HITACHI
Inspire the Next

Cisco and Hitachi Adaptive Solutions with VMware Tanzu Basic

Implementation Guide

MK-SL-228-01
February 2022

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

	Changes	Date
MK-SL-228-01	Updated with information about Cisco Intersight Capabilities with Hitachi Virtual Storage Platform.	February 21, 2022
MK-SL-228-00	Initial release.	June 30, 2021

Implementation Guide

Use Hitachi Vantara 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 Vantara 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 VMware Tanzu solution provides highly available, predictable, and high-performance infrastructures for container applications built on top of the industry's leading virtualization platform to provide customers the ability to transform their datacenters into the container-based model.

This solution feature set is comprised of VMware vSphere 7.0u2 built on top of the Cisco and Hitachi Adaptive Solution for CI with Tanzu workload management enablement to provide customers a familiar UI to orchestrate container development cycles with on-premises Kubernetes running on VMware vSphere. Enterprise VSP storage is provided by Hitachi Storage Provider for VMware vCenter to allocate persistent storage for containerized applications through Storage Policy Based Management (SPBM) backed by either virtual volumes (vVols) or Virtual Machine File System (VMFS). This document describes the implementation and best practices of Hitachi storage resources to provide persistent storage to VMware Tanzu deployment models using Cloud Native Storage (CNS).

The following figure provides a capability overview of VMware 7.0u2 with Tanzu backed by Hitachi VSP storage on top of Cisco and Hitachi Adaptive Solutions for CI.

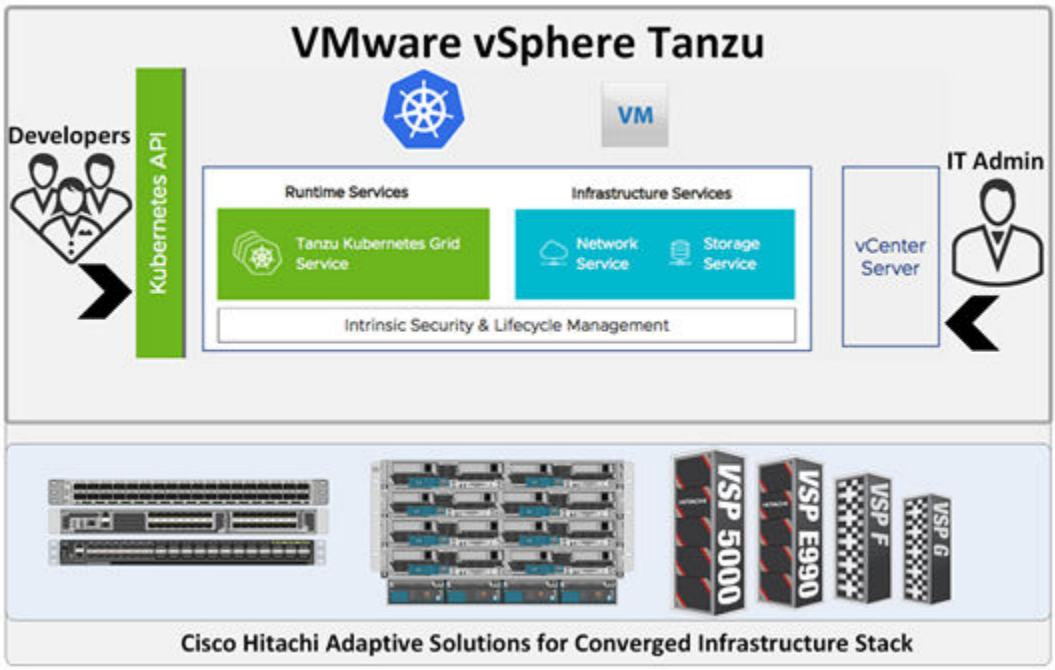

Note: Testing of these procedures was in a lab environment. Many things impact production environments beyond prediction or duplication in a lab environment. Follow the recommended practice of conducting proof-of-concept testing for acceptable results in a non-production, isolated test environment that 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 (VSP), see Related Documents (on page 82).

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 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.

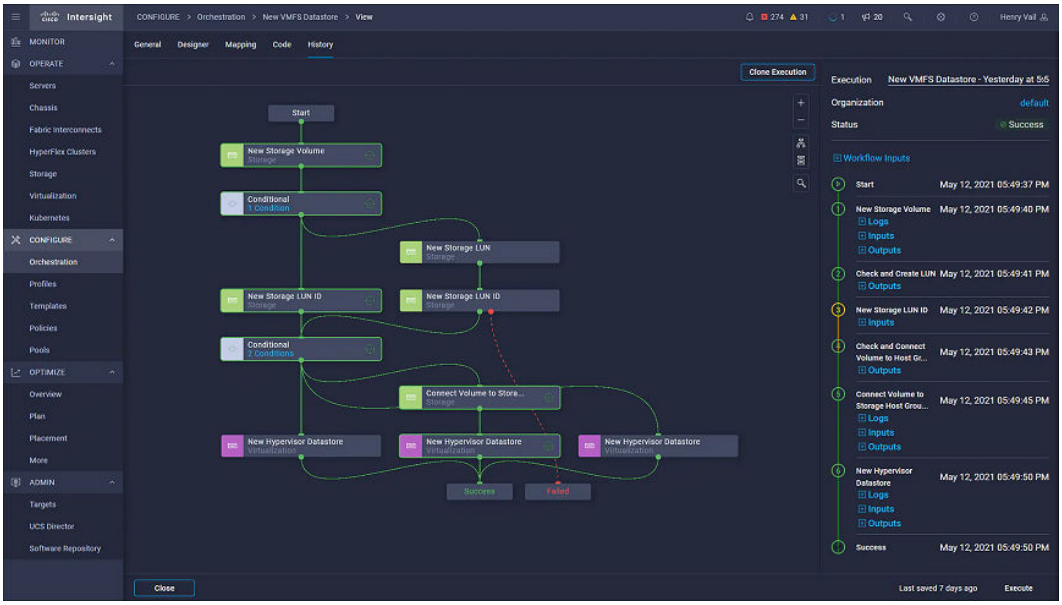

The following tables show the current capabilities of Hitachi Virtual Storage Platform (VSP) in orchestration with tasks and workflows provided by Intersight to end users.

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.

To compliment these capabilities Hitachi Vantara, in conjunction with Japan's ITpro division, have released another reference architecture (RA) guide in the series, Cisco and Hitachi Adaptive Solutions: Leveraging Hitachi Storage with Cisco Intersight for Consolidated Management and Automation at https://www.hitachivantara.com/en-us/pdf/architecture-guide/cisco-adaptive-solutions-leveraging-storage-with-cisco-intersight-for-consolidated-management-automation.pdf.

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.

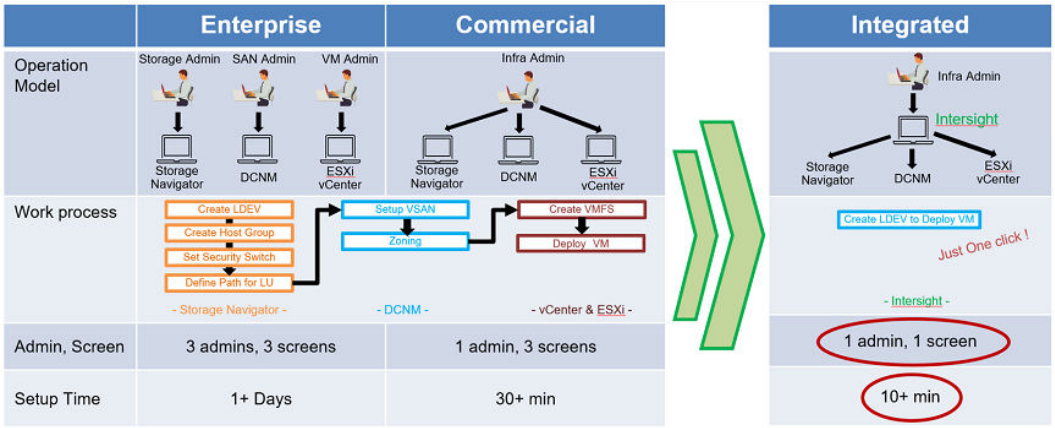

Cisco Unified Compute System Environment

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

This architecture uses a Hitachi VSP connected to Cisco MDS multilayer switches that link to the Cisco UCS Fabric Interconnects and Cisco 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 CI. The red lines represent fabric A connections, the blue lines represent fabric B connections, and the rest are port channel connections.

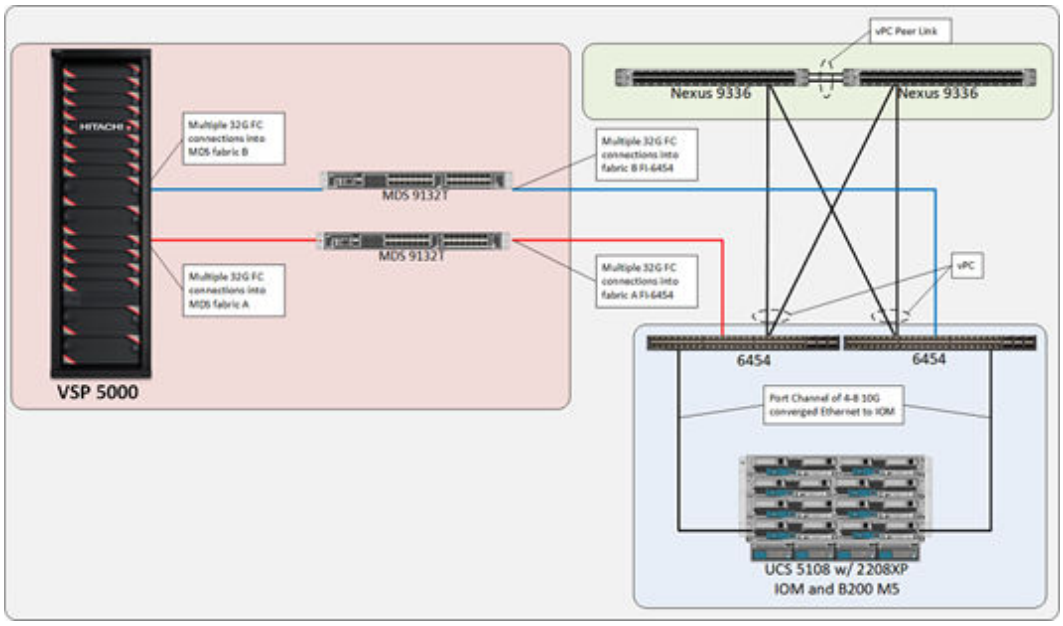

Hardware versions

This section 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 5000 series	90-06-01-00/00
Cisco MDS 9132T Fibre Channel switch	8.4(2b)
Cisco Nexus 9332-FX2 switch	NXOS 7.0(3)I7(9)
Cisco Fabric Interconnect 6454	4.1(3b)
Cisco Unified Computing System B200 M5 Blade Servers	4.1(3b)
Cisco Unified Computing System 2208XP IOM	4.1(3b)

Software versions

This section lists the software used in this solution.

Table 4 Software Versions Used for Validation

Component	Version
Hitachi Storage Provider for VMware vCenter	3.5.9 or Newer
vSphere Container Storage Interface (CSI)	2.0 or Newer
VMware vCenter Standalone (VCSA) 7.0 U2	7.0.2.00000
VMware ESXi 7.0 U1 Cisco Custom Image	7.0.1.16850804
VMware ESXi 6.7U3 nenic	1.0.35.0 or Newer
VMware ESXi 6.7U3 nfnic	4.0.0.63 or Newer
UCP Advisor	3.10

Hitachi Vantara Storage Operability with VMware Tanzu Container Management Platform

This section covers the VMware Tanzu storage supported by Hitachi storage to provide persistent storage.

The following table shows Hitachi integration points with VMware Tanzu. A Fibre Channel deployment type backed by CNS storage is covered in this guide.

Table 5 Hitachi integration points with VMware Tanzu

VMware Tanzu Container Management Platform		
Deployment Type	**Storage Type**	**Hitachi Persistent Storage Provider Compatibility**
VMware Tanzu	iSCSI	Hitachi Storage Plug-in for Containers
	FC - Cloud Native Storage (CNS)	Container Storage Interface (vVol + VMFS)

Solution components

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

VMware Tanzu

VMware Tanzu is a container management platform that allows datacenter admins and development teams to build, run, and manage Kubernetes controlled container-based applications all from a single and familiar vSphere UI. VMware provides three VMware Tanzu deployment types: Basic, Standard, and Advanced. In this implementation Tanzu Basic is implemented which allows users to run Kubernetes natively on premises within vSphere.

To learn more about Tanzu and various deployment types see Related Documents (on page 82) section.

Load Balancing

VMware Tanzu Basic is built on top of native vSphere networking which utilizes vSphere Distributed Switch (vDS). This deployment method requires admins to deploy a load balancer to carry traffic between vCenter and the Kubernetes supervisor control plane. HAProxy is a supported load balancer that can be deployed via an OVF file. During deployment admins complete a configuration wizard to supply corresponding vDS port groups for traffic routing. Management and workload network configurations are required, as well as an optional frontend network. To learn more about load balancing and HAProxy see Related Documents (on page 82).

The following figure shows the HAProxy network architecture.

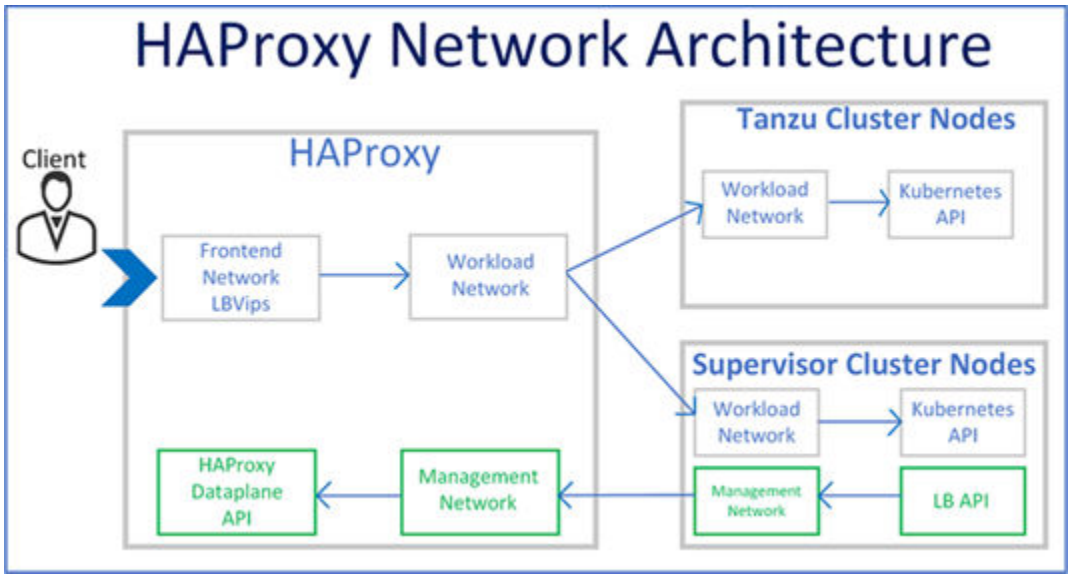

vSphere CNS

vSphere 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 and backed by Hitachi VSP storage via CNS. 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 persistent volume (PV) creation, admins can review 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 paravirtual CSI (pvCSI) is the version of the vSphere CNS-CSI driver modified for Tanzu Kubernetes clusters. The pvCSI resides in the Tanzu Kubernetes cluster and is responsible for all storage-related requests originating from the Tanzu Kubernetes cluster. The requests are delivered to the CNS-CSI, which then propagates them to CNS in vCenter Server. As a result, the pvCSI does not have direct communication with the CNS component, but instead relies on the CNS-CSI for any storage provisioning operations.

Unlike the CNS-CSI, the pvCSI does not require infrastructure credentials. It is configured with a service account in the supervisor namespace.

The following figure highlights the interaction between the CSI/pvCSI, CNS control plane, and vSphere storage.

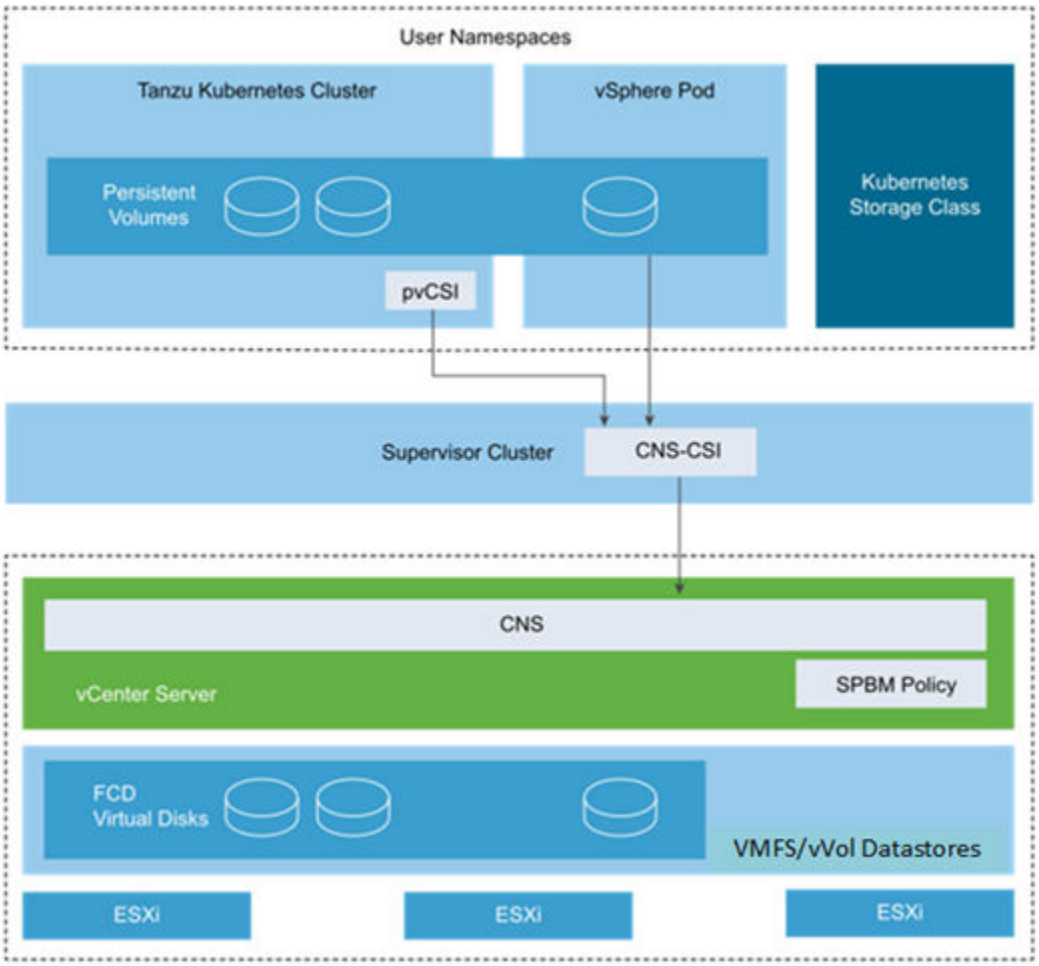

Hitachi Storage Provider for VMware vCenter

Hitachi 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. VASA makes this possible in two ways:

- VMware vSphere vVols

  This function is the VASA Provider component of VMware vVols that 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 VASA Provider 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.

Administrators define StorageClass settings that point to their respective VM storage policies, backed by either Hitachi vVols or VMFS storage that uses SPBM. 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.

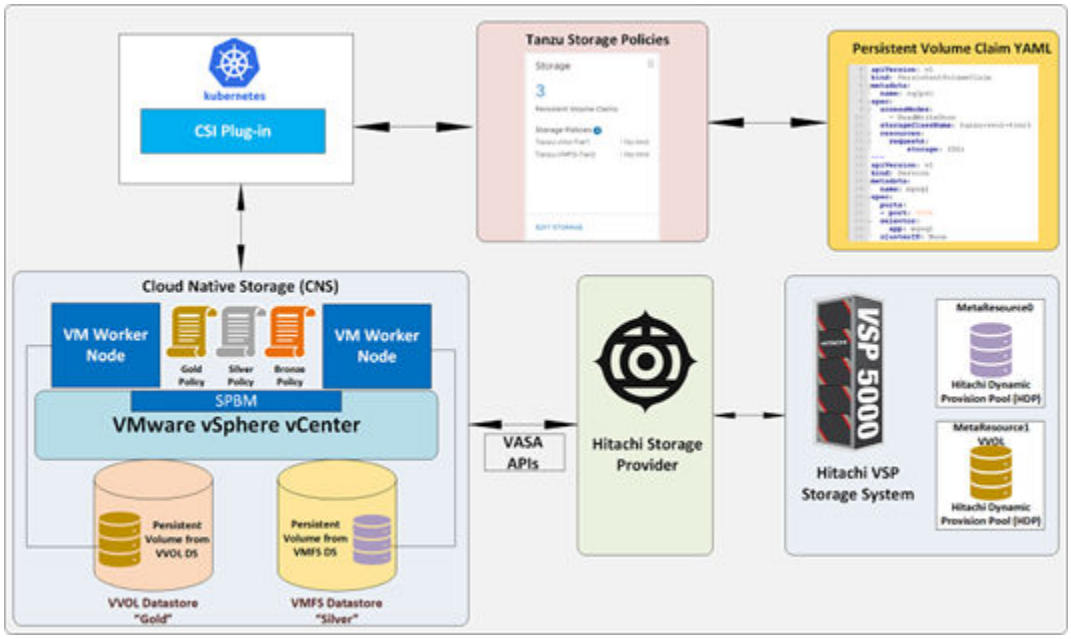

Hitachi Unified Compute Platform Advisor (Optional)

Hitachi Unified Compute Platform (UCP) Advisor is a single pane of glass management tool for converged infrastructure, providing automation for compute, network, and storage infrastructure. UCP Advisor within a UCS environment is a storage-only management tool for VMware-based virtual environments through the native vSphere web client. UCP Advisor has features that allow VMware administrators to manage native Hitachi storage through the vSphere UI and provide the capability of managing multiple VSP storage systems with a single instance.

UCP Advisor also provides deep integration with VMware management software improving administrator productivity with intuitive and intelligent operations and automation. It complements VMware vRealize software to further streamline the administration and automation of software-defined data center (SDDC). Automated workflows deliver IT agility using UCP Advisor REST APIs and vRealize Orchestrator and when used with vRealize Automation, enable self- services multi-cloud environments.

See Appendix A: UCP Advisor Storage Administration (on page 75) to review VSP pool creation as well as datastore allocation using UCP Advisor.

The following figure shows the relationship between UCP Advisor and vSphere with a Tanzu stack backed by Hitachi VSP storage.

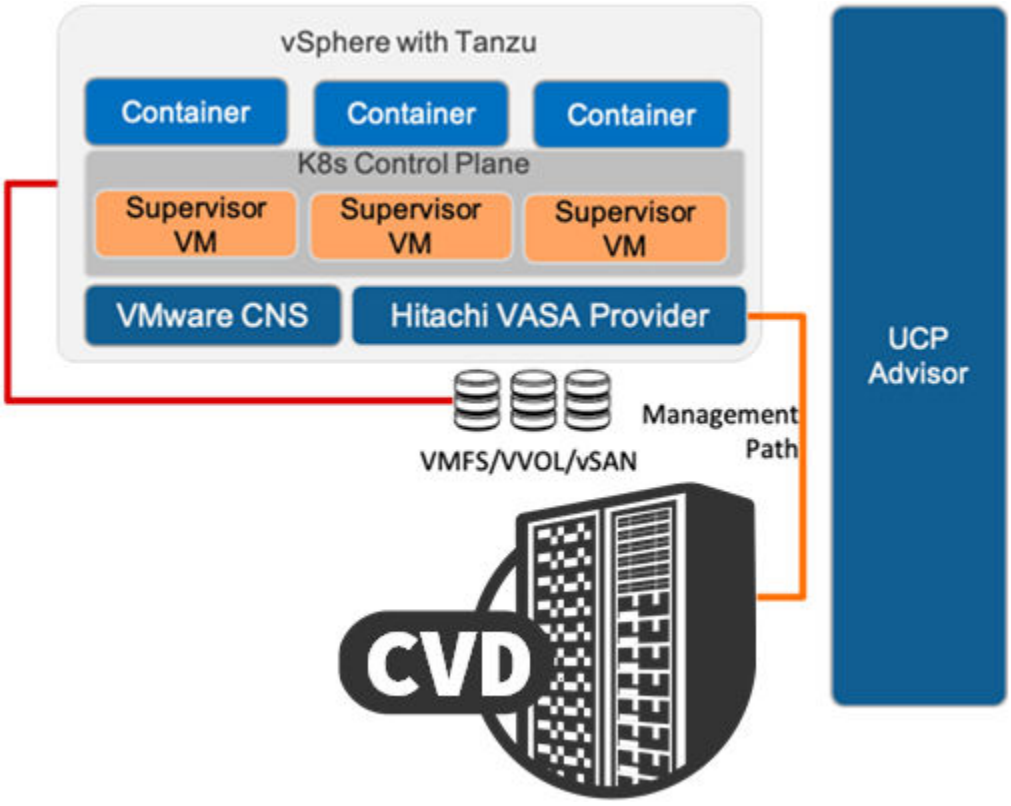

Solution Implementation with Tanzu

This section describes how to implement Hitachi VSP storage on Cisco and Hitachi Adaptive Solutions for CI with VMware to back your Tanzu virtual environment with persistent storage using vSphere CNS.

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 describes 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 VMware Tanzu environment. Additionally, Appendix A: UCP Advisor Storage Administration (on page 75) provides information regarding VSP datastore storage allocation to VMware environments using Hitachi UCP Advisor.

Create Basic LDEVs from Parity Groups

Configuration steps in this section assume that parity groups and LDEVs have been configured on the Hitachi VSP as part of the solution build configured by a Hitachi partner or Hitachi Customer Support 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.

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**.

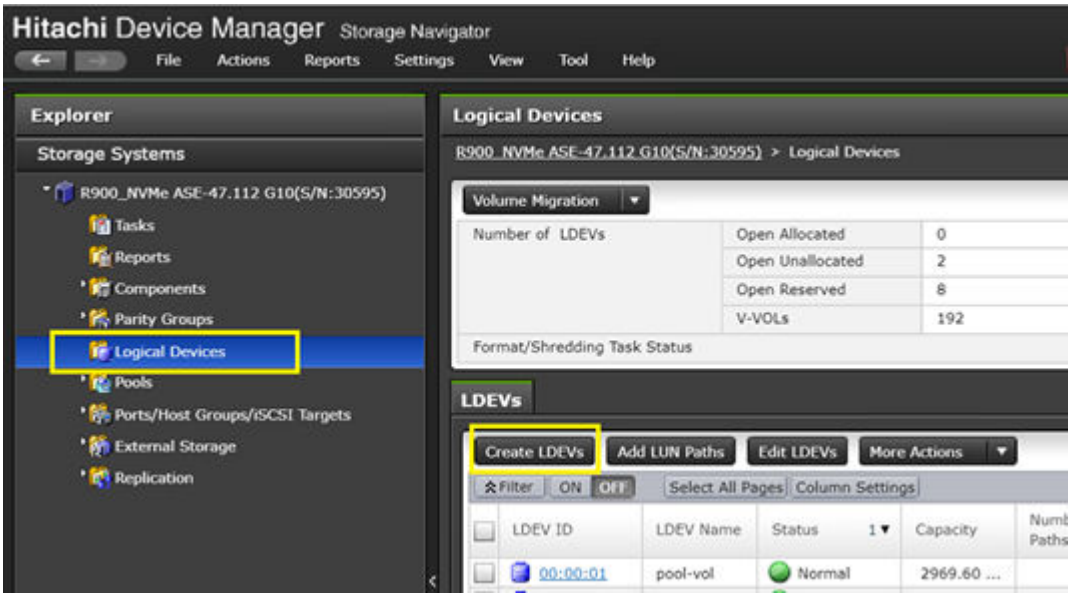

5. Configure the following items in the left pane of the Create LDEVs dialog:
   a. Select Provisioning Type: **Basic**.
   b. System Type: **Open**.
   c. Emulation Type: **OPEN-V**.
6. Click **Select Free Spaces**.
7. Select an available parity group, and then click **OK**.

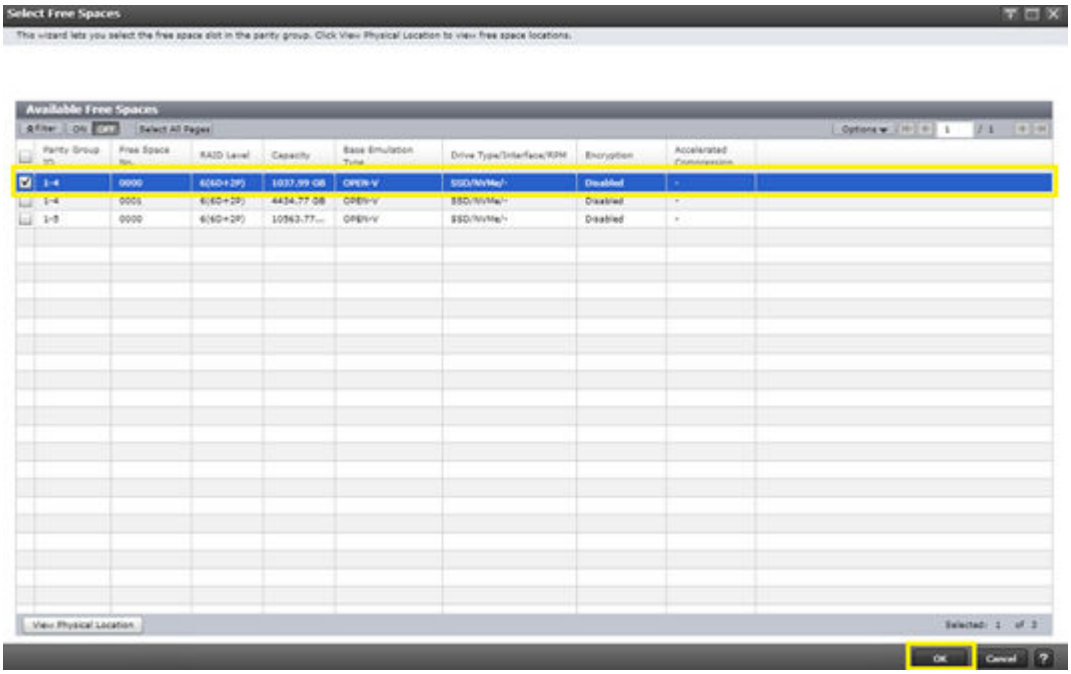

8. Configure the following items in the left pane of the Create LDEVs dialog.
   a. Define **LDEV Capacity** and select the unit size.
   b. Define **Number of LDEVs per Free Space**.
   c. Define **LDEV Name**, such as UCS_PoolVOL_VMFS, or UCS_PoolVOL_vVols
9. Click **Add**.

