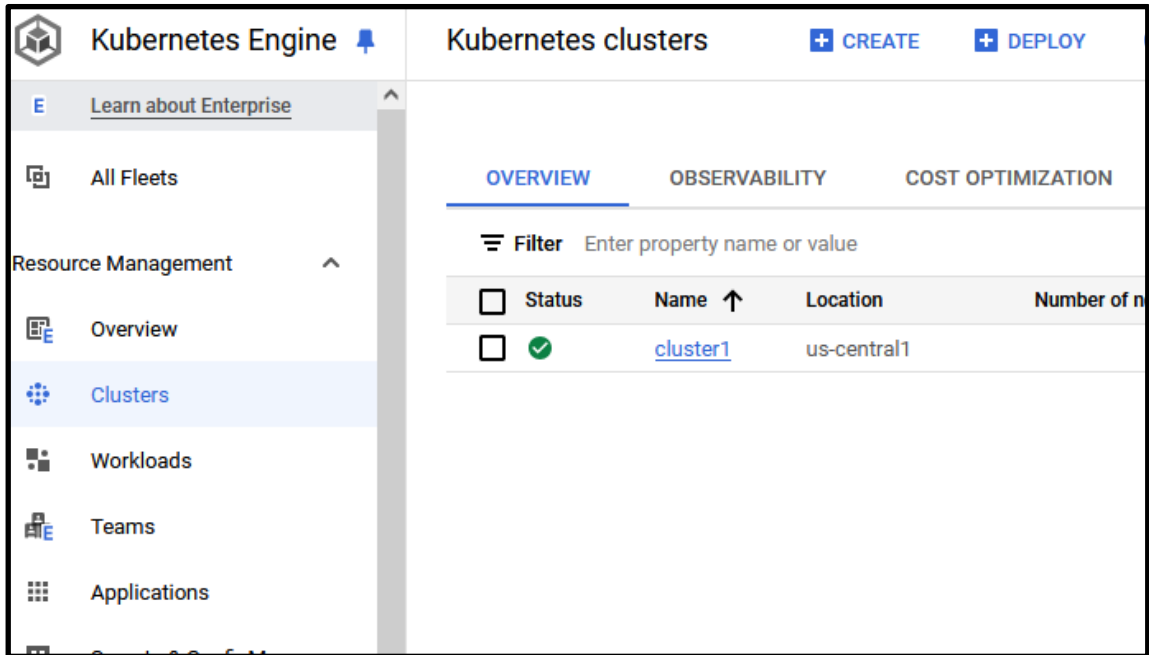


Google Cloud Platform Practices

1. Kubernetes

- Let's create a Kubernetes cluster called cluster2



- We call it "cluster2 and click on the "NEXT-FLEET REGISTRATION" button.

Apasoft

Cluster basics

Create an Autopilot cluster by specifying a name and region. After the cluster is created, you can deploy your workload through Kubernetes and we'll take care of the rest, including:

- ✓ **Nodes:** Automated node provisioning, scaling, and maintenance
- ✓ **Networking:** VPC-native traffic routing for clusters
- ✓ **Security:** Shielded GKE Nodes and Workload Identity
- ✓ **Telemetry:** Cloud Operations logging and monitoring

Name

Cluster names must start with a lowercase letter followed by up to 39 lowercase letters, numbers, or hyphens. They can't end with a hyphen. You cannot change the cluster's name once it's created.

Region

The regional location in which your cluster's control plane and nodes are located. You cannot change the cluster's region once it's created.

[NEXT: FLEET REGISTRATION](#) RESET SETTINGS

- On this screen we do not modify anything and click on the "NEXT-NETWORKING" button"

Fleet registration NEW

A fleet lets you logically group and normalize Kubernetes clusters, helping you uplevel management from individual clusters to groups of clusters. To use multi-cluster capabilities and apply consistent policies across your systems, register your cluster to a fleet. [Learn about fleets](#)

Register cluster to the fleet ?

i Your cluster will be hosted in project 'centering-oxide-437720-n3'. To register it to a fleet in another project, skip this step and use the CLI. [Learn how to register a GKE cluster](#)

[PREVIOUS](#) [NEXT: NETWORKING](#) RESET SETTINGS

- On the next screen we indicate that we are going to use the "default" network and the "default" subnet and in the node subnet we also indicate default.
- We make sure that access is public
- The rest of the options are left by default

Networking

Define how applications in this cluster communicate with each other and reach them.

