

Installation Hadoop 3.1.1 over Linux Mint 19 x64

```
$ ls10@LingoSoft:~$ echo $JAVA_HOME
/usr/lib/jvm/java-8-oracle

ls10@LingoSoft:~$ hostname
LingoSoft

ls10@LingoSoft:~$ sudo su
[sudo] password for ls10:

root@LingoSoft:/home/ls10# whoami
root

root@LingoSoft:/home/ls10# sudo addgroup hadoop

root@LingoSoft:/home/ls10# sudo adduser hduser

root@LingoSoft:/home/ls10# sudo adduser hduser hadoop

root@LingoSoft:/home/ls10# sudo visudo

    hduser ALL=(ALL) ALL

hduser@LingoSoft:~$ sudo apt-get install openssh-server

hduser@LingoSoft:~$ clear

hduser@LingoSoft:~$ ssh-keygen

hduser@LingoSoft:~$ cat ~/.ssh/id_rsa.pub >> ~/.ssh/authorized_keys

hduser@LingoSoft:~$ chmod 700 ~/.ssh/authorized_keys

hduser@LingoSoft:~$ sudo /etc/init.d/ssh restart

hduser@LingoSoft:~$ ssh localhost

hduser@LingoSoft:~$ sudo apt-get install vim

hduser@LingoSoft:~$ sudo vim /etc/sysctl.conf

    # disable ipv6
    net.ipv6.conf.all.disable_ipv6 = 1
    net.ipv6.conf.default.disable_ipv6 = 1
    net.ipv6.conf.lo.disable_ipv6 = 1

Press "i" to insert. Paste it.

Press "ESC" then write :wq then Enter. (wq -> write quit (for Save
Changes.))

hduser@LingoSoft:~$ cat /proc/sys/net/ipv6/conf/all/disable_ipv6
0
```

Should be 1, meaning that IPv6 is actually disabled.
Without reboot it would be showing you 0.

Hadoop Installation

```
hduser@LingoSoft:~/indirilenler$ sudo mv hadoop-3.1.1.tar.gz /usr/local/  
[sudo] password for hduser:
```

```
hduser@LingoSoft:~/indirilenler$ cd /usr/local/
```

```
hduser@LingoSoft:/usr/local$ ls -l
```

```
hduser@LingoSoft:/usr/local$ sudo tar -xvf hadoop-3.1.1.tar.gz
```

```
hduser@LingoSoft:/usr/local$ sudo rm hadoop-3.1.1.tar.gz
```

```
hduser@LingoSoft:/usr/local$ sudo ln -s hadoop-3.1.1 hadoop
```

```
hduser@LingoSoft:/usr/local$ sudo chown -R hduser:hadoop hadoop-3.1.1
```

```
hduser@LingoSoft:/usr/local$ sudo chmod 777 hadoop-3.1.1
```

```
hduser@LingoSoft:/usr/local$ ls -l
```

```
hduser@LingoSoft:/usr/local$ echo $JAVA_HOME  
/usr/lib/jvm/java-8-oracle
```

