

## **CLOUD COMPUTING PRACTICAL WORK**

### ***Objective:***

- To develop web applications in cloud
- Designing, Implementing and managing the issues of Cloud Computing.

### **LIST OF EXPERIMENTS:**

1. Install Virtualbox/VMware Workstation with different flavours of linux or windows OS on top of windows.
2. Install a C compiler in the virtual machine created using virtual box and execute Simple Programs
3. Find a procedure to transfer the files from one virtual machine to another virtual machine.
4. Install Google App Engine. Create hello world app and other simple web applications using python/java.
5. Install Hadoop single node cluster and run simple applications like wordcount.
6. Simulate a cloud scenario using CloudSim and run a scheduling algorithm that is not present in CloudSim.
7. Case-Study

### ***Outcome:***

On completion of this, the students will be able to:

- Configure various virtualization tools such as Virtual Box, VMware workstation.
- Design and deploy a web application in a PaaS environment.
- Learn how to simulate a cloud environment to implement new schedulers.
- Install and use a generic cloud environment that can be used as a private cloud.
- Manipulate large data sets in a parallel environment.

## EX.No:1

# Install Virtualbox/VMware Workstation with different flavours of linux or windows OS on top of windows

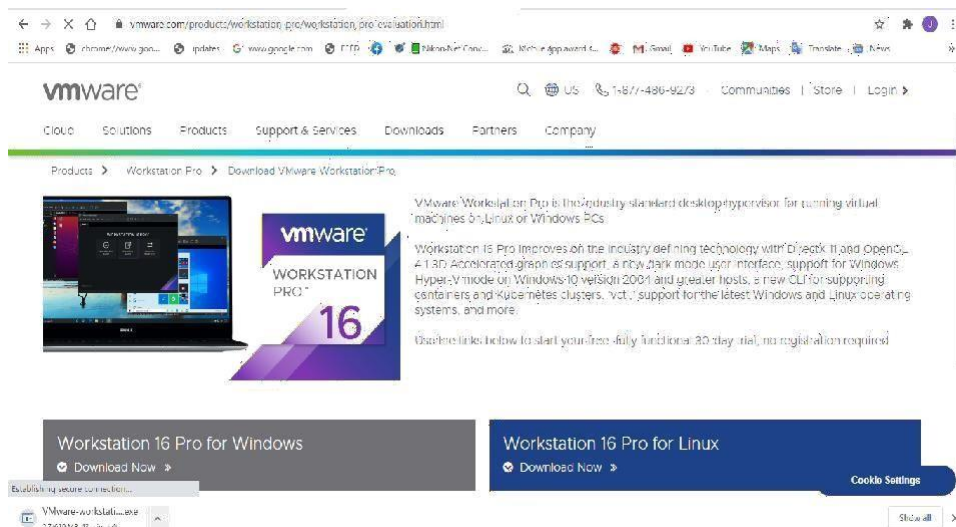
### Aim:

To install Virtualbox/VMware Workstation with different flavours of linux or windows OS on top of windows7 or 8 or 10.

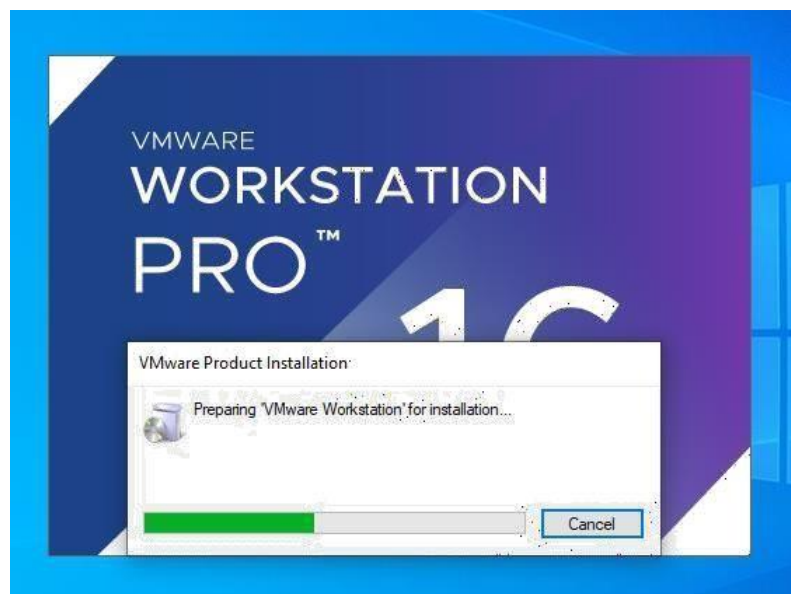
### Procedure:

#### Downloading and installing VMware

##### Step 1: Download VMware



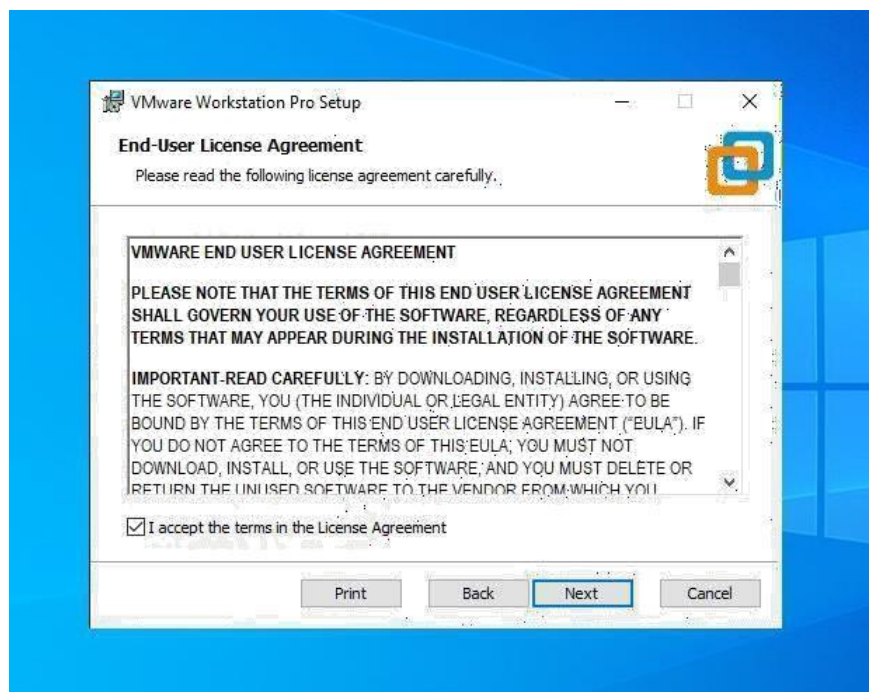
##### Step 2: Install the VMware Application



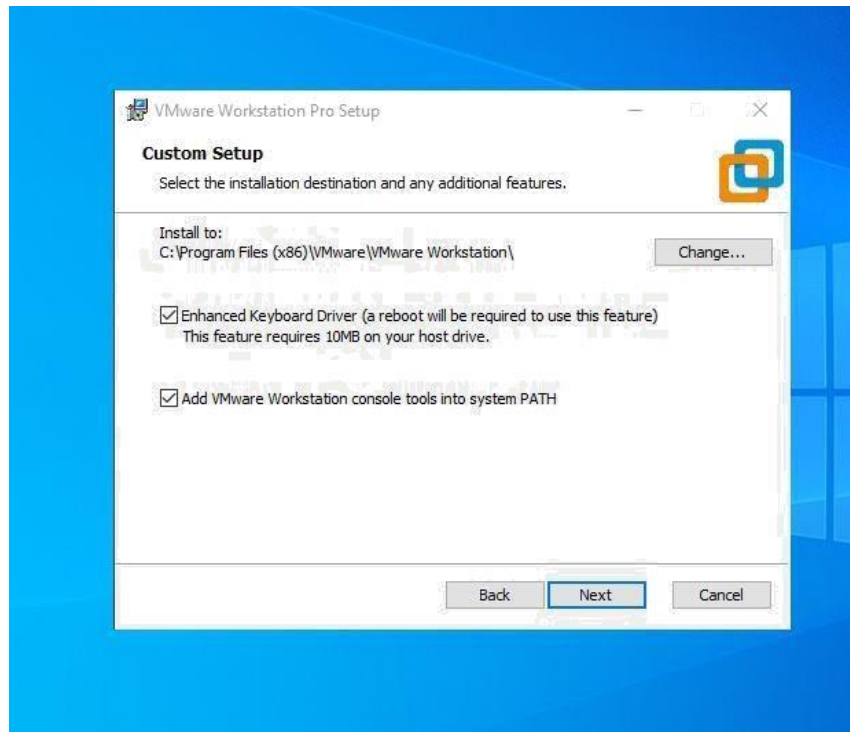
### Step 3: Click Next



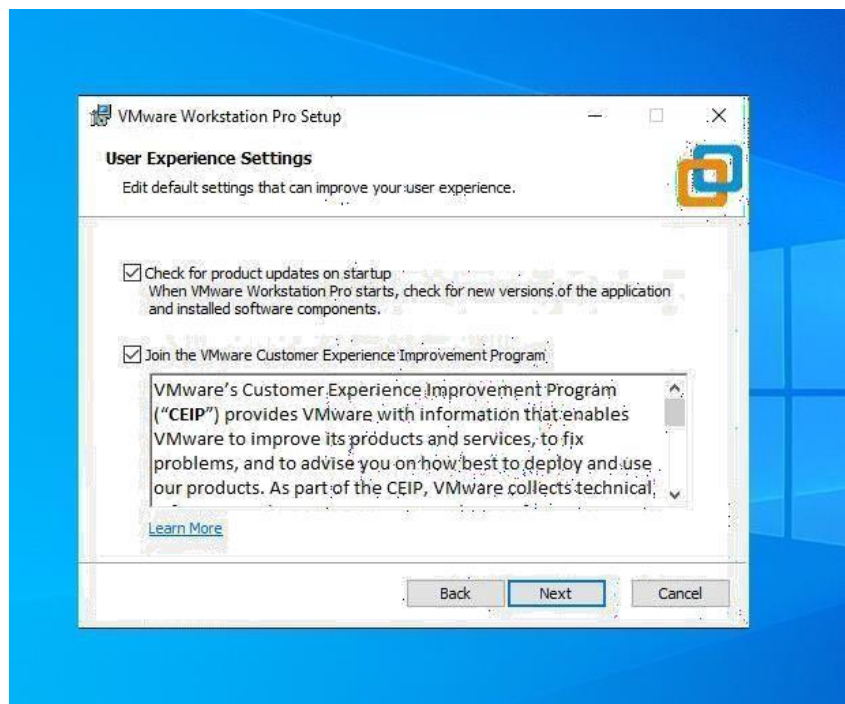
### Step 4: Accept and Click Next




## Step 5: Click Next



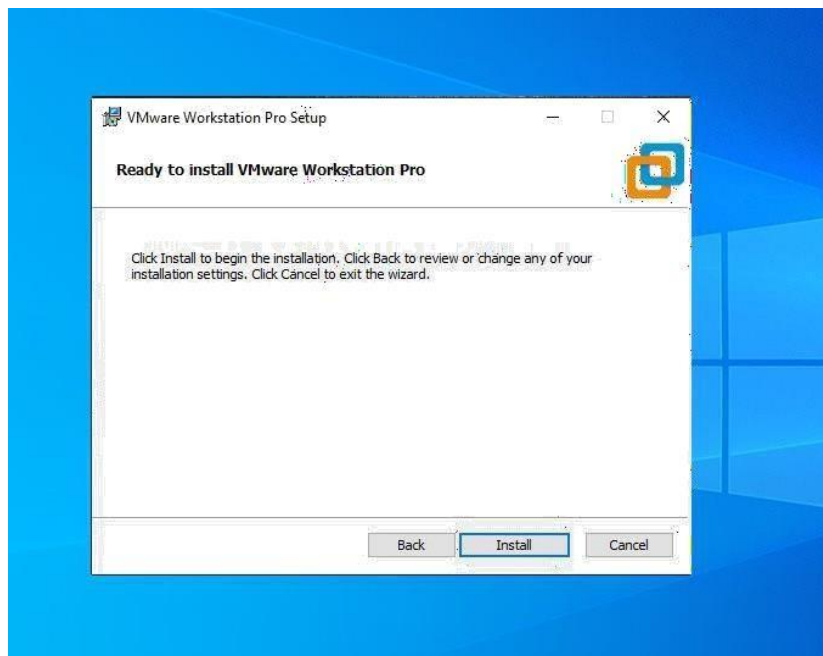
## Step 6: Click Next




## Step 7: Click Next



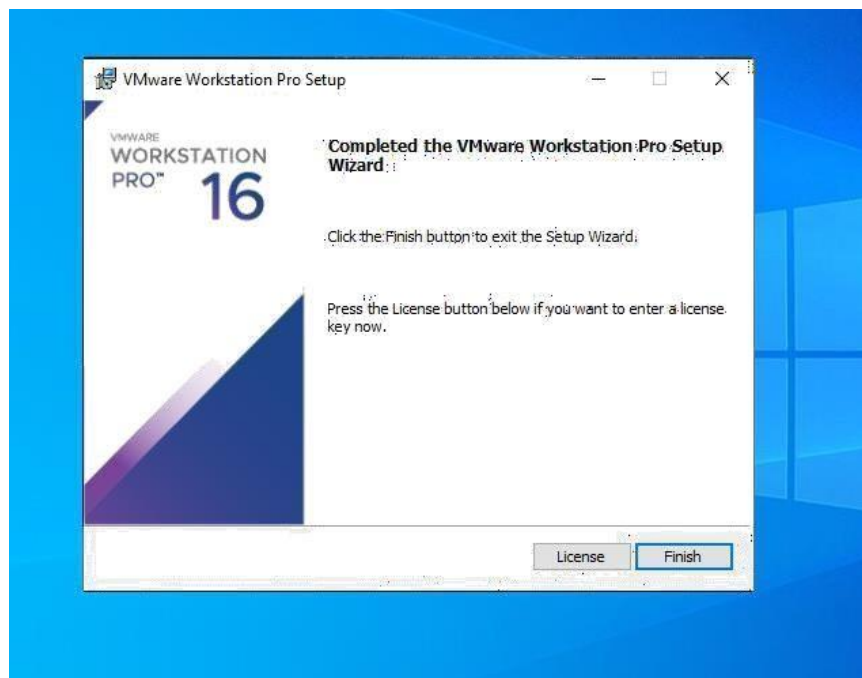
## Step 8: Click Install



## Step 9: Installing

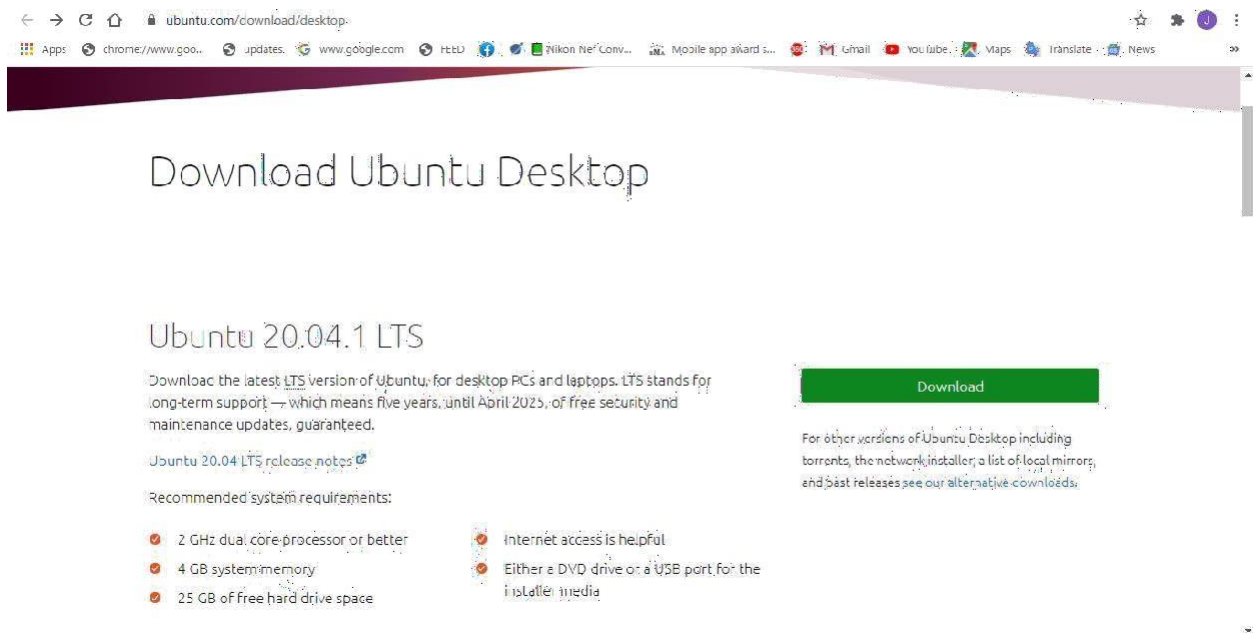


## Step 10: Click Finish



## Downloading Ubuntu

### Step 11: Download the Ubuntu OS



The screenshot shows a web browser window with the URL `ubuntu.com/download/desktop`. The page title is "Download Ubuntu Desktop". The main heading is "Download Ubuntu Desktop". Below this, it says "Ubuntu 20.04.1 LTS". A green "Download" button is visible. The page also includes a list of recommended system requirements and a note about long-term support (LTS).

Download the latest LTS version of Ubuntu, for desktop PCs and laptops. LTS stands for long-term support, which means five years, until April 2025, of free security and maintenance updates, guaranteed.

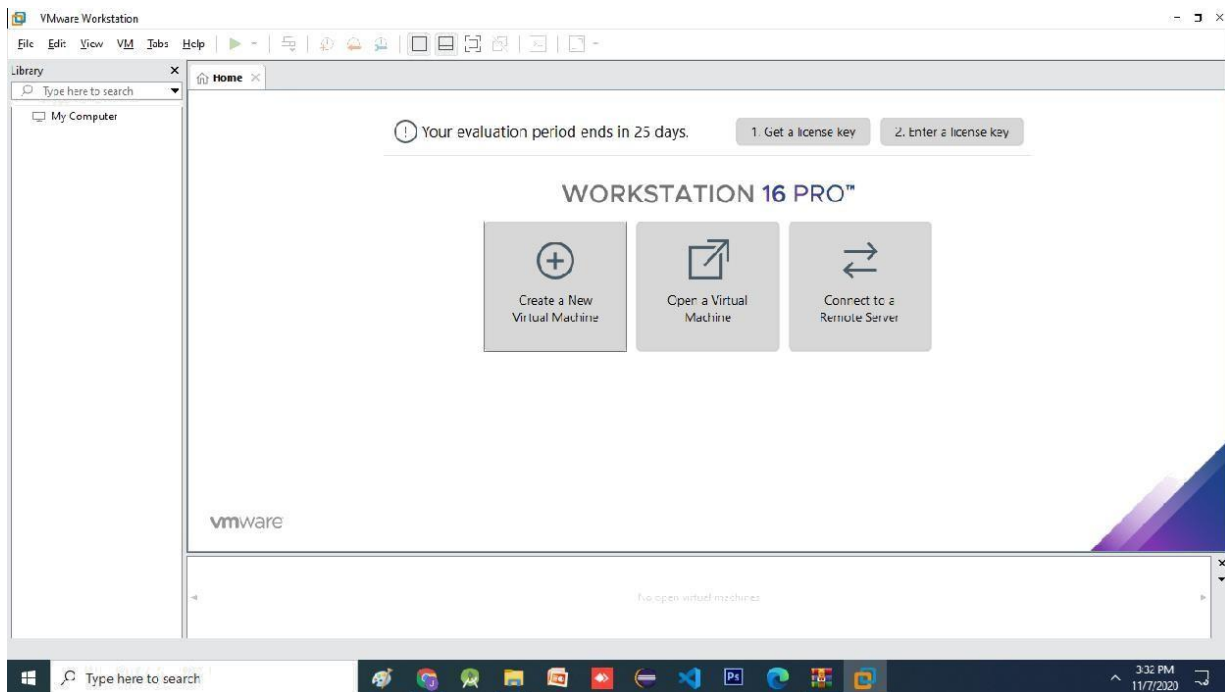
[Ubuntu 20.04 LTS release notes](#)

Recommended system requirements:

- 2 GHz dual-core processor or better
- 4 GB system memory
- 25 GB of free hard drive space
- Internet access is helpful
- Either a DVD drive or a USB port for the installer media

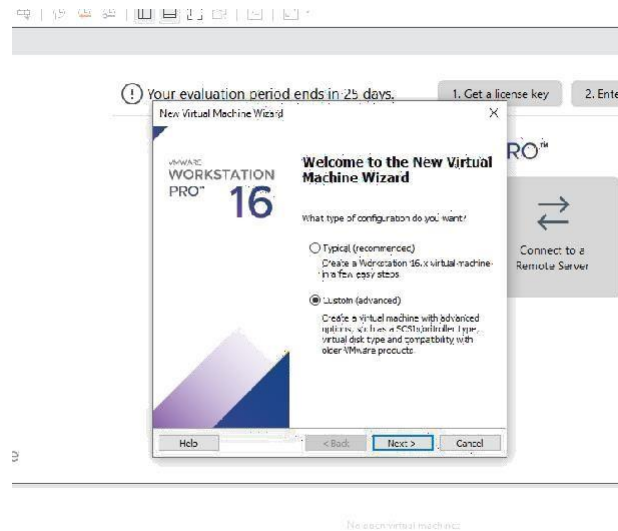
For other versions of Ubuntu Desktop including torrents, the network installer, a list of local mirrors, and past releases see our [alternative downloads](#).

### Step 12: Create new virtual machine

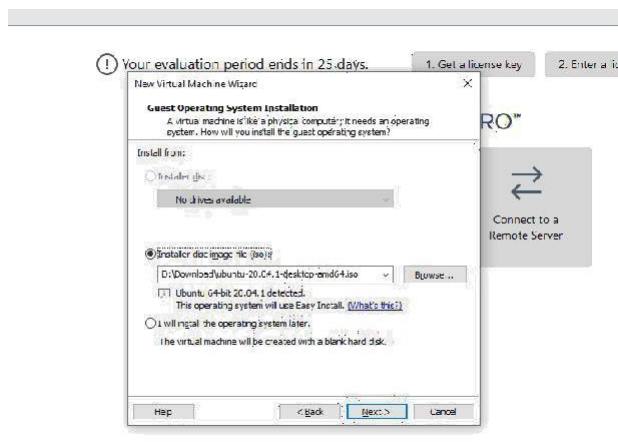


The screenshot shows the VMware Workstation 16 Pro interface. The main window displays a "Home" screen with a "Your evaluation period ends in 25 days" notification. Below the notification, there are three buttons: "Create a New Virtual Machine", "Open a Virtual Machine", and "Connect to a Remote Server". The VMware logo is visible in the bottom left corner of the main window. The Windows taskbar is visible at the bottom of the screen, showing the time as 3:32 PM on 11/7/2022.

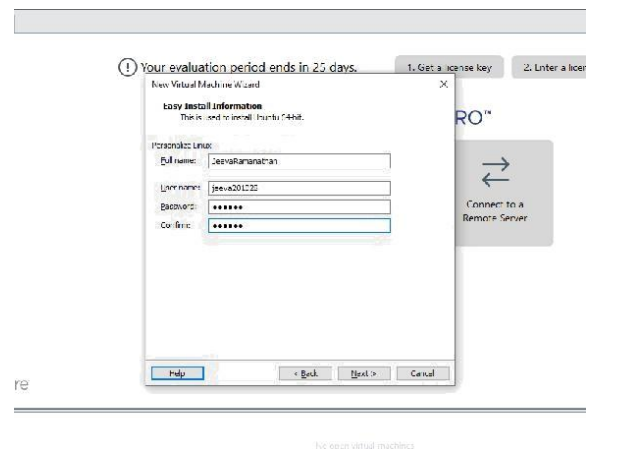
### Step 13: Click Next



### Step 14: Browse the downloaded Ubuntu file and click next

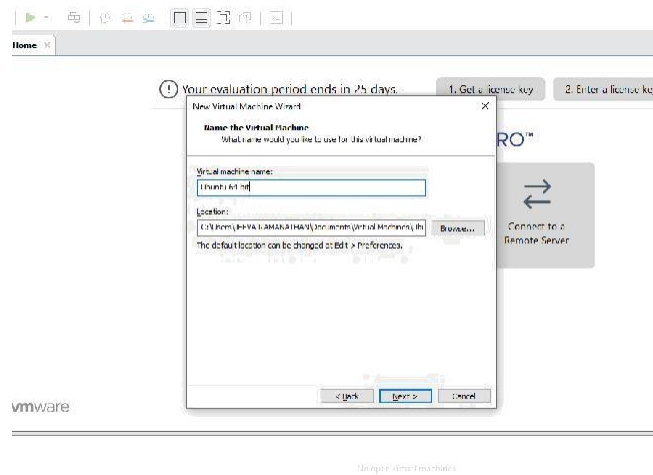


### Step 15: Create an username and password and click next

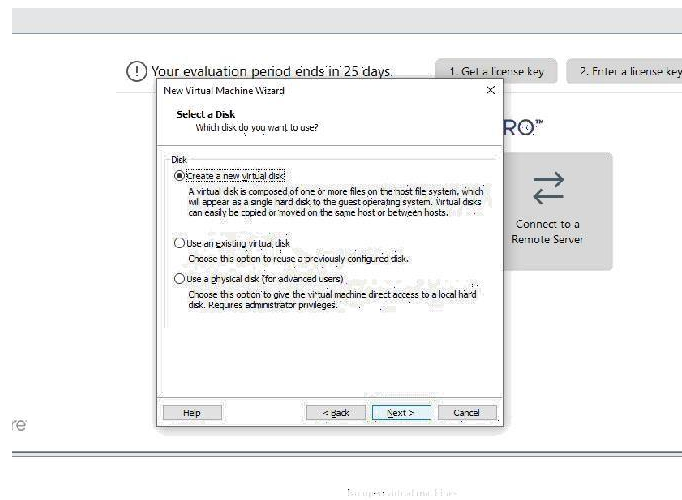




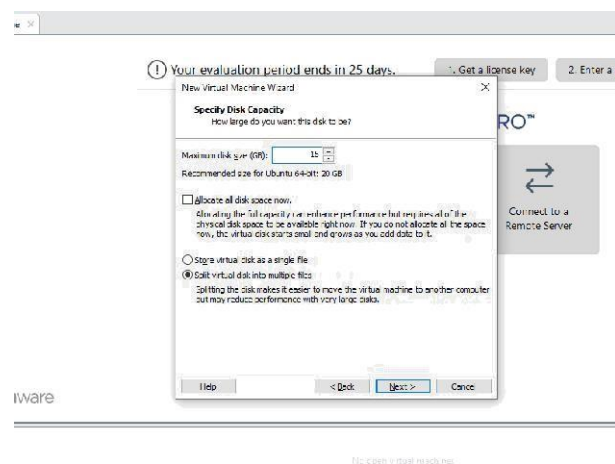
## Step 16: Choose the location to use your virtual machine and click next



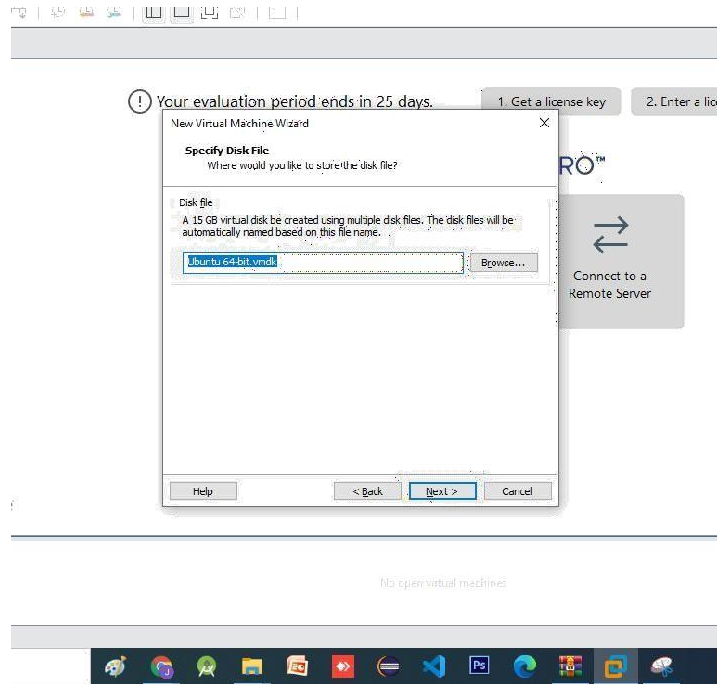
## Step 17: Select create a new virtual disk and click next



