

Google Kubernetes Engine Administrator

Course Duration: 32 Hours

Course code: GKEA

1. Course Overview

This course is designed to provide in-depth knowledge and hands-on experience in deploying, managing, and troubleshooting containerized applications using Google Kubernetes Engine (GKE). Learners will gain expertise in Kubernetes architecture, cluster management, networking, security, and scaling applications in a production environment on Google Cloud.

2. What you'll learn?

By the end of the course, you will be able to:

- Understand Kubernetes architecture and core components
- Deploy and manage GKE clusters
- Configure workloads, services, and networking in Kubernetes
- Implement security best practices in GKE
- Monitor, log, and troubleshoot Kubernetes environments
- Manage scaling and high availability of applications
- Automate deployments using CI/CD pipelines
- Optimize performance and cost of Kubernetes workloads

3. Target Audience

- Kubernetes Administrators
- Cloud Engineers
- DevOps Engineers
- System Administrators
- Application Developers

4. Pre-Requisites

Before taking this course, you should have:

- Basic knowledge of containers (Docker)
- Understanding of Kubernetes fundamentals
- Familiarity with Google Cloud Platform
- Basic command-line and networking knowledge

5. Course content

Module 1: Course Introduction

- Course objectives and structure
- Overview of containerization and Kubernetes
- Introduction to Google Kubernetes Engine

Module 2: Kubernetes Architecture Fundamentals

- Kubernetes control plane components
- Nodes and node pools
- Pods, ReplicaSets, and Deployments
- Cluster architecture overview

Module 3: GKE Cluster Deployment and Management

- Creating GKE clusters (Standard & Autopilot)
- Configuring node pools
- Cluster upgrades and maintenance
- Managing cluster lifecycle

Module 4: Workloads and Scheduling

- Deploying applications using Deployments
- StatefulSets and DaemonSets
- Job and CronJob management
- Resource requests and limits

- Pod scheduling and affinity rules

Module 5: Services and Networking in GKE

- Kubernetes Services (ClusterIP, NodePort, LoadBalancer)
- Ingress and HTTP(S) load balancing
- Service discovery and DNS
- Network policies and traffic control

Module 6: Storage and Persistent Volumes

- Persistent Volumes (PV) and Persistent Volume Claims (PVC)
- Storage classes in GKE
- Stateful application storage
- Backup and restore strategies

Module 7: Security in GKE

- Identity and Access Management (IAM)
- Role-Based Access Control (RBAC)
- Workload Identity
- Securing clusters and workloads
- Secrets and configuration management

Module 8: Monitoring, Logging, and Troubleshooting

- Google Cloud Operations Suite
- Monitoring cluster and application health
- Logging and debugging issues
- Troubleshooting common Kubernetes problems

Module 9: Scaling and High Availability

- Horizontal Pod Autoscaler (HPA)
- Vertical Pod Autoscaler (VPA)

- Cluster autoscaling
- Designing highly available applications

Module 10: CI/CD for Kubernetes

- Introduction to CI/CD pipelines
- Integrating with Cloud Build
- Continuous deployment strategies
- Managing application releases

Module 11: Configuration Management and Automation

- Using ConfigMaps and Secrets
- Infrastructure as Code (Terraform)
- Automating deployments
- GitOps practices

Module 12: Cost Optimization and Performance Tuning

- Resource optimization techniques
- Managing cluster costs
- Performance tuning strategies
- Efficient workload management

Module 13: Advanced GKE Features

- GKE Autopilot vs Standard mode
- Multi-cluster management
- Anthos integration
- Advanced networking features

Module 14: Hands-On Labs and Real-World Scenarios

- Deploying containerized applications
- Configuring networking and storage

- Implementing security policies
- Troubleshooting live scenarios

Module 15: Capstone Project

- Deploy a production-grade application on GKE
- Implement scaling, security, and monitoring
- Optimize performance and cost
- Final project presentation

