Distributed Systems

09r. Map-Reduce Programming on AWS/EMR (Part I)

Paul Krzyzanowski
TA: Long Zhao
Rutgers University
Fall 2017

© 2017 Paul Krzyzanowski

Setting Up AWS/EMR

© 2017 Paul Krzyzanowski

Set up a free AWS educate account

- Visit https://aws.amazon.com/education/awseducate/ , sign up and create your account.

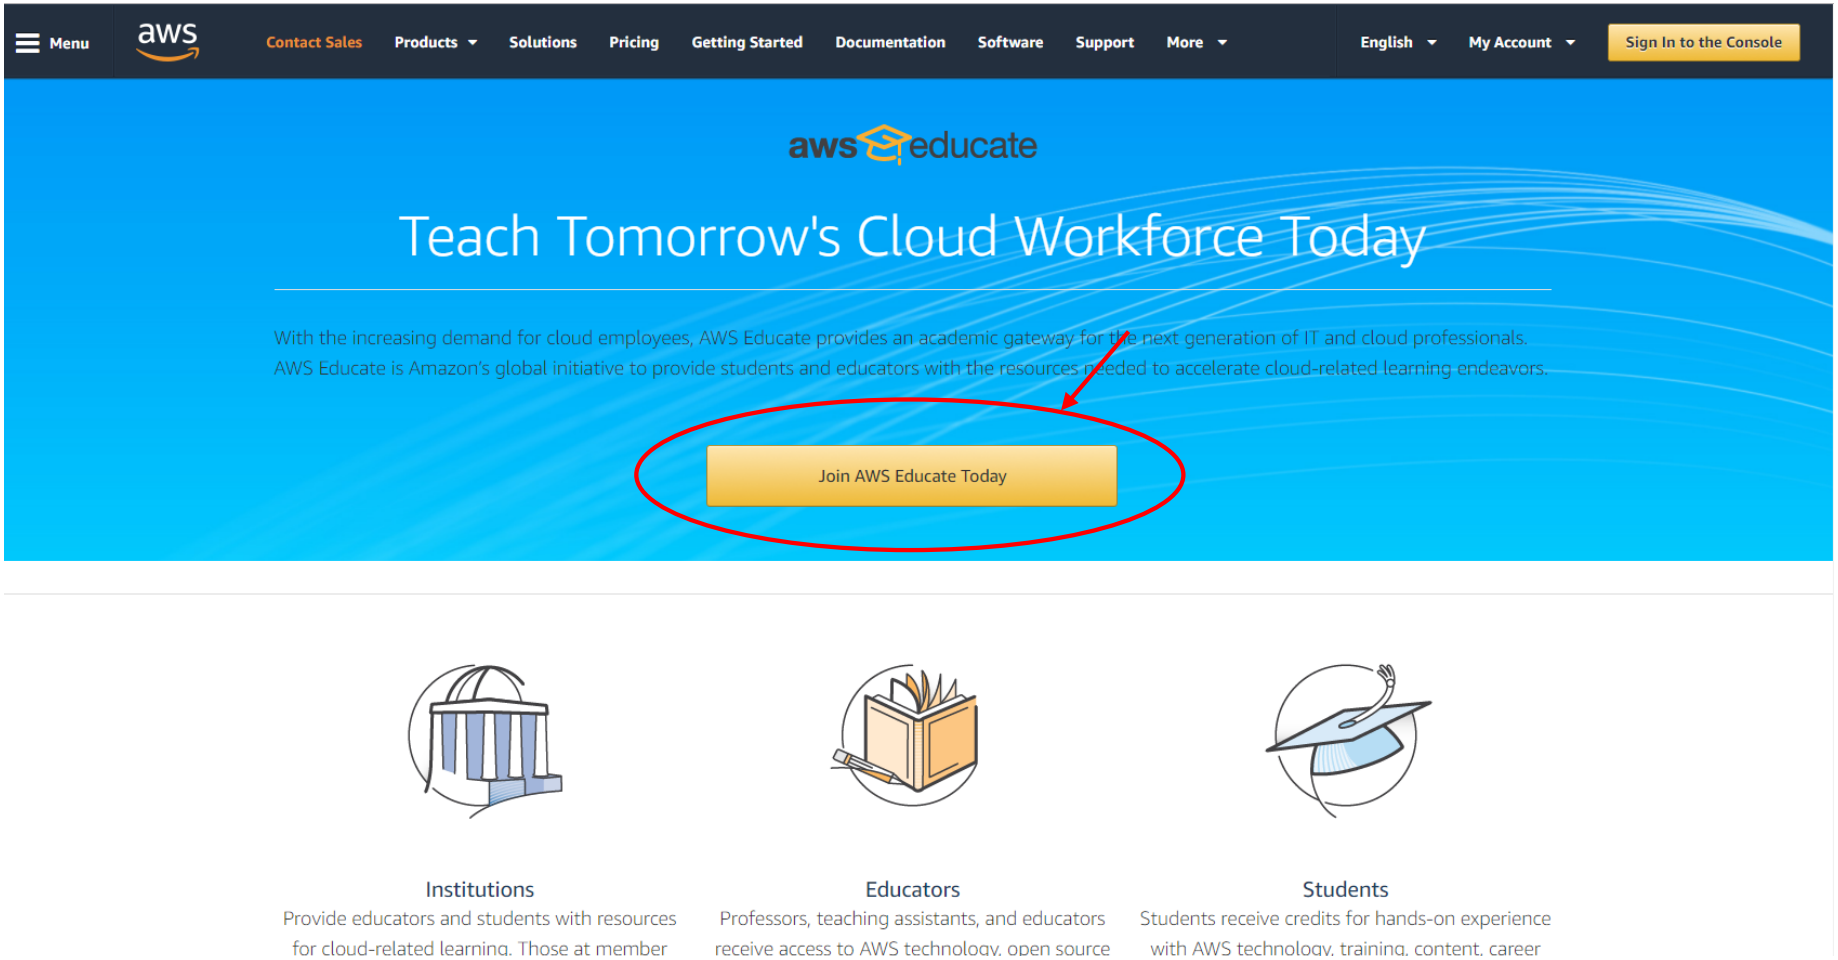

© 2017 Paul Krzyzanowski

Set up a free AWS educate account

- Step 1/3: Choose your role

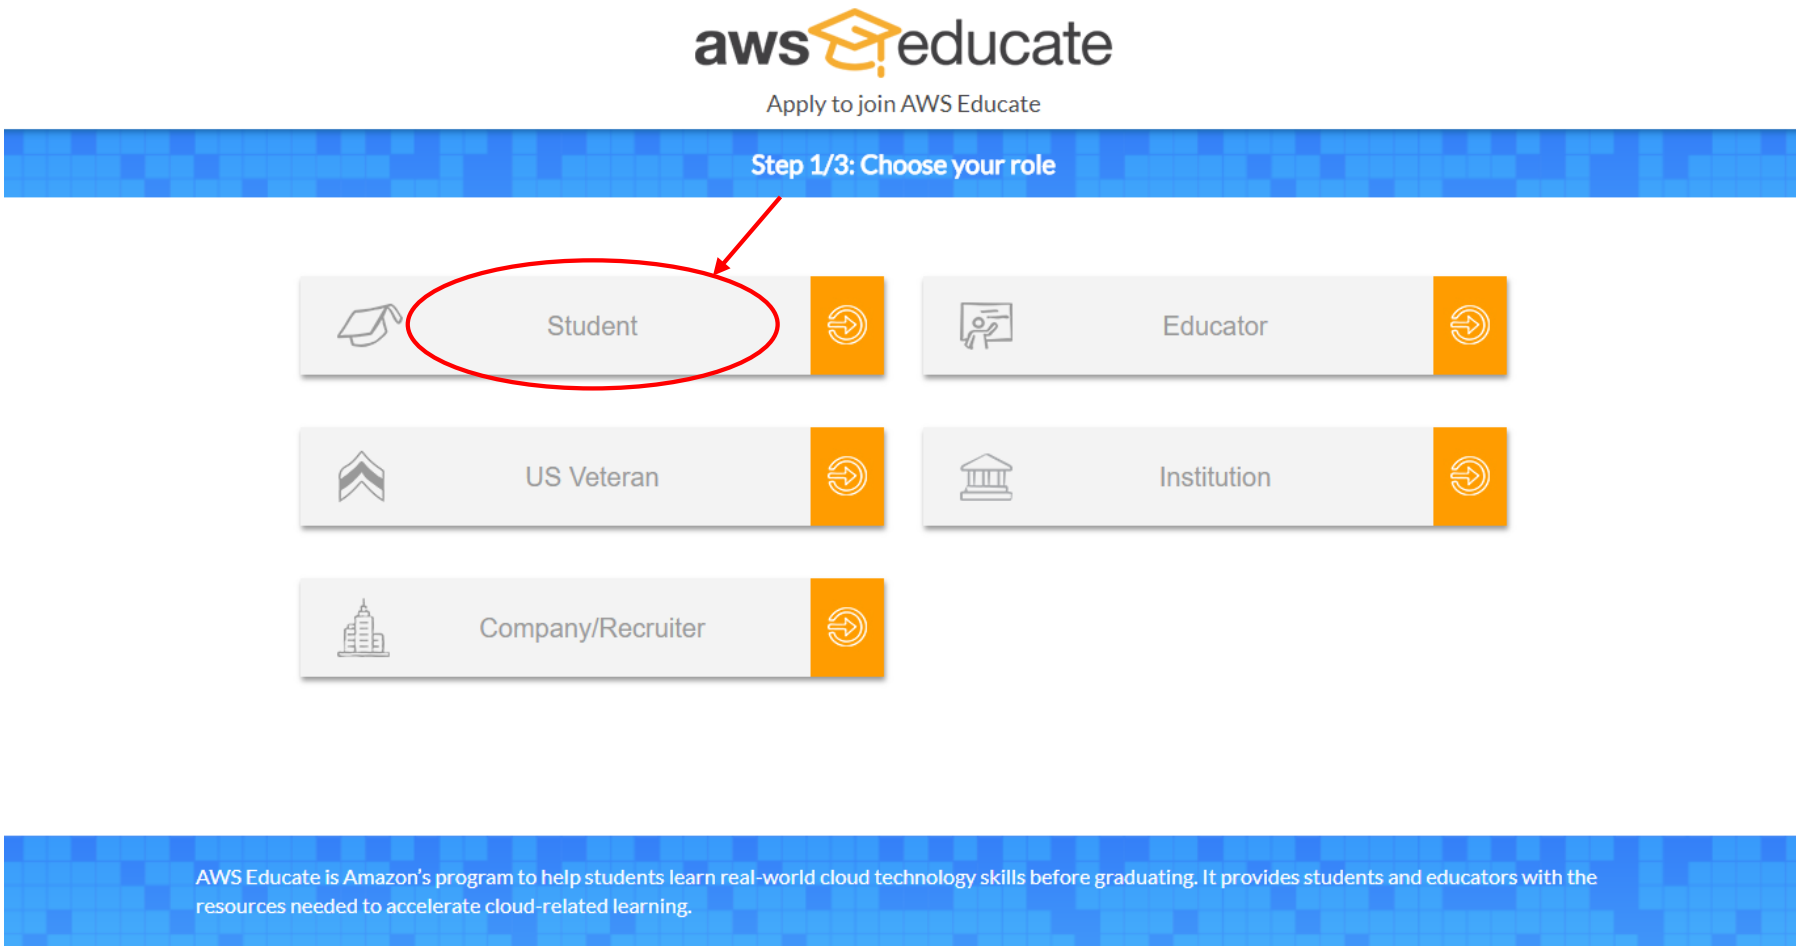

© 2017 Paul Krzyzanowski

Set up a free AWS educate account

- Step 2/3: Tell us about yourself

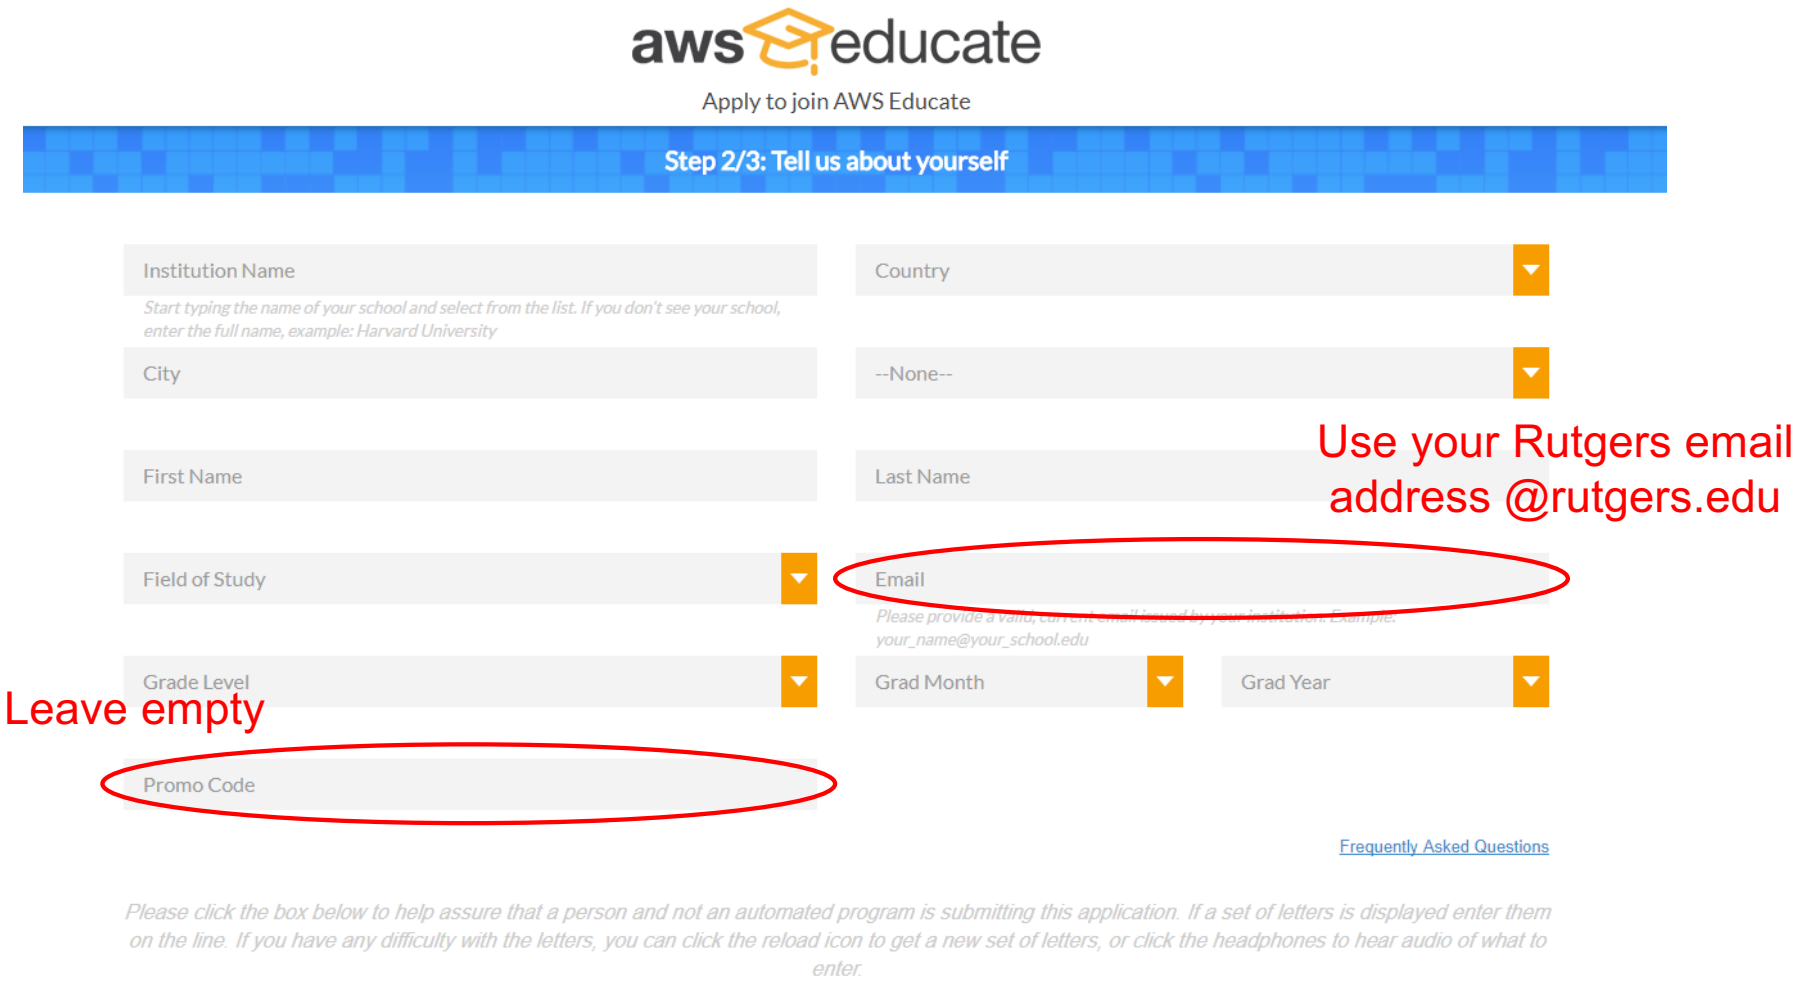

