

Manage Scalable Workloads in GKE

Discover how to modernize, manage, and observe applications at scale using Google Kubernetes Engine. This course uses lectures and hands-on labs to help you explore and deploy using Google Kubernetes Engine (GKE), GKE fleets, Cloud Service Mesh, and config controller capabilities that will enable you to work with modern applications, even when they are split among multiple clusters hosted by multiple providers.

**DURATION**

3 days

**LEVEL**

300 - Advanced

**FORMAT**

Instructor-led

What you'll learn

- Describe the challenges of designing a multi-cluster infrastructure.
- Describe the components and architecture of GKE fleets.
- Identify and describe the core components of a GKE fleet.
- Create, connect, and manage GKE clusters from multiple deployment environments.
- Describe how fleets discover and communicate with each other in GKE.
- Detail the benefits of using Service Mesh.
- Use Service Mesh to implement advanced routing and traffic management.
- Secure traffic between microservices using Cloud Service Mesh.
- Create multi-cluster networking architectures with Cloud Service Mesh.
- Use authentication to effectively manage identity in GKE at scale.
- Evaluate and apply various security measures to effectively protect and manage GKE deployments.
- Evaluate options and Google Cloud products that allow you to create scalable CI/CD implementations within a multi-cluster GKE environment.
- Explore how GKE facilitates the deployment and optimization of AI models.

Overview	13 modules · 78 videos · 9 labs · 22 classroom activities
Who this course is for	• Google Cloud practitioners, and individuals using Google Cloud to create, integrate, or modernize solutions using secure, scalable microservices architectures in hybrid environments.
Products	• Cloud Service Mesh • Cloud Deploy • Artifact Analysis • Google Kubernetes Engine • Cloud Build Catalog • Confidential Cloud Run • Security Command Center Enterprise • Config Management • Knative for Anthos • Google Kubernetes Engine (GKE) • GKE fleets • Config Sync • Policy Controller • Config Connector • Cloud Build • Cloud Run • Knative serving • Security Command Center
Prerequisites	• Having completed Google Cloud Platform Fundamentals: Core Infrastructure or having equivalent experience. • Having completed Architecting with GKE or having equivalent experience.

Module 01

Introduction to GKE at scale

Topics	• Multi-cluster overview • GKE fleets • Sameness and trust • GKE fleet management
Objectives	• Recognize the challenges of designing and building multi-environment solutions. • Explain how GKE uses fleets to streamline operations. • Describe the concepts of sameness and trust and use them to manage fleets. • Identify the features and components used to manage GKE fleets.
Activities	• Quiz

Module 02

Multi-cluster GKE architecture

Topics	• Centralized cluster management • Multi-cluster GKE • Connect and manage fleet clusters • Access GKE fleet clusters
Objectives	• Recognize how GKE can be used to centralize cluster management for multi-cluster environments. • Examine the architecture of multi-cluster GKE. • Create, connect, and manage GKE fleet clusters. • Securely access GKE fleet clusters.
Activities	• Quiz

Module 03

Fleets and teams

Topics	• GKE fleets • Example fleet solutions • Fleet team management • Fleet management
Objectives	• Define GKE fleets. • Describe how GKE fleets can solve common cluster management problems. • Manage fleets and teams in GKE. • Detail the elements of fleet management.
Activities	• Quiz • Lab: Manage Workloads at Scale with GKE Fleets and Teams

Module 04

Managing GKE configuration at scale

Topics	• Configuration management challenges • Centralized configuration management at scale • Config Sync • Policy controller • Config Connector • Blueprints
Objectives	• Recognize the challenges of scaling multi-cluster, multi-tenant configurations. • Configure a centralized configuration management using a GitOps model. • Describe the benefits and architecture of Config Sync. • Use policy controller to enforce security and compliance in GKE. • Create a standardized, reusable, and policy-driven foundation for Kubernetes deployments.
Activities	• Quiz • Lab: Automate GKE Configuration with Config Sync

Module 05

Fleet networking

Topics	• Fleet networking communications • Pod discovery in GKE • Multi-cluster Services • Configuring multi-cluster Services • Multi-cluster gateway • Configuring multi-cluster gateways
Objectives	• Explain how fleet networking works. • Describe how Pods in a Kubernetes cluster communicate with each other. • Enable multi-cluster Services. • Configure multi-cluster Services. • Outline the role of a multi-cluster gateway. • Configure a multi-cluster gateway.
Activities	• Quiz • Lab: Deploying a Multi-Cluster Gateway Across GKE Clusters

Module 06

Cloud Service Mesh

Topics	• Introducing Cloud Service Mesh. • Provisioning Cloud Service Mesh. • Handling requests. • Monitoring and supporting Cloud Service Mesh.
Objectives	• List and describe the benefits of using Cloud Service Mesh. • Install and configure Cloud Service Mesh on different clusters. • Trace the path of a request through the mesh, correctly identifying and explaining the role of key components like Envoy proxies, Mesh CA, and extensions in handling the request. • Create Service Mesh dashboards from workload telemetry including metrics, traces, and logs.
Activities	• Quiz • Lab: Installing Cloud Service Mesh on Google Kubernetes Engine

Module
07

Cloud Service Mesh routing

Topics

- Configuring Cloud Service Mesh with Istio API resources.
- Configuring VirtualService and DestinationRule.
- Configuring ServiceEntry.
- Configuring a Gateway.
- Configuring a WorkloadEntry and WorkloadGroup.
- Network resilience and testing.

