PURE VALIDATED DESIGN

VMware vSphere with Tanzu and FlashArray

Pure Validated Design

Contents

Executive Summary

Enterprises are evaluating and building container-based applications within their environments, primarily leveraging Kubernetes as the orchestration platform of choice. Application modernization enables the delivery of innovative products and services to capture market opportunities. It also enables rapid development by empowering application developers and IT operations to partner more closely. Cross-functional collaboration speeds innovation by using existing tools and common technologies to eliminate operational inefficiencies while reducing costs.

VMware Tanzu extends the value of vSphere investments by offering flexible deployment options to enable Kubernetes capabilities to application development teams. Pure Storage® delivers best-in-class solutions with deep VMware and Kubernetes integration for efficient and effortless operations.

A Pure Validated Design (PVD) means that Pure has integrated and validated their leading-edge storage technology with an industry leading application solution platform to simplify deployment, reduce risk, and free up IT resources for business-critical tasks. The PVD process validates a solution and provides design consideration and deployment best practices to accelerate deployment. The PVD process assures the chosen technologies form an integrated solution to address critical business objectives. This document provides design consideration and deployment best practices for VMware vSphere Tanzu and FlashArray™.

Introduction

This document describes the benefits of enabling Kubernetes container services on VMware vSphere with Tanzu using Pure Storage FlashArray persistent and shared storage. It is a validated design that includes considerations, deployment specifics, and configuration best practices for a developer-ready environment.

To deliver on the solution described in this document, the VMware practitioner needs to deploy a production-ready VMware vSphere environment. The vSphere environment must include a vCenter Server Appliance, ESXi cluster of three or more hosts, vSphere Distributed Switching, and Pure Storage FlashArray datastore.

Solution Overview

This Pure Validated Design is based on VMware vSphere with Tanzu Basic and Pure Storage FlashArray using vSphere Virtual Volumes (vVols) for Persistent Storage. vSphere with Tanzu utilizes core vSphere technologies and introduces Tanzu

Kubernetes services for the VMware practitioner to quickly build a Kubernetes Namespace for developers with Persistent Volumes on FlashArray.

This validated design includes an example of a stateful web application with a MySQL backend contained in a Tanzu Kubernetes cluster. WordPress is the web-based frontend application that will use persistent storage with FlashArray vVols for MySQL. Persistent storage is declared in the design yaml files using the Kubernetes' persistent volume claim with VMware Cloud Native Storage (CNS).

Figure 1 depicts a high-level framework for this Pure Validated Design.

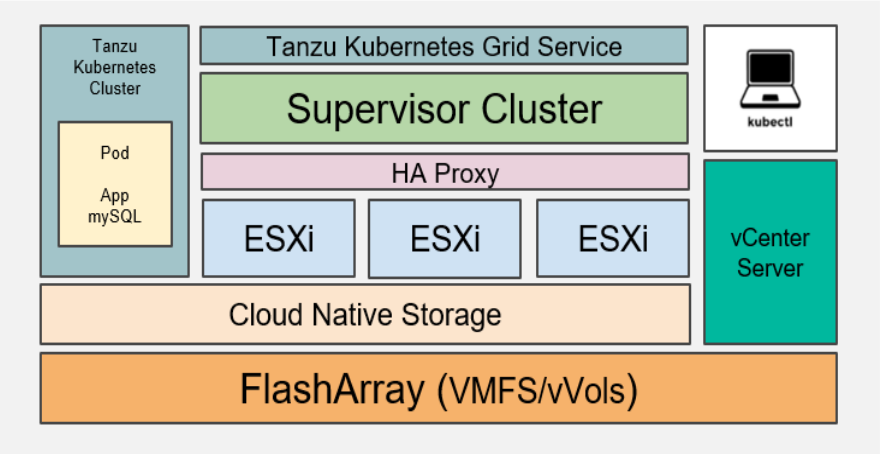

Figure 1. VMware vSphere with Tanzu and Pure Storage FlashArray datastores.

Solution Benefits

This solution enables organizations to easily evaluate and deploy Kubernetes capabilities through consistent tools with VMware vCenter Server to accelerate application modernization with Kubernetes and Persistent Storage with vVols on Pure Storage FlashArray. VMware Practitioners will be able to manage and command resources to effortlessly provide development teams with Kubernetes Namespaces.

Business Value: Accelerate app modernization and deliver innovative products and services faster with existing VMware investments while minimizing the need for adding net new resources. Pure Storage gives VMware environments simplified management, optimized data center resources, enhanced data protection, and accelerated application modernization. This solution is validated to deliver predictable outcomes and eliminate deployment risks.

Developer Operations: This Pure Validated Design eliminates deployment challenges by providing design and deployment guidance to simplify Day 1 operations, accelerating Kubernetes capabilities with persistent storage into the vSphere environment. You will quickly learn about creating Namespaces, Persistent Storage, and delivering Kubernetes services to developers while optimizing resources.

Technology Overview

VMware vSphere with Tanzu and FlashArray is a Kubernetes, developer-ready solution running vSphere 7 and high-performance storage. VMware and Pure Storage delivers a unified experience to manage, support, and deploy Kubernetes workloads. Figure 2 provides an architecture framework for the solution.

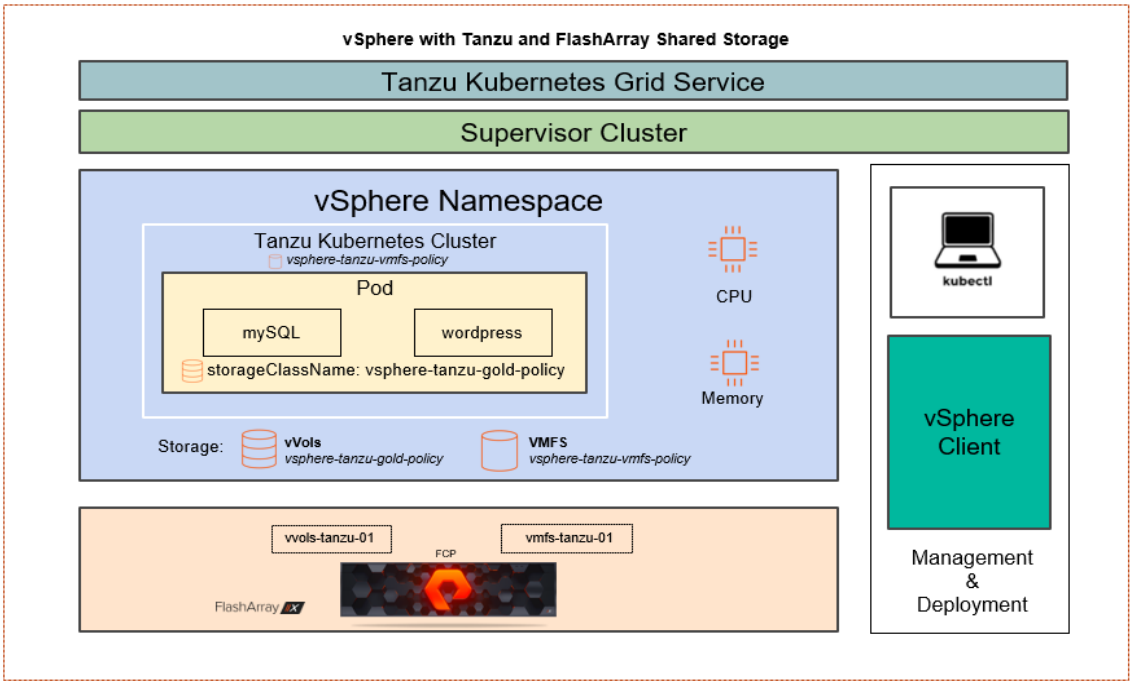

Figure 2. VMware vSphere with Tanzu and vVols with FlashArray framework

VMware vSphere with Tanzu Basic

Accelerate Kubernetes on-premises deployments and evaluations by running Tanzu Basic Edition. vSphere 7.0 Update 1 introduces support with bring-your-own virtual networking, network load-balancing solution, and shared storage.

Note: vSphere with Tanzu Basic does NOT require vSAN, VMware Cloud Foundation, or NSX-T.

vSphere with Tanzu versions:

- Tanzu Basic (Includes: Antrea/Calico CNI): Run and manage Kubernetes
- Tanzu Standard: Manage a Kubernetes multi-cloud environment
- Tanzu Advanced: Custom applications on Kubernetes
- Tanzu Enterprise: Developer velocity with automation

Note: Further reading about Tanzu Versions.

VMware Tanzu Kubernetes Grid Service: VMware Tanzu Kubernetes Grid Service delivers a consistent, upstream-compatible, Kubernetes environment for developer workloads. The Management Cluster and built-in Supervisor Cluster in vSphere 7 accelerate the construct of running Kubernetes on VMware with Tanzu Kubernetes Grid service. Manage and operate the entire lifecycle of Cloud Native Computing Foundation (CNCF) conformant Tanzu Kubernetes clusters with Cluster

API. VMware maintains the Kubernetes binaries, including testing and support to eliminate complexities and versioning challenges.

vSphere Distributed Switch: VMware vCenter Server enables virtual networking of Virtual Machines (VMs) with vSphere Distributed Switch (VDS) across multiple hosts and the entire data center.

vSphere Virtual Volumes (vVols): vVols delivers flexible storage with granular controls and storage policy based management. FlashArray simplifies provisioning of high-performance datastores and persistent storage to vSphere with Tanzu.

Content Library: VMware provides a Content Library to quickly deploy workloads and increase operational efficiency.

HAProxy: VMware's HAProxy virtual appliance delivers Load Balancing for Kubernetes Services and Applications.

Pure Storage Pure1® VM Analytics

With Pure1 VM Analytics, you can see an end-to-end view for performance and capacity across your workloads, right down to each Kubernetes node and the storage volume(s). The VM Analytics service is used to enhance other tools such as vRealize Operations when troubleshooting. VM Analytics provides detailed performance data from Pure1 and correlates that with performance data from vCenter Server, allowing you to see at a glance where problems might be. See Figure 3.

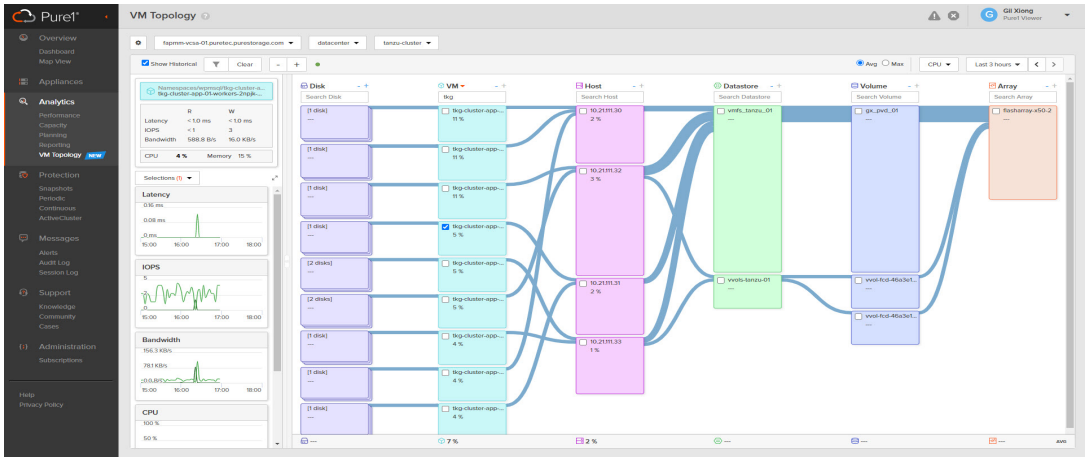

Figure 3. Pure1 VM Topology - vSphere with Tanzu and FlashArray environment

Pure Storage FlashArray //X

Pure Storage FlashArray//X, the leading 100% NVMe all-flash storage platform with deep VMware integration eliminates operational complexities for VMware and container environments. FlashArray's simplicity places it as the number one platform of VMware vVols, delivering innovation with VMware solutions. FlashArray is available as a performance optimized array (FlashArray//X) and as a capacity-optimized array (FlashArray//C) to optimize for application performance and economic requirements.

Chassis: A modular, stateless architecture engineered with Software-Defined freedom eliminates downtime and data migration.

Figure 4. Pure FlashArray Chassis Front View

DirectFlash: FlashArray moves beyond legacy SSD architectures that are architected to make flash pretend to be a hard disk. DirectFlash within Purity speaks directly to raw NAND with a super-efficient NVMe protocol. Each DirectFlash module can be hot plugged into the system to add capacity and performance (Figure 5).

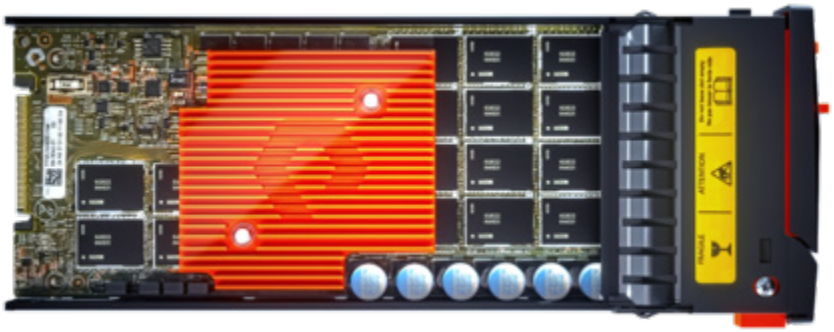

Figure 5. FlashArray - DIrectFlash module

Purity: The Pure Storage Purity operating environment is the software-defined engine of Pure Storage FlashArray. Purity is the driver that enables Pure FlashArray products, powering FlashArray//X to deliver comprehensive data services for your performance-sensitive data-center applications, and FlashArray//C for your capacity-oriented applications. Purity is built to scale non-disruptively and has broad API integration to deliver operational and business agility. Business continuity and disaster recovery are built-in, as well as industry-leading data reduction, efficient snapshots, and cloud mobility. For simplicity, the FlashArray's service for vVols is a native service running in Purity to eliminate management complexities.

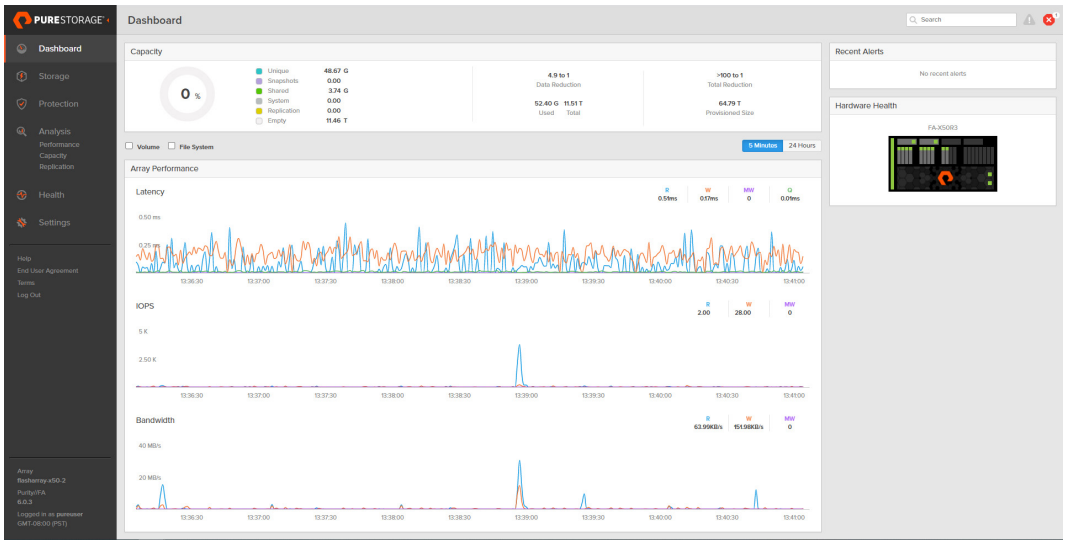

Figure 6. Purity//FA dashboard

Planning, Design, and Prework

This section of the document will cover detailed requirements and preparation before deployment. Prework of the vSphere Hosts and Environment must be carefully reviewed with preplanning the network requirements for a successful deployment.

vSphere Hosts: The solution is based on four ESXi 7.0.1 hosts with the following individual specifications:

- Dual Socket, 12 Cores per Socket
- 512GB of physical memory
- Two Fibre Channel HBA
- Two 10 GbE Single Port Network Controllers
  - vmnic0
  - vmnic1

The required physical server type is commodity x86-64 architecture with minimum specifications based on VMware Hardware Requirements. vSphere 7.0 ESXi hardware requirements and the VMware compatibility guide are located on VMware's website.

vSphere Environment

This document will showcase the vCenter Server User Interface to accomplish the required tasks to manage and operate vSphere with Tanzu. The vSphere environment for this solution must meet all the following parameters before deploying the HAProxy Load Balancer and enabling Workload Management:

- vSphere Cluster with three (3) ESXi Hosts (minimum)
  - Minimum: Version 7.0.1
  - Minimum of one (1) FlashArray//X VMFS Datastore
  - Best practice: Hostnames should use all lowercase to eliminate deployment issues

- vCenter Server Appliance
  - Minimum: Version 7.0.1
  - Running on FlashArray VMFS Datastore
- HA and DRS must be enabled for *tanzu-cluster*
  - DRS should be set to Fully Automated
- vSphere Distributed Switch
  - Minimum: Version 7
  - All ESXi hosts connected
  - Configure two port groups: "management" and "workload"
- vMotion Network
  - Best Practice: Dedicated VLAN for production workloads
- NTP configuration must be the same across ESXi hosts, vCenter Server, and HAProxy.

