

MandrakeSecurity

User Guide

MandrakeSoft

June 2001

<http://www.linux-mandrake.com/>

MandrakeSecurity: **User Guide**
by MandrakeSoft

Copyright © 1999, 2000, 2001 by **MandrakeSoft**

Table of Contents

Preface	13
1. Note from the Editor	13
2. Legal Notice	13
3. Authors and Translators	14
4. Tools Used in the Making of this Manual	15
5. About Linux-Mandrake	15
5.1. Contact Mandrake community	15
5.2. Support Mandrake	16
1. Introduction to the User Guide	19
I. Installation guide	21
2. Some words before you begin the installation	23
2.1. Welcome!	23
3. Before setup	25
3.1. Configuring your <i>BIOS</i>	25
3.2. Creating a “bootdisk”	25
3.3. Supported hardware	30
4. Disks and partitions	37
4.1. Structure of a hard disk	37
4.2. Conventions for naming the disks and partitions ..	41
5. Installation with <i>DrakX</i>	45
5.1. Introduction to the <i>MandrakeSecurity</i> installer ..	45
5.2. Disk detection and configuration	47
5.3. Configuring the Keyboard	48
5.4. Miscellaneous options	49
5.5. Partitioning your hard disk	50
5.6. Formatting partitions	54
5.7. Packages installation	55
5.8. Internal Network configuration	56
5.9. External Network configuration	60
5.10. Configure time-zone	62
5.11. Root password	63
5.12. Adding users to the system	64
5.13. Admin password	65

5.14. Boot disk	66
5.15. Installing a bootloader	67
5.16. It's finished!	69
6. <i>DiskDrake</i> : manage your partitions	71
6.1. The interface	71
6.2. In practice: resize an old partition and create a new one	73
6.3. A note about the expert mode: save the partition table	77
II. Technical overview	79
7. Security under <i>GNU/Linux</i>	81
7.1. Preamble	81
7.2. Overview	82
7.3. Physical Security	89
7.4. Local Security	95
7.5. Files and Filesystem Security	98
7.6. Password Security and Encryption	107
7.7. Kernel Security	120
7.8. Network Security	126
7.9. Security Preparation (before you go on-line)	140
7.10. What To Do During and After a Break-in	143
7.11. Security Sources	147
Security-related terms	151
7.12. Frequently Asked Questions	155
7.13. Conclusion	157
8. Firewall and Proxy Server	159
8.1. Introduction	159
8.2. What's the Purpose Behind a Firewall?	159
8.3. How Do Computers Communicate?	162
8.4. The Firewall: What Is It?	169
8.5. How Does the Firewall Work?	172
8.6. Before Implementing a Firewall	173
8.7. Nitty Gritty Firewalling Process	173
8.8. After Implementing a Firewall	175
9. Networking Overview	177

9.1. Copyright	177
9.2. How to use this Chapter.	177
9.3. General Information about Linux Networking. ...	179
9.4. Generic Network Configuration Information.	182
9.5. Ethernet Information	191
9.6. IP Related Information	197
9.7. Using common PC hardware	198
9.8. Other Network Technologies	201
9.9. Cables and Cabling	217
III. System setup and management.....	223
10. <i>MandrakeSecurity</i> Setup and Management	225
10.1. Introduction	225
10.2. Presentation of the Interface	225
10.3. Basic System Configuration	229
10.4. Internet Access	239
10.5. Access Restrictions	265
10.6. Services	291
11. Configuring Masqueraded Clients	301
11.1. Linux Box	302
11.2. Windows 95 or Windows 98 Box	305
11.3. Windows NT or Windows 2000 Box	308
11.4. DOS Box Using NCSA Telnet Package	313
11.5. Windows for Workgroup 3.11	313
11.6. MacOS Box	314
11.7. OS/2 Warp Box	318
IV. Maintenance	319
12. System Monitoring	321
12.1. System Monitoring	321
12.2. Network Traffic	323
12.3. Configuring and Consulting Logs	323
13. System Update	331
13.1. Update Software	331
13.2. Official Mirror List	332
13.3. Packages Selection	333

13.4. Personal Mirror List	334
14. Management Tools	337
14.1. Remote Secure login	337
14.2. Backup And Restore	338
A. Where To Get Documentation	341
A.1. The Documentation Included In Linux-Mandrake ...	341
A.1.1. The Manual	341
A.1.2. info Pages	342
A.1.3. <i>HOWTO</i> s	343
A.1.4. The Directory /usr/share/doc	344
A.2. <i>Internet</i>	344
A.2.1. Websites Devoted to <i>GNU/Linux</i>	344
A.2.2. Mailing Lists	346
A.2.3. Newsgroups	347
A.3. General Guidelines For Solving A Problem Under Linux-Mandrake	348
A.3.1. RTFM	348
A.3.2. Search The <i>Internet</i>	349
A.3.3. Mailing-lists And Newsgroups Archives	349
A.3.4. Questions To Mailing-Lists And Newsgroups ..	350
A.3.5. Contacting The Person In Charge Directly	350
Glossary	351
B. The GNU General Public License	387
B.1. Preamble	387
B.2. Terms and conditions for copying, distribution and modification	388
C. GNU Free Documentation License	397
0. PREAMBLE	397
1. APPLICABILITY AND DEFINITIONS	397
2. VERBATIM COPYING	399
3. COPYING IN QUANTITY	399
4. MODIFICATIONS	400
5. COMBINING DOCUMENTS	403
6. COLLECTIONS OF DOCUMENTS	404

7. AGGREGATION WITH INDEPENDENT WORKS	404
8. TRANSLATION	405
9. TERMINATION	405
10. FUTURE REVISIONS OF THIS LICENSE	405
How to use this License for your documents	406
Index	407

List of Tables

List of Figures

3-1. The dosutils directory	26
3-2. The RawWrite program	27
3-3. An example of using RawWrite	28
3-4. The <i>Windows</i> Device Manager	32
3-5. Directory structure for ISA Bus	33
3-6. Keyboard resources	34
4-1. First example of partition naming under <i>GNU/Linux</i>	42
4-2. Second example of partition naming under <i>GNU/Linux</i>	42
5-1. Choose your preferred language	45
5-2. Accept the license to proceed with the installation	46
5-3. SCSI card installation	47
5-4. Choose your keyboard	48
5-5. Miscellaneous options	50
5-6. Use free space or <i>DiskDrake</i> ?	51
5-7. Do you want to use the existing partitioning scheme?	51
5-8. Do you want to erase <i>Windows</i> ?	52
5-9. Do you want to erase the entire disk?	53
5-10. Choose which partitions to format	54
5-11. <i>DrakX</i> installing all <i>MandrakeSecurity</i> packages	56
5-12. Change default network configuration?	57
5-13. IP address configuration for this interface	57
5-14. Name information and gateway	58
5-15. Proxy servers configuration	59
5-16. Configure an <i>Internet</i> connection	61
5-17. Enter information for this dialup account	62
5-18. Choose the correct timezone	62
5-19. Choose the root password	64
5-20. Adding users to your system	65
5-21. Choose the Administrator (admin) password	66
5-22. Create a boot disk?	67

5-23. Choice of the location of the bootloader	67
5-24. Configure boot entries	68
5-25. Finish the installation	69
6-1. The <i>DiskDrake</i> main window	72
6-2. The /home partition before resizing	73
6-3. Choosing a new size	74
6-4. Defining the new partition	75
6-5. The new partition table	75
6-6. Confirm the writing of partition table	76
8-1. The TCP/IP Protocol Stack	166
8-2. The Packet's Encapsulation Process	167
8-3. The UDP Protocol Header	168
8-4. The IP Protocol Header	168
8-5. The Most Simple LAN Firewall Configuration	170
8-6. A Firewall Between a LAN and a DMZ	170
8-7. A Firewall with Three Network Interfaces	171
8-8. A Complex Firewall Configuration	172
9-1. A dynamic routing example	189
9-2. The NULL-modem cabling	218
9-3. 10base2 Ethernet Cabling	220
10-1. The Login Window to Connect to <i>MandrakeSecurity</i> ...	226
10-2. A Sample <i>MandrakeSecurity</i> Interface Screen	227
10-3. The log out menu entry	228
11-1. Reconfiguring a Network with Draknet	303
11-2. Reconfiguring the Local Network with Draknet	303
11-3. Setting up the Gateway with Draknet	304
11-4. The Network Icon under Windows 95	305
11-5. The Network Configuration Panel under Windows 95	305
11-6. The TCP/IP Configuration Panel under Windows 95	306
11-7. The Gateway Configuration Panel under Windows 95	307
11-12. Accessing the TCP/IP Control Panel	314
11-13. Automatic Configuration of Internet Access for MacOS ..	315
11-14. Manual Internet Acces Configuration for MacOS	317

Preface

1. Note from the Editor

As you may notice while browse from a chapter to another, this book is a composite document from various authors. Even though much care has been taken in insuring the technical and vocabulary consistency, the style of each author is obviously preserved.

Some of the authors write in English even though it is not their native language. Therefore, you may notice strange sentence constructions; do not hesitate to let us know if something is not clear to you.

In the open-source philosophy, contributors are much welcomed! You may provide much help to this documentation project by different means. If you have a lot of time, you can write a whole chapter. If you speak a foreign language, you can help us with the internationalization of this book. If you have ideas on how to improve its contents, let us know - even advice on typos is welcomed!

For any information about the **Linux-Mandrake** documentation project, please contact the documentation administrator (<mailto:documentation@mandrakesoft.com>).

2. Legal Notice

This manual (except the chapters listed below) is protected under **MandrakeSoft** intellectual property rights. Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.1 or any later version published by the Free Software Foundation; with the invariant sections being *About **Linux-Mandrake***, page 15, with the front-cover texts being listed below, and with no back-cover texts. A copy of the license is included in the section *GNU Free Documentation License*, page 397.

Preface

Front-Cover Texts:

MandrakeSoft April 2001

<http://www.mandrakesoft.com/>

Copyright © 1999–2001 by MandrakeSoft S.A. and MandrakeSoft Inc.

Note: The chapters listed in the table below are protected by a different license. Consult the table and links for more details about these licenses.

	Original Copyright	License
“ <i>Networking Overview</i> ”, page 177	Joshua D. {POET} Drake - linuxports (http://www.linuxports.com/)	GNU General Public License GPL (http://www.gnu.org/copyleft/gpl.html)
“ <i>Security under GNU/Linux</i> ”, page 81	Kevin Fenzi and Dave Wreski	Linux Documentation Project LDP (http://linuxdoc.org)

3. Authors and Translators

The following people contributed to the making of the **Linux-Mandrake** manuals:

- Yves Bailly
- Jay Beale
- Camille Bégnis
- Vincent Danen
- Francis Galiègue

- Fabian Mandelbaum
- Roberto Rosselli Del Turco
- Christian Roy
- Authors of whom we reproduced documents in this guide (see list at *Legal Notice*, page 13)

The following people also participated to various degrees: Renaud Chaillat, Damien Krotkine, Philippe Libat, Vincent Saugey.

4. Tools Used in the Making of this Manual

This manual was written in *DocBook*. *Perl* and *GNU Make* were used to manage the set of files involved. The SGML source files were processed by *openjade* and *jadetex* using Norman Walsh's stylesheets. Screenshots were taken using *xwd* and *GIMP* and converted with *convert* (from the *ImageMagick* package). PostScript files were produced with *dvips*. The complete software is available on your **Linux-Mandrake** distribution and all parts of it are free software.

5. About Linux-Mandrake

Linux-Mandrake is a *GNU/Linux* distribution supported by **MandrakeSoft** S.A. **MandrakeSoft** was born in the *Internet* in 1998 with the main goal to provide an easy-to-use and friendly *GNU/Linux* system. The two pillars of **MandrakeSoft** are open-source and collaborative work.

5.1. Contact Mandrake community

Following are various *Internet* links pointing you to various **Linux-Mandrake** related sources. If you wish to know more about the **MandrakeSoft** company, connect to its web site (<http://>

Preface

www.mandrakesoft.com/). There is then the site for the **Linux-Mandrake** distribution (<http://www.linux-mandrake.com/>) and all its derivatives.

First of all **MandrakeSoft** is proud to present its new open help platform. MandrakeExpert (<http://www.mandrakeexpert.com/>) isn't just another web site where people help others with their computer problems in exchange for up-front fees, payable regardless of the quality of the service received. It offers a new experience based on trust and the pleasure of rewarding others for their contributions.

In addition, MandrakeCampus (<http://www.mandrakecampus.com/>) provides the *GNU/Linux* community with open education and training courses on all open software-related technologies and issues; and teachers, tutors, and learners with a place where they can share knowledge.

There is a site for the “mandrakeholic” called Mandrake Forum (<http://www.mandrakeforum.com/>): a primary site for **Linux-Mandrake** related tips, tricks, rumors, pre-announcements, semi-official news, and more. This is also the only interactive web-site hosted by **MandrakeSoft**, so if you have something to tell us, or something you want to share with other users, search no longer: this is a place to do it!

In the philosophy of open-source, **MandrakeSoft** is offering many means of support (<http://www.Linux-Mandrake.com/en/ffreesup.php3>) for the **Linux-Mandrake** distributions. You are invited in particular to participate in the various Mailing lists (<http://www.Linux-Mandrake.com/en/flists.php3>), where the **Linux-Mandrake** community demonstrates its vivacity and keenness.

5.2. Support Mandrake

By popular request, **MandrakeSoft** proposes to its happy customers to make a donation (<http://www.linux-mandrake.com/donations/>) to support the fore-coming developments of the **Linux-**

Mandrake system. Your contribution will help **MandrakeSoft** providing its users an ever better distribution, ever safer, easier, up-to-date, and with more supported languages.

For the many talented of you, your skills will be fully useful for one of the many tasks required in the making of a **Linux-Mandrake** system:

- Packaging: a *GNU/Linux* system is mainly made of programs picked-up on the *Internet*. These programs have to be packaged so that they will hopefully work together.
- Programming: there are many many projects directly supported by **MandrakeSoft**: find the one that most appeals to you, and propose your help to the main developer.
- Internationalization: translation of the web pages, programs and their respective documentation.
- Documentation: last but not least, the book you are currently reading requires a lot of efforts to stick to the rapid evolution of the system.

Consult the contributors page (<http://www.mandrakesoft.com/labs/>) to know more about the way you can contribute to the evolution of **Linux-Mandrake**.

Preface

Chapter 1. Introduction to the User Guide

Welcome, and thank you for using *MandrakeSecurity*!

This book is divided into two parts: a **Installation guide** and a **User Guide**. The **Installation guide** will help you to install your *MandrakeSecurity* firewalling system by describing the preparation, installation and post-installation procedures. The **User Guide** will help you with the daily usage of your *MandrakeSecurity*.

The **User Guide** is itself divided into three parts, enhanced with appendices: Technical overview, System setup and management, Maintenance. Here is a summary of the **User Guide**'s chapters:

It is imperative you read the first chapter, which deals with general security issues. Before even thinking of using your computer, read this chapter! Most sections are related to specific services. Do not hesitate to read them again and again at the time of setting up those services.

The second chapter is a firewall technical introduction. Even though you are not required to be a theoretical firewall expert, we greatly encourage you to carefully read this chapter since it will help you to better tune your firewall and understand what you are really doing.

Next is a chapter about networking, of course. This reference manual deals with general network issues: protocols, physical connections, server and client configurations, etc. Once again, it is not required for you to read those boring lines, but every system administrator should at least know those very basic concepts :-)

We then tackle practical aspects. The first section describes the features of *MandrakeSecurity*'s web interface and walks you through the full configuration procedure of your server. The next one shows you how to configure the various clients of your local network in or-

der for them to connect to the Internet through your newly installed firewall.

The next part presents your system's maintenance tools, always through the web interface.

You will then find useful appendices. The first one is devoted to documentation. Apart from introducing you to the documentation available on your *GNU/Linux* system, it provides helpful links to Internet sites and, more specifically, security-related ones.

Finally, a glossary of technical terms, a copy of the GPL license applicable to the whole *MandrakeSecurity* distribution, and a short index for this book are also included.

I. Installation guide

Chapter 2. Some words before you begin the installation

2.1. Welcome!

The aim of this manual is to help you to install *MandrakeSecurity* on your computer. The setup program used is the graphical setup program: *DrakX*. If, for one reason or another, you cannot or prefer not to use the graphical install, you will be able to use a text version; how to access it is explained at the beginning of the section “*Installation with DrakX*”, page 45.

An entire section is devoted to disks and partitioning. This should help you when you have to partition your hard drive, although *DrakX* is designed to handle this automatically.

And finally we will look at the installation itself.

Chapter 2. Some words before you begin the installation

Chapter 3. Before setup

3.1. Configuring your BIOS

The *BIOS* (*Basic Input/Output System*) is used to boot up a computer. Specifically, it is used to find the device on which the operating system is located and start it up. It is also used for the initial configuration of the hardware.

The appearance of *plug'n'play* and its widespread use means that all modern *BIOS* can initialize these devices, but you still have to ask it to do so. If you have another OS that is initializing these devices instead of the *BIOS* (and want to keep it), this will need to be changed for use under *GNU/Linux*.

Changing your *BIOS* settings is usually performed by holding down the DEL key just after the computer is switched on. Unfortunately, there are many types of *BIOS*, therefore you will have to look for the appropriate option by yourself. The option to look for is often called PNP OS installed (or Plug'n'Play OS installed). Set this option to No and the *BIOS* will then initialize any *plug'n'play* devices. That can help *GNU/Linux* recognize some devices in your machine which it would not otherwise be able to initialize.

If your *BIOS* can boot from the CDROM you can also set your *BIOS* to boot from the CDROM before searching the hard disk. Look for Boot sequence in the *BIOS* features setup.

Note: It's important to note that *MandrakeSecurity* will run even on a 386 computer. Please, bear in mind that all the *plug'n'play* features may not apply to the 386 based machines because there were no 386 motherboards supporting it, however the other 386 devices are very well supported under *MandrakeSecurity*.

3.2. Creating a “bootdisk”

It may be that on an **Intel** architecture your *BIOS* cannot boot from the CDROM. In this case, you must make a “bootdisk” in order to start the installation program.

The boot images are in the `images` directory on the CDROM. For this method of installation, the significant file is named `cdrom.img`.

We use here the image `cdrom.img` when you plan to install the distribution from a CDROM. However, many other images are available to perform installs from:

- `network.img`: to install from a NFS, FTP, HTTP repository. The network configuration of the machine to be installed may be manual or automatic.
- `pcmcia.img`: if the installation media is reached through a PCMCIA card (network, CDROM, ...)

3.2.1. Under Windows

If you have *Windows* installed on your machine or a spare one, you need to use the program called `RawWrite`. This can be found in the `CD dosutils` directory (figure 3-1).

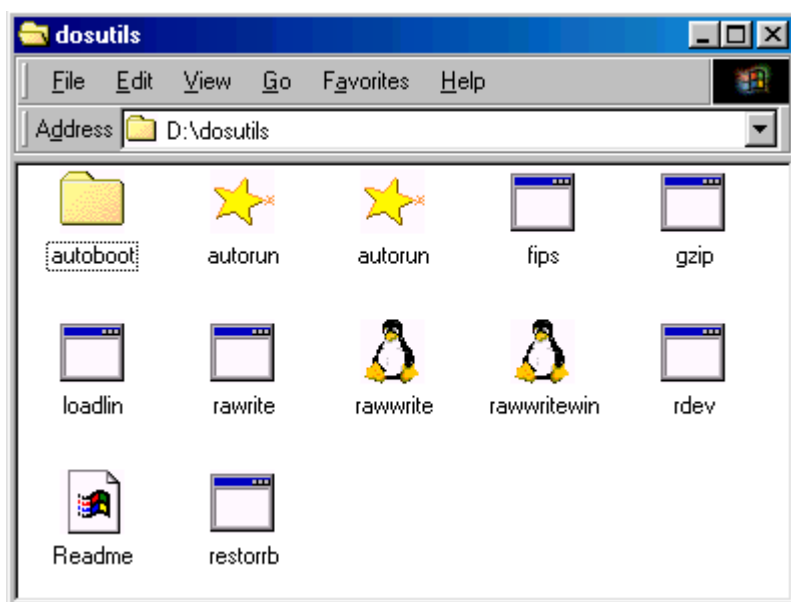


Figure 3-1. The dosutils directory

Note: In this example, the CDROM drive is designated by the letter D:; you will naturally have to choose the letter designating the CDROM drive on your own machine.

You may have noticed that there is a *DOS* version, *rawwrite*, of the same program. It is, in fact, the original version of the program: *RawWrite* is a graphical frontend to it.

Start the program, as shown in figure 3-2.

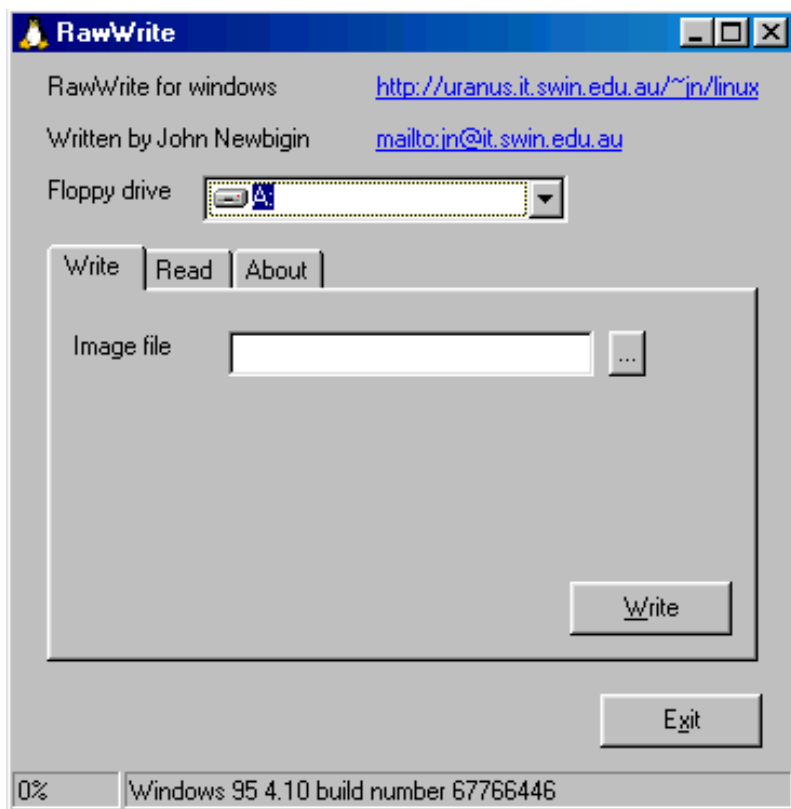


Figure 3-2. The RawWrite program

Select the boot image to copy and the target device (here A: as illustrated in figure 3-3).

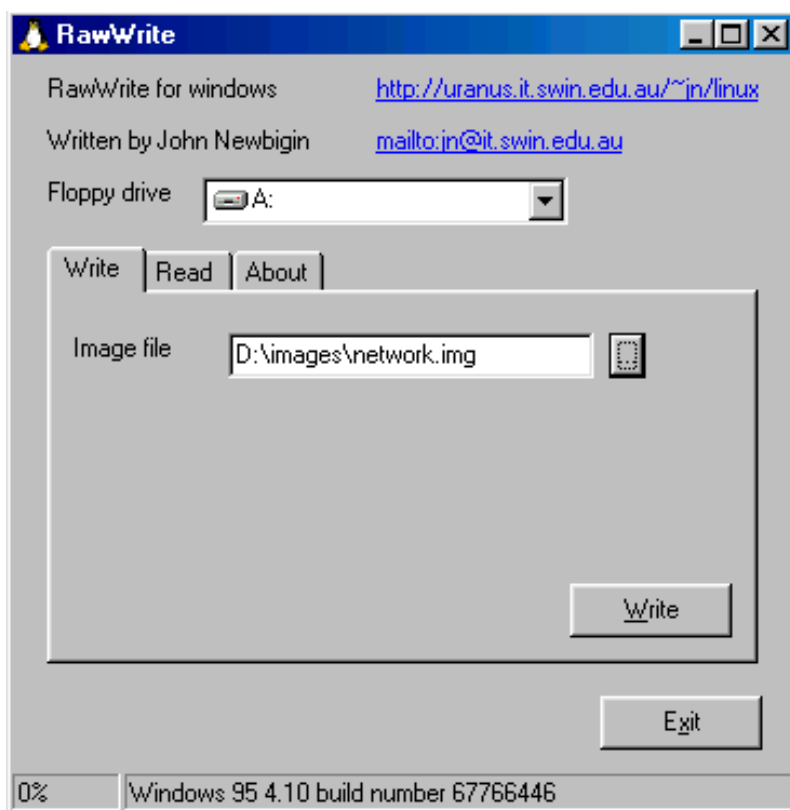


Figure 3-3. An example of using RawWrite

Then, if you haven't already done so, insert an empty disk into your chosen floppy drive and click on Write. When completed, click on Exit, you have a boot disk to install your *MandrakeSecurity* distribution.

3.2.2. Under GNU/Linux

If you already have *GNU/Linux* installed (another version, or on another machine, for example on that of a friend who has lent you his

MandrakeSecurity CD), then carry out the following steps:

1. mount the CDROM. Let us suppose that the mount point is `/mnt/cdrom`;
2. log in as root;
3. insert an empty disk into the drive and type:

```
$ cp /mnt/cdrom/images/cdrom.img /dev/fd0
```

Note: Replace `/dev/fd0` by `/dev/fd1` if you are using the second floppy drive and, of course, the name of the image with the one you want. When completed, your boot disk is ready.

3.3. Supported hardware

MandrakeSecurity can handle a large number of hardware devices, and the list is far too long to be quoted in its entirety here. Nevertheless, some of the steps described in this chapter will help you to find out if your hardware is compatible and configure some of the problematic devices.

Finally, you may consult a more up-to-date list of supported hardware on our web-site (<http://www.linux-mandrake.com/en/fhard.php3>)

Warning

Legal disclaimer: The **Linux-Mandrake** Supported Hardware List contains information about hardware devices that have been tested and/or have been reported to function properly with **Linux-Mandrake**. Due to the wide variety of system configurations, **MandrakeSoft** cannot guarantee that a specific device will work properly on your system.

3.3.1. What's not supported

Some types of hardware cannot presently be handled by *GNU/Linux*, either because the support is still in an experimental stage, because nobody has written a driver for the devices in question, or because it has been decided for valid reasons that they cannot be supported. For example:

- *winmodems*, also called controller-less modems or software modems. Support for these peripherals is currently very sparse. Drivers exist, but are binary only and for a limited range of kernel versions. The difference between a “hardware” modem and a *winmodem* is that a *winmodem* cannot function without a special driver which emulates a large number of a hardware modem's functions. You can communicate with a hardware modem by sending it a series of commands, which cannot be done with a *winmodem* without special drivers (this also explains why *GNU/Linux* does not need drivers for the modems: it only gives access to the serial port, with an external program sending the commands). If your modem is PCI, it is most likely, but not necessarily, a software modem...

If your modem is a PCI modem, as the root user look at the output of `cat /proc/pci`. This will tell you the I/O port and the IRQ of the device. Then use the `setserial` command (for our example, the I/O address is 0xb400 and the IRQ is 10) as follows:

```
setserial /dev/ttyS3 port 0xb400 irq 10 UART 16550A
```

Then see if you can query your modem using *minicom* or *Kppp*. If it doesn't work, you may have a software modem. If it does work, create the file `/etc/rc.d/rc.setserial` and place the appropriate `setserial` command line in it.

A recent project is trying to make software modems work under *GNU/Linux*. If you happen to have this type of hardware in your machine, you may have a look at *linmodems* (<http://linmodems.org/>) and *modems* and *winmodems* (<http://www.o2.net/~gromitkc/winmodem.html>).

- USB devices: support for USB is still limited. Currently the only devices fully supported by **Linux-Mandrake** are keyboards, mice, ZIP drives and printers. For other devices, a *HOWTO* (<http://linuxusbguide.sourceforge.net/USB-guide-1.0.6/book1.html>) is available. You may also consult the web-site *Linux-USB* (<http://www.linux-usb.org/>)

3.3.2. Collecting information on your hardware

GNU/Linux' hardware resources are now much better supported and, apart from the devices mentioned in the previous section, you can expect the rest of your hardware to work correctly.

Some types of devices are still problematic with *GNU/Linux*, especially ISA *plug'n'play* devices: but you can use *Windows* to discover their working configuration. If you intend to install *MandrakeSecurity* while leaving a version of *Windows* on your machine, you can ignore this section at first, and then come back here if you experience problems under *GNU/Linux*.

For this, boot under *Windows*, right-click on the My Computer icon, choose **Properties**, select the tab **Devices Manager**, then select **View devices by connection** (figure 3-4).

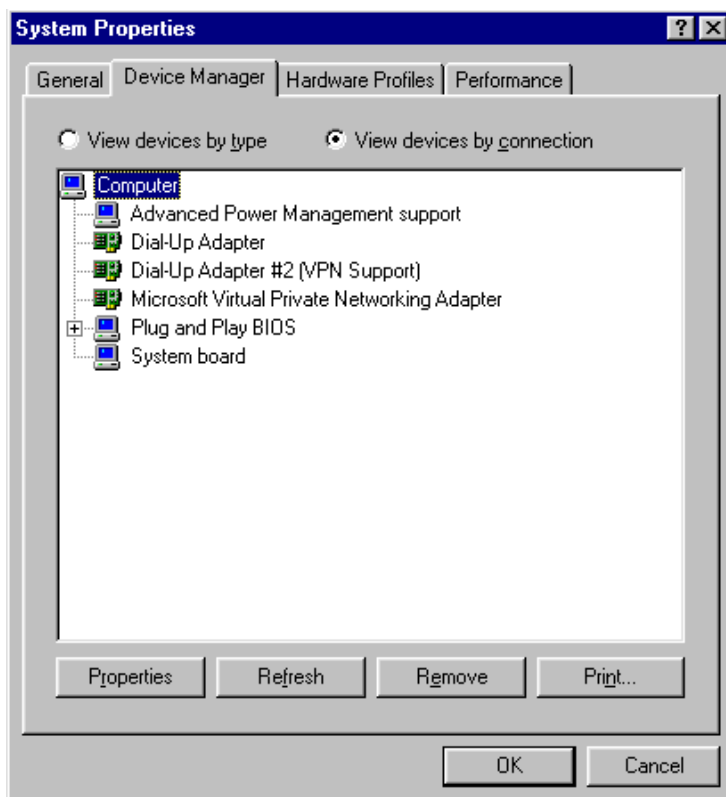


Figure 3-4. The Windows Device Manager

If you have ISA devices, you can view them when you bring up the directory structure (figure 3-5).

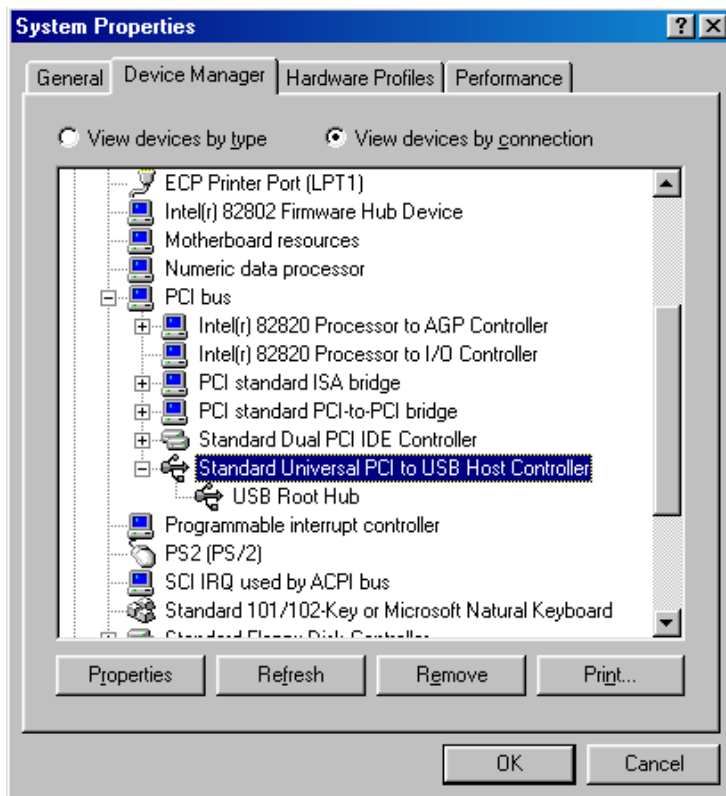


Figure 3-5. Directory structure for ISA Bus

You will be able to find the ISA devices in this part of the directory structure. If you only see one entry for the data port, ignore it. If there are devices present, and if there is no conflict, you can then select and click on the button **Properties** (figure 3-6).




Figure 3-6. Keyboard resources

You will have to write down the base address(es) (Input/output range) used, together with the IRQ(s). Also write down the DMA channel(s) used for the sound cards.

If your ISA card is *plug'n'play*, you will have to configure your *BIOS* properly, as instructed in the preceding section. But even if you do so, *GNU/Linux* may not find it. However, you can disable *plug'n'play* for the particular device. If the manufacturer has provided one, you should have a disk containing a program enabling you to re-set the card to non-*plug'n'play*. The manufacturer provides a setup program which does this. If you have this sort of program

(or can get it from the manufacturer's web-site), start it up, set the device's configuration with the parameters used by *Windows* and disable the *plug'n'play*. After that, *GNU/Linux* can then see it during the installation process.

Chapter 4. Disks and partitions

This section provides a complete description of the *PC* partitioning scheme. It is only useful to you if you intend to manually set the partitions of your hard drive using *DiskDrake* or if you're just plain curious. Anyway, you may safely ignore this section; *DrakX* will do everything automagically for you.

4.1. Structure of a hard disk

Basically, a disk is physically divided into little sectors. A sequence of sectors can form a partition. Roughly speaking, you can create as many partitions as you wish; each of them is regarded as a single hard drive.

4.1.1. Sectors

To simplify, a hard disk is merely a sequence of **sectors**. A sector is the smallest data unit on a hard disk, and its size is typically 512 bytes. The sectors on a hard disk of (*n*) sectors are numbered from (0) to (n-1).

4.1.2. Partitions

The use of multiple partitions enables you to create many virtual hard drives inside your real physical drive. This has many advantages:

- Different operating systems use different disk structures (called file system); this is the case for *Windows* and *GNU/Linux*. Having

multiple partitions on a hard drives allows you to install various operating systems on the same physical drive.

- For performance reasons, a single operating system may prefer different drives with different filesystems on them because they are used for completely different things. It is the case for *GNU/Linux* which requires a second partition called “swap” and used for virtual memory.
- Finally, it may prove very useful to separate the different parts of your OS into different partitions, even if they use the same filesystem. In the most simple configuration, you can split your files into two partitions, one for your personal data, one another for programs. That allows you to update your OS, completely erasing the programs partition while keeping the data partition safe.
- Physical errors on a hard disk are generally located at adjacent sectors and not scattered among the disk. Distributing your files into different partitions will limit data losses in case of hard disk physical damages.

Normally the partition type specifies the filesystem which the partition is supposed to contain. Each operating system recognizes some of the types, but not others.

4.1.3. Define the structure of your disk

4.1.3.1. The most simple

Is where you have just two partitions: one for the swap space, the other for the files¹.

1. the filesystem used currently for *GNU/Linux* files is called `ext2`

Tip: The rule of thumb for the swap partition size is to choose the same size as your RAM memory. However, for large memory configurations (more than 128 MB) this rule is not valid, and sizes smaller than your system RAM are preferred.

4.1.3.2. Another common scheme

Is, as we previously discussed, when you choose to separate data from programs. To be even more efficient, one usually defines a third partition called the “root” and labelled as /. It will handle the programs necessary to startup your system and the basic maintenance programs.

So we could define four partitions:

Swap

A partition of type swap, which size is roughly equivalent to the memory size.

Root: /

It is the most important partition. It not only contains the most important data and programs for the system, but will also act as a mount point for other partitions.

The needs for the root partition in terms of size are very limited, 300MB is enough.

Static data: /usr

Most packages install most of their executables and data files under /usr. The advantage of having it on a separate partition is that you can share it easily with other machines over a network.

The size depends on the packages you wish to install. It varies from 100MB for a lightweight installation to various GB for a full install. A compromise of one or two GB (depending on your disk size) generally suffices.

Home directories: `/home`

Here are kept the personal directories for all the users hosted on that machine. It also generally hosts the directories served by HTTP or FTP (respectively for web browsing and file transfers).

Here the partition size depends on the number of users (or services) hosted and their needs.

A variant to that solution is to not use a separate partition for the `/usr` files.

4.1.3.3. Exotic configurations

When setting-up your machine for specific uses such as web server or firewall, the needs are radically different than for a standard desktop machine. For example, a FTP server will probably need a big separate partition for `/home/f.t.p`, while the `/usr` will be relatively small. For such situations, you are encouraged to carefully think about your needs before even beginning the install.

Tip: If, after a period of time using your system, you notice that you should have chosen different sizes and partitions, it is possible to resize most partitions without the need to reinstall your system, it is even generally data-safe.

With a little of practice, you will even be able to move a crowded partition to another brand new hard drive. But that's another story...

4.2. Conventions for naming the disks and partitions

GNU/Linux uses a logical method for naming partitions. First, when numbering the partitions, it ignores the filesystem types of each partition that you may have. Second, it names the partitions according to the disk on which they are located. This is how the disks are named:

- the primary master and primary slave IDE devices (whether they be hard disks, CDROM drives or anything else) are called `/dev/hda` and `/dev/hdb` respectively;
- on the secondary interface, they are called `/dev/hdc` and `/dev/hdd` for the master and slave respectively;
- if your computer contains other IDE interfaces (for example, the IDE interface present in some SoundBlaster cards), the disks will then be called `/dev/hde`, `/dev/hdf`, etc.
- SCSI disks are called `/dev/sda`, `/dev/sdb`, etc., in the order of their appearance on the SCSI chain (depending on the increasing *IDs*). The SCSI CDROM drives are called `/dev/scd0`, `/dev/scd1`, always in the order of their appearance on the SCSI chain.

The partitions are named after the disk on which they are found, in the following way (in the example, we have used the case of partitions on a primary master IDE disk):

- the primary (or extended) partitions are named `/dev/hda1` through `/dev/hda4` when present;

- logical partitions, if any, are named `/dev/hda5`, `/dev/hda6`, etc. in their order of appearance in the table of logical partitions.

So *GNU/Linux* will name the partitions as follows:

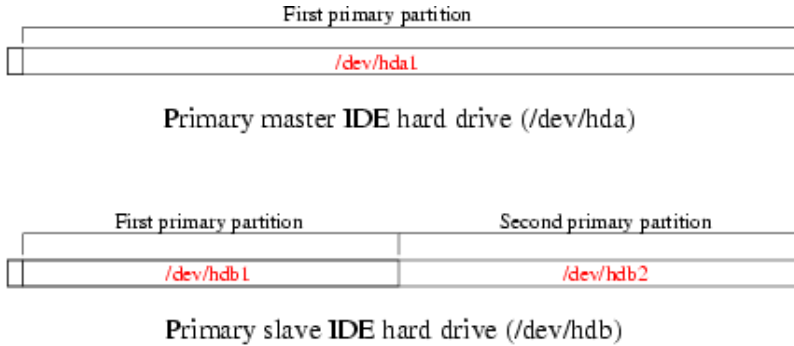


Figure 4-1. First example of partition naming under GNU/Linux

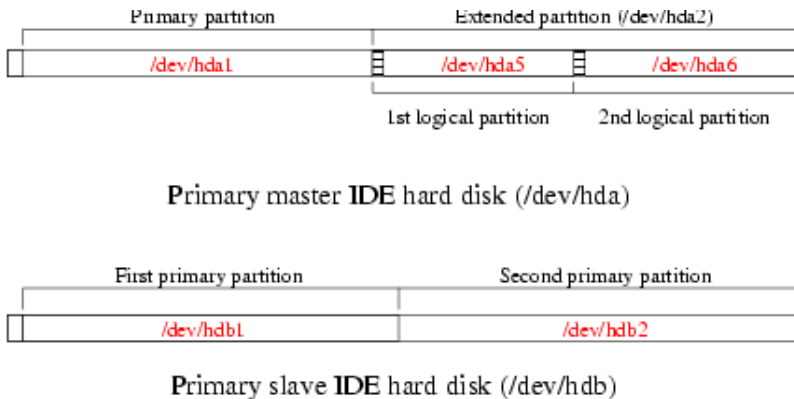


Figure 4-2. Second example of partition naming under GNU/Linux

So now you can cite the name of the various partitions and hard disks when you need to manipulate them. You will also see that *GNU/Linux* names the partitions even if it does not know how to manage them

initially (it ignores the fact that they are not native *GNU/Linux* partitions).

Chapter 5. Installation with DrakX

5.1. Introduction to the MandrakeSecurity installer

DrakX is the installation program for *MandrakeSecurity*. The text installation is not documented in this manual, but it is very similar in practice to *DrakX*.

If you want to (or have to) use this text mode install, just press F1 when the CDROM boots and type text at the prompt.

When you start the installation, the first thing to do will be to choose the language for installation and system usage, as shown in figure 5-1.

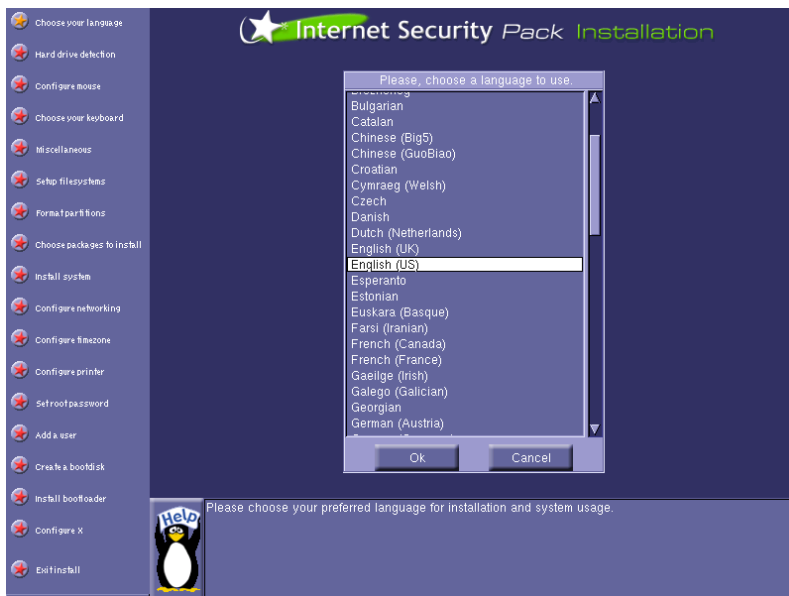


Figure 5-1. Choose your preferred language

On the left, you can see the various installation steps. Depending on the level of progress of the installation, some stages may or may not be available to you. If they are available, they will be highlighted when you move the mouse cursor over them.

The buttons representing the various stages can also have different colors:

1. red: this installation phase has not yet been carried out;
2. orange: the installation stage which is currently processing;
3. green: this installation stage has already been configured. There is nothing preventing you from going back to it if you need to.

This guide assumes that you are performing a standard, step-by-step installation, as shown in the following screenshots.

As soon as you have selected the language and confirmed by clicking the **OK** button, you will automatically go on to the next page.

Before going further, you should carefully read the terms of the license as shown in figure 5-2. It covers the whole *MandrakeSecurity* distribution, and if you do not agree with all the terms in it, click on the **Refuse** button. That'll immediately terminate the installation. To proceed with the installation, click the **Accept** button.



Figure 5-2. Accept the license to proceed with the installation

5.2. Disk detection and configuration

Then *DrakX* will go on to detect all available hard disks on your computer. It will also scan for one or more PCI SCSI card(s) on your system, if you have any. If such a device is found, *DrakX* will automatically install the right driver.

Should it fail, you are anyway asked whether you have a SCSI card or not as shown in figure 5-3. Answer **Yes** to choose your card from a list or **No** if you have no SCSI hardware.



Figure 5-3. SCSI card installation

5.3. Configuring the Keyboard

When you get to this stage, *DrakX* will have selected a keyboard matching the language which you have chosen at the beginning of the installation. So, there's normally nothing to do but clicking on OK at this point. But this may not be what you want: in this case, just select the right keyboard from the list as shown in figure 5-4.

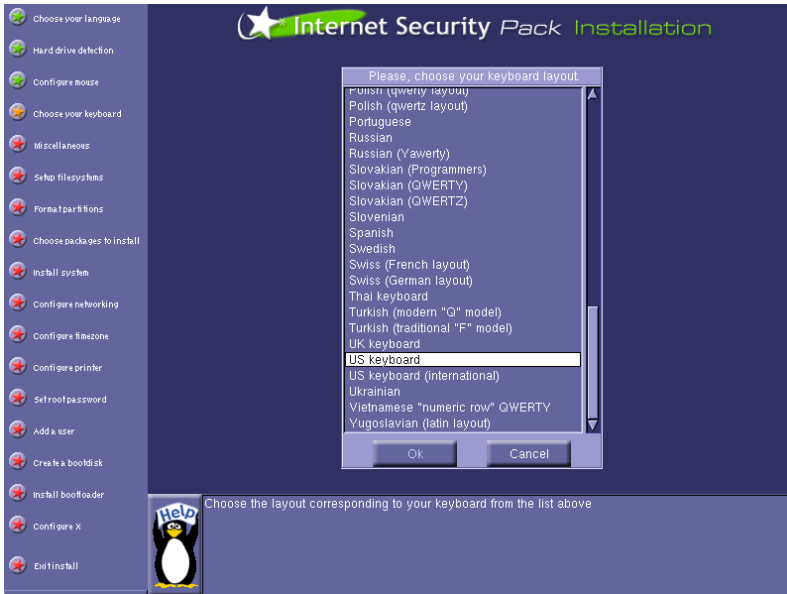


Figure 5-4. Choose your keyboard

5.4. Miscellaneous options

At this point, you may set a number of miscellaneous options for your *MandrakeSecurity* system (figure 5-5). Refer to the help available at the bottom of the screen for assistance in understanding these options.

Note: Hard disk optimizations are only relevant for IDE disks.

Warning

Since this machine is going to be your network's firewall (that is, the machine that protects your network from the outside) it is **strongly** advised that you set the security level option to High. Other levels may be suitable for an internal gateway machine inside an already secured network. Lower levels shouldn't be used at all.



Figure 5-5. Miscellaneous options

5.5. Partitioning your hard disk

It is now time to tell *DrakX* how do you want to partition your disk in order to install *MandrakeSecurity*. *DrakX* will ask you whether you want to use the *DiskDrake* graphic partitioning tool or the disk's free space to install *MandrakeSecurity* (figure 5-6).



Figure 5-6. Use free space or DiskDrake?

If you already have partitioned your disk with *GNU/Linux* type partitions *DrakX* will ask you if you want to use the existing partitions, the entire disk (thus wiping out the existing partitions and their data) or use *DiskDrake* to do the partitioning yourself (figure 5-7).



Figure 5-7. Do you want to use the existing partitioning scheme?

If you happen to have *Windows* already installed on your hard disk, *DrakX* will ask you if you want to erase it (thus using the entire disk for *MandrakeSecurity*) or if you want to use *DiskDrake* to do the partitioning yourself (figure 5-8).



Figure 5-8. Do you want to erase Windows?

Suppose that you want to use the entire disk to host *MandrakeSecurity* and don't care about your existing data, so you choose **Erase entire disk** or **Remove Windows(TM)**. In that case *DrakX* will warn you about the risk of losing all your data, if you are sure about it then answer **OK** to the question shown in figure 5-9.



Figure 5-9. Do you want to erase the entire disk?

Note: If you prefer to determine the partitioning scheme on the machine yourself, click on the use diskdrake button when asked. Then refer to the section “*DiskDrake: manage your partitions*”, page 71 to learn how to use the interface.

5.6. Formatting partitions

When you’re done setting up the partitions, *DrakX* will ask you which ones you want to format (figure 5-10). Do not select partitions containing data which you want to keep! Of course, if you selected the automatic partitioning by default, you need to reformat all partitions.



Figure 5-10. Choose which partitions to format

5.7. Packages installation



Figure 5-11. DrakX installing all MandrakeSecurity packages

DrakX will install all the packages needed by *MandrakeSecurity*. All needed packages represent a total of about 450 MB installed on your hard drive. Advanced users may choose to uninstall later unused packages.

5.8. Internal Network configuration

In this section we refer to the “internal” network part of your firewall machine. This is the network card that connects your machine to the other machines on your local network.

At this point, *DrakX* will configure your network interface(s) with default values. If you want to change these defaults, answer **Yes** to the question shown in figure 5-12



Figure 5-12. Change default network configuration?

5.8.1. Changing the default configuration

As for SCSI cards, *DrakX* will first scan the PCI bus in order to find any network card it knows about. If it finds one, it will install the driver for it, otherwise, it will ask you for the driver to try. And as for SCSI cards, do not forget the parameters you got from *Windows* if relevant.

When done configuring your card, you will first have to enter its IP address along with the netmask associated with this interface as shown in figure 5-13.

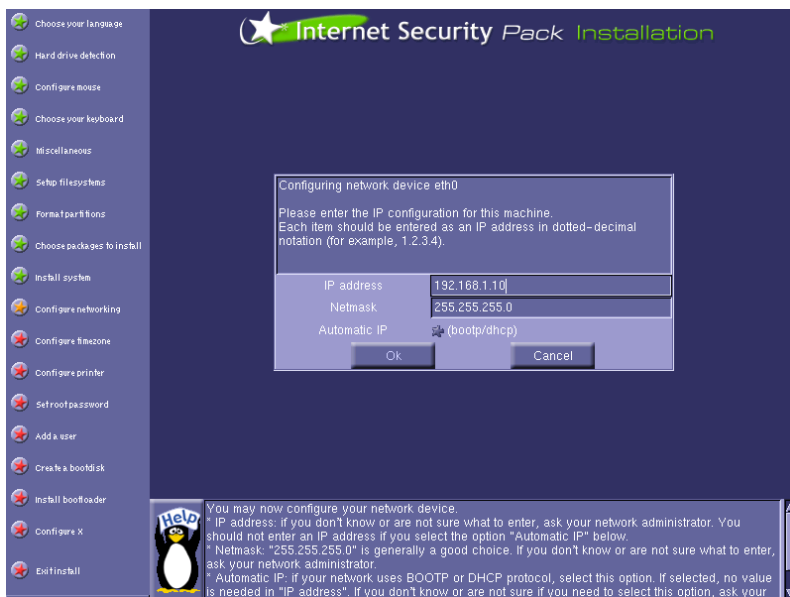


Figure 5-13. IP address configuration for this interface

You will then have to enter the full name for this machine, along with the IP address (not the name!) of the primary nameserver for this domain, and if relevant, the IP address of the gateway. Leave this field blank if there isn't any (figure 5-14), or you will have network problems later.

