

Cisco and Hitachi Adaptive Solutions with VMware Tanzu Basic

© 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 <u>https://</u>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

	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 <u>Related Documents (on page 82)</u>.

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.

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 1 List of support tasks for Hitachi VSP

Table 2 List of supported workflows for Hitachi VSP

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

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.

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

Table 4 Software Versions Used for Validation

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.

	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)				

Table 5 Hitachi integration points with VMware Tanzu

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 <u>Related Documents (on page 82)</u> 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 <u>Related Documents (on page 82)</u>.



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 <u>Appendix A: UCP Advisor Storage Administration (on page 75)</u> 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 <u>Hitachi Vantara Knowledge portal</u>.

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, <u>Appendix A: UCP Advisor Storage Administration (on page 75)</u> 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.

Hitachi Device Manager Storage I	Navigator ttings View Tool H	lelp			
Explorer	Logical Devices				
Storage Systems • 👘 R900_NVMe ASE-47.112 G10(S/N:30595)	Volume Migration	• G10(5/N: 30595) > Logical Devices		
帽 Tasks 隆 Reports *聞 Components *院 Parity Groups	Number of LDEVs	01 01 01	en Allocated en Unallocated en Reserved VOLs	0 2 8 192	
Image: Contract of Cont	Format/Shredding Tas	k Status Id LUN Paths Select All P LDEV Name	Edit LDEVs Mor ages Column Séttin Status 1 V	e Actions 🛛 🔻	Numl Paths
	00:00:01	pool-vol	Vormal	2969.60	

- 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.

A1	a Aables Free	Spaces		_	_		_			
	Party Broug	Free Space	RAID Lavel	Casechy	Base Emulation	Drive Type/Interface/80H	Deception	Accelerated	Contract of Charlenger	/1 10
2	14	0000	4040+285	1037.09.08	CONTRACT.	supported.	Developed	- III		
1	1-4	0005	4(40+20)	4454.77.08	cats-v	BOW/Wall	Disabled	and the second		
3	1-0	0000	6(60+29)	10363.77	OPEN-V	\$\$D/W/We ¹	Doubled			

- 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.

rovisioning Type:	Basic	
ystem Type:	Open Mainframe	
mulation Type:	OPEN-V	
arity Group Selection:		
Drive Type/Interface/i	PM: SSD/NVMe/-	
RAID Level:	6(6D+2P)	
	S	elect Free Spaces
Total Selected Fre	e Spaces: 1	
Total Selected Fre	e Space Capacity: 1.01 TB	
DEV Capacity:	500	GB ↓
	(0.05-1037.99)	
umber of LDEVs er Free Space:	1 (1-2)	
5 58	(1-1)	
DEV Name:	Prefix Initial Nur	mber Add
	UCS_PoolVOL_VMFS	alt aumhar
	or blank)	git number,
ormat Type:	Quick Format	
S Options		
A		

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.

Available Pool V	olames			-		Selected P	ool Volumes				
after on Cor	Select All Papen	Options w 10	4 1 /	1 (0.10)		Delect All Page	÷		110.0		Options w
LDEV ID	LDEV Name	Parity Group 10	RAID Level	Drive Type/I		LOEV ID	LOEV Name	Capacity	Parity Ontwo ID	RADD Lavel	Drive Type/#
ea.co.so	UCS_NoNOL_VMPS	14	6(60+39)	SSD./WDV-							
00-00-11	SICE_Peel/OL_VVOL	1.4	4(40+29)	\$50,7HO/-							
00.00.29	лан флука	24	6(40+29)	\$\$0,#HD/1							
					ALC .			No D	a ha		
					4 famous			NO D	ata		
-								-			
			Selected	1 03		1416				Selecte	nt 0 of 0

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

Pool Type:	- F	Dynamic F	Provisioning		Se	lected P
System Typ		(a) (Doen	Mainframe		Sel	ect All Pag
Multi-Tier Po	pol:	Enable	Disable			Pool Nan (ID)
Data Direct	Mappingi	Active F	Flash			
Pool Volume	e Selection:	🕞 Auto	 Manual 			
Γ	Drive Type	/RPM:	Mixable	•		
	RAID Level	la la	Mixable			
	Enable	Accelerated C	Compression for FMD parity group			
		UCS App	_Pool			
Pool Name:						
Pool Name:		(Max. 32 C	Characters)			
Pool Name:		(Max. 32 C	Characters)			
Pool Name:	ŧi.	(Max. 32 C	Characters)			
Pool Name:		(Max. 32 C	Characters)			
Pool Name:	<u>8</u>	(Max. 32 C	Characters)			
Vool Name:		(Max. 32 C	Characters)			

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.

Explorer	UCS_App_Pool(4)							
Storage Systems	8300 MVM ASE-47.112 G10(5/N-30395) > Pr	eds > UCS_App_Pool(4)						
* 1 R300_NVM+ ASE-47.112 G10(5/N-30595)	Status	Normal						
😭 Taska	Pool Name (ID)	UCS_App_Po	pl(4)					
🎇 Reports	Pool VOL with System Area (Name)	00-00-10(UC	S_POOLVOL_VHFS)					
' 😭 Components	Pool Type	0T						
1 👯 Parity Groups	RAID Level	6(6D+2P)	6(60+29) SSD.PHD/~ Disabled					
Carlogical Devices	Drive Type/RPM	\$\$D.PHD/-						
Contraction of the second	Encryption	Disabled						
Gold-NVMer RAID5-pool(2)	Cache Mode	• /						
Gold-NVHe-RAIDG-Pool(0)	Protect V-VOLs when I/O fails to Blocked Pool	Protect V-VOLs when I/O fails to Blocked Pool VOL						
UCS_App_Pool(4)	Protect V-VOLa when 1/O fails to Full Pool	Protect V-VOLa when 3/0 fails to Full Pool						
UCS_Intersight_App_Pool(1)	Number of Pool VOLs	Number of Pool VOLa						
UCS_Intersight_T1_Pool(3)	Number of V-VOLs	Number of V-VOLs						
💧 UCS_VVOL_Peel(S)	Number of Root VOLs	Number of Root VOLA						
* Monts/Heat Groups/GCS1 Targets	Pool Capacity (Used/Total)	Pool Capacity (Used/Total)						
15 External Storage	Total Efficiency	Total Efficiency						
1 Reductor	Saving Effect							
	V-VOL Capacity (Used/Total)							
	Subscription (Current/Limit)							
	Jines Palined Theshold (Wassiss Theshold)							
	Pool Volumes Virtual Volumes							
	Create LDEVs Add LUN Paths Espan	d V-VOLA More Actions *						
	RFitter Oh CLI Select All Pages C	olumn Settings						
Analytics		finister	Capacity					
Administration	LOEV TO LOEV TIMENA	Status Type	7144					

- 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.

s wizard lets you create an k Finish to confirm the cre	nd provision LDEVs enter nation, or click Next if you	the information for a want to add LUN p	LDEVs you want t aths for the LDEVs	o create, and I.	then click Add.
Provisioning Type:	Dynamic Provisio	ining			â
System Type:	💿 Open 🛛	Mainframe			
Data Direct Mapping:	💮 Enable 💿	Disable			
Emulation Type:	OPEN-V				
Capacity Saving:	Disabled			•	
Multi-Tier Pool:	💿 Enable 🌀	Disable			
	Active Flash				
TSE Attribute:	💮 Enable 🛞	Disable			
Pool Selection:					
Drive Type/RPM:	SSD,FMD/-			 •	
RAID Level:	6(6D+2P)				
			Select #	Pool	Add
Selected Pool Na	ame(ID):	UCS_App_F	Pool(4)		
Selected Pool Ca	epacity:	495.87 GB			
LDEV Capacity:	500		GB		
	(0.05-262144.00	1)	(
Number of LDEVs:	1			5	
	(1-63039)				
LDEV Name:	Prefix	Initi	al Number		
	VSI-OCP-DS-VSP	5000			U
	(Max. 32 character	rs total including ma	x. 9-digit number		

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 <u>Cisco and Hitachi Adaptive Solutions for Converged Infrastructure</u>.

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.

