



Anthos with NetApp

NetApp container solutions

NetApp
August 18, 2025

This PDF was generated from <https://docs.netapp.com/us-en/netapp-solutions-containers/anthos/anthos-solution-overview.html> on August 18, 2025. Always check docs.netapp.com for the latest.

Table of Contents

- Anthos with NetApp 1
 - NVA-1165: Anthos with NetApp 1
 - Use cases 1
 - Business value 1
 - Technology overview 1
 - Advanced configuration options 2
 - Current support matrix for validated releases 2
 - Learn about Anthos 2
 - Anthos Overview 2
 - Anthos Clusters on VMware 4
 - Anthos on bare metal 8
 - NetApp storage systems 12
 - NetApp Storage Overview 12
 - NetApp ONTAP 12
 - NetApp storage integrations 14
 - NetApp Storage Integration Overview 14
 - Trident Overview 16
 - Advanced configuration options 18
 - Advanced configuration options 18
 - Explore load balancer options 18
 - Solution validation and use cases 28
 - Deploy an application from the Google Cloud Console Marketplace 28
 - Where to find additional information 39

Anthos with NetApp

NVA-1165: Anthos with NetApp

Banu Sundhar and Suresh Thoppay, NetApp

This reference document provides deployment validation of the Anthos with NetApp solution by NetApp and our engineering partners when it is deployed in multiple data-center environments. It also details storage integration with NetApp storage systems by using the Trident storage orchestrator for the management of persistent storage. Lastly, we explore and document a number of solution validations and real-world use cases.

Use cases

The Anthos with NetApp solution is architected to deliver exceptional value for customers with the following use cases:

- Easy to deploy and manage Anthos environment deployed using the provided `bmctl` tool on bare metal or the `gkectl` tool on VMware vSphere.
- Combined power of enterprise container and virtualized workloads with Anthos deployed virtually on vSphere or on bare metal with [kubevirt](#).
- Real-world configuration and use cases highlighting Anthos features when used with NetApp storage and Trident, the open-source storage orchestrator for Kubernetes.

Business value

Enterprises are increasingly adopting DevOps practices to create new products, shorten release cycles, and rapidly add new features. Because of their innate agile nature, containers and microservices play a crucial role in supporting DevOps practices. However, practicing DevOps at a production scale in an enterprise environment presents its own challenges and imposes certain requirements on the underlying infrastructure, such as the following:

- High availability at all layers in the stack
- Ease of deployment procedures
- Non-disruptive operations and upgrades
- API-driven and programmable infrastructure to keep up with microservices agility
- Multitenancy with performance guarantees
- The ability to run virtualized and containerized workloads simultaneously
- The ability to scale infrastructure independently based on workload demands

The Anthos with NetApp solution acknowledges these challenges and presents a solution that helps address each concern by implementing the fully automated deployment of Anthos on prem in the customer's data center environment of choice.

Technology overview

The Anthos with NetApp solution is comprised of the following major components:

Anthos On Prem

Anthos On Prem is a fully supported enterprise Kubernetes platform that can be deployed in the VMware vSphere hypervisor, or on a bare metal infrastructure of your choosing.

For more information about Anthos, see the Anthos website located [here](#).

NetApp storage systems

NetApp has several storage systems perfect for enterprise data centers and hybrid cloud deployments. The NetApp portfolio includes NetApp ONTAP, Cloud Volumes ONTAP, Google Cloud NetApp Volumes, Azure NetApp Files, FSx ONTAP for NetApp ONTAP storage systems, all of which can provide persistent storage for containerized applications.

For more information visit the NetApp website [here](#).

NetApp storage integrations

Trident is an open-source and fully-supported storage orchestrator for containers and Kubernetes distributions, including Anthos.

For more information, visit the Trident website [here](#).

Advanced configuration options

This section is dedicated to customizations that real world users would likely need to perform when deploying this solution into production, such as creating a dedicated private image registry or deploying custom load balancer instances.

Current support matrix for validated releases

See [here](#) for the support matrix for validated releases.

Learn about Anthos

Anthos Overview

Anthos with NetApp is a verified, best-practice hybrid cloud architecture for the deployment of an on-premises Google Kubernetes Engine (GKE) environment in a reliable and dependable manner. This NetApp Verified Architecture reference document serves as both a design guide and a deployment validation of the Anthos with NetApp solution deployed to bare metal and virtual environments. The architecture described in this document has been validated by subject matter experts at NetApp and Google Cloud to provide the advantages of running Anthos within your enterprise data-center environment.

Anthos

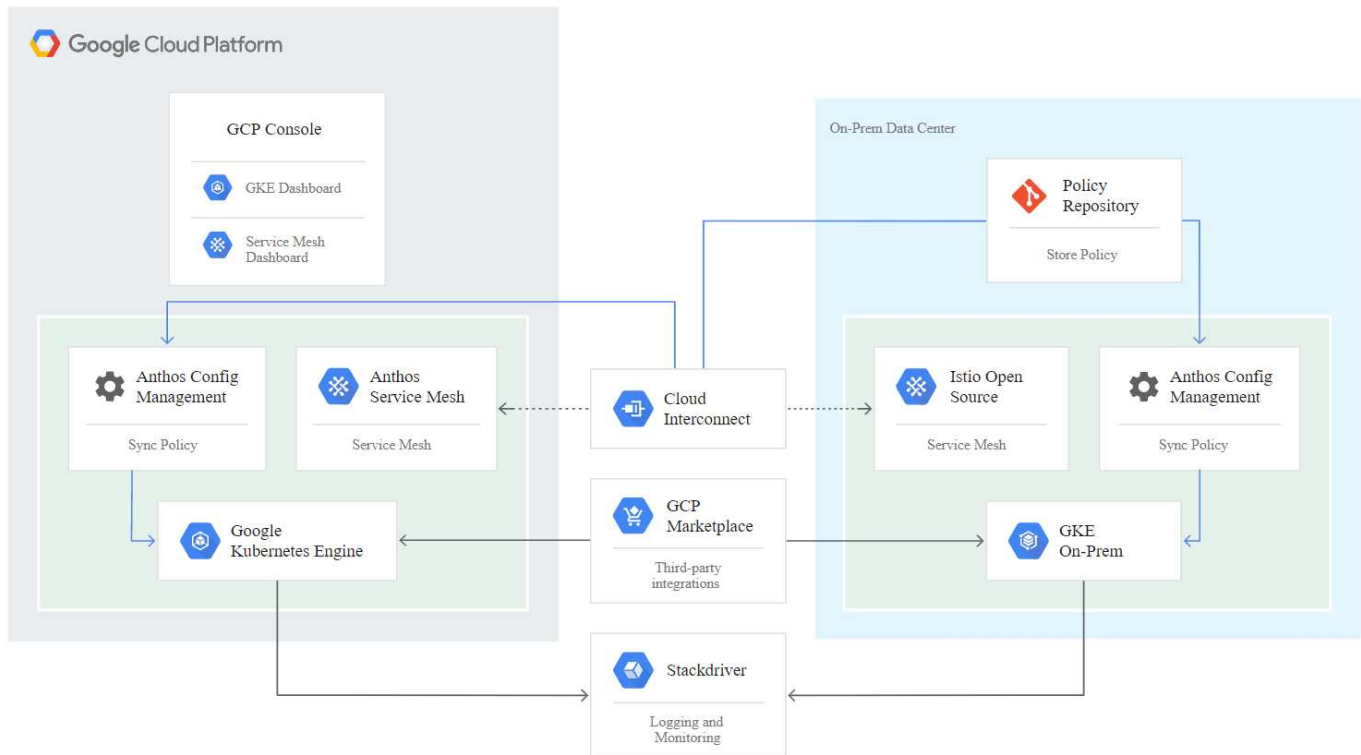
Anthos is a hybrid-cloud Kubernetes data center solution that enables organizations to construct and manage modern hybrid-cloud infrastructures while adopting agile workflows focused on application development. Anthos on VMware, a solution built on open-source technologies, runs on-premises in a VMware vSphere-

based infrastructure, which can connect and interoperate with Anthos GKE in Google Cloud. Adopting containers, service mesh, and other transformational technologies enables organizations to experience consistent application development cycles and production-ready workloads in local and cloud-based environments. The following figure depicts the Anthos solution and how a deployment in an on-premises data center interconnects with infrastructure in the cloud.

For more information about Anthos, see the Anthos website located [here](#).

Anthos provides the following features:

- **Anthos configuration management.** Automates the policy and security of hybrid Kubernetes deployments.
- **Anthos Service Mesh.** Enhances application observability, security, and control with an Istio-powered service mesh.
- **Google Cloud Marketplace for Kubernetes Applications.** A catalog of curated container applications available for easy deployment.
- **Migrate for Anthos.** Automatic migration of physical services and VMs from on-premises to the cloud.
- **Stackdriver.** Management service offered by Google for logging and monitoring cloud instances.



Deployment methods for Anthos

Anthos clusters on VMware

Anthos clusters deployed to VMware vSphere environments are easy to deploy, maintain, and scale rapidly for most end-user Kubernetes workloads.

For more information about Anthos clusters on VMware, deployed with NetApp, please visit the page [here](#).

Anthos on bare metal

Anthos clusters deployed on bare metal servers are hardware agnostic and allow you to select a compute platform optimized for your personalized use case.

For more information about Anthos on bare metal clusters deployed with NetApp, visit [here](#).

