

Linux for Data Scientists

Paul Cobbaut Andy Van Maele Thomas Parmentier
 Bert Van Vreckem

September 18, 2024

Contents

I. First Linux VM	3
1. getting Linux at home	5
1.1. download a Linux CD image	5
1.2. download Virtualbox	6
1.3. create a virtual machine	6
1.4. attach the CD image	11
1.5. install Linux	14
II. Software management; curl	15
2. package management	17
2.1. package terminology	17
2.1.1. repository	17
2.1.2. .deb packages	17
2.1.3. .rpm packages	17
2.1.4. dependency	17
2.1.5. open source	18
2.1.6. GUI software management	18
2.2. deb package management	18
2.2.1. about deb	18
2.2.2. dpkg -l	19
2.2.3. dpkg -l \$package	19
2.2.4. dpkg -S	19
2.2.5. dpkg -L	19
2.2.6. dpkg	20
2.2.7. apt-get	20
2.2.8. apt-get update	20
2.2.9. apt-get upgrade	21
2.2.10. apt-get clean	22
2.2.11. apt-cache search	22
2.2.12. apt-get install	22
2.2.13. apt-get remove	23
2.2.14. apt-get purge	24
2.2.15. apt	25
2.2.16. /etc/apt/sources.list	26
2.3. the Red Hat package manager (rpm)	26
2.3.1. dnf	27
2.3.2. dnf list	27
2.3.3. dnf search	27
2.3.4. dnf info	28
2.3.5. dnf install	28
2.3.6. dnf upgrade	30
2.3.7. dnf provides	31
2.3.8. dnf remove	31
2.3.9. dnf software groups	32
2.3.10. rpm -qa	33
2.3.11. rpm -q	34
2.3.12. rpm -ql	34

2.3.13.	rpm -Uvh	34
2.3.14.	rpm -e	35
2.3.15.	Package cache	35
2.3.16.	Configuration	35
2.3.17.	Working with multiple repositories	36
2.4.	pip, the Python package manager	37
2.4.1.	installing pip	38
2.4.2.	listing packages	38
2.4.3.	searching for packages	38
2.4.4.	installing packages	39
2.4.5.	removing packages	39
2.5.	container-based package managers	39
2.5.1.	flatpak	39
2.5.2.	snap	40
2.6.	downloading software outside the repository	41
2.6.1.	example: compiling zork	41
2.6.2.	installing from a tarball	43
2.7.	practice: package management	43
2.8.	solution: package management	43
 III. Scripting 101		 45
3. I/O redirection		47
3.1.	stdin, stdout, and stderr	47
3.2.	output redirection	47
3.2.1.	> stdout	47
3.2.2.	output file is erased	48
3.2.3.	noclobber	48
3.2.4.	overruling noclobber	49
3.2.5.	» append	49
3.3.	error redirection	49
3.3.1.	2> stderr	49
3.3.2.	2>&1	49
3.4.	output redirection and pipes	50
3.5.	joining stdout and stderr	50
3.6.	input redirection	51
3.6.1.	< stdin	51
3.6.2.	« here document	51
3.6.3.	«< here string	51
3.7.	confusing redirection	52
3.8.	quick file clear	52
3.9.	practice: input/output redirection	52
3.10.	solution: input/output redirection	53
 4. filters		 55
4.1.	cat	55
4.2.	tee	55
4.3.	grep	56
4.4.	cut	57
4.5.	tr	57
4.6.	wc	59
4.7.	sort	59
4.8.	uniq	60
4.9.	comm	60
4.10.	od	61
4.11.	sed	62
4.12.	pipe examples	62
4.12.1.	who wc	62

4.12.2. who cut sort	63
4.12.3. grep cut	63
4.13. practice: filters	63
4.14. solution: filters	64
5. shell variables	67
5.1. \$ dollar sign	67
5.2. case sensitive	67
5.3. creating variables	67
5.4. quotes	68
5.5. set	68
5.6. unset	68
5.7. \$PS1	68
5.8. \$PATH	69
5.9. env	70
5.10. export	70
5.11. delineate variables	71
5.12. unbound variables	71
5.13. practice: shell variables	71
5.14. solution: shell variables	72
6. introduction to scripting	75
6.1. introduction	75
6.2. hello world	76
6.3. she-bang	76
6.4. comments	77
6.5. extension	77
6.6. shell variables	78
6.7. variable assignment	78
6.8. unbound variables	79
6.9. sourcing a script	79
6.10. quoting	80
6.11. troubleshooting a script	81
6.12. Bash's "strict mode"	81
6.13. prevent setuid root spoofing	82
6.14. practice: introduction to scripting	82
6.15. solution: introduction to scripting	83
IV. Organising users	85
7. standard file permissions	87
7.1. file ownership	87
7.1.1. user owner and group owner	87
7.1.2. chgrp	87
7.1.3. chown	88
7.2. list of special files	88
7.3. permissions	89
7.3.1. rwx	89
7.3.2. three sets of rwx	89
7.3.3. permission examples	89
7.3.4. setting permissions with symbolic notation	90
7.3.5. setting permissions with octal notation	91
7.3.6. umask	92
7.3.7. mkdir -m	93
7.3.8. cp -p	93
7.4. practice: standard file permissions	93
7.5. solution: standard file permissions	94

8. advanced file permissions	97
8.1. sticky bit on directory	97
8.2. setgid bit on directory	97
8.3. setgid and setuid on regular files	98
8.4. setuid on sudo	99
8.5. practice: sticky, setuid and setgid bits	99
8.6. solution: sticky, setuid and setgid bits	99
9. introduction to users	101
9.1. whoami	101
9.2. who	101
9.3. who am i	101
9.4. w	102
9.5. id	102
9.6. su to another user	102
9.7. su to root	102
9.8. su as root	102
9.9. su - \$username	103
9.10. su -	103
9.11. run a program as another user	103
9.12. visudo	103
9.13. sudo su -	104
9.14. sudo logging	104
9.15. practice: introduction to users	104
9.16. solution: introduction to users	105
10. user management	107
10.1. user management	107
10.2. /etc/passwd	107
10.3. root	108
10.4. useradd	108
10.5. /etc/default/useradd	108
10.6. userdel	108
10.7. usermod	109
10.8. creating home directories	109
10.9. /etc/skel/	109
10.10. deleting home directories	109
10.11. login shell	110
10.12. chsh	110
10.13. practice: user management	110
10.14. solution: user management	111
11. user passwords	113
11.1. passwd	113
11.2. shadow file	113
11.3. encryption with passwd	114
11.4. encryption with openssl	114
11.5. encryption with crypt	115
11.6. /etc/login.defs	116
11.7. chage	116
11.8. disabling a password	117
11.9. editing local files	117
11.10. practice: user passwords	118
11.11. solution: user passwords	118
12. User profiles	121
12.1. system profile	121
12.2. ~/.bash_profile	121
12.3. ~/.bash_login	122

12.4. ~/.profile	122
12.5. ~/.bashrc	122
12.6. ~/.bash_logout	123
12.7. Debian overview	123
12.8. RHEL5 overview	124
12.9. practice: user profiles	124
12.10.solution: user profiles	124
13. groups	127
13.1. groupadd	127
13.2. group file	127
13.3. groups	128
13.4. usermod	128
13.5. groupmod	128
13.6. groupdel	128
13.7. gpasswd	129
13.8. newgrp	129
13.9. vigr	130
13.10.practice: groups	130
13.11. solution: groups	130
V. Webserver; scripting 102	133
14. apache web server	135
14.1. introduction to apache	135
14.1.1. installing on Debian	135
14.1.2. installing on RHEL/CentOS	136
14.1.3. running apache on Debian	136
14.1.4. running apache on CentOS	137
14.1.5. index file on CentOS	138
14.1.6. default website	139
14.1.7. apache configuration	139
14.2. port virtual hosts on Debian	140
14.2.1. default virtual host	140
14.2.2. three extra virtual hosts	140
14.2.3. three extra ports	141
14.2.4. three extra websites	141
14.2.5. enabling extra websites	141
14.2.6. testing the three websites	142
14.3. named virtual hosts on Debian	143
14.3.1. named virtual hosts	143
14.3.2. name resolution	144
14.3.3. enabling virtual hosts	144
14.3.4. reload and verify	144
14.4. password protected website on Debian	145
14.5. port virtual hosts on CentOS	146
14.5.1. default virtual host	146
14.5.2. three extra virtual hosts	146
14.5.3. three extra ports	146
14.5.4. SELinux guards our ports	147
14.5.5. three extra websites	147
14.5.6. enabling extra websites	147
14.5.7. testing the three websites	148
14.5.8. firewall rules	149
14.6. named virtual hosts on CentOS	149
14.6.1. named virtual hosts	149
14.6.2. name resolution	150
14.6.3. reload and verify	150

14.7. password protected website on CentOS	150
14.8. troubleshooting apache	152
14.9. virtual hosts example	153
14.10.aliases and redirects	153
14.11.more on .htaccess	153
14.12.traffic	153
14.13.self signed cert on Debian	154
14.14.self signed cert on RHEL/CentOS	156
14.15.practice: apache	158
15. scripting loops	159
15.1. test []	159
15.2. if then else	160
15.3. if then elif	160
15.4. for loop	161
15.5. while loop	161
15.6. until loop	162
15.7. practice: scripting tests and loops	162
15.8. solution: scripting tests and loops	162
16. scripting parameters	165
16.1. script parameters	165
16.2. shift through parameters	166
16.3. runtime input	166
16.4. sourcing a config file	167
16.5. get script options with getopt	167
16.6. get shell options with shopt	169
16.7. practice: parameters and options	169
16.8. solution: parameters and options	169
VI. Advanced text processing	171
17. file globbing	173
17.1. * asterisk	173
17.2. ? question mark	173
17.3. [] square brackets	174
17.4. a-z and 0-9 ranges	174
17.5. \$LANG and square brackets	175
17.6. preventing file globbing	175
17.7. practice: shell globbing	175
17.8. solution: shell globbing	176
18. regular expressions	179
18.1. regex versions	179
18.2. grep	179
18.2.1. print lines matching a pattern	179
18.2.2. concatenating characters	180
18.2.3. one or the other	180
18.2.4. one or more	181
18.2.5. match the end of a string	181
18.2.6. match the start of a string	181
18.2.7. separating words	182
18.2.8. grep features	182
18.2.9. preventing shell expansion of a regex	183
18.3. rename	183
18.3.1. the rename command	183
18.3.2. perl	183
18.3.3. well known syntax	184

18.3.4. a global replace	184
18.3.5. case insensitive replace	185
18.3.6. renaming extensions	185
18.4. sed	185
18.4.1. stream editor	185
18.4.2. interactive editor	186
18.4.3. simple back referencing	186
18.4.4. back referencing	186
18.4.5. a dot for any character	186
18.4.6. multiple back referencing	186
18.4.7. white space	187
18.4.8. optional occurrence	187
18.4.9. exactly n times	187
18.4.10.between n and m times	188
18.5. bash history	188
VII.Scripting 201; job scheduling	191
19. more scripting	193
19.1. eval	193
19.2. (())	193
19.3. let	194
19.4. case	195
19.5. shell functions	195
19.6. practice : more scripting	196
19.7. solution : more scripting	197
20.background jobs	199
20.1. background processes	199
20.1.1. jobs	199
20.1.2. control-Z	199
20.1.3. & ampersand	199
20.1.4. jobs -p	200
20.1.5. fg	200
20.1.6. bg	200
20.2.practice : background processes	201
20.3.solution : background processes	201
21. scheduling	205
21.1. one time jobs with at	205
21.1.1. at	205
21.1.2. atq	205
21.1.3. atrm	206
21.1.4. at.allow and at.deny	206
21.2. cron	206
21.2.1. crontab file	206
21.2.2. crontab command	207
21.2.3. cron.allow and cron.deny	207
21.2.4. /etc/crontab	207
21.2.5. /etc/cron.*	207
21.2.6. /etc/cron.*	207
21.3. practice : scheduling	208
21.4. solution : scheduling	208

VIISSH; Docker	211
22.ssh client and server	213
22.1. about ssh	213
22.1.1. secure shell	213
22.1.2. /etc/ssh/	213
22.1.3. ssh protocol versions	213
22.1.4. public and private keys	214
22.1.5. rsa and dsa algorithms	214
22.2. log on to a remote server	214
22.3. executing a command in remote	215
22.4. scp	215
22.5. setting up passwordless ssh	215
22.5.1. ssh-keygen	216
22.5.2. ~/.ssh	216
22.5.3. id_rsa and id_rsa.pub	216
22.5.4. copy the public key to the other computer	217
22.5.5. authorized_keys	217
22.5.6. passwordless ssh	217
22.6. X forwarding via ssh	218
22.7. troubleshooting ssh	218
22.8. sshd	218
22.9. sshd keys	219
22.10.ssh-agent	219
22.11.practice: ssh	219
22.12.solution: ssh	220
A. git	223
A.1. git	223
A.2. installing git	224
A.3. starting a project	224
A.3.1. git init	225
A.3.2. git config	225
A.3.3. git add	225
A.3.4. git commit	226
A.3.5. changing a committed file	226
A.3.6. git log	227
A.3.7. git mv	227
A.4. git branches	227
A.5. to be continued...	229
A.6. github.com	229
A.7. add your public key to github	229
A.8. practice: git	229
A.9. solution: git	230
B. Introduction to vi	231
B.1. command mode and insert mode	231
B.2. start typing (a A i l o O)	231
B.3. replace and delete a character (r x X)	232
B.4. undo, redo and repeat (u .)	232
B.5. cut, copy and paste a line (dd yy p P)	232
B.6. cut, copy and paste lines (3dd 2yy)	232
B.7. start and end of a line (0 or ^ and \$)	233
B.8. join two lines (J) and more	233
B.9. words (w b)	233
B.10.save (or not) and exit (:w :q :q!)	234
B.11. Searching (/ ?)	234
B.12. replace all (:!,\$ s/foo/bar/g)	234
B.13. reading files (:r :r !cmd)	235

B.14. text buffers	235
B.15. multiple files	235
B.16. abbreviations	235
B.17. key mappings	236
B.18. setting options	236
B.19. practice: vi(m)	236
B.20. solution: vi(m)	237
C. GNU Free Documentation License	239
C.1. PREAMBLE	239
C.2. APPLICABILITY AND DEFINITIONS	239
C.3. VERBATIM COPYING	240
C.4. COPYING IN QUANTITY	241
C.5. MODIFICATIONS	241
C.6. COMBINING DOCUMENTS	242
C.7. COLLECTIONS OF DOCUMENTS	243
C.8. AGGREGATION WITH INDEPENDENT WORKS	243
C.9. TRANSLATION	243
C.10. TERMINATION	244
C.11. FUTURE REVISIONS OF THIS LICENSE	244
C.12. RELICENSING	244

Feel free to contact the author(s):

- Paul Cobbaut (Netsec BVBA): paul.cobbaut@gmail.com, <https://cobbaut.be/>
- Bert Van Vreckem (HOGENT): <http://github.com/bertvv>

Copyright 2007-2024 Netsec BVBA, Paul Cobbaut

This copy was generated on September 18, 2024.

Permission is granted to copy, distribute and/or modify this document under the terms of the **GNU Free Documentation License**, Version 1.3 or any later version published by the Free Software Foundation; with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts. A copy of the license is included in the section entitled 'GNU Free Documentation License'. # Abstract {unnumbered}

This book is used as the syllabus for the course “Linux for Data Scientists” for the Bachelor of Applied Computer Science at the HOGENT, Belgium. The contents are based on the Linux Training book series by Paul Cobbaut, with updates and additions written by the HOGENT Linux team.

This book is aimed at students specialising in the Data Engineering track that already have some basic knowledge of Linux. Where the sibling “Linux” course for students in Operations/System Administration focuses on Linux as a server operating system, this course rather discusses how Linux can be used as a platform for task and workflow automation.

More information and free .pdf available at <https://hogenttin.github.io/linux-training-hogent/>.

Part I.

First Linux VM

1. getting Linux at home

(Written by Paul Cobbaut, <https://github.com/paulcobbaut/>)

This chapter shows a Ubuntu install in Virtualbox. Consider it legacy and use CentOS7 or Debian8 instead (each have their own chapter now).

This book assumes you have access to a working Linux computer. Most companies have one or more Linux servers, if you have already logged on to it, then you're all set (skip this chapter and go to the next).

Another option is to insert a Ubuntu Linux CD in a computer with (or without) Microsoft Windows and follow the installation. Ubuntu will resize (or create) partitions and setup a menu at boot time to choose Windows or Linux.

If you do not have access to a Linux computer at the moment, and if you are unable or unsure about installing Linux on your computer, then this chapter proposes a third option: installing Linux in a virtual machine.

Installation in a virtual machine (provided by Virtualbox) is easy and safe. Even when you make mistakes and crash everything on the virtual Linux machine, then nothing on the real computer is touched.

This chapter gives easy steps and screenshots to get a working Ubuntu server in a Virtualbox virtual machine. The steps are very similar to installing Fedora or CentOS or even Debian, and if you like you can also use VMWare instead of Virtualbox.

1.1. download a Linux CD image

Start by downloading a Linux CD image (an .ISO file) from the distribution of your choice from the Internet. Take care selecting the correct cpu architecture of your computer; choose i386 if unsure. Choosing the wrong cpu type (like x86_64 when you have an old Pentium) will almost immediately fail to boot the CD.

Home Ubuntu Business Cloud TV Download Support Project Community Partners Shop ubuntu®

Ubuntu Ubuntu Server Type to search

Download Ubuntu Server

You can download Ubuntu Server now – it's completely free.

Tweet 228 Like 2k

Download Buy CDs Ubuntu Server for ARM

1 Download Ubuntu Server

Click the big orange button to download the latest version of Ubuntu. You will need to create a CD or USB stick to install Ubuntu.

Our long-term support (LTS) releases are supported for five years on the server. Perfect for organisations that need more stability for larger deployments.

Download options

Ubuntu 11.10 – Latest version

64-bit – (recommended)

Start download

Ubuntu Server 11.10
64-bit

Direct url for this download

1. getting Linux at home


1.2. download Virtualbox

Step two (when the .ISO file has finished downloading) is to download Virtualbox. If you are currently running Microsoft Windows, then download and install Virtualbox for Windows!

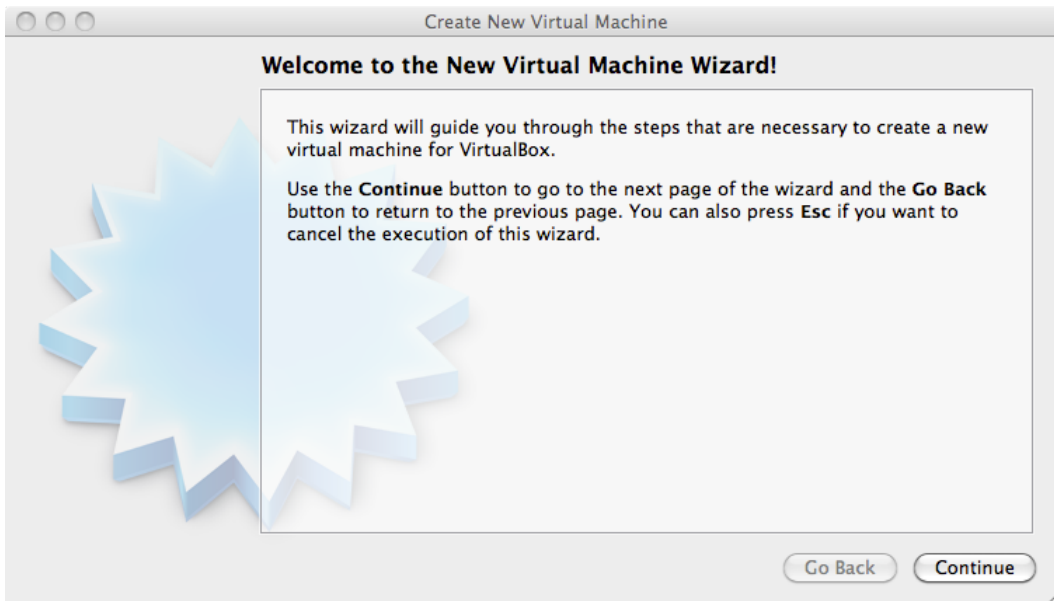


1.3. create a virtual machine

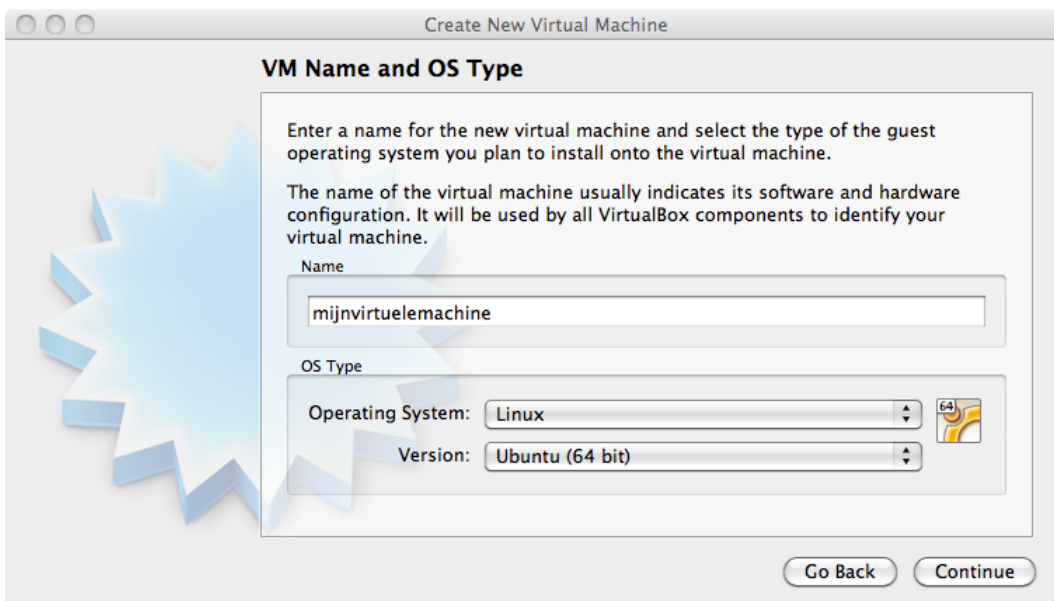
Now start Virtualbox. Contrary to the screenshot below, your left pane should be empty.



Click **New** to create a new virtual machine. We will walk together through the wizard. The screenshots below are taken on Mac OSX; they will be slightly different if you are running Microsoft Windows.




Name your virtual machine (and maybe select 32-bit or 64-bit).



Give the virtual machine some memory (512MB if you have 2GB or more, otherwise select 256MB).

1. getting Linux at home



Select to create a new disk (remember, this will be a virtual disk).



If you get the question below, choose vdi.

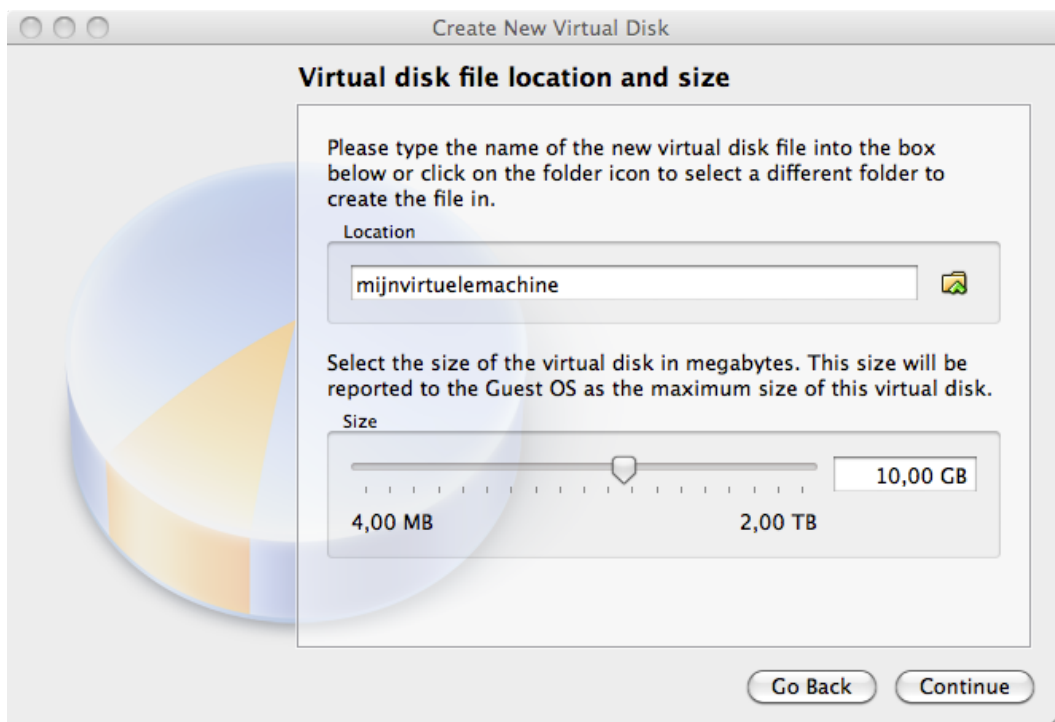


Choose dynamically allocated (fixed size is only useful in production or on really old, slow hardware).

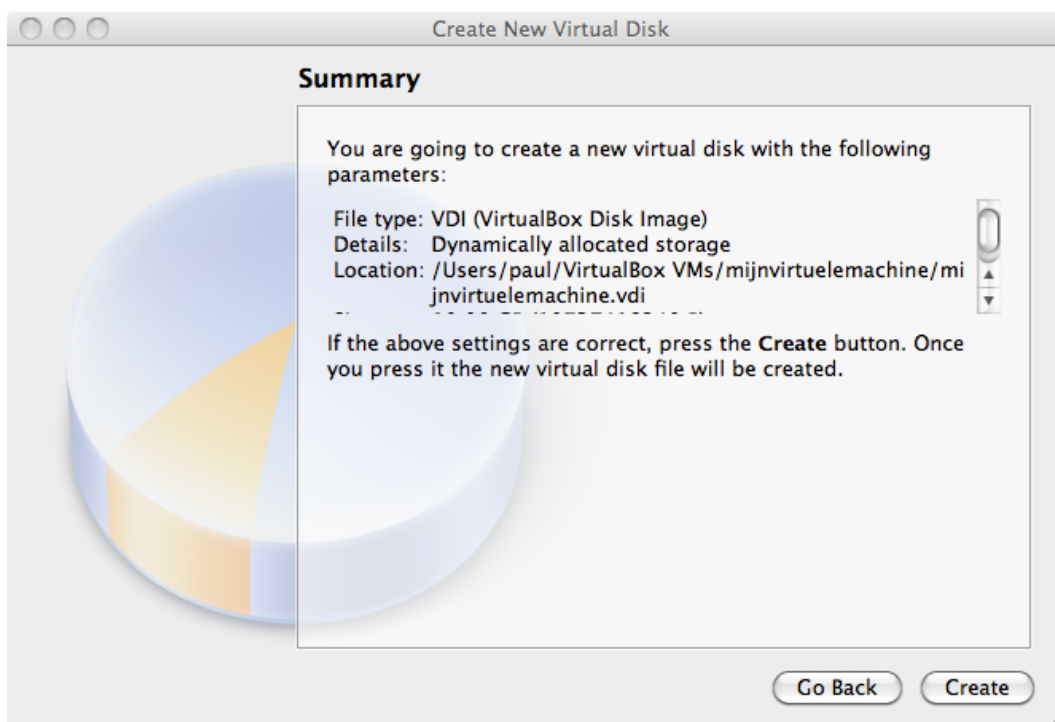


Choose between 10GB and 16GB as the disk size.


1. getting Linux at home



Click create to create the virtual disk.




Click create to create the virtual machine.



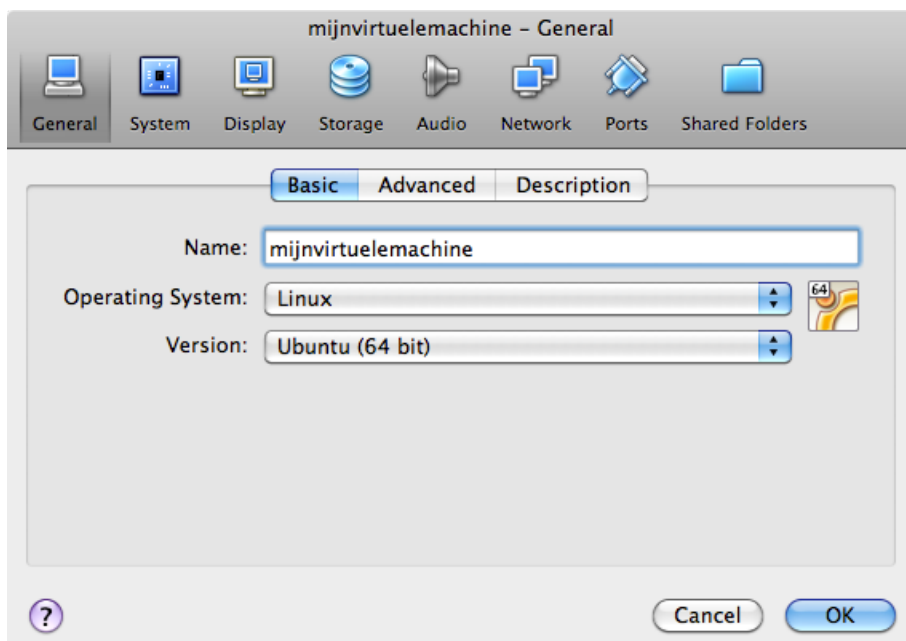
1.4. attach the CD image

Before we start the virtual computer, let us take a look at some settings (click Settings).

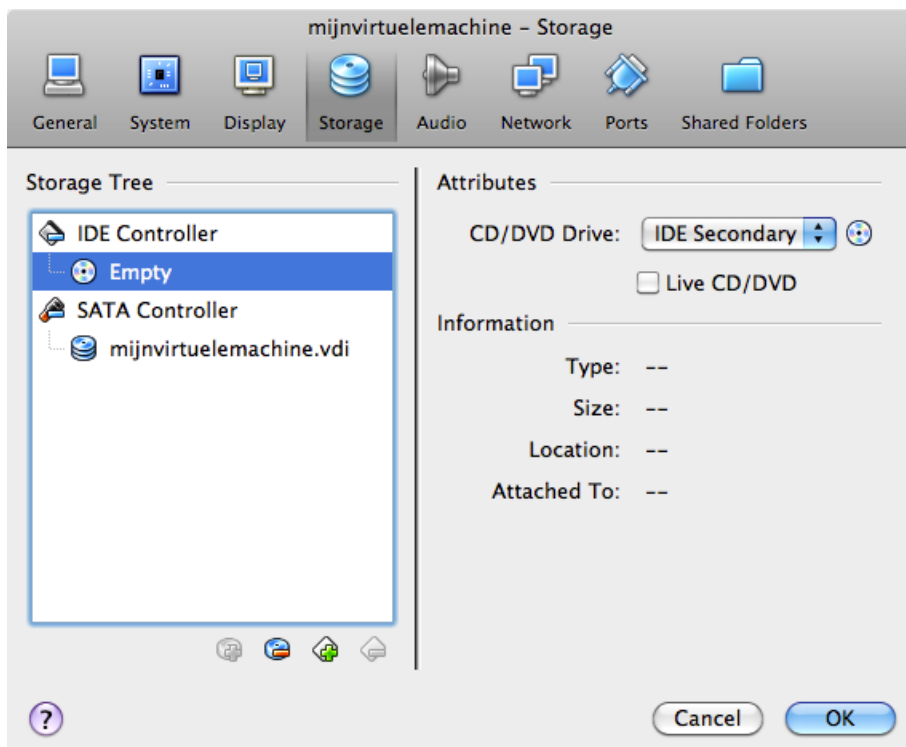


Do not worry if your screen looks different, just find the button named storage.

1. getting Linux at home

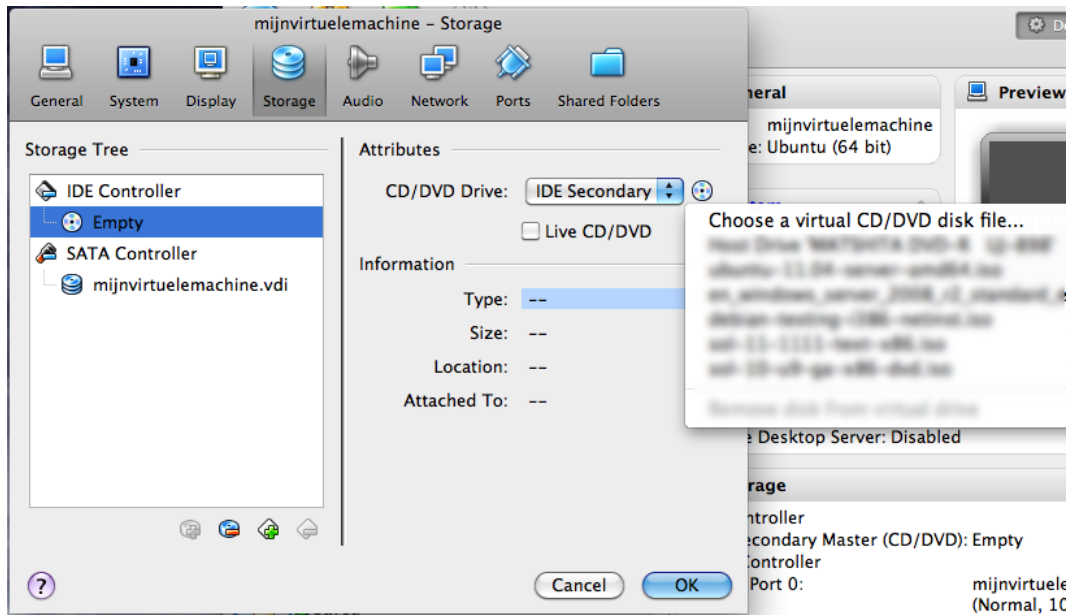


Remember the .ISO file you downloaded? Connect this .ISO file to this virtual machine by clicking on the CD icon next to Empty.




Now click on the other CD icon and attach your ISO file to this virtual CD drive.

1.4. attach the CD image

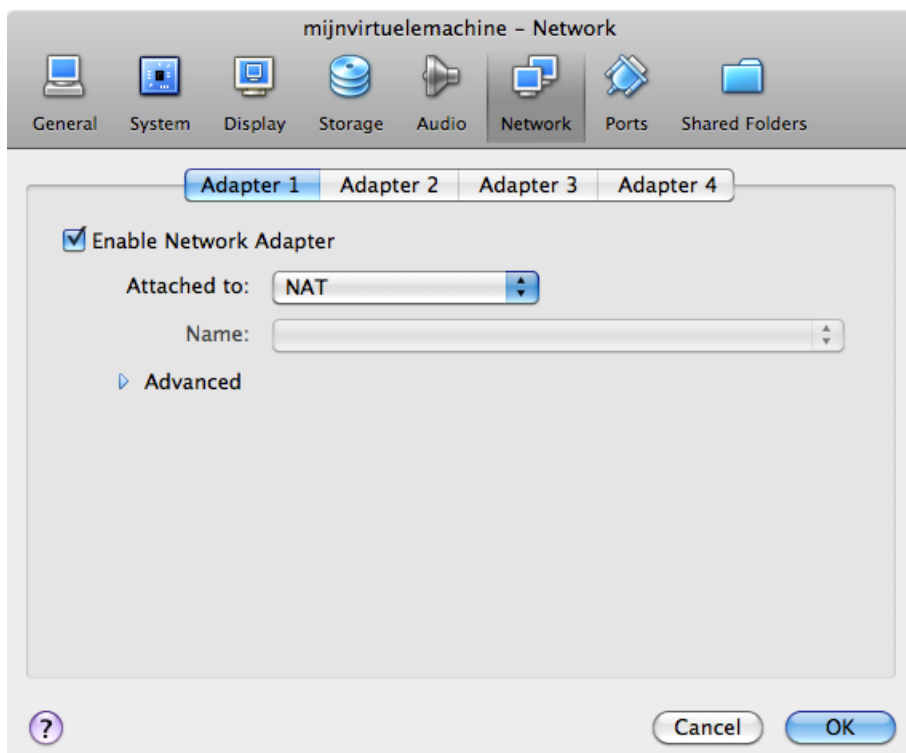


Verify that your download is accepted. If Virtualbox complains at this point, then you probably did not finish the download of the CD (try downloading it again).



It could be useful to set the network adapter to bridge instead of NAT. Bridged usually will connect your virtual computer to the Internet.

1. getting Linux at home



1.5. install Linux

The virtual machine is now ready to start. When given a choice at boot, select `install` and follow the instructions on the screen. When the installation is finished, you can log on to the machine and start practising Linux!

Part II.

Software management; curl

2. package management

(Written by Paul Cobbaut, <https://github.com/paulcobbaut/>, with contributions by: Alex M. Schapelle, <https://github.com/zero-pythagoras/>, Bert Van Vreckem <https://github.com/bertvv/>)

Most Linux distributions have a **package management** system with online **repositories** containing thousands of packages. This makes it very easy to install, update and remove applications, operating system components, documentation and much more.

We first discuss the Debian package format `.deb` and its tools `dpkg`, `apt-get` and `apt`. This should be similar on Debian, Ubuntu, Mint and all derived distributions.

Then we take a look at the Red Hat package format `.rpm` and its tools `rpm` and `dnf`. This should be similar on Red Hat, Fedora, AlmaLinux and all derived distributions.

2.1. package terminology

2.1.1. repository

A lot of software and documentation for your Linux distribution is available as **packages** in one or more centrally distributed **repositories**. The packages in such a repository are tested and very easy to install (or remove) with a graphical or command line installer.

2.1.2. .deb packages

Debian, Ubuntu, Mint and all derivatives of Debian and Ubuntu use `.deb` packages. To manage software on these systems, you can use `apt` or `apt-get`, both these tools are a front end for `dpkg`.

2.1.3. .rpm packages

Red Hat, Fedora, CentOS, OpenSUSE, Mandriva, Red Flag and others use `.rpm` packages. The tools to manage software packages on these systems are `dnf` and `rpm`.

2.1.4. dependency

Some packages need other packages to function. Tools like `apt-get`, `apt` and `dnf` will install all **dependencies** you need. When using `dpkg` or `rpm`, or when building from **source**, you will need to install dependencies yourself.

2. package management

2.1.5. open source


These repositories contain a lot of independent **open source software**. Often the source code is customized to integrate better with your distribution. Most distributions also offer this modified source code as a **package** in one or more **source repositories**.

You are free to go to the project website itself (samba.org, apache.org, github.com ...) and download the *vanilla* (= without the custom distribution changes) source code.

2.1.6. GUI software management

End users have several graphical applications available via the desktop (look for *add/remove software* or something similar).

Below a screenshot of Ubuntu Software Center running on Ubuntu 12.04. Graphical tools are not discussed in this book.



2.2. deb package management

2.2.1. about deb

Most people use apt or apt-get (APT = Advanced Package Tool) to manage their Debian/Ubuntu family of Linux distributions. Both are a front end for dpkg and are themselves a back end for *synaptic* and other graphical tools.

2.2.2. dpkg -l

The low level tool to work with .deb packages is dpkg. Among other things, you can use dpkg to list all installed packages on a Debian server.

```
student@debian:~$ dpkg -l | wc -l
365
```

Compare this to the same list on a Linux Mint system with a graphical desktop installed.

```
student@mint:~$ dpkg -l | wc -l
2118
```

2.2.3. dpkg -l \$package

Here is an example on how to get information on an individual package. The ii at the beginning means the package is installed.

```
root@debian:~# dpkg -l rsync | tail -1 | tr -s ' '
ii rsync 3.2.7-1 amd64 fast, versatile, remote (and local) file-copying tool
```

2.2.4. dpkg -S

You can find the package responsible for installing a certain file on your computer using dpkg -S. This example shows how to find the package for three files on a typical Debian server.

```
student@debian:~$ dpkg -S /usr/share/doc/tmux/ /etc/ssh/ssh_config /sbin/ifconfig
dpkg-query: no path found matching pattern /usr/share/doc/tmux/
openssh-client: /etc/ssh/ssh_config
net-tools: /sbin/ifconfig
```

2.2.5. dpkg -L

In reverse, you can also get a list of all files that have been installed by a certain program. Below is the list for the curl package.

```
student@debian:~$ dpkg -L curl
/.
/usr
/usr/bin
/usr/bin/curl
/usr/share
/usr/share/doc
/usr/share/doc/curl
/usr/share/doc/curl/changelog.Debian.gz
/usr/share/doc/curl/changelog.gz
/usr/share/doc/curl/copyright
/usr/share/man
/usr/share/man/man1
/usr/share/man/man1/curl.1.gz
/usr/share/zsh
/usr/share/zsh/vendor-completions
/usr/share/zsh/vendor-completions/_curl
```

2. package management

2.2.6. dpkg

You could use `dpkg -i` to install a package and `dpkg -r` to remove a package, but you'd have to manually download the package and keep track of dependencies. Using `apt-get` or `apt` is much easier.

2.2.7. apt-get

Debian has been using `apt-get` to manage packages since 1998. Today Debian and many Debian-based distributions still actively support `apt-get`, though some experts claim `apt`, released in 2014, is better at handling dependencies than `apt-get`.

Both commands use the same configuration files and can be used alternately; whenever you see `apt-get` in documentation, feel free to type `apt`.

We will start with `apt-get` and discuss `apt` in the next section.

2.2.8. apt-get update

When typing `apt-get update` you are downloading the names, versions and short description of all packages available on all configured repositories for your system. Remark that you need to be root to run this command.

```
student@debian:~$ apt-get update
Reading package lists ... Done
E: Could not open lock file /var/lib/apt/lists/lock - open (13: Permission denied)
E: Unable to lock directory /var/lib/apt/lists/
student@debian:~$ sudo apt-get update
Hit:1 http://security.debian.org/debian-security bookworm-security InRelease
Hit:2 http://httpredir.debian.org/debian bookworm InRelease
Hit:3 http://httpredir.debian.org/debian bookworm-updates InRelease
Reading package lists ... Done
```

In the example below you can see an interaction with an Ubuntu system. Some repositories are at the url `be.archive.ubuntu.com` because this computer was installed in Belgium. This mirror URL can be different for you.

```
student@ubuntu:~$ sudo apt-get update
Ign http://be.archive.ubuntu.com precise InRelease
Ign http://extras.ubuntu.com precise InRelease
Ign http://security.ubuntu.com precise-security InRelease
Ign http://archive.canonical.com precise InRelease
Ign http://be.archive.ubuntu.com precise-updates InRelease
...
Hit http://be.archive.ubuntu.com precise-backports/main Translation-en
Hit http://be.archive.ubuntu.com precise-backports/multiverse Translation-en
Hit http://be.archive.ubuntu.com precise-backports/restricted Translation-en
Hit http://be.archive.ubuntu.com precise-backports/universe Translation-en
Fetched 13.7 MB in 8s (1682 kB/s)
Reading package lists ... Done
student@ubuntu:~$
```

Tips:

- Run `apt-get update` every time before performing other package operations to ensure your metadata is up-to-date.
- Since the package repositories are hosted on web servers, you can open any repository URL in your browser to see how the repository is structured.

2.2.9. apt-get upgrade

One of the nicest features of `apt-get` is that it allows for a secure update of *all software currently installed* on your computer with just *one* command.

```
student@debian:~$ sudo apt-get upgrade
Reading package lists... Done
Building dependency tree
Reading state information... Done
0 upgraded, 0 newly installed, 0 to remove and 0 not upgraded.
```

The above transcript shows that all software is updated to the latest version available for my distribution. Below is an example of a system with software that can be updated. Some lines were omitted for brevity.

```
student@debian:~$ sudo apt-get upgrade
Reading package lists... Done
Building dependency tree... Done
Reading state information... Done
Calculating upgrade... Done
The following packages have been kept back:
  linux-image-amd64
The following packages will be upgraded:
  base-files bind9-dnsutils bind9-host bind9-libs cryptsetup cryptsetup-
bin libcryptsetup12 libgnutls30 libnss-systemd libpam-systemd libsystemd-
shared libsystemd0 libudev1 systemd systemd-sysv
  systemd-timesyncd tar tzdata udev usr-is-merged
20 upgraded, 0 newly installed, 0 to remove and 1 not upgraded.
Need to get 13.0 MB of archives.
After this operation, 75.8 kB of additional disk space will be used.
Do you want to continue? [Y/n] y
Get:1 http://security.debian.org/debian-security bookworm-security/main amd64 bind9-
host amd64 1:9.18.24-1 [305 kB]
[ ... ]
Get:20 http://httpredir.debian.org/debian bookworm/main amd64 cryptsetup amd64 2:2.6.1-
4~deb12u2 [213 kB]
Fetched 13.0 MB in 1s (20.3 MB/s)
Reading changelogs... Done
Preconfiguring packages ...
(Reading database ... 29205 files and directories currently installed.)
Preparing to unpack .../base-files_12.4+deb12u5_amd64.deb ...
Unpacking base-files (12.4+deb12u5) over (12.4+deb12u4) ...
Setting up base-files (12.4+deb12u5) ...
Installing new version of config file /etc/debian_version ...
[ ... ]
Preparing to unpack .../5-cryptsetup_2%3a2.6.1-4~deb12u2_amd64.deb ...
Unpacking cryptsetup (2:2.6.1-4~deb12u2) over (2:2.6.1-4~deb12u1) ...
Setting up systemd-sysv (252.22-1~deb12u1) ...
[ ... ]
Setting up bind9-dnsutils (1:9.18.24-1) ...
Processing triggers for initramfs-tools (0.142) ...
update-initramfs: Generating /boot/initrd.img-6.1.0-17-amd64
[ ... ]
Processing triggers for mailcap (3.70+nmu1) ...
```

Tip: Have you noticed that almost every time that you update software on Windows, you are asked to reboot your computer? This is **not** the case with Linux! The only time you need to reboot is when you update the kernel.