Note: vSphere 7.0.1 includes a 60-Day trial of vSphere with Tanzu Basic.

Network Requirements

This deployment is based on *two routable VLAN networks* and a single vSphere Distributed Switch with two port groups named "Management" and "Workload." The Management network supports ESXi hosts, vCenter Server, FlashArray//X, and operational management. The Workload network supports the Supervisor Cluster and Kubernetes workloads which includes access to HAProxy services. (See Figure 7.)

- Management network - vmnic0 [Example: VLAN111 w/ 192.168.111.0/24]
  - Must be routable to ESXi hosts, vCenter Server, the Supervisor Cluster, and HAProxy
  - DNS query enabled
  - Must have internet connectivity
- Workload network - vmnic1 [Example: VLAN114 w/ 192.168.114.0/24]
  - DNS query enabled
- vSphere distributed switch
  - Management port group configured to appropriate VLAN ID
  - Workload port group configured to appropriate VLAN ID
  - HAProxy LoadBalancer requires the assignment of static IP Addresses. This deployment cannot be based on DHCP.
- Management network
  - HAProxy Management - IP Address: Example 192.168.111.159/24
    - Default gateway must be on this network
    - Dataplane management port 5556 (default), which is used for Step 5 of enabling Workload Management. Example: 192.168.111.159:5556
  - Supervisor Control Plane - Five (5) IP Addresses: Example 192.168.111.160 - 192.168.111.164

- The first IP address will be used to enable Workload Management, Step 6 - Starting IP Address.

- Workload network
  - Workload IP Address: Example 192.168.114.7/24 (must be outside of LB IP Range)
    - This IP address will be used and entered in CIDR format during the HAProxy installation, Step 9 Customize Template section 2.5.
  - Load Balancer IP Ranges (CIDR): Example 192.168.114.32/27 (These are used for the TKC)
    - 10.21.114.32/27 CIDR Range = 192.168.114.32 - 192.168.114.63
    - HAProxy Installation Wizard Step 9 Section 3.1
  - Workload IP Range - Virtual Machines: Example 192.168.114.96-10.21.114.127 (must be outside of LB IP Range)
    - This will be used to enable Workload Management, adding workload network Step 7.
- VMware provides extensive networking overview and official documentation.

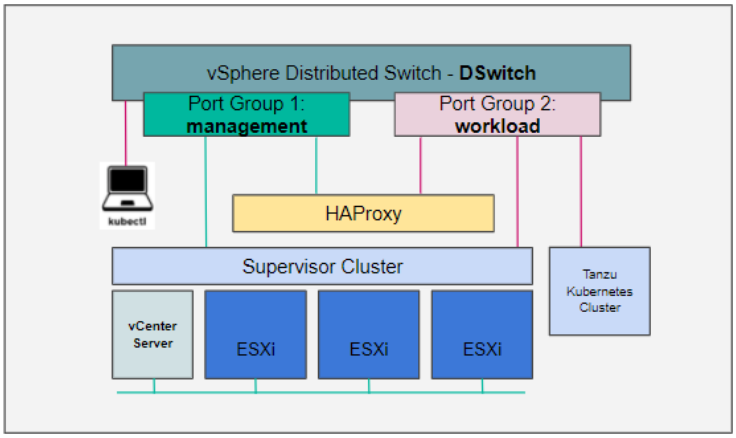

Figure 7. Virtual network

Network Worksheet

HAProxy	Management	Workload	Notes
Management IP	_____._____._____._____		root
SControlPlane VM-1	_____._____._____._____		Start
SControlPlane VM-2	_____._____._____._____		
SControlPlane VM-3	_____._____._____._____		
SControlPlane VM-4	_____._____._____._____		
SControlPlane VM-5	_____._____._____._____		End
Workload IP		_____._____._____._____	
Load Balancer IP Range of 30 IPs (min)		_____._____._____._____/	CIDR
Virtual Machines Range of 30 IPs (min)		_____._____._____._____ _____._____._____._____	Out of LB range
NTP	_____._____._____._____		
DNS	_____._____._____._____		

CIDR conversion tool: https://www.ipaddressguide.com/cidr

FlashArray

Deployment tasks in this documentation are based on the vSphere Client and the Pure Storage Plugin for the vSphere Client.

- The FlashArray is authenticated with the Pure Storage Plugin for the vSphere Client (referred to in the rest of this document as the vSphere Plugin) for ease of use.
  - There are multiple methods to install the vSphere Plugin.
    - PowerShell, vRealize Orchestrator, and the FlashArray Web UI as documented.
- Required: ESXi hosts' Personality must be set to "ESXi."
- Purity Version 5.3.10 or later.

FlashArray vVols

The following list is a set of best practices and tasks that must be completed before the creation and mounting of the vVols Datastore via vCenter.

- Required: Network port 8084 must be open and accessible from vCenter Server and ESXi hosts to the FlashArray.
- Best Practice: Run vCenter Server on VMFS.
- Best Practice: Create a dedicated local FlashArray User for VASA registration.
  - Prevent authentication issues in the event Active Directory/LDAP service is unreachable or the case of a *user account* of the storage provider being deleted/removed.
- Best practices and quick guidance points for VMware vVols with Pure.

VMware HAProxy

The VMware HAProxy Appliance must be downloaded and uploaded to the vCenter Content Library.

- Download the VMware Proxy OVA on Github.
  - Requires an FQDN and a Static IP address for direct management on the management network
  - Requires a Static IP address on the *workload* network
  - HAProxy v0.1.10

Managing the Content Library

The vSphere Content Library consists of a local HAProxy library containing the OVA and a subscription Tanzu Kubernetes Grid library for deployment. The HAProxy OVA must be uploaded on a Datastore accessible by tanzu-cluster.

- Tanzu Kubernetes Grid (subscription content library)
- HAProxy (Local content library)
  - You must Import Library Item: haproxy-v0.1.10.ova

To create vSphere Content Libraries, navigate to the Menu drop-down from vSphere Client and select Content Libraries (Figure 8).

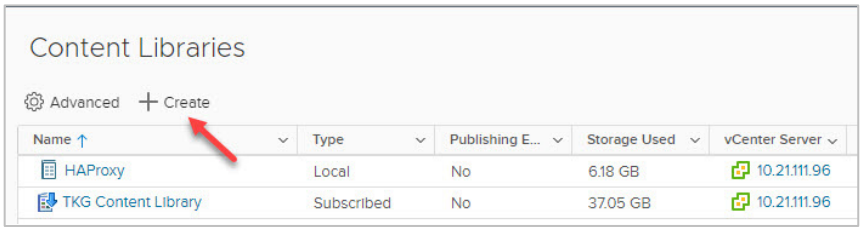

Figure 8. Content library

Deployment and Application Validation

This section describes the deployment of vSphere Workload Management, Namespaces, Tanzu Kubernetes Clusters, and persistent volume claims with FlashArray vVols Datastore. It provides guidelines for installing and configuring HAProxy, the Workload Management environment, and building a web-based application on MySQL in a Tanzu Kubernetes Cluster.

Solution validation is an operational web-application with MySQL using the Kubernetes persistent volume claim to FlashArray vVols as persistent volumes.

VMware vSphere with Tanzu is installed on dedicated physical industry-standard x86 servers and Pure FlashArray//X hardware. Pure Storage FlashArray is a high-performance platform that supports mixed workloads for shared storage efficiency. It is common to use FlashArray for hosting VMware vSphere workloads with both VMFS and vVols.

vSphere Client VASA registration of FlashArray Storage Providers

VASA registration can be accomplished in multiple ways, i.e., PowerShell, FlashArray UI, or vRealize Orchestrator. The example here uses the vSphere Plugin. Once the FlashArray Providers are registered, creating and mounting a vVols Datastore is a simple and seamless process.

1. With an existing and registered FlashArray via Plugin, navigate to the plugin.

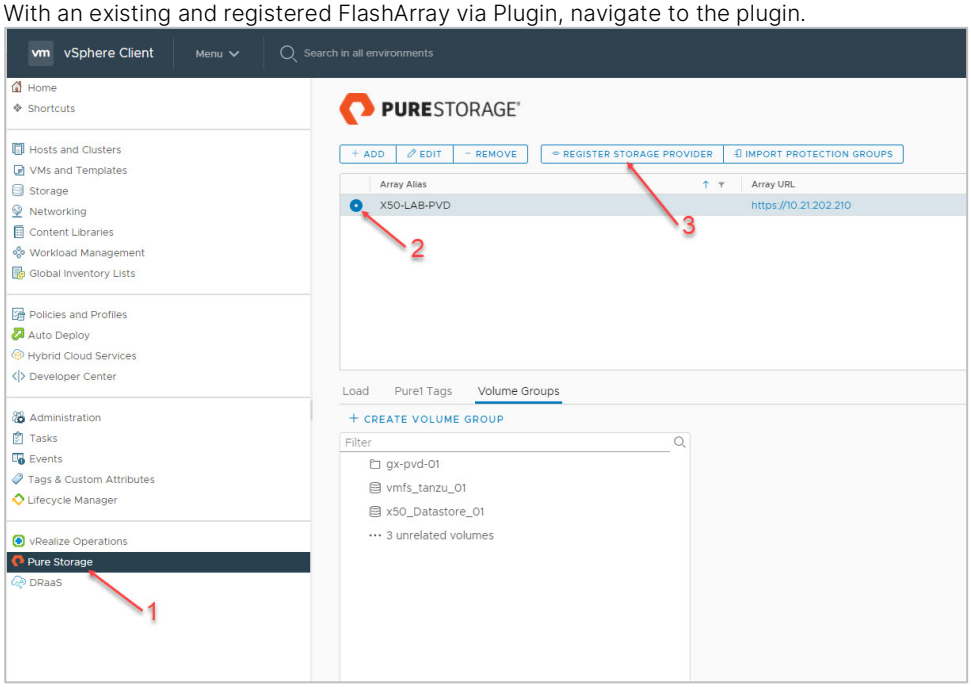

2. Once the Storage Provider is successfully registered, navigate to the Storage Providers configuration page to confirm both providers are online and healthy.

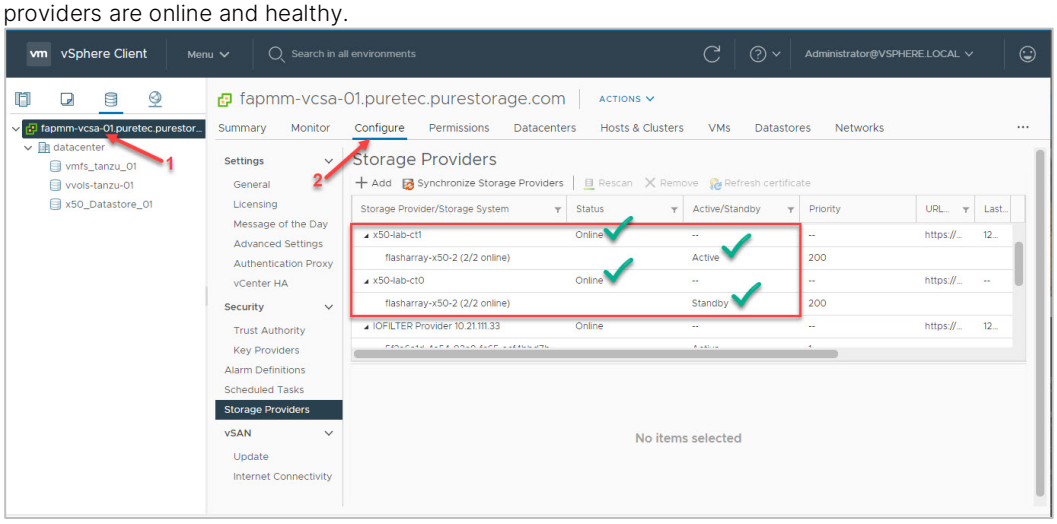

Creating a FlashArray Host Group

Creating a FlashArray host group can be accomplished in multiple ways, i.e., PowerShell, FlashArray UI, or vRealize Orchestrator. The example here uses the vSphere Plugin to streamline the creation of a host group to the cluster named *tanzu-cluster*.

1. Navigate to the tanzu-cluster and right-click to bring up the Actions menu, hover down to Pure Storage and click **Add Host Group** wizard.

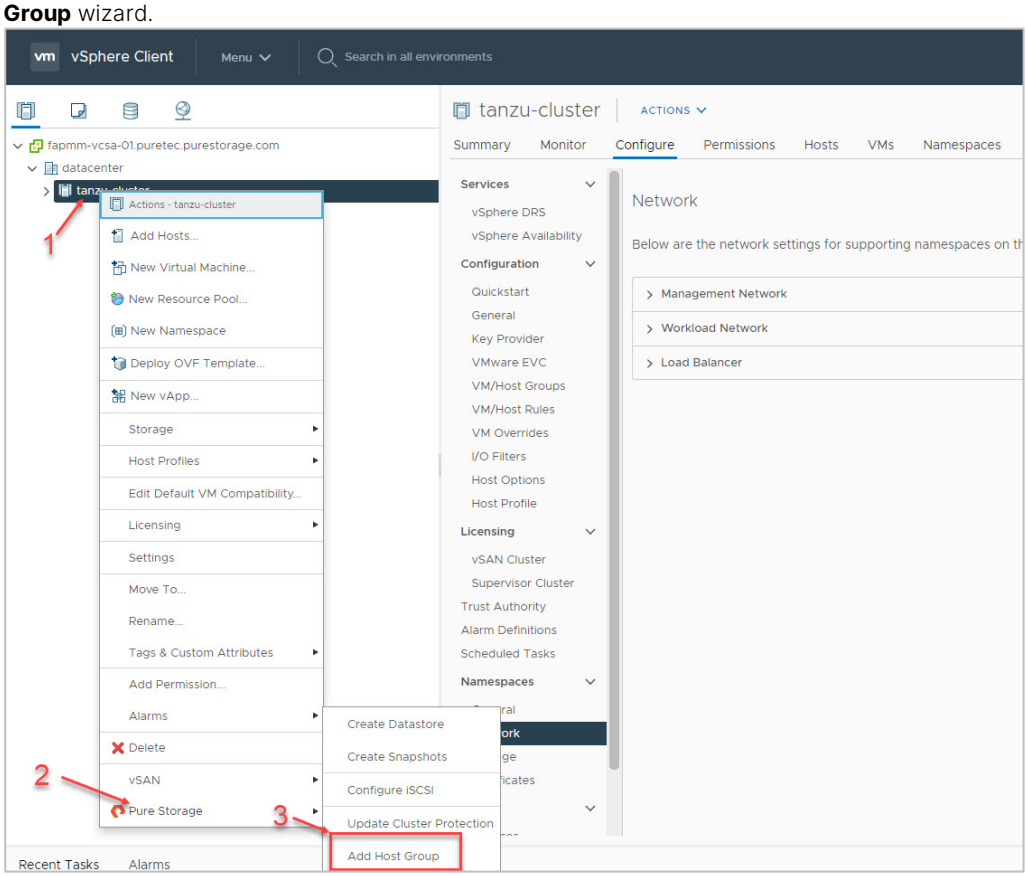

3. Select the FlashArray and the appropriate protocol that you are using for the deployment. Review and confirm the available hosts can click **Create**. Note: iSCSI deployments will default and Configure iSCSI Initiators on Hosts.

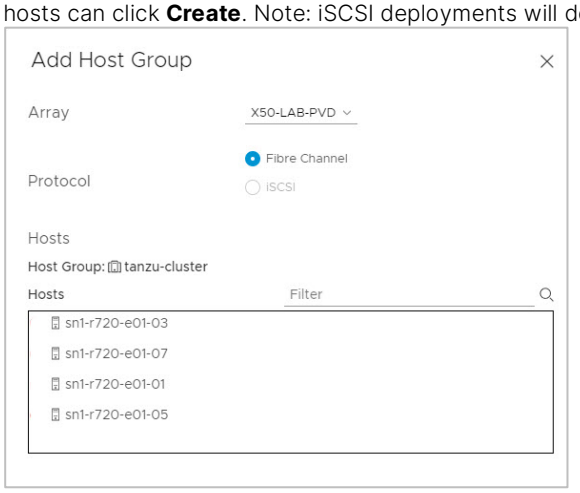

Creating and Mounting the vVols Datastore

The example here will use the vSphere Plugin to streamline the provisioning process of the vVols datastore to the *tanzu-cluster*.

1. Navigate to the tanzu-cluster and right-click to bring up the Actions menu, hover down to Pure Storage and click **Create Datastore** wizard.

