



Red Hat OpenShift Data Science self-managed 1-latest

Installing OpenShift Data Science self-managed

Install Red Hat OpenShift Data Science as an Operator to your OpenShift Container Platform cluster

Red Hat OpenShift Data Science self-managed 1-latest Installing OpenShift Data Science self-managed

Install Red Hat OpenShift Data Science as an Operator to your OpenShift Container Platform cluster

Legal Notice

Copyright © 2023 Red Hat, Inc.

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

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

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

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

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

Linux[®] is the registered trademark of Linus Torvalds in the United States and other countries.

Java[®] is a registered trademark of Oracle and/or its affiliates.

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

MySQL[®] is a registered trademark of MySQL AB in the United States, the European Union and other countries.

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

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

All other trademarks are the property of their respective owners.

Abstract

Install Red Hat OpenShift Data Science as an Operator to your OpenShift Container Platform cluster.

Table of Contents

PREFACE	3
CHAPTER 1. ARCHITECTURE OF OPENSIFT DATA SCIENCE SELF-MANAGED	4
CHAPTER 2. OVERVIEW OF INSTALLING AND DEPLOYING OPENSIFT DATA SCIENCE	6
CHAPTER 3. REQUIREMENTS FOR OPENSIFT DATA SCIENCE SELF-MANAGED	7
CHAPTER 4. ADDING ADMINISTRATIVE USERS FOR OPENSIFT CONTAINER PLATFORM	9
CHAPTER 5. INSTALLING OPENSIFT DATA SCIENCE ON OPENSIFT CONTAINER PLATFORM	10
CHAPTER 6. INSTALLING AND MANAGING VERSION 2 OF THE RED HAT OPENSIFT DATA SCIENCE OPERATOR	12
6.1. INSTALLING VERSION 2 OF THE RED HAT OPENSIFT DATA SCIENCE OPERATOR BY USING THE CLI	12
6.2. INSTALLING VERSION 2 OF THE RED HAT OPENSIFT DATA SCIENCE OPERATOR BY USING THE WEB CONSOLE	14
6.3. UNINSTALLING VERSION 2 OF THE RED HAT OPENSIFT DATA SCIENCE OPERATOR	15
CHAPTER 7. INSTALLING RED HAT OPENSIFT DATA SCIENCE COMPONENTS	18
7.1. INSTALLING RED HAT OPENSIFT DATA SCIENCE COMPONENTS BY USING THE CLI	18
7.2. INSTALLING RED HAT OPENSIFT DATA SCIENCE COMPONENTS BY USING THE WEB CONSOLE	19
CHAPTER 8. ACCESSING THE OPENSIFT DATA SCIENCE DASHBOARD	22
CHAPTER 9. ENABLING GPU SUPPORT IN OPENSIFT DATA SCIENCE	23

PREFACE

Use OperatorHub to install Red Hat OpenShift Data Science as an Operator to your OpenShift Container Platform cluster. Red Hat recommends that you install only one instance of OpenShift Data Science on your cluster. Installing the Red Hat OpenShift Data Science Operator on the same cluster as the OpenShift Data Science Add-on is not recommended or supported.

CHAPTER 1. ARCHITECTURE OF OPENSIFT DATA SCIENCE SELF-MANAGED

Red Hat OpenShift Data Science self-managed is an Operator that is available on a self-managed environment, such as Red Hat OpenShift Container Platform.

OpenShift Data Science integrates the following components and services:

- At the service layer:

OpenShift Data Science dashboard

A customer-facing dashboard that shows available and installed applications for the OpenShift Data Science environment as well as learning resources such as tutorials, quick starts, and documentation. Administrative users can access functionality to manage users, clusters, notebook images, and model-serving runtimes. Data scientists can use the dashboard to create projects to organize their data science work.

Model serving

Data scientists can deploy trained machine-learning models to serve intelligent applications in production. After deployment, applications can send requests to the model using its deployed API endpoint.

Data science pipelines

Data scientists can build portable machine learning (ML) workflows with data science pipelines, using Docker containers. This enables your data scientists to automate workflows as they develop their data science models.

Jupyter (self-managed)

A self-managed application that allows data scientists to configure their own notebook server environment and develop machine learning models in JupyterLab.

