



## **Cisco HyperFlex Systems Administration Guide for Kubernetes, Release 4.5**

**First Published:** 2021-01-22

**Last Modified:** 2023-06-05

### **Americas Headquarters**

Cisco Systems, Inc.  
170 West Tasman Drive  
San Jose, CA 95134-1706  
USA  
<http://www.cisco.com>  
Tel: 408 526-4000  
800 553-NETS (6387)  
Fax: 408 527-0883

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at [www.cisco.com/go/offices](http://www.cisco.com/go/offices).

The documentation set for this product strives to use bias-free language. For purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on standards documentation, or language that is used by a referenced third-party product.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/c/en/us/about/legal/trademarks.html>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2021–2023 Cisco Systems, Inc. All rights reserved.



## CONTENTS

[Full Cisco Trademarks with Software License](#) ?

---

### PREFACE

[Communications, Services, Bias-free Language, and Additional Information](#) v

---

### CHAPTER 1

[New and Changed](#) 1

[New and Changed Information](#) 1

---

### CHAPTER 2

[Cisco HyperFlex Kubernetes Support](#) 3

[Support Overview](#) 3

[Cisco HyperFlex CSI Interoperability Metrics](#) 3

---

### CHAPTER 3

[Cisco HyperFlex Container Storage Interface \(CSI\) for Kubernetes](#) 5

[About Cisco HyperFlex Kubernetes CSI](#) 5

[Cisco HyperFlex CSI Components](#) 5

---

### CHAPTER 4

[Configuring the Cisco HyperFlex CSI Integration for Kubernetes](#) 9

[Prerequisites](#) 9

[Administrator Host](#) 10

[Installing the Cisco HyperFlex CSI Integration for Kubernetes](#) 10

[Download the Cisco HyperFlex CSI Bundle](#) 10

[Open and Extract the Cisco HyperFlex CSI Bundle](#) 11

[Upload the Cisco HyperFlex CSI Container Image](#) 12

[Manually Import the Cisco HyperFlex CSI Container Image Directly to each Kubernetes Worker Node](#) 12

[Deploying HXCSI Using the hxcsi-setup Utility](#) 13

[Deployment Using Helm Utility](#) 14

Deploying HXCSI Using Cisco HyperFlex CSI Components 16

HXCSI Sample Pods 17

Create Cisco HyperFlex CSI Storage Class 18

Verifying Cisco HyperFlex CSI Storage Class Creation 19

---

**CHAPTER 5**

**Deploying Stateful Applications with Cisco HyperFlex CSI 21**

Prerequisites for Deploying Stateful Applications with Cisco HyperFlex CSI 21

Administrator Host 21

Deploying Stateful Applications 21

    Creating a Persistent Volume Claim 22

    Deploy Stateful Kubernetes Workload 23

---

**CHAPTER 6**

**Troubleshooting 25**

Troubleshooting 25

ImagePullBackOff Status Errors when Deploying HXCSI Pods 25

Volume Delete Fails 26

Application Pod in ContainerCreating State During Node Deletion 26

Deleted Pod Scheduled Back on the Same Node 26



## Communications, Services, Bias-free Language, and Additional Information

---

- To receive timely, relevant information from Cisco, sign up at [Cisco Profile Manager](#).
- To get the business impact you're looking for with the technologies that matter, visit [Cisco Services](#).
- To submit a service request, visit [Cisco Support](#).
- To discover and browse secure, validated enterprise-class apps, products, solutions and services, visit [Cisco Marketplace](#).
- To obtain general networking, training, and certification titles, visit [Cisco Press](#).
- To find warranty information for a specific product or product family, access [Cisco Warranty Finder](#).

### Documentation Feedback

To provide feedback about Cisco technical documentation, use the feedback form available in the right pane of every online document.

### Cisco Bug Search Tool

[Cisco Bug Search Tool](#) (BST) is a web-based tool that acts as a gateway to the Cisco bug tracking system that maintains a comprehensive list of defects and vulnerabilities in Cisco products and software. BST provides you with detailed defect information about your products and software.

### Bias-Free Language

The documentation set for this product strives to use bias-free language. For purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on standards documentation, or language that is used by a referenced third-party product.





# CHAPTER 1

## New and Changed

- [New and Changed Information, on page 1](#)

## New and Changed Information

This table summarizes the new and changed features for the Cisco HyperFlex Systems Administration Guide for Kubernetes, Release 4.5(x) and where they are documented.

Feature	Description	Date/Release Added	Where Documented
Cisco HyperFlex Systems Administration Guide for Kubernetes, HXCSI Release 1.2(3a)	End-of-Life	June 5, 2023	This document. For more information see the <a href="#">Cisco HyperFlex Container Storage Interface end-of-life and end-of-support notice</a> .
Cisco HyperFlex Systems Administration Guide for Kubernetes, Release 4.5(x)	This guide was introduced.	HX 4.5	N/A







## CHAPTER 2

# Cisco HyperFlex Kubernetes Support

- [Support Overview](#) , on page 3
- [Cisco HyperFlex CSI Interoperability Metrics](#), on page 3

## Support Overview

There are two main components that are important to consider when determining support for a Kubernetes version or distribution on Cisco HyperFlex.

- Support for the Kubernetes version or distribution with Cisco HyperFlex.
- Support for the Cisco HyperFlex Container Storage Interface (CSI) storage integration with the specific Kubernetes version or distribution.



---

**Note** Kubernetes storage special interest group (K8 SIG community) is not supported on Stretch Clusters.

---

While in general, Cisco HyperFlex supports any version or distribution of Kubernetes, there is a specific sub-set of versions and distributions that have been tested and are recommended with the Cisco HyperFlex CSI storage integration for Kubernetes. Additionally while it is possible to run Kubernetes and container-based workloads on Cisco HyperFlex without using the HyperFlex CSI storage integration, we strongly recommend that you leverage the native capability when running any stateful Kubernetes-based applications and services that require persistent storage.

