Running Google Anthos on VMware Cloud Foundation

Reference Architecture

vmware

DELL Technologies

Table of contents

Note: This solution provides general design and deployment guidelines for running Google Anthos on VMware Cloud Foundation™. It is showcased in this paper running on Dell EMC VxRail. The reference architecture applies to any compatible hardware platforms running VMware Cloud Foundation.

Executive Summary

Business Case

Containers have become a prominent technology to deliver the speed, agility, portability, and scalability required by developers without compromising the security and visibility required by IT operations. Google Anthos and VMware Software-Defined Data Center (SDDC) technologies are a great combination for delivering an enterprise container platform on a virtual infrastructure.

Google Anthos is a modern application management platform that provides a consistent development and operations experience for cloud and on-premises environments. Google Kubernetes Engine (GKE) on-prem, a key component of Anthos, is an enterprise-ready container platform that manages the lifecycle of Kubernetes clusters to improve developer productivity.

With a single architecture that is easy to deploy, VMware Cloud Foundation™ can provision compute, network, and storage on demand. VMware Cloud Foundation protects network and data with micro-segmentation and satisfies compliance requirements with data-at-rest encryption. Policy-based management delivers business-critical performance. VMware Cloud Foundation delivers flexible, consistent, secure infrastructure and operations across private and public clouds and is ideally suited to meet the demands of Anthos.

Dell VxRail™ is the only jointly engineered hyperconverged infrastructure (HCI) system with deep VMware Cloud Foundation integration, delivering a simple and direct path to modern apps and the hybrid cloud with one, complete, automated platform. Full stack integration with VMware Cloud Foundation on Dell VxRail means both hyperconverged infrastructure layer and VMware cloud software stack are managed as one, complete, automated, and turnkey hybrid cloud experience, greatly reducing risk and increasing IT operational efficiency. VMware Cloud Foundation on Dell VxRail is an ideal turnkey hybrid cloud platform for Google Anthos based modern apps enabling developers to create and push code more frequently with zero downtime.

In this solution, we provide the generic design and deployment guidelines for running Google Anthos with VMware Cloud Foundation on Dell VxRail.

Technology Overview

Solution technology components are listed below:

- VMware Cloud Foundation
    - VMware vSphere®
    - VMware vSAN™
    - VMware NSX® Data Center
    - SDDC Manager
- Dell VxRail
    - VxRail HCI System Software
- Google Anthos

VMware Cloud Foundation

VMware Cloud Foundation is an integrated software stack that combines compute virtualization (VMware vSphere), storage virtualization (VMware vSAN), network virtualization (VMware NSX), and cloud management and monitoring (VMware vRealize® Suite) into a single platform that can be deployed on-premises as a private cloud or run as a service within a public cloud. This documentation focuses on the private cloud use case. VMware Cloud Foundation

bridges the traditional administrative silos in data centers, merging compute, storage, network provisioning, and cloud management to facilitate end-to-end support for application deployment.

VMware vSphere

VMware vSphere is VMware's virtualization platform, which transforms data centers into aggregated computing infrastructures that include CPU, storage, and networking resources. vSphere manages these infrastructures as a unified operating environment and provides operators with the tools to administer the data centers that participate in that environment. The two core components of vSphere are ESXi™ and vCenter Server®. ESXi is the hypervisor platform used to create and run virtualized workloads. vCenter Server is the management plane for the hosts and workloads running on the ESXi hosts.

VMware vSAN

VMware vSAN is the industry-leading software powering VMware's software defined storage and HCI solution. vSAN helps customers evolve their data center without risk, control IT costs, and scale to tomorrow's business needs. vSAN, native to the market-leading hypervisor, delivers flash-optimized, secure storage for all of your critical vSphere workloads, and is built on industry-standard x86 servers and components that help lower TCO in comparison to traditional storage. It delivers the agility to scale IT easily and offers the industry's first native HCI encryption.

vSAN simplifies Day 1 and Day 2 operations, and customers can quickly deploy and extend cloud infrastructure and minimize maintenance disruptions. vSAN helps modernize Hyperconverged Infrastructure (HCI) by providing administrators a unified storage control plane for both block and file protocols and provides significant enhancements that make it a great solution for traditional virtual machines as well cloud-native applications. vSAN helps reduce the complexity of monitoring and maintaining infrastructure and enables administrators to rapidly provision a file share in a single workflow for Kubernetes-orchestrated cloud native applications.

VMware NSX Data Center

VMware NSX Data Center is the network virtualization and security platform that enables the virtual cloud network, a software-defined approach to networking that extends across data centers, clouds, and application frameworks. With NSX Data Center, networking and security are brought closer to the application wherever it's running, from virtual machines to containers to bare metal. Like the operational model of VMs, networks can be provisioned and managed independently of the underlying hardware. NSX Data Center reproduces the entire network model in software, enabling any network topology—from simple to complex multitier networks—to be created and provisioned in seconds. Users can create multiple virtual networks with diverse requirements, leveraging a combination of the services offered via NSX or from a broad ecosystem of third-party integrations ranging from next-generation firewalls to performance management solutions to build inherently more agile and secure environments. These services can then be extended to a variety of endpoints within and across clouds.

Kubernetes vSphere CSI Driver

Cloud Native Storage (CNS) is a vSphere and Kubernetes (K8s) feature that makes K8s aware of how to provision storage on vSphere on-demand, in a fully automated, scalable fashion as well as providing visibility for the administrator into container volumes through the CNS User Interface within vCenter. Run, monitor, and manage containers and virtual machines on the same platform—in the same way:

- Simplify your infrastructure needs, lifecycle, and operations.
- Lower costs, using a platform you already know for consistent operations across workloads and across clouds.
- Spend less time managing infrastructure and more time building apps that provide business value.

The main goal of CNS is to make vSphere and vSphere storage, including vSAN, a platform to run stateful Kubernetes workloads. vSphere has a great data path that is highly reliable, highly performant, and mature for enterprise use. CNS enables access of this data path to Kubernetes and brings an understanding of Kubernetes volume and pod abstractions to vSphere. CNS was first released in vSphere 6.7 Update 3.

