

Doing Things in GNU/Linux

A Newbie's Guide
Version 0.64

Raj Shekhar, Anirban Biswas, John Murray, Jason Barto.

Copyright © 2002 Raj Shekhar, Anirban Biswas, John Murray, Jason Barto.

Permission is granted to make and distribute verbatim copies of this manual provided the copyright notice and this permission notice are preserved on all copies.

Permission is granted to copy and distribute modified versions of this manual under the conditions for verbatim copying, provided that the entire resulting derived work is distributed under the terms of a permission notice identical to this one.

Permission is granted to copy and distribute translations of this manual into another language, under the above conditions for modified versions, except that this permission notice may be stated in a translation.

Preface

Most GNU/Linux distributions come with a ton of applications, development tools, and support for all sorts of fancy devices. But none of this is very obvious when you boot into KDE or GNOME for the first time. The menu contains a few apps but they are scattered about and don't have names that reveal what they do. The vast majority of tools on the system aren't even in the menus. This howto aims to help users to get acquainted with the tools available and have them up and running on their own. Instead of giving a list of programs you can run, this howto focus on tasks that you can do and how to get them done in GNU/Linux. People use computers to do things, not to run programs.

Intended audience

This howto is aimed at a complete newbie to GNU/Linux. The aim is to help a new user to find out how to do stuff with their shiny, new, just installed GNU/Linux operating system. Please note this document does not cover installation of Linux onto any type of computer. If your Linux distribution includes a paper manual you should consult the manual for detailed installation instructions. You can also use the Installation-HOWTO (<http://www.linuxdoc.org/HOWTO/Installation-HOWTO.html>) to get some help. If you have friends who are at ease with GNU/Linux you can ask them to help out with your installation problems.

So the rest of the HOW-TO assumes that you have a working GNU/Linux installation and can boot into your KDE or GNOME desktop.

Mailing the Authors

If you find any mistakes, please mail me ([lunatech3007\(at\)yahoo\(dot\)com](mailto:lunatech3007(at)yahoo(dot)com)). Please note that the email address has been mangled to avoid spamming. Before mailing change (at) to @ and (dot) to . . This how-to needs active participation from the readers and I welcome suggestions, praises and curses. Feel free to ask for help on a topic - just check that your question isn't answered here first. If you don't understand the any topic please tell us, so we can explain it better. General philosophy is: if you need to ask for help, then something needs to be fixed so you (and others) don't need to ask for help.

All emails (not spam) will be read. All flames will be sent to Timbuktu (that is what I call my trash can).

Our Distribution

Anirban and Raj Shekhar are using Red Hat Linux. John and Jason are using Mandrake. We have tried to keep things distribution independent as possible.

Copyrights and Other Delusion

© 2002 Raj Shekhar, Anirban Biswas, John Murray, Jason Barto. Permission is granted to make and distribute verbatim copies of this manual provided the copyright notice and this permission notice are preserved on all copies.

Permission is granted to copy and distribute modified versions of this manual under the conditions for verbatim copying, provided that the entire resulting derived work is distributed under the terms of a permission notice identical to this one.

Permission is granted to copy and distribute translations of this manual into another language, under the above conditions for modified versions, except that this permission notice may be stated in a translation.

Acknowledgments

This book owes its existence to the collaboration made possible by the Internet and the free software movement. Its four authors have yet to meet face to face, but we have been able to work closely together.

We think this book is a testament to the benefits and future possibilities of this kind of collaboration, the framework for which has been put in place by Richard Stallman and the Free Software Foundation.

Standard Disclaimer

Of course, we disavow any potential liability for the contents of this document. Use of the concepts, examples, and/or other content of this document is entirely at your own risk.

1 Some Dos and Donts

1.1 Do's

- Have fun. The computer is supposed to be your tool, not your boss.
- Be willing to read README files, and HOWTO documents. Linux is all about becoming more self reliant. Search engines are especially handy. I especially recommend Google (<http://www.google.com>).
- Once you've learned, help others with stuff that you understand. It improves the community.
- The motto "There's More Than One Way To Do It" applies to most activities in Linux. If there isn't, probably some college student somewhere is working on another way, but hasn't gotten around to releasing their code yet.

1.2 Important Donts

- Don't despair. We agree that not everyone's advances in self reliance will include becoming a programmer type. There are web sites dedicated to Linux newbies, and to specific topics as well. There are also IRC channels to talk to people across the internet live about this, and mailing lists, and newsgroups.

2 Getting Started

Menu, title bar, task bar, desktop and other GUI components have the same meaning in GNU/Linux as they do in other operating systems. Applications can be started either through menus, using the ‘Run Program dialog’ or by using the terminal emulator. I will outline each of the methods in the following sections.

2.1 Using Menus

Using menus to launch an application is the most common way users launch applications.

2.1.1 Using Menus in GNOME

The long bar at the bottom is a GNOME *Panel*, which contains a collection of menus, useful panel applets, and application launchers. The button containing the stylized foot is the Main Menu button (the most common kind of menu); just click on the Main Menu button and you’ll see a menu of preloaded applications and options, including a log out command.

You have to know where the application you want to use is located in the main menu. For example, in my computer GnomeCard, an address book manager, is located in ‘Programs >Applications >Address Book.’

2.1.2 Using Menus in KDE

The task bar in KDE is known as *Kicker* or *KDE Panel*. Kicker usually resides on the bottom of the desktop, but it can also be moved to any other border. On the KDE panel you will find the K Menu icon on the extreme left. As you might guess, it opens the main KDE Menu.

2.2 Using the Run Program Dialog Box

The *Run Program dialog* is very convenient when you need to enter just one command quickly.

2.2.1 Run Program Dialog Box in GNOME

In the main GNOME menu you will see a menu item labeled (most helpfully) ‘Run...’. Clicking on it will open the Run Program dialog box. Just go ahead and try it!

It has three buttons labeled ‘Run, Cancel’ and ‘Help’. In the text box type `gless` and click ‘Run’. You will behold a simple file viewer. If you get an error message, it means you do not have it installed on your system. Do not worry, you will have more chances to experiment with the "Run Program" dialog box as you read through this howto.

2.2.2 Run Program Dialog Box in KDE

In the main K menu you will see a menu item labeled ‘Run Command...’. Clicking on it will open the "Run Program" dialog box. It also has a button labeled ‘Options’ which will give you access to a host of options. You can safely ignore the ‘Options’ button till you get more comfortable with your new system.

2.3 Using the Terminal Emulator

A terminal emulator provides access to the UNIX shell in a GUI environment. In human language, you have the power of command line tools available to you coupled with the ease and safety of a GUI. There are some tools which you can access only from a terminal, so you should try to get at ease while using it.

2.3.1 Terminal Emulator in GNOME

The terminal emulator in GNOME is known as *gnome-terminal*. Usually, the *gnome-terminal* is on the GNOME Panel (the long bar at the bottom your desktop). Click on it to see the *gnome-terminal*. However, if you do not find it there, you can start one from the "Run Program" dialog.

Start the "Run Program" dialog box. In the text box type `gnome-terminal` and click 'Run'. You will get the *gnome-terminal*. You will see a blinking cursor where you can type commands. You can close the terminal as you would close any other GUI application (Usually this is done by clicking the **X** button on the window border). You can also close the terminal by typing `exit` at the command prompt. If you want to access it from the main menu, you can find it at 'Programs >System >GNOME Terminal'.

2.3.2 Terminal Emulator in KDE

The default terminal emulator in KDE is known as *konsole*. Usually, the *konsole* is on the task bar. However, if you do not find it there, you can start one from the "Run Program" dialog.

Start the "Run Program" dialog box. In the text box type `konsole` and click 'Run'. You will get the *konsole*. You can exit by typing `exit` at the command prompt or by clicking the **X** at the top of the window border.

Advice

In *Konsole*, you will see the 'Help' menu at the top of the window. It has a menu item labeled 'Contents'. Clicking on it will open the "The *Konsole* Handbook". At this point if you read through the first two chapters of the book you will be quite comfortable with *konsole*.

2.4 Notes on Terminal Emulators

Even though both GNOME and KDE come with their own terminal emulators, it does not mean that you cannot use *gnome-terminal* in KDE or *konsole* in GNOME. It is perfectly safe to use *gnome-terminal* in KDE desktop environment without causing any adverse side effects to other applications running in KDE (the Blue Screen of Death of Windows(TM) will be noticed by its absence).

As a newbie you may find yourself using *konsole* and then slowly shifting to *gnome-terminal* as you gain more confidence. However, eventually in your end days as a newbie you may want to try out another terminal emulator known as *rxvt*. This is a very lightweight terminal emulator. To try it out, in the "Run Program" dialog box type `rxvt` and click 'Run'.

However, in the start you may want to stick to *gnome-terminal* and *konsole*. After you are comfortable with them, you can move to *rxvt*.

In a similar vein, there is terminal emulator known as *nxterm*. In the GNOME main menu you can find it at 'Programs >System >Regular XTerm'. To run it from another terminal emulator or "Run Program", type `nxterm` and press **Enter**. There is another known as *xterm*. You can start *xterm* by typing `xterm` at a terminal emulator or "Run Program" dialog box.

Advice

Start a gnome-terminal. At command prompt type `konsole` and press enter. You will see a lot of messages and a `konsole` will come before you. In the `konsole` command prompt type `rxvt` and press enter. A `rxvt` terminal will start. Leave them running for two minutes and observe for yourself any adverse effects on your computer if you move your mouse about. You will be pleased to find that there are none. After this close the `rxvt` terminal, then the `konsole` window and finally the `gnome-terminal` window. As you can see this operating system is about choice :-). You may want to repeat this exercise while using `konsole` initially.

3 Understanding Linux Directories and Files

Files serve to store data permanently. Files are normally organized in a tree-like hierarchy where intermediate nodes are directories, which in turn are capable of grouping files and sub-directories.

3.1 /: The root directory

The root of the directory tree is `/`. This directory should contain no files (only directories) or very few files, such as the kernel image.

In the root directory you will find the following subdirectories.

Note Depending on the distribution you are using, you may or may not have all of the following directories.

<code>/bin</code>	<code>/etc</code>	<code>/lost+found</code>	<code>/proc</code>	<code>/tmp</code>
<code>/boot</code>	<code>/home</code>	<code>/mnt</code>	<code>/root</code>	<code>/usr</code>
<code>/dev</code>	<code>/lib</code>	<code>/opt</code>	<code>/sbin</code>	<code>/var</code>

Once you have a basic understanding of the GNU/Linux directory structure, it is easy to find and modify any file you wish.

3.2 The Directory Structure

The `/bin` directory has basic executables for commands, such as `ls`, `sync`, `cat`.The `/bin` directory holds programs that make the system usable, but aren't absolutely necessary to have to keep the system running.

The next directory is `/boot`. Not all distributions have this directory. If present, it usually contains kernels and boot maps for booting.

Next is the `/dev` directory, which is an interface to the devices attached to the computer. The files in `/dev` are not actually drivers. For more information, see [Section 3.3 \[/dev\], page 10](#).

Next is the directory `/etc`. `/etc` holds the configuration files for most applications. All manner of systems configuration files go in `/etc`.

Next is the directory `/home/user`. The files created by the users of their own that they need to store are usually stored in `/home/user`. Each user has a user name and a directory in `/home` with that user name to use as their home directory. They can place any kind of files there, and set up permissions to allow exactly whom they wish to see their data.

The next directory is `/lib`.The `/lib` directory contains shared object libraries that are needed by programs,which make system and library calls. The directories `/usr/X11R6/lib`, `/usr/lib` and `/usr/local/lib` are other repositories of libraries.

Next is the directory `/lost+found`. Files can be lost due to improper shutdown. The `/lost+found` directory is used for any pieces of files that can't be identified when the file system is cleaned up. If you lose some data mysteriously, check here.

`/mnt` is a mount point for temporarily mounted file systems. Usually to use floppy and CD-ROM drives, they are mounted in the `/mnt` directory. Some distributions use `/cdrom` for CD-ROMs and `/mnt` for floppies, while still others use `/mnt/cdrom` for CD-ROMs and `/mnt/floppy` for floppies.

Next is `/proc`. `/proc` is used as an interface to kernel data structure. For more information, see [Section 3.4 \[/proc\], page 11](#).

The `/root` directory is the home directory of root (the system administrator). Often, in `/root`, you will find programs or scripts written by root that perform some task for the machine.

The `/sbin` directory contains binaries that normal users generally don't need—binaries meant for root. Other sbin directories are `/usr/sbin`, `/usr/local/sbin`.

The `/tmp` directory is used as scratch space. If you are decompressing an application but the combined size of the zip file and the untarred data is too large for your `/home`, you can use `/tmp` for temporarily storing the zip file. The `/tmp` directory is meant to be used as storage space. It is usually deleted often to keep the system clean, so you shouldn't leave files laying around in there.

`/usr` contains many files for users (as the name implies) and their programs. It has sub directories similar to those in the root directory, i.e. `/usr/lib`, `/usr/bin` etc. `/usr/src` has the source code for apps and the **Linux kernel** resides here.

Finally, there is the `/var` directory. As the name implies it is used to store variable files (in size or time or content). It is used to store temporary material, like files waiting to be printed or mail waiting to be read, and you don't know how long you will need to store them (so `/tmp` isn't an option). `/var/tmp` can be used to store temporary information here instead of in `/tmp`. The advantage is that `/var/tmp` doesn't get deleted regularly.

With an understanding of where files are located in GNU/Linux, system maintenance becomes easier.

However, the two directories, `/proc` and `/dev` are not at first glance easily understandable. They are, however, powerful tools for understanding and using GNU/Linux. I have covered them separately.

3.3 `/dev`:The Directory to Devices

There are two categories of devices: *character devices* and *block devices*. Character devices deal with the input/output on a character by character basis. The most obvious example is a keyboard, where every key generates a character on the device. The mouse is another. Every motion or button click generates a character input for the computer to deal with.

Block devices read data in larger chunks. Data storage devices, such as IDE hard drives, SCSI hard drives, and CD-ROMs are block devices. Input/output interactions with block devices transact with chunks of data, which allows large quantities of data to be moved back and forth more efficiently.

Devices are often named after the equipment they represent. For example, devices name `/dev/hd*` represent IDE hard disks (hd).

In some cases, there may be multiple devices of the same type. For instance, a machine may have two ATAPI CD-ROMS. Two devices will be used – one for each device. For instance, `/dev/cdrom0` will be the first CD-ROM, and `/dev/cdrom1` will be the second.

A hard disk device name is composed of the type of disk followed by the position of the disk, and then the disk partition. The first hard disk may be called `/dev/hda`, with the "hd" part indicating that it is an IDE hard disk, and the "a" part indicating that it is the first hard disk. `/dev/hdb` would then refer to the second hard disk. Each hard disk is divided into partitions. The first partition on the first hard drive would be `/dev/hda1`, where the "1" at the end indicates the location of the partition.

SCSI hard disks use `/dev/sd` instead of `/dev/hd`, but otherwise the convention is the same.

3.4 /proc: The Directory of Processes

At any given time, Linux will have many processes running at once. Some, such as window managers, email clients, and web browsers, will be visible to the end user. Others, like servers and helper processes, are not immediately visible, but run in the background, handling tasks that do not require the users' interaction in order to manage the system, the kernel must keep track of every process running, including itself. Many user-level applications, too, must be able to find out what processes are running. The `/proc` file system is where the kernel stores information about processes. `/proc` is controlled by the kernel and does not have an underlying device. In addition to storing information about specific processes, the `/proc` file system contains a great deal of information generated by the kernel itself to describe the state of the system.

The kernel and its modules may generate files in `/proc` to provide information about their current state.

3.5 More Information on Linux Directory

Start any terminal emulator. At the command prompt type `man hier`. It has a detailed explanation on the hierarchy of the Linux directory structure. When you press `q`, you will exit from the man page.

In a terminal emulator, type `man proc 5`. This will give you a very detailed explanation of the files and directories found in the `/proc` directory. To exit the man page type `q`.

4 Files Managers

Files and directories help in bringing order to chaos. There are many file managers floating around for GNU/Linux. However I will detail how to work using only file managers which come preloaded with almost all distributions. In contrast to Microsoft Windows which comes with only Windows Explorer, there are three file managers in a standard GNU/Linux distribution. They are:

- GMC(GNU Midnight Commander)- In my computer this is the default file manager for GNOME.
- Konqueror - This is the default file manager in KDE desktop. It is also a web browser.
- mc(Midnight Commander) - This is the grand daddy of both GMC and Konqueror. It has to be run in a terminal emulator and is very powerful. Not to mention easy to use too.

4.1 GMC

This section will describe how to start GMC.If you want to know how to use GMC for managing files and folders see section "Managing files and folders".

4.1.1 Starting GMC from Menu

GMC can be found in the main menu as follows
'Programs >File Manager' By default, when GMC starts, you will be in your home directory. If your user name is raj, when GMC starts you will be in the directory '/home/raj'.

4.1.2 Starting GMC from 'Run Program'

Start the 'Run Program' dialog box, in the text box type `gmc` and click 'Run'.

4.1.3 Starting GMC from Terminal Emulator

You can do this from any terminal emulator like `gnome-terminal`,`konsole` etc. Start a terminal emulator, on the command prompt type `gmc` and press *Enter*. For tips on how to use file managers effectively .