Use your Rutgers email address @rutgers.edu

Leave empty

© 2017 Paul Krzyzanowski

Set up a free AWS educate account

- Step 3/3: Choose one of the following

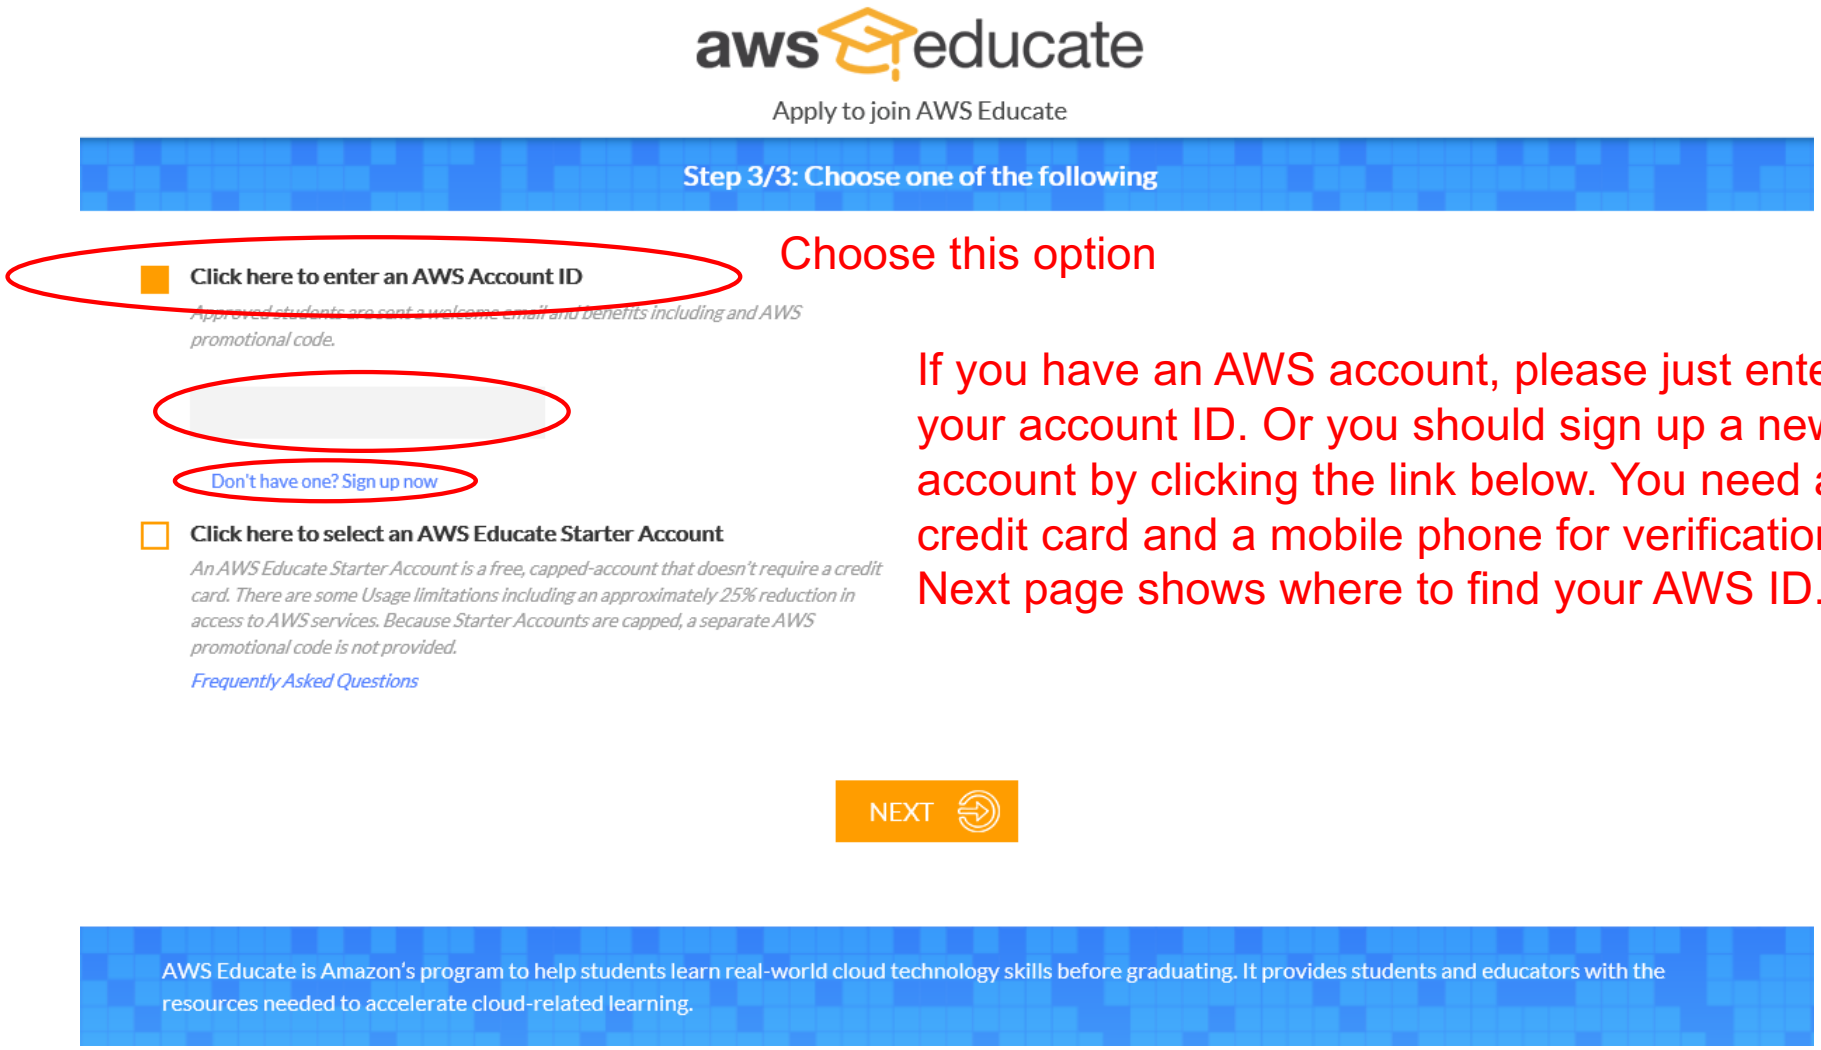

Choose this option

If you have an AWS account, please just enter your account ID. Or you should sign up a new account by clicking the link below. You need a credit card and a mobile phone for verification. Next page shows where to find your AWS ID.

© 2017 Paul Krzyzanowski

Set up a free AWS educate account

- Find your AWS account ID. First login your AWS account, then click your user name and click “My Account”.

