

Cisco and Hitachi Adaptive Solutions CI with Red Hat OpenShift Platform

Implementation Guide

© 2022 Hitachi Vantara LLC. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including copying and recording, or stored in a database or retrieval system for commercial purposes without the express written permission of Hitachi, Ltd., or Hitachi Vantara LLC (collectively "Hitachi"). Licensee may make copies of the Materials provided that any such copy is: (i) created as an essential step in utilization of the Software as licensed and is used in no other manner; or (ii) used for archival purposes. Licensee may not make any other copies of the Materials. "Materials" mean text, data, photographs, graphics, audio, video and documents.

Hitachi reserves the right to make changes to this Material at any time without notice and assumes no responsibility for its use. The Materials contain the most current information available at the time of publication.

Some of the features described in the Materials might not be currently available. Refer to the most recent product announcement for information about feature and product availability, or contact Hitachi Vantara LLC at https://support.hitachivantara.com/en_us/contact-us.html.

Notice: Hitachi products and services can be ordered only under the terms and conditions of the applicable Hitachi agreements. The use of Hitachi products is governed by the terms of your agreements with Hitachi Vantara LLC.

By using this software, you agree that you are responsible for:

1. Acquiring the relevant consents as may be required under local privacy laws or otherwise from authorized employees and other individuals; and
2. Verifying that your data continues to be held, retrieved, deleted, or otherwise processed in accordance with relevant laws.

Notice on Export Controls. The technical data and technology inherent in this Document may be subject to U.S. export control laws, including the U.S. Export Administration Act and its associated regulations, and may be subject to export or import regulations in other countries. Reader agrees to comply strictly with all such regulations and acknowledges that Reader has the responsibility to obtain licenses to export, re-export, or import the Document and any Compliant Products.

Hitachi and Lumada are trademarks or registered trademarks of Hitachi, Ltd., in the United States and other countries.

AIX, AS/400e, DB2, Domino, DS6000, DS8000, Enterprise Storage Server, eServer, FICON, FlashCopy, GDPS, HyperSwap, IBM, Lotus, MVS, OS/390, PowerHA, PowerPC, RS/6000, S/390, System z9, System z10, Tivoli, z/OS, z9, z10, z13, z14, z/VM, and z/VSE are registered trademarks or trademarks of International Business Machines Corporation.

Active Directory, ActiveX, Bing, Edge, Excel, Hyper-V, Internet Explorer, the Internet Explorer logo, Microsoft, the Microsoft corporate logo, the Microsoft Edge logo, MS-DOS, Outlook, PowerPoint, SharePoint, Silverlight, SmartScreen, SQL Server, Visual Basic, Visual C++, Visual Studio, Windows, the Windows logo, Windows Azure, Windows PowerShell, Windows Server, the Windows start button, and Windows Vista are registered trademarks or trademarks of Microsoft Corporation. Microsoft product screen shots are reprinted with permission from Microsoft Corporation.

All other trademarks, service marks, and company names in this document or website are properties of their respective owners.

Copyright and license information for third-party and open source software used in Hitachi Vantara products can be found at <https://www.hitachivantara.com/en-us/company/legal.html>.

Feedback

Hitachi Vantara welcomes your feedback. Please share your thoughts by sending an email message to SolutionLab@HitachiVantara.com. To assist the routing of this message, use the paper number in the subject and the title of this white paper in the text.

Revision history

Revision	Changes	Date
MK-SL-215-03	Included Cisco Intersight capabilities information	March 29, 2022
MK-SL-215-02	Updated with Red Hat OCP StatefulSets and failover information	March 16, 2021
MK-SL-215-01	Support for Hitachi Storage Plug-in for Containers (HSPC) on top of Red Hat OpenShift Container Platform (OCP)	February 4, 2021

Implementation Guide

Utilize Hitachi Vantara's Virtual Storage Platform (VSP) integrations with persistent storage to provide your container environments agile deployment speed for an increase in operational efficiencies and to further business outcomes. Hitachi's proven leadership and joint innovations have accelerated enterprise IT initiatives for 80% of global Fortune 100 companies. Cisco and Hitachi Adaptive Solutions for Converged Infrastructure (CI) is a pre-validated, tested, and rapidly deployable reference architecture. It is an agile data-driven foundation that supports a broad range of technologies and workloads and, when combined with continuous innovation, positions your organization to deliver better experiences and tap into new revenue streams on the same adaptable infrastructure solution provided by Hitachi and Cisco Systems.

A key element in the successful deployment of a container platform is having a robust and flexible infrastructure that can meet the wide variety of requirements in a highly dynamic environment. The Cisco and Hitachi Adaptive Solution for CI with Red Hat OpenShift Container Platform (OCP) provides highly available, predictable, and high-performance infrastructures for container applications.

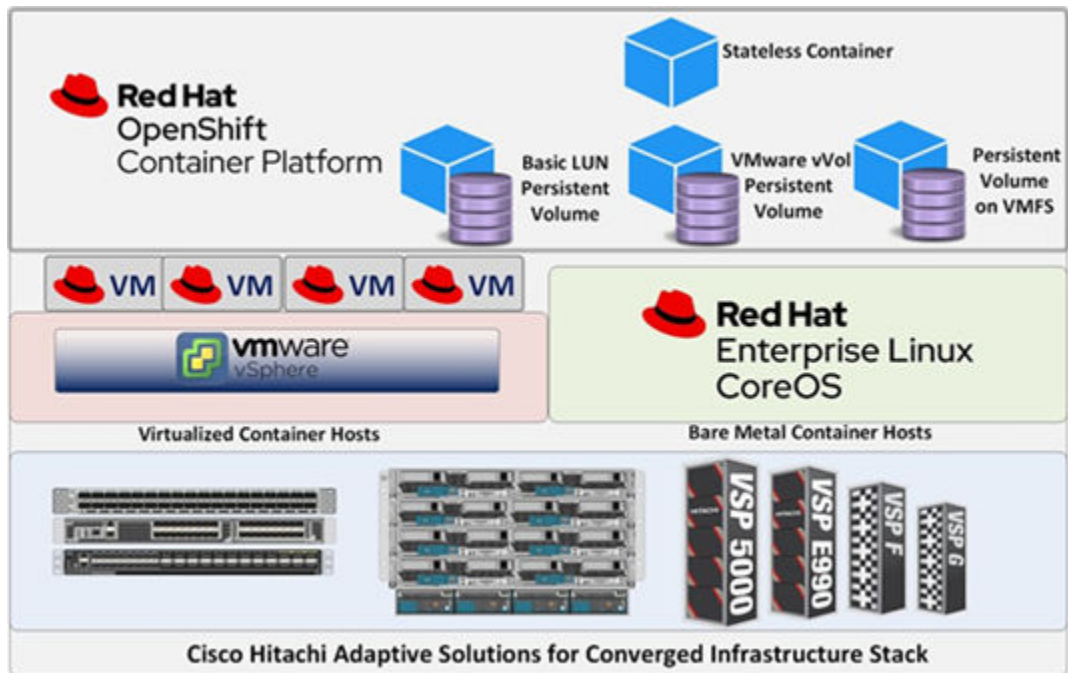
The solution feature set covered in this document is comprised of Red Hat OCP 4.6 running on VMware vSphere 6.7u3. Enterprise storage is provided by Hitachi Storage Provider for VMware vCenter to provide persistent storage for containerized applications through Storage Based Policy Management (SPBM) backed by either Virtual Volumes (vVols) or Virtual Machine File System (VMFS) as well as Hitachi Storage Plug-in for Containers (HSPC) which can provide persistent storage to Bare Metal hosts via the Fibre Channel protocol. This document covers implementation and best practices of Hitachi storage resources to provide persistent storage to hybrid OpenShift environments that are backed by virtual and Bare Metal workers within a hybrid deployment using Cloud Native Storage (CNS) and HSPC.



Note: CNS vVol integration is officially supported on vSphere 7.0 and later with CSI spec 2.x.

Installation of Red Hat OCP is not covered in this document and can be found in [Related Documents \(on page 93\)](#).

The following figure provides a capability overview of Red Hat OCP backed by Hitachi VSP storage on top of Cisco and Hitachi Adaptive Solutions for CI.



Use the procedures in this guide to enable Hitachi Vantara storage integrations with Red Hat OCP in conjunction with VMware on top of Cisco Virtual System Infrastructure. Other container management platforms as well as deployment types other than Red Hat OCP in hybrid configurations will be mentioned but not covered.



Note: Testing of these procedures was in a lab environment. Many things impact production environments beyond prediction or duplication in a lab environment. Follow the recommended practice of conducting proof-of-concept testing for acceptable results in a non-production, isolated test environment that otherwise matches your production environment before your production implementation of this solution.

For more information about validated solutions using Cisco Unified Compute System (UCS) and Hitachi Virtual Storage Platform (VSP), see [Related Documents \(on page 93\)](#).

This document is intended for the following:

- Storage administrators
- VMware administrators
- Kubernetes administrators
- Sales engineers
- Field consultants
- Professional services staff
- Validated Hitachi and Cisco resale partners

Readers of this document should have a background in or understanding of the following:

- RAID systems and their functions
- VMware ESXi and VMware vCenter environments

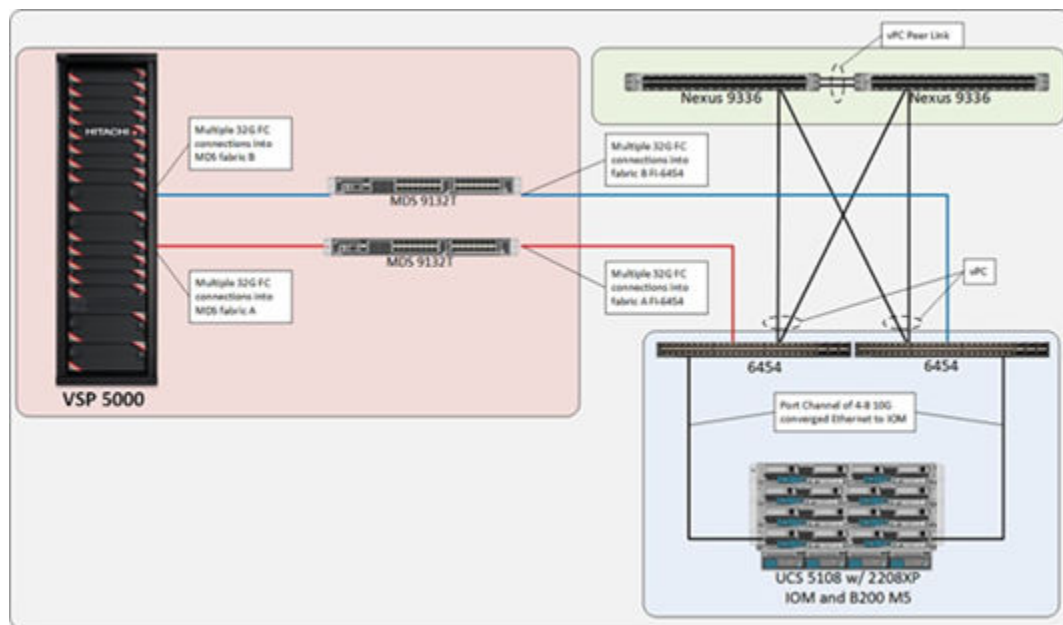
- Converged infrastructures
- Kubernetes

Cisco Unified Compute System environment

Cisco and Hitachi Adaptive Solutions for Converged Infrastructure as a virtual server infrastructure is a best-practice datacenter architecture built on collaboration between Hitachi Vantara and Cisco Systems to meet your enterprise needs using virtual server workloads.

This architecture uses a Hitachi Virtual Storage Platform (VSP) connected to Cisco MDS multilayer switches that link to the Cisco UCS Fabric Interconnects and Cisco Unified Computing System (UCS) chassis. Northbound networking is enabled through the Cisco Nexus 9000 family of switches.

The following figure shows the validated architecture for Cisco and Hitachi Adaptive Solutions for Converged Infrastructure. Red lines represent Fabric A connections, the blue lines represent Fabric B connections, and the rest are port channel connections.



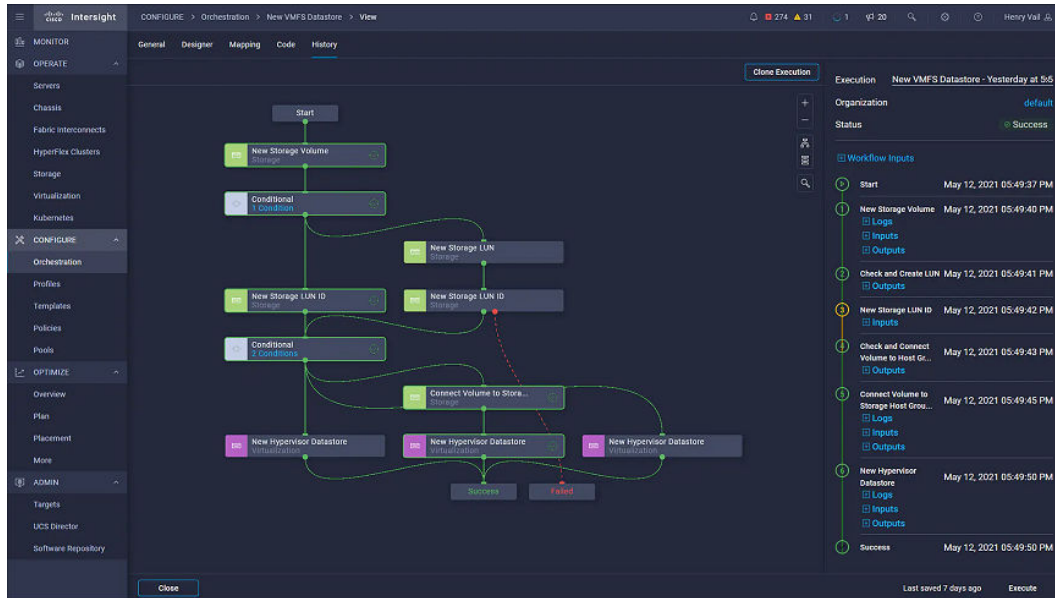
See [Cisco and Hitachi Adaptive Solutions for Converged Infrastructure](#) for more information about the the Cisco Validated Design (CVD) program.

Cisco Intersight Capabilities with Hitachi Virtual Storage Platform

Hitachi has enabled a magnitude of storage management capabilities that will now be able to be done using Cisco Intersight with the goal of saving administrators time and frustration.

Within the Cisco Intersight management platform, administrators can use the concept of tasks and workflows to easily manage their hybrid IT environments.

Tasks are essentially a library of functions that leverage API invoke calls that can be customized, or they can be provided by Cisco out of the box. These tasks can be compiled to create workflows to enable quick and easy automation of infrastructure without requiring code experts. This provides true single pane of glass orchestration through Cisco Intersight, reducing the need for datacenter administrators to host multiple screens to complete functions.



The following tables show the current capabilities of Hitachi Virtual Storage Platform (VSP) in orchestration with tasks and workflows provided by Intersight to end users.

Table 1 List of support tasks for Hitachi VSP

Tasks	Hitachi VSP
Compress Storage Pool	Y
Connect Initiators to Storage Host	Y
Connect Volume to Storage Host	Y
Copy Storage Volume	Y
Disconnect Initiators from Storage Host	Y
Disconnect Volume from Storage Host	Y
Edit Storage Pool	Y
Expand Storage Volume	Y
Expand Storage Pool	Y
Format Storage Volume	Y
New Storage Host	Y
New Storage Pool	Y
New Storage Volume	Y
Remove Storage Host	Y
Remove Storage Pool	Y
Remove Storage Volume	Y

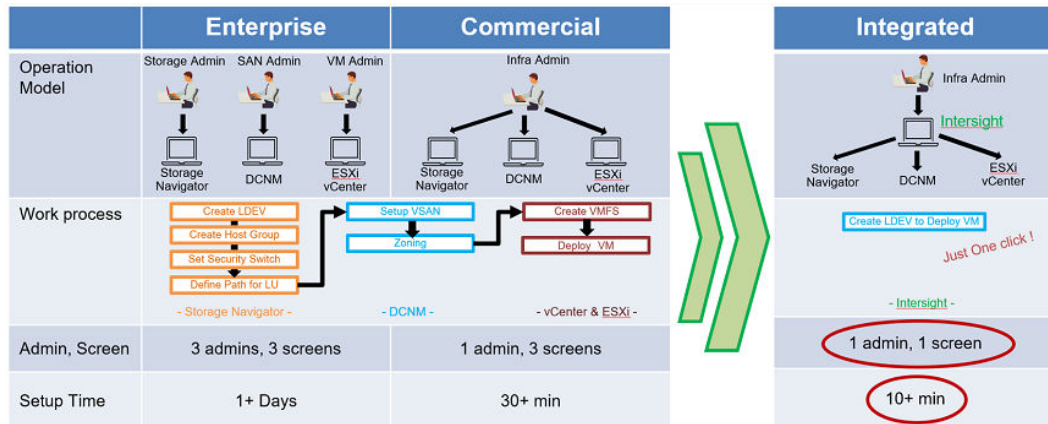
Table 2 List of supported workflows for Hitachi VSP

Storage Workflows	Hitachi VSP
New Storage Host	Y
New VMFS Datastore	Y
Remove Storage Host	Y
Update Storage Host	Y

With these capabilities administrators can complete a majority of day 0 to day N tasks to support their hybrid IT environment with Hitachi VSP storage systems.

A related reference architecture has been published as part of Cisco and Hitachi Adaptive solutions, at [Hitachi Virtual Storage Platform with Cisco Intersight Reference Architecture Guide](#). See [Related Documents \(on page 93\)](#) for more information.

This reference architecture explores the benefits of integrated management with Cisco Intersight compared to conventional methods using multiple management interfaces. When creating a virtual environment for enterprise workloads with Cisco Intersight with VSP integration, on average 50 hours of time is saved over the course of a year and 80% fewer screens are required to complete such operations, as shown in the following figure.



Hardware versions

The following table lists the hardware used to develop these procedures. Alterations can be made according to Hitachi and Cisco hardware compatibility lists.

Table 3 Hardware versions used for validation

Component	Version
Hitachi Virtual Storage Platform VSP 5000	90-03-02-00/00
Cisco MDS 9132T Fibre Channel switch	8.4(1a)
Cisco Nexus 9332-FX2 switch	NXOS 7.0(3)I7(8)
Cisco Fabric Interconnect 6454	4.1(2a)
Cisco Unified Computing System B200 M5 Blade Servers	4.1(2a)
Cisco Unified Computing System 2208XP IOM	4.1(2a)

Software versions

This section lists the software used to develop these procedures.

Table 4 Software versions used for validation

Component	Version
Hitachi Storage Provider for VMware vCenter	3.5.8 or newer

Component	Version
Hitachi Storage Plug-in for Containers	3.5 or newer
vSphere Container Storage Interface (CSI)	2.0 or newer
VMware ESXi 6.7 U3	6.7.0.14320388 or newer
VMware vCenter Standalone (VCSA) 6.7 U3	6.7.0.41000 or newer
VMware ESXi 6.7U3 nenic	1.0.33.0 or newer
VMware ESXi 6.7U3 nfnic	4.0.0.56 or newer
Red Hat OpenShift Platform (OCP)	4.6 or newer
Red Hat Core OS (RHCOS)	4.4.17 or newer
Red Hat Enterprise Linux (RHEL)	7.8
Red Hat Enterprise Linux enic	4.0.0.11
Red Hat Enterprise Linux fnic	2.0.0.63

Hitachi Vantara with container management platforms

This section covers the deployment types of container management platforms and their support by Hitachi storage to provide persistent storage. Within this guide, Red Hat OCP in an hybrid deployment is covered with Cloud Native Storage (CNS) backed by Container Storage Interface (CSI) as well as HSPC for Bare Metal nodes.



Note: CNS vVol integration is officially supported on vSphere 7.0 and later with CSI spec 2.x.

The following table shows Hitachi integration points with Red Hat OCP for various deployment types. A hybrid deployment type backed by CNS storage and HSPC via the Fibre Channel protocol is covered in this guide.

Table 5 Hitachi integration points with Red Hat openShift Container Platform

Red Hat OpenShift Container Platform		
Deployment Type	Storage Type	Hitachi Persistent Storage Provider Compatibility
VM (all virtual infrastructure)	iSCSI	Hitachi Storage Plug-in for Containers
	Cloud Native Storage (CNS)	Container Storage Interface (vVol + VMFS)

Red Hat OpenShift Container Platform		
Deployment Type	Storage Type	Hitachi Persistent Storage Provider Compatibility
Bare Metal (all physical infrastructure)	Fibre Channel (FC)	Hitachi Storage Plug-in for Containers
Hybrid (mix of physical and virtual workers)	iSCSI	Hitachi Storage Plug-in for Containers (Virtual Only)
	Fibre Channel (FC)	Hitachi Storage Plug-in for Containers(Bare Metal Only)
	Cloud Native Storage (CNS)	Container Storage Interface(vVol + VMFS)

The following components are used to implement the feature sets that are covered in this document with Red Hat OCP and Hitachi VSP storage.

Red Hat OpenShift container platform

Red Hat OCP is a container management platform for Kubernetes which is developed by Red Hat. OCP provides the control plane and data plane within the same interface. OCP provides administrator views to deploy operators, monitor container resources, manage container health, manage users, work with operators, manage pods and deployment configurations, as well as define storage resources. OCP also provides a developer view to allow users to deploy application resources from various pre-defined resources such as YAML files, Docker files, Catalogs, or GIT within user- defined namespaces. With OCP 'kubectl', a native binary of Kubernetes is complemented by the `oc` command which provides further support for OCP resources, such as deployment and build configurations, routes, image streams, and tags. OCP provides a GUI and a CLI interface.

Pods

One of the main purposes of Kubernetes is to orchestrate the life cycle of a container. A pod is a collection of containers and it is stored inside a node of a Kubernetes cluster that users can administer. With pods it is possible to have multiple containers inside of a single pod.

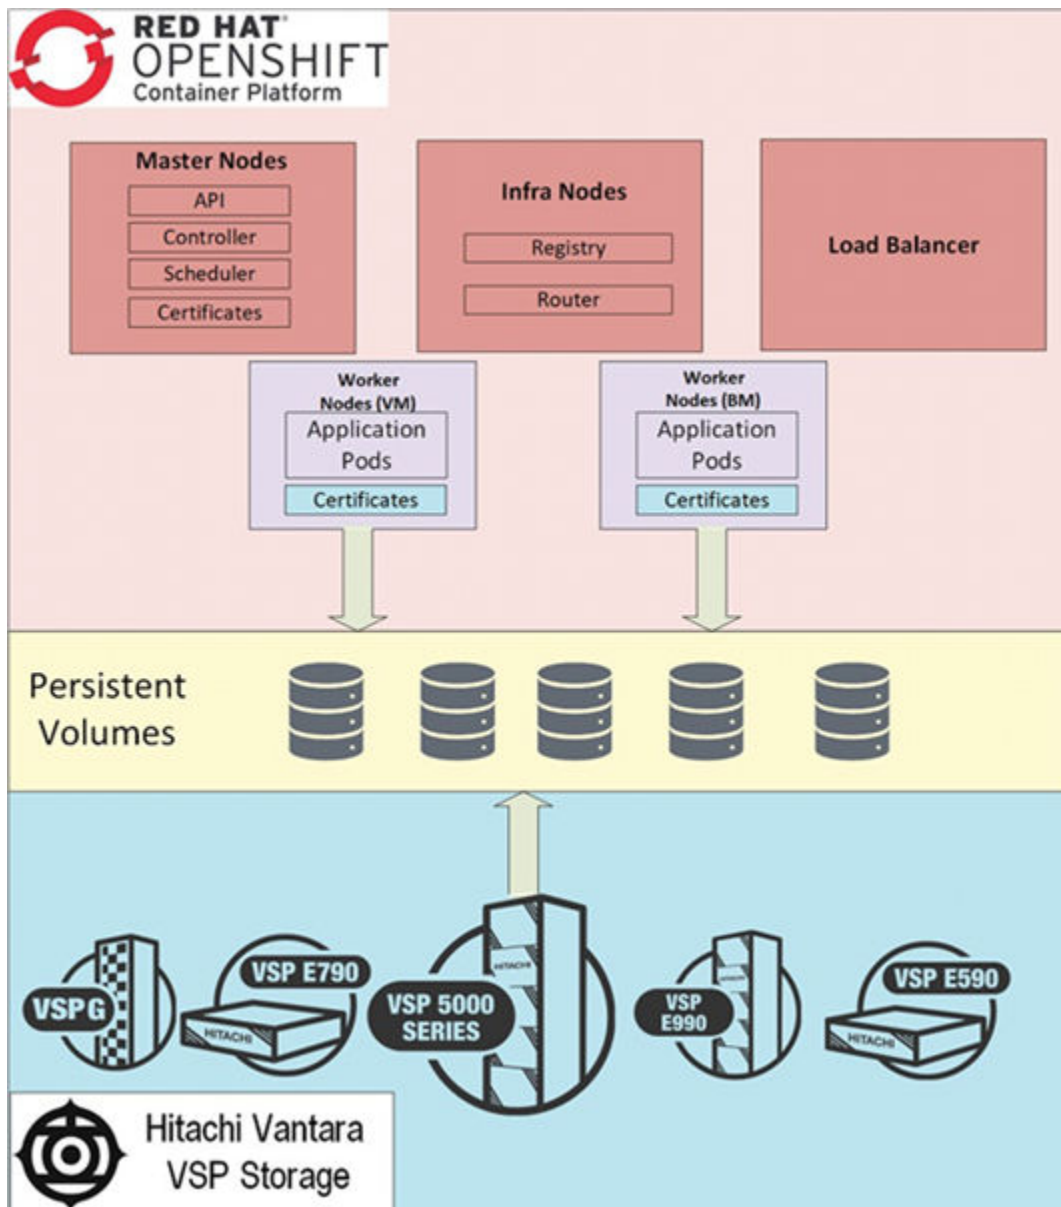


Note: Existing users of 'kubectl' can continue to use the binary with no changes to the API but should consider upgrading to `oc` in order to gain additional functionality.

Persistent storage

One of the storage resources that the OCP platform orchestrates is persistent storage via Persistent Volumes (PV). PVs are resources in the Kubernetes cluster that have a lifecycle independent of any pod that uses a PV. This is a type of volume on the host machine that stores persistent data. PVs provide storage resources in a cluster, which allows the storage resource to persist even when the pods that use them are removed. PVs can be statically or dynamically provisioned, and they can be customized for performance, size, and access mode. PVs can be attached to pods via a PersistentVolumeClaim (PVC), that is, a request for the resource that acts as a claim to check for available resources.

The following figure shows a general Red Hat OCP deployment using Hitachi VSP storage for enterprise persistent storage.



Red Hat OCP components

Within an OCP cluster there are multiple node types and roles that provide functionality to the container management platform. This section provides an overview of each node type and specific node roles within the cluster. Not all components are covered and users are encouraged to see the Red Hat documentation which can be found in [Related Documents \(on page 93\)](#).

Master Nodes

Master nodes maintain the OCP cluster configuration as well as manage nodes within the cluster and schedule pods to run on worker nodes. Master nodes consist of an API server, controller manager server, certificate, and scheduler. In the case of a master node outage, container applications will not be impacted and end users can continue using resources, but administrators of the cluster will not be able to make any changes to the cluster.

Worker Nodes

Worker nodes provide a runtime environment for containers and pods and are managed by the master nodes within the cluster. Worker nodes can either be virtual or physical based on the deployment type.

API Server

The API server, or 'kube-apiserver', provides the front end for the Kubernetes control plane by managing the interactions of cluster components via RESTful API calls. Administrators have the ability to run several instances of kube-apiserver to balance traffic among the cluster.

Scheduler

The Scheduler, or 'kube-scheduler', ensures that container applications are scheduled to run on worker nodes within the OCP cluster. The scheduler reads data from the pod and finds a node that is a good fit based on configured policies.

Controller Manager

The controller manager, or 'kube-controller-manager', provides the cluster with the necessary state changes to provide the most applicable state for a healthy cluster. The controller manager provides this functionality via the 'kube-apiserver'.

Namespace

Namespaces are intended for use in environments with many users spread across multiple teams, or projects. Namespaces provide a scope for names. Names of resources need to be unique within a namespace, but not across namespaces. Namespaces cannot be nested inside one another and each Kubernetes resource can only be in one namespace.