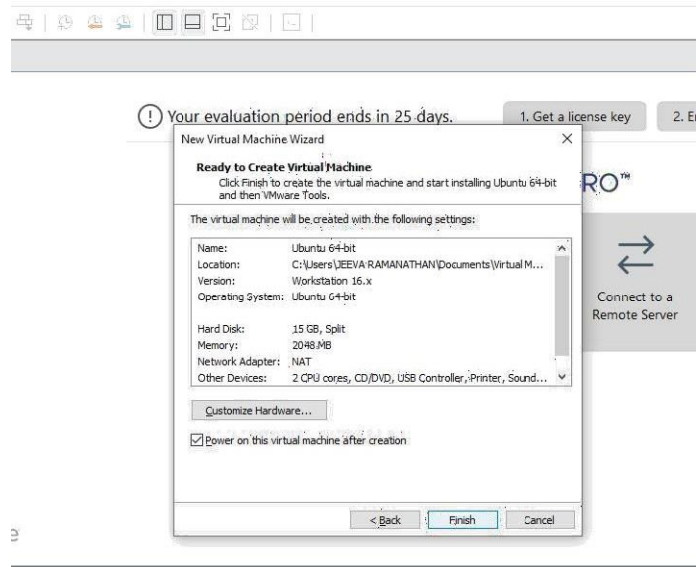
## Step 18: Specify the disk size and click next



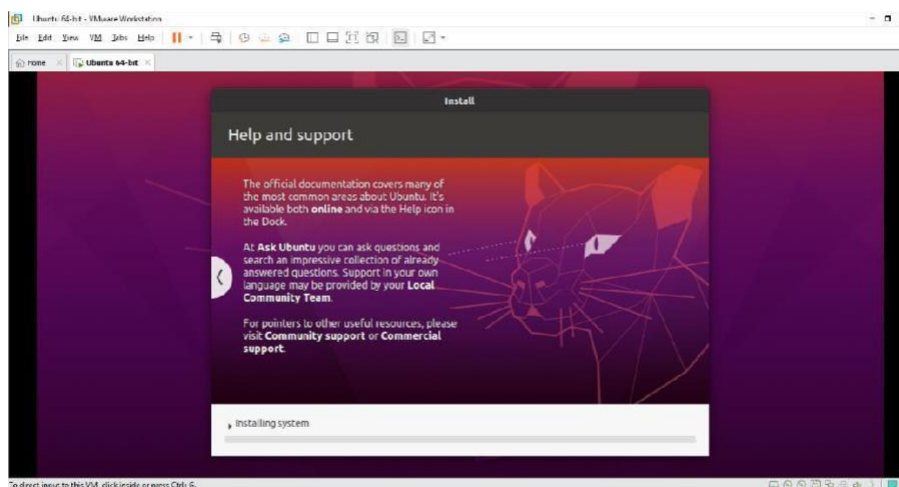
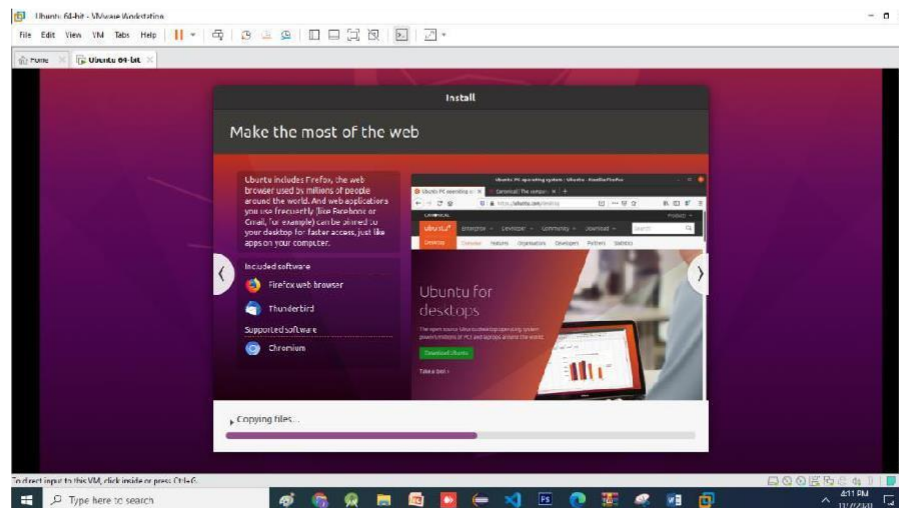
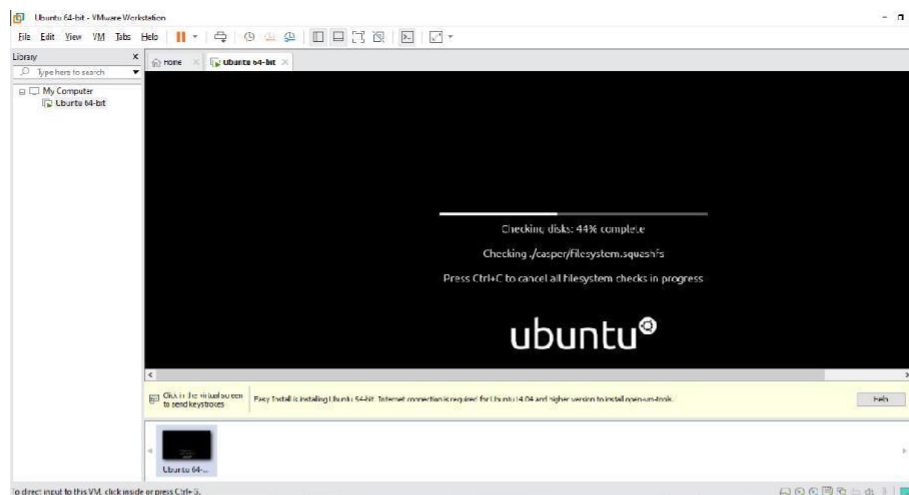
## Step 19: Click Next



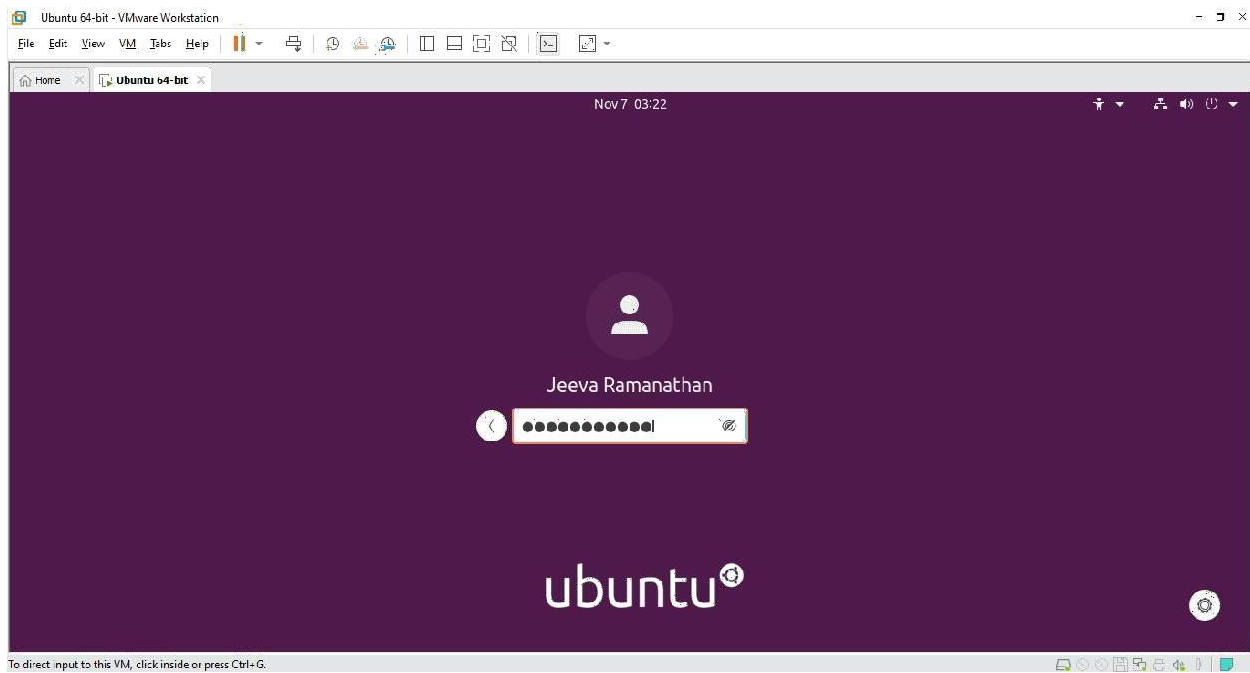
## Step 20: Click Finish



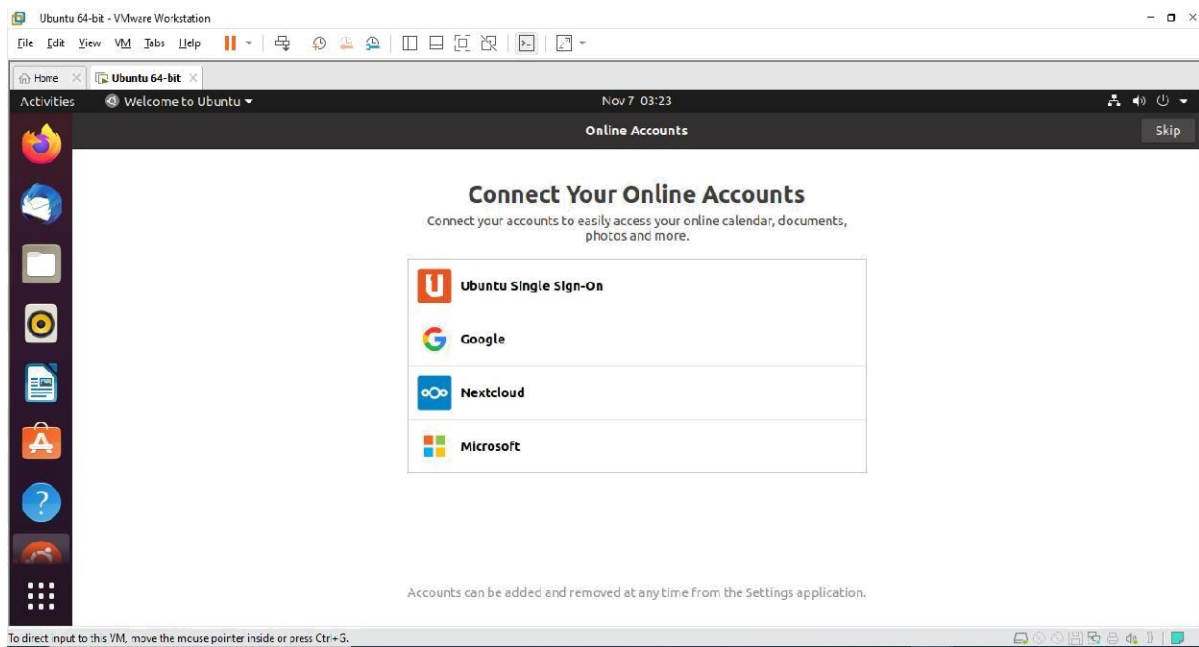
## Step 21: Installing Ubuntu on VMware and unzipping files



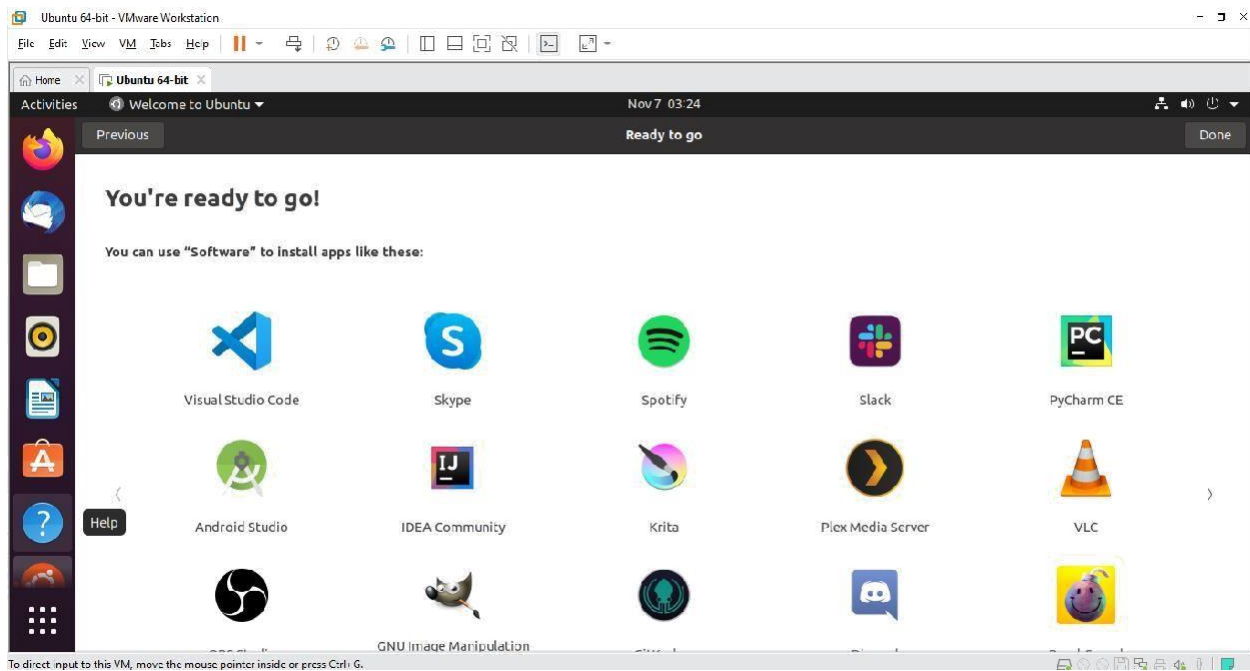
## Step 22: Login to Ubuntu



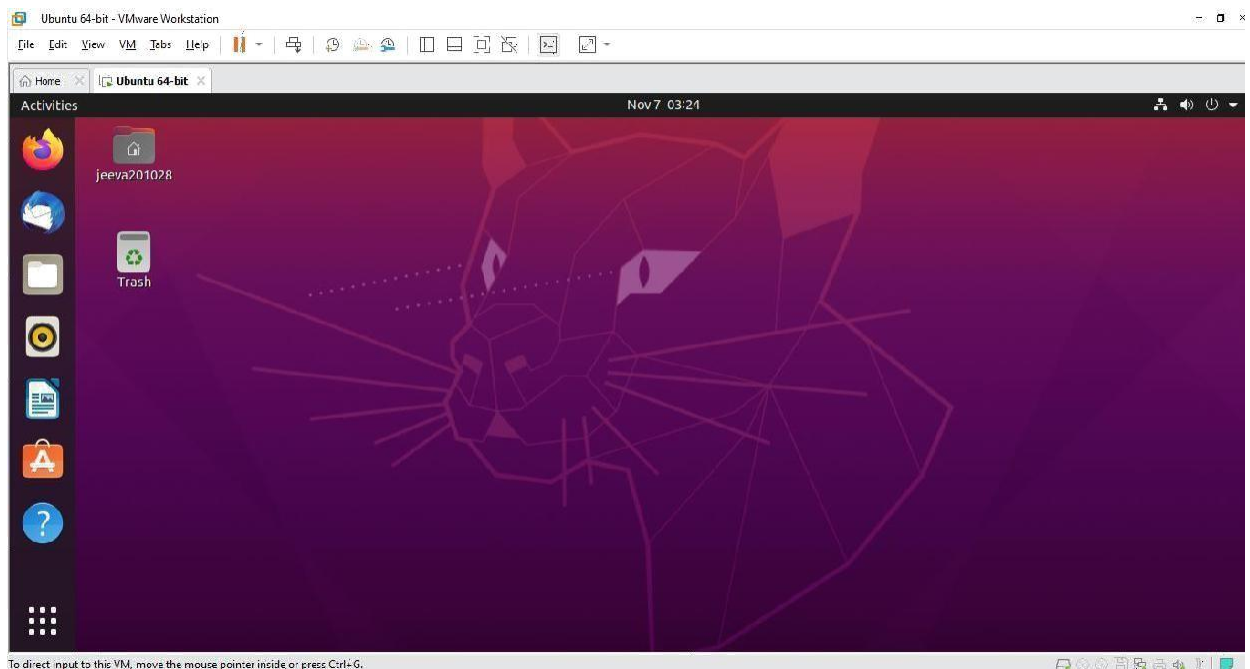
## Step 23: Skip everything



## Step 24: Click Done



Step 25: Thus we have installed VMware Workstation with different flavours of Linux on top of Windows



**Result:** Thus, VMware Workstation with different flavours of Linux or Windows OS on top of Windows 7 or 8 or 10 has been successfully installed and executed.

## EX.No:2

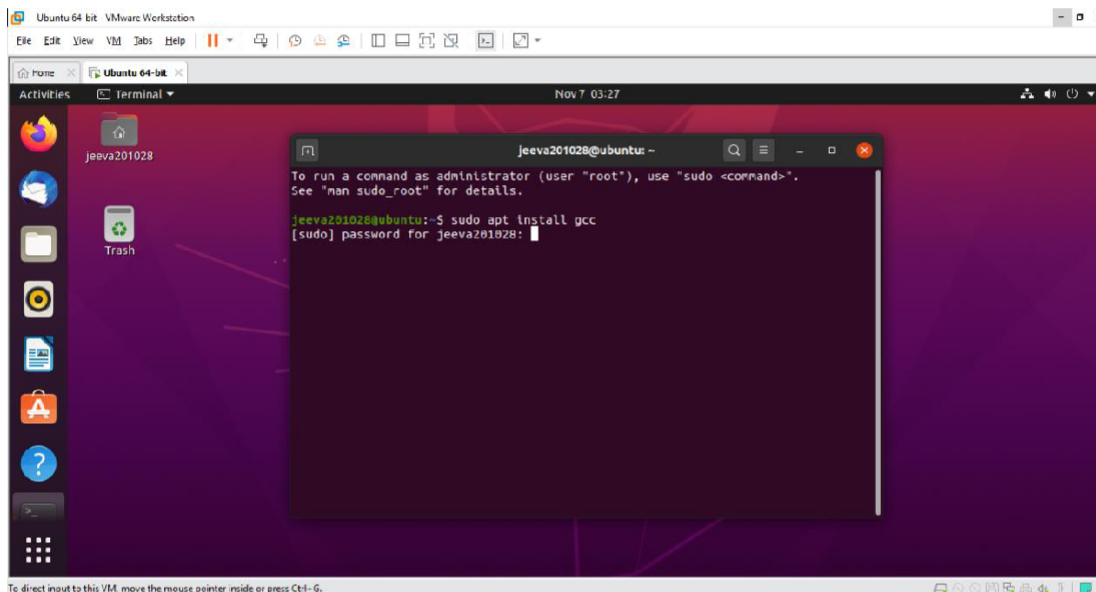
## Install a C compiler in the virtual machine created using virtual box and execute a simple program

### Aim:

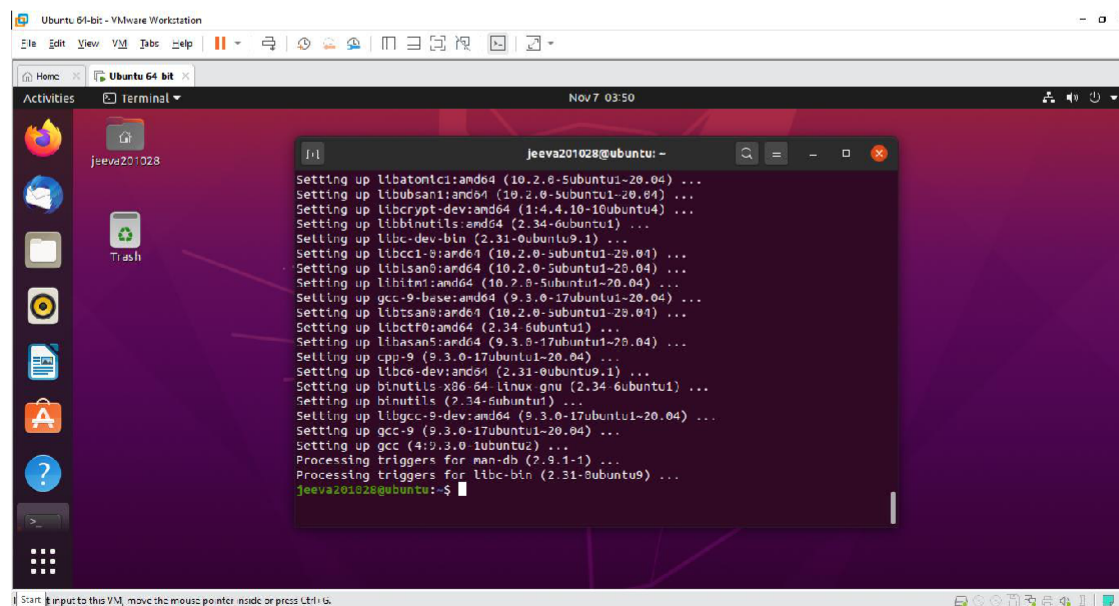
To install a C Compiler in the virtual machine created using virtual box and executes a simple C program.

### Procedure:

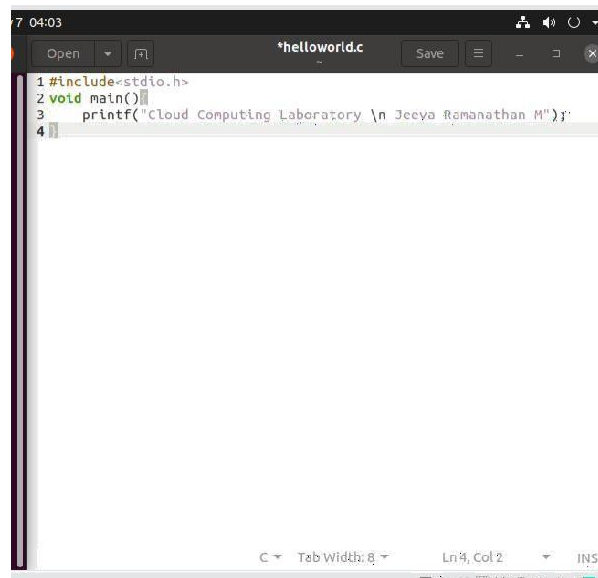
Step 1: Open the terminal on Ubuntu and install C compiler – “sudo apt install gcc”



Step 2: Once the installation finished open the editor

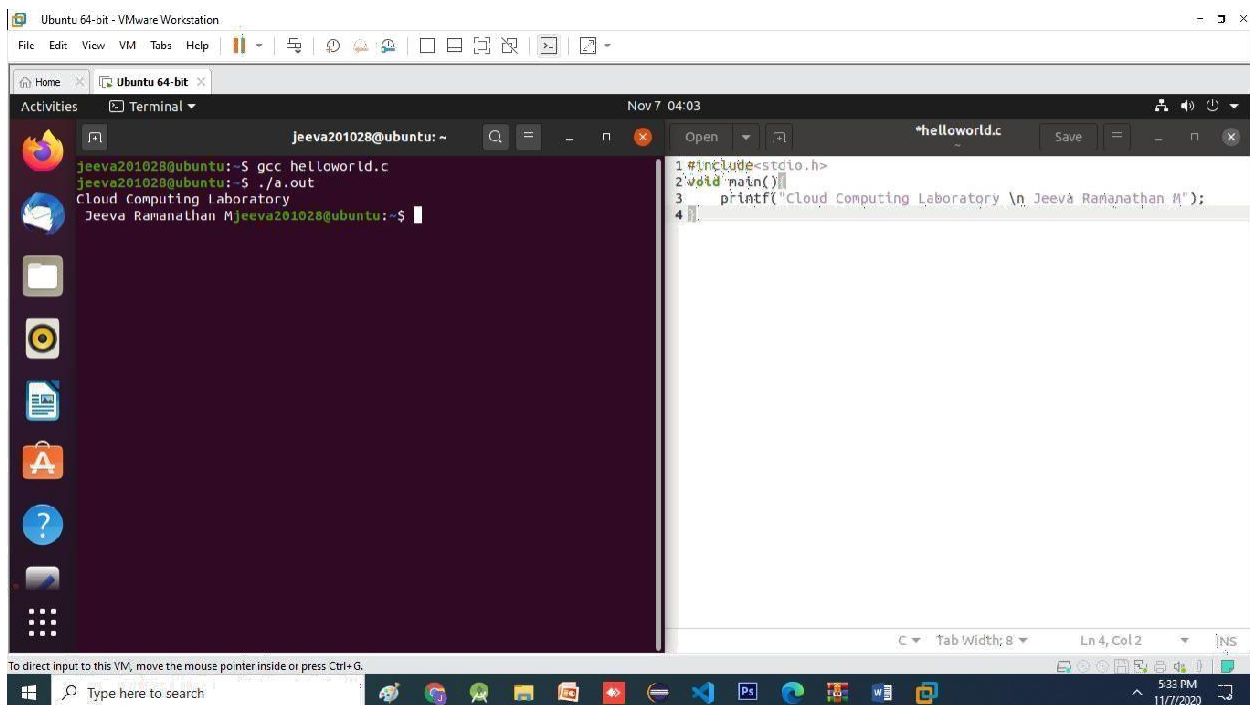


Step 3: Type a simple C program on editor and save it



```
7 04:03
*helloworld.c
1 #include <stdio.h>
2 void main()
3     printf("Cloud Computing Laboratory \n Jeeva Ramanathan M");
4
```

Step 4: Compile and run the C program



```
Ubuntu 64-bit - VMware Workstation
File Edit View VM Tabs Help
Home x Ubuntu 64-bit x
Activities Terminal Nov 7 04:03
jееva201028@ubuntu: ~
jееva201028@ubuntu:~$ gcc helloworld.c
jееva201028@ubuntu:~$ ./a.out
Cloud Computing Laboratory
Jeeva Ramanathan Mjееva201028@ubuntu:~$
```

**Result:**

Thus a C compiler is installed in the virtual machine and C program was executed and output was obtained successfully.