UCS_App_Pool(4)				
R900 NVMe ASE-47.1	12 G10(S/N:30595) >	Pools > UCS_App_Poo	l(4)		
Status			Normal		
Pool Name (ID)			UCS_App_Pool(4	4)	
Pool VOL with System	m Area (Name)		00:00:10(UCS_	POOLVOL_VMFS)	
Pool Type			DT		
RAID Level			6(6D+2P)		
Drive Type/RPM			SSD,FMD/-		
Encryption			Disabled		
Cache Mode					
Protect V-VOLs when	n I/O fails to Blocked Po	ol VOL			
Protect V-VOLs when	n I/O fails to Full Pool				
Number of Pool VOL	5				
Number of V-VOLs					
Number of Root VOL	.5				
Pool Capacity (Used	/Tota!)				
Total Efficiency					
Saving Effect					
V-VOL Capacity (Use	ed/Total)				
Subscription (Currer	nt/Limit)				
User-Defined Thresh	old (Warning/Depletion	1			
Pool Volumes	Virtual Volumes				
Create LDEVs	Add LUN Paths Expa	ind V-VOLs More A	ctions 💌		
AFilter ON OF	Select All Pages	Column Settings			
-	1000 Name	Chabus	Emulation	Capacity	
	LDEV Name	Status	Туре	Total	Reserved
V 00:00:12	VSI-OCP-DS-VSP5	000 🥥 Normal	OPEN-V CVS	500.00 G8	0.00 GB
		the second s	A CONTRACTOR OF A CONTRACTOR O	and the second se	The second se

2. Click Next.

	S	elected LDEV	/s				
	Sel	ect All Pages					Options 🛩
		LDEV 1D	LDEV Name	Parity Group ID	Pool Name (ID)	Emulation Type	Capacit
		00:00:12	VSI-OCP	•	UCS_App_Po	OPEN-V CVS	500.0
						-	
Add 🕨							
Pamaua	_						
Neilly Ye							
			-			-	
							_

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

iele	Attribute:	2.Sek	ect Hor	st Groups	s / e	SCSI T	argets > 3.View/	Tha	nge LUN Paths > Value:	4.Confirm	_				_
1	Host Group	Name	6		•	cont	ains 🔤	0	vs:					•	
2	Select Item	-		1	•	Sele	ct Item 🛛 🖛	J							
3	Select Item	į		I	•	Sele	ct Item 🛛 🖣	J							
м	atch All	• • •	onditio	ns above							Reset	Clear		Appl	Y
*	Filter ON	HF]	Sele	ict All Pa	ges	2	Options 👻	16	E 1 / 1	(+ (+)		Se	lect A	Il Pages	
	Port ID	Prio prov	rity on visionir	der for Ng			Host Group Name		Host Mode	Port ^			Port	t ID	н
	CL1-A					1	1A-G00 (00)		00 (Standard)	Tarpe					
	CL1-A					1	HDIDProvisioned		00 [Standard]	Targe					
	CL1-A					1	ds120-4590-hba		21 [VMware	Targe					
	CL2-A					1	3A-G00 (00)		00 [Standard]	Targe					
	CL3-A					1	ds120-4591-hba		21 [VMware	Targe					
	CLS-A					1	5A-G00 (00)		00 (Standard)	Targe					
	CLS-A					1	VASAProvisioned		00 [Standard]	Targe		1			
	CL5-A					1	VASAProvisioned		00 [Standard]	Targe	Add 🕨				
	CL5-A					1	spc-20000025b5		00 [Standard]	Targe					
	SLZ:A					1	7A-G00 (00)		00 [Standard]	Targe	4 Remove				
	CL1-C					5	1C-G00 (00)		00 [Standard]	Targe		9			
	CL3-C					5	3C-G00 (00)		00 [Standard]	Targe					
	CL5-C					5	5C-G00 (00)		00 [Standard]	Targe					
	SL7-C					5	7C-G00 (00)		00 [Standard]	Targe		-			
	CL7-C					5	spc-10000090fab		00 [Standard]	Targe					
	CL7-C					5	ucp2k-c2-b2-hba	2	21 [VMware	Targe					
	CL7-C					5	ucp2k-c4-b3_7C		21 [VMware	Targe					
	CL7-C					5	ucp2k-c4-b4_7C		21 [VMware	Targe					
	CL1-B					3	1B-GOO (00)		00 [Standard]	Targe					
	CL1-8					3	ds120-4592-hba		21 [VMware	Targe					
	CL3-8					3	3B-G00 (00)		00 [Standard]	Targe 🗸					
< 0					_					>		< 6	_		

- 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.

dd LUN Paths	8				The second s					TO
lalect LDKVs >	2.Select Host (Groups / HCRE To	rants > 3.Veni(Dang	n LUM Public 🗧 🔹 4.Ca	illerin (
The LUN IDs are a and select LDEVs	utometically a you want to sh	et, but you can d ange and then d	hange a LUN by clicking C lick Change LUN IDs. Click	Nanga LUN IDa. You mu Finish to confirm the L	at first select the check UN paths.	bes for the heat proup (On the table subheading) you want to change.		
UUfree .										
Added LUNS	Nonessee								and the second second second	
Afiber ON	C.L. Balac	a Al Pages						Options w	(H (H) 1 / 1	-+ -+
			LUN ID(14 Sets of I	Pathe)						
Previolating Type	Attribute	710 PI	S AVEL 0700-	S AV81_0795-	2 AVGL0700-	20.5- 20 AV82,0700-	S 8782,0700-	S 8/18_0.00-	X 8-VII_0700-	8.00
64	1.00	Deathed	28	20	39	20	29		211	29
			-							
Change LDEV	Settings 0	hange LUIV 20s					100		Selected: 7	a of a
That a second	and the second	and the second statement of						_		
								Concession in the	Finish Co	?

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

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 <u>Appendix A: UCP Advisor Storage Administration (on page 75)</u>. 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.

Explorer	Res	ource Groups			
Storage Systems	Reso	arce Groups			
Analytics	Nu	mber of Resource Grou	ps		
Administration	Re	ource Groups			
* 🕅 Resource Groups		Create Resource Groups	Edit Resource Gr	oup Delete Res	ource Groups
neta_resource (0)		Filter ON OFF	Select All Pages Co	lumn Settings	
* Cache Partitions		Resource Group Name	Number of User Groups	Number of Parity Groups	Number of LDEVs
1 Encryption Keys		& meta resource	9	5	64968
饕 License Keys		A vitels-RG	7	0	312
	<				

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

Create Resource Groups		_		
1.Create Resource Groups >	2.Confirm			
This wizard lets you create res	ource groups. Click ea	ch button to sel	ect parity groups, LDEVs, ports, h	tost groups and iSCSI target
Resource Group Name:	UCS_VVOL (Max, 32 Character	rs)		Se
Parity Group Selection:				
Total Selected Par	ity Groups:	0	Select Parity Groups	
LDEV Selection:		159		
Total Selected LD	EVs:	0	Select LDEVs	
Port Selection:				
Total Selected Por	ts:	٥	Select Ports	Add 🕨
Host Group Selection:				
Total Selected Ho	st Groups:	0	Select Host Groups	
ISCSI Target Selection:				
Total Selected iSC	SI Targets:	0	Select ISCSI Targets	
·				
				<

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**.

wailable LDEV						102	elected LDD	Vs			
arter ON CON	Select All Pages	Options w 00 0	3 7.45	+ (#)		124	wit All Pages			L. Option	
LDEV 10	LDEV Name	Party Group 10	Paul Name(10)	°.,		×	LORV ID	LDEV Name	Parity Group 10	Post Name(ID)	4
00.04-05	+:					2	00-00-11	UCS_ReeNOL_WOL	1.4	VOL.WOLA.	
00.04-0#	×					2	09-99-93	1		1	
00-04-55	+	. 6	+			2	00.03.04	10 C			
00-04-11	-		+			2	00-03-73				
00-04-12	+	4	+			2	00-03-76		101		
00-04-13	14 C		+			2	00.03.07		-	14	
00:04:54						2	00.03.78	1.0		14	
00-04-15						2	00-03-FS	- 25	1.00		
90-04-16						2	00-02-04		1	1	
00-04-57	(e)					2	00.03.01	1.00		-	
00:04:58	(*)		i*		100000		00-03-#C				
00-04-19	÷.		+		A01.9	2	00.03.90	- 10 I			
00-04-1A	+		+		1	2	00-03-41			5	
00.04:58		4			4 famova	2	00.03.49		1	1	
00:04:50	+		*			2	00-04-00	- X			
00:04:10						2	00-04-01	1 82		1	
00-04-SE			+			2	00-04-02			-	
00-04-1#			+			2	00-84-03			1	
00.04-20	+		+			2	00-04-04	100	-	+	
00-04-21			+			2	00-04-05				
00-04-22	+		+			2	00-04-06		-		
00-04-23	÷					2	00.04.07		21	10	
00-04-34	*	4	+			2	00.04.08		100	1	Ē
00.04-25	(R)	1.0				2	00.04.09		100		
60:04:28						2	00.04.04		120	12	
00-04-27	(F)		+	0		2	00-04-08	1.0	200	÷	ſ
				- "		1	All and a second second			Sector Se	1
		37	television of all a							March 12 1	ń

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

Resource Group Name: UCS WIDL			Selec	ted Resource	e Groups	_		
(Max. 32 Ch	eractien()		Select A	U Pages				111
farity Group Selection:			E Ne	ne (ID)	Number of Renty Groups	Number of LDEVs	Number of Ports	Number of Heat Groups
Total Salvated Parity Groups	Select Parity Droubs							
DEV Beleition:								
Total Balacted LORVe-	Select LDEVN							
ort Selection:								
Total Selected Ports	Belett Porte	Add B	1		N	Dat		
rest Group Salaction:								
Total Selected Host Groups:	Balact must Groups							
ICEI Tarpet Belection								
Total Salacted (ICS) Targete:	Select GCEI Terpeta							
			< Longer					

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 <u>Related Documents (on page 82)</u>.