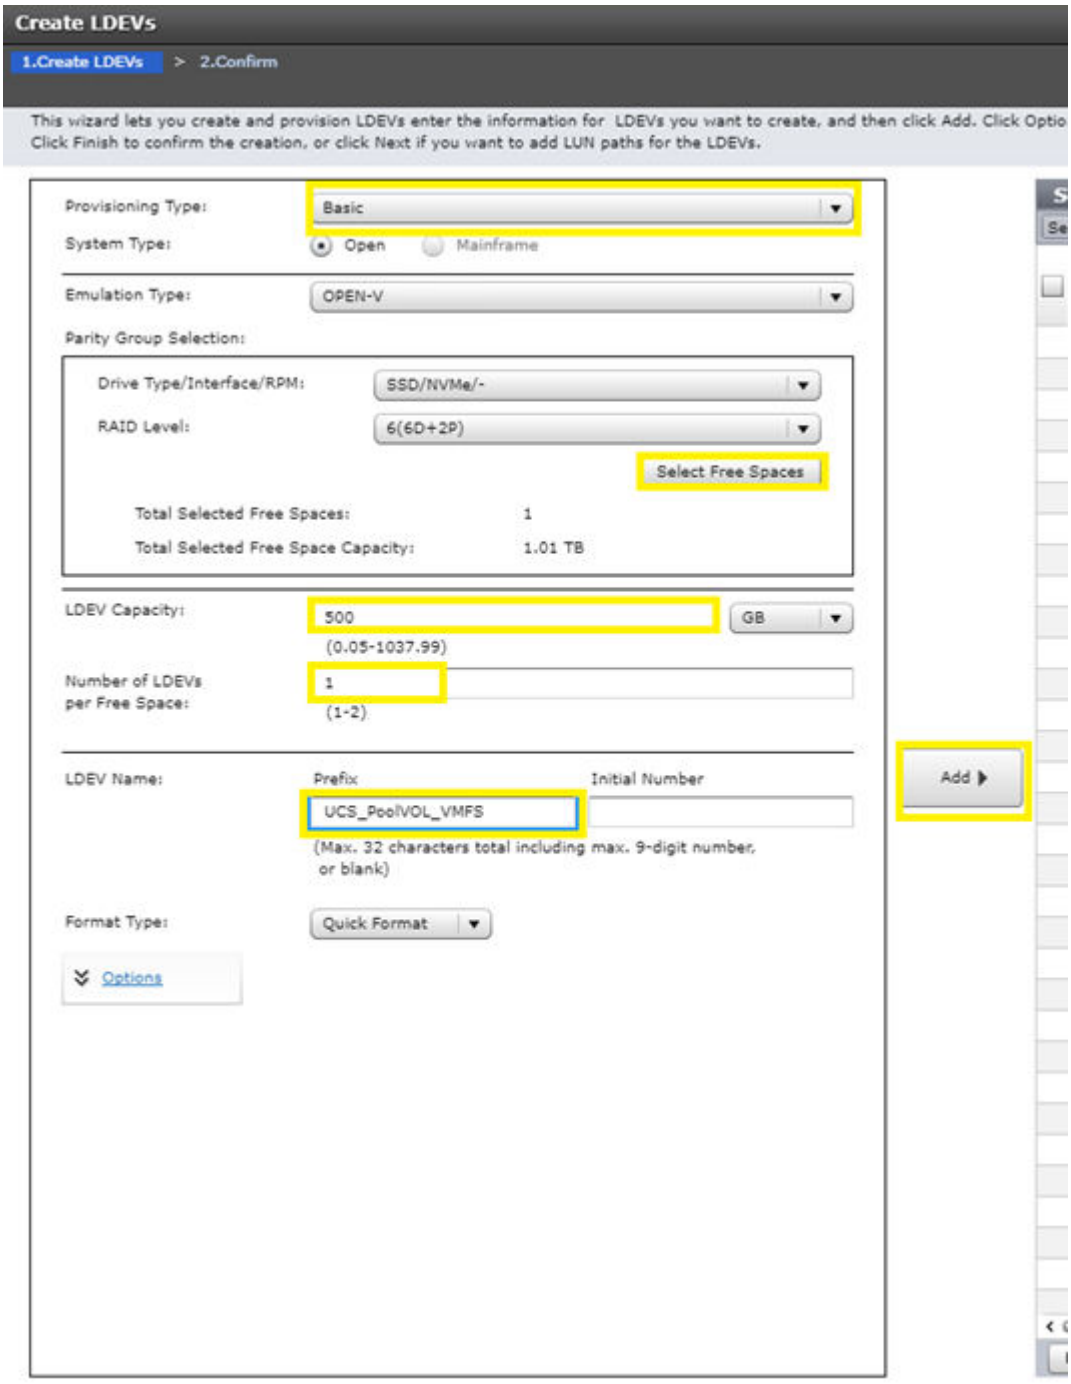

10. Click **Finish** > **Apply**.

Create Hitachi Dynamic Provisioning (HDP) Pool

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

Use the following procedure to create an 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**.

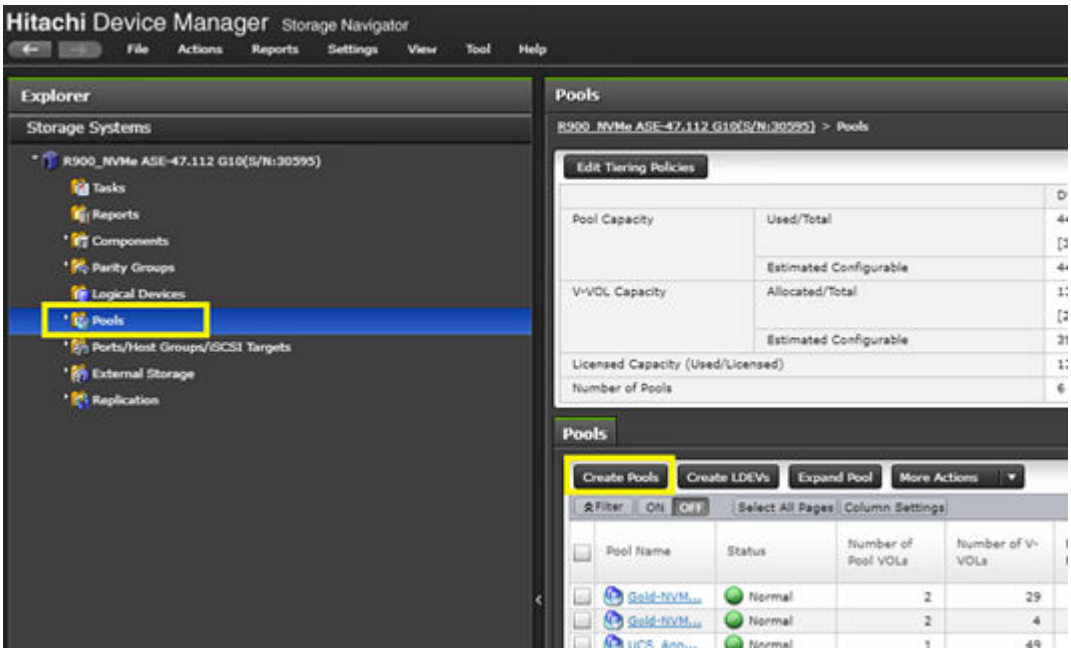

4. In the **Create LDEVs** dialog, configure the following items:
   a. Pool Type: **Dynamic Provisioning**
   b. System Type: **Open**
   c. Multi-Tier Pool: **Disable**
   d. Pool Volume Selection: **Manual**
5. Click **Select Pool VOLs**.
6. Select the applicable basic LDEV to support the HDP pool. Click **Add**.

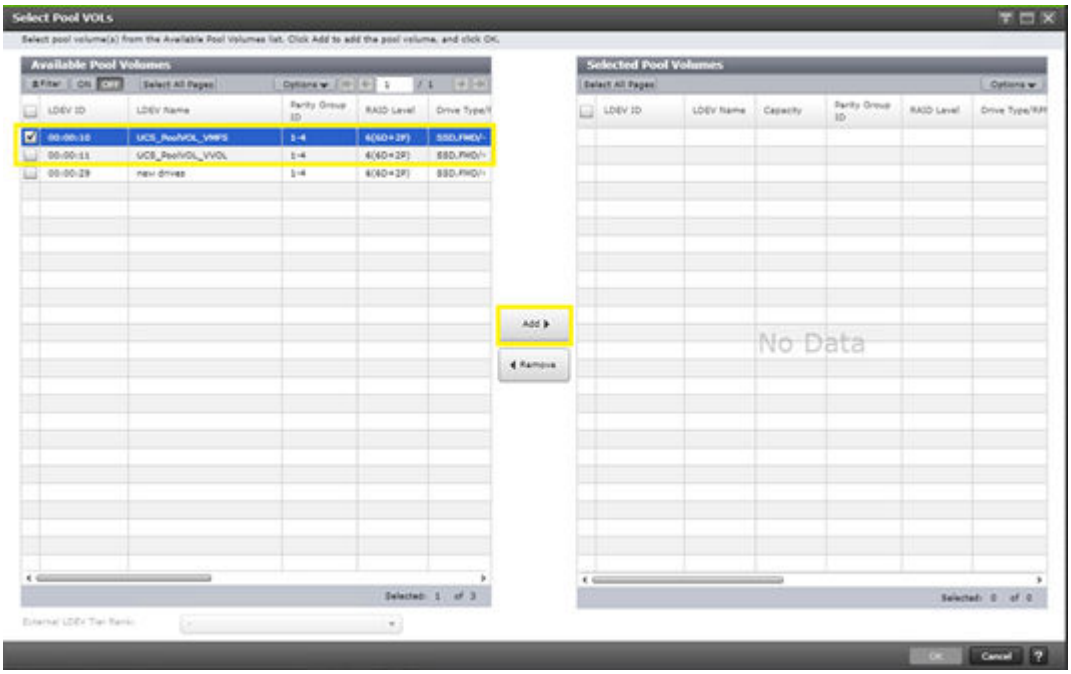

7. Click **OK**.
8. Enter the applicable Pool Name, and then click **Add**.

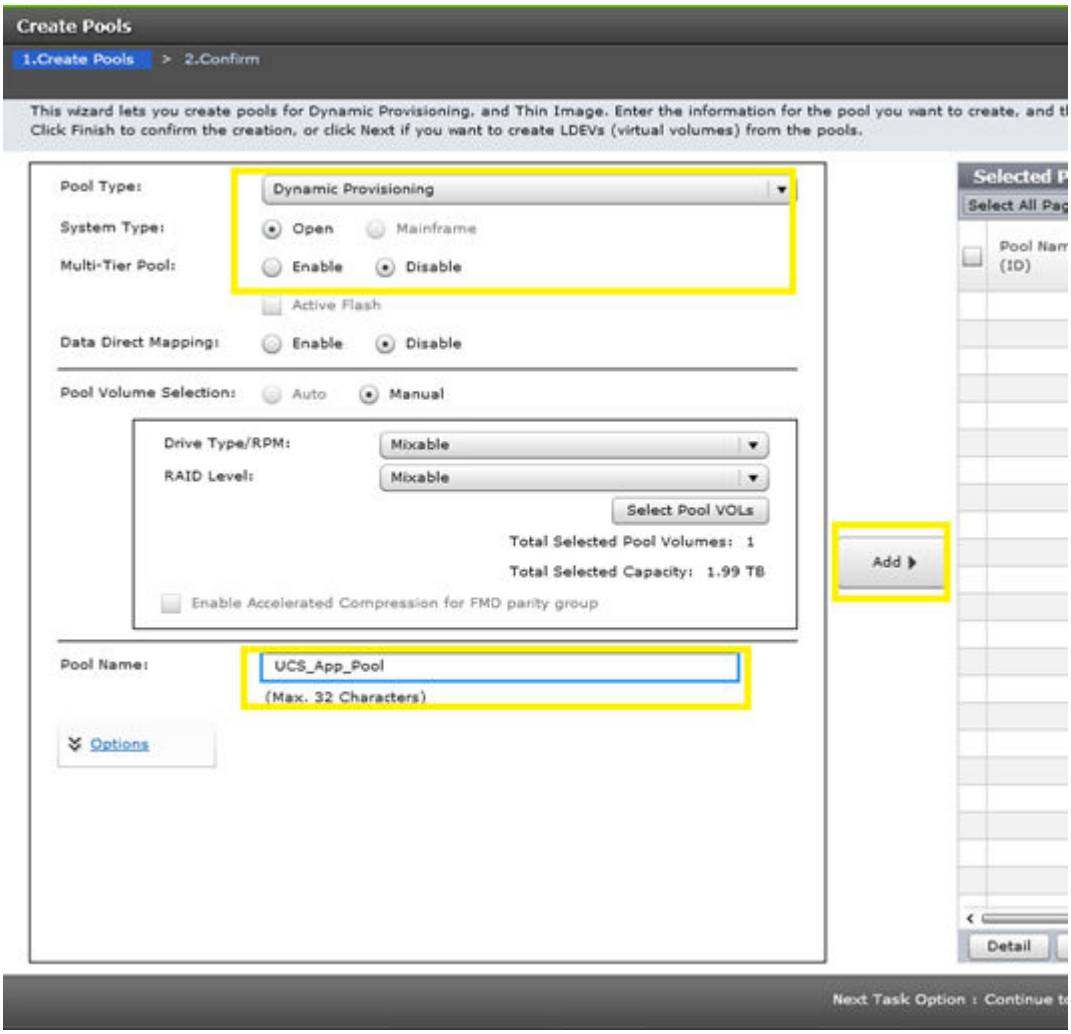

9. Click **Finish** > **Apply**.

Create LUNs to support VMFS Datastores

After the HDP pools are created, create LUNs (virtual volumes) to present as VMFS datastores to vCenter.

Use the following procedure to create a LUN from an HDP pool:

Procedure

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

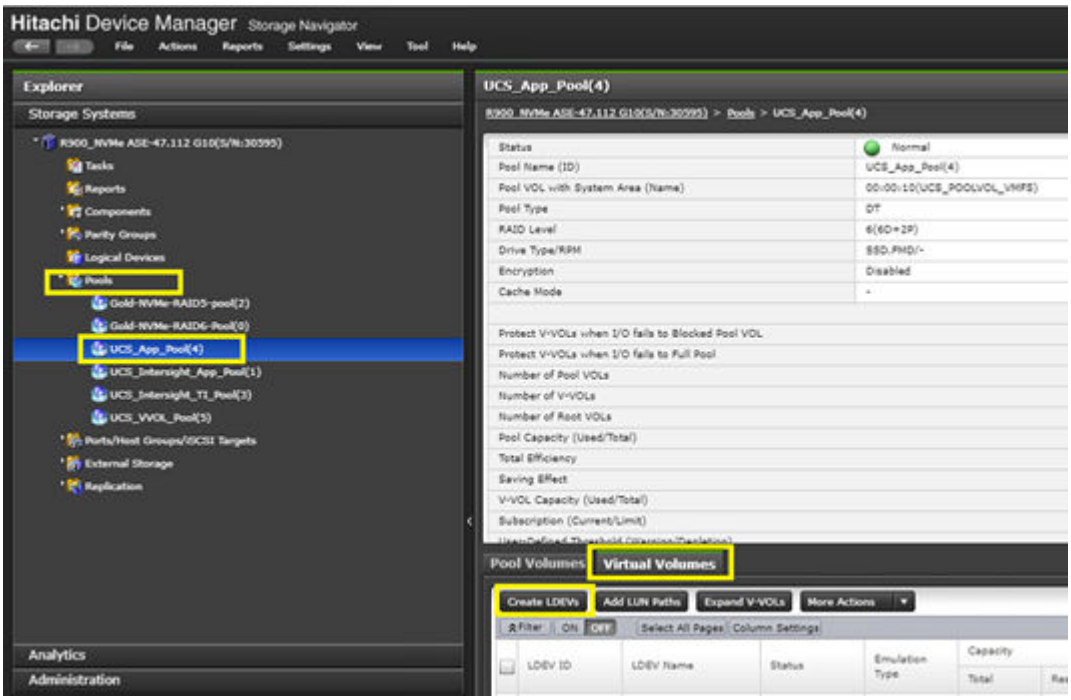

4. Configure the following items in the left pane of the Create LDEVs dialog:
   a. Provisioning Type: **Dynamic Provisioning**
   b. System Type: **Open**
   c. Emulation Type: **OPEN-V**
   d. Capacity Saving: **Disabled**
5. Also define **LDEV Capacity**, **Number of LDEVs**, and **LDEV Name**.
6.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

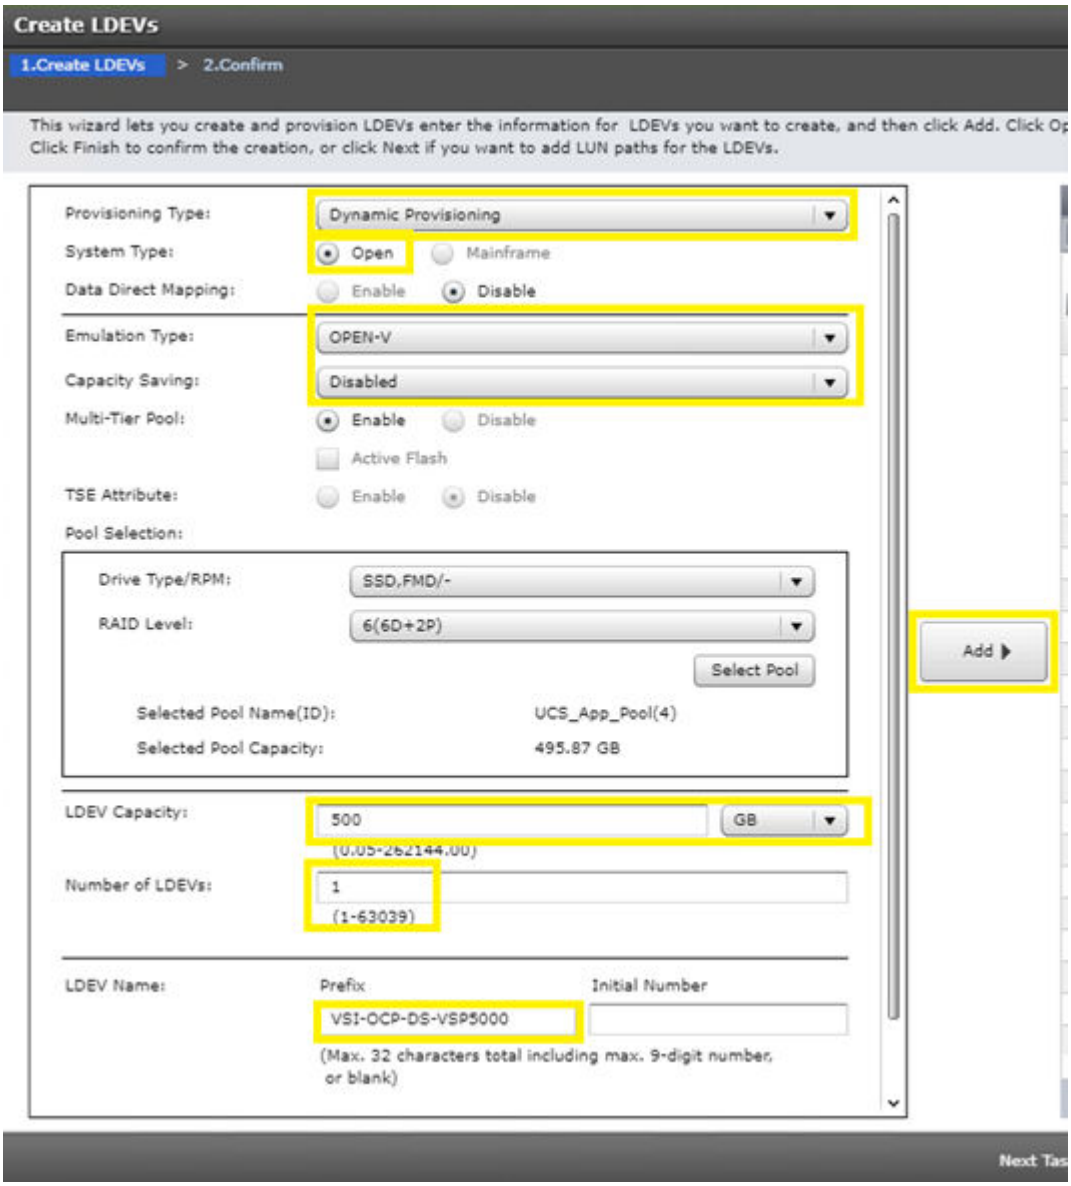

7. Click **Finish** > **Apply**.

Create host groups

After the LUNs have been created from the available HDP pools, they must be added to host groups. For host group creation, see *Create Host Groups for UCS Server vHBAs on Each Fabric* in Cisco and Hitachi Adaptive Solutions for Converged Infrastructure.

Add LDEV Paths to Host Groups

After the host groups have been defined for the underlying UCS infrastructure, add LDEV paths to the 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**.

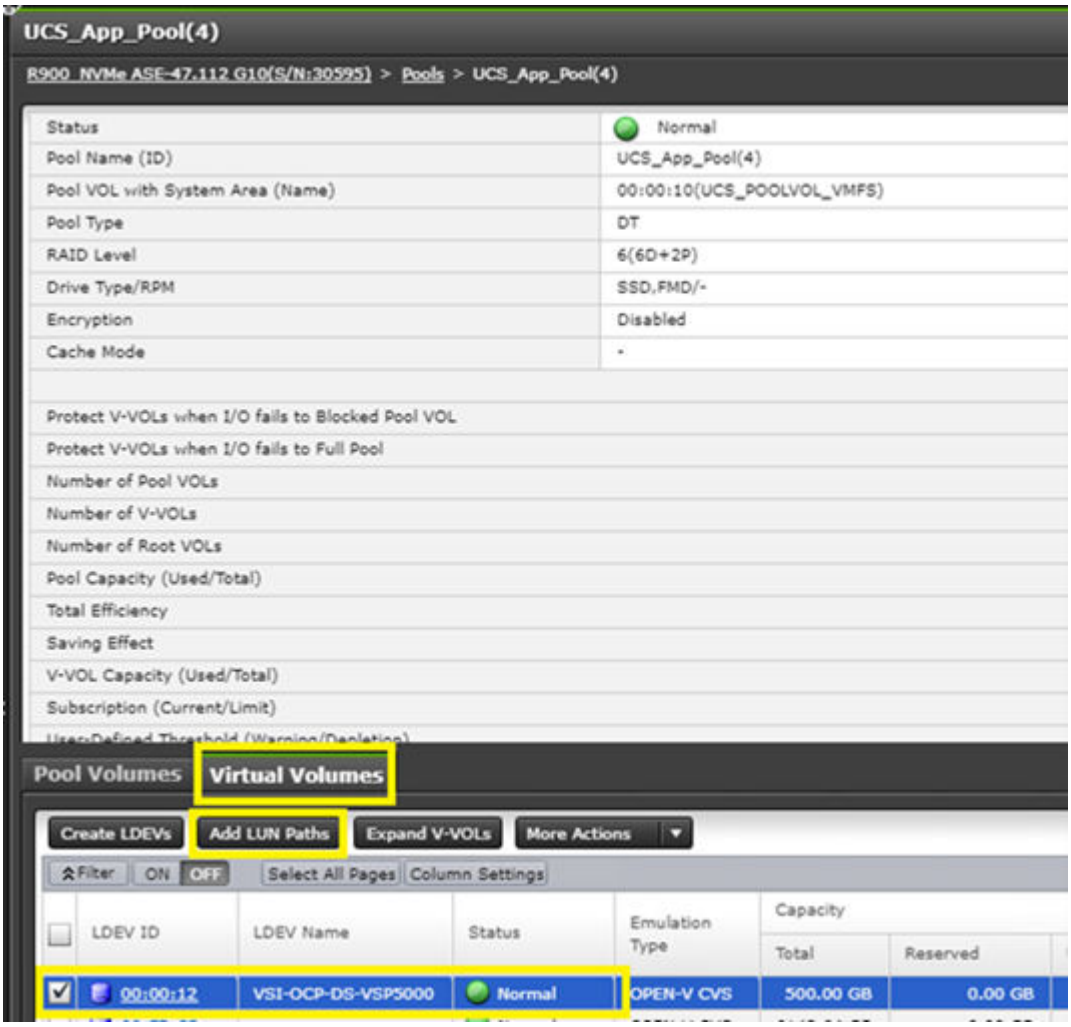

2. Click **Next**.

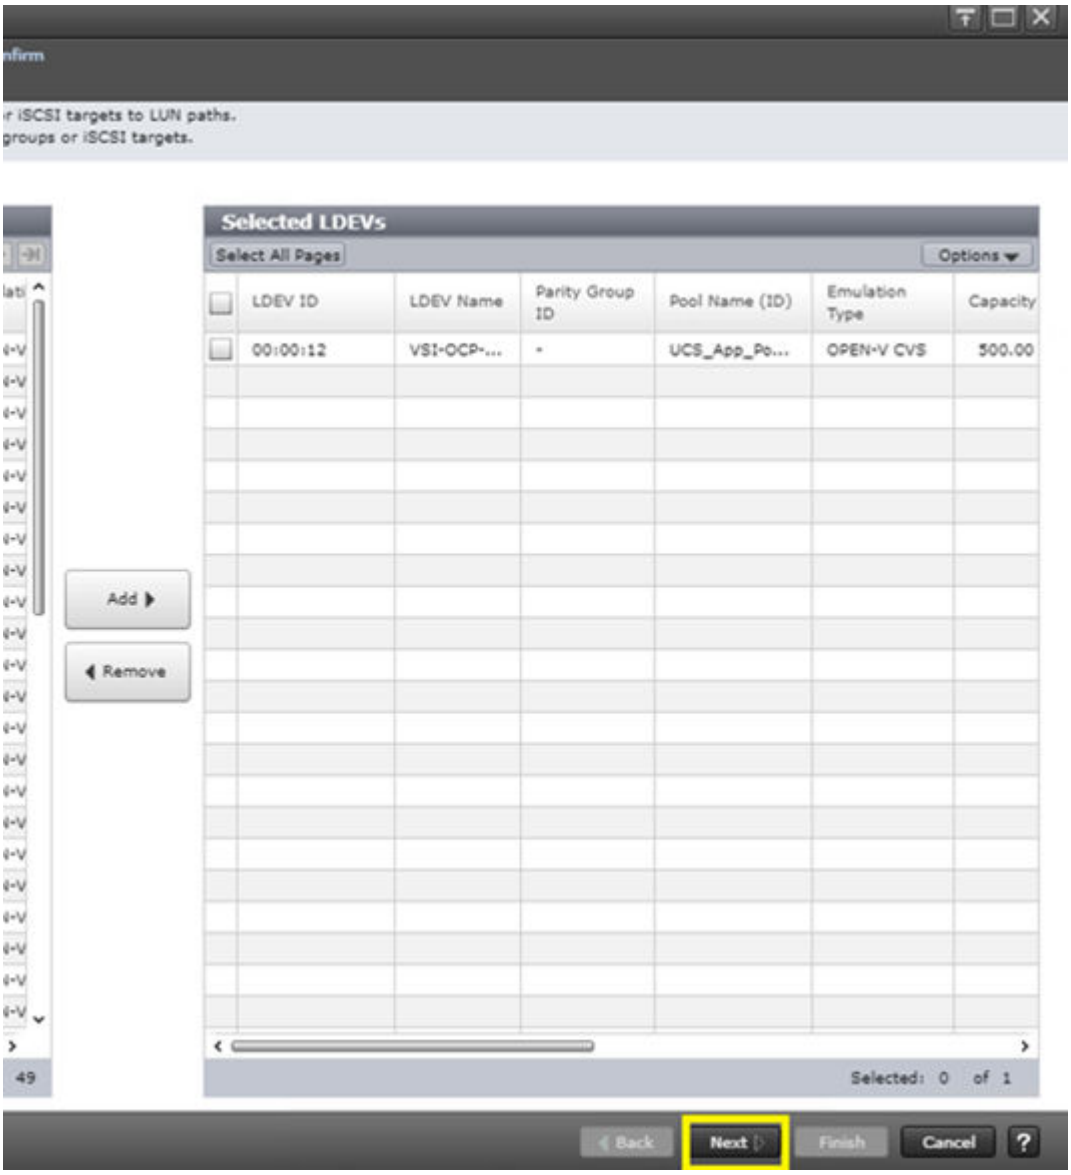

3. Using the Filter tool, search for VSI host groups based on **Host Group Name** and **Contains**.

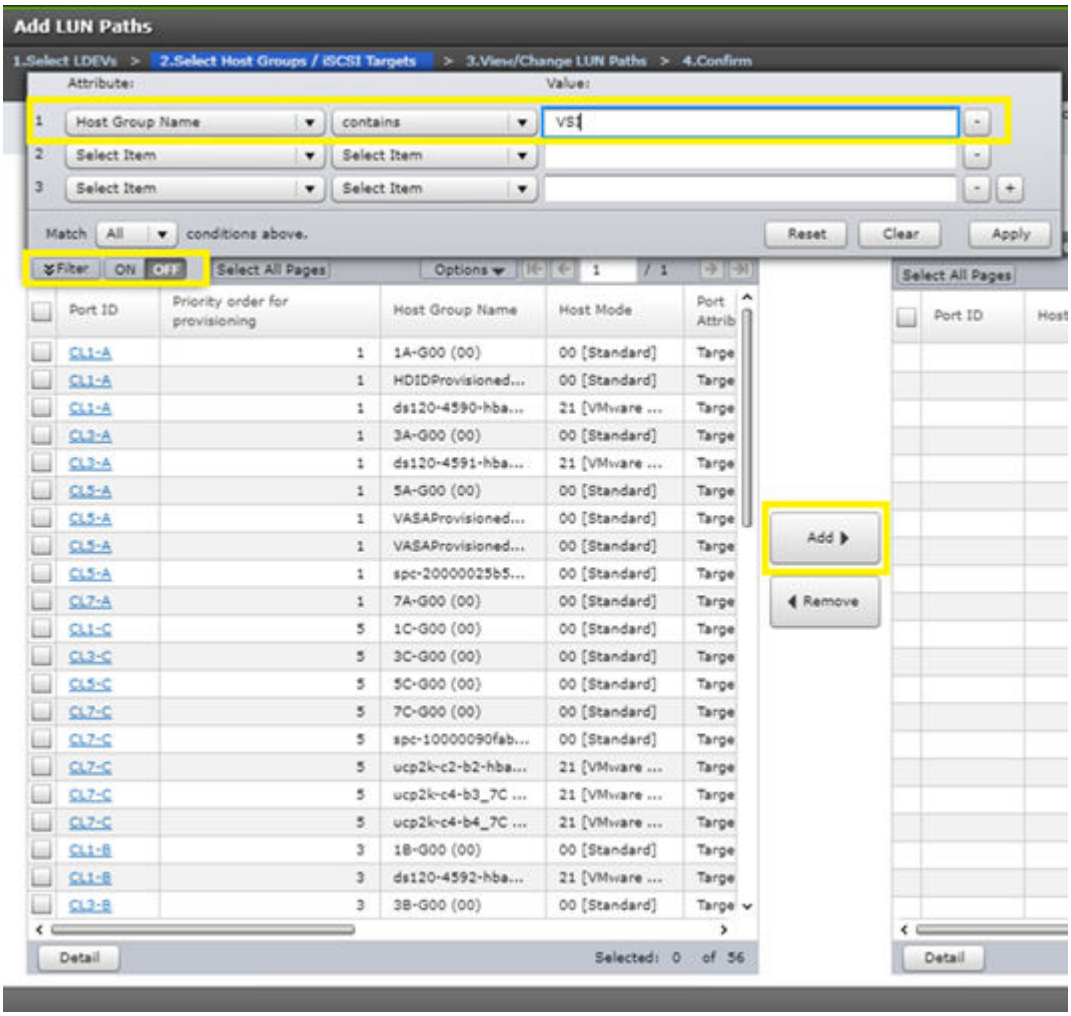

4. Select the applicable host groups, and then click **Add**.
5. Click **Next**.
6. The View/Change LUN Paths screen shows the LDEV to which you are adding paths and the associated host LUN ID that will be presented to the host on a per-path basis. Verify that the LUN alignment is correct before presenting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

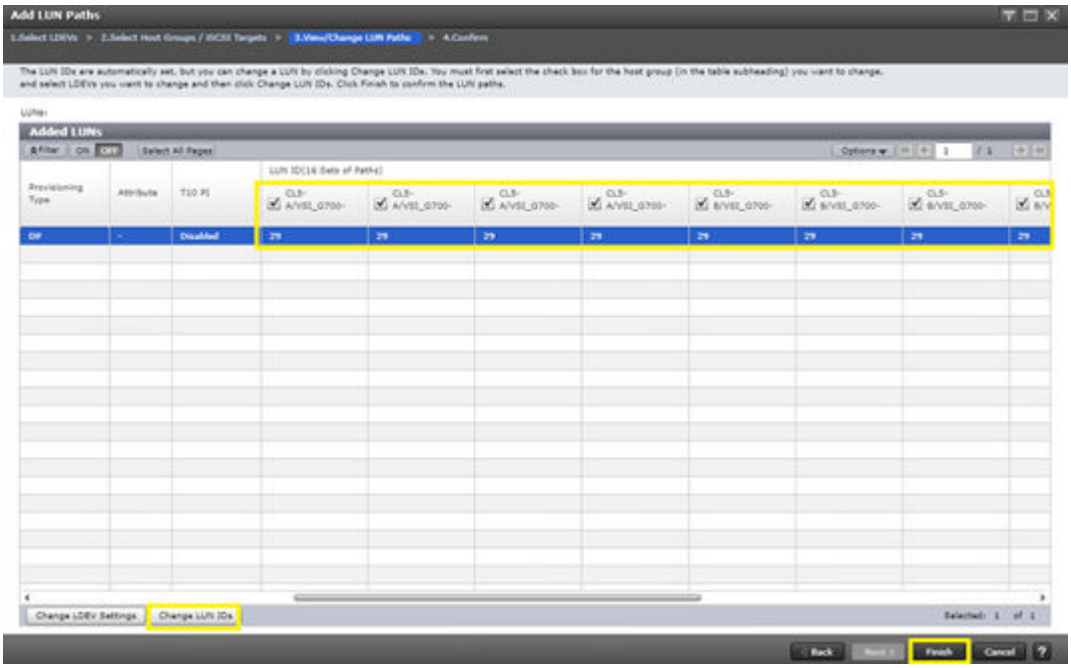

After you have added LUN paths to your LDEV, you can continue with onboarding the VMFS datastore in VMware vCenter. See Onboarding VMFS Datastores (on page 35).

VSP vVols Configuration for vVols Datastore

This section describes 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 Tanzu 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

Optionally admins can create vVols datastores using UCP Advisor using information in Appendix A: UCP Advisor Storage Administration (on page 75). Admins must verify that parity groups have been configured as well as completing the following steps to add storage pool resources to 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 you plan to 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 82).

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**.