Virtualized container hosts

In an all-virtualized deployment, deploy the Red Hat OCP master controller nodes and worker nodes on virtual machines running Red Hat Enterprise Linux CoreOS hosted on VMware vSphere hosts.

A virtualized OCP deployment is highly flexible and dynamic. Due to vSphere's ability to deploy virtual machines quickly, new virtual worker nodes can be added to the cluster quickly. No physical deployment in the data center is required to deploy a functional OCP node.

VMware vSphere can also help consolidate the OCP nodes, reducing the rack space footprint. VMware's CPU scheduling ability and memory overcommit can further consolidate nodes. VMware High Availability and VMware Distributed Resource Scheduler can also help protect the OCP cluster and enforce the SLA.

Deployment methods and types

For Red Hat OCP there are multiple deployment methods based on the hardware configuration supporting the environment. Within this guide a hybrid environment was used. For more information about deployment methods and types, refer to the [OCP](#) documentation.

The following are descriptions of OCP deployment types.

Bare Metal and Hybrid

The Bare Metal deployment type consists of physical servers while the hybrid type consists of virtual and physical servers that use the User Provisioned Infrastructure (UPI). Use the installation program to generate the assets that you require to provision the cluster infrastructure, create the cluster infrastructure, and deploy the cluster to the infrastructure that you provided. To use UPI, deploy and maintain the following:

- The underlying infrastructure for the control plane and compute machines that make up the cluster
- Load balancers
- Cluster networking, including the DNS records and required subnets
- Storage for the cluster infrastructure and applications
- The smallest footprint Red Hat OCP cluster requires is one bootstrap node, three master nodes, and a minimum of two worker nodes

Virtual

Virtual deployments use Installer Provisioned Infrastructure (IPI). With IPI, OCP can install itself, but deploying as well as managing the underlying infrastructure is needed, from compute to storage and network. For more information about IPI, see [Related Documents \(on page 93\)](#).

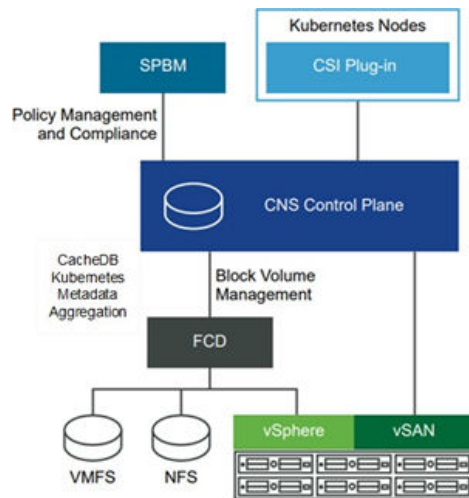
vSphere CNS

Cloud Native Storage (CNS) is an orchestration introduced in VMware vCenter 6.7u3 that provides storage data management for stateful applications. When you use CNS, you create containerized stateful applications capable of surviving restarts and outages. Stateful containers leverage storage exposed by vSphere backed by Hitachi VSP storage. With CNS, you can create persistent container volumes independent of virtual machines or containers. CNS uses several components to work with vSphere storage, this includes VMFS or vVols provided by the Hitachi Storage Provider for VMware vCenter. After you create PVs, you can review them and their backing virtual disks in the vSphere Client, and monitor their storage policy compliance.

Container Storage Interface

The Container Storage Interface (CSI) provides an industry standard interface for container orchestration to allow access to third-party storage systems. The CSI plug-in works with the CNS control plane to expose vSphere storage to containerized workloads running on Kubernetes. The CSI plug-in provides functionalities such as vSphere First Class Disk (FCD), also known as an Improved Virtual Disk (IVD), Kubernetes zones, and provisioning from multiple datastores. The CSI plug-in driver is shipped as a container image and must be deployed by the cluster administrator. For this validation, CSI v2.0 was used.

The following figure highlights the interaction between the CSI plug-in, CNS control plane, and vSphere storage.



Hitachi Storage Provider for VMware vCenter

Storage Provider for VMware vCenter allows VMware APIs for Storage Awareness (VASA) features to be used with Hitachi storage systems. Storage Provider for VMware vCenter allows policies to be made by making the storage attribute information available in vSphere. Hitachi Storage Provider makes this possible in two ways:

- VMware vSphere vVols

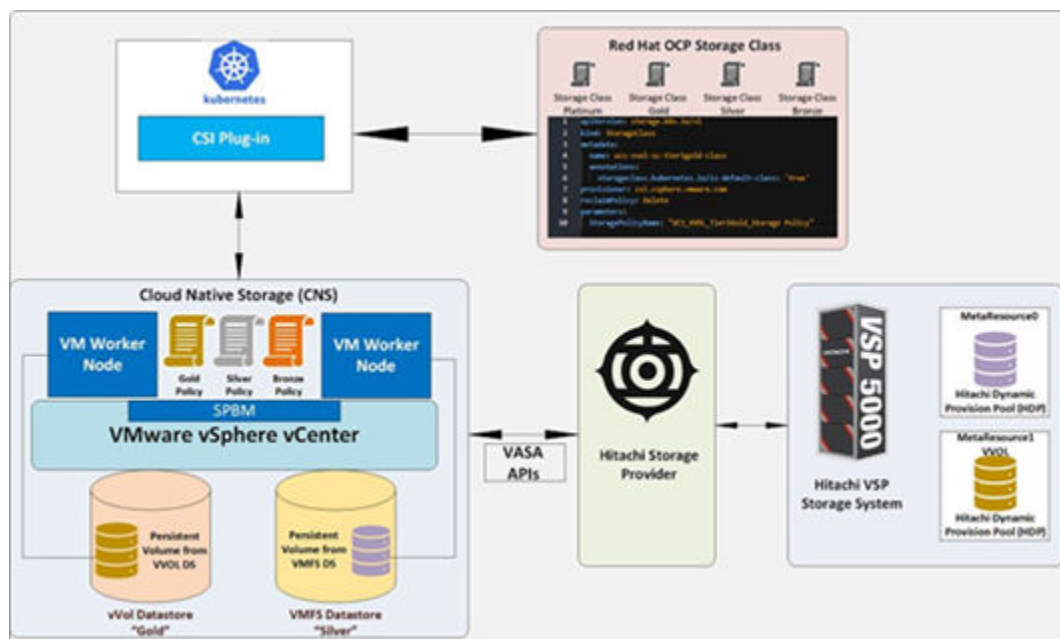
This function is the VASA Provider (VP) component of VMware vVols, which allows vVols to be used with supported Hitachi storage systems in a 1:1 mapping enabling greater insight into virtual machine performance.
- VMware VMFS

VASA allows storage capability information and alert notifications related to VMFS file systems to be generated automatically and displayed in vCenter Server

Note: You cannot register the same VSP in multiple storage providers for VMware vCenter within the same vCenter. Using different vCenters for each VP allows a storage system to be shared, a recommended best practice if sharing a VSP among two storage providers for VMware vCenter to create dedicated resource groups on the VSP.

- OCP allows these capabilities to be used through VM storage policies which are native to VMware vCenter. Administrators define StorageClass settings that point to their respective VM storage policies, backed by either Hitachi vVols or VMFS storage that use SPBM. Once these storage policies are defined and the PVCs are deployed, VASA, in conjunction with CNS, and the CSI specification provide the applicable PV based on the defined StorageClasses.

The following figure shows the relationship between VASA, CNS, CSI, and StorageClasses.

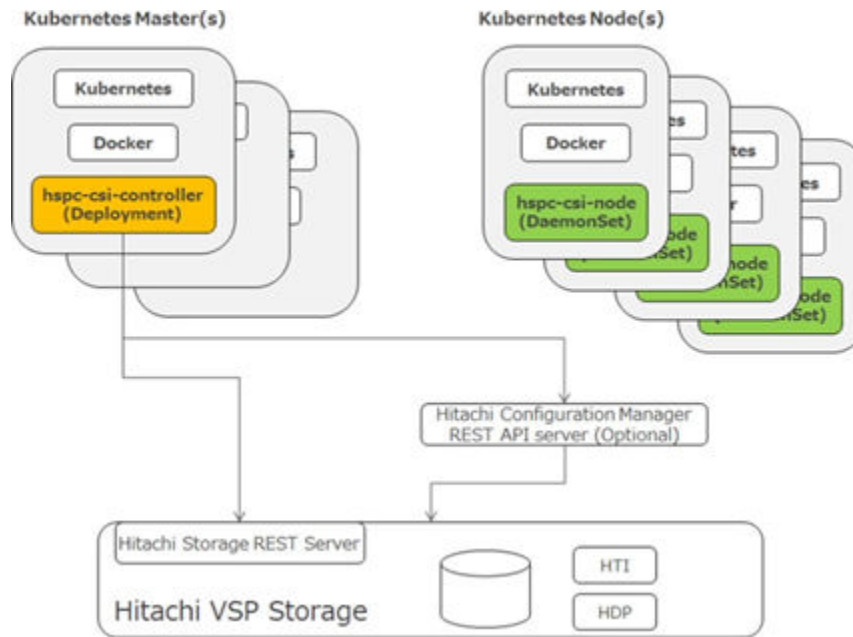


HSPC

HSPC is a software component that contains libraries, settings, and commands that you can use to create a container in order to run your stateful applications. It enables stateful applications to persist and maintain data after the lifecycle of the container has ended. Storage Plug-in for Containers provides persistent volumes from Hitachi Dynamic Provisioning (HDP) or Hitachi Thin Image (HTI) pools to Bare Metal or hybrid deployments via the Fibre Channel protocol. HSPC can also provide virtual environments for persistent storage if iSCSI is used.

HSPC uses built-in high availability to enable a Kubernetes/OCP master node to orchestrate storage tasks between hosts in a cluster as well as non-clustered environments.

The following figure shows a containerized environment backed by HSPC.



Red Hat OCP StatefulSets and Failover

PVs provided by Hitachi VSP storage systems are Read Write Once (RWO) to a single pod by design. During a disaster or failure other methods of data availability must be implemented by end users on applications utilizing persistent storage to guarantee data consistency. One method that can be used is to replicate the data across provisioned PVs that use block storage. For this, the native application must support replication and must be configured by the user. The following figure represents an example of an application utilizing the built-in replication function to guarantee data availability.

Use case 1: StatefulSet Pods with Block Storage – Product deployment



An intermediary cache between the application and storage must be implemented for applications that do not have a standard data replication utility need, for example Redis cache. With this method end users must ensure that when a Red Hat OCP StatefulSet fails, the following pod in line is designed to inherit values from the intermediary cache and provide 100% data availability guarantee. The following figure represents an example of an application using a Redis cache as the intermediary cache.

Use case 2: StatefulSet Pods with Block Storage – Custom Application Deployment**Solution implementation – Hitachi storage with VMware CNS**

The following describes how to implement Hitachi VSP storage on Cisco and Hitachi Adaptive Solutions for CI with VMware to back your Red Hat OCP virtual environment with persistent storage using VMware CNS. Red Hat OCP must be deployed prior to this implementation. Optionally, administrators can use native VMware-based adapters such as Hitachi Storage Plug-in for VMware vCenter to allocate VMFS datastores to virtual environments for more information on Storage Plug-in for VMware vCenter click [here](#). Additionally, Hitachi UCP Advisor is another native VMware orchestration tool that can streamline storage management with a native VMware interface. Operations such as pool creation, host group creation, LUN path allocation, VMFS datastore creation along with other native storage functions can be executed with UCP Advisor, for more information click [here](#).

Note: Read the release notes before installing and using any of these products. They contain requirements or restrictions that are not fully described in this document, updates, or corrections to this document. Release notes for Hitachi products are available on the [Hitachi Vantara Knowledge portal](#).

VSP storage configuration for VMFS datastore on VMware vCenter

This section covers how to configure and deploy Hitachi Virtual Storage Platform to support your virtual environment with VMFS datastores. These VMFS datastores can be used with Storage Provider for VMware vCenter in conjunction with CNS with VM storage classes to provide persistent storage for your OCP environment.

Note: The subset of directions within this guide use Hitachi Device Manager - Storage Navigator.

Create basic LDEVs from parity groups

Configuration steps in this section assume that parity groups and LDEVs have been configured on the Hitachi Virtual Storage Platform (VSP) as part of the solution build configured by a Hitachi partner or Hitachi professional services. If parity groups have not been configured on the Hitachi VSP, see the Hitachi Storage Virtualization Operating System RF (SVOS RF) documentation to create parity groups before continuing with this section.

Ensure that you have planned which parity groups and LDEVs to use for specific storage requirements. Your configuration might vary based on the types of drives ordered with your VSP storage system and the parity groups configured on them.

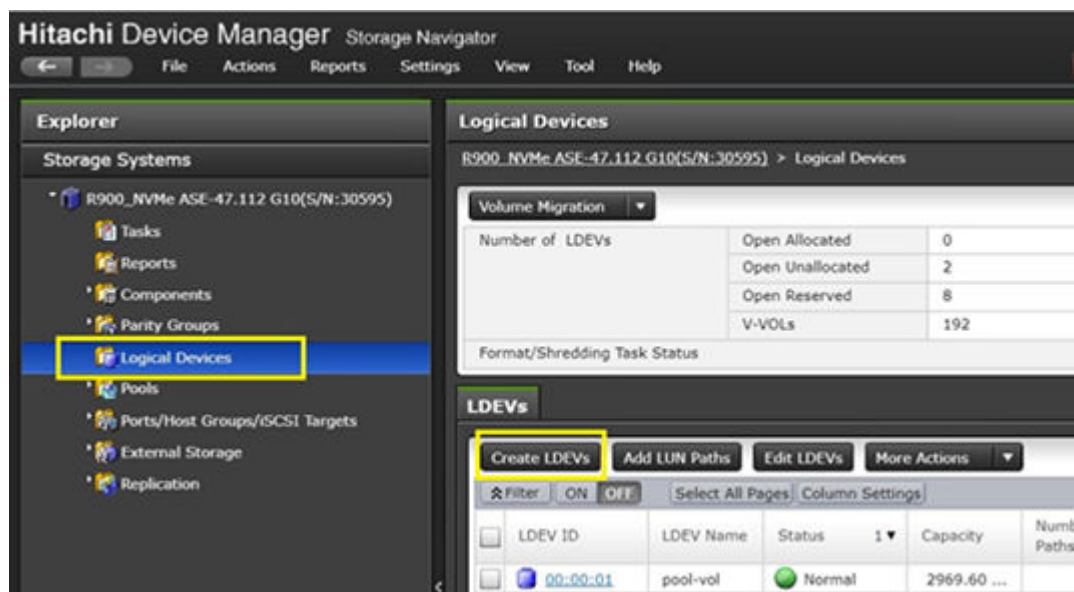
Use the following procedure to begin the provisioning process to create the basic LDEVs that will be used as pool volumes:

Procedure

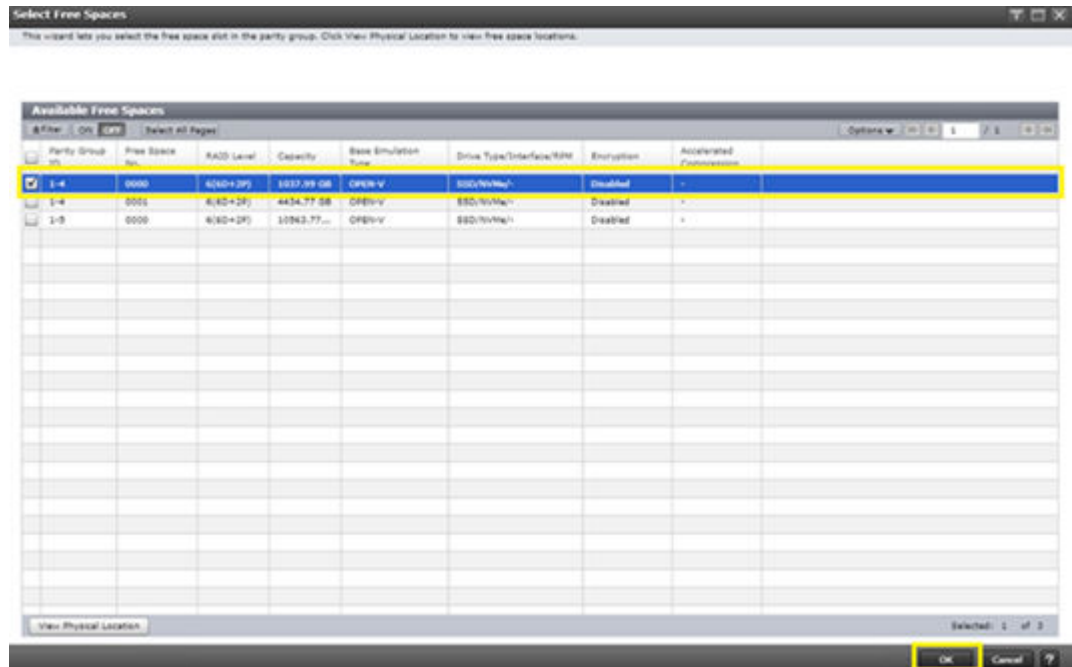
1. Log in to Hitachi Device Manager - Storage Navigator.



2. From the Explorer pane, select the **Storage Systems** tab.
3. Expand the storage system being configured, and then select **Logical Devices**.
4. Click **Create LDEVs**.



5. Configure the following items in the left pane of the Create LDEVs dialog:
6. Select Provisioning Type: **Basic**.
7. System Type: **Open**.
8. Emulation Type: **OPEN-V**.
9. Click **Select Free Spaces** to select available parities.
10. Select an available parity group, and then click **OK**.



11. Configure the following items in the left pane of the Create LDEVs dialog.
12. Define **LDEV Capacity** and select the unit size.
13. Define **Number of LDEVs per Free Space**.
14. Define **LDEV Name**, such as UCS_PoolVOL_VMFS, or UCS_PoolVOL_vVols.
15. Click **Add**.

Create LDEVs

1.Create LDEVs > 2.Confirm

This wizard lets you create and provision LDEVs enter the information for LDEVs you want to create, and then click Add. Click Options to view additional options, or click Finish to confirm the creation, or click Next if you want to add LUN paths for the LDEVs.

Provisioning Type: Basic

System Type: Open Mainframe

Emulation Type: OPEN-V

Parity Group Selection:

Drive Type/Interface/RPM: SSD/NVMe/-

RAID Level: 6(6D+2P)

Select Free Spaces

Total Selected Free Spaces: 1

Total Selected Free Space Capacity: 1.01 TB

LDEV Capacity: 500 GB (0.05-1037.99)

Number of LDEVs per Free Space: 1 (1-2)

LDEV Name: Prefix Initial Number

UCS_Pool/VOL_VMPFS

(Max. 32 characters total including max. 9-digit number, or blank)

Format Type: Quick Format

Options

Add

16. Click **Finish > Apply**.

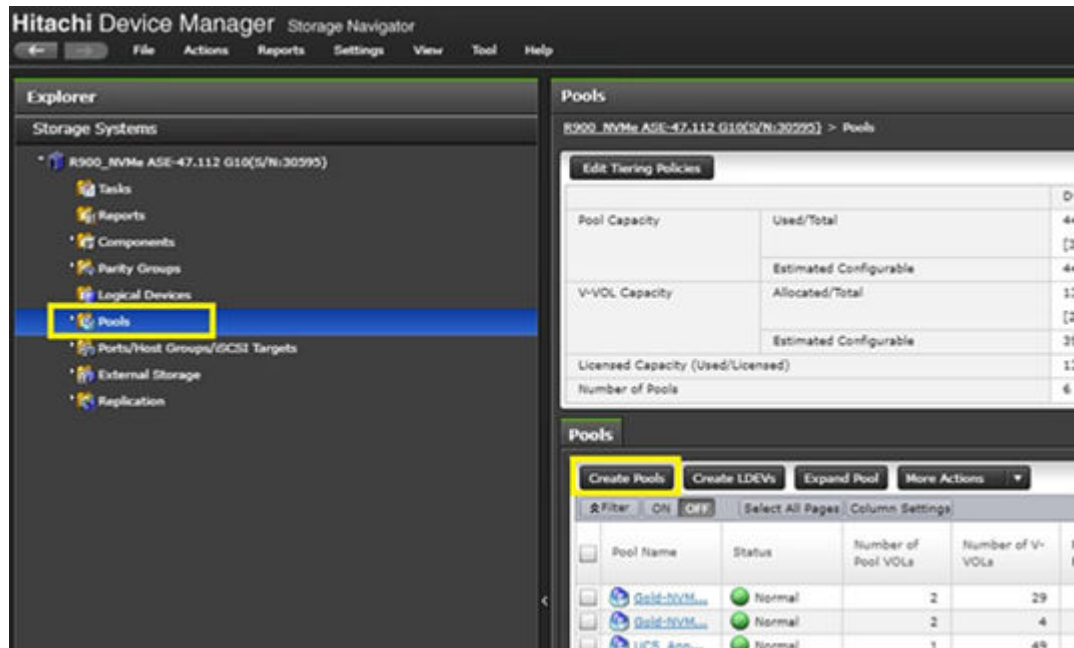
Create Hitachi Dynamic Provisioning (HDP) pools

Once you have created Basic LDEVs from available parity space, add those LDEVs to create an HDP pool.

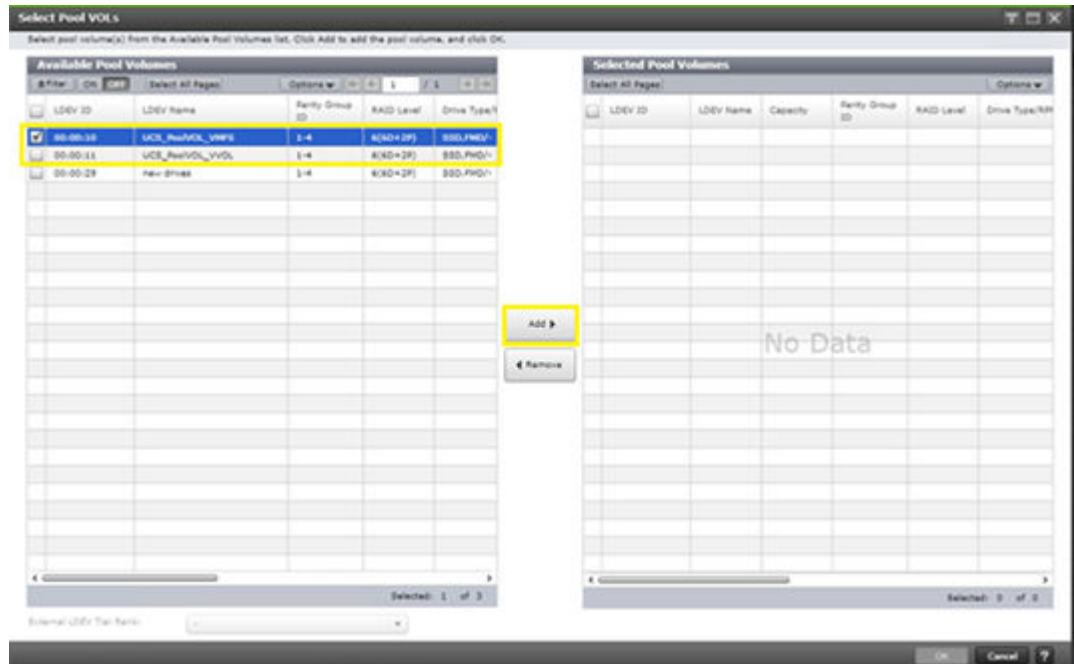
Use the following procedure to begin the provisioning process to create the HDP pool:

Procedure

1. From the Explorer pane, select the **Storage Systems** tab.
2. Expand the storage system being configured, and then select **Pools**.
3. Click **Create Pools**.



4. In the **Create LDEVs** dialog, configure the following items:
5. Pool Type: **Dynamic Provisioning**.
6. System Type: **Open**.
7. Multi-Tier Pool: **Disable**.
8. Pool Volume Selection: **Manual**.
9. Click **Select Pool VOLS**.
10. In the Select Pool VOLS wizard, select the applicable basic LDEV to support the HDP pool. Click **Add**.



11. Click **OK**.

- Enter the applicable Pool Name, and then click **Add**.

The screenshot shows the 'Create Pools' wizard with the following configuration:

- Pool Type:** Dynamic Provisioning
- System Type:** Open
- Multi-Tier Pools:** Disable
- Data Direct Mapping:** Disable
- Pool Volume Selection:** Manual
- Drive Type/RPM:** Mixable
- RAID Level:** Mixable
- Pool Name:** UCS_App_Pool (Max. 32 Characters)

The 'Add' button is highlighted in yellow, and the 'Pool Name' field is also highlighted in yellow.

- Click **Finish > Apply**.

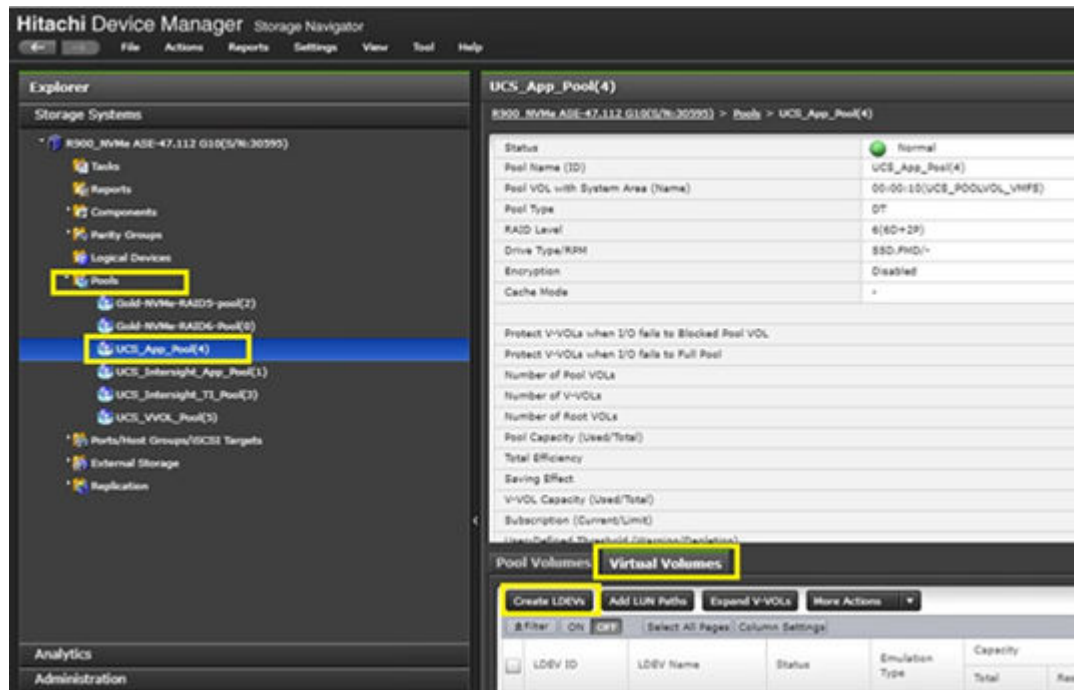
Creating LUNs to support VMFS datastores

Once applicable DP pools are created, create LUNs (Virtual Volumes) to present as VMFS datastores to vCenter.

Use the following procedure to create a LUN from a DP pool:

Procedure

- From the Explorer pane, select **Pools**.
- Select the applicable pool, and then click the **Virtual Volumes** tab.
- Click **Create LDEVs**.



4. Configure the following items in the left pane of the Create LDEVs dialog:
5. Provisioning Type: **Dynamic Provisioning**.
6. System Type: **Open**.
7. Emulation Type: **OPEN-V**.
8. Capacity Saving: Disabled.
9. Also define **LDEV Capacity, Number of LDEVs, and LDEV Name**.
10. Click **Add**.

Create LDEVs

1.Create LDEVs > 2.Confirm

This wizard lets you create and provision LDEVs enter the information for LDEVs you want to create, and then click Add. Click Finish to confirm the creation, or click Next if you want to add LUN paths for the LDEVs.

Provisioning Type:

System Type: Open Mainframe

Data Direct Mapping: Enable Disable

Emulation Type:

Capacity Saving:

Multi-Tier Pool: Enable Disable

Active Flash:

TSE Attribute: Enable Disable

Pool Selection:

Drive Type/RPM:

RAID Level:

Select Pool

Selected Pool Name(ID): UCS_App_Pool(4)

Selected Pool Capacity: 495.87 GB

LDEV Capacity:

(0.00-262144.00)

Number of LDEVs:

(1-63039)

LDEV Name: Prefix: Initial Number:

(Max. 32 characters total including max. 9-digit number, or blank)

Add

Next Tab

11. Click **Finish** > **Apply**.

Creating host groups

After the LUNs have been created from the available DP pools, they must be added to host groups. For host group creation see [Create Host Groups for UCS Server vHBAs on Each Fabric](#).

