

Introduction to Linux

Sumac Skillshare

Contents

Page 1	Contents
.....	Glossary
Page 4	Free Software and the foundations of GNU/Linux
Page 7	The Linux File system
Page 10	Shell Commands by Alphabet
Page 13	Using BASH
Page 16	Ubuntu tips & tricks

Glossary

Boot - When you start (or re-start) your computer, that's called "booting". If you want to get all technical, when you power on your computer, it's a "cold boot". Restarting it (without turning it off first) is a "warm boot".

Command Line Interface (CLI) – See Shell.

Compiz Fusion - Compiz Fusion is a compositing window manager which facilitates graphical effects for the X Window System. It uses 3D graphics hardware to create fast compositing desktop effects for window management. The effects, such as a minimization effect and a cube workspace are implemented as loadable plug-ins.

Copyleft - is a play on the word copyright and is the practice of using copyright law to remove restrictions on distributing copies and modified versions of a work for others and requiring that the same freedoms be preserved in modified versions.

Embedded System - is a special-purpose computer system designed to perform one or a few dedicated function. Embedded systems range from portable devices such as digital watches and MP3 players, to large stationary installations like traffic lights and factory controllers.

Gnome - is an international effort to build a complete desktop environment—the graphical user interface which sits on top of a computer operating system—entirely from free software. GNOME is part of the GNU Project.

Hardware - is the physical part of a computer. It is the stuff you can touch, the processor, the memory, the hard drive etc.

KDE (K Desktop Environment) - is a free software project which aims to be a powerful system for an easy-to-use desktop environment. KDE serves as an umbrella project for many standalone applications and smaller projects that are based on KDE technology, such as KOffice, KDevelop, Amarok, K3b and many more.

Kernel - is the central component of a computer operating system. It controls system resources and the communication between hardware & software.

LAMP - The acronym LAMP refers to a solution stack of software, used to run dynamic Web sites or servers.

The original expansion is as follows:

- * Linux, referring to the operating system;
- * Apache, the Web server;
- * MySQL, the database management system (or database server);
- * PHP and sometimes Perl, the programming languages.

Mount – the act of making a hard drive partition (or other media) available to the Linux filesystem. The point in the filesystem where the new medium is accessible is called the mount point.

Operating System (OS) - is the software that manages the sharing of the resources of a computer. An operating system performs tasks such as controlling and allocating memory, prioritizing system requests, controlling input and output devices, facilitating networking and managing filesystems. The operating system forms the platform for other system software and for application software. Operating systems come with an application that provides a user interface for managing the operating system, such as a command line interpreter (e.g. the Linux shell) or graphical user interface (e.g. Windows device manager & control centre).

Package Management System- is a collection of tools to automate the process of installing, upgrading, configuring, and removing software packages from your computer.

Partition – Partitioning is the dividing of a hard disk's storage space into independent parts. Therefore a single hard disk can have several partitions. In Windows these partitions appear as different drive letters e.g. c: d: etc. In Linux a partition can be mounted at any point in the filesystem.

Repositories – used by the package management system in order to give users easy control over the kinds of software that they are allowed to install on their system (and sometimes due to legal or convenience reasons on the distributors' side).

Root – the user root in Linux is also known as the superuser (SU) or system administrator. This is not to be confused with the system root (/), which is directory at the highest level in the filesystem. The user root can modify any file on the system and thus has the power to damage your installation if you're not careful. It is best practise to carry out most operations as a normal user so as to avoid potential pitfalls.

Terminal – See Shell.

Shell - is a piece of software that provides an interface to the operating system for users. The name 'shell' originates from shells being an outer layer of interface between the user and the innards of the operating system (the kernel). Bash (Bourne-Again shell) is the default shell on most Linux systems.

Software - is a general term used to describe a collection of computer programs, procedures and documentation that perform some task on a computer system. The term includes application software such as word processors which perform productive tasks for users and system software such as operating systems, which provide the necessary services for application software

Superuser – see Root

Unix - Unix (officially trademarked as UNIX®, sometimes also written as Unix or Unix® with small caps) is a

computer operating system originally developed in 1969 by a group of AT&T employees at Bell Labs.

