Learning Google Cloud Vertex AI

Build, deploy, and manage machine learning models with Vertex AI

Hemanth Kumar K

www.bpbonline.com

First published: 2024

Published by BPB Online
WeWork
119 Marylebone Road
London NW1 5PU

UK | UAE | INDIA | SINGAPORE

ISBN 978-93-55515-353

www.bpbonline.com

Dedicated to

All the people who are changing the world through technology

About the Author

Hemanth Kumar K has more than decade of experience in the corporate industry, started his career in telecom domain worked on multiple projects and clients across the globe to plan and optimize mobile networks. Got fascinated towards data & machine learning changed his focus and career path towards data science & analytics. He has worked on many data analytics projects to solve business problems in various domains such as insurance, telecom, media, health, engineering analytics. In addition to GCP he has also worked on Azure platform.

About the Reviewers

❖ **Dr. Muthu Kumaraswamy B** is an Associate Director at Searce where he leads a team of 100 + Machine learning engineers and delivers custom machine learning solutions. He has over 15 years of experience in the IT industry, and has worked on a variety of projects in the areas of software development, cloud computing, and artificial intelligence. He has a Ph.D. from the University Madras.

Dr. Muthu Kumaraswamy is an expert in a variety of programming languages, including JavaScript, Python, and C++. He is also proficient in a variety of cloud computing platforms, including Google Cloud Platform and Amazon Web Services. He is a strong advocate for using artificial intelligence to solve real-world problems, and has developed several AI-powered solutions for his clients.

Dr. Muthu Kumaraswamy is a frequent speaker at industry conferences and events. He is also a regular contributor to technical blogs and publications. He is passionate about sharing his knowledge and expertise with others, and is always looking for new ways to help people learn and grow.

Here are some of his notable achievements:

- Developed a cloud-based AI solution that helped a large manufacturing company improve its production efficiency by 15%.
- Led a team of engineers in the development of a new Machine learning platform that was used by a financial services company to automate its risk management processes.
- Published several articles in technical journals and magazines on the topics of software development, cloud computing, and artificial intelligence.
- Presented at several industry conferences on the use of AI to solve real-world problems.
- Dr. Muthu Kumaraswamy is a highly skilled and experienced Machine learning engineer with a proven track record of success in delivering high-quality solutions to his clients. He is a valuable asset to any team and would be a great addition to your company

- **Abdul Hannan Siddiqui** is a Passionate and Innovative Machine Learning Engineer with over 5 years of experience in MLOps, Time Series, Predictive Modeling, Computer Vision, and Natural Language Processing. He is a Google Certified Machine Learning Engineer with multiple Google Cloud Platform certifications, specializing in Data Engineering using Python, ETL, and effectively handling large-scale migration processes to the cloud, as well as data analytics on Google BigQuery.

  Abdul Hannan Siddiqui is currently working as a Senior Machine Learning Engineer at Argility. Prior to that, he worked as a Machine Learning Engineer at Quantiphi and as a Data Engineer at Accenture. His expertise lies in MLOps, Time Series, Predictive Modeling, Computer Vision, and Natural Language Processing, and he is highly skilled in using Python for Data Engineering and ETL processes. With his extensive experience and Google Cloud Platform certifications, Abdul Hannan Siddiqui is dedicated to leveraging his skills to drive successful machine learning and data science projects, ensuring optimal outcomes for his clients or organization.

Acknowledgements

I want to sincerely thank my family and friends for their continuous support and inspiration as I wrote this book.

I'm also appreciative of BPB Publications for their advice and know-how in helping me finish this book. This book had a lengthy revision process that benefited greatly from the participation and cooperation of editors, technical specialists, and reviewers.

Throughout my many years in the corporate industry, I would also like to recognize the important contributions of my coworkers and colleagues, who have taught me a lot and given me insightful criticism of my work.

Finally, I want to express my gratitude to all the readers who have shown interest in and support for my work. Your words of support have been

Preface

In today's world data enthusiast are expected to have expertise in one more cloud platforms. Solving business problems through data on cloud requires comprehensive understanding not just of algorithms but also cloud platforms and its components.

This book aims in helping you understand Vertex AI component of Google Cloud Platform (GCP) with real time examples. Before diving into Vertex AI, book gives you fair idea about google cloud platform and its components and also the positioning of Vertex AI. For all the examples used in this book for model training data is either uploaded to google cloud storage (GCS) or big query, so we will start with GCS, big query and also about the Identity access and management (IAM). Vertex AI has lot of components starting from datasets, Workbench, Pipelines, training, experiments, models, end points and features. The primary objective here is to learn these features of Vertex AI concepts with real-time use cases.

Chapter 1: Basics of Google Cloud Platform: starts with a fundamental understanding of Google Cloud Platform in this chapter. GCP's footprint, hierarchy, and platform, as well as how we can interact with GCP's components. GCP's several cloud service models (IAAS, PAAS, SAAS). A high-level overview of GCP's various components. We also take a look at GCP's storage capabilities, including how to set up Google Cloud storage and upload files to it. How to use Google's Big Query. We will comprehend the necessity and significance Incident and Access Management (IAM) in the platform.

Chapter 2: Introduction to Vertex AI and AutoML Tabular: In this chapter, the GCP Vertex AI component will be discussed. The integration of automl and custom model features into a single GCP component. We also comprehend the key distinctions between automl and custom models. We'll start with a walkthrough of how to create a dataset for structural data and how to train an automl model for it. We'll look at how to deploy the model and get batch and online predictions.