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.

	Welcome to VASA
Manage Storage Systems	Service Name : Hitachi Storage Prov
	Service Port Number : 50001
e Storage Containers	Database Port Number: 50003
and the Oak stars	Service Version : 03.5.8
Capability Schema	VASA Version : 1.0/2.0/3.0
Change Credentials	Supported Model : VSP 5500H
change creatinais	VSP 5100H
Update Software	VSP 5500
	VSP 5100
	VSP 5100 VSP 5x00
	VSP 5100 VSP Fx00 VSP Gx00
nance	VSP 5100 VSP Fx00 VSP Gx00 VSP F1500
nance	VSP 5100 VSP Fx00 VSP Gx00 VSP F1500 VSP G1500
enance	VSP 5100 VSP Fx00 VSP Gx00 VSP F1500 VSP G1500 VSP G1000
nance Single Sign-on Setting	VSP 5100 VSP Fx00 VSP Gx00 VSP F1500 VSP G1500 VSP G1000 HUS VM
Single Sign-on Setting	VSP 5100 VSP Fx00 VSP Gx00 VSP F1500 VSP G1500 VSP G1000 HUS VM VSP
ance Single Sign-on Setting Download Logs Update Certificate	VSP 5100 VSP Fx00 VSP Gx00 VSP F1500 VSP G1500 VSP G1000 HUS VM VSP
enance Single Sign-on Setting Download Logs Update Certificate Restart Service	VSP 5100 VSP Fx00 VSP Gx00 VSP F1500 VSP G1500 VSP G1000 HUS VM VSP

- 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.

torage System	
Storage System Type :	VSP \$100.5100.6100H-5509 +
SW # address :	17226-47110
User ID :	VAM
Password	

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.

Ma	nagement	
C	Manage Storage Systems	
C	Manage Storage Containers	
C	Capability Schema	\mathbf{D}
C	Change Credentials	\supset
C	Undata Softwara	
Mai	intenance	
Mai	intenance Single Sign-on Setting	
Mai	intenance Single Sign-on Setting Download Logs	8
Mai	intenance Single Sign-on Setting Download Logs Update Certificate	
Mai	intenance Single Sign-on Setting Download Logs Update Certificate Restart Service	

- 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.

Margh Read	y before	alian about the sharage cost	taken and resource arms.								
Same	1000	2020 (MAR 1997)	VVOL_V3P8000								
Description (reput storage) antilation description			- quine		1.1						
Haraga System Resource Grow			RAUDIOL_30909								
Capacity/Free Capacity/Free Logical Capacit Inspiritol Capacity	ly in casp loturi) p(friend) city(frie	ability profile for the DP posi Intelly InTeleS	that exists in the storage cont	ane.	GME / GME / GME /	INB INB					
2	-	Profiles									
1		Name	Descript	ini.	Pool	Pool Name	Pool Type	Total	Free	Total	Free
	- 61	(undefined)			0P.5	UCS_VVOL_Past	H07.	415.8898	455.8908		-
									0.0	IN TABLE IN	The Station Stationers

- 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**.

Sector and possible is decipited of the sequelikity ratio, and there related the regulations is the regulations Nerve Possible Postmenses Po	
Tame Tame Description Control State Description Control State Patienterse Control State Patienterse Control State Patienterse Control State Control Control Control State Control Control Control State Autor generated Control State Control State Autor generated Control State Control State Control Control State Control State	
Description Optimized Capabilities Image/Capabilities Image/Capabilities	
Managed Capabilities P Patimates (XPS - data P Deformance (Lating) - data P Control-class P Control-class P Recently fly Utilial MotorCure Integrater P O Staphilities P One Type One Speed P One Type One Speed P	
Destinances Lating:::Sea Tart Image: Sea (Sea (Sea (Sea (Sea (Sea (Sea (Sea	
Ø Avsishthy - dess Tart Cond- dess 0 Record yft What Hhamvoure Integrater 0 Stagshell Eaclog Instance - Class 0 Other Type Onio Speed -0	
Codcloss III IV Restanty by "Official Models Integrater IV IV Actor generated Copublicities IV IV University Topolowing Speed IV IV	
Raciney fty 'Uhui Infanivoure Integrater Stagsher Backap Inportance - Class Auto-generated Capabilities Drive Type Drive Speed - P O Cestaplication	
Auto-generated Expeditions - Case Auto-generated Expeditions Drive Type Drive Speed - P O Cetaplication	
Auto-ganezated Capabilities Drive Type Drive Speed -P O Deskyllication	
Dine Tge Dine Specific - P Obstylication	
Bited Type HDT (P Compression	
RADE RADEO R	
Energetion No.2	
Shaphal Na 3	
	000



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

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.

Manayement		About VASA	
		Welcom	e to VASA
Manage Sto	rage Systems		Service Name : Hitachi Storage
-		Service	Port Number : 50001
Manage Stora	age Containers	Database	Port Number: 50003
Capabilit	v Schomo	Se	ervice Version: 03.5.8
Capabilit	y schema		VASA Version: 1.0/2.0/3.0
Change (Credentials	Sup	ported Model : VSP 5500H
			VSP 5100H
	Software		VSP 5500
Update	oonnaro		
Update			VSP 5100
Update			VSP 5100 VSP Fx00
Update			VSP 5100 VSP Fx00 VSP Gx00
Maintenance			VSP 5100 VSP Fx00 VSP Gx00 VSP F1500
Maintenance			VSP 5100 VSP Fx00 VSP Gx00 VSP F1500 VSP G1500
Maintenance			VSP 5100 VSP Fx00 VSP Gx00 VSP F1500 VSP G1500 VSP G1000
Maintenance Single Sign	n-on Setting		VSP 5100 VSP Fx00 VSP Gx00 VSP F1500 VSP G1500 VSP G1000 HUS VM
Maintenance Single Sign	n-on Setting		VSP 5100 VSP Fx00 VSP Gx00 VSP F1500 VSP G1500 VSP G1000 HUS VM VSP
Maintenance Single Sign Downloo Update (n-on Setting bad Logs Certificate		VSP 5100 VSP Fx00 VSP Gx00 VSP F1500 VSP G1500 VSP G1000 HUS VM VSP
Maintenance Single Sign Downlo Update (Restart	n-on Setting bad Logs Certificate		VSP 5100 VSP Fx00 VSP Gx00 VSP F1500 VSP G1500 VSP G1000 HUS VM VSP

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

Virtual Storage					
contains •		GO OFF			
e Systems					
Storage Syste	ms		Model Type	SVP IP Address	Configuratio
-42.186-G5-U2-440138	13Poors	CLDEV's	VSP G600	172 25 42 186(172 25 42 187, 172 25 42 188)	-
e ASE-47.112 G10	(SPools	CLDEV's	VSP 5500H	172 25.47.112	-
ASE-47.112 G10	19Peols	OLDEV 's	VSP 5500H	172 25.47.112	-
	e Virtual Storage s v contains v e Systems Storage Syste 0-42.186-G5-U2-440138 e ASE-47.112 G10	e Virtual Storage s v contains v	e Virtual Storage a v contains v GO Ore e Systems Storage Systems 0-42.186-G5-U2-440138 @Pools @LDEVs e ASE-47.112 G10 @Pools @LDEVs	e Virtual Storage s v contains v 00 000 e Systems Storage Systems Model Type 0-42 186-GS-U2-440138 Becors Co.DEV:s VSP G600 e ASE-47.112 G10 Becors VSP 5500H	Wirtual Storage a v contains v GO OFF c Systems GO OFF c Systems Model Type SVP IP Address 0-42.186-G5-U2-440138 Copcorts Copcorts VSP G600 172.25.42.186(172.25.42.187, 172.25.42.188) a ASE-47.112 G10 Copcorts Copcorts VSP 5500H 172.25.47.112