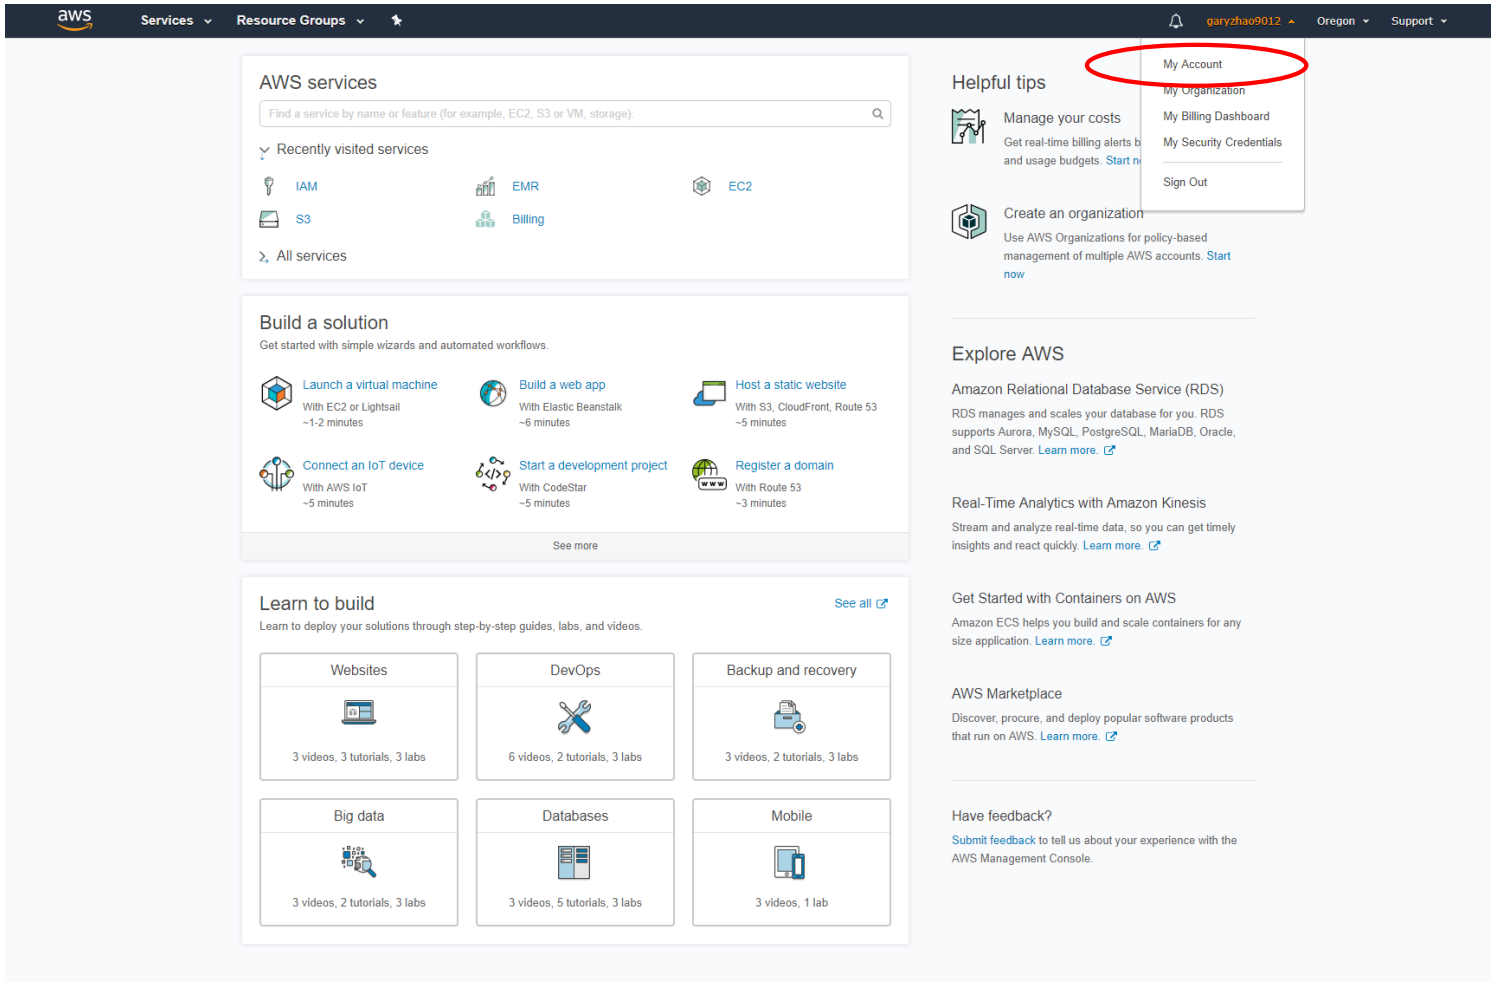

© 2017 Paul Krzyzanowski

Set up a free AWS educate account

- Find your AWS account ID

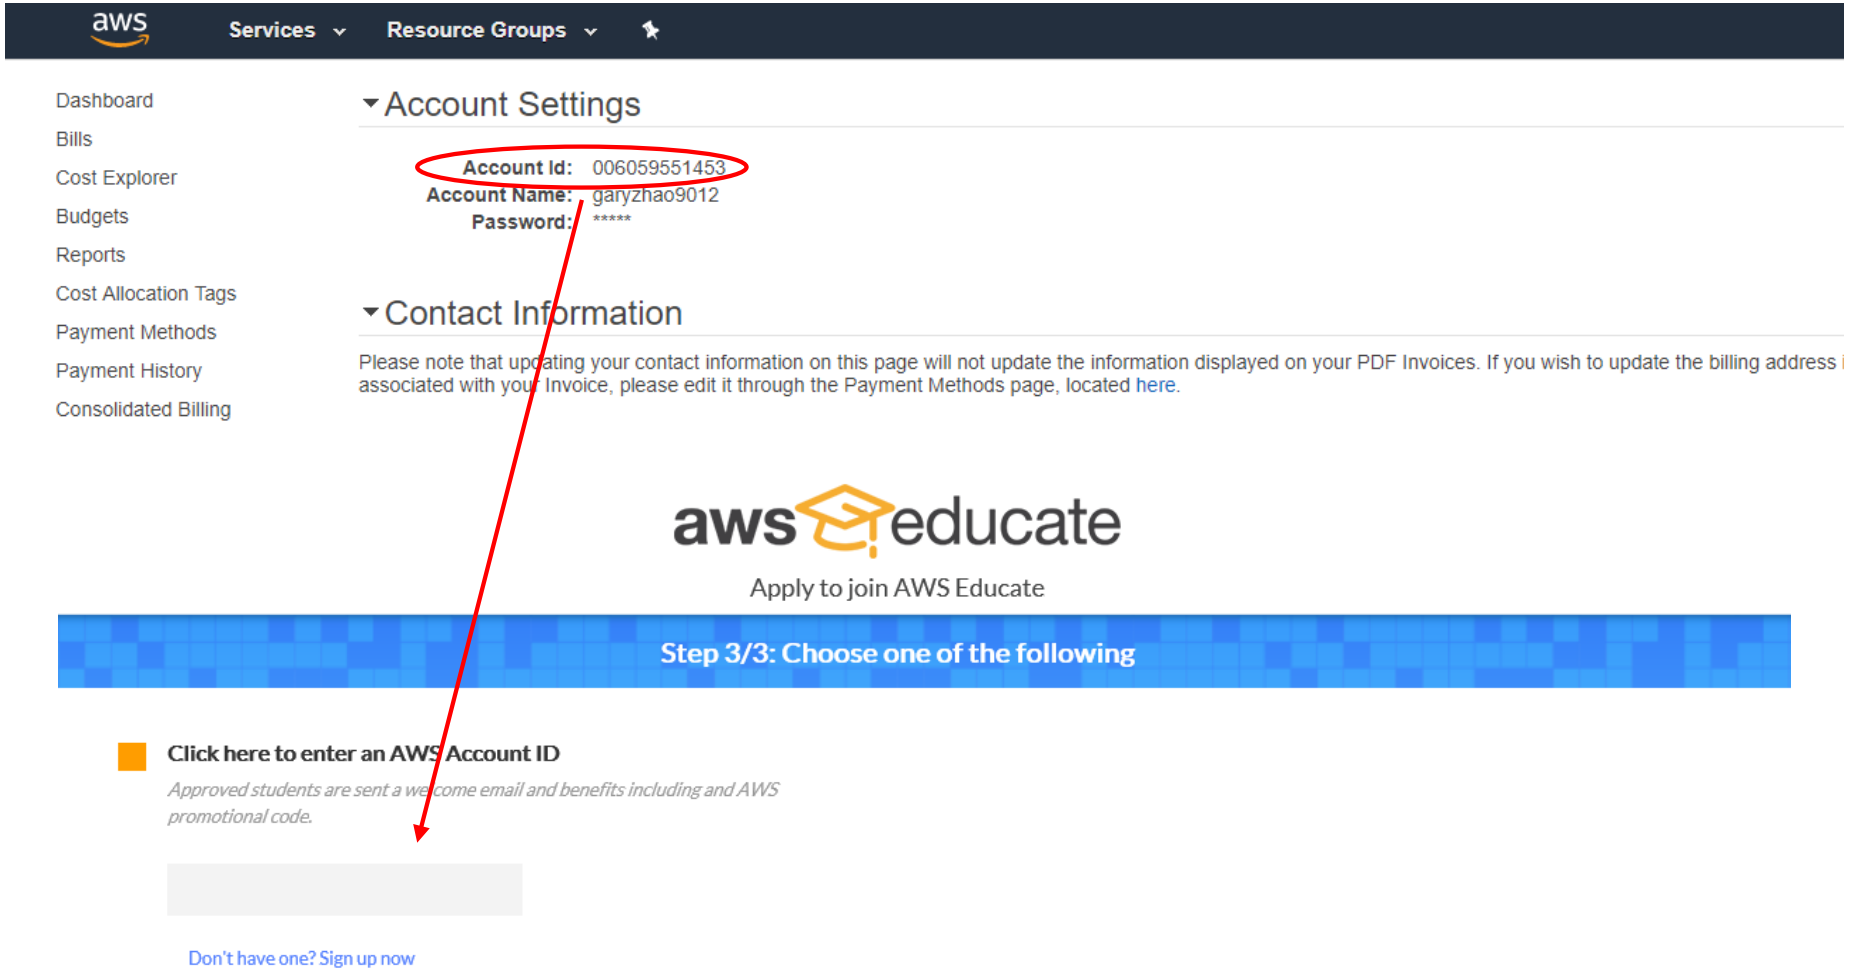

© 2017 Paul Krzyzanowski

Set up a free AWS educate account

- Find your AWS account ID

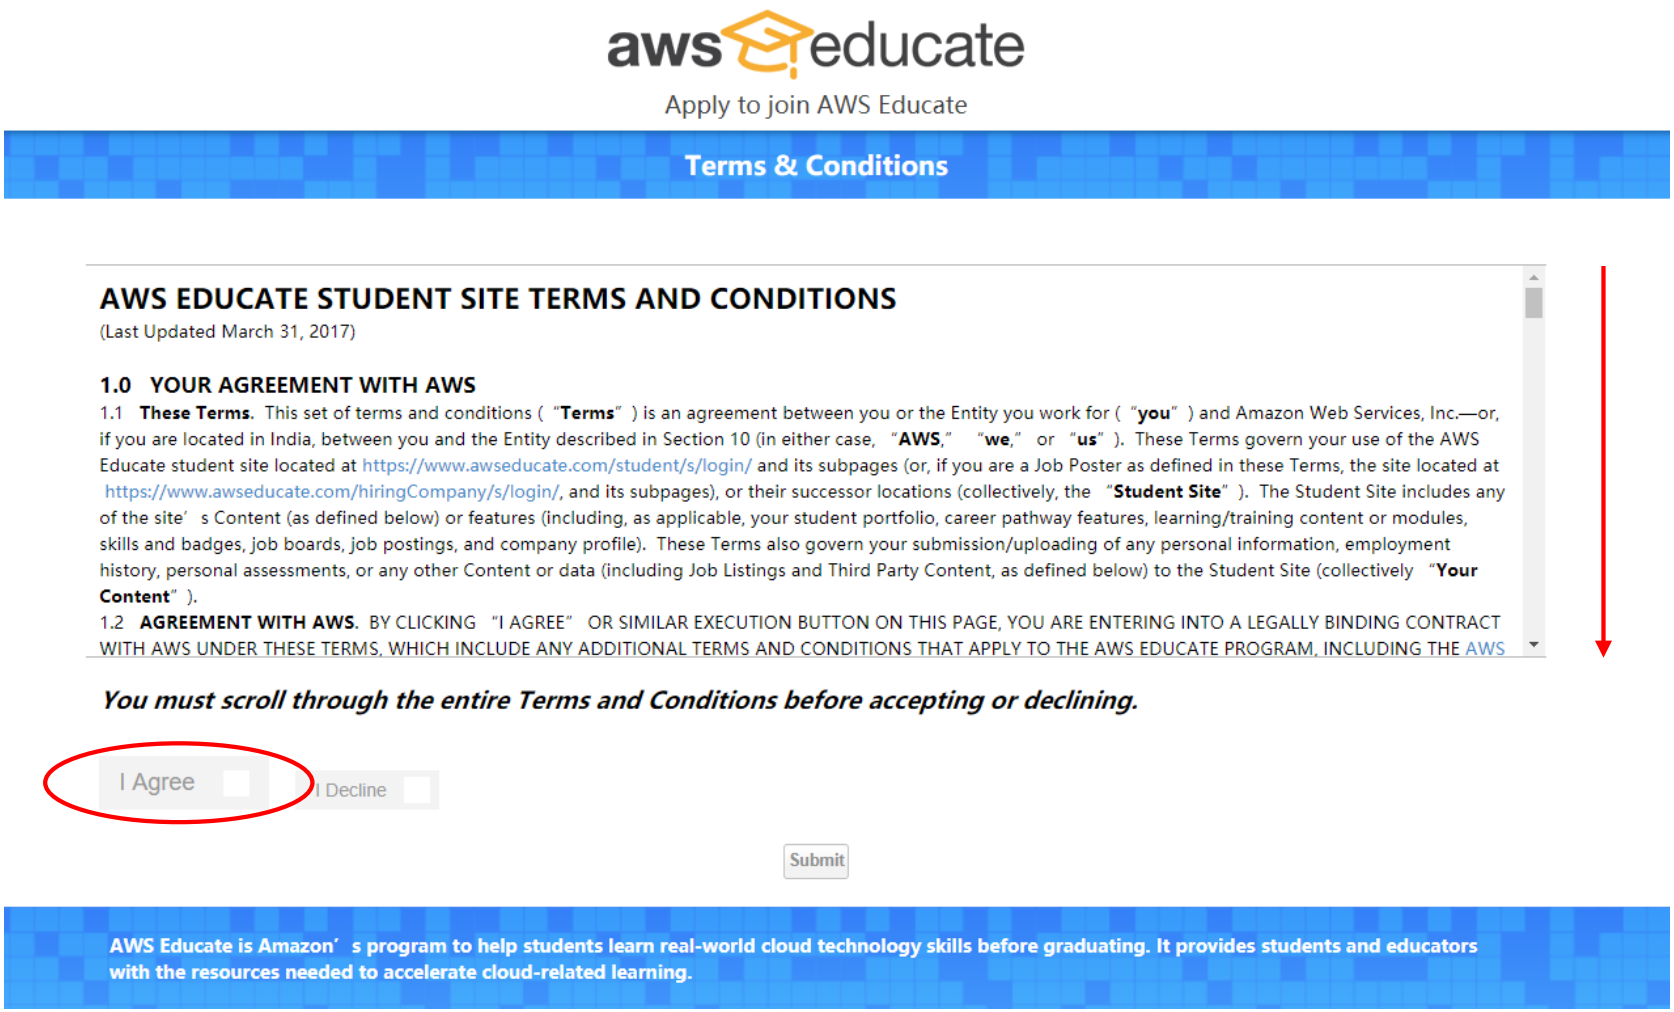

© 2017 Paul Krzyzanowski

Set up a free AWS educate account

- Then check your email, and you will find the link and the credit code.

Click the link in Step 1 and login your AWS account.
Then follow the instructions and enter your credit code.

AWS Educate Support <support@awseducate.com> 10:02 PM (32 minutes ago)
to me

Dear Long,

Congratulations!

Your AWS Educate application has been approved. As a member of the AWS Educate program, you will gain access to the benefits listed below:

AWS Educate Student Portal

The AWS Educate Student Portal is the hub for AWS Educate students around the world to find AWS content to help with classwork, connect to self-paced labs and training resources.

Click here to set your password / login to the AWS Educate Student Portal. After logging in, click AWS Account at the top of the page to access AWS services, whether you entered an AWS ID or selected Starter Account on your application. Note that Starter Accounts are not eligible for AWS Free Tier.

Bookmark the AWS Educate Student Portal for easy access, or click here to sign in directly.

You can access a video walk-through of the AWS Educate Student portal here.

AWS Promotional Credit

It's our pleasure to issue you an Amazon Web Services (AWS) promotional credit code in the amount listed below.

Credit Amount (US): $100.00
Credit Code: PC22AEZXF9F6RBU

Here's how to redeem your promotional credit:

- Step 1: Visit: https://console.aws.amazon.com/billing/home#/credits
- Step 2: Follow the instructions and enter your promo code.

Free AWS Essentials Training

To access our foundational AWS Essentials online learning class for free and find other self-paced labs, you must have either and AWS account or an Amazon ID.

- If you have an AWS account, sign in and click here to receive these benefits.
- If you do not have an AWS account, click here and follow the instructions to create an Amazon ID to access these benefits.

Once you access the Training and Certification portal, click "Find Training" and search for "Technical" to easily locate and enroll in AWS Technical Essentials on-line training. You can access AWS training any time after setting up your account by clicking here.

Thank you again for participating in AWS Educate and we hope you enjoy the program!

The AWS Educate Team

You will receive $100 for using AWS which is able to support EMR around 100+ hours.

© 2017 Paul Krzyzanowski

Set up a free AWS educate account

- Find “EMR” on the AWS console page.

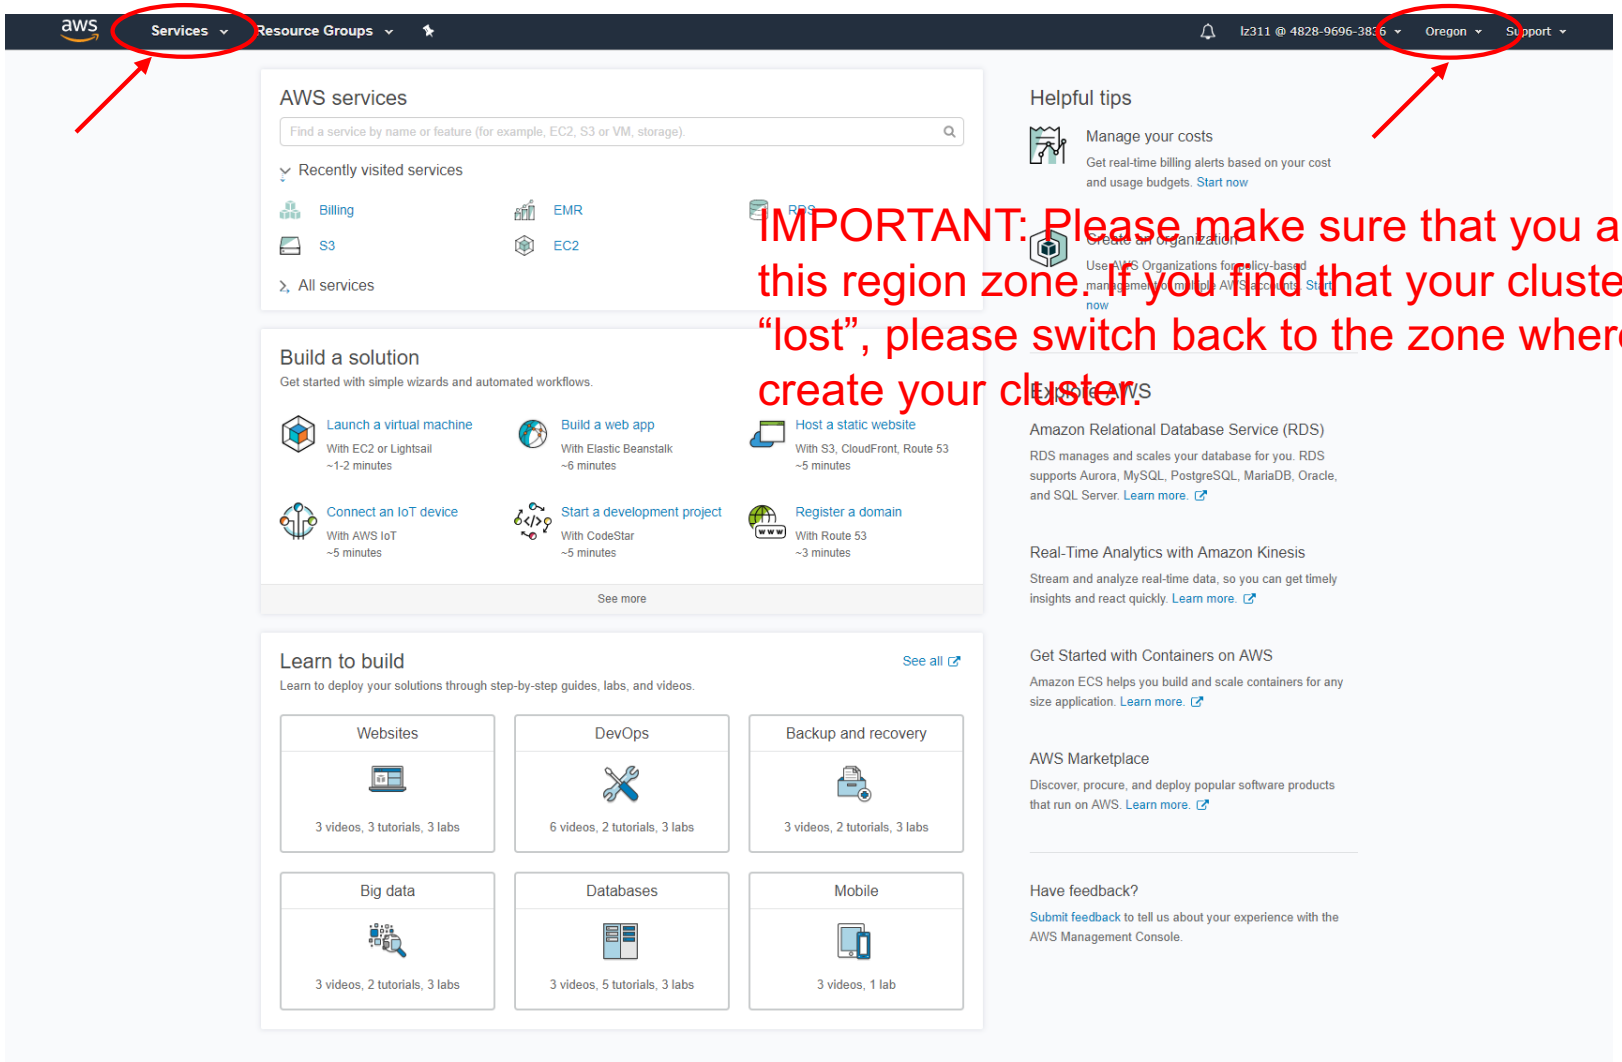

IMPORTANT: Please make sure that you are in this region zone. If you find that your cluster is “lost”, please switch back to the zone where you create your cluster.