Vi (or Vim) – is the standard Linux command line text editor. It's very powerful by takes some getting used to!

Window Manager - controls the placement and appearance of application windows. This may have an interface akin to that of Microsoft Windows or of the Macintosh (examples include Metacity in GNOME, KWin in KDE or twm the default X window manager).

X Window System (also X11 or X) - provides the basic framework for building the graphical user interface in Linux. It controls drawing and moving windows on the screen and interacting with a mouse and/or keyboard. X does not control the user interface — individual client programs handle this.

Free Software and the foundations of GNU/Linux

Free Software Definition

Free software is defined as software that can be used, studied, and modified without restriction, and which can be copied and redistributed in modified or unmodified form either without restriction, or with restrictions only to ensure that further recipients can also do these things. To make these acts possible, the human readable form of the program (called the source code) must be made available. This definition intends the word 'free' to mean "free as in free speech" and not "free as in free beer" with emphasis on the positive freedom to distribute rather than a negative freedom from cost. Free software is thus distinct from freeware which is proprietary software made available free of charge.

Alternative terms for free software have been coined in an attempt to make the use of "free" less ambiguous. The most common is "open-source software". Free software is also known as "software libre", "free, libre and open-source software" ("FLOSS"), and "free/open-source software" ("FOSS").

History of Free Software

In the 50s, 60s, and 70s, it was normal for computer users to have the freedoms provided by free software. Software was commonly shared by individuals who used computers and by hardware manufacturers who were glad that people were making software that made their hardware useful. In the 70s and early 80s, the increasing complexity of software applications drove the industry to protect their investments through application of copyright law, and they began using technical measures such as only distributing binary copies to prevent computer users from being able to study and modify the software.

In 1983, Richard Stallman launched the GNU project after becoming frustrated with the effects of the change in culture of the computer industry and users. Software development for the GNU operating system began in January 1984, and the Free Software Foundation (FSF) was founded in October 1985. He introduced the free software definition, and "copyleft", designed to ensure software freedom for all.

Free software is a huge international effort, producing software used by individuals, large organizations, and governmental administrations. Free software is very widely used in Internet server applications such as the Apache web server, MySQL database, and PHP scripting language. Completely free computing environments are available as large packages of basic system software such as the many Linux distributions and FreeBSD. Free software developers have also created free versions of almost all commonly used desktop applications such as web browsers, office productivity suites, and multimedia players.

The economic advantages of the free software model have been recognised by large corporations such as IBM, Red Hat, and Sun Microsystems. Many companies whose core business is not in the IT sector choose free software for their Internet information and sales sites, due to the lower initial capital investment and ability to freely customize the application packages. Also, some non-software industries are beginning to use techniques similar to those used in free software development for their research and development: scientists, for example, are looking towards more open development processes, and hardware such as microchips are beginning to be developed with specifications released under copyleft licenses (see the OpenCores project, for instance). Creative Commons and the free culture movement have also been largely influenced by the free software movement.

GNU/Linux

As mentioned above software development for the GNU operating system began in January 1984. Its name is a recursive acronym for GNU's Not Unix, which was chosen because its design is Unix-like, but differs from

Unix by being free software and by not containing any Unix programming code.

By 1991 the GNU project was well developed but many low level elements including the kernel were not complete. The Linux kernel was first released to the public on 17 September 1991, for the Intel x86 PC architecture. The kernel was coupled with the system utilities and libraries from the GNU project to create a fully functional free operating system, which led to an alternate term, GNU/Linux. Today the name is commonly shortened to Linux which, according to Linus Torvalds, its creator, is pronounced LIN – UKS.

Linux is released under the GNU General Public License (GNU GPL or simply GPL). The GPL is an example of a strong copyleft licensing scheme which gives every person who receives a copy of a work permission to reproduce, adapt or distribute the work as long as any resulting copies or adaptations are also bound by the same copyleft licensing scheme.

