



VMware Tanzu with NetApp

NetApp container solutions

NetApp
August 18, 2025

This PDF was generated from <https://docs.netapp.com/us-en/netapp-solutions-containers/tanzu/vtwn-solution-overview.html> on August 18, 2025. Always check docs.netapp.com for the latest.

Table of Contents

- VMware Tanzu with NetApp 1
 - NVA-1166: VMware Tanzu with NetApp 1
 - Use cases 1
 - Business value 1
 - Technology overview 2
 - Current support matrix for validated releases 3
 - VMware Tanzu product portfolio 4
 - VMware Tanzu overview 4
 - VMware Tanzu Kubernetes Grid (TKG) overview 5
 - VMware Tanzu Kubernetes Grid Service (TKGS) overview 5
 - VMware Tanzu Kubernetes Grid Integrated Edition (TKGI) overview 7
 - VMware vSphere with Tanzu overview 8
 - NetApp storage systems 10
 - NetApp storage systems overview 10
 - NetApp ONTAP 11
 - NetApp storage integrations 14
 - NetApp storage integration overview 14
 - NetApp Trident 15
 - Trident overview 15
 - NetApp ONTAP NFS configuration 18
 - NetApp ONTAP iSCSI configuration 23
 - Additional Information: VMware Tanzu with NetApp 28

VMware Tanzu with NetApp

NVA-1166: VMware Tanzu with NetApp

Alan Cowles and Nikhil M Kulkarni, NetApp

This reference document provides deployment validation of different flavors of VMware Tanzu Kubernetes solutions, deployed either as Tanzu Kubernetes Grid (TKG), Tanzu Kubernetes Grid Service (TKGS), or Tanzu Kubernetes Grid Integrated (TKGI) in several different data center environments as validated by NetApp. It also describes storage integration with NetApp storage systems and the Trident storage orchestrator for the management of persistent storage and Trident Protect for the backup and cloning of the stateful applications using that persistent storage. Lastly, the document provides video demonstrations of the solution integrations and validations.

Use cases

The VMware Tanzu with NetApp solution is architected to deliver exceptional value for customers with the following use cases:

- Easy to deploy and manage VMware Tanzu Kubernetes Grid offerings deployed on VMware vSphere and integrated with NetApp storage systems.
- The combined power of enterprise container and virtualized workloads with VMware Tanzu Kubernetes Grid offerings.
- Real world configuration and use cases highlighting the features of VMware Tanzu when used with NetApp storage and NetApp Trident suite of products.
- Application-consistent protection or migration of containerized workloads deployed on VMware Tanzu Kubernetes Grid clusters whose data resides on NetApp storage systems using Trident Protect.

Business value

Enterprises are increasingly adopting DevOps practices to create new products, shorten release cycles, and rapidly add new features. Because of their innate agile nature, containers and microservices play a crucial role in supporting DevOps practices. However, practicing DevOps at a production scale in an enterprise environment presents its own challenges and imposes certain requirements on the underlying infrastructure, such as the following:

- High availability at all layers in the stack
- Ease of deployment procedures
- Non-disruptive operations and upgrades
- API-driven and programmable infrastructure to keep up with microservices agility
- Multitenancy with performance guarantees
- Ability to run virtualized and containerized workloads simultaneously
- Ability to scale infrastructure independently based on workload demands
- Ability to deploy in a hybrid-cloud model with containers running in both on-premises data centers as well as in the cloud.

VMware Tanzu with NetApp acknowledges these challenges and presents a solution that helps address each concern by deploying VMware Tanzu Kubernetes offerings in the customer's choice of hybrid cloud environment.

Technology overview

The VMware Tanzu with NetApp solution is comprised of the following major components:

VMware Tanzu Kubernetes platforms

VMware Tanzu comes in a variety of flavors that the solutions engineering team at NetApp has validated in our labs. Each Tanzu release successfully integrates with the NetApp storage portfolio, and each can help meet certain infrastructure demands. The following bulleted highlights describe the features and offerings of each version of Tanzu described in this document.

VMware Tanzu Kubernetes Grid (TKG)

- Standard upstream Kubernetes environment deployed in a VMware vSphere environment.
- Formerly known as Essential PKS (from Heptio acquisition, Feb 2019).
- TKG is deployed with a separate management cluster instance for support on vSphere 6.7U3 onward.
- TKG deployments can be deployed in the cloud as well with AWS or Azure.
- Allows for use of Windows or Linux worker nodes (Ubuntu/Photon).
- NSX-T, HA Proxy, AVI networking, or load balancers can be used for control plane.
- TKG supports MetalLB for the application/data plane.
- Can use vSphere CSI as well as third party CSIs like NetApp Trident.

VMware Tanzu Kubernetes Grid Service (TKGS)

- Standard upstream Kubernetes environment deployed in a VMware vSphere environment.
- Formerly known as Essential PKS (from Heptio acquisition, Feb 2019).
- TKGS deployed with supervisor cluster and workload clusters only on vSphere 7.0U1 onward.
- Allows for use of Windows or Linux worker nodes (Ubuntu/Photon).
- NSX-T, HA Proxy, AVI networking, or load balancers can be used for control plane.
- TKGS supports MetalLB for application/data plane.
- Can use vSphere CSI as well as third party CSIs like NetApp Trident.
- Provides support for vSphere Pods with Tanzu, allowing pods to run directly on enabled ESXi hosts in the environment.

VMware Tanzu Kubernetes Grid Integrated (TKGI)

- Formerly known as Enterprise PKS (from Heptio acquisition, Feb 2019).
- Can use NSX-T, HA Proxy, or Avi. You can also provide your own load balancer.
- Supported from vSphere 6.7U3 onward, as well as AWS, Azure, and GCP.
- Setup via wizard to allow for ease of deployment.
- Runs Tanzu in controlled immutable VMs managed by BOSH.

- Can make use vSphere CSI as well as third party CSIs like NetApp Trident (some conditions apply).

vSphere with Tanzu (vSphere Pods)

- vSphere-native pods run in a thin, photon-based layer with prescribed virtual hardware for complete isolation.
- Requires NSX-T, but that allows for additional feature support such as a Harbor image registry.
- Deployed and managed in vSphere 7.0U1 onward using a virtual Supervisor cluster like TKGS. Runs pods directly on ESXi nodes.
- Fully vSphere integrated, highest visibility and control by vSphere administration.
- Isolated CRX-based pods for the highest level of security.
- Only supports vSphere CSI for persistent storage. No third-party storage orchestrators supported.