Anthos Clusters on VMware

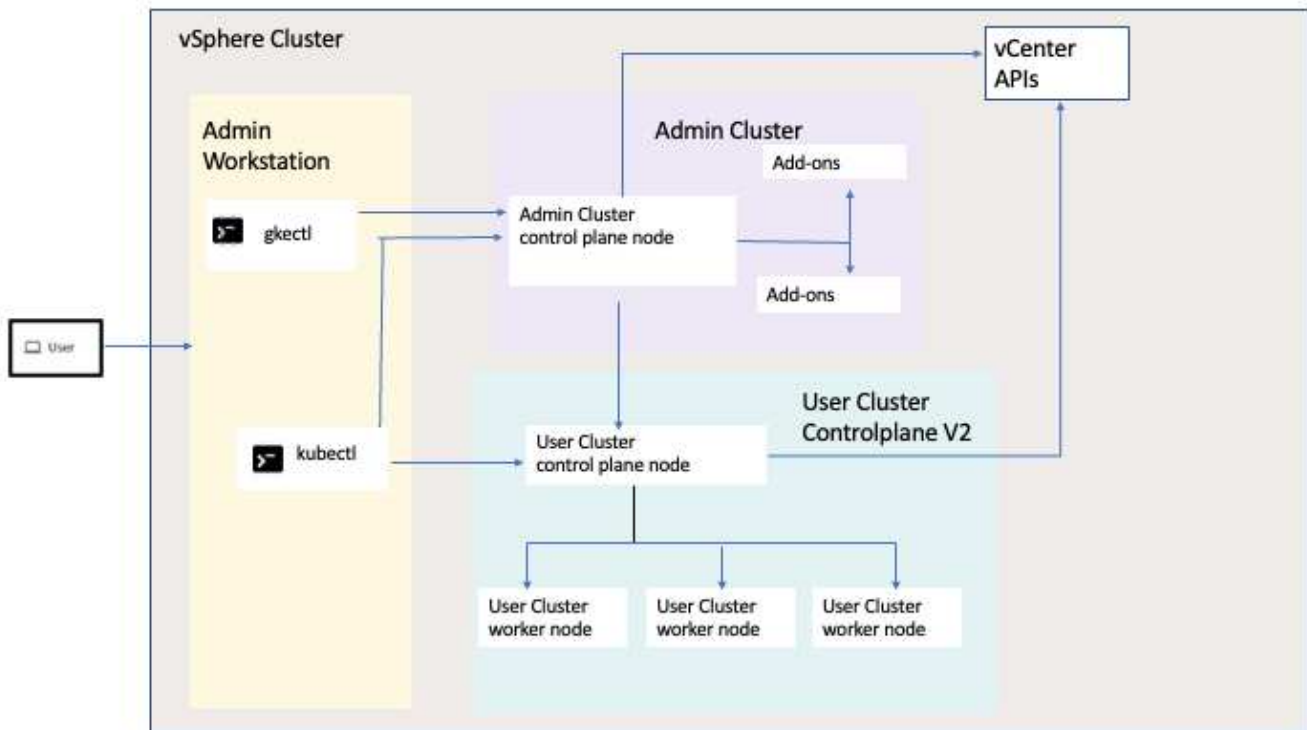
Anthos clusters on VMware is an extension of Google Kubernetes Engine that is deployed in an end user's private data center. An organization can deploy the same applications designed to run in containers in Google Cloud in Kubernetes clusters on-premises.

Anthos clusters on VMware can be deployed into an existing VMware vSphere environment in your data center, which can save on capital expenses and enable more rapid deployment and scaling operations.

The deployment of Anthos clusters on VMware includes the following components:

- **Anthos admin workstation.** A deployment host from which `gkectl` and `kubect1` commands can be run to deploy and interact with Anthos deployments.
- **Admin cluster.** The initial cluster deployed when setting up Anthos clusters on VMware. This cluster manages all subordinate user cluster actions, including deployment, scaling, and upgrade.
- **User cluster.** Each user cluster is deployed with it's own load balancer instance or partition, allowing it to act as a standalone Kubernetes cluster for individual users or groups, helping to achieve full multitenancy.

The following graphic is a description of an Anthos-clusters-on-VMware deployment.



Benefits

Anthos clusters on VMware offers the following benefits:

- **Advanced multitenancy.** Each end user can be assigned their own user cluster, deployed with the virtual resources necessary for their own development environment.
- **Cost savings.** End users can realize significant cost savings by deploying multiple user clusters to the same physical environment and utilizing their own physical resources for their application deployments instead of provisioning resources in their Google Cloud environment or on large bare-metal clusters.
- **Develop then publish.** On-premises deployments can be used while applications are in development, which allows for testing of applications in the privacy of a local data center before being made publicly available in the cloud.
- **Security requirements.** Customers with increased security concerns or sensitive data sets that cannot be stored in the public cloud are able to run their applications from the security of their own data centers, thereby meeting organizational requirements.

VMware vSphere

VMware vSphere is a virtualization platform for centrally managing a large number of virtualized servers and networks running on the ESXi hypervisor.

For more information about VMware vSphere, see the [VMware vSphere website](#).

VMware vSphere provides the following features:

- **VMware vCenter Server.** VMware vCenter Server provides unified management of all hosts and VMs from a single console and aggregates performance monitoring of clusters, hosts, and VMs.
- **VMware vSphere vMotion.** VMware vCenter allows you to hot migrate VMs between nodes in the cluster upon request in a non-disruptive manner.
- **vSphere High Availability.** To avoid disruption in the event of host failures, VMware vSphere allows hosts to be clustered and configured for high availability. VMs that are disrupted by host failure are rebooted shortly on other hosts in the cluster, restoring services.
- **Distributed Resource Scheduler (DRS).** A VMware vSphere cluster can be configured to load balance the resource needs of the VMs it is hosting. VMs with resource contentions can be hot migrated to other nodes in the cluster to make sure that enough resources are available.

Hardware requirements

Compute

Google Cloud periodically requests updated validation of partner server platforms with new releases of Anthos through their Anthos Ready platform partner program. A listing of currently validated server platforms and the versions of Anthos supported can be found [here](#).

Operating system

Anthos clusters on VMware can be deployed to both vSphere 7 and 8 environments as chosen by the customer to help match their current datacenter infrastructure.

The following table contains a list vSphere versions that have been used by NetApp and our partners to validate the solution.

Operating System	Release	Anthos Versions
VCenter	8.0.1	1.28

Additional hardware

To complete the deployment of Anthos with NetApp as a fully validated solution, additional data center components for networking and storage have been tested by NetApp and our partner engineers.

The following table includes information about these additional infrastructure components.

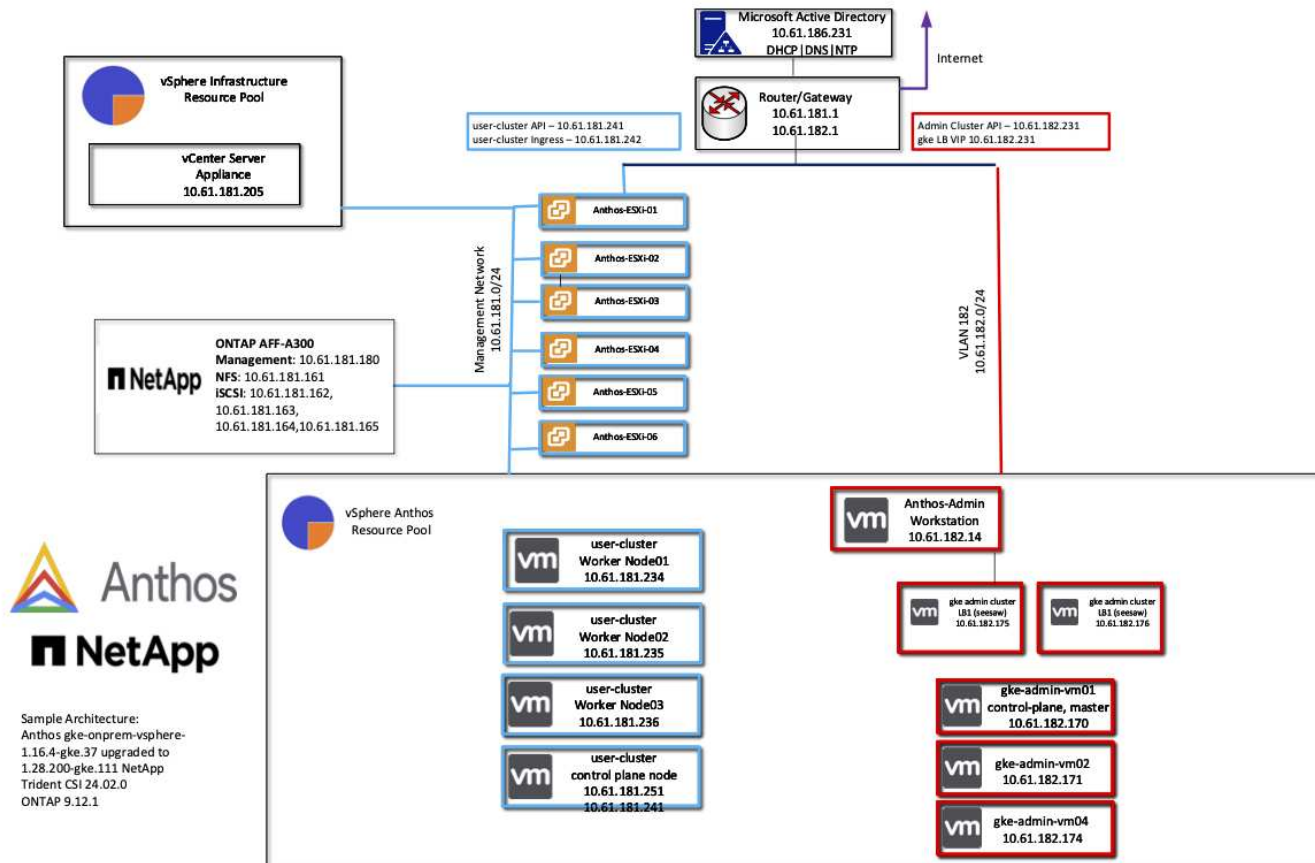
Manufacturer	Hardware Component
Mellanox	switch (data network)
Cisco	switch (management network)
NetApp	AFF Storage System

Additional software

The following table includes a list of software versions deployed in the validation environment.

Manufacturer	Software Name	Version
NetApp	ONTAP	9.12.1
NetApp	Trident	24.02.0

During the Anthos Ready platform validation performed by NetApp, the lab environment was built based on the following diagram, which allowed us to test multiple scenarios using various NetApp ONTAP storage backends.



Network infrastructure support resources

The following infrastructure should be in place prior to the deployment of Anthos:

- At least one DNS server providing full host-name resolution that is accessible from the in-band management network and the VM network.
- At least one NTP server that is accessible from the in-band management network and the VM network.
- A DHCP server available to provide network address leases on demand should clusters need to scale dynamically.
- (Optional) Outbound internet connectivity for both the in-band management network and the VM network.