Dell VxRail

The only fully integrated, pre-configured, and pre-tested VMware hyperconverged integrated system optimized for VMware vSAN and VMware Cloud Foundation, VxRail transforms HCI networking and simplifies VMware cloud adoption while meeting any HCI use case - including support for many of the most demanding workloads and applications. Powered by Dell PowerEdge server platforms and VxRail HCI System Software, VxRail features next-generation technology to future proof your infrastructure and enables deep integration across the VMware ecosystem. The advanced VMware hybrid cloud integration and automation simplifies the deployment of a secure VxRail cloud infrastructure.

VxRail HCI System Software

VxRail HCI system software is integrated software that delivers a seamless and automated operational experience, offering 100% native integration between VxRail Manager and vCenter®. Intelligent lifecycle management automates non-disruptive upgrades, patching, and node addition or retirement while keeping VxRail infrastructure in a continuously validated state to ensure that workloads are always available. The HCI System includes access to Dell CloudIQ, a SaaS multi-cluster monitoring tool used for centralized data collection and analytics that uses machine learning and AI to help customers keep their HCI stack operating at peak performance and ready for future workloads. IT teams can benefit from the actionable insights to optimize infrastructure performance, improve serviceability, and foster operational freedom.

Google Anthos

Google Anthos unifies the management of infrastructure and applications across GCP, on-premises, edge and other public clouds with a Google Cloud-backed control plane for consistent operation at scale. Anthos is a managed platform for all your application deployments, both traditional as well as cloud native so you can quickly modernize apps and establish operational consistency across them. You can accelerate developer productivity and software delivery by bringing the benefit of cloud services, containers, and serverless across your deployments with cloud-native tooling and guidance from Google. Anthos on-prem enables you to manage GKE clusters and third-party Kubernetes conformant clusters on VMware vSphere. You get a consistent managed Kubernetes experience with simple installs and upgrades validated by Google. Anthos integrates additional tools to help you manage your environment. Anthos Configuration Manager helps automate policy and security at scale across environments and uses a GitOps approach to evaluate and roll out changes. Anthos ensures the desired state is reached and corrects any configuration drift automatically. Anthos Service Mesh provides a fully managed service mesh with built-in visibility for your applications. With Anthos, security is integrated into each stage of the application life cycle—from develop to build to run—while enabling a defense-in-depth security strategy with a comprehensive portfolio of security controls across all of these deployment models.

See *https://cloud.google.com/anthos* for detailed information regarding Google Anthos.

Solution Configuration

This section introduces the resources and configurations:

- Architecture diagram
- Anthos installation
- Virtual machine placement and VMware vSphere Distributed Resource Scheduler™
- VMware vSphere High Availability
- vSAN Failures to Tolerance setting
- Hardware resources
- Software resources
- Network configuration

- vSAN configuration

Architecture Diagram

The VMware Cloud Foundation test environment was composed of a management domain and a workload domain. We deployed the Google Anthos in the workload domain, and all other infrastructure VMs were in the separate management workload domain (figure 1).