NetApp storage systems

NetApp has several storage systems perfect for enterprise data centers and hybrid cloud deployments. The NetApp portfolio includes NetApp ONTAP, NetApp Element, and NetApp e-Series storage systems, all of which can provide persistent storage for containerized applications.

For more information, visit the NetApp website [here](#).

NetApp storage integrations

Trident is an open-source, fully-supported storage orchestrator for containers and Kubernetes distributions, including VMware Tanzu.

For more information, visit the Trident website [here](#).

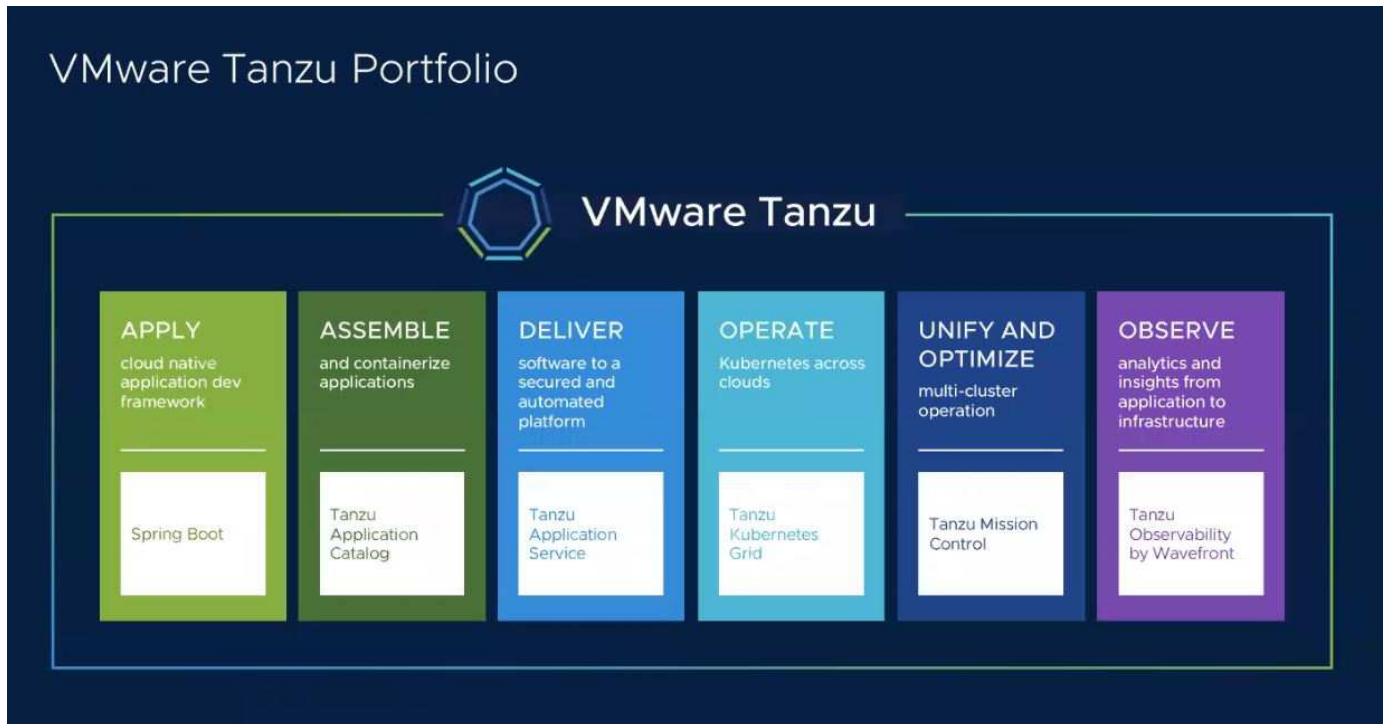
Current support matrix for validated releases

Technology	Purpose	Software version
NetApp ONTAP	Storage	9.9.1
NetApp Trident	Storage Orchestration	22.04.0
VMware Tanzu Kubernetes Grid	Container orchestration	1.4+
VMware Tanzu Kubernetes Grid Service	Container orchestration	0.0.15 [vSphere Namespaces]
		1.22.6 [Supervisor Cluster Kubernetes]
VMware Tanzu Kubernetes Grid Integrated	Container orchestration	1.13.3
VMware vSphere	Data center virtualization	7.0U3
VMware NSX-T Data Center	Networking and Security	3.1.3
VMware NSX Advanced Load Balancer	Load Balancer	20.1.3

VMware Tanzu product portfolio

VMware Tanzu overview

VMware Tanzu is a portfolio of products that enables enterprises to modernize their applications and the infrastructure they run on. VMware Tanzu's full stack of capabilities unites the development and IT operations teams on a single platform to embrace modernization in both their applications and their infrastructure consistently across on-premises and hybrid cloud environments to continuously deliver better software to production.



To understand more about the different offerings and their capabilities in the Tanzu portfolio, visit the documentation [here](#).

Regarding Tanzu's Kubernetes Operations catalog, VMware has a variety of implementations for Tanzu Kubernetes Grid, all of which provision and manage the lifecycle of Tanzu Kubernetes clusters on a variety of platforms. A Tanzu Kubernetes cluster is a full-fledged Kubernetes distribution that is built and supported by VMware.


NetApp has tested and validated the deployment and interoperability of the following products from the VMware Tanzu portfolio in its labs:

- [VMware Tanzu Kubernetes Grid \(TKG\)](#)
- [VMware Tanzu Kubernetes Grid Service \(TKGS\)](#)
- [VMware Tanzu Kubernetes Grid Integrated \(TKGI\)](#)
- [VMware vSphere with Tanzu \(vSphere Pods\)](#)

VMware Tanzu Kubernetes Grid (TKG) overview

VMware Tanzu Kubernetes Grid, also known as TKG, lets you deploy Tanzu Kubernetes clusters across hybrid cloud or public cloud environments. TKG is installed as a management cluster, which is a Kubernetes cluster itself, that deploys and operates the Tanzu Kubernetes clusters. These Tanzu Kubernetes clusters are the workload Kubernetes clusters on which the actual workload is deployed.

Tanzu Kubernetes Grid builds on a few of the promising upstream community projects and delivers a Kubernetes platform that is developed, marketed, and supported by VMware. In addition to Kubernetes distribution, Tanzu Kubernetes Grid provides additional add-ons that are essential production-grade services such as registry, load balancing, authentication, and so on. VMware TKG with management cluster is widely used in vSphere 6.7 environments, and, even though it is supported, it is not a recommended deployment for vSphere 7 environments because TKGS has native integration capabilities with vSphere 7.