Best practices for production deployments

This section lists several best practices that an organization should take into consideration before deploying this solution into production.

Deploy Anthos to an ESXi cluster of at least three nodes

Although it is possible to install Anthos in a vSphere cluster of less than three nodes for demonstration or evaluation purposes, this is not recommended for production workloads. Although two nodes allow for basic HA and fault tolerance, an Anthos cluster configuration must be modified to disable default host affinity, and this deployment method is not supported by Google Cloud.

Configure virtual machine and host affinity

Distributing Anthos cluster nodes across multiple hypervisor nodes can be achieved by enabling VM and host affinity.

Affinity or anti-affinity is a way to define rules for a set of VMs and/or hosts that determine whether the VMs run together on the same host or hosts in the group or on different hosts. It is applied to VMs by creating affinity groups that consist of VMs and/or hosts with a set of identical parameters and conditions. Depending on whether the VMs in an affinity group run on the same host or hosts in the group or separately on different hosts, the parameters of the affinity group can define either positive affinity or negative affinity.

To configure affinity groups, see the appropriate link below for your version of VMWare vSphere.

[vSphere 6.7 Documentation: Using DRS Affinity Rules.](#)

[vSphere 7.0 Documentation: Using DRS Affinity Rules.](#)



Anthos has a config option in each individual `cluster.yaml` file to automatically create node affinity rules that can be enabled or disabled based on the number of ESXi hosts in your environment.

Anthos on bare metal

The hardware-agnostic capabilities of Anthos on bare metal allow you to select a compute platform optimized for your personalized use case and also provide many additional benefits.

Benefits

The hardware-agnostic capabilities of Anthos on bare metal allow you to select a compute platform optimized for your personalized use case and also provide many additional benefits.

Examples include the following:

- **Bring your own server.** You can use servers that match your existing infrastructure to reduce capital expenditure and management costs.
- **Bring your own Linux OS.** By choosing the Linux OS that you wish to deploy your Anthos-on-bare-metal environment to, you can ensure that the Anthos environment fits neatly into your existing infrastructure and management schemes.
- **Improved performance and lowered cost.** Without the requirement of a hypervisor, Anthos-on-bare-metal clusters call for direct access to server hardware resources, including performance-optimized hardware devices like GPUs.
- **Improved network performance and lowered latency.** Because the Anthos-on-bare-metal server nodes are directly connected to your network without a virtualized abstraction layer, they can be optimized for low latency and performance.

Hardware requirements

Compute

Google Cloud periodically requests updated validation of partner server platforms with new releases of Anthos through their Anthos Ready platform partner program. A listing of currently validated server platforms and the versions of Anthos supported can be found [here](#).

The following table contains server platforms that have been tested by NetApp and NetApp partner engineers for the validation of Anthos on bare metal deployments.

Manufacturer	Make	Model
Cisco	UCS	B200 M5
HPE	Proliant	DL360

Operating System

Anthos-on-bare-metal nodes can be configured with several different Linux distributions as chosen by the customer to help match their current datacenter infrastructure.

The following table contains a list of Linux operating systems that have been used by NetApp and our partners to validate the solution.

Operating System	Release	Anthos Versions
CentOS	8.4.2105	1.14
Red Hat Enterprise Linux	8.4	1.14
Ubuntu	18.04.5 LTS (with kernel 5.4.0-81-generic)	1.14
Ubuntu	20.04.2 LTS	1.14

Additional hardware

To complete the deployment of Anthos on bare metal as a fully validated solution, additional data center components for networking and storage have been tested by NetApp and our partner engineers.

The following table includes information about these additional infrastructure components.

Manufacturer	Hardware Name	Model
Cisco	Nexus	C9336C-FX2
NetApp	AFF	A250, A220


Additional software


The following table includes a list of additional software versions deployed in the validation environment.

Manufacturer	Software name	Version
Cisco	NXOS	9.3(5)
NetApp	ONTAP	9.11.1P4
NetApp	Trident	23.01.0

During the Anthos Ready platform validation performed by NetApp and our partner team at World Wide Technology (WWT), the lab environment was built based on the following diagram, which allowed us to test the functionality of each server type, operating system, the network devices, and storage systems deployed in the solution.

Anthos BareMetal Physical Hardware and Network Diagram





This multi-OS environment shows interoperability with supported OS versions for the Anthos-on-bare-metal solution. We anticipate that customers will standardize on one or a subset of operating systems for their deployment.

Infrastructure support resources

The following infrastructure should be in place prior to the deployment of Anthos on bare metal:

- At least one DNS server that provides a full host-name resolution accessible from the management network.
- At least one NTP server that is accessible from the management network.
- (Optional) Outbound internet connectivity for both the in-band management network.

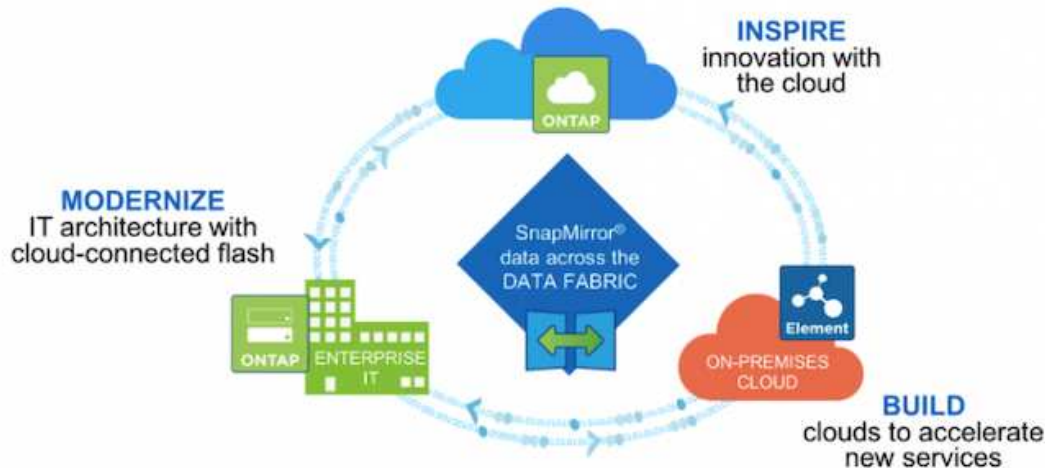


There is a demo video of an Anthos on bare metal deployment in the Videos and Demos section of this document.

NetApp storage systems

NetApp Storage Overview

NetApp has several storage platforms that are qualified with our Trident Storage Orchestrator to provision storage for applications deployed as containers.



- AFF and FAS systems run NetApp ONTAP and provide storage for both file-based (NFS) and block-based (iSCSI) use cases.
- Cloud Volumes ONTAP and ONTAP Select provide the same benefits in the cloud and virtual space respectively.
- Google Cloud NetApp Volumes (GCP) and Azure NetApp Files provide file-based storage in the cloud.
- Amazon FSx ONTAP is a fully managed service on AWS that provides storage for file-based use cases.



Each storage system in the NetApp portfolio can ease both data management and movement between on-premises sites and the cloud, ensuring that your data is where your applications are.

NetApp ONTAP

NetApp ONTAP is a powerful storage-software tool with capabilities such as an intuitive GUI, REST APIs with automation integration, AI-informed predictive analytics and corrective action, nondisruptive hardware upgrades, and cross-storage import.

For more information about the NetApp ONTAP storage system, visit the [NetApp ONTAP website](#).

ONTAP provides the following features:


- A unified storage system with simultaneous data access and management of NFS, CIFS, iSCSI, FC, FCoE, and FC-NVMe protocols.

- Different deployment models include on-premises on all-flash, hybrid, and all-HDD hardware configurations; VM-based storage platforms on a supported hypervisor such as ONTAP Select; and in the cloud as Cloud Volumes ONTAP.
- Increased data storage efficiency on ONTAP systems with support for automatic data tiering, inline data compression, deduplication, and compaction.
- Workload-based, QoS-controlled storage.
- Seamless integration with a public cloud for tiering and protection of data. ONTAP also provides robust data protection capabilities that sets it apart in any environment:
 - **NetApp Snapshot copies.** A fast, point-in-time backup of data using a minimal amount of disk space with no additional performance overhead.
 - **NetApp SnapMirror.** Mirrors the Snapshot copies of data from one storage system to another. ONTAP supports mirroring data to other physical platforms and cloud-native services as well.
 - **NetApp SnapLock.** Efficiently administration of nonrewritable data by writing it to special volumes that cannot be overwritten or erased for a designated period.
 - **NetApp SnapVault.** Backs up data from multiple storage systems to a central Snapshot copy that serves as a backup to all designated systems.
 - **NetApp SyncMirror.** Provides real-time, RAID-level mirroring of data to two different plexes of disks that are connected physically to the same controller.
 - **NetApp SnapRestore.** Provides fast restoration of backed-up data on demand from Snapshot copies.
 - **NetApp FlexClone.** Provides instantaneous provisioning of a fully readable and writeable copy of a NetApp volume based on a Snapshot copy.

For more information about ONTAP, see the [ONTAP 9 Documentation Center](#).



NetApp ONTAP is available on-premises, virtualized, or in the cloud.



NetApp platforms

NetApp AFF/FAS

NetApp provides robust all-flash (AFF) and scale-out hybrid (FAS) storage platforms that are tailor-made with low-latency performance, integrated data protection, and multiprotocol support.

Both systems are powered by NetApp ONTAP data management software, the industry's most advanced data-management software for highly-available, cloud-integrated, simplified storage management to deliver the enterprise-class speed, efficiency, and security your data fabric needs.

For more information about NETAPP AFF and FAS platforms, click [here](#).

ONTAP Select

ONTAP Select is a software-defined deployment of NetApp ONTAP that can be deployed onto a hypervisor in your environment. It can be installed on VMware vSphere or on KVM and provides the full functionality and experience of a hardware-based ONTAP system.

For more information about ONTAP Select, click [here](#).

Cloud Volumes ONTAP

NetApp Cloud Volumes ONTAP is a cloud-deployed version of NetApp ONTAP available to be deployed in a number of public clouds, including: Amazon AWS, Microsoft Azure, and Google Cloud.

For more information about Cloud Volumes ONTAP, click [here](#).