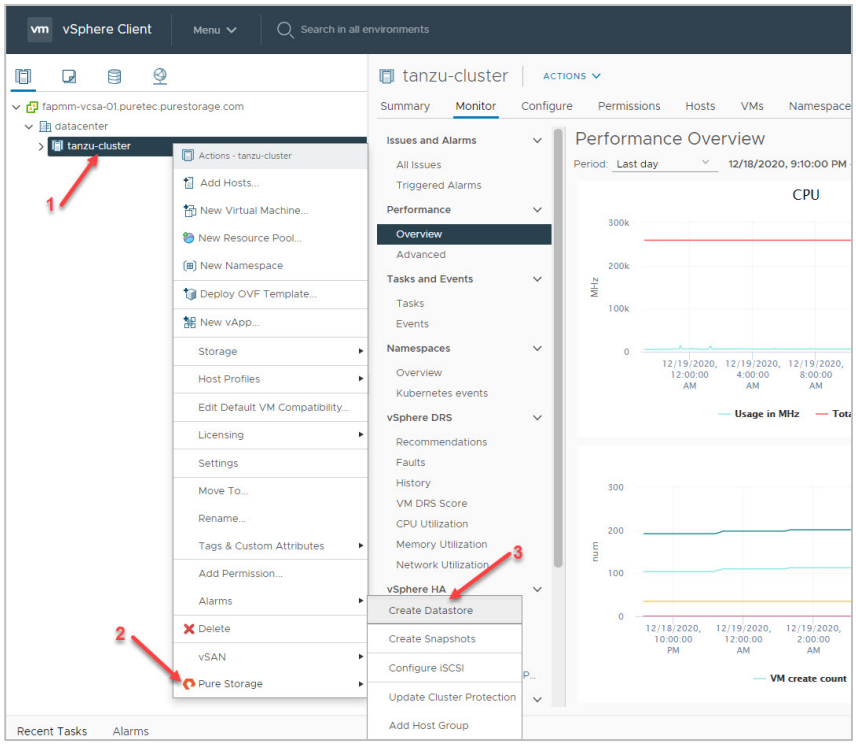

4. Select “vVol” as the Type of Datastore and proceed through the wizard. The following screens will ask for a Datastore Name, a vVol datastore is at a defaulted size of 8PB and can be modified as required. Select tanzu-cluster as the Compute Resource.

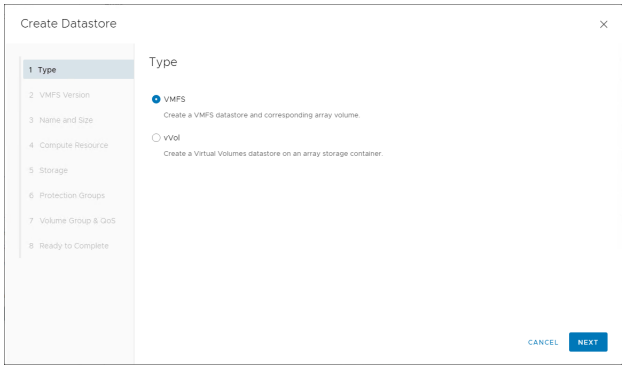

5. Select the FlashArray for the project.

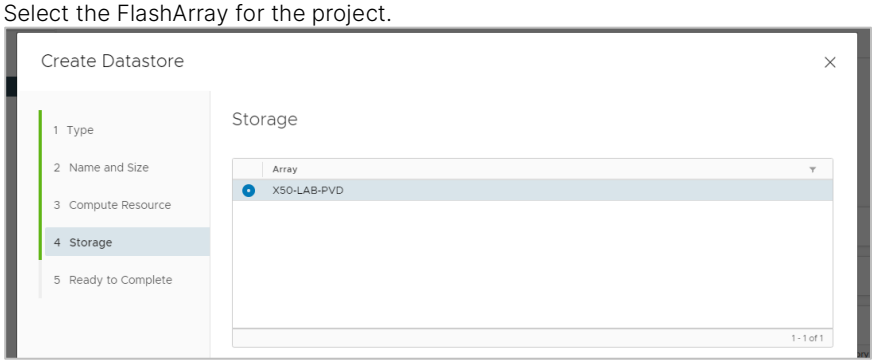

6. Review your entries, click Finish to complete the task.

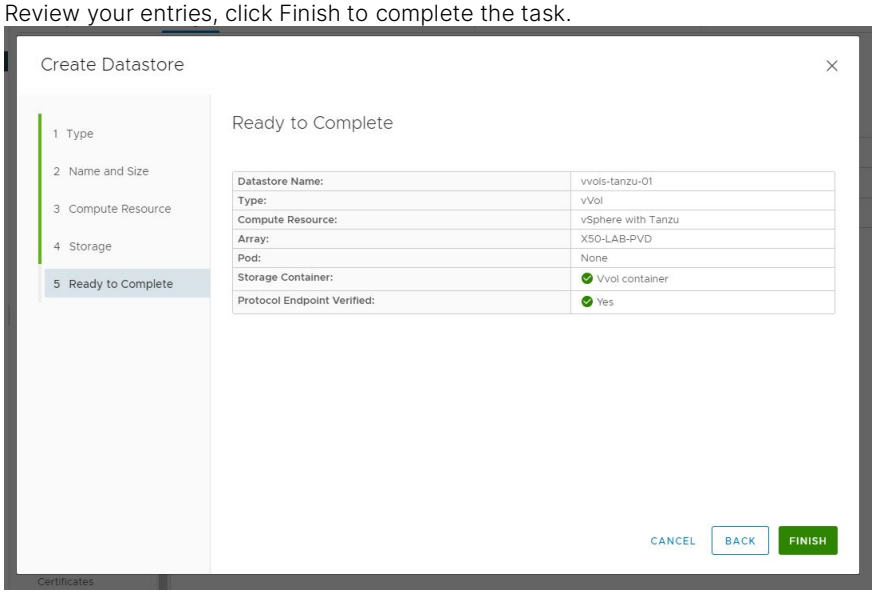

7. Navigate to the Datastore page and select the vVol Datastore for a summary.

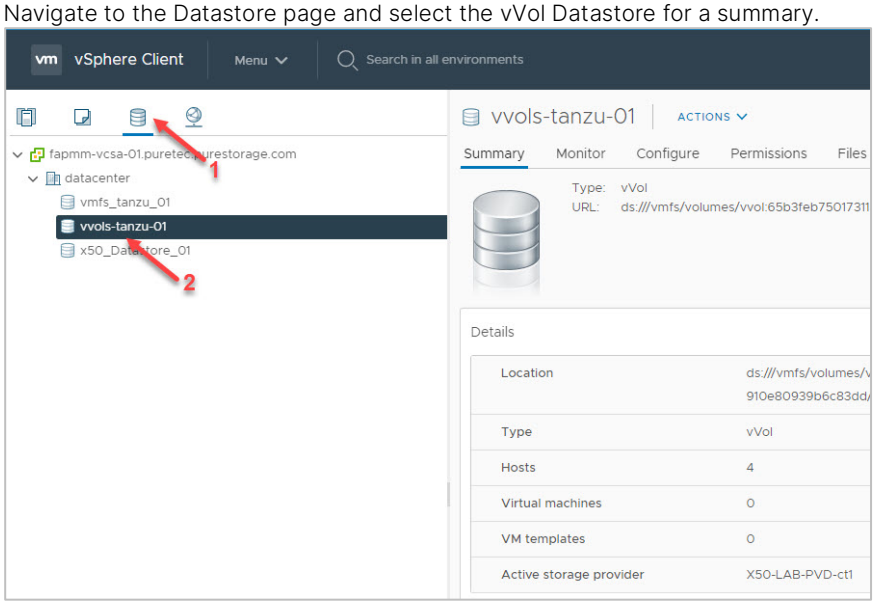

VMware vSphere Client Tasks

Creating Storage Policies

This deployment uses two storage policies. The first is a vVols policy (Example Name: vsphere-tanzu-gold-policy) for persistent storage and the second is for VMFS policy (Example Name: vsphere-tanzu-vmfs-policy) which will store VMs. These policies translate to *Storage Classes* in Kubernetes which will be inserted in the design files during the Tanzu Kubernetes Cluster, application, and database deployment.

The FlashArray vVols datastore policy will use VMware's native Storage Policy Based Management (SPBM) with a simple definition. A tag-based policy will be created for the FlashArray VMFS datastore.

Create a new Category Name with Default values. This deployment uses the name "tanzu-demo-tag-category" for ease of identification.

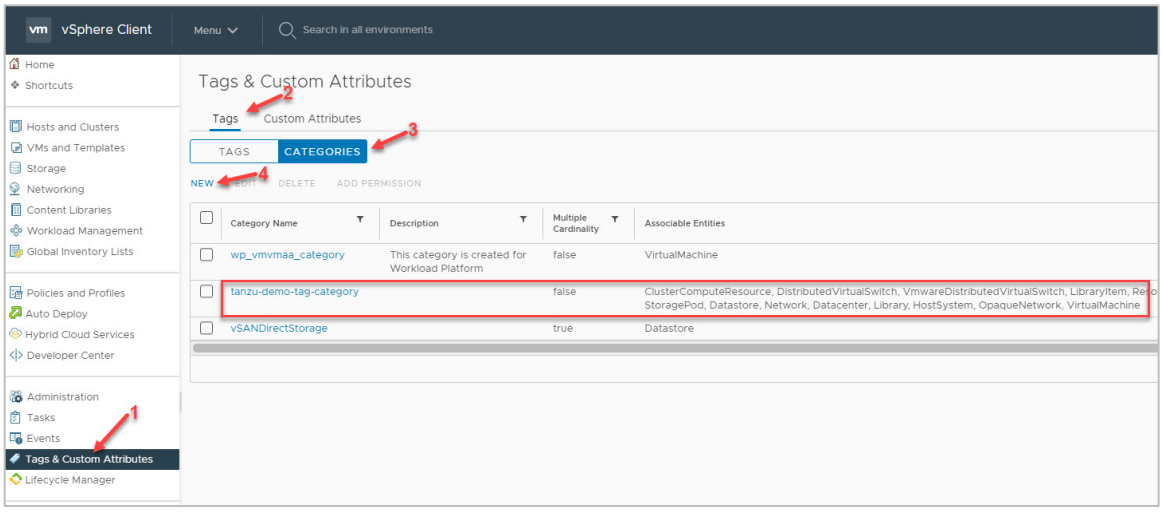

1. Additionally, create a new Tag to be assigned to the FlashArray VMFS datastore.

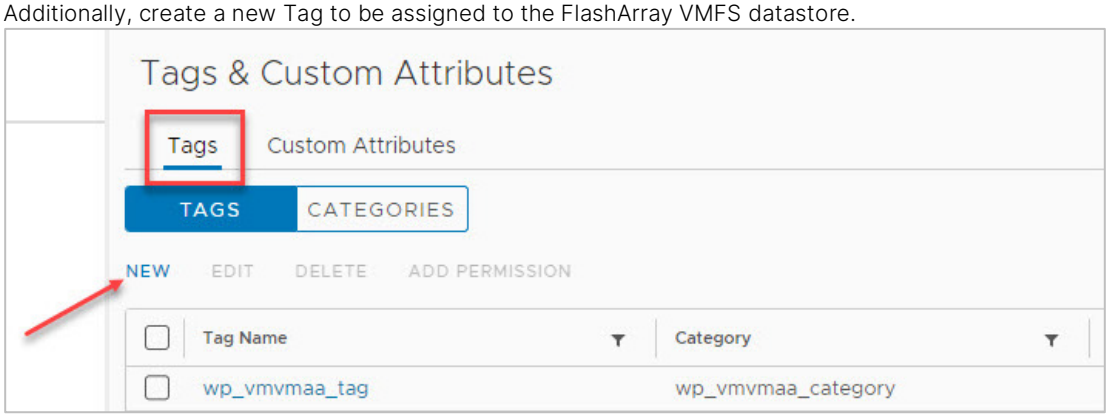

8. Navigate to the datastores and Assign the new Tag to the FlashArray **VMFS** Datastore.

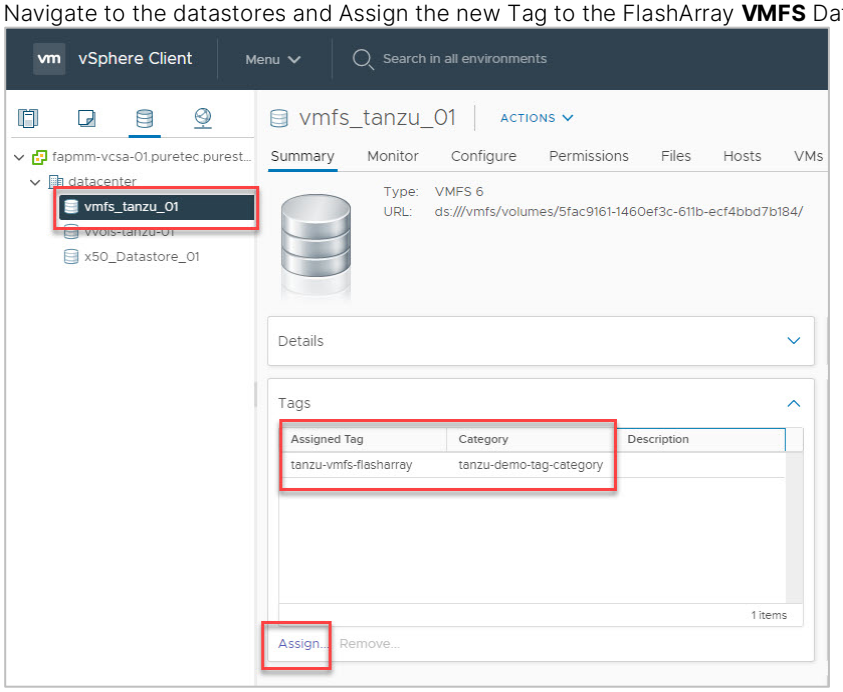

Create a Storage Policy for the FlashArray VMFS datastore.

1. Navigate to the vSphere Client Menu, click **Policies** and **Profiles**, select **VM Storage Policies**, and click **Create**.

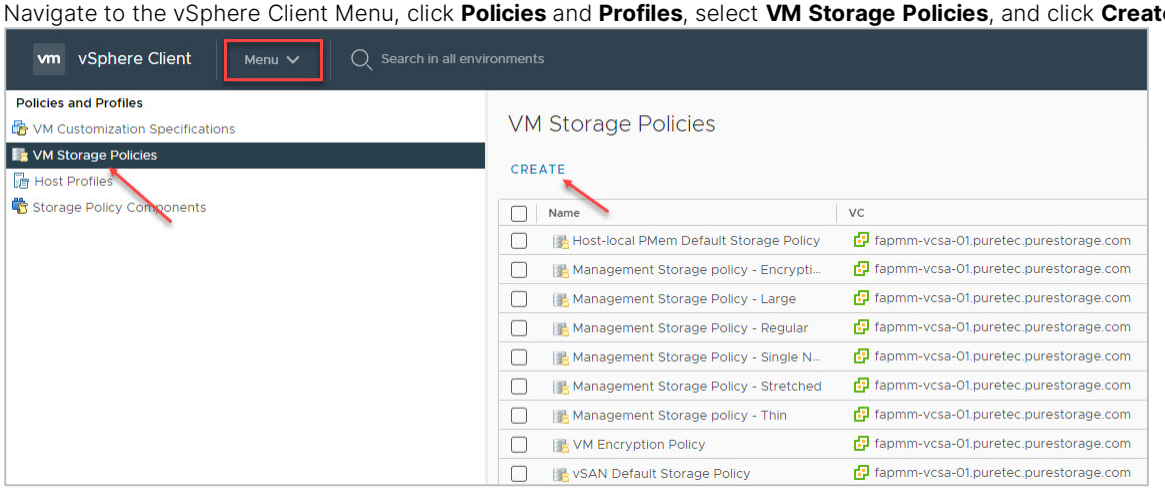

9. Assign a descriptive, add a description if preferred, and click **Next**. The policy structure will be based on datastore specific rules, check the box for Enable tag-based placement rules, and click Next.

