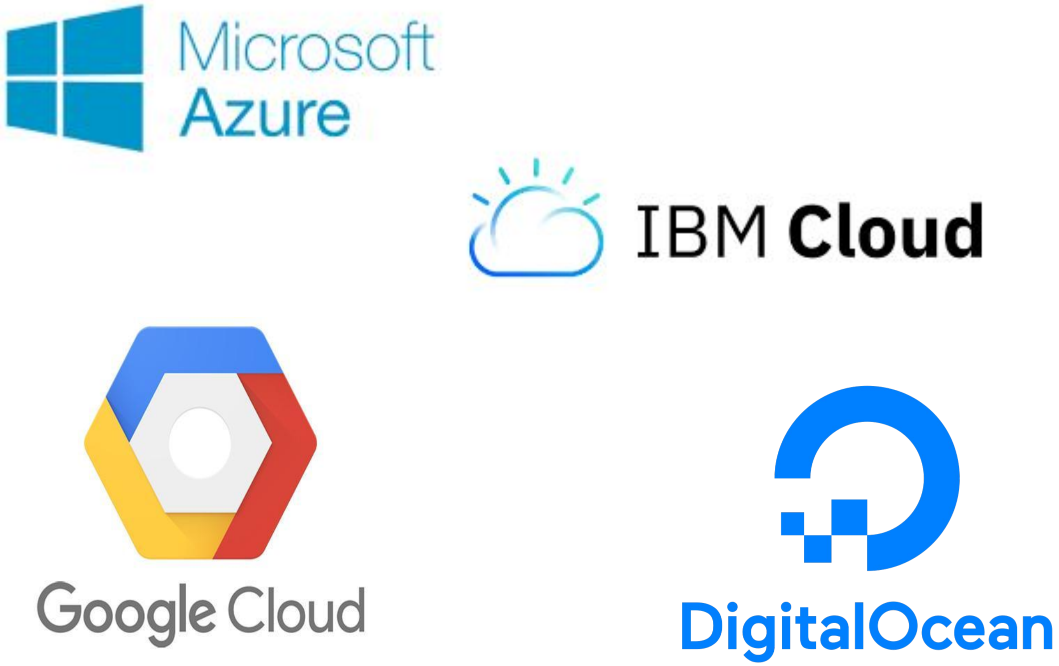

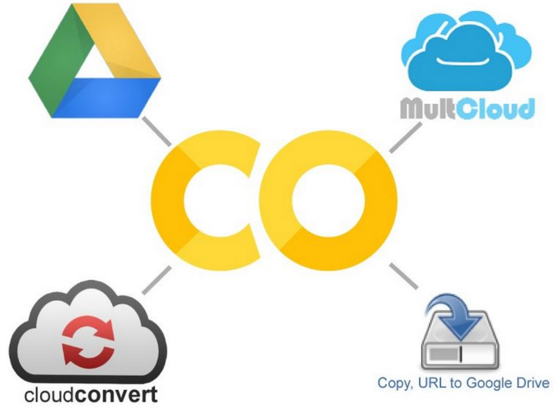

Introduction to Deep Learning
Recitation 1

What does AWS offer that will be relevant for DL?

EC2 - Compute Resources

Amazon EC2

Train the models

Jupyter Notebooks

Amazon SageMaker

S3 - Data Storage

Store training data,
models, etc

What kind of machines are available when using EC2 instances?

Instances are classified based on machine size as: nano, micro, medium, large, xlarge, 2xlarge, …, 16xlarge

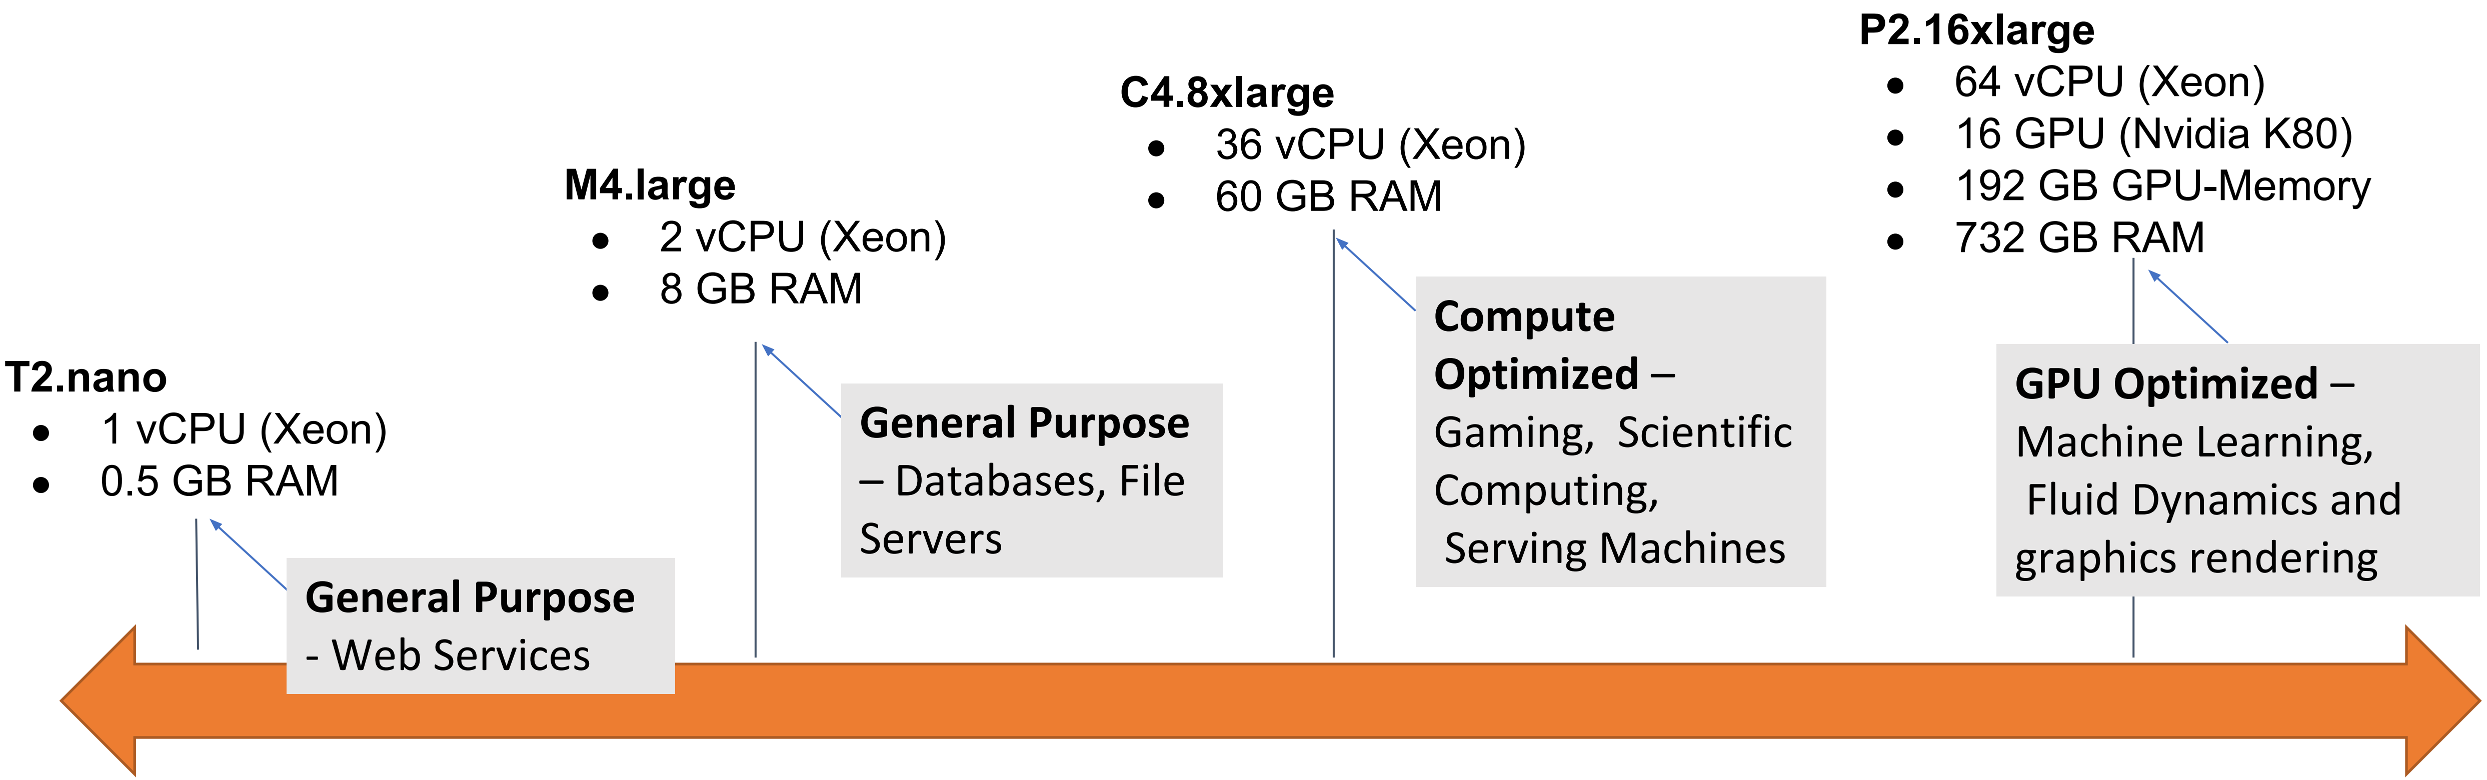

T2.nano

- 1 vCPU (Xeon)
- 0.5 GB RAM

General Purpose - Web Services

M4.large

- 2 vCPU (Xeon)
- 8 GB RAM

General Purpose – Databases, File Servers

C4.8xlarge

- 36 vCPU (Xeon)
- 60 GB RAM

Compute Optimized – Gaming, Scientific Computing, Serving Machines

P2.16xlarge

- 64 vCPU (Xeon)
- 16 GPU (Nvidia K80)
- 192 GB GPU-Memory
- 732 GB RAM

GPU Optimized – Machine Learning, Fluid Dynamics and graphics rendering

There are different types and different subtypes that you can mix and match what you need

EC2-instance types

What do we put on EC2 Instances?

- **Virtual images of existing machines**
  - You can create an image of your machine
    - Transfer it to a different machine
    - Save it as a backup
- **Use cases**
  - Software packages that are incredibly difficult to install
  - Need to create multiple different machines with the exact same data for parameters servers
  - Load balancing - create a new machine with the same AMI to be used in a different region depending on load

Why we need to know about it?

In this course we will be large amounts of data to work with neural networks and therefore:

- Anyone who doesn’t have a NVIDIA GPU
- Anyone who doesn’t have their GPU configured to be correctly used by Pytorch/Tensorflow/etc.
- GPUs are good for training **not** processing data or if the code is configured to leverage GPU

Now that you have a general idea of what you can do with

AWS

let’s start setting up an Account…

Steps you need to get started

If you sign with AWS educate for the first time, you will have an extra $100 credit to your account that might be useful later

1. Go to https://aws.amazon.com/console/

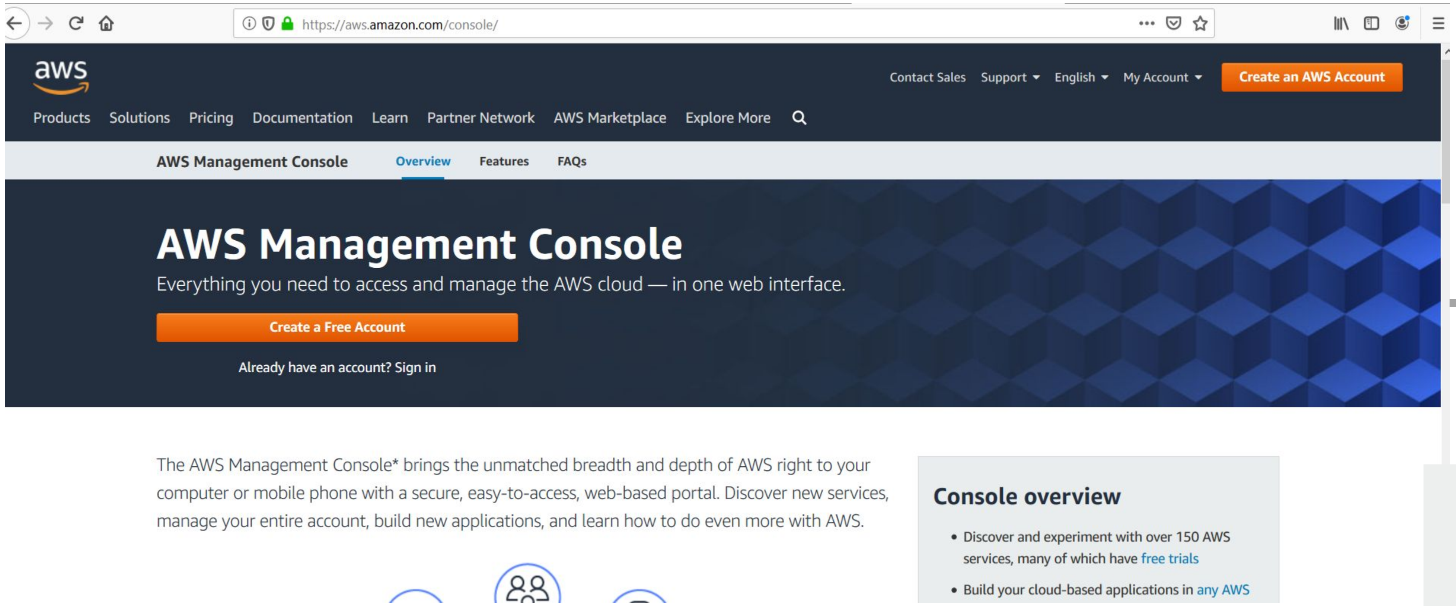

AWS Management Console

Everything you need to access and manage the AWS cloud — in one web interface.

Already have an account? Sign in

The AWS Management Console* brings the unmatched breadth and depth of AWS right to your computer or mobile phone with a secure, easy-to-access, web-based portal. Discover new services, manage your entire account, build new applications, and learn how to do even more with AWS.

Console overview

- Discover and experiment with over 150 AWS services, many of which have free trials
- Build your cloud-based applications in any AWS

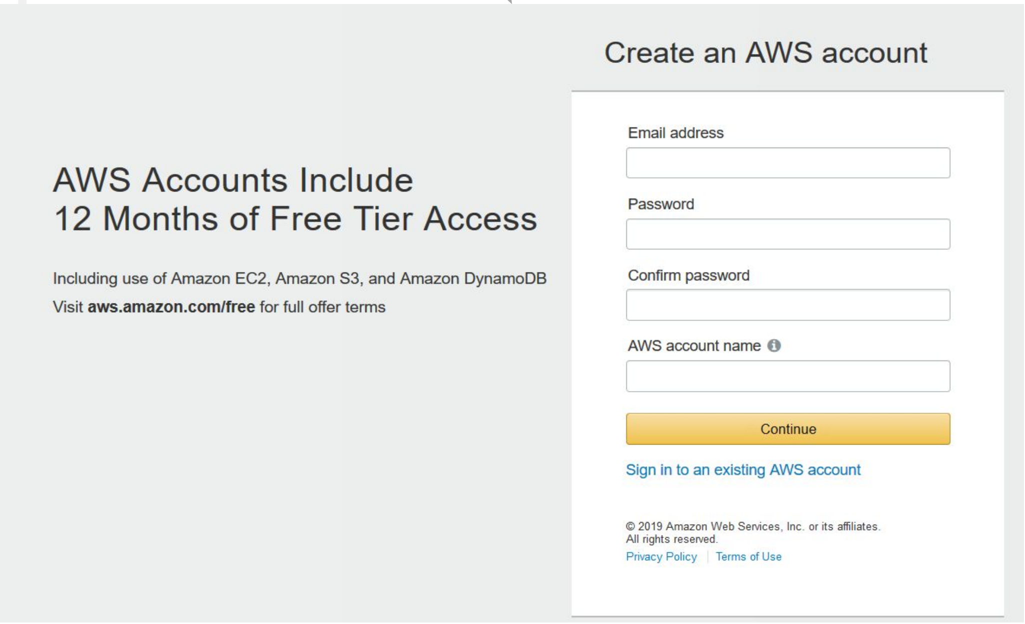

Create an AWS account

AWS Accounts Include
12 Months of Free Tier Access

Including use of Amazon EC2, Amazon S3, and Amazon DynamoDB
Visit **aws.amazon.com/free** for full offer terms

Email address

Password

Confirm password

AWS account name

© 2019 Amazon Web Services, Inc. or its affiliates.
All rights reserved.
Privacy Policy | Terms of Use

Subsequent screens you must fill:

- Personal data
- Payment information

2. Choose Region and then click on Services

2

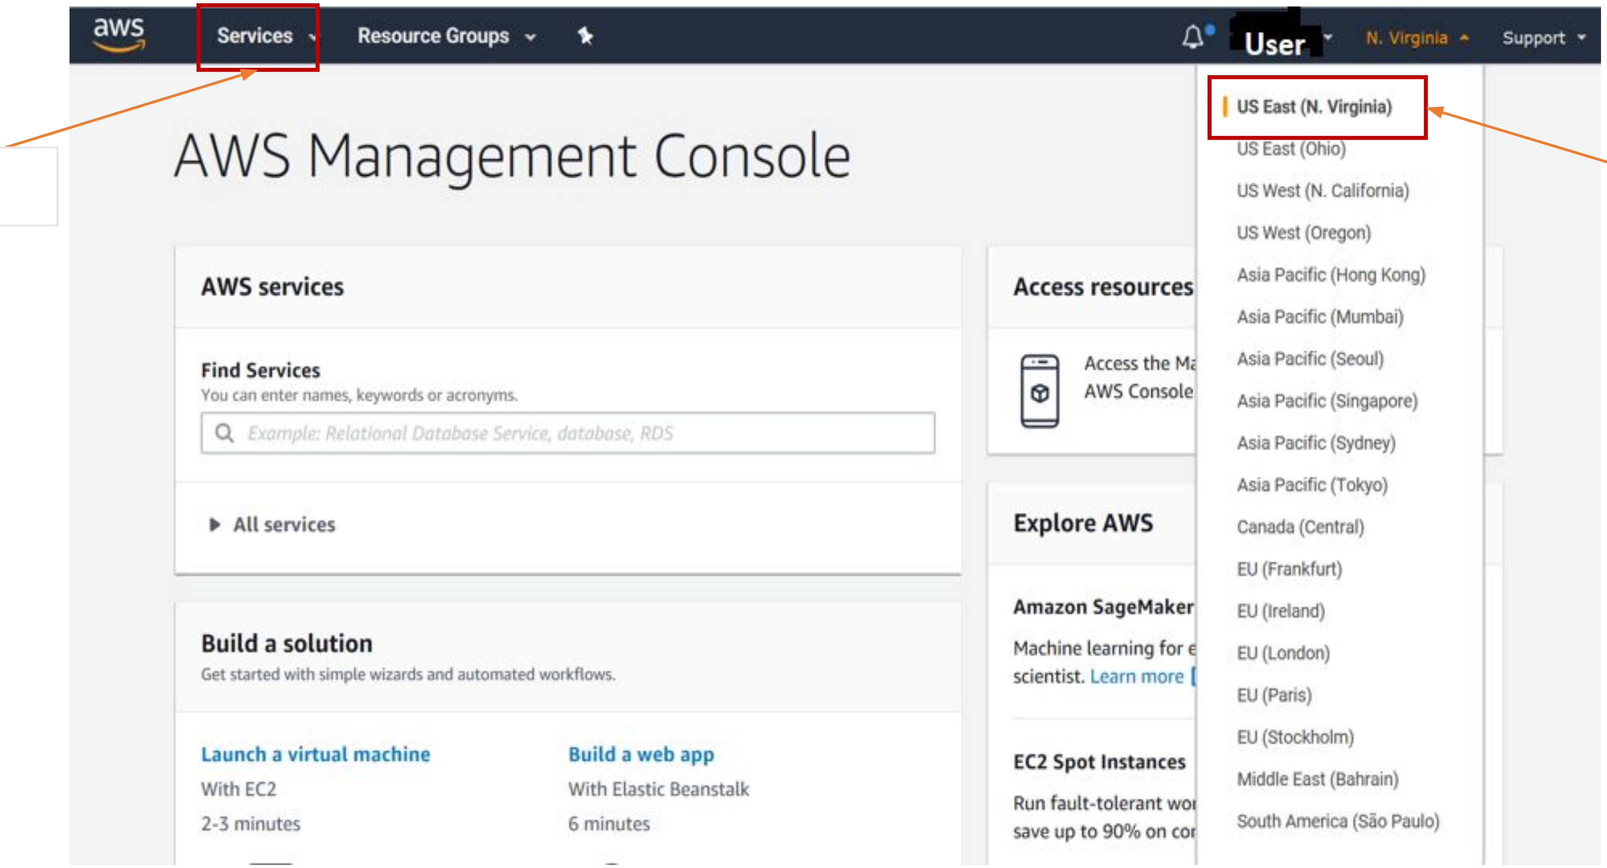

AWS Management Console

AWS services

Find Services

You can enter names, keywords or acronyms.