Chapter 3: AutoML Image, Text, and Pre-built Models: We'll work with unstructured data in this chapter. How do we make datasets out of unstructured data? (images and text data). How can an automl model be trained for picture and text classification tasks? How to use the trained model to make predictions

in batches and in real time. We'll start to see how we can use GCP's pre-trained models to solve our problem.

Chapter 4: Vertex AI Workbench and Custom Model Training: We'll begin working on how to use Vertex AI for custom model training. We'll begin with a basic model training example then progress to a more complex example. We'll work on the Vertex AI Workbench component, evaluate its various possibilities, and develop a single workbench for custom model training. We'll introduce you to the principles of Container, Docker, and Kubernetes before we go into custom Vertex AI models.

Chapter 5: Vertex AI Custom Model Hyperparameters and Deployment: We started with a modest model training exercise on Vertex AI to get a feel for custom model training. In this chapter, we'll look at a more difficult task: training a model with hyperparameter adjustment. We'll discuss how to deploy the custom model for serving predictions once it's been trained.

Chapter 6: Introduction to Pipelines and Kubeflow: In this chapter we will introduce you to the concept of MLOPS. Understanding machine learning life cycle. We will see how to build a pipeline using automl of Vertex AI with an example.

Chapter 7: Pipelines using Kubeflow for Custom Models: This chapter starts with the introduction of Kubeflow. How can we use Kubeflow components to build pipelines? How to create jobs for model training using container. CI/CD for Kubeflow pipelines.

Chapter 8: Pipelines using TensorFlow Extended: The introduction to TFX is the first part of this chapter. How can we construct pipelines using TFX components? How to use a container to construct jobs for model training. For TFX pipelines, CI/CD is used.

Chapter 9: Vertex AI Feature Store: This chapter is based on the Vertex AI Feature Store. We'll begin with an overview of the Feature Store. What is the purpose of feature stores? The fundamentals of feature stores are covered first, followed by a practical part on feature stores.

Chapter 10: Explainable AI: With Explainable AI, we'll begin investigating machine learning models in this chapter. How can we use explainable AI to better understand ML models for image categorization and how can we use explainable AI to better identify bias in ML models?

Code Bundle and Coloured Images

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

https://rebrand.ly/of74aa7

The code bundle for the book is also hosted on GitHub at **https://github.com/bpbpublications/Learning-Google-Cloud-Vertex-AI**. 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.

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

Chapter 1
Basics of Google Cloud Platform

Introduction

You will learn about the Google cloud platform in this chapter, as well as its benefits and the role it plays in today's digital revolution. Basic knowledge of cloud computing, including cloud service models, GCP account creation, footprint, range of services, and GCP hierarchy. This chapter will also introduce a few key GCP services, including storage, computation, google BigQuery and identity and access management, is then provided.

Structure

In this chapter, we will cover the following topics:

- Introduction and basics of Cloud platform
- Advantages of Cloud
- Importance of Cloud for data scientists
- Types of Cloud
- Introduction to Google Cloud platform
- Footprint of Google Cloud

- Cloud service model
- Services of GCP
- Hierarchy of GCP
- Interacting with GCP services
- Storage in GCP
- Compute in GCP
- BigQuery
- Identity and Access Management

Objectives

Before diving into the Vertex AI of the Google Cloud platform, it is very essential to grasp a few significant principles and vital services of the cloud platform. Users will have a solid understanding of the GCP components and services by the time this chapter ends. Detailed instructions for using GCP's storage, compute, and BigQuery services are included.

Introduction to Cloud

The term *Cloud* describes the applications and databases that run on servers that can be accessed over the Internet. Data centers across the globe host the cloud servers. Organizations can avoid managing physical servers or running software on their own computers by utilizing cloud computing. The cloud enables users to access the same files and applications from almost any device, because the computing and storage takes place on servers in a data center, instead of locally on the user device.

For businesses, switching to cloud computing removes some IT costs and overhead: for instance, they no longer need to update and maintain their own servers, as the cloud vendor they are using will do that.

Advantages of Cloud

There are various advantages of cloud as shown in *Figure 1.1*, and mentioned as follows:

