

Course Outline

Google Cloud Fundamentals: Core Infrastructure

Google Cloud Platform Fundamentals: Core Infrastructure covers the fundamentals of Google Cloud Platform and is the first course in the 3 specializations. This course introduces you to important concepts and terminology for working with Google Cloud Platform (GCP). You learn about, and compare, many of the computing and storage services available in Google Cloud Platform, including Google App Engine, Google Compute Engine, and Google Container Engine.

Introducing Google Cloud Platform

- Google Cloud Platform offers four main kinds of services: Compute, Storage, Big Data, and Machine Learning. This course focuses mostly on the first two, together with Google Virtual Private Cloud (VPC) networking. This module orients learners to the basics of Google Cloud Platform. It traces the evolution of cloud computing and explains what is unique about Google's approach to it. The module introduces the key structural concepts of regions and zones.

Getting Started with Google Cloud Platform

- GCP customers use projects to organize the resources they use. They use Google Cloud Identity and Access Management, also called "IAM," to control who can do what with those resources. They use any of several technologies to connect with GCP. This module covers each of these topics, and it introduces a service called Cloud Launcher that is an easy way to get started with GCP.

Virtual Machines in the Cloud

- Compute Engine lets you run virtual machines on Google's global infrastructure. This module covers how Compute Engine works, with a focus on Google virtual networking.

Storage in the Cloud

- Every application needs to store data. Different applications and workloads require different storage and database solutions. This module describes and differentiates among GCP's core storage options: Cloud Storage, Cloud SQL, Cloud Spanner, Cloud Datastore, and Google Bigtable.

Containers in the Cloud

- Containers are simple and interoperable, and they enable seamless, fine-grained scaling. Kubernetes is an orchestration layer for containers. Kubernetes Engine is Kubernetes as a service, a scalable managed offering that runs on Google's infrastructure. You direct the creation of a cluster,

and Kubernetes Engine schedules your containers into the cluster and manages them automatically, based on requirements you define. This module explains how Kubernetes Engine works and how it helps deploy applications in containers.

Applications in the Cloud

- App Engine is a Platform-as-a-Service ("PaaS") offering. The App Engine platform manages the hardware and networking infrastructure required to run your code. App Engine provides built-in services that many web applications need. This module describes how App Engine works.

Developing, Deploying and Monitoring in the Cloud

- Popular tools for development, deployment, and monitoring just work in GCP. Customers also have options for tools in each of these three areas that are tightly integrated with GCP. This module covers those tools.

Big Data and Machine Learning in the Cloud

- GCP's big-data and machine learning offerings are intended to help customers get the most out of data. These tools are intended to be simple and practical to embed in your applications. This module describes the available big-data and machine learning services and explains the usefulness of each.

Summary and Review

- This module reviews the GCP services covered in this course and reminds learners of the differences among them. The module compares GCP compute services, GCP storage services, and important Google VPC networking capabilities.