

# Preparing Your Google Cloud VM for W4187

September 20, 2016

## 1 Get a `cloud.cs.columbia.edu` account and a coupon code

1. Sign up for a **cloud** Columbia CS account using this link: <https://www.cs.columbia.edu/auth/cloud>. Note that is is an entirely new account and is **not** your normal Columbia account and is **not** your CRF account. Your **cloud** Columbia CS account will be in the form: `yourUNI@cloud.cs.columbia.edu`.
2. After you create your account, you should get an email with a temporary password to log into your **cloud** account. The email will contain a link to a guide provided by CRF. The CRF guide has some overlap with the guide you are reading right now, but we provide more detail and W4187-specific setup information.
3. Check your email for a coupon code sent by the teaching staff via Canvas (Courseworks) to you. You will use this coupon code in the next step on Google Cloud Platform (GCP).
4. **IMPORTANT STEP:** Make sure you are **not** logged into LionMail! **Log out of LionMail.** Better yet, do all these steps in incognito/private browsing mode while ensuring that you are **not** logged into LionMail.

## 2 Enter your coupon code on Google Cloud Platform (GCP)

1. **Log out of your LionMail account or use a new private browsing window.**
2. Make sure you are logged out of LionMail.
3. Ensure that you are **not** logged into your LionMail account or any other Google account. **If you try to use the coupon code on an account that is not your cloud account, it will not work.** Then login using your `cloud.cs.columbia.edu` account at <https://console.cloud.google.com/>.
4. Go to <https://console.cloud.google.com/education>.
5. **Accept** any kind of agreement confirmation that Google asks you.
6. On the “Education grants” screen, enter the coupon code you obtained in Section 1.

Google Cloud Platform Project

### Education grants

Please enter the coupon code provided to you via the Google Cloud Platform Education Grants program to receive credit for Google Cloud Platform. Get what you need to build and run your apps, websites and services.

Coupon code

Please email me updates regarding feature announcements, performance suggestions, feedback surveys and special offers.

Yes  No

Google Cloud Platform educational grants credits terms and conditions

You, on behalf of yourself and the organization you represent (“You”) agree to these terms and conditions:

The credit is valid for Google Cloud Platform products and is subject to Your acceptance of the applicable Google Cloud Platform License Agreement and any other applicable terms of service. The credit is non-transferable and may not be sold or bartered. Unused credit expires on the date indicated on the media conveying the promotion code. The credit may be issued in increments as You use the credit over the period of time during which the credit is valid. Offer void where prohibited by law.

[Additional terms for employees of public educational institutions](#)

Accept and continue Clear

7. Select **No** for “Please email me updates ...” unless you actually want email updates on new offerings.
8. Click on **Accept and continue**.
9. You will be brought to the GCP web console where you can create your first project. Click on **Project** on the top bar and then on **Create project**. Feel free to name your project whatever you want. In **Advanced options**, change the region to **us-east1**. **Ensure that your project is selected**.
10. Now, we will setup a key-pair that identifies you as you so that you can log into any virtual machine instances you start up.

### 3 Generate a private/public key-pair

To gain access to the VMs you create in your GCP project, you will need to provide to GCP your “public key” that is paired with your “private key”. Your private key, as the name implies, should be kept secret. On the other hand, your public key is broadcast publicly, and the VMs you create on GCP will use your public key to authenticate you, i.e. the holder of the private key corresponding to the known public key.

**What this means for you is that you must copy your private key to every computer that you want to use to access your Google VM.** This makes it difficult to use public computers to do your assignments but provides security. In the worst case, you can also use GCP’s browser SSH client which uses a key-pair managed internally at Google and not the key-pair you generate here.

A full guide on how to generate key-pairs can be found on the documentation for GCP at <https://cloud.google.com/compute/docs/instances/adding-removing-ssh-keys#createsshkeys>. How-

ever, not all the steps in this guide are necessary, so we reproduce below the bare minimum required steps for generating an SSH key-pair.