NetApp storage integrations

NetApp Storage Integration Overview

NetApp provides a number of products which assist our customers with orchestrating and managing persistent data in container-based environments like Anthos.

Anthos Ready storage partner program.

Google Cloud periodically requests updated validation of partner storage integrations with new releases of Anthos through their Anthos Ready storage partner program. A list of currently validated storage solutions, CSI drivers, available features, and the versions of Anthos supported can be found [here](#).

NetApp has maintained regular compliance on a quarterly basis with requests to validate our Trident CSI-compliant storage orchestrator and our ONTAP storage system with versions of Anthos.

The following table contains the Anthos versions tested by NetApp and NetApp partner engineers for validation of NetApp Trident CSI drivers and feature sets as a part of the Anthos Ready storage partner program:

Deployment Type	Version	Storage System	Trident Version	Protocol	Features
-----------------	---------	----------------	-----------------	----------	----------

VMware	1.28	ONTAP 9.12.1	24.02	NAS	Multiwriter, Volume Expansion, SnapShots, PVCDDataSource
VMware	1.28	ONTAP 9.12.1	24.02	SAN	Raw Block, Volume Expansion, SnapShots, PVCDDataSource
VMware	1.15	ONTAP 9.12.1	23.04	NAS	Multiwriter, Volume Expansion, SnapShots, PVCDDataSource
VMware	1.15	ONTAP 9.12.1	23.04	SAN	Raw Block, Volume Expansion, SnapShots, PVCDDataSource
VMware	1.14	ONTAP 9.12.1	23.01	NAS	Multiwriter, Volume Expansion, SnapShots, PVCDDataSource
VMware	1.14	ONTAP 9.12.1	23.01	SAN	Raw Block, Volume Expansion, SnapShots, PVCDDataSource
VMware	1.13	ONTAP 9.12.1	22.10	NAS	Multiwriter, Volume Expansion, SnapShots, PVCDDataSource
VMware	1.13	ONTAP 9.12.1	22.10	SAN	Raw Block, Volume Expansion, SnapShots, PVCDDataSource
VMware	1.11	ONTAP 9.9.1	22.04	NAS	Multiwriter, Volume Expansion, SnapShots
VMware	1.11	ONTAP 9.9.1	22.04	SAN	Raw Block, Volume Expansion, SnapShots

VMware	1.11	Element 12.3	22.04	SAN	Raw Block, Volume Expansion, SnapShots
bare metal	1.10	ONTAP 9.8	22.01	NAS	Multiwriter, Volume Expansion, SnapShots
bare metal	1.10	ONTAP 9.8	22.01	SAN	Raw Block, Volume Expansion, SnapShots

NetApp storage integrations

NetApp provides a number of products to help you with orchestrating and managing persistent data in container-based environments such as Anthos.

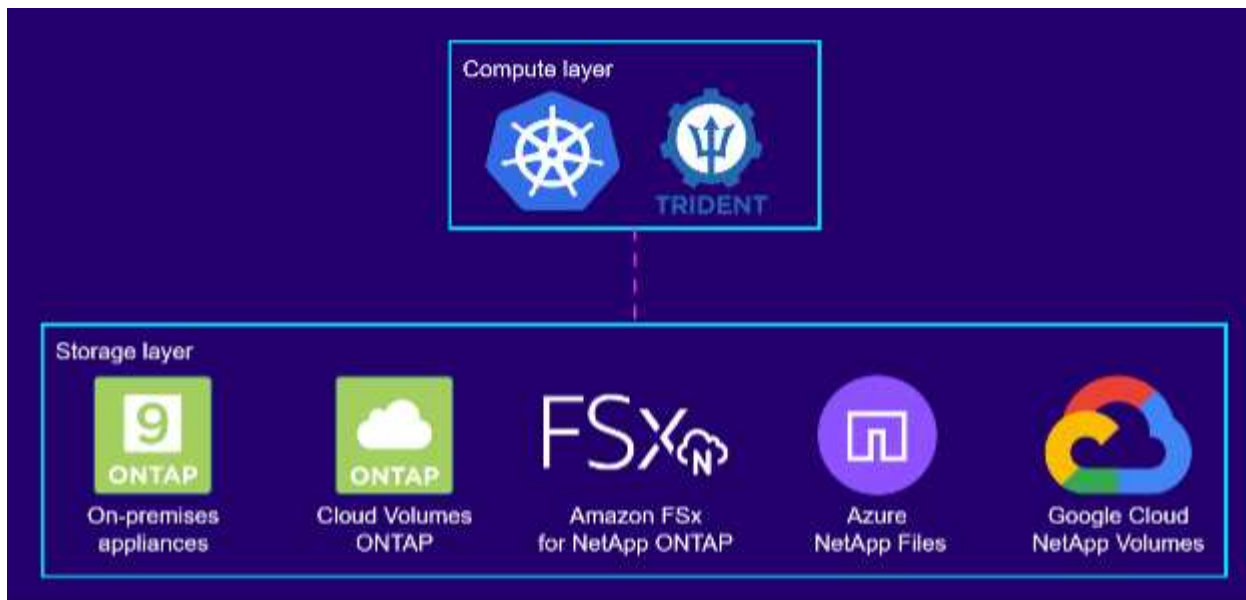
NetApp Trident is an open-source, fully-supported storage orchestrator for containers and Kubernetes distributions, including Anthos. For more information, visit the Trident website [here](#).

The following pages have additional information about the NetApp products that have been validated for application and persistent-storage management in the Anthos with NetApp solution.

Trident Overview

Trident is a fully supported, open-source storage orchestrator for containers and Kubernetes distributions, including Anthos. Trident works with the entire NetApp storage portfolio, including NetApp ONTAP, and it also supports NFS and iSCSI connections. Trident accelerates the DevOps workflow by allowing end users to provision and manage storage from their NetApp storage systems without requiring intervention from a storage administrator.

An administrator can configure a number of storage backends based on project needs and storage system models that enable advanced storage features, including compression, specific disk types, and QoS levels that guarantee a certain level of performance. After they are defined, these backends can be used by developers in their projects to create persistent volume claims (PVCs) and to attach persistent storage to their containers on demand.



Trident has a rapid development cycle and, like Kubernetes, is released four times a year.

The documentation for the latest version of Trident can be found [here](#). A support matrix for what version of Trident has been tested with which Kubernetes distribution can be found [here](#).

Starting with the 20.04 release, Trident setup is performed by the Trident operator. The operator makes large scale deployments easier and provides additional support including self healing for pods that are deployed as a part of the Trident install.

With the 22.04 release, a Helm chart was made available to ease the installation of the Trident Operator.

For Trident installation details, please see [here](#).

Create a storage-system backend

After completing the Trident Operator install, you must configure the backend for the specific NetApp storage platform you are using. Follow the link below in order to continue the setup and configuration of Trident.
[Create a backend.](#)

Create a storage class

After creating the backend, you must create a storage class that Kubernetes users will specify when they want a volume. Kubernetes users provision volumes by using persistent volume claims (PVCs) that specify a storage class by name.

Follow the link below to create a storage class.

[Create a storage class](#)

Dynamically provision a volume

You must create a Kubernetes persistent volume claim (PVC) object using the storage class to dynamically provision a volume. Follow the link below to create a PVC object.

[Create a PVC](#)

Use the volume

The volume provisioned in the above step can be used by an application by mounting the volume in the

pod. The link below shows an example.
[Mount the volume in a pod](#)

Sample provisioning

[Sample manifests for iSCSI protocol](#)

[Sample manifests for nfs protocol](#)

Advanced configuration options

Advanced configuration options

Typically, the easiest-to-deploy solution is best, but, in some cases, advanced customizations are required to meet the requirements or specifications of a specific application or the environment that solution is being deployed to. To this end, the Red Hat OpenShift with NetApp solution allows for the following customizations to meet these needs.



In this section we have documented some advanced configuration options such as using third-party load balancers or creating a private registry for hosting customized container images, both of which are prerequisites for installing the NetApp Trident Protect.

The following pages have additional information about the advanced configuration options validated in the Red Hat OpenShift with NetApp solution:

Explore load balancer options

Exploring load balancer options

An application deployed in Anthos is exposed to the world by a service that is delivered by a load balancer deployed in the Anthos on-prem environment.

The following pages have additional information about load balancer options validated in the Anthos with NetApp solution:

- [Installing F5 BIG-IP load balancers](#)
- [Installing MetalLB load balancers](#)
- [Installing SeeSaw load balancers](#)

Installing F5 BIG-IP load balancers

F5 BIG-IP is an Application Delivery Controller (ADC) that offers a broad set of advanced, production-grade traffic management and security services like L4-L7 load balancing, SSL/TLS offload, DNS, firewall, and more. These services dramatically increase the availability, security, and performance of your applications.

F5 BIG-IP can be deployed and consumed in various ways, including on dedicated hardware, in the cloud, or as a virtual appliance on-premises. Refer to the documentation [here](#) to explore and deploy F5 BIG-IP.

F5 BIG-IP was the first of the bundled load balancer solutions available with Anthos On-Prem and was used in a number of the early Anthos Ready partner validations for the Anthos with NetApp solution.



F5 BIG-IP can be deployed in standalone mode or in cluster mode. For the purpose of this validation, F5 BIG-IP was deployed in standalone mode. However, for production purposes, NetApp recommends creating a cluster of BIG-IP instances to avoid a single point of failure.



An F5 BIG-IP system can be deployed on dedicated hardware, in the cloud, or as a virtual appliance on-premises with versions greater than 12.x for it to be integrated with F5 CIS. For the purpose of this document, the F5 BIG-IP system was validated as a virtual appliance, for example using the BIG-IP VE edition.

Validated releases

This solution makes use of the virtual appliance deployed in VMware vSphere. Networking for the F5 Big-IP virtual appliance can be configured in a two-armed or three-armed configuration based on your network environment. The deployment in this document is based on the two-armed configuration. Additional details on configuring the virtual appliance for use with Anthos can be found [here](#).

