

Google Cloud Associate Cloud Engineer Certification Guide

*A comprehensive guide to implementing, managing,
and monitoring Google Cloud solutions*

Arijit Sarkar



www.bpbonline.com

First Edition 2025

Copyright © BPB Publications, India

ISBN: 978-93-65894-141

All Rights Reserved. No part of this publication may be reproduced, distributed or transmitted in any form or by any means or stored in a database or retrieval system, without the prior written permission of the publisher with the exception to the program listings which may be entered, stored and executed in a computer system, but they can not be reproduced by the means of publication, photocopy, recording, or by any electronic and mechanical means.

LIMITS OF LIABILITY AND DISCLAIMER OF WARRANTY

The information contained in this book is true and correct to the best of author's and publisher's knowledge. The author has made every effort to ensure the accuracy of these publications, but the publisher cannot be held responsible for any loss or damage arising from any information in this book.

All trademarks referred to in the book are acknowledged as properties of their respective owners but BPB Publications cannot guarantee the accuracy of this information.

To View Complete
BPB Publications Catalogue
Scan the QR Code:



Dedicated to

My wife, Sudeshna Sarkar

About the Author

Arijit Sarkar is a Data & AI enthusiast. He started his career as an ETL Developer and is currently working as a Lead Architect - Google Cloud in Data & AI consulting practice. He has been using Google Cloud since 2019 and has completed four certifications on this. He always likes to explore new technologies and platforms and implement those in real-world use cases.

About the Reviewers

- ❖ **Avik Sarkar** is a seasoned Cloud Solution Architect with extensive expertise in Google Cloud Platform (GCP) and multi-cloud environments. With over 20 years of experience in IT, he specializes in cloud migrations, enterprise application design, and digital modernization; he has led large-scale enterprise solutions across AWS, Azure, and GCP. His expertise spans multiple domains, including banking, insurance, manufacturing, and logistics, with extensive experience working in India and Switzerland.

Avik plays a pivotal role in designing and implementing solutions in financial crime detection, insurance modernization, and digital transformation. His technical acumen includes architecting high-availability, scalable, and secure cloud solutions, leveraging technologies such as Kubernetes, Kafka, Redis, and AI-driven analytics.

Avik's professional journey includes impactful tenures at Cognizant and IBM, where he contributed to cloud-based data lake solutions, enterprise data ingestion frameworks, and banking infrastructure modernization. His deep understanding of agile methodologies, project management, and security frameworks enables him to drive strategic initiatives and innovation.

Holding multiple industry-recognized certifications—including Microsoft Azure Security Engineer Associate, TOGAF 9, and Google Associate Cloud Engineer—Avik continues to push the boundaries of cloud technology. With a passion for optimizing performance and security in large-scale cloud environments, he remains a thought leader in enterprise cloud solutions.

Beyond his technical pursuits, Avik is multilingual, fluent in English, Hindi, and Bengali, and actively shares knowledge within the tech community.

- ❖ **Vibhor Malik** is a backend engineer at one of the Fortune 500 companies with a strong interest in technical literature and artificial intelligence. He holds a GCP Associate Developer certification and a Master's in Machine Learning and AI from Liverpool John Moores University, combining academic knowledge with industry experience. In his reviews, he focuses on exam-related content and certification insights to help readers better understand complex topics. Outside of work, he enjoys coding and cooking—because, as he believes, if you are good at cooking, you are probably a great coder too!

- ❖ **Karanbir Singh** is a seasoned Engineering Leader with over 8 years of experience in AI/ML, distributed systems, and backend engineering. Currently a Senior Software Engineer at Salesforce, he specializes in backend technologies and AI applications. He is a recognized thought leader, having spoken at major conferences like AAAI, ICLR, the WebConference, and Google Developer Groups, and has been featured on the Data Neighbours podcast alongside AI experts like Josh Starmer (StatQuest).

