

Google Cloud's Anthos Solution with Bare metal and VMware on PowerFlex family

May 2021

000050

White Paper

Abstract

This white paper highlights a hybrid cloud solution utilizing Dell EMC PowerFlex software-defined storage and Google Cloud's Anthos 1.6 for deploying, managing, and optimizing applications—legacy as well as cloud-native across different environments.

Dell Technologies Solutions



Google Cloud

DELLTechnologies

Copyright

The information in this publication is provided as is. Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © 2021 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Intel, the Intel logo, the Intel Inside logo and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries. Other trademarks may be trademarks of their respective owners. Published in the USA May 2021, White Paper, 000050.

Dell Inc. believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

Contents

Executive summary	5
Introduction	6
Revisions	6
Audience	6
Terminology	6
We value your feedback	7
Product overview	8
PowerFlex family	8
PowerFlex software components	8
PowerFlex consumption options	9
PowerFlex deployment architectures	10
Google Cloud's Anthos	10
Solution overview for Anthos on VMware	12
Prerequisites	12
Physical design	14
Logical design	15
Deploying Anthos on VMware	17
Prerequisites	17
Prepare vCenter	17
Deploy and configure Layer 4 network load balancer	17
Deploy Anthos GKE on-prem admin VM on PowerFlex rack	17
Create Anthos GKE on-prem admin and user cluster on PowerFlex rack	20
Solution Overview for Anthos on bare metal	23
Prerequisites	23
Logical design	24
Deploying Anthos on bare metal	26
Prerequisites	26
Prepare Linux administration machine	26
Deploy and Configure Layer 4 network load balancer	27
Creating cluster with Anthos on bare metal	28

Deploying Anthos applications on PowerFlex system..... 29
 Sample workload: NGINX web server deployment 29

Conclusion..... 32

Appendix 33
 Deploy and Configure the Seesaw load balancer 33
 Deploy and configure F5 Big-IP Virtual Edition appliance 33

Executive summary

PowerFlex is a software-defined storage platform designed to significantly reduce operational and infrastructure complexity, empowering organizations to move faster by delivering flexibility, elasticity, and simplicity with predictable performance and resiliency at scale. The PowerFlex family provides a foundation that combines compute as well as high performance storage resources in a managed, unified fabric. PowerFlex comes in flexible deployment options – integrated rack, appliance or Ready Nodes - that enable two-layer (compute and server SAN), single-layer (HCI), and/or storage only architectures. PowerFlex is ideal for high performance applications and databases, building an agile private cloud, or consolidating resources in heterogeneous environments.

Containers offer an abstraction mechanism at the application layer. Containers help lower the barrier of entry to developing microservice-based applications and enable developers to spend less time worrying about runtimes, dependencies, and differences between test, development, and production environments and more time innovating to meet the demands of today's dynamic business environment. Container Orchestrators (COs), like Kubernetes, make running containers at production scale possible by handling the complexity of managing hundreds or thousands of containers at any one time.

Google Cloud's Anthos is an integrated platform that lets you modernize how you develop, secure, and operate hybrid cloud and multi-cloud environments across the Google Cloud Platform (GCP) and on-premises infrastructure.

Working closely with Google Cloud, Dell EMC brings to market PowerFlex system configuration that is optimized to power the full Anthos stack including Anthos clusters on-prem (deployed on VMware or bare metal).

This paper talks about Anthos version 1.6.x and has been jointly validated by Google Cloud and Dell EMC on PowerFlex system. Customers can purchase Anthos knowing the PowerFlex provides an operationally efficient, life cycle managed and scalable option for on premises deployments.

Introduction

Revisions

Date	Description
April 2019	Initial release
May 2019	Branding update
July 2019	Update for Anthos Ready validation
April 2020	Updated deploying procedures and load-balancer information
June 2020	Updated VxFlex to PowerFlex as per new re-branding guidelines
Feb 2021	Updated to Anthos 1.6 with bare metal and VMware
May 2021	Updated Logical Design sections

Audience

The audience for this paper includes system engineers, field consultants, IT administrators, technical architects, and anyone else interested in configuring and deploying Anthos on Dell EMC PowerFlex family with PowerFlex as the underlying software defined storage layer.

Readers are expected to have an understanding and working knowledge of containers, Kubernetes, Google Cloud, Google Kubernetes Engine (GKE), VMware vSphere, Linux, and PowerFlex software-defined storage.

Terminology

The following table provides definitions for some of the terms that are used in this document.

Table 1. Terminology

Term	Definition
SDC	Storage Data Client for PowerFlex
SDS	Storage Data Server for PowerFlex
MDM	Meta Data Manager for PowerFlex
DNS	Domain Name System
DHCP	Dynamic Host Configuration Protocol
RCM	Release Certification Matrix
GCP	Google Cloud Platform
GKE	Google Kubernetes Engine running on GCP
GKE on-prem	Google Kubernetes Engine running on premises
GCE	Google Compute Engine

**We value your
feedback**

Dell Technologies and the authors of this document welcome your feedback on the solution and the solution documentation. Contact the Dell Technologies Solutions team by [email](#) or provide your comments by completing our [documentation survey](#).

Author: Praphul Krottapalli, Raghavendra Biligiri, and Ronald Kelley

Contributors from Google team: Quang Le, Joshua Padilla, and Denis Jatsiv

Note: For links to additional documentation for this solution, see [Dell Technologies Solutions Info Hub for PowerFlex](#).

Product overview

PowerFlex family PowerFlex is a software-defined storage platform designed to significantly reduce operational and infrastructure complexity, empowering organizations to move faster by delivering flexibility, elasticity, and simplicity with predictable performance and resiliency at scale. The PowerFlex family provides a foundation that combines compute as well as high performance storage resources in a managed unified fabric. PowerFlex comes in flexible deployment options - rack, appliance, or ready nodes-that enables disaggregated (two-layer), HCI (single-layer), or mixed architectures. PowerFlex is ideal for high performance applications and databases, building an agile private cloud, or consolidating resources in heterogeneous environments.




Figure 1. PowerFlex family

PowerFlex 3.5

This white paper is based on PowerFlex v3.5.1.2 and following are the highlights of new features in this version:

- Native asynchronous replication for two-layer or storage-only deployments.
- HTML-5 based GUI

For more details, see [PowerFlex 3.5 Release Notes](#).

PowerFlex software components

Software is the key differentiation and the “secret sauce” in the PowerFlex offering. PowerFlex software components not only provide software-defined storage services, they also help simplify infrastructure management and orchestration with comprehensive ITOM and LCM capabilities that span compute as well as storage infrastructure, from BIOS and Firmware to nodes, software and networking.

The core foundational component in the PowerFlex family that enables Software Defined Storage (SDS) services is called simply PowerFlex, to represent the core value it enables for the platform. Additionally, PowerFlex Manager is a comprehensive IT Operational

Management (ITOM) and Life Cycle Management (LCM) tool that drastically simplifies management and ongoing operation.

PowerFlex (VxFlex OS)

PowerFlex (**previously VxFlex OS**) is the software foundation of PowerFlex software-defined storage. It is a scale-out block storage service designed to deliver flexibility, elasticity, and simplicity with predictable high performance and resiliency at scale.

PowerFlex Manager

PowerFlex Manager is the software component in PowerFlex family that enables ITOM automation and life cycle management capabilities for PowerFlex systems.

PowerFlex consumption options

PowerFlex SDS platform is available in multiple consumption options to help customers meet their project and data center requirements. PowerFlex appliance and PowerFlex rack provide customers comprehensive IT Operations Management (ITOM) and life cycle management (LCM) of the entire infrastructure stack in addition to sophisticated high-performance, scalable, resilient storage services. PowerFlex appliance and PowerFlex rack are the two preferred and proactively marketed consumption options. PowerFlex is also available on VxFlex Ready Nodes without the ITOM and LCM capabilities.

Note: The brand for Ready Nodes continues to be VxFlex.