The Cisco HyperFlex Kubernetes CSI Integration allows Cisco HyperFlex to dynamically provide persistent storage to stateful Kubernetes workloads running on Cisco HyperFlex. The integration enables orchestration of the entire Persistent Volume object lifecycle to be offloaded and managed by Cisco HyperFlex, while being driven (initiated) by developers and users through standard Kubernetes Persistent Volume Claim objects. Developers and users get the benefit of leveraging Cisco HyperFlex for their Kubernetes persistent storage needs with zero additional administration overhead from their perspective.

## Cisco HyperFlex CSI Interoperability Metrics

Cisco HyperFlex CSI and Kubernetes Platform Version and Distribution Interoperability:

Hyperflex Data Platform Version	CSI Spec Version	Kubernetes Version	Cisco Qualified CCP Version	Cisco Qualified Anthos Version	Open iSCSI
4.0(1a)	1.0	1.14	5.0, 5.1	1.1, 1.2, 1.3	
4.0(2a)	1.0	1.15	5.1, 5.2	1.3	
4.0(2b)	1.0	1.16	6.0, 7.0	1.4.1	
4.0(2c)	1.0	1.17	6.0, 7.0	1.5.1	
HXDP 4.5(1a)	1.2	1.18.2	-	-	Open iSCSI - 2.0.874-5ubuntu2.10
HXDP 4.5(2a) HXCSI 1.2(1a)	1.2	1.18.2, 1.19.8	-	-	Open iSCSI - 2.0.874-5ubuntu2.10
HXDP 4.5(2b) HXDP 4.5(2d) HXCSI 1.2(1b)	1.2	1.18.2, 1.19.8	-	-	Open iSCSI - 2.0.874-5ubuntu2.10




---

**Note** HXCSI has been qualified with open-iscsi version 2.0.874-5ubuntu2.10 on Ubuntu 18.04.

---



## CHAPTER 3

# Cisco HyperFlex Container Storage Interface (CSI) for Kubernetes

---

- [About Cisco HyperFlex Kubernetes CSI , on page 5](#)
- [Cisco HyperFlex CSI Components, on page 5](#)

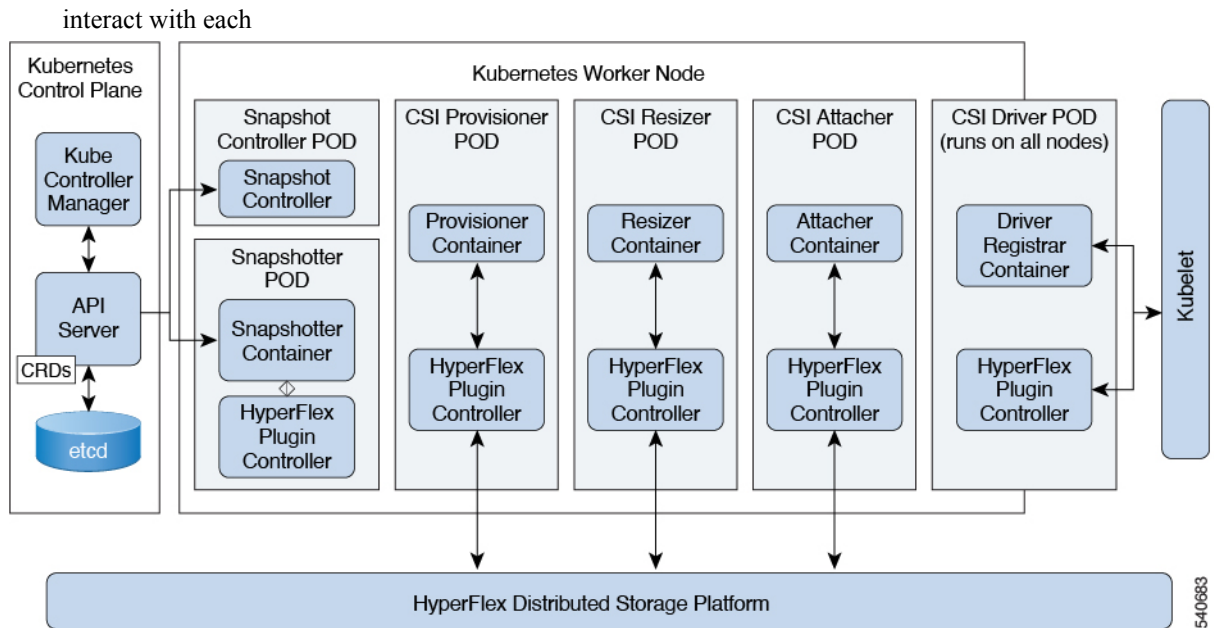
## About Cisco HyperFlex Kubernetes CSI

Cisco HyperFlex Container Storage Interface (CSI) is an out-of-tree container-based Kubernetes storage integration which is deployed and consumed through standard Kubernetes primitives such as Persistent Volume Claims and Storage Classes. Cisco HyperFlex CSI supports the following features:

- Dynamic creation and deletion of volumes
- Dynamic volume attach and detach
- Block access support
- Clone volume (when source volume is from the same Datastore)
- PV support with different filesystems (Ext4, Ext3, XFS)
- Volume space statistics reporting per CSI specs
- Multi-writer support (ReadWriteMany) for Block Mode only.
- Kubernetes 1.18, 1.19 support
- Kubernetes Cluster multitenancy target masking using dedicated initiator group
- Support for CSI 1.2 Spec APIs
- Volume resize support for block mode volumes and ext3, ext4, and xfs filesystem volumes. (expansion)
- CSI Plug-in installation and upgrade through Helm chart

## Cisco HyperFlex CSI Components

The Cisco HyperFlex CSI integration is deployed as containers on top of the target Kubernetes cluster. The following diagram shows the different components of the Cisco HyperFlex CSI deployment and how they



Deployment includes the following pods:

#### csi-attacher-hxcsi

- **Type:** StatefulSet
- **Number of Instances:** One per Kubernetes Cluster.
- **Purpose:** Required by CSI, but not currently used in Cisco deployment.

