



# Red Hat OpenShift Data Science self-managed 1.32

## Upgrading OpenShift Data Science self-managed in a disconnected environment

Learn how to upgrade Red Hat OpenShift Data Science on OpenShift Container Platform in a disconnected environment



## Red Hat OpenShift Data Science self-managed 1.32 Upgrading OpenShift Data Science self-managed in a disconnected environment

---

Learn how to upgrade Red Hat OpenShift Data Science on OpenShift Container Platform in a disconnected environment

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

## Abstract

Learn about the OpenShift Data Science Operator upgrade process in a disconnected environment.

## Table of Contents

CHAPTER 1. OVERVIEW OF UPGRADING OPENSIFT DATA SCIENCE SELF-MANAGED .....	3
CHAPTER 2. CONFIGURING THE UPGRADE STRATEGY FOR OPENSIFT DATA SCIENCE .....	4
CHAPTER 3. UPGRADING OPENSIFT DATA SCIENCE IN A DISCONNECTED ENVIRONMENT .....	5



# CHAPTER 1. OVERVIEW OF UPGRADING OPENSIFT DATA SCIENCE SELF-MANAGED

As a cluster administrator, you can configure the Red Hat OpenShift Data Science Operator to upgrade automatically or manually in a disconnected environment. A disconnected environment is a network restricted environment where Operator Lifecycle Manager (OLM) cannot access the default OperatorHub and image registries, which require Internet connectivity.



## IMPORTANT

This guide describes how to configure a *minor-version* upgrade of version 1 (that is, an upgrade from version 1.x to version 1.y). To learn how to perform a major-version upgrade from version 1 to version 2, see [Upgrading from version 1 of OpenShift Data Science](#).



## IMPORTANT

Upgrading directly from the Red Hat OpenShift Data Science self-managed Beta version to the Generally Available (GA) release is not supported. See [Uninstalling Red Hat OpenShift Data Science self-managed Beta version prior to installing a General Availability \(GA\) release](#) for more information.

- If you select automatic updates, when a new version of the Red Hat OpenShift Data Science Operator is available, and you have updated your mirror registry content, Operator Lifecycle Manager (OLM) automatically upgrades the running instance of your Operator without human intervention.
- If you select manual updates, when a new version of the Red Hat OpenShift Data Science Operator is available and you have updated your mirror registry content, OLM creates an update request. A cluster administrator must manually approve the update request to update the Operator to the new version. See [Manually approving a pending Operator upgrade](#) for more information about approving a pending Operator upgrade to remain on a supported version of OpenShift Data Science self-managed.

Red Hat supports the current release version and three previous release versions of OpenShift Data Science self-managed. For example, if version 1.24 is the latest version, then versions 1.23, 1.22, and 1.21 remain supported, and version 1.20 and older no longer receive support.

### Additional resources

- [Operator Lifecycle Manager workflow](#)
- [Configuring the upgrade strategy for OpenShift Data Science](#)

## CHAPTER 2. CONFIGURING THE UPGRADE STRATEGY FOR OPENSIFT DATA SCIENCE

As a cluster administrator, you can configure the upgrade strategy for the Red Hat OpenShift Data Science Operator.



### IMPORTANT

The following procedure describes how to configure a *minor-version* upgrade of version 1 (that is, an upgrade from version 1.x to version 1.y). To learn how to perform a major-version upgrade from version 1 to version 2, see [Upgrading from version 1 of OpenShift Data Science](#).

### Prerequisites

- Access to the OpenShift Container Platform cluster as a user with the **cluster-admin** role.
- The Red Hat OpenShift Data Science Operator is installed.

### Procedure

1. Log in to the OpenShift Container Platform cluster web console.
2. From the **Administrator** perspective, navigate to **Operators** → **Installed Operators**.
3. Click on the **Red Hat OpenShift Data Science** Operator.
4. Click the **Subscription** tab.
5. Under **Update approval**, click on the pencil icon and select one of the following update strategies:
  - **Automatic:** New updates are installed as soon as they become available.
  - **Manual:** New updates need to be manually approved before installation begins.
6. Click **Save**.

### Additional resources

- [Upgrading installed Operators](#)



## CHAPTER 3. UPGRADING OPENSIFT DATA SCIENCE IN A DISCONNECTED ENVIRONMENT

You can upgrade the Red Hat OpenShift Data Science Operator in a disconnected environment by first updating your mirror registry content and then using the OpenShift Container Platform web console to update the Operator. A disconnected environment is a network restricted environment where Operator Lifecycle Manager (OLM) cannot access the default OperatorHub and image registries, which require Internet connectivity.



### IMPORTANT

The following procedure describes how to configure a *minor-version* upgrade of version 1 (that is, an upgrade from version 1.x to version 1.y). To learn how to perform a major-version upgrade from version 1 to version 2, see [Upgrading from version 1 of OpenShift Data Science](#).



### 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.

Follow this procedure to update your mirror registry content using the *mirror registry for Red Hat OpenShift* and the **oc-mirror** OpenShift CLI (**oc**) plug-in.