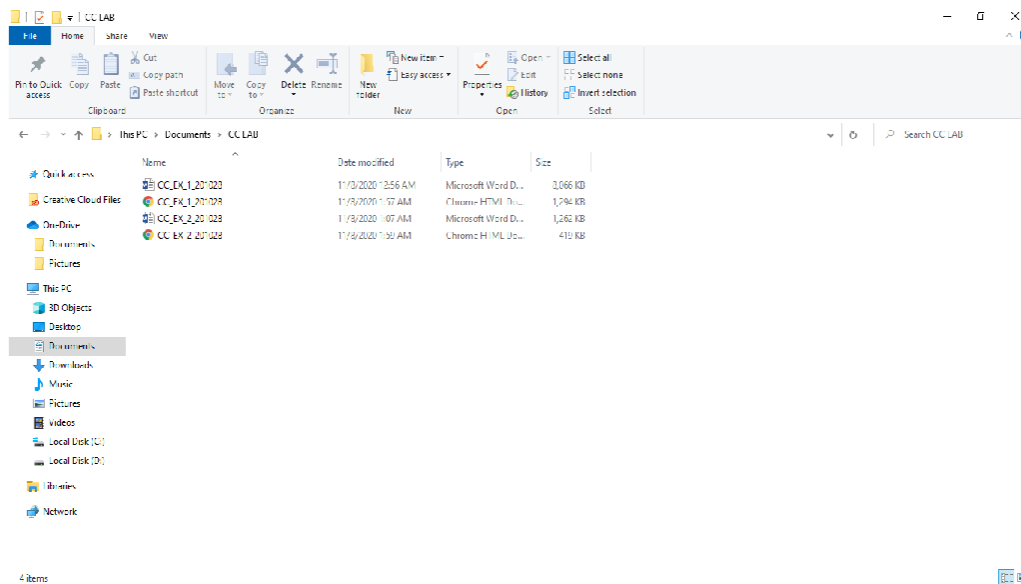
## EX.No:3      Transfer files from one host machine to another virtual machine

### Aim:

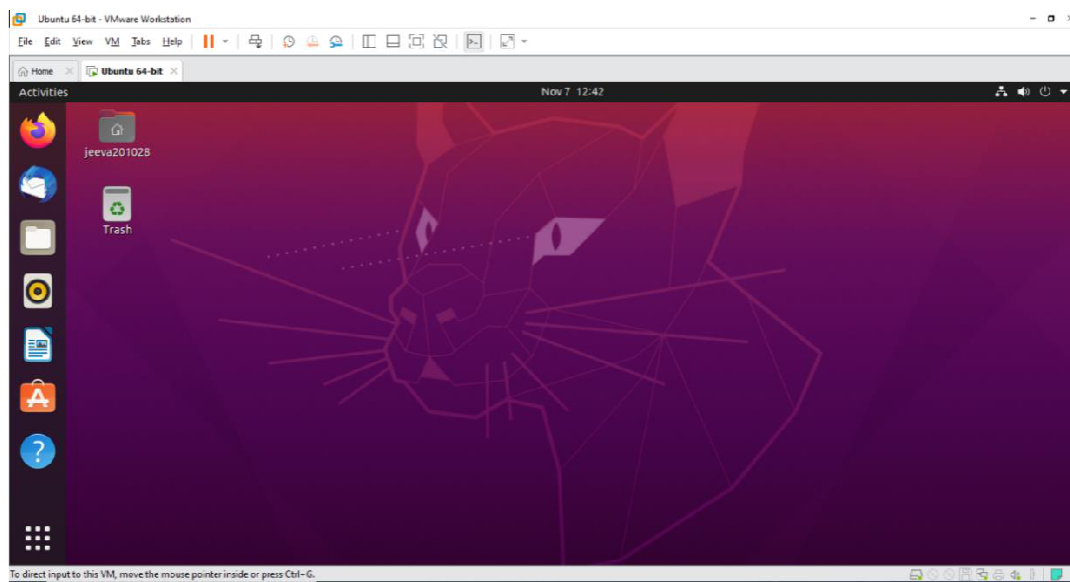
To transfer files/folders from the host machine to the virtual machine.

### Procedure:

Files that are to be shared to virtual machine

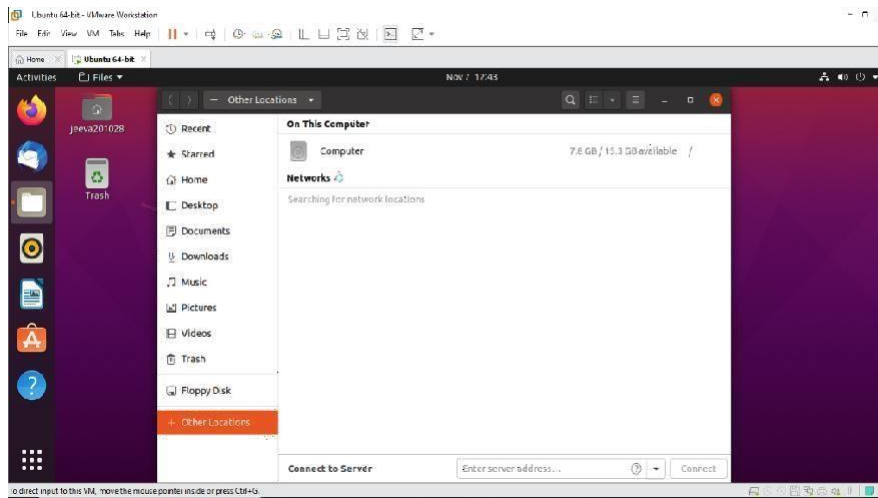


Step 1: Open the ubuntu in VM

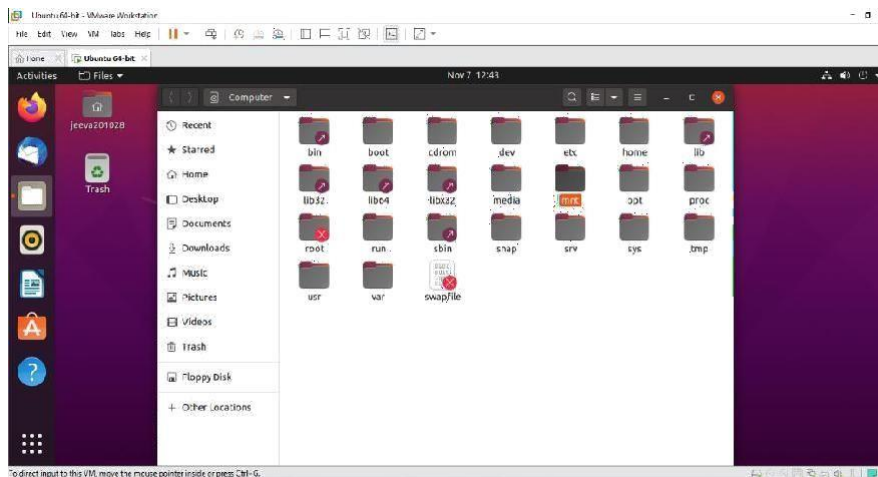




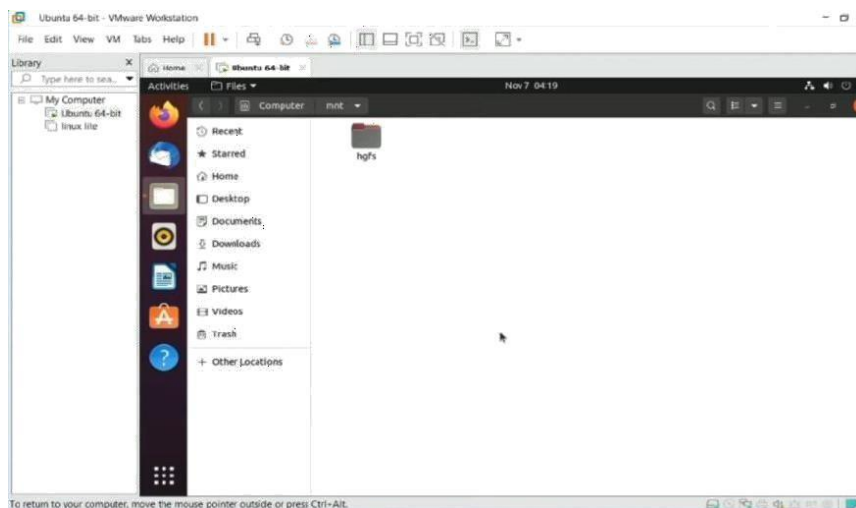
Step 2: Open the file manager and go to Other Locations



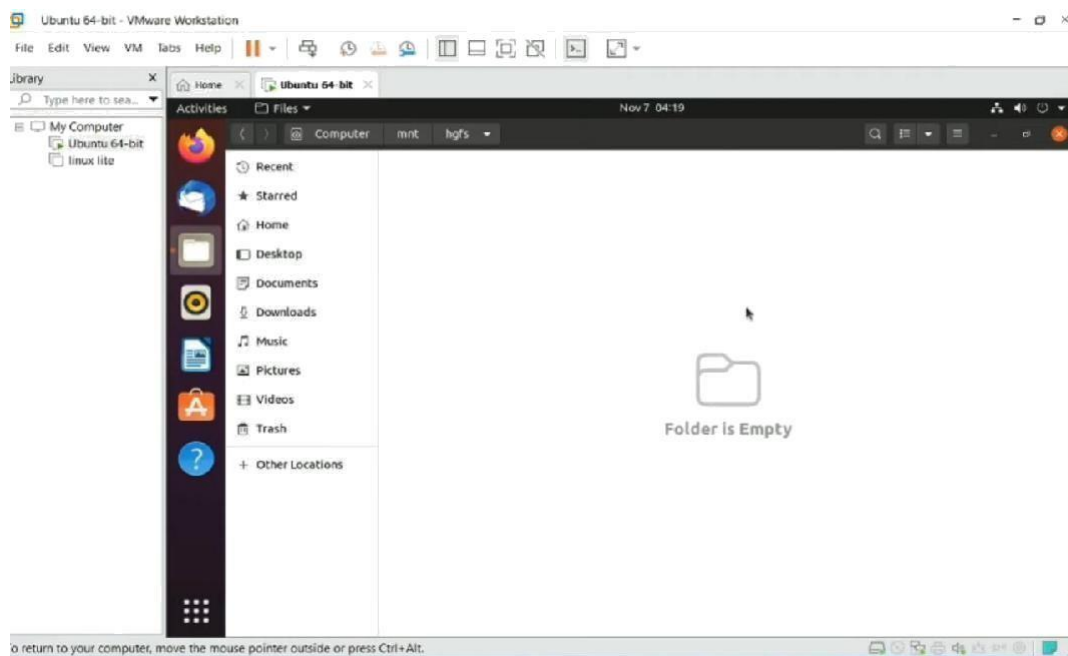
Step 3: Open the folder named “mnt”



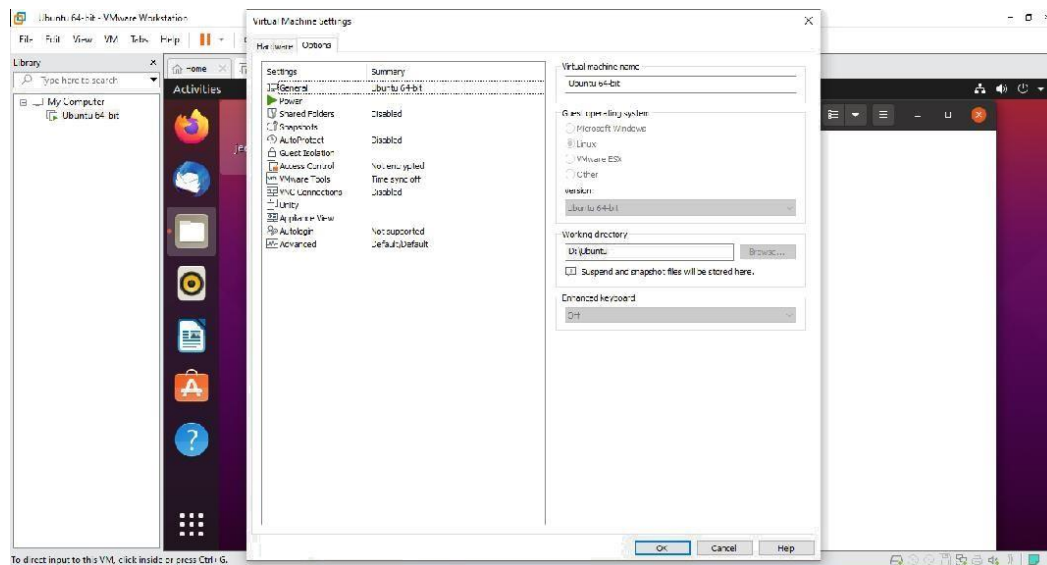
Step 4: Now open the “hgfs” folder



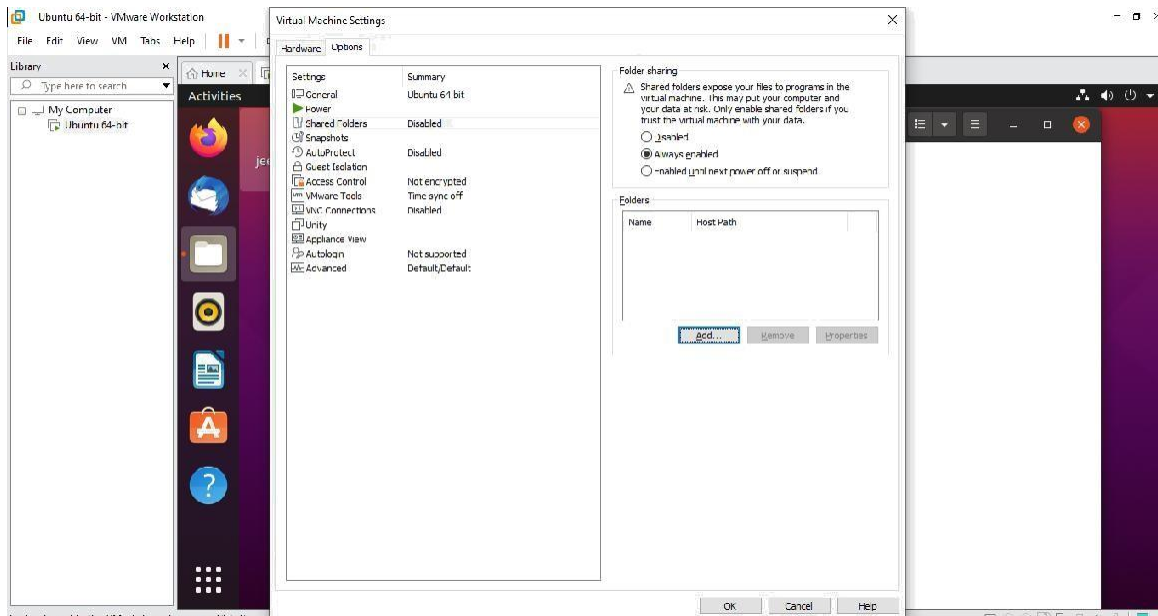
Initially the folder will be empty



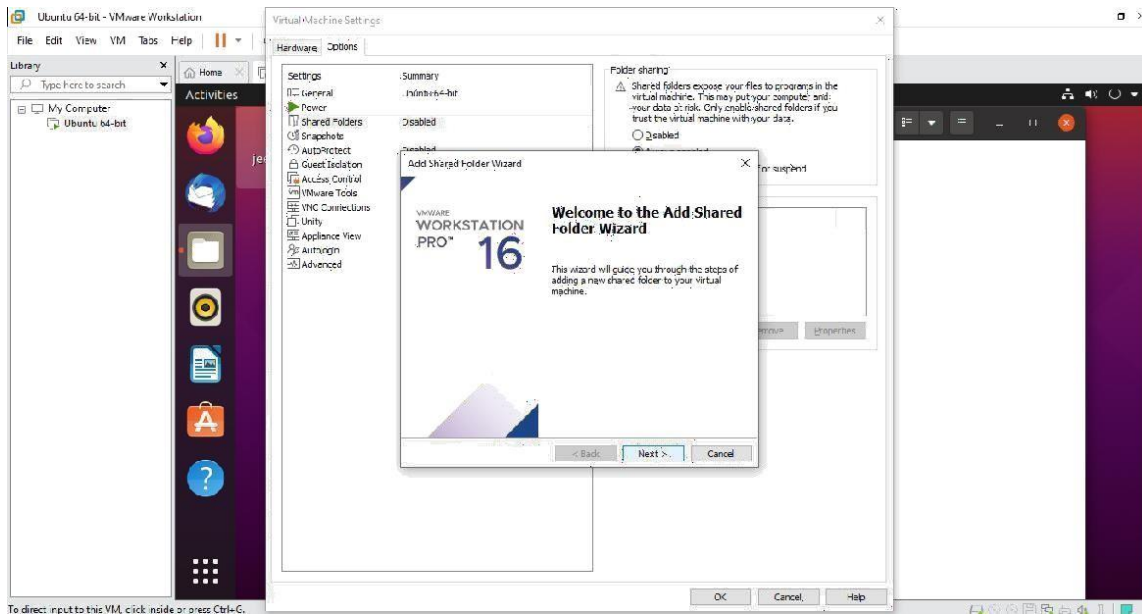
Step 5: Now right click ubuntu 64-bit(VM name) and select properties, then go to options tab



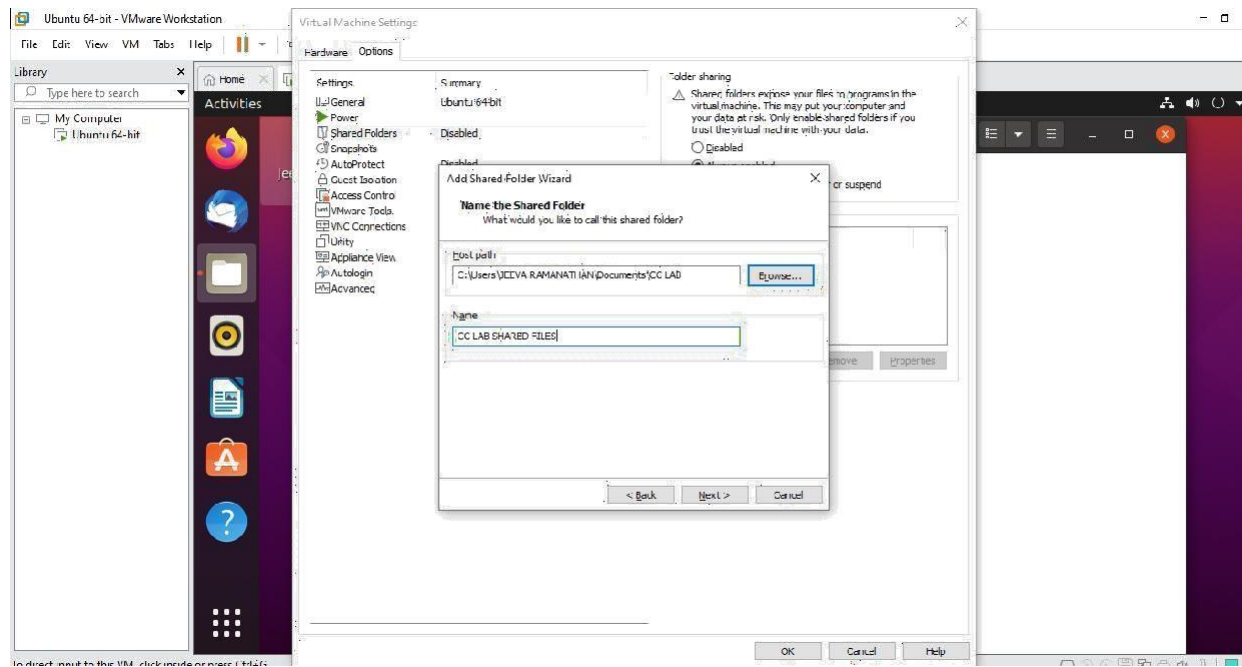
Step 6: Select the shared folders, change the radio button to “Always enabled” and click Add



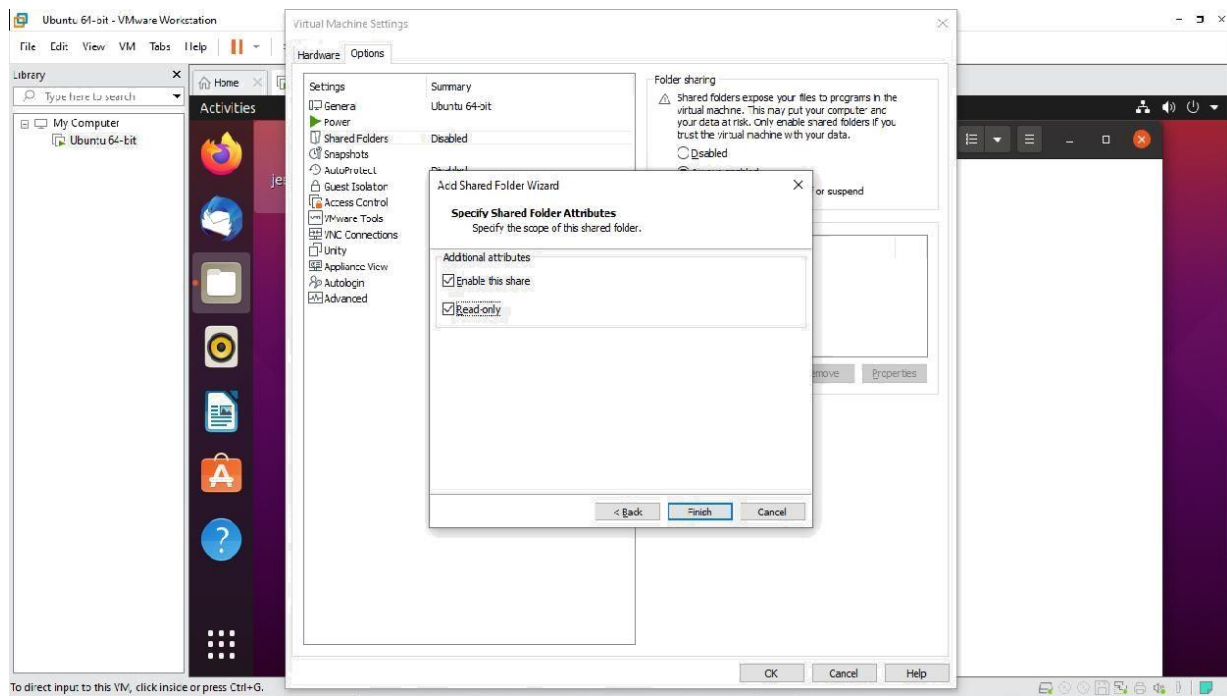
Step 7: Click Next



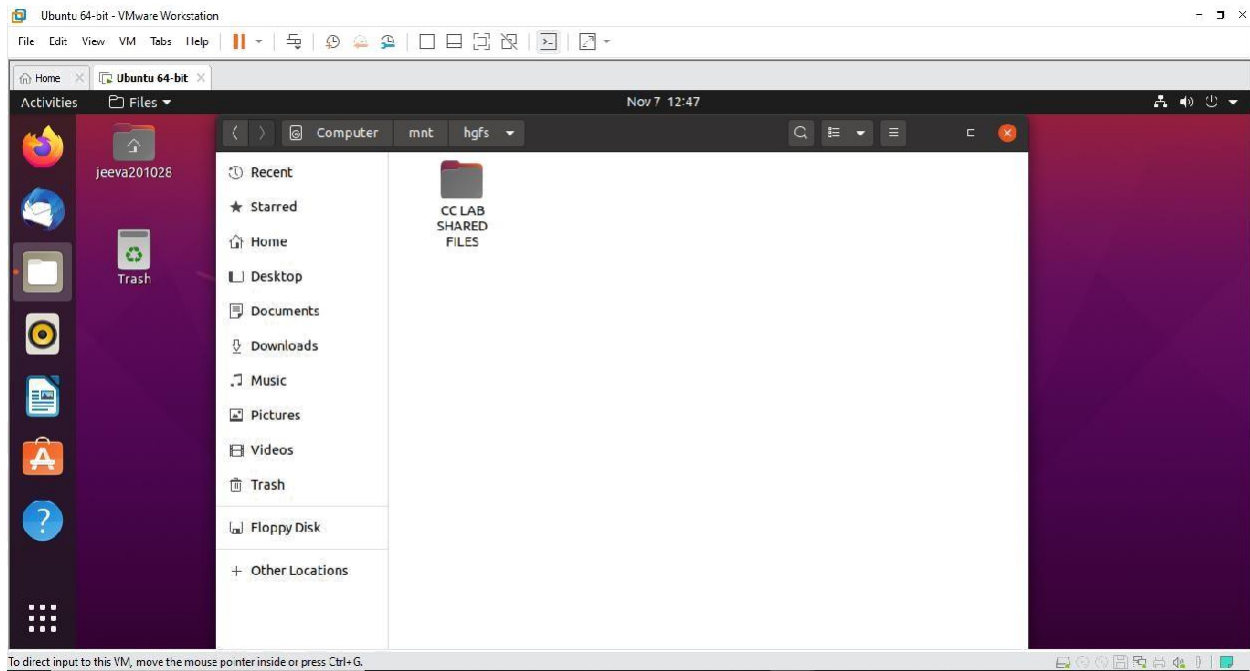
Step 8. Select the folder/file that has to be shared to VM in Host Path



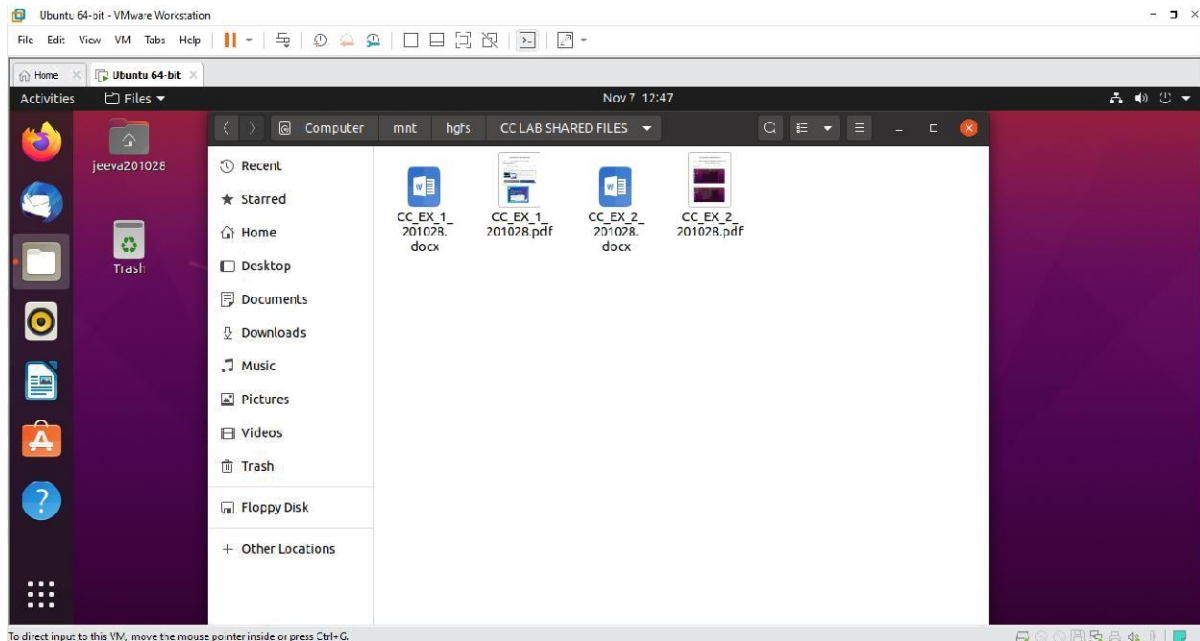
Step 9: Check Read Only and click Finish



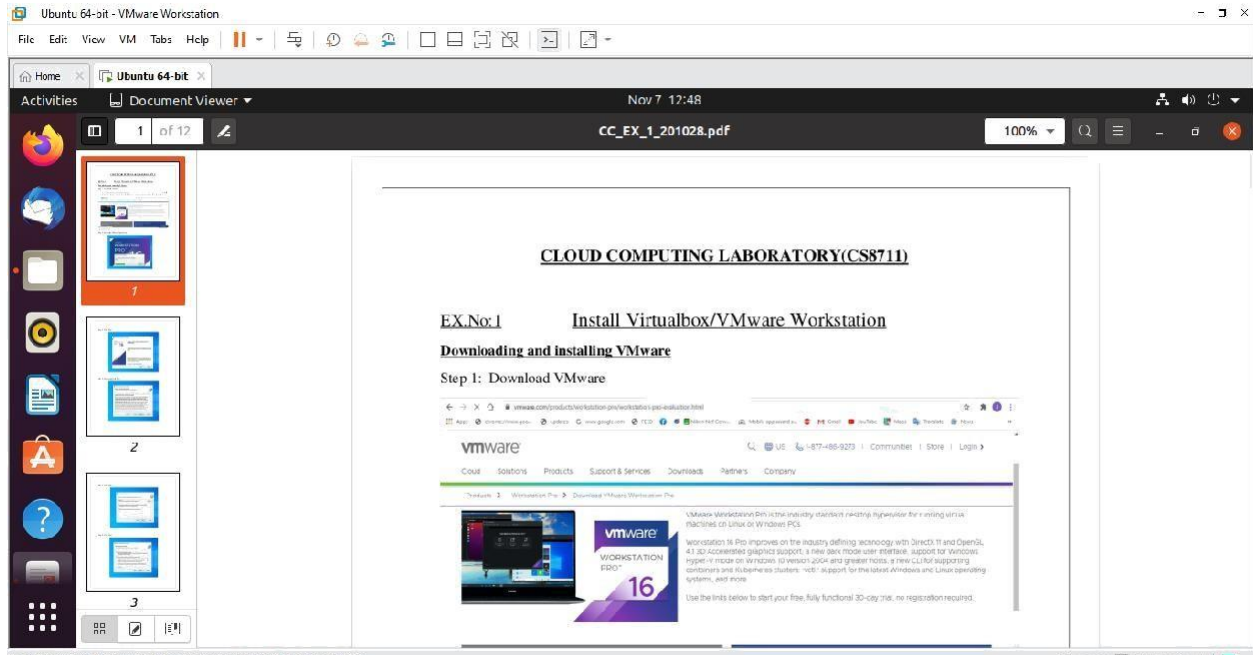
Step 10: Now in the same location “Computer>mnt>hgfs” the folder that are shared from hostis visible



Step 11:



Step 12: The files can be accessed in Read Only Mode in the VM now



**Result:**

Thus file sharing between is done between host machine and the virtual machine successfully.

