

EECS E6893 Big Data Analytics - Fall 2024

Homework Assignment 0: Intro to Big Data Analytics on GCP

Due Friday, September 20th, 2024 by 5:00pm

TL;DR

1. Setup GCP and Cloud SDK
2. Learn Dataproc (Spark), Cloud Storage, and BigQuery
3. HW: write a report containing your answers to the 3 questions and submit to Canvas

Abstract

The goals of this assignment are to (1) become familiar with one of the most popular cloud computing platforms: [Google Cloud Platform](https://cloud.google.com/) (GCP), (2) have hands-on exposure to several [big data products](#) offered by GCP: Cloud Storage, BigQuery, and Dataproc (Hadoop and Spark), (3) lay the foundations for the rest of the semester.

In this assignment, you will setup your GCP account and environment, create a Spark cluster, load data from Cloud Storage and BigQuery, and process data with Spark. Note that this assignment is “documentation heavy”, which is inevitable since you have to learn how to use tools on GCP. Later assignments should be more focused on data analysis and algorithm implementation.

Warm-up exercises

1. GCP account setup

- a. Head over to GCP and make sure you have a google account.
<https://cloud.google.com>
- b. Apply for \$300 credit for a year.
<https://cloud.google.com/free/>
- c. Go to Billing -> Account management to check your credit. You should expect \$300 in the “Promotional value”.

2. Install Cloud SDK

Google Cloud SDK is a set of tools that you can use to manage resources and applications hosted on Google Cloud Platform. It should be handy if you have this installed on your local computer.

- a. Choose a tutorial that is suitable for your local environment.
<https://cloud.google.com/sdk/docs/quickstarts>

- b. Follow the tutorial thoroughly and install SDK on your computer. Remember to **create a GCP project** first as noted in the tutorials.
- c. You should expect info pop up when you enter the command `gcloud info`
- d. You could use `gcloud init` to change the default configuration setup like account / project

3. Dataproc

Dataproc is an on-demand, fully managed cloud service for running Apache Hadoop and Spark on GCP. By using Dataproc, we don't need to maintain the underlying infrastructure, and we could easily increase / decrease the number of resources we want. It also offers built-in integration with other services on GCP like Cloud Storage and BigQuery. In this exercise, you'll create a Spark cluster, and submit a job that runs a Spark example program - "Pi calculation".

- a. Learn how to create a cluster and run a Spark job by following this tutorial. <https://cloud.google.com/dataproc/docs/quickstarts/quickstart-gcloud>
Note that GCP provides different ways to achieve the same goal, like using the API explorer or console, but we believe it is beneficial to use the command-line tool.
- b. At the end of the tutorial, you should expect the following things:
 - i. You've created a computing cluster with 3 nodes (1 master + 2 workers):
 - ii. You've submitted a Spark job:
 - iii. You've successfully calculated the value of Pi:

Optional reading:

- Create a cluster with Jupyter Notebook access: <https://cloud.google.com/dataproc/docs/tutorials/jupyter-notebook>
- Create a cluster with single node to save money: <https://cloud.google.com/dataproc/docs/concepts/configuring-clusters/single-node-clusters>
In short, you could create a single node cluster with Jupyter Notebook access by using a command.
- Take a look at the source code and examples on the official Spark site: <https://spark.apache.org/examples.html>

Remember to delete your dataproc clusters when you finish executions to save money.

4. “Word Count” using Google Cloud Storage and Spark

Cloud Storage is a file storage system built by Google for GCP. It’s an Infrastructure as a Service (IaaS), comparable to AWS S3 service. It has [pros and cons](#) compared to the standard HDFS on Hadoop. It is handy to perform operations with Cloud Storage and Spark on GCP.

- a. Study the [Spark programming guide](#).
- b. Learn how to write and execute a word-count spark program with Cloud Storage and Spark.

<https://cloud.google.com/dataproc/docs/tutorials/gcs-connector-spark-tutorial>

Optional reading:

- Learn gsutil tool to interact with Cloud Storage:
<https://cloud.google.com/storage/docs/gsutil>
- Study [Spark RDD programming guide](#) and [Spark SQL guide](#)
- Read [the Spark paper](#)

5. BigQuery

BigQuery is a data warehouse solution developed by Google for GCP. Like Apache Hive, they both are data warehouse software, a database designed for data analysis. Learn more about how they compare in [this article](#). In BigQuery, you could write SQL-like queries to interact with data from different sources.

- a. Explore BigQuery and learn how to use front-end UI to perform query.
<https://cloud.google.com/bigquery/docs/bigquery-web-ui#overview>
- b. Learn how to load data into BigQuery.
<https://cloud.google.com/bigquery/docs/loading-data>
- c. Check out the BigQuery SQL documentation as a reference.
<https://cloud.google.com/bigquery/docs/reference/standard-sql/query-syntax>

Optional reading:

- Take a look at different ways to query BigQuery data:
<https://cloud.google.com/bigquery/docs/query-overview>
- Trace the code to know how to use BigQuery with Spark:
<https://cloud.google.com/dataproc/docs/tutorials/bigquery-connector-spark-example>

Homework Submissions

1. Warm-up exercises:

For the “Pi calculation” in exercise 3 and “word count” example in exercise 4:

- (1) Provide screenshots to prove you’ve completed the exercises. (10%)
- (2) List the Spark [transformations and actions](#) involved in each exercise. Identify the RDD operation that triggers the program to execute. (20%)

2. NYC Bike expert (30%)

Download the NYC Citi bike data [here](#), upload the csv file to a Cloud Storage bucket, then load the data into BigQuery (check the [documentation](#)). You could choose “auto detect” schema when creating the table. Answer the following questions and provide screenshots of your queries:

- (1) How many stations with longitude between -73.94 and -74.04?
- (2) What’s the total number of bikes available in region_id 71?
- (3) What’s the largest *capacity* for a station? List all the *station_id* of the stations that have the largest capacity?

3. Understanding William Shakespeare (40%)

Write a Spark program to find the 5 most frequent words of the [provided shakes.txt file](#) extracted from Shakespeare’s works. Provide your results (words and its frequency) and take a screenshot of your code / jupyter notebook:

- (1) Find top 10 frequent words without any text preprocessing.
- (2) Find top 10 frequent words by first filtering out stop words provided by NLTK package. The Natural Language Toolkit, or more commonly NLTK, is a suite of libraries and programs to conduct natural language processing in Python.