Hewlett Packard Enterprise | Google Cloud

Reference Architecture

HPE REFERENCE CONFIGURATION FOR KUBERNETES DEPLOYMENT WITH GOOGLE CLOUD'S ANTHOS ON HPE NIMBLE STORAGE DHCI

Using HPE CSI driver for Kubernetes

CONTENTS

EXECUTIVE SUMMARY

At the heart of every digital enterprise are applications delivering transformational customer experience and a competitive edge. Organizations are rapidly adopting containers to build modern applications that are distributed and resilient in nature. Containers provide the necessary abstraction to package applications and their dependencies, making them ideal building blocks for wide-ranging purposes:

- Microservices-based distributed applications
- Use in continuous integration (CI) and continuous deployment (CD) pipelines
- Deployment of applications that are truly portable in hybrid cloud and multi cloud environments

As containers have come into wider use, however (particularly in enterprise environments), the limits to pure container statelessness have become all too apparent. Many of the applications now being deployed in containers were not written from scratch with containerization in mind; they are existing applications (often in the middle or early stages of their lifecycle, rather than being legacy) that might have been refactored for containers, or simply containerized. These applications are typically stateful, and they are likely to rely heavily on state data.

At the same time, the infrastructure that houses the container requirements is also becoming increasingly complex. Controlling VM sprawl is a constant battle, along with pressure to lower the cost and provide more support for both traditional and modern applications. Also, designing, building, and managing open-source Kubernetes clusters can be taxing and requires a broad skillset.

Trying to use a traditional converged solution with servers and storage requires quite a bit of manual configuration work to put all the pieces of the puzzle together: defining virtual networking, assigning the correct interface to the correct network, creating volumes, and so forth. Customers need a solution that automates most of the common tasks of setting up the infrastructure components—freeing administrators to focus solely on leveraging Google Cloud's Anthos. Anthos is an open-application modernization platform that enables you to modernize your existing applications, build new ones, and run them anywhere. Built on open source technologies pioneered by Google—including Kubernetes, Istio, and Knative—Anthos enables consistency between on-premises and cloud environments. Anthos helps accelerate application development and strategically equips your business with transformational technologies.

Now, with the announcement of HPE Nimble Storage dHCI (disaggregated hyper-converged infrastructure), you can create a complete solution combined with Anthos to help accelerate application development and strategically enable your business with transformational technologies. The HPE Nimble Storage dHCI solution, an Anthos Validated Design, dramatically simplifies the deployment, management, and scaling of VMware® and Kubernetes clusters. Using HPE primary storage and Google Cloud's Anthos, the solution provides efficient data management for containers and enables application developers and DevOps teams to leverage persistent storage into the Kubernetes construct. It offers powerful capabilities:

- A converged, flexible, flash-optimized, and API-driven architecture that enables independent scaling of compute and storage
- A cloud-native container runtime experience, powered by a primary Anthos component—Google Kubernetes Engine (GKE™)—running in a private cloud on-premises or in a public cloud
- A single-pane-of-glass view for managing container infrastructure, applications, and workloads, provided through Anthos

HPE Nimble Storage dHCI also supports integrated advanced management for persistent storage, which is a requirement for stateful container applications such as databases and message queues.

Target audience: The target audience includes IT decision makers and infrastructure and application architects who want to implement stateful containers running on Anthos GKE to build modern applications in their on-premises data centers and on Google Cloud™.

Document purpose: This document can help readers achieve the following goals:

- Gain insight into the value proposition for the HPE Nimble Storage dHCI solution with Anthos
- Better understand the requirements of HPE Nimble Storage dHCI and Anthos components
- Better understand the recommended software and features that are part of the HPE Nimble Storage dHCI solution with Anthos
- Better understand the HPE CSI driver for Kubernetes and how to integrate it with Anthos.
- Better understand the design considerations related to fault tolerance, performance, and scalability when architecting the solution

This Reference Configuration describes the solution testing performed in May 2020.

SOLUTION OVERVIEW

This section provides a brief overview of the technical aspects of the solution, including HPE Nimble Storage dHCI, Anthos, and the HPE CSI driver. Deploying the solution section explains how to deploy the Anthos cloud software and the HPE Nimble Storage dHCI solution.

HPE Nimble Storage dHCI

HPE Nimble Storage dHCI provides the flexibility to scale storage and compute independently, making it ideal for mission-critical databases and data warehouses. HPE Nimble Storage dHCI is a Converged System built on HPE InfoSight and integrated VM management, which makes it simple to manage and capable of scaling rack-to-apps in minutes.

With HPE Nimble Storage dHCI, VM administrators can now deploy their infrastructure faster, to support any application across hybrid cloud, and to create efficiency at any scale.

In this Reference Configuration, HPE Nimble Storage dHCI supports Anthos deployed on-premises. It offers a robust solution for business-critical apps and is easy to integrate with HPE CSI driver.

Google Cloud's Anthos

Anthos is an integrated cloud services platform that serves as an umbrella for the multiple features and tools offered by Google Cloud on Google Cloud Platform™. Anthos expands upon the existing Google Kubernetes Engine (GKE) and introduces new services such as Anthos GKE, Anthos Configuration Management, Anthos Service Mesh, Cloud Run for Anthos, and more to come. Anthos uses open source technologies managed by Google to modernize applications, move workloads into containers, and manage and migrate clouds in both hybrid and native environments. The unified Anthos platform simplifies the process of application development and enables customers with distributed infrastructure to manage their data consistently across both on-premises and cloud environments.

Anthos GKE deployed on-premises (GKE on-prem) provides the advantages of Kubernetes and cloud technology on-premises. It delivers the Anthos GKE experience with quick, managed, and simple installations as well as upgrades that are validated by Google. In addition, Google Cloud Console offers a single-pane-of-glass view for managing clusters across on-premises and cloud environments.

SOLUTION COMPONENTS

This section describes the hardware and software components used to build the solution.

Hardware

Table 1 summarizes the hardware components used in this solution.

TABLE 1. HPE StoreFabric 2010M switches with HPE ProLiant DL380 Gen 10 servers and HPE Nimble Storage AF40 array

Part number	Quantity	Description
Q9E63A	2	HPE StoreFabric SN2010M 25GbE 18SFP28 4QSFP28 Switch
868703-B21	3	HPE ProLiant DL380 Gen10 servers (where vSphere runs and Anthos is deployed)
Q8H41A	1	HPE Nimble Storage AF40 All Flash Dual Controller 10GBASE-T 2-port Configure-to-order Base Array (for persistent storage)

HPE M-Series 2010M

The HPE M-series SN2010M Ethernet Switch is a half-width 10/25 GbE and 40/100 GbE Ethernet switch designed for primary storage, secondary storage, and hyperconverged infrastructures. This top-of-rack solution packs 18 ports of 10/25 GbE and 4 ports of 40/100 GbE that can be reconfigured with breakout cables.

It delivers low latency for 10/25 GbE and 40/100 GbE switching. In addition, it features a robust implementation of data, control, and management planes and offers a compact form factor and low power consumption.

The switch also provides ultra-low latency of under 300 ns port-to-port. This ultra-low latency is advantageous for flash storage, which has moved latency bottlenecks from storage access to the network, as well as for the burst nature of traffic in today's software-defined and cloud data centers. Table 2 lists the components of HPE M-Series 2100M.

TABLE 2. HPE M-Series 2010M components

Part number	Quantity	Description
Q9E63A	2	HPE StoreFabric SN2010M 25GbE 18SFP28 4QSFP28 Switch
JD089B	8	HPE X120 1G SFP RJ45 T Transceiver
844477-B21	16	HPE 25Gb SFP28 to SFP28 3m Direct Attach Copper Cable
H7J34A3 W0P	2	HPE SN2010M 25GbE Switch Support
JD096C	8	HPE FlexNetwork X240 10G SFP+ to SFP+ 1.2m Direct Attach Copper Cable
JH234A	2	HPE X242 40G QSFP+ to QSFP+ 1m Direct Attach Copper Cable

HPE ProLiant DL380 Gen10 (dHCI enabled)

HPE ProLiant DL380 Gen10 servers are adaptable for diverse workloads and environments. The secure 2P 2U HPE ProLiant DL380 Gen10 server delivers world-class performance with the right balance of expandability and scalability. The combination of being designed for supreme versatility and resiliency while being backed by a comprehensive warranty makes it ideal for multiple environment—from containers to cloud to big data. HPE ProLiant DL380 Gen10 is the industry's most trusted compute platform.

TABLE 3. HPE ProLiant DL380 Gen10 (dHCI enabled) components

Part number	Quantity	Description
R0R08A	1	HPE Nimble Storage dHCI Base Configuration Tracking
868703-B21	3	HPE ProLiant DL380 Gen10 8SFF Configure-to-order Server
868703-B21 ABA	3	HPE ProLiant DL380 Gen10 8SFF CTO Server
826854-L21	3	HPE DL380 Gen10 Intel® Xeon®-Gold 5118 (2.3GHz/12-core/105W) FIO Processor Kit
826854-B21	3	HPE DL380 Gen10 Intel Xeon-Gold 5118 (2.3GHz/12-core/105W) Processor Kit
826854-B21 0D1	3	Factory Integrated
815100-B21	24	HPE 32GB (1x32GB) Dual Rank x4 DDR4-2666 CAS-19-19-19 Registered Smart Memory Kit
815100-B21 0D1	24	Factory Integrated
P18420-B21	6	HPE 240 GB SATA 6G Read Intensive SFF (2.5in) SC 3yr Wty Multi Vendor SSD
P18420-B21 0D1	6	Factory Integrated
P01366-B21	3	HPE 96W Smart Storage Battery (up to 20 Devices) with 145mm Cable Kit
P01366-B21 0D1	3	Factory Integrated
804331-B21	3	HPE Smart Array P408i-a SR Gen10 (8 Internal Lanes/2GB Cache) 12G SAS Modular Controller
804331-B21 0D1	3	Factory Integrated

Part number	Quantity	Description
867334-B21	3	HPE Ethernet 10/25Gb 2-port 622FLR-SFP28 Converged Network Adapter
867334-B21 0D1	3	Factory Integrated
865414-B21	6	HPE 800W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit
865414-B21 0D1	6	Factory Integrated
BD505A	3	HPE iLO Advanced 1-server License with 3yr Support on iLO Licensed Features
BD505A 0D1	3	Factory Integrated
733660-B21	3	HPE 2U Small Form Factor Easy Install Rail Kit
733660-B21 0D1	3	Factory Integrated
R2H14A	3	HPE Nimble Storage dHCI for DL3x0 Server with Additional Custom ESXi FIO Software
H7J34A3	1	HPE 3Y Foundation Care 24x7 SVC
H7J34A3 WAH	3	HPE DL38x Gen10 Support

The following sections highlight the configuration that has been used to successfully deploy HPE Nimble Storage dHCI with Anthos.

HPE Nimble Storage AF40 (dHCI enabled)

HPE Nimble Storage All Flash Arrays are simple to deploy, use, and manage. This product is cloud-ready, and it deploys flash on-premises or in the cloud with common data services and mobility between all-flash and adaptive flash.

HPE Nimble Storage All Flash Arrays are supported for NVMe and SCM. The offering comes with multiple benefits:

- All-inclusive software licensing
- Flat support pricing
- Replacement of all hardware for upgrades
- The option to receive a free faster controller upgrade after three years
- Always-on availability (guaranteed six-nines of storage availability)
- HPE Store More Guarantee

TABLE 4. HPE Nimble Storage AF40 (dHCI enabled) components

Part number	Quantity	Description
Q8H41A	1	HPE Nimble Storage AF40 All Flash Dual Controller 10GBASE-T 2-port Configure-to-order Base Array
Q8C18B	1	HPE Nimble Storage 4x10GbE 4-port FIO Adapter Kit
Q8G61B	1	HPE Nimble Storage AF40/60/80 All Flash Array 46TB (24x1.92TB) FIO Flash Bundle
Q8J27A	2	HPE Nimble Storage C13 to C14 250V 10Amp 1.8m Universal FIO Power Cord
R2H12A	1	HPE Nimble Storage dHCI NOS PG FIO Software
R3P91A	1	HPE Nimble Storage AF/HF Array Standard Tracking
HT6Z2A3	1	HPE NS 3Y 4H Onsite Exchange Support
HT6Z2A3 ZG3	1	HPE NS 10GbE 4p Adapter Support
HT6Z2A3 ZFF	1	HPE NS AF40 All Flash Base Array Support
HT6Z2A3 ZGN	1	HPE NS AF40/60/80 46TB Flash Support

Software

During the validation process, Hewlett Packard Enterprise used VMware vSphere® 6.7 U3 as the virtualization layer for Anthos deployed on-premises.