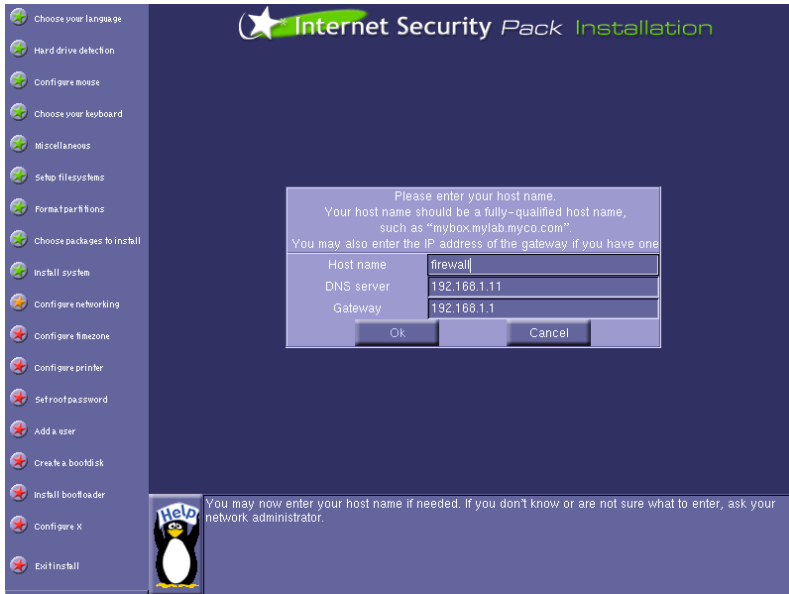


Figure 5-14. Name information and gateway

Finally, you will be able to enter the names (or IP addresses, as you see fit) for your HTTP and FTP proxy servers if necessary. Just leave the irrelevant fields blank (figure 5-15).

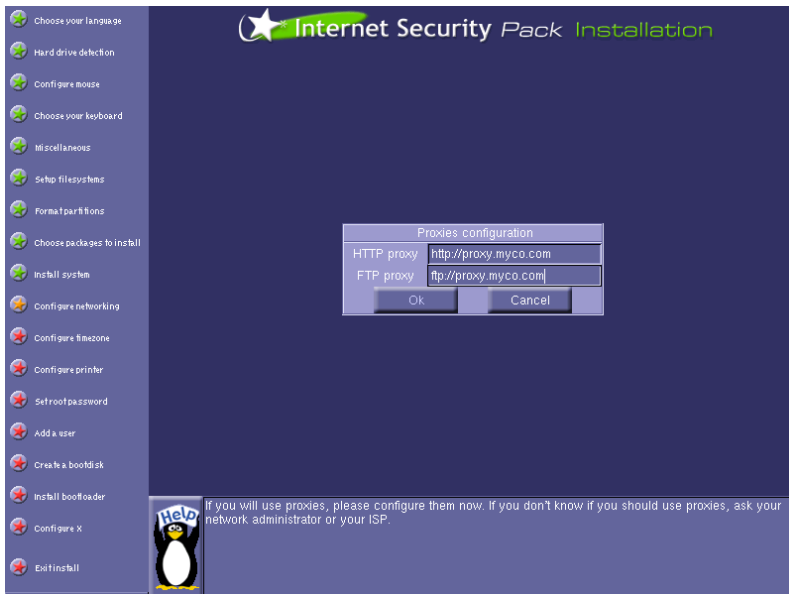


Figure 5-15. Proxy servers configuration

5.9. External Network configuration

In this section we refer to the “external” network part of your fire-wall machine. This is the part that connects your machine to other networks, most likely the *Internet*. We only talk about PPP dialup configuration but, of course, you have other options like ISDN, ADSL and cable-modem connections. Anyway, all those accesses are configurable later via the web interface.

Warning

It is strongly advised to configure the Internet access to your firewall **after** everything else is setup and secured on it.

5.9.1. Configuring a dialup connection via modem

DrakX will then present a menu with choices for your “external” (*Internet* mostly) network connection as shown in figure 5-16.

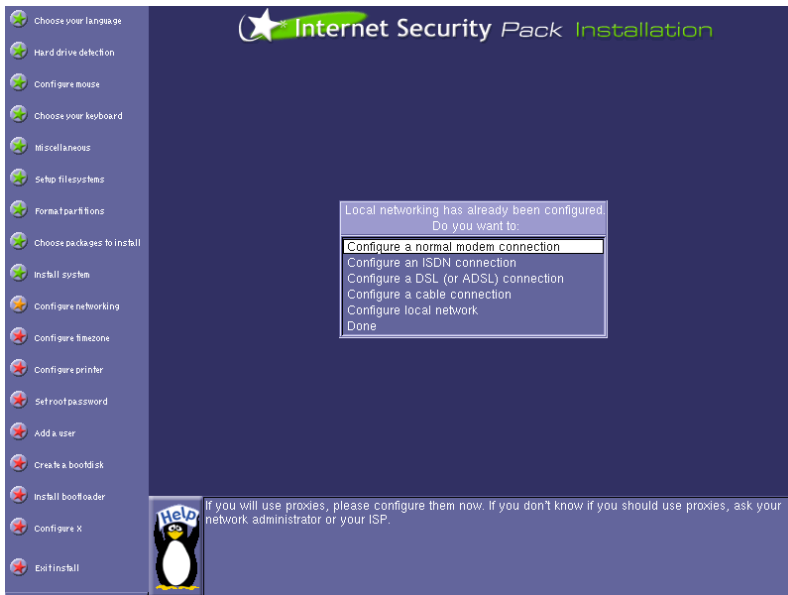


Figure 5-16. Configure an Internet connection

If you choose a dialup connection, you’ll be asked for relevant information concerning this dialup connection (figure 5-17), which you

should have obtained from your Internet Service Provider (ISP) unless you already have it.

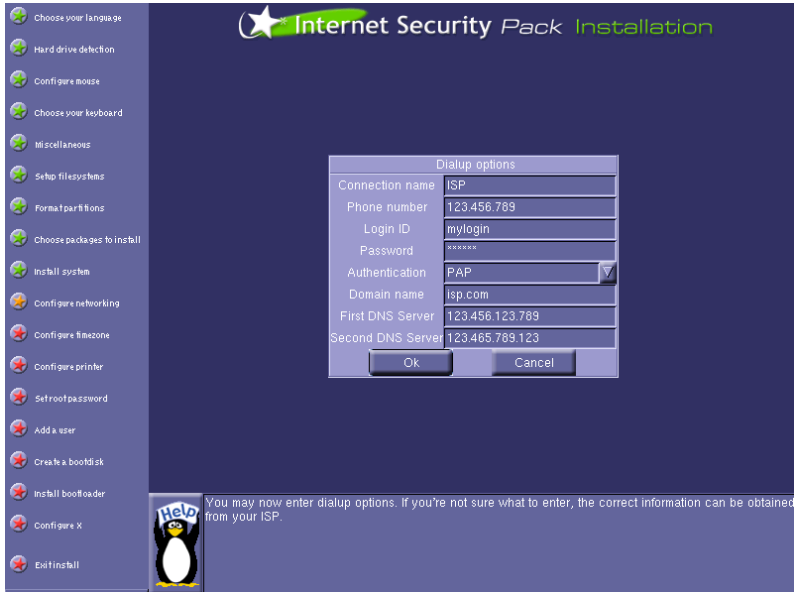


Figure 5-17. Enter information for this dialup account

5.10. Configure time-zone

DrakX, by default, guesses your timezone from the language you have chosen. But here again, as for the keyboard choice, you may not be in the country which the chosen language suggests, so you will choose the appropriate time zone in the list which appears as shown figure 5-18.

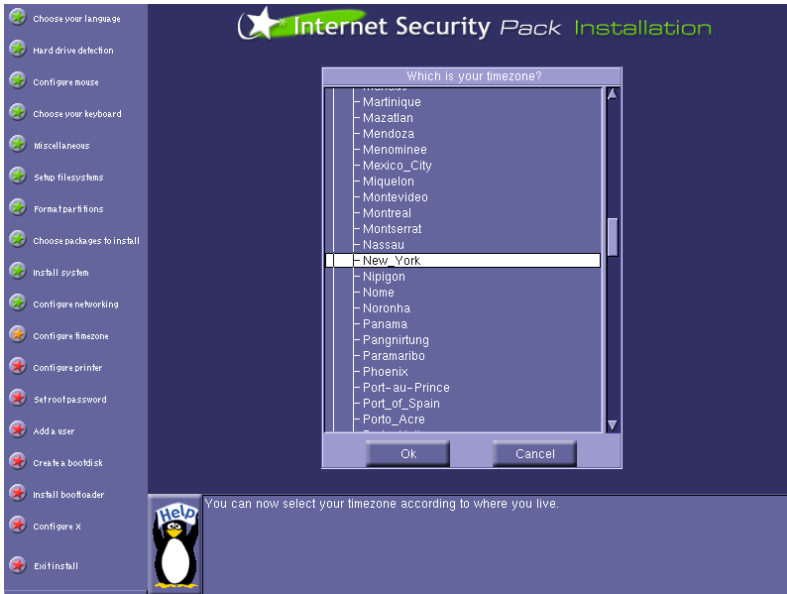


Figure 5-18. Choose the correct timezone

5.11. Root password

No need to say that you need to choose this password with care. With root being the system administrator, he can do anything he wants, including making the system unusable if the root account is used with malicious intent by another. Also bear in mind that, since this computer will be your protection from the outside, you have to be **extremely** cautious when choosing this password.

Basic recommendations for choosing a password for root are: mixed case, possibly one or more digits, ideally a few non alphabetic characters (an underscore, a space, a star – whatever) and more than 6 letters¹. Note that as *MandrakeSecurity* uses MD5 passwords, you can have a password which is more than 8 characters long (figure 5-

19). You will have to type the password twice, only to check that you didn't make a mistake that you would not reproduce when trying to connect to the system. Also, *DrakX* will warn you if your password is too weak. In fact, it will even refuse the password and ask for another one.



Figure 5-19. Choose the root password

Depending on your local network configuration, you may or may not use NIS. If you don't know, ask your system administrator. If you use NIS, check the option Use NIS.

1. Actually, the required length of the password depends on the security level chosen

5.12. Adding users to the system

At this point, you may want to add users to the system. Of course, if all user accounts are shared via a NIS server, the users you will add here will only be local to the machine you are installing on.

For each user you want to add to the system, you will have to enter their real name, their login name (the field User name) and their password. When you have filled the information for a user, you can either select **Accept user** in order to create another user, or **Done** when you're finished. A sample is shown figure 5-20.



Figure 5-20. Adding users to your system

5.13. Admin password

This user is used to connect to the firewall via the web interface. It has not so many privileges as root, but it is still a critical account. Its password must be chosen with much as care as for the root account. Carefully remember this password, or you will be unable to perform further configuration or maintenance tasks.



Figure 5-21. Choose the Administrator (admin) password

5.14. Boot disk

A boot disk can prove handy in case of a problem which prevents you from booting the machine from the disk. You may or may not want to create one (although you should), therefore your mileage may vary depending upon the answer to the question asked figure 5-22.



Figure 5-22. Create a boot disk?

5.15. Installing a bootloader

GRUB is a boot loader for *GNU/Linux*. This stage is normally totally automated. In fact, *DrakX* will analyze the disk boot sector and will act accordingly depending on what it finds here:

- if it finds a *Windows* boot sector, it will replace it with a *GRUB* boot sector so that you can start *GNU/Linux* or *Windows*;
- if it finds a *GRUB* boot sector, it will replace it with a new one;

If in doubt, *DrakX* will ask you where you want to install *GRUB* (figure 5-23).



Figure 5-23. Choice of the location of the bootloader

In most cases, you will not change the default (`/dev/hda`), but if you prefer, the bootloader can be installed on the second hard drive (`/dev/hdb`), or even on a floppy disk (`/dev/fd0`).

Beware that if you choose not to install a bootloader (by selecting **Cancel** here), you must ensure that you have a way to boot your *Man-drakeSecurity* system! Also be sure about what you are doing if you change any of the options here.

If there is another operating system installed on your machine, it'll be automatically added to the boot menu. Here you can choose to fine-tune the existing options (figure 5-24).



Figure 5-24. Configure boot entries

5.16. It's finished!

There you are. Installation is now complete and your *MandrakeSecurity* system is ready to use. Please read the notice that is shown in figure 5-25, and write down the information that is given to you since you'll need it to further configure your system.

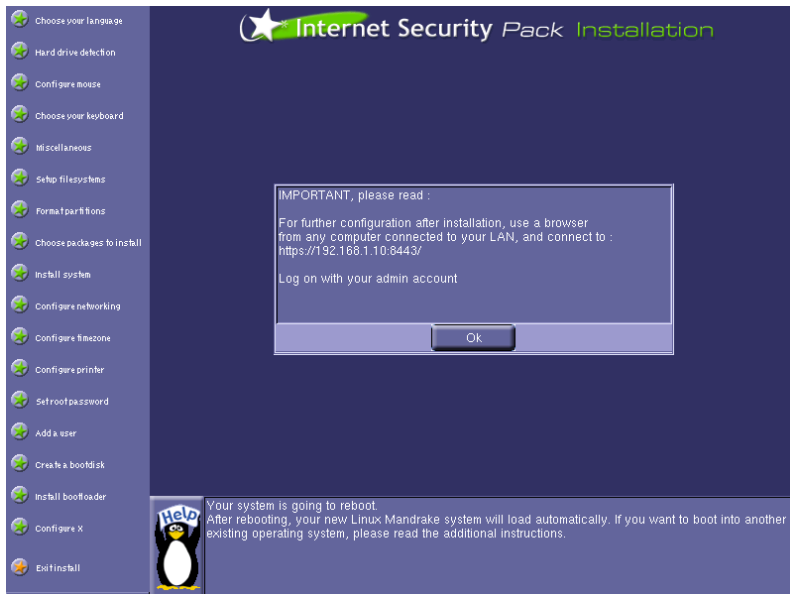


Figure 5-25. Finish the installation

Goodbye, and Thanks for using *MandrakeSecurity* hope to see you again soon :-)

Chapter 6. DiskDrake: manage your partitions

We already learned from “*Structure of a hard disk*, page 37” what partitions are used for; This section will learn you how to use the manual partitionning tool *DiskDrake*, so that you can finely tune your partitions.

Warning

DiskDrake is a very powerful, and therefore dangerous tool. Misuse of it can very easily lead to loss of data in your hard drive. Consequently, you are advised to take some measures before using it:

1. Backup your data: transferring them on another computer, ZIP disks, etc.
2. Save your current partition table (the table describing the partitions hold on your hard drive(s)) on a floppy disk (see *A note about the expert mode: save the partition table*, page 77).

6.1. The interface

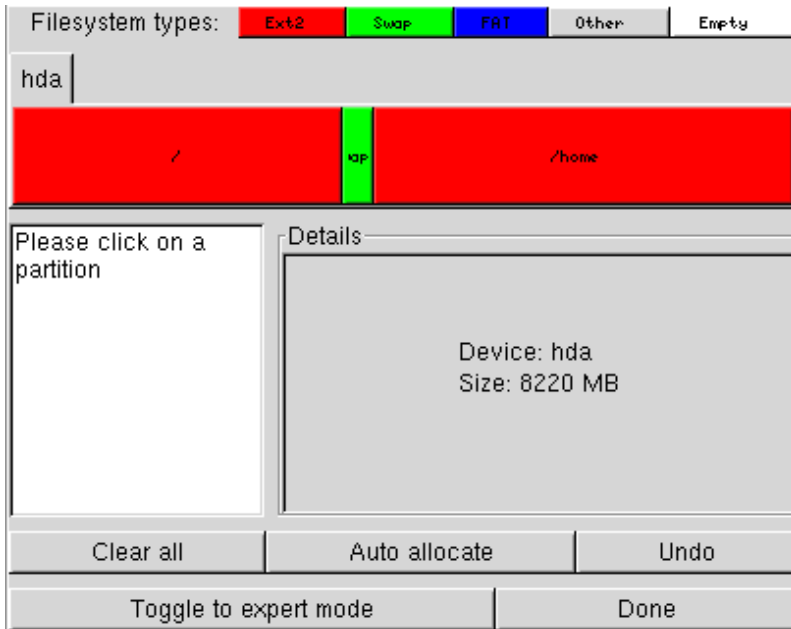


Figure 6-1. The DiskDrake main window

The main *DiskDrake* window (figure 6-1) is divided in four zones:

- On the top: The structure of your hard drive(s). When you launch *DiskDrake* it shows the current structure of the drive, and is modified in real time when you modify your partitions. Note however that changes are not effective on the drive until you press the **Done** button.
- On the left: a menu applying to the partition currently selected in the above diagram.
- On the right: a description of the selected partition.

- At the bottom: buttons for taking general actions. Note that the **Toggle to expert mode** button allows you to access expert (dangerous) functions.

Tip: It is recommended that you first erase the current partitions - if any - by clicking on them and then the Delete button that appears on the left. Then click on the Auto Allocate button at the bottom. That will pre-partition your drive so that you can begin on a good basis.

6.2. In practice: resize an old partition and create a new one

In this section, we are going to make a little exercise that will use more useful features of the tool. Let's imagine that we suddenly decide to use our machine as an FTP server. We then choose to create a separate `/home/ftp` partition in order to host the FTP files.

So this is what look like the current `/home` partition (figure 6-2), before any modification. We choose to shrink this partition in order to create the new one on the freed space.

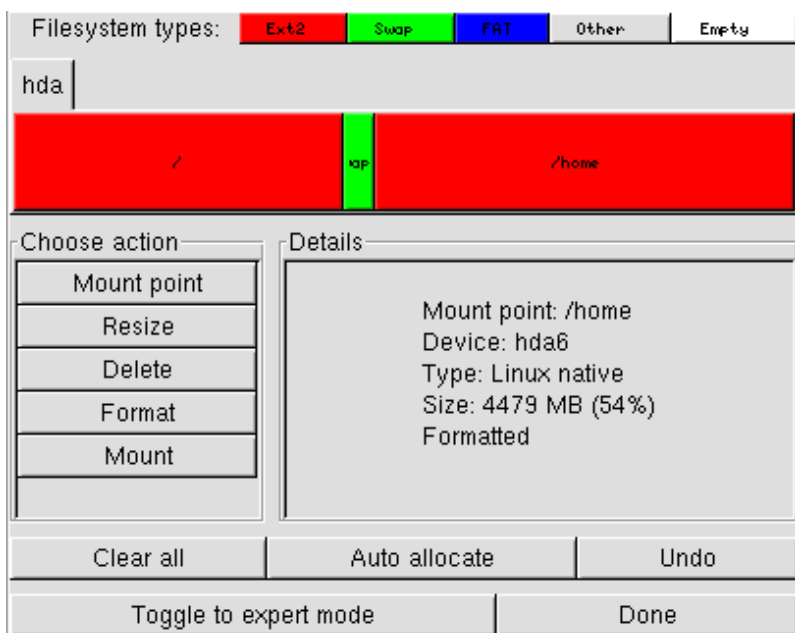


Figure 6-2. The /home partition before resizing

As you may have guessed, just click on the **Resize** button. A dialog will appear (figure 6-3), in which you choose a new size for that /home partition.




Figure 6-3. Choosing a new size

When this is done, you notice the graphic representation of your hard-

drive has changed, the `/home` partition became smaller, and an empty space appeared on the right. Click on that empty space and then on the button **Create** that just appeared. A dialog (figure 6-4) where you can choose the parameters for the new partition pops up. Change the start sector if you want to leave a new free space between the two `/home` and `/home/ftp` partitions. Define the needed size, choose the filesystem you want (generally `Linux native`) and then enter the mount point of that partition, in our case `/home/ftp`.

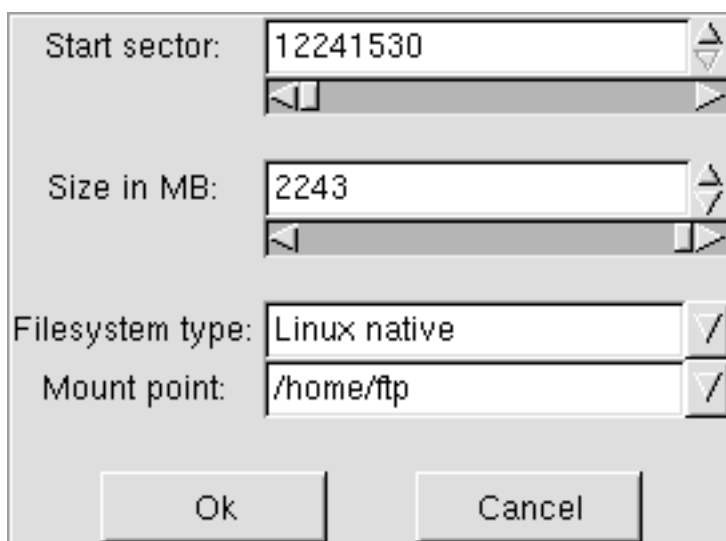


Figure 6-4. Defining the new partition

This is what our projected partition table looks like now (figure 6-5).

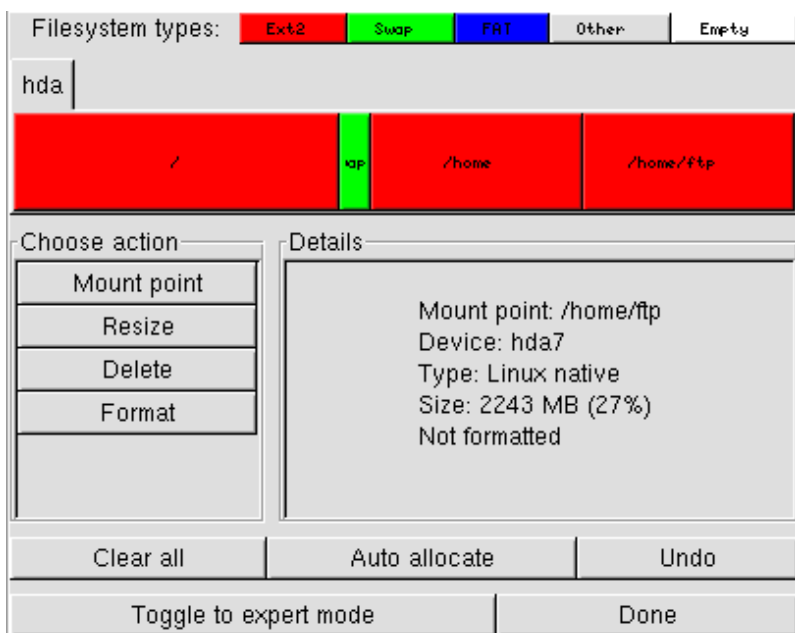


Figure 6-5. The new partition table

Warning

Up to now we did not really modify the partition table, we just redesigned and rejected it. Further steps will effectively make our changes active, so if you do not intend to modify your system, click on the Undo button until you come back to the beginning.

You finally need to format (prepare it to host files) the newly created partition: click on it, then on the **Format** button. Confirm the writing of the partition table, and then the formatting of the partition. You may be asked to reboot the computer in the meanwhile to take changes into account.




Figure 6-6. Confirm the writing of partition table

6.3. A note about the expert mode: save the partition table

Among many available features, the save and restore from file is one of the more interesting. It allows you to save the current partition table to a file on a disk (floppy for example) and then restore it in case that you totally messed up your partition table. It can prove useful as long as you did not reformat partitions, otherwise data will be lost.

II. Technical overview

Chapter 7. Security under GNU/Linux

This document is a general overview of security issues that face the administrator of *GNU/Linux* systems. It covers general security philosophy and a number of specific examples of how to better secure your *GNU/Linux* system from intruders. Also included are pointers to security-related material and programs.

Note: The original document (see below) has been adapted to **Linux-Mandrake** distribution, removing parts, changing others.

7.1. Preamble

This chapter is based on a *HOWTO* by Kevin Fenzi and Dave Wreski which original is hosted by the Linux Documentation Project (<http://linuxdoc.org>)

7.1.1. Copyright Information

This document is copyrighted (c)1998,1999,2000 Kevin Fenzi and Dave Wreski

Modifications from v1.1.1, 17 March 2000, (C)opyright 2000 MandrakeSoft

7.1.2. Introduction

This chapter covers some of the main issues that affect *GNU/Linux* security. General philosophy and net-born resources are discussed.

A number of other *HOWTO* documents overlap with security issues, and those documents have been pointed to wherever appropriate.

This chapter is **not** meant to be a up-to-date exploits document. Large numbers of new exploits happen all the time. This chapter will tell you where to look for such up-to-date information, and will give some general methods to prevent such exploits from taking place.

7.2. Overview

This chapter will attempt to explain some procedures and commonly-used software to help your *GNU/Linux* system be more secure. It is important to discuss some of the basic concepts first, and create a security foundation, before we get started.

7.2.1. Why Do We Need Security?

In the ever-changing world of global data communications, inexpensive *Internet* connections, and fast-paced software development, security is becoming more and more of an issue. Security is now a basic requirement because global computing is inherently insecure. As your data goes from point A to point B on the *Internet*, for example, it may pass through several other points along the way, giving other users the opportunity to intercept, and even alter, it. Even other users on your system may maliciously transform your data into something you did not intend. Unauthorized access to your system may be obtained by intruders, also known as “crackers”, who then use advanced knowledge to impersonate you, steal information from you, or even deny you access to your own resources. If you’re wondering what the difference is between a “Hacker” and a “Cracker”, see Eric Raymond’s document, “How to Become A Hacker”, available at <http://www.netaxs.com/~esr/faqs/hacker-howto.html> (<http://www.netaxs.com/~esr/faqs/hacker-howto.html>).

7.2.2. How Secure Is Secure?

First, keep in mind that no computer system can ever be completely secure. All you can do is make it increasingly difficult for someone to compromise your system. For the average home *GNU/Linux* user, not much is required to keep the casual cracker at bay. However, for high profile *GNU/Linux* users (banks, telecommunications companies, etc), much more work is required.

Another factor to take into account is that the more secure your system is, the more intrusive your security becomes. You need to decide where in this balancing act your system will still be usable, and yet secure for your purposes. For instance, you could require everyone dialing into your system to use a call-back modem to call them back at their home number. This is more secure, but if someone is not at home, it makes it difficult for them to login. You could also setup your *GNU/Linux* system with no network or connection to the Internet, but this limits its usefulness.

If you are a medium to large-sized site, you should establish a security policy stating how much security is required by your site and what auditing is in place to check it. You can find a well-known security policy example at <http://www.faqs.org/rfcs/rfc2196.txt> (<http://www.faqs.org/rfcs/rfc2196.txt>). It has been recently updated, and contains a great framework for establishing a security policy for your company.

7.2.3. What Are You Trying to Protect?

Before you attempt to secure your system, you should determine what level of threat you have to protect against, what risks you should or should not take, and how vulnerable your system is as a result. You should analyze your system to know what you're protecting, why you're protecting it, what value it has, and who has responsibility for your data and other assets.

- **Risk** is the possibility that an intruder may be successful in attempting to access your computer. Can an intruder read or write files, or execute programs that could cause damage? Can they delete critical data? Can they prevent you or your company from getting important work done? Don't forget: someone gaining access to your account, or your system, can also impersonate you.

Additionally, having one insecure account on your system can result in your entire network being compromised. If you allow a single user to login using a `.rhosts` file, or to use an insecure service, such as `tftp`, you risk an intruder getting "his foot in the door". Once the intruder has a user account on your system, or someone else's system, it can be used to gain access to another system, or another account.

- **Threat** is typically from someone with motivation to gain unauthorized access to your network or computer. You must decide who you trust to have access to your system, and what threat they could pose.

There are several types of intruders, and it is useful to keep their different characteristics in mind as you are securing your systems.

- **The Curious** – This type of intruder is basically interested in finding out what type of system and data you have.
- **The Malicious** – This type of intruder is out to either bring down your systems, or deface your web page, or otherwise force you to spend time and money recovering from the damage he has caused.
- **The High-Profile Intruder** – This type of intruder is trying to use your system to gain popularity and infamy. He might use your high-profile system to advertise his abilities.
- **The Competition** – This type of intruder is interested in what data you have on your system. It might be someone who thinks you

have something that could benefit him, financially or otherwise.

- **The Borrowers** – This type of intruder is interested in setting up shop on your system and using its resources for their own purposes. He typically will run chat or IRC servers, porn archive sites, or even DNS servers.
- **The Leapfrogger** – This type of intruder is only interested in your system to use it to get into other systems. If your system is well-connected or a gateway to a number of internal hosts, you may well see this type trying to compromise your system.
- Vulnerability describes how well-protected your computer is from another network, and the potential for someone to gain unauthorized access.

What's at stake if someone breaks into your system? Of course the concerns of a dial-up home user will be different from those of a company connecting their machine to the *Internet*, or another large network.

How much time would it take to retrieve/recreate any data that was lost? An initial time investment now, spent securing a system, can save far more time later if a break-in means recreating lost data. Have you checked your backup strategy, and verified your data lately?

7.2.4. Developing A Security Policy

Create a simple, generic policy for your system that your users can readily understand and follow. It should protect the data you're safeguarding as well as the privacy of the users. Some things to consider adding are: who has access to the system (Can my friend use my ac-

count?), who's allowed to install software on the system, who owns what data, disaster recovery, and appropriate use of the system.

A generally-accepted security policy starts with the phrase

“That which is not permitted is prohibited”

This means that unless you grant access to a service for a user, that user shouldn't be using that service until you do grant access. Make sure the policies work on your regular user account. Saying, “Ah, I can't figure out this permissions problem, I'll just do it as root” can lead to security holes that are very obvious, and even ones that haven't been exploited yet.

RFC 1244 (<ftp://www.faqs.org/rfcs/rfc1244.txt>) is a document that describes how to create your own network security policy.

rfc1281 (<ftp://www.faqs.org/rfcs/rfc1281.txt>) is a document that shows an example security policy with detailed descriptions of each step.

Finally, you might want to look at the COAST policy archive (<ftp://coast.cs.purdue.edu/pub/doc/policy>) to see what some real-life security policies look like.

7.2.5. Means of Securing Your Site

This section will discuss various means with which you can secure the assets you have worked hard for: your local machine, your data, your users, your network, even your reputation. What would happen to your reputation if an intruder deleted some of your users' data? Or defaced your web site? Or published your company's corporate project plan for next quarter? If you are planning a network installation, there are many factors you must take into account before adding a single machine to your network.

Even if you have a single dialup PPP account, or just a small site, this does not mean intruders won't be interested in your systems. Large,

high-profile sites are not the only targets – many intruders simply want to exploit as many sites as possible, regardless of their size. Additionally, they may use a security hole in your site to gain access to other sites you’re connected to.

Intruders have a lot of time on their hands, and can avoid guessing how you’ve obscured your system just by trying all the possibilities. There are also a number of reasons an intruder may be interested in your systems, which we will discuss later.

7.2.5.1. Host Security

Perhaps the area of security on which administrators concentrate most is host-based security. This typically involves making sure your own system is secure, and hoping everyone else on your network does the same. Choosing good passwords, securing your host’s local network services, keeping good accounting records, and upgrading programs with known security exploits are among the things the local security administrator is responsible for doing. Although this is absolutely necessary, it can become a daunting task once your network becomes larger than a few machines.

7.2.5.2. Local Network Security

Network security is as necessary as local host security. With hundreds, thousands, or more computers on the same network, you can’t rely on each one of those systems being secure. Ensuring that only authorized users can use your network, building firewalls, using strong encryption, and ensuring there are no “rogue” (that is, unsecured) machines on your network are all part of the network security administrator’s duties.

This document will discuss some of the techniques used to secure your site, and hopefully show you some of the ways to prevent an intruder from gaining access to what you are trying to protect.

7.2.5.3. Security Through Obscurity

One type of security that must be discussed is "security through obscurity". This means, for example, moving a service that has known security vulnerabilities to a non-standard port in hopes that attackers won't notice it's there and thus won't exploit it. Rest assured that they can determine that it's there and will exploit it. Security provided solely through obscurity is no security at all. Simply because you may have a small site, or a relatively low profile, does not mean an intruder won't be interested in what you have or be able to find it. We'll discuss what you're protecting in the next sections.

7.2.6. Organization of this chapter

This chapter has been divided into a number of sections. They cover several broad security issues. The first, *Physical Security*, page 89, covers how you need to protect your physical machine from tampering. The second, *Local Security*, page 95, describes how to protect your system from tampering by local users. The third, *Files and Filesystem Security*, page 98, shows you how to setup your file-systems and permissions on your files. The next, *Password Security and Encryption*, page 107, discusses how to use encryption to better secure your machine and network. *Kernel Security*, page 120 discusses what kernel options you should set or be aware of for a more secure system. *Network Security*, page 126, describes how to better secure your GNU/Linux system from network attacks. *Security Preparation (before you go on-line)*, page 140, discusses how to prepare your machine(s) before bringing them on-line. Next, *What To Do During and After a Break-in*, page 143, discusses what to do when you detect a system compromise in progress or detect one that has recently happened. In *Security Sources*, page 147, some primary security resources are enumerated. The Q and A section *Frequently Asked Questions*, page 155, answers some frequently-asked questions, and finally a conclusion in *Conclusion*, page 157.

The two main points to realize when reading this chapter are:

- Be aware of your system. Check system logs such as `/var/log/messages` and keep an eye on your system, and
- Keep your system up-to-date by making sure you have installed the current versions of software and have upgraded per security alerts. Just doing this will help make your system markedly more secure.

7.3. Physical Security

The first layer of security you need to take into account is the physical security of your computer systems. Who has direct physical access to your machine? Should they? Can you protect your machine from their tampering? Should you?

How much physical security you need on your system is very dependent on your situation, and/or budget.

If you are a home user, you probably don't need a lot (although you might need to protect your machine from tampering by children or annoying relatives). If you are in a Lab, you need considerably more, but users will still need to be able to get work done on the machines. Many of the following sections will help out. If you are in an office, you may or may not need to secure your machine off-hours or while you are away. At some companies, leaving your console unsecured is a termination offense.

Obvious physical security methods such as locks on doors, cables, locked cabinets, and video surveillance are all good ideas, but beyond the scope of this chapter. :-)

7.3.1. Computer locks

Many modern *PC* cases include a “locking” feature. Usually this will be a socket on the front of the case that allows you to turn an included key to a locked or unlocked position. Case locks can help prevent someone from stealing your *PC*, or opening up the case and directly manipulating/stealing your hardware. They can also sometimes prevent someone from rebooting your computer from their own floppy or other hardware.

These case locks do different things according to the support in the motherboard and how the case is constructed. On many *PC*’s they make it so you have to break the case to get the case open. On some others, they will not let you plug in new keyboards or mice. Check your motherboard or case instructions for more information. This can sometimes be a very useful feature, even though the locks are usually very low-quality and can easily be defeated by attackers with locksmithing.

Some machines (most notably SPARCs and macs) have a dongle on the back that, if you put a cable through attackers would have to cut the cable or break the case to get into it. Just putting a padlock or combo lock through these can be a good deterrent to someone stealing your machine.

7.3.2. BIOS Security

The *BIOS* is the lowest level of software that configures or manipulates your x86-based hardware. *GRUB* and other *GNU/Linux* boot methods access the *BIOS* to determine how to boot up your *GNU/Linux* machine. Other hardware that *GNU/Linux* runs on has similar software (OpenFirmware on Macs and new Suns, Sun boot PROM, etc...). You can use your *BIOS* to prevent attackers from rebooting your machine and manipulating your *GNU/Linux* system.

Many *PC BIOS*s let you set a boot password. This doesn’t provide all that much security (the *BIOS* can be reset, or removed if someone

can get into the case), but might be a good deterrent (i.e. it will take time and leave traces of tampering). Similarly, on *S/Linux* (*GNU/Linux* for *SPARC*(tm) processor machines), your EEPROM can be set to require a boot-up password. This might slow attackers down.

Many *x86 BIOS*s also allow you to specify various other good security settings. Check your *BIOS* manual or look at it the next time you boot up. For example, some *BIOS*s disallow booting from floppy drives and some require passwords to access some *BIOS* features.

Note: If you have a server machine, and you set up a boot password, your machine will not boot up unattended. Keep in mind that you will need to come in and supply the password in the event of a power failure. ;(

7.3.3. Boot Loader Security

Keep in mind when setting all these passwords that you need to remember them :-). Also remember that these passwords will merely slow the determined attacker. They won't prevent someone from booting from a floppy, and mounting your root partition. If you are using security in conjunction with a boot loader, you might as well disable booting from a floppy in your computer's *BIOS*, and password-protect the *BIOS*.

Note: Once again, If you have a server machine, and you set up a boot password, your machine will not boot up unattended. Keep in mind that you will need to come in and supply the password in the event of a power failure. ;(

7.3.3.1. With GRUB

The various *GNU/Linux* boot loaders also can have a boot password set. *GRUB* is quite flexible in that sense: your default config file `/boot/grub/menu.lst` may contain a line allowing the loading of a new config file with different options (this new file may contain a new password to access another third config file and so on).

So you have to add a line in your file `/boot/grub/menu.lst`, something like:

```
password very_secret /boot/grub/menu2.lst
```

and of course generate a new config file `/boot/grub/menu2.lst` where you move unsecure entries previously removed from `/boot/grub/menu.lst`.

>From the grub info page:

- Command: `password passwd new-config-file`
Disable all interactive editing control (menu entry editor and command line). If the password `PASSWD` is entered, it loads the `NEW-CONFIG-FILE` as a new config file and restarts the GRUB Stage 2.

7.3.3.2. With LILO

LILO has password and restricted settings; password requires password at boot time, whereas restricted requires a boot-time password only if you specify options (such as `single`) at the *LILO* prompt.

>From the `lilo.conf` man page:

```
password=password
    The per-image option 'password=...' (see below)
    applies to all images.

restricted
    The per-image option 'restricted' (see below)
```

applies to all images.

```
password=password
    Protect the image by a password.
```

```
restricted
    A password is only required to boot the image if
    parameters are specified on the command line
    (e.g. single).
```

7.3.4. klock and vlock

If you wander away from your machine from time to time, it is nice to be able to “lock” your console so that no one can tamper with or look at your work. Two programs that do this are: `klock` and `vlock`.

`klock` is a *X* display locker. You can run `klock` from any `xterm` on your console and it will lock the display and require your password to unlock. Most desktop environment also propose this feature in their respective menus

`vlock` is a simple little program that allows you to lock some or all of the virtual consoles on your *GNU/Linux* box. You can lock just the one you are working in or all of them. If you just lock one, others can come in and use the console; they will just not be able to use your virtual console until you unlock it.

Of course locking your console will prevent someone from tampering with your work, but won’t prevent them from rebooting your machine or otherwise disrupting your work. It also does not prevent them from accessing your machine from another machine on the network and causing problems.

More importantly, it does not prevent someone from switching out of the *X Window System* entirely, and going to a normal virtual console login prompt, or to the VC that `X11` was started from, and suspending

it, thus obtaining your privileges. To stop this latter attack, try "exec startx" in place of "startx" if you traditionally start X from the command line.

7.3.5. Detecting Physical Security Compromises

The first thing to always note is when your machine was rebooted. Since *GNU/Linux* is a robust and stable OS, the only times your machine should reboot is when **you** take it down for OS upgrades, hardware swapping, or the like. If your machine has rebooted without you doing it, that may be a sign that an intruder has compromised it. Many of the ways that your machine can be compromised require the intruder to reboot or power off your machine.

Check for signs of tampering on the case and computer area. Although many intruders clean traces of their presence out of logs, it's a good idea to check through them all and note any discrepancy.

It is also a good idea to store log data at a secure location, such as a dedicated log server within your well-protected network. Once a machine has been compromised, log data becomes of little use as it most likely has also been modified by the intruder.

The *syslog* daemon can be configured to automatically send log data to a central *syslog* server, but this is typically sent in unencrypted, allowing an intruder to view data as it is being transferred. This may reveal information about your network that is not intended to be public. There are *syslog* daemons available that encrypt the data as it is being sent.

Also be aware that faking *syslog* messages is easy – with an exploit program having been published. *syslog* even accepts net log entries claiming to come from the local host without indicating their true origin.

Some things to check for in your logs:

- Short or incomplete logs.
- Logs containing strange timestamps.
- Logs with incorrect permissions or ownership.
- Records of reboots or restarting of services.
- missing logs.
- su entries or logins from strange places.

We will discuss system log data *Keep Track of Your System Accounting Data*, page 141 in this chapter.

7.4. Local Security

The next thing to take a look at is the security in your system against attacks from local users. Did we just say **local** users? Yes!

Getting access to a local user account is one of the first things that system intruders attempt while on their way to exploiting the root account. With lax local security, they can then “escalate” their normal user access to root access using a variety of bugs and poorly setup local services. If you make sure your local security is tight, then the intruder will have another hurdle to jump.

Local users can also cause a lot of havoc with your system even (especially) if they really are who they say they are. Providing accounts to people you don’t know or for whom you have no contact information for is a very bad idea.

7.4.1. Creating New Accounts

You should make sure you provide user accounts with only the minimal requirements for the task they need to do. If you provide your son

(age 10) with an account, you might want him to only have access to a word processor or drawing program, but be unable to delete data that is not his.

Several good rules of thumb when allowing other people legitimate access to your *GNU/Linux* machine:

- Give them the minimal amount of privileges they need.
- Be aware when/where they login from, or should be logging in from.
- Make sure you remove inactive accounts.
- The use of the same `userid` on all computers and networks is advisable to ease account maintenance, and permits easier analysis of log data.
- The creation of group `userids` should be absolutely prohibited. User accounts also provide accountability, and this is not possible with group accounts.

Many local user accounts that are used in security compromises have not been used in months or years. Since no one is using them, they provide the ideal attack vehicle - an attacker is far less likely to be noticed if the real account owner isn't around to spot the extra logins!

7.4.2. `root` Security

The most sought-after account on your machine is the `root` (superuser) account. This account has authority over the entire machine, which may also include authority over other machines on the network. Remember that you should only use the `root` account for very short, specific tasks, and should mostly run as a normal user. Even small mistakes made while logged in as the `root` user can cause problems. The less time you are on with `root` privileges, the safer you will be.

Several tricks to avoid messing up your own box as root:

- When doing some complex command, try running it first in a non-destructive way...especially commands that use globbing: e.g., if you want to do `rm -f foo*.bak`, first do `ls foo*.bak` and make sure you are going to delete the files you think you are. Using `echo` in place of destructive commands also sometimes works.
- Only become root to do single specific tasks. If you find yourself trying to figure out how to do something, go back to a normal user *shell* until you are **sure** what needs to be done by root.
- The command path for the root user is very important. The command path (that is, the `PATH` environment variable) specifies the directories in which the *shell* searches for programs. Try to limit the command path for the root user as much as possible, and **never** include `.` (which means “the current directory”) in your `PATH`. Additionally, never have writable directories in your search path, as this can allow attackers to modify or place new binaries in your search path, allowing them to run as root the next time you run that command.
- Never use the `rlogin/rsh/rexec` suite of tools (called the “r-utilities”) as root. They are subject to many sorts of attacks, and are downright dangerous when run as root. Never create a `.rhosts` file for root.
- The `/etc/securetty` file contains a list of terminals that root can login from. By default this is set to only the local virtual consoles (ttys). Be very wary of adding anything else to this file. You should be able to login remotely as your regular user account and then `su` if you need to (hopefully over `ssh` or other encrypted channel), so there is no need to be able to login directly as root.
- Always be slow and deliberate running as root. Your actions could affect a lot of things. Think before you type!

If you absolutely positively need to allow someone (hopefully very trusted) to have root access to your machine, there are a few tools that can help. `sudo` allows users to use their password to access a limited set of commands as `root`. This would allow you to, for instance, let a user be able to eject and mount removable media on your GNU/Linux box, but have no other root privileges. `sudo` also keeps a log of all successful and unsuccessful `sudo` attempts, allowing you to track down who used what command to do what. For this reason `sudo` works well even in places where a number of people have root access, because it helps you keep track of changes made.

Although `sudo` can be used to give specific users specific privileges for specific tasks, it does have several shortcomings. It should be used only for a limited set of tasks, like restarting a server, or adding new users. Any program that offers a *shell* escape will give root access to a user invoking it via `sudo`. This includes most editors, for example. Also, a program as innocuous as `/bin/cat` can be used to overwrite files, which could allow root to be exploited. Consider `sudo` as a means for accountability, and don't expect it to replace the root user and still be secure.

7.5. Files and Filesystem Security

A few minutes of preparation and planning ahead before putting your systems online can help to protect them and the data stored on them.

- There should never be a reason for users' home directories to allow SUID/SGID programs to be run from there. Use the `nosuid` option in `/etc/fstab` for partitions that are writable by others than root. You may also wish to use `nodev` and `noexec` on users' home partitions, as well as `/var`, thus prohibiting execution of programs, and creation of character or block devices, which should never be necessary anyway.

- If you are exporting filesystems using NFS, be sure to configure `/etc/exports` with the most restrictive access possible. This means not using wildcards, not allowing root write access, and exporting read-only wherever possible.
- Configure your users' file-creation `umask` to be as restrictive as possible. See *umask Settings*, page 101.
- If you are mounting filesystems using a network filesystem such as NFS, be sure to configure `/etc/fstab` with suitable restrictions. Typically, using `nodev`, `nosuid`, and perhaps `noexec`, are desirable.
- Set filesystem limits instead of allowing unlimited as is the default. You can control the per-user limits using the resource-limits PAM module and `/etc/pam.d/limits.conf`. For example, limits for group users might look like this:

```
@users      hard   core    0
@users      hard   nproc   50
@users      hard   rss     5000
```

This says to prohibit the creation of core files, restrict the number of processes to 50, and restrict memory usage per user to 5MB.

- The `/var/log/wtmp` and `/var/run/utmp` files contain the login records for all users on your system. Their integrity must be maintained because they can be used to determine when and from where a user (or potential intruder) has entered your system. These files should also have 644 permissions, without affecting normal system operation.
- The immutable bit can be used to prevent accidentally deleting or overwriting a file that must be protected. It also prevents someone from creating a symbolic link to the file (such symbolic links have been the source of attacks involving deleting `/etc/passwd` or `/etc/shadow`). See the `chattr(1)` man page for information on the immutable bit.

- *suid* and SGID files on your system are a potential security risk, and should be monitored closely. Because these programs grant special privileges to the user who is executing them, it is necessary to ensure that insecure programs are not installed. A favorite trick of crackers is to exploit SUID-root programs, then leave a SUID program as a backdoor to get in the next time, even if the original hole is plugged.

Find all SUID/SGID programs on your system, and keep track of what they are, so you are aware of any changes which could indicate a potential intruder. Use the following command to find all SUID/SGID programs on your system:

```
root# find / -type f \( -perm -04000 -o -perm -02000 \)
```

You can remove the *suid* or SGID permissions on a suspicious program with *chmod*, then restore them back if you absolutely feel it is necessary.

- World-writable files, particularly system files, can be a security hole if a cracker gains access to your system and modifies them. Additionally, world-writable directories are dangerous, since they allow a cracker to add or delete files as he wishes. To locate all world-writable files on your system, use the following command:

```
root# find / -perm -2 ! -type l -ls
```

and be sure you know why those files are writable. In the normal course of operation, several files will be world-writable, including some from */dev*, and symbolic links, thus the *! -type l* which excludes these from the previous *find* command.

- Unowned files may also be an indication an intruder has accessed your system. You can locate files on your system that have no owner, or belong to no group with the command:

```
root# find / -nouser -o -nogroup -print
```

- Finding `.rhosts` files should be a part of your regular system administration duties, as these files should not be permitted on your system. Remember, a cracker only needs one insecure account to potentially gain access to your entire network. You can locate all `.rhosts` files on your system with the following command:

```
root# find /home -name .rhosts -print
```

- Finally, before changing permissions on any system files, make sure you understand what you are doing. Never change permissions on a file because it seems like the easy way to get things working. Always determine why the file has that permission before changing it.

7.5.1. `umask` Settings

The `umask` command can be used to determine the default file creation mode on your system. It is the octal complement of the desired file mode. If files are created without any regard to their permissions settings, the user could inadvertently give read or write permission to someone that should not have this permission. Typical `umask` settings include 022, 027, and 077 (which is the most restrictive). Normally the `umask` is set in `/etc/profile`, so it applies to all users on the system. The file creation mask can be calculated by subtracting the desired value from 777. In other words, a `umask` of 777 would cause newly-created files to contain no read, write or execute permission for anyone. A mask of 666 would cause newly-created files to have a mask of 111. For example, you may have a line that looks like this:

```
# Set the user's default umask
umask 033
```

Be sure to make root's umask 077, which will disable read, write, and execute permission for other users, unless explicitly changed using `chmod`. In this case, newly-created directories would have 744 permissions, obtained by subtracting 033 from 777. Newly-created files using the 033 umask would have permissions of 644.

7.5.2. File Permissions

It's important to ensure that your system files are not open for casual editing by users and groups who shouldn't be doing such system maintenance.

Unix separates access control on files and directories according to three characteristics: owner, group, and other. There is always exactly one owner, any number of members of the group, and everyone else.

A quick explanation of *Unix* permissions:

Ownership - Which user(s) and group(s) retain(s) control of the permission settings of the node and parent of the node

Permissions - Bits capable of being set or reset to allow certain types of access to it. Permissions for directories may have a different meaning than the same set of permissions on files.

Read:

- To be able to view contents of a file
- To be able to read a directory

Write:

- To be able to add to or change a file
- To be able to delete or move files in a directory

Execute:

- To be able to run a binary program or *shell* script
- To be able to search in a directory, combined with read permission

Save Text Attribute: (For directories)

The “sticky bit” also has a different meaning when applied to directories than when applied to files. If the sticky bit is set on a directory, then a user may only delete files that he owns or for which he has explicit write permission granted, even when he has write access to the directory. This is designed for directories like */tmp*, which are world-writable, but where it may not be desirable to allow any user to delete files at will. The sticky bit is seen as a *t* in a long directory listing.

suid Attribute: (For Files)

This describes set-user-id permissions on the file. When the set user *ID* access mode is set in the owner permissions, and the file is executable, processes which run it are granted access to system resources based on user who owns the file, as opposed to the user who created the process. This is the cause of many “buffer overflow” exploits.

SGID Attribute: (For Files)

If set in the group permissions, this bit controls the “*set group id*” status of a file. This behaves the same way as *suid*, except the group is affected instead. The file must be executable for this to have any effect.

SGID Attribute: (For directories)

If you set the SGID bit on a directory (with `chmod g+s directory`), files created in that directory will have their group set to the directory's group.