- At the management layer:

The Red Hat OpenShift Data Science Operator

A meta-operator that deploys and maintains all components and sub-operators that are part of OpenShift Data Science.

Monitoring services

Prometheus gathers metrics from OpenShift Data Science for monitoring purposes.

When you install the OpenShift Data Science Operator in the OpenShift Container Platform cluster, the following new projects are created:

- The **redhat-ods-operator** project contains the OpenShift Data Science operator.
- The **redhat-ods-applications** project installs the dashboard and other required components of OpenShift Data Science.
- The **redhat-ods-monitoring** project contains services for monitoring.
- The **rhods-notebooks** project is where notebook environments are deployed by default.

You or your data scientists must create additional projects for the applications that will use your machine learning models.

Do not install independent software vendor (ISV) applications in namespaces associated with OpenShift Data Science.

Additional resources

- [Installing OpenShift Data Science self-managed](#)

CHAPTER 2. OVERVIEW OF INSTALLING AND DEPLOYING OPENSIFT DATA SCIENCE

Red Hat OpenShift Data Science is a platform for data scientists and developers of artificial intelligence (AI) applications. It provides a fully supported environment that lets you rapidly develop, train, test, and deploy machine learning models on-premises and/or in the public cloud.

OpenShift Data Science is provided as a managed cloud service add-on for Red Hat OpenShift or as self-managed software that you can install on-premise or in the public cloud on OpenShift. For information on installing OpenShift Data Science as a managed cloud service add-on, see [Installing OpenShift Data Science](#).

Installing OpenShift Data Science involves the following high-level tasks:

1. Confirm that your OpenShift Container Platform cluster meets all requirements.
2. Configure an identity provider for OpenShift Container Platform.
3. Add administrative users for OpenShift Container Platform.
4. Install the OpenShift Data Science Operator.
5. Configure user and administrator groups to provide user access to OpenShift Data Science.
6. Access the OpenShift Data Science dashboard.
7. Optionally, enable graphics processing units (GPUs) in OpenShift Data Science to ensure that your data scientists can use compute-heavy workloads in their models.

CHAPTER 3. REQUIREMENTS FOR OPENSIFT DATA SCIENCE SELF-MANAGED

Your environment must meet certain requirements to receive support for Red Hat OpenShift Data Science.

Installation requirements

You must meet the following requirements before you are able to install OpenShift Data Science on your Red Hat OpenShift Container Platform cluster.

- **Product subscriptions**
 - A subscription for Red Hat OpenShift Data Science self-managed
Contact your Red Hat account manager to purchase new subscriptions. If you do not yet have an account manager, complete the form at <https://www.redhat.com/en/contact> to request one.
- **An OpenShift Container Platform cluster 4.11 or greater**
 - Use an existing cluster or create a new cluster by following the OpenShift Container Platform documentation: [OpenShift Container Platform installation overview](#).
Your cluster must have at least 2 worker nodes with at least 8 CPUs and 32 GiB RAM available for OpenShift Data Science to use when you install the Operator. The installation process fails to start and an error is displayed if this requirement is not met. To ensure that OpenShift Data Science is usable, additional cluster resources are required beyond the minimum requirements.
 - A default storage class that can be dynamically provisioned must be configured.
Confirm that a default storage class is configured by running the **oc get storageclass** command. If no storage classes are noted with **(default)** beside the name, follow the OpenShift Container Platform documentation to configure a default storage class: [Changing the default storage class](#). For more information about dynamic provisioning, see [Dynamic provisioning](#).
 - Open Data Hub must not be installed on the cluster.
For more information about managing the machines that make up an OpenShift cluster, see [Overview of machine management](#).
- **An identity provider configured for OpenShift Container Platform**
Access to the cluster as a user with the **cluster-admin** role; the **kubeadmin** user is not allowed.