### 3.1 Mac OS X and Linux

1. Open a terminal window and type the following two commands:

```
ssh-keygen -t rsa -f ~/.ssh/my_w4187_key -C [MY_UNI]
chmod 400 ~/.ssh/my_w4187_key
```

The `ssh-keygen` command will create a **private key** file at `~/.ssh/my_w4187_key` and a corresponding **public key** file at `~/.ssh/my_w4187_key.pub`.

The `chmod` command restricts the permissions on the private key—which must be kept secret—on your local machine so that only your user (and/or root) can view it.

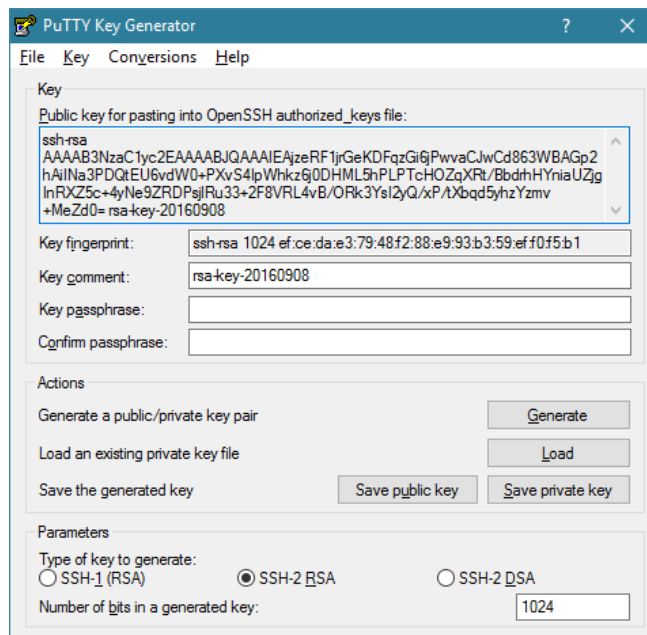
2. For convenience, add the private key to your ssh-agent:

```
ssh-add ~/.ssh/my_w4187_key
```

3. Proceed to Section 4.

### 3.2 Windows

1. Download PuTTYgen (puttygen.exe) from <http://www.chiark.greenend.org.uk/~sgtatham/putty/download.html>.
2. Run PuTTYgen. The default settings should be fine for the purposes of this class, but it is recommended to increase the number of bits from 1024 to 2048.
3. Click the Generate button. It will ask you to move your mouse around randomly (keep moving your mouse until it completes).



- When it completes, change the “Key comment” field to your UNI. This will be your login username for your VM on GCP.
- PuTTYgen will also display your public key under the header text “Public key for pasting into OpenSSH authorized.keys file”. **Copy all the text in the text box that begins with ssh-rsa, and paste it into your favorite text editor. Save this as a file with a .pub extension, such as my\_w4187\_key.pub.**<sup>1</sup> The resulting file, which is your **public key**, should look something like:

```
ssh-rsa AAAAB3NzaC1yc2EAAAABJQAAAIBwYtSgLr9GU6dHC/y/
wKEMV1x6Me1MVJ7rsHOJOHdYqKXYWLoQkYsdfVkztfj+
t1FLBI sBREpT8tX21nPI2v+ER5I6XMeQWE0nQLwBbBnxtceLd65h4na5p+
rv5DX02d3H0hQDmjM5Bg30yEHMjZw14ERX+GX50SZ0ry1go6w8w== rsa-key
-20160908
```

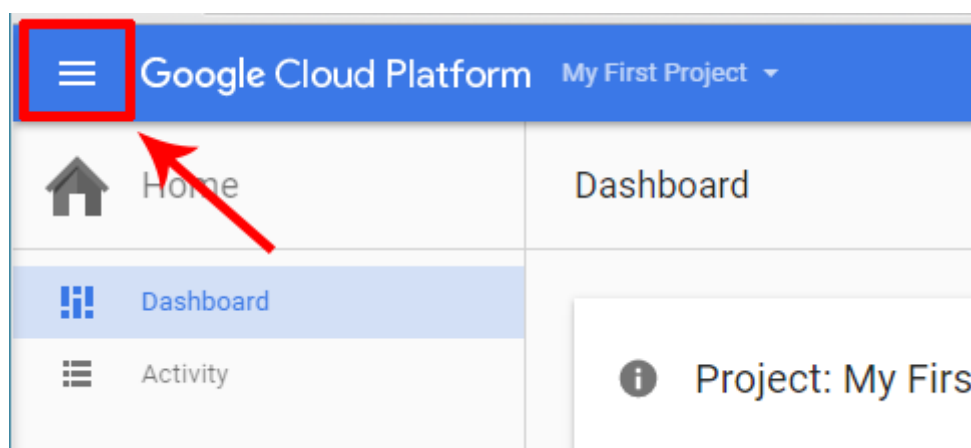
Note that the very last token of the string, the `rsa-key-20160908` part, will correspond to your login username for your Google VM. **Thus, to change your VM login username on GCP, you should change this final token.**