Example: Relational Database Service, database, RDS

▸ All services

Build a solution

Get started with simple wizards and automated workflows.

Launch a virtual machine

With EC2

2-3 minutes

Build a web app

With Elastic Beanstalk

6 minutes

Access resources

Access the Ma

AWS Console

Explore AWS

Amazon SageMaker

Machine learning for e

scientist. Learn more

EC2 Spot Instances

Run fault-tolerant wo

save up to 90% on co

1 Choose **N. Virginia or Ohio for US-East or Oregon for US-West Coast** as your location, since only certain locations have instances with GPUs

-

3. In Services

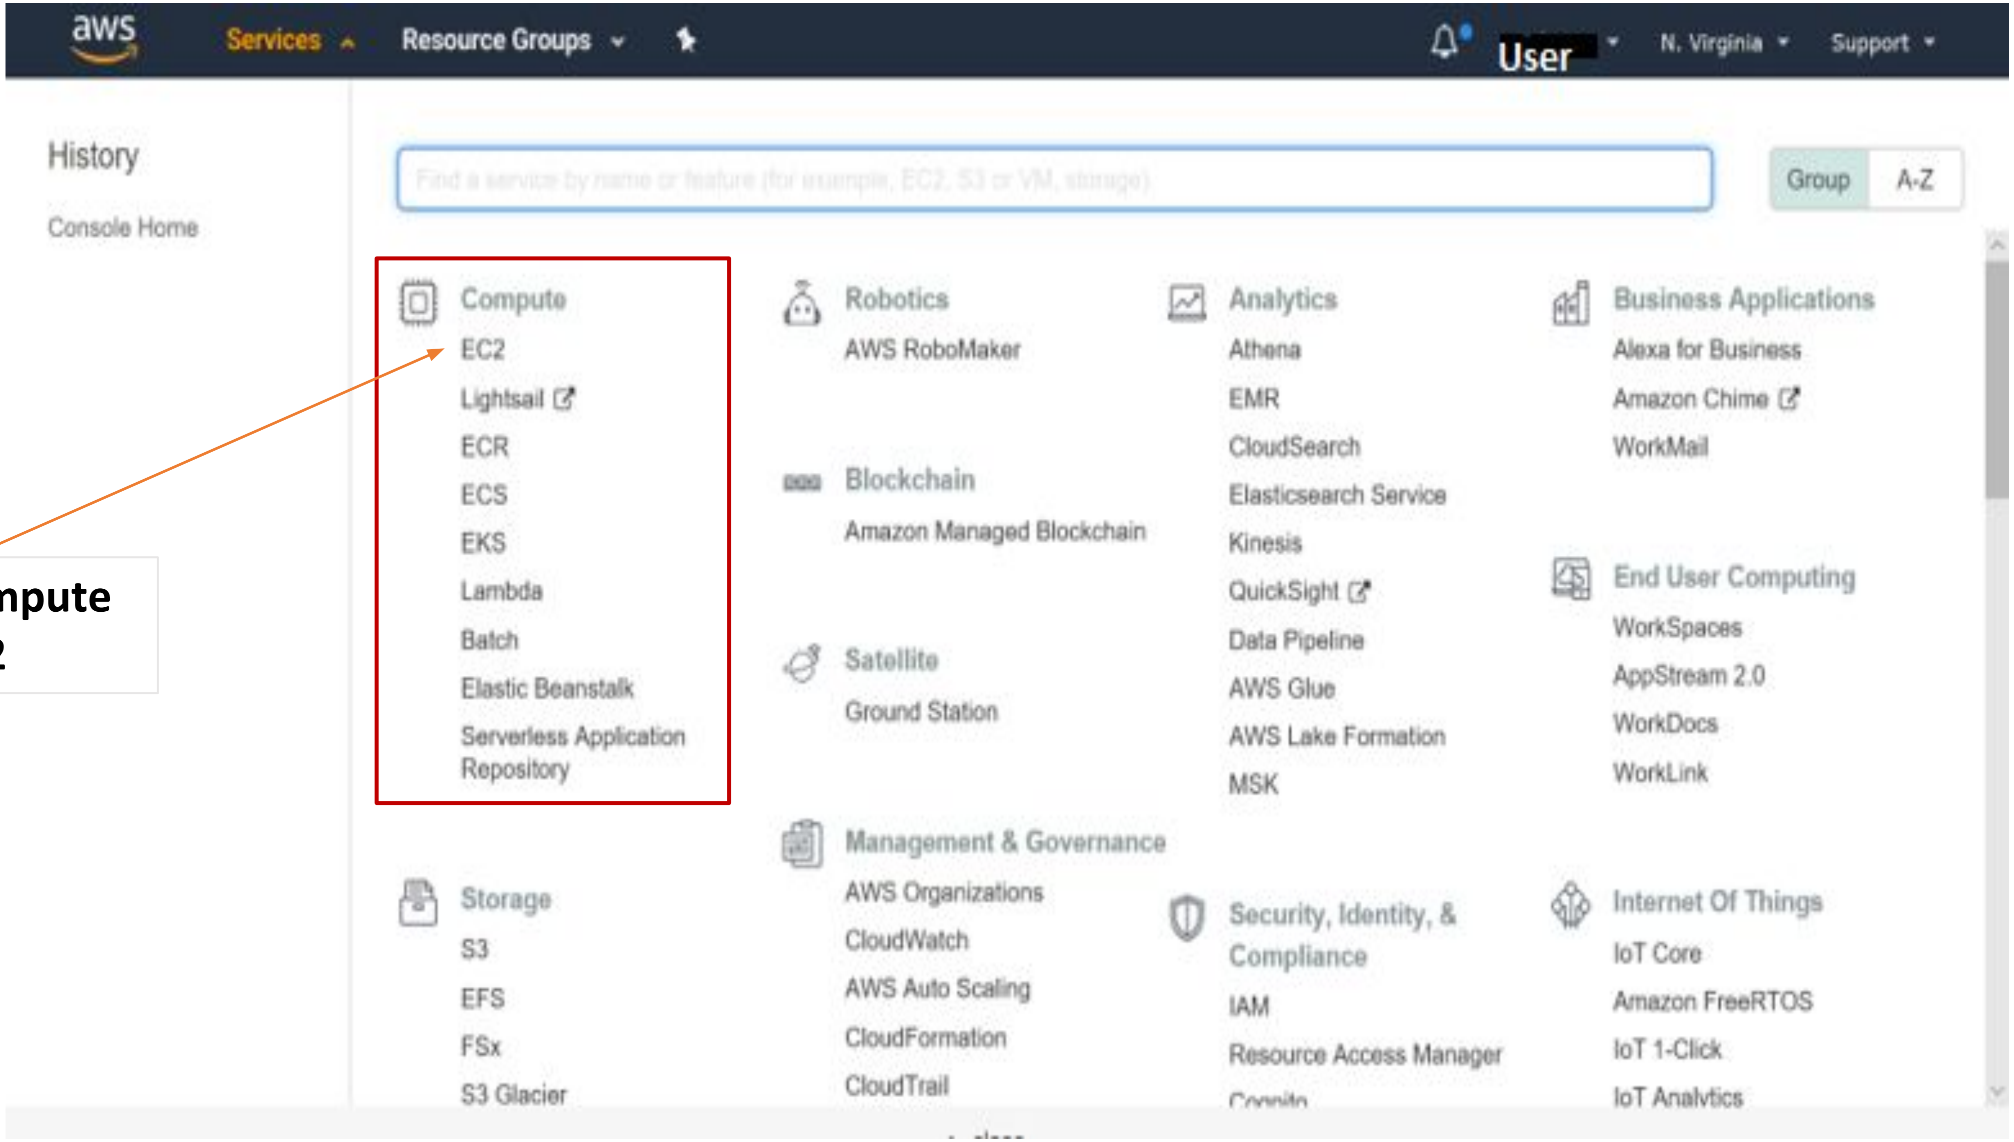

Under **Compute**
 - Click **EC2**

4. Setting key pairs to connect to instances

2 Click on Running instances if you want to launch a new instance, once launch it will appear in the **EC2 Dashboard**

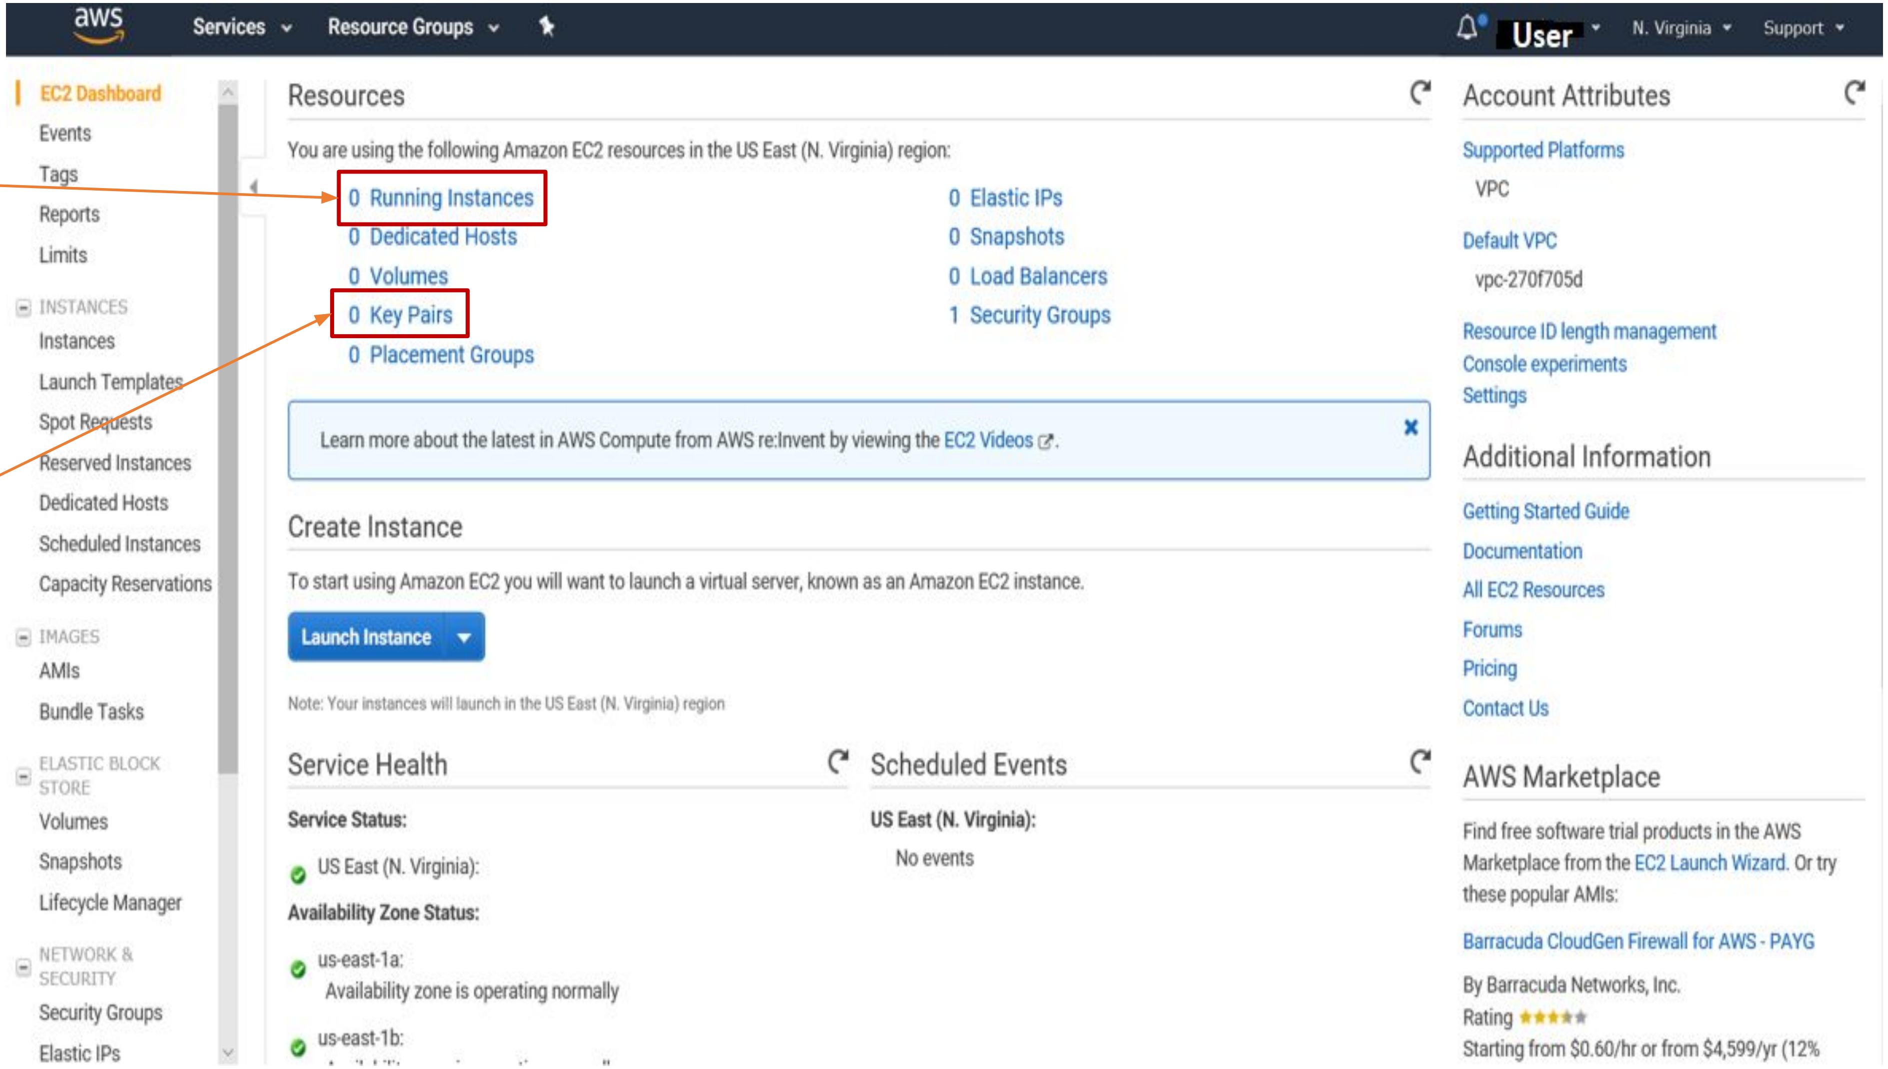

Resources

You are using the following Amazon EC2 resources in the US East (N. Virginia) region:

0 Running Instances	0 Elastic IPs
0 Dedicated Hosts	0 Snapshots
0 Volumes	0 Load Balancers
0 Key Pairs	1 Security Groups
0 Placement Groups	

Learn more about the latest in AWS Compute from AWS re:Invent by viewing the EC2 Videos.

Create Instance

To start using Amazon EC2 you will want to launch a virtual server, known as an Amazon EC2 instance.

Note: Your instances will launch in the US East (N. Virginia) region

Service Health

Service Status:

US East (N. Virginia):

Availability Zone Status:

us-east-1a:
Availability zone is operating normally

us-east-1b:

Scheduled Events

US East (N. Virginia):
No events

Account Attributes

Supported Platforms
VPC

Default VPC
vpc-270f705d

Resource ID length management
Console experiments
Settings

Additional Information

Getting Started Guide
Documentation
All EC2 Resources
Forums
Pricing
Contact Us

AWS Marketplace

Find free software trial products in the AWS Marketplace from the EC2 Launch Wizard. Or try these popular AMIs:

Barracuda CloudGen Firewall for AWS - PAYG

By Barracuda Networks, Inc.
Rating ★★★★★
Starting from $0.60/hr or from $4,599/yr (12%

1 Set your key pair, so you can connect to your instance. Once is set, it will appear in the **EC2 Dashboard**

4.1 Setting key pairs

1 You create a key pair, and once created a .pem file will download on your machine, which is the key of your connection to AWS servers

