



INSTITUTE FOR APPLIED
COMPUTATIONAL SCIENCE
AT HARVARD UNIVERSITY



HARVARD
School of Engineering
and Applied Sciences

Guide: First Access to AWS

Ignacio M. Llorente, David Sondak, Dylan Randle

v2.0 - February 3, 2020

Abstract

This is a screenshot document of how to setup your AWS environment, launch VMs with EC2 and manage storage with S3.

Notes

- First you will have to sign up for an **AWS regular account**, if you don't already have one.
- AWS recently limited the range of services available for AWS Educate accounts. Instead, you should have received AWS credits that you can apply to your regular account.
 - See this link for info about redeeming credits: <https://aws.amazon.com/awscredits/>

**First apply for a regular account and then sign up for AWS Educate as a student to receive a credit code
Do not apply for an AWS Educate Starter Account**

- Set up a **billing alert** to make sure you don't accidentally use up your free credits without noticing.
- **Stop** your instances when are done for the day to avoid incurring charges. Use your funds wisely. **Terminate** them when you are sure you are done with your instance (disk storage also costs something, and can be significant if you have a large disk footprint). Look into creating custom alarms to automatically stop your instances when they are not doing anything.
- This guide has been prepared considering that you are using Linux or Mac OS to connect to the remote instance. If you are using windows we recommend you read:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/putty.html>



- This guide describes a basic use to develop and execute the course hands-on. For further use we recommend you read the AWS guidelines to create IAM users and VPCs:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/get-set-up-for-amazon-ec2.html>

- We strongly recommend you first read this introductory Amazon EC2 guide that presents the basic concepts: instances, AMIs, security groups, root devices, regions and availability zones.

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

- We strongly recommend you first read this introductory Amazon S3 guide that presents the basic concepts: buckets and objects.

<http://docs.aws.amazon.com/AmazonS3/latest/dev/Introduction.html>

Acknowledgments

The author is grateful for constructive comments and suggestions from David Sondak, Charles Liu, Matthew Holman, Keshavamurthy Indireskumar, Kar Tong Tan, Zudi Lin and Nick Stern.

1. Create a key pair to connect to the cloud VMs

You can skip this step if you already have a key pair from other courses.

- Login to AWS, go to the EC2 dashboard, select “key pairs” on the left hand menu, and click “Create Key Pair”. We recommend you name it “CS205-key”.

EC2 > Key pairs > Create key pair

Create key pair

Key pair
A key pair, consisting of a private key and a public key, is a set of security credentials that you use to prove your identity when connecting to an instance.

Name
CS205-key
The name can be up to 255 characters long. Valid characters include _ , -, a-z, A-Z, and 0-9.

File format

pem
For use with OpenSSH

ppk
For use with PuTTY

Cancel Create key pair

- Download the private key locally and copy it to the `.ssh` folder. In my case, and for illustration:

```
$ mv $HOME/Downloads/CS205-key.pem $HOME/.ssh/CS205-key.pem
```

- Change the permission of the file


```
$ chmod 600 $HOME/.ssh/CS205-key.pem
```

- Remember the key path and name (`$HOME/.ssh/CS205-key.pem`)

2. Launch a VM

We are going to launch an Amazon EBS-backed instance (meaning that the root volume is an EBS¹ volume). We will let Amazon EC2 select an Availability Zone for us .

¹ EBS stands for “Elastic Block Storage”



- Go to the EC2 dashboard and click “Launch Instance”.

- Step 1: Select “Ubuntu Server 18.04” as AMI

- Step 2: Select “t2.micro” as instance type (“Free Tier Eligible”). It is important to select EBS backed



instances for persistency².

Step 2: Choose an Instance Type

Amazon EC2 provides a wide selection of instance types optimized to fit different use capacity, and give you the flexibility to choose the appropriate mix of resources for yo

Filter by: All instance types Current generation Show/Hide C

Currently selected: t2.micro (Variable ECUs, 1 vCPUs, 2.5 GHz, Intel Xeon Family,

	Family	Type	vCPUs
<input type="checkbox"/>	General purpose	t2.nano	1
<input checked="" type="checkbox"/>	General purpose	t2.micro Free tier eligible	1

- Step 3: Click “Next” in the bottom right corner to modify the configurations. Use default configurations for the rest of steps. It is highly recommended to revise and understand all options (especially the “Configure Security Group” option. Changing the “source” from “custom” to “My IP” is safer). At the final step, click “Launch”.
- Step 4: Select your key pair and “Launch 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

