

# AWS Recitations

By Zhe Chen (presenter) and Eason

## Some recommended EC2 Images (AMI) for deep learning

	AWS Deep Learning AMI (Ubuntu 18.04)	AWS Deep Learning Base AMI (Ubuntu 18.04)
Boot volume size	≥ 100GB	≥ 60GB
Includes	<ul style="list-style-type: none"><li>• Drivers, CuDNN, CUDA</li><li>• Anaconda environments<ul style="list-style-type: none"><li>• Machine learning packages</li><li>• Deep learning frameworks</li></ul></li></ul>	<ul style="list-style-type: none"><li>• Drivers, CuDNN, CUDA</li><li>• Python 3.6.9 environment with pip<ul style="list-style-type: none"><li>• iPython</li><li>• Jupyter notebook</li><li>• Basic ML packages</li></ul></li></ul>
Occupied boot volume space	88GB	42GB

## Recommended instance types

Type	vCores	RAM (GB)	GPU Spec	Additional Spec	Suitable for
t2.micro	1	1	-	Free tier available	Test only
c5.xlarge	4	8	-	-	Basic code development when GPU instances are not available
g4dn.xlarge	4	16	1 x Nvidia T4 16GB VRAM	125GB of additional <b>ephemeral</b> NVMe drive	Code development and model training
p3.2xlarge	8	61	1 x Nvidia Tesla V100 16GB VRAM	Expensive!	Accelerated training of large models

## On-demand vs Spot

	On-Demand Instance	Persistent Spot Instance	One-time Spot Instance
Spin up	Should always work	Subject to spot capacity in your selected region	
Interruptions	-	Can be stopped/put to hibernation by AWS	Can be terminated by AWS
Stop action	EC2 “Instances”		-
Termination action	EC2 “Instances”	EC2 “Spot Requests”	EC2 “Instances” / “Spot Requests”
Cost	Regular	Up to 90% discount	

## Different ways to stop your compute instance (and stop AWS from charging you)

	Stop	Terminate
<b>Similar to</b>	Shutting down your computer	Throwing your computer into a trash compactor
Boot EBS disk	Unaffected	Permanently Deleted
Additional EBS disks / EFS	Unaffected (AWS is still charging you for storage. Same for above)	
Ephemeral drives	Permanently Deleted	

*There is also a “hibernate” option, but our commonly used instance types do not support it.*

# Mount EFS

```
FS_ID=file-system-id REGION=us-east-2 && sudo mount  
-t nfs4 -o  
nfsvers=4.1,rsize=1048576,wsize=1048576,hard,timeo=6  
00,retrans=2,noresvport  
$FS_ID.efs.$REGION.amazonaws.com:/ ~/efs
```

- Mounts the target EFS file system on the “efs” folder in the user’s home directory

# Mount the ephemeral drive (g4dn series)

```
sudo mkfs -t xfs /dev/nvme1n1
```

```
sudo mkdir ~/data && sudo mount /dev/nvme1n1 ~/data
```

```
(cd ~/data && sudo chmod go+rw .)
```

# Advices and Tips

- Do not attempt to unzip numerous small files onto EFS. It's painfully slow.
- Optionally, use EFS to store the environment tar, data and notebook.



# Useful Links

- EC2 instance types: <https://aws.amazon.com/ec2/instance-types/>
- On-demand instance pricing: <https://aws.amazon.com/ec2/pricing/on-demand/>
- Spot instance pricing: <https://aws.amazon.com/ec2/spot/pricing/>
- My notes for using AWS: <https://chen-zhe.github.io/blog/p/aws-user-notes-for-deep-learning/>