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

					No. of Volu	mes		8	
n	stastore Na	ine w costain w		00	COT .				
-	and Mark and	and slab Dates Date	_						
	PCI VOLUEN	and click Denne Prone.							
	eline Profil								
Gaile	and Sections								
181	Volume	Datastore Name	Profile Status	Pool	Resource Group	Label	Device ID	vCenter Server	
2	99.99.12	VSI-OCP-DS_VSP5000	Custom	DP 4	meta_resource	VSI-OCP-DS-VSP5000	naa 60060+80087783000050778300000012	10.76.30.10	VSI_G700-0_Fab_8. VS
	A contractor of the			_					
	00.00.05	+	-	+	meta_resource	vVols-PE-ALU-8	naa 60060e8008778300005077830000006		ucp2k-c4-b4_8C, ucp2k
0	00.00.05	+	- Disable	- DP 2	meta_resource	VVolb-PE-ALU-8 CMD-AUTO-CREATED	naa 60060e8008778300005077830000006 naa 60060e8008778300005077830000008		ucp2k-c4-b4_8C. ucp2k- ucp2k-c4-b3_7C
0	80.00.00 80.00.00 90.00.00	+	- Disable	- DP 2	meta_resource meta_resource meta_resource	VV0b-PE-ALU-8 CMD-AUTO-CREATED VV0L-ALU1	naa 60060e8008778300005077830000006 naa 60060e8008778300005077830000008 naa 60060e80087783000050778300000008		ucp2k-c4-b4_8C, ucp2k ucp2k-c4-b3_7C ucp2k-c4-b4_8C, ucp2k
0000	00:00:05 20:00:00 00:00:00 00:00:00	+	- Disable - Disable	- DP 2 - DP 0	meta_resource meta_resource meta_resource	vVois-PE-ALU-8 CMD-AUTO-CREATED VVOL-ALU1 spc-37/3119232	mas 60060e8068778300005077830000006 mas 60060e8008778300005077830000008 mas 60060e80087783000050778300000009 mas 60060e80087783000050778300000004		ucp2k-c4-b4_8C, ucp2k ucp2k-c4-b3_7C ucp2k-c4-b4_8C, ucp2k spc-10000990fab4b714
00000	00:00:05 00:00:05 00:00:09 00:00:00 00:00:0F	+ + + +	- Disable - Disable -	- DP 2 - DP 0	meta_resource meta_resource meta_resource meta_resource	VV08-PE-ALU-8 CMD-AUTO-CREATED VV0L-ALU1 spc-37/3119232 UCS_PE_5000	naa 60050=0008778300005077830000006 naa 60060=00087783000050778300000008 naa 60060=008778300005077830000000 naa 60060=8008778300005077830000000		ucp2k-c4-b4_8C, ucp2k ucp2k-c4-b3_7C ucp2k-c4-b4_8C, ucp2k spc-10000090fab4b714 VSI_G700-0_Fab_8, VS
00000000	00:00:05 00:00:05 00:00:05 00:00:00 00:00:0F 00:00:0F	* * * * *	- Disable - Disable - Disable	- DP 2 - DP 0 - DP 0	meta_resource meta_resource meta_resource meta_resource meta_resource meta_resource	VV0Is-PE-ALU-8 CMD-AUTO-CREATED VV0L-ALU1 spc-37(3119232 UCS_PE_5000 24mme-b00	naa 60060-80087733000050778300000006 naa 60060-8008778300005077830000000 naa 60060-8008778300005077830000000 naa 60060-8008778300005077830000000 naa 60060-8008778300005077830000000 naa 60060-8008778300005077830000000		ucp2k-c4-b4_8C, ucp2k- ucp2k-c4-b3_7C ucp2k-c4-b4_8C, ucp2k- spc-1000090fab4b714- VSL_G700-0_Fab_8, VS ds120-4590 hbs1_1A, dt

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

Profile Tog					
writy select the republiky tags to be registered.					
DEV ID 08.09.12					
Aanaged Capabilities - Per Device Miler Dotneil					
Performance XXPS - Class		(fied_iOPS v)>			
Performance Latency - Class		South's Ten party of the second second second	100 In 10		
Availability - Class		Const. All Strategies and all			
Cest-Class	Cent - Class				
Recovery by Vitual Infrastructure Integrator					
C Snapshot Backup Importance - Class	TORICE.	1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1			
Auto-define Capabilities					
C Drive Type/Drive Speed	\$\$D (0	C Deduplicati	an .	Yes dr	
D Peel Type	= 10H	Compression	an .	Normal (2	
C RAD Level	RAIDER0+2PS IP				
C Encryption	No.24				
PY Building	Wate C				

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.

vm	vSphere Client Menu 🗸	Q Search in all environments		
() ~ <mark>69</mark> 10.7	B 6.30.10	Summary Monitor Configure		
	Actions - VSI_SC Add Host New Cluster New Folder Distributed Switch New Virtual Machine Deploy OVF Template	Hosts: 4 Virtual Machines: 98 Clusters: 1 Networks: 24 Datastores: 9 Custom Attributes		
	Storage •	New Datastore		
	Edit Default VM Compatibility	😵 New Datastore Cluster		
	Move To Rename Tags & Custom Attributes			
Recent Task Nan	Add Permission	Y Status		
	🔀 Delete	1		

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**.

New Datastore	Name and specify datastore	device sele	ction	ioning the datastor	0			
1 Type	Name:	VMFS_VSP50	00					
3 VMPE version	The datastore disk/LUN that disk/LUN	will be accessible to all you are interested in, 1	the hosts that might not be i	are configured with a accessible to that host	cess to the selecte Try changing the P	d dek(LUN: If you set or configure a	do not find the cossibility of that	×
4 Partition configuration	Select a host	esxi-0.vsi.hvla	blocal v					
5 Ready to complete		Select a host to o	ew its accessio	ie diskotunis				
	Name	Ŧ	LUN	T Capacity T	Hardware T	Drive Typ 🕇	Sector Fo 🔻	Cluster
	O HITACHIE	Fibre Channel Dis	100	25.00 GB	Supported	HOD	5120	Yes
	O HITACHI P	Fibre Channel Dis.	o.	5.00 GB	Supported	HOD	5t2n	Ves
	🙁 HITACH F	ibre Channel Dis	29	2.00 68	Supported	HDD	5t2n	Yes
	0							3 nems
						CANCEL	BACK	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 <u>Related Documents (on page 82)</u>.

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.

([]]) ~ <mark>(S</mark> 10.3	830.10	Summary Monitor Config
	Actions - V3L_SC Add Host New Cluster New Folder Distributed Switch New Virtual Machine Distributed Switch Distributed Switch	Hosts: Virtual Machines: Clusters: Networks: Datastores: Custom Attributes
	Storage	New Datastore
- 1	Edit Default VM Compatibility	양 New Datastore Cluster
	Another Network	🗒 Rescan Storage
Recent	Move To Rename Tags & Custom Attributes	
Task Nam	Alberta	Y Status
	S Delete	

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.

New Datastore	Name and cont	ainer selection			
1 Туре	Name: VV	Ind backing storage containe	K		
2 Name and container selection	Backing Storage Contai	iner			
3 Select hosts accessibility	Name	7	identifier	T M	simum Disk Size
4. Ready to complete	• vvol_vsps000		vvol 96cdfdd4cbce4ffc-afde3d9e3elff	e74c 60	тв
				_	
	0			_	Titer
	For SCB-backed vVid da datastore is created with	stastores. PELUNs need to be con hout configuring PELUNs, the ESX	rigured manually. Configure SCSI PE LUNis before Chost marks corresponding vivor datastore as ins	creating a da iccessible.	astore The
	Backing Storage Contai	iner Details			
	Electron property)	VSP 5500H 30595			
	Service Sec. 43(12)	the state details			
	Storage provider(s)	VASA			

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

New Datastore	Ready to comple	ete		×
	Review your settings select	ions before finishing the wizard.		
1 Type	🗸 General			
2 Name and container selection	Name	VVOL_VSP5000		
	Type:	10V/v		
3 Select hosts accessibility	V Backing storage contains	er details		
4 Ready to complete	Name	VVOL_VSP5000		
	UUID:	vvol 96cotod4cbce4ftc-atde3d9e3e8te74c		
	Storage array(s):	V\$P 5500H_30595		
	Storage provider(s)	VASA		
	V Hosts that will have acce	ss to this datatione		
	Hosts	esxi-t-vsi.tvrlab.local		
		essi-2.vsl.hvlab.local		
		ersk-Divisi.hvtab.local		
		M source and reading of the		
			CANCEL	BACK

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.

Shortcuts							
Inventories							
(])	B)		Ø	T	E	~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~	Q
Hosts and Clusters	VMs and Templates	Storage	Networking	Content Libraries	Global Inventory Lists	Workload Management	DRaaS
Monitoring							
		æ	R.	E.	\diamond	O	
Task Console	Event Console	VM Customization Specifications	VM Storage Policies	Host Profiles	Lifecycle Manager	UCP Advisor	
Administratio	n						
Q							
Licensing							

3. Select CREATE under VM Storage Policies.

Policies and Profiles VM Storage Policies	VM	1 Storage Policies	
C VM Customization Specifications	CRE	ATE	
Host Profiles	une		
🛱 Storage Policy Components		Name 1	vc
		Default-Tim	<u>6</u> 10.76.30.10
		R Host-local PMem Default Storage Policy	10.76.30.10
		😨 Management Storage policy - Encrypti	10.76.30.10
		R Management Storage Policy - Large	10.76.30.10
	0	😥 Management Storage Policy - Regular	10.76.30.10
		😰 Management Storage Policy - Single N	10.76.30.10
		R Management Storage Policy - Stretched	10.76.30.10
	0	😰 Management Storage Policy - Stretche	69 10.76.30.10

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

Create VM Storage Policy	Name and des	cription	×
1 Name and description	vCenter Server;	10.76 30 10 -	
2 Policy structure	No. of Concession, Name	Table Angle Table	_
3 Storage compatibility	name	Tangu-YMP3-Tiet2	<u> </u>
4 Review and fields	Description:		
			CANCEL

 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**.