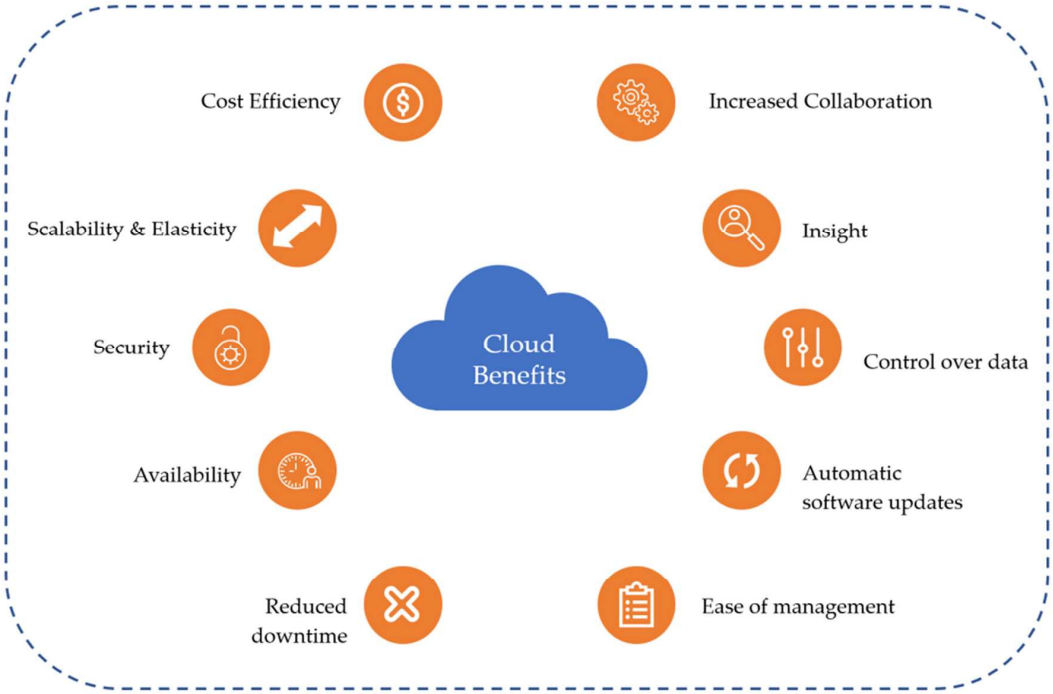

***Figure 1.1**: Advantages of Cloud platform*

- **Cost efficiency**: In terms of IT infrastructure management, cloud computing is undoubtedly the most cost-effective option. It is incredibly affordable for organizations of any size to transition from on-premises hardware to the cloud thanks to a variety of pay-as-you-go and other scalable choices. Using cloud resources instead of purchasing costly server equipment and PCs that need a lot of time to set up and maintain, such as long hours of setup and maintenance. Cloud also helps in reduced spending on compute, storage, network, operational and upgrade expenses.
- **Scalability and elasticity**: Overall, cloud hosting is more flexible than hosting on a local machine. You do not have to undertake a costly (and time-consuming) upgrade to your IT infrastructure if you need more bandwidth. This increased degree of latitude and adaptability may have a major impact on productivity.

Elasticity is only employed for a short amount of time to deal with rapid shifts in workload. This is a short-term strategy used to meet spikes in demand, whether they are unanticipated or seasonal. The static increase in workload is met through scalability. To cope with an anticipated rise in demand, a long-term approach to scalability is used.

- **Security**: Cloud platform provides a multitude of cutting-edge security measures, which ensure the safe storage and management of any data. Granular permissions and access control using federated roles are two examples of features that may help limit access to sensitive data to just those workers who have a legitimate need for it. This helps reduce the attack surface that is available to hostile actors. Authentication, access control, and encryption are some of the fundamental safeguards that providers of cloud storage put in place to secure their platforms and the data that is processed on those platforms. After that, users can implement additional security measures of their own, in addition to these precautions, to further strengthen cloud data protection and restrict access to sensitive information stored in the cloud.
- **Availability**: The vast majority of cloud service providers are quite dependable in terms of the provision of their services; in fact, the vast majority of them maintain an uptime of 99.9 percent. Moving to the cloud should be done with the intention of achieving high availability. The goal is to make your company's goods, services, and tools accessible to your clients and workers at any time of day and from any location in the world using any device that can connect to the internet.
- **Reduced downtime**: Cloud based solutions provide the ability to operate critical systems and data directly from the cloud or to restore them to any location. During a catastrophic event involving information technology, they make it easier for you to get these systems back online, reducing the amount of manual work required by conventional recovery techniques.
- **Increased Collaboration**: Developers, QA, operations, security, and product architects are all exposed to the same infrastructure and may work concurrently without tripping on one another's toes in cloud settings. To minimize disputes and misunderstanding, cloud roles and permissions provide more visibility and monitoring of who performed what and when. Different cloud environments, such as staging, QA, demo, and pre-production, may be created for specialized reasons. The cloud makes transparent collaboration simpler and promotes it.

- **Insight**: A bird's-eye perspective of your data is also provided through the integrated cloud analytics that are offered by cloud platforms. When your data is kept in the cloud, it is much simpler to put in place, monitoring systems and create individualized reports for doing information analysis throughout the whole organization. You will be able to improve efficiency and construct action plans based on these insights, which will allow your organization to fulfil its objectives.
- **Control over data**: Cloud provides you total visibility and control over your data. You have complete control over which users are granted access to which levels of specified data. This not only gives you control, but also helps simplify work by ensuring that staff members are aware of the tasks they have been allocated. Additionally, it will make working together much simpler. Because several users may make edits to the same copy of the text at the same time, there is no need that multiple copies of the document be distributed to the public.
- **Automatic software updates**: There is nothing more cumbersome than being required to wait for the installation of system upgrades, especially for those who already have a lot on their plates. Applications that are hosted in the cloud instantly refresh and update themselves, eliminating the need for an IT personnel to carry out manual updates for the whole organization. This saves critical time and money that would have been spent on consulting from other sources.
- **Ease of managing**: The use of cloud can streamline and improve IT maintenance and management capabilities through the use of agreements supported by SLA, centralized resource administration, and managed infrastructure. Users can take advantage of a simple user interface without having to worry about installing anything. In addition, users are provided with management, maintenance, and delivery of the IT services.

Importance of Cloud for data scientist

Since the beginning of the previous decade, the expansion of data has followed an exponential pattern, and this trend is expected to continue. The safe and secure storage of data should be one of the top priorities of every company. The cloud is usually the top option when it comes to storing and processing the enormous quantity of data since it has all of the advantages that were discussed above. As a consequence of this, a data scientist in today's world has to have experience with cloud computing in addition to their expertise in statistics, machine learning algorithms, and other areas.

However, due to the low processing capacity of their CPU, they are unable to carry out these responsibilities in a timely way, assuming that they are even capable of doing so at all. In addition, the memory of the machine is often incapable of storing massive datasets because of their size. It determines how quickly the assignment is performed and how well it was accomplished overall. Data scientists are now able to investigate more extensive collections of data without being constrained by the capabilities of their local workstations thanks to the cloud. Utilizing the cloud might result in a decrease in the cost of infrastructure since it eliminates the requirement for a physical server. In addition, depending on the cloud for data storage can lead to a reduction in the cost of infrastructure. In addition to offering data storage services, many cloud platforms including google cloud platform also has other services caterings to data ingestion, data processing, analytics, AI and data visualization.

