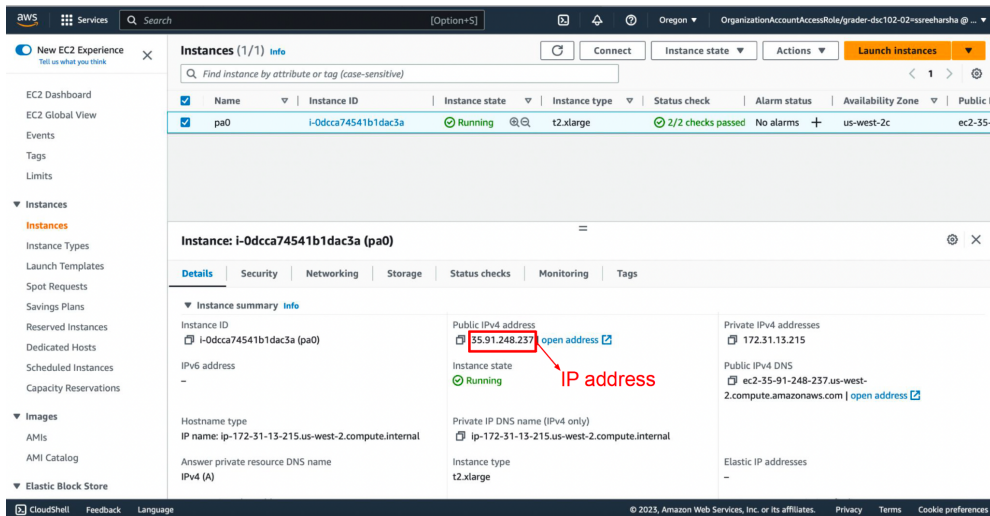


## DSC 102 PA0 System Setup Tutorial

### Step 1 : Setup Client i.e Jupyter Notebook and Port Forwarding for Jupyter Notebook onto localhost.

a) After creating your EC2 instance note down its IP address as shown below



b) Open a Terminal Window and do the following:

i) **Change permission of key file**

```
chmod 400 dask-key.pem
```

ii) **SSH Into the Scheduler EC2 Instance:**

```
ssh -i dask-key.pem ubuntu@35.91.248.237
```

iii) **Activate the Dask Environment:**

```
source dask_env/bin/activate
```

```
saisreeharsha@Sais-MacBook-Air-2 ~ % chmod 400 Downloads/dask-key.pem
saisreeharsha@Sais-MacBook-Air-2 ~ % ssh -i Downloads/dask-key.pem ubuntu@35.91.248.237
Welcome to Ubuntu 22.04.2 LTS (GNU/Linux 5.15.0-1031-aws x86_64)

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

System information as of Wed Apr 19 12:47:24 UTC 2023

System load:  0.16650390625   Processes:            128
Usage of /:   8.9% of 38.58GB Users logged in:          0
Memory usage: 2%             IPv4 address for eth0: 172.31.13.215
Swap usage:  0%

 * Introducing Expanded Security Maintenance for Applications.
   Receive updates to over 25,000 software packages with your
   Ubuntu Pro subscription. Free for personal use.

   https://ubuntu.com/aws/pro

Expanded Security Maintenance for Applications is not enabled.

14 updates can be applied immediately.
To see these additional updates run: apt list --upgradable

7 additional security updates can be applied with ESM Apps.
Learn more about enabling ESM Apps service at https://ubuntu.com/esm

*** System restart required ***
Last login: Sun Apr 16 19:22:00 2023 from 24.43.123.72
ubuntu@ip-172-31-13-215:~$ source dask_env/bin/activate
(dask_env) ubuntu@ip-172-31-13-215:~$ jupyter notebook --port=8888
```