## EX.No:4

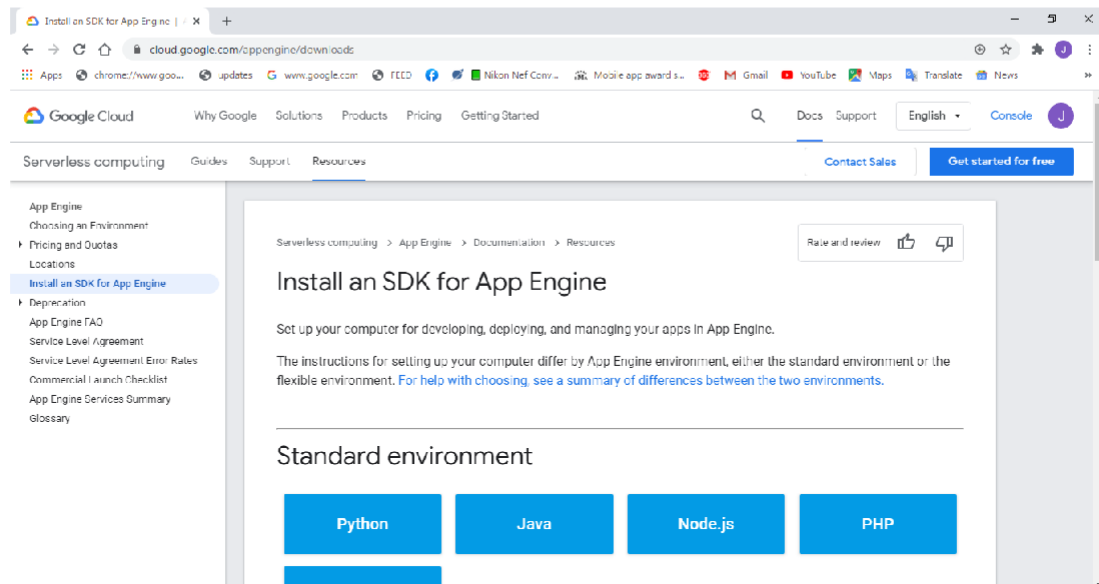
## Installation of a Google App Engine

### Aim:

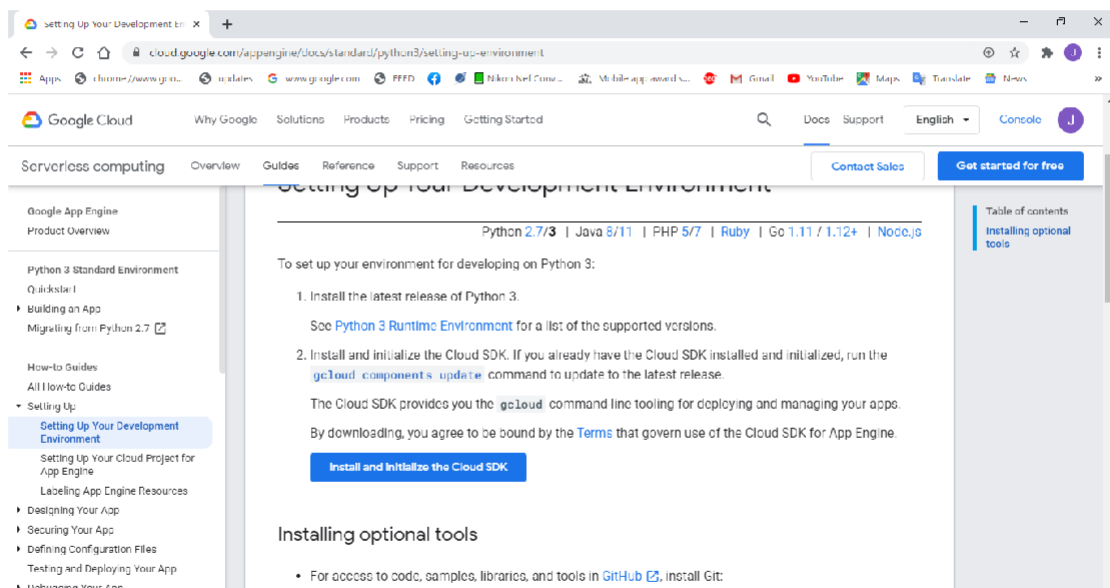
To install the google app engine in the system.

### Procedure:

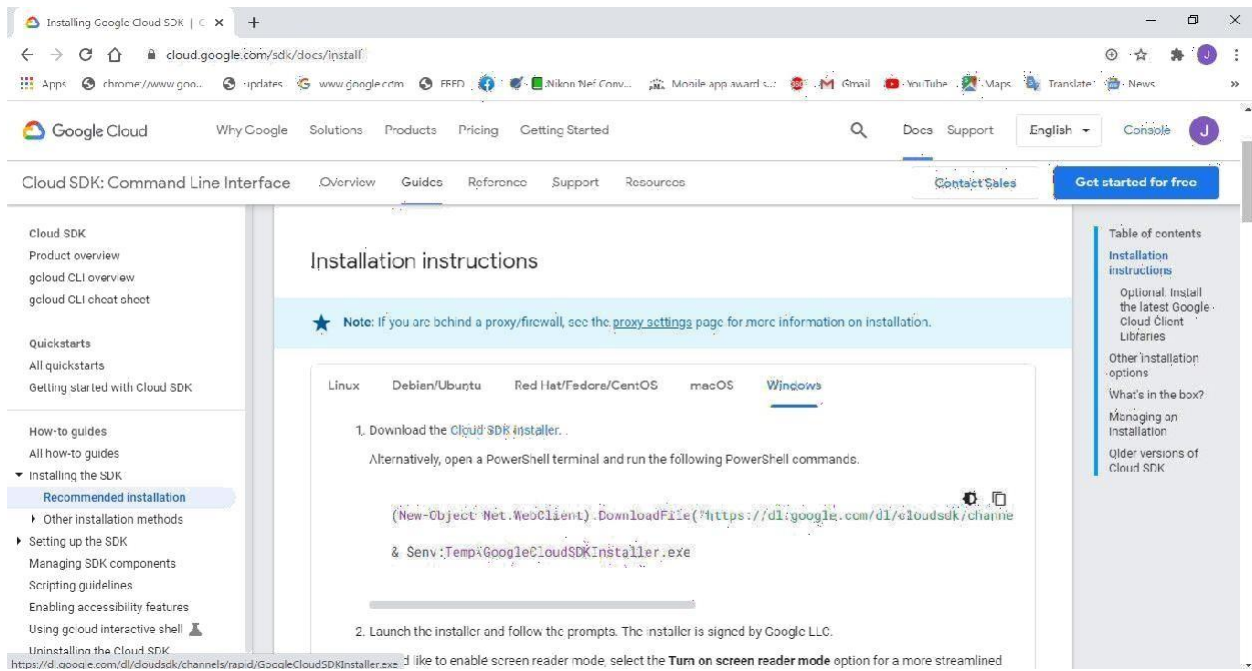
Step 1: Open the following link - <https://cloud.google.com/appengine/downloads> and clickpython.



Step 2: Select setting up your environment development and click on install the and initialize the cloud sdk.



### Step 3: Download the sdk installer and install it.



The screenshot shows a web browser window displaying the Google Cloud SDK installation instructions page. The page title is "Installation instructions" and it includes a note about proxy/firewall settings. The instructions are categorized by operating system: Linux, Debian/Ubuntu, Red Hat/Fedora/CentOS, macOS, and Windows. The Windows section is active, showing the following steps:

1. Download the [Cloud SDK installer](#).

Alternatively, open a PowerShell terminal and run the following PowerShell commands:

```
(New-Object Net.WebClient).DownloadFile(https://dl.google.com/dl/cloudsdk/channel  
& .env\Temp\GoogleCloudSDKInstaller.exe
```

2. Launch the installer and follow the prompts. The installer is signed by Google LLC.

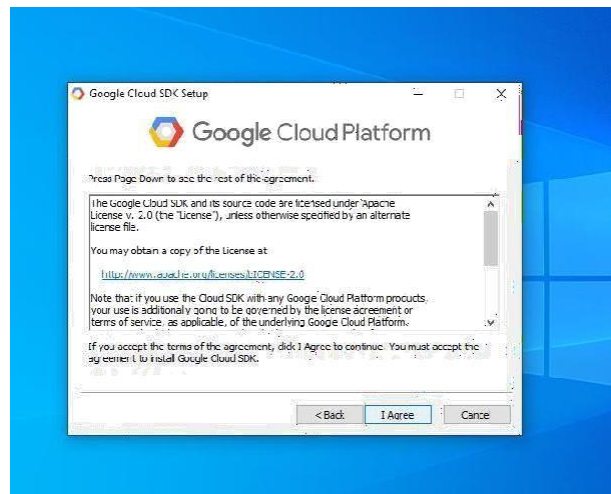
At the bottom of the page, there is a link to enable screen reader mode: <https://dl.google.com/dl/cloudsdk/channels/rapid/GoogleCloudSDKInstaller.exe> like to enable screen reader mode, select the **Turn on screen reader mode** option for a more streamlined

### Step 4: Click Next.

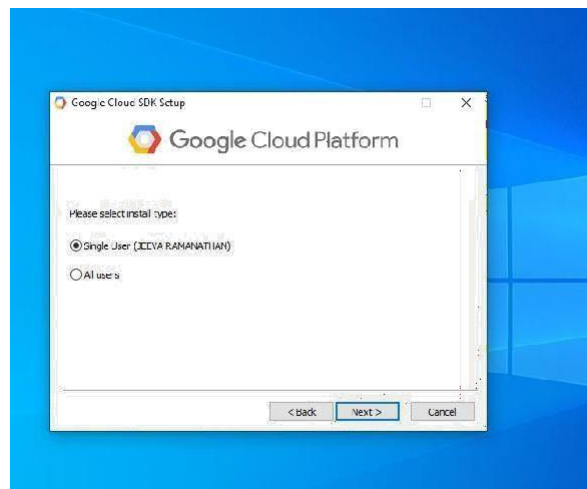




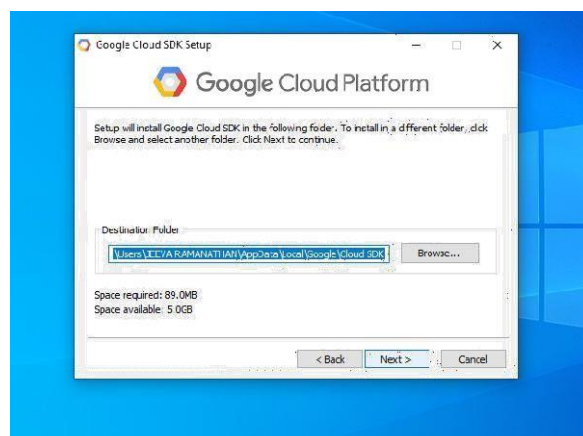
Step 5: Click I Agree.



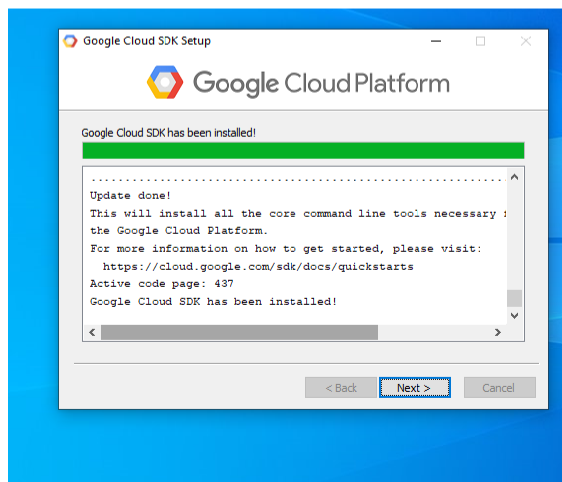
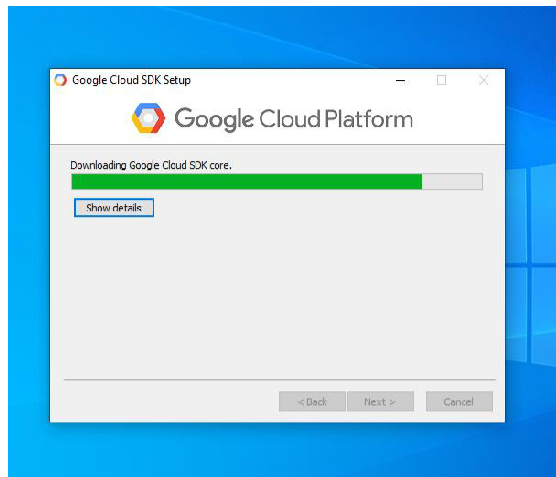
Step 6: Select single user and click Next.



Step 7: Select the destination location and click Next.



## Step 8: Downloading all the requirements and installing



## Step 9: Click Finish.



Step 10: Once successfully installed cmd line in login with your google account.

```
cmd C:\WINDOWS\SYSTEM32\cmd.exe - gcloud init
Welcome to the Google Cloud SDK! Run "gcloud -h" to get the list of available commands.
---
Welcome! This command will take you through the configuration of gcloud.

Your current configuration has been set to: [default]

You can skip diagnostics next time by using the following flag:
  gcloud init --skip-diagnostics

Network diagnostic detects and fixes local network connection issues.
Checking network connection...done.
Reachability Check passed.
Network diagnostic passed (1/1 checks passed).

You must log in to continue. Would you like to log in (Y/n)? Y
-
```

```
Administrator: Google Cloud SDK Shell
gcloud logging --help

C:\Users\JEEVA RAMANATHAN\AppData\Local\Google\Cloud SDK>gcloud auth login
Your browser has been opened to visit:

  https://accounts.google.com/o/oauth2/auth?client_id=32555940559.apps.googleusercontent.com&redirect_uri=http%3A%2F%2Flocalhost%3A8085%2F&scope=openid%2Fhttps%3A%2F%2Fwww.googleapis.com%2Fauth%2Fuserinfo.email%2Fhttps%3A%2F%2Fwww.googleapis.com%2Fauth%2Fcloud-platform%2Fhttps%3A%2F%2Fwww.googleapis.com%2Fauth%2Fappengine.admin%2Fhttps%3A%2F%2Fwww.googleapis.com%2Fauth%2Fcompute%2Fhttps%3A%2F%2Fwww.googleapis.com%2Fauth%2Faccounts.reauth&code_challenge=fr6e8PLwbenQ3uChkqT1140F00WqjIiqD6dYj9Lae0&code_challenge_method=S256&access_type=offline&response_type=code&prompt=select_account

You are now logged in as [jeevaram2000@gmail.com].
Your current project is [None]. You can change this setting by running:
  $ gcloud config set project PROJECT_ID

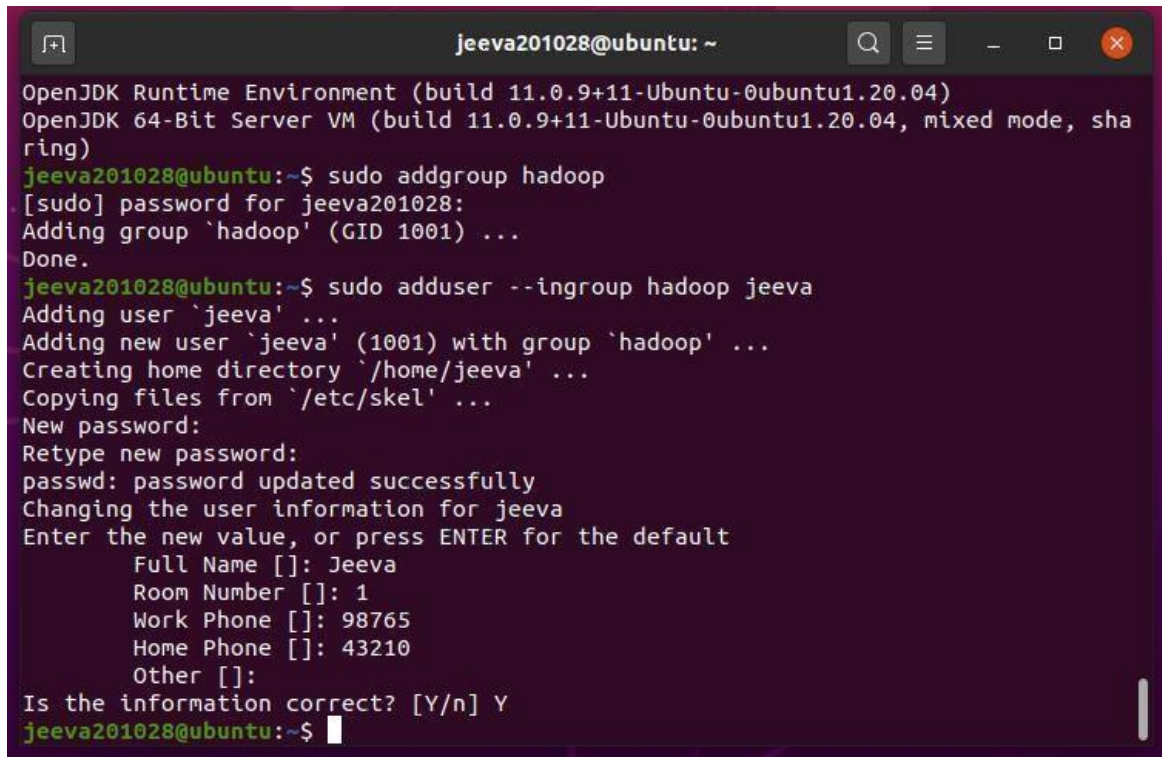
C:\Users\JEEVA RAMANATHAN\AppData\Local\Google\Cloud SDK>
```

**Result:**

Thus google app engine is installed successfully in the system.

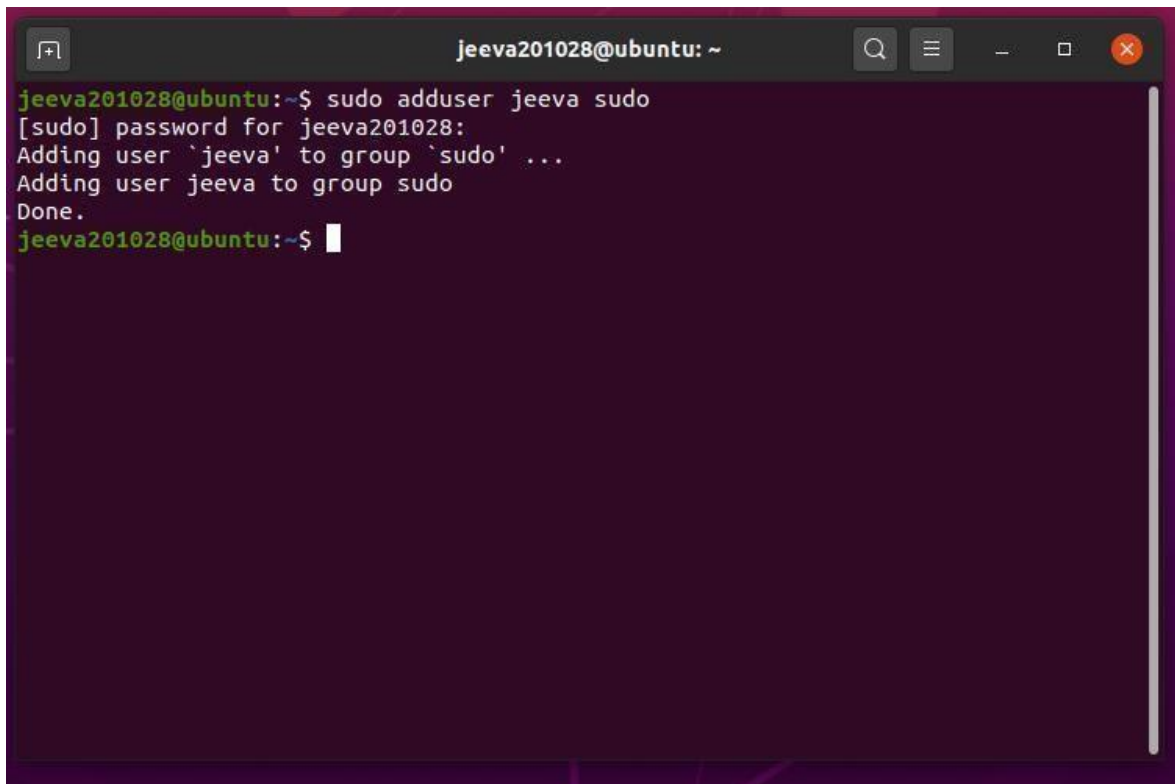


Assigning a dedicated user on hadoop to perform operations



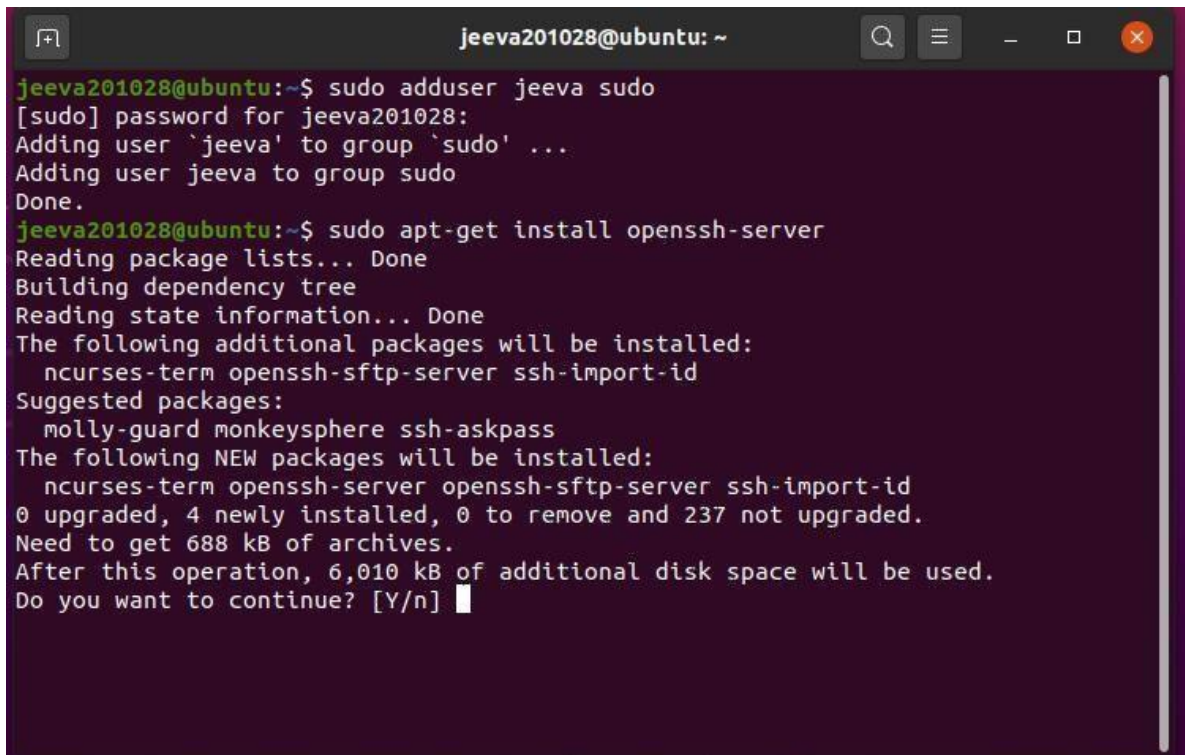
```
jeeva201028@ubuntu: ~  
OpenJDK Runtime Environment (build 11.0.9+11-Ubuntu-0ubuntu1.20.04)  
OpenJDK 64-Bit Server VM (build 11.0.9+11-Ubuntu-0ubuntu1.20.04, mixed mode, sha  
ring)  
jeeva201028@ubuntu:~$ sudo addgroup hadoop  
[sudo] password for jeeva201028:  
Adding group `hadoop' (GID 1001) ...  
Done.  
jeeva201028@ubuntu:~$ sudo adduser --ingroup hadoop jeeva  
Adding user `jeeva' ...  
Adding new user `jeeva' (1001) with group `hadoop' ...  
Creating home directory `/home/jeeva' ...  
Copying files from `/etc/skel' ...  
New password:  
Retype new password:  
passwd: password updated successfully  
Changing the user information for jeeva  
Enter the new value, or press ENTER for the default  
  Full Name []: Jeeva  
  Room Number []: 1  
  Work Phone []: 98765  
  Home Phone []: 43210  
  Other []:  
Is the information correct? [Y/n] Y  
jeeva201028@ubuntu:~$
```

Adding user to sudo list



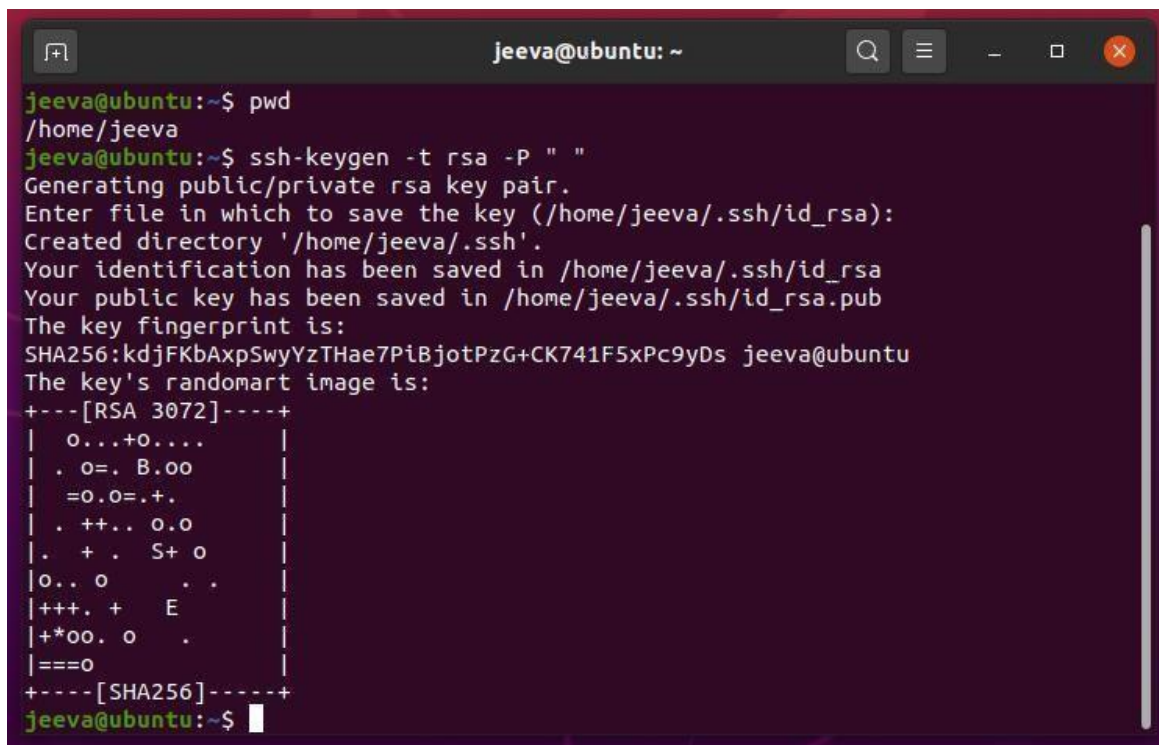
```
jeeva201028@ubuntu: ~  
jeeva201028@ubuntu:~$ sudo adduser jeeva sudo  
[sudo] password for jeeva201028:  
Adding user `jeeva' to group `sudo' ...  
Adding user jeeva to group sudo  
Done.  
jeeva201028@ubuntu:~$
```

