

# PIVOTAL CONTAINER SERVICE (PKS) ON VMWARE vSAN™

Simplicity, Agility, Resilience, and Performance

## The Rise of Cloud-Native Applications

Digital transformation is driving a new application development and deployment approach called Cloud-Native. Cloud-Native Applications empower your developers by providing the resources and environments they need to deliver applications faster, on-premises or in the cloud, without the need for ops intervention.

## Hyper-Converged Infrastructure for Cloud-Native Applications

VMware vSAN™, the market leader in Hyper-Converged Infrastructure (HCI), enables low cost and high performance next-generation HCI solutions. vSAN delivers TCO reduction benefits over traditional 3-tiered architecture by eliminating silos. vSAN allows customers to evolve without risk and to achieve infrastructure agility and scalability by converging traditional IT infrastructure silos onto industry-standard servers.

VMware Cloud Provider (VCP) enables dynamic provisioning, easy deployment, upgrade and management for storage of Cloud-Native Applications. With VMware's Cloud Provider, administrators and developers can easily use vSAN for both ephemeral or and persistent storage for stateless and stateful containerized applications.

The storage class definition in Kubernetes maps to policies defined through vSAN Storage Policy Based Management (SPBM) to achieve different levels of SLAs and QoS and can benefit from advanced vSAN data services features such as deduplication, compression and checksums.

## Why PKS on vSAN?

Pivotal Container Service (PKS) is a purpose-built container solution to operationalize Kubernetes for enterprises moving to hybrid and multi-cloud deployments as well as service providers. It significantly simplifies the deployment and management of Kubernetes clusters with Day 1 and Day 2 operations support.

vSAN's SPBM offers users flexibility to define policies on demand in VMware vCenter® and delivers ease of management of storage for containers. Data services such as snapshots, cloning, encryption, deduplication and compression are available at a container volume level of granularity.

Deep integration between PKS and vSAN means developers can consume storage as code by abstracting the complexity of the underlying storage infrastructure. With Project Hatchway and vSAN services, Cloud-Native Applications take advantage of operational benefits of hyper-converged storage and compute as well as seamless application failover and rapid recovery.

### UNIQUE CAPABILITIES

#### EASY MANAGEMENT AND OPERATION

- Easily deploy Kubernetes clusters in HCI environments.
- Integrate Kubernetes storage classes with VMware Storage Policy Based Management.
- Easily scale-out vSAN and Kubernetes clusters.
- Protect infrastructure and workload VMs with VMware vSphere® High Availability and vSAN data services.

**SIMPLE NETWORKING**

- No need to know the underlying network architecture with the help of VMware NSX-T™.
- Networking is isolated between namespaces.
- Networking can be easily managed with Kubernetes clusters - deployment, upgrade and scaling out.

**PKS 1.0 on vSAN Architecture**

We used a 4-node vSAN cluster and deployed the necessary virtual machines on it. The architecture of the cluster is depicted in Figure 1.

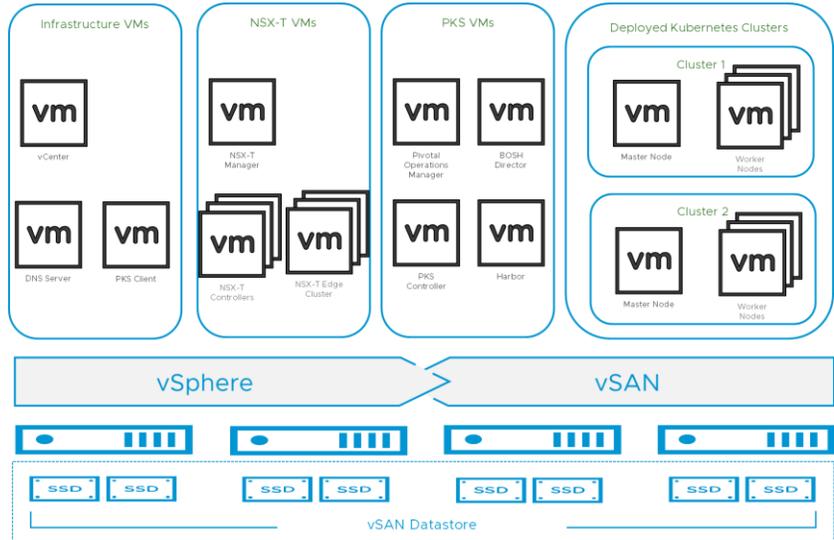


Figure 1. High-Level Architecture of the PKS on vSAN Reference Architecture

**VERIFIED SOLUTION ARCHITECTURE**

- Fast deployment.
- Deep integration with SPBM.
- Easy management and operations.
- Unified control plane.