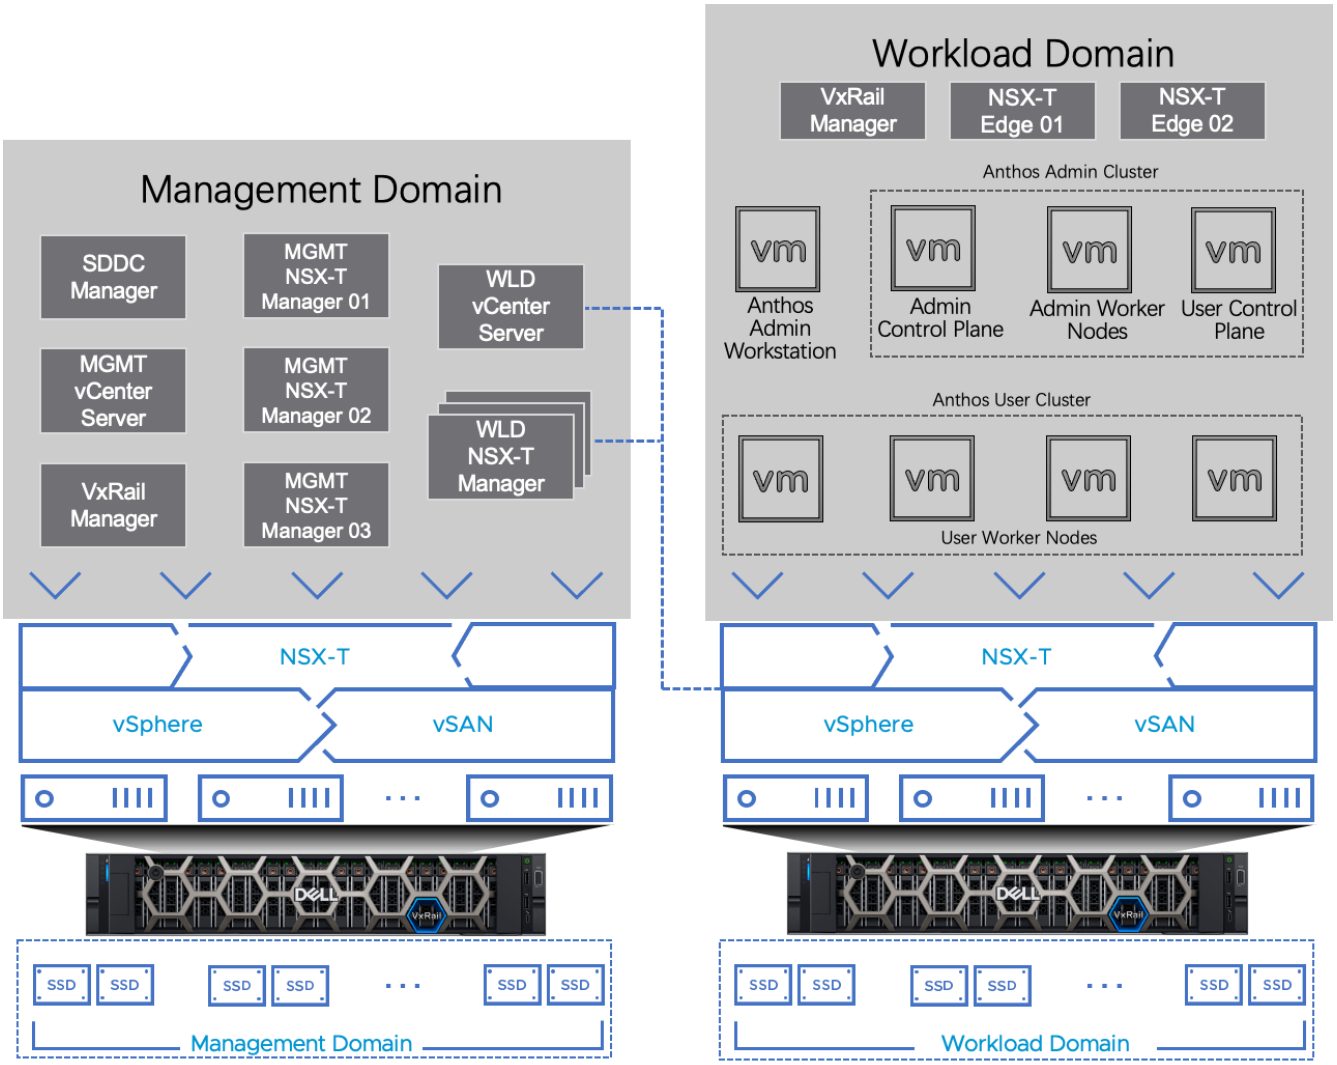

Figure 1. Google Anthos with VMware Cloud Foundation on VxRail Solution Architecture

Notation in Figure 1:

- Admin workstation: This bootstrap workstation is used to create admin and user clusters.
- Admin control plane: The admin cluster's control plane includes the Kubernetes API server, the scheduler, and several controllers for the admin cluster.
- Admin worker node: The admin worker nodes together with their control plane form the admin cluster that is responsible for lifecycle managing the user clusters.

- User control plane: For each user cluster, these are the user control plane nodes of Kubernetes deployed and managed by Anthos. Notice the user control plane nodes reside in the admin cluster and are logically segregated from their worker nodes.
- User worker node: For each user cluster, these are the user worker nodes of Kubernetes deployed and managed by Anthos. We deployed 4 worker nodes as the starting point. More worker nodes can be added on demand through the admin workstation.

In our solution, we created a 4-node VxRail P570F cluster for the VMware Cloud Foundation management domain, running management virtual machines and appliances. The management domain can be shared with other workload domains.

Table 1. Management Domain VMs

VM Role	vCPU	Memory (GB)	VM Count
Management Domain vCenter Server	4	16	1
SDDC Manager	4	16	1
Management Domain VMware NSX® Manager™	6	24	3
Workload Domain NSX Manager	12	48	3
Workload Domain vCenter Server	8	28	1
VxRail Manager Appliance	2	8	1

For the workload domain, we created another 4-node VxRail P570F cluster with a separate NSX-T Fabric, deployed an NSX Edge Cluster, and deployed the Anthos VMs in the workload domain.

Table 2 shows the deployment of the workload domain edge nodes and Anthos VMs. For the workload domain edge node, we recommend that VMware NSX Edge™ transport nodes are deployed with “Large” form factor.

Table 2. Workload Domain VMs

VM Role	Minimum vCPU	Minimum Memory (GB)	Storage	Deployment Size	VM Count
VxRail Manager Appliance	2	8	160GB	n/a	1
Workload Domain Edge Node	8	32	200GB	Large	2
Anthos Admin Workstation	4	8	50GB	n/a	1
Anthos Admin Control Plane Node	4	16	140GB	n/a	1
Anthos Admin Worker Node	4	16	40GB	n/a	2

Anthos User Control Plane Node	4	8	60GB	n/a	3
Anthos User Worker Node	4	8	40GB	n/a	4

This is a building block approach for a basic installation of Anthos with VMware Cloud Foundation on VxRail. Based on the customer demands and dataset size, we can expand the workload domain to include more physical hosts. A cluster with vSAN enabled supports up to 64 physical hosts for non-stretched cluster. With adding more hosts to the vSAN cluster, not only the capacity of CPU and memory is increased for computing but also the capacity of vSAN storage is increased accordingly. This is one of the benefits of HCI that we can increase the capacity of computing and storage at the same time and proportionally.

Anthos Installation

Listed below are the primary steps involved in the installation of Anthos clusters on VMware. For detailed installation instructions, refer to *Anthos clusters on VMware installation overview*.

1. Set up Google Cloud resources including a Cloud project and a service account.
2. Create an admin workstation. The admin workstation is a VM that has the tools to bootstrap the admin and user clusters.
3. Create an admin cluster from the admin workstation. An admin cluster is a point of management for a set of user clusters that run your workloads.
4. Create one or more user clusters from the admin workstation.
5. Deploy a workload on a user cluster.

The integration between Anthos GKE on-prem and Cloud Native Storage in vSphere through the Container Storage Interface (CSI) with VMware vSAN enables developers to provision persistent storage for Anthos workloads on vSphere on-demand in an automated fashion. It also allows IT administrators to manage container volumes through the Cloud Native Storage UI within vCenter. Developers and IT administrators can now have a consistent view of container volumes and troubleshoot at the same level.

Unlike CSI drivers provided by other storage vendors that need to be manually installed, the *vSphere CSI driver* is prebuilt in the Anthos code that it is readily to be consumed out-of-the-box as VMware vSphere is the only supported virtualized platform for Anthos GKE on-prem.

