

Server Administration Guide

Mandriva Linux 2006



<http://www.mandriva.com>

Server Administration GuideMandriva Linux 2006

Published September 2005

Copyright © 2005 Mandrakesoft SA dba Mandriva

by Camille Bégnis, Fabian Mandelbaum, Christian Roy, Fred Lepied, Nicolas Planel, Daouda Lo, François Pons, John Rye, Pascal Rigaux, Damien Chaumette, Till Kamppeter, Florent Villard, Luca Berra, Florin Grad, Frédéric Crozat, Stew Benedict, Guillaume Cottenceau, Thierry Vignaud, Jean-Michel Dault., and Lunas Moon

Legal Notice

This material may be distributed only subject to the terms and conditions set forth in the Open Publication License, v1.0 or later (the latest version is presently available at [opencontent.org \(http://www.opencontent.org/openpub/\)](http://www.opencontent.org/openpub/)).

- Distribution of substantively modified versions of this document is prohibited without the explicit permission of the copyright holder.
- Distribution of the work or derivative of the work in any standard (paper) book form is prohibited unless prior permission is obtained from the copyright holder.

“Mandriva” and “DrakX” are registered trademarks in the US and/or other countries. The related “Star logo” is also registered. All rights reserved. All other copyrights embodied in this document remain the property of their respective owners.

About the Making of this Manual

This manual is written and maintained by NeoDoc (<http://www.neodoc.biz>). Translations are ensured by NeoDoc, Mandriva and other translators.

This document was written in DocBook XML. The set of files involved were managed using the Borges Collaborative Content Creation System (C3S) (<http://sourceforge.net/projects/borges-dms>). The XML source files were processed by `xsltproc`, and `jadetex` (for the electronic version) using a customized version of Norman Walsh’s stylesheets. Screen shots were taken using `xwd` or `GIMP` and converted with `convert` (from the ImageMagick package). All these programs are free software and all of them are available in your Mandriva Linux distribution.

Table of Contents

Preface	1
1. About Mandriva Linux	1
1.1. Contacting the Mandriva Linux Community	1
1.2. Join the Club!	1
1.3. Subscribing to Mandriva Online	2
1.4. Purchasing Mandriva Products	2
1.5. Contributing to Mandriva Linux	2
2. About this Server Administration Guide	2
3. Note from the Editor	2
4. Conventions Used in this Book	3
4.1. Typing Conventions	3
4.2. General Conventions	4
I. Configuration Wizards for Common Services	7
1. Server Configuration Wizards	7
1.1. Foreword	7
1.2. DHCP Server Configuration	8
1.3. DNS Server Configuration	9
1.4. Configuring Mail with Postfix Server	10
1.5. Samba Server Configuration	11
1.6. Web Server Configuration	13
1.7. FTP Server Configuration	14
1.8. Installation Server Wizard	16
1.9. NIS and Autofs Servers Wizard	17
1.10. LDAP Configuration Wizard	17
1.11. News Server Configuration	18
1.12. Proxy Server Configuration	18
1.13. Time Configuration	20
2. Configuring Masqueraded Clients	23
2.1. Linux Box	23
2.2. Windows XP Box	24
2.3. Windows 95 or Windows 98 Box	25
2.4. Windows NT or Windows 2000 Box	27
2.5. DOS Box Using the NCSA Telnet Package	30
2.6. Windows for Workgroups 3.11	31
2.7. MacOS Box	31
2.8. OS/2 Warp Box	36
II. In-Depth Configuration of Common Services	39
3. BIND DNS Server	39
3.1. Installation and Initialization	39
3.2. Step-by-Step Configuration Example	40
3.3. Advanced Configuration and Troubleshooting	44
4. Internet and Intranet Web Server	47
4.1. Installation	47
4.2. Configuration Example	47
4.3. Advanced Configuration	50
4.4. More Documentation	52
5. Postfix Mail Server	53
5.1. SMTP Server Functions	53
5.2. Installation	53
5.3. Step-by-Step Configuration Example	53
5.4. Advanced Configuration	55
5.5. Extra Documentation	57
6. Mail Delivery Services: Pop and IMAP	59
6.1. Foreword and Installation	59
6.2. Step-by-Step Configuration Example	59
6.3. Advanced Configuration	60
7. Resource Sharing	61
7.1. Samba: Integrating Linux in a Windows Network	61

7.2. Resource Sharing: FTP	64
7.3. NFS: Exporting Directories To UNIX/Linux Hosts	67
8. The Kolab Server	69
8.1. Introduction	69
8.2. Overview	69
8.3. Installation	69
8.4. The Kolab Administration Interface	70
9. MySQL Database Server	79
9.1. Getting Started	79
9.2. Creating a User for the Database	80
9.3. Creating a Database	80
9.4. Creating a Table	80
9.5. Managing Data in a Table	81
9.6. More Documentation	81
10. NIS Client and Server	83
10.1. Installation	83
10.2. Step-by-Step Configuration	83
10.3. Advanced Client Configuration	84
10.4. Importing homes with autofs	84
III. Applied Theory and Troubleshooting	87
11. Security Under GNU/Linux	87
11.1. Preamble	87
11.2. Overview	87
11.3. Physical Security	90
11.4. Local Security	94
11.5. Files and File-System Security	95
11.6. Password Security and Encryption	100
11.7. Kernel Security	105
11.8. Network Security	108
11.9. Security Preparation (Before You Go On-Line)	114
11.10. What to Do During and After a Break-in	116
11.11. Security Sources	117
11.12. Frequently Asked Questions	119
11.13. Conclusion	121
Security-Related Terms	121
12. Networking Overview	123
12.1. Copyright	123
12.2. How to Use this Chapter	123
12.3. General Information about Linux Networking	124
12.4. Generic Network Configuration Information	125
12.5. Ethernet Information	128
12.6. IP-Related Information	129
12.7. Using Commodity PC Hardware	130
12.8. Other Network Technologies	131
12.9. Cables and Cabling	132
13. Troubleshooting	137
13.1. Introduction	137
13.2. A Boot Disk	137
13.3. Backup	138
13.4. Restore	140
13.5. Problems Arising at Boot Time	141
13.6. Bootloader Issues	142
13.7. Filesystem Issues	143
13.8. Recovering from a System Freeze	144
13.9. Killing Misbehaving Apps	145
13.10. Miscellaneous	146
13.11. Mandriva Linux's Specific Troubleshooting Tools	147
13.12. General Guidelines for Solving a Problem under Mandriva Linux	147
13.13. Final Thoughts	148
A. Glossary	149

Index.....167

List of Tables

12-1. Reserved Private Network Allocations.....	126
---	-----

Preface

1. About Mandriva Linux

Mandriva Linux is a GNU/Linux distribution supported by Mandriva S.A. which was born on the Internet in 1998. Its main goal was and still is to provide an easy-to-use and friendly GNU/Linux system. Mandriva's two pillars are open source and collaborative work.



On April 7th 2005 the Mandrakesoft company changed its name to Mandriva to reflect its merger with Brazil-based Conectiva. Its core product, Mandrakelinux, became Mandriva Linux.

1.1. Contacting the Mandriva Linux Community

The following are various Internet links pointing you to the most important Mandriva Linux-related sources. If you wish to know more about the Mandriva company, connect to our web site (<http://www.mandriva.com/>). You can also check out the Mandriva Linux distribution web site (<http://www.mandrivalinux.com/>) and all its derivatives.

Mandriva Expert (<http://www.mandrivaexpert.com/>) is Mandriva's support platform. It offers a new experience based on trust and the pleasure of rewarding others for their contributions.

We also invite you to subscribe to the various mailing lists (<http://www.mandriva.com/community/resources/newsgroups>) where the Mandriva Linux community demonstrates its vivacity and keenness.

Please also remember to connect to our security page (<http://www.mandriva.com/security>). It gathers all security-related material about Mandriva Linux distributions. You will find security and bug advisories, as well as kernel update procedures, the different security-oriented mailing lists which you can join, and Mandriva Online (<https://online.mandriva.com/>). This page is a must for any server administrator or user concerned about security.

1.2. Join the Club!

Mandriva offers a wide range of advantages through its Mandriva Club (<http://club.mandriva.com>):

- download commercial software normally only available in retail packs, such as special hardware drivers, commercial applications, freeware, and demo versions;
- vote for and propose new software through a volunteer-run RPM voting system;
- access more than 50,000 RPM packages for all Mandriva Linux distributions;
- obtain discounts for products and services on Mandriva Store (<http://store.mandriva.com>);
- access a better mirror list, exclusive to Club members;
- read multilingual forums and articles.
- access Mandriva's Knowledge Base (<http://club.mandriva.com/xwiki/bin/view/KB/>), a wiki-based site which holds documentation on many subjects such as administration, connectivity, troubleshooting, and more;
- chat with the Mandriva Linux developers on the Club Chat (<https://www.mandrivaclub.com/user.php?op=clubchat>);
- enhance your GNU/Linux knowledge through Mandriva's e-training lessons (<http://etraining.mandriva.com/>).

By financing Mandriva through the Mandriva Club you will directly enhance the Mandriva Linux distribution and help us provide the best possible GNU/Linux desktop to our users.

1.3. Subscribing to Mandriva Online

Mandriva offers a very convenient way to keep your system automatically up-to-date, keeping away bugs and fixing security holes. Visit the Mandriva Online Web site (<https://online.mandriva.com/>) to learn more about this service.

1.4. Purchasing Mandriva Products

Mandriva Linux users may purchase products on-line through the Mandriva Store (<http://store.mandriva.com/>). You will not only find Mandriva Linux software, operating systems and “live” boot CDs (such as Move), but also special subscription offers, support, third-party software and licenses, documentation, GNU/Linux-related books, as well as other Mandriva goodies.

1.5. Contributing to Mandriva Linux

The skills of the many talented folks who use Mandriva Linux can be very useful in the making of the Mandriva Linux system:

- **Packaging.** A GNU/Linux system is mainly made of programs picked up on the Internet. They have to be packaged in order to work together.
- **Programming.** There are many, many projects directly supported by Mandriva: find the one which most appeals to you and offer your help to the main developer(s).
- **Internationalization.** You can help us translate web pages, programs and their respective documentation.

Consult the development projects (<http://qa.mandriva.com/>) page to learn more about how you can contribute to the evolution of Mandriva Linux.

2. About this Server Administration Guide

We wrote this manual for users wishing to use their Mandriva Linux system as a server needing general knowledge about the configuration of the most widely-used services. Through graphical tools, we show you how to set up your server(s), run the required services and teach you how to secure your networking environment. The following is an overview of this manual’s contents.

The first part (*Configuration Wizards for Common Services*) looks at the different services you can configure through Mandriva Linux Control Center (“*Server Configuration Wizards*”, page 7). After you’re finished with this chapter, you should be able to configure services such as DHCP, DNS or Postfix.

The next chapter (“*Configuring Masqueraded Clients*”, page 23) covers the configuration of clients masqueraded through a Mandriva Linux system, allowing you to work in interconnected networks using many different platforms such as Microsoft DOS, Windows® 9x, Windows NT® and Windows® XP. In order for this chapter to be useful, you need a well configured LAN since we only focus on the gateway.

The second part (*In-Depth Configuration of Common Services*) explores the different Webmin modules which will help you to configure the most common services.

In the last part (*Applied Theory and Troubleshooting*) we discuss three topics useful as a reference: security and networking theory first, and then a very important chapter detailing how to prevent disasters and how to recover from embarrassing situations.

3. Note from the Editor

In the open-source philosophy, contributors are always welcomed! Updating the Mandriva Linux documentation pool is quite a task. You could provide help in many different ways. In fact, the documentation team is constantly looking for talented volunteers to help us out accomplish the following tasks:

- writing or updating;
- translating;

- copy editing;
- XML/XSLT programming.

If you have a lot of time, you can write or update a whole chapter; if you speak a foreign language, you can help us translate our manuals; if you have ideas on how to improve the content, let us know; if you have programming skills and would like to help us enhance the Borges Collaborative Content Creation System (C3S) (<http://sourceforge.net/projects/borges-dms>), join in. And don't hesitate to contact us if you find any mistakes in the documentation so we can correct them!

For any information about the Mandriva Linux documentation project, please contact the documentation administrator (<mailto:documentation@mandriva.com>) or visit the Mandriva Linux Documentation Project Pages (<http://qa.mandriva.com/twiki/bin/view/Main/DocumentationTask/>).



Please note that since June 2004 the Mandriva Linux documentation and the development of Borges is handled by NeoDoc (<http://www.neodoc.biz>).

4. Conventions Used in this Book

4.1. Typing Conventions

In order to clearly differentiate special words from the text flow, we use different renderings. The following table shows examples of each special word or group of words with its actual rendering, as well as its significance.

Formatted Example	Meaning
<i>inode</i>	Used to emphasize a technical term.
<code>ls -lta</code>	Used for commands and their arguments. (see <i>Commands Synopsis</i> , page 4).
<code>a_file</code>	Used for file names. It might also be used for RPM package names.
<code>ls(1)</code>	Reference to a man page. To read the page, simply type <code>man 1 ls</code> , in a command line.
<code>\$ ls *.pid</code>	Formatting used for text snapshots of what you may see on your screen including computer interactions, program listings, etc.
<code>localhost</code>	Literal data which does not generally fit in any of the previously defined categories. For example, a key word taken from a configuration file.
<code>OpenOffice.org</code>	Defines application names. Depending on context, the application and command name may be the same but formatted differently. For example, most commands are written in lowercase, while applications names usually begin with an uppercase character.
<u>Files</u>	Indicates menu entries or graphical interface labels. The underlined letter, if present, informs you of a keyboard shortcut, accessible by pressing the Alt key plus the letter in question.
<i>Le petit chaperon rouge</i>	Identifies foreign language words.
Warning!	Reserved for special warnings in order to emphasize the importance of words. Read out loud.



Highlights a note. Generally, it gives additional information about a specific area.



Represents a tip. It can be general advice on how to perform a particular action, or hints at nice features, such as shortcuts, which could make your life easier.



Be very careful when you see this icon. It always means that very important information about a specific subject will be dealt with.

4.2. General Conventions

4.2.1. Commands Synopsis

The example below shows the symbols you will see when the writer describes the arguments of a command:

```
command <non literal argument> [--option={arg1,arg2,arg3}] [optional arg ...]
```

These conventions are standard and you will find them elsewhere such as in the man pages.

The “<” (lesser than) and “>” (greater than) symbols denote a **mandatory** argument not to be copied as is, which should be replaced according to your needs. For example, <filename> refers to the actual name of a file. If this name is foo.txt you should type foo.txt, not <foo.txt> or <filename>.

The square brackets (“[]”) denote optional arguments, which you may or may not include in the command.

The ellipsis (“...”) means an arbitrary number of arguments may be included.

The curly brackets (“{ }”) contain the arguments authorized at this specific place. One of them is to be placed here.

4.2.2. Special Notations

From time to time, you will be asked to press, for example, the keys **Ctrl-R**, which means you need to press and hold the **Ctrl** key and tap the **R** character right after as well. The same applies for the **Alt** and **Shift** keys.



We use capital letters to represent the letter keys; this doesn't mean that you have to type them capitalized. However, there might be programs where typing **R** is not the same than typing **r**. You will be informed when dealing with such programs.

Regarding menus, going to menu item File→Reload user config (**Ctrl-R**) means: click on the File text displayed on the menu (generally located in the upper-left of the window). Then in the pull-down menu, click on the Reload user config item. Furthermore you are informed that you can use the **Ctrl-R** key combination (as described above) to get the same result.

4.2.3. System-Generic Users

Whenever possible, we use two generic users in our examples:

Queen Pingusa	queen	This is our default user, used through most examples in this book.
Peter Pingus	peter	This user can be created afterward by the system administrator and is sometimes used to vary the text.

Introduction to Server Wizards and Masqueraded Clients

This part is divided into two chapters : the first details the Mandriva Linux server wizards, while the second one goes deep into the configuration of masqueraded clients.

1. Introducing Server Wizards

This chapter (*“Server Configuration Wizards”*, page 7) will help you to configure services such as DNS (Domain Name Server), DHCP (Dynamic Host Configuration Protocol), Samba, HTTP, FTP, etc. using GUI-based configuration wizards.

These allow you to configure these services simply, in order for them to be smoothly integrated in a small LAN. If you wish to further configure those services, refer to the corresponding Webmin chapters (Part II).

2. Masquerading Clients

We will show you how to configure other machines on a local network in *“Configuring Masqueraded Clients”*, page 23, while using a Mandriva Linux server with masquerading set up as a gateway to the outside world. The information in this chapter covers multiple platforms and architectures.

Chapter 1. Server Configuration Wizards

1.1. Foreword

The Mandriva Linux configuration wizards are designed to configure servers located between a local network and the Internet. They give you the ability to make configurations quickly and efficiently for most common services in a local network, as well as Internet web and FTP services. In this chapter, we assume your network is as shown in figure 1-1, and that Mandriva Linux is installed on the server. Configuring and bringing up the Internet connection is beyond the scope of this chapter.

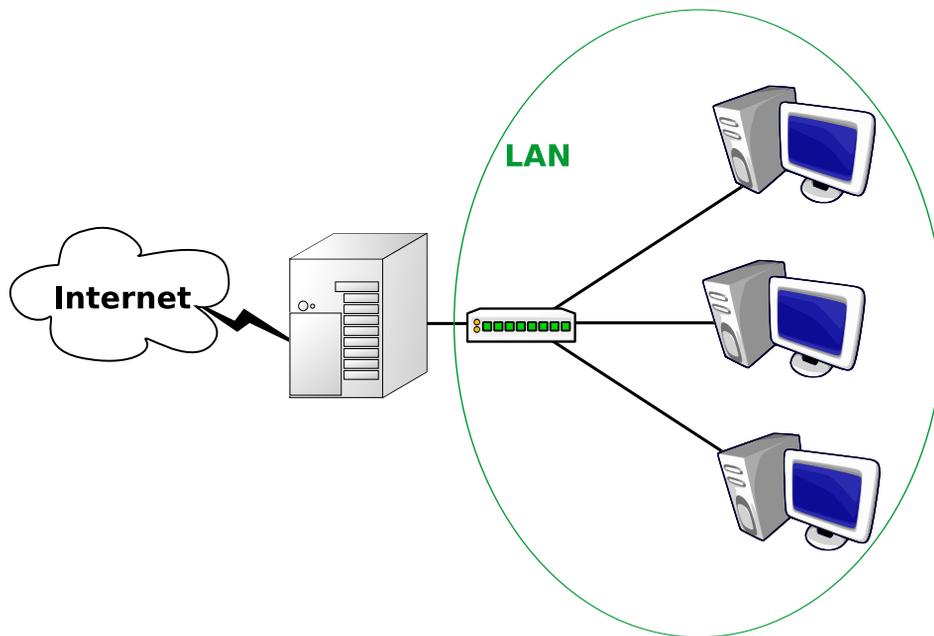


Figure 1-1. An Example of an Internal Network

The server configuration wizards are available through the Control Center when the `drakwizard` package is installed. New categories appear in the Mandriva Linux Control Center, so wizards are organized as follows:



Wizards noted as "Expert mode only" below are only accessible when the expert mode is toggled on (Options→Expert mode).

Sharing

- FTP server (*FTP Server Configuration*, page 13): configure where your FTP server should be reachable from.
- Samba server (*Samba Server Configuration*, page 11): if the server is to act as a file or print server for Windows[®] machines, this wizard helps you set up public shared files and printers, and announce their names to the Windows[®] network.
- Web server (*Web Server Configuration*, page 13): we show you how to set up your GNU/Linux box as a web server, and how to make it reachable;
- HTTP and NFS installation server (*Installation Server Wizard*, page 16): to allow your network client machines to be installed from the server, making CDs and DVDs obsolete. Expert mode only.

Network Services

- DHCP server (*DHCP Server Configuration*, page 8): your server can assign IP addresses dynamically to machines on the network.
- DNS server (*DNS Server Configuration*, page 9): to configure name resolution for machines inside and outside the private network.
- Proxy server (*Proxy Server Configuration*, page 18): configure your server to act as a web proxy cache. This speeds up web browsing while limiting the bandwidth usage on the Internet.
- Time server (*Time Configuration*, page 20): your machine can also supply time to other machines using the NTP protocol (*Network Time Protocol*).

Authentication

- Switching authentication scheme: to set up the local users authentication method: local, LDAP, NIS, Windows Domain. Expert mode only.
- NIS server (*NIS and Autofs Servers Wizard*, page 16): to set up the Network Information System, centralizing users authentication and home directories.
- LDAP server (*LDAP Configuration Wizard*, page 17): to set up a simple LDAP repository to be used as authentication mechanism.

Groupware

- News server (*News Server Configuration*, page 18): you can make your server act as a local mirror of an external news server.
- Mail server (*Configuring Mail with Postfix Server*, page 10): configure your mail domain for sending and receiving mail to and from the outside world.
- Groupware server: This wizard enables you to easily configure the Kolab server, enabling individuals from within an organization to share contacts, calendar, appointments, etc.

You can access the wizards individually by clicking on the corresponding button. In this chapter we describe wizards for the most common services in no particular order. Note that the required packages are installed by the wizard if they are not already available.



For experienced users: wizards are limited to configuring class C networks, and only the basic configuration is handled for each service. This should be enough for most situations, but if you wish for a more fine-tuned configuration, you must edit the configuration files by hand or by using another administration tool such as Webmin.

1.2. DHCP Server Configuration

DHCP stands for *Dynamic Host Configuration Protocol*. This protocol allows for machines connecting to your local network to be automatically assigned all relevant network parameters such as an IP address, the addresses of the name servers and the address of the gateway.

Figure 1-2. DHCP Server Address Range

All you have to do is specify the range of addresses¹ that you want to have available via DHCP, as shown in figure 1-2. If your server has more than one NIC, you are asked which interface the DHCP server must listen for requests on: choose the one connected to your LAN. If you wish that client computers can access the Internet, you need to provide here the IP address of the gateway. In case the DHCP server is also the gateway for your LAN fill it with the server's LAN address (for example: 192.168.0.1).



Check the Enable PXE option if you want your machine to act as an installation server for multiple machines on your LAN.

1.3. DNS Server Configuration

DNS stands for "Domain Name System". It allows you to specify a machine by its name instead of its IP address. This wizard allows you to setup a basic DNS server, master or slave.

Make sure you have a FQDN host name set for your system, otherwise the DNS wizard will refuse to start. You are given the option to run one of the following wizards:

Master DNS Server

Setup your machine as a plain DNS server. The first step allows you to provide the address of an external DNS server to which the requests that the local server cannot answer will be forwarded. It is generally the address of your ISP's DNS server.

During the second step you can specify domain names for lookups. For example if you request the IP of a machine called `kenobi`, the server adds the domain names you add here to perform the request.

1. Addresses outside this range are available for machines which need static addresses. Those machines can then be listed in the DNS configuration (*DNS Server Configuration*, page 9).

Slave DNS Server

Setup your machine as the slave server of another, master, DNS server. You only need to supply the IP address of the master server for the slave to mirror. Then clients can be configured to query both servers: if the master fails, the slave takes relay.

Add Host in DNS

If your machine is a master DNS server, you'll be able to declare all the machines with static addresses on your network so that the DNS server can answer requests about them.

Remove Host in DNS

This is used to remove a DNS entry previously entered with Add Host in DNS.



Both the Add Host in DNS and Remove Host in DNS wizards only work if your machine is set up as a master DNS server.

1.4. Configuring Mail with Postfix Server

This wizard helps you configure your internal and outgoing mail. When configured, this SMTP server will allow users of your local network to send internal and external mail through it. Likewise, if your server is referenced on the Internet public DNS as a MX server for your domain name, then it also receives and manages mail from the Internet addressed to users of your domain. In this case, make sure you open the corresponding ports in your firewall.



Your server network parameters must not be provided by DHCP for Postfix to work properly.

The first step consists of choosing whether you intend to use an external SMTP relay or not. If you can use one provided by your ISP then choose Relay mail server in the drop-down list. Otherwise, choose Main mail server. In the procedure below only the second step differs from one configuration to the other.

1. Global Postfix Configuration

Smtpd banner

The banner your server advertises when talking to other servers or clients.

Hostname

The name of your server.

Domain

The domain handled by this mail server.

Origin

The domain name that locally-posted mail appears to come from, and that locally posted mail is delivered to.

2. Relay (for Relay mail server only)

Relay host

This is where you define the mail server of the ISP responsible for relaying your outgoing messages.

Relay domains

What destination domains (and subdomains thereof) this system relays mail to. Mails sent to a domain other than the local domain that are not part of the relay domains are rejected (to prevent spam).

3. Main server Configuration (for Main mail server only)**helo required**

For security reasons you might require remote clients to identify themselves before starting communication. Choose yes in that case.

Disable verify command

The `verify` command can be used by a client to verify a specific user is actually handled by a mail server. You can disable it to prevent email harvesting by spammers.

Masquerade domains

This field is used to masquerade the domain from which internal mail appears to come from. For example: `foo.example.com example.com` directs Postfix to masquerade `toto@foo.example.com` to `toto@example.com`.

4. Message options

A few options affecting message handling you can leave at their default values.

Maximal queue life

If a message cannot be delivered after this delay it is sent back as undeliverable.

Message size limit

Messages larger than this size (bytes) are rejected.

Delay warning time

If a message cannot be delivered, the sender will receive a warning after this number of hours.

5. Network Configuration**inet interfaces**

The network interface addresses that this mail system receives mail on. By default only the local interface is listened on. Specify `all` to receive mail on all network interfaces.

my destination

The list of domains that are delivered via the local mail delivery transport. The SMTP server validates recipient addresses and rejects non-existent recipients.

my networks

The list of “trusted” SMTP clients who have more privileges than “strangers”. In particular, “trusted” SMTP clients are allowed to relay mail through Postfix. Specify a list of network addresses or network/netmask patterns, separated by commas and/or whitespace.

If a parameter is not clear to you, please refer to the Postfix Configuration Parameters (<http://www.postfix.org/postconf.5.html>).

1.5. Samba Server Configuration

Samba allows GNU/Linux to act as a file and/or printer server for Windows[®] machines. Even though this wizard can help configure primary and backup domain controllers, we will concentrate here on the most common, standalone server configuration.

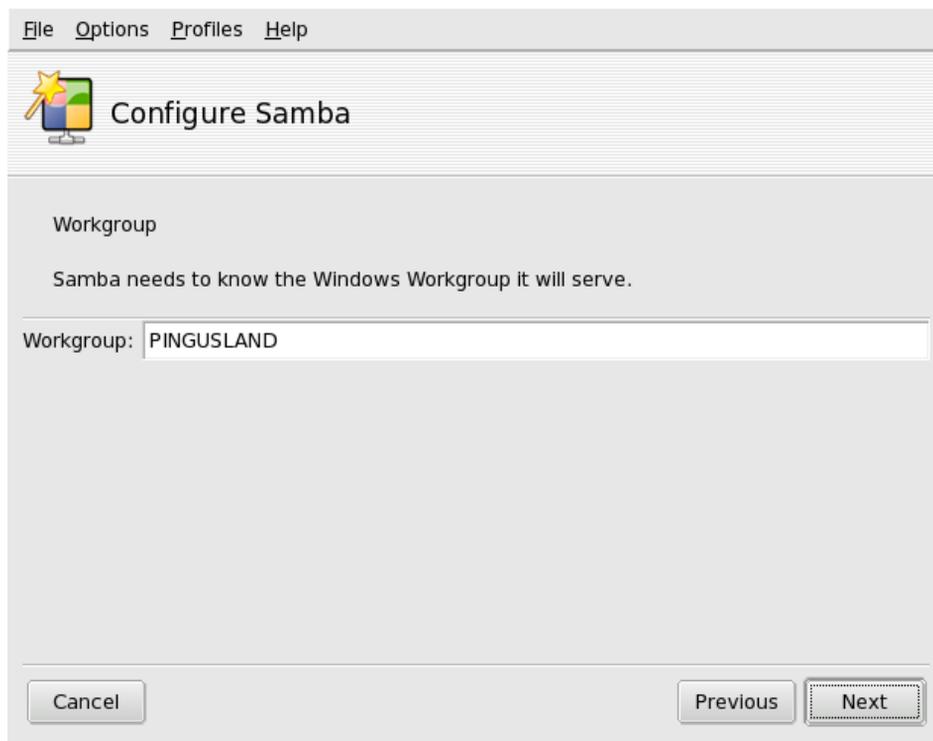


Figure 1-3. Choose the Work Group for your Shares

Now you must enter the work group for which these shares will be available (figure 1-3). You can either create a new work group or choose an existing one, but if you don't know what to do, please refer to your system administrator.



Figure 1-4. Naming your Samba Server

Then you will have to specify the name by which your Mandriva Linux server will be known to Windows[®] machines, as shown in figure 1-4. You may choose whatever name you want.

Finally you can adjust the log facility parameters. Keep the default unless you have specific needs.

When the Samba server is configured you can run `drakwizard sambashare` from a command line as root to create new shares and manage existing ones.

1.6. Web Server Configuration

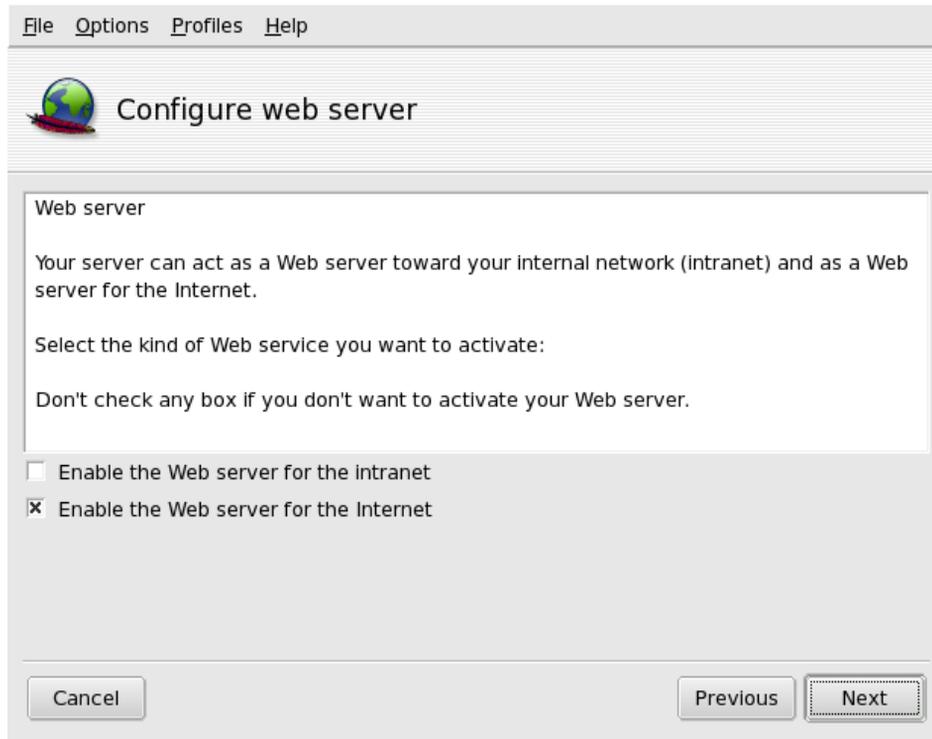


Figure 1-5. Defining Where your Web Server Should Be Visible from

This wizard will simply let you specify if your web server will be disabled, visible from the local network, from the external network (generally the Internet) or from both. Check the appropriate boxes as shown in figure 1-5.



If your server network parameters are provided by DHCP the web server might not function properly if set to be visible from the Internet.

The second step allows you to enable the feature which gives users the option to maintain their own web sites, available from the `http://servername/~user/` URL. The directory where they store their sites can also be changed if this option is checked.

The last step allows you to specify the directory where the files to be served by the server will be stored, known as the **Document Root**. To begin creating your web site, simply put the files in the chosen directory. You can connect to your web site as soon as the wizard is finished through the `http://localhost/` URL.

1.7. FTP Server Configuration

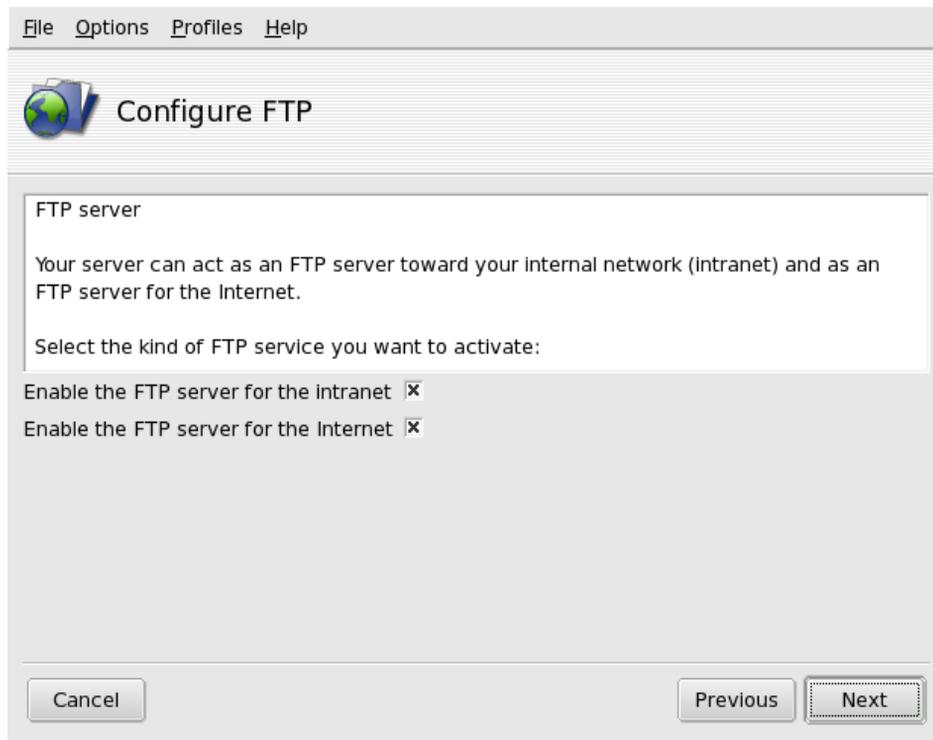


Figure 1-6. Defining Where your FTP Server Should Be Visible from



If your server network parameters are provided by DHCP the FTP server might not function properly if set to be visible from the Internet.

This wizard resembles the one used to configure a web server: It will let you specify whether FTP should be disabled, visible from the local network, from the external network, or both. A sample window is shown in figure 1-6.

Figure 1-7. FTP Server Configuration

This is the basic FTP server configuration, you are advised to provide an email address for the administrator so that he can receive alert messages.

Admin e-mail

Enter the address to which messages regarding the FTP server should be sent.

Permit root login

Check this box if you wish the root user to be allowed to login into the FTP server. If the FTP authentication is made in clear text, this option is **not** recommended.

Figure 1-8. FTP Server Options

You are then allowed to configure a few options:

FTP Port

The standard FTP port is 21, if you specify something else here FTP clients have to be configured accordingly.

Chroot home user

By checking this option, users who log into the FTP server will be “boxed inside” their home directories.

Allow FTP resume

If your server is likely to host large files, it might be prudent to allow users to resume downloads.