The Solutions Engineering Team at NetApp have validated the releases in the following table in our lab to work with deployments of Anthos On-Prem:

Make	Type	Version
F5	BIG-IP VE	15.0.1-0.0.11
F5	BIG-IP VE	16.1.0-0.0.19

Installation

To install F5 BIG-IP, complete the following steps:

1. Download the virtual application Open Virtual Appliance (OVA) file from F5 [here](#).



To download the appliance, a user must register with F5. They provide a 30-day demo license for the Big-IP Virtual Edition Load Balancer. NetApp recommends a permanent 10Gbps license for the production deployment of an appliance.

2. Right-click the Infrastructure Resource Pool and select Deploy OVF Template. A wizard launches that allows you to select the OVA file that you just downloaded in Step 1. Click Next.

Deploy OVF Template

1 Select an OVF template

2 Select a name and folder

3 Select a compute resource

4 Review details

5 Select storage

6 Ready to complete

Select an OVF template

Select an OVF template from remote URL or local file system

Enter a URL to download and install the OVF package from the Internet, or browse to a location accessible from your computer, such as a local hard drive, a network share, or a CD/DVD drive.

☐ URL

http | https://remoteserver-address/filetoinstall.ovf | .ova

☒ Local file

Choose Files

BIGIP-15.0.1-0.....ALL-vmware.ova

CANCEL

BACK

NEXT

- Click Next to continue through each step and accept the default values for each screen presented until you reach the storage selection screen. Select the VM_Datastore that you would like to deploy the virtual machine to, and then click Next.
- The next screen presented by the wizard allows you to customize the virtual networks for use in the environment. Select VM_Network for the External field and select Management_Network for the Management field. Internal and HA are used for advanced configurations for the F5 Big-IP appliance and are not configured. These parameters can be left alone, or they can be configured to connect to non-infrastructure, distributed port groups. Click Next.
- Review the summary screen for the appliance, and, if all the information is correct, click Finish to start the deployment.
- After the virtual appliance is deployed, right-click it and power it up. It should receive a DHCP address on the management network. The appliance is Linux-based, and it has VMware Tools deployed, so you can view the DHCP address it receives in the vSphere client.
- Open a web browser and connect to the appliance at the IP address from the previous step. The default login is admin/admin, and, after the first login, the appliance immediately prompts you to change the admin password. It then returns you to a screen where you must log in with the new credentials.

f5 BIG-IP Configuration Utility
F5 Networks, Inc.

Hostname
bigip1

IP Address
172.21.224.20

Username
admin

Password

Log in

Welcome to the BIG-IP Configuration Utility.

Log in with your username and password using the fields on the left.

(c) Copyright 1996-2019, F5 Networks, Inc., Seattle, Washington. All rights reserved.
[F5 Networks, Inc. Legal Notices](#)

8. The first screen prompts the user to complete the Setup Utility. Begin the utility by clicking Next.

9. The next screen prompts for activation of the license for the appliance. Click Activate to begin. When prompted on the next page, paste either the 30-day evaluation license key you received when you registered for the download or the permanent license you acquired when you purchased the appliance. Click Next.



For the device to perform activation, the network defined on the management interface must be able to reach the internet.

10. On the next screen, the End User License Agreement (EULA) is presented. If the terms in the license are acceptable, click Accept.

11. The next screen counts the elapsed time as it verifies the configuration changes that have been made so far. Click Continue to resume with the initial configuration.

12. The Configuration Change window closes, and the Setup Utility displays the Resource Provisioning menu. This window lists the features that are currently licensed and the current resource allocations for the virtual appliance and each running service.

13. Clicking the Platform menu option on the left enables additional modification of the platform. Modifications include setting the management IP address configured with DHCP, setting the host name and the time zone the appliance is installed in, and securing the appliance from SSH accessibility.
14. Next click the Network menu, which enables you to configure standard networking features. Click Next to begin the Standard Network Configuration wizard.
15. The first page of the wizard configures redundancy; leave the defaults and click Next. The next page enables you to configure an internal interface on the load balancer. Interface 1.1 maps to the VMNIC labeled Internal in the OVF deployment wizard.



The spaces in this page for Self IP Address, Netmask, and Floating IP address can be filled with a non-routable IP for use as a placeholder. They can also be filled with an internal network that has been configured as a distributed port group for virtual guests if you are deploying the three-armed configuration. They must be completed to continue with the wizard.

16. The next page enables you to configure an external network that is used to map services to the pods deployed in Kubernetes. Select a static IP from the VM_Network range, the appropriate subnet mask, and a floating IP from that same range. Interface 1.2 maps to the VMNIC labeled External in the OVF deployment wizard.
17. On the next page, you can configure an internal-HA network if you are deploying multiple virtual appliances in the environment. To proceed, you must fill the Self-IP Address and the Netmask fields, and you must select interface 1.3 as the VLAN Interface, which maps to the HA network defined by the OVF template wizard.
18. The next page enables you to configure the NTP servers. Then click Next to continue to the DNS setup. The DNS servers and domain search list should already be populated by the DHCP server. Click Next to accept the defaults and continue.
19. For the remainder of the wizard, click Next to continue through the advanced peering setup, the configuration of which is beyond the scope of this document. Then click Finish to exit the wizard.
20. Create individual partitions for the Anthos admin cluster and each user cluster deployed in the environment. Click System in the menu on the left, navigate to Users, and click Partition List.
21. The displayed screen only shows the current common partition. Click Create on the right to create the first additional partition, and name it GKE-Admin. Then click Repeat, and name the partition User-Cluster-1. Click the Repeat button again to name the next partition User-Cluster-2. Finally click Finished to complete the wizard. The Partition list screen returns with all the partitions now listed.

Integration with Anthos

There is a section in each configuration file, respectively for the admin cluster, and each user cluster that you choose to deploy to configure the load balancer so that it is managed by Anthos On Prem.

The following script is a sample from the configuration of the partition for the GKE-Admin cluster. The values that need to be uncommented and modified are placed in bold text below:

```
# (Required) Load balancer configuration
loadBalancer:
  # (Required) The VIPs to use for load balancing
  vips:
    # Used to connect to the Kubernetes API
    controlPlaneVIP: "10.61.181.230"
    # # (Optional) Used for admin cluster addons (needed for multi cluster
features). Must
    # # be the same across clusters
    # # addonsVIP: ""
  # (Required) Which load balancer to use "F5BigIP" "Seesaw" or
"ManualLB". Uncomment
  # the corresponding field below to provide the detailed spec
  kind: F5BigIP
  # # (Required when using "ManualLB" kind) Specify pre-defined nodeports
  # manualLB:
  #   # NodePort for ingress service's http (only needed for user cluster)
  #   ingressHTTPTNodePort: 0
  #   # NodePort for ingress service's https (only needed for user
cluster)
  #   ingressHTTPSNodePort: 0
  #   # NodePort for control plane service
  #   controlPlaneNodePort: 30968
  #   # NodePort for addon service (only needed for admin cluster)
  #   addonsNodePort: 31405
  # # (Required when using "F5BigIP" kind) Specify the already-existing
partition and
  # # credentials
  f5BigIP:
    address: "172.21.224.21"
    credentials:
      username: "admin"
      password: "admin-password"
    partition: "GKE-Admin"
  #   # # (Optional) Specify a pool name if using SNAT
  #   # snatPoolName: ""
  # (Required when using "Seesaw" kind) Specify the Seesaw configs
  # seesaw:
    # (Required) The absolute or relative path to the yaml file to use for
```

```

IP allocation
# for LB VMs. Must contain one or two IPs.
# ipBlockFilePath: ""
# (Required) The Virtual Router Identifier of VRRP for the Seesaw
group. Must
# be between 1-255 and unique in a VLAN.
# vrid: 0
# (Required) The IP announced by the master of Seesaw group
# masterIP: ""
# (Required) The number CPUs per machine
# cpus: 4
# (Required) Memory size in MB per machine
# memoryMB: 8192
# (Optional) Network that the LB interface of Seesaw runs in (default:
cluster
# network)
# vCenter:
# vSphere network name
# networkName: VM_Network
# (Optional) Run two LB VMs to achieve high availability (default:
false)
# enableHA: false

```

Installing MetalLB load balancers

This page lists the installation and configuration instructions for the MetalLB managed load balancer.

Installing The MetalLB Load Balancer

The MetalLB load balancer is fully integrated with Anthos Clusters on VMware and has automated deployment performed as part of the Admin and User cluster setups starting with the 1.11 release. There are blocks of text in the respective `cluster.yaml` configuration files that you must modify to provide load balancer info. It is self-hosted on your Anthos cluster instead of requiring the deployment of external resources like the other supported load balancer solutions. It also allows you to create an ip-pool that automatically assigns addresses with the creation of Kubernetes services of type load balancer in clusters that do not run on a cloud provider.

Integration with Anthos

When enabling the MetalLB load balancer for Anthos admin, you must modify a few lines in the `loadBalancer:` section that exists in the `admin-cluster.yaml` file. The only values that you must modify are to set the `controlPlaneVIP:` address and then set the `kind:` as MetalLB. See the following code snippet for an example:

```
# (Required) Load balancer configuration
loadBalancer:
  # (Required) The VIPs to use for load balancing
  vips:
    # Used to connect to the Kubernetes API
    controlPlaneVIP: "10.61.181.230"
    # # (Optional) Used for admin cluster addons (needed for multi cluster
    features). Must
    # # be the same across clusters
    # addonsVIP: ""
  # (Required) Which load balancer to use "F5BigIP" "Seesaw" "ManualLB" or
  "MetalLB".
  # Uncomment the corresponding field below to provide the detailed spec
  kind: MetalLB
```

When enabling the MetalLB load balancer for Anthos user clusters, there are two areas in each `user-cluster.yaml` file that you must update. First, in a manner similar to the `admin-cluster.yaml` file, you must modify the `controlPlaneVIP:`, `ingressVIP:`, and `kind:` values in the `loadBalancer:` section. See the following code snippet for an example:

```
loadBalancer:
  # (Required) The VIPs to use for load balancing
  vips:
    # Used to connect to the Kubernetes API
    controlPlaneVIP: "10.61.181.240"
    # Shared by all services for ingress traffic
    ingressVIP: "10.61.181.244"
  # (Required) Which load balancer to use "F5BigIP" "Seesaw" "ManualLB" or
  "MetalLB".
  # Uncomment the corresponding field below to provide the detailed spec
  kind: MetalLB
```



The ingressVIP IP address must exist within the pool of IP addresses assigned to the MetalLB load balancer later in the configuration.