Today Linux is used in numerous types of computer from embedded systems such as mobile phones to supercomputers, and has secured a place in many internet server installations with the popular LAMP application stack.

There are many different flavours of Linux that you can install on your home PC. These are called Linux distributions, or more commonly “distros”. A distribution is a project that manages a collection of Linux-based software, and facilitates installation of a Linux operating system. Distributions are maintained by individuals, loose-knit teams, volunteer organizations, and commercial entities. They include system software and application software in the form of packages, and distribution-specific software for initial system installation and configuration as well as later package upgrades and installs. A distribution is responsible for the default configuration of installed Linux systems, system security, and more generally integration of the different software packages into a coherent whole. Some well known distributions are Debian, Ubuntu, Gentoo, openSUSE, Fedora, Knoppix.

Linux is largely driven by its developer and user communities. Some vendors develop and fund their distributions on a volunteer basis, Debian being a well-known example. Others maintain a community version of their commercial distributions, as Red Hat does with Fedora.

Although Linux is generally available free of charge, several large corporations have established business models that involve selling, supporting, and contributing to Linux and free software. These include Dell, IBM, HP, Sun Microsystems, Novell, and Red Hat. The free software licenses on which Linux is based explicitly accommodate and encourage commercialization; the relationship between Linux as a whole and individual vendors may be seen as symbiotic. One common business model of commercial suppliers is charging for support, especially for business users. A number of companies also offer a specialized business version of their distribution, which adds proprietary support packages and tools to administer higher numbers of installations or to simplify administrative tasks. Another business model is to give away the software in order to sell hardware.

The Linux Filesystem

The first thing that most new users shifting from Windows will find confusing is navigating the Linux filesystem. The Linux filesystem does things a lot differently than the Windows filesystem. For starters, there is only a single hierarchical directory structure. Everything starts from the root directory, represented by '/', and then expands into sub-directories. Where Windows has various partitions (e.g. c: d:) and then directories under those partitions, Linux places all the partitions under the root directory by 'mounting' them under specific directories. Closest to root under Windows would be c:.

Under Windows, the various partitions are detected at boot and assigned a drive letter. Under Linux, unless you mount a partition or a device, the system does not know of the existence of that partition or device. This might not seem to be the easiest way to provide access to your partitions or devices but it offers great flexibility.

This kind of layout, known as the unified filesystem, does offer several advantages over the approach that Windows uses. With the Linux filesystem, you can choose to mount a directory (say /usr) off another partition or even off another machine over the network. The underlying system will not know the difference because it (/usr) appears to be a local directory that is part of the local directory structure!

Another point likely to confuse newbies is that Linux also chooses to be case sensitive. What this means is that the case, whether in capitals or not, of the characters becomes very important. So this is not the same as THIS or This for that matter. This one feature probably causes the most problems for newbies.

The standard listing of the of the root directory (/) on most Linux distributions is:

```
bin
boot
dev
etc
home
lib
lost+found
media
mnt
opt
proc
root
sbin
tmp
usr
var
```

These directories contain the following:

/bin - The bin directory contains several useful commands (binaries) that are used by both the system administrator (or superuser, SU) as well as non-privileged users. This directory usually contains the shells like bash, csh etc. as well as much used commands like cp, mv, rm, cat, ls that a user cannot do without.

`/boot` - This directory contains the Linux kernel along with the bootloader that loads the kernel into memory when the computer is started.

`/dev` - This is the devices directory. It is a very interesting directory that highlights one important characteristic of the Linux filesystem - everything is a file or a directory. This may seem strange but it will make sense if you compare the characteristics of files to that of your hardware. Both can be read from and written to. Take `/dev/dsp`, for instance. This file represents your speaker device. So any data written to this file will be re-directed to your speaker.

`/etc` - This directory contains all the configuration files for your system.