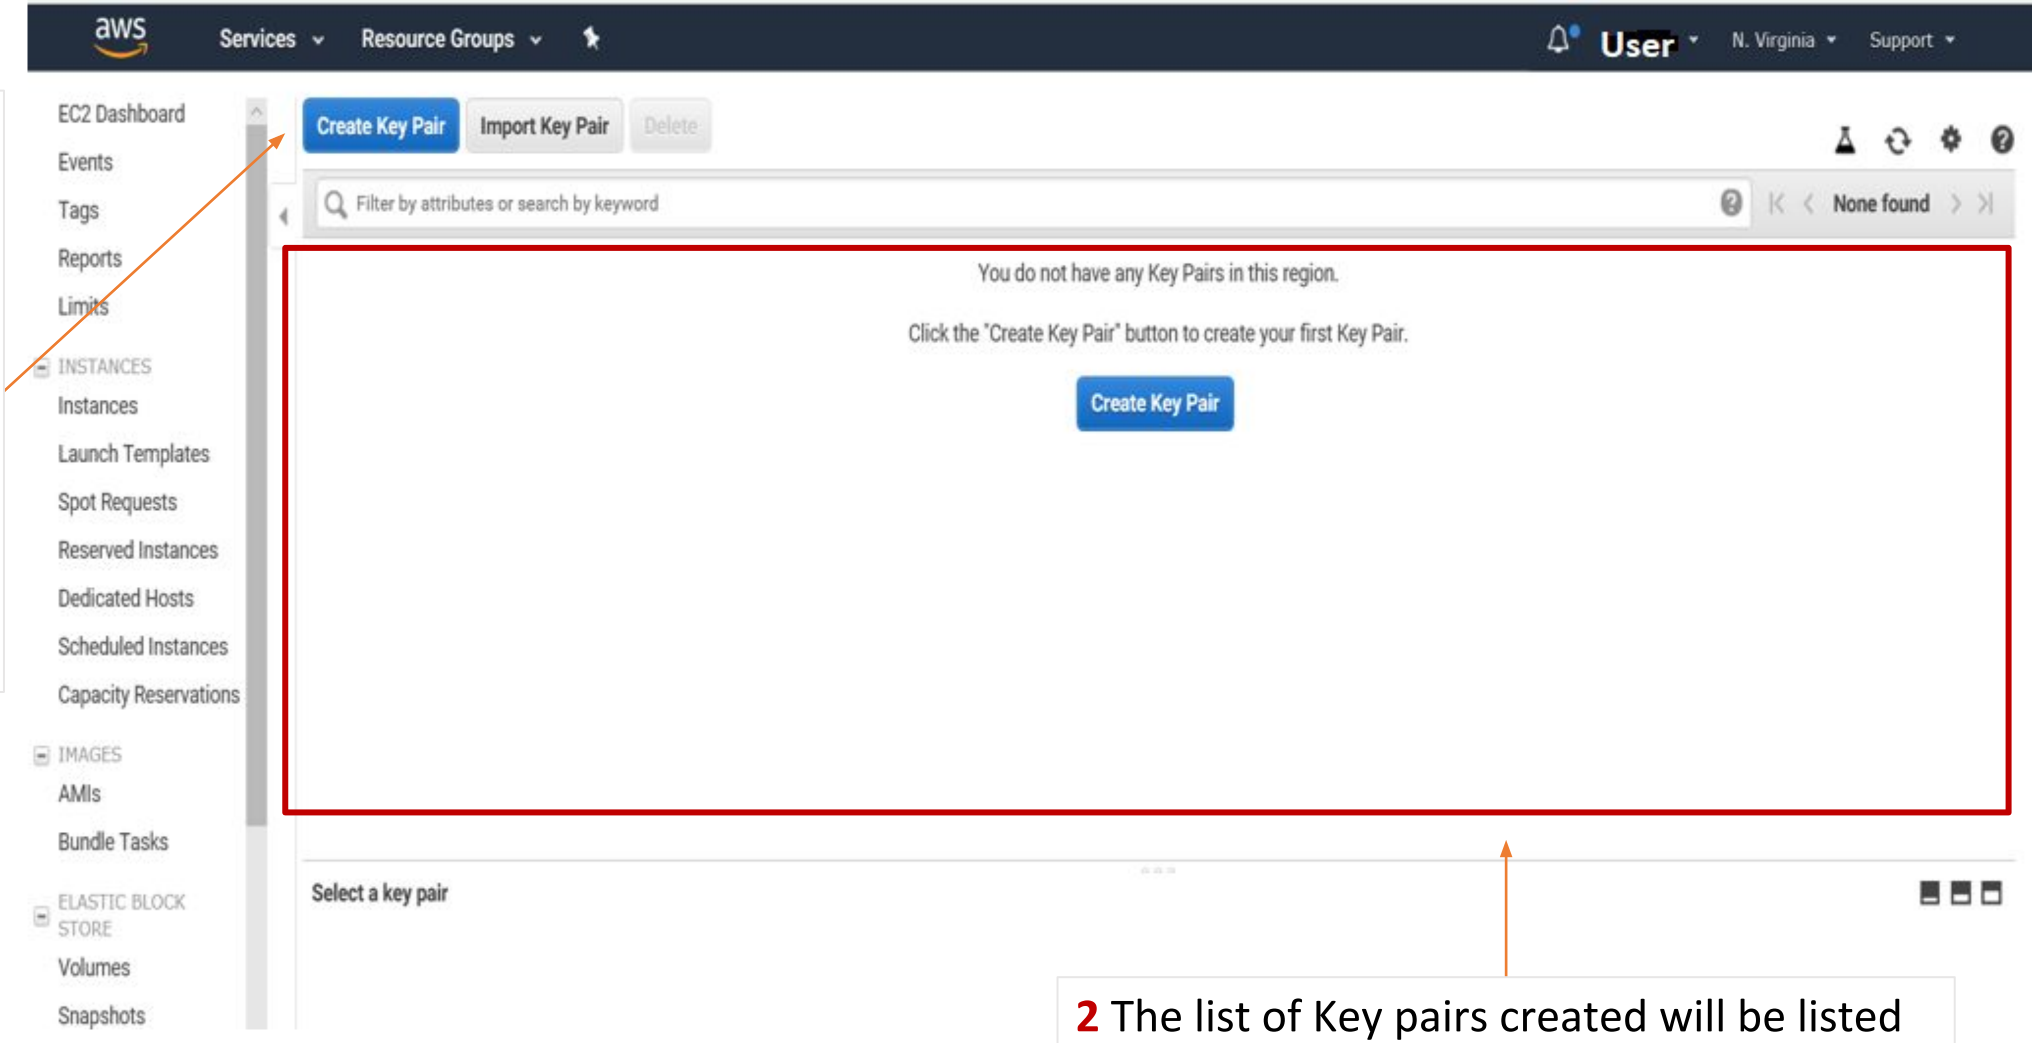

You do not have any Key Pairs in this region.

Click the "Create Key Pair" button to create your first Key Pair.

2 The list of Key pairs created will be listed on this space

4.2.2 Launch an Instance

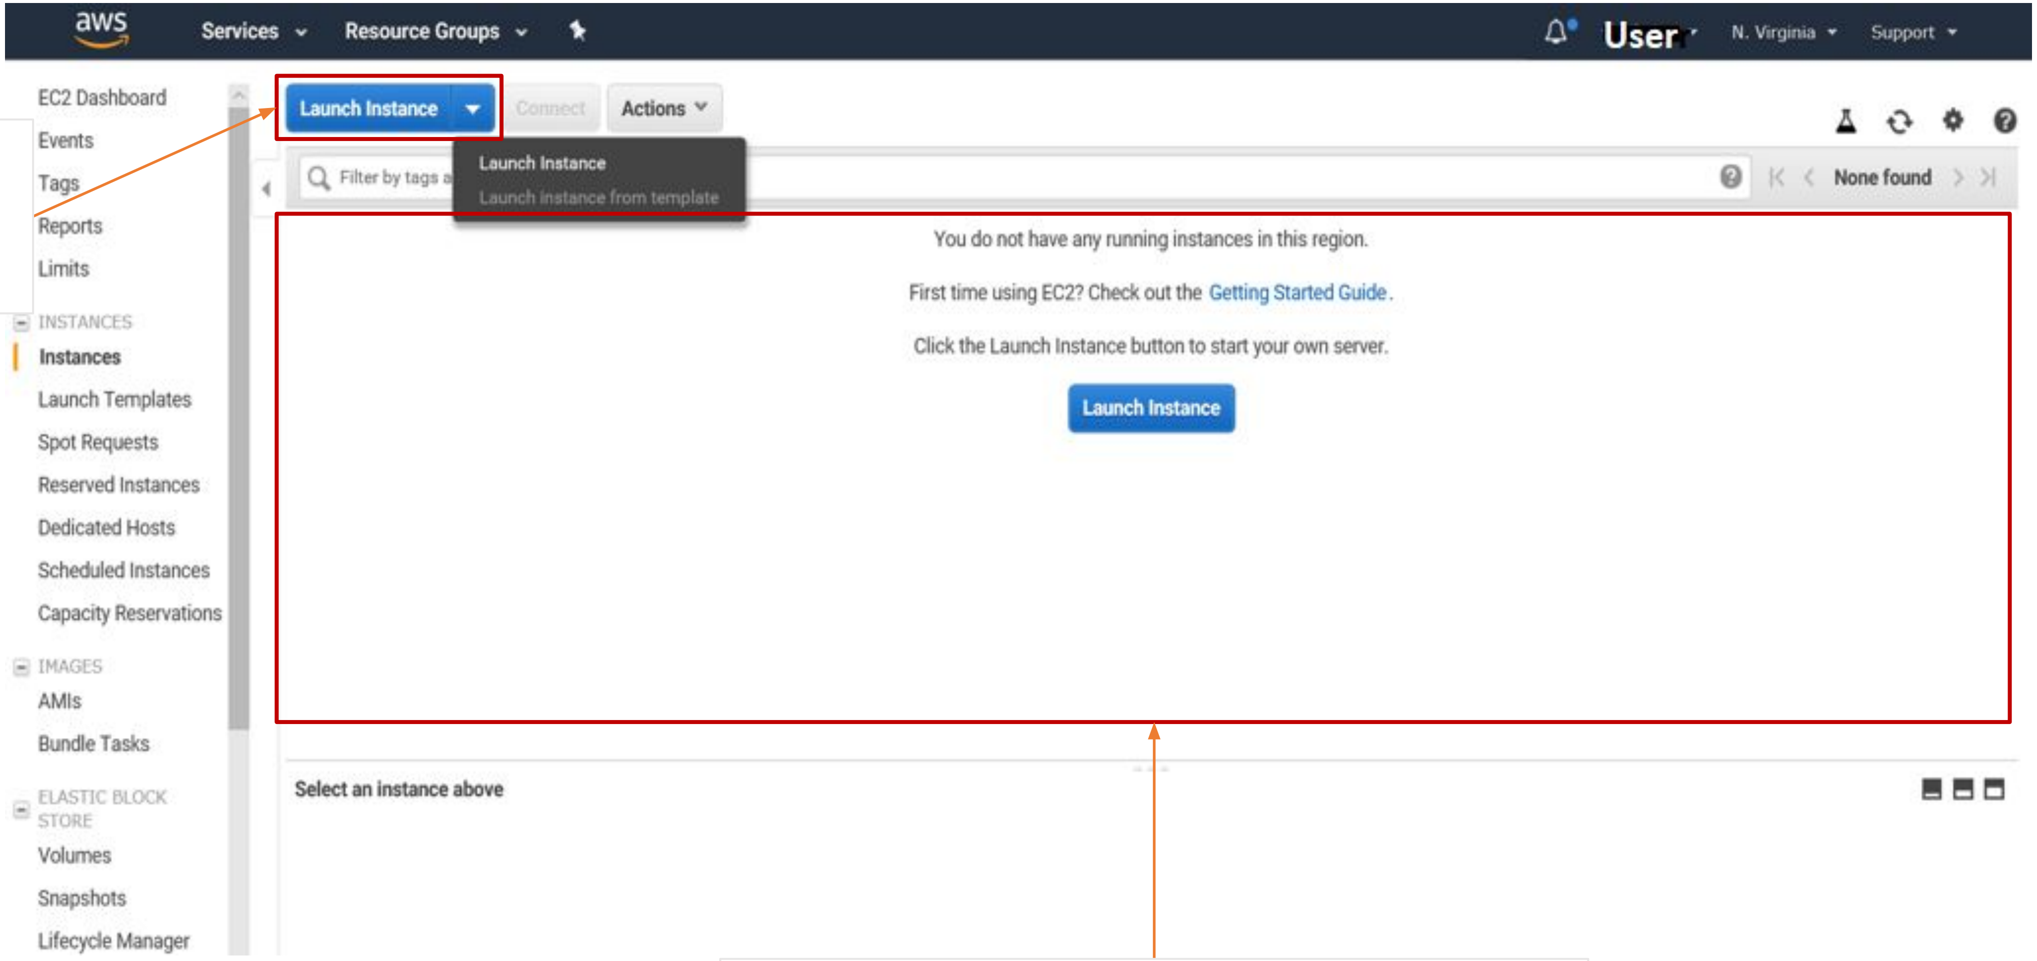

You do not have any running instances in this region.

First time using EC2? Check out the Getting Started Guide.

Click the Launch Instance button to start your own server.

1 Click on "**Launch instance**" to create a new instance

2 Set your key pair, so you can connect to your instance. Once is set, it will appear in the **EC2 Dashboard**

4.2.2 Launch an Instance (Continued)

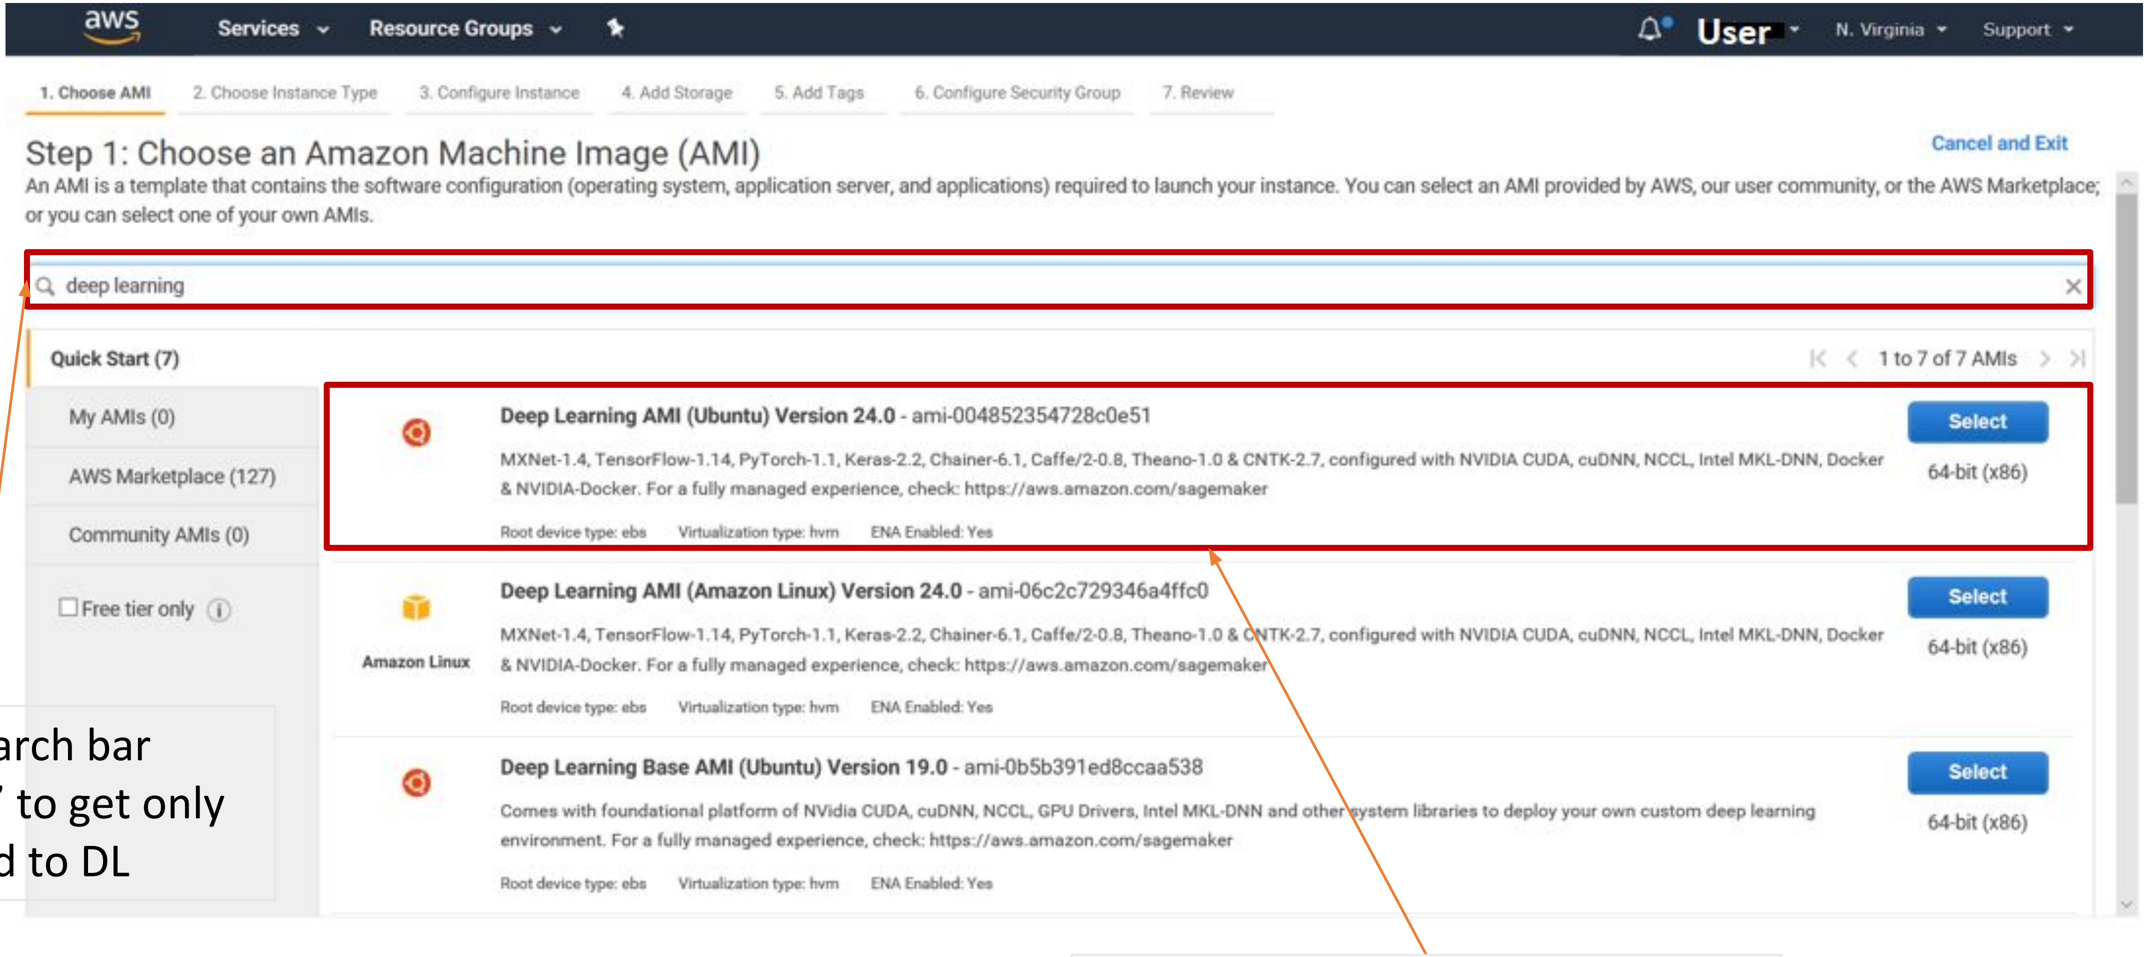

Step 1: Choose an Amazon Machine Image (AMI)

An AMI is a template that contains the software configuration (operating system, application server, and applications) required to launch your instance. You can select an AMI provided by AWS, our user community, or the AWS Marketplace; or you can select one of your own AMIs.

deep learning

Deep Learning AMI (Ubuntu) Version 24.0 - ami-004852354728c0e51

MXNet-1.4, TensorFlow-1.14, PyTorch-1.1, Keras-2.2, Chainer-6.1, Caffe/2-0.8, Theano-1.0 & CNTK-2.7, configured with NVIDIA CUDA, cuDNN, NCCL, Intel MKL-DNN, Docker & NVIDIA-Docker. For a fully managed experience, check: https://aws.amazon.com/sagemaker

Root device type: ebs Virtualization type: hvm ENA Enabled: Yes

64-bit (x86)

Deep Learning AMI (Amazon Linux) Version 24.0 - ami-06c2c729346a4ffc0

MXNet-1.4, TensorFlow-1.14, PyTorch-1.1, Keras-2.2, Chainer-6.1, Caffe/2-0.8, Theano-1.0 & CNTK-2.7, configured with NVIDIA CUDA, cuDNN, NCCL, Intel MKL-DNN, Docker & NVIDIA-Docker. For a fully managed experience, check: https://aws.amazon.com/sagemaker

Root device type: ebs Virtualization type: hvm ENA Enabled: Yes

64-bit (x86)

Deep Learning Base AMI (Ubuntu) Version 19.0 - ami-0b5b391ed8ccaa538

Comes with foundational platform of NVidia CUDA, cuDNN, NCCL, GPU Drivers, Intel MKL-DNN and other system libraries to deploy your own custom deep learning environment. For a fully managed experience, check: https://aws.amazon.com/sagemaker

Root device type: ebs Virtualization type: hvm ENA Enabled: Yes

64-bit (x86)

1 Type in the search bar "Deep Learning" to get only instances related to DL

2 Select Deep Learning AMI with Ubuntu containing Pytorch in it.

5. Select the Instance type

For the purpose of this tutorial we will use t2-micro as it is “**free tier eligible**”, the process will be the same for instances with GPUs, but we will explain that in the next slides