For more information on Tanzu Kubernetes Grid, refer to the documentation [here](#).

Depending on whether the Tanzu Kubernetes Grid is being installed on-premises on vSphere cluster or in cloud environments, prepare and deploy Tanzu Kubernetes Grid by following the installation guide [here](#).


After you have installed the management cluster for Tanzu Kubernetes Grid, deploy the user clusters or workload clusters as needed by following the documentation [here](#). VMware TKG management cluster requires that an SSH key be provided for installation and operation of Tanzu Kubernetes clusters. This key can be used to log into the cluster nodes using the `capv` user.

VMware Tanzu Kubernetes Grid Service (TKGS) overview

VMware Tanzu Kubernetes Grid Service (also known as vSphere with Tanzu) lets you create and operate Tanzu Kubernetes clusters natively in vSphere and also allows you to run some smaller workloads directly on the ESXi hosts. It allows you to transform vSphere into a platform for running containerized workloads natively on the hypervisor layer. Tanzu Kubernetes Grid Service deploys a supervisor cluster on vSphere when enabled that deploys and operates the clusters required for the workloads. It is natively integrated with vSphere 7 and leverages many reliable vSphere features like vCenter SSO, Content Library, vSphere networking, vSphere storage, vSphere HA and DRS, and

vSphere security for a more seamless Kubernetes experience.

vSphere with Tanzu offers a single platform for hybrid application environments where you can run your application components either in containers or in VMs, thus providing better visibility and ease of operations for developers, DevOps engineers, and vSphere administrators. VMware TKGS is only supported with vSphere 7 environments and is the only offering in Tanzu Kubernetes operations portfolio that allows you to run pods directly on ESXi hosts.



For more information on Tanzu Kubernetes Grid Service, follow the documentation [here](#).

There are a lot of architectural considerations regarding feature sets, networking, and so on. Depending on the architecture chosen, the prerequisites and the deployment process of Tanzu Kubernetes Grid Service differ. To deploy and configure Tanzu Kubernetes Grid Service in your environment, follow the guide [here](#). Furthermore, to log into the Tanzu Kubernetes cluster nodes deployed via TKGS, follow the procedure laid out in this [link](#).

NetApp recommends that all the production environments be deployed in multiple master deployments for fault tolerance with the choice of worker nodes' configuration to meet the requirements of the intended workloads. Thus, a recommended VM class for a highly intensive workload would have at least four vCPUs and 12GB of RAM.

When Tanzu Kubernetes clusters are created in a namespace, users with `owner` or `edit` permission can create pods directly in any namespace by using the user account. This is because users with the `owner` or `edit` permission are allotted the cluster administrator role. However, when creating deployments, daemon sets, stateful sets, or others in any namespace, you must assign a role with the required permissions to the corresponding service accounts. This is required because the deployments or daemon sets utilize service accounts to deploy the pods.


See the following example of ClusterRoleBinding to assign the cluster administrator role to all service accounts in the cluster:


```
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  name: all_sa_ca
subjects:
- kind: Group
  name: system:serviceaccounts
  namespace: default
roleRef:
  kind: ClusterRole
  name: psp:vmware-system-privileged
  apiGroup: rbac.authorization.k8s.io
```

VMware Tanzu Kubernetes Grid Integrated Edition (TKGI) overview

VMware Tanzu Kubernetes Grid Integrated (TKGI) Edition, formerly known as VMware Enterprise PKS, is a standalone container orchestration platform based on Kubernetes with capabilities such as life cycle management, cluster health monitoring, advanced networking, a container registry, and so on. TKGI provisions and manages Kubernetes clusters with the TKGI control plane, which consists of BOSH and Ops Manager.

TKGI can be installed and operated either on vSphere or OpenStack environments on-premises or in any of the major public clouds on their respective IaaS offerings. Furthermore, the integration of TKGI with NSX-T and Harbour enables wider use cases for enterprise workloads. To know more about TKGI and its capabilities, visit the documentation [here](#).



TKGI is installed in a variety of configurations on a variety of platforms based on different use-cases and designs. Follow the guide [here](#) to install and configure TKGI and its prerequisites. TKGI uses Bosh VMs as nodes for Tanzu Kubernetes clusters which run immutable configuration images and any manual changes on Bosh VMs do not remain persistent across reboots.

Important notes:

- NetApp Trident requires privileged container access. So, during TKGI installation, make sure to select the Enable Privileged Containers checkbox in the step to configure Tanzu Kubernetes cluster node plans.

The screenshot displays the configuration interface for a Tanzu Kubernetes cluster node plan. It is organized into several sections:

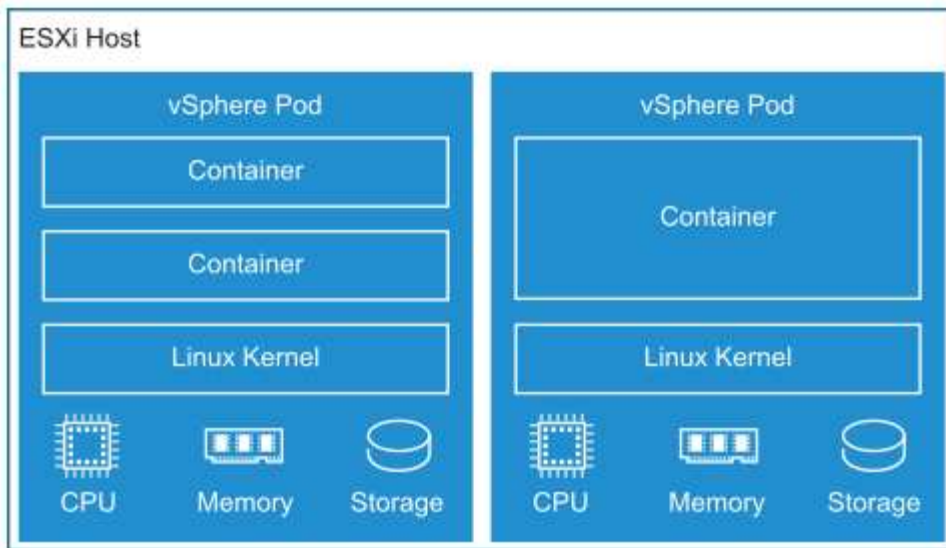
- Worker Node Instances:** A dropdown menu set to '3'.
- Worker Persistent Disk Size:** A dropdown menu set to '50 GB'.
- Worker Availability Zones:** A toggle switch is turned on, labeled 'az'.
- Worker VM Type:** A dropdown menu set to 'medium.disk (cpu: 2, ram: 4 GB, disk: 32 GB)'.
- Max Worker Node Instances:** A text input field set to '50'.
- Errand VM Type:** A dropdown menu set to 'medium.disk (cpu: 2, ram: 4 GB, disk: 32 GB)'.
- Enable Privileged Containers:** A checkbox is checked, with the text '(Use with caution)'.
- Admission Plugins:** Two toggle switches are shown: 'PodSecurityPolicy' and 'SecurityContextDeny', both of which are turned off.
- Cluster Services:** Four toggle switches are shown, all of which are turned on:
 - 'Force node to drain even if it has running pods not managed by a ReplicationController, ReplicaSet, Job, DaemonSet or Stateful Set'
 - 'Force node to drain even if it has running DaemonSet managed pods'
 - 'Force node to drain even if it has running pods using emptyDir'
 - 'Force node to drain even if pods are still running after timeout'
- Node Drain Timeout:** A text input field set to '0', with a range of '(minutes, min: 0, max: 1440)'.
- Pod Shutdown Grace Period:** A text input field set to '10', with a range of '(seconds, min: -1, max: 86400)'.