Once an Anthos cluster is deployed on vSphere and it is registered with Google Cloud Platform, it can be centrally monitored and managed within the GCP console as shown below:

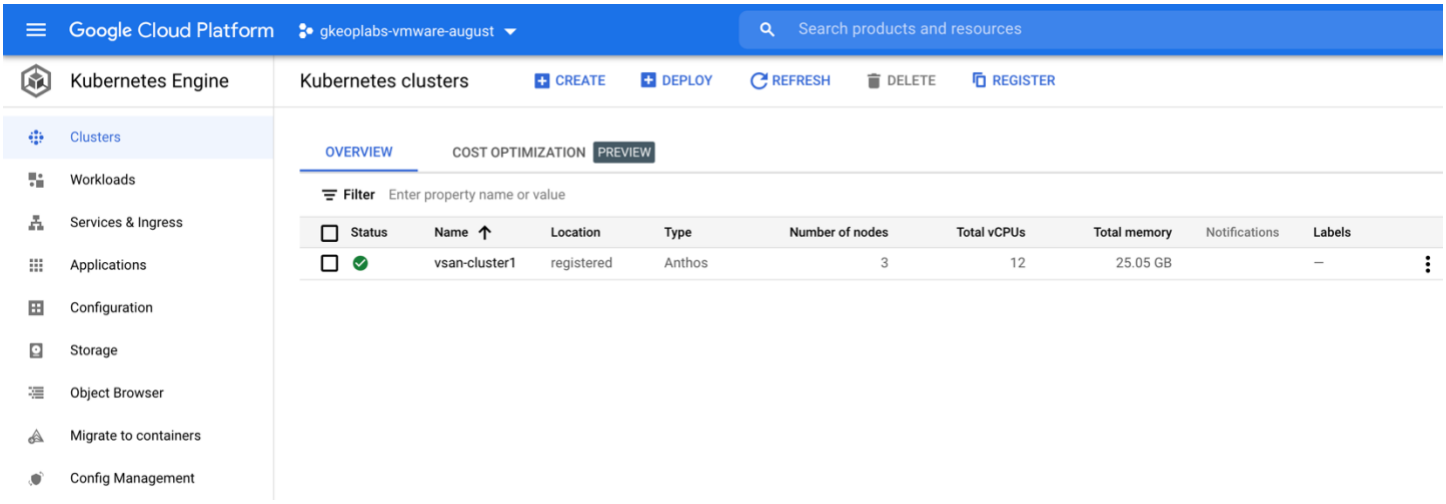

Hardware Resources

In this solution, for the workload domain of Anthos, we used a total of four VxRail R570F nodes. Each server was configured with two disk groups, and each disk group consisted of one cache-tier write-intensive SAS SSD and four capacity-tier read-intensive SAS SSDs.

Each VxRail node in the cluster had the following configuration, as shown in table 3.

Table 3. Hardware Configuration for VxRail

PROPERTY	SPECIFICATION
Server model name	VxRail P570F
CPU	2 x Intel(R) Xeon(R) Gold 6148 CPU @ 2.40GHz, 28 core each
RAM	512GB
Network adapter	2 x Broadcom BCM57414 NetXtreme-E 25Gb RDMA Ethernet Controller
Storage adapter	1 x Dell HBA330 Adapter
Disks	Cache - 2 x 800GB Write Intensive SAS SSDs Capacity - 8 x 3.84TB Read Intensive SAS SSDs

Software Resources

Table 4 shows the software resources used in this solution.

Table 4. Software Resources

SOFTWARE	VERSION	PURPOSE

VMware Cloud Foundation on VxRail	4.5	A unified SDDC platform on VxRail that brings together VMware vSphere, vSAN, NSX, and optionally, vRealize Suite components, into a natively integrated stack to deliver enterprise-ready cloud infrastructure for the private and public cloud. See *BOM of VMware Cloud Foundation on VxRail* for details.
Dell VxRail	7.0.400	Turnkey Hyperconverged Infrastructure for hybrid cloud
VMware vSphere	7.0 update 3h/MP3 (vCenter) 7.0 update 3g/EP5 (ESXi)	VMware vSphere is a suite of products: vCenter Server and ESXi.
VMware vSAN	7.0 update 3g/EP5	vSAN is the storage component in VMware Cloud Foundation to provide low-cost and high-performance next-generation HCI solutions.
NSX-T	3.2.1.2	NSX-T is the key network component in VMware Cloud Foundation on VxRail and is deployed automatically. It is designed for networking management and operation.
Anthos	1.13	The version of Anthos software being tested in this solution.

Network Configuration

Figure 2 shows the VMware vSphere Distributed Switch™ network configuration for Anthos cluster in the workload domain of the VMware Cloud Foundation on VxRail. NSX-T, which underlies the vSphere infrastructure, is used for the Anthos cluster networking. To enable external access for the Anthos cluster, an NSX-T edge cluster is required to deploy. Also it is required to configure the BGP peering and route distribution of the upstream network. For more details, refer to *VMware Cloud Foundation on VxRail Planning and Preparation Guide*.