Previously, as Engineering Manager at TrueML, he led teams deploying ML models in production, boosting development velocity, client engagement, and operational efficiency. His leadership directly influenced revenue growth and cost optimization. He also held impactful engineering roles at Lucid Motors and Poynt, contributing to product innovation and system integration.

Karanbir holds a Master's in Computer Software Engineering from San Jose State University and is a Silicon Valley Innovation Challenge winner. Passionate about mentoring, he continues to contribute to the global tech community through public speaking and knowledge sharing.

Acknowledgement

There are a few people I would really want to thank for the continued and ongoing support they have given me during the writing of this book. First and foremost, I would like to thank my wife, my brother and my colleagues for continuously encouraging me to write this book.

I am grateful to the courses, tutorials, blogs, and the companies which gave me support throughout the learning process of Google Cloud. Thank you for all the support you provided.

My sincere gratitude also goes to the team at BPB Publications for being supportive throughout the journey.

Preface

Mainly due to the exponential growth of data, organizations need to scale up their IT infrastructure to maximize their business outcome. With the emergence of cloud computing, organizations are moving their workloads to public cloud in order to reduce the capital expenditure towards on-premises IT infrastructure and Google Cloud Platform is one of the major public cloud platforms present today.

This book will guide you to prepare for Google Cloud Associate Cloud Engineer certification with detailed explanation and pictorial demonstrations of GCP services. It contains 16 Chapters and 2 Practice Tests.

Chapter 1: Cloud Computing with GCP - This will introduce the readers to cloud computing with GCP, why cloud computing is needed, the drawbacks of on-premises infrastructure, what are public, private and hybrid cloud, and what are the major public cloud providers available in the market. It will also include Google Cloud Platform's history and list of services GCP provides. It will set the foundation to learn GCP.

Chapter 2: Overview of Google Cloud Associate Cloud Engineer Certification - This will include the details about the Google Cloud Associate Cloud Engineer Certification, how much time, how many questions, questions covering the parts of services etc.

Chapter 3: Understanding GCP Console, Cloud Shell, Billing and APIs - This will introduce the readers to the GCP Console, it will make them familiar with different portions of the console and show how to use those effectively with pictorial examples. It will also show how to enable API and Billing and finally how to use Cloud SDK from the command prompt/ or Cloud Shell.

Chapter 4: Identity and Access Management in GCP - This chapter will include the Identity and Access Management in GCP. It will introduce IAM roles, service accounts, different types of access control mechanisms, IAM best practices etc.

Chapter 5: Storage Solution in GCP (Block, Network File and Objects) - This chapter will discuss Google Cloud Object Storage service, public and private access, object life cycle and versioning. Access using console and Python Programming language.

Chapter 6: Understanding Different Databases in GCP - This chapter will introduce the SQL and NoSQL databases in GCP. It will also help you to identify the correct database choice for different types of requirements.

Chapter 7: Google Compute Engine - This chapter introduces the readers to the first compute services of GCP. Here we will learn about configuring VMs and understand the managed instance group.

Chapter 8: Cloud Run - This chapter introduces the service to build and deploy scalable containerized apps written in any language (including Go, Python, Java, Node.js, .NET, and Ruby) on a fully managed platform.

Chapter 9: Google Kubernetes Engine - This chapter will include the Kubernetes service inside GCP. We will learn about creating and managing the Kubernetes cluster and deploying services into it.

Chapter 10: Cloud Functions - We will study App Engine which is a fully managed, serverless platform for developing and hosting applications at scale.

Chapter 11: App Engine - This chapter will introduce a new concept called Functions as a Service in GCP with Cloud Functions. It is a serverless scalable service that offers event driven processing.

Chapter 12: Networking in GCP – Networking is a very important topic in Computer Science. It deals with communication between two systems. In Google Cloud Platform, Networking takes place via VPC, i.e., Virtual Private Cloud. It is the backbone of various services in GCP to stay connected among each other and communicate internally or outside of GCP.