PowerFlex rack

PowerFlex rack is a software-defined storage platform designed to deliver flexibility, elasticity, and simplicity with predictable performance and resiliency at scale by combining compute as well as high performance storage resources in a managed unified network. This rack-scale engineered system, with integrated networking, enables customers to achieve the scalability and management requirements of a modern data center.

PowerFlex appliance

PowerFlex appliance is a software-defined storage platform designed to deliver flexibility, elasticity, and simplicity with predictable performance and resiliency at scale by combining compute as well as high performance storage resources in a managed unified network. This turnkey offer allows customers the flexibility and savings to bring their own compatible networking. With PowerFlex, customers deploy to match their initial needs and easily expand with massive scale potential, without having to compromise on performance and resiliency.

VxFlex Ready Nodes

VxFlex Ready Nodes are validated server building blocks configured for use with PowerFlex. They are available with thousands of configuration options and are available for customers who prefer to build their own environments.

PowerFlex deployment architectures

PowerFlex software-define storage offers flexibility of deployment architecture to help best meet the specific deployment and architectural requirements. PowerFlex can be deployed in a two-layer (Server SAN), single layer (HCI), or in storage-only architectures.

Two-layer (Server SAN) architecture

In two-layer architecture, nodes that provide storage capacity and host datasets are separated from nodes that host applications and workloads. PowerFlex manager provides LCM and IOTM for the entire infrastructure, including nodes that provide storage and nodes that host the applications. Compute and storage resources can be scaled by adding respective node to the cluster. This segregation of compute and storage resources can be helpful to minimize software licensing costs in certain situations. This architecture could be most suitable for hosting high-performance high-value databases and application workloads.

Single-layer (HCI) architecture

In this architecture, each node in the cluster contributes storage resources as well as hosts applications and workloads. This architecture allows you to scale your infrastructure uniformly and with a pre-defined building block that adds both storage and compute resources. PowerFlex Manager provides ITOM and LCM capabilities for the entire infrastructure. This architecture is most suitable for data center and workload consolidation.

Mixed architecture

Using PowerFlex storage-only nodes, a software-defined block storage environment is created that can be accessed and consumed by several applications and workloads that are hosted outside of PowerFlex cluster. PowerFlex Manager provides LCM and ITOM for the storage infrastructure. This is a suitable architecture where customer has existing computes infrastructure but needs high-performance SDS. This can be a starting point with the customer and may expand to a two-layer Server SAN deployment in the future as the external compute is migrated to PowerFlex.

Google Cloud's Anthos

Anthos provides you with a consistent platform for building and managing containerized applications across hybrid infrastructures and helps your developers become more productive across all environments. Anthos provides all the mechanisms that are required to bring your code into production reliably, securely, and consistently, while minimizing risk. Anthos is built on open-source technologies pioneered by Google Cloud, including Kubernetes, Istio, and Knative, enabling consistency between cloud and on premises environments like PowerFlex rack. Anthos GKE (on GCP and on-prem), Anthos Service Mesh, and Anthos Config Management are the core building blocks of Anthos. Anthos also includes with GCP Marketplace and integrations with platform-level services such as Stackdriver, Cloud Build, and Binary Authorization.

Anthos puts all your IT resources into a consistent development, management, and control framework, automating low-value and insecure tasks across your PowerFlex rack and GCP infrastructure.

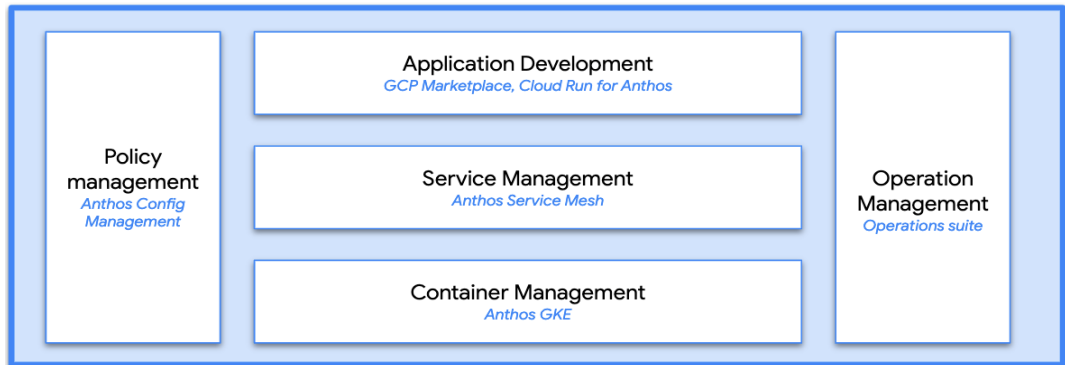


Figure 2. Google Cloud's Anthos Components

Within the context of GCP, the term hybrid cloud describes a setup in which common or interconnected services are deployed across multiple computing environments that include public cloud and on-premises. A hybrid cloud strategy lets you extend the capacity and capabilities of your IT, without upfront capital expense investments by using the public cloud and preserving your existing investments by adding one or more cloud deployments to your existing infrastructure. For more information, see [Hybrid and Multi-Cloud Architecture Patterns](#).

Solution overview for Anthos on VMware

This section provides an overview of the components that are involved in this solution from a physical and logical perspective.

Anthos on-prem is hybrid cloud software that brings Google Kubernetes Engine (GKE) to on-premises data centers. With Anthos clusters on VMware, we can create, manage, and upgrade Kubernetes clusters in your on-premises vSphere environment.

When we install Anthos clusters in vSphere environment, we create:

- An admin workstation
- An admin cluster
- One or more user cluster

The admin cluster runs the Anthos on-prem infrastructure, and the user clusters run our workloads.

Prerequisites

The following requirements are mandatory to deploy Anthos on PowerFlex rack:

- Fully configured and working PowerFlex rack system of RCM 3.5 or greater.
- Google Account, GCP Service Account, and a billing enabled GCP project. For detailed information, see [Configuring your Google Cloud project](#).
- vSphere 6.5 or 6.7 Update 3.
- At least one VMFS datastore with 2 TB capacity.
- Ability to create required DNS entries.
- A Layer 4 Network load balancer like F5 BIG-IP LTM, Citrix LB or Seesaw.
- Plan IP addresses in three networks as required, before deployment.
- Network access to Google Cloud APIs (*.googleapis.com).

For the latest supported version, see [Anthos GKE on-prem Compatibility Matrix](#).

Resource requirement for Anthos cluster

The initial installation of Anthos clusters on VMware requires these resources:

- 36 vCPU
- 98241 MB RAM
- 2280 GB virtual disk space

The details of CPU, RAM, and storage requirements for the admin workstation, admin cluster and user cluster are provided in the following tables:

Table 2. Admin workstation

Name	CPU	Memory	Storage
Admin Workstation	4	8192 MiB	50 GiB

Table 3. Admin cluster

Name	CPU	Memory	Storage	Purpose
Admin cluster control plane	4	16384 MiB	40 GiB	Runs the control plane for the admin cluster.
Add-on VMs	4	16384 MiB	40 GiB	Runs add-ons for the admin cluster.
User cluster control plane	4	8192 MiB	40 GiB	Each user cluster has its own control plane that runs on a VM called a control plane. The control plane VMs for a user cluster are nodes in the admin cluster. You can choose to create one or three control plane nodes for an individual user cluster.

Table 4. User cluster

Name	CPU	Memory	Storage	Purpose
User cluster worker nodes	4	8192 MiB	40 GiB	A user cluster node is a virtual machine where your workloads run. When you create a user cluster, you decide how many nodes it should have. The resources required for your user cluster nodes depend on the workloads you intend to run.

Scalability limits

The following limits can be affected by vSphere configuration and hardware.

- For each admin cluster, you can deploy a maximum of 20 user clusters.
- For each user cluster:
 - You can create a minimum of 3 nodes.
 - You can create a maximum of 250 nodes.

- You can create a maximum of 7500 pods:
 - Generally, each node includes an average of 30 pods. However, when Google Cloud's operations suite is enabled, each node includes an average of 20 pods.
 - Each pod can consist of one or more containers.
- For each node, you can create a maximum of 110 pods. This includes pods running add-on services.
- For bundled load balancing mode, you can create a maximum of 500 Load Balancer services. For more information, see [Scalability limits](#).