At the bottom left, there are two buttons: 'SAVE PLAN' (in blue) and 'DELETE' (in grey).

- NetApp recommends that all production environments be deployed in multiple master deployments for fault tolerance with the choice of worker nodes' configuration to meet the requirements of the intended workloads. Thus, a recommended TKGI cluster plan would consist of at least three masters and three workers with at least four vCPUs and 12GB of RAM for a highly intensive workload.

VMware vSphere with Tanzu overview

VMware vSphere with Tanzu, also known as vSphere Pods, lets you use the ESXi hypervisor nodes in your VMware vSphere environment as worker nodes in a bare metal Kubernetes environment.







A VMware vSphere with Tanzu environment is enabled under Workload Management just like a native TKGS cluster.

A virtualized Supervisor Cluster is created to provide a highly available control plane for Kubernetes, and individual Namespaces are created for each application to ensure resource isolation for users.



When VMware vSphere with Tanzu is enabled, each of the ESXi hosts have the Spherelet application installed and configured. This enables each node to act as a worker in a Kubernetes deployment and manages the pods deployed on each node.

Supervisor Cluster	
Config Status 	 Running (1)
Kubernetes Status 	 Ready
Version	0.0.15-19705778
Spherelet Version	1.3.2-19554634

Currently, VMware vSphere with Tanzu and vSphere Pods only support the local vSphere CSI driver. This works by having administrators create storage policies in the vSphere client that select from storage targets currently available to be used as vSphere datastores. These policies are used to create persistent volumes for containerized applications.



Although there currently is no support for the NetApp Trident CSI driver that allows direct connectivity to external ONTAP and Element storage arrays, these NetApp storage systems are often used to support the primary storage for the vSphere environment, and NetApp advanced data management and storage efficiency tools can be used in this manner.


If you would like to read more about VMware vSphere with Tanzu, see the documentation [here](#).

NetApp storage systems

NetApp storage systems overview

NetApp has several storage platforms that are qualified with Trident and Trident Protect to provision, protect and manage data for containerized applications and thus help in defining and maximizing DevOps throughput.

NetApp has several storage platforms that are qualified with Trident to provision, protect, and manage data for containerized applications.



- AFF and FAS systems run NetApp ONTAP and provide storage for both file-based (NFS) and block-based (iSCSI) use cases.
- Cloud Volumes ONTAP and ONTAP Select provide the same benefits in the cloud and virtual space respectively.
- Google Cloud NetApp Volumes (AWS/GCP) and Azure NetApp Files provide file-based storage in the cloud.



Each storage system in the NetApp portfolio can ease both data management and movement between on-premises sites and the cloud so that your data is where your applications are.

The following pages have additional information about the NetApp storage systems validated in the VMware Tanzu with NetApp solution:

- [NetApp ONTAP](#)

NetApp ONTAP

NetApp ONTAP is a powerful storage-software tool with capabilities such as an intuitive GUI, REST APIs with automation integration, AI-informed predictive analytics and corrective action, non-disruptive hardware upgrades, and cross-storage import.

NetApp ONTAP is a powerful storage-software tool with capabilities such as an intuitive GUI, REST APIs with automation integration, AI-informed predictive analytics and corrective action, non-disruptive hardware upgrades, and cross-storage import.

For more information about the NetApp ONTAP storage system, visit the [NetApp ONTAP website](#).

ONTAP provides the following features:


- A unified storage system with simultaneous data access and management of NFS, CIFS, iSCSI, FC, FCoE, and FC-NVMe protocols.

- Different deployment models include on-premises on all-flash, hybrid, and all-HDD hardware configurations; VM-based storage platforms on a supported hypervisor such as ONTAP Select; and in the cloud as Cloud Volumes ONTAP.
- Increased data storage efficiency on ONTAP systems with support for automatic data tiering, inline data compression, deduplication, and compaction.
- Workload-based, QoS-controlled storage.
- Seamless integration with a public cloud for tiering and protecting data. ONTAP also provides robust data protection capabilities that sets it apart in any environment:
 - **NetApp Snapshot copies.** A fast, point-in-time backup of data using a minimal amount of disk space with no additional performance overhead.
 - **NetApp SnapMirror.** Mirrors the Snapshot copies of data from one storage system to another. ONTAP supports mirroring data to other physical platforms and cloud-native services as well.
 - **NetApp SnapLock.** Efficiently administration of non-rewritable data by writing it to special volumes that cannot be overwritten or erased for a designated period.
 - **NetApp SnapVault.** Backs up data from multiple storage systems to a central Snapshot copy that serves as a backup to all designated systems.
 - **NetApp SyncMirror.** Provides real-time, RAID-level mirroring of data to two different plexes of disks that are connected physically to the same controller.
 - **NetApp SnapRestore.** Provides fast restoration of backed-up data on demand from Snapshot copies.
 - **NetApp FlexClone.** Provides instantaneous provisioning of a fully readable and writeable copy of a NetApp volume based on a Snapshot copy.

For more information about ONTAP, see the [ONTAP 9 Documentation Center](#).



NetApp ONTAP is available on-premises, virtualized, or in the cloud.



NetApp platforms

NetApp AFF/FAS

NetApp provides robust all-flash (AFF) and scale-out hybrid (FAS) storage platforms that are tailor-made with low-latency performance, integrated data protection, and multi-protocol support.