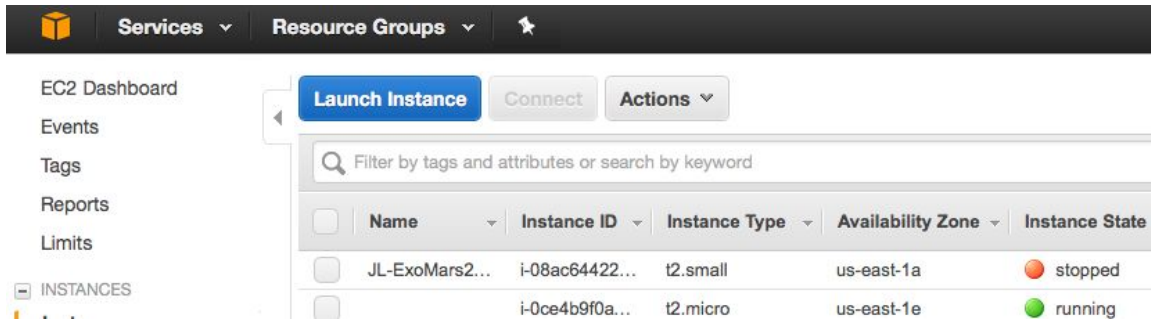
course-key

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

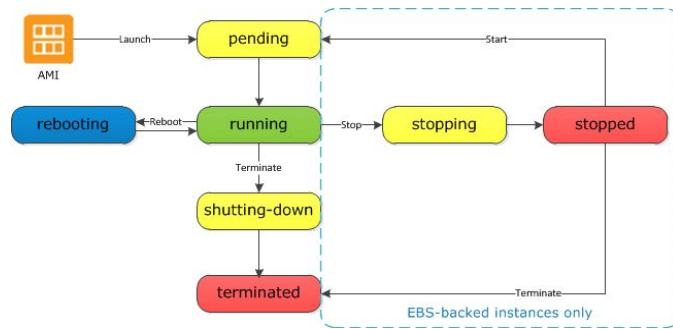
Cancel Launch Instances

- Step 5: Go to “Running Instances” in EC2 Dashboard and wait for the VM to be “running”.

² EBS will save the data on the drive after the instance is stopped.



- Step 6: Familiarize yourself with the actions and the status and monitoring information provided by the dashboard.



3. Login to the VM

After you launch your instance, you can connect to it and use it the way that you'd use a computer sitting in front of you. It can take a few minutes for the instance to be ready so that you can connect to it. Check that your instance has passed its status checks - you can view this information in the Status Checks column on the Instances page.

To connect to your Linux instance from a computer running Mac or Linux, you'll specify the `.pem` file to your SSH client with the `-i` option and the path to your private key. To connect to your Linux instance from a computer running Windows, you can use either MindTerm or PuTTY. If you plan to use PuTTY, you'll need to install it and use the following procedure to convert the `.pem` file to a `.ppk` file: [Connecting to Your Linux Instance from Windows Using PuTTY](#)

The Linux procedure is as follows:

- Select the instance, and then choose Connect.



Connect To Your 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 (course-key.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 course-key.pem
```

4. Connect to your instance using its Public IP:

```
34.230.37.255
```

Example:

```
ssh -i "course-key.pem" ubuntu@34.230.37.255
```

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](#).

Close

- Execute the SSH command to login to your VM

```
ssh -i $HOME/.ssh/CS205-key.pem ubuntu@34.230.37.25
```

```
nacho — ubuntu@ip-172-30-4-13: ~ — ssh -i ~/.ssh/course-key.pem ubuntu@34.230.37.255 — 90x25
Welcome to Ubuntu 16.04.2 LTS (GNU/Linux 4.4.0-1022-aws x86_64)

 * Documentation:  https://help.ubuntu.com
 * Management:    https://landscape.canonical.com
 * Support:       https://ubuntu.com/advantage

Get cloud support with Ubuntu Advantage Cloud Guest:
http://www.ubuntu.com/business/services/cloud

0 packages can be updated.
0 updates are security updates.

The programs included with the Ubuntu system are free software;
the exact distribution terms for each program are described in the
individual files in /usr/share/doc/*/copyright.

Ubuntu comes with ABSOLUTELY NO WARRANTY, to the extent permitted by
applicable law.

