

Dell EMC PowerProtect Data Manager protecting VMware Tanzu Kubernetes Clusters

Abstract

VMware vSphere with Tanzu transforms vSphere clusters into a platform on which Kubernetes workloads can be run directly on VMware ESXi hosts and can create Kubernetes clusters within dedicated namespaces. This document describes the architecture and configuration of VMware Tanzu Kubernetes clusters with Dell EMC PowerProtect Data Manager and explains how the VMware vSphere with Tanzu workloads are protected.

February 2021

Revisions

Date	Description
February 2021	Initial release

Acknowledgements

Author: Abhishek Shukla, Solutions Technical Marketing Team, Data Protection Domain

The information in this publication is provided “as is.” Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © 2021 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners. [24-Feb-21] [Technical Whitepaper] [H18682]

Table of contents

Revisions.....	2
Acknowledgements.....	2
Table of contents	3
Executive summary.....	5
Audience	5
Scope	5
1 Introduction.....	6
1.1 VMware vSphere with Tanzu components.....	6
1.1.1 Workload.....	6
1.1.2 Supervisor Cluster	6
1.1.3 Supervisor Namespace	6
1.1.4 Tanzu Kubernetes Cluster	6
1.1.5 vSphere Pod.....	7
1.1.6 PersistentVolume (PV) and PersistentVolumeClaim (PVC).....	7
1.1.7 Storage Class (SC).....	7
1.1.8 Custom Resource (CR)	7
1.2 PowerProtect Data Manager components	7
1.2.1 Cloud Native Data Manager	7
1.2.2 PowerProtect Controller	7
1.2.3 VMware Velero	7
1.2.4 vProxy (VM proxy).....	8
2 Prerequisites for enabling vSphere with Tanzu.....	9
2.1 vSphere Cluster.....	9
2.2 Networking Stack.....	9
2.3 Storage policy	9
2.4 Content Library	9
3 Architecture	10
4 Configuring VMWare vSphere with Tanzu	12
4.1 Create Content Library	12
4.2 Enable Workload Management	13
4.3 Create Supervisor Namespaces.....	15
4.4 Create Tanzu Kubernetes Cluster (Guest Cluster)	19
5 Configure PowerProtect Data Manager to protect Tanzu Kubernetes workloads	22
5.1 Prerequisites to Tanzu Kubernetes Guest Cluster Protection.....	21

5.2	Asset Discovery	22
5.3	VM Direct Engine.....	27
5.4	Backup Configuration	28
5.5	Replication Configuration	30
5.6	Restore Configuration.....	32
A	Technical support and resources	33
A.1	Related resources.....	33

Executive summary

Modern IT infrastructure is being transformed by Containers. Containers are similar to virtual machines but have relaxed isolation properties to share the operating system. The Container has its own filesystem, CPU, memory and process space. Agile application creation, continuous development, environmental consistency across development, application-centric management, efficient resource allocation and resource isolation are the key benefits of containers. Kubernetes is an open-source container management platform that unifies a cluster of machines into a single pool of compute resources.

VMware vSphere is the compute virtualization platform. VMware vSphere 7 rearchitected with native Kubernetes for application modernization that enable IT admins to use vCenter server to operate Kubernetes clusters through namespaces. VMware vSphere with Tanzu provides a platform for both traditional applications as well as modern applications so that both IT admins and developers can access developer-ready infrastructure, scale with simple operations.

With currently distributed container deployment, it is important to protect the workloads. Dell EMC PowerProtect Data Manager protects the workloads and ensures high availability, consistent, and reliable backup and restore for Kubernetes workload or DR situation. PowerProtect Data Manager offers centralized management, automation, multi-cloud options and advanced integration for ease and simplicity for managing workloads. PowerProtect Data Manager protects Tanzu Kubernetes cluster, pods, persistent volume claims, namespaces, and other resources.

Audience

This whitepaper is intended for customers, partners and others who want to understand how PowerProtect Data Manager software helps to protect VMware Tanzu Kubernetes clusters and the workloads.

Scope

1. VMware vSphere with Tanzu - A Tanzu edition license of vSphere 7.0.1 and above.
2. Enable workload management.
3. PowerProtect Data Manager 19.7 and above.

1 Introduction

The Cloud Native definition is an architectural philosophy for designing the applications and infrastructure 'Containers' provide a way to package and run the application. To run such applications, container orchestrator is required. Kubernetes is an open-source container orchestrator for managing containerized workloads and services, that facilitate both declarative configuration and automation. It is portable, extensible, and scalable and has a large, rapidly growing ecosystem. Kubernetes services, support, and tools are widely available and these days the applications are constructed of multiple microservices that run a large number of Kubernetes pods and VMs. VMware vSphere with Tanzu helps in creating Kubernetes control plane directly on VMware ESXi by creating Kubernetes layer within ESXi that are part of the Kubernetes cluster.

Dell EMC PowerProtect Data Manager protects existing as well as new discovered workloads. It allows IT operators and backup admins to manage VMware Tanzu clusters and its protection through a single management UI and define protection policies for Kubernetes workloads from Kubernetes APIs. The policy driven protection is defined by the Protection Policy mechanism. PowerProtect Data Manager discovers the namespaces, labels, and pods in the environment and can be protected by providing cluster credentials. Logging, Monitoring, governance, and recovery are done through PowerProtect Data Manager.

1.1 VMware vSphere with Tanzu components

VMware vSphere with Tanzu provides platform for running Kubernetes workloads natively on VMware ESXi.

1.1.1 Workload

In vSphere with Tanzu, the workload is an application deployed that consists of containers running inside vSphere Pods, VMs or both. It is an application that run inside Tanzu Kubernetes cluster that are deployed by Tanzu Kubernetes Grid service.

1.1.2 Supervisor Cluster

The Supervisor Cluster provides the management plane on which Tanzu Kubernetes clusters are built. The service called The Tanzu Kubernetes Grid (TKG) service is a controller manager that includes set of controllers which are subset of supervisor cluster. TKG service helps in provisioning Tanzu Kubernetes cluster.

1.1.3 Supervisor Namespace

When Tanzu Kubernetes clusters are provisioned, a resource pool and VM folder are created in a supervisor namespace. The resource quotas and storage policy are applied to a namespace and inherited by the Tanzu Kubernetes cluster deployed. The Tanzu Kubernetes cluster control plane and worker node VMs are placed within the resource pool and VM folder.

1.1.4 Tanzu Kubernetes Cluster

The Tanzu Kubernetes cluster is distribution of open-source Kubernetes container platform that is built, signed, and supported by VMware. Tanzu Kubernetes clusters are built on top of supervisor cluster. It is defined in the supervisor namespace using custom resource. It uses the open-source Photon OS from VMware and is integrated with underlying vSphere infrastructure including storage, network, and authentication.

1.1.5 vSphere Pod

vSphere Pod is a VM with a small footprint that runs one or more containers. It is similar to Kubernetes Pod. Each pod is sized for the workload that has explicit resource reservations for that workload. It allocated exact amount of storage, memory and CPU required for the workload to run.

1.1.6 PersistentVolume (PV) and PersistentVolumeClaim (PVC)

A PersistentVolume (PV) is a piece of storage in the cluster that has been provisioned by an administrator or dynamically provisioned using Storage Classes. It is a resource in the cluster just like a node is a cluster resource. PVs are volume plugins like Volumes but have a lifecycle independent of any individual Pod that uses the PV. This API object captures the details of the implementation of the storage, be that NFS, iSCSI, or a cloud-provider-specific storage system.

A PersistentVolumeClaim (PVC) is a request for storage by a user. It is similar to a Pod. Pods consume node resources and PVCs consume PV resources.

1.1.7 Storage Class (SC)

A Storage Class is described as the type of storage that is provisioned and allowed ranges for size and IOPS. When user creates a PVC, that specifies the storage class with size in GB and number of IOPS. A storage class is used to abstract the underlying storage platform.

1.1.8 Custom Resource (CR)

A resource in Kubernetes environment is an endpoint for API that stores a collection of API objects of a certain kind and A Custom Resource (CR) is an extension of the Kubernetes API that is not necessarily available in a default Kubernetes installation. It represents a customization of a particular Kubernetes installation.

1.2 PowerProtect Data Manager components

1.2.1 Cloud Native Data Manager

The Cloud Native Data Manager (CNDM) is in-built microservice component of PowerProtect Data Manager which communicates with the kube-apiserver of the cluster. This component is responsible for APIs for the backup and restore process.

1.2.2 PowerProtect Controller

PowerProtect controller is the component which gets installed on Kubernetes cluster when the cluster gets discovered by PowerProtect Data Manager. The backup and restore controllers that manager BackupJob CR and RestoreJob CR definitions. This component is responsible for the backup and restore of Persistent Volumes.

1.2.3 VMware Velero

VMware Velero is an open-source tool which is integrated with PowerProtect Data Manager. It is in-built and does not require to be installed separately. Velero component is pushed into the Kubernetes cluster by the PowerProtect controller pod after the same is in up and running state using Velero deployment object. It is responsible for the backup and restore of metadata.

1.2.4 vProxy (VM proxy)

The vProxy protection engine is the virtual machine data protection component within PowerProtect Data Manager. During backups, the vProxy agent creates a snapshot of virtual-machine data directly from the datastore. The snapshot is moved directly to the target storage where the backups are stored. This process uses VMware vSphere Storage API for Data Protection (VADP) which enables centralized, off-host, LAN-free backup of virtual machines.

Note: The VADP is a subset of the vSphere API that enables backup and restore applications. The snapshot-based VADP framework allows efficient, off-host, centralized backup of virtual-machine storage. After taking a snapshot to quiesce virtual disks, software can offload the backup load to the target storage.

2 Prerequisites for enabling vSphere with Tanzu

To configure or to run Kubernetes workloads natively on vSphere, Workload management is required to be enabled that creates Supervisor cluster where the vSphere pods run and to provision Tanzu Kubernetes clusters or Guest clusters. There are few prerequisites for compute, network, and storage.

2.1 vSphere Cluster

- vSphere cluster is a collection of ESXi hosts managed by vCenter server. To enable Workload Management, at least 3 ESXi hosts are a must, if you are using VSAN, then a minimum 4 ESXi hosts are required.
- vSphere cluster must be configured with High-Availability (HA) enabled
- vSphere cluster must be configured with Distributed Resource Scheduler (DRS) enabled and must be set to fully automated mode.
- The cluster must use shared storage for vSphere HA, DRS and for storage persistent volumes

2.2 Networking Stack

To enable workload management, the networking must be configured for the Supervisor Cluster. A Supervisor Cluster can either use the vSphere networking stack or VMware NSX-T™ Data Center to provide connectivity to Kubernetes control plane VMs, services, and workloads. When a Supervisor Cluster is configured with the vSphere networking stack, all hosts from the cluster are connected to a vSphere Distributed Switch (vDS) that provides connectivity to Kubernetes workloads and control plane VMs. A Supervisor Cluster that uses the vSphere networking stack requires a third-party load balancer that provides connectivity to DevOps users and external services. A Supervisor Cluster that is configured with VMware NSX-T™ Data Center, uses the software-based networks of the solution as well as an NSX Edge load balancer to provide connectivity to external services and DevOps users. However, PowerProtect Data Manager supports the protection of Tanzu Kubernetes cluster only with NSX-T configuration.