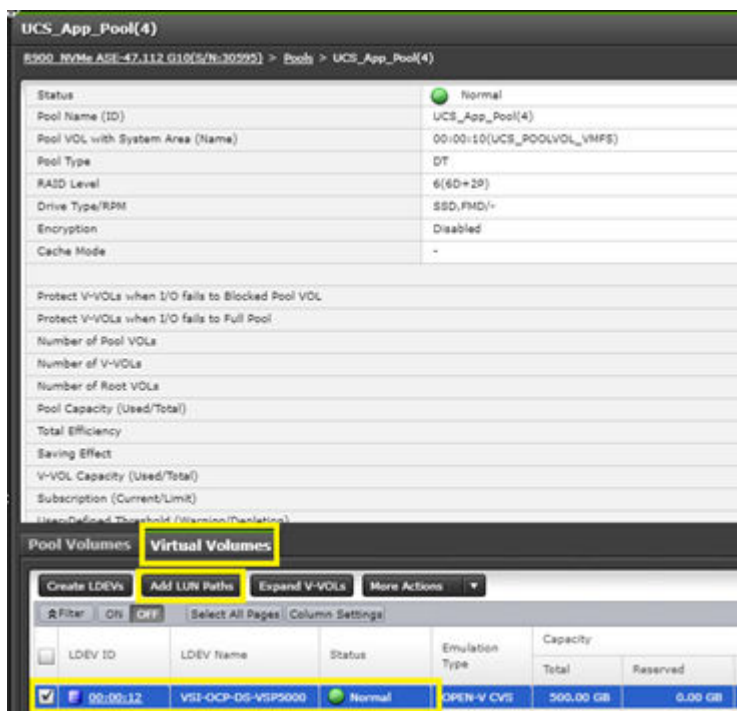
Adding LDEV paths to host groups

After the host groups have been defined for the underlying UCS infrastructure, add LDEV paths to the created LUNs so that you can onboard them as VMFS datastores on VMware vCenter.

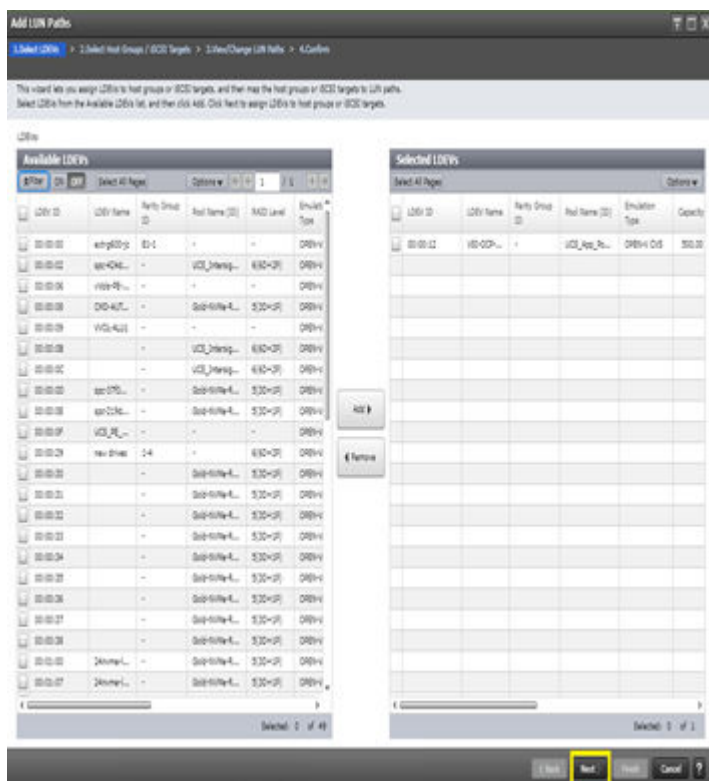
Use the following procedure to add LDEV paths:

Procedure

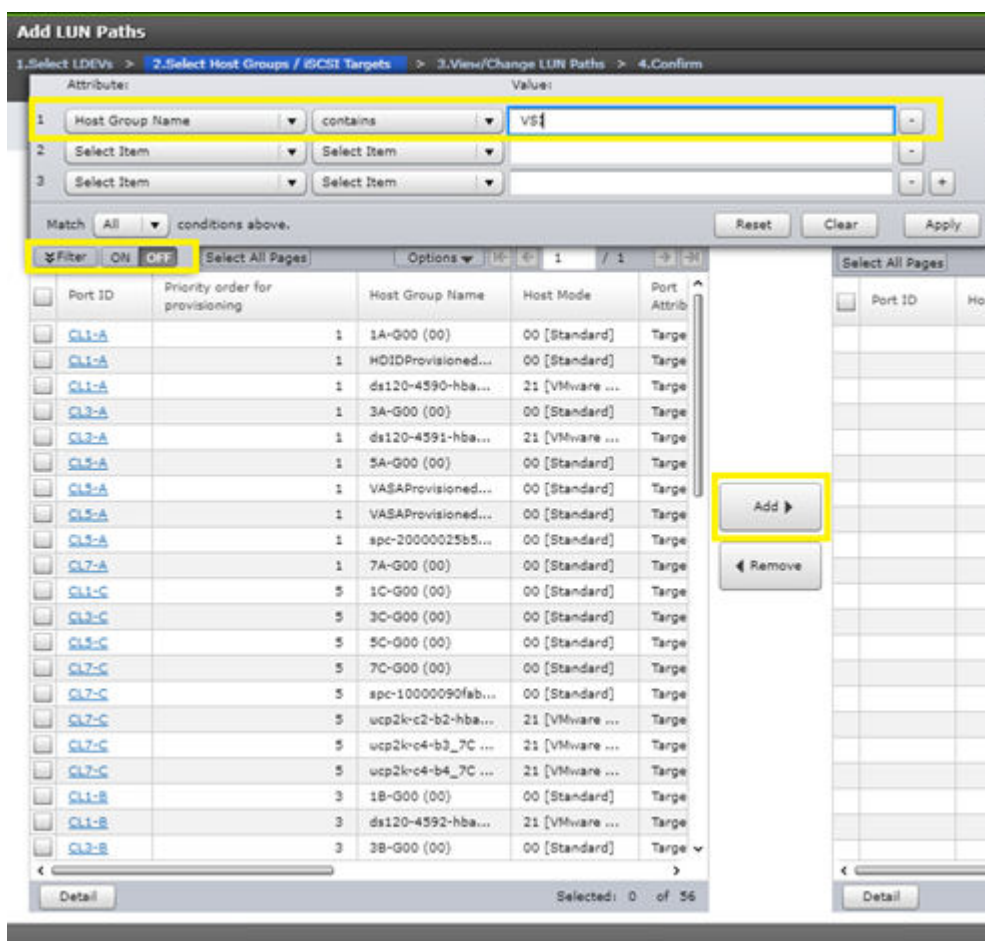
1. Select your LDEV, and then click **Add LUN Paths**.



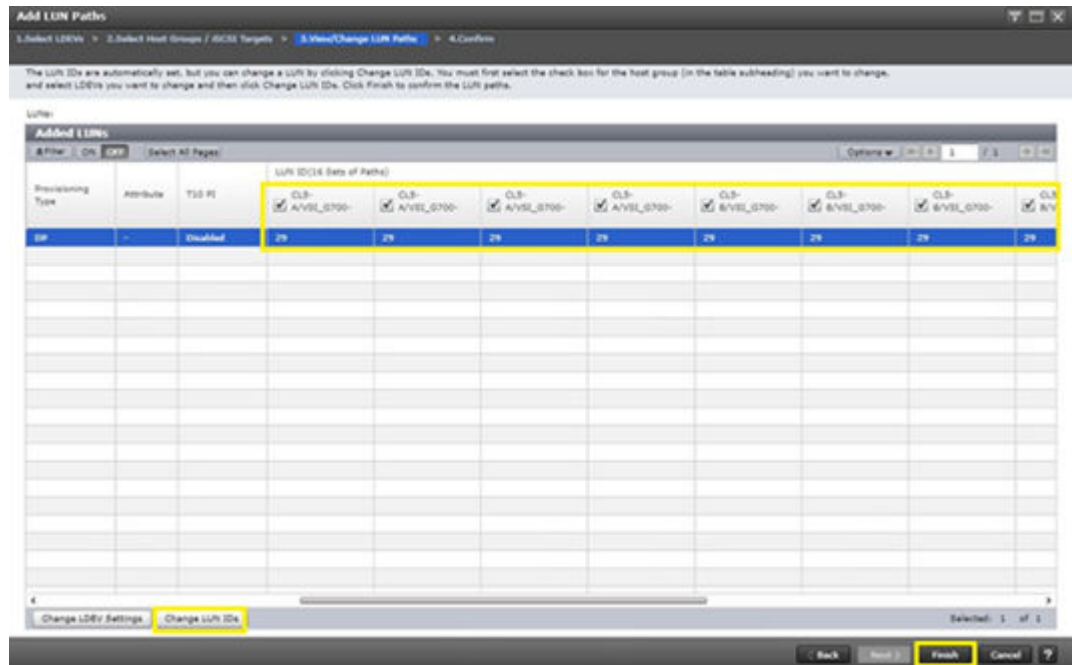
2. Using the Add LUN Paths wizard, click **Next**.



3. Using the Filter tool, search for VSI host groups based on **Host Group Name** and **Contains**.



4. Select the applicable host groups, and then click **Add**.
5. Click **Next**.
6. The View/Change LUN Paths screen shows the LDEV that you are adding the paths to and the associated host LUN ID that will be presented to the host on a per-path basis. Verify that the LUN alignment is correct before presenting the LUN to the respective VSI host groups by selecting the applicable LUN and then clicking **Change LUN IDs**. When done, click **Finish > Apply**.



7. After you have added LUN paths to your LDEV, you can continue with onboarding the VMFS datastore in VMware vCenter. See [Onboarding VMFS Datastores](#).

VSP vVols Configuration for vVols Datastore

This section covers how to configure and deploy Hitachi VSP storage to support your virtual environment with vVols datastores. vVols datastores can be used with Storage Provider for VMware vCenter in conjunction with CNS and VM storage policies to provide your OCP environment with persistent storage.

These are the prerequisites for VSP vVols configuration:

- Create Basic LDEVs from Parity Groups
- Create Hitachi Dynamic Provisioning (HDP) Pools
- Create a vVols Resource Group

Create a vVols resource group

Resource group configuration must be completed before configuring the Storage Provider for VMware vCenter. The Storage Provider for VMware vCenter uses the defined VSP resource group to provide VMware vCenter vVols storage via the VASA APIs.

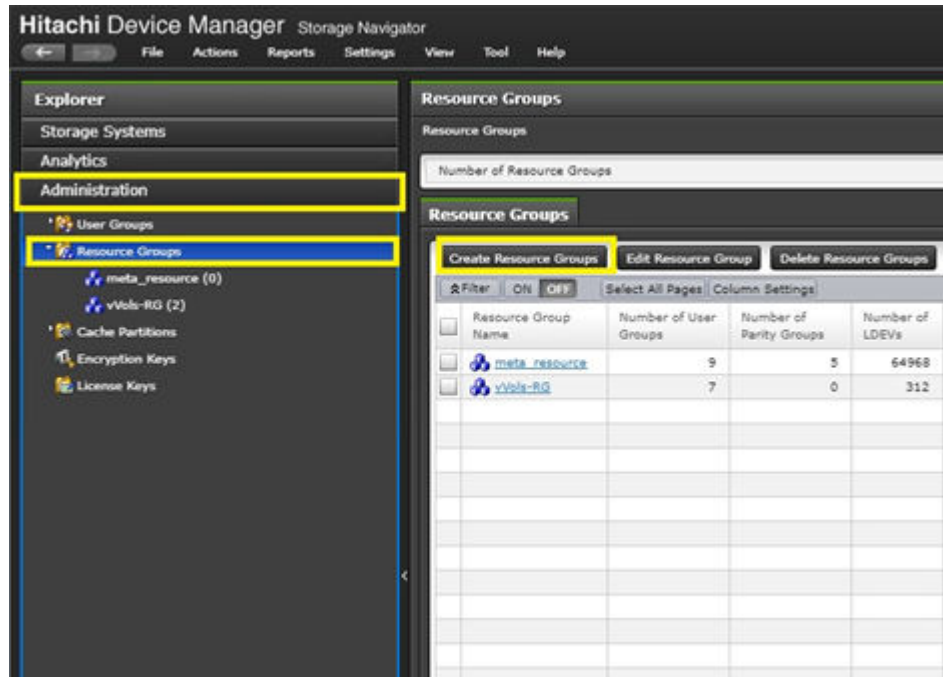
Note: If administrators will use compression and deduplication in conjunction with vVols, verify that there are enough LDEV IDs within the resource group to support the feature set. For more information, see [Related Documents \(on page 93\)](#).

Use the following procedure to configure VSP resource groups for vVols:


Procedure

1. From the Explorer pane, select **Administration > Resource Groups**.

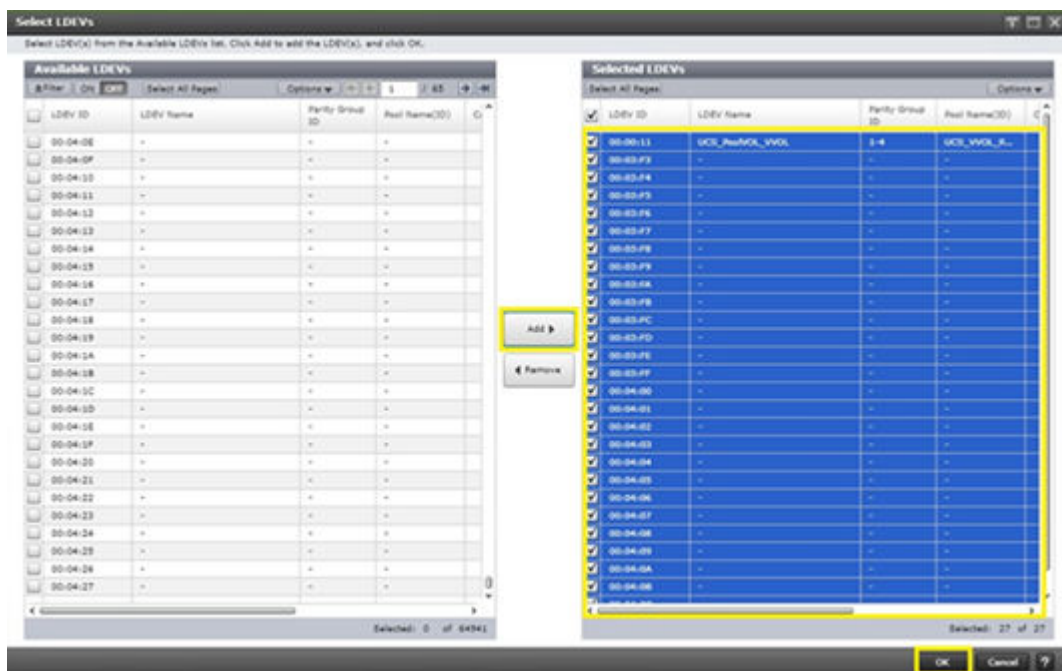
2. Click **Create Resource Groups**.



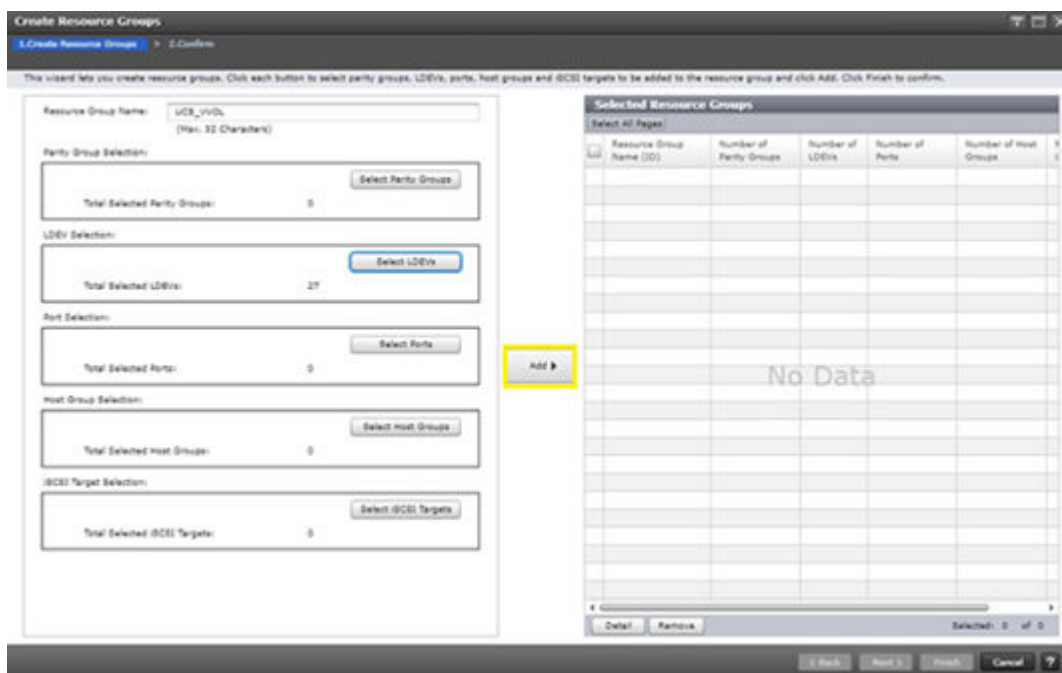
3. In the Create Resource Groups wizard, enter a **Resource Group Name**.
4. Click **Select LDEVs**.



5. In the Select LDEVs wizard, select the basic LDEV you created from the parity group that backs your vVol pool, select LDEV IDs from the Available LDEVs list, and then click **Add**.



6. Click **OK**.
7. After the LDEVs have been defined, click **Add > Finish > Apply**.



Hitachi Storage Provider for VMware vCenter storage configuration

The following section covers how to configure VSP resources using Storage Provider for VMware vCenter to be able to relay both vVols and VMFS capabilities to VMware vCenter. This enables you to apply VMware storage policies that will be used by OCP StorageClasses.

Deployment of Storage Provider for VMware vCenter and onboarding to VMware vCenter is not covered here. See the [Storage Provider for VMware vCenter Deployment Guide](#).


Onboarding Hitachi Storage

After the Storage Provider for VMware vCenter is deployed, register a VSP storage system to relay capabilities.

Use the following procedure to register a storage system:

Procedure

1. Navigate to the applicable Storage Provider for VMware vCenter IP at <https://Storage-Provider-IP:50001/VasaProviderWebUi/Views/LoginView.jsp>, and then log in using your VMware vCenter or SSO credentials.
2. Click **Manage Storage Systems**.



3. From the **Physical Storage** tab, click **Add Storage System**.
4. From **Add Storage System**, do the following (this might not apply to all storage types):
5. From the **Storage System Type** list, click the system model.
6. Click **SVP** (optional).
7. Enter the **SVP IP** of the storage system.
8. Leave the **RMI Registry Port** at 1099 (optional).
9. Enter the **User ID** and **Password** of the VSP.
10. Click **OK**.



11. Click **OK**, and then select **Reload** to view the progress of the onboarding task.

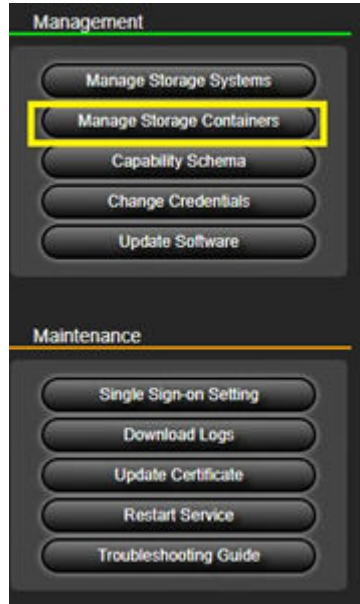
Creating storage containers and capability profiles for VSP vVols

To use vVols, you must create a storage container corresponding to the storage system's resource group and set capability profiles for each dynamic provisioning pool in the group. Profiles for storage containers push storage attributes to the VMware administrator to view within VMware vSphere.

Use the following procedure to create a storage container and define a capability profile:

Procedure

1. From the navigation tree, click **Manage Storage Containers**.

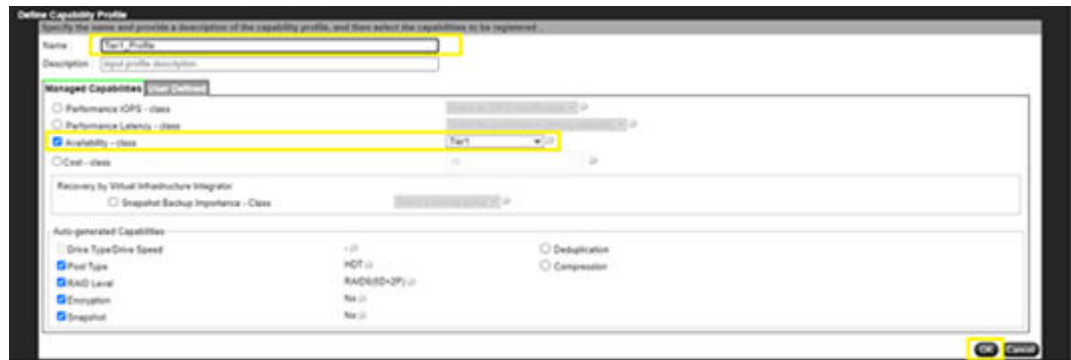


2. Click **Create Storage Container**.
3. Configure the following:
4. Define a storage container **Name**.
5. Select an onboarded **Storage System**.
6. Select the **Resource Group** configured on your VSP.
7. Select an undefined Capability profile, and then click **Define Profile**.



8. In the **Define Capability Profile** window, do the following:
9. Define a profile **Name**.

- Assign managed storage capabilities to your profile. The characteristics need to relate to your vVols resource group that is native to the registered storage system.



- Click **OK** > **Submit**.

Administrators can also define custom capabilities not natively defined within the VASA APIs under Capability Schema. For more information, see [Related Documents](#) (on page 93).

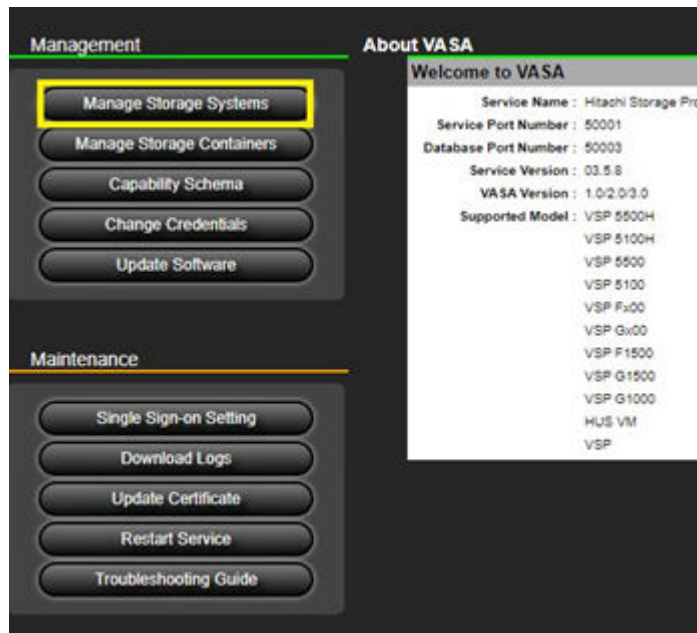
Manage storage systems for VMFS LDEVs

With the Storage Provider for VMware vCenter, attributes of the logical units supporting the VMFS datastore are passed down to VMware vSphere. For these attributes to be passed down you must define these attributes on a per-LDEV basis.

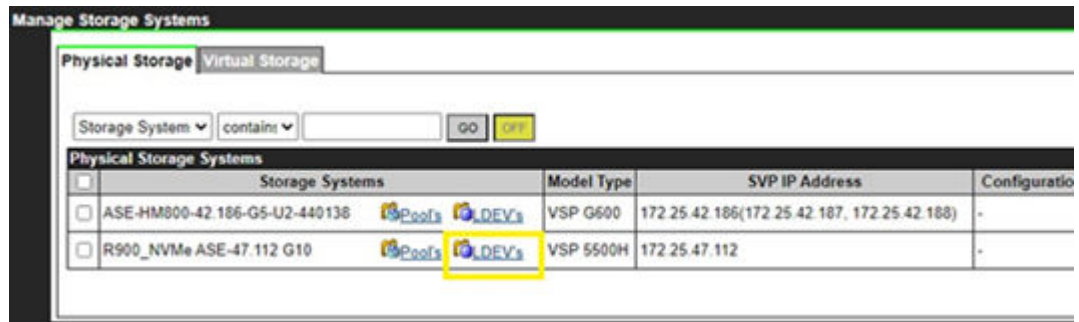
Use the following procedure to define a storage profile tag for a VMFS datastore LUN:

Procedure

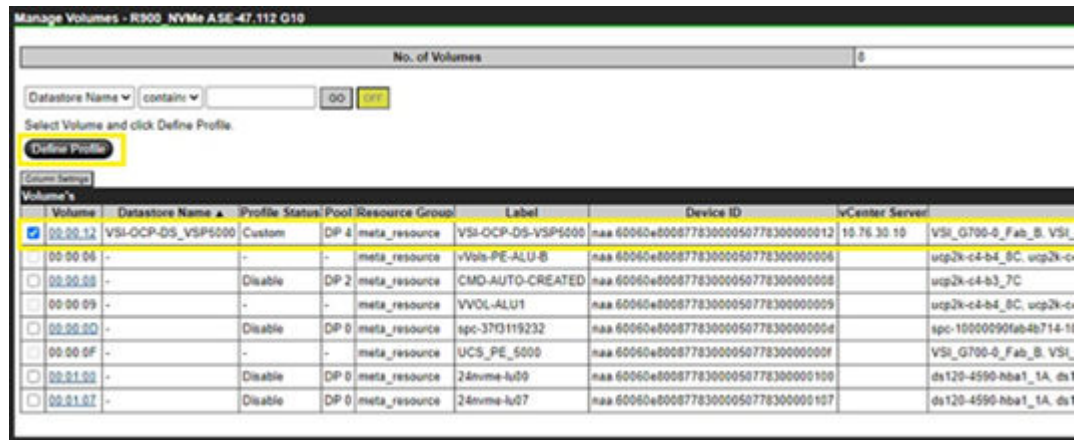
- From the navigation tree, click **Manage Storage Systems**.



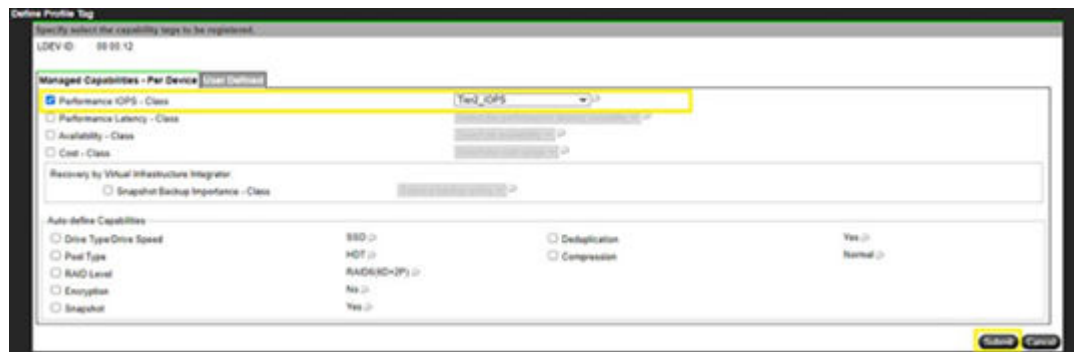
- From the Physical Storage System list, click **LDEV's** on the storage system that contains the applicable VMFS datastore LUN.



- From the storage system volume list, select the applicable **volume ID**, and then click **Define Profile**.



- In the Define Profile Tag wizard, select the applicable tags that you want relayed to the VMware administrator.
- Click **Submit**.