Types of Cloud

There are three types of cloud based on different capabilities:

- Public Cloud
- Private Cloud
- Hybrid Cloud

Public Cloud: The public cloud is a massive collection of readily available computing resources, including networking, memory, processing elements, and storage. Users can rent these resources, which are housed in one of the public cloud vendors globally dispersed and fully managed datacenters, to create your IT architecture. Using a web browser, users have access to your resources in this form of cloud. Google Cloud Platform is an example for Public Cloud.

A major advantage of the public cloud is that the underlying hardware and logic are hosted, owned, and maintained by each vendor. Customers are not responsible for purchasing or maintaining the physical components that comprise their public cloud IT solutions. In addition, **Service Level Agreements** (**SLAs**) bind each provider to a monthly uptime percentage and security guarantee in accordance with regulations.

Private Cloud: Unlike public clouds, private clouds are owned and operated only by a single organization. They have usually been housed in the company's datacenter and run on the organization 's own equipment. To host their private cloud on their equipment, however, an organization may use a third-party supplier. Even if the resources are housed in a remotely managed datacenter, private cloud has certain

characteristics with public cloud in this case. They may be able to provide certain administrative services but they would not be able to offer the full range of public cloud services.

If the private cloud is housed in your own datacenter, organization have complete control over the whole system. A self-hosted private cloud may help to comply with some of the stricter security and compliance regulations.

Hybrid Cloud: This kind of cloud computing is a blend and integration of both public and private clouds, as the name of this form of cloud computing indicates. In this manner, it will be able to provide you with the advantages associated with a variety of cloud kinds when it comes to cloud computing. It enables a larger degree of flexibility in terms of the transmission of data and expands the alternatives available to a company for its adoption. This guarantees a high level of control as well as an easy transition while giving everything at rates that are more economical.

Introduction to Google Cloud Platform

Google Cloud Platform is one of the hyper scale infrastructure providers in the industry. It is a collection of cloud computing services that are offered by Google. These services operate on the same infrastructure that Google employs for its end-user products, including YouTube, Gmail, and a number of other offerings. The Google Cloud Platform provides a wide range of services, such as computing, storage, and networking, among other things.

Google Cloud Platform was first launched in 2008, and as of now, it is the third cloud platform that sees the most widespread use. Additionally, there is a growing need for platforms that are hosted on the cloud.

The Google cloud gives us a service-centric perspective of all our environments in addition to providing a standard platform and data analysis for deployments, regardless of where they are physically located. Using the capabilities of sophisticated analytics and machine learning offered by Google Cloud, we can extract the most useful insights from our data. Users will be able to automate procedures, generate predictions, and simplify administration and operations with the support of Google's serverless data analytics and machine learning platform. The services provided by Google Cloud encrypt data while it is stored, while it is being sent, and while it is being used. Advanced security mechanisms protect the privacy of data.

Account creation on Google Cloud Platform

Users can create free GCP account from the link **https://cloud.google.com/free**.

Free account provides 300$ credit for a period of 90 days.

Steps for creating a free account are as follows:

1. Open **https://cloud.google.com/free**.
2. Click on **Get started for free**.

   The opening screen looks like *Figure 1.2*:

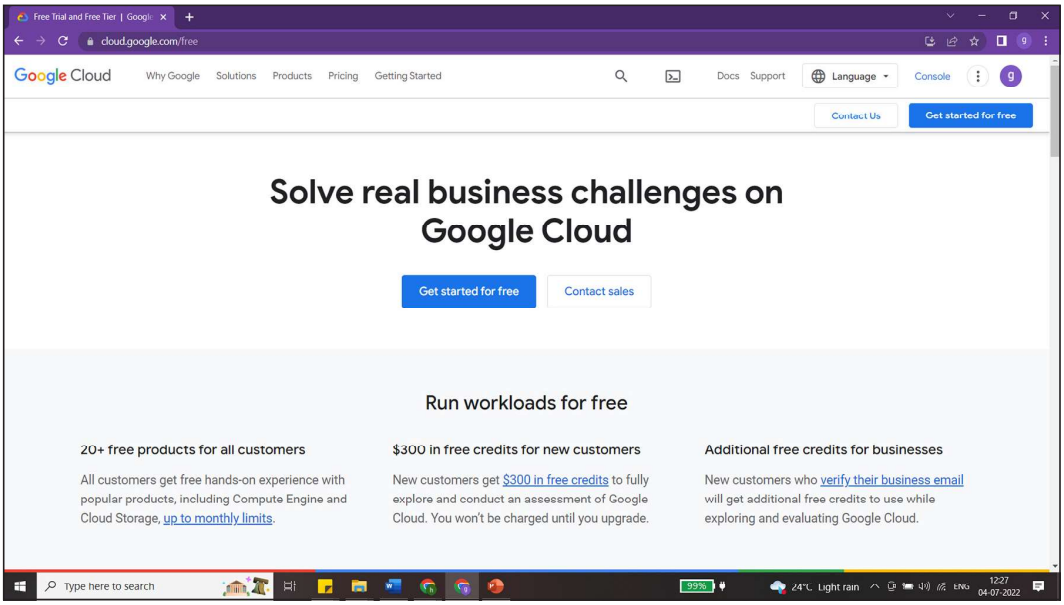

***Figure 1.2**: GCP account creation*

3. Login with your Gmail credentials, create one if you do not have. This can be seen illustrated in *Figure 1.3*:

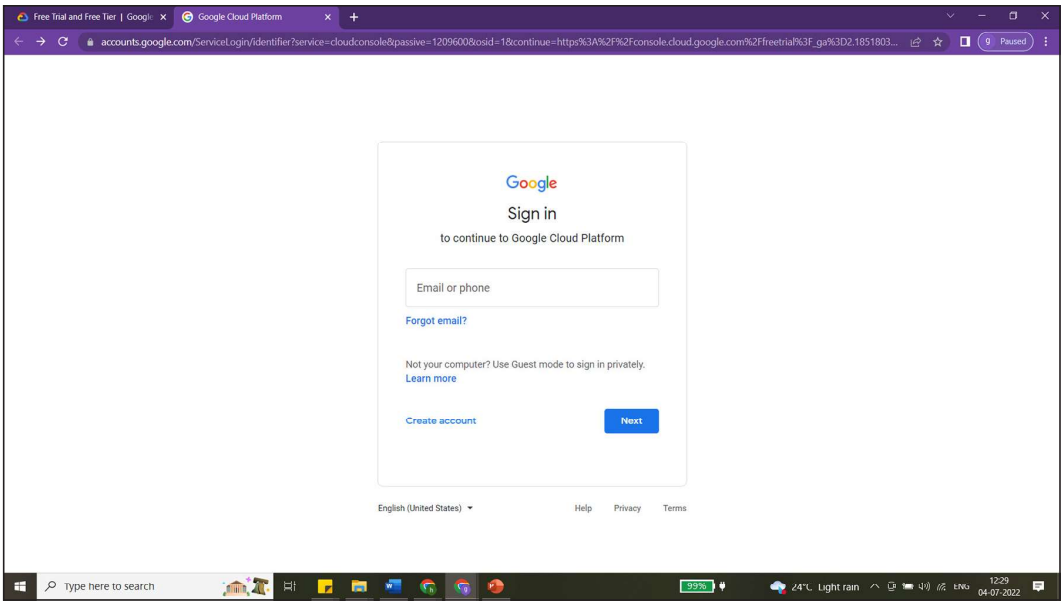

***Figure 1.3**: GCP account creation enter valid mail address*

4. Selection of **COUNTRY** and needs:

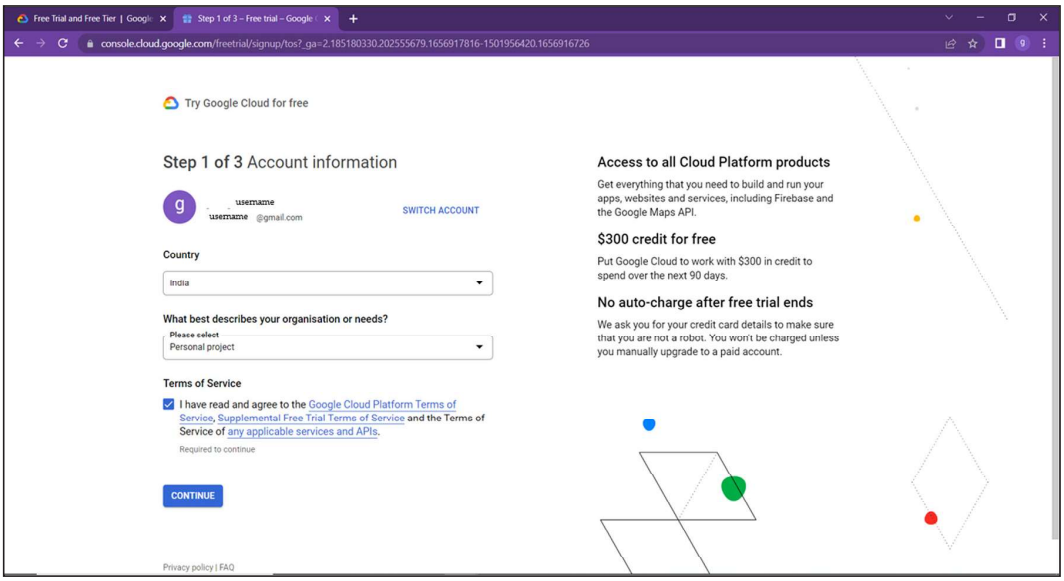

***Figure 1.4**: GCP account creation country selection*

5. Select the **Country** and project. Check the Terms of service and click on **CONTINUE**.

6. Provide phone number for the identity verification as shown in *Figure 1.5*:

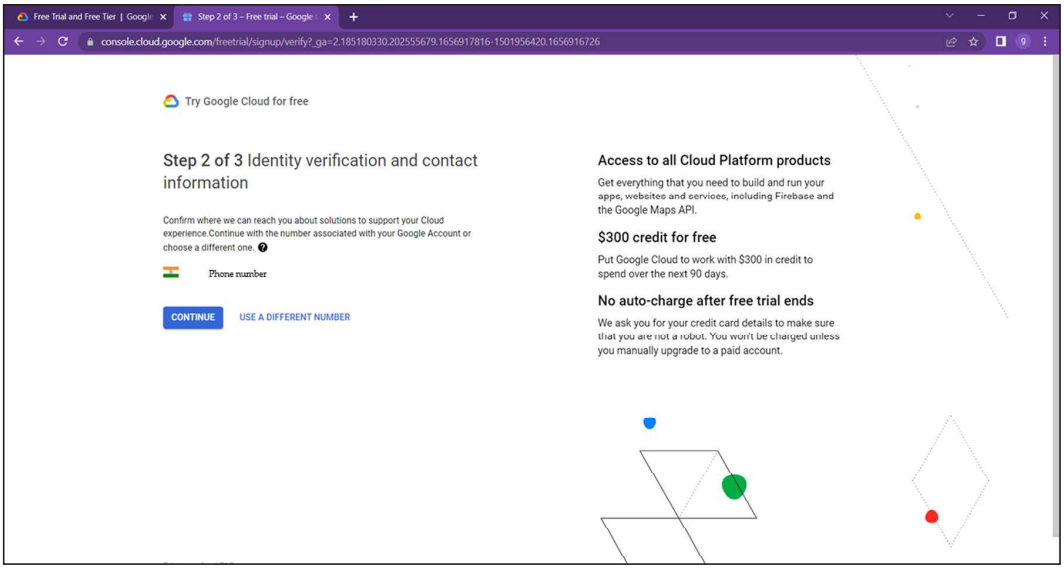

***Figure 1.5**: GCP account creation enter phone number*

7. Free accounts require a credit card. Verification costs Rs 2. Addresses must be provided. Click on **START MY FREE TRIAL** on this page:

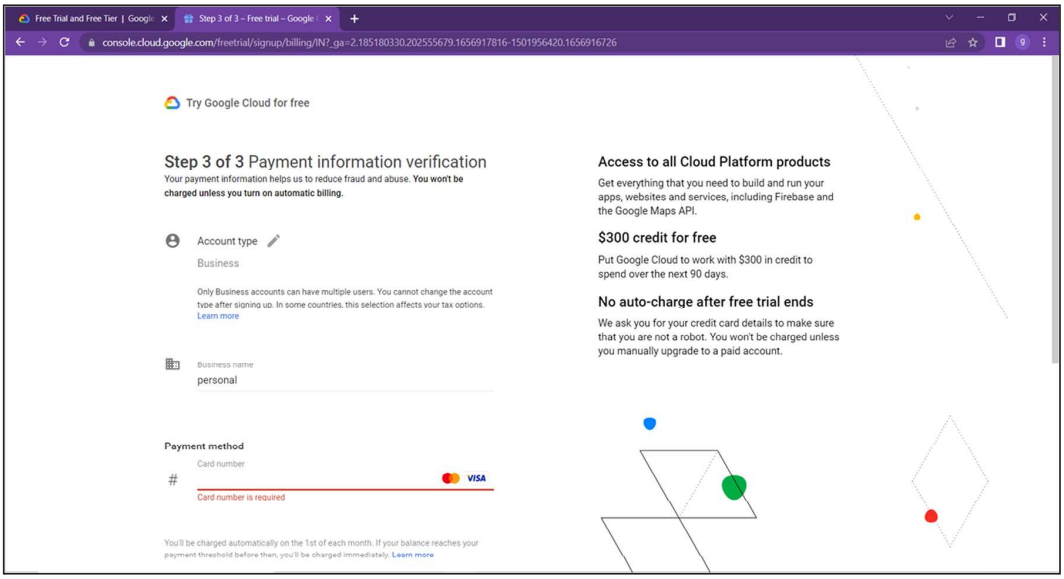