- To use NSX-T Data Center networking for the Supervisor cluster, there are system requirements and topologies to be reviewed <https://docs.vmware.com/en/VMware-vSphere/7.0/vmware-vsphere-with-tanzu/GUID-B1388E77-2EEC-41E2-8681-5AE549D50C77.html>
- vCenter and ESXi that are part of work load management cluster needs to be prepared for NSX-T. Refer VMware documentation <https://docs.vmware.com/en/VMware-vSphere/7.0/vmware-vsphere-with-tanzu/GUID-8D0E905F-9ABB-4CFB-A206-C027F847FAAC.html> to install and configure NSX-T Data Center for vSphere with Tanzu.

2.3 Storage policy

Storage policies are created for the datastore placement for Kubernetes control plane VMs, containers and images. Storage policies are associated with different storage classes. Before enabling workload management, a storage policy is created for the placement of Kubernetes control plane VMs.

- Make sure the datastore is shared between all ESXi hosts in the cluster
- VM storage policies must be configured and updated

Storage policy creation with vSphere : <https://docs.vmware.com/en/VMware-vSphere/7.0/vmware-vsphere-with-tanzu/GUID-544286A2-A403-4CA5-9C73-8EFF261545E7.html#GUID-544286A2-A403-4CA5-9C73-8EFF261545E7>

2.4 Content Library

Content library consists of distributions of Tanzu Kubernetes releases in the shape of OVA templates. You can create a **Local Content Library** where images are uploaded manually or can create **Subscribed Content Library** to pull the latest released images automatically.

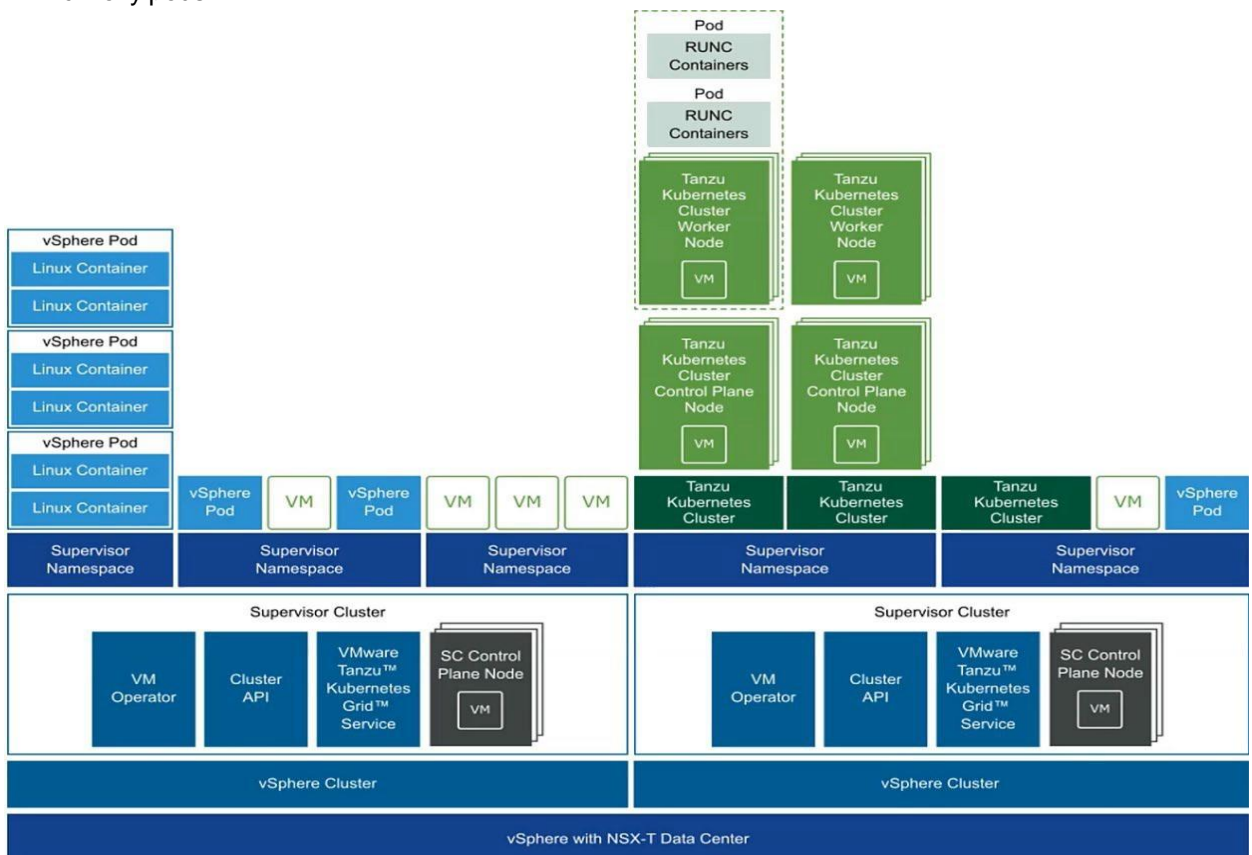
3 Architecture

PowerProtect Data Manager 19.7 introduces the ability to protect Tanzu Kubernetes Guest cluster workloads. VMware vSphere 7U1 re-architectures vSphere with native Kubernetes as its control plane. A TKG cluster is a Kubernetes cluster that runs inside the Virtual Machines on Supervisor layer which allows to run Kubernetes with consistency. It is enabled via the TKG service for VMware vSphere and is upstream-compliant with open-source Kubernetes (Guest cluster). The Guest cluster is a consistent Kubernetes cluster running on VMs and consists of control plane VM, worker nodes, pods, and containers.

PowerProtect Data Manager protects Kubernetes workloads and ensures the data is consistent and highly available. PowerProtect Data Manager is a virtual appliance that is deployed on an ESXi host using OVA and is integrated with Dell EMC PowerProtect DD series as protection target where backups are stored.

Once the discovery of cluster completes, the cluster is added as a PowerProtect Data Manager asset source and associated namespaces as assets are available to be protected. During the process of the discovery, PowerProtect Data Manager creates the two namespaces mentioned below in the cluster. The data is compressed and deduplicated at the source and sent to the target storage.

1. **Velero-ppdm:** Contains a Velero pod to backup metadata and stage to the target storage in case of a BareMetal environment. It performs PVC snapshot and metadata backup in case of VMware Cloud Native Storage (CNS).
2. **PowerProtect:** Contains a PowerProtect controller pod to drive Persistent Volume Claim snapshot and backup and push the backups to target storage using intermittently spawned cProxy pods.



Note: The pods running in the guest clusters do not have direct access to the Supervisor cluster and the persistent volumes provisioned by vSphere CSI on the guest cluster creates an FCD disk in supervisor space and mapped to the guest cluster via paravirtual CSI driver. PowerProtect Data Manager uses internal APIs to protects paravirtual volumes.

According to Tanzu Kubernetes cluster architecture, vSphere cluster (ESXi as worker node) has Supervisor clusters and Guest Clusters (TKG Clusters). The guest clusters have their own control plane VMs, management plane, worker nodes, networking, pods and namespaces and are isolated from each other. Supervisor Clusters and Guest clusters communicate via API servers.

The cProxy of PowerProtect Data Manager does not have access to Supervisor resources such as FCD that gets created as part of provisioning the PersistentVolumes in guest cluster, as it is external to the clusters, therefore, PowerProtect Data Manager does not use cProxy for backup and restore process. However, PowerProtect Data Manager utilizes the vProxy based protection solution.

The vProxy moves snapshot of the FCD created by the Velero vSphere plugin. The snapshot is moved directly to the target storage where the backups are stored. When the backup job is triggered, CNDM communicates with VM direct to find and reserve a vProxy. The vProxy should be created at the vCenter specifically for TKG clusters.

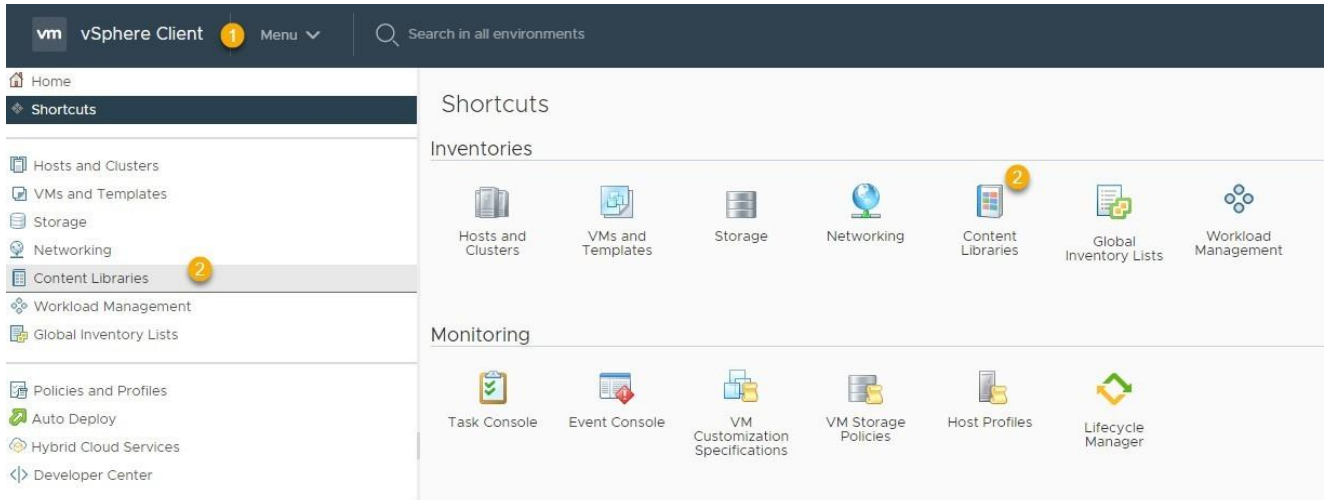
4 Configuring VMware vSphere with Tanzu

To provision Tanzu Kubernetes cluster, Content Library is required to be created in the vCenter server that manages the vSphere cluster where the Supervisor cluster runs.

4.1 Create Content Library

The content library provides the distribution of Tanzu Kubernetes releases in the shape of OVA templates.

1. Login to the **vCenter Server** with administrator credentials.
2. Select **Content Libraries** under Inventories tab.