Chapter 13: Networking in GCP Firewall Rules, Load Balancing, DNS, CDN and NAT– In this part 2 of Networking in GCP, here we are going to learn about Firewall Rules and load balancers and will see how these help in movement of network traffic in Google Cloud Platform along with Cloud DNS, CDN and NAT.

Chapter 14: Big Data Processing, AI, Deployment and Monitoring in GCP - This chapter covers the basic details around some of the services related to Big Data processing, artificial intelligence, deployment and monitoring in GCP. In various industries, Google Cloud Platform is the preferred choice for Bigdata and artificial intelligence. We are going to cover the basic understanding of these services.

Chapter 15: End-to-End Application Lifecycle in GCP Design, Build, Test, Deployment and Monitoring - This chapter covers the End-to-End Application lifecycle in Google Cloud Platform. We will start with requirement gathering and finish with monitoring of an application.

Chapter 16: Specific Topics for GCP ACE Exam - This chapter covers some of the specific topics required for GCP ACE certification exam.

Chapter 17: Practice Test 1- 20 Practice Questions on Certification exam.

Chapter 18: Practice Test 2- 20 Practice Questions on Certification exam.

Code Bundle and Coloured Images

Please follow the link to download the
Code Bundle and the *Coloured Images* of the book:

<https://rebrand.ly/wiy52el>

The code bundle for the book is also hosted on GitHub at

<https://github.com/bpbpublications/Google-Cloud-Associate-Cloud-Engineer-Certification-Guide>.

In case there's an update to the code, it will be updated on the existing GitHub repository.

We have code bundles from our rich catalogue of books and videos available at
<https://github.com/bpbpublications>. Check them out!

Errata

We take immense pride in our work at BPB Publications and follow best practices to ensure the accuracy of our content to provide with an indulging reading experience to our subscribers. Our readers are our mirrors, and we use their inputs to reflect and improve upon human errors, if any, that may have occurred during the publishing processes involved. To let us maintain the quality and help us reach out to any readers who might be having difficulties due to any unforeseen errors, please write to us at :

errata@bpbonline.com

Your support, suggestions and feedbacks are highly appreciated by the BPB Publications' Family.

Did you know that BPB offers eBook versions of every book published, with PDF and ePub files available? You can upgrade to the eBook version at www.bpbonline.com and as a print book customer, you are entitled to a discount on the eBook copy. Get in touch with us at :

business@bpbonline.com for more details.

At **www.bpbonline.com**, you can also read a collection of free technical articles, sign up for a range of free newsletters, and receive exclusive discounts and offers on BPB books and eBooks.

Piracy

If you come across any illegal copies of our works in any form on the internet, we would be grateful if you would provide us with the location address or website name. Please contact us at **business@bpbonline.com** with a link to the material.

If you are interested in becoming an author

If there is a topic that you have expertise in, and you are interested in either writing or contributing to a book, please visit **www.bpbonline.com**. We have worked with thousands of developers and tech professionals, just like you, to help them share their insights with the global tech community. You can make a general application, apply for a specific hot topic that we are recruiting an author for, or submit your own idea.

Reviews

Please leave a review. Once you have read and used this book, why not leave a review on the site that you purchased it from? Potential readers can then see and use your unbiased opinion to make purchase decisions. We at BPB can understand what you think about our products, and our authors can see your feedback on their book. Thank you!

For more information about BPB, please visit **www.bpbonline.com**.