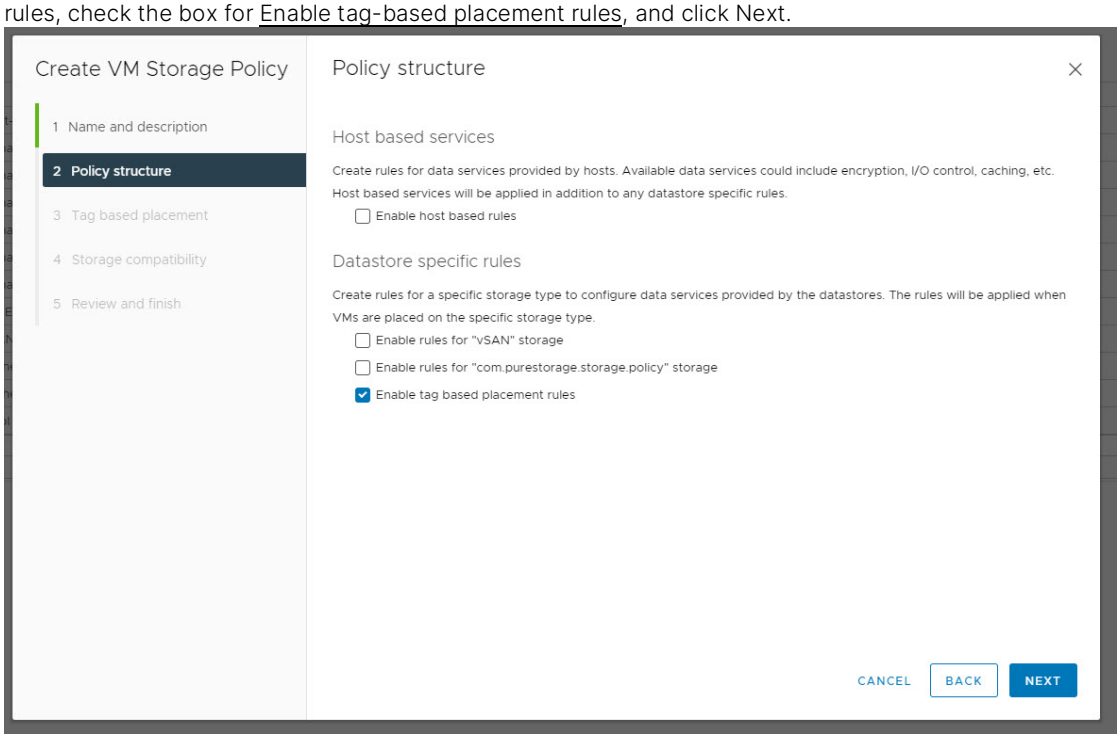

10. Select the tag category for the rule and add the appropriate tag with **browse tags**. Click **Next**.

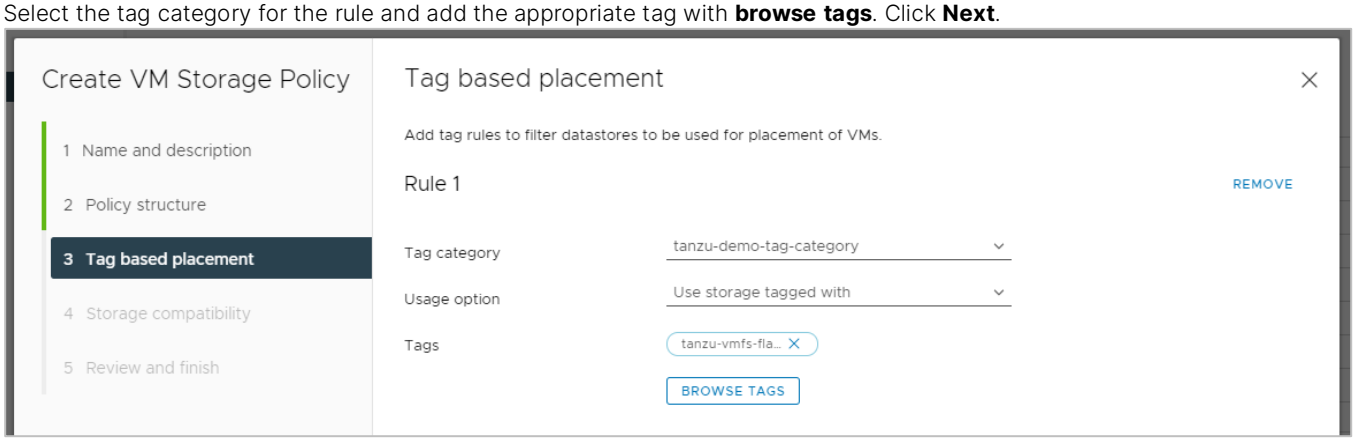

11. Review and select the compatible Datastore type and click next to review and finish.

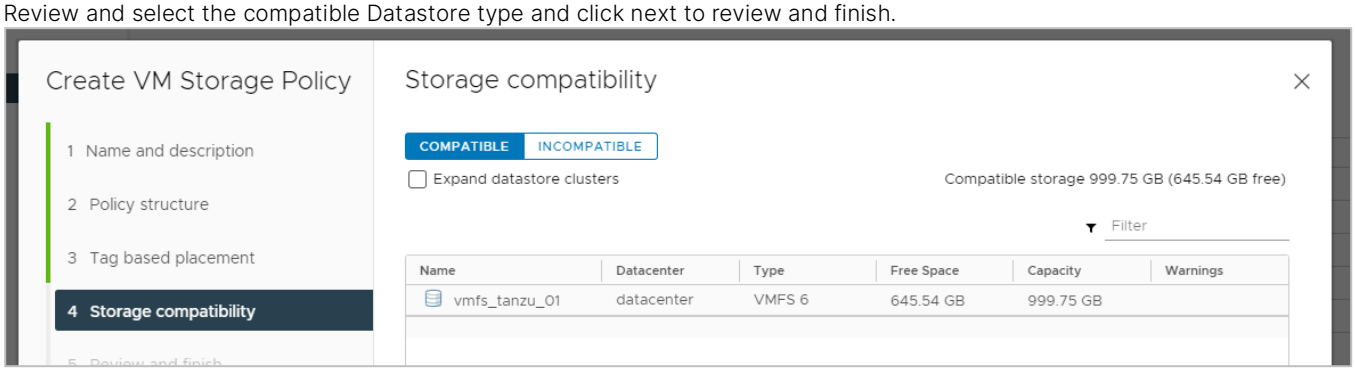

12. Review the list for the new VMFS policy.

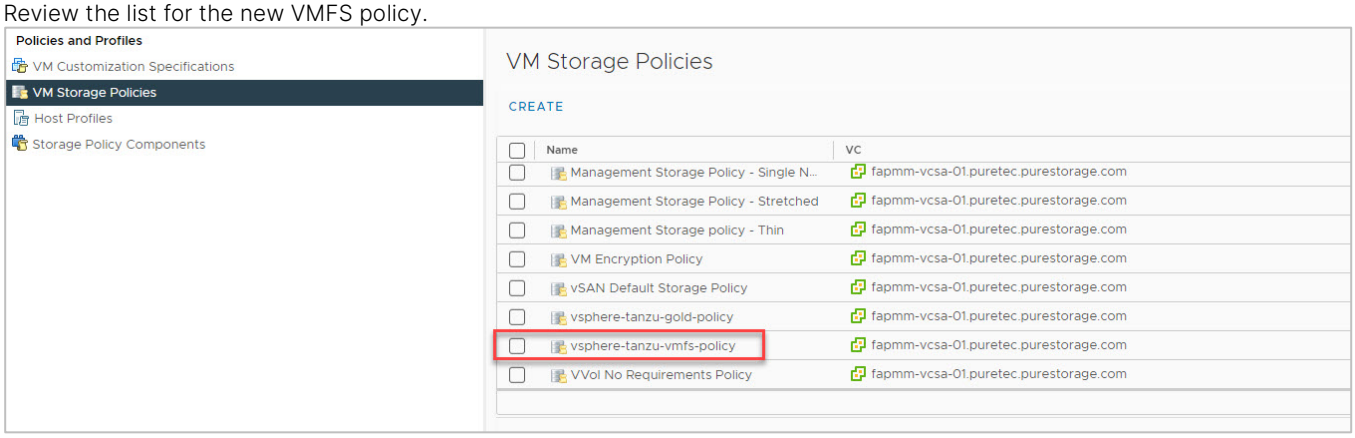

The next step is to create a storage policy for the FlashArray vVols Datastore, using Storage Policy-Based Management (SPBM) with a simple rule.

1. Navigate to the vSphere Client Menu, click **Policies and Profiles**, select **VM Storage Policies**, and click **Create**.

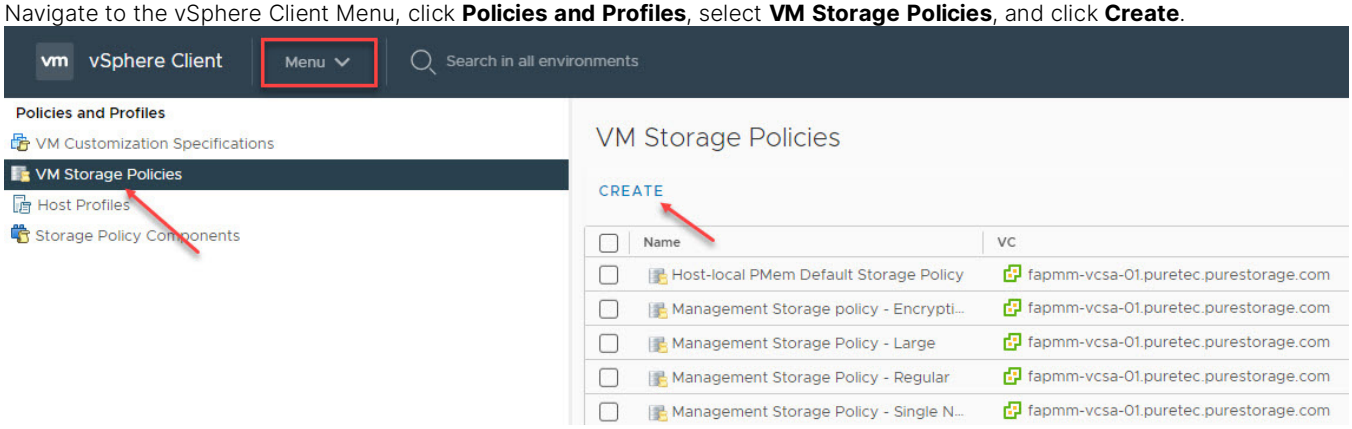

13. Use a descriptive Name with a Description (optionally) and click **Next**. The policy structure will be based on datastore specific rules, check the box for Enable rules for "com.purestorage.storage.policy" storage and click **Next**.

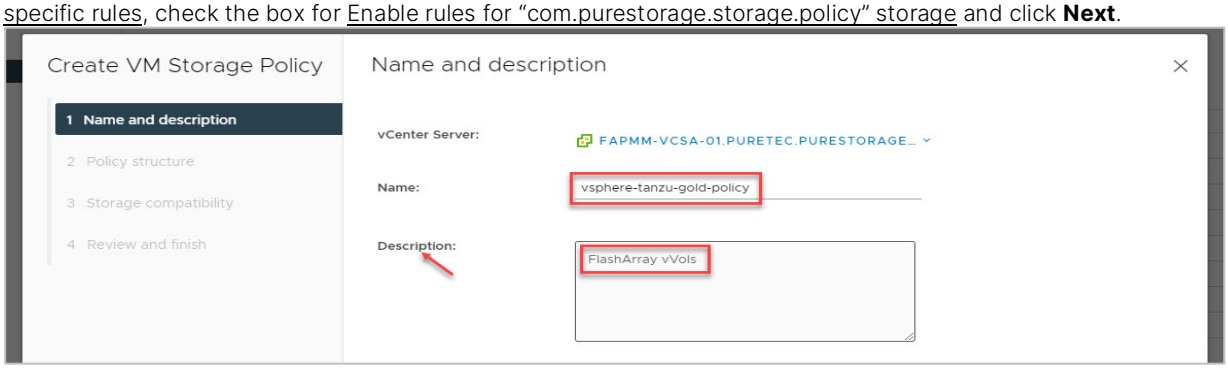

14. Using the **Add Rule** drop-down menu, select **Pure Storage FlashArray** rule with **Yes** value and click **Next**.

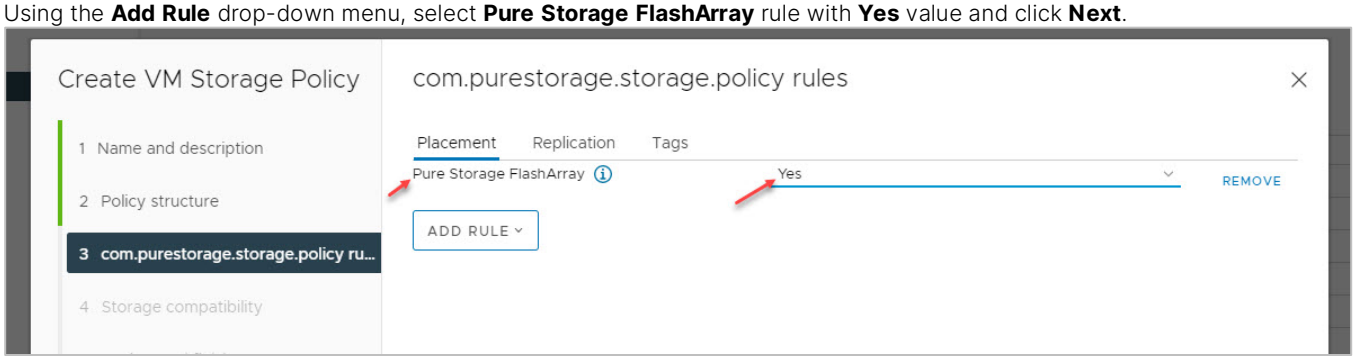

15. Select the FlashArray vVols datastore and click **Next** onto the Review screen and **Finish**.

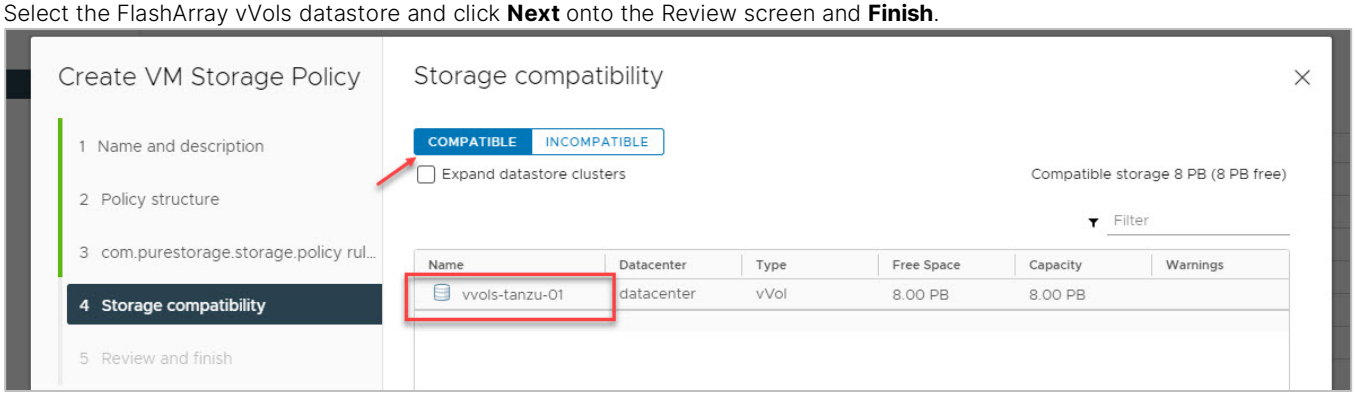

16. Review the list for the new vVols policy.

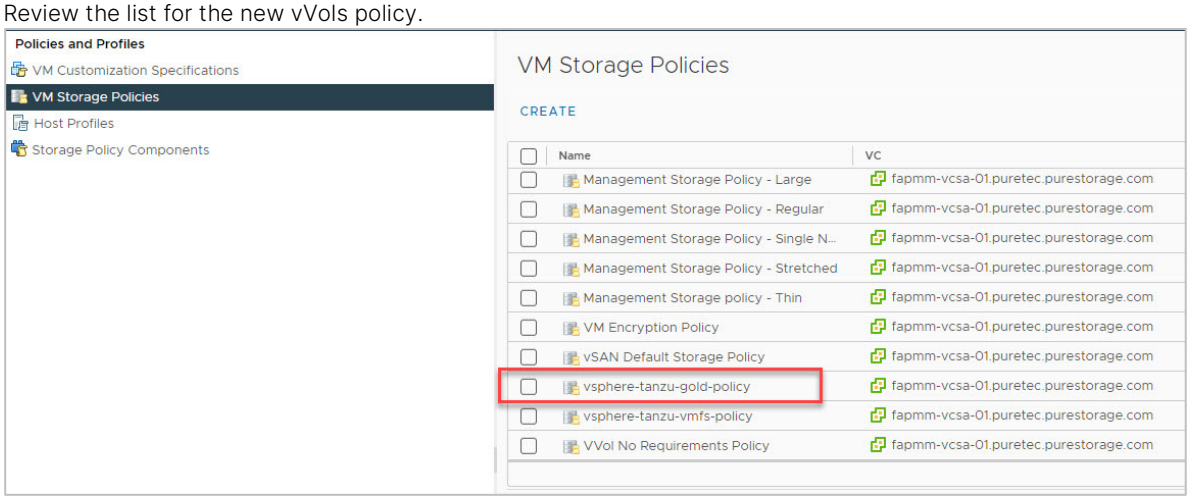

HAProxy

The deployment of the HAProxy Load Balancer is wizard-driven and requires attention with IP Address entries.

Navigate to the Content Libraries through the vCenter Server Menu drop-down and expand the HAProxy Library.

1. Select OVF & OVA Templates, select the haproxy-v0.1.10 OVA and use the Actions drop-down to create **New VM from This Template**.