Enabling controller based compression and deduplication

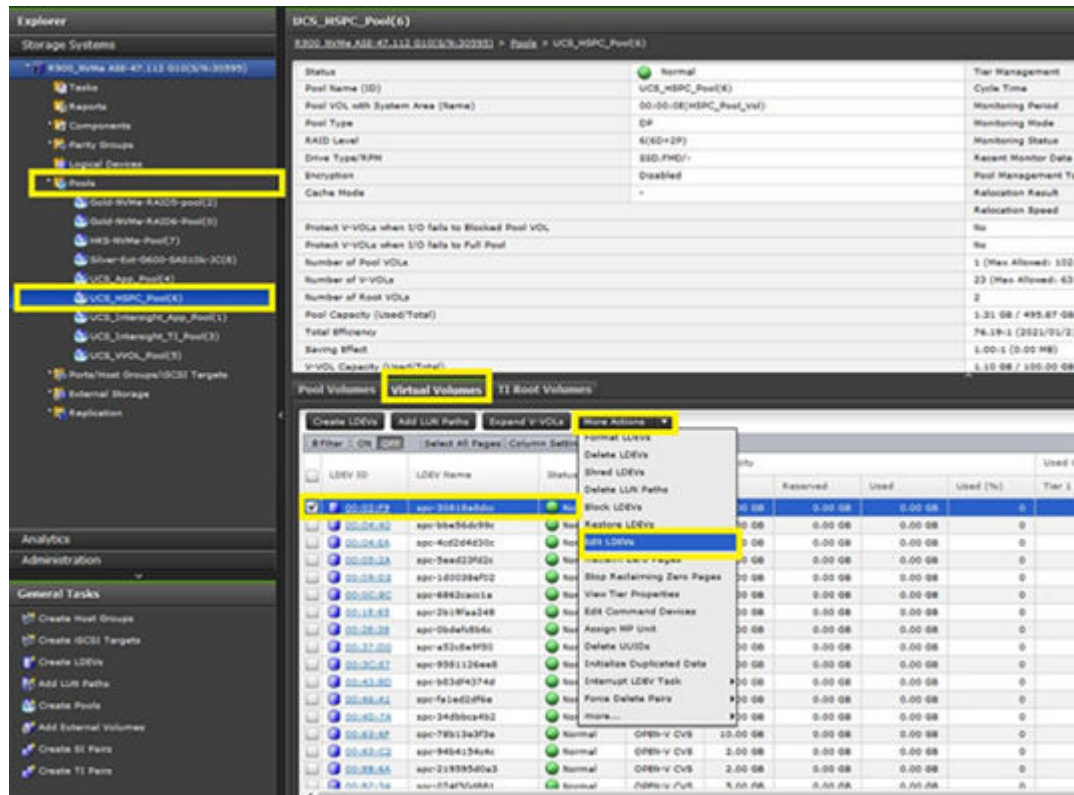
To enable compression and deduplication on a VSP storage system complete these steps.

Note: This section also applies to HSPC.

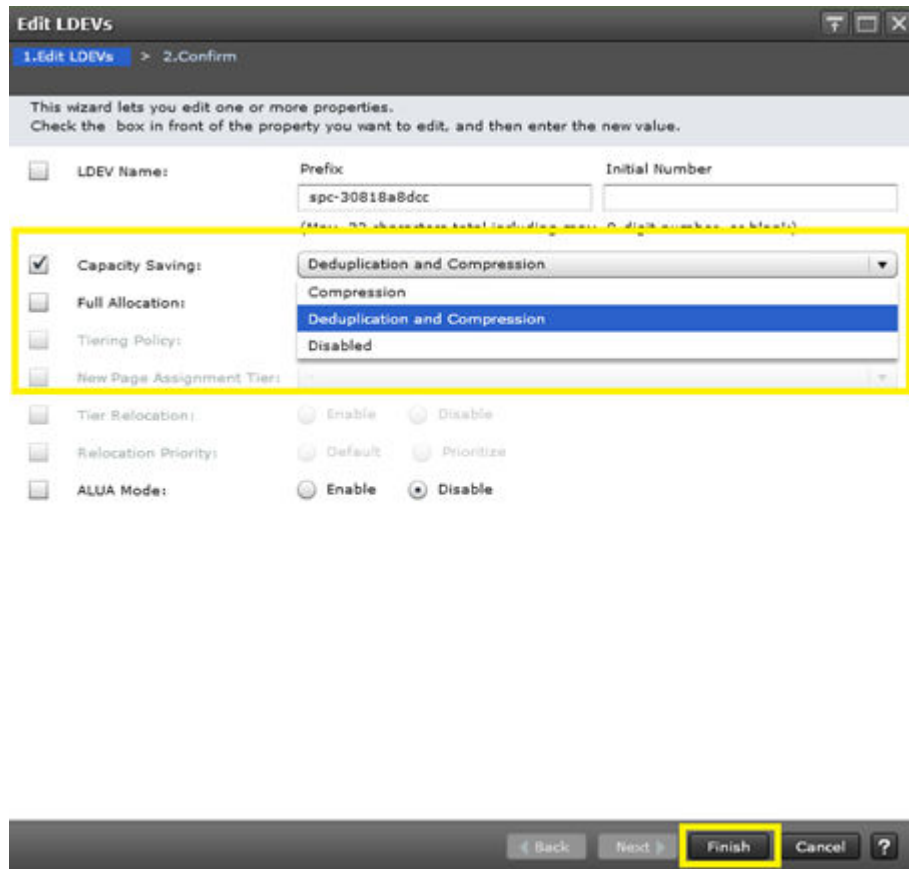
Procedure

- Log in to Hitachi Storage Navigator.
- Select the pool associated with PV allocation for CNS or HSPC.

3. Click the **LDEV** that requires enablement of deduplication.
4. From the **More Actions** drop down select **Edit LDEVs**.



5. Select the Capacity Saving drop down and select **Compression** or **Deduplication and Compression**.
6. Click **Finish**.



VMware vCenter configuration

This section covers the necessary configuration of VMware vCenter in preparation to use OCP StorageClasses in conjunction with VMware vCenter Storage Policies.

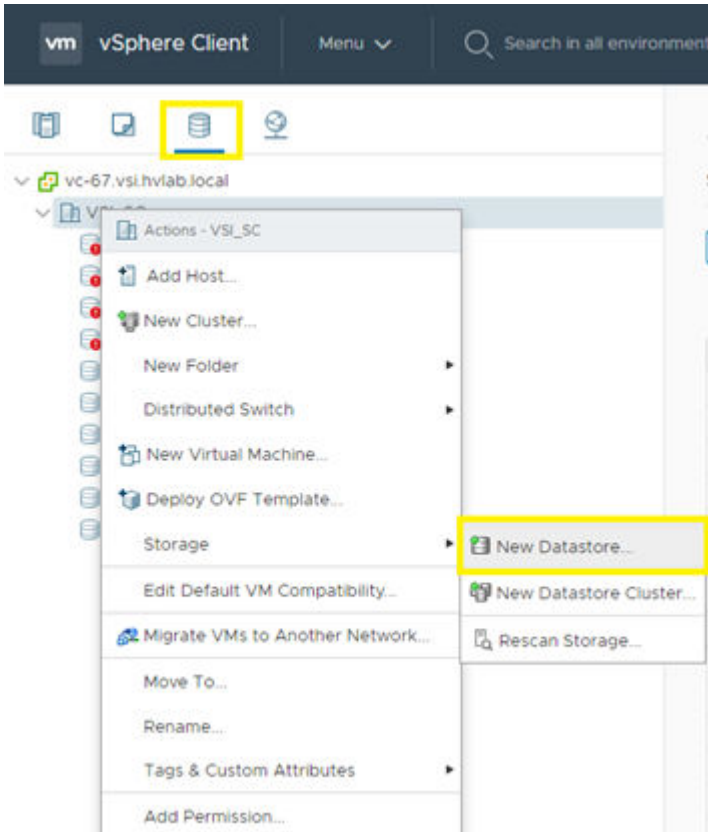
Onboarding VMFS datastores

A VMFS datastore backed by Hitachi VSP storage can be onboarded to VMware vSphere after the storage system LDEV has been allocated to the applicable system host group with a LUN ID. Verify that this has been completed before following these steps.

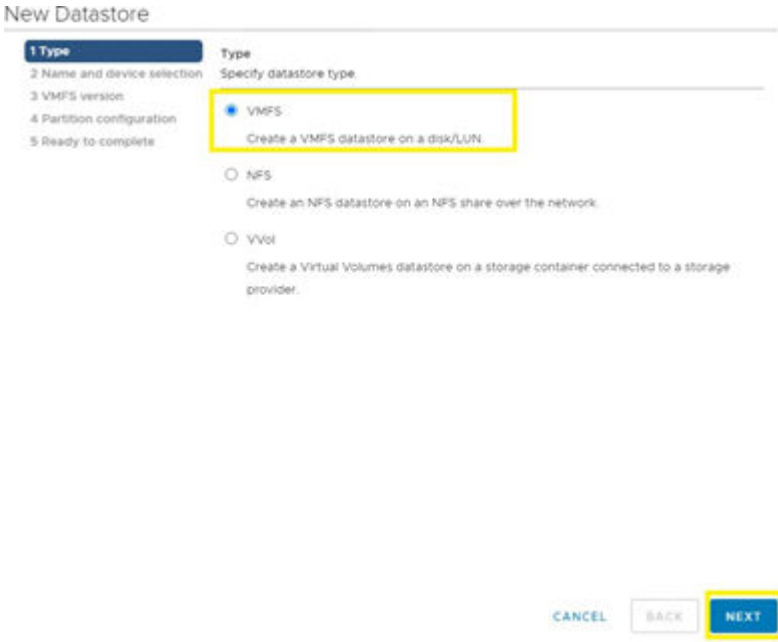
Use the following procedure to onboard a VMFS datastore in vCenter:

Procedure

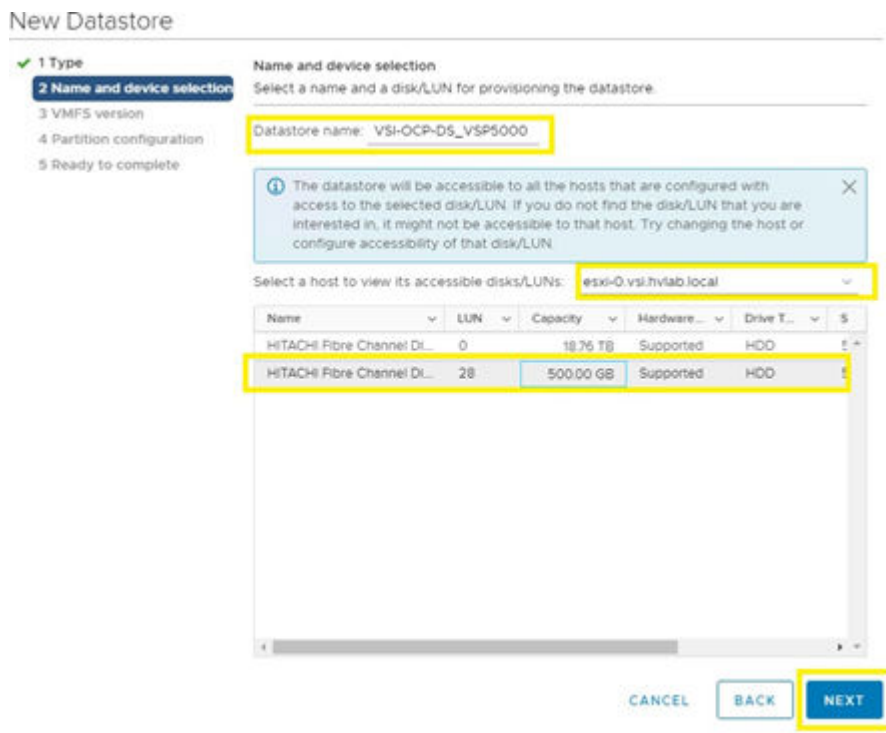
1. Log in to VMware vSphere Client.
2. Select the storage tab, and then right-click on the applicable datacenter.
3. Click **Storage > New Datastore**.



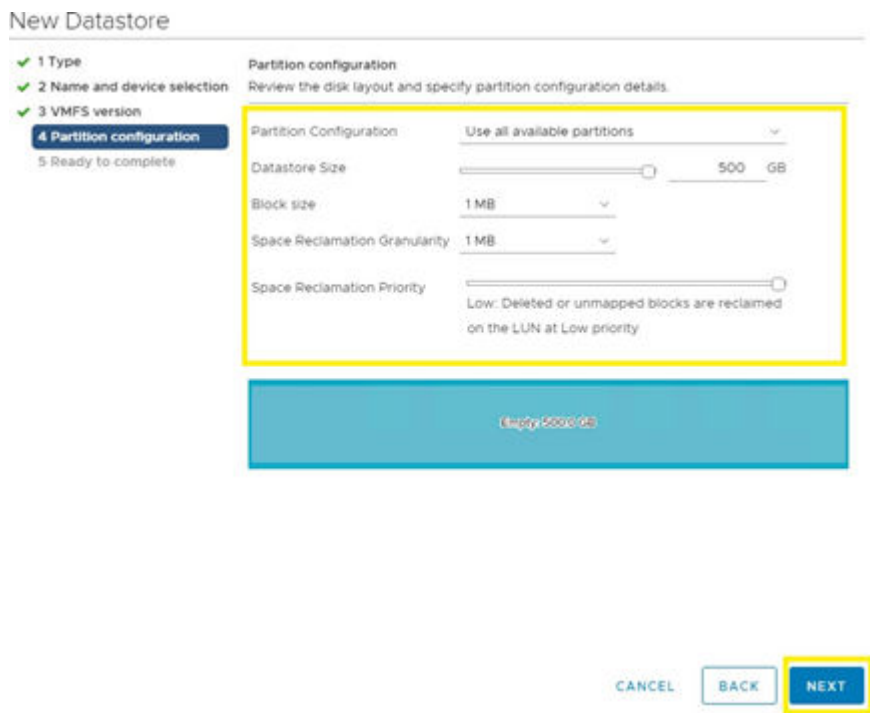
4. Select **VMFS** as the datastore type. Click **Next**.



5. Define a **Datastore name**, and then select an available ESXi host and the LUN presented by Hitachi VSP storage to onboard as the new datastore. Click **Next**.



6. Select **VMFS 6**, and then click **Next**.
7. In **Partition Configuration**, confirm that the entire capacity of the volume is claimed, and then click **Next**.



8. Review the configuration, and then click **Finish**.

New Datastore

- ✓ 1 Type
- ✓ 2 Name and device selection
- ✓ 3 VMFS version
- ✓ 4 Partition configuration
- 5 Ready to complete

Ready to complete
Review your settings selections before finishing the wizard.

General

Name: VSI-OCF-DS_VSP5000
Type: VMFS
Datastore size: 500.00 GB

Device and Formatting

Disk/LUN: HITACHI Fibre Channel Disk
(naa.60060e80087783000050778300000012)

Partition Format: GPT
VMFS Version: VMFS 6
Block Size: 1 MB
Space Reclamation: 1 MB
Granularity:
Space Reclamation Priority: Low: Deleted or unmapped blocks are reclaimed on the LUN at low priority

CANCEL
BACK
FINISH

Viewing VMFS datastore tags

After the VMFS datastores are onboarded, you can view any tags relayed by the Storage Provider for VMware vCenter on the Summary tab in the datastore view.



Note: After Storage Provider for VMware vCenter is deployed it is recommended that you do not set any manual tags within vCenter. It is best practice to pass down all desired tags from Storage Provider for VMware vCenter.

The screenshot shows the VMware vCenter interface for a VMFS datastore named 'VSI-OCF-DS_VSP5000'. The 'Summary' tab is active, displaying details about the datastore's location, type, and hosts. A 'Tags' section is highlighted with a yellow box, showing a table of assigned tags:

Assigned Tag	Category	Description
Location: Sanjo_Clean_HQ	SPSM	This storage device is located in...
Performance IOPS: Tier2	SPSM	Indicates performance class is...

Onboarding a vVols Datastore

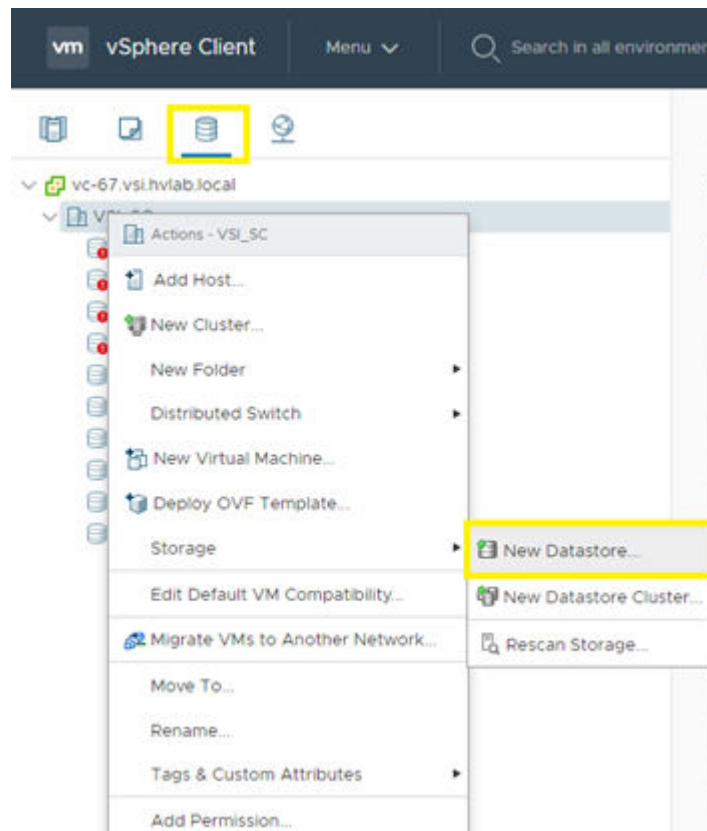
Verify that before onboarding a vVols datastore that applicable VSP storage resources groups have been configured and that the correct storage container and capability profile have been defined.

Also verify that the Storage Provider for VMware vCenter has been successfully registered in vCenter along with the VSP Administrative Logical Unit (ALU), also known as the VASA Protocol Endpoint (PE).

Use the following procedure to onboard a vVols datastore:

Procedure

1. Log in to VMware vSphere Client.
2. Select the **Storage** tab, and then right-click on the applicable datacenter.
3. Click **Storage > New Datastore**.



4. Select **VVol** as the datastore type. Click **Next**.

New Datastore

1 Type

2 Name and container sele...

3 Select hosts accessibility

4 Ready to complete

Type

Specify datastore type.

VMFS
Create a VMFS datastore on a disk/LUN.

NFS
Create an NFS datastore on an NFS share over the network.

VVol
Create a Virtual Volumes datastore on a storage container connected to a storage provider.

CANCEL BACK **NEXT**

5. Define a **Datastore name**, and then select the appropriate backing storage container.
6. Click **Next**.

New Datastore

✓ 1 Type

✓ **2 Name and container sele...**

✓ 3 Select hosts accessibility

4 Ready to complete

Name and container selection

Specify datastore name and backing storage container.

Datastore name: VVOL_DS_VSP5000

Backing Storage Container

Name	Identifier	Maxin...	Existing Datastore
VVOL	vvof1a11b9c35d4... 8b542e7e0d7212d	256 TB	VVOL_DS
VVOL_VSP5000	vvof0844f92d7b... 9e1b1b006ab9f5e9	60 TB	--

2 items

For SCSI-backed VVol datastores, PE LUNs need to be configured manually. Configure SCSI PE LUNs before creating a datastore. If the datastore is created without configuring PE LUNs, the ESXi host marks corresponding VVol datastore as inaccessible.

Backing Storage Container Details

Storage array(s): VSP 5500H_30595
Storage provider(s): VASA

CANCEL BACK **NEXT**

7. Select all of the ESXi hosts in the cluster, and then click **Next**.
8. Review the settings, and then click **Finish**.

New Datastore

- ✓ 1 Type
- ✓ 2 Name and container sele...
- ✓ 3 Select hosts accessibility
- 4 Ready to complete

Ready to complete
Review your settings selections before finishing the wizard.

General
Name: VVOL_DS_VSP5000
Type: VVol

Backing storage container details
Name: VVOL_VSP5000
UUID: vvol.08a415f2d7b54b55-9e1b16b06ab9f5e9
Storage array(s): VSP 5500H_30595
Storage provider(s): VASA

Hosts that will have access to this datastore
Hosts: esxi-1.vsi.hvlab.local
 esxi-2.vsi.hvlab.local
 esxi-0.vsi.hvlab.local

CANCEL BACK FINISH

VMware vCenter storage policies

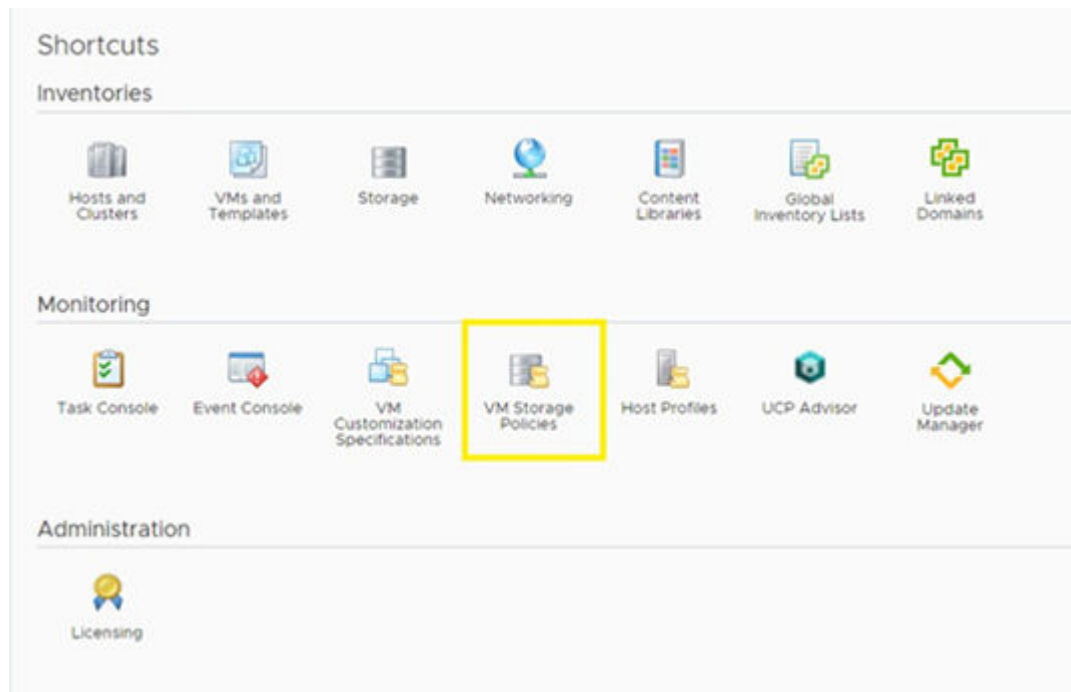
VMware storage policies must be configured before creating a storage class within Red Hat OCP. This section covers how to create storage policies for both VMFS and vVols datastores backed by Hitachi storage with capabilities translated down via the VASA APIs.

VMFS storage policy


Create a VMware storage policy for a VMFS datastore:

Procedure

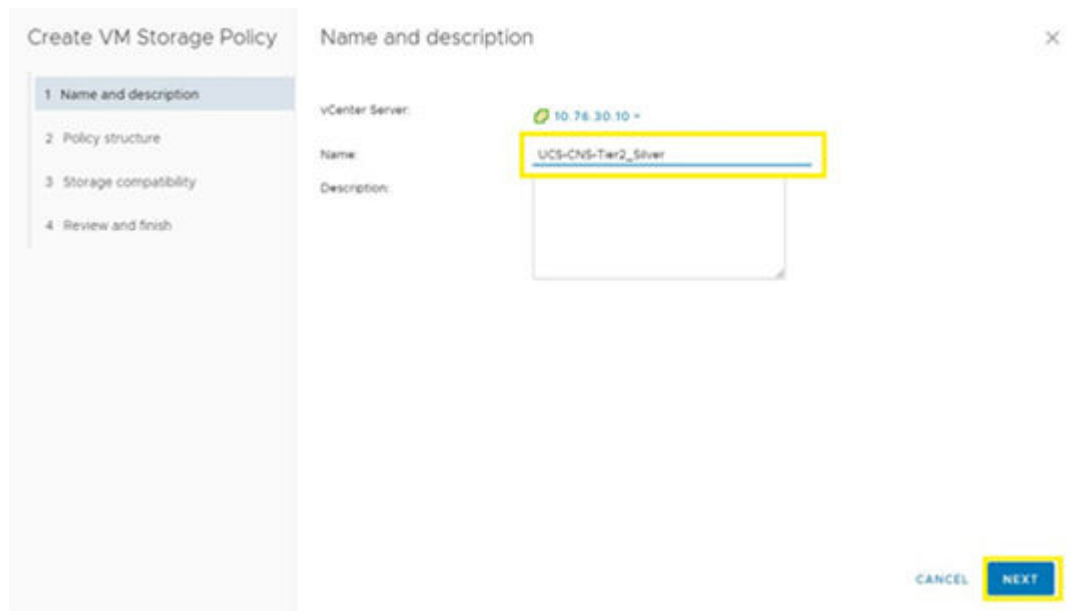
1. Log in to the VMware vSphere Client.
2. From the shortcuts directory, click **VM Storage Policies**.



3. Select **Create VM Storage Policy**.



4. Define a policy name, and then click **Next**.



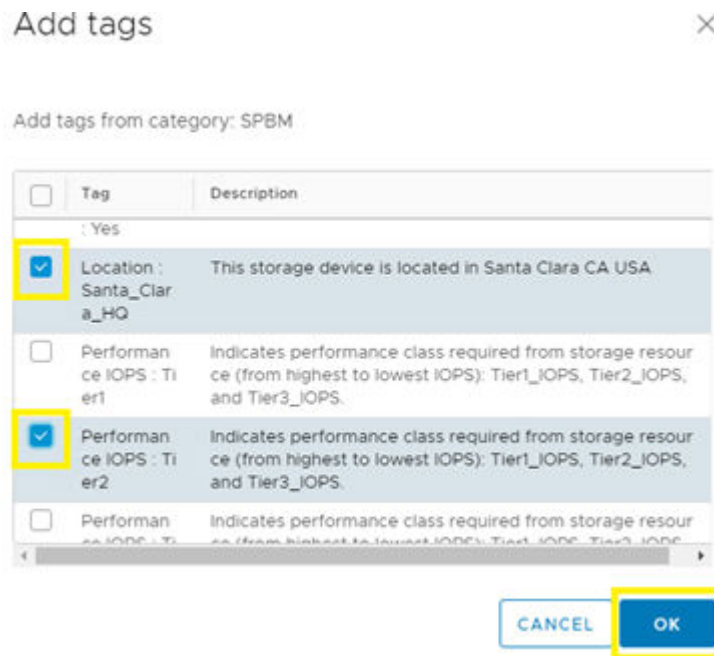
5. For Datastore specific rules, select **Enable tag-based placement rules**, and then click **Next**.

The screenshot shows the 'Create VM Storage Policy' wizard at the 'Policy structure' step. The left sidebar lists five steps: 1. Name and description, 2. Policy structure (selected), 3. Tag based placement, 4. Storage compatibility, and 5. Review and finish. The main content area is titled 'Policy structure' and contains two sections: 'Host based services' and 'Datastore specific rules'. Under 'Host based services', there is a checkbox for 'Enable host based rules' which is unchecked. Under 'Datastore specific rules', there are three checkboxes: 'Enable rules for "vSAN" storage' (unchecked), 'Enable rules for "com.hitachi.storageprovider.vvol" storage' (unchecked), and 'Enable tag based placement rules' (checked and highlighted with a yellow box). At the bottom right, there are three buttons: 'CANCEL', 'BACK', and 'NEXT' (highlighted with a yellow box).