Allow FXP

Check this option if you want your server to be able to exchange files with other FTP servers. Please note that the FXP protocol is not very secure.

To begin populating your anonymous FTP server, simply put the files in the `/var/ftp/pub/` directory. You can connect to your FTP server as soon as the wizard is finished through the `ftp://localhost/pub` URL. Home directories are also accessible by default with local password authentication. If queen wants to access her home directory she has to use the `ftp://queen@localhost` URL.

1.8. Installation Server Wizard

You are performing lots of installations and are tired of changing CDs? This wizard is for you. It configures your machine to act as an installation server, so new machines can get all needed packages directly on the network, either for initial installation or for maintenance.

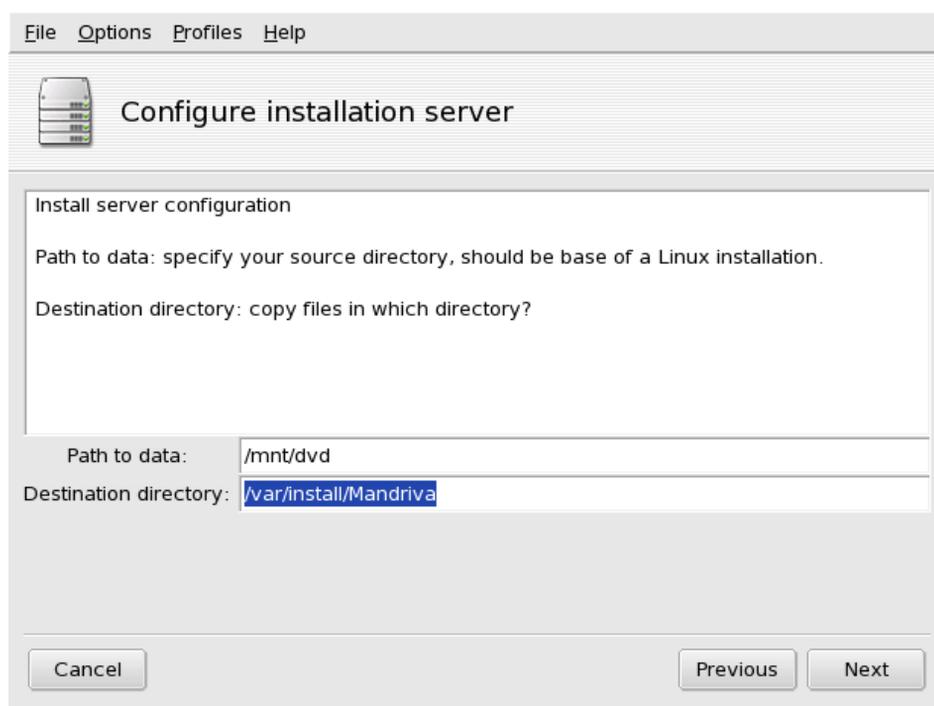


Figure 1-9. Copying Installation Sources

Specify the location to copy the CDs or DVD from, and a place on your disk where the files are to be stored.

1.9. NIS and Autofs Servers Wizard

Run this wizard if you wish to centralize your users' authentication and home directories. Doing so allows users to connect from any machine on the local network and have access to their own environment.

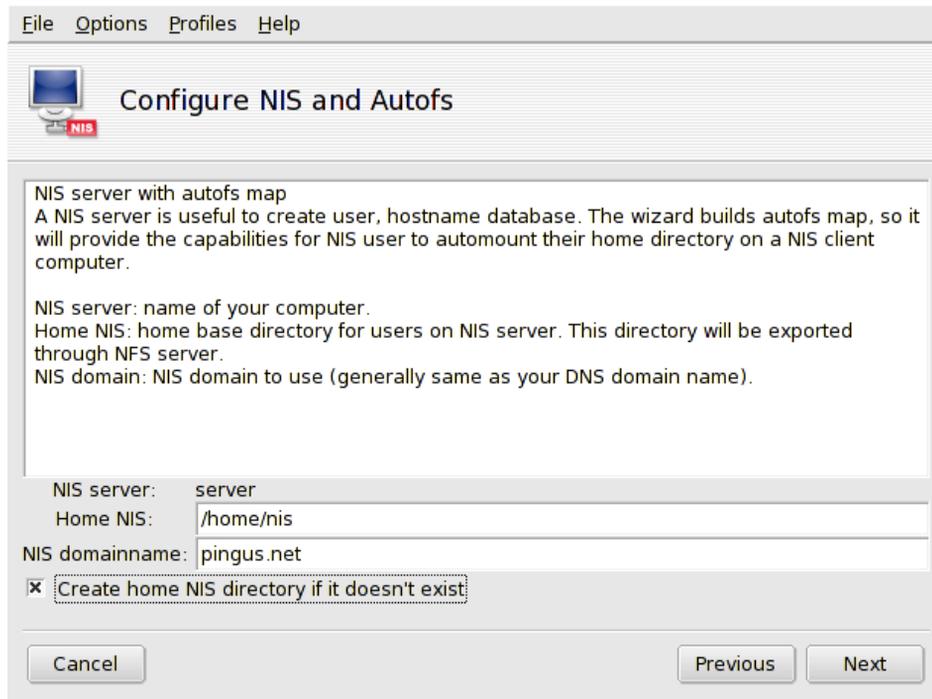


Figure 1-10. Setting NIS Server Parameters

Simply fill in the fields as directed by wizard. When configuration is done, NIS users can login from any machine on the network that is setup to connect on your NIS server. Additionally, those users have their home directories automatically mounted locally.

1.10. LDAP Configuration Wizard

This simple wizard allows you to setup the basic configuration of a LDAP server, and add users to it. This is useful to quickly setup a LDAP-based authentication mechanism.

The first time you run the wizard, you are presented the server configuration dialog.

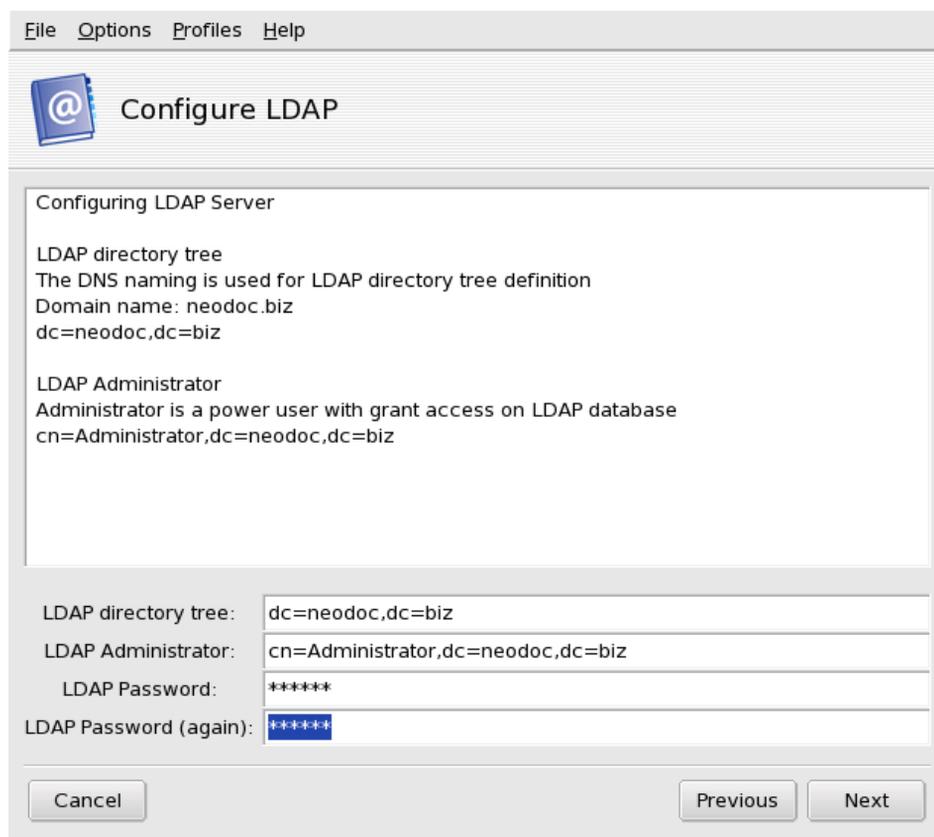


Figure 1-11. LDAP Server Configuration

When the configuration is set and the server launched, running the wizard again will suggest some options:

Show Ldap configuration

Shows current server configuration, useful to configure possible LDAP clients.

Delete Ldap configuration

Removes current server configuration and stops the server.

Add user in Ldap server

Starts a little wizard which allows you to add new users inside the users directory.

1.11. News Server Configuration

This wizard will configure a news gateway: your server will be able to fetch newsgroups from an external news server (usually your ISP's, named something like `news.domain.of.your.isp`) and make them visible to your internal network. Therefore the first step is to specify which external news server you want to use.

Then you need to specify the interval (in hours) between each refresh. You need to find a balance here: specifying too high an interval might make your server's news obsolete rapidly while specifying too low an interval might elevate your server's (inter)network load considerably. Of course, all will depend on the traffic of the news sites being mirrored and on the Internet connection speed of your server.

1.12. Proxy Server Configuration

A proxy server is very useful for a local network accessing many web pages across a slow, or relatively slow, connection. It maintains a cache of most visited pages so that they don't need to be retrieved again from the Internet if requested by different users. Squid is the proxy server used.

First of all you need to choose a port for the proxy to listen to requests on. Users will have to configure their web browsers to use this port as the proxy port and your server name as the proxy server.

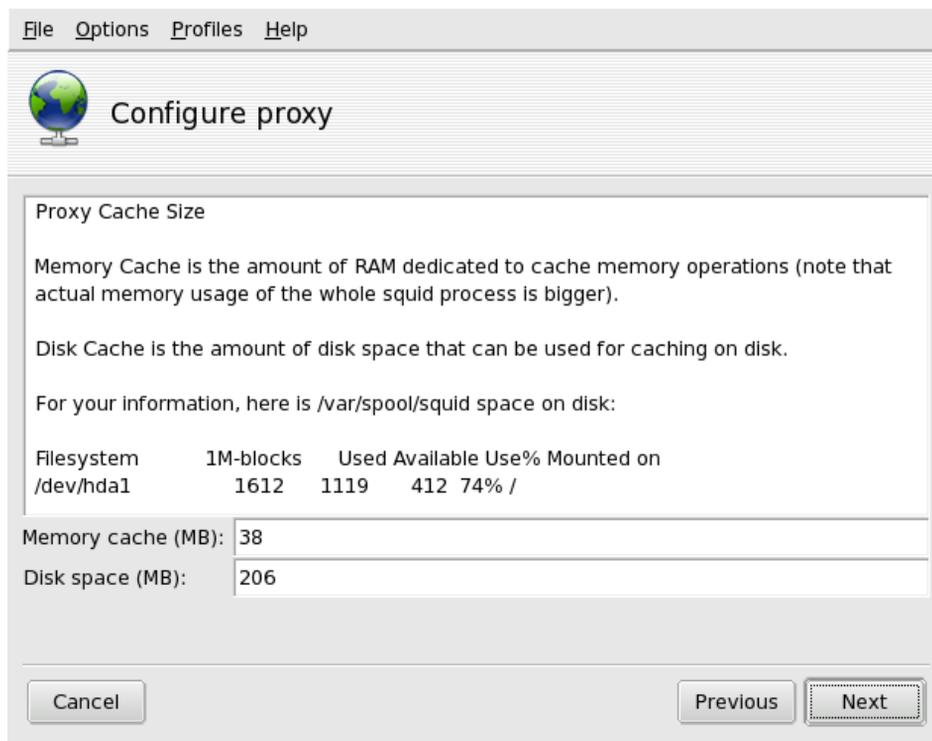


Figure 1-12. Choose the Cache Size

Depending on your machine's available memory, you can allocate more or less to the Proxy. The bigger the memory cache, the fewer disk accesses on the server. Depending on your available disk size you can allocate more or less room for cached pages. The more the space, the less accesses to the Internet. The wizard will choose appropriate values for your system, if in doubt just accept the proposed ones.

In the next step, some access levels are available for clients wishing to use the proxy:

- **All.** There is no restriction, all computers are granted access to the cache. This setting is not very secure and thus not recommended.
- **Local Network.** Only machines on the local network can access the proxy. This is the recommended setting.
- **localhost.** Only the local machine, the server, can access its own proxy.

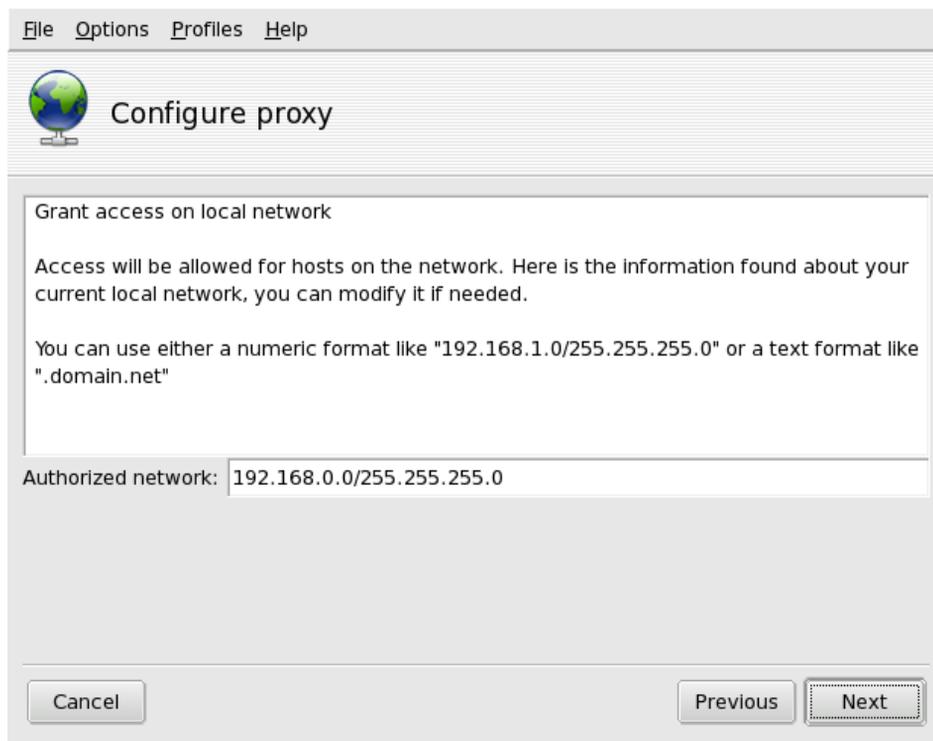


Figure 1-13. Restrict Access to a Particular Sub-network

If you have previously chosen the Local Network access policy, you can choose to restrict even more the access to a particular subnetwork or domain. The wizard will detect your LAN's network address and will offer it by default: make modifications if needed.

Finally, if your server itself has access to another larger proxy connected to the Internet, you can choose to Define an upper level proxy to which requests will be forwarded. If so, the next step will ask you for the name and port of that server.

1.13. Time Configuration

This wizard lets you set up a time server, using the NTP protocol, for your internal network. When you have set up the external time servers your own server will use to correct its internal clock, machines on your local network will be able to get the correct time from your server.

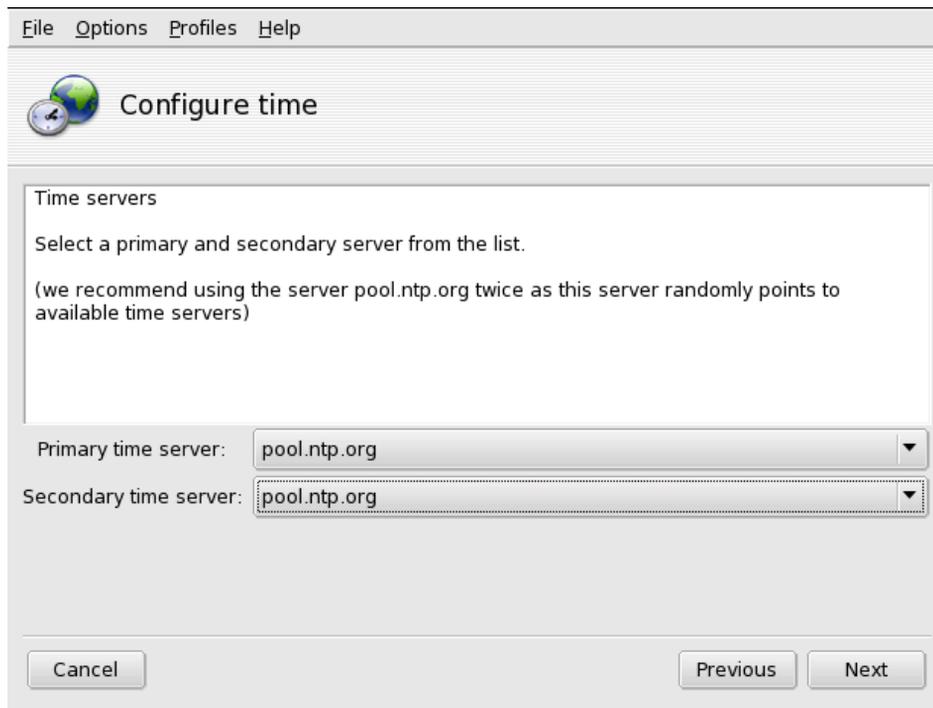


Figure 1-14. Choosing your Time Servers

Choose the time servers to query, in order of preference. It is advisable to keep the default suggested ones, otherwise choose servers which are geographically close to you. By default the time zone will be set to the one you have chosen during the server's installation.

Chapter 2. Configuring Masqueraded Clients

This chapter will show you how to make different operating systems use a Mandriva Linux box with masquerading set up as a gateway to the outside world. The configuration tests proved successful on many operating systems and architectures. Basically, any OS that has a working TCP/IP stack will work.

We will show you how to configure a few operating systems. If you are worried that your OS might not be supported, a simple way to proceed is to “just tell the OS which machine to use as a gateway”. Note that our main focus here is the **gateway** side of the network, we won’t touch on DNS, file sharing or connection schemes problems. 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 have a class C network: your different machines all have IP addresses like 192.168.0.x, with a netmask set to 255.255.255.0, and use eth0 as the network interface. We also take for granted that your gateway’s IP address is 192.168.0.1, and that your machines can each “talk” to the gateway (you can test the latter with the ping command or its equivalent in your environment).

2.1. Linux Box

2.1.1. For Any Linux Box

Here, you need to edit a configuration file. The method is different when automatic IP configuration is used.

Automatic IP configuration

Open the network interface configuration file (on a Mandriva Linux machine it is /etc/sysconfig/network-scripts/ifcfg-eth0, it may be different on yours) and make sure the BOOTPROTO parameter reads BOOTPROTO=dhcp.

Manual Configuration

You need to edit the network configuration file (on a Mandriva Linux machine it is /etc/sysconfig/network, it may be different on yours). Open it with your usual text editor, then add the following lines (adjusting values if necessary):

```
GATEWAYDEV=eth0
GATEWAY=192.168.0.1
```

You may now restart your Linux network layer. On a Mandriva Linux system this is done by executing, as root, service network restart or /etc/init.d/network restart from a terminal window.

2.1.2. On a Mandriva Linux Box



Just put the correct parameters into the Manage Connections tool in Mandriva Linux Control Center’s Network & Internet section to set the configuration automatically. Refer to *Network and Internet Connection Management* of the *Starter Guide*. You can choose to configure all network parameters manually or automatically (using DHCP).

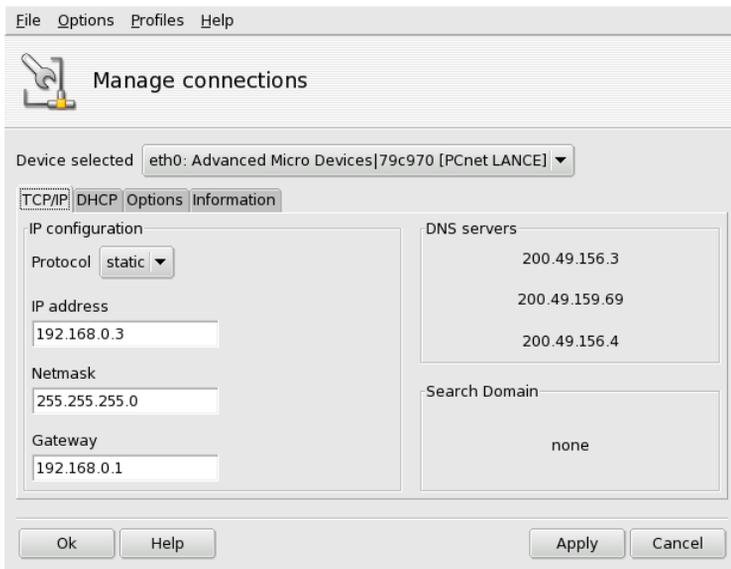


Figure 2-1. Manually Specifying Network Parameters in drakconnect

If you have a DHCP server on your local network, simply choose the DHCP entry. If you have a static IP address for your machine, enter it into the first field after making sure that the static entry is selected. Also fill the Netmask and Gateway fields according to your network configuration, as shown in figure 2-1.

Once you have applied your changes, your network is properly configured and ready to run. The configuration is now permanent.

2.2. Windows XP Box

We assume that you already have a configured network connection. figure 2-2 shows the different steps to get to the desired dialog.

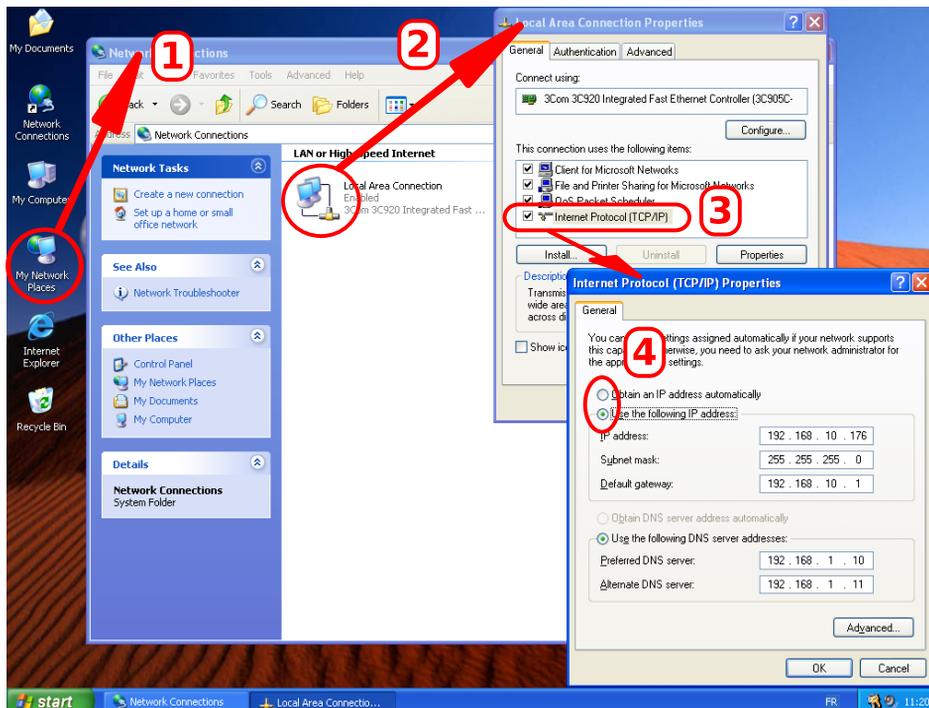


Figure 2-2. Setting up the Gateway with Windows XP

Here are the actions to step from one window to another:

1. On the desktop, right-click on the My network places icon, and select Properties in the menu that appears.
2. In the Network Connections window, do the same with the connection linked to the network where the gateway is located.
3. In the next dialog, select the Internet Protocol (TCP/IP) entry and click the Properties button.
4. In this dialog, you can choose to select the Obtain an IP address automatically option if you have a DHCP server on your network. Then, all other network parameters should also be automatically configured. If you prefer to manually set up network parameters, select the Use the following IP address option and fill in the associated fields.

2.3. Windows 95 or Windows 98 Box



Start by going in the Control Panel (Start+Settings→Control Panel) and find the network icon. Double-click on it: the network configuration panel comes up.

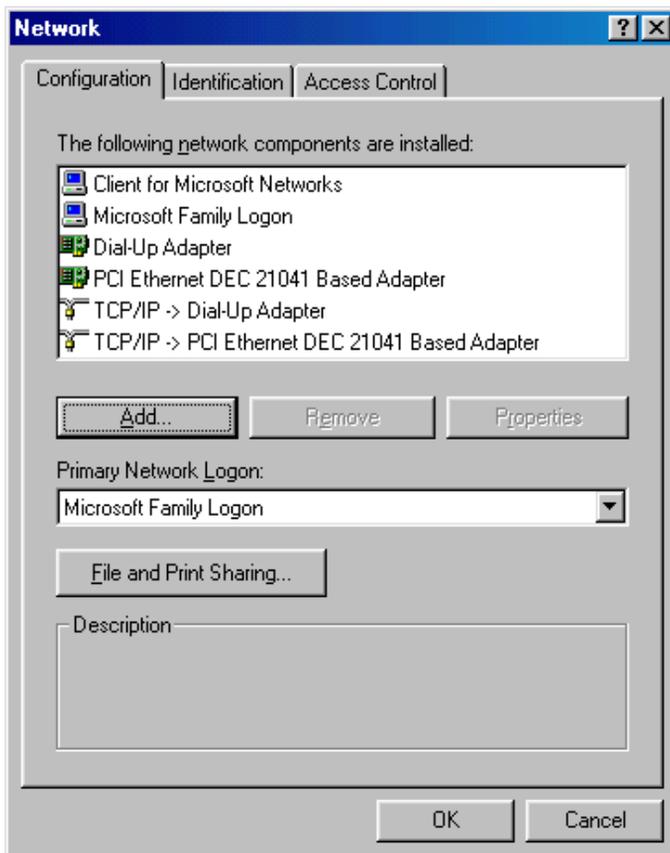


Figure 2-3. The Network Configuration Panel under Windows 9x

In the displayed list, you should find a protocol named TCP/IP and bound to your LAN adapter. If not, you will have to refer to your system documentation to find out how to install it. Select it and click on the Properties button.

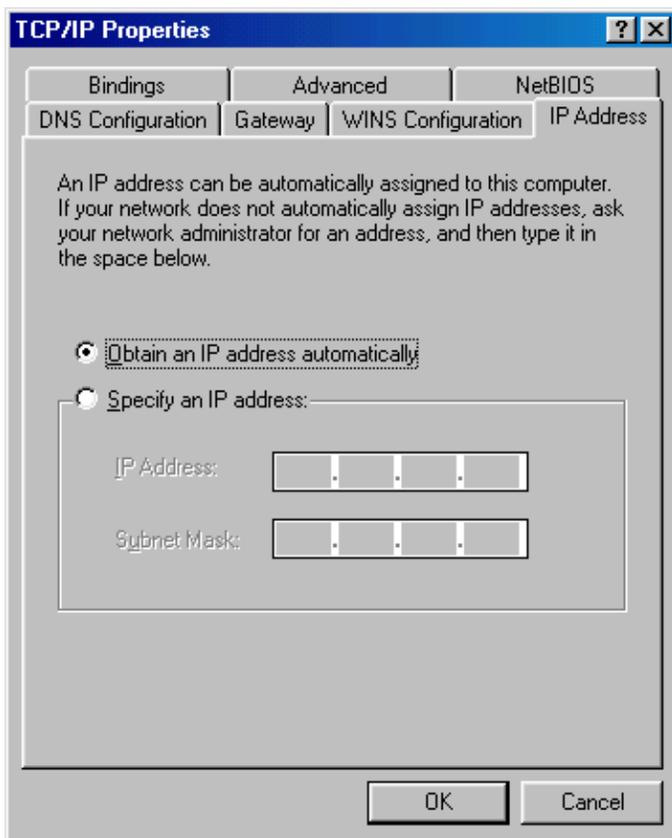


Figure 2-4. The TCP/IP Configuration Panel under Windows 9x

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 network parameters). Click on the Gateway tab.

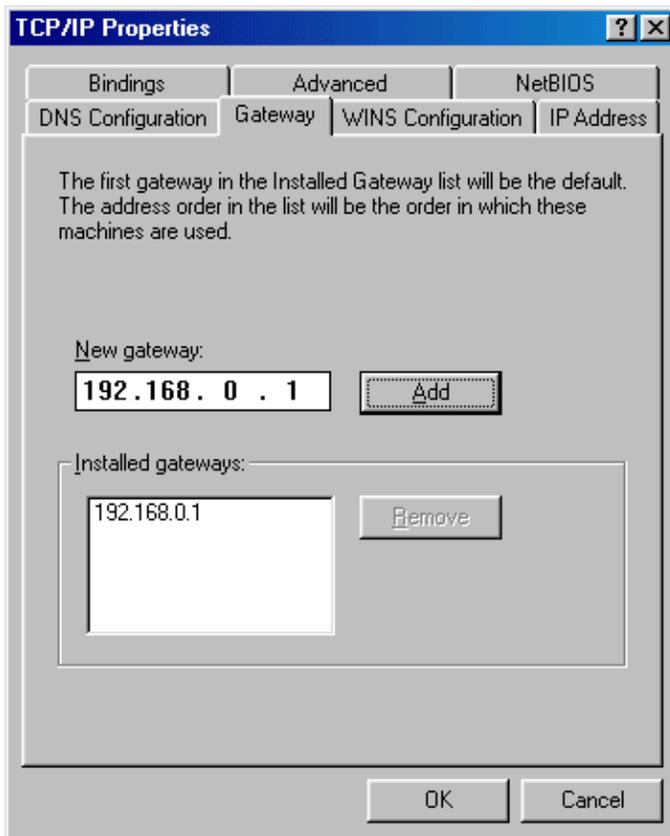


Figure 2-5. The Gateway Configuration Panel under Windows 9x

The rest is child's play! Fill in the blanks with your gateway's IP address (i.e. 192.168.0.1, in our example). Click the Add and then on the OK buttons.

You will need to reboot your computer, of course. Once this is done, check to see if you can reach the rest of the world.

2.4. Windows NT or Windows 2000 Box

To configure these OSes, follow these simple steps:

1. Go to Control Panel+Network→Protocols.



Figure 2-6. The Protocol Configuration Panel under Windows NT/2000

2. First, select the TCP/IP Protocol in the list of network protocols. Then, click on the Properties button, and select the network card connected to the local network (see figure 2-7). In this example, we show a configuration with DHCP: the Obtain an IP address from a DHCP server option is selected.

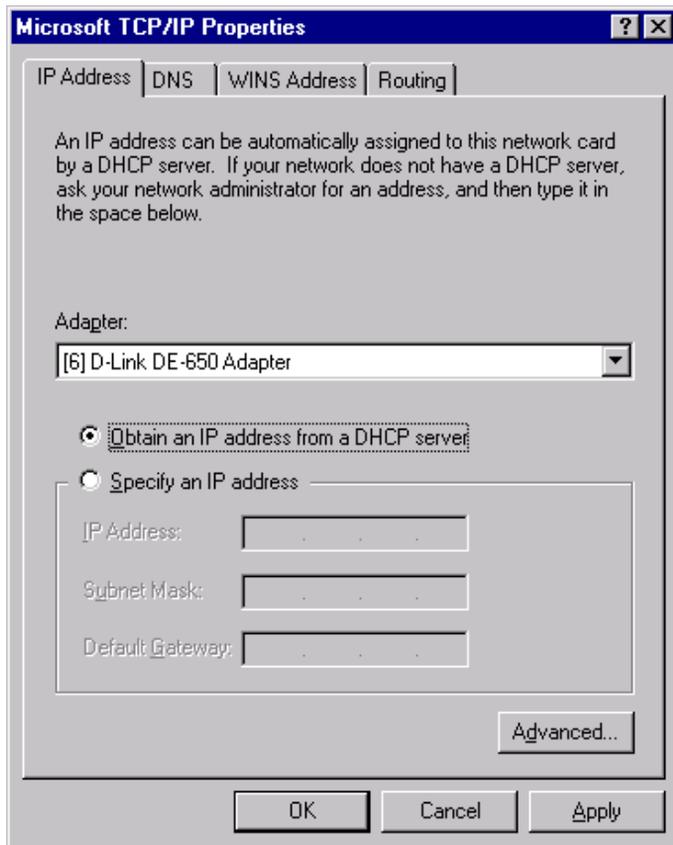


Figure 2-7. The Network Software Panel under Windows NT/2000

If this is your case, you just need to confirm these choices and reboot. Otherwise, proceed as follows.

3. To set network parameters manually, begin by selecting the Specify an IP address option (see figure 2-8).

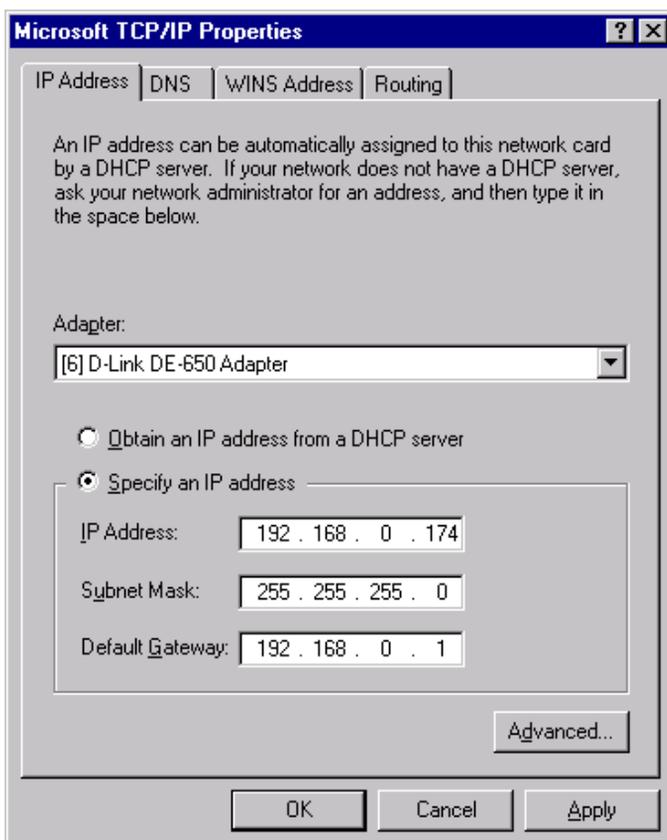


Figure 2-8. The TCP/IP Configuration Panel under Windows NT/2000

Select the appropriate adapter, and specify, if not already entered, the IP address and network mask.