You - The owner of the file

Group - The group you belong to

Everyone - Anyone on the system that is not the owner or a member of the group

File Example:

```
-rw-r--r-- 1 kevin users      114 Aug 28 1997 .zlogin
1st bit - directory?          (no)
2nd bit - read by owner?      (yes, by kevin)
3rd bit - write by owner?     (yes, by kevin)
4th bit - execute by owner?   (no)
5th bit - read by group?      (yes, by users)
6th bit - write by group?     (no)
7th bit - execute by group?   (no)
8th bit - read by everyone?   (yes, by everyone)
9th bit - write by everyone?  (no)
10th bit - execute by everyone? (no)
```

The following lines are examples of the minimum sets of permissions that are required to perform the access described. You may want to give more permission than what's listed here, but this should describe what these minimum permissions on files do:

```
--r----- Allow read access to the file by owner
--w----- Allows the owner to modify or delete the file
          (Note that anyone with write permission to the directory
           the file is in can overwrite it and thus delete it)
---x----- The owner can execute this program, but not shell scripts,
           which still need read permission
```



```
---s----- Will execute with effective User ID = to owner
-----s-   Will execute with effective Group ID = to group
-rw-----T No update of "last modified time". Usually used for swap
            files
---t----- No effect. (formerly sticky bit)
```

Directory Example:

```
drwxr-xr-x 3 kevin users          512 Sep 19 13:47 .public_html/
1st bit - directory?              (yes, it contains many files)
2nd bit - read by owner?          (yes, by kevin)
3rd bit - write by owner?         (yes, by kevin)
4th bit - execute by owner?       (yes, by kevin)
5th bit - read by group?          (yes, by users)
6th bit - write by group?         (no)
7th bit - execute by group?       (yes, by users)
8th bit - read by everyone?       (yes, by everyone)
9th bit - write by everyone?      (no)
10th bit - execute by everyone?   (yes, by everyone)
```

The following lines are examples of the minimum sets of permissions that are required to perform the access described. You may want to give more permission than what's listed, but this should describe what these minimum permissions on directories do:

```
dr----- The contents can be listed, but file attributes can't be read
d--x----- The directory can be entered, and used in full execution
            paths
dr-x----- File attributes can be read by owner
d-wx----- Files can be created/deleted, even if the directory
            isn't the current one
d-----x-t Prevents files from deletion by others with write
            access. Used on /tmp
d---s--s-- No effect
```

System configuration files (usually in /etc) are usually mode 640 (-rw-r-----), and owned by root. Depending on your site's security requirements, you might adjust this. Never leave any system files writable by a group or everyone. Some configuration files, including /etc/shadow, should only be readable by root, and directories in /etc should at least not be accessible by others.

suid shell Scripts

suid shell scripts are a serious security risk, and for this reason the kernel will not honor them. Regardless of how secure you think the *shell* script is, it can be exploited to give the cracker a *root shell*.

7.5.3. Integrity Checking

Another very good way to detect local (and also network) attacks on your system is to run an integrity checker like *Tripwire*, *Aide* or *Osiris*. These integrity checkers run number of checksums on all your important binaries and config files and compares them against a database of former, known-good values as a reference. Thus, any changes in the files will be flagged.

It's a good idea to install these sorts of programs onto a floppy, and then physically set the write protect on the floppy. This way intruders can't tamper with the integrity checker itself or change the database. Once you have something like this setup, it's a good idea to run it as part of your normal security administration duties to see if anything has changed.

You can even add a *crontab* entry to run the checker from your floppy every night and mail you the results in the morning. Something like:

```
# set mailto
MAILTO=kevin
# run Tripwire
15 05 * * * root /usr/local/adm/tcheck/tripwire
```

will mail you a report each morning at 5:15am.

Integrity checkers can be a godsend to detecting intruders before you would otherwise notice them. Since a lot of files change on the average system, you have to be careful what is cracker activity and what is your own doing.

You can find the open sourced version of Tripwire at TripWire (<http://www.tripwire.org>) free of charge. Manuals and support can be purchased.

Aide can be found at <http://www.cs.tut.fi/~rammer/aide.html> (<http://www.cs.tut.fi/~rammer/aide.html>).

Osiris can be found at <http://www.shmoo.com/osiris/> (<http://www.shmoo.com/osiris/>).

7.5.4. Trojan Horses

“Trojan Horses” are named after the fabled ploy in Homer’s “Iliad”. The idea is that a cracker distributes a program or binary that sounds great, and encourages other people to download it and run it as root. Then the program can compromise their system while they are not paying attention. While they think the binary they just pulled down does one thing (and it might very well), it also compromises their security.

You should take care of what programs you install on your machine. **MandrakeSoft** provides MD5 checksums and PGP signatures on it’s RPM files so you can verify you are installing the real thing. You should never run any unfamiliar binary, for which you don’t have the source, as root! Few attackers are willing to release source code to public scrutiny.

Although it can be complex, make sure you are getting the source for a program from its real distribution site. If the program is going to run as root, make sure either you or someone you trust has looked over the source and verified it, or at least checked the MD5 checksum and/or PGP signatures.

7.6. Password Security and Encryption

Note: Most of encryption programs described in this chapter are available by FTP for your **Linux-Mandrake** distribution. Consult <http://www.linux-mandrake.com/en/fcrypto.php3> (<http://www.linux-mandrake.com/en/fcrypto.php3>) for more informations.

One of the most important security features used today are passwords. It is important for both you and all your users to have secure, unguessable passwords. Your **Linux-Mandrake** distributions include `passwd` program that do not allow you to set a easily guessable password. Make sure your `passwd` program is up to date.

In-depth discussion of encryption is beyond the scope of this chapter, but an introduction is in order. Encryption is very useful, possibly even necessary in this day and age. There are all sorts of methods of encrypting data, each with its own set of characteristics.

Most *Unices* (and *GNU/Linux* is no exception) primarily use a one-way encryption algorithm, called DES (Data Encryption Standard) to encrypt your passwords. This encrypted password is then stored in `/etc/shadow`. When you attempt to login, the password you type in is encrypted again and compared with the entry in the file that stores your passwords. If they match, it must be the same password, and you are allowed access. Although DES is a two-way encryption algorithm (you can code and then decode a message, given the right keys), the variant that most unices use is one-way. This means that it should not be possible to reverse the encryption to get the password from the contents of `/etc/shadow`.

Brute force attacks, such as “Crack” or “John the Ripper” (see Section “*Crack*” and “*John the Ripper*”, page 117) can often guess passwords unless your password is sufficiently random. PAM modules (see below) allow you to use a different encryption routine with your passwords (MD5 or the like). You can use Crack to your advantage, as

well. Consider periodically running Crack against your own password database, to find insecure passwords. Then contact the offending user, and instruct him to change his password.

You can go to CERN (http://consult.cern.ch/writeup/security/security_3.html) for information on how to choose a good password. A few horrible passwords for an account named jack are:

1. jack
2. jack1
3. 1kcaj

All of these would be quickly found by Crack or other common password brute forcing programs.

7.6.1. PGP and Public-Key Cryptography

Public-key cryptography, such as that used for PGP, uses one key for encryption, and one key for decryption. Traditional cryptography, however, uses the same key for encryption and decryption; this key must be known to both parties, and thus somehow transferred from one to the other securely.

To alleviate the need to securely transmit the encryption key, public-key encryption uses two separate keys: a public key and a private key. Each person's public key is available by anyone to do the encryption, while at the same time each person keeps his or her private key to decrypt messages encrypted with the correct public key.

There are advantages to both public key and private key cryptography, and you can read about those differences in the RSA Cryptography FAQ (<http://www.rsa.com/rsalabs/newfaq/>), listed at the end of this section.

PGP (Pretty Good Privacy) is well-supported on *GNU/Linux*. Versions 2.6.2 and 5.0 are known to work well. For a good primer

on *PGP* and how to use it, take a look at the *PGP* FAQ: [pgp.com \(http://www.pgp.com/service/export/faq/55faq.cgi\)](http://www.pgp.com/service/export/faq/55faq.cgi)

Be sure to use the version that is applicable to your country. Due to export restrictions by the US Government, strong-encryption is prohibited from being transferred in electronic form outside the country.

US export controls are now managed by EAR (Export Administration Regulations). They are no longer governed by ITAR.

There is also a step-by-step guide for configuring *PGP* on *GNU/Linux* available at LinuxFocus (<http://mercury.chem.pitt.edu/~angel/LinuxFocus/English/November1997/article7.html>). It was written for the international version of *PGP*, but is easily adaptable to the United States version. You may also need a patch for some of the latest versions of *GNU/Linux*; the patch is available at <ftp://metalab.unc.edu/pub/Linux/apps/crypto> (<ftp://metalab.unc.edu/pub/Linux/apps/crypto>).

There is a project working on a free re-implementation of *PGP* with open source. GnuPG is a complete and free replacement for *PGP*. Because it does not use IDEA or RSA it can be used without any restrictions. *GnuPG* is nearly in compliance with OpenPGP (<http://www.faqs.org/rfcs/rfc2440.html>). See the GNU Privacy Guard web page for more information: <http://www.gnupg.org/> (<http://www.gnupg.org/>).

More information on cryptography can be found in the RSA cryptography FAQ, available at <http://www.rsa.com/rsalabs/newfaq/> (<http://www.rsa.com/rsalabs/newfaq/>). Here you will find information on such terms as “Diffie-Hellman”, “public-key cryptography”, “digital certificates”, etc.

7.6.2. SSL, S-HTTP, HTTPS and S/MIME

Often users ask about the differences between the various security and encryption protocols, and how to use them. While this isn’t an encryp-

tion document, it is a good idea to explain briefly what each protocol is, and where to find more information.

- **SSL:** - SSL, or Secure Sockets Layer, is an encryption method developed by Netscape to provide security over the *Internet*. It supports several different encryption protocols, and provides client and server authentication. SSL operates at the transport layer, creates a secure encrypted channel of data, and thus can seamlessly encrypt data of many types. This is most commonly seen when going to a secure site to view a secure online document with *Communicator*, and serves as the basis for secure communications with *Communicator*, as well as many other Netscape Communications data encryption. More information can be found at <http://www.consensus.com/security/ssl-talk-faq.html> (<http://www.consensus.com/security/ssl-talk-faq.html>). Information on Netscape's other security implementations, and a good starting point for these protocols is available at Netscape (<http://home.netscape.com/info/security-doc.html>).
- **S-HTTP:** - S-HTTP is another protocol that provides security services across the *Internet*. It was designed to provide confidentiality, authentication, integrity, and non-repudiability [cannot be mistaken for someone else] while supporting multiple key-management mechanisms and cryptographic algorithms via option negotiation between the parties involved in each transaction. S-HTTP is limited to the specific software that is implementing it, and encrypts each message individually. [From RSA Cryptography FAQ, page 138]
- **S/MIME:** - S/MIME, or Secure Multipurpose Internet Mail Extension, is an encryption standard used to encrypt electronic mail and other types of messages on the *Internet*. It is an open standard developed by RSA, so it is likely we will see it on *GNU/Linux* one day soon. More information on S/MIME can be found at netscape (<http://home.netscape.com/assist/security/smime/overview.html>).

7.6.3. IPSEC Implementations

Along with CIPE, and other forms of data encryption, there are also several other implementations of IPSEC for *GNU/Linux*. IPSEC is an effort by the IETF to create cryptographically-secure communications at the IP network level, and to provide authentication, integrity, access control, and confidentiality. Information on IPSEC and *Internet* draft can be found at <http://www.ietf.org/html.charters/ipsec-charter.html> (<http://www.ietf.org/html.charters/ipsec-charter.html>). You can also find links to other protocols involving key management, and an IPSEC mailing list and archives.

The x-kernel *GNU/Linux* implementation, which is being developed at the University of Arizona, uses an object-based framework for implementing network protocols called x-kernel, and can be found at <http://www.cs.arizona.edu/xkernel/hpcc-blue/linux.html> (<http://www.cs.arizona.edu/xkernel/hpcc-blue/linux.html>). Most simply, the x-kernel is a method of passing messages at the kernel level, which makes for an easier implementation.

Another freely-available IPSEC implementation is the *GNU/Linux* FreeS/WAN IPSEC. Their web page states, “These services allow you to build secure tunnels through untrusted networks. Everything passing through the untrusted net is encrypted by the IPSEC gateway machine and decrypted by the gateway at the other end. The result is Virtual Private Network or VPN. This is a network which is effectively private even though it includes machines at several different sites connected by the insecure *Internet*.”

It’s available for download from <http://www.xs4all.nl/~freeswan/> (<http://www.xs4all.nl/~freeswan/>), and has just reached 1.0 at the time of this writing.

As with other forms of cryptography, it is not distributed with the kernel by default due to export restrictions.

7.6.4. ssh (Secure shell) and stelnet

ssh and stelnet are suites of programs that allow you to login to remote systems and have a encrypted connection.

openssh is a suite of programs used as a secure replacement for rlogin, rsh and rcp. It uses public-key cryptography to encrypt communications between two hosts, as well as to authenticate users. It can be used to securely login to a remote host or copy data between hosts, while preventing man-in-the-middle attacks (session hijacking) and DNS spoofing. It will perform data compression on your connections, and secure X11 communications between hosts.

There are several ssh implementations now. The original commercial implementation by Data Fellows can be found at the ssh home page available at <http://www.datafellows.com> (<http://www.datafellows.com>).

The excellent *Openssh* implementation is based on a early version of the *datafellows ssh* and has been totally reworked to not include any patented or proprietary pieces. It is free and under a BSD license. It can be found at: <http://www.openssh.com> (<http://www.openssh.com>).

There is also a open source project to re-implement ssh from the ground up called “psst...”. For more information see: psst (<http://www.net.lut.ac.uk/psst/>).

You can also use ssh from your *Windows* workstation to your *GNU/Linux* ssh server. There are several freely available *Windows* client implementations, including the one at therapy ssh (<http://guardian.htu.tuwien.ac.at/therapy/ssh/>) as well as a commercial implementation from DataFellows, at datafellows (<http://www.datafellows.com>).

SSLey is a free implementation of Netscape’s Secure Sockets Layer, developed by Eric Young. It includes several applications, such as “Secure telnet”, a module for *Apache*, several databases, as well as several algorithms including DES, IDEA and “Blowfish”.

Using this library, a secure telnet replacement has been created that does encryption over a telnet connection. Unlike SSH, *steln*et uses SSL, the Secure Sockets Layer protocol developed by Netscape. You can find Secure telnet and Secure FTP by starting with the *SSLeay* FAQ, available at <http://www.psy.uq.oz.au/~ftp/Crypto/> (<http://www.psy.uq.oz.au/~ftp/Crypto/>).

Note: The *OpenSSL* Project is based on *SSLeay* and is meant to develop a robust, commercial-grade, full-featured, and Open Source toolkit implementing the Secure Sockets Layer (SSL v2/v3) and Transport Layer Security (TLS v1) protocols as well as a full-strength general purpose cryptography library. For more information about this project, consult the OpenSSL homepage (www.openssl.org). There is a large list of applications based on *OpenSSL* at OpenSSL related applications (<http://www.openssl.org/related/apps.html>).

SRP is another secure telnet/ftp implementation. From their web page:

“The SRP project is developing secure *Internet* software for free worldwide use. Starting with a fully-secure *Telnet* and *FTP* distribution, we hope to supplant weak networked authentication systems with strong replacements that do not sacrifice user-friendliness for security. Security should be the default, not an option!”

For more information, go to <http://srp.stanford.edu/srp> (<http://srp.stanford.edu/srp>).

7.6.5. PAM - Pluggable Authentication Modules

Your version of **Linux-Mandrake** distribution ships with a unified authentication scheme called PAM. PAM allows you to change your authentication methods and requirements on the fly, and encapsulate all local authentication methods without recompiling any of your

binaries. Configuration of PAM is beyond the scope of this chapter, but be sure to take a look at the PAM web site for more information. <http://www.kernel.org/pub/linux/libs/pam/index.html> (<http://www.kernel.org/pub/linux/libs/pam/index.html>).

Just a few of the things you can do with PAM:

- Use encryption other than DES for your passwords. (Making them harder to brute-force decode)
- Set resource limits on all your users so they can't perform denial-of-service attacks (number of processes, amount of memory, etc.)
- Enable shadow passwords (see below) on the fly
- allow specific users to login only at specific times from specific places

Within a few hours of installing and configuring your system, you can prevent many attacks before they even occur. For example, use PAM to disable the system-wide usage of `.rhosts` files in user's home directories by adding these lines to `/etc/pam.d/rlogin`:

```
#
# Disable rsh/rlogin/rexec for users
#
login auth required pam_rhosts_auth.so no_rhosts
```

7.6.6. Cryptographic IP Encapsulation (CIPE)

The primary goal of this software is to provide a facility for secure (against eavesdropping, including traffic analysis, and faked message injection) subnetwork interconnection across an insecure packet network such as the *Internet*.

CIPE encrypts the data at the network level. Packets traveling between hosts on the network are encrypted. The encryption engine is placed near the driver which sends and receives packets.

This is unlike SSH, which encrypts the data by connection, at the socket level. A logical connection between programs running on different hosts is encrypted.

CIPE can be used in tunneling, in order to create a Virtual Private Network. Low-level encryption has the advantage that it can be made to work transparently between the two networks connected in the VPN, without any change to application software.

Summarized from the CIPE documentation:

“The IPSEC standards define a set of protocols which can be used (among other things) to build encrypted VPNs. However, IPSEC is a rather heavyweight and complicated protocol set with a lot of options, implementations of the full protocol set are still rarely used and some issues (such as key management) are still not fully resolved. CIPE uses a simpler approach, in which many things which can be parameterized (such as the choice of the actual encryption algorithm used) are an install-time fixed choice. This limits flexibility, but allows for a simple (and therefore efficient, easy to debug...) implementation.”

Further information can be found at inka (<http://www.inka.de/~bigred/devel/cipe.html>)

As with other forms of cryptography, it is not distributed with the kernel by default due to export restrictions.

7.6.7. Kerberos

Kerberos is an authentication system developed by the Athena Project at MIT. When a user logs in, *Kerberos* authenticates that user (using a password), and provides the user with a way to prove her identity to other servers and hosts scattered around the network.

This authentication is then used by programs such as `rlogin` to allow the user to login to other hosts without a password (in place of the `.rhosts` file). This authentication method can also be used by the mail system in order to guarantee that mail is delivered to the correct person, as well as to guarantee that the sender is who he claims to be.

Kerberos and the other programs that come with it, prevent users from “spoofing” the system into believing they are someone else. Unfortunately, installing *Kerberos* is very intrusive, requiring the modification or replacement of numerous standard programs.

You can find more information about *kerberos* by looking at the *kerberos* FAQ (<http://www.cis.ohio-state.edu/hypertext/faq/usenet/kerberos-faq/general/faq.html>), and the code can be found at <http://nii.isi.edu/info/kerberos/> (<http://nii.isi.edu/info/kerberos/>).

[From: Stein, Jennifer G., Clifford Neuman, and Jeffrey L. Schiller. "Kerberos: An Authentication Service for Open Network Systems." USENIX Conference Proceedings, Dallas, Texas, Winter 1998.]

Kerberos should not be your first step in improving security of your host. It is quite involved, and not as widely used as, say, SSH.

7.6.8. “Crack” and “John the Ripper”

If for some reason your `passwd` program is not enforcing hard-to-guess passwords, you might want to run a password-cracking program and make sure your users’ passwords are secure.

Password cracking programs work on a simple idea: they try every word in the dictionary, and then variations on those words, encrypting each one and checking it against your encrypted password. If they get a match they know what your password is.

There are a number of programs out there...the two most notable of which are *Crack* and *John the Ripper* ([false.com](http://www.false.com/security/john/index.html) (<http://www.false.com/security/john/index.html>)). They will take up a lot

of your CPU time, but you should be able to tell if an attacker could get in using them by running them first yourself and notifying users with weak passwords. Note that an attacker would have to use some other hole first in order to read your `/etc/shadow` file, but such holes are more common than you might think.

Because security is only as strong as the most insecure host, it is worth mentioning that if you have any *Windows* machines on your network, you should check out *L0phtCrack*, a *Crack* implementation for *Windows*. It's available from <http://www.l0pht.com> (<http://www.l0pht.com>)

7.6.9. CFS - Cryptographic File System and TCFS - Transparent Cryptographic File System

CFS is a way of encrypting an entire directory tree and allowing users to store encrypted files on them. It uses an NFS server running on the local machine. RPMs are available at <http://www.zedz.net/redhat/> (<http://www.zedz.net/redhat/>), and more information on how it all works is at [att\(ftp://ftp.research.att.com/dist/mab/\)](ftp://ftp.research.att.com/dist/mab/).

TCFS improves on CFS by adding more integration with the file system, so that it's transparent to users that the file system is encrypted. More information at: <http://edu-gw.dia.unisa.it/tcfs/> (<http://edu-gw.dia.unisa.it/tcfs/>).

It also need not be used on entire filesystems. It works on directory trees as well.

7.6.10. X11, SVGA and display security

7.6.10.1. X11

It's important for you to secure your graphical display to prevent attackers from grabbing your passwords as you type them, reading doc-

uments or information you are reading on your screen, or even using a hole to gain root access. Running remote *X* applications over a network also can be fraught with peril, allowing sniffers to see all your interaction with the remote system.

X has a number of access-control mechanisms. The simplest of them is host-based: you use `xhost` to specify the hosts that are allowed access to your display. This is not very secure at all, because if someone has access to your machine, they can `xhost + their machine` and get in easily. Also, if you have to allow access from an untrusted machine, anyone there can compromise your display.

When using `xdm` (*X Display Manager*), or its *KDE* counterpart: *KDM*, to log in, you get a much better access method: MIT-MAGIC-COOKIE-1. A 128-bit “cookie” is generated and stored in your `.Xauthority` file. If you need to allow a remote machine access to your display, you can use the `xauth` command and the information in your `.Xauthority` file to provide access to only that connection. See the Remote-X-Apps mini-howto, available at <http://metalab.unc.edu/LDP/HOWTO/mini/Remote-X-Apps.html> (`http://metalab.unc.edu/LDP/HOWTO/mini/Remote-X-Apps.html`).

You can also use `ssh` (see *ssh (Secure shell)* and *stelnet*, page 113, above) to allow secure *X* connections. This has the advantage of also being transparent to the end user, and means that no unencrypted data flows across the network.

Take a look at the `Xsecurity` man page for more information on *X* security. The safe bet is to use `xdm` to login to your console and then use `ssh` to go to remote sites on which you wish to run *X* programs.

7.6.10.2. SVGA

*SVGA*lib programs are typically *suid-root* in order to access all your GNU/Linux machine’s video hardware. This makes them very dangerous. If they crash, you typically need to reboot your machine to get a usable console back. Make sure any *SVGA* programs you are running

are authentic, and can at least be somewhat trusted. Even better, don't run them at all.

7.6.10.3. GGI (Generic Graphics Interface project)

The *GNU/Linux* GGI project is trying to solve several of the problems with video interfaces on *GNU/Linux*. GGI will move a small piece of the video code into the *GNU/Linux* kernel, and then control access to the video system. This means GGI will be able to restore your console at any time to a known good state. They will also allow a secure attention key, so you can be sure that there is no Trojan horse login program running on your console. <http://synergy.caltech.edu/~ggi/> (<http://synergy.caltech.edu/~ggi/>)

7.7. Kernel Security

This is a description of the kernel configuration options that relate to security, and an explanation of what they do, and how to use them.

As the kernel controls your computer's networking, it is important that it be very secure, and not be compromised. To prevent some of the latest networking attacks, you should try to keep your kernel version current. You can find new kernels at <ftp://ftp.kernel.org> (<ftp://ftp.kernel.org>) or from packages updates available through Mandrake-Update.

There is also a international group providing a single unified crypto patch to the mainstream *GNU/Linux* kernel. This patch provides support for a number of cryptographic subsystems and things that cannot be included in the mainstream kernel due to export restrictions. For more information, visit their web page at: <http://www.kerneli.org> (<http://www.kerneli.org>).

7.7.1. Kernel Compile Options

For 2.2.x kernels, the following options apply. You should see these options during the kernel configuration process. Many of the comments here are from `/usr/src/linux/Documentation/Configure.help`, which is the same document that is referenced while using the Help facility during the `make config` stage of compiling the kernel.

- Network Firewalls (`CONFIG_FIREWALL`)

This option should be on if you intend to run any firewalling or masquerading on your *GNU/Linux* machine. If it's just going to be a regular client machine, it's safe to say no.

- IP: forwarding/gatewaying (`CONFIG_IP_FORWARD`)

If you enable IP forwarding, your *GNU/Linux* box essentially becomes a router. If your machine is on a network, you could be forwarding data from one network to another, and perhaps subverting a firewall that was put there to prevent this from happening. Normal dial-up users will want to disable this, and other users should concentrate on the security implications of doing this. Firewall machines will want this enabled, and used in conjunction with firewall software.

You can enable IP forwarding dynamically using the following command:

```
root# echo 1 > /proc/sys/net/ipv4/ip_forward
```

and disable it with the command:

```
root# echo 0 > /proc/sys/net/ipv4/ip_forward
```

- IP: syn cookies (`CONFIG_SYN_COOKIES`)

a “SYN Attack” is a denial of service (DoS) attack that consumes all the resources on your machine, forcing you to reboot. We can’t think of a reason you wouldn’t normally enable this. In the 2.1 kernel series this config option nearly allows syn cookies, but does not enable them. To enable them, you have to do:

```
root# echo 1 > /proc/sys/net/ipv4/tcp_syncookies
```

- IP: Firewalling (CONFIG_IP_FIREWALL)

This option is necessary if you are going to configure your machine as a firewall, do masquerading, or wish to protect your dial-up workstation from someone entering via your PPP dial-up interface.

- IP: firewall packet logging (CONFIG_IP_FIREWALL_VERBOSE)

This option gives you information about packets your firewall received, like sender, recipient, port, etc.

- IP: Drop source routed frames (CONFIG_IP_NOSR)

This option should be enabled. Source routed frames contain the entire path to their destination inside of the packet. This means that routers through which the packet goes do not need to inspect it, and just forward it on. This could lead to data entering your system that may be a potential exploit.

- IP: masquerading (CONFIG_IP_MASQUERADE)

If one of the computers on your local network for which your *GNU/Linux* box acts as a firewall wants to send something to the outside, your box can “masquerade” as that host, i.e., it forwards the traffic to the intended destination, but makes it look like it came from the

firewall box itself. See <http://www.indyramp.com/masq> (<http://www.indyramp.com/masq>) for more information.

- IP: ICMP masquerading (CONFIG_IP_MASQUERADE_ICMP)

This option adds ICMP masquerading to the previous option of only masquerading TCP or UDP traffic.

- IP: transparent proxy support (CONFIG_IP_TRANSPARENT_PROXY)

This enables your *GNU/Linux* firewall to transparently redirect any network traffic originating from the local network and destined for a remote host to a local server, called a “transparent proxy server”. This makes the local computers think they are talking to the remote end, while in fact they are connected to the local proxy. See the IP-Masquerading *HOWTO* and <http://www.indyramp.com/masq> (<http://www.indyramp.com/masq>) for more information.

- IP: always defragment (CONFIG_IP_ALWAYS_DEFRAG)

Generally this option is disabled, but if you are building a firewall or a masquerading host, you will want to enable it. When data is sent from one host to another, it does not always get sent as a single packet of data, but rather it is fragmented into several pieces. The problem with this is that the port numbers are only stored in the first fragment. This means that someone can insert information into the remaining packets that isn’t supposed to be there. It could also prevent a teardrop attack against an internal host that is not yet itself patched against it.

- Packet Signatures (CONFIG_NCPFS_PACKET_SIGNING)

This is an option that will sign NCP packets for stronger security. Normally you can leave it off, but it is there if you do need it.

- IP: Firewall packet netlink device (CONFIG_IP_FIREWALL_NETLINK)

This is a really neat option that allows you to analyze the first 128 bytes of the packets in a user-space program, to determine if you would like to accept or deny the packet, based on its validity.

- Socket Filtering (CONFIG_FILTER)

For most people, it's safe to say no to this option. This option allows you to connect a userspace filter to any socket and determine if packets should be allowed or denied. Unless you have a very specific need and are capable of programming such a filter, you should say no. Also note that as of this writing, all protocols were supported except TCP.

- Port Forwarding

Port Forwarding is an addition to IP Masquerading which allows some forwarding of packets from outside to inside a firewall on given ports. This could be useful if, for example, you want to run a web server behind the firewall or masquerading host and that web server should be accessible from the outside world. An external client sends a request to port 80 of the firewall, the firewall forwards this request to the web server, the web server handles the request and the results are sent through the firewall to the original client. The client thinks that the firewall machine itself is running the web server. This can also be used for load balancing if you have a farm of identical web servers behind the firewall. Information about this feature is available from monmouth (<http://www.monmouth.demon.co.uk/ipsubs/portforwarding.html>). For general info, please see compsoc (<ftp://ftp.compsoc.net/users/steve/ipportfw/linux21/>).

- Socket Filtering (CONFIG_FILTER)

Using this option, user-space programs can attach a filter to any socket and thereby tell the kernel that it should allow or disallow certain types of data to get through the socket. *GNU/Linux* socket filtering works on all socket types except TCP for now. See the text file `/usr/src/linux/Documentation/networking/filter.txt` for more information.

- IP: Masquerading

The 2.2 kernel masquerading has been improved. It provides additional support for masquerading special protocols, etc. Be sure to read the IP Chains *HOWTO* for more information.

7.7.2. Kernel Devices

There are a few block and character devices available on *GNU/Linux* that will also help you with security.

The two devices `/dev/random` and `/dev/urandom` are provided by the kernel to provide random data at any time.

Both `/dev/random` and `/dev/urandom` should be secure enough to use in generating *PGP* keys, `ssh` challenges, and other applications where secure random numbers are required. Attackers should be unable to predict the next number given any initial sequence of numbers from these sources. There has been a lot of effort put in to ensuring that the numbers you get from these sources are random in every sense of the word.

The only difference between the two devices, is that `/dev/random` runs out of random bytes and it makes you wait for more to be accumulated. Note that on some systems, it can block for a long time waiting for new user-generated entropy to be entered into the system. So you have

to use care before using `/dev/random`. (Perhaps the best thing to do is to use it when you're generating sensitive keying information, and you tell the user to pound on the keyboard repeatedly until you print out "OK, enough".)

`/dev/random` is high quality entropy, generated from measuring the inter-interrupt times etc. It blocks until enough bits of random data are available.

`/dev/urandom` is similar, but when the store of entropy is running low, it'll return a cryptographically strong hash of what there is. This isn't as secure, but it's enough for most applications.

You might read from the devices using something like:

```
root# head -c 6 /dev/urandom | mimencode
```

This will print six random characters on the console, suitable for password generation. You can find `mimencode` in the `metamail` package.

See `/usr/src/linux/drivers/char/random.c` for a description of the algorithm.

7.8. Network Security

Network security is becoming more and more important as people spend more and more time connected. Compromising network security is often much easier than compromising physical or local security, and is much more common.

There are a number of good tools to assist with network security, and more and more of them are shipped with your **Linux-Mandrake** distribution, either in the main CDROM, `contribs`, or through the FTP crypto server (see above).

7.8.1. Packet Sniffers

One of the most common ways intruders gain access to more systems on your network is by employing a packet sniffer on a already compromised host. This "sniffer" just listens on the *Ethernet* port for things like `passwd` and `login` and `su` in the packet stream and then logs the traffic after that. This way, attackers gain passwords for systems they are not even attempting to break into. Clear-text passwords are very vulnerable to this attack.

Example: Host A has been compromised. Attacker installs a sniffer. Sniffer picks up admin logging into Host B from Host C. It gets the admin's personal password as they login to B. Then, the admin does a `su` to fix a problem. They now have the root password for Host B. Later the admin lets someone `telnet` from his account to Host Z on another site. Now the attacker has a password/*login* on Host Z.

In this day and age, the attacker doesn't even need to compromise a system to do this: they could also bring a laptop or *PC* into a building and tap into your net.

Using `ssh` or other encrypted password methods thwarts this attack. Things like APOP for POP accounts also prevents this attack. (Normal POP logins are very vulnerable to this, as is anything that sends clear-text passwords over the network.)

7.8.2. System services and `tcp_wrappers`

Before you put your *GNU/Linux* system on **ANY** network the first thing to look at is what services you need to offer. Services that you do not need to offer should be disabled so that you have one less thing to worry about and attackers have one less place to look for a hole.

There are a number of ways to disable services under *GNU/Linux*. You can look at your `/etc/inetd.conf` file and see what services are being offered by your `inetd`. Disable any that you do not need by

commenting them out (`#` at the beginning of the line), and then restart your `inetd` service.

You can also remove (or comment out) services in your `/etc/services` file. This will mean that local clients will also be unable to find the service (i.e., if you remove `ftp`, and try and `ftp` to a remote site from that machine it will fail with an `unknown service` message). It's usually not worth the trouble to remove services from `/etc/services`, since it provides no additional security. If a local person wanted to use `ftp` even though you had commented it out, they would make their own client that use the common FTP port and would still work fine.

Some of the services you might want to leave enabled are:

- `ftp`
- `telnet` (or `ssh`)
- mail, such as `pop-3` or `imap`
- `identd`

If you know you are not going to use some particular package, you can also delete it entirely. `rpm -e packagename` will erase an entire package.

Additionally, you really want to disable the `rsh/rlogin/rcp` utilities, including `login` (used by `rlogin`), `shell` (used by `rcp`), and `exec` (used by `rsh`) from being started in `/etc/inetd.conf`. These protocols are extremely insecure and have been the cause of exploits in the past.

You should check your `/etc/rc.d/rc[0-9].d`, and see if any of the servers started in that directory are not needed. The files in that directory are actually symbolic links to files in the directory `/etc/rc.d/init.d`. Renaming the files in the `init.d` directory disables all the symbolic links that point to that file. If you only wish

to disable a service for a particular run level, rename the appropriate symbolic link by replacing the S with a K, like this:

```
root# cd /etc/rc6.d
root# mv S45dhcpd K45dhcpd
```

Note: **Linux-Mandrake** uses only runlevels 0-6, so you should normally only need to worry about directories `/etc/rc.d/rc[0-6].d`.

Note: You may also use a command line utility to do that: `chk-config` or the graphical interface under *KDE*: `ksysv`.

Your **Linux-Mandrake** distributions ships with a `tcp_wrapper` “wrapping” all your TCP services. The `tcp_wrapper` (`tcpd`) is invoked from `inetd` instead of the real server. `tcpd` then checks the host that is requesting the service, and either executes the real server, or denies access from that host. `tcpd` allows you to restrict access to your TCP services. You should edit `/etc/hosts.allow` and add in only those hosts that need to have access to your machine’s services.

If you are a home dialup user, we suggest you deny ALL. `tcpd` also logs failed attempts to access services, so this can alert you if you are under attack. If you add new services, you should be sure to configure them to use `tcp_wrappers` if they are TCP-based. For example, a normal dial-up user can prevent outsiders from connecting to his machine, yet still have the ability to retrieve mail, and make network connections to the *Internet*. To do this, you might add the following to your `/etc/hosts.allow`:

ALL: 127.

And of course `/etc/hosts.deny` would contain:

ALL: ALL

which will prevent external connections to your machine, yet still allow you from the inside to connect to servers on the *Internet*.

Keep in mind that `tcp_wrappers` only protects services executed from `inetd`, and a select few others. There very well may be other services running on your machine. You can use `netstat -ta` to find a list of all the services your machine is offering.

7.8.3. Verify Your DNS Information

Keeping up-to-date DNS information about all hosts on your network can help to increase security. If an unauthorized host becomes connected to your network, you can recognize it by its lack of a DNS entry. Many services can be configured to not accept connections from hosts that do not have valid DNS entries.

7.8.4. `identd`

`identd` is a small program that typically runs out of your `inetd` server. It keeps track of what user is running what TCP service, and then reports this to whoever requests it.

Many people misunderstand the usefulness of `identd`, and so disable it or block all off site requests for it. `identd` is not there to help out remote sites. There is no way of knowing if the data you get from the remote `identd` is correct or not. There is no authentication in `identd` requests.

Why would you want to run it then? Because it helps **you** out, and is another data-point in tracking. If your `identd` is uncompromised, then you know it's telling remote sites the user-name or UID of people using TCP services. If the admin at a remote site comes back to you

and tells you user so-and-so was trying to hack into their site, you can easily take action against that user. If you are not running `identd`, you will have to look at lots and lots of logs, figure out who was on at the time, and in general take a lot more time to track down the user.

The `identd` that ships with most distributions is more configurable than many people think. You can disable it for specific users (they can make a `.noident` file), you can log all `identd` requests (We recommend it), you can even have `identd` return a UID instead of a user name or even `NO-USER`.

7.8.5. SATAN, ISS, and Other Network Scanners

There are a number of different software packages out there that do port and service-based scanning of machines or networks. *SATAN*, *ISS*, *SAINT*, and *Nessus* are some of the more well-known ones. This software connects to the target machine (or all the target machines on a network) on all the ports they can, and try to determine what service is running there. Based on this information, you can tell if the machine is vulnerable to a specific exploit on that server.

SATAN (Security Administrator's Tool for Analyzing Networks) is a port scanner with a web interface. It can be configured to do light, medium, or strong checks on a machine or a network of machines. It's a good idea to get *SATAN* and scan your machine or network, and fix the problems it finds. Make sure you get the copy of *SATAN* from metalab (<http://metalab.unc.edu/pub/packages/security/Satan-for-Linux/>) or a reputable FTP or web site. There was a Trojan copy of *SATAN* that was distributed out on the net. trouble.org (<http://www.trouble.org/~zen/satan/satan.html>). Note that *SATAN* has not been updated in quite a while, and some of the other tools below might do a better job.

ISS (Internet Security Scanner) is another port-based scanner. It is faster than *Satan*, and thus might be better for large networks. However, *SATAN* tends to provide more information.

Abacus is a suite of tools to provide host-based security and intrusion detection. look at its home page on the web for more information. <http://www.psionic.com/abacus/> (<http://www.psionic.com/abacus>)

SAINT is a updated version of *SATAN*. It is web based and has many more up to date tests than *SATAN*. You can find out more about it at: <http://www.wwdsi.com/~saint> (<http://www.wwdsi.com/saint>)

Nessus is a free security scanner. It has a GTK graphical interface for ease of use. It is also designed with a very nice plugin setup for new port-scanning tests. For more information, take a look at: <http://www.nessus.org> (<http://www.nessus.org/>)

7.8.5.1. Detecting Port Scans

There are some tools designed to alert you to probes by *SATAN* and *ISS* and other scanning software. However, if you liberally use *tcp_wrappers*, look over your log files regularly, you should be able to notice such probes. Even on the lowest setting, *SATAN* still leaves traces in the logs.

There are also “stealth” port scanners, including *nmap*. A packet with the TCP ACK bit set (as is done with established connections) will likely get through a non-stateful packet-filtering firewall. The returned RST packet from a port that **_had no established session_** can be taken as proof of life on that port. TCP wrappers will not detect this.

7.8.6. sendmail, qmail and MTA's¹

One of the most important services you can provide is a mail server.

1. Mail Transport Agents

Unfortunately, it is also one of the most vulnerable to attack, simply due to the number of tasks it must perform and the privileges it typically needs.

If you are using *sendmail* it is very important to keep up on current versions. *sendmail* has a long long history of security exploits. Always make sure you are running the most recent version from *sendmail* (<http://www.sendmail.org/>).

Keep in mind that *sendmail* does not have to be running in order for you to send mail. If you are a home user, you can disable *sendmail* entirely, and simply use your mail client to send mail. You might also choose to remove the *-bd* flag from the *sendmail* startup file, thereby disabling incoming requests for mail. In other words, you can execute *sendmail* from your startup script using the following instead:

```
# /usr/lib/sendmail -q15m
```

This will cause *sendmail* to flush the mail queue every fifteen minutes for any messages that could not be successfully delivered on the first attempt.

Many administrators choose not to use *sendmail*, and instead choose one of the other mail transport agents. You might consider switching over to *qmail*. *qmail* was designed with security in mind from the ground up. It's fast, stable, and secure. *Qmail* can be found at *qmail* (<http://www.qmail.org>)

In direct competition to *qmail* is *postfix*, written by Wietse Venema, the author of *tcp_wrappers* and other security tools. Formerly called *vmailer*, and sponsored by **IBM**, this is also a mail transport agent written from the ground up with security in mind. You can find more information about *postfix* at *postfix* (<http://www.postfix.org>)

Note: *postfix* is the default MTA shipped with **Linux-Mandrake**.

7.8.7. Denial of Service Attacks

A “Denial of Service” (DoS) attack is one where the attacker tries to make some resource too busy to answer legitimate requests, or to deny legitimate users access to your machine.

Denial of service attacks have increased greatly in recent years. Some of the more popular and recent ones are listed below. Note that new ones show up all the time, so this is just a few examples. Read the *GNU/Linux* security lists and the bugtraq list and archives for more current information.

Note: This manual describes more DoS attacks in a chapter introducing Firewalls and Proxy Servers.

- **SYN Flooding** - SYN flooding is a network denial of service attack. It takes advantage of a “loophole” in the way TCP connections are created. The newer *GNU/Linux* kernels (2.0.30 and up) have several configurable options to prevent SYN flood attacks from denying people access to your machine or services. See *Kernel Security*, page 120 for proper kernel protection options.
- **Ping Flooding** - Ping flooding is a simple brute-force denial of service attack. The attacker sends a “flood” of ICMP packets to your machine. If they are doing this from a host with better bandwidth than yours, your machine will be unable to send anything on the network. A variation on this attack, called “smurfing”, sends ICMP packets to a host with **your** machine’s return IP, allowing them to flood you less detectably. You can find more information about the “smurf” attack at <http://www.quadrunner.com/~chuegen/smurf.txt> (<http://www.quadrunner.com/~chuegen/smurf.txt>)

If you are ever under a ping flood attack, use a tool like `tcpdump` to determine where the packets are coming from (or appear to be com-

ing from), then contact your provider with this information. Ping floods can most easily be stopped at the router level or by using a firewall.

- **Ping o' Death** - The Ping o' Death attack sends ICMP ECHO REQUEST packets that are too large to fit in the kernel data structures intended to store them. Because sending a single, large (65,510 bytes) "ping" packet to many systems will cause them to hang or even crash, this problem was quickly dubbed the "Ping o' Death". This one has long been fixed, and is no longer anything to worry about.

You can find code for most exploits, and a more in-depth description of how they work, at <http://www.rootshell.com> (<http://www.rootshell.com>) using their search engine.

7.8.8. NFS (Network File System) Security.

NFS is a very widely-used file sharing protocol. It allows servers running `nfsd` and `mountd` to "export" entire filesystems to other machines using NFS filesystem support built in to their kernels (or some other client support if they are not *GNU/Linux* machines). `mountd` keeps track of mounted filesystems in `/etc/mtab`, and can display them with `showmount`.

Many sites use NFS to serve home directories to users, so that no matter what machine in the cluster they login to, they will have all their home files.

There is some small amount of security allowed in exporting filesystems. You can make your `nfsd` map the remote root user (UID=0) to the nobody user, denying them total access to the files exported. However, since individual users have access to their own (or at least the same UID) files, the remote root user can login or `su` to their ac-

count and have total access to their files. This is only a small hindrance to an attacker that has access to mount your remote filesystems.

If you must use NFS, make sure you export to only those machines that you really need to. Never export your entire root directory; export only directories you need to export.

See the NFS *HOWTO* for more information on NFS, available at LDP (<http://metalab.unc.edu/mdw/HOWTO/NFS-HOWTO.html>)

Note: NFS is a LAN-based service, designed only to work on local networks. Please do not, for any reason, allow it to operate to/from the Internet. Configure your firewall and hosts to block such access!

7.8.9. NIS (Network Information Service) (formerly YP).

Network Information service (formerly YP) is a means of distributing information to a group of machines. The NIS master holds the information tables and converts them into NIS map files. These maps are then served over the network, allowing NIS client machines to get login, password, home directory and *shell* information (all the information in a standard `/etc/passwd` file). This allows users to change their password once and have it take effect on all the machines in the NIS domain.

NIS is not at all secure. It was never meant to be. It was meant to be handy and useful. Anyone that can guess the name of your NIS domain (anywhere on the net) can get a copy of your `passwd` file, and use *crack* and *John the Ripper* against your users' passwords. Also, it is possible to spoof NIS and do all sorts of nasty tricks. If you must use NIS, make sure you are aware of the dangers.

Note: NIS is a LAN-based service, never intended for the Internet. Please do not allow this through your firewall to/from the Internet!

There is a much more secure replacement for NIS, called *NIS+*. Check out the NIS *HOWTO* for more information: NIS-HOWTO (<http://metalab.unc.edu/mdw/HOWTO/NIS-HOWTO.html>)

7.8.10. Firewalls

Firewalls are a means of controlling what information is allowed into and out of your local network. Typically the firewall host is connected to the *Internet* and your local LAN, and the only access from your LAN to the *Internet* is through the firewall. This way the firewall can control what passes back and forth from the *Internet* and your LAN.

There are a number of types of firewalls and methods of setting them up. *GNU/Linux* machines make pretty good firewalls. Firewall code can be built right into 2.0 and higher kernels. The user-space tools *ipchains* for 2.2 kernels, allows you to change, on the fly, the types of network traffic you allow. You can also log particular types of network traffic.

Firewalls are a very useful and important technique in securing your network. However, never think that because you have a firewall, you don't need to secure the machines behind it. This is a fatal mistake. Check out the very good Firewall-HOWTO at your latest metalab archive for more information on firewalls and *GNU/Linux*. Firewall-HOWTO (<http://metalab.unc.edu/mdw/HOWTO/Firewall-HOWTO.html>)

If you have no experience with firewalls, and plan to set up one for more than just a simple security policy, the Firewalls book by O'Reilly and Associates or other online firewall document is manda-

tory reading. Check out O'Reilly site (<http://www.ora.com>) for more information. The National Institute of Standards and Technology have put together an excellent document on firewalls. Although dated 1995, it is still quite good. You can find it at nist.gov (<http://csrc.nist.gov/nistpubs/800-10/main.html>). Also of interest includes:

- The Freefire Project -- a list of freely-available firewall tools, available at freefire (http://sites.inka.de/sites/lina/freefire-1/index_en.html)
- SunWorld Firewall Design -- written by the authors of the O'Reilly book, this provides a rough introduction to the different firewall types. It's available at sunworld.com (<http://www.sunworld.com/swol-01-1996/swol-01-firewall.html>)
- Mason -- the automated firewall builder for *GNU/Linux*. This is a firewall script that learns as you do the things you need to do on your network! More info at: mason (<http://www.pobox.com/~wstearns/mason/>)

7.8.11. IP Chains - GNU/Linux **Kernel 2.2.x** Firewalling

GNU/Linux IP Firewalling Chains is an update to the 2.0 *GNU/Linux* firewalling code for the 2.2 kernel. It has many more features than previous implementations, including:

- More flexible packet manipulations
- More complex accounting
- Simple policy changes possible automatically

- Fragments can be explicitly blocked, denied, etc.
- Logs suspicious packets.
- Can handle protocols other than ICMP/TCP/UDP.

Be sure to read the IP Chains *HOWTO* for further information. It is available at rustcorp.com (<http://www.rustcorp.com/linux/ipchains/HOWTO.html>)

7.8.12. VPNs - Virtual Private Networks

VPN's are a way to establish a “virtual” network on top of some already existing network. This virtual network often is encrypted and passes traffic only to and from some known entities that have joined the network. VPN's are often used to connect someone working at home over the public *Internet* to a internal company network.

If you are running a *GNU/Linux* masquerading firewall and need to pass MS PPTP (Microsoft's VPN point-to-point product) packets, there is a Linux kernel patch out to do just that. See: [ip-masq-vpn](ftp://ftp.rubyriver.com/pub/jhardin/masquerade/ip_masq_vpn.html) (ftp://ftp.rubyriver.com/pub/jhardin/masquerade/ip_masq_vpn.html).

There are several *GNU/Linux* VPN solutions available:

- *vpnd*. See the <http://sunsite.auc.dk/vpnd/> (<http://sunsite.auc.dk/vpnd/>).
- Free S/Wan, available at <http://www.xs4all.nl/~freeswan/> (<http://www.xs4all.nl/~freeswan/>)
- *ssh* can be used to construct a VPN. See the VPN mini-howto for more information.
- *vps* (virtual private server) at <http://www.strongcrypto.com> (<http://www.strongcrypto.com>).

See also the section on IPSEC for pointers and more information.

7.9. Security Preparation (before you go on-line)

OK, so you have checked over your system, and determined it's as secure as feasible, and you're ready to put it online. There are a few things you should now do in order to prepare for an intrusion, so you can quickly disable the intruder, and get back up and running.

7.9.1. Make a Full Backup of Your Machine

Discussion of backup methods and storage is beyond the scope of this chapter, but here are a few words relating to backups and security:

If you have less than 650MB of data to store on a partition, a CD-R copy of your data is a good way to go (as it's hard to tamper with later, and if stored properly can last a long time). Tapes and other rewritable media should be write-protected as soon as your backup is complete, and then verified to prevent tampering. Make sure you store your backups in a secure off-line area. A good backup will ensure that you have a known good point to restore your system from.

7.9.2. Choosing a Good Backup Schedule

A six-tape cycle is easy to maintain. This includes four tapes for during the week, one tape for even Fridays, and one tape for odd Fridays. Perform an incremental backup every day, and a full backup on the appropriate Friday tape. If you make some particularly important changes or add some important data to your system, a full backup might well be in order.

7.9.3. Backup Your RPM File Database

In the event of an intrusion, you can use your RPM database like you would use `tripwire`, but only if you can be sure it too hasn't been modified. You should copy the RPM database to a floppy, and keep this copy off-line at all times.

The files `/var/lib/rpm/fileindex.rpm` and `/var/lib/rpm/packages.rpm` most likely won't fit on a single floppy. But if compressed, each should fit on a separate floppy.

Now, when your system is compromised, you can use the command:

```
root# rpm -Va
```

to verify each file on the system. See the `rpm` man page, as there are a few other options that can be included to make it less verbose. Keep in mind you must also be sure your RPM binary has not been compromised.

This means that every time a new RPM is added to the system, the RPM database will need to be rearchived. You will have to decide the advantages versus drawbacks.

7.9.4. Keep Track of Your System Accounting Data

It is very important that the information that comes from `syslog` has not been compromised. Making the files in `/var/log` readable and writable by only a limited number of users is a good start.

Be sure to keep an eye on what gets written there, especially under the `auth` facility. Multiple login failures, for example, can indicate an attempted break-in.

You will want to look in `/var/log` and check `messages`, `mail.log`, and others.

You might also want to configure your log-rotating script to keep logs around longer so you have time to examine them. Take a look at the `logrotate` man page.

If your log files have been tampered with, see if you can determine when the tampering started, and what sort of things appeared to be tampered with. Are there large periods of time that cannot be accounted for? Checking backup tapes (if you have any) for untampered log files is a good idea.

Intruders typically modify log files in order to cover their tracks, but they should still be checked for strange happenings. You may notice the intruder attempting to gain entrance, or exploit a program in order to obtain the root account. You might see log entries before the intruder has time to modify them.

You should also be sure to separate the `auth` facility from other log data, including attempts to switch users using `su`, login attempts, and other user accounting information.

If possible, configure `syslog` to send a copy of the most important data to a secure system. This will prevent an intruder from covering his tracks by deleting his `login/su/ftp` etc attempts. See the `syslog.conf` man page, and refer to the `@` option.