Both systems are powered by NetApp ONTAP data management software, the industry's most advanced data-management software for simplified, highly available, cloud-integrated storage management to deliver enterprise-class speed, efficiency, and security for your data fabric needs.

For more information about NETAPP AFF/FAS platforms, click [here](#).

ONTAP Select

ONTAP Select is a software-defined deployment of NetApp ONTAP that can be deployed onto a hypervisor in your environment. It can be installed on VMware vSphere or on KVM, and it provides the full functionality and experience of a hardware-based ONTAP system.

For more information about ONTAP Select, click [here](#).

Cloud Volumes ONTAP

NetApp Cloud Volumes ONTAP is a cloud-deployed version of NetApp ONTAP that can be deployed in a number of public clouds, including Amazon AWS, Microsoft Azure, and Google Cloud.

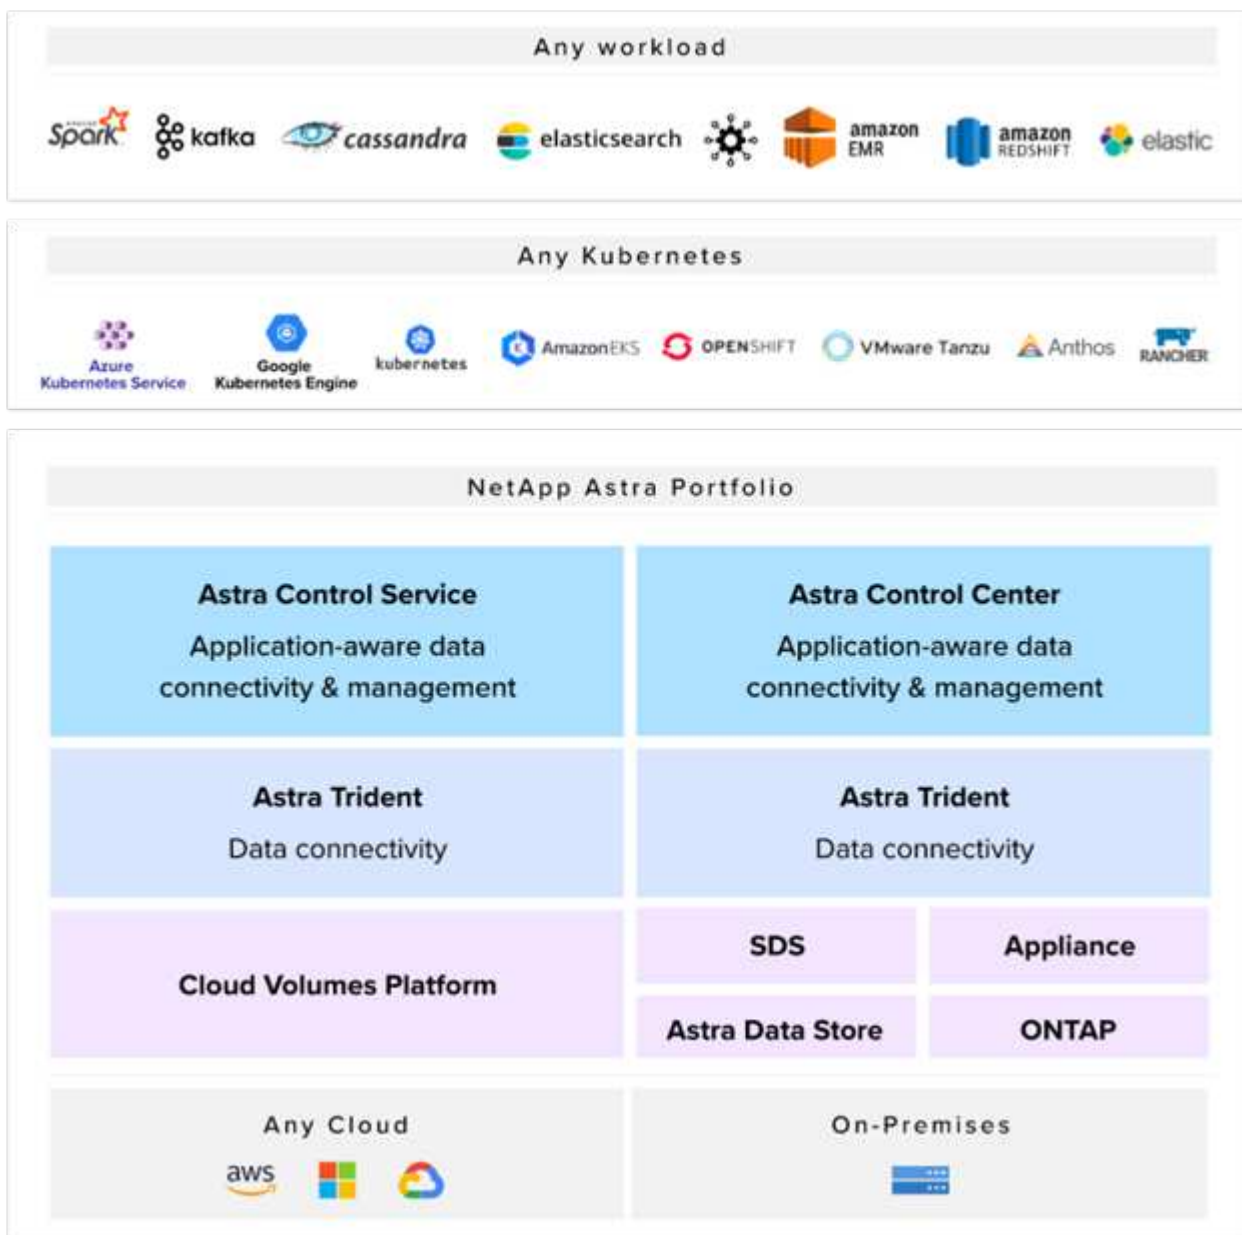
For more information about Cloud Volumes ONTAP, click [here](#).

NetApp storage integrations

NetApp storage integration overview

NetApp provides a number of products which assist our customers with orchestrating and managing persistent data in container based environments.

NetApp provides a number of products to help you orchestrate, manage, protect, and migrate stateful containerized applications and their data.



NetApp Trident is an open-source and fully-supported storage orchestrator for containers and Kubernetes

distributions like Red Hat OpenShift, Rancher, VMware Tanzu etc. For more information, visit the Trident website [here](#).