Create VM Storage Policy	Tag based plac	ement	×
1 Name and description	Add tag rules to filter dat	astores to be used for placement of VMs.	
2. Policy structure	Rule 1		REMOVE
3 Tag based placement	Tag category	SPBM Y	
4 Storage compatibility	Usage option	Use storage tagged with ~	_
1. Review and finish	Tags	(Availability (7., X.) (Performance L. X.) (Performance L.) BROWSE TAGS	×)
	ADD TAG RULE		_
		CANCEL	BACK

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

Ita S. Inconsta di datastore cluste Is_VSP5000	Tribuil ris Delacenter VSL_SC	Type VMPS 6	Compa Free Space 408.33.08	etble storage 499.7 T Ed Capacity 499.75.08	75 GB (498.33 GB fre Ter Wannings
d datastore cluste	Datacenter VS_SC	Type VMPS 0	Compa Free Space 498.23.08	Etble storage 499.7 En Capacity 499.75.08	75 08 (498.33 08 fre Ter Warnings
IS_VSP5000	Detacenter VSLSC	Type V1425 6	Free Space 498.33 GB	Capacity 499.75.08	Wainings
5_V\$P5000	V9_SC	VMPS 6	498.33.08	499.75.08	
					1 der
				CANCEL	BACK
					CANCEL

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.

Shortcuts Inventories							
() Hosts and Clusters	VMs and Templates	Storage	2 Networking	Content Libraries	Giobal Inventory Lists	Workload Management	Q DRaaS
Monitoring							
Task Console	Event Console	VM Customization Specifications	VM Storage Policies	Host Profiles	Lifecycle Manager	UCP Advisor	
Administratio	n						

3. Select CREATE under VM Storage Policies.

	VM Storage Policies	
CREATE CREATE CREATE CREATE CREATE CREATE CREATE CREATE CREATE CREATE CREATE CREATE CREATE CREATE C	VM Storage Policies	
Storage Policy Components Name ↑ Vc Default-Tim Ø 10.76.30.10 E Host-local PMem Default Storage Policy Ø 10.76.30.10 E Management Storage policy - Encryptil Ø 10.76.30.10 E Management Storage Policy - Large Ø 10.76.30.10 E Management Storage Policy - Regular Ø 10.76.30.10 E Management Storage Policy - Single N Ø 10.76.30.10 E Management Storage Policy - Single N Ø 10.76.30.10	CREATE	
Name T VC Default-Tim Ø 10.76.30.10 Host-local PMem Default Storage Policy Ø 10.76.30.10 Management Storage policy - Encrypti Ø 10.76.30.10 Management Storage Policy - Large Ø 10.76.30.10 Management Storage Policy - Single N Ø 10.76.30.10		
Image: Construction Image: Construction Image: Constrediate Construction Image: Co	R Default-Tim	10 76 30 10
Image: Management Storage policy - Encrypti. Image: Management Storage Policy - Encrypti. Image: Management Storage Policy - Large Image: Management Storage Policy - Regular Image: Management Storage Policy - Regular Image: Management Storage Policy - Single N. Image: Management Storage Policy - Single N	R Host-local PMem Default Storage Po	blicy 🙆 10.76.30.10
Image: Management Storage Policy - Large Imagement Storage Policy - Large Imagement Storage Policy - Regular Imagement Storage Policy - Regular Imagement Storage Policy - Single N. Imagement Storage Policy - Single N. Imagement Storage Policy - Single N. Imagement Storage Policy - Single N. Imagement Storage Policy - Single N.	Anagement Storage policy - Encry	pti
Image: Management Storage Policy - Regular Imagement Storage Policy - Regular Imagement Storage Policy - Single N. Imagement Storage Policy - Single N. Imagement Storage Policy - Single N. Imagement Storage Policy - Single N. Imagement Storage Policy - Single N. Imagement Storage Policy - Single N. Imagement Storage Policy - Single N.	Management Storage Policy - Large	10.76.30.10
Anagement Storage Policy - Single N. 10.76.30.10	🗌 📄 🛱 Management Storage Policy - Regul	lar 🚱 10.76.30.10
Management Storage Policy - Stratched 60 10 75 20 10	Management Storage Policy - Single	N 🚱 10.76.30.10
U Ig management atorage voircy - stretched (b) 10.70.30.10	Management Storage Policy - Streto	ched 🙆 10.76.30.10
🗌 📄 🗟 Management Storage Policy - Stretche 🛛 🥙 10.76.30.10	🗌 📄 🗟 Management Storage Policy - Streto	:he 🚱 10.76.30.10
		VM Storage Policies CREATE Name CREATE CREATE CREATE CREATE CR

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

Create VM Storage Policy	Name and des	cription	×
1 Name and description	vCenter Server:	💕 10.76.30.10 v	
2 Policy dructure	Name:	Tanzu-VMP5-Tier2	
 Starage compatibility Review and finish 	Description		
	area plant.		
			CANCEL

 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**.

Create VM Storage Policy	com.hitachi.storageprovid	der.vvol rules	
1 Name and description	Placement Tags		
2 Policy structure	Performance IOPS - Class 🔘	Tiert_IOPS	REMOVE
3 com.hitachi.storageprovider.vvol rul		Ter3_IOPS	
4 Storage competitivity	Performance Latency - Class ()	Tiert_Latency	REMOVE
3. Review and Inside		Tier2_Latency	
	Availability - Class 🕦	Tert	REMOVE
	ADD RULE -	0 H	
		CANCE	E BACK NEXT

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

1. Name and description	COMPATIBLE INCOMP	ATIBLE				
1 Henry and Metchinish	Expand datastore clust	ters		Comp	atible storage 347.	12 GB (347.12 GB fr
2 Policy structure					• Fit	er.
3 com hitachi.storageprovider.vvol ruk	Name	Datacenter	Type	Free Space	Copecity	Warnings
4 Storage compatibility	VVOL_VSP5000	vsi_sc	wwoł	347.12 08	347.12 GB	
5. Review and Trrish						
						5.8w
	a					5 Ba

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.

vm vSphere Client Menu v Q Se	arch in all environments
A Home	Llomo
Shortcuts	Home
[]] Hosts and Clusters	10.76.30.10
VMs and Templates	
E Storage	CDU
Ø Networking	CPU
Content Libraries	
% Workload Management	
Global Inventory Lists	
R Policies and Profiles	
∑ Auto Deploy	⊕ VMs
log Hybrid Cloud Services	
<>> Developer Center	
C Administration	Pow
(創) Tasks	
Events	
🛇 Tags & Custom Attributes	Objects
☆ Lifecycle Manager	
UCP Advisor	Item
@ DRaaS	HCP1_DS_C
() vRealize Operations	HCP4_DS_
	10.76.30.10
	HCP2_DS_

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

Content Libraries	1
VSI_Content_Lib_G600	Content Libraries
	Advanced + Create
	Name ↑
	VSI_Content_LIb_G600

4. Define a Name, and then click Next.

New Content Library	Name and lo Specify content libr	ocation ary name and location.	×
1 Name and location	Name:	wco-kubernetes	
2 Configure content library	Neter		
3 Add storage	Notes:		
4 Ready to complete	vCenter Server:	10.76.30.10 ~	
			CANCEL

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

New Content Library	Configure content library × Local libraries can be published externally. Subscribed libraries originate from other
1 Name and location	published libraries.
2 Configure content library	O Local content library
3 Add storage	Enable publishing Enable authentication
4 Ready to complete	Subscribed content library
	Subscription URL https://wp-content.vmware.com/v2/latest/lib.json
	Enable authentication
	Ownload content immediately when needed
	CANCEL BACK 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 <u>https://github.com/haproxytech/vmware-haproxy#download</u>.