There are several more advanced `syslogd` programs out there. Take a look at <http://www.core-sdi.com/ssyslog/> (<http://www.core-sdi.com/ssyslog/>) for Secure Syslog. Secure Syslog allows you to encrypt your `syslog` entries and make sure no one has tampered with them.

Another `syslogd` with more features is `syslog-ng` (<http://www.balabit.hu/products/syslog-ng.html>). It allows you a lot more flexibility in your logging and also can crypt your remote `syslog` streams to prevent tampering.

Finally, log files are much less useful when no one is reading them. Take some time out every once in a while to look over your log files, and get a feeling for what they look like on a normal day. Knowing

this can help make unusual things stand out.

7.9.5. Apply All New System Updates.

Due to the fast-paced nature of security fixes, new (fixed) programs are always being released. Before you connect your machine to the network, it's a good idea to run `MandrakeUpdate` (on another machine connected to the *Internet*) and get all the updated packages since you received your distribution CDRom. Many times these packages contain important security fixes, so it's a good idea to get them installed.

7.10. What To Do During and After a Break-in

So you have followed some of the advice here (or elsewhere) and have detected a break-in? The first thing to do is to remain calm. Hasty actions can cause more harm than the attacker would have.

7.10.1. Security Compromise Underway.

Spotting a security compromise under way can be a tense undertaking. How you react can have large consequences.

If the compromise you are seeing is a physical one, odds are you have spotted someone who has broken into your home, office or lab. You should notify your local authorities. In a lab, you might have spotted someone trying to open a case or reboot a machine. Depending on your authority and procedures, you might ask them to stop, or contact your local security people.

If you have detected a local user trying to compromise your security, the first thing to do is confirm they are in fact who you think they are.

Check the site they are logging in from. Is it the site they normally log in from? No? Then use a non-electronic means of getting in touch. For instance, call them on the phone or walk over to their office/house and talk to them. If they agree that they are on, you can ask them to explain what they were doing or tell them to cease doing it. If they are not on, and have no idea what you are talking about, odds are this incident requires further investigation. Look into such incidents, and have lots of information before making any accusations.

If you have detected a network compromise, the first thing to do (if you are able) is to disconnect your network. If they are connected via modem, unplug the modem cable; if they are connected via *Ethernet*, unplug the *Ethernet* cable. This will prevent them from doing any further damage, and they will probably see it as a network problem rather than detection.

If you are unable to disconnect the network (if you have a busy site, or you do not have physical control of your machines), the next best step is to use something like `tcp_wrappers` or `ipfwadm/ipchains` to deny access from the intruder's site.

If you can't deny all people from the same site as the intruder, locking the user's account will have to do. Note that locking an account is not an easy thing. You have to keep in mind `.rhosts` files, FTP access, and a host of possible backdoors.

After you have done one of the above (disconnected the network, denied access from their site, and/or disabled their account), you need to kill all their user processes and log them off.

You should monitor your site well for the next few minutes, as the attacker will try to get back in. Perhaps using a different account, and/or from a different network address.

7.10.2. Security Compromise has already

happened

So you have either detected a compromise that has already happened or you have detected it and locked (hopefully) the offending attacker out of your system. Now what?

7.10.2.1. Closing the Hole

If you are able to determine what means the attacker used to get into your system, you should try to close that hole. For instance, perhaps you see several FTP entries just before the user logged in. Disable the FTP service and check and see if there is an updated version, or if any of the lists know of a fix.

Check all your log files, and make a visit to your security lists and pages and see if there are any new common exploits you can fix. You can find your **Linux-Mandrake** security fixes by running the `MandrakeUpdate` regularly.

There is now a *GNU/Linux* security auditing project. They are methodically going through all the user-space utilities and looking for possible security exploits and overflows. From their announcement:

“We are attempting a systematic audit of *GNU/Linux* sources with a view to being as secure as *OpenBSD*. We have already uncovered (and fixed) some problems, but more help is welcome. The list is unmoderated and also a useful resource for general security discussions. The list address is: `security-audit@ferret.lmh.ox.ac.uk` To subscribe, send a mail to: `security-audit-subscribe@ferret.lmh.ox.ac.uk`”

If you don't lock the attacker out, they will likely be back. Not just back on your machine, but back somewhere on your network. If they were running a packet sniffer, odds are good they have access to other local machines.

7.10.2.2. Assessing the Damage

The first thing is to assess the damage. What has been compromised? If you are running an Integrity Checker like *Tripwire*, you can use it to perform an integrity check; and it should help to tell you what has been compromised. If not, you will have to look around at all your important data.

Since *GNU/Linux* systems are getting easier and easier to install, you might consider saving your config files, wiping your disk(s), re-installing, then restoring your user files and your config files from backups. This will ensure that you have a new, clean system. If you have to backup files from the compromised system, be especially cautious of any binaries that you restore, as they may be Trojan horses placed there by the intruder.

Re-installation should be considered mandatory upon an intruder obtaining root access. Additionally, you'd like to keep any evidence there is, so having a spare disk in the safe may make sense.

Then you have to worry about how long ago the compromise happened, and whether the backups hold any damaged work. More on backups later.

7.10.2.3. Backups, Backups, Backups!

Having regular backups is a godsend for security matters. If your system is compromised, you can restore the data you need from backups. Of course, some data is valuable to the attacker too, and they will not only destroy it, they will steal it and have their own copies; but at least you will still have the data.

You should check several backups back into the past before restoring a file that has been tampered with. The intruder could have compromised your files long ago, and you could have made many successful backups of the compromised file!

Of course, there are also a raft of security concerns with backups. Make sure you are storing them in a secure place. Know who has access to them. (If an attacker can get your backups, they can have access to all your data without you ever knowing it.)

7.10.2.4. Tracking Down the Intruder.

OK, you have locked the intruder out, and recovered your system, but you're not quite done yet. While it is unlikely that most intruders will ever be caught, you should report the attack.

You should report the attack to the admin contact at the site from which the attacker attacked your system. You can look up this contact with `whois` or the Internic database. You might send them an email with all applicable log entries and dates and times, though it would be much safer to make this contact by phone. Remember that the attacker might have compromised their e-mail systems. This illustrates what we call "out of band" communications – if you're not sure you can trust one communications medium (e-mail) consider using another! Even if you do use e-mail to make the contact, you should follow up with a phone call. If that admin in turn spots your attacker, they might be able to talk to the admin of the site where they are coming from and so on.

Good crackers often use many intermediate systems, some (or many) of which may not even know they have been compromised. Trying to track a cracker back to their home system can be difficult. Being polite to the admins you talk to can go a long way to getting help from them.

You should also notify any security organizations you are a part of (CERT (<http://www.cert.org/>) or similar), as well as **MandrakeSoft**: <http://www.linux-mandrake.com/en/security.php3> (<http://www.linux-mandrake.com/en/security.php3>)

7.11. Security Sources

There are a **lot** of good sites out there for *Unix* security in general and *GNU/Linux* security specifically. It's very important to subscribe to one (or more) of the security mailing lists and keep current on security fixes. Most of these lists are very low volume, and very informative.

7.11.1. FTP Sites

CERT is the Computer Emergency Response Team. They often send out alerts of current attacks and fixes. See <ftp://ftp.cert.org> (<ftp://ftp.cert.org>) for more information.

ZEDZ (formerly Replay) (<http://www.zedz.net> (<http://www.zedz.net>)) has archives of many security programs. Since they are outside the US, they don't need to obey US crypto restrictions.

Matt Blaze is the author of CFS and a great security advocate. Matt's archive is available at <ftp://ftp.research.att.com/pub/mab> (<ftp://ftp.research.att.com/pub/mab>)

tue.nl is a great security FTP site in the Netherlands. [tue.nl](ftp://ftp.win.tue.nl/pub/security/) (<ftp://ftp.win.tue.nl/pub/security/>)

7.11.2. web Sites

- BUGTRAQ puts out advisories on security issues: BUGTRAQ archives (<http://www.netSPACE.org/lsv-archive/bugtraq.html>)
- CERT, the Computer Emergency Response Team, puts out advisories on common attacks on *Unix* platforms: CERT home (<http://www.cert.org/>)
- The Bastille Linux project creates a security tightening program for Linux. MandrakeSoft has supported this project and is working to

include it in their distribution. The project page is at www.bastille-linux.org (<http://www.bastille-linux.org/>).

- The Hacker FAQ is a FAQ about hackers: The Hacker FAQ (<http://www.solon.com/~seebbs/faqs/hacker.html>)
- The COAST archive has a large number of *Unix* security programs and information: COAST (<http://www.cs.purdue.edu/coast/>)
- The CryptoArchive [cryptoarchive.com](http://www.cryptoarchive.com/) (<http://www.cryptoarchive.com/>) is another site to get encryption and security tools.
- SuSe Security Page: <http://www.suse.de/security/> (<http://www.suse.de/security/>)
- Rootshell.com is a great site for seeing what exploits are currently being used by crackers: <http://www.rootshell.com/> (<http://www.rootshell.com/>)
- Dan Farmer is the author of *SATAN* and many other security tools. His home site has some interesting security survey information, as well as security tools: <http://www.trouble.org> (<http://www.trouble.org>)
- The *GNU/Linux* security WWW is a good site for *GNU/Linux* security information: Linux Security WWW (<http://www.aoy.com/Linux/Security/>)
- Infilsec has a vulnerability engine that can tell you what vulnerabilities affect a specific platform: <http://www.infilsec.com/vulnerabilities/> (<http://www.infilsec.com/vulnerabilities/>)
- CIAC sends out periodic security bulletins on common exploits: <http://ciac.llnl.gov/cgi-bin/index/bulletins> (<http://ciac.llnl.gov/cgi-bin/index/bulletins>)
- A good starting point for *GNU/Linux* Pluggable Authentication modules can be found at kernel.org (<http://www.kernel.org/pub/linux/libs/pam/>).

- WWW Security FAQ, written by Lincoln Stein, is a great web security reference. Find it at <http://www.w3.org/Security/Faq/www-security-faq.html> (<http://www.w3.org/Security/Faq/www-security-faq.html>)

7.11.3. Mailing Lists

Linux-Mandrake security list: you can be informed for each security fix by subscribing to our security mailing-list: <http://www.linux-mandrake.com/en/security.php3> (<http://www.linux-mandrake.com/en/security.php3>)

Bugtraq: To subscribe to bugtraq, send mail to listserv@netspace.org containing the message body “subscribe bugtraq”. (see links above for archives).

CIAC: Send e-mail to majordomo@tholia.llnl.gov. In the BODY (not subject) of the message put: “subscribe ciac-bulletin”

7.11.4. Books – Printed Reading Material

There are a number of good security books out there. This section lists a few of them. In addition to the security specific books, security is covered in a number of other books on system administration.

References

D. Brent Chapman, Elizabeth D. Zwicky, *Building Internet Firewalls*, 1st Edition September 1995, ISBN 1-56592-124-0.

Simson Garfinkel, Gene Spafford, *Practical UNIX & Internet Security*, 2nd Edition April 1996, ISBN 1-56592-148-8.

Deborah Russell, G.T. Gangemi, Sr., *Computer Security Basics*, 1st Edition July 1991, ISBN 0-937175-71-4.

Olaf Kirch, *Linux Network Administrator's Guide*, 1st Edition January 1995, ISBN 1-56592-087-2.

Simson Garfinkel, *PGP: Pretty Good Privacy*, 1st Edition December 1994, ISBN 1-56592-098-8.

David Icove, Karl Seger, William VonStorch, *Computer Crime A Crimefighter's Handbook*, 1st Edition August 1995, ISBN 1-56592-086-4.

John S. Flowers, *Linux Security*, New Riders, March 1999, ISBN 0735700354.

Anonymous, *Maximum Linux Security : A Hacker's Guide to Protecting Your Linux Server and Network*, July 1999, ISBN 0672313413.

Terry Escamilla, *Intrusion Detection*, John Wiley and Sons, September 1998, ISBN 0471290009.

Donn Parker, *Fighting Computer Crime*, John Wiley and Sons, September 1998, ISBN 0471163783.

Security-related terms

authentication

The process of knowing that the data received is the same as the data that was sent, and that the claimed sender is in fact the actual sender.

bastion Host

A computer system that must be highly secured because it is vulnerable to attack, usually because it is exposed to the *Internet* and is a main point of contact for users of internal networks. It gets its name from the highly fortified projects on the outer walls of medieval castles. Bastions overlook critical areas of defense, usually having strong walls, room for extra troops, and the occasional useful tub of boiling hot oil for discouraging attackers. Some reasonable definition here.

buffer overflow

Common coding style is to never allocate large enough buffers, and to not check for overflows. When such buffers overflow, the executing program (daemon or set-uid program) can be tricked in doing some other things. Generally this works by overwriting a function's return address on the stack to point to another location.

denial of service

An attack that consumes the resources on your computer for things it was not intended to be doing, thus preventing normal use of your network resources for legitimate purposes.

dual-homed Host

A general-purpose computer system that has at least two network interfaces.

firewall

A component or set of components that restricts access between a protected network and the *Internet*, or between other sets of networks.

host

A computer system attached to a network.

IP spoofing

IP Spoofing is a complex technical attack that is made up of several components. It is a security exploit that works by tricking computers in a trust relationship into thinking that you are someone that you really aren't. There is an extensive paper written by

daemon9, route, and infinity in the Volume Seven, Issue Forty-Eight of Phrack Magazine.

non-repudiation

The property of a receiver being able to prove that the sender of some data did in fact send the data even though the sender might later deny ever having sent it.

packet

The fundamental unit of communication on the *Internet*.

packet filtering

The action a device takes to selectively control the flow of data to and from a network. Packet filters allow or block packets, usually while routing them from one network to another (most often from the *Internet* to an internal network, and vice-versa). To accomplish packet filtering, you set up rules that specify what types of packets (those to or from a particular IP address or port) are to be allowed and what types are to be blocked.

perimeter network

A network added between a protected network and an external network, in order to provide an additional layer of security. A perimeter network is sometimes called a DMZ.

proxy server

A program that deals with external servers on behalf of internal clients. Proxy clients talk to proxy servers, which relay approved client requests to real servers, and relay answers back to clients.

superuser

An informal name for root.

7.12. Frequently Asked Questions

Q: Is it more secure to compile driver support directly into the kernel, instead of making it a module?

A: Some people think it is better to disable the ability to load device drivers using modules, because an intruder could load a Trojan module or a module that could affect system security.

However, in order to load modules, you must be root. The module object files are also only writable by root. This means the intruder would need root access to insert a module. If the intruder gains root access, there are more serious things to worry about than whether he will load a module.

Modules are for dynamically loading support for a particular device that may be infrequently used. On server machines, or firewalls for instance, this is very unlikely to happen. For this reason, it would make more sense to compile support directly into the kernel for machines acting as a server. Modules are also slower than support compiled directly in the kernel.

Q: Why does logging in as root from a remote machine always fail?

A: See root Security, page 96. This is done intentionally to prevent remote users from attempting to connect via telnet to your machine as root, which is a serious security vulnerability, because then the root password would be transmitted, in cleartext, across the network. Don't forget: potential intruders have time on their side, and can run automated programs to find your password.

Q: How can I enable the Apache SSL extensions?

A:

1. Get `mod_ssl 0.8.0` from one of the Linux-Mandrake crypto mirrors (<http://www.linux-mandrake.com/en/fcrypto.php3>).
2. Install it! And consult the documentation at `mod_ssl` homepage (www.modssl.org).

Note: You should also consider the module `mod_ssl-sxnet` which is a plug-in for `mod_ssl` and allow the activation of the ■Thawte Secure Extranet■. `mod_ssl` encrypt communications, but `mod_ssl-sxnet` goes further and allows to securely authenticate the user of the web page thanks to a personal certificate. You can read the doc for this package (http://medcom.netrevolution.com/addon-modules/mod_ssl-sxnet/sxnet.html)

You might also try ZEDZ net (<http://www.zedz.net>) which has many pre-built packages, and is located outside of the United States.

Q: How can I manipulate user accounts, and still retain security?

A: Your Linux-Mandrake distribution, contains a great number of tools to change the properties of user accounts.

- The `pwconv` and `unpwconv` programs can be used to convert between shadow and non-shadowed passwords.
- The `pwck` and `grpck` programs can be used to verify proper organization of the `/etc/passwd` and `/etc/group` files.

- The `useradd`, `usermod`, and `userdel` programs can be used to add, delete and modify user accounts. The `groupadd`, `groupmod`, and `groupdel` programs will do the same for groups.
- Group passwords can be created using `gpasswd`.

All these programs are “shadow-aware” – that is, if you enable shadow they will use `/etc/shadow` for password information, otherwise they won’t.

Q: How can I password-protect specific HTML documents using Apache?

A: I bet you didn’t know about apacheweek (<http://www.apacheweek.com>), **did you?**

You can find information on user authentication at apacheweek (<http://www.apacheweek.com/features/userauth>) as well as other web server security tips from Apache (http://www.apache.org/docs/misc/security_tips.html)

7.13. Conclusion

By subscribing to the security alert mailing lists, and keeping current, you can do a lot towards securing your machine. If you pay attention to your log files and run something like tripwire regularly, you can do even more.

A reasonable level of computer security is not difficult to maintain on a home machine. More effort is required on business machines, but *GNU/Linux can indeed be a secure platform. Due to*

the nature of GNU/Linux development, security fixes often come out much faster than they do on commercial operating systems, making GNU/Linux an ideal platform when security is a requirement.

Chapter 8. Firewall and Proxy Server

8.1. Introduction

This chapter is designed to teach the basics of firewall systems and give you details about the underlying theory and technology on a *GNU/Linux*-based machine.

8.2. What's the Purpose Behind a Firewall?

In this section, we will discuss the main attacks that a connected computer is likely to suffer from. It is crucial for a system administrator to have good understanding of common attacks so that he/she can prevent, possibly detect and better combat them. Note that some attacks are not necessarily defended by a simple firewall without a few other mechanisms. So, let's look at those attacks.

8.2.1. Unauthorized Access

Unauthorized access is simply when someone accesses private resources for which they haven't been given explicit permission to do so. This points to a central idea in firewall design: you need to explicitly deny access to any resource that you don't want people to access. You can't just assume that no one will find your site or try to gain access: there's too many curious people out there for that to be true!

Examples of unauthorized access are: someone browsing your internal web pages or mounting your NFS volumes.

8.2.2. Eavesdropping

This is another way of accessing unauthorized information, not by actively interacting with the source, but simply by capturing the information while moving from the source to an authorized target. It is the electronic version of wiretapping.

There are many publicly available programs that listen to and capture all traffic on a specific network - some are designed to extract relevant information, depending on the needs of the cracker: logins, passwords, or even full files or HTML pages.

Before you become incredibly paranoid about this, let us clarify how this is done. When you use a non-encrypted protocol, anyone on the local network can usually eavesdrop on the content. Additionally, many of the other machines on the way between the two communicating computers can also see this traffic. The following protocols are non-encrypted and thus expose passwords: telnet, ftp, pop, imap, and possibly HTTP (web). They can be replaced by ssh, scp, apop, imap-ssl and https (secure web) respectively. Each of these use encryption, so that an eavesdropper just sees “garbage” rather than important data.

As you can see, the firewall cannot deflect this attack alone. There are only two things that you, as a firewall administrator, can do about this. First, you can educate yourself and your users about these protocols. Second, you can completely block dangerous protocols. Block telnet and tell your users about ssh. Block imap and show your users imap-ssl.

8.2.3. Denial of Service (DoS)

A cracker uses a Denial of Service attack to do just what its name states: deny some services. He might block your network connection by throwing all kinds of network garbage at it. A Denial of Service attack doesn't even have to be network-based. All an attacker has to do is starve one of your resources: memory, hard disk space, network connections, etc.

This attack differs from the others in that no confidential information is stolen, though some services are rendered inaccessible. This can still be very damaging. Many companies count minutes of web server downtime in dollars lost!

There are various Denial of Service (DoS) attack methods:

- Flood the target with many malformed packets to a particular port/service, making it unavailable to handle legitimate requests.
- Send particularly well-designed packets to an application on the target, resulting in the said application's crash. Often, this crash will disable the entire host, by consuming memory or other resources. Generally, this works because of an oversight or bug exists in the application. These packets are designed to take advantage of this oversight or bug.

To be ready to block DoS attacks, it is important that you keep up to date with the latest exploits and DoS attack types, so you can patch holes or possibly add rules to your firewall and/or Intrusion Detection System (IDS).

8.2.4. Correcting for Known Weaknesses in Programs

Most vulnerabilities are discovered at some point as bugs in various programs. Usually, the “good guys” find out about it at the same time as the “bad guys”, but the bad guys are more vigilant in using this information. The best thing you can do for your own site security is to watch BugTraq and other mailing lists, to search for vulnerabilities in your software. Patch as often as possible and understand what the vulnerabilities mean. The latter action will often let you set firewall rules or OS parameters to avoid attacks against vulnerable subsystems while the OS vendor develops a patch.

8.2.5. Spoofing

This type of attack generally exploits an inherent characteristic of TCP/IP connections: the hosts can't actually "authenticate" each other. That is, when you get a packet from host A, you can't really be sure it comes from host A. See, it might be an attacker making his computer pretend to be host A. This is partly because TCP/IP was never really designed to operate in the hostile environment that the Internet has become. It was originally conceived for a select group among military, government and research institutions.

Now, there are actually different varieties of spoofing. These include: DNS, IP and address spoofing, among others.

8.2.6. To Learn More...

If you want to learn more about network security, consider the following papers/sites:

A rather old but topical article about TCP/IP security problems (http://www.ja.net/CERT/Bellovin/TCP-IP_Security_Problems.html).

A very large glossary of security-related terms (<http://www.garlic.com/~lynn/secure.htm>).

A trio of strong security sites: SecurityFocus (<http://www.securityfocus.com>), SecurityPortal (<http://www.securityportal.com>), and LinuxSecurity (<http://www.linuxsecurity.com>).

Finally, check out MandrakeCampus (<http://www.mandrakecampus.com>) and MandrakeExpert (<http://www.mandrakeexpert.com>) to learn more about any Linux-related topic.

check: <http://disc.cba.uh.edu/~rhirsch/fall96/barba.htm>

8.3. How Do Computers Communicate?

It's important to get a basic understanding of some TCP/IP networking concepts. Firewalling is normally a difficult task. While the MandrakeSecurity Firewall removes much of this difficulty, you still need to get some background here to be fully effective. Let's start with an explanation of how two Internet-connected computers communicate.

8.3.1. IP Addresses

Before you communicate with someone, you need to know their phone number or postal address, etc. In essence, you need special data which tells the intervening network, whether it's phone or postal, how to route your information (voice or pages). On the Internet, each computer has one or more Internet, or IP, addresses. You may have seen these before: they are made up of four numbers, each from 0 to 255, which look like this: *192.168.2.45*.

There is already a rather complete explanation on IP addresses at *IP Addresses, an Explanation.*, page 183>. So we'll concentrate here on *netmasks*. A netmask is used to tell each machine on a given network what range of IP addresses belong to that network.

To correctly understand netmasks, it is necessary to think about the binary representation of IP addresses. In fact each of the four numbers is the decimal notation of eight bits. Our previous example *192.168.2.45* is the decimal representation of the IP address 11000000.10101000.00000010.00101101

To designate a range of IP addresses belonging to a given network, you'd just say how many of those bits belong to the given network. For example, to say that the network is made up of all the IP addresses

beginning with 192.168.2., you'd note that this just means the first 8 bits for the first number (192), the second 8 bits for the second number (168) and the third 8 bits for the final number (2). So, you'd want to express that 24 bits make up the network number. We'd do this like this: 11111111.11111111.11111111.00000000. This is fairly simple, until you want to start talking about slo Then, to designate a range of addresses, it suffices to specify an address and how many of the first bits of this address are common to all the addresses of the desired range.

Example

Let's say that we want to designate all the addresses from 192.168.1.128 to 192.168.1.255, in binary, 11000000.10101000.00000010.10000000 to 11000000.10101000.00000010.11111111. This implies that the first three segments remain unchanged, plus the first bit of the fourth. In total, these are $8+8+8+1=25$ bits, thus remaining 7 bits to change, which represents 128 possible combinations. So our range is finally represented by the network number 192.168.1.128/25.

Note: This system obviously limits the possible ranges of addresses, for the reasons that the range is defined by a single number of bits, counting from the left side. Therefore you cannot directly specify with a single netmask the addresses between 192.168.1.64 to 192.168.1.255 as you cannot find the common bits between all of these addresses without also including addresses from 192.168.1.0.

8.3.2. Packets, Protocols and Layers

Now that we know how to identify a computer, we can send it a message. Due to the fact that the "pipes" through which data travels can

only accept a small amount of data at a time, messages are separated into smaller chunks called *packets*.

A packet is always composed of:

- A **header** - Similar to the envelope of a letter, it contains the address of the sender, the address of the receiver, and other protocol-specific information.
- Data - The sheet of paper that's placed in our "envelope".

Each packet passes through different stages to go from one application to another. Indeed, it is important to keep in mind that messages are exchanged by applications and not simply by computers. Here are the different stages (called *layers*) through which a message passes when traveling:

1. Application layer: the end-user application such as a web browser. This is basically the "data portion" as it's specific to each application and contains the real payload.
2. Transport layer: this is the first header added to the message. This is the TCP/UDP protocol part of the equation. Remember, the protocol stack that we're describing is applied via "encapsulation". The OS takes the data and wraps it, by prepending a header for the given protocol. So, at this point, we've got some data with a TCP header prepended.
3. Internet layer: this layer helps get the packet routed through the Internet. A TCP/IP packet would then have two headers – one for the Internet protocol and one for the transport protocol.
4. Physical layer: the physical means cables and their inherent characteristics used to actually transmit the data. This is often Ethernet

at your site. In the case of Ethernet, the IP packet is encapsulated with an Ethernet header to produce an Ethernet “frame”.

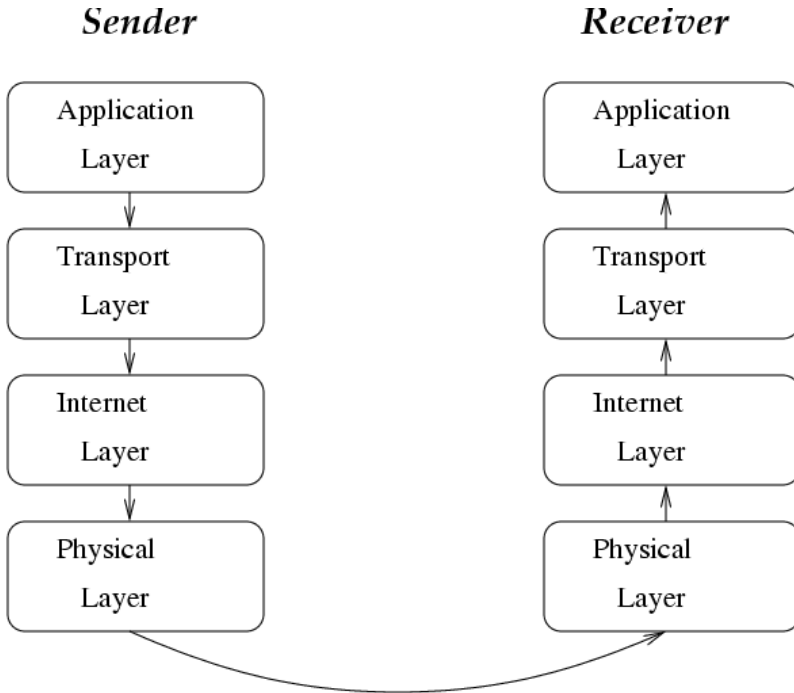


Figure 8-1. The TCP/IP Protocol Stack

Let’s talk more about this encapsulation concept. Basically, an application has some data it needs to get to an application on another machine. It transfers that data via an interface which adds a TCP header. The TCP header information is mostly there to get the data from the right program here to the right program on the other side.

Each layer is managed by one or various protocols. The application layer contains the protocols as in `/etc/services`: pop, telnet,

www, etc. The transport layer uses the protocols known as TCP, UDP, ICMP, etc.

Note: The transport and *Internet* layers are both handled by the *GNU/Linux* kernel and modules and networking low-level applications.

8.3.3. More About Headers

On the sender's side, each layer encapsulates the message it receives from the previous layer in a new header, so that the corresponding layer at the other side can understand the header. Then on the receiver's side, each layer removes the corresponding header before transmitting the packet to the next layer (figure 8-2).

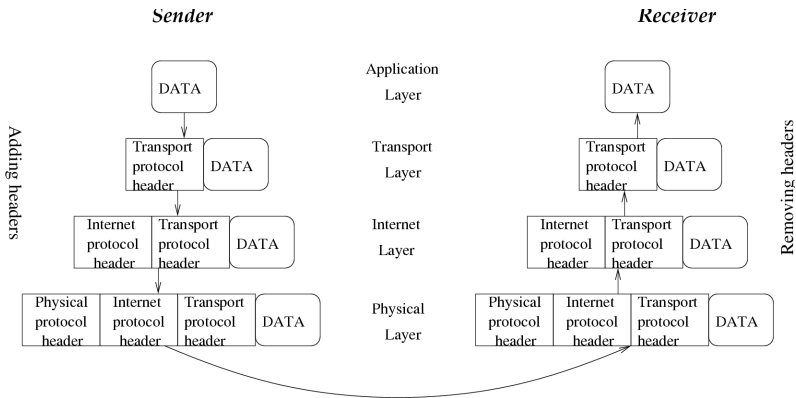


Figure 8-2. The Packet's Encapsulation Process

We won't detail the headers for all existing protocols – we will simply describe two headers generally used when doing TCP/IP firewalling: IP and UDP headers.

8.3.3.1. UDP Header

UDP is a very simple protocol which does not allow a sender to know whether its packets actually reach the target host or not. Thus the header is quite simple (figure 8-3).

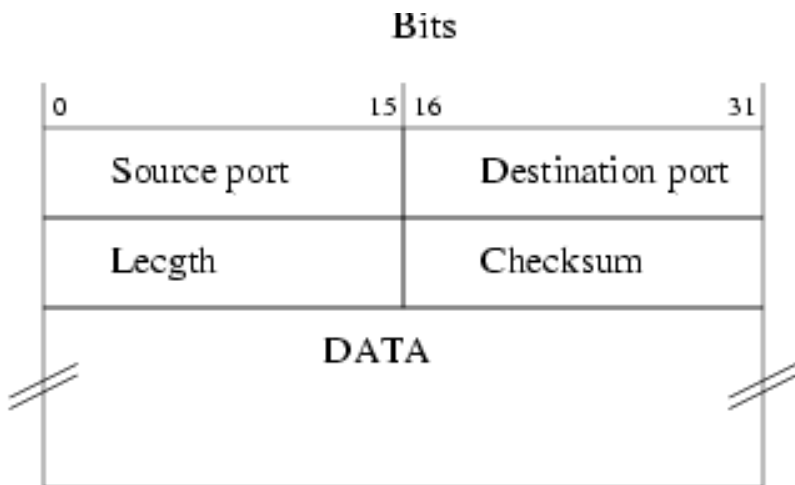


Figure 8-3. The UDP Protocol Header

The TCP protocol, on the contrary, implements a true client-server relation started with a reciprocal acknowledgement called “handshaking”. It allows for correction of errors, retransmission of lost packets and even helps the receiving host to reorder packets that arrive out of order.

8.3.3.2. IP Header

Firewalling is mainly a matter of filtering, and the more information the header contains, the more effective the filtering is.

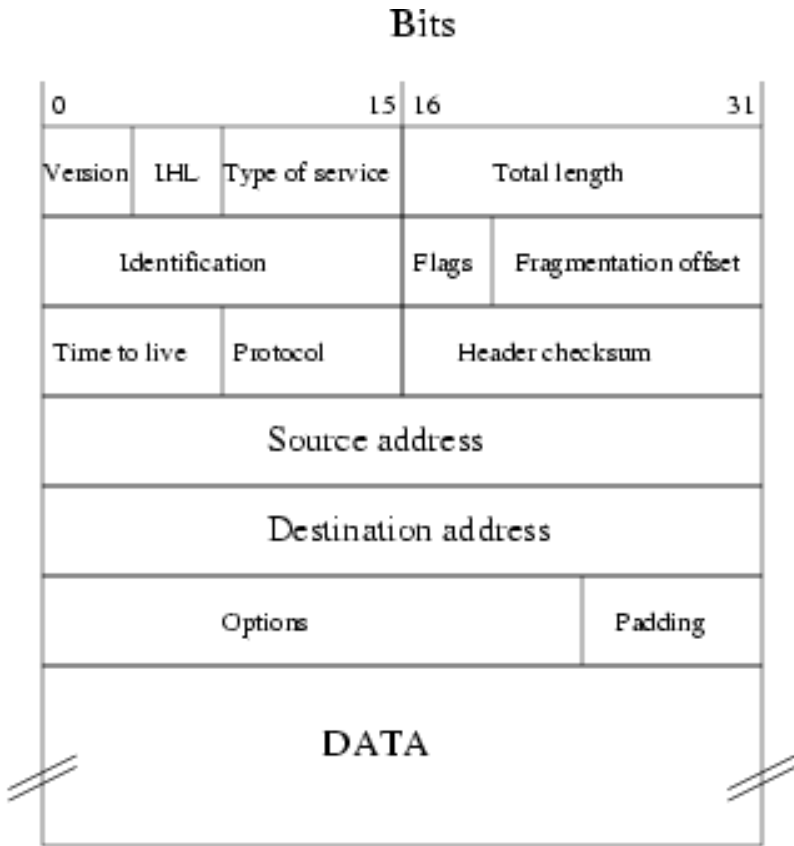


Figure 8-4. The IP Protocol Header

8.4. The Firewall: What Is It?

A firewall generally designates various things: a machine, a program, and/or a set of rules. In this section, we will focus on the first sense: what is the role of that machine and its location/connection with re-

spect to the environment.

Traditionally, a firewall is a wall separating two areas, in a building, a car, etc., to prevent fire from propagating from one area to another. By extension, it is used to separate two networks, to prevent hostile packets from one network from reaching the other. The most common firewall configuration protects a company's private network from the Internet. Firewalling traditionally operates by inspecting packet headers and discarding packets with undesirable header info.

Most simply, a firewall is placed between the wild *Internet* and a Local Area Network (LAN); this is illustrated in figure 8-5. Often, the firewall host acts as a proxy and masquerading host. In fact, a single host with a simple connection to the *Internet* and no other connection can also include a firewall.




Figure 8-5. The Most Simple LAN Firewall Configuration

To increase the complexity just slightly, you can add a DMZ (DeMilitarized Zone: figure 8-6), where you can put public servers with public information, offering services too unsafe to offer through the firewall.




Figure 8-6. A Firewall Between a LAN and a DMZ

If you wish to filter access to the DMZ, you can then add a third network interface to the machine hosting the firewall (figure 8-7). In this case, the firewall will have to manage six flows of packets (three ingoing, three outgoing).




Figure 8-7. A Firewall with Three Network Interfaces

Finally, a firewall may become a very complex machine, with functions such as:

- Network Address Translation: with private addresses (also called masquerading) or public addresses.
- Wide Area Network (WAN) connection: to provide a privileged connection with a friendly remote local network, such as local offices. This can include VPN (Virtual Private Network) services with encryption.
- Load balancing: for highly-loaded services (if your name is Yahoo! for example), to distribute requests over various similar servers.

Your firewall system's configuration can easily get out of hand, in terms of configuration maintenance.




Figure 8-8. A Complex Firewall Configuration

8.5. How Does the Firewall Work?

We will now go deeper into the firewall internals. We'll look at IP

filtering, including at what stage of the routing filtering is actually done and what criterion are used.

To understand how filtering is done, we will now look inside the kernel and examine the paths taken by packets. We'll examine the stage(s) where filtering is actually done.

8.6. Before Implementing a Firewall

It is important you sit down for a few moments and think about your firewall before charging ahead. We present quick guidelines here for designing a firewall appropriate to your specific situation, needs and environment.

8.7. Nitty Gritty Firewalling Process

These protocols provide extra capabilities and establish the possibility for a connection between two programs rather than just two computers. UDP is used to send short messages, with no real guarantees, while TCP maintains an entire connection, complete with error correction, packet re-ordering and missed/corrupted packet re-send request. But how does the computer use these protocols to facilitate communications between pairs of programs, rather than pairs of computers?

Well, TCP and UDP both accomplish this with a port number. This 16-bit number identifies the sending and receiving programs on each machine. The operating system allows each program to check out, or "bind to", one or more of these port numbers and then keeps track of which program has bound to which port, to make sure that the right

data gets to the right program. Here are some basic facts about packet filtering.

Each TCP or UDP connection is uniquely defined by the following four numbers:

- 1. Source IP Address (source computer)
- 2. Source Port (the program on a said source computer)
- 3. Destination IP Address (destination computer)
- 4. Destination Port (the program on a said destination computer)

Now, certain kinds of connections use “statically allocated” ports. For example, the telnet server basically binds to port 23 on its host machine. This means that a telnet session always looks like this: a TCP connection between two computers, from some port on the client (usually in the range 1024-65535) to port 23 on the server. Because of this fact, we can block all outgoing client connections to external telnet servers by simply blocking all outgoing TCP packets with destination port 23. We can block all incoming client connections to our internal telnet servers in a similar manner – block all incoming TCP packets with destination port 23.

Here are some other common server-side ports:

- FTP - TCP ports 20 and 21
- SSH - TCP port 22
- SMTP (e-mail between computers) - TCP port 25
- DNS - TCP and UDP ports 53
- HTTP (web) - TCP port 80
- POP (remote e-mail retrieval) - TCP port 110
- IMAP (remote e-mail retrieval) - TCP port 143

- HTTPS (encrypted web) - TCP port 443
- EXEC (rexec) - TCP port 512
- LOGIN (rlogin) - TCP port 513
- SHELL (rsh) - TCP port 514
- IRC (Internet Relay Chat) - TCP port 6667

Now, what do we use that long (yet far from comprehensive) port list for? Well, to block people on one side of the firewall from making a particular kind of connection to a machine on the other side of the firewall. We've seen how to block telnet connections and we can use the above port list to block others. There are some exceptions, in that the protocols don't follow this traditional "client initiates connection with server, some high port on the client to some fixed low port on the server" model. FTP is the most notable exception.

8.8. After Implementing a Firewall

A firewall is an incredibly dynamic system. Even if you're very careful and experienced, you'll still find yourself tuning and modifying your firewall over time to meet the needs of both your clients and a stronger security stance. Furthermore, you'll need to respond to incidents, possibly detected by the firewall, by implementing blocking rules via the firewall. It's important that you understand that this is definitely not a low-maintenance endeavor for most businesses.

Chapter 9. Networking Overview

9.1. Copyright

This chapter is based on a *HOWTO* by Joshua D. Drake {POET} which original is hosted by the LinuxPorts.com/ (<http://www.linuxports.com/>).

The NET-3/4-HOWTO, NET-3, and Networking-HOWTO, information on how to install and configure networking support for Linux. Copyright (c) 1997 Terry Dawson, 1998 Alessandro Rubini, 1999 Joshua D. Drake {POET} - LinuxPorts.com/ (<http://www.linuxports.com/>) is a FREE document. You may redistribute it under the terms of the GNU General Public License.

Modifications from v1.6.9, July 03, 2000, (C)opyright 2000,2001 MandrakeSoft

9.2. How to use this Chapter.

This document is organized top-down. The first sections include informative material and can be skipped if you are not interested; what follows is a generic discussion of networking issues, and you must ensure you understand this before proceeding to more specific parts. The rest, “technology specific” information is grouped in three main sections: Ethernet and IP-related information, technologies pertaining to widespread PC hardware and seldom-used technologies.

The suggested path through the document is thus the following:

Read the generic sections

These sections apply to every, or nearly every, technology described later and so are very important for you to understand. On

the other hand, I expect many of the readers to be already confident with this material.

Consider your network

You should know how your network is, or will be, designed and exactly what hardware and technology types you will be implementing.

Read the *Ethernet Information*, page 191 section if you are directly connected a LAN or the Internet:

This section describes basic Ethernet configuration and the various features that Linux offers for IP networks, like firewalling, advanced routing and so on.

Read the next section if you are interested in low-cost local networks or dial-up connections

The section describes PLIP, PPP, SLIP and ISDN, the widespread technologies used on personal workstations.

Read the technology specific sections related to your requirements:

If your needs differ from IP and/or common hardware, the final section covers details specific to non- IP protocols and peculiar communication hardware.

Do the configuration work

You should actually try to configure your network and take careful note of any problems you have.

Look for further help if needed

If you experience problems that this document does not help you to resolve then read the section related to where to get help or where to report bugs.

Have fun!

Networking is fun, enjoy it.

9.2.1. Conventions used in this document

No special convention is used here, but you must be warned about the way commands are shown. Following the classic Unix documentation, any command you should type to your shell is prefixed by a prompt. This howto shows “user%” as the prompt for commands that do not require superuser privileges, and “root#” as the prompt for commands that need to run as root. I chose to use “root#” instead of a plain “#” to prevent confusion with snapshots from shell scripts, where the hash mark is used to define comment lines.

When “Kernel Compile Options” are shown, they are represented in the format used by **menuconfig**. They should be understandable even if you, (like me), are not used to **menuconfig**. If you are in doubt about the options’ nesting, running the program once can’t do anything, but help.

9.3. General Information about Linux Networking.

9.3.1. Linux Networking Resources.

There are a number of places where you can find good information about Linux networking.

There are a wealth of Consultants available. A searchable listing can be found at <http://www.linuxports.com/> (<http://www.linuxports.com/>).

Alan Cox, the current maintainer of the Linux kernel networking code maintains a world wide web page that contains highlights of current and new developments in linux Networking at: www.uk.linux.org (<http://www.uk.linux.org/NetNews.html>).

There is a newsgroup in the Linux news hierarchy dedicated to networking and related matters, it is: `comp.os.linux.networking` (`news:comp.os.linux.networking`)

There is a mailing list to which you can subscribe where you may ask questions relating to Linux networking. To subscribe you should send a mail message:

```
To: majordomo@vger.rutgers.edu
Subject: anything at all
Message:
subscribe linux-net
```

Please remember when reporting any problem to include as much relevant detail about the problem as you can. Specifically you should identify the versions of software that you are using, especially the kernel version, the version of tools such as `pppd/` or `dip` and the exact nature of the problem you are experiencing. This means taking note of the exact syntax of any error messages you receive and of any commands that you are issuing.

9.3.2. Where to get some non-linux-specific network information.

If you are after some basic tutorial information on `tcp/ip` networking generally, then I recommend you take a look at the following documents:

tcp/ip introduction

This document comes as both a text version (`ftp://athos.rutgers.edu/runet/tcp-ip-intro.doc`) and a postscript version (`ftp://athos.rutgers.edu/runet/tcp-ip-intro.ps`).

tcp/ip administration

This document comes as both a text version (`ftp://athos.rutgers.edu/runet/tcp-ip-admin.doc`) and a postscript version (`ftp://athos.rutgers.edu/runet/tcp-ip-admin.ps`).

If you are after some more detailed information on tcp/ip networking then I highly recommend:

“Internetworking with TCP/IP, Volume 1: principles, protocols and architecture, by Douglas E. Comer, ISBN 0-13-227836-7, Prentice Hall publications, Third Edition, 1995.”

If you are wanting to learn about how to write network applications in a Unix compatible environment then I also highly recommend:

“Unix Network Programming, by W. Richard Stevens, ISBN 0-13-949876-1, Prentice Hall publications, 1990.”

A second edition of this book is appearing on the bookshelves; the new book is made up of three volumes: check Prentice-Hall’s web site (`http://www.phptr.com/`) for more information.

You might also try the `comp.protocols.tcp-ip` (`news:comp.protocols.tcp-ip`) newsgroup.

An important source of specific technical information relating to the Internet and the tcp/ip suite of protocols are RFC’s. RFC is an acronym for ‘Request For Comment’ and is the standard means of submitting and documenting Internet protocol standards. There are many RFC repositories. Many of these sites are ftp sites and other provide World Wide Web access with an associated search engine that allows you to search the RFC database for particular keywords.

One possible source for RFC's is at Nexor RFC database (<http://pubweb.nexor.co.uk/public/rfc/index/rfc.html>).

9.4. Generic Network Configuration Information.

The following subsections you will pretty much need to know and understand before you actually try to configure your network. They are fundamental principles that apply regardless of the exact nature of the network you wish to deploy.

9.4.1. What do I need to start ?

Before you start building or configuring your network you will need some things. The most important of these are:

9.4.1.1. Current Kernel source(Optional).

Note: Your **Linux-Mandrake** distribution comes with networking enabled, therefore it may not be required to recompile the kernel. If you are running well known hardware you should be just fine. For example: 3COM NIC, NE2000 NIC, or a Intel NIC. However if you find yourself in the position that you do need to update the kernel, the following information is provided.

Because the kernel you are running now might not yet have support for the network types or cards that you wish to use you will probably need the kernel source so that you can recompile the kernel with the appropriate options.

However, as long as you stay within the mainstream of hardware there should be no need to recompile your kernel unless there is a very specific feature that you need.

You can always obtain the latest kernel source from `ftp.cdrom.com` (`ftp://ftp.cdrom.com/pub/linux/sunsite/kernel.org/pub/linux/kernel`). This is not the official site but they have LOTS of bandwidth and capacity. The official site is `kernel.org` but please use the above if you can. Please remember that `ftp.kernel.org` is seriously overloaded. Use a mirror.

Normally the kernel source will be untarred into the `/usr/src/linux` directory. For information on how to apply patches and build the kernel you should read the Kernel-HOWTO (`Kernel-HOWTO.html`). For information on how to configure kernel modules you should read the “Modules mini-HOWTO”. Also, the `README` file found in the kernel sources and the `Documentation` directory are very informative for the brave reader.

Unless specifically stated otherwise, I recommend you stick with the standard kernel release (the one with the even number as the second digit in the version number). Development release kernels (the ones with the odd second digit) may have structural or other changes that may cause problems working with the other software on your system. If you are uncertain that you could resolve those sorts of problems in addition to the potential for there being other software errors, then don't use them.

9.4.1.2. IP Addresses, an Explanation.

Internet Protocol Addresses are composed of four bytes. The convention is to write addresses in what is called ‘dotted decimal notation’. In this form each byte is converted to a decimal number, (0-255), dropping any leading zero's unless the number is zero and written with each byte separated by a ‘.’ character. By convention each interface of a host or router has an IP address. It is legal for the same IP address to

be used on each interface of a single machine in some circumstances but usually each interface will have its own address.

Internet Protocol Networks are contiguous sequences of IP addresses. All addresses within a network have a number of digits within the address in common. The portion of the address that is common amongst all addresses within the network is called the ‘network portion’ of the address. The remaining digits are called the ‘host portion’. The number of bits that are shared by all addresses within a network is called the netmask and it is the role of the netmask to determine which addresses belong to the network it is applied to and which don’t. For example, consider the following:

Host Address	192.168.110.23
Network Mask	255.255.255.0
Network Portion	192.168.110.
Host portion	.23
Network Address	192.168.110.0
Broadcast Address	192.168.110.255

Any address that is “bitwise anded” with its netmask will reveal the address of the network it belongs to. The network address is therefore always the lowest numbered address within the range of addresses on the network and always has the host portion of the address coded all zeroes.

The broadcast address is a special address that every host on the network listens to in addition to its own unique address. This address is the one that datagrams are sent to if every host on the network is meant to receive it. Certain types of data like routing information and warning messages are transmitted to the broadcast address so that every host on the network can receive it simultaneously. There are two commonly used standards for what the broadcast address should be. The most widely accepted one is to use the highest possible address on

the network as the broadcast address. In the example above this would be 192.168.110.255. For some reason other sites have adopted the convention of using the network address as the broadcast address. In practice it doesn't matter very much which you use but you must make sure that every host on the network is configured with the same broadcast address.

For administrative reasons some time early in the development of the IP protocol some arbitrary groups of addresses were formed into networks and these networks were grouped into what are called classes. These classes provide a number of standard size networks that could be allocated. The ranges allocated are:

Network Class	Netmask	Network Addresses
A	255.0.0.0	0.0.0.0 - 127.255.255.255
B	255.255.0.0	128.0.0.0 - 191.255.255.255
C	255.255.255.0	192.0.0.0 - 223.255.255.255
Multicast	240.0.0.0	224.0.0.0 - 239.255.255.255

What addresses you should use depends on exactly what it is that you are doing. You may have to use a combination of the following activities to get all the addresses you need:

Installing a linux machine on an existing IP network

If you wish to install a linux machine onto an existing IP network then you should contact whoever administers the network and ask them for the following information:

- Host IP Address

- IP network address
- IP broadcast address
- IP netmask
- Router address
- Domain Name Server Address

You should then configure your linux network device with those details. You can not make them up and expect your configuration to work.

Building a brand new network that will never connect to the Internet

If you are building a private network and you never intend that network to be connected to the Internet then you can choose whatever addresses you like. However, for safety and consistency reasons there have been some IP network addresses that have been reserved specifically for this purpose. These are specified in RFC1597 and are as follows:

Network Class	Netmask	Network Addresses
A	255.0.0.0	10.0.0.0 - 10.255.255.255
B	255.255.0.0	172.16.0.0 - 172.31.255.255
C	255.255.255.0	192.168.0.0 - 192.168.255.255

Table 9-1. RESERVED PRIVATE NETWORK ALLOCATIONS

You should first decide how large you want your network to be

and then choose as many of the addresses as you require.

9.4.2. Routing

Routing is a big topic. It is easily possible to write large volumes of text about it. Most of you will have fairly simple routing requirements, some of you will not. I will cover some basic fundamentals of routing only. If you are interested in more detailed information then I suggest you refer to the references provided at the start of the document.

Let's start with a definition. What is IP routing ? Here is one that I'm using:

“IP Routing is the process by which a host with multiple network connections decides where to deliver IP datagrams it has received.”

It might be useful to illustrate this with an example. Imagine a typical office router, it might have a PPP link off the Internet, a number of ethernet segments feeding the workstations and another PPP link off to another office. When the router receives a datagram on any of its network connections, routing is the mechanism that it uses to determine which interface it should send the datagram to next. Simple hosts also need to route, all Internet hosts have two network devices, one is the loopback interface described above and the other is the one it uses to talk to the rest of the network, perhaps an ethernet, perhaps a PPP or SLIP serial interface.

Ok, so how does routing work ? Each host keeps a special list of routing rules, called a routing table. This table contains rows which typically contain at least three fields, the first is a destination address, the second is the name of the interface to which the datagram is to be routed and the third is optionally the IP address of another machine which

will carry the datagram on its next step through the network. In linux you can see this table by using the following command:

```
user% cat /proc/net/route
```

or by using either of the following commands:

```
user% /sbin/route -n  
user% netstat -r
```

The routing process is fairly simple: an incoming datagram is received, the destination address (who it is for) is examined and compared with each entry in the table. The entry that best matches that address is selected and the datagram is forwarded to the specified interface. If the gateway field is filled then the datagram is forwarded to that host via the specified interface, otherwise the destination address is assumed to be on the network supported by the interface.