#### csi-provisioner-hxcsi

- **Type:** StatefulSet
- **Number of Instances:** One per Kubernetes Cluster
- **Purpose:** Watches Kubernetes Persistent Volume Claim objects and triggers CreateVolume and DeleteVolume operations as part of Kubernetes CSI spec.

#### csi-nodeplugin-hxcsi

- **Type:** DaemonSet
- **Number of Instances:** One per Kubernetes Worker Node
- **Purpose:** Discovery and formatting of provisioned HyperFlex iSCSI LUNs on Kubernetes worker nodes. Implements NodePublish/NodeUnpublish Volume APIs as part of Kubernetes CSI spec.

#### csi-resizer-hxcsi

- **Type:** StatefulSet
- **Number of Instances:** One per Kubernetes Cluster

- **Purpose:** Watches Kubernetes Persistent Volume Claim objects and triggers ControllerExpandVolume and NodeExpandVolume operations as part of Kubernetes CSI spec.





## CHAPTER 4

# Configuring the Cisco HyperFlex CSI Integration for Kubernetes

---

- [Prerequisites, on page 9](#)
- [Administrator Host, on page 10](#)
- [Installing the Cisco HyperFlex CSI Integration for Kubernetes, on page 10](#)
- [Verifying Cisco HyperFlex CSI Storage Class Creation, on page 19](#)

## Prerequisites

The following prerequisites must be met prior to configuring the Cisco HyperFlex CSI Integration.

On the HyperFlex cluster:

- Cisco HyperFlex cluster is installed and running HX 4.5(1a) or later.
- Configure an iSCSI network in HX Connect before installation. For more information about configuring an iSCSI network, see the [Cisco HyperFlex Administration Guide, Release 4.5](#).

On the Kubernetes cluster:

- Verify that all Kubernetes nodes have 2.0.874-5ubuntu2.10 version or later of the `open-iscsi` package installed ahead of proceeding with HXCSI. You can do so by running the following command: `$iscsid -version`.

To install `open-iscsi` version 2.0.874-5ubuntu2.10, you can run the following command: `apt-get install -y open-iscsi=2.0.874-5ubuntu2.10`

- Verify that each Kubernetes node has either a dedicated interface on the HX iSCSI network or routable access to the HX iSCSI network.
- To ensure `iscsid` gets started on system reboot run the following command:

```
sudo systemctl enable iscsid
```

The `iscsid` status should appear (as an example):

```
$ sudo systemctl status iscsid
```

```
iscsid.service - iSCSI initiator daemon (iscid)
```

```
Loaded: loaded (/lib/systemd/system/iscid.service; enabled; vendor preset: enabled
```

- Ensure that each Kubernetes primary (also known as the "master") host system contains the file `/etc/kubernetes/manifests/kube-controller-manager.yaml` which includes the following:  
`--disable-attach-detach-reconcile-sync=true`.
- Add the following text to the `-command` section of the file:  
`--disable-attach-detach-reconcile-sync=true`

## Administrator Host

In this guide, the Administrator Host is used to refer to a Linux-based system that runs `kubectl` commands, etc. against the Kubernetes cluster. While this is typically a separate system (VM) that is not part of the Kubernetes cluster, you can use one of the Kubernetes nodes as the administrator host if you do not wish to install/manage a separate system (VM).

## Installing the Cisco HyperFlex CSI Integration for Kubernetes

To install Cisco HyperFlex CSI Integration, complete the following procedures in the order presented:

### Download the Cisco HyperFlex CSI Bundle

To download the Cisco HyperFlex CSI bundle (file) perform the following steps:

#### Procedure

- 
- Step 1** Go to <https://software.cisco.com>
- Step 2** Log in using your Cisco ID and credentials.
- Step 3** In the **Download & Upgrade** section, click **Software Download**.
- Step 4** In the **Select a Product** search field, type **HyperFlex HX Data Platform** and click **Enter**.
- Step 5** Using the Release navigation pane on the left, select the HyperFlex Data Platform software version running on the cluster.
- Cisco HyperFlex Data Platform Release 4.5(x) or later requires Cisco HyperFlex CSI integration.
- Step 6** In the main navigation pane, locate and download the “Cisco HyperFlex Kubernetes Container Storage Interface (HX-CSI) bundle (tar.gz) file to your local machine.
- Henceforth, the Cisco HyperFlex Kubernetes Container Storage Interface (HX-CSI) bundle (tar.gz) file shall be referred to as the “Cisco HyperFlex CSI bundle”
- Step 7** On the administrator host, create a new directory called `hxcsi`.
- Example:**
- ```
administrator-host:~$ mkdir hxcsi
```
- Step 8** Using secure copy (scp) or other preferred file transfer method, transfer (move or copy) the downloaded Cisco HyperFlex CSI bundle from your local machine to the “`hxcsi`” directory on the administrator host. The result should look like the following:



**Example:**

```
administrator-host:hxcsi$ ls
hxcsi-1.2.1-601.tar.gz
```

---

**What to do next**

Open and Extract the Cisco HyperFlex CSI Bundle

## Open and Extract the Cisco HyperFlex CSI Bundle

Perform the following steps to open the Cisco HyperFlex CSI bundle.:

**Before you begin**

Download the Cisco HyperFlex CSI bundle.

**Procedure**

Use the `tar` command to unarchive the HyperFlex CSI bundle (.tar.gz file).