- Click **Save public key** to write your public key to a file with a .pub extension. **This saves the public key in a different format from what you saved in the previous step, so make sure you save *both*.**
- Click **Save private key** to write your private key to a file with a .ppk extension.
- Proceed to Section 4.

## 4 Add your public key to your GCP project

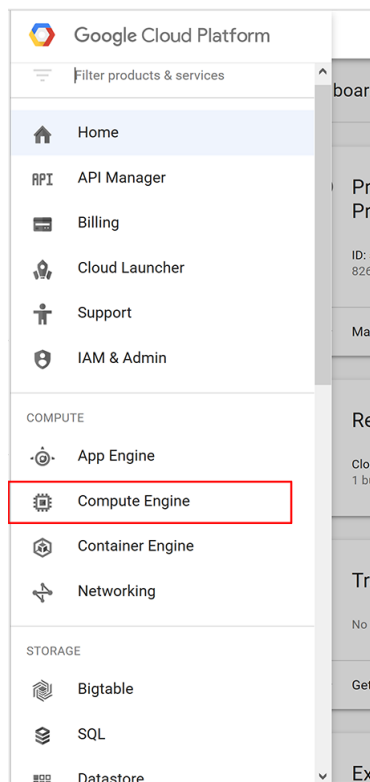
Now that you have your private/public key-pair, you need to add the **public** key to GCP. This will allow GCP VMs to authenticate anyone who has your **private key**.

- Navigate to your dashboard at <https://console.cloud.google.com/> and click on the hamburger icon on the top left as shown below.

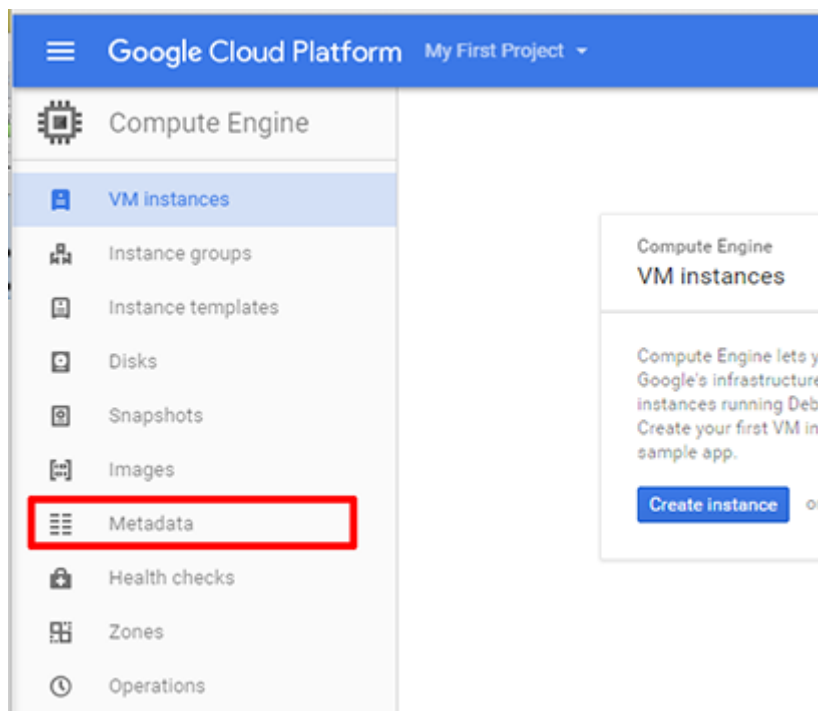


<sup>1</sup>The standard naming convention for a public/private key-pair is `id_rsa.pub` and `id_rsa`. However, in reality you can name your keys whatever you want. Here, we want to avoid having you accidentally overwrite any existing key-pairs you may have, and thus instruct you to use a custom filename.