Physical design

This solution encompasses a fully functional PowerFlex rack system which is engineered for complete fault redundancy and ability to scale out across compute and storage dimensions. At a minimum, three controller nodes for running management stack and four PowerFlex hyperconverged nodes running VMware ESXi 6.5 or greater are required to deploy GKE on-prem. For more information about the VMware ESXi version, see [Google Compatibility Matrix](#).

Additionally, at least one outbound connection to Google Cloud APIs (*.googleapis.com) is required to complete the Anthos cluster registration and deployment processes.

The following figure describes how Anthos can manage multiple Anthos clusters across geographically dispersed PowerFlex systems through a single pane of glass with GCP console:

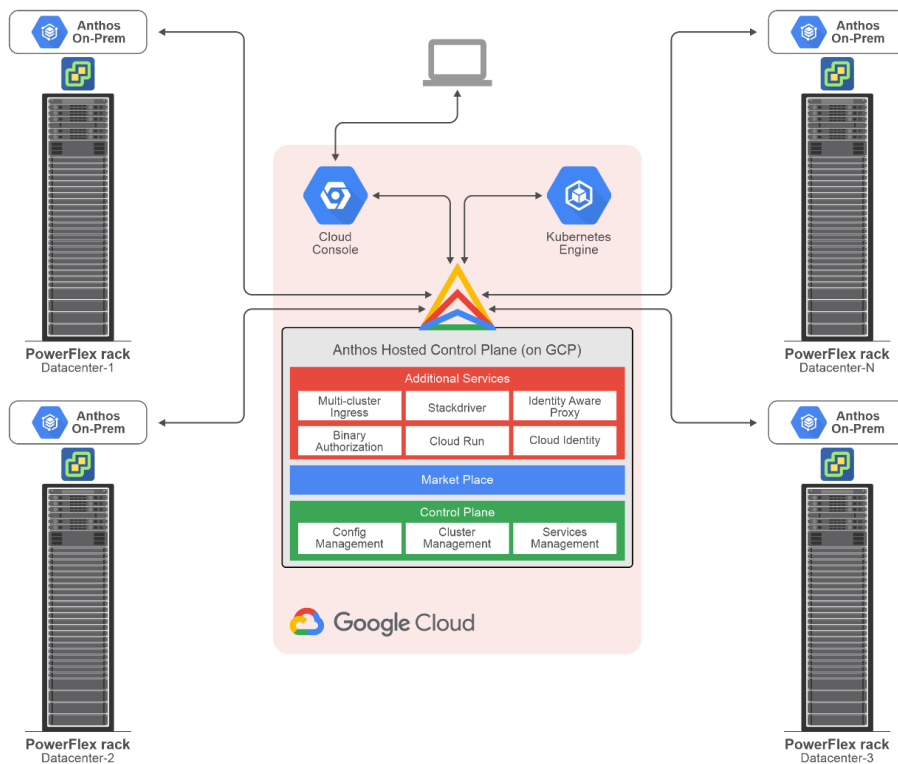


Figure 3. Physical design of Anthos with multiple PowerFlex rack systems geographically dispersed

Logical design

From a logical perspective, traffic flows into and out of the on-premises system using a load-balancer. Let us consider the example with F5 BIG-IP LTM. The F5 BIG-IP LTM creates dynamic connections between the compute nodes and the external network interfaces. Current best practice for PowerFlex rack is to create three special-purpose networks used for management, internal, and external traffic:

- Anthos on-prem admin network
- Anthos on-prem internal network
- Anthos on-prem external network

Create a new compute cluster with at least one compute server and one resource pool in the production workload vCenter to host the workloads in the production vCenter. This cluster requires VMware Dynamic Resource Scheduling (DRS) and one resource pool.

The following figure describes the logical configuration between Anthos clusters running in on-premises data centers on PowerFlex rack, Anthos clusters on VMware, and Anthos clusters on GCP:

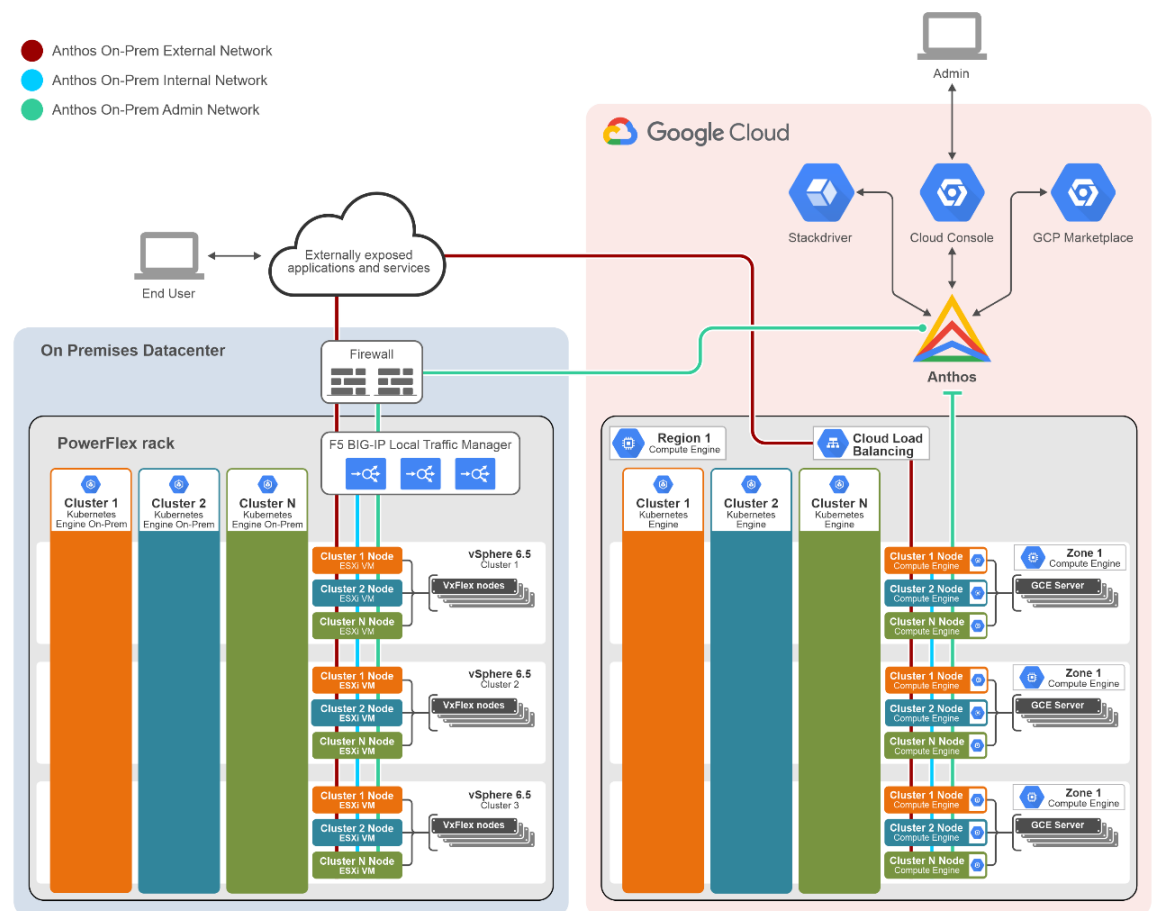


Figure 4. Logical design of Anthos with Anthos cluster deployed on PowerFlex rack

In this architecture, applications running on Anthos on-prem cluster can be exposed internally or externally to the web without traffic passing through GCP.

Note: The Anthos on-prem Admin Network connection from the on-premises data center to Anthos is outbound only.

The production vCenter server hosts multiple virtual machines in the new resource pool that consists of a virtual Anthos on-prem compute cluster. The application workloads are processes that run inside one of the Anthos on-prem compute cluster virtual machines. When an application gets deployed into the Anthos on-prem compute cluster, no additional vSphere virtual machines are created. The application runs inside the Anthos on-prem compute cluster virtual machines. If additional workload capacity is required, the Anthos on-prem compute cluster is expanded using the **gkectl** command-line utility or the Kubernetes Cluster API.