9.4.2.1. what does the routed program do ?

The routing configuration described above is best suited to simple network arrangements where there are only ever single possible paths to destinations. When you have a more complex network arrangement things get a little more complicated. Fortunately for most of you this won't be an issue.

The big problem with 'manual routing' or 'static routing' as described, is that if a machine or link fails in your network then the only way you can direct your datagrams another way, if another way exists, is by manually intervening and executing the appropriate commands. Naturally this is clumsy, slow, impractical and hazard prone. Various techniques have been developed to automatically adjust routing tables in the event of network failures where there are alternate routes, all of

these techniques are loosely grouped by the term ‘dynamic routing protocols’.

You may have heard of some of the more common dynamic routing protocols. The most common are probably RIP (Routing Information Protocol) and OSPF (Open Shortest Path First Protocol). The Routing Information Protocol is very common on small networks such as small-medium sized corporate networks or building networks. OSPF is more modern and more capable at handling large network configurations and better suited to environments where there is a large number of possible paths through the network. Common implementations of these protocols are: ‘**routed**’ - RIP and ‘**gated**’ - RIP, OSPF and others. The ‘**routed**’ program is normally supplied with your Linux distribution or is included in the ‘NetKit’ package detailed above.

An example of where and how you might use a dynamic routing protocol might look something like the figure 9-1.




Figure 9-1. A dynamic routing example

We have three routers A, B and C. Each supports one ethernet segment with a Class C IP network (netmask 255.255.255.0). Each router

also has a PPP link to each of the other routers. The network forms a triangle.

It should be clear that the routing table at router A could look like:

```
root# route add -net 192.168.1.0 netmask 255.255.255.0 eth0
root# route add -net 192.168.2.0 netmask 255.255.255.0 ppp0
root# route add -net 192.168.3.0 netmask 255.255.255.0 ppp1
```

This would work just fine until the link between router A and B should fail. If that link failed then with the routing entry shown above hosts on the ethernet segment of A could not reach hosts on the ethernet segment on B because their datagram would be directed to router A's ppp0 link which is broken. They could still continue to talk to hosts on the ethernet segment of C and hosts on the C's ethernet segment could still talk to hosts on B's ethernet segment because the link between B and C is still intact.

But wait, if A can talk to C and C can still talk to B, why shouldn't A route its datagrams for B via C and let C send them to B ? This is exactly the sort of problem that dynamic routing protocols like RIP were designed to solve. If each of the routers A, B and C were running a routing daemon then their routing tables would be automatically adjusted to reflect the new state of the network should any one of the links in the network fail. To configure such a network is simple, at each router you need only do two things. In this case for Router A:

```
root# route add -net 192.168.1.0 netmask 255.255.255.0 eth0
root# /usr/sbin/routed
```

The '**routed**' routing daemon automatically finds all active network ports when it starts and sends and listens for messages on each of the network devices to allow it to determine and update the routing table on the host.

This has been a very brief explanation of dynamic routing and where you would use it. If you want more information then you should refer to the suggested references listed at the top of the document.

The important points relating to dynamic routing are:

1. You only need to run a dynamic routing protocol daemon when your Linux machine has the possibility of selecting multiple possible routes to a destination. An example of this would be if you plan to use IP Masquerading.
2. The dynamic routing daemon will automatically modify your routing table to adjust to changes in your network.
3. RIP is suited to small to medium sized networks.

9.5. Ethernet Information

This section covers information specific to Ethernet and the configuring of Ethernet Cards.

9.5.1. Supported Ethernet Cards

9.5.1.1. 3Com

- 3Com 3c501 - ‘avoid like the plague’ (3c501 driver)
- 3Com 3c503 (3c503 driver), 3c505 (3c505 driver), 3c507 (3c507 driver), 3c509/3c509B (ISA) / 3c579 (EISA)

- 3Com Etherlink III Vortex Ethercards (3c590, 3c592, 3c595, 3c597) (PCI), 3Com Etherlink XL Boomerang (3c900, 3c905) (PCI) and Cyclone (3c905B, 3c980) Ethercards (3c59x driver) and 3Com Fast EtherLink Ethercard (3c515) (ISA) (3c515 driver)
- 3Com 3ccfe575 Cyclone Cardbus (3c59x driver)
- 3Com 3c575 series Cardbus (3c59x driver) (ALL PCMCIA ??)

9.5.1.2. AMD, ATT, Allied Telesis, Ansel, Apricot

- AMD LANCE (79C960) / PCnet-ISA/PCI (AT1500, HP J2405A, NE1500/NE2100)
- ATT GIS WaveLAN
- Allied Telesis AT1700
- Allied Telesis LA100PCI-T
- Allied Telesyn AT2400T/BT ("ne" module)
- Ansel Communications AC3200 (EISA)
- Apricot Xen-II / 82596

9.5.1.3. Cabletron, Cogent, Crystal Lan

- Cabletron E21xx
- Cogent EM110
- Crystal Lan CS8920, Cs8900

9.5.1.4. Danpex, DEC, Digi, DLink

- Danpex EN-9400
- DEC DE425 (EISA) / DE434/DE435 (PCI) / DE450/DE500 (DE4x5 driver)
- DEC DE450/DE500-XA (dc21x4x) (Tulip driver)
- DEC DEPCA and EtherWORKS
- DEC EtherWORKS 3 (DE203, DE204, DE205)
- DECchip DC21x4x "Tulip"
- DEC QSilver's (Tulip driver)
- Digi International RightSwitch
- DLink DE-220P, DE-528CT, DE-530+, DFE-500TX, DFE-530TX

9.5.1.5. Fujitsu, HP, ICL, Intel

- Fujitsu FMV-181/182/183/184
- HP PCLAN (27245 and 27xxx series)
- HP PCLAN PLUS (27247B and 27252A)
- HP 10/100VG PCLAN (J2577, J2573, 27248B, J2585) (ISA/EISA/PCI)
- ICL EtherTeam 16i / 32 (EISA)
- Intel EtherExpress
- Intel EtherExpress Pro

9.5.1.6. KTI, Macromate, NCR NE2000/1000, Netgear, New Media

- KTI ET16/P-D2, ET16/P-DC ISA (work jumperless and with hardware-configuration options)
- Macromate MN-220P (PnP or NE2000 mode)
- NCR WaveLAN
- NE2000/NE1000 (be careful with clones)
- Netgear FA-310TX (Tulip chip)
- New Media Ethernet

9.5.1.7. PureData, SEEQ, SMC

- PureData PDUC8028, PDI8023
- SEEQ 8005
- SMC Ultra / EtherEZ (ISA)
- SMC 9000 series
- SMC PCI EtherPower 10/100 (DEC Tulip driver)
- SMC EtherPower II (epic100.c driver)

9.5.1.8. Sun Lance, Sun Intel, Schneider, WD, Zenith, IBM, Enyx

- Sun LANCE adapters (kernel 2.2 and newer)
- Sun Intel adapters (kernel 2.2 and newer)
- Schneider and Koch G16
- Western Digital WD80x3
- Zenith Z-Note / IBM ThinkPad 300 built-in adapter
- Znyx 312 etherarray (Tulip driver)

9.5.2. General Ethernet Information

Ethernet devices names are `eth0`, `eth1`, `eth2` etc. The first card detected by the kernel is assigned '`eth0`' and the rest are assigned sequentially in the order they are detected.

Once you have your kernel properly built to support your ethernet card then configuration of the card is easy.

Typically you would use something like (which most distributions already do for you, if you configured them to support your ethernet):

```
root# ifconfig eth0 192.168.0.1 netmask 255.255.255.0 up
root# route add -net 192.168.0.0 netmask 255.255.255.0 eth0
```

Most of the ethernet drivers were developed by Donald Becker (<mailto:becker@CESDIS.gsfc.nasa.gov>)

9.5.3. Using 2 or more Ethernet Cards in the same machine

The module will typically can detect all of the installed cards.

Information from the detection is stored in the file: `/etc/conf.modules`.

Consider that a user has 3 NE2000 cards, one at 0x300 one at 0x240, and one at 0x220. You would add the following lines to the `/etc/conf.modules` file:

```
alias eth0 ne
alias eth1 ne
alias eth2 ne
options ne io=0x220,0x240,0x300
```

What this does is tell the program **modprobe** to look for 3 NE based cards at the following addresses. It also states in which order they should be found and the device they should be assigned.

Most ISA modules can take multiple comma separated I/O values. For example:

```
alias eth0 3c501
alias eth1 3c501
options eth0 -o 3c501-0 io=0x280 irq=5
options eth1 -o 3c501-1 io=0x300 irq=7
```

The `-o` option allows for a unique name to be assigned to each module. The reason for this is that you can not have two copies of the same module loaded.

The `irq=` option is used to specify the hardware IRQ and the `io=` to specify the different io ports.

By default, the Linux kernel only probes for one Ethernet device, you need to pass command line arguments to the kernel in order to force detection of further boards.

To learn how to make your ethernet card(s) working under Linux you should refer to the Ethernet-HOWTO ([Ethernet-HOWTO.html](#)).

9.6. IP Related Information

These sections cover information specific to IP.

9.6.1. DNS

DNS stands for Domain Name System. It is the system responsible for mapping a machine name as `www.mandrakesoft.com` with the IP address of that machine, in this case: `216.71.116.162` at the time of writing. With DNS, mapping is available in both directions; from name to IP and vice-versa.

The DNS is composed of a great number of machine all over the *Internet* responsible for a certain number of names. Each machine is attributed a DNS server to which it can ask to map a particular name with its address. If that server do not have the answer, then it asks to another one and so on. You can also have a local DNS responsible for mapping addresses on your LAN.

We can differentiate two major classes of DNS': caching DNS and master DNS server. The first one only "remembers" a previous request and then can answer it without asking again a master DNS server. These latter are servers that are really responsible as a last resort to map an address with a name - or possibly tell that this name does not map any address.

9.6.2. DHCP and DHCPD

DHCP is an acronym for Dynamic Host Configuration Protocol. The creation of DHCP has made configuring the network on multiple hosts extremely simple. Instead of having to configure each host separately you can assign all of the commonly used parameters by the hosts using a DHCP server.

Each time the host boots up it will broadcast a packet to the network. This packet is a call to any DHCP servers that are located on the same segment to configure the host.

DHCP is extremely useful in assigning items such as the IP address, Netmask, and gateway of each host.

9.7. Using common PC hardware

9.7.1. ISDN

The Integrated Services Digital Network (ISDN) is a series of standards that specify a general purpose switched digital data network. An ISDN ‘call’ creates a synchronous point to point data service to the destination. ISDN is generally delivered on a high speed link that is broken down into a number of discrete channels. There are two different types of channels, the ‘B Channels’ which will actually carry the user data and a single channel called the ‘D channel’ which is used to send control information to the ISDN exchange to establish calls and other functions. In Australia for example, ISDN may be delivered on a 2Mbps link that is broken into 30 discrete 64kbps B channels with one 64kbps D channel. Any number of channels may be used at a time and in any combination. You could for example establish 30 separate calls to 30 different destinations at 64kbps each, or you could establish 15 calls to 15 different destinations at 128kbps each (two channels used per call), or just a small number of calls and leave the rest idle. A

channel may be used for either incoming or outgoing calls. The original intention of ISDN was to allow Telecommunications companies to provide a single data service which could deliver either telephone (via digitized voice) or data services to your home or business without requiring you to make any special configuration changes.

There are a few different ways to connect your computer to an ISDN service. One way is to use a device called a ‘Terminal Adaptor’ which plugs into the Network Terminating Unit that your telecommunications carrier will have installed when you got your ISDN service and presents a number of serial interfaces. One of those interfaces is used to enter commands to establish calls and configuration and the others are actually connected to the network devices that will use the data circuits when they are established. Linux will work in this sort of configuration without modification, you just treat the port on the Terminal Adaptor like you would treat any other serial device. Another way, which is the way the kernel ISDN support is designed for allows you to install an ISDN card into your Linux machine and then has your Linux software handle the protocols and make the calls itself.

Kernel Compile Options:

```
ISDN subsystem -->
<*> ISDN support
[ ] Support synchronous PPP
[ ] Support audio via ISDN
< > ICN 2B and 4B support
< > PCBIT-D support
< > Teles/NICCY1016PC/Creatix support
```

The Linux implementation of ISDN supports a number of different types of internal ISDN cards. These are those listed in the kernel configuration options:

- ICN 2B and 4B
- Octal PCBIT-D

- Teles ISDN-cards and compatibles

Some of these cards require software to be downloaded to them to make them operational. There is a separate utility to do this with.

Full details on how to configure the Linux ISDN support is available from the `/usr/src/linux/Documentation/isdn/` directory and a FAQ dedicated to **isdn4linux** is available at www.lrz-muenchen.de (<http://www.lrz-muenchen.de/~ui161ab/www/isdn/>). (You can click on the english flag to get an english version).

Note: About PPP. The PPP suite of protocols will operate over either asynchronous or synchronous serial lines. The commonly distributed PPP daemon for Linux '**pppd**' supports only asynchronous mode. If you wish to run the PPP protocols over your ISDN service you need a specially modified version. Details of where to find it are available in the documentation referred to above.

9.7.2. PLIP

During development of the 2.1 kernel versions, support for the parallel port was changed to a better setup.

Kernel Compile Options:

```
General setup -->
  [*] Parallel port support
Network device support -->
  <*> PLIP (parallel port) support
```


The new code for PLIP behaves like the old one (use the same **ifconfig** and **route** commands as in the previous section, but initialization of the device is different due to the advanced parallel port support.

The “first” PLIP device is always called `plip0`, where first is the first device detected by the system, similarly to what happens for Ethernet devices. The actual parallel port being used is one of the available ports, as shown in `/proc/parport`. For example, if you have only one parallel port, you’ll only have a directory called `/proc/parport/0`.

If your kernel didn’t detect the IRQ number used by your port, “`insmod plip`” will fail; in this case just write the right number to `/proc/parport/0/irq` and reinvoke **insmod**.

Complete information about parallel port management is available in the file `Documentation/parport.txt`, part of your kernel sources.

9.7.3. PPP

Due to the nature of PPP, its size, complexity, and flexibility it has been moved to its own HOWTO. The PPP-HOWTO is still a Linux Documentation Project document (<http://www.linuxdoc.org>) but its official home is at the LinuxPorts.Com website (<http://www.LinuxPorts.Com>) PPP section (<http://www.linuxports.com/howto/ppp>).

9.8. Other Network Technologies

The following subsections are specific to particular network technologies. The information contained in these sections does not necessarily apply to any other type of network technology. The topics are sorted alphabetically.

9.8.1. ARCNet

ARCNet device names are `arc0e`, `arc1e`, `arc2e` etc. or `arc0s`, `arc1s`, `arc2s` etc. The first card detected by the kernel is assigned `arc0e` or `arc0s` and the rest are assigned sequentially in the order they are detected. The letter at the end signifies whether you've selected ethernet encapsulation packet format or RFC1051 packet format.

Kernel Compile Options:

```
Network device support -->
[*] Network device support
<*> ARCnet support
[ ]   Enable arc0e (ARCnet "Ether-Encap" packet format)
[ ]   Enable arc0s (ARCnet RFC1051 packet format)
```

Once you have your kernel properly built to support your ethernet card then configuration of the card is easy.

Typically you would use something like:

```
root# ifconfig arc0e 192.168.0.1 netmask 255.255.255.0 up
root# route add -net 192.168.0.0 netmask 255.255.255.0 arc0e
```

Please refer to the `/usr/src/linux/Documentation/networking/arcnet.txt` and `/usr/src/linux/Documentation/networking/arcnet-hardware.txt` files for further information.

ARCNet support was developed by Avery Pennarun, `apenwarr@foxnet.net`.

9.8.2. Appletalk (AF_APPLETALK)

The Appletalk support has no special device names as it uses existing network devices.

Kernel Compile Options:

```
Networking options -->  
<*> Appletalk DDP
```

Appletalk support allows your Linux machine to interwork with Apple networks. An important use for this is to share resources such as printers and disks between both your Linux and Apple computers. Additional software is required, this is called **netatalk**. Wesley Craig `netatalk@umich.edu` represents a team called the ‘Research Systems Unix Group’ at the University of Michigan and they have produced the **netatalk** package which provides software that implements the Appletalk protocol stack and some useful utilities. The **netatalk** package will either have been supplied with your Linux distribution, or you will have to ftp it from its home site at the University of Michigan (`ftp://terminator.rs.itd.umich.edu/unix/netatalk/`)

To build and install the package do something like:

```
user% tar xvfz ../netatalk-1.4b2.tar.Z  
user% make  
root# make install
```

You may want to edit the ‘Makefile’ before calling **make** to actually compile the software. Specifically, you might want to change the `DESTDIR` variable which defines where the files will be installed later. The default of `/usr/local/atalk` is fairly safe.

9.8.2.1. Configuring the Appletalk software.

The first thing you need to do to make it all work is to ensure that the appropriate entries in the `/etc/services` file are present. The entries you need are:

```
rtpm 1/ddp # Routing Table Maintenance Protocol  
nbp 2/ddp # Name Binding Protocol  
echo 4/ddp # AppleTalk Echo Protocol
```

```
zip 6/ddp # Zone Information Protocol
```

The next step is to create the Appletalk configuration files in the `/usr/local/atalk/etc` directory (or wherever you installed the package).

The first file to create is the `/usr/local/atalk/etc/atalkd.conf` file. Initially this file needs only one line that gives the name of the network device that supports the network that your Apple machines are on:

```
eth0
```

The Appletalk daemon program will add extra details after it is run.

9.8.2.2. Exporting a Linux filesystems via Appletalk.

You can export filesystems from your linux machine to the network so that Apple machine on the network can share them.

To do this you must configure the `/usr/local/atalk/etc/AppleVolumes.` system file. There is another configuration file called `/usr/local/atalk/etc/AppleVolumes.default` which has exactly the same format and describes which filesystems users connecting with guest privileges will receive.

Full details on how to configure these files and what the various options are can be found in the **afpd** man page.

A simple example might look like:

```
/tmp Scratch  
/home/ftp/pub "Public Area"
```

Which would export your /tmp filesystem as AppleShare Volume ‘Scratch’ and your ftp public directory as AppleShare Volume ‘Public Area’. The volume names are not mandatory, the daemon will choose some for you, but it won’t hurt to specify them anyway.

9.8.2.3. Sharing your Linux printer across Appletalk.

You can share your linux printer with your Apple machines quite simply. You need to run the **papd** program which is the Appletalk Printer Access Protocol Daemon. When you run this program it will accept requests from your Apple machines and spool the print job to your local line printer daemon for printing.

You need to edit the /usr/local/atalk/etc/papd.conf file to configure the daemon. The syntax of this file is the same as that of your usual /etc/printcap file. The name you give to the definition is registered with the Appletalk naming protocol, NBP.

A sample configuration might look like:

```
TricWriter:\n      :pr=lp:op=cg:
```

Which would make a printer named ‘TricWriter’ available to your Appletalk network and all accepted jobs would be printed to the linux printer ‘lp’ (as defined in the /etc/printcap file) using **lpd**. The entry op=cg says that the linux user cg is the operator of the printer.

9.8.2.4. Starting the appletalk software.

Ok, you should now be ready to test this basic configuration. There is an **rc.atalk** file supplied with the **netatalk** package that should work ok for you, so all you should have to do is:

```
root# /usr/local/ataalk/etc/rc.ataalk
```

and all should startup and run ok. You should see no error messages and the software will send messages to the console indicating each stage as it starts.

9.8.2.5. Testing the appletalk software.

To test that the software is functioning properly, go to one of your Apple machines, pull down the Apple menu, select the Chooser, click on AppleShare, and your Linux box should appear.

9.8.2.6. Caveats of the appletalk software.

- You may need to start the Appletalk support before you configure your IP network. If you have problems starting the Appletalk programs, or if after you start them you have trouble with your IP network, then try starting the Appletalk software before you run your `/etc/rc.d/rc.inet1` file.
- The **afpd** (Apple Filing Protocol Daemon) severely messes up your hard disk. Below the mount points it creates a couple of directories called “.AppleDesktop” and Network Trash Folder. Then, for each directory you access it will create a .AppleDouble below it so it can store resource forks, etc. So think twice before exporting /, you will have a great time cleaning up afterwards.
- The **afpd** program expects clear text passwords from the Macs. Security could be a problem, so be very careful when you run this daemon on a machine connected to the Internet, you have yourself to blame if somebody nasty does something bad.

- The existing diagnostic tools such as **netstat** and **ifconfig** don't support Appletalk. The raw information is available in the `/proc/net/` directory if you need it.

9.8.2.7. More information

For a much more detailed description of how to configure Appletalk for Linux refer to Anders Brownworth **Linux Netatalk-HOWTO** page at thehamptons.com (<http://thehamptons.com/anders/netatalk/>).

9.8.3. ATM

Werner Almesberger <werner.almesberger@lrc.di.epfl.ch> is managing a project to provide Asynchronous Transfer Mode support for Linux. Current information on the status of the project may be obtained from: [lrcwww.epfl.ch](http://lrcwww.epfl.ch/linux-atm/) (<http://lrcwww.epfl.ch/linux-atm/>).

9.8.4. AX25 (AF_AX25)

AX.25 device names are `sl0`, `sl1`, etc. in 2.0.* kernels or `ax0`, `'ax1`, etc. in 2.1.* kernels.

Kernel Compile Options:

```
Networking options -->
  [*] Amateur Radio AX.25 Level 2
```

The AX25, Netrom and Rose protocols are covered by the AX25-HOWTO (AX25-HOWTO.html). These protocols are used by Amateur Radio Operators world wide in packet radio experimentation.

Most of the work for implementation of these protocols has been done by Jonathon Naylor, jsn@cs.nott.ac.uk.

9.8.5. DECNet

Support for DECNet is now included in current stable kernel (2.4) **Mandrake** has also made it available in its 2.2 kernels.

9.8.6. FDDI

FDDI device names are fddi0, fddi1, fddi2 etc. The first card detected by the kernel is assigned fddi0 and the rest are assigned sequentially in the order they are detected.

Larry Stefani, lstefani@ultranet.com, has developed a driver for the Digital Equipment Corporation FDDI EISA and PCI cards.

Kernel Compile Options:

```
Network device support -->
  [*] FDDI driver support
  [*] Digital DEFEA and DEFPD adapter support
```

When you have your kernel built to support the FDDI driver and installed, configuration of the FDDI interface is almost identical to that of an ethernet interface. You just specify the appropriate FDDI interface name in the **ifconfig** and **route** commands.

9.8.7. Frame Relay

The Frame Relay device names are `dlci00`, `dlci01` etc for the DLCI encapsulation devices and `sdla0`, `sdla1` etc for the FRAD(s).

Frame Relay is a new networking technology that is designed to suit data communications traffic that is of a ‘bursty’ or intermittent nature. You connect to a Frame Relay network using a Frame Relay Access Device (FRAD). The Linux Frame Relay supports IP over Frame Relay as described in RFC-1490.

Kernel Compile Options:

```
Network device support -->
  <*> Frame relay DLCI support (EXPERIMENTAL)
    (24)   Max open DLCI
    (8)    Max DLCI per device
  <*>   SDLA (Sangoma S502/S508) support
```

Mike McLagan, mike.mclagan@linux.org, developed the Frame Relay support and configuration tools.

Currently the only FRAD I know of that are supported are the Sangoma Technologies (<http://www.sangoma.com/>) S502A, S502E and S508. and the Emerging Technologies. The Emerging Technologies website is here (<http://www.etinc.com/>).

To configure the FRAD and DLCI devices after you have rebuilt your kernel you will need the Frame Relay configuration tools. These are available from [ftp.invlogic.com](ftp://ftp.invlogic.com) (<ftp://ftp.invlogic.com/pub/linux/fr/frad-0.15.tgz>).

Compiling and installing the tools is straightforward, but the lack of a top level Makefile makes it a fairly manual process:

```
user% tar xvfz ../frad-0.15.tgz
user% cd frad-0.15
user% for i in common dlci frad; do make -C $i clean; make -C $i; done
root# mkdir /etc/frad
root# install -m 644 -o root -g root bin/*.sfm /etc/frad
root# install -m 700 -o root -g root frad/fradcfg /sbin
```

Chapter 9. Networking Overview

```
rppt# install -m 700 -o root -g root dlci/dlcicfg /sbin
```

Note that the previous commands use **sh** syntax, if you use a **cs**h flavor instead (like **tcsh**), the **for** loop will look different.

After installing the tools you need to create an `/etc/frad/router.conf` file. You can use this template, which is a modified version of one of the example files:

```
# /etc/frad/router.conf
# This is a template configuration for frame relay.
# All tags are included. The default values are based on the code
# supplied with the DOS drivers for the Sangoma S502A card.
#
# A '#' anywhere in a line constitutes a comment
# Blanks are ignored (you can indent with tabs too)
# Unknown [] entries and unknown keys are ignored
#

[Devices]
Count=1                # number of devices to configure
Dev_1=sdl1a0           # the name of a device
Dev_2=sdl1a1           # the name of a device

# Specified here, these are applied to all devices and can be overridden for
# each individual board.
#
Access=CPE
Clock=Internal
KBaud=64
Flags=TX
#
# MTU=1500              # Maximum transmit IFrame length, default is 4096
# T391=10               # T391 value      5 - 30, default is 10
# T392=15               # T392 value      5 - 30, default is 15
# N391=6                # N391 value      1 - 255, default is 6
# N392=3                # N392 value      1 - 10, default is 3
# N393=4                # N393 value      1 - 10, default is 4

# Specified here, these set the defaults for all boards
# CIRfwd=16             # CIR forward     1 - 64
# Bc_fwd=16             # Bc forward      1 - 512
# Be_fwd=0              # Be forward      0 - 511
# CIRbak=16             # CIR backward    1 - 64
# Bc_bak=16             # Bc backward     1 - 512
```

```
# Be_bak=0                # Be backward    0 - 511

#
#
# Device specific configuration
#
#

#
# The first device is a Sangoma S502E
#
[sdla0]
Type=Sangoma              # Type of the device to configure, currently only
                          # SANGOMA is recognized

#
# These keys are specific to the 'Sangoma' type
#
# The type of Sangoma board - S502A, S502E, S508
Board=S502E
#
# The name of the test firmware for the Sangoma board
# Testware=/usr/src/frad-0.10/bin/sdla_tst.502
#
# The name of the FR firmware
# Firmware=/usr/src/frad-0.10/bin/frm_rel.502
#
Port=360                  # Port for this particular card
Mem=C8                    # Address of memory window, A0-EE, depending on card
IRQ=5                     # IRQ number, do not supply for S502A
DLCIs=1                   # Number of DLCI's attached to this device
DLCI_1=16                 # DLCI #1's number, 16 - 991
# DLCI_2=17
# DLCI_3=18
# DLCI_4=19
# DLCI_5=20
#
# Specified here, these apply to this device only,
# and override defaults from above
#
# Access=CPE              # CPE or NODE, default is CPE
# Flags=TXIgnore,RXIgnore,BufferFrames,DropAborted,Stats,MCI,AutoDLCI
# Clock=Internal          # External or Internal, default is Internal
# Baud=128                 # Specified baud rate of attached CSU/DSU
# MTU=2048                 # Maximum transmit IFrame length, default is 4096
# T391=10                  # T391 value      5 - 30, default is 10
# T392=15                  # T392 value      5 - 30, default is 15
# N391=6                   # N391 value      1 - 255, default is 6
# N392=3                   # N392 value      1 - 10, default is 3
# N393=4                   # N393 value      1 - 10, default is 4
```

Chapter 9. Networking Overview

```
#
# The second device is some other card
#
# [sdla1]
# Type=FancyCard          # Type of the device to configure.
# Board=                  # Type of Sangoma board
# Key=Value               # values specific to this type of device

#
# DLCI Default configuration parameters
# These may be overridden in the DLCI specific configurations
#
CIRfwd=64                 # CIR forward    1 - 64
# Bc_fwd=16               # Bc forward    1 - 512
# Be_fwd=0                # Be forward    0 - 511
# CIRbak=16               # CIR backward 1 - 64
# Bc_bak=16               # Bc backward 1 - 512
# Be_bak=0                # Be backward 0 - 511

#
# DLCI Configuration
# These are all optional. The naming convention is
# [DLCI_D<devicenum>_<DLCI_Num>]
#

[DLCI_D1_16]
# IP=
# Net=
# Mask=
# Flags defined by Sangoma: TXIgnore,RXIgnore,BufferFrames
# DLCIFlags=TXIgnore,RXIgnore,BufferFrames
# CIRfwd=64
# Bc_fwd=512
# Be_fwd=0
# CIRbak=64
# Bc_bak=512
# Be_bak=0

[DLCI_D2_16]
# IP=
# Net=
# Mask=
# Flags defined by Sangoma: TXIgnore,RXIgnore,BufferFrames
# DLCIFlags=TXIgnore,RXIgnore,BufferFrames
# CIRfwd=16
# Bc_fwd=16
# Be_fwd=0
# CIRbak=16
# Bc_bak=16
# Be_bak=0
```

When you've built your `/etc/frad/router.conf` file the only step remaining is to configure the actual devices themselves. This is only a little trickier than a normal network device configuration, you need to remember to bring up the FRAD device before the DLCI encapsulation devices. These commands are best hosted in a shell script, due to their number:

```
#!/bin/sh
# Configure the frad hardware and the DLCI parameters
/sbin/fradcfg /etc/frad/router.conf || exit 1
/sbin/dlcicfg file /etc/frad/router.conf
#
# Bring up the FRAD device
ifconfig sdla0 up
#
# Configure the DLCI encapsulation interfaces and routing
ifconfig dlci00 192.168.10.1 pointopoint 192.168.10.2 up
route add -net 192.168.10.0 netmask 255.255.255.0 dlci00
#
ifconfig dlci01 192.168.11.1 pointopoint 192.168.11.2 up
route add -net 192.168.11.0 netmask 255.255.255.0 dlci00
#
route add default dev dlci00
#
```

9.8.8. IPX (AF_IPX)

The IPX protocol is most commonly utilized in Novell NetWare(tm) local area network environments. Linux includes support for this protocol and may be configured to act as a network endpoint, or as a router for IPX.

Kernel Compile Options:

```
Networking options -->
  [*] The IPX protocol
  [ ] Full internal IPX network
```

The IPX protocol and the NCPFS are covered in greater depth in the IPX-HOWTO (IPX-HOWTO.html).

9.8.9. NetRom (AF_NETROM)

NetRom device names are nr0, nr1, etc.

Kernel Compile Options:

```
Networking options -->
  [*] Amateur Radio AX.25 Level 2
  [*] Amateur Radio NET/ROM
```

The AX25, Netrom and Rose protocols are covered by the AX25-HOWTO (AX25-HOWTO.html). These protocols are used by Amateur Radio Operators world wide in packet radio experimentation.

Most of the work for implementation of these protocols has been done by Jonathon Naylor, jsn@cs.nott.ac.uk.

9.8.10. Rose protocol (AF_ROSE)

Rose device names are rs0, rs1, etc. in 2.1.* kernels. Rose is available in the 2.1.* kernels.

Kernel Compile Options:

```
Networking options -->
  [*] Amateur Radio AX.25 Level 2
  <*> Amateur Radio X.25 PLP (Rose)
```

The AX25, Netrom and Rose protocols are covered by the AX25-HOWTO (AX25-HOWTO.html). These protocols are used by Amateur Radio Operators world wide in packet radio experimentation.

Most of the work for implementation of these protocols has been done by Jonathon Naylor, jsn@cs.nott.ac.uk.

9.8.11. SAMBA - ‘NetBEUI’, ‘NetBios’, ‘CIFS’ support.

SAMBA is an implementation of the Session Management Block protocol. Samba allows Windows and other systems to mount and use your disks and printers.

SAMBA and its configuration are covered in detail in the SMB-HOWTO ([SMB-HOWTO.html](#)).

9.8.12. STRIP support (Starmode Radio IP)

STRIP device names are ‘st0’, ‘st1’, etc.

Kernel Compile Options:

```
Network device support -->
[*] Network device support
....
[*] Radio network interfaces
< > STRIP (Metricom starmode radio IP)
```

STRIP is a protocol designed specifically for a range of Metricom radio modems for a research project being conducted by Stanford University called the MosquitoNet Project (<http://mosquitonet.Stanford.EDU/mosquitonet.html>). There is a lot of interesting reading here, even if you aren’t directly interested in the project.

The Metricom radios connect to a serial port, employ spread spectrum technology and are typically capable of about 100kbps. Information on the Metricom radios is available from the: Metricom Web Server (<http://www.metricom.com/>).

At present the standard network tools and utilities do not support the STRIP driver, so you will have to download some customized tools from the MosquitoNet web server. Details on what software you need is available at the: MosquitoNet STRIP Page (<http://mosquitonet.Stanford.EDU/strip.html>).

A summary of configuration is that you use a modified **slattach** program to set the line discipline of a serial tty device to STRIP and then configure the resulting 'st[0-9]' device as you would for ethernet with one important exception, for technical reasons STRIP does not support the ARP protocol, so you must manually configure the ARP entries for each of the hosts on your subnet. This shouldn't prove too onerous.

9.8.13. Token Ring

Token ring device names are 'tr0', 'tr1' etc. Token Ring is an IBM standard LAN protocol that avoids collisions by providing a mechanism that allows only one station on the LAN the right to transmit at a time. A 'token' is held by one station at a time and the station holding the token is the only station allowed to transmit. When it has transmitted its data it passes the token onto the next station. The token loops amongst all active stations, hence the name 'Token Ring'.

Kernel Compile Options:

```
Network device support -->
[*] Network device support
....
[*] Token Ring driver support
< > IBM Tropic chipset based adaptor support
```

Configuration of token ring is identical to that of ethernet with the exception of the network device name to configure.

9.8.14. X.25

X.25 is a circuit based packet switching protocol defined by the C.C.I.T.T. (a standards body recognized by Telecommunications companies in most parts of the world). An implementation of X.25 and LAPB are being worked on and recent kernels (from 2.1.*) include the work in progress.

Jonathon Naylor jsn@cs.nott.ac.uk is leading the development and a mailing list has been established to discuss Linux X.25 related matters. To subscribe send a message to: majordomo@vger.rutgers.edu with the text `subscribe linux-x25` in the body of the message.

Early versions of the configuration tools may be obtained from Jonathon's ftp site at [ftp.cs.nott.ac.uk](ftp://ftp.cs.nott.ac.uk/jsn/) (<ftp://ftp.cs.nott.ac.uk/jsn/>).

9.8.15. WaveLan Card

Wavelan device names are `eth0`, `eth1`, etc.

Kernel Compile Options:

```
Network device support -->
[*] Network device support
....
[*] Radio network interfaces
....
<*> WaveLAN support
```

The WaveLAN card is a spread spectrum wireless lan card. The card looks very like an ethernet card in practice and is configured in much the same way.

You can get information on the Wavelan card from [Wavelan.com](http://www.wavelan.com) (<http://www.wavelan.com/>).

9.9. Cables and Cabling

Those of you handy with a soldering iron may want to build your own cables to interconnect two linux machines. The following cabling diagrams should assist you in this.

9.9.1. Serial NULL Modem cable

Not all NULL modem cables are alike. Many null modem cables do little more than trick your computer into thinking all the appropriate signals are present and swap transmit and receive data. This is ok but means that you must use software flow control (XON/XOFF) which is less efficient than hardware flow control. The following cable provides the best possible signalling between machines and allows you to use hardware (RTS/CTS) flow control.




Figure 9-2. The NULL-modem cabling

9.9.2. Parallel port cable (PLIP cable)

If you intend to use the PLIP protocol between two machines then this cable will work for you irrespective of what sort of parallel ports you

have installed.

Pin Name	pin	pin
STROBE	1*	
D0->ERROR	2	----- 15
D1->SLCT	3	----- 13
D2->PAPOUT	4	----- 12
D3->ACK	5	----- 10
D4->BUSY	6	----- 11
D5	7*	
D6	8*	
D7	9*	
ACK->D3	10	----- 5
BUSY->D4	11	----- 6
PAPOUT->D2	12	----- 4
SLCT->D1	13	----- 3
FEED	14*	
ERROR->D0	15	----- 2
INIT	16*	
SLCTIN	17*	
GROUND	25	----- 25

Notes:

- Do not connect the pins marked with an asterisk ‘*’.
- Extra grounds are 18,19,20,21,22,23 and 24.
- If the cable you are using has a metallic shield, it should be connected to the metallic DB-25 shell at **one end only**.

Warning

A miswired PLIP cable can destroy your controller card. Be very careful and double check every connection to ensure you don't cause yourself any unnecessary work or heartache.

While you may be able to run PLIP cables for long distances, you should avoid it if you can. The specifications for the cable allow for

a cable length of about 1 meter or so. Please be very careful when running long PLIP cables as sources of strong electromagnetic fields such as lightning, power lines and radio transmitters can interfere with and sometimes even damage your controller. If you really want to connect two of your computers over a large distance you really should be looking at obtaining a pair of thin-net ethernet cards and running some coaxial cable.

9.9.3. 10base2 (thin coax) Ethernet Cabling

10base2 is an ethernet cabling standard that specifies the use of 50 ohm coaxial cable with a diameter of about 5 millimeters. There are a couple of important rules to remember when interconnecting machines with 10base2 cabling. The first is that you must use terminators at **both ends** of the cabling. A terminator is a 50 ohm resistor that helps to ensure that the signal is absorbed and not reflected when it reaches the end of the cable. Without a terminator at each end of the cabling you may find that the ethernet is unreliable or doesn't work at all. Normally you'd use 'T pieces' to interconnect the machines, so that you end up with something that looks like:




Figure 9-3. 10base2 Ethernet Cabling

where the ‘|’ at either end represents a terminator, the ‘=====’ represents a length of coaxial cable with BNC plugs at either end and the ‘T’ represents a ‘T piece’ connector. You should keep the length of cable between the ‘T piece’ and the actual ethernet card in the PC as short as possible, ideally the ‘T piece’ will be plugged directly into the ethernet card.

9.9.4. Twisted Pair Ethernet Cable

If you have only two twisted pair ethernet cards and you wish to connect them you do not require a hub. You can cable the two cards directly together. A diagram showing how to do this is included in the Ethernet-HOWTO ([Ethernet-HOWTO.html](#))

III. System setup and management

Chapter 10. MandrakeSecurity Setup and Management

10.1. Introduction

This chapter is dedicated to the use of the web administration tool which allows you to remotely control your firewall from any machine of your local network. We will first describe the interface and then go through the various available screens. However, two sections of the interface are treated in the maintenance part of this book: part IV in *MandrakeSecurity*. We focus here on initial setup and later reconfigurations and tuning of the services.

As a matter of fact, we already configured the minimum needed for the firewall to work during the *MandrakeSecurity* installation. We will now go into configuration details since the firewall is pretty useless in its current state. The different sections should be read and applied chronologically.

10.2. Presentation of the Interface

We will briefly present the interface and how to navigate through it. It is basically made of a menu leading to configuration wizards.

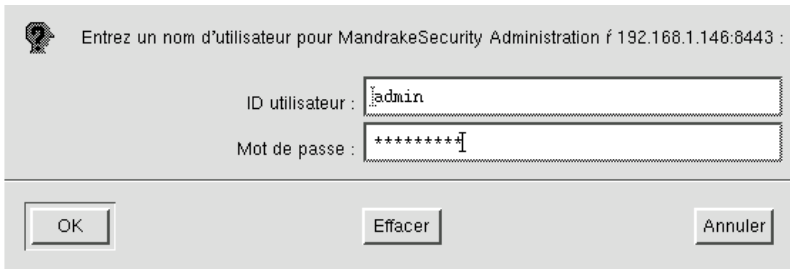
10.2.1. Connecting

The connection to the firewall server from any client can be done through any modern graphical web browser. The communication is entirely encrypted in order for no one to intercept the transferred information, especially passwords.

To initiate the session, type in the location field of your browser the URL that was given to you in the last screen of the installation procedure. It is something like

`https://192.168.1.160:8443/`

where 192.168.1.160 is the IP address of the firewall you choose in the LAN. You will then get some screens about a new certificate, accept it. Finally the welcome screen of *MandrakeSecurity* appears. Click on the Login button and an identification windows appears (figure 10-1).



Entrez un nom d'utilisateur pour MandrakeSecurity Administration f 192.168.1.146:8443 :

ID utilisateur : admin

Mot de passe : *****

OK Effacer Annuler

Figure 10-1. The Login Window to Connect to MandrakeSecurity

Fill it with the admin login and password as defined during the installation. Whenever you are asked to identify to connect to the interface, always use the admin login.

10.2.2. The Interface








Figure 10-2. A Sample MandrakeSecurity Interface Screen

The interface is designed in a traditional way with a two-level menu on the left and a content frame on the right. The content frame is organized in various tabs, each one corresponding to the second-level entries of the selected menu.

Then each of the tabs contains a wizard dedicated to the configuration of a particular aspect of the server. Each page of the wizard is made of:

- Informative text: what is that screen about.
- User-entry fields: to fill or select according to your choices.
- Buttons: to perform special actions.

You will also see icons. Here are the most important ones:

	You will get a pop-up windows displaying help on that particular screen, informing you on the signification of the various elements present in it.
	This button discards all changes made since the beginning of the wizard and goes back to the first step.
	Goes back to the previous wizard step.
	Goes to the next wizard step.
	At a wizard summary's last screen, confirms the choices and applies them to the system.

10.2.3. Logout

It is very important to explicitly log out of the interface when you are done with all your tasks, or whenever you leave your display. Actually, simply closing the browser is not enough since the server has no mean to know that you want to close your session, and someone else using your computer right after you could take your session where you left off.



Figure 10-3. The log out menu entry

Whenever you finish a session, simply click on that icon. Next time you try to reconnect, you will be asked to identify again.

10.3. Basic System Configuration

This section is for basic server setup. It also allows the administrator to change his password in order to access the interface.

10.3.1. General System Configuration



Here the system will be attributed a name which will be allocated to a local network. Which parameters will be entered at this point, depend on whether or not you have a permanent access to the Internet with a fixed IP address.

System Name	machine.domain.net
--------------------	--------------------

This field holds the full host-name of your machine: the machine name followed by the domain name.

Domain Name	domain.net
--------------------	------------

This field holds the domain name of the machine. If you hold a domain name and have the required DNSs to point to your IP address, use it here. Otherwise use the domain name of your Internet Service Provider.

10.3.2. Changing The Administrator's Password



This form will enable you to modify the admin login password. We recommend it be changed periodically.

Login Name	admin
New Password	*****
New Password (again)	*****

You need to choose a safe password. When done, click on the "Change" button. A dialog box will appear asking you for the new password. Enter it there and again on the welcome page.

10.3.3. Ethernet Cards Configuration



This screen lists the Network Interface Cards (NIC) currently configured on your machine. It will enable you to select a particular card and reconfigure it, or add another card.

```
IP Address Subnet Mask On Boot Protocol
Eth0 192.168.1.160 255.255.255.0 yes boot admin suppress
Eth1 10.0.0.1 255.0.0.0 no boot admin suppress
```

Each line corresponds to a physical NIC in your computer.

- To reconfigure it, select its name (Ethx)
- To activate it on each boot, select "Boot"

- To allow the network, associated to this interface, to connect to the web interface, click on "Admin" (see "Administration interface" below).
- To suppress it, select "suppress"

Administration Interface	Eth0
---------------------------------	------

Indicates the interface through which administration connections are allowed. This signifies that your firewall will have to be administered from a computer connected to the sub-network which is associated to the aforementioned card.

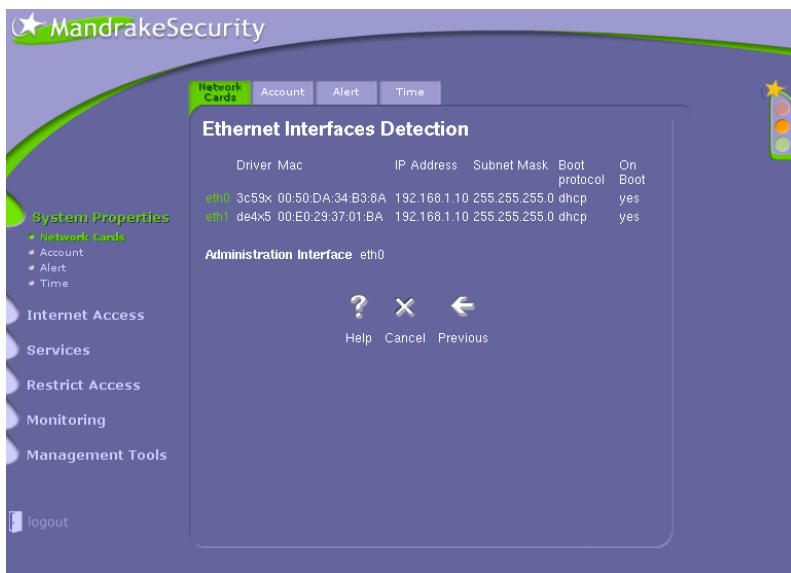


: Clicking on that icon will launch a NIC auto-detection process. Use it if you previously installed a new NIC on your computer. Note: after your click, it may take some time for the next screen to appear as the computer is detecting new cards.



Should the previous action fail, you can manually configure your card by clicking on that icon.

10.3.3.1. Ethernet Interfaces Detection

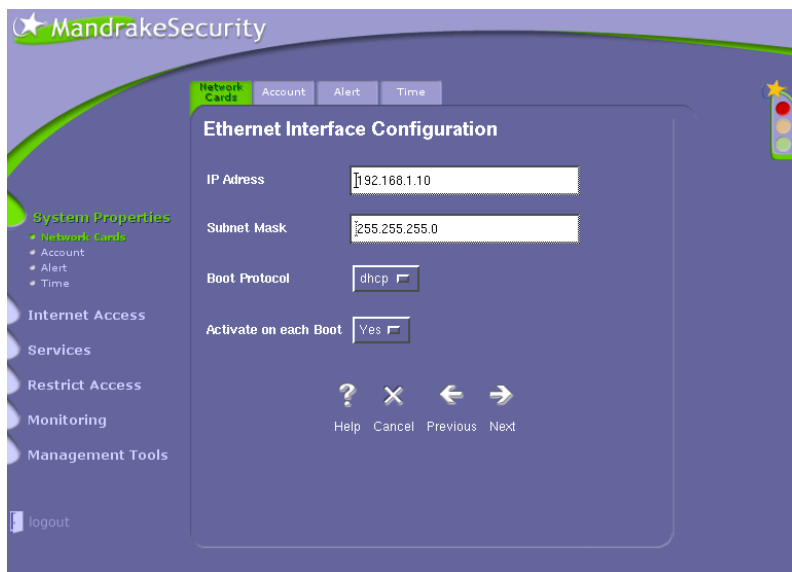


This screen shows the NICs which has just been automatically detected on your machine. If the card you wish to configure does not appear here, come back to the previous page and choose "Add a NIC manually".

```
Driver Mac IP Address Subnet Mask On Boot
Eth0 ne2k-pci 00:40:05:E2:55:F6 192.168.1.160 255.255.255.0 yes
```

Each line corresponds to a physical NIC in your computer. To select and configure it as your local network access interface, click on its name (ETHx)

10.3.3.2. Ethernet Interface Configuration For Your Local Network(s)



Here you will have to define the parameters of the interface card necessary to map the needs of your local network(s). Some of the parameters may have been chosen already during the install or a previous configuration and/or filled out with standard values. Make the necessary modifications to answer your present needs.

IP Address	192.168.1.1
-------------------	-------------

Fill out this field if you have a static IP address for that interface. This address is essential because it is your server's, that is the one client systems will have to refer to.

Subnet Mask (ex: 255.0.0.0)	255.255.255.0
------------------------------------	---------------

In this field, enter the name of the subnet mask related to the network to which this interface is connected.

Now set the boot protocol to be used when this interface is initialized. This depends on the protocol used by your ISP. Select the right checkbox, i.e. one of the following:

- static
- dhcp
- bootp

Finally, decide whether or not you want this interface to be activated on each boot.

10.3.4. Time Configuration



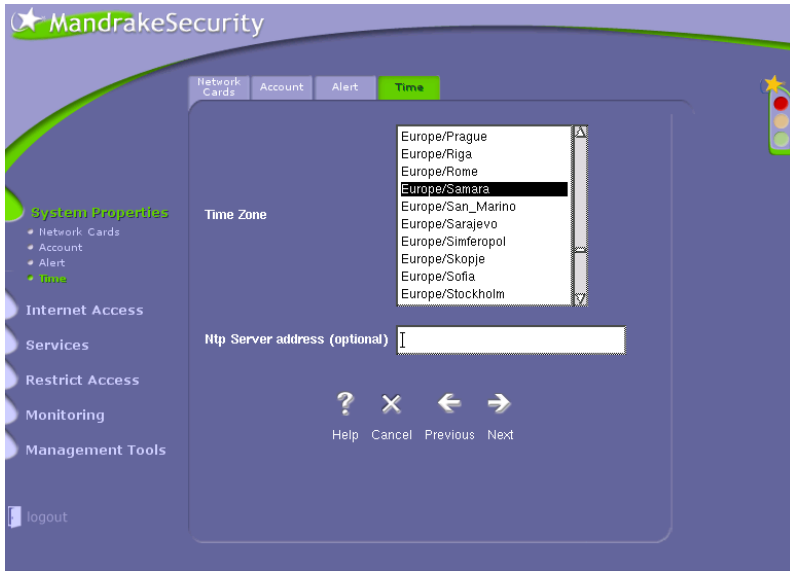
As a first step, the wizard will make two suggestions as to the internal time configuration.



Click on the "Modify" icon of what you want to setup:

- Time zone and NTP server address: to indicate the physical location of the server, and eventually set up a time server which would automatically set the system's date and time.
- Date and time: if you have no NTP server, click on the "Modify" button to manually set the current date and time on the machine.

10.3.4.1. Time Zone And NTP Server Configuration



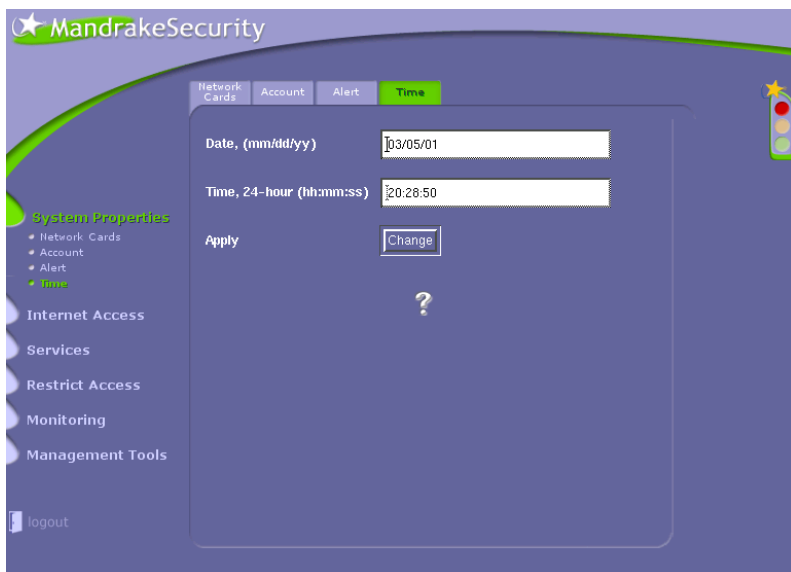
You need to pick the time zone of your physical location and indicate the presence of an eventual NTP server.