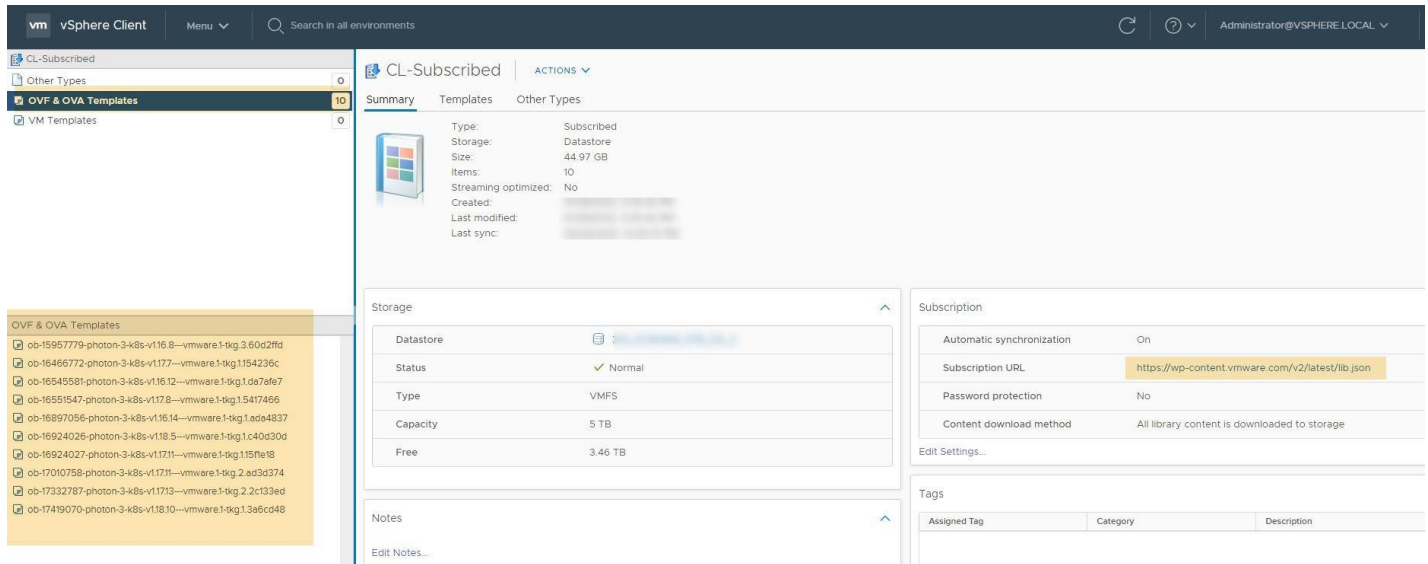
3. Click on **+ Create** and enter the required details.
 - **Name:** Specify the name
 - **Notes:** Optional
 - **vCenter Server:** Select the vCenter from the dropdown list (if incase there are multiple vCenter servers)

4. Click **NEXT**

Note: You can create a **Subscribed Content Library** to automatically pull the latest released images or you can a **Local Content Library** and upload the images manually

Subscription URL: <https://wp-content.vmware.com/v2/latest/lib.json>

5. Click **NEXT**
6. Verify the identity of the subscription host and click **YES** to proceed
7. Select the storage location for the library contents and Click **NEXT**
8. Review content library settings and click **FINISH**
9. Library is created and is available under **Content Libraries** section
10. Click on recent created library and observe the details and OVAs available



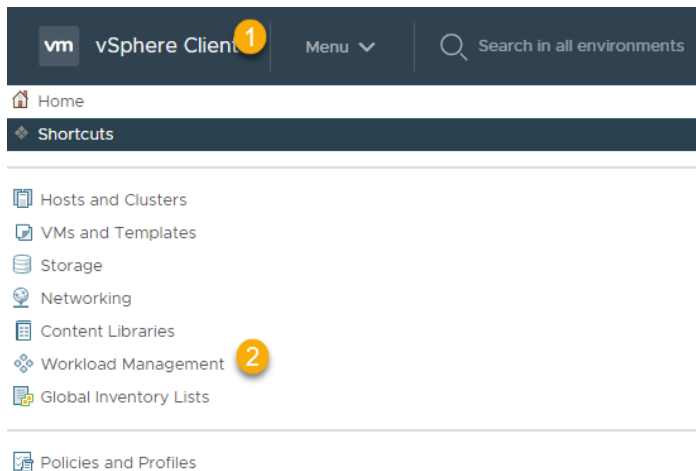
4.2 Enable Workload Management

Enabling workload management on a vSphere cluster creates Supervisor cluster. Workload Management enables deploying and managing Kubernetes workloads in vSphere. By using workload management, you can leverage both Kubernetes and vSphere functionality. Once vSphere cluster for workload management is configured, namespaces can be created which provides compute networking and storage resources for Kubernetes applications.

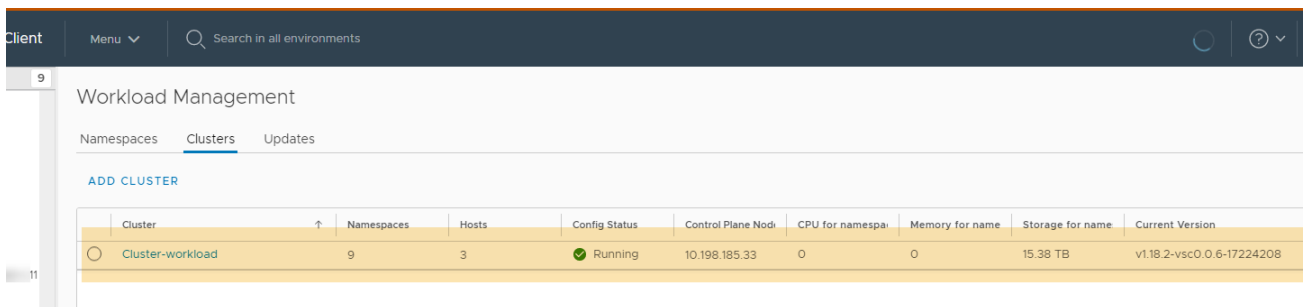
- **Network Support:** You can select between two networking stacks when configuring workload management such as NSX-T and vCenter server networks. You can check the checklist for the same by clicking **Menu > Workload Management > Network Support**
- **HA and DRS Support:** HA and DRS must be enabled on the vSphere cluster in fully automated mode on the cluster where you set up workload management.
- **Storage Policy:** Storage policies must be created that determines the datastore placement of the Kubernetes control plane VMs, containers and the images.
- **Load Balancer:** A Supervisor Cluster that is configured with VMware NSX-T™ Data Center, uses the software-based networks of the solution as well as an NSX Edge load balancer to provide connectivity to external service. If the vCenter Server network is used, a load balancer must be configured to support the network connectivity to workloads from client networks and for load balancing the traffic between Tanzu Kubernetes clusters.
- **Tanzu Kubernetes Grid:** The content library must be created on the vCenter server system. The VM image that is used for creating the nodes of Tanzu Kubernetes clusters is pulled from that library. This library will contain the latest distributions for Kubernetes and another OS. (<https://wp-content.vmware.com/v2/latest/lib.json>)

Steps to enable workload management

1. Login to the **vCenter server** with administrator credentials.
2. Select **Workload Management**



3. Click on **GET STARTED**
4. **vCenter Server and Network:** Select a vCenter and then select a networking stack option and click NEXT
5. **Select a Cluster:** Select the compatible cluster listed in the cluster details and click NEXT.
6. **Control Plane Size:** Allocate capacity for the Kubernetes control plane VMs. The amount of resources that you allocate to the control plane VMs determine the amount of Kubernetes workloads the cluster can support. Select from resource allocation size and click NEXT.
7. **Storage:** Select the storage policy to be used for datastore placement of Kubernetes control plane VMs and containers. This policy is associated with a datastore on the vSphere environment.
8. **Management Network:** The workload management consists of three Kubernetes control plane VMs and the spherelet process on each host, which allows the host to be joined to a Kubernetes cluster. The cluster where the workload management is connected to management network supporting traffic vCenter server.
9. **Workload Network configuration:** Configure the NSX-T capable vDS switch, NSX-T edge cluster, POD CIDR, Service CIDR, Ingress CIDR and Egress CIDR for TKG guest cluster VMs.
10. **TKG configuration:** Add content library to give the access to the workloads
11. **Review and Confirm:** Review all the details before confirming the setup for workload management on the cluster.
12. Click **FINISH**



13. Cluster is available under **Menu > Workload Management > Clusters**

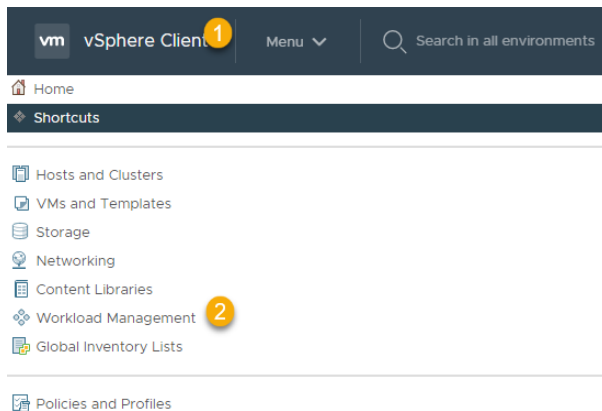
Note: As mentioned above, the workload management consists of three Kubernetes control plane VMs which allows the ESXi hosts (Kubernetes nodes) to be joined in the Kubernetes cluster. Once the workload cluster is created, you would observe three **SupervisorControlPlaneVM** are created. These are the control plane VMs and interact with vSphere infrastructure to provide the services and capabilities for vSphere with Tanzu.