2. package management

2.2.10. apt-get clean

`apt-get` keeps a copy of downloaded packages in `/var/cache/apt/archives`, as can be seen in this screenshot.

```
student@debian:~$ ls /var/cache/apt/archives/ | head
base-files_12.4+deb12u5_amd64.deb
bind9-dnsutils_1%3a9.18.24-1_amd64.deb
bind9-host_1%3a9.18.24-1_amd64.deb
bind9-libs_1%3a9.18.24-1_amd64.deb
cryptsetup_2%3a2.6.1-4~deb12u2_amd64.deb
cryptsetup-bin_2%3a2.6.1-4~deb12u2_amd64.deb
libcryptsetup12_2%3a2.6.1-4~deb12u2_amd64.deb
libgnutls30_3.7.9-2+deb12u2_amd64.deb
libnss-systemd_252.22-1~deb12u1_amd64.deb
libpam-systemd_252.22-1~deb12u1_amd64.deb
```

Running `apt-get clean` removes all `.deb` files from that directory.

```
student@debian:~$ sudo apt-get clean
student@debian:~$ ls /var/cache/apt/archives/*.deb
ls: cannot access /var/cache/apt/archives/*.deb: No such file or directory
```

2.2.11. apt-cache search

Use `apt-cache search` to search for availability of a package. Here we look for `rsync`.

```
student@debian:~$ apt-cache search rsync | grep '^rsync'
rsync - fast, versatile, remote (and local) file-copying tool
rsyncrypto - rsync friendly encryption
```

2.2.12. apt-get install

You can install one or more applications by appending their name behind `apt-get install`. The following example shows how to install the `tftpd-hpa` package (a TFTP server).

```
student@debian:~$ sudo apt-get install tftpd-hpa
Reading package lists... Done
Building dependency tree... Done
Reading state information... Done
Suggested packages:
  pxelinux
The following NEW packages will be installed:
  tftpd-hpa
0 upgraded, 1 newly installed, 0 to remove and 1 not upgraded.
Need to get 41.9 kB of archives.
After this operation, 117 kB of additional disk space will be used.
Get:1 http://htpredirect.debian.org/debian bookworm/main amd64 tftpd-hpa amd64 5.2+20150808-1.4 [41.9 kB]
Fetched 41.9 kB in 0s (241 kB/s)
Preconfiguring packages ...
Selecting previously unselected package tftpd-hpa.
(Reading database ... 29179 files and directories currently installed.)
Preparing to unpack .../tftpd-hpa_5.2+20150808-1.4_amd64.deb ...
```

```
Unpacking tftpd-hpa (5.2+20150808-1.4) ...
Setting up tftpd-hpa (5.2+20150808-1.4) ...
Processing triggers for man-db (2.11.2-2) ...
```

The `apt-get` command will ask the user to confirm the installation of the package by pressing “y” and ENTER. You can use the `-y` option to automatically answer yes to all questions.

The following example installs the `vim` package (VI iMproved, a powerful text editor for the terminal). **Remark** that some additional packages are installed as dependencies!

```
student@debian:~$ sudo apt-get install -y vim
Reading package lists... Done
Building dependency tree... Done
Reading state information... Done
The following additional packages will be installed:
  libgpm2 libsodium23 vim-runtime
Suggested packages:
  gpm ctags vim-doc vim-scripts
The following NEW packages will be installed:
  libgpm2 libsodium23 vim vim-runtime
0 upgraded, 4 newly installed, 0 to remove and 1 not upgraded.
Need to get 8,768 kB of archives.
After this operation, 41.5 MB of additional disk space will be used.
[ ... ]
Setting up libsodium23:amd64 (1.0.18-1) ...
Setting up libgpm2:amd64 (1.20.7-10+b1) ...
Setting up vim-runtime (2:9.0.1378-2) ...
Setting up vim (2:9.0.1378-2) ...
[ ... ]
Processing triggers for man-db (2.11.2-2) ...
Processing triggers for libc-bin (2.36-9+deb12u4) ...
```

2.2.13. apt-get remove

You can remove one or more applications by appending their name behind `apt-get remove`.

```
student@debian:~$ sudo apt-get remove tftpd-hpa
Reading package lists... Done
Building dependency tree... Done
Reading state information... Done
The following packages will be REMOVED:
  tftpd-hpa
0 upgraded, 0 newly installed, 1 to remove and 1 not upgraded.
After this operation, 117 kB disk space will be freed.
Do you want to continue? [Y/n] y
(Reading database ... 29194 files and directories currently installed.)
Removing tftpd-hpa (5.2+20150808-1.4) ...
Processing triggers for man-db (2.11.2-2) ...
```

If we use `dpkg -l` to check the status of the `tftpd-hpa` package, we see that it is removed, some configuration (`rc`) files are left on the system. Indeed, the configuration file `/etc/init/tftpd-hpa.conf` is not removed! We'll solve this in the next section.

2. package management

```
student@debian:~$ dpkg -l tftpd-hpa | tail -1
rc tftpd-hpa      5.2+20150808-1.4 amd64      HPA's tftp server
student@debian:~$ ls -l /etc/init/tftpd-hpa.conf
-rw-r--r-- 1 root root 980 Oct 25  2022 /etc/init/tftpd-hpa.conf
```

The example below shows how to remove the vim package. Note that dependencies are **not** removed! You can execute `sudo apt autoremove` afterwards (as is suggested by the output of the command!) to remove those as well.

```
student@debian:~$ sudo apt-get remove vim
Reading package lists... Done
Building dependency tree... Done
Reading state information... Done
The following packages were automatically installed and are no longer required:
  libsodium23 vim-runtime
Use 'sudo apt autoremove' to remove them.
The following packages will be REMOVED:
  vim
0 upgraded, 0 newly installed, 1 to remove and 1 not upgraded.
After this operation, 3,738 kB disk space will be freed.
Do you want to continue? [Y/n] y
(Reading database ... 31257 files and directories currently installed.)
Removing vim (2:9.0.1378-2) ...
[ ... ]
student@debian:~$ sudo apt-get autoremove
Reading package lists... Done
Building dependency tree... Done
Reading state information... Done
The following packages will be REMOVED:
  libsodium23 vim-runtime
0 upgraded, 0 newly installed, 2 to remove and 1 not upgraded.
After this operation, 37.7 MB disk space will be freed.
Do you want to continue? [Y/n] y
(Reading database ... 31247 files and directories currently installed.)
Removing libsodium23:amd64 (1.0.18-1) ...
Removing vim-runtime (2:9.0.1378-2) ...
Removing 'diversion of /usr/share/vim/vim90/doc/help.txt to /usr/share/vim/vim90/doc/help.tiny by vim-runtime'
Removing 'diversion of /usr/share/vim/vim90/doc/tags to /usr/share/vim/vim90/doc/tags.vim-tiny by vim-runtime'
Processing triggers for man-db (2.11.2-2) ...
Processing triggers for libc-bin (2.36-9+deb12u4) ...
```

2.2.14. apt-get purge

You can purge one or more applications by appending their name behind `apt-get purge`. Purging will also remove all existing configuration files related to that application. The screenshot shows how to purge the `tftpd-hpa` package.

```
student@debian:~$ ls -l /etc/init/tftpd-hpa.conf
-rw-r--r-- 1 root root 980 Oct 25  2022 /etc/init/tftpd-hpa.conf
student@debian:~$ sudo apt-get purge tftpd-hpa
Reading package lists... Done
Building dependency tree... Done
Reading state information... Done
The following packages will be REMOVED:
```

```
tftpd-hpa*
0 upgraded, 0 newly installed, 1 to remove and 1 not upgraded.
After this operation, 0 B of additional disk space will be used.
Do you want to continue? [Y/n] y
(Reading database ... 29182 files and directories currently installed.)
Purging configuration files for tftpd-hpa (5.2+20150808-1.4) ...
student@debian:~$ ls -l /etc/init/tftpd-hpa.conf
ls: cannot access '/etc/init/tftpd-hpa.conf': No such file or directory
```

Note that dpkg has no information about a purged package!

```
student@debian:~$ dpkg -l tftpd-hpa | tail -1 | tr -s ' '
dpkg-query: no packages found matching tftpd-hpa
```

2.2.15. apt

Nowadays, most people use apt for package management on Debian, Mint and Ubuntu systems. That does not mean that apt-get is no longer useful. In scripts, it is actually recommended to use apt-get because its options and behaviour are more stable and predictable than apt. For interactive use, apt is more user-friendly.

To synchronize with the repositories.

```
sudo apt update
```

To patch and upgrade all software to the latest version on Debian.

```
sudo apt upgrade
```

To patch and upgrade all software to the latest version on Ubuntu and Mint.

```
sudo apt safe-upgrade
```

To install an application with all dependencies.

```
sudo apt install $package
```

To search the repositories for applications that contain a certain string in their name or description.

```
apt search $string
```

To remove an application.

```
sudo apt remove $package
```

To remove an application and all configuration files.

```
sudo apt purge $package
```

2.2.16. /etc/apt/sources.list

Both `apt-get` and `apt` use the same configuration information in `/etc/apt/`. The main configuration file is `/etc/apt/sources.list` and the directory `/etc/apt/sources.list.d/` contains additional files. These contain a list of `http` or `ftp` sources where packages for the distribution can be downloaded. Third party software vendors may provide their own package repositories for Debian or Ubuntu. These repositories are typically added through a new file in `/etc/apt/sources.list.d/`.

This is what that list looks like on a Debian server system shortly after installation.

```
student@debian:~$ cat /etc/apt/sources.list
deb http://httpredir.debian.org/debian/ bookworm main non-free-firmware
deb-src http://httpredir.debian.org/debian/ bookworm main non-free-firmware

deb http://security.debian.org/debian-security bookworm-security main non-free-firmware
deb-src http://security.debian.org/debian-security bookworm-security main non-free-firmware

# bookworm-updates, to get updates before a point release is made;
deb http://httpredir.debian.org/debian/ bookworm-updates main non-free-firmware
deb-src http://httpredir.debian.org/debian/ bookworm-updates main non-free-firmware
```

If you use Linux as a daily driver, you may end up with a repository list with many more entries, like on this Ubuntu system:

```
student@ubuntu:~$ wc -l /etc/apt/sources.list
63 /etc/apt/sources.list
```

There is much more to learn about `apt`, explore commands like `add-apt-repository`, `apt-key` and `apropos apt`.

2.3. the Red Hat package manager (rpm)

On Red Hat and other distros of that family, the *Red Hat package manager* (RPM) is used to install, upgrade and remove software. There's a basic command, `rpm`, and a more advanced tool, `dnf` (comparable with the situation on Debian-based systems, where `dpkg` is the basic tool and `apt` the more advanced one). When you install a graphical desktop, there's also a GUI tool for package management, but we won't be discussing that here.

Software distributed in the `rpm` format will have a file name following this format: `package-version-release.architecture.rpm`. For example, the package name `openssh-server-8.7p1-34.el9.x86_64.rpm` has the following components:

- package name: `openssh-server`
- version: `8.7p1`
- release: `34.el9` (`el9` stands for Enterprise Linux 9, indicating it is compatible with RHEL 9)
- architecture: `x86_64` (suitable for a 64-bit Intel/AMD processor)

We will start with discussing the `dnf` command, since that one is most commonly used. After that, we'll show how to use the `rpm` command.

2.3.1. dnf

The name of the `dnf` command has a bit of a convoluted history. It stands for “Dandified Yum”, and is a fork/improvement of the `yum` package manager command. Yum stands for *Yellowdog Updater, Modified*, and was originally developed for the now defunct Yellow Dog Linux distribution (for the IBM POWER7 processor). Red Hat started using it in RHEL 5 and it was the default package manager for Red Hat and its derivatives for many years. However, more recently, they developed `dnf` to replace `yum` with the former now being the default package manager for Fedora, Red Hat Enterprise Linux and its derivatives.

The `dnf` command works quite similarly to the `apt` command on Debian-based systems. It has similar subcommands, which we will discuss in the next sections. However, an equivalent for `apt update` does *not* exist. The `dnf` command will automatically update its package database whenever you execute it.

2.3.2. dnf list

Issue `dnf list` to see a list of all packages that DNF knows about.

```
[student@el ~]$ dnf list | wc -l
6751
[student@el ~]$ dnf list --all | wc -l
6751
```

Add the option `--available` or `--installed` to see only the packages that are available for installation or installed on the system.

```
[student@el ~]$ dnf list --available | wc -l
6392
[student@el ~]$ dnf list --installed | wc -l
353
```

Issue `dnf list $package` to get all versions (in different repositories) of one package.

```
[student@el ~]$ dnf list kernel
Last metadata expiration check: 0:12:15 ago on Sun 25 Feb 2024 07:16:59 PM UTC.
Installed Packages
kernel.x86_64                5.14.0-362.8.1.el9_3      @anaconda
kernel.x86_64                5.14.0-362.13.1.el9_3     @baseos
Available Packages
kernel.x86_64                5.14.0-362.18.1.el9_3     baseos
```

2.3.3. dnf search

To search for a package containing a certain string in the description or name use `dnf search $string`.

```
[student@el ~]$ dnf search openssh
Last metadata expiration check: 0:15:35 ago on Sun 25 Feb 2024 07:16:59 PM UTC.
===== Name Exactly Matched: openssh =====
openssh.x86_64 : An open source implementation of SSH protocol version 2
===== Name & Summary Matched: openssh =====
openssh-askpass.x86_64 : A passphrase dialog for OpenSSH and X
openssh-keycat.x86_64 : A mls keycat backend for openssh
===== Name Matched: openssh =====
```

2. package management

```
openssh-clients.x86_64 : An open source SSH client applications
openssh-server.x86_64 : An open source SSH server daemon
[student@el ~]$ dnf search epel
Last metadata expiration check: 0:18:51 ago on Sun 25 Feb 2024 07:16:59 PM UTC.
===== Name Matched: epel =====
epel-release.noarch : Extra Packages for Enterprise Linux repository configuration
```

2.3.4. dnf info

Information about a specific package can be obtained with `dnf info $package`.

```
[student@el ~]$ dnf info epel-release
Last metadata expiration check: 1:15:53 ago on Sun 25 Feb 2024 07:55:24 PM UTC.
Installed Packages
Name           : epel-release
Version        : 9
Release        : 7.el9
Architecture   : noarch
Size           : 26 k
Source         : epel-release-9-7.el9.src.rpm
Repository     : @System
From repo      : epel
Summary        : Extra Packages for Enterprise Linux repository configuration
URL            : http://download.fedoraproject.org/pub/epel
License        : GPLv2
Description    : This package contains the Extra Packages for Enterprise Linux
                : (EPEL) repository GPG key as well as configuration for yum.
```

This gives you a lot of information about the package, including the version, release, architecture, size, source, repository, summary, link to the project website, license and description.

If the repository is indicated as `@System`, it means that the package is installed. Otherwise, it would show the name of the repository from which the package would be installed.

```
[student@el ~]$ dnf info zork
Last metadata expiration check: 1:19:14 ago on Sun 25 Feb 2024 07:55:24 PM UTC.
Available Packages
Name           : zork
Version        : 1.0.3
Release        : 5.el9
Architecture   : x86_64
Size           : 179 k
Source         : zork-1.0.3-5.el9.src.rpm
Repository     : epel
Summary        : Public Domain original DUNGEON game (Zork I)
URL            : https://github.com/devshane/zork
License        : Public Domain
Description    : Public Domain source code to the original DUNGEON game (Zork I).
[ ... ]
```

2.3.5. dnf install

To install an application, use `dnf install $package`. Naturally, `dnf` will install all the necessary dependencies.

2.3. the Red Hat package manager (rpm)

```
[student@el ~]$ sudo dnf install epel-release
Last metadata expiration check: 2:07:04 ago on Sun 25 Feb 2024 05:32:50 PM UTC.
Dependencies resolved.
```

```
=====
Package           Architecture    Version          Repository      Size
=====
Installing:
epel-release      noarch         9-5.el9         extras          18 k
=====
```

Transaction Summary

```
=====
Install 1 Package
=====
```

Total download size: 18 k

Installed size: 25 k

Is this ok [y/N]: y

Downloading Packages:

```
epel-release-9-5.el9.noarch.rpm          62 kB/s | 18 kB    00:00
-----
```

```
-----
Total                                     23 kB/s | 18 kB    00:00
```

Running transaction check

Transaction check succeeded.

Running transaction test

Transaction test succeeded.

Running transaction

```
  Preparing      :                               1/1
```

```
  Installing     : epel-release-9-5.el9.noarch  1/1
```

```
  Running scriptlet: epel-release-9-5.el9.noarch 1/1
```

Many EPEL packages require the CodeReady Builder (CRB) repository.

It is recommended that you run `/usr/bin/crb enable` to enable the CRB repository.

```
  Verifying      : epel-release-9-5.el9.noarch  1/1
```

Installed:

```
epel-release-9-5.el9.noarch
```

Complete!

Add the option `-y` to skip confirmation. If the package is already installed, `install` will upgrade the package to the latest version.

```
[student@el ~]$ sudo dnf install -y sudo
Last metadata expiration check: 0:01:45 ago on Sun 25 Feb 2024 07:43:07 PM UTC.
Package sudo-1.9.5p2-9.el9.x86_64 is already installed.
Dependencies resolved.
```

```
=====
Package           Architecture    Version          Repository      Size
=====
Upgrading:
sudo              x86_64         1.9.5p2-10.el9_3 baseos          1.0 M
=====
```

Transaction Summary

```
=====
Upgrade 1 Package
=====
```

Total download size: 1.0 M

Downloading Packages:

2. package management

```
sudo-1.9.5p2-10.el9_3.x86_64.rpm          3.0 MB/s | 1.0 MB      00:00
-----
Total                                    1.3 MB/s | 1.0 MB      00:00
Running transaction check
Transaction check succeeded.
Running transaction test
Transaction test succeeded.
Running transaction
  Preparing      :                                1/1
  Upgrading      : sudo-1.9.5p2-10.el9_3.x86_64  1/2
  Running scriptlet: sudo-1.9.5p2-10.el9_3.x86_64  1/2
  Cleanup        : sudo-1.9.5p2-9.el9.x86_64     2/2
  Running scriptlet: sudo-1.9.5p2-9.el9.x86_64     2/2
  Verifying      : sudo-1.9.5p2-10.el9_3.x86_64  1/2
  Verifying      : sudo-1.9.5p2-9.el9.x86_64     2/2
```

```
Upgraded:
  sudo-1.9.5p2-10.el9_3.x86_64
```

Complete!

You can add more than one parameter here.

```
[student@el ~]$ sudo dnf install httpd mod_ssl mariadb-server php php-mysqldb
```

2.3.6. dnf upgrade

To bring all applications up to date by downloading and installing them, issue `dnf upgrade`. All software that was installed via `dnf` will be updated to the latest version that is available in the repository.

```
[student@el ~]$ sudo dnf upgrade
Last metadata expiration check: 0:05:19 ago on Sun 25 Feb 2024 07:43:07 PM UTC.
Dependencies resolved.
```

```
=====
Package                Arch    Version                               Repository  Size
=====
Installing:
kernel                  x86_64  5.14.0-362.18.1.el9_3                baseos      9.4 k
Upgrading:
epel-release            noarch  9-7.el9                               epel        19 k
gnutls                  x86_64  3.7.6-23.el9_3.3                     baseos      1.0 M
[ ... ]
```

```
Transaction Summary
```

```
=====
Install  10 Packages
Upgrade  12 Packages
```

```
Total download size: 89 M
```

```
Is this ok [y/N]: y
```

```
Downloading Packages:
```

```
(1/22): graphite2-1.3.14-9.el9.x86_64.rpm    189 kB/s | 94 kB      00:00
(2/22): freetype-2.10.4-9.el9.x86_64.rpm    752 kB/s | 387 kB     00:00
```

```
[ ... ]
```

Complete!

If you only want to update one package, use `dnf upgrade $package`. It behaves the same as `dnf install $package`.

2.3.7. dnf provides

To search for a package containing a certain file use `dnf provides $filename` (or globbing pattern). This is especially useful if you want to install a specific command that has a different name than the package name. For example, say that you've heard about the `ag` command that is a faster alternative to `grep`. The command `dnf search ag` spews out too much output, so no useful results:

```
[student@el ~]$ dnf search ag | wc -l
Last metadata expiration check: 0:02:48 ago on Sun 25 Feb 2024 07:55:24 PM UTC.
2979
```

Listing available packages with `ag` shows that there is no such package:

```
[student@el ~]$ dnf list --available ag
Last metadata expiration check: 0:04:05 ago on Sun 25 Feb 2024 07:55:24 PM UTC.
Error: No matching Packages to list
[student@el ~]$ dnf list --available ag*
Last metadata expiration check: 0:04:09 ago on Sun 25 Feb 2024 07:55:24 PM UTC.
Available Packages
Agda.x86_64                2.6.2.2-36.el9                epel
Agda-common.noarch        2.6.2.2-36.el9                epel
aggregate6.noarch        1.0.12-2.el9                  epel
agrep.x86_64              0.8.0-34.20140228gitc2f5d13.el9 epel
```

The last package looks promising, but it's not the one we're looking for. So let's use `dnf provides` to find out which package contains the `ag` command. If the command is `ag`, we expect that it is installed in one of the `bin/` directories, i.e. `/bin/`, `/usr/bin/`, `/sbin/`, `/usr/sbin/`, `/usr/local/bin/`, `/usr/local/sbin/`, `/usr/local/bin/`, `/usr/local/sbin/`. We can summarize the possible path names with globbing pattern `*bin/ag`:

```
[student@el ~]$ dnf provides *bin/ag
Last metadata expiration check: 0:07:13 ago on Sun 25 Feb 2024 07:55:24 PM UTC.
the_silver_searcher-2.2.0^2020704.5a1c8d8-3.el9.x86_64 : Super-fast text
                                                         : searching tool (ag)

Repo          : epel
Matched from:
Other         : *bin/ag
[student@el ~]$ sudo dnf install -y the_silver_searcher
```

So the name of the package is `the_silver_searcher` (`ag` being the chemical symbol for silver) and it is provided by the EPEL repository (Extra Packages for Enterprise Linux). We can install it with `dnf install the_silver_searcher`.

2.3.8. dnf remove

Removing a package is done with `dnf remove $package`. This will remove the package and all its dependencies that are not needed by other packages.

2. package management

```
[student@el ~]$ sudo dnf remove net-tools
Dependencies resolved.
```

```
=====
Package      Arch      Version                               Repository      Size
=====
Removing:
net-tools    x86_64    2.0-0.62.20160912git.el9             @anaconda      912 k
=====
```

Transaction Summary

```
=====
Remove 1 Package
=====
```

Freed space: 912 k

Is this ok [y/N]: y

Running transaction check

Transaction check succeeded.

Running transaction test

Transaction test succeeded.

Running transaction

```
  Preparing      :                               1/1
  Erasing       : net-tools-2.0-0.62.20160912git.el9.x86_64 1/1
  Verifying     : net-tools-2.0-0.62.20160912git.el9.x86_64 1/1
```

Removed:

```
net-tools-2.0-0.62.20160912git.el9.x86_64
```

Complete!

By the way, this package, `net-tools`, contains commands that are considered to be obsolete and have been replaced by other, newer implementations. You don't really need it, so it's a good example for this section. If you removed it, feel free to reinstall it if you want!

2.3.9. dnf software groups

Issue `dnf grouplist` to see a list of all available software groups.

```
[student@el ~]$ dnf grouplist
```

Last metadata expiration check: 1:00:37 ago on Sun 25 Feb 2024 07:55:24 PM UTC.

Available Environment Groups:

```
  Server with GUI
  Server
  Minimal Install
  Workstation
  KDE Plasma Workspaces
  Virtualization Host
  Custom Operating System
```

Available Groups:

```
  RPM Development Tools
  .NET Development
  Container Management
  Console Internet Tools
  Graphical Administration Tools
  Scientific Support
  Headless Management
  Smart Card Support
  Legacy UNIX Compatibility
  Security Tools
```

```

Network Servers
System Tools
Development Tools
Fedora Packager
VideoLAN Client
Xfce

```

To install a set of applications, brought together via a group, use `yum groupinstall $group-name`.

```
[student@el ~]$ sudo dnf groupinstall 'Security Tools'
```

```
Last metadata expiration check: 1:00:35 ago on Sun 25 Feb 2024 08:03:34 PM UTC.
Dependencies resolved.
```

```

=====
Package                Arch      Version                               Repository  Size
=====
Installing group/module packages:
scap-security-guide    noarch   0.1.69-3.el9_3.alma.1               appstream   813 k
Installing dependencies:
libtool-ltdl           x86_64   2.4.6-45.el9                        appstream   36 k
libxslt                x86_64   1.1.34-9.el9                        appstream   240 k
openscap               x86_64   1:1.3.8-1.el9_2.alma.2             appstream   1.9 M
openscap-scanner       x86_64   1:1.3.8-1.el9_2.alma.2             appstream   57 k
xml-common             noarch   0.6.3-58.el9                        appstream   31 k
xmlsec1                x86_64   1.2.29-9.el9                        appstream   189 k
xmlsec1-openssl       x86_64   1.2.29-9.el9                        appstream   90 k
Installing Groups:
Security Tools

```

```
Transaction Summary
```

```
=====
Install 8 Packages
```

```
Total download size: 3.3 M
```

```
Installed size: 103 M
```

```
Is this ok [y/N]:
```

```
[ ... ]
```

Read the manual page of `dnf` for more information about managing groups in `dnf`. In practice, chances are that you won't need this feature very often.

2.3.10. rpm -qa

In the following sections, we'll show what you can do with the `rpm` command.

To obtain a list of all installed software, use the `rpm -qa` command.

```

[student@el ~]$ rpm -qa | grep ssh
libssh-config-0.10.4-11.el9.noarch
libssh-0.10.4-11.el9.x86_64
openssh-8.7p1-34.el9.x86_64
openssh-clients-8.7p1-34.el9.x86_64
openssh-server-8.7p1-34.el9.x86_64

```

2. package management

2.3.11. rpm -q

To verify whether one package is installed, use `rpm -q`.

```
[student@el ~]$ rpm -q vim-enhanced
package vim-enhanced is not installed
[student@el ~]$ rpm -q vim-minimal
vim-minimal-8.2.2637-20.el9_1.x86_64
[student@el ~]$ rpm -q kernel
kernel-5.14.0-362.8.1.el9_3.x86_64
kernel-5.14.0-362.13.1.el9_3.x86_64
kernel-5.14.0-362.18.1.el9_3.x86_64
```

2.3.12. rpm -ql

To see which files are installed by a package, use `rpm -ql`.

```
[student@el ~]$ rpm -ql vim-minimal
/etc/virc
/usr/bin/ex
/usr/bin/rvi
/usr/bin/rview
/usr/bin/vi
/usr/bin/view
/usr/lib/.build-id
/usr/lib/.build-id/c6
/usr/lib/.build-id/c6/aa3d8d79f09dd48e99475c332bed4df39d76e1
/usr/libexec/vi
/usr/share/man/man1/ex.1.gz
/usr/share/man/man1/rvi.1.gz
/usr/share/man/man1/rview.1.gz
/usr/share/man/man1/vi.1.gz
/usr/share/man/man1/view.1.gz
/usr/share/man/man5/virc.5.gz
```

2.3.13. rpm -Uvh

To install or upgrade a package, use the `-Uvh` switches. The `-U` switch is the same as `-i` for install, except that older versions of the software are removed. The `-vh` switches are for nicer output.

You would typically use this command to install an `.rpm` package that you have downloaded from the internet. Beware, though, that `rpm` does not resolve dependencies, so you might need to install other packages first.

```
[student@el ~]$ sudo rpm -Uvh ./htop-3.3.0-1.el9.x86_64.rpm
error: Failed dependencies:
    libhwloc.so.15()(64bit) is needed by htop-3.3.0-1.el9.x86_64
```

2.3.14. rpm -e

To remove a package, use the `-e` switch.

```
[student@el ~]$ rpm -q net-tools
net-tools-2.0-0.62.20160912git.el9.x86_64
[student@el ~]$ sudo rpm -e net-tools
[student@el ~]$ rpm -q net-tools
package net-tools is not installed
```

`rpm -e` verifies dependencies, and thus will prevent you from accidentally erasing packages that are needed by other packages.

```
[student@el ~]$ sudo rpm -e slang
error: Failed dependencies:
    libslang.so.2()(64bit) is needed by (installed) newt-0.52.21-
11.el9.x86_64
    libslang.so.2(SLANG2)(64bit) is needed by (installed) newt-0.52.21-
11.el9.x86_64
```

2.3.15. Package cache

When `dnf` installs or upgrades a package, it will download the package from the repository and store it temporarily in the cache. The cache also contains repository metadata. The default location of the cache is `/var/cache/dnf`. You can clean the cache with `dnf clean all`.

```
[student@el ~]$ dnf clean all
51 files removed
```

Remark that `.rpm` files will normally be removed automatically after they were installed successfully. You can change this behavior in `/etc/dnf/dnf.conf` by setting `keepcache=1`.

2.3.16. Configuration

The main configuration file for `dnf` is `/etc/dnf/dnf.conf`. This file contains a few basic settings. The location of package repositories that are available to the system are kept in the directory `/etc/yum.repos.d/`. Each repository has its own file, with a `.repo` extension.

```
[student@el ~]$ ls /etc/yum.repos.d/
almalinux-appstream.repo      almalinux-resilientstorage.repo
almalinux-baseos.repo        almalinux-rt.repo
almalinux-crb.repo           almalinux-saphana.repo
almalinux-extras.repo       almalinux-sap.repo
almalinux-highavailability.repo epel-cisco-openh264.repo
almalinux-nfv.repo           epel.repo
almalinux-plus.repo          epel-testing.repo
```

A `repo` file is a text file in the INI format, and contains information about the repository, such as the name, the base URL, the GPG key, etc. Here's an example with part of the contents of the `epel.repo` file:

2. package management

```
1 [epel]
2 name=Extra Packages for Enterprise Linux $releasever - $basearch
3 ## It is much more secure to use the metalink, but if you wish to use a local
   ↪ mirror
4 ## place its address here.
5 #baseurl=https://download.example/pub/epel/$releasever/Everything/$basearch/
   ↪
6 metalink=https://mirrors.fedoraproject.org/metalink?repo=epel-
   ↪ $releasever&arch=$basearch&infra=$infra&content=$contentdir
7 enabled=1
8 gpgcheck=1
9 countme=1
10 gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-EPEL-$releasever
```

2.3.17. Working with multiple repositories

You can get a list of the currently enabled repositories with `dnf repolist`.

```
[student@el ~]$ dnf repolist
repo id           repo name
appstream        AlmaLinux 9 - AppStream
baseos           AlmaLinux 9 - BaseOS
epel             Extra Packages for Enterprise Linux 9 - x86_64
epel-cisco-openh264 Extra Packages for Enterprise Linux 9 openh264 (From Cisco) -
x86_64
extras          AlmaLinux 9 - Extras
```

And specific information about a repository with `dnf repoinfo $repo`.

```
[student@el ~]$ dnf repoinfo epel-testing
Last metadata expiration check: 0:02:40 ago on Sun 25 Feb 2024 10:32:04 PM UTC.
Repo-id           : epel-testing
Repo-name         : Extra Packages for Enterprise Linux 9 - Testing - x86_64
Repo-status       : disabled
Repo-metalink     : https://mirrors.fedoraproject.org/metalink?repo=testing-
epel9&arch=x86_64&infra=$infra&content=$contentdir
Repo-expire       : 172,800 second(s) (last: unknown)
Repo-filename     : /etc/yum.repos.d/epel-testing.repo
Total packages: 0
```

One important flag for `dnf` is `--enablerepo`. Use this command if you want to use a repository that is not enabled by default. For example, let's say you want to install the latest version of `fail2ban`, but the one in the "normal" repository is too old:

```
[student@el ~]$ dnf list --available fail2ban
Last metadata expiration check: 0:01:44 ago on Sun 25 Feb 2024 10:32:04 PM UTC.
Available Packages
fail2ban.noarch          1.0.2-7.el9          epel
```

Maybe `epel-testing` has a newer version:

```
[student@el ~]$ dnf list --available --repo epel-testing fail2ban
Last metadata expiration check: 0:06:10 ago on Sun 25 Feb 2024 10:30:33 PM UTC.
Available Packages
fail2ban.noarch          1.0.2-12.el9        epel-testing
```


It does, but you won't be able to install it due to the fact that `epel-testing` is disabled. However, you can temporarily enable it with the `--enablerepo` flag:

```
[student@el ~]$ sudo dnf install --enablerepo=epel-testing fail2ban
[sudo] password for student:
Last metadata expiration check: 0:13:34 ago on Sun 25 Feb 2024 10:24:12 PM UTC.
Dependencies resolved.
=====
Package                Arch      Version      Repository    Size
=====
Installing:
fail2ban                noarch    1.0.2-12.el9 epel-testing  8.8 k
Installing dependencies:
esmtplib                x86_64    1.2-19.el9   epel          52 k
fail2ban-firewalld     noarch    1.0.2-12.el9 epel-testing  8.9 k
fail2ban-selinux       noarch    1.0.2-12.el9 epel-testing  29 k
fail2ban-sendmail      noarch    1.0.2-12.el9 epel-testing  12 k
fail2ban-server        noarch    1.0.2-12.el9 epel-testing 444 k
libesmtplib            x86_64    1.0.6-24.el9 epel          66 k
libblockfile           x86_64    1.14-10.el9  baseos       28 k
=====
Transaction Summary
=====
Install 8 Packages

Total download size: 647 k
Installed size: 1.8 M
Is this ok [y/N]:
```

2.4. pip, the Python package manager

Some programming languages, a.o. Python, have their own package management system that allows you to install applications and/or libraries. In the case of Python, the package manager is called `pip`. It is used to install Python packages from the Python Package Index (PyPI). In fact, there are multiple package managers for Python (a.o. `easy_install`, `conda`, etc.), but `pip` is the most widely used.

As a system administrator, or as an end user, this sometimes puts you in a difficult position. Some widely known and used Python libraries can be installed both through your distribution's package manager, and through `pip`. Which one to choose is not always clear. In general, it is best to use the distribution's package manager, as it will integrate the package into the system and will be updated when the system is updated. However, some packages are not available in the distribution's repositories, or the version you get with `pip` is more recent. In that case, you can use `pip` to install the package.

Another thing to note is that `pip` can be used as a normal user, or as root, and in each case it will install the package in a different location. When you install a package as a normal user, it will be installed in your home directory, and will only be available to you. When you install a package as root, it will be installed system-wide, and will be available to all users. However, if you install a package as root, you will get a warning message:

WARNING: Running `pip` as the 'root' user can result in broken permissions and conflicting beha

A virtual environment is a way to create an isolated environment for a Python project, where you can install packages without affecting the system's Python installation. This is especially useful when you are developing Python applications, and you want to make sure that the

2. package management

libraries you use are the same as the ones used in production. Using and managing virtual environments is beyond the scope of this course, but you can find more information in the Python documentation.

As general guidelines, we suggest the following:

- If the library or application is available in the distribution's repositories, use the distribution's package manager to install it.
- Avoid installing Python libraries or applications system-wide as root using `pip`.
- Normal users may use `pip` to install Python libraries or applications in their home directory.

2.4.1. installing pip

`pip` may not be installed by default on your system. You can install it using your distribution's package manager. For example, on Debian-based systems, you can install it using `apt`:

```
1 student@debian:~$ sudo apt install python3-pip
```

On Red Hat-based systems, you can install it using `dnf`:

```
1 student@el ~$ sudo dnf install python3-pip
```

2.4.2. listing packages

You can list the packages installed with `pip` using the `list` command:

```
student@linux:~$ pip list
Package          Version
-----
dbus-python      1.2.18
distro           1.5.0
gpg              1.15.1
libcomps         0.1.18
nftables         0.1
pip              21.2.3
PyGObject        3.40.1
python-dateutil 2.8.1
PyYAML           5.4.1
rpm              4.16.1.3
selinux          3.5
sepol            3.5
setools          4.4.3
setuptools       53.0.0
six              1.15.0
systemd-python  234
```

2.4.3. searching for packages

Searching for packages can **NOT** be done on the command line. To search for packages, you can use the Python Package Index website instead. If you try `pip search`, you will get an error message:

```
student@linux:~$ pip search ansible
ERROR: XMLRPC request failed [code: -32500]
RuntimeError: PyPI no longer supports 'pip search' (or XML-RPC search). Please use https://pypi.org/search/ for more information.
```

2.4.4. installing packages

You can install a package using the `install` command:

```
student@linux:~$ pip install ansible
```

Just like `apt` and `dnf`, `pip` will install the package and its dependencies.

2.4.5. removing packages

Uninstalling a package is done with the `uninstall` command:

```
student@linux:~$ pip uninstall ansible
```

Unfortunately, dependencies are not removed when you uninstall a package with `pip`.

2.5. container-based package managers

With the release of Docker, container-based virtualization has become very popular as a method of distributing and deploying applications on servers. One of the advantages of containers is that they offer a sandbox environment for applications, meaning the application and its dependencies are isolated from the rest of the system. This makes it possible to run applications with different dependencies on the same server, without the risk of conflicts. Containers are also very lightweight, they don't impose much overhead on the host system.

Now, there is no reason why containers can't be used to deploy applications on desktop systems as well. In fact, there are several container-based package managers that allow you to install and run applications in containers on your desktop. The advantage is that third party software vendors can distribute their applications independent of the Linux distribution, so they don't need to maintain different packages for (each family of) distribution(s). The disadvantage is that each application comes with their own dependencies, so you lose the advantage of sharing libraries between applications. Also, since the application is running in a container, it may not integrate well with the rest of the system, or may have only limited permissions to access files or other resources on your computer.

As with many Linux-based technologies, there are multiple tools to choose from. The most popular ones are Flatpak and Snap.

2.5.1. flatpak

Flatpak is a container-based package manager developed by an independent community of contributors, volunteers and supporting organizations. It is available for most Linux distributions and is supported by a large number of third party software vendors. Red Hat was one of the first to endorse Flatpak, and many others followed. Fedora Silverblue is a variant of Fedora that uses Flatpak as its primary package manager. Linux Mint also has Flatpak support enabled by default: in the Software Manager, some applications like Bitwarden, Slack, VS Code, etc. are available as Flatpaks.

If you want to use a container based package manager, Flatpak is probably the best choice for any Linux distribution other than Ubuntu.

In the following example, we'll install the open source password manager Bitwarden with Flatpak on a Linux Mint system. Remark that you don't need to be root to install Flatpak applications!

2. package management

```
student@mint:~$ flatpak search Bitwarden
Name      Description                               Application ID      Version Branch Remotes
Bitwarden A secure and free password manager for com.bitwarden.desktop 2024.2.0 stable flathub
Goldwarden A Bitwarden compatible desktop client com.quexten.Goldwarden 0.2.13 stable flathub
student@mint:~$ flatpak install Bitwarden
Looking for matches...
Found ref 'app/com.bitwarden.desktop/x86_64/stable' in remote 'flathub' (system).
Use this ref? [Y/n]: y
Required runtime for com.bitwarden.desktop/x86_64/stable (runtime/org.freedesktop.Platform/x86_64/23.08)
Do you want to install it? [Y/n]: y
```

```
com.bitwarden.desktop permissions:
ipc          network          wayland         x11          dri          file access [1]
dbus access [2]      system dbus access [3]

[1] xdg-download
[2] com.canonical.AppMenu.Registrars, org.freedesktop.Notifications, org.freedesktop.Secret
[3] org.freedesktop.login1
```

	ID	Branch	Op	Remote	Download
1.	[v] com.bitwarden.desktop.Locale	stable	i	flathub	300.7 kB / 9.8 MB
2.	[v] org.freedesktop.Platform.GL.default	23.08	i	flathub	162.0 MB / 162.3 MB
3.	[v] org.freedesktop.Platform.GL.default	23.08-extra	i	flathub	17.9 MB / 162.3 MB
4.	[v] org.freedesktop.Platform.Locale	23.08	i	flathub	17.9 kB / 359.9 MB
5.	[v] org.freedesktop.Platform	23.08	i	flathub	171.6 MB / 225.6 MB
6.	[v] com.bitwarden.desktop	stable	i	flathub	132.5 MB / 133.4 MB

Installation complete.

To remove a Flatpak application, you can use the `uninstall` command:

```
student@mint:~$ flatpak uninstall Bitwarden
Found installed ref 'app/com.bitwarden.desktop/x86_64/stable' (system). Is this correct? [Y/n]: y
```

	ID	Branch	Op
1.	[-] com.bitwarden.desktop	stable	r
2.	[-] com.bitwarden.desktop.Locale	stable	r

Uninstall complete.

2.5.2. snap

Snap was developed by Canonical and is installed by default on Ubuntu. It is also available for other distributions (like the official Ubuntu derivatives, Solus and Zorin OS), but it is not as widely supported as Flatpak. Snap was also designed to work for cloud applications and Internet of Things devices.

In the following example, we'll install Grafana on an Ubuntu Server system.

```
student@ubuntu:~$ snap search grafana
Name      Version Publisher Notes Summary
grafana   6.7.4 canonical✓ -   feature rich metrics dashboard and graph editor
grafana-agent 0.35.4 0x12b -   Telemetry Agent
[ ... ]
student@ubuntu:~$ sudo snap install grafana
grafana 6.7.4 from Canonical✓ installed
```

To uninstall a Snap application, you can use the `remove` command:

```
student@ubuntu:~$ sudo snap remove grafana
grafana removed
```

2.6. downloading software outside the repository

These days, the case where you need software that is not available as a binary package has become exceedingly rare. However, *if* you want to install some experimental tool that hasn't been packaged yet, or you want to test the very latest experimental version of an application, you may have to download the source code and compile it yourself. Usually, the source code is available on the project's website or on a code hosting platform like GitHub, GitLab or Bitbucket. You then either download the source code as a `tgz`, `.tar.gz`, `.tar.bz2`, `tar.xz` file (also called a *tarball*) or you can clone the repository using `git`.

In the example below, we assume that you have downloaded the source code of an application written in C or C++, as is common for many Linux applications. Remark that in order to be able to compile the source code, you need to have the C compiler `gcc` and the build tool `make` installed on your system. You can install these using your distribution's package manager. Also, many applications depend on other libraries, which also have to be installed as source.

2.6.1. example: compiling zork

As an example, we will download the source code for Zork, an ancient text based adventure game, and compile it on a Fedora system. The source code is available on GitHub. We have installed `git`, `gcc` and `make` beforehand.

```
[student@fedora ~]$ git clone https://github.com/devshane/zork.git
Cloning into 'zork' ...
remote: Enumerating objects: 79, done.
remote: Total 79 (delta 0), reused 0 (delta 0), pack-reused 79
Receiving objects: 100% (79/79), 241.70 KiB | 2.14 MiB/s, done.
Resolving deltas: 100% (20/20), done.
[student@fedora ~]$ cd zork/
[student@fedora zork]$ ls
actors.c demons.c dmain.c dso3.c dso6.c dtextc.dat dverb2.c history Makefile np2.c n
ballop.c dgame.c dso1.c dso4.c dso7.c dungeon.6 funcs.h lightp.c nobjs.c np3.c ob
clockr.c dinit.c dso2.c dso5.c dsub.c dverb1.c gdt.c local.c np1.c np.c parse
[student@fedora ~]$ make
cc -g -c -o actors.o actors.c
cc -g -c -o ballop.o ballop.c
cc -g -c -o clockr.o clockr.c
[ ... etc ... ]
cc -g -o zork actors.o ballop.o clockr.o demons.o dgame.o dinit.o dmain.o dso1.o dso2.o dso3
ltermcap
/usr/bin/ld: cannot find -ltermcap: No such file or directory
collect2: error: ld returned 1 exit status
make: *** [Makefile:69: dungeon] Error 1
```

As you can see, the `make` command fails because it cannot find the `termcap` library. This is a library that is used to control the terminal, and it is not installed on our system. This is a common problem when you try to install packages from source. You need to install these dependencies yourself and these are not always easy to find. In this case, we can install the `ncurses-devel` library, which is a modern replacement for `termcap`. How did we now that?

2. package management

We used `dnf provides` to find library files that contain the string `termcap` (remark that the command took a long time to finish):

```
[student@fedora zork]$ dnf provides '*libtermcap.so*'
Last metadata expiration check: 1:56:05 ago on Mon 26 Feb 2024 05:46:43 PM UTC.
ncurses-devel-6.4-7.20230520.fc39.i686 : Development files for the ncurses library
Repo           : fedora
Matched from:
Other          : *libtermcap.so*
[student@fedora ~]$ sudo dnf install ncurses-devel
[ ... etc ... ]
```

Let's try to compile again:

```
[student@fedora zork]$ make
cc -g -c -o actors.o actors.c
cc -g -c -o ballop.o ballop.c
[ ... etc ... ]
cc -g -c -o villns.o villns.c
cc -g -o zork actors.o ballop.o clockr.o demons.o dgame.o dinit.o dmain.o dso1.o dso2.o dso3.o
ltermcap
[student@fedora zork]$
```

The command seems to have succeeded. The current directory now contains a new file called `zork`. This is the compiled application and it has execute permissions. You can run it by typing `./zork`:

```
[student@fedora zork]$ ls -l zork
-rwxr-xr-x. 1 vagrant vagrant 400968 Feb 26 19:45 zork
[student@fedora zork]$ file zork
zork: ELF 64-bit LSB executable, x86-64, version 1 (SYSV), dynamically linked, interpreter /lib64/ld-linux-x86-64.so.2, BuildID[sha1]=3089e3cb1c1a7fc1cc1db41c3aa578c0b52f83f3, for GNU/Linux 3.2+
[student@fedora zork]$ ./zork
Welcome to Dungeon.                This version created 11-MAR-91.
You are in an open field west of a big white house with a boarded
front door.
There is a small mailbox here.
>
```

In this case, installing the game is as simple as copying the `zork` file to a directory in your `PATH`, like `/usr/local/bin` or (for a computer game) `/usr/local/games`. However, most Makefiles provide a way to install the application in the system, usually by running `make install`. This will copy the executable, manual pages and other documentation to the correct location.

```
[student@fedora zork]$ sudo make install
mkdir -p /usr/games /usr/share/man/man6
cp zork /usr/games
cp dtextc.dat /usr/games/lib
cp dungeon.6 /usr/share/man/man6/
```

Remark that the “official” location where manually installed applications belong in a Linux directory structure is `/usr/local` (for applications that follow the Filesystem Hierarchy Standard) or `/opt` (for applications that want to keep all files in a single directory).