Time Zone	Europe/Paris
Ntp Server address (optional)	ntp.myco.com

In the list, select the time zone and then the city closest to you.

Eventually, you can enter the name of a NTP (Network Time Protocol) server, which automatically sets up and checks your clock periodically. If your company has its own server, use it. Otherwise, you can find a public server, e.g.: <http://www.eecis.udel.edu/~mills/ntp/clock2.htm>

10.3.4.2. Time And Date Setup



Simply enter the current date and time in the respective fields:

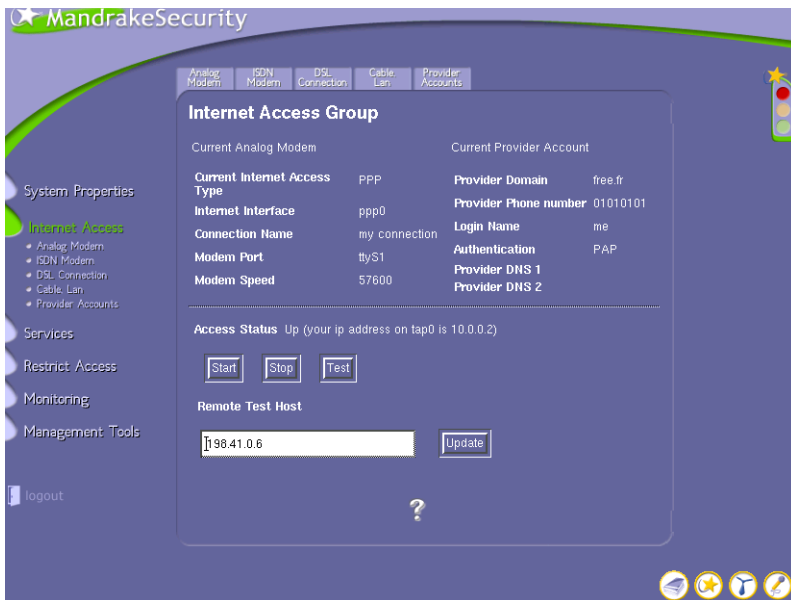
Date, (mm/dd/yy)	02/17/01
Time, 24-hour (hh:mm:ss)	17:07:58

Then apply any modifications made by clicking on the "Change" button.

10.4. Internet Access

This section allows you to configure how your server will access the *Internet*. It enables the configuration of interfaces with most common protocols, as well as the definition of all provider accounts.

10.4.1. Internet Access Group



This is the central page for Internet access: It summarizes current Internet access configuration and allows to bring up or down this connection.

First of all, you are reminded the type of Internet connection currently used, and all parameters related to that configuration.

Access Status	up
----------------------	----

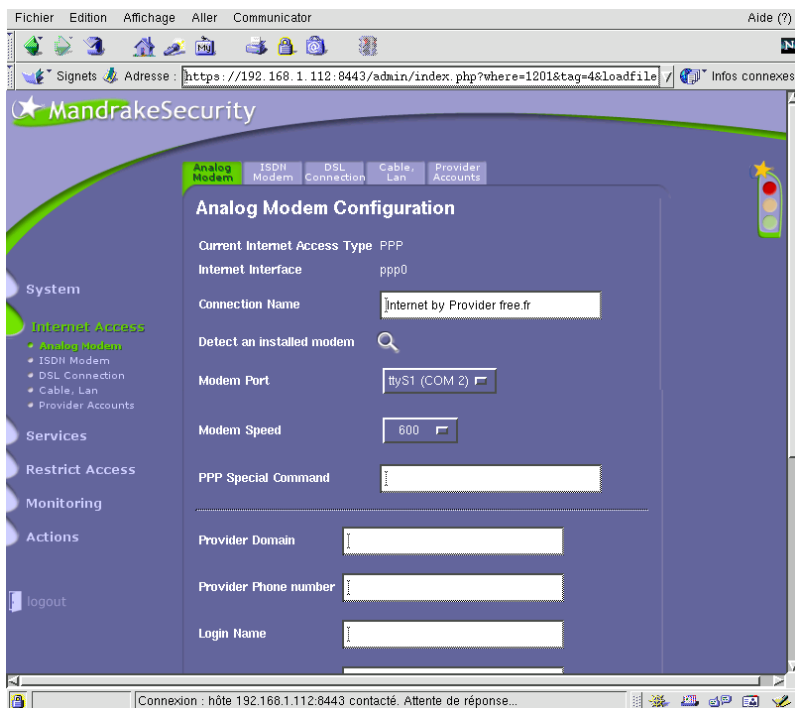
You are then informed of the current status of the connection: either "Up" or "Down". Then you can choose pressing the buttons to:

- Start: brings up the Internet connection. Note that this action override the possible connection schedule settings.
- Stop: brings down the Internet connection. Note that this action override the possible connection schedule settings.
- Test: tries to reach a specific host on the Internet and updates the link status.

Remote Test Host	198.41.0.6
-------------------------	------------

Here, you can choose the test host used to determine whether the Internet is reachable or not. When you change that, click on the "Update" button to make the change effective.

10.4.2. Analog Modem Configuration



This form contains all the parameters required to configure a standard analog modem connection to the Internet. Make sure you have all parameters from your ISP.

First there may be some reminders about the current Internet connection configuration..

Connection Name	Internet by Provider free.fr
------------------------	------------------------------

Fill this field out with any name that fits the configuration so that you can remember which connection it is relevant to.

You can then try to auto-detect the modem connected to your machine,

by clicking on the "Detect" icon:



Detected Modem List	ttyS0
----------------------------	-------

This list contains all the modems detected on the ports of your machine (The first serial port in this example). Choose the one you wish to use for this connection.

Modem Port	ttyS1 (COM 2)
-------------------	---------------

If your modem couldn't be detected, you can always manually select the port to which it is attached in this list.

Modem Speed	57600
--------------------	-------

Simply choose the maximum transfer speed of your modem (in bits/second).

PPP Special Command	
----------------------------	--

In the case that your connection need to pass special options to the pppd daemon, you may put them here. You shouldn't need to write anything here for most cases.

Provider Domain	free.fr
------------------------	---------

The domain of your provider.

Provider Phone number	0123456587
------------------------------	------------

The dial-in number of the Internet Provider.

Login name	foo
Password	*****
Password (confirm)	*****

Carefully enter here the login name and password as provided by your ISP.

Authentication	PAP
-----------------------	-----

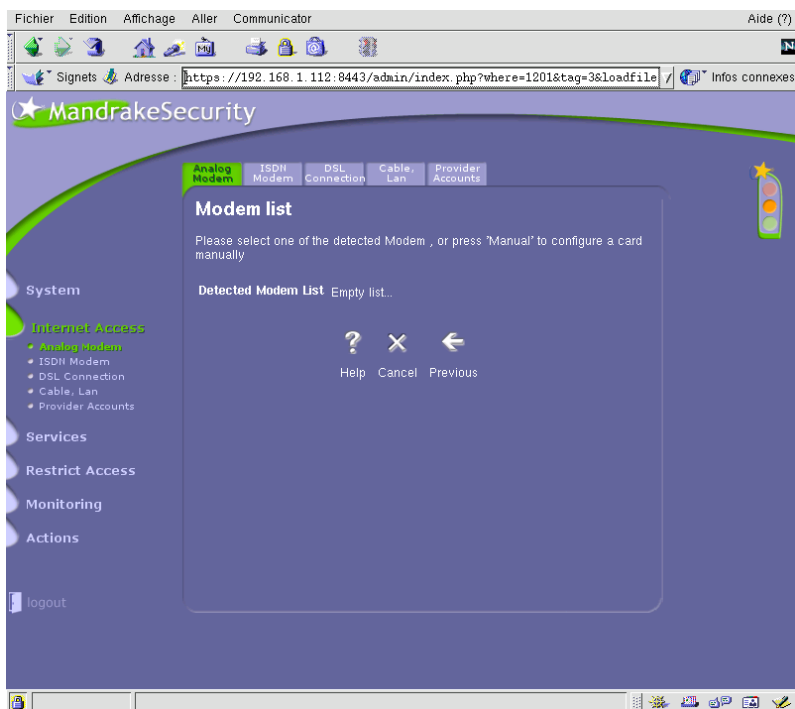
The authentication mechanism used by your ISP. Generally PAP.

Provider DNS 1	123.456.789.122
Provider DNS 2	123.456.789.123

The Domain Name Servers of your ISP.

When you are done with all fields, go on to next step. You will be able to review all parameters and then confirm your choices. The connection will be configured immediately.

10.4.2.1. Analog Modem List



This page displays all modems detected on your machine. Make sure modems are correctly connected and powered on before opening this page.

Detected Modem List	ttyS0 (COM1)
----------------------------	--------------

In the pop-down menu, simply choose the port into which the requested modem is plugged, and go on to the next step.

10.4.3. Configure your ISDN Internet Access

10.4.3.1. Choose The Type Of ISDN Card



This first step of the ISDN configuration wizard provides you with various options:

- Detect an internal card: The selection of this icon will bring up a list of the ISDN cards detected on your machine. If you have an internal card try this first. Otherwise you will need to either:
- Select an internal card: a list of supported ISDN cards will be shown if the previous step has failed.

- Configure an external modem: Select this icon if you have an external modem and not an internal ISDN card.

10.4.3.2. Choose The ISDN Card



You are here presented the list of the ISDN cards detected on your machine.

Simply select the card you wish to use for your Internet connection and go on to the next step. If your card is not listed, click on "Configure an internal card manually".

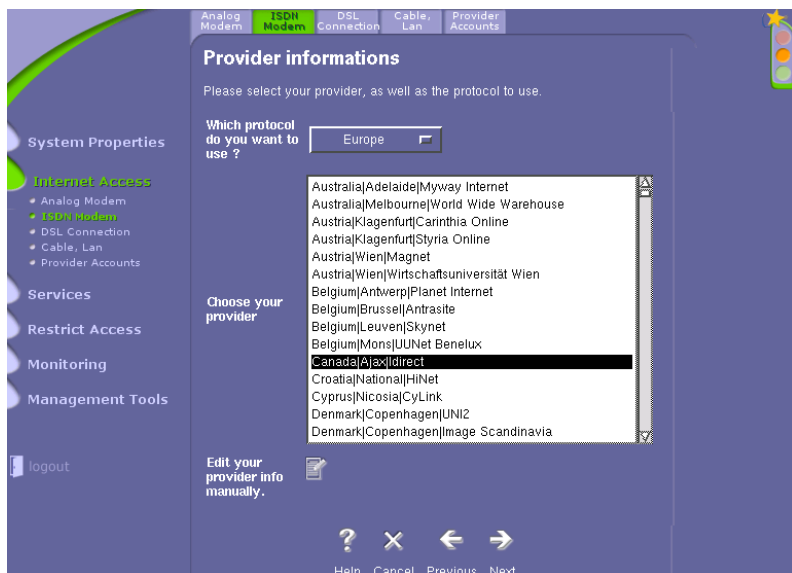
10.4.3.3. Choose The ISDN Card Model



Simply select the name of your card model on the list of suggested models, and go on to the next step.

If your card model is not listed, find out which model is compatible with it in the documentation you were given with your card.

10.4.3.4. Choose The Provider And Protocol For The ISDN Access



You are here presented with an extensive list of the providers existing around the world. If yours is absent, you will need to configure it manually.

You first need to indicate which protocol to use, this depending on your location:

- Europe
- Rest of the world

Then, you need to configure your provider in one of the following way:

- Select your provider: find your provider in the list which is organized by country, city and finally names of providers. If yours is there, great! just select it and go on to the next step.
- Edit your provider's info manually: If your provider's name did not appear in the above listing, select that icon.

10.4.3.5. ISDN Access Configuration

This form lists all the parameters required to configure an ISDN connection to the Internet. Make sure you have all the necessary parameters or inquire with your ISP.

If your provider is listed, simply fill in the blank fields.

Your ISDN Login	foo
Your ISDN Password	*****
Your ISDN Password (confirm)	*****

Here, carefully enter the login name and password as provided by your ISP.

Your Personal Phone Number	01.40.41.42.43
-----------------------------------	----------------

What is required here is the number of the phone line you use to connect to the Internet through ISDN.

Provider name	My favorite ISDN provider
Provider Phone Number	01.12.56.89.23

A simple string to identify first your provider and then the phone number you have to dial to connect to the ISDN service of that provider.

Provider DNS 1	123.456.789.122
Provider DNS 2	123.456.789.123

The Domain Name Servers of your ISP.

ISDN Card Description	ELSA Quickstep 1000 (PCI)
------------------------------	---------------------------

This indicates the name of the card being configured.

Dialing Mode	Automatic/manual
---------------------	------------------

Select how you will connect to the Internet:

- **Automatic:** Whenever the server receives an Internet request compatible with outgoing firewall rules, the connection will automatically be effected.
- **Manual:** This option will necessitate the intervention of the administrator to manually connect and disconnect that connection when required.

ISDN Card IRQ	12
ISDN Card I/O	0x300

If your card could not be detected, you will need to provide that information. Otherwise leave the fields unchanged.

When all fields have been filled out or left blank as needed, go on to the next step. You will be able to review all parameters and then confirm your choices. The connection will be configured immediately.

10.4.4. ADSL Connection Setup

10.4.4.1. Configure A DSL (ADSL) Connection



This is the first screen of the wizard that will guide you through the process of configuring a xDSL connection to the Internet. First of all, select the NIC used for this purpose.

In the list of suggestions, click on the name of the interface you want to use for the xDSL connection. If your specific card seems absent, try detecting it by clicking on the button "Detect".

10.4.4.2. Add Ethernet Interface



This page shows the interfaces detected on your firewall system.

Simply select the name (ETHx) of the proper card.

10.4.4.3. Configuration Of The ADSL Protocol Type

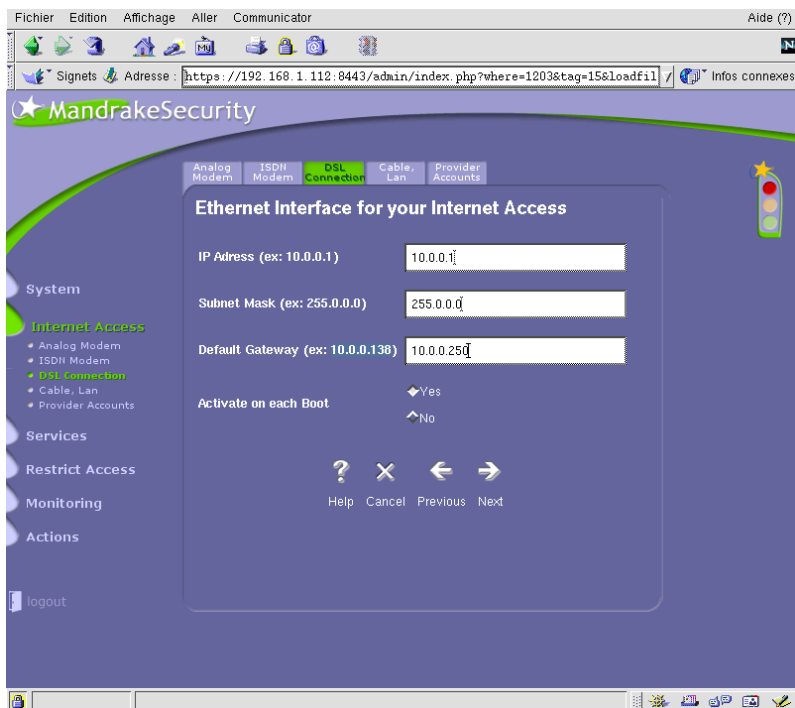


Here, select the specific protocol used by your Internet Service provider (ISP).

Choose the appropriate protocol in the pop-down list. If in doubts, ask your ISP.

- Point-to-Point Tunnelling Protocol (PPTP)
- Point-to-Point Over Ethernet (PPPOE)
- Dynamic Host Configuration Protocol (DHCP)

10.4.4.4. Ethernet Interface Configuration For Your xDSL Access



What are defined here are the parameters of the interface card necessary to map out the parameters of your xDSL access. Most of the parameters will have been chosen or filled out with standard values already: simply verify that they correspond to your needs.

IP Address (ex: 10.0.0.1)	10.0.0.1
----------------------------------	----------

Fill out this field if you have a static IP address for that interface. Be sure it is the one assigned to you.

Subnet Mask (ex: 255.0.0.0)	255.255.255.0
------------------------------------	---------------

Fill out this field with the subnet mask of the network this interface is connected to. Make sure it is the one you have been assigned.

Default Gateway (ex: 10.0.0.138)	10.0.0.250
---	------------

This is the gateway through which your Internet requests will pass. This parameter is crucial for your firewall machine to reach the Internet. In case of an external xDSL modem using PPTp, this is simply the IP address of the modem.

Finally, you can decide whether this interface will be activated on each boot or not.

10.4.4.5. Internet Account Configuration For Your

xDSL Access

The screenshot shows the MandrakeSecurity web interface. At the top, there's a navigation bar with tabs: Analog Modem, ISDN Modem, DSL Connection (highlighted in green), Cable, Lan, and Provider Accounts. Below this is the 'Internet Account Configuration' form. On the left, a sidebar menu lists 'System', 'Internet Access' (with sub-items: Analog Modem, ISDN Modem, DSL Connection, Cable, Lan, Provider Accounts), 'Services', 'Restrict Access', 'Monitoring', and 'Actions'. The main form area has a title 'Internet Account Configuration' and a prompt 'Please fill the form'. It contains several input fields: 'Username' (filled with 'foo'), 'Password' (masked with '*****'), 'Password (confirm)' (masked with '*****'), 'Provider Name (ex: provider.net)' (filled with 'My favorite ADSL provider'), 'Provider DNS 1 (ex: 198.41.0.4)' (filled with '123.456.789.123'), and 'Provider DNS 2' (filled with '123.456.789.123'). At the bottom of the form are navigation buttons: '?', 'X', '<', and '>', with labels 'Help', 'Cancel', 'Previous', and 'Next' respectively. A 'logout' link is visible in the bottom left corner of the sidebar.

To be authenticated as a user by your provider, you need to give out your account information. The necessary parameters should have been provided by your ISP.

Username	foo
Password	*****
Password (confirm)	*****

Carefully enter the login name and password provided by your ISP.

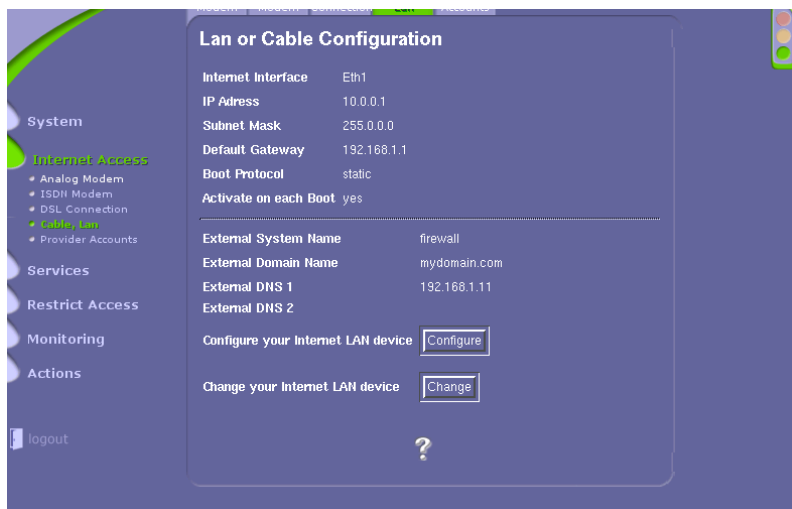
Provider name	My favorite ADSL provider
Provider DNS 1	123.456.789.122
Provider DNS 2	123.456.789.123

A simple string which first identifies your provider and then the Domain Name Servers of your ISP.

Once all fields are filled out, go on to next step. You will have the opportunity to review all parameters before confirming your choices. The connection will be configured immediately.

10.4.5. Cable/LAN Connection Setup

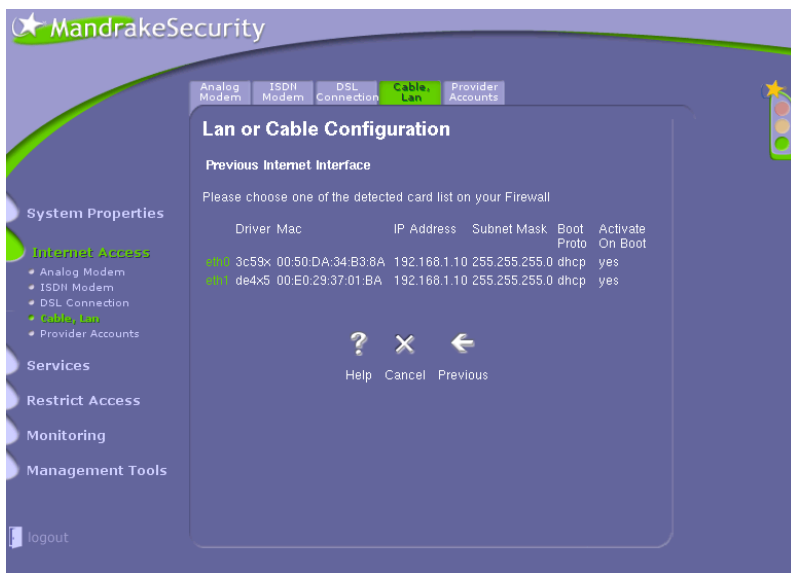
10.4.5.1. Configure A Cable Or LAN Connection



This screen appears when you have previously configured that type of connection to the Internet. It sums up the current configuration.

- Click on the "Change" button if you wish to use another NIC for use with that Internet access.
- Select the "Configure" button if you wish to reconfigure the selected NIC.

10.4.5.2. Configure A Cable Or LAN Connection



This screen is the first of the wizard that will guide you through the

process of configuring a Cable/LAN connection to the Internet. Those two types of connections are basically identical. This is the reason why they are treated together. First, select the NIC used for this purpose.

Select the name of the interface you wish to use for the Cable/LAN connection in the list of suggestions.

10.4.5.3. Ethernet Interface Configuration For Your Internet Access

Ethernet interface for your Internet Access

Internet Interface: eth0

IP Address: 192.168.1.10

Subnet Mask: 255.255.255.0

Default Gateway: 192.168.1.1

Boot Protocol: dhcp

Activate on each Boot: Yes

External System Name: firewall

External Domain Name: mydomain.com

External DNS 1: 192.168.1.11

External DNS 2:

? X ← →

What you will define here are the parameters of the interface card necessary to map the needs of your Cable/LAN access. Most of the parameters will already have been selected and the fields filled out with standard values. Start by checking they are correct.

IP Address	10.0.0.1
-------------------	----------

Fill this field out if you have a static IP address for that interface. Make sure it is the one you have been assigned.

Subnet Mask	255.255.255.0
--------------------	---------------

Fill this field with the subnet mask corresponding to the network to which this interface is connected. Make sure it is the one you have been assigned.

Default Gateway	10.0.0.250
------------------------	------------

This is the gateway through which your Internet requests will pass. This parameter is crucial to enable your firewall machine to reach the Internet.

Then, you will have to indicate which boot protocol is to be used when this interface is initialized. This depends on the protocol used by your ISP. Select one of the following:

- static
- dhcp
- bootp

Finally, you can decide whether or not to automatically activate this interface on each boot.

Then comes the configuration of your host as a member of the Internet.

External System Name	www.myco.org
-----------------------------	--------------

External Domain Name	myco.org
-----------------------------	----------

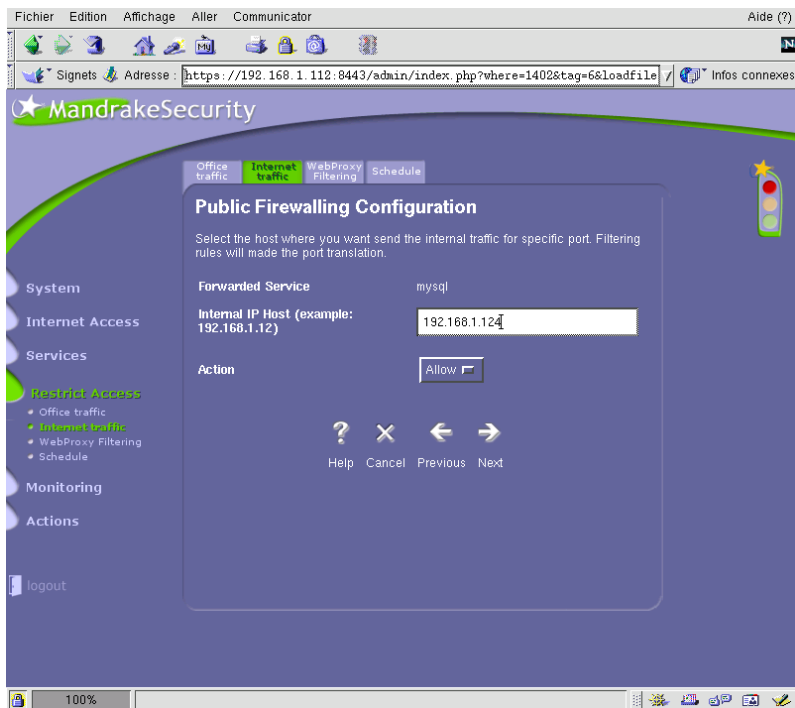
Fill out those fields with the external identification of your firewall machine.

External DNS 1	123.456.789.122
External DNS 2	123.456.789.123

Those DNS IPs generally corresponds to the Domain Name Servers of your ISP.

10.4.6. Public Firewalling Configuration:

Modifications



This form allows you to change the behaviour of the Firewall for any specific service: possible forwarding and accept policy.

Forwarded Service	mysql
--------------------------	-------

This reminds you which service is presently being configured.

Internal IP Host	192.168.1.112
-------------------------	---------------

Provide this information to forward requests on that service to another host of your internal network. Otherwise, the Firewall machine will try and handle the requests on this service on its own.

Action	Allow
---------------	-------

Select either "Allow" or "Deny", whether you want to enable incoming requests for this service or not.

When all modifications have been made, go on to the next page, the General Public Firewalling Configuration Page.

10.5. Access Restrictions



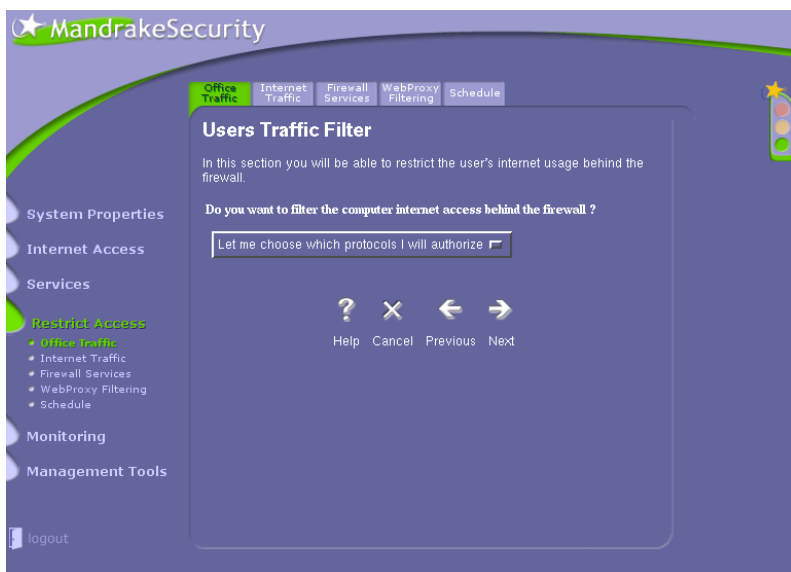
This page allows you to activate or de-activate the main access restrictions provided with MandrakeSecurity: Incoming and outgoing traffic, logs.

Select "on" if you want to activate the following features, "off" otherwise.

- N.A.T: (Network Address Translation) activating this feature will enable internal users to access the Internet through the firewall.
- Firewall Rules: Activate this to filter the traffic passing through your firewall. If you previously chose "off", your internal network will no longer be protected from attacks directed from the Internet!

You can now review all parameters and then, confirm your choice.

10.5.1. Users Traffic Filter



This option will enable you to either restrict your local users' access to the Internet or leave it somewhat or completely unrestricted.

- Leave all protocols open: no restriction on the protocols local users are authorized to use with the Internet.
- Let me choose which protocols I will authorize: later, you will be asked to decide whether any particular protocol is accessible or re-

stricted, e.g. choose to allow HTTP access while forbidding access-
es to IRC and FTP.

Check the boxes of your choice and go on to the next step.

10.5.1.1. Internal Traffic Firewalling Configuration



This step will enable you to choose what basic level of filtering you want for outgoing traffic.

Simply choose a level between the following two extremes:

- Disable: all protocols asked for by internal users are allowed,
- Full: No traffic is allowed from the local network to the Internet.

When this is done, you can choose various between three configuration levels:

- Preset only: depending on the firewalling level previously chosen, the services allowed or not are already presetted. You will be able to review them on next step.
- Predefined: You will be able on next sep to choose various types of services which will lead to the opening of the corresponding ports on the firewall.
- Custom: reserved for experts, allows to precisely tell the firewall which ports for which protocols have to be opened.

Make your selection and go on to the next step.

10.5.1.2. Predefined Internal Traffic Firewalling Configuration



Here you can choose what type of traffic is allowed for local users.

Check the boxes of your choice according to which protocols will be authorized for your local users. Note that one checkbox may in fact correspond to various similar protocols.

Make your choices and go on to the next step. To have a closer control on restrictions, click on the button for "Custom" configurations.

10.5.1.3. Custom Internal Traffic Firewalling Configuration



What we will do here is choose the precise protocols that local users will be allowed to use.

You will need to list precisely which services (in accordance with the associated protocol) will be enabled for use by local users when they access content on the Internet. The list of existing ports is available in the file `/etc/services`. Either specify the port number or the generic protocol name. To enable all of the services from a specific protocol, just insert "all" in the particular field. You can also specify a port range like that: "2048-300"; or "2048-" for all ports above 2048.

Enter TCP port to be opened :	http,https,pop-3,imap,ssh
Enter UDP port to be opened :	all

Make your selection and go on to the next step. If you are lost or if you do not know the names of the ports, click on the "Pre-defined" configuration button.

10.5.2. Public Firewalling Configuration



This page lists the services which are currently allowed to come through the firewall. It will also enable you to add new services or reconfigure current ones.

You are shown two lists: one for TCP services, the other for UDP services. The following is an example of allowed TCP services:

```
forward    action
ssh --      Allow  Modify  Suppress
www 127.0.0.1 Allow  Modify  Suppress
```


- forward: Indicates the possible internal machine to which the requests to this service will be forwarded.
- action: "Allow" or "Deny", according to the policy you will choose for the particular incoming service.
- Modify: Select to modify the behaviour of the firewall whenever requests are made to that service: port forwarding or accept policy.
- Suppress: Select to remove an incoming service permanently, thus blocking all further requests to that service.



To add a new service, click on the "Add" button:

When you have gone through all the possible modifications on this page, go on to the next step: review choices and apply.

10.5.2.1. Public Traffic Firewalling Configuration

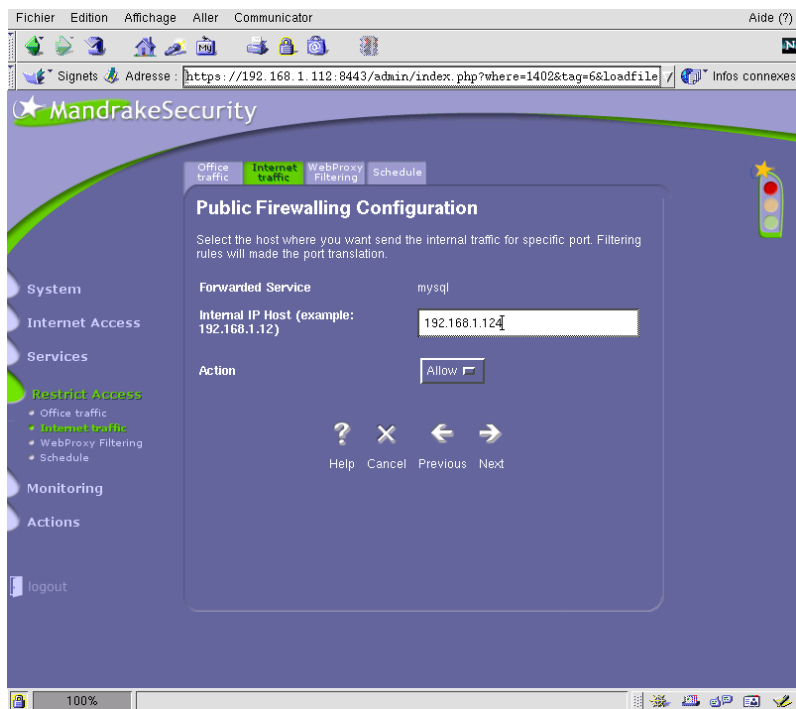


This form will enable you to define any new incoming service with the following options: forwarding and accept policy.

- Select the service you wish to enable: make your selection from the predefined TCP services list, or add any customized service in the "Port" field and choose the corresponding protocol in the "Protocol" pop-down menu.
- To forward requests made on this particular service to a specialized host of the internal network, enter the latter's address in the "IP Host" field. If you leave this field blank, the Firewall will try and manage incoming requests by itself.
- Go on to the next step: the General Public Firewalling Configuration page.

10.5.2.2. Public Firewalling Configuration:

Modifications



This page allows you to make changes to the configuration associated with a particular incoming service.

Forwarded Service	mysql
--------------------------	-------

This first field indicates the service which is being configured.

Internal IP Host	192.168.1.124
-------------------------	---------------

The internal host to which requests to this particular service will be forwarded.

Action	Allow
---------------	-------

Select "Allow" or "Deny" whether or not you want to accept incoming requests of that service.

When all your modifications have been entered, go on to the next step, the General Public Firewalling Configuration Page.

10.5.3. Custom Internal Traffic Firewalling

Configuration



What we will do here is choose the precise protocols that local users will be allowed to use.

You will need to list precisely which services (in accordance with the associated protocol) will be enabled for use by local users when they access content on the Internet. The list of existing ports is available in the file `/etc/services`. Either specify the port number or the generic protocol name. To enable all of the services from a specific protocol, just insert "all" in the particular field. You can also specify a port range like that: "2048-300"; or "2048-" for all ports above 2048.

Enter TCP port to be opened :	http,https,pop-3,imap,ssh
--------------------------------------	---------------------------

Enter UDP port to be opened :	all
-------------------------------	-----

Make your selection and go on to the next step. If you are lost or if you do not know the names of the ports, click on the "Pre-defined" configuration button.

10.5.4. WebProxy Filtering URLs



You have activated Proxy-Guard and this page will enable you to configure it. This is the first screen of the wizard. It will make suggestions to configure the various aspects of the filtering.

First choose the section to be configured in the pop-down list and go on to the next step:

- **Privileged source IPs:** enter the IPs of the privileged machines of your local network. Those machines will be freed from any restrictions imposed by the filter to other hosts.
- **Banned source IPs:** designates those machines which have no authorization whatsoever to use the proxy.
- **Allowed Lan Source:** designates the sub-network authorized to access the proxy services.
- **Banned sites:** Enter the URLs or whole domains for which all access should be blocked.
- **Advertising:** Enter the URLs or whole domains of advertising sites. The images proceeding from those sites will not be forwarded to the clients.
- **Time Restriction:** Allows you to define the connection schedule, i.e. when people are allowed to connect or not.


Whenever you are finished configuring one of those preceding sections, you will be brought back to this page. When you have updated all sections, remember to restart the proxy server Squid by clicking on the "Restart Squid" button.

10.5.4.1. Privileged IPs




This form will enable you to add or remove IPs of the privileged machines of your local network. Those machines will be freed from any restrictions imposed by the filter to other hosts.

Enter a new privileged IP address	192.168.1.111
--	---------------

In the field, enter the full IP address of the privileged host. Then click on the "Add" button: . The IP will appear in the list at the bottom of the page.

To suppress an IP's privileges on the list, simply select it and click on

the "Suppress" button: .


When you have gone through that list, go on to the next step. This will bring you back to the main WebProxy Filtering Page.


10.5.4.2. Banned IPs



This form will enable you to add or remove the IPs of those machines not authorized to use the proxy at all. This means that if, from the Local network, there is no other gateway to the Internet, these machine users will not be able to browse the Web.

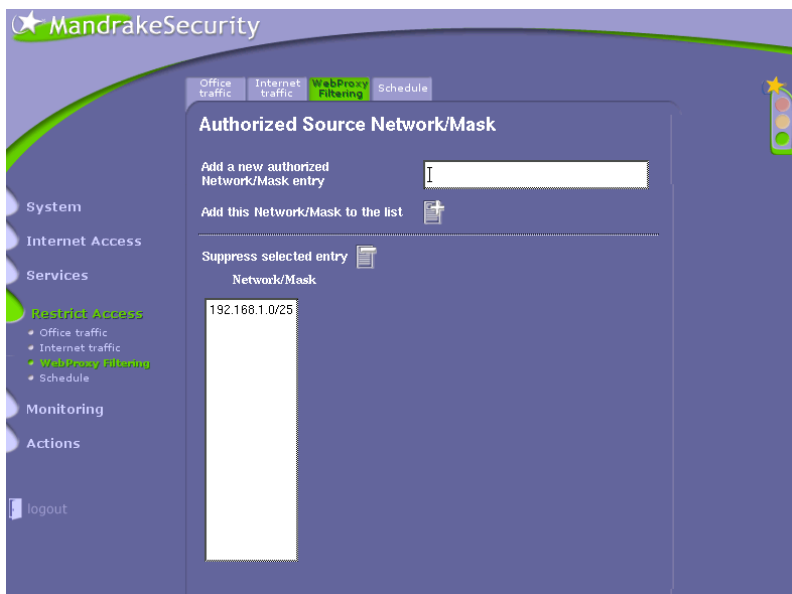
Enter a new banned IP address	192.168.1.110
--------------------------------------	---------------

In the field, enter the full IP address of the banned host. Then click on the "Add" button: . The IP will appear in the list at the bottom of the page.

If you wish to suppress an IP from the list, simply select it and click on the "Suppress" button: .


When you have gone through that list, go on to the next step. This will bring you back to the main WebProxy Filtering Page.

10.5.4.3. Authorized Source Network/Mask



This form will enable you to specify which sub-networks are allowed access to the proxy services. If different classes of machines are to be specified on your network (one for people authorized access to the Web, the other for unauthorized people), you will have to create a sub-network for those machines who are authorized and assign Ips in accordance with the authorization status the particular computer has been granted.

Add a new authorized Network/Mask entry	192.168.1.0/25
--	----------------

In the field, enter the Network/Mask IP address of the sub-network (the example shown designates the IP range from 192.168.1.1 to 192.168.1.128). Then click on the "Add" button: . The address will appear in the list at the bottom of the page.

If you wish to suppress an IP range from the list, simply select it and click on the "Suppress" button: .

When you have gone through list, go on to the next step which will bring you back to the main WebProxy Filtering Page.

10.5.4.4. Banned Sites



This form offers you three different ways to filter the pages viewed within the local network. These three types of filtering depend on those pages' URLs.

New Keyword	microsoft
--------------------	-----------

All URLs containing this word will be blocked.

New Banned Domain	msn.com
--------------------------	---------

All pages dependent on a server whose name ends with this domain

will be blocked. In our example: "http://eshop.msn.com/category.asp?catId=212" will not be displayed.

New banned URL	www.XXX.com/index_ns.html
-----------------------	---------------------------

This specific URL will not be displayed. In our example: "http://www.XXX.com/ad" won't be discarded.

Then come the lists for the three categories. You can select an item on those lists and delete it by selecting it and clicking the "Suppress"

button: .

When you have gone through those lists, go on to the next step, this will bring you back to the main WebProxy Filtering Page.

10.5.4.5. Advertising Domains/URLs



How can we get rid of those bothersome ads on our favorite web sites ? Let's examine two examples: freshmeat.net and yahoo.com.

On the freshmeat.net web site: right click on the advertising picture and then on copy image location at the bottom of the menu. Then go to the New Banned Advertising Domain section, click on the middle mouse button or on both of the mouse buttons, simultaneously. Erase the last part of the URL and you will get ads.freshmeat.net. Now add this domain to the list.

New Banned Advertising Domain	doubleclick.net
--------------------------------------	-----------------

On the yahoo.com web site: right click on the advertising picture and then on copy image location at the bottom of the menu. Then go to the New Banned Advertising URL section and click on the middle mouse button or on both of the mouse buttons, simultaneously, to paste the information. Erase the last part of the URL and you will get us.a1.yimg.com/us.yimg.com/a/pr/promo/anchor. Now add this URL to the list. If you click on your browser's "Refresh" button several times and copy the image location, you will get several different URLs : us.a1.yimg.com/us.yimg.com/a/ya/promo/anchor, us.a1.yimg.com/us.yimg.com/a/an/promo/anchor, us.a1.yimg.com/us.yimg.com/a/ us.a1.yimg.com/us.yimg.com/a/ya/yahoo_auctions, and so on...

New banned URL	www.ads.com/bigad.jpg
-----------------------	-----------------------

The ad corresponding to this particular URL will not be displayed.

The lists for the two categories will then appear. Select any item in those lists and delete it by selecting it and clicking the "Suppress"

button: .

When you have gone through those lists, go on to the next step. This will get you back to the main WebProxy filtering page.

10.5.4.6. Time Restriction

Time Restriction

Do you want to:

Define the authorized time zones. Note that each day is split into two zones: AM and PM.
Example: Sun AM 09h-12h and Sun PM 13h-22h

Sun AM	<input type="text" value="09:30-12:00"/>	Sun PM	<input type="text" value="13:00-19:00"/>
Mon AM	<input type="text" value="09:00-12:00"/>	Mon PM	<input type="text" value="13:00-19:00"/>
Tue AM	<input type="text" value="09:00-11:00"/>	Tue PM	<input type="text" value="12:00-19:00"/>
Wed AM	<input type="text" value="09:00-12:00"/>	Wed PM	<input type="text" value="12:00-18:00"/>
Thu AM	<input type="text" value="09:00-13:00"/>	Thu PM	<input type="text" value="13:00-18:00"/>
Fri AM	<input type="text" value="09:00-12:00"/>	Fri PM	<input type="text" value="13:30-18:00"/>
Sat AM	<input type="text" value="08:20-13:00"/>	Sat PM	<input type="text" value="13:30-19:00"/>

? X ← →

This form will enable you to define time periods within which access to the proxy will be permitted. Note that this does not affect privileged hosts of your local network. Outside these time periods, restricted machines will not be able to browse the Web.

You first need to choose to enable or disable this feature. If you enable it, you will have to define the time periods in question: there are two periods per weekday.

Sun AM	<input type="text" value="09:00-13:00"/>
---------------	--

Make sure to strictly respect the time format as illustrated: HH:MM-

HH:MM. Modify all periods at your convenience.

When you are done with all periods, go on to the next step. You will be shown what choices you have made. Review them and go on to the next step. This will bring you back to the main WebProxy Filtering Page.

10.5.5. Time Restriction



In case you have a non permanent connection, this page will enable you to define your Internet connection schemes. For each of the three time periods defined, you will be given five options for your connection.

- Dialup Connect Office: Define the connection schemes during office hours (8:00 AM to 6:00 PM).
- Dialup Connect Outside: Define the connection schemes outside office hours (6:00 PM to 8:00 AM).
- Dialup Connect Weekend: Define the connection schemes during the weekend (Saturday, Sunday).

For each of these periods, choose one of the following policy:

- No connection: Connection is down during that period.
- Short connect times: Connections are made on demand, and the link cut out whenever requests stop. [relevant only for analog and ISDN modem type links]
- Medium connect times: Connections are made on demand, and the link is cut out a little while after requests have stopped. [irrelevant for permanent type links]
- Long connect times: Connections are made on demand, and the link is cut out much longer after requests have stopped, average connection delays are thus minimized. [irrelevant for permanent type links]
- Continuous connection: The Internet link is maintained during that period.

When you have gone through the three different time periods. The next step will show you the choices you have just made. Review them and go on to the next step. This will bring you back to the main WebProxy filtering page.



10.6. Services

This section controls the use of other services, mainly DHCP, DNS and proxy services.

10.6.1. DHCP Server



You will need to setup a DHCP server on your firewall to enable the dynamic configuration of new machines connected to your Local Area Network. When those machines are configured to use a DHCP server on bootup, they will automatically be set up with all the network parameters they need to integrate into the LAN. All you then need to do is configure the client to use a DHCP server. This feature is available on most modern operating systems.

Just click on the check-box of your choice and go on to next step.

10.6.1.1. DHCP Server Configuration



You will now need to inform the connection scripts in order to be authenticated by your ISP. They must have provided you with all the necessary information.

Interface that the DHCP Should Listen To:	eth0
--	------

This field holds the name of the Interface connected to the LAN. Only those computers which share the same sub-network with that address will get a response from the DHCP server.

Start of the IP Range:	65
End of the IP Range:	254

Those fields contain the IP addresses range allowed for client DHCP hosts. The example given is for class C sub-network. Make sure that in the range, you do not include the first IP (0 in that case) nor the last (255) which are reserved. Note that the first addresses are generally reserved for static IP hosts, while last ones are used by DHCP servers.

Default time interval (21600 = 6h)	21600
Max time interval (43200 = 12h)	43200

The assignment of an IP to a host is always limited in time. When the client does not set the needed leasing period, the server will intervene and reassign an IP to the host each "Default time interval". However, a client's request for a specific leasing period inferior to the "Max time interval" will be honored. Otherwise, an IP will be reassigned automatically after that "Max time interval".

10.6.2. Squid Proxy Server



To be able to cache HTTP and FTP requests made from inside your LAN to the Internet, you will need to set up a proxy server on your firewall. This would permit a page to be requested by two different users to be retrieved only once from the Internet, thus dramatically fastening access to this page while saving bandwidth.

Mandrake security has chosen the proxy server Squid. The latter acts as an agent, accepting requests from clients (such as browsers) and passing them on to the appropriate Internet server. It then stores a copy of the returned data in an on-disk cache.

Choose between four options before going on to the next step:

- deactivate the proxy server: If you choose not to use the proxy, requests from users will be directly forwarded to the outside.
- activate transparent proxy: Activates the proxy and configures it to act as a transparent proxy, i.e. users will not need to configure their clients to enable them to use the proxy: all requests are automatically intercepted and managed by the proxy.
- activate manual proxy: Same as previous, but client web browsers will need to be reconfigured to explicitly use the proxy server installed on your MandrakeSecurity server.
- activate manual proxy with user level authentication: Same as previous. WARNING: create accounts on the firewall linux box for the users who are authorized to connect to the net.

10.6.2.1. Proxy Main Configuration

The screenshot shows the 'Squid default parameters' configuration window. On the left is a sidebar with a tree view containing 'System Properties', 'Internet Access', 'Services' (with sub-items 'DHCP Server', 'WebProxy', and 'Properties'), 'Restrict Access', 'Monitoring', and 'Management Tools'. At the bottom of the sidebar is a 'logout' button. The main content area is titled 'Squid default parameters' and contains three input fields: 'Squid Port (the default 3328)' with the value '3328', 'Squid Cache Size (the default is 100 Mib)' with the value '100', and 'Squid Admin Email' with the value 'admin@yourdomain.com'. Below these fields is a section titled 'Filtering URLs, Domains or IPs' which contains a paragraph of text describing squidGuard. At the bottom of this section is a question 'Do you want to use squidGuard Filtering ?' with a 'No' button and a checked checkbox.

Squid default parameters

Squid Port (the default 3328)

Squid Cache Size (the default is 100 Mib)

Squid Admin Email

Filtering URLs, Domains or IPs

squidGuard is a combined filter, redirector and access controller plugin for Squid. It is: free, very flexible, extremely fast, easily installed, portable. squidGuard can be used to: limit the web access for some users to a list of accepted/well known web servers and/or URLs only, block access to some listed or blacklisted web servers and/or URLs for some users, block access to URLs matching a list of regular expressions or words for some users, enforce the use of domainnames/prohibit the use of IP address in URLs, redirect blocked URLs to an "intelligent" CGI based info page, redirect unregistered user to a registration form, redirect popular downloads like Netscape, MSIE etc. to local copies, redirect banners to an empty GIF, have different access rules based on time of day, day of the week, date etc, filter/censor/edit text inside documents.

Neither squidGuard nor Squid can be used to :

- filter/censor/edit text inside documents
- filter/censor/edit embedded scripting languages like JavaScript or VBscript inside HTML

Do you want to use squidGuard Filtering ? ☐ No ☒

The proxy parameters will here be configured. After deciding on a few common parameters, you have the option to activate the Web filtering or not.

Squid Port (we recommend 3328)	3328
---------------------------------------	------

This is the port on the firewall machine on which Squid will listen for requests. There is no need to make any changes here unless this port is to be used by another service.

Squid Cache Size (the default is 100 Mb)	100
---	-----

Here, you can control the amount of cached data the Squid can store and manage. The more users on your LAN using the proxy, the more space needed for your cache to be efficient. It may vary between 10 Mb and 10Gb or more.

The Web filtering can now be activated. This feature will enable you to deny or restrict access to certain pages on the Internet, depending on their URLs. It is useful to block access to ad banners or some types of content for children.

Select the check-box of your choice and go on to the next step.

10.6.2.2. Proxy Guard Configuration



This page simply shows current settings. Click on "Previous" to change them, or "Next" to go on to the next step.

10.6.3. Services Properties



This page lists the services present on your machine. You will be given the opportunity to enable or disable them.

```
Status
[...]
Gpm Running reload restart stop start Details
Httpd-naat Running reload restart stop start Details
Squid Running reload restart stop start Details
[...]
```

The first column of the table lists the name of the service and its present status:

- **Running:** The service is installed and accepting connections.

- Stopped: The service is installed on the firewall but is currently disabled.
- Unknown: For some reason the Interface was unable to determine the status of that service.

The parameters of these services may then be modified:

- reload: allows the reloading of that service configuration without interruption. To be used when a parameter of that service has just been modified.
- restart: stops and restarts the service.
- stop: the service will refuse further connections and terminate current ones.
- start: the service will accept further connections.
- Details: brings up another page with more information about that particular service.