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.

vm vSphere	e Client Menu 🗸 🔾	Search in all environme
	3 0	
VSI_SC		
VSI_Clus	ter	
🗐 esxi-([] Actions - VSI_Cluster	
🖡 esxi-1	📑 Add Hosts	
🖡 esxi-1	🔂 New Virtual Machine	
> @ Cisco	🧭 New Resource Pool	
> 🕢 Hitac	(#) New Namespace	
> @ HKS-L > @ SAP	🕏 Deploy OVF Template	1
⊖ Tanzi	B New vApp	
✓ ⊖ тк-u В тк	Import VMs	
С ТК	Storage	•
🔂 ТК Ю ТК	Host Profiles	•
🕃 тк	Edit Default VM Compatibility	.2
⊘ vSph	Licensing	F
ලා Benc- කි Centr	Settings	

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



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

Deploy OVF Template	Select a name ar	nd folder			
	Specify a unique name and	target location			
1 Select an OVF template	Virtual machine name:	haproxy-Tanzu			
2 Select a name and folder	Select a location for the virt	ual machine.			
3 Select a compute resource	✓ ∅ 10.76.30.10				
4 Review details	→ E vs_sc				
5 Select storage					
6 Ready to complete					
			CANCE	BACK	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**.

eploy OVF Template	Select networks select a destination network for each s	ource network.		
1 Select an OVF template	Factoriation	- Factorial and the second		
2 Select a name and folder	Management	Acc-30		
3 Select a compute resource	Workload	VLAN64	~	
Review details	Frontend	App-25	÷	
License agreements	0			3 Au
Configuration	IP Allocation Settings			
Select storage	IP allocation:	Static - Manual		
Select networks	IP protocol:	iPv4		
Customate template				
3 Ready to complete				

- **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.
 - I. Define a HAProxy user and password.
- 12. Click Next.

Deploy OVF Template	Customize template	tine andere is landstation out and appendices wherein maintagenticities with
1 Select an OVF template		Oroup in CIDH format (Big. sp/subnet mask bits). This cannot be DHCP.
2 Select a name and folder		10.76.30.78/24
3 Select a compute resource	2.4. Management Gateway	The gateway address for the workload network. This is also the default gateway for the appliance.
4 Review details		10.76.30.1
5 License agreements	2.5. Workoad IP	The static IP address for the appliance on the Workload Port Group in CIDR format (Eg. lp/subnet mask bits). This IP must be outside of the Load Balancer IP Range
6 Configuration		102-168-64-2/24
7 Select storage	2.6 Workload Gateway	The gateway address for the workload network
R Select naturalis		102 168 64 1
e succession	v 3. Load Balancing	4 settings
9 Customize template	3.1 Load Balancer IP Ranges, comma-	The IP ranges the load balancer will use for Kubernetes Services
10 Ready to complete	separated in CIDR format (Eg 1.2.3.4/28.5.6.7.8/28)	and Control Planes. The Appliance will currently respond to ALL the IPs in these ranges whether they're assigned or not. As such, these ranges must not overlap with the IPs assigned for the appliance or any other VMs on the network. 192368.64.16/28
		CANCEL BACK NEX

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

Deploy OVF Template	Ready to com	plete
	Storage mapping	1
1 Select an OVF template	All disks	Datastore: UCS-VMPS-G600_10K; Format: Thick provision lazy zeroed
2 Select a name and folder	Network mapping	3
3 Select a compute resource	Management	App-30
	Workload	VLAN64
4 Review details	Frontend	App-25
5 License agreements	IP allocation settings	
6 Configuration	IP protocol	IPV4
e congrana	IP allocation	Static - Manual
7 Select storage	Properties	12. Permit Root Login « True
8 Select networks		1.3. TLS Certificate Authonity Certificate (ca.crt) = 1.4. TLS Certificate Authority Private Key (ca.key) = 2.1. Host Name = Buschardenical
9. Outomite template		2.2. DNS = 10.76.31.254
5 costomze template		2.3. Management IP = 10.76.30.78/24 2.4. Management Gateway = 10.76.30.1
10 Ready to complete		2.5. Workload IP + 192.168.64.2/24 2.6. Workload Gateway a 192.168.64.1
	-	31 Load Balancer IP Ranges, comma-separated in CIOR format (Eg 12.3.4/28.5.6.7.8/28) = 192.1
		68.64.16/28 3.2 Datables ACI Management Dwt = 5556
		3.3 HADroay Litter ID = admin

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.

1 vCenter Server and Network	Select a vCenter Server and a network to enable a cluster	
C 1/2 > A You must config Management of	une a total balance instance with your rightene environment before you set up Wonkload Management. You cannot complete the Wonkload high without a load balancer instance. Supported load balancers are HAProxy and Au.	Learnmore 🗙
To enable Workland Management on a cluste	e, select the vCenter Barver system that nods the cluster	
Select a vCenter	🔘	
Select the networking stack that will provide	contractivity to the Workload Management platform	
Select a networking stack option	NSK-T (NCE Assister) Substration States (Subserved States) vCenter Server Network Supports Tarou Kubernetes clusters	
SEXT		

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

Cluster Details 10.76.30.7 COMPATIBLE INCOMPATIB	0 1			
Civeler Name	 Number of Hosts 	· Available CPU	y Analable Memory	
C VS_Outer	4	195.87 GHz	540.49 GB	
0				t

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

÷	3.	Control Plane Size	Select the size and resources available for control plane VM on this cluster
		Alocate capacity for the Ru support.	ubernetes control plane VMs. The amount of resources that you abocate to the control plane VMs determines the amount of Kubernetes workloads the cluster can
		Resource allocation	O tay +2 CPUL 8 08 Memory 16 08 Storage
			Small - 4 CPUs, to GB Memory, 10 GB Storage Orderstone, 18 CPUs, 18 CB Memory, 19 CB Storage
			O Large - N CPUI, 32 Gil Memory, N Gill Storage
	-	_	
	Ľ	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.

Configure a load balancer to support the re-	twolv connectivity to workloads from claimt networks and to load balance traffic between Tanzs Kuberhates cluste	rs Supported load balancers are HAProxy and Ave.
Name 💿	35004y	A VIEW NETWORK TOPO
Type ()		
Data Plane API Address(an)	10.76.30.78.8664	
Duer namer @	Adva.	
Fesseord ()		
Vitual IP Address Banges 💿	112 188.84 17 113 168 64 10 Tapanite multiple project by utervise	
ierver Cortificate Authority 🕕	stangi nani Palyang si naga palya kaya katala kaya palya kaya nana kaya kaya kaya kaya kaya ka	

- **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.

The Workload Management consists of time	Kubernelies control plane VMs and the Sphe	which process on each host, which allows the hosts to be jun	ned in a Kultermetes studier. The cluster where you set up Worklaat
Management is converted to a management (Age 32		A VEW NETWORK TOPOLOGY
Durling IP Address (1)	10.7630.147		
Subret Mark	288.265.259.0		
Gatzway ()	10.76.201		
ONS Server (D)	10.78.37.254		
ONS Search Osmains (Optional)	(a) more than (
NITP Server (1)	10.76.301		

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

and the second se								
Services P address Details for service P detail	ts and the Jefs	of vece must to be abled here.					A VEW NETWOR	K TOPOLOSY
P Address for Services			0.160.023					
DNS Servers			10.71.21.254					
Workload Network								
Workload Network This can left workload net 400 IIIII I	prove of prove	n la pour estimada in Baterinan C	Later. This we arrive for more security	peranetars between werecade.				
Workload Network You can add workload net 4000 I 2017 - 1 Name		n is your worklass in Eusenvisor C while Destinated Sector	when the we arrow for more becarily	peranetters between activities.	•	Subset 🗸	IP datasets barget	
Washined Network The Carl and Holician ref 4000 (2017 - 1) Aprile		n to your exercises in futer-sole of	Unter This will allow for more security	(a ymethe Bathaan a chwraat). • Anteany	•	Subject +	P Address barges	
Norkisad Network This can all with Coal Art 400 Start I	icities to acception icities t	n to your exercises in Euterward wheelbentschet (with No Networks	the first area for new security	eranetters between workvees.	•	salaar •	и алтан түүрэ	

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.

Services IP address								I come antiques
Details for service IP details and the default value	e input to be added here.							A VAN MATWON
IP Address for Services		10 96 0 0/23						
ONS Servers		10.76.31.254						
Workbaad Network This can add workbad networks to assign to you ADD EDIT REPROVE	e workloads in Supervisor Custe	r. Thy will alow	for more security pa	rameters beh	een workoads.			
Note y	Virtual Distributed Seriels		Part Group	۲	Ontenay	 tubret	,	P Address Ranges
O neteorist (foren)	Approxim-Dilwitch		VDAN64		192,968,64.1	255 255 255.0		192168.6432-192168.8450
NEXT								

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

8.	Tanzu Kubernetes Grid Configuration	Service	Set up the Tanzu Kubernetes Grid Service to enable self-service of Tanzu Kubernetes clusters for your developers.
Cont	tent Library	wcp-kubernetes	sor .
N	EXT		

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.

vm vSphere Client Menu v
分 Home
& Shortcuts
[]] Hosts and Clusters
B VMs and Templates
Storage
Content Libraries
🗞 Workload Management
🐻 Global Inventory Lists
Policies and Profiles
Auto Deploy
Hybrid Cloud Services

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

Namespaces	Clusters Updates	
You have	successfully enabled Workload Management	
This is your gio	ar view of namespaces across your clusters.	
Next Steps		
To get start	d with namespace centric management, create your first namespace. Learn more	e
You can also	add a content library to support Tanzu Kubernetes clusters. Go to Content Libra	iry
CREATE NA	IESPACE	
	NI	

- 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.