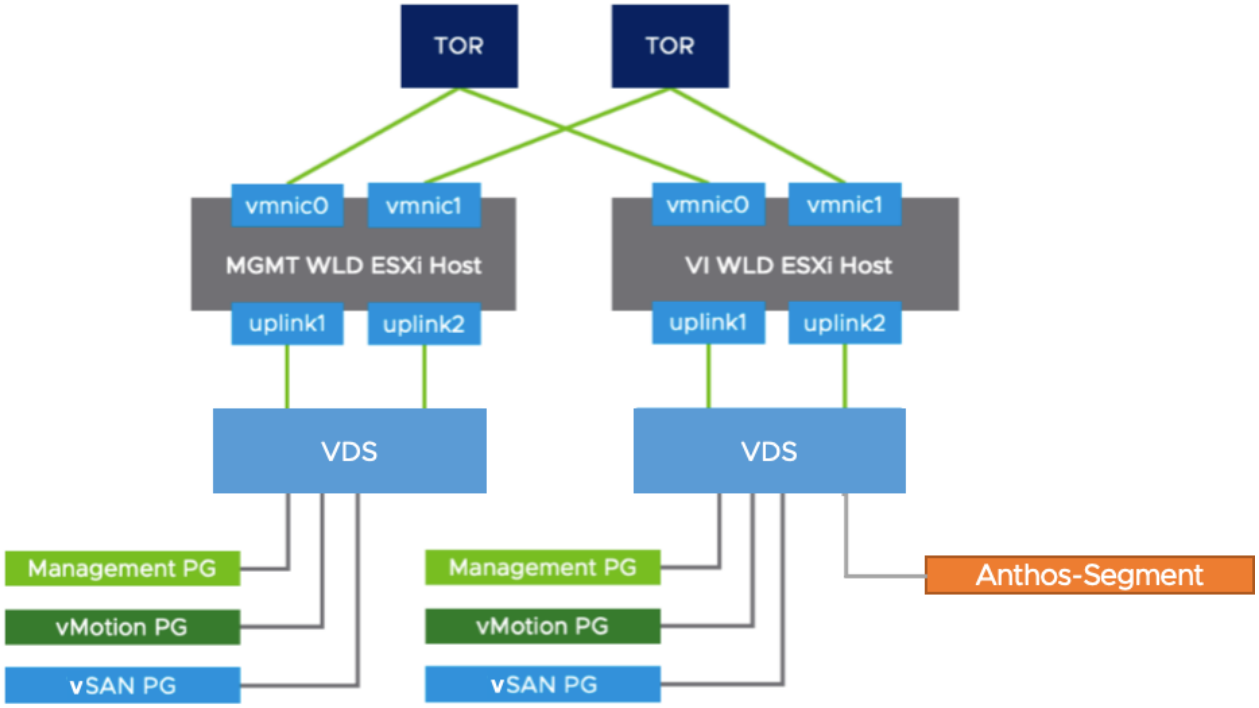

Figure 2. The Overall NSX-T Networking Architecture

Figure 2 shows the VMware vSphere Distributed Switches configuration for both management domain and workload domain of the VMware Cloud Foundation. For each domain, two 25 GbE vmnics were used and configured with teaming policies. The management domain can be shared among different workloads.

The NSX-T controllers resided in the management domain. The Anthos virtual machines were configured with a VM network called 'Anthos-Segment' on an NSX-T segment. VMware vSphere vMotion®, vSAN, and TEP for NSX-T had another dedicated segment created. In the workload domain's uplink setting, we used a dedicated physical NIC for vSAN traffic. The other physical NIC was dedicated to Anthos virtual machines' traffic. The reason is that both vSAN and Anthos may use high network traffic, so both of them need the dedicated NIC.

Jumbo Frame (MTU=9000) was enabled on the physical switches, vSAN VMkernel, and all the virtual switches to improve performance.

NSX-T managers and edges have more than one instance to form NSX clusters to achieve HA and better load balancing. Besides, based on workloads, the vCPU and memory may be adjusted to achieve better performance. Table 5 shows the configuration of the NSX-T managers and edge nodes virtual machines. The NSX-T managers reside in the management workload domain, so it will not cost the compute resources for Anthos VMs. However, the NSX-T edge nodes reside in the Anthos workload domain, and it will cost some CPU and memory resources. This should be taken into consideration while doing the sizing of the cluster before Anthos is deployed.

Table 5. NSX-T VM Configuration

NSX-T VM ROLE	INSTANCE	VCPU	MEMORY (GB)	VM NAME	VIRTUAL DISK SIZE	OPERATING SYSTEM
NSX-T Manager	3	12	48	NSX-unified-appliance-<version>	200GB	Ubuntu
NSX-T Edge Nodes	2	4	8	Edge-<UUID>	120GB	Ubuntu

vSAN Configuration

The solution validation was based on a 4-node vSAN cluster as a building block.

The validation tests were conducted using the default vSAN datastore storage policy of RAID 1 FTT=1, checksums enabled, deduplication and compression deactivated, and no encryption. In the below sections, we explained the detailed configurations of the vSAN cluster and some items in the Storage Policy Based Management (SPBM).

Deduplication and Compression

The 'Deduplication and Compression' option was configured on the cluster level and it can be enabled or deactivated for the whole vSAN cluster. While in our testing, we deactivated it; by enabling it, we can reduce the vSAN storage usage but induce higher latencies for the Anthos application. This is a tradeoff for customers' choices.

Failures to Tolerance (FTT)

Failures to Tolerance (FTT) is a configuration item in vSAN's storage policy. For the 'StorageClass' in Anthos and the corresponding vSAN's storage policy, we recommended setting vSAN's Failures to Tolerate (FTT) to 1. In our testing, we set FTT to 1 as the baseline. Do not set the FTT to 0 in an Anthos cluster with vSAN deployment because FTT=0 may possibly cause the data of the replications of the same pod to be stored in the same physical disk. This may cause data loss in case of a physical disk failure.

In the case of using RAID 1 in vSAN policy, there are two copies for each piece of data in vSAN. So, the estimated database capacity requirement should not exceed half of the vSAN's overall capacity. In the case of RAID 5, vSAN consumes 1.33 times of the raw capacity, and you can calculate the storage usage accordingly. If the capacity increase is needed, the additional machines can be added to the cluster, and vSAN can increase the data capacity storage for Anthos online without the service interruption to Anthos users.

Checksum

Checksum is a configuration item in vSAN's storage policy. We compared the Kubernetes performance between enabling and disabling checksum. By disabling vSAN's checksum, there is barely any performance improvement for applications deployed by Anthos, while by enabling it, we can ensure the data is correct from the vSAN storage hardware level. So, we recommend keeping the checksum enabled, which is the default value.

Erasure Coding (RAID 1 vs. RAID 5)

Erasure Coding is a configuration item in vSAN's storage policy. It is also known as configuring RAID 5 or RAID 6 for vSAN objects. With FTT=1 and RAID 1, the data in vSAN is mirrored and the capacity cost would be 2 times of the raw capacity. With FTT=1 and RAID 5, the data is stored as RAID 5 and the capacity cost would be 1.33 times of the raw capacity.

In our testing, we used FTT=1 without Erasure Coding (RAID 1). By enabling Erasure Coding, we could save some vSAN storage spaces but induce higher latencies for the Kubernetes applications. Again, this is a tradeoff for customers' choices.