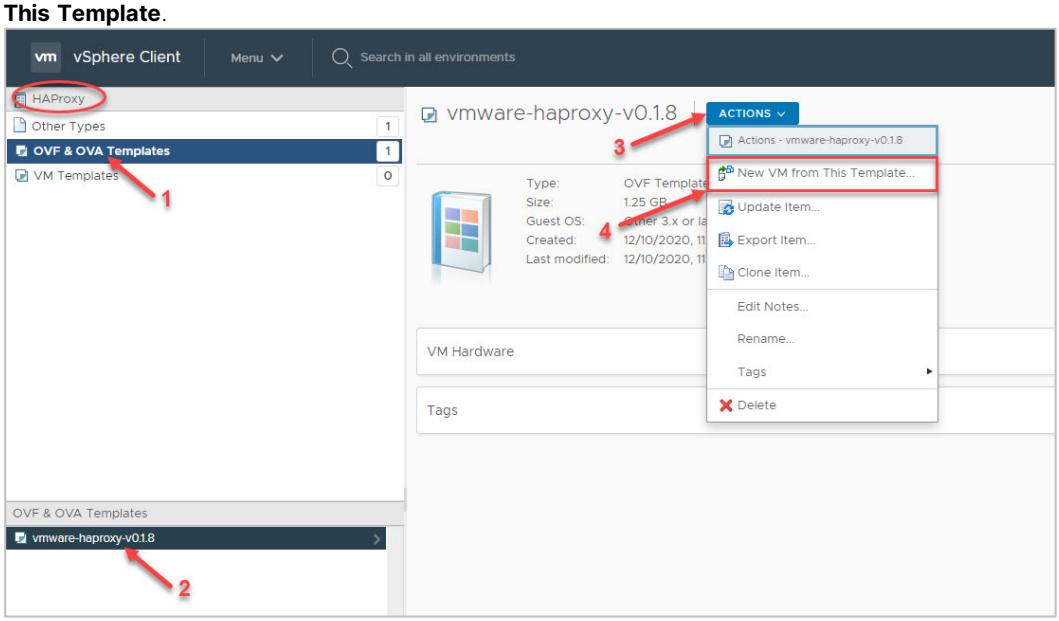

2. Enter a virtual machine name for your haproxy and select the location.
3. Select destination cluster for the project: tanzu-cluster
4. Review details
5. Accept the license agreements

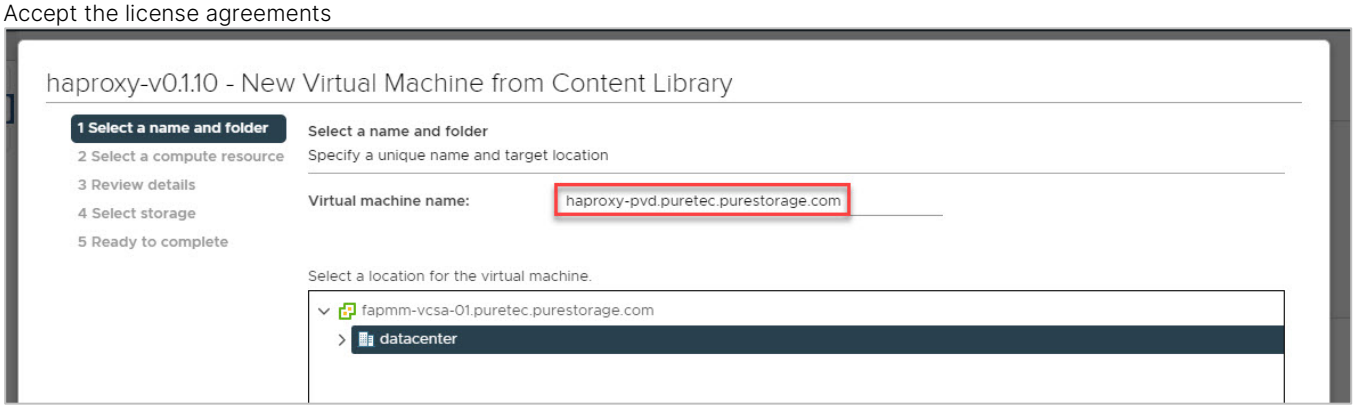

6. Select the **Default** deployment configuration.

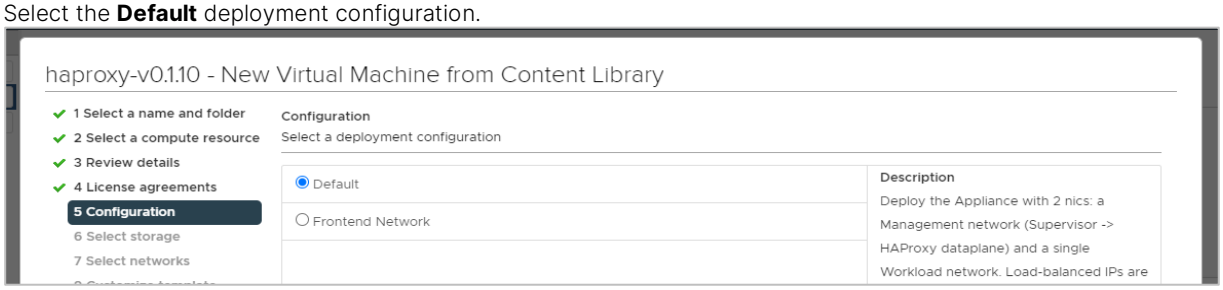

7. Select the **vmfs_tanzu_01** Datastore storage option for the HAProxy.

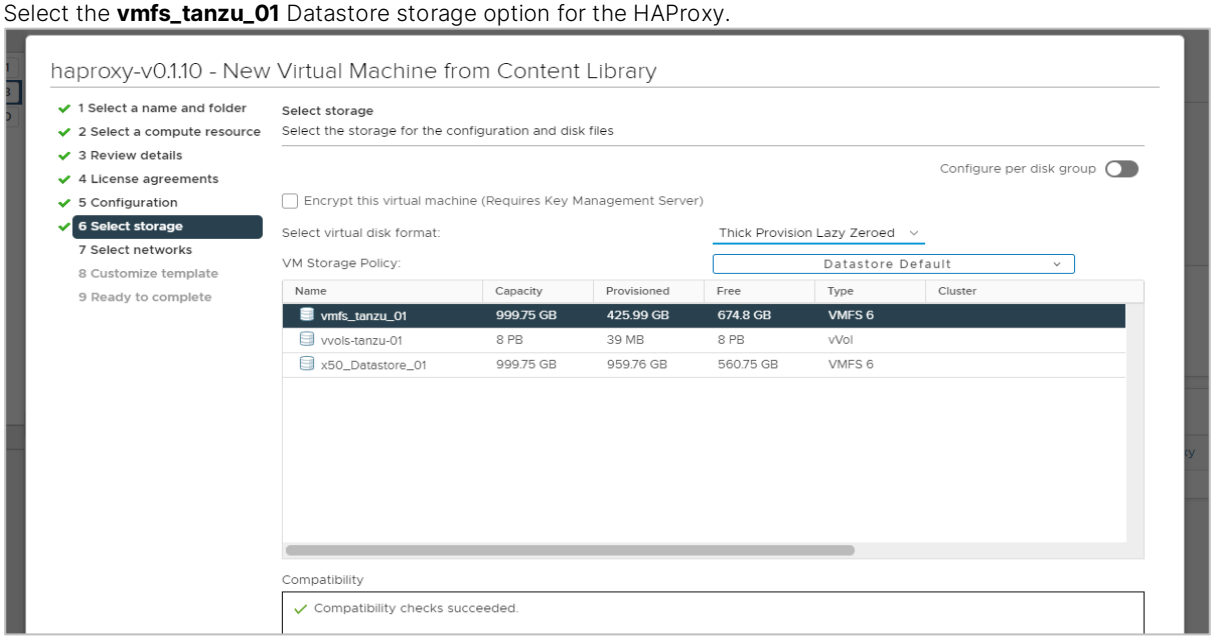

8. Select the appropriate Destination Network for each Source Network.

Note: Frontend Source Network does not require a change of Destination Network selection.

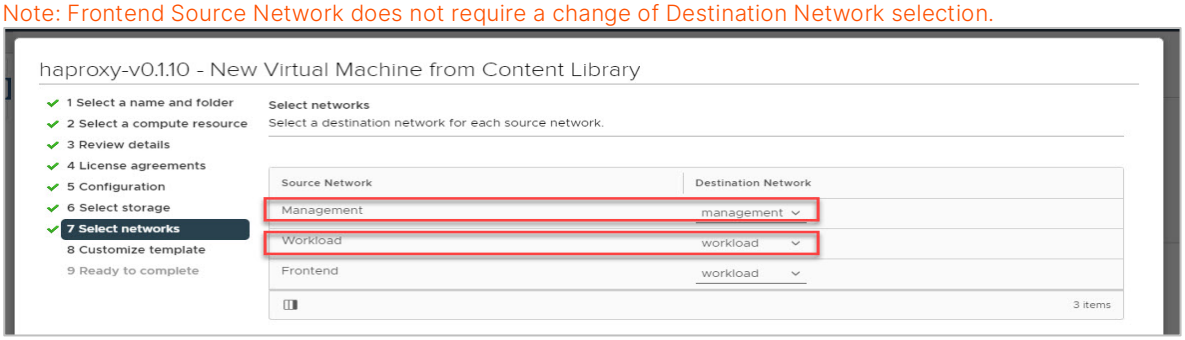

9. Appliance Configuration

1.1: Enter a password for the root user to manage the HAProxy via Management network. Take note of this password, it will be used with the *root* user to retrieve the Server Certificate Authority.

1.2: Permit Root Login: Default as Enabled.

1.3–1.4: Does not require entry.

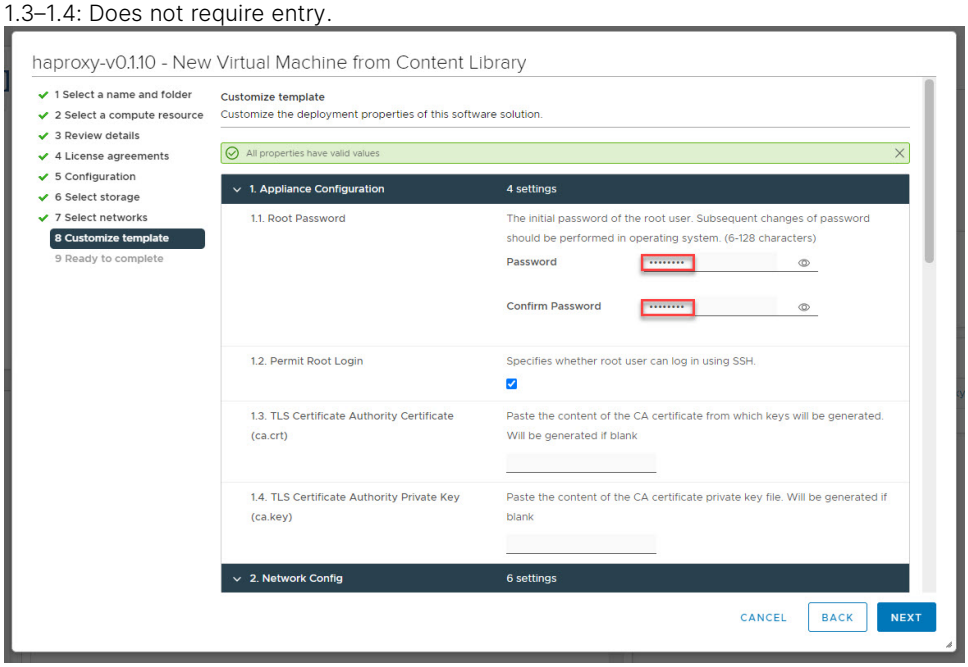

10. Network Configuration entries here must be entered and reviewed thoroughly to ensure accuracy based on the worksheet. It is important to note that entries for sections 2.3 and 2.5 require the proper CIDR format based on the IP Address.

Example: 10.21.111.159/24 (E.g., IP/subnet mask bits), /24 is a subnet mask of 255.255.255.0

2.1: It is best practice to use an FQDN.

2.3–2.5: The Management and Workload IP address entry must include the proper CIDR format.

2.4–2.6: The Gateway IP Address entry does not require CIDR format.

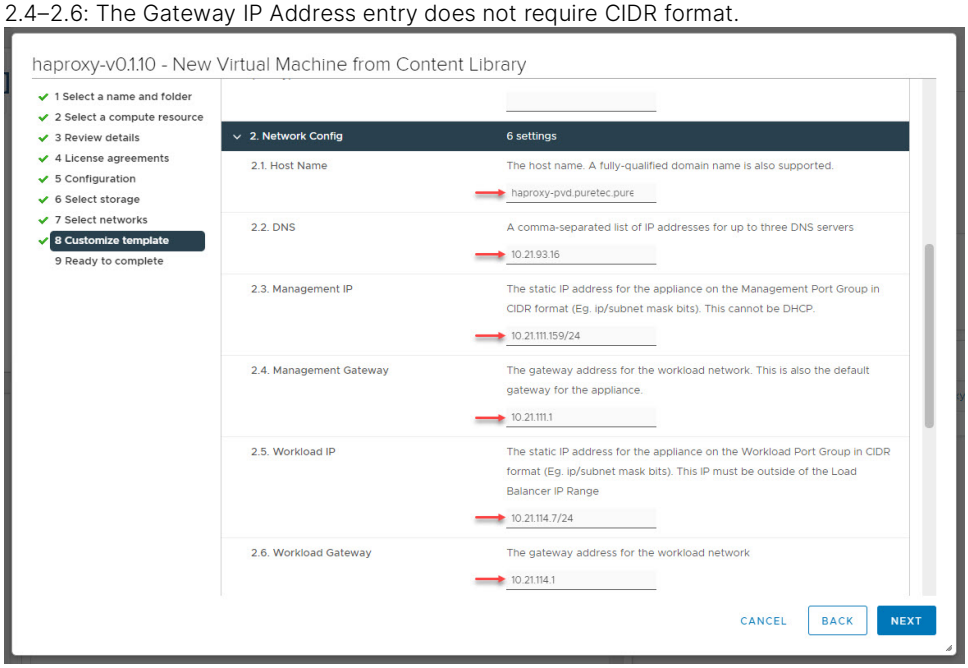

17. Load Balancing entries here must be entered and reviewed thoroughly to ensure accuracy based on the worksheet.

    a. Entry for section 3.1 Load Balancer IP Ranges requires the proper CIDR format based on the **first** IP Address of the Load Balancer IP Range from the worksheet in CIDR format.

    b. Example: 192.168.114.32/27 with /27 CIDR range of 192.168.114.32-195.168.114.63

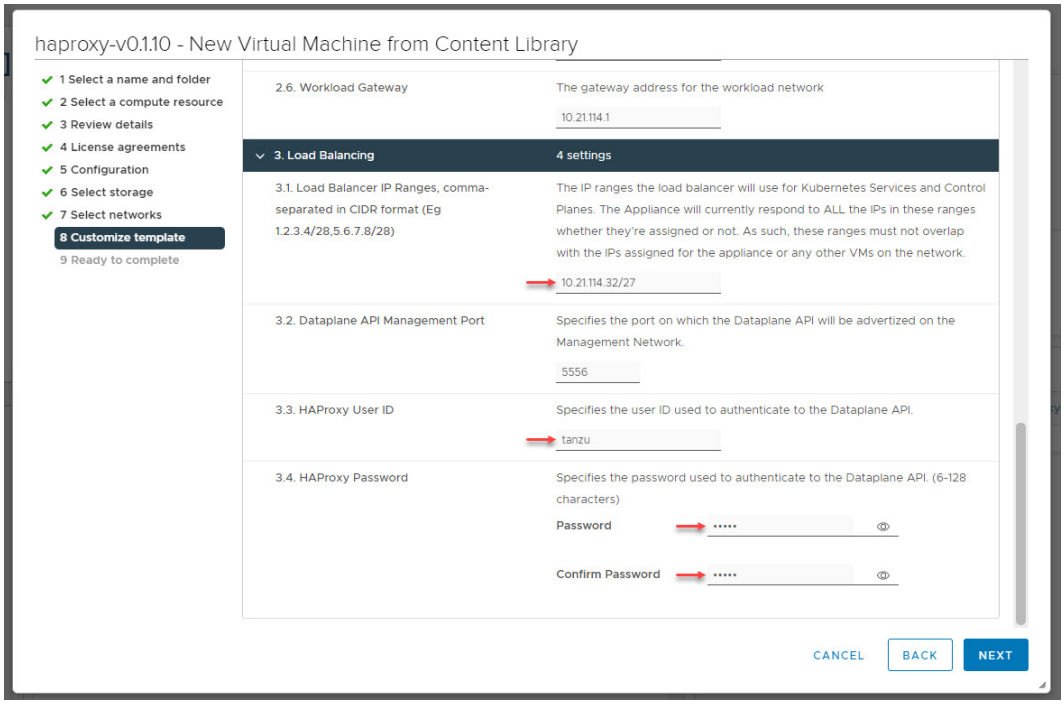

18. Review your entries and take note of the IP Address on Section 2.3, Management IP. It will be used to log into the HAProxy to confirm the installation and pull the certificate authority to enable Workload Management. Review thoroughly and click Finish to start the HAProxy build.

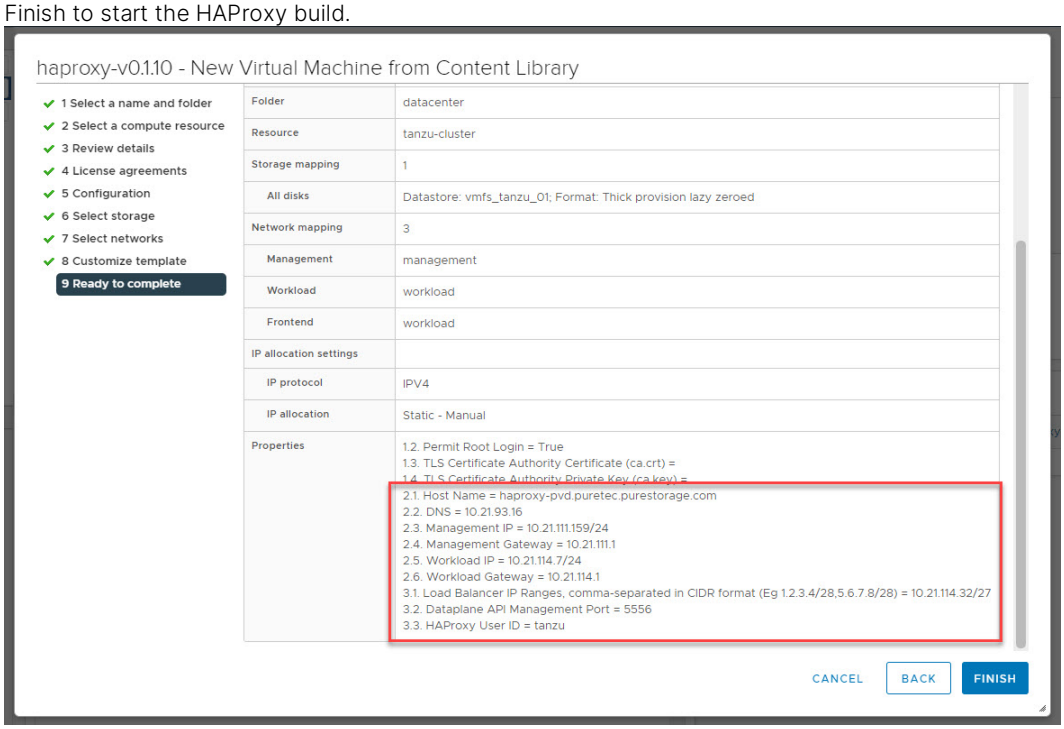

Check the status of the OVA build in vCenter, Power-On HAProxy once the build is complete.

