

Tutorial: Google Cloud for Beginners: Architecture, Storage, and Computation

Lewis Tseng, Haochen Pan, Yingjian Wu
Boston College
Boston, USA
{lewis.tseng, haochen.pan, wuit}@bc.edu

TUTORIAL ABSTRACT

Cloud computing is a powerful paradigm to support pervasive applications that require always-on computation and interaction. In this tutorial, we focus on the Google Cloud Platform (GCP). Even though it is not currently the leader in the cloud industry, GCP is one of the best choices for deploying machine learning-based applications. This is due to the advancement of Google's deep learning technology, e.g., TensorFlow and TPU.

This tutorial consists of a series of hands-on exercises for beginners that have minimum or even no experience in cloud computing platforms. We will first give an overview of cloud computing. Then we cover some high-level architecture of GCP and how they affect application development and deployment. Finally, we will walk through exercises on using several services in GCP with the focus on storage and computation services. If time permits, we will build a simple application using TensorFlow.

TUTORIAL OUTLINE

The tutorial will follow the outline below:

- Brief introduction to cloud computing
- Brief introduction to Google Cloud Platform (GCP)
- Basic computation and storage service provided by GCP
- VMs and containers in GCP
- TensorFlow/ML applications (if time permits)

AUDIENCE

The tutorial will be accessible to anyone with a background of basic knowledge on algorithms and programming. We do *not* assume any background in cloud computing or distributed systems. Some knowledge of Linux would help.

We welcome everyone to attend, but the tutorial is mainly designed for people who are interested in (i) architecting computation and storage solutions using GCP, and (ii) developing and deploying customized machine learning-based applications in GCP.

One does not need to have a GCP account. As of February 2020, GCP provides a free tier “A 12-month free trial with \$300 credit to use with any Google Cloud services.”¹

¹<https://cloud.google.com/free/docs/gcp-free-tier>

LEARNING GOALS

After the tutorial, you are expected to learn:

- Deploy simple GCP storage services and own applications on GCP
- Use big data solutions such as Spanner and Pub/Sub
- Build simple deep learning models using TensorFlow (if time permits)

BIO



Lewis Tseng is currently an assistant professor in Computer Science department at Boston College. Before that, he spent a year and a half as a researcher at Toyota

InfoTechnology Center. He received a B.S. and a Ph.D. degree both in Computer Science from the University of Illinois at Urbana-Champaign (UIUC) in 2010 and 2016, respectively. His research broadly lies in the intersection of fault-tolerant computing and distributed computing. He won the best paper award in the International Symposium on Stabilization, Safety, and Security of Distributed Systems (SSS) 2017.



Haochen (Roger) Pan is a junior at Boston College pursuing his dual degree in Computer Science and Mathematics. His research interests include distributed computing & systems, Blockchain-based systems, and vehicular ad-hoc networks. He received the Sophomore Scholar

Award, the Advanced Study Grant, and the Undergraduate Research Fellowship from Boston College.



Yingjian (Steven) Wu

is a senior at Boston College pursuing his dual degree in Computer Science and Mathematics. His research interests include distributed computing & systems, Blockchain-based systems, and Big Data infrastructure internal

design. He received IEEE Comsoc Student Travel Grant for Globecom 2019. He received the Advanced Study Grant for thesis research and Undergraduate Research Fellowship from Boston College.