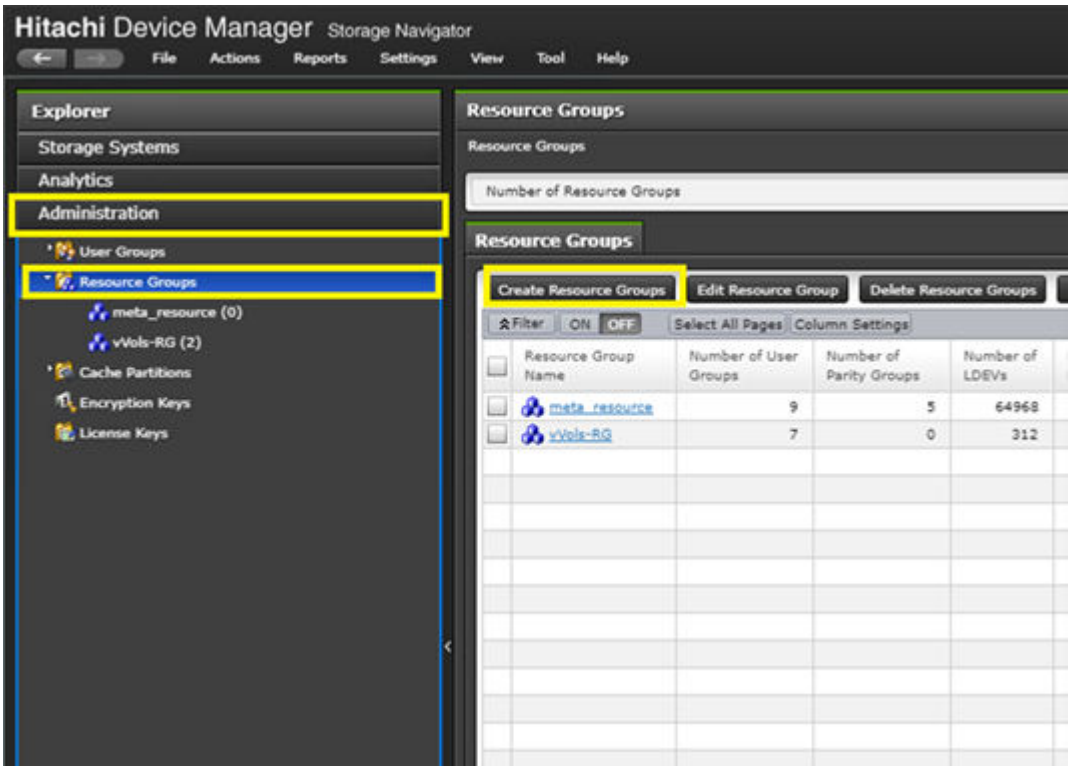

3. Enter a **Resource Group Name**.
4. Click **Select LDEVs**.

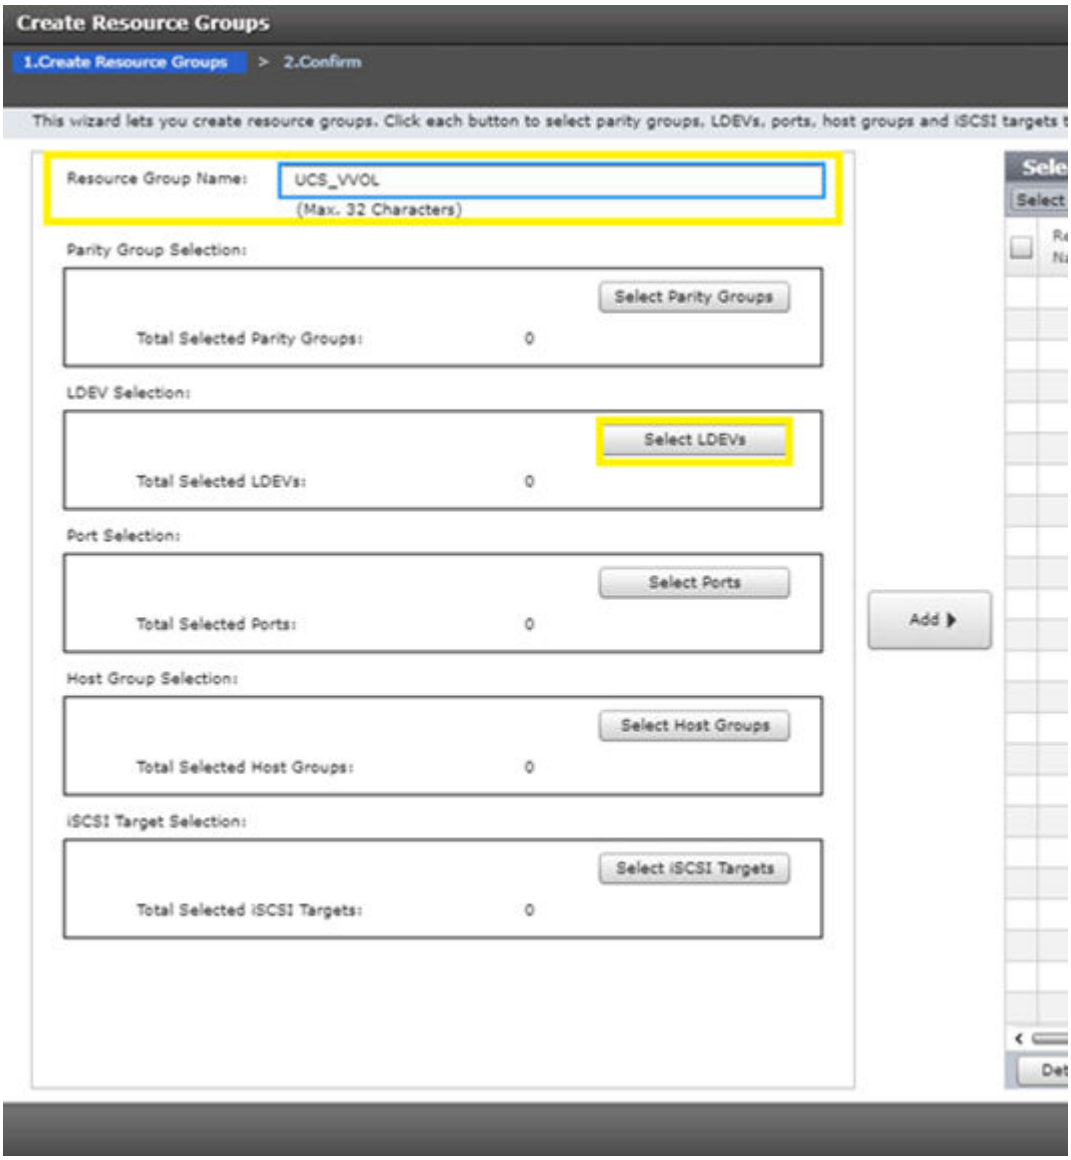

5. Select the basic LDEV you created from the parity group that backs your vVol pool, select LDEV IDs from the Available LDEVs list, and then click **Add**.

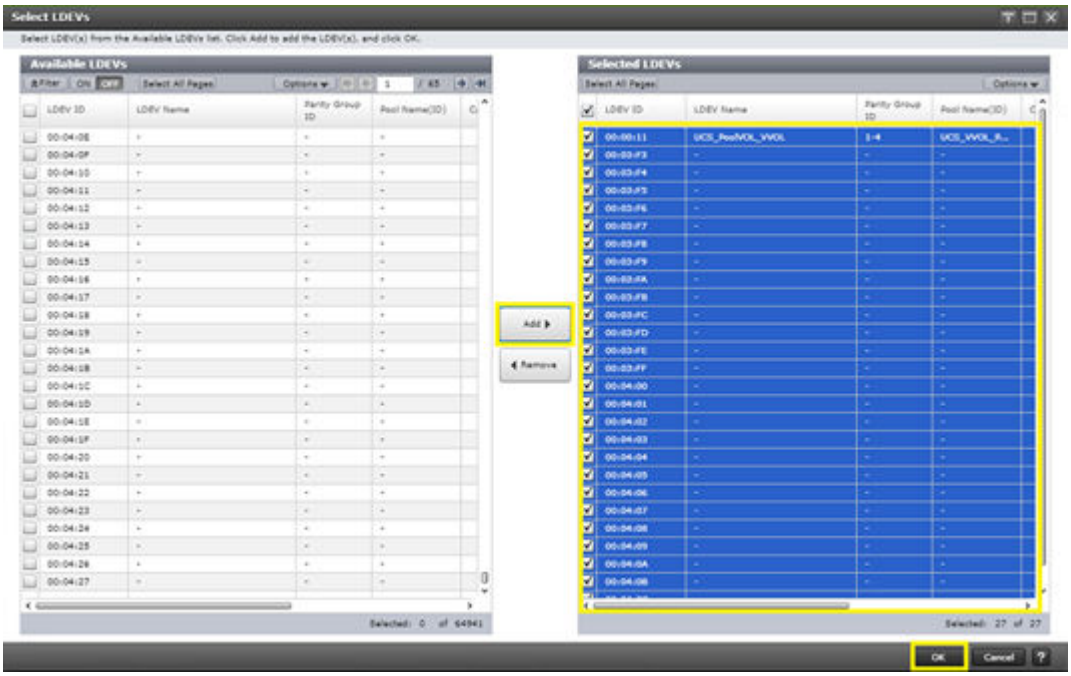

6. Click **OK**.
7. After the LDEVs have been defined, click **Add** > **Finish** > **Apply**.

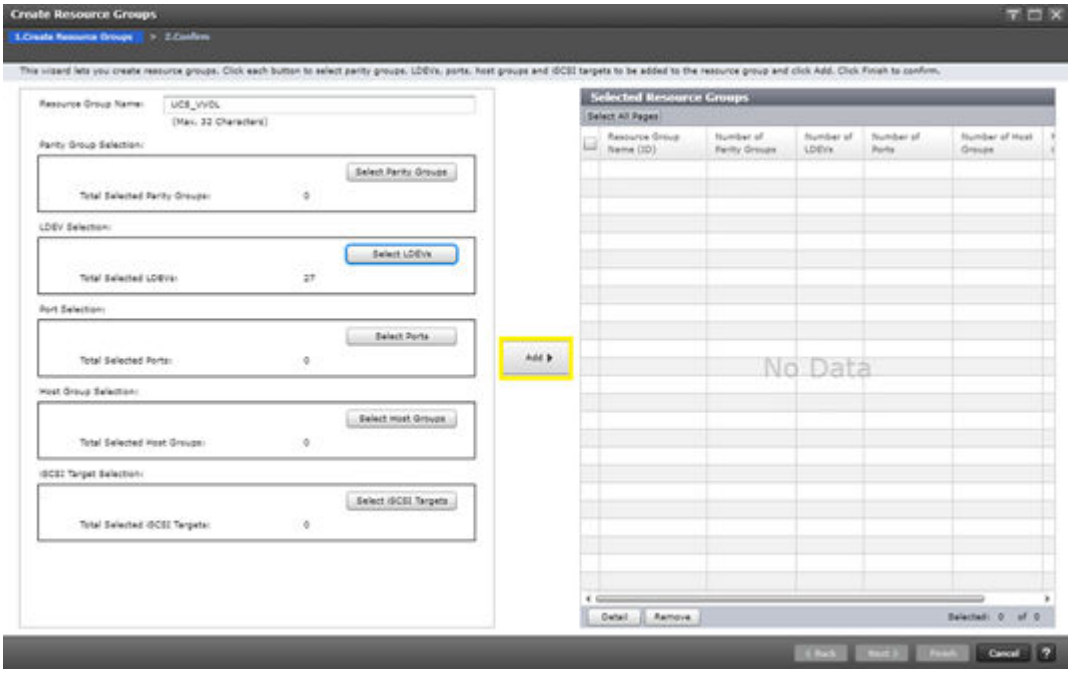

Hitachi Storage Provider for VMware vCenter storage configuration

The following section describes 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 polices that will be used by StorageClasses.

Storage Provider for VMware vCenter deployment and onboarding to VMware vCenter is not covered in this guide. See Related Documents (on page 82).

Onboard Hitachi storage

After the Storage Provider for VMware vCenter is deployed, 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.

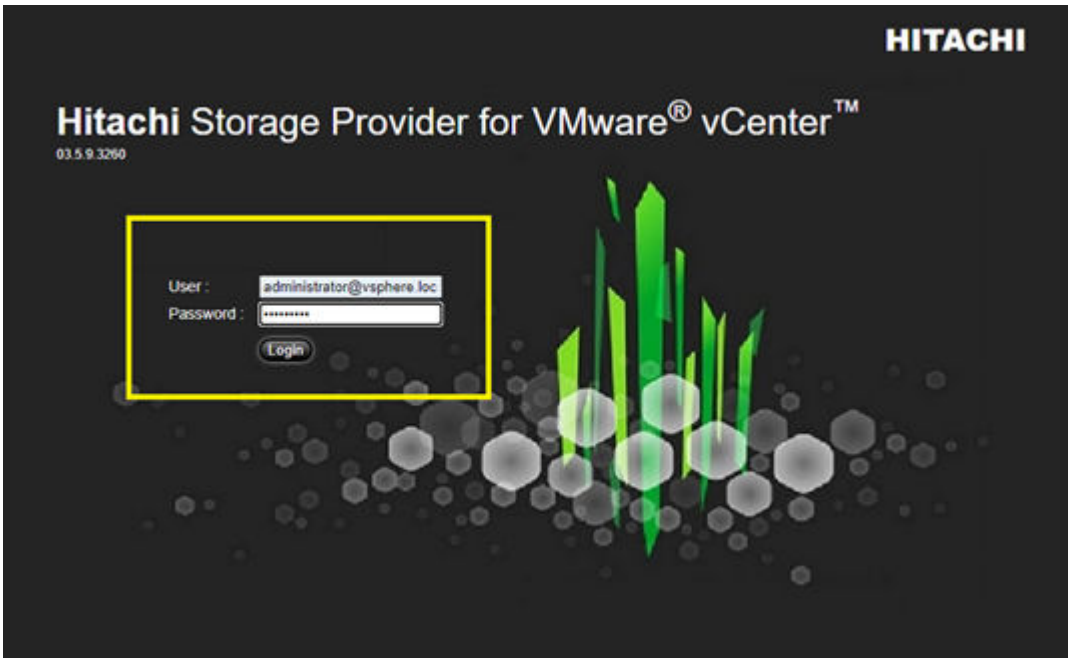

2. Click **Manage Storage Systems**.

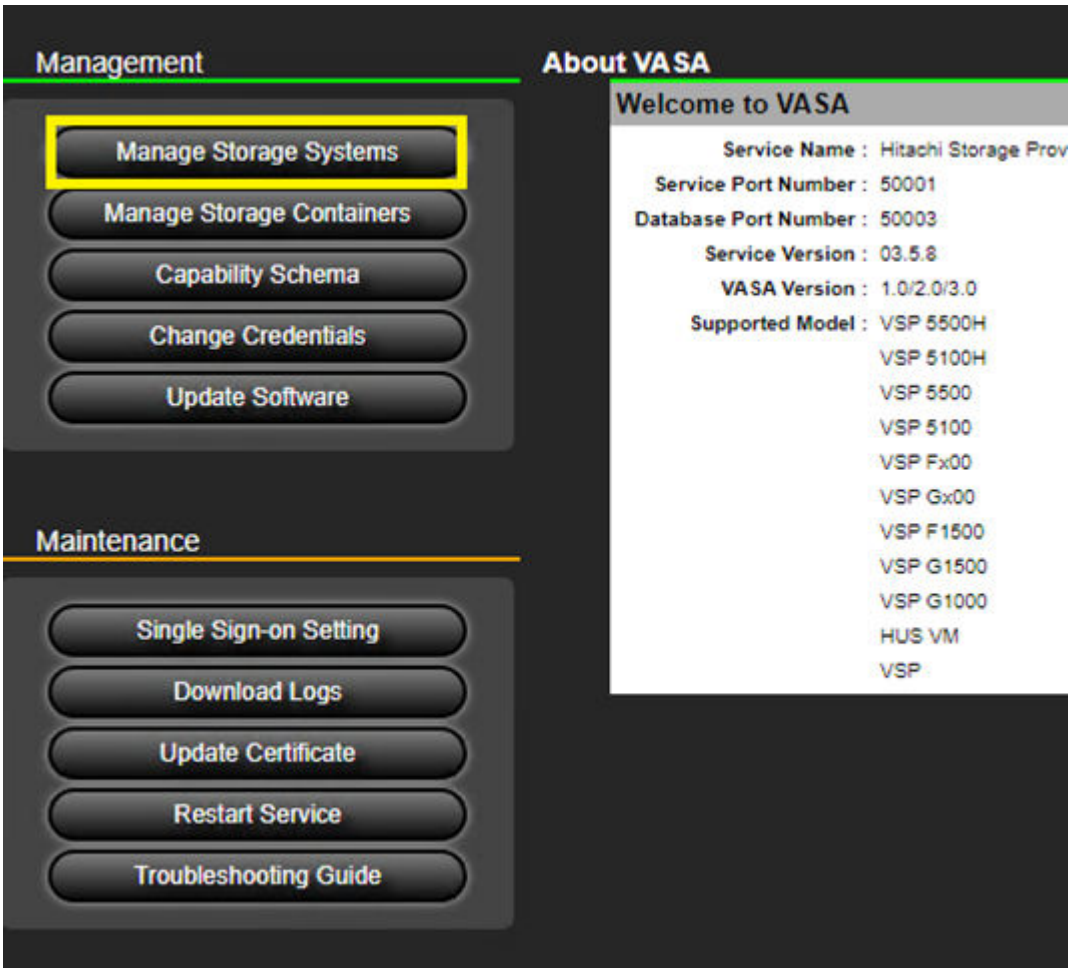

3. From the **Physical Storage** tab, click **Add Storage System**.
4. From **Add Storage System**, do the following (this might not apply to all storage types):
   a. From the **Storage System Type** list, click the system model.
   b. Click **SVP** (optional).
   c. Enter the **SVP IP** of the storage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

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

Create storage containers and capability profile

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**.

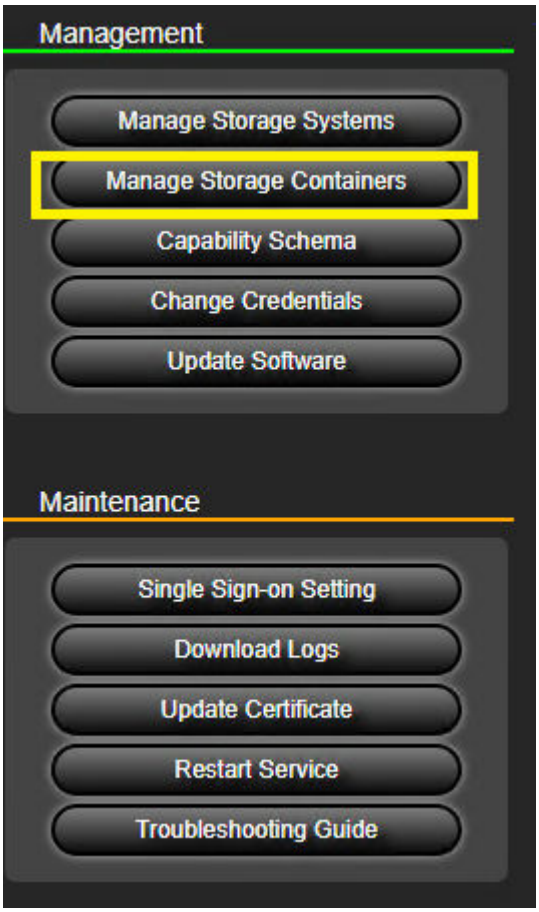

2. Click **Create Storage Container**.
3. Configure the following:
   a. Define a storage container **Name**.
   b. Select an onboarded **Storage System**.
   c. Select the **Resource Group** configured on your VSP.
4. Select an undefined Capability profile, and then click **Define Profile**.

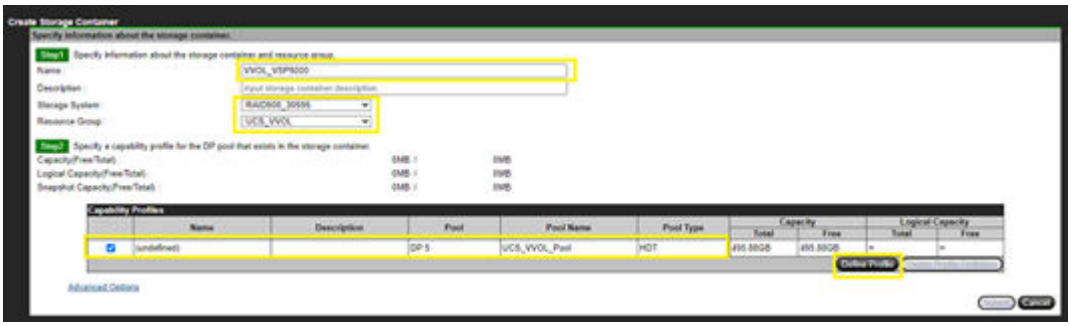

5. In the **Define Capability Profile** window, do the following:
   a. Define a profile **Name**.
   b. b. Assign managed storage capabilities to your profile. The characteristics need to relate to your vVols resource group that is native to the registered storage system.
   c. Click **OK** > **Submit**.

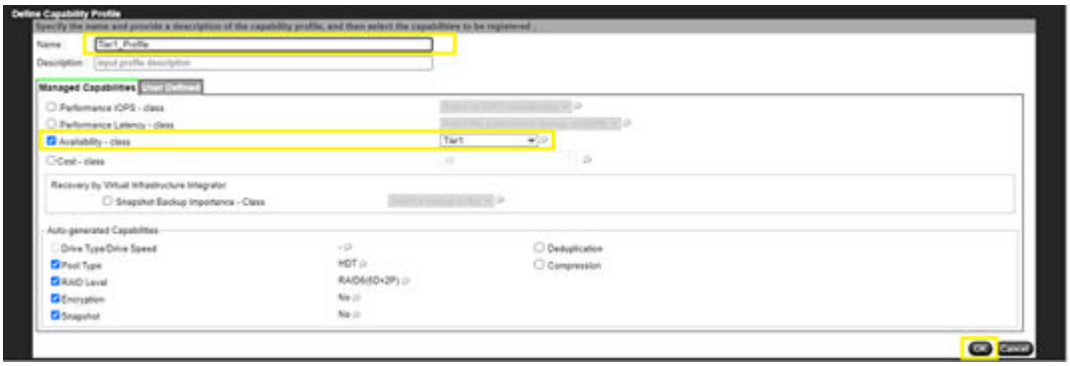

Note: Administrators can also define custom capabilities not natively defined within the VASA APIs under Capability Schema. For more information see Related Documents (on page 82).

Manage storage systems for VMFS LDEVs

With Storage Provider for VMware vCenter, attributes of the logical units supporting the VMFS datastore are passed down to VMware vSphere. 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

1. From the Management list, click **Manage Storage Systems**.

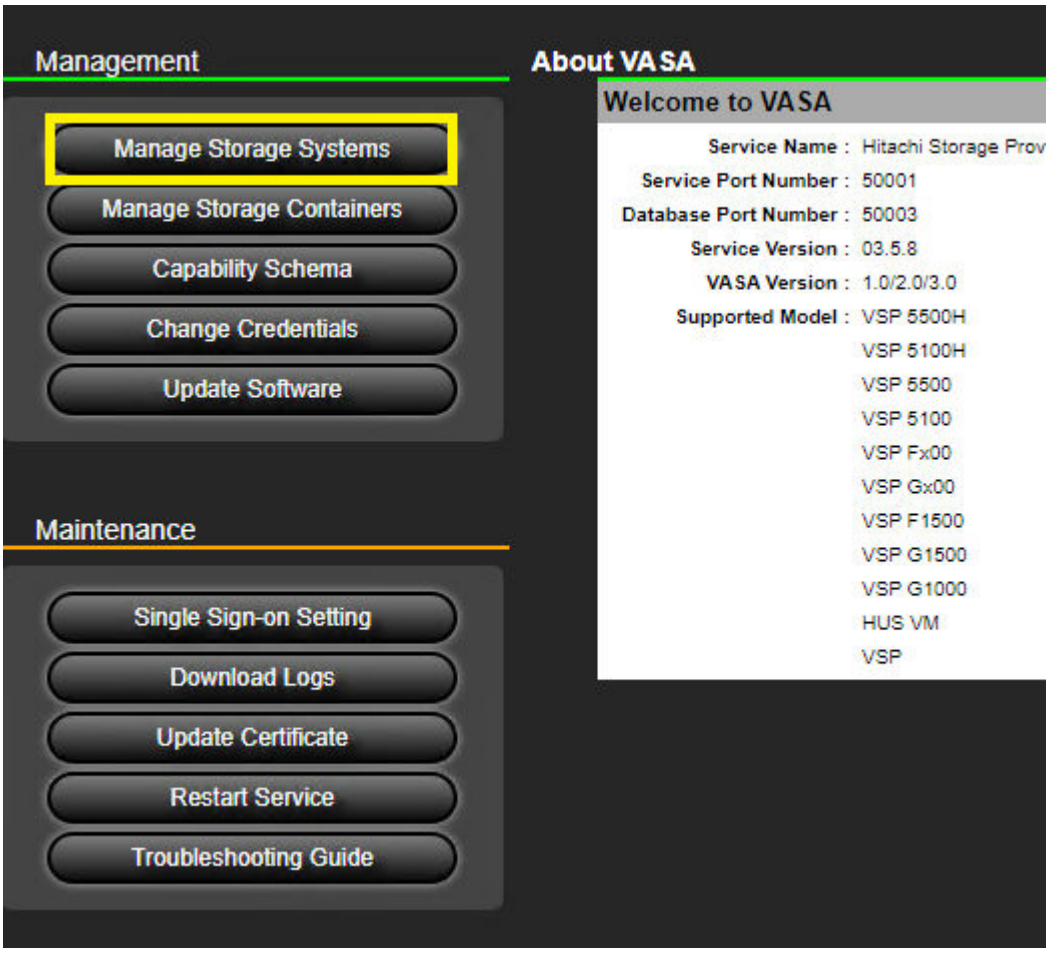

2. From the Physical Storage System list, click **LDEVs** on the storage system that contains the applicable VMFS datastore LUN.

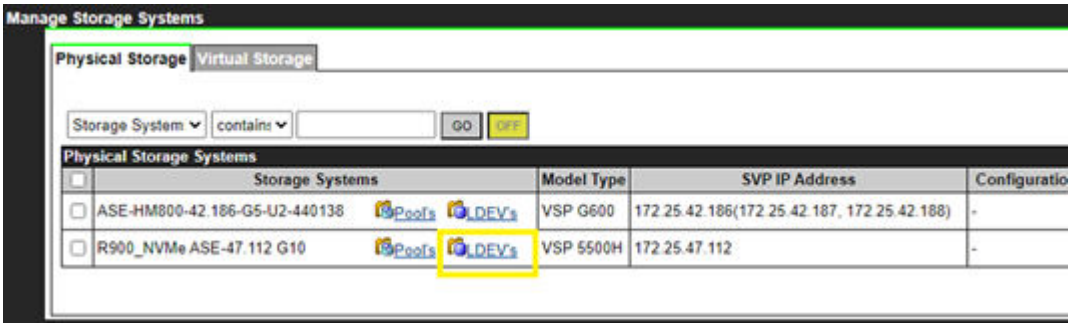

3. From the storage system volume list, select the applicable **volume ID**, and then click **Define Profile**.

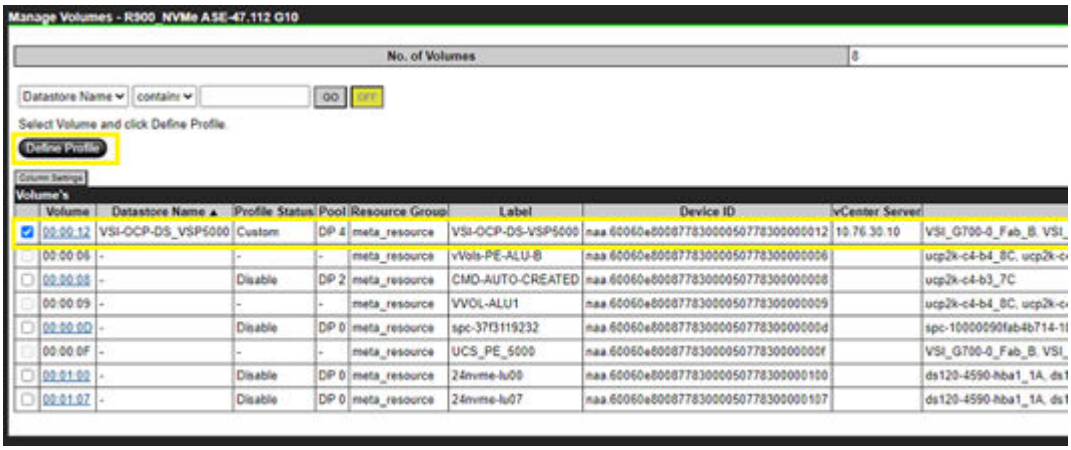

4. Select the applicable tags that you want relayed to the VMware administrator.
5. Click **Submit**.

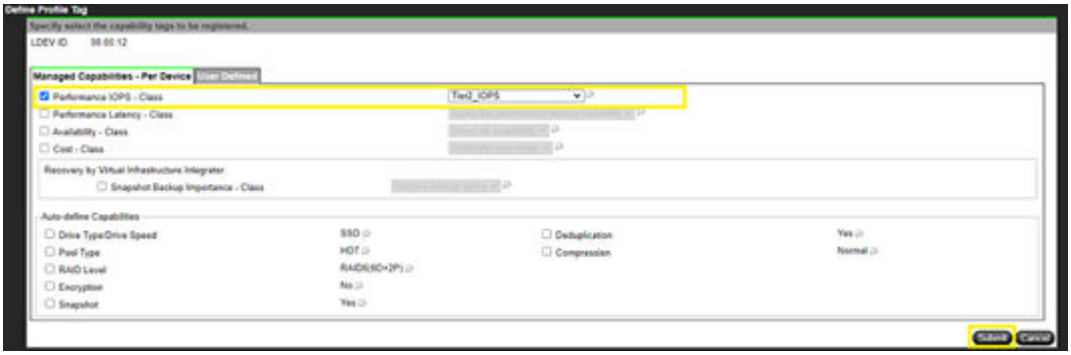

VMware vCenter configuration

This section describes the configuration of VMware vCenter in preparation to use 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 icon, and then right-click on the applicable datacenter.
3. Click **Storage** > **New Datastore**.

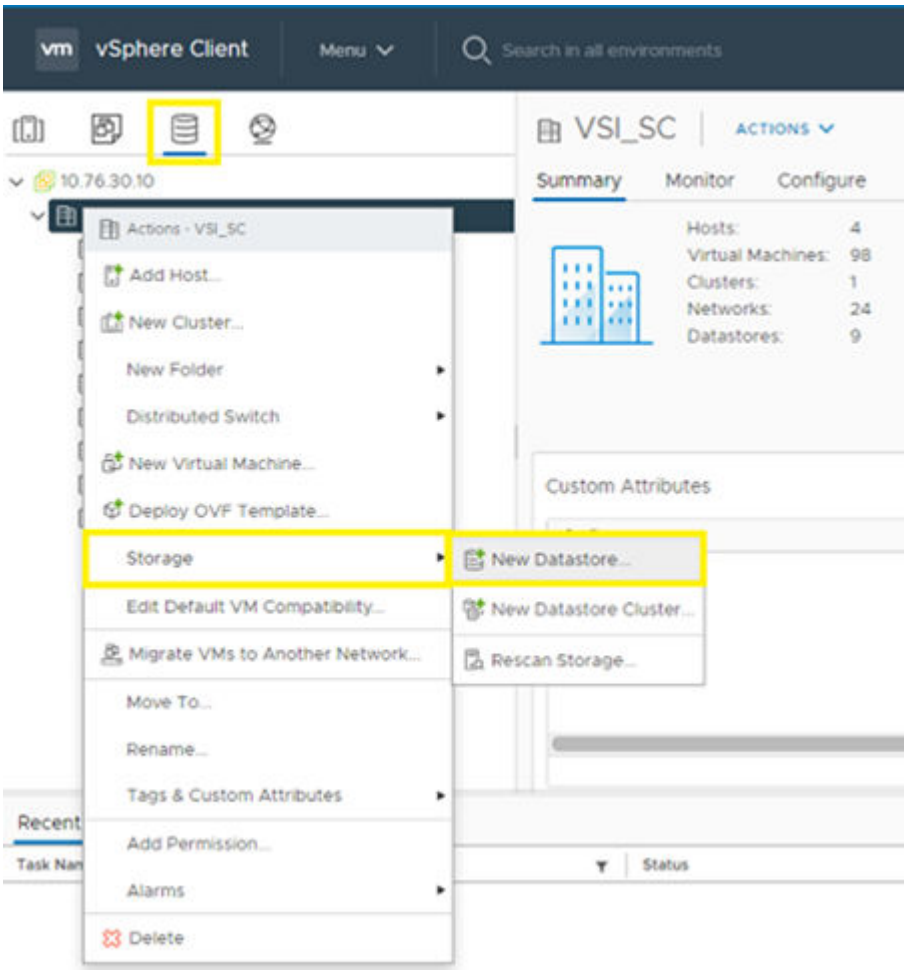

4. Select **VMFS** as the datastore type. Click **Next**.

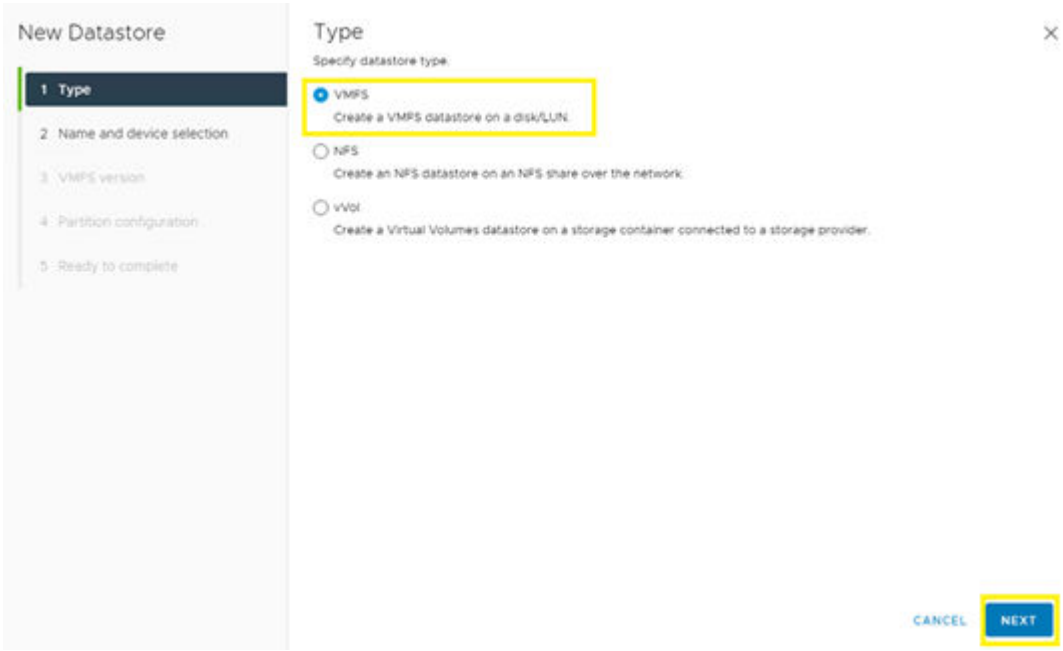

5. Define a **Datastore name**, and then select an available ESXi host and the LUN presented by Hitachi VSP storage to onboard as the new datastore. Click **Next**.