For more information about the version of vSphere supported by Google, see the Google Cloud documentation page.

Application software

During the validation process, Hewlett Packard Enterprise used Anthos 1.3 for GKE on-prem.

For more information about the version of Anthos supported by Google, see the Google Cloud documentation page.

BEST PRACTICES AND CONFIGURATION GUIDANCE FOR THE PRIVATE CLOUD

This section summarizes the best practices for configuring the solution for an on-premises private cloud.

DNS

Use the same DNS server across the servers and the HPE Nimble Storage array configuration to ensure that all relevant host names resolve to the same IP addresses.

NTP

Use the same NTP server across the servers and the HPE Nimble Storage array configuration to ensure that the time is set the same for all components.

VLANs

To work properly, the HPE Nimble Storage dHCI solution requires at least three VLANs. The iSCSI VLAN must be configured in access mode only. The management VLAN must be the native VLAN in the trunk port (untagged).

Table 5 lists the VLAN requirements for the solution. The GKE on-prem configuration has a VLAN for Kubernetes node called the **GKE admin cluster** and a VLAN for user cluster called **GKE user1 cluster**. The **MLAG** VLAN interconnects the two switches.

TABLE 5. VLAN requirements for HPE Nimble Storage dHCI with Anthos

VLAN description	VLAN ID	VLAN mode	Note
Management/vMotion	40	mgmt_vlan = native	Native VLAN on ESXi port 1 and 2
iSCSI1 IP address range	1000	iscsi1_vlan = access	Native VLAN only
iSCSI2 IP address range	1001	Iscsi2_vlan = access	Native VLAN only
GKE admin cluster	60	GKE-admin = trunk	Trunk on ESXi port 1 and 2
GKE user1 cluster	500	GKE-user1 = trunk	Trunk on ESXi port 1 and 2
MLAG	4094		Port 20 and 21 used for MLAG

IP address ranges

During deployment, you must provide four different ranges of IP addresses. Use the information listed in Table 6 to determine how many IP addresses are needed.

TABLE 6. Determining the required number of contiguous IP addresses

IP address range description	VLAN ID	Number of contiguous IP addresses
Management IP address range	40	Each HPE ProLiant server requires two IP addresses: one for iLO and one for the ESXi management interface.
iSCSI1 IP address range	1000	Each HPE ProLiant server requires one IP address.
iSCSI2 IP address range	1001	Each HPE ProLiant server requires one IP address.
GKE admin cluster	60	Reserves at least a /24
GKE user1 cluster	500	Reserves at least a /24

Network switches configuration

Network is a key component of a successful GKE on-prem deployment. Before deploying the solution, be sure to review the networking requirements described in the Google Cloud data center requirements. Figure 1 shows the network switch configuration that was used in the Hewlett Packard Enterprise lab to accommodate GKE on-prem and HPE Nimble Storage dHCI.

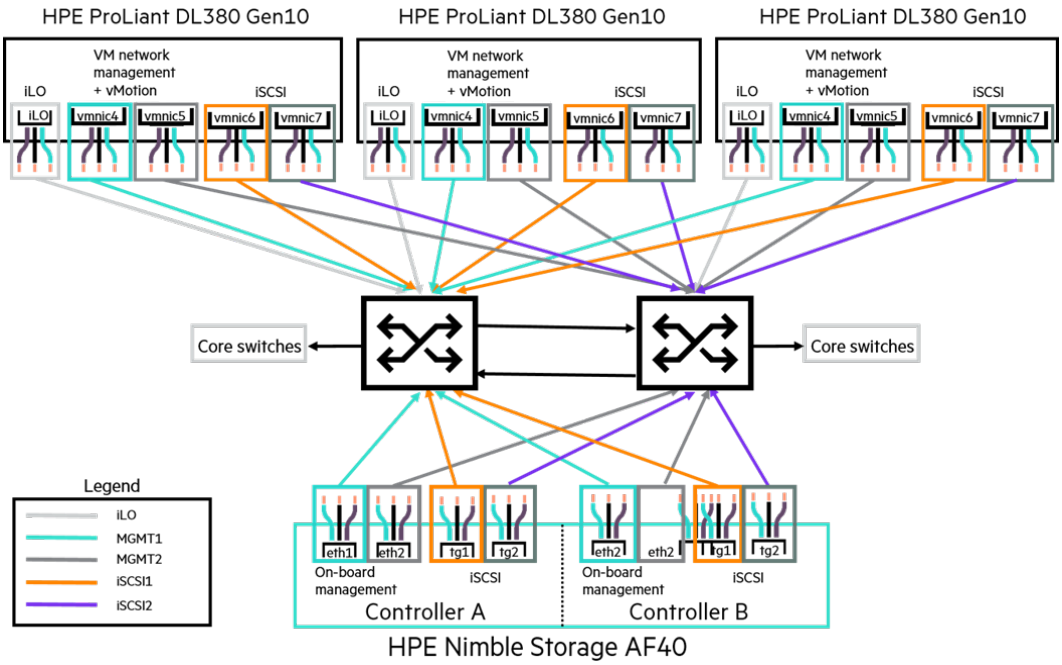

FIGURE 1. GKE-OP switch configuration

The following list categorizes the specific components of the configuration under storage, server, and network switches:

- Storage
  - HPE Nimble Storage AF40 with 2 x 10 Gbps per controller for iSCSI
  - Onboard management ports configured in **access mode**
  - iSCSI port #1 configured in **trunk mode**
    - iSCSI VLAN ID 1000 = Native
    - GKE user1 VLAN ID 500 = Trunk
  - iSCSI Port #2 configured in **trunk mode**
    - iSCSI VLAN ID 1001 = Native
    - GKE user1 cluster VLAN ID 500 = Trunk
- Server
  - HPE DL380 Gen10 with 4 x 10 Gbps
  - iLO configured in **access mode**
  - 2 ports used for management, vMotion, GKE admin cluster, and GKE user1 cluster
    - Management VLAN ID 40 = Native
    - GKE admin cluster VLAN ID 60 = Trunk
    - GKE user1 cluster VLAN ID 500 = Trunk
  - iSCSI port #1 in **access mode**
    - iSCSI VLAN ID 1000
  - iSCSI port #2 In **access mode**
    - iSCSI VLAN ID 1001
- Network switches
  - HPE StoreFabric SN2010M
  - 18 x 10/25 Gbps ports per switch
  - MLAG configured using 2 x 40 Gbps ports
  - 2 uplinks to HPE core switches

For more information, see Appendix B: HPE M-Series switches configuration, which provides the HPE M-Series configuration example that was used in testing.

Router configuration

To accommodate GKE on-prem, internet access is required. A dedicated router/firewall provides internet access to the GKE on-prem environment to route the traffic between subnet and internet.

Figure 2 shows the routing services that are used in the configuration. The router can route traffic in the following ways:

- From the GKE admin cluster and the GKE user1 cluster to the internet
- From the GKE admin cluster and the GKE user1 cluster subnet to the management subnet
- From the management subnet to the GKE admin cluster and the GKE user1 cluster

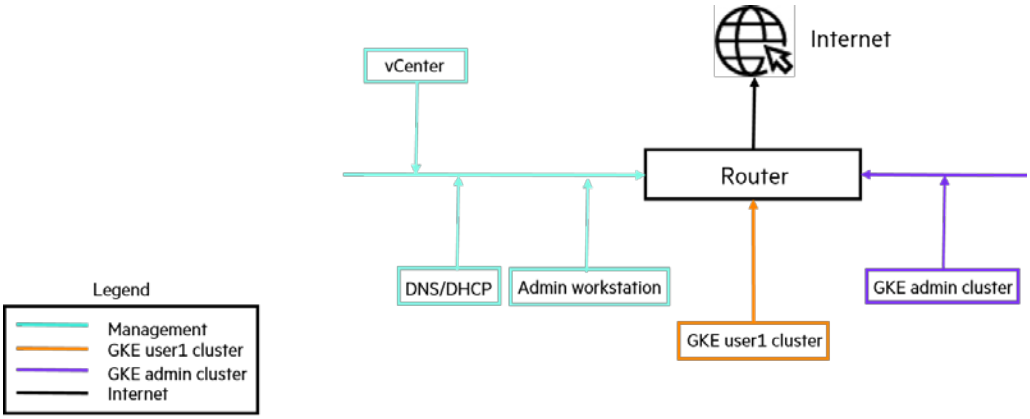

FIGURE 2. Routing services

Seesaw configuration

GKE on-prem requires a load balancer; any layer 4 load balancer can be used. With Anthos 1.3, you can now deploy the Seesaw load balancer as part of the overall deployment. During testing, we deployed Seesaw in high-availability (HA) mode. This choice ensured that GKE would take care of configuring the Seesaw load balancer when the service type **load balancer** was created within Kubernetes.

As Figure 3 shows, the Seesaw VM uses two interfaces:

- One in the GKE subnet to which it belongs
- One for HA

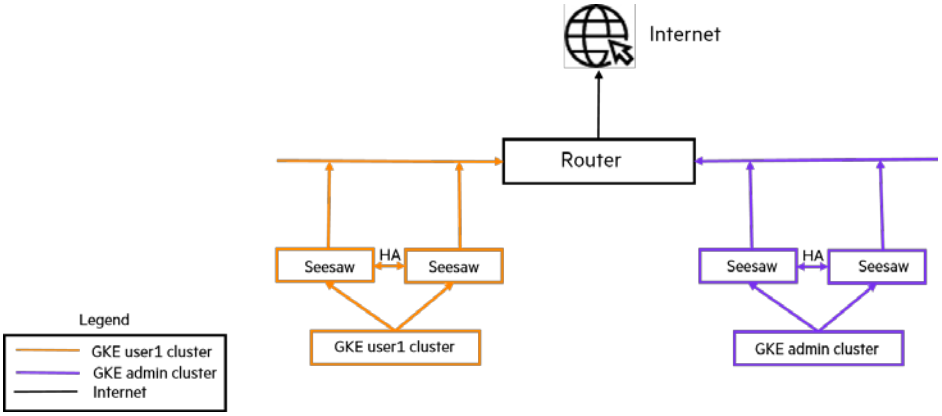

FIGURE 3. Seesaw router configuration

For more information about deployment and Anthos requirements, see the Google Cloud data center requirements.

CAPACITY AND SIZING

The hardware required for Anthos deployed on-premises varies depending on the workload and user requirements. Factors that affect the requirements include the following conditions:

- Types of workload
- Performance and scalability required
- Container images to run
- Deployment type (development, testing, staging, or production)

Anthos deployed on-premises runs as a VM cluster, in which Kubernetes runs inside the VMs. Table 7 lists the minimum VM requirements for running Anthos deployed on-premises. The number of worker nodes of Kubernetes you might require will vary depending on the workload that you plan to run.

TABLE 7. Components of dHCI-enabled HPE Nimble Storage AF40

Name of requirement	Number of VMs	Specifications	Description
Admin cluster master	1	• 4 vCPU • 16 GB RAM • 40 GB hard disk	Runs the admin control plane.
Add-on VMs	2	• 4 vCPU • 16 GB RAM • 40 GB hard disk	Runs the admin control plane's add-on.
User cluster master	1	• 4 vCPU • 8 GB RAM • 40 GB hard disk	User cluster control plane. The number of VMs varies depending on the number of user clusters and on whether high availability is required.
User cluster worker node	3	• 4 vCPU • 8 GB RAM • 40 GB hard disk	User cluster work node. The number of VMs varies depending on the workloads.

This Reference Configuration provides four recommended hardware configurations based on the following assumptions:

- RAM requirement per pod: 2 GB
- Maximum number of pods per worker node: 100
- 12 microservices per physical core

No over commitment or over allocation of RAM and CPU

Appendix A: Bills of materials describes the hardware configurations that Hewlett Packard Enterprise recommends based on the workload estimates:

- Entry-level configuration (up to 300 pods)
- Medium configuration (up to 1000 pods)
- Large configuration (up to 2500 pods)
- Extra-large configuration (up to 5000 pods)

DEPLOYING THE SOLUTION

The following sections show how to deploy the HPE Nimble Storage dHCI solution, including the Anthos cloud software.