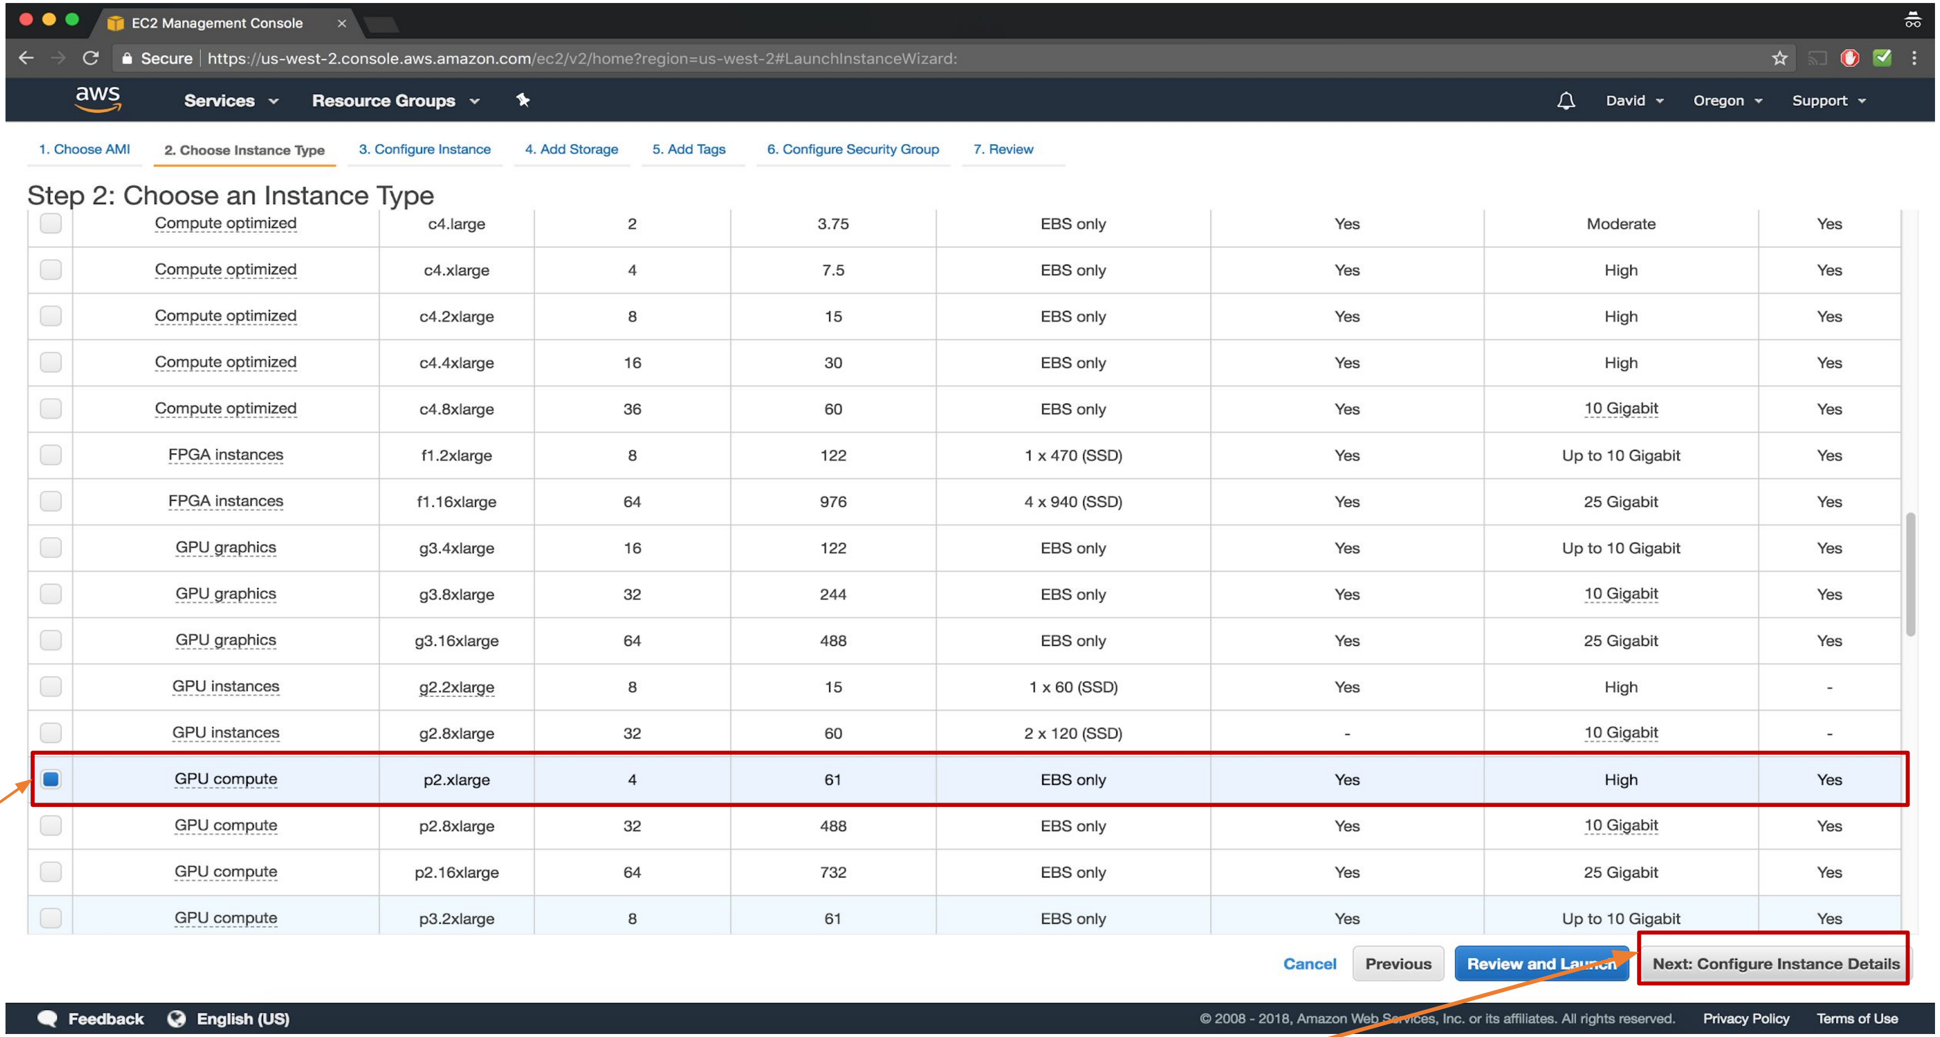
EC2 Management Console
Secure | https://us-west-2.console.aws.amazon.com/ec2/v2/home?region=us-west-2#LaunchInstanceWizard:
aws
Services
Resource Groups
David
Oregon
Support
1. Choose AMI
2. Choose Instance Type
3. Configure Instance
4. Add Storage
5. Add Tags
6. Configure Security Group
7. Review
Cancel
Previous
Review and Launch
Next: Configure Instance Details
Feedback
English (US)
© 2008 - 2018, Amazon Web Services, Inc. or its affiliates. All rights reserved.
Privacy Policy
Terms of Use

Step 2: Choose an Instance Type

□	Compute optimized	c4.large	2	3.75	EBS only	Yes	Moderate	Yes
□	Compute optimized	c4.xlarge	4	7.5	EBS only	Yes	High	Yes
□	Compute optimized	c4.2xlarge	8	15	EBS only	Yes	High	Yes
□	Compute optimized	c4.4xlarge	16	30	EBS only	Yes	High	Yes
□	Compute optimized	c4.8xlarge	36	60	EBS only	Yes	10 Gigabit	Yes
□	FPGA instances	f1.2xlarge	8	122	1 x 470 (SSD)	Yes	Up to 10 Gigabit	Yes
□	FPGA instances	f1.16xlarge	64	976	4 x 940 (SSD)	Yes	25 Gigabit	Yes
□	GPU graphics	g3.4xlarge	16	122	EBS only	Yes	Up to 10 Gigabit	Yes
□	GPU graphics	g3.8xlarge	32	244	EBS only	Yes	10 Gigabit	Yes
□	GPU graphics	g3.16xlarge	64	488	EBS only	Yes	25 Gigabit	Yes
□	GPU instances	g2.2xlarge	8	15	1 x 60 (SSD)	Yes	High	-
□	GPU instances	g2.8xlarge	32	60	2 x 120 (SSD)	-	10 Gigabit	-
☑	GPU compute	p2.xlarge	4	61	EBS only	Yes	High	Yes
□	GPU compute	p2.8xlarge	32	488	EBS only	Yes	10 Gigabit	Yes
□	GPU compute	p2.16xlarge	64	732	EBS only	Yes	25 Gigabit	Yes
□	GPU compute	p3.2xlarge	8	61	EBS only	Yes	Up to 10 Gigabit	Yes

1 Select **p2.xlarge**

2 after selecting the instance type, hit next

5.1 You Will Need Permission to Get an Instance with a GPU

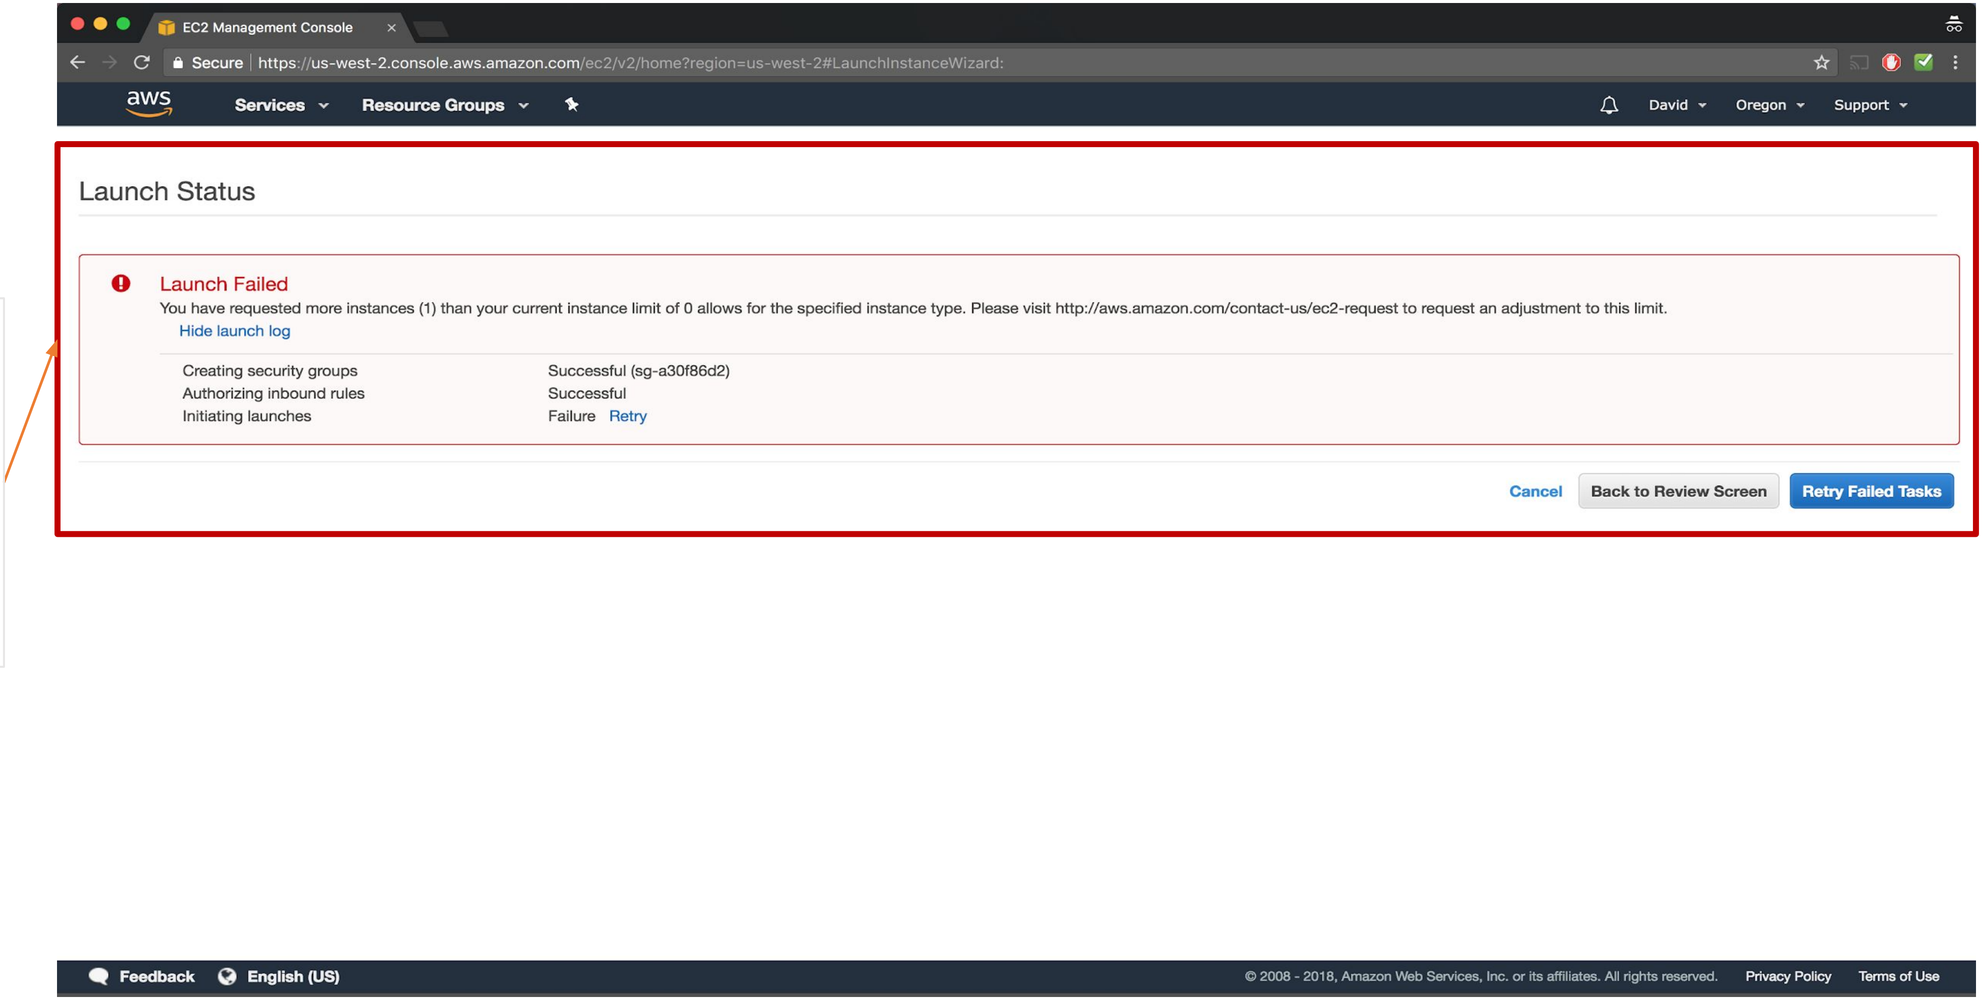

EC2 Management Console
Secure | https://us-west-2.console.aws.amazon.com/ec2/v2/home?region=us-west-2#LaunchInstanceWizard:
aws
Services
Resource Groups
David
Oregon
Support

Launch Status

Launch Failed

You have requested more instances (1) than your current instance limit of 0 allows for the specified instance type. Please visit http://aws.amazon.com/contact-us/ec2-request to request an adjustment to this limit.

Hide launch log

```
Creating security groups       Successful (sg-a30f86d2)
Authorizing inbound rules      Successful
Initiating launches            Failure  Retry
```

If you haven't requested an AWS instance with GPU before, you will run into a launch error like this

Visit the link provided by the error page in the previous slide https:aws.amazon.com/contact-us/ec2-request …

5.1.1 Follow the Instructions and…

Fill out the information requested on the ticket

You will need to provide a reason as to why you are requesting a GPU. You can customize the following: “***I need GPUs for some Deep Learning projects that I do on the side and don’t have a GPU of my own. Will be using it for Kaggle competitions***”

Make sure to provide a contact information, so they can contact you to set approve your request

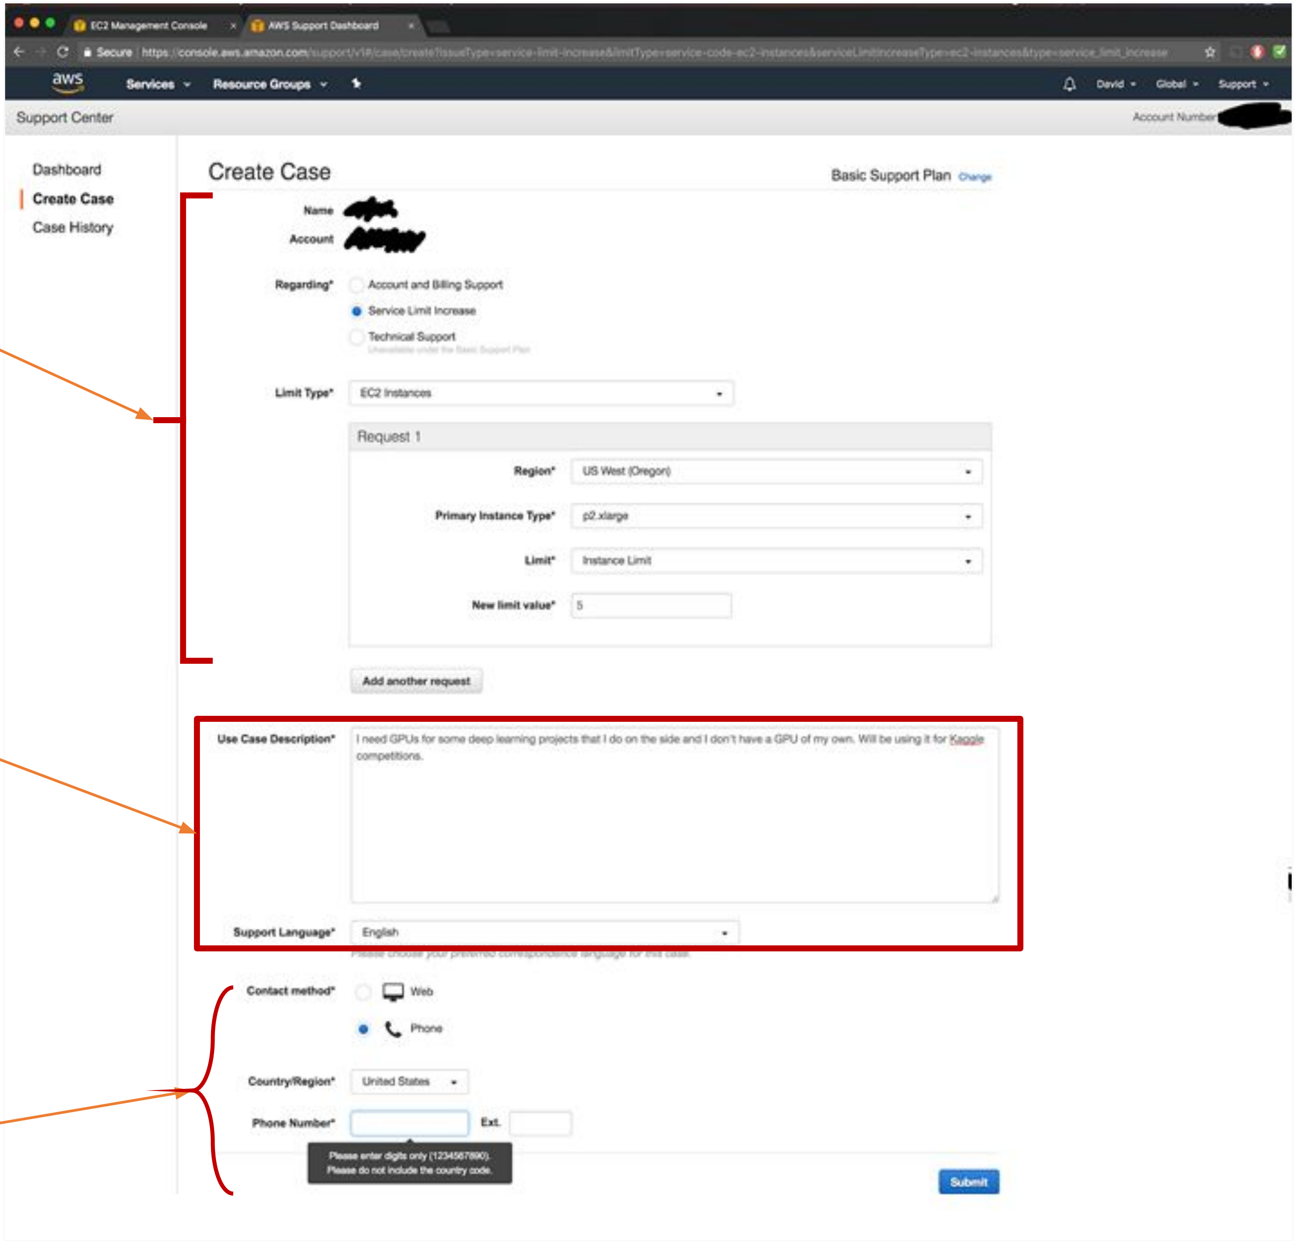
EC2 Management Console
AWS Support Dashboard
Secure https://console.aws.amazon.com/support/v1#/case/create?issueType=service-limit-increase&limitType=service-code-ec2-instances&serviceLimitIncreaseType=ec2-instances&type=service_limit_increase
aws
Services
Resource Groups
David
Global
Support
Support Center
Dashboard
Create Case
Case History

Account Number [redacted]

Create Case

Basic Support Plan Change

Name [redacted]

Account [redacted]

Regarding*
○ Account and Billing Support
● Service Limit Increase
○ Technical Support
Unavailable under the Basic Support Plan

Limit Type* EC2 Instances

Request 1

Region* US West (Oregon)

Primary Instance Type* p2.xlarge

Limit* Instance Limit

New limit value* 5

Add another request

Use Case Description* I need GPUs for some deep learning projects that I do on the side and I don't have a GPU of my own. Will be using it for Kaggle competitions.