4.1.4 Starting GMC from Menu KDE

In KDE, GMC can be found in the main menu as follows
'GNOME Programs > File Manager'. However, if you start GMC from KDE desktop, GMC "pollutes" your desktop with its own set of icons. Though harmless, it leads to a cluttered looking desktop. You can go ahead and try running GMC from KDE to see what I mean. Do not worry though, it's harmless .

4.1.5 Starting GMC from 'Run Program' and terminal emulator in KDE

Same as in GNOME. See [Section 4.1 \[GMC\], page 13](#).

4.2 Konqueror

This section will describe how to start Konqueror. If you want to know how to use Konqueror for managing files and folders See [Chapter 5 \[Managing Files\]](#), page 17. By default, when Konqueror starts, you will be in your home directory. If your user name is raj, when Konqueror starts you will be in the directory `‘/home/raj’`.

4.2.1 Starting Konqueror from Menu in GNOME

You can start Konqueror in the main GNOME menu as follows
`‘KDE Menus > Home Directory’`

4.2.2 Starting Konqueror from ‘Run Program’ in GNOME

Start the ‘Run Program’ dialog box, in the text box type `konqueror` and click ‘Run’. In my opinion, if you want to use Konqueror in GNOME you should use this method. To know the reason, read the next section.

4.2.3 Starting Konqueror from a terminal emulator in GNOME

You can do this from any terminal emulator like `gnome-terminal`, `konsole` etc. Start a terminal emulator, on the command prompt type `konqueror` and press *Enter*.

You will see a lot of messages flash by which will "pollute" your terminal. Even after you close Konqueror, some more messages will keep coming to your terminal which can be quite distracting. Go ahead and try it! Start `konsole` and type `konqueror` at the command prompt. You will see a lot of messages. Now, close Konqueror and you will see some more messages. This can be distracting when you are trying to get things done.

That is why I advise you to use ‘Run Program’ dialog box to run Konqueror when using GNOME desktop. You will not see those distracting messages.

See [Section 4.4 \[Notes on File Managers\]](#), page 15, for tips on how to use file managers effectively .

4.2.4 Starting Konqueror from Menu in KDE

In the main menu, click on ‘Home Directoy’.

4.2.5 Starting Konqueror from ‘Run Program’ in KDE

Start the ‘Run Program’ dialog box, in the text box type `konqueror` and click ‘Run’.

4.2.6 Starting Konqueror from terminal emulators in KDE

You can do this from any terminal emulator like `gnome-terminal`, `konsole` etc. Start a terminal emulator, on the command prompt type `konqueror` and press `enter`.

4.3 MC

The Midnight Commander is a directory browser/file manager for Unix-like operating systems. It can be used only through a terminal emulator and supports mouse functions like clicking and double clicking for selection of menu items and files. Midnight Commander provides a view of two directories at the same time. You can run Midnight Commander in any terminal emulator and in either GNOME or in KDE.

4.3.1 Starting Midnight Commander

Midnight Commander can be used in any terminal emulator and in any desktop environment. Start a terminal emulator. At the command prompt type `mc` and press enter. Midnight Commander will show up on the terminal emulator. At this point you should go ahead and try starting the Midnight Commander.

At the top you will see a menu bar with various options on it. To use the menu bar click on it with your mouse. If that does not work, you can access the menu items by pressing F9 key. You can exit Midnight Commander (MC from here on) by clicking on File menu and selecting Exit. You can also exit by using the F10 key.

See [Section 4.4 \[Notes on File Managers\], page 15](#), to learn more on how to manage files and folders with Midnight Commander.

4.4 Notes on File Managers

Unless otherwise stated, all the points stated below apply to all the three file managers.

If you want to start your file manager in some particular directory, say `‘/usr/share’` instead of the default home directory you can do that too. In the 'Run Program' dialog box type `konqueror /usr/share` and behold! Konqueror starts in the directory `/usr/share`. You can do the same thing from a terminal emulator too. This is very useful when you have to go to some particular directory directly. However, this cannot be done while using menus.

If you enter the name of a non-existent directory, you will start in your home directory.
Advice Start a `gnome-terminal`. At the command prompt type `gmc /.` You will start in the root directory. Close `GMC`. Now in the `gnome-terminal` type `gmc /foo`. `GMC` will start in your home directory.

5 Managing Files

Files and directories help in organizing the information present in the computer. Usually the following are the functions you can perform on a file:

1. Delete
2. Copy
3. Move
4. Search

5.1 Managing Files using GMC

See [Section 4.1 \[GMC\], page 13](#), if you do not know what GMC is . There are two main windows within the GNOME File Manager. On the left is the tree view, which represents all of the directories on your system by their hierarchical position. On the right is the directory window, which will show you the contents of the directory which you have selected in the tree view.

To select a directory in the tree view simply use a single mouse click. This will change the main directory view, showing the files in the directory you have chosen. If there is a plus sign to the left of a directory name in the tree window, you may click on it to show sub directories. At this point the plus sign will have changed to a minus sign. If you click on the minus sign, the sub directory once again becomes hidden in the tree view and the plus sign reappears.

5.1.1 Selecting files

Select a file by clicking on it. To select multiple files, press *Ctrl* and click on files you want to select.

You can also select by clicking and dragging the mouse cursor around several files. This action will produce a small dotted line to show you the area in which files will be selected.

5.1.2 Deleting files

Select a file by clicking on it. Left click and from the pop out menu select ‘delete’. You will be asked for confirmation before sending the file to trash can. You can delete multiple files.

5.1.3 Copying Files

Select files which you want to copy. Right click on any of the file after selecting them. From the pop up menu select ‘Copy’. A dialog box will pop up asking you the destination for copying the file. Select the destination and click ‘OK’.

5.1.4 Drag and Drop for Copying

You can also use drag and drop feature of GMC for copying files between various locations. The default action for dragging and dropping files between the directories of GMC is to move them. But you can also use drag and drop to copy a file by pressing the *Shift* key while dragging the file(s). You can also drag a file within the file manager to another directory. Select the file in the main window and drag it to the desired directory in the tree window.

You may also bring up a menu which lets you decide what action to perform with a drag by using the middle mouse button to drag a file or pressing the *Alt* key while dragging a file. Once

you release the drag you will get a pop-up menu which contains the options Copy, Move, Link, and Cancel Drag.

The motto is "There's More Than One Way To Do It"!

5.1.5 Moving Files

Select files which you want to move. Right click on any of the file after selecting them. From the pop up menu select 'Move'. A dialog box will pop up asking you the destination for moving the file. Select the destination and click 'OK'.

5.1.6 Searching for Files

In the **Commands** menu you will find a menu item labeled **Find File**. This brings up a dialog which allows you to search for particular files on your system. There will be three text boxes in the dialog box to help in defining your search criteria.

1. **Start at** - The sub directory from where to start searching the file. If you want to search all the partitions enter '/' in this text box.
2. **File Name** - The name of the file which you want to search. You can use wild cards here.
3. **Content** - If you want to look into the contents of files for some particular text, enter that text in this text box.

5.2 Managing Files using Konqueror

See [Section 4.2 \[Konqueror\], page 14](#), if you do not know what Konqueror is . When Konqueror starts, it usually has two main windows. The left window is known as *Sidebar*. If you want a better view of the directory structure, or easy access to your Bookmarks, History and Network, then the 'Window->Show Sidebar' option splits the window into two views.

5.2.1 Selecting Files

Select a file by clicking on it. To select multiple files, press *Ctrl* and click on files you want to select.

You can also select by clicking and dragging the mouse cursor around several files.

You can also select files using some *wild cards*.

- The wild card characters *, which matches any number of characters
- ? which matches a single character.

Use the Menu bar 'Edit->Select...' item. Press OK and Konqueror will highlight all files with matching names. For example;

- flag*.png will select all file names starting with the letters "flag" and ending with ".png".
- memo?.txt will select memo1.txt and memo9.txt but not memo99.txt.

5.2.2 Deleting Files

Konqueror gives you three ways to dispose of an unwanted file or directory:

- You can move it to the *Trash directory*, which is the safest method as you can get it back when you realize that you have made a mistake.
- You can just plain Delete it, which removes the entry from the directory and adds the disk area occupied by the file(s) to the system's list of free disk areas.

- Or you can *Shred* it, which is like Delete, but also overwrites the disk area occupied by the file(s) with random data. Useful for disposing of incriminating departmental memos or similar files when you don't want any trace of them to remain on your system.

The simplest way of removing a file or directory is to position the mouse pointer over its name or icon and press the right mouse button mouse button, which will bring up a menu containing the options 'Move to Trash, Delete' and 'Shred'.

Or, if you have "selected" the item, the Menu bar 'Edit' menu will give you the choice of 'Move to Trash', 'Delete' and 'Shred' options.

You can use your keyboard to delete.

- *Del* will move a selected item or items to Trash.
- *Shift+Del* will really, truly and irrevocably delete the selected item or items.
- *Ctrl+Shift+Del* will shred the selected item or items.

5.2.3 Copying Files

Position the mouse pointer over the file or directory you want to copy. Press the right mouse button, which will bring up a menu containing the 'Copy' option. Choose that. Navigate to the directory you want to copy the item into then 'Paste' the item into the new directory by using the Tool bar Paste button or the Menu bar 'Edit->Paste' option or the *Ctrl+V* hot key, or by moving the mouse pointer to a clear area of the window and holding the right mouse button down to bring up a menu containing the 'Paste' option.

Or if the item is "selected" you can use the 'Copy' button on the Tool bar or the Menu bar 'Edit->Copy' item, or the *Ctrl+C* on your keyboard.

Navigate to the directory you want to copy the item into then 'Paste' the item into the new directory.

5.2.4 Moving Files

Moving a file or sub-directory between directories can be done in the same way as copying, except that you choose the 'Cut' option or *Ctrl+X* instead of Copy. The item that you have Cut will be removed from the original directory when you do the 'Paste' into the new directory.

5.2.5 Searching Files

In the Menu bar click on 'Tools->Find file' option. This will embed the file finder application *KFind* into Konqueror's window.

Select the *Name & Location* tab, and enter the file name. Wild cards may be used as needed. The search begins in your home directory, but you can select any starting directory you wish by clicking on the 'Look in' or 'Browse' To start searching, click the icon (which looks like a magnifying glass over a sheet of paper) on the left of the tool bar. After a moment, a list of files will appear in the search results window. If they do not appear, you started the search in the wrong directory, made a spelling mistake in the name field, or no files meeting your search criteria.

You can supply more restrictions on your search criteria if you have more details about the file you are searching. There are more tabs where you can do this.

- **Date modified**-Here, you can specify that you only want to see files which were last touched in a given period of time. You can also specify that you only want to see files that were touched since a specified number of months or days ago.
- **File type**-If you know that the file was of a special type (e.g., a tar/gzip archive or a jpeg picture), you can tell *KFind* to find only this type of file.
- **Text string**-You can specify text that the file must contain.
- **Size**-If you know the file size, you can limit your search in this regard, as well.

5.3 Managing Files using Midnight Commander

See [Section 4.3 \[MC\], page 14](#), if you do not know what Midnight Commander(MC) is. The screen of the Midnight Commander is divided into four parts. Almost all of the screen space is taken up by two directory panels. By default, the second bottom most line of the screen is the shell command line, and the bottom line shows the function key labels. The topmost line is the menu bar line. The menu bar line may not be visible, but appears if you click the topmost line with the mouse or press the *F9* key.

The Midnight Commander provides a view of two directories at the same time.

You can execute commands from the Midnight Commander by simply typing them. Everything you type will appear on the shell command line, and when you press *Enter* the Midnight Commander will execute the command line you typed.

In this manual we will use the following abbreviations:

Cntrl-char means hold the *Control* key while typing the character *char*. Thus *Cntrl-f* would be: hold the *Control* key and type *f*.

5.3.1 Using Mouse

The Midnight Commander comes with mouse support. Most terminal emulators (konsole,gnome-terminal,xterm etc.) can support mouse and if you use them to run MC, it will provide mouse support too.

When you left click on a file in the directory panels, that file is selected; if you click with the right button, the file is marked (or unmarked, depending on the previous state). Marking multiple files is equivalent to selecting multiple files.

Double-clicking on a file will try to execute the command if it is an executable program; and if the extension file has a program specified for the file's extension, the specified program is executed.

5.3.2 Navigating

In the directory list, you can use the following commands for moving around

Up arrow, Cntrl-p : Moves one line backward.

Down arrow, Cntrl-n : Moves one line forward.

Prev Page, Page Up, Alt-v : Moves one pageful backward.

Next Page, Page Down, Cntrl-v : Moves one pageful forward.

Home : Moves to the beginning.

End : Move to the end.

There are two directory panels in MC. If you want to switch from one panel to to another, i.e. if you are working on left panel and want to move to the right panel, you can do it by:

- Using the mouse, left click on the panel you want to make as your new panel. If you are in the left directory panel, clicking anywhere on the right directory panel will make the right panel your new panel. The selection bar moves from the old current panel to the new current panel.
- *Tab*
- *Cntrl-i*

5.3.3 Using Menus

The menu bar pops up when you press *F9* or click the mouse on the top row of the screen. The menu bar has five menus: 'Left', 'File', 'Command', 'Options' and 'Right'.

The Left and Right Menus allow you to modify the appearance of the left and right directory panels.

The File Menu lists the actions you can perform on the currently selected file or the tagged files.

The Command Menu lists the actions which are more general and bear no relation to the currently selected file or the tagged files.

5.3.4 Selecting Files

When you left click on a file in the directory panels, that file is selected.

Selecting multiple files is known as *tagging/marking* in MC. To mark a file click it with the right button, the file is marked (or unmarked, depending on the previous state). You can also tag files by using the *Insert* key or the *Control-t*.

Re tagging an already marked file untags it.

5.3.5 Deleting Files

To delete the currently selected file or the tagged files in the currently selected panel use *F8*. During the process, you can press *Ctrl-c* or *Esc* to abort the operation.

5.3.6 Copying Files

Use *F5* on your keyboard to pop up an input dialog with destination that defaults to the directory in the non-selected panel and copies the currently selected file (or the tagged files, if there is at least one file tagged) to the directory specified by the user in the input dialog. During this process, you can press *Ctrl-c* or *ESC* to abort the operation.

Note Some file operations like Rename and Copy by default use the directory of the unselected panel as a destination (don't worry, they always ask you for confirmation first).

5.3.7 Moving Files

To move files use *F6* to pop up an input dialog that defaults to the directory in the non-selected panel and moves the currently selected file (or the tagged files if there is at least one tagged file) to the directory specified by the user in the input dialog.

During the process, you can press *Ctrl-c* or *Esc* to abort the operation.

5.3.8 Searching Files

Alt-? executes the *Find file* command. (Note that the key sequence is *Alt +Shift+./*.)

The Find File feature first asks for the start directory for the search and the file name to be searched for. By pressing the 'Tree' button you can select the start directory from the directory tree figure.

6 Connecting to the Internet

Many Linux distros allow you to enter network details during installation; doing so makes a working connection available as soon as the installation is complete. Or, you can add or change an internet connection at any time. This is particularly easy to do using a KDE tool named *kppp*. But first, you'll need to know a little about modems....

6.1 Hardware Requirements

The hardware required is a PC with a modem of course :-). Basically there are two types of modem - hardware modems and software modems (also known as winmodems). Hardware modems will generally work with any operating system, and require no drivers or configuration. Software modems (winmodems) are designed to be used with Windows, and utilise the CPU to perform many of the functions normally done by the modem. This type of modem requires suitable drivers to be installed before they will work. Generally speaking, external modems are the hardware type, while most of the cheaper internal modems are winmodems, though internal hardware modems are also available. While there are Linux drivers available for some winmodems, not all can be made to work. For this reason, if you are buying a new modem it's best to get the hardware type. Refer to this list of modems (<http://www.o2.net/~gromitkc/winmodem.html>) which work/don't work under Linux if you plan on buying a new modem. The linmodem (<http://linmodems.org/>) website has details on which software modems can be made to work, and how to set them up.

If you still have trouble configuring your modem, read Modem-HOWTO (<http://metalab.unc.edu/LDP/HOWTO/Modem-HOWTO.html>), for more information.

6.2 Setting up *kppp* from KDE

Setting up *kppp* is very easy - in fact it is easier than setting up an internet connection in Windows. First run *kppp*. In the main KDE menu 'Internet > Internet dialup tool' ie. open KDE main menu clicking on the gear in the KDE panel, now select internet and then 'Internet Dialup Tool'.

Now hit 'setup' (For the first time you run *kppp* there will be no phone number in the 'Connect To' list).

Select 'Device', here just specify where your modem is located (e.g. *com1* or *com2*) in the 'Modem Device' drop down list. Now, Linux sees *com1* & *com2* in a different way. It refers to *com1* as */dev/ttyS0*, *com2* as */dev/ttyS1* and so on. So select the suitable device name (i.e. */dev/ttyS0* if *com1*) here. If you are using an internal modem select '*/dev/modem*' from the drop down list.

Next, select 'Modem' from the tab and click on 'Modem Query', your modem will be queried. If everything goes right you will see a pop up window with some text in it, just hit 'Okay' here. If you don't get any response from the modem, check that you have the Device section described above properly configured - you can use trial and error to find the right setting if you don't know which port your modem is on.

Now select 'Account' from the tab window and hit 'New'. You will get a message box asking you 'Do you want wizard' or 'Dialog base setup'. Choose 'Dialog base setup' if you are an

Indian (as I am). You can try the wizard setup if you can find your country or network there or else just use the dialog base setup.