Encryption

vSAN can perform data at rest encryption. Data is encrypted after all other processing, such as deduplication. Data at rest encryption protects data on the storage devices.

Encryption is not used in our testing. Use encryption as per your company's Information Security requirements.

Solution Validation

Test Tools

We leveraged the following monitoring and benchmark tools in this solution.

Monitoring Tools

vSAN Performance Service

vSAN Performance Service is used to monitor the performance of the vSAN environment through the vSphere Client. The performance service collects and analyzes performance statistics and displays the data in a graphical format. You can use the performance charts to manage your workload and determine the root cause of the problems.

vSAN Health Check

vSAN Health Check delivers a simplified troubleshooting and monitoring experience of all things related to vSAN. Through the vSphere client, it offers multiple health checks specifically for vSAN including cluster, hardware compatibility, data, limits, and physical disks. It is used to check the vSAN health before the mixed-workload environment deployment.

This is only for vSAN health check. We can also enable VxRail cluster health monitoring for overall health monitoring.

Application Validation Tools

Confluent Kafka

Apache Kafka is a community distributed event streaming platform capable of handling trillions of events a day. Initially conceived as a messaging queue, Kafka is based on an abstraction of a distributed commit log. Since being created and open sourced by LinkedIn in 2011, Kafka has quickly evolved from a messaging queue to a full-fledged event streaming platform.

Founded by the original developers of Apache Kafka, Confluent delivers the most complete distribution of Kafka with the Confluent Platform. In addition, the Confluent Platform improves Kafka with additional community and commercial features designed to enhance the streaming experience of both operators and developers in production at a massive scale.

We used Confluent Kafka as one of the applications for performance validation that Anthos is on a par with other Kubernetes distributions.

Jenkins

Jenkins is an open-source automation server that lets you flexibly orchestrate your build, test, and deployment pipelines. A Kubernetes cluster adds a new automation layer to Jenkins. Kubernetes makes sure that resources are used effectively and that your servers and the underlying infrastructure are not overloaded. In addition, Kubernetes' ability to orchestrate container deployment ensures that Jenkins always has the right amount of resources available.

We used Jenkins as one of the applications for Anthos's functional validation. Listed below are the primary steps involved in deploying Jenkins with persistent volume on a user cluster:

1. Execute **kubectl apply –f jenkins-pvc.yaml** to create a persistent volume claim (PVC) using the default Anthos storage class "standard-rwo".

```
# jenkins-pvc.yaml
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  annotations:
  name: jenkins-pvc
spec:
  accessModes:
  - ReadWriteOnce
  resources:
    requests:
      storage: 80Gi
  storageClassName: standard-rwo
```

2. Execute **kubectl apply –f jenkins-deployment.yaml** to create a Kubernetes deployment that leverages the persistent volume created above to store the persistent data in $JENKINS_HOME.

```
# jenkins-deployment.yaml
apiVersion: apps/v1
kind: Deployment
metadata:
  name: jenkins
spec:
  replicas: 1
  selector:
    matchLabels:
      app: jenkins
  template:
    metadata:
```

```
      labels:
        app: jenkins
    spec:
      containers:
      - image: jenkins/jenkins:lts
        name: jenkins
        ports:
          - name: http-port
            containerPort: 8080
          - name: jnlp-port
            containerPort: 50000
        securityContext:
          allowPrivilegeEscalation: true
          privileged: true
          readOnlyRootFilesystem: false
          runAsUser: 0
        volumeMounts:
          - mountPath: /var/jenkins_home
            name: jenkins-vol
      volumes:
      - name: jenkins-vol
        persistentVolumeClaim:
          claimName: jenkins-pvc
```

3. Execute **kubectl –f jenkins-services.yaml** to create Kubernetes services that expose the deployment through NodePort 30000 on a worker node.

```
# jenkins-services.yaml
apiVersion: v1
kind: Service
metadata:
  name: jenkins
spec:
  type: NodePort
  ports:
    - port: 8080
      targetPort: 8080
      nodePort: 30000
```

```
  selector:
    app: jenkins
---
apiVersion: v1
kind: Service
metadata:
  name: jenkins-jnlp
spec:
  type: ClusterIP
  ports:
    - port: 50000
      targetPort: 50000
  selector:
    app: jenkins
```

4. Execute **kubectl get all** to verify the deployment and services are running.

```
ubuntu@gke-admin-ws1:~$ kubectl get all
NAME                              READY   STATUS    RESTARTS   AGE
pod/jenkins-69454ccfc6-knbdj      1/1     Running   0          35m

NAME                    TYPE        CLUSTER-IP      EXTERNAL-IP   PORT(S)          AGE
service/jenkins         NodePort    10.96.13.162    <none>        8080:30000/TCP   6h4m
service/jenkins-jnlp    ClusterIP   10.96.1.195     <none>        50000/TCP        6h4m
service/kubernetes      ClusterIP   10.96.0.1       <none>        443/TCP          11d

NAME                       READY   UP-TO-DATE   AVAILABLE   AGE
deployment.apps/jenkins    1/1     1            1           35m

NAME                                  DESIRED   CURRENT   READY   AGE
replicaset.apps/jenkins-69454ccfc6    1         1         1       35m
```

5. In vSphere UI, select your vSphere cluster in the left pane and navigate to ***Monitor***->***Cloud Native Storage***->***Container Volumes*** to examine if the PVC is successfully mapped to a Kubernetes pod.