Join our book's Discord space

Join the book's Discord Workspace for Latest updates, Offers, Tech happenings around the world, New Release and Sessions with the Authors:

<https://discord.bpbonline.com>



Table of Contents

1. Cloud Computing with GCP	1
Introduction	1
Structure.....	1
Objectives	2
Cloud computing	2
On premise infrastructure and its drawbacks.....	2
Public, private, and hybrid cloud	2
Major public cloud providers	3
IaaS, PaaS, and SaaS.....	3
Introduction to GCP.....	3
Google Cloud Platform resource hierarchy.....	3
Brief introduction about different GCP services.....	4
<i>Storage related services</i>	<i>4</i>
<i>Compute related services</i>	<i>5</i>
<i>Database related services</i>	<i>5</i>
<i>Network services</i>	<i>6</i>
<i>Data analytics services</i>	<i>7</i>
<i>Security and identity services</i>	<i>8</i>
<i>AI and machine learning services</i>	<i>9</i>
<i>Operations related services.....</i>	<i>10</i>
<i>API related services</i>	<i>10</i>
<i>Developer tools</i>	<i>10</i>
Conclusion.....	11
2. Overview of Google Cloud Associate Cloud Engineer Certification	13
Introduction.....	13
Structure.....	13
Objectives	14
Cloud engineer	14
Exam details	14
Target audience	14
Mode of exam	14

Syllabus.....	15
<i>Setting up cloud environment.....</i>	15
<i>Planning and configuring a cloud based solution.....</i>	15
<i>Deploying and implementing a cloud based solution</i>	16
<i>Ensuring successful operation of a cloud based solution</i>	17
<i>Configuring access and security</i>	18
How to set up an exam.....	19
Preparation strategy.....	21
Conclusion.....	22
Sample questions.....	22
<i>Multiple choice single select questions.....</i>	22
<i>Answers.....</i>	22
<i>Multiple choice multi select questions</i>	23
<i>Answers.....</i>	23
3. Understanding GCP Console, Cloud Shell, Billing, and APIs	25
Introduction.....	25
Structure.....	25
Objectives	26
Creating a GCP project	26
Understanding GCP console	29
Cloud Shell	31
gcloud Command Line Interface	32
Regions and zones.....	32
Cloud Billing	33
Cloud APIs	35
<i>Example usage.....</i>	36
<i>Installation</i>	36
Cloud SDK.....	36
<i>Installation in MacOS</i>	37
<i>Installation in Windows.....</i>	37
<i>Initializing the gcloud CLI.....</i>	37
<i>Authorize the gcloud CLI.....</i>	37
Conclusion.....	38
Multiple choice questions	38
<i>Answers.....</i>	38

4. Identity and Access Management in GCP	39
Introduction.....	39
Structure.....	39
Objectives	40
Principals	40
<i>Google account</i>	40
<i>Service account</i>	40
<i>Google group</i>	40
<i>Google workspace account</i>	40
<i>Cloud Identity domain</i>	40
<i>All authenticated users</i>	41
<i>All users</i>	41
Resources, permissions, roles, and policies	41
<i>Resources</i>	41
<i>Permissions</i>	41
<i>Roles</i>	41
<i>Policies</i>	42
Inside IAM in GCP console.....	42
Service accounts.....	43
Custom role	47
IAM best practices.....	51
Conclusion.....	51
Multiples choice questions.....	51
<i>Answers</i>	52
 5. Storage Solutions in GCP (Block, Network File and Object).....	53
Introduction.....	53
Structure.....	53
Objectives	54
Block storage: Local SSDs and persistent disks	54
<i>Local SSDs</i>	54
<i>Persistent disks</i>	54
<i>Comparison between Local SSDs and persistent disks</i>	54
Network file storage: Cloud Filestore	55
Comparison between block and file storage	55