Next install a package ssh(secured shell login)



```
jeeva201028@ubuntu: ~  
jeeva201028@ubuntu:~$ sudo adduser jeeva sudo  
[sudo] password for jeeva201028:  
Adding user `jeeva' to group `sudo' ...  
Adding user jeeva to group sudo  
Done.  
jeeva201028@ubuntu:~$ sudo apt-get install openssh-server  
Reading package lists... Done  
Building dependency tree  
Reading state information... Done  
The following additional packages will be installed:  
  ncurses-term openssh-sftp-server ssh-import-id  
Suggested packages:  
  molly-guard monkeysphere ssh-askpass  
The following NEW packages will be installed:  
  ncurses-term openssh-server openssh-sftp-server ssh-import-id  
0 upgraded, 4 newly installed, 0 to remove and 237 not upgraded.  
Need to get 688 kB of archives.  
After this operation, 6,010 kB of additional disk space will be used.  
Do you want to continue? [Y/n]
```

Next step is key generation and add a key to the file

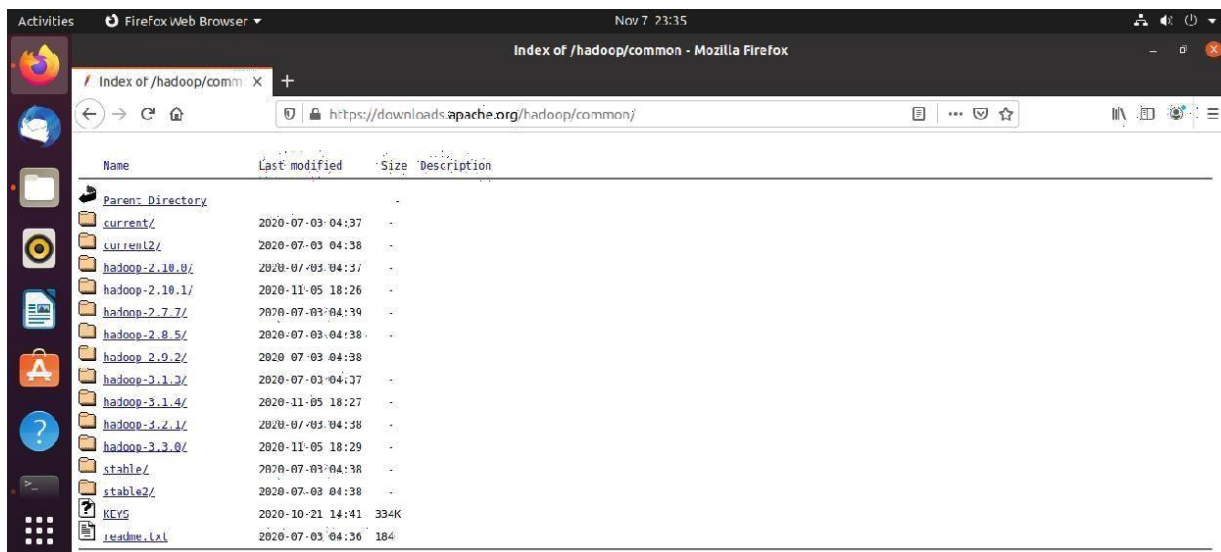


```
jeeva@ubuntu: ~  
jeeva@ubuntu:~$ pwd  
/home/jeeva  
jeeva@ubuntu:~$ ssh-keygen -t rsa -P ""  
Generating public/private rsa key pair.  
Enter file in which to save the key (/home/jeeva/.ssh/id_rsa):  
Created directory '/home/jeeva/.ssh'.  
Your identification has been saved in /home/jeeva/.ssh/id_rsa  
Your public key has been saved in /home/jeeva/.ssh/id_rsa.pub  
The key fingerprint is:  
SHA256:kdjFKbAxpSwyYzTHae7PiBjotPzG+CK741F5xPc9yDs jeeva@ubuntu  
The key's randomart image is:  
+---[RSA 3072]-----+  
|  o...+o....      |  
| . o= . B.oo      |  
| =o.o=.+.         |  
| . ++.. o.o       |  
| . + . S+ o       |  
|o.. o . .         |  
|+++ . + E         |  
|+*oo. o .         |  
|===o              |  
+----[SHA256]-----+  
jeeva@ubuntu:~$
```

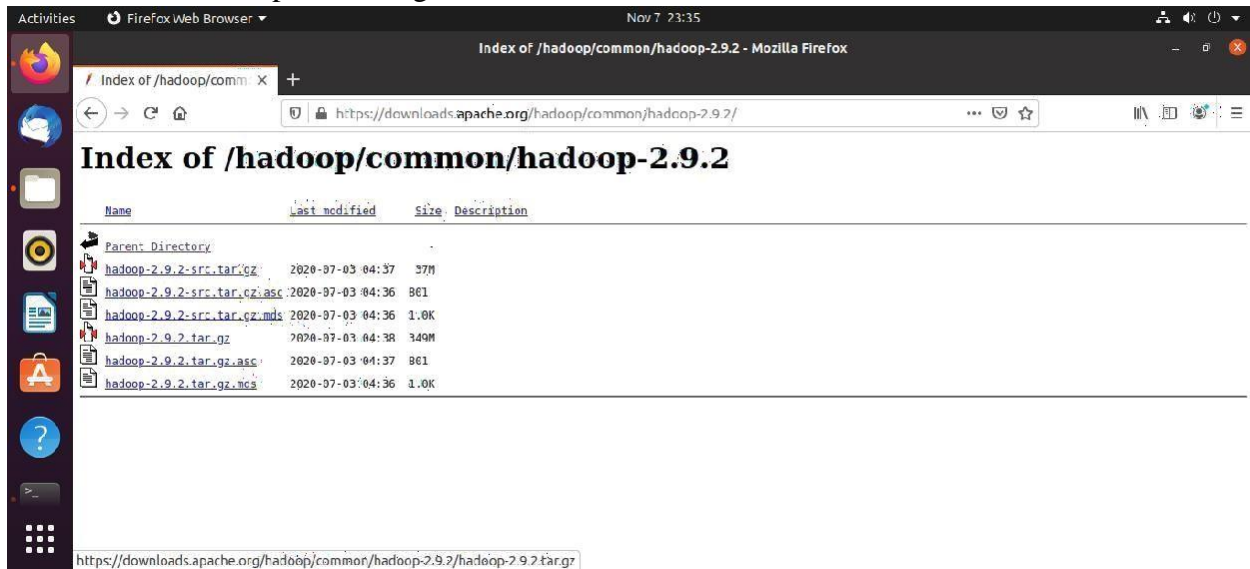
To check whether ssh is installed properly by logging in and after that exit from it

```
jeeva201028@ubuntu: ~  
jeeva201028@ubuntu:~$ ssh localhost  
The authenticity of host 'localhost (127.0.0.1)' can't be established.  
ECDSA key fingerprint is SHA256:3w08Zvw2Q8dtjtAaSLLWahjenZ7zywgGfYjr0Im8eYM.  
Are you sure you want to continue connecting (yes/no/[fingerprint])? yes  
Warning: Permanently added 'localhost' (ECDSA) to the list of known hosts.  
jeeva201028@localhost's password:  
Welcome to Ubuntu 20.04.1 LTS (GNU/Linux 5.4.0-52-generic x86_64)  
  
* Documentation:  https://help.ubuntu.com  
* Management:    https://landscape.canonical.com  
* Support:       https://ubuntu.com/advantage  
  
237 updates can be installed immediately.  
100 of these updates are security updates.  
To see these additional updates run: apt list --upgradable  
  
Your Hardware Enablement Stack (HWE) is supported until April 2025.  
  
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.  
  
jeeva201028@ubuntu:~$ exit  
Logout  
Connection to localhost closed.
```


Download Hadoop using this link : <https://downloads.apache.org/hadoop/common/>



## Download the hadoop-2.9.2.tar.gz file

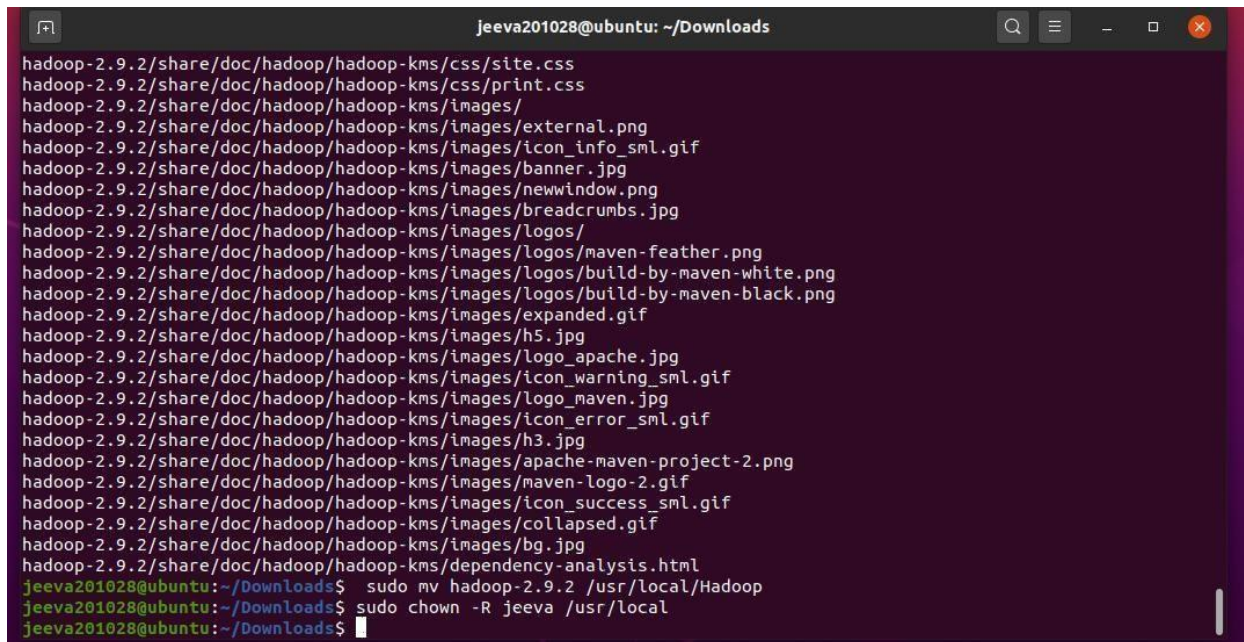


## Now extract the tar file



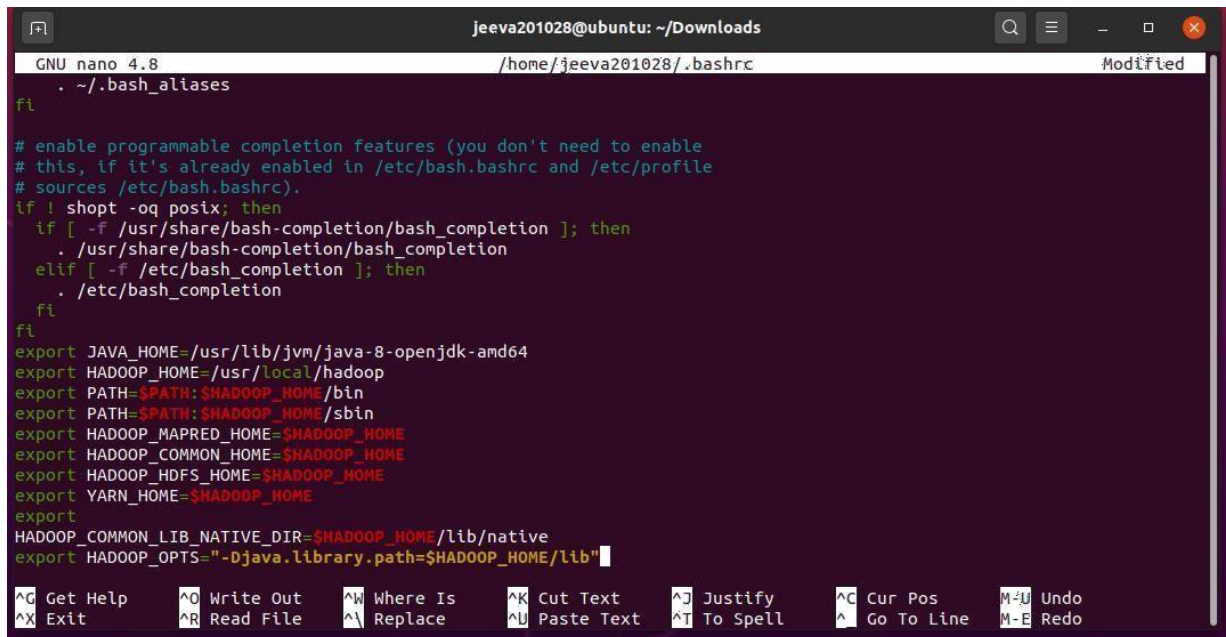


Move the file to local and Change the ownership of hadoop folder



```
jeeva201028@ubuntu: ~/Downloads
hadoop-2.9.2/share/doc/hadoop/hadoop-kms/css/site.css
hadoop-2.9.2/share/doc/hadoop/hadoop-kms/css/print.css
hadoop-2.9.2/share/doc/hadoop/hadoop-kms/images/
hadoop-2.9.2/share/doc/hadoop/hadoop-kms/images/external.png
hadoop-2.9.2/share/doc/hadoop/hadoop-kms/images/icon_info_sml.gif
hadoop-2.9.2/share/doc/hadoop/hadoop-kms/images/banner.jpg
hadoop-2.9.2/share/doc/hadoop/hadoop-kms/images/newwindow.png
hadoop-2.9.2/share/doc/hadoop/hadoop-kms/images/breadcrumbs.jpg
hadoop-2.9.2/share/doc/hadoop/hadoop-kms/images/logos/
hadoop-2.9.2/share/doc/hadoop/hadoop-kms/images/logos/maven-feather.png
hadoop-2.9.2/share/doc/hadoop/hadoop-kms/images/logos/build-by-maven-white.png
hadoop-2.9.2/share/doc/hadoop/hadoop-kms/images/logos/build-by-maven-black.png
hadoop-2.9.2/share/doc/hadoop/hadoop-kms/images/expanded.gif
hadoop-2.9.2/share/doc/hadoop/hadoop-kms/images/h5.jpg
hadoop-2.9.2/share/doc/hadoop/hadoop-kms/images/logo_apache.jpg
hadoop-2.9.2/share/doc/hadoop/hadoop-kms/images/icon_warning_sml.gif
hadoop-2.9.2/share/doc/hadoop/hadoop-kms/images/logo_maven.jpg
hadoop-2.9.2/share/doc/hadoop/hadoop-kms/images/icon_error_sml.gif
hadoop-2.9.2/share/doc/hadoop/hadoop-kms/images/h3.jpg
hadoop-2.9.2/share/doc/hadoop/hadoop-kms/images/apache-maven-project-2.png
hadoop-2.9.2/share/doc/hadoop/hadoop-kms/images/maven-logo-2.gif
hadoop-2.9.2/share/doc/hadoop/hadoop-kms/images/icon_success_sml.gif
hadoop-2.9.2/share/doc/hadoop/hadoop-kms/images/collapsed.gif
hadoop-2.9.2/share/doc/hadoop/hadoop-kms/images/bg.jpg
hadoop-2.9.2/share/doc/hadoop/hadoop-kms/dependency-analysis.html
jeeva201028@ubuntu:~/Downloads$ sudo mv hadoop-2.9.2 /usr/local/Hadoop
jeeva201028@ubuntu:~/Downloads$ sudo chown -R jeeva /usr/local
jeeva201028@ubuntu:~/Downloads$
```

In .bashrc file add the following and make it at source



```
GNU nano 4.8 /home/jeeva201028/.bashrc Modified
. ~/.bash_aliases
fi

# enable programmable completion features (you don't need to enable
# this, if it's already enabled in /etc/bash.bashrc and /etc/profile
# sources /etc/bash.bashrc).
if ! shopt -oq posix; then
  if [ -f /usr/share/bash-completion/bash_completion ]; then
    . /usr/share/bash-completion/bash_completion
  elif [ -f /etc/bash_completion ]; then
    . /etc/bash_completion
  fi
fi

export JAVA_HOME=/usr/lib/jvm/java-8-openjdk-amd64
export HADOOP_HOME=/usr/local/hadoop
export PATH=$PATH:$HADOOP_HOME/bin
export PATH=$PATH:$HADOOP_HOME/sbin
export HADOOP_MAPRED_HOME=$HADOOP_HOME
export HADOOP_COMMON_HOME=$HADOOP_HOME
export HADOOP_HDFS_HOME=$HADOOP_HOME
export YARN_HOME=$HADOOP_HOME
export
HADOOP_COMMON_LIB_NATIVE_DIR=$HADOOP_HOME/lib/native
export HADOOP_OPTS="-Djava.library.path=$HADOOP_HOME/lib"
```

```
jeeva201028@ubuntu: ~/Downloads
hadoop-2.9.2/share/doc/hadoop/hadoop-kms/images/collapsed.gif
hadoop-2.9.2/share/doc/hadoop/hadoop-kms/images/bg.jpg
hadoop-2.9.2/share/doc/hadoop/hadoop-kms/dependency-analysis.html
jeeva201028@ubuntu:~/Downloads$ sudo mv hadoop-2.9.2 /usr/local/Hadoop
jeeva201028@ubuntu:~/Downloads$ sudo chown -R jeeva /usr/local
jeeva201028@ubuntu:~/Downloads$ sudo nano ~/.bashrc
jeeva201028@ubuntu:~/Downloads$ source ~/.bashrc
declare -x CLUTTER_IM_MODULE="ibus"
declare -x COLORTERM="truecolor"
declare -x DBUS_SESSION_BUS_ADDRESS="unix:path=/run/user/1000/bus"
declare -x DESKTOP_SESSION="ubuntu"
declare -x DISPLAY=":0"
declare -x GDMSESSION="ubuntu"
declare -x GJS_DEBUG_OUTPUT="stderr"
declare -x GJS_DEBUG_TOPICS="JS ERROR;JS LOG"
declare -x GNOME_DESKTOP_SESSION_ID="this-is-deprecated"
declare -x GNOME_SHELL_SESSION_MODE="ubuntu"
declare -x GNOME_TERMINAL_SCREEN="/org/gnome/Terminal/screen/c98c3178_c4bb_40a4_9647_3c0ea63ea95a"
declare -x GNOME_TERMINAL_SERVICE=":1.128"
declare -x GPG_AGENT_INFO="/run/user/1000/gnupg/S.gpg-agent:0:1"
declare -x GTK_IM_MODULE="ibus"
declare -x GTK_MODULES="gail:atk-bridge"
declare -x HADOOP_COMMON_HOME="/usr/local/hadoop"
declare -x HADOOP_HDFS_HOME="/usr/local/hadoop"
declare -x HADOOP_HOME="/usr/local/hadoop"
declare -x HADOOP_MAPRED_HOME="/usr/local/hadoop"
declare -x HOME="/home/jeeva201028"
declare -x IM_CONFIG_PHASE="1"
```

Edit the hadoop-env.sh file as following

```
jeeva201028@ubuntu: /usr/local/Hadoop/etc/hadoop
GNU nano 4.8 hadoop-env.sh
# limitations under the License.

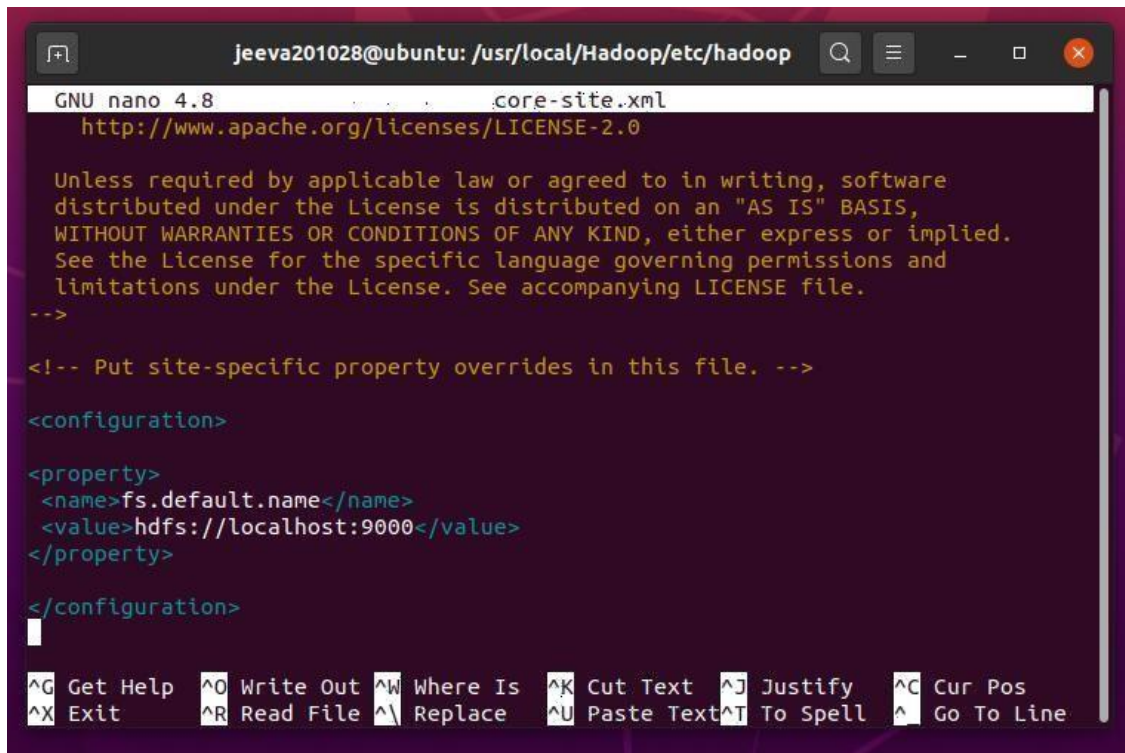
# Set Hadoop-specific environment variables here.

# The only required environment variable is JAVA_HOME. All others are
# optional. When running a distributed configuration it is best to
# set JAVA_HOME in this file, so that it is correctly defined on
# remote nodes.

# The java implementation to use.
#export JAVA_HOME=${JAVA_HOME}
export JAVA_HOME=/usr/lib/jvm/java-11-openjdk-amd64

# The jsvc implementation to use. Jsvc is required to run secure datanodes
# that bind to privileged ports to provide authentication of data transfer
# protocol. Jsvc is not required if SASL is configured for authentication of
# data transfer protocol using non-privileged ports.
#export JSVC_HOME=${JSVC_HOME}
[ line 28/121 (23%), col 52/52 (100%), char 1175/5024 (23%) ]
^G Get Help ^O Write Out ^W Where Is ^K Cut Text ^J Justify ^C Cur Pos
^X Exit ^R Read File ^\ Replace ^U Paste Text ^T To Spell ^_ Go To Line
```

## Edit the core-site.xml

A screenshot of a terminal window showing the nano text editor editing the file /usr/local/Hadoop/etc/hadoop/core-site.xml. The window title is 'jeeva201028@ubuntu: /usr/local/Hadoop/etc/hadoop'. The editor shows the following content:

```
GNU nano 4.8 /usr/local/Hadoop/etc/hadoop/core-site.xml
http://www.apache.org/licenses/LICENSE-2.0

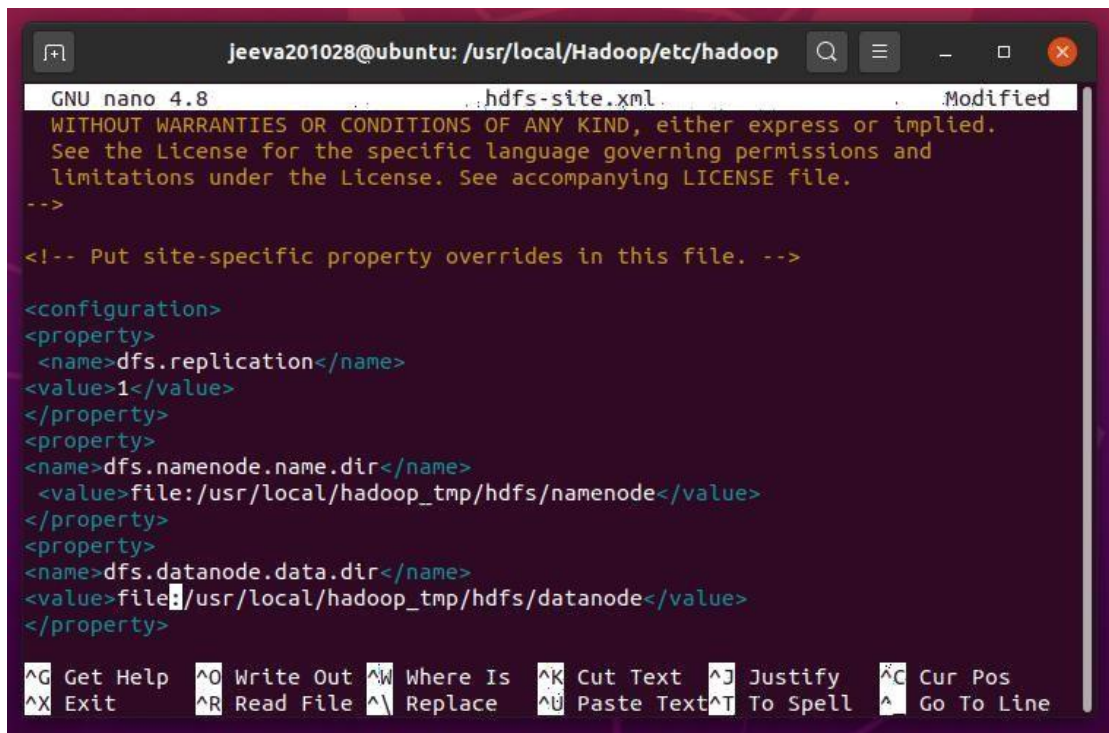
Unless required by applicable law or agreed to in writing, software
distributed under the License is distributed on an "AS IS" BASIS,
WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.
See the License for the specific language governing permissions and
limitations under the License. See accompanying LICENSE file.
-->

<!-- Put site-specific property overrides in this file. -->

<configuration>
<property>
  <name>fs.default.name</name>
  <value>hdfs://localhost:9000</value>
</property>
</configuration>
```

The bottom of the screen shows the nano editor's command palette with options like '^G Get Help', '^O Write Out', '^W Where Is', '^K Cut Text', '^J Justify', '^C Cur Pos', '^X Exit', '^R Read File', '^\_ Replace', '^U Paste Text', '^T To Spell', and '^\_ Go To Line'.