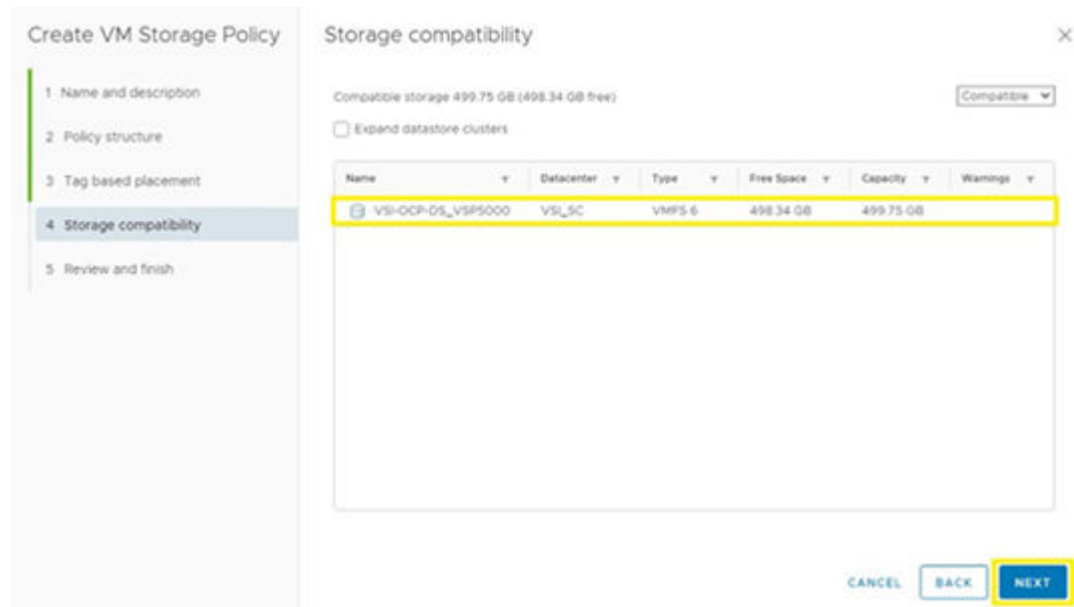
6. Under Tag-based placement select the following:
 - a. Tag category: **SPBM**
 - b. Usage: **Use storage tagged with**
 - c. Select **BROWSE TAGS**

The screenshot shows the 'Create VM Storage Policy' wizard at the 'Tag based placement' step. The left sidebar lists five steps: 1. Name and description, 2. Policy structure, 3. Tag based placement (selected), 4. Storage compatibility, and 5. Review and finish. The main content area is titled 'Tag based placement' and contains a section for 'Rule 1'. Under 'Rule 1', there are three dropdown menus: 'Tag category' (set to 'SPBM' and highlighted with a yellow box), 'Usage option' (set to 'Use storage tagged with' and highlighted with a yellow box), and 'Tags' (set to 'BROWSE TAGS' and highlighted with a yellow box). There is a 'REMOVE' button to the right of 'Rule 1'. At the bottom left, there is an 'ADD TAG RULE' button. At the bottom right, there are three buttons: 'CANCEL', 'BACK', and 'NEXT' (highlighted with a yellow box).

7. Select the applicable storage tags that the storage administrator has defined via the Storage Provider for VMware vCenter, and then click **OK**.



8. Click **Next**.
9. From the Storage compatibility window, you will see datastores that match the tags that you enabled in the previous step. Click **Next**.



10. Click **Finish**.

Red Hat OCP Configuration

Red Hat OCP provides a GUI and a CLI for administrators to deploy PVs with containerized applications. This section describes how to create prerequisite storage operations to provide persistent storage to the environment backed by Hitachi VSP storage via the OCP GUI. For additional CLI operations see [Related Documents \(on page 93\)](#).

Verify that you have completed the previous procedures in this guide before continuing.

StorageClass

A StorageClass provides a way for administrators to describe the classes of storage they offer which can be requested through the OCP interface. Each class contains fields for administrators to define the provisioner, parameter, and reclaim policy which are used for PV creation via PVCs. The provisioner parameter in a virtual environment backed by Hitachi storage on top of VMware would use the CSI provisioner `csi.vsphere.vmware.com`. StorageClasses also have specific names and are called out when making PVCs. When administrators create StorageClass objects, these objects cannot be updated once they have been created.



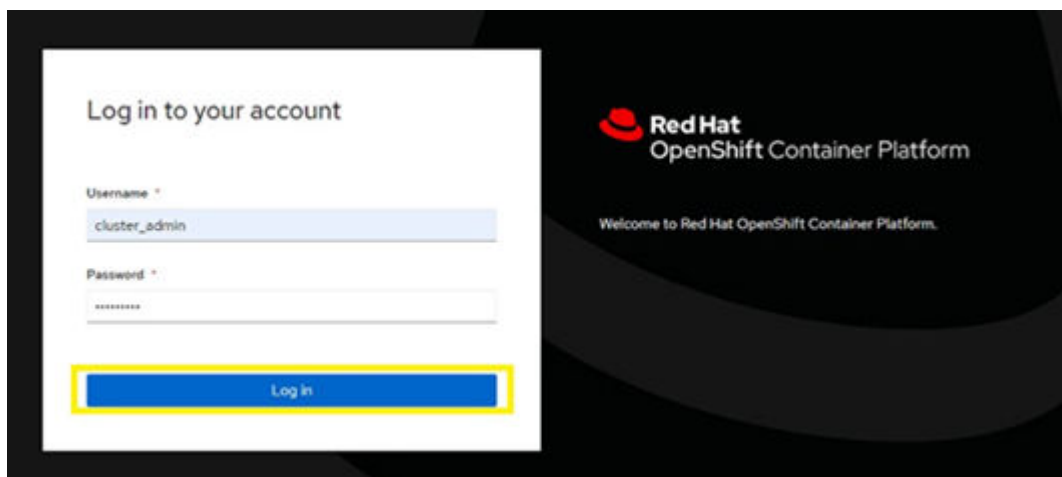
Note: PVs can have various reclaim policies, including Retain, Recycle, and Delete. For dynamically provisioned PVs, the default reclaim policy is Delete.

vVols Storage Class

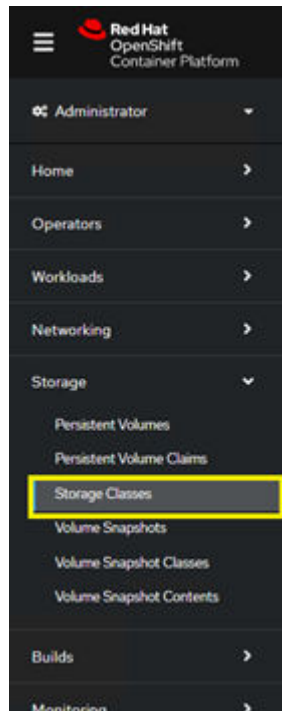
Use the following procedure to create a Storage Class from the OCP GUI:

Procedure

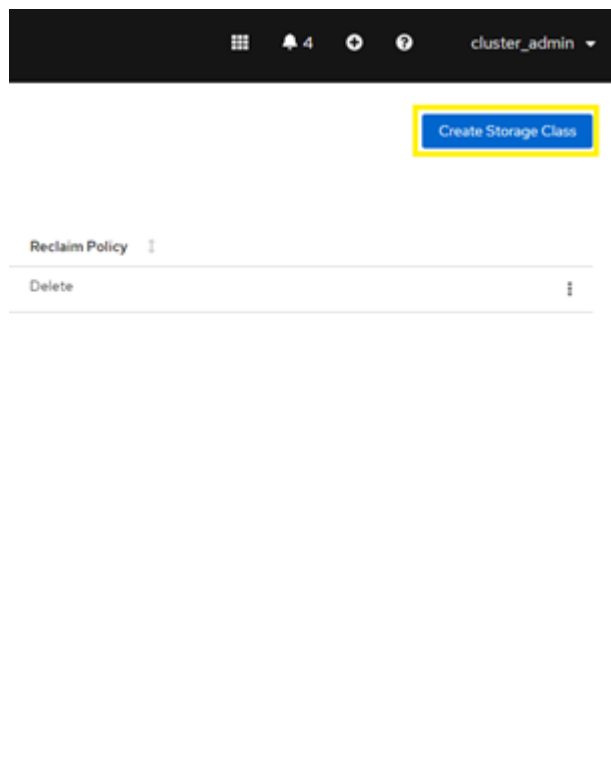
1. Confirm that the VMware storage policy backed by a VMFS datastore is created and has compatible storage.
2. Log in to the OCP GUI using valid user credentials.



3. From the navigation tree, select **Storage > Storage Classes**.



4. Click **Create Storage Class**.



5. Click **Edit YAML**.

Create Storage Class Edit YAML

Name *

Description

Reclaim Policy *

Delete

Determines what happens to persistent volumes when the associated persistent volume claim is deleted. Defaults to 'Delete'

Provisioner *

Select Provisioner

Determines what volume plugin is used for provisioning persistent volumes.

Create Cancel

6. Enter the following parameters on the YAML file:
 - a. Define the API version used.
 - b. Define the *kind* parameter as StorageClass.
 - c. Select an applicable StorageClass name.
 - d. Optionally select whether this is a default StorageClass.
 - e. Define the provisioner used. For CNS, the default is *csi.vsphere.vmware.com*.
 - f. Define the reclaim Policy as *Delete*.
 - g. Point to the VMware storage policy within vCenter that utilizes VSP vVols.

```

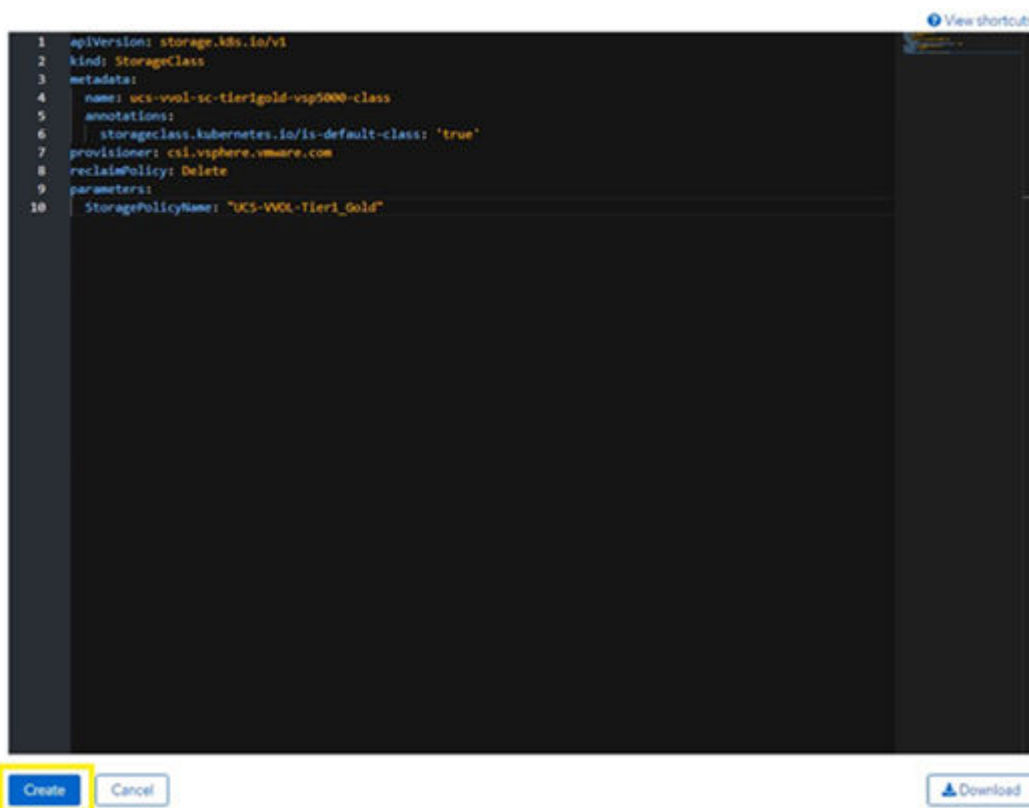
apiVersion: storage.k8s.io/v1           (1)
kind: StorageClass                      (2)
metadata:
  name: ucs-vvol-sc-tier1gold-vsp5000-class (3)
  annotations:
    storageclass.kubernetes.io/is-default-class: 'true' (4)
  provisioner: csi.vsphere.vmware.com (5)
  reclaimPolicy: Delete (6)
  parameters:
    StoragePolicyName: "UCS-VVOL-Tier1_Gold" (7)

```

7. Click **Create**.

Create Storage Class

Create by manually entering YAML, or JSON definitions, or by dragging and dropping a file into the editor.



```

1  apiVersion: storage.k8s.io/v1
2  kind: StorageClass
3  metadata:
4    name: ucs-vvol-sc-tier1gold-vsp5000-class
5  annotations:
6    storageclass.kubernetes.io/is-default-class: 'true'
7  provisioner: csi.vsphere.vmware.com
8  reclaimPolicy: Delete
9  parameters:
10   StoragePolicyName: "UCS-VVOL-Tier1_Gold"

```

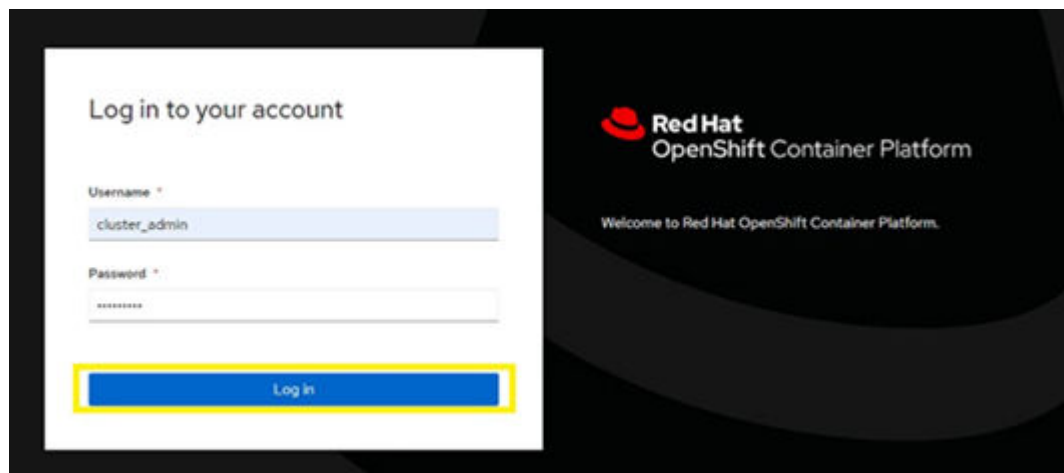
Buttons: Create (highlighted), Cancel, Download

VMFS StorageClass

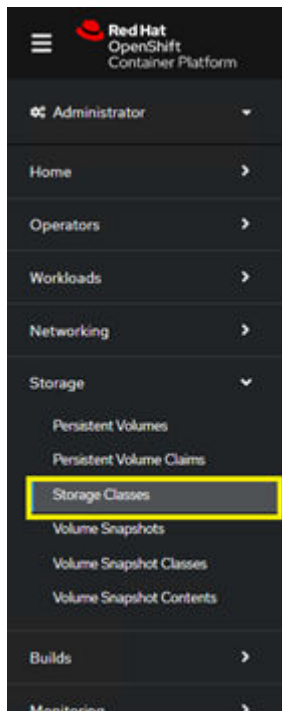
Create a StorageClass using the OCP GUI:

Procedure


1. Confirm that the VMware storage policy backed by a VMFS datastore is created and has compatible storage.
2. Log in to the OCP GUI using valid user credentials.



- From the navigation tree, select **Storage > Storage Classes**.



- Click **CreateStorage Class**.



- Select **Edit YAML** in the upper right-hand corner.

Create Storage Class Edit YAML

Name *

Description

Reclaim Policy *

Delete

Determines what happens to persistent volumes when the associated persistent volume claim is deleted. Defaults to 'Delete'

Provisioner *

Select Provisioner

Determines what volume plugin is used for provisioning persistent volumes.

Create Cancel

6. Enter the following parameters on the YAML file:
 - a. Define the *kind* parameter as StorageClass.
 - b. Define the API version used.
 - c. Select an applicable StorageClass name.
 - d. Optionally select whether this is a default StorageClass.
 - e. Define the provisioner used. For CNS, the default is *csi.vsphere.vmware.com*.
 - f. Point to the VMware storage policy within vCenter that utilizes VSP VMFS storage.
 - g. Define reclaimPolicy as *Delete*.

```

kind: StorageClass          (1)
apiVersion: storage.k8s.io/v1      (2)
metadata:
  name: csi-cns-ucs-vsp5000-sc      (3)
  annotations:
    storageclass.kubernetes.io/is-default-class: 'false'      (4)
  provisioner: csi.vsphere.vmware.com      (5)
parameters:
  StoragePolicyName: UCS-CNS-Tier2_Silver      (6)
  reclaimPolicy: Delete      (7)

```

7. Click **Create**.

Create Storage Class

Create by manually entering YAML, or JSON definitions, or by dragging and dropping a file into the editor.

```

1 kind: StorageClass
2 apiVersion: storage.k8s.io/v1
3 metadata:
4   name: csi-cns-ucs-vsp5000-sc
5   annotations:
6     storageclass.kubernetes.io/is-default-class: 'false'
7   provisioner: csi.vsphere.vmware.com
8 parameters:
9   StoragePolicyName: UCS-CNS-Tier2_Silver
10 reclaimPolicy: Delete
11 volumeBindingMode: Immediate

```

Buttons: Create (highlighted), Cancel, Download

Persistent Volume Claim (PVC)

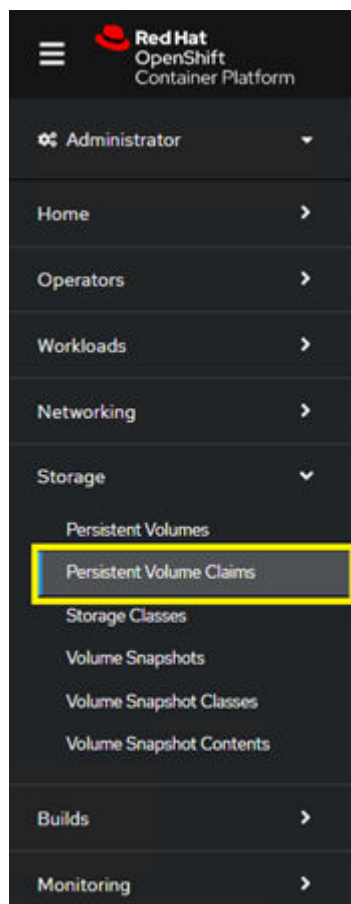
A PVC is a request for storage by a user. PVCs are dynamically provisioned using StorageClass parameters. PVCs are namespace objects, mounting claims with multiple modes (ROX, RWX) is only possible within one namespace.

vVols PVC

Before creating a PVC verify that an appropriate StorageClass exists. Use the following procedure to deploy a PVC using vVols:

Procedure

1. From the navigation tree, select **Storage > Persistent Volume Claims**.



2. Verify that you are using the correct namespace, then and click **Create Persistent Volume Claim**.



3. From the Create Persistent Volume Claim wizard, select the appropriate **Storage Class** that uses vVols from VMware CNS.
4. Define a **Persistent Volume Claim Name**.
5. Define the **Access Mode** parameters.
6. Enter an appropriate **Size**.
7. Click **Create**.

Project: engineering ▾

Create Persistent Volume Claim

[Edit YAML](#)

Storage Class

SC ucs-vvol-sc-tier1gold-vsp5000-class ▾

Storage class for the new claim

Persistent Volume Claim Name *

mysql-dev-pv

A unique name for the storage claim within the project

Access Mode *

 Single User (RWO) Shared Access (RWX) Read Only (ROX)

Permissions to the mounted drive

Size *

100 GiB ▾

Desired storage capacity

 Use label selectors to request storage

Use label selectors to define how storage is created

Create

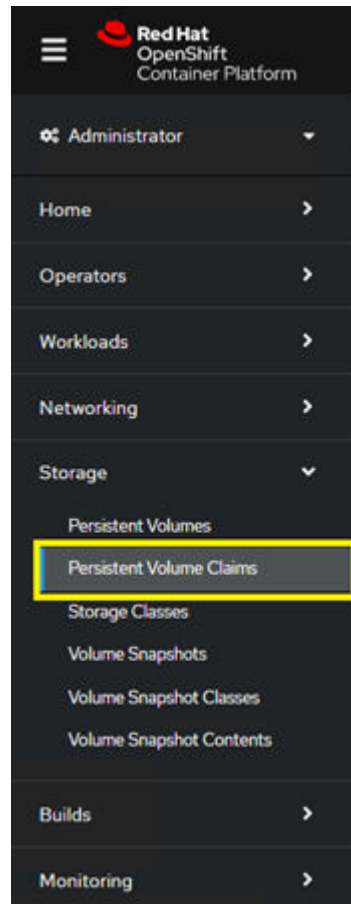
Cancel

VMFS PVC

Before creating a PVC verify that an appropriate StorageClass exists. Use the following procedure to deploy a PVC using vVols:

Procedure

1. From the navigation tree, select **Storage > Persistent Volume Claims**.



2. Verify that you are using the correct namespace, then and click **Create Persistent Volume Claim**.



3. In the Create Persistent Volume Claim wizard, select the appropriate **Storage Class** that utilizes VMFS from VMware CNS.
4. Define a **Persistent Volume Claim Name**.
5. Define the **Access Mode** parameters.
6. Enter an appropriate **Size**.
7. Click **Create**.

Project: engineering ▾

Create Persistent Volume Claim Edit YAML

Storage Class

▾
Storage class for the new claim

Persistent Volume Claim Name *

A unique name for the storage claim within the project

Access Mode *

Single User (RWO) Shared Access (RWX) Read Only (ROX)
Permissions to the mounted drive

Size *

▾
Desired storage capacity

Use label selectors to request storage
Use label selectors to define how storage is created


Viewing persistent storage on VMware vCenter

After you have deployed the PVCs, you can view them natively within VMware vCenter. From this vantage point, administrators can view other information about the object such as PVC ID, PVC name, as well as namespace information which relates to the OCP deployment.

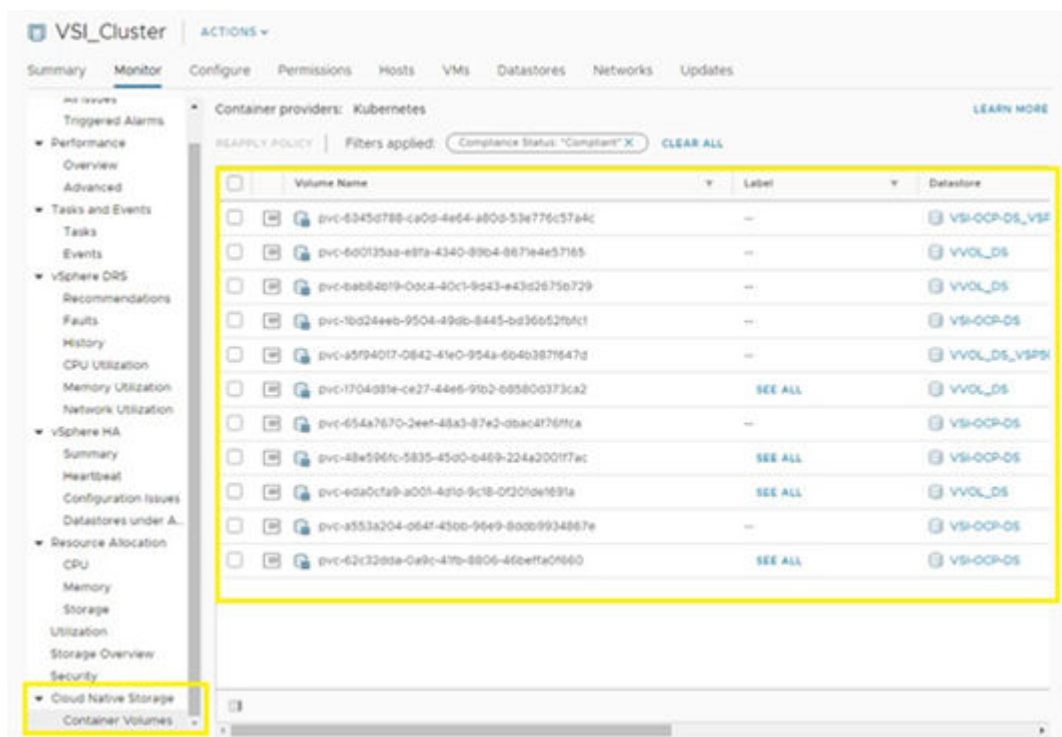
Use the following procedure to view container volumes within vCenter:

Procedure

1. Log on to VMware vSphere Client.
2. Click **Hosts and Clusters**.
3. Click your vCenter cluster.
4. Select the **Monitor** tab.



5. From the Monitor navigation tree, select **Cloud Native Storage > Container Volumes**.
6. The workspace presents the PVCs deployed via OCP, and you can view compliance state, datastore, volume ID, and capacity information.



Deleting PVCs

Deleting a PVC will also remove any associated PVs that are deployed.

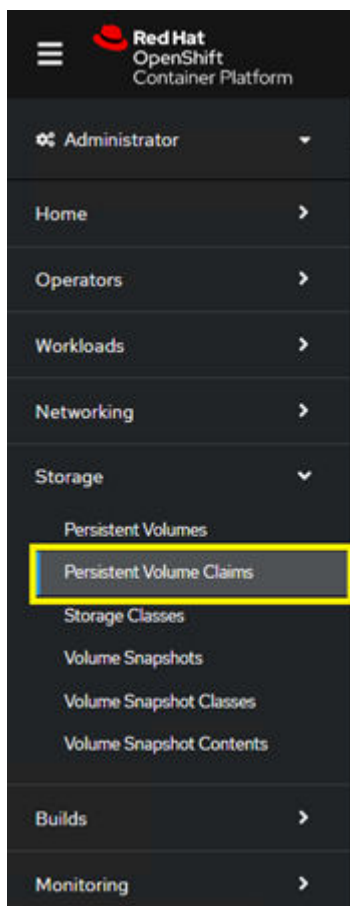


Note: To delete PVs, it is recommended that you also delete the associated PVC.

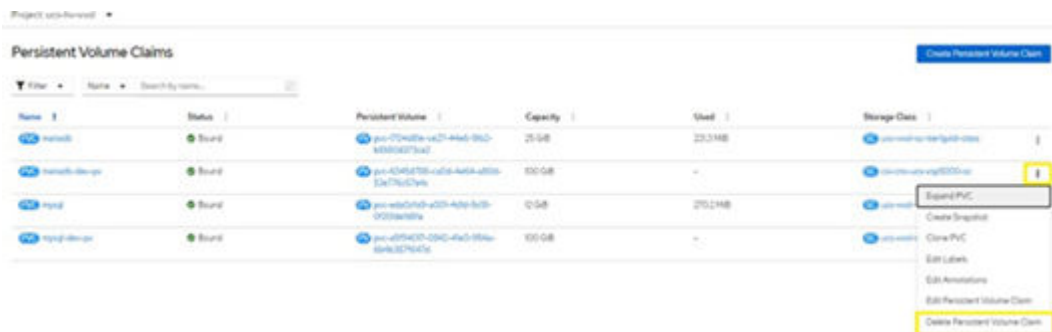
Use the following procedure to delete a PVC:

Procedure

1. From the navigation tree, select **Storage > Persistent Volume Claims**.



2. Verify that you are using the correct namespace, and then click the More Options next to the PVC that you want to delete.



3. On the confirmation window, click **Delete**.

Delete Persistent Volume Claim

Are you sure you want to delete mariadb-dev-pv Persistent Volume Claim?



Solution implementation - Bare Metal Cisco UCS with HSPC

The following describes how to implement Hitachi VSP storage on Cisco and Hitachi Adaptive Solutions for CI with VMware to back your Red Hat OCP hybrid environment with persistent storage using HSPC. Red Hat OCP must be deployed before this implementation.



Note: Hybrid environments consist of Bare Metal servers in which HSPC provides persistent storage via the Fibre Channel protocol.



Note: HSPC only supports RHEL as the OS of the worker node.

Configuration steps in this section assume that parity groups and LDEVs have been configured on the Hitachi Virtual Storage Platform (VSP) as part of the solution build configured by a Hitachi partner or Hitachi professional services. If parity groups have not been configured on the Hitachi VSP, see the Hitachi Storage Virtualization Operating System RF (SVOS RF) documentation to create parity groups before continuing with this section.

Ensure that you have planned which parity groups and LDEVs to use for specific storage requirements. Your configuration might vary based on the types of drives ordered with your VSP and the parity groups configured on them. This section also assumes HDP pools have been created on the VSP storage system. For information on how to create HDP pools refer to [Create Hitachi Dynamic Provisioning \(HDP\) pools \(on page 20\)](#).



Note: Verify that the VSP backing pool is Hitachi Dynamic Provision (HDP) and is not an Hitachi Dynamic Tiered (HDT) type pool.



Note: To avoid data corruption, only a single node should have read/write data operations to the PV.