4. Fill in the Default Gateway field with 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 in figure 2-9.

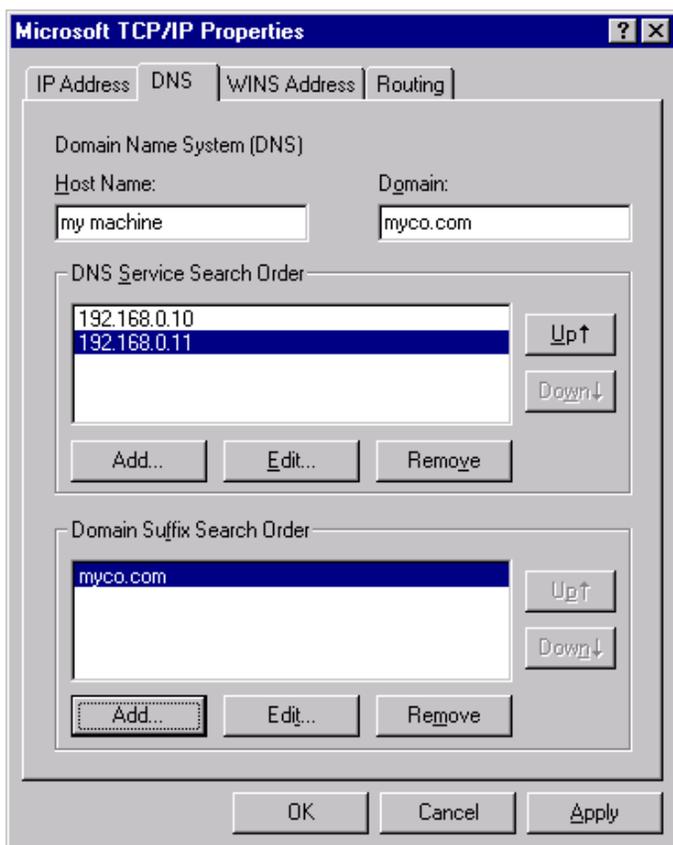


Figure 2-9. The DNS Configuration Panel under Windows NT/2000

You must also provide a host name and an associated domain name.



Unless you know exactly what you are 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 servers;
- do not place a check in the Enable IP Forwarding field unless your NT/2000 machine is used for routing and, once again, you know exactly what you are doing;
- 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.

2.5. DOS Box Using the NCSA Telnet Package

In the directory which hosts the NCSA package, you will find a file called `config.tel`. Edit it with your favorite editor and add the following lines:

```
name=default
host=your_linux_host_name
```

```
hostip=192.168.0.1
gateway=1
```

Change `your_linux_host_name` to the real host name of your Linux gateway.

Now save the file, try to `telnet` your Linux box, then to a machine somewhere out there...

2.6. Windows for Workgroups 3.11

The TCP/IP 32b package should already be installed. Go to the Main+Windows Setup+Network Setup→Drivers menu entry 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 *Windows NT or Windows 2000 Box*, page 27.

2.7. MacOS Box

2.7.1. MacOS X

The configuration consists of setting the correct parameters for the Ethernet interface connected to your gateway.



Figure 2-10. MacOS X Dock

First of all, you need to open the System Preferences window by clicking on its icon on the system dock.

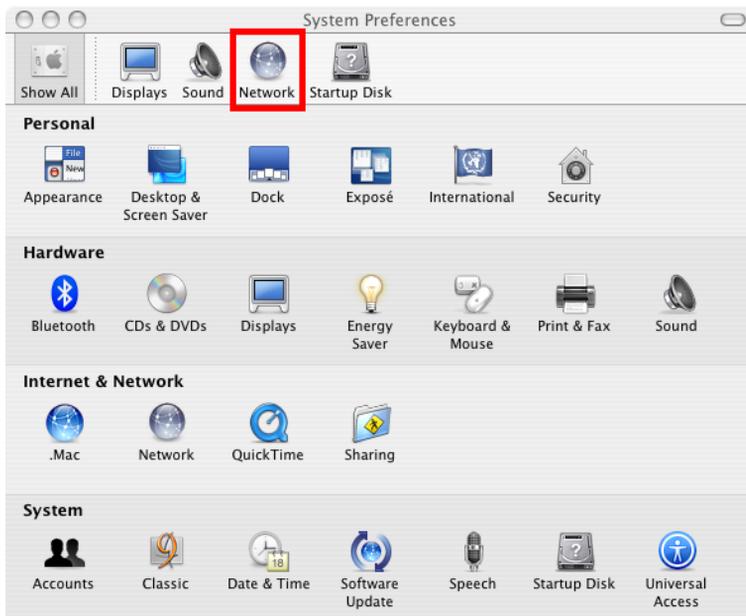


Figure 2-11. MacOS X System Preferences

2.7.1.1. With an Automatic DHCP Configuration



Figure 2-12. Automatic Configuration of Internet Access For MacOS X

In the dialog that appears, Select the Built-In Ethernet or the interface actually connected to the gateway, and then select Using DHCP in the TCP/IP tab as shown in figure 2-12. Then click on Apply Now, and if all goes well, the Router field should now show the IP of your gateway.

2.7.1.2. For a Manual Configuration

If you do not have a DHCP server on your local network, follow this procedure:

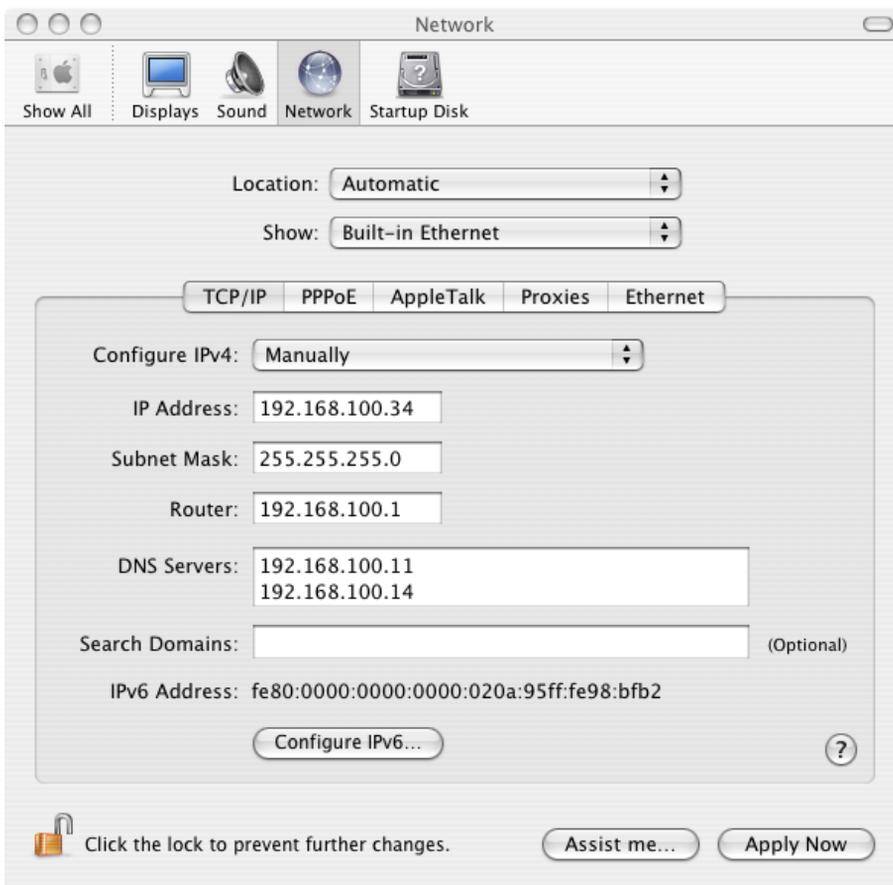


Figure 2-13. Manual Configuration of Internet Access For MacOS X

In the dialog that appears fill the fields as shown here:

- Configure: Manually;
- IP address: 192.168.100.34 (The client machine IP address);
- Subnet Mask: 255.255.255.0 (The local network subnet mask);
- Router Address: 192.168.100.1 (the address of the gateway server);
- DNS Servers: 192.168.100.11; 192.168.100.14 (The IP addresses of the DNS servers)



The name server's addresses may be the addresses of the internal DNSs or those of your Internet Service Provider's servers.

When this is done, click on Apply Now, and if all goes well, you should be able to surf the Internet.

2.7.2. MacOS 8/9

First of all, you need to open the TCP/IP Control Panel as shown below in the Apple menu.



Figure 2-14. Accessing The TCP/IP Control Panel

2.7.2.1. With an Automatic DHCP Configuration

If you have configured your firewall to be a DHCP server, follow this procedure, otherwise go to the next section.

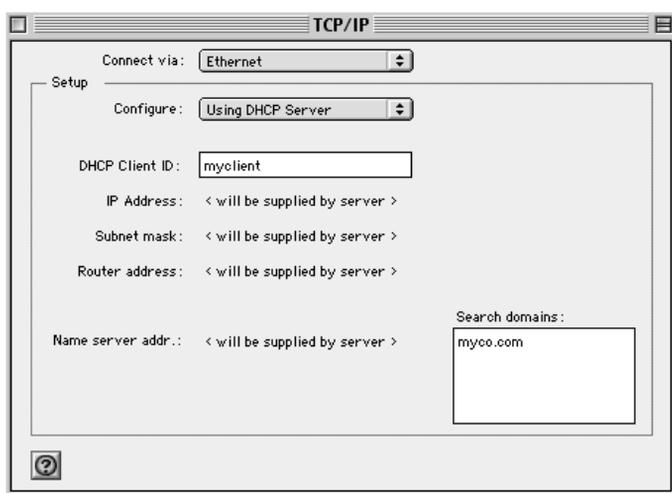


Figure 2-15. Automatic Configuration of Internet Access For MacOS

In the dialog which appears, fill the fields as shown here:

- Connect via: Ethernet;
- Configure: Using DHCP server;
- DHCP Client ID: 192.168.0.1.

2.7.2.2. For a Manual Configuration

If you do not have a DHCP server on your local network, follow this procedure:

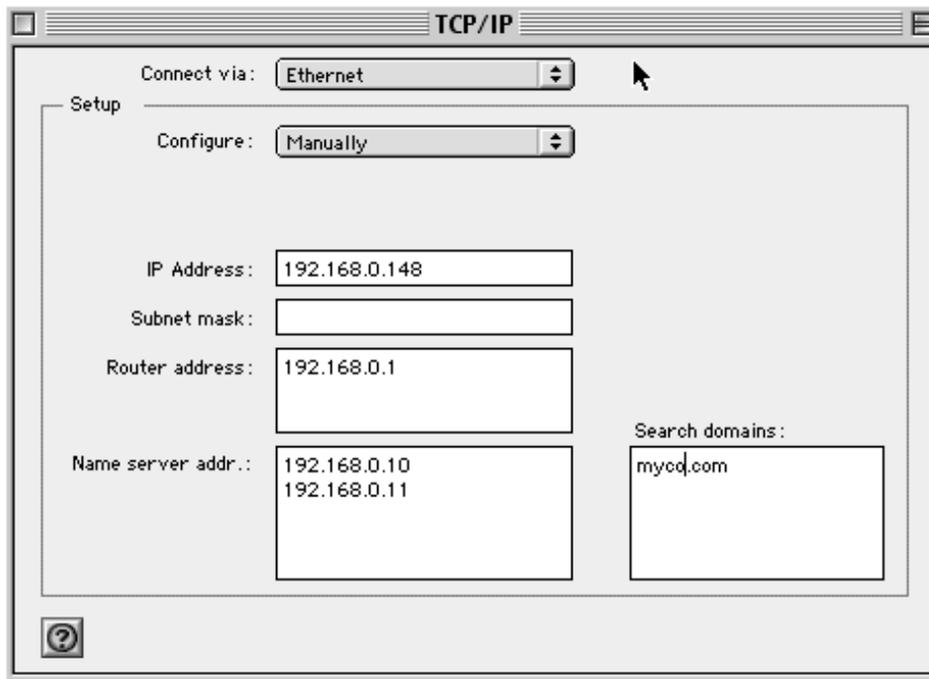


Figure 2-16. Manual Configuration of Internet Access For MacOS

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



The name server's addresses may be the addresses of the internal DNS or those of your Internet Service Provider's servers.

2.7.3. 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.

2.8. OS/2 Warp Box

The TCP/IP protocol should already be installed. If not, install it.

1. Go to Programs, then TCP/IP (LAN), then TCP/IP Settings.
2. Under Routing, choose Add. In Type, select default.
3. Fill the Router address field with the address of your Linux box sharing the Internet connection (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.

Introduction to Services Configuration

This part details the most common services a system administrator may need for both Internet and Intranet use. We document the services which should be of interest for mid-sized companies. Almost every service is configured using the Webmin tool.

1. Introduction to Webmin

The Webmin tool allows you to administer your machines remotely through a web interface using only a web browser which supports the HTTPS (HTTP over SSL) protocol. This facilitates easy and in-depth remote administration, while ensuring security.

This makes Webmin ideal for system administrators because all major platforms have web browsers which meet or exceed the above requirements. Moreover, Webmin has its own “web server” so it doesn’t need 3rd party software (such as a web server) to work. Everything is included.

1.1. Accessing and Using the Interface

First of all, make sure the webmin package is installed. You should also make sure it’s running through the drakxservices application. Once this is done, you can access it through any web browser, either locally or from another machine in the same local network. Point your browser to `https://ServerNameOrIP:10000` to access the interface. Accept the certificate and you will be presented with a login screen. Enter the `root` login and password, and click on Login to access the main Webmin screen.



Figure 31. Webmin Interface

The interface presents tabs or sections, each tab giving access to a number of modules specializing in a particular configuration aspect of the machine.

As a useful exercise, you should now go to the IP Access Control option of the Webmin Configuration module.

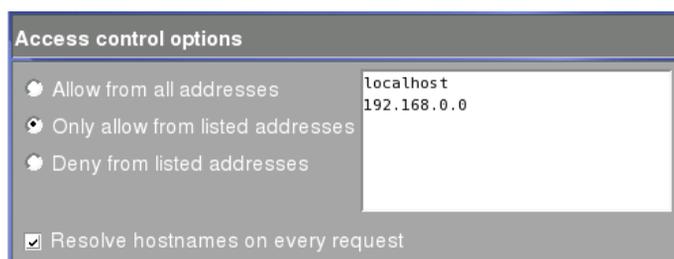


Figure 32. Webmin Access Control

Note that if your server is accessible from the Internet, this step is mandatory to narrow the possibility of people tampering with your system. Of course that doesn't prevent you from setting up a firewall through Mandriva Linux Control Center, Webmin (Networking, Linux Firewall) or any other firewall setup application.

Select the Only allow from listed addresses option. Then list all the machines and/or networks allowed to connect to Webmin in the text field, one entry per line. For increased security also check Resolve hostnames on every request.

1.2. Getting Help

Many module parameter names are available as hyperlinks. Clicking on them opens a pop-up window containing on-disk help concerning that parameter. It may give the meaning, the syntax to use and possibly examples.

Also, most modules present a Search Docs link on the upper-right corner of the screen. Clicking on the link launches a search on the module's name in local documents, Google™ (<http://www.google.com>), HOWTO documents, Package documentation, and more. This way you have instant access to relevant documents for the module you are currently configuring. You can have access to the search form and options in the System+System Documentation module.

The Webmin home site (<http://www.webmin.com/index2.html>) also offers various resources, including a link to Joe Cooper's very good The Book of Webmin (<http://www.swelltech.com/support/documentation.html>).

Finally, at the end of each service chapter, we include a list of interesting resources to get additional information on the configuration of that particular service.

2. Services

The services covered in this part are:

- Domain Name System (DNS). We discuss the BIND name server in *"BIND DNS Server"*, page 39.
- Internet/Intranet web site hosting (HTTP). The Apache web server is discussed in *"Internet and Intranet Web Server"*, page 47.
- The mail management (SMTP) chapter (*"Postfix Mail Server"*, page 53) focuses on sending mail with the Postfix mail server.
- We talk about mail delivery services using the POP and IMAP protocols with the IMAP mail server (see *"Mail Delivery Services: Pop and IMAP"*, page 59).
- Sharing resources such as files and printers is the main topic of the next chapter (*"Resource Sharing"*, page 61), using NFS, Samba or ProFTPD FTP server.
- We document the server part of the Kroupware groupware solution in *"The Kolab Server"*, page 69.
- We detail the usage of the MySQL database server in the next chapter (*"MySQL Database Server"*, page 79).
- Distributed user management or home hosting (NIS) is the main subject of the final chapter of the services configuration section (*"NIS Client and Server"*, page 83).

Chapter 3. BIND DNS Server

A DNS server allows you to associate an IP address to a name and vice versa. For example: `www.mandriva.com` (“Name”) is currently associated to `212.85.150.181` (“Address”). To make a comparison, a name server acts somewhat like a telephone directory: you provide it a name and it gives you the number which allows you to connect to your correspondent. However this mechanism is generally transparent to the end user: he never needs to remember or type IP addresses thanks to the DNS servers.

In this chapter we briefly show you how to configure the global server options, and how to declare new zones (basically domain names) to be handled by your DNS server. This means other machines on the local network, and possibly on the Internet, will be able to access the machines and services associated to your own domain names.

The Webmin BIND DNS Server module creates and edits domains, DNS records and BIND options for the 8.x and 9.x releases. BIND (Berkeley Internet Name Domain) is an implementation of the Domain Name System (DNS) protocol and provides an open, redistributable reference implementation of the major components of the Domain Name System.

BIND is very useful for simple configurations, but there are a few differences between the 8.x and 9.x releases. However, you must be careful with this Webmin module because not all BIND 9.x options are supported yet. Therefore if you try to use the advanced options, you have to look at the log files more carefully than with the other Webmin modules to make sure BIND works properly.



In this chapter we cover the configuration of a DNS server for local queries. We don't explicitly cover the configuration of a public name server.

3.1. Installation and Initialization

Make sure the `bind` package is installed.



At this point Mandriva Linux allows you to automatically configure your name server in a very specific case. If you don't wish to serve your own specific domains with your name server, but just want it to act as a forwarder for local network clients, you can simply install the `caching-nameserver` package. It will build a basic BIND configuration, allowing your server to answer local DNS requests for Internet addresses. The server must have access to the Internet.

Once this package is installed, start the server by issuing the service named `restart` command as `root`. You can now configure your local machines (*Configuring the Client*, page 44) to make DNS requests on your server.

In order to use the Webmin module, you have to select the **Servers** category and then the BIND DNS Server button (with the number 8 in the icon).

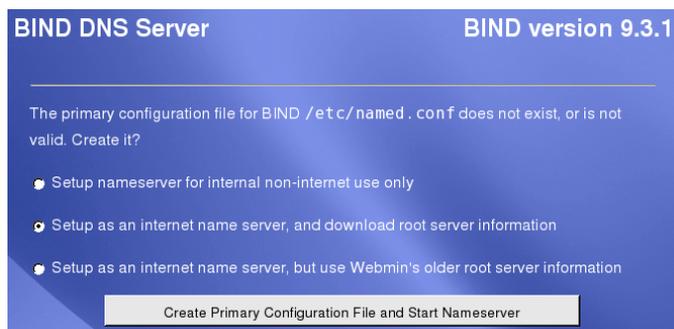


Figure 3-1. Creating a Primary Bind Configuration File

The first time you open this module, and provided you have not configured BIND beforehand, you have to choose which kind of usage you plan for your name server.

Setup nameserver for internal non-Internet use only

Select this option if you plan to use the name server solely to serve requests from the local network for machines also located in the local network. This is only useful if your network isn't connected to the Internet.

Setup as an Internet name server, and download root server information

Select this option if your server is meant to answer requests from or to the Internet, and if the server currently has access to the Internet.

Setup as an Internet name server, but use Webmin's older root server information

Select this option if your server is meant to answer requests from or to the Internet, and if the server currently **does not** have access to the Internet.

Then click on Create Primary Configuration File and Start Nameserver to follow on with the configuration.

The main screen is divided into two parts: the Global Server Options and the Existing DNS zone, which allows you to access each zone already defined and represented by an icon, as well as to create new zones.



Whenever you change parameters, whether in global server options or in zones, don't forget to click the Apply Changes button for the server to reload the configuration.

3.2. Step-by-Step Configuration Example

If you chose to set up your nameserver as an Internet nameserver, there is already one Existing DNS Zone: the **Root Zone**. It's used by the DNS server to contact the root servers on the Internet so it can resolve domain names not handled by your own DNS server. Unless your DNS server is used on an internal network (no access to the Internet) or if you're forwarding all queries to another server, you should **not** delete this root zone.

3.2.1. Basic Server Configuration and Security

Several parameters can be adjusted to optimize and secure your DNS server.

3.2.1.1. Defining DNS Forwarders

The Forwarding and Transfers screen allows you to list nearby nameservers to which requests will be forwarded if the local server cannot reply directly.

Global forwarding and zone transfer options		
Servers to forward queries to	IP address	Port (optional)
	192.168.0.1	
	212.27.39.134	
	212.27.39.135	
Lookup directly if forwarders cannot?	<input type="radio"/> Yes <input type="radio"/> No <input checked="" type="radio"/> Default	
Maximum zone transfer time	<input checked="" type="radio"/> Default <input type="radio"/> <input type="text"/> minutes	
Zone transfer format	<input type="radio"/> One at a time <input type="radio"/> Many <input checked="" type="radio"/> Default	
Maximum concurrent zone transfers	<input checked="" type="radio"/> Default <input type="radio"/> <input type="text"/>	

Figure 3-2. DNS Forwarders

In the Servers to forward queries to field, you should list the IPs of other possible nameservers in your local network, and at least two Internet nameservers: your ISP's. This is likely to lower the server load and accelerate response time.

3.2.1.2. Securing the Server

The Addresses and Topology screen permits you to define which addresses the server will listen on. It's safer to listen only on internal interfaces if the server isn't meant to answer outside requests.

Global address and topology options			
Ports and addresses to listen on	<input checked="" type="radio"/> Default <input type="radio"/> Listed below..		
	Port	Addresses	
	<input checked="" type="radio"/> Default <input type="radio"/> <input type="text"/>	192.168.0.10 127.0.0.1	
Source IP address for queries	<input type="text"/>	Source port for queries	<input checked="" type="radio"/> Default <input type="radio"/> <input type="text"/>
Nameserver choice topology	<input checked="" type="radio"/> Default <input type="radio"/> Listed ..		
	<input type="text"/>		

Figure 3-3. DNS Addresses to Listen on

In the first Addresses field, enter the list (separated by spaces) of the server addresses on which the DNS server will listen to queries on. Requests made through other addresses will be ignored.



Of course this measure doesn't prevent you from setting up a firewall which is nowadays mandatory for every security-conscious network installation.

3.2.2. Setting up Basic DNS Zones

In order to use each network service properly, you need to create zones for each of the domains the server will deal with. Now let's concentrate on the Existing DNS Zones part. If you chose for your DNS server to act as an Internet nameserver at service initialization (figure 3-1), the Root Zone has already been created. It allows your nameserver to answer queries concerning public Internet names by forwarding them to the appropriate servers on the Internet.

First you need to create a minimal Master 127.0.0 zone to describe the loopback network. This is useful for security reasons and to use the server as a caching server. This is done in two simple steps: reverse master zone creation and host configuration. Click on the Create master zone link in the main screen.

Figure 3-4. Creating a Forward Master Zone

Select Reverse Zone type, and fill the Domain name / Network field with the local host network address: 127.0.0 (no final dot). Use localhost as Master server and root@localhost (or whatever you like) as administrator Email Address. Clicking on Create to save your configuration.

You are then directed to the corresponding Edit Master Zone screen. Click on the Reverse Address Icon.

Figure 3-5. Creating a localhost Reverse Record

Fill the Address and Hostname fields to match your localhost configuration. Then click on Create. Your local zone is now created and you can forget about it. Click on the Return to zone list link to go back to the main screen.

3.2.3. Setting up Specific DNS Zones

We now have the two basic zones: root and 127.0.0. Our next task is to create a new master zone which will describe all of our local network machines. Select Create master zone and complete the page with your own values, as shown in figure 3-6.

Figure 3-6. Creating a New Master Zone

A new page with many icons is displayed: most of them can be ignored if you don't need advanced configuration. You can add all network machine names through this page, but first, you should create the reverse part of your master zone. In fact, a DNS zone is composed of two parts: one for name-to-address conversion (i.e.: forward), and another one for address-to-name conversion (i.e.: reverse).

Then, select Return to the zone list and choose Create a master zone once more but this time, you must change the selection from Forward to Reverse. Instead of entering your domain name, enter the network class: for a 192.168.1.0/24 network, you should write 192.168.1.

New master zone options

Zone type: Forward (Names to Addresses) Reverse (Addresses to Names)

Domain name / Network: 192.168.1

Records file: Automatic

Master server: mycomputer.mydomain.test Add NS record for master server?

Email address: myemail@mydomain.test

Use zone template? Yes No IP address for template records: []

Refresh time: 10800 seconds Transfer retry time: 3600 seconds

Expiry time: 604800 seconds Default time-to-live: 38400 seconds

Figure 3-7. Creating a Reverse Master Zone

Remember to click on the Create button. Your zone is now ready to host new records for machines or services of our local network.

3.2.4. Recording your Network's Computers

This is the only step you have to repeat each time you add a new machine in your network. All other parameters are configured only once, as long as your network doesn't change and you don't add other DNS servers.

Return to zone list and now your Existing DNS Zones should show four zones.



Figure 3-8. All DNS zones

Click on the mydomain.test zone, and then on the Address icon. This is where you actually define all the existing machine names in your network and record their IP addresses.

In mydomain.test

Add Address Record

Name: Time-To-Live: Default seconds

Address:

Update reverse? Yes Yes (and replace existing) No

Name	TTL	Address	Name	TTL	Address
machine1.mydomain.test.	Default	192.168.0.11	machine3.mydomain.test.	Default	192.168.0.13
machine2.mydomain.test.	Default	192.168.0.12			

Figure 3-9. Adding Machine Names

You can now add as many machine Names as your IP class allows (254 machine names in our example). Click on Create to add the new record, and you are then prompted to fill a new one. Note that the Update reverse? option is selected by default. With this option, the Reverse part of your DNS is updated automatically.

3.2.5. Starting the Service

We created a very simple, yet complete DNS service. To start it and load the new configuration, Return to zone list and click on the Apply Changes button.

Webmin checks the parameters you enter so it shouldn't be possible to provide BIND with a corrupt configuration. However if the button is replaced by a new one named Start Name Server, then the server didn't start because of a configuration error. In this case, you should read *Advanced Configuration and Troubleshooting*, page 44.

3.2.6. Configuring the Client

To use your local network to resolve Internet addresses, you have to configure the client to access the DNS. This can be done either through the Mandriva Linux Control Center network configuration tool, or through Webmin: go to the Networking tab and click on the Network Configuration icon. Then, select the DNS Client and type your DNS's IP if it's a remote client, or 127.0.0.1 if you're on the server.

DNS Client Options

Hostname: Resolution order:

Update hostname in host addresses if changed?

DNS servers:

Search domains: None Listed...

Figure 3-10. Configuring the Client

3.3. Advanced Configuration and Troubleshooting

3.3.1. How to Debug

If the service didn't start, you should look at the `/var/log/messages` file to read the debug output of BIND.

If you don't find the error, you can use the `named-checkconf` and `named-checkzone` programs to check your configuration.

With the `bind-utils` package, you can use many utilities and especially the `dig` command to perform advanced queries on DNS servers. For example to query your local server about `machine2.mydomain.test`, you could run:

```
$ dig machine2.mydomain.test @127.0.0.1

; <<>> DiG 9.2.3 <<>> machine2.mydomain.test @127.0.0.1
;; global options: printcmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 3287
;; flags: qr aa rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 1, ADDITIONAL: 0

;; QUESTION SECTION:
;machine2.mydomain.test.          IN      A

;; ANSWER SECTION:
machine2.mydomain.test. 38400   IN      A      192.168.1.12

;; AUTHORITY SECTION:
mydomain.test.          38400   IN      NS     mycomputer.mydomain.test.

;; Query time: 14 msec
;; SERVER: 127.0.0.1#53(127.0.0.1)
;; WHEN: Fri Jan 23 14:11:10 2004
;; MSG SIZE rcvd: 81
```

3.3.2. Declaring your Mail Server

If your mail server is meant to handle addresses for your own domains, handled by your own DNS server, then this mail server should be declared in your domain configuration. This way, other mail servers know which machine is responsible for delivering messages addressed to users of your domain.

In the main screen of your domain's zone, click on the Mail Server icon.

Add Mail Server Record			
Name	<input type="text" value="mydomain.test."/>	Time-To-Live	<input checked="" type="radio" value="Default"/> <input type="text" value=""/> <input type="text" value="seconds"/>
Mail Server	<input type="text" value="smtp.mydomain.test"/>	Priority	<input type="text" value="1"/>
			<input type="button" value="Create"/>

Figure 3-11. Declaring a Mail Server

Fill the Name field with the domain served by the mail server (the same as the zone domain) and write the server name in the Mail Server field. Make sure this name is also defined as a local domain machine if it's part of the local network. Repeat this operation for each of the mail servers.



The domain Name in this form must be terminated by a dot as in our example.



The Priority field (useful when you have more than one mail server for the same domain) defines the order in which servers should be contacted should the ones of highest priority (lower priority number) be unreachable.

3.3.3. More Documentation

An extensive chapter is dedicated to BIND in Joe Cooper's *The Book of Webmin* (<http://www.swelltech.com/support/webminguide/ch08.html>). Explanations for almost every available options in Webmin's BIND are included.

If you want to learn more about BIND, we strongly recommend you read the BIND 9 Administrator Reference Manual (<http://www.bind9.net/Bv9ARM.html>) which is available locally on your machine: click on Search docs in the upper-right corner of Webmin's BIND DNS Server page. It shows many local and Internet related links. Note that the *Reference Manual* is available in HTML if you click on [bind-9.3.1/html/Bv9ARM.html](http://www.bind9.net/html/Bv9ARM.html). This manual is also available in PDF format (<http://www.nominum.com/content/documents/bind9arm.pdf>). Finally don't hesitate to browse the official BIND web site (<http://www.bind9.net/>).

Chapter 4. Internet and Intranet Web Server

Apache allows you or your organization to host web sites and serve web pages to client browsers such as Firefox. Apache can serve static or dynamic sites using a number of technologies like PHP, SSL, etc.

4.1. Installation

The first step is to check that the Apache web server is installed on your computer. If it isn't please use Rpmrake or type `urpmi apache-mpm-prefork` in a terminal to install it.

Mandriva Linux's highly modular distribution of Apache allows you to have support only for the technologies you need, no more no less. The module's package names are `apache-mod_XXX`, where `XXX` is the name of the module in question. Make sure you also install the corresponding modules you need.

The server configuration is done through the Apache Webserver button which is in the Servers category.

4.2. Configuration Example



Figure 4-1. Apache Module's Start-Up Screen

You can host several different sites on a single Apache server: this feature is known as "Virtual Servers". Your "main site" is the Default Server. Global Configuration options apply to all virtual servers. Each virtual servers' options are found in the Virtual Servers section. We will defer global options to *Advanced Configuration*, page 49.

Apache's configuration files are stored in the `/etc/httpd/conf/` directory. The main Apache options you need to set are located in the `/etc/httpd/conf/httpd.conf` file and can be accessed by clicking on Default Server.

4.2.1. Default Server General Options

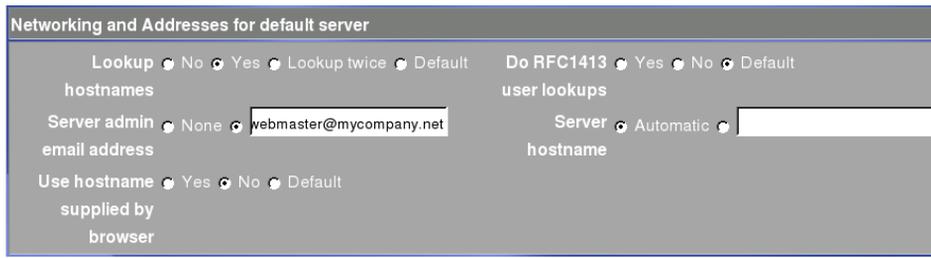


Figure 4-2. Networking and Addresses Options Section

In the Networking and Addresses section you can specify the webmaster’s e-mail address in the Server admin email address field. To avoid “false” requests on your web server, set the Lookup hostnames option to Yes and the Use hostname supplied by browser to No (see figure 4-2).



Activating the hostname lookups in order to make your web site(s) more secure has its penalty: performance. A DNS lookup is performed each time a request is made to your web server.

Accesses, errors and other operations are logged in files under the /var/log/httpd/ directory. Logging options are available under the Log Files section: where to log to (the centralized system log, a specific log file or program), the format for the log entries, etc. Feel free to investigate the different options.

Now click on the Document Options icon.

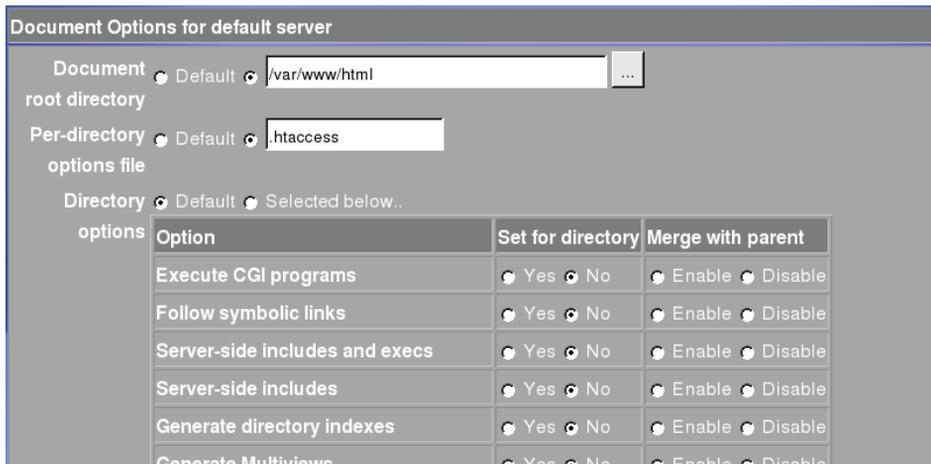


Figure 4-3. Document Options Section

The Document root is the path to the directory containing the web pages and their associated files. Directory options contains some options like the ability to execute a CGI program (Execute CGI programs) or to automatically generate an index of files in a directory if no explicit index exists (Generate directory indexes).

If your web site’s structure is large, a few directories deep and you want to simplify navigation, you can create aliases in the Aliases and Redirects section.

Aliases and Redirects for default server			
Document directory aliases	From	To	
	/icons/	/var/www/icons/	
	/error/	/var/www/error/	
	/morestuff/	/var/www/html/foo/bar/again/and/more	
Regexp document directory aliases	From	To	
URL redirects	From	Status	To
Regexp URL redirects	From	Status	To

Figure 4-4. Aliases and Redirections Section

The example in figure 4-4 shows how to let point your browser to `http://www.example.com/morestuff` instead of `http://www.example.com/foo/bar/again/and/more`. The second part of the screen is dedicated to redirections, which let you redirect a part of your web address to a specific web page or directory.

4.2.2. Common Technologies Specific Options

4.2.2.1. CGI

If you plan to use Common Gateway Interface programs, the CGI Programs section lets you specify which directory contains your CGI scripts, and configure some variables passed to the executables. The default settings allow you to directly use your CGI scripts in your web site.

4.2.2.2. SSL

SSL provides encrypted communications facilities which makes access to your site(s) secure. You can enable SSL by simply installing the `apache-mod_ssl`.

Your site(s) will then be ready to use encrypted communications channels with SSL, using the `https://` prefix instead of `http://`.



At the time of writing, the standard version of Apache only support one SSL-enabled site per IP address. Should you need to host several secured sites on the same IP address, you need to install the `apache-ssl` package instead. However note that the architecture of this server (notably configuration files wise) differs from the default server.

