

GOOGLE KUBERNETES ENGINE (GKE) ON-PREM ON VMWARE vSAN™

Simplicity, Agility, Resilience, and Performance

Google Cloud's Anthos on VMware vSAN

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 IT operations intervention.

Hyperconverged Infrastructure for Cloud-Native Applications

VMware, Hyperconverged Infrastructure (HCI), powered by vSAN™, is the market leader in low cost and high performance next-generation HCI solutions. vSAN delivers TCO reduction benefits over traditional 3-tiered architecture by eliminating infrastructure silos. vSAN also allows customers to achieve infrastructure agility and scalability by converging traditional compute and storage silos onto industry-standard servers, dramatically simplifying operations.

VMware vSphere Cloud Provider (VCP) enables dynamic provisioning, easy deployment, upgrade and management for storage of Cloud-native applications. With VMware's vSphere Cloud Provider, administrators and developers can easily use vSAN for both ephemeral 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 GKE On-Prem on vSAN?

Google Cloud's Anthos enables you build and manage modern hybrid applications on existing on-premises investments or in the public cloud. Built on open source technologies pioneered by Google—including Kubernetes, Istio, and Knative—Anthos enables consistency between on-premises and cloud environments. GKE On-Prem is a key component of Anthos. GKE On-Prem enables you take advantage of Kubernetes and cloud technology in your data center. You get Google Kubernetes Engine (GKE) experience with quick, managed, and simple installs as well as upgrades validated by Google Cloud.

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

Integration between GKE On-Prem and vSAN means developers can consume storage as code by abstracting the complexity of the underlying storage infrastructure.

UNIQUE CAPABILITIES

- 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.

GKE ON-PREM FEATURE HIGHLIGHTS

- Workload portability between on-premises and cloud.
- Stateful application support.
- Centralized multi-cluster management.

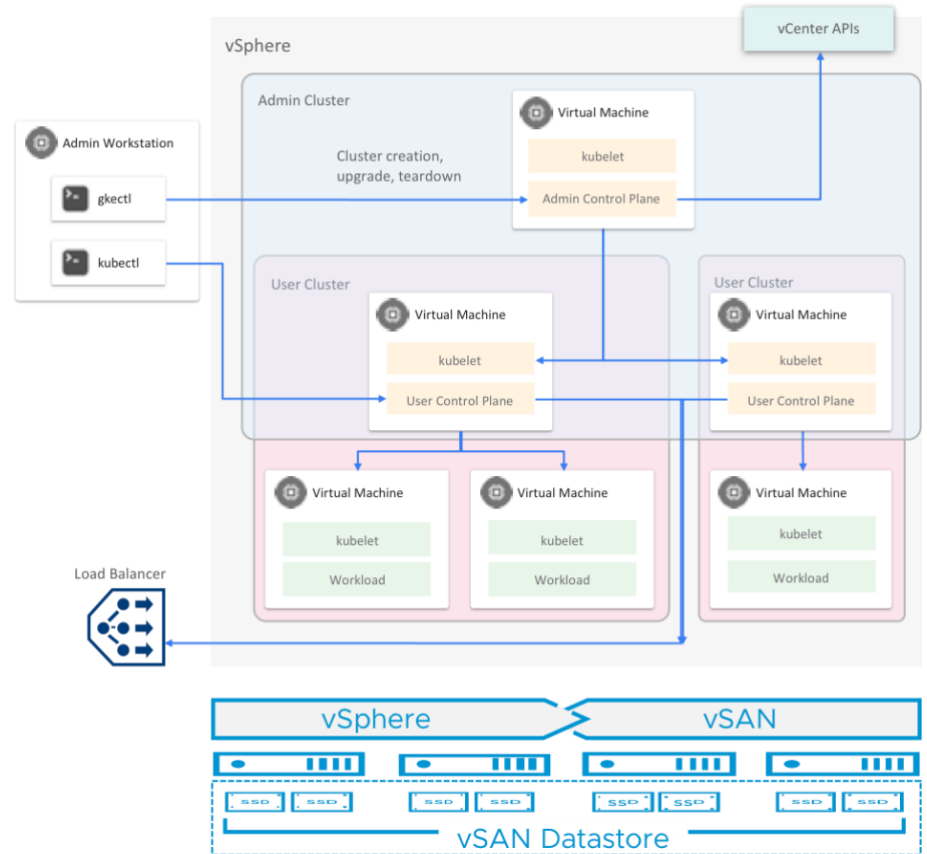
VERIFIED SOLUTION ARCHITECTURE

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

With Project Hatchway and vSAN services, Cloud-native applications take advantage of the operational benefits of hyperconverged infrastructure as well as seamless application failover and rapid recovery.

Google Kubernetes Engine (GKE) On-Prem on vSAN Architecture

The high-level architecture of GKE On-Prem on vSAN is depicted in figure below.



GKE On-Prem uses both Kubernetes admin clusters and user clusters. An admin cluster is a cluster from which user clusters are created and managed. User clusters are where you actually deploy and run your Kubernetes workloads. You can think of an admin cluster as a "parent" cluster, and user clusters as "child" clusters.

All virtual machines running in both Kubernetes admin and user clusters can be deployed on a vSAN datastore, the same datastore on which their Persistent Volumes (PV) are provisioned. Each Kubernetes cluster contains one virtual machine acting as the master node and multiple virtual machines acting as the worker nodes. You can deploy more than one Kubernetes user cluster per vSAN cluster.

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 and vSAN.
- 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 GOOGLE CLOUD'S ANTHOS, GKE ON-PREM AND vSAN

- [GKE On-Prem](#) – Product overview
- [Anthos](#) – Component overview
- [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.

Solution Validation

Deployment

Admin clusters run the following GKE On-Prem components:

- **Admin control plane:** includes the Kubernetes control plane components for the admin cluster. This control plane specifically handles all API calls to and from GKE On-Prem.
- **User cluster control plane:** includes the Kubernetes control plane components for a user cluster. Each user cluster you create has its own control plane. This control plane handles network load balancing and routes API requests to user cluster nodes.

User control planes are managed by the admin cluster. They run on nodes in admin clusters, not in the user clusters themselves. To manage user control planes, admin cluster needs to:

- Manage the machines that run the user cluster control planes
- Create, update, and delete the control plane components
- Expose the Kubernetes API server to the user cluster
- Manage cluster certificates

Integration with Storage Policy Based Management (SPBM)

Storage class definition in Kubernetes makes it simple to instantiate Persistent Volume Claim (PVC) that consumes vSAN storage using the storage policy and other related parameters defined in the storage class as shown in the example below:

Create a storage policy in vCenter:

The screenshot shows the vSphere Client interface. The top navigation bar includes the 'vm' logo, 'vSphere Client', a 'Menu' dropdown, and a search bar with the text 'Search in all environments'. The main content area is divided into two sections. On the left, under 'Policies and Profiles', there is a list of items: 'VM Customization Specifications', 'VM Storage Policies' (which is highlighted), 'Host Profiles', and 'Storage Policy Components'. On the right, the 'VM Storage Policies' section is expanded, showing a 'Create VM Storage Policy' button and a table of existing policies. The table has a 'Name' column and lists the following policies: 'Host-local PMem Default Storage Policy', 'raid5-policy', 'VM Encryption Policy', 'vSAN Default Storage Policy', and 'VVol No Requirements Policy'.

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 seamless integration between GKE On-Prem and vSAN, you can easily provision 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 GKE On-Prem on vSAN an agile cloud native solution, which is perfect for Kubernetes administrators and developers.