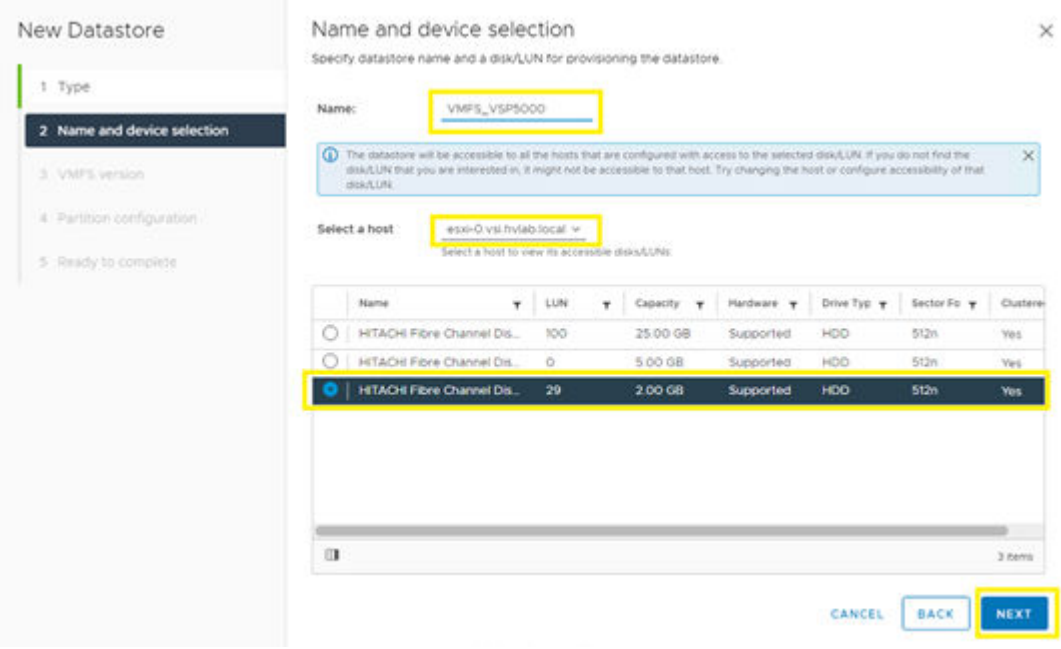

6. Select **VMFS 6**, and then click **Next**.
7. In **Partition Configuration**, confirm that the entire capacity of the volume is claimed, and then click **Next**.

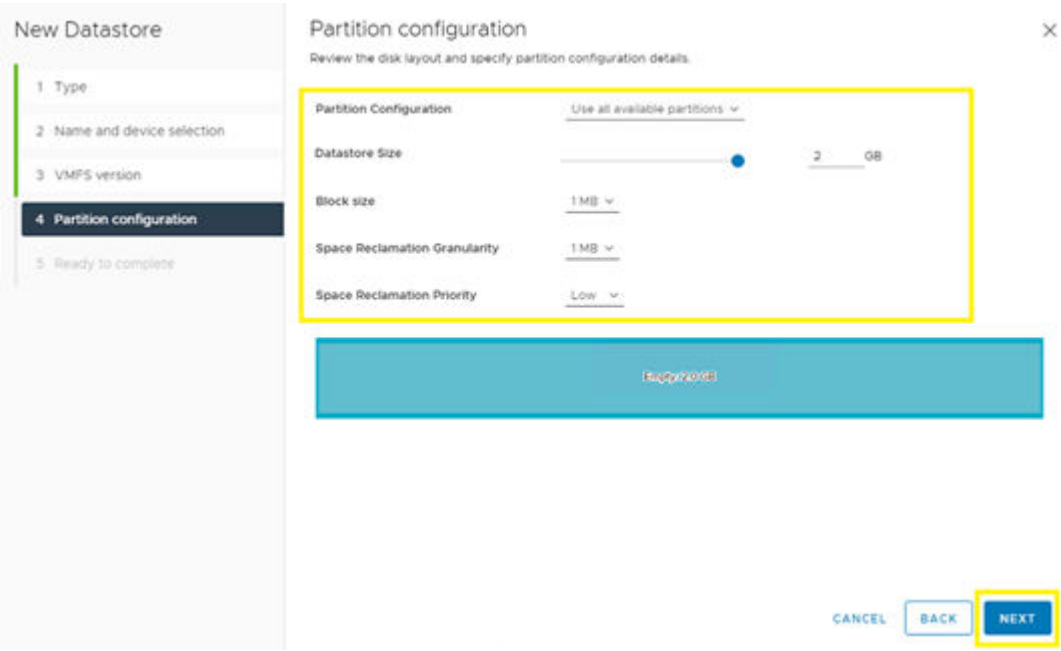

8. Review the configuration, and then click **Finish**.

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 a best practice to pass down all desired tags from Storage Provider for VMware vCenter.

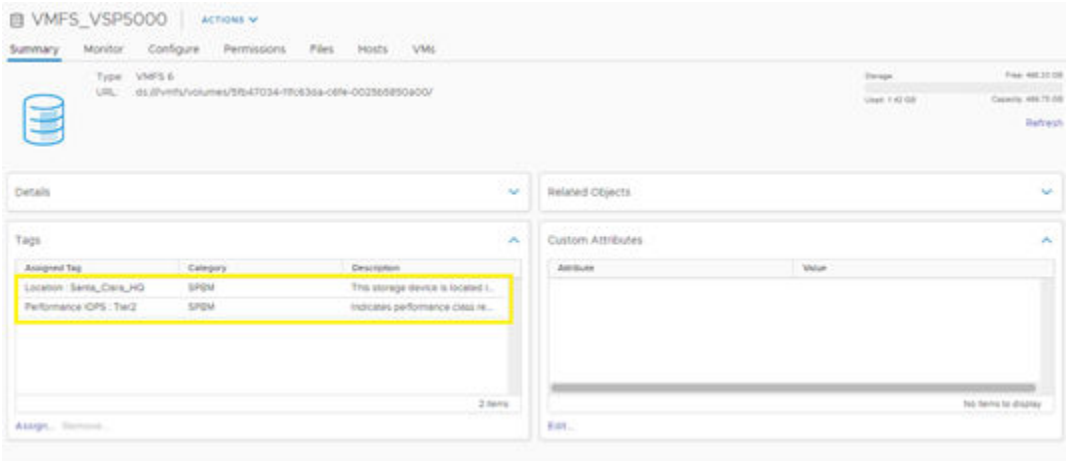

Onboarding a vVols datastore

Prior to onboarding a vVols datastore verify that applicable VSP storage resource groups have been configured and that the correct storage container and capability profile have been defined in Hitachi Storage Provider for VMware vCenter.

Also verify that the Storage Provider for VMware vCenter has been registered in vCenter along with the VSP Administrative Logical Unit (ALU), also known as the VASA Protocol Endpoint (PE). For more information on deployment, see Related Documents (on page 82).

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**.

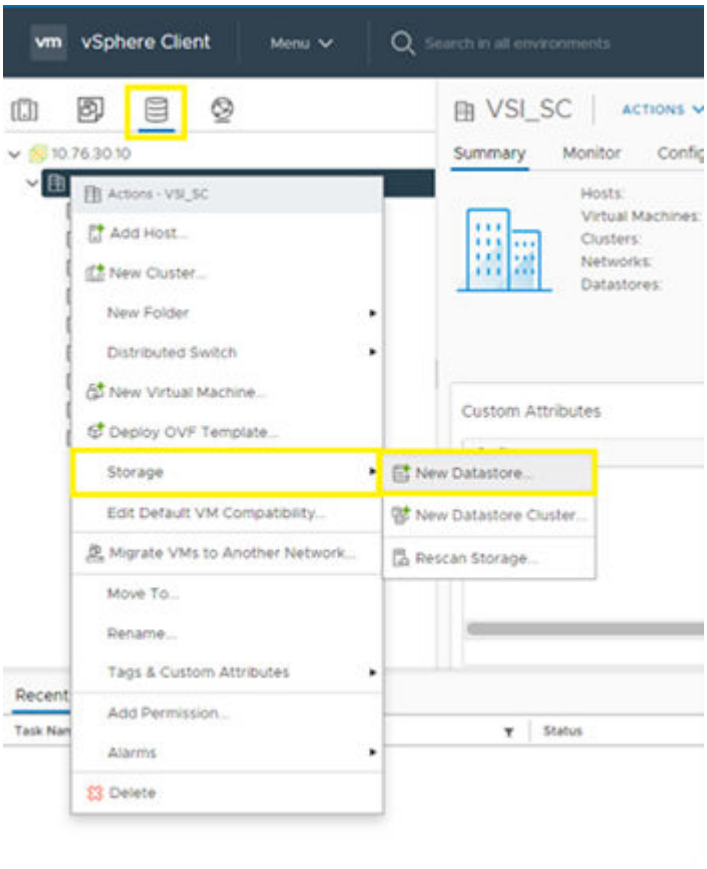

4. Select **vVol** as the datastore type. Click **Next**.

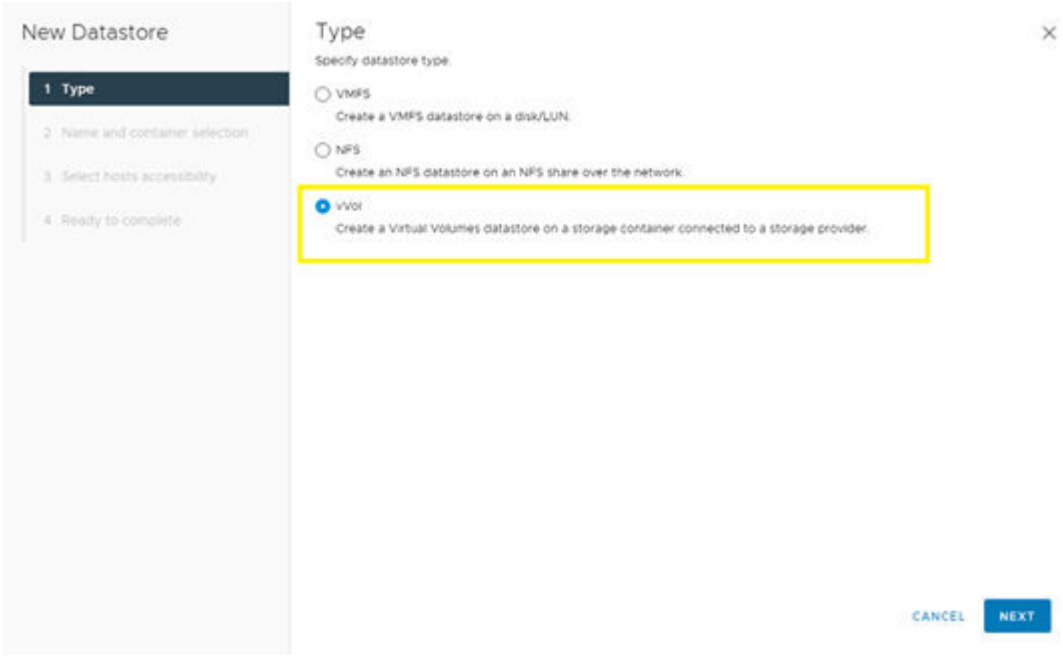

5. Define a **Datastore name**, and then select the appropriate backing storage container.
6. Click **Next**.

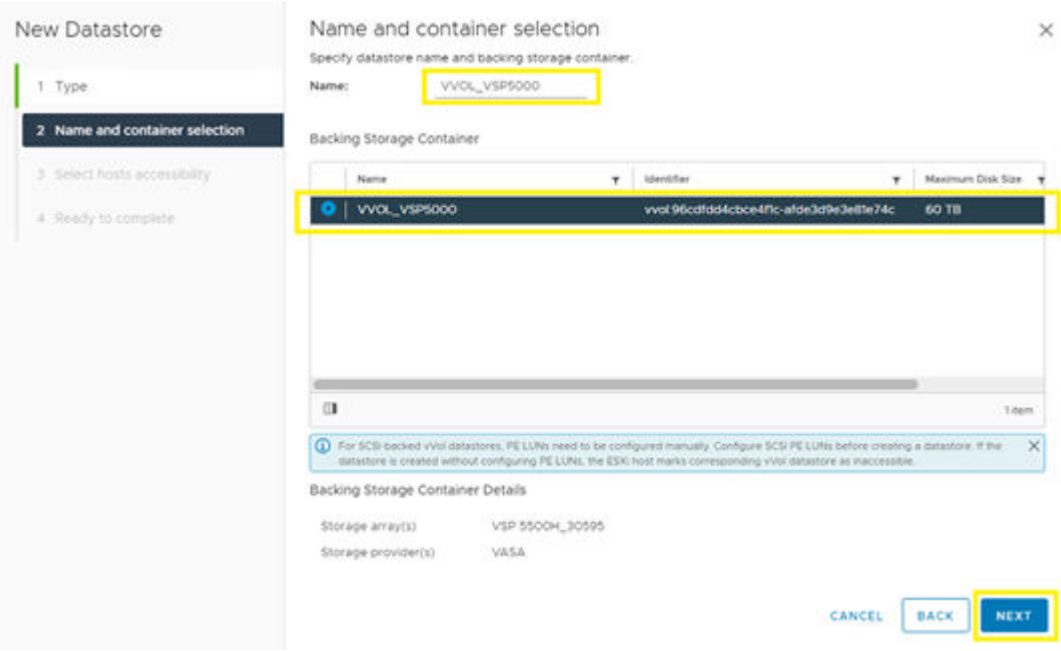

7. Select all ESXi hosts in the cluster, and then click **Next**.
8. Review settings, and then click **Finish**.

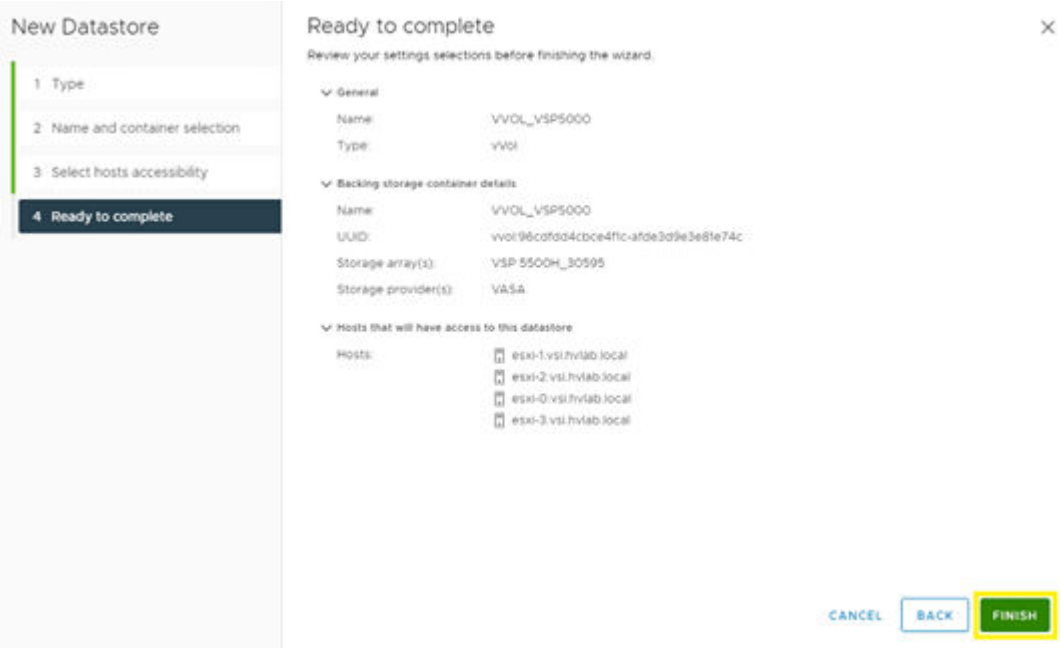

VMware vCenter storage policies

VMware storage policies must be configured prior to creating a StorageClass. This section describes how to create storage policies for both VMFS and vVols datastores backed by Hitachi storage with capabilities passed down from the VASA APIs.

VMFS storage policy

To create a VMware storage policy for a VMFS datastore:

Procedure

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

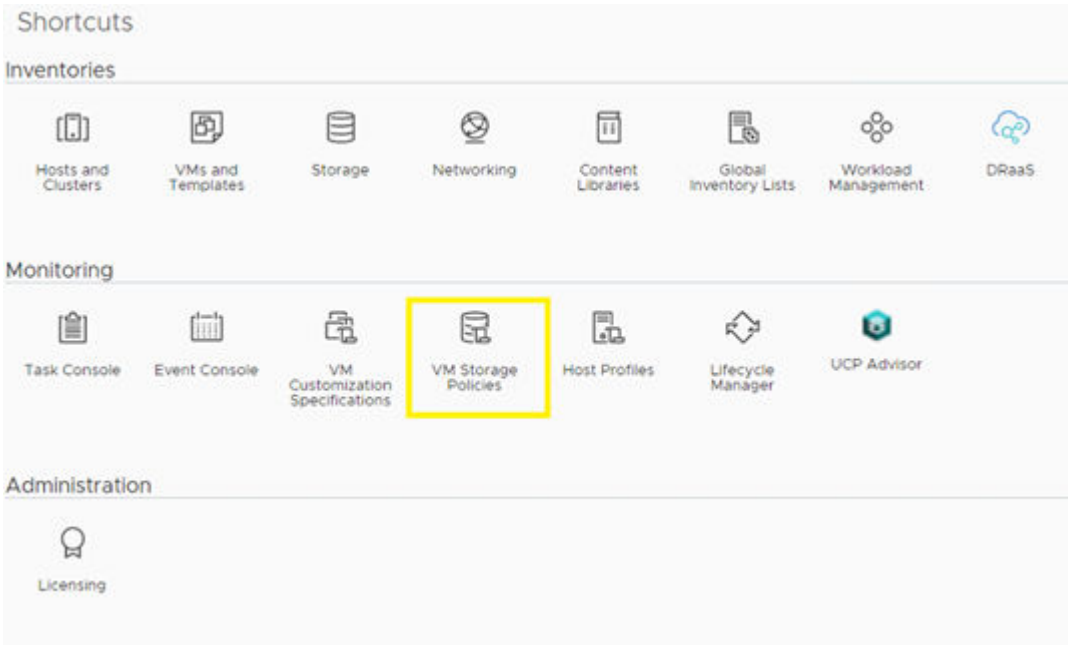

3. Select **CREATE** under **VM Storage Policies**.

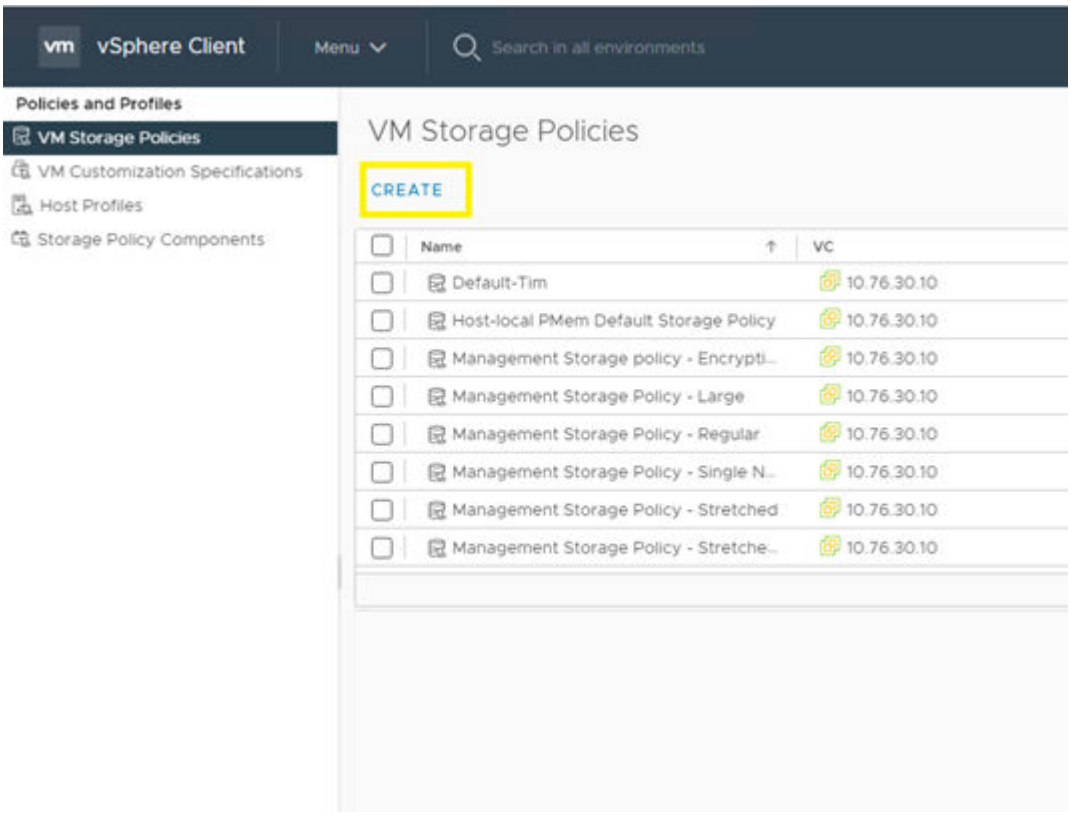

4. Define the **policy name**, and then click **Next**.

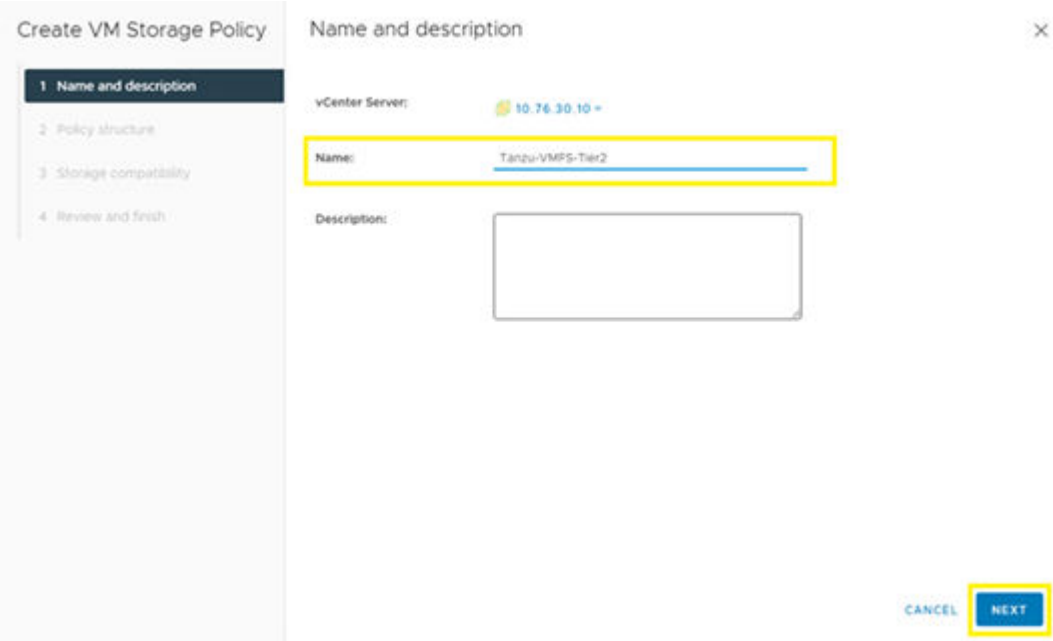

5. For Datastore specific rules, select **Enable rules for "com.hitachi.storageprovider.vvol" storage**, and then click **Next**.

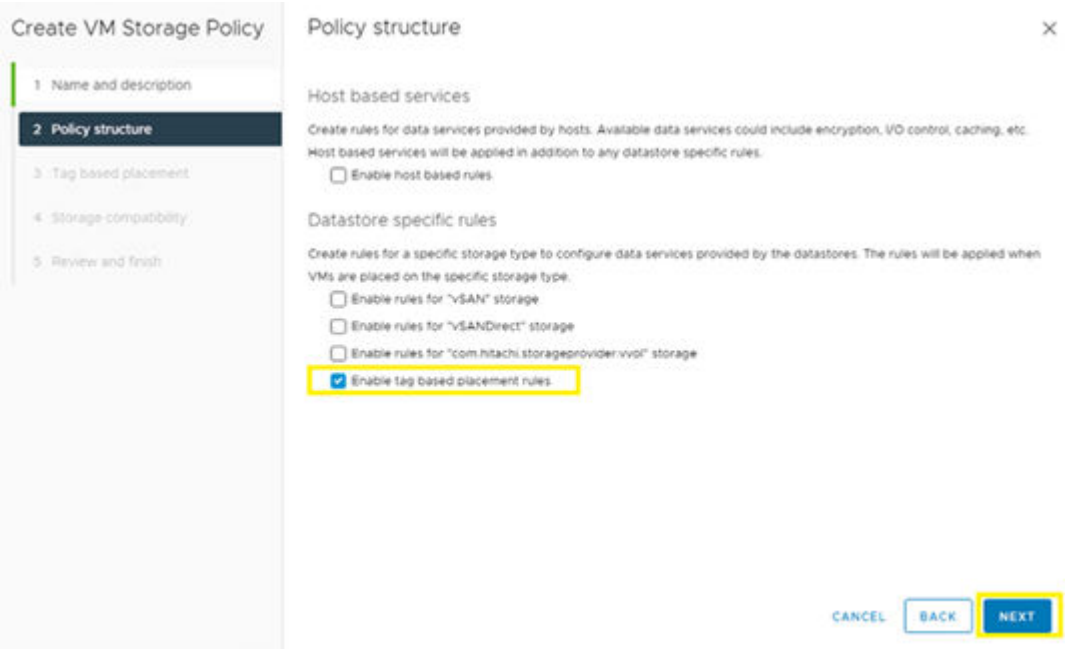

6. Under **Tag-based placement** select the following:
   a. Tag category: **SPBM**
   b. Usage: **Use storage tagged with**
   c. Select **BROWSE TAGS**
7. Select the storage tags that the storage administer has defined via the Storage Provider for VMware vCenter, and then click **OK**. Click **Next**.

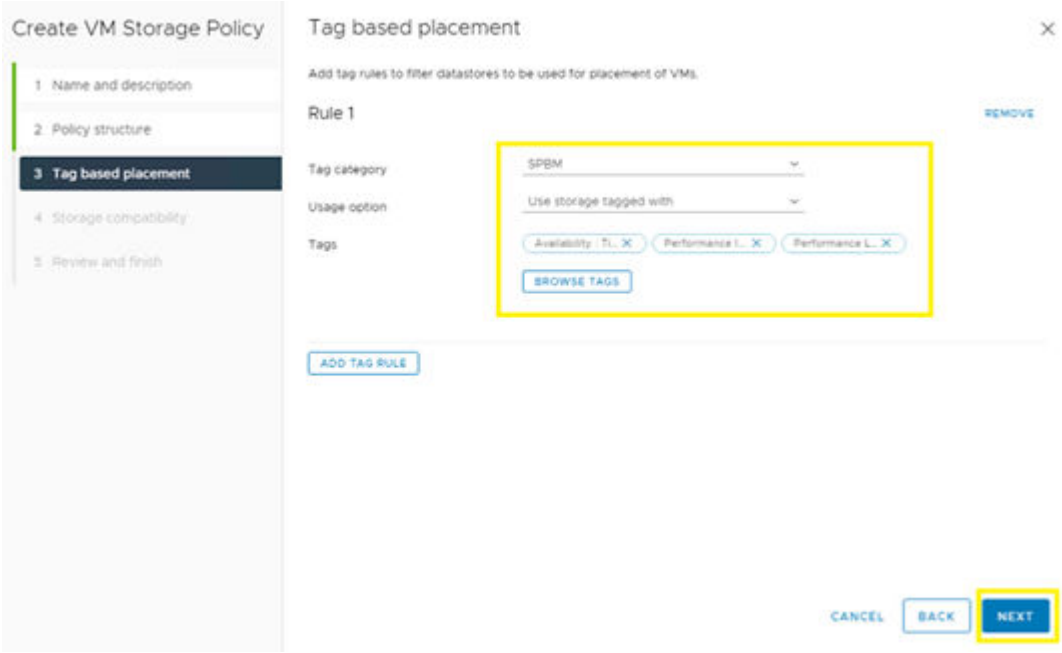

8. The Storage compatibility window displays the datastores that match the tags that you enabled in the previous step. Click **Next**.

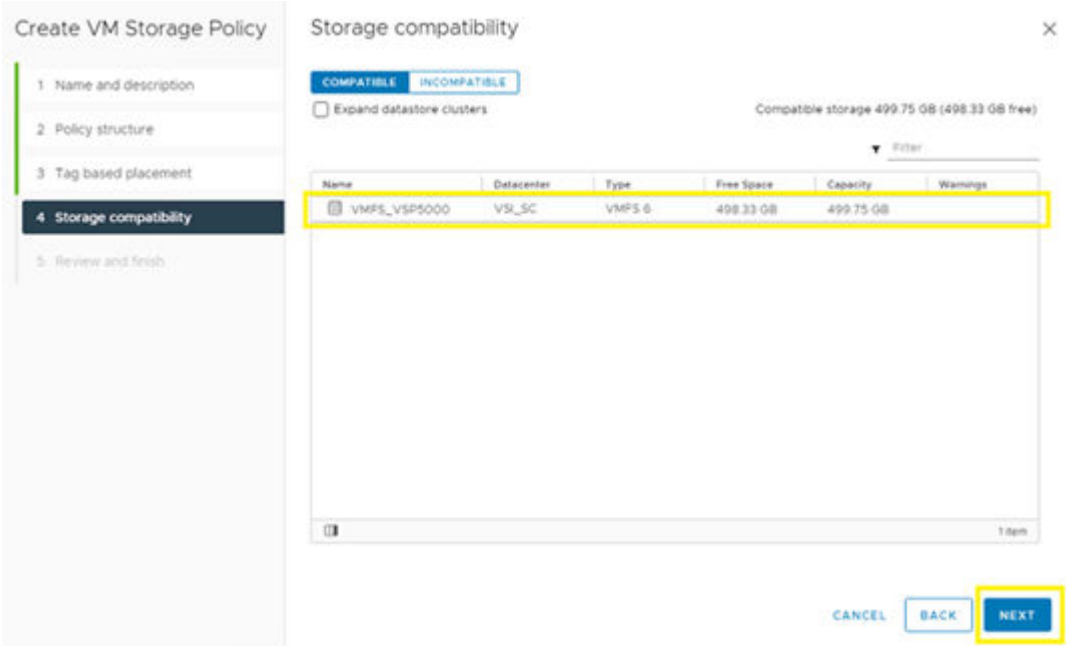

9. Click **Finish**.

vVols storage policy

To 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**.