4.2.2.3. PHP

Simply install the `apache-mod_php` package for your PHP pages to be interpreted.

4.2.3. Per-Directory Options

If you click on a directory name under this part of the server configuration you can specify general options for that directory. For example, you can configure specific Mime Types for your site's download directory, Access Control for a specific directory, etc.

4.3. Advanced Configuration

The options discussed in this section are accessible from the Webmin Apache module index.

4.3.1. Processes and Limits

Figure 4-5. Limiting the Number of Apache Processes

You can fine-tune Apache's usage of your systems' resources by imposing limits on the number of initial instances of Apache (Initial server processes) and the maximum number of processes those are going to start, if needed (Maximum spare server processes); the number of clients per process (Maximum requests per server process) and the header's size (Maximum headers in request, Maximum request header size and Maximum request line size).

4.3.2. Networking and Addresses

Figure 4-6. Changing the Ports Apache Listens On

Figure 4-6 shows how you can specify the ports Apache listens on for both regular (80 by default, 8080 in the example) and encrypted sessions (443 by default, 4433 in the example).

4.3.3. Controlling Access With Basic Authentication

Authentication usually involves a username and a password, but can include any other method of demonstrating identity. You can control access to certain parts (directories) of your site(s) using a password file which acts as an authoritative listing of usernames and passwords. Here's the procedure to implement such a mechanism:

1. Create a passwords list file and fill it.
2. Protect a specific directory by creating special configuration directives for that directory, which refers to the passwords file.

Let's imagine you need to control access to the `/var/www/html/restricted/` directory.

To create the passwords file, type `htpasswd -c -m path_to_the_password_file username` in a console, as root. The `-c` is used only the first time to create the file.

```
# htpasswd -c -m /etc/httpd/.htpass queen
New password: verySecret
Re-type new password: verySecret
Adding password for user queen
```

The above example creates the `/etc/httpd/.htpass` file containing the password (`verySecret`) for user `queen`. Of course, `verySecret` will be encrypted.



To minimize security risks, it's a good idea to store the generated password file **outside** of the documents directory and make sure its file-access rights are as tight as possible.

Once you have the password file populated, you must instruct Apache to use it. Enter the server which is meant to serve your protected directory (Default Server for example) in Apache's main screen (figure 4-1). A little form enables to create per-directory options at the bottom of the server screen.

Figure 4-7. Per-Directory Options Directive

Once this is done, a new icon Directory `/var/www/html/restricted` appears in the Per-Directory Options section. Click on it and then on the Access Control button of the page. You can then fill the form using the example shown in figure 4-8 as a guide. Click on the Save button to record your settings.

Figure 4-8. Per-Directory Options

4.3.4. Handling Multiple Domains With One Web Server



When setting up virtual hosts, the first one catches all requests which don't match other virtual hosts.

Using the Virtual servers section, you can directly set up a multi-domain web server using the form at the bottom.

Figure 4-9. Creating a New Virtual Server Based on an Existing One

For example, your company owns the `foo.com` and `bar.net` domains. You just have to specify the document root (where your site's files are stored), and the name of the virtual server. If you manage multiple sites, you can copy configuration directives from other virtual servers (see figure 4-9). This can save you lots of time.



We are doing here "Name based virtual hosts", meaning that we are hosting different servers on the same IP address. For this to work, you need to add a special directive in Apache's main configuration file with the following command:

```
# echo "NameVirtualHost *:80" >> /etc/httpd/conf/httpd.conf
```

Don't forget to restart the server by clicking on the Apply Changes button.



Of course your nameserver ("*BIND DNS Server*", page 39) must be configured so that when clients request the virtual host name (`www.foo.com`), they are directed to the machine which hosts your web server.

Each virtual server possesses similar options to the ones described in the preceding sections, but they all share a common Apache parent process.

4.4. More Documentation

Joe Cooper dedicates a long chapter to Apache in *The Book of Webmin* (<http://www.swelltech.com/support/webminguide/ch07.html>). You will find explanations for almost every option available in the Webmin's Apache module, though it might be a little outdated.

Browsing the Apache Documentation Project (<http://httpd.apache.org/docs-project/>) is also a good idea. If you install the `apache-doc` RPM, you can alternatively access Apache documentation on your own installation under the `/usr/share/doc/apache-doc-*/` directory, or by browsing to `http://localhost/manual/` (`http://localhost/manual/`).

Chapter 5. Postfix Mail Server

With Postfix, you can set up and configure a mail server to send and receive mail. This server can communicate directly with other mail servers on the Internet through the SMTP protocol. With the right configuration, Postfix can handle all mail sent to your company's domain.

5.1. SMTP Server Functions

An SMTP (Simple Mail Transfer Protocol) server can be compared to a postal sorting office. The office receives letters from the neighborhood, and sorts them: if a letter is addressed to someone else in the neighborhood, it's stored in his or her mailbox. Otherwise the letter is sent to the postal office corresponding to the recipient's address. The same happens for letters relayed by another post office.

The operations of a standard Postfix mail server are very similar: it receives e-mail from the local network users and from other mail servers which have identified your mail server as responsible for handling the e-mail addressed to a specific domain name. The server reads the recipient address, and then:

- If the domain name corresponds to the one locally served, the mail is stored on the corresponding local mailbox. Then users pick up their messages through a mail client. The actual delivery of the messages to the user is done through another protocol (*"Mail Delivery Services: Pop and IMAP"*, page 59).
- Otherwise the server looks on the Internet for the server responsible for handling that domain name's addresses, and forwards the message to it.

5.2. Installation

Make sure the `postfix` package is installed on your system.

The server configuration is done through the Postfix Configuration button which is located in the Servers category.



If the mail server is to receive messages for a specific domain from other servers, it must be marked as such in the DNS configuration, either on your local DNS server (*"BIND DNS Server"*, page 39) or at your registrar's.

5.3. Step-by-Step Configuration Example



Each Postfix option in the Webmin module is documented. Just click on the option's name and a pop-up window appears, explaining the relevant option.

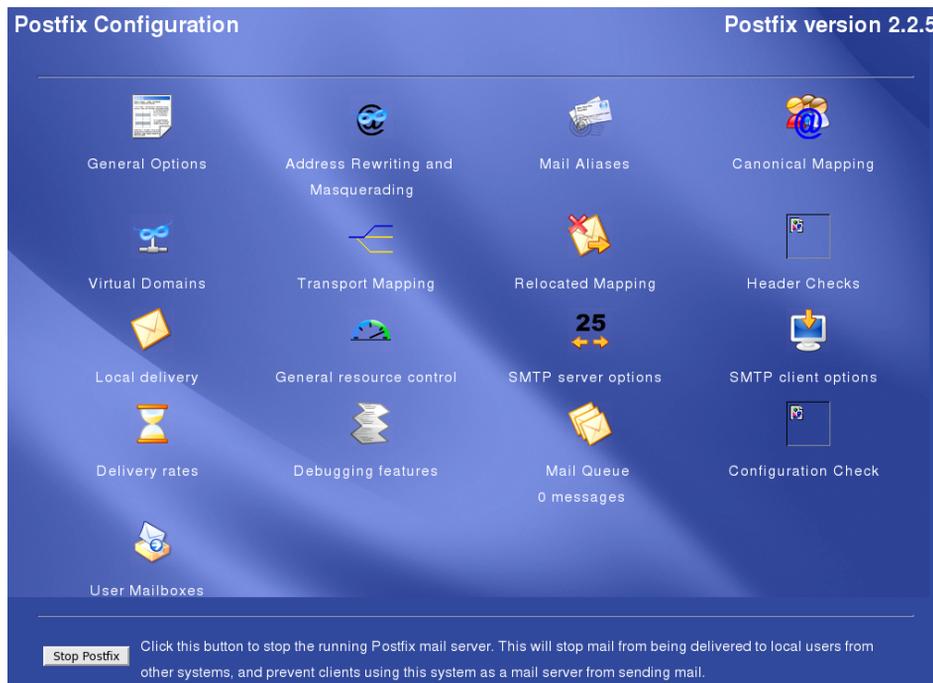


Figure 5-1. Postfix Start-Up Screen

Postfix’s configuration files are stored in the `/etc/postfix` directory. The main options to configure are located in the `/etc/postfix/main.cf` file and can be configured by clicking on the General Options icon.

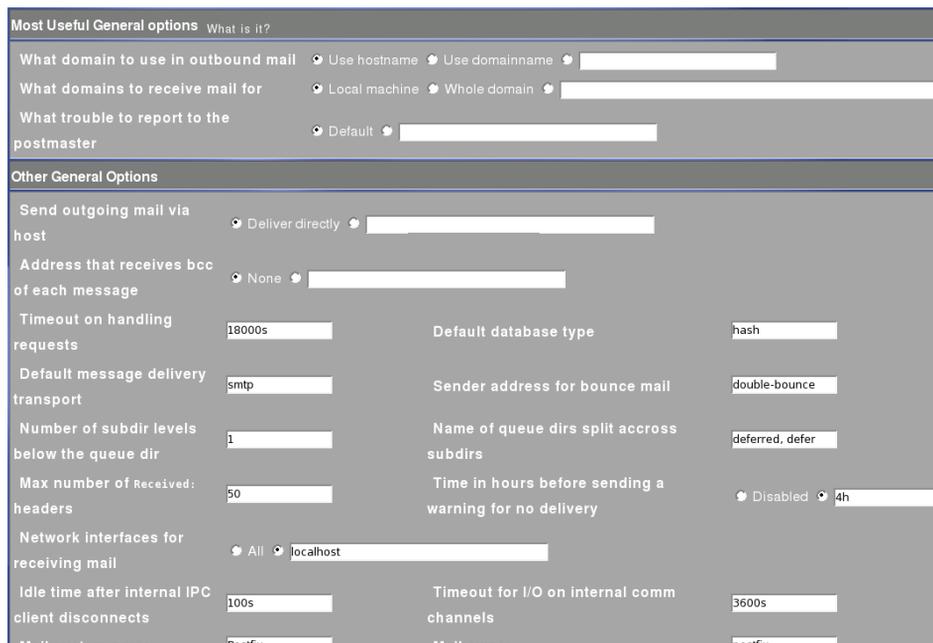


Figure 5-2. Postfix’s Main Configuration Screen

In this section we must configure the following main options:

- What domain to use in outbound mail. This option concerns outgoing mail. You should specify the mail domain. Leave it as Use domainname if the computer’s domain name has the same value as your mail domain name.
- What domains to receive mail for. This parameter deals with incoming mail. You should specify the mail domains for which this server is responsible. Set it to Whole domain if the computer’s domain name has the same value as your mail domain name. Otherwise select the third radio button and enter the comma-separated list of handled domains, as well as all the names the server could have, including `$myhostname` and `localhost.$mydomain`.

- Send outgoing mail via host. This option is useful if you access the Internet through an ISP which provides you with a mail server. If this is the case, you can choose to use it as a relay to dispatch your own messages to their actual recipients. That will reduce the load of your server but you must trust your ISP's integrity. Then provide the name of your ISP's mail server: `smtp.provider.net` for example.
- Internet hostname of this mail system. This parameter specifies the Internet hostname of this mail system. The default is to use the Fully-Qualified Domain Name (FQDN). For example, `gateway.example.com`. Leave the "Default (provided by system)" value if your hostname looks like `computer_name.mx_domainname`.
- Local internet domain name. This option specifies the local Internet domain name. The default is to use `$myhostname` minus the first component. For our example it would be `example.com`. Leave the "Default (provided by system)" value if your hostname looks like `computer_name.mx_domainname`.
- Local networks. This parameter is used to identify which machines are trusted for relaying through your mail server. Messages coming from those machines and directed to other servers on the Net will be accepted and forwarded without restrictions. Typical values, for example, are `192.168.1.0/24`, `127.0.0.0/8`. This means that you allow relaying from the localhost and systems for which the addresses are in the `192.168.1.1-254` range. Please be sure to specify the correct networks to avoid becoming a spam victim.

Back to the module index, the Mail Aliases section helps you configure the mail redirection to valid existing mailboxes. As you can see in the table, many default aliases exist which all converge, possibly after various hops, to the `postfix` address. We recommend that you add an alias for this address to point to your personal account or address, in order for messages sent to one of the defined aliases (including root messages for system alerts) are actually forwarded to you instead of being stored locally in the `postfix` user's mailbox.

Figure 5-3. Defining a New Mail Alias

To conclude with general server configuration, it may be interesting to check the General resource control page. Two options are interesting:

Max size of a message

Configures the maximum size (in bytes) of e-mail accepted by Postfix. This number is the size of the whole message, including the headers and any attachments. Consider that mail software encodes non-text attachments, so the total size of a message is actually greater than the on-disk size of attachments.

Max size of bounced message

If Postfix cannot deliver mail to its final destination, it sends a non-delivery notification message to the original sender. This notification contains the cause of the error and a configurable amount of text from the original message. This parameter indicates how much text (in bytes) should be included.

5.4. Advanced Configuration

Here are some of the more interesting General options ones. If you want to keep trace of all the e-mails which pass through the server, put an address in the Address that receives bcc of each message field. You can also specify the Time in hours before sending a warning for non-delivery. The other options are system-specific, don't change these parameters unless you fully understand their meaning.

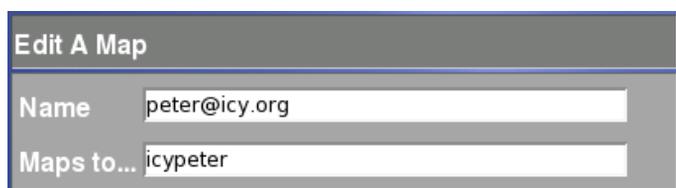
5.4.1. Address Mapping

In the Canonical mapping section, you can specify mapping table files which are used to rewrite e-mail headers managed by Postfix. For example, in the Address mapping lookup tables field, you could associate the name of employees with their e-mail address: `John.Doe@example.com` and `jdoh@example.com`.

5.4.2. Virtual Domains User Mapping

If you're serving more than one mail domain with your mail server, you may need to configure the Virtual domains section. Indeed if you handle mails for both the `pingus.org` and `icy.org` domains, e-mails sent to `peter@pingus.org` and `peter@icy.org` are delivered to the same `peter` mailbox. If this is actually the same person, it's not a problem. But if they're two different persons, the messages must be delivered to two different mail accounts.

First you need to specify the Domain mapping lookup tables. You must type `hash:/etc/postfix/virtual` and click on Save and Apply. Next create a new mapping by clicking on New mapping. For example, e-mails addressed to `peter@icy.org` to be redirected to local user `icypeter`, instead of `peter`.



Edit A Map	
Name	<input type="text" value="peter@icy.org"/>
Maps to...	<input type="text" value="icypeter"/>

Figure 5-4. Postfix: Configuring Virtual Domains

This way you can redirect mail to any local user, users of any other domain or server, and even whole domains to a single user.

5.4.3. More Miscellaneous Options

The Local delivery section contains options to help you configure e-mail handling after Postfix receives them.

In the SMTP server options section, you can prevent receiving spam mail by configuring the DNS domains for blacklist lookups field. Some Internet servers run public DNSs with blacklisted hosts. These mail hosts relay spam mail. Configuring this option allows Postfix to look in these databases before accepting e-mail. All Postfix responses at the bottom of the page should be kept to the default values.

If LDAP is installed on your system, you could access and configure options in the LDAP lookups section.

You can use the Address rewriting and masquerading section to hide all hosts inside a domain behind their mail gateway, and to make it appear as if the mail comes from the gateway itself instead of from individual machines. Note that this option is activated by default.

5.4.4. Mailbox Access

Finally you may have noted the User mailboxes section which allows you to browse local mailboxes to possibly perform maintenance tasks on them by manipulating messages.

It may be interesting to have access to a user's messages if he has problems accessing his mailbox, or to remove huge attachments which overload the server's disk.



When using this feature, bear in mind your system administrator ethics and your user's rights. Don't abuse your rights by reading other users' messages.

5.5. Extra Documentation

Joe Cooper's The Book of Webmin (<http://www.swelltech.com/support/webminguide/ch10.html>) contains a long chapter dedicated to Postfix. You will find explanations for almost all options available in Webmin's Postfix module.

It's also a good idea to browse the Postfix Documentation pages (<http://www.postfix.org/docs.html>). Alternatively, you may access the Postfix documentation on your own installation in `/usr/share/doc/postfix-*/`.

Chapter 6. Mail Delivery Services: Pop and IMAP

By using POP (Post Office Protocol) and IMAP (Internet Message Access Protocol), users can access their electronic mailboxes and get their e-mails to read them on their machines.

6.1. Foreword and Installation

If you did a standard Mandriva Linux installation, mail access servers (POP3 or IMAP) are launched on demand: when a connection is made to the POP or IMAP ports, the appropriate program to answer that request is launched.

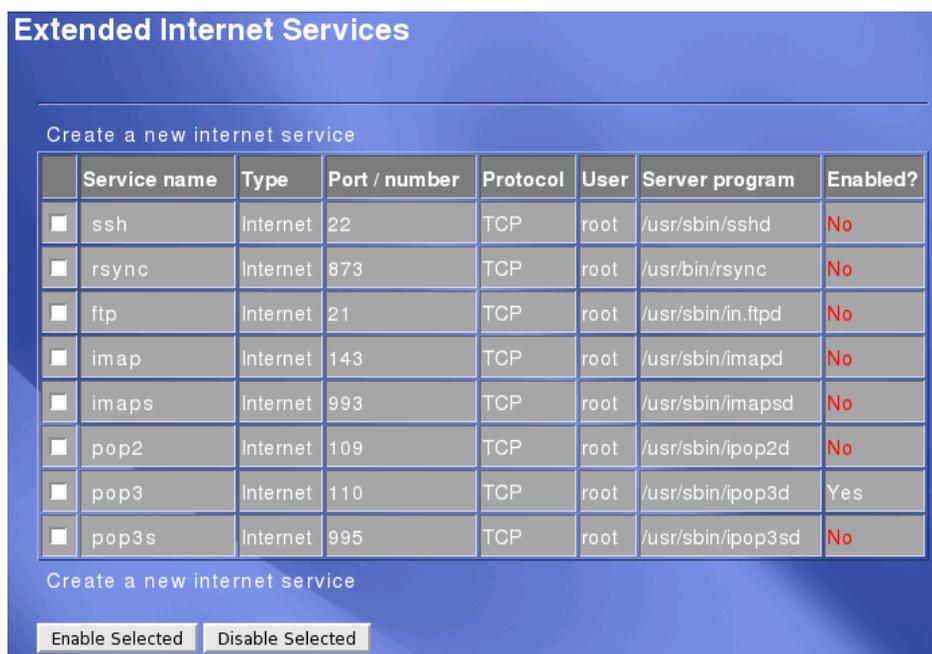
A POP3 user fetches messages on his computer and reads them with a mail reader such as KMail or Evolution, while the IMAP protocol allows users to leave their messages on the server and manage them remotely, making it ideal for mobile users. However since messages can consume a lot of disk space, system administrators should regularly check their servers or set up quota policies.

Make sure the `imap` package is installed.

6.2. Step-by-Step Configuration Example

It's better to configure only the services which you need to use and to close others: know if your users need POP, IMAP, or both.

Click on the Extended Internet Services icon in the Networking category for the program to list all the accessible services on your computer which are managed by `xinetd`. These services can be up (activated) or down (stopped).



	Service name	Type	Port / number	Protocol	User	Server program	Enabled?
<input type="checkbox"/>	ssh	Internet	22	TCP	root	/usr/sbin/sshd	No
<input type="checkbox"/>	rsync	Internet	873	TCP	root	/usr/bin/rsync	No
<input type="checkbox"/>	ftp	Internet	21	TCP	root	/usr/sbin/in.ftpd	No
<input type="checkbox"/>	imap	Internet	143	TCP	root	/usr/sbin/imapd	No
<input type="checkbox"/>	imaps	Internet	993	TCP	root	/usr/sbin/imapd	No
<input type="checkbox"/>	pop2	Internet	109	TCP	root	/usr/sbin/ipop2d	No
<input type="checkbox"/>	pop3	Internet	110	TCP	root	/usr/sbin/ipop3d	Yes
<input type="checkbox"/>	pop3s	Internet	995	TCP	root	/usr/sbin/ipop3sd	No

Figure 6-1. Start-up Screen of the `xinetd` Module

Both POP and IMAP have a secure protocol associated with them: POP3S and IMAPS which encrypt data flow. If your users access their e-mail through the Internet it's safer to use these secured services and to deactivate others. Make sure the clients they use actually support those secure protocols.

The `imap` package automatically activates the POP service with standard options. Click on a service name to activate or deactivate it and to configure it.

Service network options	
Service name	pop3
Service enabled?	<input checked="" type="radio"/> Yes <input type="radio"/> No
Bind to address	<input checked="" type="radio"/> All <input type="radio"/> []
Port number	<input checked="" type="radio"/> Standard <input type="radio"/> []
Socket type	Stream
Protocol	Default
Service program options	
Service handled by	<input checked="" type="radio"/> Internal to Xinetd
	<input type="radio"/> Server program /usr/sbin/pop3d
	<input type="radio"/> Redirect to host [] port []
Run as user	root
Run as group	<input checked="" type="radio"/> From user []
Wait until complete?	<input type="radio"/> Yes <input checked="" type="radio"/> No
Max concurrent servers	<input checked="" type="radio"/> Unlimited []
Nice level for server	<input checked="" type="radio"/> Default []
Maximum connections per second	<input checked="" type="radio"/> Unlimited []
Delay if maximum is reached	[] seconds
Service access control	
Allow access from	<input checked="" type="radio"/> All hosts <input type="radio"/> Only listed hosts..
Deny access from	<input checked="" type="radio"/> No hosts <input type="radio"/> Only listed hosts..
Allow access at times	<input checked="" type="radio"/> Any time <input type="radio"/> []

Figure 6-2. POP3 Configuration Module

In Service enabled? select the Yes option for the service to be accessible. You can then restrict access to it in the Service access control section. Select the Only listed hosts option and enter the IP addresses of computers allowed to retrieve mail in the Allow access from box.

Save your changes and click on Apply changes so the xinetd daemon applies the new configuration.



This configuration screen is the same for all services managed with xinetd.

6.3. Advanced Configuration

There are many other options which are not mandatory for a standard configuration. In the Service network options section, the Bind to address and Port number options allow you to force the daemon to listen on a specific address-and-port pair. If you have many network interfaces and you want mail traffic to pass only through a specific one, you can specify it in this section by entering that interface's IP address.

In the Service program options section, you can choose to redirect all your requests to another computer. In Service handled by, select the Redirect to host option and enter the IP address and port of the machine to which you want requests to be redirected to. The Run as user and Run as group options both allow the service to be run as a specific user.

xinetd allows you to set up specific limits for each service. The Max concurrent servers option specifies the maximum number of daemon instances which can be running at the same time. Maximum connection per second specifies the number of connection requests the server can handle. If the maximum is reached, then the Delay if maximum is reached option specifies the time interval until that instance of the service daemon will be reachable again. In the POP3 example, you can specify that only three POP3 servers can be launched and respond to five connection requests per second. This can be useful to minimize server overloading.

The last useful option is Nice level for server, which indicates the program's system priority. If various services are available on the same server and some are more critical than others, this option allows you to tell the system to assign more resources to the more critical processes. The Nice level is at 0 by default and you may choose values between -20 (the highest priority) and 19 (the lowest). If you consider the POP3 service to be of less importance with respect to the other services hosted on your server, then you could assign it a nice level of 10 for example.

Chapter 7. Resource Sharing

7.1. Samba: Integrating Linux in a Windows Network

The Samba server allows you to easily integrate a Mandriva Linux computer in a mixed GNU/Linux & Windows® network. Through Samba, your computer can appear in other people’s network neighborhood and act as a Windows® server sharing files and printers, remote user accounts etc.

7.1.1. Installing Samba

Make sure the `samba-server` package is installed.

The server configuration is done through the Samba Windows File Sharing configuration button located under the Servers category.

7.1.2. Step-by-Step Configuration Example

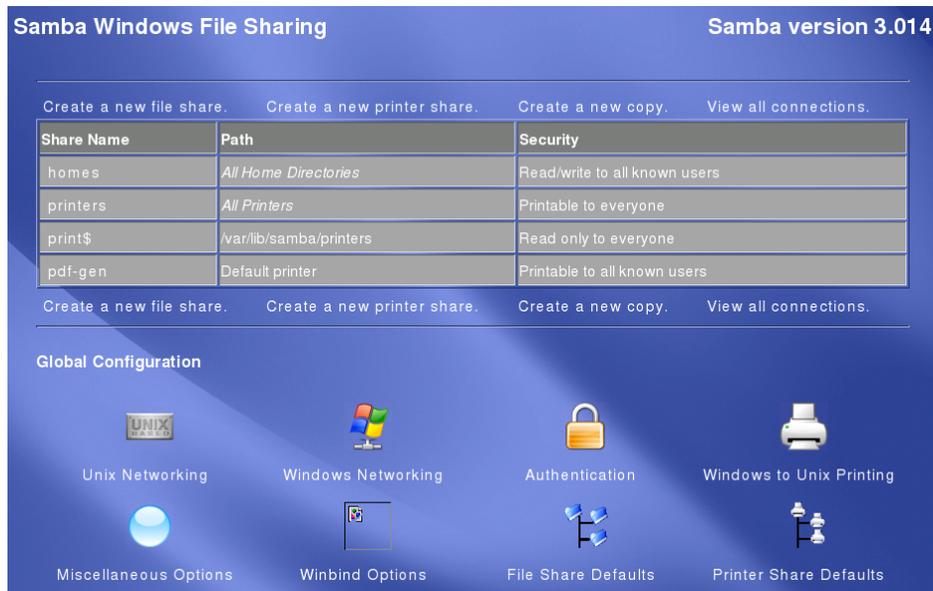


Figure 7-1. The Samba Module’s Main Window

Samba’s configuration files are stored in the `/etc/samba` directory. The main Samba options you need to set are located in the `/etc/samba/smb.conf` file and are accessed by clicking on Windows Networking.



Samba automatically reloads its configuration every minute, so there’s no need to constantly restart the Samba server for your changes to the settings to be effective.

7.1.2.1. General Settings

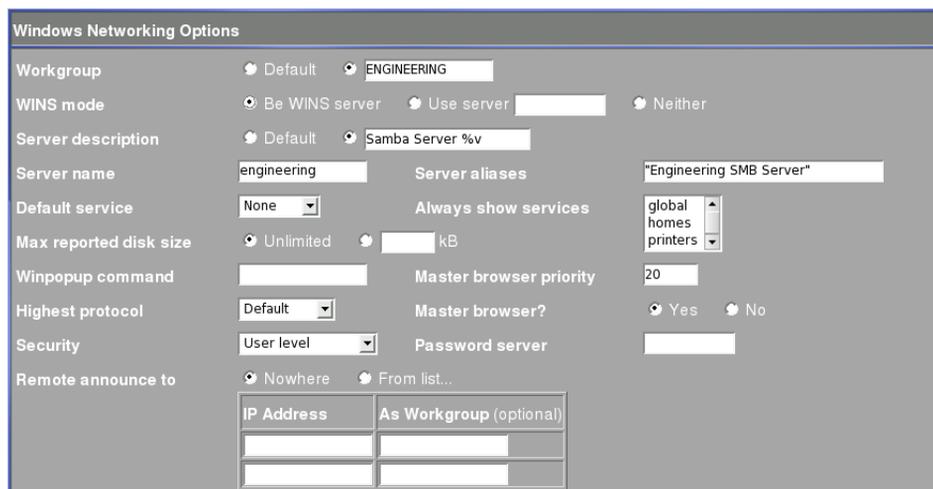


Figure 7-2. Configuring The Common Networking Options

Define a Workgroup for your server (ENGINEERING in our example). You can also set the Server name and the Server aliases. You can set the Samba server to be the WINS server of your network with the WINS mode option¹. Set Security to `User level` and validate your choices by clicking on Save (see figure 7-2).

7.1.2.2. Authentication Method

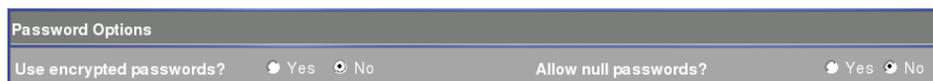


Figure 7-3. Setting The Authentication Method for Windows 95 Clients

If you have Windows[®] 95 clients on your network, click on the Authentication button and change the Use encrypted passwords? option to No as shown in figure 7-3.

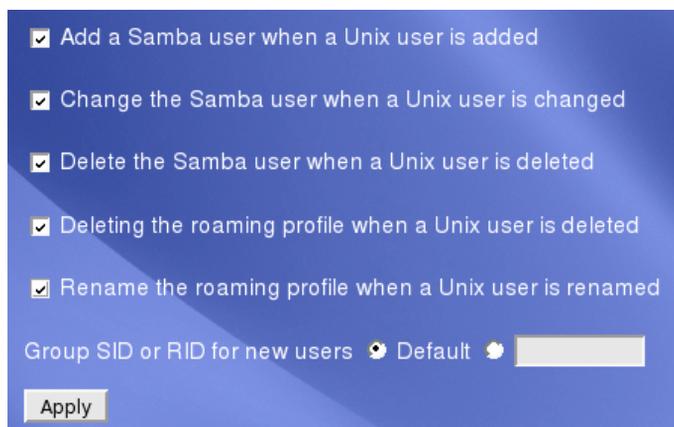


Figure 7-4. Synchronize Samba and Unix Users

If there are no Windows[®] 95 clients on your network click on the Configure automatic Unix and Samba user synchronisation link, put a check mark in all options as shown in figure 7-4 and then click on the Apply button.

To add the current Linux users on your system as Samba users, click on the Convert Unix users to Samba users link, make your changes to the settings or accept the default ones, and click on the Convert Users button. After adding users, you should click on Edit Samba users and passwords to modify and/or remove unwanted users.

1. You shouldn't mix Windows[®] and Samba WINS servers on your network.

7.1.2.3. Adding Shares

Share Information

Share name Public Home Directories Share

Directory to share ...

Automatically create directory? Yes No Create with owner ...

Available? Yes No Browseable? Yes No

Share Comment

Figure 7-5. Configuring a Public Share

To create a public share where **any** user can read and write files, click on the Create a new file share link and fill the form like the one shown in figure 7-5. Then click on the share name (`Public` in our example) and on the Security and Access Control button changing the Writable? and Guest Access? options to Yes. Save your changes and repeat the process to add other shared folder entries, setting access control appropriately.

Please bear in mind that all directories configured as shares will have to have proper Linux access rights in order to be readable/browsable/writable by Windows[®] users.

7.1.3. Advanced Configuration

7.1.3.1. Share Access

Security and Access Control

Writable? Yes No Guest Access? None Yes Guest only

Guest Unix user ... Limit to possible list? Yes No

Hosts to allow All Only allow:

Hosts to deny None Only deny:

Revalidate users? Yes No

Valid users ...

Valid groups ...

Invalid users ...

Invalid groups ...

Possible users ... Possible groups ...

Read only users ... Read only groups ...

Read/write users ... Read/write groups ...

Figure 7-6. Limiting Access

Select a share to be edited in the share list and click on the Security and Access Control button. Use the Hosts to allow and Hosts to deny options to specify a space-separated list of the IP addresses of the hosts allowed to connect, or not, to this share. If you set the Limit to possible list? option to Yes, then you must fill the Possible users and Possible groups fields with a space-separated list of the users/groups in question. See figure 7-6 for an example.

7.1.3.2. Default Printer



Figure 7-7. Printer Share Default Options

Even if all the Samba server's printers are available, you might want to set the Printer Share Defaults (see figure 7-7). Use the Unix Printer pull-down list to select the default printer and specify whether the printer will be available or not, the spooling directory (leave blank for default), a comment and security and access control options. Click on Save when you are satisfied with your settings.

7.1.4. Extra Documentation

Browsing the Samba Documentation (<http://samba.org/samba/docs/>) is a good idea. If you installed the `samba-doc` RPM, you can alternatively access Samba documentation on your own installation in the `/usr/share/doc/samba-doc-*/` directory.

7.2. Resource Sharing: FTP

ProFTPD allows you to create and set up an FTP server. With the latter, your company can share files with people connected to the Internet (or to your Intranet). Depending on your configuration, they could also upload files on to your server.

7.2.1. Installation

Make sure that the `proftpd` package is installed.

The server's configuration is done through Webmin's ProFTPD Server module located under the Servers category.

7.2.2. Configuration Examples



Using FTP is insecure because login names and passwords are not encrypted while transferring them to the server. Use non-anonymous FTP connections with logins and passwords only within trusted networks.

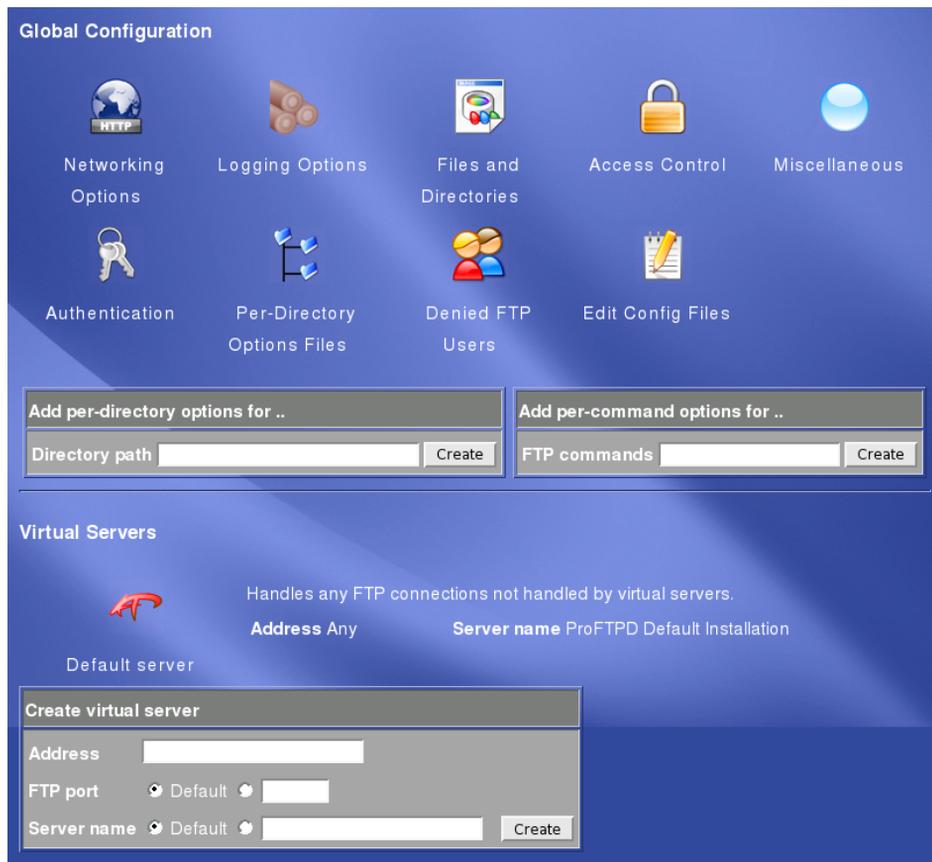


Figure 7-8. ProFTPD Module's Start-Up Screen

ProFTPD's configuration files are stored in the `/etc/proftpd.conf` file. ProFTPD's default configuration only allows users of your system with valid accounts to connect via FTP. In the following sections we present common FTP server configuration examples.

