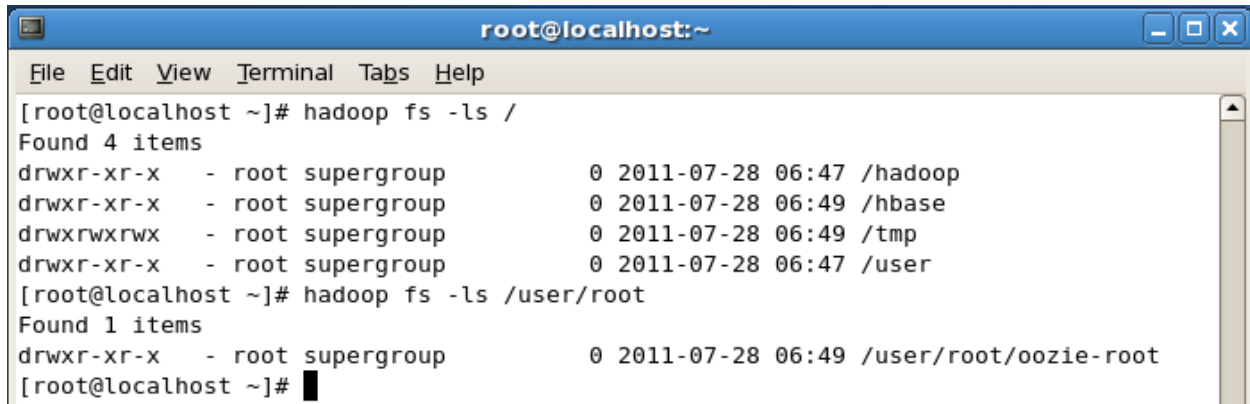


1. In order to work with HDFS you need to use the **hadoop fs** command. For example to list the / and /user/root directories you need to input the following commands:

```
> hadoop fs -ls /  
> hadoop fs -ls /user/root
```



```
root@localhost:~  
File Edit View Terminal Tabs Help  
[root@localhost ~]# hadoop fs -ls /  
Found 4 items  
drwxr-xr-x - root supergroup          0 2011-07-28 06:47 /hadoop  
drwxr-xr-x - root supergroup          0 2011-07-28 06:49 /hbase  
drwxrwxrwx - root supergroup          0 2011-07-28 06:49 /tmp  
drwxr-xr-x - root supergroup          0 2011-07-28 06:47 /user  
[root@localhost ~]# hadoop fs -ls /user/root  
Found 1 items  
drwxr-xr-x - root supergroup          0 2011-07-28 06:49 /user/root/oozie-root  
[root@localhost ~]# █
```

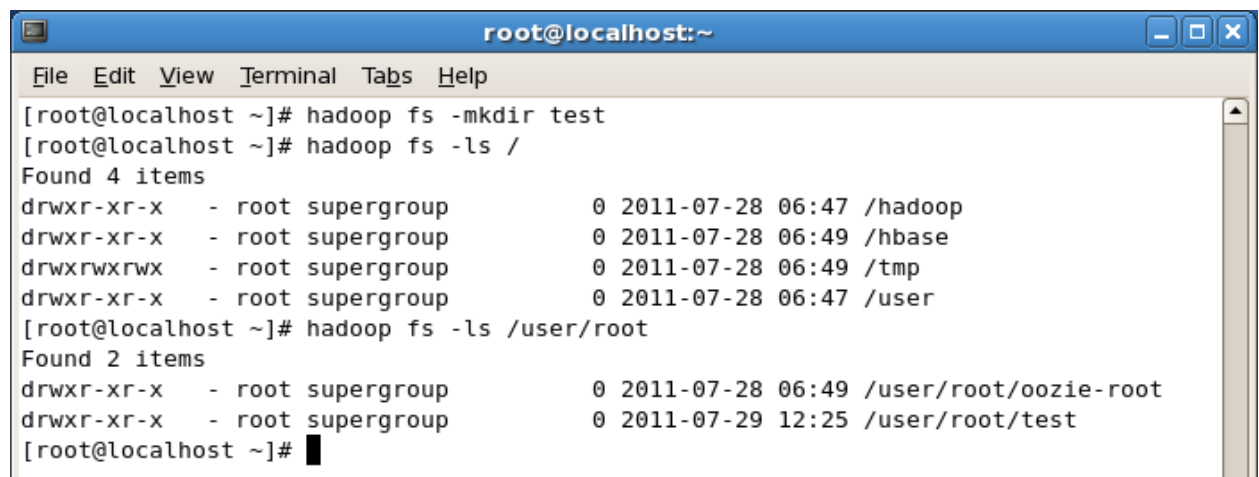
2. There are many commands you can run within the Hadoop filesystem. For example to make the directory *test* you can issue the following command:

```
> hadoop fs -mkdir test
```

Now let's see the directory we've created:

```
> hadoop fs -ls /  
> hadoop fs -ls /user/root
```

You will notice that the test directory got created under the /user/root directory. This is because as the root user, your default path is /user/root and thus if you don't specify an absolute path all HDFS commands work out of /user/root (this will be your default working directory).