HSPC configuration

HSPC allows admins to create containers and run stateful applications inside those containers by using the VSP series volumes as dynamically provisioned persistent volumes.

HSPC VSP Host Groups

Storage Provider for VMware vCenter automatically searches host group targets based on their name and if none exist, HSPC will automatically create a host group based off the ports designated within the StorageClass. HSPC uses the following naming rule when creating host groups:

```
"spc-<wwn1>-<wwn2>-<wwn3>"
```


If host groups already exist delete or rename them based on the naming rule.

Node labeling

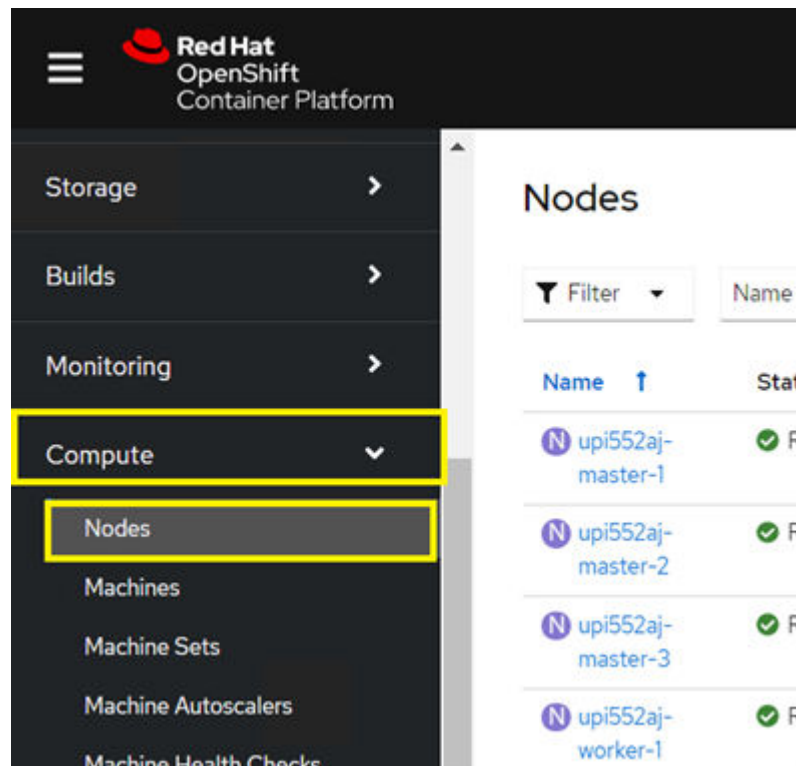
With hybrid OCP deployments operators must distinguish node types between physical and virtual workers. Labeling nodes allows granular deployment of resources to the correct compute host. To label nodes within a hybrid cluster follow these steps:

Procedure

1. Log in to the OCP GUI using valid user credentials.



2. From the navigation tree, select **Compute > Nodes**.



3. From the Nodes list click the ellipsis icon on a physical or virtual node and select **Edit Labels**.

Nodes

Name	Status	Role	Pods	Mem.	CPU	Filesy...	Creat...
upi552aj-master-1	Ready	master	31	3.47 GiB / 15.66 GiB	0.484 cores	81.46 GiB / 358.9 GiB	Nov 13, 11:04 am
upi552aj-master-2	Ready	master	26	3.55 GiB / 15.66 GiB	0.501 cores	80.22 GiB / 358.9 GiB	Nov 13, 11:10 am
upi552aj-master-3	Ready	master	50	5.2 GiB / 15.66 GiB	0.770 cores	87.99 GiB / 358.9 GiB	Nov 13, 11:14 am
upi552aj-worker-1	Ready	worker	32	6.56 GiB / 15.66 GiB	0.897 cores	196 GiB / 358.9 GiB	Nov 13, 11:25 am
upi552aj-worker-2	Ready	worker	15	1.99 GiB / 15.66 GiB	0.164 cores	64.39 GiB / 358.9 GiB	Nov 13, 11:26 am
upi552aj-worker-3	Ready	worker	11	24.43 GiB / 754.3 GiB	0.169 cores	24.44 GiB / 195.9 GiB	Nov 13, 11:26 am
upi552aj-worker-4	Ready	worker	15	13.38 GiB / 376.4 GiB	0.237 cores	23.09 GiB / 195.9 GiB	Nov 13, 11:26 am

- If a virtual node is selected define **node-type=vm-cns**, and then click **Save**.

Edit Labels

Labels help you organize and select resources. Adding labels below will let you query for objects that have similar, overlapping or dissimilar labels.

Labels for upi552aj-worker-1

beta.kubernetes.io/arch=amd64
 beta.kubernetes.io/instance-type=vsphere-vm.cpu-4.mem-16gb.os-cores
 beta.kubernetes.io/os=linux
 kubernetes.io/arch=amd64
 kubernetes.io/hostname=upi552aj-worker-1
 kubernetes.io/os=linux
 node.openshift.io/os_id=rhcos
 node-role.kubernetes.io/worker

node-type=vm-cns

Cancel

Save

- If a physical node is selected define **node-type=hspc-fc**, and then click **Save**.

Edit Labels

Labels help you organize and select resources. Adding labels below will let you query for objects that have similar, overlapping or dissimilar labels.

Labels for upi552aj-worker-3

beta.kubernetes.io/arch=amd64
 beta.kubernetes.io/os=linux
 kubernetes.io/arch=amd64
 kubernetes.io/hostname=upi552aj-worker-3
 kubernetes.io/os=linux
 node.openshift.io/os_id=rhel
 node-role.kubernetes.io/worker
 node-type=hspc-fc

Cancel

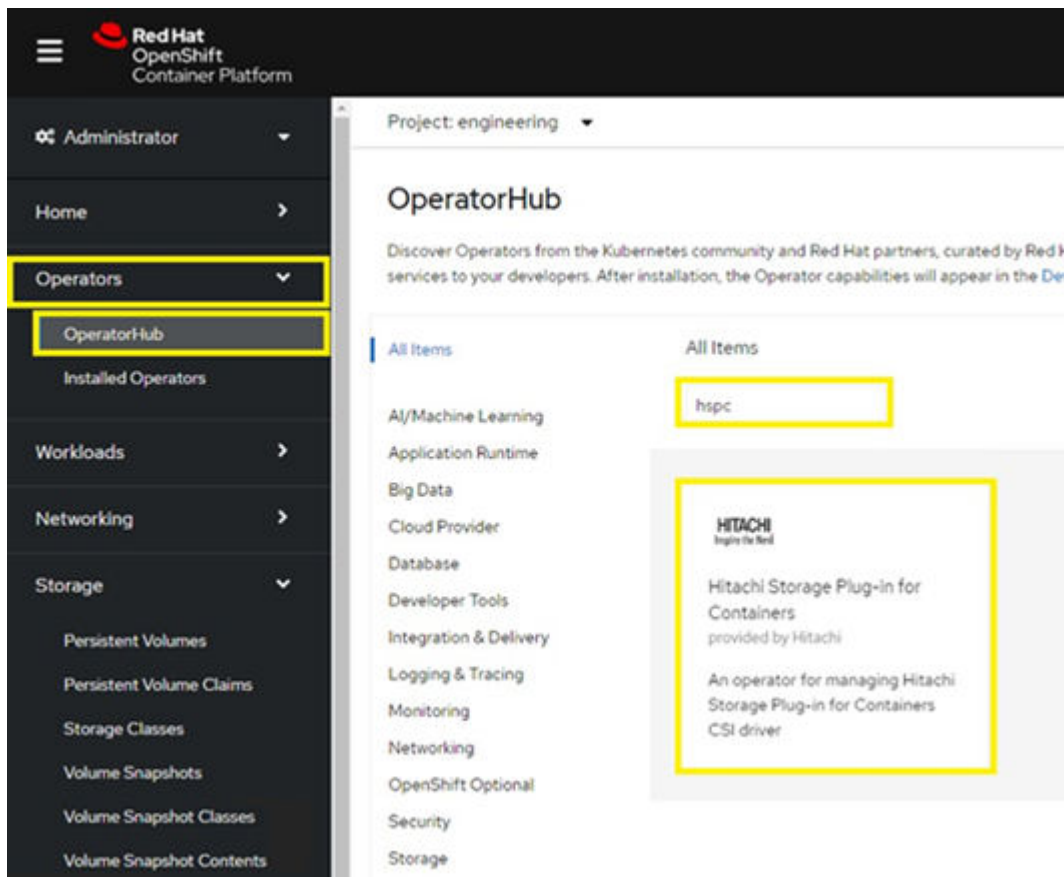
Save

Installing HSPC

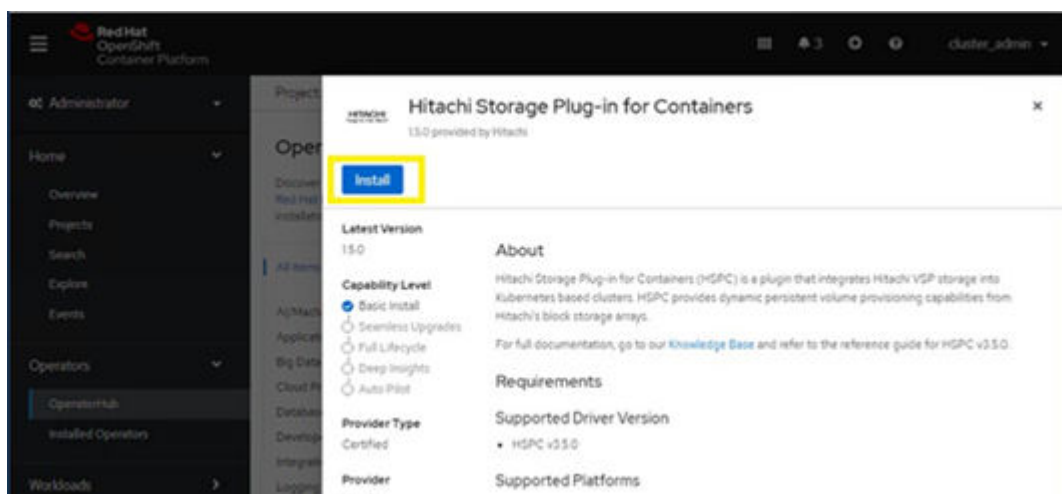
To install HSPC within an OCP Environment follow these steps:

Procedure

1. From the navigation tree, select **Operators > OperatorHub**.
2. Within the OperatorHub search bar type **HSPC**.
3. In the search results select **Hitachi Storage Plug-in for Containers**.

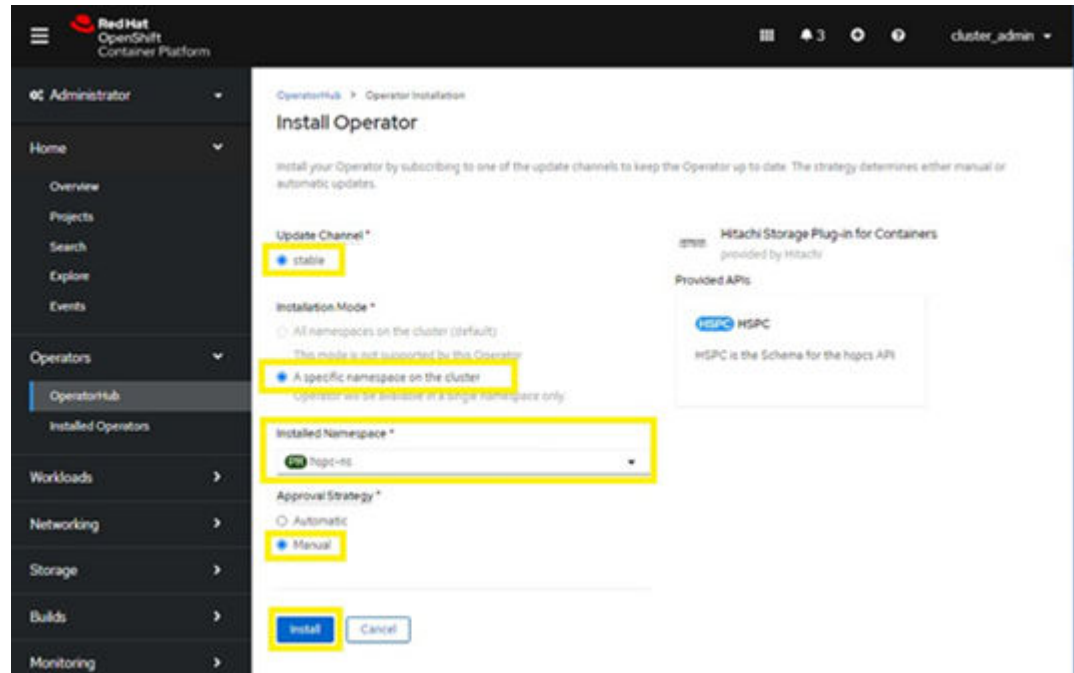


4. From the Install window prompt confirm the HSPC version, and then click **Install**.

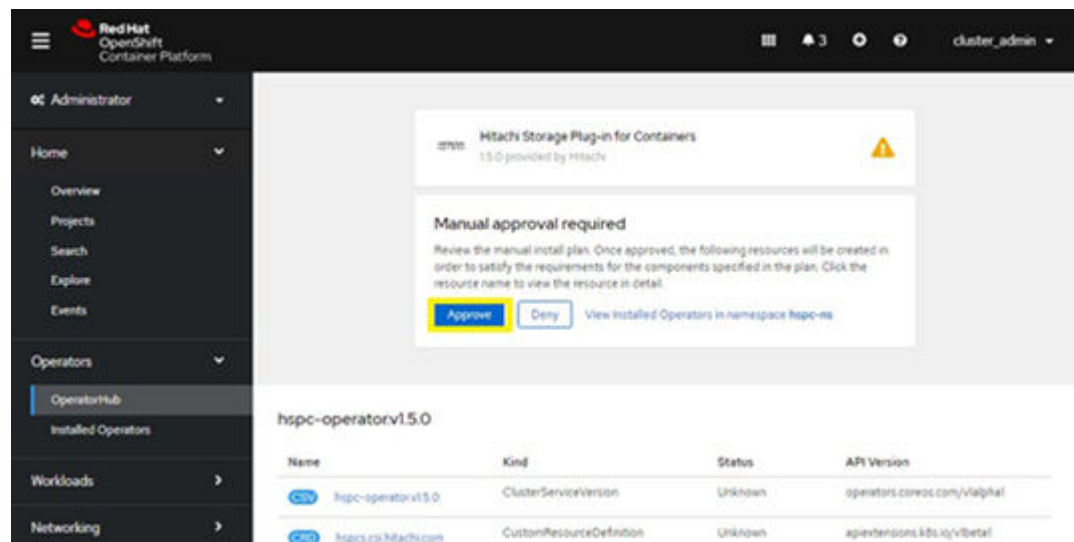


5. In the Install Operator window, select the following:
6. Update Channel – **Stable**.
7. Installation Mode – **A Specific namespace on the cluster**.

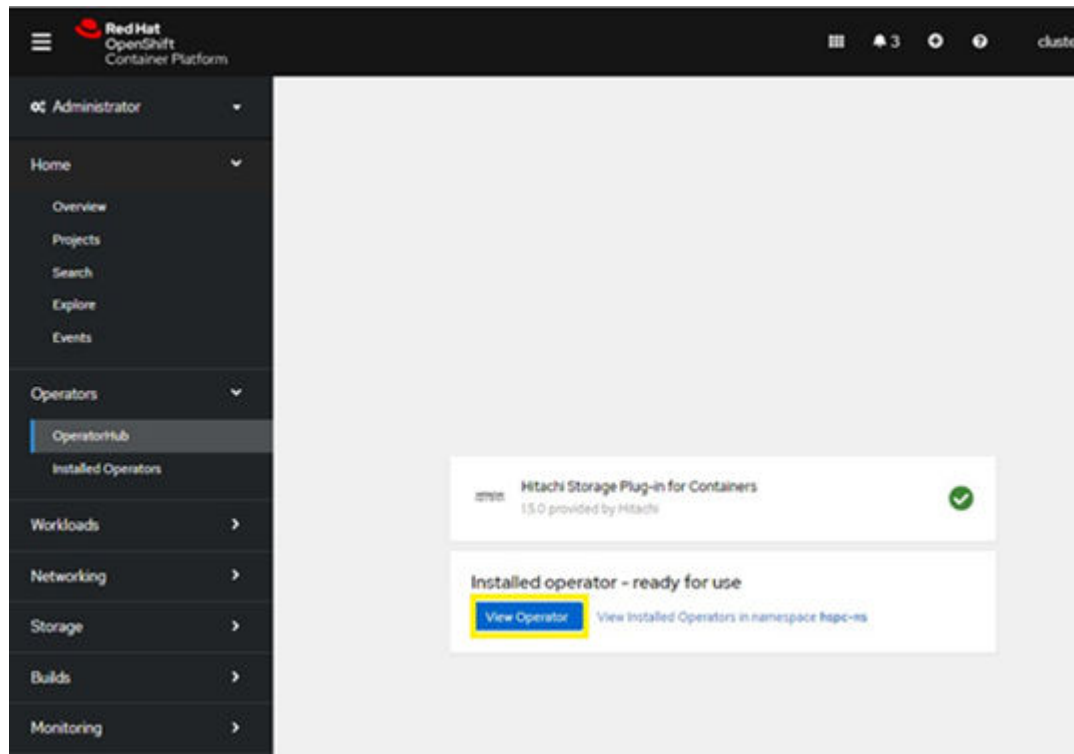
8. Installed Namespace – **Select an applicable namespace**, in this example hspc-ns.
9. Approval Strategy – **Manual**.
10. Click **Install**.



11. Click **Approve** on the **Manual approval required** prompt. HSPC installation starts.



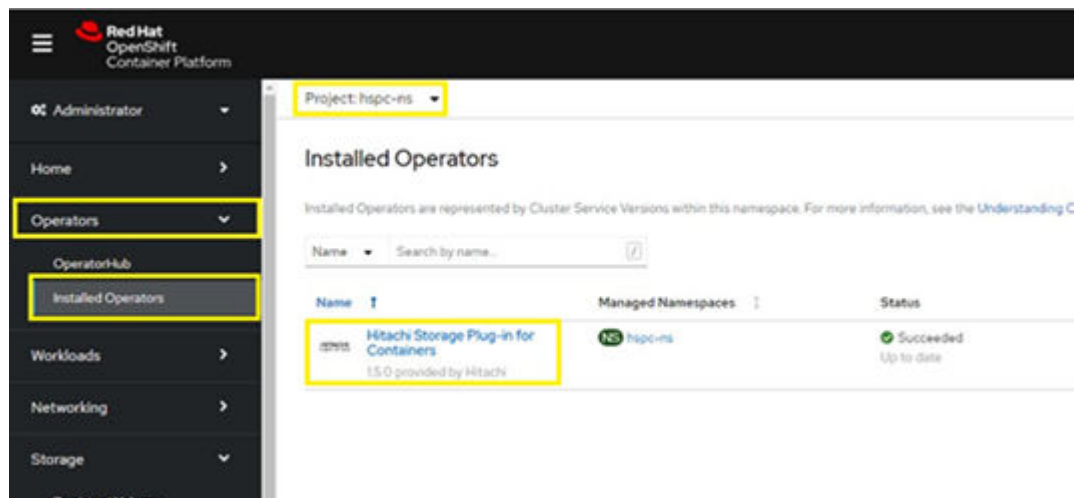
12. Once completed you will be notified that the operator is ready for use. Click **View Operator** and proceed with [Deploying HSPC instances \(on page 63\)](#).



Deploying HSPC instances

Procedure

1. From the navigation tree, select **Operators > Installed Operators**.
2. Select the **namespace** where HSPC operator is installed, in this example hspc-ns.
3. Click **Hitachi Storage Plug-in for Containers**.



4. From the Details tab, click **Create Instance** under Provided APIs.



Hitachi Storage Plug-in for Containers

1.5.0 provided by Hitachi

[Details](#)
[YAML](#)
[Subscription](#)
[Events](#)
[HSPC](#)

Provided APIs

HSPC HSPC

HSPC is the Schema for the hspcs API

[+ Create Instance](#)

- Define the HSPC instance name within the applicable namespace, and then click **Create**.

[Hitachi Storage Plug-in for Containers](#) > [Create HSPC](#)

Create HSPC

Create by completing the form. Default values may be provided by the Operator authors.

Configure via: Form View YAML View

Note: Some fields may not be represented in this form view. Please select "YAML view" for full control.

Name *

hspc-sc-ucs

Labels

app=frontend

Image Pull Secrets


ImagePullSecrets for pulling images from RedHat registries

[Create](#)

[Cancel](#)


- Under the HSPC tab you will now see the running HSPC instance.


Installed Operators > Operator Details

 **Hitachi Storage Plug-in for Containers**
1.5.0 provided by Hitachi

Details YAML Subscription Events **HSPC**

HSPCs

Name ▾ Search by name... 

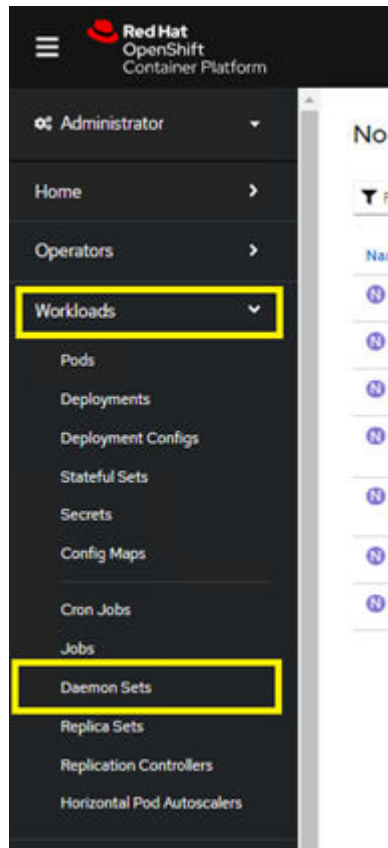
Name ↑	Kind ↓	Status ↓
 hspc-sc-ucs	HSPC	-

Node selector daemon set configuration

Node selector daemon configuration allows deployment of resources to the correct worker compute node, before following these steps verify that nodes within the cluster have been labeled as described in [Node labeling \(on page 58\)](#). To configure the HSPC daemon follow these steps:

Procedure

1. From the navigation tree, select **Workloads > Daemon Sets**.



- From the Daemon Sets window verify that you are in the correct **namespace**, and then select **hspc-csi-node**.



- Select the YAML tab.
- Within the YAML file define node selector as the following 2 lines:
nodeSelector:
node type: hspc-fc

Daemon Sets > Daemon Set Details

DS hspc-csi-node
Managed by **HSPC** hspc-sc-ucs

Details **YAML** Pods Environment Events

```

374 ownerReferences:
375   - apiVersion: csi.hitachi.com/v1
376     kind: HSPC
377     name: hspc-sc-ucs
378     uid: 34b2b671-bf73-4352-9512-cbe2655a3915
379     controller: true
380     blockOwnerDeletion: true
381 spec:
382   selector:
383     matchLabels:
384       app: hspc-csi-driver
385   template:
386     metadata:
387       creationTimestamp: null
388     labels:
389       app: hspc-csi-driver
390   spec:
391     nodeSelector:
392       node-type: hspc-fd
393     restartPolicy: Always
394     serviceAccountName: hspc-csi-sa
395     schedulerName: default-scheduler
396     hostNetwork: true
397     terminationGracePeriodSeconds: 30
398     securityContext: {}
399     containers:
400     - resources:
401       limits:

```

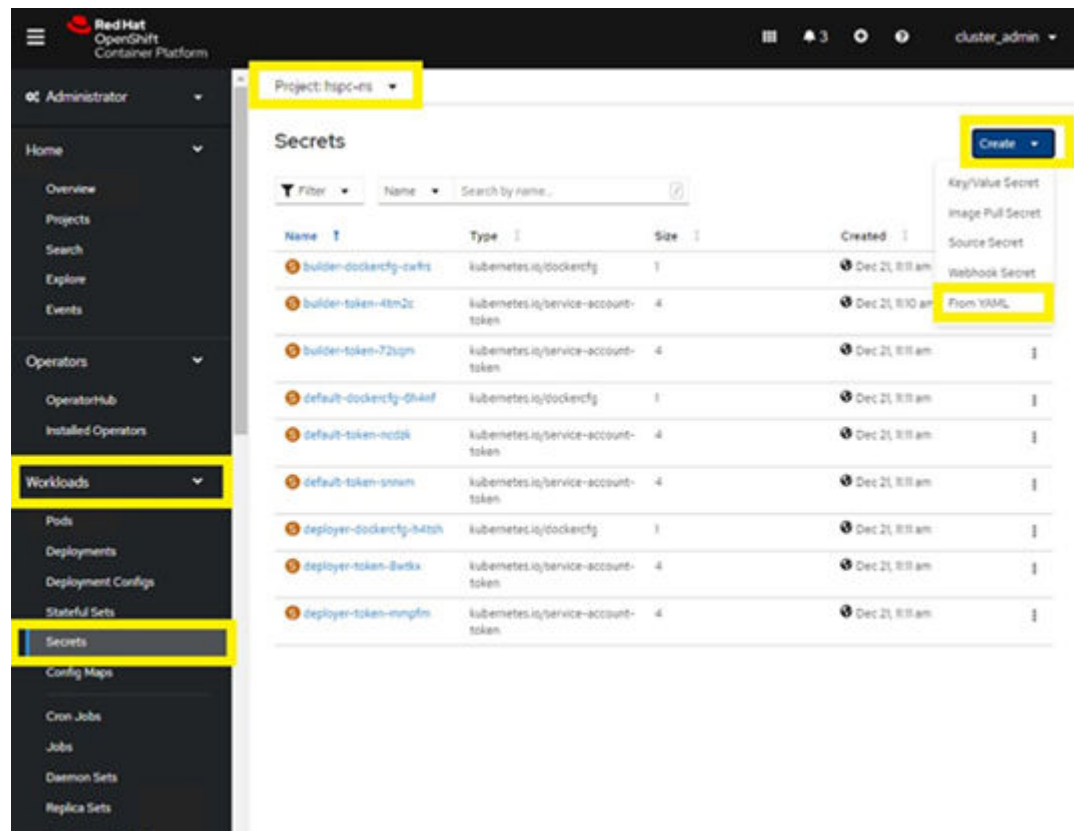
5. Click **Save**.
6. Acknowledge the Managed resources prompt and click **Save**.

Secret settings

The Secret file contains the storage URL, username, and password settings that are necessary for Storage Provider for VMware vCenter to work with your environment. From an available Linux machine admins must encode the URL, username, and password to access Hitachi Virtual Storage Platform (VSP). To create a secret file, follow these steps:

Procedure

1. From the navigation tree, select **Workloads > Secrets**.
2. Confirm the namespace in which you are creating the secret, in this example hspc-ns.
3. Select **Create**.
4. Click **From YAML**.