The summary of the correlation between vSphere VMs and Anthos on-prem cluster servers are as follows:

- Anthos on-prem cluster is a collection of vSphere VMs running as a collective instance.
- Anthos on-prem cluster's performance profile (for example: RAM, CPU, storage, etc.) is the additive sum of all vSphere VMs running in the GKE on-prem cluster.
- Multiple Anthos on-prem clusters can exist in a single vCenter deployment.

Deploying Anthos on VMware

Deploying the Anthos software requires knowledge about vCenter, Linux command-line skills, and an understanding of the existing network topology. The time to deploy the code varies based on network speed and familiarity with the tools that are discussed in this section but can generally be completed within a few hours.

Note: The instructions captured in this section are relevant to GKE on-prem 1.6.x. For latest information, see [Anthos GKE on-prem documentation](#).

Prerequisites

Before starting Anthos deployment, ensure that all requirements mentioned in the [Prerequisites](#) section are ready for deployment.

Prepare vCenter

Complete the following instructions to configure the vCenters of PowerFlex rack:

1. Log in to the PowerFlex management controller vCSA and Customer vCSA using the Administrator credentials.
2. Create three special-purpose networks used for management, internal, and external traffic on both the vCenters as defined in the [logical design](#) section.
3. Deploy and configure the load-balancer and deployment workstation on the PowerFlex management controller vCSA.
4. Create a resource pool for VMs deployed by Anthos on the customer vCSA. Configure the resources in the resource pool as per the environment where the Anthos is being deployed.

Deploy and configure Layer 4 network load balancer

Anthos requires a Layer 4 network load balancer. By default, Anthos integrates with F5 BIG-IP load balancer. You can also choose to enable manual load balancing and use your own L4 load balancer.

- To deploy F5 BIG-IP load balancer, see [Deployment of F5 BIG-IP](#).
- To deploy Seesaw load balancer, see [Deployment of Seesaw](#).

Deploy Anthos GKE on-prem admin VM on PowerFlex rack

Post physical environment setup, to create a deployment workstation on the PowerFlex management controller vCSA, complete the following steps on any Linux workstation:

1. Install Cloud SDK and related tools.
 - a. To install Cloud SDK, complete the steps mentioned in [Installing Google Cloud SDK](#).
 - b. Install `anthos-auth` and `kubectl` components by performing the following commands:

```
gcloud components install kubectl
gcloud components install anthos-auth
```

2. Configure Google Cloud Project.

- a. Log in to Google Cloud using your account credentials and set the default project ID:

```
gcloud auth login
gcloud config set project [PROJECT ID]
```

where [PROJECT_ID] is your project ID.

3. Enable the required APIs in the Google Cloud project.

```
gcloud services enable \
  anthos.googleapis.com \
  anthosgke.googleapis.com \
  anthosaudit.googleapis.com \
  cloudresourcemanager.googleapis.com \
  container.googleapis.com \
  gkeconnect.googleapis.com \
  gkehub.googleapis.com \
  serviceusage.googleapis.com \
  stackdriver.googleapis.com \
  monitoring.googleapis.com \
  logging.googleapis.com
```

4. Grant IAM roles - `resourcemanager.projectIamAdmin` and `serviceusage.serviceUsageAdmin`, to your Google account.

```
gcloud projects add-iam-policy-binding PROJECT_ID \
  --member="user:ACCOUNT" \
  --role="roles/resourcemanager.projectIamAdmin"

gcloud projects add-iam-policy-binding PROJECT_ID \
  --member="user:ACCOUNT" \
  --role="roles/serviceusage.serviceUsageAdmin"
```

where ACCOUNT is your Google Cloud account and [PROJECT_ID] is the associated project ID.

5. Create a component access service account which is used to download Anthos components from Container Registry.

```
gcloud iam service-accounts create component-access-sa \
  --display-name "Component Access Service Account" \
  --project PROJECT_ID
```

6. Create JSON key for your component access service account.

```
gcloud iam service-accounts keys create component-access-
key.json \
  --iam-account component-access-
sa@[PROJECT_ID].iam.gserviceaccount.com
```

7. Grant the below IAM roles to your component access service account so that Anthos clusters on VMware can do preflight checks:

- serviceusage.serviceUsageViewer
- iam.serviceAccountCreator
- iam.roleViewer

```
gcloud projects add-iam-policy-binding PROJECT_ID \
  --member "serviceAccount:component-access-
sa@[PROJECT_ID].iam.gserviceaccount.com" \
  --role "roles/serviceusage.serviceUsageViewer"

gcloud projects add-iam-policy-binding PROJECT_ID \
  --member "serviceAccount:component-access-
sa@[PROJECT_ID].iam.gserviceaccount.com" \
  --role "roles/iam.serviceAccountCreator"

gcloud projects add-iam-policy-binding PROJECT_ID \
  --member "serviceAccount:component-access-
sa@[PROJECT_ID].iam.gserviceaccount.com" \
  --role "roles/iam.roleViewer"
```

8. Download **gkeadm** utility and make it executable.

```
gsutil cp gs://gke-on-prem-release/gkeadm/1.6.0-
gke.7/linux/gkeadm ./
chmod +x gkeadm
```

9. Generate configuration file templates using the **gkeadm** tool.

```
./gkeadm create config
```

The preceding command creates the below files in your current directory:

- **credential.yaml**: Populate vCenter credentials. For detailed information, see [Fill in credential.yaml](#).
- **admin-ws-config.yaml**: Populate admin workstation configurations. For detailed information, see [Admin workstation configuration file](#).

10. Create your admin workstation and a set of service account automatically.

```
./gkeadm create admin-workstation --auto-create-service-
accounts
```

11. On completion of the admin workstation, a command displayed on the console which can be used to create SSH connection to the admin workstation.

```
ssh -i /usr/local/google/home/me/.ssh/gke-admin-workstation
ubuntu@<admin_workstation_ip>
```

If you list the files on the admin workstation, you can see two cluster configuration files, your CA certificate file, and the JSON key files for your service accounts:

Example output from “**ls -l**” command:

```
admin-cluster.yaml
connect-agent-sa-2007081316.json
connect-register-sa-2007081316.json
log-mon-sa-2007081316.json
user-cluster.yaml
vcenter-ca-cert.pem
component-access-key.json
```

Create Anthos GKE on-prem admin and user cluster on PowerFlex rack

To create the admin and user cluster, complete the following steps:

1. SSH to the admin workstation from the deployment workstation.
2. For the admin cluster, modify the **admin-cluster.yaml** file to match your environment-specific variable. For detailed information, see [Admin cluster configuration file](#).

3. Validate the admin cluster configuration file using **gkectl check-config** command.

```
gkectl check-config --config admin-cluster.yaml
```

4. Initialize your vSphere environment using **gkectl prepare** command.

```
gkectl prepare --config admin-cluster.yaml --skip-validation-all
```

5. Create the admin cluster using the **gkectl create admin** command.

```
gkectl create admin --config admin-cluster.yaml
```

6. For the user cluster, modify the **user-cluster.yaml** file to match your environment-specific variable. For detailed information, see [User cluster configuration file](#).

7. Validate the user cluster configuration file using **gkectl check-config** command.

```
gkectl check-config --kubeconfig [ADMIN_CLUSTER_KUBECONFIG]
--config user-cluster.yaml
```

where, [ADMIN_CLUSTER_KUBECONFIG] is the path of the kubeconfig file for your admin cluster.

8. Create the user cluster using the **gkectl** create cluster command.

```
gkectl create cluster --kubeconfig  
[ADMIN_CLUSTER_KUBECONFIG] --config user-cluster.yaml
```

9. Verify the successful creation of admin cluster.

```
kubect1 get nodes --kubeconfig [ADMIN_CLUSTER_KUBECONFIG]
```

10. Verify the successful creation of user cluster.

```
kubect1 get nodes --kubeconfig [USER_CLUSTER_KUBECONFIG]
```

Results: The user cluster would be registered with Google Cloud once the cluster is created successfully.