For PKS 1.0, we use one Availability Zone (AZ) to reflect one vSAN cluster. As shown in the figure, the virtual machines are divided into four groups:

- Infrastructure VMs
- NSX-T VMs
- PKS VMs
- Deployed Kubernetes Clusters

Kubernetes clusters are deployed by PKS. The virtual machines are grouped by Kubernetes cluster. Each Kubernetes cluster contains one virtual machine acting as the master node and multiple virtual machines acting as the worker nodes. We can deploy more than one Kubernetes cluster per vSAN cluster.

**Solution Validation**

**Deployment**

After the deployment of the PKS virtual machines, we run the following commands in the 'PKS Client' virtual machine to create a cluster for validation:

```
$ pks create-cluster cluster1 --external-hostname cluster1.vsphere.local --plan small --num-nodes 5
```

The plan 'small' is configured with 1 master node and 3 worker nodes. Both master and worker nodes are configured with 1 vCPU, 2GB memory, an 8GB OS disk, and a 10GB persistent disk.

## AGILITY PERSONIFIED

### ADAPT TO CHANGING NEEDS

- Adopt and integrate the very latest hardware technologies like 3D XPoint NVMe devices into a cluster.
- Adopt the very latest and well proven software including VMware vCenter, vSphere, vSAN and NSX-T.
- Scale up or out incrementally, as needed by an organization.
- Maintain full independence of storage from demands of other clusters. Just as with compute and memory, vSAN storage is a cluster resource that remains independent from other clusters.

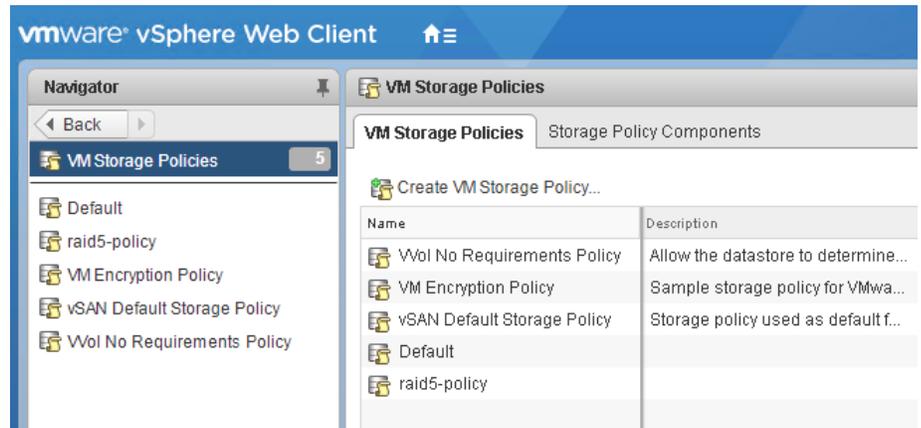
## LEARN MORE ABOUT PKS AND vSAN

- [Pivotal Container Service](#)
- [VirtualBlocks](#) – VMware’s blog site for all topics related to storage and availability.
- [StorageHub](#) – The one-stop location for all documentation on storage and availability.

## Integration with SPBM

For PKS 1.0, we used one Availability Zone with one vSAN cluster. This makes the storage class definition and Persistent Volume Claim (PVC) simple.

Create a storage policy in vCenter:



Use the storage policy in the Kubernetes storage class definition:

```
vmware@UbuntuJump:~$ cat storage-class-raid5.yml
kind: StorageClass
apiVersion: storage.k8s.io/v1
metadata:
  name: vsan-default-sc
provisioner: kubernetes.io/vsphere-volume
parameters:
  diskformat: thin
  storagePolicyName: raid5-policy
  datastore: vsanDatastore
```

This configuration file can be used to create a Kubernetes storage class and PVC. This operation creates a corresponding VMDK on the vSAN datastore with associated storage policy.

## Summary

With the tight integration of PKS, NSX-T, and vSAN, you can easily provision networks for containers in Kubernetes clusters, manage ephemeral and persistent storage as well as benefit from vSAN’s availability and data service features. In addition, you can protect virtual machines against physical server failure by using vSphere HA and VMware vSphere Fault Tolerance. The combination of these technologies makes PKS a complete solution, which is perfect for Kubernetes administrators and developers.

For more details, stay tuned for our complete solution paper.

