

A blue-tinted background image of a space shuttle launch. The shuttle is angled upwards from the bottom left towards the top right. The main body of the shuttle is in the foreground, with the external tank and boosters visible. The background shows a bright, hazy sky. The overall tone is professional and technical.

Launch Jupyter to the Cloud

With the power of Terraform and Docker

By Cheuk Ting Ho



Cheuk Tink Ho

@cheukting_ho

- Base in London
- Love Python
- Love Comedy
- Love Beer
- Love Travel


Data scientist at



Organizer of

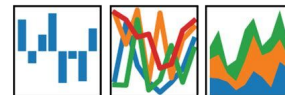


Member of



Contributor of
pandas


$$y_{it} = \beta' x_{it} + \mu_i + \epsilon_{it}$$



What is Docker Container? (a.k.a the cute whale with boxes)

“ A container image is a lightweight, stand-alone, executable package of a piece of software that includes everything needed to run it: code, runtime, system tools, system libraries, settings.....

software will always run the same, regardless of the environment. ”



What is Terraform? (infrastructure as code?)

“Infrastructure as code (IaC) is the process of managing and provisioning computer data centers through machine-readable definition files, rather than physical hardware configuration or interactive configuration tools.”

In cloud computing, which means we can launch an instance by code rather than clicking on the console



The mission - Coursera deep learning assignment: train a CNN model

Computing power

A laptop's CPU is estimated to spend 7 hrs to train a large neural network.

Nvidia GPU?

(tensorflow-gpu like Nvidia)

Price

It cost £200 - 1000 for a Nvidia GPU

GCP GPU accelerator:
\$0.22 - 1.46 USD per GPU per hr

N1-standard-4 GCE:
\$0.04 - 0.19 USD per hr

Customization

Coursera robo-grader needs a very specific environment

Luckily there's already a Docker image for the course

Just need some customization

Action Plan for launching **Jupyter** to **Google Cloud Platform**

Task 1

Write Terraform Script

You will also need:

- Terraform installed
- Service account key for Google API
- SSH key setup for your project

Task 2

Write startup script to install driver and Docker

- Put all installation in code
- Pull and run Docker image

Task 3 (Optional)

Build Docker Image

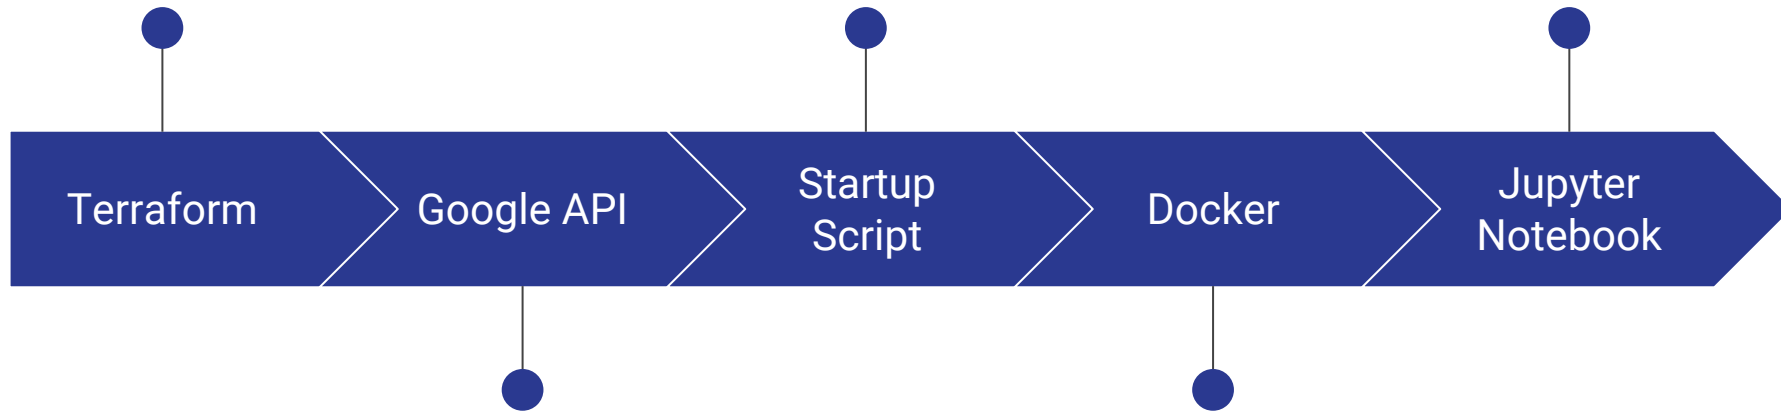
- You want to make build one yourself
- Git clone your repo

What's going on?

Send the structure and the startup script to the API

Run script when the instance is created, install GPU driver and Docker

The Jupyter Notebook with the requirements were set up by the Docker image



Launch an instance with the structure

Startup script also use docker to pull and run the image

Advantage

Why use Terraform and Docker?

Reproducible:

- Collaboration: You can share it with your friend
- Easily modify for different “mission”

Mistake-proof:

- Everything is code, no “hidden” settings
 - Reduce manual error
 - Easy to start all over again
-

Demo...(do I have enough time?)

Reference

Resources to help you get started

Documentations:

- Terraform:
<https://www.terraform.io/docs/providers/google/index.html>
- GCP GCE:
<https://cloud.google.com/compute/docs/>

My Github repo:

- <https://github.com/Cheukting/jupyter-cloud-demo>
- <https://github.com/Cheukting/GCP-GPU-Jupyter>

Other tutorials:

- <http://container-solutions.com/simple-gce-setup-terraform/>
- <https://blog.elasticbyte.net/getting-started-with-terraform-and-google-compute-engine/>