© 2017 Paul Krzyzanowski

Set up a free AWS educate account

- Find “EMR” on the AWS console page.

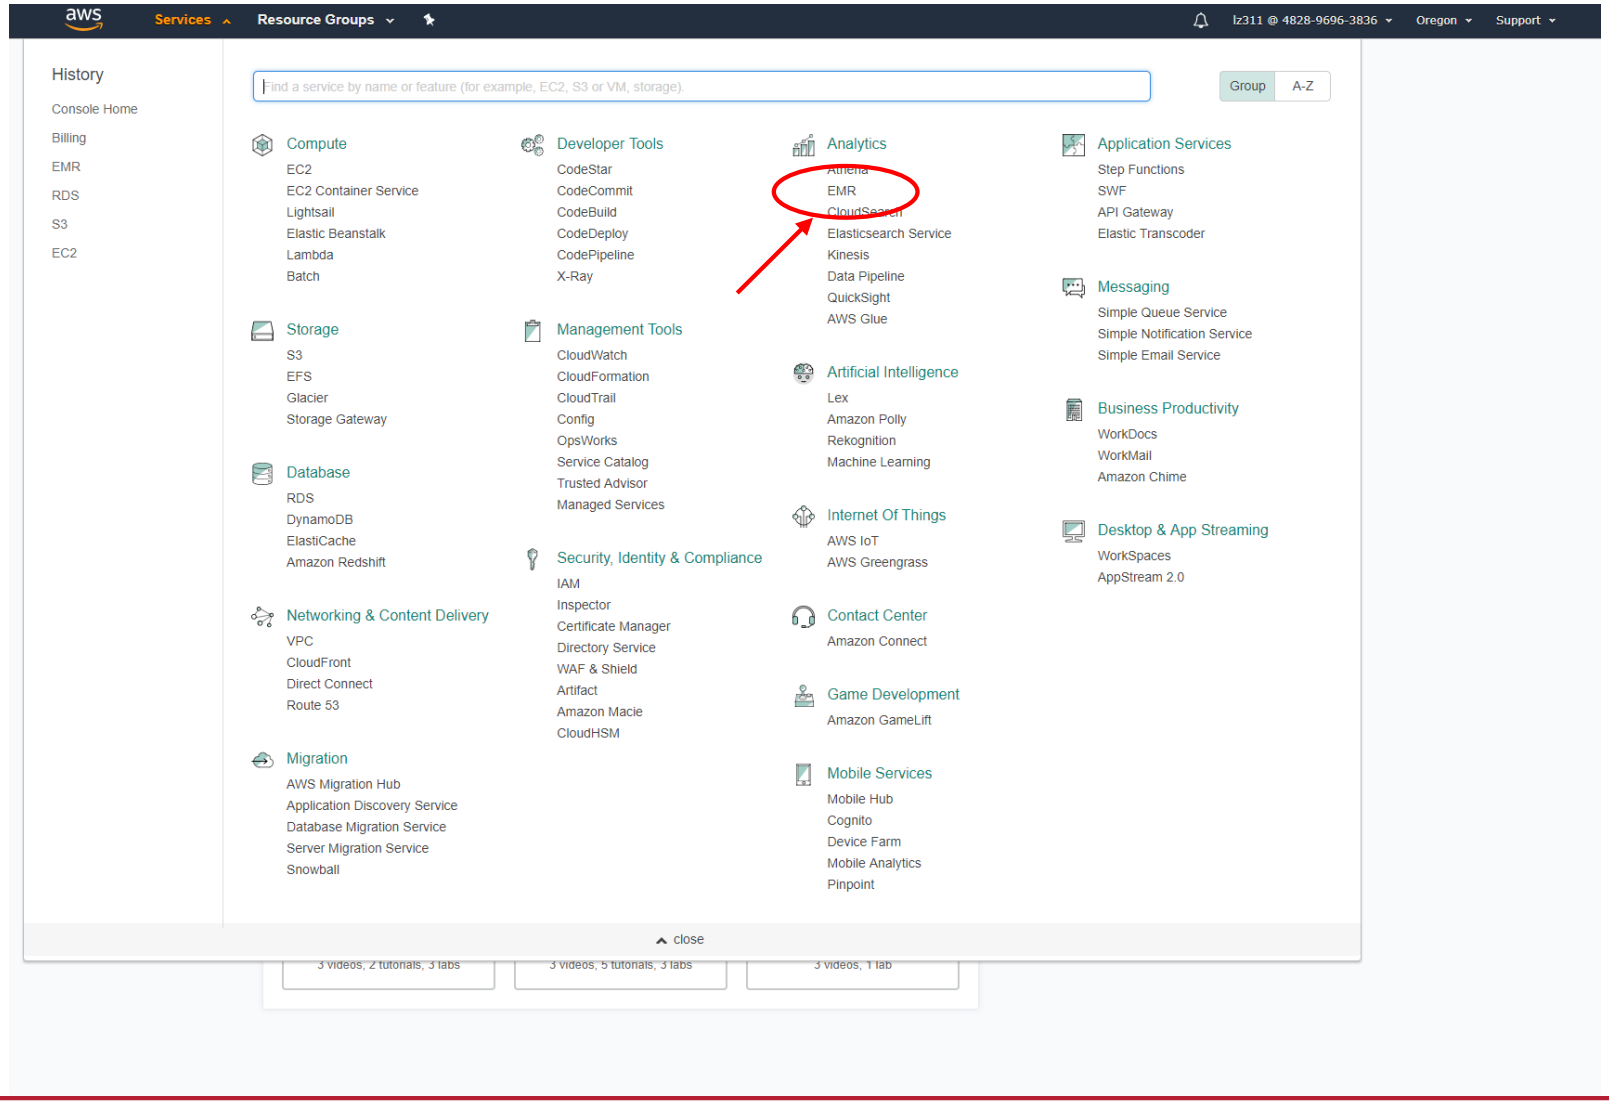

© 2017 Paul Krzyzanowski

Set up a free AWS educate account

- Create your cluster.

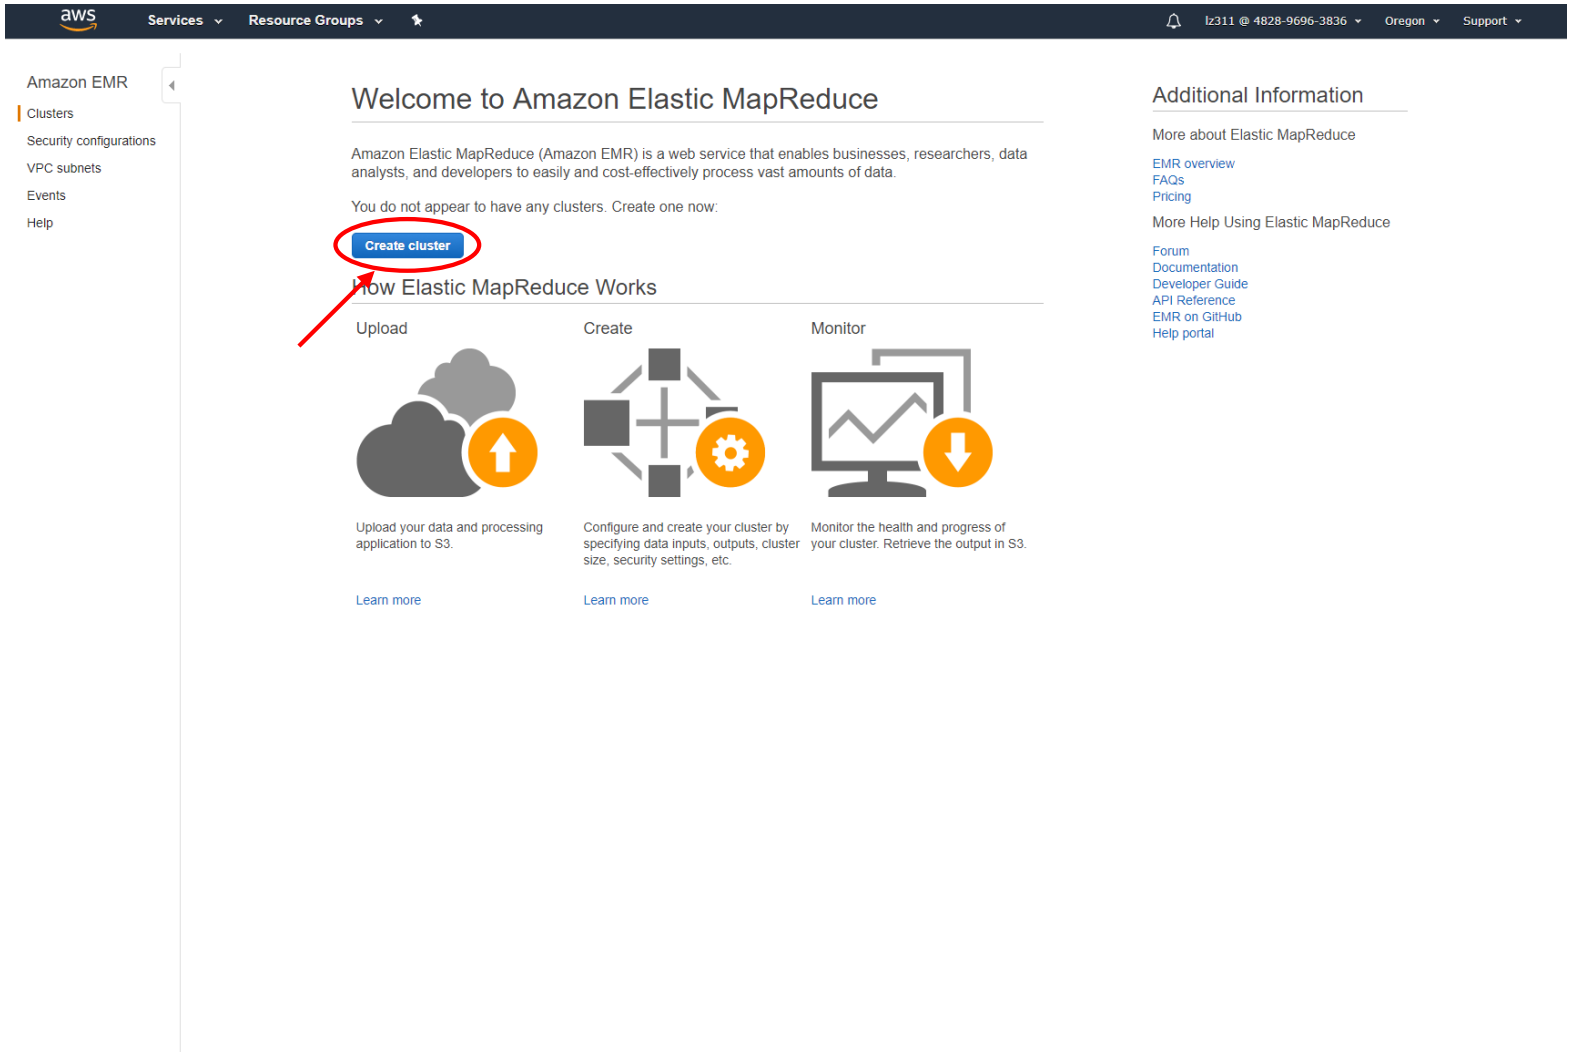

© 2017 Paul Krzyzanowski

Set up a free AWS educate account

- Create your cluster.

aws Services Resource Groups garyzhao9012 Oregon Support

Create Cluster - Quick Options Go to advanced options

General Configuration

Cluster name: My cluster

Logging

S3 folder: s3://aws-logs-006059551453-us-west-2/elasticmapred

Launch mode: Cluster Step execution

Software configuration

Release: emr-5.9.0

Applications:
- Core Hadoop: Hadoop 2.7.3 with Ganglia 3.7.2, Hive 2.3.0, Hue 4.0.1, Mahout 0.13.0, Pig 0.17.0, and Tez 0.8.4
- HBase: HBase 1.3.1 with Ganglia 3.7.2, Hadoop 2.7.3, Hive 2.3.0, Hue 4.0.1, Phoenix 4.11.0, and ZooKeeper 3.4.10
- Presto: Presto 0.170 with Hadoop 2.7.3 HDFS and Hive 2.3.0 Metastore
- Spark: Spark 2.2.0 on Hadoop 2.7.3 YARN with Ganglia 3.7.2 and Zeppelin 0.7.2

Use AWS Glue Data Catalog for table metadata

Hardware configuration

Instance type: m3.xlarge

Number of instances: 3 (1 master and 2 core nodes)

Security and access

EC2 key pair: Choose an option Learn how to create an EC2 key pair.

Permissions: Default Custom

Use default IAM roles. If roles are not present, they will be automatically created for you with managed policies for automatic policy updates.

EMR role: EMR_DefaultRole

EC2 instance profile: EMR_EC2_DefaultRole

Cancel Create cluster

© 2017 Paul Krzyzanowski

Set up a free AWS educate account

- Create AWS key-pair.