To run a command as administrator (user "root"), use "sudo <command>".
See "man sudo_root" for details.

ubuntu@ip-172-30-4-13:~$ █
```




In most cases, the public IP address is associated with the instance until it's stopped or terminated, after which it's no longer available for you to use. If you require a persistent public IP address that you can associate and disassociate at will, use an Elastic IP address (EIP) instead. You can allocate your own EIP, and associate it to your instance after launch.

4. Not able to login?

If you could not login, then make sure that the security groups (firewalls) of the VM opens the port 22 to the outside world. Under the security access of the cluster console, check the security groups.

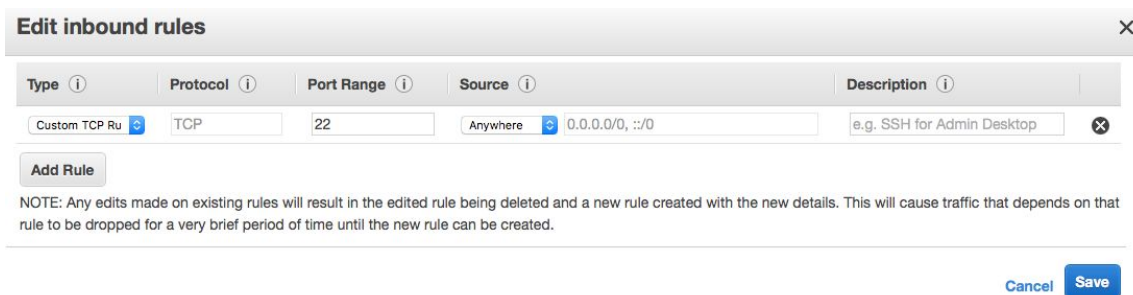
- Click on the Security groups of the VM, id:"sg-9d898fed" (your security groups will have different IDs but it starts with "sg-"). This can be found on the left pane under "Network and Security".



- Select the Inbound tab and click on Edit



- Click on "Add Rule", a new row will be created and fill out the new row with the following:



- Click "Save"

If you still cannot login, make sure your VPC route table is configured to allow traffic to and from the Internet

- Open the Amazon **VPC** console.



CS205: Computing Foundations for Computational Science, Spring 2020

- In the navigation pane, choose **Route Tables** and then select your VPC route table from the list.
- On the Routes tab, ensure that you have a default route pointing to your **Internet gateway** (IGW).
- If you do not see this, choose Internet Gateways from the navigation pane and copy the ID of your Internet gateway. If you do not have an Internet gateway, create one and attach it to your VPC. Be sure to copy the ID of the new IGW
- Go back to Route Tables and select the Routes tab.
- Edit and create a route that points `0.0.0.0/0` to your Internet gateway ID.
- Save the route table.

5. Upload/download files to/from the VM

After you launch your instance, you can use the SCP command upload and download files to/from your client computer.

- Copy the file "foobar.txt" from the local host to a remote host

```
scp -i $HOME/.ssh/CS205-key.pem foobar.txt ubuntu@34.230.37.25:/home/ubuntu
```

- Copy the file "foobar.txt" from a remote host to the local host

