf-type** will determine the value of `PACKAGE_VERSION`.

The environmental variables can also be configured manually to override the values.

Name	Description	Default
<code>PACKAGE_VERSION</code>	Version of Carvel package.	Release version
<code>DATAFLOW_VERSION</code>	Version of Spring Cloud Data Flow	2.10.3
<code>SKIPPER_VERSION</code>	Version of Spring Cloud Skipper	2.9.3
<code>REGISTRY</code>	Url and repository of package registry. Format <code><private-registry-host/repo-name></code> . This will be used to prefix the carvel repo and package.	<code>docker.io/springcloud</code>
<code>BROKER</code>	One of <code>kafka</code> or <code>rabbitmq</code>	<code>rabbitmq</code>
<code>DATABASE</code>	One of <code>mariadb</code> or <code>postgresql</code> . The default is <code>postgresql</code> . This will only apply when you <code>deploy-local-database.sh</code>	<code>postgresql</code>
<code>NS</code>	A Kubernetes namespace other than <code>default</code> .	<code>scdf</code>
<code>SCDF_TYPE</code>	One of <code>oss</code> or <code>pro</code> .	<code>oss</code>

Prepare Configuration file

Create a file name `scdf-values.yml` by executing:

```
./carvel/carvel-use-template.sh
```

Edit the file as needed to configure the deployment. The `deploy-local-` scripts will

Uses scdf-type previously selected.

Prepare cluster and add repository

Login to docker and optionally registry.pivotal.io for Spring Cloud Data Flow Pro.

```
# When deploying SCDF Pro.
export TANZU_DOCKER_USERNAME=<tanzu-net-username>
export TANZU_DOCKER_PASSWORD=<tanzu-net-password>
docker login --username $TANZU_DOCKER_USERNAME --password $TANZU_DOCKER_PASSWORD
registry.pivotal.io

# Always required to ensure you don't experience rate limiting with Docker HUB
export DOCKER_HUB_USERNAME=<docker-hub-username>
export DOCKER_HUB_PASSWORD=<docker-hub-password>
docker login --username $DOCKER_HUB_USERNAME --password $DOCKER_HUB_PASSWORD
index.docker.io
```

Install carvel kapp-controller, secretgen-controller and certmanager

```
./carvel/prepare-cluster.sh
```

Load scdf repo package for the *scdf-type*

```
./carvel/setup-scdf-repo.sh
```

Install supporting services

In a production environment you should be using supported database and broker services or operators along with shared observability tools.

For local development or demonstration the following can be used to install database, broker and prometheus.

Deploy local database.

```
./carvel/deploy-local-database.sh <database> ①
```

① `database` must be one of `postgresql` or `mariadb`. Default is `postgresql` or configure in `DATABASE` using `start-deploy.sh`.

NOTE | This script updates `scdf-values.yml` with the correct secret name.

Deploy local message broker.

```
./carvel/deploy-local-broker.sh
```

Deploy local Prometheus and proxy.

```
./carvel/deploy-local-prometheus.sh
```

This script also configures the Grafana endpoint in `scdf-values.yml`

Configure Prometheus proxy

In the case where an existing prometheus and prometheus proxy is deployed the proxy can be configured using:

```
./carvel/configure-prometheus-proxy.sh <host> <port> [step]
```

Deploy Spring Cloud Data Flow

```
./carvel/deploy-scdf.sh  
source ./carvel/export-dataflow-ip.sh  
# expected output: Dataflow URL: <url-to-access-dataflow>  
./carvel/register-apps.sh
```

Update deployed application.

You can modify the values file used during installation and then update the deployment using `./carvel/update-scdf.sh`

Configure Kubernetes for local development or testing

Prerequisites

You will need to install kubectl and then kind or minikube for a local cluster.

All the examples assume you have cloned the `spring-cloud-dataflow` repository and are executing the scripts from `deploy/k8s`.

On macOS you may need to install `realpath` from `Macports` or `brew install realpath`

NOTE

The scripts require a shell like `bash` or `zsh` and should work on Linux, WSL 2 or macOS.

Steps

- Choose Kubernetes provider. Kind, Minikube or remote GKE or TMC.
- Decide the namespace to use for deployment if not `default`.
- Configure Kubernetes and loadbalancer.
- Choose Broker with `export BROKER=kafka|rabbitmq`
- Build or Pull container images for Skipper and Data Flow Server.
- Deploy and Launch Spring Cloud Data Flow.
- Export Data Flow Server address to env.

Kubernetes Provider

How do I choose between minikube and kind? kind will generally provide quicker setup and teardown time than Minikube. There is little to choose in terms of performance between the 2 apart from being able to configure limits on CPUs and memory when deploying minikube. So in the case where you have memory constraints or need to enforce memory limitations Minikube will be a better option.

Kubectl

You will need to [install](#) kubectl in order to configure the Kubernetes cluster

Kind

Kind is Kubernetes in docker and ideal for local development.

- [Installation](#)
- [LoadBalancer](#)

The LoadBalancer will be installed by the `configure-k8s.sh` script by will require an update to a yaml file to provide the address range available to the LoadBalancer.

Minikube

Minikube uses one of a selection of drivers to provide a virtualization environment.

- [Installation](#)
- [LoadBalancer](#)

NOTE	Delete existing Minikube installation if you have any. <code>minikube delete</code>
-------------	---

Remote TMC Cluster

[Tanzu Mission Control](#)

Building and loading containers.

For local development you need control of the containers used in the local environment.