4.3 Create Supervisor Namespaces

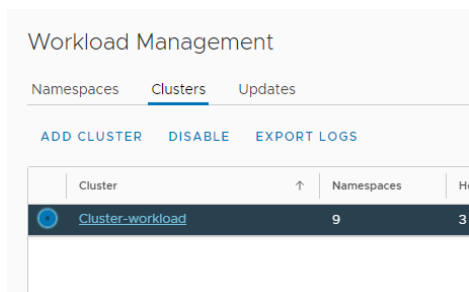
Once the Supervisor cluster is deployed, configured and licensed, the Supervisor namespace can be deployed on the Supervisor cluster to run Kubernetes applications

Steps to create Supervisor namespaces

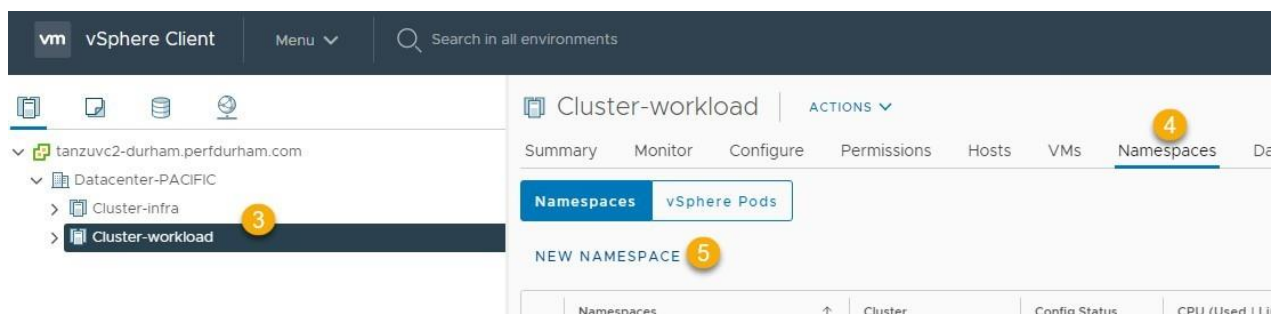
1. Login to the **vCenter server** with administrator credentials.
2. Select **Workload Management**



3. Select the Supervisor cluster created

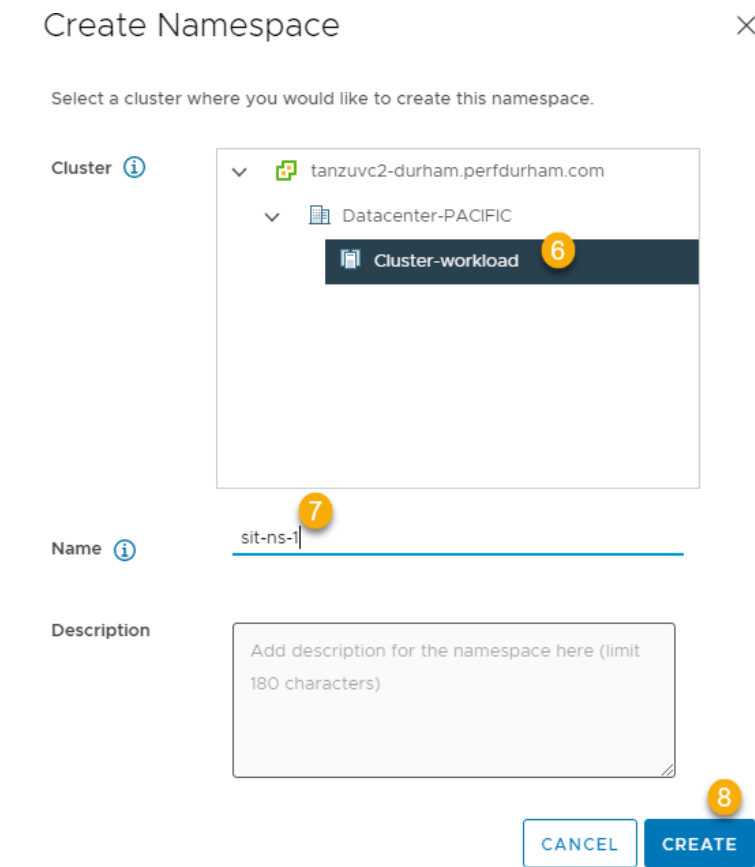


4. Click on **Namespaces**

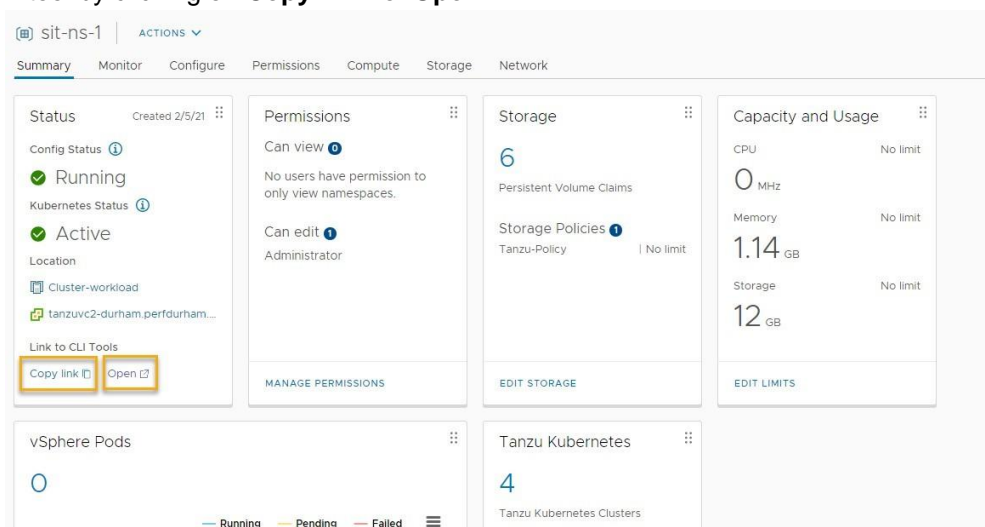


5. Click on **NEW NAMESPACE**

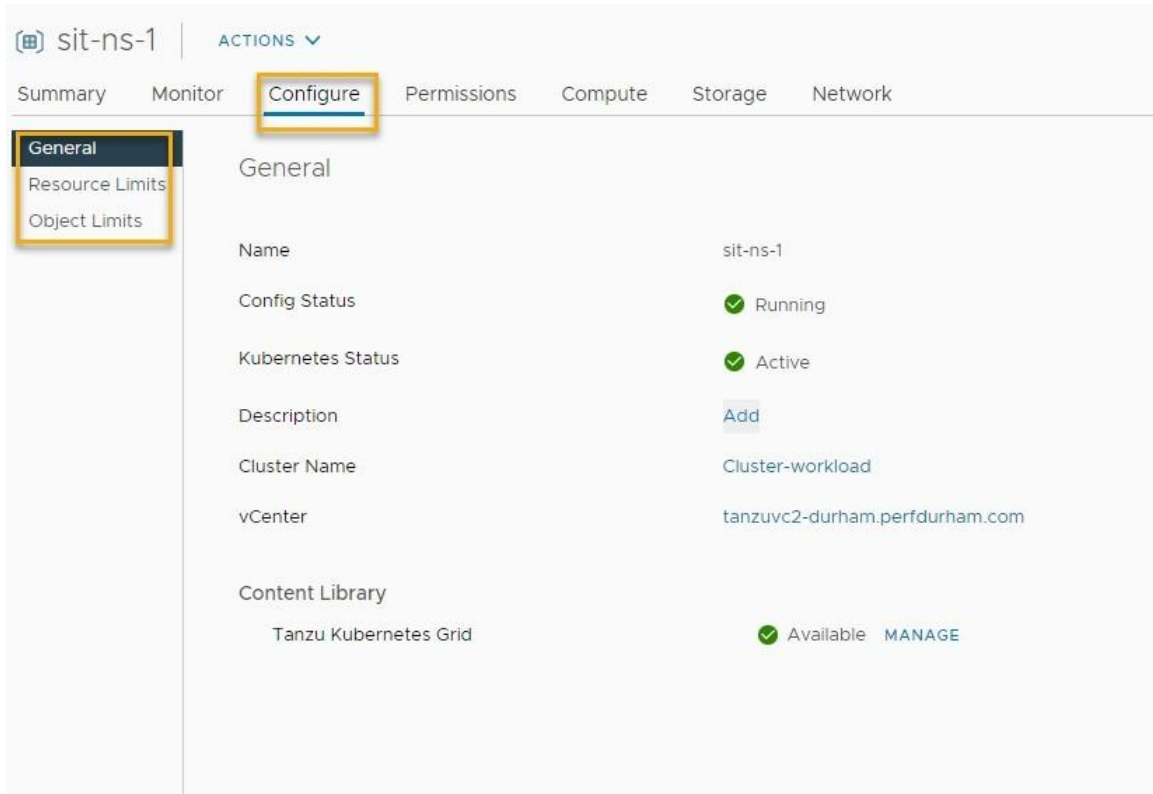
6. Select a cluster where you would like to create the namespace
7. Provide the name



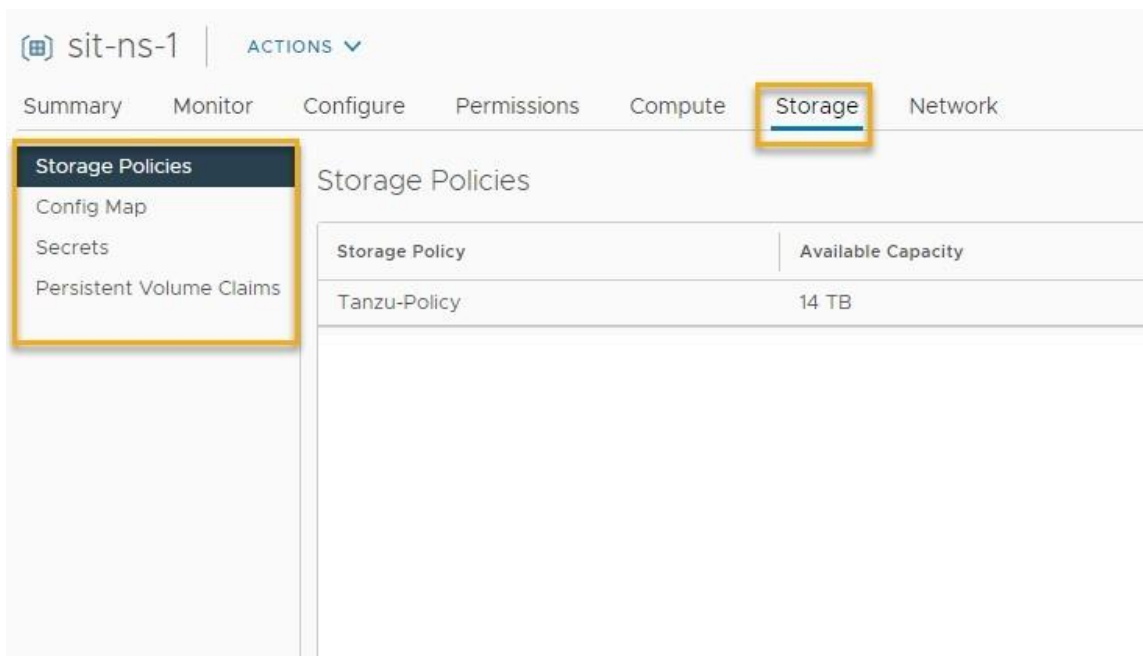
8. Click on **CREATE**
9. The namespace has been created and available under **Menu > Workload Management > Namespaces**
10. To access the namespace, you must have Kubernetes CLI tool installed as plugin. You can get that CLI tool by clicking on **Copy Link** or **Open**.