A feature called **Stack Setup** simplified the deployment of the VMware environment during testing. This key component of the solution automates the following required tasks:

- Configuring the HPE Nimble Storage array and the HPE ProLiant server
- Deploying a new VMware vCenter Server® instance
- Creating a new data center, cluster, and datastore

The following list of required deployment tasks shows how the user interface appears at each step in the process:

1. Create a new VMware vCenter Server instance.

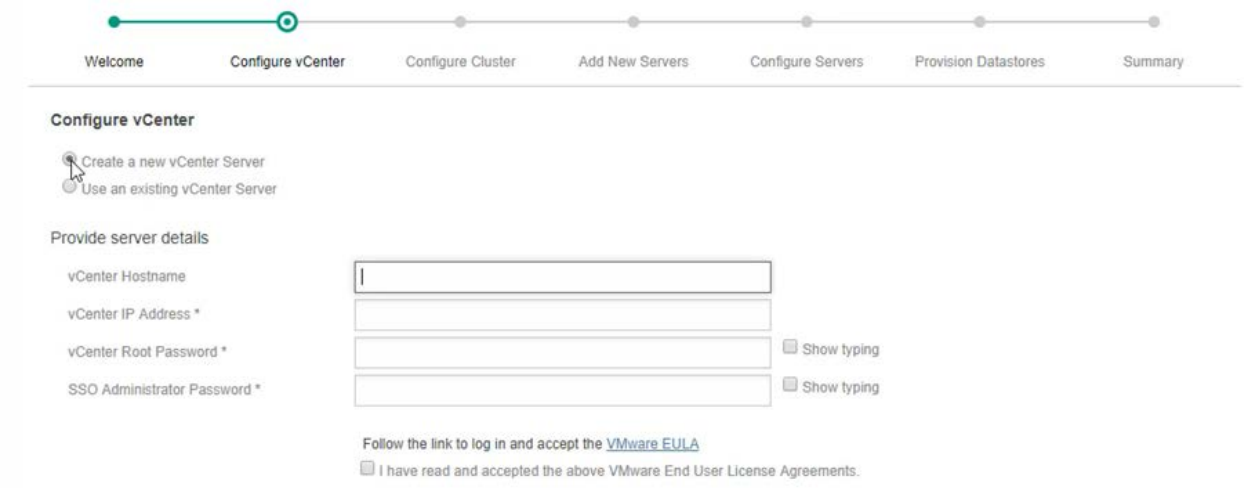

FIGURE 4. Configure vCenter

2. Create a new data center and a new cluster.

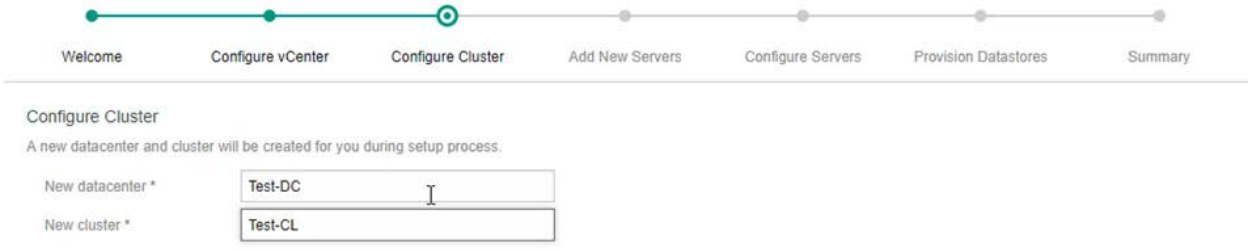

FIGURE 5. Configure the cluster

3. Select the servers to use for the VMware cluster.

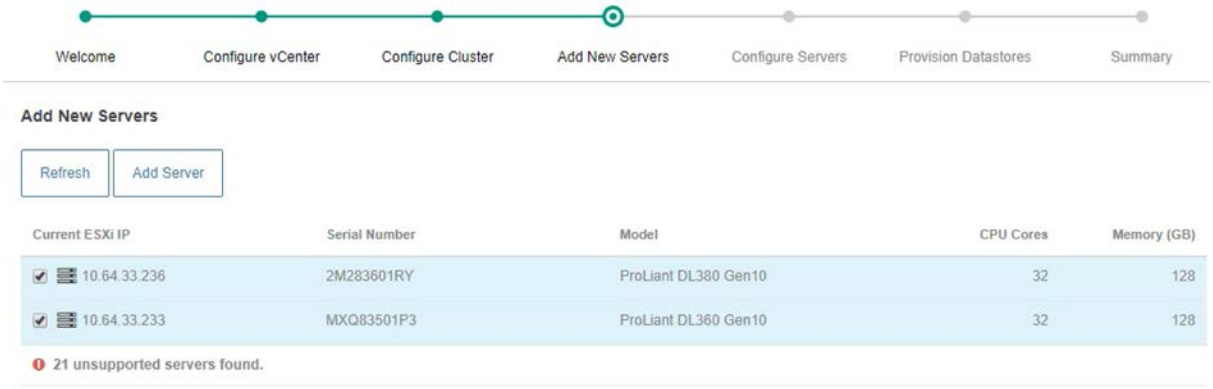

FIGURE 6. Add new servers

4. Create a VMFS datastore on the vSphere cluster where GKE on-prem is deployed.

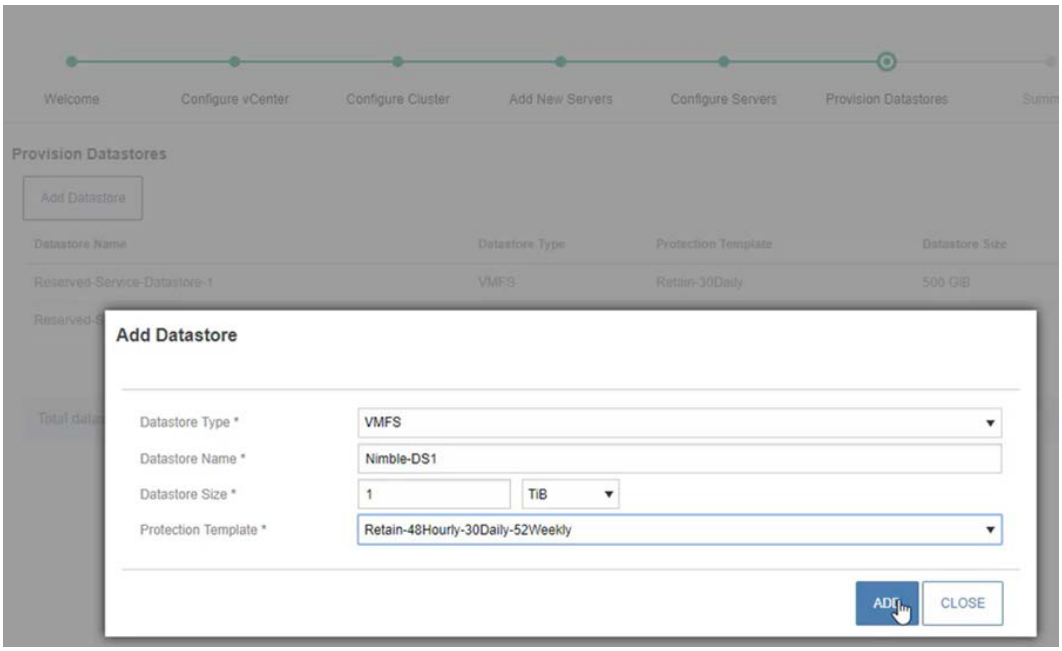

FIGURE 7. Provision the datastores

After the HPE Nimble Storage dHCI environment was deployed, three modifications were required before moving forward with the deployment of GKE on-prem:

- Create a new VM network called **GKE admin cluster** to provide a dedicated network for the GKE admin cluster.
- Create a new VM network called **GKE user1 cluster** to provide a dedicated network for the GKE user1 cluster.
- Create a resource pool for GKE on-prem.

Figure 8 shows where the network and the pool appear.

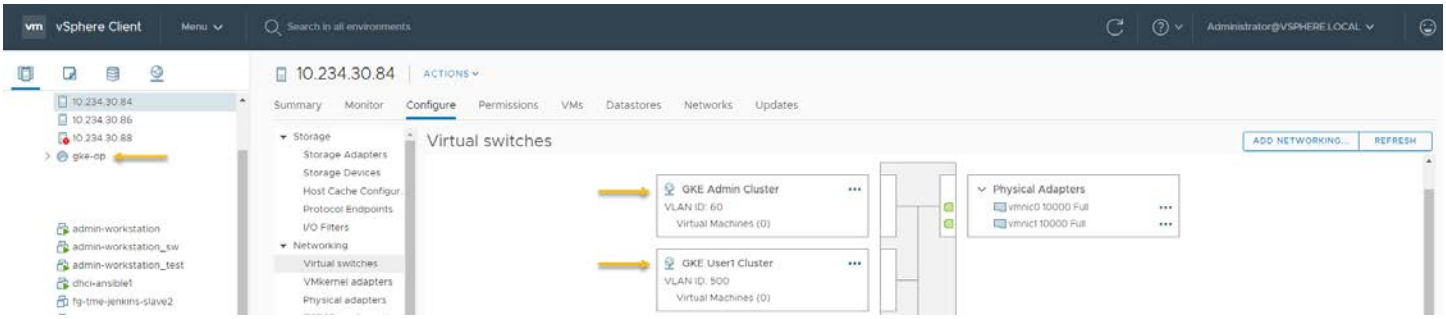

FIGURE 8. Create a dedicated network and resource pool for Anthos

After the completion of these few steps, the VMware environment was ready for GKE on-prem deployment. The process took less than 30 minutes and was far easier than the old way of having to deploy each component manually.

Deploying GKE on-prem

Google provides all required documentation and tools for deploying GKE on-prem. Because there are many different ways to deploy the GKE on-prem environment, Hewlett Packard Enterprise recommends that you refer to the latest Google documentation. This section simply highlights the tasks that must be performed.

You can choose from many options, some simple and some complex, for installing GKE on-prem and creating user clusters. For your initial installation, Hewlett Packard Enterprise recommends that you choose the simpler options. At the end of the installation, these options give you an admin cluster and a small user cluster that you can use to deploy Kubernetes workloads. After you gain experience with your small cluster, you might choose to perform a second installation to include some advanced features and also potentially create a larger user cluster.

For the initial installation, Hewlett Packard Enterprise recommends the following choices:

- Use the Seesaw load balancer bundled options so that the initial installation topics do not have to discuss the various load-balancing options.
- Use static IP addresses for cluster nodes and the admin workstation.
- Use the Google maintained container registry to host Anthos on VMware container images instead of a private Docker registry.
- Create a user cluster that has three worker nodes and one control plane node.

The following tasks are required for installing GKE on-prem and creating your first user cluster:

1. Get an Anthos subscription.
2. Prepare your vSphere environment.
3. Set up networking and load balancing.
4. Configure your Google Cloud project.
5. Create an admin workstation.
6. Create an admin cluster and a user cluster.

After GKE on-prem is fully deployed, it displays a report resembling the following one:

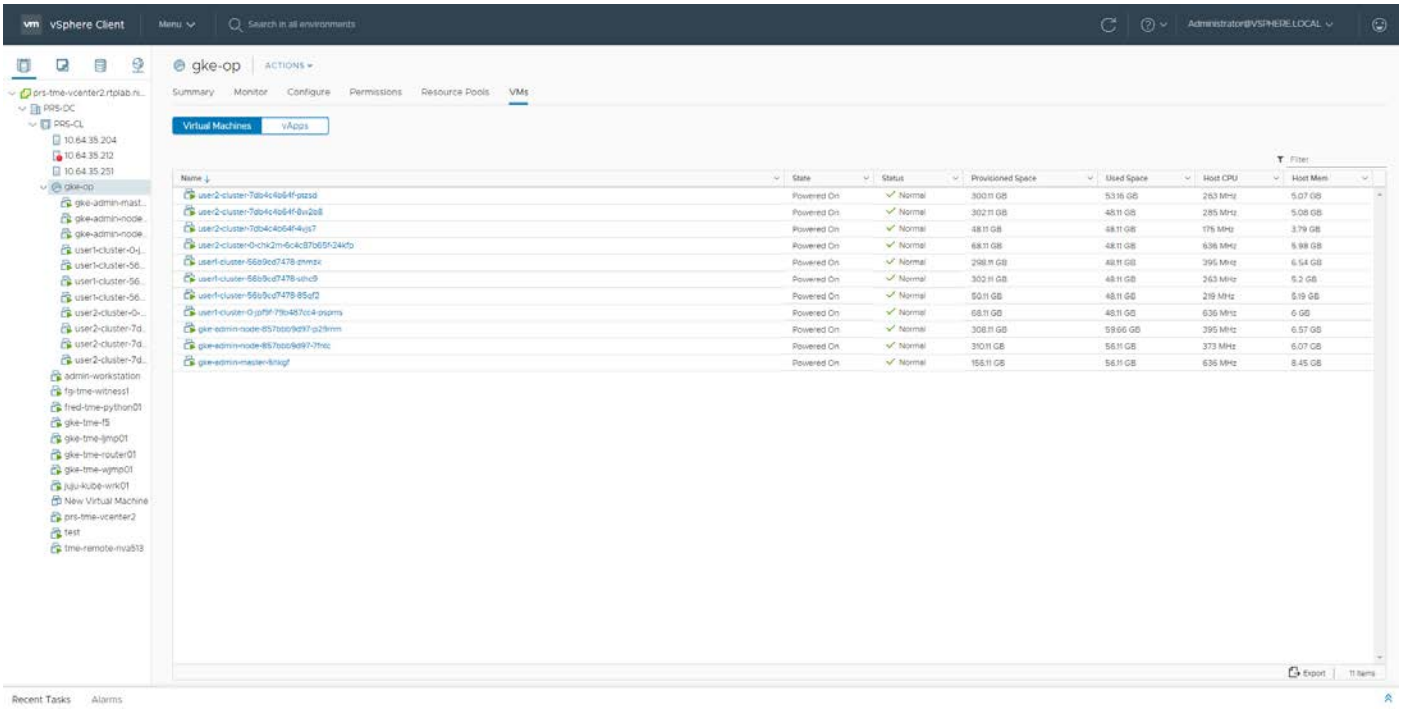

FIGURE 9. GKE on-prem deployed

At this point, the Kubernetes environment is ready, and you can use the `kubectl get node` command to check the status of your node.

```
ubuntu@gke-tme-wrks01:~$ kubectl get node --kubeconfig user1-cluster-dhci-kubeconfig
NAME            STATUS   ROLES    AGE   VERSION
user1-node1sw   Ready    <none>   20d   v1.15.7-gke.32
user1-node2sw   Ready    <none>   20d   v1.15.7-gke.32
user1-node3sw   Ready    <none>   20d   v1.15.7-gke.32
```