**Example:**

```
administrator-host:hxcsi$ tar -xf hxcsi-1.2.0-601.tar.gz
```

---

Once completed, the following directory structure should exist:

- **examples (directory)** – includes some example YAML files for using the HXCSI integration
- **images (directory)** – includes HXCSI docker container image for the HXCSI integration. It also includes the base CSI images for the Provisioner, Attacher, Node-driver and Resizer.
- **setup (directory)** – includes the setup script for deploying the HXCSI integration
- **support(directory)** – includes the script for collecting useful logs to help debugging.
- **hxcsi-1.2.1.tgz (file)**– this is the HELM chart package for this release of the HXCSI

**Example**

```
administrator-host:hxcsi$ ls -l
total 133196
-rw-r--r--  1 ubuntu ubuntu    6791 May 10 11:23 hxcsi-1.2.1.tgz
drwxr-xr-x  2 ubuntu ubuntu    4096 May 10 11:23 support
drwxr-xr-x  2 ubuntu ubuntu    4096 May 10 11:23 setup
drwxr-xr-x  2 ubuntu ubuntu    4096 May 10 11:23 images
drwxr-xr-x 11 ubuntu ubuntu    4096 May 10 11:23 examples
```

**What to do next**

Upload the Cisco HyperFlex CSI Container Image

**Upload the Cisco HyperFlex CSI Container Image**

The Cisco HyperFlex CSI integration components are deployed from a single container image provided in the “images” directory of the Cisco HyperFlex CSI bundle. The hxcsi container image leverages the other four base CSI images in the same directory. Before you can deploy the container image, move the container image to a location that is accessible to Docker running on the Kubernetes cluster worker nodes.

**Manually Import the Cisco HyperFlex CSI Container Image Directly to each Kubernetes Worker Node**

To add the Cisco HyperFlex CSI container image directly to each Kubernetes worker node, perform the following steps:

**Before you begin**

Open the Cisco HyperFlex CSI Bundle

**Procedure**

**Step 1** On the administrator host, copy the Cisco HyperFlex CSI container image (.tar) file, located in the “images” directory, to the /tmp directory on each Kubernetes worker node.

**Example:**

```
administrator-host:hxcsi$ scp ./images/hxcsi-1.2.1-601.tar k8s-worker1:/tmp
```

```
administrator-host:hxcsi$ scp ./images/hxcsi-1.2.1-601.tar k8s-worker2:/tmp
```

```
administrator-host:hxcsi$ scp ./images/hxcsi-1.2.1-601.tar k8s-workerN:/tmp
```

**Step 2** Copy the other base CSI container image files on each Kubernetes node:

**Step 3** On each Kubernetes worker node, use the `docker load --input` command to load the Cisco HyperFlex CSI container image.

**Example:**

```
k8s-worker1:/tmp# docker load -input ./hxcsi-1.2.2a-626.tar
Loaded image: hxcsi:hxcsi-1.2.1-601
```

```
k8s-worker2:/tmp# docker load -input ./hxcsi-1.2.1-601.tar Loaded image: hxcsi:hxcsi-1.2.1-601
```

```
k8s-workerN:/tmp# docker load -input ./hxcsi-1.2.1-601.tar Loaded image: hxcsi:hxcsi-1.2.1-601
```

**Step 4** Docker load the other base CSI container image files on each Kubernetes node:

```
csi-attacher-3.0.2-cisool.tar,csi-node-driver-registrar-2.0.1-cisool.tar,csi-resizer-1.0.1-cisool.tar,csi-provisioner-2.0.4-cisool.tar
```

**What to do next**

Install Cisco HyperFlex CSI.

**Deploying HXCSI Using the hxcsi-setup Utility**

In order to deploy the Cisco HyperFlex CSI integration, you must run the `hxcsi-setup` script. The `hxcsi-setup` script resides in the “setup” directory and automatically generates the necessary YAML files or Helm chart that then get applied (submitted) to the Kubernetes cluster to deploy the Cisco HyperFlex CSI components.

The following table describes parameters that you can provide with the `hxcsi-setup` command.

**Table 1: hxcsi-setup Parameters**

| Parameter                         | Required or Optional                                                    | Description                                                                                                                                                                                                                            |
|-----------------------------------|-------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>-clientId</code>            | Optional                                                                | Client ID for the Tenant.<br><br><b>Note</b> You can create multiple Kubernetes clusters to request storage from the same HX cluster. The “clientId” parameter helps isolate the storage allocation for each of these clients/tenants. |
| <code>-cluster-name</code>        | Required                                                                | Provides a name to uniquely identify this specific Kubernetes cluster                                                                                                                                                                  |
| <code>-helm-chart</code>          | Optional                                                                | Generates helm chart for helm install (default will generate YAML files)                                                                                                                                                               |
| <code>-hx-csi-image string</code> | Required                                                                | Name and location of the Cisco HyperFlex CSI container image. This tells Kubernetes where the Cisco HyperFlex CSI container image should be pulled.<br><br><a href="#">1</a>                                                           |
| <code>-iscsi-url string</code>    | Required                                                                | HyperFlex iSCSI cluster IP address of the eth-iscsi1:0 interface on the HyperFlex cluster. For more information, see <a href="#">Prerequisites, on page 9</a> .                                                                        |
| <code>-output-dir string</code>   | Optional                                                                | Output directory (default “./hxcsi-deploy/”)                                                                                                                                                                                           |
| <code>-password string</code>     | Required<br><br>(if not entered initially, you will be prompted for it) | password to HX cluster API                                                                                                                                                                                                             |
| <code>-token string</code>        | Optional                                                                | Service Authentication Token. You can create the token out of band before invoking the <code>hxcsi-setup</code> .                                                                                                                      |
| <code>-url string</code>          | Required                                                                | HyperFlex cluster management IP address. This IP is used for volume provisioning.                                                                                                                                                      |

| Parameter        | Required or Optional | Description                                           |
|------------------|----------------------|-------------------------------------------------------|
| -username string | Required             | user name to the HX cluster API (i.e., "admin")       |
| -docker-registry | Optional             | Docker registry name (e.g. mydockerhub.com/hx-docker) |

<sup>1</sup> If the Cisco HyperFlex CSI container image was imported directly into docker on each Kubernetes worker node, then the format for this parameter should be entered as **<repository\_name>:<tag>**.