Support Language* English
Please choose your preferred correspondence language for this case.

Contact method*
○ Web
● Phone

Country/Region* United States

Phone Number* Ext.

Please enter digits only (1234567890).
Please do not include the country code.

Submit

6. Configuring the Instance

After getting access to the instance with GPUs…

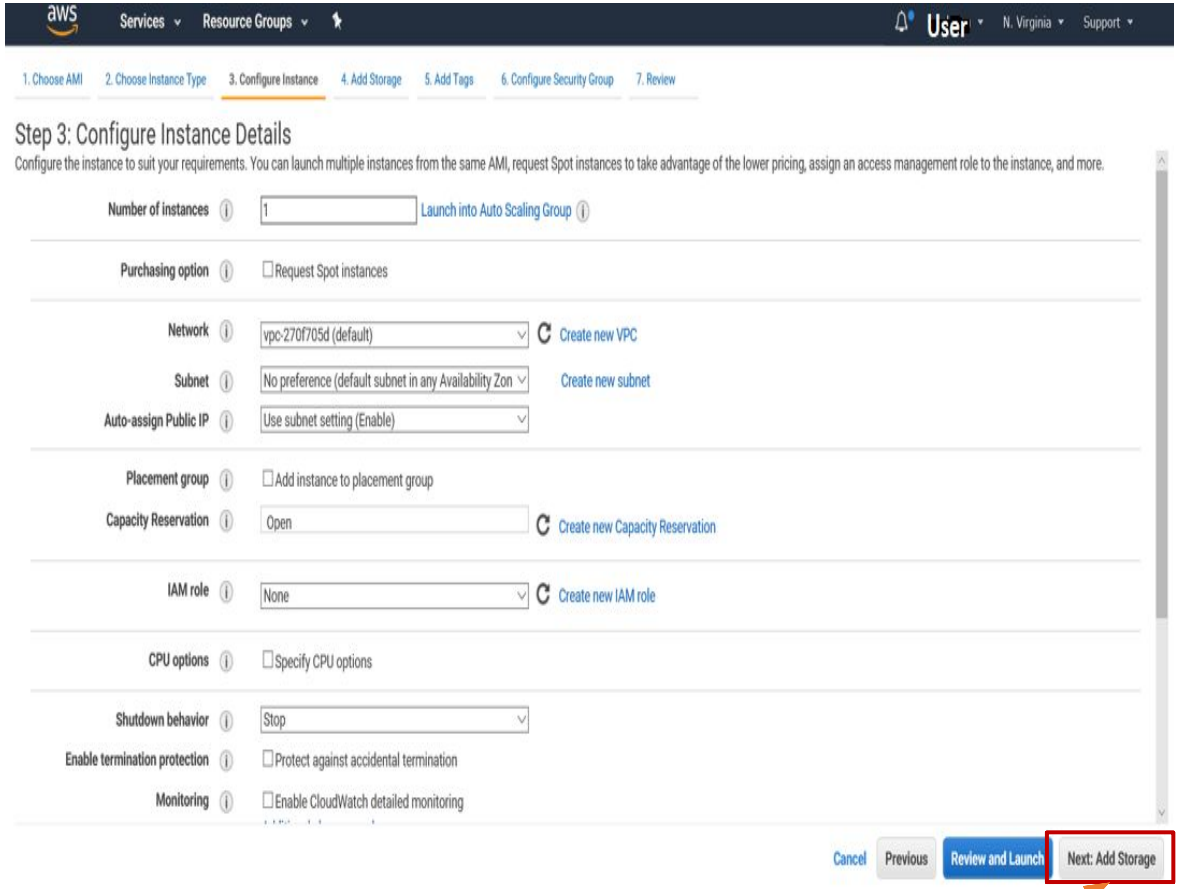

1. Choose AMI 2. Choose Instance Type 3. Configure Instance 4. Add Storage 5. Add Tags 6. Configure Security Group 7. Review

Step 3: Configure Instance Details

Configure the instance to suit your requirements. You can launch multiple instances from the same AMI, request Spot instances to take advantage of the lower pricing, assign an access management role to the instance, and more.

Number of instances	1 Launch into Auto Scaling Group
Purchasing option	□ Request Spot instances
Network	vpc-270f705d (default) Create new VPC
Subnet	No preference (default subnet in any Availability Zon Create new subnet
Auto-assign Public IP	Use subnet setting (Enable)
Placement group	□ Add instance to placement group
Capacity Reservation	Open Create new Capacity Reservation
IAM role	None Create new IAM role
CPU options	□ Specify CPU options
Shutdown behavior	Stop
Enable termination protection	□ Protect against accidental termination
Monitoring	□ Enable CloudWatch detailed monitoring

1 You will see "**Configure Instance Details**", we don't change the configuration and instead just hit **Next**

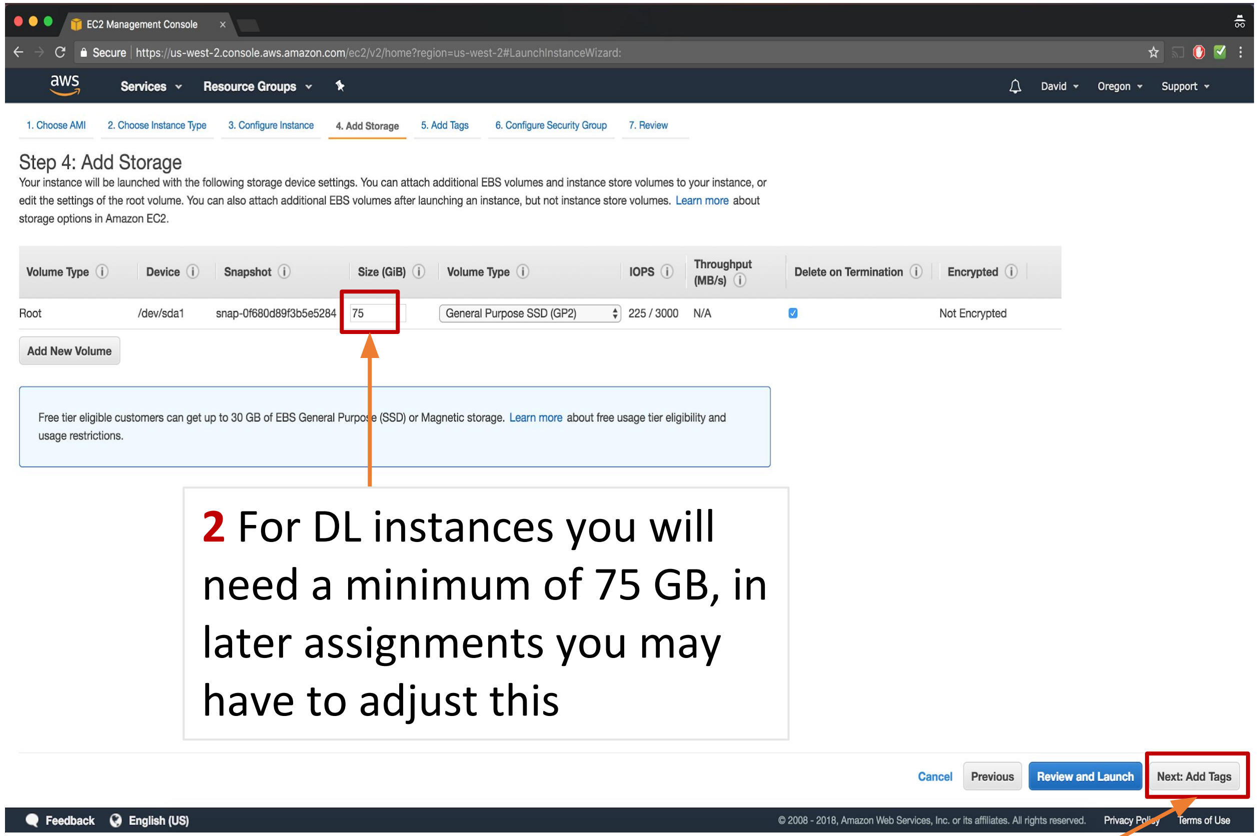

Secure | https://us-west-2.console.aws.amazon.com/ec2/v2/home?region=us-west-2#LaunchInstanceWizard:

1. Choose AMI 2. Choose Instance Type 3. Configure Instance 4. Add Storage 5. Add Tags 6. Configure Security Group 7. Review

Step 4: Add Storage

Your instance will be launched with the following storage device settings. You can attach additional EBS volumes and instance store volumes to your instance, or edit the settings of the root volume. You can also attach additional EBS volumes after launching an instance, but not instance store volumes. Learn more about storage options in Amazon EC2.

Volume Type	Device	Snapshot	Size (GiB)	Volume Type	IOPS	Throughput (MB/s)	Delete on Termination	Encrypted
Root	/dev/sda1	snap-0f680d89f3b5e5284	75	General Purpose SSD (GP2)	225 / 3000	N/A	☑	Not Encrypted

Free tier eligible customers can get up to 30 GB of EBS General Purpose (SSD) or Magnetic storage. Learn more about free usage tier eligibility and usage restrictions.

2 For DL instances you will need a minimum of 75 GB, in later assignments you may have to adjust this

3 Then just hit **Next**

Note.- Steps 5 and 6 of the configuration we don't really use them, so just hit **next**

6.1 Configuring the Instance

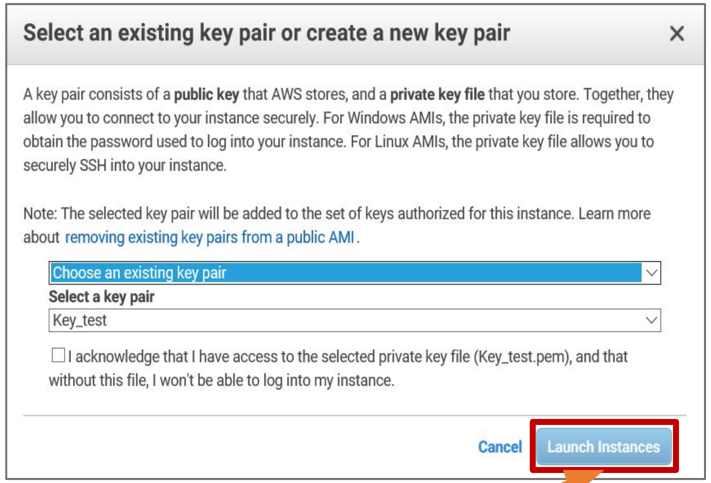

Select an existing key pair or create a new key pair ✕

A key pair consists of a **public key** that AWS stores, and a **private key file** that you store. Together, they allow you to connect to your instance securely. For Windows AMIs, the private key file is required to obtain the password used to log into your instance. For Linux AMIs, the private key file allows you to securely SSH into your instance.

Note: The selected key pair will be added to the set of keys authorized for this instance. Learn more about removing existing key pairs from a public AMI .

Choose an existing key pair

Select a key pair

Key_test

□ I acknowledge that I have access to the selected private key file (Key_test.pem), and that without this file, I won't be able to log into my instance.

1 Select the Key pair created on step **4.1** and click **Launch Instance**

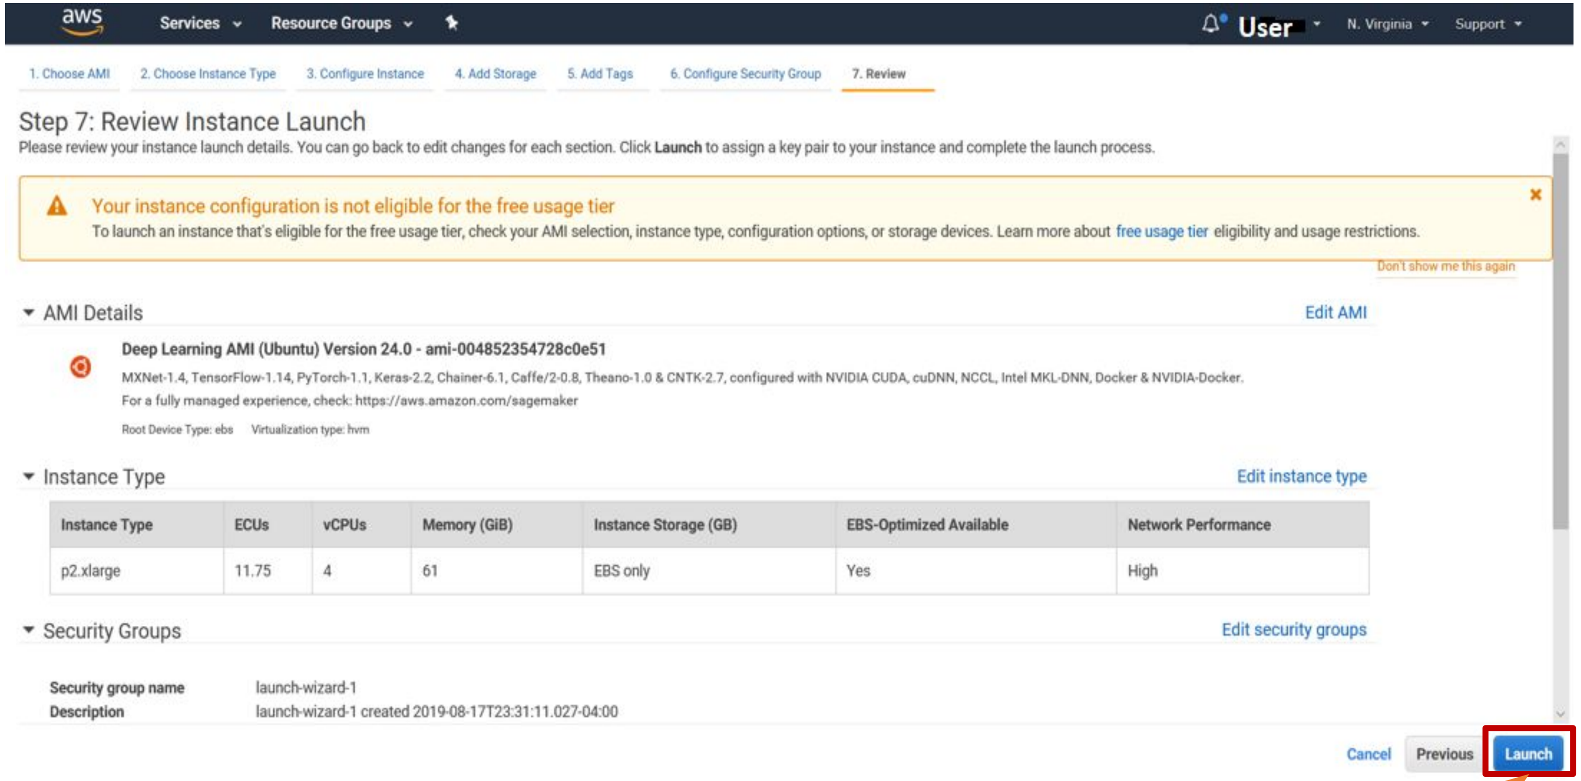

Step 7: Review Instance Launch

Please review your instance launch details. You can go back to edit changes for each section. Click **Launch** to assign a key pair to your instance and complete the launch process.

⚠ Your instance configuration is not eligible for the free usage tier

To launch an instance that's eligible for the free usage tier, check your AMI selection, instance type, configuration options, or storage devices. Learn more about free usage tier eligibility and usage restrictions.

Don't show me this again

▾ AMI Details — Edit AMI

Deep Learning AMI (Ubuntu) Version 24.0 - ami-004852354728c0e51

MXNet-1.4, TensorFlow-1.14, PyTorch-1.1, Keras-2.2, Chainer-6.1, Caffe/2-0.8, Theano-1.0 & CNTK-2.7, configured with NVIDIA CUDA, cuDNN, NCCL, Intel MKL-DNN, Docker & NVIDIA-Docker.

For a fully managed experience, check: https://aws.amazon.com/sagemaker

Root Device Type: ebs Virtualization type: hvm

▾ Instance Type — Edit instance type

Instance Type	ECUs	vCPUs	Memory (GiB)	Instance Storage (GB)	EBS-Optimized Available	Network Performance
p2.xlarge	11.75	4	61	EBS only	Yes	High

▾ Security Groups — Edit security groups

Security group name: launch-wizard-1

Description: launch-wizard-1 created 2019-08-17T23:31:11.027-04:00

Steps 5 and 6 of the configuration we don't really use them, so just hit **next** and finally hit **Launch**

7. Launched Instance

Once the instance has been configured and launch successfully

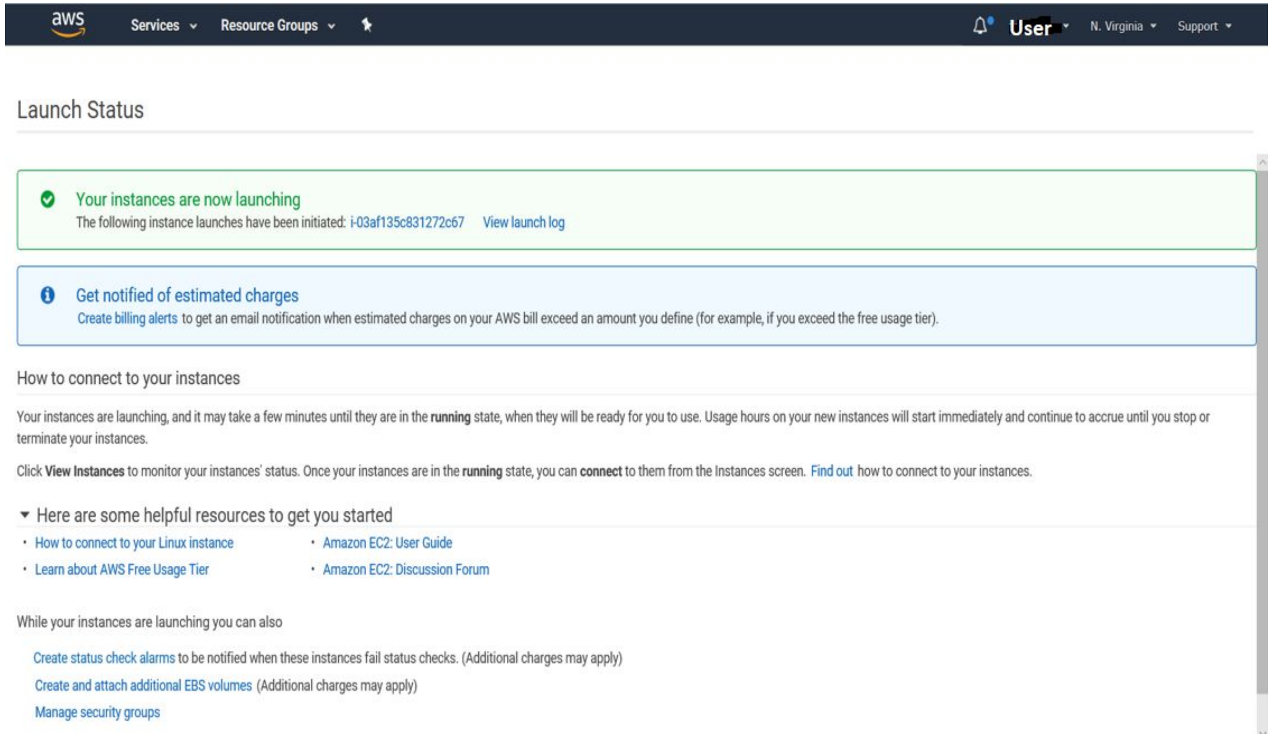

Launch Status

Your instances are now launching
The following instance launches have been initiated: i-03af135c831272c67 View launch log

Get notified of estimated charges
Create billing alerts to get an email notification when estimated charges on your AWS bill exceed an amount you define (for example, if you exceed the free usage tier).

How to connect to your instances

Your instances are launching, and it may take a few minutes until they are in the **running** state, when they will be ready for you to use. Usage hours on your new instances will start immediately and continue to accrue until you stop or terminate your instances.

Click **View Instances** to monitor your instances' status. Once your instances are in the **running** state, you can **connect** to them from the Instances screen. Find out how to connect to your instances.