Object storage: Google Cloud Storage.....	56
<i>Tools for Google Cloud Storage</i>	56
Working with Google Cloud Storage	57
<i>Access Google Cloud Storage using Python</i>	65
<i>Access Google Cloud Storage using gsutil commands</i>	65
<i>Cost calculation</i>	66
<i>Data Encryption in Cloud Storage</i>	67
Conclusion.....	67
Multiple choice questions	67
<i>Answers</i>	68
6. Understanding Different Databases in GCP	69
Introduction.....	69
Structure.....	69
Objectives	70
SQL versus NoSQL databases	70
<i>Comparative study between SQL and NoSQL databases</i>	70
SQL databases in GCP	71
Cloud SQL.....	71
<i>Use cases for Cloud SQL</i>	72
<i>Key features</i>	72
AlloyDB for PostgreSQL	74
<i>Use cases for AlloyDB for PostgreSQL</i>	74
<i>Key features</i>	74
Cloud Spanner.....	75
<i>Use cases for Cloud Spanner</i>	75
<i>Key features</i>	75
NoSQL databases in GCP	76
Firestore.....	76
<i>Use cases for Firestore</i>	77
<i>Key features</i>	77
Cloud Bigtable	78
<i>Use cases for Cloud Bigtable</i>	78
<i>Key features</i>	78
In-memory database in GCP.....	79
<i>Use cases for Memorystore</i>	79

<i>Key features</i>	79
Database for analytics.....	80
<i>Use cases for BigQuery</i>	80
<i>Key features</i>	80
Conclusion.....	81
Multiple choice questions	81
<i>Answers</i>	82
7. Google Compute Engine.....	83
Introduction.....	83
Structure.....	83
Objectives	84
Understanding Google Compute Engine	84
<i>Key features</i>	84
Creation of Compute Engine VM instance.....	85
<i>Using GCP console</i>	85
<i>Using command line</i>	90
Custom SSH Keys.....	92
Managed instance groups	93
<i>Scenarios</i>	93
<i>Benefits</i>	93
Working in a Compute Engine VM instance.....	94
Conclusion.....	95
Multiple choice questions	95
<i>Answers</i>	96
8. Cloud Run.....	97
Introduction.....	97
Structure.....	97
Objectives	98
Understanding platform as a service	98
Difference between PaaS and IaaS.....	98
Introduction to Cloud Run.....	99
<i>Ways of using Cloud Run</i>	99
<i>Features of Cloud Run</i>	100
Introduction to containerized applications	100

Container or artifact registry	101
<i>Features</i>	101
Creating a REST API using Cloud Run.....	101
Conclusion.....	115
Multiple choice questions	115
<i>Answers</i>	115
9. Google Kubernetes Engine	117
Introduction.....	117
Structure.....	117
Objectives	118
Container orchestration.....	118
<i>Functioning of container orchestration</i>	118
Understanding Kubernetes.....	118
<i>Features of Kubernetes</i>	119
<i>Kubernetes terminologies</i>	119
Google Kubernetes Engine.....	121
<i>Cloud Run vs Google Kubernetes Engine</i>	121
<i>Cloud Run for Anthos</i>	122
Working with Google Kubernetes Engine	122
<i>gcloud commands</i>	142
<i>kubectrl commands</i>	142
Conclusion.....	143
Multiple choice questions	143
<i>Answers</i>	144
10. Cloud Functions	145
Introduction.....	145
Structure.....	145
Objectives	146
Functions as a service	146
Function as a service vs serverless applications	146
Cloud Functions	146
<i>Some of the key features of Cloud Functions</i>	147
Cloud Functions triggers.....	147
<i>Triggers supported in 1st generation</i>	147

<i>Triggers supported in 2nd generation</i>	147
Cloud Functions configurations	148
Working with Cloud Functions	149
Conclusion	159
Multiple choice questions	159
<i>Answers</i>	160
11. App Engine	161
Introduction	161
Structure	161
Objectives	162
App Engine	162
Types of App Engine	162
<i>Comparison between standard and flexible types</i>	163
App Engine components	164
Creating an application using App Engine	164
Traffic splitting between versions	170
Multiple services inside one application	173
Request routing to applications, services, and versions	175
Conclusion	175
Multiple choice questions	175
<i>Answers</i>	176
12. Networking in GCP	177
Introduction	177
Structure	177
Objectives	178
Computer networking	178
<i>Working of computer network</i>	178
<i>Basic terminologies of computer networks</i>	178
Classless Inter-Domain Routing	179
<i>Working of CIDR</i>	179
VPC and subnets	179
<i>Custom mode and auto mode VPC</i>	180
<i>Creating resources using VPC and subnets</i>	180
<i>Working VPC and subnets</i>	180

