



Red Hat OpenShift Data Foundation 4.9

Deploying OpenShift Data Foundation using IBM Cloud

Instructions on deploying Red Hat OpenShift Data Foundation using IBM Cloud

Red Hat OpenShift Data Foundation 4.9 Deploying OpenShift Data Foundation using IBM Cloud

Instructions on deploying Red Hat OpenShift Data Foundation using IBM Cloud

Legal Notice

Copyright © 2022 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

Read this document for instructions about how to install Red Hat OpenShift Data Foundation using Red Hat OpenShift Container Platform on IBM cloud clusters.

Table of Contents

MAKING OPEN SOURCE MORE INCLUSIVE	3
PROVIDING FEEDBACK ON RED HAT DOCUMENTATION	4
CHAPTER 1. DEPLOYING OPENSIFT DATA FOUNDATION USING IBM CLOUD	5
1.1. DEPLOYING ON IBM CLOUD PUBLIC	5
1.1.1. Deploying on classic infrastructure in IBM Cloud	5
1.1.2. Deploying on VPC infrastructure in IBM Cloud	5
1.2. DEPLOYING ON IBM CLOUD SATELLITE	6
1.2.1. Using OpenShift Data Foundation with the local storage present on your worker nodes in IBM Cloud Satellite	6
1.2.2. Using OpenShift Data Foundation with remote, dynamically provisioned storage volumes in IBM Cloud Satellite	6

MAKING OPEN SOURCE MORE INCLUSIVE

Red Hat is committed to replacing problematic language in our code, documentation, and web properties. We are beginning with these four terms: master, slave, blacklist, and whitelist. Because of the enormity of this endeavor, these changes will be implemented gradually over several upcoming releases. For more details, see [our CTO Chris Wright's message](#).

PROVIDING FEEDBACK ON RED HAT DOCUMENTATION

We appreciate your input on our documentation. Do let us know how we can make it better. To give feedback:

- For simple comments on specific passages:
 1. Make sure you are viewing the documentation in the *Multi-page HTML* format. In addition, ensure you see the **Feedback** button in the upper right corner of the document.
 2. Use your mouse cursor to highlight the part of text that you want to comment on.
 3. Click the **Add Feedback** pop-up that appears below the highlighted text.
 4. Follow the displayed instructions.
- For submitting more complex feedback, create a Bugzilla ticket:
 1. Go to the [Bugzilla](#) website.
 2. In the **Component** section, choose **documentation**.
 3. Fill in the **Description** field with your suggestion for improvement. Include a link to the relevant part(s) of documentation.
 4. Click **Submit Bug**.

Red Hat OpenShift Data Foundation 4.9 supports deployment of Red Hat OpenShift on IBM Cloud clusters in connected environments.

CHAPTER 1. DEPLOYING OPENSIFT DATA FOUNDATION USING IBM CLOUD

You can use Red Hat OpenShift Data Foundation for your workloads that run in IBM Cloud. These workloads might run in Red Hat OpenShift on IBM Cloud clusters that are in the public cloud or in your own IBM Cloud Satellite location.

1.1. DEPLOYING ON IBM CLOUD PUBLIC

When you create a Red Hat OpenShift on IBM Cloud cluster, you can choose between classic or Virtual Private Cloud (VPC) infrastructure. The Red Hat OpenShift Data Foundation managed cluster add-on supports both infrastructure providers. For classic clusters, the add-on deploys the OpenShift Data Foundation operator with the Local Storage operator. For VPC clusters, the add-on deploys the OpenShift Data Foundation operator which you can use with IBM Cloud Block Storage on VPC storage volumes.

Benefits of using the OpenShift Data Foundation managed cluster add-on to install OpenShift Data Foundation instead of installing from OperatorHub

- Deploy OpenShift Data Foundation from a single CRD instead of manually creating separate resources. For example, in the single CRD that add-on enables, you configure the namespaces, storagecluster, and other resources you need to run OpenShift Data Foundation.
- **Classic** - Automatically create PVs using the storage devices that you specify in your OpenShift Data Foundation CRD.
- **VPC** - Dynamically provision IBM Cloud Block Storage on VPC storage volumes for your OpenShift Data Foundation storage cluster.
- Get patch updates automatically for the managed add-on.
- Update the OpenShift Data Foundation version by modifying a single field in the CRD.
- Integrate with IBM Cloud Object Storage by providing credentials in the CRD.