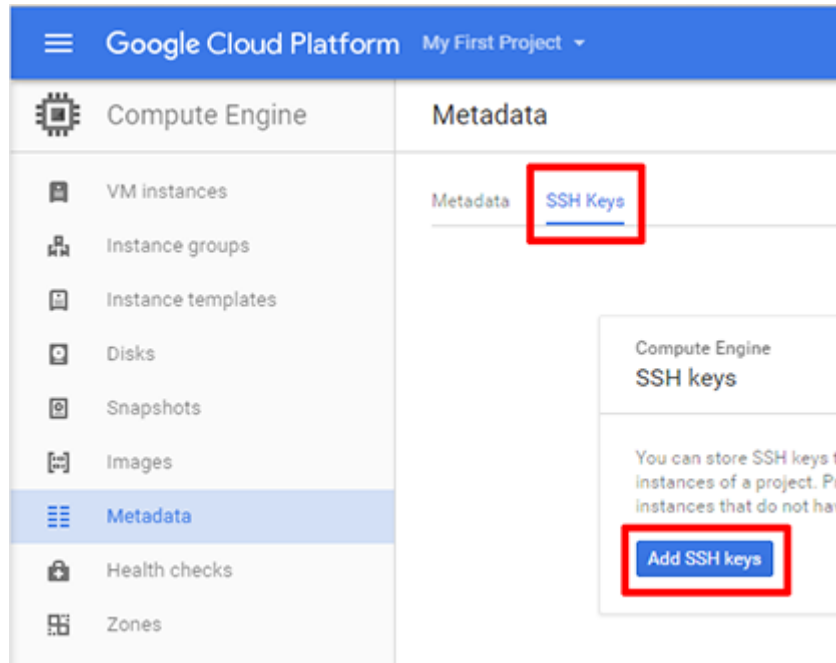
2. In the menu that expands, click on **Compute Engine** as depicted below.



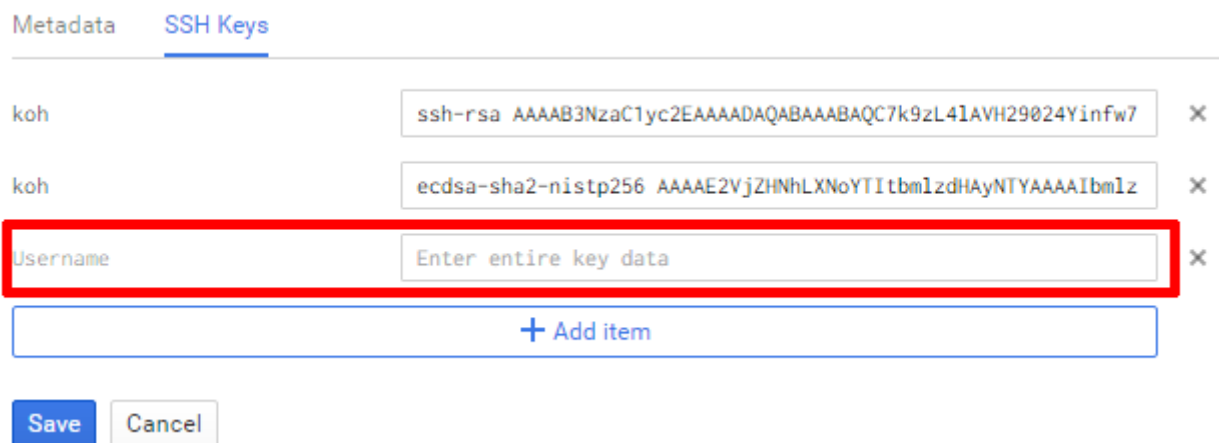
3. You will now be on the Compute Engine portion of GCP which has its own menu on the left side. Click on **Metadata** on the menu.