Deploying the HPE CSI driver for Kubernetes

Container Storage Interface (CSI) is a specification that allows container orchestration systems to implement a standardized interface to interact with storage. Kubernetes is one of the container orchestration implementations that supports CSI. The CSI specification has evolved at a rapid pace, steadily adding new features and capabilities. The Kubernetes community declared CSI stable and made it generally available (GA) in Kubernetes 1.13, which was released early in 2019.

CSI improves the quality of life for both development and operations staff. The developer gets a very consistent interface that makes it possible to consume storage resources. The operations person benefits from a very consistent deployment model because CSI drivers and relevant sidecar containers are simply Kubernetes workloads.

The new HPE CSI driver architecture introduces a sidecar deployment called a Container Storage Provider (CSP). The CSP is unique per storage platform (backend), and it responds to a minimal set of APIs that interface with the storage platform. The HPE CSI driver does all the heavy lifting on the nodes themselves that is common across storage platforms, such as attaching or detaching a block device or mounting a remote filesystem. The key here is that using the new HPE CSI driver to introduce a new storage platform to Kubernetes will now be quite trivial because you need to respond only at the API endpoint. No knowledge of CSI or Kubernetes is required.

The CSP can be written in any language that is capable of providing a RESTful API on a network port. Some vendors might choose to keep the CSP proprietary to protect their IP address because it is just a microservice and is part of a larger system that happens to be open source. Running the CSP on Kubernetes is also not a requirement because it might be an external service as part of an appliance or delivered as SaaS.

The HPE CSI driver provides scalable and persistent storage for stateful applications. For more information about the HPE CSI driver, see the HPE Nimble Storage section of the HPE Storage Container Orchestrator Documentation.

The following procedure shows how to deploy the plugin by using a YAML file for Kubernetes. The procedure uses the configuration example that was created in the lab.

1. Add the GKE user1 subnet on the HPE Nimble Storage array.

   Adding the subnet provides iSCSI connectivity to the user cluster where the HPE CSI driver is deployed.

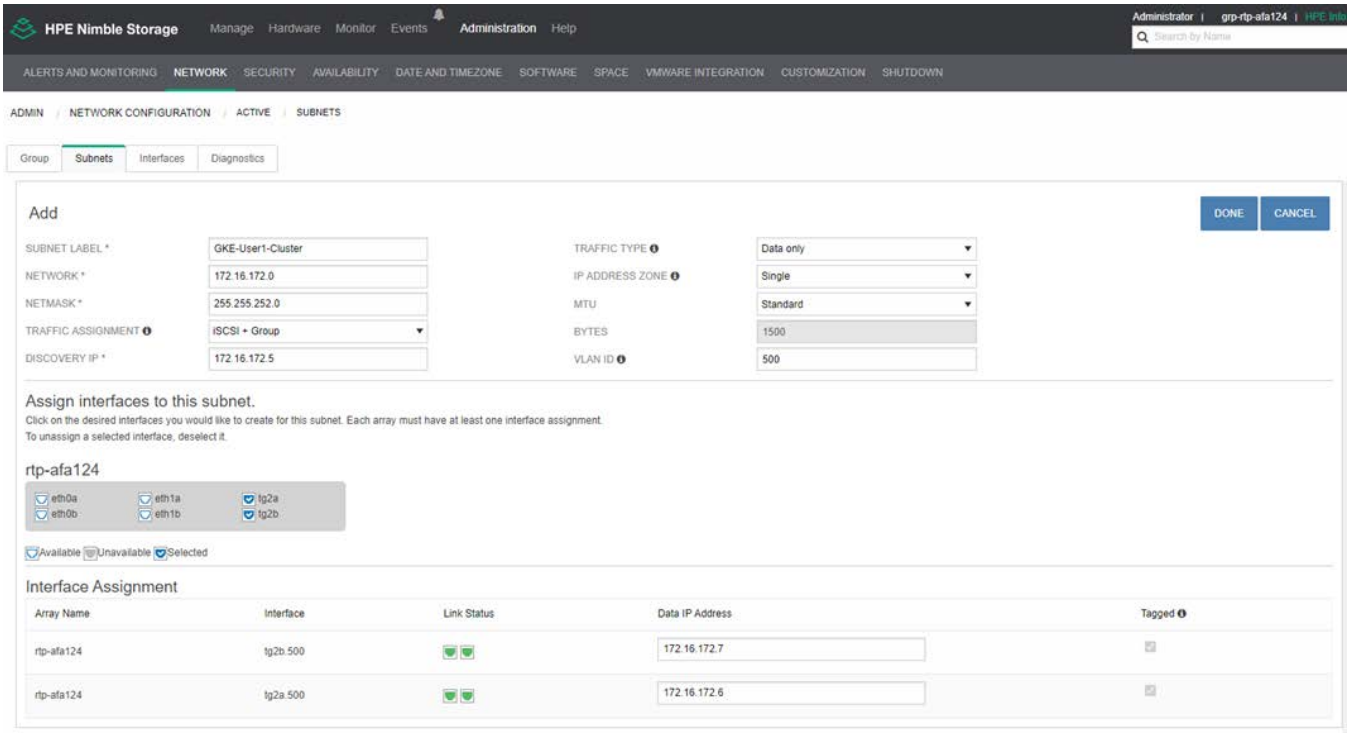

FIGURE 10. Adding the GKE user1 subnet

2. Create `secret.yaml`, add your HPE Nimble Storage array credential, and use `kubectl` to create a secret.

```
ubuntu@gke-tme-wrks01:~$ vi secret.yaml
apiVersion: v1
kind: Secret
metadata:
  name: nimble-secret
  namespace: kube-system
stringData:
  serviceName: nimble-csp-svc
  servicePort: "8080"
  backend: 10.234.30.71
  username: admin
data:
  # echo -n "admin" | base64
  password: YWRtaW4=

ubuntu@gke-tme-wrks01:~$ kubectl apply -f secret.yaml
```

3. Install the HPE Nimble Storage CSP driver for Kubernetes.

```
ubuntu@gke-tme-wrks01:~$ kubectl create -f https://raw.githubusercontent.com/hpe-storage/co-
deployments/master/yaml/csi-driver/v1.1.0/hpe-linux-config.yaml
ubuntu@gke-tme-wrks01:~$ kubectl create -f https://raw.githubusercontent.com/hpe-storage/co-
deployments/master/yaml/csi-driver/v1.1.0/nimble-csp.yaml
```

4. Install the HPE CSI driver for Kubernetes.

   Anthos 1.3 uses Kubernetes version 1.15.

```
ubuntu@gke-tme-wrks01:~$ kubectl create -f https://raw.githubusercontent.com/hpe-storage/co-
deployments/master/yaml/csi-driver/v1.1.0/hpe-csi-k8s-1.15.yaml
```

5. Check the status of the HPE Nimble Storage CSP for Kubernetes and the HPE CSI driver for Kubernetes.

```
ubuntu@admin-workstation:~$ kubectl get pod -n kube-system
NAME                                    READY   STATUS    RESTARTS   AGE
anthos-metering-6498df6d97-7rx24        1/1     Running   2          22d
anthos-metering-6498df6d97-npjw6        1/1     Running   2          22d
anthos-metering-6498df6d97-z4cwh        1/1     Running   0          22d
calico-node-8q2st                       1/1     Running   1          22d
calico-node-fwb8c                       1/1     Running   0          22d
calico-node-gqzqj                       1/1     Running   0          22d
calico-typha-864957894d-lql26           1/1     Running   0          22d
citadel-5696d8ff88-kl9rw                1/1     Running   0          22d
clientconfig-operator-5865bc94df-rwjnx  1/1     Running   1          22d
config-management-operator-786bdd5df-vg2zq  1/1  Running  1         22d
csp-service-5b655b5dc7-wnpjp              1/1     Running   0          4m38s
hpe-csi-controller-5d5db668b9-dqckr       4/4     Running   0          89s
hpe-csi-node-h4prn                        2/2     Running   0          90s
hpe-csi-node-nhsw5                        2/2     Running   0          90s
hpe-csi-node-rfl8d                       2/2     Running   0          90s
kube-dns-86d4b6c5c6-m69d8               3/3     Running   0          22d
kube-dns-86d4b6c5c6-znq49               3/3     Running   0          22d
kube-dns-autoscaler-64d974dc98-wkntd    1/1     Running   0          22d
kube-proxy-55lrq                        1/1     Running   1          22d
kube-proxy-jvht6                        1/1     Running   0          22d
kube-proxy-krqr6                        1/1     Running   0          22d
kube-state-metrics-5b866c97f8-dggss     2/2     Running   0          22d
logging-operator-7cf4b5567-xpcmc        1/1     Running   0          22d
metrics-server-6b4f754bdc-9clcs         1/1     Running   0          22d
node-exporter-7xw5p                     2/2     Running   0          22d
node-exporter-kmdrd                     2/2     Running   0          22d
node-exporter-tjf9c                     2/2     Running   2          22d
stackdriver-operator-64c5fd88f5-dnjmv   1/1     Running   1          22d
```

MANAGING THE HPE CSI DRIVER FOR KUBERNETES

During testing, a stateful application was deployed in the lab to validate the solution and test the HPE CSI driver. This application uses WordPress and MySQL as the back-end database. Figure 11 shows the testing configuration.

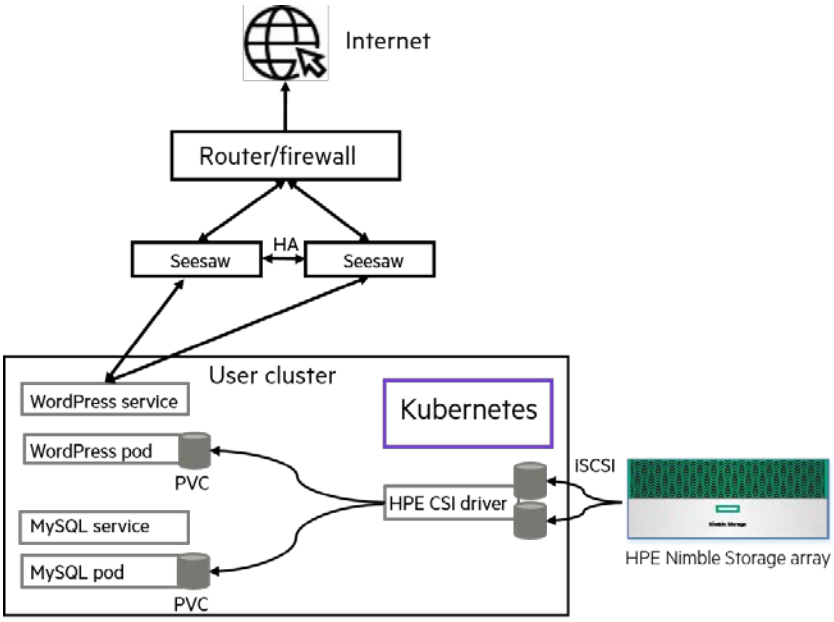

FIGURE 11. Lab testing configuration for the HPE CSI driver

WordPress Pod and MySQL Pod use persistent storage from HPE Nimble Storage to store data. For WordPress, a service of the **load balancer** type was created to expose the service through the Seesaw load balancer.

