



Red Hat OpenShift Container Storage 4.6

Deploying OpenShift Container Storage using IBM Cloud

How to install

Red Hat OpenShift Container Storage 4.6 Deploying OpenShift Container Storage using IBM Cloud

How to install

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

Read this document for instructions on installing Red Hat OpenShift Container Storage on IBM Cloud. Deploying and managing OpenShift Container Storage in IBM Cloud is a Technology Preview feature. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

Table of Contents

PREFACE	3
CHAPTER 1. DEPLOYING OPENSIFT CONTAINER STORAGE USING IBM CLOUD	4
1.1. DEPLOYING ON IBM CLOUD PUBLIC	4
1.1.1. Deploying on classic infrastructure in IBM Cloud	4
1.1.2. Deploying on VPC infrastructure in IBM Cloud	4
1.2. DEPLOYING ON IBM CLOUD SATELLITE	5
1.2.1. Using OpenShift Container Storage with the local storage present on your worker nodes in IBM Cloud Satellite	5
1.2.2. Using OpenShift Container Storage with remote, dynamically provisioned storage volumes in IBM Cloud Satellite	5

PREFACE

Red Hat OpenShift Container Storage 4.6 supports deployment on Red Hat OpenShift on IBM Cloud clusters in connected environments.

CHAPTER 1. DEPLOYING OPENSIFT CONTAINER STORAGE USING IBM CLOUD

You can use Red Hat OpenShift Container Storage 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.

- [Deploying on IBM Cloud public](#)
- [Deploying on IBM Cloud Satellite](#)

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 Container Storage managed cluster add-on supports both infrastructure providers. For classic clusters, the add-on deploys the OpenShift Container Storage operator with the Local Storage operator. For VPC clusters, the add-on deploys the OpenShift Container Storage operator which you can use with IBM Cloud Block Storage on VPC storage volumes.

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

- Deploy OpenShift Container Storage 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 Container Storage.
- **Classic** - Automatically create PVs using the storage devices that you specify in your OpenShift Container Storage CRD.
- **VPC** - Dynamically provision IBM Cloud Block Storage on VPC storage volumes for your OpenShift Container Storage storage cluster.
- Get patch updates automatically for the managed add-on.
- Update the OpenShift Container Storage 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 Container Storage on IBM Cloud classic clusters by using the managed cluster add-on to install the OpenShift Container Storage operator and the Local Storage operator. After you install the OpenShift Container Storage 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 Container Storage](#).

1.1.2. Deploying on VPC infrastructure in IBM Cloud

You can deploy OpenShift Container Storage on IBM Cloud VPC clusters by using the managed cluster add-on to install the OpenShift Container Storage operator. After you install the OpenShift Container Storage 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 Container Storage storage devices.

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

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 Container Storage, you can use Satellite storage templates to consistently install OpenShift Container Storage across the clusters in your Satellite location. The templates help you create a Satellite configuration of the various OpenShift Container Storage 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 Container Storage.

Benefits of using Satellite storage to install OpenShift Container Storage instead of installing from OperatorHub

- Create versions your OpenShift Container Storage configuration to install across multiple clusters or expand your existing configuration.
- Update OpenShift Container Storage 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 Container Storage cluster using local disks on your worker nodes or an OpenShift Container Storage 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 Container Storage with the local storage present on your worker nodes in IBM Cloud Satellite

For an OpenShift Container Storage configuration that uses the local storage present on your worker nodes, you can use a Satellite template to configure your OpenShift Container Storage 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 Container Storage 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 Container Storage storage documentation](#).

1.2.2. Using OpenShift Container Storage with remote, dynamically provisioned storage volumes in IBM Cloud Satellite

For an OpenShift Container Storage 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 Container Storage 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 Container Storage-remote storage template when you want to use dynamically provisioned remote volumes from your storage provider in your OpenShift Container Storage configuration.

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