Another tabbed window will appear , select 'Dial' from tab, give the name of the connection (this can be anything as it is just to allow you to identify the account), then the phone number which must be dialed to get your ISP & the authentication type. See your ISP manual for the type of authentication (usually PAP), for VSNL user choose terminal based.

Then select 'DNS' tab here you have to provide the DNS server domain name or ip addresses. Now if you do not know the ip address then choose automatic & give the domain name (it is generally the part after the "sign in your email address which you obtain from your ISP. i.e if your email address is xyz@cal2.vsnl.net.in then the DNS server's domain name is cal2.vsnl.net.in .)

If you know the IP address (it is generally provided with your ISP manual) select manual & add each of them (for cal2.vsnl.net.in server i.e the VSNL server in Kolkata for cal2 user are 202.54.30.1 & 202.1.30). You may also provide both i.e. domain name as well as the ip addresses.

You will also see an "Accounting" tab. This gives you the option of keeping an account of the time or data you use with your ISP, and can be especially useful if you are billed by the hour or megabyte. If you decide to use this feature, click on the "Account" tab and you will see a list of different countries. You will be surprised to find that even India is listed. Tick the checkbox to enable accounting (you won't find this even in WinXP). Select it and you will see different types of accounts available in India (or your country). Choose accordingly and then choose the type of accounting (time, volume etc.)to suit your ISPs billing arrangement. Now hit 'Ok'.

Your account is set up & you will see the entry in the list. Click 'Ok' again. Now you are back to the starting window.

What are you waiting for? Hit 'Connect', and after some screaming of the modem a connection will be made (you can turn the volume down by selecting the "Modem" tab and adjusting the slider in the setup process). A pop up window will appear. Give your user name and password and hit connect. The window will disappear after some time (actually it will be minimized)

Congratulations, you are now connected to the Internet! Select your favourite browser or email client from the "Internet" list in the main menu. Mozilla and Konqueror are popular Linux web browsers, while kmail and Evolution are good email clients. There is also an email client built into Mozilla.

6.2.1 Solving Permission Restrictions with kppp under RedHat

Most distros (except RedHat) will allow anyone to connect to the 'net via kppp, but RedHat users will probably find that kppp will only run when you are logged in as root. There are a couple of ways to work around this; the first and simplest method only requires setting the suid ("substitute user id") of kppp to root. To do this log in as root, then open an xterm and type:

```
chmod a+s /usr/bin/kppp
```

This method will allow any user to connect with kppp.

The other method is more complex, but it has the advantage of allowing you to specify who can or cannot connect. It entails making a group, giving this group permission to connect, and then adding the users who will be permitted 'net access to the group. Again, you'll have to be running as root to do this:

- 1.Add a group called dialout (or any other funky name you can think of)
- 2.Add the specific users (who want to use kppp) accounts to this group.
- 3.Now execute the following in a terminal:

```
chown root.dialout /usr/bin/kppp
then
chmod 4750 /usr/bin/kppp
```

4. Finally, create a file named `/etc/kppp.allow` and put in the users who are in the new group (dialout), one user name per line.

The first two steps can be done graphically using a GUI user manager tool - you will find such a program inside the kontrol panel at the desktop or in Main menu—>system—>user manager. To learn how use the interface click on its "help".

7 Chatting / Instant Messaging with Linux

7.1 History of chatting

Many normal users think that GNU/Linux is not for them, but only for geeks. One reason for this is they believe they can not do the basic things like chatting under Linux. They believe that there are no Instant Messenger clients for Yahoo, MSN, and ICQ in GNU/Linux, but this is absolutely wrong - chatting was actually first implemented in UNIX. (Linux is a UNIX like free/open operating system).

`talk` was the first chatting program developed from UNIX long ago when there was no MS trying to capture the Internet. The computing world was a free land then, and you could share any program with any one, you could change them too to suit your needs - much like what Free Software (<http://www.gnu.org/fsf/fsf.html>) is trying to do. `talk` is still available with UNIX & GNU/Linux .

From `talk` other chatting concepts were developed. IRC was the first to be developed, then other companies came and hence ICQ, Yahoo, MSN, Jabber, AIM etc. chat systems were developed.

I shall try to touch on each of the chat systems here.

7.2 talk

It is the most basic chat system but still popular in LAN's. If you are in a college or office and can access only a UNIX or Linux terminal then you can chat with your friends. First of all the `talk` server daemon must running.

To chat with your friend then all you have to do the following
`[anirban@anirban anirban]$ talk <username>@host <tty>`
 i.e. if the user name is raj (it will the same as the login name to the system) and his host computer is `www.anyhost.com` then it will be
`[anirban@anirban anirban]$ talk raj@www.anyhost.com`

You may be wondering what `tty` is? Suppose your friend opened many terminals - the terminal in which you want to send the message is specified by the `tty` number. Numbers start from 0 and only integers are allowed.

You can do the above with `writetoo`.
`[anirban@anirban anirban]$ write <username@host> <tty>`

If you want not to receive any chat invitation or any chatting then you have to give the command.

```
[anirban@anirban anirban]$ mesg n
to remove the blocking you have to do
[anirban@anirban anirban]$ mesg y
```

If you a GUI lover and a heavy Yahoo or MSN chatter then you may not like this kind of chatting, but for many of us who like GNU/Linux this old system is still gold.

7.3 IRC or Internet Relay Chat

After `talk` came *IRC* and it is still popular. I think that if you really want a high class chatting experience without the flooding or other bad stuff of Yahoo and MSN then IRC is the thing for you. Also there are many rooms (known as channels in IRC) from where you can get really good help on GNU/Linux, C/C++ programming, maintaining your Linux box and much more. (As an aside, in my personal experience I was a Yahoo chatter and did not get any thing more than flooding and 4 or 5 guys running for a girl in the room.) But from IRC I received really good help when I was stuck - IRC can really be a great source of help.

7.3.1 Basic concept of IRC

IRC is a little bit different from Yahoo or MSN chat since IRC is not owned by any company. It is free/open like Linux and generally run by volunteers.

The main difference is that you do not have to sign up to get a ID or password. So what do you do instead? Choose a nick name and a host to connect to. Since it is not run by a single company you have to know the host address, like you have to know the URL to visit a page on Internet. You can get the addresses of different hosts from the internet and also to which topic it is dedicated; for example

`irc.openprojects.net` is dedicated to the betterment of open source projects and open source developers.

So you have to provide your nick name and the host you want to connect to. If the nick name you pick is already taken then you'll have to provide another nickname.

IRC newbies should check out the IRC Primer (<http://www.irchelp.org/irchelp/ircprimer.html>) before using IRC for the first time.

7.3.2 Software for IRC

There are many IRC chat applications available but I think `Xchat` is the best. Most distributions provide it with their installation CDs and it is often included in default installations. If it isn't already installed fear not; you can find it in the installation CDs or else download it from <http://www.xchat.org>. You will generally find it in RPM format so installation will not be difficult. See [Chapter 9 \[Installing and Managing Software with RPM\]](#), [page 33](#), to learn how to manage software with RPM.

7.3.3 Configuring Xchat

After installation type `xchat` in a terminal or click the `xchat` icon (you will find it in 'Main Menu > Internet > Chat').

The first window of `Xchat` will appear. Provide the nickname you would like. You can provide more than one. In case a nickname is already taken in a room `xchat` will use the other nickname you provided, else it will pick the first nickname in the list). You also can provide your real name and as which user you want to use it - (generally you do not have to provide all these, the system guesses it for you from your system login name and real name).

Now choose a host from the list of hosts and double click it or click on 'Connect' at the bottom. A new window will open with some text flowing in it. It will take a little time to connect, then after connecting it will show the rules you should follow to chat in this host. Since

IRC generally is a volunteer effort by good-at-heart people and not by any company please try to follow it or else you may get banned. Maintainers of IRC chat rooms are very strict about the rules (that is why chatting experience is much better here than yahoo or MSN).

Now you will see a single line text box where you can write what you want to chat and also commands to navigate.

Commands all start with a / (ie. slash). To get the list of rooms (or channels) in the host type `/list`. You will see all the rooms, choose the one which suits you and then type `/join #<roomname>` and then click 'Enter'. Please note that you have to always give the number sign (#) before any room name.

Now you will enter that room and start chatting. At the extreme right there will be the list of all users/chatters in that room; selecting any one will get info about him /her. You will find many buttons at the right side of your chat window, by selecting a user and clicking the buttons you can ban or block a user, get info about him/her, invite him/her in a personal chat or even transfer files in IRC.

So I think you will now be able to chat in IRC. Some day you may even meet me in IRC. I generally live in the host `irc.openprojects.net` and in the room `linux` (you have to give a number sign ie. # before joining ie. `/join #linux`).

7.4 Instant Messaging

7.4.1 ICQ

There are several ICQ (<http://www.icq.com/>) chat clients in GNU/Linux, and *Licq* is one of the most popular. You will find it inside the internet or network sections of the main menu i.e. KDE main menu (so 'Main Menu > Internet/Networking > Instant messenger') or just type `licq` in a terminal. After starting for the first time it will want you to register to their server to get a ID and password. Then you can login with that ID and password as you do with most of the Window's versions of ICQ clients.

7.4.2 Yahoo!

Yahoo provides its own yahoo messenger for Linux and it is similar to the windows version except that you may find some features missing. To get more details go to yahoo (<http://www.yahoo.com>)

Since it is similar to the windows version you will find 'Add Friends, Your Status, Ban' etc buttons in their usual places, generally as part of a menu at the top of Yahoo Messenger. Currently you can send files,invite people to group chat and get email notifications.

7.4.3 AIM

Kit (<http://kitclient.sourceforge.net>) is the AIM client for Linux (KDE). You will find it in the main menu under network / internet or just type 'kit' in a terminal. At the first startup it will ask you to create a profile and if you do not have a account in AOL it will ask you to create one by going to their site. The current version of Netscape also has a built-in AIM client.

7.5 All-In-One IM Clients

7.5.1 Everybuddy

Do you use or need several different instant messaging clients? Everybuddy is an Open Source IM client that supports AIM, ICQ, MSN, Yahoo!, and Jabber chat, as well as having some file transfer capabilities. In other words, a single Everybuddy client can take the place of several single-purpose clients. It is included (and often installed by default) with some distros, or you can download it from the Everybuddy homepage (<http://freshmeat.net/projects/everybuddy/>).

7.5.2 GAIM

GAIM is another all-in-one client resembling AIM that works with AIM, ICQ, MSN, Yahoo! and more. If you don't have it already installed, check out your installation CDs or go to the GAIM website (<http://sourceforge.net/projects/gaim/>).

8 Finding Help in GNU/Linux

Most of the applications come bundled with their own help files, FAQ (Frequently Asked Questions) and other documentation. Being able to access information efficiently is just as important as having the information available in the first place. This chapter will guide you to uncover the documentation that comes with the applications.

8.1 Man Pages

man is short for manual. Man pages are the mainstay of reference information. These pages are primarily reference material. Man pages can be read using terminal emulators. (See [Section 2.3 \[Using the Terminal Emulator\], page 6](#), for more on terminal emulators). Start a terminal emulator, and at the command prompt type

```
man man
```

You will get the manual page for the *man* command. you can press *space bar* to advance a page, *Enter* or *Return* to advance a line, *b* to go back a page. When done you can press *q* to quit.

You have to know the application or program for which you are looking up the man page. The syntax is `man <topic>`. For example, `ls` command is used to list the files in a directory. To find more information about `ls`, type `man ls`. You will find a synopsis, options, information on reporting bug etc. Most of the man pages are usually in this format.

8.2 Info Pages

Although man pages are the backbone of the reference material in the GNU/Linux world, information is presented increasingly in GNU *info format*. The syntax for viewing the info page for any program is `info <topic>`. Info comes with its own info page and a tutorial. Start a terminal emulator and at command prompt type

```
info info
```

Go through the tutorial and you will be quite comfortable with the info help system.

8.3 More Documentation

Literally every program installed in your computer places more documentation in `‘/usr/doc/<program_name>’` or `‘/usr/share/<program_name>’` on your system. This is an additional treasure trove of information. It usually contains a list of authors, FAQs, additional tips and anything extra that the developers care to throw in. Most programs will have a file called `‘README’` or `‘READ.ME’` or `‘readme.txt’` or something similar. This is a road map of the program’s documentation. By ancient convention, if you find this file you should read this first.

8.4 Programs with no Documentation

What if `info` and `man` tells me that the program `foobar` is undocumented. Here are some ways to get information on undocumented programs.

In a terminal emulator type
`foobar -? or`

```
foobar -h or  
foobar --help
```

One of them should show you the usage for the program.

8.5 Getting help From LUGs

Linux User Group (LUG) in your city will let you interact with other users of Linux. Here you can meet people who can help you out of your woes. However make sure you cannot find your answers in the documentation that accompanies the software i.e. in its info,man and other resource pages.

You can find a LUG nearest to you at Groups of Linux Users Everywhere (GLUE) (<http://www.ssc.com/glue/groups/>).

8.6 Tips on info and man

If you prefer using a GUI, here are a couple of painless (and also colorful) ways of viewing man or info pages. If you want to read man pages Konqueror makes it easy. For example type `man:ls` or `#ls` into the address bar to see the page for the `ls` command.

To browse through Info pages type `info:/dir` into the address bar. Konqueror takes you to Info's top level directory then it's just a matter of clicking on the right links to find the page you want.

Gnome help browser too can be used to read through info or man pages. To start Gnome help browser, start a terminal emulator and at the command prompt type

```
gnome-help-browser
```

In its address bar type `man:<topic>` to get your man pages. And to view info files, in the address bar type `info:<topic>`. For example to read the info page for Emacs type `info:emacs` in the address bar.

The advantage with viewing man pages like this is that you can switch to a particular man page you see listed there by just clicking on it.

9 Installing and Managing Software with RPM

9.1 What is RPM?

RPM stands for **R**ed **H**at **P**ackage **M**anager, and is an easy to use and widely adopted tool for installing, deleting, upgrading, querying and building software packages. There are other systems in use (Debian's DEB for example), though RPM is by far the most popular, and is what you'll get with Red Hat, Mandrake, SuSE and some others.

9.2 What is a Software Package?

Under Linux, programs are often distributed as single files known as RPM packages. These packages contain the actual program files, its documentation or manual pages, a summary of what the program does, menu entries and icons, plus information on where each file in the package should be installed. The package also contains information on any other files requires to run the program (*dependencies*), disk space required and so on. It's not unusual for a program to consist of a hundred or more files, so you can see that packaging them all into a single RPM file greatly simplifies the tasks of adding or removing programs. When you install an RPM package, it is uncompressed and broken down into its individual files which are then put into their correct places. RPM also checks to see whether other files necessary to run the new package (dependencies) are present. Another feature of RPM is the building and maintenance of a database of all packages installed on your computer. This means that you can quickly check to see which packages are installed, the files belonging to a particular package, or the package that provides a particular file.

9.3 Using RPM

You can use RPM from the command line, or if you prefer point'n'click, there are several graphical tools available. KDE has a particularly good one named *kpackage*; there are similar apps for users of other desktop environments, and distribution builders such as Mandrake have their own RPM front ends. I tend to use *kpackage* for removing un-needed packages to free up disk space, and the command line for everything else; but it doesn't really matter which tool you use. There are significant advantages to becoming familiar with using RPM from the command line; firstly it will be available on any RPM based system you may encounter, regardless of the desktop environment installed, plus it allows you to manage packages on machines that can't or don't run X. The ability to use wild cards (eg.*) to install multiple packages from a group at once is another feature of the command line - for example `rpm -ivh mysql*`

Just remember that generally speaking, you'll need to be running as root to install, upgrade or delete packages. However any user can run queries.

9.4 Installing Packages

In all the examples below, we'll use the Mozilla web browser package as a sample. To install it, first navigate into the directory holding the package, either from the command line or a graphical file manager. This directory might be one of your Linux installation CDs, or your home folder if you've downloaded the package. Using the command line, you would type:

`rpm -ivh mozilla-0.9.8-10mdk.i586.rpm` (your version of Mozilla might be different..)

Note that you only need to type in the complete file name when you are working with a package that is not yet installed, otherwise the basic package name (in this case "mozilla") is enough. And don't forget the tab key to auto-complete the long file names. If you prefer graphical tools, clicking on the RPM file in most file managers (Konqueror for example) will open the appropriate tool, or you can right click and use the "Open With" dialog box. It's then just a matter of clicking the 'Install' button.

9.5 Updating Packages

Updating an existing program with a later version is done in almost exactly the same way as installing. From the command line:

```
rpm-Uvh mozilla-0.9.8-10mdk.i586.rpm (note the uppercase "U")
```

Or click the 'Update' button from kpackage or similar tools.

9.6 Downgrading Packages

What if you upgrade, and then find that you preferred the older version? You can use the "--oldpackage" option from the command line like this:

```
rpm -Uvh --oldpackage mozilla-0.9.8-10mdk.i586.rpm.
```

9.7 Uninstalling Packages

To uninstall from the command line:

```
rpm -e mozilla(the complete package name isn't required)
```

Or start your graphical tool, either from the menu or a terminal. You'll see a list of all the installed packages. Click on the package you want to remove, then click the 'Uninstall' button. Note that if there are other packages installed that require files from the one you are deleting, a warning will appear and the uninstall won't go ahead. You can override this by using

```
rpm -e --nodeps mozilla (command line),
```

or selecting "Ignore Dependencies" (GUI tool), but be aware that this will break the other programs.