You must be able to ping the Management IP, Workload IP, and the Load Balancer IP Range to enable Workload Management.

- If you can ping all the IP addresses, you have successfully installed the HAProxy and will be able to enable Workload Management.
  - Retrieve the Server Certificate Authority from the HAProxy.
    - SSH to the HAProxy Management IP Address using root user and the password from step 1.1
    - Once logged in as root run cat /etc/haproxy/ca.crt, copy the entire contents and save for the Workload Management installation.
- If you are not able to ping all IP addresses, give the HAProxy a little more time to start all of its services while rechecking your entries for accuracy with CIDR entry requirements.
  - Validate physical core networking for any restrictions in the network.
  - Check vDS DSwitch port group settings for appropriate VLAN ID.

Enable Workload Management and Create a Namespace

Enabling the vSphere Workload Management requires a fully functional HAProxy, which is the ability to ping the Management IP, Workload IP, and the entire Load Balancer IP Range(s). The Namespace can be created once the Workload Management is enabled.

Enable Workload Management

Navigate to vSphere Client Menu and select **Workload Management**. Click **Get Started.**

1. Select vCenter Server Network and click **Next**.
2. Select a Compatible Cluster and click **Next**.

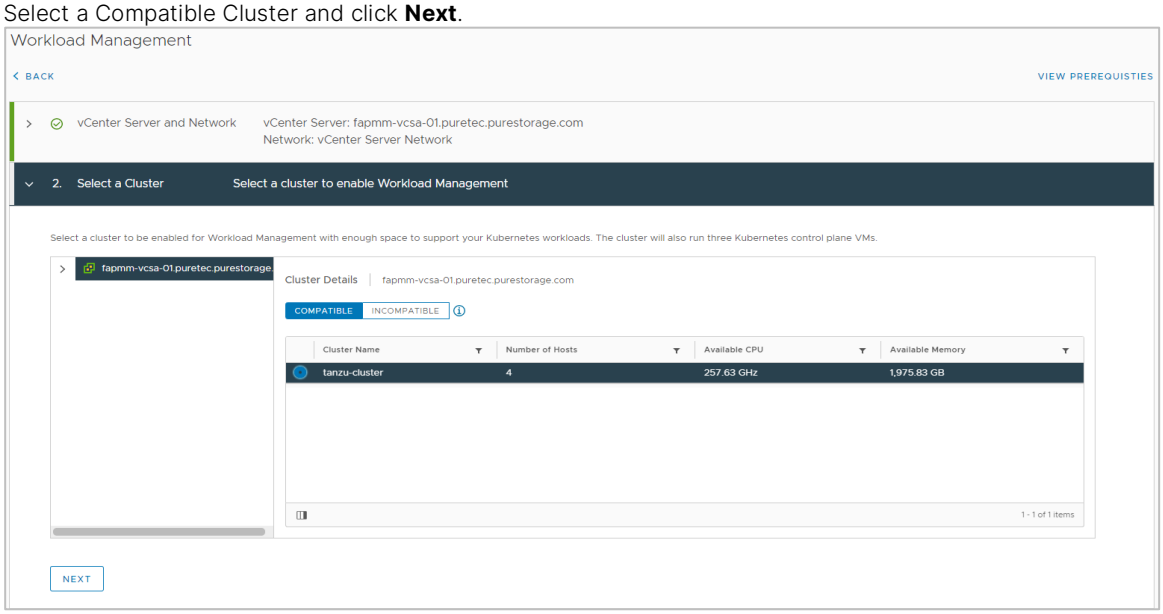

3. Select a Deployment Size and click **Next**.
4. Using the drop-down menu, select vsphere-tanzu-vmfs-policy for the Control Plane VMs and click **Next**.

5. Configure the Load Balancer

   a. IP Address Range is based on your CIDR LB Range Example: 10.21.114.32/27 would be 10.21.114.32-10.21.114.63

   b. Server Certificate Authority: Must retrieve from HAProxy via HAProxy Management IP Address as Root. Run cat `/etc/haproxy/ca.crt`

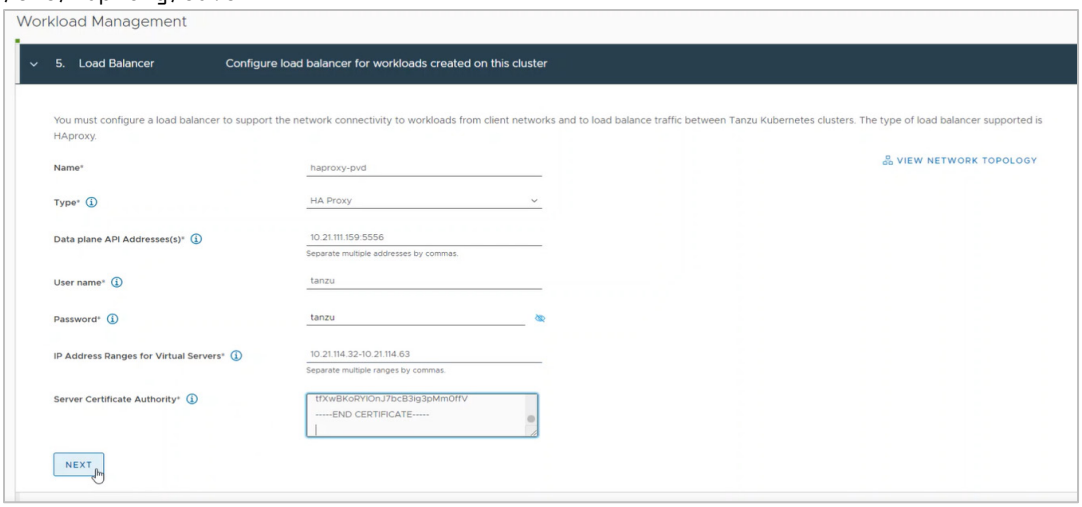

19. Enter management network data of Supervisor Control Plane virtual machines from the worksheet.

   a. **Starting IP Address** entry is the first IP Address of the 5 Supervisor Control Plane IP Addresses

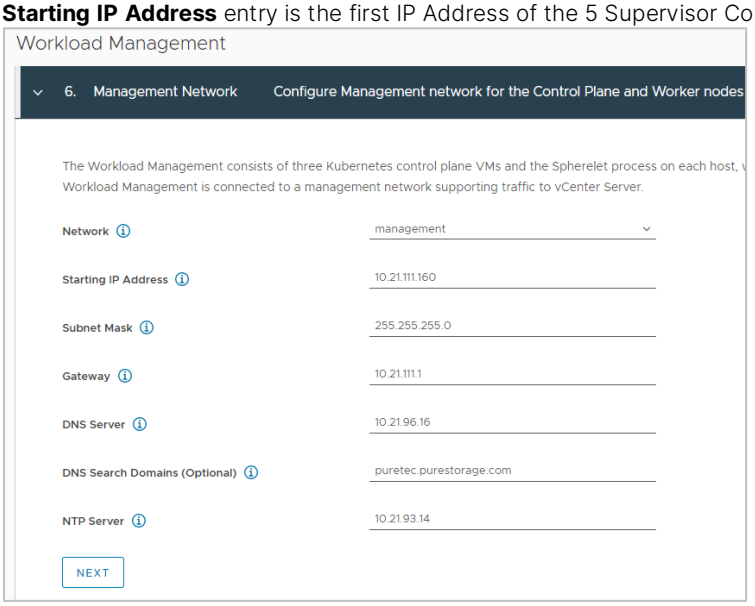

20. Workload Network, use default IP Address for services, enter DNS Server IP and add Workload Network

   a. Workload Network - Create a name, select the workload portgroup and provide the gateway and subnet of the workload network. The “IP Address Ranges” is a range for Virtual Machines and must be outside of the Load Balancer Range. The

total number of IP addresses for this "IP Address Ranges" is 30. **Save** and continue.

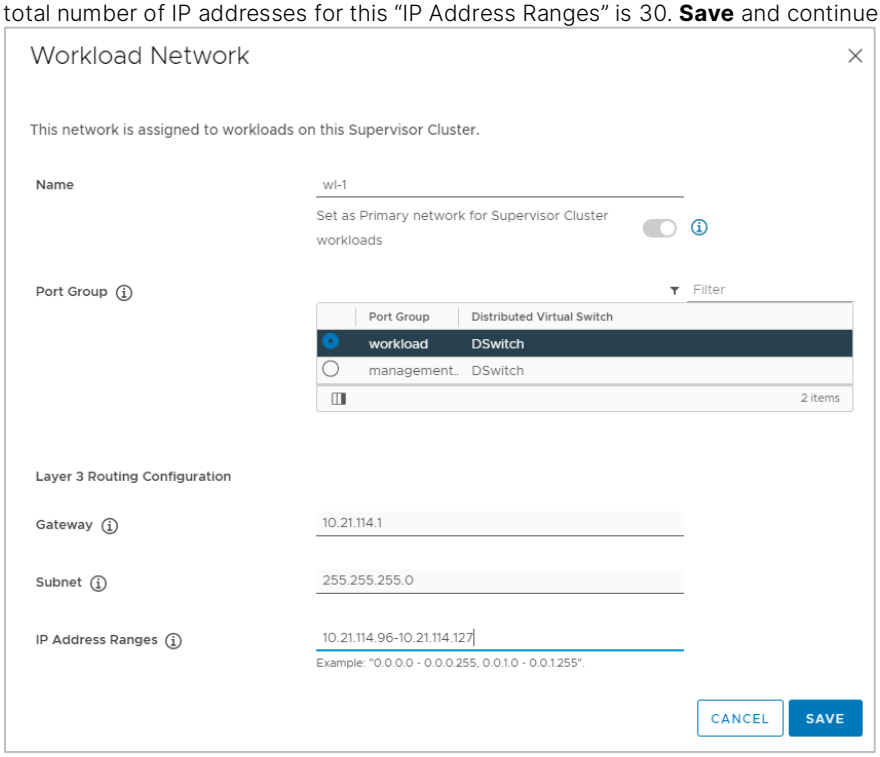

21. Select the TKG Library. Click **OK** to continue.

22. Review and Confirm, select **Finish** to enable workload management.

Create and Configure the vSphere Namespace

Navigate to vCenter Server Menu and select Workload Management. Select Namespaces to Create Namespace.

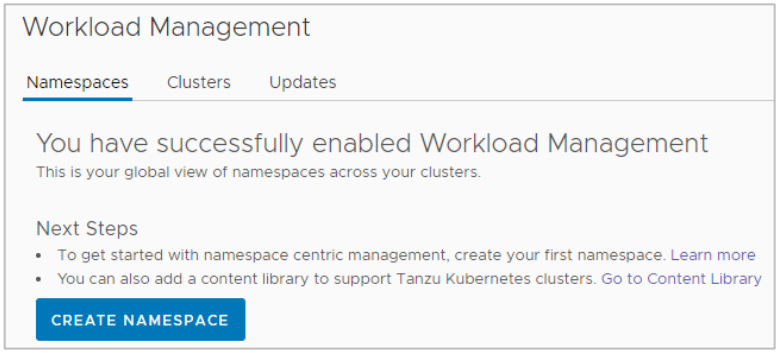

1. Select the designated vSphere Cluster, assign a name and choose the workload network to support Tanzu workloads. Click **Create**.

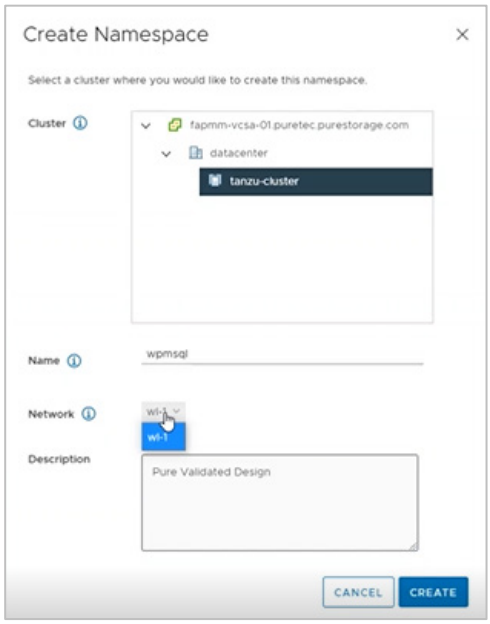

23. Add administrator@vsphere.local Permissions and assign Storage Policies.

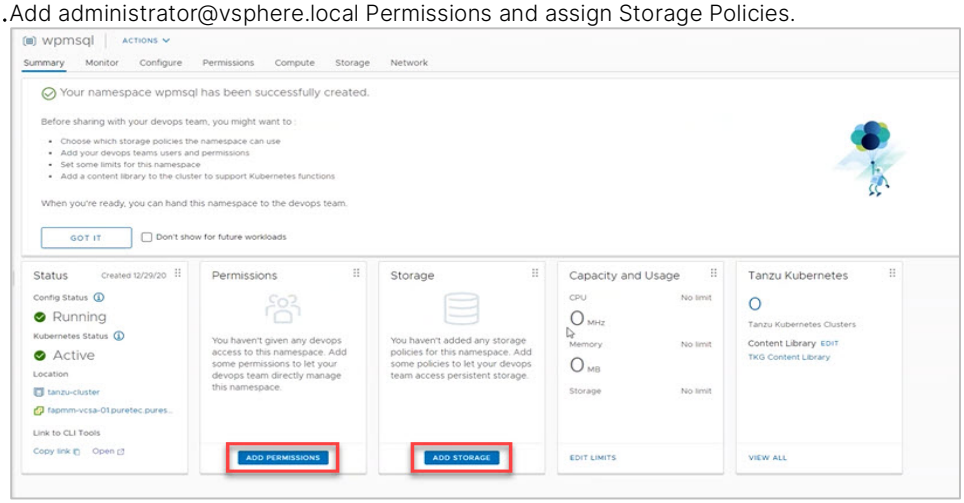

24. Add appropriate User as authorized and applicable for the project, authentication to the Namespace is via vSphere Plugin for the kubectl CLI Tool.

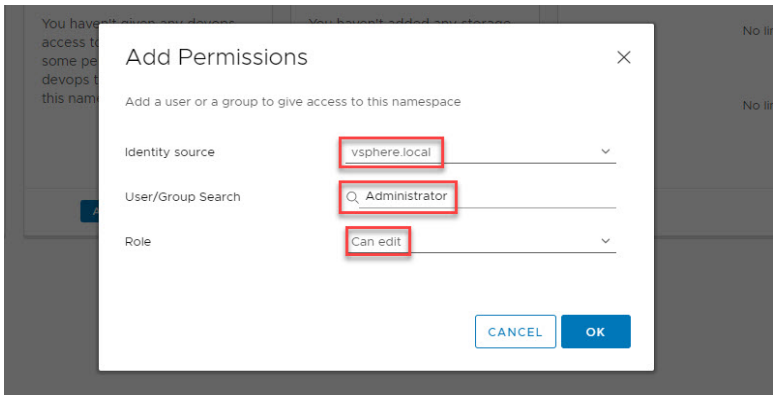

25. Edit namespace storage to Allocate the appropriate Storage Polices for workloads to be provisioned.

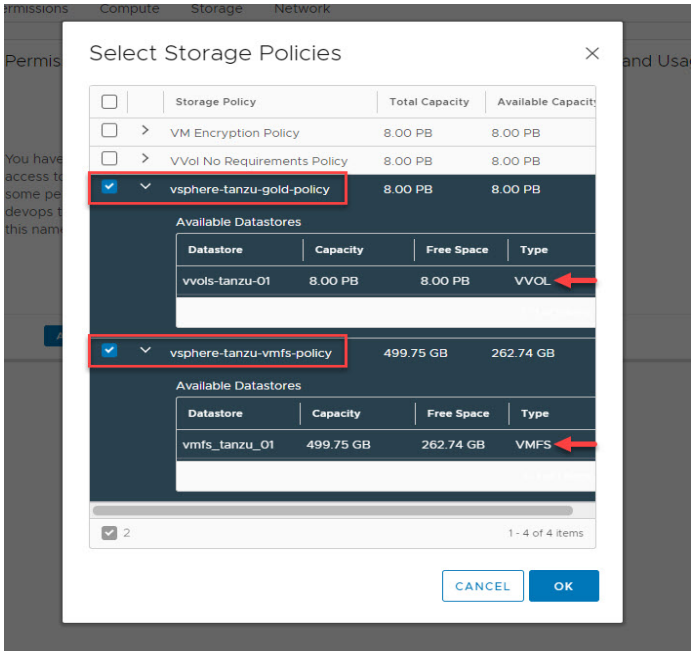

Tanzu Kubernetes Cluster and Workload Deployment

Tanzu Kubernetes Cluster

An instance of Tanzu Kubernetes Cluster will be deployed to the wpmsql workspace using PowerShell with the kubectl CLI tool.

