



Red Hat OpenShift Data Science 1

Installing OpenShift Data Science

Use Red Hat OpenShift Cluster Manager to install Red Hat OpenShift Data Science as an Add-on to your OpenShift Dedicated cluster

Red Hat OpenShift Data Science 1 Installing OpenShift Data Science

Use Red Hat OpenShift Cluster Manager to install Red Hat OpenShift Data Science as an Add-on to your OpenShift Dedicated cluster

Legal Notice

Copyright © 2021 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

Use Red Hat Cluster Manager to install Red Hat OpenShift Data Science as an Add-on to your OpenShift Dedicated cluster.

Table of Contents

PREFACE	3
CHAPTER 1. OVERVIEW OF DEPLOYING OPENSIFT DATA SCIENCE	4
CHAPTER 2. REQUIREMENTS FOR OPENSIFT DATA SCIENCE	5
CHAPTER 3. CONFIGURING AN IDENTITY PROVIDER FOR OPENSIFT DEDICATED	6
3.1. IDENTITY MANAGEMENT OPTIONS FOR OPENSIFT DATA SCIENCE	7
CHAPTER 4. ADDING ADMINISTRATIVE USERS FOR OPENSIFT DEDICATED	8
CHAPTER 5. INSTALLING OPENSIFT DATA SCIENCE ON OPENSIFT DEDICATED	9
CHAPTER 6. SHARING THE INSTANCE ADDRESS WITH USERS	11
CHAPTER 7. ADDITIONAL RESOURCES	12

PREFACE

This documentation is provided for the Field Trial release of Red Hat OpenShift Data Science.

See the following documents for service and life cycle information related to this Field Trial release:

- [OpenShift Data Science Service Definition](#)
- [OpenShift Data Science Life Cycle](#)

CHAPTER 1. OVERVIEW OF DEPLOYING OPENSIFT DATA SCIENCE

Read this section to understand how to deploy Red Hat OpenShift Data Science as a development and testing environment for data scientists.

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

1. Confirm that your OpenShift Dedicated cluster meets all requirements.
2. Configure an identity provider for OpenShift Dedicated.
3. Add administrative users for OpenShift Dedicated.
4. Install the OpenShift Data Science Add-on.
5. Configure user and administrator groups to provide user access to OpenShift Data Science.
6. Provide your users with the URL for the OpenShift Data Science instance.

Additional resources

- [Installing OpenShift Data Science on OpenShift Dedicated](#)
- [Adding users for OpenShift Data Science](#)

CHAPTER 2. REQUIREMENTS FOR OPENSIFT DATA SCIENCE

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

Installation requirements

You need to meet the following requirements before you are able to install OpenShift Data Science on your OpenShift Dedicated instance.

- **A Red Hat customer account**

Go to console.redhat.com and log in or register for a new account.

- **Product subscriptions**

Subscriptions for the following product and Add-on:

- Red Hat OpenShift Dedicated
- Red Hat OpenShift Data Science Add-on

Contact your Red Hat account manager to purchase new subscriptions. If you do not yet have an account manager, complete the form at <https://cloud.redhat.com/products/dedicated/contact/> to request one.


- **An OpenShift Dedicated cluster**

Use an existing cluster or create a new cluster by following the OpenShift Dedicated documentation: [Creating your cluster](#).

Your cluster must have at least 2 worker nodes with at least 8 CPUs and 32 GiB RAM (instance type **m5.2xlarge** or larger) available for OpenShift Data Science use when you install the Add-on. The installation process fails to start and an error is displayed if this requirement is not met.

CHAPTER 3. CONFIGURING AN IDENTITY PROVIDER FOR OPENSIFT DEDICATED

Configure an identity provider for your OpenShift Dedicated cluster to manage users and groups.



IMPORTANT

Adding more than one OpenShift Identity Provider can create problems when the same user name exists in multiple providers.

When **mappingMethod** is set to **claim** (the default mapping method for identity providers) and multiple providers have credentials associated with the same user name, the first provider used to log in to OpenShift is the one that works for that user, regardless of the order in which identity providers are configured.

Refer to [Identity provider parameters](#) in the OpenShift Dedicated documentation for more information about mapping methods.

Prerequisites

- Credentials for OpenShift Cluster Manager (<https://console.redhat.com/openshift/>).
- An existing OpenShift Dedicated cluster.