2.6.2. installing from a tarball

Before unpacking a tarball, it's useful to check its contents:

```
student@linux:~$ tar tf $downloadedFile.tgz
```

The `t` option lists the content of the archive, `f` should be followed by the filename of the tarball. For `.tgz`, you may add option `z` and for `.tar.bz2` option `j`. However, the `tar` command should recognize the compression method automatically.

Check whether the package archive unpacks in a subdirectory (which is the preferred case) or in the current directory and create a subdirectory yourself if necessary. After that, you can unpack the tarball:

```
student@linux:~$ tar xf $downloadedFile.tgz
```

Now, be sure to read the README file carefully! Normally the readme will explain what to do after download.

Usually the steps are always the same three:

1. running a script `./configure`. It will gather information about your system that is needed to compile the software so that it can actually run on your system
2. executing the command `make` (which is the actual compiling)
3. finally, executing `make install` to copy the files to their proper location.

2.7. practice: package management

1. Verify whether `gcc`, `sudo` and `zork` are installed.
2. Use `dnf` or `apt` to search for and install the `scp`, `tmux`, and `man-pages` packages. Did you find them all?
3. Search the internet for 'webmin' and figure out how to install it.
4. If time permits, search for and install `samba` including the `samba docs pdf` files (thousands of pages in two pdf's).

2.8. solution: package management

1. Verify whether `gcc`, `sudo` and `zork` are installed.

On Enterprise Linux:

```
rpm -qa | grep gcc
rpm -qa | grep sudo
rpm -qa | grep zork
```

On Debian/Ubuntu:

```
dpkg -l | grep gcc
dpkg -l | grep sudo
dpkg -l | grep zork
```

2. Use `dnf` or `apt` to search for and install the `scp`, `tmux`, and `man-pages` packages. Did you find them all ?

On Red Hat/CentOS:

2. *package management*

```
dnf search scp
dnf search tmux
dnf search man-pages
```

On Debian/Ubuntu:

```
apt search scp
apt search tmux
apt search man-pages
```

3. Search the internet for 'webmin' and figure out how to install it.

Google should point you to webmin.com. The download page helps you to download a repository file so you can install webmin with your package manager. The latest Webmin distribution is available in various package formats for download, a.o. .rpm, .deb, etc.

4. If time permits, search for and install samba including the samba docs pdf files (thousands of pages in two pdf's).

Part III.

Scripting 101

3. I/O redirection

(Written by Paul Cobbaut, <https://github.com/paulcobbaut/>, with contributions by: Alex M. Schapelle, <https://github.com/zero-pythagoras/>)


One of the powers of the Unix command line is the use of `input/output redirection` and `pipes`.

This chapter explains `redirection` of `input`, `output` and `error` streams.


3.1. `stdin`, `stdout`, and `stderr`

The `bash` shell has three basic streams; it takes input from `stdin` (stream 0), it sends output to `stdout` (stream 1) and it sends error messages to `stderr` (stream 2).

The drawing below has a graphical interpretation of these three streams.



The keyboard often serves as `stdin`, whereas `stdout` and `stderr` both go to the display. This can be confusing to new Linux users because there is no obvious way to recognize `stdout` from `stderr`. Experienced users know that separating output from errors can be very useful.




The next sections will explain how to redirect these streams.

3.2. output redirection

3.2.1. `> stdout`

`stdout` can be redirected with a `greater` than sign. While scanning the line, the shell will see the `>` sign and will clear the file.

3. I/O redirection



The > notation is in fact the abbreviation of 1> (stdout being referred to as stream 1).

```
[student@linux ~]$ echo It is cold today!  
It is cold today!  
[student@linux ~]$ echo It is cold today! > winter.txt  
[student@linux ~]$ cat winter.txt  
It is cold today!  
[student@linux ~]$
```

Note that the bash shell effectively **removes** the redirection from the command line before argument 0 is executed. This means that in the case of this command:

```
echo hello > greetings.txt
```

the shell only counts two arguments (echo = argument 0, hello = argument 1). The redirection is removed before the argument counting takes place.

3.2.2. output file is erased

While scanning the line, the shell will see the > sign and will **clear** the file! Since this happens before resolving argument 0, this means that even when the command fails, the file will have been cleared!

```
[student@linux ~]$ cat winter.txt  
It is cold today!  
[student@linux ~]$ zcho It is cold today! > winter.txt  
-bash: zcho: command not found  
[student@linux ~]$ cat winter.txt  
[student@linux ~]$
```

3.2.3. noclobber

Erasing a file while using > can be prevented by setting the `noclobber` option.

```
[student@linux ~]$ cat winter.txt  
It is cold today!  
[student@linux ~]$ set -o noclobber  
[student@linux ~]$ echo It is cold today! > winter.txt  
-bash: winter.txt: cannot overwrite existing file  
[student@linux ~]$ set +o noclobber  
[student@linux ~]$
```

3.2.4. overruling noclobber

The `noclobber` can be overruled with `>|`.

```
[student@linux ~]$ set -o noclobber
[student@linux ~]$ echo It is cold today! > winter.txt
-bash: winter.txt: cannot overwrite existing file
[student@linux ~]$ echo It is very cold today! >| winter.txt
[student@linux ~]$ cat winter.txt
It is very cold today!
[student@linux ~]$
```

3.2.5. » append


Use `>>` to append output to a file.

```
[student@linux ~]$ echo It is cold today! > winter.txt
[student@linux ~]$ cat winter.txt
It is cold today!
[student@linux ~]$ echo Where is the summer ? >> winter.txt
[student@linux ~]$ cat winter.txt
It is cold today!
Where is the summer ?
[student@linux ~]$
```

3.3. error redirection

3.3.1. 2> stderr

Redirecting `stderr` is done with `2>`. This can be very useful to prevent error messages from cluttering your screen.



The screenshot below shows redirection of `stdout` to a file, and `stderr` to `/dev/null`. Writing `1>` is the same as `>`.

```
[student@linux ~]$ find / > allfiles.txt 2> /dev/null
[student@linux ~]$
```

3.3.2. 2>&1

To redirect both `stdout` and `stderr` to the same file, use `2>&1`.

```
[student@linux ~]$ find / > allfiles_and_errors.txt 2>&1
[student@linux ~]$
```

Note that the order of redirections is significant. For example, the command

3. I/O redirection

```
ls > dirlist 2>&1
```

directs both standard output (file descriptor 1) and standard error (file descriptor 2) to the file dirlist, while the command

```
ls 2>&1 > dirlist
```

directs only the standard output to file dirlist, because the standard error made a copy of the standard output before the standard output was redirected to dirlist.

3.4. output redirection and pipes

By default you cannot grep inside stderr when using pipes on the command line, because only stdout is passed.

```
student@linux:~$ rm file42 file33 file1201 | grep file42
rm: cannot remove 'file42': No such file or directory
rm: cannot remove 'file33': No such file or directory
rm: cannot remove 'file1201': No such file or directory
```

With 2>&1 you can force stderr to go to stdout. This enables the next command in the pipe to act on both streams.

```
student@linux:~$ rm file42 file33 file1201 2>&1 | grep file42
rm: cannot remove 'file42': No such file or directory
```

You cannot use both 1>&2 and 2>&1 to switch stdout and stderr.

```
student@linux:~$ rm file42 file33 file1201 2>&1 1>&2 | grep file42
rm: cannot remove 'file42': No such file or directory
student@linux:~$ echo file42 2>&1 1>&2 | sed 's/file42/FILE42/'
FILE42
```

You need a third stream to switch stdout and stderr after a pipe symbol.

```
student@linux:~$ echo file42 3>&1 1>&2 2>&3 | sed 's/file42/FILE42/'
file42
student@linux:~$ rm file42 3>&1 1>&2 2>&3 | sed 's/file42/FILE42/'
rm: cannot remove 'FILE42': No such file or directory
```

3.5. joining stdout and stderr

The &> construction will put both stdout and stderr in one stream (to a file).

```
student@linux:~$ rm file42 &> out_and_err
student@linux:~$ cat out_and_err
rm: cannot remove 'file42': No such file or directory
student@linux:~$ echo file42 &> out_and_err
student@linux:~$ cat out_and_err
file42
student@linux:~$
```

3.6. input redirection

3.6.1. < stdin

Redirecting `stdin` is done with `<` (short for `0<`).

```
[student@linux ~]$ cat < text.txt
one
two
[student@linux ~]$ tr 'onetw' 'ONEZZ' < text.txt
ONE
ZZO
[student@linux ~]$
```

3.6.2. « here document

The `here document` (sometimes called `here-is-document`) is a way to append input until a certain sequence (usually EOF) is encountered. The EOF marker can be typed literally or can be called with `Ctrl-D`.

```
[student@linux ~]$ cat <<EOF > text.txt
> one
> two
> EOF
[student@linux ~]$ cat text.txt
one
two
[student@linux ~]$ cat <<bro1 > text.txt
> bro1
[student@linux ~]$ cat text.txt
bro1
[student@linux ~]$
```

3.6.3. «< here string

The `here string` can be used to directly pass strings to a command. The result is the same as using `echo string | command` (but you have one less process running).

```
student@linux~$ base64 <<< linux-training.be
bGludXgtHJhaW5pbmYmUK
student@linux~$ base64 -d <<< bGludXgtHJhaW5pbmYmUK
linux-training.be
```

See [rfc 3548](#) for more information about `base64`.

3. I/O redirection

3.7. confusing redirection

The shell will scan the whole line before applying redirection. The following command line is very readable and is correct.

```
cat winter.txt > snow.txt 2> errors.txt
```

But this one is also correct, but less readable.

```
2> errors.txt cat winter.txt > snow.txt
```

Even this will be understood perfectly by the shell.

```
< winter.txt > snow.txt 2> errors.txt cat
```

3.8. quick file clear

So what is the quickest way to clear a file ?

```
>foo
```

And what is the quickest way to clear a file when the `noclobber` option is set ?

```
>|bar
```

3.9. practice: input/output redirection

1. Activate the `noclobber` shell option.
2. Verify that `noclobber` is active by repeating an `ls` on `/etc/` with redirected output to a file.
3. When listing all shell options, which character represents the `noclobber` option ?
4. Deactivate the `noclobber` option.
5. Make sure you have two shells open on the same computer. Create an empty `tailing.txt` file. Then type `tail -f tailing.txt`. Use the second shell to append a line of text to that file. Verify that the first shell displays this line.
6. Create a file that contains the names of five people. Use `cat` and output redirection to create the file and use a `here` document to end the input.

3.10. solution: input/output redirection

1. Activate the noclobber shell option.

```
set -o noclobber
set -C
```

2. Verify that noclobber is active by repeating an `ls` on `/etc/` with redirected output to a file.

```
ls /etc > etc.txt
ls /etc > etc.txt (should not work)
```

3. When listing all shell options, which character represents the noclobber option ?

```
echo $- (noclobber is visible as C)
```

4. Deactivate the noclobber option.

```
set +o noclobber
```

5. Make sure you have two shells open on the same computer. Create an empty `tailing.txt` file. Then type `tail -f tailing.txt`. Use the second shell to append a line of text to that file. Verify that the first shell displays this line.

```
student@linux:~$ > tailing.txt
student@linux:~$ tail -f tailing.txt
hello
world
```

in the other shell:

```
student@linux:~$ echo hello >> tailing.txt
student@linux:~$ echo world >> tailing.txt
```

6. Create a file that contains the names of five people. Use `cat` and output redirection to create the file and use a `here` document to end the input.

```
student@linux:~$ cat > tennis.txt << ace
> Justine Henin
> Venus Williams
> Serena Williams
> Martina Hingis
> Kim Clijsters
> ace
student@linux:~$ cat tennis.txt
Justine Henin
Venus Williams
Serena Williams
Martina Hingis
Kim Clijsters
student@linux:~$
```


4. filters

(Written by Paul Cobbaut, <https://github.com/paulcobbaut/>, with contributions by: Alex M. Schapelle, <https://github.com/zero-pythagoras/>)

Commands that are created to be used with a pipe are often called **filters**. These filters are very small programs that do one specific thing very efficiently. They can be used as building blocks.

This chapter will introduce you to the most common **filters**. The combination of simple commands and filters in a long **pipe** allows you to design elegant solutions.

4.1. cat

When between two pipes, the **cat** command does nothing (except putting **stdin** on **stdout**).

```
[student@linux pipes]$ tac count.txt | cat | cat | cat | cat | cat
five
four
three
two
one
[student@linux pipes]$
```

4.2. tee

Writing long pipes in Unix is fun, but sometimes you may want intermediate results. This is where **tee** comes in handy. The **tee** filter puts **stdin** on **stdout** and also into a file. So **tee** is almost the same as **cat**, except that it has two identical outputs.

```
[student@linux pipes]$ tac count.txt | tee temp.txt | tac
one
two
three
four
five
[student@linux pipes]$ cat temp.txt
five
four
three
two
one
[student@linux pipes]$
```

4. filters

4.3. grep

The `grep` filter is famous among Unix users. The most common use of `grep` is to filter lines of text containing (or not containing) a certain string.

```
[student@linux pipes]$ cat tennis.txt
Amelie Mauresmo, Fra
Kim Clijsters, BEL
Justine Henin, Bel
Serena Williams, usa
Venus Williams, USA
[student@linux pipes]$ cat tennis.txt | grep Williams
Serena Williams, usa
Venus Williams, USA
```

You can write this without the `cat`.

```
[student@linux pipes]$ grep Williams tennis.txt
Serena Williams, usa
Venus Williams, USA
```

One of the most useful options of `grep` is `grep -i` which filters in a case insensitive way.

```
[student@linux pipes]$ grep Bel tennis.txt
Justine Henin, Bel
[student@linux pipes]$ grep -i Bel tennis.txt
Kim Clijsters, BEL
Justine Henin, Bel
[student@linux pipes]$
```

Another very useful option is `grep -v` which outputs lines not matching the string.

```
[student@linux pipes]$ grep -v Fra tennis.txt
Kim Clijsters, BEL
Justine Henin, Bel
Serena Williams, usa
Venus Williams, USA
[student@linux pipes]$
```

And of course, both options can be combined to filter all lines not containing a case insensitive string.

```
[student@linux pipes]$ grep -vi usa tennis.txt
Amelie Mauresmo, Fra
Kim Clijsters, BEL
Justine Henin, Bel
[student@linux pipes]$
```

With `grep -A1` one line after the result is also displayed.

```
student@linux:~/pipes$ grep -A1 Henin tennis.txt
Justine Henin, Bel
Serena Williams, usa
```

With `grep -B1` one line before the result is also displayed.

```
student@linux:~/pipes$ grep -B1 Henin tennis.txt
Kim Clijsters, BEL
Justine Henin, Bel
```

With `grep -C1` (context) one line before and one after are also displayed. All three options (A,B, and C) can display any number of lines (using e.g. A2, B4 or C20).

```
student@linux:~/pipes$ grep -C1 Henin tennis.txt
Kim Clijsters, BEL
Justine Henin, Bel
Serena Williams, usa
```

4.4. cut

The `cut` filter can select columns from files, depending on a delimiter or a count of bytes. The screenshot below uses `cut` to filter for the username and userid in the `/etc/passwd` file. It uses the colon as a delimiter, and selects fields 1 and 3.

```
[[student@linux pipes]$ cut -d: -f1,3 /etc/passwd | tail -4
Figo:510
Pfaff:511
Harry:516
Hermione:517
[student@linux pipes]$
```

When using a space as the delimiter for `cut`, you have to quote the space.

```
[student@linux pipes]$ cut -d" " -f1 tennis.txt
Amelie
Kim
Justine
Serena
Venus
[student@linux pipes]$
```

This example uses `cut` to display the second to the seventh character of `/etc/passwd`.

```
[student@linux pipes]$ cut -c2-7 /etc/passwd | tail -4
igo:x:
faff:x
arry:x
ermion
[student@linux pipes]$
```

4.5. tr

You can translate characters with `tr`. The screenshot shows the translation of all occurrences of `e` to `E`.

4. filters

```
[student@linux pipes]$ cat tennis.txt | tr 'e' 'E'  
AmELiE MaurESmo, Fra  
Kim CliJstErs, BEL  
JustinE HEnin, BEL  
SErEna Williams, usa  
VENus Williams, USA
```

Here we set all letters to uppercase by defining two ranges.

```
[student@linux pipes]$ cat tennis.txt | tr 'a-z' 'A-Z'  
AMELIE MAURESMO, FRA  
KIM CLIJSTERS, BEL  
JUSTINE HENIN, BEL  
SERENA WILLIAMS, USA  
VENUS WILLIAMS, USA  
[student@linux pipes]$
```

Here we translate all newlines to spaces.

```
[student@linux pipes]$ cat count.txt  
one  
two  
three  
four  
five  
[student@linux pipes]$ cat count.txt | tr '\n' ' '  
one two three four five [student@linux pipes]$
```

The `tr -s` filter can also be used to squeeze multiple occurrences of a character to one.

```
[student@linux pipes]$ cat spaces.txt  
one two three  
four five six  
[student@linux pipes]$ cat spaces.txt | tr -s ' '  
one two three  
four five six  
[student@linux pipes]$
```

You can also use `tr` to 'encrypt' texts with `rot13`.

```
[student@linux pipes]$ cat count.txt | tr 'a-z' 'nopqrstuvwxyzabcdefghijklm'  
bar  
gjb  
guerr  
sbhe  
svir  
[student@linux pipes]$ cat count.txt | tr 'a-z' 'n-za-m'  
bar  
gjb  
guerr  
sbhe  
svir  
[student@linux pipes]$
```

This last example uses `tr -d` to delete characters.

```
student@linux:~/pipes$ cat tennis.txt | tr -d e
Amlı Maursmo, Fra
Kim Clijstrs, BEL
Justin Hnin, Bl
Srna Williams, usa
Vnus Williams, USA
```

4.6. wc

Counting words, lines and characters is easy with wc.

```
[student@linux pipes]$ wc tennis.txt
 5 15 100 tennis.txt
[student@linux pipes]$ wc -l tennis.txt
5 tennis.txt
[student@linux pipes]$ wc -w tennis.txt
15 tennis.txt
[student@linux pipes]$ wc -c tennis.txt
100 tennis.txt
[student@linux pipes]$
```

4.7. sort

The sort filter will default to an alphabetical sort.

```
student@linux:~/pipes$ cat music.txt
Queen
Brel
Led Zeppelin
Abba
student@linux:~/pipes$ sort music.txt
Abba
Brel
Led Zeppelin
Queen
```

But the sort filter has many options to tweak its usage. This example shows sorting different columns (column 1 or column 2).

```
[student@linux pipes]$ sort -k1 country.txt
Belgium, Brussels, 10
France, Paris, 60
Germany, Berlin, 100
Iran, Teheran, 70
Italy, Rome, 50
[student@linux pipes]$ sort -k2 country.txt
Germany, Berlin, 100
Belgium, Brussels, 10
France, Paris, 60
Italy, Rome, 50
Iran, Teheran, 70
```

4. filters

The screenshot below shows the difference between an alphabetical sort and a numerical sort (both on the third column).

```
[student@linux pipes]$ sort -k3 country.txt
Belgium, Brussels, 10
Germany, Berlin, 100
Italy, Rome, 50
France, Paris, 60
Iran, Teheran, 70
[student@linux pipes]$ sort -n -k3 country.txt
Belgium, Brussels, 10
Italy, Rome, 50
France, Paris, 60
Iran, Teheran, 70
Germany, Berlin, 100
```

4.8. uniq

With `uniq` you can remove duplicates from a sorted list.

```
student@linux:~/pipes$ cat music.txt
Queen
Brel
Queen
Abba
student@linux:~/pipes$ sort music.txt
Abba
Brel
Queen
Queen
student@linux:~/pipes$ sort music.txt |uniq
Abba
Brel
Queen
```

`uniq` can also count occurrences with the `-c` option.

```
student@linux:~/pipes$ sort music.txt |uniq -c
 1 Abba
 1 Brel
 2 Queen
```

4.9. comm

Comparing streams (or files) can be done with the `comm`. By default `comm` will output three columns. In this example, Abba, Cure and Queen are in both lists, Bowie and Sweet are only in the first file, Turner is only in the second.

```
student@linux:~/pipes$ cat > list1.txt
Abba
Bowie
Cure
Queen
```



```

Sweet
student@linux:~/pipes$ cat > list2.txt
Abba
Cure
Queen
Turner
student@linux:~/pipes$ comm list1.txt list2.txt
                Abba
Bowie
                Cure
                Queen
Sweet
                Turner

```

The output of `comm` can be easier to read when outputting only a single column. The digits point out which output columns should not be displayed.

```

student@linux:~/pipes$ comm -12 list1.txt list2.txt
Abba
Cure
Queen
student@linux:~/pipes$ comm -13 list1.txt list2.txt
Turner
student@linux:~/pipes$ comm -23 list1.txt list2.txt
Bowie
Sweet

```

4.10. od

European humans like to work with ascii characters, but computers store files in bytes. The example below creates a simple file, and then uses `od` to show the contents of the file in hexadecimal bytes

```

student@linux:~/test$ cat > text.txt
abcdefg
1234567
student@linux:~/test$ od -t x1 text.txt
0000000 61 62 63 64 65 66 67 0a 31 32 33 34 35 36 37 0a
0000020

```

The same file can also be displayed in octal bytes.

```

student@linux:~/test$ od -b text.txt
0000000 141 142 143 144 145 146 147 012 061 062 063 064 065 066 067 012
0000020

```

And here is the file in ascii (or backslashed) characters.

```

student@linux:~/test$ od -c text.txt
0000000 a b c d e f g \n 1 2 3 4 5 6 7 \n
0000020

```

4. filters

4.11. sed

The stream editor sed can perform editing functions in the stream, using regular expressions.

```
student@linux:~/pipes$ echo level5 | sed 's/5/42/'
level42
student@linux:~/pipes$ echo level5 | sed 's/level/jump/'
jump5
```

Add g for global replacements (all occurrences of the string per line).

```
student@linux:~/pipes$ echo level5 level7 | sed 's/level/jump/'
jump5 level7
student@linux:~/pipes$ echo level5 level7 | sed 's/level/jump/g'
jump5 jump7
```

With d you can remove lines from a stream containing a character.

```
student@linux:~/test42$ cat tennis.txt
Venus Williams, USA
Martina Hingis, SUI
Justine Henin, BE
Serena williams, USA
Kim Clijsters, BE
Yanina Wickmayer, BE
student@linux:~/test42$ cat tennis.txt | sed '/BE/d'
Venus Williams, USA
Martina Hingis, SUI
Serena williams, USA
```

4.12. pipe examples

4.12.1. who | wc

How many users are logged on to this system ?

```
[student@linux pipes]$ who
root    tty1          Jul 25 10:50
paul    pts/0         Jul 25 09:29 (laika)
Harry   pts/1         Jul 25 12:26 (barry)
paul    pts/2         Jul 25 12:26 (pasha)
[student@linux pipes]$ who | wc -l
4
```

4.12.2. who | cut | sort

Display a sorted list of logged on users.

```
[student@linux pipes]$ who | cut -d' ' -f1 | sort
Harry
paul
paul
root
```

Display a sorted list of logged on users, but every user only once .

```
[student@linux pipes]$ who | cut -d' ' -f1 | sort | uniq
Harry
paul
root
```

4.12.3. grep | cut

Display a list of all bash user accounts on this computer. Users accounts are explained in detail later.

```
student@linux:~$ grep bash /etc/passwd
root:x:0:0:root:/root:/bin/bash
paul:x:1000:1000:paul,,,:/home/paul:/bin/bash
serena:x:1001:1001::/home/serena:/bin/bash
student@linux:~$ grep bash /etc/passwd | cut -d: -f1
root
paul
serena
```

4.13. practice: filters

1. Put a sorted list of all bash users in bashusers.txt.
2. Put a sorted list of all logged on users in onlineusers.txt.
3. Make a list of all filenames in /etc that contain the string conf in their filename.
4. Make a sorted list of all files in /etc that contain the case insensitive string conf in their filename.
5. Look at the output of /sbin/ifconfig. Write a line that displays only ip address and the subnet mask.
6. Write a line that removes all non-letters from a stream.
7. Write a line that receives a text file, and outputs all words on a separate line.
8. Write a spell checker on the command line. (There may be a dictionary in /usr/share/dict/.)

4. filters

4.14. solution: filters

1. Put a sorted list of all bash users in bashusers.txt.

```
grep bash /etc/passwd | cut -d: -f1 | sort > bashusers.txt
```

2. Put a sorted list of all logged on users in onlineusers.txt.

```
who | cut -d' ' -f1 | sort > onlineusers.txt
```

3. Make a list of all filenames in /etc that contain the string conf in their filename.

```
ls /etc | grep conf
```

4. Make a sorted list of all files in /etc that contain the case insensitive string conf in their filename.

```
ls /etc | grep -i conf | sort
```

5. Look at the output of /sbin/ifconfig. Write a line that displays only ip address and the subnet mask.

```
/sbin/ifconfig | head -2 | grep 'inet ' | tr -s ' ' | cut -d' ' -f3,5
```

6. Write a line that removes all non-letters from a stream.

```
student@linux:~$ cat text
This is, yes really! , a text with ?&* too many str$ange# characters ;-)
student@linux:~$ cat text | tr -d ',!$?.*&^%#@;()-'
This is yes really a text with too many strange characters
```

7. Write a line that receives a text file, and outputs all words on a separate line.

```
student@linux:~$ cat text2
it is very cold today without the sun

student@linux:~$ cat text2 | tr ' ' '\n'
it
is
very
cold
today
without
the
sun
```

8. Write a spell checker on the command line. (There may be a dictionary in /usr/share/dict/.)

```
student@linux ~$ echo "The zun is shining today" > text
```

```
student@linux ~$ cat > DICT
```

```
is  
shining  
sun  
the  
today
```

```
student@linux ~$ cat text | tr 'A-Z ' 'a-z\n' | sort | uniq | comm -23 - DICT  
zun
```

You could also add the solution from question number 6 to remove non-letters, and `tr -s ' '` to remove redundant spaces.

5. shell variables

(Written by Paul Cobbaut, <https://github.com/paulcobbaut/>, with contributions by: Alex M. Schapelle, <https://github.com/zero-pythagoras/>)

In this chapter we learn to manage environment variables in the shell. These variables are often needed by applications.

5.1. \$ dollar sign

Another important character interpreted by the shell is the dollar sign `$`. The shell will look for an environment variable named like the string following the dollar sign and replace it with the value of the variable (or with nothing if the variable does not exist).

These are some examples using `$HOSTNAME`, `$USER`, `$UID`, `$SHELL`, and `$HOME`.

```
[student@linux ~]$ echo This is the $SHELL shell
This is the /bin/bash shell
[student@linux ~]$ echo This is $SHELL on computer $HOSTNAME
This is /bin/bash on computer RHELv8u3.localdomain
[student@linux ~]$ echo The userid of $USER is $UID
The userid of paul is 500
[student@linux ~]$ echo My homedir is $HOME
My homedir is /home/paul
```

5.2. case sensitive

This example shows that shell variables are case sensitive!

```
[student@linux ~]$ echo Hello $USER
Hello paul
[student@linux ~]$ echo Hello $user
Hello
```

5.3. creating variables

This example creates the variable `$MyVar` and sets its value. It then uses `echo` to verify the value.

```
[student@linux gen]$ MyVar=555
[student@linux gen]$ echo $MyVar
555
[student@linux gen]$
```

5. shell variables

5.4. quotes

Notice that double quotes still allow the parsing of variables, whereas single quotes prevent this.

```
[student@linux ~]$ MyVar=555
[student@linux ~]$ echo $MyVar
555
[student@linux ~]$ echo "$MyVar"
555
[student@linux ~]$ echo '$MyVar'
$MyVar
```

The bash shell will replace variables with their value in double quoted lines, but not in single quoted lines.

```
student@linux:~$ city=Burtonville
student@linux:~$ echo "We are in $city today."
We are in Burtonville today.
student@linux:~$ echo 'We are in $city today.'
We are in $city today.
```

5.5. set

You can use the `set` command to display a list of environment variables. On Ubuntu and Debian systems, the `set` command will also list shell functions after the shell variables. Use `set | more` to see the variables then.

5.6. unset

Use the `unset` command to remove a variable from your shell environment.

```
[student@linux ~]$ MyVar=8472
[student@linux ~]$ echo $MyVar
8472
[student@linux ~]$ unset MyVar
[student@linux ~]$ echo $MyVar

[student@linux ~]$
```

5.7. \$PS1

The `$PS1` variable determines your shell prompt. You can use backslash escaped special characters like `\u` for the username or `\w` for the working directory. The bash manual has a complete reference.

In this example we change the value of `$PS1` a couple of times.


```
student@linux:~$ PS1=prompt
prompt
promptPS1='prompt '
prompt
prompt PS1='> '
>
> PS1='\u@\h$ '
student@linux$
student@linux$ PS1='\u@\h:\w$'
student@linux:~$
```

To avoid unrecoverable mistakes, you can set normal user prompts to green and the root prompt to red. Add the following to your `.bashrc` for a green user prompt:

```
# color prompt by paul
RED='\[\033[01;31m\]'
WHITE='\[\033[01;00m\]'
GREEN='\[\033[01;32m\]'
BLUE='\[\033[01;34m\]'
export PS1="${debian_chroot:+($debian_chroot)}$GREEN\u$WHITE@$BLUE\h$WHITE\w\$ "
```

5.8. \$PATH

The `$PATH` variable determines where the shell is looking for commands to execute (unless the command is builtin or aliased). This variable contains a list of directories, separated by colons.

```
[student@linux ~]$ echo $PATH
/usr/kerberos/bin:/usr/local/bin:/bin:/usr/bin:
```

The shell will not look in the current directory for commands to execute! (Looking for executables in the current directory provided an easy way to hack PC-DOS computers). If you want the shell to look in the current directory, then add a `.` at the end of your `$PATH`.

```
[student@linux ~]$ PATH=$PATH:..
[student@linux ~]$ echo $PATH
/usr/kerberos/bin:/usr/local/bin:/bin:/usr/bin:..
[student@linux ~]$
```

Your path might be different when using `su` instead of `su -` – because the latter will take on the environment of the target user. The root user typically has `/sbin` directories added to the `$PATH` variable.

```
[student@linux ~]$ su
Password:
[root@linux paul]# echo $PATH
/usr/local/bin:/bin:/usr/bin:/usr/X11R6/bin
[root@linux paul]# exit
[student@linux ~]$ su -
Password:
[root@linux ~]# echo $PATH
/usr/local/sbin:/usr/local/bin:/sbin:/bin:/usr/sbin:/usr/bin:
[root@linux ~]#
```

5.9. env

The `env` command without options will display a list of exported variables. The difference with `set` with options is that `set` lists all variables, including those not exported to child shells.

But `env` can also be used to start a clean shell (a shell without any inherited environment). The `env -i` command clears the environment for the subshell.

Notice in this screenshot that `bash` will set the `$SHELL` variable on startup.

```
[student@linux ~]$ bash -c 'echo $SHELL $HOME $USER'
/bin/bash /home/paul paul
[student@linux ~]$ env -i bash -c 'echo $SHELL $HOME $USER'
/bin/bash
[student@linux ~]$
```

You can use the `env` command to set the `$LANG`, or any other, variable for just one instance of `bash` with one command. The example below uses this to show the influence of the `$LANG` variable on file globbing (see the chapter on file globbing).

```
[student@linux test]$ env LANG=C bash -c 'ls File[a-z]'
Filea Fileb
[student@linux test]$ env LANG=en_US.UTF-8 bash -c 'ls File[a-z]'
Filea FileA Fileb FileB
[student@linux test]$
```

5.10. export

You can export shell variables to other shells with the `export` command. This will export the variable to child shells.

```
[student@linux ~]$ var3=three
[student@linux ~]$ var4=four
[student@linux ~]$ export var4
[student@linux ~]$ echo $var3 $var4
three four
[student@linux ~]$ bash
[student@linux ~]$ echo $var3 $var4
four
```

But it will not export to the parent shell (previous screenshot continued).

```
[student@linux ~]$ export var5=five
[student@linux ~]$ echo $var3 $var4 $var5
four five
[student@linux ~]$ exit
exit
[student@linux ~]$ echo $var3 $var4 $var5
three four
[student@linux ~]$
```

5.11. delineate variables

Until now, we have seen that bash interprets a variable starting from a dollar sign, continuing until the first occurrence of a non-alphanumeric character that is not an underscore. In some situations, this can be a problem. This issue can be resolved with curly braces like in this example.

```
[student@linux ~]$ prefix=Super
[student@linux ~]$ echo Hello $prefixman and $prefixgirl
Hello  and
[student@linux ~]$ echo Hello ${prefix}man and ${prefix}girl
Hello Superman and Supergirl
[student@linux ~]$
```

5.12. unbound variables

The example below tries to display the value of the `$MyVar` variable, but it fails because the variable does not exist. By default the shell will display nothing when a variable is unbound (does not exist).

```
[student@linux gen]$ echo $MyVar

[student@linux gen]$
```

There is, however, the `nounset` shell option that you can use to generate an error when a variable does not exist.

```
student@linux:~$ set -u
student@linux:~$ echo $Myvar
bash: Myvar: unbound variable
student@linux:~$ set +u
student@linux:~$ echo $Myvar

student@linux:~$
```

In the bash shell `set -u` is identical to `set -o nounset` and likewise `set +u` is identical to `set +o nounset`.

5.13. practice: shell variables

1. Use `echo` to display Hello followed by your username. (use a bash variable!)
2. Create a variable `answer` with a value of 42.
3. Copy the value of `$LANG` to `$MyLANG`.
4. List all current shell variables.
5. List all exported shell variables.
6. Do the `env` and `set` commands display your variable?
6. Destroy your `answer` variable.
7. Create two variables, and export one of them.

5. shell variables

8. Display the exported variable in an interactive child shell.
9. Create a variable, give it the value 'Dumb', create another variable with value 'do'. Use echo and the two variables to echo Dumbledore.
10. Find the list of backslash escaped characters in the manual of bash. Add the time to your PS1 prompt.

5.14. solution: shell variables

1. Use echo to display Hello followed by your username. (use a bash variable!)

```
echo Hello $USER
```

2. Create a variable answer with a value of 42.

```
answer=42
```

3. Copy the value of \$LANG to \$MyLANG.

```
MyLANG=$LANG
```

4. List all current shell variables.

```
set
```

```
set | more on Ubuntu/Debian
```

5. List all exported shell variables.

```
env  
export  
declare -x
```

6. Do the env and set commands display your variable ?

```
env | more  
set | more
```

6. Destroy your answer variable.

```
unset answer
```

7. Create two variables, and export one of them.

```
var1=1; export var2=2
```

8. Display the exported variable in an interactive child shell.

```
bash  
echo $var2
```

9. Create a variable, give it the value 'Dumb', create another variable with value 'do'. Use echo and the two variables to echo Dumbledore.

```
varx=Dumb; vary=do
```

```
echo ${varx}le${vary}re
```

```
solution by Yves from Dexia : echo $varx'le'$vary're'
```

```
solution by Erwin from Telenet : echo "$varx"le"$vary"re
```

10. Find the list of backslash escaped characters in the manual of bash. Add the time to your PS1 prompt.

```
PS1='\t \u@\h \W$ '
```


6. introduction to scripting

(Written by Paul Cobbaut, <https://github.com/paulcobbaut/>, with contributions by: Alex M. Schapelle, <https://github.com/zero-pythagoras/>, Bert Van Vreckem <https://github.com/bertvw/>)

The goal of this chapter is to give you all the information in order to read, write and understand small, long and complex shell scripts.

You should have read and understood part III shell expansion and part IV pipes and commands before starting this chapter.

6.1. introduction

When you open a terminal and type a command, you are using a *shell*, an interactive environment that interprets your commands, executes them, and shows you the output the command generates. Most Linux distributions have Bash (the “Bourne Again Shell”) as the default, but there are others as well: the original “Bourne shell” (sh), the “Debian Amquist Shell” (dash, a modern implementation of sh), the “Korn shell” (ksh), the “C shell” (csh), and the “Z shell” (zsh), to name a few.

A sequence of commands can be saved in a file and executed as a single command. This is called a *script*. Shell scripts are used to automate tasks, and are an essential tool for system administrators and developers. Subsequently, this means that system administrators or SysOps also need solid knowledge of *scripting* to understand how their servers and their applications are started, updated, upgraded, patched, maintained, configured and removed, and also to understand how a user environment is built.

Shells have also support for programming constructs (like loops, functions, variables, etc.) so that you can write more complex scripts. This makes a scripting language basically as powerful as a programming language. Scripting languages are often interpreted, rather than compiled.

If you copy a script to one of the bin directories (e.g. /usr/local/bin), you can execute it from the command line just like any other command. In fact, many UNIX/Linux commands are essentially scripts. You can check this for yourself by executing the file command on the executables in the /bin directory. For example:

```
student@linux:~$ file /usr/bin/* | awk '{ print($2, $3, $4) }' \
| sort | uniq -c | sort -nr
 466 ELF 64-bit LSB
 168 symbolic link to
  74 POSIX shell script,
  71 Perl script text
  14 Python script, ASCII
  10 setuid ELF 64-bit
   7 setgid ELF 64-bit
   6 Bourne-Again shell script,
   2 Python script, Unicode
   1 Python script, ISO-8859
```

6. introduction to scripting

We find POSIX (Bourne), Bash, Perl and Python scripts, as well as ELF binaries (compiled programs). This shows that a significant portion of the commands in a typical Linux system are actually scripts.

Bash scripting is a valuable skill for any Linux user, but these days, its applications are no longer limited to Linux. Bash is also present on macOS (albeit an older version), and with the advent of Windows Subsystem for Linux (WSL), Bash is now available for Windows users as well. Moreover, Git Bash, a Bash shell for Windows, is also available.

6.2. hello world

Just like in every programming course, we start with a simple `hello_world` script. The following script will output `Hello World`.

```
1 echo Hello World
```

After creating this simple script in `nano`, `vi`, or with `echo`, you'll have to `chmod +x hello_world` to make it executable. And unless you add the scripts directory to your path, you'll have to type the path to the script for the shell to be able to find it.

```
student@linux:~$ echo echo Hello World > hello_world
student@linux:~$ chmod +x hello_world
student@linux:~$ ./hello_world
Hello World
student@linux:~$
```

6.3. she-bang

Let's expand our example a little further by putting `#!/bin/bash` on the first line of the script. The `#!` is called a she-bang (sometimes called sha-bang), where the she-bang is the first two characters of the script.

Open the file with `nano hello_world` or `vi hello_world` and add the following line at the top of the file.

```
1 #!/bin/bash
2 echo Hello World
```

You can never be sure which (interactive) shell a user is running. A script that works flawlessly in bash might not work in `ksh`, `csh`, or `dash`. To instruct a shell to run your script with a specific interpreter, you should start your script with a she-bang followed by the absolute path to the executable of the interpreter.

This script will run in a bash shell.

```
1 #!/bin/bash
2 echo -n hello
3 echo A bash subshell $(echo -n hello)
```

This script will be interpreted by Python:

```
1 #!/usr/bin/env python3
2 print("Hello World!")
```

The following script will run in a Korn shell (unless `/bin/ksh` is a hard link to `/bin/bash`). The `/etc/shells` file contains a list of shells available on your system. Check it to see which ones are available to you


```

1  #!/bin/ksh
2  echo -n hello
3  echo a Korn subshell $(echo -n hello)

```

If you're not sure in which `bin` directory the shell executable is located, you can use `env`. The command `env` is normally used to print environment variables, but in the context of a script, it is used to launch the correct interpreter.

```

1  #!/usr/bin/env bash
2  echo -n hello
3  echo A bash subshell $(echo -n hello)

```

This is particularly useful for macOS users: out-of-the-box, a macOS system has a very old version of `bash` in `/bin/bash`. If you want to use a more recent version, you can install it with Homebrew, that will put it in `/usr/local/bin/bash`. If you use `#!/usr/bin/env bash` in your scripts, the newer version will be used.

6.4. comments

When writing Bash scripts, it is always a good practice to make your code clean and easily understandable. Organizing your code in blocks, indenting, giving variables and functions descriptive names are several ways to do this. Another way to improve the readability of your code is by using comments. A comment is a human-readable explanation or annotation that is written in the shell script.

Let's expand our example a little further by adding comment lines.

```

1  #!/usr/bin/env bash
2  #
3  # hello_world.sh -- My first script
4  #
5  echo Hello World
6
7  # this is old way of calling for subshell with backtick ``
8  echo A bash subshell `echo -n hello`
9
10 # this is more modern way of calling for subshell with dollar and brackets
    ↪ $( )
11 echo A bash subshell $(echo -n hello)
12
13 #NOTICE: backtick might not work in future versions of bash shell

```

6.5. extension

A general convention is to give files an extension that indicates the file type. On a Linux system, this is not strictly necessary. Remember that you can always use the `file` command to determine the type of a file by scanning its contents. The system will not care if you call your script `hello_world.sh` or `hello_world`. However, it is a good practice to use an extension, as it makes it easier to identify the type of file.

We recommend to always give your scripts the `.sh` extension, but to remove the extension when you install it in a `bin` directory as a command.

6.6. shell variables

Here is a simple example of a shell variable used inside a script.

```
1 #!/bin/bash
2 # hello-user.sh -- example of a shell variable in a script
3 echo "Hello ${USER}"
```

In Bash, you can access the value of a variable by prefixing the variable name with the `$` sign. The braces are not mandatory in this case, but they are a good practice to avoid ambiguity. In some cases they are required, so it's best to be consistent in your coding style.

The variable `${USER}` is a shell variable that is defined by the system when you log in.

```
student@linux:~$ chmod +x hello-user.sh
student@linux:~$ ./hello-user.sh
Hello student
```

6.7. variable assignment

Assigning a variable is done by using the `=` operator. The variable name must start with a letter or an underscore, and can contain only letters, digits, or underscores. Remark that spaces are not allowed around the `=` sign!

```
1 #!/bin/bash
2 # hello-var.sh -- example of variable assignment
3 user="Tux"
4
5 echo "Hello ${user}"
```

Because variable names are case-sensitive, this variable `${user}` is different from `${USER}` in the previous example!

Tip: naming convention. You can use any name for a variable, but it is a good practice to use all uppercase letters for environment variables (e.g. `${USER}`) and constants and all lowercase letters for local variables (e.g. `${user}`). This is also recommended by the Google Shell Style Guide. If a variable consists of multiple words, use underscores to separate them (e.g. `${current_user}`).

Running the script:

```
student@linux:~$ chmod +x hello-var.sh
student@linux:~$ ./hello-var.sh
Hello Tux
```

Scripts can contain variables, but since scripts are run in their own subshell, the variables do not survive the end of the script.

```
student@linux:~$ echo $user

student@linux:~$ ./hello-var.sh
Hello Tux
student@linux:~$ echo $user

student@linux:~$
```

6.8. unbound variables

Remove the line `user="Tux"` from the script, or comment out the line and run it again. What do you expect to happen if the variable `user` is not assigned, but we try to use it in the script?

```
student@linux:~$ ./hello-var.sh
Hello
```

Bash will not complain if you use a variable that is not assigned, but it will simply replace the variable with an empty string. This can lead to unexpected results and is a common cause of bugs that can be hard to find. However, you can change the behavior of the shell by starting your scripts with the command `set -o nounset` (or shorter: `set -u`). This will cause the script to exit with an error if you try to use an unassigned variable.

Add the line to the script, right below the comment lines and try again!

```
1  #!/bin/bash
2  # hello-var.sh -- example of variable assignment
3
4  set -o nounset
5
6  echo "Hello ${user}"
```

Running the script:

```
student@linux:~$ ./hello-var.sh
./hello-var.sh: line 6: user: unbound variable
```

This is what you want to see. The script exits with an error, and you can see the line number where the error occurred and which variable is unbound. Start all your scripts with `set -o nounset` to prevent this kind of error!

6.9. sourcing a script

Luckily, you can force a script to run in the same shell; this is called *sourcing* a script.

```
student@linux:~$ source hello-var.sh
Hello Tux
student@linux:~$ echo $name
Tux
```

Instead of `source`, you can use the `.` (dot) command.

```
student@linux:~$ . hello-var.sh
Hello Tux
student@linux:~$ echo $name
Tux
```

6.10. quoting

Go back to `hello-user.sh` and replace the double quotes with single quotes:

```
1 #!/bin/bash
2 # hello-user.sh -- example of a shell variable in a script
3 echo 'Hello ${USER}'
```

Run the script again:

```
student@linux:~$ ./hello-user.sh
Hello ${USER}
```

What happened? By using single quotes, we turned off the shell's variable expansion. The shell will not replace `${USER}` with the value of the `USER` variable. This is why you should use double quotes when you want to use a variable.

Using quotes is important. Most of the times, when you reference the value of a variable, you should enclose it in double quotes. To illustrate this, write the following script:

```
1 #!/bin/bash
2 # create-file.sh -- example of using quotes
3 file="my file.txt"
4 touch $file
```

What we expect is that the script will create a file called `my file.txt`. However, when we run the script:

```
student@linux:~$ ./create-file.sh
student@linux:~$ ls -l
total 4
-rwxr-xr-x 1 student student 88 Mar 6 16:20 create-file.sh
-rw-r--r-- 1 student student 0 Mar 6 16:20 file.txt
-rw-r--r-- 1 student student 0 Mar 6 16:20 my
```

So actually two files were created, one named `my` and the other `file.txt`. The reason has to do with the way Bash interprets a command and how it substitutes variables. The line

```
1 touch $file
```

is expanded to

```
1 touch my file.txt
```

without the quotes. The `touch` command now sees two arguments, `my` and `file.txt`, and creates two files. To fix this, you should always use double quotes:

```
1 #!/bin/bash
2 # create-file.sh -- example of using quotes
3 file="my file.txt"
4 touch "${file}"
```

Now the expansion of the variable is done within the quotes, and the `touch` command sees only one argument.

```
student@linux:~$ ./create-file.sh
student@linux:~$ ls -l
total 4
-rwxr-xr-x 1 student student 92 Mar 6 16:20 create-file.sh
-rw-r--r-- 1 student student 0 Mar 6 16:20 'my file.txt'
```

6.11. troubleshooting a script

Another way to run a script in a separate shell is by typing `bash` with the name of the script as a parameter. Expanding this to `bash -x` allows you to see the commands that the shell is executing (after shell expansion).