### Before you begin

Upload the Cisco HyperFlex CSI Container Image and the base CSI container images.

### Procedure

On the administrator host, use the `hxcsi-setup` command in the "setup" directory to create the required Cisco HyperFlex CSI deployment files.

**Note** Prior to running the `hxcsi-setup` command, ensure that the IPs specified in the `url` and the `iscsi-url` parameters are reachable from the Kubernetes nodes.

### Example:

The following example shows a Cisco HyperFlex CSI container image deployment that has been uploaded on each node. The image name is `hxcsi` and the tag name is `hxcsi-1.2.1-601`.

```
administrator-host:hxcsi$ ./setup/hxcsi-setup -cluster-name demo-hxcsi
-hx-csi-image hxcsi-1.2.1-601
-iscsi-url 10.2.17.18 -url 10.2.17.13 -username admin

password for [admin] at [10.2.17.13]: *****
wrote config to hxcsi-deploy/hxcsi-config.yaml
wrote config to hxcsi-deploy/csi-attacher-hxcsi.yaml
wrote config to hxcsi-deploy/csi-nodeplugin-hxcsi.yaml
wrote config to hxcsi-deploy/csi-provisioner-hxcsi.yaml
wrote config to hxcsi-deploy/csi-attacher-rbac.yaml
wrote config to hxcsi-deploy/csi-nodeplugin-rbac.yaml
wrote config to hxcsi-deploy/csi-provisioner-rbac.yaml
wrote config to hxcsi-deploy/csi-resizer-rbac.yaml
wrote config to hxcsi-deploy/csi-resizer-hxcsi.yaml
```

### What to do next

Deploy the Cisco HyperFlex CSI Components

## Deployment Using Helm Utility

In order to deploy the Cisco HyperFlex CSI integration using the Helm chart package, you must run the `helm` utility (version v3.5.2 or later). The following table describes the parameters that you can provide with the `helm` command.

Table 2: HELM Command Parameters

| Parameter Name               | Required or Optional | Description                                                                                                                                          |
|------------------------------|----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------|
| Hx.clientId string           | Optional             | ClientID for the Tenant.<br>[Refer to note in Table-1]                                                                                               |
| Hx.iscsiUrl string           | Required             | HyperFlex iSCSI cluster IP address of the eth-iscsi1:0 interface on the HyperFlex cluster. For more information, see <a href="#">Prerequisites</a> . |
| Hx.token string              | Required             | Service Authentication Token. You must create the token out of band before invoking the helm install.                                                |
| Hx.url string                | Required             | HyperFlex cluster management IPaddress. This IP is used for volume provisioning.                                                                     |
| hx.dockerRegistryName string | Required             | Docker registry name (e.g. mydockerhub.com/hx-docker)                                                                                                |



**Note** One of the ways you can create a token is to use the `hxcsi-setup` utility described previously. The token is located in the generated “`hxcsi-deploy/hxcsi-config.yaml`” file. This token needs to be base64 decoded before passing to the `helm` utility.



**Note** Prior to running the following command, ensure that the IPs specified in the `url` and the `iscsi-url` parameters are reachable from the Kubernetes nodes.

## Procedure

The following example shows a Cisco HyperFlex CSI container image deployment that has been uploaded on the docker registry accessible from each Kubernetes node.

### Example:

```
administrator-host:hxcsi$ helm install hxcsi hxcsi-1.2.1.tgz --set hx.url=10.2.17.13 \
--set hx.iscsiUrl=10.2.17.18 --set hx.clientId=myClientId \
--set hx.dockerRegistryName=mydockerhub.com/hx-docker --set hx.token=myToken123456
administrator-host:hxcsi$
NAME: hxcsi
LAST DEPLOYED: Wed May 26 15:01:59 2021
NAMESPACE: default
STATUS: deployed
REVISION: 1
TEST SUITE: None
administrator-host:hxcsi$ kubectl get pods
NAME READY STATUS RESTARTS AGE
csi-attacher-hxcsi-0 2/2 Running 0 9s
csi-nodeplugin-hxcsi-rm4h4 2/2 Running 0 9s
csi-nodeplugin-hxcsi-w5fxt 2/2 Running 0 9s
```

```
csi-provisioner-hxcsi-0 2/2 Running 0 9s
csi-resizer-hxcsi-0 2/2 Running 0 9s
administrator-host:hxcsi$
administrator-host:hxcsi$ helm status hxcsi
NAME: hxcsi
LAST DEPLOYED: Wed May 26 15:01:59 2021
NAMESPACE: default
STATUS: deployed
REVISION: 1
```

---

### What to do next

Create the Cisco HyperFlex CSI Storage Class.

## Deploying HXCSI Using Cisco HyperFlex CSI Components

After running the `hxcsi-setup` script and generating the Cisco HyperFlex CSI deployment files, a new “`hxcsi-deploy`” directory is created on the administrator host.

```
root@administrator-host:hxcsi$ ls
examples hxcsi-1.2.1-601.tar.gz hxcsi-1.2.1.tgz hxcsi-deploy images setup support
```

### Before you begin

Create the Cisco HyperFlex CSI Deployment files.

### Procedure

---

- Step 1** On the administrator host, use the `kubectl create -f` command to deploy the Cisco HyperFlex CSI components.