Deploying WordPress and MySQL

The following procedure shows how to deploy the application.

1. Create a storage class with the following parameters:

```
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name:  nimble-gke
provisioner: csi.hpe.com
parameters:
  csi.storage.k8s.io/fstype: xfs
  csi.storage.k8s.io/provisioner-secret-name: nimble-secret
  csi.storage.k8s.io/provisioner-secret-namespace: kube-system
  csi.storage.k8s.io/controller-publish-secret-name: nimble-secret
  csi.storage.k8s.io/controller-publish-secret-namespace: kube-system
  csi.storage.k8s.io/node-stage-secret-name: nimble-secret
  csi.storage.k8s.io/node-stage-secret-namespace: kube-system
  csi.storage.k8s.io/node-publish-secret-name: nimble-secret
  csi.storage.k8s.io/node-publish-secret-namespace: kube-system
  accessProtocol: "iscsi"
  destroyOnDelete: "true"
  description: "Volume from csi"
  dedupeEnabled: "true"
  limitIops: "76800"
  allowOverrides: "cloneOf"
```

2. Use the following `kubectl` command to create the storage class:

```
ubuntu@gke-tme-wrks01:~$ kubectl create -f sc.yaml
```

3. Use the following `kubectl` command to check that the storage class is created:

```
ubuntu@gke-tme-wrks01:~$ kubectl get sc
NAME                 PROVISIONER                    AGE
nimble-gke           csi.hpe.com                    22h
standard (default)   kubernetes.io/vsphere-volume   11d
```

4. Create the following file to create a secret for MySQL:

```
ubuntu@gke-tme-wrks01:~$ cat <<EOF >./kustomization.yaml
secretGenerator:
- name: mysql-pass
  literals:
  - password=YOUR_PASSWORD
EOF
```

5. Create the file `mysql-deployment.yaml` for MySQL deployment:

   This deployment uses the storage class that was created in step 2 to create a persistent volume claim (PVC).

```
apiVersion: v1
kind: Service
metadata:
  name: wordpress-mysql
  labels:
    app: wordpress
spec:
  ports:
    - port: 3306
  selector:
    app: wordpress
```

```
    tier: mysql
  clusterIP: None
---
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: mysql-pv-claim
  labels:
    app: wordpress
spec:
  accessModes:
    - ReadWriteOnce
  resources:
    requests:
      storage: 20Gi
  storageClassName: nimble-gke
---
apiVersion: apps/v1
kind: StatefulSet
metadata:
  name: wordpress-mysql
  labels:
    app: wordpress
spec:
  selector:
    matchLabels:
      app: wordpress
      tier: mysql
  serviceName: "wordpress-mysql"
  template:
    metadata:
      labels:
        app: wordpress
        tier: mysql
    spec:
      containers:
      - image: mysql:5.6
        name: mysql
        env:
        - name: MYSQL_ROOT_PASSWORD
          valueFrom:
            secretKeyRef:
              name: mysql-pass
              key: password
        ports:
        - containerPort: 3306
          name: mysql
        volumeMounts:
        - name: mysql-persistent-storage
          mountPath: /var/lib/mysql
      volumes:
      - name: mysql-persistent-storage
        persistentVolumeClaim:
          claimName: mysql-pv-claim
```

6. Create the file `wordpress-deployment.yaml` for MySQL deployment.

```
apiVersion: v1
kind: Service
metadata:
  name: wordpress
  labels:
    app: wordpress
spec:
  ports:
    - port: 80
  selector:
    app: wordpress
    tier: frontend
  type: LoadBalancer
  loadBalancerIP: 10.234.28.234
---
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: wp-pv-claim
  labels:
    app: wordpress
spec:
  accessModes:
    - ReadWriteOnce
  resources:
    requests:
      storage: 20Gi
  storageClassName: nimble-gke
---
apiVersion: apps/v1
kind: StatefulSet
metadata:
  name: wordpress
  labels:
    app: wordpress
spec:
  selector:
    matchLabels:
      app: wordpress
      tier: frontend
  serviceName: ”wordpress”
  template:
    metadata:
      labels:
        app: wordpress
        tier: frontend
    spec:
      containers:
      - image: wordpress:4.8-apache
        name: wordpress
        env:
        - name: WORDPRESS_DB_HOST
          value: wordpress-mysql
        - name: WORDPRESS_DB_PASSWORD
          valueFrom:
            secretKeyRef:
```

```
            name: mysql-pass
            key: password
      ports:
      - containerPort: 80
        name: wordpress
      volumeMounts:
      - name: wordpress-persistent-storage
        mountPath: /var/www/html
    volumes:
    - name: wordpress-persistent-storage
      persistentVolumeClaim:
        claimName: wp-pv-claim
```

7. Add the **mysql-deployment.yaml** and **wordpress-deployment.yaml** to kustomization.yaml.

```
ubuntu@gke-tme-wrks01:~$ cat <<EOF >>./kustomization.yaml
resources:
  - mysql-deployment.yaml
  - wordpress-deployment.yaml
EOF
```

8. Deploy the application.

```
ubuntu@gke-tme-wrks01:~$ kubectl apply -k kustomization.yaml
```

9. Verify that the application has been successfully deployed.

```
ubuntu@gke-tme-wrks01:~$ kubectl get pvc
NAME                STATUS   VOLUME     CAPACITY   ACCESS MODES     STORAGECLASS     AGE
mysql-pv-claim      Bound    pvc-4e0e...   20Gi       RWO            nimble-gke     2d19h
wp-pv-claim         Bound    pvc-f8a0...   20Gi       RWO            nimble-gke     2d19h
```

```
ubuntu@gke-tme-wrks01:~$ kubectl get pod
NAME                                READY   STATUS    RESTARTS   AGE
wordpress-0                         1/1     Running   0          2d19h
wordpress-mysql-0                   1/1     Running   0          2d19h
```

```
ubuntu@gke-tme-wrks01:~$ kubectl get svc
NAME                TYPE           CLUSTER-IP       EXTERNAL-IP     PORT(S)          AGE
kubernetes          ClusterIP      10.96.0.1        <none>          443/TCP          13d
wordpress           LoadBalancer   10.107.178.167   10.234.28.234   80:31231/TCP     2d19h
wordpress-mysql     ClusterIP      None             <none>          3306/TCP         2d19h
```

10. Open a web browser and enter the load balancer IP address to verify that you can reach the service.

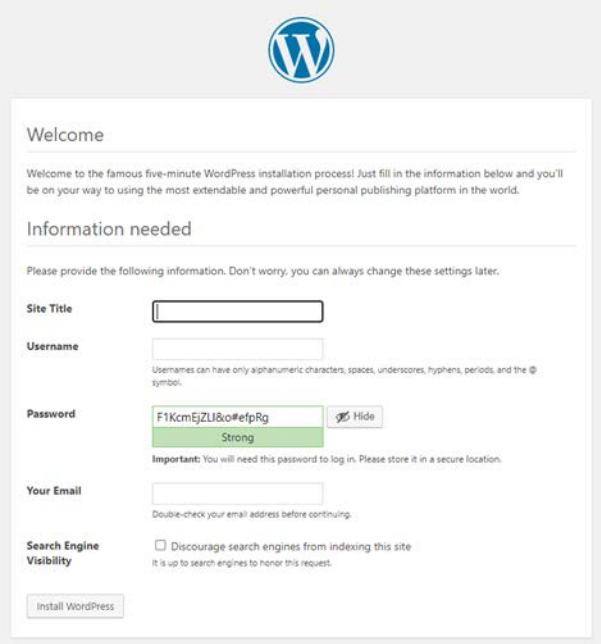

FIGURE 12. Verifying that you can reach the service

11. Install WordPress and make sure you can reach your website.

Cloning WordPress and MySQL

The following procedure shows how to clone the application by using the clone feature of the HPE CSI driver.

1. Get the PVC name for WordPress and the MySQL pod.

```
kubectl get pvc
NAME                 STATUS   VOLUME         CAPACITY   ACCESS MODES   STORAGECLASS   AGE
mysql-pv-claim       Bound    pvc-c0bb6f4c...   20Gi       RWO            nimble-gke     5m5s
wp-pv-claim          Bound    pvc-d83fefab...   20Gi       RWO            nimble-gke     5m5s
```

2. Create the following file to create a secret for MySQL:

```
ubuntu@gke-tme-wrks01:~$ cat <<EOF >./kustomization.yaml
secretGenerator:
- name: mysql-pass-dev
  literals:
  - password=YOUR_PASSWORD
EOF
```

3. Create the file `mysql-deployment.yaml` for MySQL deployment.

   This deployment creates a clone of the PVC that was used during the deployment of the first application.

```
apiVersion: v1
kind: Service
metadata:
  name: wordpress-mysql-dev
  labels:
    app: wordpress-dev
spec:
  ports:
    - port: 3306
  selector:
    app: wordpress-dev
    tier: mysql-dev
  clusterIP: None
---
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: mysql-dev-pv-claim
  labels:
    app: wordpress-dev
  annotations:
    csi.hpe.com/cloneOf: pvc-c0bb6f4c-7475-4e6b-9931-920c5526f8e5
spec:
  accessModes:
    - ReadWriteOnce
  resources:
    requests:
      storage: 20Gi
  storageClassName: nimble-gke
---
apiVersion: apps/v1
kind: StatefulSet
metadata:
  name: wordpress-mysql-dev
  labels:
    app: wordpress-dev
spec:
  selector:
    matchLabels:
      app: wordpress-dev
      tier: mysql-dev
  serviceName: "wordpress-mysql-dev"
  template:
    metadata:
      labels:
        app: wordpress-dev
        tier: mysql-dev
    spec:
      containers:
      - image: mysql:5.6
        name: mysql-dev
        env:
        - name: MYSQL_ROOT_PASSWORD
          valueFrom:
            secretKeyRef:
              name: mysql-pass-dev
              key: password
```

```
          ports:
        - containerPort: 3306
        name: mysql
        volumeMounts:
        - name: mysql-persistent-storage
          mountPath: /var/lib/mysql
      volumes:
      - name: mysql-persistent-storage
```

4. Create the file `wordpress-deployment.yaml` for MySQL deployment.

```
apiVersion: v1
kind: Service
metadata:
  name: wordpress-dev
  labels:
    app: wordpress-dev
spec:
  ports:
    - port: 80
  selector:
    app: wordpress-dev
    tier: frontend
  type: LoadBalancer
  loadBalancerIP: 10.234.28.235
---
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: wp-dev-pv-claim
  labels:
    app: wordpress-dev
  annotations:
    csi.hpe.com/cloneOf: pvc-d83fefab-d109-40ea-8091-97c5e8468913
spec:
  accessModes:
    - ReadWriteOnce
  resources:
    requests:
      storage: 20Gi
  storageClassName: nimble-gke
---
apiVersion: apps/v1
kind: StatefulSet
metadata:
  name: wordpress-dev
  labels:
    app: wordpress-dev
spec:
  selector:
    matchLabels:
      app: wordpress-dev
      tier: frontend
  serviceName: "wordpress-dev"
  template:
    metadata:
      labels:
        app: wordpress-dev
        tier: frontend
```

```
    spec:
      containers:
      - image: wordpress:4.8-apache
        name: wordpress-dev
        env:
        - name: WORDPRESS_DB_HOST
          value: wordpress-mysql-dev
        - name: WORDPRESS_DB_PASSWORD
          valueFrom:
            secretKeyRef:
              name: mysql-pass-dev
              key: password
        ports:
        - containerPort: 80
          name: wordpress
        volumeMounts:
        - name: wordpress-persistent-storage
          mountPath: /var/www/html
      volumes:
      - name: wordpress-persistent-storage
        persistentVolumeClaim:
          claimName: wp-dev-pv-claim
```

5. Add `mysql-deployment.yaml` and `wordpress-deployment.yaml` to `kustomization.yaml`.

```
ubuntu@gke-tme-wrks01:~$ cat <<EOF >>./kustomization.yaml
resources:
  - mysql-deployment.yaml
  - wordpress-deployment.yaml
EOF
```

6. Deploy the application.

```
ubuntu@gke-tme-wrks01:~$ kubectl apply -k ./
```

7. Verify that the application has been successfully deployed.

```
ubuntu@gke-tme-wrks01:~$ kubectl get pvc
NAME                 STATUS   VOLUME           CAPACITY   ACCESS MODES    STORAGECLASS   AGE
mysql-dev-pv-claim    Bound    pvc-13414a82-...   20Gi       RWO            nimble-gke     48s
mysql-pv-claim       Bound    pvc-c0bb6f4c-...   20Gi       RWO            nimble-gke     20m
wp-dev-pv-claim       Bound    pvc-7611b6fa-...   20Gi       RWO            nimble-gke     48s
wp-pv-claim          Bound    pvc-d83fefab-...   20Gi       RWO            nimble-gke     20m
```

```
ubuntu@gke-tme-wrks01:~$ kubectl get pod
NAME                                READY   STATUS    RESTARTS   AGE
wordpress-0                         1/1     Running   0          2d19h
wordpress-mysql-0                   1/1     Running   0          2d19h
wordpress-dev-0                     1/1     Running   0          2m6s
wordpress-mysql-dev-0               1/1     Running   0          2m6s
```

```
ubuntu@gke-tme-wrks01:~$ kubectl get svc
NAME                 TYPE           CLUSTER-IP       EXTERNAL-IP     PORT(S)          AGE
kubernetes           ClusterIP      10.96.0.1        <none>          443/TCP          13d
wordpress            LoadBalancer   10.107.178.167   10.234.28.234   80:31231/TCP     2d19h
wordpress-dev        LoadBalancer   10.98.43.242     10.234.28.235   80:32650/TCP     2m51s
wordpress-mysql      ClusterIP      None             <none>          3306/TCP         2d19h
wordpress-mysql-dev  ClusterIP      None             <none>          3306/TCP         2m51s
```

8. Open a web browser and enter the load balancer IP address to verify that you can reach the service.

   Because this deployment is cloned from the existing deployment, the same website is displayed.