Like Apache, ProFTPD can be used to host several FTP sites on a single server by means of virtual servers. Global Configuration options will apply to all virtual servers. Each virtual servers' options are found in the Virtual Servers section.



The default values are already good for most configurations, so please do not change them unless you know what you are doing.

7.2.2.1. Simple Anonymous FTP Server

Simply install the `proftpd-anonymous` package. It adds a configuration file with everything needed for anonymous connections to the FTP server.

To test the anonymous FTP server access:

1. Connect to the hostname or IP of the FTP server using your favourite FTP client.
2. Enter `anonymous` as the login name and your e-mail address as the password.
3. If all went well, all files located in the `ftp` user directory root (the `/var/ftp` directory) and its subdirectories will be available for anonymous downloads.

7.2.2.2. Anonymous FTP Server With Uploads Directory

If you want to enable people to upload their files, you need to create a special place for those files to be stored. Let us add that functionality to our server.

Type the following in a terminal, as `root`, to create the storage space and change its access rights and owner.

```
# mkdir /var/ftp/uploads
# chown ftp:ftp /var/ftp/uploads
# chmod g+w /var/ftp/uploads
```

Click on the Default server link and then on the Anonymous FTP button. Fill the Directory path field of the Add per-directory options for.. table with uploads/* and click on the Create button.

We want users from our trusted network to read and store data in this directory. The following per-command options must be set:

- Enter STOR in the Add per-command options for table and click on the Create button, then click on Access Control and fill the form with the values as shown in figure 7-9.

Figure 7-9. Access Control For Commands in the Directory Page

In our example we allow anonymous storage of files into this directory only for network 192.168.0. and deny from any other networks. You can also allow/deny access for certain hostnames, IP addresses, users and groups of users.

- Click on the Return to per-directory options link and create the READ command and follow the instructions above, but this time to allow/deny reading of the uploads/* directory.

Return to the main menu and click on the Apply Changes button to restart the FTP server. All files located in the /var/ftp/uploads/ directory will be available for anonymous uploads and downloads.

You can repeat the above test procedure, but for the /var/ftp/uploads directory.

7.2.3. Advanced Configuration

Here is a list of useful options to configure:

- Under the Miscellaneous section of any of the virtual servers you can change the Server administrator's email address (ftp_admin@company.net, for example), and the Server name displayed to users (for example Main FTP server of my Company).
- Under the Networking Options section you may change the value of Maximum concurrent logins from the default (10) to an unlimited number or the number of connections you specify². You can also set the Login error message to a meaningful value, for example: Sorry, max %m users connected -- try again later., where %m will be replaced by the number set before. You can also set the port that virtual server listens to (21 for the main server).
- Under the Files and Directories section you can specify the Initial login directory. This is the default root directory and can be used to put users into a restricted environment. You can also use Limit users to directories to restrict all users to their home directories or to a directory of your choice.

2. Setting this value to 0 will result in the denial of all connections to the FTP server.

You can always use Edit Config files in the main screen of the ProFTPD module (see figure 7-8) if you need to edit ProFTPD's configuration file by hand.

7.2.4. More Documentation

Browse the ProFTPD Documentation (<http://www.proftpd.org/docs/>) where you'll also find configuration examples.

7.3. NFS: Exporting Directories To UNIX/Linux Hosts

NFS allows you to easily export directories of your computer to others through the network, thus enabling file sharing between several users. This type of sharing is easier to set up than Samba, but it is only used on GNU/Linux and UNIX[®] systems. NFS is very insecure and must be used only in a secure local network.

7.3.1. Installation

Make sure that the `nfs-utils` and `nfs-utils-clients` packages are installed.

7.3.2. Configuration Example

Exported directories configuration is done using the NFS Exports button. You will find it under the Networking category. Click on the Add a new export link to create a new export.



Simply click on a parameter you do not understand and a help window explaining that parameter will pop-up.

The screenshot shows the 'Export details' configuration window. It includes the following fields and options:

- NFS Version:** Radio buttons for 4 and 3 (or lower).
- NFSv4 Pseudofilesystem to export:** A text input field.
- Directory to export:** A text input field with a browse button, showing '/home/exports' in 'home/exports'.
- Active?:** Radio buttons for Yes and No.
- Export to.. (with or without Authentication):**
 - Radio buttons for Everyone, Host(s), WebNFS clients, NIS Netgroup, sys, IPv4 Network, IPv6 Address, krb5, lipkey, and spkm-3.
 - Under IPv4 Network, the IP address is '192.168.1.0' and the Netmask is '255.255.255.0'.
- Security level:** Radio buttons for None, Integrity, and Privacy (including Integrity).

Figure 7-10. Adding a New NFS-exported Directory

The new export creation window (see figure 7-10) is divided in two sections. Under Export details you should select NFS Version 3, to ensure optimum compatibility. Then specify the Directory to export and where you want to Export to. By default, the export is accessible to Everyone, this should be changed to the sub-network you actually use (for example: 192.168.1.0/255.255.255.0) or a netgroup.

Under Export security you can restrict the access to your exported directories even more. For example, you can choose which UIDs and GIDs to trust, whether to allow read only or read/write access, etc.

Once you are satisfied with your settings, click on the Save button to return to the main NFS exports configuration screen. You will now see the export entry you just added. You can add more exports or edit the existing

ones by clicking on the proper link in the Exported to.. column. Click on the Apply Changes button to make the exported directories available to clients.

7.3.3. Accessing the Exported Directories

You will have to configure the client computer(s) to mount the NFS-exported directories created. The settings vary from one OS to another, we will concentrate here on a Mandriva Linux client configuration.

The simplest way is to use the Mandriva Linux Control Center Set NFS mount points subsection of the Mount Points section.

You can also do it the Webmin way by entering the Disk and Network Filesystems interface into the System section. There you will be able to Add mount of type Network filesystem (nfs)

7.3.4. Extra Documentation

Joe Cooper dedicates a chapter to NFS exports in The Book of Webmin (<http://www.swelltech.com/support/webminguide/ch13.html#exports>). Though based on an old version of Webmin it might prove to be an interesting read.

Chapter 8. The Kolab Server

8.1. Introduction

Kolab is the server part of Kroupware, the KDE groupware solution. Kolab stores synchronization information such as addresses, calendar information and files which are useful to groups of users. Information stored by the Kolab server may be accessed using Kontact, the client part of the Kroupware project. Kontact is a combination of KMail, KOrganizer, KAddressbook, KNotes, KPilot and KNode. This chapter will give a technical overview of the Kolab server, then explain its installation, configuration and administration. User information for the Kontact client can be found in Mandriva Linux's *Starter Guide*. Information on the Kroupware project can be found on the Kroupware web site (<http://kroupware.kde.org>).

8.2. Overview

The Kolab administration interface is hosted on the enhanced Apache HTTP server. Login on to the administration module is only possible using a secure connection.

Kolab allows three levels of users:

- user. Can change personal user data.
- maintainer. User rights plus administrative rights over users, groups and shared folders.
- administrator. Maintainer rights plus rights to the whole LDAP tree, rights to toggle legacy services and view logs.

All user level login to the Kolab administration module using the same web interface. After the login has been authenticated using LDAP credentials, the user is presented with the choice of web forms permitted at his level. All new users are automatically created as regular users.

The LDAP server is hosted on the same physical machine as Kolab. All data used by Kolab, both user and configuration, is stored on the LDAP server, which is configured using the Kolab bootstrap script, `kolab_bootstrap -b`. By regularly backing up LDAP data on a separate machine it is possible to restore the state of the Kolab server even after a hardware disaster.

Kolab uses a special account **manager**, created at installation time for administrative users to manipulate all possible data on the LDAP server. Maintainers and users have different access rights to the LDAP tree. This means a potential security attacker, gaining control of the Apache server would not be able to manipulate account data on the machine so long as the **manager** password is unknown.

8.3. Installation

To install the Kolab server, use Rpm Drake (see *Rpm Drake: Package Management* in the *Starter Guide*) from the Mandriva Linux Control Center. You can also issue the `urpmi --auto kolab-server` command, on a terminal as `root`, to install it.

When Kolab is installed, the server has to be configured. This is done by launching the `/usr/sbin/kolab_bootstrap -b` script. The script will configure the LDAP server used to store Kolab configuration information and user data. While initializing the server, `kolab_bootstrap` will ask you to provide a password for the LDAP server manager account. You must keep this information, this account doubles as the Kolab administrator account and is used to login to the administration interface for the first time.



You will need a valid DNS entry for your hostname, as well as a valid MX record for your maildomain. If you only use `/etc/hosts`, you will be able to access most of the features, but you won't be able to receive mail from the outside. One solution could be to use `fetchmail` or a similar utility to retrieve mail from an outside POP server, but the ideal solution is to have your own domain name and fixed IP address accessible from the outside.

When the bootstrap script has run, issue the `service kolab-server start` command to start the Kolab server. Kolab is now running and ready for the first login.

In order to configure the Kontact client to work in conjunction with Kolab, you need essential LDAP connection information which you'll find in the `/etc/openldap/slapd.conf` file. For example:

```
# grep suffix /etc/openldap/slapd.conf
suffix          "dc=kolab,dc=yourdomain,dc=com"
```

Gives you the Base DN string: `dc=kolab,dc=yourdomain,dc=com`. Then for user `peter`, the Bind DN string will be `cn=peter,dc=kolab,dc=yourdomain,dc=com`.

8.4. The Kolab Administration Interface

All users log in to the Kolab administration module using the same web interface. If your Internet browser is on the same machine as the Kolab server, this will be found at `https://localhost/kolab/`. If you are logging in over the network the URL would be `https://hostname/kolab/`.



Since we're using a dummy, self-signed SSL certificate, your browser will complain that the "certificate failed the authenticity test". Don't worry, this is normal. Simply click Continue, then Accept forever and you're set. Depending on your DNS configuration, you may have to do this step twice.

The first login to the Kolab server has to be done using the `manager` account whose password was chosen when the `kolab_bootstrap` script was run. Once logged in the administrator can create Kolab user, maintainer and administrator accounts. For more information on how this is done, see *Maintainers*, page 72.



If the LDAP database is already populated, the administrator can view the existing user information. Existing LDAP users will not be able to log in to the Kolab administration module automatically. Kolab users have to be created explicitly using the administration module.

Once a new account is created, it is active. The Kolab user can log into the administration module.

After logging into the administration module, users are presented with the Kolab welcome page. A list of web forms available to the user is presented on the left-hand frame. The following sections explain the forms available for the different user types and the information necessary to use them correctly.

8.4.1. Users

Members of the regular groupware users group have the right to:

- modify personal user data;
- add an additional e-mail address for his/her account;
- activate a vacation messaging service;
- activate an e-mail forwarding service.



Figure 8-1. The Kolab Server User Interface



The vacation message service and e-mail forwarding services are mutually exclusive. The user has to explicitly deactivate one service to activate the other.

8.4.1.1. Changing User Data

Members of the regular users group cannot:

- change their user name;
- change their unique user identification (UID);
- change their primary e-mail address;

to change user data.

Modify Existing User

Attribute	Value	Comment
First Name	the	
Last Name	user	
Password	<input type="password" value="*****"/>	Required
Verify Password	<input type="password" value="*****"/>	Required
UserName	userone	Cannot be modified
Title	<input type="text" value="Mr"/>	
E-Mail Alias	<input type="text" value="userone@myorg.org"/>	
Organization	<input type="text" value="myorg"/>	
Organizational Unit	<input type="text"/>	
Room Number	<input type="text"/>	
Street Address	<input type="text"/>	
Postbox	<input type="text"/>	
Postal Code	<input type="text"/>	
City	<input type="text"/>	
Country	<input type="text"/>	
Telephone Number	<input type="text"/>	
Fax Number	<input type="text"/>	

Figure 8-2. Modify Existing User Data

1. Login to the Kolab Administration Interface.
2. Click **My User Settings** in the left-hand panel of the Kolab administration interface.
3. In the Modify Existing User web form change the user information desired.

4. Click OK.

8.4.1.2. Activate Vacation Settings

On the occasions when a user is on vacation and cannot or will not check their e-mail, it is helpful to inform other users who are to be contacted in their absence. This can be done using the Kolab vacation message service.

To activate the Kolab vacation service:

1. Login to the Kolab administration interface.
2. Click My User Settings in the left-hand panel of the Kolab administration interface.
3. Click Vacation. The User Vacations Settings page will open.
4. Select the duration of your vacation in the drop-down combo box.
5. Type the message you wish to be shown to other users in the Vacation text box.
6. Click OK.

8.4.1.3. Activate E-mail Forward Settings

Kolab supplies an e-mail forwarding service to assist on occasions when users are on external missions or vacation and cannot contact the standard e-mail of their organization but have the option to use a different e-mail address. When the e-mail forwarding service is activated, e-mails are forwarded to the desired address. Optionally a copy of these e-mails can be stored on the organizations e-mail server.

To activate the Kolab e-mail forwarding service:

1. Login to the Kolab administration interface.
2. Click My User Settings in the left hand panel of the Kolab Administration Interface.
3. Click Forward E-Mail and the User Forward Settings page will open.
4. Type the address the e-mail is to be forwarded to.
5. If you wish to keep a copy of your e-mail on the local server, click the Keep check-box.
6. Click OK.

8.4.2. Maintainers

The role of the maintainer group is to administer users and shared folders on the Kolab server. The following rights are available to maintainers in addition to basic user rights:

- add, modify and delete Kroupware users;
- add, modify and delete address-book users, that is to say those users in the LDAP directory who are not registered on the Kolab server;
- add modify and delete shared folders.



Activities related to the maintainer groups basic user rights are documented in *Users*, page 70.



Figure 8-3. The Kolab Server Maintainer Interface

8.4.2.1. Kroupware Users

Kroupware users are fully registered Kolab users and profit from all Kolab groupware services. Kolab maintainers can manage Kroupware user data and create new address-book users.

8.4.2.1.1. Creating New Users.

To add a new user:

1. Click Create New User in the Users section of the left-hand frame.
2. Fill in the necessary user information in the Create New User web form.
3. Click OK.



Select the Addressbook check-box to allow the Kroupware user address information to be visible in the address book.

8.4.2.1.2. Managing Existing Users

When a member of the maintainers group clicks the Users button on the left hand-frame of the Kolab administration interface, an alphabetical list of all current Kroupware users is presented in the right-hand side. For each user there is a choice to modify or delete the user.

- To modify user information, click the Modify button. This will lead you to the Modify Existing User web page. This page is explained in the Users section (see *Users*, page 70).
- To delete a user, click the Delete button.

8.4.2.2. Address-Book Users

Address-book users are users which exist in the LDAP directory but are not registered Kolab users. The latter can access their address-book data; Kolab maintainers can manage address-book user data and create new address-book users.

8.4.2.2.1. Creating New Users

Attribute	Value	Comment
First Name	<input type="text"/>	Required
Last Name	<input type="text"/>	Required
Title	<input type="text"/>	
Primary E-Mail Address	<input type="text"/>	
E-Mail Alias	<input type="text"/>	
Organization	<input type="text"/>	
Organizational Unit	<input type="text"/>	
Room Number	<input type="text"/>	
Street Address	<input type="text"/>	
Postbox	<input type="text"/>	
Postal Code	<input type="text"/>	
City	<input type="text"/>	
Country	<input type="text"/>	
Telephone Number	<input type="text"/>	
Fax Number	<input type="text"/>	

Figure 8-4. The Create New Address Book Form

To add a new address-book user:

1. Click Create new VCard in the Users section of the left-hand frame.
2. Fill in the necessary user information in the Create New Address Book Entry web form.
3. Click OK.

8.4.2.2.2. Managing Existing Users

When a member of the maintainers group clicks the Addressbook button on the left-hand side of the Kolab administration module, an alphabetical list of all current address-book users will be presented in the right-hand frame. For each user there is a choice to modify or delete the user.

Manage Address Book
(only external addresses without a kolab user account)
[\[A-F\]](#) [\[G-L\]](#) [\[M-S\]](#) [\[T-Z\]](#) [\[all\]](#)

(No external address book entries in the category [A-F])

Figure 8-5. The Manage Address Book Users Table

The following actions can be done on the Manage Address Book Users form:

- To modify user information:
 1. Click the Modify button.
 2. In the Modify Address Book Users page change the information desired.
 3. Click OK.
- To delete a user, click the Delete button.

8.4.2.3. Managing Shared Folders

Maintainers can reconfigure existing shared folders and create new shared folders.

8.4.2.3.1. Creating New Shared Folders

Attribute	Value	Option	Comment
Folder Name	<input type="text"/>		
Permissions for UID:	<input type="text"/>	none read post append write all	
Quota Limit (KByte)	<input type="text"/>		

Figure 8-6. The Create New Shared Folder Form

To add a new shared folder:

1. Click Add Folder in the Shared Folder section of the left-hand frame.
2. Fill in the necessary user information in the Create New Shared Folder web form.
3. Click OK.

8.4.2.3.2. Configuring Shared Folders

When a member of the maintainers group clicks the Shared Folder button on the left-hand frame of the Kolab administration module, an alphabetical list of all current shared folders is presented in the right-hand frame. For each folder there is a choice for it to be modified or deleted.

- To modify a folder:
 1. Click the Modify button.
 2. In the Modify Shared folder page change the information desired.
 3. Click OK.
- To delete a folder, click the Delete button.

8.4.3. Administrators

The administrator group has full control over all objects on the LDAP server and non secure legacy services such as FTP, HTTP, IMAP and POP3. Administrators can fulfill all the functions of the maintainers group. In addition being able to add, modify or delete groupware users, address-book users and shared folders, administrators have the following rights:

- add, modify or delete accounts for users in the maintainer and administrator groups;
- change server settings (hostname and mail domain);
- toggle non secure legacy services (FTP, HTTP, IMAP and POP3).



Activities related to the Administrator groups maintainer and basic user rights are documented in *Users*, page 70, and *Maintainers*, page 72.

8.4.3.1. Managing Maintainer Accounts

Administrators have the right to modify or delete an existing maintainer account and add a new maintainer account.

8.4.3.1.1. Creating New Maintainers

To add a new maintainer:

1. Click Add Maintainer in the Maintainer section of the left hand frame.
2. Fill in the necessary user information in the web form.
3. Click OK.

8.4.3.1.2. Managing Existing Maintainers

When a member of the administrators group clicks the Maintainers button on the left-hand frame of the Kolab administration module, an alphabetical list of all current maintainer users is presented in the right-side frame. For each user there is a choice to modify or delete the user.

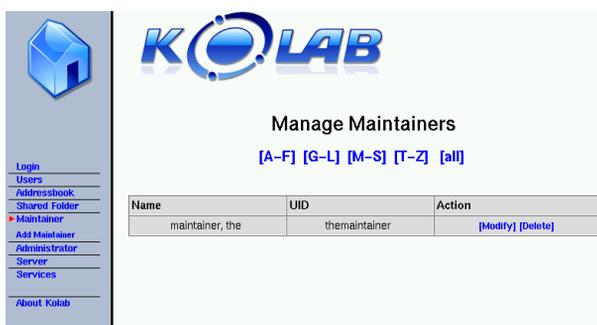


Figure 8-7. The Manage Maintainers Table

- To modify maintainer information:
 1. Click the Modify button for the maintainer required.
 2. This will open the Modify Existing User page.
 3. Update the information desired.
 4. Click OK.
- To delete a maintainer, click the Delete button.

8.4.3.2. Creating New Administrators

It is always safer to have at least one back-up administrator for any server. Kolab allows the creation of multiple administrators.

Attribute	Value	Comment
First Name	<input type="text"/>	Required
Last Name	<input type="text"/>	Required
Password	<input type="text"/>	Required
Verify Password	<input type="text"/>	Required
Unique UserID	<input type="text"/>	Required

Figure 8-8. The Create New Administrator Form

To add a new administrator:

1. Click Administrator in the left-hand pane of the administration module.
2. Click Add Administrator.
3. Fill in the new administrator information.
4. Click OK.

8.4.3.3. Changing Server Settings

Members of the administer group can change the hostname and domain name of the e-mail domain on the host server.



Changes in these settings can directly affect the mail transport system and can result in e-mail delivery problems.

To change server settings click Server in the left-hand frame of the administration module, fill in the necessary information as shown in the screen shot below and click OK.

Attribute	Value	Comment
Hostname	<input type="text" value="localhost"/>	This hostname will be given by the Mail Server and the IMAP Server to the clients
E-Mail Domain	<input type="text" value="localhost"/>	Be advised that renaming the E-Mail Domain affects all E-Mail Addresses!

Figure 8-9. The Server Settings Form

8.4.3.4. Toggling Services

The Kolab administrator has the right to enable or disable the following services on the host machine:

- POP3;
- POP3/TLS service (TCP port 995);
- IMAP/TLS service (TCP port 993);
- FTP free-busy service;
- HTTP free-busy service.

To enable or disable the service required, click on Services in the left-hand pane of the administration module. You will see the following web form.

Enable or Disable individual Services

Using legacy services poses a security threat due to leakage of cleartext passwords, lack of authenticity and privacy.

The legacy Freebusy Support (FTP and HTTP) is only required for Outlook2000 clients. Under all other circumstances it is advised to use the secure [WebDAV](#) over TLS instead (WebDAV is enabled by default and may not be deactivated).

Further details with regards to security considerations are available on the internet at the [Kolab](#) webserver.

Service	Status	Action
POP3 service	active	disable pop3
POP3/SSL service (TCP port 995)	active	disable pop3s
IMAP service	active	(may not be deactivated)
IMAP/SSL service (TCP port 993)	active	disable imaps
Sieve service (TCP port 2000)	active	disable sieve
FTP free-busy service	disabled	activate ftp
HTTP free-busy service	disabled	activate http

Figure 8-10. The Services Form

In the form you can see the status of the different services. To activate or deactivate, click the required URL in the action column.

Chapter 9. MySQL Database Server

Databases are responsible for storing data (mainly text and numbers) and delivering it back in an efficient manner. They are generally used by other applications which need to quickly access data to process or display it.

MySQL is a true multiuser, multithreaded SQL (Structured Query Language) database server. MySQL is a client/server implementation which consists of a server daemon (`mysqld`) and many different client programs and libraries. The main goals of MySQL are speed, robustness and ease of use.

9.1. Getting Started

We cover basic MySQL configuration and usage through the Webmin interface. Make sure that the MySQL package is installed.

The MySQL Database Server configuration button is in the Servers index. If you just installed MySQL, you will be asked to start it: click on the Start MySQL Server button.



If you get an error when trying to start or stop the server, open the Module Config tab, and make sure the Command to start/stop MySQL server actually uses the `/etc/rc.d/init.d/mysqld` command.

In the main screen notice that there are three default databases (`mysql`, `test` and `tmp`). You should **not** modify, nor erase them.

Your first task is to set up the administrator's password. **This is mandatory to prevent other users on the machine from having unlimited access to the database.** To do so, click on the User Permissions icon then perform the following operation for the `root` links in the users table.

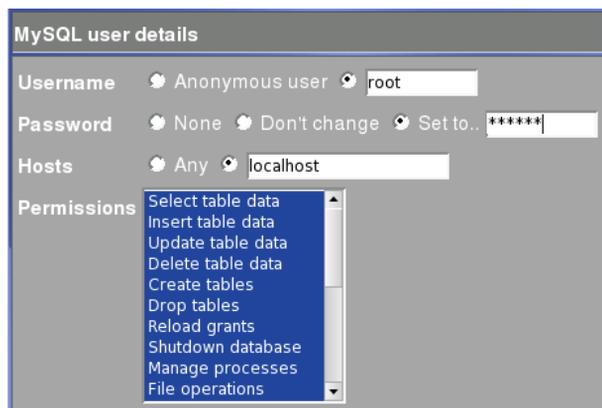


Figure 9-1. Setting the Administrator's Password

Select the Set to option of the Password line and enter the password for the administrator. Confirm your changes by clicking on the Save button. Once you set the password, you have to reconnect to the MySQL server, go to the Module Index tab and enter `root` as the user name and the password you've just set.



For security reasons, network access to the MySQL database is disabled by default. If you need applications outside of the system on which MySQL is installed to access the database, remove the `/etc/sysconfig/mysqld` file from your system.

In some cases, local applications are doing network requests to query the database, and are therefore blocked even though the application is actually running on the same machine.

9.2. Creating a User for the Database

A database user has nothing to do with a UNIX user. Therefore you have to manage them differently. From the index page of the MySQL Database Server component, click on User Permissions and then on Create new user. The database users may have specific permissions, which are all listed in that window. Select from the list the permissions that the user will be granted from the specified host.

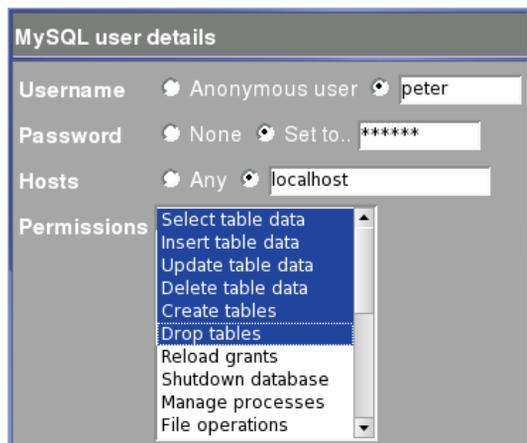


Figure 9-2. Creating a MySQL User

For security reasons don't leave the Hosts value set to **Any**. Do specify the hostname(s) from which the user is to be granted access to the MySQL database.

9.3. Creating a Database

Let's create the database which will hold our tables and data. Click on Create a new database from the main page and write in a name for your database.

Field name	Data type	Type width	Key?	Autoinc?	Allow nulls?	Unsigned?	Default value
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="checkbox"/> Yes	<input type="checkbox"/> Yes	<input checked="" type="checkbox"/> Yes	<input type="checkbox"/> Yes	<input type="text"/>
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="checkbox"/> Yes	<input type="checkbox"/> Yes	<input checked="" type="checkbox"/> Yes	<input type="checkbox"/> Yes	<input type="text"/>
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="checkbox"/> Yes	<input type="checkbox"/> Yes	<input checked="" type="checkbox"/> Yes	<input type="checkbox"/> Yes	<input type="text"/>
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="checkbox"/> Yes	<input type="checkbox"/> Yes	<input checked="" type="checkbox"/> Yes	<input type="checkbox"/> Yes	<input type="text"/>

Figure 9-3. Creating a MySQL Database

You can add a first table to that new database and define up to four fields in that table at this point. For our example we decided to do that in the next step. Note that if you intend to use your database with a third-party application, such as a web frontend, everything you need should be available at this step: a user and its password plus a running database, ready to store data.

9.4. Creating a Table

Once the database is created, it is possible to define its structure manually with Webmin. Click on its icon to access the Edit Database page. Note that there are no tables for the moment, but from this step we can Create a new table, Drop Database (delete it) or Backup Database. You can also Execute SQL directly by entering SQL commands or uploading an SQL command file. You can select the number of fields the new table should contain (4 by default) before clicking on the Create a new table button.

In the **Create Table** page, you must write the table's name and the field parameters.

Figure 9-4. Creating a MySQL Table

If you want to modify the table parameters or add new fields, click on the table's name from the Edit Database page.

Field name	Type	Allow nulls?	Key	Default value	Extras
Number	int(11)	No	Primary		auto_increment
Name	tinytext	Yes	None		
birth	date	Yes	None		

Figure 9-5. Modifying a MySQL Table

From here you can either modify or delete an existing field, or create a new one, view the table's data or delete (Drop Table), the whole table and its data.

9.5. Managing Data in a Table

We created everything needed to start filling the MySQL database. There are a number of client programs you can use to connect to the MySQL server as well as many programs which use such databases. You can also use Webmin to manage your data. In the Edit Table page, click on the View Data button to add, modify, or remove data.

	Number	Name	birth
<input type="checkbox"/>	1	William	1974-05-18
<input type="checkbox"/>	2	Arnold	1246-02-15
<input type="checkbox"/>	3	Michael	1845-09-01

Figure 9-6. Managing your Data

9.6. More Documentation

You will find an extensive set of documents on the MySQL web site (<http://www.mysql.com/documentation/>), including documentation translated into many languages.

Chapter 10. NIS Client and Server

To simplify user management on your local network, you can centralize network information such as user name and password lists on a NIS (Network Information System) domain.

With NIS users can connect on any computer using the same login and password. Information sharing allows you to distribute files such as `/etc/passwd`, `/etc/shadow` or `/etc/hosts` to share machine passwords or aliases. To distribute the data, you have to configure a Resource Sharing server such as NFS (see *NFS: Exporting Directories To UNIX/Linux Hosts*, page 67) or Samba (see *Samba: Integrating Linux in a Windows Network*, page 61) and use autofs (*Importing homes with autofs*, page 84).

10.1. Installation

First, you need to check that the `ypserv` server is installed on your computer. If it isn't use the `rpm` application or type `urpmi ypserv` in a terminal to install it.

The configuration of the NIS server is done in two steps: the first one is the configuration of the NIS tables¹ of the server; the second step is the configuration of each client.

- On the server, you need to install the `ypserv` RPM package.
- For each client, install `portmap`, `yp-tools` and `ypbind`.

To use Webmin's NIS Client and Server module, you must select the Networking category, then click on NIS Client and Server.

10.2. Step-by-Step Configuration

10.2.1. NIS Server

After clicking on the NIS Server icon you need to configure the NIS domain using your domain name (such as `mydomain.test`). Then choose the NIS tables to serve. For our example, we selected the `passwd`, `group` and `shadow` files (use the **Ctrl** key to select several files in the list).

NIS server options	
Enable NIS server?	<input checked="" type="radio"/> Yes <input type="radio"/> No
Serve NIS domain	<input checked="" type="radio"/> Same as client <input type="radio"/> Different <input type="text" value="mydomain.test"/>
Server type	<input checked="" type="radio"/> Master NIS server <input type="radio"/> Slave of server <input type="text"/>

Master NIS server options	
Lookup missing hosts in DNS?	<input checked="" type="radio"/> Yes <input type="radio"/> No
Push updates to slaves?	<input checked="" type="radio"/> Yes <input type="radio"/> No
NIS tables to serve	<input type="text" value="passwd group hosts rpc services"/>
Minimum UID for 'Unix user' table records	<input type="text" value="500"/>
Minimum GID for 'Unix group' table records	<input type="text" value="500"/>

Slave servers	
<input type="text"/>	

Figure 10-1. NIS Server

You don't need to modify the file description in the Master NIS files section. On the NIS Client and Server menu, the NIS Tables icon allows you to modify the tables which are served. The Server Security icon allows you to select the clients you want to serve.

1. The NIS tables are the files you chose to export.

10.2.2. NIS Client

For each client, go to the NIS module and enter the NIS client configuration screen. You have to configure the NIS `domain` name parameter with the domain name used by the server. You must also enter its IP address. That's all.



Figure 10-2. NIS Client

10.3. Advanced Client Configuration

Amongst all the exported data, some may be redundant with regard to the local configuration. You can choose the priority to be allocated to every source (local, NIS or other). To do so, use the Client Services button on each NIS client. It allows you to choose the preferred order to search for data. For example, you can choose to resolve the host's addresses using:

1. the `/etc/hosts` file;
2. the NIS hosts served (if you selected it in the NIS table);
3. and finally (if the client cannot resolve any more), use the DNS server.

To verify that the client communicates with the server, you can use the `ypcat passwd` command to read the password data served by the server.

10.4. Importing homes with autofs

If you export your user's directories (using NFS for example), and if you configure the **autofs** service on your NIS client, users' home directories are automatically mounted when they log onto the client. This way everybody can automatically log on every client, and have all their personal data and configuration files available to them.

Security, Network and Troubleshooting Issues

To help you deal with more complex issues, we added the following three chapters to complete your Mandriva Linux know-how.

- The chapter which discusses security (*"Security Under GNU/Linux"*, page 87) is a **must read** for any system administrator. Even though you can make your Mandriva Linux system quite secure with default tools, efficient security can only be achieved through active administration, taking care of physical, logical and global system security. This chapter will help you choose the appropriate security policy for your systems, the right means to secure your network infrastructure, how to determine if your system has been altered or compromised, and more.
- The networks' configuration and some concepts of the TCP/IP protocol — the most commonly used ones — are tackled in (*"Networking Overview"*, page 123). Whenever possible pointers to information about other network protocols is provided.
- In an effort to try to fight Murphy's law, we wrote the troubleshooting chapter (*"Troubleshooting"*, page 137) to save you sleepless nights. It also covers disaster prevention, so don't read this part when it's already too late!

Chapter 11. Security Under GNU/Linux

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



The original document has been adapted to the Mandriva Linux distribution, removing some parts, changing others, etc.

11.1. Preamble

This chapter is based on a HOWTO by Kevin Fenzi and Dave Wreski. The Linux Documentation Project (<http://www.tldp.org>) hosts the original document.

11.1.1. Copyright Information

This document is copyrighted (c) 1998 - 2000 Kevin Fenzi and Dave Wreski.

Modifications from v2.0, 11 June 2002, (C)opyright 2000-2004 Mandriva and distributed under the following terms:

- Linux HOWTO documents may be reproduced and distributed in whole or in part, in any medium, physical or electronic, as long as this copyright notice is retained on all copies. Commercial redistribution is allowed and encouraged; however, the authors would like to be notified of any such distributions.
- All translations, derivative works, or aggregate works incorporating any Linux HOWTO documents must be covered under this copyright notice. That is, you may not produce a derivative work from a HOWTO and impose additional restrictions on its distribution. Exceptions to these rules may be granted under certain conditions; please contact the Linux HOWTO coordinator at the address given below.
- If you have questions, please contact the Linux HOWTO coordinator Tim Bynum (<mailto:tjbynum@metalab.unc.edu>).

11.1.2. Introduction

This chapter covers some of the main issues which affect GNU/Linux security. General philosophy and net-born resources are also 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 an 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 you some general methods to prevent such exploits from taking place.

11.2. 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.

11.2.1. 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, or worst, make alterations. 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 could then use advanced knowledge to impersonate you, steal information from you, or even deny you access to your own resources. See Eric Raymond’s How to Become a Hacker (<http://www.catb.org/~esr/faqs/hacker-howto.html>) if you’re wondering what the difference between a “hacker” and a “cracker” is.

11.2.2. How Secure Is Secure?

First, bear 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 log in. You could also set up your GNU/Linux system with no network, nor connection to the Internet, but this limits its usefulness.

If you are a medium-to-large-size 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 [faqs.org](http://www.faqs.org/rfcs/rfc2196.html) (<http://www.faqs.org/rfcs/rfc2196.html>). It contains a great framework for establishing a security policy for your company.

11.2.3. 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 log in using a `.rhosts` file, or to use an insecure service, such as `tfptp`, 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 level of threat they might pose.

There are several types of intruder, 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 purpose. 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 using your system to get access to 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 dynamic PPP 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 can save ten times more time later if you have to recreate data that was lost. Have you checked your backup strategy and verified your data lately?