Red Hat OpenShift Data Science supports the same authentication systems as Red Hat OpenShift Container Platform. See [Understanding identity provider configuration](#) for more information on configuring identity providers.
- **Internet access**
Along with Internet access, the following domains must be accessible during the installation of OpenShift Data Science self-managed:
 - cdn.redhat.com
 - subscription.rhn.redhat.com
 - registry.access.redhat.com

- registry.redhat.io
- quay.io

For CUDA-based images, the following domains must be accessible:

- ngc.download.nvidia.cn
- developer.download.nvidia.com

- **OpenShift Pipelines operator installation**

- The Red Hat OpenShift Pipelines operator enables support for installation of pipelines in a self-managed environment.
Before you use data science pipelines in OpenShift Data Science, you must install the Red Hat OpenShift Pipelines Operator. For more information, see [Installing OpenShift Pipelines](#). If your deployment is in a disconnected self-managed environment, see [Red Hat OpenShift Pipelines Operator in a restricted environment](#).
- Before you can execute a pipeline in a disconnected environment, you must mirror any images used by your pipelines to a private registry.
- You can store your pipeline artifacts in an Amazon Web Services (AWS) Simple Storage Service (S3) bucket to ensure that you do not consume local storage. To do this, you must first configure write access to your S3 bucket on your AWS account.
If you do not have access to Amazon S3 storage, you must configure your own storage solution for use with pipelines.

CHAPTER 4. ADDING ADMINISTRATIVE USERS FOR OPENSIFT CONTAINER PLATFORM

Before you can install and configure OpenShift Data Science for your data scientist users, you must define administrative users. Only users with the **cluster-admin** role can install and configure OpenShift Data Science.

For more information about creating a cluster admin user, see [Creating a cluster admin](#).

CHAPTER 5. INSTALLING OPENSIFT DATA SCIENCE ON OPENSIFT CONTAINER PLATFORM

You can install the Red Hat OpenShift Data Science Operator to your Red Hat OpenShift Container Platform cluster using the OpenShift Container Platform web console.



IMPORTANT

Upgrading directly from the Red Hat OpenShift Data Science self-managed Beta version to the Generally Available (GA) release is not supported. To install the OpenShift Data Science self-managed GA release, you must remove the Beta version first and then proceed with the following procedure. See [Uninstalling Red Hat OpenShift Data Science self-managed Beta version prior to installing a General Availability \(GA\) release](#) for more information.



NOTE

If your OpenShift cluster uses a proxy to access the Internet, you can configure the proxy settings for the Red Hat OpenShift Data Science Operator. See [Overriding proxy settings of an Operator](#) for more information.

Prerequisites

- Purchase entitlements for OpenShift Data Science self-managed.
- A running OpenShift Container Platform cluster, version 4.11 or greater, configured with a default storage class that can be dynamically provisioned.
- Open Data Hub must not be installed on the same OpenShift cluster.
- Access to the OpenShift Container Platform cluster as a user with the **cluster-admin** role.

Procedure

1. Log in to the OpenShift Container Platform cluster web console.
2. Click **Operators** → **OperatorHub**.
The **OperatorHub** page opens.
3. Locate the Red Hat OpenShift Data Science Operator.
 - a. Scroll through available Operators or type Red Hat OpenShift Data Science into the **Filter by keyword** box to find the Red Hat OpenShift Data Science Operator.
4. Select the Operator to display additional information.
5. Read the information about the Operator and click **Install**.
The **Install Operator** page opens.
6. For **Installation mode**, select the **All namespaces on the cluster (default)** to install the Operator in the default **redhat-ods-operator** namespace and make it available to all namespaces in the cluster.
Note: The option to select **A specific namespace on the cluster** is not available.
7. Under **Update approval**, select either **Automatic** or **Manual**.

8. Click **Install**.

Verification

- In the OpenShift Container Platform web console, click **Operators** → **Installed Operators** and confirm that the Red Hat OpenShift Data Science Operator shows one of the following statuses:
 - **Installing** - installation is in progress; wait for this to change to **Succeeded**. This takes around 10 minutes.
 - **Succeeded** - installation is successful.
- In OpenShift Container Platform, click **Home** → **Projects** and confirm that the following project namespaces are visible and listed as **Active**:
 - **redhat-ods-applications**
 - **redhat-ods-monitoring**
 - **redhat-ods-operator**
 - **rhods-notebooks**

Additional resources

- [Adding users for OpenShift Data Science](#)
- [Adding Operators to a cluster](#)

CHAPTER 6. INSTALLING AND MANAGING VERSION 2 OF THE RED HAT OPENSIFT DATA SCIENCE OPERATOR

This section shows to install version 2 (that is, version 2.2 or 2.1) of the Red Hat OpenShift Data Science Operator on your OpenShift Container Platform cluster using the command-line interface (CLI) and the OpenShift web console. The section also shows how to uninstall the Operator.