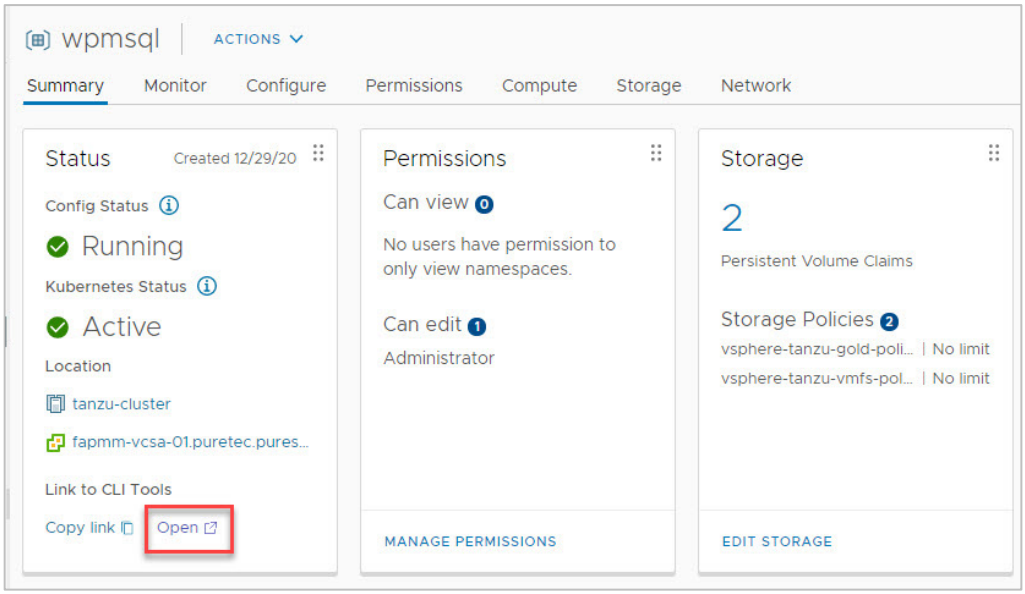

1. Once Kubernetes CLI Tools is installed, log into the Namespace using PowerShell. Example:‘ `kubectl vsphere login --server=10.21.114.32 --vsphere-username administrator@vsphere.local --insecure-skip-tls-verify`
‘

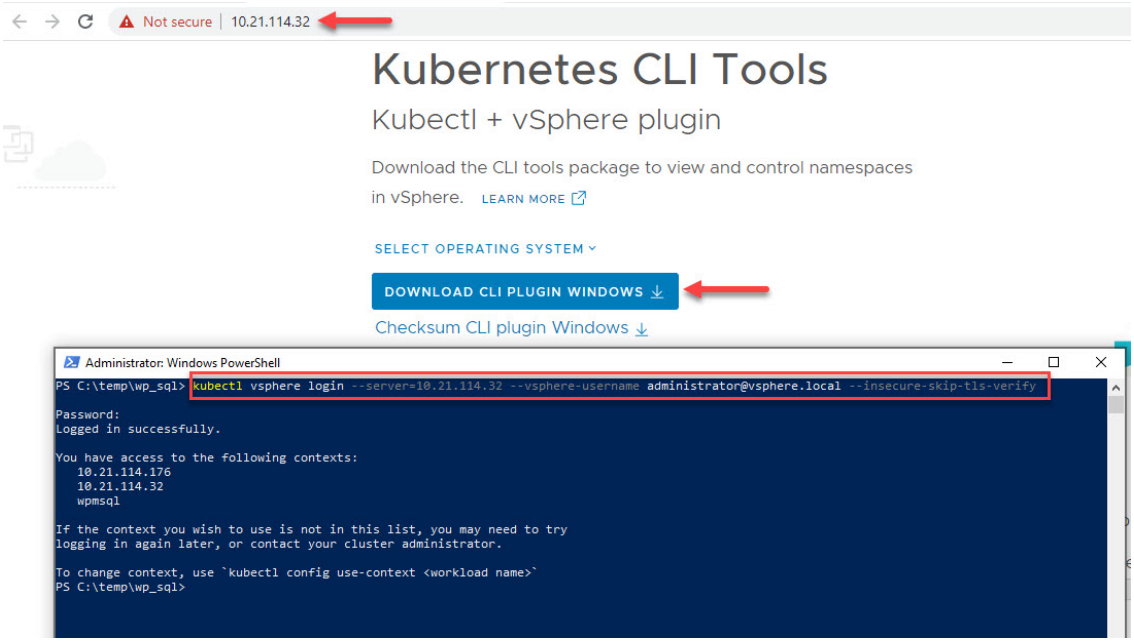

26. Use wpmsql Namespace by running Example: `kubectl config use-context wpmsql`. Run Example: `kubectl config get-contexts`. The asterisk indicates the current context in which you are logged into.

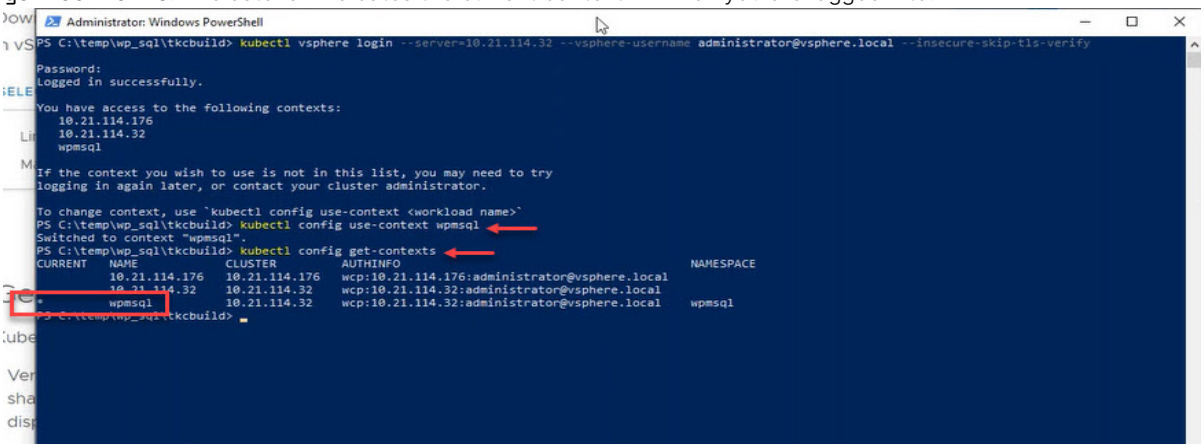

27. Use `kubectl get vmimage` to see the list of available Tanzu Kubernetes Cluster versions.

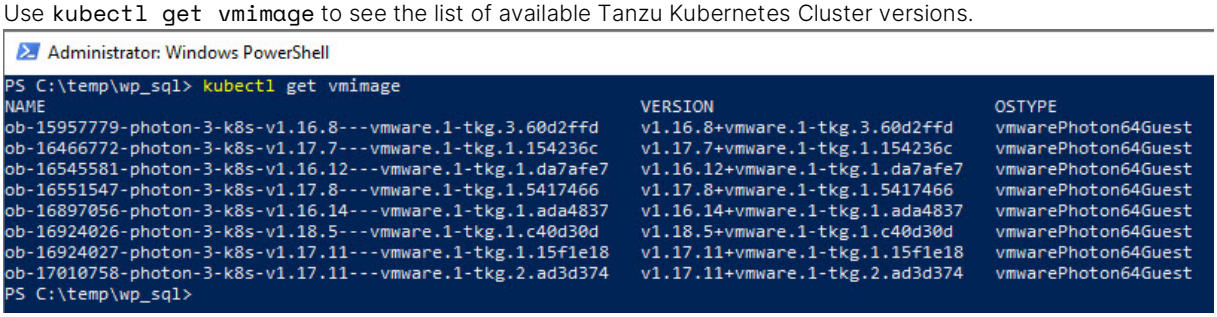

28. Using a source-code editor such as Visual Studio Code, build a design yaml file to deploy the Tanzu Kubernetes Cluster with `storageClass: vsphere-vmfs-policy`.

Note: Tanzu Kubernetes Cluster deployments are currently only supported on FlashArray VMFS Datastores.

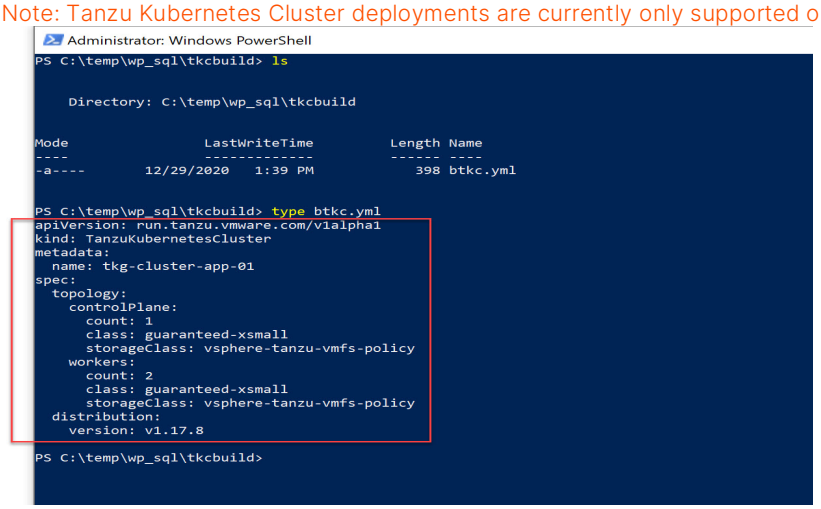

29. Tanzu Kubernetes Cluster design yaml file example:

```
apiVersion: run.tanzu.vmware.com/v1alpha1
kind: TanzuKubernetesCluster
metadata:
  name: tkg-cluster-app-01
spec:
  topology:
    controlPlane:
      count: 1
      class: guaranteed-xsmall
      storageClass: vsphere-tanzu-vmfs-policy
    workers:
      count: 2
      class: guaranteed-xsmall
      storageClass: vsphere-tanzu-vmfs-policy
  distribution:
    version: v1.17.8
```

30. From the directory of the TKC design yaml file manifest the deployment, Example: `kubectl apply -f btkc.yml`

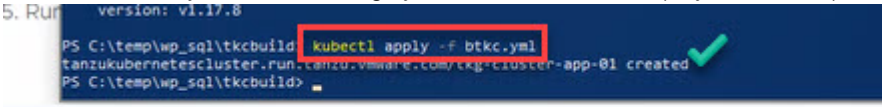

31. Check the status of the Tanzu Kubernetes Cluster deployment, `kubectl get tkc`‘. Once the TKC is running, proceed to the next step.

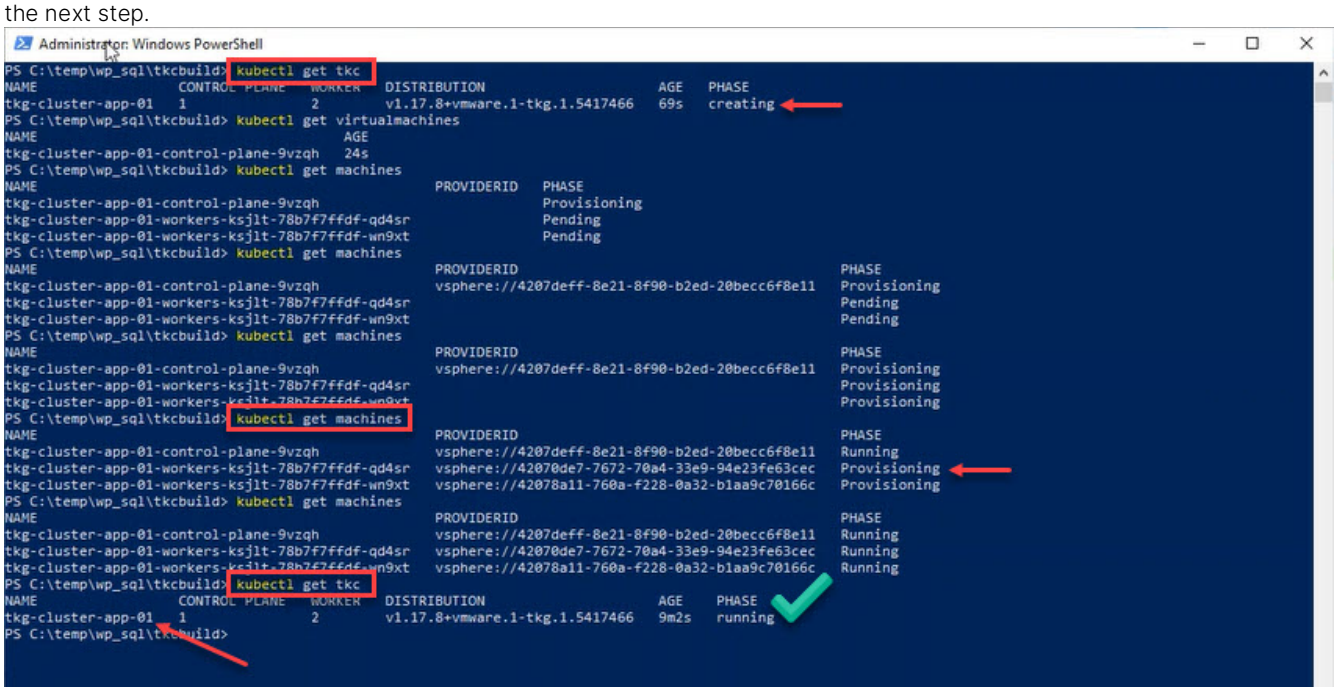

32. Log into tkg-cluster-app-01 Example: ‘kubectl vsphere login --server=10.21.114.32 --insecure-skip-tls-verify --tanzu-kubernetes-cluster-namespace=wpmsql --tanzu-kubernetes-cluster-name=tkg-cluster-app-01 ‘

    a. Use `kubectl config get-contexts` ‘ to validate with the asterisk that you are logged into the Tanzu Kubernetes Cluster.

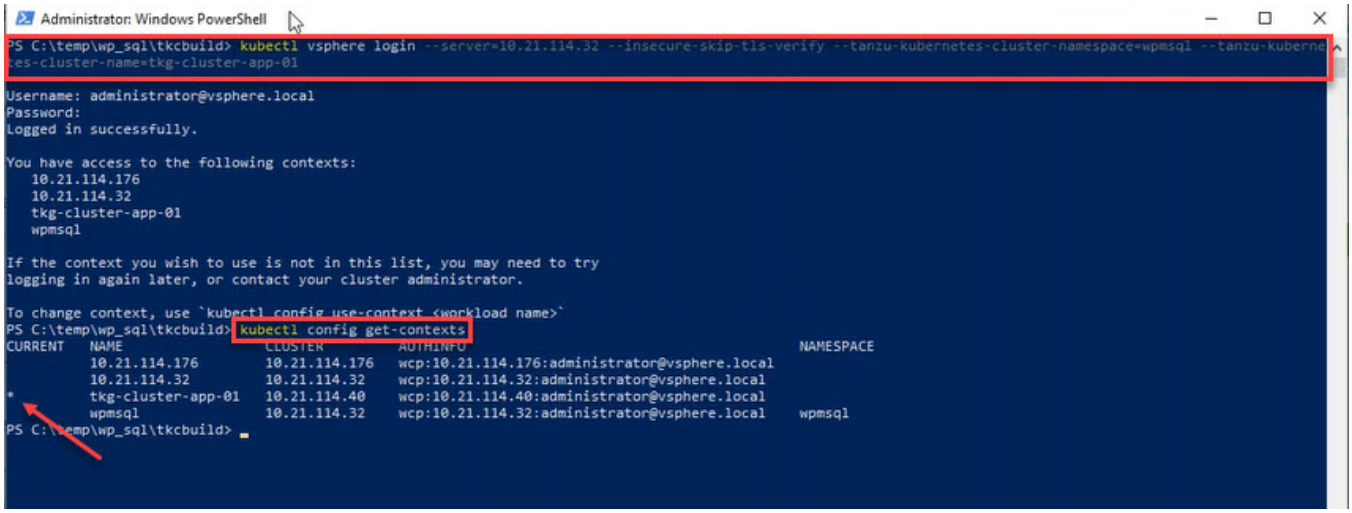

WordPress and MySQL

The final element of this solution is a frontend web application example with WordPress and a MySQL backend.

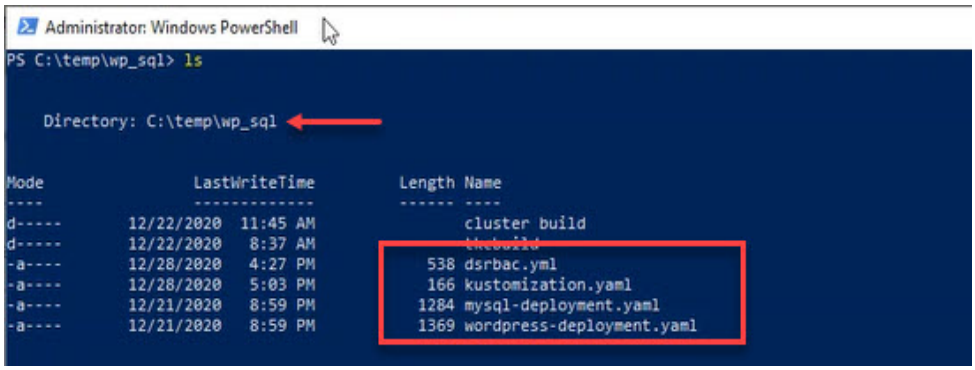

1. Download the following configuration files to the local directory.
   b. mysql-deployment.yaml
   c. wordpress-deployment.yaml
   d. allow-runasnonroot-clusterrole.yaml

33. Using a source-code editor, open and modify both WordPress and MySQL configuration files to include a persistent volume claim with the Storage Class using vVols.
   a. Identify kind: PersistentVolumeClaim and add "storageClassName: <storage-class-containing-flasharray-vvols>" Note: Entry is case sensitive. Example:

```
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: wp-pv-claim
  labels:
    app: wordpress
spec:
  accessModes:
    - ReadWriteOnce
  storageClassName: vsphere-tanzu-gold-policy
  resources:
    requests:
      storage: 20Gi
```

34. Using a source-code editor, create a new yaml configuration file named kustomization.yaml to modify RBAC, generate a Secret, and to run mysql-deployment.yaml and wordpress-deployment.yaml design files. The kustomization.yaml must be in the same directory as the dsrbac, WordPress, and MySQL deployment files. kustomization.yaml. An example is below:

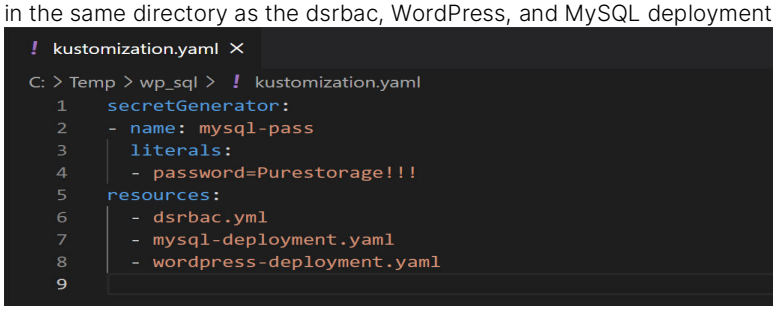

35. Confirm that all four files are in the same directory. Run `kubectl apply -k ./ to Modify RBAC`, Generate the Secret, deploy WordPress and MySQL with persistent volume claims.