Here are some helpful resources to get you started

- How to connect to your Linux instance
- Learn about AWS Free Usage Tier
- Amazon EC2: User Guide
- Amazon EC2: Discussion Forum

While your instances are launching you can also

Create status check alarms to be notified when these instances fail status checks. (Additional charges may apply)
Create and attach additional EBS volumes (Additional charges may apply)
Manage security groups

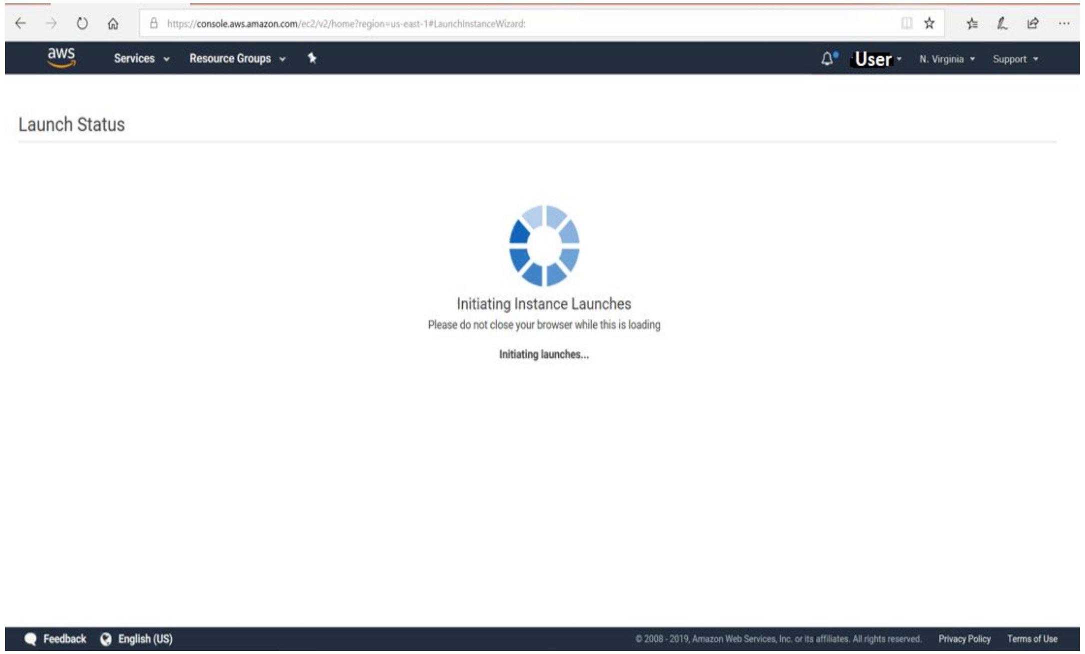
https://console.aws.amazon.com/ec2/v2/home?region=us-east-1#LaunchInstanceWizard:
aws
Services
Resource Groups
User
N. Virginia
Support

Launch Status

Initiating Instance Launches
Please do not close your browser while this is loading

Initiating launches...

Feedback
English (US)
© 2008 - 2019, Amazon Web Services, Inc. or its affiliates. All rights reserved.
Privacy Policy
Terms of Use

Left side display a successful instance launched.
Right side display instance launching/loading in the dashboard.

8. Connecting to the Instance

1 You can initialize the created instance by clicking on “**Actions**”

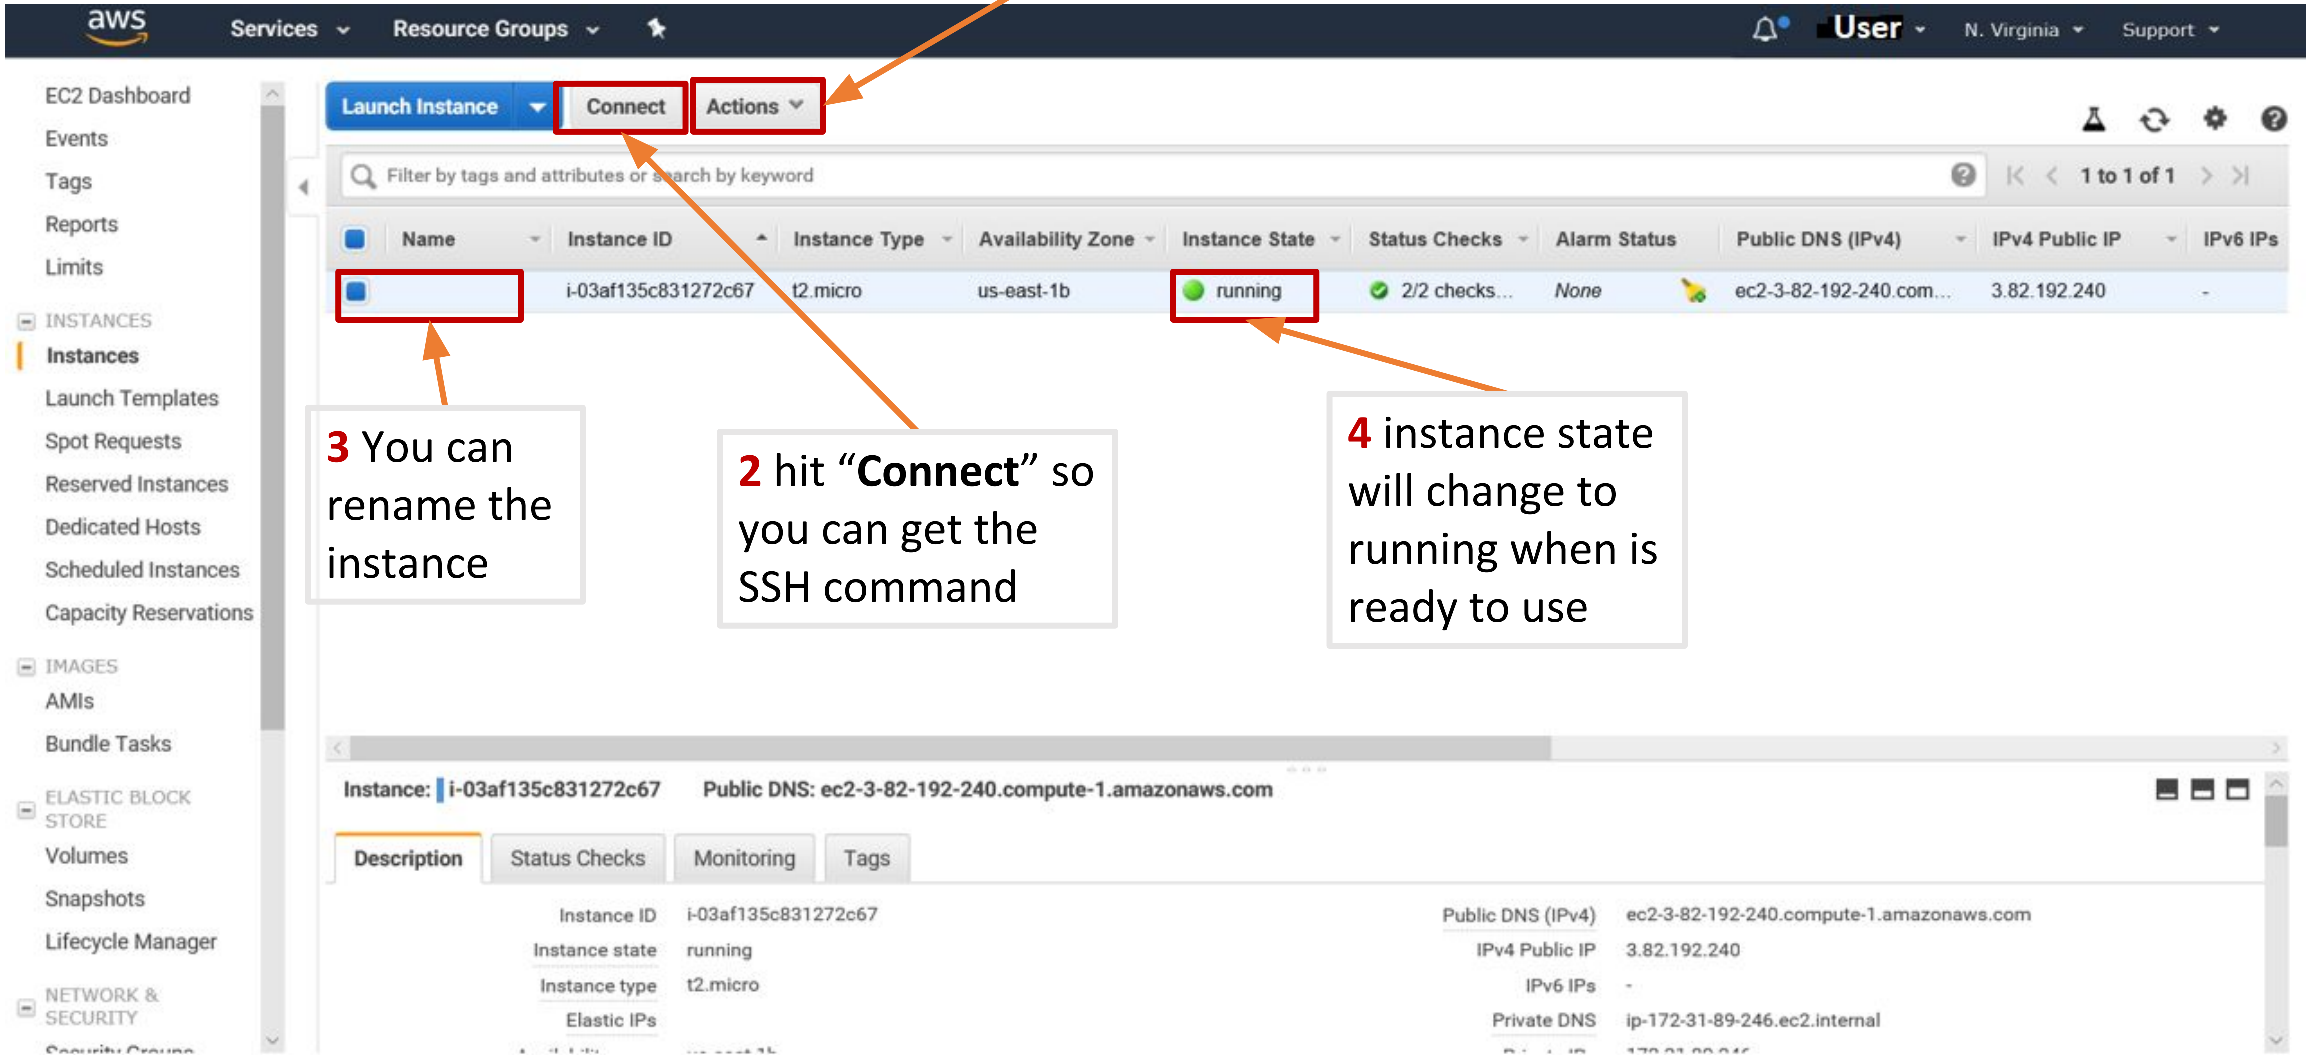

Name	Instance ID	Instance Type	Availability Zone	Instance State	Status Checks	Alarm Status	Public DNS (IPv4)	IPv4 Public IP	IPv6 IPs
	i-03af135c831272c67	t2.micro	us-east-1b	running	2/2 checks...	None	ec2-3-82-192-240.com...	3.82.192.240	-

Instance: i-03af135c831272c67 Public DNS: ec2-3-82-192-240.compute-1.amazonaws.com

Instance ID	i-03af135c831272c67	Public DNS (IPv4)	ec2-3-82-192-240.compute-1.amazonaws.com
Instance state	running	IPv4 Public IP	3.82.192.240
Instance type	t2.micro	IPv6 IPs	-
Elastic IPs		Private DNS	ip-172-31-89-246.ec2.internal

3 You can rename the instance

2 hit “**Connect**” so you can get the SSH command

4 instance state will change to running when is ready to use

8.1 Connecting to the Instance

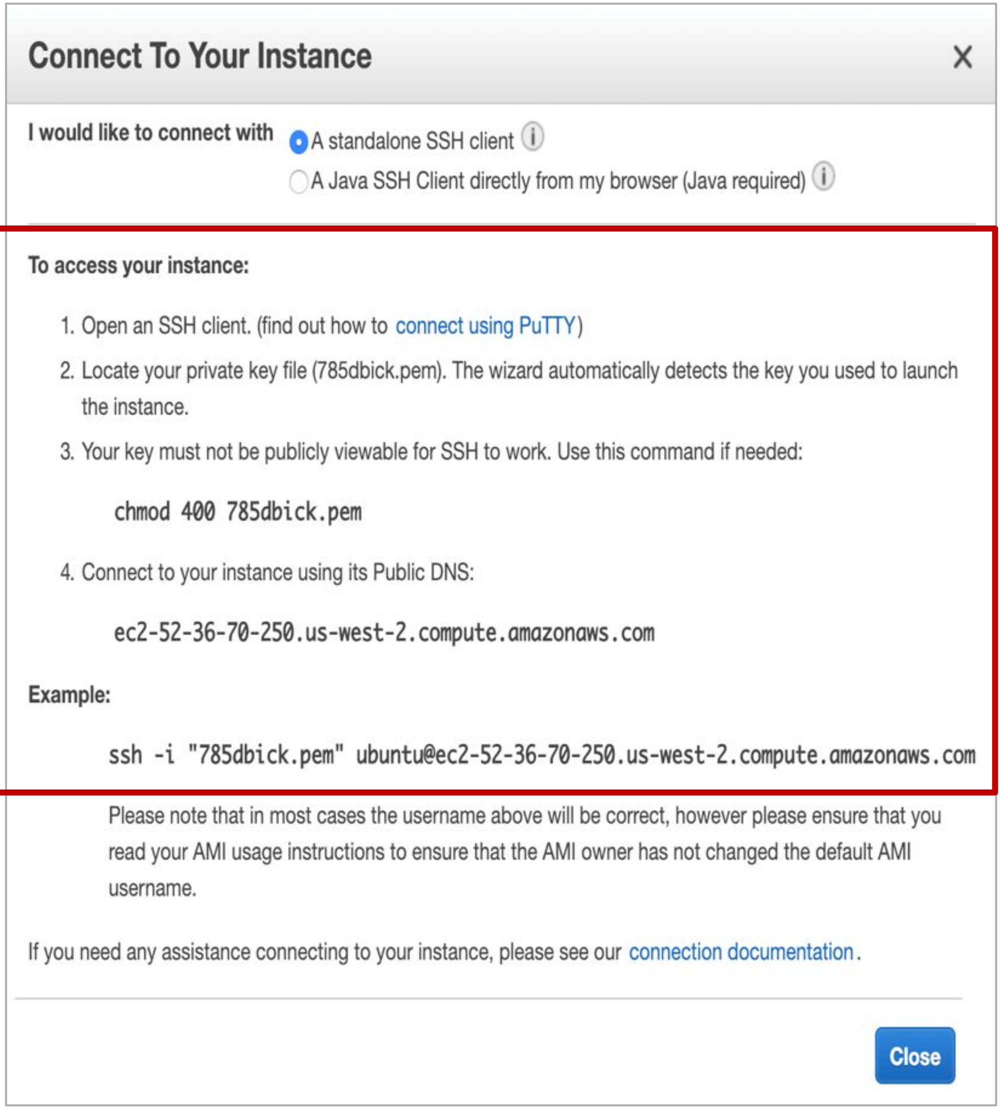

I would like to connect with ● A standalone SSH client ○ A Java SSH Client directly from my browser (Java required)

To access your instance:

1. Open an SSH client. (find out how to connect using PuTTY)
2. Locate your private key file (785dbick.pem). The wizard automatically detects the key you used to launch the instance.
3. Your key must not be publicly viewable for SSH to work. Use this command if needed:

```
chmod 400 785dbick.pem
```

4. Connect to your instance using its Public DNS:

```
ec2-52-36-70-250.us-west-2.compute.amazonaws.com
```

Example:

```
ssh -i "785dbick.pem" ubuntu@ec2-52-36-70-250.us-west-2.compute.amazonaws.com
```

Please note that in most cases the username above will be correct, however please ensure that you read your AMI usage instructions to ensure that the AMI owner has not changed the default AMI username.

If you need any assistance connecting to your instance, please see our connection documentation.