You then need to navigate to the `metalLB:` subsection and modify the `addressPools:` section by naming the pool in the `- name:` variable. You must also create a pool of ip-addresses that MetalLB can assign to services of type LoadBalancer by providing a range to the `addresses:` variable.

```
# # (Required when using "MetalLB" kind in user clusters) Specify the
MetalLB config
  metalLB:
    # # (Required) A list of non-overlapping IP pools used by load balancer
typed services.
    # # Must include ingressVIP of the cluster.
    addressPools:
      # # (Required) Name of the address pool
      - name: "default"
      # # (Required) The addresses that are part of this pool. Each address
must be either
      # # in the CIDR form (1.2.3.0/24) or range form (1.2.3.1-1.2.3.5).
      addresses:
        - "10.61.181.244-10.61.181.249"
```



The address pool can be provided as a range like in the example, limiting it to a number of addresses in a particular subnet, or it can be provided as a CIDR notation if the entire subnet is made available.

1. When Kubernetes services of type LoadBalancer are created, MetalLB automatically assigns an externalIP to the services and advertises the IP address by responding to ARP requests.

Installing SeeSaw load balancers

This page lists the installation and configuration instructions for the SeeSaw managed load balancer.

Seesaw is the default managed network load balancer installed in an Anthos Clusters on VMware environment from versions 1.6 to 1.10.

Installing The SeeSaw load balancer

The SeeSaw load balancer is fully integrated with Anthos Clusters on VMware and has automated deployment performed as part of the Admin and User cluster setups. There are blocks of text in the `cluster.yaml` configuration files that must be modified to provide load balancer info, and then there is an additional step prior to cluster deployment to deploy the load balancer using the built in `gkectl` tool.



SeeSaw load balancers can be deployed in HA or non-HA mode. For the purpose of this validation, the SeeSaw load balancer was deployed in non-HA mode, which is the default setting. For production purposes, NetApp recommends deploying SeeSaw in an HA configuration for fault tolerance and reliability.

Integration with Anthos

There is a section in each configuration file, respectively for the admin cluster, and in each user cluster that you choose to deploy to configure the load balancer so that it is managed by Anthos On-Prem.

The following text is a sample from the configuration of the partition for the GKE-Admin cluster. The values that need to be uncommented and modified are placed in bold text below:

loadBalancer:

(Required) The VIPs to use for load balancing

vips:

Used to connect to the Kubernetes API

controlPlaneVIP: "10.61.181.230"

(Optional) Used for admin cluster addons (needed for multi cluster features). Must

be the same across clusters

addonsVIP: ""

(Required) Which load balancer to use "F5BigIP" "Seesaw" or "ManualLB". Uncomment

the corresponding field below to provide the detailed spec

kind: Seesaw

(Required when using "ManualLB" kind) Specify pre-defined nodeports

manualLB:

NodePort for ingress service's http (only needed for user cluster)

ingressHTTPTNodePort: 0

NodePort for ingress service's https (only needed for user

cluster)

ingressHTTPSNodePort: 0

NodePort for control plane service

controlPlaneNodePort: 30968

NodePort for addon service (only needed for admin cluster)

addonsNodePort: 31405

(Required when using "F5BigIP" kind) Specify the already-existing partition and

credentials

f5BigIP:

address:

credentials:

username:

password:

partition:

(Optional) Specify a pool name if using SNAT

snatPoolName: ""

(Required when using "Seesaw" kind) Specify the Seesaw configs

seesaw:

(Required) The absolute or relative path to the yaml file to use for IP allocation

for LB VMs. Must contain one or two IPs.

ipBlockFilePath: "admin-seesaw-block.yaml"

(Required) The Virtual Router Identifier of VRRP for the Seesaw group. Must

be between 1-255 and unique in a VLAN.

vrid: 100

```
# (Required) The IP announced by the master of Seesaw group
masterIP: "10.61.181.236"
# (Required) The number CPUs per machine
cpus: 1
# (Required) Memory size in MB per machine
memoryMB: 2048
# (Optional) Network that the LB interface of Seesaw runs in (default:
cluster
# network)
vCenter:
# vSphere network name
networkName: VM_Network
# (Optional) Run two LB VMs to achieve high availability (default:
false)
enableHA: false
```

The SeeSaw load balancer also has a separate static `seesaw-block.yaml` file that you must provide for each cluster deployment. This file must be located in the same directory relative to the `cluster.yaml` deployment file, or the full path must be specified in the section above.

A sample of the `admin-seesaw-block.yaml` file looks like the following script:

```
blocks:
- netmask: "255.255.255.0"
  gateway: "10.63.172.1"
  ips:
- ip: "10.63.172.152"
  hostname: "admin-seesaw-vm"
```



This file provides the gateway and netmask for the network that the load balancer provides to the underlying cluster, as well as the management IP and hostname for the virtual machine that is deployed to run the load balancer.

Solution validation and use cases

Deploy an application from the Google Cloud Console Marketplace

This section details how to deploy an application to your Anthos GKE cluster on prem, using the Google Cloud Console.

Prerequisites

- An Anthos cluster deployed on premises and registered with Google Cloud Console
- A MetalLB load balancer configured in your Anthos cluster
- An account with permissions to deploy applications to the cluster

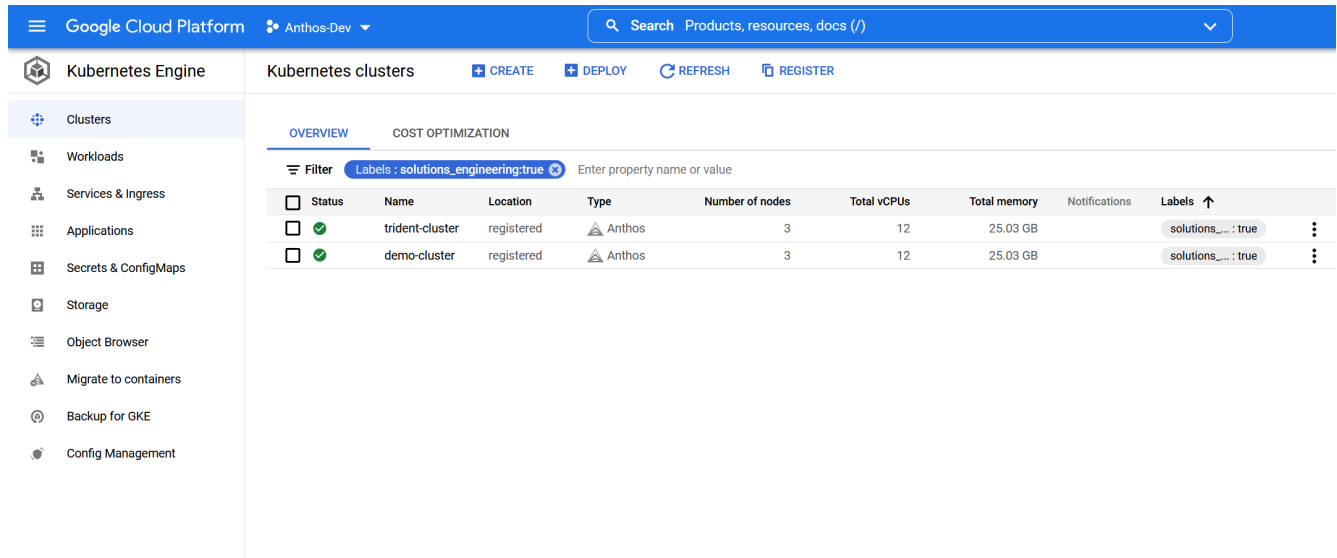
- A billing account with Google Cloud if you choose an application with associated costs (optional)

Deploying an application

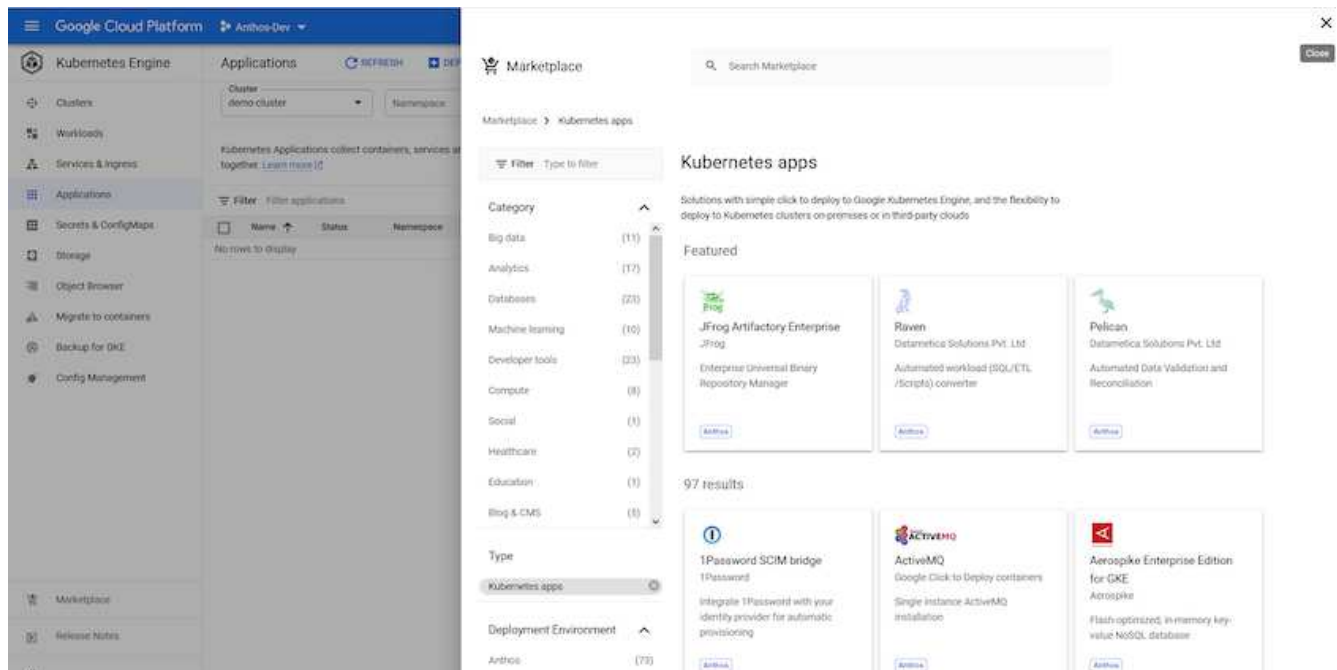
For this use case, we deploy a simple WordPress application to one of our Anthos clusters using the Google Cloud Console. The deployment uses persistent storage provided by NetApp ONTAP in a predefined storageclass. We then demonstrate two different methods to modify the applications default service so that the MetalLB load balancer supplies it with an IP address and exposes it to the world.