IMPORTANT

Version 2 of the Red Hat OpenShift Data Science Operator is a Limited Availability feature. Limited Availability means that you can install and receive support for the feature only with specific approval from Red Hat. Without such approval, the feature is unsupported.

6.1. INSTALLING VERSION 2 OF THE RED HAT OPENSIFT DATA SCIENCE OPERATOR BY USING THE CLI

The following procedure shows how to use the OpenShift command-line interface (CLI) to install version 2 (that is, version 2.2 or 2.1) of the Red Hat OpenShift Data Science Operator on your OpenShift Container Platform cluster. The steps describe how to perform a basic installation of the Operator without installing any OpenShift Data Science components.

Prerequisites

- You are an IBM watsonx user or Red Hat has granted you installation entitlements for this version of the Operator.
- You have a running OpenShift Container Platform cluster, version 4.11 or greater, configured with a default storage class that can be dynamically provisioned.
- You have cluster administrator privileges for your OpenShift Container Platform cluster.
- You have [downloaded and installed](#) the OpenShift command-line interface (CLI).

Procedure

1. Open a new terminal window.
2. In the OpenShift command-line interface (CLI), log in to your OpenShift Container Platform cluster as a cluster administrator, as shown in the following example:

```
$ oc login <openshift_cluster_url> -u system:admin
```

3. Create a namespace for installation of the Operator by performing the following actions:
 - a. Create a namespace YAML file, for example, **rhods-operator-namespace.yaml**.

```
apiVersion: v1
kind: Namespace
metadata:
  name: redhat-ods-operator 1
```

- 1** **redhat-ods-operator** is the recommended namespace for the Operator.

- b. Create the namespace in your OpenShift Container Platform cluster.

```
$ oc create -f rhods-operator-namespace.yaml
```

You see output that resembles the following:

```
namespace/redhat-ods-operator created
```

4. Create an operator group for installation of the Operator by performing the following actions:

- a. Create an **OperatorGroup** object custom resource (CR) file, for example, **rhods-operator-group.yaml**.

```
apiVersion: operators.coreos.com/v1
kind: OperatorGroup
metadata:
  name: rhods-operator
  namespace: redhat-ods-operator 1
```

1 You must specify the same namespace that you created earlier in this procedure.

- b. Create the **OperatorGroup** object in your OpenShift Container Platform cluster.

```
$ oc create -f rhods-operator-group.yaml
```

You see output that resembles the following:

```
operatorgroup.operators.coreos.com/rhods-operator created
```

5. Create a subscription for installation of the Operator by performing the following actions:

- a. Create a **Subscription** object CR file, for example, **rhods-operator-subscription.yaml**.

```
apiVersion: operators.coreos.com/v1alpha1
kind: Subscription
metadata:
  name: rhods-operator
  namespace: redhat-ods-operator 1
spec:
  name: rhods-operator
  channel: alpha 2
  source: redhat-operators
  sourceNamespace: openshift-marketplace
```

1 You must specify the same namespace that you created earlier in this procedure.

2 To install version 2.2 of the Operator, specify a channel called **alpha**. To install version 2.1, specify a channel called **embedded**.

- b. As described in the preceding step, ensure that the channel name is correct for the version of the Operator that you want to install.

- c. Create the **Subscription** object in your OpenShift Container Platform cluster to install the Operator.

```
$ oc create -f rhods-operator-subscription.yaml
```

You see output that resembles the following:

```
subscription.operators.coreos.com/rhods-operator created
```

Verification

- In the OpenShift Container Platform web console, click **Operators** → **Installed Operators** and confirm that the Red Hat OpenShift Data Science Operator shows one of the following statuses:
 - **Installing** - installation is in progress; wait for this to change to **Succeeded**. This might take several minutes.
 - **Succeeded** - installation is successful.
- In the web console, click **Home** → **Projects** and confirm that the following project namespaces are visible and listed as **Active**:
 - **redhat-ods-applications**
 - **redhat-ods-monitoring**
 - **redhat-ods-operator**

6.2. INSTALLING VERSION 2 OF THE RED HAT OPENSIFT DATA SCIENCE OPERATOR BY USING THE WEB CONSOLE

The following procedure shows how to use the OpenShift Container Platform web console to install version 2 (that is, version 2.2 or 2.1) of the Red Hat OpenShift Data Science Operator on your cluster. The steps describe how to perform a basic installation of the Operator without installing any OpenShift Data Science components.