Network *
default

Node subnet *
default

IPv4 network access

Choose the type of network you want to allow to access your cluster's workloads.

Public cluster

Choose a public cluster to configure access from public networks to your cluster's workloads. Routes aren't created automatically. You cannot change the type of network after the cluster is created.

Private cluster

Choose a private cluster to assign internal IP addresses to Pods and Services. You cannot change the type of network after the cluster is created.

Override control plane's default private endpoint subnet ?

- We press the "NETX-ADVANCED SETTINGS" button.
- In this window we do not modify anything either.

Advanced settings

Target release channel

Regular (recommended) ▼ ?

Versions in the Regular channel have been qualified over a longer period. They offer a balance of feature availability and release stability. We recommend the Regular channel for most users. For known issues and workarounds, review [release notes](#).

cAdvisor and Kubelet Metrics will be enabled by default in GKE Autopilot clusters with the version in the Rapid channel. See more details in the **Advanced settings > Operations** section.

[DISMISS](#)

Automation ▼

Service Mesh ▼

Backup plan ▼

Security ▼

AI and Machine Learning ▼

Operations ▼

Metadata ▼

CREATE
CANCEL
Equivalent REST or

- Click on the CREATE button
- We verify that the cluster has been created

Kubernetes clusters + CREATE + DEPLOY REFRESH

OVERVIEW OBSERVABILITY COST OPTIMIZATION

Filter Enter property name or value

<input type="checkbox"/>	Status	Name ↑	Location	Number of nodes	Total
<input type="checkbox"/>	✓	cluster1	us-central1		
<input type="checkbox"/>	✓	cluster2	us-central1		

- We enter the cluster and click the deploy button.

← Clusters DELETE + DEPLOY CONNECT DUPLICATE

✓ cluster2

DETAILS STORAGE OBSERVABILITY LOGS APP ERRORS (2)

Cluster basics

Name	cluster2
Location type	Regional
Region	us-central1

- Como We write in the name field: "blog1". The rest of the options are left unchanged. We press the button "NEXT CONTAINER DETAILS".

Deployment configuration

A deployment is a configuration which defines how Kubernetes deploys, manages, and scales your container image. Kubernetes will ensure your system matches this configuration. Three replicas will be created by default.

Deployment name *

Namespace *

Labels

Use Kubernetes labels to control how workloads are scheduled to your nodes. Labels are applied to all nodes in this node pool.

Key 1 * Value 1

[+ ADD KUBERNETES LABEL](#)

Cluster

Kubernetes Cluster

Cluster in which the deployment will be created.

- We indicate the image Apasoft/blog.
- As we have already seen in the Cloud Run section, this image allowed you to change the title and color of the main banner.
- We add the following two variables.

Container details

^ **New container**

Existing container image

New container image

Image path * SELECT

Enter your image path, or choose from Google Container Registry. You can also enter the name of any public image from Docker Hub, such as nginx:latest.
[Explore Docker Hub](#)

Environment variables

Key 1 * <input type="text" value="BLOG_SITE_NAME"/>	Value 1 * <input type="text" value="MY FIRST KUBERNETES CLUSTE"/>
Key 2 * <input type="text" value="BLOG_BANNER_COLOR"/>	Value 2 * <input type="text" value="yellow"/>

+ ADD ENVIRONMENT VARIABLE

- Click on the "NEXT EXPOSE" button.
- On the next screen we check the option to create a service

Expose

This operation creates a Kubernetes Service. A service lets your deployment receive traffic and defines how your deployment is exposed. You can always expose your deployment later from the Deployment details page.

Expose deployment as a new service

Port mapping

Item 1 🗑️

Port 1 ?

Target port 1 ?

Protocol 1 ▼ ?

[+ ADD PORT MAPPING](#)

Service type ▼ ?

- Click on the "DEPLOY" button and wait for the process to finish

✔️ **blog1**

OVERVIEW | DETAILS | OBSERVABILITY | REVISION HISTORY | EVENTS | LOGS

CPU ? ≅ 📉 ⋮ Memory ?