Itatus = bearen Stillon Running usermens Status Active coation VS_Course () 10,263010 () 20,263010 () 10,263010 () 10,26301 () 10,26301 () 10,2630 () 10,263	You haven't given any devops access to this namespace. Add some permissions to lat your devops team directly manage this namespace.	Vou heven't added any storage policies for the namespace. Add some policies to list your devaps team access pensistent storage	Capacity and Usage III CPU Notes Memory Notes D MB Biorage Notes	Tanzu Kubernetes Grid Service Tanzu Kubernetes clusters Content Library 601 with Kubernetes
CONTRACT COMM	ADD REPRESENCES	400 \$30840F	and course	VIEW ALL

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

Add Permissions

 \times

Add a user or a group to give access to this namespace

dentity source	vsphere.local		~
Jser/Group Search	Q_Administrator		
Role	Can edit		~
		CANCEL	ок

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

(#) DS1 ACTIONS ~ Summary Monitor Configure	Permissions Compute Storage	Network		
Status E Created 5/5/21 Config Status () P Running Kubernetes Status () Active Location () V& Custer () 10/26/30/10	Permissions [] Can view () 36 Januari, Ruw (permission to only view namespaces. Can edit () Administrator	Storage	Capacity and Usage III CPU No.com O Mrs Memory No.com O Mu Storage No.com	Tanzu Kubernetes Grid Service Tanzu Kubernetes Outers Content Library 901 wcb-kubernetes
Link to CLI Tools Copy link () Open ()	MANAGE PERMISSIONS	ADD STORAGE	BOIT LIMITE	VEW ALL

CANCEL

OK

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

>	Tanzu-VMFS-Tier2	499.75 GB	354.08 GB
>	Tanzu-vVol-Tier1	347.12 GB	347.12 GB
>	VVol No Requirements	1.04 TB	1.02 TB
>	UCS-CNS-Tier2_Silver	499.75 GB	354.08 GB
>	UCS_VVOL_Tier1Gold	1.04 TB	1.02 TB
>	Default-Tim	2.00 TB	704.48 GB
>	VM Encryption Policy	26.32 TB	6.42 TB
>	UCS-CNS-StoragePolicy	2.00 TB	704.48 GB
	Storage Policy	Total Capacity	Available Capacity

Select Storage Policies

 \times

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.

F 🖂 📒 🖛			Compressed Folder Tools	Downloads		
File Home 5	Share	View	Extract			
← → ~ ↑ 🕹	> This	PC > Do	wnloads >			
 Quick access Desktop Documents Pictures This PC Desktop Documents Documents Documents Pictures Videos Local Disk (C:) Network 	> This	PC > Do Name	vnloads > ^ ere-plugin	Date modified	Type Compressed frinn ow s Defender	Size 28,582 KE
				Rename		÷

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

A REAL PROPERTY AND A REAL						
Organize • New fold	er				168 ·	- (
Cuick access Cuick access Covenioads Covenioads Pictures This PC Documents Covenioads Music Pictures Videos	Name HORCM PerfLogs Program Files Vogram Files (x86) Users Windows	Date modified 2/9/2021 8:50 PM 7/16/2016 6:23 AM 5/5/2021 10:53 AM 2/28/2020 6:54 PM 5/5/2021 6:44 PM	Type File folder File folder File folder File folder File folder File folder	Size		
Local Disk (C:)						
Network Y						

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.

Power Options	
Event Viewer	
System	
Device Manager	
Network Connections	
Disk Management	
Computer Management	
Command Prompt	
Command Prompt (Admin)	
Task Manager	
Control Panel	
File Explorer	
Search	
Run	
Shut down or sign out	>
Desktop	

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



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

System Properties	×
Computer Name Hardware Advanced Remote	
You must be logged on as an Administrator to make most of these changes	\$.
Visual effects, processor scheduling, memory usage, and virtual memory	
Settings	
User Profiles	
Desktop settings related to your sign-in	
Settings	
Startup and Recovery	
System startup, system failure, and debugging information	
Settings	
Environment Variables.]
OK Cancel Apply	

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

	Value					
Path	%USERPROFILE%\AppData\Local\Microsoft\WindowsApps;					
TEMP	%USERPROFILE%\AppData\Local\Temp					
TMP %USERPROFILE%\AppData\Local\Temp						
	New Edit Delete					
stern variables						
Variable	Value					
ComSpec	C:\Windows\system32\cmd.exe					
NUMBER_OF_PROCESSORS	2					
US	Windows NI					
	C\Windows\system32C\WindowsC\Windows\System32\Wberr;					
Path	COMPACTION AND AND AND ADDRESS OF AN ADDRESS AND A ADDRESS AND A ADDRESS AND A ADDRESS					
Path ANTURYT PROCESSOR ARCHITECTURE	AMD64					
Path PROCESSOR_ARCHITECTURE PROCESSOR_IDENTIFIER	AMD64 Intel64 Family 6 Model 85 Stepping 4. GenuineIntel					

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

omputer Name Har	dware Advanced	Remote	
You must be logged	on as an Administr	rator to make most of these cha	nger
Performance			
Visual effects, proc	essor scheduling,	memory usage, and virtual mem	ory
		Settings.	0
User Profiles			
Desktop settings re	lated to your sign-i	n	
		Settings.	
Startup and Recov	ery		
System startup, sys	tem failure, and de	bugging information	
		Settings.	
		Environment Variab	les

- 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 <u>Assigning User</u> <u>Roles and Permissions to Supervisor Cluster (on page 58)</u>.
- 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 <u>Creating a Namespace (on page 57)</u> using the following command:

kubectl config use-context <namespace>
To change context, use `kubectl config use-context <workload name>
PS C:\Users\Administrator> kubectl config use-context ns1
Switched to context "ns1"
Site con

Tanzu Kubernetes Cluster deployment

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

See the <u>Workflow for Provisioning Tanzu Kubernetes Clusters</u> 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.
- Log in to the Supervisor Control Plane IP. When prompted provide the vSphere username and password enabled in <u>Assigning Users Roles and Permissions to</u> <u>Supervisor Cluster (on page 58)</u>.

```
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 <u>Assigning Storage Policies (on page 59)</u>.

apiVersion: run.tanzu.vmware.com/vlalpha1 #TKGS API kind: TanzuKubernetesCluster #required

#TKGS API endpoint
#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

PE California Administra	tor Decktool	Vite Scheet Casely of	Part Sand				
tanzukubernetescluste	- run, tanzu, v	muare.com/tkos-cluster	-1 created				
P5 C:\Users\Administr	stor\Desktop\	YAML> kubect1 get tkc					
NAME CONTI	IOL PLANE W	ORKER DISTRIBUTION		AGE	PHASE	TKR COMPATIBLE	UPDATES AVAILABLE
tkgs-cluster-1 1		v1.19.7+vm/are	.1-tkg.1.fc82c41	38335	creating	True	[1.20.2+vmware.1-tkg.1.1d4f79a]
P5 C:\Users\Administra	tor\Desktop\	YAML> kubect1 get tkc					
NAME CONTI	IOL PLANE W	ORKER DISTRIBUTION		AGE	PHASE	TKR COMPATIBLE	UPDATES AVAILABLE
tkgs-cluster-1 1	10 C 10 C 10	v1.19.7+vm/are	.1-tkg.1.fc82c41	76175	running	True	[1.20.2+vmware.1-tkg.1.1d4f79a]
PS C:\Users\Administr	stor\Desktop\	YAMLS _					

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 <u>Assigning Users Roles and Permissions to</u> <u>Supervisor Cluster (on page 58)</u>.

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

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

apiVersion: run.tanzu.vmware.com/vlalphal	#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 <u>Assigning Users Roles and Permissions to Supervisor Cluster (on page 58)</u>.