***Figure 1.6**: GCP account creation enter valid credit card details*

8. Users will land into this page once the free trail has started. The welcome page can be seen in *Figure 1.7*:

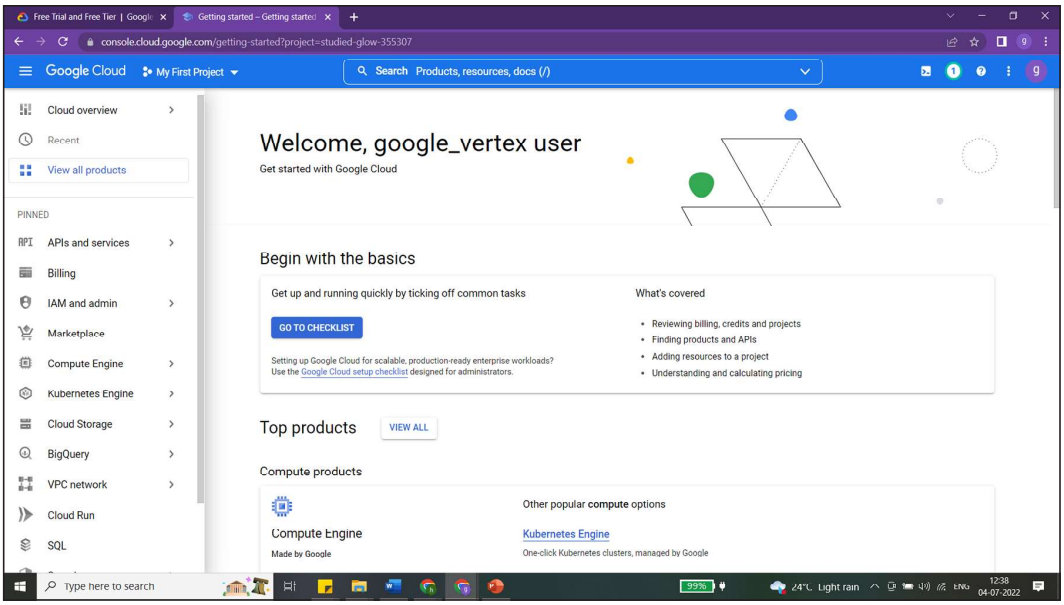

***Figure 1.7**: Landing page of GCP*

Footprint of Google Cloud Platform

Independent geographical areas are known as regions, while zones make up regions. Zones and regions are logical abstractions of the underlying physical resources that are offered in one or more datacenters physically located throughout the world. Within a region, the Google Cloud resources are deployed to specific locations referred to as zones. It is important that zones are seen as a single failure area within a region. *Figure 1.8* shows the footprint of GCP:

***Figure 1.8**: Footprint of GCP*

The time this book was written there were about 34 regions, 103 zones and 147 network edge location across 200+ countries. GCP is constantly increasing its presence across the globe, please check the link mentioned below to get the latest numbers.