To deploy an application in this manner, complete the following steps:


1. Verify that the cluster you want to deploy to is accessible in Google Cloud Console.




2. Select Applications from the left-side menu, select the three-dot option menu at the top, and select Deploy from Marketplace, which brings up a new window from which you can select an application from the Google Cloud Marketplace.



3. Search for the application you want to install, in this case WordPress.



 Marketplace

Marketplace > "wordpress" > Kubernetes apps

Filter

Type to filter

Type

Kubernetes apps

Deployment Environment

Anthos (2)

GKE (2)

GKE on-prem (2)


Price

Free (1)


Paid (1)

Kubernetes apps

2 results



WordPress
Google Click to Deploy containers
WordPress is an open source publishing platform for creating websites and blogs. This application supports GKE On Prem deployment and provides Prometheus metrics, and supports Stackdriver integration.
[Anthos](#)



Bitpoke App
Bitpoke
The Bitpoke App provides a dashboard to host, deploy, scale, manage and monitor WordPress or WooCommerce sites in a Kubernetes environment. It integrates seamlessly with industry-standard tools: git, Docker, Bedrock, Composer, and Grafana - Prometheus. Further integration of other tools is possible. The solution is designed for enterprise WordPress agencies, publishers, shops, and hosting companies with millions of users.
[Anthos](#)

4. After you select the WordPress application, an overview screen displays. Click the Configure button.



WordPress

Version: 5.9 ▼

Google Click to Deploy containers

Web publishing platform for websites and blogs

CONFIGURE

Click to to launch configuration page

OVERVIEW

PRICING

DOCUMENTATION

SUPPORT

Overview

WordPress is an open source publishing platform for creating websites and blogs.

This application supports [GKE On Prem](#) deployment.

This application provides Prometheus metrics, and supports Stackdriver integration.

[Learn more](#)

About Google Click to Deploy containers

Popular open stacks packaged for containers by Google. The images serve as base images for building applications on [App Engine Flexible Environment](#), [Kubernetes Engine](#), or other Docker hosts.

About Kubernetes apps

[Google Kubernetes Engine](#) is a managed, production-ready environment for deploying containerized applications. Kubernetes apps are prepackaged applications that can be deployed to Google Kubernetes Engine in minutes.

5. On the next page you must select the cluster to deploy to, in our case Demo-Cluster. Select or create a new namespace and application instance name, and select which storage classes and persistent volume sizes you need for both the WordPress application and its backing MariaDB database. In both cases, we chose the ONTAP-NAS-CSI storage class.

Deploy WordPress

CLICK TO DEPLOY ON GKE

DEPLOY VIA COMMAND LINE

Existing Kubernetes Cluster
demo-cluster

OR CREATE A NEW CLUSTER

Namespace
anthos-wp

The namespace in which to deploy the application

App instance name *
wordpress

StorageClass for WordPress Application
ontap-nas-csi

Storage size for persistent volumes in WordPress Application
5Gi

StorageClass for MySQL Application
ontap-nas-csi

Storage size for persistent volumes in MySQL Application
5Gi

WordPress admin e-mail address *
alan.cowles@netapp.com

☐ Enable public IP access

☐ Enable Stackdriver Metrics Exporter

DEPLOY



WordPress Overview

Solution provided by Google Click to Deploy containers

Pricing

Note: There is no usage fee for this product. Charges will apply for the use of Google Kubernetes Engine. Please refer to [GCP Price List](#) for the latest pricing.

Documentation

- [User Guide](#)
Get started with Google Cloud Platform's WordPress Kubernetes application
- [Getting Started with WordPress](#)
Official WordPress documentation

Terms of Service

By deploying the software or accessing the service you are agreeing to comply with the [Google Click to Deploy containers terms of service](#), [GCP Marketplace terms of service](#) and the terms of applicable open source software licenses bundled with the software or service. Please review these terms and licenses carefully for details about any obligations you may have related to the software or service. To the limited extent an open source software license related to the software or service expressly supersedes the GCP Marketplace Terms of Service, that open source software license governs your use of that software or service.

By using this product, you understand that certain account and usage information may be shared with Google Click to Deploy containers for the purposes of financial accounting, sales attribution, performance analysis, and support.

Google is providing this software or service "as-is" and any support for this software or service will be provided by Google Click to Deploy containers under their terms of service.



Do not select Enable public IP Access. Doing so creates a service of type NodePort that is not accessible from an on-premises Anthos deployment.

- After you click the Deploy button, you get a page providing application details. You can refresh this page or log into your cluster using the CLI to check the status of the deployment.

Google Cloud Platform Anthos-Dev

Search Products, resources, docs (/)

Kubernetes Engine

Application details REFRESH EDIT DELETE HIDE INFO PANEL HELP ASSISTANT Application info

Clusters Workloads Services & Ingress Applications Secrets & ConfigMaps Storage Object Browser Migrate to containers Backup for GKE Config Management

Marketplace

wordpress

DETAILS EVENTS YAML VERSION HISTORY

Cluster	demo-cluster
Namespace	anthos-wp
Created	May 12, 2022, 12:38:34 PM
Labels	No labels set
Annotations	Not set

Components

Type	Name ↑	Status
No rows to display		

7. The CLI can be used to check the status of the application as it is being deployed by running the command to fetch pod info in our application namespace: `kubectl get pods -n anthos-wp`.

```
acowles@ac-rhel7:~$ kubectl get pods -n anthos-wp
NAME                                READY   STATUS    RESTARTS   AGE
wordpress-deployer--1-1h2bz        0/1     Error     0           28s
wordpress-mysql-0                   0/2     ContainerCreating  0           11s
wordpress-wordpress-0               0/2     ContainerCreating  0           11s
ubuntu@gke-admin-ws-2022-05-03:~$
```



Notice in this screenshot that there is a deployer pod in an error state. This is normal. This pod is a helper pod used by the Google Cloud Console to deploy the application that self-terminates after the other pods have begun their initialization process.

8. After a few moments, verify that your application is running.

```
acowles@ac-rhel7:~  
ubuntu@gke-admin-ws-2022-05-03:~$ kubectl get pods -n anthos-wp  
NAME                                READY   STATUS    RESTARTS   AGE  
wordpress-deployer--1-1h2bz        0/1     Error     0           4m20s  
wordpress-mysql-0                  2/2     Running   0           4m3s  
wordpress-wordpress-0              2/2     Running   1 (2m22s ago) 4m3s  
ubuntu@gke-admin-ws-2022-05-03:~$
```

Exposing the application

After the application is deployed, you have two methods to assign it a world-reachable IP.

Using the Google Cloud Console

You can expose the application by using the Google Cloud Console and editing the YAML output for the services in a browser to set a publicly reachable IP. To do so, perform the following steps:

1. In the Google Cloud Console, click Services and Ingress in the left-side menu.

Google Cloud Platform | Anthos-Dev | Search | Products, resources, docs (/)

Kubernetes Engine | Services & Ingress | REFRESH | CREATE INGRESS | DELETE

Cluster: demo-cluster | Namespace: anthos-wp | RESET | SAVE

SERVICES | INGRESS

Services are sets of Pods with a network endpoint that can be used for discovery and load balancing. Ingresses are collections of rules for routing external HTTP(S) traffic to Services.

Filter: Is system object: False | Filter services and ingresses

<input type="checkbox"/>	Name ↑	Status	Type	Endpoints	Pods	Namespace	Clusters
<input type="checkbox"/>	wordpress-apache-exporter-svc	OK	Cluster IP	None	1/1	anthos-wp	demo-cluster
<input type="checkbox"/>	wordpress-mysql-svc	OK	Cluster IP	None	1/1	anthos-wp	demo-cluster
<input type="checkbox"/>	wordpress-mysqld-exporter-svc	OK	Cluster IP	None	1/1	anthos-wp	demo-cluster
<input type="checkbox"/>	wordpress-wordpress-svc	OK	Cluster IP	10.96.8.66	1/1	anthos-wp	demo-cluster

- Click the `wordpress-wordpress-svc` service. This opens the Service Details screen. Click the Edit button at the top.

Google Cloud Platform | Anthos-Dev | Search | Products, resources, docs (/)

Kubernetes Engine | Service details | REFRESH | EDIT | DELETE | OPERATIONS

✓ wordpress-wordpress-svc

OVERVIEW | DETAILS | EVENTS | LOGS | YAML

Select the Cloud Monitoring account to see charts.

Cluster: demo-cluster
 Namespace: anthos-wp
 Labels: app.kubernetes.io/com...: wordpress-webserver, app.kubernetes.io/name: wordpress
 Type: ClusterIP

Cluster IP

Cluster IP: 10.96.8.66

Serving pods

Name	Status	Endpoints	Restarts	Created on ↑
wordpress-wordpress-0	Running	192.168.1.18	1	May 19, 2022, 11:18:58 AM

- The Editing Service Details page opens containing the YAML information for the service. Scroll down until

you see the `spec:` section and the `type:` value, which is set to `ClusterIP`. Change this value to `LoadBalancer` and click the `Save` button.