9.8 Querying packages

Listing all installed packages is easy. From the command line type:

```
rpm -qa
```

If the list is too big to view you can pipe it through less like this:

```
rpm -qa | less
```

Graphical tools will usually show the list of installed packages on start up.

9.8.1 Listing all of the files installed by a package

This is done with the `rpm -ql` command. Using mozilla again as an example we would type:

```
rpm -ql mozilla
```

Under a GUI tool just select the package and then click on the "File List" (or equivalent) button. Listing all of the files supplied by a package not yet installed - is done with the `rpm`

`-qp1` command. Requires the complete file name.
Eg. `rpm -qp1 mozilla-0.9.8-10mdk.i586.rpm`.

9.8.2 Listing a description of an installed package

can be performed with the `rpm -qi` command. For example:
`rpm -qi mozilla`.
Using a GUI tool, just click on the desired package. To see a description and other information of an package not yet installed, use the `rpm -qpi` command. You'll need to use the complete file name for this one. For example:
`rpm -qpi mozilla-0.9.8-10mdk.i586.rpm`.
With a GUI, select the package, or just click on the file from within your file manager (eg. Konqueror)

9.8.3 To find the package to which a file belongs

To find the package to which a file belongs is done with the `rpm -q --whatprovides` command. Example:
`rpm -q --whatprovides/usr/lib/mozilla/xpicleanup`

9.8.4 To list all the packages a package depends on

Use `rpm -qR` like this:
`rpm -qR -mozilla`
(For a package not yet installed use `rpm -qpR`, with the full file name)

9.9 Resolving Dependency Problems

The most common problem encountered when installing software is an unsatisfied dependency. You might be already familiar with this problem if you've installed new software under Windows and then found it refuses to run, with a "missing ***DLL" error message.

Linux is subject to the same problems, except that RPM will advise you of the problem before the program is installed. Many problems can be avoided when you install Linux - selecting Gnome and KDE for installation will help, even if you don't intend to run them, as many other programs use the same libraries.

So what do you do when RPM complains that a package can't be installed, because of missing packages or files. Write down the missing package/file names, and check your installation CD-ROMs for packages with similar names to the ones required. You can use the `rpm-qp1` command to view the files supplied by a not-yet-installed package. Often it is just a matter of installing these packages to resolve the problem. Sometimes, though, it leads to even more dependencies, so it can be a rather lengthy process.

9.10 Using RPMFind and RPMBone

There is another extremely useful tool for finding files and packages, and that's the RPMFind website (<http://rpmfind.net/linux/>). Type the name of the needed file or package into the box on the main page and click on the search button; you'll then be presented with some relevant information and links to the package. Often you'll already have the package on your Linux CDROM - using the information from RPMFind you'll know which one to look for. RPMFind can provide a list of dependencies for a package; the file names in this list link

back to their parent package. Usually, there is also a link to the package's home site. RPM-Bone (<http://rpm.pbone.net/>) is another site that's useful for finding and downloading RPM packages, and is somewhat similar to RPMFind. RPMBone has a more flexible search function; you can narrow down your search to only give results for a certain distribution or architecture for example. It also provides links to a huge number of ftp servers for downloading. If you need to find a package containing a particular file to satisfy a dependency however, RPMFind should be your first stop.

9.11 Circular Dependencies

Occasionally, you might come across a circular dependency. This is when package A won't install because package B is missing, but when you try to install package B, RPM complains that package A is missing. What you do here is use the `'--nodeps'` option. For example:

```
rpm -ivh --nodeps mozilla-0.9.8-10mdk.i586.rpm.
```

If you are using a GUI tool, click on the "Ignore Dependencies" button.

9.12 Library Version Problems

Sometimes a package will refuse to install because it requires a library file version later than the one already installed. This is easily fixed by upgrading the package to which the file belongs. While library files are usually backwards compatible, occasionally a package will refuse to install because a certain version of a file is missing, even though a later version is present. While you could downgrade the library package, this might well break other programs. Try creating a symbolic link with the name of the older library file that points to the existing newer version. If for example the package you are trying to install insists on having `foo.so.3`, and you already have `foo.so.4` installed in `/usr/lib`, do this (as root): `ln -s /usr/lib/foo.so.3 /usr/lib/foo.so.4`

9.13 Automatic Dependency Resolution Tools

Automatic Dependency Resolution Tools are available with some distributions. Mandrake for example has `'urpmi'`, RedHat has `'up2date'`, while Ximian has RedCarpet (<http://www.ximian.com/products/redcarpet/>). There are also tools like apt4rpm (<http://apt4rpm.sourceforge.net/faq.html>). These can automatically download and install required packages for you. See your distributions documentation or the relevant website for details.

9.14 Miscellaneous

9.14.1 RPM Version Incompatibilities

You probably won't have to worry about this unless your Linux installation is fairly old. Earlier Linux distributions were packaged with version 3.x of RPM, and are unable to handle the later 4.x series of packages. The exception to this seems to be version 3.05, you could update RPM to this version, or just replace your distribution with something newer. The RPM 4.x series is backwards compatible with the older series.

9.14.2 Midnight Commander

Occasionally you might want to copy files from an RPM package without actually installing it. You can do this with a file manager known as *mc* (for Midnight Commander). Despite being somewhat ugly, it is extremely capable. It is supplied (though not always installed by default) with most distributions, and can be started from a terminal window by typing `mc`. You can then navigate around the package as if it was a normal folder, and copy individual files from it.

9.14.3 Learning More

This article only covers the bare basics of RPM, if you'd like to learn more you could read the RPM manual page (type `man rpm`), or follow the links below:

RPM One Liners (<http://www.linuxlaboratory.org/modules.php?op=modload&name=Sections&file=>

- A concise guide by Brian Jones, worth downloading and printing for reference.

The RPM How-To (<http://www.tldp.org/HOWTO/RPM-HOWTO/index.html>) - The "official" how-to at the Linux Documentation Project.

Maximum RPM (<http://www.rpm.org/max-rpm/>) - An extremely thorough guide to just about anything that can be done with RPM. (All these resources were used in the writing of this article)

10 Printing

10.1 Which Printer to Use

First check to see if your printer is supported by Linux. Most Epson, HP, and Canon printers will be supported, though there are some cheap printers which have less hardware hence need special software to simulate the hardware. This special software isn't generally available for Linux, so you cannot use these printers. But I would not recommend you buy these types of printer anyway since their performance is less than normal printers.

You can find the list of printers that are supported by GNU/Linux at [linuxprinting.org](http://www.linuxprinting.org) (<http://www.linuxprinting.org/database.html>). I have used RH 7.3 and a HP 810c Printer here as an example.

10.2 Connecting the Printer to the Computer

After choosing the printer, check how it connects to the computer; that is which interface it uses, eg.USB (universal serial bus) or parallel port. Most printers use parallel port, but modern printers generally have both options. My printer (HP 810c) can be connected to the computer by parallel port or USB - I chose the parallel port interface. After connecting the printer comes the software part.

10.3 Installing Printer

A standard installation of Red Hat (any version from 6.1) will contain the required software for installing the printer, but old versions of the software can be difficult to configure. Here I will focus on RH 7.1 to Rh 7.3 using KDE, though it can also be done from GNOME.

10.3.1 Configuration Program

Now to install printer do the following:

1. In KDE click the 'kontrol panel' on the desktop and then click 'Printer Configuration'.
2. Click 'New' and a printer configuration wizard will appear (In RH 7.1 and 7.2 a wizard will not appear but a new window will appear, however the procedure is almost the same).
3. Next you have to specify the kind of printer that you want to add, that is 'Network Printer' or 'local Printer'. Choose local printer since the printer is attached to the machine from where you are configuring it. You have to also specify a name to identify the printer. A name must only have alphabet characters, numbers, "-" (i.e underscore) and "-"(i.e hyphen). Now click 'New'.
4. If it is a normal printer the interface it uses will be automatically detected. If not, specify it. For example if it is on the first parallel port it will be /dev/lp0, if it is in the second parallel port it will be /dev/lp1 and so on. Now click 'Next'.
5. Next choose the driver for your printer from the given list, you will find drivers for most normal printers. Different brands of printers (like HP, Canon, Epson, etc) are listed; choose the appropriate brand and then expand it by clicking the 'arrow' at the left side of the brand name or double clicking the brand name.

6. Choose your printer from the list and expand it to find the driver. Now, you may find more than one driver has been created for your printer by different people. Generally choose the driver supported by your brand, eg. the driver named "hpijs" is supported by HP and although other drivers also work "hpijs" works better.
7. Next click 'Finish' and you will be back to the main window. Now click 'Apply' and then choose 'Save Changes' from File in the main menu. (i.e 'File>Save Changes').
8. Now choose 'Restart lpd' from File in the main menu (i.e 'File >Restart lpd'). It will restart the printer daemon or process and your system will be ready for printing.
9. You can test if the installation of printer was successful by selecting 'test' in the main menu.

If you have trouble setting up the printer, you can use The Linux Printing HOWTO (<http://www.linuxprinting.org/>) for help in troubleshooting.

11 Multimedia Part 1 - Playing and Recording Music

It's easy to enjoy music with Linux, whether you are playing an audio CD, an mp3 or OGG files you recorded yourself onto CDR or hard disk. You can download tracks, or copy them from your own audio cds. While there are plenty of tools for audio work under Linux in both command line and GUI form, I'll be mainly concentrating on the command line as these tools are available on nearly all Linux distributions. Familiarity with the command line tools will also make configuring GUI programs much easier. I'll assume you already have a sound card installed and working.

Warning

Breach of copyright is taken very seriously in most parts of the world - this article in no way encourages users to break the law.

11.1 The Basics

Since much of this story involves CDs, perhaps we should start with a brief look at both audio and data CDs.

Ordinary audio CDs like the ones you'd play in your home stereo differ from data CDs in that the music is recorded onto the disk as raw data, that is, there is no file system on the disk. That's why if you put an ordinary audio CD into your CD drive and try to read the contents in a file manager, you won't find anything. Your computer is looking for a file system where there is none. An audio CD doesn't need to be mounted to be read or burnt - unlike data disks.

Data CDs on the other hand use a file system to organize the way in which the data is written to and read from the disk, similar to the file system on a hard disk. Music files in formats such as .mp3, wav, or ogg are written onto data CDs using a file system just like any other CDROM. These CDs can be opened in a file manager or from the command line, and the music played using the appropriate program.

11.2 Playing Audio Cds

There are several GUI tools available for playing audio CDs. For example Gnome has `gtcd`, KDE has `kscd`, and `xmms` can also play CDs if you have the audio CD plugin enabled. These can be started from the Multimedia section of your menu. From the command line you could try the `cdplay` program, though it's not very intuitive. Read the manual page (`man cdplay`) to find out more. Or you can simply use the 'play/skip/stop' buttons on your CD drive to play audio CDs.

11.3 Playing MP3 Files

The mp3 format is a hugely popular way of storing and sharing music. One reason for its popularity is its compact size compared to other formats or conventional audio CDs. A typical mp3 file is usually only about a tenth of the size of the same file in .wav or audio CD form. This means you can fit the equivalent of ten audio CDs on a single CD using the mp3 format. Other advantages include the reduction in space used when storing music on a hard disk, and the smaller file size also makes transferring files over a network much more practical. The disadvantage is that mp3 CDs can't be played on most normal CD players (although mp3 compatibility is starting to appear on some portable Walkman type players), so you can only

play them on your computer. The most popular player for Linux is probably `xmms`, an excellent clone of the Windows Winamp player.

11.4 Using XMMS

Xmms (**X** MultiMedia System) is a widely used multi purpose sound file player that is included with most common distributions. It's mostly used for playing mp3 files, but it can do much more than that. It is also capable of playing wav files, *ogg-vorbis* files (an open source alternative to mp3), streaming audio etc. Starting xmms can be done from the menus (look under "Sound" or "Multimedia"), or from a command line just type `xmms`. The interface is much like that of a CD or tape player, with buttons and sliders to control starting, stopping, pausing, skip, repeat, volume, balance and so on. It also includes an equalizer function (the 'eq' button) and allows you to set up play lists. To choose a track to play, hit the *L* key or press the eject ("^") button. This brings up a window allowing you navigate to the folder holding your music files. Once there, you can select a track or tracks to play, or you can choose to play every file in the folder. As well as the audio options, xmms also has visual options, and different skins can be selected to change the appearance of the player. It can even use Windows Winamp skins. Despite the multitude of options, xmms is exceptionally easy to use. If you want to explore its options and capabilities, click on the small *O* on the left hand side of the display.

11.5 Recording (or ripping) Tracks from Audio CDs

There are several tools for recording audio CDs to hard disk. You can record a single track, selected tracks or the entire CD at once. The music will be converted to a file format that can be read by your computer (usually `.wav`) as it is recorded. While there are both command line and graphical tools for the job, my favorite is the command line program `cdparanoia`. If you prefer GUI tools, you might like to check out `grip`. One of the things I particularly like about `cdparanoia` is the way it can correct jitters or skips on marked or scratched disks. Here are some examples of how to record tracks from an audio CD using `cdparanoia`:

To record a single track type:

```
cdparanoia n
```

'n' specifies the track number to record. By default the track will be recorded to a file named `cdda.wav`. If `cdda.wav` already exists it will be overwritten, so be careful if you are recording several tracks! You can specify your own file name like this:

```
cdparanoia n filename.wav
```

To record the entire CD type: `cdparanoia -B`

The `-B` in the above command simply ensures that the tracks are put into separate files (`track1.wav`, `track2.wav` etc.). `Cdparanoia` has many more options and an easy to understand manual page; type `man cdparanoia` to read it.

11.6 Converting .wav files to mp3

If you intend to burn files to an audio CD, you should leave them in `.wav` format. On the other hand, if you want to play them from your hard drive, or burn a data CD that you'll play

from your computer, you'll probably want to convert them to the mp3 format to save space. One of the most popular tools for this is *bladeenc*. To convert a .wav file into a .mp3 use this command:

```
bladeenc filename.wav
```

This will produce a file with the same name as the source file, but with the .mp3 suffix. If you want to specify a destination filename you can add it to the end like this:

```
bladeenc filename.wav filename.mp3
```

By default, bladeenc will encode the file at 128kbit/sec, this is the most commonly used bitrate and results in a very compact file of reasonable quality. Higher rates can be specified, giving a better sound quality at the expense of a slightly bigger file size, though it's hard to detect any improvement in sound quality using sampling rates above 160kbits/sec. To convert a file at 160kbits/sec use:

```
bladeenc -160 filename.wav
```

11.7 The Ogg-Vorbis Format

Ogg-Vorbis is a completely free and open alternative to the mp3 format. The sound quality is at least as good as mp3, and ogg files can be played on players such as xmms. You'll need the vorbis-toolspackage (http://www.vorbis.com/download_unix.psp) (check your distributions installation CDs) to convert .wav files to .ogg. Converting is easy:

```
oggenc filename.wav
```

As with bladeenc, the sampling rate (and sound quality) can be specified. This is done by using the following command:

```
oggenc -q n filename.wav (where n is the quality level)
```

The default level is 3, but can be any number between 1 and 10. Level 5 seems to be roughly equivalent to an mp3 encoded at 160kbits/sec.

11.8 Converting .mp3 files into .wav format

Audio CDs are usually burned from a collection of .wav or .cdr files - you can't directly burn mp3s to an audio cd unless you convert them to one of these formats. The *mpg123* program can do this for you and is often installed by default with many distributions. To convert an .mp3 to a .wav, type:

```
mpg123 -w filename.wav filename.mp3 (note - the destination filename comes first)
```

Note also that there is some slight loss of sound quality when a .wav file is converted to mp3 format, and this isn't regained when converting back to .wav - so if possible, you should try to use .wav files that have been ripped from an audio CD rather than converting back from mp3s.

11.9 To Normalize a group of .wav files

Creating an audio CD using tracks from different sources can lead to variations in volume amongst the tracks. By using a program named *normalize*, you can equalize the volume level of a group of files. You'd normally do this to a group of .wav files before burning them to CD. Normalize is a command line tool; to equalize the volume of all the .wav files in a folder type:

```
normalize -m /path/to/files/*.wav
```

11.10 Recording (or burning) an Audio CD

I'll assume that you have a CDR or CDRW drive installed and configured already - if you don't, see the links section at the end of this article for more information on set-up details. I'll also assume that you'll be using *cdrecord* to burn your disks - it's the most popular tool for this and is also what's used by most graphical front-ends like *XCDRoast* etc. Your files will need to be in .wav or .cdr format; most likely they will be .wavs. Put all the files you want to burn into a separate folder to simplify the burning process, and make sure that they will fit onto the disk (you can check by changing into the folder and running the *du* command). Now it's just a matter of typing this command:

```
cdrecord -v speed=4 dev=0,0,0 -audio -pad *.wav
```

Of course, your speed and device numbers might be different - you can use *cdrecord -scanbus* to find the device address, and the speed setting will depend on your CD burners' speed rating. In general, burning will be more reliable at slower speeds, especially on older machines.