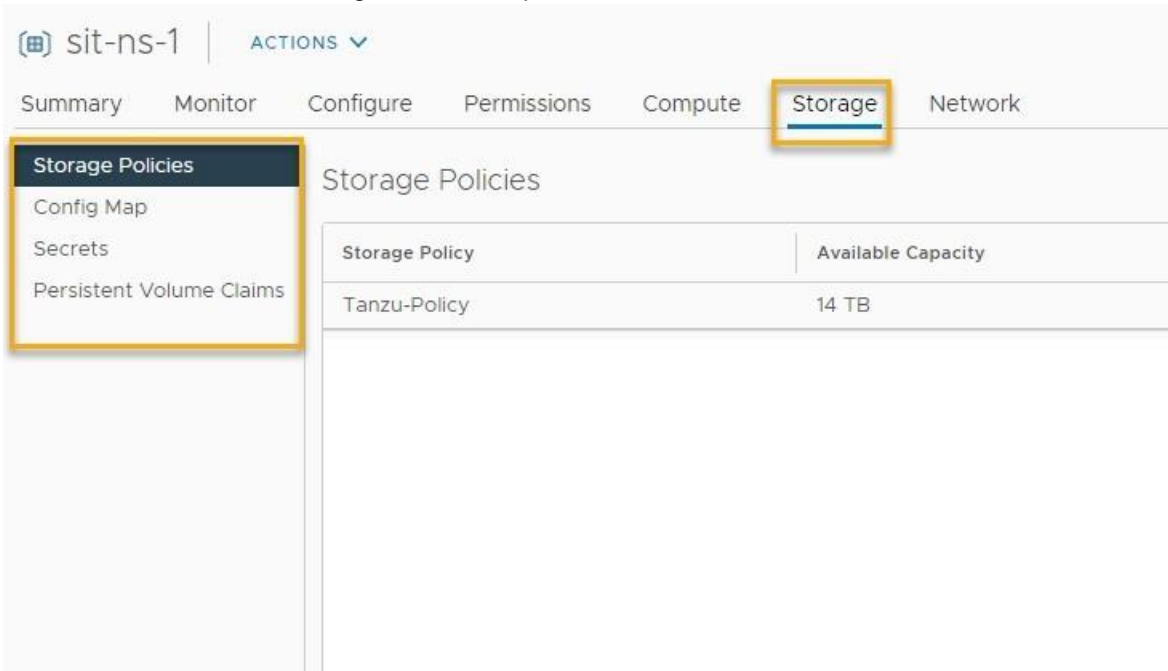
11. The resource limits and Object limits information are available through vCenter server under **Configure** section.



12. Storage Policies, Config Map, Secrets and Persistent Volume Claims are available under **Storage** section.



13. Network Policies, Services, Ingress and Endpoint information is available under **Network** section.



14. Download CLI plugin as per the operating system.

Kubernetes CLI Tools

Kubectl + vSphere plugin

Download the CLI tools package to view and control namespaces in vSphere. [LEARN MORE](#)

SELECT OPERATING SYSTEM

DOWNLOAD CLI PLUGIN WINDOWS

Checksum CLI plugin Windows



Get started with CLI Plugin for vSphere

Kubernetes CLI tool lets you manage your namespaces. Below are a few steps that will help you get started.

1. Verify that the SHA256 checksum of `vsphere*-plugin.zip` matches the checksum in the provided file `sha256sum.txt`. In Powershell run command `Get-FileHash -Algorithm SHA256 -Path vsphere*-plugin.zip` to display the checksum
2. Put the contents of the .zip file in your OS's executable search path
3. Run command `kubectl vsphere login --server=<IP_or_master_hostname>` to log in to server
4. Run command `kubectl config get-contexts` to view a list of your Namespaces
5. Run command `kubectl config use-context <context>` to choose your default context

15. You can access the namespaces and create guest clusters using CLI tool.

4.4 Create Tanzu Kubernetes Cluster (Guest Cluster)

Tanzu Kubernetes cluster is created by invoking Tanzu Kubernetes Grid service declarative API. Once the cluster is created, you can manage and deploy workloads to it by using `kubectl` command.

Steps to create TKG clusters

1. Download and install [Kubernetes CLI tool](#) for vSphere as mentioned in previous section.
2. Login to the namespace context using below command.

```
kubectl-vsphere.exe login --insecure-skip-tls-verify --vsphere-username  
<USERNAME> --server=<ip-address>
```

3. Verify the control plane and storage class.

```
kubectl get nodes  
kubectl get sc  
kubectl get virtualmachineimages
```

4. Switch context to the supervisor Namespace where you decide to provision Tanzu Kubernetes Cluster.

```
kubectl config get-contexts  
kubectl config use-context <SUPERVISOR-NAMESPACE>
```

5. Construct the YAML file for provisioning Tanzu Kubernetes Cluster and save it as *<cluster-name.yaml>*. The storageClass is populated with the previously configured storage policy and is backed by a datastore. For example:

```
apiVersion: run.tanzu.vmware.com/v1  
kind: TanzuKubernetesCluster  
metadata:  
  name: tkg-cluster-01  
  namespace: test-ns-1  
spec:  
  distribution:  
    version: v1.18.5  
  topology:  
    controlPlane:  
      count: 1  
      class: best-effort-small  
      storageClass: vwk-storage-policy  
    workers:  
      count: 3  
      class: best-effort-small  
      storageClass: vwk-storage-policy
```


6. Provision the cluster by running *<apply>* command.

```
kubectl apply -f <cluster-name>.yaml
```

7. Verify the cluster provisioned using below commands.

```
kubectl get tanzukubernetesclusters  
kubectl describe tanzukubernetescluster CLUSTER-NAME
```


8. At step 5, yaml file describes control plane count is 1 and worker nodes are 3. This can be verified at vCenter under **Namespaces**



5 Configure PowerProtect Data Manager to protect Tanzu Kubernetes workloads

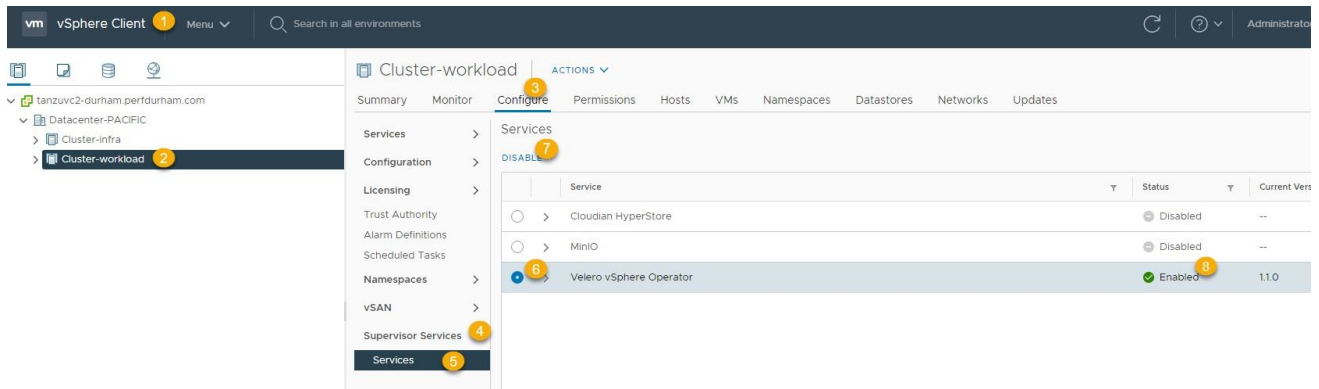
5.1 Prerequisites to Tanzu Kubernetes Guest Cluster Protection

Before adding a Tanzu Kubernetes guest cluster in PowerProtect Data Manager for namespace and PVC protection, the following one-time configuration is required to set up the Supervisor cluster:



5.1.1 Steps to Enable Velero vSphere Operator, Changed Block Tracking (CBT) and to install Velero and Velero-vSphere plugin

1. Login to the **vCenter server** with administrator credentials
2. Select the workload cluster
3. Click on **Configure**
4. Expand **Supervisor Services**
5. Click on **Services**
6. Select **Velero vSphere Operator**
7. Click on **Enable**



8. Verify if Velero vSphere Operator service is **Enabled**. Once enabled, the new Kubernetes namespace gets created automatically with its own vSphere pods running with Supervisor affinity. This allows the Supervisor cluster to perform backups using the FCD snapshot.
9. Create a new namespace "velero" in workload cluster. Velero components will be installed in this namespace.
10. Select **Menu > Workload Management** to view the namespaces running in the Supervisor cluster. For a selected namespace, click the **Compute** tab in the right pane to display the Tanzu guest clusters.
11. Download the command line binary velero-vsphere from the following location:
<https://github.com/vmware-tanzu/velero-plugin-for-vsphere/releases/download/v1.1.0/velero-vsphere-1.1.0-linuxamd64.tar.gz>

12. Use the Velero vSphere command line to enable changed block tracking (CBT) in the guest clusters:

```
# velero-vsphere configure --enable-cbt-in-guests
```

Once enabled, this setting is applied to the current cluster and all incoming guest clusters.

13. Using the same command line, install Velero and the Velero plug-in for the vSphere Client:

```
# velero-vsphere install --namespace velero --plugins vsphereveleroplugin/velero-  
pluginfor-vsphere:1.1.0 --no-secret --no-default-backup-location --use-volume-  
snapshots=false
```

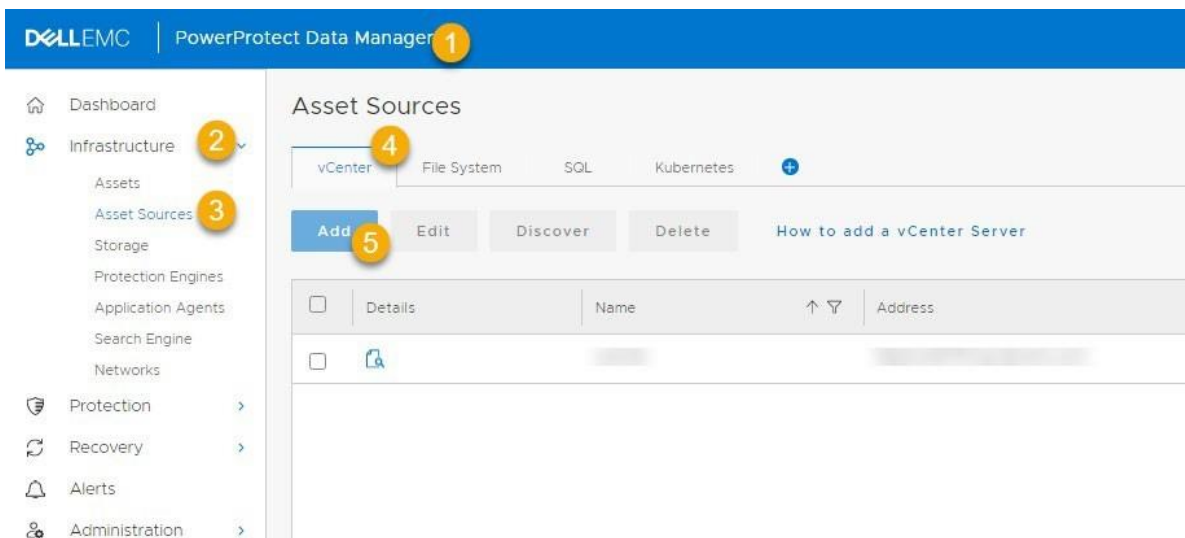
5.2 Asset Discovery

PowerProtect Data Manager discovers the Tanzu Kubernetes clusters by using IP address or Fully Qualified Domain Name (FQDN) of the vCenter server where the Tanzu Kubernetes clusters reside. PowerProtect Data Manager uses the discovery service account and the token (kubeconfig file) to integrate with kube-APIServer of Kubernetes cluster.