Chapter 11. Configuring Masqueraded Clients

This chapter will show you how to make different operating systems use a *GNU/Linux* box with masquerading set up as a gateway to the outside world. The configuration tests on the following operating systems all proved successful :

- Apple Macintosh, with MacTCP or Open Transport;
- Commodore Amiga, with AmiTCP or AS225-stack;
- Digital VAX Stations 3520 and 3100, with UCX (TCP/IP stack for VMS);
- Digital Alpha/AXP, with Linux/Red Hat;
- IBM AIX (on RS/6000), OS/2 (including Warp 3) and OS400 (on OS/400);
- Linux (of course!): any kernel release since 1.2.x;
- Microsoft DOS (with NCSA Telnet package, partial DOS Trumpet support), Windows 3.1 (with the Netmanage Chameleon package) and Windows for Workgroup 3.11 (with TCP/IP package);
- Microsoft Windows 95, Windows 95 OSR2, Windows 98, Windows 98se;
- Microsoft Windows NT 3.51, 4.0 and 2000 (both workstation and server);
- Novell Netware 4.01 Server, with the TCP/IP service;
- SCO OpenServer (v3.2., 4.2 and 5);
- Sun Solaris 2.51, 2.6 and 7.

Let's go through the configuration of a few of them. If your system is not listed, a simple way to proceed is to "just tell the OS which computer to use as a gateway". Note that our main focus here is the

gateway side of the network: therefore, we won't touch on DNS problems, file sharing or connection schemes. Thus, for this chapter to be of any use to you, you need a well-configured local network. Refer to your system's documentation to set it up properly, paying special attention to the DNS settings.

What follows assumes that you are set up on a class C network: your different machines all have IP addresses like 192.168.0.x, with a net-mask set to 255.255.255.0, and use eth0 as the network interface. We also take for granted that: your gateway has its IP address set as 192.168.0.1; and that your computers can each “talk” to the gateway (test the latter with the ping command or its equivalent in your environment).

11.1. Linux Box

There are (at least) three ways to go about this.

11.1.1. On-the-fly Configuration

This is probably the fastest way to proceed. However, when you next restart your network layer or your whole system, any configuration change you will have made will have disappeared!

If eth0 is the network interface through which you access the gateway, (as root) issue the simple command:

```
route add default gw 192.168.0.1 eth0
```

That's it! If the gateway is properly configured and connected to the Internet, the whole world is now within your reach thanks to your favorite web browser.

11.1.2. Permanent, Manual Configuration

To maintain the configuration each time the system is shut down and restarted, we need to edit a config file. Its name is `/etc/sysconfig/network` on a **Linux-Mandrake** machine (it may be different on yours). Open it with your usual text editor, then add the following lines:

```
GATEWAYDEV="eth0" GATEWAY="192.168.0.1"
```

You may now restart your network layer with: `/etc/rc.d/init.d/network restart`

11.1.3. Permanent, Automatic Configuration

To install the configuration automatically, we need *DrakNet*, a great program to be found in the *DrakConf* panel. When you click on Internet and Network Configuration, this window will appear:

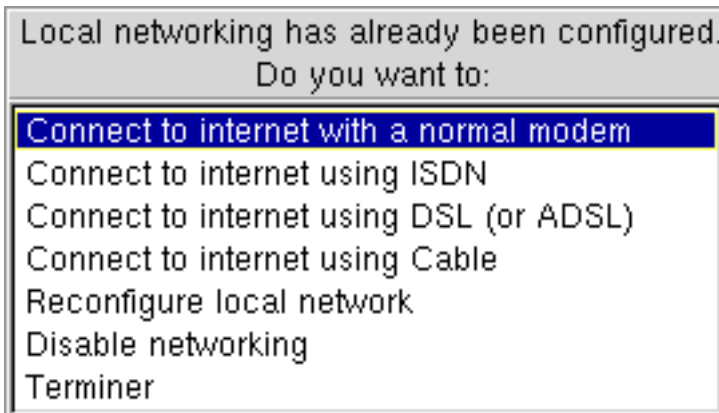


Figure 11-1. Reconfiguring a Network with Draknet

Please enter the IP configuration for this machine. Each item should be entered as an IP address in dotted-decimal notation (for example, 1.2.3.4).	
IP address	192.168.0.186
Netmask	255.255.255.0
Automatic IP	<input type="checkbox"/> (bootp/dhcp)
<div>Ok</div> <div>Cancel</div>	

Figure 11-2. Reconfiguring the Local Network with Draknet

Then click on **Reconfigure Local Network**: a dialog box appears asking for the static address of that *GNU/Linux* host and the associated netmask, or suggesting the use of a DHCP/bootp configuration protocol. If your firewall is configured to act as a DHCP server, leave the fields blank, click the appropriate check box and then **OK**.

Please enter your host name. Your host name should be a fully-qualified host name, such as "mybox.mylab.myco.com". You may also enter the IP address of the gateway if you have one	
Host name	localhost.localdomain
DNS server	192.168.1.11
Gateway	192.168.1.1
<div>Ok</div> <div>Cancel</div>	

Figure 11-3. Setting up the Gateway with Draknet

Once this is done, answer the various questions. If the automatic configuration is not set up, follow the instructions which tells you to enter the gateway. Then just fill it in with: 192.168.0.1. Confirm when asked to. That's it! Your network is properly configured and ready to run.

11.2. Windows 95 or Windows 98 Box



Figure 11-4. The Network Icon under Windows 95

Start by going in the Control Panel: **Start+Settings→Control panel** and find the network icon as shown. Double-click on it: the network configuration panel comes up.

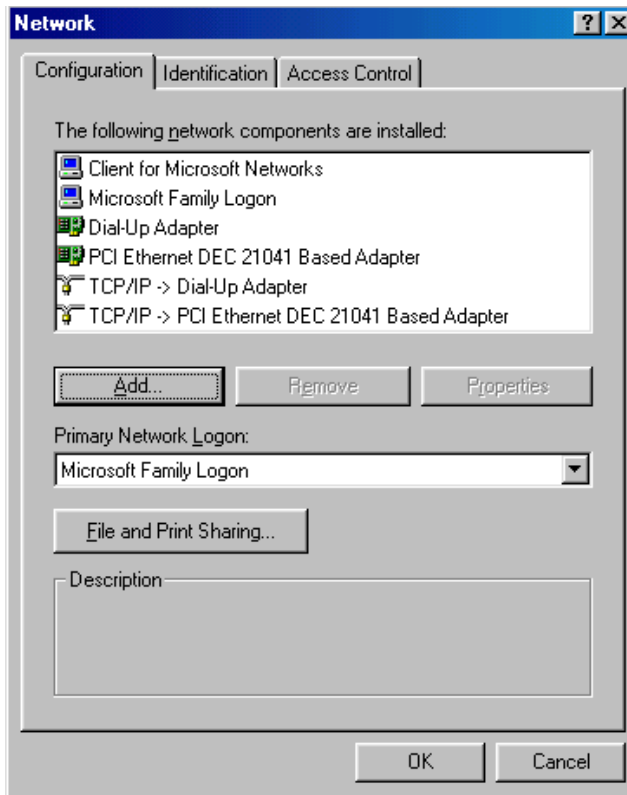


Figure 11-5. The Network Configuration Panel under Windows 95

In the list which appears, you should find a protocol named TCP/IP . If not, you will have to refer to your system documentation to find out how to install it. If it is already there, select it and click on “Properties”.

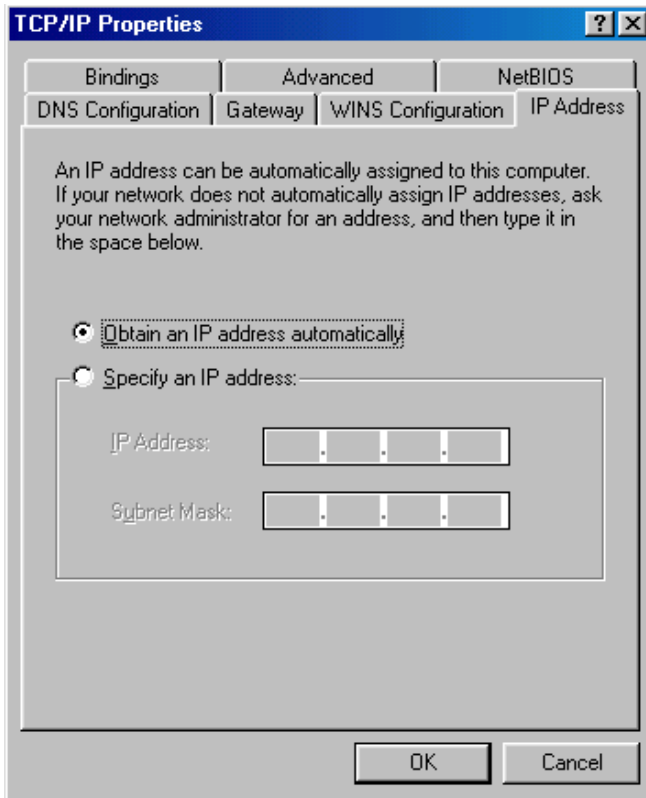


Figure 11-6. The TCP/IP Configuration Panel under Windows 95

This window will enable you to set up your TCP/IP parameters. Your system administrator will tell you if you have a static IP address or if you are using DHCP (automatic IP address). Click on the Gateway tab.

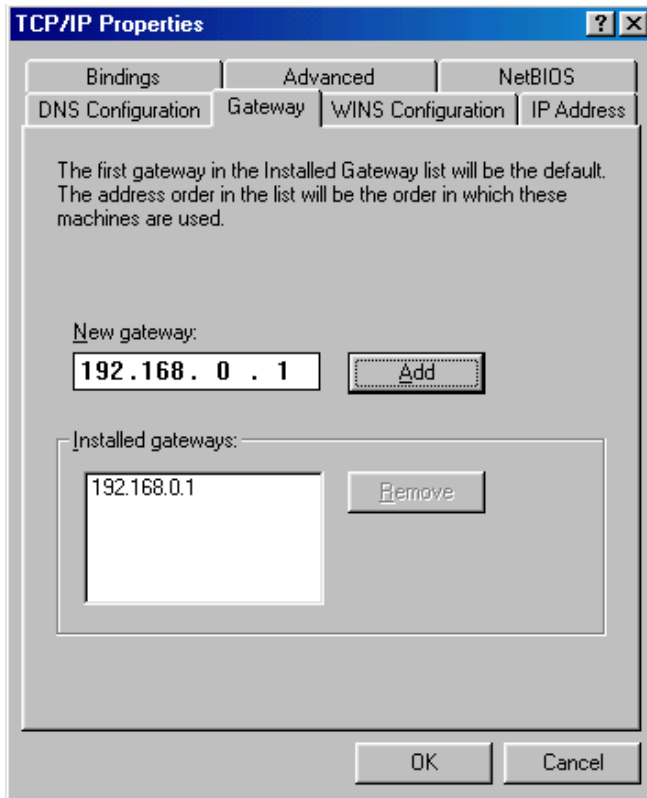


Figure 11-7. The Gateway Configuration Panel under Windows 95

The rest is child's play! Fill in the blanks with your gateway's IP address (192.168.0.1, in our example). Click the Add then the OK buttons.

You will need to reboot your computer, of course. Once this is done, find out if you can reach the rest of the world.

11.3. Windows NT or Windows 2000 Box

To configure these OSs, follow these simple steps:

1. Go to: Control Panel+Network→Protocol.

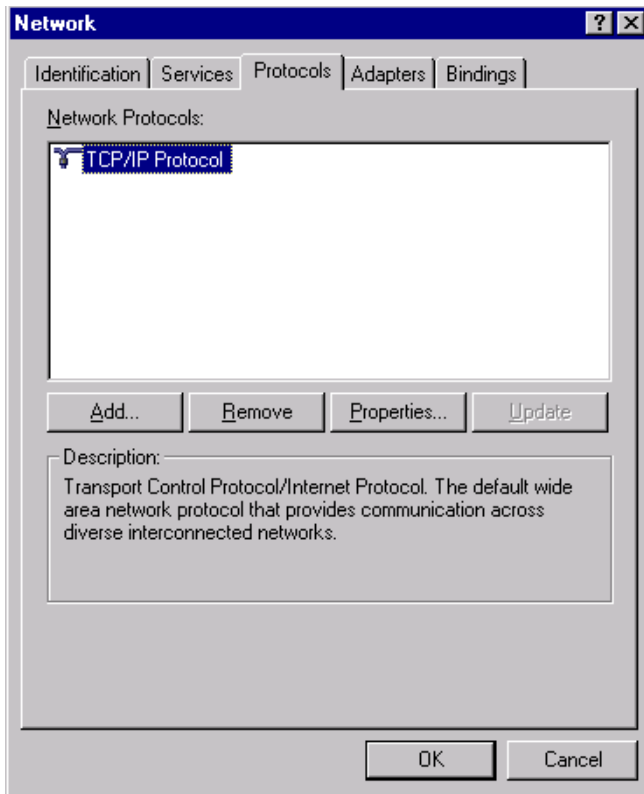


Figure 11-8. The Protocol Configuration Panel under Windows NT

2. First select the TCP/IP Protocol in the list of network protocols. Then, click the Properties button and select the network

card connected to the local network (figure 11-9). In this example, we show a configuration with the DHCP server activated on the *MandrakeSecurity* server: the Obtain an IP address from a DHCP server option is checked.

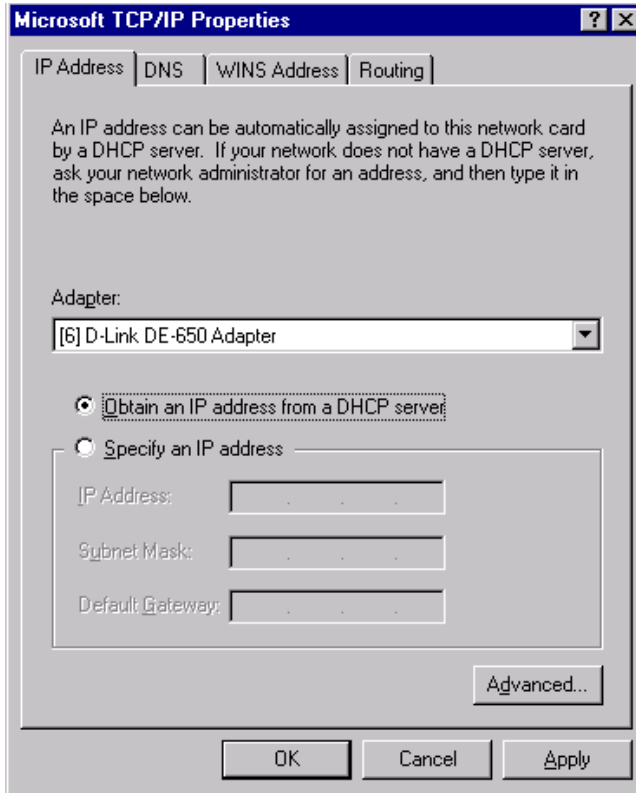


Figure 11-9. The Network Software Panel under Windows NT

If this is your case, you just need to confirm all those choices and reboot. Otherwise, follow these steps.

3. If you don't have a DHCP server, you need to manually set all parameters. Begin by checking the Specify an IP address option (figure 11-10).

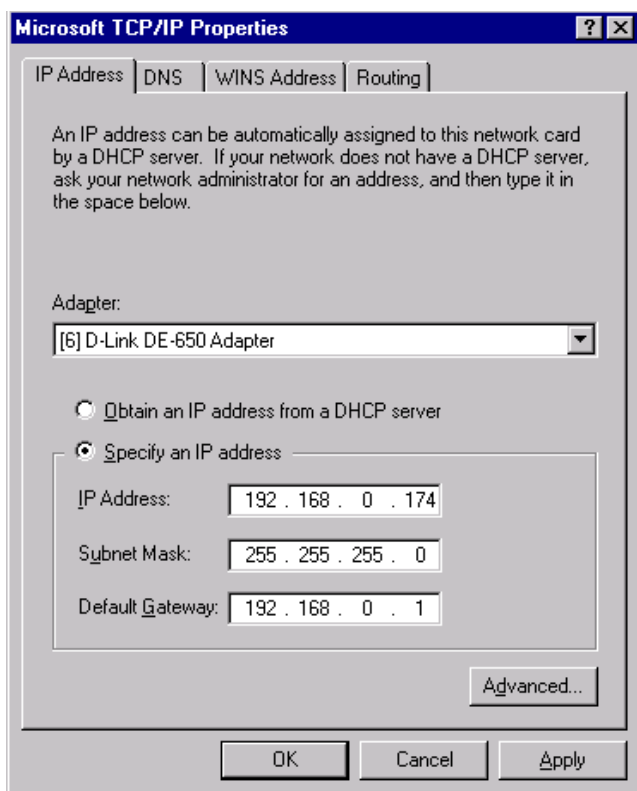


Figure 11-10. The TCP/IP Configuration Panel under Windows NT

Select the appropriate adapter, the IP address should already be correct.

4. Simply fill **Default Gateway** to 192.168.0.1 (the address of the Linux box sharing the connection in our example).
5. Finally, you will need to specify the DNS servers you use in the DNS tab as shown figure 11-11.

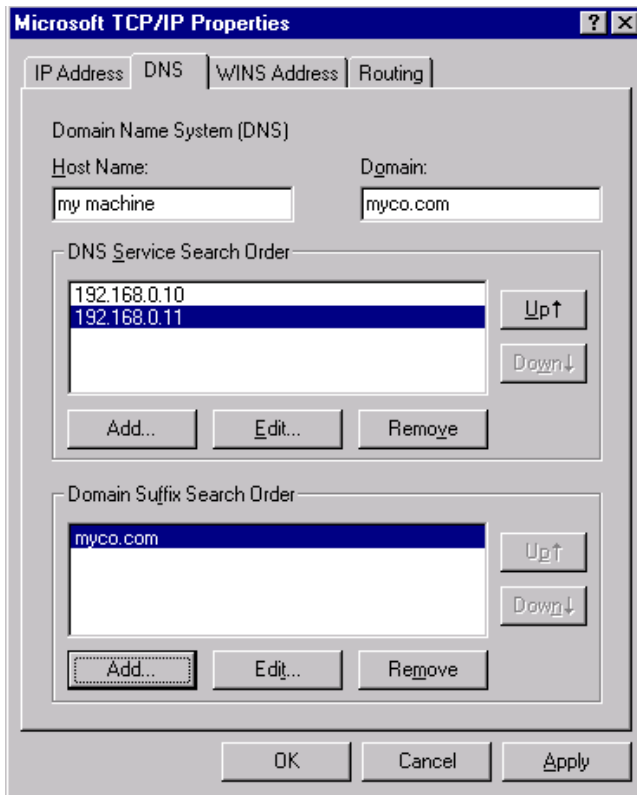


Figure 11-11. The DNS Configuration Panel under Windows NT

You must also provide a host name and associated domain name.

Warning! Unless you know exactly what you're doing, proceed with utmost care with the following steps:

- leave the Automatic DHCP configuration field blank unless you have a DHCP server somewhere on your network;
- leave all the WINS Server fields blank as well unless you have one or more WINS server(s);

- do not place a check mark in the **Enable IP Forwardings** field unless your NT machine is used for routing and, once again, you know perfectly what you are doing;
- please disable DNS for Windows Name Resolution and Enable LMHOSTS lookup.

Click on OK in the dialog boxes which then appear and restart your computer to test the configuration.

11.4. DOS Box Using NCSA Telnet Package

In the directory which hosts the NCSA package, you will find a file called `config.tel`. Edit this file with your favorite editor and add lines like these:

```
name=default
host=yourlinuxhostname
hostip=192.168.0.1
gateway=1
```

Of course, write the name of your Linux box instead of `yourlinux-hostname` and change the gateway address given here (`192.168.0.1`) which is only an example.

Now save the file, try to telnet your Linux box, then a machine somewhere out there...

11.5. Windows for Workgroup 3.11

You should already have the TCP/IP 32b package installed. Go to the menu entry: **Main+Windows Setup+Network Setup→Drivers** and select **Microsoft TCP/IP-32 3.11b** in the **Network Drivers** section, then click **Setup**.

From here, the procedure is quite similar to the one described in the Windows NT section.

11.6. MacOS Box

11.6.1. MacOS 8/9

First of all, you need to open the TCP/IP control panel as shown below in the Apple menu.

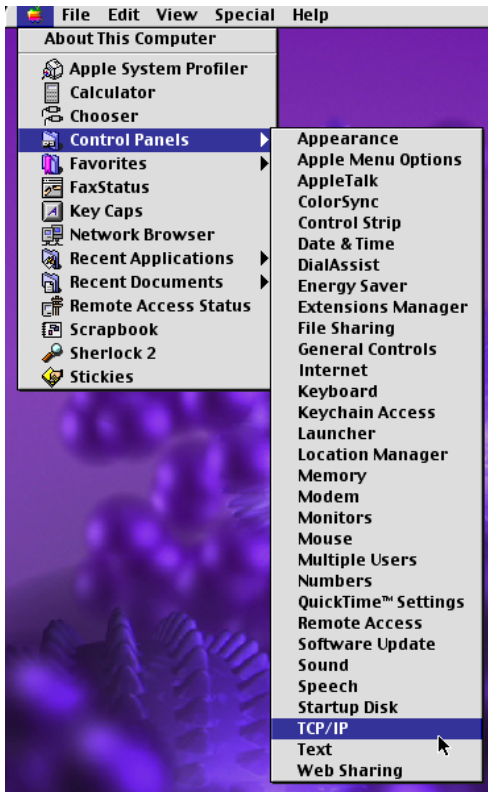


Figure 11-12. Accessing the TCP/IP Control Panel

11.6.1.1. With an Automatic DHCP Configuration

If you configured your firewall to be a DHCP server, follow this very procedure, otherwise go to the next section.

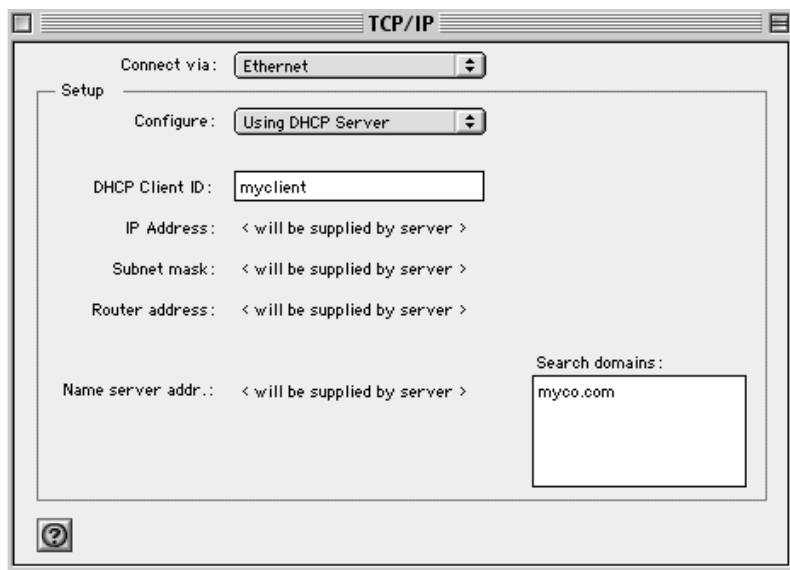


Figure 11-13. Automatic Configuration of Internet Access for Mac OS

In the dialog that appears, fill the fields as shown hereafter:

- Connect via: Ethernet
- Configure: Using DHCP server
- DHCP Client ID: 192.168.0.1

11.6.1.2. Manual Configuration

If you don't have a DHCP server on your local network, follow this procedure:

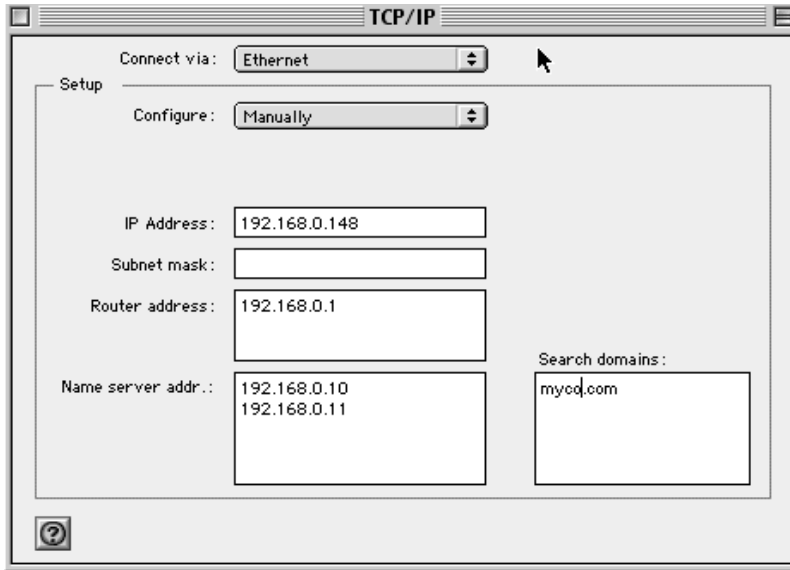


Figure 11-14. Manual Internet Access Configuration for MacOS

In the dialog that appears, fill the fields as shown here:

- Connect via: Ethernet
- Configure: Manually
- IP address: 192.168.0.248
- Subnet Mask: 255.255.255.0
- Router Address: 192.168.0.1
- Name Servers Addresses: 192.168.0.10 192.168.0.11
- Search Domain: myco.com

Note: The name servers addresses may be the internal DNS's or your ISP's.

11.6.2. If You're Running MacTCP

1. In the MacTCP Control Panel, select the Ethernet network driver (caution, it's not EtherTalk) then click the **More...** button.
2. Under **Gateway Address**, enter the address of the Linux box sharing the connection (192.168.0.1 in our example).
3. Click **OK** to save the settings. You may have to restart your system to test these settings.

11.7. OS/2 Warp Box

You should already have the TCP/IP protocol installed. If not, proceed to install it.

1. Go in **Programs**, then **TCP/IP (LAN)** and **TCP/IP Settings**.
2. Under **Routing**, choose **Add. In Type**, select **Default**.
3. Fill the **Router Address** field with the one of your Linux box sharing the Internet connection (i.e. 192.168.0.1 in our example).
4. Now close the TCP/IP control panel, answer "Yes" to all questions, then reboot your system before testing the settings.

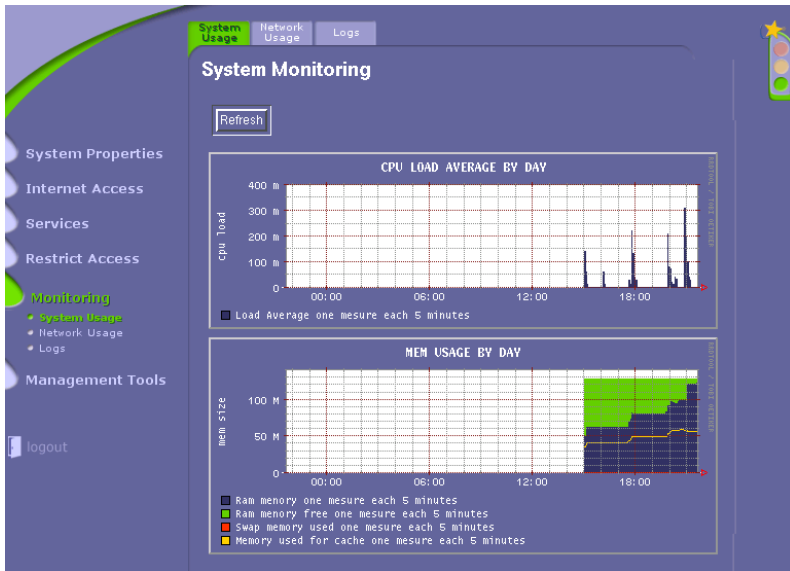
IV. Maintenance

Chapter 12. System Monitoring

In this chapter, we will explain the usage of the monitoring tools provided through the web interface. These tools will allow you to access graphical usage statistics and text logs for your system.

Note: You may have noticed a traffic light on the right of the interface. It is an overall indicator of the system's state: the light will switch to orange or red if at least one system indicator detected a dangerous or critical issue. To get more details on the problem, just click on the traffic light. It will pop-up a new window containing system information and indicators.

12.1. System Monitoring

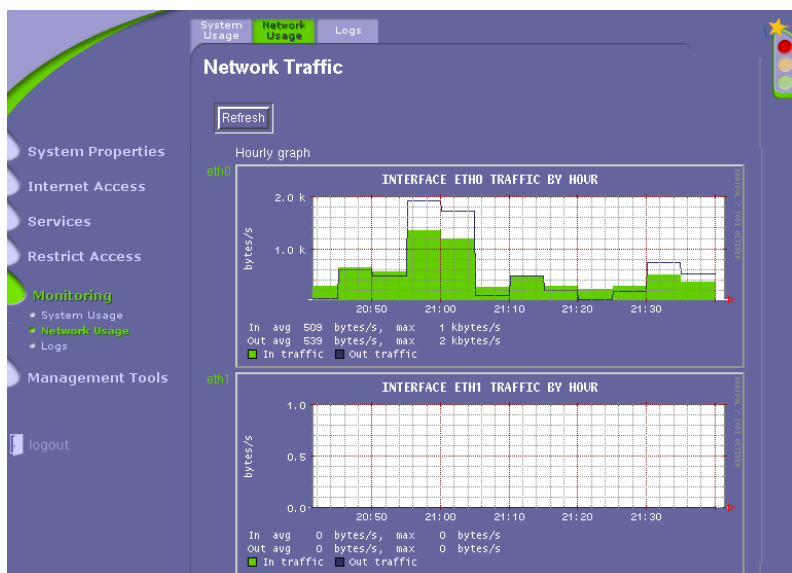


The two graphs shown here inform you about your system load. They are good indicators of the suitability of your system with its actual use.

- CPU load: representation of the CPU usage for the last 24 hours. The unit used roughly indicates the number of processes trying to access the CPU at one point. A normal load should remain below 2. Your system is a rather busy one if the load is between 2 and 5. Above 6, you should consider upgrading the hardware.
- Memory usage: indicates the usage made of your main RAM memory in megabytes. The different colors give more precise informa-

tion about the way the memory is used.

12.2. Network Traffic



Two graphs will appear to inform you about the present traffic manageable by your interfaces. They are good indicators of the suitability of your system with respect to its actual use.

This first page informs you of the traffic for all interfaces during the last hour. You can get more graphs by clicking on the name of the interface (i.e. "eth0") at the left of each graph.

12.3. Configuring and Consulting Logs

12.3.1. Audit Incoming Public Services



This form will enable you to choose which public services will be audited for traffic.

Follow these steps to add a service to the list of logged traffic:

- Determine whether the service is TCP, UDP or ICMP,
- Click on the relevant selection list, and release the mouse button on the required service.

- Click on the "Add" button:



To cancel the audit on a specific service, follow the same steps and

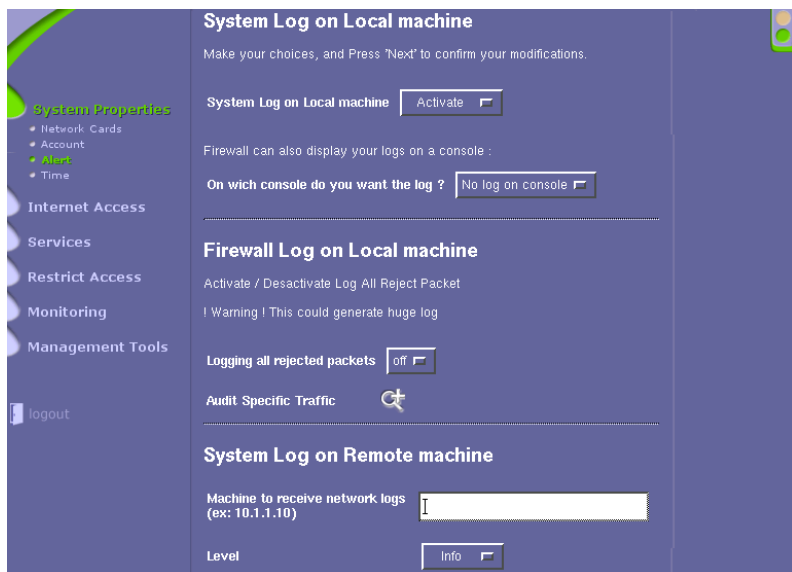
then click on the corresponding "Stop" button:



As you select or disable the items on the list of audited services, these will show up or disappear from the lists at the bottom of the page.

When you are done, go on to the next step, review your choices, "accept" them or come back to the beginning of the page.

12.3.2. System Log On Local/Remote Machines



Logs are an essential part of a security critical system like a firewall.

Not only does it give out information in real time on what is happening on the system, but it also retraces its history, e.g. when something goes wrong in the system - a crash or an intrusion - it will find out why it happened and most generally figure out a solution.

First of all, you have the choice to activate (or not) the logging system on the local machine (the firewall itself). This, of course, will only be relevant if a display is directly attached to the firewalling machine. It will be possible to control:

Level	info
--------------	------


This parameter controls the amount of info that will be displayed, from:

- Info: outputs every single message on the firewall, from normal operations to critical messages.
- Panic: outputs only those critical messages which generally lead to system failure.

Which console do you want the log on?	Console 12 (tty12)
--	--------------------

Here, you can choose the virtual console on which some selected messages will be displayed, e.g. you can switch to Console 12 by pressing the keys (Ctrl-Alt-F12).

Then you can choose whether to log Reject Packet by the firewall. If you decides so, put the switch "on" and click on the "Audit specif-

ic traffic" button: , that will lead you to a form where you can choose the exact services to be audited.

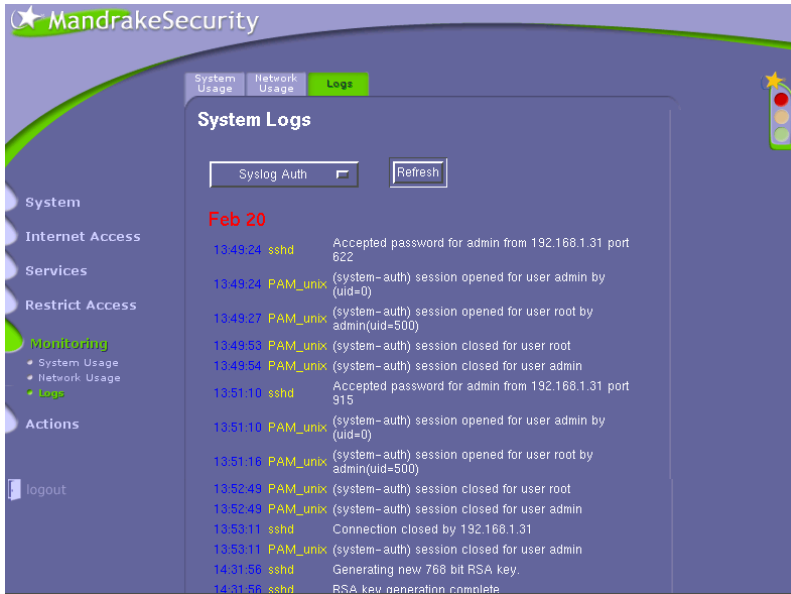
Then, a suggestion will be made to have the logs sent to a log server. This is an interesting feature for machines likely to be attacked by intruders. Having logs stored on another machine will enable you to keep a secure copy of the logs less likely to be modified by crackers wanting to erase their tracks.

Machine to receive network logs (ex: 10.1.1.10)	192.168.1.11
--	--------------

This field holds the IP address of a possible syslog server on your secured network. To enable the server to receive logs from your firewall, you need to modify the file `/etc/init.d/syslog`. Simply replace the line `"daemon syslogd -m 0"` by `"daemon syslogd -m 0 -r -l firewall"`. (provided your firewall is called "firewall")

Finally choose the amount of info that will be sent to the log server.

12.3.3. System Logs



This page allows you to look up the most important logs for the current day.

First choose the log you wish to look at in the pop-down list:

- Syslog all messages: all system messages logged by the facility.
- Filtered packets: list of logged packets as defined in the "System -> Alert" wizard.
- Syslog messages: filters all syslog messages to extract most relevant ones.

- Syslog auth: displays the messages related to connection authorizations accepted or rejected by the system.
- Squid reports: summarizes proxy use: clients, visited sites, etc. Choosing this option will display a list of reports. Clicking on a report pops up a new window with all information displayed in it.

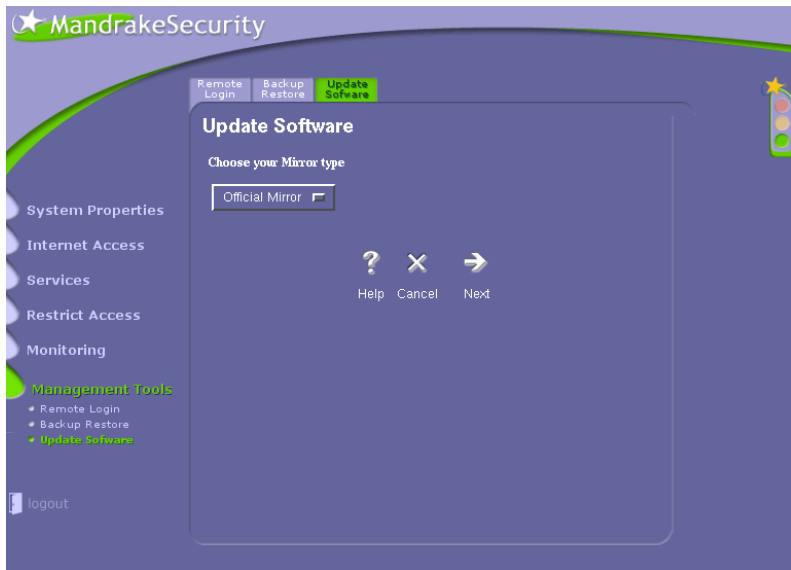
Note that some logs might not be immediately available.

Click on "Refresh" to get the latest entries.

Chapter 13. System Update

One of *MandrakeSecurity*'s very convenient features lies in the facilitated updates of old package versions. These updates are essential for security critical systems as they significantly reduce the risks associated with security holes.

13.1. Update Software



This wizard will enable you to update the packages installed on your system. For the security of your system, it is essential that you regularly check for updates of the software installed on a firewall.

You are now asked to choose between two different types of mirror host:

- **Official Mirror:** on the following page, you are given a list of all Linux-Mandrake official updates mirrors.
- **Personal Mirror:** on the next page, you will be able to manually enter the URL of the FTP site hosting the updates for your system.

13.2. Official Mirror List



This page contains the list of all Linux-Mandrake official updates mirrors around the world.

Simply choose a mirror in the list. As a rule of thumb, try selecting the one nearest your physical location to get the best transfer speed.

13.3. Packages Selection



This page contains the list of updates available for your current installation.

```
version installed_version size
bind 8.2.3-1.1mdk 8.2.2P7-1.1mdk 1782540 description
php 4.0.4pl1-1.2mdk 4.0.3pl1-1mdk 506943 description
[...]
```

What follows is the description of the fields available for each package:

- version: the version of the update available on the Mirror.
- installed_version: the version of the package currently installed on your system.
- size: the size of the package.
- description: click on this link to display a full page of informations about that package. The importance of updates will be made clear to you at this point.

Check the boxes corresponding to the packages you wish to update after reviewing the description of that update. Then go on to next step where you will be able to review your choices. After that the package will be downloaded (that may last some time...) and installed after one more confirmation..

13.4. Personal Mirror List



Here you will be able to manually enter the URL of a non official Linux-Mandrake updates mirror.

Enter your Mirror	ftp://MyMirror/updates/MandrakeSecurity
--------------------------	---

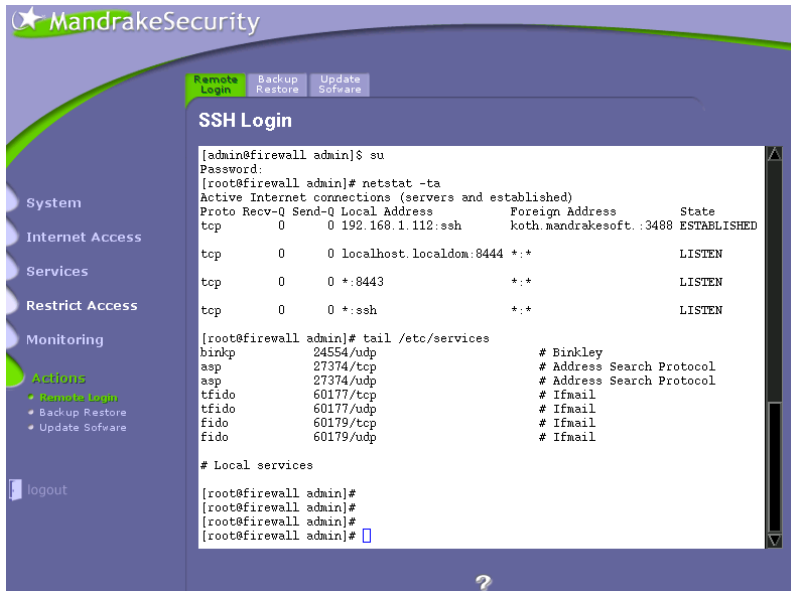
It might be a public FTP site available on the Internet or a local FTP server available on your local network.

WARNING: It is a potential security risk to install updates from a non official Linux-Mandrake updates mirror. Make sure you fully trust the location before installing such updates.

Chapter 14. Management Tools

MandrakeSecurity provides a web interface to the secure shell (ssh). Options for server configuration backup and restoration are also available.

14.1. Remote Secure login



This page contains a SSH console which will enable you to perform a secured shell connection to the firewall system.

At the prompt, enter "admin" and its password:

```
firewall login:admin
```

```
admin@firewall's password: *****
```

You will then enter The Linux shell, where you will be able to perform any administration task as if you were directly using the machine.

WARNING: the actions performed by the backend of the web interface may cause your modifications to configuration files to be lost. In fact, if again you happen to modify files through the web interface, the manual modifications will be lost. Thus, make sure to use the web interface whenever possible rather than the secure shell connection.

Finally, note that if you have a SSH client installed on your machine, you can connect to the firewalling host directly and bypass the web interface. Careful! the previous warning is then still valid.

Warning

We shall not describe the many tasks feasible through the command line. However, be warned that if you possess little or no experience with it, this chapter will be meaningless to you.

14.2. Backup And Restore



This feature backs up the entire configuration of your server enabling you to recover rapidly from a major system failure or to reconfigure a new server easily.

First of all, get to the configuration file by clicking the "Backup" button. A file containing the entire configuration file appears. It is highly recommended to make regular back-ups of this file on a floppy disk which should be stored in a secure place.

WARNING: This backup will only contain the parameters handled by the web interface. If modifications were made to some files by other means, e.g. the secure shell connection, those specific files will have to be backed up manually.

The second part restores the backed-up configuration. Click on the "Choose" button. A navigation window will be displayed where you can select the local copy of the configuration file to be restored on the firewall.

Choose The Configuration File	/mnt/floppy/ConfigurationBackup
--------------------------------------	---------------------------------

When done, click the "Upload" button. A confirmation page appears where the "Apply" option will enable you to activate the chosen configuration.

Appendix A. Where To Get Documentation

Apart from the manuals included with **Linux-Mandrake**, documentation is available from many sources. The next few pages will offer you some suggestions which you might find useful.

A.1. The Documentation Included In Linux-Mandrake

A.1.1. The Manual

This is a primary source of information on a day-to-day basis. To each command corresponds a manual page, or almost so. Plus, certain configuration files, library functions for programmers and others, also possess their own manual pages.

The content of the manual is arranged in different sections. References to these sections are made in the following manner :for example, “`man 2 open`”, “`man 5 fstab`” will respectively refer to the open page in section 2 and the `fstab` page in section 5.

To display a manual page, type `man`. Its syntax will be as follows :

```
man [options] [section] <manual page>
```

Even for `man` itself, a command is available :`man man`. Manual pages are formatted then displayed using the `less` *pager* by default. There you go! you can now browse through and quit a manual page :-)

The names of the manual page and of its relevant section appear at the top of each page. At the bottom, are given references to other pages with related subjects (in general in the **SEE ALSO** section).

You can start by consulting the pages related to the different commands covered in this manual :`man ls`, `man chmod`, etc.

If you can't find the right manual page - for example, you want to use the function `mknod` in one of your programs but you end up on the `mknod` command page -, make sure you spell out the section explicitly. In our example `:man 2 mknod`). If you have forgotten the exact section, `man -a mknod` will read through all the sections looking for pages named `mknod`.

A.1.2. info Pages

`info` pages complete the documentation included in the manual pages. The command for accessing `info` pages is `info`.

The `info` pages have a tree structure, the top of which is called `dir`. From there, you can access all of its `info` pages.

`info` may be called up in two ways :either by omitting any argument, which will place you at the very top of the tree structure, or by adding a command or a package name, which will open the relevant page, if it exists. For example :

```
info emacs
```

In the `info` pages, such a text :

```
* Buffers::
```

will indicate a link. Moving the cursor to this link (using the arrow keys) and pressing `Enter` will take you to the corresponding `info` page.

You may use the following keyboard shortcuts :

- **u** :for *Up* takes you up one level;
- **n** :for *Next* brings you to the next `info` page on the same level of the tree structure;
- **p** :for *Prev* takes you back to the previous `info` page.

A great number of commands may be listed by typing ; ?.

A.1.3. HOWTOs

HOWTOs, published by the LDP (*Linux Documentation Project*) and available in many languages, will help you with the configurations of the many aspects of your system. As long as the proper packages have been installed (the `howto-html-en` package for the English edition), *HOWTOs* will provide you with an answer to a specific question or a solution to a problem on your hard disk. The documentation is located in the directory `/usr/share/doc/HOWTO/HTML/en/`. The documents are text files in their primary form although they are also readable in HTML with a web browser, and printable with *PostScript*.

The list is quite exhaustive. Get an idea of its length by consulting the index from the main menu : Documentation→Howtos English. When met with a complex problem, start by reading the corresponding *HOWTO* (if it exists of course!). Not only will you be given a solution to your problem but you will also learn a great deal at the same time. Among others, examples of what is covered range from networking (NET-3-HOWTO), sound card configuration (Sound-HOWTO), the writing of CD media (CD-Writing-HOWTO) and NIS and NFS configuration.

An important step is to check the modification dates of the *HOWTO* documents - i.e. the publication date situated at the beginning of the document - to make sure they are up to date. Otherwise their contents may be invalid. Watch out for old *HOWTO* relating to hardware configuration especially, as *GNU/Linux* evolves very fast in that specific area. Remember also that, in the world of free software, the term “old” carries even more weight than in IT in general : free software may be considered old after being around for fifteen days!

Note: *HOWTO* are available online at the LDP (<http://linuxdoc.org/>) and likely to be slightly more up-to-date there. Have a look at the following as well : (classified by categories (<http://linuxdoc.org/>))

[//linuxdoc.org/HOWTO/HOWTO-INDEX/categories.html](http://linuxdoc.org/HOWTO/HOWTO-INDEX/categories.html))); and
FAQs (<http://linuxdoc.org/docs.html> \hyper@hashfaq).

A.1.4. The Directory `/usr/share/doc`

Some packages include their own documentation located in a subdirectory of `/usr/share/doc` and named after the specific package.

A.2. Internet

Internet information sources are widespread and websites devoted to *GNU/Linux* and its use or configuration numerous. There are however other places than websites.

Your preferred source of information about **Linux-Mandrake** should be the official web-site (<http://linux-mandrake.com>). In particular, check out the support (<http://linux-mandrake.com/en/ffreesup.php3>) section.

A.2.1. Websites Devoted to GNU/Linux

A.2.1.1. MUO

MandrakeUser.Org (MUO) is **the** data base for **Linux-Mandrake** users. With over 200 pages and growing, it is arguably the largest collection of **Linux-Mandrake** related documentation on the web. Apart from the online version of our present wonderful handbook, that is... :-).

MUO collects submissions by **Linux-Mandrake** users, features a discussion forum and a community newsletter. The articles are targeted

towards beginners to semi-advanced users and do not simply repeat what may be read somewhere else. They are written in a “hands-on” manner. In short, they do their job! :-).

Topics range from administrative issues, like the handling of the shell, to the tweaking of the performance of *X*, *GNU/Linux*. graphical subsystem.

It may be found at MUO (<http://mandrakeuser.org/>)

A.2.1.2. Demos And Tutorials

A specific section of the **Linux-Mandrake** website is devoted to numerous demos and tutorials (<http://www.linux-mandrake.com/en/demos/>), among many :installation and graphical environment, many aspects of the configuration of your system such as network, packages maintenance, server configuration, etc...

A.2.1.3. Security Related Websites

SecurityPortal (<http://www.securityportal.com/>)

This site is devoted to the general issue of security on the Internet and contains some very interesting articles covering many aspects of the question. A weekly newsletter is also available.

Another service (with a charge) is the customized sending of security threats.

SecurityFocus (<http://www.securityfocus.com/>)

A very well organized site which reviews current attacks, gives out vulnerability advisories for a remarkably great number of products, including **Linux-Mandrake**.

LinuxSecurity (<http://www.linuxsecurity.com/>)

This one is entirely devoted to *GNU/Linux* and includes news, advisories, newsletters, and many resources such as documentation, forums, tools, etc...

Linux.com (<http://secure.linux.com/>)

An excellent site regularly fed with numerous articles on present security issues. A must for anybody who is in charge of *GNU/Linux* security.

A.2.1.4. Other GNU/Linux Websites

Out of the multiple existing websites, here are some of the most exhaustive :

- [linux.org/](http://www.linux.org/) (<http://www.linux.org/>) :one of the very first sites devoted to *GNU/Linux* which contains a whole slew of links to other useful sites;
- Freshmeat (<http://freshmeat.net/>) :this is the place to visit to get the latest applications available in the *GNU/Linux* world;
- [linux-howto/](http://www.linux-howto.com/) (<http://www.linux-howto.com/>): info and more info :-)
- Linux Gazette (<http://www.linuxgazette.com/>) :a well done online publication with interesting articles on new projects and present issues, tutorials, etc..;

And, of course, do not forget your favorite search engine. It often is the most practical tool in your search for information. A few carefully chosen keywords in a search engine will often produce the needed answers to your specific problem.

A.2.2. Mailing Lists

Mailing lists still remain very popular in spite of the multiplication of other means of communication. Every piece of *GNU/Linux* software, every project generally has its own mailing list geared towards users, developers or announcers, etc...

The **Linux-Mandrake** project has its own support lists (<http://www.Linux-Mandrake.com/en/flists.php3>).

Here, we cannot give out a list of addresses but bear in mind that it very often is the best mean to get in touch with the best experts on a particular subject. Some pieces of advice, however :

- Do not post questions which are off-topic. Carefully read the guidelines generally posted when you first subscribed or where you found the address of the list. We also recommend that you read this version of the Netiquette (<http://www.iwillfollow.com/email.htm>), where some precious hints are available. If you have spare time you may also consider reading the corresponding RFC (<http://www.cis.ohio-state.edu/htbin/rfc/rfc1855.htm>).
- Respect the general rules applicable to e-mails :in particular, do **not** send HTML messages :text only.
- Mailing lists usually have archives :check them out! your question may have been debated just before you subscribed to the list.
- God helps those who help themselves.