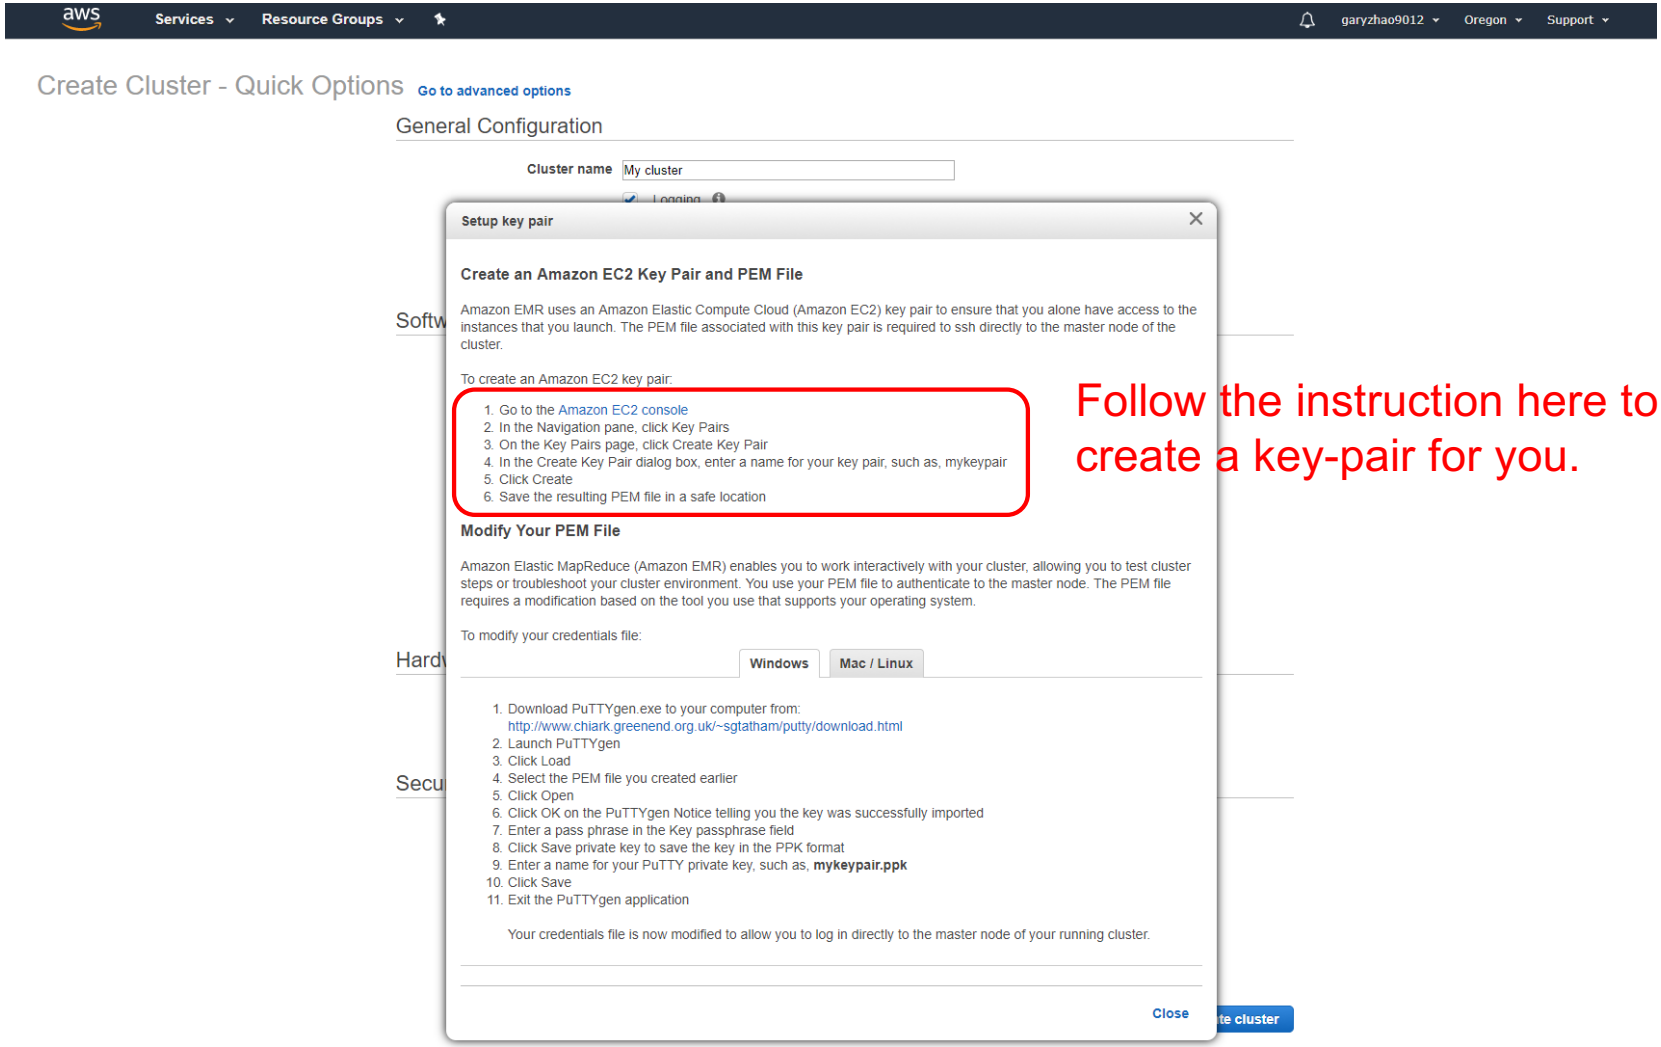

© 2017 Paul Krzyzanowski

Set up a free AWS educate account

- Create AWS key-pair.

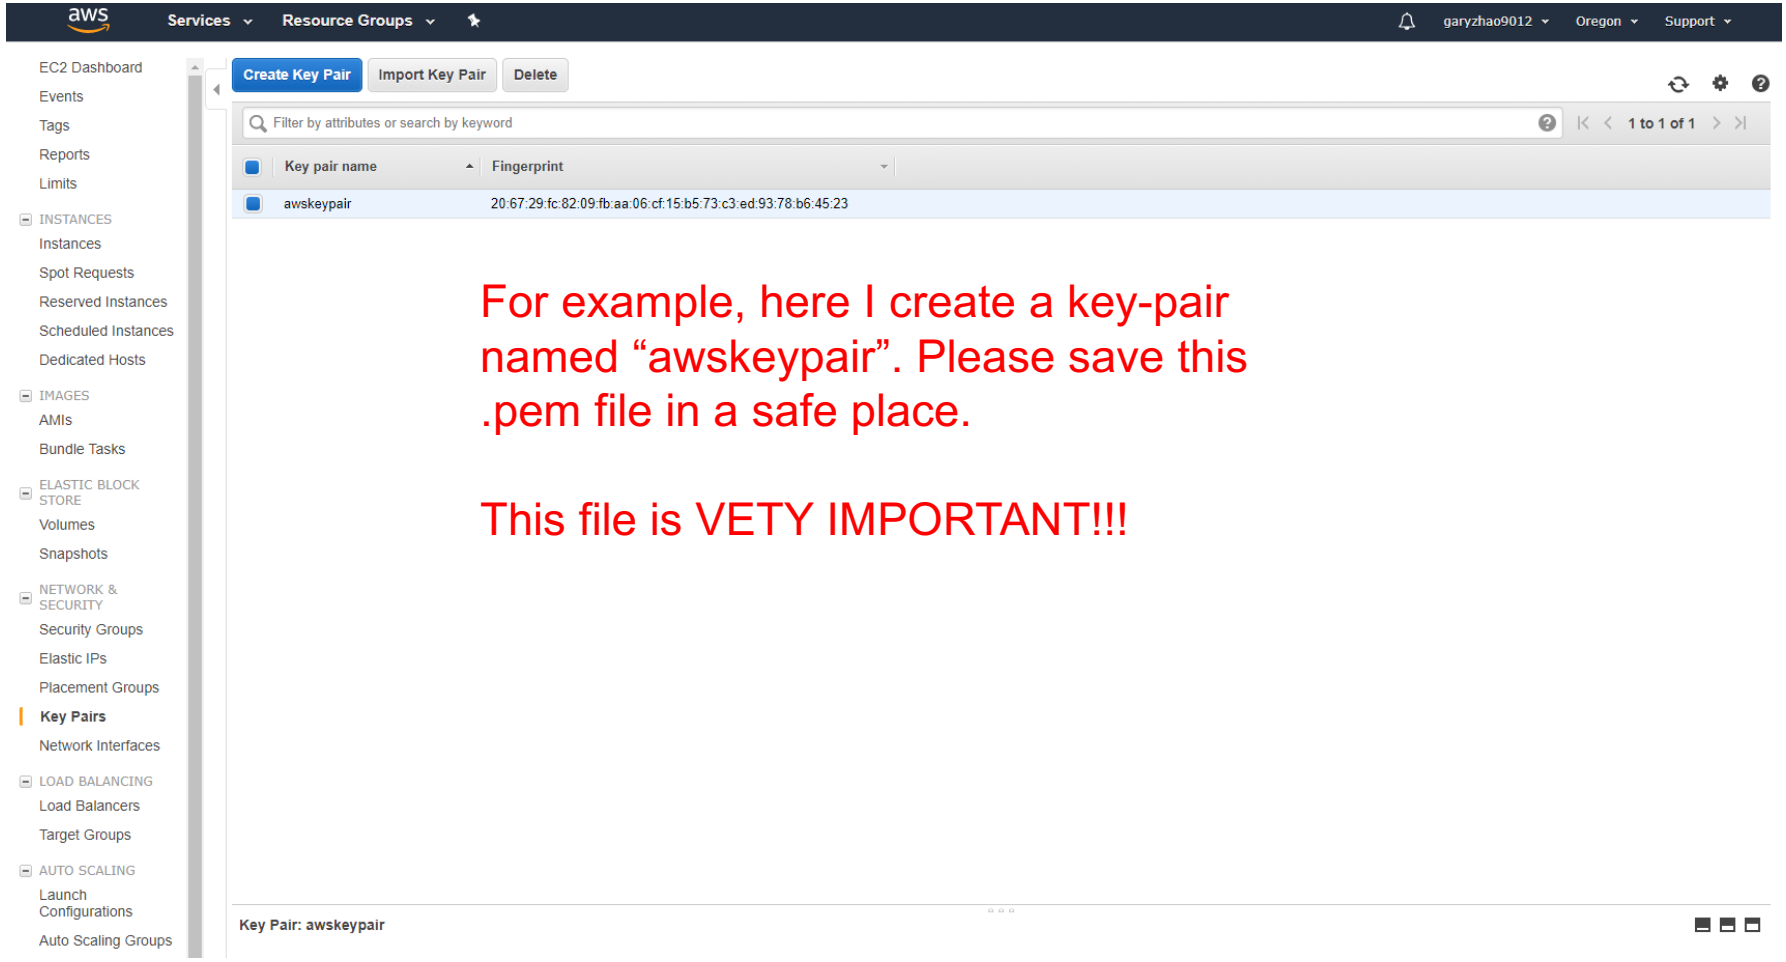

For example, here I create a key-pair named “awskeypair”. Please save this .pem file in a safe place.

This file is VETY IMPORTANT!!!

© 2017 Paul Krzyzanowski

Set up a free AWS educate account

- Then go back to the “EWR” page, choose the key-pair you just created and then “create cluster”.

aws Services Resource Groups garyzhao9012 Oregon Support

Create Cluster - Quick Options Go to advanced options

General Configuration

Cluster name My cluster

Logging

S3 folder s3://aws-logs-006059551453-us-west-2/elasticmapred

Launch mode Cluster Step execution

Software configuration

Release emr-5.9.0

Applications Core Hadoop: Hadoop 2.7.3 with Ganglia 3.7.2, Hive 2.3.0, Hue 4.0.1, Mahout 0.13.0, Pig 0.17.0, and Tez 0.8.4

HBase: HBase 1.3.1 with Ganglia 3.7.2, Hadoop 2.7.3, Hive 2.3.0, Hue 4.0.1, Phoenix 4.11.0, and ZooKeeper 3.4.10

Presto: Presto 0.170 with Hadoop 2.7.3 HDFS and Hive 2.3.0 Metastore

Spark: Spark 2.2.0 on Hadoop 2.7.3 YARN with Ganglia 3.7.2 and Zeppelin 0.7.2

Use AWS Glue Data Catalog for table metadata

Hardware configuration

Instance type m3.xlarge

Number of instances 3 (1 master and 2 core nodes)

Security and access

EC2 key pair awskeypair Learn how to create an EC2 key pair.

Permissions Default Custom

Use default IAM roles. If roles are not present, they will be automatically created for you with managed policies for automatic policy updates.

EMR role EMR_DefaultRole

EC2 instance profile EMR_EC2_DefaultRole

Cancel Create cluster

© 2017 Paul Krzyzanowski

Set up a free AWS educate account

- Wait several minutes until the cluster is created.

