

Unix Commands

Notes

- Commands preceded with “\$” imply that you should execute the command as a general user - not as root.
- Commands preceded with “#” imply that you should be working as root with “sudo”
- Commands with more specific command lines (e.g. “RTR-GW>” or “mysql>”) imply that you are executing commands on remote equipment, or within another program.

Goals

- Initial VM Configuration
- We could do this for you, but it's important to understand how some of this software work with the tools you will be installing this week.

Exercises

Log in as the *sysadm* user.

If you have been allocated a virtual machine by the instructor, you will log in as SSH. If this is a machine running inside VirtualBox on your laptop, you will probably log in directly on the console.

```
username: sysadm
password: `<given in class>`
```

Become the root user

At the command prompt type the following command:

```
$ sudo -s
```

Enter the class user's password when prompted

Now that you are root the command prompt will change. We indicate this using the “#” symbol.

You are now the superuser - be careful!

Ok, exit the root account:

```
# exit
$
```

Update your software package repository

```
$ sudo apt-get update
```

This might take a few moments if everyone in class is doing this at the same moment.

Install the "nano" editor package

```
$ sudo apt-get install nano
```

The nano editor package is simpler to use than vi. Try using the editor to create a new file in your sysadm home directory:

```
$ cd  
$ nano newfile.txt
```

Type in some text for practice. You can type `ctrl-g` to see a list of nano editor command, that is "press the ctrl key and the g key".

You can save and exit from the file by typing `ctrl-x`.

Setting time to UTC, Updating time and install Network Time Protocol service

In order to manage and monitor your network it is *critical* that all devices and servers maintain the same, consistent time. To achieve this you can, for example, select a single time zone, use the `ntpd` command to set your server's clock exactly and install the NTP (Network Time Protocol) service to maintain your server's clock with precise time.

First, let's set your server's clock to use UTC time (Coordinated Universal Time). At the command line type:

```
$ sudo dpkg-reconfigure tzdata
```

- Scroll to the bottom of the list and select "None of the above"
- Scroll down the list and select "UTC"
- Use the tab key to select `<0k>` and press `<ENTER>`

Now your server is using UTC time. Next be sure the time is precise by using `ntpd`:

```
$ sudo ntpdate -s ntp.ubuntu.com
```

You can always type:

```
$ date
```

to see your server's current timezone (UTC, which is technically a standard), date and time.

Finally, let's install the NTP service to ensure that our server's clock maintains precise time.

```
$ sudo apt-get install ntp
```

At this point the default configuration should be acceptable for our case. You may wish to read up on *ntp* upon returning home and edit the file `/etc/ntp.conf` to select different time servers, or update settings to your local *ntp* service configuration.

In addition, *ntp* has been part of several security warnings the past few years. You should sign up for the Ubuntu Security mailing list at:

<https://lists.ubuntu.com/mailman/listinfo/ubuntu-security-announce>

You should do this whether you run *ntp* or not. And, as *ntp* is so critical to proper network instrumentation, this is one service that should be run on any server that will be running network monitoring or management software or that will be monitored and on all your network devices.

If you would like to see the status of your local *ntp* service you can type:

```
$ sudo ntpq -p
```

and you should see something like:

remote	refid	st	t	when	poll	reach	delay	offset
jitter								
=====								
==								
ntp1.dpacguam.n .GPS.		1	u	7	64	1	53.568	3.605
0.000								
ntp4.dpacguam.n .GPS.		1	u	6	64	1	52.234	3.368
0.000								
ns2.unico.com.a .INIT.		16	u	-	64	0	0.000	0.000
0.000								
ntp3.dpacguam.n .GPS.		1	u	6	64	1	50.163	3.833
0.000								
golem.canonical .INIT.		16	u	-	64	0	0.000	0.000
0.000								

For a reasonable discussion of what this output means see:

<http://tech.kulish.com/2007/10/30/ntp-ntpq-output-explained/>

Install the postfix mailer software and additional utilities

At the command line type:

```
$ sudo apt-get install postfix mutt mailutils dnsutils traceroute man-db
```

You will see several prompts. If you are using putty ssh sometimes the screens can be harder to read.

When you are prompted (a fair number of packages will be installed):

- At the initial prompt press <Enter> for <Ok>
- When available, select “Internet Site” (use tab key to move to <Ok> and press <Enter> to continue)
- Accept the hostname presented (tab to <Ok> then press <Enter>)

For fun you can practice restarting a service by restarting the *postfix* mailserver. Note that the service was started as soon as installation was completed:

```
$ sudo service postfix restart
```

You might do this if you changed a *postfix* configuration file.

Viewing log files in real time

Log files are critical to solve problems. They reside (largely) in the `/var/log/` directory.

Some popular log files include:

- `/var/log/syslog`
- `/var/log/apache2/access.log`
- `/var/log/mail.log`

and many more.

To view the last entry in a log file, such as the system log file, type:

```
$ sudo tail /var/log/syslog
```

Some log files may require that you use “`sudo tail logfilename`” to view their contents.