Shared VPC and VPC peering	188
Hybrid connectivity	188
<i>Cloud interconnect</i>	188
<i>Cloud VPN</i>	189
Conclusion.....	189
Multiple choice questions	189
<i>Answers</i>	190
13. Networking in GCP Firewall Rules, Load Balancing, DNS, CDN and NAT	191
Introduction.....	191
Structure.....	191
Objectives	191
Firewall rules.....	192
<i>Best practices for firewall rules</i>	192
Cloud load balancing.....	197
<i>Networking protocols</i>	197
<i>Working with load balancer in GCP</i>	197
Cloud Domain Name System	204
Cloud Content Delivery Network	204
Cloud Network Address Translation	204
Conclusion.....	204
Multiple choice questions	204
<i>Answers</i>	205
14. Big Data Processing, AI, Deployment and Monitoring in GCP	207
Introduction.....	207
Structure.....	207
Objectives	208
Pub/Sub.....	208
<i>Components of Pub/Sub</i>	208
<i>Workflow of the Pub/Sub service</i>	209
Big Data.....	209
Dataflow, Dataproc, and composer.....	210
<i>Dataflow</i>	210
<i>Dataproc</i>	211
<i>Composer</i>	212

<i>Workflows, DAGs, and tasks</i>	212
<i>Components of cloud composer</i>	212
Artificial intelligence and machine learning	212
Vertex AI	213
<i>Machine learning workflow</i>	214
Cloud Build and Deployment Manager	214
<i>Cloud Build</i>	214
<i>Deployment Manager</i>	215
Cloud Monitoring and Cloud Logging	215
<i>Cloud Monitoring</i>	215
<i>Cloud Logging</i>	216
Conclusion	216
Multiple choice questions	216
<i>Answers</i>	217
15. End-to-End Application Lifecycle in GCP Design, Build, Test, Deployment and	
Monitoring	219
Introduction	219
Structure	219
Objectives	219
Requirement	220
Design and architecture	220
Building, unit testing, and deployment	220
Functionality testing	236
Monitoring	237
Conclusion	239
16. Specific Topics for GCP ACE Exam	241
Introduction	241
Structure	241
Objectives	242
GCP database product choice	242
Installing and configuring Ops Agent on GCP VM	243
Adding a GPU to a new instance	244
Compute engine VM snapshots and images	245
Cloud Run or Cloud Run for Anthos	248

Horizontal Pod Autoscaling for GKE.....	248
Vertical Pod Autoscaling for GKE.....	248
Cloud Foundation Toolkit.....	248
Connecting On-Prem and GCP using Cloud VPN.....	249
Managing users and groups in Cloud Identity.....	250
Provisioning and setting up products in Google Cloud’s operations suite.....	251
Assessing compute quotas and requesting increases	254
Creating and managing short-lived service account credentials	255
Conclusion.....	256
17. Practice Test 1	257
Practice questions.....	257
<i>Answers</i>	262
18. Practice Test 2	263
Practice questions.....	263
<i>Answers</i>	269
Index	271-275

CHAPTER 1

Cloud Computing with GCP

Introduction

This chapter will cover the introduction of cloud computing with **Google Cloud Platform (GCP)**. We will understand why cloud computing is needed, the drawbacks of on-premises infrastructure, the meaning of public, private, and hybrid cloud infrastructure, and the major public cloud providers available in the market. It will also include GCP's history and the types of services it provides. We will also learn a brief description of most of the services in GCP. This chapter will set the foundation for learning the GCP.