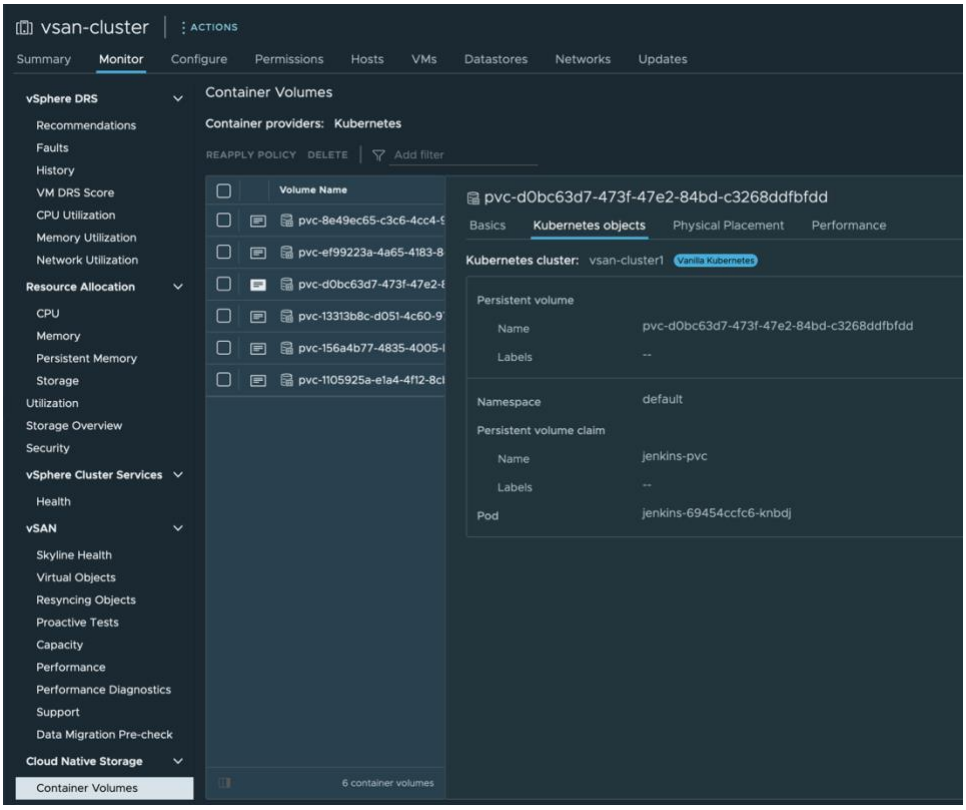

6. Recall that the Jenkins service is exposed via NodePort 30000 on a worker node. Open a browser using URL *http://WorkerNodeIPAddress:30000* to access the Jenkins interface. You should see the *Getting Started* screen as shown below. To retrieve the initial admin password, you can run **kubectl logs <jenkins-pod-name>**.