Edit hadoop-env.sh and configure Java

```
hduser@LingoSoft:/usr/local$ sudo vim  
/usr/local/hadoop/etc/hadoop/hadoop-env.sh
```

```
export HADOOP_OPTS=-Djava.net.preferIPv4Stack=true  
export HADOOP_HOME_WARN_SUPPRESS="TRUE"  
export JAVA_HOME=/usr/lib/jvm/java-8-oracle
```

```
hduser@LingoSoft:/usr/local$ vim ~/.bashrc
```

```
# Set Hadoop-related environment variables  
export HADOOP_HOME=/usr/local/hadoop  
export HADOOP_MAPRED_HOME=${HADOOP_HOME}  
export HADOOP_COMMON_HOME=${HADOOP_HOME}  
export HADOOP_HDFS_HOME=${HADOOP_HOME}  
export HADOOP_YARN_HOME=${HADOOP_HOME}  
export HADOOP_CONF_DIR=${HADOOP_HOME}/etc/hadoop
```

```
# Native Path  
export HADOOP_COMMON_LIB_NATIVE_DIR=${HADOOP_PREFIX}/lib/native  
export HADOOP_OPTS="-Djava.library.path=${HADOOP_PREFIX}/lib"
```

```
# Set JAVA_HOME (we will also configure JAVA_HOME directly for  
Hadoop later on)  
export JAVA_HOME=/usr/lib/jvm/java-8-oracle
```

```
# Some convenient aliases and functions for running Hadoop-related  
commands
```

```
unalias fs&> /dev/null  
alias fs="hadoop fs"  
unalias ls&> /dev/null  
alias ls="fs -ls"
```

```
export
PATH=$PATH:$HADOOP_HOME/bin:$PATH:$JAVA_HOME/bin:$HADOOP_HOME/sbin
```

```
hduser@LingoSoft:/usr/local$ cd ~
```

```
hduser@LingoSoft:~$ sudo vim /usr/local/hadoop/etc/hadoop/yarn-site.xml
<property>
  <name>yarn.nodemanager.aux-services</name>
  <value>mapreduce_shuffle</value>
</property>
<property>
  <name>yarn.nodemanager.aux-services.mapreduce.shuffle.class</name>
  <value>org.apache.hadoop.mapred.ShuffleHandler</value>
</property>
```

```
hduser@LingoSoft:~$ sudo vim /usr/local/hadoop/etc/hadoop/core-site.xml
<property>
  <name>hadoop.tmp.dir</name>
  <value>/app/hadoop/tmp</value>
  <description>A base for other temporary directories.</description>
</property>

<property>
  <name>fs.default.name</name>
  <value>hdfs://localhost:9000</value>
  <description>default host and port</description>
</property>
```

```
hduser@LingoSoft:~$ sudo mkdir -p /app/hadoop/tmp
```

```
hduser@LingoSoft:~$ sudo chown hduser:hadoop -R /app/hadoop/tmp
```

```
hduser@LingoSoft:~$ sudo chmod 750 /app/hadoop/tmp
```

```
hduser@LingoSoft:~$ sudo vim /usr/local/hadoop/etc/hadoop/mapred-site.xml
<property>
  <name>mapreduce.framework.name</name>
  <value>yarn</value>
</property>
<property>
  <name>mapreduce.jobhistory.address</name>
  <value>localhost:10020</value>
  <description>Host and port for Job History Server (default
0.0.0.0:10020)</description>
</property>
<property>
  <name>yarn.app.mapreduce.am.env</name>
  <value>HADOOP_MAPRED_HOME=${HADOOP_HOME}</value>
</property>
<property>
  <name>mapreduce.map.env</name>
  <value>HADOOP_MAPRED_HOME=${HADOOP_HOME}</value>
</property>
<property>
  <name>mapreduce.reduce.env</name>
  <value>HADOOP_MAPRED_HOME=${HADOOP_HOME}</value>
</property>
```

```
sudo mkdir -p /usr/local/hadoop_tmp/hdfs/namenode
sudo mkdir -p /usr/local/hadoop_tmp/hdfs/datanode
sudo chown hduser:hadoop -R /usr/local/hadoop_tmp/
```

```
hduser@LingoSoft:~$ sudo vim /usr/local/hadoop/etc/hadoop/hdfs-site.xml
<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>
```

```
hduser@LingoSoft:~$ hadoop namenode -format
```

```
hduser@LingoSoft:~$ start-dfs.sh -- Starts NN,SNN,DN
hduser@LingoSoft:~$ start-yarn.sh -- Starts
NodeManager,ResourceManager
hduser@LingoSoft:~$ start-dfs.sh && start-yarn.sh -- iki komutu
tek satırda yazmak için
```

```
hduser@LingoSoft:~$ start-all.sh
```

```
hduser@LingoSoft: mr-jobhistory-daemon.sh start historyserver
-If you want to stop- mr-jobhistory-daemon.sh stop historyserver
```

```
hduser@LingoSoft:~$ hadoop fs -ls
```

```
hduser@LingoSoft:~$ hadoop fs -mkdir -p /user/hduser (Deprecated)
hduser@LingoSoft:~$ hdfs dfs -mkdir -p /user/hduser (Use this)
```

```
hduser@LingoSoft:~$ hdfs dfs -ls
```

```
-----
127.0.0.1 / localhost
-----
```

```
NameNode http://127.0.0.1:9870
```

```
ResourceManager http://127.0.0.1:8088
```

```
MapReduce JobHistory Server http://127.0.0.1:19888
```