11.2.4. 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 are 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 account?), who is 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 should not be using that service until you do grant access. Make sure the policies work on your regular user account. Saying, “Ah, I cannot figure out this permissions problem, I'll just do it as root”, can lead to security holes that are very obvious, and even ones which haven't been exploited yet.

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

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

To see what some real-life security policies look like, you can take a look at the COAST policy archive (<ftp://coast.cs.purdue.edu/pub/doc/policy>).

11.2.5. Means of Securing your Site

This section discusses various means by which you can secure the assets you have worked hard for: your local computer, 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 the next quarter? If you are planning a network installation, there are many factors you must take into account before adding a single computer to your network.

Even if you have a single dial up PPP account, or just a small site, this does not mean intruders will not 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 are connected to.

Intruders often have a lot of time on their hands, and can guess how you have 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.

11.2.5.1. 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. Although this is absolutely necessary, it can become a daunting task once your network becomes larger than a few computers.

11.2.5.2. Local Network Security

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

This document discusses some of the techniques which can be used to secure your site, and hopefully shows you some of the ways to prevent intruders from gaining access to that which you are trying to protect.

11.2.5.3. 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 the hope that an attacker won't notice it's there and thus won't exploit it. Rest assured that they can determine that it is there and will exploit it. Security through obscurity is no security at all. Simply because you have a small site, or a relatively low profile, does not mean an intruder will not be interested in what you have. We will discuss what you are protecting in the next sections.

11.2.6. 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 90, covers how to protect your physical machine from tampering. The second, *Local Security*, page 93, describes how to protect your system from tampering by local users. The third, *Files and File-System Security*, page 95, shows you how to set up your file systems and permissions on your files. The next, *Password Security and Encryption*, page 99, discusses how to use encryption to better secure your machine and network. *Kernel Security*, page 105 discusses what kernel options you should set or be aware of for a more secure system. *Network Security*, page 108, describes how to better secure your GNU/Linux system from network attacks. *Security Preparation (Before You Go On-Line)*, page 114, discusses how to prepare your machine(s) before bringing them on-line. Next, *What to Do During and After a Break-in*, page 115, discusses what to do when you detect a system compromise in progress or detect one that has recently happened. In *Security Sources*, page 117, some primary security resources are enumerated. The Q & A section *Frequently Asked Questions*, page 119, answers some frequently-asked questions, and finally, a conclusion in *Conclusion*, page 120.

The two main points to understand 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;
- 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.

11.3. 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 computer? Should they? Can you protect your computer 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 computer 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 computers. Many of the following sections will help out. If you are in an office, you may or may not need to secure your computer after-hours or while you are away. In 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.

11.3.1. Computer Locks

Many modern computer 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 computer, or opening up the case and directly manipulating or stealing your hardware. Some locks can even 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 computers, 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 skills.

Some computers (most notably SPARCs and Macs) have a dongle on the back: 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 computer.

11.3.2. BIOS Security

The BIOS is the lowest level of software to configure or manipulate your x86-based hardware. LILO and other GNU/Linux boot methods access the BIOS to determine how to boot your GNU/Linux computer. Other hardware that GNU/Linux runs on has similar software (Open Firmware on Macs and new Suns, Sun boot PROM, etc...). You can use your BIOS to prevent attackers from rebooting your computer and manipulating your GNU/Linux system.

Many PC BIOSes let you set a boot password. This does not provide much security (the BIOS can be reset, or removed if someone can get into the case), but may be a good deterrent (i.e. it will take time and leave traces of tampering). This may slow attackers down.

Another risk of trusting BIOS passwords to secure your system is the default password problem. Most BIOS makers do not expect people to open up their computer and disconnect batteries if they forget their password and have equipped their BIOSes with default passwords which work regardless of your chosen password. Some of the more common passwords include:

```
j262
AWARD_SW
AWARD_PW
lkwpeter
Biostar
AMI
Award
bios
BIOS
setup
cmos
AMI!SW1
AMI?SW1
password
hewittrand
shift + s y x z
```

I tested an Award BIOS and AWARD_PW worked. These passwords are quite easily available from manufacturers' web sites and astalavista (<http://astalavista.box.sk>) and as such a BIOS password cannot be considered adequate protection from a knowledgeable attacker.

Many x86 BIOSes 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 BIOSes disallow booting from floppy drives and some require passwords to access some BIOS features.



If you have a server computer, and you set a boot password, your computer will not boot up unattended. Bear in mind that you will need to come in and supply the password in the event of a power failure.

11.3.3. Bootloader Security

Bear in mind when setting these passwords that you need to remember them. Also remember that these passwords will only slow the determined attacker. They will not prevent someone from booting from a floppy and mounting your root partition.

If you are using security in conjunction with a bootloader, you might as well disable booting from a floppy in your computer's BIOS, and password-protect the BIOS.

If you are using security in conjunction with a bootloader, you might as well password-protect the PROM.



Once again, if you have a server computer, and you set up a boot password, your computer will not boot up unattended. Bear in mind that you will need to come in and supply the password in the event of a power failure!

11.3.3.1. With LILO

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

Please refer to `lilo.conf(5)` for more information on the `password` and `restricted settings`.

Also bear in mind that the `/etc/lilo.conf` will need to be mode `600` (readable and writing for root only), or others will be able to read your boot passwords!

11.3.3.2. With GRUB

GRUB is quite flexible when it comes to password setting: your default configuration file, `/boot/grub/menu.lst`, may contain a line allowing the loading of a new configuration file with different options (this new file may contain a new password to access another third configuration file and so on).

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

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

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

Please refer to the GRUB `info` page for more information.

11.3.4. xlock and vlock

If you wander away from your computer 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: `xlock` and `vlock`.

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

`vlock` is a simple little program which 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 computer or otherwise disrupting your work. It also does not prevent them from accessing your computer from another computer 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. For this reason, you might consider only using it while under control of KDM (or other login manager).

11.3.5. Security of Local Devices

If you have a webcam or a microphone attached to your system, you should consider if there is some danger of an attacker gaining access to those devices. When not in use, unplugging or removing such devices might be an option. Otherwise you should carefully read and look at any software which provides access to such devices.

11.3.6. Detecting Physical Security Compromises

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

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 discrepancies.

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 computer 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 form, allowing an intruder to view data as it is being transferred. This may reveal information about your network which is not intended to be public. There are `syslog` daemons available which 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 in *Keep Track of your System Accounting Data*, page 115.

11.4. 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 “upgrade” their normal user access to `root` access using a variety of bugs and poorly set up 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 do not know or for whom you have no contact information is a very bad idea.

11.4.1. 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 computer:

- give them the minimal amount of privileges they need;
- be aware when/where they log in from, or should be logging in from;
- make sure you remove inactive accounts, which you can determine by using the `last` command and/or checking log files for any activity by the user;
- 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 `user-ids` should be absolutely prohibited. User accounts also provide accountability, and this is not possible with group accounts.

Many local user accounts which are utilized in security compromises have not been used in months or years. Since no one is using them, they provide the ideal attack vehicle.

11.4.2. Root Security

The most sought-after account on your computer is the `root` (superuser) account. It has authority over the entire computer, which may also include authority over other computers 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 deleting. Using `echo` in place of destructive commands sometimes works too.
- 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 kinds 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 log in 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 computer, there are a few tools which can help. `sudo` allows users to use their own password to access a limited set of commands as `root`. This could 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 to keep track of changes made.

Although `sudo` can be used to give specific users special privileges for particular tasks, it does have several shortcomings. It should be used only for a limited set of tasks, such as 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.

11.5. Files and File-System Security

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

- There should never be a reason for users' home directories to allow SUID/SGID programs to be run from them. Use the `nosuid` option in `/etc/fstab` for partitions which are writable by users other 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 file systems using NFS, be sure to configure the `/etc/exports` file 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 96.
- If you are mounting file systems using a network file system such as NFS, be sure to configure `/etc/fstab` with suitable restrictions. Typically, using `nODEV`, `nosuid`, and perhaps `noexec`, are desirable.
- Set file system limits instead of allowing unlimited as 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, to restrict the number of processes to 50, and to restrict memory usage per user to 5MB.

You can also use the `/etc/login.defs` configuration file to set the same limits.

- 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 which must be protected. It also prevents someone from creating a hard link to the file. See `chattr(1)` 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 back door 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 if you feel it is absolutely 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.

- Un-owned files may also be an indication that 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 they 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.

11.5.1. 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 permission settings, the user could inadvertently give read or write permission to someone who should not have it. 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`.



In Mandriva Linux, it is only necessary to use `002` for a `umask`. This is due to the fact that the default configuration is one user per group.

11.5.2. File Permissions

It is important to ensure that your system files are not open for casual editing by users and groups who should not 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 a file. Permissions for directories may have a different meaning to the same set of permissions on a file.

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 had explicit write permission granted, even when he has write access to the directory. This is designed for directories such as /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 the 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 in the same way as `suid`, except that 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 queen users      114 Aug 28 1997 .zlogin
1st bit - directory?          (no)
2nd bit - read by owner?      (yes, by queen)
3rd bit - write by owner?     (yes, by queen)
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 required to perform the access described. You may want to give more permission than those 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 queen users      512 Sep 19 13:47 .public_html/
1st bit - directory?          (yes, it contains many files)
2nd bit - read by owner?      (yes, by queen)
3rd bit - write by owner?     (yes, by queen)
4th bit - execute by owner?   (yes, by queen)
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 required to perform the access described. You may want to give more permission than those 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 the `/etc` directory) are usually mode `640` (`-rw-r-----`), and owned by `root`. Depending on your site's security requirements, you might want to adjust this. Never leave any system files writable by a group or everyone. Some configuration files, including the `/etc/shadow` one, 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 still be exploited to give the cracker a `root` shell.

11.5.3. Integrity Checking

Another very good way to detect local (and also network) attacks on your system is to run an integrity checker such as Tripwire, Aide or Osiris. These integrity checkers run a number of checksums on all your important binaries and configuration 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 is a good idea to install these types of program onto a floppy, and then physically set the write protect on the floppy. This way intruders cannot tamper with the integrity checker itself or change the database. Once you have something like this setup, it is 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=queen
# 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 to determine which is cracker activity and which is your own doing.

You can find the freely available unsupported version of Tripwire on the TripWire web site (<http://www.tripwire.org>) free of charge. Manuals and support can be purchased.

Aide can be found on Sourceforge (<http://sourceforge.net/projects/aide>).

OSIRIS can be found on the OSIRIS web site (<http://osiris.shmoo.com/>).

11.5.4. 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 systems while they are not paying attention. While they think the binary they just pulled down does one thing (and it might very well do so), it also compromises their security.

You should take care of what programs you install on your computer. Mandriva provides MD5 checksums and PGP signatures on its 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.

11.6. Password Security and Encryption



Most of the encryption programs described in this chapter are available in your Mandriva Linux distribution.

One of the most important security features used today are passwords. It is important for both you and all of your users to have secure, unguessable passwords. Your Mandriva Linux distributions include a `passwd` program that does not allow you to set an easy to guess 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 UNIX® systems (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 Unixes 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 “Crack” and “John the Ripper”, page 104) 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.

To obtain information on how to choose a good password, check the CERN web site (http://consult.cern.ch/writeup/security/security_3.html).

11.6.1. 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, 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 different PGP FAQs available on the Internet FAQ Archives (<http://www.faqs.org/faqs/pgp-faq/>).

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/~sasha/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 metalab (<ftp://metalab.unc.edu/pub/Linux/apps/crypto>).

There is a project maintaining 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 in compliance with OpenPGP (<http://www.faqs.org/rfcs/rfc2440.html>). See the GNU Privacy Guard web page (<http://www.gnupg.org>) for more information.

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

11.6.2. SSL, S-HTTP 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 encryption 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 on-line 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 on the OpenSSL web site (<http://www.openssl.org>). Information on Netscape's other security implementations, and a good starting point for these protocols is available at Netscape Security Center (<http://wp.netscape.com/security/index.html>). It's also worth noting that the SSL protocol can be used to pass many other common protocols, “wrapping” them for security. See the sslwrap web page (<http://www.quiltaholic.com/rickk/sslwrap/>).
- **S-HTTP:** - S-HTTP is another protocol which 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 rfc2311 (<http://www.ietf.org/rfc/rfc2311.txt>).

11.6.3. 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 the ipsec Charter (<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 was being developed at the University of Arizona, uses an object-based framework for implementing network protocols called x-kernel. Most simply, the x-kernel is a method of passing messages at the kernel level, which makes for an easier implementation. This project is now closed, but contact information can be found on The x-Kernel Project (<http://openresource.com/openres/orgs/DP/P/x-Kernel.shtml>) web site.

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 computer 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 computers at several different sites connected by the insecure Internet.”

It's available for download from the Linux FreeS/WAN web site (<http://www.freeswan.org/>).

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

11.6.4. ssh (Secure SHell) and stelnet

`ssh` and `stelnet` are suites of programs that allow you to login to remote systems and to 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 to 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 on the Datafellows web site (<http://www.datafellows.com>).

The excellent OpenSSH implementation is based on a early version of the DataFellows `ssh` and has been totally reworked so as not to include any patented or proprietary parts. It is free and released under a BSD license. It can be found on the OpenSSH web site (<http://www.openssh.com>).

There is also a open source project to re-implement `ssh` from the ground up called "lsh". For more information see the LSH (<http://www.lysator.liu.se/~nisse/lsh/>) web site.

You can also use `ssh` from your Windows[®] workstation to your GNU/Linux `ssh` server. There are several freely available Windows[®] clients, including PuTTY (<http://www.chiark.greenend.org.uk/~sgtatham/putty/>) and a commercial implementation from DataFellows on the Datafellows web site (<http://www.datafellows.com>).

SSLeay (outdated, see OpenSSL below) 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 which does encryption over a `telnet` connection. Unlike SSH, `stelnet` uses SSL, the Secure Sockets Layer protocol developed by Netscape. You can find Secure telnet and Secure FTP by starting with the SSLeay and SSLapps FAQ (<http://www.psy.uq.oz.au/~ftp/Crypto/>).



The OpenSSL Project is based on SSLeay and is intended 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 home page (<http://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, visit the Stanford University web site (<http://srp.stanford.edu/>).

11.6.5. PAM - Pluggable Authentication Modules

Your version of the Mandriva Linux 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 (<http://www.kernel.org/pub/linux/libs/pam/index.html>) for more information.

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 cannot 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
```

11.6.6. 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 of the 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, and easy to debug...) implementation.”

At the CIPE project (<http://sites.inka.de/sites/bigred/devel/cipe.html>) web site more information can be found.

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

11.6.7. 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 their 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.faqs.org/faqs/kerberos-faq/general/>), and the code can be found on the Kerberos web site (<http://web.mit.edu/kerberos/www/>).

[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.

11.6.8. Shadow Passwords

Shadow passwords are a means of keeping your encrypted password information secret from normal users. Your Mandriva Linux system uses shadow passwords by default, but on other systems, encrypted passwords are stored in the `/etc/passwd` file for all to read. Anyone can then run password-guesser programs on them and attempt to determine what they are. Shadow passwords, by contrast, are saved in the `/etc/shadow` file, which only privileged users can read. You can refer to the Shadow-Password HOWTO (<http://www.tldp.org/HOWTO/Shadow-Password-HOWTO.html>) for further information if necessary. It is rather dated now, and will not be required for distributions supporting PAM, like your Mandriva Linux system.

11.6.9. "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 (See OpenWall (<http://www.openwall.com/john/>)). 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[®] computers on your network, you should check out L0phtCrack, a Crack implementation for Windows[®]. Check out the @stake LC 4 web site (<http://www.atstake.com/research/lc3/>).

11.6.10. 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 computer. More information and source code is available on the AT&T web site (<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 when the file system is encrypted. More information is available on the TCFS web site (<http://www.tcfs.it/>).

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

11.6.11. X11, SVGA And Display Security

11.6.11.1. X11

It's important for you to secure your graphical display to prevent attackers from grabbing your passwords as you type them, reading documents or information you are reading on your screen, or even using a hole to gain `root` access. Running remote X applications over a network can also 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 computer, they can `xhost + their computer` and get in easily. Also, if you have to allow access from an untrusted computer, 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 computer 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 The Linux Documentation Project (<http://www.tldp.org/HOWTO/Remote-X-Apps.html>) web site.

You can also use `ssh` (see *ssh (Secure SHell) and stelnet*, page 101) 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.

You can also disable any remote connections to your X server by using the `-nolisten tcp` options to your X server. This will prevent any network connections to your server over tcp sockets.

Take a look at Xsecurity(7x) for more information on X security. The safest 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.

11.6.11.2. SVGA

`SVGALib` programs are typically `suid-root` in order to access all your GNU/Linux computer’s video hardware. This makes them very dangerous. If they crash, you typically need to reboot your computer 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.

11.6.11.3. 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. More information available at the GGI Project (<http://www.ggi-project.org/>) web site.

11.7. Kernel Security

This is a description of the kernel configuration options which 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 kernel dot org (<ftp://ftp.kernel.org>) or from packages updates available through `MandrivaUpdate`.

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

11.7.1. Kernel Compile Options

When this document was written, kernel 2.2 was state-of-the-art. Still today, most firewalls still run 2.2. However, with kernel 2.4, a lot of things have changed. Most of the compile options in this chapter are still valid, but the Masquerading and port forwarding have been replaced by `iptables`. You can find more information on `iptables` on the Linux Guruz web site (<http://www.linuxguruz.org/iptables/howto/iptables-HOWTO.html>).

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. Please consult the chapter Compiling and Installing New Kernels of the *Reference Manual* for a full description of the compilation of a brand new kernel.

- Network Firewalls (CONFIG_FIREWALL)

This option should be on if you intend to run any firewalling or masquerading on your GNU/Linux computer. If it's just going to be a regular client computer, 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 computer 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 computers 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 which consumes all the resources on your computer, 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 merely allows syn cookies, but does not enable them. To enable them, you have to do:

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

- IP: Firewalling (CONFIG_IP_FIREWALL)

This option is necessary if you are going to configure your computer 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 receives, such as 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 which may be potential exploits.

- 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 the Indyramp web site (<http://www.indyramp.com/masq>) and "Configuring Masqueraded Clients", page 23 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 Masquerade HOWTO (<http://www.tldp.org/HOWTO/IP-Masquerade-HOWTO/index.html>) 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 which isn't supposed to be there. It could also prevent a teardrop attack against an internal host which is not 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 which 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 user-space 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 computer 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>).

- 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.

11.7.2. Kernel Devices

There are a few block and character devices available on GNU/Linux which can 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` can run out of random bytes and 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.

11.8. 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 Mandriva Linux distribution, either in the main CD-ROM, `contribs`, or through the FTP crypto server (see above).

11.8.1. Packet Sniffers

One of the most common ways intruders gain access to 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 computer (portable or not) 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.

11.8.2. 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 computer 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` or `urpme 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 which 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
```



You may also use a command line utility to do that: `chkconfig` or the graphical interface under KDE: `ksysv`.

Your Mandriva Linux distribution 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 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 which need to have access to your computer’s services.

If you are a home dial up 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 computer, 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 computer, yet still allow you from the inside to connect to servers on the Internet.

Bear 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 computer. You can use `netstat -ta` to find a list of all the services your computer is offering.

11.8.3. Verify Your DNS Information

Keeping up-to-date DNS information about all of the 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 reject connections from hosts that do not have valid DNS entries.

11.8.4. identd

`identd` is a small program that typically runs out of your `inetd` server. It keeps track of which user is running which TCP services, and can then report this to whoever requests this information.

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 not compromised, 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 (Which we recommend), you can even have `identd` return a UID instead of a user name or even `NO-USER`.

11.8.5. Configuring and Securing the Postfix MTA

The Postfix mail server was written by Wietse Venema, author of Postfix and several other staple Internet security products, as an “attempt to provide an alternative to the widely-used Sendmail program. Postfix attempts to be fast, easy to administer, and hopefully secure, while at the same time being Sendmail-compatible enough to not upset your users.”

Further information on postfix can be found at the Postfix web site (<http://www.postfix.org>) and in *Configuring and Securing Postfix* (http://www.linuxsecurity.com/feature_stories/feature_story-91.html).

11.8.6. SATAN, ISS, and Other Network Scanners

There are a number of software packages out there that do port and service-based scanning of computers or networks. SATAN, ISS, SAINT, and Nessus are some of the more well-known ones. This software connects to the target computer (or all the target computers on a network) on as many of the ports as they can, and try to determine what services are running. Based on this information, you can tell if the computer is vulnerable to specific exploits 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 computer or a network of computers. It’s a good idea to get SATAN and scan your computer or network, and fix any 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 on the net (see the Trouble web site (<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 may do a better job.

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

SAINT™ is an 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 the SAINT (<http://www.wdsi.com/saint>) home page.

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 the Nessus web site (<http://www.nessus.org/>).

11.8.6.1. 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. A packet with the TCP ACK bit set (as is done with established connections) will likely get through a packet-filtering firewall. The returned RST packet from a port which **has no established session** can be taken as proof of life on that port. I don’t think TCP wrappers will detect this.

You might also look at SNORT™ (<http://www.snort.org>), which is a free IDS (Intrusion Detection System), which can detect other network intrusions.

11.8.7. Sendmail, qmail and MTA's¹

One of the most important services you can provide is a mail server. 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 with 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/>).

Bear 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 which 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 to `qmail`. `qmail` was designed with security in mind from the ground up. It's fast, stable, and secure. `Qmail` can be found on the `qmail` web site (<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` on the `Postfix` web site (<http://www.postfix.org>).



`Postfix` is the default MTA shipped with Mandriva Linux. Please refer to "*Postfix Mail Server*", page 53.

11.8.8. Denial of Service (DoS) 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 a computer or network.

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.

- **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 computer or services. See *Kernel Security*, page 105 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 computer. If they are doing this from a host with better bandwidth than yours, your computer will be unable to send anything on the network. A variation on this attack, called "smurfing", sends ICMP packets to a host with **your** computer's return IP, allowing them to flood you less detectably. You can find more information about the "smurf" attack on the Linux Security web site (http://www.linuxsecurity.com/articles/network_security_article-4258.html).

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 coming 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 for 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.

1. Mail Transport Agents

You can find code for most exploits, and a more in-depth descriptions of how they work, on the Insecure web site (<http://www.insecure.org/sploits.html>) using their search engine.

11.8.9. NFS (Network File System) Security

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

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

There is some small amount of security allowed in exporting file systems. 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 account and have total access to their files. This is only a small hindrance to an attacker that has access to mount your remote file systems.

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

See the NFS HOWTO (<http://www.tldp.org/HOWTO/NFS-HOWTO/>) for more information on NFS.

11.8.10. NIS (Network Information Service)

Network Information Service (formerly YP) is a means of distributing information to a group of computers. 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 computers 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 computers 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.

There is a much more secure replacement for NIS, called NIS+. Check out the NIS HOWTO (<http://www.tldp.org/HOWTO/NIS-HOWTO/>) for more information.

11.8.11. 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 computers 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, and `iptables` for 2.4 kernels, allow 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 computers behind it. This is a fatal mistake. Check out the very good Firewall HOWTO (<http://www.ibiblio.org/mdw/HOWTO/Firewall-HOWTO.html>) for more information on firewalls and GNU/Linux.

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 on-line firewall documents are mandatory reading. Check out the 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 on the Computer Security Resource Center (CSRC) web site (<http://cs-www.ncsl.nist.gov/publications/nistpubs/800-10/main.html>). Also of interest are:

- The Freefire Project — a list of freely-available firewall tools, available at freefire (http://sites.inka.de/sites/lina/freefire-1/index_en.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.stearns.org/mason/>).

11.8.12. 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 (<http://www.tldp.org/HOWTO/IPCHAINS-HOWTO.html>) for further information.

11.8.13. Netfilter – Linux Kernel 2.4.x Firewalling

In yet another set of advancements to the kernel IP packet filtering code, netfilter allows users to set up, maintain, and inspect the packet filtering rules in the new 2.4 kernel.

The netfilter subsystem is a complete rewrite of previous packet filtering implementations including ipchains and ipfwadm. Netfilter provides a large number of improvements, and it has now become an even more mature and robust solution for protecting corporate networks.

`iptables` is the command-line interface used to manipulate the firewall tables within the kernel.

Netfilter provides a raw framework for manipulating packets as they traverse the various parts of the kernel. Part of this framework includes support for masquerading, standard packet filtering, and now more complete network address translation. It even includes improved support for load balancing requests for a particular service between a group of servers behind the firewall.

The stateful inspection features are especially powerful. Stateful inspection provides the ability to track and control the flow of communication passing through the filter. The ability to keep track of state and context information about a session not only makes rules simpler but also helps to better interpret higher-level protocols.

Additionally, small modules can be developed to perform additional specific functions, such as passing packets to programs in userspace for processing then reinjecting back into the normal packet flow. The ability to develop these programs in userspace reduces the level of complexity that was previously associated with having to make changes directly at the kernel level.

Other IP Tables references include:

- Oskar Andreasson IP Tables Tutorial (http://www.linuxsecurity.com/feature_stories/feature_story-94.html) — Oskar Andreasson speaks with LinuxSecurity.com about his comprehensive IP Tables tutorial and how this document can be used to build a robust firewall for your organization.
- Hal Burgiss Introduces Linux Security Quick-Start Guides (http://www.linuxsecurity.com/feature_stories/feature_story-93.html) — Hal Burgiss has written two authoritative guides on securing Linux, including managing firewalling.
- Netfilter Home page (<http://www.netfilter.org>) — The netfilter/iptables home page.

- Linux Kernel 2.4 Firewalling Matures: netfilter (http://www.linuxsecurity.com/feature_stories/kernel-netfilter.html) — This LinuxSecurity.com article describes the basics of packet filtering, how to get started using iptables, and a list of the new features available in the latest generation of firewalling for Linux.

11.8.14. VPNs: Virtual Private Networks

VPNs are a way to establish a “virtual” network on top of some already existing network. This virtual network is often 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.

There are several GNU/Linux VPN solutions available:

- `vpnd`. See the VPN Daemon (<http://sunsite.dk/vpnd/>) web site.
- Free S/Wan, available on the Linux FreeS/WAN web site (<http://www.freeswan.org/>).
- `ssh` can be used to construct a VPN. See the VPN PPP-SSH mini-HOWTO (<http://www.tldp.org/HOWTO/ppp-ssh/index.html>) for more information.
- `vps` (virtual private server) at strongcrypto (<http://www.strongcrypto.com>).
- `vtun` (virtual tunnel) at sourceforge (<http://vtun.sourceforge.net/>).
- `yavipin` (<http://yavipin.sourceforge.net>).

Also see the section on IPSEC for more pointers and information.

11.9. Security Preparation (Before You Go On-Line)

Ok, so you have checked over your system, and determined that it is as secure as feasible, and you’re ready to put it on-line. 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.

11.9.1. Make a Full Backup of Your Computer

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), you will of course need at least 650MB of space to make the image. Tapes and other re-writable 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.

11.9.2. 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.

11.9.3. Testing your Backups

You should do periodic tests of your backups to make sure they are working as you might expect them to. Restores of files and checking against the real data, sizes and listings of backups, and reading old backups should be done on a regular basis.

11.9.4. 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 that it has not 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` probably will not 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 which can be included to make it less verbose. Bear in mind that 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 re-archived. You will have to decide the advantages versus drawbacks.

11.9.5. 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 may also want to configure your log-rotating script to keep logs around longer so you have time to examine them. Take a look at `logrotate(8)`.

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 which 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 the SDSC Syslog Home Page (<http://security.sdsc.edu/software/sdsc-syslog/>) 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.com/products/syslog_ng/). It allows you a lot more flexibility in your logging and can also encrypt 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.

11.9.6. Apply All New System Updates

Most users install their systems from CD-ROM. Due to the fast-paced nature of security fixes, new (fixed) programs are always being released. Before you connect your computer to the network, it's a good idea to run `MandrivaUpdate` (on another computer connected to the Internet) and get all the updated packages since you received your distribution CD-ROM. Many times these packages contain important security fixes, so it's a good idea to get them installed.

11.10. 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.

11.10.1. 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 computer. Depending on your own level of 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 computers), the next best step is to use something like `tcp_wrappers` or `ipfwadm` to deny access from the intruder's site.

If you can't deny all people from the same site as the intruder, locking that 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 back doors.

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.

11.10.2. 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?

11.10.2.1. 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 to 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 Mandriva Linux 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 computer, but back somewhere on your network. If they were running a packet sniffer, odds are good they have access to other local computers.

11.10.2.2. 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 a file integrity check; this should help to show 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), reinstalling, 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 should 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.

11.10.2.3. 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.

11.10.2.4. 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. If you spotted anything else distinctive about your intruder, you might mention that too. After sending the email, you should (if you are so inclined) 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).

11.11. Security Sources

There are a **lot** of good sites out there for UNIX[®] security in general and GNU/Linux security specifically. It is 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.

11.11.1. LinuxSecurity.com References

The Linux Security (<http://www.linuxsecurity.com>) web site has numerous Linux and open source security references written by the Linux Security staff and people collectively around the world.

- Linux Advisory Watch (<http://www.linuxsecurity.com/vuln-newsletter.html>) — A comprehensive newsletter that outlines the security vulnerabilities that have been announced throughout the week. It includes pointers to updated packages and descriptions of each vulnerability.
- Linux Security Week (<http://www.linuxsecurity.com/newsletter.html>) — The purpose of this document is to provide readers with a quick summary of each week's most relevant Linux security headlines.
- Linux Security Discussion List (<http://www.linuxsecurity.com/general/maillinglists.html>) — This mailing list is for general security-related questions and comments.
- Linux Security Newsletters (<http://www.linuxsecurity.com/general/maillinglists.html>) — Subscription information for all newsletters.
- comp.os.linux.security FAQ (<http://www.linuxsecurity.com/docs/colsfaq.html>) — Frequently Asked Questions with answers for the comp.os.linux.security newsgroup.
- Linux Security Documentation (<http://www.linuxsecurity.com/docs/>) — A great starting point for information pertaining to Linux and Open Source security.

11.11.2. FTP Sites

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

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

Matt Blaze is the author of CFS and a great security advocate. Matt's archive is available on the AT&T web site (<ftp://ftp.research.att.com/pub/mab>).

11.11.3. Web Sites

- The Hacker FAQ is a FAQ about hackers available on the Plethora web site (<http://www.plethora.net/~seebs/faqs/hacker.html>).
- The COAST archive has a large number of UNIX[®] security programs and information. It is available at CERIAS (<http://www.cerias.purdue.edu>).
- SuSE Security Page available at SuSE (<http://www.suse.de/de/security/>).
- BUGTRAQ puts out advisories on security issues, available on Security Focus (<http://www.securityfocus.com/archive/1>).
- CERT, the Computer Emergency Response Team, puts out advisories on common attacks on UNIX[®] platforms available at CERT (<http://www.cert.org/>).
- 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 available on the Trouble web site (<http://www.trouble.org/security>).
- CIAC sends out periodic security bulletins on common exploits. See CIAC (<http://ciac.llnl.gov/cgi-bin/index/bulletins>) for more information.
- A good starting point for GNU/Linux Pluggable Authentication modules can be found on The Linux Kernel Archives (<http://www.kernel.org/pub/linux/libs/pam/>).
- WWW Security FAQ, written by Lincoln Stein, is a great web security reference. Find it on the W3C web site (<http://www.w3.org/Security/Faq/www-security-faq.html>).
- Mandriva Security Advisories (<http://www.mandriva.com/security>) is Mandriva Linux official security page, notably holding new advisories, End of life product support policy, etc.

11.11.4. Mailing Lists

Mandriva Linux security lists. You can be informed for each security fix by subscribing to our security mailing-lists (<http://www.mandriva.com/community/resources/newsgroups>).

`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`”

11.11.5. 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 Icové, 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.

11.12. 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 which 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 computers, 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 servers. 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 94. This is done intentionally to prevent remote users from attempting to connect via `telnet` to your computer as `root`, which is a serious security vulnerability, because then the root password would be transmitted, in clear text, across the network. Don't forget: potential intruders have time on their side, and can run automated programs to find your password. Additionally, this is done to keep a clear record of who logged in, not just root.

Q: How can I enable the Apache SSL extensions?

A: Simply install the package `mod_ssl`, and consult the documentation at `mod_ssl` home page (www.modssl.org).



You should also consider the `mod_sxnet` module, which is a plug-in for `mod_ssl` and allows the activation of the "Thawte Secure Extranet". `mod_ssl` encrypt communications, but `mod_ssl-sxnet` goes further and allows you to securely authenticate the user of the web page thanks to a personal certificate. You have more info on this application on Thawte (<http://www.thawte.com/certs/strongextranet/>) or install the `mod_sxnet` module from your Mandriva Linux distribution and read the included package documentation.

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 Mandriva Linux distribution contains a large number of tools to change the properties of user accounts.

- The `pwconv` and `unpwconv` programs can be used to convert between shadowed 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 will not.

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

A: I bet you did not know about Apache Week (<http://www.apacheweek.com>), did you?

You can find information on user authentication on Apache Week (<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).

11.13. Conclusion

By subscribing to the security alert mailing lists, and keeping current, you can do a lot towards securing your computer. 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 computer. More effort is required on business computers, 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.

Security-Related Terms

Included below are several of the most frequently used terms in computer security. A comprehensive dictionary of computer security terms is available in the Linux Security Dictionary (<http://www.linuxsecurity.com/dictionary/>)

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 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 which consumes the resources of 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 which has at least two network interfaces.

firewall

A component or set of components which 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 which is made up of several components. It is a security exploit that works by tricking computers in a trust relationship into thinking that you really are someone you are not. 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 which 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 which 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`.

Chapter 12. Networking Overview

12.1. Copyright

This chapter is based on a HOWTO by Joshua D. Drake {POET}, the original of which is hosted by the The Linux Review (<http://www.thelinuxreview.com/>) web site.

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} , CommandPrompt, Inc. is a FREE document. You may redistribute it under the terms of the GNU General Public License.

Modifications from v1.7.0, 2000, (C)opyright 2000 - 2005 Mandriva.

12.2. How to Use this Chapter

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

The suggested path through the document follows:

Read the generic sections

These sections apply to almost every technology described, and they are very important for you to understand. Most readers should be comfortable with this material.

Consider your network

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

Read *Ethernet Information*, page 128 if you are directly connected to a LAN or the Internet:

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

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

This section describes PLIP, PPP, SLIP and ISDN, the most 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, *Other Network Technologies*, page 131 provides some pointers to information on non-IP protocols and other communications hardware.

Do the configuration work

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

Look for further help if needed

If you experience problems that this chapter 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.

12.2.1. Conventions Used in this chapter

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

12.3. General Information about Linux Networking

12.3.1. Linux Networking Resources

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

There is a wealth of consultants available. A searchable listing can be found on the linuxports (<http://www.linuxports.com/>) web site. There is also a newsgroup in the Linux news hierarchy dedicated to networking and related matters, it is: `comp.os.linux.networking` (`news:comp.os.linux.networking`).

When reporting problems, please remember 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 you issued.

12.3.2. Where to Get some non Linux-Specific Network Information

If you are after more basic tutorial information on TCP/IP networking, then the following documents are recommended:

TCP/IP Introduction

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

TCP/IP Administration

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

If you are looking for more detailed information on TCP/IP networking, then *Inter networking with TCP/IP, Volume 1: principles, protocols and architecture*, by Douglas E. Comer, ISBN 0-13-227836-7, Prentice Hall publications, 3rd Edition or newer, is highly recommended.