The Tanzu Kubernetes clusters are hosted by vCenter server and hence the vCenter server should be registered with PowerProtect Data Manager. PowerProtect Data Manager discovers the namespaces as assets.

Add vCenter as Asset Source

1. Login to **PowerProtect Data Manager UI** with administrator credentials.
2. On the left pane of the PowerProtect Data Manager UI, click **Infrastructure**.
3. Click on **Asset Sources**



4. Select **vCenter**
5. Click on **Add**
 - **Name:** Specify the name
 - **FQDN/IP:** Specify the IP address or fully qualified domain name
 - **Port:** 443
 - **Host Credentials:** Select credentials from the drop down if already added, if not, click on **Add Credentials**
 - Provide the FQDN of vCenter server
 - Username

- Password
- **vSphere Plugin:** Check box to install vSphere plugin
- **Scheduled Discovery:** This is optional and toggle if you need to specify automated discovery at given time schedule
- **Certificate:** Click on verify to authenticate

Edit vCenter ✕

Name

FQDN/IP

Port

Host Credentials

Schedule Discovery

Discovery Time (hour) (minute)

Certificate ✔ **Verified** C=US, CN=i

Cancel

Save

6. Click on Save

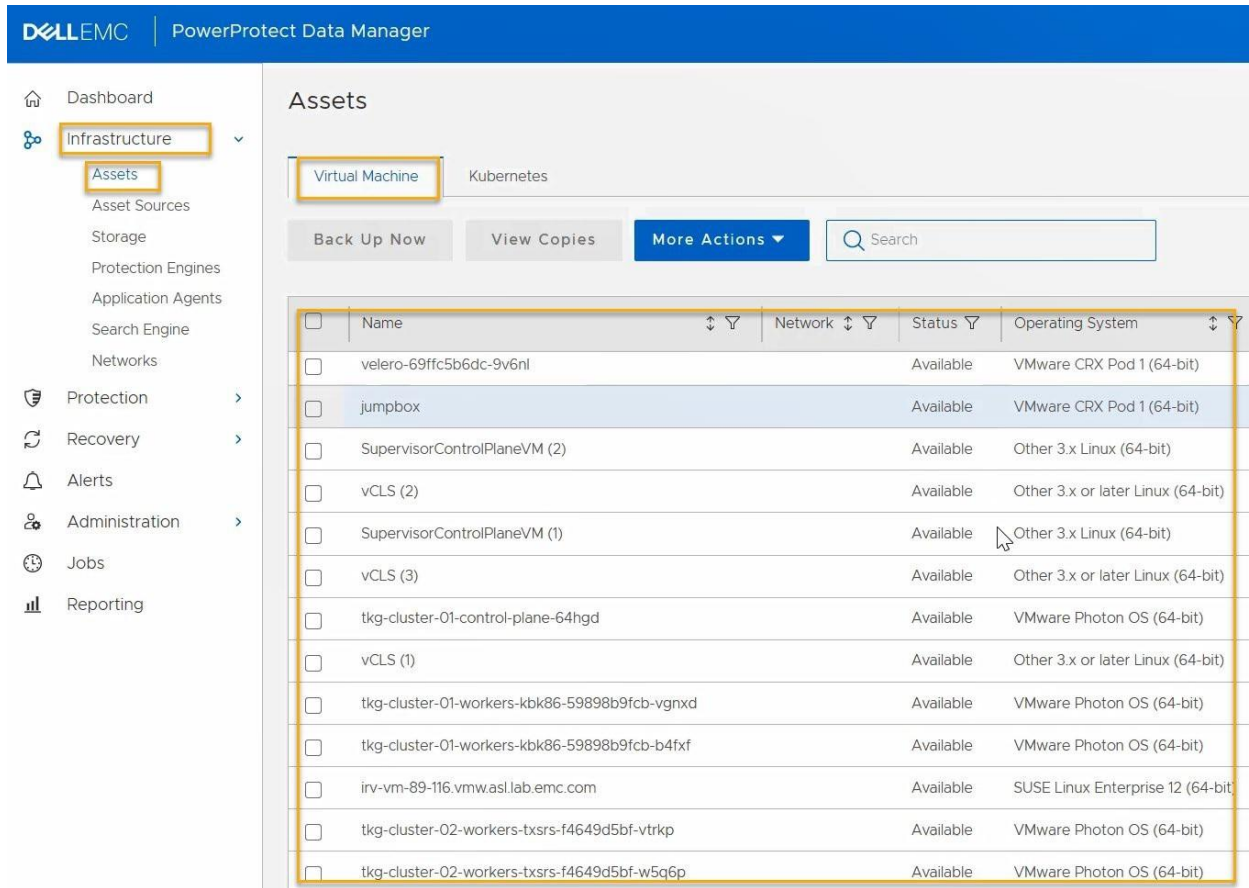
Asset Sources

vCenter Kubernetes +

Add Edit Discover Delete [How to add a vCenter Server](#)

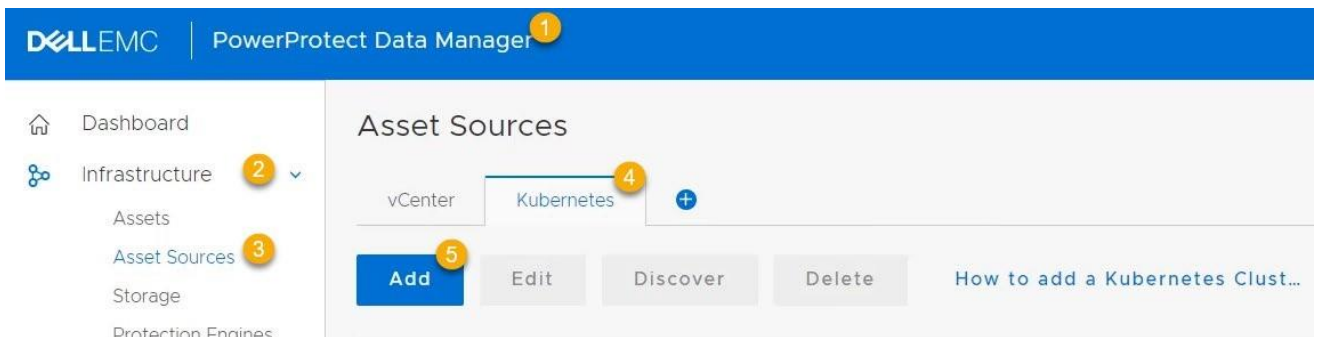
<input checked="" type="checkbox"/>	Details	Name	Address	PPDM Host	Version	Discovery Status
<input checked="" type="checkbox"/>		irv-vm-89-95.vmw.asl.lab.emc.com	irv-vm-89-95.vmw.asl.lab.emc.com	No	7.0.1	Ok

- vCenter server is successfully registered with PowerProtect Data Manager and is available under **Infrastructure > Asset Sources**



Add vSphere with Tanzu TKG Kubernetes clusters

- Login to **PowerProtect Data Manager UI** with administrator credentials
- On the left pane of the PowerProtect Data Manager UI, click **Infrastructure**
- Click on **Asset Sources**
- Click on **Kubernetes**
- Click on **Add**



Specify the details for the Tanzu Kubernetes Cluster

Tanzu Cluster

Select vCenter i

Name

FQDN/IP

Port

Host Credentials

Schedule Discovery

Discovery Time (hour) (minute)

Certificate

- **Tanzu Cluster:** Click on **Toggle button** to enable Tanzu Cluster
- **Select vCenter:** Select the vCenter that contains the Tanzu Kubernetes Cluster
- **Name:** Specify the name/IP address
- **FQDN/IP:** Provide the IP address or fully qualified domain name of the cluster

Note: Provide the IP address of the Tanzu Kubernetes Cluster's load balancer IP address. This can be obtained from vCenter at **Namespaces-->Namespace->Compute-->VMware Resources-->Tanzu Kubernetes-->Control Plane Address**.

Use the IP address and provide the token from the service account with required privileges. Service account token is provided in the Host credentials section.

- **Port:** 6443 (can be changed as per the configuration)
- **Host Credentials:** On the **Host Credentials**, click **Add** to add the service account token for the Kubernetes cluster, and click **Save**.

The service account must have the following privileges:

- Get/Create/Update/List CustomResourceDefinitions
- Get/Create/Update ClusterRoleBinding for 'cluster-admin' role
- Create/Update 'powerprotect' namespace
- Get/List/Create/Update/Delete all kinds of resources inside 'powerprotect' namespace
- Get/List/Watch all namespaces in the cluster as well as PV, PVC, and pods in all these namespaces

Note: A service account can be created that is bound to a cluster role that contains these privileges, and then provide the token of this service account. The cluster-admin role contains all these privileges.