0.5

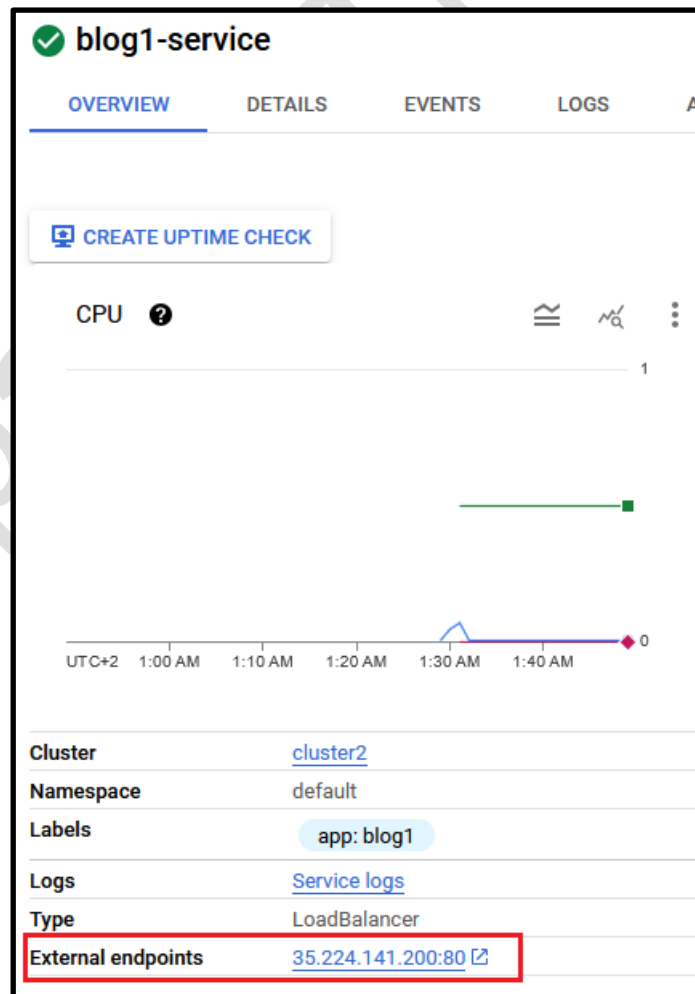
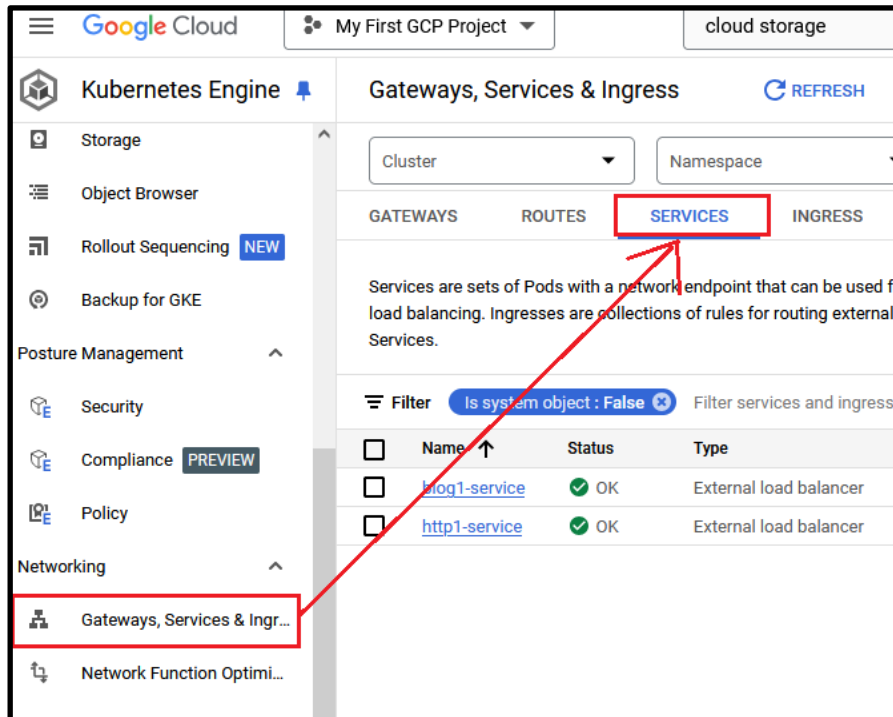
0