## Edit the hdfs-site.xml

A screenshot of a terminal window showing the nano text editor editing the file /usr/local/Hadoop/etc/hadoop/hdfs-site.xml. The window title is 'jeeva201028@ubuntu: /usr/local/Hadoop/etc/hadoop'. The editor shows the following content:

```
GNU nano 4.8 /usr/local/Hadoop/etc/hadoop/hdfs-site.xml Modified
WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.
See the License for the specific language governing permissions and
limitations under the License. See accompanying LICENSE file.
-->

<!-- Put site-specific property overrides in this file. -->

<configuration>
<property>
  <name>dfs.replication</name>
  <value>1</value>
</property>
<property>
  <name>dfs.namenode.name.dir</name>
  <value>file:/usr/local/hadoop_tmp/hdfs/namenode</value>
</property>
<property>
  <name>dfs.datanode.data.dir</name>
  <value>file:/usr/local/hadoop_tmp/hdfs/datanode</value>
</property>
```

The bottom of the screen shows the nano editor's command palette with options like '^G Get Help', '^O Write Out', '^W Where Is', '^K Cut Text', '^J Justify', '^C Cur Pos', '^X Exit', '^R Read File', '^\_ Replace', '^U Paste Text', '^T To Spell', and '^\_ Go To Line'.



```
jeeva201028@ubuntu: /usr/local/Hadoop/etc/hadoop
GNU nano 4.8 mapred-site.xml Modified
Licensed under the Apache License, Version 2.0 (the "License");
you may not use this file except in compliance with the License.
You may obtain a copy of the License at

http://www.apache.org/licenses/LICENSE-2.0

Unless required by applicable law or agreed to in writing, software
distributed under the License is distributed on an "AS IS" BASIS,
WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.
See the License for the specific language governing permissions and
limitations under the License. See accompanying LICENSE file.
-->

<!-- Put site-specific property overrides in this file. -->

<configuration>
<property>
    <name>mapreduce.framework.name</name>
    <value>yarn</value>
</property>

```

^G Get Help   ^O Write Out   ^W Where Is   ^K Cut Text   ^J Justify   ^C Cur Pos  
^X Exit   ^R Read File   ^\ Replace   ^U Paste Text   ^T To Spell   ^\_ Go To Line

Create the following directory

```
jeeva201028@ubuntu: ~/Desktop
jeeva201028@ubuntu: /usr/local/Hadoop/etc/hadoop$
jeeva201028@ubuntu: /usr/local/Hadoop/etc/hadoop$
jeeva201028@ubuntu: /usr/local/Hadoop/etc/hadoop$
jeeva201028@ubuntu: /usr/local/Hadoop/etc/hadoop$
jeeva201028@ubuntu: /usr/local/Hadoop/etc/hadoop$
jeeva201028@ubuntu: /usr/local/Hadoop/etc/hadoop$
jeeva201028@ubuntu: /usr/local/Hadoop/etc/hadoop$
jeeva201028@ubuntu: /usr/local/Hadoop/etc/hadoop$
jeeva201028@ubuntu: /usr/local/Hadoop/etc/hadoop$
jeeva201028@ubuntu: /usr/local/Hadoop/etc/hadoop$
jeeva201028@ubuntu: /usr/local/Hadoop/etc/hadoop$
jeeva201028@ubuntu: /usr/local/Hadoop/etc/hadoop$ sudo nano mapred-site.xml
jeeva201028@ubuntu: /usr/local/Hadoop/etc/hadoop$ sudo cp mapred-site.xml.templat
e mapred-site.xml
jeeva201028@ubuntu: /usr/local/Hadoop/etc/hadoop$ sudo nano mapred-site.xml
jeeva201028@ubuntu: /usr/local/Hadoop/etc/hadoop$ cd
jeeva201028@ubuntu: ~$ cd Desktop/
jeeva201028@ubuntu: ~/Desktop$ sudo mkdir -p /usr/local/hadoop_tmp
jeeva201028@ubuntu: ~/Desktop$ sudo mkdir -p /usr/local/hadoop_tmp/hdfs/namenode
jeeva201028@ubuntu: ~/Desktop$ sudo mkdir -p /usr/local/hadoop_tmp/hdfs/datanode
jeeva201028@ubuntu: ~/Desktop$ sudo chown -R jeeva /usr/local/hadoop_tmp
jeeva201028@ubuntu: ~/Desktop$
```

## Format hdfs namenode

```
jeeva201028@ubuntu: ~  
jeeva201028@ubuntu:~$ hdfs  
Usage: hdfs [--config confdir] [--loglevel loglevel] COMMAND  
       where COMMAND is one of:  
    dfs                run a filesystem command on the file systems supported in Hadoop.  
    classpath          prints the classpath  
    namenode -format   format the DFS filesystem  
    secondarynamenode run the DFS secondary namenode  
    namenode           run the DFS namenode  
    journalnode        run the DFS journalnode  
    zkfc               run the ZK Failover Controller daemon  
    datanode           run a DFS datanode  
    debug              run a Debug Admin to execute HDFS debug commands  
    dfsadmin           run a DFS admin client  
    dfsrouter          run the DFS router  
    dfsrouteradmin     manage Router-based federation  
    haadmin            run a DFS HA admin client  
    fsck               run a DFS filesystem checking utility  
    balancer           run a cluster balancing utility  
    jmxget             get JMX exported values from NameNode or DataNode.  
    mover              run a utility to move block replicas across  
                       storage types  
    oiv                apply the offline fsimage viewer to an fsimage  
    oiv_legacy         apply the offline fsimage viewer to an legacy fsimage  
    oev               apply the offline edits viewer to an edits file  
    fetchdt           fetch a delegation token from the NameNode  
    getconf           get config values from configuration
```

```
balancer           run a cluster balancing utility  
jmxget             get JMX exported values from NameNode or DataNode.  
mover              run a utility to move block replicas across  
                  storage types  
oiv                apply the offline fsimage viewer to an fsimage  
oiv_legacy         apply the offline fsimage viewer to an legacy fsimage  
oev               apply the offline edits viewer to an edits file  
fetchdt           fetch a delegation token from the NameNode  
getconf           get config values from configuration  
groups            get the groups which users belong to  
snapshotdiff      diff two snapshots of a directory or diff the  
                  current directory contents with a snapshot  
lsSnapshottableDir list all snapshottable dirs owned by the current user  
                  use -help to see options  
portmap           run a portmap service  
nfs3              run an NFS version 3 gateway  
cacheadmin        configure the HDFS cache  
crypto            configure HDFS encryption zones  
storagepolicies   list/get/set block storage policies  
version           print the version  
Most commands print help when invoked w/o parameters.  
jeeva201028@ubuntu:~$ hdfs namenode -format
```

```
jeeva201028@ubuntu: ~  
20/11/08 04:19:35 INFO namenode.FSNamesystem: Retry cache will use 0.03 of total  
heap and retry cache entry expiry time is 600000 millis  
20/11/08 04:19:35 INFO util.GSet: Computing capacity for map NameNodeRetryCache  
20/11/08 04:19:35 INFO util.GSet: VM type = 64-bit  
20/11/08 04:19:35 INFO util.GSet: 0.029999999329447746% max memory 1000 MB = 307  
.2 KB  
20/11/08 04:19:35 INFO util.GSet: capacity = 2^15 = 32768 entries  
20/11/08 04:19:35 INFO namenode.FSImage: Allocated new BlockPoolId: BP-725375456  
-127.0.1.1-1604837975223  
20/11/08 04:19:35 INFO common.Storage: Storage directory /usr/local/hadoop_tmp/h  
dfs/namenode has been successfully formatted.  
20/11/08 04:19:35 INFO namenode.FSImageFormatProtobuf: Saving image file /usr/lo  
cal/hadoop_tmp/hdfs/namenode/current/fsimage.ckpt_000000000000000000 using no c  
ompression  
20/11/08 04:19:35 INFO namenode.FSImageFormatProtobuf: Image file /usr/local/had  
oop_tmp/hdfs/namenode/current/fsimage.ckpt_000000000000000000 of size 329 bytes  
saved in 0 seconds .  
20/11/08 04:19:35 INFO namenode.NNStorageRetentionManager: Going to retain 1 ima  
ges with txid >= 0  
20/11/08 04:19:35 INFO namenode.NameNode: SHUTDOWN_MSG:  
/*****  
SHUTDOWN_MSG: Shutting down NameNode at ubuntu/127.0.1.1  
*****/  
jeeva201028@ubuntu:~$
```

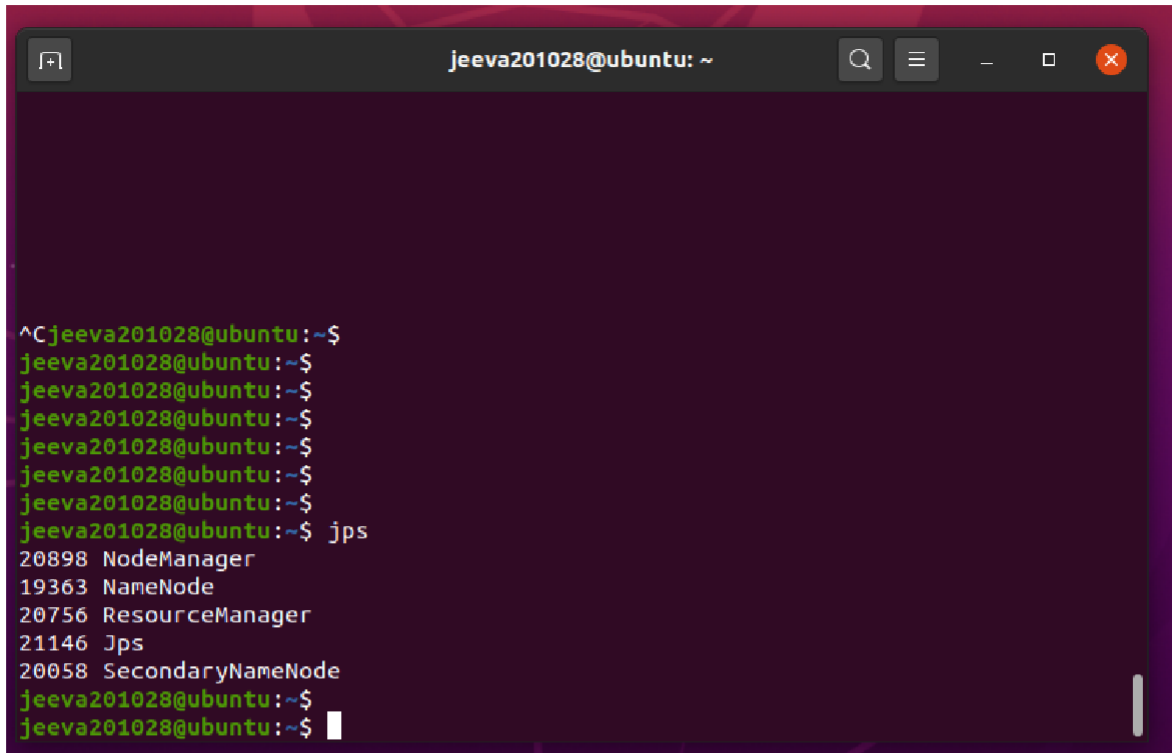
```
$ start-dfs.sh
```

```
$ start-yarn.sh
```

To check whether hadoop is correctly installed or not: \$ jps

```
jeeva201028@ubuntu: ~  
doop_tmp/hdfs/namenode/current/fsimage.ckpt_000000000000000000 of size 330 bytes saved in 0 seconds .  
20/11/08 12:04:47 INFO namenode.NNStorageRetentionManager: Going to retain 1 images with txid >= 0  
20/11/08 12:04:47 INFO namenode.NameNode: SHUTDOWN_MSG:  
/*****  
SHUTDOWN_MSG: Shutting down NameNode at ubuntu/127.0.1.1  
*****/  
jeeva201028@ubuntu:~/usr/local/hadoop$ cd  
jeeva201028@ubuntu:~$ start-dfs.sh  
WARNING: An illegal reflective access operation has occurred  
WARNING: Illegal reflective access by org.apache.hadoop.security.authentication.util.KerberosUtil (file:/usr/local/hadoop/share/hadoop/common/lib/hadoop-auth-2.9.2.jar) to method sun.security.krb5.Config.getInstance()  
WARNING: Please consider reporting this to the maintainers of org.apache.hadoop.security.authentication.util.KerberosUtil  
WARNING: Use --illegal-access=warn to enable warnings of further illegal reflective access operations  
WARNING: All illegal access operations will be denied in a future release  
20/11/08 12:06:10 WARN util.NativeCodeLoader: Unable to load native-hadoop library for your platform... using builtin-java classes where applicable  
Starting namenodes on [localhost]  
jeeva201028@localhost's password: █
```

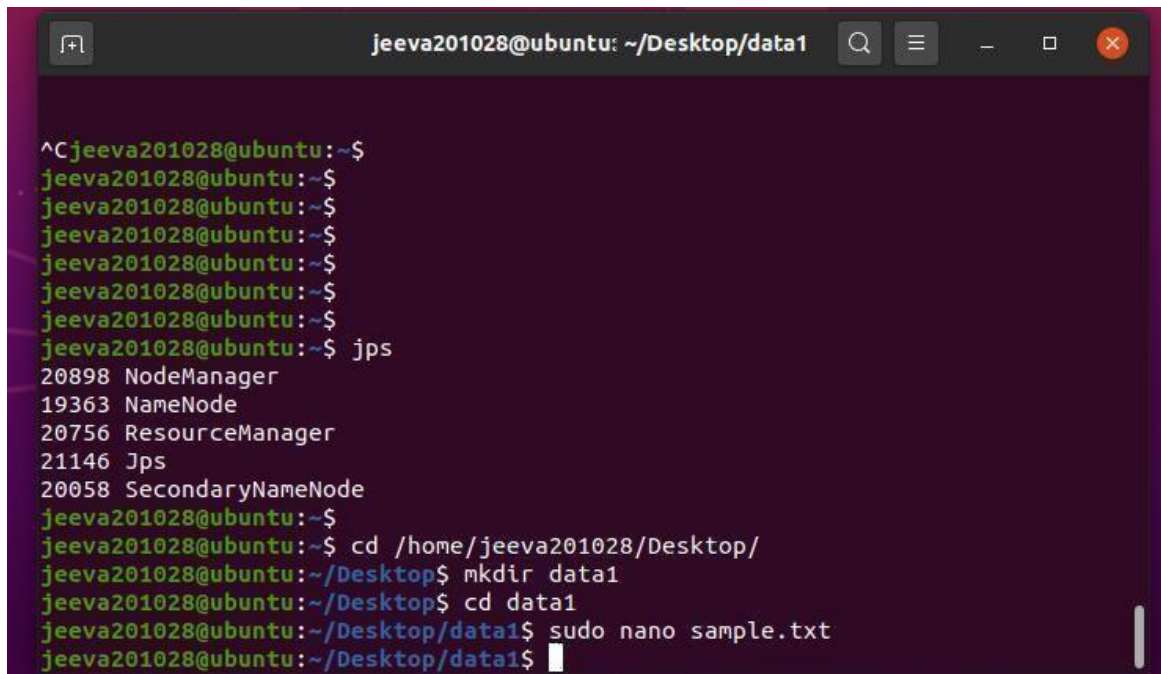
```
jeeva201028@ubuntu: ~  
028-datanode-ubuntu.out  
Starting secondary namenodes [0.0.0.0]  
jeeva201028@0.0.0.0's password:  
0.0.0.0: secondarynamenode running as process 20058. Stop it first.  
WARNING: An illegal reflective access operation has occurred  
WARNING: Illegal reflective access by org.apache.hadoop.security.authentication.util.KerberosUtil (file:/usr/local/hadoop/share/hadoop/common/lib/hadoop-auth-2.9.2.jar) to method sun.security.krb5.Config.getInstance()  
WARNING: Please consider reporting this to the maintainers of org.apache.hadoop.security.authentication.util.KerberosUtil  
WARNING: Use --illegal-access=warn to enable warnings of further illegal reflective access operations  
WARNING: All illegal access operations will be denied in a future release  
20/11/08 12:09:32 WARN util.NativeCodeLoader: Unable to load native-hadoop library for your platform... using builtin-java classes where applicable  
jeeva201028@ubuntu:~$ start-yarn.sh  
starting yarn daemons  
starting resourcemanager, logging to /usr/local/hadoop/logs/yarn-jeeva201028-resourcemanager-ubuntu.out  
jeeva201028@localhost's password:  
localhost: starting nodemanager, logging to /usr/local/hadoop/logs/yarn-jeeva201028-nodemanager-ubuntu.out  
jeeva201028@ubuntu:~$ █
```



```
^Cjeeva201028@ubuntu:~$
jeeva201028@ubuntu:~$
jeeva201028@ubuntu:~$
jeeva201028@ubuntu:~$
jeeva201028@ubuntu:~$
jeeva201028@ubuntu:~$
jeeva201028@ubuntu:~$
jeeva201028@ubuntu:~$ jps
20898 NodeManager
19363 NameNode
20756 ResourceManager
21146 Jps
20058 SecondaryNameNode
jeeva201028@ubuntu:~$
jeeva201028@ubuntu:~$
```

To execute word count program:

Create a directory and add a text file in it named “sample.txt”



```
^Cjeeva201028@ubuntu:~$
jeeva201028@ubuntu:~$
jeeva201028@ubuntu:~$
jeeva201028@ubuntu:~$
jeeva201028@ubuntu:~$
jeeva201028@ubuntu:~$
jeeva201028@ubuntu:~$ jps
20898 NodeManager
19363 NameNode
20756 ResourceManager
21146 Jps
20058 SecondaryNameNode
jeeva201028@ubuntu:~$
jeeva201028@ubuntu:~$ cd /home/jeeva201028/Desktop/
jeeva201028@ubuntu:~/Desktop$ mkdir data1
jeeva201028@ubuntu:~/Desktop$ cd data1
jeeva201028@ubuntu:~/Desktop/data1$ sudo nano sample.txt
jeeva201028@ubuntu:~/Desktop/data1$
```



Input: 'sample.txt' file

```
hello
how are you
communication
components
computational
computer
computing
coordinate
compiler
distributed file system
hadoop
single node cluster
word count
```

AG Get Help   AO Write Out   AW Where Is   AK Cut Text   AJ Justify   AC Cur Pos   MU Undo   MA Mark Text   M-] To Bracket  
AX Exit   AR Read File   AL Replace   AU Paste Text   AT To Spell   A Go To Line   ME Redo   M-A Copy Text   M-? Where Was

Now run bin/hdfs dfs -mkdir /user

```
jeeva201028@ubuntu: /usr/local/hadoop
```

```
20058 SecondaryNameNode
jeeva201028@ubuntu:~$
jeeva201028@ubuntu:~$ cd /home/jeeva201028/Desktop/
jeeva201028@ubuntu:~/Desktop$ mkdir data1
jeeva201028@ubuntu:~/Desktop$ cd data1
jeeva201028@ubuntu:~/Desktop/data1$ sudo nano sample.txt
jeeva201028@ubuntu:~/Desktop/data1$ cd
jeeva201028@ubuntu:~$ cd /usr/local/hadoop
jeeva201028@ubuntu:/usr/local/hadoop$ bin/hdfs dfs -mkdir /user
WARNING: An illegal reflective access operation has occurred
WARNING: Illegal reflective access by org.apache.hadoop.security.authentication.util.KerberosUtil (file:/usr/local/hadoop/share/hadoop/common/lib/hadoop-auth-2.9.2.jar) to method sun.security.krb5.Config.getInstance()
WARNING: Please consider reporting this to the maintainers of org.apache.hadoop.security.authentication.util.KerberosUtil
WARNING: Use --illegal-access=warn to enable warnings of further illegal reflective access operations
WARNING: All illegal access operations will be denied in a future release
20/11/08 12:14:30 WARN util.NativeCodeLoader: Unable to load native-hadoop library for your platform... using builtin-java classes where applicable
```

Then run `bin/hdfs dfs -mkdir /ypm`

```
Jeeva201028@ubuntu: /usr/local/hadoop
file:/usr/local/hadoop/share/hadoop/common/lib/hadoop-auth-2.9.2.jar) to method sun.security.krb5.Con
fig.getInstance()
WARNING: Please consider reporting this to the maintainers of org.apache.hadoop.security.authenticat
ion.util.KerberosUtil
WARNING: Use --illegal-access=warn to enable warnings of further illegal reflective access operatio
n
WARNING: All illegal access operations will be denied in a future release
20/11/08 12:14:30 WARN util.NativeCodeLoader: Unable to load native-hadoop library for your platform
... using builtin-java classes where applicable
Jeeva201028@ubuntu: /usr/local/hadoop$ bin/hdfs dfs -mkdir /ypm
WARNING: An illegal reflective access operation has occurred
WARNING: Illegal reflective access by org.apache.hadoop.security.authentication.util.KerberosUtil (f
ile:/usr/local/hadoop/share/hadoop/common/lib/hadoop-auth-2.9.2.jar) to method sun.security.krb5.Con
fig.getInstance()
WARNING: Please consider reporting this to the maintainers of org.apache.hadoop.security.authenticat
ion.util.KerberosUtil
WARNING: Use --illegal-access=warn to enable warnings of further illegal reflective access operatio
n
WARNING: All illegal access operations will be denied in a future release
20/11/08 12:15:49 WARN util.NativeCodeLoader: Unable to load native-hadoop library for your platform
... using builtin-java classes where applicable
Jeeva201028@ubuntu: /usr/local/hadoop$
```

Run the program //word count prgm will be in the jar file by default which we are using now.

`$ bin/hadoop jar share/hadoop/mapreduce/hadoop-mapreduce-examples-2.9.2.jar wordcount /user/input output`

```
Jeeva201028@ubuntu: /usr/local/hadoop
Jeeva201028@ubuntu: /usr/local/hadoop$
Jeeva201028@ubuntu: /usr/local/hadoop$
Jeeva201028@ubuntu: /usr/local/hadoop$
Jeeva201028@ubuntu: /usr/local/hadoop$
Jeeva201028@ubuntu: /usr/local/hadoop$
Jeeva201028@ubuntu: /usr/local/hadoop$
Jeeva201028@ubuntu: /usr/local/hadoop$ bin/hadoop jar share/hadoop/mapreduce/hadoop-mapreduce-exampl
e-2.9.2.jar wordcount /user/input output
WARNING: An illegal reflective access operation has occurred
WARNING: Illegal reflective access by org.apache.hadoop.security.authentication.util.KerberosUtil (f
ile:/usr/local/hadoop/share/hadoop/common/lib/hadoop-auth-2.9.2.jar) to method sun.security.krb5.Con
fig.getInstance()
WARNING: Please consider reporting this to the maintainers of org.apache.hadoop.security.authenticat
ion.util.KerberosUtil
WARNING: Use --illegal-access=warn to enable warnings of further illegal reflective access operatio
n
WARNING: All illegal access operations will be denied in a future release
20/11/08 12:24:53 WARN util.NativeCodeLoader: Unable to load native-hadoop library for your platform
... using builtin-java classes where applicable
20/11/08 12:24:54 INFO client.RMProxy: Connecting to ResourceManager at /0.0.0.0:8032
```

```
Jeeva201028@ubuntu: /usr/local/hadoop
Jeeva201028@ubuntu: /usr/local/hadoop$
Jeeva201028@ubuntu: /usr/local/hadoop$
Jeeva201028@ubuntu: /usr/local/hadoop$
Jeeva201028@ubuntu: /usr/local/hadoop$
Jeeva201028@ubuntu: /usr/local/hadoop$
Jeeva201028@ubuntu: /usr/local/hadoop$
Jeeva201028@ubuntu: /usr/local/hadoop$
Jeeva201028@ubuntu: /usr/local/hadoop$
Jeeva201028@ubuntu: /usr/local/hadoop$ bin/hdfs dfs -put /home/jeeva/Desktop/data/user/input
WARNING: An illegal reflective access operation has occurred
WARNING: Illegal reflective access by org.apache.hadoop.security.authentication.util.KerberosUtil (f
ile:/usr/local/hadoop/share/hadoop/common/lib/hadoop-auth-2.9.2.jar) to method sun.security.krb5.Con
fig.getInstance()
WARNING: Please consider reporting this to the maintainers of org.apache.hadoop.security.authenticat
ion.util.KerberosUtil
WARNING: Use --illegal-access=warn to enable warnings of further illegal reflective access operatio
n
WARNING: All illegal access operations will be denied in a future release
20/11/08 12:37:51 WARN util.NativeCodeLoader: Unable to load native-hadoop library for your platform
... using builtin-java classes where applicable
```

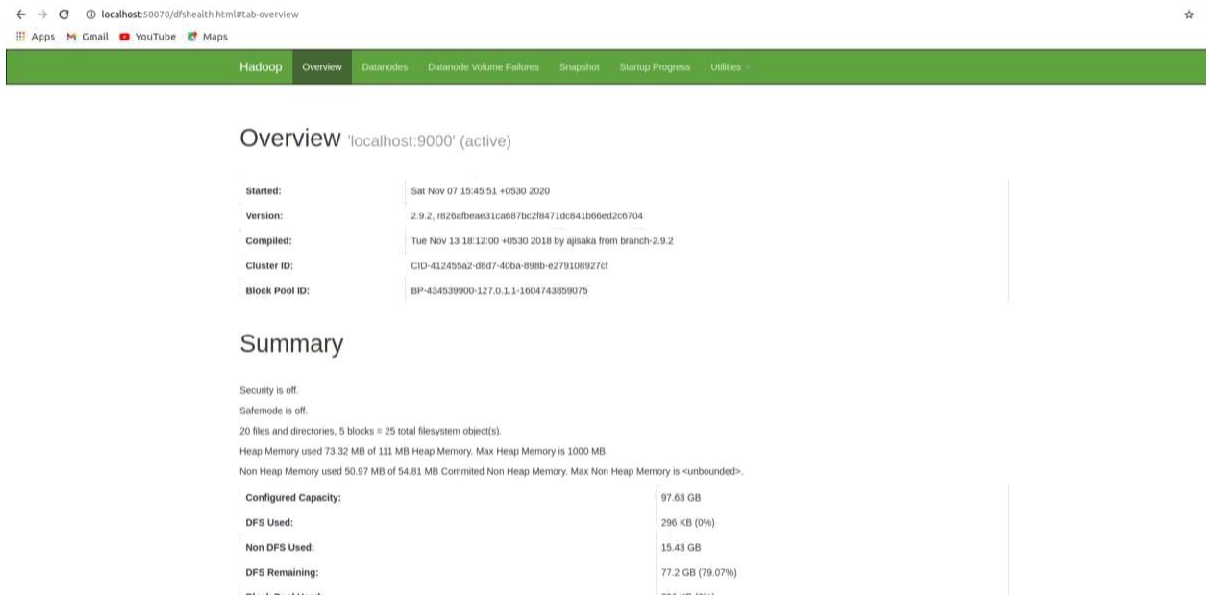
```
jeeva201028@ubuntu: ~  
Total megabyte-millisecond taken by all reduce tasks=2942976  
Map-Reduce Framework  
  Map input records=16  
  Map output records=23  
  Map output bytes=274  
  Map output materialized bytes=285  
  Input split bytes=108  
  Combine input records=23  
  Combine output records=20  
  Reduce input groups=20  
  Reduce shuffle bytes=285  
  Reduce input records=20  
  Reduce output records=20  
  Spilled Records=40  
  Shuffled Maps =1  
  Failed Shuffles=0  
  Merged Map outputs=1  
  GC time elapsed (ms)=62  
  CPU time spent (ms)=1510  
  Physical memory (bytes) snapshot=455553024  
  Virtual memory (bytes) snapshot=4149194752  
  Total committed heap usage (bytes)=273678336  
Shuffle Errors  
  BAD_IO=0  
  CONNECTION=0  
  IO_ERROR=0  
  WRONG_LENGTH=0  
  WRONG_MAP=0  
  WRONG_REDUCE=0  
File Input Format Counters  
  Bytes Read=186  
File Output Format Counters  
  Bytes Written=192
```