`/home` - Linux is a multi-user environment so each user is also assigned a specific directory which is accessible only to them and the system administrator. These are the user home directories, which can be found under `/home/username`. This directory also contains the user specific settings for programs like firefox etc.

`/lib` - This contains all the shared libraries that are required by system programs. Windows equivalent to a shared library would be a DLL file.

`/lost+found` - Linux should always go through a proper shutdown. Sometimes your system might crash or a power failure might take the machine down. Either way, at the next boot, a lengthy filesystem check using `fsck` will be done. `fsck` will go through the system and try to recover any corrupt files that it finds. The result of this recovery operation will be placed in this directory. The files recovered are not likely to be complete or make much sense but there always is a chance that something worthwhile is recovered.

`/mnt` - This is a generic mount point under which you mount your filesystems or devices. Mounting is the process by which you make a filesystem available to the system. After mounting your files will be accessible under the mount-point. This directory usually contains mount points or sub-directories where you mount your floppy and your CD. You can also create additional mount-points here if you want. There is no limitation to creating a mount-point anywhere on your system but convention says that you do not litter your filesystem with mount-points.

`/media` - lots of modern distributions now mount attached media in the media directory so if you are looking for your cdrom or USB flash drive it may well be in here.

`/opt` - This directory contains all the software and add-on packages that are not part of the default installation. This directory is not used very often as it's mostly a standard in Unix installations.

`/proc` - The Proc pseudo filesystem is a real time, memory resident filesystem that tracks the processes running on your machine and the state of your system. The contents of the `/proc` filesystem are used by many utilities which grab the data from the particular `/proc` directory and display it. Typing `cat /proc/cpuinfo` will tell you all about your cpu!

`/root` - This is the home directory of the user root or superuser. This is not to be confused with the system root (`/`), which is directory at the highest level in the filesystem.

`/tmp` - This directory contains mostly files that are required temporarily. Many programs use this to create lock files and for temporary storage of data. On some systems, this directory is cleared out at boot or at shutdown.

`/usr` - This is one of the most important directories in the system as it contains all the user binaries (`/usr/bin`). User programs like firefox, OpenOffice etc are also placed here. `/usr/doc` contains useful system documentation. `/usr/src/linux` contains the source code for the Linux kernel.

`/var` - This directory contains variable data like the system logs which are kept in `/var/log/messages`.

Shell Commands by Alphabet

alias Create an alias
apropos Search Help manual pages (man -k)

cal Display a calendar
cat Display the contents of a file
cd Change Directory
chgrp Change group ownership
chmod Change access permissions
chown Change file owner and group
chroot Run a command with a different root directory
clear Clear terminal screen
cp Copy one or more files to another location

date Display or change the date & time
df Display free disk space
diff Display the differences between two files
diff3 Show differences among three files
du Estimate file space usage

echo Display message on screen
eject Eject removable media
env List environment variables
ethtool Ethernet card settings
exit Exit the shell

fdisk Partition table manipulator for Linux
find Search for files that meet a desired criteria
free Display memory usage
fsck filesystem consistency check and repair

grep Search file(s) for lines that match a given pattern
groups Print group names a user is in
gzip Compress or decompress named file(s)

head Output the first part of file(s)
history Command History
hostname Print or set system name

id Print user and group id's
ifconfig Configure a network interface

kill Stop a process from running

less Display output one screen at a time
locate Find files
logname Print current login name

logout Exit a login shell
ls List information about file(s)
lsof List open files

make Recompile a group of programs
man Help manual
mkdir Create new folder(s)
more Display output one screen at a time
mount Mount a filesystem
mv Move or rename files or directories

netstat Networking information
nice Set the priority of a command or job
nl Number lines and write files

passwd Modify a user password
ping Test a network connection
ps Process status
pwd Print Working Directory

rm Remove files
rmdir Remove folder(s)
rsync Remote file copy (Synchronize file trees)