If you want to learn about how to write network applications in a UNIX-compatible environment, then *Unix Network Programming*, by W. Richard Stevens, ISBN 0-13-949876-1, Prentice Hall publications, 2nd Edition or newer, is highly recommended. Check the Prentice-Hall Professional Technical Reference (<http://www.phptr.com/>) web site for more information.

You might also try the TCP-IP Protocol Newsgroup (`news:comp.protocols.tcp-ip`).

An important source of specific technical information relating to the Internet and the TCP/IP suite of protocols are RFCs. 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 and some allow you to search the RFC database for particular keywords. One such repository with a search engine is available at NEXOR dot com (http://www.nexor.com/rfc_search.htm).

12.4. Generic Network Configuration Information

The following sections present concepts that need to be known and understood before actually trying to configure a network. These fundamental principles apply regardless of the exact nature of the network you wish to deploy.

12.4.1. What Do I Need to Start?

Before you start building or configuring your network, the most important things needed are:

12.4.1.1. Kernel Support for the Network Hardware and Protocols

Your Mandriva Linux distribution comes with networking enabled, and with a kernel which supports most well-known networking hardware (for example: 3COM, RealTek, NE2000 or Intel NICs) and protocols. Please refer to your hardware documentation for more information.

12.4.1.2. An Explanation of IP Addresses

Internet Protocol (IP) addresses are composed of four bytes¹. The convention is to write addresses in what is called “dotted decimal notation”. In this form, each byte is represented in decimal notation (0-255), dropping any leading zeroes, unless the number is zero, and written with each byte separated by a “.” character. By convention, each host or router interface has an IP address. In some circumstances, the same IP address could be used on each interface of a single machine, 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 which is common to all addresses within the network is called the “network portion”. The remaining digits are called the “host portion”. The number of bits that are shared by all addresses within a network form the “netmask” whose role is 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 with zeroes.

The broadcast address is a special one which every host on the network listens to (in addition to its own unique address). Datagrams (like certain type of routing information and warning messages) that are to be received by all hosts on the network are sent to this address. The standard is to use one of the addresses at either extreme of the network address range as the broadcast address, preferring the one at the high end of the range. In the example above, this would be 192.168.110.255. In practice, it doesn’t matter very much which one you use, just 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 can be allocated. The ranges allocated are:

Network Class	Netmask	Network Addresses
A	255.0.0.0	0.0.0.0 – 127.255.255.255

1. for version 4 of IP, a.k.a IPv4

Network Class	Netmask	Network Addresses
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

The addresses you should use depend 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 that 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 just make them up and then expect your configuration to work.

Building a brand new network which 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 12-1. Reserved Private Network Allocations

You should first decide how large you want your network to be, and then choose as many addresses as you require.

12.4.2. Routing

Routing is a big topic. Whole books are written about it. Most of you will have fairly simple routing requirements, some of you will not. Only the basic fundamentals of routing are covered.

What is IP routing? Here is a possible definition: “IP routing is the process by which a host with multiple network connections decides where to deliver the 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 to 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 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, 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. You can see this table by using the following command:

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

or by using either one 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.

To manipulate this table, the `route` command is used. Please refer to `route(8)` for more information.

12.4.2.1. What Does the Routed Program Do?

The routing configuration described above is best suited for simple network arrangements where there is only one possible path to a determined destination. When you have a more complex network arrangement, things get a little more complicated.

The big problem with “manual routing”, or “static routing” as described, is that if a machine or link fails on 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 error-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). The Routing Information Protocol is very common on small networks such as small-to-medium size corporate networks or building networks. OSPF is more modern and more capable of handling large network configurations and better suited to environments where there are 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 supplied with your Linux distribution.

An example of where and how you might use a dynamic routing protocol may look something like figure 12-1.

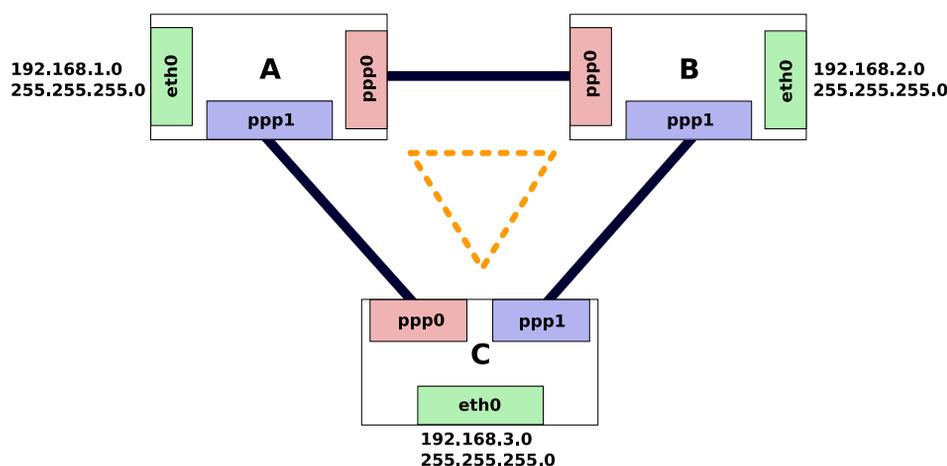


Figure 12-1. A Dynamic Routing Example

There are three routers: A, B and C, supporting one Ethernet segment with a Class C IP network (netmask 255.255.255.0). All routers also have a PPP link to each other. 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 fails. If that link fails, then with the routing entry shown above, hosts on the A Ethernet segment could not reach hosts on the B segment 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 C segment and hosts on the C segment could still talk to hosts on the B 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 on 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. For each router, you only need to do the following (for example, 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 need more information, then you should refer to the suggested references listed in *General Information about Linux Networking*, page 124.

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 capability 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 for small-to-medium size networks.

12.5. Ethernet Information

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

12.5.1. Supported Ethernet Cards

GNU/Linux supports almost every known network card. It would be pointless to list them all here. If you have trouble setting your network card up, please refer to its accompanying documentation (if any) or to its manufacturer's web site. You can also check your kernel's specific documentation for some NICs.

12.5.2. General Ethernet Information

Ethernet device 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 in which they are detected.

Card configuration is easy. Typically, you would use something like (which most distributions already do for you, if you configured them to support your Ethernet) this:

```
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
```

12.5.3. Using Multiple Ethernet Cards in The Same Machine

If you have a computer with three NE2000 cards, one at 0x300, one at 0x240, and one at 0x220 then you would add the following lines to the `/etc/modprobe.conf` file:

```
alias eth0 ne
alias eth1 ne
alias eth2 ne
options ne io=0x220,0x240,0x300
```

What this does is tell the **modprobe** program to look for three NE-based cards at the listed I/O addresses. It also states in which order they should be found and the device to which 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 so that you cannot load two copies of the same module.

The `irq=` option is used to specify the hardware IRQ, and the `io=` to specify the different I/O ports.

Please refer to the Ethernet-HOWTO (<http://www.tldp.org/HOWTO/Ethernet-HOWTO.html>) for more information about Ethernet.

12.6. IP-Related Information

12.6.1. DNS

DNS stands for Domain Name System. It is responsible of mapping a machine name such as `www.mandriva.com` with the IP address of that machine, in this case: `212.85.150.181` at the time of writing. Mapping is available in both directions, that is from name to IP and vice versa (called reverse DNS lookup).

The DNS is composed of a great number of machines 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 does not have the answer, then it asks another one and so on. You can also have a local DNS responsible for mapping addresses on your LAN.

We can differentiate two major DNS classes: caching DNS and master DNS servers. The first only “remembers” a previous request and then it can answer without asking a master DNS server again. The latter servers are really responsible as a last resort to map an address with a name — or possibly to specify that a given name does not map to any known address.

12.6.2. DHCP

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 configure them all automatically, assigning all common parameters through 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 located on the same segment to configure the host.

12.6.3. IP Aliasing

There are some applications where being able to configure multiple IP addresses to a single network device is useful. Internet Service Providers often use this facility to provide a “customized” feature to their World Wide Web and FTP offerings for their customers. You can refer to the IP-Alias mini-HOWTO (<http://www.tldp.org/HOWTO/IP-Alias/>) for more information.

12.6.4. IP Firewall

IP Firewall and Firewalling issues are covered in more depth in the Firewall-HOWTO (<http://tldp.org/HOWTO/Firewall-HOWTO.html>). IP Firewalling allows you to secure your machine or network against unauthorized network access by filtering or allowing datagrams to or from IP addresses which you nominate. There are three different classes of rules; incoming filtering, outgoing filtering, and forwarding filtering. Incoming rules are applied to datagrams that are received by a network device. Outgoing rules are applied to datagrams that are to be transmitted from a network device. Forwarding rules are applied to datagrams that are received but are not intended for this machine (ie. datagrams that would be routed).

12.6.5. IP Masquerade

Many people have a single link to connect to the Internet being assigned a single IP address by the Internet Service Provider. This is enough to allow only one host full access to the Internet. IP Masquerade is a technique enabling many machines make use of that one IP address, causing the other hosts to “look like” (be masqueraded by) the machine having the Internet connection. There is one caveat: the masquerade function usually works in only one direction. That is, the masqueraded hosts can make calls out, but they cannot accept or receive network connections from remote hosts. This means that some network services do not work (such as talk), and others (such as FTP) must be configured in passive (PASV) mode to operate. Fortunately, the most common network services work just fine.

12.6.6. IPv6

Just when you thought you were beginning to understand IP networking, the rules get changed! IPv6 is the shorthand notation for version 6 of the Internet Protocol. IPv6 was developed primarily to overcome the concern that address space wouldn't be sufficient, making the address 16 bytes long (128 bits). IPv6 incorporates a number of other changes, mostly simplifications, that will make IPv6 networks more manageable than IPv4 networks.

Please refer to the Linux IPv6 HOWTO (<http://www.tldp.org/HOWTO/Linux+IPv6-HOWTO/>) for more information.

12.6.7. Traffic Shaping

Traffic shaping is used to limit the traffic to given maximum values. This technique is useful to “partition” the available bandwidth among network interfaces. The traffic shaper creates new interface devices which rely on physical network devices for actual transmission, and can be used as outgoing routes for network traffic.

Please refer to the Traffic Control HOWTO (<http://www.tldp.org/HOWTO/Traffic-Control-HOWTO/>) for more information on traffic control in general, and traffic shaping in particular.

12.7. Using Commodity PC Hardware

12.7.1. ISDN

The Integrated Services Digital Network (ISDN) is a series of standards which 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 establish calls and other functions. In Australia for example, ISDN may be delivered on a 2Mbps link that is broken into 30 discrete 64Kbps data channels with one 64Kbps control 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 Digital Network Terminal that your telecommunications carrier will have installed when you got your ISDN service and presents a number of serial interfaces. One of these interfaces is used to enter commands to establish calls and configuration while the others are actually connected to the network devices that will use the data circuits when they are established. This sort of configuration doesn’t need special setup: 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 was designed, allows you to install an ISDN card into your machine and then have your Linux software handle the protocols and make the calls itself.

Both passive and active internal ISDN cards are supported. We will not list them all here.

Some of these cards may require software to be downloaded in order to make them operational. There are separate utilities to do this.

Please refer to the included ISDN documentation and to the FAQ available on the ISDN4Linux (<http://www.isdn4linux.de/faq/>) web site, for more information.



About PPP. The PPP suite of protocols will operate over both asynchronous or synchronous serial lines. The commonly distributed PPP daemon for Linux, `pppd`, supports only asynchronous mode. If you wish to run the PPP protocol over your ISDN service, you need a specially modified version. Details of where to find it are available in the documentation referred to above.

12.7.2. PLIP

The new code for PLIP behaves just 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, similar 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 the IRQ number used by your port wasn’t detected, PLIP will fail to work. In this case, just write the right number to `/proc/parport/0/irq` and load PLIP again.

You can refer to the PLIP mini-HOWTO (<http://www.tldp.org/HOWTO/PLIP.html>) for more information.

12.7.3. PPP

Due to the nature of PPP, its size, complexity, and flexibility, it has its own HOWTO. The Linux PPP HOWTO (<http://tldp.org/HOWTO/PPP-HOWTO/>) is available, but its official home is now on The Linux Review (<http://www.thelinuxreview.com/howto/ppp/>) web site.

12.8. Other Network Technologies

GNU/Linux networking is not limited to Ethernet and IP. Almost all existing networking hardware and protocol is supported. The following subsections will list some of the “more common” ones, introduced in no particular order.

12.8.1. Appletalk

Appletalk support allows your GNU/Linux machine to connect with Apple networks to share resources such as printers and disks. Install the `netatalk` package and refer to the Networking Applet Macintosh through Open Source (<http://netatalk.sourceforge.net/>) web site, for more information.

12.8.2. IPX

The IPX (Internetwork Packet Exchange) protocol is most commonly utilized in Novell NetWare™ local area network environments. Linux includes support for this protocol and may be configured to act as a network end point, or as a router for IPX.

The IPX protocol and the NCPFS (Netware Core Protocol File System) are covered in greater depth in the Linux IPX-HOWTO (<http://www.tldp.org/HOWTO/IPX-HOWTO.html>).

12.8.3. NetBEUI, NetBios, CIFS

Samba is an implementation of the Session Management Block protocol (SMB). It allows Windows® and other systems to mount and use your disks and printers, as well as access disks and printers served by Windows® servers.

Please refer to the SMB-HOWTO (<http://www.tldp.org/HOWTO/SMB-HOWTO.html>) and the Samba - opening windows to a wider world (<http://www.samba.org>) web sites for more information.

12.8.4. Token Ring

Token Ring is an IBM standard LAN protocol which avoids data collisions by providing a mechanism which 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 on to the next station. The token loops amongst all active stations, hence the name “Token Ring”.

The configuration of Token Ring is identical to that of Ethernet, with the exception of the network device name to configure. Token ring device names are `tr0`, `tr1` etc.

12.9. 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. The PLIP cable is done using D-25 male connectors on both ends, while the serial NULL-modem cable can be done using D-25 or DB-9 female connectors on both ends.

12.9.1. 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 lines. This is OK, but it means you must use software flow control (XON/XOFF) which is less efficient than the preferred hardware (RTS/CTS) flow control. The following cable provides the best possible signalling between machines and allows you to use hardware flow control.

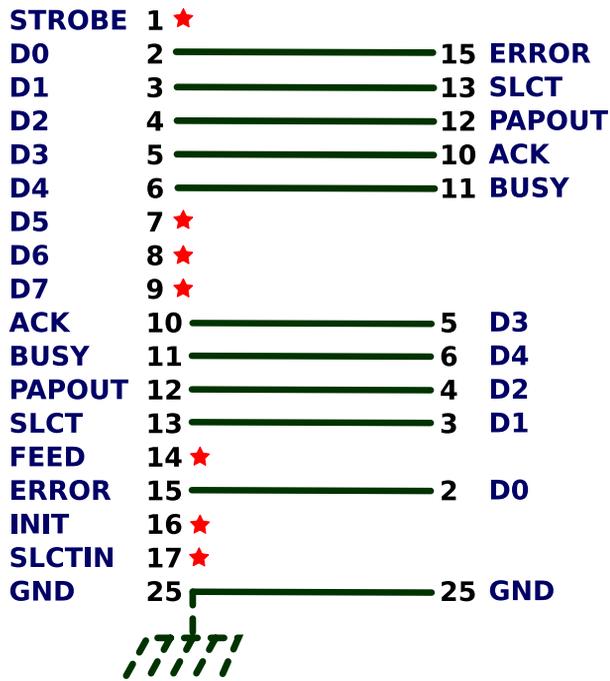


Figure 12-3. PLIP Cable Cabling



- Do not connect the pins marked with a star;
- extra grounds are on pins 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**.



A badly wired PLIP cable can destroy the computer's parallel port. Be very careful and double-check every connection to ensure you don't cause yourself any unnecessary problems.

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.

12.9.3. 10base2 (Thin Coax) Ethernet Cabling

10base2 is an Ethernet cabling standard which 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 should use "T pieces" to interconnect the machines, so that you end up with something that looks like what's shown in figure 12-4.



Figure 12-4. 10base2 Ethernet Cabling

Both ends have a terminator piece, the coaxial cables have BNC plugs at either end, and “middle” computers are connected using 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.

12.9.4. Twisted-Pair Ethernet Cable

If you only have two computers with twisted pair Ethernet cards and you wish to connect them, you do not require a hub. You can cable the two cards directly together using the cabling scheme shown in figure 12-5.

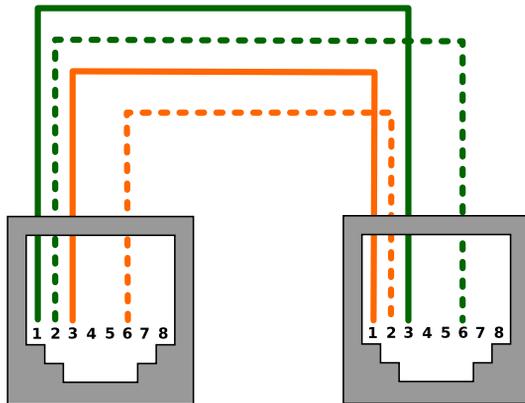


Figure 12-5. Twisted-Pair “NULL-modem” Cabling

Chapter 13. Troubleshooting

This chapter guides you through some troubleshooting basics, that is: what to do when everything goes wrong or, better yet, what to do to be **prepared** if something goes wrong and then how to fix it.

13.1. Introduction

Making backup copies of your data, fixing little problems, recompiling the kernel, installing software, and tweaking configuration files are not uncommon scenarios in every day GNU/Linux life: even if you don't do it all the time, some day you may want or need to. Those tasks can be managed without any hassle if you use a little common sense and follow some practices and guidelines we discuss in this chapter.



Many of the examples and tools presented in this chapter deal with the command line. Usually, restoration of a damaged system to a working state can only be done this way. We assume that you feel comfortable enough using this powerful tool.

So, on to the basic things you need to have ready...

13.2. A Boot Disk

The very first thing you need when your system cannot boot from the hard disk is a boot disk. It allows you to boot your system up and, in a matter of minutes, enable you to undo what rendered your system unusable.

13.2.1. Using the Mandriva Linux CD Rescue Mode

To access Mandriva Linux's rescue mode (available on the first CD-ROM), boot from the CD-ROM, and press the **F1** key, then type `rescue` and hit **Enter**. The system boots in rescue mode (see figure 13-1).

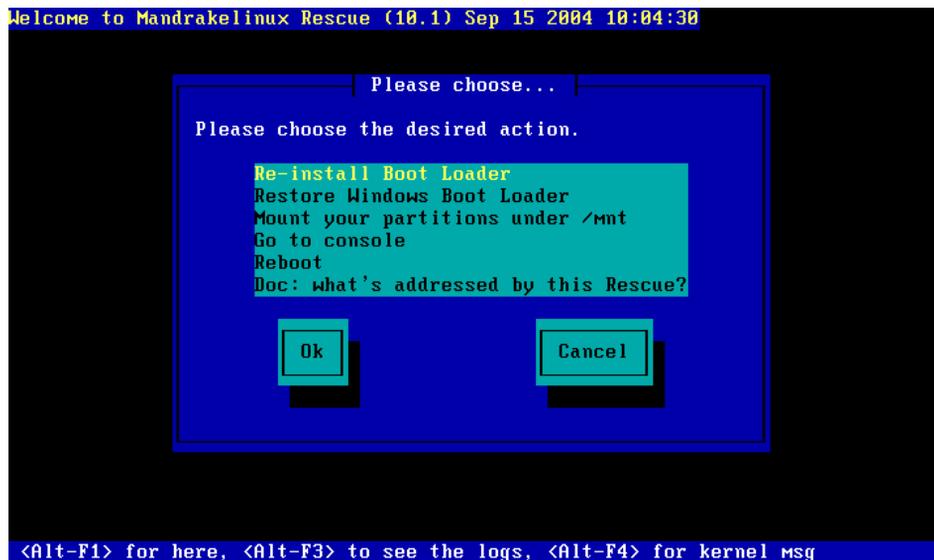


Figure 13-1. Available Rescue Mode Actions

You navigate through the actions with the arrow keys and execute the selected action by pressing **Enter**. The actions available are:

Re-install Bootloader

Use this option to restore the Linux bootloader to the disk's MBR. The former bootloader configuration will be activated again. This is handy, for example, if you dual boot with Windows® and activated a virus which corrupted your disk's MBR leading to an unbootable system.

Restore Windows Bootloader

Use this option to restore the Windows® bootloader to the disk's MBR. This can be used to completely clean the Linux bootloader information and leave Windows® only "as if Linux had never been installed". Press **Enter** to perform the action, or **N** followed by **Enter** to cancel the action.



You won't be able to boot Linux after performing this action. Note however, that this doesn't erase the Linux partitions and system from your harddisk.

Mount your partitions under /mnt

Use this option to mount all available partitions under the /mnt directory. Each partition will be mounted in its own directory, with the same name it would have if mounted by the original system. This option is very useful when you need to access data on your partitions, for example to back it up. You will probably need to mount partitions before accessing the console, for example.

Go to Console

Use this option to access the console where you can perform further operations, for example loading network card drivers, copying files, formatting partitions, etc. A very basic Linux system is available with a few consoles which you can switch between using the **Alt-F<n>** key sequence.



You can return to the rescue mode actions menu by issuing the `rescue-gui` command.

Once you have finished using the console you can issue the `reboot` command to restart the system.

Reboot

Restarts the machine. Take the CD-ROM out if you want the system to boot as usual. You won't be asked for confirmation, the system reboots immediately.

Doc: What's addressed by this Rescue?

Shows a few pages of help text, with brief explanations on what the rescue mode addresses. Navigate through the pages using the **Page Up** and **Page Down** or the arrow keys and press the **Q** key followed by the **Enter** key to return to the rescue actions menu.

13.3. Backup

13.3.1. Why Backup?

Backing up your system is the **only** means of being able to repair it if it suffers severe damage, if you accidentally delete some important system files, or if someone breaks into your system and intentionally deletes some files. You should also back up your personal data (compressed audio, images, office documents, e-mails, address book, etc.) to be safe.

You should make your backups using an appropriate medium and keep them in a safe place. Such a place should be outside the place you usually work in, if possible. You can even have two backups, one on-site, and one outside. Generally speaking, you should make sure that you will be able to restore those backups if you want all this to be really useful.

13.3.2. Preparing your System

You probably have everything you need already installed in your system. You should also keep a boot disk near at hand (you **created** one, didn't you?). Actually, you can make backups using only `tar` and, optionally, a compression tool such as `gzip` or `bzip2`. See an example in *Backup Example Using tar*, page 139.

As an alternative, you can use specialized backup programs, such as Taper, Time Navigator, Arkeia, or Mandriva Linux's own Drakbackup.

13.3.3. What to Backup?

Well, this might be the single most difficult question every system administrator asks himself when the time to backup comes. The answer depends on issues such as: are you only backing up your personal data, your configuration files, or your whole system? How much time or space is it going to take? Will you be restoring your backup on the same machine/OS version, or on a different one?

Since this is a troubleshooting chapter, we try to focus on making a backup which allows you to quickly restore your system to the state it was before that terrible thing which rendered it unusable happened. Of course, you need to make a backup of your personal data if you don't want to lose it.

As a rule of thumb, you should back up the following directories: `/etc`, `/home`, `/root` and `/var`. If you do a complete backup of these directories, you have saved not only your system configuration, but your data as well. Please bear in mind that a backup can take a **long** time to complete, but it's the safest bet.

A more sophisticated scheme would be to backup only those files which have changed, skipping the ones which haven't. This will take more planning time, but will lead to quicker backups (and quicker restores, too). They will also be "easier" to port from one machine/OS version to another.

To summarize, back up all the configuration files of the programs you use and all of the configuration files you have changed. Also back up all your personal (and your system's users) data files. You won't regret it.

13.3.4. Where to Back Up?

The other big question to answer. This depends on how much you want to back up, how fast you want to make your backups, how easy it is to access the backup media, and a large list of etceteras.

Generally speaking, you need media that is at least as large as the amount of information you want to back up, and fast enough so the whole process won't take forever to complete.

Available backup media options vary in capacity, reliability, and speed. You can combine backup media according to your backup strategy, for example: tapes and CD-R/DVD+RW, hard disk and tapes, hard disk and CD-R/DVD+RW, etc., but bear in mind that your backup software may or may not support all of these options.

13.3.5. When to Back Up?

There are many policies for backup schedules. We discuss a few in this chapter but remember that these are not mandatory, nor the best ones, nor the only ones. These are just guidelines you may want to follow in rolling out your own backup schedule.

The many backup strategies out there depend on the media you use, on how often your data changes, and on how critical that data is to you or your organization. For example, one strategy states that you should make a full backup each weekend, and an incremental (changed data only) backup every day. Then make a full backup every month and store that one in at least two places. This strategy might prove useful for a company, but not for a personal computer. For your personal backups you can think of something like this: make a weekly backup of your files on your disk drive and each month transfer those backups to CD-R/DVD+RW or tape.

13.3.6. Backup Example Using tar

Next, we introduce you to a little backup script which uses `tar` and `bzip2` to make a compressed backup of the list of directories you provide. Please read the script's comments for tips on its usage.



You need read permission on the files, and read and execute permissions on the directories you are going to back up. Otherwise the backup operation will fail.

```
#!/bin/bash

# Create a compressed backup of all the directories specified and put the
# resulting file in a directory of your choice.

BACKUP_DIRS="$HOME /etc /var"
BACKUP_FILENAME='date +%b%d%Y'
BACKUP_DEST_DIR="/backups"

# Uncomment the following line for GZipped backups, comment for
# BZipped backups

#tar cvzf $BACKUP_DEST_DIR/$BACKUP_FILENAME.tar.gz $BACKUP_DIRS

# We do a BZipped backup here...
# Comment the following line for GZipped backups, uncomment for
# BZipped backups

tar cvjf $BACKUP_DEST_DIR/$BACKUP_FILENAME.tar.bz2 $BACKUP_DIRS
```

Use `BACKUP_DIRS` to specify the directories you want to include in the backup and `BACKUP_DEST_DIR` to specify the destination directory where the backup is going to be stored. Make the script executable: open a terminal and run `chmod 700 backup.sh`.

Of course, you can always move the resulting `tar.bz2` or `tar.gz` file to any media you want later. You can even backup directly to the media you want by mounting it and changing the variable `BACKUP_DEST_DIR` of the script accordingly. Feel free to enhance this script and make it as flexible as you want.

To restore the backups made this way, please look at *Restore Example Using tar*, page 140.

13.4. Restore

The restoration of a backup depends on which program, media, and schedule you used to make it. We won't cover all the restore cases, but only mention that in order to recover your settings and data files, make sure that you restore the files and/or directories to the same places they were in when you made the backup.

13.4.1. Restore Example Using tar

Now, we introduce a little script to restore the backup we made with `tar` using the script introduced earlier in *Backup Example Using tar*, page 139.



You need write permissions on the files and directories you are going to restore. Otherwise the restore operation will fail.

```
#!/bin/bash

# Extract a compressed backup of all the directories specified
# putting the backed up files into their original places.

BACKUP_SOURCE_DIR="/backups"
RESTORE_FILENAME=$1

# Uncomment the following line if you are restoring GZipped
# backups

#tar xvzf $BACKUP_SOURCE_DIR/$RESTORE_FILENAME
```

```
# Restore a BZipped backup here...
tar xvjf $BACKUP_SOURCE_DIR/$RESTORE_FILENAME
```

As you can see, this script is simple enough. All we have to do is to pass it the file name of the backup we want to restore as a parameter (just the file name, not the full path), and it restores the backed up files into their original locations. Make sure the script is executable: open a terminal and run `chmod 700 restore.sh`.

13.4.2. Making a Recovery CD-ROM

There is a way to be prepared in case of “total disaster”, and that is by making a **full** backup of your system. Programs such as mkCDrec can be very useful to get you up and running in a matter of minutes. You can find it, together with its documentation on the mkCDrec web site (<http://mkcdrec.ota.be>).

mkCDrec allows you to do multiple-CD-ROM volumes, disk cloning (copying the full contents of a disk or partition to another with similar characteristics — at least the same size), and many more.

In order to restore a system with mkCDrec you only have to boot with the first CD-ROM of the multiple-CD-ROM volume and then follow the on-screen instructions.

13.5. Problems Arising at Boot Time

It could happen that your system hangs during boot up. If so, don’t panic, just keep reading.



The next sections are not introduced in any particular order.

13.5.1. System Hanging During Boot

If your system hangs during Rebuilding RPM database or Finding module dependencies, just press **Ctrl-C**. This allows the system to skip this step and continue to boot. Once booted, execute `rpm --rebuilddb` as root if the system hang was at the Rebuilding RPM database phase. If the system hang was at the Finding module dependencies phase you have most likely been through a kernel upgrade, but haven’t done it correctly. Check if the files in the `/boot` and `/lib/modules` directories match the current kernel version (i.e., have the current version number attached).

If the boot process hangs at `RAMDISK: Compressed image found at block 0` you have a corrupted `initrd` image. Either try to boot another boot entry or boot an emergency system and remove or change the `initrd=` section in `/etc/lilo.conf`

13.5.2. Filesystem Check on Boot Fails



The information below applies to ext2 and ext3 filesystems only. If you use a different filesystem, please check its documentation.

If, for any reason, you didn’t shut your box down properly, the system runs a routine filesystem check during the next boot. Sometimes it may fail to do this on its own and asks for the root password and drops you to a console. Execute `e2fsck -py [device]` where `[device]` is the name of the partition on which the automatic check failed. The `-p` switch tells `e2fsck` to make all the necessary repairs without asking, `-y` assumes you answer `yes` to all questions. When the check and repair phase is over, press **Ctrl-D** to leave the emergency console. The system will reboot.

If you get this error regularly, there may be bad blocks on your disk. Execute `e2fsck -c [device]` to find out. This command automatically marks any bad blocks and prevents the filesystem from storing data in these blocks. `e2fsck` checks the file system automatically only if it hasn’t been unmounted properly during

the previous system shutdown; or if the `maximal mount count` has been reached. To force a check, use the `-f` option.



The verification for bad blocks on a disk should only be done on unmounted file systems, and can take a long time to complete.

13.5.3. X Doesn't Start

If you boot into X by default and have managed to break your X configuration somehow, and cannot enter X anymore, you can log into a console and use XFdrake to reconfigure X. You can also boot into a different runlevel, fix X's configuration with XFdrake and reboot into X.

13.5.3.1. Booting Into a Different Run Level

The default run level GNU/Linux boots to is defined in the `/etc/inittab` file. Look for an entry like `id:5:initdefault:.` To boot into run level 3 (the console), you have to define that run level on the boot prompt. Under LILO, press the **Esc** key once and type `linux init 3`. Under GRUB, press the **E** key twice, add `init 3`, press the **Enter** key and then the **B** key to boot.

For a more detailed description about run levels, please refer to the *The Start-Up Files: init sysv* chapter of Mandriva Linux's *Reference Manual*.

13.5.3.2. Configuring X From The Console

To reconfigure X using XFdrake from the console, simply type `XFdrake` as `root`.

Using XFdrake is no different to the graphical environment except that you won't have nice icons and may not be able to use the mouse pointer. To move down you have to press the right or down arrow keys on your keyboard; to move up press the left or up keys on your keyboard. You can also use the **Tab** key to move between the different options/buttons. The text on the currently selected button/option will be highlighted with a different color. Press the **Enter** key to activate it.

13.6. Bootloader Issues

13.6.1. Bootloader Reinstall

It may happen that you make a mistake and wipe your disk's MBR (Master Boot Record), or some misbehaving program erases it, or you dual boot with Windows® and catch a virus which suppresses it. So, you think you won't be able to boot your system anymore, right? **Wrong!** There are many ways to recover the boot record.

To recover your bootloader you **need** a boot disk. Without a boot disk of some kind you might be completely lost, unless you made a backup of your MBR, see *Backing Up and Restoring the MBR*, page 143.

Reboot your computer using the boot disk. What you do next varies according whether you use LILO or GRUB. No matter which bootloader you use, all the commands you must execute need to be run as `root`.

13.6.1.1. With LILO

If you use LILO, you only need to issue the following at the command prompt: `/sbin/lilo`. This command reinstalls LILO on your disk's boot sector and fixes the problem.

13.6.1.2. With GRUB

If you use GRUB things are a little bit different to that of LILO.



In the following example we assume that you are trying to install GRUB in the MBR of your first IDE drive, and that the file `stage1` is in the `/boot/grub/` directory.

First, invoke GRUB's shell by issuing the `grub` command. Once there, issue the following command: `root (hd0,0)`. This will tell GRUB that the files it needs are in the first partition (0) of your first hard disk (`hd0`). Then issue the following command: `setup (hd0)`. This installs GRUB in the MBR of your first hard disk. That's it!

You can also try to use `grub-install /dev/hda` to install GRUB on your first hard drive's MBR, but the method described above is the preferred one.

13.6.1.3. Some Considerations for Dual-Booting Systems

Windows 9x, NT, 2000 and XP upgrades. If you run a dual-boot system, be very careful to always have a GNU/Linux boot disk prepared. If you don't have a boot disk, and you (re)install Windows® (all versions) you won't be able to boot GNU/Linux after the Windows® upgrade because Windows® rewrites the MBR **without any warning at all**.

13.6.2. Backing Up and Restoring the MBR

To make a backup copy of your hard disk's MBR, insert a blank floppy in your floppy disk drive and issue the following:

```
# dd if=/dev/hda of=/dev/fd0/mbr.bin bs=512 count=1
```

If you want to restore a backed up copy of your hard disk's MBR, insert the floppy containing it into your floppy disk drive and issue the following:

```
# dd if=/dev/fd0/mbr.bin of=/dev/hda bs=512
```



The above examples assume that the MBR of your first IDE hard disk (`/dev/hda`) is backed up to a file named `mbr.bin` on your first floppy diskette drive (`/dev/fd0`) and should be run as the `root` user.

13.7. Filesystem Issues

13.7.1. Repairing a Damaged Superblock



The information below only applies to `ext2` and `ext3` filesystems. If you use a different filesystem, please check its documentation.

The superblock is the first block of each `ext2FS`/`ext3FS` partition. It contains important data about the file system, such as its size, free space, etc. (it is similar to the method used by FAT partitions). A partition with a damaged superblock cannot be mounted. Fortunately, `ext2FS`/`ext3FS` keeps several superblock backup copies scattered over the partition.

Boot your system with a boot disk. The backup copies' location depends on the block size of the filesystem. For filesystems with 1 KB block sizes it is at the beginning of each 8 KB (8192 bytes) block. For filesystems with 2 KB sizes it is at the beginning of each 16 KB (16384 bytes) block, and so on. You can use the `mke2fs`

`-n [your_disk_device_name]` command to find out at which byte positions the superblock copies are. Assuming a 1 KB block size, the first backup copy is in byte number 8193. To restore the superblock from this copy, execute `e2fsck -b 8193 /dev/hda4`; change `hda4` accordingly to reflect the name of your damaged partition. If that block also happens to be damaged, try the next one at byte number 16385, and so on until you find a suitable one. Reboot your system to activate the changes.