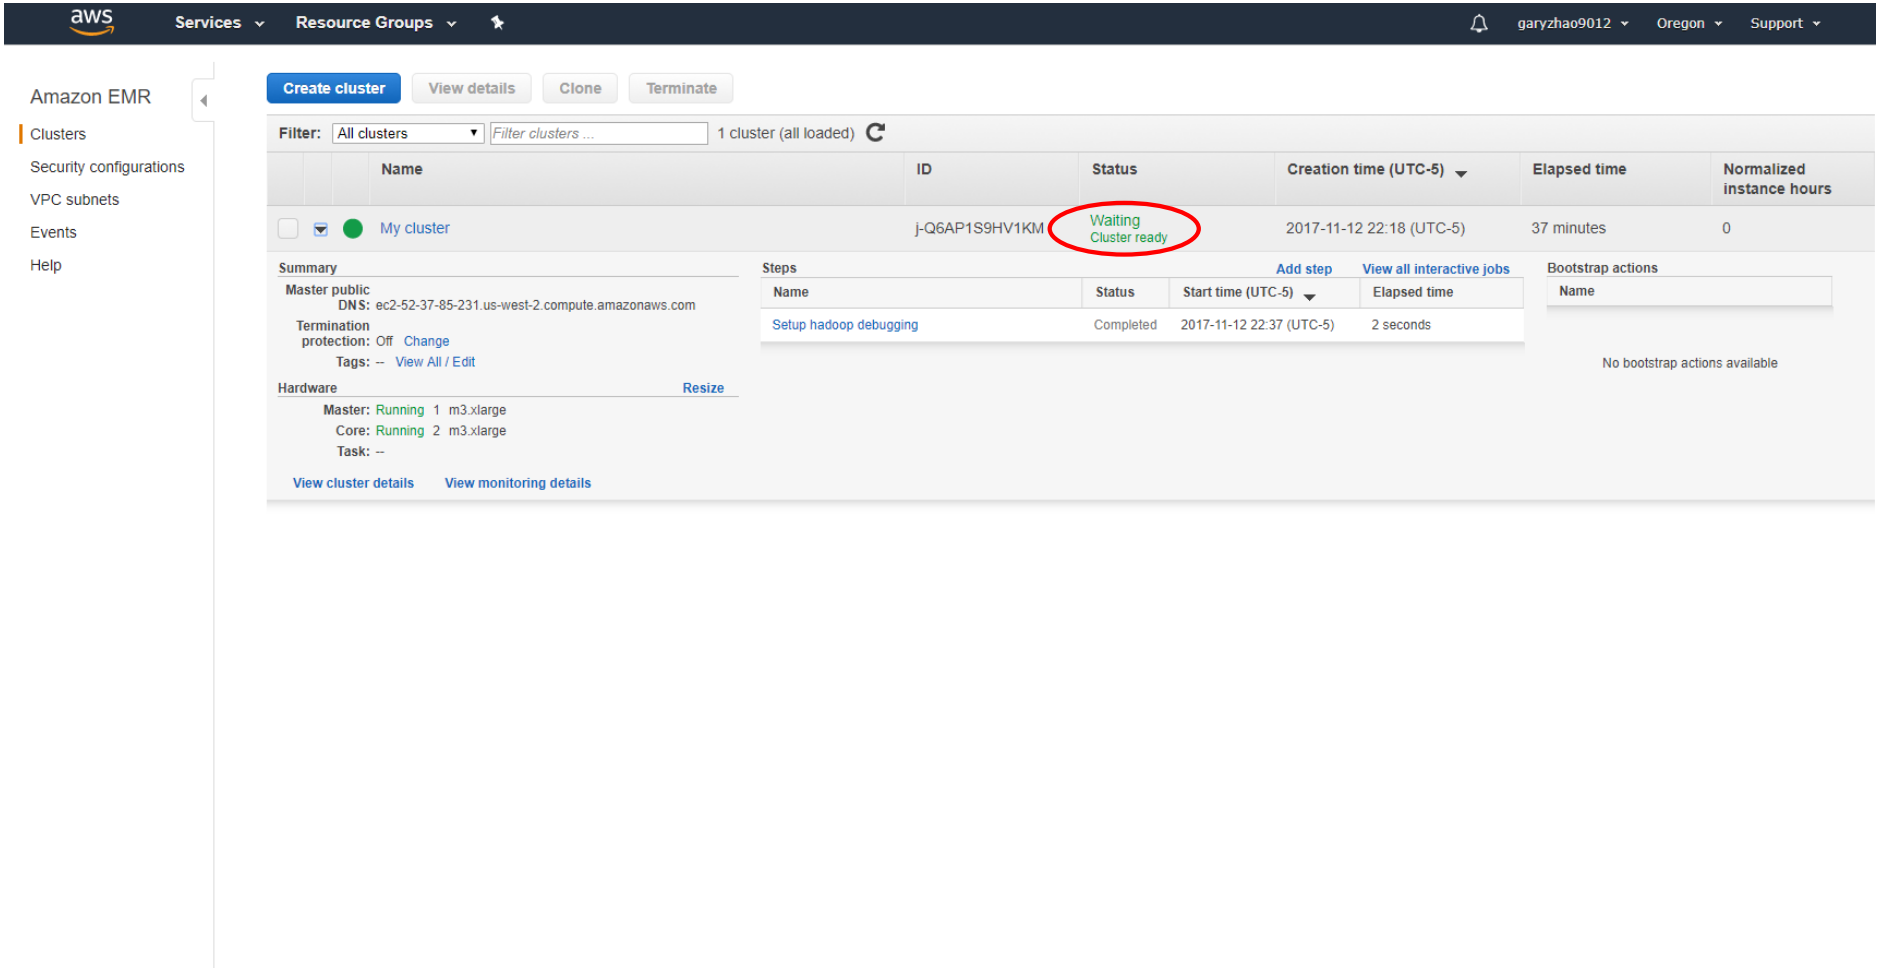

© 2017 Paul Krzyzanowski

Set up a free AWS educate account

- Configure Security Groups. Click “My cluster”.

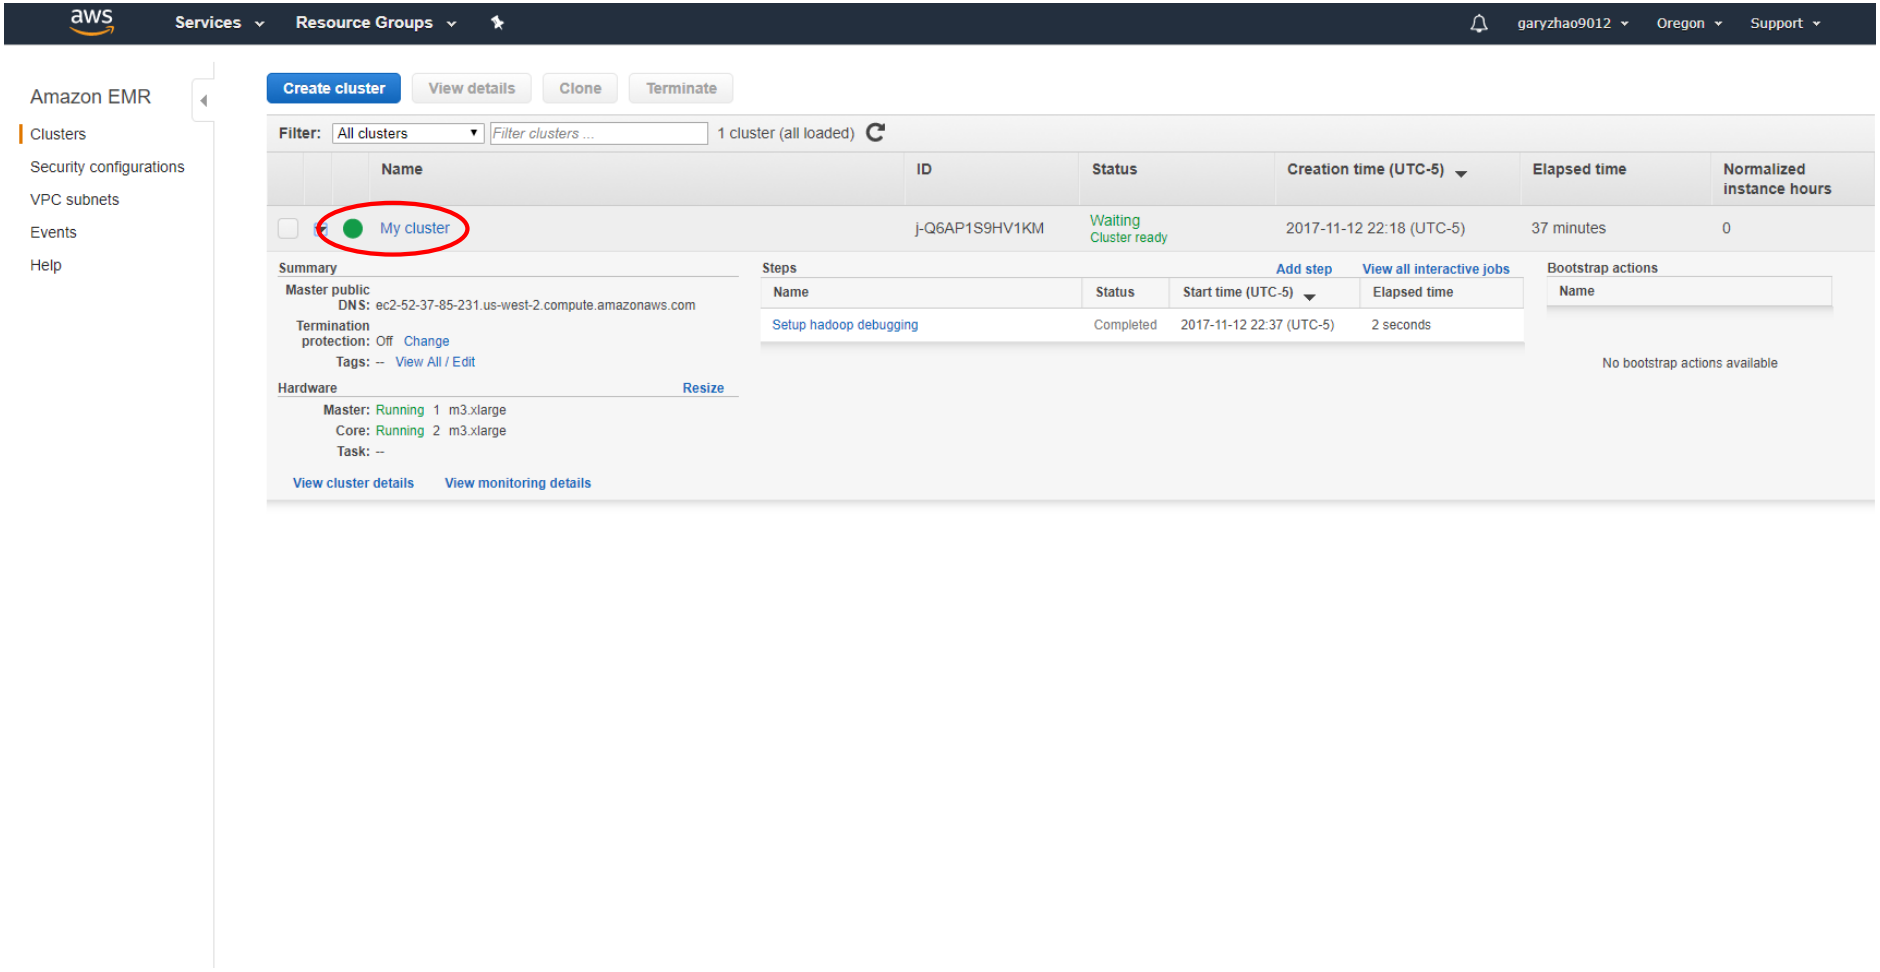

© 2017 Paul Krzyzanowski

Set up a free AWS educate account

- Choose “Security Groups” for the master.

aws Services Resource Groups garyzhao9012 Oregon Support

Amazon EMR
- Clusters
- Security configurations
- VPC subnets
- Events
- Help

Clone | Terminate | AWS CLI export

Cluster: My cluster Waiting Cluster ready after last step completed.

Summary | Application history | Monitoring | Hardware | Events | Steps | Configurations | Bootstrap actions

Connections: Enable Web Connection – Hue, Ganglia, Resource Manager ... (View All)
Master public DNS: ec2-52-37-85-231.us-west-2.compute.amazonaws.com SSH
Tags: -- View All / Edit

Summary
ID: j-Q6AP1S9HV1KM
Creation date: 2017-11-12 22:18 (UTC-5)
Elapsed time: 55 minutes
Auto-terminate: No
Termination protection: Off Change

Configuration details
Release label: emr-5.9.0
Hadoop distribution: Amazon 2.7.3
Applications: Ganglia 3.7.2, Hive 2.3.0, Hue 4.0.1, Mahout 0.13.0, Pig 0.17.0, Tez 0.8.4
Log URI: s3://aws-logs-006059551453-us-west-2/elasticmapreduce/
EMRFS consistent view: Disabled
Custom AMI ID: --

Network and hardware
Availability zone: us-west-2b
Subnet ID: subnet-f83958b0
Master: Running 1 m3.xlarge
Core: Running 2 m3.xlarge
Task: --

Security and access
Key name: awskeypair
EC2 instance profile: EMR_EC2_DefaultRole
EMR role: EMR_DefaultRole
Visible to all users: All Change
Security groups for Master: sg-0bc2c576 (ElasticMapReduce-master)
Security groups for Core & Task: sg-d3393eae (ElasticMapReduce-slave)

© 2017 Paul Krzyzanowski

Set up a free AWS educate account

- Choose “Security Groups” for the master and “Inbound”.

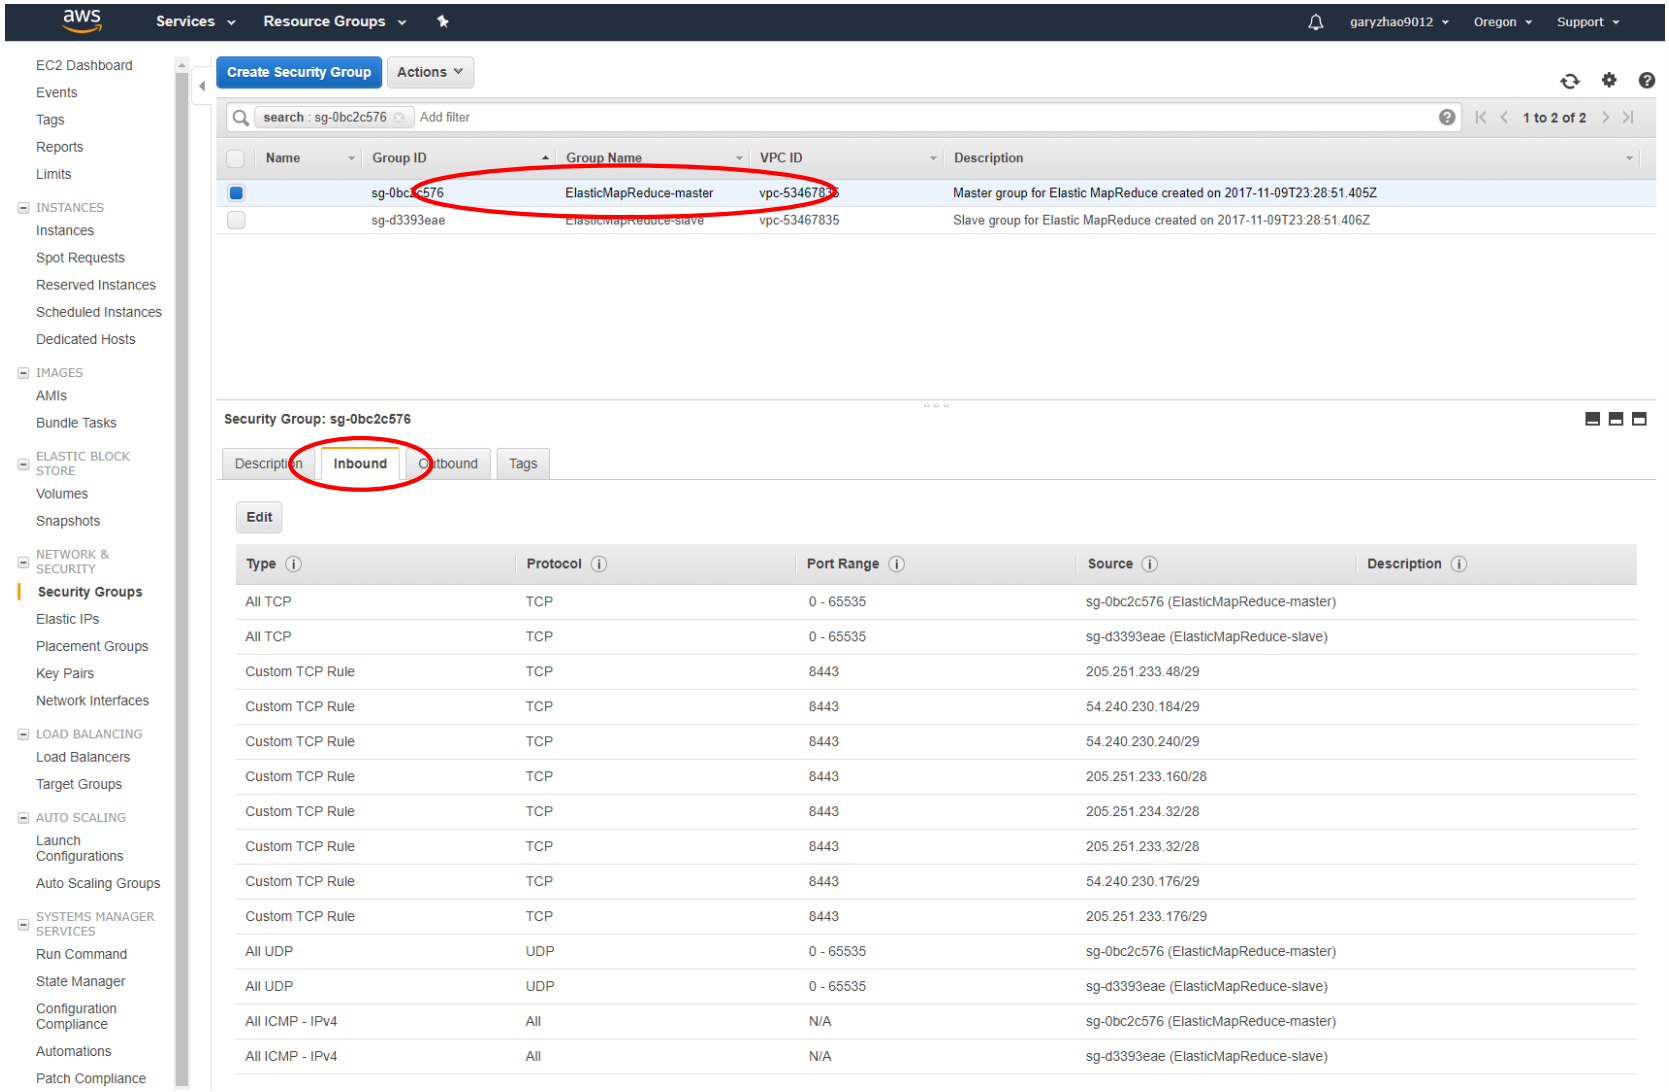

© 2017 Paul Krzyzanowski

Set up a free AWS educate account

- Add rules to make “TCP” and “ICMP” can be used form “Anywhere”.