```
scp -i $HOME/.ssh/CS205-key.pem ubuntu@34.230.37.25:/home/ubuntu/foobar.txt .
```

- You can also use the tool `sftp`, which is integrated into many graphical tools

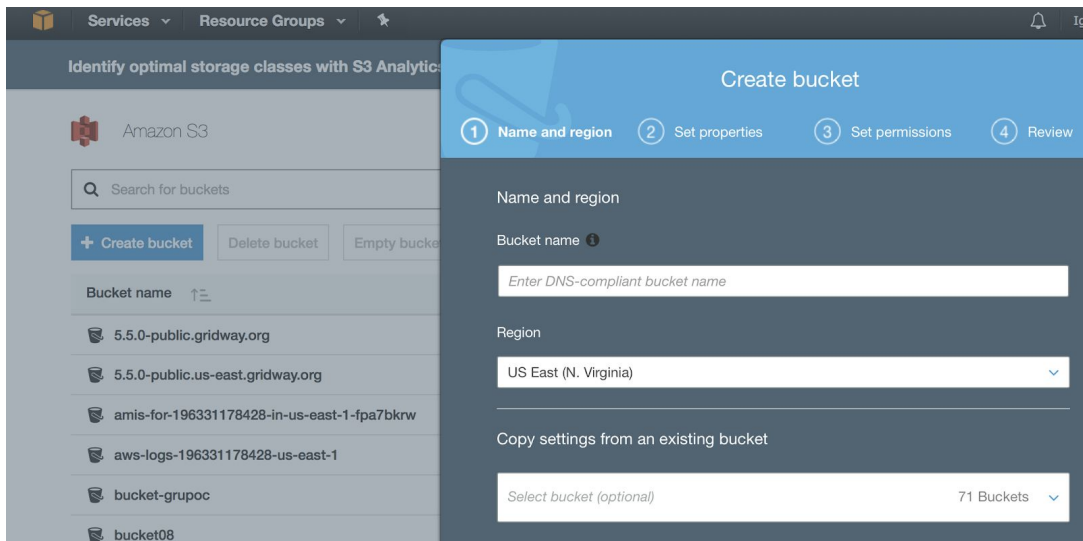
```
sftp -i $HOME/.ssh/CS205-key.pem ubuntu@34.230.37.25
```

6. Bucket and object management in S3

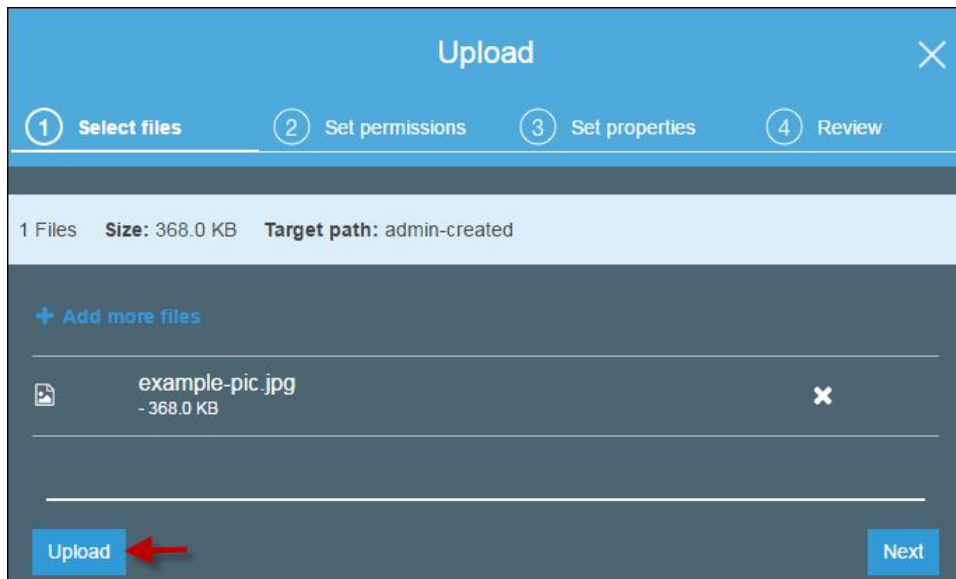
Amazon S3 has a simple web services interface that you can use to store and retrieve any amount of data, at any time, from anywhere on the web.

- Every object in Amazon S3 is stored in a bucket. Before you can store data in Amazon S3, you must create a bucket. Go to the S3 dashboard and click "Create Bucket" in our default region "US East". The name must be unique across all existing bucket names in Amazon S3. After you create the bucket you cannot change the name, so choose wisely. Choose a bucket name that reflects the objects in the bucket because the bucket name is visible in the URL that points to the objects that you're going to put in your bucket.

Proceed through the four steps. You don't need to update the permissions in step 3 at the moment.



- To upload an object to a bucket, choose the name of the bucket that you want to upload your object to, choose Upload, and then choose a file to upload and Open.



- You can then easily download, copy, delete or move the object, and delete the bucket



The screenshot shows the AWS S3 console interface. On the left, the 'Objects' tab is active, displaying a list of objects. The object 'example-pic.jpg' is selected, indicated by a red arrow pointing to its checkbox. The main panel shows the details for 'example-pic.jpg', including its key, size (368.0 KB), expiration date, ETag, last write time, and link. The 'Properties' section shows the storage class as 'Standard' and encryption as 'None'. The 'Permissions' section shows the owner and the number of grantees for read and write permissions. A red arrow points to the 'Download' button at the bottom right of the details panel.

Object	Key	Size	Expiration date	Expiration rule	ETag	Last write	Link	Tags
example-pic.jpg	example-pic.jpg	368.0 KB	N/A	N/A	a6a880d5a3feb1b7780bdfaba9b057a2	Mar 11, 2017 5:40:08 PM GMT-0800	https://s3-us-west-2.amazonaws.com/admin-created/example-pic.jpg	0 Tags

Properties	Storage class	Encryption
	Standard	None

Permissions	Owner	Object	Read	Write
			1 Grantees	1 Grantees

Stop your instances when are done for the day to avoid incurring charges
Terminate them when you are sure you are done with your instance