11.11 Recording a Data CD (mp3 or ogg)

If you only plan to play a music CD on your computer or other mp3 capable device, you can burn mp3 or ogg files in exactly the same way as an ordinary data CD. Because data CDs use a file system, we'll use *mkisofs* (to create the file system) and *cdrecord* to burn the disk. As in the audio CD example above, put all the files into a separate folder. The two operations can be combined into a single command like this:

```
mkisofs -R /path/to/folder_to_record/ | cdrecord -v speed=4 dev=0,0,0 -
```

Don't forget the hyphen at the end! As in the example for burning audio CDs, you might have to use different speed and dev numbers. Older or slower computers might have difficulties running both *mkisofs* and *cdrecord* at once - if so you can do it as two separate operations like this:

```
mkisofs -R -o cdimage.raw /path/to/folder_to_record/
```

This creates an image named *cdimage.raw*. Then burn the disk:

```
cdrecord -v speed=4 dev=0,0,0 cdimage.raw (using suitable speed and device settings..)
```

11.12 Some detailed information on related topics:

The Linux MP3 How-To (<http://www.tldp.org/HOWTO/MP3-HOWTO.html>)

The Linux MP3 CD Burning How-To (<http://www.tldp.org/HOWTO/mini/MP3-CD-Burning/index.html>)

The SOX Homepage (<http://sox.sourceforge.net>)- The swiss army knife of Linux sound tools.

The Normalize Homepage (<http://www.cs.columbia.edu/~cvail/normalize>)

Installing and Setting Up a CDR/CDRW (<http://www.linuxgazette.com/issue57/feenberg.html>)

- ***Note*** Modern desktop distributions can usually detect and setup a CD burner without any manual configuration required. This page may be useful however for older/difficult distributions that require manual installation.

The OggVorbis Homepage (<http://www.vorbis.com/>)

The Bladeenc Homepage (<http://bladeenc.mp3.no/>)

The CDRRecord Homepage (<http://www.fokus.gmd.de/research/cc/glone/employees/joerg.scchillin>)

The CDParanoia Homepage (<http://www.xiph.org/paranoia>)
The mpg123 Homepage (<http://www.mpg123.org/>)

12 Multimedia Part 2 - Linux and Video

12.1 Introduction

Like most other Linux programs, the video-media playing applications have matured greatly in the last couple of years, and it's now possible to view any reasonably common type of video file under Linux. But don't think that viewing video clips is as simple as say, playing music files; while the vast majority of music is in two or three common formats (mp3, vorbis, wav), video files are commonly produced with a variety of formats and codecs (more on codecs later).

This means that those users who view lots of video from a variety of sources may have to occasionally download and install additional software in order to play the various formats. On the other hand, casual users will probably find that the distro-supplied multi-purpose media apps will satisfy their requirements.

In the paragraphs below, I'll try to explain some of the jargon you're likely to come across when dealing with video files and media players. Don't worry too much about this if (for example) you just want to play the occasional mpeg, you can safely skip ahead to the Linux Media Players section.

12.2 File Formats and CODECs

12.2.1 File Formats

Anyone trying to play or download video clips from the internet will soon discover that there is no standard (or even defacto) file format. Unlike music files, which are dominated by mp3s, it's not unusual to find video files in mpg, wmv, mov and rv formats, as well as several others. To complicate matters even more, there can be a variety of file types even within a single format; so just because you can play some .mov files (for example), doesn't necessarily mean you can play all of them. In other words, instead of referring to a specific filetype, filename extensions can be a wrapper for a variety of filetypes. And this is where codecs come in...

12.2.2 Codecs

In the context of multimedia, codec is short for compressor-decompressor (it's also used as an abbreviation for coder-decoder, but that's a different story...). In simple terms then, we could say that the codec is the algorithm or software routine used to compress or decompress a media file. The important point to remember is that if a particular codec is used to create a file, that same codec will be required by your media player to play the file. And as we mentioned earlier, there may be several types of codec that can be used with a particular file format, so not all files of the same format will necessarily be playable by your media player. In practice, you could probably think of codecs as being the library files or plug-ins that you download and install to enable your media player to handle a particular type of file. Most players will come with a variety of codecs included, while others are available on the net. Some players can even make use of MS Windows DLLs.

12.3 Streaming Media

Streaming video and audio is becoming increasingly common on the web, and streaming simply refers to the way the file content (or data stream) is fed directly into a media player or browser plugin instead of being first downloaded to disk. The advantage is that the contents

are displayed immediately, so that you don't have to wait for the entire file to be downloaded before you can view it. In effect, you view the file as it is being downloaded. There is an obvious disadvantage to this however; you need to be able to download the file at least as quickly as your viewing program plays it. And because of this, there are limitations to streaming media files over a dial-up connection.

12.4 Some Common File Formats

The list below includes the most commonly encountered formats, but is by no means exhaustive. All the formats mentioned can be played under Linux, though some may require closed or proprietary software.

12.4.1 mpg (or mpeg)

The group of file types known as mpegs include the type 1 and 2, which are very widely used video files that are quite compact. Type 3 mpegs are better known as the ubiquitous mp3 sound files, while the newer type 4 mpegs are high quality, streamable files that are closely related to the DivX codec .

12.4.2 AVI (Audio Video Interleaved)

The original avi file is an older, uncompressed format not seen so often on the net now, but still used on some CDROMs. Basically a slideshow of images, with audio data interleaved. Not streamable. Note that some other codecs such as DivX also now use the .avi suffix.

12.4.3 DivX (.avi)

A patented codec from DivXNetworks and based on Mpeg-4 standards. DivX files are of a high visual quality and highly compressed. DivX codecs are available for Linux media players.

12.4.4 Mov (Quicktime)

There are several variations of Quicktime, perhaps the most common now being the Sorenson codec. Until fairly recently, it was difficult or impossible to play files encoded with this proprietary codec under Linux, though now players such as mplayer are able to utilise MS Windows .dlls to play them without problems.

12.4.5 RealMedia Video (.rm)

This streamable format from RealNetworks is fairly popular. The Linux version of RealPlayer has all the functions of the Windows version. RealServer is also available for Linux.

12.4.6 Windows Media Video (.wmv)

This is a Microsoft format that uses the Windows Media Player. Files can be downloaded and played from disk or streamed.

12.4.7 ASF

Microsofts Advanced Streaming Format is similar to the wmv above, but specially designed for streaming applications.

12.4.8 Shockwave Flash (Flash, .swf)

The Shockwave Flash format is an extremely popular way to make web animations, cartoons, games etc.

12.5 Linux Media Players

These are only some of the players you are likely to encounter - there are many more available.

12.5.1 KDE Media Player (noatun)

A basic KDE media player that can handle a few audio and video filetypes, including mp3, wav and mpeg-1 video. Some other codecs can be played by installing plugins. Often installed by default, noatun is part of the kdemultimedia package.

12.5.2 Xine

This is a very popular player that can be extended with plugins to handle nearly any of the common filetypes. Xine is often installed by default, or at least included on the distro CDROMs. As well as being able to play many video formats, xine is widely used to play DVD movies, though you'll probably have to install some additional components in order to play encrypted DVDs. Xine comprises several separate packages; you need to install at least two of these in order to use it, though installation should be reasonably easy. The two main packages are *xine-lib*, which is the engine of xine, and *xine-ui*, the user interface. There are other xine related packages to allow the use of extra codecs and plugins, among other things. These should be already installed on systems where xine is installed by default. In general, xine is a very capable player, though I did find its controls a little confusing. Some other players use the xine engine, but with their own GUI.

12.5.3 XMovie

This is another player that's often included with current distros, and can play most mpeg-1 and mpeg-2 video files, as well as a few other video and audio files. Designed to be a player of high quality, high resolution movies, it can also play VCD files and function as a rudimentary DVD player. While it's somewhat limited when compared to xine or mplayer, it does have a clear, easy to understand interface, and very good audio capabilities. For many users who only need to view the occasional mpeg and so on, xmovie will work nicely.

12.5.4 XMMS

Not just a music player, XMMS can also play some mpeg video files when the smpeg plugin is installed, as it often is with default installations. Click on *Preferences > Audio I/O Plugins* to see the list of installed / enabled plugins. If it's not there, look for a package named *xmms-smpeg* on your distros installation CDROMs. There is more information regarding other playable/installable codecs at the XMMS homepage.

12.5.5 MPlayer

A versatile, heavy duty app, mplayer can handle nearly any type of video. In fact it'll play anything that runs under Windows Media Player, plus more. Mplayer is also widely used for playing DVD movies, and while it's not usually included with distros (some include a stripped down version), it is readily available from the web. The officially supported version is available only as source code; though there are links to Redhat RPMs as well at the main site. For those

not comfortable with compiling from source there are pre-compiled packages available from other sites. I found RPMs for my Mandrake box at the PLF site; other sites such as RPMFind etc. also have binaries. The only problem with installing mplayer from RPMs is that you'll probably have to track down and install quite a few other packages in order to satisfy dependencies - I actually found it much easier and quicker to compile from source. Compiling mplayer is easy, and only requires the three usual steps (assuming you've already unpacked the archive into a temporary directory, and cd'd into that directory):

```
./configure  
make  
and then, (as root):  
make install
```

You should also install the win32-codecs package before building mplayer to enable the use of many more codecs; the win32 RPM package is fine.

The instructions above are for the the bare-bones command-line version; if you want or need a GUI you'll need to take a couple of extra steps, though it's not complicated or difficult. See the excellent documentation at the website for details. You can play a video file with the command line version with:

```
mplayer filename
```

or setup your file manager to start mplayer whenever you click on a video file. In short, if you need to view a wide variety of video file types, mplayer is hard to beat.

12.5.6 Mtv / mtvp

Mtv is a shareware mpeg player, and while mtv itself isn't free, there is a free (for personal use) command line player named **mtvp** in the mtv package. Mtv is fast, light, and uncluttered, features that make it a favourite for older or less powerful hardware.

12.5.7 RealPlayer

The Linux version of RealPlayer is fully featured, and can play streaming video and audio as well as files from disk. In addition to Reals' own rm and rv file formats, it can also play a few other commonly found files such as mp3, wav, midi, some mpegs and some swf files.

12.5.8 Xanim

This one goes back quite a few years, and is not usually supplied with current distros. It can handle some AVI and older Quicktime codecs, plus there are other codecs packages available. Some other players are also able to make use of the xanim codec packages. Xanim is the engine behind Aktion, an older KDE app.

12.5.9 Macromedia Flash Player

A little different to the other players mentioned above, this one is more likely to be used to play animations, cartoons or even interactive games, rather than video clips. There are two versions available; one is a web browser plugin while the other is a stand-alone program (gflashplayer) that can play downloaded swf files. Installation is easy, and consists of unpacking the archive and then running the install script. Versions 6.0 and later are much improved and seem to be much more reliably detected, so update if you have an older version.

12.6 Miscellaneous

As you can see, video playback isn't as simple or straightforward as many other tasks, but this complexity isn't something peculiar to Linux; it could even be argued that some of the Linux players are more capable than those for other platforms. You'll probably also notice that some of media players that are included with Linux distros ship with only some of their usable codecs or features installed. This is because some of these features require closed or proprietary components that can't be included because of licensing issues. Other things - encrypted DVD playback being a well known example - may be subject to legal issues when used under certain conditions. And because the distro makers are unable to control how these tools are used, it's left up to the end user to download these components and install them. You can often find the relevant links at the respective media players homepage, or you can try the links below. And don't forget Google; there's a huge amount of info out there...

12.6.1 Additional Resources

The Xine home page (<http://xinehq.de/>)
MPlayer's Home page (<http://www.mplayerhq.hu/homepage/>)
Mtv homepage (<http://www.mpegTV.com/>)
Xmovie home (<http://heroinewarrior.com/xmovie.php3>)
The XMMS Homepage (<http://www.xmms.org/>)
The RealPlayer Unix/Linux download page (<http://scopes.real.com/real/player/unix/unix.html>)

Xanim Homepage (<http://xanim.polter.net/>)
Macromedia's download page (<http://sdc.shockwave.com/shockwave/downloads/alternates/>)

PLF (Penguin Liberation Front) (<http://plf.zarb.org/>) have a selection of files for various versions of Mandrake that can't be included with the official Mandrake distro for whatever reason.
LinuxApps Multimedia listing (<http://linuxapps.com/?page=category&category=multimedia>)

LinuxApps MPEG listing (<http://linuxapps.com/?page=category&category=mpeg>)

13 Linux in the Office

Office applications for Linux are now quite mature. Linux desktop productivity tools are in fact so capable and feature rich that corporations are beginning to look at alternatives to MS Office with its high TCO, and leaning towards office suites like OpenOffice and StarOffice. OpenOffice in fact is a part of discussion being held between multiple companies including Boeing Aeronautics, a major international technology contractor, to begin to define a standard for Office document formats. Allowing greater portability of documents between office suites; XML of course is being discussed as the most viable vehicle for the mission. But I digress, to sum it up, if someone wishes to do all of their word processing, spreadsheets, and so on on the Linux desktop they would find themselves very satisfied with today's applications.

13.1 Word Processing

User's choices for word processing on Linux are varied and diverse. To list several applications would probably only begin to scratch the surface of what is available. So in an effort to simplify things I will include a review of those applications with which I have experience and list a few more with which I do not. In addition all of the Word Processors I have used in the past are Microsoft Word compatible - meaning that they can both read and write MS Word documents. This will come in handy for all those who are afraid they will never be able to open a '.Doc' (Word Document) once they move over to Linux.

13.1.1 StarOffice Star Writer

StarOffice is an office suite written for Unix / Linux and developed by Sun Microsystems. Until recently StarOffice was a freely available download but recently with their newest version (I believe StarOffice 6.0) they have begun charging a license fee. I haven't used StarOffice 6.0 but I am familiar with the final release just before it. StarOffice is a functional Office suite with many additional features. When you first open StarOffice you are presented with a screen very similar in appearance to the MS Windows desktop complete with 'Start' bar. StarOffice provides the full suite of functionality including word processing, spreadsheets, email, and MS PowerPoint-like presentations. And again, any and all documents written in StarOffice can be saved in the equivalent MS Office format so you lose no compatibility with co-workers / family members / unconverted Windows- but-soon-to-be-Linux users. StarOffice, along with all the other applications I review here provides a very similar interface to MS Word. So there is little to any learning curve involved with using it. In fact the only real difference between the list of applications reviewed here and those in MS Office is how well the applications can read and write in the MS format. StarOffice does a very adequate job of processing MS Word documents. The only area where StarOffice runs into trouble is reading and writing MS Word documents that have tables embedded in them or those containing forms. However if it is merely straight text, such as a report there is typically no problem involved. But it is my unconfirmed suspicion that even this has changed now that you can pay for StarOffice. Again I have not checked this first hand but I believe the reason Sun now charges for StarOffice is because they paid Microsoft for the APIs that allow StarOffice to read and write MS Office documents. Up until now the formats have merely been reverse-engineered, kind of a best guess at how to interpret the symbols in a MS Word document. For more information, and to confirm / deny my crazy allegations check out StarOffice at <http://www.staroffice.com>.

13.1.2 OpenOffice Writer

OpenOffice is a spinoff of Sun's StarOffice (as the name may suggest). And like StarOffice, OpenOffice also provides a suite of applications including word processing, spreadsheets, and MS

PowerPoint-like presentations. OpenOffice also supports the reading and writing of MS Office documents. Recently I rewrote my resume (and being a Linux-only kinda guy I of course couldn't use MS Word) using OpenOffice. This consisted of multiple fonts and font sizes, the embedding of tables so as to properly position the many elements of my resume, and also included bullets. After completing my resume I proceeded to save it in both the native OpenOffice format as well as the MS Word 2000 format. Of course before shipping it out to employers I wanted to check to see how it would look in MS Word. So when I went to work the next day I proceeded to open it using (the very expensive) MS Word 2000. Much to my surprise, with the exception of some bullets the resume had made it through quite well. All the tables were properly in place, the fonts well represented in their multiple sizes, and the only thing wrong with the bullets was that instead of the '>' arrow I had originally had, it was replaced with a round bullet (I guess perhaps that MS Word didn't support the particular type of bullet I had specified.) So OpenOffice (if you absolutely refuse to pay for software) will do very well for your Office and Word Processing needs. More can be learned about OpenOffice (you can download a copy of OpenOffice from here too) at <http://www.openoffice.org>.

13.1.3 AbiWord

The only complaint I can really make concerning AbiWord can be summed up in one word - tables. While AbiWord does support tables the interface and handling of tables has a long way to go. Otherwise AbiWord is much like StarOffice and OpenOffice. Reads and writes simple MS Word documents, very MS-like interface, etc. Another nice feature of AbiWord is its support of Gnome themes; a feature that neither OpenOffice or StarOffice provide.

13.2 Others Word Processors

13.2.1 Kword

Kword is part of the KOffice office suite for the KDE desktop. It has all the usual bells and whistles, frames, numbering, bullets, tables, paragraph alignment, etc. However from what I can see at on the webpage KOffice does not support reading and writing of MS Word documents. If you would like to learn more point your browser to <http://www.koffice.org>

13.2.2 Corel WordPerfect

WordPerfect was once the dominant word processor for PCs, and the latest available version for Linux is WordPerfect2000. It is a fully featured application and is unusual in that it is not Linux native, but is essentially the Windows binaries running under a built-in version of Wine. For this reason, it may not be as stable or fast as some of the others. You can find out more about Corel WordPerfect 2000 at: http://linux.corel.com/products/wpo2000_linux/index.htm

13.3 Spreadsheets

Spreadsheets are possibly the most widely used Office program, and as with word processors, Linux users have several quality offerings to choose from.

13.3.1 Gnumeric

Gnumeric is the GNOME projects spreadsheet, and it is a mature and stable program. Compatibility with MS Excel files is quite good, and gnumeric is often installed by default with many distros, or at least is available on the installation CDs. A good choice for those who don't want to install a big, fullblown suite like Star/Open Office.