#### Example:

```
administrator-host:hxcsi$ kubectl create -f ./hxcsi-deploy/
service/csi-attacher-hxcsi created
statefulset.apps/csi-attacher-hxcsi created
serviceaccount/csi-attacher created
clusterrole.rbac.authorization.k8s.io/external-attacher-runner created
clusterrolebinding.rbac.authorization.k8s.io/csi-attacher-role created
daemonset.apps/csi-nodeplugin-hxcsi created
serviceaccount/csi-nodeplugin created
clusterrole.rbac.authorization.k8s.io/csi-nodeplugin created
clusterrolebinding.rbac.authorization.k8s.io/csi-nodeplugin created
service/csi-provisioner-hxcsi created
statefulset.apps/csi-provisioner-hxcsi created
serviceaccount/csi-provisioner created
clusterrole.rbac.authorization.k8s.io/external-provisioner-runner created
clusterrolebinding.rbac.authorization.k8s.io/csi-provisioner-role created
deployment.apps/csi-resizer-hxcsi created
serviceaccount/csi-resizer created
clusterrole.rbac.authorization.k8s.io/external-resizer-runner created
clusterrolebinding.rbac.authorization.k8s.io/csi-resizer-role created
role.rbac.authorization.k8s.io/external-resizer-cfg created
rolebinding.rbac.authorization.k8s.io/csi-resizer-role-cfg created
secret/hxcsitoken created
configmap/hxcsi-config created
```

**Step 2** On the administrator host, use the `kubectl get pods` command to verify the HXCSI components have been deployed and have a status of `Running`.

**Example:**

**Note** There should be one instance of the `csi-attacher-hxcsi` pod, one instance of the `csi-provisioner-hxcsi` pod, one instance of the `csi-resizer-hxcsi` pod, and then a single instance of the `csi-nodeplugin-hxcsi` pod for each Kubernetes worker node. Therefore, if you have a total of two Kubernetes worker nodes, you should see two instances of the `csi-nodeplugin-hxcsi` pod as shown in the following example:

```
administrator-host:hxcsi$ kubectl get pods
NAME                                READY   STATUS    RESTARTS   AGE
csi-attacher-hxcsi-0                2/2    Running   0           37h
csi-nodeplugin-hxcsi-2nsfq          2/2    Running   2           37h
csi-nodeplugin-hxcsi-qjh9n         2/2    Running   2           37h
csi-provisioner-hxcsi-0             2/2    Running   0           37h
csi-resizer-hxcsi-0                 2/2    Running   0           37h
```

**What to do next**

Create Cisco HyperFlex CSI Storage Class.

## HXCSI Sample Pods

The HXCSI package includes several examples to create Pods.

**Table 3: HXCSI Package Samples**

| # | Directory Name           | Description                                                                                                                            |
|---|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------|
| 1 | sample-hxcsi             | Pod running <code>nginx</code> . Basic example creates a PVC named <code>'hxpvcclaim-default'</code>                                   |
| 2 | sample-hxcsi-csi-clone   | Pod that creates a clone from the <code>'sample-hxcsi'</code> , showing <code>'dataSource'</code> as <code>'hxpvcclaim-default'</code> |
| 3 | sample-hxcsi-ds          | Pod using named datastore <code>'test-ds'</code>                                                                                       |
| 4 | sample-hxcsi-fs          | Pod using default datastore and file system type <code>'xfs'</code>                                                                    |
| 5 | sample-hxcsi-no-ds       | Pod using default datastore and default file system.                                                                                   |
| 6 | sample-hxcsi-no-ds-clone | Clone from the sample <code>'sample-hxcsi-no-ds'</code> where datastore name was specified.                                            |
| 7 | sample-resize-block      | Resize the block volume, uses attribute <code>- allowVolumeExpansion: true</code>                                                      |
| 8 | sample-resize-fs         | Resize the default file-system <code>'ext4'</code> , uses <code>allowVolumeExpansion: true</code>                                      |
| 9 | sample-resize-clone      | Clone from the <code>hxpvcclaim-default-resize</code> <code>'sample-resize-fs'</code> example Pod.                                     |



**Note** To resize a volume:

1. Change the size of the volume in the PVC's yaml file.
2. Run the `kubectl apply -f <pvc.yaml>` command to apply the new size setting.

## Create Cisco HyperFlex CSI Storage Class

Once the components are up and running, you must need to now create a Storage Class that allows developers to consume storage through the Cisco HyperFlex CSI integration.

### Before you begin

Deploy the Cisco HyperFlex CSI Components

### Procedure

**Step 1** On the administrator host, create a file named “`hxcsi-storage-class.yaml`” with the following contents:

#### Example:

```
kind: StorageClass
apiVersion: storage.k8s.io/v1
metadata:
  name: csi-hxcsi-default
provisioner: csi-hxcsi
parameters:
  datastore: default-ds
  datastoreSize: "20000000000000"
```

You can specify the datastore name and size in the parameter section as above. You can optionally choose to make this the default Storage Class, which means that the Cisco HyperFlex CSI storage integration will be used by default for any Persistent Volume Claims that do not otherwise specify any other Storage Class to use. If you choose to make the Cisco HyperFlex CSI Storage Class the default storage class, then your “`hxcsi-storage-class.yaml`” file should contain the following contents:

#### Example:

```
kind: StorageClass
apiVersion: storage.k8s.io/v1
metadata:
  name: csi-hxcsi-default
  annotations:
    storageclass.kubernetes.io/is-default-class: "true"
provisioner: csi-hxcsi
parameters:
```

**Note** If a datastore does not already exist, a new datastore will be created. If you do not specify a datastore name, a default datastore with the name “`iscsiDs`” is created.

**Note** Always create a datastore that is larger than the volumes to be created.



**Step 2** On the administrator host, use the `kubectl create -f` command to create the Cisco HyperFlex CSI Storage Class.

**Example:**

```
root@administrator-host:hxcsi$ kubectl create -f ./hxcsi-storage-class.yaml
storageclass.storage.k8s.io/csi-hxcsi-default created
```

**Sample Storage Class for Resize Volume**

**Example:**

```
kind: StorageClass
apiVersion: storage.k8s.io/v1
metadata:
  name: csi-hxcsi-default-resize
provisioner: csi-hxcsi
parameters:
  datastore: default-ds
  datastoreSize: "20000000000000"
allowVolumeExpansion: true
```

Note that resizing the volume only allows volumes that are provisioned against this storageclass to support resizing. To change the actual size of the volume, you will need to edit the PVC specifications to change to the new size, for example, edit the PVC YAML file, and run `kubectl apply -f <pvc-yaml>`

**Sample Storage Class for File System**

**Example:**

```
kind: StorageClass
apiVersion: storage.k8s.io/v1
metadata:
  name: csi-hxcsi-default-fs
provisioner: csi-hxcsi
parameters:
  fsType: xfs
```

Note: Default file system is “ext4”

---

**What to do next**

Verify Cisco HyperFlex CSI Storage Class Creation.