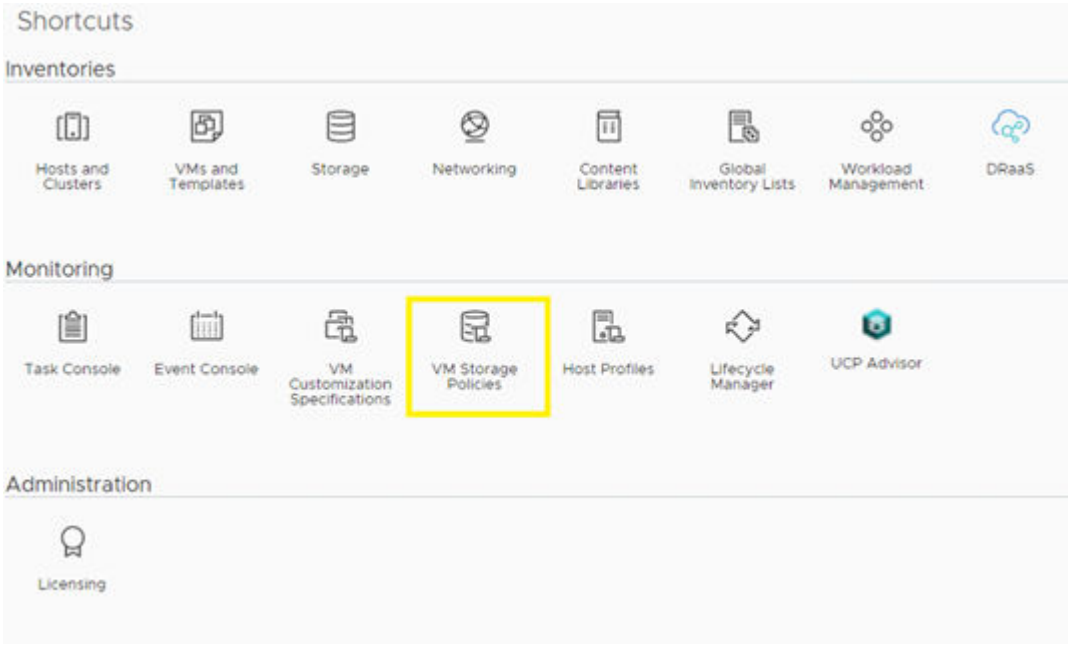

3. Select **CREATE** under **VM Storage Policies**.

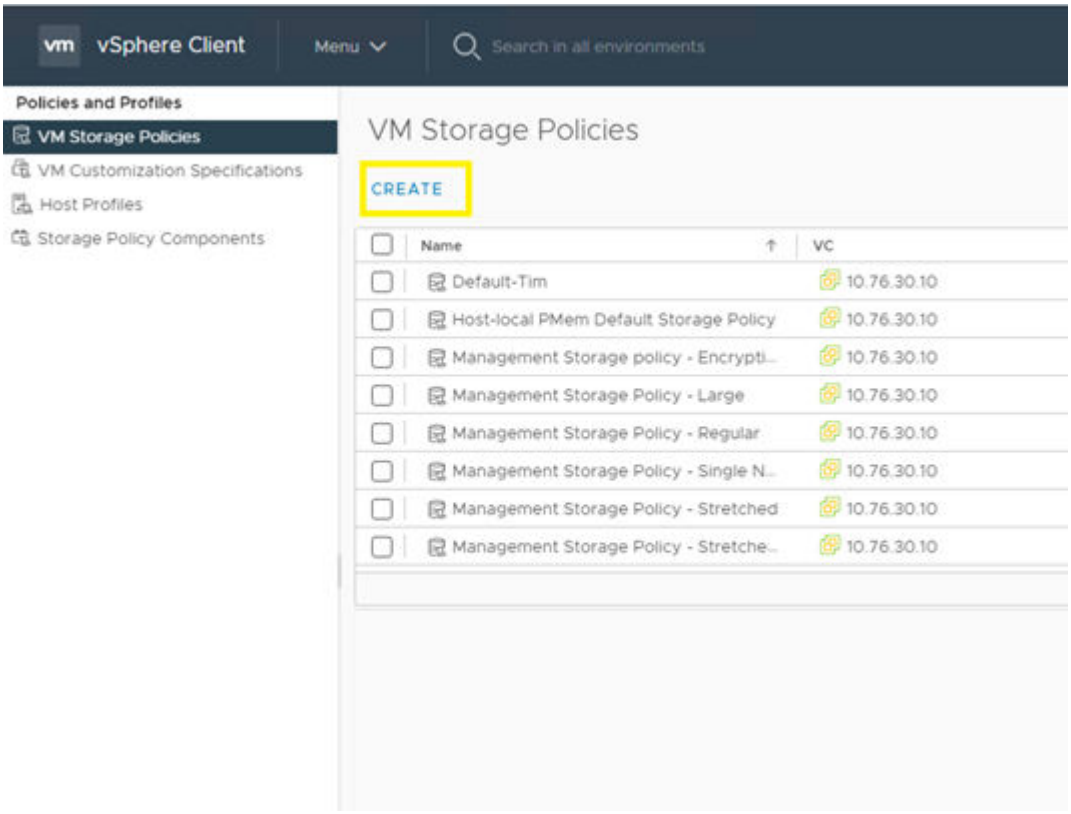

4. Define the **policy name**, and then click **Next**.

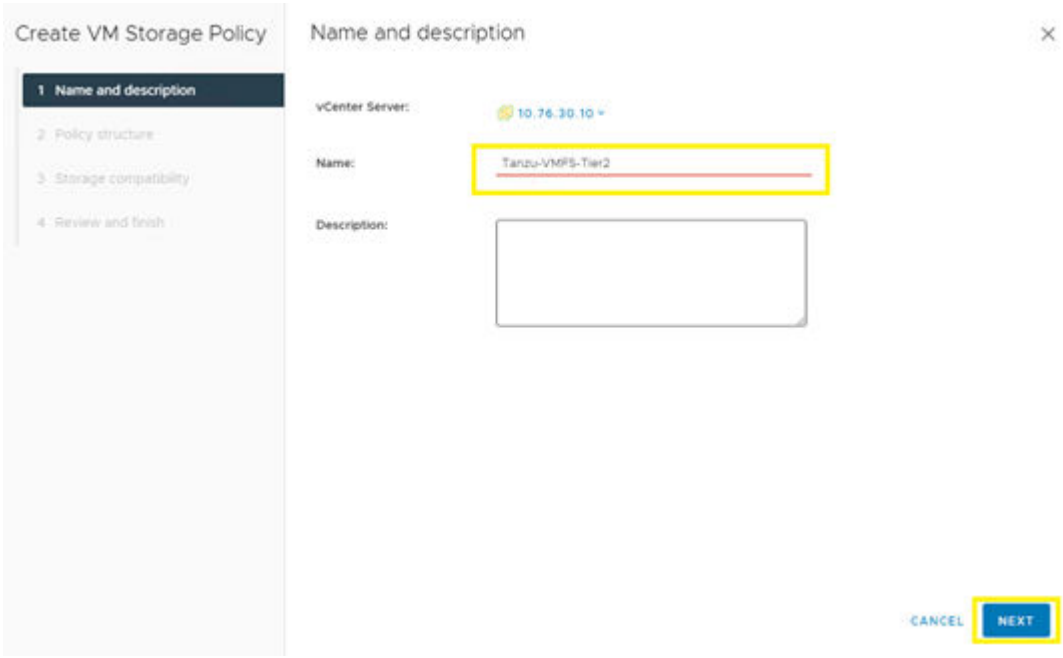

5. For Datastore specific rules, select **Enable rules for "com.hitachi.storageprovider.vvol" storage**, and then click **Next**.

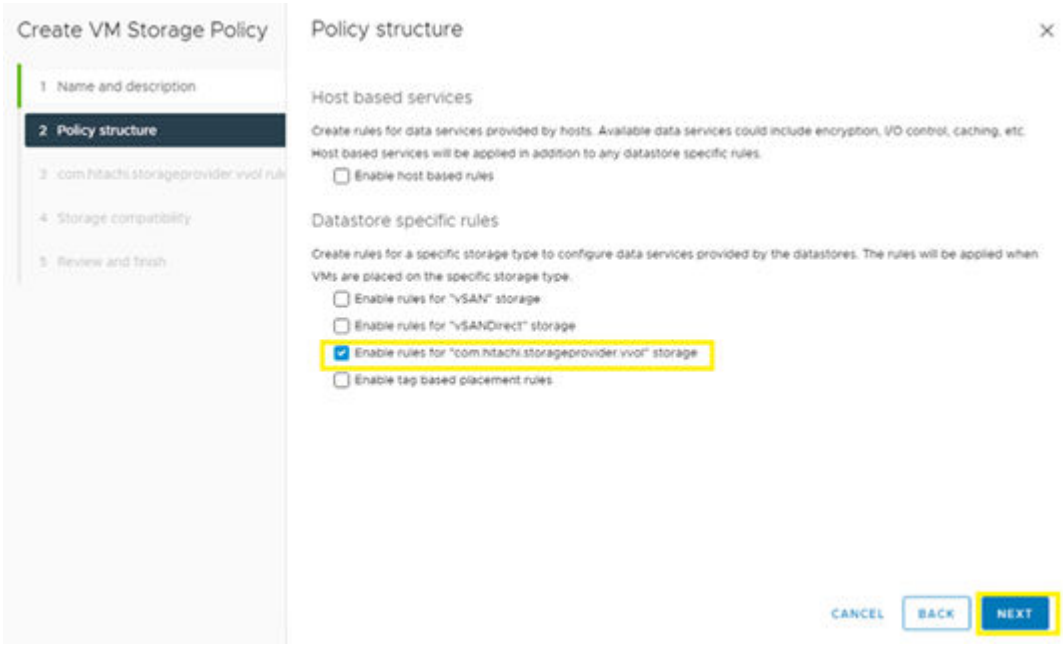

6. Click **ADD RULE**.
7. From the ADD RULE list, select the appropriate rule passed down from the Storage Provider for VMware vCenter, and then click **Next**.

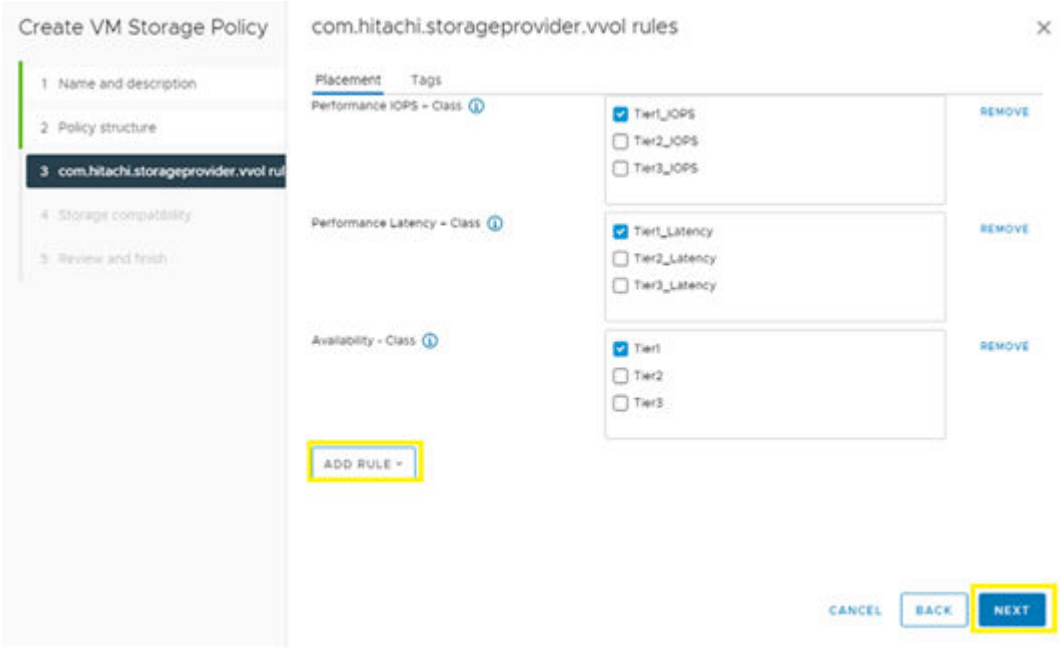

8. The Storage compatibility window displays the vVols datastores that match the capabilities that you enabled previously. Click **Next**.

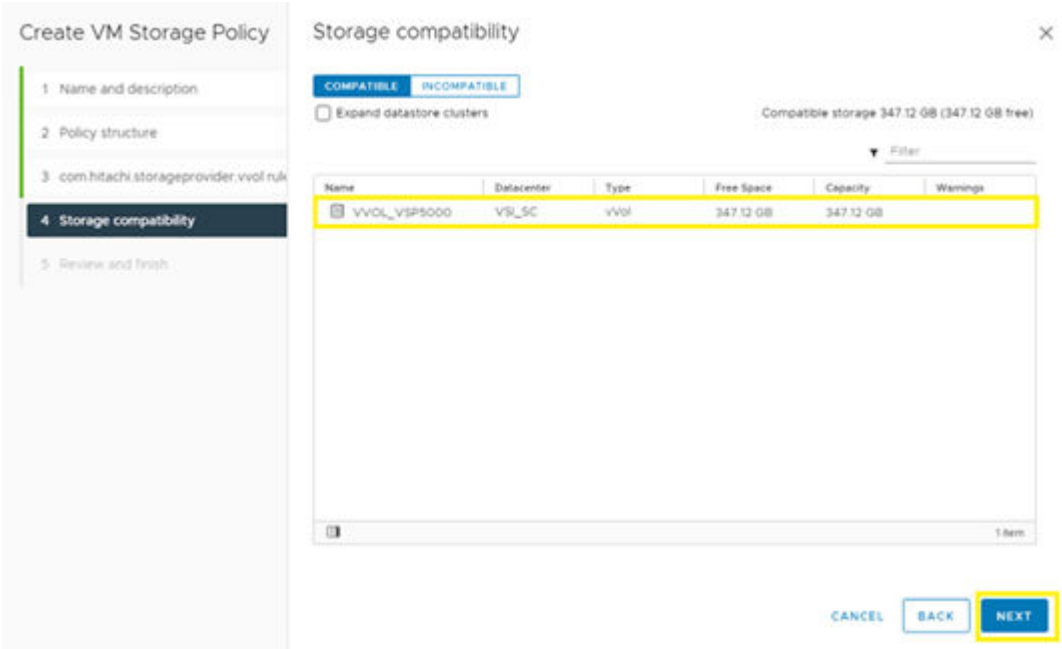

9. Click **Finish**.

Tanzu content library subscription

Before deploying a Tanzu Kubernetes cluster, a content library must be created pointing to the latest available images provided by VMware.

To create and link a content library:

Procedure

1. Log in to VMware vSphere Client.
2. From the Home directory, select **Content Libraries**.

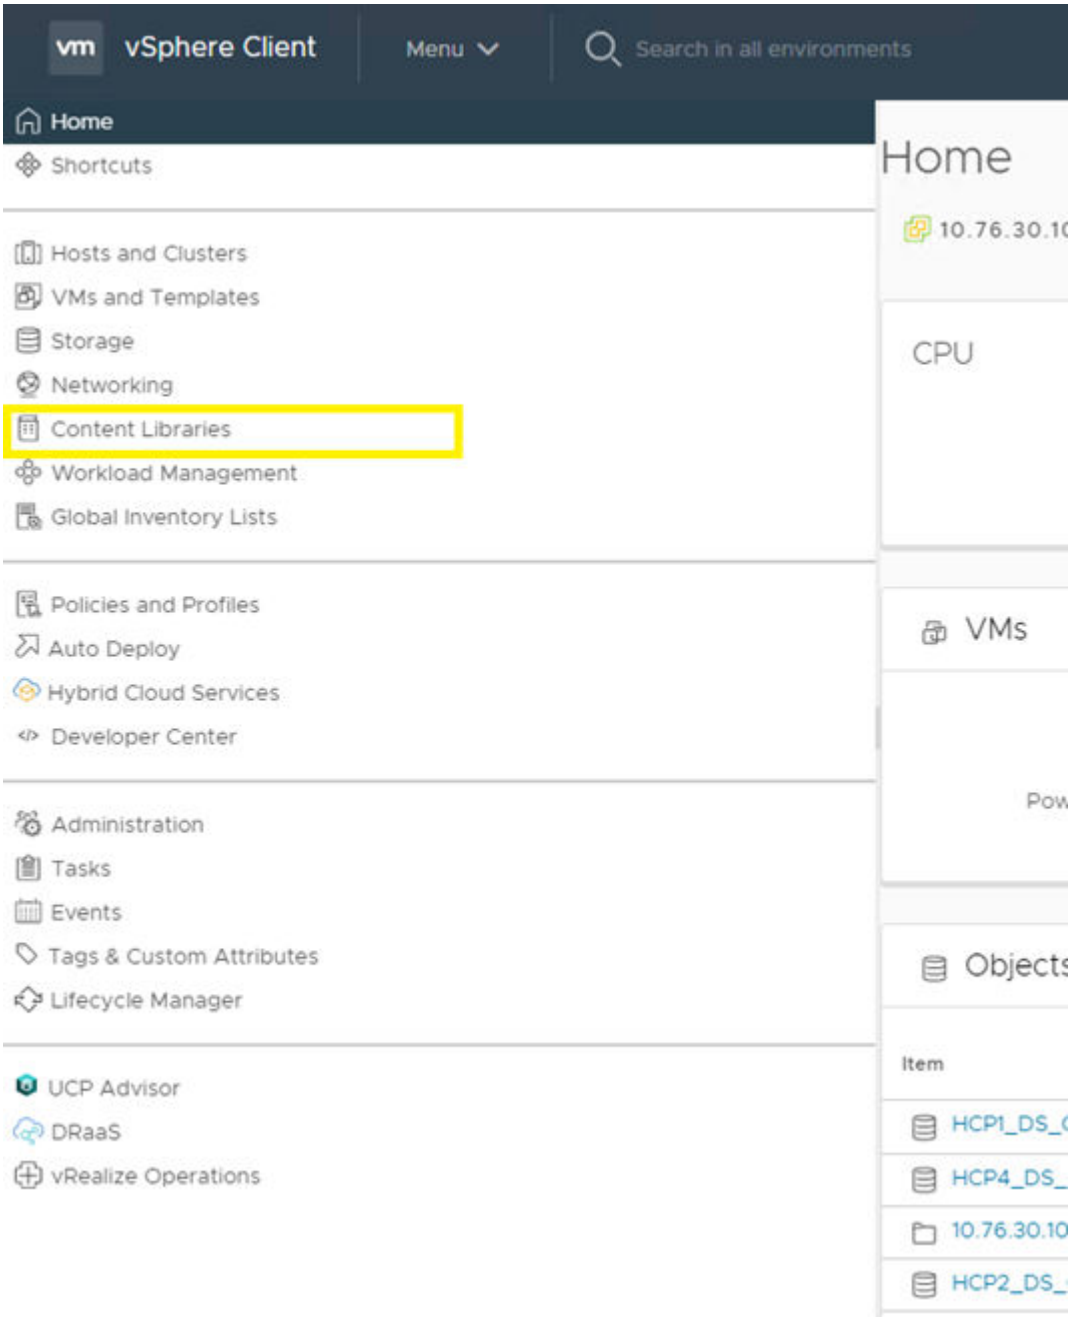

3. Click **+ Create** to create a new content library.

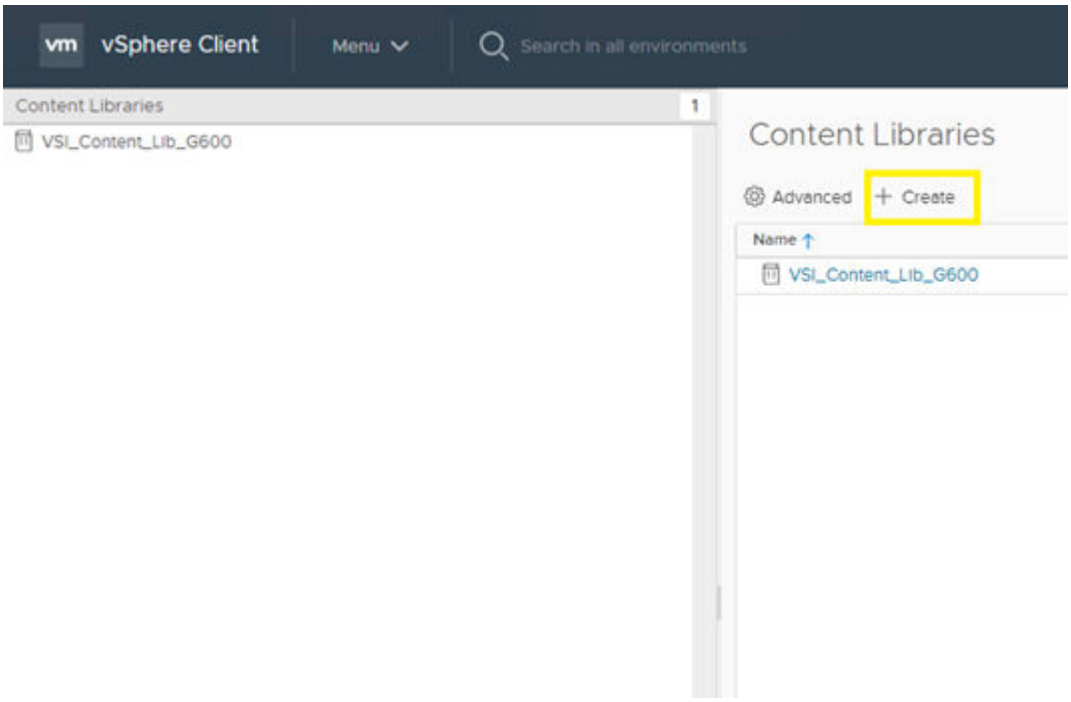

4. Define a **Name**, and then click **Next**.

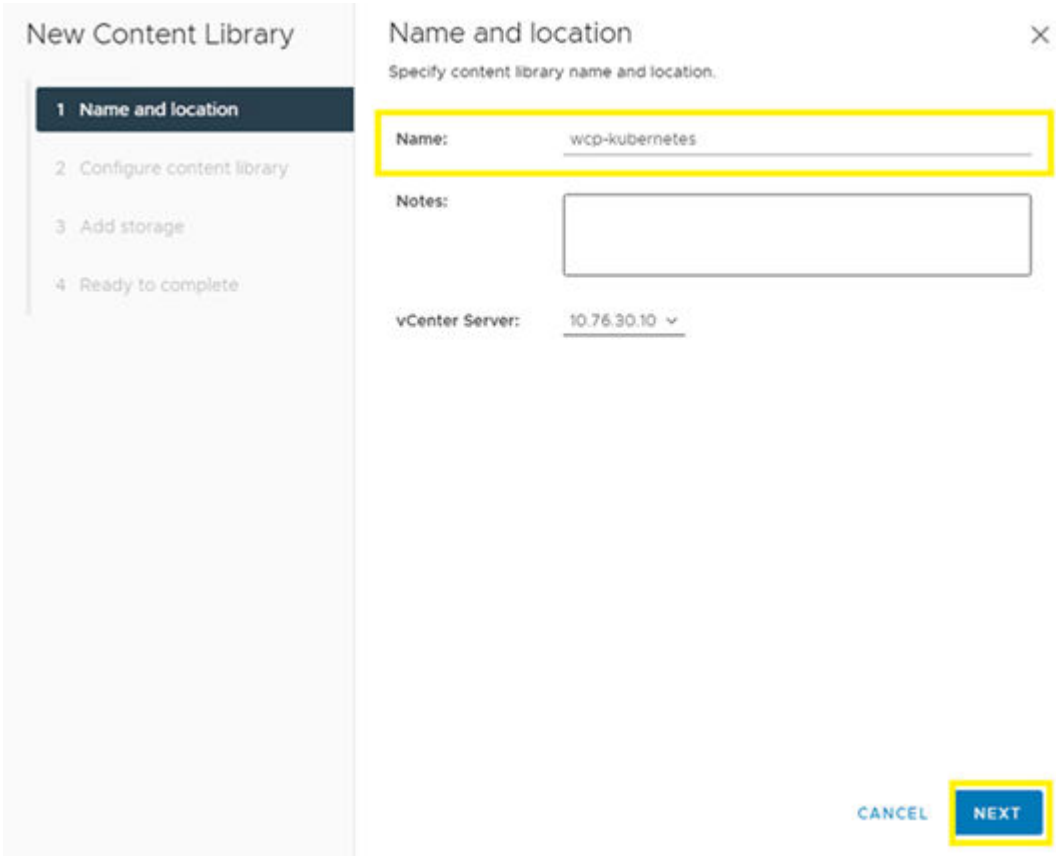

5. Select **Subscribed content library**, define the subscription URL as https://wp-content.vmware.com/v2/latest/lib.json.
6. Under Download content select **immediately**. Click **Next**.

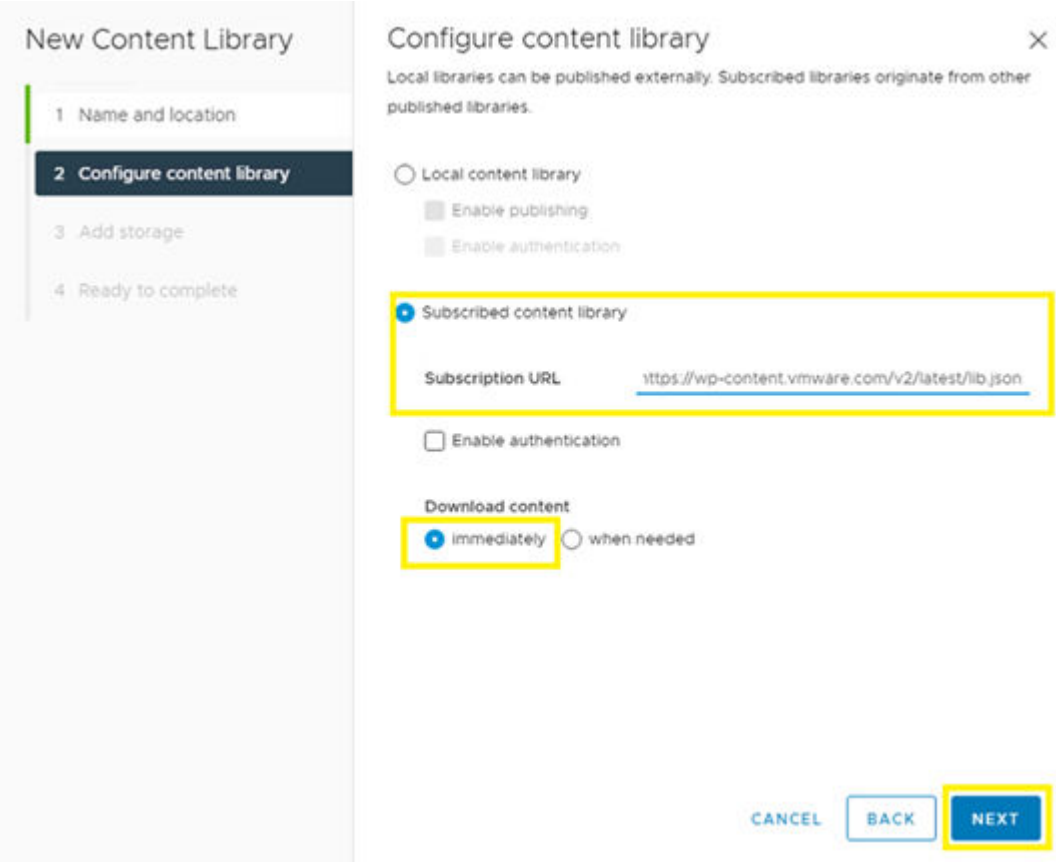

7. Click **YES** to confirm in the verification popup.
8. Select an available environment datastore to assign to the content library. Click **Next**.
9. Click **Finish**.

HAProxy deployment and configuration

Before configuring workload management HAProxy must be deployed and configured. To obtain the latest version of HAProxy go to https://github.com/haproxytech/vmware-haproxy#download.

Note: HAProxy v1.10 was used for the following example.

To deploy and configure HAProxy follow these steps:

Procedure

1. Log in to VMware vSphere Client.
2. From the **Host and Clusters** view, right click and select **Deploy OVF Template**.

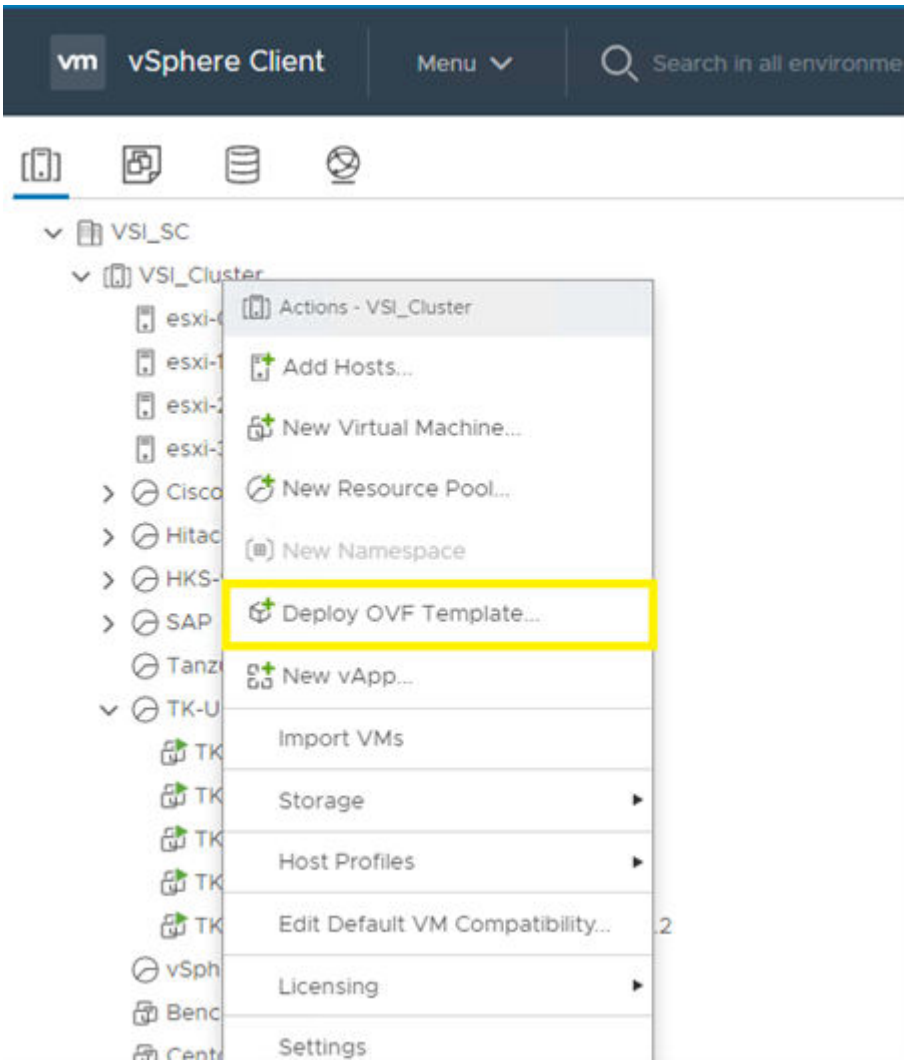

3. Select **Local File**, and upload the HAProxy OVA. Click **Next**.

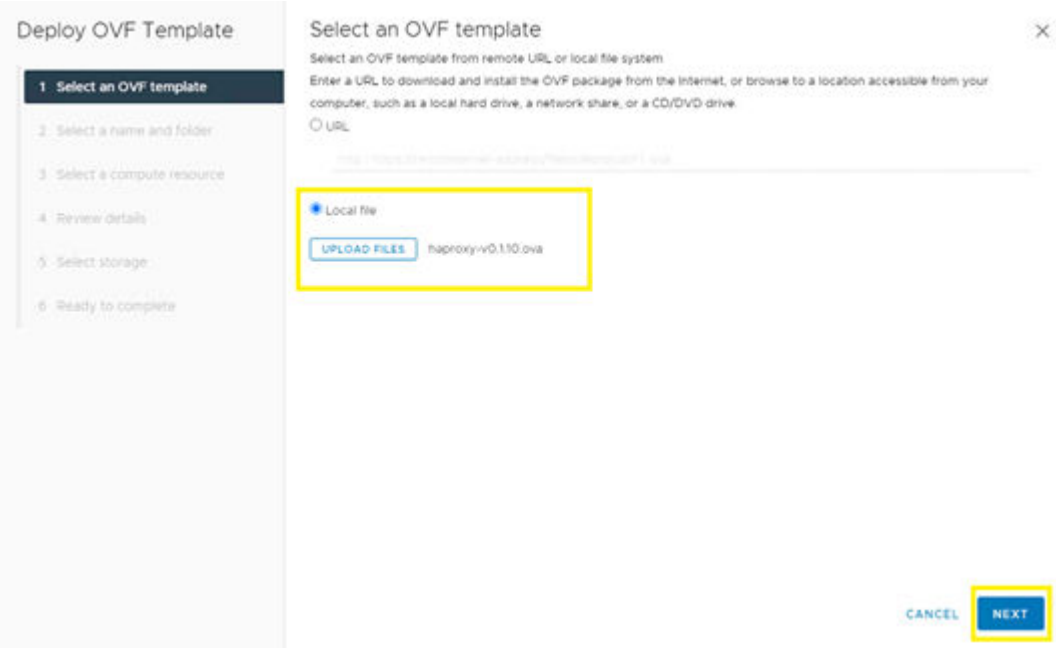

4. Define a **VM name** and select the associated datacenter. Click **Next**.

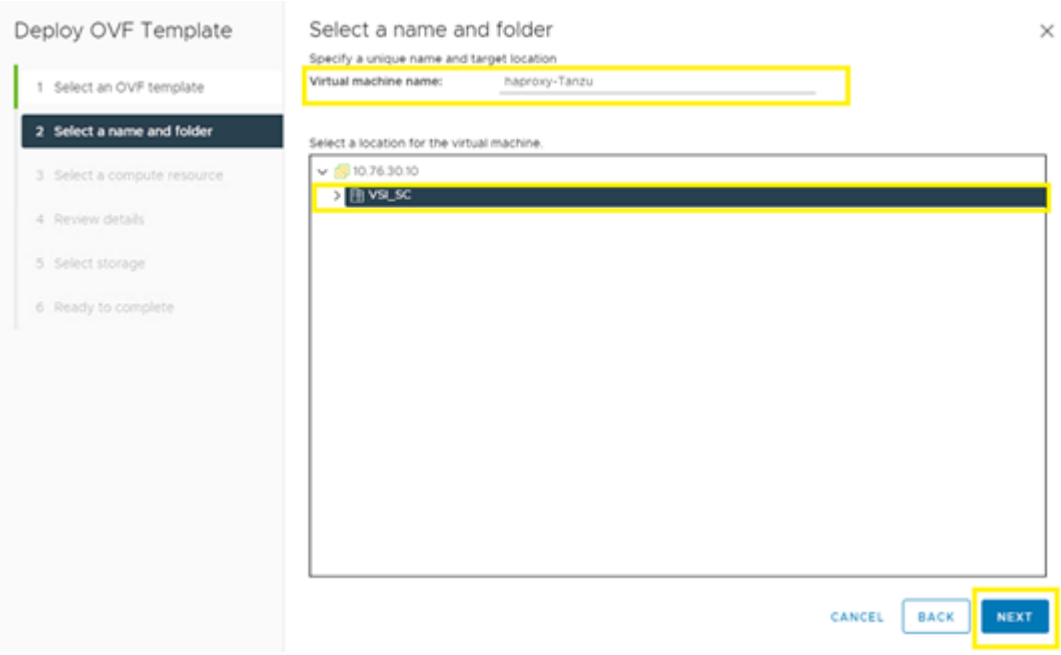

5. Select the compute resource, and then click **Next**.
6. Via the Review details window click **Next**.
7. Select **I accept all license agreements**, and then click **Next**.
8. From the configuration menu select **Default** or **Frontend Network**, and then click **Next**.