13.3.2 StarOffice/OpenOffice Calc

The Calc spreadsheet is another very competent office tool, with very good Excel compatibility. Possibly the best choice for heavy duty spreadsheet users.

13.3.3 kspread

Koffices' spreadsheet, kspread, is a good looking, powerful app, however its Excel compatibility is somewhat limited, so if this is important to you, perhaps one of the others would be a better choice. Otherwise an excellent spreadsheet.

13.4 Other Office Applications

The number and quality of office type apps. has grown rapidly in the last couple of years. Some of these are listed below, along with a brief description and links. Most people will probably find these will meet their needs, though some may find they are dependant on certain features of Microsoft Office apps. that just aren't available under Linux yet. For these users, a proprietary product known as Codeweavers Crossover Office (<http://www.codeweavers.com/products/office/>) allows MS Office (as well as some others) to be installed and run directly from Linux. I've only listed the more well known programs here, and some of these are probably already installed on your computer. If you need to install them, most of these packages can be found on your installation CDs, otherwise just follow the links. The KDE apps listed here are mainly included in the koffice (<http://www.koffice.org/>) package, while the GNOME programs are usually separate packages.

13.4.1 Address Books

GNOME has 'gnomecard' (part of the gnome-pim package), KDE uses 'kaddressbook'.

13.4.2 Fax Apps

There is 'kfax' with KDE, gfax (<http://www.gnome.org/gnome-office/gfax.shtml>) for GNOME. Programs like hylafax and mgetty+sendfax are also popular.

13.4.3 Email/PIM

Outlook users will probably be most interested in Ximians' Evolution (<http://www.gnome.org/gnome-ofi>) a fully featured email/PIM program. There is also a proprietary add-on for Evolution named Connector, and this can enable Evolution to function as an MS Exchange client. As well as email, it has address book, calendar and task-scheduling/alarm features.

13.4.4 Drawing/Graphics

Dia (<http://www.gnome.org/gnome-office/dia.shtml>) is a structured diagrams program similar to Visio, while Sketch (<http://www.gnome.org/gnome-office/sketch.shtml>) is a vector drawing package.

KDE has Kontour (<http://www.koffice.org/kontour/>) (another vector drawing tool), Kivio (<http://www.thekompany.com/projects/kivio/>) for flowcharts, and KChart (<http://www.koffice.org/kchart>) for drawing charts/graphs.

13.4.5 Financial

Gnucash (<http://www.gnome.org/gnome-office/gnucash.shtml>) is a popular personal finance manager, though there are several others. And if you just can't survive without Quicken, you'll be pleased to know it will run under Linux using Codeweavers Crossover Office.

13.4.6 Database

postgreSQL (<http://www.postgresql.org>) is included with distros such as Mandrake and RedHat, also there is MySQL (<http://www.MySQL.com/>), a somewhat simpler database.

As well, there are databases such as Interbase (<http://www.borland.com/interbase/>), and Firebird (http://www.ibphoenix.com/main.nfs?a=ibphoenix&page=ibp_development) a free, open source version of Interbase.

The big names like Oracle (http://www.oracle.com/ip/dep/otn/database/oracle9i/index.html?db_linux.html) and IBM (<http://www-3.ibm.com/software/data/db2/linux/>) (DB2 for Linux) support Linux too.

13.4.7 Presentation Apps.

The major office suites (StarOffice, OpenOffice, Applix, KOffice) all have functional presentation programs. The StarOffice and OpenOffice versions can handle MS PowerPoint format files.

13.4.8 Organizers

If you are looking for something a little lighter than Evolution, KDE has 'korganizer', and GNOME uses 'gnomecal' (part of the gnome-pim package).

13.4.9 Calculators

RedHat distribution installs three calculators by default. *Xcalc*, *GNOME Calculator*, *KCalc*.

Xcalc is a scientific calculator desktop accessory that can emulate a TI-30 or an HP-10C. *Xcalc* can be started from a terminal emulator or from the Run dialog box by typing `xcalc`. It takes the following command line argument (among others)

-rpn This option indicates that Reverse Polish Notation should be used. In this mode the calculator will look and behave like an HP-10C. Without this flag, it will emulate a TI-30.

GNOME Calculator is a double precision calculator application. *GNOME Calculator* is included in the `gnome-utils` package, which is part of the GNOME desktop environment. It is intended as a GNOME replacement for *xcalc*. To run *GNOME Calculator*, select `gcalc` from the Utilities submenu of the Main Menu in GNOME, or type `gcalc` on the command line in a terminal emulator or Run Program dialog box.

KCalc can be started by typing `kcalc` on the command prompt or in the Run Program Dialog box.

13.5 PDF Files Viewer

PDF (Adobe's Portable Document Format) is a format for transferring documents with formatting (including fonts, sizes, etc) with a few more extra features (such as URLs). It is a quite common format for publishing documents - it is generally quite difficult to edit such a document, but relatively easy to show it, as it already contains an exact definition of the document (somewhat similar to postscript). There are a number of viewers for PDF documents under GNU/Linux.

13.5.1 XPdf

'xpdf' supports most of PDFs features, including LZW-compressed images URLs and encryption. It can be started from command prompt by typing `xpdf`. xpdf's home is at <http://www.foolabs.com/xpdf/> . It also comes with a number of distros, including RedHat, Mandrake and SuSE.

13.5.2 Adobe Acrobat Reader

This is not free software (although it can be used free of charge for non-commercial use). Free software gives you permission to use, copy, study and improve the software. You can learn more about Free Software here (<http://www.gnu.org/philosophy/free-sw.html>).

You can get Adobe Acrobat Reader from: <http://www.adobe.com/products/acrobat/readstep2.html>

13.6 Links

It's impossible to cover all the available office type programs in just a few paragraphs; if you need to know more try the links below:

The Linux-Office Site (<http://www.Linux-Office.net/>) is a very useful resource for Linux office apps.

The KOffice website (<http://www.koffice.org/>)

The Gnome-Office website (<http://www.gnome.org/gnome-office/index.shtml>)

Codeweavers Crossover Office (<http://www.codeweavers.com/products/office/>) can run Windows apps like MS Office, Lotus Notes and others under Linux.

14 Gaming With Linux

Okay, you've got your word processor and spreadsheet set up under Linux, plus a web browser, an email client and a million other boring programs. But what about the important stuff? *Where are the games?* Most people probably wouldn't think of Linux as being a gamer's platform, and it's true that the real hard-core gamer might need to stick with a dual boot system for the time being at least. But for the rest of us, Linux can offer a great environment for playing games. There are plenty of good ones on offer, and accelerated 3D is no longer a pain to set up for many common cards. And now that lots of Windows games are playable using emulators like WineX, we've never had so much choice.

14.1 Where to Get Them

Most distros come with a variety of games, and you probably have some already installed. Look in your GNOME or KDE menus under "Games" or "Amusements". If you don't have any installed, check your installation CDs for packages named "kdegames" and "gnome-games". These packages include a wide variety of games ranging from Arcade style games (Tetris and Jezzball clones, Snakerace etc.), board games (Chess, Mahjongg, Reversi and so on), to card games, plus games to test your strategic skills and much more. As well as the ones in the KDE and GNOME packages, some distros include others like Maelstrom, Bzflag (a popular tank game), FrozenBubble (one of my favourites), and even 3D games such as Tuxracer and Chromium. Browse the package directory of your distro CDs to see what's available. There are also lots of games freely available from the internet, plus some commercially produced/ported titles for sale.

14.2 Commercial Games

A few companies make or have made games available to Linux users. Perhaps the best known of these was Loki, who are sadly no longer in business. Loki ported quite a few popular titles to linux, (QuakeIII Arena, HeavyGearII, Descent III etc), and you might even find some of them still available for sale. Probably the best way to find out what's available is to check out online stores like TuxGames (<http://www.tuxgames.com/>).

14.3 Hardware and Software Requirements

Broadly speaking, games can be split into two groups; those that require accelerated 3D support, and those that don't. The first group would include 3D games such as QuakeIII, UnrealTournament, Tuxracer and so on, while the second group includes the 2D style of games such as those found in the GNOME and KDE game packages (and of course the old style text games would be in this group too). Games in the 2D group don't need anything special to run them; if you can run GNOME or KDE you'll have no problems. The 3D games however are much more fussy about what they'll run on; as well as having enough RAM and CPU power, you'll also need a Linux-supported 3D graphics card (or on-board chip). Individual game requirements vary widely, but as rough guide, the recommended minimum for QuakeIII is a 233Mhz CPU with an 8meg graphics card and 64 meg. of RAM. Keep in mind that this is the bare minimum required just to run the game; you'd probably need to double those figures to get reasonable performance.

Setting up 3D graphics with Linux used to be a bit tricky, but now many modern distros will set up the appropriate drivers during installation, giving accelerated 3D out of the box. When you are setting up your machine, keep in mind that it isn't the brand of graphics card you have that is important, but rather the brand of chipset it uses. In other words, you would use ATI drivers for a card with an ATI chipset, regardless of its brand. Currently, most Linux gamers seem to prefer nVidia based cards, and with good reason. NVidia write their own (closed source) drivers for Linux; these are easy to install and set up and their performance is generally on a par with their Windows counterparts. ATI based cards are also popular, and ATI have recently released unified drivers for Linux users with their higher end cards. Check out this site (<http://linuxgamers.net/lsd/hardwaretalk/x42.html>) to see what cards are supported. As well as suitable hardware, you'll also want to use a recent version (>4.0) of XFree86. Later versions have much better 3D support, so if you are having problems an XFree86 upgrade should be one of your first steps.

14.4 Setting up NVidia Based Graphics Cards

As I mentioned earlier, nVidia based cards have become a favourite amongst Linux gamers. While these cards will usually work perfectly out of the box for normal 2D work, you'll probably have to install nVidia's drivers to get accelerated 3D. Some recent distros will install these for you during installation of Linux; even so, you might want to read on so you can update to the current drivers. These drivers are "unified", ie. the same drivers are used for all versions of nVidia based cards. Before you start, you should check that you are running a reasonably recent version of XFree86. There are two drivers that will need to be installed, the NVIDIA_kernel package and the NVIDIA_GLX package. The kernel package is available in several versions to suit most common distros; if there isn't one to suit your distro you can also get tarballs. And if you aren't sure which package to get there is a script you can download from nVidia that will advise you of the best package to use.

Once you've downloaded the packages, you should exit X (not strictly necessary, but it makes recovery easier if things go wrong..) and install the kernel package and then the GLX package. If you are upgrading rather than installing, nVidia recommend removing the old GLX package first instead of upgrading over it. Now all you need to do is edit a couple of lines in your XF86 configuration file (usually this will be /etc/X11/XF86Config-4). Assuming you already have an XF86Config file working with a different driver (such as the 'nv' or 'vesa' driver that is installed by default), then all you need to do is find the relevant Device section and replace the line:

```
Driver "nv" (or Driver "vesa")
```

with

```
Driver "nvidia".
```

In the Module section, make sure you have:

```
Load "glx"
```

You should also remove the following lines:

```
Load "dri"
```

```
Load "GLcore"
```

if they exist. Now restart X to use the new drivers. If you have any problems, check the 'XF86' log file (named ' /var/log/XFree86.0.log' or similar) for clues. Also read the documentation on the nVidia website and in the README file included in the NVIDIA_GLX package.

14.5 Playing Windows Games With Linux

Some well known games produced for windows have Linux binaries available (Return To Castle Wolfenstein etc). The Linux binaries allow you to install the data files from your Windows

game CD, and then run the game directly from Linux. Some games include the Linux binaries on the CD (rare, but hopefully this will become commonplace), or you may have to download them.

Another way to run Windows games is to use an emulator like Wine, or WineX. The list of programs that will run well under Wine is growing steadily, though for gaming you'll probably be more interested in WineX by Transgaming. WineX is a commercial offshoot of the Wine project, and while Wine aims to enable Windows programs in general to be run under Linux, WineX focusses exclusively on games. Many Windows games install and play perfectly with WineX, including Max Payne, Warcraft III, Diablo II, The Sims etc. There is a list of games at the TransGaming (<http://www.transgaming.com/>) website, however I have found that there are some games not listed that will still play under WineX. Try searching Google (<http://google.com>) for *name of game + winex* for help on unlisted games. You can download the WineX source from the CVS tree for free, but compiling and configuring can be confusing for a newbie. Much better is the precompiled packages that are available to subscribers. Subscriptions cost US\$5 per month, with a 3 month minimum. There are some other benefits to subscribers, though I think the binaries alone are worth the price.

14.6 Links

Obviously this has been no more than a very brief overview of Linux gaming; see the sites listed below for more info.

The Linux Gamers HOWTO (<http://www.tldp.org/HOWTO/Linux-Gamers-HOWTO/index.html>) - I can't recommend this one highly enough; if you are serious about gaming with Linux, read this doc!

Linux for Kids (<http://www.linuxforkids.com/>) - This site has lots of links and info about games and educational apps. You don't have to be a kid to enjoy this stuff - adults will probably find some good stuff here too.

The Linux Game FAQ (<http://www.icculus.org/lgfaq/>) - A comprehensive list of Frequently Asked Questions about Linux gaming.

The Linux Game Tome (<http://www.happypenguin.org/news>) - Definitely worth a look!

New Breed Software (<http://www.newbreedsoftware.com/linux/>) - Bill Kendrick and co. have written some good games, mainly for kids.

Racer (<http://www.racer.nl/>) - is a promising race car game with extremely good graphics and physics. Not finished yet, but still playable, and makes a nice change from the shooters.

Transgamings Winex Homepage (<http://www.transgaming.com/>)

LinuxGamers (<http://linuxgamers.net/>) is another interesting game site.

15 Creating Archives and Zipped Files

An *archive* is a single file holding a collection of other files in a structure that makes it possible to retrieve the original individual files (called "members" of the archive). The standard tool for creating and extracting archives in Linux is named 'tar', and it is commonly used in conjunction with a compression tool named 'gzip'. A *Compressed Archive* is an archive which has been compressed (zipped), similar to the compressed archives made by PKZIP or Winzip under Windows.

Why do you need an archive? Archives/compressed archives are useful for taking backups and distributing files over the Internet.

In this chapter we will see how to create and extract an archive and a compressed archive.

15.1 Creating Archives

Tar is used to make archives—it's like PKZIP or Winzip. You have to run the `tar` command from a terminal emulator as in these examples:

To make a new archive:

```
tar -cvf <archive_name.tar> <file> [file...]
```

For example you are in the directory '/home/raj/temp'. This has three files 'barney.txt, betty.txt, fred.txt' and you want to create an archive of these files. Here is how you will do it:

```
tar -cvf flints.tar barney.txt fred.txt betty.txt
```

15.2 Creating Compressed Archives

Using `tar` with the 'z' option compresses the files with gzip and creates a compressed archive (or tarball); this is probably the most common use for tar. For example:

```
tar -czvf flints.tgz barney.txt fred.txt betty.txt
```

will create a compressed archive named 'flints.tgz'.

Note that you can use the tar command on a folder or folders as well as individual files. For example, if I wanted to create a compressed archive containing the folder '/home/raj/rubble/' and all its contents, I would do this (from the directory holding the folder I want to compress/archive):

```
tar -czvf rubble/
```

15.3 Extracting Files from Archives

To extract files from an archive use:

```
tar -xvf <archive_name.tar> [filename.tar...]
```

For example to extract all the files of the archive 'flints.tar' (created in the previous section)

```
tar -xvf flints.tar
```

This will extract the files 'fred.txt barney.txt betty.txt'

To extract only 'fred.txt'

```
'tar -xvf flints.tar fred.txt'
```

This will extract the files 'fred.txt'

15.4 Extracting Compressed Archives (Tarballs)

You can use the 'z' option to extract compressed archives - these are the most common and have .tgz or .tar.gz filename extensions. For example:

```
tar -zxvf flints.tgz
```

15.5 Viewing Archive Contents

What if you only want to view the contents of the archive without extracting them.

```
tar -tf <file.tar.gz>
```

lets you do that.

For example to view the members of the archive 'flints.tar'

```
tar -tf flints.tar
```

This will show you the following

```
fred.txt
```

```
barney.txt
```

```
betty.txt
```

15.6 Compressing Files with gzip

Gzip reduces the size of the files. Whenever possible, each file is replaced by one with the extension .gz. Though there are other file compressing utilities like *compress*, files compressed using *gzip* are most common.

The command `gzip <file>` compresses the file. For example if I have two files '/home/crypto/temp/foo.txt' and '/home/crypto/temp/bar.txt'. Here is how to compress them

```
gzip /home/crypto/temp/foo.txt /home/crypto/temp/bar.txt
```

If you are in the directory '/home/crypto/temp/' you do not have to specify the full path to the file.

The files are replaced with one with the extension .gz. For example, after running the above example you will find 'bar.txt.gz, foo.txt.gz' instead of 'bar.txt' and 'foo.txt'.

Gzip takes following command line options (among other)

- v Verbose. Display the name and percentage reduction for each file compressed.
- h Print an informative help message describing the options then quit.

You might have noticed that *gzip* does not lump the files together. Each file is compressed separately. What if you want all your files in one single archive? Just use the 'z' option with *tar* as mentioned earlier.