11. Log in to **Google cloud > Kubernetes Engine > Clusters** to verify the results.

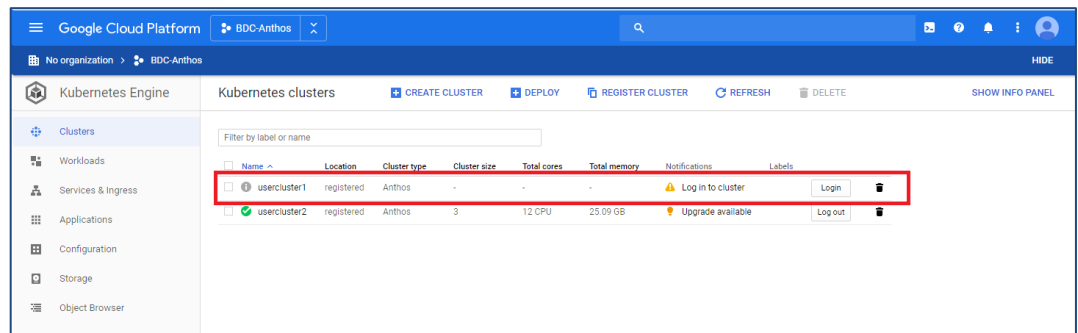


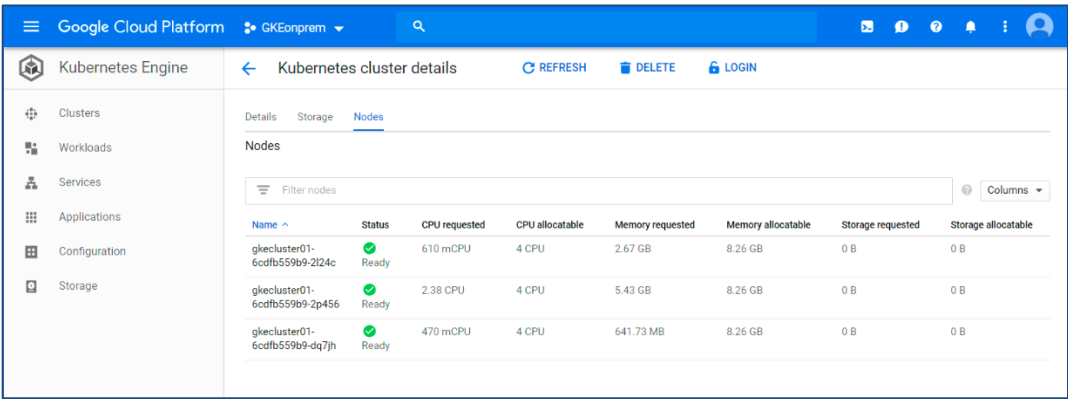
Figure 5. User cluster registered with GCP after successful deployment

12. Select a **cluster** and click **Login**. You are prompted to choose a method you want to use for authentication of the cluster.
- Token
 - Basic authentication
 - Authenticate with Identity Provider configured for the cluster

Select the method of your choice and login to view the cluster overview.

For detailed information, see [Logging in to a cluster from Cloud Console](#).

13. For complete detail of the cluster, select **Kubernetes cluster details > Nodes**.



Name	Status	CPU requested	CPU allocatable	Memory requested	Memory allocatable	Storage requested	Storage allocatable
gkecluster01-6cdfb559b9-2l24c	Ready	610 mCPU	4 CPU	2.67 GB	8.26 GB	0 B	0 B
gkecluster01-6cdfb559b9-2p456	Ready	2.38 CPU	4 CPU	5.43 GB	8.26 GB	0 B	0 B
gkecluster01-6cdfb559b9-dq7jh	Ready	470 mCPU	4 CPU	641.73 MB	8.26 GB	0 B	0 B

Figure 6. Kubernetes cluster node details in GCP

Note: The Kubernetes cluster nodes for Anthos on-prem clusters are running as vSphere VMs on PowerFlex system.

14. **(Optional)** Check if the Kubernetes cluster nodes in the GCP console has the same name as the worker node VMs running on PowerFlex rack.




Figure 7. Kubernetes cluster nodes in VMware vCenter

For detailed information, see [Creating an admin cluster](#) and [Creating a user clusters](#).

Solution Overview for Anthos on bare metal

Anthos on bare metal allows us to install Google Kubernetes Engine (GKE) clusters to on-premises data centers without the need for a vSphere environment.

Anthos on bare metal works reliably and efficiently in multiple computing environments and without the need for a virtual machine. For customers looking for increased compute efficiency, this option provides greater available CPU resources and may be a preferred cost optimization.

We can define four types of clusters with Anthos on bare metal:

- **admin** — A cluster used to manage user clusters.
- **user** — A cluster used to run workloads.
- **hybrid** — A single cluster for both admin and workloads, that can also manage user clusters.
- **standalone** — A single cluster that can administer itself, and that can also run workloads, but can't create or manage other user clusters.

Prerequisites

The following requirements are mandatory to create a cluster using Anthos on bare metal:

- PowerFlex solution – PowerFlex Rack / Appliance running PowerFlex (previously VxFlex OS) 3.5.1 or later configured in two-layer architecture.
- Google Account, GCP Service Account, and a billing enabled GCP project.
- A Linux administration machine with `gcloud`, `gsutil`, `kubectl`, `docker` and `bmctl` utilities installed.
- Network access to Google Cloud APIs (`servicemanagement.googleapis.com` and `servicecontrol.googleapis.com`).
- Operating System: RHEL 8.1/8.2, CentOS 8.1/8.2, Ubuntu 18.04/20.04 LTS
- A Layer 4 Network load balancer.

Table 5. Minimum and recommended hardware requirements for running Anthos on bare metal

Resource	Minimum	Recommended
CPUs / vCPUs	4 cores	8 cores
RAM	32 GiB	64 GiB
Storage	128 GiB	256 GiB

Logical design

Anthos on bare metal only requires L3 connectivity for worker nodes, but requires L2 connectivity among control plane nodes to announce VIPs. The load balancer can be run on a separate node pool or on the control plane node pool. The same L2 network requirements apply if the load balancer is run on the control plane node pool

Control plane traffic all goes through the control plane VIP, so load balancer setup is required during cluster startup.

The requirements for load balancer machines are:

- All load balancers are in the same L2 subnet.
- All VIPs must be in the load balancer machine subnet and routable to the gateway of the subnet.

The following figure describes the logical configuration of an Anthos on bare metal environment with a non-HA control plane:

