

# Architecting Hybrid Cloud Infrastructure with Anthos

4jours / 28h

## Learning outcomes

- Connect and manage Anthos GKE clusters from GCP Console whether clusters are part of Anthos on Google Cloud or Anthos deployed on VMware.
- Understand how service mesh proxies are installed, configured and managed.
- Configure centralized logging, monitoring, tracing, and service visualizations wherever the Anthos GKE clusters are hosted.
- Understand and configure fine-grained traffic management.
- Use service mesh security features for service-service authentication, user authentication, and policy-based service authorization.
- Install a multi-service application spanning multiple clusters in a hybrid environment.
- Understand how services communicate across clusters.
- Migrate services between clusters.
- Install Anthos Config Management, use it to enforce policies, and explain how it can be used across multiple clusters.

## Target audience

Technical employees using GCP, including customer companies, partners and system integrators: deployment engineers, cloud architects, cloud administrators, system engineers, and SysOps/DevOps engineers. Individuals using GCP to create, integrate, or modernize solutions using secure, scalable microservices architectures in hybrid

environments.

## Prerequisites

completed the Architecting with Google Kubernetes Engine course and its prerequisites, or have equivalent experience.

## Course Outline

The course includes presentations and hands-on labs.

### Module 1: Anthos Overview

Describe challenges of hybrid cloud

Discuss modern solutions

Describe the Anthos Technology Stack

### Module 2: Managing Hybrid Clusters using Kubernetes Engine

Understand Anthos GKE hybrid environments, with Admin and User clusters

Register and authenticate remote Anthos GKE clusters in GKE Hub

View and manage registered clusters, in cloud and on-premises, using GKE Hub

View workloads in all clusters from GKE Hub

Lab: Managing Hybrid Clusters using Kubernetes Engine

### Module 3: Introduction to Service Mesh

Understand service mesh, and problems it solves

Understand Istio architecture and components

Explain Istio on GKE add on and its lifecycle, vs OSS Istio

Understand request network traffic flow in a service mesh

Create a GKE cluster, with a service mesh

Configure a multi-service application with service mesh

Enable external access using an ingress gateway

Explain the multi-service example applications: Hipster Shop, and Bookinfo

Lab: Installing Open Source Istio on Kubernetes Engine

Lab: Installing the Istio on GKE Add-On with Kubernetes Engine

### Module 4: Observing Services using Service Mesh Adapters

- Understand service mesh flexible adapter model
- Understand service mesh telemetry processing
- Explain Stackdriver configurations for logging and monitoring
- Compare telemetry defaults for cloud and on-premises environments
- Configure and view custom metrics using service mesh
- View cluster and service metrics with pre-configured dashboards
- Trace microservice calls with timing data using service mesh adapters
- Visualize and discover service attributes with service mesh

Lab: Telemetry and Observability with Istio

## Module 5: Managing Traffic Routing with Service Mesh

- Understand the service mesh abstract model for traffic management
- Understand service mesh service discovery and load balancing
- Review and compare traffic management use cases and configurations
- Understand ingress configuration using service mesh
- Visualize traffic routing with live generated requests
- Configure a service mesh gateway to allow access to services from outside the mesh
- Apply virtual services and destination rules for version-specific routing
- Route traffic based on application-layer configuration
- Shift traffic from one service version to another, with fine-grained control, like a canary deployment

Lab: Managing Traffic Routing with Istio and Envoy

## Module 6: Managing Policies and Security with Service Mesh

- Understand authentication and authorization in service mesh
- Explain mTLS flow for service to service communication
- Adopt mutual TLS authentication across the service mesh incrementally
- Enable end-user authentication for the frontend service
- Use service mesh access control policies to secure access to the frontend service

Lab: Managing Policies and Security with Service Mesh

## Module 7: Managing Policies using Anthos Config Management

- Understand the challenge of managing resources across multiple clusters
- Understand how a Git repository is as a configuration source of truth
- Explain the Anthos Config Management components, and object lifecycle
- Install and configure Anthos Config Management, operators, tools, and related Git repository
- Verify cluster configuration compliance and drift management

Update workload configuration using repo changes

Lab: Managing Policies in Kubernetes Engine using Anthos Config

## Module 8: Configuring Anthos GKE for Multi-Cluster Operation

Understand how multiple clusters work together using DNS, root CA, and service discovery

Explain service mesh control-plane architectures for multi-cluster

Configure a multi-service application using service mesh across multiple clusters with multiple control-planes

Configure a multi-service application using service mesh across multiple clusters with a shared control-plane

Configure service naming/discovery between clusters

Review ServiceEntries for cross-cluster service discovery

Migrate workload from a remote cluster to an Anthos GKE cluster

Lab: Configuring GKE for Multi-Cluster Operation with Istio

Lab: Configuring GKE for Shared Control Plane Multi-Cluster Operation