## iv) Launch Jupyter Notebook on the EC2:

```
jupyter notebook --port=8888
```

```
(dask_env) ubuntu@ip-172-31-13-215:~$ jupyter notebook --port=8888
[I 12:49:50.842 NotebookApp] Writing notebook server cookie secret to /home/ubuntu/.local/share/jupyter/runtime/notebook_cookie_secret

      |
      |      .--.
      |     /  \
      |    /    \
      |   /      \
      |  /        \
      | /          \
      |/_          _\
      |  \        /
      |   \      /
      |    \    /
      |     \  /
      |      .--.
      |
      |

Read the migration plan to Notebook 7 to learn about the new features and the actions to take if you are using extensions.

https://jupyter-notebook.readthedocs.io/en/latest/migrate_to_notebook7.html

Please note that updating to Notebook 7 might break some of your extensions.

[I 12:49:54.030 NotebookApp] Serving notebooks from local directory: /home/ubuntu
[I 12:49:54.030 NotebookApp] Jupyter Notebook 6.5.4 is running at:
[I 12:49:54.030 NotebookApp] http://localhost:8888/?token=ace5c6b61bf24461ee067412cea642df40a5809e38cbe756
[I 12:49:54.030 NotebookApp] or http://127.0.0.1:8888/?token=ace5c6b61bf24461ee067412cea642df40a5809e38cbe756
[I 12:49:54.030 NotebookApp] Use Control-C to stop this server and shut down all kernels (twice to skip confirmation).
[W 12:49:54.035 NotebookApp] No web browser found: could not locate runnable browser.
[C 12:49:54.035 NotebookApp]

To access the notebook, open this file in a browser:
    file:///home/ubuntu/.local/share/jupyter/runtime/nbserver-7696-open.html
Or copy and paste one of these URLs:
    http://localhost:8888/?token=ace5c6b61bf24461ee067412cea642df40a5809e38cbe756
    or http://127.0.0.1:8888/?token=ace5c6b61bf24461ee067412cea642df40a5809e38cbe756
```

## v) Copy the link to the Jupyter Server

(shown in the last line of the above screenshot)

You can paste it in the browser AFTER performing step (c) below

c) Open **New Terminal Window** and run the following command:

## i) Port Forwarding Jupyter Notebook running on port 8888 on the EC2 to port 8888 on local system:

```
ssh -i dask-key.pem ubuntu@35.91.248.237 -L 8888:localhost:8888
```

```
saisreeharsha@Sais-MacBook-Air-2 ~ % ssh -i Downloads/dask-key.pem ubuntu@35.91.248.237 -L 8888:localhost:8888
Welcome to Ubuntu 22.04.2 LTS (GNU/Linux 5.15.0-1031-aws x86_64)

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

System information as of Wed Apr 19 13:02:06 UTC 2023

System load:  0.0               Processes:    124
Usage of /:   8.9% of 38.58GB    Users logged in: 1
Memory usage: 2%               IPv4 address for eth0: 172.31.13.215
Swap usage:   0%

 * Introducing Expanded Security Maintenance for Applications.
   Receive updates to over 25,000 software packages with your
   Ubuntu Pro subscription. Free for personal use.

   https://ubuntu.com/aws/pro

Expanded Security Maintenance for Applications is not enabled.

14 updates can be applied immediately.
To see these additional updates run: apt list --upgradable

7 additional security updates can be applied with ESM Apps.
Learn more about enabling ESM Apps service at https://ubuntu.com/esm

*** System restart required ***
Last login: Wed Apr 19 12:47:26 2023 from 24.43.123.72
ubuntu@ip-172-31-13-215:~$
```