## Verifying Cisco HyperFlex CSI Storage Class Creation

To verify the storage class creation perform the following step:




---

**Note** If setting the Cisco HyperFlex CSI Storage Class as the default, verify that “(default)” is present next to the Storage Class name.

---

**Before you begin**

Create Cisco HyperFlex CSI Storage Class.

## Procedure

---

On the administrator host, use the `kubectl get sc` command to verify the Cisco HyperFlex CSI Storage Class was created.

### Example:

```
root@administrator-host:hxcsi$ kubectl get sc
NAME          PROVISIONER  AGE
csi-hxcsi    (default)   csi-hxcsi   67s
```

---



## CHAPTER 5

# Deploying Stateful Applications with Cisco HyperFlex CSI

---

- [Prerequisites for Deploying Stateful Applications with Cisco HyperFlex CSI, on page 21](#)
- [Administrator Host, on page 21](#)
- [Deploying Stateful Applications, on page 21](#)

## Prerequisites for Deploying Stateful Applications with Cisco HyperFlex CSI

The following prerequisites must be met prior to deploying stateful applications using the HyperFlex CSI storage integration.

- Cisco HyperFlex cluster is installed and running HX 4.5(2a) or later
- The Cisco HyperFlex CSI integration has been deployed.
- From the iSCSI tab in HX Connect, you must first create an iSCSI network. For more information, see the [Cisco HyperFlex Administration Guide, Release 4.5](#).

## Administrator Host

In this guide, the Administrator Host is used to refer to a Linux-based system that runs kubectl commands, etc. against the Kubernetes cluster. While this is typically a separate system (VM) that is not part of the Kubernetes cluster, you can use one of the Kubernetes nodes as the administrator host if you do not wish to install/manage a separate system (VM).

## Deploying Stateful Applications

To deploy stateful applications, perform the following procedures:

## Creating a Persistent Volume Claim

A Persistent Volume Claim is simply a request for storage by a user. Users specify their storage requirements, the size or capacity of the storage required, and other options. Depending on the associated Storage Class, the storage requirements are routed to the appropriate provisioner which knows how to provision the requested storage, and make it available to Kubernetes.




---

**Note** The maximum PVC size is 64Ti. The minimum PVC size supported is 1 Gi.

---




---

**Note** You can create volumes that are protected by CHAP. The limit that you can create with one storage class for each target is 255 volumes (Persistent Volume Claims).

---

### Procedure

---

**Step 1** On the administrator host, create a file named “message-board-pvc.yaml” with the following contents

**Example:**

```
administrator-host:hxcsi$ cat ./message-board-pvc.yaml
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: message-board-pvc
spec:
  storageClassName: csi-hxcsi-default
  accessModes:
    - ReadWriteOnce
  resources:
    requests:
      storage: 10Gi
```

**Step 2** On the administrator host, use the `kubectl create -f` command to create the Persistent Volume Claim.

**Example:**

```
administrator-host:hxcsi$ kubectl create -f ./message-board-pvc.yaml

persistentvolumeclaim/message-board-pvc created
```

**Step 3** On the administrator host, use the `kubectl get pvc` command to verify the Persistent Volume Claim was created and is successfully bound to a Persistent Volume.

**Example:**

```
administrator-host:hxcsi$ kubectl get pvc

NAME          STATUS  VOLUME          CAPACITY  ACCESS  MODES  STORAGECLASS  AGE
message-board-pvc  BOUND  pvc-8069462e-662c-11e9-a163-005056a086d9  10Gi  RWO
csi-hxcsi-default  20s
```

---

## Deploy Stateful Kubernetes Workload

Kubernetes workloads come in various forms, such as Pods and Deployments regardless of the type of Kubernetes workload, each can leverage persistent storage using the Cisco HyperFlex CSI integration and Persistent Volume Claims. The following shows the deployment of a sample open source application called Cisco Message Board that can be used to test the Cisco HyperFlex CSI integration. You can also test with your own applications following the same methodology and procedures.

### Procedure

**Step 1** On the administrator host, create the YAML file which defines the workload to be deployed.

#### Example:

The following shows the YAML file for the example Cisco Message Board application which will create both a Kubernetes Deployment and a Kubernetes Service which will allow for connecting to the deployed Cisco Message Board application through a NodePort.

**Note** That we are referencing the Persistent Volume Claim name in the “volumes” section of the Kubernetes Deployment definition. In this example, the Persistent Volume bound to the “message-board-pvc” Persistent Volume Claim will be mounted inside the “message\_board:version1” container at the “/sqlldb” location (path)

```
administrator-host:hxcsi$ cat ./message-board-deployment.yaml
apiVersion: apps/v1
kind: Deployment
metadata:
  name: message-board
  labels:
    app: message-board
spec:
  replicas: 1
  selector:
    matchLabels:
      app: message-board
  template:
    metadata:
      labels:
        app: message-board
        name: message-board
    spec:
      volumes:
        - name: demovolume1
          persistentVolumeClaim:
            claimName: message-board-pvc
      containers:
        - name: message-board
          image: michzimm/message_board:version1
          ports:
            - containerPort: 5000
          volumeMounts:
            - mountPath: "/sqlldb"
              name: demovolume1
---
apiVersion: v1
kind: Service
metadata:
  name: message-board
  labels:
```

```

    name: message-board
  namespace: default
spec:
  type: NodePort
  ports:
  - port: 5000
    nodePort: 30002
  selector:
    name: message-board