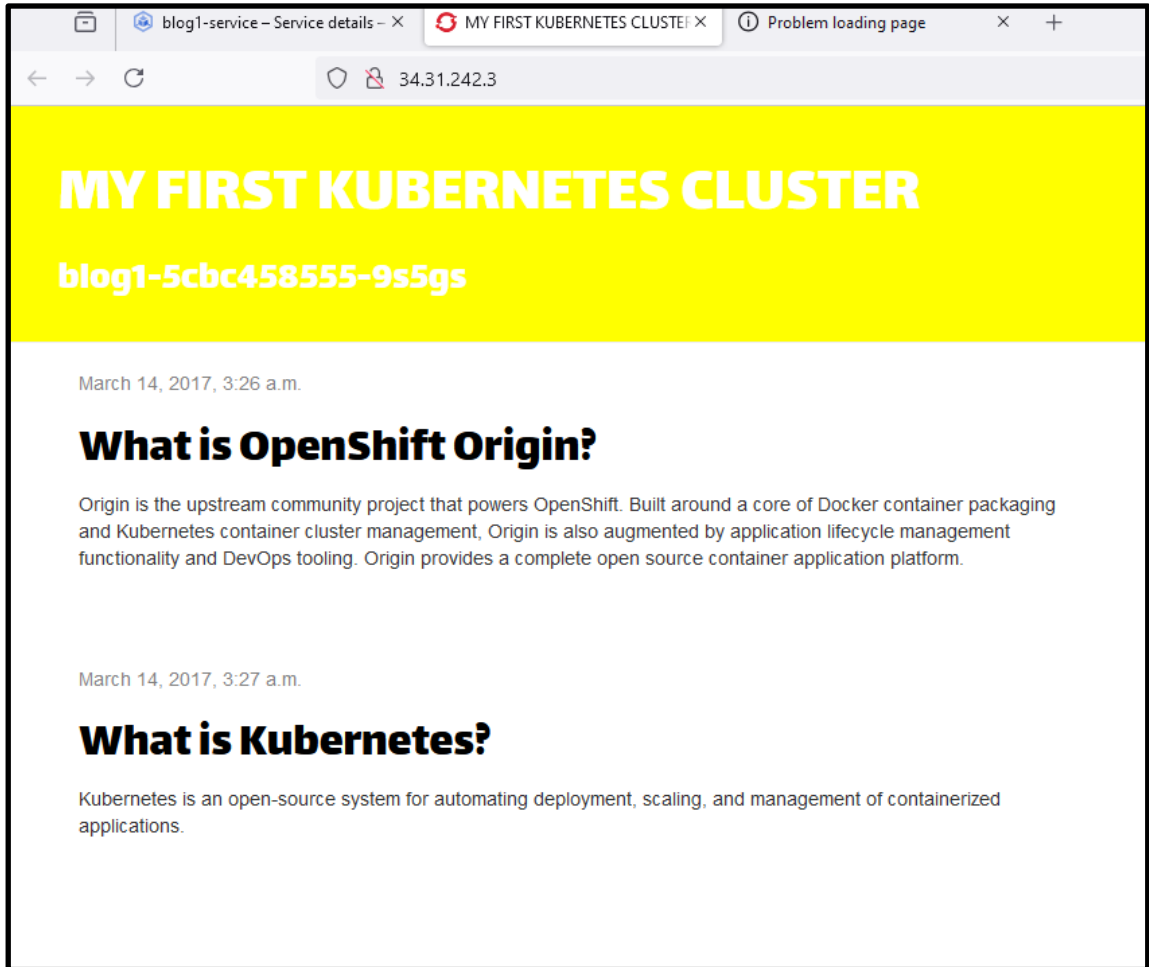
UTC+2 12:50 AM 1:00 AM 1:10 AM 1:20 AM 1:30 AM UTC+2 12:50 AM 1:00 AM 1:10 AM

Cluster [cluster2](#)

-

- Finally, we access the service and test it.





- Delete the cluster