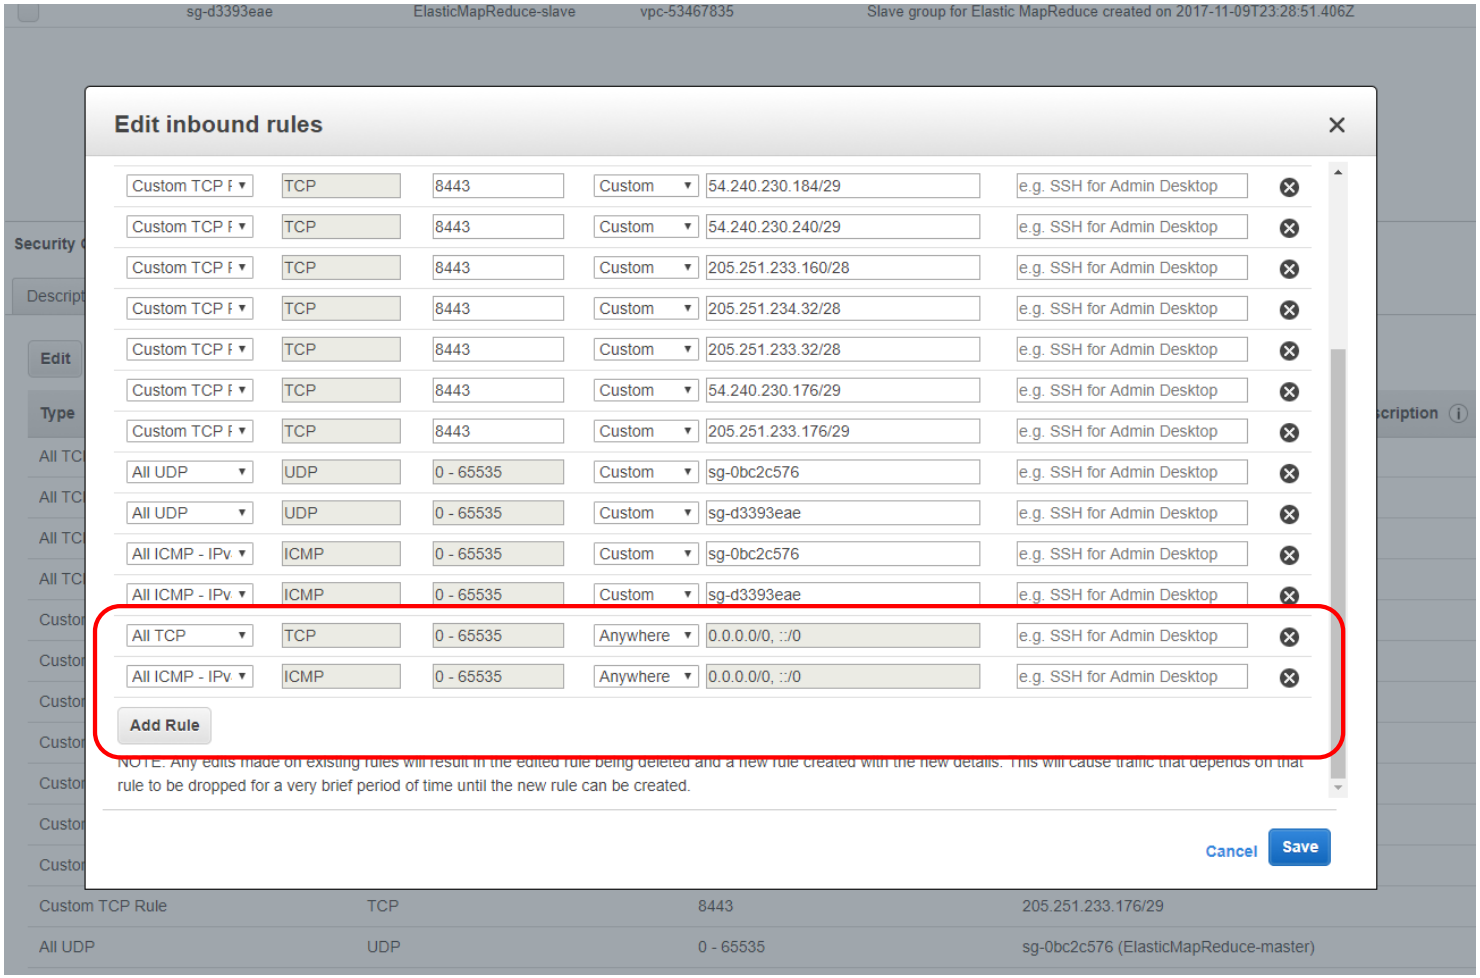

© 2017 Paul Krzyzanowski

Set up a free AWS educate account

- Add another rule for SSH. Then “Save”.

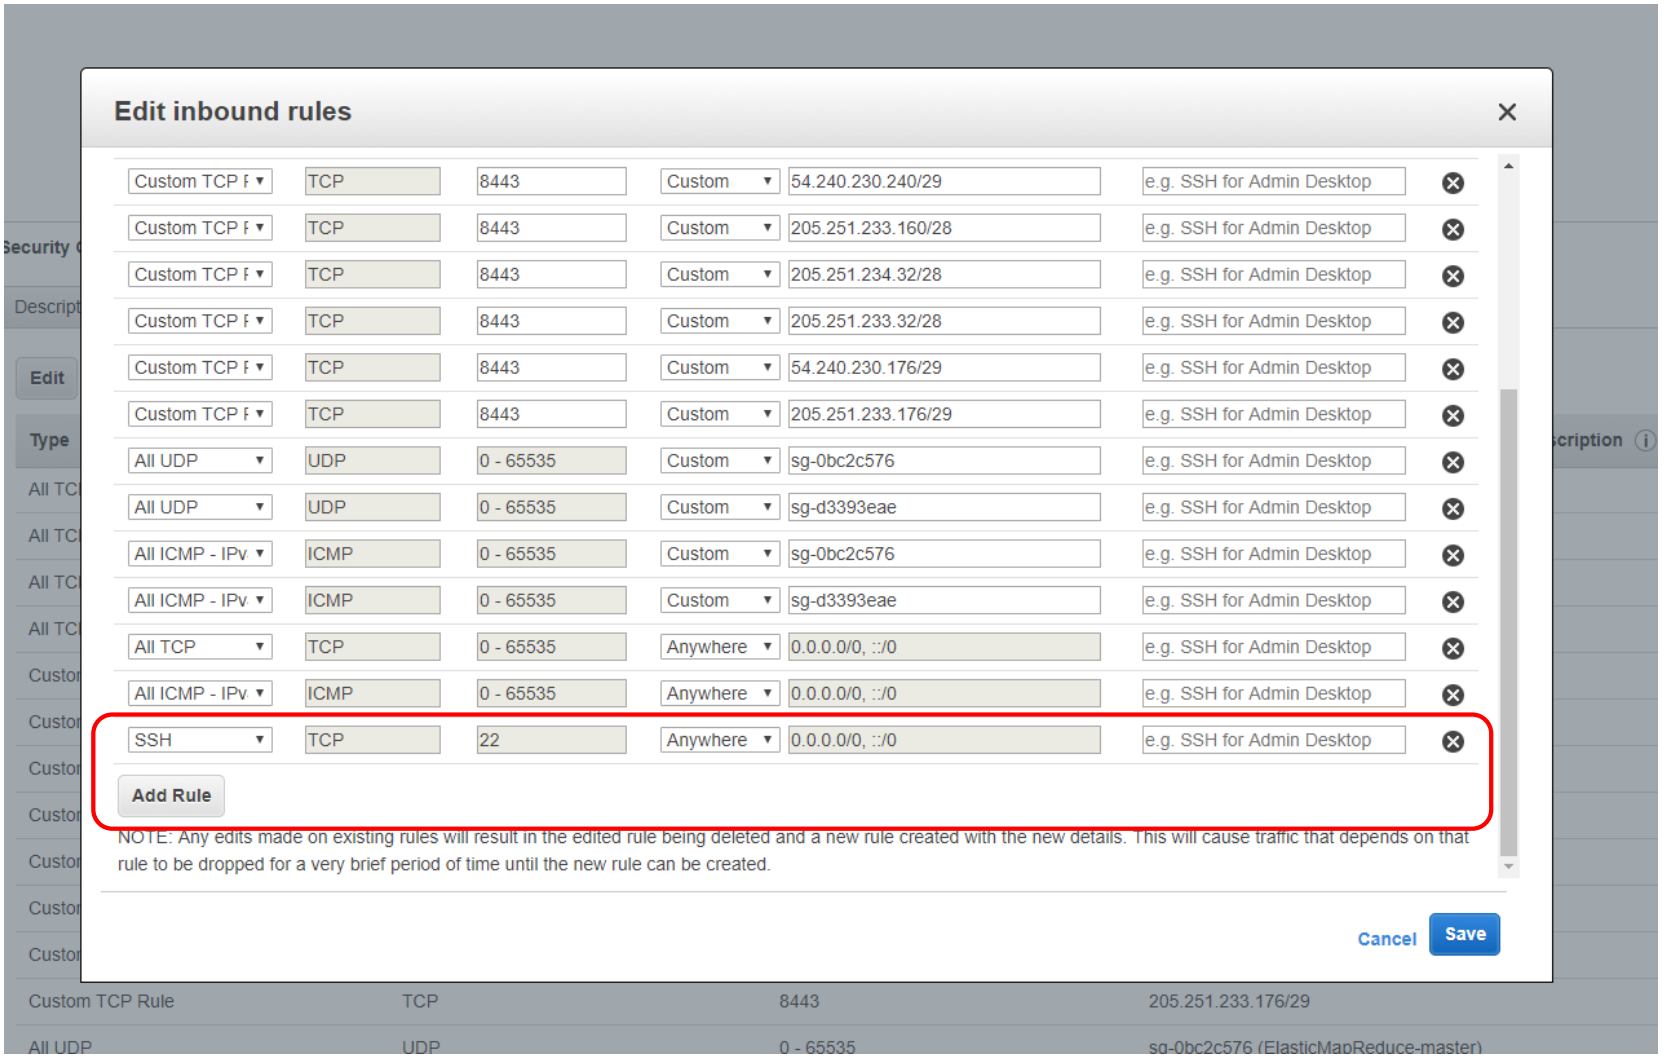

© 2017 Paul Krzyzanowski

Set up a free AWS educate account

- Then go back to “EMR” page to check your DNS of the master node.

aws Services Resource Groups garyzhao9012 Oregon Support

Amazon EMR
Clusters
Security configurations
VPC subnets
Events
Help

Create cluster | View details | Clone | Terminate

Filter: All clusters | Filter clusters ... | 1 cluster (all loaded)

Name	ID	Status	Creation time (UTC-5)	Elapsed time	Normalized instance hours
My cluster	j-Q6AP1S9HV1KM	Waiting Cluster ready	2017-11-12 22:18 (UTC-5)	1 hour, 13 minutes	0

Summary
Master public DNS: ec2-52-37-85-231.us-west-2.compute.amazonaws.com
Termination protection: Off Change
Tags: -- View All / Edit

Hardware Resize
Master: Running 1 m3.xlarge
Core: Running 2 m3.xlarge
Task: --

View cluster details View monitoring details

Steps Add step View all interactive jobs

Name	Status	Start time (UTC-5)	Elapsed time
Setup hadoop debugging	Completed	2017-11-12 22:37 (UTC-5)	2 seconds

Bootstrap actions

Name

No bootstrap actions available

© 2017 Paul Krzyzanowski

Set up a free AWS educate account

- Then we need to check if the cluster can be visited by us. You have do all the operations on a Linux machine or Mac OS. The iLab machine is highly recommended.

1. Ping the master DNS: Open the terminal, and type the command “ping <dns>”, where <dns> is the DNS of your master node.
2. Log into the master via SSH: In the terminal type the command “ssh -i <path-to-pem> hadoop@<dns>”, where <path-to-pem> is the path to the “.pem” key-pair file you have saved.