A.2.3. Newsgroups

Before asking for help on newsgroups, it is advisable to find out if your problem has already been covered (or solved) on *Dejanews* (http://www.deja.com/home_ps.shtml). If nothing is relevant to your question, access this newsgroup entirely devoted to **Linux-Mandrake**

(`news:alt.os.linux.mandrake`). Or you may also join many other groups in the `comp.os.linux.*` “hierarchy”:

- `comp.os.linux.setup` (`news:comp.os.linux.setup`): questions on *GNU/Linux* configuration (devices, configuration of applications) and resolution of miscellaneous problems.
- `comp.os.linux.misc` (`news:comp.os.linux.misc`): whatever does not fit in any other group.
- and others...

Before posting to one of these groups, make certain that you have done your homework and read the available documentation on your specific issue. If you have not, you will most likely get the following answer: RTFM. And nothing more!

A.3. General Guidelines For Solving A Problem Under Linux-Mandrake

Here are the different means available to you in your problem solving quest. To start with, try the first option and only then, if that did not work, try the second, and so on. As a last resort only, and only if nothing else has worked for you, start thinking about the possibility of... throwing your machine through the window.

A.3.1. RTFM

“READ THE FU**INK MANUAL!” (we could not resist, sorry! :-)
:-):-))

“The manual” means that very manual AND all the manuals and literature available on that subject. Our previous sections offer you some good starting points. Only when all these resources have been exhausted, you may start thinking you have indeed stumbled over a real problem.

A.3.2. Search The Internet

The various sites *Internet*, previously mentioned, are another good starting point. They will deal with the more general to the more specific aspects of your problem. Finally, try a general search engine such as Google (<http://www.google.com>). And do not hesitate using the Advanced search (http://www.google.com/advanced_search).

A.3.3. Mailing-lists And Newsgroups Archives

The previous searches may lead you to general answers which hide the results to your specific inquiry among many others. To refine your search, this is what you should do :

First, try to find a list which seems specifically linked to your problem, then perform a search in its archive pages.

Example

You noticed a strange behavior while trying to use *GRUB* with a Minix partition.

A search with “grub mailing list” keywords in Google gives as a second result the link to an archive’s message of the GRUB mailing-list July 1999 (<http://mail.gnu.org/pipermail/bug-grub/1999-July/003129.html>). Once there, you get the URL for the archive’s root :GRUB mailing-list archive ([http://www.mail-archive.com/bug-grub%](http://www.mail-archive.com/bug-grub%25)

40gnu.org/). This archive even suggests a search engine. Thus, searching for “Minix” will lead you directly to a patch.

Note however that all archives do not propose an embedded search engine. However, you can easily use the field `return results from the site` to limit your search to the specific site hosting the archive.

For a newsgroups search, this reference, *Dejanews* (<http://www.dejanews.com>), holds the archives for an amazing number of newsgroup channels.

A.3.4. Questions To Mailing-Lists And Newsgroups

See the related section above :*Mailing Lists*, page 346 and *Newsgroups*, page 347

A.3.5. Contacting The Person In Charge Directly

Use this option as a very last resort and in really extreme situations - unless you want to offer your collaboration:-). Software developers generally receive mountains of e-mails therefore your anguished question on the use of the `cd` command will most likely... be ignored!

The addresses will be found either on the home pages of the projects' sites or in the software documentation.

That's all for now! A last word however :do not underestimate your neighbors' skills or those of your local LUG (Linux Users group). And, please, do not throw your computer through the window as of yet. If your problem is not fixed today, it will be tomorrow :-)...

Glossary

account

on a *Unix* system, a *login* is a combination of a name, a personal directory, a password and a *shell* which allows a person to connect to this system.

alias

mechanism used in a *shell* in order to make it substitute one string for another before executing the command. You can see all aliases defined in the current session by typing *alias* at the prompt.

APM

Advanced Power Management. A feature used by some *BIOSes* in order to make the machine enter a standby state after a given period of inactivity. On laptops, APM is also responsible for reporting the battery status and, if it is supported, the estimated remaining battery life.

arp

Address Resolution Protocol. The Internet protocol used to dynamically map an Internet address to physical (hardware) addresses on local area networks. This is limited to networks that support hardware broadcasting.

ASCII

American Standard Code for Information Interchange. The standard code used for storing characters, including control characters, on a computer. Many 8-bit codes (such as ISO 8859-1, the Linux default character set) contain ASCII as their lower half

(See *ISO 8859*).

	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F
0																
1																
2		!	"	#	\$	%	&	'	()	*	+	,	-	.	/
3	0	1	2	3	4	5	6	7	8	9	:	;	<	=	>	?
4	@	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O
5	P	Q	R	S	T	U	V	W	X	Y	Z	[\]	^	_
6	'	a	b	c	d	e	f	g	h	i	j	k	l	m	n	o
7	p	q	r	s	t	u	v	w	x	y	z	{		}	~	

Figure 1. ASCII-Table

assembly language

is the programming language that is closest to the computer, thus it's called a "low level" programming language. Assembly has the advantage of speed since assembly programs are written in terms of processor instructions so little or no translation is needed when generating executables. It's main disadvantage is that it is processor (or architecture) dependent. Writing complex programs is very time-consuming as well. So, assembly is the fastest programming language, but it isn't portable between architectures.

ATAPI

("AT Attachment Packet Interface") An extension to the ATA specification ("Advanced Technology Attachment", more commonly known as IDE, *Integrated Drive Electronics*) which provides additional commands to control CDROM drives and magnetic tape drives. IDE controllers equipped with this extension are also referred to as EIDE (*Enhanced IDE*) controllers.

ATM

This is an acronym for **Asynchronous Transfer Mode**. An ATM network packages data into standard size blocks (53 bytes: 48 for the data and 5 for the header) which it can convey efficiently from point to point. ATM is a circuit switched packet network technology oriented towards high speed (multi-megabits) optical networks.

atomic

a set of operations is said to be atomic when it executes all at once, and cannot be preempted.

background

in *shell* context, a process is running in the background if you can type commands while this process is running.

See Also: job, foreground.

backup

is a means of saving your important data to a safe medium and location. Backups should be done regularly, especially with more critical information and configuration files (the prime directories to backup are */etc*, */home* and */usr/local*). Traditionally, many people use *tar* with *gzip* or *bzip2* to backup directories and files. You can use these tools or programs like *dump* and *restore*, along with many other free or commercial backup solutions.

batch

is a processing mode where jobs are submitted to the processor, and then the processor executes them one after the other till it executes the last one and it's ready for another list of processes.

beep

is the little noise your computer's speaker does to warn you of some ambiguous situation when you're using command completion and, for example, there's more than one possible choice for completion. There might be other programs that make beeps to let you know of some particular situation.

beta testing

is the name given to the process of testing the beta version of a program. Programs usually get released in alpha and beta states for testing prior to final release.

bit

stands for *BI*nary *uni*T. A single digit which can take the values 0 or 1, because calculation is done in base two.

block mode files

files whose contents are buffered. All read/write operations for such files go through buffers, which allows for asynchronous writes on the underlying hardware, and for reads, not to read again what is already in a buffer.

See Also: buffer, buffer cache.

boot

the procedure taking place when a computer is switched on, where peripherals are recognized one after the other, and where the operating system is loaded into memory.

bootdisk

a bootable floppy disk containing the code necessary to load the operating system from the hard disk (sometimes it is self-sufficient).

bootloader

is a program that starts the operating system. Many bootloaders give you the opportunity to load more than one operating system by letting you choose between them at a boot menu. Bootloaders like *GRUB* are popular because of this feature and are very useful in dual- or multi-boot systems.

BSD

Berkeley Software Distribution. A *Unix* variant developed at the Berkeley University computing department. This version has always been considered more advanced technically than the others, and has brought many innovations to the computing world in general and to *Unix* in particular.

buffer

a small portion of memory with a fixed size, which can be associated with a block mode file, a system table, a process and so on. The coherency of all buffers is maintained by the buffer cache. See **buffer cache**.

buffer cache

a crucial part of an operating system kernel, it is in charge of keeping all buffers up-to-date, shrinking the cache when needed, clearing unneeded buffers and more. See **buffers**.

bug

illogical or incoherent behavior of a program in a special case, or a behavior which does not follow the documentation or accepted standards issued for the program. Often, new features introduce new bugs in a program. Historically, this term comes from the old days of punch cards: a bug (the insect!) slipped into a hole of a punch card and, as a consequence, the program misbehaved. Ada Lovelace, having discovered this, declared "It's a bug!", and since then the term has remained.

byte

eight consecutive bits, interpreted in base two as a number between 0 and 255. See **bits**.

case

when taken in the context of strings, the case is the difference between lowercase letters and uppercase (or capital) letters.

CHAP

Challenge-Handshake Authentication Protocol: protocol used by ISPs to authenticate their clients. In this scheme, a value is sent to the client (the machine who connects), the client calculates a *hash* from this value which it sends to the server, and the server compares the *hash* with the one it has calculated. It is different from PAP in that it re-authenticates on a periodic basis after the initial authentication.

See Also: PAP.

character mode files

files whose content is not buffered. When associated to physical devices, all input/output on these devices is performed immediately. Some special character devices are created by the operating system (/dev/zero, /dev/null and others). They correspond to data flows.

See Also: block mode files.

CIFS

Common Internet FileSystem The predecessor of the SMB filesystem, used on *DOS* systems.

client

program or computer that periodically connects to another program or computer to give it orders or ask for information. In the case of **peer to peer** systems such as **slip** or **ppp** the client is taken to be the end that initiates the connection and the remote end, being called, is taken to be the server. It is one of the components of a **client/server system**.

client/server system

system or protocol consisting of a **server** and one or several **clients**.

command line

what is provided by a shell and allows the user to type commands directly. Also subject of an eternal “flame war” between its supporters and its detractors :-)

command mode

under *VI* or one of its clones, it is the state of the program in which pressing a key (this above all regards letters) will not insert the character in the file being edited, but instead perform an action specific to the said key (unless the clone has remappable commands and you have customized your configuration). You may get out of it typing one of the “back to insertion mode” commands: **i**, **I**, **a**, **A**, **s**, **S**, **o**, **O**, **c**, **C**, ...

compilation

is the process of translating source code that is human readable (well, with some training) and that is written in some programming language (*C*, for example) into a binary file that is machine readable.

completion

ability of a *shell* to automatically expand a substring to a filename, user name or other, as long as there is a match.

compression

is a way to shrink files or decrease the number of characters sent over a communications connection. Some file compression programs include *compress*, *zip*, *gzip*, and *bzip2*.

console

is the name given to what used to be called terminals. They were the users machines (a screen plus a keyboard) connected to one big central mainframe. On *PC*s, the physical terminal is the keyboard and screen.

See Also: virtual console.

cookies

temporary files written on the local hard disk by a remote web server. It allows for the server to be aware of a user's preferences when this user connects again.

datagram

A datagram is a discrete package of data and headers which contain addresses, which is the basic unit of transmission across an IP network. You might also hear this called a "packet".

dependencies

are the stages of compilation that need to be satisfied before going on to other compilation stages in order to successfully compile a program.

desktop

If you're using the X Window System, the desktop is the place on the screen inside which you work and upon which your windows and icons are displayed. It is also called the background, and is usually filled with a simple color, a gradient color or even an image.

See Also: virtual desktops.

DHCP

Dynamic Host Configuration Protocol. A protocol designed for machines on a local network to dynamically get an IP address from a DHCP server.

directory

Part of the filesystem structure. Within a directory, files or other directories are stored. Sometimes there are sub-directories (or branches) within a directory. This is often referred to as a directory tree. If you want to see what's inside another directory, you will either have to list it or change to it. Files inside a directory are referred to as leaves while sub-directories are referred to as branches. Directories follow the same restrictions as files although the permissions mean different things. The special directories '.' and '..' refer to the directory itself and to the parent directory respectively.

discrete values

are values that are non-continuous. That is, there's some kind of "spacing" between two consecutive values.

distribution

is a term used to distinguish one *GNU/Linux* vendor product from another. A distribution is made up of the core Linux kernel and utilities, as well as installation programs, third-party programs, and sometimes proprietary software.

DLCI

The DLCI is the Data Link Connection Identifier and is used to identify a unique virtual point to point connection via a Frame

Relay network. The DLCI's are normally assigned by the Frame Relay network provider.

DMA

Direct Memory Access. A facility used on the *PC* architecture which allows for a peripheral to read or write from main memory without the help of the CPU. PCI peripherals use bus mastering and do not need DMA.

DNS

Domain Name System. The distributed name/address mechanism used in the Internet. This mechanism allows you to map a domain name to an IP address, which is what lets you look up a site by domain name without knowing the IP address of the site. DNS also allows reverse lookup, that is you can get a machine's IP address from its name.

DPMS

Display Power Management System. Protocol used by all modern monitors in order to manage power saving features. Monitors supporting these features are commonly called "green monitors".

echo

is when the characters you type in a username entry field, for example, are shown "as is", instead of showing "*" for each one you type.

editor

is a term typically used for programs that edit text files (aka text editor). The most well-known *GNU/Linux* editors are the GNU Emacs (*Emacs*) editor and the *Unix* editor *VI*.

ELF

Executable and Linking Format. This is the binary format used by most *GNU/Linux* distributions nowadays.

email

stands for Electronic Mail. This is a way to send messages electronically between people on the same network. Similar to regular mail (aka snail mail), email needs a destination and sender

address to be sent properly. The sender must have an address like “sender@senders.domain” and the recipient must have an address like “recipient@recipients.domain.” Email is a very fast method of communication and typically only takes a few minutes to reach anyone, regardless of where in the world they are located. In order to write email, you need an email client like *Pine* or *mutt* which are text-mode clients, or GUI clients like *KMail*.

environment

is the execution context of a process. It includes all the information that the operating system needs to manage the process and what the processor needs to execute the process properly.

See Also: process.

environment variables

a part of a process’ environment. Environment variables are directly viewable from the *shell*.

See Also: process.

escape

in the shell context, is the action of surrounding some string between quotes to prevent the shell from interpreting that string. For example, when you need to use spaces in some command line and pipe the results to some other command you have to put the first command between quotes (“escape” the command) otherwise the shell will interpret it wrong and won’t work as expected.

ext2

short for the “Extended 2 filesystem”. This is *GNU/Linux*’ native filesystem and has all characteristics of any *Unix* filesystem: support for special files (character devices, symbolic links...), file permissions and ownership, and so on.

FAQ

Frequently Asked Questions. A document containing a series of questions/answers about a specific topic. Historically, FAQs appeared in newsgroups, but this sort of document now appears on

various web sites, and even commercial products have their FAQ. Generally, they are very good sources of information.

FAT

File Allocation Table. Filesystem used by *DOS* and *Windows*.

FDDI

Fiber Distributed Digital Interface. A high-speed network physical layer, which uses optical fiber for communication. Only used on big networks, mainly because of its price.

FHS

Filesystem Hierarchy Standard. A document containing guidelines for a coherent file tree organization on *Unix* systems. **Linux-Mandrake** complies with this standard in most aspects.

FIFO

First In, First Out. A data structure or hardware buffer from which items are taken out in the order they were put in. *Unix* pipes are the most common examples of FIFOs.

filesystem

scheme used to store files on a physical media (hard drive, floppy) in a consistent manner. Examples of filesystems are FAT, *GNU/Linux*' ext2fs, iso9660 (used by CDROMs) and so on. An example of a virtual filesystem is the */proc* filesystem.

firewall

a machine or a dedicated piece of hardware which, in the topology of a local network, is the unique connecting point to the outside network, and which filters, or controls the activity on some ports, or makes sure only some specific interfaces may have access to them.

flag

is an indicator (usually a bit) that it's used to signal some condition to a program. For example, a filesystem has, among others, a flag indicating if it has to be dumped in a backup, so when the

flag is active the filesystem gets backed up, and when it's inactive it doesn't.

focus

the state for a window to receive keyboard events (such as key-presses, key-releases and mouse clicks) unless they are trapped by the window manager.

foreground

in shell context, the process in the foreground is the one which is currently running. You have to wait for such a process to finish in order to be able to type commands again.

See Also: job, background.

framebuffer

projection of a video card's RAM into the machine's address space. This allows for applications to access the video RAM without the chore of having to talk to the card. All high-end graphical workstations use framebuffers, for example.

Frame Relay

Frame Relay is a network technology ideally suited to carrying traffic that is of bursty or sporadic nature. Network costs are reduced by having many Frame Relay customer sharing the same network capacity and relying on them wanting to make use of the network at slightly different times.

FTP

File Transfer Protocol. This is the standard *Internet* protocol used to transfer files from one machine to another.

full-screen

This term is used to refer to applications that take up the whole visible area of your display.

gateway

link connecting two IP networks.

GIF

Graphics Interchange Format. An image file format, widely used on the web. GIF images may be compressed or animated. Due to copyright problems it is a bad idea to use them, replace them as much as possible by the far advanced PNG format instead.

globbing

in the *shell*, the ability to group a certain set of filenames with a globbing pattern. See **globbing pattern**.

globbing pattern

a string made of normal characters and special characters. Special characters are interpreted and expanded by the *shell*.

GNU

GNU's Not Unix. The GNU project has been initiated by Richard Stallman at the beginning of the 80s, and aimed at developing a free operating system (“free” as in “free speech”). Currently, all tools are there, except... the kernel. The GNU project kernel, *Hurd*, is not rock solid yet. *GNU/Linux* borrows, among others, two things from GNU: its *C* compiler, *gcc*, and its license, the GPL. See **GPL**.

GPL

General Public License. The license of the *GNU/Linux* kernel, it goes the opposite way of all proprietary licenses in that it gives no restriction as to copying, modifying and redistributing the software, as long as the source code is made available. The only restriction, if one can call it that, is that the persons to which you redistribute it must also benefit from the same rights.

GUI

Graphical User Interface. Interface to a computer consisting of windows with menus, buttons, icons and so on. The vast majority prefer a GUI over a CLI (*Command Line Interface*) for ease of use, even though the latter is more versatile.

hardware address

This is a number that uniquely identifies a host in a physical network at the media access layer. Examples of this are **Ethernet Addresses** and **AX.25 Addresses**.

hidden file

is a file which can't be "seen" when doing a `ls` command with no options. Hidden files' filenames begin with a `.` and are used to store the user's personal preferences and configurations for the different programs (s)he uses. For example, *Bash*'s command history is saved into `.bash_history`, which is a hidden file.

home directory

often abbreviated by "home", this is the name for the personal directory of a given user. See also **account**.

host

refers to a computer and is commonly used when talking about computers that are connected on a network.

HTML

HyperText Markup Language. The language used to create web documents.

HTTP

HyperText Transfer Protocol. The protocol used to connect to websites and retrieve HTML documents or files.

icon

is a little drawing (normally sized 16x16, 32x32, 48x48, and sometimes 64x64 pixels) which represents, under a graphical environment, a document, a file or a program.

IDE

Integrated Drive Electronics. The most widely used bus on today's PCs for hard disks. An IDE bus can contain up to two devices, and the speed of the bus is limited by the device on the bus which has the slower command queue (and not the slower transfer rate!).
See Also: ATAPI.

inode

entry point leading to the contents of a file on a *Unix*-like filesystem. An inode is identified in a unique way by a number, and contains meta-information about the file it refers to, such as its access times, its type, its size, **but not its name!**

insert mode

under *VI* or one of its clones, it is the state of the program in which pressing a key will insert that character in the file being edited (except pathological cases like the completion of an abbreviation, right justify at the end of the line, ...). One gets out of it pressing the key Esc (or Ctrl-L).

Internet

is a huge network that connects computers around the world.

IP address

is a numeric address consisting of four parts which identifies your computer on the Internet. IP addresses are structured in a hierarchical manner, with top level and national domains, domains, subdomains and each machine's personal address. An IP address would look something like 192.168.0.1. A machine's personal address can be one of two types: static or dynamic. Static IP addresses are addresses that never change, but rather are permanent. Dynamic IP addresses mean your IP address will change with each new connection to the network. Dial-up and cable modem users typically have dynamic IP addresses while some DSL and other high-speed connections provide static IP addresses.

IP masquerading

is when you use a firewall to hide your computer's true IP address from the outside. Typically any outside network connections you make beyond the firewall will inherit the firewall's IP address. This is useful in situations where you may have a fast Internet connection with only one IP address but wish to use more than one computer that have internal network IP addresses assigned.

IRC

Internet Relay Chat. One of the few *Internet* standards for live speech. It allows for channel creation, private talks, and also file exchange. It is also designed to be able to make servers connect to each other, which is why several IRC networks exist today: **Undernet**, **DALnet**, **EFnet** to name a few.

IRC channels

are the “places” inside IRC servers where you can chat with other people. Channels are created in IRC servers and users join those channels so they can communicate with each other. Messages written on an channel are only visible to those people connected to that channel. Two or more users can also create a “private” channel so they don’t get disturbed by other users. Channel names begin with a #.

ISA

Industry Standard Architecture. The very first bus used on PCs, it is slowly being abandoned in favor of the PCI bus. Some hardware manufacturers still use it, though. It is still very common that SCSI cards supplied with scanners, CD writers, ... are ISA. Too bad.

ISDN

Integrated Services Digital Network. A set of communication standards for allowing a single wire or optical fiber to carry voice, digital network services and video. It has been designed in order to eventually replace the current phone system, known as PSTN (*Public Switched Telephone Network*). Technically ISDN is a circuit switched data network.

ISO

International Standards Organization. A group of companies, consultants, universities and other sources which enumerates standards in various topics, including computing. The papers describing standards are numbered. The standard number iso9660, for example, describes the filesystem used on CDROMs.

ISO 8859

The ISO 8859 standard includes several 8-bit extensions to the ASCII character set (see *ASCII*). Especially important is ISO 8859-1, the "Latin Alphabet No. 1", which has become widely implemented and may already be seen as the de facto standard ASCII replacement.

ISO 8859-1 (figure 2) supports the following languages: Afrikaans, Basque, Catalan, Danish, Dutch, English, Faroese, Finnish, French, Galician, German, Icelandic, Irish, Italian, Norwegian, Portuguese, Scottish, Spanish, and Swedish.

	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F
8																
9																
A		¡	¢	£	¤	¥	¦	§	¨	©	ª	«	¬	­	®	¯
B	°	±	²	³	´	µ	¶	·	¸	¹	º	»	¼	½	¾	¿
C	À	Á	Â	Ã	Ä	Å	Æ	Ç	È	É	Ê	Ë	Ì	Í	Î	Ï
D	Ð	Ñ	Ò	Ó	Ô	Õ	Ö	×	Ø	Ù	Ú	Û	Ü	Ý	Þ	ß
E	à	á	â	ã	ä	å	æ	ç	è	é	ê	ë	ì	í	î	ï
F	ð	ñ	ò	ó	ô	õ	ö	÷	ø	ù	ú	û	ü	ý	þ	ÿ

Figure 2. ISO-8859-1 Table

Note that the ISO 8859-1 characters are also the first 256 characters of ISO 10646 (Unicode). However, it lacks the EURO symbol and does not fully cover Finnish and French. ISO 8859-15 (figure 3) is a modification of ISO 8859-1 that covers these needs.

The full set of ISO 8859 alphabets includes:

Name	Language(s)
ISO 8859-1	west European languages (Latin-1)

Name	Language(s)
ISO 8859-2	east European languages (Latin-2)
ISO 8859-3	southeast European and miscellaneous (Latin-3)
ISO 8859-4	Scandinavian/Baltic languages (Latin-4)
ISO 8859-5	Latin/Cyrillic
ISO 8859-6	Latin/Arabic
ISO 8859-7	Latin/Greek
ISO 8859-8	Latin/Hebrew
ISO 8859-9	Latin-1 modification for Turkish (Latin-5)
ISO 8859-10	Lappish/Nordic/Eskimo languages (Latin-6)
ISO 8859-11	Thai
ISO 8859-13	Baltic Rim languages (Latin-7)
ISO 8859-14	Celtic (Latin-8)
ISO 8859-15	west European languages with Euro symbol (Latin-9)

	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F
8																
9																
A		ı	ç	£	€	¥	Š	š	š	©	ª	«	¬	–	®	-
B	º	±	²	³	Ž	µ	¶	·	ž	¹	º	»	œ	æ	ÿ	ı
C	À	Á	Â	Ã	Ä	Å	Æ	Ç	È	É	Ê	Ë	Ì	Í	Î	Ï
D	Ð	Ñ	Ò	Ó	Ô	Õ	Ö	×	Ø	Ù	Ú	Û	Ü	Ý	Þ	ß
E	à	á	â	ã	ä	å	æ	ç	è	é	ê	ë	ì	í	î	ï
F	ð	ñ	ò	ó	ô	õ	ö	÷	ø	ù	ú	û	ü	ý	þ	ÿ

Figure 3. ISO-8859-15 Table

ISP

Internet Service Provider. A company which sells *Internet* access to its customers, whether the access is over telephone lines or dedicated lines.

job

in *shell* context, a job is a process running in the background. You can have several jobs in the same shell and control these jobs. See also **background**, **foreground**.

JPEG

Joint Photographic Experts Group. Another very common image file format. JPEG is mostly suited for compressing real-world scenes, and does not work very well on non-realistic images.

kernel

is the guts of the operating system. The kernel is responsible for allocating resources and separating processes from each other. It handles all of the low-level operations that allow programs to talk directly to the hardware on your computer, manages the buffer cache and so on.

kill ring

under *Emacs*, it is the set of text areas cut or copied since the beginning of the editor, which may be recalled to be inserted again, and which is organized like a ring.

LAN

Local Area Network. Generic name given to a network of machines connected to the same physical wire.

launch

is the action of invoking, or starting, a program.

LDP

Linux Documentation Project. A nonprofit organization which maintains *GNU/Linux* documentation. Its mostly known documents are *HOWTOs*, but it also maintains FAQs, and even a few books.

library

is a collection of procedures and functions in binary form to be used by programmers in their programs (as long as the library's license allows them to do so). The program in charge of loading shared libraries at run time is called the dynamic linker.

link

reference to an inode in a directory, therefore giving a (file) name to the inode. Examples of inodes which don't have a link (and hence have no name) are: anonymous pipes (as used by the shell), sockets (aka network connections), network devices and so on.

linkage

last stage of the compile process, which consists in linking together all object files in order to produce an executable file, and matches unresolved symbols with dynamic libraries (unless a static linkage has been asked, in which case the code of these symbols will be included in the executable).

Linux

is a *Unix*-like operating system which runs on a variety of different computers, and is free for anyone to use and modify. Linux (the kernel) was written by Linus Torvalds.

login

connection name for a user on a *Unix* system, and the action to connect.

lookup table

is a table that puts in correspondance codes (or tags) and their meaning. It is often a data file used by a program to get further information about a particular item.

For example, *HardDrake* uses such a table to know what a manufacturer's product code means. This is one line from the table, giving information about item CTL0001

```
CTL0001 sound sb Creative Labs SB16 \
HAS_OPL3|HAS_MPU401|HAS_DMA16|HAS_JOYSTICK
```

loopback

virtual network interface of a machine to itself, allowing the running programs not to have to take into account the special case where two network entities are in fact the same machine.

major

number specific to the device class.

manual page

a small document containing the definition of a command and its usage, to be consulted with the `man` command. The first thing one should (learn how to) read when hearing of a command he doesn't know :-)

MBR

Master Boot Record. Name given to the first sector of a bootable hard drive. The MBR contains the code used to load the operating system into memory or a bootloader (such as *LILLO*), and the partition table of that hard drive.

MIME

Multipurpose Internet Mail Extensions. A string of the form type/subtype describing the contents of a file attached in an e-mail. This allows MIME-aware mail clients to define actions depending on the type of the file.

minor

number identifying the specific device we are talking about.

mount point

is the directory where a partition or another device is attached to the *GNU/Linux* filesystem. For example, your CDROM is mounted in the `/mnt/cdrom` directory, from where you can explore the contents of any mounted CDs.

mounted

A device is mounted when it is attached to the *GNU/Linux* filesystem. When you mount a device you can browse its con-

tents. This term is partly obsolete as with the “supermount” feature, users do not need any more to manually mount removable medias.

See Also: mount point.

MPEG

Moving Pictures Experts Group. An ISO committee which generates standards for video and audio compression. MPEG is also the name of their algorithms. Unfortunately, the license for this format is very restrictive, and as a consequence there are still no *Open Source* MPEG players...

MSS

The Maximum Segment Size (**MSS**) is the largest quantity of data that can be transmitted at one time. If you want to prevent local fragmentation MSS would equal MTU-IP header.

MTU

The Maximum Transmission Unit (**MTU**) is a parameter that determines the largest datagram than can be transmitted by an IP interface without it needing to be broken down into smaller units. The MTU should be larger than the largest datagram you wish to transmit unfragmented. Note, this only prevents fragmentation locally, some other link in the path may have a smaller MTU and the datagram will be fragmented there. Typical values are 1500 bytes for an ethernet interface, or 576 bytes for a SLIP interface.

multitasking

the ability for an operating system to share CPU time between several processes. At low level, this is also known as multiprogramming. Switching from one process to another requires that all the current process context be saved and restored when this process is elected again. This operation is called context switch, and on Intel, is done 100 times per second; therefore it's fast enough so that a user has the illusion that the operating system runs several applications at the same time. There are two types of multitasking: preemptive multitasking is where the operating

system is responsible for taking away the CPU and pass it to another process; cooperative multitasking is where the process itself gives back the CPU. The first variant is, obviously, the better choice because no program can consume the entire CPU time and block other processes. *GNU/Linux* does preemptive multitasking. The policy to select which process should be run, depending on several parameters, is called scheduling.

multiuser

is used to describe an operating system which allows multiple users to log into and use the system at the exact same time, each being able to do their own work independent of other users. A multitasking operating system is required to provide multiuser support. *GNU/Linux* is both a multitasking and multiuser operating system, as any *Unix* system for that matter.

named pipe

a *Unix* pipe which is linked, as opposed to pipes used in shells. See also **pipe**, **link**.

naming

a word commonly used in computing for a method to identify objects. You will often hear of “naming conventions” for files, functions in a program and so on.

NCP

NetWare Core Protocol. A protocol defined by **Novell** to access Novell NetWare file and print services.

newsgroups

discussion and news areas that can be accessed by a news or USENET client to read and write messages specific to the topic of the newsgroup. For example, the newsgroup `alt.os.linux.mandrake` is an alternate newsgroup (alt) dealing with the Operating System (os) *GNU/Linux*, and specifically, **Linux-Mandrake** (mandrake). Newsgroups are broken down in this fashion to make it easier to search for a particular topic.

NFS

Network FileSystem. A network filesystem created by **Sun Microsystems** in order to share files across a network in a transparent way.

NIC

Network Interface Controller. An adapter installed in a computer which provides a physical connection to a network, such as an *Ethernet* card.

NIS

Network Information System. NIS was also known as “Yellow Pages”, but **British Telecom** holds a copyright on this name. NIS is a protocol designed by **Sun Microsystems** in order to share common information across a NIS **domain**, which can gather a whole LAN, part of this LAN or several LANs. It can export password databases, service databases, groups information and more.

null, character

the character or byte number 0, it is used to mark the end of a string.

object code

is the code generated by the compilation process to be linked with other object codes and libraries to form an executable file. Object code is machine readable.

See Also: compilation, linkage.

on the fly

Something is said to be done “on the fly” when it’s done along with something else, without you noticing it or explicitly asking for it.

open source

is the name given to free source code of a program that is made available to development community and public at large. The theory behind this is that allowing source code to be used and modified by a broader group of programmers will ultimately produce

a more useful product for everyone. Some popular open source programs include *Apache*, *sendmail* and *GNU/Linux*.

operating system

is the interface between the applications and the underlying hardware. The tasks for any operating system are primarily to manage all of the machine specific resources. On a *GNU/Linux* system, this is done by the kernel and loadable modules. Other well-known operating systems include *AmigaOS*, *MacOS*, *FreeBSD*, *OS/2*, *Unix*, *Windows NT*, and *Windows 9x*.

owner

in the context of users and their files, the owner of a file is the user who created that file.

owner group

in the context of groups and their files, the owner group of a file is the group to which the user who created that file belongs to.

pager

program displaying a text file one screenful at a time, and making it easy to move back and forth and search for strings in this file. We advise you to use `less`.

PAP

Password Authentication Protocol. A protocol used by many ISPs to authenticate their clients. In this scheme, the client (you) sends an identifier/password pair to the server, which is not encrypted. See also **CHAP**.

password

is a secret word or combination of words or letters that is used to secure something. Passwords are used in conjunction with user logins to multi-user operating systems, web sites, FTP sites, and so forth. Passwords should be hard-to-guess phrases or alphanumeric combinations, and should never be based on common dictionary words. Passwords ensure that other people cannot log into a computer or site with your account.

patch, to patch

file holding a list of corrections to issue to a source code in order to add new features, to remove bugs, or to modify it according to one's wishes and needs. The action consisting of the application of these corrections to the archive of source code (aka "patching").

path

is an assignment for files and directories to the filesystem. The different layers of a path are separated by the "slash" or '/' character. There are two types of paths on *GNU/Linux* systems. The **relative** path is the position of a file or directory in relation to the current directory. The **absolute** path is the position of a file or directory in relation to the root directory.

PCI

Peripheral Components Interconnect. A bus created by **Intel** and which is today the standard bus for *PC* architectures, but other architectures use it too. It is the successor of ISA, and it offers numerous services: device identification, configuration information, IRQ sharing, bus mastering and more.

PCMCIA

Personal Computer Memory Card International Association. More and more commonly called "PC Card" for simplicity reasons, this is the standard for external cards attached to a laptop: modems, hard disks, memory cards, *Ethernet* cards, and more. The acronym is sometimes humorously expanded to *People Cannot Memorize Computer Industry Acronyms...*

pipe

a special *Unix* file type. One program writes data into the pipe, and another program reads the data at the other end. *Unix* pipes are FIFOs, so the data is read at the other end in the order it was sent. Very widely used with the shell. See also **named pipe**.

pixmap

is an acronym for “pixel map”. It’s another way of referring to bitmapped images.

plugin

add-on program used to display or play some multimedia content found on a web document. It can usually be easily downloaded if your browser is not yet able to display or play that kind of information.

PNG

Portable Network Graphics. Image file format created mainly for web use, it has been designed as a patent-free replacement for GIF and also has some additional features.

PNP

Plug’N’Play. First an add-on for ISA in order to add configuration information for devices, it has become a more widespread term which groups all devices able to report their configuration parameters. As such, all PCI devices are Plug’N’Play.

POP

Post Office Protocol. The common protocol used for downloading mail from an ISP.

porting

a program is translating that program in such a way that it can be used in a system it was not originally intended for, or it can be used in “similar” systems. For example, to be able to run a *Windows*-native program under *GNU/Linux* (natively), it must first be ported to *GNU/Linux*.

PPP

Point to Point Protocol. This is the protocol used to send data over serial lines. It is commonly used to send IP packets to the Internet, but it can also be used with other protocols such as Novell’s IPX protocol.

precedence

dictates the order of evaluation of operands in an expression. For example: If you have $4 + 3 * 2$ you get 10 as the result, since the product has more precedence than the addition. If you want to evaluate the addition first, then you have to add parenthesis like this $(4 + 3) * 2$, and you get 14 as the result since the parenthesis have more precedence than the addition and the product, so the operations in parenthesis get evaluated first.

preprocessors

are compilation directives that instruct the compiler to replace those directives for code in the programming language used in the source file. Examples of *C*'s preprocessors are `#include`, `#define`, etc.

process

in the operating system context, a process is an instance of a program being executed along with its environment.

prompt

in a *shell*, this is the string before the cursor. When you see it, you can type your commands.

protocol

Protocols organize the communication between different machines across a network, either using hardware or software. They define the format of transferred data, whether one machine controls another, etc. Many well-known protocols include HTTP, FTP, TCP, and UDP.

proxy

a machine which sits between a network and the *Internet*, whose role is to speed up data transfers for the most widely used protocols (HTTP and FTP, for example). It maintains a cache of previous demands, which avoids the cost of asking for the file again if another machine asks for the same thing. Proxies are very useful on low bandwidth networks (such as modem connections). Sometimes the proxy is the only machine able to access outside the network.

pull-down menu

it is a menu that is “rolled” with a button in some of its corners. When you press that button, the menu “unrolls” itself showing you the full menu.

quota

is a method for restricting disk usage and limits for users. Administrators can restrict the size of home directories for a user by setting quota limits on specific filesystems.

RAID

Redundant Array of Independent Disks. A project initiated at the computing science department of Berkeley University, in which the storage of data is spread along an array of disks using different schemes. At first, this was implemented using floppy drives, which is why the acronym originally stood for *Redundant Array of Inexpensive Disks*.

RAM

Random Access Memory. Term used to identify a computer’s main memory. The “Random” here means that any part of the memory can be directly accessed...

read-only mode

for a file means that the file cannot be written to. You can read its contents but you can’t them and you can’t erase the file.

read-write mode

for a file, it means that the file can be written to. You can read its contents and modify them. If you have this kind of permission on a file, you can also erase that file.

regular expression

a powerful theoretical tool which is used to search and match text strings. It lets one specify patterns these strings must obey. Many *Unix* utilities use it: *sed*, *awk*, *grep*, *Perl* among others.

RFC

Request For Comments. RFCs are the official *Internet* standard documents, published by the IETF (*Internet Engineering Task Force*). They describe all protocols, their usage, their requirements and so on. When you want to learn how a protocol works, pick up the corresponding RFC.

root

is the superuser of any *Unix* system. Typically root (aka the system administrator) is the person responsible for maintaining and supervising the *Unix* system. This person also has complete access to everything on the system.

root directory

This is the top level directory of a filesystem. This directory has no parent directory, thus `'..'` for root points back to itself. The root directory is written as `'/'`.

root filesystem

This is the top level filesystem. This is the filesystem where *GNU/Linux* mounts its root directory tree. It is necessary for the root filesystem to reside in a partition of its own, as it is the basis for the whole system. It holds the root directory.

route

Is the path that your datagrams take through the network to reach their destination. Is the path between one machine and another in a network.

RPM

Redhat Package Manager. A packaging format developed by **Red Hat** in order to create software packages, it is used in many *GNU/Linux* distributions, including **Linux-Mandrake**.

run level

is a configuration of the system software that only allows certain selected processes to exist. Allowed processes are defined, for each runlevel, in the file `/etc/inittab`. There are eight defined

runlevels: 0, 1, 2, 3, 4, 5, 6, S and switching among them can only be achieved by a privileged user by means of executing the commands `init` and `telinit`.

script

shell scripts are sequences of commands to be executed as if they were entered in the console one after the other. *shell* scripts are *Unix*'s (somewhat) equivalent of *DOS* batch files.

SCSI

Small Computers System Interface. A bus with a high throughput designed to allow for several types of peripherals. Unlike IDE, a SCSI bus is not limited by the speed at which the peripherals accept commands. Only high-end machines integrate a SCSI bus directly on the motherboard, *PC*s need add-on cards.

security levels

Linux-Mandrake's unique feature that allows you to set different levels of restrictions according to how secure you want to make your system. There are 6 predefined levels ranging from 0 to 5, where 5 is the tightest security. You can also define your own security level.

server

program or computer that provides a feature or service and awaits the connections from **clients** to execute their orders or give them the information they ask. In the case of **peer to peer** systems such as **slip** or **ppp** the server is taken to be the end of the link that is called and the end calling is taken to be the client. It is one of the components of a **client/ server system**.

shadow passwords

a password management suite on *Unix* systems in which the file containing the encrypted passwords is not world-readable, whereas it is when using the normal password system. It also offers other features such as password aging.

shell

The *shell* is the basic interface to the operating system kernel and is what provides the command line where users enter commands to run programs and system commands. All shells provide a scripting language which can be used to automate tasks or simplify often-used complex tasks. These *shell* scripts are similar to batch files from the *DOS* operating system, but are much more powerful. Some example shells are *Bash*, *sh*, and *tcsh*.

single user

is used to describe a state of an operating system, or even an operating system itself, that only allows a single user to log into and use the system at any time.

site dependent

means that the information used by programs like *Imake* and *make* to compile some source file depends on the site, the computer architecture, the computer's installed libraries, and so on.

SMB

Server Message Block. Protocol used by *Windows* machines (*9x* or *NT*) for file and printer sharing across a network. See also **CIFS**.

SMTP

Simple Mail Transfer Protocol. This is the common protocol for transferring email. Mail Transfer Agents such as *sendmail* or *postfix* use SMTP. They are sometimes also called SMTP servers.

socket

file type corresponding to any network connection.

soft links

see “symbolic links”.

standard error

the file descriptor number 2, opened by every process, used by convention to print error messages to the terminal screen by default. See also **standard input**, **standard output**.

standard input

the file descriptor number 0, opened by every process, used by convention as the file descriptor from which the process receives data. See also **standard error**, **standard output**.

standard output

the file descriptor number 1, opened by every process, used by convention as the file descriptor in which the process prints its output. See also **standard error**, **standard input**.

streamer

is a device that takes “streams” (not interrupted or divided in shorter chunks) of characters as its input. A typical streamer is a tape drive.

SVGA

Super Video Graphics Array. The video display standard defined by VESA for the *PC* architecture. The resolution is 800x600 x 16 colors.

switch

Switches are used to change the behavior of programs, and are also called command-line options or arguments. To determine if a program has optional switches that can be used, read the *man* pages or try to pass the `-help` switch to the program (ie. `program -help`).

symbolic links

special files, containing nothing but a string that makes reference to another file. Any access to them is the same as accessing the file whose name is the referenced string, which may or may not exist, and the path to which can be given in a relative or an absolute way.

target

is the object of compilation, i.e. the binary file to be generated by the compiler.

TCP

Transmission Control Protocol. This is the most common reliable protocol that uses IP to transfer network packets. TCP adds the necessary checks on top of IP to make sure that packets are delivered. Unlike UDP, TCP works in connected mode, which means that two machines must establish a connection before exchanging data.

telnet

creates a connection to a remote host and allows you to log into the machine, provided you have an account. Telnet is the most widely-used method of remote logins, however there are better and more secure alternatives, like ssh.

theme-able

a graphical application is theme-able if it is able to change its appearance in real time. Many window managers are theme-able as well.

traverse

for a directory on a *Unix* system, this means that the user is allowed to go through this directory, and possibly to directories under it. This requires that the user has the execute permission on this directory.

URL

Uniform Resource Locator. A string with a special format used to identify a resource on the *Internet* in a unique way. The resource can be a file, a server or other item. The syntax for a URL is `protocol://server.name[:port]/path/to/resource`. When only a machine name is given and the protocol is `http://`, it defaults to retrieving the file `index.html` on the server.

username

is a name (or more generally a word) that identifies a user in a system. Each username is attached to a unique and single UID (user ID)

See Also: login.

variables

are strings that are used in `Makefile` files to be replaced by their value each time they appear. Usually they are set at the beginning of the `Makefile`. They are used to simplify `Makefile` and source files tree management.

More generally, variables in programming, are words that refer to other entities (numbers, strings, tables, etc.) that are likely to vary while the program is executing.

verbose

For commands, the verbose mode means that the command reports to standard (or possibly error) output all the actions it performs and the results of those actions. Sometimes, commands have a way to define the “verbosity level”, which means that the amount of information that the command will report can be controlled.

VESA

Video Electronics Standards Association. An industry standards association aimed at the *PC* architecture. It is the author of the *SVGA* standard, for example.

virtual console

is the name given to what used to be called terminals. On *GNU/Linux* systems, you have what are called virtual consoles which enable you to use one screen or monitor for many independently running sessions. By default, you have six virtual consoles which can be reached by pressing **ALT-F1** through **ALT-F6**. There is a seventh virtual console by default, **ALT-F7**, which will permit you to reach a running X Window System. In X, you can reach the text console by pressing **CTRL-ALT-F1** through **CTRL-ALT-F6**.

See Also: console.

virtual desktops

In the X Window System, the *window manager* may provide you several *desktops*. This handy feature allows you to organize your windows, avoiding the problem of having dozens of them stacked

on top of each other. It works as if you had several screens. You can switch from one virtual desktop to another in a manner that depends on the window manager you're using.

WAN

Wide Area Network. This network, although similar to a LAN connects computers on a network that is not physically connected to the same wires and are separated by a greater distance.

wildcard

The '*' and '?' characters are used as wildcard characters and can represent anything. The '*' represents any number of characters, including no characters. The '?' represents exactly one character. Wildcards are often used in regular expressions.

window

In networking, the **window** is the largest amount of data that the receiving end can accept at a given point in time.

window manager

the program responsible for the “look and feel” of a graphical environment, dealing with window bars, frames, buttons, root menus, and some keyboard shortcuts. Without it, it would be hard or impossible to have virtual desktops, to resize windows on the fly, to move them around, ...

Appendix B. The GNU General Public License

The following text is the GPL license that applies to most programs found in **Linux-Mandrake** distributions.

Version 2, June 1991 Copyright (C) 1989, 1991 Free Software Foundation, Inc. 59 Temple Place - Suite 330, Boston, MA 02111-1307, USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

B.1. Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software – to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation’s software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Library General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These

restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps:

1. copyright the software, and
2. offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

B.2. Terms and conditions for copying, distribution and modification

- 0. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be dis-

tributed under the terms of this General Public License. The “Program”, below, refers to any such program or work, and a “work based on the Program” means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term “modification”.) Each licensee is addressed as “you”.

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

- 1. You may copy and distribute verbatim copies of the Program’s source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

- 2. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:
 - 1. You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.

2. You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.
3. If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

- 3. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:
 1. Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
 2. Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
 3. Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled

to copy the source along with the object code.

- 4. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.
- 5. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.
- 6. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.
- 7. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program

by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

- 8. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.
- 9. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Pro-

gram specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.

- 10. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

- 11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.
- 12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT

HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

Appendix C. GNU Free Documentation License

Version 1.1, March 2000

Copyright (C) 2000 Free Software Foundation, Inc. 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

0. PREAMBLE

The purpose of this License is to make a manual, textbook, or other written document "free" in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or noncommercially. Secondly, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of "copyleft", which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

1. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. The "Document", below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as "you".

A "Modified Version" of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A "Secondary Section" is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document's overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (For example, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The "Invariant Sections" are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this License.

The "Cover Texts" are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License.

A "Transparent" copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, whose contents can be viewed and edited directly and straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format

whose markup has been designed to thwart or discourage subsequent modification by readers is not Transparent. A copy that is not "Transparent" is called "Opaque".

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML designed for human modification. Opaque formats include PostScript, PDF, proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML produced by some word processors for output purposes only.

The "Title Page" means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, "Title Page" means the text near the most prominent appearance of the work's title, preceding the beginning of the body of the text.

2. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

3. COPYING IN QUANTITY

If you publish printed copies of the Document numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a publicly-accessible computer-network location containing a complete Transparent copy of the Document, free of added material, which the general network-using public has access to download anonymously at no charge using public-standard network protocols. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

4. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

- A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.
- B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has less than five).
- C. State on the Title page the name of the publisher of the Modified Version, as the publisher.
- D. Preserve all the copyright notices of the Document.
- E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
- F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.
- G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
- H. Include an unaltered copy of this License.
- I. Preserve the section entitled "History", and its title, and add to it an item stating at least the title, year, new authors, and publisher

of the Modified Version as given on the Title Page. If there is no section entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.

- J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.
- K. In any section entitled "Acknowledgements" or "Dedications", preserve the section's title, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.
- L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.
- M. Delete any section entitled "Endorsements". Such a section may not be included in the Modified Version.
- N. Do not retitle any existing section as "Endorsements" or to conflict in title with any Invariant Section.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles.

You may add a section entitled "Endorsements", provided it contains nothing but endorsements of your Modified Version by various parties—for example, statements of peer review or that the text has been

approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

5. COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections entitled "History" in the various original documents, forming one section entitled "History"; likewise combine any sections entitled "Acknowledgements", and any sections entitled "Dedications". You must delete all sections entitled "Endorsements."

6. COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

7. AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, does not as a whole count as a Modified Version of the Document, provided no compilation copyright is claimed for the compilation. Such a compilation is called an "aggregate", and this License does not apply to the other self-contained works thus compiled with the Document, on account of their being thus compiled, if they are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one quarter of the

entire aggregate, the Document's Cover Texts may be placed on covers that surround only the Document within the aggregate. Otherwise they must appear on covers around the whole aggregate.

8. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License provided that you also include the original English version of this License. In case of a disagreement between the translation and the original English version of this License, the original English version will prevail.

9. TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided for under this License. Any other attempt to copy, modify, sublicense or distribute the Document is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

10. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but

may differ in detail to address new problems or concerns. See <http://www.gnu.org/copyleft/> (<http://www.gnu.org/copyleft/>).

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License "or any later version" applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation.

How to use this License for your documents

To use this License in a document you have written, include a copy of the License in the document and put the following copyright and license notices just after the title page:

Copyright (c) YEAR YOUR NAME. Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.1 or any later version published by the Free Software Foundation; with the Invariant Sections being LIST THEIR TITLES, with the Front-Cover Texts being LIST, and with the Back-Cover Texts being LIST. A copy of the license is included in the section entitled "GNU Free Documentation License".

If you have no Invariant Sections, write "with no Invariant Sections" instead of saying which ones are invariant. If you have no Front-Cover Texts, write "no Front-Cover Texts" instead of "Front-Cover Texts being LIST"; likewise for Back-Cover Texts.

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license, such as the GNU General Public License, to permit their use in free software.

Index

- /dev/hda, 41
- /dev/sda, 41
- admin
 - password, 66
- ADSL, 253
- Appletalk, 202
- ARCNet, 202
- ATM, 207
- AX25, 207
- backup, 339
- BIOS, 25
- bootdisk, 26, 66
- bootloader, 67
- Cable, 259
 - null modem, 218
 - parallel, 218
 - PLIP, 218
 - twisted pair, 221
- command
 - rawrite, 26
- DECNet, 208
- DHCP, 198, 292
- DMA channel, 35
- DNS, 197
- domain name, 229
- DOS, 313
- eth0, 195
- Ethernet
 - card, 191
 - configuration, 232
 - detection, 234
 - editing, 234
- FDDI, 208

- Frame Relay, 209
- gateway, 301
- GFDL, 397
- GPL, 387
- I/O address, 35
- IP
 - address, 183
 - routing, 187
- IPX, 213
- IRQ, ??
- ISA
 - Plug and Play, 35
- ISDN, 198, 246
- login, 225
- logout, 228
- MacOS, 314
- modem, 242
- monitoring
 - logs, 324
 - network, 323
 - system, 323
- netiquette, 347
- netmask, 185
- NetRom, 214
- network
 - cable, 218
- network, 179
 - class, 185
 - configuration, 182
 - external, 60
 - internal, 56
 - private, 186
- OS/2, 318
- partition, 37
- PLIP, 200

- PPP, 200, 201
- proxy, 296
- restore, 339
- RFC, 181
- root
 - password, 63
- routing, 187
- SAMBA, 215
- sector, 37
- security
 - level, 50
- services
 - activation, 299
- squid, 296
- ssh, 338
- swap, 38
 - size, 39
- system name, 229
- TCP/IP, 180
- timezone, 62
- Token Ring, 216, 217
- update, 331
- USB, 32
- wavelan, 217
- windows 3.11, 313
- windows 95/98, 305
- windows NT/2000, 309
- winmodem, 31