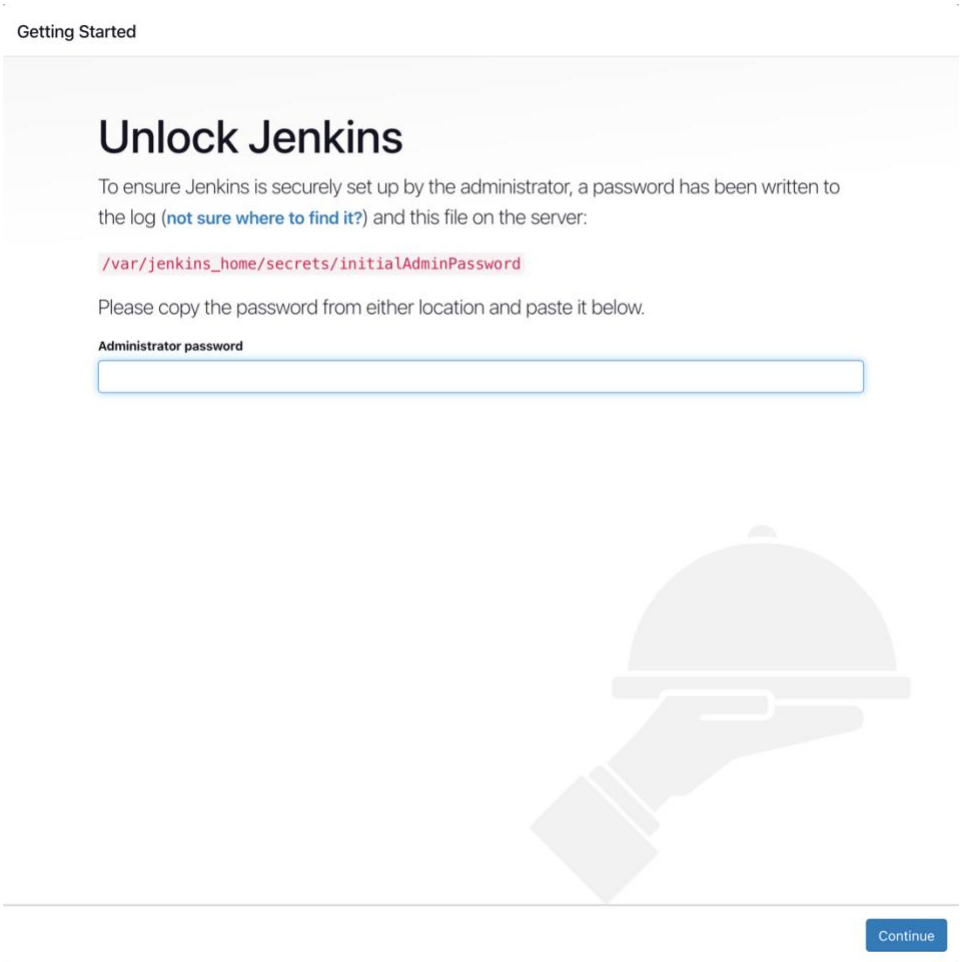

7. After you answer a few questions in the initial setup wizard, you can navigate to the Jenkins dashboard as follows.

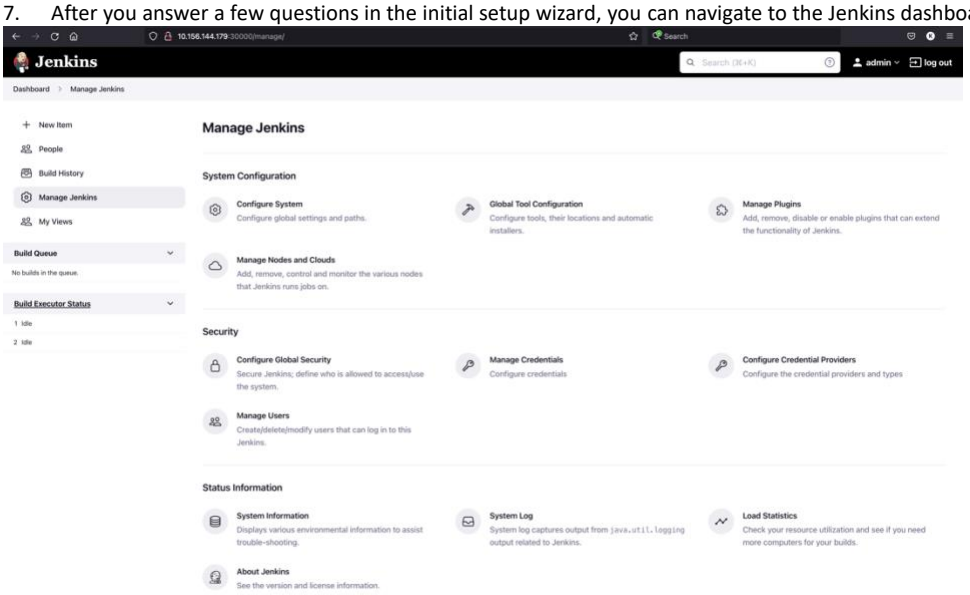

Gitlab

GitLab is the community version and a GitHub like service that organizations can use to provide internal management of git repositories. It is a self-hosted Git-repository management system that keeps the user code private and can easily deploy the changes of the code. It manages projects, not tools. With GitLab, you get an open DevOps platform delivered as a single application—one interface, one conversation thread, and one data store.

We used Gitlab as one of the applications for Anthos's functional validation.

Refer to this *Gitlab's document* for installing Gitlab on Kubernetes.

Failure Testing

This section introduces the failure scenarios and the behavior of failover and failback. This section includes:

- Physical host failure
- Physical cache disk failure
- Physical capacity disk failure

These are related to vSAN storage functional testing under hardware failure scenarios. Under each failure condition, we saw data in vSAN stay intact and no data is lost. For more detailed information, refer to vSAN Operations Guide.

Best Practices

vSphere

- Use the same server model for the physical hosts in the workload domain.
- Enable vSphere HA in the cluster.

  We recommend enabling vSphere High Availability for the workload domain cluster.

  If vSphere HA is enabled, in case of a physical host failure and there are enough remaining resources to satisfy the resource reservation like having a spare host, vSphere can automatically power on the impacted virtual machines on the other surviving hosts.

  In case of a physical host failure and if there are not enough remaining resources to satisfy the resource reservation, vSphere HA would not restart the impacted virtual machines, which is by design. Because forcing a virtual machine restart on a surviving host may cause resource contention and imbalanced performance among the Anthos nodes. We suggest that the resource reservation should at least be set to all the control plane nodes.
- Enable vSphere DRS in the cluster.

  VMware vSphere Distributed Resource Scheduler (DRS) is a feature included in the vSphere Enterprise Plus. In this solution, if DRS is enabled in the cluster, the rule of thumb is:

  - Place the control plane nodes on different physical hosts to accommodate one host failure. Anthos can automatically create DRS Anti-Affinity rules to separate the VMs to different physical hosts.
  - Let DRS do the automatic placement of compute node virtual machines.

For DRS Anti-Affinity rules, see *the DRS documentation.*

vSAN

- Enable Jumbo Frame on the physical switches. Use Jumbo Frames on the vSAN VMKernel and all virtual switches.

- Set Failures to Tolerate (FTT) to at least 1 in vSAN's storage policy for data protection.
- Enable vSAN's checksum.

Anthos

- Follow the guidelines from *Anthos documentation* for the detailed deployment and optimization items.
- Followed the *VMware CSI driver documentation* for the list of CSI driver features.

Conclusion

VMware Cloud Foundation on VxRail delivers flexible, consistent, secure infrastructure and operations across private and public clouds. It is ideally suited to meet the demands of modern applications running on Google Anthos in a virtualized environment.

With VMware Cloud Foundation, we can easily manage the lifecycle of the hybrid cloud environment. Besides, we have a unified management plane for all applications including Anthos. With VMware Cloud Foundation, we can leverage the leading virtualization technologies including vSphere, NSX-T, and vSAN.

In this solution paper, we demonstrated the architecture of running Google Anthos with VMware Cloud Foundation on VxRail. We showed the configuration details, the hardware resources, and the software resources used in the solution validation. We showed the various configuration options in addition to the best practices. VxRail Manager and VMware Cloud Foundation SDDC Manager provided the lifecycle management. vSAN provides reliable, high-performance, and flexible storage to Anthos. NSX-T provided the fine-grained, secured, and high-performance virtual networking infrastructure to Anthos. Also, vSphere DRS and vSphere HA provided efficient resource usage and high availability. All the above lead to a consolidated solution of running Google Anthos with VMware Cloud Foundation on VxRail.

References

- *VMware Cloud Foundation*
  - *Announcing General Availability of VMware Cloud Foundation 4.5*
  - *Get the Facts of VMware Cloud Foundation – Part 6*
- *VMware vSphere*
- *VMware vSAN*
- *VMware NSX*
- *Dell VxRail*
- *VMware Cloud Foundation on Dell EMC VxRail Administration Guide*
- *VMware Cloud Foundation 4.x on VxRail Architecture Guide*
- *Google Anthos*

Appendix

Sample YAML files of Jenkins Deployment during Solution Validation

https://github.com/vsphere-tmm/Anthos-on-VCF/

About the Author

Ka Kit Wong, Staff Solutions Architect in the Solutions Architecture team of the Cloud Infrastructure Business Group (CIBG) at VMware, wrote the original version of this paper.

The following reviewers also contributed to the paper contents:

- Chen Wei, Director in the Workload Technical Marketing team at VMware
- Catherine Xu, Senior Manager in the Workload Technical Marketing team at VMware
- Raja Jadeja, Product Manager at Google
- Lisa Shen, Outbound Product Manager at Google
- Denis Jatsiv, Cloud Partner Engineering at Google
- Jason Marques, Sr. Principal Engineer of VxRail Technical Marketing at Dell Technologies