SUMMARY

This Reference Configuration offers practical insights into how Hewlett Packard Enterprise brings together the key components of its portfolio to simply and easily create a containerized environment with unified management. The HPE Nimble Storage dHCI solution combines the powerful HPE ProLiant platform with the capable HPE Nimble Storage array to deliver unified setup, configuration, and management. These benefits are further enhanced by the advanced data management and services provided by the HPE CSI driver, coupled with Google Anthos integration. As one of the earliest adopters of containers for persistent storage, Hewlett Packard Enterprise has gained the depth of experience and capability to provide a true enterprise-grade solution for customers who want to leverage containerized stateful applications.

The solution enables several critical use cases for hybrid cloud environments:

- Dev/test workflows to develop in the cloud and deploy on-premises
- Instant availability of the right data for DevOps and CI/CD
- Ability to lift and shift applications between virtualized and bare metal environments
- Sophisticated analytics to analyze data in the cloud without disrupting the on-premises workloads

The detailed example provided in this Reference Configuration illustrates the configuration through which Hewlett Packard Enterprise has enabled industry-leading data management for containers that is developer-friendly, optimized for enterprise-grade scale, resiliency, and performance.

- Enterprise-grade and advanced data services and APIs
- Consistent user experience
- Everything as a service

The solution described in this Reference Configuration is just the beginning of the capabilities and integrations that Hewlett Packard Enterprise plans to offer in the future. Through integration with Anthos, this newly created solution will become the launching point for the ability to manage all of your containerized workloads, both on-premises and in the cloud. Hewlett Packard Enterprise will continue to provide insights and best practices for stateful and stateless applications though simplified setup, configuration, and monitoring.

APPENDIX A: BILLS OF MATERIALS

The following bills of materials contain electronic license to use (E-LTU) parts. Electronic software license delivery is available in most countries. Hewlett Packard Enterprise recommends purchasing electronic products rather than physical products (when available) for faster delivery and for the convenience of not having to track and manage confidential paper licenses. For more information about ordering, consult your Hewlett Packard Enterprise Reseller or Hewlett Packard Enterprise Sales Representative. For more details, see hpe.com/us/en/services/consulting.html.

NOTE

The part numbers listed were accurate at the time of testing and publication, but they are subject to change. The bill of materials does not include complete support options or other rack and power requirements. For more information about ordering, consult your Hewlett Packard Enterprise Reseller or Hewlett Packard Enterprise Sales Representative. For more details, see hpe.com/us/en/services/consulting.html.

TABLE A1. Bill of materials for an entry-level configuration

Part number	Quantity	Description
R0R08A	1	HPE Nimble Storage dHCI Base Configuration Tracking
868703-B21	3	HPE ProLiant DL380 Gen10 8SFF Configure-to-order Server
868703-B21 ABA	3	HPE DL380 Gen10 8SFF CTO Server
P02496-L21	3	HPE DL380 Gen10 Intel Xeon-Gold 5215 (2.5GHz/10-core/85W) FIO Processor Kit
P02496-B21	3	HPE DL380 Gen10 Intel Xeon-Gold 5215 (2.5GHz/10-core/85W) Processor Kit
P02496-B21 0D1	3	Factory Integrated

Part number	Quantity	Description
P00922-B21	48	HPE 16GB (1x16GB) Dual Rank x8 DDR4-2933 CAS-21-21-21 Registered Smart Memory Kit
P00922-B21 0D1	48	Factory Integrated
P18420-B21	6	HPE 240GB SATA 6G Read Intensive SFF (2.5in) SC 3yr Wty Multi Vendor SSD
P18420-B21 0D1	6	Factory Integrated
817718-B21	3	HPE Ethernet 10/25Gb 2-port 631SFP28 Adapter
817718-B21 0D1	3	Factory Integrated
P01366-B21	3	HPE 96W Smart Storage Battery (up to 20 Devices) with 145mm Cable Kit
P01366-B21 0D1	3	Factory Integrated
804331-B21	3	HPE Smart Array P408i-a SR Gen10 (8 Internal Lanes/2GB Cache) 12G SAS Modular Controller
804331-B21 0D1	3	Factory Integrated
817709-B21	3	HPE Ethernet 10/25Gb 2-port 631FLR-SFP28 Adapter
817709-B21 0D1	3	Factory Integrated
865414-B21	6	HPE 800W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit
865414-B21 0D1	6	Factory Integrated
BD505A	3	HPE iLO Advanced 1-server License with 3yr Support on iLO Licensed Features
BD505A 0D1	3	Factory Integrated
733660-B21	3	HPE 2U Small Form Factor Easy Install Rail Kit
733660-B21 0D1	3	Factory Integrated
R2H14A	3	HPE Nimble Storage dHCI for DL3x0 Server with Additional Custom ESXi FIO Software
Q9E63A	2	HPE StoreFabric SN2010M 25GbE 18SFP28 4QSFP28 Switch
JD089B	10	HPE X120 1G SFP RJ45 T Transceiver
844477-B21	12	HPE 25Gb SFP28 to SFP28 3m Direct Attach Copper Cable
845406-B21	2	HPE 100Gb QSFP28 to QSFP28 3m Direct Attach Copper Cable
H7J34A3	1	HPE 3Y Foundation Care 24x7 SVC
H7J34A3 WAH	3	HPE DL380 Gen10 Support
H7J34A3 W0P	2	HPE SN2010M 25GbE Switch Support
487655-B21	12	HPE BladeSystem c-Class 10GbE SFP+ to SFP+ 3m Direct Attach Copper Cable
HA124A1	1	HPE Technical Installation Startup SVC
HA124A1 5WX	1	HPE Nimble Storage dHCI Startup SVC
HB983A1	1	HPE Installation Comm Svrs Hourly SVC
HF385A1	2	HPE Training Credit Servers/HybridIT SVC
Q8H74A	1	HPE Nimble Storage AF20 All Flash Dual Controller 10GBASE-T 2-port Configure-to-order Base Array
Q8B74B	1	HPE Nimble Storage AF20 All Flash Array R2 23TB (24x960GB) FIO Flash Bundle
Q8C17B	1	HPE Nimble Storage 2x10GbE 4-port FIO Adapter Kit
Q8J27A	2	HPE Nimble Storage C13 to C14 250V 10Amp 1.8m Universal FIO Power Cord
R2H12A	1	HPE Nimble Storage dHCI NOS PG FIO Software
R3P91A	1	HPE Nimble Storage AF/HF Array Standard Tracking
HT6Z2A3	1	HPE NS 3Y 4H Onsite Exchange Support
HT6Z2A3 ZF4	1	HPE NS AF 23TB Flash Supp
HT6Z2A3 ZEA	1	HPE NS AF20 All Flash Base Array Supp
HT6Z2A3 ZG0	1	HPE NS 2x10GbE 4p Adptr Supp

TABLE A2. Bill of materials for a medium-sized configuration

Part number	Quantity	Description
R0R08A	1	HPE Nimble Storage dHCI Base Configuration Tracking
868703-B21	6	HPE ProLiant DL380 Gen10 8SFF Configure-to-order Server
868703-B21 ABA	6	HPE DL380 Gen10 8SFF CTO Server
P02504-L21	6	HPE DL380 Gen10 Intel Xeon-Gold 6238 (2.1GHz/22-core/140W) FIO Processor Kit
P02504-B21	6	HPE DL380 Gen10 Intel Xeon-Gold 6238 (2.1GHz/22-core/140W) Processor Kit
P02504-B21 0D1	6	Factory Integrated
P00924-B21	96	HPE 32GB (1x32GB) Dual Rank x4 DDR4-2933 CAS-21-21-21 Registered Smart Memory Kit
P00924-B21 0D1	96	Factory Integrated
P18420-B21	12	HPE 240GB SATA 6G Read Intensive SFF (2.5in) SC 3yr Wty Multi Vendor SSD
P18420-B21 0D1	12	Factory Integrated
817718-B21	6	HPE Ethernet 10/25Gb 2-port 631SFP28 Adapter
817718-B21 0D1	6	Factory Integrated
P01366-B21	6	HPE 96W Smart Storage Battery (up to 20 Devices) with 145mm Cable Kit
P01366-B21 0D1	6	Factory Integrated
804331-B21	6	HPE Smart Array P408i-a SR Gen10 (8 Internal Lanes/2GB Cache) 12G SAS Modular Controller
804331-B21 0D1	6	Factory Integrated
817709-B21	6	HPE Ethernet 10/25Gb 2-port 631FLR-SFP28 Adapter
817709-B21 0D1	6	Factory Integrated
865414-B21	12	HPE 800W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit
865414-B21 0D1	12	Factory Integrated
BD505A	6	HPE iLO Advanced 1-server License with 3yr Support on iLO Licensed Features
BD505A 0D1	6	Factory Integrated
733660-B21	6	HPE 2U Small Form Factor Easy Install Rail Kit
733660-B21 0D1	6	Factory Integrated
R2H14A	6	HPE Nimble Storage dHCI for DL3x0 Server with Additional Custom ESXi FIO Software
Q2F22A	2	HPE StoreFabric SN2410M 25GbE 48SFP28 8QSFP28 Switch
JD089B	12	HPE X120 1G SFP RJ45 T Transceiver
844477-B21	24	HPE 25Gb SFP28 to SFP28 3m Direct Attach Copper Cable
845406-B21	2	HPE 100Gb QSFP28 to QSFP28 3m Direct Attach Copper Cable
H7J34A3	1	HPE 3Y Foundation Care 24x7 SVC
H7J34A3 WAH	6	HPE DL38x Gen10 Support
H7J34A3 RCD	2	HPE SN2410M Storage Switch Support
487655-B21	12	HPE BladeSystem c-Class 10GbE SFP+ to SFP+ 3m Direct Attach Copper Cable
HA124A1	1	HPE Technical Installation Startup SVC
HA124A1 5WX	1	HPE Nimble Storage dHCI Startup SVC
HB983A1	4	HPE Installation Comm Svrs Hourly SVC
HF385A1	2	HPE Training Credit Servers/HybridIT SVC
Q8H41A	1	HPE Nimble Storage AF40 All Flash Dual Controller 10GBASE-T 2-port Configure-to-order Base Array
Q8C17B	1	HPE Nimble Storage 2x10GbE 4-port FIO Adapter Kit
Q8G61B	1	HPE Nimble Storage AF40/60/80 All Flash Array 46TB (24x1.92TB) FIO Flash Bundle
Q8J27A	2	HPE Nimble Storage C13 to C14 250V 10Amp 1.8m Universal FIO Power Cord

Part number	Quantity	Description
R2H12A	1	HPE Nimble Storage dHCI NOS PG FIO Software
R3P91A	1	HPE Nimble Storage AF/HF Array Standard Tracking
HT6Z2A3	1	HPE NS 3Y 4H Onsite Exchange Support
HT6Z2A3 ZG0	1	HPE NS 2x10GbE 4p Adptr Supp
HT6Z2A3 ZFF	1	HPE NS AF40 All Flash Base Array Supp
HT6Z2A3 ZGN	1	HPE NS AF40/60/80 46TB Flash Supp