Try this with the `create-file.sh` script! The incorrect version without the quotes:

```
$ bash -x create-file.sh
+ file='my file.txt'
+ touch my file.txt
```

Notice the absence of the commented (`#`) line, and the replacement of the variable in the argument `touch`.

After the fix, you get:

```
$ bash -x create-file.sh
+ file='my file.txt'
+ touch 'my file.txt'
```

Do you notice the difference?

In longer scripts, this setting produces a lot of output, which may be hard to read. You can limit the output to a specific problematic part of your script by using `set -x` and `set +x` to turn the debugging on and off.

```
1 #!/bin/bash
2 # create-file.sh -- example of using quotes
3 file="my file.txt"
4
5 set -x
6 touch "${file}"
7 set +x
```

6.12. Bash's "strict mode"

Apart from the `nounset` shell option, there are two other options that are very useful for debugging scripts: `set -o errexit` (or `set -e`) and `set -o pipefail`. The first option causes the script to exit with an error if any command fails. The second option gives better error messages when a command in a pipeline fails.

Start all your scripts with the following lines to prevent errors and to make debugging easier:

```
1 #!/bin/bash --
2 set -o nounset
3 set -o errexit
4 set -o pipefail
```

This is called "strict mode" by some. You can write this shorter in one line as `set -euo pipefail`, but this is less readable.

6.13. prevent setuid root spoofing

Some user may try to perform setuid based script root spoofing. This is a rare but possible attack. To improve script security and to avoid interpreter spoofing, you need to add `--` after the `#!/bin/bash`, which disables further option processing so the shell will not accept any options.

```
1 #!/usr/bin/env bash -
2 or
3 #!/usr/bin/env bash --
```

Any arguments after the `--` are treated as filenames and arguments. An argument of `-` is equivalent to `--`.

6.14. practice: introduction to scripting

1. Write a Python “Hello World” script, give it a shebang and make it executable. Execute it like you would a shell script and verify that this works.
2. What would happen if you remove the shebang and try to execute the script again?
3. Create a Bash script `greeting.sh` that says hello to the user (make use of the shell variable with the current user’s login name), prints the current date and time, and prints a quote, e.g.:

```
student@linux:~$ ./greeting.sh
Hello student, today is:
Wed Mar  6 09:04:19 PM UTC 2024
Quote of the day:
```

```
-----
/ Having nothing, nothing can he lose. \
|                                     |
\ -- William Shakespeare, "Henry VI" /
-----
      ^ ^
      (oo)\_____)\ \
      (__) \         )\ \
           || -----w |
           ||         ||
```

Ensure that you apply the shell settings to make your script easier to debug.

4. Copy the script to `/usr/local/bin` without the extension and verify that you can run it from any directory as a command.
5. Take another look at the script `hello-var.sh` where we printed a variable that was not assigned:

```
1 #!/bin/bash
2 # hello-var.sh -- example of variable assignment
3 # user="Tux" # Remark: this line is commented out
4
5 echo "Hello ${user}"
```

What happens if you assign the value `Tux` to the variable `user` on the interactive shell and then run the script? What do we have to do to make sure the variable is available in the script?

6. What if we change the value of the variable `user` in the script? Will this change affect the value of the variable in the interactive shell after the script is finished?

6.15. solution: introduction to scripting

1. Write a Python Hello World script, give it a shebang and make it executable.

```

1 #!/usr/bin/python3
2 print("Hello, World!")

$ chmod +x hello.py
$ ./hello.py
Hello, World!

```

2. What would happen if you remove the shebang and try to execute the script again?

The script will be executed by the default interpreter, in this case, the Bash shell, which will not understand the Python syntax.

```

$ ./hello.py
./hello.py: line 1: syntax error near unexpected token `"Hello world!'"
./hello.py: line 1: `print("Hello world!")'

```

3. Create a Bash script `greeting.sh` that says hello to the user (make use of the shell variable with the current user's login name), prints the current date and time, and prints a quote. Ensure that you apply the shell settings to make your script easier to debug.

```

1 #! /bin/bash --
2
3 set -o nounset
4 set -o errexit
5 set -o pipefail
6
7 echo "Hello ${USER}, today is:"
8 date
9 echo "Quote of the day:"
10 fortune | cowsay

```

4. Copy the script to `/usr/local/bin` without the extension and verify that you can run it from any directory as a command.

```

student@linux:~$ sudo cp greeting.sh /usr/local/bin/greeting
student@linux:~$ greeting
Hello student, today is:
Wed Mar  6 09:17:00 PM UTC 2024
Quote of the day:

```

```

-----
/ You plan things that you do not even \
| attempt because of your extreme      |
\ caution.                             /
-----

```

```

\      ^__^
 (oo)\_____)
 (_____)  )\/\
      ||----w |
      ||     ||

```

```

student@linux:~$ cd /tmp
student@linux:/tmp$ greeting
Hello student, today is:
Wed Mar  6 09:17:08 PM UTC 2024
Quote of the day:


```

```

-----
< You will be successful in love. >
-----

```

6. introduction to scripting



5. Take another look at the script `hello-var.sh` where we printed a variable that was not assigned. What happens if you assign the value `Tux` to the variable `user` on the interactive shell and then run the script? What do we have to do to make sure the variable is available in the script?

```
student@linux:~$ ./hello-var.sh
Hello
student@linux:~$ user=Tux
student@linux:~$ ./hello-var.sh
Hello
student@linux:~$ export user
student@linux:~$ ./hello-var.sh
Hello Tux
```

6. What if we change the value of the variable `user` in the script? Will this change affect the value of the variable in the interactive shell after the script is finished?

We change the script to:

```
1 #!/bin/bash
2 # hello-var.sh -- example of variable assignment
3 user="Linus"
4
5 echo "Hello ${user}"
```

And execute it:

```
student@linux:~$ export user=Tux
student@linux:~$ echo $user
Tux
student@linux:~$ ./hello-var.sh
Hello Linus
student@linux:~$ echo $user
Tux
```

The change in the script does not affect the value of the variable in the interactive shell after the script is finished!

Part IV.

Organising users

7. standard file permissions

(Written by Paul Cobbaut, <https://github.com/paulcobbaut/>, with contributions by: Alex M. Schapelle, <https://github.com/zero-pythagoras/>, Bert Van Vreckem, <https://github.com/bertvv/>)

This chapter contains details about basic file security through *file ownership* and *file permissions*.

7.1. file ownership

7.1.1. user owner and group owner

The *users* and *groups* of a system can be locally managed in `/etc/passwd` and `/etc/group`, or they can be in a NIS, LDAP, or Samba domain. These users and groups can *own* files. Actually, every file has a *user owner* and a *group owner*, as can be seen in the following example.

```
student@linux:~/owners$ ls -lh
total 636K
-rw-r--r--. 1 student snooker 1.1K Apr  8 18:47 data.odt
-rw-r--r--. 1 student student 626K Apr  8 18:46 file1
-rw-r--r--. 1 student tennis  185 Apr  8 18:46 file2
-rw-rw-r--. 1 root root      0 Apr  8 18:47 stuff.txt
```

User `student` owns three files: `file1` has `student` as *user owner* and has the group `student` as *group owner*, `data.odt` is *group owned* by the group `snooker`, `file2` by the group `tennis`.

The last file is called `stuff.txt` and is owned by the `root` user and the `root` group.

7.1.2. chgrp

You can change the group owner of a file using the `chgrp` command. You must have root privileges to do this.

```
root@linux:/home/student/owners# ls -l file2
-rw-r--r--. 1 root tennis 185 Apr  8 18:46 file2
root@linux:/home/student/owners# chgrp snooker file2
root@linux:/home/student/owners# ls -l file2
-rw-r--r--. 1 root snooker 185 Apr  8 18:46 file2
root@linux:/home/student/owners#
```

7. standard file permissions

7.1.3. chown

The user owner of a file can be changed with `chown` command. You must have root privileges to do this. In the following example, the user owner of `file2` is changed from `root` to `student`.

```
root@linux:/home/student# ls -l FileForStudent
-rw-r--r-- 1 root student 0 2008-08-06 14:11 FileForStudent
root@linux:/home/student# chown student FileForStudent
root@linux:/home/student# ls -l FileForStudent
-rw-r--r-- 1 student student 0 2008-08-06 14:11 FileForStudent
```

You can also use `chown user:group` to change both the user owner and the group owner.

```
root@linux:/home/student# ls -l FileForStudent
-rw-r--r-- 1 student student 0 2008-08-06 14:11 FileForStudent
root@linux:/home/student# chown root:project42 FileForStudent
root@linux:/home/student# ls -l FileForStudent
-rw-r--r-- 1 root project42 0 2008-08-06 14:11 FileForStudent
```

7.2. list of special files

When you use `ls -l`, for each file you can see ten characters before the user and group owner. The first character tells us the type of file. Regular files get a `-`, directories get a `d`, symbolic links are shown with an `l`, pipes get a `p`, character devices a `c`, block devices a `b`, and sockets an `s`.

first character	file type
-	normal file
d	directory
l	symbolic link
p	named pipe
b	block device
c	character device
s	socket

Below an example of a character device (the console) and a block device (the hard disk).

```
student@linux:~$ ls -l /dev/console /dev/sda
crw--w---- 1 root tty 5, 1 Mar 8 08:32 /dev/console
brw-rw---- 1 root disk 8, 0 Mar 8 08:32 /dev/sda
```

And here you can see a directory, a regular file and a symbolic link.

```
student@linux:~$ ls -ld /etc /etc/hosts /etc/os-release
drwxr-xr-x 81 root root 4096 Mar 8 08:32 /etc
-rw-r--r-- 1 root root 186 Feb 26 14:58 /etc/hosts
lrwxrwxrwx 1 root root 21 Dec 9 21:08 /etc/os-release -> ../usr/lib/os-release
```

7.3. permissions

7.3.1. rwx

The nine characters following the file type denote the permissions in three triplets. A permission can be **r** for **r**ead access, **w** for **w**rite access, and **x** for **e**xecute. You need the **r** permission to list (**ls**) the contents of a directory. You need the **x** permission to enter (**cd**) a directory. You need the **w** permission to create files in or remove files from a directory.

permission	on a file	on a directory
r ead	read file contents (cat)	read directory contents (ls)
w rite	change file contents	create/delete files (touch , rm)
x ecute	execute the file	enter the directory (cd)

7.3.2. three sets of rwx

We already know that the output of **ls -l** starts with ten characters for each file. This example shows a regular file (because the first character is a **-**).

```
student@linux:~/test$ ls -l proc42.sh
-rwxr-xr-- 1 student proj 984 Feb 6 12:01 proc42.sh
```

Below is a table describing the function of all ten characters.

position	characters	function
1	-	file type
2-4	rwx	permissions for the <i>user owner</i>
5-7	r-x	permissions for the <i>group owner</i>
8-10	r--	permissions for <i>others</i>

When you are the *user owner* of a file, then the *user owner permissions* apply to you. The rest of the permissions have no influence on your access to the file.

When you belong to the *group* that is the *group owner* of a file, then the *group owner permissions* apply to you. The rest of the permissions have no influence on your access to the file.

When you are not the *user owner* of a file and you do not belong to the *group owner*, then the *others permissions* apply to you. The rest of the permissions have no influence on your access to the file.

7.3.3. permission examples

Some example combinations on files and directories are seen in this example. The name of the file explains the permissions.

```
student@linux:~/perms$ ls -lh
total 12K
drwxr-xr-x 2 student student 4.0K 2007-02-07 22:26 AllEnter_UserCreateDelete
-rwxrwxrwx 1 student student 0 2007-02-07 22:21 EveryoneFullControl.txt
-r--r----- 1 student student 0 2007-02-07 22:21 OnlyOwnersRead.txt
-rwxrwx--- 1 student student 0 2007-02-07 22:21 OwnersAll_RestNothing.txt
dr-xr-x--- 2 student student 4.0K 2007-02-07 22:25 UserAndGroupEnter
dr-x----- 2 student student 4.0K 2007-02-07 22:25 OnlyUserEnter
```

7. standard file permissions

To summarise, the first `rwX` triplet represents the permissions for the *user owner*. The second triplet corresponds to the *group owner*; it specifies permissions for all members of that group. The third triplet defines permissions for all *other* users that are not the *user owner* and are not a member of the *group owner*. The root user ignores all restrictions and can do anything with any file.

7.3.4. setting permissions with symbolic notation

Permissions can be changed with `chmod MODE FILE ...`. You need to be the owner of the file to do this. The first example gives (+) the *user owner* (u) execute (x) permissions.

```
student@linux:~/perms$ ls -l permissions.txt
-rw-r--r-- 1 student student 0 2007-02-07 22:34 permissions.txt
student@linux:~/perms$ chmod u+x permissions.txt
student@linux:~/perms$ ls -l permissions.txt
-rwxr--r-- 1 student student 0 2007-02-07 22:34 permissions.txt
```

This example removes (-) the group owner's (g) read (r) permission.

```
student@linux:~/perms$ chmod g-r permissions.txt
student@linux:~/perms$ ls -l permissions.txt
-rwx---r-- 1 student student 0 2007-02-07 22:34 permissions.txt
```

This example removes (-) the other's (o) read (r) permission.

```
student@linux:~/perms$ chmod o-r permissions.txt
student@linux:~/perms$ ls -l permissions.txt
-rwx----- 1 student student 0 2007-02-07 22:34 permissions.txt
```

This example gives (+) all (a) of them the write (w) permission.

```
student@linux:~/perms$ chmod a+w permissions.txt
student@linux:~/perms$ ls -l permissions.txt
-rwx-w--w- 1 student student 0 2007-02-07 22:34 permissions.txt
```

You don't even have to type the a.

```
student@linux:~/perms$ chmod +x permissions.txt
student@linux:~/perms$ ls -l permissions.txt
-rwx-wx-wx 1 student student 0 2007-02-07 22:34 permissions.txt
```

You can also set explicit permissions with =.

```
student@linux:~/perms$ chmod u=rw permissions.txt
student@linux:~/perms$ ls -l permissions.txt
-rw--wx-wx 1 student student 0 2007-02-07 22:34 permissions.txt
```

Feel free to make any kind of combination, separating them with a comma. Remark that spaces are **not** allowed!

```
student@linux:~/perms$ chmod u=rw,g=rw,o=r permissions.txt
student@linux:~/perms$ ls -l permissions.txt
-rw-rw-r-- 1 student student 0 2007-02-07 22:34 permissions.txt
```

Even fishy combinations are accepted by `chmod`.

```
student@linux:~/perms$ chmod u=rwx,ug+rw,o=r permissions.txt
student@linux:~/perms$ ls -l permissions.txt
-rwxrw-r-- 1 student student 0 2007-02-07 22:34 permissions.txt
```

Summarized, in order to change permissions with `chmod` using symbolic notation:

- first specify who the permissions are for: `u` for the user owner, `g` for the group owner, `o` for others, and `a` for all. `a` is the default and can be omitted.
- then specify the operation: `+` to add permissions, `-` to remove permissions, and `=` to set permissions.
- finally specify the permission(s): `r` for read, `w` for write, and `x` for execute.
- multiple operations can be combined with a comma (no spaces!)

7.3.5. setting permissions with octal notation

Most Unix administrators will use the “old school” octal system to talk about and set permissions. Consider the triplet to be a binary number with 0 indicating the permission is not set and 1 indicating the permission is set. You then have $2^3 = 8$ possible combinations, hence the name *octal*. You can then convert the binary number to an octal number, equating `r` to 4, `w` to 2, and `x` to 1.

permission	binary	octal
---	000	0
--x	001	1
-w-	010	2
-wx	011	3
r--	100	4
r-x	101	5
rw-	110	6
rwx	111	7

Since we have three triplets, we can use three octal digits to represent the permissions. This makes 777 equal to `rwxrwxrwx` and by the same logic, 654 mean `rw-r-xr--`. The `chmod` command will accept these numbers.

```
student@linux:~/perms$ chmod 777 permissions.txt
student@linux:~/perms$ ls -l permissions.txt
-rwxrwxrwx 1 student student 0 2007-02-07 22:34 permissions.txt
student@linux:~/perms$ chmod 664 permissions.txt
student@linux:~/perms$ ls -l permissions.txt
-rw-rw-r-- 1 student student 0 2007-02-07 22:34 permissions.txt
student@linux:~/perms$ chmod 750 permissions.txt
student@linux:~/perms$ ls -l permissions.txt
-rwxr-x--- 1 student student 0 2007-02-07 22:34 permissions.txt
```

Remark that in practice, some combinations will never occur:

- The permissions of a user will never be smaller than the permissions of the group owner or others. Consequently, the digits will always be in descending order.
- Setting the write or execute permission without read access is useless. Consequently, you will never use 1, 2, or 3 in an octal permission code

7. standard file permissions

- A directory will always have the read and execute permission set or unset together. It is useless to allow a user to read the directory contents, but not let them cd into that directory. Allowing cd without read access is also useless. The permission code for a directory will therefore always be odd.

Here's a little tip: you can print the permissions of a file in either octal or symbolic notation with the stat command (check the man page of stat to see how this works).

```
[student@linux ~]$ stat -c '%A %a' /etc/passwd
-rw-r--r-- 644
[student@linux ~]$ stat -c '%A %a' /etc/shadow
----- 0
[student@linux ~]$ stat -c '%A %a' /bin/ls
-rwxr-xr-x 755
```

7.3.6. umask

When creating a file or directory, a set of default permissions are applied. These default permissions are determined by the umask value. The umask specifies permissions that you do not want set on by default. You can display the umask with the umask command.

```
[student@linux ~]$ umask
0002
[student@linux ~]$ touch test
[student@linux ~]$ ls -l test
-rw-rw-r-- 1 student student 0 Jul 24 06:03 test
[student@linux ~]$
```

As you can also see, the file is also not executable by default. This is a general security feature among Unixes; newly created files are never executable by default. You have to explicitly do a `chmod +x` to make a file executable. This also means that the 1 bit in the umask has no meaning. A umask value of 0022 has the same effect as 0033.

In practice, you will only use umask values:

- 0: don't take away any permissions
- 2: take away write permissions
- 7: take away all permissions

You can set the umask value to a new value with the umask command. The umask value is a four-digit octal number. The first digit is for special permissions (and is always zero), the second for the user permissions (is in practice always 0, since there is no use in taking away the user's permissions), the third for the group owner (sometimes 0, but usually 2 or 7), and the last for others (usually 2 or 7, 0 is very uncommon and can be considered to be a security risk).

The umask value is subtracted from 777 to get the default permissions and in the case of a file, the execute bit is removed.

```
[student@linux ~]$ umask 0002
[student@linux ~]$ touch file0002
[student@linux ~]$ mkdir dir0002
[student@linux ~]$ ls -ld *0002
drwxrwxr-x. 2 student student 6 Mar  8 10:48 dir0002
-rw-rw-r--. 1 student student 0 Mar  8 10:47 file0002
[student@linux ~]$ umask 0027
[student@linux ~]$ touch file0027
[student@linux ~]$ mkdir dir0027
```



```
[student@linux ~]$ ls -ld *0027
drwxr-x---. 2 student student 6 Mar  8 10:48 dir0027
-rw-r-----. 1 student student 0 Mar  8 10:48 file0027
[student@linux ~]$ umask 0077
[student@linux ~]$ touch file0077
[student@linux ~]$ mkdir dir0077
[student@linux ~]$ ls -ld *0077
drwx-----. 2 student student 6 Mar  8 10:51 dir0077
-rw-----. 1 student student 0 Mar  8 10:51 file0077
```

7.3.7. mkdir -m

When creating directories with `mkdir` you can use the `-m` option to set the mode. This example explains.

```
student@linux~$ mkdir -m 700 MyDir
student@linux~$ mkdir -m 777 Public
student@linux~$ ls -dl MyDir/ Public/
drwx----- 2 student student 4096 2011-10-16 19:16 MyDir/
drwxrwxrwx 2 student student 4096 2011-10-16 19:16 Public/
```

7.3.8. cp -p

To preserve permissions and time stamps from source files, use `cp -p`.

```
student@linux:~/perms$ cp file* cp
student@linux:~/perms$ cp -p file* cpp
student@linux:~/perms$ ll *
-rwx----- 1 student student 0 2008-08-25 13:26 file33
-rwxr-x--- 1 student student 0 2008-08-25 13:26 file42

cp:
total 0
-rwx----- 1 student student 0 2008-08-25 13:34 file33
-rwxr-x--- 1 student student 0 2008-08-25 13:34 file42

cpp:
total 0
-rwx----- 1 student student 0 2008-08-25 13:26 file33
-rwxr-x--- 1 student student 0 2008-08-25 13:26 file42
```

7.4. practice: standard file permissions

1. As normal user, create a directory `~/permissions`. Create a file owned by yourself in there.
2. Copy a file owned by root from `/etc/` to your permissions dir, who owns this file now ?
3. As root, create a file in the users `~/permissions` directory.
4. As normal user, look at who owns this file created by root.
5. Change the ownership of all files in `~/permissions` to yourself.
6. Delete the file created by root. Is this possible?

7. standard file permissions

7. With chmod, is 770 the same as `rw-rwx---`?
8. With chmod, is 664 the same as `r-xr-xr--`?
9. With chmod, is 400 the same as `r-----`?
10. With chmod, is 734 the same as `rw-r-xr--`?
11. Display the umask value in octal and in symbolic form.
12. Set the umask to 0077, but use the symbolic format to set it. Verify that this works.
13. Create a file as root, give only read to others. Can a normal user read this file? Test writing to this file with `vi` or `nano`.
14. Create a file as a normal user, take away all permissions for the group owner and others. Can you still read the file? Can root read the file? Can root write to the file?
15. Create a directory that belongs to group users, where every member of that group can read and write to files, and create files. Make sure that people can only delete their own files.

7.5. solution: standard file permissions

1. As normal user, create a directory `~/permissions`. Create a file owned by yourself in there.

```
[student@linux ~]$ mkdir permissions
[student@linux ~]$ touch permissions/myfile.txt
[student@linux ~]$ ls -l permissions/
total 0
-rw-r--r--. 1 student student 0 Mar  8 10:59 myfile.txt
```

2. Copy a file owned by root from `/etc/` to your `permissions` dir, who owns this file now ?

```
[student@linux ~]$ ls -l /etc/hosts
-rw-r--r--. 1 root root 174 Feb 26 15:05 /etc/hosts
[student@linux ~]$ cp /etc/hosts ~/permissions/
[student@linux ~]$ ls -l permissions/hosts
-rw-r--r--. 1 student student 174 Mar  8 11:00 permissions/hosts
```

The copy is owned by you.

3. As root, create a file in the `users ~/permissions` directory.

```
[student@linux ~]$ sudo touch permissions/rootfile.txt
[sudo] password for student:
```

4. As normal user, look at who owns this file created by root.

```
[student@linux ~]$ ls -l permissions/*.txt
-rw-r--r--. 1 student student 0 Mar  8 10:59 permissions/myfile.txt
-rw-r--r--. 1 root    root    0 Mar  8 11:02 permissions/rootfile.txt
```

The file created by root is owned by root.

5. Change the ownership of all files in `~/permissions` to yourself.

```
[student@linux ~]$ chown student ~/permissions/*
chown: changing ownership of '/home/student/permissions/rootfile.txt': Operation not p
```

You cannot become owner of the file that belongs to root. Root must change the ownership.

6. Delete the file created by root. Is this possible?

```
[student@linux ~]$ rm ~/permissions/rootfile.txt
rm: remove write-protected regular empty file '/home/student/permissions/rootfile.txt'
[student@linux ~]$ ls -l permissions/*.txt
-rw-r--r--. 1 student student 0 Mar  8 10:59 permissions/myfile.txt
```

You can delete the file since you have write permission on the directory!

7. With chmod, is 770 the same as `rw-rwx---`?

yes

8. With chmod, is 664 the same as `r-xr-xr--`?

no, `rw-rw-r--` is 664 and `r-xr-xr--` is 774

9. With chmod, is 400 the same as `r-----`?

yes

10. With chmod, is 734 the same as `rw-r-xr--`?

no, `rw-r-xr--` is 754 and `rw-x-wr--` is 734

11. Display the umask in octal and in symbolic form.

```
umask and umask -S
```

12. Set the umask to 0077, but use the symbolic format to set it. Verify that this works.

```
[student@linux ~]$ umask -S u=rwx,g=,o=
u=rwx,g=,o=
[student@linux ~]$ umask
0077
```

13. Create a file as root, give only read to others. Can a normal user read this file? Test writing to this file with vi or nano.

```
[student@linux ~]$ sudo vi permissions/rootfile.txt
[student@linux ~]$ sudo chmod 644 permissions/rootfile.txt
[student@linux ~]$ ls -l permissions/*.txt
-rw-r--r--. 1 student student 0 Mar  8 10:59 permissions/myfile.txt
-rw-r--r--. 1 root    root    6 Mar  8 13:53 permissions/rootfile.txt
[student@linux ~]$ cat permissions/rootfile.txt
hello
[student@linux ~]$ echo " world" >> permissions/rootfile.txt
-bash: permissions/rootfile.txt: Permission denied
```

Yes, a normal user can read the file, but not write to it.

14. Create a file as a normal user, take away all permissions for the group and others. Can you still read the file? Can root read the file? Can root write to the file?

```
[student@linux ~]$ vi permissions/privatefile.txt
... (editing the file) ...
[student@linux ~]$ cat permissions/privatefile.txt
hello
[student@linux ~]$ chmod 600 permissions/privatefile.txt
[student@linux ~]$ ls -l permissions/privatefile.txt
-rw-----. 1 student student 0 Mar  8 16:06 permissions/privatefile.txt
[student@linux ~]$ cat permissions/privatefile.txt
hello
```

Of course, the owner can still read (and write to) the file.

7. standard file permissions

```
[student@linux ~]$ sudo vi permissions/privatefile.txt
[sudo] password for student:
... (editing the file) ...
[student@linux ~]$ cat permissions/privatefile.txt
hello world
```

Root can read and write to the file. In fact, root ignores all file permissions and can do anything with any file.

15. Create a directory `shared/` that belongs to group `users`, where every member of that group can read and write to files, and create files.

```
[student@linux ~]$ mkdir shared
[student@linux ~]$ sudo chgrp users shared
[student@linux ~]$ chmod 775 shared/
[student@linux ~]$ ls -ld shared/
drwxrwxr-x. 2 student users 6 Mar  8 18:26 shared/
```

8. advanced file permissions

(Written by Paul Cobbaut, <https://github.com/paulcobbaut/>, with contributions by: Alex M. Schapelle, <https://github.com/zero-pythagoras/>)

8.1. sticky bit on directory

You can set the `sticky` bit on a directory to prevent users from removing files that they do not own as a user owner. The sticky bit is displayed at the same location as the `x` permission for others. The sticky bit is represented by a `t` (meaning `x` is also there) or a `T` (when there is no `x` for others).

```
root@linux:~# mkdir /project55
root@linux:~# ls -ld /project55
drwxr-xr-x  2 root root 4096 Feb  7 17:38 /project55
root@linux:~# chmod +t /project55/
root@linux:~# ls -ld /project55
drwxr-xr-t  2 root root 4096 Feb  7 17:38 /project55
root@linux:~#
```

The `sticky` bit can also be set with octal permissions, it is binary 1 in the first of four triplets.

```
root@linux:~# chmod 1775 /project55/
root@linux:~# ls -ld /project55
drwxrwxr-t  2 root root 4096 Feb  7 17:38 /project55
root@linux:~#
```

You will typically find the `sticky` bit on the `/tmp` directory.

```
root@linux:~# ls -ld /tmp
drwxrwxrwt 6 root root 4096 2009-06-04 19:02 /tmp
```

8.2. setgid bit on directory

`setgid` can be used on directories to make sure that all files inside the directory are owned by the group owner of the directory. The `setgid` bit is displayed at the same location as the `x` permission for group owner. The `setgid` bit is represented by an `s` (meaning `x` is also there) or a `S` (when there is no `x` for the group owner). As this example shows, even though `root` does not belong to the group `proj55`, the files created by `root` in `/project55` will belong to `proj55` since the `setgid` is set.

8. advanced file permissions

```
root@linux:~# groupadd proj55
root@linux:~# chown root:proj55 /project55/
root@linux:~# chmod 2775 /project55/
root@linux:~# touch /project55/fromroot.txt
root@linux:~# ls -ld /project55/
drwxrwsr-x 2 root proj55 4096 Feb  7 17:45 /project55/
root@linux:~# ls -l /project55/
total 4
-rw-r--r-- 1 root proj55 0 Feb  7 17:45 fromroot.txt
root@linux:~#
```

You can use the `find` command to find all `setgid` directories.

```
student@linux:~$ find / -type d -perm -2000 2> /dev/null
/var/log/mysql
/var/log/news
/var/local
...
```

8.3. setgid and setuid on regular files

These two permissions cause an executable file to be executed with the permissions of the file owner instead of the executing owner. This means that if any user executes a program that belongs to the `root` user, and the `setuid` bit is set on that program, then the program runs as `root`. This can be dangerous, but sometimes this is good for security.

Take the example of passwords; they are stored in `/etc/shadow` which is only readable by `root`. (The `root` user never needs permissions anyway.)

```
root@linux:~# ls -l /etc/shadow
-r----- 1 root root 1260 Jan 21 07:49 /etc/shadow
```

Changing your password requires an update of this file, so how can normal non-root users do this? Let's take a look at the permissions on the `/usr/bin/passwd`.

```
root@linux:~# ls -l /usr/bin/passwd
-r-s--x--x 1 root root 21200 Jun 17 2005 /usr/bin/passwd
```

When running the `passwd` program, you are executing it with `root` credentials.

You can use the `find` command to find all `setuid` programs.

```
student@linux:~$ find /usr/bin -type f -perm -04000
/usr/bin/arping
/usr/bin/kgrantpty
/usr/bin/newgrp
/usr/bin/chfn
/usr/bin/sudo
/usr/bin/fping6
/usr/bin/passwd
/usr/bin/gpasswd
...
```

In most cases, setting the `setuid` bit on executables is sufficient. Setting the `setgid` bit will result in these programs to run with the credentials of their group owner.

8.4. setuid on sudo

The sudo binary has the setuid bit set, so any user can run it with the effective userid of root.

```
student@linux:~$ ls -l $(which sudo)
---s--x--x. 1 root root 123832 Oct  7  2013 /usr/bin/sudo
student@linux:~$
```

8.5. practice: sticky, setuid and setgid bits

1a. Set up a directory, owned by the group sports.

1b. Members of the sports group should be able to create files in this directory.

1c. All files created in this directory should be group-owned by the sports group.

1d. Users should be able to delete only their own user-owned files.

1e. Test that this works!

2. Verify the permissions on `/usr/bin/passwd`. Remove the setuid, then try changing your password as a normal user. Reset the permissions back and try again.

3. If time permits (or if you are waiting for other students to finish this practice), read about file attributes in the man page of `chattr` and `lsattr`. Try setting the `i` attribute on a file and test that it works.

8.6. solution: sticky, setuid and setgid bits

1a. Set up a directory, owned by the group sports.

```
groupadd sports
```

```
mkdir /home/sports
```

```
chown root:sports /home/sports
```

1b. Members of the sports group should be able to create files in this directory.

```
chmod 770 /home/sports
```

1c. All files created in this directory should be group-owned by the sports group.

```
chmod 2770 /home/sports
```

1d. Users should be able to delete only their own user-owned files.

```
chmod +t /home/sports
```

8. advanced file permissions

1e. Test that this works!

Log in with different users (group members and others and root), create files and watch the permissions. Try changing and deleting files...

2. Verify the permissions on `/usr/bin/passwd`. Remove the `setuid`, then try changing your password as a normal user. Reset the permissions back and try again.

```
root@linux:~# ls -l /usr/bin/passwd
-rwsr-xr-x 1 root root 31704 2009-11-14 15:41 /usr/bin/passwd
root@linux:~# chmod 755 /usr/bin/passwd
root@linux:~# ls -l /usr/bin/passwd
-rwxr-xr-x 1 root root 31704 2009-11-14 15:41 /usr/bin/passwd
```

A normal user cannot change password now.

```
root@linux:~# chmod 4755 /usr/bin/passwd
root@linux:~# ls -l /usr/bin/passwd
-rwsr-xr-x 1 root root 31704 2009-11-14 15:41 /usr/bin/passwd
```

3. If time permits (or if you are waiting for other students to finish this practice), read about file attributes in the man page of `chattr` and `lsattr`. Try setting the `i` attribute on a file and test that it works.

```
student@linux:~$ sudo su -
[sudo] password for paul:
root@linux:~# mkdir attr
root@linux:~# cd attr/
root@linux:~/attr# touch file42
root@linux:~/attr# lsattr
----- ./file42
root@linux:~/attr# chattr +i file42
root@linux:~/attr# lsattr
----i----- ./file42
root@linux:~/attr# rm -rf file42
rm: cannot remove `file42': Operation not permitted
root@linux:~/attr# chattr -i file42
root@linux:~/attr# rm -rf file42
root@linux:~/attr#
```


9. introduction to users

(Written by Paul Cobbaut, <https://github.com/paulcobbaut/>, with contributions by: Alex M. Schapelle, <https://github.com/zero-pythagoras/>)

This little chapter will teach you how to identify your user account on a Unix computer using commands like `who`, `am i`, `id`, and more.

In a second part you will learn how to become another user with the `su` command.

And you will learn how to run a program as another user with `sudo`.

9.1. whoami

The `whoami` command tells you your username.

```
[student@linux ~]$ whoami
paul
[student@linux ~]$
```

9.2. who

The `who` command will give you information about who is logged on the system.

```
[student@linux ~]$ who
root    pts/0      2014-10-10 23:07 (10.104.33.101)
paul    pts/1      2014-10-10 23:30 (10.104.33.101)
laura   pts/2      2014-10-10 23:34 (10.104.33.96)
tania   pts/3      2014-10-10 23:39 (10.104.33.91)
[student@linux ~]$
```

9.3. who am i

With `who am i` the `who` command will display only the line pointing to your current session.

```
[student@linux ~]$ who am i
paul    pts/1      2014-10-10 23:30 (10.104.33.101)
[student@linux ~]$
```

9.4. w

The `w` command shows you who is logged on and what they are doing.

```
[student@linux ~]$ w
 23:34:07 up 31 min,  2 users,  load average: 0.00, 0.01, 0.02
USER      TTY      LOGIN@  IDLE   JCPU   PCPU WHAT
root      pts/0    23:07   15.00s 0.01s  0.01s top
paul      pts/1    23:30    7.00s 0.00s  0.00s w
[student@linux ~]$
```

9.5. id

The `id` command will give you your user id, primary group id, and a list of the groups that you belong to.

```
student@linux:~$ id
uid=1000(paul) gid=1000(paul) groups=1000(paul)
```

On RHEL/CentOS you will also get SELinux context information with this command.

```
[root@linux ~]# id
uid=0(root) gid=0(root) groups=0(root) context=unconfined_u:unconfined_r\
:unconfined_t:s0-s0:c0.c1023
```

9.6. su to another user

The `su` command allows a user to run a shell as another user.

```
laura@linux:~$ su tania
Password:
tania@linux:/home/laura$
```

9.7. su to root

Yes you can also `su` to become `root`, when you know the `root` password.

```
laura@linux:~$ su root
Password:
root@linux:/home/laura#
```

9.8. su as root

You need to know the password of the user you want to substitute to, unless your are logged in as `root`. The `root` user can become any existing user without knowing that user's password.

```
root@linux:~# id
uid=0(root) gid=0(root) groups=0(root)
root@linux:~# su - valentina
valentina@linux:~$
```

9.9. su - \$username

By default, the su command maintains the same shell environment. To become another user and also get the target user's environment, issue the su - command followed by the target username.

```
root@linux:~# su laura
laura@linux:/root$ exit
exit
root@linux:~# su - laura
laura@linux:~$ pwd
/home/laura
```

9.10. su -

When no username is provided to su or su -, the command will assume root is the target.

```
tania@linux:~$ su -
Password:
root@linux:~#
```

9.11. run a program as another user

The sudo program allows a user to start a program with the credentials of another user. Before this works, the system administrator has to set up the /etc/sudoers file. This can be useful to delegate administrative tasks to another user (without giving the root password).

The screenshot below shows the usage of sudo. User paul received the right to run useradd with the credentials of root. This allows paul to create new users on the system without becoming root and without knowing the root password.

First the command fails for paul.

```
student@linux:~$ /usr/sbin/useradd -m valentina
useradd: Permission denied.
useradd: cannot lock /etc/passwd; try again later.
```

But with sudo it works.

```
student@linux:~$ sudo /usr/sbin/useradd -m valentina
[sudo] password for paul:
student@linux:~$
```

9.12. visudo

Check the man page of visudo before playing with the /etc/sudoers file. Editing the sudoers is out of scope for this fundamentals book.

```
student@linux:~$ apropos visudo
visudo          (8) - edit the sudoers file
student@linux:~$
```

9.13. sudo su -

On some Linux systems like Ubuntu and Xubuntu, the `root` user does not have a password set. This means that it is not possible to login as `root` (extra security). To perform tasks as `root`, the first user is given all `sudo` rights via the `/etc/sudoers`. In fact all users that are members of the `admin` group can use `sudo` to run all commands as `root`.

```
root@linux:~# grep admin /etc/sudoers
# Members of the admin group may gain root privileges
%admin ALL=(ALL) ALL
```

The end result of this is that the user can type `sudo su -` and become `root` without having to enter the `root` password. The `sudo` command does require you to enter your own password. Thus the password prompt in the screenshot below is for `sudo`, not for `su`.

```
student@linux:~$ sudo su -
Password:
root@linux:~#
```

9.14. sudo logging

Using `sudo` without authorization will result in a severe warning:

```
student@linux:~$ sudo su -
```

We trust you have received the usual lecture from the local System Administrator. It usually boils down to these three things:

- #1) Respect the privacy of others.
- #2) Think before you type.
- #3) With great power comes great responsibility.

```
[sudo] password for paul:
paul is not in the sudoers file. This incident will be reported.
student@linux:~$
```

The `root` user can see this in the `/var/log/secure` on Red Hat and in `/var/log/auth.log` on Debian).

```
root@linux:~# tail /var/log/secure | grep sudo | tr -s ' '
Apr 13 16:03:42 rhel65 sudo: paul : user NOT in sudoers ; TTY=pts/0 ; PWD=\
/home/paul ; USER=root ; COMMAND=/bin/su -
root@linux:~#
```

9.15. practice: introduction to users

1. Run a command that displays only your currently logged on user name.
2. Display a list of all logged on users.
3. Display a list of all logged on users including the command they are running at this very moment.
4. Display your user name and your unique user identification (userid).

5. Use `su` to switch to another user account (unless you are root, you will need the password of the other account). And get back to the previous account.

6. Now use `su -` to switch to another user and notice the difference.

Note that `su -` gets you into the home directory of Tania.

7. Try to create a new user account (when using your normal user account). this should fail. (Details on adding user accounts are explained in the next chapter.)

8. Now try the same, but with `sudo` before your command.

9.16. solution: introduction to users

1. Run a command that displays only your currently logged on user name.

```
laura@linux:~$ whoami
laura
laura@linux:~$ echo $USER
laura
```

2. Display a list of all logged on users.

```
laura@linux:~$ who
laura pts/0 2014-10-13 07:22 (10.104.33.101)
laura@linux:~$
```

3. Display a list of all logged on users including the command they are running at this very moment.

```
laura@linux:~$ w
 07:47:02 up 16 min,  2 users,  load average: 0.00, 0.00, 0.00
USER      TTY      FROM          LOGIN@   IDLE   JCPU   PCPU   WHAT
root      pts/0    10.104.33.101 07:30    6.00s  0.04s  0.00s  w
root      pts/1    10.104.33.101 07:46    6.00s  0.01s  0.00s  sleep 42
laura@linux:~$
```

4. Display your user name and your unique user identification (userid).

```
laura@linux:~$ id
uid=1005(laura) gid=1007(laura) groups=1007(laura)
laura@linux:~$
```

5. Use `su` to switch to another user account (unless you are root, you will need the password of the other account). And get back to the previous account.

```
laura@linux:~$ su tania
Password:
tania@linux:/home/laura$ id
uid=1006(tania) gid=1008(tania) groups=1008(tania)
tania@linux:/home/laura$ exit
laura@linux:~$
```

6. Now use `su -` to switch to another user and notice the difference.

9. introduction to users

```
laura@linux:~$ su - tania
Password:
tania@linux:~$ pwd
/home/tania
tania@linux:~$ logout
laura@linux:~$
```

Note that `su -` gets you into the home directory of Tania.

7. Try to create a new user account (when using your normal user account). this should fail. (Details on adding user accounts are explained in the next chapter.)

```
laura@linux:~$ useradd valentina
-su: useradd: command not found
laura@linux:~$ /usr/sbin/useradd valentina
useradd: Permission denied.
useradd: cannot lock /etc/passwd; try again later.
```

It is possible that `useradd` is located in `/sbin/useradd` on your computer.

8. Now try the same, but with `sudo` before your command.

```
laura@linux:~$ sudo /usr/sbin/useradd valentina
[sudo] password for laura:
laura is not in the sudoers file. This incident will be reported.
laura@linux:~$
```

Notice that `laura` has no permission to use the `sudo` on this system.

10. user management

(Written by Paul Cobbaut, <https://github.com/paulcobbaut/>, with contributions by: Alex M. Schapelle, <https://github.com/zero-pythagoras/>)

This chapter will teach you how to use `useradd`, `usermod` and `userdel` to create, modify and remove user accounts.

You will need `root` access on a Linux computer to complete this chapter.

10.1. user management

User management on Linux can be done in three complementary ways. You can use the graphical tools provided by your distribution. These tools have a look and feel that depends on the distribution. If you are a novice Linux user on your home system, then use the graphical tool that is provided by your distribution. This will make sure that you do not run into problems.

Another option is to use command line tools like `useradd`, `usermod`, `gpasswd`, `passwd` and others. Server administrators are likely to use these tools, since they are familiar and very similar across many different distributions. This chapter will focus on these command line tools.

A third and rather extremist way is to edit the local configuration files directly using `vi` (or `vim`/`vi`). Do not attempt this as a novice on production systems!

10.2. `/etc/passwd`

The local user database on Linux (and on most Unixes) is `/etc/passwd`.

```
[root@linux ~]# tail /etc/passwd
inge:x:518:524:art dealer:/home/inge:/bin/ksh
ann:x:519:525:flute player:/home/ann:/bin/bash
frederik:x:520:526:rubius poet:/home/frederik:/bin/bash
steven:x:521:527:roman emperor:/home/steven:/bin/bash
pascale:x:522:528:artist:/home/pascale:/bin/ksh
geert:x:524:530:kernel developer:/home/geert:/bin/bash
wim:x:525:531:master damuti:/home/wim:/bin/bash
sandra:x:526:532:radish stresser:/home/sandra:/bin/bash
annelies:x:527:533:sword fighter:/home/annelies:/bin/bash
laura:x:528:534:art dealer:/home/laura:/bin/ksh
```

As you can see, this file contains seven columns separated by a colon. The columns contain the username, an `x`, the user id, the primary group id, a description, the name of the home directory, and the login shell.

More information can be found by typing `man 5 passwd`.

```
[root@linux ~]# man 5 passwd
```

10.3. root

The `root` user also called the `superuser` is the most powerful account on your Linux system. This user can do almost anything, including the creation of other users. The `root` user always has `userid 0` (regardless of the name of the account).

```
[root@linux ~]# head -1 /etc/passwd
root:x:0:0:root:/root:/bin/bash
```

10.4. useradd

You can add users with the `useradd` command. The example below shows how to add a user named `yanina` (last parameter) and at the same time forcing the creation of the home directory (`-m`), setting the name of the home directory (`-d`), and setting a description (`-c`).

```
[root@linux ~]# useradd -m -d /home/yanina -c "yanina wickmayer" yanina
[root@linux ~]# tail -1 /etc/passwd
yanina:x:529:529:yanina wickmayer:/home/yanina:/bin/bash
```

The user named `yanina` received `userid 529` and `primary group id 529`.

10.5. /etc/default/useradd

Both Red Hat Enterprise Linux and Debian/Ubuntu have a file called `/etc/default/useradd` that contains some default user options. Besides using `cat` to display this file, you can also use `useradd -D`.

```
[root@RHEL4 ~]# useradd -D
GROUP=100
HOME=/home
INACTIVE=-1
EXPIRE=
SHELL=/bin/bash
SKEL=/etc/skel
```

10.6. userdel

You can delete the user `yanina` with `userdel`. The `-r` option of `userdel` will also remove the home directory.

```
[root@linux ~]# userdel -r yanina
```


10.7. usermod

You can modify the properties of a user with the `usermod` command. This example uses `usermod` to change the description of the user `harry`.

```
[root@RHEL4 ~]# tail -1 /etc/passwd
harry:x:516:520:harry potter:/home/harry:/bin/bash
[root@RHEL4 ~]# usermod -c 'wizard' harry
[root@RHEL4 ~]# tail -1 /etc/passwd
harry:x:516:520:wizard:/home/harry:/bin/bash
```

10.8. creating home directories

The easiest way to create a home directory is to supply the `-m` option with `useradd` (it is likely set as a default option on Linux).