4. Click on **SSH Keys** in the “Metadata” content area and then click on **Add SSH keys**.



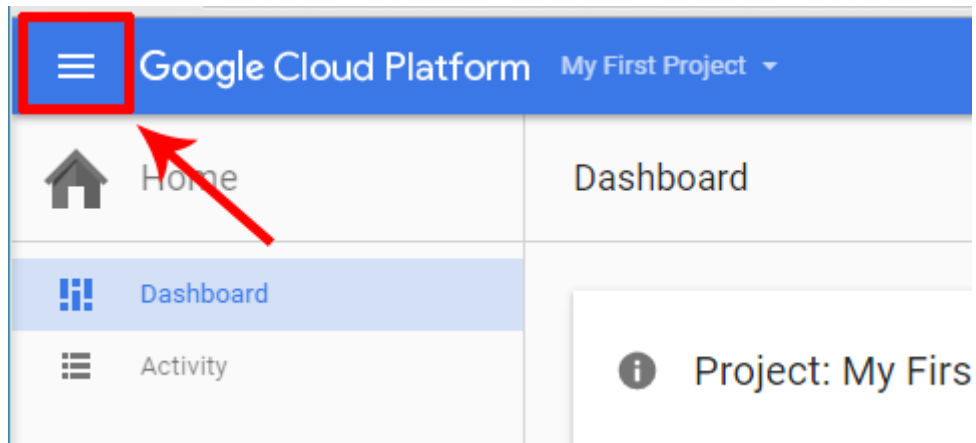
5. Copy and paste the entire contents of the **public key** (`my_w4187_key.pub`) you generated in Section 3 into the **“Enter entire key data”** field. (You can open the public key in a text editor to view its contents.) The key you paste should begin with the text `ssh-rsa`.
6. When you paste your public key, it automatically fills in the “Username” field. The “Username” field is the login username associated with the public key. **You can change the “Username” field to your UNI by modifying the very last part of the public key that you pasted.**
7. Click **Save**.



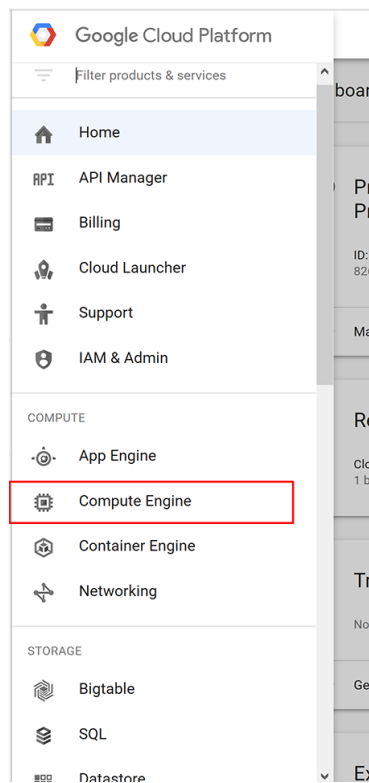
## 5 Create your Google virtual machine instance

Now you will configure and create a virtual machine instance on GCP'S Compute Engine. This virtual machine is where you should test your programming assignments. **You must exactly follow these instructions. Otherwise, your environment may differ from the graders' VMs and cause issues with running your assignment submissions.**