Now run the following command to see the output

```
$ bin/hdfs dfs -cat output/*
```

```
jeeva201028@ubuntu: /usr/local/hadoop$  
jeeva201028@ubuntu: /usr/local/hadoop$  
jeeva201028@ubuntu: /usr/local/hadoop$  
jeeva201028@ubuntu: /usr/local/hadoop$  
jeeva201028@ubuntu: /usr/local/hadoop$ bin/hdfs dfs -cat output/*  
WARNING: An illegal reflective access operation has occurred  
WARNING: Illegal reflective access by org.apache.hadoop.security.authentication.util.KerberosUtil (file:/usr/local/hadoop/share/hadoop/common/lib/hadoop-auth-2.9.2.jar) to method sun.security.krb5.Config.getInstance()  
WARNING: Please consider reporting this to the maintainers of org.apache.hadoop.security.authentication.util.KerberosUtil  
WARNING: Use --illegal-access=warn to enable warnings of further illegal reflective access operations  
WARNING: All illegal access operations will be denied in a future release  
26/11/07 17:25:31 WARN util.NativeCodeLoader: Unable to load native-hadoop library for your platform... using builtin-java classes where applicable  
are 1  
cluster 1  
communication 1  
compiler 1  
components 1  
computational 1  
computer 2  
computing 1  
coordinate 1  
count 1  
distributed 1  
file 1  
hadoop 3  
hello 1  
how 1  
node 1  
single 1  
system 1  
word 1  
you 1
```

Using gui: In browser open port 50700, <http://localhost/50070> (Hadoop must be running)

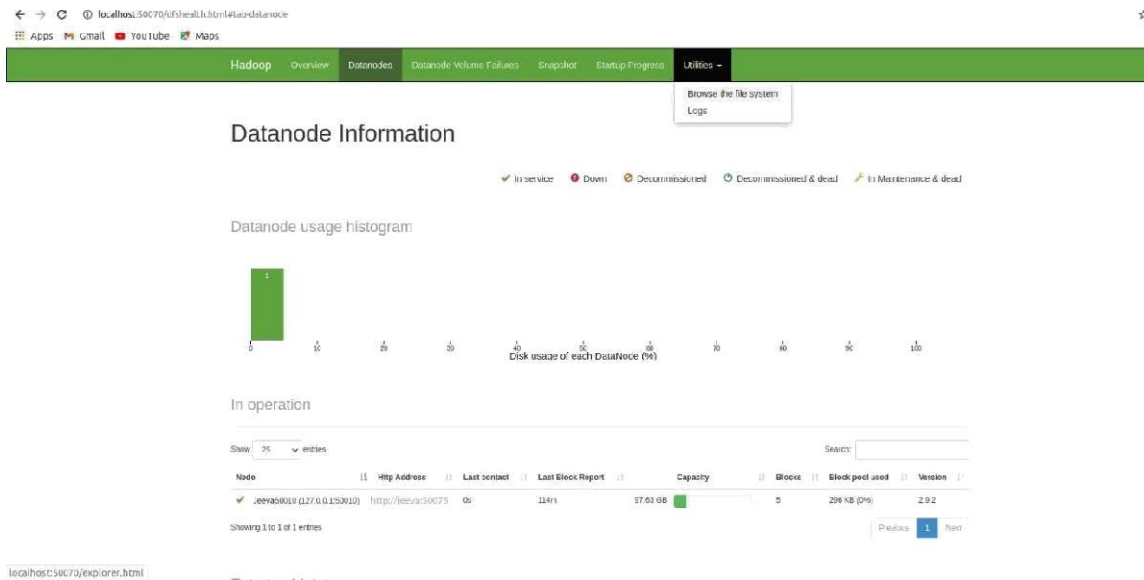


The screenshot shows the Hadoop Overview page in a browser. The address bar shows `localhost:50070/dfshealth.html#tab-overview`. The navigation bar includes `Hadoop`, `Overview`, `Datanodes`, `Datanode Volume Failures`, `Snapshot`, `Startup Progress`, and `Utilities`. The main content area is titled `Overview 'localhost:9000' (active)`. It displays the following information:

<b>Started:</b>	Sat Nov 07 15:45:51 +0530 2020
<b>Version:</b>	2.9.2, f820afbea031ca687bc2f8471dc841b66e2c0704
<b>Compiled:</b>	Tue Nov 13 18:12:00 +0530 2018 by ajsaka from branch-2.9.2
<b>Cluster ID:</b>	C10-412455a2-d8d7-4c6a-898b-e279108927c1
<b>Block Pool ID:</b>	BP-434538900-127.0.1.1-1004743359075

Below the overview is a **Summary** section. It states: `Security is off.`, `Safemode is off.`, `20 files and directories, 5 blocks = 25 total filesystem object(s).`, `Heap Memory used 73.32 MB of 111 MB Heap Memory. Max Heap Memory is 1000 MB`, and `Non Heap Memory used 50.67 MB of 54.61 MB Committed Non Heap Memory. Max Non Heap Memory is <unbounded>.`

<b>Configured Capacity:</b>	97.63 GB
<b>DFS Used:</b>	296 KB (0%)
<b>Non DFS Used:</b>	15.43 GB
<b>DFS Remaining:</b>	77.2 GB (79.07%)



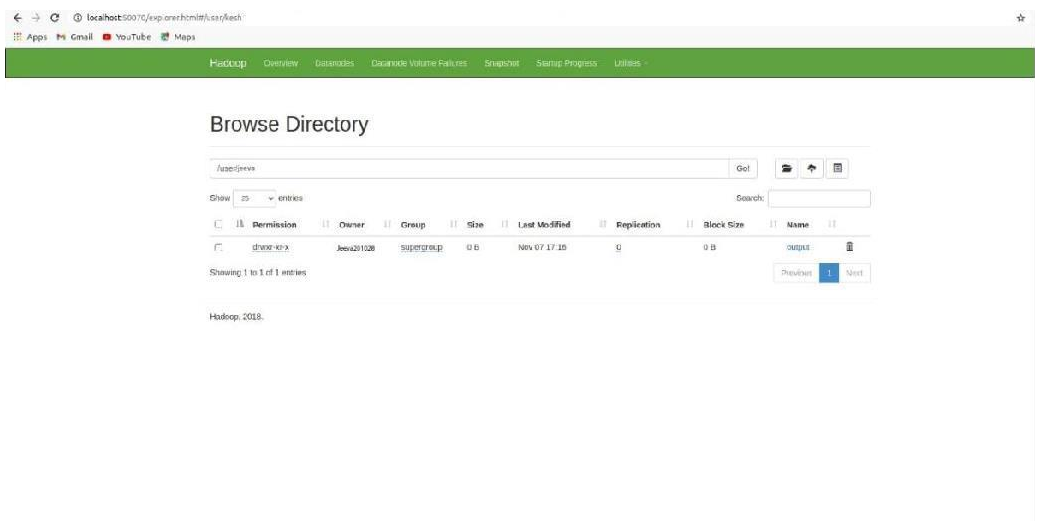
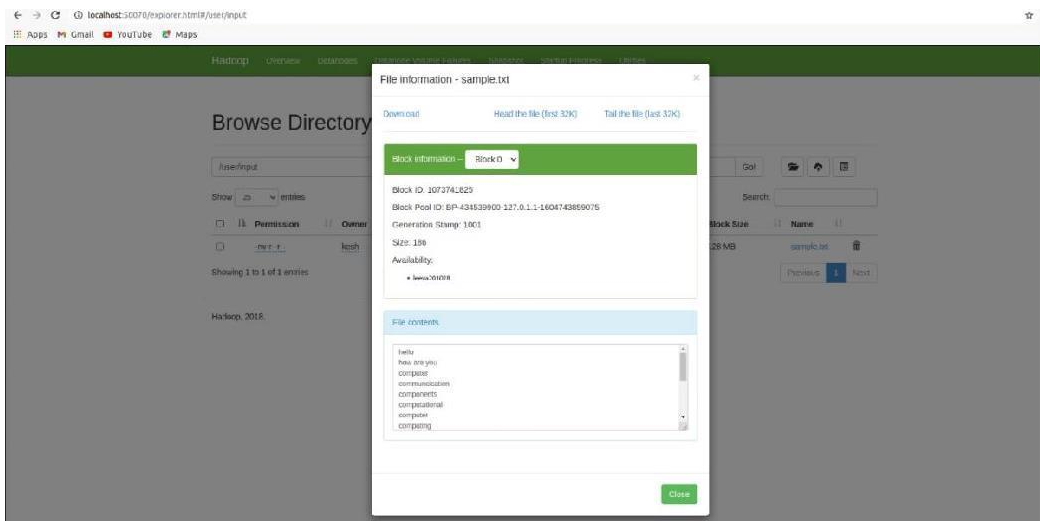
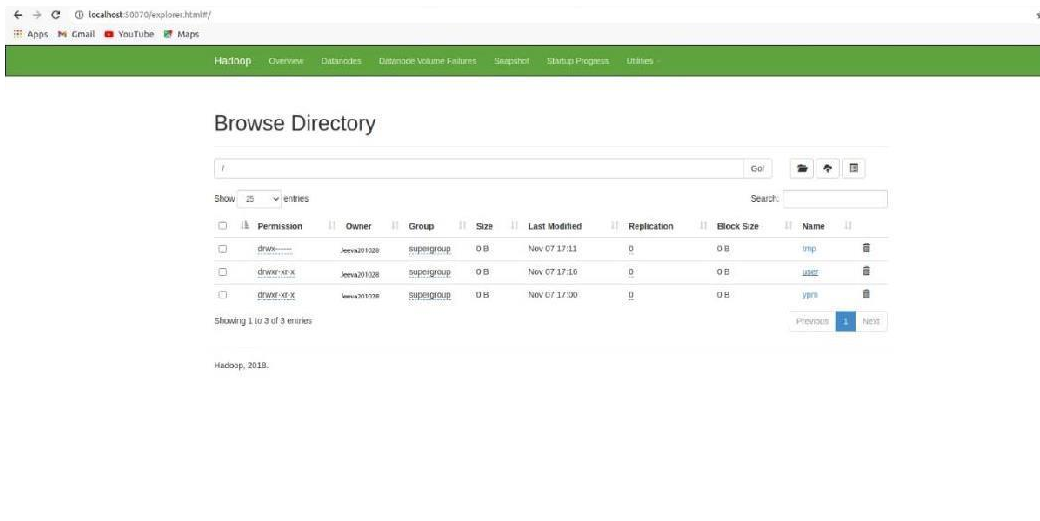
The screenshot shows the Hadoop Datanode Information page. The address bar shows `localhost:50070/dfshealth.html#tab-datanode`. The navigation bar includes `Hadoop`, `Overview`, `Datanodes`, `Datanode Volume Failures`, `Snapshot`, `Startup Progress`, and `Utilities`. The main content area is titled **Datanode Information**. It features a legend for node status: `✓ In service`, `● Down`, `⊗ Decommissioned`, `⊖ Decommissioned & dead`, and `⚠ Maintenance & dead`. Below the legend is a **Datanode usage histogram** showing a single bar for disk usage of each Datanode (0% to 100%).

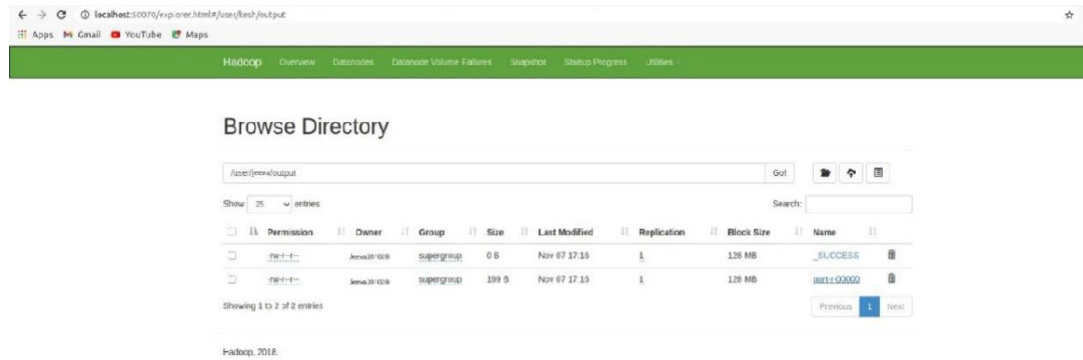
Under the histogram is the **In operation** section, which displays a table of datanodes. The table has columns for `Node`, `Http Address`, `Last contact`, `Last Block Report`, `Capacity`, `Blocks`, `Block pool used`, and `Version`.

Node	Http Address	Last contact	Last Block Report	Capacity	Blocks	Block pool used	Version
✓ jeeva30019 (127.0.0.1253019)	<a href="http://jeeva30075">http://jeeva30075</a>	0s	1181s	97.63 GB	0	296 KB (0%)	2.9.2

At the bottom of the table, it says `Showing 1 to 1 of 1 entries` and includes `Previous` and `Next` buttons.

`localhost:50070/exp/over.html`





After completing stop hadoop \$ stop-all.sh

```
jeeva201028@ubuntu: /usr/local/hadoop
jeeva201028@ubuntu: /usr/local/hadoop$
jeeva201028@ubuntu: /usr/local/hadoop$
jeeva201028@ubuntu: /usr/local/hadoop$
jeeva201028@ubuntu: /usr/local/hadoop$
jeeva201028@ubuntu: /usr/local/hadoop$
jeeva201028@ubuntu: /usr/local/hadoop$
jeeva201028@ubuntu: /usr/local/hadoop$
jeeva201028@ubuntu: /usr/local/hadoop$
jeeva201028@ubuntu: /usr/local/hadoop$
jeeva201028@ubuntu: /usr/local/hadoop$ stop-all.sh
This script is Deprecated. Instead use stop-dfs.sh and stop-yarn.sh
WARNING: An illegal reflective access operation has occurred
WARNING: Illegal reflective access by org.apache.hadoop.security.authentication.
util.KerberosUtil (file:/usr/local/hadoop/share/hadoop/common/lib/hadoop-auth-2.
9.2.jar) to method sun.security.krb5.Config.getInstance()
WARNING: Please consider reporting this to the maintainers of org.apache.hadoop.
security.authentication.util.KerberosUtil
WARNING: Use --illegal-access=warn to enable warnings of further illegal reflect
ive access operations
WARNING: All illegal access operations will be denied in a future release
20/11/09 00:21:44 WARN util.NativeCodeLoader: Unable to load native-hadoop libra
ry for your platform... using builtin-java classes where applicable
Stopping namenodes on [localhost]
jeeva201028@localhost's password: █
```

**Result:**

Thus the installation of hadoop single mode cluster and execution of word count program is done and the output is obtained successfully.

## EX.No:6

# Simulate a cloud scenario using CloudSim and run a scheduling algorithm

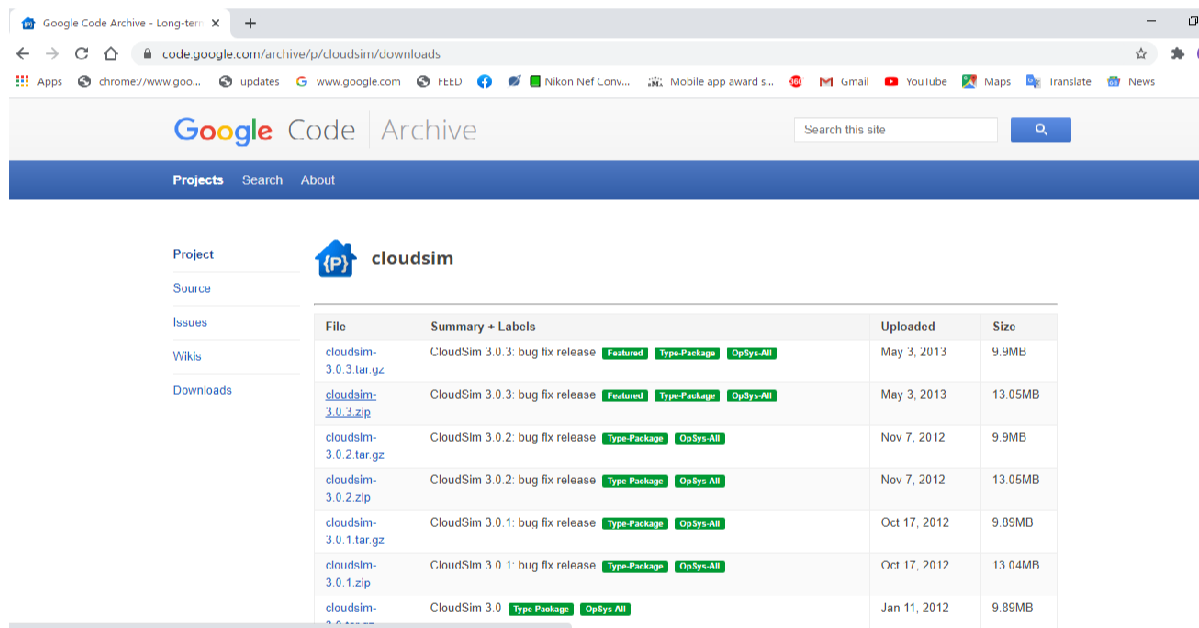
### Aim:

To simulate a cloud scenario using CloudSim and run a scheduling algorithm that is not present in CloudSim.

### Procedure:


Step 1: Download CloudSim installable files from:

<https://code.google.com/p/cloudsim/downloads/list> and unzip the download.

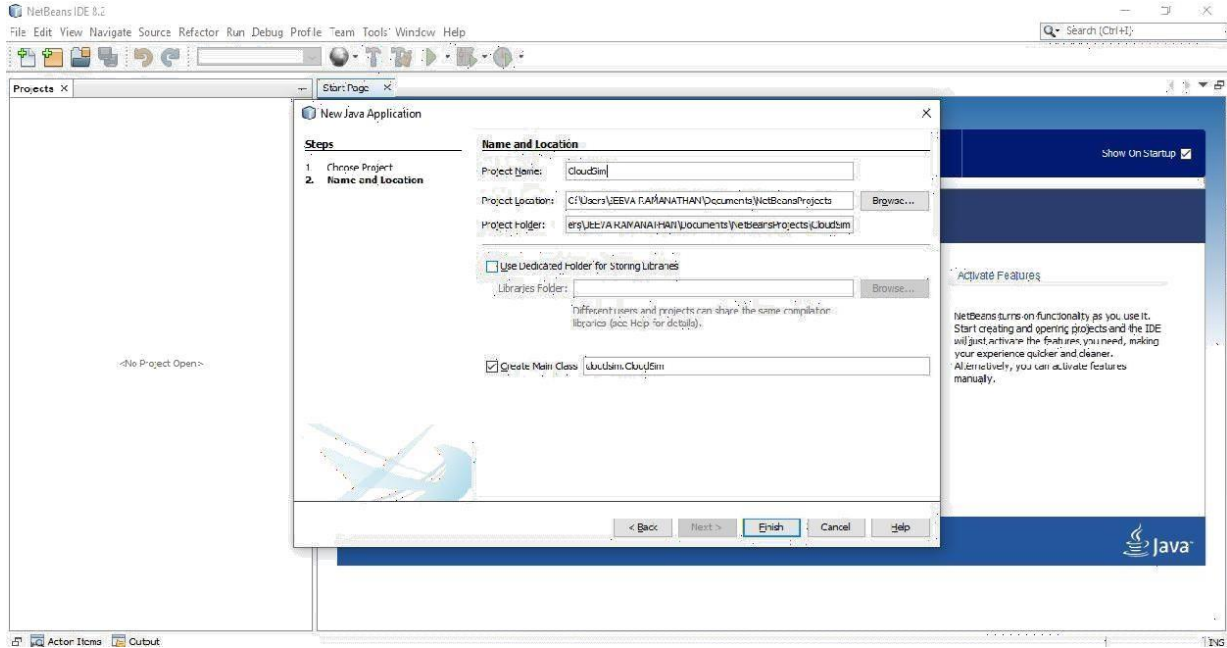
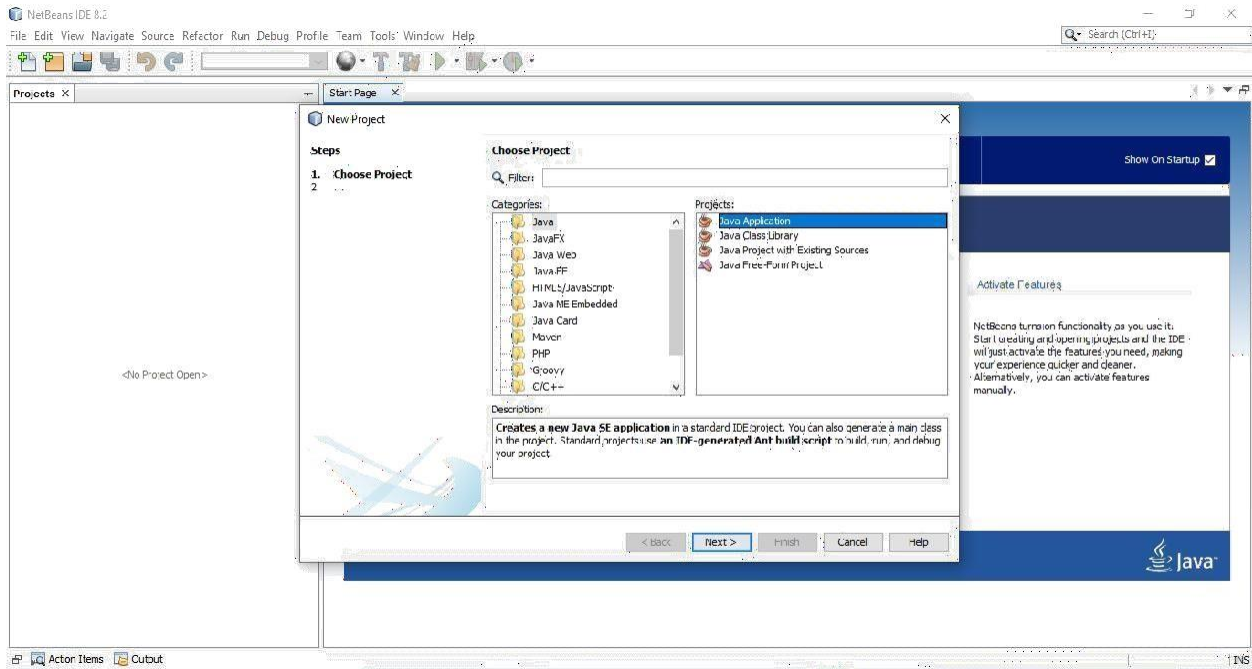


The screenshot shows the Google Code Archive page for the CloudSim project. The page lists several download files with their respective summaries, upload dates, and sizes. The files are organized into a table with columns for File, Summary + Labels, Uploaded, and Size.

File	Summary + Labels	Uploaded	Size
cloudsim-3.0.3.tar.gz	CloudSim 3.0.3: bug fix release <span>Featured</span> <span>Type:Package</span> <span>OpSys:All</span>	May 3, 2013	9.9MB
cloudsim-3.0.3.zip	CloudSim 3.0.3: bug fix release <span>Featured</span> <span>Type:Package</span> <span>OpSys:All</span>	May 3, 2013	13.05MB
cloudsim-3.0.2.tar.gz	CloudSim 3.0.2: bug fix release <span>Type:Package</span> <span>OpSys:All</span>	Nov 7, 2012	9.9MB
cloudsim-3.0.2.zip	CloudSim 3.0.2: bug fix release <span>Type:Package</span> <span>OpSys:All</span>	Nov 7, 2012	13.05MB
cloudsim-3.0.1.tar.gz	CloudSim 3.0.1: bug fix release <span>Type:Package</span> <span>OpSys:All</span>	Oct 17, 2012	9.95MB
cloudsim-3.0.1.zip	CloudSim 3.0.1: bug fix release <span>Type:Package</span> <span>OpSys:All</span>	Oct 17, 2012	13.04MB
cloudsim-3.0.0.tar.gz	CloudSim 3.0.0 <span>Type:Package</span> <span>OpSys:All</span>	Jan 11, 2012	9.85MB

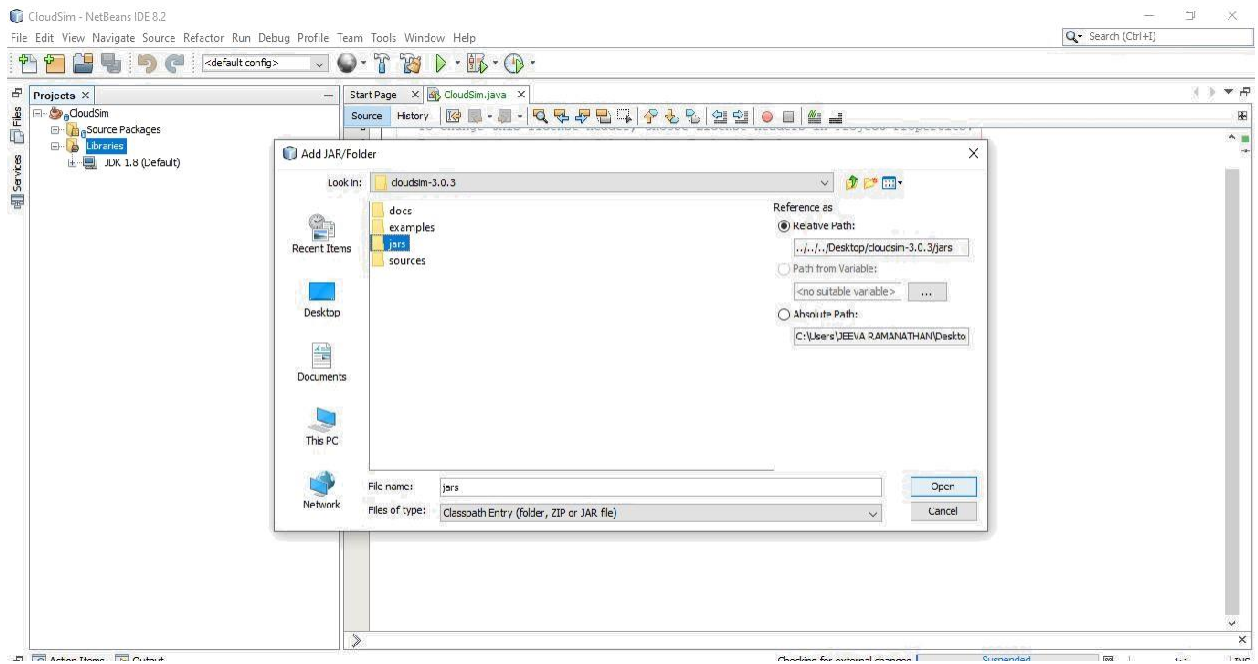
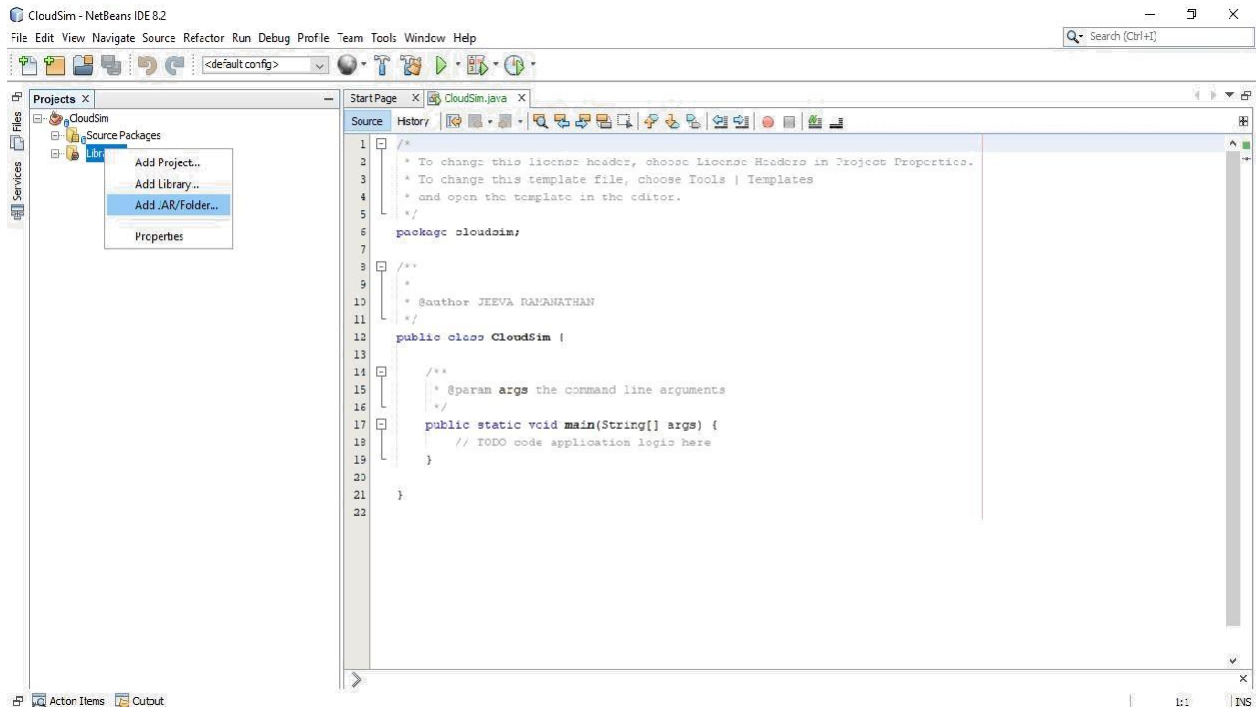