A less easy way is to create a home directory manually with `mkdir` which also requires setting the owner and the permissions on the directory with `chmod` and `chown` (both commands are discussed in detail in another chapter).

```
[root@linux ~]# mkdir /home/laura
[root@linux ~]# chown laura:laura /home/laura
[root@linux ~]# chmod 700 /home/laura
[root@linux ~]# ls -ld /home/laura/
drwx----- 2 laura laura 4096 Jun 24 15:17 /home/laura/
```

10.9. /etc/skel/

When using `useradd` the `-m` option, the `/etc/skel/` directory is copied to the newly created home directory. The `/etc/skel/` directory contains some (usually hidden) files that contain profile settings and default values for applications. In this way `/etc/skel/` serves as a default home directory and as a default user profile.

```
[root@linux ~]# ls -la /etc/skel/
total 48
drwxr-xr-x  2 root root  4096 Apr  1 00:11 .
drwxr-xr-x 97 root root 12288 Jun 24 15:36 ..
-rw-r--r--  1 root root    24 Jul 12  2006 .bash_logout
-rw-r--r--  1 root root   176 Jul 12  2006 .bash_profile
-rw-r--r--  1 root root   124 Jul 12  2006 .bashrc
```

10.10. deleting home directories

The `-r` option of `userdel` will make sure that the home directory is deleted together with the user account.

```
[root@linux ~]# ls -ld /home/wim/
drwx----- 2 wim wim 4096 Jun 24 15:19 /home/wim/
[root@linux ~]# userdel -r wim
[root@linux ~]# ls -ld /home/wim/
ls: /home/wim/: No such file or directory
```

10.11. login shell

The `/etc/passwd` file specifies the `login shell` for the user. In the screenshot below you can see that user `annelies` will log in with the `/bin/bash` shell, and user `laura` with the `/bin/ksh` shell.

```
[root@linux ~]# tail -2 /etc/passwd
annelies:x:527:533:sword fighter:/home/annelies:/bin/bash
laura:x:528:534:art dealer:/home/laura:/bin/ksh
```

You can use the `usermod` command to change the shell for a user.

```
[root@linux ~]# usermod -s /bin/bash laura
[root@linux ~]# tail -1 /etc/passwd
laura:x:528:534:art dealer:/home/laura:/bin/bash
```

10.12. chsh

Users can change their login shell with the `chsh` command. First, user `harry` obtains a list of available shells (he could also have done a `cat /etc/shells`) and then changes his login shell to the Korn shell (`/bin/ksh`). At the next login, `harry` will default into `ksh` instead of `bash`.

```
[laura@linux ~]$ chsh -l
/bin/sh
/bin/bash
/sbin/nologin
/usr/bin/sh
/usr/bin/bash
/usr/sbin/nologin
/bin/ksh
/bin/tcsh
/bin/csh
[laura@linux ~]$
```

Note that the `-l` option does not exist on Debian and that the above screenshot assumes that `ksh` and `csh` shells are installed.

The screenshot below shows how `laura` can change her default shell (active on next login).

```
[laura@linux ~]$ chsh -s /bin/ksh
Changing shell for laura.
Password:
Shell changed.
```

10.13. practice: user management

1. Create a user account named `serena`, including a home directory and a description (or comment) that reads `Serena Williams`. Do all this in one single command.
2. Create a user named `venus`, including home directory, `bash` shell, a description that reads `Venus Williams` all in one single command.
3. Verify that both users have correct entries in `/etc/passwd`, `/etc/shadow` and `/etc/group`.

4. Verify that their home directory was created.
5. Create a user named `einstime` with `/bin/date` as his default logon shell.
6. What happens when you log on with the `einstime` user ? Can you think of a useful real world example for changing a user's login shell to an application ?
7. Create a file named `welcome.txt` and make sure every new user will see this file in their home directory.
8. Verify this setup by creating (and deleting) a test user account.
9. Change the default login shell for the `serena` user to `/bin/bash`. Verify before and after you make this change.

10.14. solution: user management

1. Create a user account named `serena`, including a home directory and a description (or comment) that reads `Serena Williams`. Do all this in one single command.

```
root@linux:~# useradd -m -c 'Serena Williams' serena
```

2. Create a user named `venus`, including home directory, bash shell, a description that reads `Venus Williams` all in one single command.

```
root@linux:~# useradd -m -c "Venus Williams" -s /bin/bash venus
```

3. Verify that both users have correct entries in `/etc/passwd`, `/etc/shadow` and `/etc/group`.

```
root@linux:~# tail -2 /etc/passwd
serena:x:1008:1010:Serena Williams:/home/serena:/bin/sh
venus:x:1009:1011:Venus Williams:/home/venus:/bin/bash
root@linux:~# tail -2 /etc/shadow
serena:!:16358:0:99999:7:::
venus:!:16358:0:99999:7:::
root@linux:~# tail -2 /etc/group
serena:x:1010:
venus:x:1011:
```

4. Verify that their home directory was created.

```
root@linux:~# ls -lrtl /home | tail -2
drwxr-xr-x 2 serena  serena  4096 Oct 15 10:50 serena
drwxr-xr-x 2 venus   venus   4096 Oct 15 10:59 venus
root@linux:~#
```

5. Create a user named `einstime` with `/bin/date` as his default logon shell.

```
root@linux:~# useradd -s /bin/date einstime
```

Or even better:

```
root@linux:~# useradd -s $(which date) einstime
```

6. What happens when you log on with the `einstime` user ? Can you think of a useful real world example for changing a user's login shell to an application ?

10. user management

```
root@linux:~# su - einstime
Wed Oct 15 11:05:56 UTC 2014    # You get the output of the date command
root@linux:~#
```

It can be useful when users need to access only one application on the server. Just logging in opens the application for them, and closing the application automatically logs them out.

7. Create a file named `welcome.txt` and make sure every new user will see this file in their home directory.

```
root@linux:~# echo Hello > /etc/skel/welcome.txt
```

8. Verify this setup by creating (and deleting) a test user account.

```
root@linux:~# useradd -m test
root@linux:~# ls -l /home/test
total 4
-rw-r--r-- 1 test test 6 Oct 15 11:16 welcome.txt
root@linux:~# userdel -r test
root@linux:~#
```

9. Change the default login shell for the `serena` user to `/bin/bash`. Verify before and after you make this change.

```
root@linux:~# grep serena /etc/passwd
serena:x:1008:1010:Serena Williams:/home/serena:/bin/sh
root@linux:~# usermod -s /bin/bash serena
root@linux:~# grep serena /etc/passwd
serena:x:1008:1010:Serena Williams:/home/serena:/bin/bash
root@linux:~#
```

11. user passwords

(Written by Paul Cobbaut, <https://github.com/paulcobbaut/>, with contributions by: Alex M. Schapelle, <https://github.com/zero-pythagoras/>)

This chapter will tell you more about passwords for local users.

Three methods for setting passwords are explained; using the `passwd` command, using `openssl passwd`, and using the `crypt` function in a C program.

The chapter will also discuss password settings and disabling, suspending or locking accounts.

11.1. passwd

Passwords of users can be set with the `passwd` command. Users will have to provide their old password before twice entering the new one.

```
[tania@linux ~]$ passwd
Changing password for user tania.
Changing password for tania.
(current) UNIX password:
New password:
BAD PASSWORD: The password is shorter than 8 characters
New password:
BAD PASSWORD: The password is a palindrome
New password:
BAD PASSWORD: The password is too similar to the old one
passwd: Have exhausted maximum number of retries for service
```

As you can see, the `passwd` tool will do some basic verification to prevent users from using too simple passwords. The `root` user does not have to follow these rules (there will be a warning though). The `root` user also does not have to provide the old password before entering the new password twice.

```
root@linux:~# passwd tania
Enter new UNIX password:
Retype new UNIX password:
passwd: password updated successfully
```

11.2. shadow file

User passwords are encrypted and kept in `/etc/shadow`. The `/etc/shadow` file is read only and can only be read by root. We will see in the file permissions section how it is possible for users to change their password. For now, you will have to know that users can change their password with the `/usr/bin/passwd` command.

11. user passwords

```
[root@linux ~]# tail -4 /etc/shadow
paul:$6$ikp2Xta5BT.Tml.p$2TZjNn0YNNQKpwLJqoGJbVsZG5/Fti8ovBRd.VzRbiDSl7TEq\
IaSMH.TeBKntS/SjLMruW8qffc0JNORW.BTW1:16338:0:99999:7 :::
tania:$6$8Z/zovxj$9qvoqT8i9KIrmN.k4EQwAF5ryz5yzNwEvYjAa9L5XVXQu.z4DlvpMREH\
eQpQzvRnqFdKkVj17H5ST.c79HDZw0:16356:0:99999:7 :::
laura:$6$glDuTY5e$/NYYWLxfHgZFWeoujaXSMcR.Mz.lG0xtcxFocFVJNb98nbTPhWFXfKWG\
SyYh1WCv6763Wq54.w24Yr3uAZB0m/:16356:0:99999:7 :::
valentina:$6$jRZa6PVI$1uQgqR6En9mZB6mKJ3LXRB4CnFko6LRhbh.v4iqUk9MVreui1lv7\
GxHOUDSKA0N55ZRNhGHa6T2ouFnVno/0o1:16356:0:99999:7 :::
[root@linux ~]#
```

The `/etc/shadow` file contains nine colon separated columns. The nine fields contain (from left to right) the user name, the encrypted password (note that only inge and laura have an encrypted password), the day the password was last changed (day 1 is January 1, 1970), number of days the password must be left unchanged, password expiry day, warning number of days before password expiry, number of days after expiry before disabling the account, and the day the account was disabled (again, since 1970). The last field has no meaning yet.

All the passwords in the screenshot above are hashes of `hunter2`.

11.3. encryption with `passwd`

Passwords are stored in an encrypted format. This encryption is done by the `crypt` function. The easiest (and recommended) way to add a user with a password to the system is to add the user with the `useradd -m user` command, and then set the user's password with `passwd`.

```
[root@RHEL4 ~]# useradd -m xavier
[root@RHEL4 ~]# passwd xavier
Changing password for user xavier.
New UNIX password:
Retype new UNIX password:
passwd: all authentication tokens updated successfully.
[root@RHEL4 ~]#
```

11.4. encryption with `openssl`

Another way to create users with a password is to use the `-p` option of `useradd`, but that option requires an encrypted password. You can generate this encrypted password with the `openssl passwd` command.

The `openssl passwd` command will generate several distinct hashes for the same password, for this it uses a salt.

```
student@linux:~$ openssl passwd hunter2
86jcUNlnGDFpY
student@linux:~$ openssl passwd hunter2
Yj7mD090Anvq6
student@linux:~$ openssl passwd hunter2
YqDcJeGoDbzKA
student@linux:~$
```

This salt can be chosen and is visible as the first two characters of the hash.

```
student@linux:~$ openssl passwd -salt 42 hunter2
42ZrbtP1Ze8G.
student@linux:~$ openssl passwd -salt 42 hunter2
42ZrbtP1Ze8G.
student@linux:~$ openssl passwd -salt 42 hunter2
42ZrbtP1Ze8G.
student@linux:~$
```

This example shows how to create a user with password.

```
root@linux:~# useradd -m -p $(openssl passwd hunter2) mohamed
```

Note that this command puts the password in your command history!

11.5. encryption with crypt

A third option is to create your own C program using the crypt function, and compile this into a command.

```
student@linux:~$ cat MyCrypt.c
#include <stdio.h>
#define __USE_XOPEN
#include <unistd.h>

int main(int argc, char** argv)
{
    if(argc==3)
    {
        printf("%s\n", crypt(argv[1],argv[2]));
    }
    else
    {
        printf("Usage: MyCrypt $password $salt\n" );
    }
    return 0;
}
```

This little program can be compiled with gcc like this.

```
student@linux:~$ gcc MyCrypt.c -o MyCrypt -lcrypt
```

To use it, we need to give two parameters to MyCrypt. The first is the unencrypted password, the second is the salt. The salt is used to perturb the encryption algorithm in one of 4096 different ways. This variation prevents two users with the same password from having the same entry in /etc/shadow.

```
student@linux:~$ ./MyCrypt hunter2 42
42ZrbtP1Ze8G.
student@linux:~$ ./MyCrypt hunter2 33
33d6taYSiEUXI
```

Did you notice that the first two characters of the password are the salt?

The standard output of the crypt function is using the DES algorithm which is old and can be cracked in minutes. A better method is to use md5 passwords which can be recognized by a salt starting with \$1\$.

11. user passwords

```
student@linux:~$ ./MyCrypt hunter2 '$1$42'  
$1$42$7l6Y3xT5282XmZrtD0F9f0  
student@linux:~$ ./MyCrypt hunter2 '$6$42'  
$6$42$0qFFAVnI3gTSYG0yI9TZWX9cpyQzwIop7HwpG1LLEsNBiMr4w60vLX1KDa./UpwXfrFk1i ...
```

The md5 salt can be up to eight characters long. The salt is displayed in `/etc/shadow` between the second and third \$, so never use the password as the salt!

```
student@linux:~$ ./MyCrypt hunter2 '$1$hunter2'  
$1$hunter2$YVxrxDmidq7Xf8Gdt6qM2.
```

11.6. /etc/login.defs

The `/etc/login.defs` file contains some default settings for user passwords like password aging and length settings. (You will also find the numerical limits of user ids and group ids and whether or not a home directory should be created by default).

```
root@linux:~# grep ^PASS /etc/login.defs  
PASS_MAX_DAYS    99999  
PASS_MIN_DAYS    0  
PASS_MIN_LEN     5  
PASS_WARN_AGE    7
```

Debian also has this file.

```
root@linux:~# grep PASS /etc/login.defs  
# PASS_MAX_DAYS    Maximum number of days a password may be used.  
# PASS_MIN_DAYS    Minimum number of days allowed between password changes.  
# PASS_WARN_AGE    Number of days warning given before a password expires.  
PASS_MAX_DAYS    99999  
PASS_MIN_DAYS    0  
PASS_WARN_AGE    7  
#PASS_CHANGE_TRIES  
#PASS_ALWAYS_WARN  
#PASS_MIN_LEN  
#PASS_MAX_LEN  
# NO_PASSWORD_CONSOLE  
root@linux:~#
```

11.7. chage

The `chage` command can be used to set an expiration date for a user account (`-E`), set a minimum (`-m`) and maximum (`-M`) password age, a password expiration date, and set the number of warning days before the password expiration date. Much of this functionality is also available from the `passwd` command. The `-l` option of `chage` will list these settings for a user.

```
root@linux:~# chage -l paul  
Last password change           : Mar 27, 2014  
Password expires               : never  
Password inactive              : never  
Account expires                : never  
Minimum number of days between password change : 0
```



```
Maximum number of days between password change      : 99999
Number of days of warning before password expires   : 7
root@linux:~#
```

11.8. disabling a password

Passwords in `/etc/shadow` cannot begin with an exclamation mark. When the second field in `/etc/passwd` starts with an exclamation mark, then the password can not be used.

Using this feature is often called `locking`, `disabling`, or `suspending` a user account. Besides `vi` (or `vipw`) you can also accomplish this with `usermod`.

The first command in the next screenshot will show the hashed password of `laura` in `/etc/shadow`. The next command disables the password of `laura`, making it impossible for `laura` to authenticate using this password.

```
root@linux:~# grep laura /etc/shadow | cut -c1-70
laura:$6$JYj4JZqp$stwwWACp30tE1R2aZuE87j.nbW.puDkNUYVvk7mCHfCVMa3CoDUJV
root@linux:~# usermod -L laura
```

As you can see below, the password hash is simply preceded with an exclamation mark.

```
root@linux:~# grep laura /etc/shadow | cut -c1-70
laura: !$6$JYj4JZqp$stwwWACp30tE1R2aZuE87j.nbW.puDkNUYVvk7mCHfCVMa3CoDUJ
root@linux:~#
```

The root user (and users with `sudo` rights on `su`) still will be able to `su` into the `laura` account (because the password is not needed here). Also note that `laura` will still be able to login if she has set up passwordless `ssh`!

```
root@linux:~# su - laura
laura@linux:~$
```

You can unlock the account again with `usermod -U`.

```
root@linux:~# usermod -U laura
root@linux:~# grep laura /etc/shadow | cut -c1-70
laura:$6$JYj4JZqp$stwwWACp30tE1R2aZuE87j.nbW.puDkNUYVvk7mCHfCVMa3CoDUJV
```

Watch out for tiny differences in the command line options of `passwd`, `usermod`, and `useradd` on different Linux distributions. Verify the local files when using features like "disabling, suspending, or locking" on user accounts and their passwords.

11.9. editing local files

If you still want to manually edit the `/etc/passwd` or `/etc/shadow`, after knowing these commands for password management, then use `vipw` instead of `vi(m)` directly. The `vipw` tool will do proper locking of the file.

```
[root@linux ~]# vipw /etc/passwd
vipw: the password file is busy (/etc/ptmp present)
```

11.10. practice: user passwords

1. Set the password for serena to hunter2.
2. Also set a password for venus and then lock the venus user account with `usermod`. Verify the locking in `/etc/shadow` before and after you lock it.
3. Use `passwd -d` to disable the serena password. Verify the serena line in `/etc/shadow` before and after disabling.
4. What is the difference between locking a user account and disabling a user account's password like we just did with `usermod -L` and `passwd -d`?
5. Try changing the password of serena to serena as serena.
6. Make sure serena has to change her password in 10 days.
7. Make sure every new user needs to change their password every 10 days.
8. Take a backup as root of `/etc/shadow`. Use `vi` to copy an encrypted hunter2 hash from venus to serena. Can serena now log on with hunter2 as a password ?
9. Why use `vipw` instead of `vi` ? What could be the problem when using `vi` or `vim` ?
10. Use `chsh` to list all shells (only works on RHEL/CentOS/Fedora), and compare to `cat /etc/shells`.
11. Which `useradd` option allows you to name a home directory ?
12. How can you see whether the password of user serena is locked or unlocked ? Give a solution with `grep` and a solution with `passwd`.

11.11. solution: user passwords

1. Set the password for serena to hunter2.

```
root@linux:~# passwd serena
Enter new UNIX password:
Retype new UNIX password:
passwd: password updated successfully
```

2. Also set a password for venus and then lock the venus user account with `usermod`. Verify the locking in `/etc/shadow` before and after you lock it.

```
root@linux:~# passwd venus
Enter new UNIX password:
Retype new UNIX password:
passwd: password updated successfully
root@linux:~# grep venus /etc/shadow | cut -c1-70
venus:$6$gswzXICW$uSnKFV1kFKZmTPaMVS4AvNA/K0270xN0v5LHdV9ed0gTyXrjUeM/
root@linux:~# usermod -L venus
root@linux:~# grep venus /etc/shadow | cut -c1-70
venus:!$6$gswzXICW$uSnKFV1kFKZmTPaMVS4AvNA/K0270xN0v5LHdV9ed0gTyXrjUeM
```

Note that `usermod -L` precedes the password hash with an exclamation mark (!).

3. Use `passwd -d` to disable the serena password. Verify the serena line in `/etc/shadow` before and after disabling.

```

root@linux:~# grep serena /etc/shadow | cut -c1-70
serena:$6$Es/omrPE$F2Ypu8kpLrfKdW0v/UIwA5jrYyBD2nwZ/dt.i/IypRgiPZSdB/B
root@linux:~# passwd -d serena
passwd: password expiry information changed.
root@linux:~# grep serena /etc/shadow
serena::16358:0:99999:7:::
root@linux:~#

```

4. What is the difference between locking a user account and disabling a user account's password like we just did with `usermod -L` and `passwd -d`?

Locking will prevent the user from logging on to the system with his password by putting a `!` in front of the password in `/etc/shadow`.

Disabling with `passwd` will erase the password from `/etc/shadow`.

5. Try changing the password of serena to serena as serena.

log on as serena, then execute: `passwd serena ...` it should fail!

6. Make sure serena has to change her password in 10 days.

```
chage -M 10 serena
```

7. Make sure every new user needs to change their password every 10 days.

```
vi /etc/login.defs (and change PASS_MAX_DAYS to 10)
```

8. Take a backup as root of `/etc/shadow`. Use `vi` to copy an encrypted `hunter2` hash from `venus` to `serena`. Can serena now log on with `hunter2` as a password?

Yes.

9. Why use `vipw` instead of `vi`? What could be the problem when using `vi` or `vim`?

`vipw` will give a warning when someone else is already using that file (with `vipw`).

10. Use `chsh` to list all shells (only works on RHEL/CentOS/Fedora), and compare to `cat /etc/shells`.

```
chsh -l
cat /etc/shells
```

11. Which `useradd` option allows you to name a home directory?

```
-d
```

12. How can you see whether the password of user `serena` is locked or unlocked? Give a solution with `grep` and a solution with `passwd`.

```
grep serena /etc/shadow
```

```
passwd -S serena
```


12. User profiles

(Written by Paul Cobbaut, <https://github.com/paulcobbaut/>, with contributions by: Alex M. Schapelle, <https://github.com/zero-pythagoras/>)

Logged on users have a number of preset (and customized) aliases, variables, and functions, but where do they come from? The shell uses a number of startup files that are executed (or rather sourced) whenever the shell is invoked. What follows is an overview of startup scripts.

12.1. system profile

Both the bash and the ksh shell will verify the existence of `/etc/profile` and source it if it exists.

When reading this script, you will notice (both on Debian and on Red Hat Enterprise Linux) that it builds the `PATH` environment variable (among others). The script might also change the `PS1` variable, set the `HOSTNAME` and execute even more scripts like `/etc/inputrc`

This screenshot uses `grep` to show `PATH` manipulation in `/etc/profile` on Debian.

```
root@linux:~# grep PATH /etc/profile
  PATH="/usr/local/sbin:/usr/local/bin:/usr/sbin:/usr/bin:/sbin:/bin"
  PATH="/usr/local/bin:/usr/bin:/bin:/usr/local/games:/usr/games"
export PATH
root@linux:~#
```

This screenshot uses `grep` to show `PATH` manipulation in `/etc/profile` on RHEL7/CentOS7.

```
[root@linux ~]# grep PATH /etc/profile
  case ":{PATH}:" in
    PATH=$PATH:$1
    PATH=$1:$PATH
export PATH USER LOGNAME MAIL HOSTNAME HISTSIZE HISTCONTROL
[root@linux ~]#
```

The `root` user can use this script to set aliases, functions, and variables for every user on the system.

12.2. `~/.bash_profile`

When this file exists in the home directory, then bash will source it. On Debian Linux 5/6/7 this file does not exist by default.

RHEL7/CentOS7 uses a small `~/.bash_profile` where it checks for the existence of `~/.bashrc` and then sources it. It also adds `$HOME/bin` to the `$PATH` variable.

12. User profiles

```
[root@linux ~]# cat /home/paul/.bash_profile
# .bash_profile

# Get the aliases and functions
if [ -f ~/.bashrc ]; then
    . ~/.bashrc
fi

# User specific environment and startup programs

PATH=$PATH:$HOME/.local/bin:$HOME/bin

export PATH
[root@linux ~]#
```

12.3. ~/.bash_login

When `.bash_profile` does not exist, then `bash` will check for `~/.bash_login` and source it.

Neither Debian nor Red Hat have this file by default.

12.4. ~/.profile

When neither `~/.bash_profile` and `~/.bash_login` exist, then `bash` will verify the existence of `~/.profile` and execute it. This file does not exist by default on Red Hat.

On Debian this script can execute `~/.bashrc` and will add `$HOME/bin` to the `$PATH` variable.

```
root@linux:~# tail -11 /home/paul/.profile
if [ -n "$BASH_VERSION" ]; then
    # include .bashrc if it exists
    if [ -f "$HOME/.bashrc" ]; then
        . "$HOME/.bashrc"
    fi
fi

# set PATH so it includes user's private bin if it exists
if [ -d "$HOME/bin" ] ; then
    PATH="$HOME/bin:$PATH"
fi
```

RHEL/CentOS does not have this file by default.

12.5. ~/.bashrc

The `~/.bashrc` script is often sourced by other scripts. Let us take a look at what it does by default.

Red Hat uses a very simple `~/.bashrc`, checking for `/etc/bashrc` and sourcing it. It also leaves room for custom aliases and functions.

```
[root@linux ~]# cat /home/paul/.bashrc
# .bashrc

# Source global definitions
if [ -f /etc/bashrc ]; then
    . /etc/bashrc
fi

# Uncomment the following line if you don't like systemctl's auto-
# paging feature:
# export SYSTEMD_PAGER=

# User specific aliases and functions
```

On Debian this script is quite a bit longer and configures \$PS1, some history variables and a number of active and inactive aliases.

```
root@linux:~# wc -l /home/paul/.bashrc
110 /home/paul/.bashrc
```

12.6. ~/.bash_logout

When exiting bash, it can execute ~/.bash_logout.

Debian use this opportunity to clear the console screen.

```
serena@linux:~$ cat .bash_logout
# ~/.bash_logout: executed by bash(1) when login shell exits.

# when leaving the console clear the screen to increase privacy

if [ "$SHLVL" = 1 ]; then
    [ -x /usr/bin/clear_console ] && /usr/bin/clear_console -q
fi
```

Red Hat Enterprise Linux 5 will simply call the /usr/bin/clear command in this script.

```
[serena@linux ~]$ cat .bash_logout
# ~/.bash_logout

/usr/bin/clear
```

Red Hat Enterprise Linux 6 and 7 create this file, but leave it empty (except for a comment).

```
student@linux:~$ cat .bash_logout
# ~/.bash_logout
```

12.7. Debian overview

Below is a table overview of when Debian is running any of these bash startup scripts.

12. User profiles

Table 12.1.: Debian User Environment

script	su	su -	ssh	gdm
~/bashrc	no	yes	yes	yes
~/profile	no	yes	yes	yes
/etc/profile	no	yes	yes	yes
/etc/bash.bashrc	yes	no	no	yes

12.8. RHEL5 overview

Below is a table overview of when Red Hat Enterprise Linux 5 is running any of these bash startup scripts.

Table 12.2.: Red Hat User Environment

script	su	su -	ssh	gdm
~/bashrc	yes	yes	yes	yes
~/bash_profile	no	yes	yes	yes
/etc/profile	no	yes	yes	yes
/etc/bashrc	yes	yes	yes	yes

12.9. practice: user profiles

1. Make a list of all the profile files on your system.
2. Read the contents of each of these, often they source extra scripts.
3. Put a unique variable, alias and function in each of those files.
4. Try several different ways to obtain a shell (su, su -, ssh, tmux, gnome-terminal, Ctrl-alt-F1, ...) and verify which of your custom variables, aliases and function are present in your environment.
5. Do you also know the order in which they are executed?
6. When an application depends on a setting in \$HOME/profile, does it matter whether \$HOME/bash_profile exists or not ?

12.10. solution: user profiles

1. Make a list of all the profile files on your system.

```
ls -a ~ ; ls -l /etc/pro* /etc/bash*
```

2. Read the contents of each of these, often they source extra scripts.
3. Put a unique variable, alias and function in each of those files.
4. Try several different ways to obtain a shell (su, su -, ssh, tmux, gnome-terminal, Ctrl-alt-F1, ...) and verify which of your custom variables, aliases and function are present in your environment.
5. Do you also know the order in which they are executed?

same name aliases, functions and variables will overwrite each other

6. When an application depends on a setting in `$HOME/.profile`, does it matter whether `$HOME/.bash_profile` exists or not?

Yes it does matter. (man bash /INVOCATION)

13. groups

(Written by Paul Cobbaut, <https://github.com/paulcobbaut/>, with contributions by: Alex M. Schapelle, <https://github.com/zero-pythagoras/>)

Users can be listed in groups. Groups allow you to set permissions on the group level instead of having to set permissions for every individual user.

Every Unix or Linux distribution will have a graphical tool to manage groups. Novice users are advised to use this graphical tool. More experienced users can use command line tools to manage users, but be careful: Some distributions do not allow the mixed use of GUI and CLI tools to manage groups (YaST in Novell Suse). Senior administrators can edit the relevant files directly with `vi` or `vigr`.

13.1. groupadd

Groups can be created with the `groupadd` command. The example below shows the creation of five (empty) groups.

```
root@linux:~# groupadd tennis
root@linux:~# groupadd football
root@linux:~# groupadd snooker
root@linux:~# groupadd formula1
root@linux:~# groupadd salsa
```

13.2. group file

Users can be a member of several groups. Group membership is defined by the `/etc/group` file.

```
root@linux:~# tail -5 /etc/group
tennis:x:1006:
football:x:1007:
snooker:x:1008:
formula1:x:1009:
salsa:x:1010:
root@linux:~#
```

The first field is the group's name. The second field is the group's (encrypted) password (can be empty). The third field is the group identification or `GID`. The fourth field is the list of members, these groups have no members.

13.3. groups

A user can type the `groups` command to see a list of groups where the user belongs to.

```
[harry@linux ~]$ groups
harry sports
[harry@linux ~]$
```

13.4. usermod

Group membership can be modified with the `useradd` or `usermod` command.

```
root@linux:~# usermod -a -G tennis inge
root@linux:~# usermod -a -G tennis katrien
root@linux:~# usermod -a -G salsa katrien
root@linux:~# usermod -a -G snooker sandra
root@linux:~# usermod -a -G formula1 annelies
root@linux:~# tail -5 /etc/group
tennis:x:1006:inge,katrien
football:x:1007:
snooker:x:1008:sandra
formula1:x:1009:annelies
salsa:x:1010:katrien
root@linux:~#
```

Be careful when using `usermod` to add users to groups. By default, the `usermod` command will remove the user from every group of which he is a member if the group is not listed in the command! Using the `-a` (append) switch prevents this behaviour.

13.5. groupmod

You can change the group name with the `groupmod` command.

```
root@linux:~# groupmod -n darts snooker
root@linux:~# tail -5 /etc/group
tennis:x:1006:inge,katrien
football:x:1007:
formula1:x:1009:annelies
salsa:x:1010:katrien
darts:x:1008:sandra
```

13.6. groupdel

You can permanently remove a group with the `groupdel` command.

```
root@linux:~# groupdel tennis
root@linux:~#
```

13.7. gpasswd

You can delegate control of group membership to another user with the `gpasswd` command. In the example below we delegate permissions to add and remove group members to `serena` for the `sports` group. Then we `su` to `serena` and add `harry` to the `sports` group.

```
[root@linux ~]# gpasswd -A serena sports
[root@linux ~]# su - serena
[serena@linux ~]$ id harry
uid=516(harry) gid=520(harry) groups=520(harry)
[serena@linux ~]$ gpasswd -a harry sports
Adding user harry to group sports
[serena@linux ~]$ id harry
uid=516(harry) gid=520(harry) groups=520(harry),522(sports)
[serena@linux ~]$ tail -1 /etc/group
sports:x:522:serena,venus,harry
[serena@linux ~]$
```

Group administrators do not have to be a member of the group. They can remove themselves from a group, but this does not influence their ability to add or remove members.

```
[serena@linux ~]$ gpasswd -d serena sports
Removing user serena from group sports
[serena@linux ~]$ exit
```

Information about group administrators is kept in the `/etc/gshadow` file.

```
[root@linux ~]# tail -1 /etc/gshadow
sports:!:serena:venus,harry
[root@linux ~]#
```

To remove all group administrators from a group, use the `gpasswd` command to set an empty administrators list.

```
[root@linux ~]# gpasswd -A "" sports
```

13.8. newgrp

You can start a child shell with a new temporary `primary` group using the `newgrp` command.

```
root@linux:~# mkdir prigroup
root@linux:~# cd prigroup/
root@linux:~/prigroup# touch standard.txt
root@linux:~/prigroup# ls -l
total 0
-rw-r--r--. 1 root root 0 Apr 13 17:49 standard.txt
root@linux:~/prigroup# echo $SHLVL
1
root@linux:~/prigroup# newgrp tennis
root@linux:~/prigroup# echo $SHLVL
2
root@linux:~/prigroup# touch newgrp.txt
root@linux:~/prigroup# ls -l
```

13. groups

```
total 0
-rw-r--r--. 1 root tennis 0 Apr 13 17:49 newgrp.txt
-rw-r--r--. 1 root root    0 Apr 13 17:49 standard.txt
root@linux:~/prigroup# exit
exit
root@linux:~/prigroup#
```

13.9. vigr

Similar to vipw, the `vigr` command can be used to manually edit the `/etc/group` file, since it will do proper locking of the file. Only experienced senior administrators should use `vi` or `vigr` to manage groups.

13.10. practice: groups

1. Create the groups tennis, football and sports.
2. In one command, make venus a member of tennis and sports.
3. Rename the football group to foot.
4. Use `vi` to add serena to the tennis group.
5. Use the `id` command to verify that serena is a member of tennis.
6. Make someone responsible for managing group membership of foot and sports. Test that it works.

13.11. solution: groups

1. Create the groups tennis, football and sports.

```
groupadd tennis ; groupadd football ; groupadd sports
```

2. In one command, make venus a member of tennis and sports.

```
usermod -a -G tennis,sports venus
```

3. Rename the football group to foot.

```
groupmod -n foot football
```

4. Use `vi` to add serena to the tennis group.

```
vi /etc/group
```

5. Use the `id` command to verify that serena is a member of tennis.

```
id (and after logoff logon serena should be member)
```

6. Make someone responsible for managing group membership of foot and sports. Test that it works.

`gpsswd -A` (to make manager)

`gpsswd -a` (to add member)

Part V.

Webserver; scripting 102

14. apache web server

(Written by Paul Cobbaut, <https://github.com/paulcobbaut/>, with contributions by: Hans Roes, <https://github.com/Blokker-1999/>, Alex M. Schapelle, <https://github.com/zero-pythagoras/>)

In this chapter we learn how to setup a web server with the apache software.

According to NetCraft (http://news.netcraft.com/archives/web_server_survey.html) about seventy percent of all web servers are running on Apache. The name is derived from a patchy web server, because of all the patches people wrote for the NCSA httpd server.

Later chapters will expand this web server into a LAMP stack (Linux, Apache, MySQL, Perl/PHP/Python).

14.1. introduction to apache

14.1.1. installing on Debian

This screenshot shows that there is no apache server installed, nor does the `/var/www` directory exist.

```
root@linux:~# ls -l /var/www
ls: cannot access /var/www: No such file or directory
root@linux:~# dpkg -l | grep apache
```

To install apache on Debian:

```
root@linux:~# aptitude install apache2
The following NEW packages will be installed:
  apache2 apache2-mpm-worker{a} apache2-utils{a} apache2.2-bin{a} apache2.2-
com\
mon{a} libapr1{a} libaprutil1{a} libaprutil1-dbd-sqlite3{a} libaprutil1-
ldap{a}\
ssl-cert{a}
0 packages upgraded, 10 newly installed, 0 to remove and 0 not upgraded.
Need to get 1,487 kB of archives. After unpacking 5,673 kB will be used.
Do you want to continue? [Y/n/?]
```

After installation, the same two commands as above will yield a different result:

```
root@linux:~# ls -l /var/www
total 4
-rw-r--r-- 1 root root 177 Apr 29 11:55 index.html
root@linux:~# dpkg -l | grep apache | tr -s ' '
ii apache2 2.2.22-13+deb7u1 amd64 Apache HTTP Server metapackage
ii apache2-mpm-worker 2.2.22-13+deb7u1 amd64 Apache HTTP Server - high speed th\
readed model
ii apache2-utils 2.2.22-13+deb7u1 amd64 utility programs for webservers
ii apache2.2-bin 2.2.22-13+deb7u1 amd64 Apache HTTP Server common binary files
ii apache2.2-common 2.2.22-13+deb7u1 amd64 Apache HTTP Server common files
```

14.1.2. installing on RHEL/CentOS

Note that Red Hat derived distributions use `httpd` as package and process name instead of `apache`.

To verify whether `apache` is installed in CentOS/RHEL:

```
[root@linux ~]# rpm -q httpd
package httpd is not installed
[root@linux ~]# ls -l /var/www
ls: cannot access /var/www: No such file or directory
```

To install `apache` on CentOS:

```
[root@linux ~]# yum install httpd
```

After running the `yum install httpd` command, the Centos 6.5 server has `apache` installed and the `/var/www` directory exists.

```
[root@linux ~]# rpm -q httpd
httpd-2.2.15-30.el6.centos.x86_64
[root@linux ~]# ls -l /var/www
total 16
drwxr-xr-x. 2 root root 4096 Apr  3 23:57 cgi-bin
drwxr-xr-x. 3 root root 4096 May  6 13:08 error
drwxr-xr-x. 2 root root 4096 Apr  3 23:57 html
drwxr-xr-x. 3 root root 4096 May  6 13:08 icons
[root@linux ~]#
```

14.1.3. running apache on Debian

This is how you start `apache2` on Debian.

```
root@linux:~# service apache2 status
Apache2 is NOT running.
root@linux:~# service apache2 start
Starting web server: apache2apache2: Could not reliably determine the server's \
fully qualified domain name, using 127.0.1.1 for ServerName
.
```

To verify, run the `service apache2 status` command again or use `ps`.

```
root@linux:~# service apache2 status
Apache2 is running (pid 3680).
root@linux:~# ps -C apache2
  PID TTY          TIME CMD
 3680 ?            00:00:00 apache2
 3683 ?            00:00:00 apache2
 3684 ?            00:00:00 apache2
 3685 ?            00:00:00 apache2
root@linux:~#
```

Or use `wget` and `file` to verify that your web server serves an `html` document.

```

root@linux:~# wget 127.0.0.1
--2014-05-06 13:27:02-- http://127.0.0.1/
Connecting to 127.0.0.1:80... connected.
HTTP request sent, awaiting response... 200 OK
Length: 177 [text/html]
Saving to: `index.html'

100%[=====>] 177          --.-
K/s   in 0s

2014-05-06 13:27:02 (15.8 MB/s) - `index.html' saved [177/177]

root@linux:~# file index.html
index.html: HTML document, ASCII text
root@linux:~#

```

Or verify that apache is running by opening a web browser, and browse to the ip-address of your server. An Apache test page should be shown.

You can do the following to quickly avoid the 'could not reliably determine the fqdn' message when restarting apache.

```

root@linux:~# echo ServerName debian10 >> /etc/apache2/apache2.conf
root@linux:~# service apache2 restart
Restarting web server: apache2 ... waiting .
root@linux:~#

```

14.1.4. running apache on CentOS

Starting the httpd on RHEL/CentOS is done with the service command.

```

[root@linux ~]# service httpd status
httpd is stopped
[root@linux ~]# service httpd start
Starting httpd: httpd: Could not reliably determine the server's fully qualifie\
d domain name, using 127.0.0.1 for ServerName
[ OK ]
[root@linux ~]#

```

To verify that apache is running, use ps or issue the service httpd status command again.

```

[root@linux ~]# service httpd status
httpd (pid 2410) is running...
[root@linux ~]# ps -C httpd
  PID TTY          TIME CMD
 2410 ?            00:00:00 httpd
 2412 ?            00:00:00 httpd
 2413 ?            00:00:00 httpd
 2414 ?            00:00:00 httpd
 2415 ?            00:00:00 httpd
 2416 ?            00:00:00 httpd
 2417 ?            00:00:00 httpd
 2418 ?            00:00:00 httpd
 2419 ?            00:00:00 httpd
[root@linux ~]#

```

14. apache web server

To prevent the 'Could not reliably determine the fqdn' message, issue the following command.

```
[root@linux ~]# echo ServerName Centos65 >> /etc/httpd/conf/httpd.conf
[root@linux ~]# service httpd restart
Stopping httpd:          [ OK ]
Starting httpd:         [ OK ]
[root@linux ~]#
```

14.1.5. index file on CentOS

CentOS does not provide a standard index.html or index.php file. A simple wget gives an error.

```
[root@linux ~]# wget 127.0.0.1
--2014-05-06 15:10:22-- http://127.0.0.1/
Connecting to 127.0.0.1:80 ... connected.
HTTP request sent, awaiting response ... 403 Forbidden
2014-05-06 15:10:22 ERROR 403: Forbidden.
```

Instead when visiting the ip-address of your server in a web browser you get a noindex.html page. You can verify this using wget.

```
[root@linux ~]# wget http://127.0.0.1/error/noindex.html
--2014-05-06 15:16:05-- http://127.0.0.1/error/noindex.html
Connecting to 127.0.0.1:80 ... connected.
HTTP request sent, awaiting response ... 200 OK
Length: 5039 (4.9K) [text/html]
Saving to: "noindex.html"

100%[=====] 5,039      --.-K/s  in 0s

2014-05-06 15:16:05 (289 MB/s) - "noindex.html" saved [5039/5039]

[root@linux ~]# file noindex.html
noindex.html: HTML document text
[root@linux ~]#
```

Any custom index.html file in /var/www/html will immediately serve as an index for this web server.

```
[root@linux ~]# echo 'Welcome to my website' > /var/www/html/index.html
[root@linux ~]# wget http://127.0.0.1
--2014-05-06 15:19:16-- http://127.0.0.1/
Connecting to 127.0.0.1:80 ... connected.
HTTP request sent, awaiting response ... 200 OK
Length: 22 [text/html]
Saving to: "index.html"

100%[=====] 22          --.-K/s  in 0s

2014-05-06 15:19:16 (1.95 MB/s) - "index.html" saved [22/22]

[root@linux ~]# cat index.html
Welcome to my website
```

14.1.6. default website

Changing the default website of a freshly installed apache web server is easy. All you need to do is create (or change) an index.html file in the DocumentRoot directory.

To locate the DocumentRoot directory on Debian:

```
root@linux:~# grep DocumentRoot /etc/apache2/sites-available/default
DocumentRoot /var/www
```

This means that /var/www/index.html is the default web site.

```
root@linux:~# cat /var/www/index.html
<html><body><h1>It works!</h1>
<p>This is the default web page for this server.</p>
<p>The web server software is running but no content has been added, yet.</p>
</body></html>
root@linux:~#
```

This screenshot shows how to locate the DocumentRoot directory on RHEL/CentOS.

```
[root@linux ~]# grep ^DocumentRoot /etc/httpd/conf/httpd.conf
DocumentRoot "/var/www/html"
```

RHEL/CentOS have no default web page (only the noindex.html error page mentioned before). But an index.html file created in /var/www/html/ will automatically be used as default page.

```
[root@linux ~]# echo '<html><head><title>Default website</title></head><body>\
><p>A new web page</p></body></html>' > /var/www/html/index.html
[root@linux ~]# cat /var/www/html/index.html
<html><head><title>Default website</title></head><body><p>A new web page</p></b\
ody></html>
[root@linux ~]#
```

14.1.7. apache configuration

There are many similarities, but also a couple of differences when configuring apache on Debian or on CentOS. Both Linux families will get their own chapters with examples.

All configuration on RHEL/CentOS is done in /etc/httpd.

```
[root@linux ~]# ls -l /etc/httpd/
total 8
drwxr-xr-x. 2 root root 4096 May  6 13:08 conf
drwxr-xr-x. 2 root root 4096 May  6 13:08 conf.d
lrwxrwxrwx. 1 root root   19 May  6 13:08 logs -> ../.. /var/log/httpd
lrwxrwxrwx. 1 root root  29 May  6 13:08 modules -> ../.. /usr/lib64/httpd/modu\
les
lrwxrwxrwx. 1 root root   19 May  6 13:08 run -> ../.. /var/run/httpd
[root@linux ~]#
```

Debian (and ubuntu/mint/...) use /etc/apache2.

14. apache web server

```
root@linux:~# ls -l /etc/apache2/
total 72
-rw-r--r-- 1 root root  9659 May  6 14:23 apache2.conf
drwxr-xr-x 2 root root  4096 May  6 13:19 conf.d
-rw-r--r-- 1 root root  1465 Jan 31 18:35 envvars
-rw-r--r-- 1 root root 31063 Jul 20  2013 magic
drwxr-xr-x 2 root root  4096 May  6 13:19 mods-available
drwxr-xr-x 2 root root  4096 May  6 13:19 mods-enabled
-rw-r--r-- 1 root root   750 Jan 26 12:13 ports.conf
drwxr-xr-x 2 root root  4096 May  6 13:19 sites-available
drwxr-xr-x 2 root root  4096 May  6 13:19 sites-enabled
root@linux:~#
```

14.2. port virtual hosts on Debian

14.2.1. default virtual host

Debian has a virtualhost configuration file for its default website in `/etc/apache2/sites-available/default`.

```
root@linux:~# head -2 /etc/apache2/sites-available/default
<VirtualHost *:80>
    ServerAdmin webmaster@localhost
```

14.2.2. three extra virtual hosts

In this scenario we create three additional websites for three customers that share a clubhouse and want to jointly hire you. They are a model train club named Choo Choo, a chess club named Chess Club 42 and a hackerspace named hunter2.

One way to put three websites on one web server, is to put each website on a different port. This screenshot shows three newly created virtual hosts, one for each customer.

```
root@linux:~# vi /etc/apache2/sites-available/choochoo
root@linux:~# cat /etc/apache2/sites-available/choochoo
<VirtualHost *:7000>
    ServerAdmin webmaster@localhost
    DocumentRoot /var/www/choochoo
</VirtualHost>
root@linux:~# vi /etc/apache2/sites-available/chessclub42
root@linux:~# cat /etc/apache2/sites-available/chessclub42
<VirtualHost *:8000>
    ServerAdmin webmaster@localhost
    DocumentRoot /var/www/chessclub42
</VirtualHost>
root@linux:~# vi /etc/apache2/sites-available/hunter2
root@linux:~# cat /etc/apache2/sites-available/hunter2
<VirtualHost *:9000>
    ServerAdmin webmaster@localhost
    DocumentRoot /var/www/hunter2
</VirtualHost>
```

Notice the different port numbers 7000, 8000 and 9000. Notice also that we specified a unique DocumentRoot for each website.

Are you using Ubuntu or Mint, then these configfiles need to end in `.conf`.

14.2.3. three extra ports

We need to enable these three ports on apache in the `ports.conf` file. Open this file with `vi` and add three lines to listen on three extra ports.

```
root@linux:~# vi /etc/apache2/ports.conf
```

Verify with `grep` that the `Listen` directives are added correctly.

```
root@linux:~# grep ^Listen /etc/apache2/ports.conf
Listen 80
Listen 7000
Listen 8000
Listen 9000
```

14.2.4. three extra websites

Next we need to create three `DocumentRoot` directories.

```
root@linux:~# mkdir /var/www/choochoo
root@linux:~# mkdir /var/www/chessclub42
root@linux:~# mkdir /var/www/hunter2
```

And we have to put some really simple website in those directories.

```
root@linux:~# echo 'Choo Choo model train Choo Choo' > /var/www/choochoo/index.html
root@linux:~# echo 'Welcome to chess club 42' > /var/www/chessclub42/index.html
root@linux:~# echo 'HaCkInG iS fUn At HuNtEr2' > /var/www/hunter2/index.html
```

14.2.5. enabling extra websites

The last step is to enable the websites with the `a2ensite` command. This command will create links in `sites-enabled`.