13.7.2. Recovering Deleted Files

We discuss ways of recovering deleted files and directories. Please bear in mind that recovery tools are not magical, and they will only work depending on how recently you deleted the file(s) you are trying to recover.

You might be wondering how to recover files you accidentally deleted. There are some utilities designed for GNU/Linux's ext2 filesystem which allow you to recover deleted files and directories. However they won't recover the files you deleted a few months ago because of disk usage, space marked as "free" will have been overwritten. So the **best** way to protect against accidental or not so accidental deletions is by making backups.



There are not (as yet) tools to recover files deleted on `reiserfs` file systems. Keep in touch with the ReiserFS home page (<http://www.namesys.com>) for the latest news about it.

One recovery tool is `Recover`. It's an interactive tool. You can find it in the `contribs` CD-ROM or on the `Rpmfind` web site (<http://www.rpmfind.net>). Once you have the RPM, install it. Then run it with `recover` and answer the questions it asks you. The questions help you to set a time span to look for deleted files and directories to minimize the time it takes to do the search¹.

Once the tool finishes its search, it asks you where you want to save the recovered files and directories. Pick a directory of your choice, and you have all the files and directories recovered into it. Note that you won't be able to recover the file names, just their contents, but you can inspect them or try to rename them with different names until you get the right one. This is better than nothing.



There are also mini-HOWTOs related to "undeletion" for ext2, look at `Ext2fs-Undeletion` (<http://www.tldp.org/HOWTO/mini/Ext2fs-Undeletion.html>) and undeletion of whole directory structures (<http://www.tldp.org/HOWTO/mini/Ext2fs-Undeletion-Dir-Struct/index.html>).

13.8. Recovering from a System Freeze

When stuck in a "freeze", your computer doesn't respond to commands anymore and input devices such as keyboard and mouse seem to be blocked. This is a worst-case scenario and could mean that you have a very severe error in either your configuration, your software or your hardware. We will show you how to deal with this annoying situation.

In the case of a system freeze, your top priority should be trying to shutdown your system properly. We assume you are running under X. So try these steps consecutively:

1. Try to kill the X server by pressing the **Ctrl-Alt-Backspace** keys.
2. Try to switch to another console by pressing the **Ctrl-Alt-Fn** keys (where `n` is the console number, from 1 to 6). If you succeed, login as `root` and issue the command: `kill -15 $(pidof X)` or the command `kill -9 $(pidof X)`, if the first command shows no effect. Check with `top` to see if X is still running.
3. If you are part of a local network, try to use `ssh` to connect into your machine from another. It is advisable to `ssh` into the remote machine as an unprivileged user and then use the `su` command to become `root`.

1. You can search for **all** deleted files too by appending the `-a` option, but it takes much longer...

4. If the system doesn't respond to any of these steps, you have to go through the SysRq (System Request) sequence. The SysRq sequence involves pressing and holding three keys at once: the left **Alt** key, the **SysRq** key (labeled **Print Screen** on older keyboards) and a letter key.
 - a. **Alt-SysRq-R** puts the keyboard in "raw" mode. Now try pressing **Alt-Ctrl-Backspace** again to kill X. If that doesn't work, carry on.
 - b. **Alt-SysRq-S** attempts to write all unsaved data to disk ("sync" the disk).
 - c. **Alt-SysRq-E** sends a termination signal to all processes, except for `init`.
 - d. **Alt-SysRq-I** sends a kill signal to all processes, except for `init`.
 - e. **Alt-SysRq-U** attempts to re-mount all mounted filesystems read-only. This removes the "dirty flag" and prevents a filesystem check upon reboot.
 - f. **Alt-SysRq-B** reboots the system. You might just as well press the "reset" button on your machine.



Remember that this is a sequence, i.e. you have to press one combination after the other in the right order: **R**aw, **S**ync, **tE**rm, **k**ill, **U**mount, **rE**boot². Read the kernel documentation for more information on this feature.

5. If none of the above helps, cross your fingers and press the "reset" switch on your machine. If you are lucky, GNU/Linux will just run a disk check upon reboot.

By all means, try to find out what causes these lockups because they can do severe damage to the filesystem. You might also want to consider using one of the journaling filesystems included in Mandriva Linux: `ext3`, `reiserfs`, etc. which handle such failures more gracefully. However, replacing `ext2FS` with `reiserfs` requires reformatting your partitions. You can use `tune2fs -j /dev/hdaN` to convert the filesystem in the Nth partition of the first IDE disk from `ext2FS` to `ext3FS`.

13.9. Killing Misbehaving Apps

Well, this one is not so hard after all. You have many ways to do it. You can do it by finding the PID of the program which stopped responding, and then using the `kill` command to terminate it, or you can use the `xkill` tool or other graphical tools such as the ones that show the process tree.

13.9.1. From the Console

The first thing to do to terminate a misbehaving program is to find its PID, or process ID. To do so, execute the following from a console: `ps aux | grep mozilla-firefox-bin`, supposing that Firefox is the misbehaving program. You will get something like the following, which tells you among other things that Firefox was started by user `peter` and that its PID is 3505:

```
peter  3505  1.7  5.0  82208 25804 ?          Sl   09:30   0:01 /usr/lib/mozilla-firefox-1.0.6/mozilla-firefox-bin
```

Now that we have the PID of the misbehaving program, we can execute the `kill` command to terminate it. So we execute the following: `kill -9 3505`, and that's it! Firefox is killed. Note that this is **only** to be used when the program doesn't respond to your input anymore. **Do not** use it as a standard means of exiting from applications.

Actually, we sent the `KILL` signal to the process number 3505. The `kill` command accepts other signals besides `KILL`, so you can have greater control over your processes. For more info, see `kill(1)`.

13.9.2. Using Graphical Monitoring Tools

You can also use the graphical process' status tools (such as KPM, KSysGuard, and GTOP to name a few) which allow you to point to the process name and with one click send that process a signal or just kill that process.



If you are using KDE, you can press the **Ctrl-Alt-Esc** keys: the pointer changes to a skull with crossed bones and you can simply click on the window of the misbehaving application to kill it.

13.10. Miscellaneous

Some considerations on newer hardware such as legacy-free systems, nVidia® and ATI 3D® graphics accelerator cards, winmodems and other things that don't fit in the preceding sections.

13.10.1. Legacy-Free Systems

Hardware manufacturers have recently introduced what they call "legacy-free systems", mainly on laptops³, but there are also legacy-free desktop computers. This basically means that the BIOS has been considerably reduced to allow you only to choose which media to boot from. Mandriva Linux will be able to configure everything properly.

13.10.2. nVidia and ATI 3D Graphics Cards

Computers with nVidia or ATI graphics cards need a patched kernel to be able to use OpenGL hardware 3D acceleration on OpenGL-compatible applications. If you own a Mandriva Linux — PowerPack Edition, the kernel should have been installed by DrakX. If this is not your case, please obtain and install the corresponding packages. You can visit nVidia's web site (<http://www.nvidia.com>) and ATI's web site (<http://www.ati.com>) and download the appropriate drivers, or you can download the RPM packages from Mandriva Club (<http://club.mandriva.com>). Then run Mandriva Linux Control Center and re-configure X from there.

13.10.3. Winmodems

winmodems are also called controller-less modems or software modems. Support for these peripherals is improving. Drivers do exist, but most of them are in binary form and available only for newer kernel versions.

If you have a PCI modem, look at the output of `cat /proc/pci` run as the `root` user from a terminal window. It tells you the device's I/O port and IRQ. Then use the `setserial` command (for our example, the I/O address is `0xb400`, the IRQ is `10` and the modem is the 4th serial device) as follows:

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

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

If you happen to have a software modem in your machine, and you have a Mandriva Club account, you might find an RPM package that supports your modem (try searching on the `ltmodem` package for instance). You should also take a look at the web site of your modem's manufacturer and at the `linmodems` (<http://linmodems.org/>) and the Winmodems are not modems (<http://start.at/modem/>) web sites.

3. Refer to the great Linux on Laptops (<http://www.linux-laptop.net>) web site for more information on your laptop make/model.

13.10.4. My Computer is “slow”

If you notice your computer is really slow, or significantly slower than with other GNU/Linux versions, you might overcome this “problem” by disabling ACPI support. To do so, add the following to your `/etc/lilo.conf` file:

```
append="acpi=off"
```

If the file already has an `append=` line, only add `acpi=off` at its end. Running `lilo -v` as `root` and rebooting your computer will make the changes effective.

13.11. Mandriva Linux’s Specific Troubleshooting Tools

Each administration tool (the ones started from Mandriva Linux Control Center) is a potential trouble fixing tool. You can use all these tools to revert configuration changes, to add or remove software, to update your system with the latest fixes from Mandriva, etc.

If you think you have found a bug in any of our tools, please feel free to submit a bug report using Drakbug, our automated bug report tool.

13.12. General Guidelines for Solving a Problem under Mandriva Linux

Here are the different means available to you in your problem-solving quest. Try the first option and only then, if that does not work, try the second, and so on.

13.12.1. Search the Internet

The various Internet sites previously mentioned are excellent starting points. They deal with general **and** very specific aspects of your potential problems. Finally, try a general search engine such as Google™ or, as mentioned above, the Linux-specific Google™ search engine. And do not hesitate to use the Advanced search (http://www.google.com/advanced_search) option with very detailed questions, such as the error message you are receiving.

13.12.2. Mailing Lists and Newsgroups Archives

The previous searches may lead you to general answers which hide the results of your specific question amongst many other answers. To refine your search, you can try the following.

First, try to find a list which seems specifically geared to your problem, then perform a search in its archive pages.

Example

You’ve noticed some strange behavior while trying to use GRUB with a minix partition.

One of the results of a search using the “grub mailing list” keywords in Google™ is a link to the *GRUB mailing list archive* (<http://mail.gnu.org/archive/html/bug-grub/>). It even offers a search engine, which when searched for “Minix” leads you directly to a patch.



Note that not all archives have an embedded search engine. However, using Google™ as an example, you can easily use the advanced field `domain` to limit your search to the specific site hosting the archive. This strategy may also be used to exclude sites which keep returning garbage.

For a newsgroups search, Google Groups™ (<http://groups.google.com/>) maintains an archive of an amazingly large number of newsgroup channels.

13.12.3. Directly Contacting the Person in Charge

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, so your anguished question on the use of the `cd` command will most likely... be ignored!

The addresses will be found either on the home page of a project's site or in the software documentation.

A last word: 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. If your problem isn't fixed today, it may be tomorrow...

13.12.4. Mandriva Business Services

Finally, when facing a really challenging situation, corporate users (especially) might consider hiring one of Mandriva's consultants to address their specific needs.

This is one of the strong suits of open-source products: we have the source, we have the power! Therefore, almost any problem, no matter how complex, specific or high level, may be solved right in the heart of the software.

You might also want to customize your Linux environment to meet very precise goals. For example, you could use Mandriva Linux as a custom routing application on special devices. Know that Mandriva consulting services (<http://www.mandriva.com/enterprise/products/>) can help you.

13.13. Final Thoughts

As you have seen there are many more ways to recover from an emergency than by re-installing the whole system again⁴. Sure, you need a little expertise in applying some of the techniques described in this chapter, but with a little practice you will gain such expertise. However, we hope that you will never need to really master these techniques ... although it does not hurt to know them. We hope that the instructions and examples given will be useful when you are in need. Good luck recovering from an emergency!

4. The usual way to fix things in some other operating systems...

Appendix A. Glossary

account

On a UNIX[®] system, the combination of a name, a personal directory, a password and a shell which allows a user to connect to the system.

alias

A mechanism used in a shell in order to substitute one string for another before executing a command. You can see all aliases defined in the current session by typing `alias` at the prompt.

ACPI

Advanced Configuration and Power Interface. A feature used to recognize and configure hardware and for power management. Unlike APM, which relies on the BIOS only, ACPI also relies on the operating system, making its control more simple for the user. ACPI also brings power management capabilities to servers and workstations.

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 supported) the estimated remaining battery life. However, newer laptops are based on ACPI rather than APM.

See Also: ACPI.

ARP

Address Resolution Protocol. The Internet protocol used to dynamically map an Internet address to a physical (hardware) address on a local area network. This is limited to networks which 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, usually the Linux default character set, unless you have chosen to use something like UTF-8) contain ASCII as their lower half.

See Also: ISO 8859, UTF-8.

assembly language

Is the programming language that is closest to the computer, which is why 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. Its 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 CD-ROM drives and magnetic tape drives. IDE controllers equipped with this extension are also referred to as EIDE (*Enhanced IDE*) controllers.

See Also: IDE.

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 can be conveyed efficiently from point to point. ATM is a circuit switched packet network technology oriented towards high speed (multi-megabit) networks.

atomic

A set of operations is said to be atomic when they execute all at once and cannot be preempted. It is commonly used for an "all or nothing" set: either all of the operations perform successfully or none of them are taken into account. It might also be used for essential or very simple operations, like the sum of two integral numbers.

background

In shell context, a process is running in the background if you can type commands that are captured by the process while it is running. It is the opposite of a foreground process.

See Also: job, foreground.

backup

A means of saving important data to a safe medium and location. Backups should be made regularly, especially with more critical information and configuration files (the most important 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

A processing mode where jobs or instructions which are submitted to the CPU are executed sequentially until all have been processed.

beep

The little noise your computer's speaker emits 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

The name given to the process of testing the beta version of a program. Programs usually get released in "alpha", "beta" and "release candidate" states for testing prior to final release.

binary

In the context of programming, binaries are the compiled, executable code.

bit

Stands for *Binary digit*. A single digit which can take the values 0 or 1, because calculation is done in base two. It is the most basic unit of digital information.

block mode files

Files whose contents are buffered. All read/write operations for such files go through buffers, which allow for asynchronous reads and writes to the underlying hardware, which prevents the system from making disk accesses if the data is already in a buffer.

See Also: buffer, buffer cache, character mode files.

boot

The procedure taking place when a computer is first switched on, where peripherals are recognized sequentially and where the operating system is loaded into memory.

boot disk

A bootable disk (floppy, CD, DVD, or any other device) containing the code necessary to load the operating system from the hard disk (sometimes it is self-sufficient).

bootloader

This is a program which starts the operating system. Many bootloaders give you the opportunity to load more than one operating system by allowing you choose between them from a menu. Bootloaders such as GRUB and LILO 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 technically advanced 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 of fixed size, which can be associated with a block mode file, a system table, a process and so on. The buffer cache maintains coherency of all buffers.

See Also: 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 Also: buffer.

bug

Illogical or incoherent behavior of a program in a special case, or a behavior that 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. Admiral Grace Hopper, having discovered this, declared "It's a bug!", and since then the term has remained. Note that this is only one of the many stories which attempt to explain the term *bug*.

byte

A sequence of, usually, eight consecutive bits, which when converted to base ten result in an integer number between 0 and 255. A byte is always "atomic" on the system, meaning that it is the smallest addressable unit.

See Also: bit.

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: A protocol used by ISPs to authenticate their clients. In this scheme, a value is sent to the client (the machine making the connection), which it uses to calculate a hash based on the value. The client sends the hash back to the server for comparison to the hash calculated by the server. This authentication method is different to 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 with 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 File System. The successor to the SMB file system, used on DOS systems.

See Also: SMB.

client

A program or computer which sporadically connects, for a given period of time, 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 which initiates the connection, the remote end receiving the call is designated as the server. It is one of the components of a **client/server system**.

See Also: server.

client/server system

System or protocol consisting of a **server** and one or more **clients**.

command line

Provided by a shell and which 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 its clones, it is the state of the program in which pressing a key does not insert the character into the file being edited, but instead performs an action specific to the 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 which is human readable (well, with some training) and which is written in some programming language (C, for example) into a binary file which is machine readable.

completion

The ability of a shell to automatically expand a substring to a filename, user name or other item, as long as there is a match.

compression

A way to shrink files or decrease the number of characters sent over a communications link. File compression programs include `compress`, `zip`, `gzip`, and `bzip2`.

console

This is the name given to what used to be called terminals. They were the machines (a screen plus a keyboard) connected to one big central mainframe. On PCs, 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. They allow 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. It is the basic unit of transmission across an IP network. You might also hear this called a "packet".

dependencies

The stages of compilation which need to be satisfied before going on to other compilation stages in order to successfully compile a program. This term is also used where one set of programs you wish to install are dependent on other programs which may or may not be installed on your system, in which case you may get a message telling you that the system needs to "satisfy dependencies" in order to continue the installation.

desktop

If you're using the X Window System, the desktop is the place on the screen where 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 and other network settings from a server.

directory

Part of the file system structure. Files or other directories can be stored within a directory. Sometimes there are subdirectories (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 subdirectories 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. In graphical environments it is also known as a folder.

discrete values

Are values which are non-continuous. That is, there's some kind of "spacing" between consecutive values.

distribution

Is a term used to distinguish one GNU/Linux manufacturer's 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 DLCIs are normally assigned by the Frame Relay network provider.

DMA

Direct Memory Access. A facility used in the PC architecture which allows 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. Bus mastering allows a controller to talk to other devices without going through the CPU.

DNS

Domain Name System. The distributed name and address mechanism used in the Internet. This mechanism allows you to map a domain name to an IP address, allowing you to look up a site by domain name without knowing the IP address of the site. DNS also allows reverse lookup, allowing you to obtain a machine's IP address from its name.

DPMS

Display Power Management System. Protocol used by all modern monitors to manage power saving features. Monitors supporting these features are commonly called “green” monitors.

echo

Occurs when the characters you type are shown on the screen, such as in a user name entry field, for example. Some programs may also mask what is typed for security reasons. The example is a password prompt showing an *, or even nothing at all, for each typed char instead of the character itself.

editor

Is a term typically used for programs which 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.

email

Stands for Electronic Mail. This is a way to send messages electronically. 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 such as pine or mutt which are text-mode clients, or GUI clients such as 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 a string with quotes to prevent the shell from interpreting that string. For example, when you need to use spaces in a command line and then pipe the results to some other command you have to put the first command between quotes or precede the spaces with a \ (“escape” the command) otherwise the shell will interpret it incorrectly and your command won’t work as expected.

ext2

Short for the “Extended 2 file system”. This is GNU/Linux’s native file system and has the characteristics of any UNIX[®] file system: support for special files (character devices, symbolic links, etc), file permissions and ownership, and other features.

FAQ

Frequently Asked Questions. A document containing a series of questions and 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 FAQs. Generally, they are very good sources of information.

FAT

File Allocation Table. File system used by DOS and Windows[®].

FDDI

Fiber Distributed Digital Interface. A high-speed network physical layer, which uses optical fiber for communication instead of wire. Mostly used on large networks, mainly because of its price. It is rarely seen as a means of connection between a PC and a network switch.

FHS

File system Hierarchy Standard. A document containing guidelines for a coherent file tree organization on UNIX[®] systems. Mandriva Linux complies with this standard in most aspects.

FIFO

First In, First Out. A data structure or hardware buffer where 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 physical media (hard drive, floppy, etc.) in a consistent manner. Examples of file systems are FAT, GNU/Linux' ext2fs, ISO9660 (used by CD-ROMs) 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 single connection point to the outside network, and which filters and controls the activity on some ports, or makes sure that only some specific interfaces may have access to, or can be accessed from, the outside world.

flag

Is an indicator (usually a bit) that is 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 of 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 that is currently running and has keyboard and screen control. You have to wait for such a process to finish in order to be able to type commands again. *See Also:* job, background.

Frame Relay

Frame Relay is a network technology ideally suited to carrying traffic which is of a bursty or sporadic nature. Network costs are reduced by having many Frame Relay customers sharing the same network capacity and relying on them wanting to make use of the network at slightly different times.

framebuffer

Projection of a video card's RAM into the machine's address space. This allows applications to access the video RAM without the chore of having to talk to the card. All high-end graphical workstations use frame buffers.

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 entire visible area of your display.

gateway

Machine or device giving a local network access to an outside network.

GFDL

The GNU Free Documentation License. The license which applies to all Mandriva Linux documentation.

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, the recommended solution is to replace them as much as possible by the PNG format.

See Also: PNG.

globbing

In the shell, the ability to group a certain set of file names with a globbing pattern.

See Also: 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 was initiated by Richard Stallman at the beginning of the 1980s, 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. Linux borrows, among others, two things from GNU: its C compiler, `gcc`, and its license, the GPL.

See Also: GPL.

GPL

General Public License. The license of the GNU/Linux kernel, it goes the opposite way to all proprietary licenses in that it applies no restrictions as to copying, modifying and redistributing the software, as long as the source code is made available. The only restriction is that the persons to whom 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. A great majority of users prefer a GUI to a CLI (*Command Line Interface*) for ease of use, even though the latter is far more versatile.

guru

An expert. Used to qualify someone particularly skilled, but also of valuable help for others.

hardware address

This is a number which 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

A file which can't be "seen" when doing a `ls` command without options. The names of hidden files begin with a `.` and are used to store the user's personal preferences and configurations for the different programs he uses. For example, bash's command history is saved into `.bash_history`, a hidden file.

home directory

Often abbreviated as "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 which are connected to a network.

HTML

HyperText Markup Language. The language used to create web documents.

HTTP

HyperText Transfer Protocol. The protocol used to connect to web sites and retrieve HTML documents or files.

icon

Is a little drawing (normally sized 16×16, 32×32, 48×48 and sometimes 64× 64 pixels) which in a graphical environment represents 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 may contain up to two devices, and the speed of the bus is limited by the device on the bus with the slowest command queue (and not the slowest transfer rate!).

See Also: ATAPI, SATA, S-ATA.

IMAP

Internet Message Access Protocol. A protocol which allows you to access your email messages on a remote server, without the need to transfer them locally first; as opposed to the POP mail retrieval protocol.

See Also: POP.

inode

Entry point leading to the contents of a file on a UNIX[®]-like filesystem. An inode is identified in a unique way with 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 any 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 such as the completion of an abbreviation, right justify at the end of the line, ...). One gets out of it pressing the **Esc** key, (or **Ctrl-I**).

Internet

Is a huge network which connects computers around the world.

IP address

Is a numeric address consisting (in version 4, also called IPv4) of four parts which identifies your computer on a network. IP addresses are structured in a hierarchical manner, with top level and national domains, domains, sub-domains and each machine's personal address. An IP address will 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 which never change, they are permanently assigned. Dynamic IP addresses mean that an IP address will change with each new connection to the network. Most home users typically have dynamic IP addresses while most corporate users typically have static IP addresses.

IP masquerading

This is a technique where a firewall is used to hide your computer's true IP address from the outside. Typically, any outside network connections you make through 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 on your internal network.

IRC

Internet Relay Chat. One of the few Internet standards for live speech. It allows for channel creation, private talks and file exchange. It also allows servers to 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 one channel are only visible to the people connected to that channel. Two or more users can 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. ISA is still commonly found on SCSI cards supplied with scanners, CD writers and some other older hardware.

ISDN

Integrated Services Digital Network. A set of communication standards for voice, digital network services and video. It has been designed to eventually replace the current phone system, known as PSTN (*Public Switched Telephone Network*) or POTS (*Plain Old Telephone Service*). ISDN is known as a circuit switched data network.

ISO

International Standards Organization. A group of companies, consultants, universities and other sources which enumerate standards in various disciplines, including computing. The papers describing standards are numbered. The standard number iso9660, for example, describes the file system used on CD-ROMs.

ISO 8859

The ISO 8859 standard includes several 8-bit extensions to the ASCII character set. 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 supports the following languages: Afrikaans, Basque, Catalan, Danish, Dutch, English, Faroese, Finnish, French, Galician, German, Icelandic, Irish, Italian, Norwegian, Portuguese, Scottish, Spanish, and Swedish.

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 is a modification of ISO 8859-1 to covers these needs.

See Also: ASCII, UTF-8.

ISP

Internet Service Provider. A company which sells Internet access to customers, either over telephone lines or high-bandwidth circuits such as dedicated T-1 circuits, DSL or cable.

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.

job

In a shell context, a job is a process running in the background. You can have several jobs running in the same shell and control each job independently.

See Also: foreground, background.

journaling

Journaling adds robustness to a file system, by making it transactional. Thus, instead of physically writing data at the moment it's asked for, a journal of the writes is kept, and data is written "in a block" at a later time which also has a great impact on performance and on the time needed to analyze and fix the file system, if needed.

kernel

Is the core 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 which 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 editor was started. The text areas may be recalled to be inserted again, and the structure is ring-like.

LAN

Local Area Network. Generic name given to a network of machines connected to the same physical wiring in a reduced geographical area, such as the same office or building.

See Also: WAN.

launch

Is the action of invoking, or starting, a program.

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

The last stage of the compilation process, consisting of linking together all object files in order to produce an executable file, and matching unresolved symbols with dynamic libraries (unless a static linkage has been requested, 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 which stores corresponding codes (or tags) and their meanings. 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 store a manufacturer's product codes and associated configuration information. This is one line from that table, giving information about item CTL0001

```
"CTL0001"      "sb"      "Creative Labs|SB16"      "sound" "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

Small documents containing the definitions of a command and its usage, to be consulted with the `man` command. The first thing one should (learn how to) read when learning about a command one isn't familiar with.

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 LILO), 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.

MPEG

Moving Picture 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...

mount point

Is the place or directory where a partition or another device is attached to the GNU/Linux filesystem. For example, your CD-ROM 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 contents. This term is partly obsolete due to the "supermount" feature, so users do not need to manually mount removable media.

See Also: mount point.

MSS

The *Maximum Segment Size* is the largest quantity of data which can be transmitted at one time across an interface. If you want to prevent local fragmentation MSS would equal the MTU IP header.

MTU

The *Maximum Transmission Unit* is a parameter which determines the size of the largest datagram which 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 without fragmentation. 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 PPP interface.

multitasking

The ability of an operating system to share CPU time between several processes. At a 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 runs again. This operation is called a context switch, and is done several times per second, thereby making it 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: in preemptive multitasking the operating system is responsible for taking away the CPU and passing it to another process; cooperative multitasking is where the process itself gives back the CPU. The first variant, used by GNU/Linux, is obviously the better choice because no program can consume the entire CPU time and block other processes. 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 user 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 is 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.

NFS

Network File System. A network file system created by **Sun Microsystems** in order to share files across a network in a transparent way.

newsgroups

Discussion and news areas which 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 (linux), and specifically, Mandriva Linux (mandrake). Newsgroups are broken down in this fashion to make it easier to search for a particular topic.

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 may consist of an entire LAN, or just a part of it. 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 which is made available to the 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 Amiga[®]OS, Mac OS[®], FreeBSD[®], OS/2[®], UNIX[®], and Windows[®] in all its variants.

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.

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, but none of the information is encrypted. CHAP is a more secure, and thus preferred, authentication protocol.
See Also: CHAP.

pager

A program which displays a text file one screen at a time, making it easy to move back and forth and search for strings in this file. We suggest you to use `less`.

password

Is a secret word or combination of words or letters which 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

A file containing a list of corrections to issue to 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** (or **full**) path is the position of a file or directory in relation to the root directory.

PCI

Peripheral Component Interconnect. A bus created by **Intel** which today is the standard bus for PC and other architectures. It is the successor to 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 from 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 bitmap 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. All PCI devices are Plug'N'Play.

POP

Post Office Protocol. One common protocol used for retrieving mail from an ISP. IMAP is an example of another remote-access mail protocol.
See Also: IMAP.

porting

One of two ways to run a program on a system it was not originally intended for. 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 multiplication has higher precedence than the addition. If you want to evaluate the addition first, then you have to add parenthesis like this: $(4 + 3) * 2$. When you do this, you'll get 14 as the result since the parenthesis have higher precedence than the addition and the multiplication, so the operations in parenthesis get evaluated first.

preprocessors

Are compilation directives which 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 communications 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 (for example, HTTP and FTP). It maintains a cache of previous requests, so that a machine which makes a request for something which is already cached will receive it quickly, because it will get the information from the local cache. Proxies are very useful on low bandwidth networks (such as modem connections). Sometimes the proxy is the only machine able to access the outside network.

pull-down menu

Is a menu that is "rolled" with a button in one of its corners. When you press that button, the menu "unrolls" itself, showing you the full menu.

quota

Is a method of restricting disk usage and place limits on users. Administrators can restrict the size of home directories for a user by setting quota limits on specific file systems.

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 across an array of disks using different schemes. At first, this was implemented using low-cost, older, 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 may be directly accessed.

read-only mode

For a file means that the file cannot be written to. You may read its content but you cannot modify it.
See Also: read-write mode.

read-write mode

For a file, it means that the file can be written to. You may read its content and modify them.
See Also: read-only mode.

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` and 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 contains the root directory.

route

Is the path which your datagrams take through the network to reach their destination. It is the path between one machine and another in a network.

RPM

RPM Package Manager. A packaging format developed by **Red Hat** in order to create software packages, it is used in many GNU/Linux distributions, including Mandriva Linux.

run level

Is a configuration of the system software which only allows certain selected processes to exist. Allowed processes are defined, for each runlevel, in the file `/etc/inittab`. Usually, there are seven defined runlevels: 0, 1, 2, 3, 4, 5, 6 and switching between them can only be achieved by a privileged user by means of executing the commands `init` and `telinit`.

SATA, S-ATA

Serial ATA. The successor to the ATA specification. First generation SATA has a bandwidth of 1.5Gbps, but the serial link and underlying technologies allow for much greater bandwidths, while parallel ATA has reached its practical limits with UDMA133.

See Also: ATAPI, IDE.

script

shell scripts are sequences of commands to be executed as if they were sequentially entered in the console. 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 to be connected to it. Unlike IDE, a SCSI bus is not limited by the speed at which the peripherals accept commands. Usually only high-end machines integrate a SCSI bus directly on the motherboard, therefore most PCs need add-on cards.

security levels

Mandriva Linux's unique feature which allows you to set different levels of restriction 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.

segmentation fault

A segmentation fault occurs when a program tries to access memory that is not allocated to it. This generally causes the program to stop immediately.

server

A program or computer which provides a feature or service and awaits connections from **clients** to execute their orders or give them the information they ask for. 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**.

See Also: client, client/server system.

shadow passwords

A password management suite on UNIX[®] systems in which the file containing the encrypted passwords is not world-readable, unlike that usually found with a 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 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, which only allows a single user to log into and use the system at any one time.

site dependent

Means that the information used by programs such as `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 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 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 as the file descriptor to which the process writes errors. It is usually the computer's screen.

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. It is usually the computer's keyboard.

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. It is usually the computer's screen.

See Also: standard error, standard input.

streamer

Is a device which takes "streams" (not interrupted or divided into 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 was at first 800x 600 x 16 colors, quickly extended to 1024x768 x 16 colors, and beyond.

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 which may be used, read the `man` pages or try to pass the `--help` switch to the program (i.e.. `program --help`).

symbolic links

Are special files, containing nothing but a string which references 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 which 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 have established 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, such as `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.

TLDP

The Linux Documentation Project. A nonprofit organization which maintains GNU/Linux documentation. It's mostly known for documents such as HOWTOs, but it also maintains FAQs, and even a few books. *See Also:* FAQ.

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 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 may be a file, a server or other item. The syntax for a URL is `protocol://user:password@server.name[:port]/path/to/resource`. When only a machine name is given and the protocol is `http://`, it defaults to retrieving the file that the server is configured to show by default, usually it is the `index.html` file.

username

Is a name (or more generally a word) which identifies a user on a system. Each username is attached to a unique and single UID (user ID) *See Also:* `login`.

UTF-8

Unicode Transformation Format 8. It is an octet (8-bit) lossless encoding of Unicode characters. UTF-8 encodes each Unicode character as a variable number of 1 to 4 octets, where the number of octets depends on the integer value assigned to the Unicode character. It is an efficient encoding of Unicode documents which mostly use US-ASCII characters because it represents each character in the range U+0000 through U+007F as a single octet. UTF-8 is the default encoding for XML. *See Also:* ISO 8859, ASCII.

variables

Are strings which 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 which 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. For example, it is the author of the SVGA standard.

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, **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 with 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 which depends on the window manager you're using.

See Also: window manager, desktop.

WAN

Wide Area Network. This network, although similar to a LAN, connects computers on networks which are not physically connected to the same wiring and are separated by a greater distance.

See Also: LAN.

wildcard

The '*' and '?' characters are used as wildcard characters and may 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.

In the context of a graphical user environment, a window is the rectangle which occupies a given running application which usually contains a title, a menu, a status bar, and the application's work area.

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, ...

workspace switcher

A little applet which allows you to switch between the available virtual desktops. It is also known as pager.

See Also: virtual desktops.

Index

- Appletalk, 132
- applications
 - kill misbehaving apps, 145
 - killing, 146
 - troubleshooting tools, 147
- ATI 3D graphics cards
 - OpenGL, 146
- backup, 138
 - Master Boot Record, 143
 - restore, 141
 - tar, 140, 140
- boot
 - different run level, 142
 - file-system, 141
 - system hanging, 141
- boot disk, 137
 - Master Boot Record, 143
- bootloader
 - dual-boot, 143
 - reinstall, 142
- Borges, ??
- cable
 - null modem, 132
 - parallel, 133
 - PLIP, 133
 - twisted pair, 135
- CIFS, 132
- commands
 - Kppp, 146
 - minicom, 146
 - tar, 140
- console
 - switch to another, 144
- development, 2
- DHCP, 129
- DNS, 129
- DocBook, ??
- documentation
 - Mandriva Linux, 3
- DOS, 30
- eth0, 128
- Ethernet
 - card, 128
- file
 - deletion recovery, 144
- filesystem
 - damaged superblock, 143
- gateway, 23
- GRUB
 - reinstall, 143
- IMAP, 59
- internationalization, 2
- IP
 - address, 125
 - routing, 126
- IPX, 132
- ISDN, 130
- legacy-free

- desktops, 146
- laptops, 146
- LILO
 - reinstall, 142
- MacOS, 31, 33
- Mandriva Expert, 1
- Mandriva Club, 1
- Mandriva Linux, 147
 - mailing lists, 1
 - security, 1
- Mandriva Store, 2
- modems
 - linmodems, 146
 - winmodem, 146
- MySQL, 79
- NetBEUI, 132
- NetBIOS, 132
- netmask, 125
- network
 - cable, 132
 - class, 125
 - configuration, 125
 - Network Information System, 83
 - private, 126
- networking, 123
- NIS, 83
- nVidia 3D graphics cards
 - OpenGL, 146
- openGL
 - ATI 3D Graphics Cards, 146
 - nVidia 3D Graphics Cards, 146
- OS/2, 36
- packaging, 2
- Peter Pingus, 4
- PLIP, 131
- POP, 59
- PPP, 131, 131
- programming, 2
- Queen Pingusa, 4
- RFC, 124
- routing, 126
- Samba, 132
- server
 - Network Information System, 83
- services
 - mail delivery, 59
- superblock
 - repairing, 143
- synopsis
 - command, 4
- system request, 145
- TCP/IP, 124
- Token Ring, 132
- troubleshooting, 137, 147
 - computer is slow, 147
 - filesystem, 143
 - Mandriva Linux, 147
- users
 - generic, 4
- Webmin, 37

- windows 3.11, 31
- windows NT/2000, 27
- windows 95/98, 25
- Windows XP, 24
- X, 142
 - configuration, 142
- x server
 - kill, 144