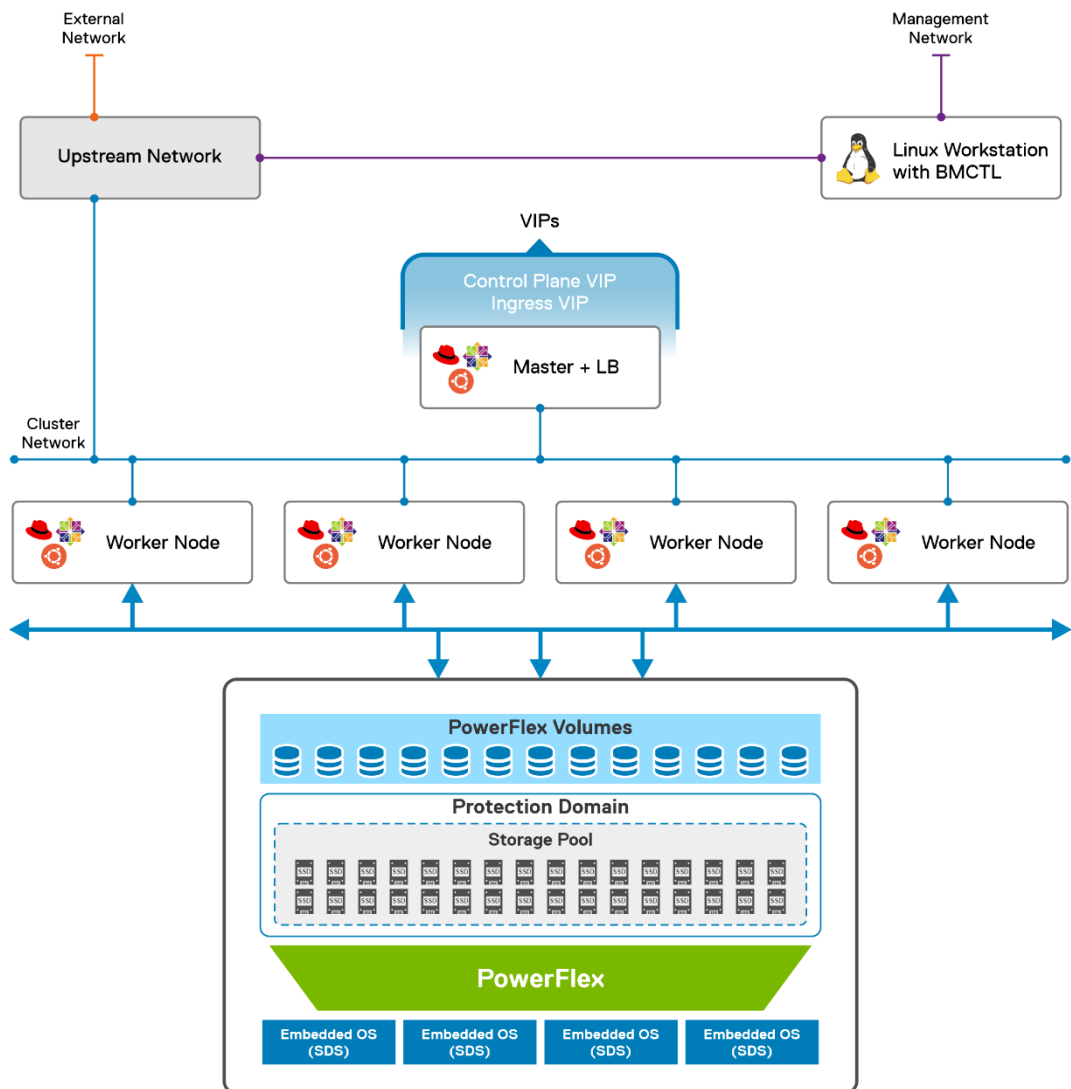


Figure 8. Bare metal environment with non-HA control plane

- The control plane nodes run load balancers, and they are all on the same L2 network, while other connections, including worker nodes, only require L3 connectivity.
- Configuration files define IP addresses for worker node pools, virtual IP addresses for services, ingress and control plane (Kubernetes API) access.
- A connection to Google Cloud is also required.

The following figure describes the logical configuration of bare metal environment with a HA control plane:

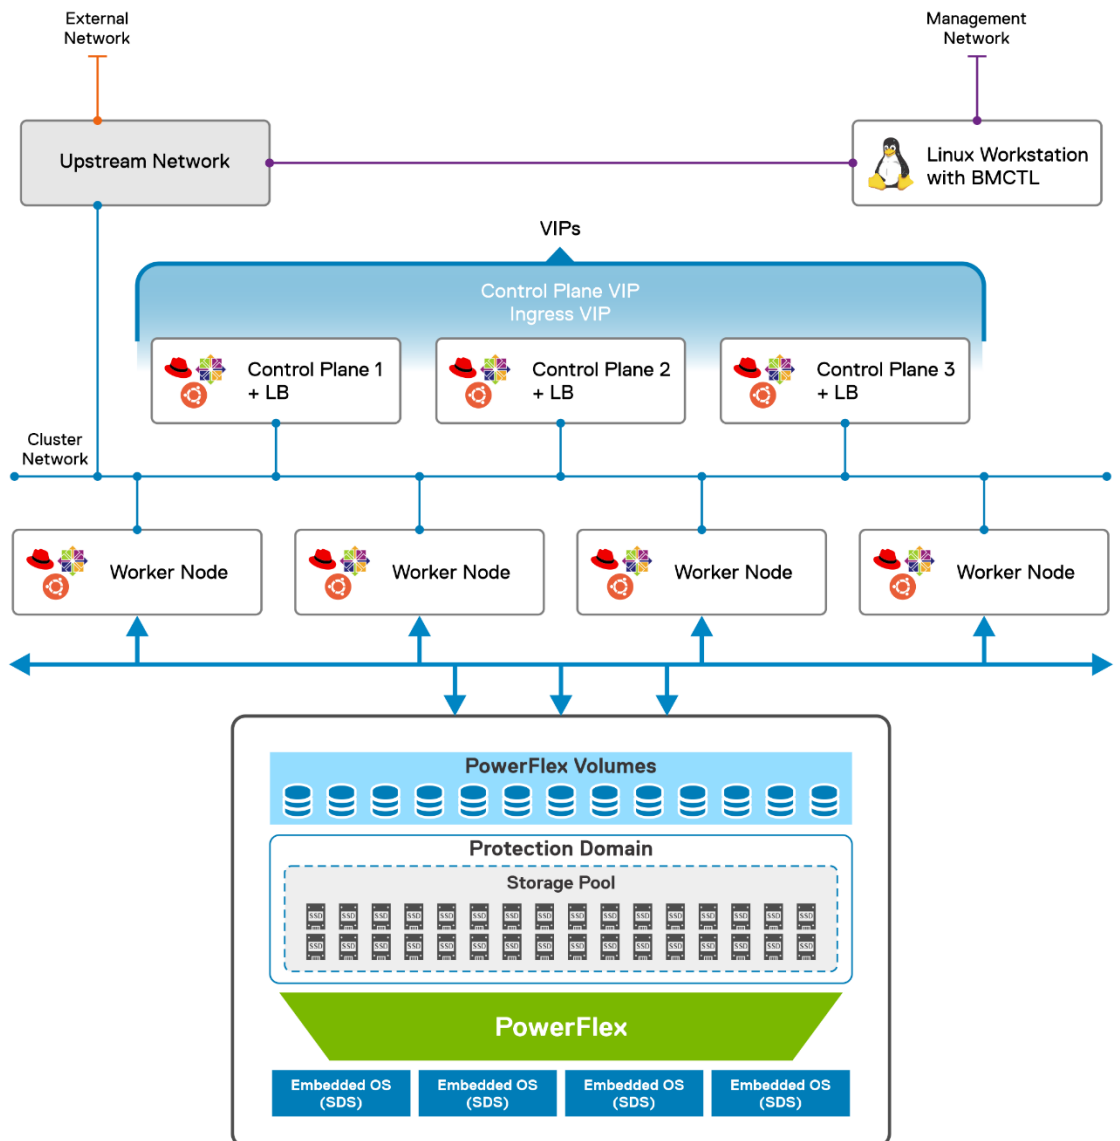


Figure 9. Bare metal environment with HA control plane

Deploying Anthos on bare metal

Deploying Anthos on bare metal requires knowledge with Linux command-line skills, and an understanding of the existing network topology. The time to deploy the code varies based on network speed and familiarity with the tools that are discussed in this section but can generally be completed within a few hours.

Prerequisites

Before starting Anthos deployment, ensure that all prerequisites mentioned in [Solution Overview](#) section are ready for deployment.

Prepare Linux administration machine

Complete the following instructions to configure the Linux administration machine:

1. Install `gcloud` and `gsutil` tools which are included in the Cloud SDK. For detailed information, see [Installing Google Cloud SDK](#).

2. Update the Cloud SDK.

```
gcloud components update
```

3. Log in with your Google account so that you can manage services and service accounts.

```
gcloud auth login
```

4. Set the default project. Replace `PROJECT_ID` with your Google Cloud project ID.

```
gcloud config set project PROJECT_ID
```

5. Use `gcloud` to install `kubectl`.

```
gcloud components install kubectl
```

6. Configure service accounts like the Connect-agent service account, Connect-register service account and Logging-monitoring service account to connect the cluster to Google cloud, manage, and monitor the cluster from Cloud Console. For detailed information, see [Enabling Google services and service accounts](#).