> **Note:** Within this guide, a default network configuration was used where the appliance was deployed with 2 NICs: A Management network (Supervisor -> HAProxy dataplane) and a single Workload network. Load-balanced IP addresses are assigned on the workload network.

9. Select an available datastore for HAProxy deployment, and then click **Next**.
10. From the Select networks menu, choose **Management** and **Workload** networks. Click **Next**.

> **Note:** For a default deployment, the frontend destination network will be ignored.

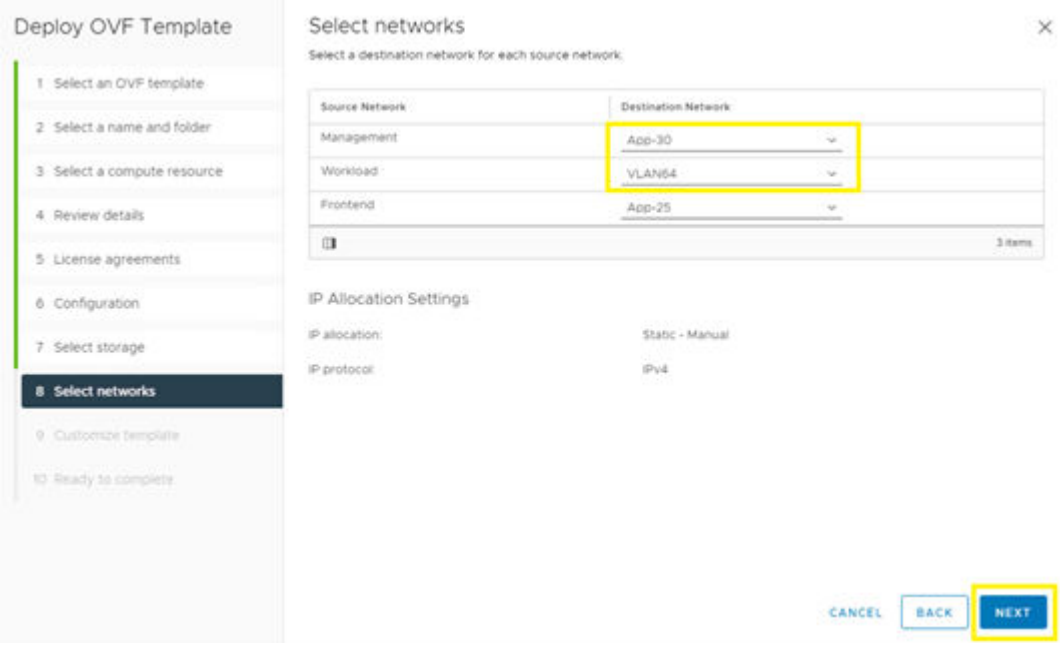

11. Via Customize template configure the following:
    a. Enter a root **password**.
    b. Select the **enable root login** option.
    c. Leave **TLS Certificate Authority Certificate/Private** key blank.
    d. Define a **hostname**.
    e. Enter an appropriate **DNS** server.
    f. Enter a **management IP address in CIRD format**.
    g. Define the **management IP gateway**.
    h. Define the **static workload IP in CIRD format**.
    i. Define the **workload gateway**.
    j. Define the **load balancer IP range in CIRD format**.
    k. Leave the **data plane management port set to 5556**.
    l. Define a **HAProxy user and password**.
12. Click **Next**.

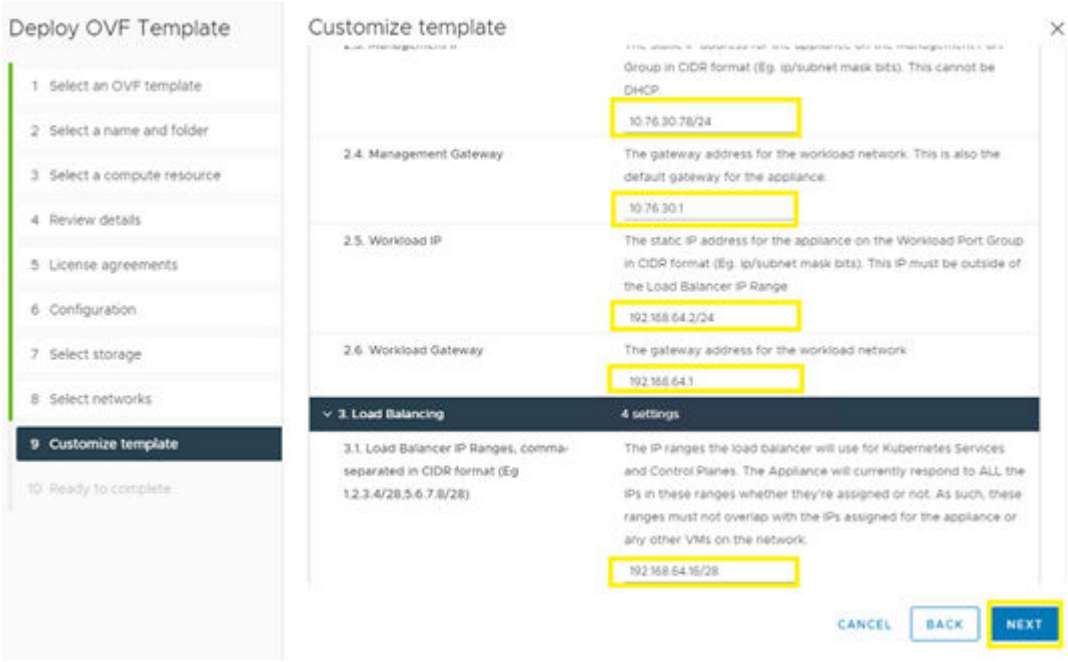

13. Review settings, and then click **Finish**.

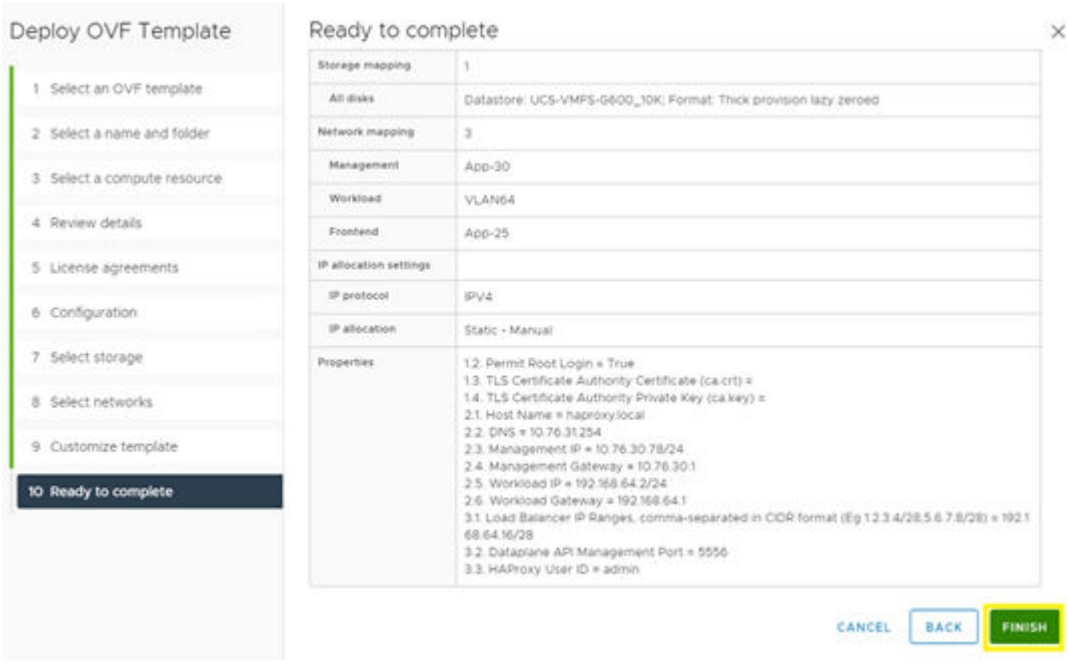

14. Power on the VM. and continue to the next section.

VMware workload management configuration

This section describes how to deploy a workload management cluster via VMware vSphere 7.0u2.

Tanzu supervisor cluster deployment

To deploy a workload management supervisor cluster, follow these steps:

Procedure

1. Log in to VMware vSphere Client.
2. From the Home menu click **Workload Management**.
3. If the evaluation license is used, enter activation information, and then click **GET STARTED**.
4. Select the **vCenter server** and **vCenter Server Network**. Click **Next**.

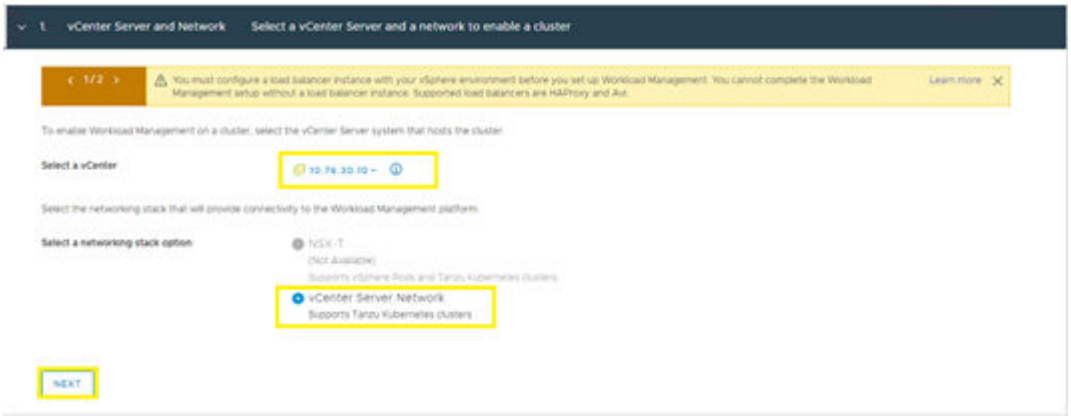

5. Select an available cluster, and then click **Next**.

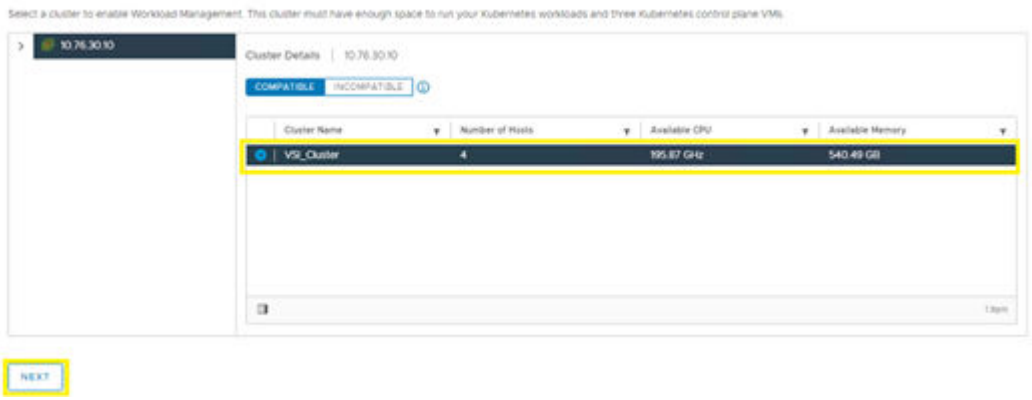

6. Select the supervisor cluster resource settings. Click **Next**.

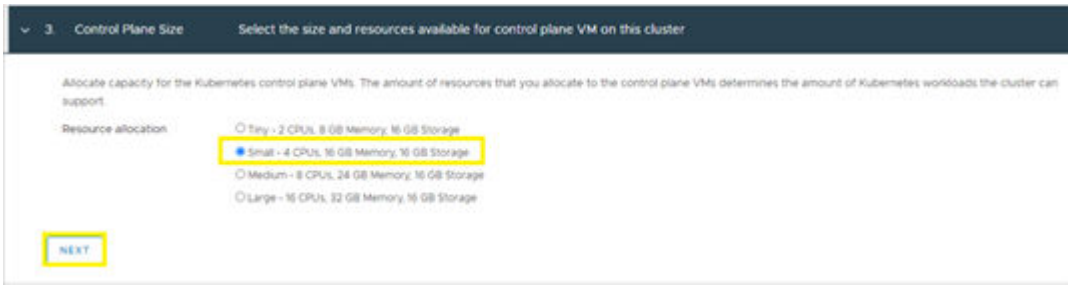

7. Select a storage policy. Click **Next**.

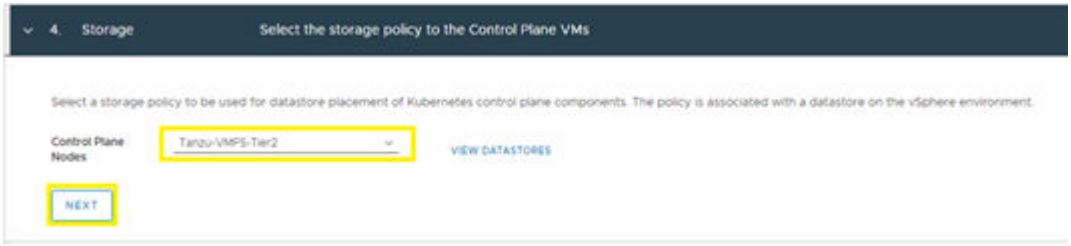

8. For the load balancer, configure the following:
   a. Define a load balancer **Name**.
   b. From the **Type** list select **HAProxy**.
   c. Enter the **Data Plane IP API Address(es)** followed by port **5556**.
   d. Enter the **username** and **password** defined during HAProxy deployment.
   e. Define the **Virtual IP address range**, this is the same range defined in CIRD format during HAProxy deployment.
   f. SSH as the root user to the static IP of the HAProxy VM using the password defined during deployment.
   g. Run *cat /etc/haproxy/ca.crt* to copy the certificate authority and past it into **Server Certificate Authority**.
9. Click **Next**.

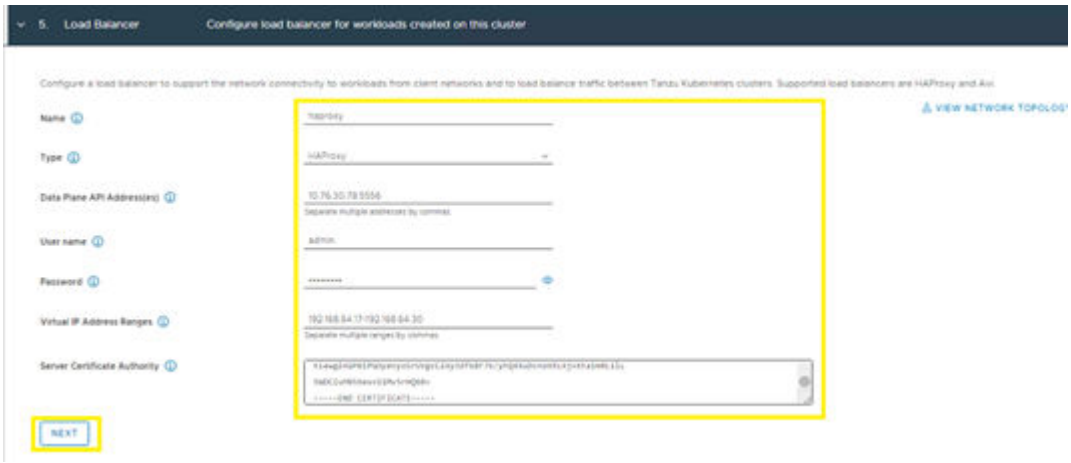

10. Via the management network configuration define the following:
    a. Select the **Management Network** from the **Network** list.
    b. Enter the **Starting IP Address** for the management network for the supervisor cluster.
    c. Enter the **Subnet Mask**.
    d. Define the **Gateway**.
    e. Enter an applicable **DNS server**.
    f. Define an **NTP** server.
11. Click **Next**.

12. From the Workload Network configuration window, enter an available **DNS server**, and then click **Add** to define the workload network.

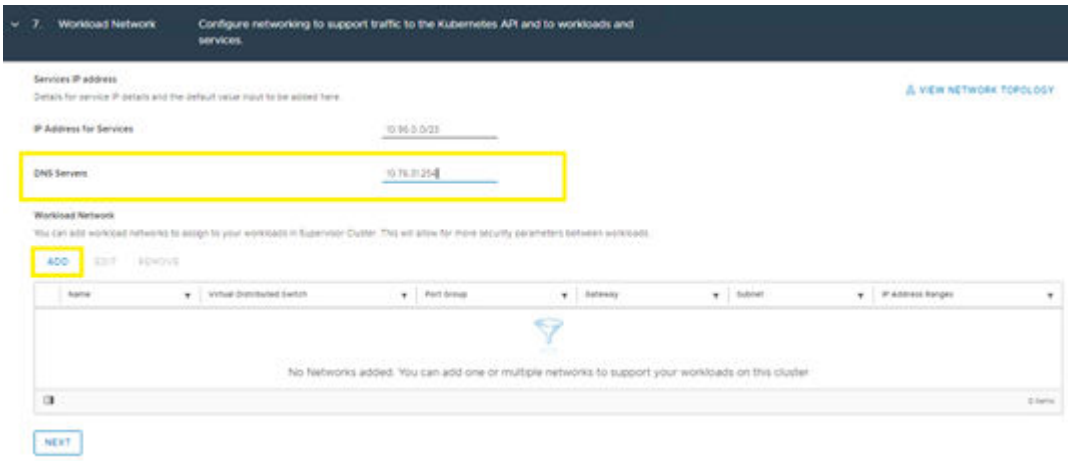

13. Define the following:

a. Enter a **Name**.

b. Select the workload network **Port Group**.

c. Define the workload network **Gateway**.

d. Define the **IP Address Ranges** for the workload IP address range.

e. Click **Save**.

14. Click **Next**.

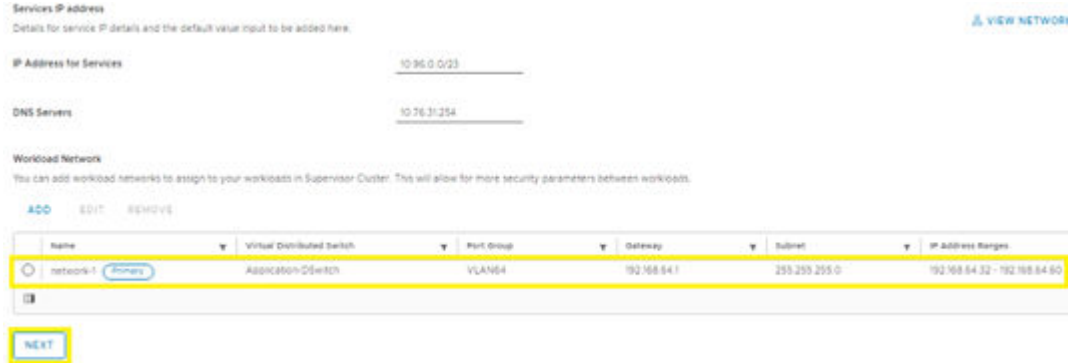

15. Via Tanzu Kubernetes Grid Service Configuration, click **Add** and select the subscribed content library covered in section Tanzu Content Library Subscription (on page 46).
16. Click **Next**.

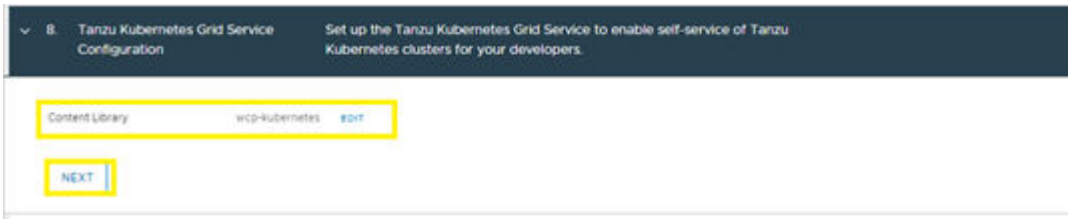

17. Click **Finish** to deploy the Tanzu Supervisor Cluster.

Create a Namespace

Administrators can continue with Namespace creation once the Tanzu supervisor management cluster has been successfully deployed and configured.

To create a Namespace, follow these steps:

Procedure

1. Log in to VMware vSphere Client.
2. From the Home directory select **Workload Management**.

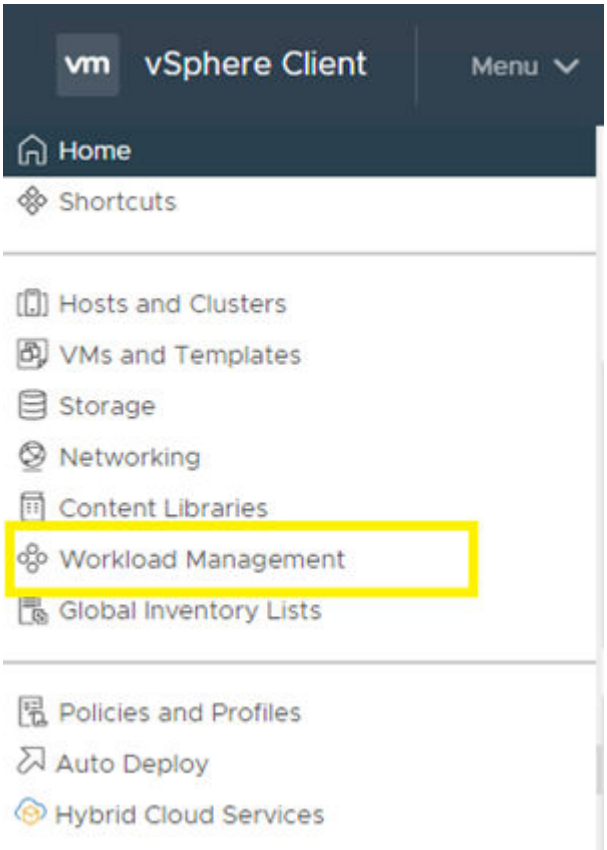

3. Select the **Namespace tab** at the top of the screen.
4. Click **CREATE NAMESPACE**.

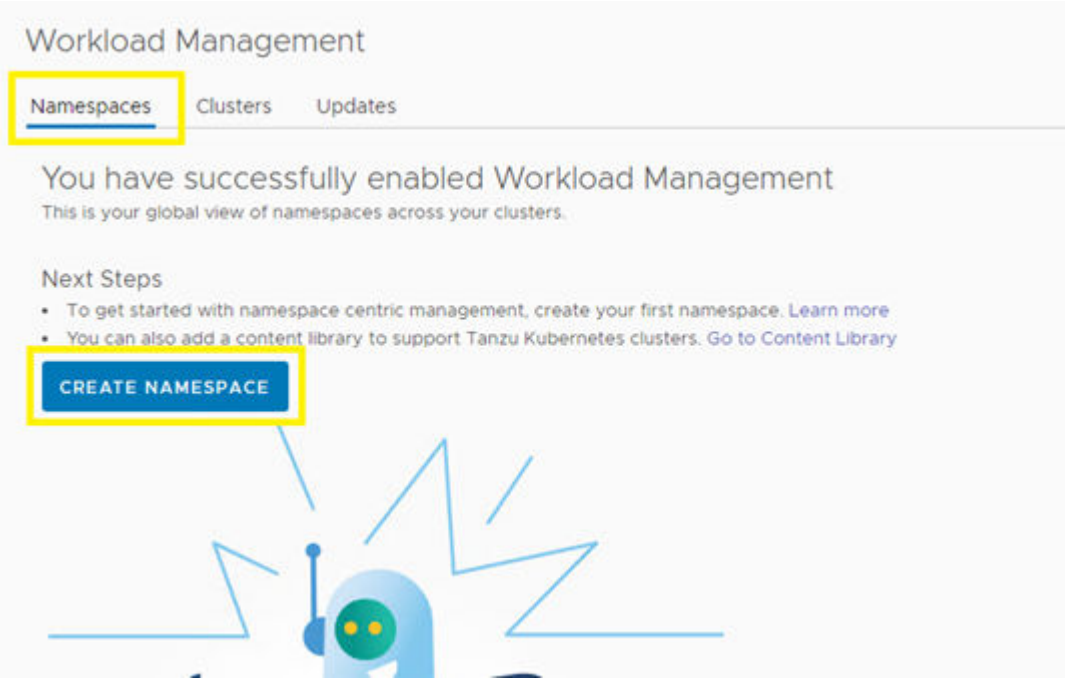

5. Select a **Cluster**, enter a **Name**, and select the workload **Network**. Click **Create**.
6. Click **GOT IT** to begin working with the namespace.

Assign user roles and permissions to supervisor cluster

Once a namespace has been created, the vSphere administrator must assign permissions for themselves as well as any additional developers that plan to use workload management resources. Permissions utilize built-in vSphere users. Prior to configuring permissions, confirm that the above steps have been completed and a namespace has been created.

To begin assigning roles and permissions within a namespace follow these steps:

Procedure

1. Select the appropriate **namespace**.
2. Select the **Summary** tab.
3. Click **ADD PERMISSIONS**.

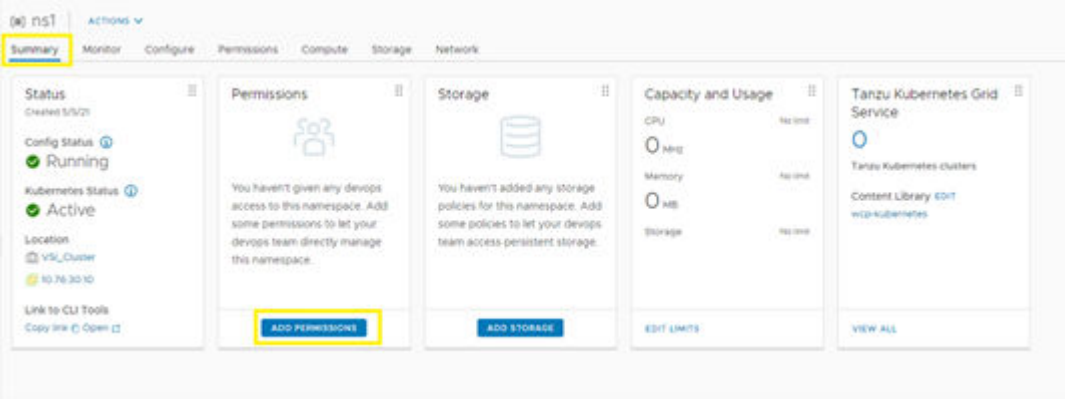

4. Select the following:
   - **vsphere.local** as the **Identity Source**
   - **User** as **Administrator**
   - **Role** with **Can edit permissions**
5. Click **OK**.

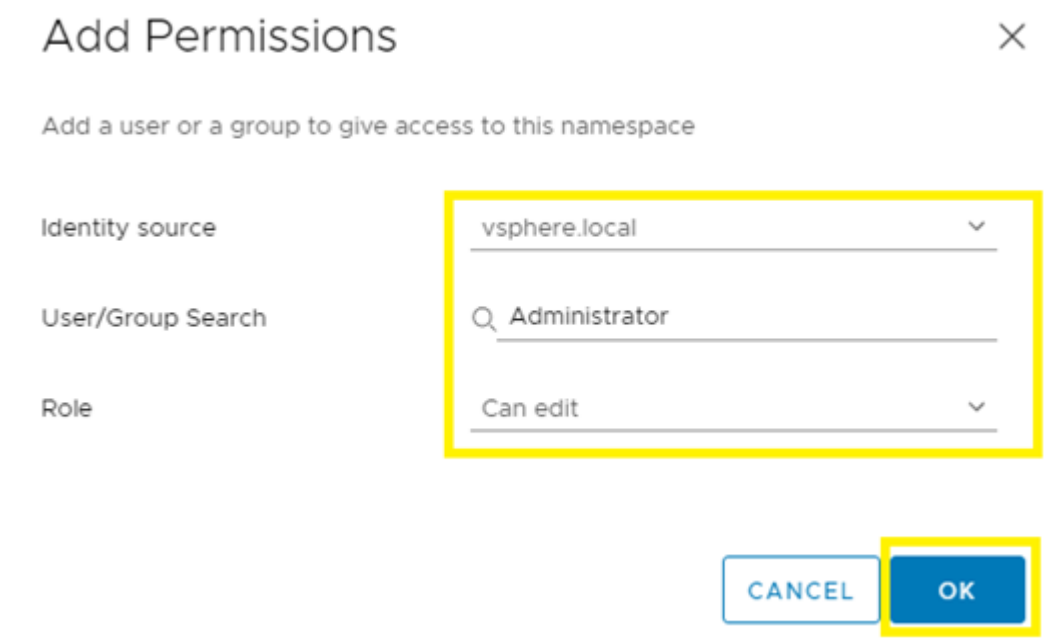

Assign storage policies

Before deploying a Kubernetes cluster, storage policies must be assigned to the namespace.

To assign storage policies to a namespace, follow these steps:

Procedure

1. Select the applicable **Namespace**.
2. Select the **Summary** tab.
3. Click **ADD STORAGE**

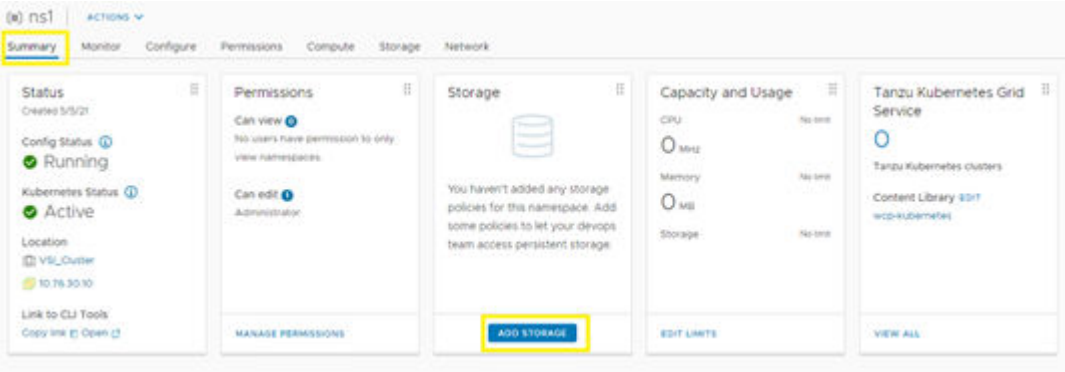

4. Select the CNS polices that were created previously in VMFS Storage Policy (on page 40) and vVols Storage Policy (on page 43). Click **OK**.