TABLE A3. Bill of materials for a large configuration

Part number	Quantity	Description
R0R08A	1	HPE Nimble Storage dHCI Base Configuration Tracking
868703-B21	12	HPE ProLiant DL380 Gen10 8SFF Configure-to-order Server
868703-B21 ABA	12	HPE DL380 Gen10 8SFF CTO Server
P02504-L21	12	HPE DL380 Gen10 Intel Xeon-Gold 6238 (2.1GHz/22-core/140W) FIO Processor Kit
P02504-B21	12	HPE DL380 Gen10 Intel Xeon-Gold 6238 (2.1GHz/22-core/140W) Processor Kit
P02504-B21 0D1	12	Factory Integrated
P00924-B21	240	HPE 32GB (1x32GB) Dual Rank x4 DDR4-2933 CAS-21-21-21 Registered Smart Memory Kit
P00924-B21 0D1	240	Factory Integrated
P18420-B21	24	HPE 240GB SATA 6G Read Intensive SFF (2.5in) SC 3yr Wty Multi Vendor SSD
P18420-B21 0D1	24	Factory Integrated
817718-B21	12	HPE Ethernet 10/25Gb 2-port 631SFP28 Adapter
817718-B21 0D1	12	Factory Integrated
P01366-B21	12	HPE 96W Smart Storage Battery (up to 20 Devices) with 145mm Cable Kit
P01366-B21 0D1	12	Factory Integrated
804331-B21	12	HPE Smart Array P408i-a SR Gen10 (8 Internal Lanes/2GB Cache) 12G SAS Modular Controller
804331-B21 0D1	12	Factory Integrated
817709-B21	12	HPE Ethernet 10/25Gb 2-port 631FLR-SFP28 Adapter
817709-B21 0D1	12	Factory Integrated
865414-B21	24	HPE 800W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit
865414-B21 0D1	24	Factory Integrated
BD505A	12	HPE iLO Advanced 1-server License with 3yr Support on iLO Licensed Features
BD505A 0D1	12	Factory Integrated
733660-B21	12	HPE 2U Small Form Factor Easy Install Rail Kit
733660-B21 0D1	12	Factory Integrated
R2H14A	12	HPE Nimble Storage dHCI for DL3x0 Server with Additional Custom ESXi FIO Software
Q2F22A	2	HPE StoreFabric SN2410M 25GbE 48SFP28 8QSFP28 Switch
JD089B	12	HPE X120 1G SFP RJ45 T Transceiver
844477-B21	24	HPE 25Gb SFP28 to SFP28 3m Direct Attach Copper Cable
845406-B21	2	HPE 100Gb QSFP28 to QSFP28 3m Direct Attach Copper Cable
H7J34A3	1	HPE 3Y Foundation Care 24x7 SVC
H7J34A3 WAH	12	HPE DL38x Gen10 Support
H7J34A3 RCD	2	HPE SN2410M Storage Switch Support

Part number	Quantity	Description
487655-B21	12	HPE BladeSystem c-Class 10GbE SFP+ to SFP+ 3m Direct Attach Copper Cable
HA124A1	1	HPE Technical Installation Startup SVC
HA124A1 5WX	1	HPE Nimble Storage dHCI Startup SVC
HB983A1	10	HPE Installation Comm Svrs Hourly SVC
HF385A1	2	HPE Training Credit Servers/HybridIT SVC
Q8H42A	1	HPE Nimble Storage AF60 All Flash Dual Controller 10GBASE-T 2-port Configure-to-order Base Array
Q8C17B	1	HPE Nimble Storage 2x10GbE 4-port FIO Adapter Kit
Q8G61B	1	HPE Nimble Storage AF40/60/80 All Flash Array 46TB (24x1.92TB) FIO Flash Bundle
Q8J14A	2	HPE Nimble Storage C19 to C20 250V 16Amp 1.8m PDU Base Array FIO Power Cord
R2H12A	1	HPE Nimble Storage dHCI NOS PG FIO Software
R3P91A	1	HPE Nimble Storage AF/HF Array Standard Tracking
HT6Z2A3	1	HPE NS 3Y 4H Onsite Exchange Support
HT6Z2A3 ZG0	1	HPE NS 2x10GbE 4p Adptr Supp
HT6Z2A3 ZFJ	1	HPE NS AF60 All Flash Base Array Supp
HT6Z2A3 ZGN	1	HPE NS AF40/60/80 46TB Flash Supp

TABLE A4. Bill of materials for an extra-large configuration

Part number	Quantity	Description
R0R08A	1	HPE Nimble Storage dHCI Base Configuration Tracking
868703-B21	16	HPE ProLiant DL380 Gen10 8SFF Configure-to-order Server
868703-B21 ABA	16	HPE DL380 Gen10 8SFF CTO Server
P02504-L21	16	HPE DL380 Gen10 Intel Xeon-Gold 6238 (2.1GHz/22-core/140W) FIO Processor Kit
P02504-B21	16	HPE DL380 Gen10 Intel Xeon-Gold 6238 (2.1GHz/22-core/140W) Processor Kit
P02504-B21 0D1	16	Factory Integrated
P00926-B21	256	HPE 64GB (1x64GB) Quad Rank x4 DDR4-2933 CAS-21-21-21 Load Reduced Smart Memory Kit
P00926-B21 0D1	256	Factory Integrated
P18420-B21	32	HPE 240GB SATA 6G Read Intensive SFF (2.5in) SC 3yr Wty Multi Vendor SSD
P18420-B21 0D1	32	Factory Integrated
817718-B21	16	HPE Ethernet 10/25Gb 2-port 631SFP28 Adapter
817718-B21 0D1	16	Factory Integrated
P01366-B21	16	HPE 96W Smart Storage Battery (up to 20 Devices) with 145mm Cable Kit
P01366-B21 0D1	16	Factory Integrated
804331-B21	16	HPE Smart Array P408i-a SR Gen10 (8 Internal Lanes/2GB Cache) 12G SAS Modular Controller
804331-B21 0D1	16	Factory Integrated
817709-B21	16	HPE Ethernet 10/25Gb 2-port 631FLR-SFP28 Adapter
817709-B21 0D1	16	Factory Integrated
865414-B21	32	HPE 800W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit
865414-B21 0D1	32	Factory Integrated
BD505A	16	HPE iLO Advanced 1-server License with 3yr Support on iLO Licensed Features
BD505A 0D1	16	Factory Integrated
733660-B21	16	HPE 2U Small Form Factor Easy Install Rail Kit

Part number	Quantity	Description
733660-B21 0D1	16	Factory Integrated
R2H14A	16	HPE Nimble Storage dHCI for DL3x0 Server with Additional Custom ESXi FIO Software
Q2F22A	2	HPE StoreFabric SN2410M 25GbE 48SFP28 8QSFP28 Switch
JD089B	20	HPE X120 1G SFP RJ45 T Transceiver
Q6J40AAE	2	HPE StoreFabric SN2410M 25GbE 24-port Upgrade E-LT
844477-B21	64	HPE 25Gb SFP28 to SFP28 3m Direct Attach Copper Cable
845406-B21	2	HPE 100Gb QSFP28 to QSFP28 3m Direct Attach Copper Cable
H7J34A3	1	HPE 3Y Foundation Care 24x7 SVC
H7J34A3 WAH	16	HPE DL38x Gen10 Support
H7J34A3 RCD	2	HPE SN2410M Storage Switch Support
487655-B21	12	HPE BladeSystem c-Class 10GbE SFP+ to SFP+ 3m Direct Attach Copper Cable
HA124A1	1	HPE Technical Installation Startup SVC
HA124A1 5WX	1	HPE Nimble Storage dHCI Startup SVC
HB983A1	14	HPE Installation Comm Svrs Hourly SVC
HF385A1	2	HPE Training Credit Servers/HybridIT SVC
Q8H43A	1	HPE Nimble Storage AF80 All Flash Dual Controller 10GBASE-T 2-port Configure-to-order Base Array
Q8C17B	1	HPE Nimble Storage 2x10GbE 4-port FIO Adapter Kit
Q8G62B	1	HPE Nimble Storage AF40/60/80 All Flash Array 92TB (24x3.84TB) FIO Flash Bundle
Q8J14A	2	HPE Nimble Storage C19 to C20 250V 16Amp 1.8m PDU Base Array FIO Power Cord
R2H12A	1	HPE Nimble Storage dHCI NOS PG FIO Software
R3P91A	1	HPE Nimble Storage AF/HF Array Standard Tracking
HT6Z2A3	1	HPE NS 3Y 4H Onsite Exchange Support
HT6Z2A3 ZG0	1	HPE NS 2x10GbE 4p Adptr Supp
HT6Z2A3 ZFM	1	HPE NS AF80 All Flash Base Array Supp
HT6Z2A3 ZGP	1	HPE NS AF40/60/80 92TB Flash Supp