1.1.1. Deploying on classic infrastructure in IBM Cloud

You can deploy OpenShift Data Foundation on IBM Cloud classic clusters by using the managed cluster add-on to install the OpenShift Data Foundation operator and the Local Storage operator. After you install the OpenShift Data Foundation add-on in your IBM Cloud classic cluster, you create a single custom resource definition that contains your storage device configuration details.

For more information, see the [Preparing your cluster for OpenShift Data Foundation](#).

1.1.2. Deploying on VPC infrastructure in IBM Cloud

You can deploy OpenShift Data Foundation on IBM Cloud VPC clusters by using the managed cluster add-on to install the OpenShift Data Foundation operator. After you install the OpenShift Data Foundation add-on in your IBM Cloud classic cluster, you create a custom resource definition that contains your worker node information and the IBM Cloud Block Storage for VPC storage classes that you want to use to dynamically provision the OpenShift Data Foundation storage devices.

For more information, see the [Preparing your cluster OpenShift Data Foundation](#).

1.2. DEPLOYING ON IBM CLOUD SATELLITE

With IBM Cloud Satellite, you can create a location with your own infrastructure, such as an on-premises data center or another cloud provider, to bring IBM Cloud services anywhere, including where your data resides. If you store your data by using Red Hat OpenShift Data Foundation, you can use Satellite storage templates to consistently install OpenShift Data Foundation across the clusters in your Satellite location. The templates help you create a Satellite configuration of the various OpenShift Data Foundation parameters, such as the device paths to your local disks or the storage classes that you want to use to dynamically provision volumes. Then, you assign the Satellite configuration to the clusters where you want to install OpenShift Data Foundation.

Benefits of using Satellite storage to install OpenShift Data Foundation instead of installing from OperatorHub

- Create versions your OpenShift Data Foundation configuration to install across multiple clusters or expand your existing configuration.
- Update OpenShift Data Foundation across multiple clusters consistently.
- Standardize storage classes that developers can use for persistent storage across clusters.
- Use a similar deployment pattern for your apps with Satellite Config.
- Choose from templates for an OpenShift Data Foundation cluster using local disks on your worker nodes or an OpenShift Data Foundation cluster that uses dynamically provisioned volumes from your storage provider.
- Integrate with IBM Cloud Object Storage by providing credentials in the template.

1.2.1. Using OpenShift Data Foundation with the local storage present on your worker nodes in IBM Cloud Satellite

For an OpenShift Data Foundation configuration that uses the local storage present on your worker nodes, you can use a Satellite template to configure your OpenShift Data Foundation configuration. Your cluster must meet certain requirements, such as CPU and memory requirements and size requirements of the available raw unformatted, unmounted disks. Choose a local OpenShift Data Foundation configuration when you want to use the local storage devices already present on your worker nodes, or statically provisioned raw volumes that you attach to your worker nodes.

For more information, see the [IBM Cloud Satellite local OpenShift Data Foundation storage documentation](#).

1.2.2. Using OpenShift Data Foundation with remote, dynamically provisioned storage volumes in IBM Cloud Satellite

For an OpenShift Data Foundation configuration that uses remote, dynamically provisioned storage volumes from your preferred storage provider, you can use a Satellite storage template to create your storage configuration. In your OpenShift Data Foundation configuration, you specify the storage classes that you want use and the volume sizes that you want to provision. Your cluster must meet certain requirements, such as CPU and memory requirements. Choose the OpenShift Data Foundation-remote storage template when you want to use dynamically provisioned remote volumes from your storage provider in your OpenShift Data Foundation configuration.

For more information, see the [IBM Cloud Satellite remote OpenShift Data Foundation storage documentation](#).