Image source: **https://cloud.google.com/about/locations**

The services and resources offered by Google Cloud may either be handled on a zonal or regional level, or they can be managed centrally by Google across various regions.:

- **Zonal resources**: The resources in a zone only work in that zone. When a zone goes down, some or all of the resources in that zone can be affected.
- **Regional resources**: They are spread across multiple zones in a region to make sure they are always available.
- **Multiregional resources**: Google manages a number of Google Cloud services to be redundant and spread both inside and between regions. These services improve resource efficiency, performance, and availability.
- **Global resources**: Any resource within the same project has access to global resources from any zone. There is no requirement to specify a scope when creating a global resource.

Network edge locations are helpful for hosting static material that is well-liked by the user base of the hosting service. The material is temporarily cached on these edge nodes, which enables users to get the information from a place that is much closer to where they are located. Users will have a more positive experience as a result of this.

There are few benefits associated with the GCP's regions and zones. When it comes to ensuring high availability, high redundancy, and high dependability, the notion of regions and zones is helpful. Obey the laws and regulations that have been established by the government. Data rules might vary greatly from one nation to the next.

Cloud Service Model

The cloud platform offers a variety of services, all of which may be roughly placed into one of three distinct categories:

- **Infrastructure as a service (IAAS)**
- **Platform as a service (PAAS)**
- **Software as a service (SAAS)**

The difference between cloud service models is illustrated in the *Figure 1.9*

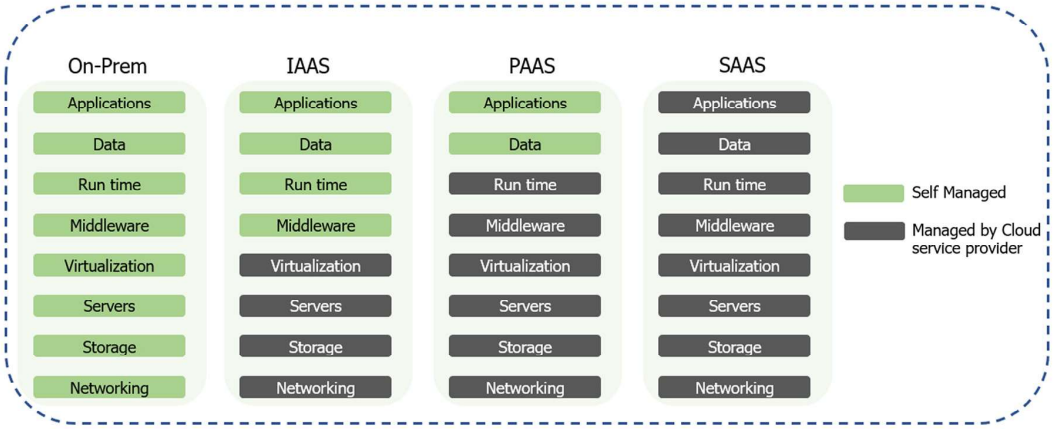

***Figure 1.9**: Cloud Service Model*

Let us imagine we are working on an application and hosting it at the same time on a server that is located on our premises. In this particular circumstance, it is our obligation to own and maintain the appropriate infrastructure, as well as the appropriate platforms, and of course, our application.

- **Infrastructure as a service**: In the case of IAAS, it will be the cloud's obligation to provide the necessary infrastructure, which may include virtual machines, networking, and storage devices. We are still responsible for ensuring that we have the appropriate platform for development and deployment. We have no other option for exercising control over the underlying infrastructure but to make use of it. One example of an infrastructure as a service provided by Google with its compute engine and Kubernetes engine.
- **Platform as a service**: In the case of PAAS, the responsibility for providing the appropriate platform for development and deployment, such as an operating system and tools for the environment in which programming languages are used, lies with the cloud service provider. This responsibility exists in addition to the infrastructure responsibility to provide such things. One example of a PAAS platform is Google App Engine.
- **Software as a service**: In the instance of SAAS, a cloud service provider will rent out to their customers apps that are theirs to run on their infrastructure. The maintenance of the software applications will also fall within the purview of the cloud service provider, in addition to the platform and the underlying infrastructure. These software programs are accessible to us on

whatever device we choose by way of web browsers, app browsers, and so on. Email (Gmail) and cloud storage (Google Drive) are two excellent instances of SAAS.

- **Data as a service (DAAS)**: DAAS is a service that is now starting to gain broad use, in contrast to the three service models that were mentioned before, which have been popular for more than a decade. This is partly owing to the fact that general cloud computing services were not originally built for the management of enormous data workloads; rather, they catered to the hosting of applications and basic data storage needs (as opposed to data integration, analytics, and processing).

SaaS eliminates the need to install and administer software on a local computer. Similarly, Data-as-a-Service methodology centers on the on-demand delivery of data from a number of sources using **application programming interfaces** (**APIs**). It is intended to make data access more straightforward, and provide curated datasets or streams of data that can be consumed in a variety of forms. These formats are often unified via the use of data virtualization. In fact, a DaaS architecture may consist of a wide variety of data management technologies such as data virtualization, data services, self-service analytics.