Structure

The chapter covers the following topics:

- Cloud computing
- On-premise infrastructure and its drawbacks
- Public, private, and hybrid cloud
- Major public cloud providers
- IaaS, PaaS, and SaaS
- Introduction to GCP

- Google Cloud Platform resource hierarchy
- Brief introduction about different GCP services

Objectives

By the end of this chapter, you will be able to understand the cloud computing concepts, the difference between cloud computing and on-premise-based computing, the drawbacks of on-premise-based infrastructure, the difference between **Infrastructure as a Service (IaaS)**, **Platform as a Service (PaaS)**, **Software as a Service (SaaS)**, a brief about various GCP services solving challenges around storage, computing, artificial intelligence, etc.

Cloud computing

Cloud computing is the on-demand availability of computing resources like storage, infrastructure, etc., without actually owning any of these. It is very much like renting the services which are required. In this way, the capital expenditure towards hardware can be minimized, and users can focus more on the application than on managing and monitoring the hardware infrastructure.

On premise infrastructure and its drawbacks

On-premise infrastructure means owning the computing resources. It requires upfront investment towards hardware, which leads to higher capital expenditure before even focusing on the actual product. Secondly, managing the infrastructure and monitoring these needs specialized resources, which again increases IT expenditure. Upgradation and licensing of this hardware again fall under the IT department of the organization, along with the need for scalability with the exponential growth of data and digital initiatives.

Public, private, and hybrid cloud

In cloud computing, when we are using multi-tenant platforms (hardware) for storage, infrastructure, etc., then it is called a public cloud. To be precise, multiple organizations use the same shared platform (hardware).

When the usage of the cloud service provider's hardware is single-tenant, then it is a private cloud; that is, only a single organization is using the hardware that is rented, and no other organizations are using it.

Sometimes, organizations need to keep their data movement within on-premises hardware due to security restrictions or the governing body's policy. At that time, both on-premises and public cloud are used to deliver a particular business application or solution. This type of arrangement is called a hybrid cloud.

Major public cloud providers

There are many cloud service providers at this moment, but the majority of the market share is with AWS from Amazon, Azure from Microsoft, and GCP from Google. Take a look at the following figure:



Figure 1.1: Major public cloud providers

IaaS, PaaS, and SaaS

Starting with IaaS. It is a cloud computing model where you rent the infrastructure. The infrastructure can be **virtual machine (VMs)**, storage, network, etc. The service provider takes care of hosting and maintaining the infrastructure, and you need to take care of the operating system, data, and applications.

When it comes to PaaS, Cloud service providers will provide a platform, and they will manage the backend infrastructure completely. You do not have to worry about the infrastructure. You just need to write the code for your application and manage the data and applications.

SaaS is a software distribution model that connects its customers mostly by internet. It normally comes with a subscription-based software services. Examples of SaaS are very common to us, like Dropbox, Workplace, Netflix, YouTube, Salesforce etc.

Introduction to GCP

GCP is a cloud computing platform offered by Google. It started its journey in April 2008 with **App Engine** as its first service which was generally available in November 2011.

GCP offers various services, which can be categorized as storage, computing, database, networking, data analytics, security, AI, monitoring and logging, API, etc. It has 11% of the market share at the start of 2023, and it is growing daily.

Google Cloud Platform resource hierarchy

For GCP resources, the hierarchy is **Organization | Folders | Projects | Services**. The organization is like a company, say *abc.com*. Folders are specific to the area of services the company provides or the actual products of the company or features. Under a folder, there are different projects, mostly for various environments. These can be defined as per the organization's needs and choices. Take a look at the following figure:

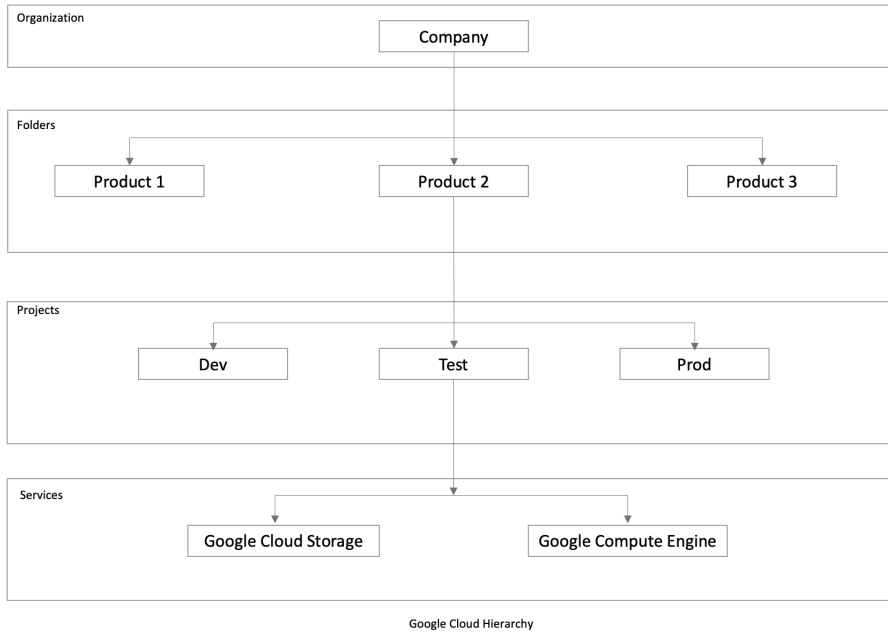


Figure 1.2: GCP resource hierarchy

Brief introduction about different GCP services

Let us understand each service one by one. For ease of understanding, we are going to categorize the services in different areas like storage, compute, network, etc.

Storage related services

The storage services are as follows:

- **Google Cloud Storage:** Google Cloud Storage is an object storage service. It can store any type or any amount of data. It is good for building Data Lakes and static content delivery websites like images, video streaming, etc. You can retrieve the data any time you want. It can also be used as a low-cost archival solution.
- **Persistent Disk:** It is a block storage service mainly used with Compute Engines/ VMs and GCP's Kubernetes service **Google Kubernetes Engine (GKE)**. It can be a storage solution attached to VMs as Disks. It can also be used as storage for installed databases.
- **Filestore:** It is a fully managed service mainly used for file storage for application migration.

- **Data transfer services:** It is a service that transfers data between one cloud service provider to GCP or from on-premises to GCP. The data transfer services are as follows:
- **Transfer appliance:** It is used to collect and move data from limited connectivity areas or bandwidth-constrained locations.
- **Local SSD:** It is an ephemeral locally attached block storage for VMs and containers.
- **Google Cloud Backup and disaster recovery (DR):** It is a service for managed backup and DR to protect file systems, databases, and VMs. It can also be used to keep a backup of on-premise workloads to GCP.
- **Artifact registry:** It is used as a centralized repository of container images, OS, and language packages. It is mostly used with CI/CD pipelines in various environments.

Compute related services

The compute services are as follows:

- **Compute Engine:** It is a general-purpose VM. It can be used for web server hosting, **Extract, Transform, Load (ETL)** and application servers. It can integrate with various GCP services. It can be used as a replica of an on-premises VM.
- **Google Kubernetes Engine:** It is a Google-managed implementation of Kubernetes, which is actually an open-source container orchestration platform. It is a reliable and secure way to deploy containerized applications.
- **Cloud Run:** It is a highly scalable service to deploy containerized applications. It comes with two versions, one is fully managed Cloud Run, and another is Cloud Run for Anthos. We will discuss this in more detail in later chapters.
- **App Engine:** It is a fully managed and serverless platform to build applications using common programming languages.
- **Batch:** It is a fully managed and dynamically scalable batch service to schedule, queue, and execute containerized or script-based batch jobs on Compute Engine.

Database related services

The database services are as follows:

- **Cloud SQL:** A fully managed MySQL or PostgreSQL or SQL Server database. It can be used as a regular RDBMS application database.
- **AlloyDB for PostgreSQL:** It is a PostgreSQL compatible GCP's database service.
- **Spanner:** Cloud Spanner is a very highly scalable global cloud native database with 99.999% of availability.