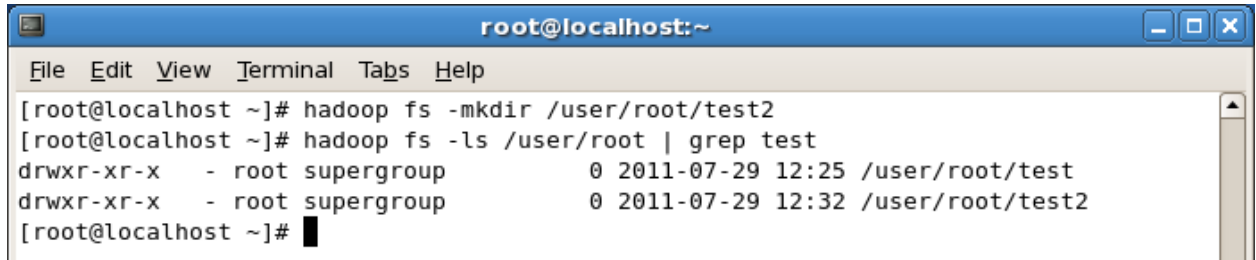


```
root@localhost:~  
File Edit View Terminal Tabs Help  
[root@localhost ~]# hadoop fs -mkdir test  
[root@localhost ~]# hadoop fs -ls /  
Found 4 items  
drwxr-xr-x - root supergroup          0 2011-07-28 06:47 /hadoop  
drwxr-xr-x - root supergroup          0 2011-07-28 06:49 /hbase  
drwxrwxrwx - root supergroup          0 2011-07-28 06:49 /tmp  
drwxr-xr-x - root supergroup          0 2011-07-28 06:47 /user  
[root@localhost ~]# hadoop fs -ls /user/root  
Found 2 items  
drwxr-xr-x - root supergroup          0 2011-07-28 06:49 /user/root/oozie-root  
drwxr-xr-x - root supergroup          0 2011-07-29 12:25 /user/root/test  
[root@localhost ~]# █
```

3. You should be aware that you can pipe (using the `|` character) any HDFS command to be used with the Linux shell. For example, you can easily use *grep* with HDFS by doing the following:

```
> hadoop fs -mkdir /user/root/test2  
> hadoop fs -ls /user/root | grep test
```

As you can see the *grep* command only returned the lines which had *test* in them (thus removing the “Found x items” line and *oozie-root* directory from the listing).



The screenshot shows a terminal window titled "root@localhost:~". The terminal content is as follows:

```
File Edit View Terminal Tabs Help  
[root@localhost ~]# hadoop fs -mkdir /user/root/test2  
[root@localhost ~]# hadoop fs -ls /user/root | grep test  
drwxr-xr-x - root supergroup 0 2011-07-29 12:25 /user/root/test  
drwxr-xr-x - root supergroup 0 2011-07-29 12:32 /user/root/test2  
[root@localhost ~]# █
```

4. In order to move files between your regular linux filesystem and HDFS you will likely use the *put* and *get* commands. First, move a single file to the hadoop filesystem.

```
> hadoop fs -put /BigDataUniversity/README README  
> hadoop fs -ls /user/root
```

You should now see a new file called `/user/root/README` listed. In order to view the contents of this file we will use the `-cat` command as follows:

```
> hadoop fs -cat README
```

You should see the output of the `README` file (that is stored in HDFS). We can also use the linux *diff* command to see if the file we put on HDFS is actually the same as the original on the local filesystem. You can do this as follows:

```
> diff <( hadoop fs -cat README ) /BigData/README
```

Since the *diff* command produces no output we know that the files are the same (the *diff* command prints all the lines in the files that differ).

```
root@localhost:~
File Edit View Terminal Tabs Help
[root@localhost ~]# hadoop fs -put /BigDataUniversity/README README
[root@localhost ~]# hadoop fs -ls /user/root
Found 4 items
-rw-r--r--   3 root supergroup      20 2011-07-29 12:41 /user/root/README
drwxr-xr-x   - root supergroup      0 2011-07-28 06:49 /user/root/oozie-root
drwxr-xr-x   - root supergroup      0 2011-07-29 12:25 /user/root/test
drwxr-xr-x   - root supergroup      0 2011-07-29 12:32 /user/root/test2
[root@localhost ~]# hadoop fs -cat README
I am a README file

[root@localhost ~]# diff <( hadoop fs -cat README ) /BigDataUniversity/README
[root@localhost ~]# █
```

Some more Hadoop Filesystem commands

1. In order to use HDFS commands recursively generally you add an “r” to the HDFS command (In the Linux shell this is generally done with the “-R” argument) For example, to do a recursive listing we’ll use the `-lsr` command rather than just `-ls`. Try this:

```
> hadoop fs -ls /user
> hadoop fs -lsr /user
```

```
root@localhost:~
File Edit View Terminal Tabs Help
[root@localhost ~]# hadoop fs -ls /user
Found 2 items
drwxr-xr-x   - root supergroup      0 2011-07-28 06:43 /user/hive
drwxr-xr-x   - root supergroup      0 2011-07-29 12:41 /user/root
[root@localhost ~]# hadoop fs -lsr /user
drwxr-xr-x   - root supergroup      0 2011-07-28 06:43 /user/hive
drwxrwxr-x   - root supergroup      0 2011-07-28 06:48 /user/hive/warehouse
drwxr-xr-x   - root supergroup      0 2011-07-29 12:41 /user/root
-rw-r--r--   3 root supergroup      20 2011-07-29 12:41 /user/root/README
drwxr-xr-x   - root supergroup      0 2011-07-28 06:49 /user/root/oozie-root
drwxr-xr-x   - root supergroup      0 2011-07-28 06:50 /user/root/oozie-root
/0000000-110728064455684-oozie-root-W
drwxr-xr-x   - root supergroup      0 2011-07-29 12:25 /user/root/test
drwxr-xr-x   - root supergroup      0 2011-07-29 12:32 /user/root/test2
[root@localhost ~]#
```