The following pages have additional information about the NetApp products that have been validated for application and persistent storage management in the VMware Tanzu with NetApp solution:

- [NetApp Trident](#)

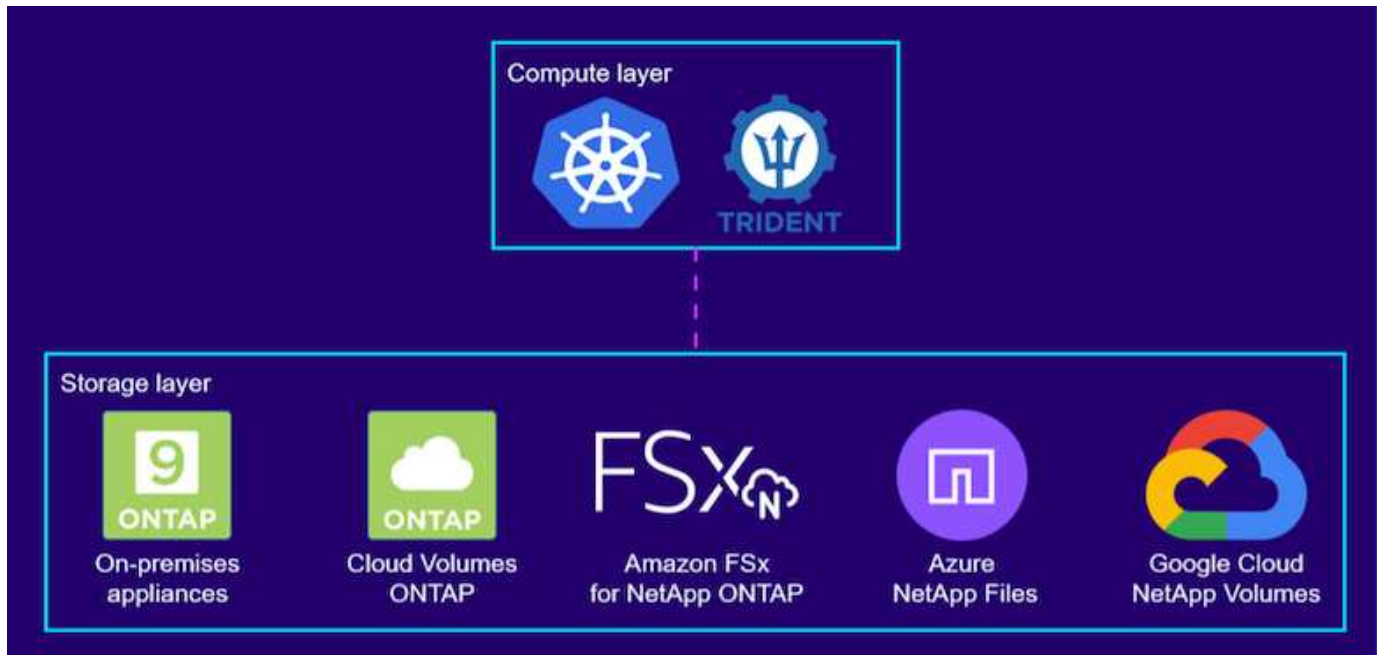
NetApp Trident

Trident overview

Trident is an open-source and fully-supported storage orchestrator for containers and Kubernetes distributions, including VMware Tanzu.

Trident is an open-source, fully supported storage orchestrator for containers and Kubernetes distributions like Red Hat OpenShift, VMware Tanzu, Anthos by Google Cloud, Rancher etc. Trident works with the entire NetApp storage portfolio, including the NetApp ONTAP and Element storage systems, and it also supports NFS and iSCSI connections. Trident accelerates the DevOps workflow by allowing end users to provision and manage storage from their NetApp storage systems without requiring intervention from a storage administrator.

An administrator can configure a number of storage backends based on project needs and storage system models that enable advanced storage features, including compression, specific disk types, or QoS levels that guarantee a certain level of performance. After they are defined, these backends can be used by developers in their projects to create persistent volume claims (PVCs) and to attach persistent storage to their containers on demand.



Trident has a rapid development cycle and, like Kubernetes, is released four times a year.

The latest version of Trident is 22.04 released in April 2022. A support matrix for what version of Trident has been tested with which Kubernetes distribution can be found [here](#).

Starting with the 20.04 release, Trident setup is performed by the Trident operator. The operator makes large

scale deployments easier and provides additional support, including self healing for pods that are deployed as a part of the Trident install.

With the 21.01 release, a Helm chart was made available to ease the installation of the Trident Operator.

Deploy Trident operator using Helm

1. First set the location of the user cluster's `kubeconfig` file as an environment variable so that you don't have to reference it, because Trident has no option to pass this file.

```
[netapp-user@rhel7]$ export KUBECONFIG=~/.tanzu-install/auth/kubeconfig
```

2. Add the NetApp Trident helm repository.

```
[netapp-user@rhel7]$ helm repo add netapp-trident
https://netapp.github.io/trident-helm-chart
"netapp-trident" has been added to your repositories
```

3. Update the helm repositories.

```
[netapp-user@rhel7]$ helm repo update
Hang tight while we grab the latest from your chart repositories...
...Successfully got an update from the "netapp-trident" chart repository
...Successfully got an update from the "bitnami" chart repository
Update Complete. ☐Happy Helming!☐
```

4. Create a new namespace for the installation of Trident.

```
[netapp-user@rhel7]$ kubectl create ns trident
```

5. Create a secret with DockerHub credentials to download the Trident images.

```
[netapp-user@rhel7]$ kubectl create secret docker-registry docker-
registry-cred --docker-server=docker.io --docker-username=netapp
-solutions-tme --docker-password=xxxxxxx -n trident
```

6. For user or workload clusters managed by TKGS (vSphere with Tanzu) or TKG with management cluster deployments, complete the following procedure to install Trident:
 - a. Ensure that the logged in user has the permissions to create service accounts in trident namespace and that the service accounts in trident namespace have the permissions to create pods.
 - b. Run the below helm command to install Trident operator in the namespace created.

```
[netapp-user@rhel7]$ helm install trident netapp-trident/trident-operator -n trident --set imagePullSecrets[0]=docker-registry-cred
```

7. For a user or workload cluster managed by TKGI deployments, run the following helm command to install Trident operator in the namespace created.

```
[netapp-user@rhel7]$ helm install trident netapp-trident/trident-operator -n trident --set imagePullSecrets[0]=docker-registry-cred,kubeletDir="/var/vcap/data/kubelet"
```

8. Verify that the Trident pods are up and running.

NAME	READY	STATUS	RESTARTS
AGE			
trident-csi-6vv62	2/2	Running	0
14m			
trident-csi-cfd844bcc-sqhcg	6/6	Running	0
12m			
trident-csi-dfcmz	2/2	Running	0
14m			
trident-csi-pb2n7	2/2	Running	0
14m			
trident-csi-qsw6z	2/2	Running	0
14m			
trident-operator-67c94c4768-xw978	1/1	Running	0
14m			

```
[netapp-user@rhel7]$ ./tridentctl -n trident version
+-----+-----+
| SERVER VERSION | CLIENT VERSION |
+-----+-----+
| 22.04.0        | 22.04.0        |
+-----+-----+
```

Create storage-system backends

After completing the Trident Operator install, you must configure the backend for the specific NetApp storage platform you are using. Follow the links below to continue the setup and configuration of Trident.