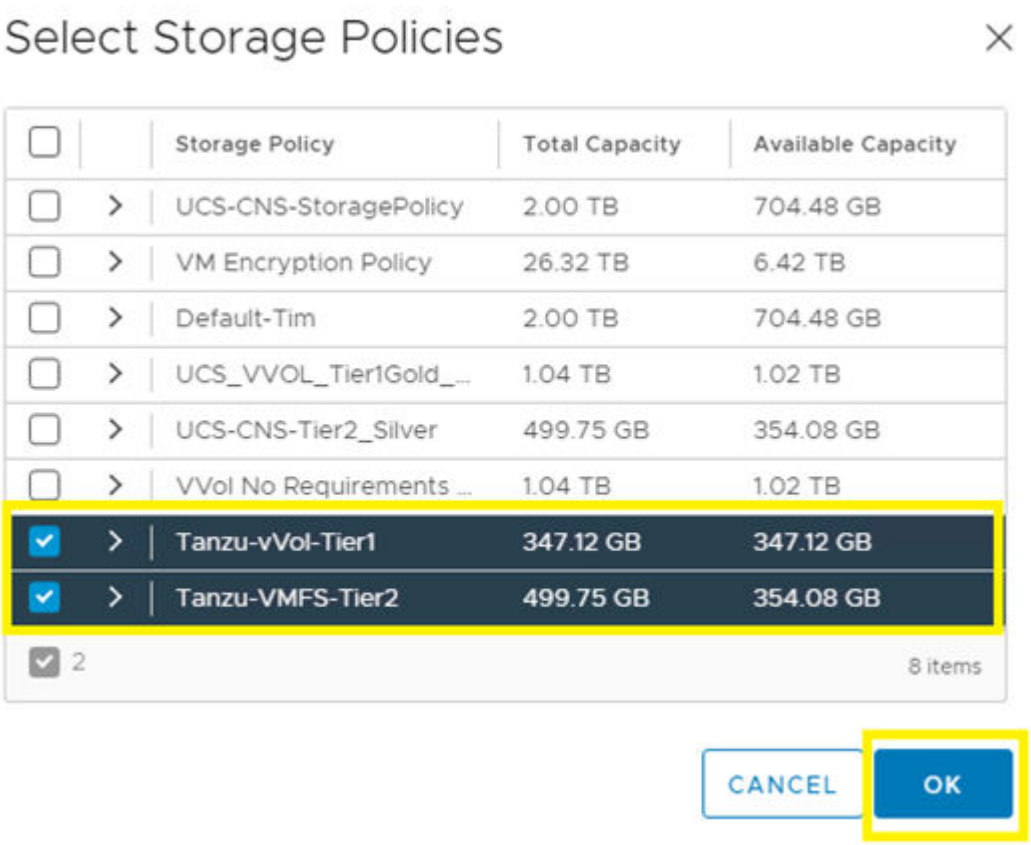

Download operating system CLI tools

Before deploying or working with a Tanzu Kubernetes cluster, administrators as well as developers must download and install Kubernetes CLI Tools for their respective Windows, Linux, or Mac operating system.

To download and install Kubernetes CLI Tools, follow these steps:

Note: The directions in this guide cover the installation of CLI tools using the Windows operating system.

Procedure

1. Select the **Namespace**.
2. Select the **Summary** tab.
3. On the **Status** pane **Link to CLI Tools**, and then click **Open**.

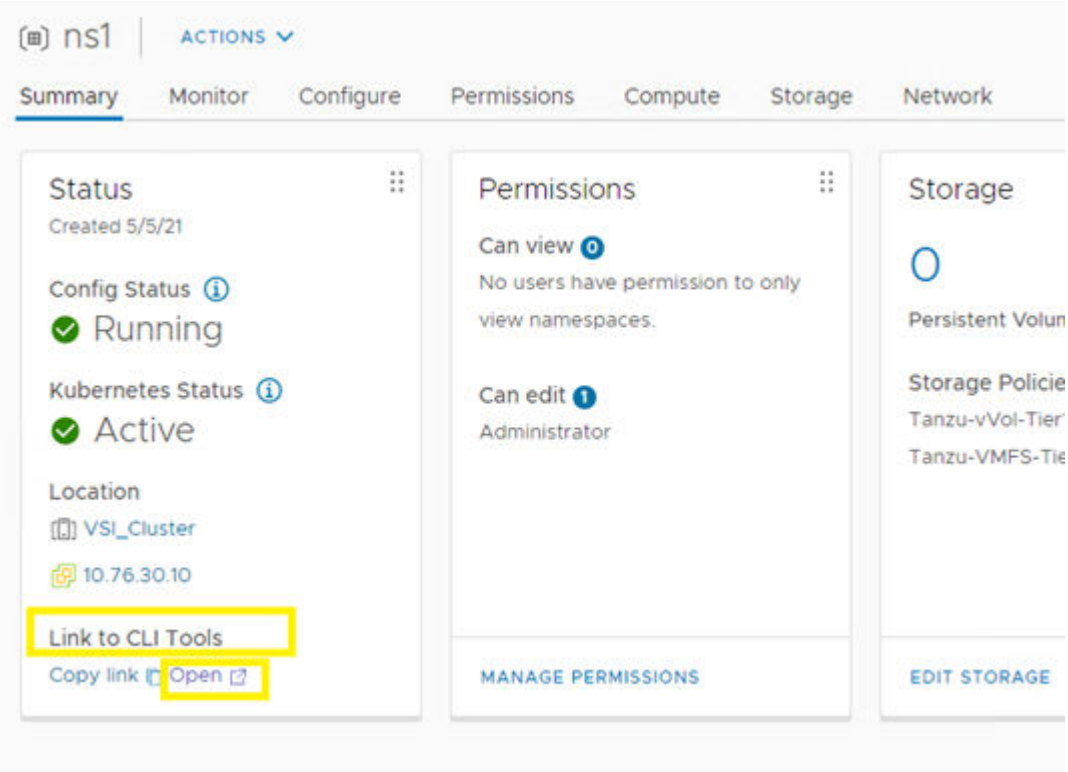

4. Click **DOWNLOAD CLI PLUGIN WINDOWS**.
5. Right click on the downloaded compressed folder and then select **Extract All**.

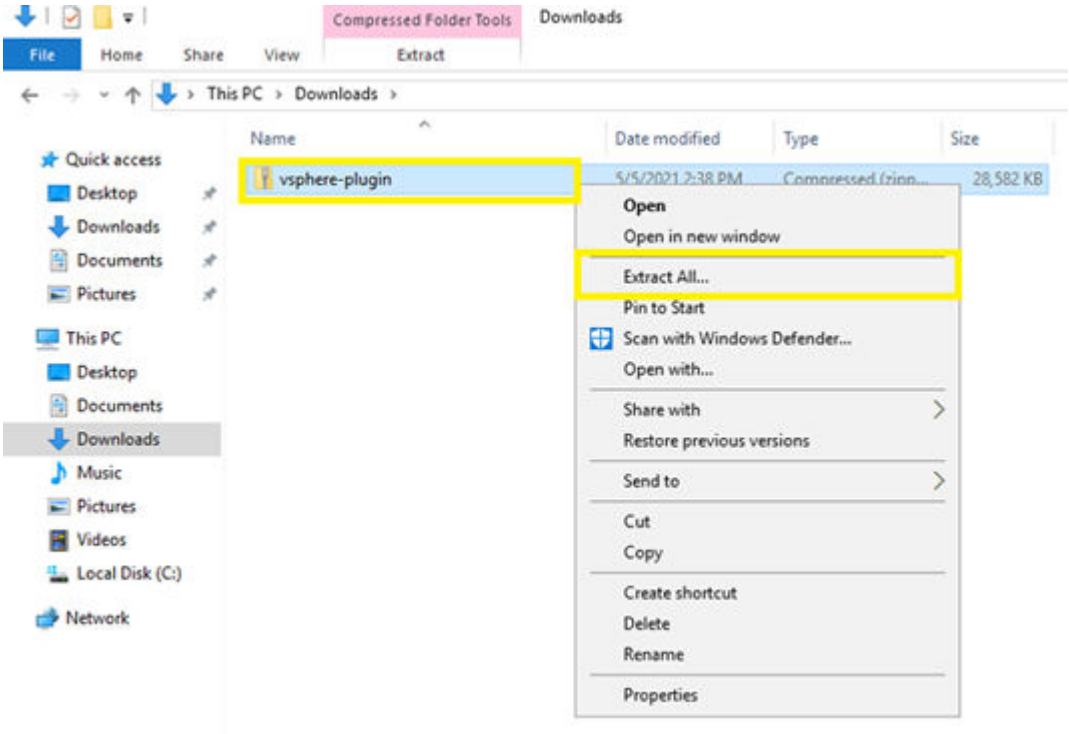

6. Click **Browse** to update the extraction location.
7. Select **Local Disk C Drive**.
8. Click **Select Folder**.

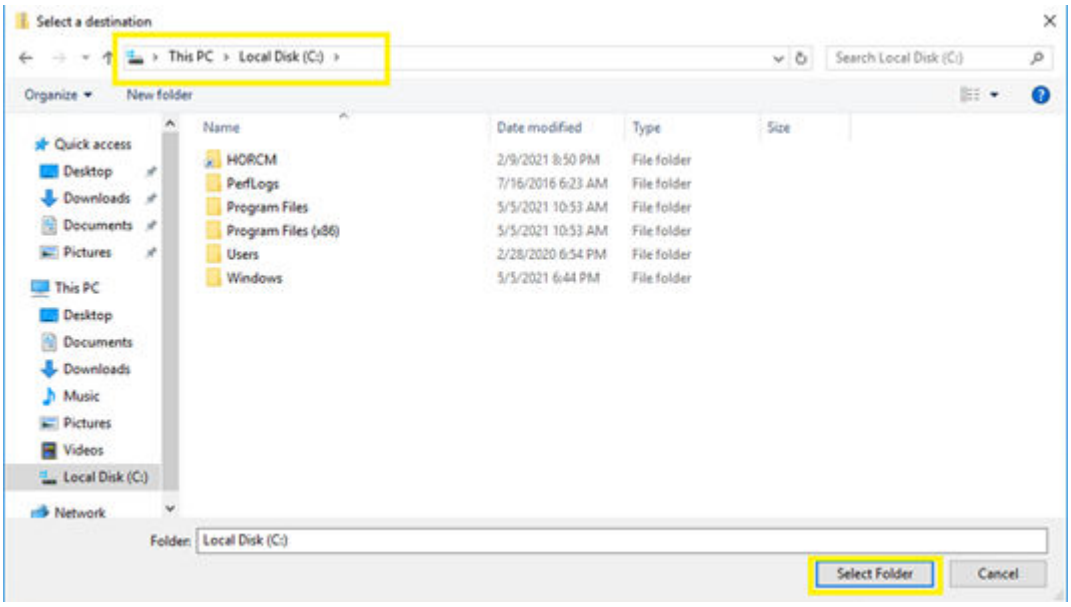

9. Click **Extract**.
   Continue to next section to add binaries to the OS path.

Add binaries to the operating system path

Before you begin

Verify that CLI binaries have been downloaded and extracted.

To add binaries to operating system path, follow these steps:

Note: This procedure applies to the Windows operating system.

Procedure

1. Right click on the windows icon and select **System**.

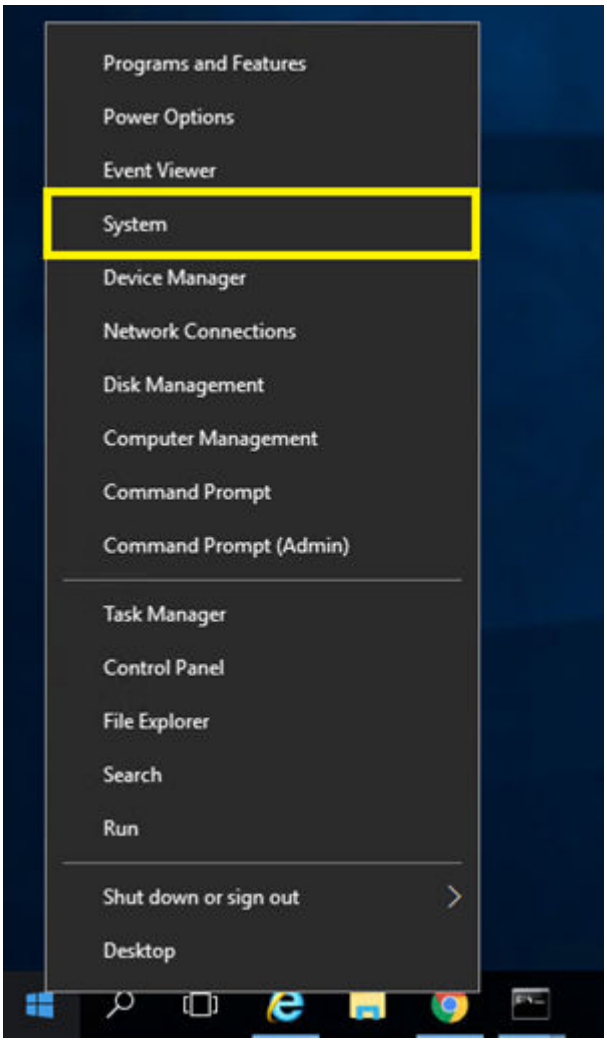

2. From the System menu, click **Advanced system settings**.

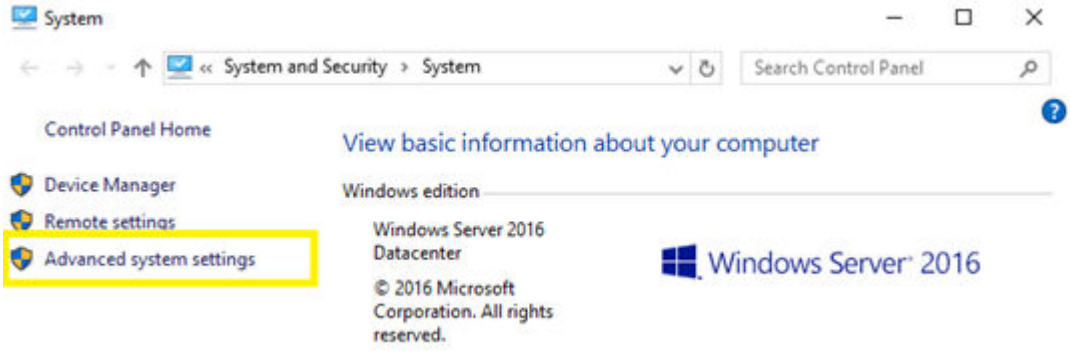

3. On the **System Properties** > **Advanced** tab click **Environment Variables**.

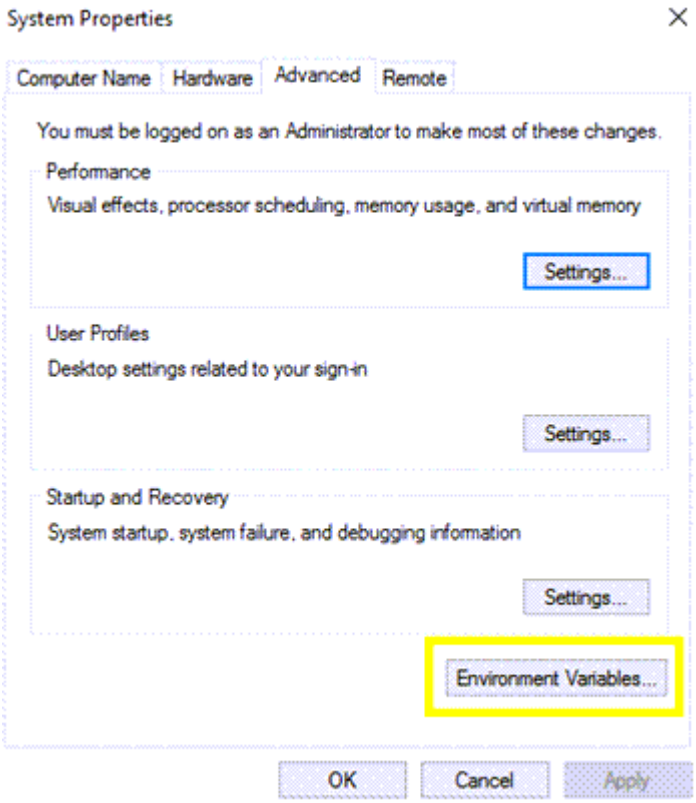

4. In the **System Variables** pane, select **Path**, and then click **Edit**.

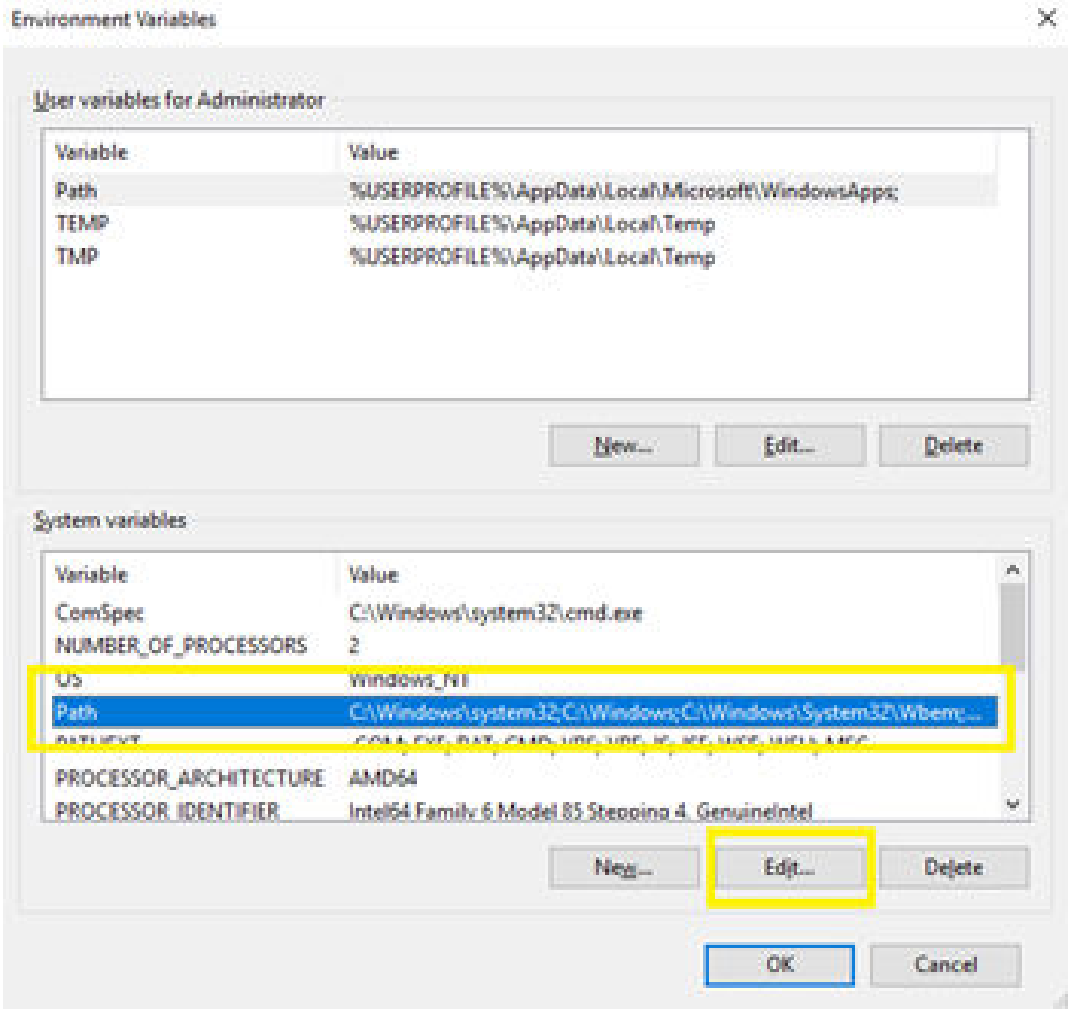

5. In the **Edit environment variables** window, click **New**.
6. Enter the directory to the bin folder that was extracted previously.
7. Click **OK**.

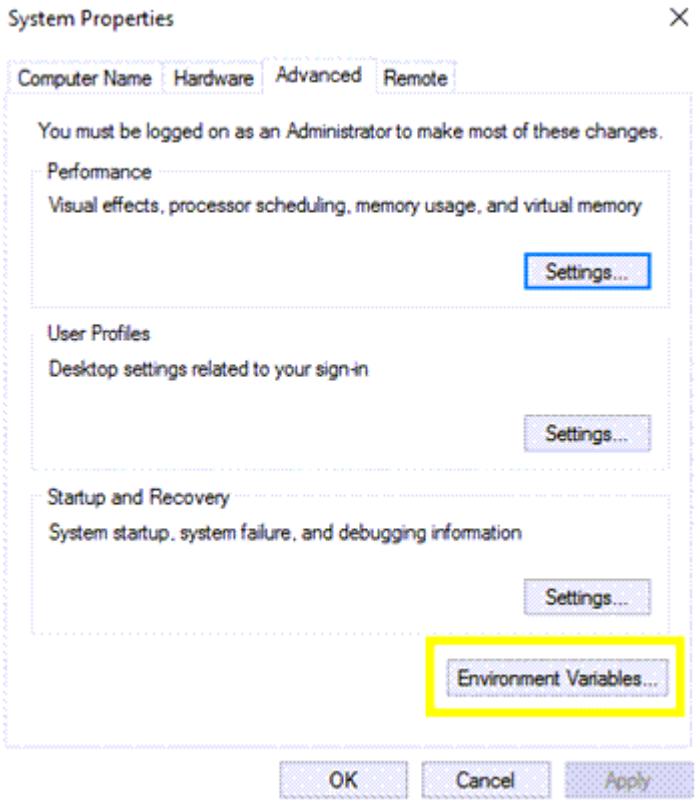

8. Click **OK**.
9. To verify that the path is correctly placed, open Windows PowerShell and enter **kubectl** and **kubectl-vsphere** to confirm command acceptance.

Log in to the Tanzu supervisor cluster from the CLI

Once the binaries have been downloaded and added to the operating system path, Administrators can begin logging in to the supervisor cluster (also known as the control plane node IP address) and begin deployment of a Tanzu Kubernetes Cluster.

To log in to the control plane node, follow these steps:

Procedure

1. Open **Windows PowerShell**.
2. Run the following command:

```
kubectl vsphere login --server=<control plane node IP> --insecure-skip-tls-verify
```

3. Enter the vSphere local user that was enabled with edit permissions in Assigning User Roles and Permissions to Supervisor Cluster (on page 58).
4. Enter the respective **password** for the local vSphere user.
5. If successful, a list of contexts within the cluster will be listed.

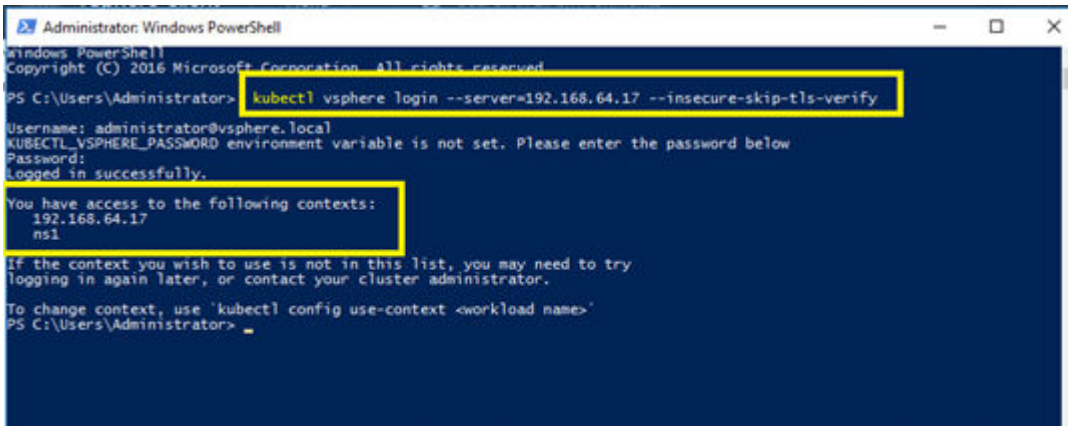

6. Update context to the namespace created in Creating a Namespace (on page 57) using the following command:

```
kubectl config use-context <namespace>
```

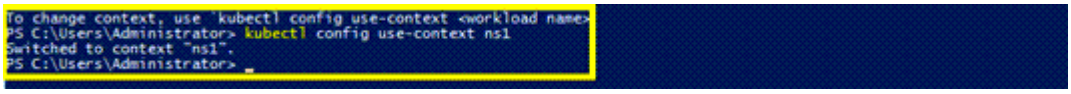

Tanzu Kubernetes Cluster deployment

Before deploying a Tanzu Kubernetes cluster (TKC), verify that all previous steps within this guide have been completed. VMware documentation regarding cluster operational commands must be reviewed by the administrator before continuing with this section.

See the Workflow for Provisioning Tanzu Kubernetes Clusters for example YAML files and operational flows.

TKC deployment with VMFS

To deploy a Tanzu Kubernetes Cluster using Hitachi VSP storage as a VMFS, follow these steps:

Procedure

1. Open **Windows PowerShell**.
2. Log in to the **Supervisor Control Plane IP**. When prompted provide the vSphere username and password enabled in Assigning Users Roles and Permissions to Supervisor Cluster (on page 58).

```
kubectl vsphere login --server=<control plane node IP> --insecure-skip-tls-verify
```

3. Create and save a deployment YAML file that points to the respective storage policy that uses a VSP VMFS datastore defined in Assigning Storage Policies (on page 59).

```
apiVersion: run.tanzu.vmware.com/v1alpha1      #TKGS API endpoint
kind: TanzuKubernetesCluster                   #required parameter
```

```
metadata:
  name: tkgs-cluster-1                          #cluster name, user
defined
  namespace: ns1                                #vsphere namespace
spec:
  distribution:
    version: v1.19.7                            #Resolves to the latest
v1.19 image
  topology:
    controlPlane:
count: 1                                    #number of control plane nodes
      class: best-effort-small                  #vmclass for control
plane nodes
      storageClass: tanzu-vmfs-tier2            #storageclass for
control plane
    workers:
      count: 3                                  #number of worker nodes
      class: best-effort-small                  #vmclass for worker
nodes
      storageClass: tanzu-vmfs-tier2            #storageclass for
worker nodes
```

Note: These examples use a YAML file names tkc.yaml. When creating the YAML file verify that all values including StorageClass definitions are lowercase.

4. After creating and saving the YAML file, use PowerShell to navigate to the location of the YAML file and run the following command:

```
kubectl apply -f tkc.yaml
```

5. Run the following command to view cluster creation status:

```
kubectl get tkc
```

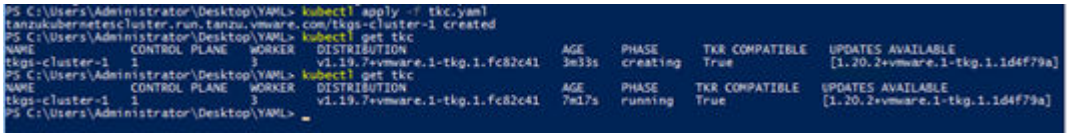

TKC deployment with vVols

To deploy a Tanzu Kubernetes Cluster using Hitachi VSP storage with vVols follow these steps:

Procedure

1. Open **Windows PowerShell**.

2. Log in to the **Supervisor Control Plane IP**. When prompted provide the vSphere username and password enabled in Assigning Users Roles and Permissions to Supervisor Cluster (on page 58).

```
kubectl vsphere login --server=<control plane node IP> --insecure-skip-
tls-verify
```

3. Create and save a deployment YAML file that points the respective storage policy that uses a VSP vVols datastore which was defined in Assigning Storage Policies (on page 59).

```
apiVersion: run.tanzu.vmware.com/v1alpha1      #TKGS API endpoint
kind: TanzuKubernetesCluster                   #required parameter
metadata:
  name: tkgs-cluster-2                         #cluster name, user
defined
  namespace: ns1                               #vsphere namespace
spec:
  distribution:
    version: v1.19.7                           #Resolves to the latest
v1.19 image
  topology:
    controlPlane:
      count: 1                                 #number of control
plane nodes
      class: best-effort-small                 #vmclass for control
plane nodes
      storageClass: tanzu-vvol-tier1           #storageclass for
control plane
    workers:
      count: 3                                 #number of worker nodes
      class: best-effort-small                 #vmclass for worker
nodes
      storageClass: tanzu-vvol-tier1           #storageclass for
worker nodes
```

Note: This example uses a YAML file named tkc2.yaml. When creating the YAML file verify that all values including StorageClass definitions are lowercase.

4. After creating and saving the YAML file, use PowerShell to navigate to the location of the YAML file and run the following command:

```
kubectl apply -f tkc2.yaml
```

5. Run the following command to view cluster creation status:

```
kubectl get tkc
```

Log in to Tanzu Kubernetes Cluster

Once a Tanzu Kubernetes Cluster (TKC) has been created using YAML, administrators and developers can log in to the respective resource using the user names defined within the workload management.

To login to a TKC, follow these steps:

Procedure

1. Open **Windows PowerShell**.
2. Log in to the TKC using the **Supervisor Control Plane IP** along with **cluster name** and **namespace**. When prompted, provide the vSphere username and password enabled in Assigning Users Roles and Permissions to Supervisor Cluster (on page 58).

```
kubectl vsphere login --server=<Supervisor Control Plane IP> --
insecure-skip-tls-verify --tanzu-kubernetes-cluster-name <TKC cluster
name> --tanzu-kubernetes-cluster-namespace <namespace>
```

Stateful applications and persistent volumes

Once all previous steps within this guide have been verified, administrators and developers can deploy stateful applications backed by persistent volumes supplied by the Hitachi VSP from the Hitachi Storage Provider in the form of VMFS vVols.

Verify StorageClasses

To verify storageClass configuration, follow these steps:

Procedure

1. Open **Windows PowerShell**.
2. Log in to the TKC using the **Supervisor Control Plane IP** along with **cluster name** and **namespace**. When prompted provide the vSphere username and password enabled in Assigning Users Roles and Permissions to Supervisor Cluster (on page 58).
3. Run the following command:

```
kubectl get storageclasses
```

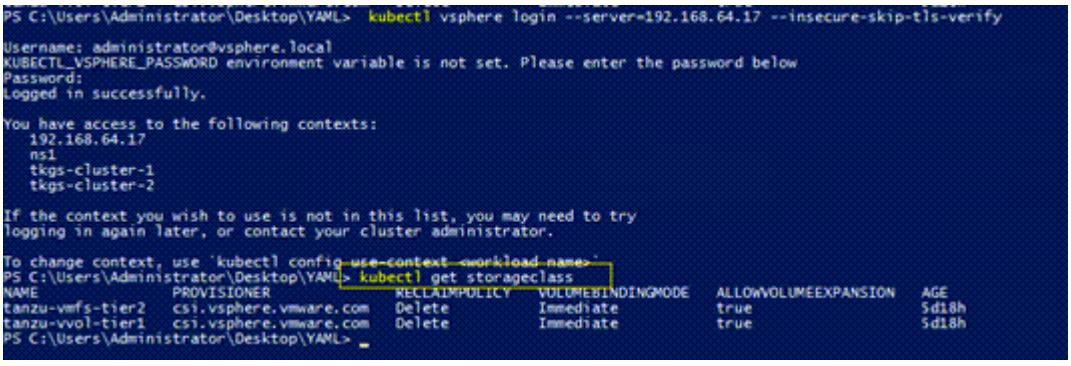

4. The policies defined in Assigning Storage Policies (on page 59) should be listed as available storageClasses.

Deploy a PVC

This procedure applies to both VMFS and vVols.

To deploy a persistent volume claim (PVC) follow these steps:

Procedure

1. Open **Windows PowerShell**.
2. Log in to the **Supervisor Control Plane IP**. When prompted provide the vSphere username and password enabled in Assigning Users Roles and Permissions to Supervisor Cluster (on page 58).
3. Run the following command to change to the namespace defined in workload management:

```
kubectl config use-context <namespace>
```

4. Create a **PVC YAML** file:

```
apiVersion: v1
kind: PersistentVolumeClaim           #kind type
metadata:
  name: my-pvc
spec:
  accessModes:
    - ReadWriteOnce                   #access mode
  storageClassName: <storage policy>  #define vvol or vmfs storage
policy
  resources:
    requests:
        storage: 5Gi                  #capacity
```

Note: A YAML file named pvc.yaml was used in this example.