screen Terminal window manager
scp Secure copy (remote file copy)
sdiff Merge two files interactively
sftp Secure File Transfer Program
shutdown Shutdown or restart linux
ssh Secure Shell client (remote login program)
su Substitute user identity
sync Synchronize data on disk with memory

tail Output the last part of files
tar Tape ARchiver
touch Change file timestamps
top List processes running on the system
traceroute Trace Route to Host
tty Print filename of terminal on stdin

umount Unmount a device
uname Print system information
useradd Create new user account
usermod Modify user account
users List users currently logged in

vi Text Editor

wc Print byte, word, and line counts
which Locate a program file in the user's path.
who Print all usernames currently logged in
whoami Print the current user id and name ('id -un')
wget Retrieve web pages or files via HTTP, HTTPS or FTP

xargs Execute utility, passing constructed argument list(s)

Using BASH

BASH has many shortcuts and useful commands that can make it quite fun to use! Here I will attempt to cover a few that I use regularly. Like any other aspect of Linux there is a wealth of information available on the internet for further shortcuts, tips & tricks should you wish for more. Happy Googling!

Tab completion – This is by far the coolest aspect of BASH. Hitting TAB while entering a command or filename will attempt to complete the expression for you. If there is more than one option hitting TAB twice will display a list of possible options.

Shortcuts:

up arrow – scroll through previous commands entered

Ctrl + C – cancel (kill) the current command

Ctrl + Z – stop the current command but don't kill it. (followed by fg or bg to restart the command in the foreground or the background. A command is normally started in the background by adding an & after the command e.g. `firefox &`

Ctrl + L - clear the screen

Ctrl + D – exit the current terminal session

Shift + Ins – insert the contents of the clipboard

Ctrl + Alt + F1 – switch to terminal number 1. There are normally 6 terminals F1-F6. F7 is normally reserved for the X-window system.

Operators:

& - using an & after a command makes it run in the background

| - pipes the output to another command e.g. `ls |more` or `dmesg |grep -i error`

> - redirects the output to a file e.g. `ls >listoffiles`

Useful commands:

man – This is by far the most useful command of all giving you access to the extensive manual pages available for shell commands. Entering man followed by the command you wish to learn about will provide you with most of what you wish to know about a command and is especially useful for syntax. For example `man mp1ayer` will keep you up for a few nights!

sudo – if you use Ubuntu as your Linux distribution (which I personally recommend as a good distribution to start with) you will need to become familiar with the use of sudo. sudo is used to carry out actions as the superuser (aka root). This is required for any any administrator functions such as adding new software. Attempting to add new software as a normal user without prefixing the command with sudo will result in an error. Other distributions that do not use sudo require you to log in as root to carry out the same tasks (it will make sense, don't worry!).

lspci - is a utility for displaying information about all PCI buses in the system and all devices connected to them – try it and find out!

uname – tells you system information including which version of the Linux kernel you're running.

dmesg – prints out kernel messages

grep - searches the input file for lines containing a match to the given pattern.

top - The top program provides a dynamic real-time view of a running system. It can display system summary information as well as a list of tasks currently being managed by the Linux kernel.

ps – list the processes currently running on the computer

kill – stop a process from the command line

df – display free disk space

free – display free memory

File management

ls – list the contents of the current directory

cd – change directory (eg `cd /` changes to the root directory `cd` changes to your home directory)

mkdir – make a new directory

rmdir – remove an empty directory

mv – move a file or directory (also rename a file)

cp – copy a file or directory

rm – remove a file or directory (`rm -r` remove a directory and all files in it)

pwd – tells you the present working directory

Ubuntu tips & tricks

Firstly & foremostly check out www.ubuntuguide.org for lots of good advice as to how to set things up on your ubuntu machine. Secondly, if things don't work as you expected then you need to ask someone else how to fix it. The way to do that is to use the forums www.ubuntuforums.org DON'T just blindly ask a question though. Do a search through the forums and see if your questions has already been answered it almost always has been. If you're setting up ubuntu on your PC for the first time there may well be things that don't work 100% out of the box. Linux by default does not include proprietary drivers so if you have a Nvidia display adaptor you will need to install the driver manually to get maximum performance from your card. This is one of the aspect of using Linux that makes it more challenging than using Windows but once you're set up you won't have to spend any time worrying about the viruses and spyware that trouble your friends computers!!!