- **Scheduled Discovery:** This is optional and toggle if you need to specify automated discovery at given time schedule.
- **Certificate:** Click on Verify to authenticate the credentials
- Click on **Save** to add Tanzu Kubernetes cluster

The screenshot shows the Dell EMC PowerProtect Data Manager interface. The left sidebar contains navigation options: Dashboard, Infrastructure (selected), Assets (selected), Asset Sources, Storage, Protection Engines, Application Agents, Search Engine, Networks, Protection, Recovery, Alerts, Administration, Jobs, and Reporting. The main content area is titled 'Assets' and has a sub-tab 'Kubernetes' selected. Below the sub-tab are buttons for 'Back Up Now', 'View Copies', and 'More Actions', along with a search box. A table displays a list of namespaces:

<input type="checkbox"/>	Details	Namespace ↑ ▾	Network ↓ ▾	Status ▾	Labels ↓ ▾	Age ↓ ▾	Protection
<input type="checkbox"/>		dk-k8s-4		Available		2 weeks	PLC-test-
<input type="checkbox"/>		dk-k8s-5		Available		1 month	PLC-tkg0
<input type="checkbox"/>		dk-k8s-5		Available		2 weeks	PLC-test-
<input type="checkbox"/>		dk-k8s-6		Available		1 month	PLC-tkg0
<input type="checkbox"/>		dk-k8s-6		Available		2 weeks	PLC-test-
<input type="checkbox"/>		dk-k8s-7		Available		1 month	PLC-tkg0
<input type="checkbox"/>		dk-k8s-7		Available		2 weeks	PLC-test-
<input type="checkbox"/>		dk-k8s-8		Available		1 month	PLC-tkg0
<input type="checkbox"/>		dk-k8s-8		Available		2 weeks	PLC-test-
<input type="checkbox"/>		dk-k8s-9		Available		1 month	PLC-tkg0
<input type="checkbox"/>		dk-k8s-9		Available		2 weeks	PLC-test-

6. Verify that the Tanzu Kubernetes Cluster is added successfully, and the namespaces are available under **Assets > Kubernetes**

Note: The vCenter hosting Tanzu Cluster is registered with PowerProtect Data Manager and the namespaces are available to be protected. At the time of integration between PowerProtect Data Manager and Tanzu Kubernetes cluster, **velero-ppdm** and **powerprotect**, the two namespaces are created on Tanzu Kubernetes Cluster.

kube-public	Active	16d
kube-system	Active	16d
mysql-alter-1	Active	12d
mysql-ns-1	Active	14d
mysql-ns-multi-3	Active	14d
powerprotect	Active	3d3h
recover-dk-k8s-8	Active	11d
tkg01-new	Active	5d9h
velero-ppdm	Active	3d3h
velero-vsphere-plugin-backupdriver	Active	14d
vmware-system-auth	Active	16d
vmware-system-cloud-provider	Active	16d
vmware-system-csi	Active	16d

The **powerprotect** namespace contains powerprotect-controller pod and **velero-ppdm** namespace contains velero pod and backup-driver pod. The backup-driver pod is responsible to take and delete snapshots for paravirtualized CSI volumes of Tanzu Kubernetes Clusters.

5.3 VM Direct Engine

VM direct is the protection engine within PowerProtect Data Manager which is used to manage and protect VM assets. VM direct engine can be created from PowerProtect Data Manager UI.

1. Log in to **PowerProtect Data Manager UI** with administrator credentials.
2. On the left pane of the PowerProtect Data Manager UI, click on **Infrastructure**
3. Click on **Protection Engines**
4. Click on **Add** to add new protection engine
5. Specify the details for the protection engine
 - **Hostname:** Provide the hostname
 - **Gateway:** Specify the Gateway
 - **IP address:** Provide the valid IPv4 or IPv6 address
 - **Netmask:** Provide the valid netmask address
 - **vCenter to Deploy:** Select the vCenter on which the VM is deployed
 - **ESX Host/ Cluster:** Select the ESX host or cluster on which the VM is deployed
 - **Supported Protection Type:** Click on the dropdown and select **Kubernetes**
 - **Primary DNS:** Mention the primary DNS
 - **Secondary DNS:** Mention the Secondary DNS (if any)
 - **Tertiary DNS:** Mention the third DNS (if any)
 - **Network:** Click on dropdown and select the required VM Network
 - **Data Store:** Select the datastore
 - **Transport Mode:** Select the required transport mode i.e. hot add, Network Block or Hot add, Network Block Device
6. Specify the details about network configuration
7. Verify the summary and click on **Save**
8. The available VM direct engines are visible under VM Direct Protection Engines tab

5.4 Backup Configuration

Backup of an asset can be taken manually as well as can be scheduled, a protection policy must be created to run the backup.

1. Login to **PowerProtect Data Manager UI** with administrator credentials.
2. On the left pane of the PowerProtect Data Manager UI, click on **Protection** dropdown
3. Click on **Protection Policies**
4. Click on **Add** to add new policy
 - **Type:** Specify the necessary details as
 - a. Name:
 - b. Description:
 - c. Type: Select **Kubernetes**
 - d. Click **Next**
 - **Purpose:** Select the option from below according to the purpose of the backup,
 - a. **Crash Consistent:** Select this option to snapshot persistent volumes bound to the persistent volume claims in the namespace and back them up to the storage target
Or
 - b. **Exclusion:** Select this option to exclude in this group from protection activities and protection rule assistant.
 - c. Click **Next**
 - **Assets**
 - a. Select the Asset (s) to be backed up from the list or particular asset can be searched by typing in **Find More Assets** box
 - b. Click **Next**
 - **Schedule:** In this section, a Backup schedule is created, Replicated, Edited and Deleted
 - a. Click on **+Backup**
 - b. Fill the required details under **Add Primary Backup** section
 - Select **Recurrence** as Hourly, Daily, Weekly or Monthly and fill the details accordingly
 - Create Every:
 - Keep For:
 - Start time:
 - End Time:
 - Click **Ok**
 - c. The schedule is ready, SLA can be added as per the requirement in this section
 - Click on dropdown under **SLA**
 - Click on **Add** and fill the Name and select the objective (s)
 - i. SLA name:
 - ii. Check box to specify **Recovery Point Option** (minutes, hours, days, weeks, months or years)
 - iii. Check box if any **Compliance Window** is to be mentioned
 - iv. Check box to **Verify expired copies are deleted**
 - v. Check box to specify **Retention Time Objective** (days, weeks, months or years)
 - vi. Check box to **Verify Retention Lock is enabled for all copies**
 - vii. Click on **Save**
5. Click **Next**
6. Verify the provided information is correct under **Summary** section. If yes, click on **Finish**
7. Click on **Go to Jobs** to check the status of the policy.

The Protection policy triggers the backup at the scheduled time. When the protection policy is created successfully. There are options to modify the existing policy i.e. **Edit, Disable, Export** and **Protect Now**

1. **Edit:** To edit the information or to change the schedule
2. **Disable:** Backup Schedule is disabled with this option so backup would not be taken
3. **Export:** Downloadable file which contains the information about the asset protection
4. **Protect Now:** This Option allows to take a backup manually at ad-hoc basis.
 - **Asset Selection:** It has further two options to select the assets i.e.
 - a. All assets defined in the protection policy
 - b. **Choose some of the assets defined in the policy:** This option allows to select namespaces within the cluster
 - **Configuration:**
 - a. This allows to **select type of backup**
 - **Full:** Backs up the namespace metadata and persistent volumes and creates a new full backup
 - **Synthetic full:** Backs up namespace metadata, change blocks for persistent volumes on VMware first class disks, all other persistent volumes and creates a new full backup
 - b. Keep For: In days
 - c. Click **Next**
 - **Summary:** Verify the information
 - Click on **Backup**
 - Monitor the backup job by Clicking **Go to Jobs**

Manually Protecting Kubernetes - TME_Pacific - PROTECTION - Full

The screenshot shows the backup job progress and details. At the top, a summary box displays: 1 Assets, 42.0 MB Data Transferred, and 00:01:52 Duration. To the right, performance metrics are listed: Storage Target (irv-12-196.as1.lab.emc.com), Data Size (42.0 MB), Average Throughput (335.7 Mbit/s), Reduction % (99.6%), Next Scheduled, Data Change (41.9 MB), Total Compression Factor (227.1x), and Factor.

Below the summary, a status bar shows 0 Critical, 0 Warning, 1 Success, 0 Cancelled, and 1 Total. Action buttons include Restart, Export Log, and Cancel, along with a search field.

The main table lists the backup job details:

Details	Task Name	Status	Asset Name	Start Time
	Protecting Kubernetes - dk-k8s-demo	Success	dk-k8s-demo	Dec 21, 2020, 10:34:24

The right-hand pane shows the job steps and a summary of the backup parameters:

```

Step 1: Protect KUBERNETES for dk-k8s-demo (00:01:24)
Summary:
{
  "parameters": {
    "assetName": "dk-k8s-demo",
    "storage": "irv-12-196.as1.lab.emc.com",
  }
}
  
```

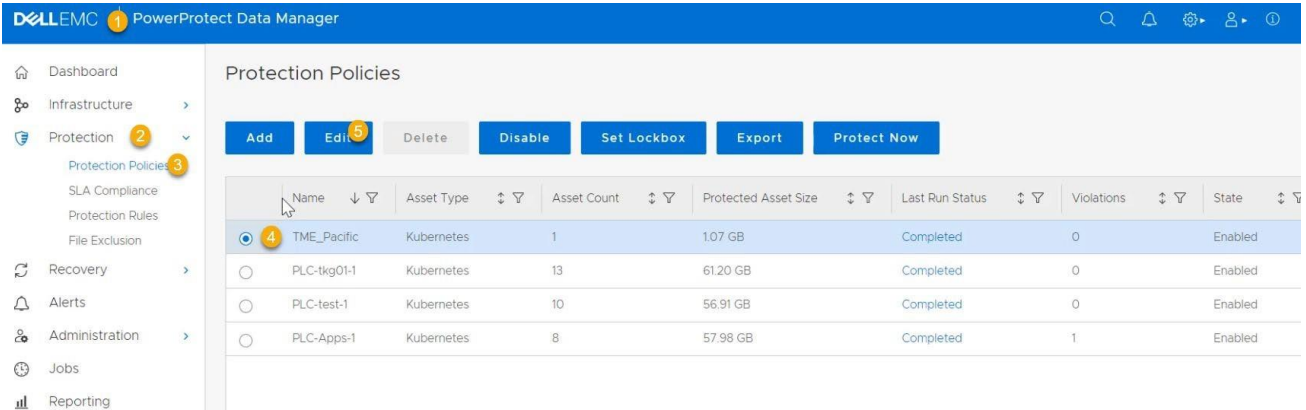
The backup job completed successfully, and the details like task ID, vProxy, throughput, storage etc are available in the backup job details.