Prerequisites

- You are an IBM watsonx user or Red Hat has granted you installation entitlements for this version of the Operator.
- You have a running OpenShift Container Platform cluster, version 4.11 or greater, configured with a default storage class that can be dynamically provisioned.
- You have cluster administrator privileges for your OpenShift Container Platform cluster.

Procedure

1. Log in to the OpenShift Container Platform web console as a cluster administrator.
2. In the web console, click **Operators** → **OperatorHub**.
The **OperatorHub** page opens.
3. Locate the Red Hat OpenShift Data Science Operator.

- a. Scroll through available Operators or type **Red Hat OpenShift Data Science** into the **Filter by keyword** box to find the Red Hat OpenShift Data Science Operator.
4. Select the Operator to display additional information.
5. Read the information about the Operator and click **Install**.
The **Install Operator** page opens.
6. For **Update channel**, select a value as follows:
 - a. To install version 2.2 of the Operator, select **alpha**.
 - b. To install version 2.1 of the Operator, select **embedded**.
7. For **Installation mode**, observe that the only available value is **All namespaces on the cluster (default)**. This installation mode makes the Operator available to all namespaces in the cluster.
8. For **Installed Namespace**, select **redhat-ods-operator (Operator recommended)**.
9. Under **Update approval**, select either **Automatic** or **Manual**.
10. Click **Install**.
An installation pane opens. When the installation finishes, a check mark appears beside the Operator name in the installation pane.

Verification

- In the OpenShift Container Platform web console, click **Operators → Installed Operators** and confirm that the Red Hat OpenShift Data Science Operator shows one of the following statuses:
 - **Installing** - installation is in progress; wait for this to change to **Succeeded**. This might take several minutes.
 - **Succeeded** - installation is successful.
- In the web console, click **Home → Projects** and confirm that the following project namespaces are visible and listed as **Active**:
 - **redhat-ods-applications**
 - **redhat-ods-monitoring**
 - **redhat-ods-operator**

6.3. UNINSTALLING VERSION 2 OF THE RED HAT OPENSIFT DATA SCIENCE OPERATOR

The following procedure shows how to use the OpenShift command-line interface (CLI) to uninstall version 2 (that is, version 2.2 or 2.1) of the Red Hat OpenShift Data Science Operator and any OpenShift Data Science components installed and managed by the Operator. Using the CLI is the recommended way to perform this uninstallation.

Prerequisites

- You have cluster administrator privileges for your OpenShift Container Platform cluster.

- You have [downloaded and installed](#) the OpenShift command-line interface (CLI).

Procedure

- Open a new terminal window.
- In the OpenShift command-line interface (CLI), log in to your OpenShift Container Platform cluster as a cluster administrator, as shown in the following example:

```
$ oc login <openshift_cluster_url> -u system:admin
```

- Optional: If you created a **DataScienceCluster** object to install OpenShift Data Science components, delete the **DataScienceCluster** object.

```
$ oc delete datasciencecluster $(oc get datasciencecluster --no-headers | awk '{print $1}')
```

You see output that resembles the following:

```
datasciencecluster.datasciencecluster.opendatahub.io "default" deleted
```



NOTE

Deleting the **DataScienceCluster** object also deletes the pods for any OpenShift Data Science components that you installed. This removes the OpenShift Data Science components.

- Delete the **DSCInitialization** object that the Operator created during installation.

```
$ oc delete dscinitialization $(oc get dscinitialization --no-headers | awk '{print $1}')
```

You see the following output:

```
dscinitialization.dscinitialization.opendatahub.io "default" deleted
```

- Delete the **Subscription** object that you created to install the Operator.

```
$ oc delete subscription <subscription_name> -n <namespace_name>
```

In the following example, the command shown deletes the **rhods-operator** subscription from the **redhat-ods-operator** namespace:

```
$ oc delete subscription rhods-operator -n redhat-ods-operator
```

You see output that resembles the following:

```
subscription.operators.coreos.com "rhods-operator" deleted
```

- Navigate to the directory that contains the YAML file you used to create a namespace to install the Operator.
- Delete the namespace that you created to install the Operator.

```
$ oc delete -f <namespace_file_name>.yaml
```

You see output that resembles the following:

```
namespace "redhat-ods-operator" deleted
```



NOTE

Deleting the namespace also deletes the **OperatorGroup** object that you created during installation of the Operator.