5. Enter the following parameters on the YAML file:
 - a. Define apiVersion as V1.
 - b. Define kind as Secret.
 - c. Select an applicable secret name.
 - d. Base64 encode the URL of the VSP storage system.

```
echo -n "http://172.16.1.1" | base64
```

- e. Base64 encode the username to log in to the VSP.

```
echo -n "User01" | base64
```

- f. Base64 encode the password associated with the VSP username.

```
echo -n "Password01" | base64
```

```
apiVersion: v1 (1)
```

```
kind: Secret (2)
```

```
metadata:
```

```
  name: secret-vsp5000 (3)
```

```
type: Opaque
```

```

data:

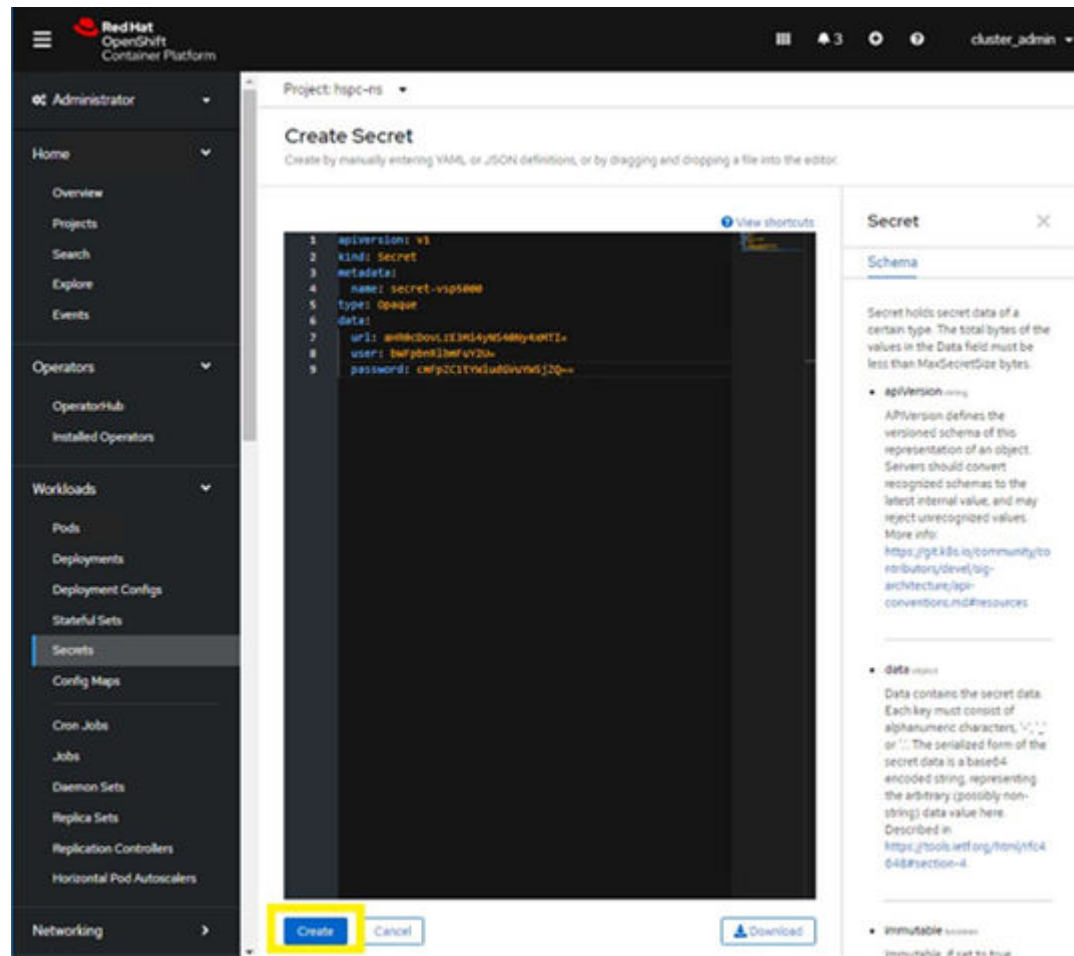
  url: aHR0cDovLzE3Mi4yNS40Ny4xMTI=      (4)

  user: bWFpbnRlbnFuY2U=                 (5)

  password: cmFpZC1tYWludGVuYW5jZQ==    (6)

```

6. Click **Create**.



HSPC usage

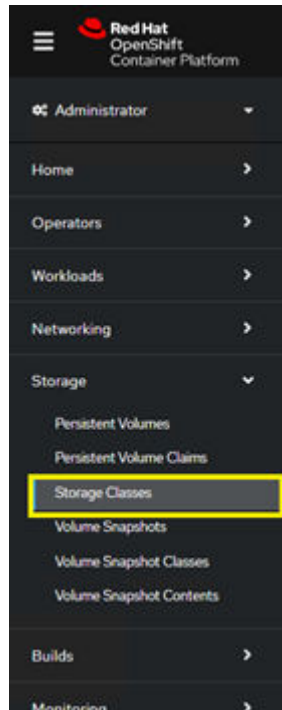
This section covers using HSPC to provide persistent storage to an OCP cluster.

StorageClass

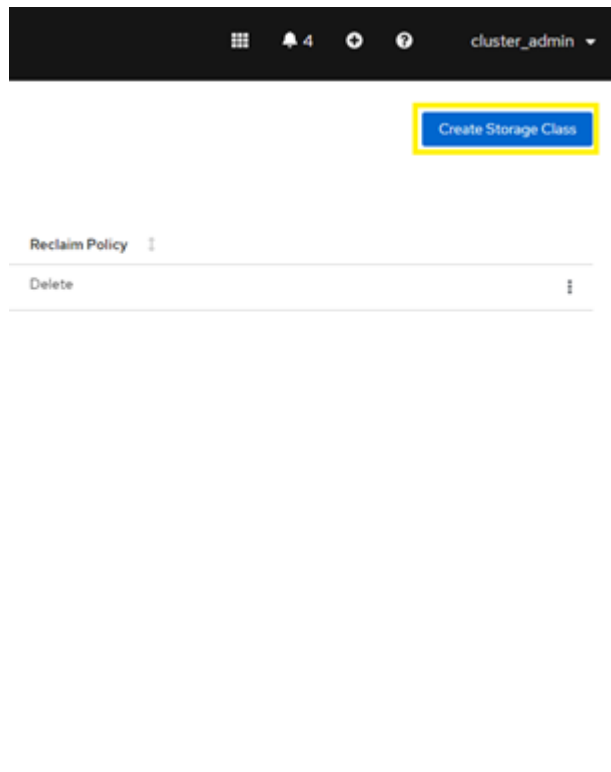
Before creating a StorageClass verify that VSP resources such as ports, parity groups, and pools have been allocated to meet your specifications. To create a HSPC StorageClass follow these steps:

Procedure

1. From the navigation tree, select **Storage > Storage Classes**.



2. Click **Create Storage Class**.



3. Click **Edit YAML**.

Create Storage Class Edit YAML

Name *

Description

Reclaim Policy *
Delete

Determines what happens to persistent volumes when the associated persistent volume claim is deleted. Defaults to 'Delete'

Provisioner *
Select Provisioner

Determines what volume plugin is used for provisioning persistent volumes.

Create Cancel

4. Enter the following parameters on the YAML file:
 - a. Define the API version used.
 - b. Define the *kind* parameter as StorageClass.
 - c. Select an applicable StorageClass name.
 - d. Define the provisioner used. For HSPC, the default is *hspc.csi.hitachi.com*.
 - e. Define the reclaimPolicy as *Delete*.
 - f. Define volumeBindingMode as *Immediate*.
 - g. Define allowVolumeExpansion as *true*.
 - h. Define the VSP serial number.
 - i. Define the pool on the VSP used to carve dynamically provided persistent volumes.
 - j. Define the VSP storage ports for allocation.
 - k. Define connection type as FC.
 - l. Define filesystem type as ext4.
 - m. Define VSP secret name.
 - n. Define operator namespace.

```

apiVersion: storage.k8s.io/v1      (1)
kind: StorageClass                (2)
metadata:
  name: sc-hspc-vsp5000          (3)
  annotations:

  kubernetes.io/description: Hitachi Storage Plug-in for Containers
  provisioner: hspc.csi.hitachi.com (4)
  reclaimPolicy: Delete          (5)

  volumeBindingMode: Immediate   (6)

```

```

allowVolumeExpansion: true      (7)

parameters:

serialNumber: "30595"          (8)

poolID: "6"                    (9)

portID : CL5-A,CL6-A,CL5-B,CL6-B  (10)

connectionType: fc             (11)

csi.storage.k8s.io/fstype: ext4   (12)

csi.storage.k8s.io/node-publish-secret-name: "secret-
vsp5000"                        (13)

csi.storage.k8s.io/node-publish-secret-namespace: "hspc-ns"  (14)

csi.storage.k8s.io/provisioner-secret-name: "secret-
vsp5000"                        (13)

csi.storage.k8s.io/provisioner-secret-namespace: "hspc-ns"  (14)

csi.storage.k8s.io/controller-publish-secret-name: "secret-
vsp5000"                        (13)

csi.storage.k8s.io/controller-publish-secret-namespace: "hspc-
ns"                              (14)

csi.storage.k8s.io/node-stage-secret-name: "secret-vsp5000"  (13)

csi.storage.k8s.io/node-stage-secret-namespace: "hspc-ns"  (14)

csi.storage.k8s.io/controller-expand-secret-name: "secret-
vsp5000"                        (13)

csi.storage.k8s.io/controller-expand-secret-namespace: "hspc-
ns"                              (14)

```

5. Click Create.

Create Storage Class

Create by manually entering YAML or JSON definitions, or by dragging and dropping a file into the editor.

```

1  apiVersion: storage.k8s.io/v1
2  kind: StorageClass
3  metadata:
4    name: sc-hpc-vsp5000
5  annotations:
6    kubernetes.io/description: Hitachi Storage Plug-in for Containers
7  provisioner: hpc.csi.hitachi.com
8  reclaimPolicy: Delete
9  volumeBindingMode: Immediate
10 allowVolumeExpansion: true
11 parameters:
12   serialNumber: "10555"
13   poolID: "s"
14   portID: CLS-A,CL6-A,CLS-B,CL6-B
15   connectionType: fc
16   csi.storage.k8s.io/ftype: ext4
17   csi.storage.k8s.io/node-publish-secret-name: "secret-vsp5000"
18   csi.storage.k8s.io/node-publish-secret-namespace: "hpc-ns"
19   csi.storage.k8s.io/provisioner-secret-name: "secret-vsp5000"
20   csi.storage.k8s.io/provisioner-secret-namespace: "hpc-ns"
21   csi.storage.k8s.io/controller-publish-secret-name: "secret-vsp5000"
22   csi.storage.k8s.io/controller-publish-secret-namespace: "hpc-ns"
23   csi.storage.k8s.io/node-stage-secret-name: "secret-vsp5000"
24   csi.storage.k8s.io/node-stage-secret-namespace: "hpc-ns"
25   csi.storage.k8s.io/controller-expand-secret-name: "secret-vsp5000"
26   csi.storage.k8s.io/controller-expand-secret-namespace: "hpc-ns"

```

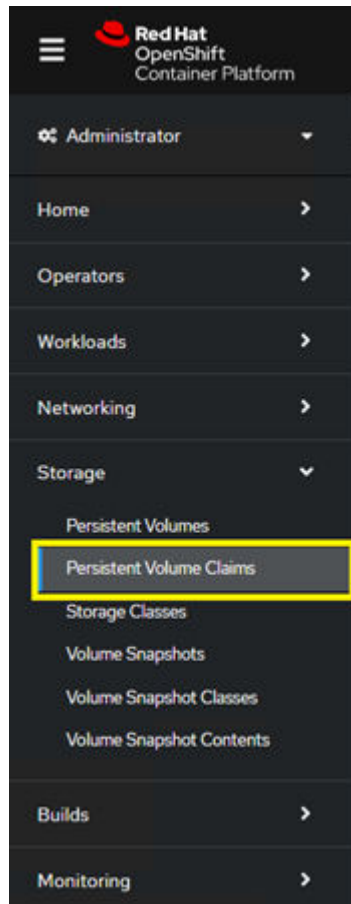
Buttons: Create, Cancel, Download

Persistent Volume Claim (PVC)

Before creating a PVC verify that an appropriate Secret and StorageClass exist for the VSP. Use the following procedure to deploy a PVC using HSPC:

Procedure

1. From the navigation tree, select **Storage > Persistent Volume Claims**.



2. Verify that you are using the correct namespace, then and click **Create Persistent Volume Claim**.



3. In the Create Persistent Volume Claim wizard, select the **Storage Class** that uses HSPC provisioner.
4. Define a **Persistent Volume Claim Name**.
5. Define the **Access Mode** parameters.
6. Enter an appropriate **Size**.
7. Click **Create**.

Project: hpc-ns

Create Persistent Volume Claim Edit YAML

Storage Class

 Storage class for the new claim

Persistent Volume Claim Name *

 A unique name for the storage claim within the project

Access Mode *
 Single User (RWX) Shared Access (RWX) Read Only (ROX)
 Permissions to the mounted drive

Size *
 GiB
 Desired storage capacity
 Use label selectors to request storage
 Use label selectors to define how storage is created

Viewing LDEV information using OCP

Once a PV has been bound, admins can view the associated YAML file to correlate PV IDs to VSP logical device designations. To view LDEV information follow these steps:

Procedure

1. From the navigation tree, select **Storage > Persistent Volumes**.
2. Select the deployed PV that is associated with the PVC that uses HSPC operator.

Name	Status	Claim
efs-registry-pv	Bound	pvc-registry-pvc
pvc-0ee33c2-5404-4089-86a2-c52a5cb83ee	Bound	pvc-hpc
pvc-0ed47ad-2521-4fac-91e1-f0eba5e75e55	Bound	pvc-ucs-hpc
pvc-0b4826d5-4224-463a-99c5-3e77900f6e87	Bound	pvc-ucs-hpc-snapshot2-restore
pvc-96ac91d1-7388-4563-b08f-bc2965e6c4b	Bound	maradb
pvc-9e23043e-6243-4b6b-9bc1-da73eec2b07	Bound	pvc-hpc-new

3. Click the **YAML** tab.
4. View YAML parameters for **ldevIDHex**.

Persistent Volumes > Persistent Volume Details

PV pvc-0eef31c2-5404-4089-86a2-c512a5cb83ee Bound Actions

Details **YAML**

```

80     "f:phase": {}
81   spec:
82     capacity:
83       storage: 10Gi
84     csi:
85       driver: hspc.csi.hitachi.com
86       volumeHandle: 60060e800677830000507783000009d3--spc-1d0038ef02
87       fsType: ext4
88       volumeAttributes:
89         hostModeOption: '91'
90         size: 10Gi
91         mode: normal
92         nickname: spc-1d0038ef02
93         autoHG: 'true'
94         ports: 'CL5-A,CL6-A,CL5-B,CL6-B'
95         ldevIDHex: '09:D3'
96       connectionType: fc
97       storage.kubernetes.io/csiProvisionerIdentity: 1610648289335-8081-hspc.csi.hitachi.co
98       ldevIDDec: '2515'
99       controllerPublishSecretRef:
100        name: secret-vsp5000

```

- The associated LDEV can be viewed using Hitachi Storage Navigator by selecting the Virtual Volumes tab of the associated pool used for HSPC.

LDEV ID	LDEV Name	Status	Emulation Type	Capacity			Used Capacity	
				Total	Reserved	Used	Used (%)	Tier 1
09:09:D3	spc-1d0038ef02	Normal	OPEN-V CVS	10.00 GB	0.00 GB	0.00 GB	0	-
09:09:EC	spc-6842ca0c1a	Normal	OPEN-V CVS	3.00 GB	0.00 GB	0.00 GB	0	-
09:09:ES	spc-2b19fae248	Normal	OPEN-V CVS	5.00 GB	0.00 GB	0.00 GB	0	-

Expanding PVs

HPSC allows volume expansion of deployed PVs. To expand the capacity of a deployed PV using OCP follow these steps. Before expanding a PV, confirm that the StorageClass has volume expansion parameter `allowVolumeExpansion: true`.

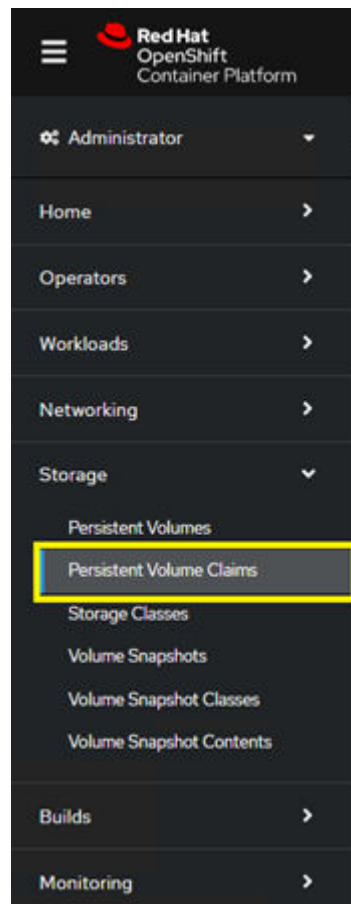
Note: The maximum additional capacity for volume expansion is 7 TB or pool capacity that backs the HSPC operator defined within StorageClass. The minimum size for expansion is 1Gb.

Note: PV capacity cannot be reduced.

Note: When expanding volume capacity there is no need to delete and redeploy Pods.

Procedure

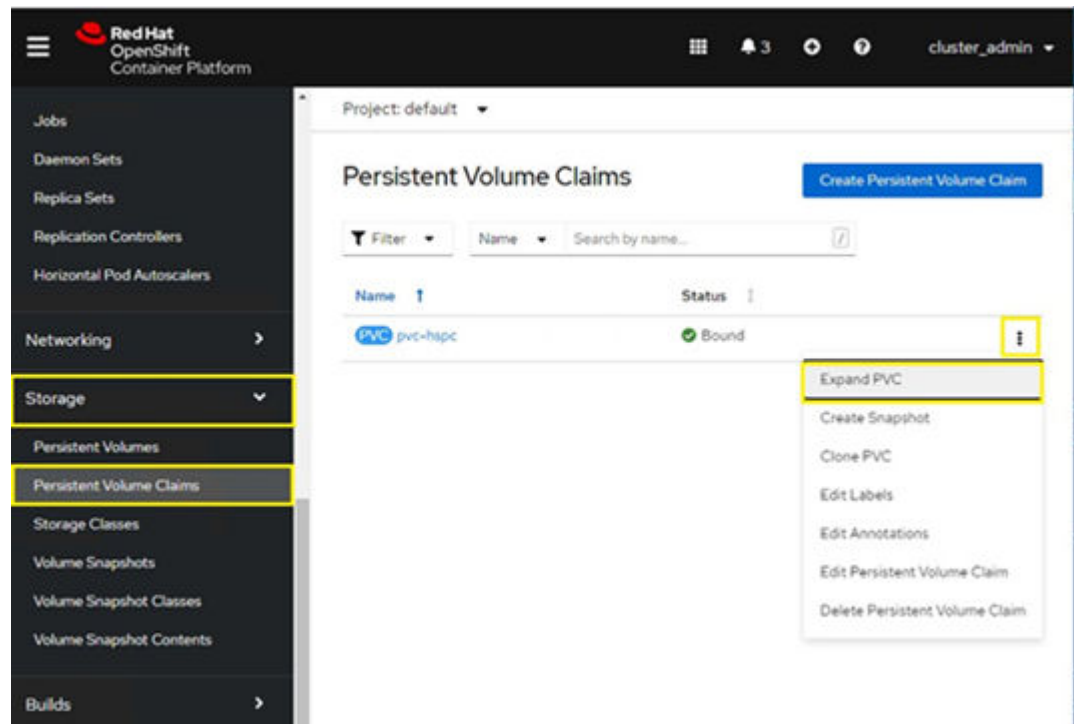
- From the navigation tree, select **Storage > Persistent Volume Claims**.



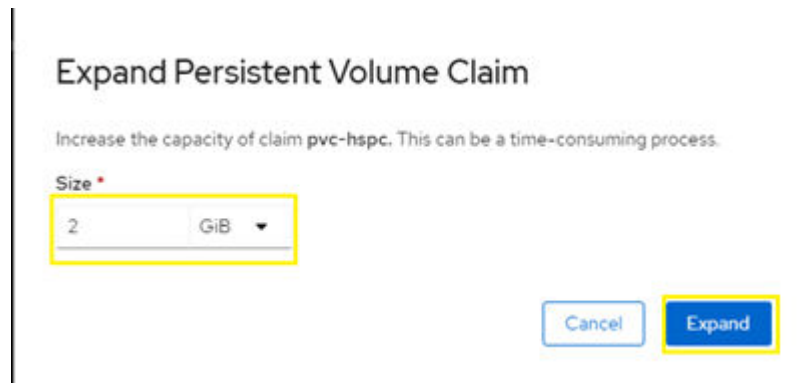
2. Verify that you are using the correct namespace.



3. Click the **ellipsis** icon associated with the PV that needs expansion.
4. Click **Expand PVC**.



5. In the **Expand Persistent Volume Claim** dialog box, enter the total capacity of the PV.
6. Click **Expand**.



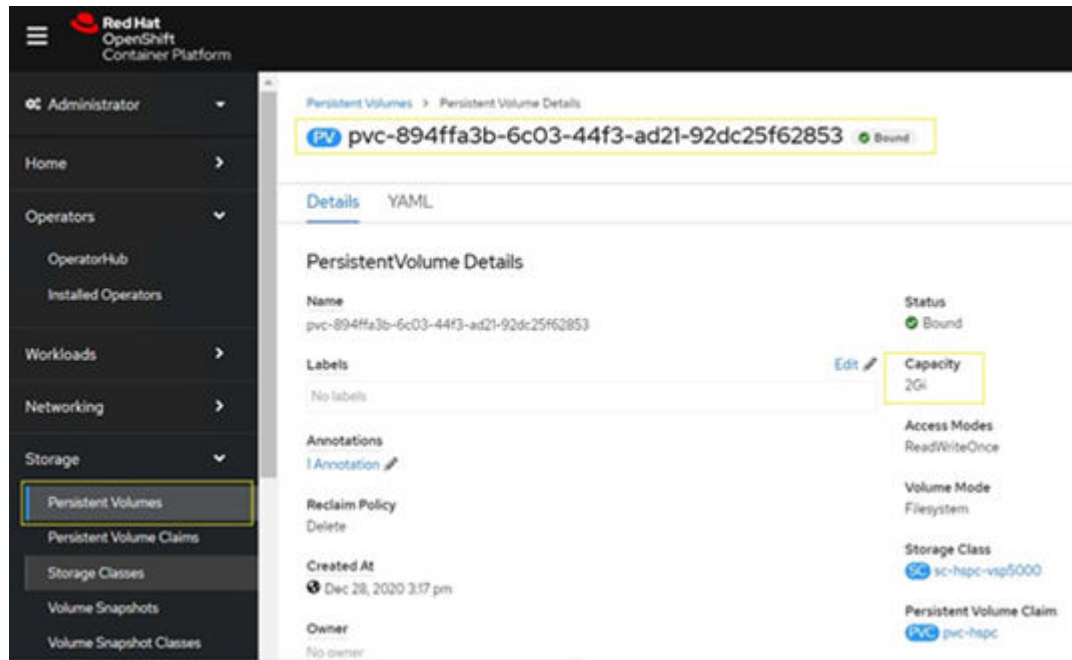
7. Once the PV is expanded, its status can be viewed from the OCP CLI using the following command: `oc get pv <PV_NAME>`.

```

[ocp@rhel119p1552a3-admin-va ~]$ oc get pv pvc-93d72a3b-4003-11e7-a021-92d021f428b3
NAME                CAPACITY  ACCESS MODES  RECLAIM POLICY  STATUS  CLAIM                STORAGECLASS  REASON  AGE
pvc-93d72a3b-4003-11e7-a021-92d021f428b3  2Gi      RWO           Delete          Bound   default/pvc-hspc    ec-hpco-vgp5000  21d
[ocp@rhel119p1552a3-admin-va ~]$

```

8. PV volumes can also be viewed from the OCP GUI by selecting the **Persistent Volume** tab from the navigation tree.




PV clone

This HSPC feature can create an exact clone of an existing PV within an OCP environment. To clone a PV using the OCP GUI follow these steps. If the volume is expanded, confirm for completion before starting. Also, before PV clone verify that StorageClass has been created and the PVC has been defined.

Procedure

1. From the navigation tree, select **Storage > Persistent Volume Claims**.



2. Verify that you are using the correct namespace, then and click **Create Persistent Volume Claim**.



3. On the Create Persistent Volume Claim window select **Edit YAML**.

Create Persistent Volume Claim

[Edit YAML](#)

Storage Class

SC sc-hspc-vsp5000

Storage class for the new claim

Persistent Volume Claim Name *

my-storage-claim

A unique name for the storage claim within the project

Access Mode *

Single User (RWO) Shared Access (RWX) Read Only (ROX)

Permissions to the mounted drive

Size *

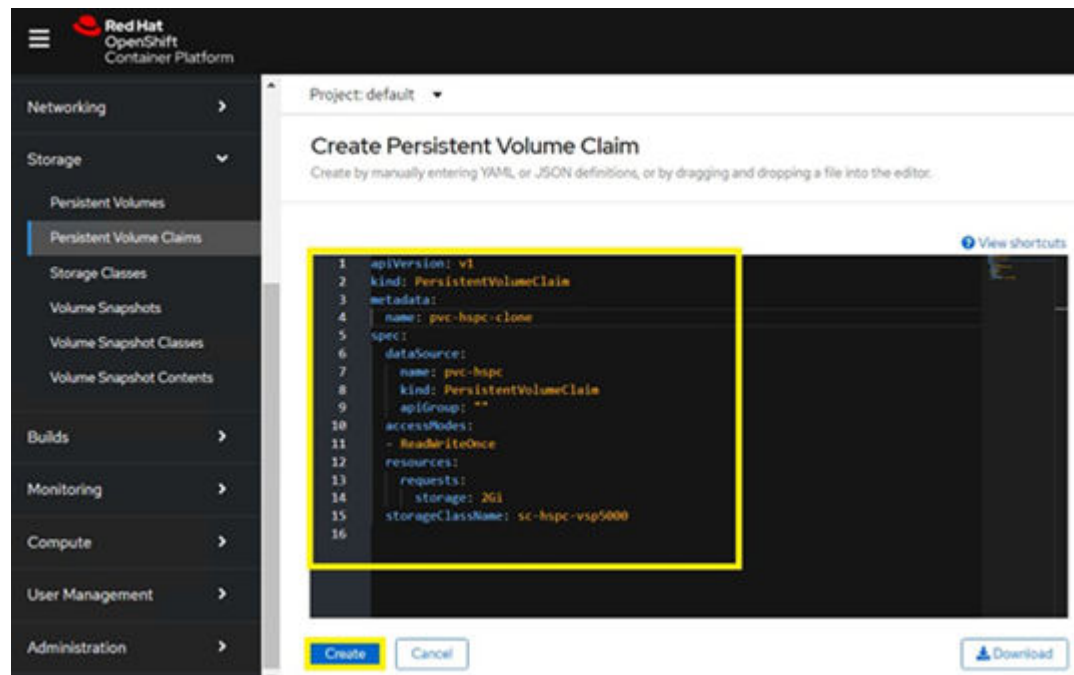
4. Enter the following parameters on the YAML file:
 - a. Define the API version used.
 - b. Define the *kind* parameter as PersistentVolumeClaim.
 - c. Select an applicable PVC clone name.
 - d. Define the name of the source PV.
 - e. Define source kind as PersistentVolumeClaim.
 - f. Define the exact clone storage capacity, verify that it is the same as the source PV.
 - g. Define StorageClass name.

```

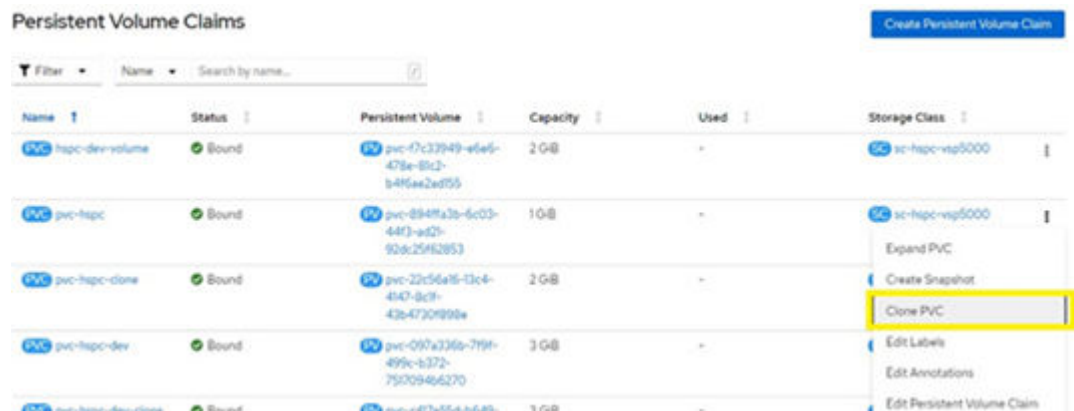
apiVersion: v1      (1)
kind: PersistentVolumeClaim    (2)
metadata:
  name: pvc-hspc-clone    (3)
spec:
  dataSource:
    name: pvc-hspc    (4)
    kind: PersistentVolumeClaim    (5)
  apiGroup: ""
  accessModes:
  - ReadWriteOnce
  resources:
  requests:
    storage: 2Gi    (6)
    storageClassName: sc-hspc-vsp5000    (7)

```

5. Click **Create**.



6. Alternatively, PVCs can be cloned by selecting the ellipsis icon on the Persistent Volume Claim list.
7. Click **Clone PVC**.