What's more effective is to watch a log file as you perform some action on your system. To do this open another ssh session to your server now, log in as user *sysadm* and in that other window type:

```
$ sudo tail -f /var/log/syslog
```

Now in your other window try restarting the *ntp* service you recently installed:

```
$ sudo service ntp restart
```

You should see several log messages appear in your other ssh window. These are real-time messages coming from the *ntp* service. We'll talk about logging more later in the week, but viewing your log files to debug issues is often the only way to solve a problem.

In the window where you typed `sudo tail -f /var/log/syslog` you can press `ctrl-c` to exit from the `tail` command.

Practice using the man command

To get help on common commands you can simply type (don't do this now):

```
$ man command-name
```

Let's imagine you wanted to know more about the ssh command. To get help type (do this now):

```
$ man ssh
```

Now you can move around the help screen quickly by using some editing tricks. Note that these tricks work if you are using the less command as well.

Try doing the following:

- Search for “ports” by typing /ports and press <ENTER>
- Press **n** to go to the next occurrence of **port** – do this several times.
- Press **N** to search backwards.
- Press **p** to go to the start.
- Search on “/-p” and see what you find.
- Press **h** for all the keyboard shortcuts.
- Press **q** (twice in this case) to quit from the man page.

Look at the network configuration of your host

```
$ cat /etc/network/interfaces
```

The IP configuration of your host is either done using DHCP, or configured statically. Which is it in your case ?

“cat” is for “concatenate” and is one way to view what is in a file.

List files

Use `ls` to list files:

```
$ cd      [go to your home directory]
$ ls
```

Do you see anything? Try this instead:

```
$ ls -lah
```

What's inside one of these files?

```
$ cat .profile
```

```
$ less .profile
```

Press q to get out of the less display.

Another command:

```
$ clear
```

If you don't understand what cat, clear or less do, then type:

```
$ man cat  
$ man clear  
$ man less
```

Working with the command prompt

You can recall previous commands by using the up-arrow and down-arrow keys. Give this a try now.

Alternately, try typing this command:

```
$ history
```

If you wish to execute one of the commands in the list you saw type:

```
$ !nn
```

Where nn is the number of the command in the history list. This is useful if you want to run a past command that was long and/or complicated.

Command completion

With the bash shell you can auto-complete commands using the tab key. This means, if you type part of a command, once you have a unique string if you press the TAB key the command will complete. If you press the TAB key twice you'll see all your available options. Your instructor will demonstrate this, but give it a try by doing:

```
$ hist <TAB>  
$ del <TAB><TAB>  
$ rm <TAB><TAB>      [Include the space after the "rm"]
```

Working with pipes

What if you wanted to have the results of `ls /sbin` and put this information in a file called `sbin.txt` and have the results sorted alphabetically? You can do the following:

```
$ cd
$ ls /sbin | sort > sbin.txt
```

Now view the contents of what is in sbin.txt to verify that this worked.

```
$ less sbin.txt
```

Press the “q” key to quit viewing the contents.

Finding text strings

Use the command `grep` to print lines matching a pattern in a data stream (such as a file). For example, view the entry for the *nsr*c account in the system `passwd` file:

```
$ sudo grep sysadm /etc/passwd
```

You should see something like:

```
sysadm:x:1001:1001:NSRC Workshop Participant,,,:/home/sysadm:/bin/bash
```

The previous items above are:

```
userid:passwd:uid:gid:Name,extrastuff,,:HomeDir:LoginShell
```

`grep` is often used with a pipe to FILTER the output of commands. For instance:

```
$ history | grep ls
```

Will display your previous use of the `ls` command from exercise 2.

Editing the command line revisited

It is particularly useful to realize that you can edit a command just as you would a line of text in a file. For instance, you can:

- Use your up-arrow to select a previous command to work with, or history as noted below.
- Use your back-arrow (left) and forward-arrow (right) keys to change text in a command.
- Use the Home and End keys to go to the start and the end of a command:
 - `ctrl-a` = start
 - `ctrl-e` = end

NOTE: you do not need to go to the end of a command before pressing <ENTER> to execute the command.

- You can use the `history` command with `grep` to find a previous command.
- You can copy and paste this command, then edit it to make adjustments.
- For long commands this can save considerable time.

- To terminate a command without executing it press ctrl-c

Alternatively you can use the reverse-search feature of bash:

1. Press ctrl-r
2. Type the term you are searching for.
3. Press ctrl-r again to cycle through all occurrences of the term in your history.
4. Press the right or left-arrow, HOME or END key to start editing the command.

Let's give some of these editing rules a try:

```
$ ls -lah /usr/lib/ | grep postfix
```

Then, let's look for postfix:

1. Press ctrl-r, type `postfix`, then press left arrow.
2. Edit the previous command (which you should now have) and change `/usr/lib/` to `/usr/sbin/`.
3. Use the left+right arrow key to move, and backspace to erase.

You should now have:

```
$ ls -lah /usr/sbin/ | grep postfix
```

With your cursor just past the `/` in `/sbin/`, press <ENTER> to execute the command.

From:

<https://workshops.nsrc.org/dokuwiki/> - **Workshops**

Permanent link:

https://workshops.nsrc.org/dokuwiki/2016/nsrc-tein-bdren/unix_commands

Last update: **2016/10/13 00:16**