- [NetApp ONTAP NFS](#)
- [NetApp ONTAP iSCSI](#)

NetApp ONTAP NFS configuration

To enable Trident integration with the NetApp ONTAP storage system via NFS, you must create a backend that enables communication with the storage system. We configure a basic backend in this solution, but if you are looking for more customized options, visit the documentation [here](#).

Create an SVM in ONTAP

1. Log into ONTAP System Manager, navigate to Storage > Storage VMs, and click Add.
2. Enter a name for the SVM, enable the NFS protocol, check the Allow NFS Client Access checkbox, and add the subnets that your worker nodes are on in the export policy rules for allowing the volumes to be mounted as PVs in your workload clusters.

Add Storage VM



STORAGE VM NAME

trident_svm

Access Protocol

☒ SMB/CIFS, NFS, S3

iSCSI

☐ Enable SMB/CIFS

☒ Enable NFS

☒ Allow NFS client access

Add at least one rule to allow NFS clients to access volumes in this storage VM. [?](#)

EXPORT POLICY

Default

RULES

Rule Index	Clients	Access Protocols	Read-Only Rule	Read/Wr
	0.0.0.0/0	Any	Any	Any



If you are using NAT'ed deployment of user clusters or workload clusters with NSX-T, you need to add the Egress subnet (in the case of TKGS0 or the Floating IP subnet (in the case of TKGI) to the export policy rules.

3. Provide the details for data LIFs and the details for SVM administration account, and then click Save.

NETWORK INTERFACE

Use multiple network interfaces when client traffic is high.

K8s-Ontap-01

IP ADDRESS

172.21.252.180

SUBNET MASK

24

GATEWAY

172.21.252.1



BROADCAST DOMAIN

Default



Storage VM Administration

☒ Manage administrator account

USER NAME

vsadmin

PASSWORD

.....

CONFIRM PASSWORD

.....

☐ Add a network interface for storage VM management.

4. Assign the aggregates to an SVM. Navigate to Storage > Storage VMs, click the ellipsis next to the newly created SVM and then click Edit. Check the Limit Volume Creation to Preferred Local Tiers checkbox and attach the required aggregates to it.

Edit Storage VM



STORAGE VM NAME

trident_svm

DEFAULT LANGUAGE

c.utf_8



DELETED VOLUME RETENTION PERIOD 

12

HOURS

Resource Allocation

☒ Limit volume creation to preferred local tiers

LOCAL TIERS

K8s_Ontap_01_SSD_1 

Cancel

Save

5. In case of NAT'ed deployments of user or workload clusters on which Trident is to be installed, the storage mount request might arrive from a non-standard port due to SNAT. By default, ONTAP only allows the volume mount requests when originated from root port. Thus, log into ONTAP CLI and modify the setting to

allow mount requests from non-standard ports.

```
ontap-01> vserver nfs modify -vserver tanzu_svm -mount-rootonly disabled
```

Create backends and StorageClasses

1. For NetApp ONTAP systems serving NFS, create a backend config file on the jumphost with the backendName, managementLIF, dataLIF, svm, username, password, and other details.

```
{
  "version": 1,
  "storageDriverName": "ontap-nas",
  "backendName": "ontap-nas+10.61.181.221",
  "managementLIF": "172.21.224.201",
  "dataLIF": "10.61.181.221",
  "svm": "trident_svm",
  "username": "admin",
  "password": "password"
}
```



It is a best practice to define the custom backendName value as a combination of the storageDriverName and the dataLIF that is serving NFS for easy identification.

2. Create the Trident backend by running the following command.

```
[netapp-user@rhel7]$ ./tridentctl -n trident create backend -f backend-ontap-nas.json
+-----+-----+
+-----+-----+-----+
|          NAME          | STORAGE DRIVER |                UUID                |
| STATE | VOLUMES | |
+-----+-----+-----+
+-----+-----+-----+
| ontap-nas+10.61.181.221 | ontap-nas      | be7a619d-c81d-445c-b80c-5c87a73c5b1e |
| online |         | 0 |
+-----+-----+-----+
+-----+-----+-----+
```

3. With the backend created, you must next create a storage class. The following sample storage class definition highlights the required and basic fields. The parameter backendType should reflect the storage driver from the newly created Trident backend.

```

apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: ontap-nfs
provisioner: csi.trident.netapp.io
parameters:
  backendType: "ontap-nas"

```

4. Create the storage class by running the kubectl command.

```

[netapp-user@rhel7 trident-installer]$ kubectl create -f storage-class-nfs.yaml
storageclass.storage.k8s.io/ontap-nfs created

```

5. With the storage class created, you must then create the first persistent volume claim (PVC). A sample PVC definition is given below. Make sure that the `storageClassName` field matches the name of the storage class just created. The PVC definition can be further customized as required depending upon the workload to be provisioned.

```

kind: PersistentVolumeClaim
apiVersion: v1
metadata:
  name: basic
spec:
  accessModes:
    - ReadWriteOnce
  resources:
    requests:
      storage: 1Gi
  storageClassName: ontap-nfs

```

6. Create the PVC by issuing the kubectl command. Creation can take some time depending on the size of the backing volume being created, so you can watch the process as it completes.

```

[netapp-user@rhel7 trident-installer]$ kubectl create -f pvc-basic.yaml
persistentvolumeclaim/basic created

```

```

[netapp-user@rhel7 trident-installer]$ kubectl get pvc

```

NAME	STATUS	VOLUME	CAPACITY
ACCESS MODES		STORAGECLASS	AGE
basic	Bound	pvc-b4370d37-0fa4-4c17-bd86-94f96c94b42d	1Gi
RWO		ontap-nfs	7s

NetApp ONTAP iSCSI configuration

To integrate NetApp ONTAP storage system with VMware Tanzu Kubernetes clusters for persistent volumes via iSCSI, the first step is to prepare the nodes by logging into each node and configuring the iSCSI utilities or packages to mount iSCSI volumes. To do so, follow the procedure laid out in this [link](#).