1 These are the options to connect to the instance once you hit "**Connect**" on step **2** in the previous slide

9. Stopping the Instance

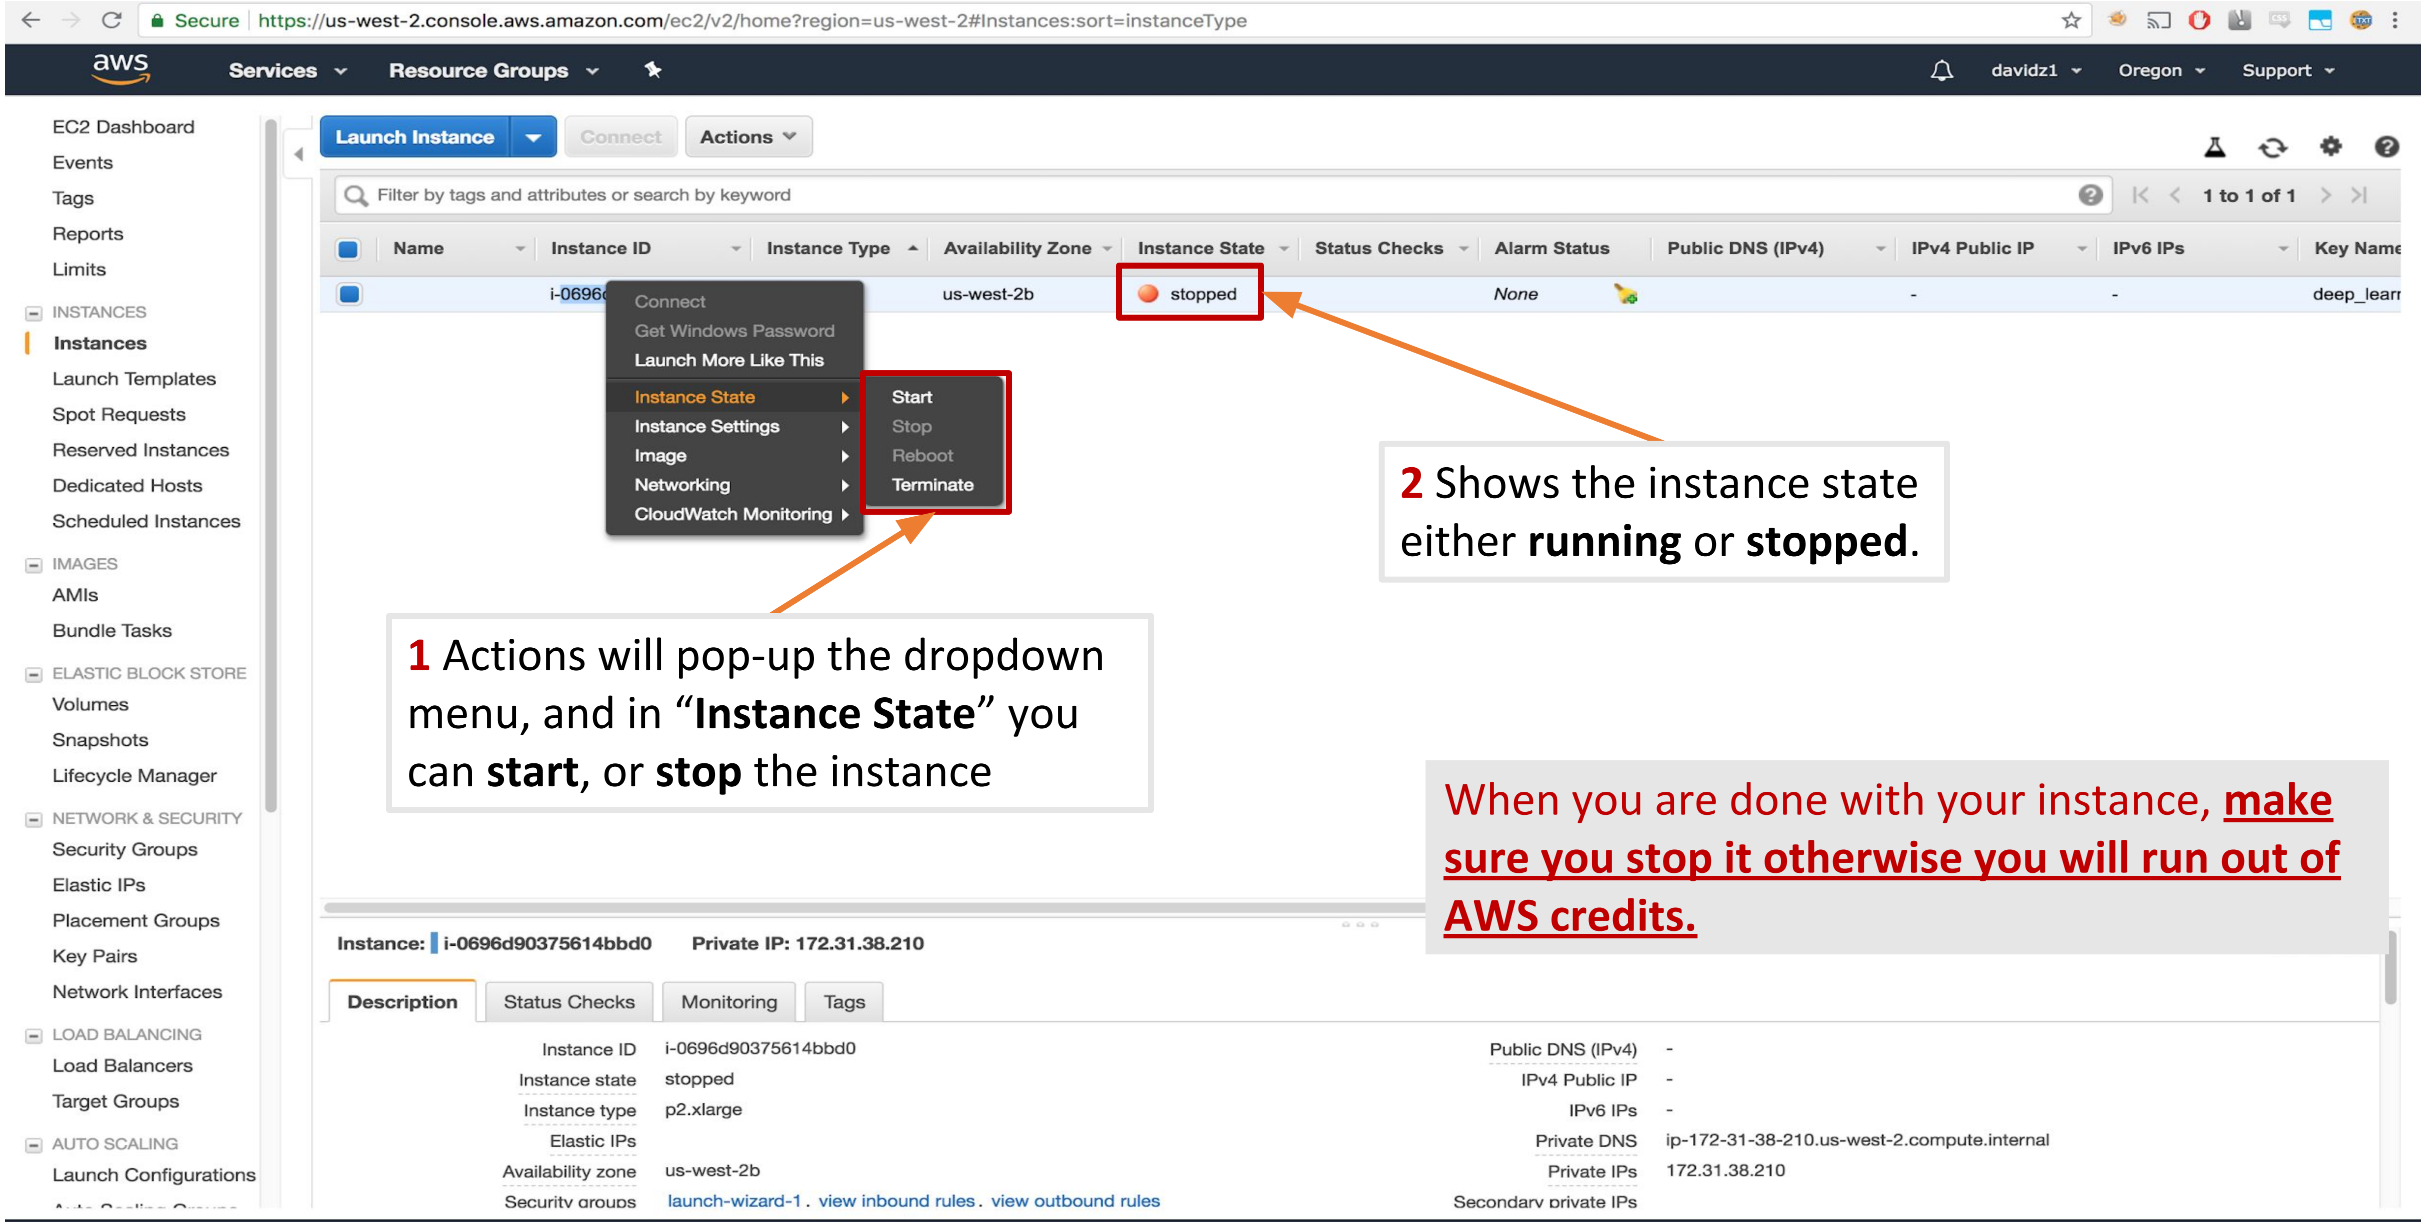

Name	Instance ID	Instance Type	Availability Zone	Instance State	Status Checks	Alarm Status	Public DNS (IPv4)	IPv4 Public IP	IPv6 IPs	Key Name
	i-0696d		us-west-2b	● stopped		None		-	-	deep_learr

Instance: i-0696d90375614bbd0 Private IP: 172.31.38.210

Instance ID	i-0696d90375614bbd0	Public DNS (IPv4)	-
Instance state	stopped	IPv4 Public IP	-
Instance type	p2.xlarge	IPv6 IPs	-
Elastic IPs		Private DNS	ip-172-31-38-210.us-west-2.compute.internal
Availability zone	us-west-2b	Private IPs	172.31.38.210
Security groups	launch-wizard-1 . view inbound rules . view outbound rules	Secondary private IPs	

1 Actions will pop-up the dropdown menu, and in "**Instance State**" you can **start**, or **stop** the instance

2 Shows the instance state either **running** or **stopped**.

When you are done with your instance, **make sure you stop it otherwise you will run out of AWS credits.**

AWS Billing and Coupons

Don’t go Broke!

- Amazon charges you for a lot of services, the most significant is having running instances (GPU time is NOT cheap).

- Terminating an instance **deletes the entire machine only do this when you completely done with whatever you are doing with your AWS instance**
- Stopping an instance is basically shutting down a computer, the saved files persist, etc.

- **We will provide you with AWS tokens**, you need to redeem them from Billing.

1. Billing

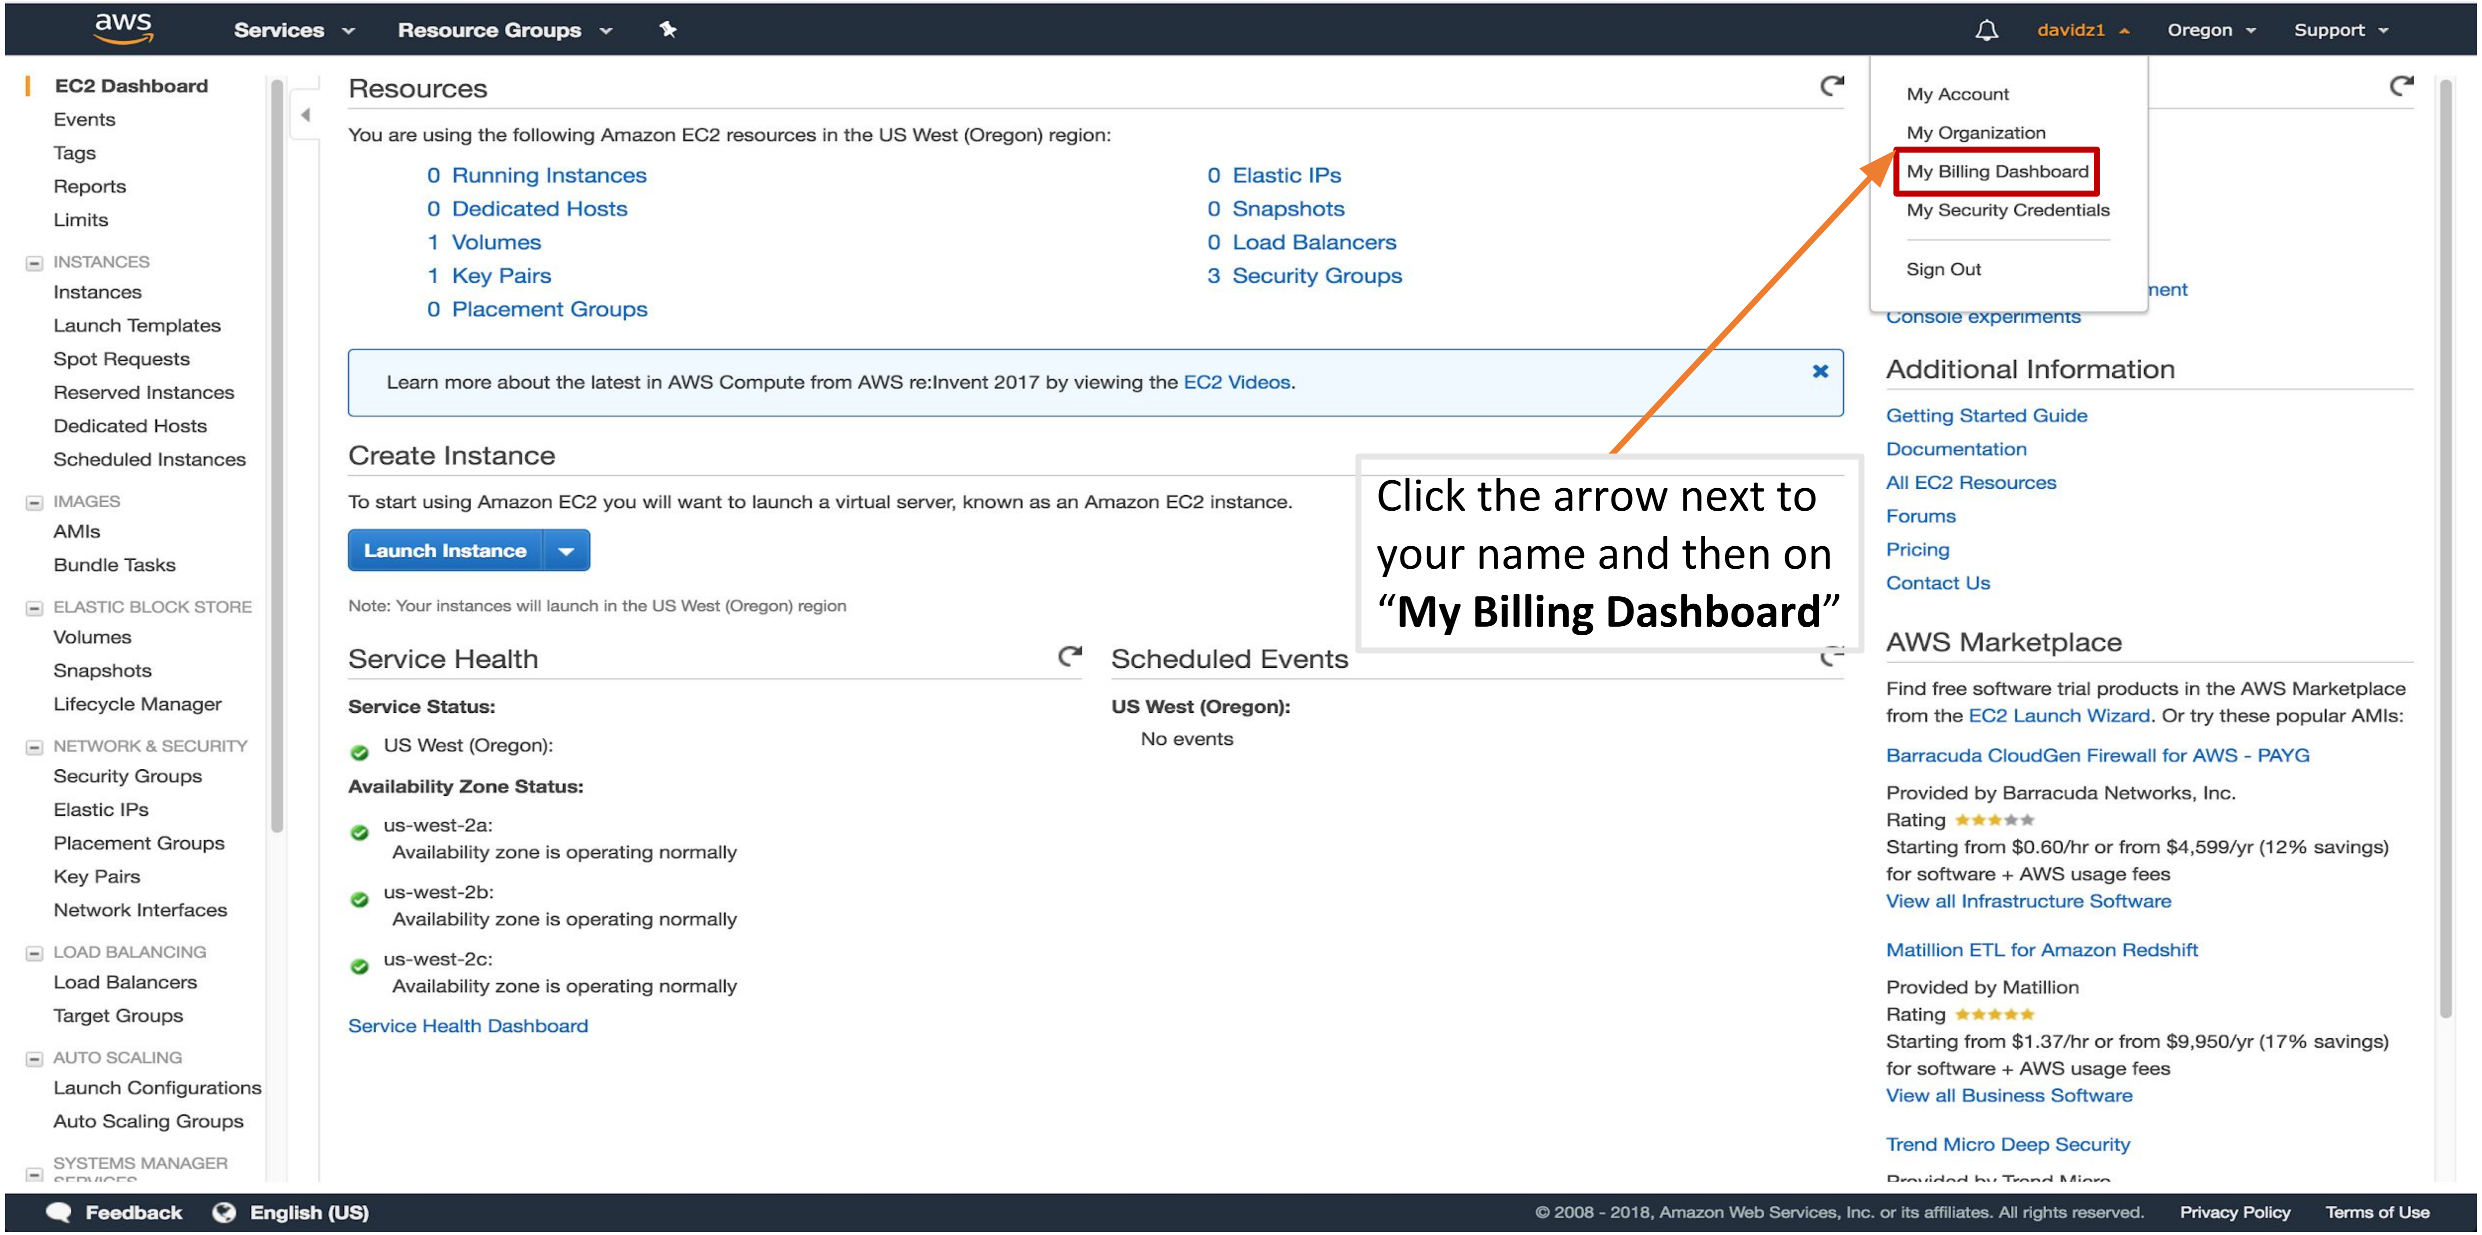

aws Services ▾ Resource Groups ▾ davidz1 ▴ Oregon ▾ Support ▾

My Account

My Organization

My Billing Dashboard

My Security Credentials

Sign Out

EC2 Dashboard

Events

Tags

Reports

Limits

INSTANCES

Instances

Launch Templates

Spot Requests

Reserved Instances

Dedicated Hosts

Scheduled Instances

IMAGES

AMIs

Bundle Tasks

ELASTIC BLOCK STORE

Volumes

Snapshots

Lifecycle Manager

NETWORK & SECURITY

Security Groups

Elastic IPs

Placement Groups

Key Pairs

Network Interfaces

LOAD BALANCING

Load Balancers

Target Groups

AUTO SCALING

Launch Configurations

Auto Scaling Groups

SYSTEMS MANAGER SERVICES

Resources

You are using the following Amazon EC2 resources in the US West (Oregon) region:

0 Running Instances	0 Elastic IPs
0 Dedicated Hosts	0 Snapshots
1 Volumes	0 Load Balancers
1 Key Pairs	3 Security Groups
0 Placement Groups	

Learn more about the latest in AWS Compute from AWS re:Invent 2017 by viewing the EC2 Videos. ✖

Create Instance

To start using Amazon EC2 you will want to launch a virtual server, known as an Amazon EC2 instance.