The links are not there yet...

```
root@linux:~# cd /etc/apache2/
root@linux:/etc/apache2# ls sites-available/
chessclub42 choochoo default default-ssl hunter2
root@linux:/etc/apache2# ls sites-enabled/
000-default
```

So we run the `a2ensite` command for all websites.

```
root@linux:/etc/apache2# a2ensite choochoo
Enabling site choochoo.
To activate the new configuration, you need to run:
  service apache2 reload
root@linux:/etc/apache2# a2ensite chessclub42
Enabling site chessclub42.
To activate the new configuration, you need to run:
  service apache2 reload
```

14. apache web server

```
root@linux:/etc/apache2# a2ensite hunter2
Enabling site hunter2.
To activate the new configuration, you need to run:
  service apache2 reload
```

The links are created, so we can tell apache.

```
root@linux:/etc/apache2# ls sites-enabled/
000-default chessclub42 choochoo hunter2
root@linux:/etc/apache2# service apache2 reload
Reloading web server config: apache2.
root@linux:/etc/apache2#
```

14.2.6. testing the three websites

Testing the model train club named Choo Choo on port 7000.

```
root@linux:/etc/apache2# wget 127.0.0.1:7000
--2014-05-06 21:16:03-- http://127.0.0.1:7000/
Connecting to 127.0.0.1:7000 ... connected.
HTTP request sent, awaiting response ... 200 OK
Length: 32 [text/html]
Saving to: `index.html'

100%[=====>] 32          --.-K/s   in 0s

2014-05-06 21:16:03 (2.92 MB/s) - `index.html' saved [32/32]

root@linux:/etc/apache2# cat index.html
Choo Choo model train Choo Choo
```

Testing the chess club named Chess Club 42 on port 8000.

```
root@linux:/etc/apache2# wget 127.0.0.1:8000
--2014-05-06 21:16:20-- http://127.0.0.1:8000/
Connecting to 127.0.0.1:8000 ... connected.
HTTP request sent, awaiting response ... 200 OK
Length: 25 [text/html]
Saving to: `index.html.1'

100%[=====>] 25          --.-K/s   in 0s

2014-05-06 21:16:20 (2.16 MB/s) - `index.html.1' saved [25/25]

root@linux:/etc/apache2# cat index.html.1
Welcome to chess club 42
```

Testing the hacker club named hunter2 on port 9000.

```
root@linux:/etc/apache2# wget 127.0.0.1:9000
--2014-05-06 21:16:30-- http://127.0.0.1:9000/
Connecting to 127.0.0.1:9000 ... connected.
HTTP request sent, awaiting response ... 200 OK
Length: 26 [text/html]
Saving to: `index.html.2'
```

```
100%[=====] 26      --.-K/s   in 0s
2014-05-06 21:16:30 (2.01 MB/s) - `index.html.2' saved [26/26]

root@linux:/etc/apache2# cat index.html.2
HaCkInG iS fUn At HuNtEr2
```

Cleaning up the temporary files.

```
root@linux:/etc/apache2# rm index.html index.html.1 index.html.2
```

Try testing from another computer using the ip-address of your server.

14.3. named virtual hosts on Debian

14.3.1. named virtual hosts

The chess club and the model train club find the port numbers too hard to remember. They would prefer to have their website accessible by name.

We continue work on the same server that has three websites on three ports. We need to make sure those websites are accessible using the names `choochoo.local`, `chessclub42.local` and `hunter2.local`.

We start by creating three new virtualhosts.

```
root@linux:/etc/apache2/sites-available# vi choochoo.local
root@linux:/etc/apache2/sites-available# vi chessclub42.local
root@linux:/etc/apache2/sites-available# vi hunter2.local
root@linux:/etc/apache2/sites-available# cat choochoo.local
<VirtualHost *:80>
    ServerAdmin webmaster@localhost
    ServerName choochoo.local
    DocumentRoot /var/www/choochoo
</VirtualHost>
root@linux:/etc/apache2/sites-available# cat chessclub42.local
<VirtualHost *:80>
    ServerAdmin webmaster@localhost
    ServerName chessclub42.local
    DocumentRoot /var/www/chessclub42
</VirtualHost>
root@linux:/etc/apache2/sites-available# cat hunter2.local
<VirtualHost *:80>
    ServerAdmin webmaster@localhost
    ServerName hunter2.local
    DocumentRoot /var/www/hunter2
</VirtualHost>
root@linux:/etc/apache2/sites-available#
```

Notice that they all listen on port `80` and have an extra `ServerName` directive.

14.3.2. name resolution

We need some way to resolve names. This can be done with DNS, which is discussed in another chapter. For this demo it is also possible to quickly add the three names to the /etc/hosts file.

```
root@linux:/etc/apache2/sites-available# grep ^192 /etc/hosts
192.168.42.50 choochoo.local
192.168.42.50 chessclub42.local
192.168.42.50 hunter2.local
```

Note that you may have another ip address...

14.3.3. enabling virtual hosts

Next we enable them with a2ensite.

```
root@linux:/etc/apache2/sites-available# a2ensite choochoo.local
Enabling site choochoo.local.
To activate the new configuration, you need to run:
  service apache2 reload
root@linux:/etc/apache2/sites-available# a2ensite chessclub42.local
Enabling site chessclub42.local.
To activate the new configuration, you need to run:
  service apache2 reload
root@linux:/etc/apache2/sites-available# a2ensite hunter2.local
Enabling site hunter2.local.
To activate the new configuration, you need to run:
  service apache2 reload
```

14.3.4. reload and verify

After a service apache2 reload the websites should be available by name.

```
root@linux:/etc/apache2/sites-available# service apache2 reload
Reloading web server config: apache2.
root@linux:/etc/apache2/sites-available# wget chessclub42.local
--2014-05-06 21:37:13-- http://chessclub42.local/
Resolving chessclub42.local (chessclub42.local) ... 192.168.42.50
Connecting to chessclub42.local (chessclub42.local)|192.168.42.50|:80 ... connecte\
cted.
HTTP request sent, awaiting response ... 200 OK
Length: 25 [text/html]
Saving to: `index.html'

100%[=====] 25      --.-K/s  in 0s

2014-05-06 21:37:13 (2.06 MB/s) - `index.html' saved [25/25]

root@linux:/etc/apache2/sites-available# cat index.html
Welcome to chess club 42
```

14.4. password protected website on Debian

You can secure files and directories in your website with a `.htaccess` file that refers to a `.htpasswd` file. The `htpasswd` command can create a `.htpasswd` file that contains a userid and an (encrypted) password.

This screenshot creates a user and password for the hacker named `cliff` and uses the `-c` flag to create the `.htpasswd` file.

```
root@linux:~# htpasswd -c /var/www/.htpasswd cliff
New password:
Re-type new password:
Adding password for user cliff
root@linux:~# cat /var/www/.htpasswd
cliff:$apr1$vuJll0KL$./SZ4w9q0swhX93pQ0PVp.
```

Hacker `rob` also wants access, this screenshot shows how to add a second user and password to `.htpasswd`.

```
root@linux:~# htpasswd /var/www/.htpasswd rob
New password:
Re-type new password:
Adding password for user rob
root@linux:~# cat /var/www/.htpasswd
cliff:$apr1$vuJll0KL$./SZ4w9q0swhX93pQ0PVp.
rob:$apr1$HNln1FFt$nRlpF0H.IW11/1DRq4lQo0
```

Both Cliff and Rob chose the same password (`hunter2`), but that is not visible in the `.htpasswd` file because of the different salts.

Next we need to create a `.htaccess` file in the DocumentRoot of the website we want to protect. This screenshot shows an example.

```
root@linux:~# cd /var/www/hunter2/
root@linux:/var/www/hunter2# cat .htaccess
AuthUserFile /var/www/.htpasswd
AuthName "Members only!"
AuthType Basic
require valid-user
```

Note that we are protecting the website on port `9000` that we created earlier.

And because we put the website for the Hackerspace named `hunter2` in a subdirectory of the default website, we will need to adjust the `AllowOverride` parameter in `/etc/apache2/sites-available/default` as this screenshot shows (with line numbers on `debian10`, your may vary).

```
9      <Directory /var/www/>
10          Options Indexes FollowSymLinks MultiViews
11          AllowOverride Authconfig
12          Order allow,deny
13          allow from all
14      </Directory
```

Now restart the `apache2` server and test that it works!

14.5. port virtual hosts on CentOS

14.5.1. default virtual host

Unlike Debian, CentOS has no virtualHost configuration file for its default website. Instead the default configuration will throw a standard error page when no index file can be found in the default location (/var/www/html).

14.5.2. three extra virtual hosts

In this scenario we create three additional websites for three customers that share a clubhouse and want to jointly hire you. They are a model train club named Choo Choo, a chess club named Chess Club 42 and a hackerspace named hunter2.

One way to put three websites on one web server, is to put each website on a different port. This screenshot shows three newly created virtual hosts, one for each customer.

```
[root@CentOS65 ~]# vi /etc/httpd/conf.d/choochoo.conf
[root@CentOS65 ~]# cat /etc/httpd/conf.d/choochoo.conf
<VirtualHost *:7000>
    ServerAdmin webmaster@localhost
    DocumentRoot /var/www/html/choochoo
</VirtualHost>
[root@CentOS65 ~]# vi /etc/httpd/conf.d/chessclub42.conf
[root@CentOS65 ~]# cat /etc/httpd/conf.d/chessclub42.conf
<VirtualHost *:8000>
    ServerAdmin webmaster@localhost
    DocumentRoot /var/www/html/chessclub42
</VirtualHost>
[root@CentOS65 ~]# vi /etc/httpd/conf.d/hunter2.conf
[root@CentOS65 ~]# cat /etc/httpd/conf.d/hunter2.conf
<VirtualHost *:9000>
    ServerAdmin webmaster@localhost
    DocumentRoot /var/www/html/hunter2
</VirtualHost>
```

Notice the different port numbers 7000, 8000 and 9000. Notice also that we specified a unique DocumentRoot for each website.

14.5.3. three extra ports

We need to enable these three ports on apache in the httpd.conf file.

```
[root@CentOS65 ~]# vi /etc/httpd/conf/httpd.conf
root@linux:~# grep ^Listen /etc/httpd/conf/httpd.conf
Listen 80
Listen 7000
Listen 8000
Listen 9000
```

14.5.4. SELinux guards our ports

If we try to restart our server, we will notice the following error:

```
[root@CentOS65 ~]# service httpd restart
Stopping httpd:                               [ OK ]
Starting httpd:
  (13)Permission denied: make_sock: could not bind to address 0.0.0.0:7000
no listening sockets available, shutting down
                                           [FAILED]
```

This is due to SELinux reserving ports 7000 and 8000 for other uses. We need to tell SELinux we want to use these ports for http traffic

```
[root@CentOS65 ~]# semanage port -m -t http_port_t -p tcp 7000
[root@CentOS65 ~]# semanage port -m -t http_port_t -p tcp 8000
[root@CentOS65 ~]# service httpd restart
Stopping httpd:                               [ OK ]
Starting httpd:                               [ OK ]
```

14.5.5. three extra websites

Next we need to create three DocumentRoot directories.

```
[root@CentOS65 ~]# mkdir /var/www/html/choochoo
[root@CentOS65 ~]# mkdir /var/www/html/chessclub42
[root@CentOS65 ~]# mkdir /var/www/html/hunter2
```

And we have to put some really simple website in those directories.

```
[root@CentOS65 ~]# echo 'Choo Choo model train Choo Choo' > /var/www/html/chooc\
hoo/index.html
[root@CentOS65 ~]# echo 'Welcome to chess club 42' > /var/www/html/chessclub42/\
index.html
[root@CentOS65 ~]# echo 'HaCkInG iS fUn At HuNtEr2' > /var/www/html/hunter2/ind\
ex.html
```

14.5.6. enabling extra websites

The only way to enable or disable configurations in RHEL/CentOS is by renaming or moving the configuration files. Any file in /etc/httpd/conf.d ending on .conf will be loaded by Apache. To disable a site we can either rename the file or move it to another directory.

The files are created, so we can tell apache.

```
[root@CentOS65 ~]# ls /etc/httpd/conf.d/
chessclub42.conf choochoo.conf hunter2.conf  README  welcome.conf
[root@CentOS65 ~]# service httpd reload
Reloading httpd:
```

14.5.7. testing the three websites

Testing the model train club named Choo Choo on port 7000.

```
[root@CentOS65 ~]# wget 127.0.0.1:7000
--2014-05-11 11:59:36-- http://127.0.0.1:7000/
Connecting to 127.0.0.1:7000 ... connected.
HTTP request sent, awaiting response ... 200 OK
Length: 32 [text/html]
Saving to: `index.html'

100%[=====] 32          --.-K/s  in 0s

2014-05-11 11:59:36 (4.47 MB/s) - `index.html' saved [32/32]

[root@CentOS65 ~]# cat index.html
Choo Choo model train Choo Choo
```

Testing the chess club named Chess Club 42 on port 8000.

```
[root@CentOS65 ~]# wget 127.0.0.1:8000
--2014-05-11 12:01:30-- http://127.0.0.1:8000/
Connecting to 127.0.0.1:8000 ... connected.
HTTP request sent, awaiting response ... 200 OK
Length: 25 [text/html]
Saving to: `index.html.1'

100%[=====] 25          --.-K/s  in 0s

2014-05-11 12:01:30 (4.25 MB/s) - `index.html.1' saved [25/25]

root@linux:/etc/apache2# cat index.html.1
Welcome to chess club 42
```

Testing the hacker club named hunter2 on port 9000.

```
[root@CentOS65 ~]# wget 127.0.0.1:9000
--2014-05-11 12:02:37-- http://127.0.0.1:9000/
Connecting to 127.0.0.1:9000 ... connected.
HTTP request sent, awaiting response ... 200 OK
Length: 26 [text/html]
Saving to: `index.html.2'

100%[=====] 26          --.-K/s  in 0s

2014-05-11 12:02:37 (4.49 MB/s) - `index.html.2' saved [26/26]

root@linux:/etc/apache2# cat index.html.2
HaCkInG iS fUn At HuNtEr2
```

Cleaning up the temporary files.

```
[root@CentOS65 ~]# rm index.html index.html.1 index.html.2
```


14.5.8. firewall rules

If we attempt to access the site from another machine however, we will not be able to view the website yet. The firewall is blocking incoming connections. We need to open these incoming ports first

```
[root@CentOS65 ~]# iptables -I INPUT -p tcp --dport 80 -j ACCEPT
[root@CentOS65 ~]# iptables -I INPUT -p tcp --dport 7000 -j ACCEPT
[root@CentOS65 ~]# iptables -I INPUT -p tcp --dport 8000 -j ACCEPT
[root@CentOS65 ~]# iptables -I INPUT -p tcp --dport 9000 -j ACCEPT
```

And if we want these rules to remain active after a reboot, we need to save them

```
[root@CentOS65 ~]# service iptables save
iptables: Saving firewall rules to /etc/sysconfig/iptables:[ OK ]
```

14.6. named virtual hosts on CentOS

14.6.1. named virtual hosts

The chess club and the model train club find the port numbers too hard to remember. They would prefer to have their website accessible by name.

We continue work on the same server that has three websites on three ports. We need to make sure those websites are accessible using the names `choochoo.local`, `chessclub42.local` and `hunter2.local`.

First, we need to enable named virtual hosts in the configuration

```
[root@CentOS65 ~]# vi /etc/httpd/conf/httpd.conf
[root@CentOS65 ~]# grep ^NameVirtualHost /etc/httpd/conf/httpd.conf
NameVirtualHost *:80
[root@CentOS65 ~]#
```

Next we need to create three new virtualhosts.

```
[root@CentOS65 ~]# vi /etc/httpd/conf.d/choochoo.local.conf
[root@CentOS65 ~]# vi /etc/httpd/conf.d/chessclub42.local.conf
[root@CentOS65 ~]# vi /etc/httpd/conf.d/hunter2.local.conf
[root@CentOS65 ~]# cat /etc/httpd/conf.d/choochoo.local.conf
<VirtualHost *:80>
    ServerAdmin webmaster@localhost
    ServerName choochoo.local
    DocumentRoot /var/www/html/choochoo
</VirtualHost>
[root@CentOS65 ~]# cat /etc/httpd/conf.d/chessclub42.local.conf
<VirtualHost *:80>
    ServerAdmin webmaster@localhost
    ServerName chessclub42.local
    DocumentRoot /var/www/html/chessclub42
</VirtualHost>
[root@CentOS65 ~]# cat /etc/httpd/conf.d/hunter2.local.conf
<VirtualHost *:80>
    ServerAdmin webmaster@localhost
    ServerName hunter2.local
```

14. apache web server

```
        DocumentRoot /var/www/html/hunter2
</VirtualHost>
[root@CentOS65 ~]#
```

Notice that they all listen on port 80 and have an extra `ServerName` directive.

14.6.2. name resolution

We need some way to resolve names. This can be done with DNS, which is discussed in another chapter. For this demo it is also possible to quickly add the three names to the `/etc/hosts` file.

```
[root@CentOS65 ~]# grep ^192 /etc/hosts
192.168.1.225 choochoo.local
192.168.1.225 chessclub42.local
192.168.1.225 hunter2.local
```

Note that you may have another ip address...

14.6.3. reload and verify

After a service `httpd` reload the websites should be available by name.

```
[root@CentOS65 ~]# service httpd reload
Reloading httpd:
[root@CentOS65 ~]# wget chessclub42.local
--2014-05-25 16:59:14-- http://chessclub42.local/
Resolving chessclub42.local ... 192.168.1.225
Connecting to chessclub42.local|192.168.1.225|:80 ... connected.
HTTP request sent, awaiting response ... 200 OK
Length: 25 [text/html]
Saving to: 'index.html'

100%[=====] 25          --.-K/s   in 0s

2014-05-25 16:59:15 (1014 KB/s) - 'index.html' saved [25/25]

[root@CentOS65 ~]# cat index.html
Welcome to chess club 42
```

14.7. password protected website on CentOS

You can secure files and directories in your website with a `.htaccess` file that refers to a `.htpasswd` file. The `htpasswd` command can create a `.htpasswd` file that contains a userid and an (encrypted) password.

This screenshot creates a user and password for the hacker named `cliff` and uses the `-c` flag to create the `.htpasswd` file.

```
[root@CentOS65 ~]# htpasswd -c /var/www/.htpasswd cliff
New password:
Re-type new password:
Adding password for user cliff
[root@CentOS65 ~]# cat /var/www/.htpasswd
cliff:QNwTrymMLBctU
```

Hacker rob also wants access, this screenshot shows how to add a second user and password to .htpasswd.

```
[root@CentOS65 ~]# htpasswd /var/www/.htpasswd rob
New password:
Re-type new password:
Adding password for user rob
[root@CentOS65 ~]# cat /var/www/.htpasswd
cliff:QNwTrymMLBctU
rob:EC2v0CcrMXDoM
[root@CentOS65 ~]#
```

Both Cliff and Rob chose the same password (hunter2), but that is not visible in the .htpasswd file because of the different salts.

Next we need to create a .htaccess file in the DocumentRoot of the website we want to protect. This screenshot shows an example.

```
[root@CentOS65 ~]# cat /var/www/html/hunter2/.htaccess
AuthUserFile /var/www/.htpasswd
AuthName "Members only!"
AuthType Basic
require valid-user
```

Note that we are protecting the website on port 9000 that we created earlier.

And because we put the website for the Hackerspace named hunter2 in a subdirectory of the default website, we will need to adjust the AllowOverride parameter in /etc/httpd/conf/httpd.conf under the <Directory "/var/www/html"> directive as this screenshot shows.

```
[root@CentOS65 ~]# vi /etc/httpd/conf/httpd.conf

<Directory "/var/www/html">

#
# Possible values for the Options directive are "None", "All",
# or any combination of:
#   Indexes Includes FollowSymLinks SymLinksifOwnerMatch ExecCGI MultiViews
#
# Note that "MultiViews" must be named *explicitly* --- "Options All"
# doesn't give it to you.
#
# The Options directive is both complicated and important. Please see
# http://httpd.apache.org/docs/2.2/mod/core.html#options
# for more information.
#
    Options Indexes FollowSymLinks

#
# AllowOverride controls what directives may be placed in .htaccess files.
```

14. apache web server

```
# It can be "All", "None", or any combination of the keywords:
# Options FileInfo AuthConfig Limit
#
#     AllowOverride Authconfig
#
# Controls who can get stuff from this server.
#
#     Order allow,deny
#     Allow from all
</Directory>
```

Now restart the apache2 server and test that it works!

14.8. troubleshooting apache

When apache restarts, it will verify the syntax of files in the configuration folder /etc/apache2 on debian or /etc/httpd on CentOS and it will tell you the name of the faulty file, the line number and an explanation of the error.

```
root@linux:~# service apache2 restart
apache2: Syntax error on line 268 of /etc/apache2/apache2.conf: Syntax error on
line 1 of /etc/apache2/sites-enabled/chessclub42: /etc/apache2/sites-
enabled\
/chessclub42:4: <VirtualHost> was not closed.\n/etc/apache2/sites-enabled/ches\
sclub42:1: <VirtualHost> was not closed.
Action 'configtest' failed.
The Apache error log may have more information.
failed!
```

Below you see the problem... a missing / before on line 4.

```
root@linux:~# cat /etc/apache2/sites-available/chessclub42
<VirtualHost *:8000>
    ServerAdmin webmaster@localhost
    DocumentRoot /var/www/chessclub42
<VirtualHost>
```

Let us force another error by renaming the directory of one of our websites:

```
root@linux:~# mv /var/www/choochoo/ /var/www/chooshoo
root@linux:~# !ser
service apache2 restart
Restarting web server: apache2Warning: DocumentRoot [/var/www/choochoo] does n\
ot exist
Warning: DocumentRoot [/var/www/choochoo] does not exist
... waiting Warning: DocumentRoot [/var/www/choochoo] does not exist
Warning: DocumentRoot [/var/www/choochoo] does not exist
.
```

As you can see, apache will tell you exactly what is wrong.

You can also troubleshoot by connecting to the website via a browser and then checking the apache log files in /var/log/apache.

14.9. virtual hosts example

Below is a sample virtual host configuration. This virtual hosts overrides the default Apache `ErrorDocument` directive.

```
<VirtualHost 83.217.76.245:80>
ServerName cobbaut.be
ServerAlias www.cobbaut.be
DocumentRoot /home/paul/public_html
ErrorLog /home/paul/logs/error_log
CustomLog /home/paul/logs/access_log common
ScriptAlias /cgi-bin/ /home/paul/cgi-bin/
<Directory /home/paul/public_html>
    Options Indexes IncludesNOEXEC FollowSymLinks
    allow from all
</Directory>
ErrorDocument 404 http://www.cobbaut.be/cobbaut.php
</VirtualHost>
```

14.10. aliases and redirects

Apache supports aliases for directories, like this example shows.

```
Alias /paul/ "/home/paul/public_html/"
```

Similarly, content can be redirected to another website or web server.

```
Redirect permanent /foo http://www.foo.com/bar
```

14.11. more on .htaccess

You can do much more with `.htaccess`. One example is to use `.htaccess` to prevent people from certain domains to access your website. Like in this case, where a number of referer spammers are blocked from the website.

```
student@linux:~/cobbaut.be$ cat .htaccess
# Options +FollowSymlinks
RewriteEngine On
RewriteCond %{HTTP_REFERER} ^http://(www\.)?buy-adipex.fw.nu.*$ [OR]
RewriteCond %{HTTP_REFERER} ^http://(www\.)?buy-levitra.asso.ws.*$ [NC,OR]
RewriteCond %{HTTP_REFERER} ^http://(www\.)?buy-tramadol.fw.nu.*$ [NC,OR]
RewriteCond %{HTTP_REFERER} ^http://(www\.)?buy-viagra.lookin.at.*$ [NC,OR]
...
RewriteCond %{HTTP_REFERER} ^http://(www\.)?www.healthinsurancehelp.net.*$ [NC]
RewriteRule .* - [F,L]
student@linux:~/cobbaut.be$
```

14.12. traffic

Apache keeps a log of all visitors. The `webalizer` is often used to parse this log into nice html statistics.

14.13. self signed cert on Debian

Below is a very quick guide on setting up Apache2 on Debian 7 with a self-signed certificate.

Chances are these packages are already installed.

```
root@linux:~# aptitude install apache2 openssl
No packages will be installed, upgraded, or removed.
0 packages upgraded, 0 newly installed, 0 to remove and 0 not upgraded.
Need to get 0 B of archives. After unpacking 0 B will be used.
```

Create a directory to store the certs, and use openssl to create a self signed cert that is valid for 999 days.

```
root@linux:~# mkdir /etc/ssl/localcerts
root@linux:~# openssl req -new -x509 -days 999 -nodes -out /etc/ssl/local\
certs/apache.pem -keyout /etc/ssl/localcerts/apache.key
Generating a 2048 bit RSA private key
...
...
writing new private key to '/etc/ssl/localcerts/apache.key'
-----
You are about to be asked to enter information that will be incorporated
into your certificate request.
What you are about to enter is what is called a Distinguished Name or a DN.
There are quite a few fields but you can leave some blank
For some fields there will be a default value,
If you enter '.', the field will be left blank.
-----
Country Name (2 letter code) [AU]:BE
State or Province Name (full name) [Some-State]:Antwerp
Locality Name (eg, city) []:Antwerp
Organization Name (eg, company) [Internet Widgits Pty Ltd]:linux-training.be
Organizational Unit Name (eg, section) []:
Common Name (e.g. server FQDN or YOUR name) []:Paul
Email Address []:
```

A little security never hurt anyone.

```
root@linux:~# ls -l /etc/ssl/localcerts/
total 8
-rw-r--r-- 1 root root 1704 Sep 16 18:24 apache.key
-rw-r--r-- 1 root root 1302 Sep 16 18:24 apache.pem
root@linux:~# chmod 600 /etc/ssl/localcerts/*
root@linux:~# ls -l /etc/ssl/localcerts/
total 8
-rw----- 1 root root 1704 Sep 16 18:24 apache.key
-rw----- 1 root root 1302 Sep 16 18:24 apache.pem
```

Enable the apache ssl mod.

```
root@linux:~# a2enmod ssl
Enabling module ssl.
See /usr/share/doc/apache2.2-common/README.Debian.gz on how to configure SSL\
and create self-signed certificates.
To activate the new configuration, you need to run:
    service apache2 restart
```

Create the website configuration.

```
root@linux:~# vi /etc/apache2/sites-available/choochoos
root@linux:~# cat /etc/apache2/sites-available/choochoos
<VirtualHost *:7000>
    ServerAdmin webmaster@localhost
    DocumentRoot /var/www/choochoos
    SSLEngine On
    SSLCertificateFile /etc/ssl/localcerts/apache.pem
    SSLCertificateKeyFile /etc/ssl/localcerts/apache.key
</VirtualHost>
root@linux:~#
```

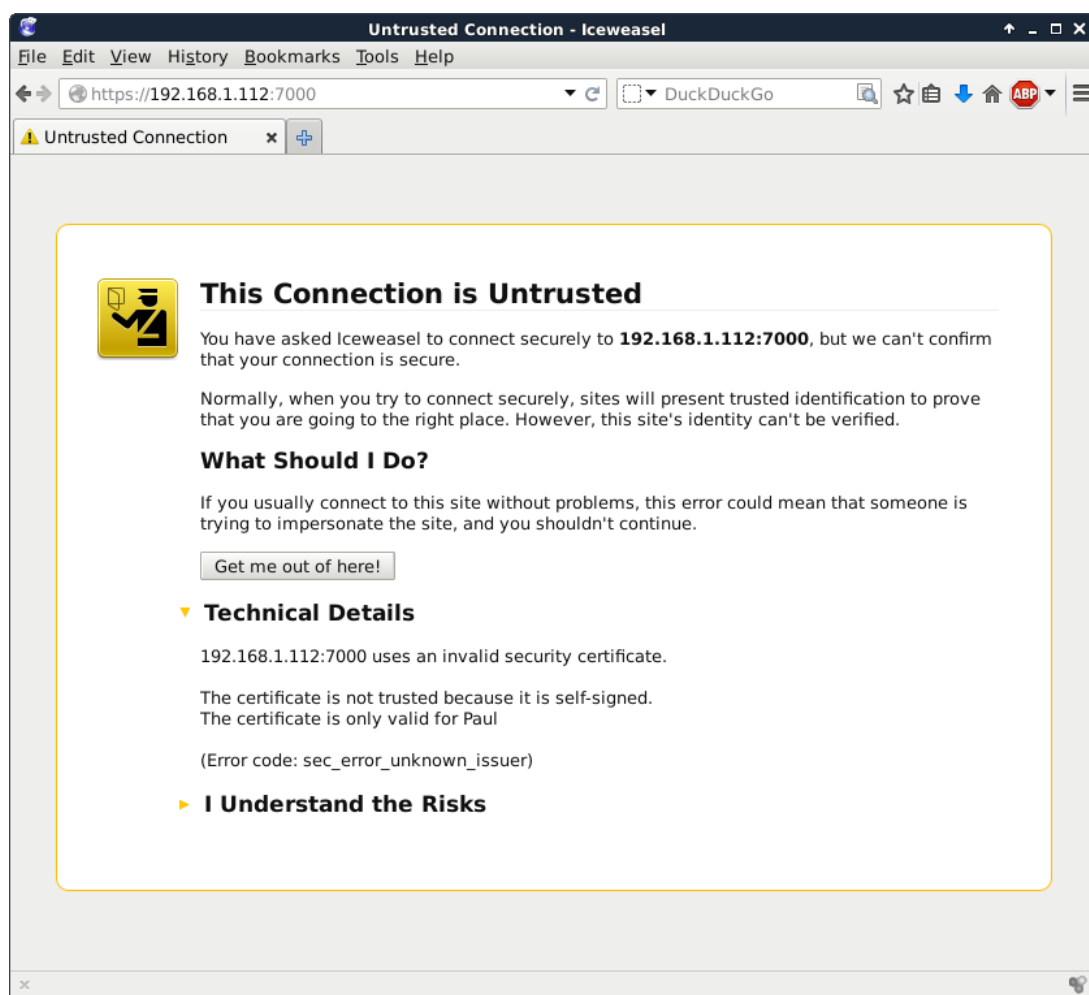
And create the website itself.

```
root@linux:/var/www/choochoos# vi index.html
root@linux:/var/www/choochoos# cat index.html
Choo Choo HTTPS secured model train Choo Choo
```

Enable the website and restart (or reload) apache2.

```
root@linux:/var/www/choochoos# a2ensite choochoos
Enabling site choochoos.
To activate the new configuration, you need to run:
  service apache2 reload
root@linux:/var/www/choochoos# service apache2 restart
Restarting web server: apache2 ... waiting .
```

Chances are your browser will warn you about the self signed certificate.



14.14. self signed cert on RHEL/CentOS

Below is a quick way to create a self signed cert for https on RHEL/CentOS. You may need these packages:

```
[root@paulserver ~]# yum install httpd openssl mod_ssl
Loaded plugins: fastestmirror
Loading mirror speeds from cached hostfile
 * base: ftp.belnet.be
 * extras: ftp.belnet.be
 * updates: mirrors.vooservers.com
base | 3.7 kB 00:00
Setting up Install Process
Package httpd-2.2.15-31.el6.centos.x86_64 already installed and latest version
Package openssl-1.0.1e-16.el6_5.15.x86_64 already installed and latest version
Package 1:mod_ssl-2.2.15-31.el6.centos.x86_64 already ins ... and latest version
Nothing to do
```

We use openssl to create the certificate.

```
[root@paulserver ~]# mkdir certs
[root@paulserver ~]# cd certs
[root@paulserver certs]# openssl genrsa -out ca.key 2048
```


Generating RSA private key, 2048 bit long modulus

```
.....+++
.....+++
```

e is 65537 (0x10001)

```
[root@paulserver certs]# openssl req -new -key ca.key -out ca.csr
```

You are about to be asked to enter information that will be incorporated into your certificate request.

What you are about to enter is what is called a Distinguished Name or a DN.

There are quite a few fields but you can leave some blank

For some fields there will be a default value,

If you enter '.', the field will be left blank.

```
-----
```

Country Name (2 letter code) [XX]:BE

State or Province Name (full name) []:antwerp

Locality Name (eg, city) [Default City]:antwerp

Organization Name (eg, company) [Default Company Ltd]:antwerp

Organizational Unit Name (eg, section) []:

Common Name (eg, your name or your server's hostname) []:paulserver

Email Address []:

Please enter the following 'extra' attributes

to be sent with your certificate request

A challenge password []:

An optional company name []:

```
[root@paulserver certs]# openssl x509 -req -days 365 -in ca.csr -signkey ca.ke\
y -out ca.crt
```

Signature ok

subject=/C=BE/ST=antwerp/L=antwerp/O=antwerp/CN=paulserver

Getting Private key

We copy the keys to the right location (You may be missing SELinux info here).

```
[root@paulserver certs]# cp ca.crt /etc/pki/tls/certs/
```

```
[root@paulserver certs]# cp ca.key ca.csr /etc/pki/tls/private/
```

We add the location of our keys to this file, and also add the NameVirtualHost *:443 directive.

```
[root@paulserver certs]# vi /etc/httpd/conf.d/ssl.conf
```

```
[root@paulserver certs]# grep ^SSLCerti /etc/httpd/conf.d/ssl.conf
```

```
SSLCertificateFile /etc/pki/tls/certs/ca.crt
```

```
SSLCertificateKeyFile /etc/pki/tls/private/ca.key
```

Create a website configuration.

```
[root@paulserver certs]# vi /etc/httpd/conf.d/choochoos.conf
```

```
[root@paulserver certs]# cat /etc/httpd/conf.d/choochoos.conf
```

```
<VirtualHost *:443>
```

```
    SSLEngine on
```

```
    SSLCertificateFile /etc/pki/tls/certs/ca.crt
```

```
    SSLCertificateKeyFile /etc/pki/tls/private/ca.key
```

```
    DocumentRoot /var/www/choochoos
```

```
    ServerName paulserver
```

```
</VirtualHost>
```

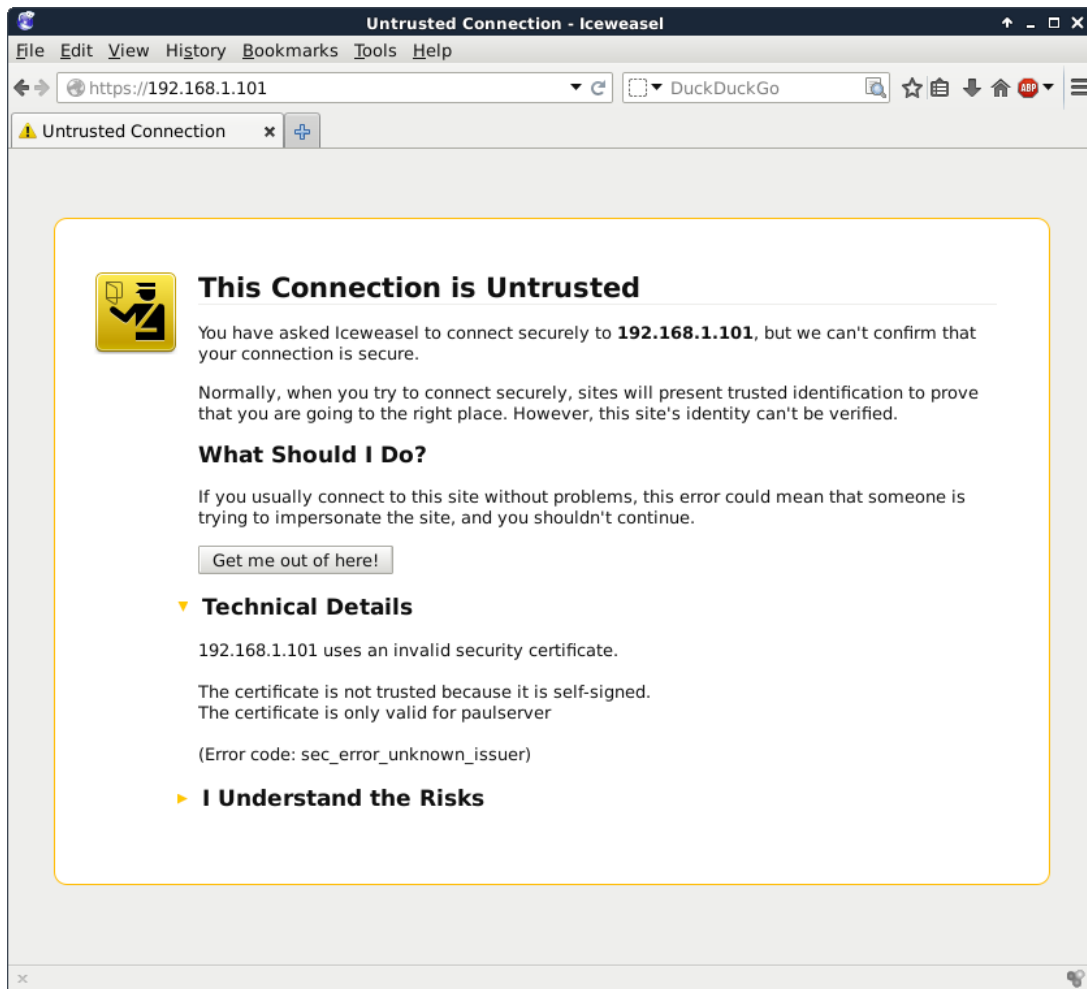
```
[root@paulserver certs]#
```

Create a simple website and restart apache.

14. apache web server

```
[root@paulserver certs]# mkdir /var/www/choochoos
[root@paulserver certs]# echo HTTPS model train choochoos > /var/www/choochoos/\
index.html
[root@paulserver httpd]# service httpd restart
Stopping httpd:          [ OK ]
Starting httpd:         [ OK ]
```

And your browser will probably warn you that this certificate is self signed.



14.15. practice: apache

1. Verify that Apache is installed and running.
2. Browse to the Apache HTML manual.
3. Create three virtual hosts that listen on ports 8472, 31337 and 1201. Test that it all works.
4. Create three named virtual hosts startrek.local, starwars.local and stargate.local. Test that it all works.
5. Create a virtual hosts that listens on another ip-address.
6. Protect one of your websites with a user/password combo.

15. scripting loops

(Written by Paul Cobbaut, <https://github.com/paulcobbaut/>, with contributions by: Alex M. Schapelle, <https://github.com/zero-pythagoras/>)

15.1. test []

The test command can test whether something is true or false. Let's start by testing whether 10 is greater than 55.

```
[student@linux ~]$ test 10 -gt 55 ; echo $?
1
[student@linux ~]$
```

The test command returns 1 if the test fails. And as you see in the next screenshot, test returns 0 when a test succeeds.

```
[student@linux ~]$ test 56 -gt 55 ; echo $?
0
[student@linux ~]$
```

If you prefer true and false, then write the test like this.

```
[student@linux ~]$ test 56 -gt 55 && echo true || echo false
true
[student@linux ~]$ test 6 -gt 55 && echo true || echo false
false
```

The test command can also be written as square brackets, the screenshot below is identical to the one above.

```
[student@linux ~]$ [ 56 -gt 55 ] && echo true || echo false
true
[student@linux ~]$ [ 6 -gt 55 ] && echo true || echo false
false
```

Below are some example tests. Take a look at `man test` to see more options for tests.

[-d foo]	Does the directory foo exist ?
[-e bar]	Does the file bar exist ?
['/etc' = \$PWD]	Is the string /etc equal to the variable \$PWD ?
[\$1 ≠ 'secret']	Is the first parameter different from secret ?
[55 -lt \$bar]	Is 55 less than the value of \$bar ?
[\$foo -ge 1000]	Is the value of \$foo greater or equal to 1000 ?
["abc" < \$bar]	Does abc sort before the value of \$bar ?
[-f foo]	Is foo a regular file ?
[-r bar]	Is bar a readable file ?
[foo -nt bar]	Is file foo newer than file bar ?
[-o nounset]	Is the shell option nounset set ?

15. scripting loops

Tests can be combined with logical AND and OR.

```
student@linux:~$ [ 66 -gt 55 -a 66 -lt 500 ] && echo true || echo false
true
student@linux:~$ [ 66 -gt 55 -a 660 -lt 500 ] && echo true || echo false
false
student@linux:~$ [ 66 -gt 55 -o 660 -lt 500 ] && echo true || echo false
true
```

15.2. if then else

The `if then else` construction is about choice. If a certain condition is met, then execute something, else execute something else. The example below tests whether a file exists, and if the file exists then a proper message is echoed.

```
#!/bin/bash

if [ -f isit.txt ]
then echo isit.txt exists!
else echo isit.txt not found!
fi
```

If we name the above script 'choice', then it executes like this.

```
[student@linux scripts]$ ./choice
isit.txt not found!
[student@linux scripts]$ touch isit.txt
[student@linux scripts]$ ./choice
isit.txt exists!
[student@linux scripts]$
```

15.3. if then elif

You can nest a new `if` inside an `else` with `elif`. This is a simple example.

```
#!/bin/bash
count=42
if [ $count -eq 42 ]
then
    echo "42 is correct."
elif [ $count -gt 42 ]
then
    echo "Too much."
else
    echo "Not enough."
fi
```

15.4. for loop

The example below shows the syntax of a classical `for` loop in bash.

```
for i in 1 2 4
do
    echo $i
done
```

An example of a `for` loop combined with an embedded shell.

```
#!/bin/ksh
for counter in `seq 1 20`
do
    echo counting from 1 to 20, now at $counter
    sleep 1
done
```

The same example as above can be written without the embedded shell using the bash `{from..to}` shorthand.

```
#!/bin/bash
for counter in {1..20}
do
    echo counting from 1 to 20, now at $counter
    sleep 1
done
```

This `for` loop uses file globbing (from the shell expansion). Putting the instruction on the command line has identical functionality.

```
kahlan@solexp11$ ls
count.ksh  go.ksh
kahlan@solexp11$ for file in *.ksh ; do cp $file $file.backup ; done
kahlan@solexp11$ ls
count.ksh  count.ksh.backup  go.ksh  go.ksh.backup
```

15.5. while loop

Below a simple example of a `while` loop.

```
i=100;
while [ $i -ge 0 ] ;
do
    echo Counting down, from 100 to 0, now at $i;
    let i--;
done
```

Endless loops can be made with `while true` or `while :`, where the colon is the equivalent of `no` operation in the Korn and bash shells.

15. scripting loops

```
#!/bin/ksh
# endless loop
while :
do
  echo hello
  sleep 1
done
```

15.6. until loop

Below a simple example of an until loop.

```
let i=100;
until [ $i -le 0 ] ;
do
  echo Counting down, from 100 to 1, now at $i;
  let i--;
done
```

15.7. practice: scripting tests and loops

1. Write a script that uses a `for` loop to count from 3 to 7.
2. Write a script that uses a `for` loop to count from 1 to 17000.
3. Write a script that uses a `while` loop to count from 3 to 7.
4. Write a script that uses an `until` loop to count down from 8 to 4.
5. Write a script that counts the number of files ending in `.txt` in the current directory.
6. Wrap an `if` statement around the script so it is also correct when there are zero files ending in `.txt`.

15.8. solution: scripting tests and loops

1. Write a script that uses a `for` loop to count from 3 to 7.

```
1  #!/bin/bash
2
3  for i in 3 4 5 6 7
4  do
5    echo "Counting from 3 to 7, now at ${i}"
6  done
```

2. Write a script that uses a `for` loop to count from 1 to 17000.

```
1  #!/bin/bash
2
3  for i in `seq 1 17000`
4  do
5    echo "Counting from 1 to 17000, now at ${i}"
6  done
```

3. Write a script that uses a `while` loop to count from 3 to 7.

```

1  #!/bin/bash
2
3  i=3
4  while [ $i -le 7 ]
5  do
6    echo "Counting from 3 to 7, now at ${i}"
7    let i=i+1
8  done

```

4. Write a script that uses an until loop to count down from 8 to 4.

```

1  #!/bin/bash
2
3  i=8
4  until [ $i -lt 4 ]
5  do
6    echo "Counting down from 8 to 4, now at ${i}"
7    let i=i-1
8  done

```

5. Write a script that counts the number of files ending in .txt in the current directory.

```

1  #!/bin/bash
2
3  let i=0
4  for file in *.txt
5  do
6    let i++
7  done
8  echo "There are ${i} files ending in .txt"

```

6. Wrap an if statement around the script so it is also correct when there are zero files ending in .txt.

```

1  #!/bin/bash
2
3  ls *.txt > /dev/null 2>&1
4  if [ $? -ne 0 ]
5  then echo "There are 0 files ending in .txt"
6  else
7    let i=0
8    for file in *.txt
9    do
10     let i++
11   done
12   echo "There are ${i} files ending in .txt"
13 fi

```


16. scripting parameters

(Written by Paul Cobbaut, <https://github.com/paulcobbaut/>, with contributions by: Alex M. Schapelle, <https://github.com/zero-pythagoras/>)

16.1. script parameters

A bash shell script can have parameters. The numbering you see in the script below continues if you have more parameters. You also have special parameters containing the number of parameters, a string of all of them, and also the process id, and the last return code. The man page of bash has a full list.

```
#!/bin/bash
echo The first argument is $1
echo The second argument is $2
echo The third argument is $3

echo \$ $$ PID of the script
echo \# $# count arguments
echo \? $? last return code
echo \* $* all the arguments
```

Below is the output of the script above in action.

```
[student@linux scripts]$ ./pars one two three
The first argument is one
The second argument is two
The third argument is three
$ 5610 PID of the script
# 3 count arguments
? 0 last return code
* one two three all the arguments
```

Once more the same script, but with only two parameters.

```
[student@linux scripts]$ ./pars 1 2
The first argument is 1
The second argument is 2
The third argument is
$ 5612 PID of the script
# 2 count arguments
? 0 last return code
* 1 2 all the arguments
[student@linux scripts]$
```

Here is another example, where we use \$0. The \$0 parameter contains the name of the script.