Objectives

- Explain how Cloud Service Mesh learns the network from Kubernetes.
- Configure traffic behavior using VirtualService and DestinationRule.
- Manage traffic for services running outside the mesh using ServiceEntry.
- Configure Gateways to manage inbound and outbound traffic to the mesh.
- Explore how WorkloadEntry and WorkloadGroup onboards and manages non-Kubernetes workloads.
- Evaluate network resilience and test it through fault injection and traffic mirroring.

Activities

- Quiz
 - Lab: Managing Traffic Flow with Cloud Service Mesh
-

Module 08

Service Mesh security

Topics

- Authentication and encryption
- Service authentication in the mesh
- End-user authentication in Cloud Service Mesh
- Authorization in Cloud Service Mesh

Objectives

- Encrypt traffic between microservices to prevent anyone in the network from gaining access to private information.
- Authorize services and requests, ensuring that services only access the information that is allowed access from other services.
- Authenticate and authorize services and requests to verify trust among services in the mesh and among end users.
- Limit service access in the network so that granular controls over the communication can be established.

Activities

- Quiz
- Lab: Secure Cloud Service Mesh with Policy Controller and mTLS

Module 09

Multi-cluster networking with Cloud Service Mesh

Topics

- Single network east-west routing
- Multiple network east-west routing

Objectives

- Set up a multi-cluster mesh with a single subnet in a single VPC network. Account for variations like multi-region clusters, multiple projects, shared VPC, and private clusters.
- Enable communication between GKE clusters on different networks using an east-west gateway and attached clusters.

Activities

- Quiz
 - Lab: Manage and Secure Distributed Services with GKE Managed Service Mesh
-

Module 10

Manage identity in GKE with authentication

Topics

- Introduction to GKE Identity Service
- Connect gateway overview
- Configuring connect gateway for authentication and authorization
- Accessing clusters with GKE Identity Service
- Authenticating third-party identities with GKE Identity Service
- Fleet Workload Identity

Objectives

- Summarize the differences between authentication methods for GKE clusters and explain when to use each.
- Summarize the key features of connect gateway and explain how it simplifies and secures connections to GKE fleet member clusters.
- Configure connect gateway for authentication and authorization.
- Securely access clusters and provide authentication using OpenID Connect (OIDC) and third-party identity providers (IdPs).
- Given a GKE cluster and a third-party identity provider (IdP), configure GKE Identity Service to enable authentication and authorization for users.

Activities

- Quiz
 - Lab: Manage Authentication at Scale with Connect Gateway
-

Module 11 Security posture, compliance, and preventative controls

Topics	• GKE security posture overview • Security posture dashboard • Implementing node security • Vulnerability scanning • Additional security services
Objectives	• Describe GKE security posture. • Navigate and interpret the GKE security posture dashboard to identify security issues. • Implement node security measures to protect GKE worker nodes from potential threats. • Describe the process of vulnerability scanning in GKE. • Explain the roles and capabilities of Google Cloud's Artifact Analysis and Security Command Center in enhancing GKE security.
Activities	• Quiz

Module
12

CI/CD at scale in GKE

Topics

- CI/CD in Google Cloud
- Cloud Build and GKE
- Cloud Deploy and GKE
- Cloud Deploy: Policies, deployments, and security
- Cloud Run and Knative serving
- Cloud Deploy and Knative serving
- CI/CD in a private network
- Software supply chain security

Objectives

- Describe the core components of Google Cloud's CI/CD pipeline and how they address common challenges in application modernization.
- Analyze how Cloud Deploy integrates with GKE to manage Kubernetes manifests and control deployments.
- Compare and contrast the deployment strategies for Knative serving within GKE.
- Explain the steps required to establish a peered VPC connection for secure CI/CD in a private network.
- Evaluate the various security measures and tools available within Google Cloud for securing the software supply chain.

Activities

- Quiz
 - Lab: Creating CI/CD Pipelines for GKE Clusters
-

Module 13

GKE and AI

Topics

- AI and GKE overview
- AI model training on GKE
- AI model serving on GKE
- AI cost management on GKE

Objectives

- Explain how GKE serves as a suitable platform for large language models and the increasing demand for hardware accelerators.
- Describe the high-level architecture of a GKE-based training platform for AI models.
- Outline the architecture for a GKE-based model serving platform.
- Outline different cost management strategies available when using GKE for AI/ML workloads.

Activities

- Quiz