5.5 Replication Configuration

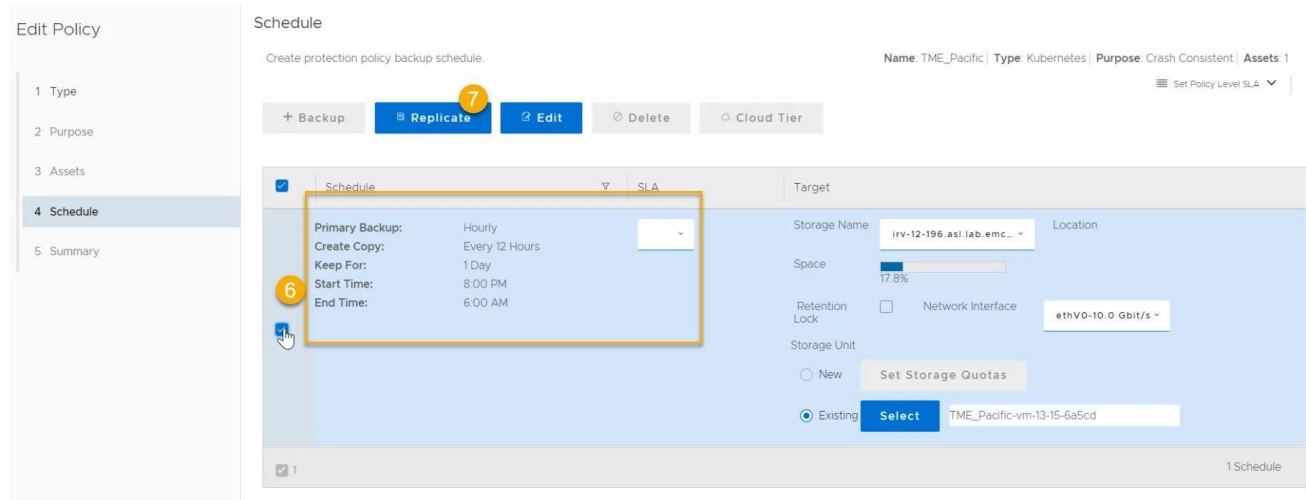
The replication is configured with existing backup policy or can be created a new policy.

Steps to configure replication on existing backup policy

1. Login to **PowerProtect Data Manager UI** with administrator credentials.
2. On the left pane of the PowerProtect Data Manager UI, click on **Protection** dropdown
3. Click on **Protection Policies**

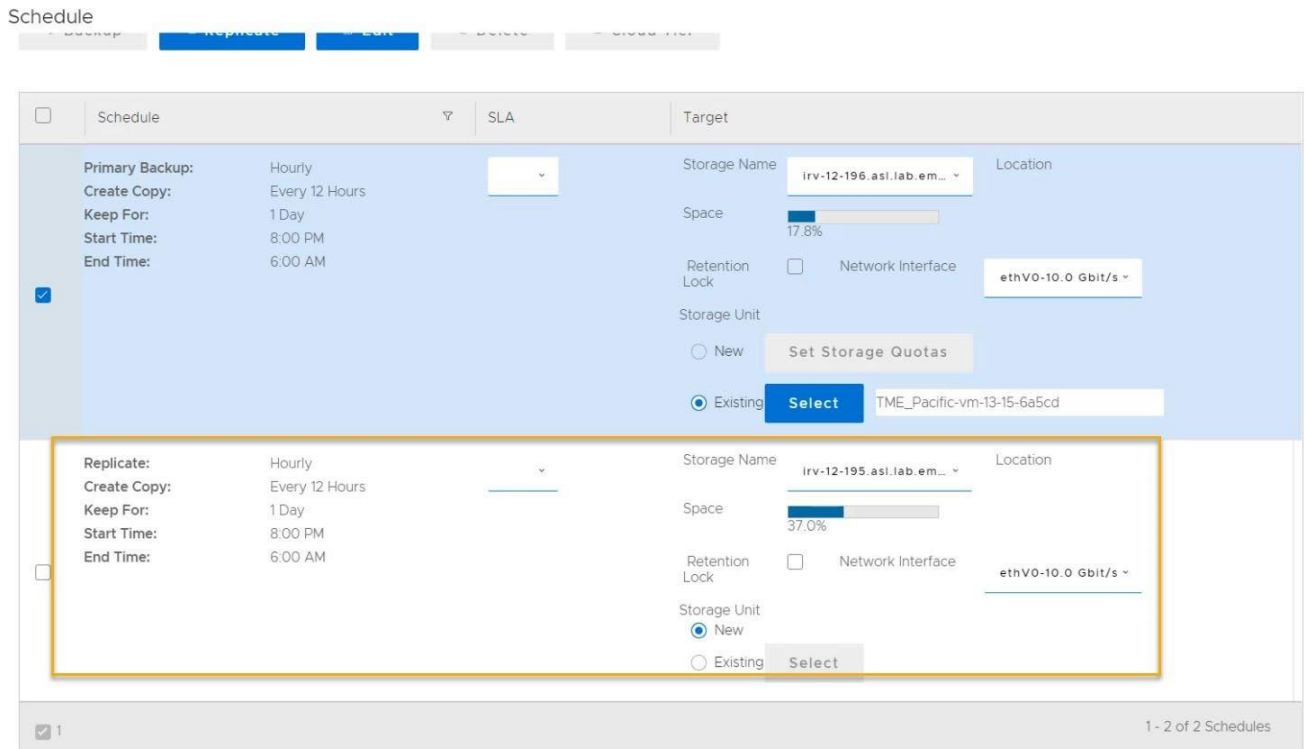


4. Select an existing backup policy
5. Click on **Edit** and Click on **Back** button



6. Select primary backup **Schedule**
7. click on **Replicate**
8. Add Primary Replication details
 - **Replicate Every:** Provide the appropriate time
 - **Keep for:** Specify the number of days
 - **Start Time**
 - **End Time**

9. Click **OK**

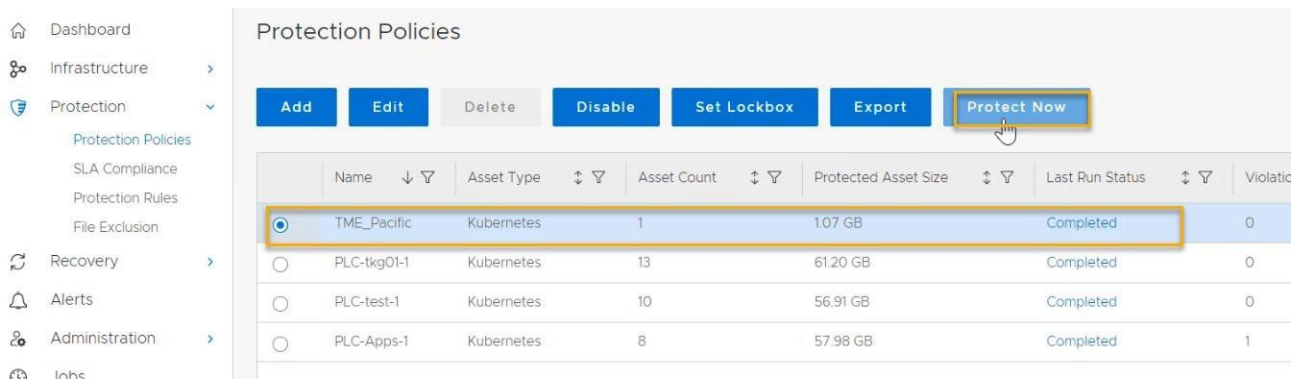


10. Observe that replicate schedule is created.

11. Click on **Finish**

12. Verify the replication job is created under **Jobs** section


13. To run the replication now, select existing policies from **Protection > Protection Policies**



14. Click on **Protect Now**

- **Asset Selection:** Choose one option for ad-hoc protection
- **Configuration:** Select **Replicate Now** option and check box to select replication storage
- **Summary:** Click on **Protect Now** to start replication

15. Click on **Go to Jobs** for the progress and the details



16. Verify the replication job is successfully completed, click on details button for detailed results.

5.6 Restore Configuration

The Recovery of the assets is a manual process. The restoration of the Kubernetes namespaces includes pods, stateful sets, PVCs and other resources is done at the same cluster. With PowerProtect Data Manager, there are options to recover the Kubernetes namespaces to the **same** as well as an **alternate cluster**.

1. Log in to **PowerProtect Data Manager UI** with admin credentials.
2. On the left pane of the PowerProtect Data Manager UI, click on **Recovery**
3. Click on **Assets**
4. Click on **Kubernetes** on top and select the namespaces to be restored
5. Click on **Restore**
 - **Select Copy:**
 - a. Select the restore copy. The most recent copy will be used as default. To change from default copy, click on **Change Copy**
 - b. Click **Ok**
 - c. Click **Next**
 - **Cluster:** This provides the option to select the cluster on which assets to be restored
 - a. **Restore to Original Cluster:** The Assets are restored to the source cluster from which the backup is taken.
 - b. **Restore to Alternate Cluster:** The Assets are restored on the alternate cluster. To utilize this option, alternate cluster is added as an asset source to the managing PowerProtect Data Manager.
 - **Purpose:** Select the option what is to be restored
 - a. **Restore Namespace and Select PVCs:** This option restores the namespace and a subset of PVCs in the namespace. The namespace resources including pods, services, secrets and deployments will not be overwritten during a restore. All other resources that do not currently exists in the namespace will be restored.
 - b. **Restore PVCs only:** This Option will restore only PVCs
 - **Restore Type:** Restore type has different options depending on the **purpose** of the restore
 - a. If purpose is to Restore namespaces and PVCs, then options are
 - Restore to Original Namespace,
 - Restore to New Namespace and
 - Restore to an Existing Namespace
 - b. And if purpose if to restore PVCs only, then options are
 - Restore to Original Namespace and
 - Restore to an Existing Namespace
 - **PVCs:** Select PVCs to be restored to the namespace, options are to
 - a. Overwrite content of existing PVCs
 - b. Skip restores of existing PVCs
 - **Summary:** Verify the information and click on **Restore**

Note: When the restore job starts, a new namespace is created at Kubernetes cluster. Once the restore job completes, pod and PVC get created with details of pod and Persistent Volume bound to PVC.

A Technical support and resources

[Dell.com/support](https://www.dell.com/support) is focused on meeting customer needs with proven services and support.

[Storage technical documents and videos](#) provide expertise that helps to ensure customer success on Dell Technologies storage platforms.

A.1 Related resources

- vSphere with Tanzu brief: <https://d1fto35gciffzn.cloudfront.net/tanzu/VMware-Tanzu-Basic-Solution-Brief.pdf>
- vSphere with Tanzu configuration: <https://docs.vmware.com/en/VMware-vSphere/7.0/vmware-vsphere-with-tanzu/GUID-152BE7D2-E227-4DAA-B527-557B564D9718.html>
- Storage policy creation with vSphere : <https://docs.vmware.com/en/VMware-vSphere/7.0/vmware-vsphere-with-tanzu/GUID-544286A2-A403-4CA5-9C73-8EFF261545E7.html#GUID-544286A2-A403-4CA5-9C73-8EFF261545E7>
- Kubernetes CLI tool <https://docs.vmware.com/en/VMware-vSphere/7.0/vmware-vsphere-with-tanzu/GUID-0F6E45C4-3CB1-4562-9370-686668519FCA.html>
- Dell EMC PowerProtect Data Manager Data Sheet: <https://www.dellemc.com/en-me/collaterals/unauth/data-sheets/products/data-protection/h17691-dellemc-powerprotect-software-ds.pdf>
- Dell EMC PowerProtect Data Manager: <https://www.delltechnologies.com/en-in/data-protection/powerprotect-data-manager.htm#scroll=off>