

Very Basic Introduction to Amazon EC2

CMU 16-785 Integrated Intelligence in Robotics (Spring 2019)

Allan Wang

Overview

- Preparations
- AMI and Instances
- Storage: EBS
- Backup & Data transfer: Snapshot
- Pricing

Preparations

- <https://aws.amazon.com/>
- Redeem coupon code
- Set region

AMI and Instances

- Instances are virtual machines.
- Instances created from Amazon Machine Image (AMI).
- AMI are templates that have OS and various software preinstalled.
- Instance type determines your hardware configuration.

AMI recommendations

- For code development, some basic debugging etc. (Free)
 - **Amazon Linux AMI 2018.03.0 (HVM), SSD Volume Type**
- For training, evaluating, inferencing deep learning models
 - **Deep Learning AMI (Ubuntu) Version 21.2**

Instance type recommendation

- t2: Basic CPU performance
- p2: General purpose GPU applications
- t2.micro
 - 750 hours of free usage per month
 - Use this for writing codes or debug some parts of your code
- p2.xlarge
 - Only use this when you need a GPU
 - 1 NVIDIA Tesla-K80 (12GB memory)
 - \$0.9 per hour

Key Pair

- Need the key file to log into instances or transfer files.
- Only need one.
- Save it at a secure location, don't lost it.

Log into instance (Linux)

- `chmod 0400 <path-to-key-file>`
- `ssh -i <path-to-key-file> <user-name>@<public-DNS-from-EC2>`
- User name:
 - `ec2-user` (**Amazon Linux AMI 2018.03.0 (HVM), SSD Volume Type**)
 - `ubuntu` (**Deep Learning AMI (Ubuntu) Version 21.2**)

Log into instance (Windows)

- Have PuTTY and PuTTYgen installed.
- Open PuTTYgen
 - Load an existing private key file (Select “All Files” when loading)
 - In “Parameters”, select a type that has “RSA”
 - Save private key. (This generates a .ppk file)
 - Click yes/ok on all warnings
- Open PuTTY
 - Host Name is <user-name>@<public-DNS-from-EC2>
 - Menu on the left: Connection -> SSH -> Auth
 - Browse Private key file for authentication.
 - Click Open

Storage

- When an instance is created, a root storage is also created.
- Root storage stores OS, installed software etc.
- Do not store your codes or datasets on this.
- Root storage will be deleted when the instance is terminated.
- This will be charged additionally (EBS pricing).

Amazon Elastic Block Store(EBS)

- Blocks of storage that can be attached to or detached from instances.
- Generally use “General Purpose SSD (gp2)”
- This is where you store your codes and datasets.
- Can be initialized with snapshot.
- 30 GB and under is free (including root storage)

Make EBS available to use

- Create EBS volume
 - Size can be increased later, but can not be decreased.
 - Make sure it's in the same subnet region as the instance
- Attach/detach EBS volume on Amazon EC2 console.

Make EBS available to use

- After Attach:

- If the volume is created for the first time:

- lsblk (*shows all attached devices, find name of attached EBS volume, e.g. xvdf*)
 - sudo file -s /dev/<device_name> (*shows if the volume has a file system.*)
 - sudo mkfs -t ext4 /dev/<device_name> (*creates the ext4 file system on the volume*)

(**Important:** only do this when the previous command returns “data”! Otherwise, all the data in this volume will be erased!)

- sudo mkdir <mount_point> (*creates a folder as the mounting point for the volume*)
 - sudo chmod 0777 -R <mount_point> (*add permissions to the folder*)
 - sudo mount /dev/<device_name> <mount_point> (*mount the volume*)
 - df -h (*show usage on the devices and volumes*)

Make EBS available to use

- Remember to save all your work in the mounted directory.
- Before detach:
 - lsblk
 - `umount -d /dev/<device_name>` (*unmounts the volume*)

Snapshots

- Backups. In case you lost your volume by accident.
- More importantly, snapshots can be shared with others.
- Uses Amazon S3 cloud storage service.
- 1 GB of free storage

Snapshots

- Snapshots can be created from EBS volumes.
- EBS volumes can be created from snapshots.
- Share snapshots by modifying permissions on EC2 console.
- Teammate's AWS ID needed.
- Teammate can then find the snapshot under "Private Snapshots"
(Not "Owned By Me")
- Teammate can create EBS volumes from the snapshot and get all the data contained within

Note on p2.xlarge

- Need to set up virtual environment first.
- Environments named as <deep_learning_package>_p<27/36>
- source activate <environment>
 - For example, source activate tensorflow_p27
- Check README once logged in.

Pricing

| | Instance Running | Instance stopped | Instance terminated |
|--------------------|------------------|------------------|---------------------|
| Instance charge | Charged | Not charged | Deleted |
| Root device charge | Charged | Charged | Deleted |
| EBS charge | Charged | Charged | Charged |

Pricing

- Instance type charge:

<https://aws.amazon.com/ec2/pricing/on-demand/>

- EBS charge:

<https://aws.amazon.com/ebs/pricing/>

- Note: EBS snapshots to Amazon S3 is charged differently from Amazon S3 storage charge.

Detailed documentation

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/concepts.html>