7. Enable the required Google services (`servicemanagement.googleapis.com` and `servicecontrol.googleapis.com`) in your Cloud project.

```
gcloud services enable servicemanagement.googleapis.com
servicecontrol.googleapis.com
```

8. Create a service account to access Cloud Storage bucket.

```
gcloud iam service-accounts create baremetal-gcr
```

9. Download the service account key as a JSON file.

```
gcloud iam service-accounts keys create gcr.json \
  --iam-account=baremetal-
gcr@PROJECT_ID.iam.gserviceaccount.com \
  --project=PROJECT_ID
```

10. Create a directory and copy the JSON service account file to it.

```
cd ~
mkdir baremetal
cd baremetal
gcloud auth activate-service-account --key-file=gcr.json
```

11. Download `bmctl` and make it executable.

bmctl, is a utility to streamline the creation of Anthos clusters on bare metal. This utility automates the setting up of Google service accounts and API's which are required to setup the Anthos clusters on bare metal.

```
gsutil cp gs://anthos-baremetal-release/bmctl/0.5.0-
gke.2/linux/bmctl .
chmod a+x bmctl
```

12. Switch the `gcloud` active account back to your personal account.

```
gcloud config set account EMAIL_ADDRESS
```

where `EMAIL_ADDRESS` is the email address of the account `gcloud` should use for authentication.

13. Verify if `bmctl` is installed correctly by viewing the help information.

```
./bmctl -h
```

Deploy and Configure Layer 4 network load balancer

Anthos on bare metal requires a Layer 4 network load balancer. It supports two load balancer options: bundled and manual.

In bundled load balancing, Anthos on bare metal deploys Layer 4 load balancers at the time of cluster installation. External load balancer is not required in this case. The load balancers can run on a dedicated pool of worker nodes, or they can be located on the same nodes as the control plane. Since the load balancers broadcast ARP messages to announce VIPs, all the nodes must be on the same Layer 2 subnet.

For detailed information, see [Configuring bundled load balancing](#).

In manual load balancing, Anthos on bare metal does not deploy load balancers for clusters. We must configure an external load balancer before installing the cluster.

Before the cluster installation, we must configure our clusters control plane VIP on an external load balancer. After installation, we must configure the load balancer for the Kubernetes services and ingresses in the cluster.

For detailed information, see [Configuring manual load balancing](#).

Creating cluster with Anthos on bare metal

As mentioned in [Solution Overview](#) section, we can define four types of cluster with Anthos on bare metal - admin cluster, user cluster, hybrid cluster, and stand-alone cluster.

This section provides instructions to create a hybrid cluster.

Prerequisites

1. Download and install `bmctl` tool on the Linux administration machine.
2. The Linux administration machine running `bmctl` tool should have network connectivity to all the nodes of the target hybrid cluster as well as the control plane VIP.
3. Root user SSH keys or SUDO user access on all nodes in the target hybrid cluster.

Deployment

To create the hybrid cluster, complete the following steps:

1. SSH to the Linux administration machine and change to the bare metal directory created in [Prepare Linux administration machine](#).

2. Log in to `gcloud` as a user.

```
gcloud auth application-default login
```

3. Create the cluster configuration file.

```
bmctl create config -c <HYBRID_CLUSTER_NAME>
```

4. Modify the configuration file parameters to match your environment-specific variable. For detailed information, see [Edit the cluster config file](#).

5. Create the cluster.

```
./bmctl create cluster -c CLUSTER_NAME
```

The `bmctl` command runs preflight checks on the cluster config file before it creates a cluster. If they are successful, the cluster is created. The `bmctl` logs, as well as the preflight check and node installation logs, are in the directory: `baremetal/bmctl-workspace/<CLUSTER_NAME>/log`.

6. During cluster creation, the `bmctl` command writes a kubeconfig file which is written to `bmctl-workspace/<CLUSTER_NAME>/<CLUSTER_NAME>-kubeconfig`. Verify the successful creation of cluster.

```
kubectl --kubeconfig bmctl-workspace/<CLUSTER_NAME>/<CLUSTER_NAME>-kubeconfig get nodes
```

Deploying Anthos applications on PowerFlex system

After the Anthos cluster has been deployed and registered, you can easily deploy workloads using the Google Cloud Console tool.

Note: Deploying applications to the Anthos on-prem cluster running on PowerFlex system is no different than deploying applications to GKE running on GCP. Anthos GKE on-prem clusters can be selected as the workload destination cluster in the same way GCP data center would be selected. End users use a single, simplified interface to deploy applications without requiring in-depth knowledge of the underlying infrastructure.

Sample workload: NGINX web server deployment

In this example, we are deploying NGINX - a free, open-source HTTP server and reverse proxy that is well known for its high performance, stability, rich feature set, simple configuration, and low resource consumption. Since it is already in the Google market place, deploying NGINX requires just a few mouse clicks.

For complete information about NGINX, see [NGINX Wiki](#).

Deploying NGINX

1. Log in to **Google Cloud Platform**.
2. Select **Kubernetes Engine > Workloads > Deploy**.

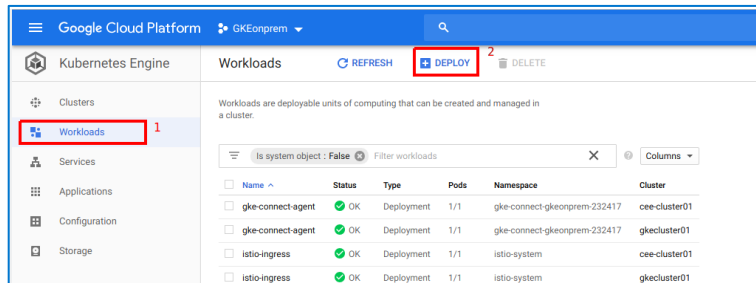


Figure 10. Workload deployment wizard

3. Enter appropriate container and configuration details in **Create a Deployment** window.

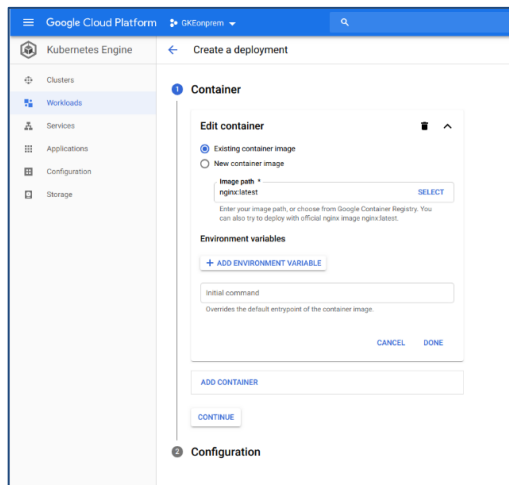


Figure 11. Container configuration

Providing examples to enter details:

- Container image — `nginx:latest`
 - Application Name — `nginx-test`
 - Namespace — `nginx-deployment`
 - Destination Cluster — `gke-on-prem-US`
 - Pods — Enter minimum pod number to start.
4. Click **Deploy** to start the deployment.
 5. Verify if the deployment has been completed, select **Kubernetes Engine > Workloads**.