### Prerequisites

- A running OpenShift cluster, version 4.10 or greater.
- Access to the OpenShift cluster as a user with the **cluster-admin** role.
- A mirror registry that is populated with the initial image set. See [Mirroring images to a private registry for a disconnected installation](#) for more information.
- You have used the **oc-mirror** plug-in to mirror the initial image set to your mirror registry.
- You have access to the storage backend that was used for the initial execution of the **oc-mirror** plug-in.
- You have a GitHub account linked to a verified email address.
- If you plan to use the OpenShift Data Science disconnected installer helper, you have installed the following:
  - [Bash](#) (version 4.0 or above)
  - [yq](#)
  - [jq](#)
  - [skopeo](#)

- If you plan to use the OpenShift Data Science disconnected installer helper, you have cloned the [OpenShift Data Science disconnected installer helper](#) repository. For more information about how to clone a GitHub repository, see [Cloning a repository](#)

## Procedure

1. Run the OpenShift Data Science disconnected installer helper to obtain the values for your image set configuration.



### IMPORTANT

If you decide not to use the OpenShift Data Science disconnected installer helper, you can instead use an example image set configuration file (**rhods-`<version>`.md**) from the [OpenShift Data Science disconnected installer helper](#) repository.

The example image set configurations are for demonstration purposes only and might need further alterations depending on your deployment.

To identify the attributes most suitable for your deployment, examine the documentation and use cases in [Mirroring images for a disconnected installation using the oc-mirror plugin](#).

Open the relevant **rhods-`<version>`.md** file and skip to step 2.

- a. At a command-line terminal, change to the directory that contains the OpenShift Data Science disconnected installer helper repository.
- b. Enter the following command to run the OpenShift Data Science disconnected installer helper:

```
./rhods-disconnected-helper.sh -v rhods-<version>
```

Replace **version** with the version of OpenShift Data Science that you are upgrading to.

This example command generates a file containing an image set configuration for OpenShift Data Science version 1.31:

```
./rhods-disconnected-helper.sh -v rhods-1.31
```

The OpenShift Data Science disconnected installer helper generates a file (**rhods-`<version>`.md**) that contains an example image set configuration along with a separate list of notebook image values.

2. Open the **rhods-`<version>`.md** file in a text editor and examine its contents.
3. Open your **imageset-config.yaml** file and populate it with values suitable for the image set configuration in your deployment. Use the example image set configuration generated by the OpenShift Data Science disconnected installer helper. You might need to make additional alterations to the example image set configuration that are suitable for your deployment.
4. From the mirror registry host, run the **oc-mirror** command to mirror the specified image set configuration to disk:

```
$ oc mirror --config=./imageset-config.yaml file://mirror-rhods
```

- Replace **mirror-rhods** with the target directory where you want to output the image set file.
- The target directory path must start with **file://**.

5. Verify that the image set **.tar** file was created:

```
$ ls mirror-rhods
mirror_seq2_000000.tar
```

If an **archiveSize** value was specified in the image set configuration file, the image set might be separated into multiple **.tar** files.

6. Mirror the contents of the generated image set to the target mirror registry:

```
$ oc mirror --from=./mirror-rhods docker://registry.example.com:5000
```

- Replace **mirror-rhods** with the directory that contains your image set **.tar** files.
- Replace **registry.example.com:5000** with your mirror registry.

7. Verify that the YAML files are present for the **ImageContentSourcePolicy** and **CatalogSource** resources:

```
$ ls oc-mirror-workspace/results-1639608409/
catalogSource-rhods-operator-live-catalog.yaml
charts
imageContentSourcePolicy.yaml
mapping.txt
release-signatures
```

Replace **results-1639608409** with the name of your results directory.

8. Log in to the OpenShift CLI as a user with the cluster-admin role.

9. Install the generated **ImageContentSourcePolicy** and **CatalogSource** resources into the cluster:

```
$ oc apply -f ./oc-mirror-workspace/results-1639608409/imageContentSourcePolicy.yaml
$ oc apply -f ./oc-mirror-workspace/results-1639608409/catalogSource-rhods-operator-live-catalog.yaml
```

Replace **results-1639608409** with the name of your results directory.

10. Update the OpenShift Data Science Operator:

- If your OpenShift Data Science Operator is set to update automatically, the Operator Lifecycle Manager (OLM) automatically upgrades the running instance of your Operator without human intervention.
- If your OpenShift Data Science Operator is set to update manually, OLM creates an update request. A cluster administrator must manually approve the update request to update the Operator to the new version.

## Verification

- Verify the current version of the Red Hat OpenShift Data Science Operator:
  - a. Log in to the OpenShift Container Platform cluster web console.
  - b. Click **Operators** → **Installed Operators**.  
The **Installed Operators** page opens.
  - c. Locate the Red Hat OpenShift Data Science Operator and verify the version number.

#### **Additional resources**

- [Keeping your mirror registry content updated](#)
- [Updating your mirror registry content](#)