15.7 Zip and Unzip

Compressed archives in the Linux world are mainly of the tgz or tarball type, while Windows users are probably more familiar with the 'zip' format. Despite this, it's quite easy to share archives with Windows users. The tool for unzipping or expanding zip files has the logical name of 'unzip', and it's installed by default with most distros. *Unzip* has a few options (see 'man zip'). but usually it's just a matter of using `unzip filename.zip`. For example, assuming I have a file named 'betty.zip' into my home directory. I'd use this to unzip it:

```
unzip betty.zip
```

You can also create zip files under Linux with a tool named 'zip'. Unlike *unzip*, *zip* isn't always installed by default, though you should be able to find the zip package on your installation CDs.

Zip has quite a few options (see ‘`man zip`’), but usually you’ll only need to use `zip filename file1 file2 file3 etc..`. For example, if I wanted to create a zip archive named ‘`rock.zip`’ that contained the files ‘`dino.txt bambam.txt wilma.txt`’, I’d do this:

```
zip rock dino.txt bambam.txt wilma.txt
```

Note that zip will add the `.zip` extension to the filename. Zip can compress folders and their contents, eg:

```
zip myzip myfiles/*
```

It will also compress folders recursively with the ‘`-r`’ option. Files created with zip are compatible with Windows tools such as WinZip and PKZIP, and files created with WinZip/PKZIP etc can be expanded with unzip. I should also mention that reasonably recent versions of WinZip etc can unpack `.tgz` archives, though if you are making archives for Windows users you might want to stick with the zip format to avoid confusion.

15.8 More on Archives

- The `info` pages for `tar` comes with a tutorial which teaches the even most intricate details in an easy to follow manner.(see [Section 8.2 \[Info Pages\]](#), page 31 to learn howto use Info)
- The `info` page for `gzip` is also good to learn more.

16 The Post-Install Tune-up

Many Linux distributions install lots of stuff that many people never use. Often there are services or daemons set to run that aren't needed, and system configuration settings are conservative so as to run on the widest range of hardware. All these things can detract from your Linux desktops performance, and that's what this article is about; getting more performance from your box by performing a post-install tune-up. Here's what we'll be doing:

- Tidying up the boot-up process
- Freeing resources by removing un-needed services/daemons from boot scripts
- Freeing hard disk space
- Speeding up hard disk performance with `hdparm`
- Recompiling the kernel for better system performance

You don't need to do **all** these things by the way, only those that you need or want to, the rest are optional. Be aware that if your box is reasonably configured to begin with (and most distributions are, out of the box), you are unlikely to get a dramatic improvement from any one of these things.

However, by doing some or all of them you should end up with a system that boots more quickly, has more disk space and slightly more free memory, and a small but noticeable improvement in performance. The one thing you can do that will have a profound effect on performance is run lightweight, efficient software, so you should make that your first priority when building a fast Linux desktop.

16.1 Disclaimer

I don't guarantee the accuracy of anything that follows, so use this information at your own risk. In other words, if by following this guide you trash your computer, don't blame me.

16.2 Before You Start

I'll assume you'll be running a *SysV type system*, as this is the most common and what you'll have if you are running a Red Hat-type distribution. SysV simply refers to the way services etc are started at boot time. If you are running some other system, you can still clean up the boot process; check your distribution's documentation for details. It's a good idea to browse through any documentation that came with your distribution anyway. This might be in the form of HTML files or a printed manual, and with many modern distributions is very comprehensive. The documentation should be able to provide you with details of any variations to the boot process used by your particular distribution, though I think the common distributions are pretty much all the same in this regard.

If this is a fresh installation, you should make sure all your hardware is properly configured first. Linux has really come a long way as far as hardware recognition goes, and chances are you won't have to do anything, though things like sound cards sometimes have to be setup manually. Once you are sure everything is going to work, you can continue with the tuning

16.3 The Boot-Up

We probably should start with a brief description of what actually happens when you boot your Linux box. You can skip this bit if you want to, but I think an understanding of what goes on at boot-up can often be helpful, so stick around

After the kernel is kicked into life by GRUB or LILO or whatever, the following steps occur (with possible minor variations):

1. The kernel gets it's own internal systems set up.
2. The `init` program is started.
3. `init` reads the `/etc/inittab` file. This file provides `init` with the default run level for the system (eg console, graphical, single-user etc.) Take a look at `/etc/inittab` yourself so you understand what the various run levels do.
4. `init` then runs a script (usually `/etc/rc.d/rc.modules`) to load auto loaded kernel modules.
5. Depending on the default run level (from `/etc/inittab`), `init` then starts or stops all the services in that particular run level directory. For example, if **runlevel 5** is the default according to `/etc/inittab`, then all the scripts in `/etc/rc.d/rc5.d/` are run.
6. `init` then runs another script (usually `/etc/rc.d/rc.local`) . This is where the user can put stuff that he/she wants to be started automatically at boot-up. You might want to start your OSS sound driver from here for example. Users of older versions of Mandrake can edit this file to get rid of that damn ugly penguin

This is a bit over-simplified, but I hope you get the idea. If you take a look at the scripts in `/etc/rc.d/rc5.d` (or whatever your default runlevel is), you'll see that the names of the scripts all start with an **S**(for Start) or a **K**(for Kill or stop) followed by a number. The number determines the order in which the scripts are run. Most distributions start a diverse range of services or daemons at boot time, and while this automatically covers the needs of the majority of users, it also means that there will probably be several processes started that aren't required. This results in longer boot-up times, increased memory usage, and more potential security holes. Stripping un-needed stuff from the start-up scripts is easy; the hard part is determining what does what, and what you do and don't need. Hopefully the listing below will be of some help, it notes some of the most commonly found services and gives a brief description of what they do. And don't forget to make backups or notes of your changes, just in case you find you really did need to have that daemon started after all (this list is courtesy of Stan and Peter Klimas' Linux Newbie Administrators Guide)

<code>anacron</code>	checks <code>cron</code> jobs that were left out due to down time and executes them. Useful if you have <code>cron</code> jobs scheduled but don't run your machine all the time— <code>anacron</code> will detect that during bootup.
<code>amd</code>	automount daemon (automatically mounts removable media).
<code>apmd</code>	Advanced Power Management BIOS daemon. For use on machines, especially laptops, that support <code>apm</code> .
<code>arpwatch</code>	keeps watch for ethernet/ip address pairings.
<code>atd</code>	runs jobs queued by the "at" command.
<code>autofs</code>	control the operation of automount daemons (competition to <code>amd</code>).
<code>bootparamd</code>	server process that provides information to diskless clients necessary for booting.
<code>crond</code>	automatic task scheduler. Manages the execution of tasks that are executed at regular but infrequent intervals, such as rotating log files, cleaning up <code>/tmp</code> directories, etc.

<code>cupsd</code>	the <i>Common UNIX Printing System (CUPS)</i> daemon. CUPS is an advanced printer spooling system which allows setting of printer options and automatic availability of a printer configured on one server in the whole network. The default printing system of Linux Mandrake.
<code>dhcpd</code>	implements the Dynamic Host Configuration Protocol (DHCP) and the Internet Bootstrap Protocol (BOOTP).
<code>gated</code>	routing daemon that handles multiple routing protocols and replaces <code>routed</code> and <code>egppup</code> .
<code>gpm</code>	useful mouse server for applications running on the Linux text console.
<code>httpd</code>	daemon for the Apache webserver.
<code>inetd</code>	listens for service requests on network connections, particularly dial-in services. This daemon can automatically load and unload other daemons (<code>ftpd</code> , <code>telnetd</code> , etc.), thereby economizing on system resources. Newer systems use <code>xinetd</code> instead.
<code>isdn4linux</code>	for users of ISDN cards.
<code>kerneld</code>	automatically loads and unloads kernel modules.
<code>klogd</code>	the daemon that intercepts and displays/logs the kernel messages depending on the priority level of the messages. The messages typically go to the appropriately named files in the directory <code>/var/log/kernel</code> .
<code>kudzu</code>	detects and configures new or changed hardware during boot.
<code>keytable</code>	loads selected keyboard map.
<code>linuxconf</code>	the linuxconf configuration tool. The automated part is run if you want linuxconf to perform various tasks at boottime to maintain the system configuration.
<code>lpd</code>	printing daemon.
<code>mcserv</code>	server program for the Midnight Commander networking file system. It provides access to the host file system to clients running the Midnight file system (currently, only the Midnight Commander file manager). If the program is run as root the program will try to get a reserved port otherwise it will use 9876 as the port. If the system has a portmapper running, then the port will be registered with the portmapper and thus clients will automatically connect to the right port. If the system does not have a portmapper, then a port should be manually specified with the <code>-p</code> option (see below).
<code>named</code>	the Internet Domain Name Server (DNS) daemon.
<code>netfs</code>	network filesystem mounter. Used for mounting <code>nfs</code> , <code>smb</code> and <code>ncp</code> shares on boot.
<code>network</code>	activates all network interfaces at boot time by calling scripts in <code>/etc/sysconfig/network-scripts</code> .
<code>nfsd</code>	used for exporting <code>nfs</code> shares when requested by remote systems.
<code>nfslock</code>	starts and stops <code>nfs</code> file locking service.
<code>numlock</code>	locks numlock key at init runlevel change.
<code>pcmcia</code>	generic services for pcmcia cards in laptops.
<code>portmap</code>	needed for Remote Procedure Calls. Most likely, you need it for running network.
<code>postfix</code>	mail transport agent which is a replacement for <code>sendmail</code> . Now the default on desktop installations of Mandrake (RedHat uses <code>sendmail</code> instead).

<code>random</code>	saves and restores the "entropy" pool for higher quality random number generation.
<code>routed</code>	daemon that manages routing tables.
<code>rstatd</code>	kernel statistics server.
<code>rusersd, rwalld</code>	identification of users and "wall" messaging services for remote users.
<code>rwhod</code>	server which maintains the database used by the <code>rwho(1)</code> and <code>ruptime(1)</code> programs. Its operation depends on the ability to broadcast messages on a network.
<code>sendmail</code>	mail transfer agent. This is the agent that comes with Red Hat.
<code>smbd</code>	the SAMBA (or smb) daemon, a network connectivity services to MS Windows computers on your network (hard drive sharing, printers, etc).
<code>squid</code>	An http proxy with caching. Proxies relay requests from clients to the outside world, and return the results. You would use this particular proxy if you wanted to use your linux computer as a gateway to the Internet for other computer on your network. Another (and probably safer at home) way to do it, is to set up masquarading.
<code>syslogd</code>	manages system activity logging. The configuration file is <code>‘/etc/syslog.conf’</code> .
<code>smtpd</code>	Simple Mail Transfer Protocol, designed for the exchange of electronic mail messages. Several daemons that support SMTP are available, including <code>sendmail</code> , <code>smtpd</code> , <code>rsmtpd</code> , <code>qmail</code> , <code>zmail</code> , etc.
<code>usb</code>	daemon for devices on Universal Serial Bus.
<code>xfs</code>	X font server.
<code>xntpd</code>	finds the server for a NIS domain and stores the information about it in a binding file.
<code>yppbind</code>	NIS binder. Needed if computer is part of Network Information Service domain.

Many users will find they have a lot of unnecessary stuff in their `/etc/rc.d/rc*.d/` folders. If you aren't sure if you need something or not, just move it somewhere else temporarily (but **don't delete** it), re-boot and see how things go. If you find that you do need it, just move it back and re-boot.

16.4 How To Make The Changes

I usually just fire up a file manager and make a new folder in `/etc/rc.d` called `JunkFromRc5` or something similar. Then I just drag the unneeded scripts from `/etc/rc.d/rc5.d/` into the new folder (obviously my default runlevel is 5, you might need to use something different..). Alternatively, you can use a graphical tool like `tksysv`, or perhaps your distribution has it's own tool. You might also want to edit `‘/etc/rc.d/rc/local’`. Apart from user's entries, this often has a few lines that overwrite `‘/etc/issue’` with some system specs (and/or that hideous penguin), and the contents of the `‘/etc/issue’` file are displayed just before the login screen. Many people prefer to delete this bit and insert a line to display a fortune here instead, eg. `/usr/games/fortune > /etc/issue`. As usual, if you aren't sure what you are doing, make a backup copy of the file first.

16.5 Re-Claiming Hard Disk Space

This bit is easy, if a little time consuming. I usually start by getting rid of unneeded software packages. Fire up your favorite package management tool (like `kpackage`) and spend some time browsing through the list of installed programs. Tools like `kpackage` are ideal for this kind of work as they can easily show you the size of each package, a summary (so you know what the package is for), and any related dependencies.

Do you really need six editors, four file managers, five shells, three ftp clients etc.? Don't be surprised to get rid of a hundred megs or more of stuff. Packages like the Tex related ones, Emacs/Xemacs, and various emulators are never used by many, yet they occupy lots of space. If you are doubtful about removing some packages, keep some notes so you can re install them later if you have to.

Many distributions also install lots of documentation (check out `/usr/doc` or `/usr/share/doc`). You'll probably find that there are only a few files in there worth keeping, and remember most of this stuff is available on the Web anyway. The `du` tool is invaluable for finding disk hogs. Also look for core files left over from crashes; these are only really useful to debuggers and can be deleted.

16.6 Hard Disk Tuning

I've seen a few articles claiming huge performance gains from using `hdparm`, a command line tool for setting (IDE) hard disk parameters. The claimed gains are sometimes in the order of several hundred percent. While I'm not doubting these figures, I have to wonder whether they just indicate that the disk was horribly mis-configured to start with. I've tried using `hdparm` on a few disks, and have found modest gains in performance. Keep in mind that disk performance is only one factor in overall system performance, and even a fairly big jump in disk performance might not make a perceptible difference to the overall speed of your system. If a disk was noticeably slow, I would certainly give `hdparm` a try, but otherwise I wouldn't worry about it too much. Some common distributions already set some of the optimizing parameters at boot time anyway, so as I said before, unless you think there is a problem, you could probably just leave well enough alone. If you do decide to give it a go, make sure you read (and understand) the man page (at a terminal emulator type `man hdparm`), and be aware that with some of the adjustments there is a small but real risk that things can go spectacularly wrong, ie. corruption of the file system. If you'd like to give `hdparm` a try, here's the basic usage:

```
hdparm [-flag] device
```

Running `hdparm` without any flags (or with the `-v` flag) will display the current settings. To see the current settings for my first hard disk (`/dev/hda`) for example, I would use: `hdparm /dev/hda`. To do a basic check of the speed of the first hard disk I would use: `hdparm -Tt /dev/hda`. Some more commonly used flags:

- `-c3` Enables IDE 32 bit I/O support
- `-a [sectcount]`
 Get/set sector count for read ahead
- `-m16` Sets multi-sector I/O (in this example 16 sectors, you may need to experiment to find the optimal number for your disk)
- `-u1` Unmasks interrupts
- `-d1 -X34` Unable DMA mode2 transfers
- `-d1 -X66` Enable UltraDMA mode2 transfers

Read the man page for more options.

I guess the logical way to use `hdparm` would be to find out what your disk supports, then set `hdparm` accordingly. More commonly though, trial and error is used, changing one setting at a time and measuring the performance after each change. Don't use settings recommended by someone else ; while they may have worked perfectly on that persons disk, your disk might be completely different and the results may not be good. There are several tools available for testing disk performance, one of the better known ones is *bonnie*. And remember the changes will be lost when you re-boot, so if you want to make them permanent, you'll have to add them to a boot script like `‘/etc/rc.d/rc.local’`.

16.7 Filesystem

Linux updates a *last access time* attribute every time you open a file. If you are after all the speed you can get, and you are sure you don't need this feature, you can add `noatime` to the mount options listed in `‘/etc/fstab’`. For example:

In the `‘/etc/fstab’` file add the line `/dev/hda5/ ext2 defaults,noatime 11` if you do not wish to update last access time on the files in `/dev/hda5` **partition**.

16.7.1 Alternative Filesystems

You might have tried (or read about) alternatives to the traditional ext2 file system, and at present the most common seem to be ReiserFS and Ext3. These have some advantages over ext2, including quicker performance, so if you are about to start a new Linux installation you should certainly consider using the Reiser file system. However, as with `hdparm`, unless you are doing something unusually disk intensive, the gains are likely to be minor, and if your current system is doing the job I'd stick with it.

16.8 Kernel Recompile

This is another one of those things that is often recommended in the Linux tune-up guides. While it may have been important years ago, it is probably of questionable value now that modular kernels are the norm. So unless you need to compile in some special feature, or you are using a pre-historic non-modular kernel (in which case you could probably benefit from updating your Linux installation), I wouldn't bother. Most recent distributions come with a variety of optimized kernels, and automatically install the one that best suits your system. Of course, you might want to recompile just for its sheer geek entertainment value, and I guess that's as good a reason as any.... I won't go into the details of kernel compilation here, check your distribution's documentation or the Kernel How-To (<http://www.linuxdoc.org/HOWTO/Kernel-HOWTO.html>) for details.

16.9 Miscellaneous Tips

This article has only covered the very basic stuff - if you are interested in reading some much more detailed info about configuring Linux read the Configuration How-To (<http://www.linuxdoc.org/HOWTO/Config-HOWTO/index.html>).

If you are serious about tuning your Linux box, you'll need some benchmarking tools. To get started, take a look at these two sites:

The Linux Benchmarking Project and bench (<http://www.tux.org/bench/>).

Obviously you'll be aiming to conserve memory as much as possible. Use the `free` command from a terminal emulator to see memory usage details. Ideally, you'll be able to balance usage against available memory so that swap isn't used.

You can save some memory by using a plain background on your desktop, rather than an image file.