Launch Instance ▾

Note: Your instances will launch in the US West (Oregon) region

Service Health

Service Status:

✓ US West (Oregon):

Availability Zone Status:

✓ us-west-2a:
Availability zone is operating normally

✓ us-west-2b:
Availability zone is operating normally

✓ us-west-2c:
Availability zone is operating normally

Service Health Dashboard

Scheduled Events

US West (Oregon):

No events

...nent

Console experiments

Additional Information

Getting Started Guide

Documentation

All EC2 Resources

Forums

Pricing

Contact Us

AWS Marketplace

Find free software trial products in the AWS Marketplace from the EC2 Launch Wizard. Or try these popular AMIs:

Barracuda CloudGen Firewall for AWS - PAYG

Provided by Barracuda Networks, Inc.
Rating ★★★★★
Starting from $0.60/hr or from $4,599/yr (12% savings) for software + AWS usage fees
View all Infrastructure Software

Matillion ETL for Amazon Redshift

Provided by Matillion
Rating ★★★★★
Starting from $1.37/hr or from $9,950/yr (17% savings) for software + AWS usage fees
View all Business Software

Trend Micro Deep Security

Provided by Trend Micro

Feedback English (US) © 2008 - 2018, Amazon Web Services, Inc. or its affiliates. All rights reserved. Privacy Policy Terms of Use

Click the arrow next to your name and then on “**My Billing Dashboard**”

2. Redemption Page

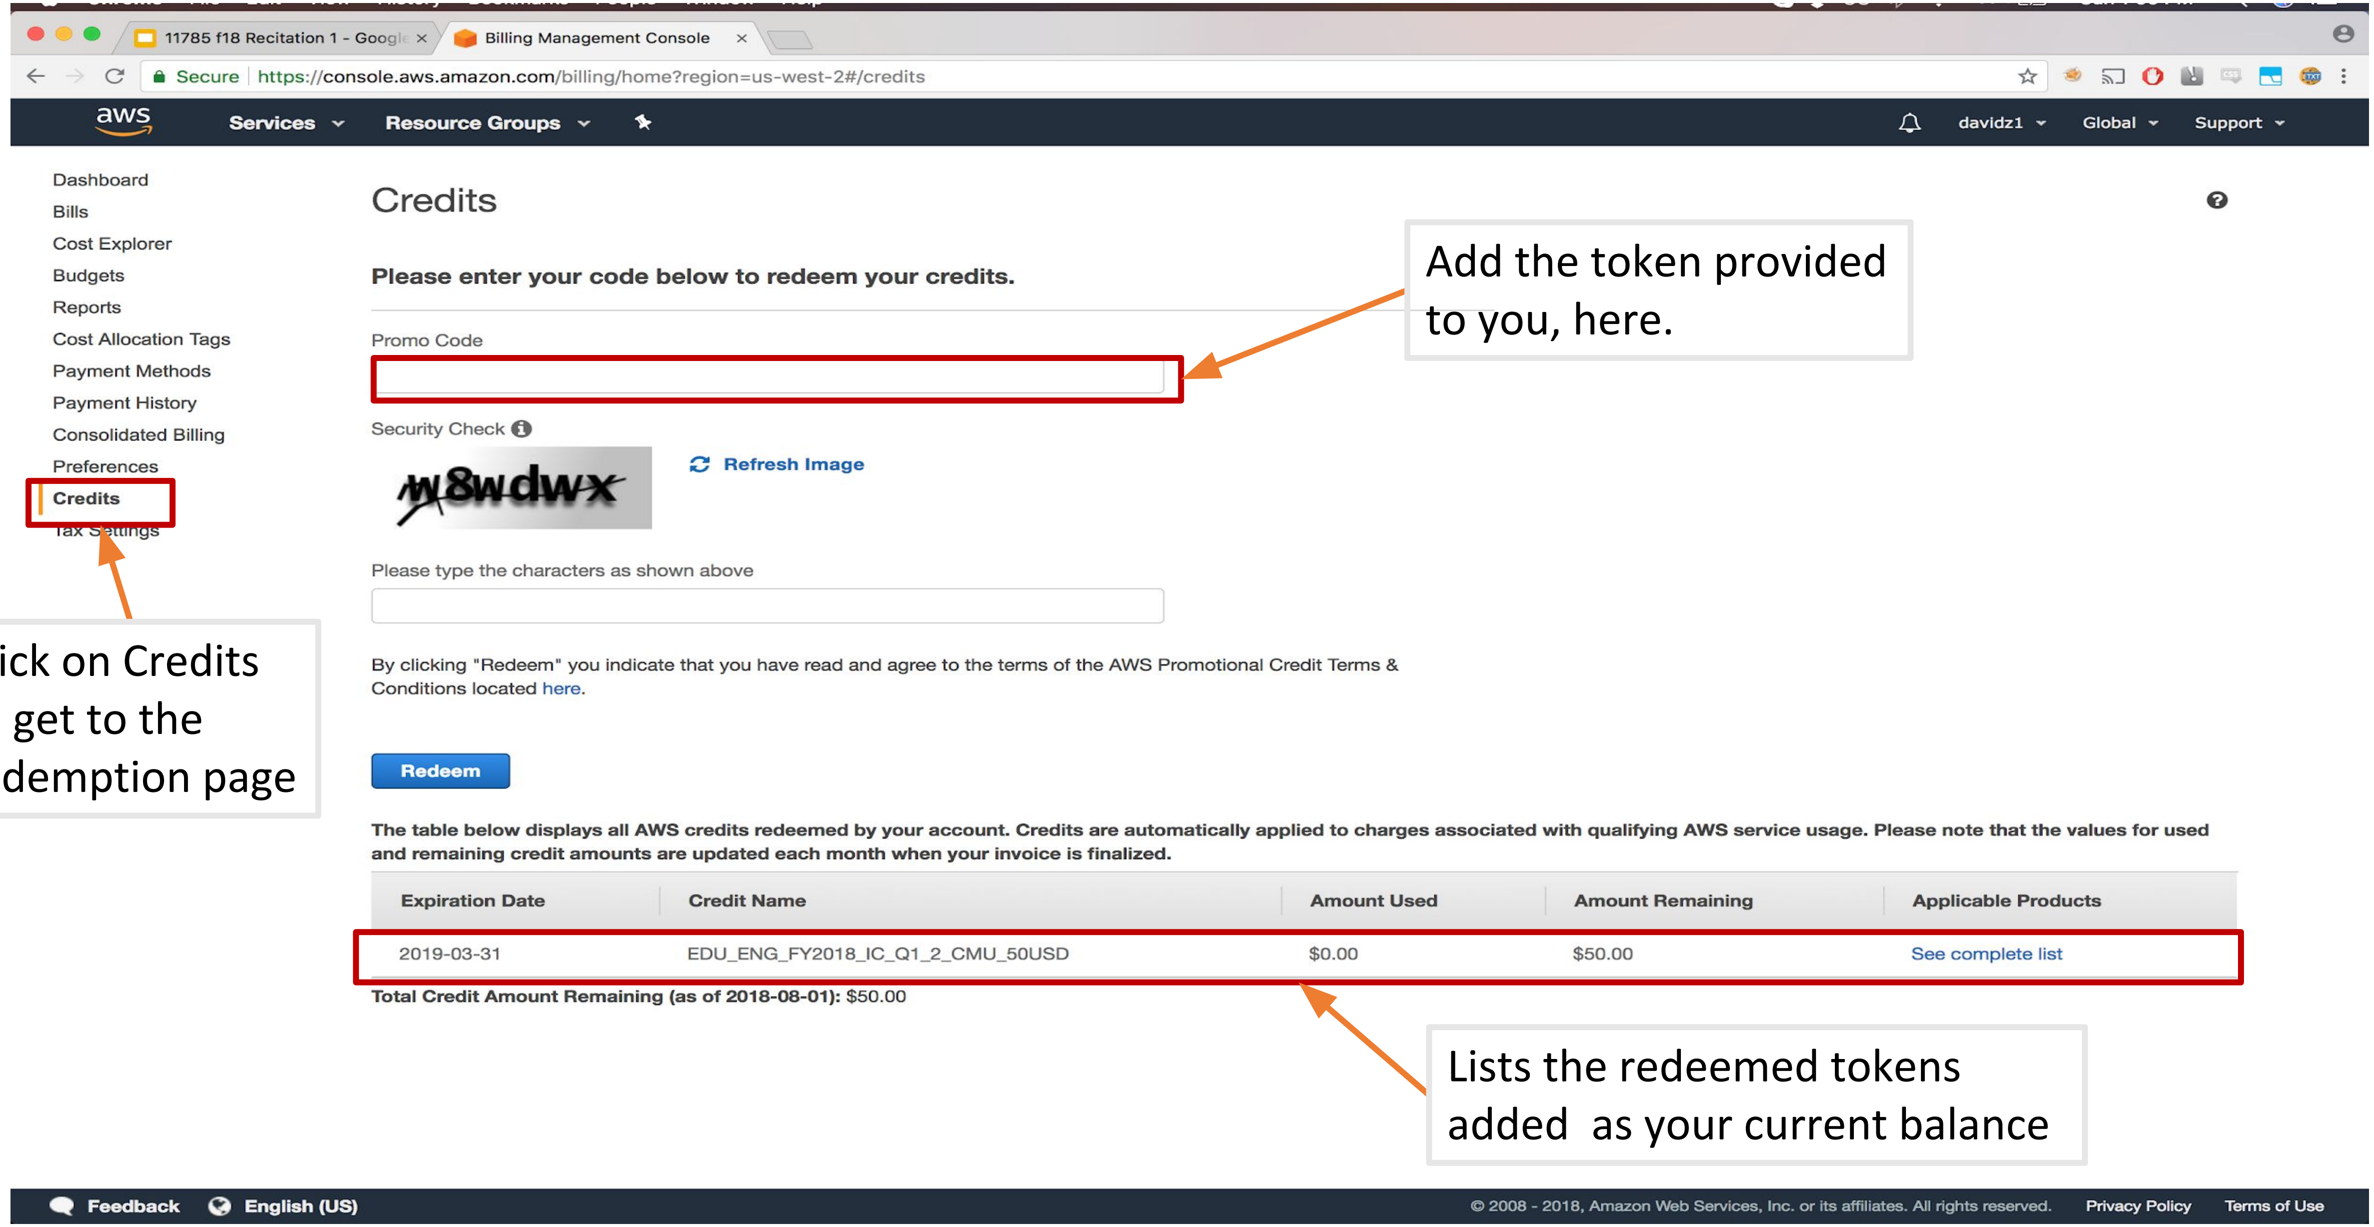

11785 f18 Recitation 1 - Google
Billing Management Console
Secure https://console.aws.amazon.com/billing/home?region=us-west-2#/credits
aws
Services
Resource Groups
davidz1
Global
Support

- Dashboard
- Bills
- Cost Explorer
- Budgets
- Reports
- Cost Allocation Tags
- Payment Methods
- Payment History
- Consolidated Billing
- Preferences
- **Credits**
- Tax Settings

Credits

Please enter your code below to redeem your credits.

Promo Code

Security Check

w8wdwx

Refresh Image

Please type the characters as shown above

By clicking "Redeem" you indicate that you have read and agree to the terms of the AWS Promotional Credit Terms & Conditions located here.

Redeem

The table below displays all AWS credits redeemed by your account. Credits are automatically applied to charges associated with qualifying AWS service usage. Please note that the values for used and remaining credit amounts are updated each month when your invoice is finalized.

Expiration Date	Credit Name	Amount Used	Amount Remaining	Applicable Products
2019-03-31	EDU_ENG_FY2018_IC_Q1_2_CMU_50USD	$0.00	$50.00	See complete list

Total Credit Amount Remaining (as of 2018-08-01): $50.00

Feedback
English (US)
© 2008 - 2018, Amazon Web Services, Inc. or its affiliates. All rights reserved.
Privacy Policy
Terms of Use

Add the token provided to you, here.

Click on Credits to get to the redemption page

Lists the redeemed tokens added as your current balance

Useful Tools and Tips while working with AWS

How to Open your Jupyter Notebook from Your Instance?

1. Connect to your instance through SSH:

```
ssh -i KeyTest.pem -L 8000:localhost:8888 ubuntu@ec2-34-227-222-100.compute-1.amazonaws.com
```

2. Open a new terminal window, and repeat step **1** then,

Type:

```
source activate pytorch_p36
```

3. Once in the environment Call the Jupyter Notebook by executing:

```
>$ jupyter notebook --no-browser --port=8888
```

4. Finally, go to browser and enter

```
localhost:8888
```

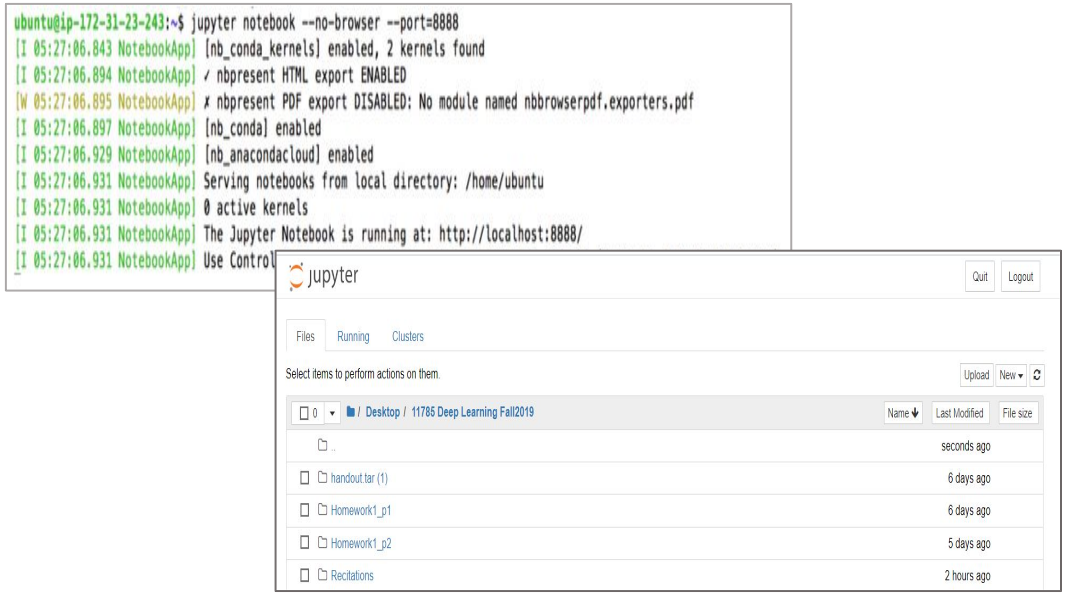

```
ubuntu@ip-172-31-23-243:~$ jupyter notebook --no-browser --port=8888
[I 05:27:06.843 NotebookApp] [nb_conda_kernels] enabled, 2 kernels found
[I 05:27:06.894 NotebookApp] ✓ nbpresent HTML export ENABLED
[W 05:27:06.895 NotebookApp] ✗ nbpresent PDF export DISABLED: No module named nbbrowserpdf.exporters.pdf
[I 05:27:06.897 NotebookApp] [nb_conda] enabled
[I 05:27:06.929 NotebookApp] [nb_anacondacloud] enabled
[I 05:27:06.931 NotebookApp] Serving notebooks from local directory: /home/ubuntu
[I 05:27:06.931 NotebookApp] 0 active kernels
[I 05:27:06.931 NotebookApp] The Jupyter Notebook is running at: http://localhost:8888/
[I 05:27:06.931 NotebookApp] Use Control
```

Select items to perform actions on them.

□ 0 ▾ / Desktop / 11785 Deep Learning Fall2019	Name ↓	Last Modified	File size
..		seconds ago	
□ handout.tar (1)		6 days ago	
□ Homework1_p1		6 days ago	
□ Homework1_p2		5 days ago	
□ Recitations		2 hours ago	

So, you should be able to see the home directory requesting for a password or token, enter the token id for the notebook, so you will be able to see the home directory and the current projects under. the first window will run the notebook, and the second window you can use it for everything else.

Reference:

For Jupyter Notebooks

- You can also use SageMaker another service from AWS that was explained in the previous recitation just be aware that the price is bit more than using EC2 instances.

How to Use Screens?

1. Install Linux screen by using the command:

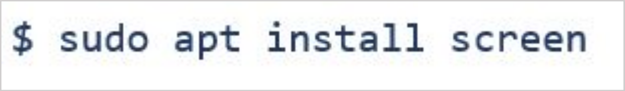

```
$ sudo apt install screen
```

2. To start a screen simply type:

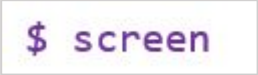

```
$ screen
```

This will open a screen session, create a new window and start a shell in that window. You can get a list of commands on how to use screens by typing:

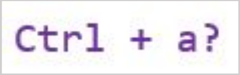

```
Ctrl + a?
```

3. You can start a named session by using:

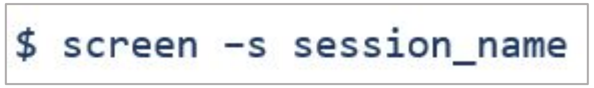

```
$ screen -s session_name
```

By using screens you can create multiple sessions without creating multiple windows (very useful when using the notebooks on EC2).

Reference: https://linuxize.com

Some useful commands used with Screens

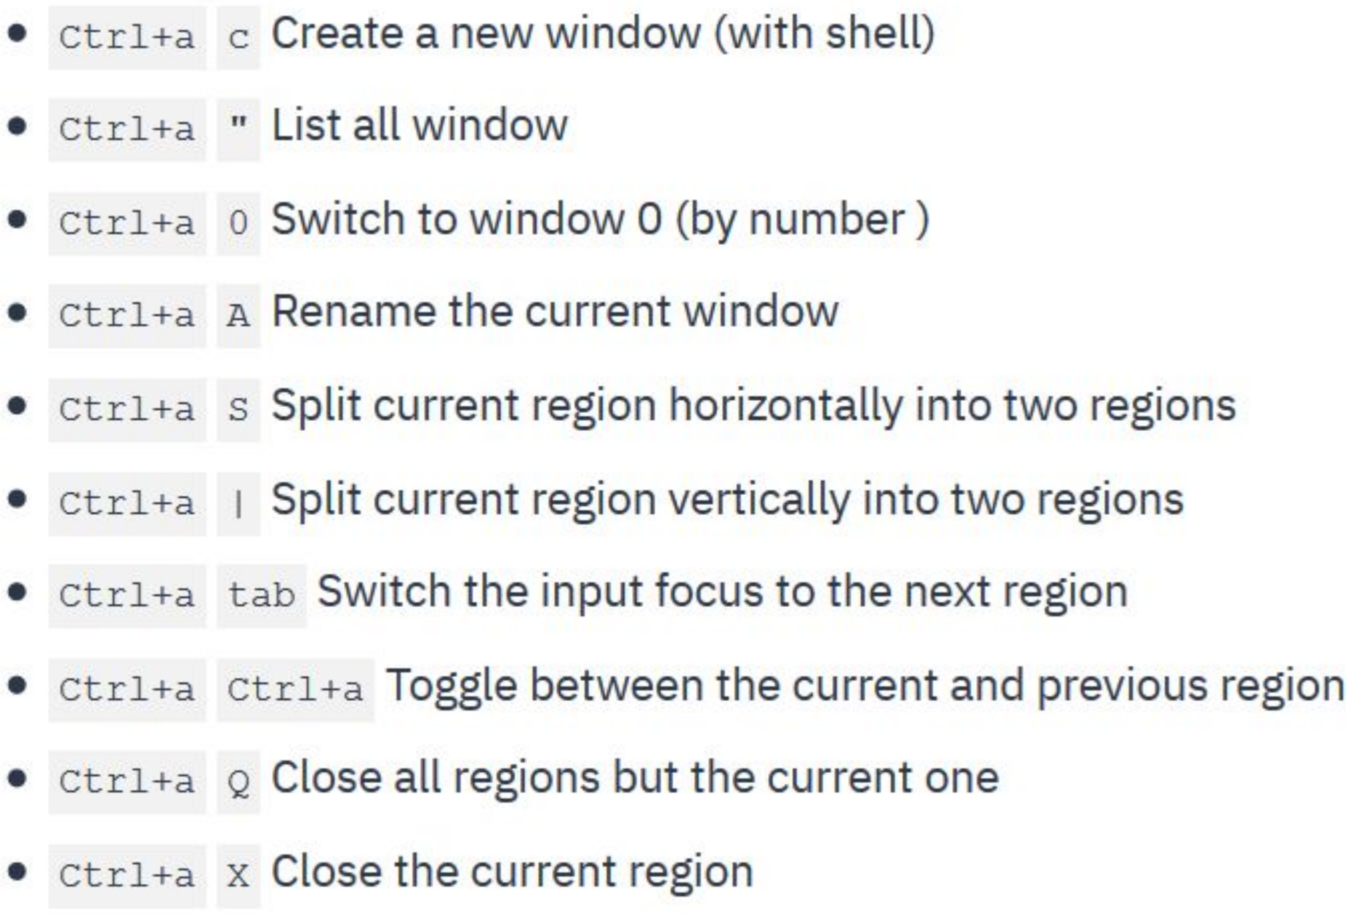

- `Ctrl+a` `c` Create a new window (with shell)
- `Ctrl+a` `"` List all window
- `Ctrl+a` `0` Switch to window 0 (by number)
- `Ctrl+a` `A` Rename the current window
- `Ctrl+a` `S` Split current region horizontally into two regions
- `Ctrl+a` `|` Split current region vertically into two regions
- `Ctrl+a` `tab` Switch the input focus to the next region
- `Ctrl+a` `Ctrl+a` Toggle between the current and previous region
- `Ctrl+a` `Q` Close all regions but the current one
- `Ctrl+a` `X` Close the current region

Reference: https://linuxize.com

Using TMUX

1. Use tmux if you don’t want to wait on your computer like its your child
2. After SSHing into an AWS instance you can type in the command tmux
3. Then run anything as would normally
4. Usually when you run something on SSH if your connection dies (ie when your computer sleeps/etc) the program you are running in the SSH session dies too
5. tmux prevents this (you can close out the session whenever)
6. To reconnect just SSH back in and then do: tmux ls
7. And then you should see your session (your first unnamed session is 0)
8. tmux attach-session -t 0

TMUX makes your life easier!!

File transfering

- From Local Machine to AWS EC2 instance

  ```
  scp -i path/to/key file/to/copy user@ec2-xx-xx-xxx-xxx.compute-1.amazonaws.com:path/to/file
  ```

- From AWS EC2 instance to Local Machine

  ```
  scp -i path/to/key user@ec2-xx-xx-xxx-xxx.compute-1.amazonaws.com:path/to/file file/to/copy
  ```

- FTP services like FileZilla can also be used

- Secure copy (scp) is a linux/unix command. Windows user can either use Ubuntu shell (download from Windows Store) or can use Git Bash (recommended)

Colab 101