8. Delete the namespaces that the Operator created during installation.

```
$ oc delete ns -l opendatahub.io/generated-namespace
```

For the default Operator installation, you see the following output:

```
namespace "redhat-ods-applications" deleted
namespace "redhat-ods-monitoring" deleted
```



NOTE

If you installed the workbenches component of OpenShift Data Science, the preceding command also deletes the **rhods-notebooks** namespace created during that installation.

Verification

- On the command line, get the list of current Operator subscriptions in all namespaces.

```
$ oc get subscriptions --all-namespaces
```

Confirm that the subscription for the Red Hat OpenShift Data Science Operator (for example, **rhods-operator**) is not listed.

CHAPTER 7. INSTALLING RED HAT OPENSIFT DATA SCIENCE COMPONENTS

The following procedures show how to use the command-line interface (CLI) and OpenShift Container Platform web console to install components of Red Hat OpenShift Data Science when version 2.1 of the Red Hat OpenShift Data Science Operator is already installed on your cluster.



IMPORTANT

Version 2.1 of the Red Hat OpenShift Data Science Operator is a Limited Availability feature. Limited Availability means that you can install and receive support for the feature only with specific approval from Red Hat. Without such approval, the feature is unsupported.

7.1. INSTALLING RED HAT OPENSIFT DATA SCIENCE COMPONENTS BY USING THE CLI

The following procedure shows how to use the OpenShift command-line interface (CLI) to install specific components of Red Hat OpenShift Data Science on your OpenShift Container Platform cluster when version 2 of the Red Hat OpenShift Data Science Operator is already installed on the cluster.

Prerequisites

- Version 2 of the Red Hat OpenShift Data Science Operator is [installed](#) on your OpenShift Container Platform cluster.
- You have cluster administrator privileges for your OpenShift Container Platform cluster.
- You have [downloaded and installed](#) the OpenShift command-line interface (CLI).

Procedure

1. Open a new terminal window.
2. In the OpenShift command-line interface (CLI), log in to your on your OpenShift Container Platform cluster as a cluster administrator, as shown in the following example:

```
$ oc login <openshift_cluster_url> -u system:admin
```

3. Create a **DataScienceCluster** object custom resource (CR) file, for example, **rhods-operator-dsc.yaml**.

```
apiVersion: datasciencecluster.opendatahub.io/v1
kind: DataScienceCluster
metadata:
  name: default
spec:
  components:
    codeflare:
      managementState: "Removed"
    dashboard:
      managementState: "Removed"
    datasciencepipelines:
```

```

managementState: "Removed"
kserve:
  managementState: "Removed"
modelmeshserving:
  managementState: "Removed"
ray:
  managementState: "Removed"
workbenches:
  managementState: "Removed"

```

- In the **spec.components** section of the CR, for each OpenShift Data Science component shown, set the value of the **managementState** field to either **Managed** or **Removed**. These values are defined as follows:

Managed

The Operator actively manages the component, installs it, and tries to keep it active. The Operator will upgrade the component only if it is safe to do so.

Removed

The Operator actively manages the component but does not install it. If the component is already installed, the Operator will try to remove it.

- Create the **DataScienceCluster** object in your OpenShift Container Platform cluster to install the specified OpenShift Data Science components.

```
$ oc create -f rhods-operator-dsc.yaml
```

You see output that resembles the following:

```
datasciencecluster.datasciencecluster.opendatahub.io/default created
```

Verification

- In the OpenShift Container Platform web console, click **Workloads** → **Pods**. In the **Project** list at the top of the page, select **redhat-ods-applications**. In the applications namespace, confirm that there are running pods for each of the OpenShift Data Science components that you installed.
- In the web console, click **Operators** → **Installed Operators** and then perform the following actions:
 - Click the Red Hat OpenShift Data Science Operator.
 - Click the **Data Science Cluster** tab and select the default **DataScienceCluster** object shown on the page.
 - Select the **YAML** tab.
 - In the **installedComponents** section, confirm that the components you installed have a status value of **true**.