In its most basic form, DaaS enables organizations to have access to the ever-growing quantity and sophistication of the data sources at their disposal in order to give consumers with the most relevant insights. The democratization of data is something that is absolutely necessary for every company that wants to get actual value from its data. It gives a significant potential to monetize an organization's data and acquire a competitive edge by adopting a more data-centric approach to the operations and procedures.

Services offered by GCP

Users may make use of a comprehensive selection of services provided by Google Cloud Platform. Every one of the services may be placed into one of the categories that are shown in the *Figure 1.10*:

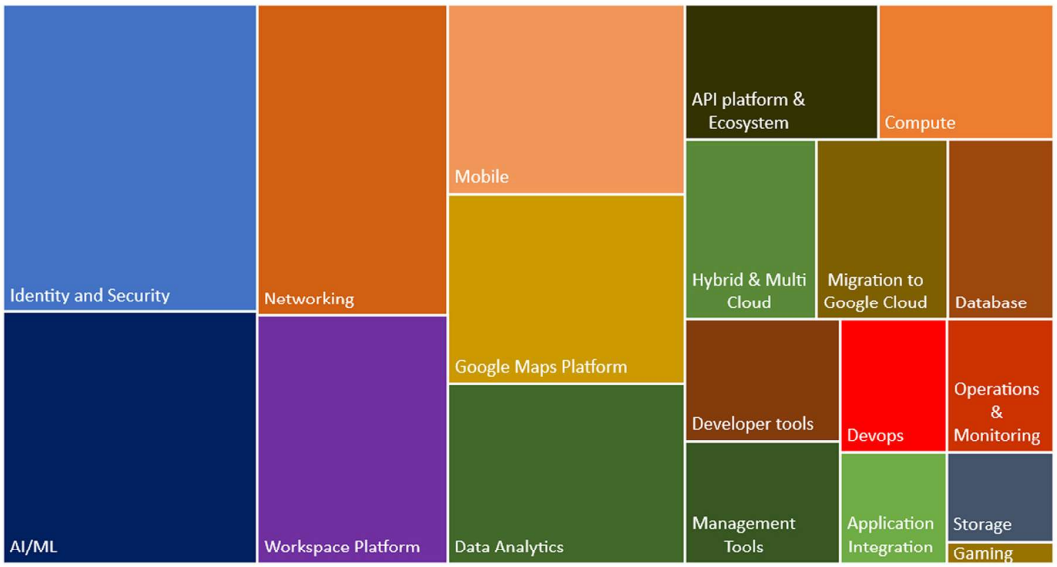

***Figure 1.10**: Services of GCP*

- Google offers Cloud Storage for storing unstructured items, Filestore for sharing files in the traditional manner, and persistent disk for virtual machines in the storage space. Compute Engine, App Engine, Cloud Run, Kubernetes Engine, and Cloud Functions are the core computing services that Google Cloud Platform provides.
- Cloud SQL, supports MySQL, PostgreSQL, and Microsoft SQL Server; and Cloud Spanner, which is a massively scalable database that is capable of running on a global scale. These are the relational database services that GCP provides.
- Bigtable, Firestore, Memorystore, and Firebase Realtime Database are different NoSQL services that Google provides. When dealing with massive amounts of analytical work, Bigtable is the most effective solution. Firestore is well suited for use in the construction of client-side web and mobile apps. Firebase Realtime Database is ideal for real-time data synchronization between users, such as is required for collaborative app development. Memorystore is a kind of datastore that operates entirely inside memory and is generally used to accelerate application performance by caching data that is frequently accessed.
- BigQuery is the name of the data warehouse service offered by Google.

- A **Virtual Private Cloud** (**VPC**) is an on-premises network on GCP. By using VPC Network Peering, virtual private clouds may be linked to one another. Users may utilize Cloud VPN, which operates over the internet, to establish a protected connection between a VPC and an on-premises network. Alternatively, users can establish a dedicated, private connection by using either Cloud Interconnect or Peering. To facilitate the migration of applications and data sets to its platform, the platform provides a diverse set of options. Offers Anthos as an alternative for hybrid cloud deployments.
- The field of data analytics is one in which Google excels in particular. Pub/Sub is used as a buffer for services that may not be able to deal with large surges in the amount of data coming in. Dataproc is a Hadoop and Spark implementation that is controlled by Dataproc. Apache Beam is the underlying technology that powers Dataflow, a managed implementation. You can do data processing using Dataprep even if you do not know how to write code, and it leverages Dataflow behind the scenes. Users may use google looker studio to visualize or show your data using graphs, charts, and other such graphical representations.
- Platform provides AI and ML services for a diverse group of customers. Vertex AI provides AUTOML option for the novices, for more experienced users, it provides trained models that make predictions via API and also provides various options for the advanced AI practitioners.
- Cloud Build enables you to develop continuous integration / continuous deployment pipelines. Private Git repositories that are hosted on GCP are known as Cloud Source Repositories. Artifact Registry expands on the capabilities of Container Registry and is the recommended container registry for Google Cloud. It provides a single location for storing and managing your language packages and Docker container images.
- **IAM** stands for **Identity and Access Management,** and it enables users and apps to have roles assigned to them. Everything you store in the GCP is by default encrypted. Companies now have the ability to control their encryption keys thanks to Cloud Key Management. Your API keys, passwords, certificates, and any other sensitive information may be safely stored in the Secret Manager.
- The Monitoring, Logging, Error Reporting, Trace, Debugger, and Profiler functions are all included in the Cloud Operations suite. The Active Security Threats and Vulnerabilities, as well as Compliance Infractions, are Presented to You by the Security Command Center. The development of Google Cloud Platform resources may be automated with the help of Cloud Deployment Manager.