```
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	
PS C:\Users\Administrator\Desktop\YAML> kubect1 vsphere loginserve	r=192.168.64.17insecure-skip-tls-verify
Username: administrator@vsphere.local KUBECTL_VSPHERE_PASSWORD environment variable is not set. Please enter Password: Logged in successfully.	the password below
You have access to the following contexts: 192.168.64.17 nsl tkgs-cluster-1 tkgs-cluster-2	
If the context you wish to use is not in this list, you may need to tr logging in again later, or contact your cluster administrator.	y .
To change context, use 'kubectl config use-context workload name' PS C:\Users\Administrator\Desktop\YAML> kubectl get storageclass NAME PROVISIONER RELIAINFULLY VULWEBING tanzu-vwfs-tier2 csi.vsphere.vmware.com Delete Immediate tanzu-vv0-tier1 csi.vsphere.vmware.com Delete Immediate PS C:\Users\Administrator\Desktop\YAML>	DINGMODE ALLOWVOLUMEEXPANSION AGE true Sd18h true Sd18h

4. The policies defined in <u>Assigning Storage Policies (on page 59)</u> 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 <u>Assigning Users Roles and Permissions to</u> <u>Supervisor Cluster (on page 58)</u>.
- **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

and the second second second second	COLOR BOARD COLORS COL	Contactory of the second	a no como o				
o change context,	use kubecti config use-	context work loa	d names				
S C:\Users\Adminis	trator\Desktop\YAML> kub	ectl get storage	class				
AME	PROVISIONER	RECLAIMPOLICY	VOLUMEBINDI	ENGNODE	ALLOW	OLUMEEXPANSION	AGE
anzu-vmfs-tier2	csi.vsphere.vmware.com	Delete	Inmediate		true		5d18h
anzu-vvol-tier1	csi, vsphere, vmware, com	Delete	Innediate		true		Sd18h
C:\Users\Adminis	trator Deskton VAM > but	ectl apply -f m	in vani				
and the study of the second	in a contract constant	receipting of br	ci yuur				
ersistentvoruneera	im/my-pvc created	and the second					
S C:\Users\Adminis	trator\Desktop\TAML> kub	ecti get pvc	Constant and the second se			CONTRACTOR OF A REAL PROPERTY.	
AME STATUS \	OLUME		CAPACITY	ACCESS	MODES	STORAGECLASS	AG
y-pyc Bound p	vc-82b92beb-eaf2-42d4-b6	09-eSba51f0695a	SGi	RWO.		tanzu-vmfs-tier2	75
S C:\llsecs\Adminis	trator\Deskton\YAML>						

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 <u>Assign user roles and permissions to supervisor cluster (on page 58)</u>.
- 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 <u>Example Role Bindings for</u> <u>Pod Security Policy</u>.

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 <u>Assigning Users Roles and Permissions to</u> <u>Supervisor Cluster (on page 58)</u>.
- 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
```



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:



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.

vm vSphere Client Menu v O	Search in all environments
[]] & e ø	I VSI_Cluster
 ✓ 個 10.76.30.10 ✓ 目 VSL_SC 	Summary Monitor Configure
VSI_Cluster esxi-0.vsi.hvlab.local esxi-1.vsi.hvlab.local esxi-2.vsi.hvlab.local	VM DRS Score CPU Utilization Memory Utilization

- 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.

Webby VM DRS Score DRU Unitation Memory Ubitation	Cont	ainer p	Volumes woolen: Kubernetes communication () 12 <u>augus</u>	Ther			(Even when
Network USAbatien vSphere HA V	0		Voture Name	Line .	B pvc-9c7300fa-04	20-452c-b943-5a1cf5556a81	×
Summary received	0	80	E motostalija tiet-44.	-	Seales Rubermeters on	prote Block	
Configuration Issues Datastores under APO or P.	0		篇 mc-8c73009s-0428-4	-	Volume ID	46e45204-1946-4524-1986-3e5656545ced	
Resource Adocation 🛛 🗸 🗸					Volume Backing Object (O	Tela 800805315077635077830000000000HH	
(PU Henory					Vourse Path	[VV0L_V5P5000] rea.60060028507783800778380000000453785836446464646464646408077866758677867581 wrstk	10
Bitorage					VM.	18 tilgs Custor 2 workers 62spn 554stellud when	
Storage Overview					Datastore	El ANOLTANADO	
Security					Elorade Policy	Tariza vV6i Tiert	- 84
Cloud Native Storage V	0		That I servery a	-	Compliance Status	Compliant	1

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 <u>Related Documents</u> (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.

A Home	5					
Shortcuts	Shortcuts					
Hosts and Clusters VMs and Templates Storage Networking Content Libraries	Inventories	VMs and Templates	Storage	O Networking	Content Libraries	Global Inventory Lists
Workload Management Goobal Inventory Lista	Monitoring					
Policies and Profiles Auto Deploy Hydrid Cloud Services Developer Center	Task Console	Event Console	VM Customization Specifications	VM Storage Potcies	Host Profiles	Lifecycle Manager
¹ ² 6 Administration ■ Tasks ■ Events S Tags & Custom Attributes © Lifecycle Manager	Plugins					
UCP Advisor	Administratio	m				
VRealize Operations	Q					

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

Username (usernar	ne@domain)	
administrator@	vsphere.local	
Password		
		×

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.

Add System

×

System Name	
UCS-StorageOnly	
Model	
UCP CI	~
Serial Number	
UCP-CI-000001	
Gateway Address	
10.76.30.20	
	_
Cancel	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: Before onboarding the VSP storage system, verify command device allocation to the UCP Advisor gateway.

Procedure

Ξ

- 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.

Add Storage	×
Serial Number	
440138	
Address	
172.25.42.186	
Username	
maintenance	
Password	
Controller 1 Address (Optional)	
Controller 2 Address (Optional)	
Cancel	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.

/SP G600 (440138)	VSP 5500H (3059	5)
Management Address Serial	Management Address	Serial
172.25.42.106 440138	172 25 47 112	30595
Allocated / Free Total	Allocated / Free	Total
99.4 TB / 6.16 TB 107.56 TE	21.79 TB / 11.16 TB	32.95 TB

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

Create Stora	ge Pool				Storage Pool	\$	
							\ \ Filter
Storage Pool	Name	Туре	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	нті	0	0.8	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
0							

- 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 Storage Pool

JCS_UCPA_Pool				
source Group (Op	otional)			
Select resource gro	pup			~
ol Type				
200 A 110 A 10 A 10				
Dynamic Pool (HD	P)			~
Dynamic Pool (HD Dol Volumes Parity Group	P) Drive Type	Size (G	B)	Ŷ
Dynamic Pool (HD Dol Volumes Parity Group 1-4 (3.33 TB V	P) Drive Type SSD	Size (G 200	B)	↓ Volumes
Dynamic Pool (HD Dol Volumes Parity Group 1-4 (3.33 TB ~ Parity Group	P) Drive Type SSD Free Space	Size (G 200 Drive Type	B) Add Size (GB)	Volumes Action

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.

VSP G600 (44013	8)	VSP 5500H (305	95)
Management Address 172 25 42 106	Serial	Management Address 172 25 47 112	Serial
Allocated / Free	Total	Allocated / Free	Total
99.4 TB / 8.16 TB	107.56 TB	21.79 TB / 11.16 TB	32 95 TB

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

					Select a Storage	detail	
Configure D	atastore				Datastores		
							₽F
Name	LUNs/NFS Export	ESXi Host	Allocated Capacity	Free Capacity	Total Capacity	Status	Туре
VMFS_V	18	esxi-2.vsi	193.61 GB	306.14 GB	499.75 G8	Mounted	VMFS

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

lame and Capacity	Name and Capacity	
torage Configuration	Datastore Name*	
ost or Cluster	UCPA_DS	\odot
	Max length: 32	
	Datastore Capacity*	
	100	\odot
	Capacity in GB, only integers allowed	
	O Single Datastore Creation	
	O 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.

Name and Capacity	Storage Configuration	
Storage Configuration	Storage System*	
Host or Cluster	30595	¥
	Resource Group	
	Select	×
	Storage Pool*	
	UCS_UCPA_Pool (HDP)	×
	Capacity Saving Mode	
	Select	¥

- 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.

· · · · · · · · · · · · · · · · · · ·	Provision Datastore	
omit Cancel	Subr	

k. Click submit.

Related Documents

This document references the following documentation:

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

Solution References

For more information on Hitachi solutions and products, go to <u>HitachiVantara.com</u> and see the following solution references.

- Network
 - <u>Cisco Nexus 9000 Series Switches Data Sheets</u>
 - <u>Cisco MDS 9000 Series Multilayer Switches</u>
- Compute
 - <u>Cisco Unified Computing</u>
 - <u>Cisco UCS 6400 Series Fabric Interconnects Data Sheet</u>
 - <u>Cisco UCS 5100 Series Blade Server Chassis Data Sheet</u>
 - <u>Cisco UCS VIC 1440 Adapter Data Sheet</u>
 - <u>Cisco UCS Manager</u>
- Storage
 - <u>Hitachi Virtual Storage Platform 5000 Series</u>
 - <u>Hitachi Virtual Storage Platform F Series All-Flash Enterprise Cloud Solutions</u>
 - <u>Hitachi Virtual Storage Platform G Series Hybrid-Flash Midrange Cloud Solutions</u>
- Virtualization Layer
 - <u>VMware vCenter Server</u>
 - <u>VMware vSphere</u>
- Compatibility Matrixes
 - Hitachi Interoperability Reports
 - <u>VMware Compatibility Guide</u>
 - <u>Cisco UCS Hardware and Software Compatibility</u>

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 <u>Customer Contact Us</u>.
- To open a new support case, see the <u>How to Create a New Case on the Support Website</u>.

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