```

**Step 2** On the administrator host, use the `kubectl create -f` command to create the Deployment and Service.

**Example:**

```

administrator-host:hxcsi$ kubectl create -f ./message-board-deployment.yaml
deployment.apps/message-board created
service/message-board created

```

**Step 3** On the administrator host, use the `kubectl get pods` command to check the status of the deployed Pods.

**Example:**

```

administrator-host:hxcsi$ kubectl get pods
NAME                                READY   STATUS    RESTARTS   AGE
csi-attacher-hxcsi-0                2/2     Running   0           3h51m
csi-nodeplugin-hxcsi-9fgsf          2/2     Running   0           3h51m
csi-nodeplugin-hxcsi-qqvwj         2/2     Running   0           3h51m
csi-provisioner-hxcsi-0             2/2     Running   0           3h51m
csi-resizer-hxcsi-5b444c8478-6qxws 2/2     Running   0           3h51m
message-board-6df65d6b59-49xhq     1/1     Running   0           95s

```

**Example:**

**Step 4** On the administrator host, use the `kubectl get services` command to check the status of the deployed Service.

**Example:**

```

root@administrator-host:hxcsi$ kubectl get services
NAME                                TYPE          CLUSTER-IP      EXTERNAL-IP   PORT(S)          AGE
csi-attacher-hxcsi                  ClusterIP     10.98.79.159    <none>        12346/TCP        3h53m
csi-provisioner-hxcsi                ClusterIP     10.99.73.185    <none>        12345/TCP        3h53m
kubernetes                           ClusterIP     10.96.0.1       <none>        443/TCP          4h24m
message-board                         NodePort      10.107.227.152 <none>        5000:30002/TCP  2m59s

```

For the sample Cisco Message Board application, the service is configured using “NodePort” and port “30002” meaning the application should be a up and running and accessible by pointing your web browser to any Kubernetes node IP address and port “30002”. For example: `http://<k8s-worker1>:30002`



## CHAPTER 6

# Troubleshooting

- [Troubleshooting](#), on page 25
- [ImagePullBackOff Status Errors when Deploying HXCSI Pods](#), on page 25
- [Volume Delete Fails](#), on page 26
- [Application Pod in ContainerCreating State During Node Deletion](#), on page 26
- [Deleted Pod Scheduled Back on the Same Node](#), on page 26

## Troubleshooting

The following section highlights common issues seen when installing and using the HyperFlex CSI integration. The information provided includes symptoms to help diagnose the issue as well as a solution to resolve the issue.

## ImagePullBackOff Status Errors when Deploying HXCSI Pods

- **Symptom 1:** Running the command “`kubectl get pods [-n <namespace>]`” shows that the HXCSI pods are showing a status of “ImagePullBackOff”.
- **Symptom 2:** Running the command “`kubectl describe pod <csi-pod_name>`” shows a message containing the following error: “Error: ErrImagePull” and “Back-off pulling image...”

### Solution:

- **Solution 1:** Ensure the HXCSI container image name provided to the `hxcsi-setup` script is correct
- **Solution 2:** Ensure the HXCSI container image exists, either directly within docker on each Kubernetes worker node or on the local container image registry.
- **Solution 3:** Ensure that the “`imagePullPolicy`” lines in the following YAML files generated by the `hxcsi-setup` script are set to “`IfNotPresent`”: `csi-attacher-hxcsi.yaml`, `csi-nodeplugin-hxcsi.yaml`, `csi-provisioner-hxcsi.yaml`
- **Solution 4:** Ensure that the following images are loaded on the local container image registry on each Kubernetes node:  
`csi-attacher-3.0.2-cisco1.tar`, `csi-node-driver-registrar-2.0.1-cisco1.tar`,  
`csi-resizer-1.0.1-cisco1.tar`, `csi-provisioner-2.0.4-cisco1.tar`

## Volume Delete Fails

Volume delete fails due to stale volume attachments present even after NodeUnpublish succeeds and the volume is unmounted. This happens when the delete volumeattachment kubernetes api is lost during etcd leader election. Even after nodeUnpublish completes and the volume is successfully unmounted from the node, delete volume fails.

The log in the external-provisioner appears as:

```
volume deletion failed: persistentvolume <pv-name> is still attached to node <node-name>.
```

The log in the external-attacher appears as:

```
<Volume-attachment> is already attached.
```

### Solution:

Delete the stale volumeattachments using the following command:

```
kubectl delete volumeattachments <VA-name>
```

After few seconds when the provisioner retries, pv is deleted.

Alternatively, you can delete manually using the following command:

```
kubectl delete pv <pv-name>
```

## Application Pod in ContainerCreating State During Node Deletion

Application pods that get stuck in a ContainerCreating state or multi attach error state and are not able to mount the volumes. This may happen during the remove or delete a k8s worker node from the cluster and pods migrate to a new worker node.

The recommended way to delete a K8s worker node is by using the following commands:

```
kubectl drain <node-name>  
kubectl delete node <node-name>
```

For more information, see [How to gracefully remove a node from Kubernetes?](#).

## Deleted Pod Scheduled Back on the Same Node

A pod in **Running** state may get recreated after deleting using the `kubectl delete pod` command and then get stuck in **Terminating** state on deleting its namespace.

Instead of using the `kubectl delete pod` command on any running pod, the following best practices way is recommended:

1. Note the node name on which pod to be deleted is running on.

```
kubectl get pods -o wide --all-namespaces
```

2. Cordon off the node on which the pod is running.



```
kubectl cordon <node-name>
```

**3. Delete the pod.**

```
kubectl delete pod <pod-name>
```

**4. Check that the deleted pod is scheduled on a different node using the following command:**

```
kubectl get pods -o wide --all-namespaces
```

**5. Uncordon the cordoned node.**

```
kubectl uncordon <node-name>
```