NetApp does not recommend this procedure for NAT'ed deployments of VMware Tanzu Kubernetes clusters.



TKGI uses Bosh VMs as nodes for Tanzu Kubernetes clusters that run immutable configuration images, and any manual changes of iSCSI packages on Bosh VMs do not remain persistent across reboots. Therefore, NetApp recommends using NFS volumes for persistent storage for Tanzu Kubernetes clusters deployed and operated by TKGI.

After the cluster nodes are prepared for iSCSI volumes, you must create a backend that enables communication with the storage system. We configured a basic backend in this solution, but, if you are looking for more customized options, visit the documentation [here](#).

Create an SVM in ONTAP

To create an SVM in ONTAP, complete the following steps:

1. Log into ONTAP System Manager, navigate to Storage > Storage VMs, and click Add.
2. Enter a name for the SVM, enable the iSCSI protocol, and then provide details for the data LIFs.

Add Storage VM



STORAGE VM NAME

trident_svm_iscsi

Access Protocol

SMB/CIFS, NFS, S3

iSCSI

☒ Enable iSCSI

NETWORK INTERFACE

K8s-Ontap-01

IP ADDRESS	SUBNET MASK	GATEWAY	BROADCAST DOMAIN
10.61.181.231	24	10.61.181.1	Defa...

☐ Use the same subnet mask, gateway, and broadcast domain for all of the following interfaces

IP ADDRESS	SUBNET MASK	GATEWAY	BROADCAST DOMAIN
10.61.181.232	24	10.61.181.1	Defa...

3. Enter the details for the SVM administration account, and then click Save.

Storage VM Administration

☒ Manage administrator account

USER NAME

vsadmin

PASSWORD

CONFIRM PASSWORD

☐ Add a network interface for storage VM management.

Save

Cancel

4. To assign the aggregates to the SVM, navigate to Storage > Storage VMs, click the ellipsis next to the newly created SVM, and then click Edit. Check the Limit Volume Creation to Preferred Local Tiers checkbox, and attach the required aggregates to it.

Edit Storage VM



STORAGE VM NAME

trident_svm_iscsi

DEFAULT LANGUAGE

c.utf_8



DELETED VOLUME RETENTION PERIOD 


12

HOURS

Resource Allocation

☒ Limit volume creation to preferred local tiers

LOCAL TIERS

K8s_Ontap_01_SSD_1 

Cancel

Save

Create backends and StorageClasses

1. For NetApp ONTAP systems serving NFS, create a backend config file on the jumphost with the backendName, managementLIF, dataLIF, svm, username, password, and other details.

```
{
  "version": 1,
  "storageDriverName": "ontap-san",
  "backendName": "ontap-san+10.61.181.231",
  "managementLIF": "172.21.224.201",
  "dataLIF": "10.61.181.231",
  "svm": "trident_svm_iscsi",
  "username": "admin",
  "password": "password"
}
```

2. Create the Trident backend by running the following command.

```
[netapp-user@rhel7 trident-installer]$ ./tridentctl -n trident create
backend -f backend-ontap-san.json
+-----+-----+-----+
+-----+-----+-----+-----+
|          NAME          | STORAGE DRIVER |          UUID          |
| STATE | VOLUMES | |          |          |
+-----+-----+-----+-----+
+-----+-----+-----+-----+
| ontap-san+10.61.181.231 | ontap-san      | 6788533c-7fea-4a35-b797- |
| fb9bb3322b91 | online |      0 |          |
+-----+-----+-----+-----+
+-----+-----+-----+-----+
```

3. After you create a backend, you must next create a storage class. The following sample storage class definition highlights the required and basic fields. The parameter `backendType` should reflect the storage driver from the newly created Trident backend. Also note the name-field value, which must be referenced in a later step.

```
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: ontap-iscsi
provisioner: csi.trident.netapp.io
parameters:
  backendType: "ontap-san"
```



There is an optional field called `fsType` that is defined in this file. In iSCSI backends, this value can be set to a specific Linux filesystem type (XFS, ext4, and so on) or can be deleted to allow Tanzu Kubernetes clusters to decide what filesystem to use.

4. Create the storage class by running the kubectl command.

```
[netapp-user@rhel7 trident-installer]$ kubectl create -f storage-class-iscsi.yaml
storageclass.storage.k8s.io/ontap-iscsi created
```

5. With the storage class created, you must then create the first persistent volume claim (PVC). A sample PVC definition is given below. Make sure that the `storageClassName` field matches the name of the storage class just created. The PVC definition can be further customized as required depending upon the workload to be provisioned.

```
kind: PersistentVolumeClaim
apiVersion: v1
metadata:
  name: basic
spec:
  accessModes:
    - ReadWriteOnce
  resources:
    requests:
      storage: 1Gi
  storageClassName: ontap-iscsi
```

6. Create the PVC by issuing the kubectl command. Creation can take some time depending on the size of the backing volume being created, so you can watch the process as it completes.

```
[netapp-user@rhel7 trident-installer]$ kubectl create -f pvc-basic.yaml
persistentvolumeclaim/basic created
```

```
[netapp-user@rhel7 trident-installer]$ kubectl get pvc
```

NAME	STATUS	VOLUME	CAPACITY
ACCESS MODES	STORAGECLASS	AGE	
basic	Bound	pvc-7ceac1ba-0189-43c7-8f98-094719f7956c	1Gi
RWO		ontap-iscsi	3s

Additional Information: VMware Tanzu with NetApp

To learn more about the information described in this document, review the following websites:

- NetApp Documentation

<https://docs.netapp.com/>

- Trident Documentation

<https://docs.netapp.com/us-en/trident/>

- Ansible Documentation

<https://docs.ansible.com/>

- VMware Tanzu Documentation

<https://docs.vmware.com/en/VMware-Tanzu/index.html>

- VMware Tanzu Kubernetes Grid Documentation

<https://docs.vmware.com/en/VMware-Tanzu-Kubernetes-Grid/1.5/vmware-tanzu-kubernetes-grid-15/GUID-index.html>

- VMware Tanzu Kubernetes Grid Service Documentation

<https://docs.vmware.com/en/VMware-vSphere/7.0/vmware-vsphere-with-tanzu/GUID-152BE7D2-E227-4DAA-B527-557B564D9718.html>

- VMware Tanzu Kubernetes Grid Integrated Edition Documentation

<https://docs.vmware.com/en/VMware-Tanzu-Kubernetes-Grid-Integrated-Edition/index.html>

Copyright information

Copyright © 2025 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.