### Step 2 : Dask UI Port forwarding

a) Open a **New Terminal Window** and run the following command:

i) **Port Forwarding the Dask dashboard UI running on port 8787 on the EC2 to port 8001 on local system:**

```
ssh -i dask-key.pem ubuntu@35.91.248.237 -L 8001:localhost:8787
```

```
saisreeharsha@Sais-MacBook-Air-2 ~ % ssh -i Downloads/dask-key.pem ubuntu@35.91.248.237 -L 8001:localhost:8787
Welcome to Ubuntu 22.04.2 LTS (GNU/Linux 5.15.0-1031-aws x86_64)

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

System information as of Wed Apr 19 13:13:16 UTC 2023

System load:  0.0          Processes:      127
Usage of /:   8.9% of 38.58GB Users logged in: 1
Memory usage: 2%          IPv4 address for eth0: 172.31.13.215
Swap usage:  0%

 * Introducing Expanded Security Maintenance for Applications.
   Receive updates to over 25,000 software packages with your
   Ubuntu Pro subscription. Free for personal use.

   https://ubuntu.com/aws/pro

Expanded Security Maintenance for Applications is not enabled.

14 updates can be applied immediately.
To see these additional updates run: apt list --upgradable

7 additional security updates can be applied with ESM Apps.
Learn more about enabling ESM Apps service at https://ubuntu.com/esm

*** System restart required ***
Last login: Wed Apr 19 13:02:07 2023 from 24.43.123.72
ubuntu@ip-172-31-13-215:~$
```

### Step 3 : Download data from S3

a) In the most recently opened Terminal Window:

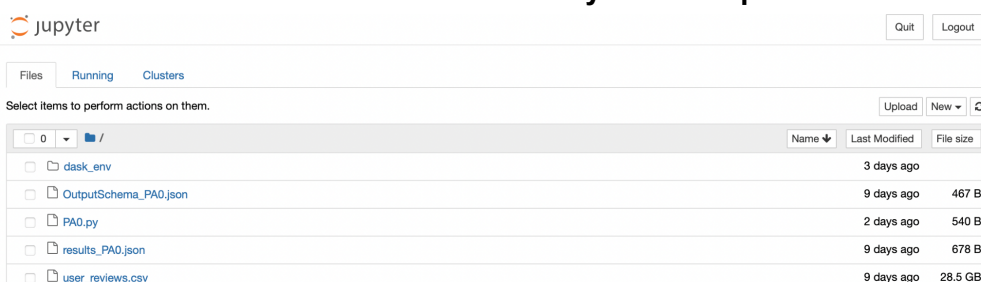
i) **Copy and paste the AWS ACCESS KEY ID, AWS SECRET ACCESS KEY, and AWS SESSION TOKEN**

ii) **Download all the files from the S3 :**

```
aws s3 sync s3://dsc102-public /home/ubuntu/
```

```
ubuntu@ip-172-31-13-215:~$ aws s3 sync s3://dsc102-public /home/ubuntu/
download: s3://dsc102-public/PA0.py to ./PA0.py
download: s3://dsc102-public/OutputSchema_PA0.json to ./OutputSchema_PA0.json
download: s3://dsc102-public/results_PA0.json to ./results_PA0.json
download: s3://dsc102-public/user_reviews.csv to ./user_reviews.csv
ubuntu@ip-172-31-13-215:~$
```

Now, on navigating to the link copied at the end of Step 1 b), you should see the following. You can now create a new notebook and are ready to code up.



The image shows the JupyterLab interface. At the top, there's a 'jupyter' logo and 'Quit' and 'Logout' buttons. Below that, there are tabs for 'Files', 'Running', and 'Clusters'. The 'Files' tab is active, showing a file browser. At the top of the file browser, it says 'Select items to perform actions on them.' and has 'Upload', 'New', and a refresh icon. Below this is a table of files:

	Name	Last Modified	File size
<input type="checkbox"/>	/		
<input type="checkbox"/>	dask_env	3 days ago	
<input type="checkbox"/>	OutputSchema_PA0.json	9 days ago	467 B
<input type="checkbox"/>	PA0.py	2 days ago	540 B
<input type="checkbox"/>	results_PA0.json	9 days ago	678 B
<input type="checkbox"/>	user_reviews.csv	9 days ago	28.5 GB