Procedure

1. Log in to OpenShift Cluster Manager (<https://console.redhat.com/openshift/>).
2. Click **Clusters**. The **Clusters** page opens.
3. Click the name of the cluster to configure.
4. Click the **Access control** tab.
5. Click **Add identity provider**.
 - a. Select your provider from the **Identity Provider** list.
 - b. Enter a unique **Name** for the identity provider. You cannot change this name later.
 - c. Select a **Mapping method** from the drop-down menu.
 - d. Enter a **Client ID** and **Client secret**.
 - e. Optional: Enter a **Hostname** to connect to either GitHub or GitHub Enterprise.
 - f. Select **Use organizations** or **Use teams** to restrict access to a particular GitHub organization or a GitHub team.
 - g. Enter the name of the organization or team you want to restrict access to. Click **Add more** to specify multiple organizations or teams that users can be a member of.
 - h. Click **Confirm**.

Verification

- The configured identity providers are visible on the **Access control** tab of the **Cluster details** page.

Additional resources

- [Configuring identity providers](#)
- [Identity management options for OpenShift Data Science](#)
- [Syncing LDAP groups](#)

3.1. IDENTITY MANAGEMENT OPTIONS FOR OPENSIFT DATA SCIENCE

Red Hat OpenShift Data Science supports the same authentication systems as Red Hat OpenShift Dedicated and Red Hat OpenShift Service on Amazon Web Services (AWS).

Check the appropriate documentation for your cluster for more information.

Additional resources

- [Supported identity providers on OpenShift Dedicated](#)
- [Supported identity providers for Red Hat OpenShift Service on AWS](#)

CHAPTER 4. ADDING ADMINISTRATIVE USERS FOR OPENSIFT DEDICATED


You need administrative access to install and configure OpenShift Data Science for your data scientist users.

Prerequisites

- Credentials for OpenShift Cluster Manager (<https://console.redhat.com/openshift/>).
- An existing OpenShift Dedicated cluster with an identity provider configured.

Procedure

1. Log in to OpenShift Cluster Manager (<https://console.redhat.com/openshift/>).
2. Click **Clusters**. The **Clusters** page opens.
3. Click the name of the cluster to configure.
4. Click the **Access control** tab.
5. Under **Cluster administrative users** click the **Add user** button. The **Add cluster user** popover appears.
6. Enter the user name in the **User ID** field.
7. Select an appropriate **Group** for the user.



IMPORTANT

If this user needs to use existing groups in an identity provider to control OpenShift Data Science access, select **cluster-admins**.

Check [Administering your cluster](#) in the OpenShift Dedicated documentation for more information about these user types.

8. Click **Add user**.

Verification

- The user name and selected group are visible in the list of **Cluster administrative users**.

Additional resources

- [Administering your OpenShift Dedicated cluster](#)

CHAPTER 5. INSTALLING OPENSIFT DATA SCIENCE ON OPENSIFT DEDICATED

You can install Red Hat OpenShift Data Science as an Add-on to your Red Hat OpenShift Dedicated cluster using Red Hat OpenShift Cluster Manager.

Prerequisites

- Purchase entitlements for OpenShift Data Science.
- Credentials for OpenShift Cluster Manager (<https://console.redhat.com/openshift/>).
- Administrator access to the OpenShift Dedicated cluster.

Procedure

1. Log in to OpenShift Cluster Manager (<https://console.redhat.com/openshift/>).
2. Click **Clusters**.
The **Clusters** page opens.
3. Click the name of the cluster you want to install OpenShift Data Science on.
The **Details** page for the cluster opens.
4. Click the **Add-ons** tab and locate the **Red Hat OpenShift Data Science** card.
5. Click **Install**. The **Configure Red Hat OpenShift Data Science** pane appears.
6. In the **Notification email** field, enter any email addresses that you want to receive important alerts about the state of Red Hat OpenShift Data Science, such as outage alerts.
7. Click **Install**.

Verification

- In OpenShift Cluster Manager, under the **Add-ons** tab for the cluster, confirm that the OpenShift Data Science card shows one of the following states:
 - **Installing** - installation is in progress; wait for this to change to **Installed**. This takes around 30 minutes.
 - **Installed** - installation is complete; verify that the **View in console** button is visible.
- In OpenShift Dedicated, 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](#)


CHAPTER 6. SHARING THE INSTANCE ADDRESS WITH USERS

After you have added users to Red Hat OpenShift Data Science, share the instance address with those users to let them log in and work on their data models.

Prerequisites

- You have installed OpenShift Data Science on your OpenShift Dedicated instance.
- You have added at least one user to the user group for OpenShift Data Science.

Procedure

1. Log in to OpenShift Dedicated 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 using the instance URL.

Additional resources

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

CHAPTER 7. ADDITIONAL RESOURCES

- [Identity management options for OpenShift Data Science](#)
- [Adding users for OpenShift Data Science](#)