Note: This manifest must be executed in the Tanzu Kubernetes Cluster tkg-cluster-app-01.

```
PS C:\temp\wp_sql> ls

    Directory: C:\temp\wp_sql

Mode                 LastWriteTime         Length Name
----                 -------------         ------ ----
d-----       12/22/2020  11:45 AM                cluster build
d-----       12/22/2020   8:37 AM                tkcbuild
-a----       12/28/2020   4:27 PM            538 dsrbac.yml
-a----       12/28/2020   5:03 PM            166 kustomization.yaml
-a----       12/21/2020   8:59 PM           1284 mysql-deployment.yaml
-a----       12/21/2020   8:59 PM           1369 wordpress-deployment.yaml


PS C:\temp\wp_sql> kubectl apply -k ./
clusterrole.rbac.authorization.k8s.io/psp:privileged created
clusterrolebinding.rbac.authorization.k8s.io/all:psp:privileged created
secret/mysql-pass-d7c88cb268 created
service/wordpress-mysql created
service/wordpress created
deployment.apps/wordpress-mysql created
deployment.apps/wordpress created
persistentvolumeclaim/mysql-pv-claim created
persistentvolumeclaim/wp-pv-claim created
PS C:\temp\wp_sql>
```

```
                10.21.114.32      wcp:10.21
*  tkg-cluster-app-01  10.21.114.40      wcp:10.21
                mysql             10.21.114.32      wcp:10.21
PS C:\temp\wp_sql\tkcbuild>
```

36. Confirm the deployment, persistent volume claim and identify the assigned IP address of the application to set up the administrator account and WordPress installation. Examples are

‘ kubectl get deployments ‘

‘ kubectl get pvc ‘

‘ kubectl get svc wordpress ‘

```
Administrator: Windows PowerShell
PS C:\temp\wp_sql> kubectl get deployments
NAME              READY   UP-TO-DATE   AVAILABLE   AGE
wordpress         1/1     1            1           6d2h
wordpress-mysql   1/1     1            1           6d2h
PS C:\temp\wp_sql> kubectl get pvc
NAME             STATUS   VOLUME                                     CAPACITY   ACCESS MODES   STORAGECLASS
mysql-pv-claim   Bound    pvc-4a1ce744-3eb6-4773-b032-3ee47a444f33   20Gi       RWO            vsphere-tanzu-gold-policy
wp-pv-claim      Bound    pvc-c42e3f93-a0d4-4232-873e-29caa81c13c9   20Gi       RWO            vsphere-tanzu-gold-policy
PS C:\temp\wp_sql> kubectl get services wordpress
NAME        TYPE           CLUSTER-IP      EXTERNAL-IP    PORT(S)        AGE
wordpress   LoadBalancer   10.111.122.196  10.21.114.39   80:30150/TCP   6d2h
PS C:\temp\wp_sql>
```

37. Using a web browser, go to **http://<external-ip>** to create the administrator account to install WordPress.

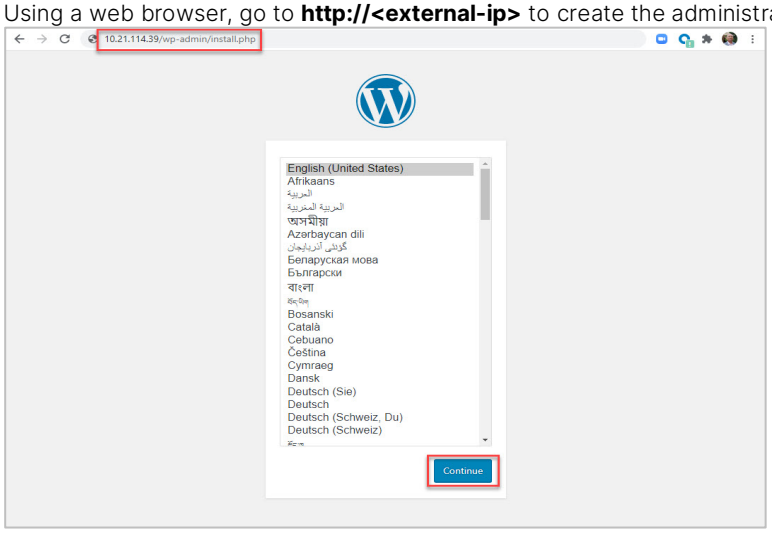

38. Once the administrator account is set up for WordPress, you can customize and publish the web application. WordPress is a fully functional web application that is highly customizable and can be published immediately. You've deployed a Kubernetes Stateful application with persistent storage. Example:

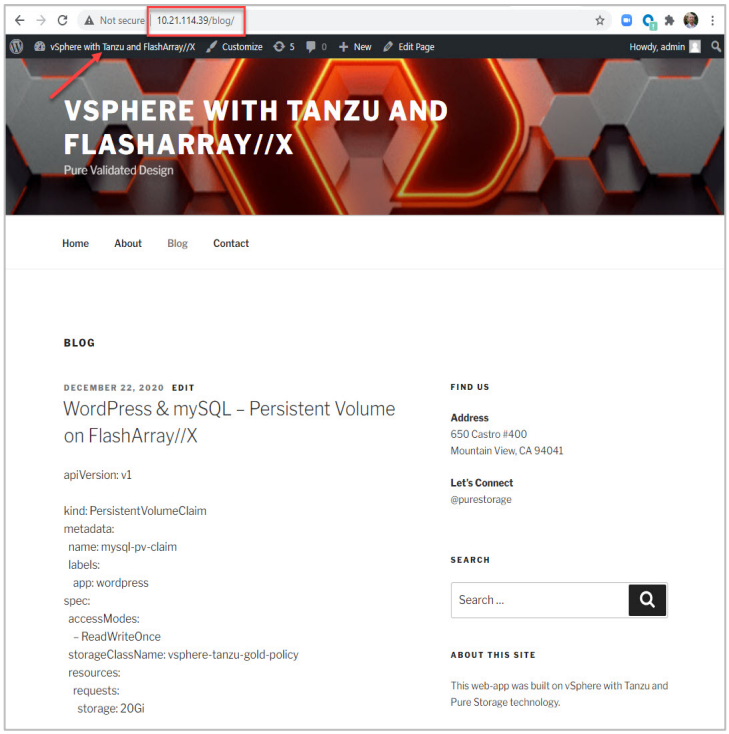

Managing and Monitoring

Monitoring with vSphere Client

Namespaces Dashboard: vSphere Client is the single pane of glass UI to monitor the health of Namespaces and utilization. On the Compute and Storage tab, you can review additional information about Tanzu Kubernetes clusters and persistent volume claims.

Workload Management: Namespace dashboard provide an easy way to identify the current state of the environment, as seen in Figure 9:

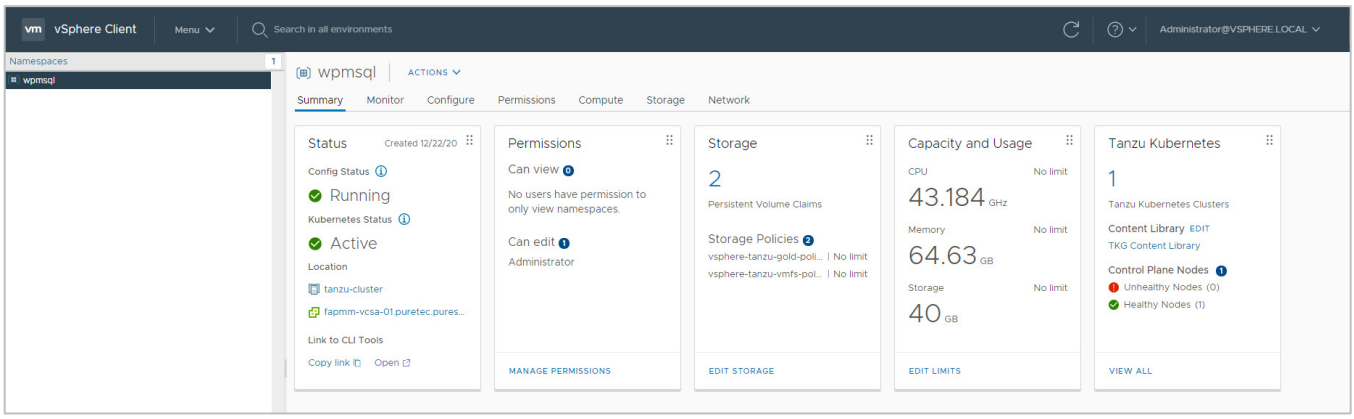

Figure 9. Namespaces dashboard

Accessing Resources: For the Tanzu Kubernetes Cluster, you can use vSphere Client to identify the VMs and Persistent Volumes supporting your application.

Select Compute > Virtual Machines to display Virtual Machines, Name, Status, and VM Class (Figure 10).

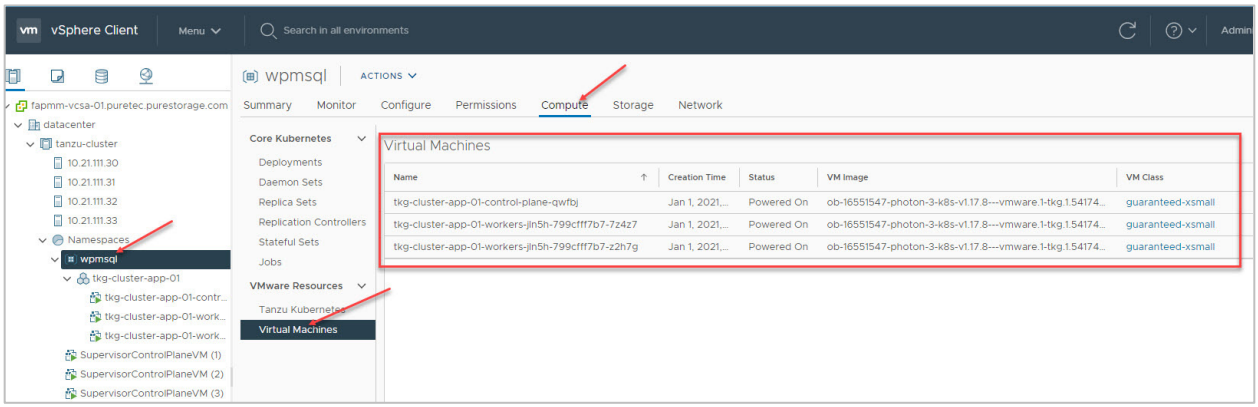

Figure 10. vCenter UI Namespaces - Virtual Machines; Control Plane and Workers

Select Storage > Persistent Volume Claims to view persistent volume claims (Figure 11).

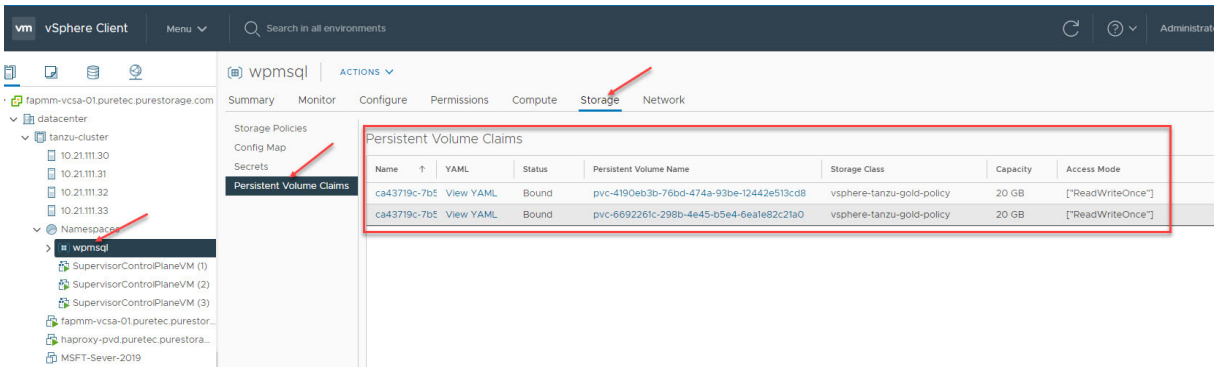

Figure 11. vCenter UI Namespaces - Persistent Volume Claims

Monitoring with Pure1 VM Analytics

VM Topology

Pure1 Analytics provides powerful correlation capabilities. Simply search VMs by name, select the resources you wish to analyze, and hover-over objects for statistics as shown in Figure 12.

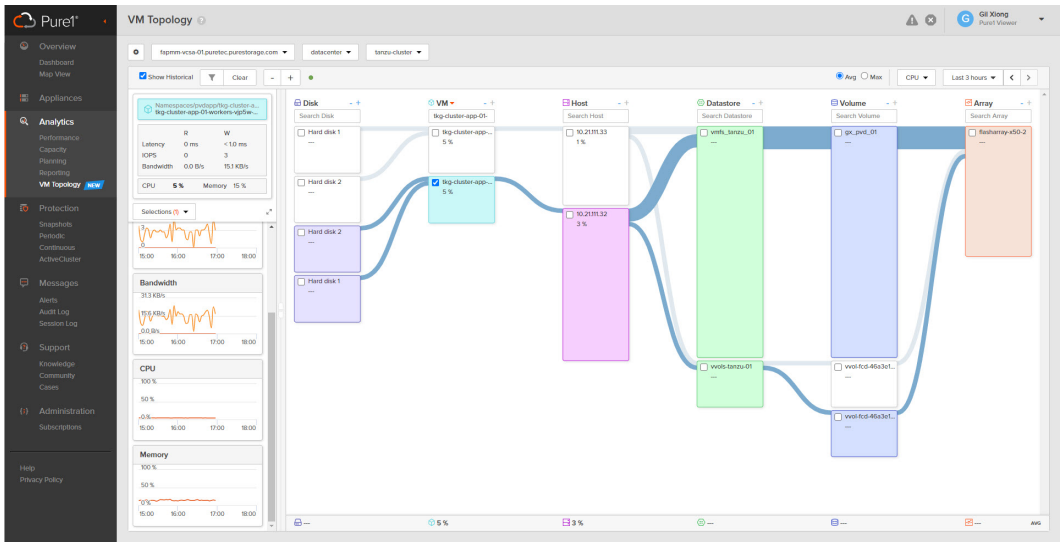

Figure 12. Pure1 VM Topology - Using search function.

The VM Analytics **panels and charts** deliver a single-pane-of-glass to quickly identify metrics of the TKG-Cluster-Worker Nodes:

- CPU and Memory Utilization
- Latency
- IOPS
- Bandwidth
- Capacity
- Capacity churn

Conclusion

VMware and Pure Storage deliver an integrated and validated Tanzu Kubernetes solution to support application modernization. vSphere with Tanzu and FlashArray enables organizations to easily pilot and deploy Containers and orchestrate with Kubernetes with minimal effort and upfront investment. vSphere with Tanzu is accessible, offering multiple editions to meet your requirements from early adoption to a fully automated application platform. Tanzu and FlashArray shared storage is simple and fully integrated with vVols to provide Persistent Volumes for Stateful Applications. FlashArray's predictable and consistent performance is critical to support dynamic Kubernetes workloads, building a foundation of developer-ready infrastructure.

Product Support

As a VMware Ready certified partner, VMware and Pure are committed to jointly support our customers to deliver the best customer experience. Pure Storage and VMware will support their customers following each respective company's standard support process. If Pure Storage needs to engage VMware, we will work together to solve the issue. Pure Storage offers support over the phone, by email, and through our web portal.

Contact Pure Storage Support

- Web pure1.purestorage.com/support
- Email support@purestorage.com
- Phone (US) +1 (866) 244-7121 or +1 (650) 729-4088
- Phone Numbers (International) support.purestorage.com/pure1/support

Professional Services

Projects involving containers, Kubernetes, and VMware Tanzu can be complex. Pure's Professional Services offer workshops and advisory services that can help you plan and execute these projects. Our services provide proven and cost-effective methods, processes, and tools.

Additional Documentation and Reference

- VMware VSphere with Tanzu POC Guide
- Pure Storage VMware Platform Guide
- VMware Pure Storage Partnership

Document Updates

We are always looking to improve the quality of our content and documentation and welcome your feedback. Please send us your comments at pvd-documents@purestorage.com.

Document Revisions

Rev #	Description	Date
1.0	Initial Publication	February 2021

Pure Storage, Inc.
650 Castro Street, #400
Mountain View, CA 94041