16. scripting parameters

```
student@linux~$ cat myname
echo this script is called $0
student@linux~$ ./myname
this script is called ./myname
student@linux~$ mv myname test42
student@linux~$ ./test42
this script is called ./test42
```

16.2. shift through parameters

The shift statement can parse all parameters one by one. This is a sample script.

```
kahlan@solexp11$ cat shift.ksh
#!/bin/ksh

if [ "$#" = "0" ]
then
    echo You have to give at least one parameter.
    exit 1
fi

while (( $# ))
do
    echo You gave me $1
    shift
done
```

Below is some sample output of the script above.

```
kahlan@solexp11$ ./shift.ksh one
You gave me one
kahlan@solexp11$ ./shift.ksh one two three 1201 "33 42"
You gave me one
You gave me two
You gave me three
You gave me 1201
You gave me 33 42
kahlan@solexp11$ ./shift.ksh
You have to give at least one parameter.
```

16.3. runtime input

You can ask the user for input with the read command in a script.

```
#!/bin/bash
echo -n Enter a number:
read number
```

16.4. sourcing a config file

The source (as seen in the shell chapters) can be used to source a configuration file.

Below a sample configuration file for an application.

```
[student@linux scripts]$ cat myApp.conf
# The config file of myApp

# Enter the path here
myAppPath=/var/myApp

# Enter the number of quines here
quines=5
```

And here an application that uses this file.

```
[student@linux scripts]$ cat myApp.bash
#!/bin/bash
#
# Welcome to the myApp application
#

. ./myApp.conf

echo There are $quines quines
```

The running application can use the values inside the sourced configuration file.

```
[student@linux scripts]$ ./myApp.bash
There are 5 quines
[student@linux scripts]$
```

16.5. get script options with getopt

The getopt function allows you to parse options given to a command. The following script allows for any combination of the options a, f and z.

```
kahlan@solexp11$ cat options.ksh
#!/bin/ksh

while getopt ":afz" option;
do
  case $option in
    a)
      echo received -a
      ;;
    f)
      echo received -f
      ;;
    z)
      echo received -z
      ;;
    *)
      echo "invalid option -$OPTARG"
```

16. scripting parameters

```
;;
esac
done
```

This is sample output from the script above. First we use correct options, then we enter twice an invalid option.

```
kahlan@solexp11$ ./options.ksh
kahlan@solexp11$ ./options.ksh -af
received -a
received -f
kahlan@solexp11$ ./options.ksh -zfg
received -z
received -f
invalid option -g
kahlan@solexp11$ ./options.ksh -a -b -z
received -a
invalid option -b
received -z
```

You can also check for options that need an argument, as this example shows.

```
kahlan@solexp11$ cat argoptions.ksh
#!/bin/ksh

while getopts ":af:z" option;
do
  case $option in
    a)
      echo received -a
      ;;
    f)
      echo received -f with $OPTARG
      ;;
    z)
      echo received -z
      ;;
    :)
      echo "option -$OPTARG needs an argument"
      ;;
    *)
      echo "invalid option -$OPTARG"
      ;;
  esac
done
```

This is sample output from the script above.

```
kahlan@solexp11$ ./argoptions.ksh -a -f hello -z
received -a
received -f with hello
received -z
kahlan@solexp11$ ./argoptions.ksh -zaf 42
received -z
received -a
received -f with 42
kahlan@solexp11$ ./argoptions.ksh -zf
received -z
option -f needs an argument
```

16.6. get shell options with shopt

You can toggle the values of variables controlling optional shell behaviour with the `shopt` built-in shell command. The example below first verifies whether the `cdspell` option is set; it is not. The next `shopt` command sets the value, and the third `shopt` command verifies that the option really is set. You can now use minor spelling mistakes in the `cd` command. The man page of `bash` has a complete list of options.

```
student@linux:~$ shopt -q cdspell ; echo $?
1
student@linux:~$ shopt -s cdspell
student@linux:~$ shopt -q cdspell ; echo $?
0
student@linux:~$ cd /Etc
/etc
```

16.7. practice: parameters and options

1. Write a script that receives four parameters, and outputs them in reverse order.
2. Write a script that receives two parameters (two filenames) and outputs whether those files exist.
3. Write a script that asks for a filename. Verify existence of the file, then verify that you own the file, and whether it is writable. If not, then make it writable.
4. Make a configuration file for the previous script. Put a logging switch in the config file, logging means writing detailed output of everything the script does to a log file in `/tmp`.

16.8. solution: parameters and options

1. Write a script that receives four parameters, and outputs them in reverse order.

```
echo $4 $3 $2 $1
```

2. Write a script that receives two parameters (two filenames) and outputs whether those files exist.

```
#!/bin/bash

if [ -f $1 ]
then echo $1 exists!
else echo $1 not found!
fi

if [ -f $2 ]
then echo $2 exists!
else echo $2 not found!
fi
```

3. Write a script that asks for a filename. Verify existence of the file, then verify that you own the file, and whether it is writable. If not, then make it writable.
4. Make a configuration file for the previous script. Put a logging switch in the config file, logging means writing detailed output of everything the script does to a log file in `/tmp`.

Part VI.

Advanced text processing

17. file globbing

(Written by Paul Cobbaut, <https://github.com/paulcobbaut/>, with contributions by: Alex M. Schapelle, <https://github.com/zero-pythagoras/>)

Typing `man 7 glob` (on Debian) will tell you that long ago there was a program called `/etc/glob` that would expand wildcard patterns.

Today the shell is responsible for file globbing (or dynamic filename generation). This chapter will explain file globbing.

17.1. * asterisk

The asterisk `*` is interpreted by the shell as a sign to generate filenames, matching the asterisk to any combination of characters (even none). When no path is given, the shell will use filenames in the current directory. See the man page of `glob(7)` for more information. (This is part of LPI topic 1.103.3.)

```
[student@linux gen]$ ls
file1 file2 file3 File4 File55 FileA fileab Fileab FileAB fileabc
[student@linux gen]$ ls File*
File4 File55 FileA Fileab FileAB
[student@linux gen]$ ls file*
file1 file2 file3 fileab fileabc
[student@linux gen]$ ls *ile55
File55
[student@linux gen]$ ls F*ile55
File55
[student@linux gen]$ ls F*55
File55
[student@linux gen]$
```

17.2. ? question mark

Similar to the asterisk, the question mark `?` is interpreted by the shell as a sign to generate filenames, matching the question mark with exactly one character.

```
[student@linux gen]$ ls
file1 file2 file3 File4 File55 FileA fileab Fileab FileAB fileabc
[student@linux gen]$ ls File?
File4 FileA
[student@linux gen]$ ls Fil?4
File4
[student@linux gen]$ ls Fil??
File4 FileA
[student@linux gen]$ ls File??
File55 Fileab FileAB
[student@linux gen]$
```

17.3. [] square brackets

The square bracket [is interpreted by the shell as a sign to generate filenames, matching any of the characters between [and the first subsequent]. The order in this list between the brackets is not important. Each pair of brackets is replaced by exactly one character.

```
[student@linux gen]$ ls
file1 file2 file3 File4 File55 FileA fileab Fileab FileAB fileabc
[student@linux gen]$ ls File[5A]
FileA
[student@linux gen]$ ls File[A5]
FileA
[student@linux gen]$ ls File[A5][5b]
File55
[student@linux gen]$ ls File[a5][5b]
File55 Fileab
[student@linux gen]$ ls File[a5][5b][abcdefghijklm]
ls: File[a5][5b][abcdefghijklm]: No such file or directory
[student@linux gen]$ ls file[a5][5b][abcdefghijklm]
fileabc
[student@linux gen]$
```

You can also exclude characters from a list between square brackets with the exclamation mark !. And you are allowed to make combinations of these wild cards.

```
[student@linux gen]$ ls
file1 file2 file3 File4 File55 FileA fileab Fileab FileAB fileabc
[student@linux gen]$ ls file[a5][!Z]
fileab
[student@linux gen]$ ls file[!5]*
file1 file2 file3 fileab fileabc
[student@linux gen]$ ls file[!5]?
fileab
[student@linux gen]$
```

17.4. a-z and 0-9 ranges

The bash shell will also understand ranges of characters between brackets.

```
[student@linux gen]$ ls
file1 file3 File55 fileab FileAB fileabc
file2 File4 FileA Fileab fileab2
[student@linux gen]$ ls file[a-z]*
fileab fileab2 fileabc
[student@linux gen]$ ls file[0-9]
file1 file2 file3
[student@linux gen]$ ls file[a-z][a-z][0-9]*
fileab2
[student@linux gen]$
```

17.5. \$LANG and square brackets

But, don't forget the influence of the LANG variable. Some languages include lower case letters in an upper case range (and vice versa).

```
student@linux:~/test$ ls [A-Z]ile?
file1 file2 file3 File4
student@linux:~/test$ ls [a-z]ile?
file1 file2 file3 File4
student@linux:~/test$ echo $LANG
en_US.UTF-8
student@linux:~/test$ LANG=C
student@linux:~/test$ echo $LANG
C
student@linux:~/test$ ls [a-z]ile?
file1 file2 file3
student@linux:~/test$ ls [A-Z]ile?
File4
student@linux:~/test$
```

If \$LC_ALL is set, then this will also need to be reset to prevent file globbing.

17.6. preventing file globbing

The screenshot below should be no surprise. The echo * will echo a * when in an empty directory. And it will echo the names of all files when the directory is not empty.

```
student@linux:~$ mkdir test42
student@linux:~$ cd test42
student@linux:~/test42$ echo *
*
student@linux:~/test42$ touch file42 file33
student@linux:~/test42$ echo *
file33 file42
```

Globbering can be prevented using quotes or by escaping the special characters, as shown in this screenshot.

```
student@linux:~/test42$ echo *
file33 file42
student@linux:~/test42$ echo \*
*
student@linux:~/test42$ echo '*'
*
student@linux:~/test42$ echo "*"
*
```

17.7. practice: shell globbing

1. Create a test directory and enter it.
2. Create the following files :

17. file globbing

```
file1
file10
file11
file2
File2
File3
file33
fileAB
filea
fileA
fileAAA
file(
file 2
```

(the last one has 6 characters including a space)

3. List (with ls) all files starting with file
4. List (with ls) all files starting with File
5. List (with ls) all files starting with file and ending in a number.
6. List (with ls) all files starting with file and ending with a letter
7. List (with ls) all files starting with File and having a digit as fifth character.
8. List (with ls) all files starting with File and having a digit as fifth character and nothing else.
9. List (with ls) all files starting with a letter and ending in a number.
10. List (with ls) all files that have exactly five characters.
11. List (with ls) all files that start with f or F and end with 3 or A.
12. List (with ls) all files that start with f have i or R as second character and end in a number.
13. List all files that do not start with the letter F.
14. Copy the value of \$LANG to \$MyLANG.
15. Show the influence of \$LANG in listing A-Z or a-z ranges.
16. You receive information that one of your servers was cracked, the cracker probably replaced the ls command. You know that the echo command is safe to use. Can echo replace ls ? How can you list the files in the current directory with echo ?
17. Is there another command besides cd to change directories ?

17.8. solution: shell globbing

1. Create a test directory and enter it.

```
mkdir testdir; cd testdir
```

2. Create the following files :

```

file1
file10
file11
file2
File2
File3
file33
fileAB
filea
fileA
fileAAA
file(
file 2

```

(the last one has 6 characters including a space)

```

touch file1 file10 file11 file2 File2 File3
touch file33 fileAB filea fileA fileAAA
touch "file("
touch "file 2"

```

3. List (with ls) all files starting with file

```
ls file*
```

4. List (with ls) all files starting with File

```
ls File*
```

5. List (with ls) all files starting with file and ending in a number.

```
ls file*[0-9]
```

6. List (with ls) all files starting with file and ending with a letter

```
ls file*[a-z]
```

7. List (with ls) all files starting with File and having a digit as fifth character.

```
ls File[0-9]*
```

8. List (with ls) all files starting with File and having a digit as fifth character and nothing else.

```
ls File[0-9]
```

9. List (with ls) all files starting with a letter and ending in a number.

```
ls [a-z]*[0-9]
```

10. List (with ls) all files that have exactly five characters.

```
ls ?????
```

17. *file globbing*

11. List (with `ls`) all files that start with `f` or `F` and end with `3` or `A`.

```
ls [fF]*[3A]
```

12. List (with `ls`) all files that start with `f` have `i` or `R` as second character and end in a number.

```
ls f[iR]*[0-9]
```

13. List all files that do not start with the letter `F`.

```
ls [!F]*
```

14. Copy the value of `$LANG` to `$MyLANG`.

```
MyLANG=$LANG
```

15. Show the influence of `$LANG` in listing `A-Z` or `a-z` ranges.

see example in book

16. You receive information that one of your servers was cracked, the cracker probably replaced the `ls` command. You know that the `echo` command is safe to use. Can `echo` replace `ls`? How can you list the files in the current directory with `echo`?

```
echo *
```

17. Is there another command besides `cd` to change directories?

```
pushd popd
```

18. regular expressions

(Written by Paul Cobbaut, <https://github.com/paulcobbaut/>, with contributions by: Alex M. Schapelle, <https://github.com/zero-pythagoras/>)

Regular expressions are a very powerful tool in Linux. They can be used with a variety of programs like `bash`, `vi`, `rename`, `grep`, `sed`, and more.

This chapter introduces you to the basics of regular expressions.

18.1. regex versions

There are three different versions of regular expression syntax:

BRE: Basic Regular Expressions
ERE: Extended Regular Expressions
PRCE: Perl Regular Expressions

Depending on the tool being used, one or more of these syntaxes can be used.

For example the `grep` tool has the `-E` option to force a string to be read as ERE while `-G` forces BRE and `-P` forces PRCE.

Note that `grep` also has `-F` to force the string to be read literally.

The `sed` tool also has options to choose a regex syntax.

Read the manual of the tools you use!

18.2. grep

18.2.1. print lines matching a pattern

`grep` is a popular Linux tool to search for lines that match a certain pattern. Below are some examples of the simplest regular expressions.

This is the contents of the test file. This file contains three lines (or three newline characters).

```
student@linux:~$ cat names
Tania
Laura
Valentina
```

When grepping for a single character, only the lines containing that character are returned.

18. regular expressions

```
student@linux:~$ grep u names
Laura
student@linux:~$ grep e names
Valentina
student@linux:~$ grep i names
Tania
Valentina
```

The pattern matching in this example should be very straightforward; if the given character occurs on a line, then `grep` will return that line.

18.2.2. concatenating characters

Two concatenated characters will have to be concatenated in the same way to have a match.

This example demonstrates that `ia` will match `Tania` but not `Valentina` and `in` will match `Valentina` but not `Tania`.

```
student@linux:~$ grep a names
Tania
Laura
Valentina
student@linux:~$ grep ia names
Tania
student@linux:~$ grep in names
Valentina
student@linux:~$
```

18.2.3. one or the other

PRCE and ERE both use the pipe symbol to signify OR. In this example we `grep` for lines containing the letter `i` or the letter `a`.

```
student@linux:~$ cat list
Tania
Laura
student@linux:~$ grep -E 'i|a' list
Tania
Laura
```

Note that we use the `-E` switch of `grep` to force interpretation of our string as an ERE.

We need to escape the pipe symbol in a BRE to get the same logical OR.

```
student@linux:~$ grep -G 'i|a' list
student@linux:~$ grep -G 'i\|a' list
Tania
Laura
```


18.2.4. one or more

The `*` signifies zero, one or more occurrences of the previous and the `+` signifies one or more of the previous.

```
student@linux:~$ cat list2
ll
lol
lool
loool
student@linux:~$ grep -E 'o*' list2
ll
lol
lool
loool
student@linux:~$ grep -E 'o+' list2
lol
lool
loool
student@linux:~$
```

18.2.5. match the end of a string

For the following examples, we will use this file.

```
student@linux:~$ cat names
Tania
Laura
Valentina
Fleur
Floor
```

The two examples below show how to use the `dollar` character to match the end of a string.

```
student@linux:~$ grep a$ names
Tania
Laura
Valentina
student@linux:~$ grep r$ names
Fleur
Floor
```

18.2.6. match the start of a string

The `caret` character (`^`) will match a string at the start (or the beginning) of a line.

Given the same file as above, here are two examples.

```
student@linux:~$ grep ^Val names
Valentina
student@linux:~$ grep ^F names
Fleur
Floor
```

Both the dollar sign and the little hat are called anchors in a regex.

18.2.7. separating words

Regular expressions use a `\b` sequence to reference a word separator. Take for example this file:

```
student@linux:~$ cat text
The governer is governing.
The winter is over.
Can you get over there?
```

Simply grepping for over will give too many results.

```
student@linux:~$ grep over text
The governer is governing.
The winter is over.
Can you get over there?
```

Surrounding the searched word with spaces is not a good solution (because other characters can be word separators). This screenshot below show how to use `\b` to find only the searched word:

```
student@linux:~$ grep '\bover\b' text
The winter is over.
Can you get over there?
student@linux:~$
```

Note that `grep` also has a `-w` option to `grep` for words.

```
student@linux:~$ cat text
The governer is governing.
The winter is over.
Can you get over there?
student@linux:~$ grep -w over text
The winter is over.
Can you get over there?
student@linux:~$
```

18.2.8. grep features

Sometimes it is easier to combine a simple regex with `grep` options, than it is to write a more complex regex. These options where discussed before:

```
grep -i
grep -v
grep -w
grep -A5
grep -B5
grep -C5
```

18.2.9. preventing shell expansion of a regex

The dollar sign is a special character, both for the regex and also for the shell (remember variables and embedded shells). Therefore it is advised to always quote the regex, this prevents shell expansion.

```
student@linux:~$ grep 'r$' names
Fleur
Floor
```

18.3. rename

18.3.1. the rename command

On Debian Linux the `/usr/bin/rename` command is a link to `/usr/bin/prename` installed by the perl package.

```
student@linux ~ $ dpkg -S $(readlink -f $(which rename))
perl: /usr/bin/prename
```

Red Hat derived systems do not install the same `rename` command, so this section does not describe `rename` on Red Hat (unless you copy the perl script manually).

There is often confusion on the internet about the `rename` command because solutions that work fine in Debian (and Ubuntu, xubuntu, Mint, ...) cannot be used in Red Hat (and CentOS, Fedora, ...).

18.3.2. perl

The `rename` command is actually a perl script that uses perl regular expressions. The complete manual for these can be found by typing `perldoc perlrequick` (after installing `perldoc`).

```
root@linux:~# aptitude install perl-doc
The following NEW packages will be installed:
  perl-doc
0 packages upgraded, 1 newly installed, 0 to remove and 0 not upgraded.
Need to get 8,170 kB of archives. After unpacking 13.2 MB will be used.
Get: 1 http://mirrordirector.raspbian.org/raspbian/ wheezy/main perl-do ...
Fetched 8,170 kB in 19s (412 kB/s)
Selecting previously unselected package perl-doc.
(Reading database ... 67121 files and directories currently installed.)
Unpacking perl-doc (from .../perl-doc_5.14.2-21+rpi2_all.deb) ...
Adding 'diversion of /usr/bin/perldoc to /usr/bin/perldoc.stub by perl-doc'
Processing triggers for man-db ...
Setting up perl-doc (5.14.2-21+rpi2) ...

root@linux:~# perldoc perlrequick
```

18.3.3. well known syntax

The most common use of the `rename` is to search for filenames matching a certain string and replacing this string with an other string.

This is often presented as `s/string/other string/` as seen in this example:

```
student@linux ~ $ ls
abc      allfiles.TXT  bllfiles.TXT  Scratch      tennis2.TXT
abc.conf backup        cllfiles.TXT  temp.TXT     tennis.TXT
student@linux ~ $ rename 's/TXT/text/' *
student@linux ~ $ ls
abc      allfiles.text  bllfiles.text  Scratch      tennis2.text
abc.conf backup        cllfiles.text  temp.text    tennis.text
```

And here is another example that uses `rename` with the well know syntax to change the extensions of the same files once more:

```
student@linux ~ $ ls
abc      allfiles.text  bllfiles.text  Scratch      tennis2.text
abc.conf backup        cllfiles.text  temp.text    tennis.text
student@linux ~ $ rename 's/text/txt/' *.text
student@linux ~ $ ls
abc      allfiles.txt  bllfiles.txt  Scratch      tennis2.txt
abc.conf backup        cllfiles.txt  temp.txt     tennis.txt
student@linux ~ $
```

These two examples appear to work because the strings we used only exist at the end of the filename. Remember that file extensions have no meaning in the bash shell.

The next example shows what can go wrong with this syntax.

```
student@linux ~ $ touch atxt.txt
student@linux ~ $ rename 's/txt/problem/' atxt.txt
student@linux ~ $ ls
abc      allfiles.txt  backup        cllfiles.txt  temp.txt     tennis.txt
abc.conf aproblem.txt  bllfiles.txt  Scratch       tennis2.txt
student@linux ~ $
```

Only the first occurrence of the searched string is replaced.

18.3.4. a global replace

The syntax used in the previous example can be described as `s/regex/replacement/`. This is simple and straightforward, you enter a `regex` between the first two slashes and a `replacement` string between the last two.

This example expands this syntax only a little, by adding a `modifier`.

```
student@linux ~ $ rename -n 's/TXT/txt/g' aTXT.TXT
aTXT.TXT renamed as atxt.txt
student@linux ~ $
```

The syntax we use now can be described as `s/regex/replacement/g` where `s` signifies `switch` and `g` stands for `global`.

Note that this example used the `-n` switch to show what is being done (instead of actually renaming the file).

18.3.5. case insensitive replace

Another modifier that can be useful is `i`. this example shows how to replace a case insensitive string with another string.

```
student@linux:~/files$ ls
file1.text file2.TEXT file3.txt
student@linux:~/files$ rename 's/.text/.txt/i' *
student@linux:~/files$ ls
file1.txt file2.txt file3.txt
student@linux:~/files$
```

18.3.6. renaming extensions

Command line Linux has no knowledge of MS-DOS like extensions, but many end users and graphical application do use them.

Here is an example on how to use `rename` to only rename the file extension. It uses the dollar sign to mark the ending of the filename.

```
student@linux ~ $ ls *.txt
allfiles.txt bllfiles.txt cllfiles.txt really.txt.txt temp.txt tennis.txt
student@linux ~ $ rename 's/.txt$/.TXT/' *.txt
student@linux ~ $ ls *.TXT
allfiles.TXT bllfiles.TXT cllfiles.TXT really.txt.TXT
temp.TXT tennis.TXT
student@linux ~ $
```

Note that the dollar sign in the regex means at the end. Without the dollar sign this command would fail on the `really.txt.txt` file.

18.4. sed

18.4.1. stream editor

The stream editor or short `sed` uses regex for stream editing.

In this example `sed` is used to replace a string.

```
echo Sunday | sed 's/Sun/Mon/'
Monday
```

The slashes can be replaced by a couple of other characters, which can be handy in some cases to improve readability.

```
echo Sunday | sed 's:Sun:Mon:'
Monday
echo Sunday | sed 's_Sun_Mon_'
Monday
echo Sunday | sed 's|Sun|Mon|'
Monday
```

18.4.2. interactive editor

While sed is meant to be used in a stream, it can also be used interactively on a file.

```
student@linux:~/files$ echo Sunday > today
student@linux:~/files$ cat today
Sunday
student@linux:~/files$ sed -i 's/Sun/Mon/' today
student@linux:~/files$ cat today
Monday
```

18.4.3. simple back referencing

The ampersand character can be used to reference the searched (and found) string. In this example the ampersand is used to double the occurrence of the found string.

```
echo Sunday | sed 's/Sun/&&/'
SunSunday
echo Sunday | sed 's/day/&&/'
Sundayday
```

18.4.4. back referencing

Parentheses (often called round brackets) are used to group sections of the regex so they can later be referenced.

Consider this simple example:

```
student@linux:~$ echo Sunday | sed 's_\(Sun\)_\1ny_'
Sunnyday
student@linux:~$ echo Sunday | sed 's_\(Sun\)_\1ny \1_'
Sunny Sunday
```

18.4.5. a dot for any character

In a regex a simple dot can signify any character.

```
student@linux:~$ echo 2014-04-01 | sed 's/.....-..-../YYYY-MM-DD/'
YYYY-MM-DD
student@linux:~$ echo abcd-ef-gh | sed 's/.....-..-../YYYY-MM-DD/'
YYYY-MM-DD
```

18.4.6. multiple back referencing

When more than one pair of parentheses is used, each of them can be referenced separately by consecutive numbers.

```
student@linux:~$ echo 2014-04-01 | sed 's/\(....\)-(..)-\(..)/\1+\2+\3/'
2014+04+01
student@linux:~$ echo 2014-04-01 | sed 's/\(....\)-(..)-\(..)/\3:\2:\1/'
01:04:2014
```

This feature is called grouping.

18.4.7. white space

The `\s` can refer to white space such as a space or a tab.

This example looks for white spaces (`\s`) globally and replaces them with 1 space.

```
student@linux:~$ echo -e 'today\tis\twarm'
today  is      warm
student@linux:~$ echo -e 'today\tis\twarm' | sed 's_\s_ _g'
today is warm
```

18.4.8. optional occurrence

A question mark signifies that the previous is optional.

The example below searches for three consecutive letter o, but the third o is optional.

```
student@linux:~$ cat list2
ll
lol
lool
loool
student@linux:~$ grep -E 'ooo?' list2
lool
loool
student@linux:~$ cat list2 | sed 's/ooo\?/A/'
ll
lol
lAl
lAl
```

18.4.9. exactly n times

You can demand an exact number of times the oprevious has to occur.

This example wants exactly three o's.

```
student@linux:~$ cat list2
ll
lol
lool
loool
student@linux:~$ grep -E 'o{3}' list2
loool
student@linux:~$ cat list2 | sed 's/o\{3\}/A/'
ll
lol
lool
lAl
student@linux:~$
```

18.4.10. between n and m times

And here we demand exactly from minimum 2 to maximum 3 times.

```
student@linux:~$ cat list2
ll
lol
lool
loool
student@linux:~$ grep -E 'o{2,3}' list2
lool
loool
student@linux:~$ grep 'o\{2,3\}' list2
lool
loool
student@linux:~$ cat list2 | sed 's/o\{2,3\}/A/'
ll
lol
lAl
lAl
student@linux:~$
```

18.5. bash history

The bash shell can also interpret some regular expressions.

This example shows how to manipulate the exclamation mark history feature of the bash shell.

```
student@linux:~$ mkdir hist
student@linux:~$ cd hist/
student@linux:~/hist$ touch file1 file2 file3
student@linux:~/hist$ ls -l file1
-rw-r--r-- 1 paul paul 0 Apr 15 22:07 file1
student@linux:~/hist$ !l
ls -l file1
-rw-r--r-- 1 paul paul 0 Apr 15 22:07 file1
student@linux:~/hist$ !l:s/1/3
ls -l file3
-rw-r--r-- 1 paul paul 0 Apr 15 22:07 file3
student@linux:~/hist$
```

This also works with the history numbers in bash.

```
student@linux:~/hist$ history 6
2089  mkdir hist
2090  cd hist/
2091  touch file1 file2 file3
2092  ls -l file1
2093  ls -l file3
2094  history 6
student@linux:~/hist$ !2092
ls -l file1
-rw-r--r-- 1 paul paul 0 Apr 15 22:07 file1
student@linux:~/hist$ !2092:s/1/2
ls -l file2
```



```
-rw-r--r-- 1 paul paul 0 Apr 15 22:07 file2  
student@linux:~/hist$
```


Part VII.

Scripting 201; job scheduling

19. more scripting

(Written by Paul Cobbaut, <https://github.com/paulcobbaut/>, with contributions by: Alex M. Schapelle, <https://github.com/zero-pythagoras/>)

19.1. eval

`eval` reads arguments as input to the shell (the resulting commands are executed). This allows using the value of a variable as a variable.

```
student@linux:~/test42$ answer=42
student@linux:~/test42$ word=answer
student@linux:~/test42$ eval x=\$$word ; echo $x
42
```

Both in `bash` and `Korn` the arguments can be quoted.

```
kahlan@solexp11$ answer=42
kahlan@solexp11$ word=answer
kahlan@solexp11$ eval "y=\$$word" ; echo $y
42
```

Sometimes the `eval` is needed to have correct parsing of arguments. Consider this example where the `date` command receives one parameter `1 week ago`.

```
student@linux~$ date --date="1 week ago"
Thu Mar  8 21:36:25 CET 2012
```

When we set this command in a variable, then executing that variable fails unless we use `eval`.

```
student@linux~$ lastweek='date --date="1 week ago"'
student@linux~$ $lastweek
date: extra operand `ago"'
Try `date --help' for more information.
student@linux~$ eval $lastweek
Thu Mar  8 21:36:39 CET 2012
```

19.2. (())

The `(())` allows for evaluation of numerical expressions.

19. more scripting

```
student@linux:~/test42$ (( 42 > 33 )) && echo true || echo false
true
student@linux:~/test42$ (( 42 > 1201 )) && echo true || echo false
false
student@linux:~/test42$ var42=42
student@linux:~/test42$ (( 42 = var42 )) && echo true || echo false
true
student@linux:~/test42$ (( 42 = $var42 )) && echo true || echo false
true
student@linux:~/test42$ var42=33
student@linux:~/test42$ (( 42 = var42 )) && echo true || echo false
false
```

19.3. let

The `let` built-in shell function instructs the shell to perform an evaluation of arithmetic expressions. It will return 0 unless the last arithmetic expression evaluates to 0.

```
[student@linux ~]$ let x="3 + 4" ; echo $x
7
[student@linux ~]$ let x="10 + 100/10" ; echo $x
20
[student@linux ~]$ let x="10-2+100/10" ; echo $x
18
[student@linux ~]$ let x="10*2+100/10" ; echo $x
30
```

The shell can also convert between different bases.

```
[student@linux ~]$ let x="0xFF" ; echo $x
255
[student@linux ~]$ let x="0xC0" ; echo $x
192
[student@linux ~]$ let x="0xA8" ; echo $x
168
[student@linux ~]$ let x="8#70" ; echo $x
56
[student@linux ~]$ let x="8#77" ; echo $x
63
[student@linux ~]$ let x="16#c0" ; echo $x
192
```

There is a difference between assigning a variable directly, or using `let` to evaluate the arithmetic expressions (even if it is just assigning a value).

```
kahlan@solexp11$ dec=15 ; oct=017 ; hex=0x0f
kahlan@solexp11$ echo $dec $oct $hex
15 017 0x0f
kahlan@solexp11$ let dec=15 ; let oct=017 ; let hex=0x0f
kahlan@solexp11$ echo $dec $oct $hex
15 15 15
```

19.4. case

You can sometimes simplify nested if statements with a case construct.

```
[student@linux ~]$ ./help
What animal did you see ? lion
You better start running fast!
[student@linux ~]$ ./help
What animal did you see ? dog
Don't worry, give it a cookie.
[student@linux ~]$ cat help
#!/bin/bash
#
# Wild Animals Helpdesk Advice
#
echo -n "What animal did you see ? "
read animal
case $animal in
    "lion" | "tiger")
        echo "You better start running fast!"
        ;;
    "cat")
        echo "Let that mouse go... "
        ;;
    "dog")
        echo "Don't worry, give it a cookie."
        ;;
    "chicken" | "goose" | "duck" )
        echo "Eggs for breakfast!"
        ;;
    "liger")
        echo "Approach and say 'Ah you big fluffy kitty... '."
        ;;
    "babelfish")
        echo "Did it fall out your ear ?"
        ;;
    *)
        echo "You discovered an unknown animal, name it!"
        ;;
esac
[student@linux ~]$
```

19.5. shell functions

Shell functions can be used to group commands in a logical way.

```
kahlan@solexp11$ cat funcs.ksh
#!/bin/ksh

function greetings {
echo Hello World!
echo and hello to $USER to!
}
```

19. more scripting

```
echo We will now call a function
greetings
echo The end
```

This is sample output from this script with a function.

```
kahlan@solexp11$ ./funcs.ksh
We will now call a function
Hello World!
and hello to kahlan to!
The end
```

A shell function can also receive parameters.

```
kahlan@solexp11$ cat addfunc.ksh
#!/bin/ksh

function plus {
let result="$1 + $2"
echo $1 + $2 = $result
}

plus 3 10
plus 20 13
plus 20 22
```

This script produces the following output.

```
kahlan@solexp11$ ./addfunc.ksh
3 + 10 = 13
20 + 13 = 33
20 + 22 = 42
```

19.6. practice : more scripting

1. Write a script that asks for two numbers, and outputs the sum and product (as shown here).

```
Enter a number: 5
Enter another number: 2
```

```
Sum:          5 + 2 = 7
Product:      5 x 2 = 10
```

2. Improve the previous script to test that the numbers are between 1 and 100, exit with an error if necessary.

3. Improve the previous script to congratulate the user if the sum equals the product.

4. Write a script with a case insensitive case statement, using the `shopt nocasematch` option. The `nocasematch` option is reset to the value it had before the scripts started.

5. If time permits (or if you are waiting for other students to finish this practice), take a look at Linux system scripts in `/etc/init.d` and `/etc/rc.d` and try to understand them. Where does execution of a script start in `/etc/init.d/samba`? There are also some hidden scripts in `~`, we will discuss them later.

19.7. solution : more scripting

1. Write a script that asks for two numbers, and outputs the sum and product (as shown here).

```
Enter a number: 5
Enter another number: 2
```

```
Sum:      5 + 2 = 7
Product:  5 x 2 = 10
```

```
#!/bin/bash

echo -n "Enter a number : "
read n1

echo -n "Enter another number : "
read n2

let sum="$n1+$n2"
let pro="$n1*$n2"

echo -e "Sum\t: $n1 + $n2 = $sum"
echo -e "Product\t: $n1 * $n2 = $pro"
```

2. Improve the previous script to test that the numbers are between 1 and 100, exit with an error if necessary.

```
echo -n "Enter a number between 1 and 100 : "
read n1

if [ $n1 -lt 1 -o $n1 -gt 100 ]
then
    echo Wrong number ...
    exit 1
fi
```

3. Improve the previous script to congratulate the user if the sum equals the product.

```
if [ $sum -eq $pro ]
then echo Congratulations $sum == $pro
fi
```

4. Write a script with a case insensitive case statement, using the shopt nocasematch option. The nocasematch option is reset to the value it had before the scripts started.

```
#!/bin/bash
#
# Wild Animals Case Insensitive Helpdesk Advice
#

if shopt -q nocasematch; then
    nocase=yes;
else
    nocase=no;
```

19. more scripting

```
    shopt -s nocasematch;
fi

echo -n "What animal did you see ? "
read animal

case $animal in
    "lion" | "tiger")
        echo "You better start running fast!"
        ;;
    "cat")
        echo "Let that mouse go... "
        ;;
    "dog")
        echo "Don't worry, give it a cookie."
        ;;
    "chicken" | "goose" | "duck" )
        echo "Eggs for breakfast!"
        ;;
    "liger")
        echo "Approach and say 'Ah you big fluffy kitty.'"
        ;;
    "babelfish")
        echo "Did it fall out your ear ?"
        ;;
    *)
        echo "You discovered an unknown animal, name it!"
        ;;
esac

if [ nocase = yes ] ; then
    shopt -s nocasematch;
else
    shopt -u nocasematch;
fi
```

5. If time permits (or if you are waiting for other students to finish this practice), take a look at Linux system scripts in `/etc/init.d` and `/etc/rc.d` and try to understand them. Where does execution of a script start in `/etc/init.d/samba` ? There are also some hidden scripts in `~`, we will discuss them later.

20. background jobs

(Written by Paul Cobbaut, <https://github.com/paulcobbaut/>, with contributions by: Alex M. Schapelle, <https://github.com/zero-pythagoras/>)

20.1. background processes

20.1.1. jobs

Stuff that runs in background of your current shell can be displayed with the `jobs` command. By default you will not have any jobs running in background.

```
root@linux ~# jobs
root@linux ~#
```

This `jobs` command will be used several times in this section.

20.1.2. control-Z

Some processes can be suspended with the `Ctrl-Z` key combination. This sends a `SIGSTOP` signal to the Linux kernel, effectively freezing the operation of the process.

When doing this in `vi(m)`, then `vi(m)` goes to the background. The background `vi(m)` can be seen with the `jobs` command.

```
[student@linux ~]$ vi procdemo.txt
[5]+ Stopped                  vim procdemo.txt
[student@linux ~]$ jobs
[5]+ Stopped                  vim procdemo.txt
```

20.1.3. & ampersand

Processes that are started in background using the `&` character at the end of the command line are also visible with the `jobs` command.

```
[student@linux ~]$ find / > allfiles.txt 2> /dev/null &
[6] 5230
[student@linux ~]$ jobs
[5]+ Stopped                  vim procdemo.txt
[6]- Running                  find / >allfiles.txt 2>/dev/null &
[student@linux ~]$
```

20. background jobs

20.1.4. jobs -p

An interesting option is `jobs -p` to see the process id of background processes.

```
[student@linux ~]$ sleep 500 &
[1] 4902
[student@linux ~]$ sleep 400 &
[2] 4903
[student@linux ~]$ jobs -p
4902
4903
[student@linux ~]$ ps `jobs -p`
  PID TTY          STAT       TIME COMMAND
  4902 pts/0    S           0:00 sleep 500
  4903 pts/0    S           0:00 sleep 400
[student@linux ~]$
```

20.1.5. fg

Running the `fg` command will bring a background job to the foreground. The number of the background job to bring forward is the parameter of `fg`.

```
[student@linux ~]$ jobs
[1]  Running                sleep 1000 &
[2]-  Running                sleep 1000 &
[3]+  Running                sleep 2000 &
[student@linux ~]$ fg 3
sleep 2000
```

20.1.6. bg

Jobs that are suspended in background can be started in background with `bg`. The `bg` will send a `SIGCONT` signal.

Below an example of the `sleep` command (suspended with `Ctrl-Z`) being reactivated in background with `bg`.

```
[student@linux ~]$ jobs
[student@linux ~]$ sleep 5000 &
[1] 6702
[student@linux ~]$ sleep 3000

[2]+  Stopped                sleep 3000
[student@linux ~]$ jobs
[1]-  Running                sleep 5000 &
[2]+  Stopped                sleep 3000
[student@linux ~]$ bg 2
[2]+ sleep 3000 &
[student@linux ~]$ jobs
[1]-  Running                sleep 5000 &
[2]+  Running                sleep 3000 &
[student@linux ~]$
```

20.2. practice : background processes

1. Use the `jobs` command to verify whether you have any processes running in background.
2. Use `vi` to create a little text file. Suspend `vi` in background.
3. Verify with `jobs` that `vi` is suspended in background.
4. Start `find / > allfiles.txt 2>/dev/null` in foreground. Suspend it in background before it finishes.
5. Start two long `sleep` processes in background.
6. Display all jobs in background.
7. Use the `kill` command to suspend the last `sleep` process.
8. Continue the `find` process in background (make sure it runs again).
9. Put one of the `sleep` commands back in foreground.
10. (if time permits, a general review question...) Explain in detail where the numbers come from in the next screenshot. When are the variables replaced by their value ? By which shell ?

```
[student@linux ~]$ echo $$ $PPID
4224 4223
[student@linux ~]$ bash -c "echo $$ $PPID"
4224 4223
[student@linux ~]$ bash -c 'echo $$ $PPID'
5059 4224
[student@linux ~]$ bash -c `echo $$ $PPID`
4223: 4224: command not found
```

20.3. solution : background processes

1. Use the `jobs` command to verify whether you have any processes running in background.

```
jobs (maybe the catfun is still running?)
```

2. Use `vi` to create a little text file. Suspend `vi` in background.

```
vi text.txt
(inside vi press ctrl-z)
```

3. Verify with `jobs` that `vi` is suspended in background.

```
[student@linux ~]$ jobs
[1]+  Stopped                  vim text.txt
```

4. Start `find / > allfiles.txt 2>/dev/null` in foreground. Suspend it in background before it finishes.

```
[student@linux ~]$ find / > allfiles.txt 2>/dev/null
      (press ctrl-z)
[2]+  Stopped                  find / > allfiles.txt 2> /dev/null
```

20. background jobs

5. Start two long sleep processes in background.

```
sleep 4000 & ; sleep 5000 &
```

6. Display all jobs in background.

```
[student@linux ~]$ jobs
[1]-  Stopped                  vim text.txt
[2]+  Stopped                  find / > allfiles.txt 2> /dev/null
[3]   Running                  sleep 4000 &
[4]   Running                  sleep 5000 &
```

7. Use the kill command to suspend the last sleep process.

```
[student@linux ~]$ kill -SIGSTOP 4519
[student@linux ~]$ jobs
[1]   Stopped                  vim text.txt
[2]-  Stopped                  find / > allfiles.txt 2> /dev/null
[3]   Running                  sleep 4000 &
[4]+  Stopped                  sleep 5000
```

8. Continue the find process in background (make sure it runs again).

```
bg 2 (verify the job-id in your jobs list)
```

9. Put one of the sleep commands back in foreground.

```
fg 3 (again verify your job-id)
```

10. (if time permits, a general review question...) Explain in detail where the numbers come from in the next screenshot. When are the variables replaced by their value ? By which shell ?

```
[student@linux ~]$ echo $$ $PPID
4224 4223
[student@linux ~]$ bash -c "echo $$ $PPID"
4224 4223
[student@linux ~]$ bash -c 'echo $$ $PPID'
5059 4224
[student@linux ~]$ bash -c `echo $$ $PPID`
4223: 4224: command not found
```

The current bash shell will replace the \$\$ and \$PPID while scanning the line, and before executing the echo command.

```
[student@linux ~]$ echo $$ $PPID
4224 4223
```

The variables are now double quoted, but the current bash shell will replace \$\$ and \$PPID while scanning the line, and before executing the bash -c command.

```
[student@linux ~]$ bash -c "echo $$ $PPID"  
4224 4223
```

The variables are now single quoted. The current bash shell will not replace the \$\$ and the \$PPID. The bash -c command will be executed before the variables replaced with their value. This latter bash is the one replacing the \$\$ and \$PPID with their value.

```
[student@linux ~]$ bash -c 'echo $$ $PPID'  
5059 4224
```

With backticks the shell will still replace both variable before the embedded echo is executed. The result of this echo is the two process id's. These are given as commands to bash -c. But two numbers are not commands!

```
[student@linux ~]$ bash -c `echo $$ $PPID`  
4223: 4224: command not found
```


21. scheduling

(Written by Paul Cobbaut, <https://github.com/paulcobbaut/>)

Linux administrators use the `at` to schedule one time jobs. Recurring jobs are better scheduled with `cron`. The next two sections will discuss both tools.

21.1. one time jobs with `at`

21.1.1. `at`

Simple scheduling can be done with the `at` command. This screenshot shows the scheduling of the `date` command at 22:01 and the `sleep` command at 22:03.

```
root@linux:~# at 22:01
at> date
at> <EOT>
job 1 at Wed Aug  1 22:01:00 2007
root@linux:~# at 22:03
at> sleep 10
at> <EOT>
job 2 at Wed Aug  1 22:03:00 2007
root@linux:~#
```

In real life you will hopefully be scheduling more useful commands ;-)

21.1.2. `atq`

It is easy to check when jobs are scheduled with the `atq` or `at -l` commands.

```
root@linux:~# atq
1      Wed Aug  1 22:01:00 2007 a root
2      Wed Aug  1 22:03:00 2007 a root
root@linux:~# at -l
1      Wed Aug  1 22:01:00 2007 a root
2      Wed Aug  1 22:03:00 2007 a root
root@linux:~#
```

The `at` command understands English words like `tomorrow` and `teatime` to schedule commands the next day and at four in the afternoon.

```
root@linux:~# at 10:05 tomorrow
at> sleep 100
at> <EOT>
job 5 at Thu Aug  2 10:05:00 2007
root@linux:~# at teatime tomorrow
at> tea
```

21. scheduling

```
at> <EOT>
job 6 at Thu Aug  2 16:00:00 2007
root@linux:~# atq
6      Thu Aug  2 16:00:00 2007 a root
5      Thu Aug  2 10:05:00 2007 a root
root@linux:~#
```

21.1.3. atrm

Jobs in the at queue can be removed with `atrm`.

```
root@linux:~# atq
6      Thu Aug  2 16:00:00 2007 a root
5      Thu Aug  2 10:05:00 2007 a root
root@linux:~# atrm 5
root@linux:~# atq
6      Thu Aug  2 16:00:00 2007 a root
root@linux:~#
```

21.1.4. at.allow and at.deny

You can also use the `/etc/at.allow` and `/etc/at.deny` files to manage who can schedule jobs with `at`.

The `/etc/at.allow` file can contain a list of users that are allowed to schedule `at` jobs. When `/etc/at.allow` does not exist, then everyone can use `at` unless their username is listed in `/etc/at.deny`.

If none of these files exist, then everyone can use `at`.

21.2. cron

21.2.1. crontab file

The `crontab(1)` command can be used to maintain the `crontab(5)` file. Each user can have their own crontab file to schedule jobs at a specific time. This time can be specified with five fields in this order: minute, hour, day of the month, month and day of the week. If a field contains an asterisk (*), then this means all values of that field.

The following example means: run `script42` eight minutes after two, every day of the month, every month and every day of the week.

```
8 14 * * * script42
```

Run `script8472` every month on the first of the month at 25 past midnight.

```
25 0 1 * * script8472
```

Run this `script33` every two minutes on Sunday (both 0 and 7 refer to Sunday).

```
*/2 * * * 0
```

Instead of these five fields, you can also type one of these: `@reboot`, `@yearly` or `@annually`, `@monthly`, `@weekly`, `@daily` or `@midnight`, and `@hourly`.