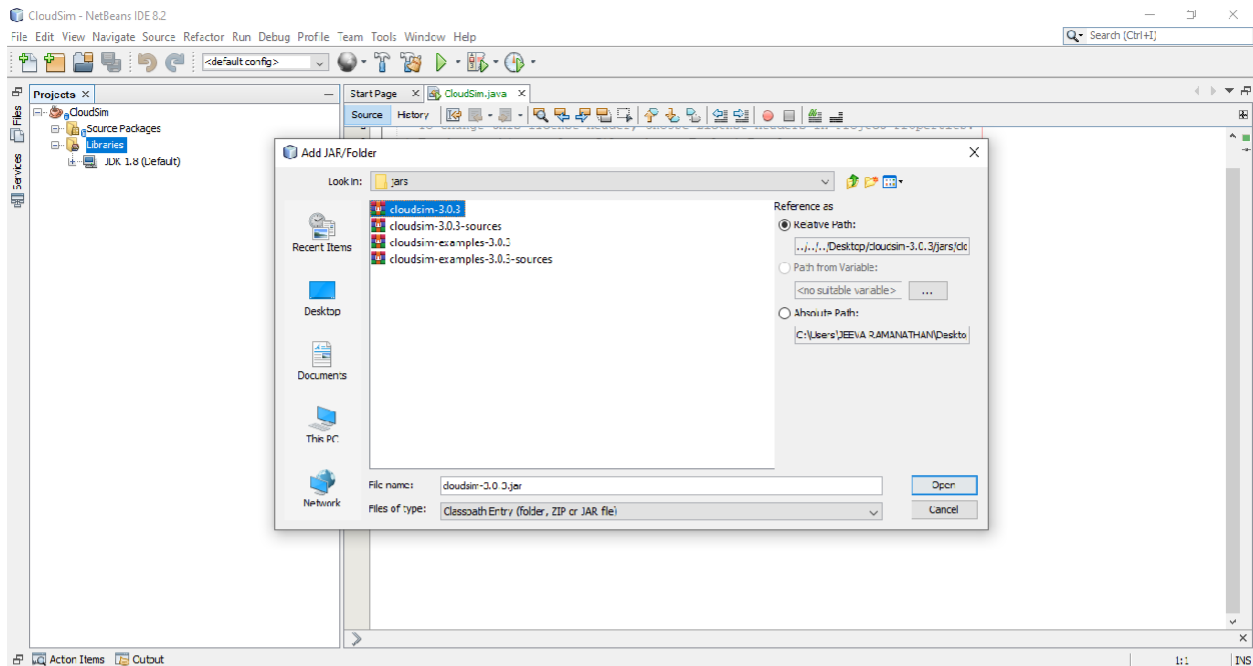
## Step 2: Open Netbeans and create a new project named "Clousim"



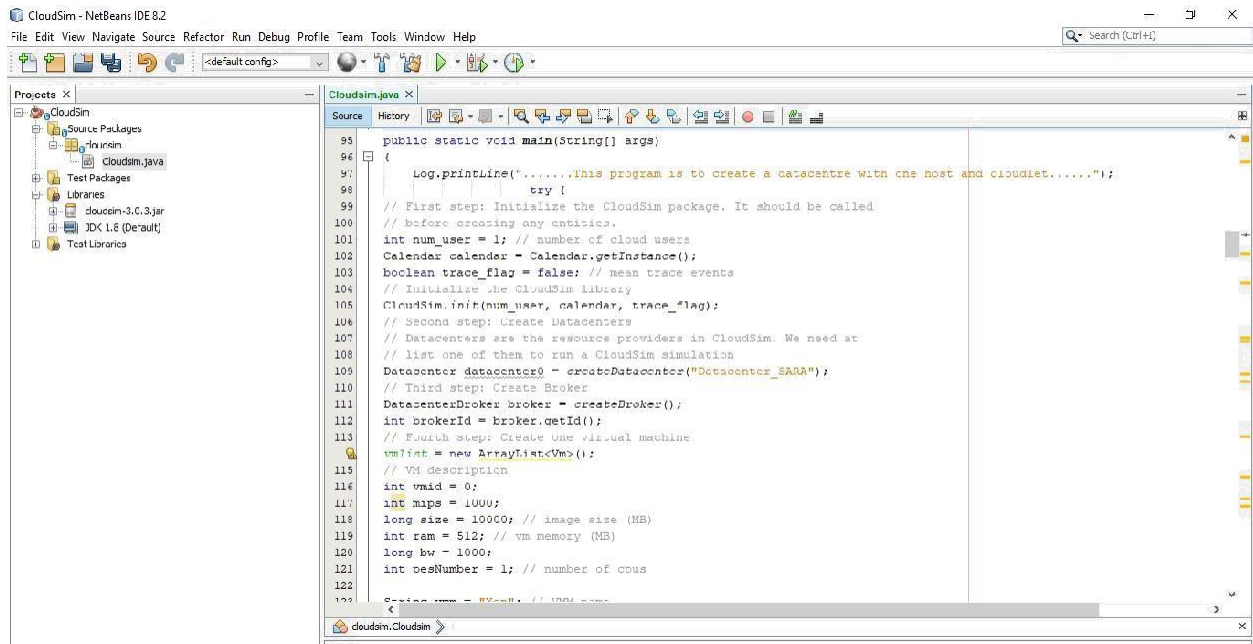


Step 3: Type the code and add the jar file to the libraries. The jar file will be in the extracted cloudsims.

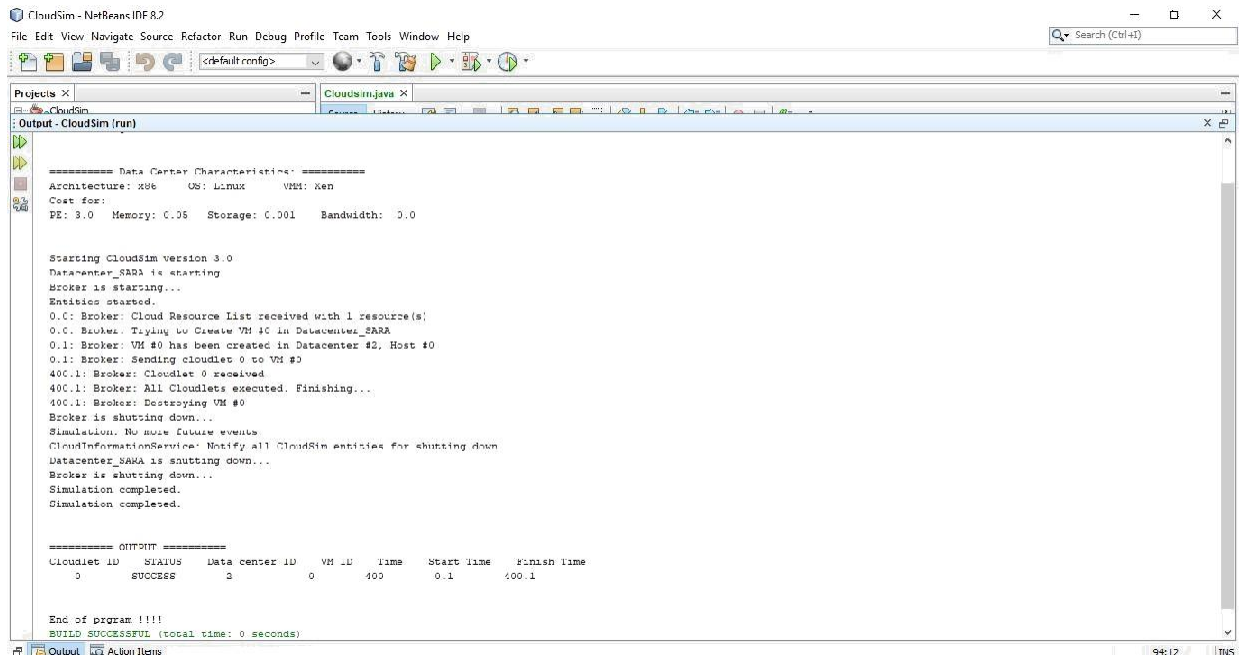




## Step 4: Run the application

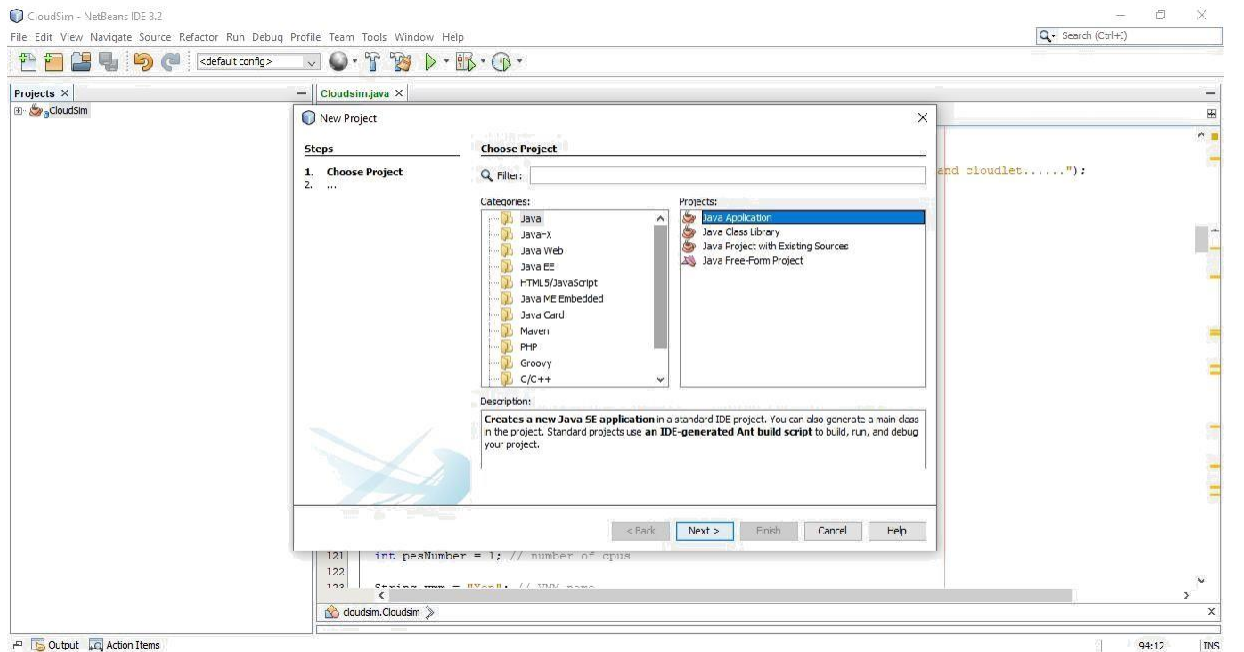


## OUTPUT:

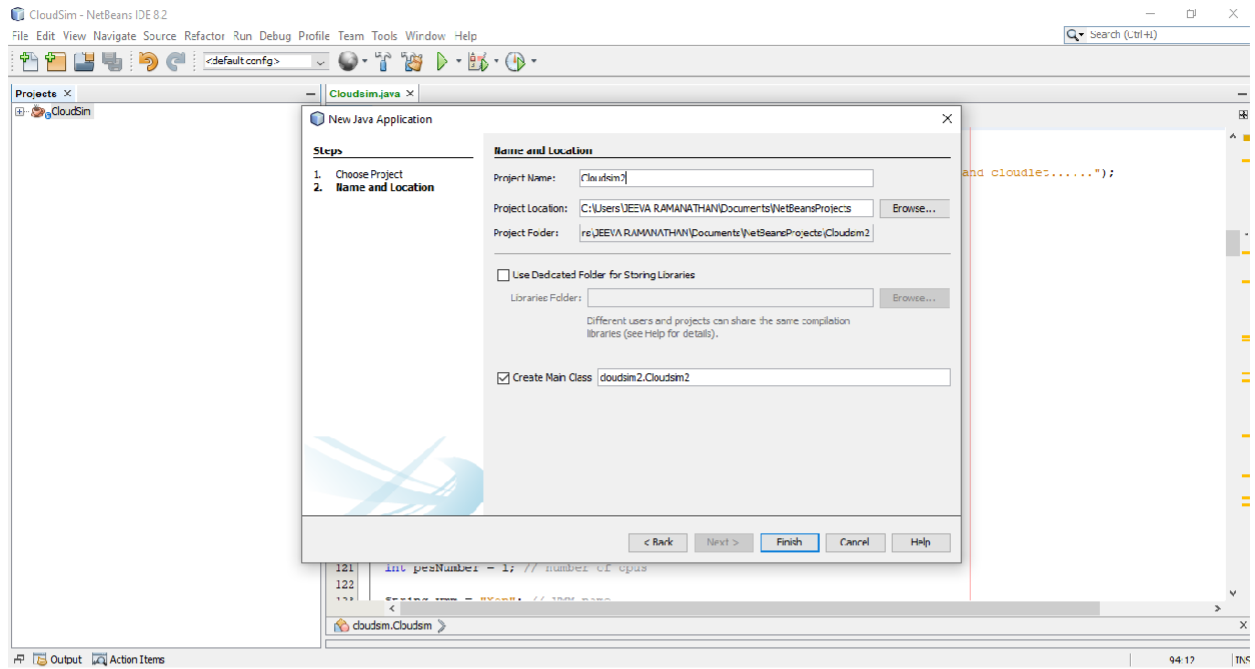


```
=====  
Data Center Characteristics: =====  
Architecture: x86 OS: Linux VM: Ken  
Cost for:  
PE: 3.0 Memory: 0.06 Storage: 0.001 Bandwidth: 0.0  
  
Starting CloudSim version 3.0  
Datacenter_SARA is starting  
Broker is starting...  
Entities started.  
0.0: Broker: Cloud Resource List received with 1 resource(s)  
0.0: Broker: Trying to create VM 10 in Datacenter_SARA  
0.1: Broker: VM #0 has been created in Datacenter #2, Host #0  
0.1: Broker: Sending cloudlet 0 to VM #0  
400.1: Broker: Cloudlet 0 received  
400.1: Broker: All Cloudlets executed. Finishing...  
400.1: Broker: Destroying VM #0  
Broker is shutting down...  
Simulation: No more future events  
CloudInformationService: Notify all CloudSim entities for shutting down  
Datacenter_SARA is shutting down...  
Broker is shutting down...  
Simulation completed.  
Simulation completed.  
  
=====  
Cloudlet ID STATUS Data center ID VM ID Time Start time Finish time  
0 SUCCESS 2 0 400 0.1 400.1  
  
End of program !!!!!  
BUILD SUCCESSFUL (total time: 0 seconds)
```

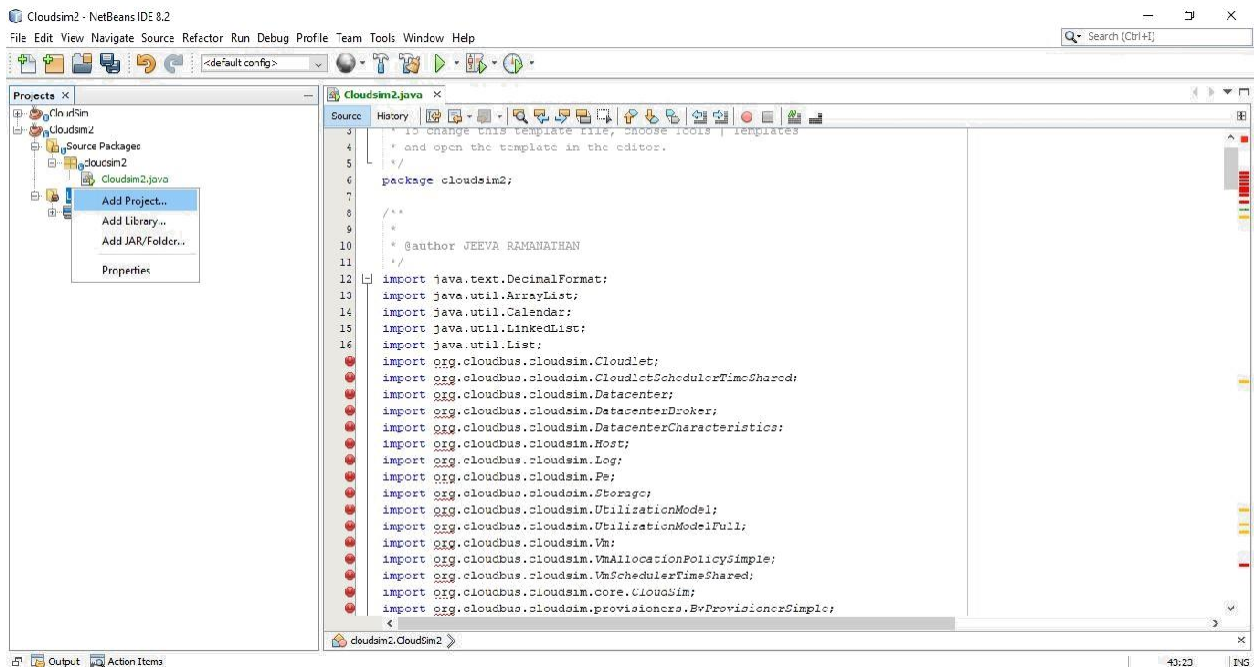
## Step 5: Again create a new project and name it as “Cloudsim2”

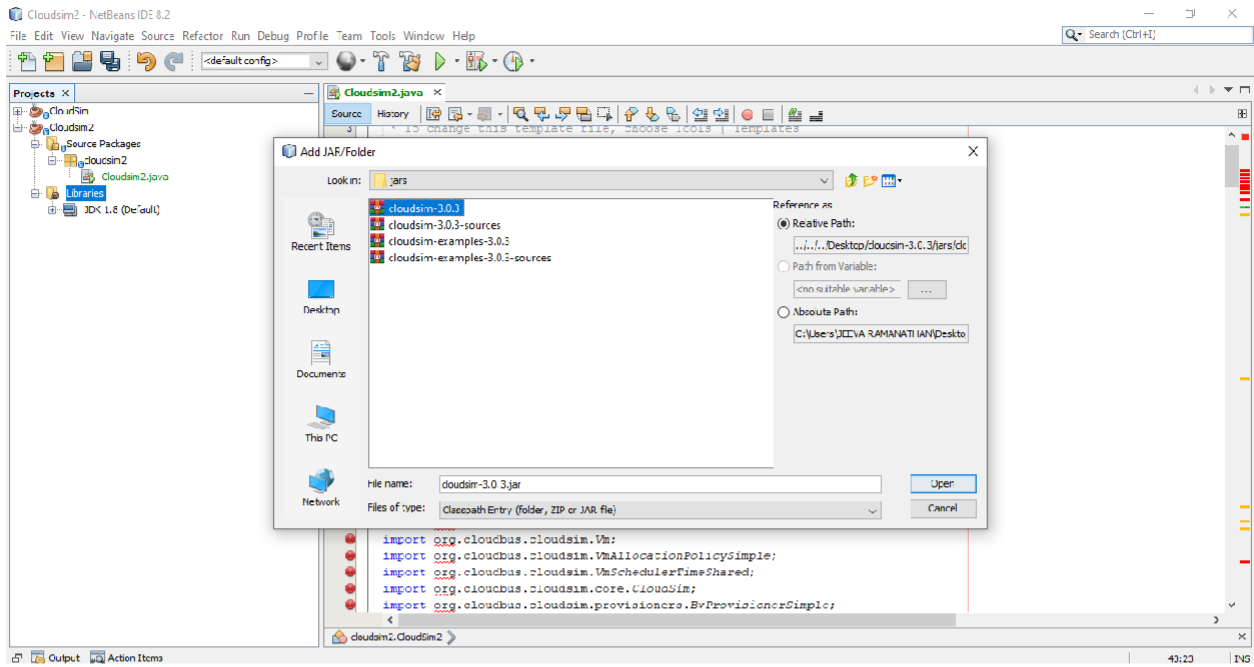


The screenshot shows the NetBeans IDE interface with the 'New Project' dialog box open. The 'Steps' section on the left lists '1. Choose Project' and '2. ...'. The 'Choose Project' section has a search filter and a list of categories. The 'Projects' list on the right shows 'Java Application' selected. The description at the bottom states: 'Creates a new Java SE application in a standard IDE project. You can also generate a main class in the project. Standard projects use an IDE-generated Ant build script to build, run, and debug your project.'

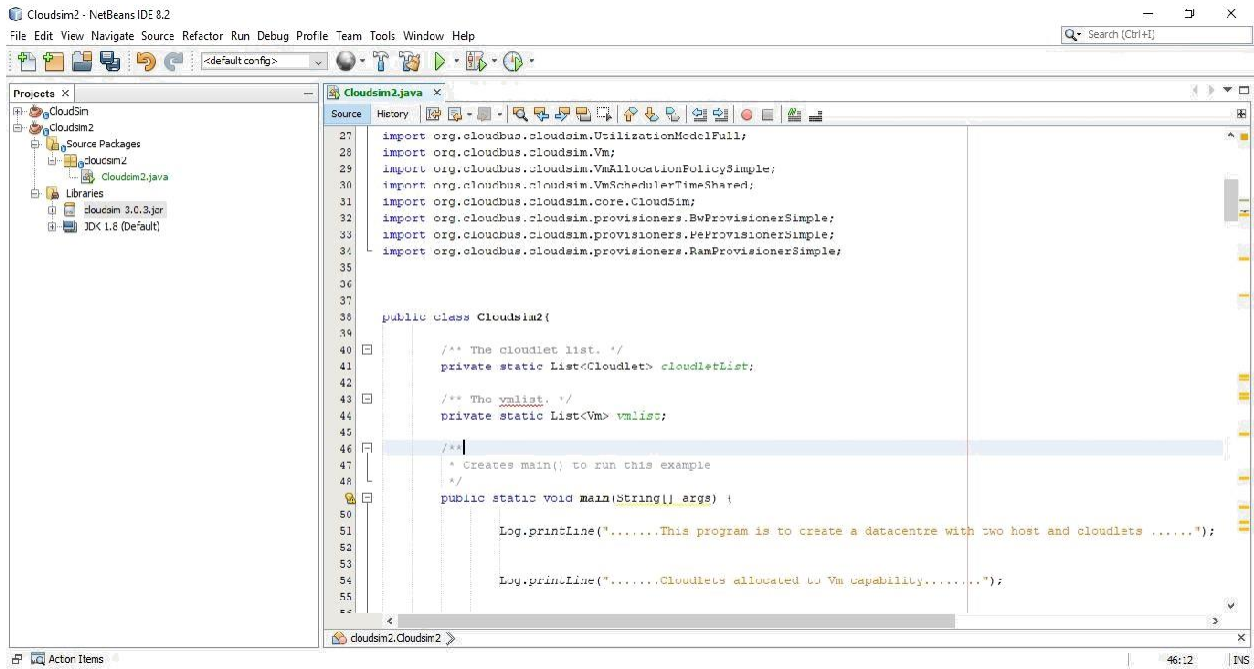


## Step 6: Type the code and include the jar file

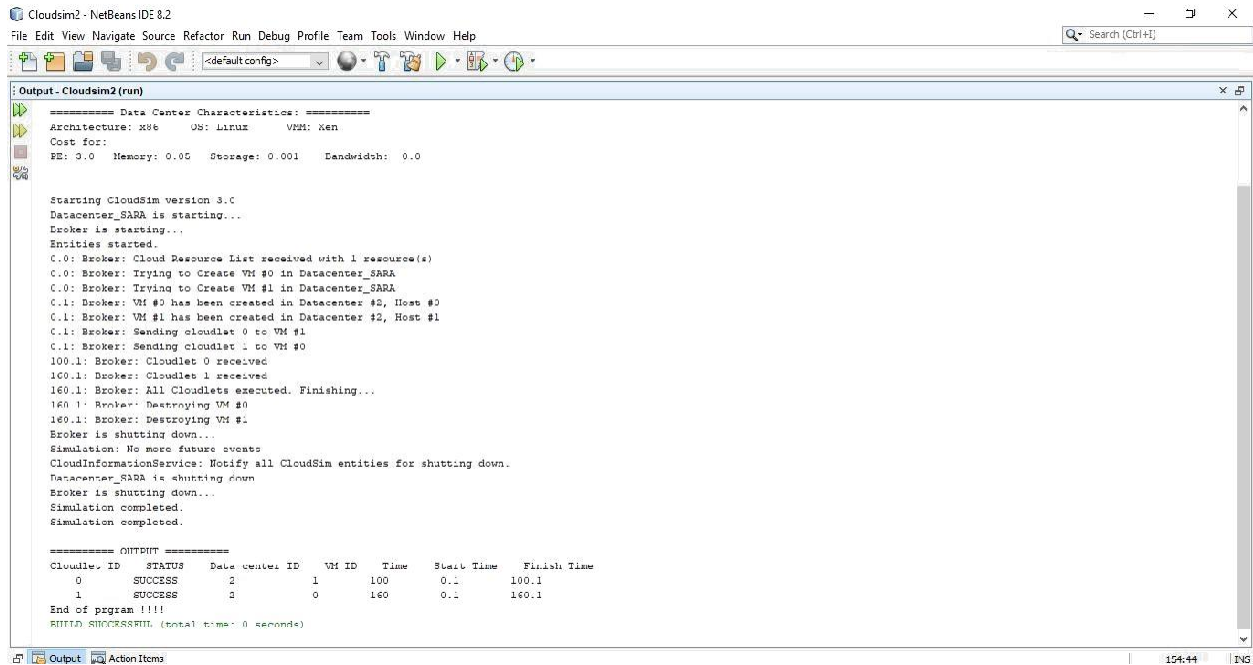




## Step 7: Run the application



## OUTPUT:



```
===== Data Center Characteristics: =====
Architecture: x86   OS: Linux   VMs: Ken
Cost for:
RE: 0.0  Memory: 0.05  Storage: 0.001  Bandwidth: 0.0

Starting CloudSim version 3.0
Datacenter_SARA is starting...
Broker is starting...
Entities started.
C.0: Broker: Cloud Resource List received with 1 resource(s)
C.0: Broker: Trying to Create VM #0 in Datacenter_SARA
C.0: Broker: Trying to Create VM #1 in Datacenter_SARA
C.1: Broker: VM #0 has been created in Datacenter #2, Host #2
C.1: Broker: VM #1 has been created in Datacenter #2, Host #1
C.1: Broker: Sending cloudlet 0 to VM #1
C.1: Broker: Sending cloudlet 1 to VM #0
100.1: Broker: Cloudlet 0 received
100.1: Broker: Cloudlet 1 received
160.1: Broker: All Cloudlets executed. Finishing...
160.1: Broker: Destroying VM #0
160.1: Broker: Destroying VM #1
Broker is shutting down...
Simulation: No more future events
CloudInformationService: Notify all CloudSim entities for shutting down.
Datacenter_SARA is shutting down
Broker is shutting down...
Simulation completed.
Simulation completed.

===== OUTPUT =====
Cloudlet ID  STATUS  Data center ID  VM ID  Time  Start Time  Finish Time
0           SUCCESS  2               1      100   0.1         100.1
1           SUCCESS  2               0      160   0.1         160.1

End of program !!!!
RMSTD SUCCESSFUL (total time: 0 seconds)
```

## Result:

Thus the simulation a cloud scenario using CloudSim and run a scheduling algorithm that is not present in CloudSim is done and the output is obtained successfully.