8. From the Clone wizard, click **Clone**.

Clone

Name *

pvc-hspc-clone

Size *

2

GiB ▾

PVC Details

Namespace

 default

Requested Capacity

2 GiB

Access Mode

Single User (RWO)

Storage Class

 sc-hspc-vsp5000

Used Capacity

-

Volume Mode

Filesystem

Cancel

Clone

PV snapshot

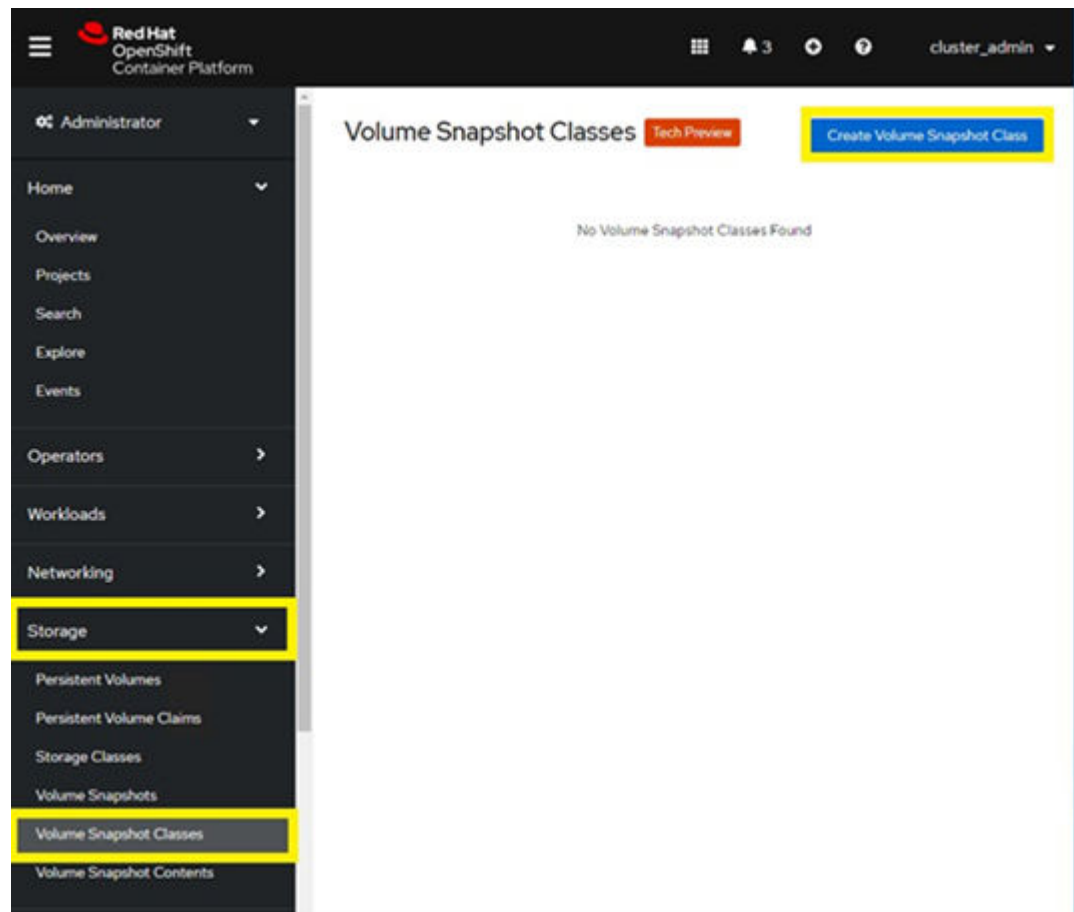
This HSPC feature creates a snapshot that is a point-in-time image of a volume. A snapshot can be used to duplicate the previous state of an existing volume.

PV Snapshot Classes

Before taking a PV snapshot, a Snapshot Class must be defined, to create a Snapshot Class follow these steps:

Procedure

1. From the navigation tree, select **Storage > Volume Snapshot Classes**.
2. Click **Create Volume Snapshot Class**.



3. Enter the following parameters on the YAML file:
 - a. Define the API version as `snapshot.storage.k8s.io/v1beta1`.
 - b. Define the `kind` parameter as `VolumeSnapshotClass`.
 - c. Select an applicable snapshot class name.
 - d. Define the driver as `hspc.csi.hitachi.com`.
 - e. Define deletion policy as `Delete`.
 - f. Define the VSP pool ID used for the snapshot.
 - g. Define the snapshotter secret name.
 - h. Define the snapshotter secret namespace.

```

apiVersion: snapshot.storage.k8s.io/v1beta1    (1)
kind: VolumeSnapshotClass                      (2)
metadata:
  name: snapshotclass-vsp5000                  (3)
  driver: hspc.csi.hitachi.com                 (4)
  deletionPolicy: Delete                       (5)
parameters:
  poolID: "6"                                  (6)
  csi.storage.k8s.io/snapshotter-secret-name: "secret-vsp5000" (7)
  csi.storage.k8s.io/snapshotter-secret-namespace: "default"   (8)

```

4. Click **Create**.

Create Volume Snapshot Class

Create by manually entering YAML, or JSON definitions, or by dragging and dropping a file into the editor.

```

1  apiVersion: snapshot.storage.k8s.io/v1beta1
2  kind: VolumeSnapshotClass
3  metadata:
4    name: snapshotclass-vsp5000
5  driver: hspc.csi.hitachi.com
6  deletionPolicy: Delete
7  parameters:
8    poolID: "6"
9  csi.storage.k8s.io/snapshotter-secret-name: "secret-vsp5000"
10 csi.storage.k8s.io/snapshotter-secret-namespace: "default"

```

Buttons: Create, Cancel, Download

Creating PV snapshots

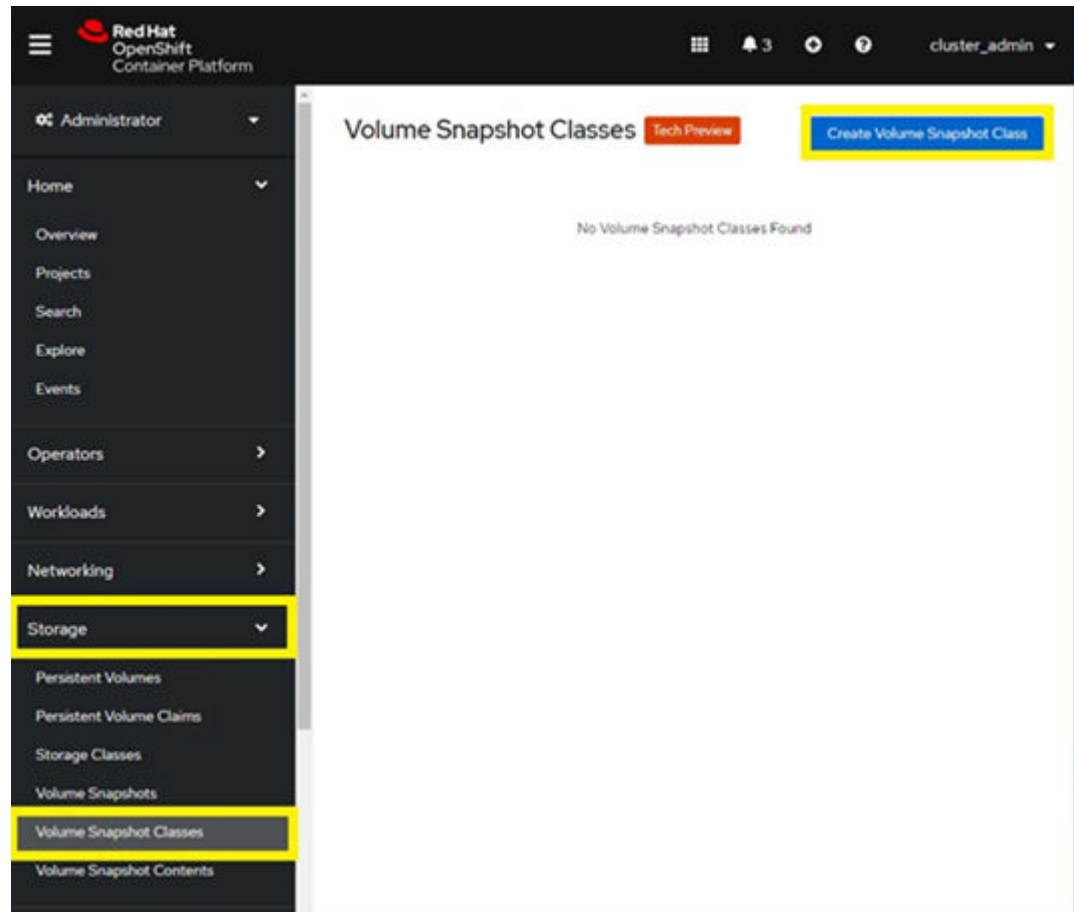
To create a snapshot of an existing PV using OCP, follow these steps.



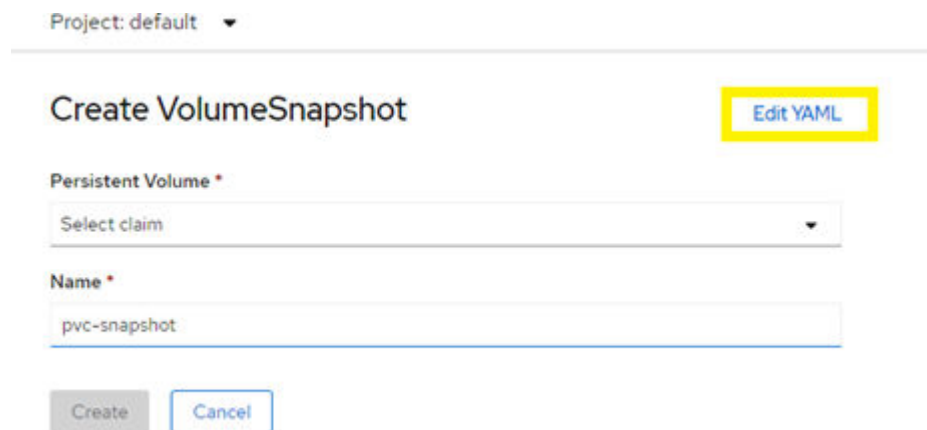
Note: Before creating a snapshot, verify that a Volume Snapshot Class is created.

Procedure

1. From the navigation tree, select **Storage > Volume Snapshots**.
2. Click **Create Volume Snapshot**.



3. From the Create VolumeSnapshot wizard, select **Edit YAML**.



4. Enter the following parameters on the YAML file:
 - a. Define the API version as `snaphost.storage.k8s.io/v1beta1`.
 - b. Define the `kind` parameter as `VolumeSnapshot`.
 - c. Select an applicable snapshot name.
 - d. Define the `volumeSnapshotClassName`.

- e. Define the source PVC to take a point in time snapshot.

```

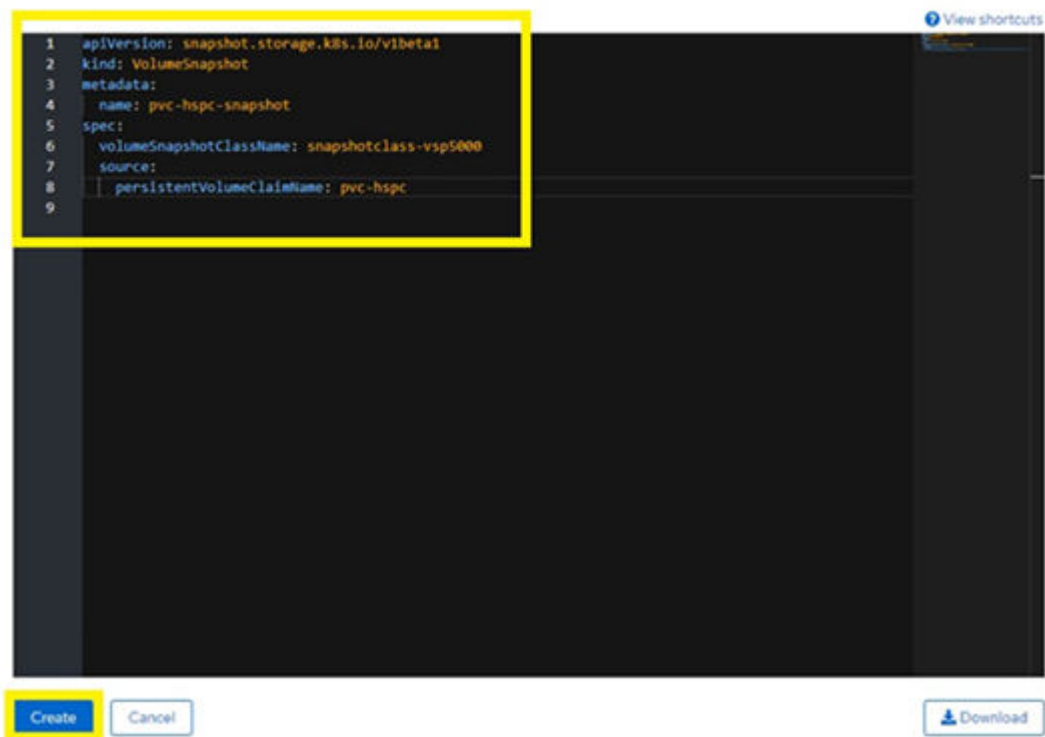
apiVersion: snapshot.storage.k8s.io/v1beta1    (1)
kind: VolumeSnapshot    (2)
metadata:
  name: pvc-hspc-snapshot    (3)
spec:
  volumeSnapshotClassName: snapshotclass-vsp5000    (4)
  source:
    persistentVolumeClaimName: pvc-hspc    (5)

```

5. Click **Create**.

Create Volume Snapshot

Create by manually entering YAML or JSON definitions, or by dragging and dropping a file into the editor.



6. Alternatively, PVCs can be snapshotted by selecting the **ellipsis** icon in the Persistent Volume Claim list.
7. Click **Create Snapshot**.

Persistent Volume Claims Create Persistent Volume Claim

Filter Name Search by name...

Name	Status	Persistent Volume	Capacity	Used	Storage Class
hspc-dev-volume	Bound	pvc-f7c33949-e6e6-478e-8c2-b46ae2ad855	2 GiB	-	sc-hspc-vsp5000
pvc-hspc	Bound	pvc-894ff42b-6c03-44f3-ad21-92dc25f62853	1 GiB	-	sc-hspc-vsp5000
pvc-hspc-clone	Bound	pvc-22c56a16-03e4-4e47-8c7f-43b4730f898x	2 GiB	-	sc-hspc-vsp5000
pvc-hspc-clone2	Bound	pvc-423ca4c8-c958-48f5-b20f-40c8c452e8c9	2 GiB	-	sc-hspc-vsp5000
pvc-hspc-dev	Bound	pvc-097a336b-79f1-	3 GiB	-	sc-hspc-vsp5000

Actions for pvc-hspc: Expand PVC, **Create Snapshot**, Clone PVC, Edit Labels, Edit Annotations, Edit Persistent Volume Claim

8. Enter a **Name**.
9. Select the appropriate **Snapshot Class**.
10. Click **Create**.

Create VolumeSnapshot Edit YAML

Creating snapshot for claim **pvc-hspc**

Name *
pvc-hspc-snapshot

Snapshot Class *
 snapshotclass-vsp5000

Create **Cancel**

Persistent Volume Claim Details

Name
 pvc-hspc

Namespace
 default

Status
 Bound

Storage Class
 sc-hspc-vsp5000

Requested Capacity
2 GiB

Access Mode
Single User (RWO)

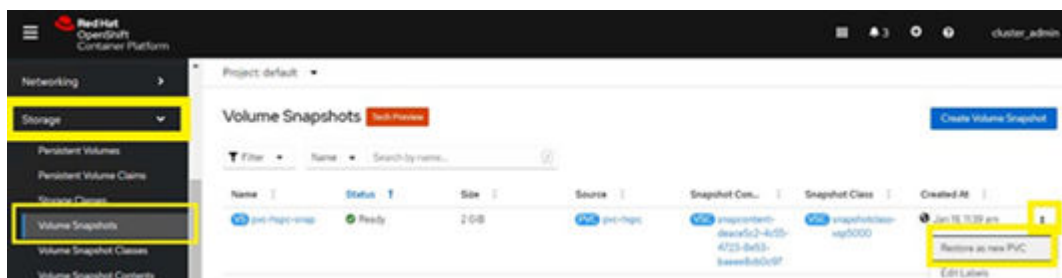
Volume Mode
Filesystem

PV snapshot restore

To restore a point in time snapshot using OCP follow these steps:

Procedure

1. From the navigation tree, select **Storage > Volume Snapshots**.
2. Select the **ellipsis** icon next to the snapshot to be restored.
3. Click **Restore as new PVC**.



4. Enter a **Name**.
5. Select an appropriate **Storage Class**.
6. Verify the snapshot **capacity**.
7. Click **Restore**.

Restore as new PVC

When restore action for snapshot `pvc-hspc-snap` is finished a new crash-consistent PVC copy will be created.

Name *

Storage Class

Size *

Volume Snapshot details

Created At 🕒 Jan 19, 11:39 am	Namespace 🏠 NS default
Status ✅ Ready	API Version snapshot.storage.k8s.io/v1beta1
Size 2 GiB	

Deleting PVCs

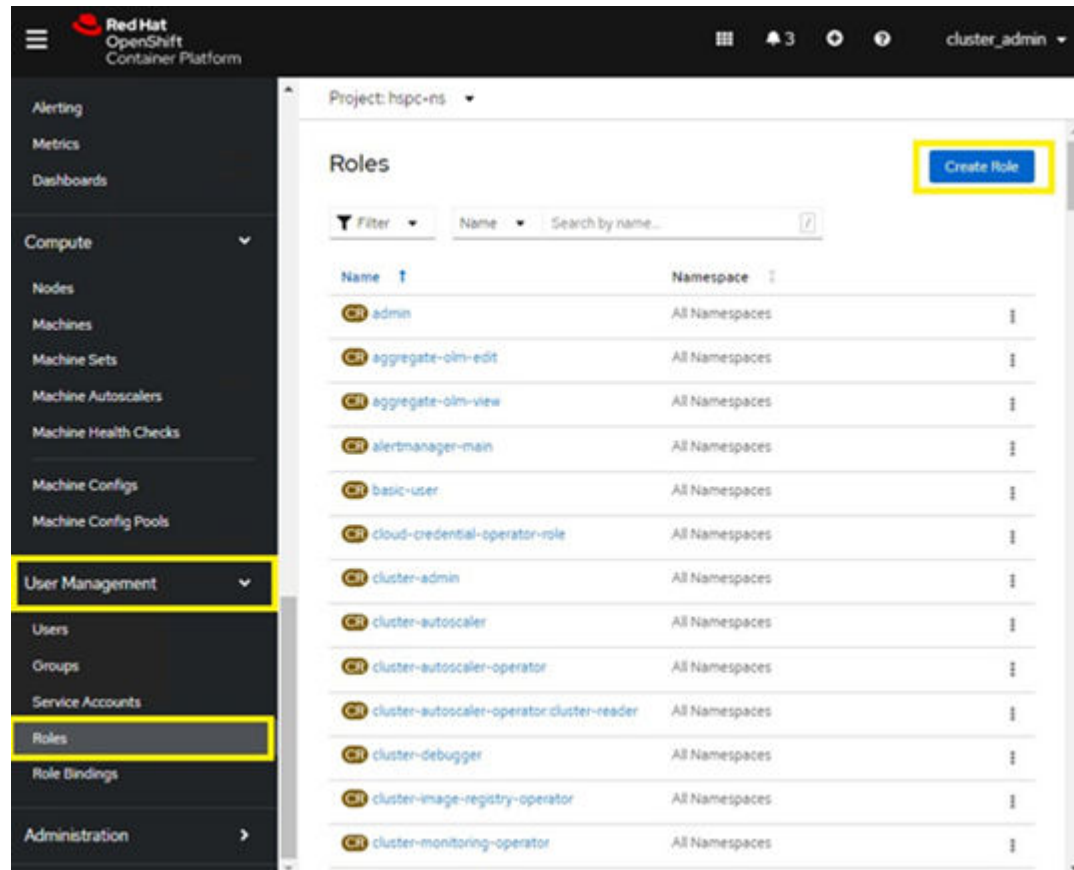
To delete PVCs, see [Deleting PVCs \(on page 56\)](#).

Roles

Admins can use the Role function within OCP to provide volume access to users. Once a role has been defined it can be bound to a specific namespace which consists of a user or given multiuser access using ClusterRoles. To create a role within OCP follow these steps:

Procedure

1. From the navigation tree, select **User Management > Roles**.
2. Click **Create Role**.



3. Enter the following parameters on the YAML file:
 - a. Define the API version as `rbac.authorization.k8s.io/v1`.
 - b. Define the `kind` parameter as `Role`. Alternatively, this can be defined as `ClusterRole` to provide cross-namespace access.
 - c. Define the Role name.
 - d. Define the associated namespace for the role.
 - e. Define the `resourceName` for the PVC that will use the role. If `resourceName` is left undefined the role will apply to all resources within the namespace.

```

apiVersion: rbac.authorization.k8s.io/v1    (1)
kind: Role    (2)
metadata:
  name: DevRoleAccess    (3)
  namespace: hspc-ns    (4)

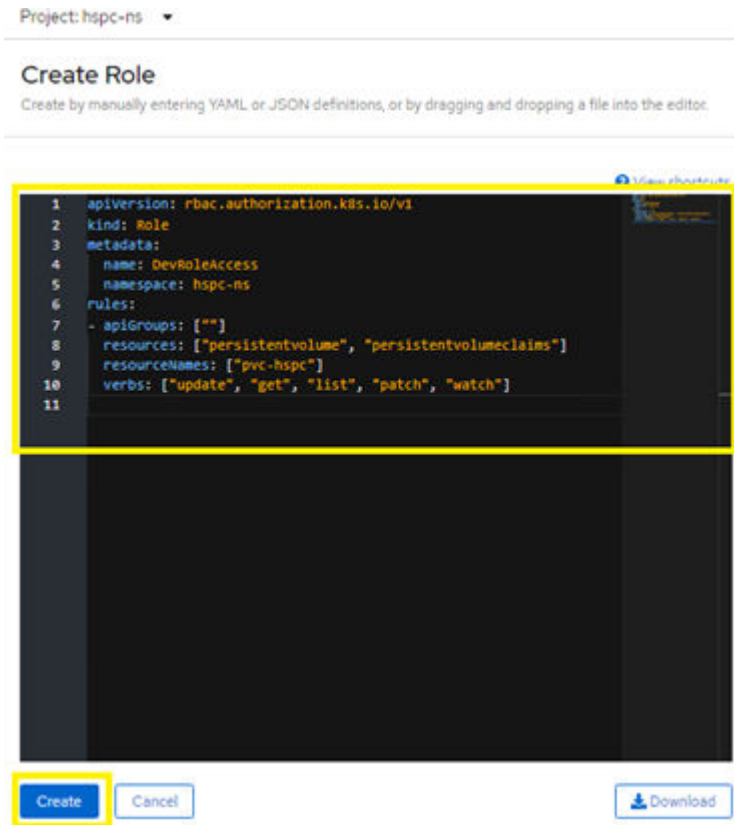
```

```

rules:
- apiGroups: [""]
resources: ["persistentvolume", "persistentvolumeclaims"]
resourceNames: ["pvc-hspc"]      (5)
verbs: ["update", "get", "list", "patch", "watch"]

```

4. Click Create.

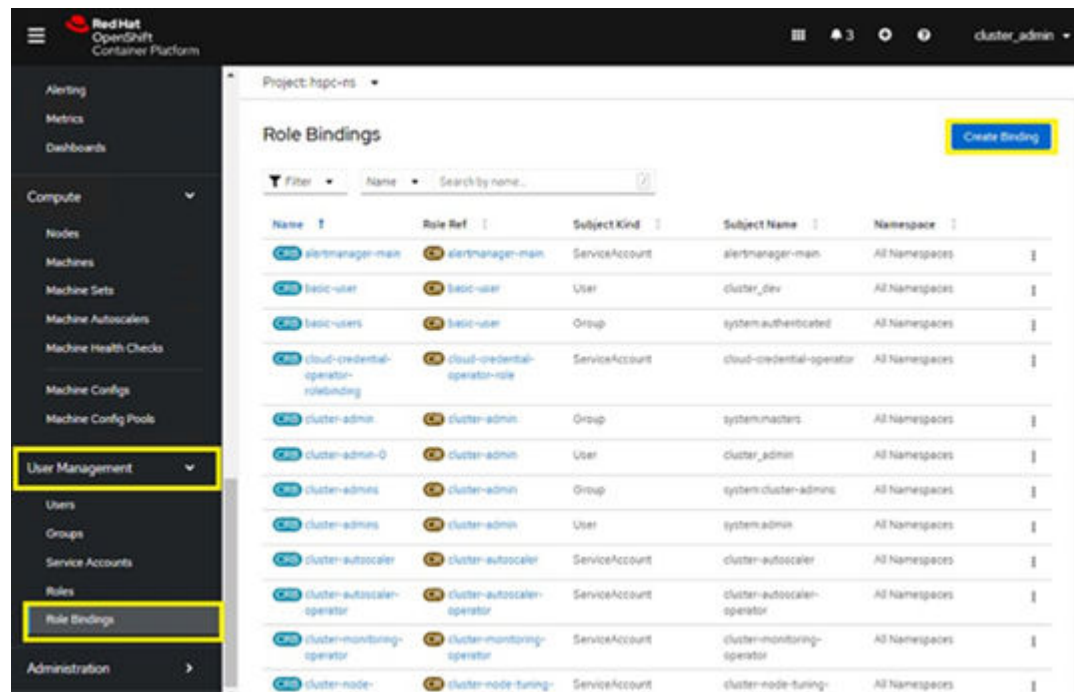


Role bindings

Once a role has been defined it needs to be bound to a user. To create a role binding follow these steps:

Procedure

1. From the navigation tree, select **User Management > Role Bindings**.
2. Click **Create Role**.



3. Define the Role Binding **Name**.
4. Select **Role**.
5. Select a **User**, **Group**, or **Service** account based on desired binding.
6. Click **Create**.

Project: hspc-ns ▾

Create Role Binding

Associate a user/group to the selected role to define the type of access and resources that are allowed.

Role Binding

Name *

Namespace *
 hspc-ns

Role

Role Name *

Subject

User
 Group
 Service Account

Subject Name *

Related Documents

This guide references the following documentation:

- [Hitachi Virtual Storage Platform Documentation](#)
- [Hitachi Storage Provider for VMware vCenter Documentation](#)
- [Red Hat OpenShift Container Platform Documentation](#)
- [Cisco and Hitachi Adaptive Solutions for Converged Infrastructure Design Guide](#)
- [Cisco and Hitachi Adaptive Solutions for Converged Infrastructure Deployment Guide](#)
- [Hitachi Virtual Storage Platform with Cisco Intersight Reference Architecture Guide](#)

Solution References

For more information on Hitachi solutions and products, go to Hitachi Vantara.com and see the following solution references.

Network

- [Cisco Nexus 9000 Series Switches Data Sheets](#)
- [Cisco MDS 9000 Series Multilayer Switches](#)

Compute

- [Cisco Unified Computing](#)
- [Cisco UCS 6400 Series Fabric Interconnects Data Sheet](#)
- [Cisco UCS 5100 Series Blade Server Chassis Data Sheet](#)
- [Cisco UCS VIC 1440 Adapter Data Sheet](#)
- [Cisco UCS Manager](#)

Storage

- [Hitachi Virtual Storage Platform 5000 Series](#)
- [Hitachi Virtual Storage Platform F Series All-Flash Enterprise Cloud Solutions](#)
- [Hitachi Virtual Storage Platform G Series Hybrid-Flash Midrange Cloud Solutions](#)

Virtualization Layer

- [VMware vCenter Server](#)
- [VMware vSphere](#)

Compatibility Matrixes

- [Hitachi Interoperability Reports](#)
- [VMware Compatibility Guide](#)
- [Cisco UCS Hardware and Software Compatibility](#)

Getting Help

Hitachi Vantara Support is the destination for technical support of products and solutions sold by Hitachi Vantara.

- To contact technical support, log on to Hitachi Vantara Support Connect for contact information see [Customer Contact Us](#) .
- To open a new support case, see [How to Create a New Case on the Support Website](#).

Hitachi Vantara



Corporate Headquarters
2535 Augustine Drive
Santa Clara, CA 95054 USA
HitachiVantara.com | community.HitachiVantara.com

Contact Information
USA: 1-800-446-0744
Global: 1-858-547-4526
HitachiVantara.com/contact