21.2.2. crontab command

Users should not edit the crontab file directly, instead they should type `crontab -e` which will use the editor defined in the `EDITOR` or `VISUAL` environment variable. Users can display their cron table with `crontab -l`.

21.2.3. cron.allow and cron.deny

The cron daemon `crond` is reading the cron tables, taking into account the `/etc/cron.allow` and `/etc/cron.deny` files.

These files work in the same way as `at.allow` and `at.deny`. When the `cron.allow` file exists, then your username has to be in it, otherwise you cannot use cron. When the `cron.allow` file does not exist, then your username cannot be in the `cron.deny` file if you want to use cron.

21.2.4. /etc/crontab

The `/etc/crontab` file contains entries for when to run hourly/daily/weekly/monthly tasks. It will look similar to this output.

```
SHELL=/bin/sh
PATH=/usr/local/sbin:/usr/local/bin:/sbin:/bin:/usr/sbin:/usr/bin

20 3 * * *      root    run-parts --report /etc/cron.daily
40 3 * * 7      root    run-parts --report /etc/cron.weekly
55 3 1 * *      root    run-parts --report /etc/cron.monthly
```

21.2.5. /etc/cron.*

The directories shown in the next screenshot contain the tasks that are run at the times scheduled in `/etc/crontab`. The `/etc/cron.d` directory is for special cases, to schedule jobs that require finer control than hourly/daily/weekly/monthly.

```
student@linux:~$ ls -ld /etc/cron.*
drwxr-xr-x 2 root root 4096 2008-04-11 09:14 /etc/cron.d
drwxr-xr-x 2 root root 4096 2008-04-19 15:04 /etc/cron.daily
drwxr-xr-x 2 root root 4096 2008-04-11 09:14 /etc/cron.hourly
drwxr-xr-x 2 root root 4096 2008-04-11 09:14 /etc/cron.monthly
drwxr-xr-x 2 root root 4096 2008-04-11 09:14 /etc/cron.weekly
```

21.2.6. /etc/cron.*

Note that Red Hat uses `anacron` to schedule daily, weekly and monthly cron jobs.

```
root@linux:/etc# cat anacrontab
# /etc/anacrontab: configuration file for anacron

# See anacron(8) and anacrontab(5) for details.

SHELL=/bin/sh
PATH=/sbin:/bin:/usr/sbin:/usr/bin
MAILTO=root
```

21. scheduling

```
# the maximal random delay added to the base delay of the jobs
RANDOM_DELAY=45
# the jobs will be started during the following hours only
START_HOURS_RANGE=3-22

#period in days  delay in minutes  job-identifier  command
1      5      cron.daily      nice run-parts /etc/cron.daily
7     25     cron.weekly     nice run-parts /etc/cron.weekly
@monthly 45     cron.monthly    nice run-parts /etc/cron.monthly
root@linux:/etc#
```

21.3. practice : scheduling

1. Schedule two jobs with at, display the at queue and remove a job.
2. As normal user, use crontab -e to schedule a script to run every four minutes.
3. As root, display the crontab file of your normal user.
4. As the normal user again, remove your crontab file.
5. Take a look at the cron files and directories in /etc and understand them. What is the run-parts command doing ?

21.4. solution : scheduling

1. Schedule two jobs with at, display the at queue and remove a job.

```
root@linux ~# at 9pm today
at> echo go to bed >> /root/todo.txt
at> <EOT>
job 1 at 2010-11-14 21:00
root@linux ~# at 17h31 today
at> echo go to lunch >> /root/todo.txt
at> <EOT>
job 2 at 2010-11-14 17:31
root@linux ~# atq
2 2010-11-14 17:31 a root
1 2010-11-14 21:00 a root
root@linux ~# atrm 1
root@linux ~# atq
2 2010-11-14 17:31 a root
root@linux ~# date
Sun Nov 14 17:31:01 CET 2010
root@linux ~# cat /root/todo.txt
go to lunch
```

2. As normal user, use crontab -e to schedule a script to run every four minutes.

```
student@linux ~$ crontab -e
no crontab for paul - using an empty one
crontab: installing new crontab
```

3. As root, display the crontab file of your normal user.

```
root@linux ~# crontab -l -u paul
*/4 * * * * echo `date` >> /home/paul/crontest.txt
```

4. As the normal user again, remove your crontab file.

```
student@linux ~$ crontab -r
student@linux ~$ crontab -l
no crontab for paul
```

5. Take a look at the cron files and directories in /etc and understand them. What is the run-parts command doing ?

run-parts runs a script in a directory

Part VIII.
SSH; Docker

22. ssh client and server

(Written by Paul Cobbaut, <https://github.com/paulcobbaut/>, with contributions by: Alex M. Schapelle, <https://github.com/zero-pythagoras/>)

The secure shell or ssh is a collection of tools using a secure protocol for communications with remote Linux computers.

This chapter gives an overview of the most common commands related to the use of the sshd server and the ssh client.

22.1. about ssh

22.1.1. secure shell

Avoid using telnet, rlogin and rsh to remotely connect to your servers. These older protocols do not encrypt the login session, which means your user id and password can be sniffed by tools like wireshark or tcpdump. To securely connect to your servers, use ssh.

The ssh protocol is secure in two ways. Firstly the connection is encrypted and secondly the connection is authenticated both ways.

An ssh connection always starts with a cryptographic handshake, followed by encryption of the transport layer using a symmetric cypher. In other words, the tunnel is encrypted before you start typing anything.

Then authentication takes place (using user id/password or public/private keys) and communication can begin over the encrypted connection.

The ssh protocol will remember the servers it connected to (and warn you in case something suspicious happened).

The openssh package is maintained by the OpenBSD people and is distributed with a lot of operating systems (it may even be the most popular package in the world).

22.1.2. /etc/ssh/

Configuration of ssh client and server is done in the /etc/ssh directory. In the next sections we will discuss most of the files found in /etc/ssh/.

22.1.3. ssh protocol versions

The ssh protocol has two versions (1 and 2). Avoid using version 1 anywhere, since it contains some known vulnerabilities. You can control the protocol version via /etc/ssh/ssh_config for the client side and /etc/ssh/sshd_config for the openssh-server daemon.

```
student@linux:/etc/ssh$ grep Protocol ssh_config
# Protocol 2,1
student@linux:/etc/ssh$ grep Protocol sshd_config
Protocol 2
```

22.1.4. public and private keys

The ssh protocol uses the well known system of `public` and `private` keys. The below explanation is succinct, more information can be found on wikipedia.

http://en.wikipedia.org/wiki/Public-key_cryptography

Imagine Alice and Bob, two people that like to communicate with each other. Using `public` and `private` keys they can communicate with encryption and with authentication.

When Alice wants to send an encrypted message to Bob, she uses the `public` key of Bob. Bob shares his `public` key with Alice, but keeps his `private` key private! Since Bob is the only one to have Bob's `private` key, Alice is sure that Bob is the only one that can read the encrypted message.

When Bob wants to verify that the message came from Alice, Bob uses the `public` key of Alice to verify that Alice signed the message with her `private` key. Since Alice is the only one to have Alice's `private` key, Bob is sure the message came from Alice.

22.1.5. rsa and dsa algorithms

This chapter does not explain the technical implementation of cryptographic algorithms, it only explains how to use the ssh tools with `rsa` and `dsa`. More information about these algorithms can be found here:

[http://en.wikipedia.org/wiki/RSA_\(algorithm\)](http://en.wikipedia.org/wiki/RSA_(algorithm))
http://en.wikipedia.org/wiki/Digital_Signature_Algorithm

22.2. log on to a remote server

The following screenshot shows how to use ssh to log on to a remote computer running Linux. The local user is named `paul` and he is logging on as user `admin42` on the remote system.

```
student@linux:~$ ssh admin42@192.168.1.30
The authenticity of host '192.168.1.30 (192.168.1.30)' can't be established.
RSA key fingerprint is b5:fb:3c:53:50:b4:ab:81:f3:cd:2e:bb:ba:44:d3:75.
Are you sure you want to continue connecting (yes/no)?
```

As you can see, the user `paul` is presented with an `rsa` authentication fingerprint from the remote system. The user can accept this by typing `yes`. We will see later that an entry will be added to the `~/.ssh/known_hosts` file.

```
student@linux:~$ ssh admin42@192.168.1.30
The authenticity of host '192.168.1.30 (192.168.1.30)' can't be established.
RSA key fingerprint is b5:fb:3c:53:50:b4:ab:81:f3:cd:2e:bb:ba:44:d3:75.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '192.168.1.30' (RSA) to the list of known hosts.
admin42@192.168.1.30's password:
Welcome to Ubuntu 12.04 LTS (GNU/Linux 3.2.0-26-generic-pae i686)
```

```
* Documentation:  https://help.ubuntu.com/
```

```
1 package can be updated.
```

```
0 updates are security updates.
```

```
Last login: Wed Jun 6 19:25:57 2012 from 172.28.0.131
admin42@ubuserver:~$
```

The user can get log out of the remote server by typing `exit` or by using `Ctrl-d`.

```
admin42@ubuserver:~$ exit
logout
Connection to 192.168.1.30 closed.
student@linux:~$
```

22.3. executing a command in remote

This screenshot shows how to execute the `pwd` command on the remote server. There is no need to exit the server manually.

```
student@linux:~$ ssh admin42@192.168.1.30 pwd
admin42@192.168.1.30's password:
/home/admin42
student@linux:~$
```

22.4. scp

The `scp` command works just like `cp`, but allows the source and destination of the copy to be behind `ssh`. Here is an example where we copy the `/etc/hosts` file from the remote server to the home directory of user `paul`.

```
student@linux:~$ scp admin42@192.168.1.30:/etc/hosts /home/paul/serverhosts
admin42@192.168.1.30's password:
hosts                               100% 809      0.8KB/s   00:00
```

Here is an example of the reverse, copying a local file to a remote server.

```
student@linux:~$ scp ~/serverhosts admin42@192.168.1.30:/etc/hosts.new
admin42@192.168.1.30's password:
serverhosts                          100% 809      0.8KB/s   00:00
```

22.5. setting up passwordless ssh

To set up passwordless `ssh` authentication through public/private keys, use `ssh-keygen` to generate a key pair without a passphrase, and then copy your public key to the destination server. Let's do this step by step.

In the example that follows, we will set up `ssh` without password between Alice and Bob. Alice has an account on a Red Hat Enterprise Linux server, Bob is using Ubuntu on his laptop. Bob wants to give Alice access using `ssh` and the public and private key system. This means that even if Bob changes his password on his laptop, Alice will still have access.

22.5.1. ssh-keygen

The example below shows how Alice uses ssh-keygen to generate a key pair. Alice does not enter a passphrase.

```
[alice@linux ~]$ ssh-keygen -t rsa
Generating public/private rsa key pair.
Enter file in which to save the key (/home/alice/.ssh/id_rsa):
Created directory '/home/alice/.ssh'.
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /home/alice/.ssh/id_rsa.
Your public key has been saved in /home/alice/.ssh/id_rsa.pub.
The key fingerprint is:
9b:ac:ac:56:c2:98:e5:d9:18:c4:2a:51:72:bb:45:eb alice@linux
[alice@linux ~]$
```

You can use ssh-keygen -t dsa in the same way.

22.5.2. ~/.ssh

While ssh-keygen generates a public and a private key, it will also create a hidden .ssh directory with proper permissions. If you create the .ssh directory manually, then you need to chmod 700 it! Otherwise ssh will refuse to use the keys (world readable private keys are not secure!).

As you can see, the .ssh directory is secure in Alice's home directory.

```
[alice@linux ~]$ ls -ld .ssh
drwx----- 2 alice alice 4096 May  1 07:38 .ssh
[alice@linux ~]$
```

Bob is using Ubuntu at home. He decides to manually create the .ssh directory, so he needs to manually secure it.

```
bob@linux:~$ mkdir .ssh
bob@linux:~$ ls -ld .ssh
drwxr-xr-x 2 bob bob 4096 2008-05-14 16:53 .ssh
bob@linux:~$ chmod 700 .ssh/
bob@linux:~$
```

22.5.3. id_rsa and id_rsa.pub

The ssh-keygen command generate two keys in .ssh. The public key is named ~/.ssh/id_rsa.pub. The private key is named ~/.ssh/id_rsa.

```
[alice@linux ~]$ ls -l .ssh/
total 16
-rw----- 1 alice alice 1671 May  1 07:38 id_rsa
-rw-r--r-- 1 alice alice  393 May  1 07:38 id_rsa.pub
```

The files will be named id_dsa and id_dsa.pub when using dsa instead of rsa.

22.5.4. copy the public key to the other computer

To copy the public key from Alice's server to Bob's laptop, Alice decides to use scp.

```
[alice@linux .ssh]$ scp id_rsa.pub bob@192.168.48.92:~/.ssh/authorized_keys
bob@192.168.48.92's password:
id_rsa.pub                                100% 393      0.4KB/s   00:00
```

Be careful when copying a second key! Do not overwrite the first key, instead append the key to the same `~/.ssh/authorized_keys` file!

```
cat id_rsa.pub >> ~/.ssh/authorized_keys
```

Alice could also have used `ssh-copy-id` like in this example.

```
ssh-copy-id -i .ssh/id_rsa.pub bob@192.168.48.92
```

22.5.5. authorized_keys

In your `~/.ssh` directory, you can create a file called `authorized_keys`. This file can contain one or more public keys from people you trust. Those trusted people can use their private keys to prove their identity and gain access to your account via ssh (without password). The example shows Bob's `authorized_keys` file containing the public key of Alice.

```
bob@linux:~$ cat .ssh/authorized_keys
ssh-rsa AAAAB3NzaC1yc2EAAAABIwAAAQEApcQ9xzyLzJes1sR+hPyqW2vyzt1D4zTLqk\
MDWBR4mMFuUZD/0583I3Lg/Q+JIq0RSksNzaL/BNLDou1jMpBe2Dmf/u22u4KmqLJBfDhe\
yTmGSBzeNYCYRSMq78CT9l9a+y6x/shucwhaILsy8A2XfJ9VCgkVtu7XlWFDL2cum08/0\
mRFwVrfc/uPsAn5XkkTscL4g21mQbnp9wJC40pGSJXXMuF0k8MgCb5ieSnpKFniAKM+tEo\
/vjDGSi3F/bxu691jscrU0VUdIo0So98HUFef7jKBRikxGAC7I4HLA+/zX730IvRFAb2hv\
tUhn6RHRbtUJUjbsGgiYeFTLdfctQ= alice@linux
```

22.5.6. passwordless ssh

Alice can now use ssh to connect passwordless to Bob's laptop. In combination with ssh's capability to execute commands on the remote host, this can be useful in pipes across different machines.

```
[alice@linux ~]$ ssh bob@192.168.48.92 "ls -l .ssh"
total 4
-rw-r--r-- 1 bob bob 393 2008-05-14 17:03 authorized_keys
[alice@linux ~]$
```

22.6. X forwarding via ssh

Another popular feature of ssh is called X11 forwarding and is implemented with `ssh -X`.

Below an example of X forwarding: user paul logs in as user greet on her computer to start the graphical application mozilla-thunderbird. Although the application will run on the remote computer from greet, it will be displayed on the screen attached locally to paul's computer.

```
student@linux:~/PDF$ ssh -X greet@greet.dyndns.org -p 55555
Warning: Permanently added the RSA host key for IP address \
'81.240.174.161' to the list of known hosts.
Password:
Linux raika 2.6.8-2-686 #1 Tue Aug 16 13:22:48 UTC 2005 i686 GNU/Linux

Last login: Thu Jan 18 12:35:56 2007
greet@raika:~$ ps fax | grep thun
greet@raika:~$ mozilla-thunderbird &
[1] 30336
```

22.7. troubleshooting ssh

Use `ssh -v` to get debug information about the ssh connection attempt.

```
student@linux:~$ ssh -v bert@192.168.1.192
OpenSSH_4.3p2 Debian-8ubuntu1, OpenSSL 0.9.8c 05 Sep 2006
debug1: Reading configuration data /home/paul/.ssh/config
debug1: Reading configuration data /etc/ssh/ssh_config
debug1: Applying options for *
debug1: Connecting to 192.168.1.192 [192.168.1.192] port 22.
debug1: Connection established.
debug1: identity file /home/paul/.ssh/identity type -1
debug1: identity file /home/paul/.ssh/id_rsa type 1
debug1: identity file /home/paul/.ssh/id_dsa type -1
debug1: Remote protocol version 1.99, remote software version OpenSSH_3
debug1: match: OpenSSH_3.9p1 pat OpenSSH_3.*
debug1: Enabling compatibility mode for protocol 2.0
...
```

22.8. sshd

The ssh server is called `sshd` and is provided by the `openssh-server` package.

```
root@linux~# dpkg -l openssh-server | tail -1
ii openssh-server 1:5.9p1-5ubuntu1 secure shell (SSH) server, ...
```

22.9. sshd keys

The public keys used by the sshd server are located in `/etc/ssh` and are world readable. The private keys are only readable by root.

```
root@linux~# ls -l /etc/ssh/ssh_host_*
-rw----- 1 root root  668 Jun  7  2011 /etc/ssh/ssh_host_dsa_key
-rw-r--r-- 1 root root  598 Jun  7  2011 /etc/ssh/ssh_host_dsa_key.pub
-rw----- 1 root root 1679 Jun  7  2011 /etc/ssh/ssh_host_rsa_key
-rw-r--r-- 1 root root  390 Jun  7  2011 /etc/ssh/ssh_host_rsa_key.pub
```

22.10. ssh-agent

When generating keys with `ssh-keygen`, you have the option to enter a passphrase to protect access to the keys. To avoid having to type this passphrase every time, you can add the key to `ssh-agent` using `ssh-add`.

Most Linux distributions will start the `ssh-agent` automatically when you log on.

```
root@linux~# ps -ef | grep ssh-agent
paul      2405  2365  0 08:13 ?        00:00:00 /usr/bin/ssh-agent ...
```

This clipped screenshot shows how to use `ssh-add` to list the keys that are currently added to the `ssh-agent`

```
student@linux:~$ ssh-add -L
ssh-rsa AAAAB3NzaC1yc2EAAAABIwAAAQEAvgI+Vx5UrIsusZPl8da8URHGsxG7yivv3/\
...
wMGqa48Kelwom8TGb4Sgcwpp/V0/ldA5m+BGCw== student@linux
```

22.11. practice: ssh

0. Make sure that you have access to two Linux computers, or work together with a partner for this exercise. For this practice, we will name one of the machines the server.

1. Install `sshd` on the server
2. Verify in the `ssh` configuration files that only protocol version 2 is allowed.
3. Use `ssh` to log on to the server, show your current directory and then exit the server.
4. Use `scp` to copy a file from your computer to the server.
5. Use `scp` to copy a file from the server to your computer.
6. (optional, only works when you have a graphical install of Linux) Install the `xeyes` package on the server and use `ssh` to run `xeyes` on the server, but display it on your client.
7. (optional, same as previous) Create a bookmark in `firefox`, then quit `firefox` on client and server. Use `ssh -X` to run `firefox` on your display, but on your neighbour's computer. Do you see your neighbour's bookmark?
8. Use `ssh-keygen` to create a key pair without passphrase. Setup passwordless `ssh` between you and your neighbour. (or between your client and your server)
9. Verify that the permissions on the server key files are correct; world readable for the public keys and only root access for the private keys.

22. ssh client and server

10. Verify that the ssh-agent is running.

11. (optional) Protect your keypair with a passphrase, then add this key to the ssh-agent and test your passwordless ssh to the server.

22.12. solution: ssh

0. Make sure that you have access to two Linux computers, or work together with a partner for this exercise. For this practice, we will name one of the machines the server.

1. Install sshd on the server

```
apt-get install openssh-server (on Ubuntu/Debian)
yum -y install openssh-server (on Centos/Fedora/Red Hat)
```

2. Verify in the ssh configuration files that only protocol version 2 is allowed.

```
grep Protocol /etc/ssh/ssh*_config
```

3. Use ssh to log on to the server, show your current directory and then exit the server.

```
user@client$ ssh user@server-ip-address
user@server$ pwd
/home/user
user@server$ exit
```

4. Use scp to copy a file from your computer to the server.

```
scp localfile user@server:~
```

5. Use scp to copy a file from the server to your computer.

```
scp user@server:~/serverfile .
```

6. (optional, only works when you have a graphical install of Linux) Install the xeyes package on the server and use ssh to run xeyes on the server, but display it on your client.

```
on the server:
apt-get install xeyes
on the client:
ssh -X user@server-ip
xeyes
```

7. (optional, same as previous) Create a bookmark in firefox, then quit firefox on client and server. Use ssh -X to run firefox on your display, but on your neighbour's computer. Do you see your neighbour's bookmark ?

8. Use ssh-keygen to create a key pair without passphrase. Setup passwordless ssh between you and your neighbour. (or between your client and your server)

See solution in book "setting up passwordless ssh"

9. Verify that the permissions on the server key files are correct; world readable for the public keys and only root access for the private keys.


```
ls -l /etc/ssh/ssh_host_*
```

10. Verify that the ssh-agent is running.

```
ps fax | grep ssh-agent
```

11. (optional) Protect your keypair with a passphrase, then add this key to the ssh-agent and test your passwordless ssh to the server.

```
man ssh-keygen  
man ssh-agent  
man ssh-add
```


A. git

(Written by Paul Cobbaut, <https://github.com/paulcobbaut/>, with contributions by: Alex M. Schapelle, <https://github.com/zero-pythagoras/>)

This chapter is an introduction to using `git` on the command line. The `git` repository is hosted by `github`, but you are free to choose another server (or create your own).

There are many excellent online tutorials for `git`. This list can save you one Google query:

<http://gitimmersion.com/>
<http://git-scm.com/book>


A.1. git

Linus Torvalds created `git` back in 2005 when Bitkeeper changed its license and the Linux kernel developers were no longer able to use it for free.

`git` quickly became popular and is now the most widely used distributed version control system in the world.

Geek and Poke demonstrates why we need version control (image property of Geek and Poke CCA 3.0).

SIMPLY EXPLAINED



VERSION CONTROL

Besides source code for software, you can also find German and Icelandic law on github (and probably much more by the time you are reading this).

A.2. installing git

We install git with aptitude `install git` as seen in this screenshot on Debian 6.

```
root@linux:~# aptitude install git
The following NEW packages will be installed:
  git libcurl3-gnutls{a} liberror-perl{a}
0 packages upgraded, 3 newly installed, 0 to remove and 0 not upgraded.
...
Processing triggers for man-db ...
Setting up libcurl3-gnutls (7.21.0-2.1+squeeze2) ...
Setting up liberror-perl (0.17-1) ...
Setting up git (1:1.7.2.5-3) ...
```

A.3. starting a project

First we create a project directory, with a simple file in it.

```
student@linux~$ mkdir project42
student@linux~$ cd project42/
student@linux~/project42$ echo "echo The answer is 42." >> question.sh
```

A.3.1. git init

Then we tell git to create an empty git repository in this directory.

```
student@linux~/project42$ ls -la
total 12
drwxrwxr-x  2 paul paul 4096 Dec  8 16:41 .
drwxr-xr-x 46 paul paul 4096 Dec  8 16:41 ..
-rw-rw-r--  1 paul paul  23 Dec  8 16:41 question.sh
student@linux~/project42$ git init
Initialized empty Git repository in /home/paul/project42/.git/
student@linux~/project42$ ls -la
total 16
drwxrwxr-x  3 paul paul 4096 Dec  8 16:44 .
drwxr-xr-x 46 paul paul 4096 Dec  8 16:41 ..
drwxrwxr-x  7 paul paul 4096 Dec  8 16:44 .git
-rw-rw-r--  1 paul paul  23 Dec  8 16:41 question.sh
```

A.3.2. git config

Next we use `git config` to set some global options.

```
student@linux$ git config --global user.name Paul
student@linux$ git config --global user.email "paul.cobbaut@gmail.com"
student@linux$ git config --global core.editor vi
```

We can verify this config in `~/.gitconfig`:

```
student@linux~/project42$ cat ~/.gitconfig
[user]
  name = Paul
  email = paul.cobbaut@gmail.com
[core]
  editor = vi
```

A.3.3. git add

Time now to add file to our project with `git add`, and verify that it is added with `git status`.

```
student@linux~/project42$ git add question.sh
student@linux~/project42$ git status
# On branch master
#
# Initial commit
#
# Changes to be committed:
#   (use "git rm --cached <file> ..." to unstage)
#
#   new file:   question.sh
#
```

A. git

The `git status` tells us there is a new file ready to be committed.

A.3.4. git commit

With `git commit` you force git to record all added files (and all changes to those files) permanently.

```
student@linux~/project42$ git commit -m "starting a project"
[master (root-commit) 5c10768] starting a project
 1 file changed, 1 insertion(+)
 create mode 100644 question.sh
student@linux~/project42$ git status
# On branch master
nothing to commit (working directory clean)
```

A.3.5. changing a committed file

The screenshots below show several steps. First we change a file:

```
student@linux~/project42$ git status
# On branch master
nothing to commit (working directory clean)
student@linux~/project42$ vi question.sh
```

Then we verify the status and see that it is modified:

```
student@linux~/project42$ git status
# On branch master
# Changes not staged for commit:
#   (use "git add <file> ..." to update what will be committed)
#   (use "git checkout -- <file> ..." to discard changes in working directory)
#
#   modified:   question.sh
#
no changes added to commit (use "git add" and/or "git commit -a")
```

Next we add it to the git repository.

```
student@linux~/project42$ git add question.sh
student@linux~/project42$ git commit -m "adding a she-bang to the main script"
[master 86b8347] adding a she-bang to the main script
 1 file changed, 1 insertion(+)
student@linux~/project42$ git status
# On branch master
nothing to commit (working directory clean)
```

A.3.6. git log

We can see all our commits again using `git log`.

```
student@linux~/project42$ git log
commit 86b8347192ea025815df7a8e628d99474b41fb6c
Author: Paul <paul.cobbaut@gmail.com>
Date: Sat Dec 8 17:12:24 2012 +0100
```

adding a she-bang to the main script

```
commit 5c10768f29aecc16161fb197765e0f14383f7bca
Author: Paul <paul.cobbaut@gmail.com>
Date: Sat Dec 8 17:09:29 2012 +0100
```

starting a project

The log format can be changed.

```
student@linux~/project42$ git log --pretty=oneline
86b8347192ea025815df7a8e628d99474b41fb6c adding a she-bang to the main script
5c10768f29aecc16161fb197765e0f14383f7bca starting a project
```

The log format can be customized a lot.

```
student@linux~/project42$ git log --pretty=format:"%an: %ar :%s"
Paul: 8 minutes ago :adding a she-bang to the main script
Paul: 11 minutes ago :starting a project
```

A.3.7. git mv

Renaming a file can be done with `mv` followed by a `git remove` and a `git add` of the new filename. But it can be done easier and in one command using `git mv`.

```
student@linux~/project42$ git mv question.sh thequestion.sh
student@linux~/project42$ git status
# On branch master
# Changes to be committed:
#   (use "git reset HEAD <file> ..." to unstage)
#
#   renamed:    question.sh -> thequestion.sh
#
student@linux~/project42$ git commit -m "improved naming scheme"
[master 69b2c8b] improved naming scheme
 1 file changed, 0 insertions(+), 0 deletions(-)
 rename question.sh => thequestion.sh (100%)
```

A.4. git branches

Working on the project can be done in one or more `git` branches. Here we create a new branch that will make changes to the script. We will merge this branch with the `master` branch when we are sure the script works. (It can be useful to add `git status` commands when practicing).

A. git

```
student@linux~/project42$ git branch
* master
student@linux~/project42$ git checkout -b newheader
Switched to a new branch 'newheader'
student@linux~/project42$ vi thequestion.sh
student@linux~/project42$ git add thequestion.sh
student@linux~/project42$ source thequestion.sh
The answer is 42.
```

It seems to work, so we commit in this branch.

```
student@linux~/project42$ git commit -m "adding a new company header"
[newheader 730a22b] adding a new company header
1 file changed, 4 insertions(+)
student@linux~/project42$ git branch
  master
* newheader
student@linux~/project42$ cat thequestion.sh
#!/bin/bash
#
# copyright linux-training.be
#
```

```
echo The answer is 42.
```

Let us go back to the master branch and see what happened there.

```
student@linux~/project42$ git checkout master
Switched to branch 'master'
student@linux~/project42$ cat thequestion.sh
#!/bin/bash
echo The answer is 42.
```

Nothing happened in the master branch, because we worked in another branch.

When we are sure the branch is ready for production, then we merge it into the master branch.

```
student@linux~/project42$ cat thequestion.sh
#!/bin/bash
echo The answer is 42.
student@linux~/project42$ git merge newheader
Updating 69b2c8b..730a22b
Fast-forward
 thequestion.sh | 4 ++++
1 file changed, 4 insertions(+)
student@linux~/project42$ cat thequestion.sh
#!/bin/bash
#
# copyright linux-training.be
#
```

```
echo The answer is 42.
```

The newheader branch can now be deleted.


```
student@linux~/project42$ git branch
* master
  newheader
student@linux~/project42$ git branch -d newheader
Deleted branch newheader (was 730a22b).
student@linux~/project42$ git branch
* master
```

A.5. to be continued...

The git story is not finished.

There are many excellent online tutorials for git. This list can save you one Google query:

```
http://gitimmersion.com/
http://git-scm.com/book
```

A.6. github.com

Create an account on `github.com`. This website is a frontend for an immense git server with over two and a half million users and almost five million projects (including Fedora, Linux kernel, Android, Ruby on Rails, Wine, X.org, VLC...)

```
https://github.com/signup/free
```

This account is free of charge, we will use it in the examples below.

A.7. add your public key to github

I prefer to use github with a `public` key, so it probably is a good idea that you also upload your public key to github.com.

You can upload your own key via the web interface:

```
https://github.com/settings/ssh
```

Please do not forget to protect your `private` key!

A.8. practice: git

1. Create local project called `git_practice`.
2. Create a project on `gitlab.com` to host a local project that you have created.
3. The project should have `README.md` file as well as `TODO.md` file in it.
4. Write in `README.md` file description of the project and what you think it might be.
5. Initialize your project with `git` command, setup your username, mail and remote server.
6. Use `git push -u origin master` to send project saves to remote host.

A. *git*

7. Verify on `gitlab.com` that the project has been setup and is updated with `README.md` and `TODO.md`.

8. Add `git_hello.sh` script that prints hello to username from its current location.

9. Push the script to gitlab repository.

A.9. solution: git

1. Create local project called `git_practice`.

```
aschappelle@vaio3:~$ mkdir git_practice; cd git_practice
```

2. Create a project on `gitlab.com` to host a local project that you have created.

3. The project should have `README.md` file as well as `TODO.md` file in it.

```
aschappelle@vaio3:~/git_practice$ touch README.md TODO.md
```

4. Write in `README.md` file description of the project and what you think it might be.

```
aschappelle@vaio3:~/git_practice$ echo "This is readme file for git_practice project" > README.md
aschappelle@vaio3:~/git_practice$ echo "This is todo file for git_practice project" > TODO.md
```

5. Initialize your project with `git` command, setup your username, mail and remote server.

```
aschappelle@vaio3:~/git_practice$ git init
aschappelle@vaio3:~/git_practice$ git config user.name alex.schappelle
aschappelle@vaio3:~/git_practice$ git config user.mail alex@vaiolabs.com
aschappelle@vaio3:~/git_practice$ git remote add origin https://gitlab.com/url_to_your_repo
```

6. Use `git push -u origin master` to send project saves to remote host.

```
aschappelle@vaio3:~/git_practice$ git push -u origin master
```

7. Verify on `gitlab.com` that the project has been setup and is updated with `README.md` and `TODO.md`.

8. Add `git_hello.sh` script that prints hello to username from its current location.

```
aschappelle@vaio3:~/git_practice$ git push -u origin master
```

9. Push the script to gitlab repository.

```
aschappelle@vaio3:~/git_practice$ git push -u origin master
```

B. Introduction to vi

(Written by Paul Cobbaut, <https://github.com/paulcobbaut/>)

The `vi` editor is installed on almost every Unix. Linux will very often install `vim` (`vi improved`) which is similar. Every system administrator should know `vi(m)`, because it is an easy tool to solve problems.

The `vi` editor is not intuitive, but once you get to know it, `vi` becomes a very powerful application. Most Linux distributions will include the `vimtutor` which is a 45 minute lesson in `vi(m)`.

B.1. command mode and insert mode

The `vi` editor starts in `command mode`. In `command mode`, you can type commands. Some commands will bring you into `insert mode`. In `insert mode`, you can type text. The escape key will return you to `command mode`.

Table B.1.: getting to command mode

key	action
Esc	set <code>vi(m)</code> in <code>command mode</code> .

B.2. start typing (a A i l o O)

The difference between `a A i l o` and `O` is the location where you can start typing. `a` will append after the current character and `A` will append at the end of the line. `i` will insert before the current character and `l` will insert at the beginning of the line. `o` will put you in a new line after the current line and `O` will put you in a new line before the current line.

Table B.2.: switch to insert mode

command	action
<code>a</code>	start typing after the current character
<code>A</code>	start typing at the end of the current line
<code>i</code>	start typing before the current character
<code>l</code>	start typing at the start of the current line
<code>o</code>	start typing on a new line after the current line
<code>O</code>	start typing on a new line before the current line

B.3. replace and delete a character (r x X)

When in command mode (it doesn't hurt to hit the escape key more than once) you can use the x key to delete the current character. The big X key (or shift x) will delete the character left of the cursor. Also when in command mode, you can use the r key to replace one single character. The r key will bring you in insert mode for just one key press, and will return you immediately to command mode.

Table B.3.: replace and delete

command	action
x	delete the character below the cursor
X	delete the character before the cursor
r	replace the character below the cursor
p	paste after the cursor (here the last deleted character)
xp	switch two characters

B.4. undo, redo and repeat (u .)

When in command mode, you can undo your mistakes with u. Use `ctrl-r` to redo the undo.

You can do your mistakes twice with . (in other words, the . will repeat your last command).

Table B.4.: undo and repeat

command	action
u	undo the last action
ctrl-r	redo the last undo
.	repeat the last action

B.5. cut, copy and paste a line (dd yy p P)

When in command mode, dd will cut the current line. yy will copy the current line. You can paste the last copied or cut line after (p) or before (P) the current line.

Table B.5.: cut, copy and paste a line

command	action
dd	cut the current line
yy	(yank yank) copy the current line
p	paste after the current line
P	paste before the current line

B.6. cut, copy and paste lines (3dd 2yy)

When in command mode, before typing dd or yy, you can type a number to repeat the command a number of times. Thus, 5dd will cut 5 lines and 4yy will copy (yank) 4 lines. That last one will be noted by vi in the bottom left corner as "4 line yanked".

Table B.6.: cut, copy and paste lines

command	action
3dd	cut three lines
4yy	copy four lines

B.7. start and end of a line (0 or ^ and \$)

When in command mode, the 0 and the caret ^ will bring you to the start of the current line, whereas the \$ will put the cursor at the end of the current line. You can add 0 and \$ to the d command, d0 will delete every character between the current character and the start of the line. Likewise d\$ will delete everything from the current character till the end of the line. Similarly y0 and y\$ will yank till start and end of the current line.

Table B.7.: start and end of line

command	action
0	jump to start of current line
^	jump to start of current line
\$	jump to end of current line
d0	delete until start of line
d\$	delete until end of line

B.8. join two lines (J) and more

When in command mode, pressing J will append the next line to the current line. With yyp you duplicate a line and with ddp you switch two lines.

Table B.8.: join two lines

command	action
J	join two lines
yyp	duplicate a line
ddp	switch two lines

B.9. words (w b)

When in command mode, w will jump to the next word and b will move to the previous word. w and b can also be combined with d and y to copy and cut words (dw db yw yb).

Table B.9.: words

command	action
w	forward one word
b	back one word
3w	forward three words
dw	delete one word
yw	yank (copy) one word

command	action
5yb	yank five words back
7dw	delete seven words

B.10. save (or not) and exit (:w :q :q!)

Pressing the colon `:` will allow you to give instructions to vi (technically speaking, typing the colon will open the ex editor). `:w` will write (save) the file, `:q` will quit an unchanged file without saving, and `:q!` will quit vi discarding any changes. `:wq` will save and quit and is the same as typing ZZ in command mode.

Table B.10.: save and exit vi

command	action
<code>:w</code>	save (write)
<code>:w fname</code>	save as fname
<code>:q</code>	quit
<code>:wq</code>	save and quit
ZZ	save and quit
<code>:q!</code>	quit (discarding your changes)
<code>:w!</code>	save (and write to non-writable file!)

The last one is a bit special. With `:w!` vi will try to `chmod` the file to get write permission (this works when you are the owner) and will `chmod` it back when the write succeeds. This should always work when you are root (and the file system is writable).

B.11. Searching (/ ?)

When in command mode typing `/` will allow you to search in vi for strings (can be a regular expression). Typing `/foo` will do a forward search for the string `foo` and typing `?bar` will do a backward search for `bar`.

Table B.11.: searching

command	action
<code>/string</code>	forward search for string
<code>?string</code>	backward search for string
<code>n</code>	go to next occurrence of search string
<code>/^string</code>	forward search string at beginning of line
<code>/string\$</code>	forward search string at end of line
<code>/br[aeio]l</code>	search for bral brel bril and brol
<code>^<he\></code>	search for the word he (and not for here or the)

B.12. replace all (:1,\$ s/foo/bar/g)

To replace all occurrences of the string `foo` with `bar`, first switch to ex mode with `:`. Then tell vi which lines to use, for example `1,$` will do the replace all from the first to the last line. You can write `1,5` to only process the first five lines. The `s/foo/bar/g` will replace all occurrences of `foo` with `bar`.

Table B.12.: replace

command	action
:4,8 s/foo/bar/g	replace foo with bar on lines 4 to 8
:1,\$ s/foo/bar/g	replace foo with bar on all lines

B.13. reading files (:r :r !cmd)

When in command mode, :r foo will read the file named foo, :r !foo will execute the command foo. The result will be put at the current location. Thus :r !ls will put a listing of the current directory in your text file.

Table B.13.: read files and input

command	action
:r fname	(read) file fname and paste contents
:r !cmd	execute cmd and paste its output

B.14. text buffers

There are 36 buffers in vi to store text. You can use them with the " character.

Table B.14.: text buffers

command	action
"add	delete current line and put text in buffer a
"g7yy	copy seven lines into buffer g
"ap	paste from buffer a

B.15. multiple files

You can edit multiple files with vi. Here are some tips.

Table B.15.: multiple files

command	action
vi file1 file2 file3	start editing three files
:args	lists files and marks active file
:n	start editing the next file
:e	toggle with last edited file
:rew	rewind file pointer to first file

B.16. abbreviations

With :ab you can put abbreviations in vi. Use :una to undo the abbreviation.

Table B.16.: abbreviations

command	action
:ab str long string	abbreviate str to be 'long string'
:una str	un-abbreviate str

B.17. key mappings

Similarly to their abbreviations, you can use mappings with `:map` for command mode and `:map!` for insert mode.

This example shows how to set the F6 function key to toggle between `set number` and `set nonumber`. The `<bar>` separates the two commands, `set number!` toggles the state and `set number?` reports the current state.

```
:map <F6> :set number!<bar>set number?<CR>
```

B.18. setting options

Some options that you can set in vim.

```
:set number ( also try :se nu )
:set nonumber
:syntax on
:syntax off
:set all (list all options)
:set tabstop=8
:set tx (CR/LF style endings)
:set notx
```

You can set these options (and much more) in `~/.vimrc` for vim or in `~/.exrc` for standard vi.

```
student@linux:~$ cat ~/.vimrc
set number
set tabstop=8
set textwidth=78
map <F6> :set number!<bar>set number?<CR>
student@linux:~$
```

B.19. practice: vi(m)

1. Start the vimtutor and do some or all of the exercises. You might need to run `aptitude install vim` on xubuntu.
2. What 3 key sequence in command mode will duplicate the current line.
3. What 3 key sequence in command mode will switch two lines' place (line five becomes line six and line six becomes line five).

4. What 2 key sequence in command mode will switch a character's place with the next one.
5. vi can understand macro's. A macro can be recorded with q followed by the name of the macro. So qa will record the macro named a. Pressing q again will end the recording. You can recall the macro with @ followed by the name of the macro. Try this example: i l 'Escape Key' qa yyp 'Ctrl a' q 5@a (Ctrl a will increase the number with one).
6. Copy /etc/passwd to your ~/passwd. Open the last one in vi and press Ctrl v. Use the arrow keys to select a Visual Block, you can copy this with y or delete it with d. Try pasting it.
7. What does dwwP do when you are at the beginning of a word in a sentence ?

B.20. solution: vi(m)

1. Start the vintutor and do some or all of the exercises. You might need to run `aptitude install vim` on xubuntu.

vintutor

2. What 3 key sequence in command mode will duplicate the current line.

yyp

3. What 3 key sequence in command mode will switch two lines' place (line five becomes line six and line six becomes line five).

ddp

4. What 2 key sequence in command mode will switch a character's place with the next one.

xp

5. vi can understand macro's. A macro can be recorded with q followed by the name of the macro. So qa will record the macro named a. Pressing q again will end the recording. You can recall the macro with @ followed by the name of the macro. Try this example: i l 'Escape Key' qa yyp 'Ctrl a' q 5@a (Ctrl a will increase the number with one).

6. Copy /etc/passwd to your ~/passwd. Open the last one in vi and press Ctrl v. Use the arrow keys to select a Visual Block, you can copy this with y or delete it with d. Try pasting it.

```
cp /etc/passwd ~
vi passwd
(press Ctrl-V)
```

7. What does dwwP do when you are at the beginning of a word in a sentence ?

dwwP can switch the current word with the next word.

C. GNU Free Documentation License

Version 1.3, 3 November 2008

Copyright © 2000, 2001, 2002, 2007, 2008 Free Software Foundation, Inc.

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

C.1. PREAMBLE

The purpose of this License is to make a manual, textbook, or other functional and useful document “free” in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or noncommercially. Secondly, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of “copyleft”, which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

C.2. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world-wide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The “Document”, below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as “you”. You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A “Modified Version” of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A “Secondary Section” is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document’s overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The “Invariant Sections” are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this

C. GNU Free Documentation License

License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The “Cover Texts” are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License. A Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words.

A “Transparent” copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. A copy that is not “Transparent” is called “Opaque”.

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The “Title Page” means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, “Title Page” means the text near the most prominent appearance of the work’s title, preceding the beginning of the body of the text.

The “publisher” means any person or entity that distributes copies of the Document to the public.

A section “Entitled XYZ” means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as “Acknowledgements”, “Dedications”, “Endorsements”, or “History”.) To “Preserve the Title” of such a section when you modify the Document means that it remains a section “Entitled XYZ” according to this definition.

The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

C.3. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

C.4. COPYING IN QUANTITY

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer-network location from which the general network-using public has access to download using public-standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

C.5. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

- A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.
- B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.
- C. State on the Title page the name of the publisher of the Modified Version, as the publisher.
- D. Preserve all the copyright notices of the Document.
- E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.

C. GNU Free Documentation License

- F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.
- G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
- H. Include an unaltered copy of this License.
- I. Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.
- J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.
- K. For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.
- L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.
- M. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.
- N. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.
- O. Preserve any Warranty Disclaimers.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles.

You may add a section Entitled "Endorsements", provided it contains nothing but endorsements of your Modified Version by various parties—for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

C.6. COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections Entitled “History” in the various original documents, forming one section Entitled “History”; likewise combine any sections Entitled “Acknowledgements”, and any sections Entitled “Dedications”. You must delete all sections Entitled “Endorsements”.

C.7. COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

C.8. AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called an “aggregate” if the copyright resulting from the compilation is not used to limit the legal rights of the compilation’s users beyond what the individual works permit. When the Document is included in an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document’s Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

C.9. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

If a section in the Document is Entitled “Acknowledgements”, “Dedications”, or “History”, the requirement (section 4) to Preserve its Title (section 1) will typically require changing the actual title.

C.10. TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, or distribute it is void, and will automatically terminate your rights under this License.

However, if you cease all violation of this License, then your license from a particular copyright holder is reinstated (a) provisionally, unless and until the copyright holder explicitly and finally terminates your license, and (b) permanently, if the copyright holder fails to notify you of the violation by some reasonable means prior to 60 days after the cessation.

Moreover, your license from a particular copyright holder is reinstated permanently if the copyright holder notifies you of the violation by some reasonable means, this is the first time you have received notice of violation of this License (for any work) from that copyright holder, and you cure the violation prior to 30 days after your receipt of the notice.

Termination of your rights under this section does not terminate the licenses of parties who have received copies or rights from you under this License. If your rights have been terminated and not permanently reinstated, receipt of a copy of some or all of the same material does not give you any rights to use it.

C.11. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See <http://www.gnu.org/copyleft/>.

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License “or any later version” applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation. If the Document specifies that a proxy can decide which future versions of this License can be used, that proxy’s public statement of acceptance of a version permanently authorizes you to choose that version for the Document.

C.12. RELICENSING

“Massive Multiauthor Collaboration Site” (or “MMC Site”) means any World Wide Web server that publishes copyrightable works and also provides prominent facilities for anybody to edit those works. A public wiki that anybody can edit is an example of such a server. A “Massive Multiauthor Collaboration” (or “MMC”) contained in the site means any set of copyrightable works thus published on the MMC site.

“CC-BY-SA” means the Creative Commons Attribution-Share Alike 3.0 license published by Creative Commons Corporation, a not-for-profit corporation with a principal place of business in San Francisco, California, as well as future copyleft versions of that license published by that same organization.

“Incorporate” means to publish or republish a Document, in whole or in part, as part of another Document.

An MMC is “eligible for relicensing” if it is licensed under this License, and if all works that were first published under this License somewhere other than this MMC, and subsequently

incorporated in whole or in part into the MMC, (1) had no cover texts or invariant sections, and (2) were thus incorporated prior to November 1, 2008.

The operator of an MMC Site may republish an MMC contained in the site under CC-BY-SA on the same site at any time before August 1, 2009, provided the MMC is eligible for relicensing.