Name	Status	Type	Pods	Namespace	Cluster
gke-connect-agent	OK	Deployment	1/1	gke-connect-gkeonprem-232417	gkecluster01
istio-ingress	OK	Deployment	1/1	istio-system	gkecluster01
istio-pilot	OK	Deployment	1/1	istio-system	gkecluster01
nginx-1	OK	Deployment	3/3	default	gkecluster01
nginx-1-deployer	OK	Job	0/0	demo-nginx	gkecluster01
nginx-1-nginx	OK	Stateful Set	5/5	demo-nginx	gkecluster01

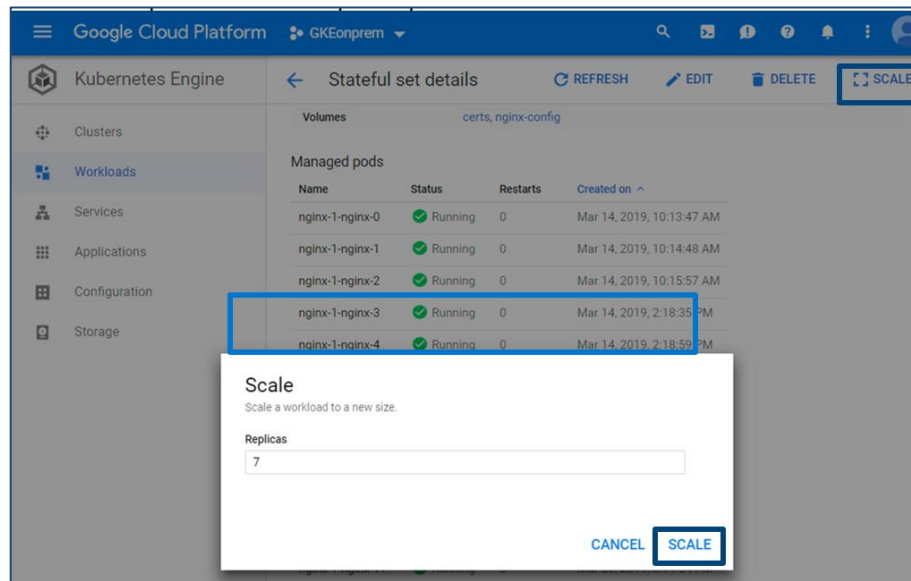
Figure 12. Successful deployment of Nginx workload

Note: Since the NGINX pods run inside the GKE on-prem cluster virtual machines, you will not see any additional virtual machines in vCenter.

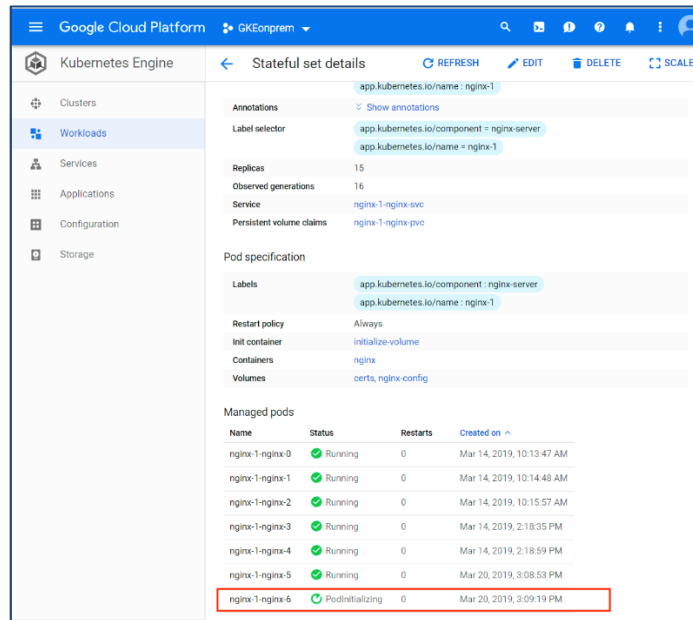
Scaling NGINX deployment

Once deployed, scaling the number of pods in the deployment to meet the demands of your application is accomplished easily with a few clicks.

1. Log in to **Google Cloud Platform**.
2. Select **Kubernetes Engine > Workloads > SCALE**.
3. Enter the required number of pod replicas and click **SCALE**.



4. Refresh console window to view new pods.



Note: Refresh immediately else new pods passes the **PodInitializing** status and shows status as **Running**.

Conclusion

Engineered together with the Google Cloud team, this hybrid cloud solution provides you with best in class software-defined platform from Dell EMC and managed Kubernetes by Google Cloud, helping simplify hybrid cloud workload management while enabling organizations to respond to changing business needs rapidly.

PowerFlex offers enterprises extreme flexibility, massive scalability, predictable performance, and simple management for their on-premises private cloud infrastructure. With Google Cloud's Anthos, enterprises can complement their geographically dispersed on-premises PowerFlex deployments and Google Cloud deployments with a simple hybrid cloud management framework for their next-generation cloud-native applications.

Appendix

Deploy and Configure the Seesaw load balancer

The bundled load balancer that Anthos clusters on VMware provides is the Seesaw load balancer. Seesaw load balancer is available for Anthos bare metal as well.

With bundled load balancing, Anthos clusters on VMware provides and manages the load balancer. You do not have to get a license for a load balancer, and minimal effort is required to setup this load balancer.

To configure Seesaw load balancer, in the configuration file of admin cluster and user cluster, set `loadBalancer.kind` to 'Seesaw' and update the load balancer section to match your environment.

For detailed information, see [Bundled load balancing with Seesaw](#).

Deploy and configure F5 Big-IP Virtual Edition appliance

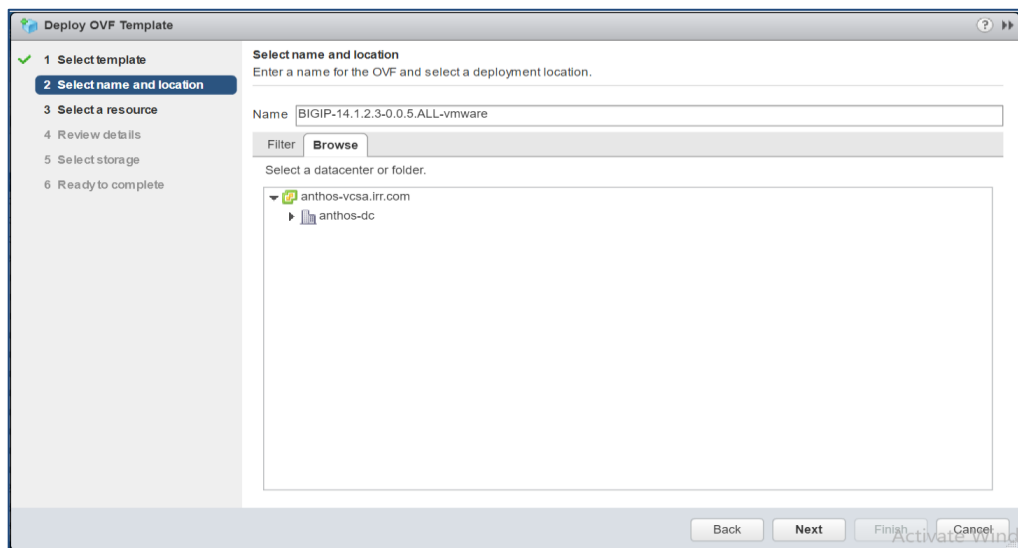
Anthos clusters on VMware require a Layer 4 network load balancer and F5 BIG-IP load balancer is an option.

Prerequisite

Register, login, and download the Big-IP Virtual Edition OVA from BIG-IQ Virtual Editions [official site](#). Use this file to deploy OVF template.

Procedure

1. Log in to the PowerFlex Management Controller vCSA.
2. Right-click the cluster and select **Deploy OVF Template**. The Deploy OVF Template wizard appears.



3. On the **Select template** page, select the **Big-IP Virtual Edition OVA** file and click **Next**.

4. Ensure to complete update all mandatory fields in the following pages:
 - Select name and location
 - Select a resource
 - Review details
 - Accept license agreements
 - Select configuration
 - Select storage: Select vsan-datastore.
 - Select network: Select the appropriate Networks for the Internal, External, and Management networks as defined in the [Logical design](#) section.
 - Ready to complete
5. Click **Finish** to start the deployment. For more details, see [Deploy an OVF or OVA Template](#).
6. Post-deployment of the virtual appliance, power the appliance up from the vCenter.
 - If the management network is configured with DHCP, then it should receive a DHCP address.
 - If the management is configured for static address then you must configure the Management Interface using VM console. See, [Deploying BIG-IP Virtual Edition in ESXi](#).

Result: You must successfully set up Management Interface.

7. Log in to the management interface of the load balancer using a browser.
8. Go to [Try F5](#) official site to get F5 trial license.
9. Configure the self-IPs, one for external interface and the other for internal interface.
10. Create Administrative and User partition.

For more details, see [Installing F5 BIG-IP ADC for Anthos GKE on-prem](#).