7.2. INSTALLING RED HAT OPENSIFT DATA SCIENCE COMPONENTS BY USING THE WEB CONSOLE

The following procedure shows how to use the OpenShift Container Platform web console to install specific components of Red Hat OpenShift Data Science on your cluster when version 2 of the Red Hat OpenShift Data Science Operator is already installed on the cluster.

Prerequisites

- Version 2 of the Red Hat OpenShift Data Science Operator is [installed](#) on your OpenShift Container Platform cluster.
- You have cluster administrator privileges for your OpenShift Container Platform cluster.

Procedure

1. Log in to the OpenShift Container Platform web console as a cluster administrator.
2. In the web console, click **Operators** → **Installed Operators** and then click the Red Hat OpenShift Data Science Operator.
The **Operator details** page opens.
3. Create a **DataScienceCluster** object to install OpenShift Data Science components by performing the following actions:
 - a. Click the **Data Science Cluster** tab.
 - b. Click **Create DataScienceCluster**.
 - c. For **Configure via**, select **YAML view**.
An embedded YAML editor opens showing a default custom resource (CR) for the **DataScienceCluster** object.
 - d. In the **spec.components** section of the CR, for each OpenShift Data Science component shown, set the value of the **managementState** field to either **Managed** or **Removed**. These values are defined as follows:

Managed

The Operator actively manages the component, installs it, and tries to keep it active. The Operator will upgrade the component only if it is safe to do so.

Removed

The Operator actively manages the component but does not install it. If the component is already installed, the Operator will try to remove it.

4. Click **Create**.
The **Data Science Cluster** tab reopens.

Verification

- In the OpenShift Container Platform web console, click **Workloads** → **Pods**. In the **Project** list at the top of the page, select **redhat-ods-applications**. In the applications namespace, confirm that there are running pods for each of the OpenShift Data Science components that you installed.
- In the OpenShift Container Platform web console, click **Operators** → **Installed Operators** and then perform the following actions:
 - Click the Red Hat OpenShift Data Science Operator.

- Click the **Data Science Cluster** tab and select the default **DataScienceCluster** object shown on the page.
- Select the **YAML** tab.
- In the **installedComponents** section, confirm that the components you installed have a status value of **true**.


CHAPTER 8. ACCESSING THE OPENSIFT DATA SCIENCE DASHBOARD

After you have installed OpenShift Data Science and added users, you can access the URL for your OpenShift Data Science console and share the URL with the users to let them log in and work on their data models.

Prerequisites

- You have installed OpenShift Data Science on your OpenShift Container Platform cluster.
- You have added at least one user to the user group for OpenShift Data Science.

Procedure

1. Log in to OpenShift Container Platform web console.
2. Click the application launcher ().
3. Right-click on **Red Hat OpenShift Data Science** and copy the URL for your OpenShift Data Science instance.
4. Provide this instance URL to your data scientists to let them log in to OpenShift Data Science.

Verification

- Confirm that you and your users can log in to OpenShift Data Science by using the instance URL.

Additional resources

- [Logging in to OpenShift Data Science](#)
- [Adding users for OpenShift Data Science](#)

CHAPTER 9. ENABLING GPU SUPPORT IN OPENSIFT DATA SCIENCE

Optionally, to ensure that your data scientists can use compute-heavy workloads in their models, you can enable graphics processing units (GPUs) in OpenShift Data Science. To enable GPUs on OpenShift, you must install the NVIDIA GPU Operator. As a prerequisite to installing the NVIDIA GPU Operator, you must install the Node Feature Discovery (NFD) Operator. For information about how to install these operators, see [NVIDIA GPU Operator on Red Hat OpenShift Container Platform](#).



IMPORTANT

Follow the instructions in this chapter only if you want to enable GPU support in an unrestricted self-managed environment. To enable GPU support in a disconnected self-managed environment, see [Enabling GPU support in OpenShift Data Science](#) instead.

Additional resources

- [NVIDIA GPU Operator on Red Hat OpenShift Container Platform](#) in the NVIDIA documentation.
- [Node Feature Discovery Operator](#)