5. Deploy the PVC by running the following command:

```
kubectl apply -f pvc.yaml
```

6. Verify its creation by running the following command:

```
kubectl get pvc
```

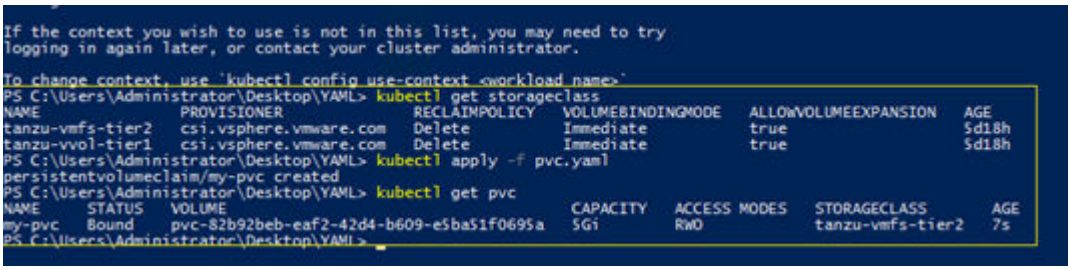

Delete a PVC

To delete a PVC, follow these steps:

Procedure

1. Open **Windows PowerShell**.
2. Log in to the **Supervisor Control Plane IP**. When prompted provide the vSphere username and password enabled in Assign user roles and permissions to supervisor cluster (on page 58).
3. Change into the **context**, which is the TKC you want to use:

```
kubectl config use-context <TKC cluster>
```

4. Within this context, run the following command:

```
kubectl get pvc
```

5. To delete a PVC, run the following command:

```
kubectl delete pvc <PVC name>
```

Deploy a stateful application with PVC

The example in this section applies both to VMFS and vVols storage. Before deploying, confirm that a Pod Security Policy (PSP) has been configured; if this is not addressed, pods will not be deployed. For information on how to set up a PSP see Example Role Bindings for Pod Security Policy.

To deploy a stateful application using a PVC follow these steps:

Procedure

1. Open **Windows PowerShell**.
2. Log in to the **Supervisor Control Plane IP**. When prompted, provide the vSphere username and password enabled in Assigning Users Roles and Permissions to Supervisor Cluster (on page 58).
3. Change into the **context**, which is the TKC you want to use:

```
kubectl config use-context <TKC cluster>
```

4. Create a **deployment YAML** file:

```
apiVersion: v1
kind: PersistentVolumeClaim                  #kind type
metadata:
  name: sqlpvc                               #PVC name
spec:
  accessModes:
    - ReadWriteOnce
  storageClassName: <storage policy>         # define vvol or vmfs
storage policy
  resources:
    requests:
        storage: 10Gi                        #storage capacity
---
apiVersion: v1
kind: Service
metadata:
  name: mysql
spec:
  ports:
  - port: 3306
  selector:
    app: mysql
  clusterIP: None
---
apiVersion: apps/v1
kind: Deployment
metadata:
  name: mysql
spec:
  selector:
    matchLabels:
      app: mysql
  strategy:
    type: Recreate
  template:
    metadata:
      labels:
        app: mysql
    spec:
      containers:
      - image: mysql:5.6
        name: mysql
        env:
          # Use secret in real usage
        - name: MYSQL_ROOT_PASSWORD
          value: password
        ports:
        - containerPort: 3306
          name: mysql
```

```
        volumeMounts:
        - name: mysql-persistent-storage
          mountPath: /var/lib/mysql               #PV mount directory
      volumes:
      - name: mysql-persistent-storage
        persistentVolumeClaim:
          claimName: sqlpvc                       #PVC name
```

Note: The file named *app.yaml* was used in this example.

5. Deploy using the following command:

```
kubectl apply -f app.yaml
```

6. Verify the deployment using the following commands:

```
kubectl get all
kubectl get pvc
```

```
PS C:\Users\Administrator\Desktop\YAML> kubectl get all
NAME                         READY   STATUS    RESTARTS   AGE
pod/mysql-5d5764876c-n5g24   1/1     Running   0          8m15s

NAME                 TYPE        CLUSTER-IP   EXTERNAL-IP   PORT(S)    AGE
service/kubernetes   ClusterIP   10.96.0.1    <none>        443/TCP    4d21h
service/mysql        ClusterIP   None         <none>        3306/TCP   8m15s
service/supervisor   ClusterIP   None         <none>        6443/TCP   4d21h

NAME                    READY   UP-TO-DATE   AVAILABLE   AGE
deployment.apps/mysql   1/1     1            1           8m15s

NAME                               DESIRED   CURRENT   READY   AGE
replicaset.apps/mysql-5d5764876c   1         1         1       8m15s
PS C:\Users\Administrator\Desktop\YAML> kubectl get pvc
NAME     STATUS   VOLUME                                     CAPACITY   ACCESS MODES   STORAGECLASS       AGE
sqlpvc   Bound    pvc-bafdb3a5-0453-4d8f-80b4-f7e4174772d7   10Gi       RWO            tanzu-vvol-tier1   8m17s
PS C:\Users\Administrator\Desktop\YAML>
```

View persistent storage on VMware vCenter

After PVCs are deployed, 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.

To view PV information in vCenter, follow these steps:

Procedure

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

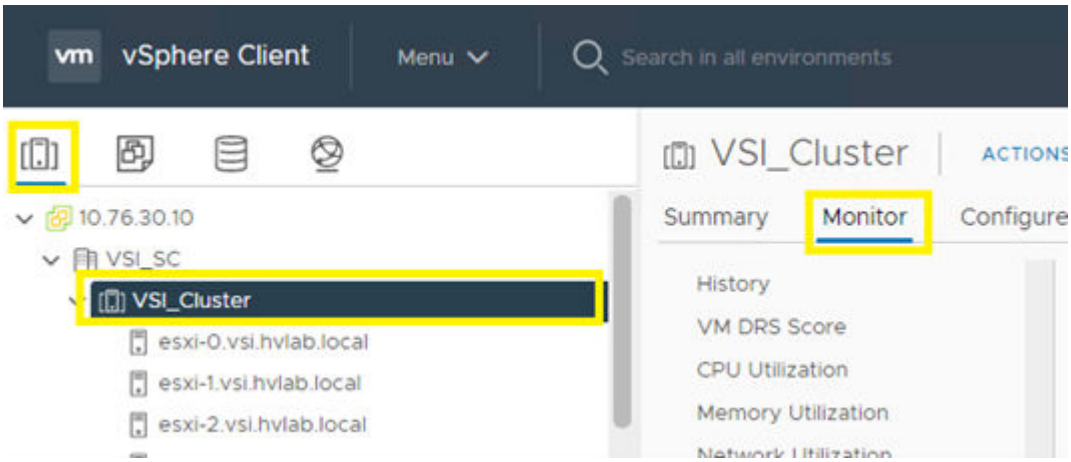

5. From the **Monitor** tab, select **Cloud Native Storage** > **Container Volumes**.
6. The workspace presents the PVCs deployed from workload management, and you can view compliance state, datastore, volume ID, relative worker information, and capacity.

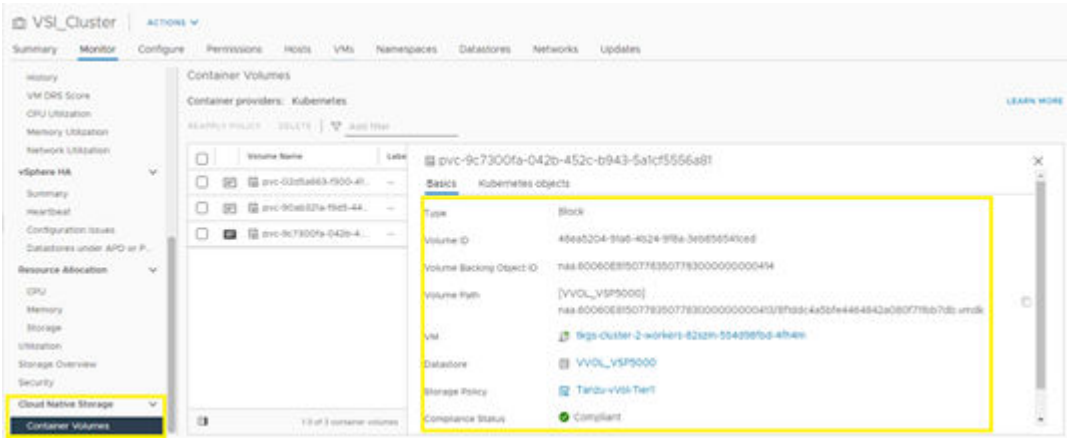

Appendix A: UCP Advisor Storage Administration

This section covers basic storage administration to assist in deploying VMware Tanzu. Installation of UCP A is not covered in this guide and can be found in the Related Documents (on page 82) section.

Log in to UCP Advisor

To log in to Hitachi UCP Advisor follow these steps:

Procedure

1. Log in to the **vSphere web client**.
2. Select **Shortcuts** from the navigation tree.
3. Under **Plugins** select **UCP Advisor**.

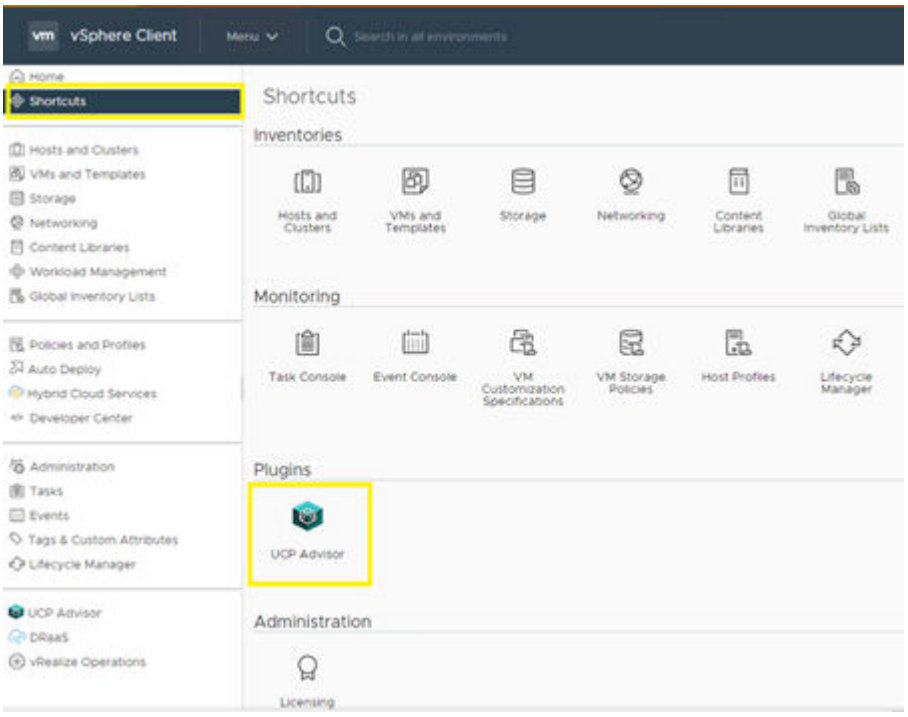

4. Enter the login credentials, and then click **Log In**.

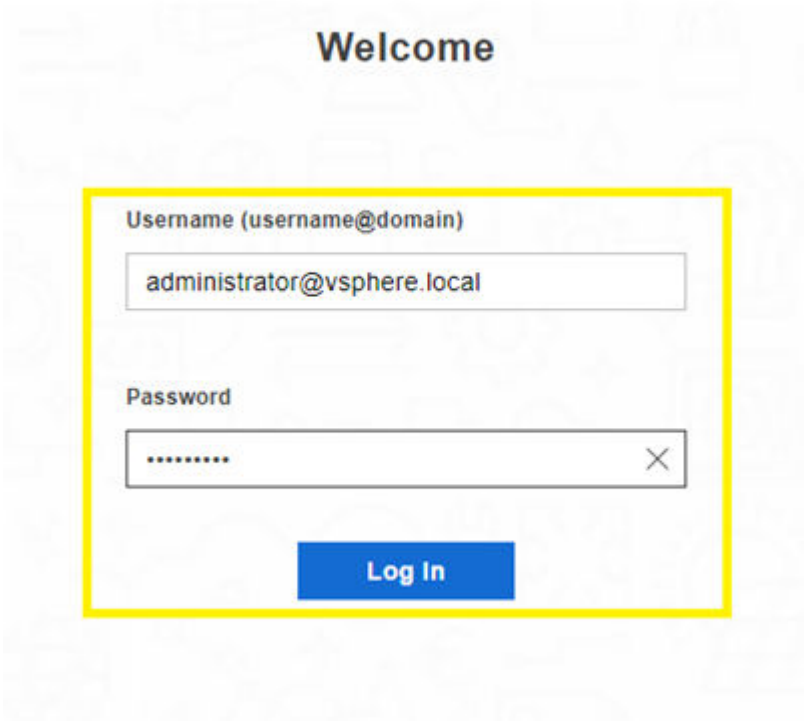

Register a UCP CI system

Before onboarding a storage system with UCP Advisor, a system must be defined. To create a system for a storage-only environment follow these steps:

Procedure

1. Log in to **UCP Advisor**.
2. Select **System** > **Add System**.
   a. Enter a **System Name**.
   b. Select **UCP CI** as the **Model**.
   c. Enter an arbitrary **Serial Number** or leave the default value.
   d. Enter the **Gateway Address** of the Gateway VM deployed during UCP Advisor installation.
   e. Click **Submit**.

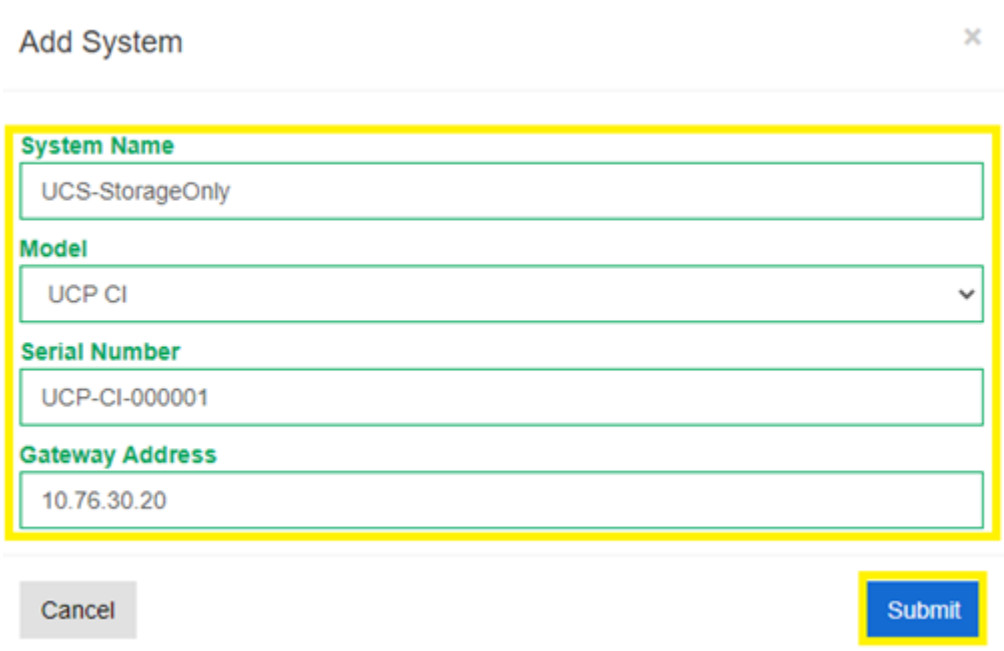

Onboard a VSP to UCP Advisor

To onboard a VSP to UCP Advisor, follow these steps:

Note: Verify that the UCP CI system has been created.

Note: Verify that User Authentication settings are enabled on the respective Virtual Storage Platform (VSP) command device.

Note: Before onboarding the VSP storage system, verify command device allocation to the UCP Advisor gateway.

Procedure

1. Log in to **UCP Advisor**.
2. Under the onboarded UCP CI system, click **Add Storage**.
   a. Enter the VSP **serial number**.

b. Enter the VSP **SVP IP address**.
c. Enter the VSP **username and password**.
d. Optionally, if you are using a VSP that does not have an SVP, enter the **CTL1** and **CTL2 IP** addresses.
e. Click **Submit**.

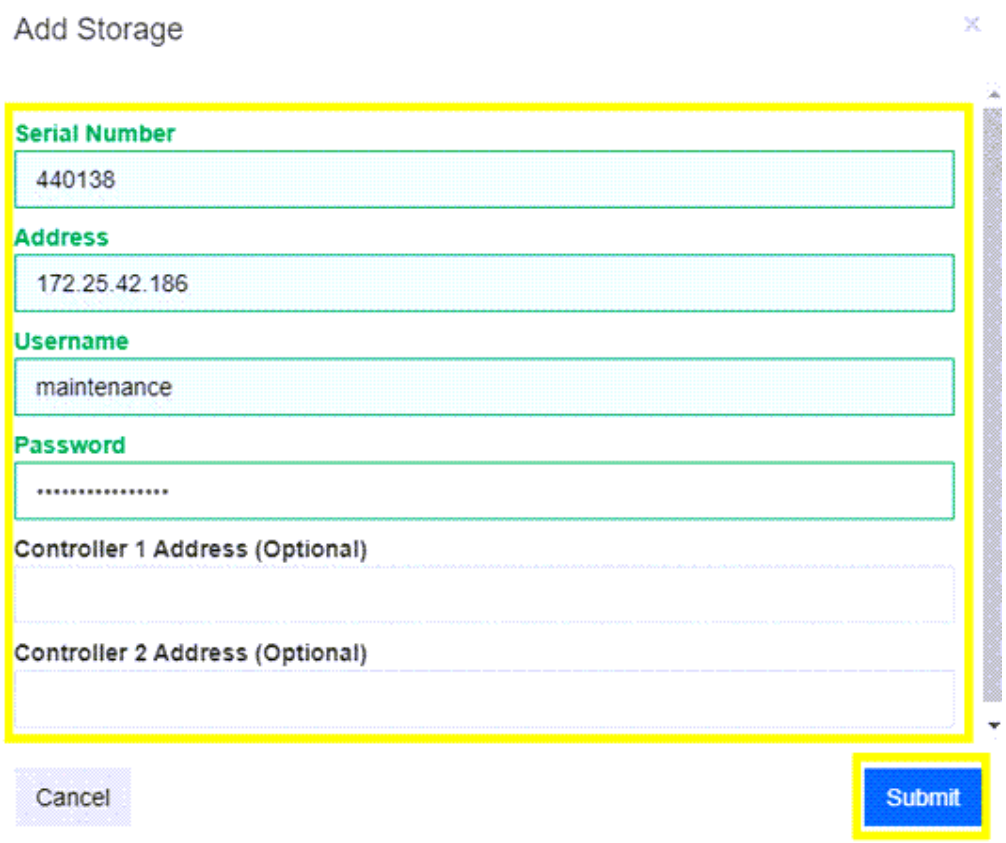

Create a Hitachi Dynamic Provisioning (HDP) Pool

Before you begin

Confirm parity group configuration.

To create an HDP pool from UCP Advisor, follow these steps:

Procedure

1. Log in to **UCP Advisor**.
2. Select **Storage** under the applicable UCP CI system.
3. Select a VSP system.

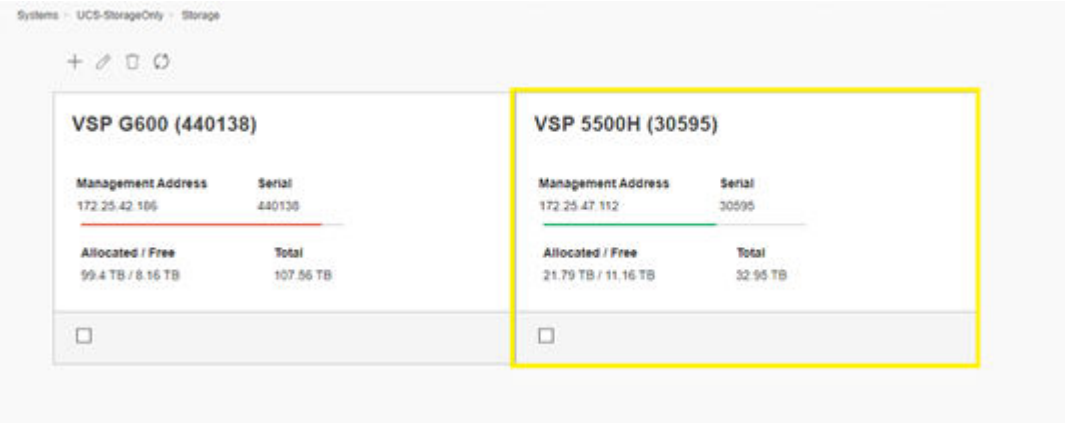

4. From the **Select the storage detail** list select **Storage Pools**.
5. Click **Create Storage Pool**.

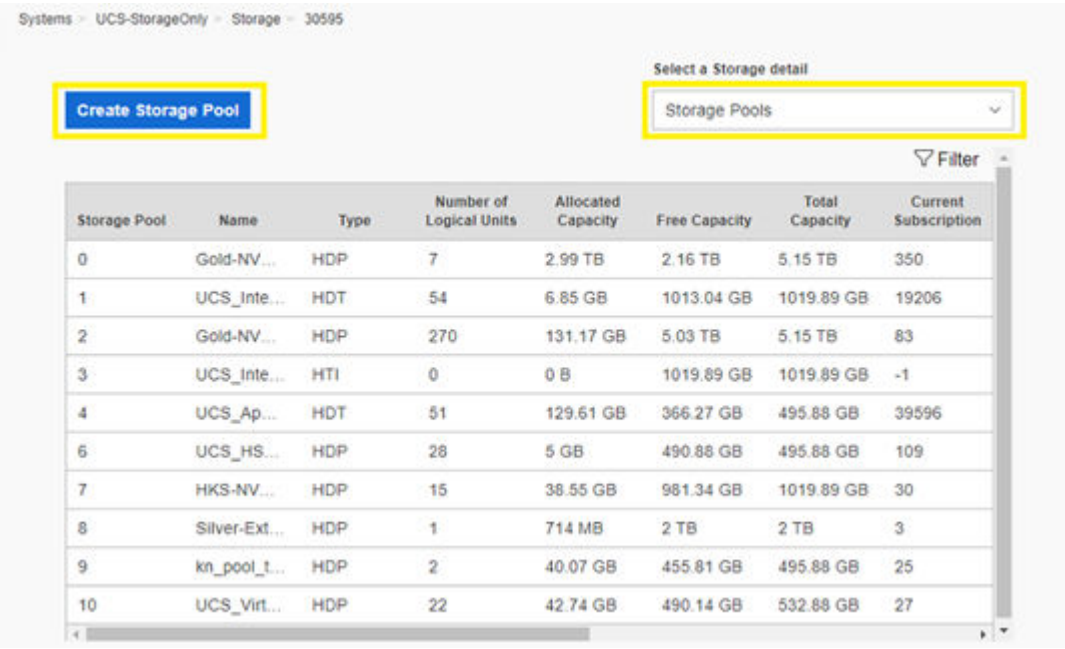

Storage Pool	Name	Type	Number of Logical Units	Allocated Capacity	Free Capacity	Total Capacity	Current Subscription
0	Gold-NV...	HDP	7	2.99 TB	2.16 TB	5.15 TB	350
1	UCS_Inte...	HDT	54	6.85 GB	1013.04 GB	1019.89 GB	19206
2	Gold-NV...	HDP	270	131.17 GB	5.03 TB	5.15 TB	83
3	UCS_Inte...	HTI	0	0 B	1019.89 GB	1019.89 GB	-1
4	UCS_Ap...	HDT	51	129.61 GB	366.27 GB	495.88 GB	39596
6	UCS_HS...	HDP	28	5 GB	490.88 GB	495.88 GB	109
7	HKS-NV...	HDP	15	38.55 GB	981.34 GB	1019.89 GB	30
8	Silver-Ext...	HDP	1	714 MB	2 TB	2 TB	3
9	kn_pool_t...	HDP	2	40.07 GB	455.81 GB	495.88 GB	25
10	UCS_Virt...	HDP	22	42.74 GB	490.14 GB	532.88 GB	27

6. Enter a **Pool Name**.
7. Select **Pool Type as Dynamic Pool (HDP)**.
8. Select an available **Parity Group** and define a **Size (GB)**.
9. Click **Add Volumes**.
10. Enter a **Warning** and **Depletion Threshold**.
11. Click **Submit**.

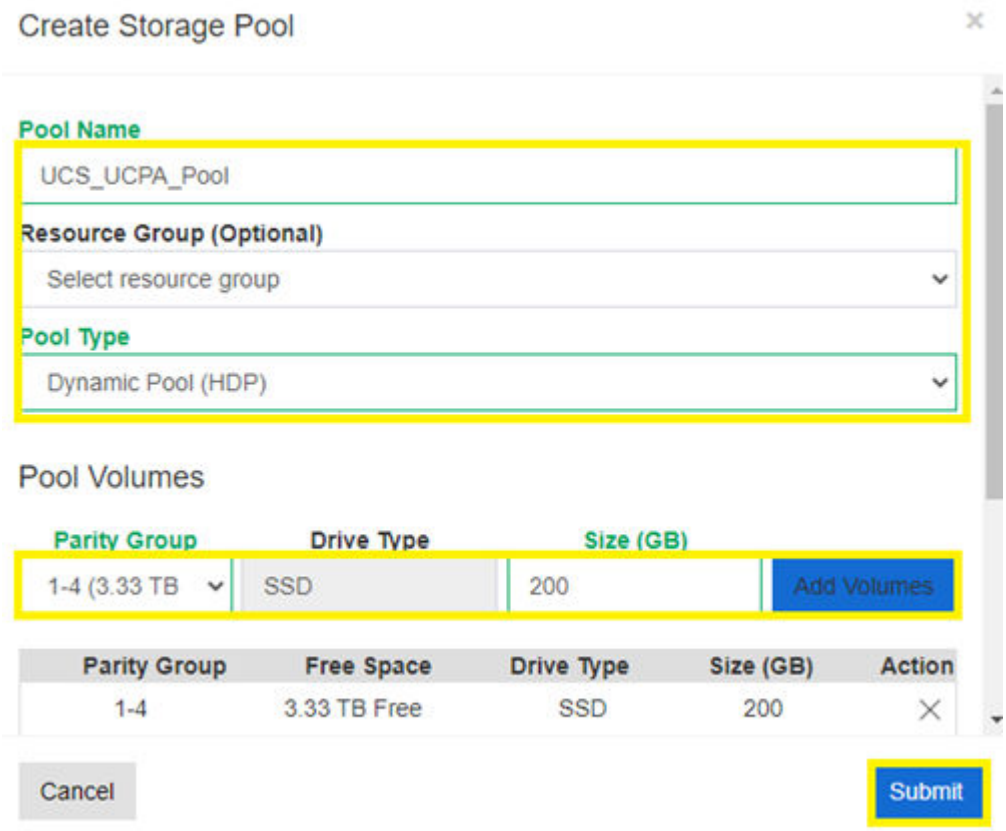

Create a VMFS datastore

To create a VMFS datastore from UCP Advisor, follow these steps:

Procedure

1. Log in to **UCP Advisor**.
2. Select **Storage** under the applicable UCP CI system.
3. Select a **VSP system**.

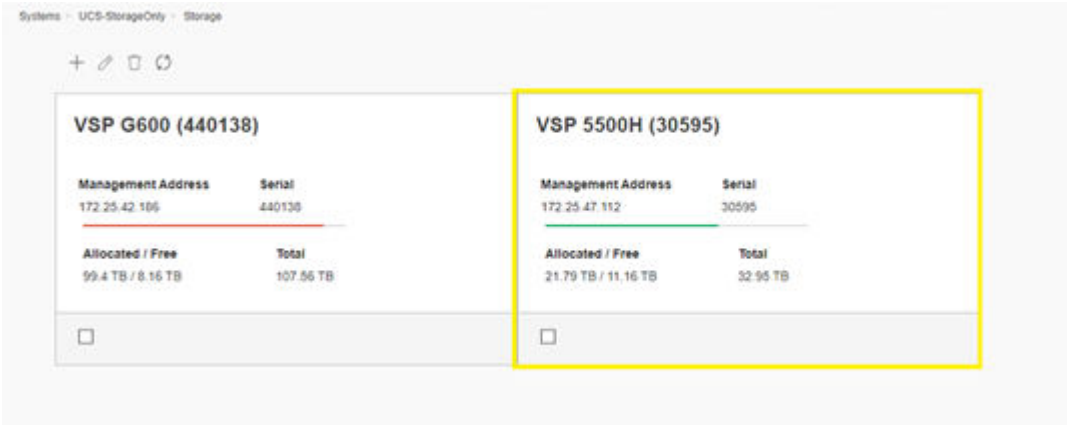

4. From the **Select a Storage detail** list select **Datastores**.
5. Click **Configure Datastore**.

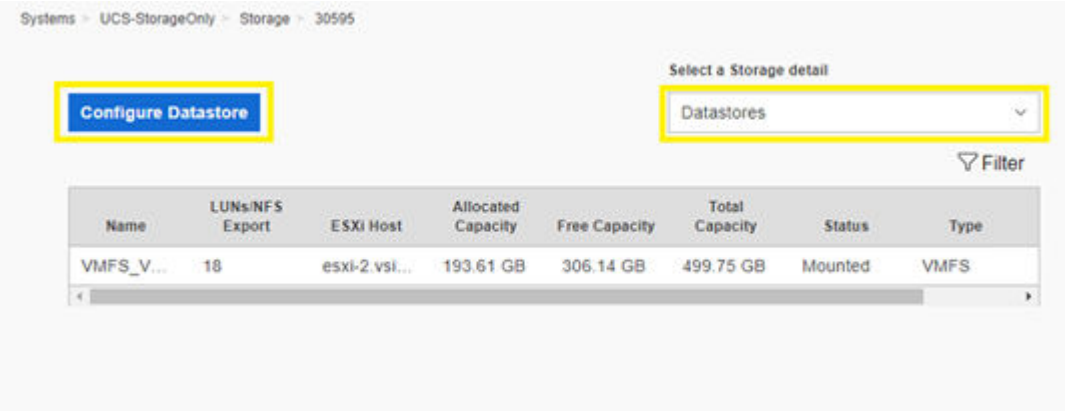

6. Complete the following:
   a. Define a **Datastore Name**.
   b. Enter a **Capacity**.
   c. Select **Single** or **Multiple Datastore Creation**.

   d. Select the applicable **VSP storage system Serial Number**.
   e. Optionally, select a **Resource Group**.
   f. Select an available **HDP Pool**.
   g. Optionally, select a **Capacity Saving Mode**.

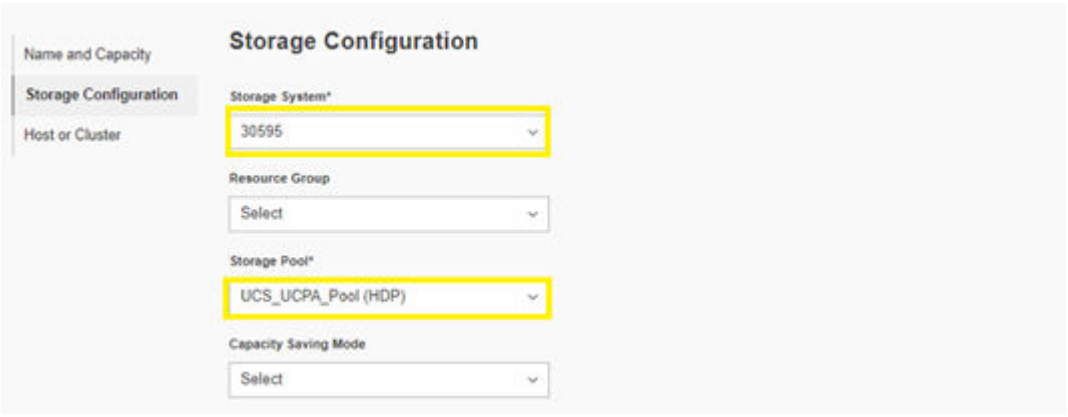

   h. Select the **datacenter cluster**.
   i. Click **Advanced Options**.

j. Under each ESXi host, select the host groups that will provide the VMFS datastores to both UCS SAN fabrics.

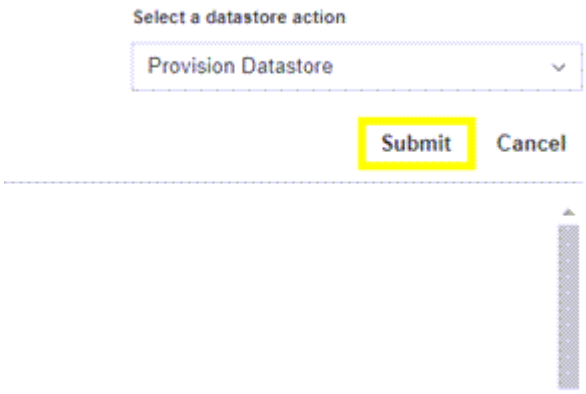

k. Click **submit**.

Related Documents

This document references the following documentation:

- Hitachi Virtual Storage Platform Documentation
- Hitachi Storage Provider for VMware vCenter Documentation
- Hitachi UCP Advisor Documentation
- VMware Tanzu Documentation
- VMware Tanzu Kubernetes Grid Documentation
- vSphere with Tanzu Configuration and Management Documentation
- Cisco and Hitachi Adaptive Solutions for Converged Infrastructure Design Guide
- Cisco and Hitachi Adaptive Solutions for Converged Infrastructure Deployment Guide

Solution References

For more information on Hitachi solutions and products, go to HitachiVantara.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 in to Hitachi Vantara Support Connect. For contact information see Customer Contact Us.
- To open a new support case, see the 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