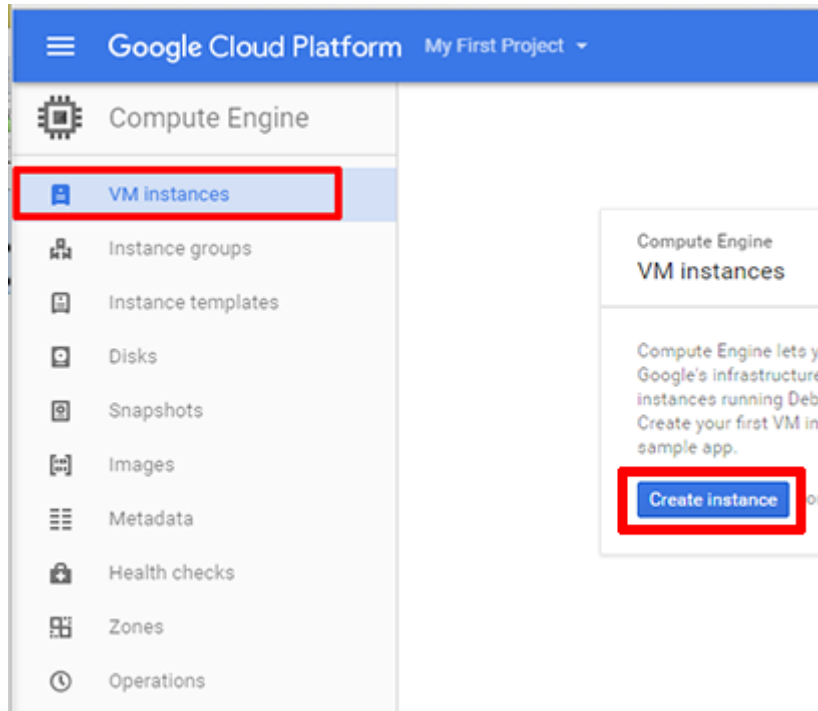
1. Navigate to your dashboard at <https://console.cloud.google.com/> and click on the hamburger icon on the top left as shown below.



2. In the menu that expands, click on **Compute Engine** as depicted below.



3. In the Compute Engine section, click on **VM instances** on the left menu. At this point, you may need to click on **Enable billing**, which will then ask you to choose a project to enable billing on. Select the appropriate project, and then click on **Create instance**.



4. Use the settings shown in the picture below. The settings are:

- Zone: us-east1-d
- Machine type: small (1 shared vCPU)
- Boot disk: Ubuntu 14.04 LTS

Then click on **Create**.



← Create an instance

Name <sup>?</sup>  
instance-1

Zone <sup>?</sup>  
us-east1-d

Machine type  
small (1 shared vCPU) 1.7 GB memory [Customize](#)

Boot disk <sup>?</sup>  
New 10 GB standard persistent disk  
Image  
Ubuntu 14.04 LTS [Change](#)

Identity and API access <sup>?</sup>  
Service account <sup>?</sup>  
Compute Engine default service account

Access scopes <sup>?</sup>  
 Allow default access  
 Allow full access to all Cloud APIs  
 Set access for each API

Firewall <sup>?</sup>  
Add tags and firewall rules to allow specific network traffic from the Internet  
 Allow HTTP traffic  
 Allow HTTPS traffic  
[Management, disk, networking, SSH keys](#)

You will be billed for this instance. [Learn more](#)

[Create](#) [Cancel](#)

Equivalent [REST](#) or [command line](#)

5. You will be brought to the “VM instances” console which shows various statistics and information about your instances. It will take some time for your newly created instance to boot up. When

it is booted up, there will be a green checkmark icon and an “External IP” visible.

## 6 Accessing your virtual machine

You can access your GCP VM via SSH similar to how you access a CLIC machine via SSH. Some important points:

- The username you use **must** be the “Username” that you entered for your SSH key in Section 4.
- The IP address you connect to is the “External IP” for your VM instance. This IP will change every time you restart your instance, so take note.
- You must use the private key that you created in Section 4.
- Example: On Linux/Mac OS X/Windows 10 with Ubuntu (Windows’ Bash), I would use the following command to SSH into the VM instance that was setup in this guide:

```
ssh rsa-key-20160908@104.196.107.142
```

Note that `rsa-key-20160908` is the username that GCP parsed out from the public key that was pasted into their console, and that `104.196.107.142` is the external IP address of the VM instance that was launched.

For generic SSH client details, see GCP’s guide at <https://cloud.google.com/compute/docs/instances/connecting-to-instance#standardssh>.

For quick access to look around, you can also use GCP’s **browser SSH client** which you launch simply by clicking on the **SSH** button on the “VM instances” console.