The screenshot shows the Google Cloud Platform console interface. The left sidebar contains navigation links for Kubernetes Engine, Clusters, Workloads, Services & Ingress (selected), Applications, Secrets & ConfigMaps, Storage, Object Browser, Migrate to containers, Backup for GKE, Config Management, Marketplace, and Release Notes. The main content area is titled 'Editing Service details' and shows the 'wordpress-wordpress-svc' service. The 'YAML' tab is active, displaying the service's configuration. The 'spec' section is visible, and the 'type: ClusterIP' line is circled in red. The 'status' section shows 'loadBalancer: {}'.

```
58 | uid: 1748ce66-bf0b-4f3e-8af0-46852cf1cf01
59 | resourceVersion: "7972605"
60 | uid: da707297-96d3-4135-a565-408029710b71
61 | spec:
62 |   clusterIP: 10.96.8.66
63 |   clusterIPs:
64 |   - 10.96.8.66
65 |   internalTrafficPolicy: Cluster
66 |   ipFamilies:
67 |   - IPv4
68 |   ipFamilyPolicy: SingleStack
69 |   ports:
70 |   - name: http
71 |     port: 80
72 |     protocol: TCP
73 |     targetPort: http
74 |   selector:
75 |     app.kubernetes.io/component: wordpress-webserver
76 |     app.kubernetes.io/name: wordpress
77 |   sessionAffinity: None
78 |   type: ClusterIP
79 | status:
80 |   loadBalancer: {}
81 |
```

The screenshot shows the Google Cloud Platform console interface, similar to the previous one. The 'YAML' tab is active, and the 'spec' section is visible. The 'type: LoadBalancer' line in the 'spec' section is circled in red. The 'status' section shows 'loadBalancer: {}'.

```
58 | uid: 1748ce66-bf0b-4f3e-8af0-46852cf1cf01
59 | resourceVersion: "7972605"
60 | uid: da707297-96d3-4135-a565-408029710b71
61 | spec:
62 |   clusterIP: 10.96.8.66
63 |   clusterIPs:
64 |   - 10.96.8.66
65 |   internalTrafficPolicy: Cluster
66 |   ipFamilies:
67 |   - IPv4
68 |   ipFamilyPolicy: SingleStack
69 |   ports:
70 |   - name: http
71 |     port: 80
72 |     protocol: TCP
73 |     targetPort: http
74 |   selector:
75 |     app.kubernetes.io/component: wordpress-webserver
76 |     app.kubernetes.io/name: wordpress
77 |   sessionAffinity: None
78 |   type: LoadBalancer
79 | status:
80 |   loadBalancer: {}
81 |
```


- When you return to the Service Details page, the `Type`: now lists `LoadBalancer` and the `External endpoints`: field lists an assigned IP address from the MetalLB pool and the port through which the application is accessible.

The screenshot shows the Google Cloud Platform console interface. On the left is a navigation menu with options like Clusters, Workloads, Services & Ingress, Applications, Secrets & ConfigMaps, Storage, Object Browser, Migrate to containers, Backup for GKE, and Config Management. The main panel displays the 'Service details' for 'wordpress-wordpress-svc'. It includes tabs for OVERVIEW, DETAILS, EVENTS, LOGS, and YAML. A message at the top says 'Select the Cloud Monitoring account to see charts.' Below this, a table lists service attributes: Cluster (demo-cluster), Namespace (anthos-wp), Labels (app.kubernetes.io/com...: wordpress-webserver, app.kubernetes.io/name: wordpress), Type (LoadBalancer), and External endpoints (10.61.181.245:80). The 'Type' and 'External endpoints' fields are circled in red. Under the 'Load Balancer' section, it shows Cluster IP (10.96.8.66), Load balancer IP (10.61.181.245), and Load balancer ID (ada70729796d34135a565408029710b7). At the bottom, there is a 'Serving pods' table with columns for Name, Status, Endpoints, Restarts, and Created on.

Patching the service with Kubectl

You can expose the application by using the CLI and the `kubectl patch` command to modify your deployment and set a publicly reachable IP. To do so, complete the following steps:

- List the services associated with the pods in your namespace with the `kubectl get services -n anthos-wp` command.

```
acowles@ac-rhel7:~  
ubuntu@gke-admin-ws-2022-05-03:~$ kubectl get services -n anthos-wp  
NAME                                TYPE      CLUSTER-IP    EXTERNAL-IP    PORT(S)    AGE  
wordpress-apache-exporter-svc      ClusterIP  None          <none>         9117/TCP   57m  
wordpress-mysql-svc               ClusterIP  None          <none>         3306/TCP   57m  
wordpress-mysqld-exporter-svc      ClusterIP  None          <none>         9104/TCP   57m  
wordpress-wordpress-svc           ClusterIP  10.96.8.66    <none>         80/TCP     57m  
ubuntu@gke-admin-ws-2022-05-03:~$
```

2. Modify the service type from ClusterIP to type Loadbalancer using the following command:

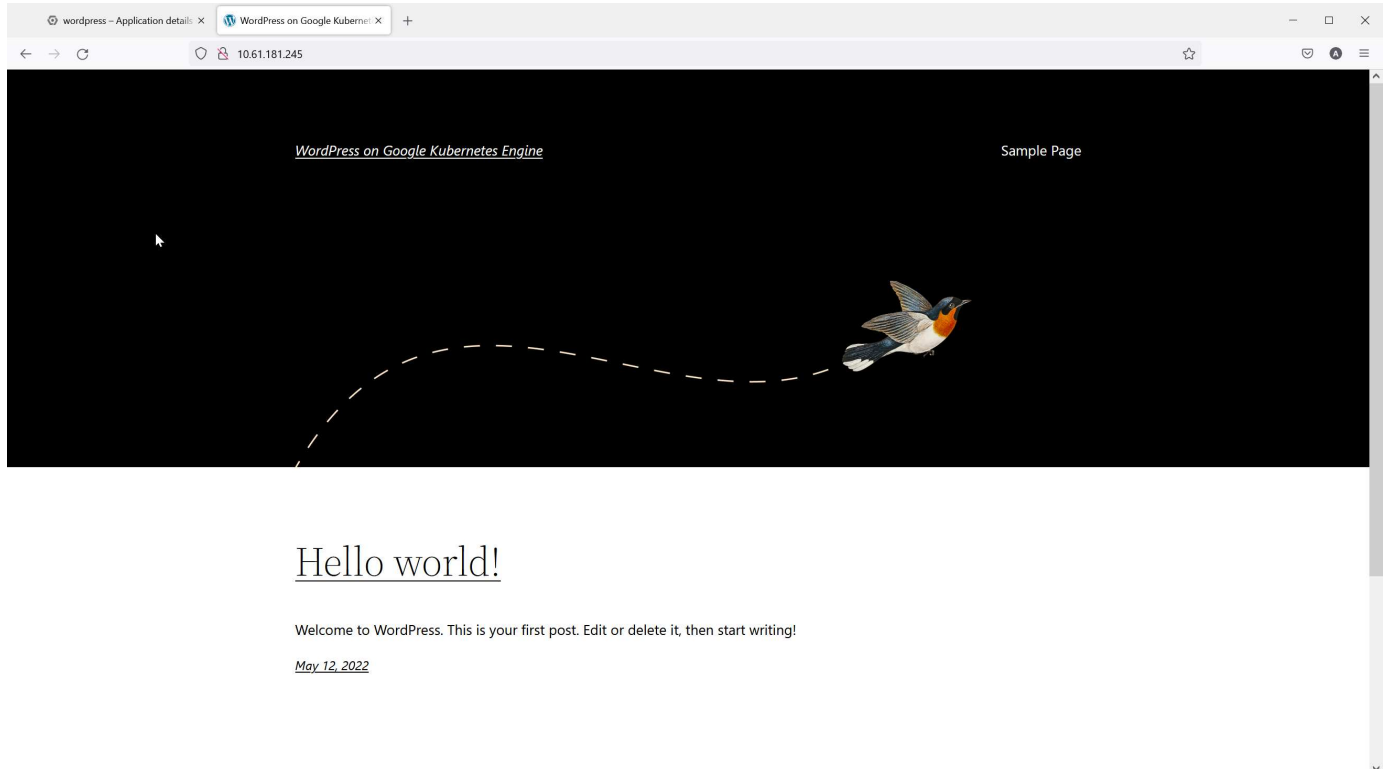
```
kubectl patch svc wordpress-wordpress-svc -p '{"spec": {"type":  
"LoadBalancer"}}' -n anthos-wp'.
```

This new service type is automatically assigned an available IP address from the MetalLB pool.

```
acowles@ac-rhel7:~  
ubuntu@gke-admin-ws-2022-05-03:~$ kubectl get services -n anthos-wp  
NAME                                TYPE      CLUSTER-IP    EXTERNAL-IP    PORT(S)    AGE  
wordpress-apache-exporter-svc      ClusterIP  None          <none>         9117/TCP   119m  
wordpress-mysql-svc               ClusterIP  None          <none>         3306/TCP   119m  
wordpress-mysqld-exporter-svc      ClusterIP  None          <none>         9104/TCP   119m  
wordpress-wordpress-svc           ClusterIP  10.96.8.66    <none>         80/TCP     119m  
ubuntu@gke-admin-ws-2022-05-03:~$ kubectl patch svc wordpress-wordpress-svc -p '{"spec": {"type": "LoadBalancer"}}' -n anthos-wp  
service/wordpress-wordpress-svc patched  
ubuntu@gke-admin-ws-2022-05-03:~$ kubectl get services -n anthos-wp  
NAME                                TYPE      CLUSTER-IP    EXTERNAL-IP    PORT(S)    AGE  
wordpress-apache-exporter-svc      ClusterIP  None          <none>         9117/TCP   120m  
wordpress-mysql-svc               ClusterIP  None          <none>         3306/TCP   120m  
wordpress-mysqld-exporter-svc      ClusterIP  None          <none>         9104/TCP   120m  
wordpress-wordpress-svc           LoadBalancer  10.96.8.66    10.61.181.245  80:30836/TCP 120m  
ubuntu@gke-admin-ws-2022-05-03:~$
```


Visit the application at the exposed external IP

Now that you have the application exposed with a publicly reachable IP address, you can visit your WordPress instance using a browser.



Where to find additional information

To learn more about the information described in this document, review the following websites:

- NetApp Documentation
<https://docs.netapp.com/>
- NetApp Trident Documentation
<https://docs.netapp.com/us-en/trident/index.html>
- Anthos Clusters on VMware Documentation
<https://cloud.google.com/anthos/clusters/docs/on-prem/latest/overview>
- Anthos on bare metal Documentation
<https://cloud.google.com/anthos/clusters/docs/bare-metal/latest>
- VMware vSphere Documentation
<https://docs.vmware.com/>

Copyright information

Copyright © 2025 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.