Other useful tools are `ps -aux` (shows details of running processes), and `top` (similar to `ps` but continually updates).

Help reduce the time it takes X to update the screen on low-end machines by not using a greater colour depth than necessary, eg. use 16bit instead of 32 bit. You can check X's performance with `x11bench`, which is often installed by default.

17 Recovering from Crashed Programs

Linux is renowned for its stability, and it's true that the operating system itself seldom crashes or hangs. The programs you run on it however, might not be so sturdy, and some programs can and do crash. Unlike some other operating systems, Linux itself is rarely affected by crashing programs. Recovery is normally fairly easy, and reboots are very rarely required.

17.1 Terminating Misbehaving Programs

17.1.1 XKill

XKill is a handy tool for terminating hung or frozen apps that refuse to respond to the File>Quit button, or the windows close button. Simply start `xkill`, either from your menus or an Xterminal window, and the mouse cursor will change into something resembling a gunsight. Move the cursor over the offending window and click on it; this will usually shut down the problem program.

17.1.2 The `ps` and `kill` Commands

The `ps` command is an extremely useful tool that is run from the command line to list all the processes currently running. It can be used in conjunction with the `kill` command to stop crashed or runaway programs. To use it, open an xterminal and type in: `ps -aux`. The `-aux` option just makes `ps` show some additional information. You'll be presented with a list of all the running processes, their owners, their resource usage and so on. On the extreme left side you'll see the identifying number (or PID) of each process. This number can be used with the `kill` command to stop rogue programs.

As an example, we'll assume Netscape has crashed, and can't be shutdown normally. Running `ps -aux` shows up Netscape in the list of running processes, along with Netscapes PID or process identifying number (let's say the PID is 2345). Typing in this:

```
kill 2345
```

should stop Netscape immediately. Run `ps -aux` again to see if it has indeed been killed. What if it's still there? Run the `kill` command again, but this time with the `-9` option, eg.

```
kill -9 2345
```

 should do the trick.

The `kill` command (without the `-9`) should always be tried first, as it gives the process a chance to close any files it may have open and generally tidy up after itself, whereas `kill -9` is not so gentle and just kills the process instantly. Note that the 2345 PID I've used is just an example, and that each process is allocated a PID number when it is started.

There is another version of `kill`, and that is `killall`. The `killall` command is used with the processes name, rather than PID. For example:

```
killall netscape
```

or

```
killall -9 netscape
```

could be used when you know the processes name.

The `ps -aux` command is also handy for spotting resource hogs. Let's say for example you are working at your computer when you notice things are slowing to a crawl. Run `ps -aux` and check the columns showing the `cpu` and `memory` usage to spot the offending program. You can then use `kill` to terminate it if you wish.

17.1.3 Stopping Command Line Programs

Occasionally you may start a program from the command line, only to find that it hangs, or runs away eating up resources. In these cases, you can usually stop it simply by hitting `<ctrl-c>`, ie. holding down the Ctrl and the c key at the same time.

17.2 X Related Problems

In Linux, the GUI system (ie. the graphical display, mouse cursor etc) are commonly referred to simply as X. What do you do if a program hangs or crashes, and takes your mouse cursor with it? You can't use `xkill`, and you mightn't be able to open a terminal window either. What you can do is hit `Alt-Ctrl-Backspace` to restart X and regain mouse control. If you do this, you should probably run `ps -aux` when X restarts just to make sure the crashed program didn't leave any leftover processes. These should be killed before restarting them in the new X session.

17.3 Rebooting - The Last Resort

Unlike other systems, Linux very rarely requires a reboot. However, if a program locks up so badly that neither the mouse or keyboard respond, rebooting may be the simplest solution. If you can somehow reboot normally then this is preferable to simply hitting the power button, but sometimes you just have no other option. For those using a journalling filesystem such as `ext3` or `reiserfs` (these are often installed by default with recent distros) this shouldn't cause any problems; `ext2` users will probably find the system will complain about the rude shutdown and automatically run a filesystem check at the next bootup. Usually though, everything will be OK.

Index

- /
- /dev directory 10
- /proc directory 11

- 2**
- 2D games 59

- 3**
- 3D games 59

- A**
- Abiword 54
- Acknowledgments 1
- address books 55
- AIM 29
- archive 63
- archive-extract 63
- archive-view 63
- audience of this howto 1

- B**
- Basic concept of IRC 28
- bladeenc 42
- block devices 10
- boot up 68

- C**
- cdparanoia 42
- changing directories in MC 20
- character devices 10
- Chatting / Instant Messaging with Linux 27
- chatting history 27
- codecs 47
- commercial games 59
- compressed archives 63
- Compressed Archives, Creating 63
- Compressing files 64
- Configuring Xchat 28
- Connecting to Internet 23
- Connector 55
- copying files in GMC 17
- copying files in Konqueror 19
- copying files in MC 21
- copyright 1

- Corel WordPerfect 54
- crash recovery 75
- creating Archives 63

- D**
- databases 56
- deleting files in GMC 17
- deleting files in Konqueror 18
- deleting files in MC 21
- directory 9
- documentation 31
- downgrading packages 34
- Drag and Drop in GMC 17
- drawing/graphics 56

- E**
- email clients 55
- Everybuddy 30
- Evolution 55
- Extracting compressed archives 64
- extracting files from archives 63
- Extracting tarballs 64

- F**
- fax programs 55
- file formats, video 47
- Flash 50

- G**
- GAIM 30
- game links 61
- games 59
- games hardware requirements 59
- games software requirements 59
- getting games 59
- GMC 13, 17
- GMC in GNOME 13
- GMC in Konqueror 13
- GNOME Panel 5
- gnome-help-browser 32
- gnome-terminal 6
- Gnucash 56
- Gnumeric 55
- gtcd 41
- gzip 64

H

Hard Disk Tuning	71
hdparm	71
help on no documentation	31
history of chatting	27
how to use RPMs	33

I

ICQ	29
info page	31
Installing packages	33
installing printer	39
Installing Printer	39
Instant Messaging	29
Internet Relay Chat	28
introduction	1
IRC	28
IRC - basic concept	28

K

KDE Panel	5
Kernel Recompilation	72
Kicker	5
kill	75
killall	75
kit	29
Konqueror	14
konsole	6
kscd	41
kspread	55
Kword	54

L

Linux User Group	32
LUGs	32

M

mailing guidelines	1
man pages	31
Managing files in Konqueror	18
media players	49
more documentation	31
mouse support in MC	20
Moving files in GMC	18
moving files in Konqueror	19

moving files in MC	21
MPlayer	49
MS Exchange	55
MS Outlook	55
mtv / mtvp	50
multimedia, video	47
music CDs	41

N

name of devices	10
name of hard disks	10
navigating in MC	20

O

Office Applications	53
Ogg-Vorbis	42
Ogg-Vorbis Format	43
OpenOffice Writer	53
organizers	56

P

PDF Files Viewer	57
PIMs	56
Playing and Recording Music	41
Playing Audio Cds	41
Playing MP3 Files	41
Playing Windows Games With Linux	60
presentation apps	56
printer	39
printing in linux	39
Programs FAQs	31
ps	75

R

Re-Claiming Hard Disk Space	71
reading info	31
reading info from konqueror	32
reading info using gnome-help	32
reading man from konqueror	32
reading man pages	31
reading man using gnome-help	32
RealPlayer	50
rebooting	76
Recovering from Crashes	75
RPM	33

RPM Version Incompatibilities	36
Run Program dialog	5
Run Program dialog box	5
rxvt	6

S

searching files in Konqueror	19
Searching in GMC	18
selecting files in konqueror	18
selecting files in MC	21
selecting multiple files in GMC	17
Setting up kppp from KDE	23
Setting up NVidia Based Graphics Cards	60
Shockwave	50
shred	18
Software Package	33
Spreadsheets	54
StarOffice Star Writer	53
StarOffice/OpenOffice Calc	55
starting Xmms	42
streaming media	47
switching directory panels in MC	20
SysV type system	67

T

talk	27
tar	63
terminal emulators	6
the root directory	9
Trash	18
Tuning your GNU/Linux system	67

U

unzip	64
Updating installed package	34
using data CDs	41

Using MC for browsing RPMs	37
Using menus	5
using menus in MC	20
using mouse in MC	20
Using Run Program Dialog Box	5
Using terminal emulators	5

V

view the members of the archive	63
vorbis- tools package	43

W

wav files to mp3	42
What is RPM	33
which printer to be used	39
wild cards	18
Word Processing	53

X

X - crash	76
X - restarting	76
Xanim	50
Xchat - Configuring	28
Xine	49
xkill	75
xmms	41
Xmms	42
XMMS (video)	49
XMovie	49

Y

Yahoo Messenger	29
-----------------------	----

Z

zip	64
-----------	----

Table of Contents

Preface	1
Intended audience	1
Mailing the Authors	1
Our Distribution	1
Copyrights and Other Delusion	1
Acknowledgments	2
Standard Disclaimer	2
1 Some Dos and Donts	3
1.1 Do's	3
1.2 Important Donts	3
2 Getting Started	5
2.1 Using Menus	5
2.1.1 Using Menus in GNOME	5
2.1.2 Using Menus in KDE	5
2.2 Using the Run Program Dialog Box	5
2.2.1 Run Program Dialog Box in GNOME	5
2.2.2 Run Program Dialog Box in KDE	5
2.3 Using the Terminal Emulator	6
2.3.1 Terminal Emulator in GNOME	6
2.3.2 Terminal Emulator in KDE	6
2.4 Notes on Terminal Emulators	6
3 Understanding Linux Directories and Files	9
3.1 /: The root directory	9
3.2 The Directory Structure	9
3.3 /dev:The Directory to Devices	10
3.4 /proc: The Directory of Processes	11
3.5 More Information on Linux Directory	11
4 Files Managers	13
4.1 GMC	13
4.1.1 Starting GMC from Menu	13
4.1.2 Starting GMC from 'Run Program'	13
4.1.3 Starting GMC from Terminal Emulator	13
4.1.4 Starting GMC from Menu KDE	13
4.1.5 Starting GMC from 'Run Program' and terminal emulator in KDE	13
4.2 Konqueror	14
4.2.1 Starting Konqueror from Menu in GNOME	14
4.2.2 Starting Konqueror from 'Run Program' in GNOME	14
4.2.3 Starting Konqueror from a terminal emulator in GNOME	14
4.2.4 Starting Konqueror from Menu in KDE	14
4.2.5 Starting Konqueror from 'Run Program' in KDE	14
4.2.6 Starting Konqueror from terminal emulators in KDE ..	14
4.3 MC	14

4.3.1	Starting Midnight Commander	15
4.4	Notes on File Managers	15
5	Managing Files	17
5.1	Managing Files using GMC	17
5.1.1	Selecting files	17
5.1.2	Deleting files	17
5.1.3	Copying Files	17
5.1.4	Drag and Drop for Copying	17
5.1.5	Moving Files	18
5.1.6	Searching for Files	18
5.2	Managing Files using Konqueror	18
5.2.1	Selecting Files	18
5.2.2	Deleting Files	18
5.2.3	Copying Files	19
5.2.4	Moving Files	19
5.2.5	Searching Files	19
5.3	Managing Files using Midnight Commander	20
5.3.1	Using Mouse	20
5.3.2	Navigating	20
5.3.3	Using Menus	20
5.3.4	Selecting Files	21
5.3.5	Deleting Files	21
5.3.6	Copying Files	21
5.3.7	Moving Files	21
5.3.8	Searching Files	21
6	Connecting to the Internet	23
6.1	Hardware Requirements	23
6.2	Setting up kppp from KDE	23
6.2.1	Solving Permission Restrictions with kppp under RedHat	24
7	Chatting / Instant Messaging with Linux	27
7.1	History of chatting	27
7.2	talk	27
7.3	IRC or Internet Relay Chat	28
7.3.1	Basic concept of IRC	28
7.3.2	Software for IRC	28
7.3.3	Configuring Xchat	28
7.4	Instant Messaging	29
7.4.1	ICQ	29
7.4.2	Yahoo!	29
7.4.3	AIM	29
7.5	All-In-One IM Clients	30
7.5.1	Everybuddy	30
7.5.2	GAIM	30
8	Finding Help in GNU/Linux	31
8.1	Man Pages	31
8.2	Info Pages	31
8.3	More Documentation	31
8.4	Programs with no Documentation	31
8.5	Getting help From LUGs	32
8.6	Tips on info and man	32

9	Installing and Managing Software with RPM..	33
9.1	What is RPM?	33
9.2	What is a Software Package?	33
9.3	Using RPM	33
9.4	Installing Packages	33
9.5	Updating Packages	34
9.6	Downgrading Packages	34
9.7	Uninstalling Packages	34
9.8	Querying packages	34
9.8.1	Listing all of the files installed by a package	34
9.8.2	Listing a description of an installed package	35
9.8.3	To find the package to which a file belongs	35
9.8.4	To list all the packages a package depends on	35
9.9	Resolving Dependency Problems	35
9.10	Using RPMFind and RPMBone	35
9.11	Circular Dependencies	36
9.12	Library Version Problems	36
9.13	Automatic Dependency Resolution Tools	36
9.14	Miscellaneous	36
9.14.1	RPM Version Incompatibilities	36
9.14.2	Midnight Commander	37
9.14.3	Learning More	37
10	Printing	39
10.1	Which Printer to Use	39
10.2	Connecting the Printer to the Computer	39
10.3	Installing Printer	39
10.3.1	Configuration Program	39
11	Multimedia Part 1 - Playing and Recording Music	41
11.1	The Basics	41
11.2	Playing Audio Cds	41
11.3	Playing MP3 Files	41
11.4	Using XMMS	42
11.5	Recording (or ripping) Tracks from Audio CDs	42
11.6	Converting .wav files to mp3	42
11.7	The Ogg-Vorbis Format	43
11.8	Converting .mp3 files into .wav format	43
11.9	To Normalize a group of .wav files	43
11.10	Recording (or burning) an Audio CD	44
11.11	Recording a Data CD (mp3 or ogg)	44
11.12	Some detailed information on related topics:	44

12	Multimedia Part 2 - Linux and Video	47
12.1	Introduction	47
12.2	File Formats and CODECs	47
12.2.1	File Formats	47
12.2.2	Codecs	47
12.3	Streaming Media	47
12.4	Some Common File Formats	48
12.4.1	mpg (or mpeg)	48
12.4.2	AVI (Audio Video Interleaved)	48
12.4.3	DivX (.avi)	48
12.4.4	Mov (Quicktime)	48
12.4.5	RealMedia Video (.rm)	48
12.4.6	Windows Media Video (.wmv)	48
12.4.7	ASF	48
12.4.8	Shockwave Flash (Flash, .swf)	49
12.5	Linux Media Players	49
12.5.1	KDE Media Player (noatun)	49
12.5.2	Xine	49
12.5.3	XMovie	49
12.5.4	XMMS	49
12.5.5	MPlayer	49
12.5.6	Mtv / mtvp	50
12.5.7	RealPlayer	50
12.5.8	Xanim	50
12.5.9	Macromedia Flash Player	50
12.6	Miscellaneous	51
12.6.1	Additional Resources	51
13	Linux in the Office	53
13.1	Word Processing	53
13.1.1	StarOffice Star Writer	53
13.1.2	OpenOffice Writer	53
13.1.3	AbiWord	54
13.2	Others Word Processors	54
13.2.1	Kword	54
13.2.2	Corel WordPerfect	54
13.3	Spreadsheets	54
13.3.1	Gnumeric	55
13.3.2	StarOffice/OpenOffice Calc	55
13.3.3	kspread	55
13.4	Other Office Applications	55
13.4.1	Address Books	55
13.4.2	Fax Apps	55
13.4.3	Email/PIM	55
13.4.4	Drawing/Graphics	56
13.4.5	Financial	56
13.4.6	Database	56
13.4.7	Presentation Apps	56
13.4.8	Organizers	56
13.4.9	Calculators	56
13.5	PDF Files Viewer	57
13.5.1	XPdf	57
13.5.2	Adobe Acrobat Reader	57
13.6	Links	57

14	Gaming With Linux	59
14.1	Where to Get Them	59
14.2	Commercial Games	59
14.3	Hardware and Software Requirements	59
14.4	Setting up NVidia Based Graphics Cards	60
14.5	Playing Windows Games With Linux	60
14.6	Links	61
15	Creating Archives and Zipped Files	63
15.1	Creating Archives	63
15.2	Creating Compressed Archives	63
15.3	Extracting Files from Archives	63
15.4	Extracting Compressed Archives (Tarballs)	64
15.5	Viewing Archive Contents	64
15.6	Compressing Files with gzip	64
15.7	Zip and Unzip	64
15.8	More on Archives	65
16	The Post-Install Tune-up	67
16.1	Disclaimer	67
16.2	Before You Start	67
16.3	The Boot-Up	68
16.4	How To Make The Changes	70
16.5	Re-Claiming Hard Disk Space	71
16.6	Hard Disk Tuning	71
16.7	Filesystem	72
16.7.1	Alternative Filesystems	72
16.8	Kernel Recompilation	72
16.9	Miscellaneous Tips	72
17	Recovering from Crashed Programs	75
17.1	Terminating Misbehaving Programs	75
17.1.1	XKill	75
17.1.2	The ps and kill Commands	75
17.1.3	Stopping Command Line Programs	76
17.2	X Related Problems	76
17.3	Rebooting - The Last Resort	76
	Index	77