APPENDIX B: HPE M-SERIES SWITCHES CONFIGURATION

The following script provides a detailed view of an example HPE M-Series switches configuration.

```
## Running database "initial"
## Generated at 2019/10/22 14:18:12 +0000
## Hostname: rtp-b6-mlnx-3
##

##
## Running-config temporary prefix mode setting
##
```

```
no cli default prefix-modes enable

##
## Interface Ethernet buffer configuration
##
   traffic pool iscsi type lossless
   traffic pool iscsi memory percent 80.00
   traffic pool lossy-default memory percent 10.00
   traffic pool iscsi map switch-priority 4

##
## MLAG protocol
##
   protocol mlag

##
## Interface Ethernet configuration
##
   interface port-channel 10
   interface ethernet 1/1-1/22 speed 10G force
   interface ethernet 1/1-1/2 mtu 9216 force
   interface ethernet 1/4 mtu 9216 force
   interface ethernet 1/6-1/8 mtu 9216 force
   interface ethernet 1/10 mtu 9216 force
   interface ethernet 1/12 mtu 9216 force
   interface ethernet 1/1-1/2 flowcontrol receive on force
   interface ethernet 1/1-1/2 flowcontrol send on force
   interface ethernet 1/4 flowcontrol receive on force
   interface ethernet 1/4 flowcontrol send on force
   interface ethernet 1/6-1/8 flowcontrol receive on force
   interface ethernet 1/6-1/8 flowcontrol send on force
   interface ethernet 1/1-1/3 switchport mode hybrid
   interface ethernet 1/5 switchport mode hybrid
   interface ethernet 1/7-1/10 switchport mode hybrid
   interface ethernet 1/16 switchport mode trunk
   interface ethernet 1/21-1/22 channel-group 10 mode active
   interface ethernet 1/1 description Nimble CTRLA Port #1
   interface ethernet 1/2 description Nimble CTRLB Port  #1
   interface ethernet 1/3 description rtp-hpe-tme24 Mgmt Port #1
   interface ethernet 1/4 description rtp-hpe-tme24 Data Port #1
   interface ethernet 1/5 description rtp-hpe-tme26 Mgmt Port #1
   interface ethernet 1/6 description rtp-hpe-tme26 Data Port #1
   interface ethernet 1/16 description Port Trunk to Cisco Nexus 5K
   interface ethernet 1/21 description mlag ipl
   interface ethernet 1/22 description mlag ipl
   interface ethernet 1/16 shutdown
##
## LAG configuration
##
   lacp

##
## VLAN configuration
##
   vlan 20
   vlan 30
   vlan 40
   vlan 50
```

```
   vlan 60
   vlan 100
   vlan 500
   vlan 632
   vlan 1000-1001
   vlan 2069
   vlan 4094
   interface ethernet 1/16 switchport trunk allowed-vlan none
   interface ethernet 1/1-1/2 switchport access vlan 1000
   interface ethernet 1/3 switchport access vlan 632
   interface ethernet 1/4 switchport access vlan 1000
   interface ethernet 1/5 switchport access vlan 632
   interface ethernet 1/6 switchport access vlan 1000
   interface ethernet 1/7-1/8 switchport access vlan 1001
   interface ethernet 1/9 switchport access vlan 632
   interface ethernet 1/10 switchport access vlan 1000
   interface ethernet 1/12 switchport access vlan 1000
   vlan 30 name "Test"
   vlan 40 name "Island Network"
   vlan 50 name "Public Network"
   vlan 60 name "GKE Admin Network"
   vlan 100 name "Test1"
   vlan 500 name "GKE User1 Network"
   vlan 1000 name "Data #1"
   vlan 1001 name "Data #2"
   vlan 2069 name "Test9"
   interface ethernet 1/1 switchport hybrid allowed-vlan add 40
   interface ethernet 1/2 switchport hybrid allowed-vlan add 40
   interface ethernet 1/3 switchport hybrid allowed-vlan add 40
   interface ethernet 1/3 switchport hybrid allowed-vlan add 50
   interface ethernet 1/3 switchport hybrid allowed-vlan add 60
   interface ethernet 1/3 switchport hybrid allowed-vlan add 500
   interface ethernet 1/5 switchport hybrid allowed-vlan add 40
   interface ethernet 1/5 switchport hybrid allowed-vlan add 50
   interface ethernet 1/5 switchport hybrid allowed-vlan add 60
   interface ethernet 1/5 switchport hybrid allowed-vlan add 500
   interface ethernet 1/7 switchport hybrid allowed-vlan 40
   interface ethernet 1/8 switchport hybrid allowed-vlan 40
   interface ethernet 1/9 switchport hybrid allowed-vlan add 40
   interface ethernet 1/9 switchport hybrid allowed-vlan add 50
   interface ethernet 1/9 switchport hybrid allowed-vlan add 60
   interface ethernet 1/10 switchport hybrid allowed-vlan add 40
   interface ethernet 1/16 switchport trunk allowed-vlan add 632
   interface ethernet 1/16 switchport trunk allowed-vlan remove 1
##
## STP configuration
##
   spanning-tree mode rpvst
   spanning-tree port type edge default
   interface ethernet 1/1-1/18 spanning-tree port type edge

##
## L3 configuration
##
   vrf definition test
   interface vlan 4094
   interface vlan 4094 ip address 10.10.10.22/30 primary
##
```

```
## DCBX PFC configuration
##
   dcb priority-flow-control enable force
   dcb priority-flow-control priority 4 enable
   interface ethernet 1/1-1/2 dcb priority-flow-control mode off force
   interface ethernet 1/7-1/8 dcb priority-flow-control mode off force
   interface port-channel 10 dcb priority-flow-control mode on force

[7mlines 69-136 [27m[K[K##
## LLDP configuration
##
   lldp

##
## QoS switch configuration
##
   interface ethernet 1/1-1/15 qos trust port
   interface ethernet 1/1-1/15 qos default switch-priority 4
   interface ethernet 1/1-1/15 qos rewrite pcp
 ##
## DCBX Application Priority configuration
##
   dcb application-priority tcp iscsi 4

##
## DHCP relay configuration
##
   ip dhcp relay instance 1 vrf default
   ip dhcp relay instance 1 vrf-auto-helper
   ip dhcp relay instance 1 address 172.16.60.9
   ip dhcp relay instance 1 address 172.16.172.9
   ip dhcp relay instance 1 always-on

##
## MLAG configurations
##
   mlag-vip MLAG ip 10.64.12.220 /22 force
no mlag shutdown
   interface port-channel 10 ipl 1
   interface vlan 4094 ipl 1 peer-address 10.10.10.21
##
## Network interface configuration
##
no interface mgmt0 dhcp
   interface mgmt0 ip address 10.64.12.145 /22

##
## Other IP configuration
##
   hostname rtp-b6-mlnx-3
   ip domain-list rtplab.nimblestorage.com
   ip name-server 10.234.4.5
   ip route vrf default 0.0.0.0/0 10.64.12.2
##
## Other IPv6 configuration
##
no ipv6 enable
##
```

```
## Local user account configuration
##
   username admin password 7
$6$M.tHicu5$GNb6lOQlDcx64IukpehP7H/fZcQ4kXYfTydxJ.VFT9FpA5qruYvSLHFVhetapSXV7d2aC1iR4tLObi7URYiOw/
   username test capability admin
no username test disable
   username test full-name ""
   username test password 7
$6$w9Gsj1Pr$HGDVcl3VKPXhZI4pNKvyteFwGRODXb4QuisZyvbf6Nz6hfd/DWXhENkSXdcg7oP2SD/tm4flHM3jLZ7DLMSmm.

##
## AAA remote server configuration
##
# ldap bind-password ********
# radius-server key ********
# tacacs-server key ********
[7mlines 137-204 [27m[K[K
##
## Network management configuration
##
# web proxy auth basic password ********
no ntp server 10.80.100.16 disable
   ntp server 10.80.100.16 keyID 0
no ntp server 10.80.100.16 trusted-enable
   ntp server 10.80.100.16 version 4
   web http enable
no web https ssl secure-cookie enable

##
## Virtualization configuration
##
   virtual-machine enable

##
## X.509 certificates configuration
##
#
# Certificate name system-self-signed, ID a6b3c0fa80589ef566b049dde6150f2de0a77266
# (public-cert config omitted since private-key config is hidden)

##
## Persistent prefix mode setting
##
cli default prefix-modes enable
```

APPENDIX C: THE CONFIG.YAML FILE

The following `config.yaml` file was used in deploying the tested solution.

```
# Absolute path to a GKE bundle on disk
bundlepath: "/var/lib/gke/bundles/gke-onprem-vsphere-1.3.1-gke.0-full.tgz"
# Specify which vCenter resources to use for deployment
vcenter:
  # The credentials and address GKE should use to connect to vCenter
  credentials:
    address: "tme-vcenter1.rtplab.nimblestorage.com"
    username: "administrator@vsphere.local"
```

```
    password: ”
  datacenter: ”PRS-DC”
  datastore: ”Nimble-DS1”
  cluster: ”PRS-CL”
  network: ”GKE Admin Cluster”
  resourcepool: ”gke-op”
  # Provide the name for the persistent disk to be used by the deployment (ending
  # in .vmdk). Any directory in the supplied path must be created before deployment.
  # Not required when adding additional user clusters
  datadisk: ”mydisk-gke.vmdk”
  # Provide the path to vCenter CA certificate pub key for SSL verification
  cacertpath: ”/home/ubuntu/vcenter.pem”
# Specify the proxy configuration.
proxy:
  # The URL of the proxy
  url: ”http://web-proxy.corp.hpecorp.net:8080”
  # The domains and IP addresses excluded from proxying
  noproxy: ”127.0.0.1,localhost,.nimblestorage.com,10.0.0.0/8,192.0.0.0/8,172.0.0.0/8,.hpe.com”
# Specify admin cluster settings for a fresh GKE On-Prem deployment. Omit this section
# and use the --adminconfig flag when adding a new user cluster to an existing deployment
admincluster:
  # In-Cluster vCenter configuration
  vcenter:
    # If specified it overwrites the network field in global vcenter configuration
    network: ””
  # # The absolute or relative path to the yaml file to use for static IP allocation.
  # # Do not include if using DHCP
  ipblockfilepath: ”/home/ubuntu/master-hostconfig.yaml”
  # # Specify pre-defined nodeports if using ”manual” load balancer mode
  manuallbspec:
    ingresshttpnodeport: 32002
    ingresshttpsnodeport: 32003
    controlplanenodeport: 32001
    #addonsnodeport: 32004
  # Specify the already-existing partition and credentials to use with F5
  loadbalancerconfig:
    ipblockfilepath: ”/home/ubuntu/admin-hostconfig.yaml”
    vrid: 124
    vip: 10.234.28.239
    cpus: 8
    memorymb: 8192
    enableha: true
    antiaffinitygroups:
      enabled: true
   network: ”GKE Admin Cluster”
  # The VIPs to use for load balancing
  vips:
    # Used to connect to the Kubernetes API
    controlplanevip: ”10.234.28.243”
    # Shared by all services for ingress traffic
    ingressvip: ”10.234.28.244”
    # # Used for admin cluster addons (needed for multi cluster features). Must be the same
    # # across clusters
    #addonsvip: ”MASTER_ADDON_VIPS”
  # The Kubernetes service CIDR range for the cluster. Must not overlap with the pod
  # CIDR range
  serviceiprange: 10.96.232.0/24
  # The Kubernetes pod CIDR range for the cluster. Must not overlap with the service
```

```
  # CIDR range
  podiprange: 192.168.0.0/16
# Specify settings when deploying a new user cluster. Used both with a fresh deployment
# or when adding a new cluster to an existing deployment.
usercluster:
# In-Cluster vCenter configuration
  vcenter:
    # If specified it overwrites the network field in global vcenter configuration
    network: ""
  # # The absolute or relative path to the yaml file to use for static IP allocation.
  # # Do not include if using DHCP
  ipblockfilepath: "/home/ubuntu/usercluster1-hostconfig.yaml"
  # # Specify pre-defined nodeports if using "manual" load balancer mode
  manuallbspec:
    ingresshttpnodeport: 31002
     ingresshttpsnodeport: 31003
     controlplanenodeport: 31001
     addonsnodeport: 0
   # Specify the already-existing partition and credentials to use with F5
  loadbalancerconfig:
    ipblockfilepath: "/home/ubuntu/user-hostconfig.yaml"
    vrid: 125
    vip: 10.234.30.242
    cpus: 8
    memorymb: 8192
    enableha: true
    antiaffinitygroups:
      enabled: true
    network: "GKE User1 Cluster"
  # The VIPs to use for load balancing
  vips:
    # Used to connect to the Kubernetes API
    controlplanevip: "10.234.30.245"
    # Shared by all services for ingress traffic
    ingressvip: "10.234.30.246"
  # A unique name for this cluster
  clustername: "user1-cluster-dhci"
  # User cluster master nodes must have either 1 or 3 replicas
  masternode:
    cpus: 4
    memorymb: 8192
    # How many machines of this type to deploy
    replicas: 1
  # The number of worker nodes to deploy and their size. Min. 2 replicas
  workernode:
    cpus: 4
    memorymb: 8192
    # How many machines of this type to deploy
    replicas: 3
  # The Kubernetes service CIDR range for the cluster
  serviceiprange: 10.96.0.0/12
  # The Kubernetes pod CIDR range for the cluster
  podiprange: 192.168.0.0/16
# Which load balancer mode to use "Manual" or "Integrated"
lbmode: "Bundled"
# Specify which GCP project to connect your GKE clusters to
gkeconnect:
  projectid: ""
```

```
  # The absolute or relative path to the key file for a GCP service account used to
  # register the cluster
  registerserviceaccountkeypath: "/home/ubuntu/register-key.json"
  # The absolute or relative path to the key file for a GCP service account used by
  # the GKE connect agent
  agentserviceaccountkeypath: "/home/ubuntu/connect-key.json"
# Specify which GCP project to connect your logs and metrics to
#stackdriver:
  #projectid: ""
  # A GCP region where you would like to store logs and metrics for this cluster.
  #clusterlocation: ""
  #enablevpc: false
  # The absolute or relative path to the key file for a GCP service account used to
  # send logs and metrics from the cluster
  #serviceaccountkeypath: ""
# # Optionally use a private Docker registry to host GKE images
# privateregistryconfig:
#   # Do not include the scheme with your registry address
#   credentials:
#     address: ""
#     username: ""
#     password: ""
#   # The absolute or relative path to the CA certificate for this registry
#   cacertpath: ""
# The absolute or relative path to the GCP service account key that will be used to
# pull GKE images
gcrkeypath: "/home/ubuntu/access-key.json"
# Configure kubernetes apiserver audit logging
#cloudauditlogging:
  #projectid: ""
  # A GCP region where you would like to store audit logs for this cluster.
  #clusterlocation: ""
  # The absolute or relative path to the key file for a GCP service account used to
  # send audit logs from the cluster
  #serviceaccountkeypath: ""
```

RESOURCES AND ADDITIONAL LINKS

David Wang blog post about the HPE Nimble Storage dHCI announcement, community.hpe.com/t5/Around-the-Storage-Block/Give-me-morehyperconverged-please/ba-p/7050469#.XQzV-OhKg2w

HPE DEV, hpedev.io (Platforms ⟶ 3PAR/Primera/Nimble Storage)

HPE DEV, Slack (https://slack.hpedev.io)

HPE and Google Cloud Partnership, hpe.com/us/en/alliance/google-cloud.html#Solutions

HPE InfoSight, infosight.hpe.com

HPE Networking, hpe.com/networking

HPE Reference Architectures, https://www.hpe.com/docs/reference-architecture

HPE Servers, hpe.com/servers

HPE Solution Brief: Intelligent Data Management for Containers, hpe.com/v2/Getdocument.aspx?docname=a00096698ENW

HPE Storage, hpe.com/storage

HPE Storage Container Orchestrator Documentation, scod.hpedev.io

HPE GreenLake Advisory and Professional Services, https://www.hpe.com/us/en/services/consulting.html

LinkedIn, linkedin.com/showcase/hpestorage

To help us improve our documents, please provide feedback at hpe.com/contact/feedback.

a50002388enw, July 2021

Hewlett Packard Enterprise