© 2017 Paul Krzyzanowski

Set up a free AWS educate account

- For me, if the “awskeypair.pem” file is in the current folder, I would type “ssh -i ./awskeypair.pem hadoop@ec2-52-37-85-231.us-west-2.compute.amazonaws.com”. If you see login information similar to the screenshot below, it means you have successfully set up the EMR cluster.

```
Long@GARY-DESKTOP /cygdrive/c/Users/Long/Dropbox
$ ssh -i ./awskeypair.pem hadoop@ec2-52-37-85-231.us-west-2.compute.amazonaws.com
The authenticity of host 'ec2-52-37-85-231.us-west-2.compute.amazonaws.com (52.37.85.231)' can't be established.
ECDSA key fingerprint is SHA256:EzDGyxqqmxVEpujCXU6nUaOg4ctW3bFT1yoUHQ6y/Q8.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added 'ec2-52-37-85-231.us-west-2.compute.amazonaws.com,52.37.85.231' (ECDSA) to the list of known hosts.
Last login: Mon Nov 13 04:38:26 2017

       __|  __|_  )
       _|  (     /   Amazon Linux AMI
      ___|\___|___|

https://aws.amazon.com/amazon-linux-ami/2017.03-release-notes/
2 package(s) needed for security, out of 216 available
Run "sudo yum update" to apply all updates.
Amazon Linux version 2017.09 is available.

EEEEEEEEEEEEEEEEEEEE MMMMMMMM           MMMMMMMM RRRRRRRRRRRRRRR
E::::::::::::::::::E M::::::M         M::::::M R::::::::::::::R
EE::::::EEEEEEEEE:::E M::::::::M       M::::::::M R::::::RRRRRR:::::R
  E::::E       EEEEE M:::::::::M     M:::::::::M RR::::R      R::::R
  E::::E             M::::::M:::M   M:::M::::::M   R:::R      R::::R
  E::::::EEEEEEEEEE  M::::::M M:::M M:::M M::::::M   R:::RRRRRR:::::R
  E::::::::::::::E   M::::::M  M:::M:::M  M::::::M   R::::::::::::RR
  E::::::EEEEEEEEEE  M::::::M   M:::::M   M::::::M   R:::RRRRRR:::::R
  E::::E             M::::::M    M:::M    M::::::M   R:::R      R::::R
  E::::E       EEEEE M::::::M     MMM     M::::::M   R:::R      R::::R
EE::::::EEEEEEEE:::E M::::::M             M::::::M   R:::R      R::::R
E::::::::::::::::::E M::::::M             M::::::M RR::::R      R::::R
EEEEEEEEEEEEEEEEEEEE MMMMMMM              MMMMMMM RRRRRRR      RRRRRR

[hadoop@ip-172-31-37-101 ~]$ |
```


© 2017 Paul Krzyzanowski

Set up a free AWS educate account

- The following tips are **VERY IMPORTANT**. If you finished using your cluster, please remember to do the following steps. Or there will be a service charge for EMR, once $100 credits run out. (You will receive another $100 each year.)

1. Terminate the cluster

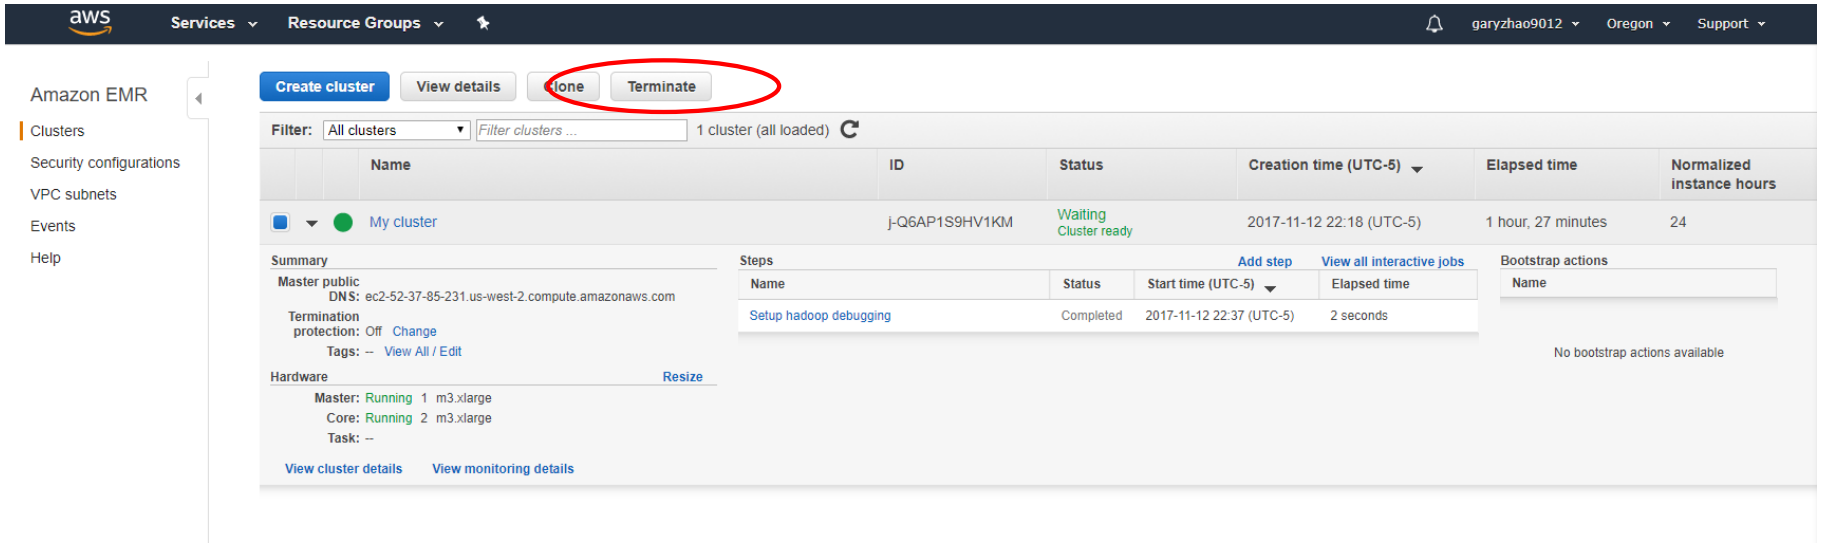

© 2017 Paul Krzyzanowski

Set up a free AWS educate account

- The following tips are **VERY IMPORTANT**. If you finished using your cluster, please remember to do the following steps. Or there will be a service charge for EMR, once $100 credits run out. If you need EMR again, please just repeat the above steps to create another cluster.

2. Delete S3 storage for the cluster by click "Services", then "S3".

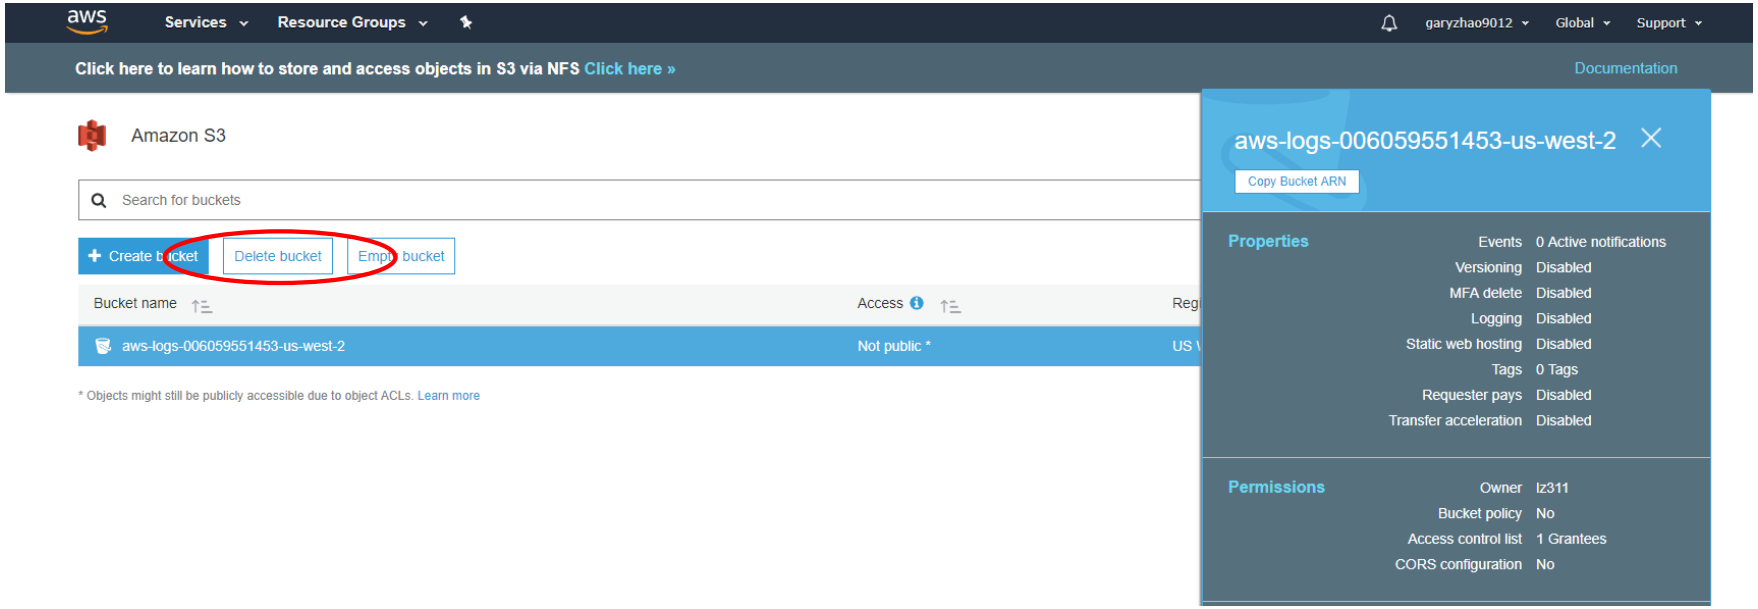

© 2017 Paul Krzyzanowski

Introduction to HDFS

© 2017 Paul Krzyzanowski

What is HDFS?

HDFS is an implementation of the Google File System (GFS) within the Apache Hadoop project – it is a large-scale distributed, parallel, fault-tolerant Java-based file system.

1. HDFS is a distributed file system that is fault tolerant, scalable and extremely easy to expand.
2. HDFS is the primary distributed storage for Hadoop applications.
3. HDFS provides interfaces for applications to move themselves closer to data.
4. HDFS is designed to ‘just work’, however a working knowledge helps in diagnostics and improvements.

© 2017 Paul Krzyzanowski

Components of HDFS

There are two (and a half) types of machines in a HDFS cluster:

- **NameNode** – is the heart of an HDFS filesystem. It maintains and manages the file system metadata; e.g., what blocks make up a file, and on which datanodes those blocks are stored.
- **DataNode** – where HDFS stores the actual data. There are usually quite a few of these.

© 2017 Paul Krzyzanowski

HDFS – Data Organization

1. Each file written into HDFS is split into data blocks
2. Each block is stored on one or more nodes
3. Each copy of the block is called replica
4. Block placement policy
   - First replica is placed on the local node
   - Second replica is placed in a different rack
   - Third replica is placed in the same rack as the second replica

© 2017 Paul Krzyzanowski

Interfaces to HDFS

- Java API `(DistributedFileSystem)`
- C wrapper `(libhdfs)`
- HTTP protocol
- WebDAV protocol
- Shell Commands*

*However, the command line is one of the simplest and most familiar.

© 2017 Paul Krzyzanowski

HDFS – Shell Commands

There are two types of shell commands:

- User Commands
  - `hdfs dfs` runs filesystem commands on the HDFS
  - `hdfs fsck` runs a HDFS filesystem checking command
- Administration Commands
  - `hdfs dfsadmin` runs HDFS administration commands

© 2017 Paul Krzyzanowski

HDFS – User Commands (dfs)

- List directory contents

```
>> hdfs dfs -ls
>> hdfs dfs -ls /
>> hdfs dfs -ls -R /var
```

- Display the disk space used by files

```
>> hdfs dfs -du -h /
>> hdfs dfs -du /hbase/data/hbase/
>> hdfs dfs -du -h /hbase/data/hbase/
>> hdfs dfs -du -s /hbase/data/hbase/
```


© 2017 Paul Krzyzanowski

HDFS – User Commands (dfs)

- Copy data to HDFS

```
>> hdfs dfs -mkdir tdata
>> hdfs dfs -ls
>> hdfs dfs -copyFromLocal tutorials/data/geneva.csv tdata
>> hdfs dfs -ls -R
```

- Copy the file back to local filesystem

```
>> cd tutorials/data/
>> hdfs dfs -copyToLocal tdata/geneva.csv
geneva.csv.hdfs
>> md5sum geneva.csv geneva.csv.hdfs
```


© 2017 Paul Krzyzanowski

HDFS – User Commands (dfs)

- List acl for a file

```
>> hdfs dfs -getfacl tdata/geneva.csv
```

- List the file statistics – (%r – replication factor)

```
>> hdfs dfs -stat "%r" tdata/geneva.csv
```

- Write to hdfs reading from stdin

```
>> echo "blah blah blah" | hdfs dfs -put -
tdataset/tfile.txt
>> hdfs dfs -ls -R
>> hdfs dfs -cat tdataset/tfile.txt
```


© 2017 Paul Krzyzanowski

HDFS – User Commands (fsck)

- Removing a file

```
>> hdfs dfs -rm tdataset/tfile.txt
>> hdfs dfs -ls -R
```

- List the blocks of a file and their locations

```
>> hdfs fsck /user/cloudera/tdata/geneva.csv
-files -blocks -locations
```

- Print missing blocks and the files they belong to

```
>> hdfs fsck / -list-corruptfileblocks
```


© 2017 Paul Krzyzanowski

HDFS – Adminstration Commands

- Comprehensive status report of HDFS cluster

```
>> hdfs dfsadmin -report
```

- Prints a tree of racks and their nodes

```
>> hdfs dfsadmin -printTopology
```

- Get the information for a given datanode (like ping)

```
>> hdfs dfsadmin -getDatanodeInfo
localhost:50020
```


© 2017 Paul Krzyzanowski

HDFS – Adminstration Commands

- Get a list of namenodes in the Hadoop cluster

```
>> hdfs getconf -namenodes
```

- Dump the NameNode fsimage to XML file

```
>> cd /var/lib/hadoop-
hdfs/cache/hdfs/dfs/name/current
hdfs oiv -i fsimage_0000000000000003388 -o
/tmp/fsimage.xml -p XML
```

The general command line syntax is

hdfs command [genericOptions] [commandOptions]

© 2017 Paul Krzyzanowski

Other Interfaces to HDFS

- HTTP Interface

```
http://<dns>:50070
```

Hadoop | Overview | Datanodes | Datanode Volume Failures | Snapshot | Startup Progress | Utilities

Overview 'ip-172-31-37-101.us-west-2.compute.internal:8020' (active)

Started:	Mon Nov 13 03:31:39 UTC 2017
Version:	2.7.3-amzn-4, r5f779390180acb8f1d86999c0a0294917976289f
Compiled:	2017-09-22T21:15Z by ec2-user from (HEAD detached at 5f779390)
Cluster ID:	CID-5f0e0b29-4159-4cad-a6d6-c96e05ca974e
Block Pool ID:	BP-2118224034-172.31.37.101-1510543894759

Summary

Security is off.

Safemode is off.

1051 files and directories, 1014 blocks = 2065 total filesystem object(s).

Heap Memory used 196.21 MB of 503 MB Heap Memory. Max Heap Memory is 1.51 GB.

Non Heap Memory used 60.85 MB of 62.19 MB Commited Non Heap Memory. Max Non Heap Memory is -1 B.

Configured Capacity:	137.9 GB
DFS Used:	770.69 MB (0.55%)
Non DFS Used:	344.26 MB
DFS Remaining:	136.81 GB (99.21%)
Block Pool Used:	770.69 MB (0.55%)

© 2017 Paul Krzyzanowski

Other Useful Links

- iLab: https://www.cs.rutgers.edu/resources/instructional-lab
- Amazon EMR Official Documentation: https://aws.amazon.com/documentation/emr/
- HDFS Architecture Guide: http://hadoop.apache.org/docs/stable/hadoop-project-dist/hadoop-hdfs/HdfsDesign.html
- HDFS File System Shell Guide: http://hadoop.apache.org/docs/stable/hadoop-project-dist/hadoop-common/FileSystemShell.html

© 2017 Paul Krzyzanowski

The end

© 2017 Paul Krzyzanowski