How to configure X

```
dpkg-reconfigure xserver-xorg
```

How to setup extra repositories

Adding extra repositories - menu method:
System-->Administration-->Software Sources

Check the repositories you think you will need (main, universe, restricted, multiverse). You probably won't need the 'sources' repository.

Add any third-party repositories. Such repositories are not monitored in any way. Some are quite popular, however. Use any third-party repository at your own risk.

System-->Administration-->Software Sources-->Third-party software-->Add

Add the name of your repository. In this example, we will use Medibuntu, a popular third-party repository not affiliated with Ubuntu in any way.

APT line: `deb http://packages.medibuntu.org/ gutsy free non-free`

Download any needed gpg keys and add them to the keylist. This key verifies the repository to your system. The Medibuntu repository (not affiliated with Ubuntu) example is shown:

```
wget -q http://packages.medibuntu.org/medibuntu-key.gpg -O- | sudo apt-key add -
```

Adding extra repositories - manual method:
Create a backup of your current list of sources.

```
sudo cp -p /etc/apt/sources.list /etc/apt/sources.list_backup
```

Open the list of sources in a text editor

```
sudo gedit /etc/apt/sources.list
```

make your sources.list look like this

(# at the start of the line marks it as a comment field and not data to process)

```
deb http://us.archive.ubuntu.com/ubuntu/ gutsy main restricted
#deb-src http://us.archive.ubuntu.com/ubuntu/ gutsy main restricted

## Major bug fix updates produced after the final release of the
## distribution.

deb http://us.archive.ubuntu.com/ubuntu/ gutsy-updates main restricted
#deb-src http://us.archive.ubuntu.com/ubuntu/ gutsy-updates main restricted

##Universe

deb http://us.archive.ubuntu.com/ubuntu/ gutsy universe
#deb-src http://us.archive.ubuntu.com/ubuntu/ gutsy universe
deb http://us.archive.ubuntu.com/ubuntu/ gutsy-updates universe
#deb-src http://us.archive.ubuntu.com/ubuntu/ gutsy-updates universe

## Multiverse

deb http://us.archive.ubuntu.com/ubuntu/ gutsy multiverse
#deb-src http://us.archive.ubuntu.com/ubuntu/ gutsy multiverse
deb http://us.archive.ubuntu.com/ubuntu/ gutsy-updates multiverse
#deb-src http://us.archive.ubuntu.com/ubuntu/ gutsy-updates multiverse

## Backports

deb http://us.archive.ubuntu.com/ubuntu/ gutsy-backports main restricted universe multiverse
#deb-src http://us.archive.ubuntu.com/ubuntu/ gutsy-backports main restricted universe multiverse

## Canonical Partner Repository

deb http://archive.canonical.com/ubuntu gutsy partner
#deb-src http://archive.canonical.com/ubuntu gutsy partner
deb http://security.ubuntu.com/ubuntu gutsy-security main restricted
#deb-src http://security.ubuntu.com/ubuntu gutsy-security main restricted
deb http://security.ubuntu.com/ubuntu gutsy-security universe
#deb-src http://security.ubuntu.com/ubuntu gutsy-security universe
deb http://security.ubuntu.com/ubuntu gutsy-security multiverse
#deb-src http://security.ubuntu.com/ubuntu gutsy-security multiverse

## PLF REPOSITORY (Unsupported. May contain illegal packages. Use at own risk.)
deb http://packages.medibuntu.org/ gutsy free non-free
```

Save the edited file then add the GPG key and update apt-get to include the new lists:

```
wget -q http://packages.medibuntu.org/medibuntu-key.gpg -O- | sudo apt-key add -
sudo apt-get update
```