In order to ensure to manage the specific versions of data flow and skipper containers you can set `SKIPPER_VERSION` and `DATAFLOW_VERSION` environmental variable and then invoke `./images/pull-dataflow.sh` and `./images/pull-skipper.sh` or if you want to use a locally built application you can invoke `./images/build-skipper-image.sh` and `./images/build-dataflow.sh`

Configure k8s environment

You can invoke one of the following scripts to choose the type of installation you are targeting:

```
./k8s/use-kind.sh [<namespace>] [<database>] [<broker>]
./k8s/use-mk-docker.sh [<namespace>] [<database>] [<broker>]
./k8s/use-mk-kvm2.sh [<namespace>] [<database>] [<broker>]
./k8s/use-mk.sh <driver> [<namespace>] [<database>] [<broker>] ①
./k8s/use-tmc.sh <cluster-name> [<namespace>] [<database>] [<broker>]
./k8s/use-gke.sh <cluster-name> [<namespace>] [<database>] [<broker>]
```

① <driver> must be one of `kvm2`, `docker`, `vmware`, `virtualbox`, `vmwarefusion` or `hyperkit`. `docker` is the recommended option for local development.

NOTE

<namespace> will be `default` if not provided. The default <database> is `postgresql` and the default <broker> is `kafka`.

Since these scripts export environmental variable they need to be executes as in the following example:

```
source ./k8s/use-mk-docker.sh test-ns postgresql rabbitmq
```

TMC or GKE Cluster in Cloud

The cluster must exist before use and you should use the relevant cli to login before executing `source ./k8s/use-gke.sh`

Create Local Cluster.

The following script will create the local cluster.

```
# Optionally add to control cpu and memory allocation.
export MK_ARGS="--cpus=8 --memory=12g"
./k8s/configure-k8s.sh
```

- For **kind** follow instruction to update `./k8s/yaml/metallb-configmap.yaml` and then apply using `kubectl apply -f ./k8s/yaml/metallb-configmap.yaml`
- For **minikube** launch a new shell and execute `minikube tunnel`

Deploy Spring Cloud Data Flow.

The **use-*** scripts will configure the values of BROKER and DATABASE.

Configure Broker

```
export BROKER=<broker> ①
```

① <broker> one of **kafka** or **rabbitmq**

Configure Database

```
export DATABASE=<database> ①
```

① <database> one of **mariadb** or **postgresql**

Docker credentials need to be configured for Kubernetes to pull the various container images.

For Docker Hub you can create a personal free account and use a personal access token as your password.

Test your docker login using **./k8s/docker-login.sh**

```
export DOCKER_SERVER=https://docker.io
export DOCKER_USER=<docker-userid>
export DOCKER_PASSWORD=<docker-password>
export DOCKER_EMAIL=<email-of-docker-use>
```

Before you can install SCDF you will need to pull the following images to ensure they are present for uploading to the k8s cluster.

Use:

```
./images/pull-app-images.sh
./images/pull-dataflow.sh
./images/pull-skipper.sh
./images/pull-composed-task-runner.sh
```

```
./k8s/install-scdf.sh
source ./k8s/export-dataflow-ip.sh
```

Delete the deployment from the cluster.

```
./k8s/delete-scdf.sh
```


Delete the cluster

This script will also delete the TMC cluster if you have configured one.

```
./k8s/destroy-k8s.sh
```

Utilities

The following list of utilities may prove useful.

Name	Description
k9s	k9s is a text based monitor to explore the Kubernetes cluster.
kail	Extra and tail the logs of various pods based on various naming criteria.

kail

- Using kail to log activity related to a specific stream.

```
kail --label=spring-group-id=<stream-name>
```

- Using kail to log all pods in specific namespace.

```
kail --ns=<namespace>
```

Scripts

Some of the scripts apply to local containers as well and can be found in [src/local](#), the Kubernetes specific scripts are in [deploy/k8s](#)

Script	Description
./images/build-app-images.sh	Build all images of Restaurant Sample Stream Apps
./images/pull-app-images.sh	Pull all images of Restaurant Sample Stream Apps from Docker Hub
./images/pull-dataflow.sh	Pull dataflow from DockerHub based on DATAFLOW_VERSION .
./images/pull-scdf-pro.sh	Pull Dataflow Pro from Tanzu Network based on SCDF_PRO_VERSION .
./images/pull-skipper.sh	Pull Skipper from DockerHub base on the SKIPPER_VERSION .
./images/build-dataflow-image.sh	Build a docker image from the local repo of Dataflow
./images/build-scdf-pro-image.sh	Build a docker image from the local repo of Dataflow Pro. Set USE_PRO=true in environment to use Dataflow Pro

Script	Description
<code>./images/build-skipper-image.sh</code>	Build a docker image from the local repo of Skipper.
<code>./k8s/configure-k8s.sh</code>	Configure the Kubernetes environment based on your configuration of K8S_DRIVER.
<code>./k8s/delete-scdf.sh</code>	Delete all Kubernetes resources create by the deployment.
<code>./k8s/destroy-k8s.sh</code>	Delete cluster, kind or minikube.
<code>./k8s/export-dataflow-ip.sh</code>	Export the url of the data flow server to <code>DATAFLOW_IP</code>
<code>./k8s/export-http-url.sh</code>	Export the url of an http source of a specific flow by name to <code>HTTP_APP_URL</code>
<code>./k8s/install-scdf.sh</code>	Configure and deploy all the containers for Spring Cloud Dataflow
<code>./k8s/load-images.sh</code>	Load all container images required by tests into kind or minikube to ensure you have control over what is used.
<code>./k8s/load-image.sh</code>	Load a specific container image into local kind or minikube.
<code>src/local/local-k8s-acceptance-tests.sh</code>	Execute acceptance tests against cluster where <code>DATAFLOW_IP</code> is pointing.
<code>./k8s/register-apps.sh</code>	Register the Task and Stream apps used by the unit tests.

IMPORTANT

Please report any errors with the scripts along with detail information about the relevant environment.