The screenshot shows the Google Cloud Platform VM instances console. At the top, there are buttons for "CREATE INSTANCE", "CREATE INSTANCE GROUP", "REFRESH", and a trash icon. Below that is a search bar "Filter by label or name" and buttons for "Columns", "Labels", and "Recommendations". A "CPU utilization" dropdown is set to "1 hour". A line graph shows CPU usage over time, with a peak of 16.24% at 12:45 PM. Below the graph is a table of VM instances. The table has columns: Name, Zone, Machine type, Recommendation, In use by, Internal IP, External IP, and Connect. The first instance, "instance-1", has an external IP of "104.196.107.142" and an SSH button next to it. A red box highlights the "External IP" and "SSH" columns for the first instance.

Name	Zone	Machine type	Recommendation	In use by	Internal IP	External IP	Connect
<input type="checkbox"/> instance-1	us-east1-d	1 vCPU, 1.7 GB			10.142.0.2	104.196.107.142	SSH

## 7 Setting up your development environment

For consistency, we are providing a Bash script for you to run which will install the basic libraries you need for C/C++ development.

1. SSH into your VM instance.
2. Run either these commands:

```
sudo apt-get update
sudo apt-get -y upgrade
sudo apt-get -y install build-essential manpages-dev autoconf automake
sudo apt-get -y install libevent-dev libssl-dev
sudo apt-get -y install rsync
```

or this command:

```
curl -fsS https://www.cs.columbia.edu/~smb/classes/f16/gcp_vm_setup_script.sh | bash
```

Either of these choices installs the basic requirements on your VM for this course. You can verify this by running the following commands and checking that they correctly output version numbers:

```
gcc -v
g++ -v
make -v
```

## 8 Copying files to and from your Google VM

There are many different ways to copy files to and from your VM. Here is a non-exhaustive list of possible methods you can use:

- scp - (Linux, Mac OS X) Secure copy which supports copying over the network.
- SFTP - (Windows, Linux, Mac OS X) SSH FTP which is similar in functionality to FTP but uses SSH. You can use FTP clients such as FileZilla (or WinSCP) and connect to your VM on port 22.
- sshfs - (Linux, Mac OS X) SSH File System which lets you mount an SSH-accessible device as if it were a hard disk. If you use this method, **make sure that you test your code on the VM, not on your local machine.**

You can do your assignments locally on your own computers, but you **must** make sure to test your code on your Google VM to ensure that it works. You should also make sure to save frequently and to make backups of your work.

## 9 Backing up your assignments

You can create backups of your files in many ways. One way is to initialize your assignment as a Git repository and to push your code to private online repositories such as BitBucket.

We are also exposing a central class VM managed by the teaching staff on which you can backup your assignments. The IP address of the central class VM is:

104.196.149.20

After you create your public key in Step 3, if you send the teaching staff your public key, we can add you as a user to the central class VM. **Then, you can SSH into the central class VM or use SCP or Rsync to synchronize your files on your personal VM with your files on the central class VM.** Your SSH/SCP/RSync connection string would be `yourUNI@104.196.149.20`.

Note: We may occasionally need to take down or reboot the central class VM, so keep this in mind.

## 10 Creating a new user on the VM

For some assignments, you will require multiple users on the VMs. To create a new user with password authentication:

1. SSH into your VM.
2. Type the following command:

```
sudo adduser --gecos "" your_new_username
```

Replace `your_new_username` with your desired username. Also, enter your desired password when it prompts **Enter new UNIX password** and **Retype new UNIX password**.

3. Switch to your newly created user by typing:

```
su your_new_username
```

This command will prompt you for the password of the user. (Replace `your_new_username` with your desired username.) If you want, you can also run `sudo su your_new_username` to skip the password prompt.

## 11 Other important information

- Try to shut off your instance when you aren't using it for long periods of time.
- Do **not** delete your instance if you have saved work on it. Deleting your instance will delete all its contents.
- Your files will persist through instance shutdowns—**not instance deletes**—but you should consider backing up your files elsewhere just in case. The cloud is highly resilient and highly available, but failures are still possible.
- Consider using `git` or `svn` to version control your work. If you use `git`, it also becomes easy to push your code to **private** online repositories such as Bitbucket as a backup solution.